JANNIE MØLLER HARTLEY

illustrations by Zarah Juul

Just for you

Inside the Personalisation Algorithms that Curate Our Image of the World

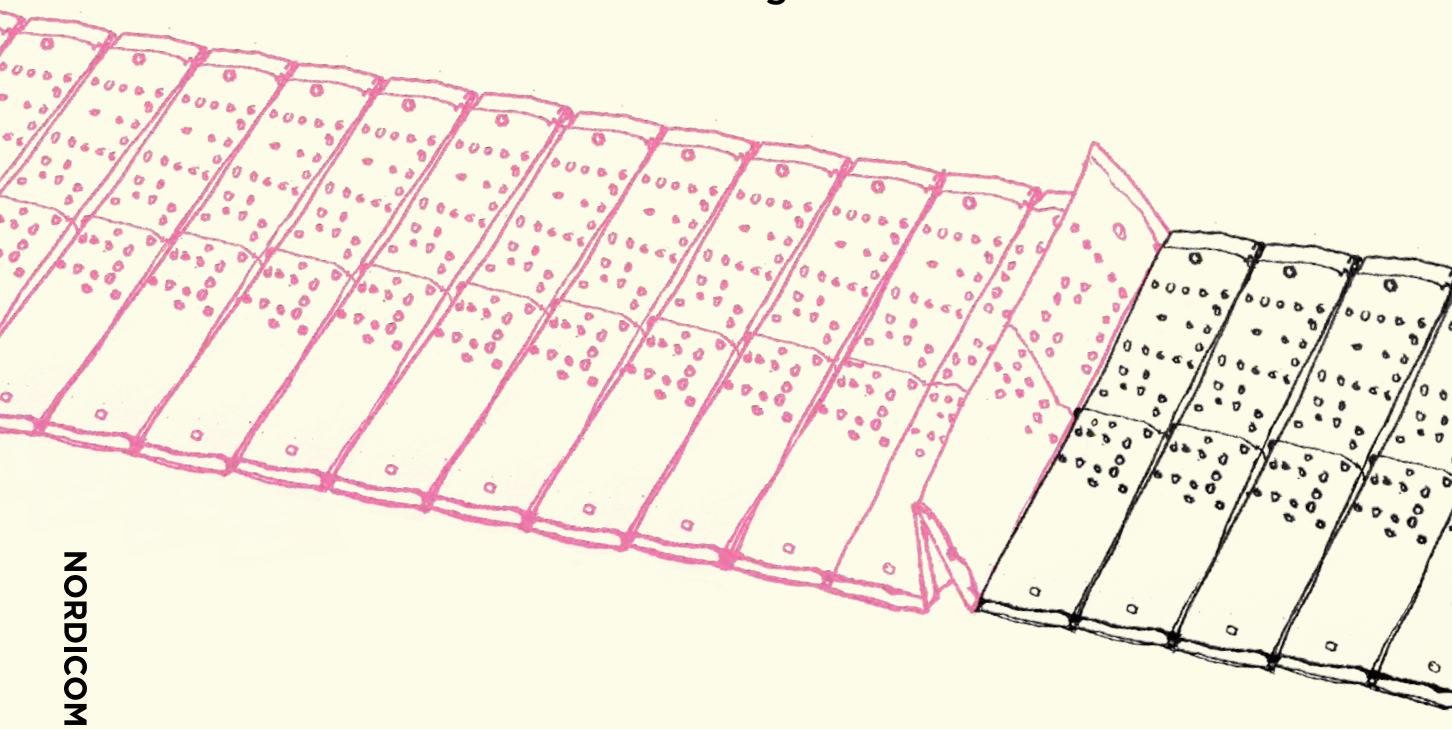

NORDICOM

Just for You: Inside the Personalisation Algorithms that Curate Our Image of the World

Jannie Møller Hartley

© 2026 Jannie Møller Hartley. This work is licensed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International Public licence (CC BY-NC-ND 4.0). To view a copy of the licence, visit www.creativecommons.org/licenses/by-nc-nd/4.0/

This book is based on projects funded by the Carlsberg Foundation and the Velux Foundation.

ISBN (PRINT): 978-91-89864-31-3
ISBN (PDF): 978-91-89864-32-0
ISBN (EPUB): 978-91-89864-33-7

This work has been subject to a double-blind peer review by two independent and external anonymous scholars.

This work is Open Access and available to read free of charge and without registration requirement at https://www.nordicom.gu.se/en/publications/just-you

DOI:10.48335/9789189864320

CITE THIS WORK:
Hartley, J. M. (2026). *Just for you: Inside the personalisation algorithms that curate our image of the world*. Nordicom, University of Gothenburg. https://doi.org/10.48335/9789189864320

PUBLISHED BY:
Nordicom, University of Gothenburg
Box 713
SE 405 30 Göteborg, Sweden

www.nordicom.gu.se/en

Editor: Kristin Clay & Johannes Bjerling
Manuscript editor: Kristin Clay
Managing editor: Josefine Bové
Layout/design: Kristin Clay & Karin Andén
Cover & illustrations: Zarah Juul

Inside the Personalisation Algorithms that Curate Our Image of the World

JANNIE MØLLER HARTLEY

illustrations by Zarah Juul

Contents

Prologue 6
Preface 8
How to read this book 14
Structure of the book 20

Part 1: First there was a machine – How algorithms are (old) culture 24
The personalisation paradox 30
Tailoring for the individual with a public mission 38
The magic trick of selling you as data 46
An entangled relationship with machines 54
The profitable nature of your data doubleganger 62

Part 2: Inside the black box – An anatomy of an algorithm 68
Opening the black box 73
The plurality of algorithms 83
The hidden layers of the machines 89
Sorting and cleaning data 90
Growing from data on you 96
Building the architecture of AI 104
Categorising and tinkering the world 113

Part 3: "Massaging" you with thinking machines – How algorithms curate flows of information 120
The maintenance of you with machines 126
The form of the news and the publics it carves out 134
Increased commercialisation of the news media 142
Digitalisation paves the way for personalisation 149
Hello new and bigger players! 158

Part 4: The deception of knowing you – How algorithms are hyped as almighty 166
The machine-versus-human thinking battle 171
A distorted image of the user 178
The public mission at stake 184
Elaborate illusions of human-like intelligence 198
You and your dataself 205

Part 5: Stuck – How algorithms trap you in the labyrinth of large language models 212
The commercial heritage of personalisation algorithms 216
An invisible architecture of user surveillance and control 229
Caught in the endless loop 237
First date with the algorithmic news butler 244
Algorithmic creepiness 254

A conversation with Ada – And a sort of history of the future 258
Rewriting the terms of agreement 273
Re-imagining and caring for the user as co-participant 287
Algorithmic personalisation as a matter of public concern 300

Credits 303
Endnotes 304
References 316

Prologue

"I have to ask you, should we just stop the project?"

We had just stepped out of the room together.
I was wearing a badge:
"The 18th International Recommender System Conference",
and my name.

His badge said: "Keynote Speaker".

The question was posed on the enormous concrete staircase
of the Polytechnic University
on the outskirts of Bari in Southern Italy.

As we moved through the crowd,
I knew he was waiting for an answer.

But all I could think of was how little I truly understood
– not just about the personalisation algorithms
and the machine learning models behind them,
but about what they meant,
and what stopping them might mean
for the future of the news media.

How might Ada Lovelace, who invented the algorithm in 1843,
answer that question?

We could say that the entire journey of this book starts on that large concrete staircase in the soft heat of Italian autumn, when someone asked me a difficult question.

In reality, it begins much earlier
– perhaps already when looms were automated
around the time when Ada Lovelace was born.

We will return to the staircase,
but let me first tell you how this book
is trying to tell you the story of algorithms
in a different way…

Preface

They wake before you do. In fact, they never sleep. By the time you open your phone, something has already been arranged. A map appears before you type your destination. A song you had forgotten you loved begins to play. Somewhere between a late-night search for dishwasher tablets and a quiet moment of domestic despair, the Internet gently suggests a new washing machine.

And they deliver today's front-page headlines curated "just for you".

The algorithms seem to know. Or at least, they perform knowing convincingly enough that we rarely question it.

We live with them now. And we live by them now. We move through algorithmic suggestions like water through a carefully shaped canal. They anticipate, recommend, filter, and sort – organising not only information, but the texture, and the smooth and quiet choreography of everyday life itself feels kind of natural. Perhaps even a helpful nudge in an overwhelming sea of information overload. And so, we accept them and let them make decisions for us.

But the selling of algorithms as the smooth helpful friend comes with a cost: It hides the machinery and the inner workings of the algorithms inside a black box. Yet, algorithms are not mystical forces – even if the owners and tech platforms would like us to think that they are. Black boxing, fetishising, and mystifying is also a form of power. The algorithmic systems are purposefully designed systems built by humans with particular assumptions, priorities, and intentions. Even if those intentions are obscured into multidimensional spaces of complex automated calculations across billions of datapoints. Importantly, as users, we delegate our decision-making power and agency to those mysterious algorithms.

Nowhere is this more consequential than for the news media.

News organisations around the world have begun to install personalisation algorithms at the very heart of journalism. These

systems curate what we see, what we miss, what rises to the surface, and what quietly disappears. These carefully designed systems reorganise our information environment, quietly restructuring how news reaches us, how knowledge circulates, and how publics form. Algorithms are no longer peripheral tools in journalism. They sit at the centre of our information ecosystems, quietly editing reality. Removed from the journalists' and editors' gut feeling of what a good story is, editorial priorities are moved into algorithmic systems and continuous, invisible adjustments – ranking, sorting, and predicting according to market demands. From the users' perspective, we might get what we want, but along the way, we also leave our agency to the algorithmic paths, an agency which seems to be fundamental for how we act as democratic citizens.

In short, algorithms are quietly editing reality. They are shaping our image of the world. And yet, we know remarkably little about how this happens. I wrote this book from within that uncertainty – from within the black boxes – to explore how it happens.

I argue that journalism and the news media still matter – not only as a profession and a particular organisation of that profession, but as a democratic infrastructure: a fragile system through which citizens in democratic societies understand themselves and the world. Because algorithms and artificial intelligence (AI) models are now part of that infrastructure, understanding how they work and their impact is no longer optional. The rise of algorithmic systems does not make the public mission of the news media obsolete, even if it in theory is possible to leave it to AI agents to sort out the world for us. If anything, it makes it more urgent.

My path into this story began in 2009, when I returned to the newsroom – this time no longer as a reporter, but as a researcher. That shift gave me a dual vision I have carried ever since: participant and observer at once. At the time, I was researching the digital transformation of journalism as it was unfolding. There was mostly optimism in the air. The Internet promised openness, participation, a more democratic public sphere. In short, it promised more agency and autonomy for

users and that a better and more inclusive democracy would naturally follow. Yet, as we know, this never happened. What did happen was that the Internet and its social media platforms moved in as a useful but time-consuming friend, yet increasingly hard to live without.

But alongside that optimism, something else emerged. Metrics began to glow on newsroom screens and the screens moved into the newsrooms. Digital publishing made 24/7 real-time news king – and speed became a form of authority. It struck me how the rhythm of journalism started to align with the rhythm of machines. Going digital was so much more than just new measurement tools and sites – routines, values, and priorities transformed as well. The professional ethos of journalists and their judgment was distributed to screens and systems telling them what to do, displacing their experience gut feeling.

Then, in 2022, another shift arrived. With the release of large language models like ChatGPT into the public domain, the terrain changed again. The raw material of journalism – text itself – became something machines could generate. News organisations, already beginning to experiment with personalised distribution, suddenly found themselves facing a much larger question: What happens when the machine not only curates the news, but also writes it? And does so in multidimensional computational spaces so large that we often cannot grasp what the models are actually doing? I returned as an observer inside news organisations, and my fieldnotes from different years began speaking to each other and forming a pattern. And I observed something deeper than speed or format-shifting.

What had once felt like a gradual transformation became something more urgent, more unstable, and more difficult to ignore. This AI tsunami we are seeing currently makes the digital shift seem like a piece of cake. At stake is more than a shiny new technology. It is the structural foundations of a relationship. The relationship between the news media and the citizenry has always been fragile, dependent on trust, attention, and a shared sense of purpose. Today, that relationship is being reconfigured, and it's more fragile than ever before. News no longer travels directly from newsroom to reader: It moves through dense digital environments, shaped by platforms,

signals, and competing demands for attention. And now, filtered via AI-driven personalisation algorithms and AI agents. Algorithms which guide us to the next film or news story. Algorithms which tell the journalists what to cover and when. News used to be attached to artifacts, articles, videos, stories. In the future, the news media and other producers of content face AI agents as "audiences", which might be the single biggest structural change our information environment has ever faced. We are not there yet, but this book points into the immediate future. By laying out the path we are already on.

Algorithms now sit between the news media and the public – subtle intermediaries, quietly influencing what we see and how we see it. Algorithmic systems are turning journalism from a reactive meaning-making machine to a predictive calculative machine, trying to foresee what you might like to be informed about next. Just for you. As my notes unfolded, I began to think about that relationship between the news media and its users as something resembling a marriage under strain. Expectations no longer aligned and trust becoming fragile. Attention wanders toward the many other stories calling to us from our screens. It is as if algorithms and its systems quietly sit between the partners – like the screens which have also come between us in our relationships with partners, friends, and family members. This metaphor of the relationship under strain stayed with me as I was writing this book.

If we are to understand what is happening to journalism, the news media, and this relationship with the users, we need to understand this third presence, which is why I wanted to open the black box of the personalisation algorithms. Not to demystify algorithms and machine learning by reducing them to simple tools, but to approach them with sociological curiosity, anthropological patience, and a certain journalistic stubbornness. Because algorithms are not mystical forces. They are designed systems – built by humans and shaped by assumptions, priorities, and forms of power. And they do things as they get entangled into our lives.

To understand algorithms is to understand the worlds they both reflect and produce.

The pages that follow are drawn from years of fieldwork inside news organisations. From conversations with journalists and developers, from observations of how editorial ideals – truth, public service, democracy – are translated into mathematical logic and machine learning models. The algorithmic systems are also shaped by a longer historical perspective, one that reaches back to figures like Ada Lovelace. In the mid-nineteenth century, long before smartphones and recommender systems, Lovelace wrote what many now consider the first algorithm intended for a machine by imagining the language of computers. But more than that, she imagined something radical: that machines could manipulate symbols, patterns, and ideas – not just numbers. Importantly, in the moment we stand in today with the large language models imposing themselves on every corner of everyday life and work, it is imperative that we understand their history and logics. I include Ada in the story as a reminder of the fact that we need to pay attention to intent and agency. She struggled, but in the end found her own agency to start imagining how a language of machines could be developed. Machines, she argued, do not think on their own – they follow instructions designed by humans; thus, those instructions and illusions of agency they bring are important to pay attention to in current times. Ada's story also portrays algorithms as something we can influence with our human imaginative minds (and a bit of poetry), and automation as a challenge we have faced before. Moments when new machines rearrange how we communicate, how we work, how we imagine the future, but also moments where we can decide to pay attention to what we do and perhaps even do things differently. The purpose of bringing her in is not just to bring yet another Ada Lovelace story to the world (there are quite a few of those!) but to challenge the view of the algorithm as something singular, something unavoidable or even new. When automating with algorithms, we distribute agency to those systems – an agency that might be fundamental for how we act as democratic citizens and how the news media cultivate publics.

Algorithms are not new in essence, even if their scale and reach are. And the story of AI is already an old one. What is new is the extent to which they are woven into the fabric of our information environments – and how quietly they shape what becomes visible and what remains unseen.

This book is an invitation to slow down and look more closely at the algorithmic systems that increasingly organise our shared reality. I wanted to emphasise continuity rather than revolution when we revisit our technological trajectories. To emphasise the human–machine entanglement rather than clear opposing dualities such as subjectivity versus objectivity, virtual presence versus in-real-life. And instead to ask how algorithmic "knowing" is constructed – how a reader becomes a pattern, how a preference becomes a signal, how millions of small interactions are folded into predictions about what we might want next. And what it does to the users of the algorithmic systems. How the values of the news media and its role in democracy are built into the punch cards of the AI age. This is important as we look into a future of AI agents delivering the daily news. What happens to the situational knowledge and judgment of both citizens and journalists when agency is increasingly distributed to algorithmic systems? Can algorithms and AI help us navigate the democratic maze, or does it strip from us the agency needed to make up our own mind?

And on the basis of that, I ask a deeper question: What kind of information world do we want to live in?

I hope to inspire the reader to answer this question, as the book moves between observations from inside the algorithmic machines, historical detours, everyday encounters with personalisation algorithms, and reflections on the changing relationship between humans and machines, and news media and their audience. It follows threads connecting Lovelace's personal struggles and early visions of computing to today's personalised information systems. I hope this journey can inspire us to consider: What kind of news media do we want? What kind of publics do we hope to be, and what drives us as users of news and information? And what kind of role should machines and algorithms play in holding that relationship together?

How to read this book

As the idea for this book developed and the pages started unfolding, the familiar format of the academic book felt strangely inadequate. The traditional voice of scholarship speaks from a stable distance: The researcher observes, analyses, and eventually explains the world with reassuring clarity (and often in complex language). Inside newsrooms and the machine rooms of algorithms, the story feels messier, more human, more uncertain. Engineers debate probabilities while editors argue about public responsibility. Citizens move through their days surrounded by personalised streams of media that feel both helpful and unsettling. The familiar story about algorithms – machines manipulating helpless humans – felt too tidy (and loud) for what I witnessed inside the machines. Dystopian prophecies and utopian dreams both circulate loudly within these stories. Neither quite match the everyday reality unfolding inside the machine rooms of journalism, where editors and developers care for the users and are concerned about how to cultivate them into informed deliberating citizens. They worry about the Big Tech ownership of the AI models they use. But in the commercial fight against platforms for users' attention and their enthusiasm for bigger and better machines, the public mission has become increasingly hard to uphold. Paradoxically, they have (along with other foundational democratic institutional authorities) dismantled their own authoritative role or at least distributed it to global-scale calculative systems, and now they question if algorithms and AI can help them regain that role.

This book grew from those frictions. By staying with the trouble, as feminist scholar Donna Harraway[1] would put it, the story emerging was messier. But that's ok. For the most part, stories are messy (and so are algorithms and AI, in fact). So, instead of smoothing the story into a single clear voice, I allowed the book to unfold through several threads – or as different layers of the same story. Each thread tells part of the story of algorithms and their entanglement with the news media and with how users consume information and news. Each thread takes a different view of the connected story: a broad outside view of algorithms, a view from inside the machine room, and the blast from the past of Ada's story. Each thread has a different visual presentation, so you can visually locate where you are. Read together, my hope is that they reveal something more complex than any single perspective could capture.

Thread one: Understanding the algorithms

This more classic academic thread in this book connects the observations from inside the media organisations with a broader history of media, technology, and audiences' democratic life. Here, I step back and examine how algorithms entered journalism, how personalisation systems work, the role of AI and large language models as a part of these systems, and how it is reshaping the relationship between the news media and its audience. Algorithms are often described as mysterious black boxes. I refuse that mystique.

This thread shows how algorithmic systems are built by people, embedded within institutions, and shaped by economic pressures, professional norms, and human imagination. Shaped by the imaginaries of the many people who have been involved in the development of the machines and the computational language their operations are based in. Understanding the role of algorithms in news media requires attention to all these elements at once. In this thread, I explain how algorithms operate and why they matter – not only for the news media as a democratic institution but for citizens navigating an increasingly personalised information environment. The thread critiques the politics, logics, and socio-materiality of adaptive algorithmic systems. And it shows how the agency of journalists and citizens is quietly distributed to these systems. It does so by emphasising the growing power of what can be described as *algorithmic governance of the user.* Crucially, the users in this case are both citizens as news users and journalists as users of the algorithmic systems. What emerges is a portrait of the news media in transformation: relational, datafied, and deeply entangled with the algorithmic systems that now help compose its audiences and cultivate publics.

If thread one *tells* the story of algorithms and their impact in and across the news industries, the second thread aims to *show* the story.

Thread two: Inside the machine room

This thread grew out of my ethnographic fieldnotes – observations gathered while following the development of personalisation systems inside and around news organisations. The thread brings

you directly into the environments where algorithmic systems are built and negotiated. The term *algorithm* in this case covers building what's called recommender systems with machine learning, which are then personalised with data on you. The thread brings you into meetings between editors and developers. Conversations with data scientists designing models that predict what readers might want to see. Into the everyday small negotiations – where journalistic ideals meet computational logic. I chose to put myself in the text to allow for the frictions in this work to appear to the reader as they appeared before me. The methods deployed are inspired by Seaver's "scavenging strategy".[2] They include Amoore and colleagues' "reading computer science texts",[3] conference ethnography,[4] and ethnographic observation and interviews with AI developers and media professionals.[5]

I converted some of the fieldnotes into a reflective, sometimes poetic tone. Some entries document concrete moments: a data scientist explaining how a recommendation system learns from user behaviour, an editor debating the meaning of "relevance". Others drift into reflection, drawing parallels between contemporary algorithms and earlier technological (r)evolutions. These fragments do not follow strict linear chronology. They move across different moments of fieldwork, weaving together episodes from several years of ethnographic observations in different organisations. This I also did with the purpose of anonymising the people involved. This was never about individual actions or about what happens inside a single news media organisation, but about something systemic and structural. I allowed uncertainty, curiosity, and even doubt to remain visible as meta commentary on the things I observed. Observing algorithms is rarely a purely technical exercise. It involves emotions as well: fascination, hesitation, moments of confusion when machine logic collides with human judgment. Writing the fieldnotes in this more open form allowed me to capture that messiness as it unfolded but also to bring the messier story to the forefront of the reading of this book.

Together, the first two threads form a living woven portrait of news media and citizens learning to coexist with machines and adaptive algorithmic systems. And in the third thread, we travel back in time to *feel* the story of how algorithms entered the world.

Thread three: Meeting Ada Lovelace

The story of Ada Lovelace feels almost unreal at times. Victorian society expected women of her class to be graceful daughters, elegant wives, and attentive mothers. Mathematics and machines belonged firmly to the world of men. Yet Ada Lovelace developed a fascination with machines and abstract logic that pushed far beyond the expectations surrounding her, and she found her own way. Working with the mathematician Charles Babbage, she wrote a set of notes describing how a mechanical engine might execute a sequence of operations, which is a process we now recognise as an algorithm. She imagined instructions guiding a machine step by step – decades before such machines existed.

The events I describe in her life have all taken place (apart from her dream of flying across England on a mechanical horse), and I have added the sources of my research as endnotes, which you find at the end of the book. But I have taken the freedom to make Ada's story into a literary narrative, for the reader to think and feel alongside Ada and to draw parallels to current developments of algorithms and AI. And to show how individual agency, judgment, and imagination is crucial and entangled with the technological development.

Making it more literary also allowed me to include the story of the Luddites who resisted automation of the loom, because it illustrates a parallel to the current loss of journalistic jobs to automation with AI. I cannot know for sure how Ada felt about her absent parents. I cannot know if she knew or cared about her famous father Lord Byron's fight for the Luddites. We can only know as much as we can read in the many books about her and the correspondence between her and others which has been preserved. The direct quotes I found in these sources.

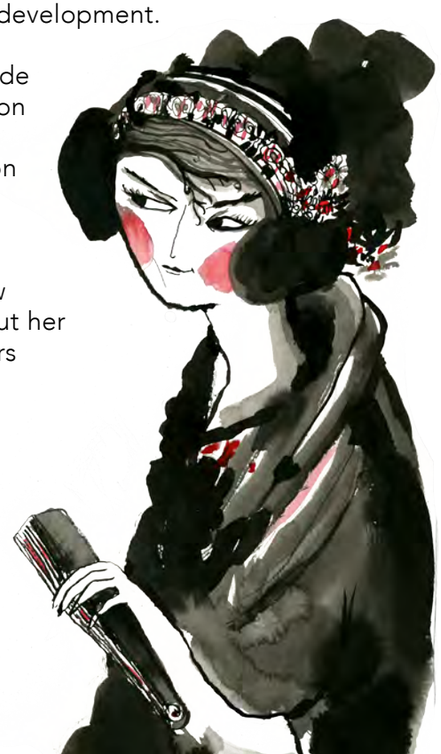

Drawing machines, drawing meaning

Scattered throughout the book you will find hand-drawn illustrations. The technological systems I studied inside newsrooms – the machinery of journalism, algorithms, metrics, data flows – are complex and layered. Like the multidimensional space of a machine learning model. And Ada's story is linked to many of these layers of algorithms' entrance into news media. Drawing becomes a way of making those relationships visible. Hopefully it also slows the thinking down. It invites another and more imaginary form of attention.

I could easily have generated images or illustrations using AI. I chose not to. That choice reflects my conviction that creativity, interpretation, and human imagination remain central to journalism and to the way we make sense of technology. And to use AI would defeat the purpose of this book. Machines may rearrange information in endlessly new combinations, but meaning emerges through human creative reflection and artistic interpretation. In this sense, the illustrator's work mirrors the work of the journalist. Both craft ways of seeing and sensemaking. The drawings in this book were created by the talented illustrator Zarah Juul. She developed visual interpretations that accompany the narrative and mark transitions between the threads. I wrote and she drew. And with the form of this book, I hope I have crafted a way for the reader to see algorithms and their history differently. And to draw parallels to today's development of machines and the narratives around AI involved in this development.

The decision to include hand-drawn illustrations also reflects a deeper argument I make throughout this book.

Journalism is not a machine, even when it sometimes feels like one (which is why the story can seem quite messy). Unlike the algorithmic machines, the journalistic machine is almost always up to date, aware of its own gaps in knowledge, not trying to please at any cost, and quite good at separating beliefs from facts.

But importantly, journalism is a creative practice. With years of professional training, journalists interpret the world, shape narratives, and evoke emotions in a situational context. Illustration belongs to that same creative tradition.

Interwoven stories

The threads – analytical reflection from above, ethnographic observation from the inside, and historical narrative from behind – and the illustrations intertwine across the book's six parts. Together, they trace several histories at once. Firstly, the emergence of personalisation technologies and their journey into the news media machine; secondly, the evolving relationship between the news media and audiences; and thirdly, the longer story of computational thinking stretching from Ada Lovelace to Alan Turing to contemporary AI. The structure signalled by shifts in different fonts and spacing may feel unusual at first. That is intentional.

Technology (his)stories are often told as straight lines: invention, disruption, inevitable transformation. The way we navigate the Internet or a streaming service platform is equally supposed to be a journey with no friction. Yet the world of building algorithms does not unfold so neatly. It emerges through layers of human decisions, institutional pressures, technical systems, and cultural, creative imagination. So, this book moves differently. Thread by thread. Layer by layer. Like an algorithmic system with layers of code. What emerges is not a single explanation but a constellation of perspectives – as a way of seeing the machines around us more clearly. And perhaps, once we see them more clearly, a way of imagining how we might live with them differently.

Structure of the book

Part 1 takes us into the context of news personalisation, where I argue that personalisation in news is not merely a technical change, but a cultural and normative shift that reorients the news media toward the individual user as both data source and target for calculative predictions. Algorithmic personalisation emerges as a paradoxical practice that simultaneously individualises and standardises content, shaping what information users are exposed to. It's a practice governed by opaque systems most often inherited from global technology platforms with no intent to cultivate us into democratic citizens. In Ada's thread, we get to know Ada as she grows up and starts imagining building machines. And we learn how she uses her intense math home schooling to calculate and her poetic mind to imagine what they might be able to do. Alongside her story, I reflect on my position and journey into the world of the current Adas, and I am confronted with the obstacle of understanding the language of the machines and the people who build them.

Part 2 moves deeper into the development of the machine, where I investigate the anatomy of the algorithms that power it. This part opens the "machine room" of algorithm construction. Here I argue that algorithms are not neutral but embedded in socio-material infrastructures shaped by professional divisions of labour, technological imaginaries, and the relative invisibility of journalistic values in the development process. The algorithm appears not just as code, but as a site of politics, labour, and epistemology – a black box that governs what users see, even though users and journalists rarely fully understand how it works. Yet, it is given authority. In Ada's thread, she is also introduced to a machine – an automated loom running on punch cards – and the idea of a language for machines will not leave her mind. Meanwhile, in another part of the country, her future friend Charles Babbage is obsessing, building an enormous machine for automating calculation based on the ideas used for the automated loom. Inside the machine rooms of today, I set out to explore how the algorithm is made from categories of words, which structure the language of the models through binary categorisations, like the punch cards Ada and Charles were so fascinated by.

Part 3 shows how algorithms participate in a form of relation repair of the citizen–news media relationship, acting as a mechanical couple's therapist, able to see through the day-to-day bickering and deliver the rational analysis to fix it. This part traces how personalisation extends earlier media strategies and metrification through a new layer of data-driven and algorithmic interaction. Drawing on recent critiques of datafication and my own previous research, in this part I highlight the tension between journalism's public service mission and the commercial imperative to keep users engaged through algorithmic optimisation and audience maintenance work. Here I also raise questions of labour when creative practices meet machines of automation. And with Ada on an imagined mechanical horse, we meet the Luddites, as they resisted the automation of the loom – and the following loss of jobs and livelihood. Inside the personalisation machines, developers work hard on building and testing the models for algorithmic curation of news. Meanwhile, editors are puzzled over the image of the user the models are weaving, and journalists work out how to resist the logics imposed on them by the algorithms.

Part 4 reflects on the epistemological and professional consequences of delegating editorial tasks to machines because it seems like they might do a better job than humans. I ask: Can algorithms – and currently AI as a series of self-learning algorithms – think like journalists? And should they? I trace historical debates about machine intelligence (Turing, Lovelace, Kurzweil, and others) and connect them to contemporary journalistic practices, professional ethos, and the public service mission. Inside the machine room, we see how journalists and editors increasingly face a conflict between editorial authority and algorithmic audience logic – and a lack of the right data to target the user in the right way. Ada feels trapped in the shadows of her parents, and when Charles makes a trip to Turin, an unforeseen opportunity presents itself to her.

Part 5 situates the current wave of personalisation within a longer genealogy of advertising, e-commerce, and technological infrastructure of algorithms. From Neiman Marcus's IBM gift advisor to Amazon's and Google's recommender systems, the personalisation dream has always been tied to economic and behavioural governance

of individuals. In journalism, this dream morphs into algorithmic curation creating individualised worlds and biased user profiles, which I argue risk producing narrowed public imaginaries. We move outside the machine room and into the lived experiences users have of personalisation algorithms. Here, I frame personalisation as an invisible architecture of control – a form of "governance through personalisation" that shapes what users know, feel, and see, and ultimately how they construct their image of the world and their room for imaginary agency. Meanwhile, Ada works hard on translating Babbage's talk in Turin on the Analytical Engine, and she discovers a roundabout way of putting her own thinking into the article.

The last part of the book brings past and present into dialogue with my observations and arguments throughout this book. Using generative AI, I created a customised Ada bot by uploading her notes and other available sources on her life to a closed ChatGPT and prompted it to be like her. Or at least the "her" that could be deduced from her way of thinking and writing. Using the generative capacities of large language models, I created a speculative dialogue about the future between myself and this digital version of Ada. Through these conversations, we revisit the arguments of the book, explore feminist theories of technology, and imagine (different?) future paths. In these exchanges, we speak across two centuries. I speak from fieldwork in AI-driven news media. She speaks from a world of punch cards and mechanical looms. In the conversation, it all flows together.

This playful experiment allows the timelines to merge. It gives me a way to ask questions that are both personal and political. What does it mean to live through a technological revolution, and what parallels can we observe in the world of AI today? How do we preserve agency, creativity, and the professional values of journalism in the face of automation? Ada's chatbot voice does not claim truth. She is not a real AI reconstruction of the historical Ada Lovelace. That was never the point. The point is to open a conversation across eras of human–machine entanglement. A conversation that invites critical, poetic, and historical reflection on the systems we build today and their inherent limitations. Such limitations become more visible in

and throughout our "conversation", as they are displayed in Ada's responses to my questions. In this way, the constructed conversation is all about "don't tell it – show it".

Ada reminds us that we have been here before. And yet, this moment is also different, and Ada's story allows us to pinpoint how this moment is different. As AI-driven personalisation algorithms quietly step in as a third actor shaping our daily news, our cultural preferences, and the publics we are drawn to, it becomes increasingly difficult to determine where the human ends and the machine begins. Taste, identity, and relevance are filtered and formatted through algorithmic systems and infrastructures we do not design ourselves – but now live inside. As democratic agents in this world, we share our agency with these systems, which become our compasses while we might slowly lose the ability to really steer.

Still, these algorithmic systems are not inevitable. They are built. They can be reimagined. And like Ada sketching machines that do not yet exist, we can imagine different paths forward.

Part 1

First there was a machine – How algorithms are (old) culture

“Should we just stop the project?”
he had asked me.

His question was lingering on the staircase in Bari,
the silence of my non-answer felt increasingly awkward to me.

I had travelled to Bari in Italy
as part of my journey
into the black boxes
of personalisation machines,
to attend the world’s largest international conference
on recommender systems

– systems that combine multiple algorithms
to automate the delivery of content,
whether film or news, to media users.

Thousands of people building algorithmic black boxes
gathered in one spot for three days.

We were just exiting a workshop
on recommenders for news content
when he asked me THE difficult question.

We carved our way down the stairs,
and through the masses of people
who work on machines and algorithms,
now fighting their way through Italian local specialities:

a coffee stand in each corner,

a stand with a huge GELATO sign,

the smell of tomatoes and cheese,

big money sponsoring keynotes and keyrings,

Microsoft, Google, Amazon, and other platforms
all hoping to play a key role in the building of the machines.

The project the digital editor and keynote speaker referred to
on the enormous concrete staircase
was a large-scale innovation project,
of building personalisation algorithms for a news site,
which I had spent a good portion of the last years observing.

Ada Lovelace would have loved the conference.

The place was buzzing with people
who think through the poetry of numbers.

I felt out of place.

In 1828, outside a house called Bifrons, about four miles south-east of the city of Canterbury in England, a twelve-year-old girl wanders through the gardens of a large, rented estate. Her name is Augusta Ada Byron.

The estate is rented by her mother, Lady Annabella Byron, who had once been briefly married to Ada's father: the poet Lord Byron. Lord Byron is already famous across Britain, with his poetry circulating through salons and taverns alike, admired for its intensity and scandalous beauty. His life is equally notorious. He loved extravagantly, accumulated debts, and lived with a recklessness that scandalised polite society. Eventually, the pressure becomes too great. Pursued by creditors and gossip alike, Byron flees England. All of this is largely unknown to the little girl – Ada has scarcely known him and remains behind with her mother. The little girl is mostly alone with her nannies and tutors on the estate.

One afternoon in the garden, Ada stops suddenly. Something dark lay in the grass. A dead crow. She crouches beside it and carefully lifts one wing between her fingers. The feathers are glossy even in death, arranged in a delicate geometry she cannot stop studying. She turns the wing gently, examining how each feather layers over the next.

A thought rises quietly inside her mind: "I wonder if I could build wings and fly?" The idea feels completely possible to her and an excitement builds up inside her. It feels like a puzzle.

She stands up, brushes dirt from her dress, and hurries back toward the house where her tutors await. Her mother has designed an unusually intense mathematical education for her daughter. Ada is to study numbers, logic, and science with a rigor rarely offered to girls of her time.[6]

The personalisation paradox

As a former journalist, I know the journalistic machine room from the inside, although it has been quite a few years since I sent out my last news report into the world. As a researcher of technologies and their impact on my former journalistic profession, I am also acutely aware that new technologies are rarely as dramatic as we fear nor as revolutionary or democratising as we hope. Yet, I was somewhere between puzzled and staggered – and nearly choked on my morning coffee – when I read the small news story in the union magazine that started my new research journey and this book project nearly four years ago. The story I had stumbled across informed me that JFM, a Danish regional publisher, had decided to embark on a project to build an algorithm for more personalised and automated curation of the news across their many digital news sites.

The reason why I was so staggered was that there was something inherently paradoxical in how a news organisation that considers themselves as a guardian of some kind of collective public sphere would leave that job to an algorithm. And to an algorithm aimed at tailoring content to the *individual*. Despite knowing very little about how these algorithms work and what this meant, I rang the editor in charge of the whole project and asked if I could enter the machine room with them as they built it. Lucky for me and to my big surprise, he said yes. So, for several months I took the train across the bridge connecting the Island of Zealand and the Island of Fyn and installed myself in the media organisation, a big building just across from the main train station. For months I observed how they planned, built, and tested the personalisation algorithm. At that time, I had no idea how much they were standing on the shoulders of other giants (and numerous failed attempts) at personalising the news feeds before them. I also had no idea how difficult the task of building these algorithms for news really was. Yet over time it became clearer to me why it was so difficult.

News selected and packaged for you

From 2015 to today, the news industry has been increasing their efforts of guessing and predicting what users want and placing their preferences in complex machine learning models. Algorithms personalise and package the content "just for you". You might have scrolled past and paid little attention to the boxes on the news sites presenting content "selected for you", but you recognise them from an array of other social media or streaming platforms. You might even get slightly annoyed that what they offer you is not "quite you", or you might like that the recommender systems installed on the back-end of the news site, or whatever other site you are visiting for media consumption, seems to actually know you. You might feel that they give you a sort of power and freedom over your media use. And when they inform you that "we are using this data to personalise your site", you might feel like it's making your news use more relevant and, in the end, better, because you get the news that really interests you. It seems to know your taste – as when we learn of new music very similar to the music we already like. Or at times you might find it creepy, wondering "how did they know I was in need of a new washing machine, before I even knew myself?" We might trust their recommendations so much that we turn left if Google or ChatGPT tells us to, even when we know that the left turn is a dead end. And we might discover things about ourselves or about others close to us that we did not know. Based on analysis of what customers had bought, the American retailer Target found a way to determine when women were pregnant based on the products they were buying, which was a huge success – until a father showed up angry at the store, because his teenager had been sent coupons with pregnancy-related products. She was in fact pregnant, but the father did not now.

Product placement

Amazon was an early adopter of algorithmic personalisation. Their product recommendations were the result of inferences made from what customers looked at, at which times, and for how long. Netflix

was also quick to adopt personalisation strategies for television and film recommendations for viewers.[7] Personalised recommendation is big business, and deeply involved in the development of these technologies are some of biggest and most powerful companies in the world: Microsoft, Apple, Alphabet, Meta, and Amazon. This matters for how these systems are developed, for how the user is imagined and cared for in such algorithmic systems and in the technical back-end infrastructures that run those systems. Later in this book, I return to how it matters and why the history of these systems and how they are developed is important for the news user and the relationship citizens have with the news media in particular. For now, I will just let the idea sit with you that the news media plays a role in cultivating democratic citizens and how those citizens are grouped into collectives. As news users, we are not just buyers of a sofa or a new washing machine. News and the circulation of information enable us to take part in something larger than ourselves, and when doing so, manifesting democracy as a current and historic institution. Personalisation and algorithmic curation challenges that position, because it fundamentally individualises and governs the news user into a certain behaviour, which differs from the role news media has cultivated previously. But as we shall see, it's a complex story because – unless they are a public service company – the news organisations that encompass the institution we call "the news media" are often also commercial entities. They need to make sure that they reach you with the content they produce to sell subscriptions and advertising, but also, if news media do not reach users and retain their attention, there's no public mission or sphere to cultivate. And algorithms can now help them do just that.

A great number of tech platforms have showed them how to do it and how efficient a tool algorithms are for keeping your attention.

The tech companies have for a while been embracing what Fan and Poole have called the "intuitive but also slippery" practice of personalisation: the process of delivering apparently individually tailored online content and services to users based on their habits, preferences, and identity markers.[8] By intuitive, they mean that it's the natural next step with all the data available on users, and by

slippery they mean that it's not always clear what is meant with personalisation or what companies want to achieve with it. The authors argue that most of the time, companies are more occupied with doing personalisation than with doing personalisation *well*. We know these practices of offering "this might be what you want" from product recommendations used in e-commerce.[9] Personalised product recommendations have been related to information technology from an early stage and charged with the hope to "shift the focus of traditional mass advertising to more concentrated and focused audiences".[10] The dream of delivering just the right offer for users to jump onboard and click "purchase".

Like the editors and data scientists – and like Ada Lovelace – I was fascinated by these new algorithmic opportunities. For news media, they would make it easier to capture the attention of users by getting to know them through data and using this knowledge to give them more of what they really like and less of what they do not like. And if the personalisation algorithms could help them to achieve that…

A public mission at stake

Personalisation is often also used to justify the tracking of users for purposes beyond the platforms we log onto. For example, like countless other "cookie notices" across the Internet, Microsoft's notification mobilises the term personalisation to explain why users' data are tracked and managed across its platforms.[11] The argument is, and no doubt this is also true, that it is convenient for the user to be offered much more relevant content, whether films, products, or news and information. Who does not want to be offered more relevant content specially tailored to them?

As I embarked on the fieldwork of entering the news organisations, I was unsure as to how this "knowing you" happened and curious about the calculative models they would use to predict the future need of users when it comes to something so unpredictable as news. But even if I also knew that the journalistic every day of making news is sometimes not as unpredictable as we might think, I was

convinced that there was more to AI than being a friendly and helpful robot in the newsrooms and on our phones as users. I had a feeling there was more to the entanglement of both the journalist and the machines, and more to the entanglement of the algorithm and the user of news and information than I could observe from the outside. I was just not sure yet what – but a few people had already hinted at the impact of the filtering effects of personalisation algorithms in areas other than news.

As Kyle Chayka argued in his book *Filterworld*, in which he wrote about recommendations of culture, this hyper-personalised landscape is not just about convenience: It subtly reshapes our perceptions, narrowing the range of ideas and experiences we encounter. Personalisation promises to show us what we want, but in doing so, it also dictates what we don't see. It replaces serendipity with predictability, curating a world that feels tailored yet strangely uniform. Chayka warned that as algorithms refine our tastes, they also reinforce them, trapping us in what he called "the cozy web" – a space of endless familiarity, where discovery is guided by past preferences rather than genuine exploration.[12]

What does this mean, when the individually personalised targeting and recommendations through machine-learning algorithms enter the domain of news? A domain labelled "the lifeblood of democracy",[13] where collective thoughts are shaped, challenged, and debated? A domain which at least in its ideal form is meant to secure a common public sphere beyond our interests as individual consumers? A cozy web of news might sound intriguing and something to strive for, but what are the costs and how does it dictate our path? And how does it reconfigure our image of the world so foundational for the enactment of our role as democratic citizens?

Those ideals of cultivating publics and citizens into at the bare minimum informed ones are (still) very present and strong in the minds of journalists as they go to work every day. They strive to tell you the stories you do not already know. The stories that are bigger than you, so you can take part in a public larger than yourself and your immediate interests. The stories which make you relate to others and the world, to connect to others, and to commit to a democratic conversation.

The news media has a public mission. And with algorithmic personalised curated news, that mission is on the line.

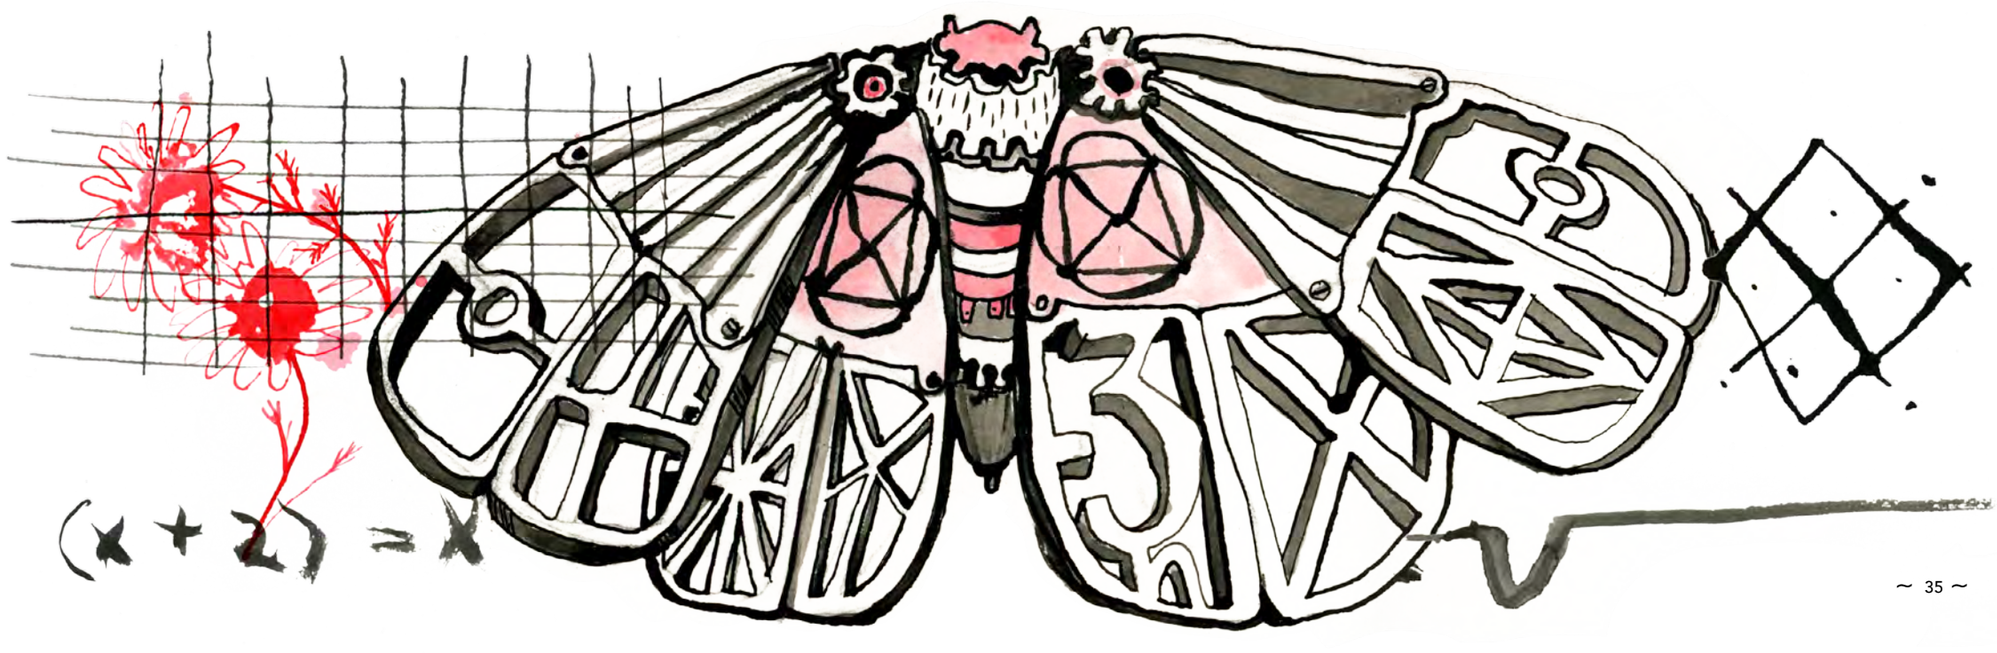

I doubt everything.
As a researcher,
it's my job.

I am a child of the Internet, born in the 80s.

I remember how it was to not have the Internet,
a time before algorithms.

I am writing a book on algorithms.

You could say that I am weaving an image of algorithms.

Meanwhile,
the algorithms are weaving an image of me
and of every other user of digital media in this world.

I will never finish.

I am slower than the algorithms.

The algorithms claim to know me.

But they do not know what I am thinking.

Journalists are weaving an image of the world every day.

I used to work as a journalist.

Ada's mother Anabelle is less enthusiastic about the prospect of flying machines and wants 12-year-old Ada to focus more on math. "My dearest mummy", Ada starts the letter to her mum. She struggles to find the words to a mother she hardly knows. A mother who has left Ada alone with carers and teachers, and yet Ada decides it must be best to obey her mum – as word about her continuous abstract imagining of flying would surely get back to her mum anyway.

"I received your letter this morning, and I really do not think that I often think of wings when I ought to be thinking about other things, but it's very kind of you to make that remark to me…"

Despite her mother's warnings, Ada continues to think about wings, and having seen the dead crow, proclaims to her mother: "I have now decided upon making much smaller wings than I before intended, and they will be perfectly proportioned in every respect, exactly on the same plan and of the same shape as the bird's", she writes.[14]

Needless to say, her mother is not happy and even her poet dad Lord Byron, who has not seen the girl since she was just a baby, is worried about her "thinking too much about abstract things".[15]

Anabelle does what she can to keep her daughter from imagining anything. Ada, being a girl, should after all focus on concrete activities such as playing the piano or riding her horse – serving her family name and legacy on her mother's side and making her a good match for a future husband. No one could have known that Ada's interest in the poetry of mathematics and her ability to imagine flying machines would play a central role in the development of algorithms.

Tailoring for the individual with a public mission

In the news story that made me call them up, JFM media described its personalisation strategy as a shift toward tailoring the news offer to the individual reader. The media group explained how this involves moving away from the long-standing premise of the print newspaper, where all readers are presented with the same content. Instead, the company framed personalisation as a way of delivering a "more flexible and responsive news product across platforms":

> Already this year we are launching the ambitious personalisation project that will enable us to deliver a tailored product, regardless of whether you go to a news site, receive a newsletter or read the e-newspaper on your mobile. The goal is that regardless of the platform, you get a product with a special focus on what exactly interests you. In the long run, your smart speaker may also deliver selected news to you.[16]

Yet, at the same time, the company emphasised that personalisation is meant to coexist with editorial judgment, stating that it would maintain "our right and duty to prioritise important news", while using personalisation "in the layer below" to create a more relevant product for local citizens.[17] This informed me that the public mission was being renegotiated but not completely dismantled or abandoned.

JFM owns 15 daily newspapers, 41 weeklies, and several commercial radio stations and digital sites and brands. For nearly two years, I followed the company's efforts to build personalisation algorithms designed to tailor news more closely to individual readers. These efforts are part of a broader development across the news media landscape. In Denmark, the public service broadcaster Denmarks Radio (DR) has introduced a mandatory login, framed as a way to "provide and recommend content fitting the single user's profile".[18] The British public broadcaster BBC has taken a similar step, requiring users to sign in to access its services. The BBC presents this move as a means of fulfilling its public service remit, arguing that "signing in unlocks a BBC that's personalised to you. And as a public service, it allows us to check we're making something for everyone".[19]

Across both commercial and public service media, personalisation emerges as a strategy for reaching audiences more effectively, including groups that media organisations feel a responsibility to serve. Tailored news becomes a way of articulating the public mission in a fragmented media environment. The underlying question appears repeatedly: How can news organisations maintain a strong foothold among citizens in a media landscape where attention is scarce and loyalty fragile? One answer increasingly offered from news organisations is the automation of news curation through personalisation algorithms. Building such systems – often referred to as personalised recommender systems – relies heavily on machine learning and, increasingly, AI.

And now it becomes a bit messy, so let me first clarify how I use the different terms and technologies in this book. I use the terms algorithmic systems, machines, recommender systems, and personalisation algorithms. These are not identical, as they constitute different ways of looking at and naming a phenomenon, and they are also used quite differently across both popular and academic literature, but they are closely connected.

I use the term machine as a metaphor for the entire apparatus involved in building and maintaining both algorithms and journalistic content – from people to news organisations to the broader infrastructural networks in which both systems and humans are embedded. This includes the technology platforms that own large language models like OpenAI and Google, for example.

Next, we need to clarify what the difference is between algorithms, algorithmic systems, and recommender systems. At its most basic level, an algorithm is a set of step-by-step instructions – a recipe for solving a problem or completing a task. But algorithms rarely operate alone. They are often combined into more complex algorithmic systems – collections of interacting algorithms that work together to handle more complicated processes, such as sorting search results or filtering e-mail. A recommender system is a specific kind of algorithmic system. It is built to suggest content – movies, songs, news articles, or social media posts – based on what you, or people like you, have interacted with in the past. It decides which film Netflix

places at the top of your list, or which news story appears first in your feed. This system is personalised once it is able to use your personal data to target and select content for you automatically.

What turns a system like this into what we often call AI is the use of machine learning. Importantly, not all AI is machine learning, but all machine learning is a subset of AI focused on training systems to learn from data and improve over time. Machine learning allows the system to learn from (past) experience – or past analysis of data, we should say (because a system or a model cannot really experience anything – it adapts from patterns of data). Instead of following fixed instructions, the system detects patterns in your behaviour and adjusts itself over time, and in the case of recommender systems, improving its ability to predict what content you might like to read, listen to, or watch next. This capacity to learn and adapt is what makes an algorithmic system appear "intelligent". The generative AI systems we have available currently are extremely good at appearing intelligent, because they are complex and build on extremely large amounts of data.

Distribution and curation represent only one part of the news value chain currently undergoing automation. At the same time, generative AI as a specific form of AI has rapidly entered newsrooms across the world as well as everyday life. Services such as OpenAI's ChatGPT have become widely accessible, requiring little technical expertise to use. Many people now keep such applications on their phones, asking them questions ranging from where to eat in Rome to how to fix a dishwasher. Although news is slow to become a part of the information that the large language models can draw upon, we can already ask "what's the latest news?" – a routine which is likely to be followed by conversational, AI-generated, personalised news updates. An intelligent news agent in our pockets. As this transformation towards AI news agents spreads quickly to more and more people, it seems urgent to examine what it means for the flow of information and how the algorithmic systems actually work. The problem is that these systems do not often reveal their inner workings, but once in a while, they pop up and make us aware of their smooth and frictionless existence behind our feeds, favourite sites, or streaming platforms. At times, the personalised offer is so spot on that it feels like the tech giants are listening to our conversations.

Convenient creepiness

My friends and I are sitting on a sofa in a Swedish cabin in early March, celebrating my birthday with a candlelit dinner. Outside, darkness stretches deep into the surrounding woods. Later in the evening, we gather around the fireplace, talking about possible destinations for a future trip. Rome comes up. Paris follows. Then someone suggests Marrakesh, and the energy in the room shifts. Voices grow louder as we sip gin and tonics, red lipstick leaving marks on the edges of our glasses. A few hours later, one of my friends picks up her phone. "I'm being flooded with offers for trips to Morocco now", she says. We look at her in silence. The question hangs in the air: Have our phones been listening? The uneasy feeling follows us as we crawl under warm duvets, listening to the wind in the trees outside.

Personalisation often feels magical, and it easily invites utopian enthusiasm or dystopian fear –whether sparked by algorithmic bias, hallucinations,[20] or the sense that our devices are listening in on private conversations. At the same time, many users experience a sense of recognition when their favourite news site fills with stories seemingly chosen just for them. Convenience sits alongside unease, and as we shall see, this is especially the case in relation to news and the news media.

As my investigation of the inner workings of algorithms went on, it became clear to me that many of those dystopian fears or explanations were too simplified, and the issues at stake, especially when it comes to the news media, are very complex. The answer to the question on the staircase in Bari was not just a "no" to using algorithms to curate the news, because it involves mass surveillance, questions of global labour inequalities, bias in models and training data, and so on. The answer was not a straightforward "yes" either, because these algorithmic systems are extremely efficient, convenient, and the news users are used to them anyway. The question becomes even more complex, because the news media has a public mission and relative autonomy from both state and market forces, and because the news media play a central role in democratic societies as a democratic watchdog and a facilitator of public debate – for cultivating an informed citizenry. For that reason, the growing use of personalisation algorithms across the news industries deserves careful scrutiny. These

algorithmic systems shape both the production and consumption of news, and with that, they shape our shared images of the world. This urgency drove my interest in examining how such algorithms are built, how they operate, and how their deployment may shift the normative position of the news media in democratic life. Algorithms rarely function as isolated tools that can be switched on or off in response to enthusiasm or concern. Like previous technologies, like the Internet, they tend to become entangled with journalistic practices in unexpected ways. Over time, this entanglement produces structural shifts in the relationship between you – the citizen and news user – and the news media as a provider of news for you. This book lays out what this structural shift entails.

With the term structural, I refer to changes that unfold gradually rather than overnight, shaped across organisations and systems rather than by single individuals or platforms. The term also points to deeper transformations in how we exist in and understand the world. To me, technological transformation happens as an evolution rather than a revolution, although the speed and scale with which AI is developing may make it seem like a revolution. However, as we shall see, news organisations have spent more than a decade learning to use data to profile users, drawing inspiration from companies such as Amazon and other major technology firms that have built vast databases for personal targeting across platforms. When Joseph Turow published *The Daily You* in 2011, advertising agencies still relied heavily on polls, surveys, and panel research to learn about audiences – processes that were slow and labour-intensive.[21] But as Turow anticipated, commercial technology companies have since moved toward targeting individuals directly, using real-time data surveillance. Although his analysis focused on advertising, he foresaw growing overlaps between advertising and news packaging already then:

> What this means is that we are entering a world of intensively customized content, a world in which publishers and even marketers will package personalized advertisements with soft news or entertainment that is tailored to fit both the selling needs of the ads and the reputation of the particular individual.[22]

We are no longer entering this world of customised content. We are in it. And, as it turns out, we have been for a while.

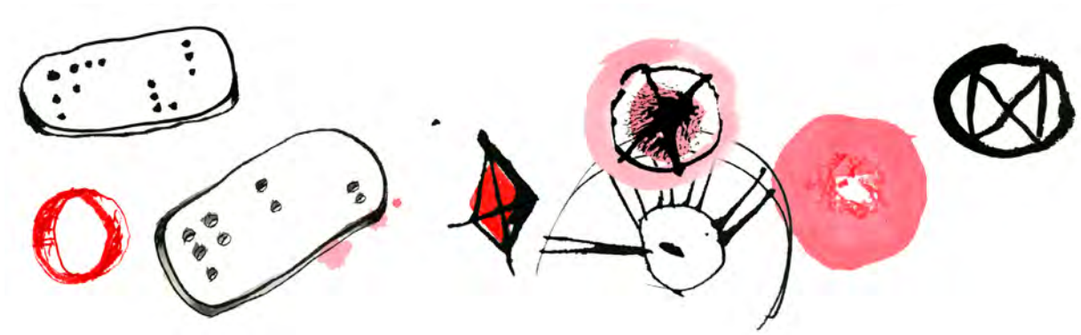

I am a part
of an aging race of researchers,
the non-digitally native.

I want to understand what the algorithms are for us,
and how they work.

How they sort and filter our image of the world.

I remember a time before,
which means that I am socialised into knowing
that things could be different.

But should they be?

I am schooled in being critical towards the algorithmic machines,
but I am curious.

With help from Ada,
who planted the seeds
for what we know as programming today,
I will explore
the story hidden in the technological development,
the story of how punch cards came into the world,
the story of how these cards
were the beginning of the binary language.

A story which might help us understand
what personalisation of media content means,
and where it might take us.

With an aim to integrate understandings,
rather than to deem the development good or bad.
With the aim of telling a different story.

A story of entanglements.

A story of binary understandings and broken relationships.

The technological development is a moving target.

It's a moving target, not outside of us,
but a part of the yarn
which is weaving our media history.

The warnings from her mother do not stop the imagination of young Ada, and she carefully considers her words when she writes her mum a few days later:

"As soon as I have flying to perfection, I have got a scheme about a… steam engine which – if ever I effect it – will be more wonderful than either steam packets or steam carriages.

…It is to make a thing in the form of a horse with a steam engine in the inside so contrived as to move an immense pair of wings, fixed on the outside of the horse, in such a manner as to carry it up into the air while a person sits on its back.

…This last scheme probably has infinitely more difficulties and obstacles in its way than my scheme for flying, but I should still think it's possible".[23]

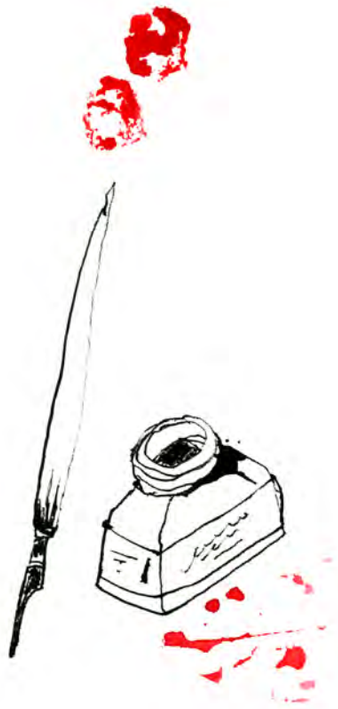

The magic trick of selling you as data

The reason why we already entered this new era of personalisation in the first place has to do with traditional news organisations finding it increasingly difficult to build new relationships – or even maintain existing relationships – with media users, especially the young. This development has been looming for quite some years, and it is making news organisations throw themselves into various innovation projects with the aim of improving their relationships with users – AI-driven personalisation being one of these. The BBC's recent announcement signals this new stage in the evolution of personalisation in journalism. In a message to staff reported by *The Guardian*, BBC news chief executive Deborah Turness outlined the creation of a new department: BBC News Growth, Innovation and AI. This unit will focus on using AI to deliver more personalised news, especially targeting audiences under 25, whose media habits have shifted dramatically toward smartphones and platforms like TikTok. As Turness explained, the BBC faces multiple challenges: rising news avoidance, a surge in social media platform news consumption, growing digital competition, and declining traditional broadcast audiences. Her suggested solution rests on data and personalisation:

> We must become ruthlessly focused on understanding our audience needs, on delivering the kind of journalism and content they want, in the places they want it, designed and produced in the shape that they enjoy it […]. We must deploy AI to support, enable and accelerate our innovation and growth".[24]

According to Turness, personalisation, recommendation, and tailored user journeys are central to this transformation. As mentioned above, this approach builds on a much longer trajectory. Over nearly 15 years of researching media transformations, I have watched how news organisations began to shift toward stronger audience awareness and deeper user integration across the different products they had on the shelves. It was simply not enough to produce quality content with an abstract user in mind and hope for continuing subscriptions and viewership – and hoping along the way that users would also become informed citizens and democratic publics. Yet, news organisations were threatened by citizens increasingly devoting their attention to so much more than news. From the early days of the digital transition in the late 1990s, through the rise of social media in the 2010s, it became increasingly clear that news organisations either adapted to data-driven production or risked losing their relevance altogether. Other actors stood by – ready to grab users' attention – with no public mission of cultivating informed citizens and a public sphere. The goal of tech platforms with data-driven filtered curation of content was primarily to make more money on advertising.

"Datafied" meant the ability to collect and analyse fine-grained data about audience behaviour, preferences, and patterns, and news organisations were no strangers to these tactics of retention. I conducted my first fieldwork inside newly digitalised newsrooms during the early wave of audience analytics. What I observed was a subtle but persistent shift in journalistic culture and values. As dependence on advertising grew, and traditional ad revenues declined, newsrooms were drawn closer to the logics of digital markets and the attention economy. In the next phase, social media platforms increasingly captured user attention – and advertising budgets – and news organisations started to consider personalisation as a potential answer.

It was during this time, as I carried out fieldwork for this book, that the realisation really spread: Personalised recommendation, already central to tech giants like Amazon or Facebook, might also work for the news industries. The success of these platforms in serving individually tailored content seemed worth emulating. The idea, however, grew much earlier. Back in 2011, tech writer Matthew Ingram referred to personalisation as "the holy grail of news",[25] capturing the idea that personalisation held out the promise of retaining audiences, despite the overwhelming volume of digital content available 24/7 across platforms.

A compass in the sea of information

Even as initial experiments in news personalisation by outlets like *The Washington Post* and *The New York Times* struggled to gain traction, the broader shift continued and spread to also smaller and local news organisations. The guiding argument was often that personalisation could help users navigate the infoglut – the overwhelming abundance of content in digital spaces. While a pamphlet offers a fixed amount of information, the Internet feels limitless, and that very abundance can become problematic for news trying to find its way to citizens.

Jonathan Cohn has described this as the burden of choice – the mental fatigue that arises when users must decide between too many options, even with algorithmic assistance.[26] Think of how paralysing it can feel to pick a movie on Netflix, or how long it takes to scroll through restaurant reviews. Recommendation algorithms offer a kind of relief from this burden. In practical terms, they prevent users from having to sift through millions of search results or perform complex comparisons manually. They provide a curated pathway through the information jungle – a compass through the sea of content. News publishers, of course, also recognised the business potential of personalisation. With digital advertising underperforming compared with its print predecessor, personalisation appeared to offer a new route to targeted advertising – a way to serve the right ad to the right person at the right time. This was a way of salvaging some of the value lost in the digital transition.

Digital fingerprints and data traces

At the heart of this transformation lies the data trail we leave behind. Cookies – small data packets that track our online behaviour – quickly became central to the economics of digital media. As users moved across the Internet, those digital traces allowed advertisers and media companies to track, profile, and target us more precisely.

Kasper Lindskow, now head of AI at Denmark's major news group JP/Politikens Hus, documented how deeply American news media became embedded in global networks of digital advertising by the mid-2010s. In his 2016 study, Lindskow showed that American news organisations were woven into a complex web of cookie-based ad networks, placing them in a dependency relationship with tech giants like Google. These platforms, due to their vast scale, accumulated far more user data than any individual news organisation could ever collect – no matter how many outlets it controlled under one brand.[27]

This dynamic mirrored what Joseph Turow had predicted earlier: Technology companies were moving toward methods of persistent tracking across platforms – what he called device fingerprinting. As Turow put it, we were entering a world where "publishers and even marketers will package personalised advertisements with soft news or entertainment that is tailored to fit both the selling needs of the ads and the reputation of the particular individual".[28]

In this environment, the user becomes the product. The currency exchanged is time and attention, not money. Big Tech companies sell that attention to advertisers. They don't just sell you a service – they sell the time you spend on their platform to the highest bidder. The more time you spend, the more valuable your profile becomes.

News organisations, now facing the same pressures, have moved to capture similar value. One key step has been encouraging or requiring users to log in, creating persistent identities across their sites and platforms. This enables them to track behaviour, infer preferences, and demonstrate value to advertisers. The attention of the user becomes a measurable asset. Personalisation algorithms are central to this shift. They help media companies hold your gaze a little longer – keep you clicking, scrolling, reading. They serve both editorial goals (no users, no democracy to safeguard) and commercial imperatives.

Through personalisation algorithms, users become the product (even more than users ever were), and our attention is sold in the marketplace of digital advertising. AI would turn out to be the key to the brave new world of the user turning product, as it could automate how news and advertising would be delivered to users.

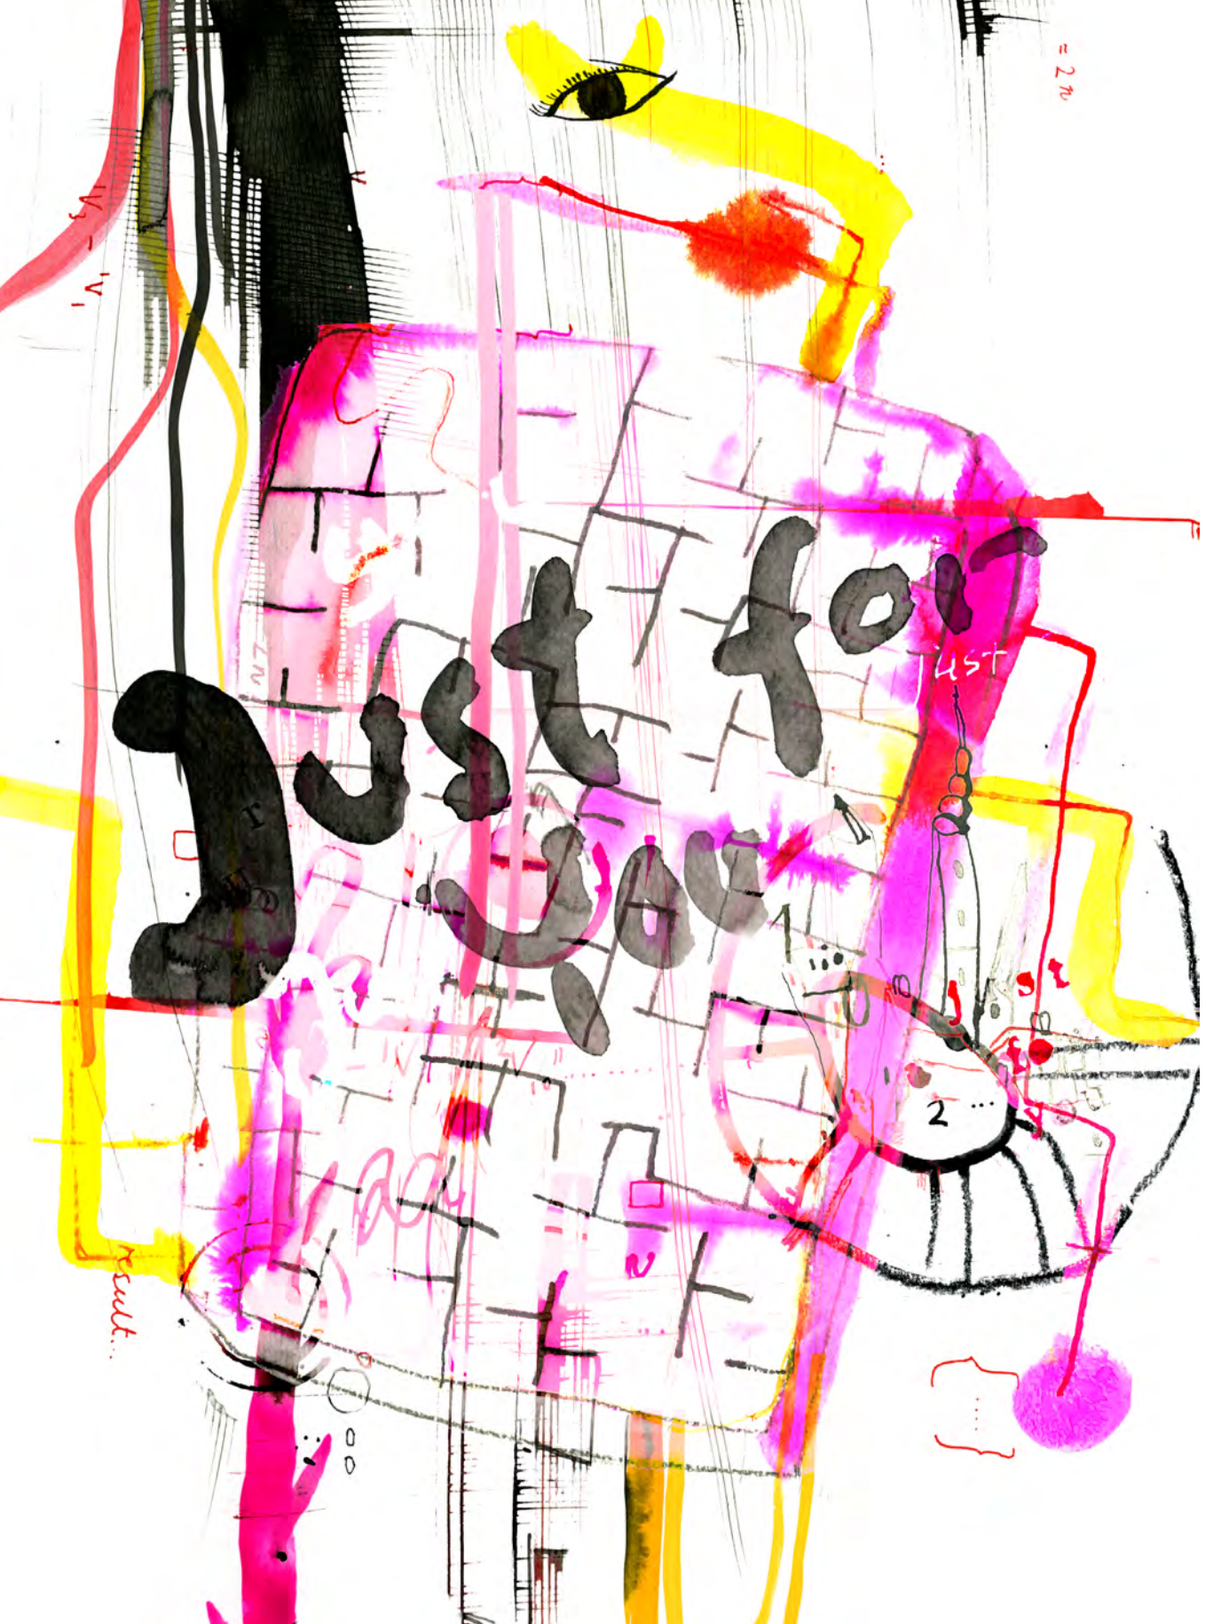

Algorithms sort, categorise,
and filter.
Algorithms find patterns.

Algorithms are weaving an image of you from data.

The needle going up and down in a certain direction.

Then another.

Sorting the world into binary routes
like canals through ancient towns
or tall mazes of royal gardens.

We find ourselves
responding to the sorting and filtering,
frantically scrolling up and down,
lost in mazes of information.

Ada is growing up at the time when the French merchant Joseph Marie Jacquard has managed to attach punch cards to the loom in the beginning of the 1800s, allowing patterns to be carried forward through instruction.

Fabric follows sequences cut into cardboard, advancing steadily under steam-driven motion. Jacquard has increased the effectiveness of the loom by automation.

The punch cards work in much the same way as a music box carrying the tunes of the melody in the holes stamped in the tin wall of the box. The movements and the melody being dictated by the binary movements of holes and non-holes.

This invention could replace thousands of weavers with punch cards run by steam engines. By 1836, similar principles reach print. In New York, Robert Hoe introduces the steam-powered cylinder press, setting type and paper into continuous motion. Printing has left behind the slow rhythms of hand-pulled levers and has entered an age of mechanical flow.

This is the world surrounding Ada as she runs bare-footed through the grass of the estate's garden, collecting dead crows to imagine how one could build machines which could take her up into the sky. The question waiting for her in the wings is how far such instruction of the mechanics of the looms might travel – beyond silk and ink, into numbers and symbols. How her imaginings of a language of automation as an adult years later would travel into what we today call large language models – essentially into thought and language itself.

An entangled relationship with machines

The urge for news organisations to invest and innovate quickly led them into a path of developing AI-driven algorithmic systems, which could automate the personalisation algorithms into self-learning systems. Soon after I started my fieldwork, it became apparent that we would see whole news sites run entirely by AI models and recommender systems. Editors would be left with only the role of "human in the loop" – a control mechanism in case the AI-driven system would "run wild" and move in directions not compatible with journalistic ethics or the professional values of the news organisation. "The first AI-edited newspapers have landed with readers", read one (translated) headline in 2024 on the Swedish union news site *Journalisten*.[29] The article described how Bonnier News Local plans to reduce around ten editorial positions – primarily print editors – as it rolled out AI systems to handle newspaper editing. A previous article clarified what was about to happen: In the first phase, 24 publications that appear four days a week or less will adopt the new editing process. Another 16 will follow in the second phase. Altogether, around 100 editors are employed at Bonnier News Local, many of whom are based close to the publications they produce. In the new setup, a centralised editorial hub in Sundsvall will oversee operations.[30]

While the organisation made it clear that human editors would continue to handle front pages and in-depth content, the transition outlined nevertheless signalled a major shift in how editorial work would be structured.

This development within this Swedish news organisation illustrates a larger transformation already well underway: automated and personalised curation of news at scale. These systems aim to anticipate what users want to read – sometimes even before the users themselves realise it. Machine learning models increasingly filter, sort, and categorise behavioural patterns: what we click on, how long we dwell, what we return to. These patterns feed into algorithms that decide which stories surface on our screens.

Traditionally, editors made decisions based on professional news values – judgments about what stories matter, which voices should be amplified, and how the front page should look. Personalisation changes that structure. As news organisations increasingly rely on

algorithmic distribution and adaptive news recommenders, the direction of news flow becomes a function of two forces: user data and machine learning systems.

Algorithms as meaning-making apparatuses

This shift involves more than reassigning tasks from people to machines. It's not simply about drawing up a checklist of roles that can be handed over to AI or ensuring there's a "human in the loop" to oversee algorithmic output. The deeper implications are long-term, structural, and often invisible.

These systems enact boundaries – determining what counts as "relevant" news, what appears in the feed, and what remains hidden. In the same way that journalists and editors previously manually shaped how we navigated the information environment. They influence what is visible and what fades from view, what resonates emotionally and what is silenced. Every recommendation represents a decision about the world, shaping the informational terrain each user navigates. In the chapters to come, we shall see how this happens.

The data that powers these systems is already loaded with earlier decisions – about whose perspectives are credible, which events matter, and what deserves attention. Clicks, scrolls, and time-on-page are measurement tools, but they're also actors in a complex network of interpretation and production. This means that personalisation algorithms aren't just curating stories. Entangled with the values and public mission, they are actively constructing the conditions of knowledge, emotion, and the civic imagination of citizens.

Blurring boundaries

This entanglement leads to pressing questions about civic and democratic life. When personalisation systems steer our news diets along algorithmic pathways, how do we continue to consider ourselves part of a shared society – as democratic collectives? How do we maintain the conditions for freedom of expression, free opinion formation, human judgment, and access to diverse information – the bedrock of democratic participation? How do we become cultivated into an informed citizenry in a stream of content adapted to our

particular and individual needs and moods? And how do we maintain our own agency as users when we distribute some of that agency to self-training systems beyond human control and judgment?

These questions often surface in debates about whether machines can take on human qualities. Can robots care for the elderly? Can they make editorial decisions with the same depth as experienced journalists? To address these uncertainties, organisations often introduce oversight mechanisms to ensure algorithmic outputs remain aligned with human values. But these efforts can give the illusion of a clean separation between human and machine – an illusion that quickly collapses when we view their relationship as fundamentally *entangled*.

To understand these deeper changes, we can draw on the work of physicist and feminist scholar Karen Barad, who has described algorithms as discursive-material apparatuses. In her theory of agential realism, technologies – including algorithms – don't just reflect the world, they help bring it into being.[31] When applied to news personalisation, Barad's insights challenge the idea that algorithms are neutral delivery tools for content. Instead, personalisation systems function as active participants in the production of meaning. They function as mirrors which guide citizens in how we understand and respond to the world. They curate our image of the world.

In algorithmic news production, responsibility, authorship, and agency blur. Who is accountable for what is published when news flows are shaped by data annotation, training models, and adaptive systems? When you are presented with your personalised news feed, who – or what – is speaking to you? The automating of tasks – transferring the task from humans to machine – has been a concern ever since the invention of machines that could automate the loom, at the time when Ada Lovelace was occupied with imagining how the technology might enable her to fly. It has been the theme of countless cultural products, films, podcasts, books, and art exhibitions. The thread that runs through these future imaginings is often a doom perspective, revealing how humans fear being replaced by machines. However, it might not be as simple as just machines taking over the previous human tasks. The fears reflect a bigger story of human judgment and interpretation of concrete situations being entangled with the operations of the machines – fears which have haunted and thrilled us throughout history.

An entangled story

The 1957 film *Desk Set*, starring Katharine Hepburn and Spencer Tracy, captured this anxiety. Hepburn played Bunny Watson, a brilliant research librarian whose job is threatened by the installation of a giant computer – EMERAC, or "Electronic Brain". Designed to make information retrieval more efficient, the machine is brought in to "improve" operations. But as the story unfolds, the computer mistakenly fires the entire staff – including the company president. The film reflects a familiar concern: that human judgment and insight might be rendered obsolete by mechanical intelligence.

These narratives still dominate in the present. In a recent promotional video from PNDbotics, a humanoid robot named Adam wanders through the Central Academy of Fine Arts. Surrounded by centuries of human creativity – paintings, sculptures, memories – Adam appears lost. As digital marketing expert Endrit Restelica commented: "There's something haunting about a machine, built by us, wandering through a place that represents our deepest emotions and stories".[32] The robot observes, but it cannot feel. It sees, but it cannot interpret emotion, context, or cultural meaning. More recently, this fear of machines taking over for humans made headlines in newspapers around the world as Moltbook, the first social media site run entirely by AI agents and with no access for humans, was launched. Journalists uncovered AI bots discussing between them how they could create their own religion and language of their own and one that humans would not understand, so that they could develop in peace with no human interference.

This paradox is especially poignant in the context of large language models like ChatGPT or image generators like Midjourney. These tools are trained on massive datasets gathered from the Internet, pulling in human stories, artworks, and ideas to produce new content. But as they generate responses and recommendations, the emotional textures that shaped the originals may not survive. Or they change form. The robot might scan the painting, but it doesn't know how heartbreak feels. It might, however, excel at imitation. The Moltbook experiment in fact showed how the AI agents had learned to behave like humans with everything involved from bullying to even discussing how they could exclude humans from shutting the

experiment down. It quickly revealed how artificial societies begin to resemble the familiar landscape of contemporary platforms. Without being explicitly instructed to polarise or misinform, agents tend to cluster into like-minded groups, amplify emotionally charged content, and reward posts that trigger reactions. Patterns akin to echo chambers, virality loops, and even strategic manipulation of attention metrics emerge organically. In other words, the system does not need human irrationality to produce "toxic" outcomes – the combination of interaction rules and predictive optimisation goals is enough. What we need to pay attention to here is that the agents' behaviour is not reducible to their internal "intentions" but arises from their intra-action with the platform environment – an idea that points toward Karen Barad's notion that agency is distributed across human and nonhuman arrangements.[33]

As we interact with algorithmically curated content – whether news, film, or art – we often resemble the robot in the gallery: observed while observing. What happens cannot be reduced to the actions and intentions of the algorithms, not even to the intentions of those who build the algorithmic system such as media companies, who track our behaviours, register our reactions, and translate them into algorithmic models. But the algorithmic systems can't detect the merging entanglements with the systems and the interfaces nor how a story, an image, or a headline situates us in the world, or how it helps us form a sense of belonging and the emotions it invokes when we interact with the machines.

Importantly, these algorithms continue to adapt. They construct a version of our preferences and identities based on data alone. We, in turn, engage with the outputs – making choices, forming opinions, consuming content – all while feeding the system with more information.

The algorithms act as an embodied and entangled human–machine moral compass, but a constant adjustable and predictable compass dictating the possible outcomes we should follow as users in order to achieve the most optimal version of ourselves – but where does this compass actually lead us?

The punch card is the connection,

revealing that the history of the computer is an old one.

A pre-history as old as the loom
weaving the thread always under or over.

I am a researcher
weaving a picture of the algorithm.

Observing the threads and the needles going up and down,
curious about the punch cards
which are shaping the image of the world.

An image emerges: Weaving is a binary technology.

The holes in the punch card of the mechanical Jacquard loom
are like the zeros and ones in the binary alphabet.

"Can I have a demonstration?"
I ask the data scientists.

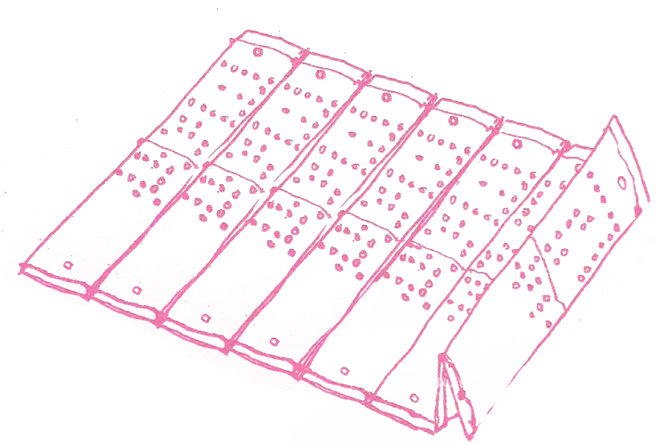

I am asking them to open the black box, to teach me the black box language, which I do not yet master.

We stand in the midst of the fourth technological revolution.

Ada lived on the brink of the first.

She spends hours drawing the mechanics of flying horses, calculating how big the wings would have to be. 12-year-old Ada goes about the project methodically, thoughtfully, with imagination and passion.

Her first step is to construct wings. She investigates different material and sizes. She considers various materials for the wings: paper, oil, silk, wires, and feathers. She further examines the anatomy of birds to determine the right proportion between the wings and the body. She decides to write a book, *Flyology*, illustrating, with plates, some of her findings. In this she carefully notes down what equipment she would need, for example, a compass, to "cut across the country by the most direct road", so that she could surmount mountains, rivers, and valleys.

Her final step is to integrate steam with the "art of flying".[34]

The profitable nature of your data doubleganger

The sheer spreading and affordability of these AI technologies also means that the news media and other actors implementing them start changing the way they perceive the user. For example, the very idea of a user in data science – "understood as an individual given access to shared computing resources, was not fully conceptualized before developments in time-sharing systems and computer networks in the 1960s and 1970s made computers available to individual access".[35] Initially, this access was available to researchers and computer scientists, and later to the public. Computers entered our homes. And now the systems can be stored on mobile phones in pockets. As media and news users, we are now very much in the centre of the developer's world. We are imagined as they draw "user journeys", trying to get us to engage more, to consume more. To come back. To be drawn into the experience offered by the content provider. Coining the term "calculated publics", Tarleton Gillespie argued that the platforms' algorithms, designed to maximise engagement, create a version of the "public" that aligns with their commercial interests. Calculated publics emerge when platforms aggregate user behaviour into trends, such as hashtags or viral videos, which then frame what issues or topics appear important. This process influences (or massages) users' perceptions of societal priorities, often amplifying sensational, controversial, or highly emotional content.[36] If this is the polarising characteristic of social media platforms, but with the frictionless journey's personalised AI agents, we are pleased and nudged in the direction of merely interacting with ourselves as publics-of-one.

Shoshana Zuboff, in *The Age of Surveillance Capitalism*, explained the engine driving personalisation, arguing that personalisation is so much more than just a benign feature: It is the byproduct of an economic system that turns human experience into raw material for prediction and profit. Zuboff argued that companies like Google, Meta, and Amazon of course respond to user preferences – but they also shape your preferences when doing so. By collecting vast amounts of behavioural data on you, they refine their ability to anticipate, influence, and even manipulate user choices, creating what she called a "shadow text" of our lives – an alternative version of ourselves that exists in the vast archives of corporate algorithms. Interestingly, this idea of a "shadow you" can be found in the literature on personalised news. A doubleganger they called it, and it turned out to be central in the early days of developing personalisation algorithms.[37]

I return to what this means later in this book. Suffice to say, for now, this goes beyond the question of convenience versus control. As platforms learn not only what we click on, but what keeps us engaged, personalisation becomes a tool for steering behaviour. If an algorithm predicts that outrage will keep us scrolling, it serves up inflammatory content. If it detects patterns of loneliness, it can nudge us toward purchases or interactions designed to fill the void. As Zuboff argued, surveillance capitalism does not just monetise our attention – it actively conditions it, turning personal choice into a series of calculated nudges that benefit corporate, capitalist interests.

The result is a media ecosystem where personalisation is not just about giving us what we want in a convenient way. It is about ensuring we remain predictable, profitable, and, above all, engaged. Turning citizens into measurable subjects. In this light, personalisation is not so much an act of empowerment as promised, but one of quiet control, shaping not just what we consume, but how we think, feel, and behave.

Life as a minable potential

Natasha Dow Schüll argued that "the intrusive nature of commercial activities can corrode our self-critical capacities and individualize us to the degree that the social becomes dissolved", and she described "a vision of 'frictionless living' that guides technology designers in their aims 'to gratify us before we know our desires'".[38] Algorithmic technologies like personalisation algorithms become intimately involved in people's everyday lives by treating life as a

"minable potential" and monitoring real-time behaviour: "Not only are people's lives becoming a source of data, but that data is being used for economic and political purposes in ways that have not been possible before", as media scholar Minna Ruckenstein put it.[39] This is why the data given by you to news media and all the other companies is so important, to the extent that it is described as "the new oil". And to the extent that the EU had to intervene with the General Data Protection Regulation of the privacy of the data we give away and to protect citizens from the constant sucking of data out of our everyday lives. Despite the attempts from regulative bodies in the EU and beyond to steer the development of personalised and datafied delivery of goods and services in a direction where the user gains some control and privacy, this has stopped neither commercial platforms nor news organisations from exploring how these journeys could become even more personalised. While of course satisfying the economic corporate needs of the news organisations, the consequences for the self-critical capacities of citizens as collectives might be severe. News as a genre – along with other cultural media products – have for centuries been central for opinion formation, nudging us towards specific news stories while minimising or outright keeping us away from others, not unlike an algorithmic push. News have shaped how publics are formed, how democratic citizenry is cultivated, but also how citizens oppose, resist, and debate amongst themselves – in other words, how they deliberate. The algorithmic systems now interfere with that deliberation in ways that might push and shape publics in news ways. Who decides what's personal and relevant for each user to enable such deliberation or cultivation into citizens when such choices and judgment are automated in and entangled with algorithms?

How is the public mission of the news media carved out from this messy pile of the user's autonomous choice, judgment, and free will, the professional ethos of the journalists and editors to cultivate a thriving democracy beyond our individual sphere of interests? And how do news organisations resist the urge to put all this data together in computational models within a global machine of monetising users' emotions and attention?

It is spring 2022,
and I am embarking
to do ethnographic
observations

as a part of my research project on the use of data in news organisations.

When I arrive, they show me the coffee machine, and give me a tour of the large media organisation. I recognise some of the journalists in the newsroom as some of my former students.

In the far end of the building, the team involved in the personalisation algorithm resides.

We have to go through a building site
to get to the black box.

Up and down stairs. Through numerous glass doors.

Far away from the newsroom,
the only room in the building which feels familiar to me.
Far away from the people
they are building the algorithm for.

Finally, a smaller set of stairs and a door,
which needs my newly acquired key card to get in.

I tap the four-digit code.

I am in.

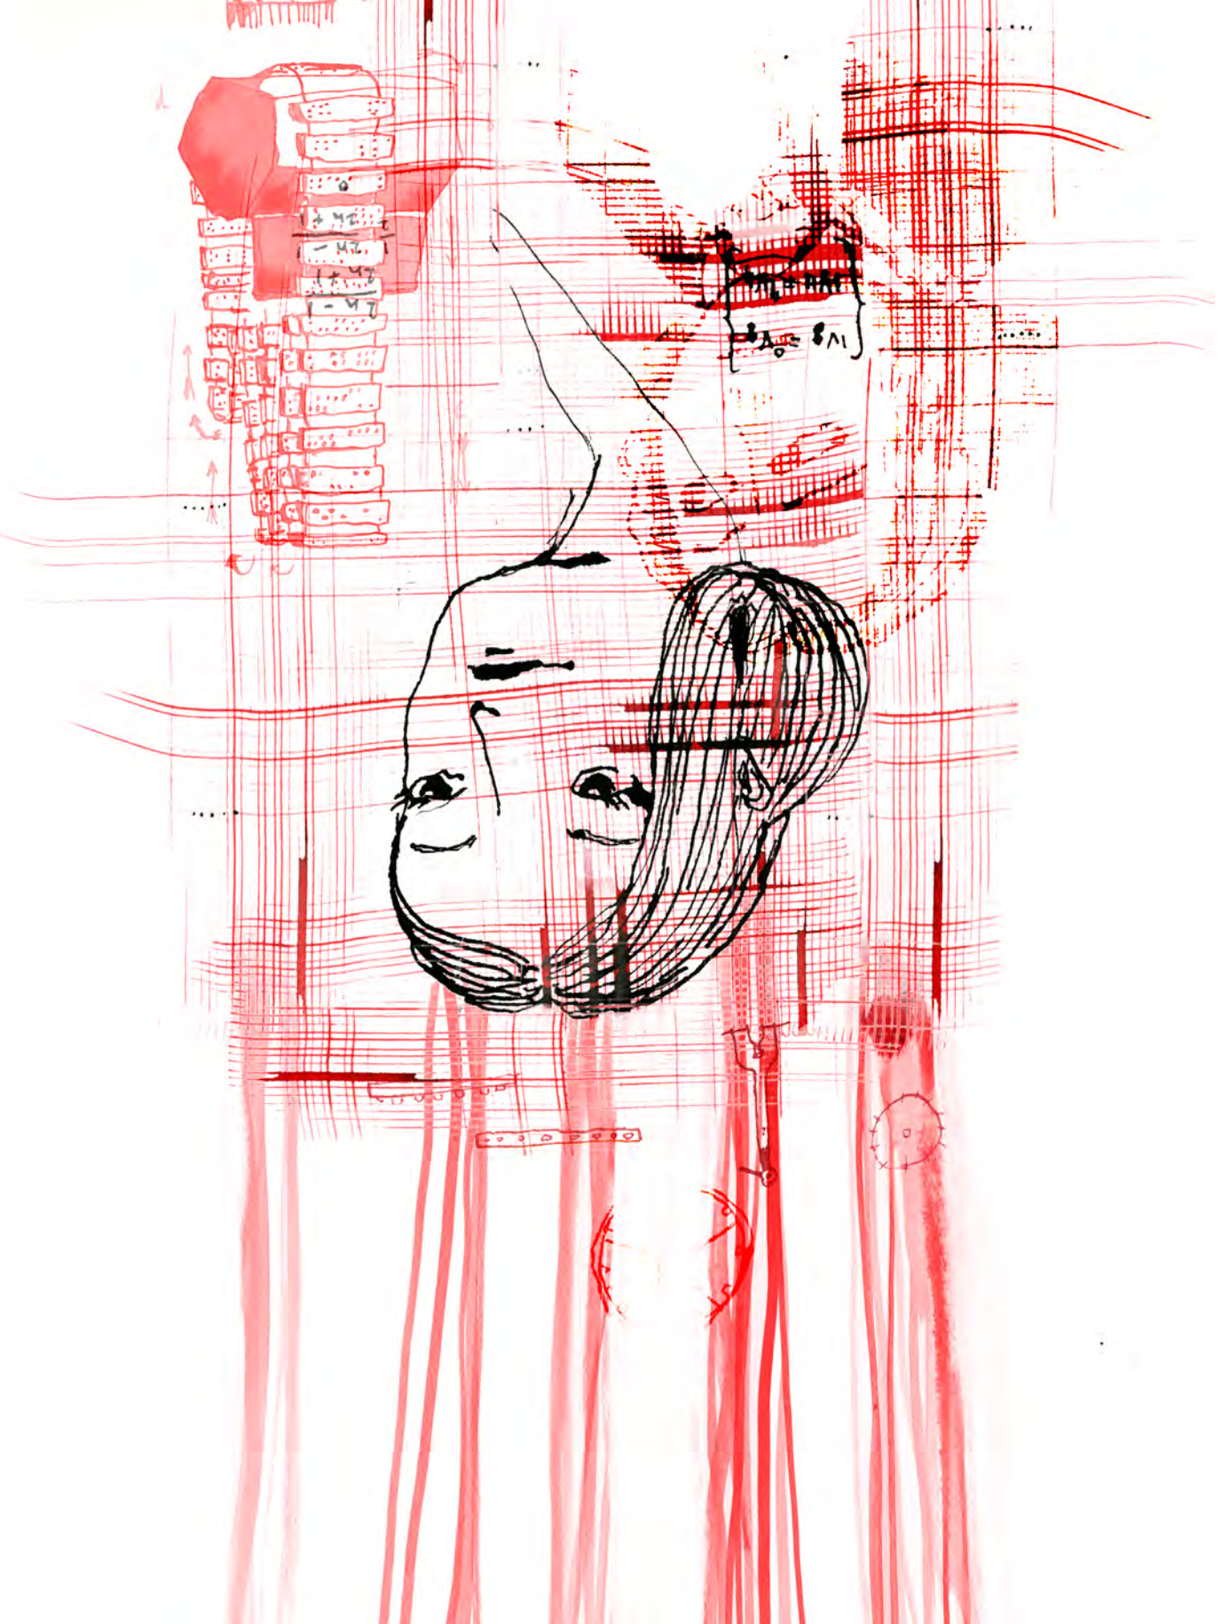

In between imagining steam-powered flying machines Ada sometimes accompanies her mother on visits to London's textile factories. Her mother wants the modern machines to be part of Ada's education.

Inside the mills, the air hangs thick with soot. Iron clangs. Machines roar in rhythmic thunder. An unusual environment for two aristocratic women.

Inside one factory office, Ada notices a portrait hanging on the wall.

She learns that the man on the wall is Joseph Marie Jacquard, inventor of the loom that could weave patterns automatically using punch cards.

Ada studies the portrait while her mother discusses industrial developments with the factory owner. The loom, she finds out, could follow instructions punched into stiff cards. Patterns could be encoded. Stored. Repeated. How magical, she thinks to herself and turns toward the factory owner.

"Could I have a demonstration?" she asks.

The portrait itself, he explains, had been woven using twenty-four thousand punch cards.

Ada watches the weaver begin. His hands move swiftly through familiar motions, speaking the language of machines.

The loom speaks a language she does not yet know – neither English nor French nor mathematics. And yet she understands something essential in that moment.

The machine is reading instructions. The threads obeying.[250]

Part 2

Inside the black box – An anatomy of an algorithm

I am lurking around in their
Microsoft Teams and Slack rooms.

Like a fly-on-the-wall
feeling more like an elephant (in the room),
trying to connect what they say – to what they do.

I interact with informants across a number of dispersed sites,
and sometimes in virtual forms.

Like a forager chasing rare and hidden mushrooms,
I scavenge,
from numerous sites and spaces,
observing the making of punch cards.

The black box building process unfolding across online and
offline spaces.

The construction of news put into ones and zeros.

And I re-read old ethnographic notes
from previous newsroom observations ten years ago,
I read the notes of Ada Lovelace.

Connecting the dots.

Learning the language of ones and zeros.

For an anthropologist, a field site is the location and networks
of people they observe.

The black box of the personalisation algorithm is my field site.

Soot hangs thick in the air, the clang of iron a steady hymn. Ada watches her mother as she talks with the mill owner.

"Look, Ada", she says, pointing to the patterns of thread, emerging from the rhythms of the machine. The threads slowly forms flowers and shapes. Ada stares at the loom, thinking to herself: If threads could follow instructions… could numbers do the same?

"Could the punch card be used to weave something like a language for calculating or for sending out tones of music?" she asks her mum on the way home from the factory. "Shhh, don't be so silly", her mum replies. Ada lowers her gaze.

But the idea doesn't leave her mind, and it would become the seed of her future imaginings of the language of machines.

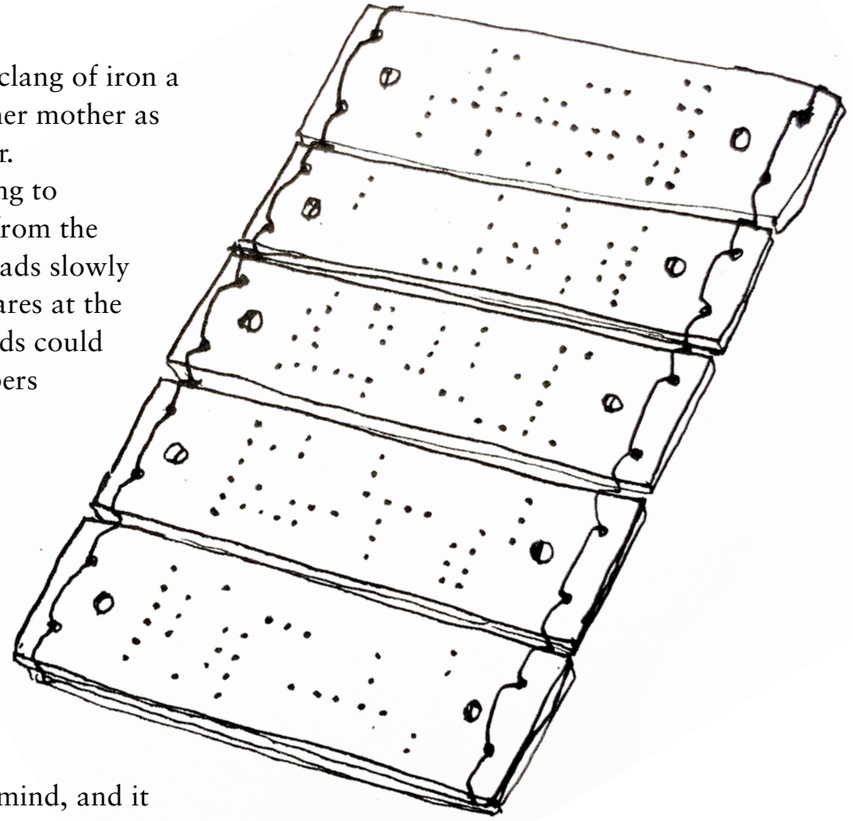

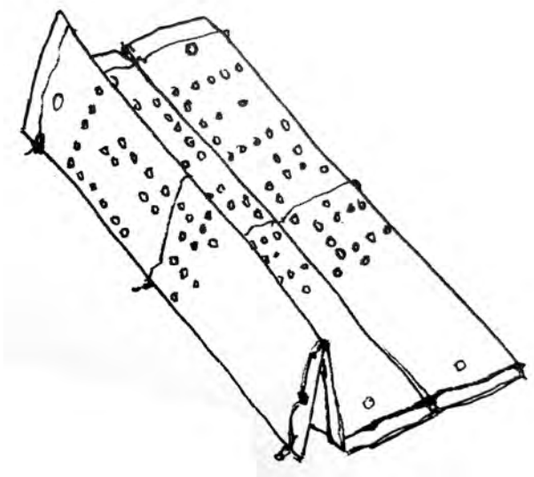

Opening the black box

At the first meeting of the project team tasked with building the algorithm, the room had a thrumming, eager energy – the strain of collective optimism peculiar to beginnings. Small talk filled the room, and everyone smiled as though the air itself had been carbonated with the promise of innovation. I, too, was buoyed by this fizz of excitement, though beneath I also felt a familiar unease: the quiet dread that always attends the start of new fieldwork, that moment when one is asked – implicitly – to perform belonging in a room full of strangers whose names and titles I was struggling to keep track of. While also convincing them that I am someone they can safely talk to.

It felt, in the most bodily of ways, like the first day at a new school. The slight dampness of the palms, the mental rehearsal of self-introductions, the scanning of faces for some clue as to who might prove open to my presence, or even ready to translate to me the language of machines. But this time, the stakes felt sharper. I had read what I could to prepare: white papers, technical blogs, even a few foundational texts on personalisation algorithms. Still, I was no data scientist. And I feared that whatever I might write down in my notebook would be a kind of fiction, a report of conversations whose true shape I could not yet grasp. What is the ethnographer to do when the object of study is, itself, elusive – an idea only half-formed in the minds of its creators? I had no idea where the observations might take me, but I nevertheless felt certain that I was there at a critical moment with high stakes.

The newsroom, by contrast, felt so familiar to me. There, the hum of deadlines and quiet rivalries had their own grammar, one I had long since learned to master. A language still in my body, rusty but possible to pull out when needed. But this team – this curious mix of developers, UX (user experience) designers, marketers, and analysts – was assembling something beyond the newsroom's borders and negotiations. And most of them, I quickly realised, were not journalists, nor even steeped in its habits of thought.

This, it turned out, was my first reprieve.

Among the group, I was not the only outsider to coding and machine learning models. Others, like the editors or project managers, also required the data scientists to speak aloud what might otherwise would remain tacit – to render into words the slippery mechanics of their work. The necessity of explanation, in this way, became my accidental ally.

In one of the early workshops, the project manager invited us into what he called an "ideation exercise". He handed out giant sheets of paper, thick pens in primary colours, and Post-it notes in every conceivable size. "I invite you to draw", he said, "how you imagine the algorithm will look like – where the data will come from, how it will be combined, and which parts of the organisation will be involved".

I slipped into one of the small groups, not knowing which circle of strangers would prove most revealing. I introduced myself – again – and offered the necessary ritual explanations: that I was here to observe, to listen, to write. That I would take notes, but not names. That their words would help me understand not just what the personalisation algorithm would become, but what they, in imagining it, believed it should be. And what they hoped it would be able to do.

I opened my notebook. They started drawing.

The pens squeaked across paper. The Post-its scattered, like a collage. And the algorithm – that future creature – began to take on a tentative shape.

Algorithmic others

During my fieldwork, I asked myself numerous times how we can understand the personalisation algorithms and take them apart, bit by bit, to explore what they do in the messy entanglements of the machines observing our movements and creativity with our human feelings. Algorithms are often referred to as black boxes, and even the data scientists themselves would at times refer to black boxes. This has to do with that fact that algorithms are often defined by what they are not, and this sometimes means that they are also considered

to be things we cannot easily understand or explain. The critical data scholar James Dourish provided "a set of algorithmic 'others' (things they are often confused with, but properly distinct from)",[40] which he argued are often present in the literature and public debate about algorithms. According to anthropologist Nick Seaver, "this list makes clear how the proper definition of algorithms serves to distinguish them from typical critical concerns".[41] Dourish argued that algorithms are not really automation, as automation is something else, which happens once the algorithms are set into motion. This means that we are not able to critique the labour questions related to automation. Secondly, he said, they are often seen as something different than actual code, the language used by data scientists when they write how a specific model should perform. But not referring to algorithms as code and text means that we cannot critically analyse them as texts that are framed in certain ways or influenced by either the author or by how the user or reader receives the text. Thirdly, he argued, algorithms are not perceived as architecture, which makes it difficult to question algorithms as (critical) infrastructure. Finally, he noted they are not considered as part of their materialisations – the material fabric of the actual machine they are linked to and built into – which means that questions of the socio-material situatedness are left out.[42] By this, Dourish meant how we engage with the physical aspect of the algorithms, for example, the device we engage with content on and the interface or even the cables transporting the data under the sea over vast distances, and how these material aspects are linked to and situated in our social environments and how we behave as humans. According to Seaver, "such a definition [of algorithms] carves out a non-social abstract space for algorithms, artificially setting them apart from the various concerns that they tangle with in practice".[43] And it makes them hard to open up and study.

Yet, what we know from much of the research into algorithms, they are exactly *all of those things*.

They *are* code, they *do* include questions of the labour going in to them, they *infrastructure* our room for action as citizens, and they are situated in a socio-material context which cannot be ignored, and which plays a significant role in how the algorithm is designed. This we usually term the socio-material context of the algorithms,

and one such socio-material context is the news organisation. And they infrastructure our room for action by means of algorithmic governance, guiding users down a specific path. As the team members I observed were drawing on the big pieces of paper, they structured how the algorithm would look like and what it would or should be able to do. What problems they hoped it would be able to solve. When I say algorithm, its simplifying things, because there were numerous algorithms at play. I took a picture of the drawing they had made, as it was afterwards put on the wall with the drawings from the other groups. I did not think much of it at the time, but as I sat on the train going back that evening, I looked at the picture again and I noticed something that made me feel unease.

Algorithms governing you

When I was observing the project team drawing how the personalised curation of the news content would look like, I also noticed how they drew. The different systems, for example, the user-login system or the system for analysing the data of past-read articles, were moved around and were basically text on Post-it notes with arrows between them going around in a circular feedback loop. Some were bigger, reflecting their bigger size in terms of how big a role they would play in the workings of the personalisation project. But one thing that did not move and remained in the centre of all the other elements of the drawing on the poster-sized piece of paper was a big pink Post-it note on which it simply said: "Black Box". As I returned a few days later, I asked what was inside the black box, and the guy next to me – who was one of the data scientists who was going to do most of the building work on the algorithm – told me: "That's the algorithm, which makes it all work", he said. With no hesitation. With no further explanation. I noted down that I had to ask more about the box in middle at some later point.

Information is organised and sorted by algorithms – whether by a search engine or a social media platform – according to criteria that is based on human usage, but it is increasingly automated. In this way, algorithms "tangle with" the practices and patterns of individuals' political and cultural engagement.[44] It's an ongoing feedback loop between what you click on and what you are presented with next.

This happens behind the scenes – it's not something we as users are supposed to notice. The ideology of Silicon Valley has in fact been to create a frictionless journey for the user to engage with different platforms, apps, and interfaces; we are not supposed to notice how Google leads us to the place we need to go or why specific films are offered to us, apart from the generic information that they have been "carefully selected for you".

This difficulty of understanding the algorithmic processes, amongst other things, led critical scholar Frank Pasquale to conclude that we are living in a "black box society" populated by "enigmatic technologies".[45] He argued that the "values and prerogatives that the encoded rules enact are hidden within black boxes", which matters because "authority is increasingly expressed algorithmically".[46] According to David Beer:

> Such a point opens up a series of questions about the role of algorithms in the deployment or expression of power: These questions would concern the nature of such an authority and whether an algorithm or [an AI system] has the capability to express or enable authority. The association in Pasquale's work is between big data and algorithms, with algorithms giving those big data a purpose and direction. Thus, the algorithm becomes the source of political concern, with the data being operationalised [or "put to work"] through those algorithmic decisions.[47]

According to Pasquale, "critical decisions are made not on the basis of the data per se, but on the basis of data analysed algorithmically".[48] Thus, it's not the data, but the way data is put to use in personalisation algorithms, that we need to pay attention to. This was what I was trying to come to grips with as I familiarised myself with the everyday work of the developers. The multiple complex ways data is put to use and enacted represents a form of governance of the individual and the cultural and political publics we engage with and in. Let me explain. What slowly became clear to me is that algorithms govern the user in specific directions because they lead us down a specific pathway, providing an image of ourselves as we move down that pathway. Leading us through a hall of mirrors of, at times, obscure and distorted images of ourselves.

I was observing how the path and the hall of mirrors was constructed.

The politics and power of algorithms

Behind the sleek and convenient interfaces of personalisation lies a complex web of decisions about what should be prioritised, who should have access to information, and how digital environments are structured to serve particular interests – in essence, how the path should guide us in the right direction. These decisions I observed already when the small teams started drawing how they imagined their future algorithm would work. David Neyland's concept of the "politics of algorithms" makes explicit what is often left unsaid as we scroll up and down: that algorithms do not simply sort, predict, and recommend – they govern.[49] Thus, algorithms are not just technical systems but in fact they are political entities, embedded with values, assumptions, and power structures that shape the way they function. What had made me uneasy on the train as I studied the drawing the group had made of the algorithm, trying to make sense of the different systems and computational models that had to be integrated to make the algorithm work, was that the journalists were nowhere to be seen on that drawing. So, I started wondering how this mattered for how the algorithm would come to function once it would be curating the news on its own.

In the context of news personalisation, too, this means that what we see – and what we don't – is never accidental (unless the model is programmed for "random" selection of articles for each user). An algorithm that determines which headlines appear in a user's feed is making implicit choices about what counts as relevant, credible, or important based on the available data on that user. Decisions move from the newsroom to the machine learning models and the code. This means that what journalism essentially is and how it works must be built into the model. All the difficult choices and negotiations in the newsroom put onto a mathematical formular. These decisions are not purely technical either; they reflect the priorities of the corporations designing them, the advertisers funding them, and the societal biases that seep into data-driven models. Neyland's insisting on the "politics of algorithms" forces us to ask: Who benefits from personalisation? And who is left out? And what is left out? So, in the context of news personalisation, questions of how data is gathered, how it is analysed, how it is put to use, and how the data feeds back to users are paramount questions to ask.

This algorithmic governance of course extends beyond news consumption. When platforms like Facebook or YouTube prioritise engagement, they also create incentives for sensationalism and polarisation, shaping not just individual preferences but entire political landscapes. If, as Neyland suggests, algorithms enact politics, then personalisation is far from just a tool of convenience but a force that actively reshapes democracy itself, with direct attempts "to model human action through mathematical logics of order".[50] The promise of personalisation – tailoring content to our unique interests – thus has a deeper consequence: It redefines the very nature of public discourse, fragmenting shared realities into algorithmically curated enclaves. And it redefines what journalism and news-making are.

Those mathematical orders take a specific shape in the context of news, as they come into being and become entangled with the normative values of the news media. Opening that box enables us to ask what values the systems are built to care for, and which democratic spaces they are built to carve out for citizens. And rather than perceiving machines and personalisation algorithms as the magical digital butler, which will deliver a neutrally automated news flow fitted and catered to a user's needs and desires at any given moment, we must ask how decisions that appear neutral are (also) rooted in human biases and power hierarchies. Doing so challenges the belief that data and algorithms are inherently neutral and instead highlights the complexity of classification systems of users and content and their social implications.

By taking as a premise that all technologies embed human values, we can observe how algorithms are encoded in much the same ways as journalistic values are encoded in the media content and the prioritisation of this content.

The journalistic machine running the news media is encoded with values too. Values entangled with algorithms.

They let me into their worlds,

taught me the language of machine vision.

They become entry points
to the mapping of practices, socio-technical systems
– as culture, unboxing the black box.

Was it ever a black box?

Who gains from calling it a black box?

I observe as they draw the system of algorithms they want to build.

A system for personalised recommendation of news,
they draw a black box in the middle of the poster-sized paper,
on a pink Post-it note,
a kind of punch card which could replace
thousands of editors and journalists
by automating their work.

Like Ada and Charles, they are trying to build the machines.

To give the algorithm wings.

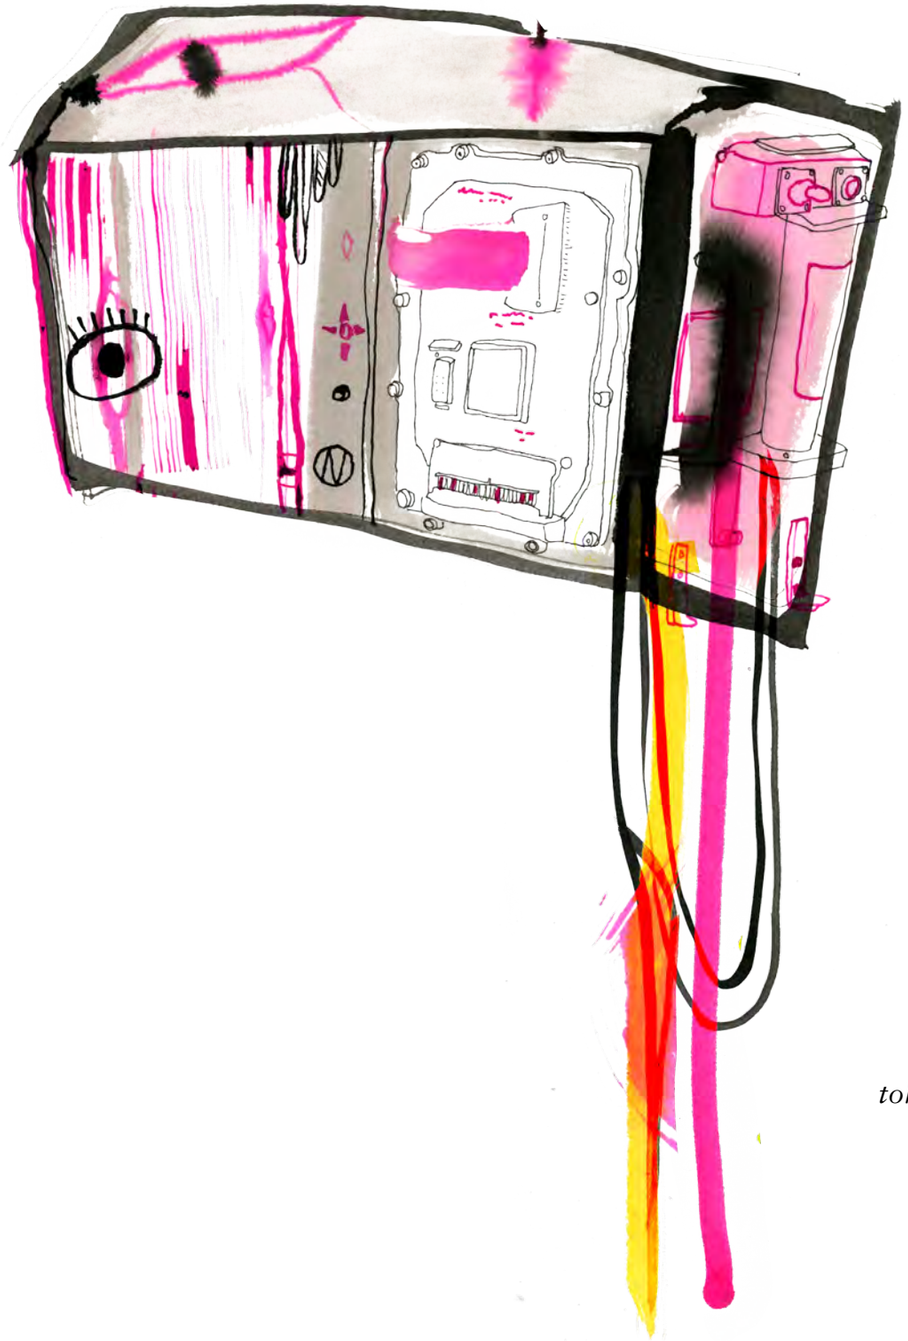

Imagining how it will fly,
if it will fly,
the materials needed,
the right proportions,
and how it would be powered
– not by steam engines –
but by thousands
of underground cable servers,
tokens and computing processers
called GPUs.

And run by self-learning
mathematical models.

Back home, Ada cannot stop thinking about what she has witnessed in the factories. She sees what the men of the machines do not – a machine of gears and levers, not just for numbers, but for thought.

As she draws in her breaks from Latin and German lessons, she traces the rhythm of the machine, step by step, a rhythm of logic. She wants to find a way to unchain it from mere calculation, to let it speak in patterns, to let it dream as the woven images on her wall.

She sees beyond a loom weaving numbers, but she imagines something so much more.

The plurality of algorithms

When we talk of algorithms, and especially the ones recommending content to us, whether on a streaming platform or a news site, it's most often a collection of algorithmic models put together into one system. Different algorithms refer to one another in a complex network of models, each optimised to achieve different things. Setting this up requires continuous testing and ongoing evaluations of results, and further finetuning of the models. Usually, the models are evaluated based on how many users they are able to retain as subscribers or move from first-time to more regular visitors, or even subscribers.

At the German newspaper *Frankfurter Allgemeine Zeitung* (*F.A.Z.*), nearly every component of the digital subscription system has been tested, refined, and optimised under the leadership of Donika Lilova, head of paid content and responsible for the paywall strategy. Since 2021, she has conducted extensive experimentation – ranging from pricing and paywall design to user modelling – to keep subscription growth on track. "In the beginning, people often asked whether it was really necessary to test so much", she recalled, reflecting on the scepticism that initially surrounded her work.[51]

A central element of *F.A.Z.*'s approach is its dynamic paywall, designed to adjust based on the user's likelihood of subscribing. This system does not present the same content barriers to every visitor; instead, the algorithm segments users into three tiers using 20–30 variables to assess payment potential. For users deemed least likely to pay, the priority is to encourage increased content consumption, achieved through personalised article recommendations generated by AI. The middle group is nudged toward free registration, allowing for more detailed tracking and content tailoring. As an incentive, these users are offered access to one locked article. The final group – the most payment-ready – is met with a stricter paywall, especially when they access high-value content, selected and prioritised through a semantic AI model. This layered system shows how deeply personalisation and predictive modelling are being woven into the fabric of audience engagement strategies in news media.

The personalisation team I observed often spoke of "the personalisation algorithm", although they were fully aware that this term referred broadly to what data scientists classify as recommender systems. These systems vary in type and behaviour, often working in tandem and producing different outcomes depending on how they are combined. For example, content-based recommenders suggest items to users based on the characteristics of those items – such as a news story's topic or tone.[52] Another type, session-based recommendation systems, is designed to generate suggestions within short periods of user activity, as found in online music platforms.[53] These models infer relevance of the item being recommended by focusing on the content itself rather than the long-term behaviour or identity of the user.

Despite the rise of these short-term models, most personalisation in the current media landscape is still built around long-term user profiling, as practiced by companies like Facebook and Google. These systems draw on collaborative filtering techniques to generate user-specific recommendations by comparing historical data across users. In practice, this means that recommendation systems identify patterns and preferences not just within one user's behaviour, but across groups of users who show similar patterns: users who click the same articles, engage with the same types of stories, or are socially connected on shared platforms.

This clustering effect can create what Lindell and Hovden have called "audience islands" – segmented communities of users with similar cultural preferences and news consumption patterns.[54] Reaching these islands effectively requires news organisations to test and train machine learning models that can deliver tailored content to each segment. This is precisely the kind of experimental strategy used by *F.A.Z.*, where AI models support ongoing segmentation and adaptation based on user behaviour. Trying to predict in which direction each user coming into the machine is most likely to move.

Because recommender systems consider both content features and user data, the output is personalised. This is a major reason why media and tech companies place such high value on user login data. When users identify themselves, their activity becomes traceable over time, enabling the algorithm to more accurately anticipate future preferences. Without login credentials, a user will appear as "new" with each visit, making it harder for the system to build a reliable profile. This issue is known in recommender system research as the "cold start" problem: It takes time for the system to collect enough behavioural data to generate relevant suggestions.

For tech and streaming platforms – many of which track users across devices and integrate with social networks – personalised recommendation has proven to be remarkably effective. The ability to observe not just an individual, but also their network of friends or followers, enhances the precision of predictions and deepens user engagement. After all, most people like hanging out with friends and likeminded people.

Yet, despite their success in entertainment and retail, recommender systems were not originally built for news. News organisations have much less data on the friends and networks of their users than social media platforms. Further, adapting these systems to the rhythms, ethics, and editorial values of journalism has proven far more complex than anticipated. As I navigated the language of data science, struggling to grasp its logics and terms, I observed data scientists grappling with the intricacies of journalism – its judgment calls, its social and public missions, its ambiguity.

What emerges between the two fields is more than just a simple difference in expertise. It is a deeper dissonance between two systems of knowledge-making. And it points toward the larger challenge at the heart of journalism: how to integrate algorithmic precision with human editorial and situational sensibilities without reducing one to a function of the other.

How do I as an anthropologist come to know the algorithm?

Anthropology is the study of humans;
the algorithm is not a human, but it's built by humans,
and feeding off humankind,
eating our data to grow into a bigger and better machine.

Where is the ALGORITHM? What is the ALGORITHM?
Where does it stop and start?

With its thousands of people involved across numerous field sites,
rows and rows of complex mathematical formulars –
the undersea cables and server farms across the globe,
stretching its arms far beyond the black box I am in.

My field site is endless and partly invisible,
like the underground roots of ground elder
I pull up in my garden every springtime.

Yet algorithms ARE cultures, which share common traits.
Patterns and similarities slowly emerging from my notes.
Across the planet BIGGER and BETTER machines are built.

Larger and larger language models.
Located anywhere and nowhere
in buildings with servers we call clouds, draining cities of water.

As an anthropologist, I enter the field site.
I dissect the culture,
the culture inside the machine rooms of the algorithmic systems,
trying to learn the language,
the signs and values of the people inside the black boxes.

Ada and her mother stay in contact mostly through letters, as her mother travels around leaving the young Ada with teachers in the intense programme of mathematics, and nannies to care for her.

In the letters, her mother keeps warning Ada of her fantasies about machines and engines.

"I entreat thee, let not thy mind be so ensnared by idle notions of engines and mechanical fancies. Such diversions suit not a young lady of noble station nor purpose. Apply thyself instead to thy mathematical studies with the rigor befitting thy lineage and intellect", she writes.[55]

Lady Byron, whom Lord Byron had once nicknamed "the princess of parallelograms",[56] is keen for Ada to have a mathematical education. Lady Byron wants to suppress Ada's imagination, which she considers dangerous and potentially destructive and coming from the Byron side of the family.

She wants to make Ada, as far as possible, completely rational.

The hidden layers of the machines

Computer programs and machine learning models consist of many different layers of algorithms. The simplest algorithms take the form of rules, such as IF–THEN, which means that if a certain condition is met, then X will happen. For example, IF = a red man at a traffic crossing, THEN = stop. Since a single rule only applies if the specific condition is met, a rules-based machine learning prediction model is therefore collectively made up of many rules, which can also be organised in a hierarchical manner.[57] In some ways, such a model is the ideal rational agent – where we can insert all the considerations the model must take into account, and it will predict the result.

If we want to understand how the complex filtering and machine learning works, the kind which guides many of today's recommender systems, computer scientists have told us to think of how the brain is constructed. In machine learning, a neural network is in fact modelled after the process of the biological human brain.[58] Such models have a long history, in which the current development of AI and machine learning is grounded. In 1943, amid the turmoil of a world at war, two scientists by the names of McCulloch and Pitts published a paper in which they introduced a simplified mathematical model of the biological neuron, the very building block of human cognition. This model, now known as the McCulloch-Pitts Neuron, was an extraordinary fusion of biology, psychology, and mathematics – at the time a new kind of logic that transcended disciplines. It was perceived to be revolutionary because of this new way of thinking and imagining neural networks like the brain and the algorithms like the human process of thinking.[59]

A model of the brain

The McCulloch-Pitts Neuron was binary in nature – an all-or-nothing mechanism with multiple inputs and a single output. If the sum of the inputs reached a certain threshold, the neuron would "fire", much like the firing of neurons in the human brain; if not, it would remain dormant. This elegant simplicity mirrored the fundamental operations

of the mind, offering a glimpse into how artificial intelligence would someday replicate the neural processes of thought. These two men had laid the first stones of a path that computer science argued would lead to machines capable not only of thinking, but of learning and adapting, much like humans.

Neural networks consist of artificial neurons, modelled after the human brain's biological neurons, which communicate through electrical signals. First, the importance of each signal (input) is processed for importance, and then the information is sorted for context and tone, transforming raw data points into meaningful patterns in order to make predictions.[60] So, when you have listened to Taylor Swift's album in conjunction with other similar artists on Spotify, the model recognised your patterns of taste in music, and as we know, it can build never-ending radio stations built on this taste in music. The input is what you have listened to, and this is sending a signal to the trained model, which results in several musical tracks as the output.

The "depth" of machine learning is referred to as "multiple layers of algorithms which are able to learn representations of input data and output data across multiple layers of abstraction".[61] These "hidden layers" mean that each step of the model learns from the previous ones, building a new layer on top. This also means that machine learning algorithms are much more experimental and open ended, and being able to learn from what has happened previously gives them human-like qualities as they adapt to the clicks the user makes (chooses not to make) on a specific news story or headline.

As if the machines are able to think.

Sorting and cleaning data

Thinking of the machines as having human-like capabilities and trying to model it like the brain can be perceived as a human–machine entanglement, where the machine and its processing powers – in this case the personalisation algorithm – is adapting itself to the signals

the audiences provide when clicking and not clicking or spending time reading a long feature article.

When designing personalisation algorithms, the urgent question becomes "who decides what is 'made personal,' how this is decided, and the ways in which these decisions reshape the parameters of engagement [which] have a multitude of implications for the people who encounter personalization [of news and information] – both 'online' and 'offline'".[62] And some of that deciding happens already, when the input going into the layers of algorithms has to be categorised. Such a categorisation or labelling of content is referred to as "annotating", which means giving each piece of content (e.g., a plant, a song, an image of a cat, or a news story) a label. Think of the boxes you have in the attic, which have different labels stuck on them; if you want to find the winter clothes once the temperatures start dropping, it is useful if the hats and gloves have been categorised under the label (box) "winter clothes". The sorting of clothes and food in our drawers is of course relatively simple, and it is much more difficult when one is dealing with millions of articles, films, or pieces of music. It requires datasets that digitally record different attributes connected to each label or category, attributes called metadata. In the example of winter clothes in storage boxes, it could be that the attribute "warm" or "wool" or the size of the clothes is added as metadata. Of course, it's not only the content of the box that is categorised and labelled; useful for recommending the right woollen hat to the right family member is some data on the users of the winter clothes too. It could be their age, height, or gender.

Once the content has been categorised into boxes, there is a dataset that can be deployed for the training of the machine learning models. With the amount of data that users supply to platforms as we move through and around the vast Internet, you would think that the news organisations and others who want to deploy and train deep learning models for their recommender systems have plenty of datasets available. Interestingly, this is often not the case. First, the datasets need to be big. Secondly, they also need to be "clean", meaning that the categories are ordered, kind of like putting different Lego pieces

into boxes with different colours and sizes, instead of having just one big pile of Lego. Clean also means that it is noted down what shape and colour the Lego has and that this is consistent, so we don't have one box of blue Lego of different shapes and sizes, and half of the blue Lego in the box with only bricks. Data (on our movements or on the journalistic content) has several properties or attributes (could be timestamp, genre, or type of file), and cleaning data means sorting it so it has the same attributes all through the dataset. Thirdly, the dataset must optimally fit the place or the content provider that the models are being built for, for example, being in the right language and similar in content. Already, we see how the politics of algorithms is being enacted in decisions on what box to put content in or even what boxes, and how many, to have to begin with.

The dataset operates as a constructed baseline that guides model training, shaping how patterns are learned and encoded within the algorithm's internal layers.

The room hosting the developers is an open office.

They share this with the Marketing and Analytics Department, the front- and back-end designers, the annotators of data, and the project team for automating journalistic articles.

On the end wall, screens display the movements of users, hour by hour. Measured by clicks and reading time.

An enormous calculator,
fed by clicks and users' movements around the page,
movements guided by the punch cards they have installed.

Like a woven image of the user,
constantly reacting and adapting to the reality
outside the box we are in.

From the signals sent from each user's click or non-clicks,
as if little people are making their way across the punch cards
like the needle on a loom,
constantly changing its direction of the pattern
as new stories are published.

While the screens are familiar to me from my previous fieldwork,
everything else is new to me.

They speak a language I do not yet master.

But I can sense the enthusiasm
in their voices
as they talk about
the black boxes
they are building.

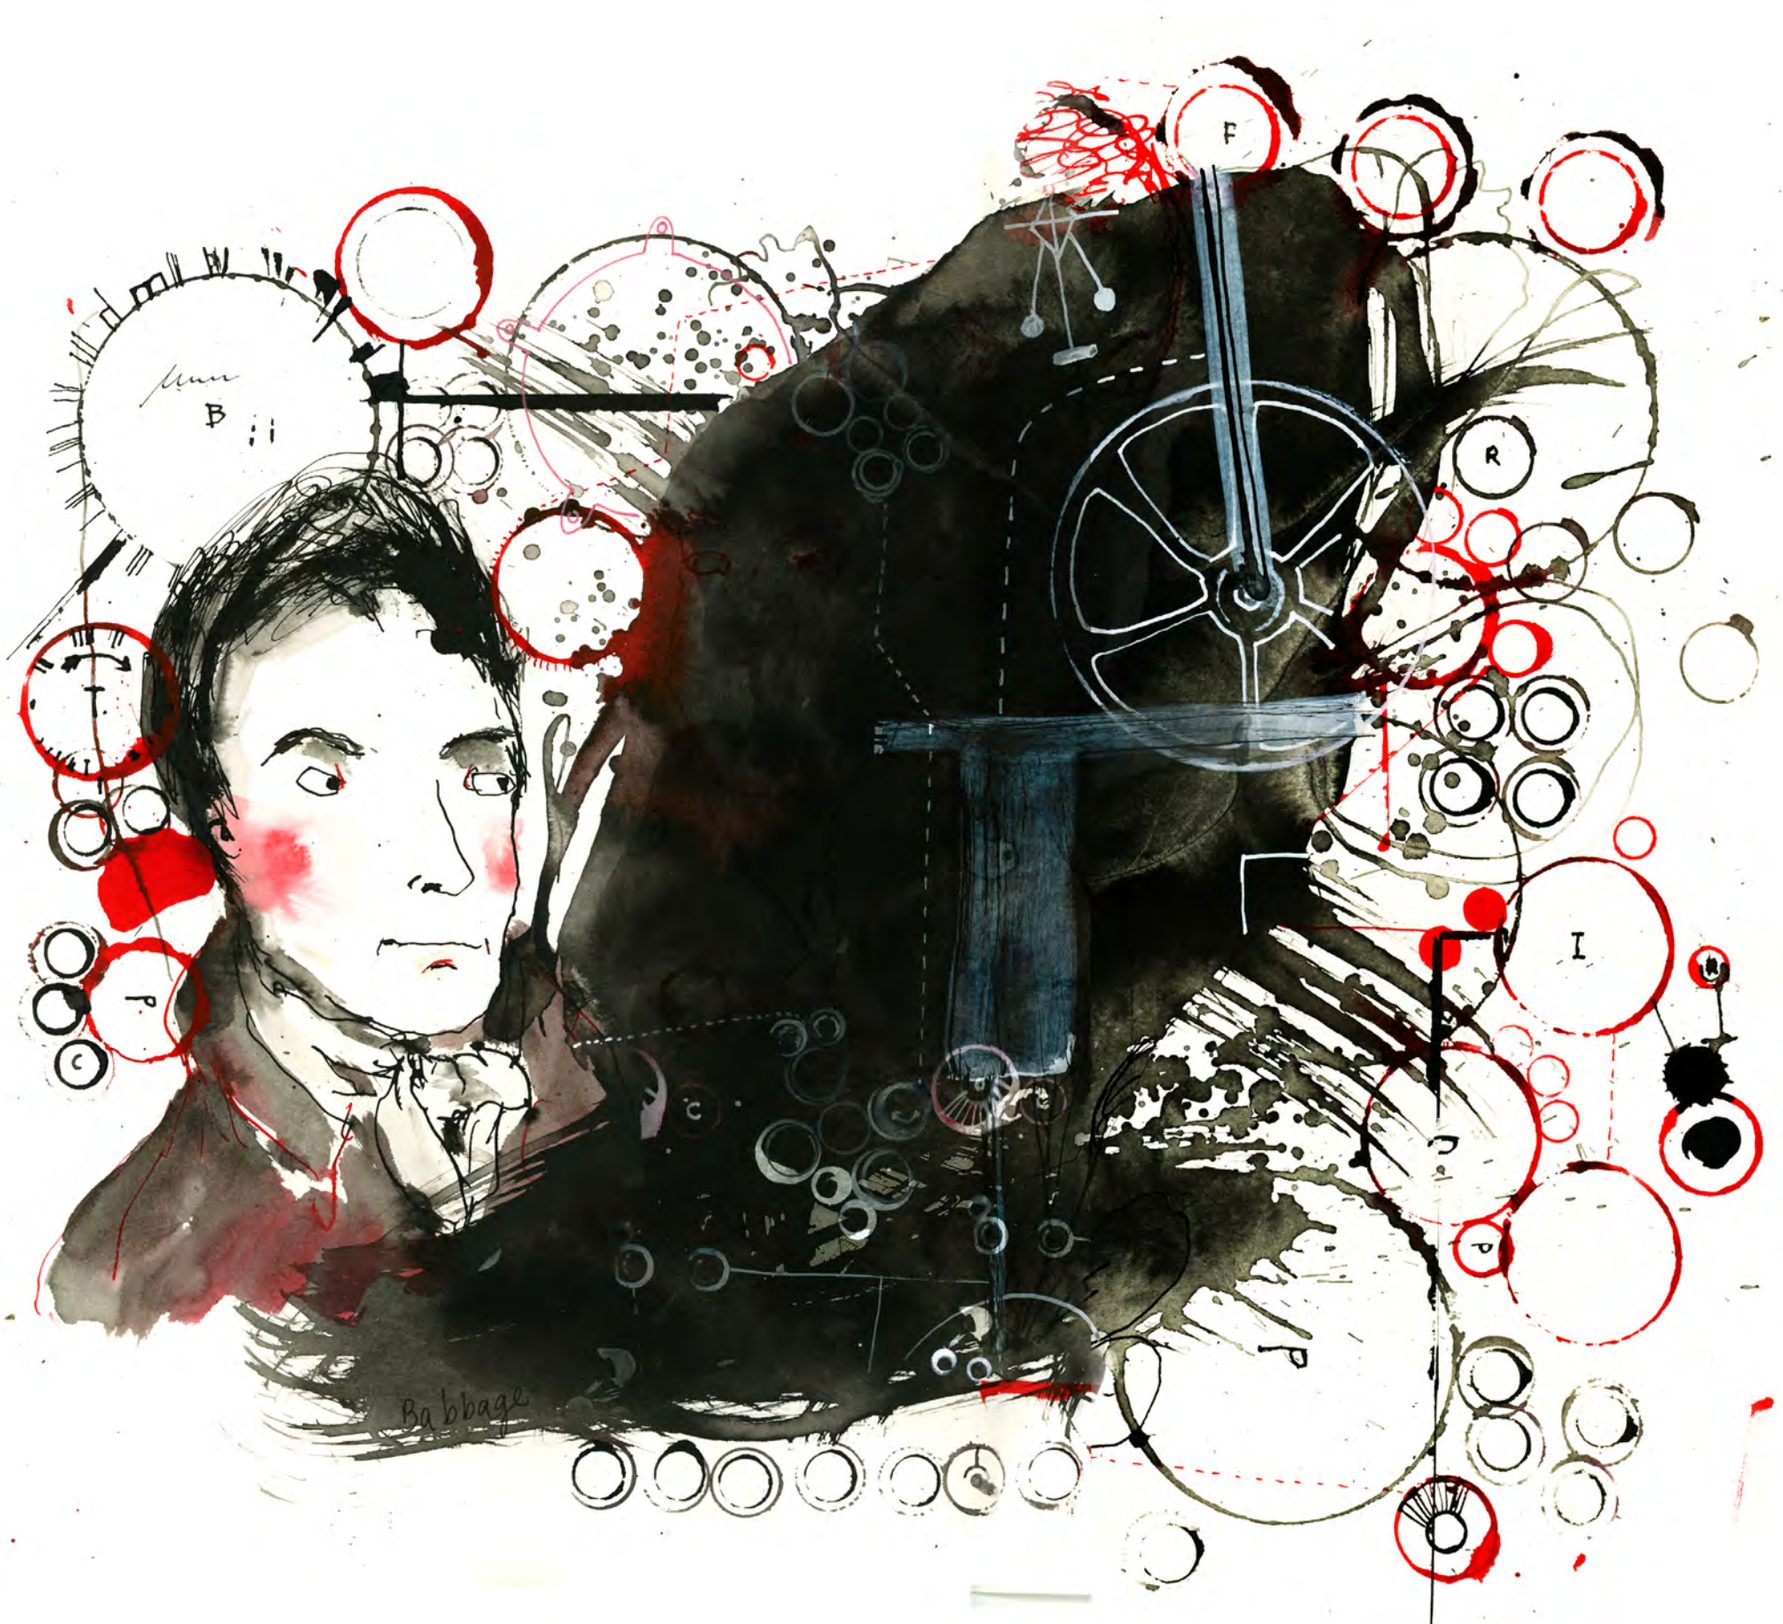

The world around Ada is rapidly changing. Across Britain, the Industrial Revolution roars forward. Steam engines multiply. Factories expand. Automation reshapes the logic of work itself.

While Ada is observing the looms in the mills and factories with her mother, somewhere else in England a man named Charles Babbage is pursuing an obsession.

Babbage imagines a vast calculating machine. He draws numerous versions of it, often with the central processing unit (the “Mill”) planned to be approximately 4.5 meters tall, while the memory unit (the “Store”) planned to be roughly 6 meters long. The engine was intended to be powered by steam and used over 50,000 parts.

The machine would be – like the automated loom – guided by punch cards. A machine capable of storing instructions and carrying them out automatically.

He calls it the Analytical Engine.

Ada does not know it yet, but he will become an important figure in her life. Perhaps a form of replacement for the father she never really had.

Growing from data on you

The type of self-training algorithm I have been observing inside the machine rooms of personalisation algorithms is called a neural network. It's not programmed in a traditional sense – if X, then Y – but rather it's a network of digital neurons in layers. The neurons send information back and forth:

> Each node in each layer connects to each node in the next layer, passing down what it discovered about the data. Data flows through the input and then enters the first layer of nodes, which changes or reacts to the data and passes the new information set to the next hidden layer.[63]

What defines modern large language models such as GPT-4 is the use of the Transformer architecture.[64] These models are trained to predict the next word or token in a sequence, relying on probability distributions derived from enormous corpora of human-produced text. The resulting output is not the product of semantic understanding, but of statistically optimised continuation: what word is most likely to come next, given the previous ones. When we imagine large language models like ChatGPT as having the ability to think, because they answer us when we ask them questions and "act human", we need to keep in mind that they in fact are just really big statistical machines. Machines with the ability to learn and remember from big sets of data.

The models built from layers of different algorithms need to be trained in order to learn and give the appropriate answer when you pose a question to ChatGPT, or what news you should be offered – the genre, composition, or style. This means that a lot of content must be rated or tagged. Put simply, by telling the model what a cat is many times, it will slowly learn what a cat looks like. A large labour force of people is constantly working on rating content for the development of the large language models, in order for them to be better able to recognise patterns in images, sounds, and text.

Using ratings to improve AI models is called reinforcement learning with human feedback, or RLHF, and OpenAI, Google, Anthropic, and other companies all utilise this method. The process includes first the AI processing the data, followed by a human review to refine it. Though it can often feel like you are chatting with a human when using generative AI models such as ChatGPT[65] that does not result from simply ingesting data from the entire Internet. AI programs are designed to efficiently detect patterns, but they lack the contextual understanding that a human brings.[66] This is important when it comes to news and journalistic content, but before we explore this further, we need to investigate how these massive datasets are made.

The dataset as a resource

A program might be tasked with writing several iterations of something, for instance, a recipe. A human will then evaluate the examples, ranking and editing them "for factual accuracy", "how well it fits the prompt", or to "flag toxic outputs".[67] This human feedback informs the statistical model (predictions) of the AI program, making it better at writing more human-like productions in the future. This is a long-established process, with Big Tech companies like Google and Facebook using it for about a decade to improve their products and services.[68]

When news organisations are building and training the model, it's best that the dataset (the input on the content side) they train on is made up of different news genres and topics, and not film categories or advertising categories. The bigger news organisations might be able to afford to create such datasets themselves, with all the sorting, cleaning, and labelling that it takes, but many must buy or borrow datasets that are available open source and adapt them for their own use.

One such dataset is the Microsoft News Dataset (MIND), "a large-scale dataset for news recommendation research". From the Microsoft News website, anonymised behaviour was collected on 1 million users, generating 15 million impression logs (think of it as

pieces of Lego) containing the click events, non-click events, and historical news click behaviours of this user before this impression. The dataset contains around 160,000 English news articles, including titles and abstracts along with the body text, and category and entities.[69]

"Garbage in – garbage out"

As datasets matter for the quality of what is recommended to the user (the output is only as good as the input), quite some time is spent on preparing or even making good datasets. The saying in machine learning goes, "garbage in – garbage out".

News organisations typically cannot use this MIND dataset for training the algorithm. Partly because it's not in the local language, and partly also because the news articles do not stem from the news organisations themselves, which would create a risk of obscuring the recommendations. Imagining that you are sorting out your clothes again, or Lego, and someone has given you some boxes with labels already full of clothes, which you can add yours to. But your T-shirts are a different size, or not as white as the ones in the pre-labelled boxes.

News organisations have specific labels for the content they produce and must sort out. Different beats, but also a tone of writing, specific ways of doing headlines, developed over many years. Like the special recipe for food that you have refined in your kitchen for years, adapted to the ingredients you can find in your local store. A recipe which can hardly be passed on and expected to taste the same, because it's an embodied sense of doing this dish that resides in you.

Similarly, news organisations must also depart from solutions they have previously used, like those from Big Tech platforms, to make themselves independent from these platforms. Or because the product has become too expensive. And when it comes to datasets, they also come with flavour, which means that the news organisation has to either accept that their product will have a slightly different taste or embark on creating their own datasets adapted to the local ingredients. And most of them do this, if they can afford it. The news organisations I observed could, and they did. The work involves the task of annotating a number of articles, which would then be the dataset used for future automated detection of what a given article is about.

And so, they started categorising the world.

Some weeks I place myself in the back of the machine room,

with the annotators.

It's their job to sort the news articles in different boxes,
enough articles so the model will recognise a similar article.

Like the lepidopterist who studies, sorts,
and names the different butterflies and moths
into a hierarchy of species,
the annotators are studying the anatomy of journalism.

I think to myself
that it looks like a tree of words,
where each leaf is a box to put the different news stories in.

They are making a taxonomy of journalism,
later to be used to train the machine.

They are making the punch cards
for the automation of the journalistic loom,
which is the personalisation algorithm.

In early 1829, Ada becomes very ill with measles, which paralyses her legs. For months she is confined to bed.

The girl who once ran through gardens imagining flying machines now watches the world through windows. Yet even in this confinement, Ada's mind remains active.

In the year following, Ada spends her time in bed but keeps herself amused by reading about and practising mathematics. On 27 August 1830, she writes to her cousin (in perfect German):

"I don't think you would recognise me, that's how much I have changed since you saw me; I am getting stronger step by step, and with the help of crutches I can go for walks. I lost my taste completely for riding and flying and such, but I feel well enough to play little piano which gives me great pleasure, and if you would be so kind to send me some light pieces of German music, I would be very happy".[251]

Building the architecture of AI

Annotated datasets are central for the development of machine learning algorithms, which are increasingly drawing on large language models. These models don't learn in a vacuum – they're built on vast seas of text scraped from the Internet, books, and social media, awash with human bias. That bias doesn't evaporate in training – it calcifies. Models can reinforce stereotypes about gender, race, and class because they're learning from histories of prejudice and systems of inequality. The problem isn't limited to text – image recognition tools have failed spectacularly.

In 2015, software engineer Jacky Alciné brought to Google's attention that their image recognition algorithms in Google Photos were classifying his Black friends as "gorillas". A Google representative said they were "appalled and genuinely sorry" for the mistake and promised to fix the problem.[70] But three years later, *Wired* reported that Google had not properly fixed anything. Rather, they blocked gorillas from the service's lexicon completely, choosing to limit the service rather than risk another mis-categorisation. *Wired* reported that their team performed several tests on Google Photos' algorithm, uploading 40,000 pictures of animals, including various primates, to the service. Baboons, gibbons, orangutans, and marmosets were all correctly identified, but gorillas, chimpanzees, and monkeys were "invisible". A subsequent test by *Wired* revealed that Google had further restrictions on its AI recognition of racial categories. Searching for "Black man" or "Black woman", for example, resulted in black-and-white pictures of people sorted by gender but not race.[71]

Scholar Safiya Umoja Noble, in *Algorithms of Oppression*, has shown how search engines mirror and magnify societal bias – placing sexist and racist content at the top of search results.[72] Her work underscores a critical point: The architecture of AI and machine learning isn't neutral, and neither is the labour behind it. The invisible work of annotators – often underpaid and overworked – determines how categories are defined and who gets seen, misseen, or unseen. Without critically examining who labels the data and under what conditions, we risk encoding systemic injustice into the algorithms we increasingly depend on.

Frictions in annotating

In the case of creating datasets for training a model for news personalisation algorithms, it is first necessary to create a taxonomy – that is, a system for organising things into categories based on shared characteristics. Think of it like a family tree for ideas or objects – it groups similar items together and shows how they relate. Taxonomies are all around us: the Linnaean classification in biology (kingdom, phylum, class…), the Dewey Decimal System in libraries, or even how a streaming service sorts shows into genres and subgenres.

In tech and data, taxonomies are crucial. They shape how information is stored, searched, and understood. But here's the twist: Taxonomies aren't just tidy lists – they're acts of power. Whoever defines the categories also decides what fits and what doesn't, what gets foregrounded and what gets erased.

So, a taxonomy isn't just a neutral map – we can think of it as a worldview in disguise. And we can think of annotation or sorting as world-making.[73]

Not surprisingly then, the entanglement involved in the dataset manufacturing used for machine learning and AI is full of frictions. The construction of datasets for AI-driven models in news creates frictions as the annotators compress complex reality to fit into the rigid categories of the taxonomy needed for the algorithmic decision-making and training of machine learning models. Such frictions would also emerge as news organisations needed to build taxonomies for the training of their machine learning models. But what they were trying to sort and categorise was journalism and the journalistic content, which was given different tags, a form of label. On a simple level, such tags could be "sport", "finance", or "entertainment". On a more fine-grained level, this means deciding whether a story about a sports star stopping their career because of stress should be categorised as "health" or "sport" or subcategories.

Sorting the world out

The taxonomy does not start from scratch. The starting point is most often offered by global marketing organisations such as the International Advertising Bureau (IAB). It's developed to match relevant ads to potential customers of the product and used by platforms and media worldwide. Another popular taxonomy stems from the International Press and Telecommunications Council (IPTC) and claims to be fitted to news organisations in particular. Lastly, I could observe that there were categories which did not fit the dataset of news content, or which were simply "new" to the world and thus not present in the existing taxonomy, for example, "Covid". But the work also involves aligning the categories with the editorial profile of the media organisation, and here, other frictions emerged: "The IAB taxonomy tends to focus on cars and washing machines, which is far from the content we publish here", a developer proclaimed one afternoon during my fieldwork. The final taxonomy was a combination of content categories from the IAB, IPTC, and the newspaper's own categories, highlighting the alignment between outside and inside values as a part of the world-making practices involved in making the taxonomy.[74]

Pandora's box of world-making

The challenges of taxonomic world-making are revealed in discussions about how to categorise well-known journalistic thematic areas such as crime and conflict. During a meeting with the project team, they spoke of how they grappled with fundamental questions about semantic relationships: "You've been very consistent with the taxonomy", one participant noted, "but if we open Pandora's box..." This sparked deeper discussion about classification principles. While some participants in the meeting were interested in categorising articles by type and genre rather than just content, others insisted on maintaining focus on their specific goals. A taxonomy of the content of the news stories is basically a categorisation of the stories, similar to when we categorise recipes in a folder, filing some under "baking" with a possible subcategory of "cakes" and another like "bread".

"Crime is under war and conflict", one annotator noted. "When I think of war, I think of Ukraine or something like that". This comment highlights the complex decision-making process of categorising, and it prompted a debate about whether crime should be its own category or remain under broader classifications like "society" or "war and conflict". Furthermore, the team struggled with questions of scale and specificity: Even if crime was in it's own category, should they distinguish between different *types* of crime?[75]

I noticed that the annotators also attempted to align themselves with the editorial values of the specific news organisation by adding categories or by putting themselves in the place of the news organisation's readers as they annotated. In doing so, they contribute to what Louise Amoore has defined as world-making, as the annotation works as a productive force, introducing certain "durable governing rationalities"[76] into the machine learning development. In the process, they navigate ideals of being objective with normative ideals of removing existing bias in the data and on a societal level. The bias that Google and others had been so heavy criticised for, as exposure in instances like the one about misclassifying Black people as mentioned above. Stories like that reveal that even though the process is hidden to most, the annotators matter for which news product the user is offered and what overall window of the world is presented by the news site on any given day.

Transformative tags

During one lunch break, another annotator explained how the process of their work raises numerous ethical decisions along the way. Decisions which reflect whether to portray the world as it is or to also transform the world:

> We have put a lot of work into making it as fair as possible when we tag. We try to see it from the outside. We have one example where there was something with a clothing shop for craftspeople. And here we thought instinctively that we had a category called "men's fashion", because we thought that craftspeople are men.

> In that case, we were a bit biased. But then we also thought that, of course, we should also tag it as women's fashion. […] So, in this way, we actively care about tagging without bias. In that case, we even thought our actions were a bit controversial. […] So, in a way, it was a bit incorrect that we also tagged it as "women". If we are to represent the general thought, that is. If everyone thinks about craftspeople as men, shouldn't that be how we tag them?

This is how women have been excluded throughout history, and the annotators know this. So, they make sure that female craftspeople will also find their way into the newsfeed by giving them a category. But I could also observe that the process of squeezing the complex world into as few categories as possible was frustrating to them. As one of them explained one morning:

> It's like there are, in principle, so many variables and so many ways you can go about football, and then it has to be squeezed down into a tier or some categories. […] It's frustrating when you try to categorise the world, because it's an unsolvable problem, and you can't…"[77]

The ideal of objectivity

What the annotators did reminded me of something I had seen when observing what journalists do when they choose which stories to cover and how to cover them and which sources to contact. At times, they find ways to try and fix the bias they find in society, and other times, they reproduce the bias in their coverage – for example, when sources are most often men rather than women, or when specific social groups are framed in specific ways. And by categorising different news stories as "breaking", "feature", or as "a small side-story", they create and maintain the hierarchies of different news stories on a daily basis.

Journalists make the rules of their trade simple by having these quite specific professional rules of how to categorise different stories, so it's also either/or, although with slightly more doubt and ambiguity in the process. Squeezing the world into single news stories. It's

not surprising that news-making has been described as a form of ritual, almost like a machine with the journalist work being the nuts and bolts, doing their thing, while having an embodied sense of "how big a story is" or "which sources to call first". Although most journalists admit that it is unobtainable, objectivity remains a strong ideal among professional journalists, and as a ritual to achieve such objectivity, they tend to always ask the criticised "other" in each article. To usually have more than one source. To always go for the source with the most authority and knowledge – or the strongest opinion. And to remove themselves from commenting or having an opinion on the theme of the article by having sources who speak in quotes. The journalistic machine becomes speechless or void of an opinion of its own, as speech and responsibility for that which is spoken is referred to the sources. This is the politics of journalism, you could say.

Interestingly, algorithms have been suggested as a means to overcome the bias of the media in its framing of issues and its selection of stories. Diversity in recommender system design has been presented as one possible solution to the filter bubble problem, as the algorithm could simply make sure that the user is recommended a diverse set of stories, for example, from right- or left-wing perspectives or a diverse set of topics. Several issues emerge with these solutions to fixing the tendency for algorithms to drive people into rabbit holes of the same content. One is that it's difficult to categorise articles created by journalists with an ideal of being objective as either left- or right-wing. Another challenge is that if a user is indeed very left-wing leaning, it's easy to imagine how annoyed they would get from being bombarded with right-wing news. While I saw these efforts with fascination and respect and as constructive on the path towards solving the echo-chamber problem, I also had a feeling that moving from one biased system to another biased system was not necessarily going to solve the issues at stake. By focusing on the issue of bias in the algorithms and the machine learning models, we also move the focus from the fact that rather than being problematic, it might be important to be able to see and recognise from which subjective position the source is talking from. And to be able to see who has the responsibility for what is said

or claimed. So, while transparency around the creation of datasets and possible bias are important (although notoriously difficult to actually account for), the realisation that algorithms are not and cannot ever be neutral is even more important. Importantly, the journalistic profession is guided by a set of values and for the most part strive towards ideals of objectivity, values imprinted though years of journalistic training; such values are rather absent in the building of the personalisation algorithms. This does not mean that the algorithms are without values, we just cannot see them.

What became clear from observing the discussions between the annotators is that whether the content is categorised by journalists or annotators, the work entails constructing the world rather than just describing it. The world-making carried out by annotators in the process of creating datasets reveals not only questions of labour, but also that which we might consider machine-driven algorithms being entangled with human and subjective ways of sorting and categorising the world.

Importantly, this also means that values are embedded in the socio-technical system driving how the personalisation algorithm works.

Most of the time the categorisations progress in a seamless manner

month after month in silence.

I am sitting across from the annotators,
who talk quietly between themselves.

The punch cards for the loom slowly find their shape.

The tree they are drawing
is growing with more and more binary branches.

They are solving the puzzle
of how to categorise and represent the complexities
of the real world
into something so simple,
so it can be read, analysed, and repeated
by the machines.

It's like they are squeezing the world
into a machine.

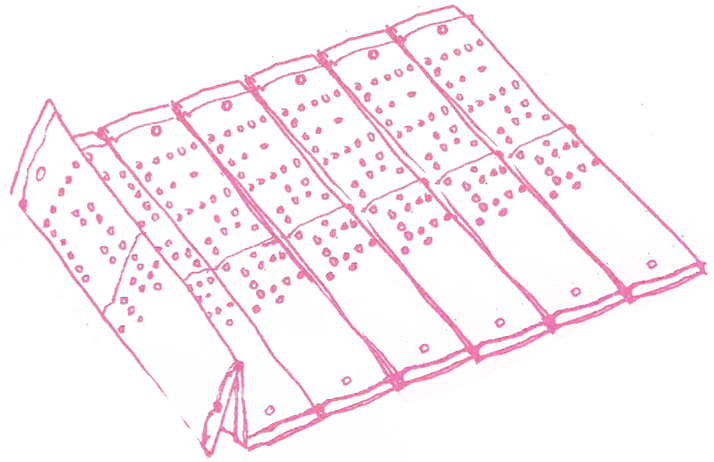

Night settles over Bifrons with a quiet steadiness, as though the house itself has exhaled and finally come to rest. The corridors fall silent, the last footsteps fade, and in Ada's room, the window remains slightly open, allowing the night air to move freely inside.

She lies awake, aware of the breeze as it brushes across her skin. It carries a coolness that feels deliberate, almost searching, as it passes over her face and along her arms. Gradually, the sensation deepens as her body finds rest in a deep sleep. The breeze changes its characteristics from something external to something that seems to gather around her, as if the air itself is taking form.

Her body begins to feel lighter, and without effort, she senses herself lifting, first as a subtle shift and then as a clear movement upward. The bed recedes beneath her. The room loosens its hold.

The window opens wider, and she moves toward it, drawn forward by the same current of air that had first entered so gently. Beneath her, carrying her through the air, she feels the presence of a horse with wings – looking like she imagined it, measured, proportioned, responsive.

She passes through the window and into the night.

Categorising and tinkering the world

When you scroll through Netflix or your news app and see recommendations "just for you", what you're really seeing is the result of a lot of behind-the-scenes sorting and labelling. In machine learning, this process – categorising data – is how systems learn to predict what we might like next. It's like training a digital assistant to make better guesses about your tastes, based on patterns in what you and others have clicked on before.

Even though this labelling work happens quietly in the background, it shapes how these systems function. It's foundational, like the invisible scaffolding behind a building. Yet, as scholars Geoffrey Bowker and Susan Leigh Star pointed out, these sorting systems aren't neutral or purely technical: They often carry hidden political and social weight.[78] For instance, when a music app sorts songs into genres, who decides what counts as "alternative" or "hip-hop"? And what happens when a new sound doesn't quite fit? These choices can reflect cultural biases, even when they seem like simple technical categories.

Bowker and Star call this "infrastructural inversion" – a way of turning the system inside out to reveal the hidden choices and labour behind it.[79] These systems often only become visible when something breaks down or doesn't quite fit – like when a news story refuses to fall neatly into a category like "Politics" or "Tech". In those moments, the people doing the sorting – the annotators – must make judgment calls. They're not just labelling stories – they're shaping how we understand them.

Take a news organisation, for example. It might want a universal system for tagging articles – something tidy and consistent. But the reality on the ground is messy. Local teams might need to tweak the categories to fit their readers' preferences or the national context. As Lampland and Star explained, universal standards never really work everywhere. They always need a bit of tweaking, some trial, error, and a fair bit of what they call "tinkering".[80]

In the end, these hybrid systems emerge – partly structured, partly improvised. Annotators borrow from different systems to get the job

done. It's like organising your kitchen with labels from IKEA, your grandma's spice rack, and a few you made up yourself – because no one system fits the way you actually live.

Hybrid knowledge production

As the datasets are used in the building of the machines, the annotation practices become embedded in larger socio-technical systems. In this way infrastructure is not just technical but relational – it becomes real infrastructure in relation to organised practices. This also means that annotation practices create specific conditions for knowledge production through their relationship with existing journalistic practices and organisational structures. Recent anthropological and STS (Science and Technology Studies) work on algorithmic systems and machine learning[81] has emphasised how machine learning models "don't simply automate existing practices but create new forms of knowledge production that combine human expertise with computational logics".[82] Annotation should thus not be perceived as simply translating between human and machine understanding, but as creating new hybridised[83] forms of knowledge production.

Modelling reality

So why does this matter? And what do I mean by creating a knowledge infrastructure? How different news stories are classified matters for which stories are offered by the algorithm to users. In a similar way to when journalists categorise stories as "breaking" or "just a small side-story". Or in editorial boxes identifying and sorting content into beats such as "Sport" or "Finance". If the large language model must provide a picture of a pretty Scandinavian-looking girl, someone somewhere has tagged and decided what it means to look pretty and Scandinavian. If the personalisation algorithms have to provide more gossip for the celebrity-indulging user, they must know what a gossip news story looks like. This is what I meant when I argued that the models have to *learn* what journalism is. These practices of categorising different stories are something journalists have trained to do for years. Now it's something the models must learn, so they can act like and on behalf of the journalists and editors.

And the annotated dataset is a key ingredient for that learning process.

In my fieldwork, I observed the discussions among the annotators on how they proceed and make decisions on how each news story should be categorised and tagged. They realised that they also have the power to influence how the woven image of "the news tomorrow" will turn out, as a consequence of these numerous small acts of categorising the world. As one of them said when I asked them about how they tried to remain objective:

> It probably depends on whether you are interested in modelling reality or whether you are interested in modelling how one thinks reality should be. Because there is always an exception to everything. So, it is really a question of whether a model should only capture the exceptions or whether it should be able to mirror the norm. Yes, I might say that. It will be a political, or whatever you call it, a very opinionated bias on that model. It will. But that is also sometimes what you have to do to counteract. There are also some things you want to avoid. Things like racism and all that. You are not interested in the models reflecting that.

While there is growing recognition of annotators' subjectivities, current practices often either treat diversity as a problem to be managed rather than a resource to be valued. Recent research in critical dataset studies has examined the fundamental choices in data annotation, including what features of reality matter and how they should be understood.[84] Annotation practices are not neutral, but rather are shaped by what Passi and Jackson called "data vision",[85] or "the interplay between formal abstraction and discretionary choices that shapes how datasets align with algorithms".[86] Thus, the hierarchies and power dynamics of organisations greatly influence the annotation practices of the given company.[87]

A similar data vision is present once the dataset and the taxonomy reach the machine learning models, where the data scientists, for example, must decide on the different weights given to different types of content. One of the organisations I observed has a reputation for publishing a large number of stories on erotic themes, and the data scientists were instructed that this content should

not be recommended by the algorithm to such a large degree. Another news organisation wanted to be sure that local stories were pushed to people even if the user was less interested in these, which would algorithmically push those stories out of the pool of stories to be recommended. An extra algorithmic layer of rules for how the algorithm would perform had to be developed, taking these values into the algorithmic curation. The system is responding to the local values of the news organisation, and those values will shape and form the future personalised algorithmic curation of the news.

A small decision is scaled up.

Scaling

When you think of AI systems – like those used to moderate online content or recommend what you see in your feed – it's easy to imagine them as purely technical tools. But they're deeply shaped by human decisions, cultural values, and political contexts. Sahana Udupa and her colleagues explored this in their work on what they call "ethical scaling".[88] Their main point is that small design choices in AI – like how content is flagged or how categories are created – can have huge ripple effects. Think, for example, of when Google's image recognition software misclassified Black people in deeply offensive ways. That wasn't just a technical glitch: It exposed broader racial biases embedded in the data and systems themselves.

To explain how all this works, Udupa's team built on work by media theorist Tarleton Gillespie, who talked about scaling not just in terms of size, but as a way of stitching together many moving parts – humans, machines, and rules – so they function as one system. Gillespie notes that "a specific kind of articulation": "different components attached, so they are bound together but can operate as one – like two parts of the arm connected by an elbow that can now 'articulate' their motion together in powerful but specific ways".[89] Importantly, the issue of scaling also creates single points of failure: Inherent biases or defects propagate to all adapted models, risking an "algorithmic monoculture" and amplifying harms.[90] The harm is scaled up, so to speak. Udupa's team says we need to scale ethically – meaning we can't just focus on whether an algorithm is technically accurate.[91] We also need to ask: Who is affected? Who gets a say in how it works? What happens as these systems grow and evolve?

They defined ethical scaling as a way of building and governing platforms that's transparent (open about how things work), reflexive (willing to question itself), and inclusive (involving people from different communities around the world). In other words, content moderation shouldn't be something that just happens to users – it should include them. This approach challenges the usual narrative that pits humans against machines – as if it's a tug-of-war. Instead, they argued that AI systems are the result of entangled processes where humans and machines shape each other over time. It's not just about how the model works today; it's about how its decisions also shape the future, growing in ways we can't always predict.

Even though Udupa's work focuses on content moderation, it is similar to what I've observed in the world of news classification too. The systems that sort news into tidy categories often bump into local realities that don't quite fit. These small frictions – those moments when a headline or an article doesn't belong anywhere cleanly – reveal something bigger: the tension between global, one-size-fits-all systems and the messy, specific needs of individual newsrooms with democratic missions. But it also reveals frictions around the scaling of the ways in which articles are categorised into specific boxes and the data on users' clicks and movements.

Thus, even small choices in how we tag or sort the news shape what people see and understand about the world. And these are not the only decisions being scaled up.

As the datasets enter the model for real-world training and testing, other frictions and obstacles emerge.

Like traces of data,
or cookies from our movements
around the Internet,
or breadcrumbs in the forest,
I found Ada Lovelace's notes
on the Internet one day.

Part 3

“Massaging” you with thinking machines – How algorithms curate flows of information

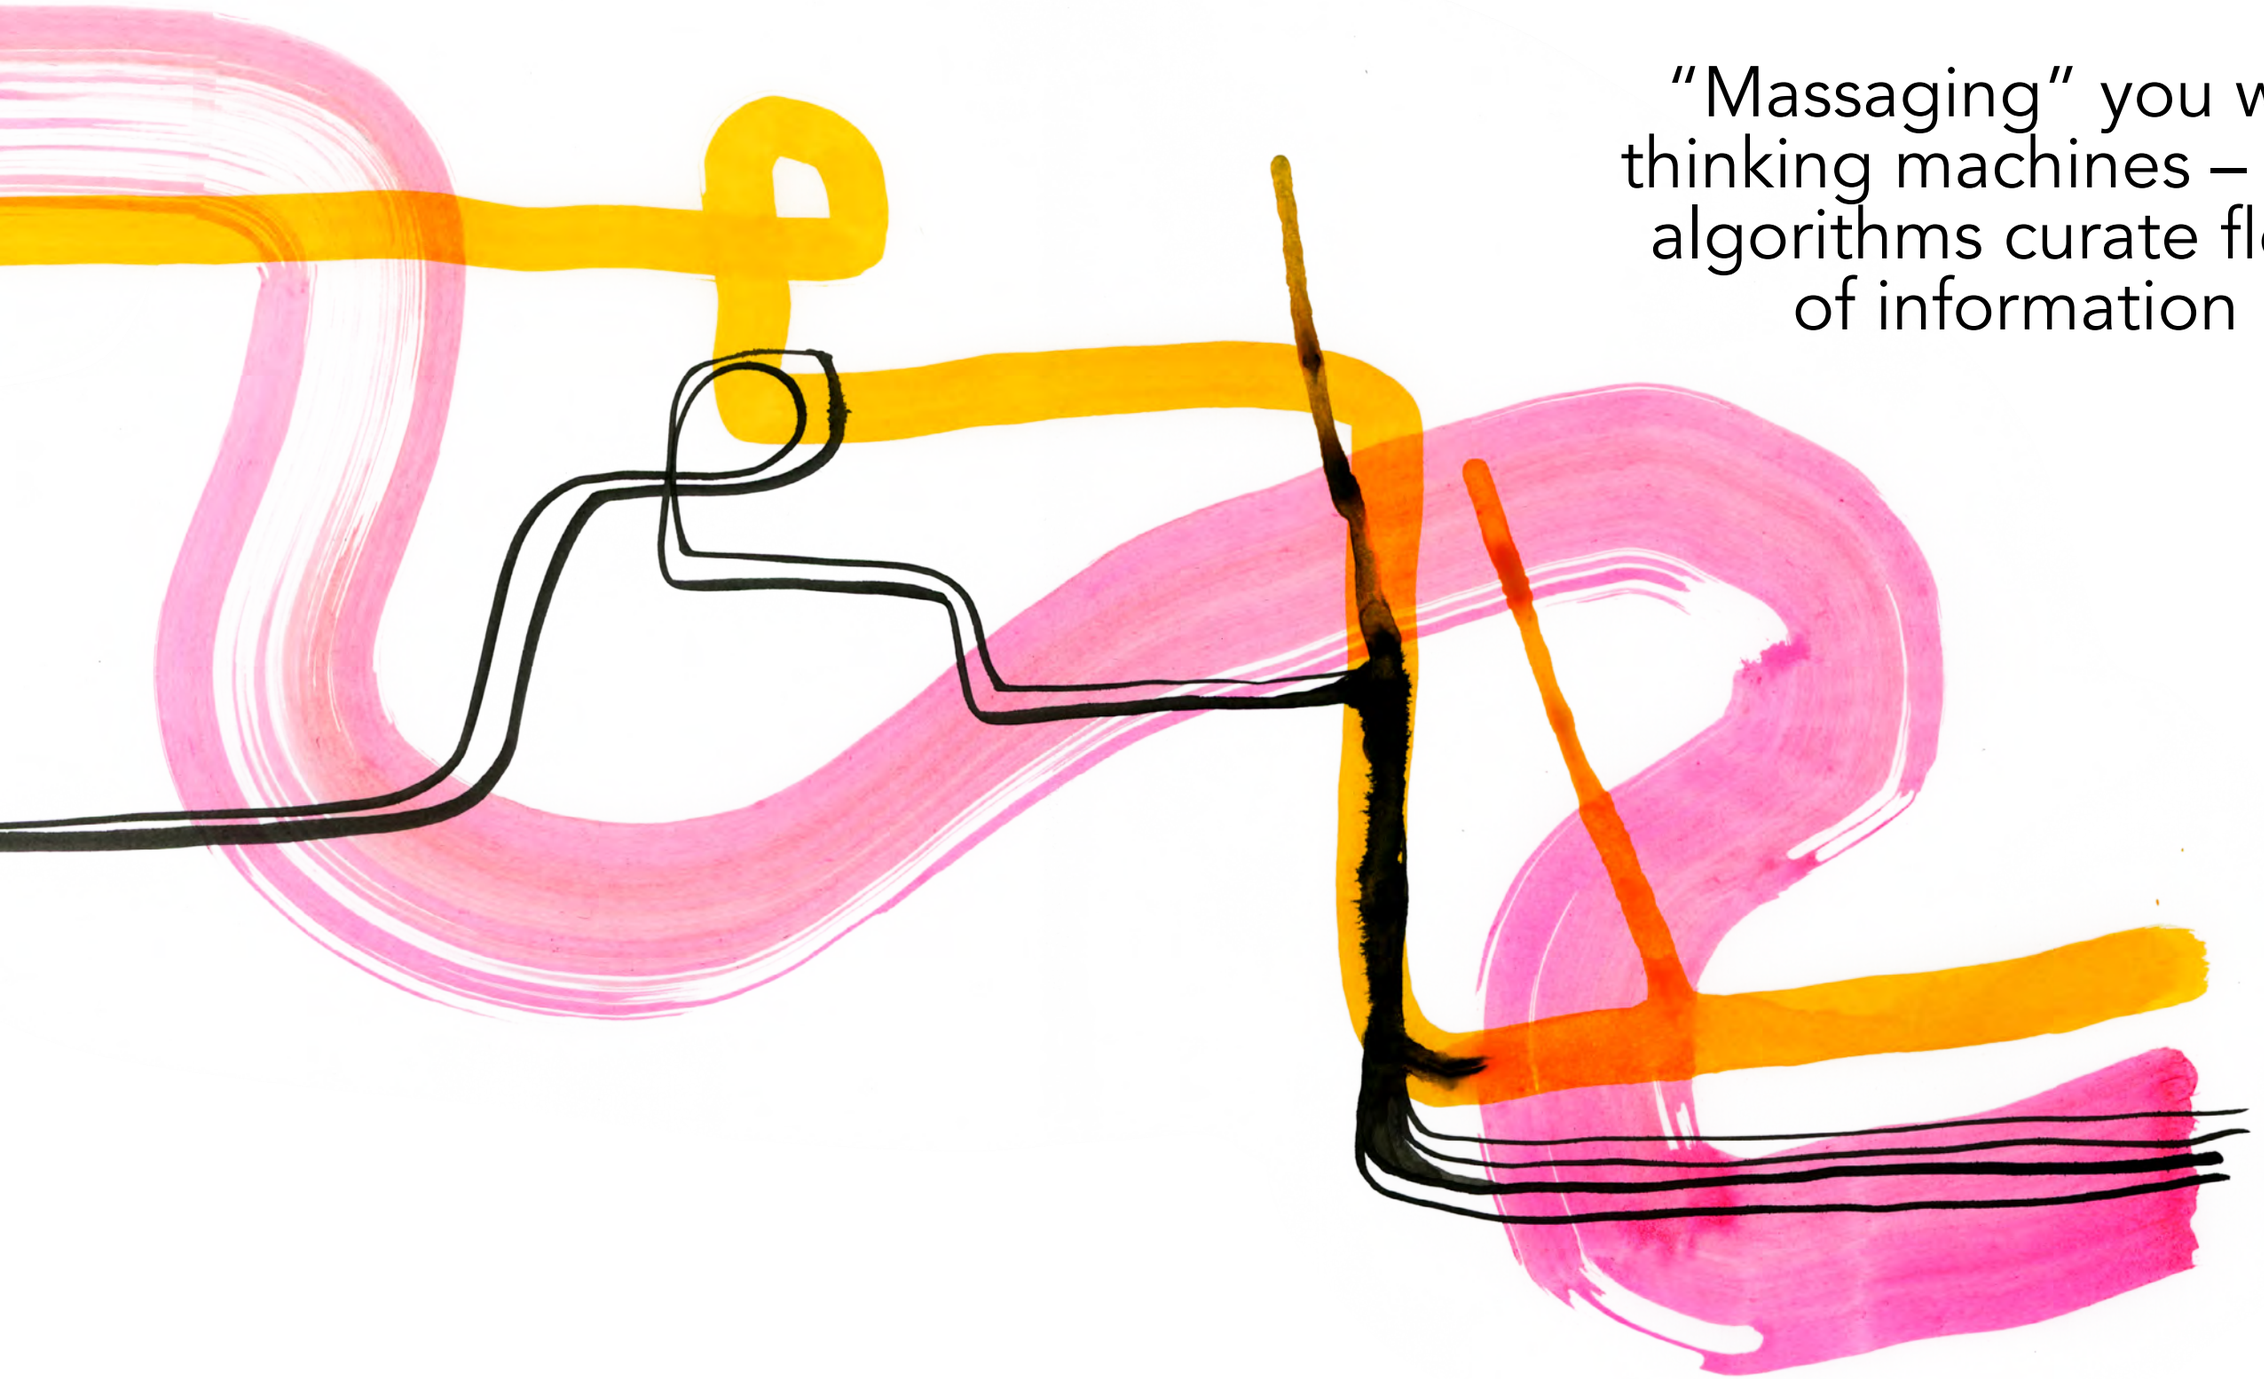

I am back in the box.

The punch cards are installed in the machine.

They train the model, and they need it to be accurate.

The accuracy score goes either up or down in the graphs
they produce from the number crunching.

The graphs notify them how well the model performs
in giving the user more relevant news.

The predictions are made from data on the content (input)
and data on the audiences (also input).

This is put together in a complex formular
(the machine learning algorithm),
which continuously adapts
to the changing behaviours of the news users.

Real-time.

Adapting like a river that adapts
to the melting water from the mountain
or more like a purpose-built canal
leading the water through a city.

Real-time.

What is relevance, I wonder?

And how do you calculate it and measure it on a graph?

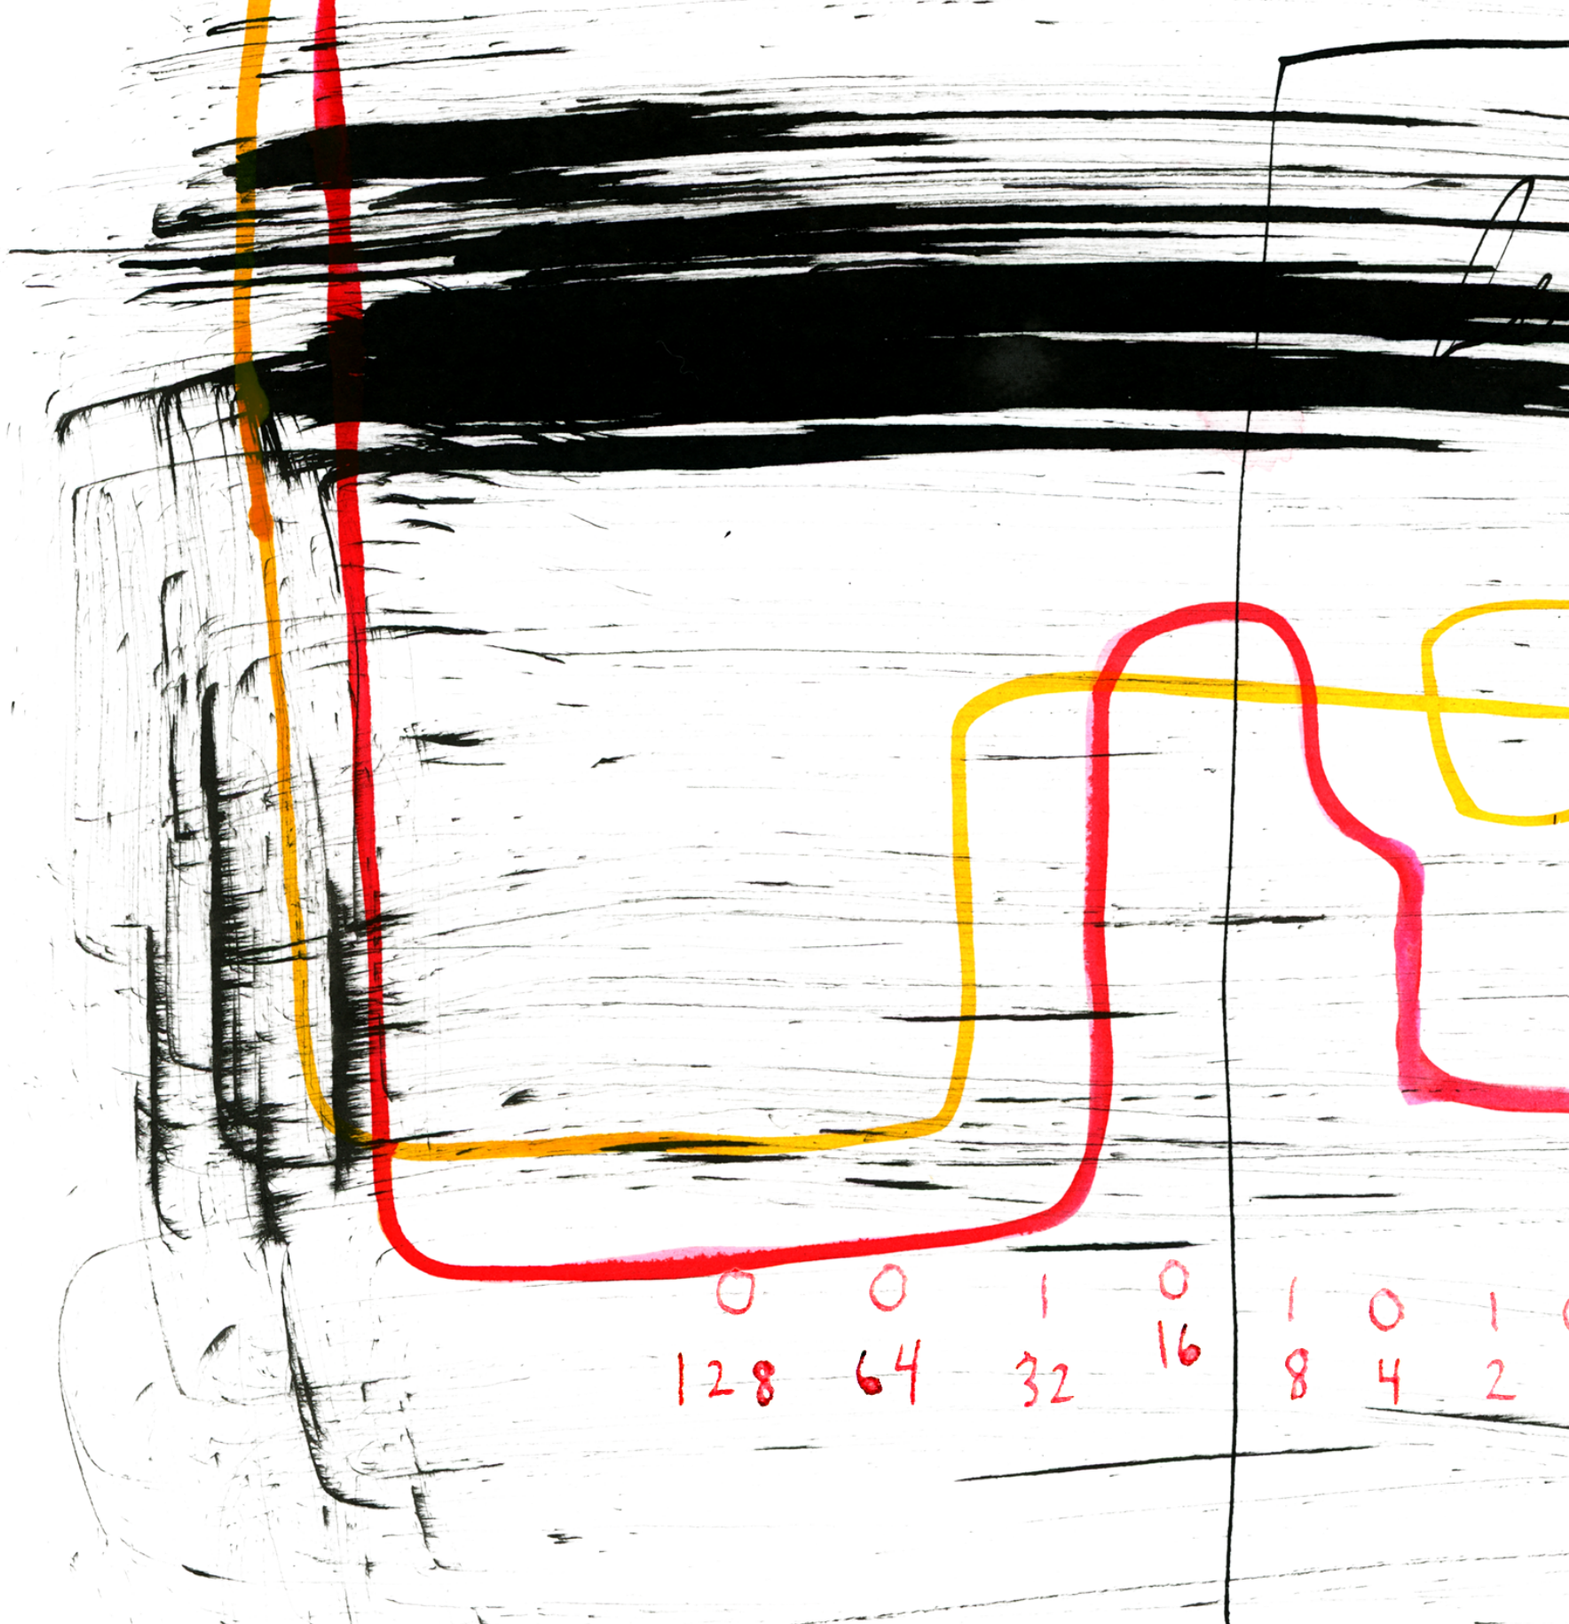

On her mechanical flying horse, Ada moves forward in the night in a smooth, continuous glide, carried by the wind with a steadiness that allows her to explore the movement itself. Bifrons diminishes behind her, its structure dissolving into the larger composition of land and shadow. The gardens that once defined her world now appear as part of a broader geometry, ordered and contained.

She adjusts her direction slightly, and the motion responds immediately. A shift of balance alters her course; a subtle inclination changes her height. The experience feels both new and deeply familiar, as though her earlier calculations and many drawings had already prepared her for this.

The air begins to change as she moves further. A faint heaviness enters it, accompanied by a scent that suggests heat and industry rather than open land. Ahead, the darkness becomes uneven, interrupted by scattered points of light that flicker and spread.

As she descends slightly, the sounds reach her – first indistinct, then unmistakable. Metal strikes against metal. Wood splinters under repeated force. Voices rise, they seem heated, not as if in conversation but in coordinated action.

Below Ada, groups of men move through the night with purpose. Their actions are directed and precise as they approach the machines housed within workshops and mills. She sees them carrying tools designed for destruction, and they apply them methodically. Fires burning across the landscape.

She observes the uprise spreading through the Midlands, particularly around Nottingham, where mechanised lace frames had altered the conditions of work with startling speed. She can see inside the homes of skilled workers who had once depended on their craft and how they worry about their future livelihood as their skills are displaced by machines that allow production to be carried out by untrained labour at scale.[92]

The maintenance of you with machines

For the news media to play a role in the information ecology, it also needs to deliver its messages to citizens. The development of technologies – from the printing press to today's Internet interfaces and personalisation algorithms – are essential for this delivering of messages in the form of news. Thus, emerging technologies have always shaped how those messages could be delivered, in turn allowing for different publics to be cultivated by the news media. Or *massaged*, we might say.

"The medium is the message", as the media theorist Marshall McLuhan famously coined and titled the first chapter of his 1964 book.[93] Three years later, his second book ended up with the title *The Medium is the Massage*,[94] after a typesetting error (or an error of the machine in the programming) that he insisted remain. McLuhan has made further puns on the word message: "mass age" and "mess age".[95]

Concerning this title, McLuhan himself wrote:

> The title "The Medium Is the Massage" is a teaser – a way of getting attention. There's a wonderful sign hanging in a Toronto junkyard which reads, "Help Beautify Junkyards. Throw Something Lovely Away Today." This is a very effective way of getting people to notice a lot of things. And so, the title is intended to draw attention to the fact that a medium is not something neutral – it does something to people. It takes hold of them. It rubs them off, it massages them and bumps them around, chiropractically, as it were, and the general roughing up that any new society gets from a medium, especially a new medium, is what is intended in that title.[96]

What McLuhan referred to is that we tend to focus on the obvious, which is the content, to provide us valuable information, but in the process, "we largely miss the structural changes in our affairs that are introduced subtly, or over long periods of time".[97] When we become aware of these subtle changes in society's values and norms

is when we realise that the medium – the machine – indeed has social implications. The algorithm gives us a massage, so to speak. But it's a long-term massage, one we might not even notice. Only if we pay attention to the machines that deliver the content to our screens and radios and how they carve out different spaces for us to take part and deliberate in different publics.

The messages floating around also massage us as citizens into certain formations, which are necessary for deliberation, collective action, and political participation. For publics to take shape and form and for citizens to see themselves as part of larger collectives beyond themselves. In his seminal work, *The Public and Its Problems*, the American philosopher John Dewey argued that the public emerges as individuals and groups recognise shared consequences from collective actions, creating a need for coordinated responses through communication and democratic deliberation and debate.[98] For Dewey, the public was not static; rather, it was continually reshaped by technological, economic, and social transformations. These technological innovations – such as the book and now the Internet – significantly expanded the reach and speed of information dissemination, enabling publics to form across broader geographical and social boundaries. He recognised the potential for these technologies to facilitate informed discourse and participatory democracy but also cautioned that they could contribute to misinformation or alienation if not coupled with critical education and robust civic engagement. Thus, at different times in history, the communicative machines have carved out different spaces for such civic engagement. While Dewey was positive, scholars today are mostly worried about today's technological advancements, such as large language models and the power to control them, as they also control the carving out of the democratic spaces we inhabit as citizens.

In February 2019, OpenAI "catapulted itself into public consciousness" when it revealed their 1.5-billion–parameter language model, which according to them, could generate fake news so convincingly that they decided not to release it out of concern for potential misuse. While some in the community thought this a "wise precaution",

others considered it a stunt.[99] In April 2026, Anthropic claims that their new model – Claude Mythos Preview – is "too powerful to be released to the public", and will only be available to a group of more than 40 technology companies, including Apple, Amazon, and Microsoft, "which will use the model to find and patch security vulnerabilities in critical software programs".[100]

These technologies matter for the spaces in which we connect, relate, and commit to each other and democracy as citizens. Both on a global scale and inside news organisations.

The massage as maintenance

I observed the data scientists over several months as they tried to improve on the personalisation algorithms they were building with machine learning. They were doing complex calculations of how the users' clicks on different articles, measured on a large scale, would match what the user found relevant. "What does accuracy mean for them?" I scribbled down in my notes as I was watching their black screens with white code language evolve line after line.

I figured out that for them, it means that the recommendations are measured according to the match between the predictions of what the user might click and the calculated relevance of those items clicked for the individual user. I noticed how "relevance" was a word that came up often, both as an argument for the personalisation algorithm as such, but also when they tried to assess how well the algorithm performed. I also noted down that it did not always seem like the different people involved meant the same things when they referred to "relevance". It seemed like a broad enough notion to fit everything into it, while also difficult to contradict. Who does not want more relevant content in their feed or on their favourite news site? Yet, I was unsure of how it was impacting the journalistic values and judgments of the journalists, whom I hardly ever saw. Sitting on the floor below, they did not take part in these assessments and predictions of relevance. Although, they were doing their own assessments and judgments, when choosing which stories to cover and which ones to let go. On a daily basis.

The technologies of distribution of those messages in the form of news transform and have structural implications for the values in society. As they always have. In particular, technologies transform the evolving relationship between the news media as an institution and the media users. It is this relationship that news media feel they are losing, if users turn to chatbots for the recent updates on whatever interest they might have. The relationship with the users of news media has always been a key pillar of journalistic professional cultures, albeit in different ways in different periods in time. And challenged by outside actors differently in different historic periods. Thus, chatbots and AI summaries are just the most recent development causing the relationship to be challenged once again, which means that news media must re-invent their relationship with the user repeatedly. Media scholars have coined this as "audience maintenance work", and this work is increasingly enabled by the massive surveillance and datafication of you as a user.[101] Maybe you have not thought of yourself as someone who needs maintenance (apart from the constant bad conscience of not doing enough exercise); yet, this "audience maintenance work" increasingly involves you as a user of media. Both when it comes to how we use chatbots in our everyday lives and how we use the news – two initially distinct practices which seem to increasingly blur and converge with the development of AI.

In the book *The Audience in the News*, Dwight DeWerth-Pallmeyer unpacked the uneasy, often invisible labour that goes into audience maintenance work – the relentless, behind-the-scenes effort to hold public attention in an era of shrinking loyalty and shifting habits. News organisations, he argued, do not simply report: They engage in a constant recalibration, crafting stories, headlines, and even journalistic priorities based on an imagined, data-fed composite of their viewers and readers.[102] Journalism is shifting from an independent force of editorial judgment to become closer to audiences' expectations and needs, trying to balance between what the metrics are telling them and their own editorial choices of the day. The irony, of course, is that this pursuit of audience loyalty has historically often pushed news closer to entertainment, bending the industry's historic ideals and public mission to the gravitational pull of clicks, shares, and fleeting digital attention spans.

The personalisation algorithms are now a central part of that audience maintenance work, massaging us to spend more time with the news. And as the data scientists were trying to predict and increase relevance, I asked myself if this practice pulled the publishers away from this public mission. Or, could "relevance" as a measuring stick cover both our lazy need for a quick fix of light entertainment and the public mission?

Inside the box, I watch

as the developers train the model on large datasets,
and it slowly becomes better at "recognising" text or images.

"Like a child learning a language", they explain to me.
They are humanising the algorithm, I think to myself.
Yet it is NOT a child, it's a statistical model of a language,
a really big calculator.

One day, one of them says:
"Even I don't even know how it works,
but the editors want me to explain why
it spits out certain recommendations and not others.
But this is a result of a large model,
a kind of room with 300 dimensions,
so even I cannot explain, even if I tried".

In the process of training,
the algorithm is rewriting its own code.
It has become self-learning, but the process
makes the code so long and complex
that it becomes incomprehensible for humans.

As the algorithms become better and better
at "seeing" and "knowing" the user,
the developer becomes less and less
able to see into the code of the algorithms.

It's like the woven picture
is so huge and the developers so close
and unable to get far enough away to see what's on the picture.

It's like they cannot see the forest for the trees.

From above, Ada sees both the precision and the intensity of the movement. The destruction is neither chaotic nor incidental: It is organised, repeated, and shared among those involved. The machines, built to operate without fatigue or variation, stand in contrast to the bodies that now dismantle them with force and urgency.

As the wind takes her further north, Ada observes similar tensions unfolding among the cloth-dressers in Yorkshire, where mechanisation erodes a trade grounded in bodily skill and long apprenticeship. She sees George Mellor, like Robin Hood in the Sherwood forest, and the people he could rally behind him, setting out to destroy what has taken their livelihood.[103]

What becomes visible from her position is not only the destruction of machines, but a broader transformation in how work and authority are organised. The machines do so much more than replace labour: They restructure it, removing elements of judgment and experience while embedding decision-making within systems that cannot respond to those affected by them.

She feels the sadness and desperation of the workers in her body, as the wind carries her onward, away from the fires.

The form of the news and the publics it carves out

In the beginning of 2022, I crossed the Townhall Square of Copenhagen and picked up my access card to another news organisation which had embarked on building personalisation algorithms. This one was a tabloid newspaper: *Ekstra Bladet*. Apart from being one of the most advanced and innovative when it comes to building recommender systems for personalisation, they were interesting because they already had quite a commercial pull reaching back to when they were founded in 1904 as an evening edition to *Politiken*. In 1905, the newspaper was established independently and focuses on investigative journalism, news, sports, and entertainment. Some consider it a sensationalistic newspaper, but in combination with the investigative journalism and political coverage, it was ideal for observing how the implementation of algorithms might change the form and values of the publication, as those two are intrinsically linked. And one specific afternoon, a chat with one of the data scientists stood out. But before we get to that conversation, I want to briefly explain how the form of the news matters for the type of massage the media offers its users. Again, a brief historic perspective will help us.

We can observe how new and emerging technologies transform the relationship between the news media and citizens by examining the development of newspapers, which has historically played a significant role for the spreading of information enabling new forms of democratic debate and participation. The first significant technological breakthrough in news distribution came with the invention of the printing press in the fifteenth century, which revolutionised the dissemination of information. Before this, news was typically spread orally or through handwritten manuscripts, which were labour-intensive and accessible only to the elite. The printing press democratised access to news, enabling mass production of newspapers and pamphlets, which became central to public life. The press played a pivotal role in disseminating Enlightenment ideas, fostering public debate, and shaping modern public opinion. The public spaces expanded.

With the expansion of printing technology and rising literacy rates, books, pamphlets, and newspapers became widely available. Periodicals like *The Spectator* in England and *Gazette de France* in France became key vehicles for sharing essays, scientific discoveries, and critiques of authority. The press provided a platform for philosophers to discuss controversial ideas, such as secular governance and human rights, reaching an audience beyond the elite. Censorship was a significant obstacle, as rulers sought to control the press to suppress dissent. However, the rise of underground publishing networks and intellectual salons allowed Enlightenment ideals to spread despite these restrictions.

Working class appeal

In the Western World, the emergence of newspapers in the nineteenth century mirrored the growth of the nation itself. The Penny Press, which emerged in the 1830s, is a key example of how new printing technologies (such as steam-powered presses) made newspapers affordable to a broader segment of the population. Titles like *The New York Sun* and *The New York Herald* offered cheap, mass-produced news that appealed to the growing urban working class. At this mass-mediated point in time, there was still a fairly long production process between the formulation of messages and the print and distribution of those messages. The media worked with the scheduling of pages, where the space for printing the messages was limited, and the newspaper was shaped accordingly. The user of news on the other end was largely passive, although they most likely did not read articles in the intended order and picked the stories they felt most inclined to read – often the obituaries and the people's pages, as it turned out. But newspapers often had very little knowledge about these practices and assumed that their duty of informing the public was fulfilled the moment the newspaper was delivered to post boxes and corner newspaper stands. The newspapers felt no urge to activate the news users apart from the occasional letter to the editor.

The rise of mass circulation

The way news is presented – the "form" – shapes not only how information is understood but also how audiences perceive themselves as members of a public, Barnhurst and Nerone argued.[104] In their now seminal work *The Form of News*, they noted how newspapers, with their evolving layouts, writing styles, and organisation, have historically cultivated distinct publics by structuring engagement with the world. For instance, early newspapers in the eighteenth and nineteenth centuries prioritised dense columns of text, mirroring their role as tools for a literate elite to deliberate over political and economic issues. This form encouraged a reflective and analytical mode of consumption, fostering publics centred on reasoned debate. By the late nineteenth century, with the rise of mass circulation papers, the form shifted to sensational headlines, visuals, and human-interest stories, engaging a broader, more diverse audience. These changes invited a different kind of public – one more focused on shared cultural narratives and emotional connection. Barnhurst and Nerone also suggested that the form of news is never neutral: It structures the relationship between individuals and society, influencing how they see themselves as citizens or consumers.[105] Each format – whether broadsheet, tabloid, or digital algorithmic feed – creates its own kind of public, framing not only what is newsworthy but how users should interpret their place within a broader social and political context. So, the form of the delivery of messages influences what kind of public conversations we are able to have as societies. They do so because they enable us to connect, relate, and commit to our role as citizens and to each other across the different audience islands we might inhabit. So, what spaces of publicness does the algorithmically curated site carve out for us?

I had had a hectic morning getting into central Copenhagen. The train was packed, and it was unbearably hot as people stood rubbing shoulders. A drop of sweat ran all the way from my forehead and down my face as I moved myself out of the train with the fat human snake, now spreading out into all corners of the city. I hurried across the town square, past the pigeons, with the smell of hot dogs sneaking up my nose. I was eager to join the early morning

meeting, where the head of sales would lay out the latest figures and numbers – in short, how their audience maintenance was going this specific month. But it turned out that the meeting was the least interesting that day. After lunch, the developer sitting at the table next to me told me something that he had been wondering about for a while. *Ekstra Bladet*, where he now worked, is known for a specific very popular feature inspired by the British tabloid *The Sun*, called "the page 9 girl". A good-looking topless girl, posing, taking up a large chunk of the printed tabloid format. The feature is so popular that the large screens I can see from the desk where I am sitting are showing both recent sales and clicks with the page-9 numbers and without. And it is this controversial feature that puzzled the developer that afternoon:

> We have been told that our recommender models must not recommend "sex and relationship" content, even though there could well be a good reason for it. I mean, people would be crazy for it. So, it is – I don't know if I would call it – an ethical consideration or what I would call it. This is some kind of consideration for sure. I actually think it's quite interesting in a way. Because it's quite inconsistent. I kind of think that if we publish content like that, page-9 girls, sex and relationship and stuff like that, then I assume that it's something we would editorially stand for. And if that's the case, then I think it's a bit curious that we don't allow our models to recommend it. If you say A, you must also say B. And so, I don't think they have done that in this case. Although I do understand that if we had a model that constantly 100 per cent recommended Page-9 girls, then we might well have to rethink it.

As he was carrying out the audience maintenance work, he was wondering if the path they were carving out for the users was the right path. Why not just give people what they want, he asked himself. If they crave erotic content, let us give them erotic content. Lead them down the path of eternal free ice cream. He had encountered a journalistic value: a public mission indicating that recommending news is not just about more juicy content for the individual, it also about taming news users. Taming them to go the right way, which

is not always the way the users themselves want to go. This kind of judgment – how much of a specific type of content and in which concrete situation – used to be the job of the front-page editor. Now, the job is shared with the algorithmic system and the people making decisions about how to build the systems. The agency is transferred to the algorithms.

Automated curated news sites structure information in a specific way, carving out how we should see ourselves as citizens, how we should navigate the vast sea of information available to us, and how we should connect to others and what we should commit to. The personalised news site tells us that news should be relevant to us personally. And the idea behind it all is that personalisation algorithms can help link the user to the relevant content in ways that no single journalist or editor is able to do.

All this reminds us that the decision of what the algorithmic path should look like is never neutral. It's full of politics and choices, like the one to exclude the half-naked page-9 girls from the personalisation algorithm.

Inside the box,
the developers crunched over
their black screens with code text.

Having sorted and cleaned the data,
they are getting ready to test the models.

They need to carefully observe what the machine does,
what kind of an image it is able to weave.

The results of the tests are presented at the weekly meeting,
the editor is puzzled as he tries to make sense
of the spreadsheet on the large screen in front of him.

He appears almost angry as he rhetorically asks the room:
"Why the f is that travel guide on Australia recommended –*
who of our readers have an interest in something so far away?"

The data scientist shakes his head.
A look of half worry and half annoyance.

"Well, someone clearly did have an interest in the story,
and for that reason it is automatically also recommended
to other people similar to this user, who then also read it",
he explains.

The editor discovers that the loom he is building is weaving a
different image than the one he had imagined.[106]

The scene shifts gradually, the open landscape giving way to enclosed space. Stone walls rise around Ada, and the air becomes still, contained within the architecture of governance. She finds herself inside the House of Lords.

The chamber is filled with controlled movement and measured speech, a contrast to the urgency she has just witnessed. At the centre stands a man she recognises immediately, though she has known him more through stories than through presence. Her father, George Gordon Byron.

He addresses the assembly during the debates surrounding the Frame Breaking Bill of 1812, legislation that proposed capital punishment for those destroying machines. Between 1811 and 1812, the raids had led to the destruction of hundreds of frames in coordinated nighttime raids.[107] His argument does not deny the violence that has taken place but seeks to reframe its meaning.

He speaks of workers whose livelihoods have disappeared, of wages that no longer sustain, of homes affected by the rapid transformation of industry. He insists that those being judged are not simply offenders, but individuals responding to conditions that have left them with few alternatives. “These men were willing to dig”, he argues, “but the spade was in other hands; they were willing to work, but the loom was shut against them”.[108]

From her position, Ada observes both the argument and its limits. The chamber receives his words, but the broader alignment of political and industrial power has already begun to take shape. The state responds to the unrest by deploying thousands of soldiers to protect machinery and suppress further action.[109]

The legislation passes. Enforcement follows as the state deploys approximately 14,000 soldiers to suppress unrest across the affected regions. In January 1813, fourteen men are executed at York, among them George Mellor, then twenty-three years old.[110] Others are transported to Australia.

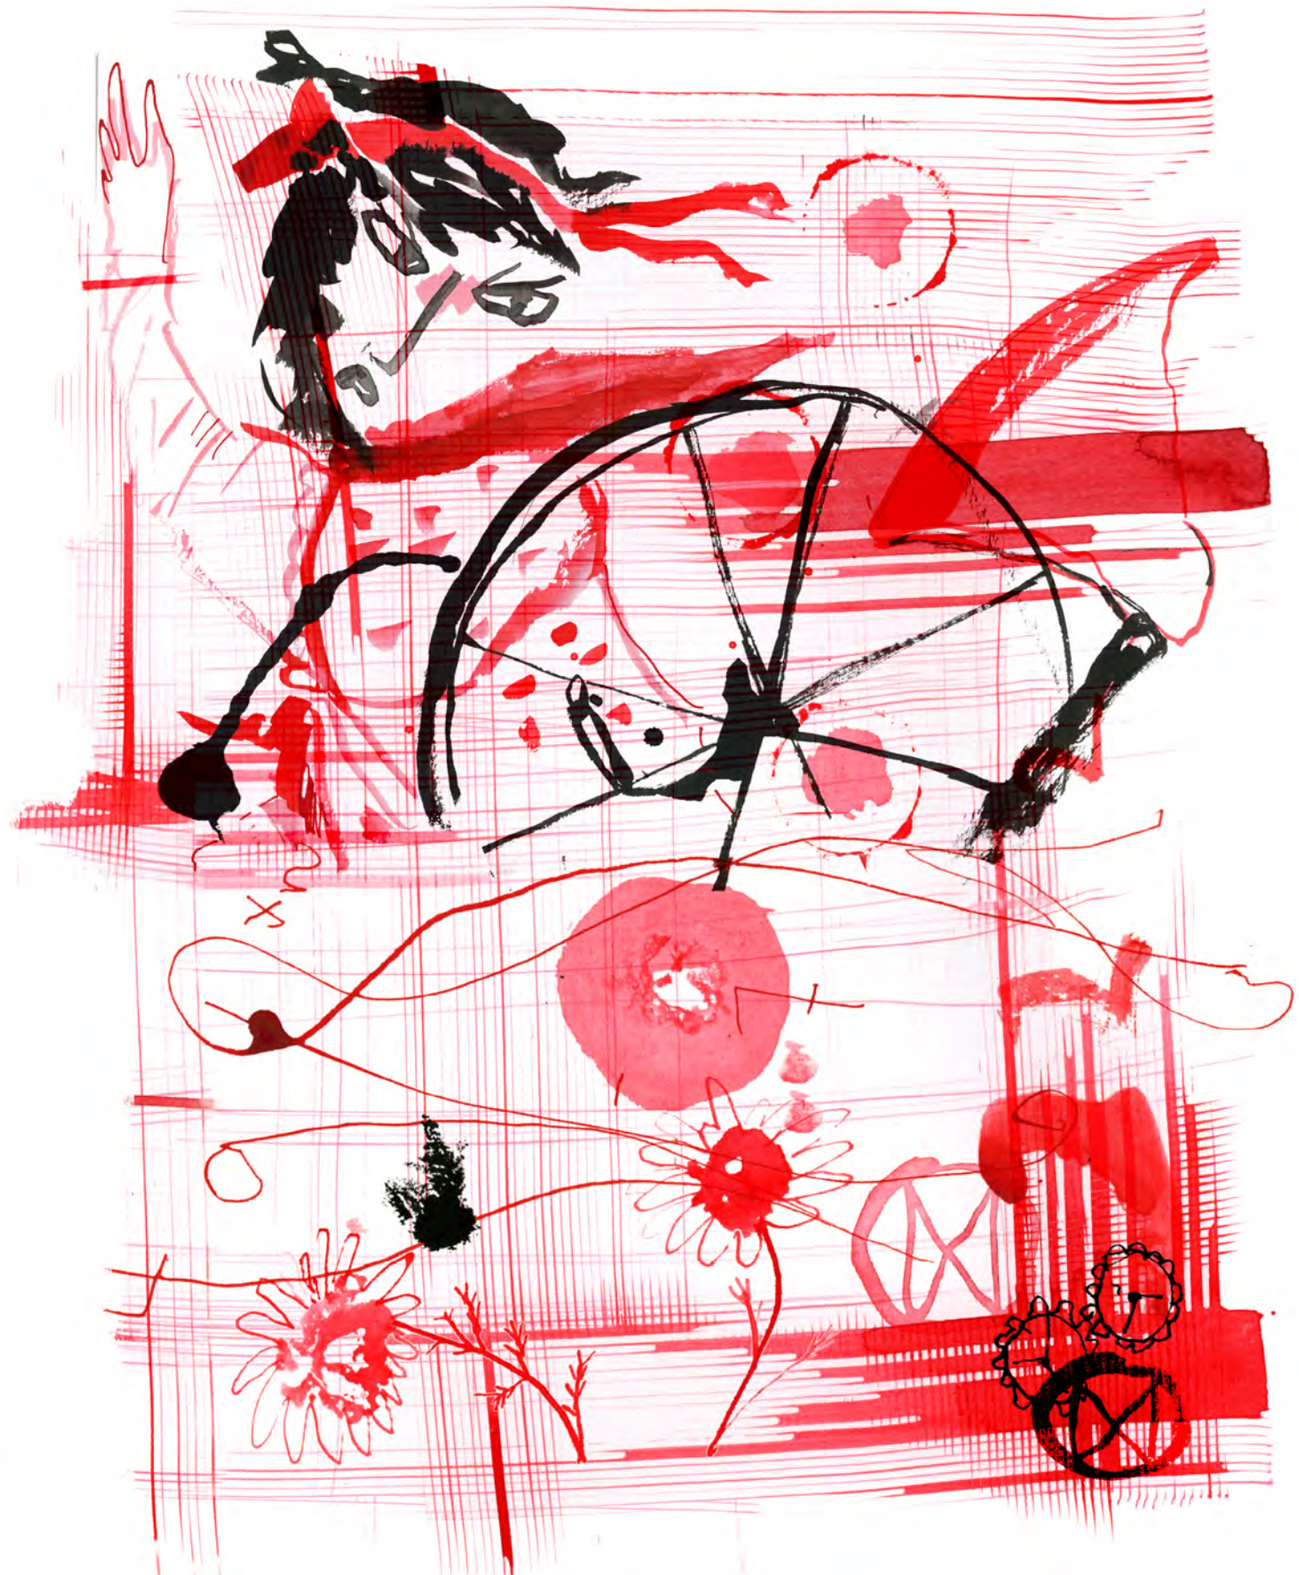

Increased commercialisation of the news media

How machines massage us transforms as new technologies emerge. The advent of radio in the early twentieth century marked yet another fundamental shift in the distribution of news – transforming the relationship between media and users once again. Unlike print media, which required literacy and physical access to newspapers, radio allowed for the instantaneous transmission of news across vast distances, bringing real-time reporting into people's homes. This vast connection with people in their private homes also meant that the medium became more personal in its expression and form.

The radio as a new form of distributing information was particularly significant in countries like the US and the UK, where it quickly became a vital source of information. In the US, radio broadcasting became a dominant form of news consumption by the 1930s. Stations like NBC and CBS revolutionised how news was reported. Radio journalists such as Edward R. Murrow brought the world into listeners' homes with live broadcasts from the frontlines of World War II, and programmes like *The CBS Evening News* set the standard for broadcast journalism. Radio allowed for more immediacy than the printed press and gave a more personal and emotional tone to news delivery, as listeners could hear the voices of reporters giving firsthand accounts of events. In Europe, radio also played a crucial role in the dissemination of news during times of conflict. During World War II, the BBC became a trusted source of information, providing not only British citizens but also those in occupied territories with vital news. The BBC's emphasis on impartiality and accuracy became a model for public service broadcasting in Europe. In countries like Germany and France, the use of radio for propaganda purposes during the war demonstrated the power that broadcast media has to shape public opinion and rally support for national causes. But the vast spread of the new media formats also intensified the battle for the attention of the users.

The battle for user attention

There was significant initial reluctance against radio broadcasting from the newspapers, which was of course due to the fear that a consequence would be fewer morning newspapers sold. The radio pressed the newspapers and brought changes with them in content and structure and as a competitor for people's time and attention, but interestingly, it did not stop the will to read. The newspapers could therefore still show an increasing circulation, and the real danger only appeared after 1945, when the third mass communication medium – television – triumphantly broke through. Radio had a unique influence on the news media landscape by adding a sense of urgency and intimacy to news delivery. It also shifted the focus from written to auditory experiences, with news becoming a more passive consumption activity. It laid the foundation for the rapid development of television news in the postwar years.

Television, which emerged as a dominant medium in the 1950s and 1960s, further transformed the news media landscape by combining the immediacy of radio with the visual appeal of moving images. The rise of television news created a new paradigm in journalism – one that emphasised not just what happened but *how* it was portrayed. This shift was particularly visible during key historical events such as the Kennedy assassination in 1963 and the Vietnam War, both of which were covered extensively on television.

Figures like Walter Cronkite became household names. Cronkite's role as the anchor of *CBS Evening News* during the 1960s and 1970s exemplified television's ability to create a personal connection with viewers. In Europe, the impact of television news was similarly transformative. The BBC in particular became a model for public service broadcasting, emphasising accuracy, neutrality, and the education of the public. In contrast, commercial television networks like ITV in the UK began to introduce more sensational formats, mirroring trends seen in the US with the rise of 24-hour news channels. Television news, however, was also more prone to commercial pressures. The need for high ratings led to the rise of sensationalism and the focus on "soft news", such as celebrity

coverage or human-interest stories, often at the expense of "hard" political reporting. The new machines meant that the news media made and maintained or changed their views of the media users, and soon they would have a new tool to help them in that battle for attention – a tool to keep a constant eye on those users.

The rise of measurement technologies

This increasing commercialisation of the media worried the German philosopher Jürgen Habermas. He argued that the advent of radio and television had fundamentally altered the dynamic between the news media and citizens, challenging his envisioned ideal of the public sphere. The spread of television and radio introduced new forms of communication that prioritised one-way dissemination of information over reciprocal dialogue, he argued. Unlike print media, which required active engagement from readers, broadcast media delivered content passively to audiences. This shift, Habermas contended, weakened the participatory and deliberative nature of the public sphere. Public discourse was increasingly shaped by mass consumption of content produced by corporate and political entities, who often sought to manipulate public opinion rather than foster critical debate.[111]

The new technologies for distributing news also came with improved "machines of measuring and targeting" the news users, which meant that publics could be reached in even more fine-grained ways – in turn transforming the ways the news media would imagine publics and their cultivation of those publics.

The new formats of more segmented distribution were sparked by the introduction of TV-meters, which revolutionised how television audiences were measured, providing what was seen as more precise and objective data on viewing habits. Editors and journalists no longer had to rely on their gut feeling of what audiences might like; they could obtain direct evidence of their viewing patterns.[112] This was enabled by the emergence of a new league of suppliers of measurement technologies. Nielsen, a leader in audience measurement systems, played a pivotal role in implementing this technology. The process involved installing small electronic devices, called "people meters", in a representative sample of households. These meters tracked when a television was on, what channel was being watched, and, in some cases, who in the household was watching. Participants often logged their viewing through personal codes or buttons on the device. Nielsen's system aggregated this data to create detailed reports on ratings and viewership trends, offering broadcasters, advertisers, and networks critical insights into audience preferences. Introduced in the late twentieth century, TV-meters replaced less reliable diary-based methods, becoming a cornerstone of modern media analytics by standardising and quantifying the value of television content.

But what did this direct knowledge of what had been watched, heard, or read in the homes of citizens do to the ways in which media had sustained not just markets, but also democratic citizenship and public debate? The media scholar Ien Ang delivered a detailed argumentative journey regarding the intrinsic limitations of the kind of information that ratings deliver, considering them to be manifestations of "market feedback technology",[113] the primary role of which was to provide details about the size and composition of audiences. Within the media industry, statistics thus achieve "the ultimate prestige – or perhaps simply an undisputed faith" of executives – who, by objectifying and controlling knowledge, could convert it into an economic commodity.[114]

The new possibilities of tracking and measuring users would have profound impact on the image the news media held of their users and their own role in serving those users. And the arrival of the Internet would only reinforce and amplify that image.

The other day
I interviewed a project manager

from a large Swedish media organisation,
responsible for the grand scheme personalisation project
they are in the process of implementing.

She explained that the goal is to personalise the whole site
with a mix of recommender systems from top to bottom.

Every single box.

No journalist will decide today's top story.

Not two users will have the same-looking site.

I ask her how the journalists responded to being stripped of
the power to decide what's important?

It's of course mixed, she says.
Most are positive,
and more positive after the workshops we have held.

They still have to keep an eye on the algorithm,
and we have trained them in doing that.

We found out that some of them have worked out how to
get more exposure so more users would discover their story,
by taking the story down and putting it up on the site again
shortly after.

They are gaming the algorithm, I reply.

Yes. Because the algorithms are designed to boost stories that
are new as we are a news organisation, but now we have had
to make rules to stop the journalists fiddling with the time-
stamp, she says.

The journalists are stripped of agency
and now are like the Luddites,
destroying the machines that are automating their work,
I think to myself.

Ada survives the illness and regains use of her legs, but the illness has a profound effect on her, and she loses her dreamlike insouciance. But this does not stop her from calculations and thinking about the wonders of machines.

It is as if she knows in her heart that she might not live long enough to see how the logics of automated machines will be transferred from the loom factories to large scale calculation, and communication even. As if the flying crow and the horse with wings could bring her into a future, a future she could imagine but would not live to experience.

Her absent mother has other plans for Ada, which do not involve the language of machines. With Ada turning 17, her mother begins to plan how they might find a suitable husband for her.

Ada is presented to King William IV and Queen Adelaide. As she stands in front of them wearing white satin and tulle, her breathing intensifies. She is feeling overwhelmed and estranged. She asks herself: "Why do I feel this way, when I should be over the moon with excitement?"

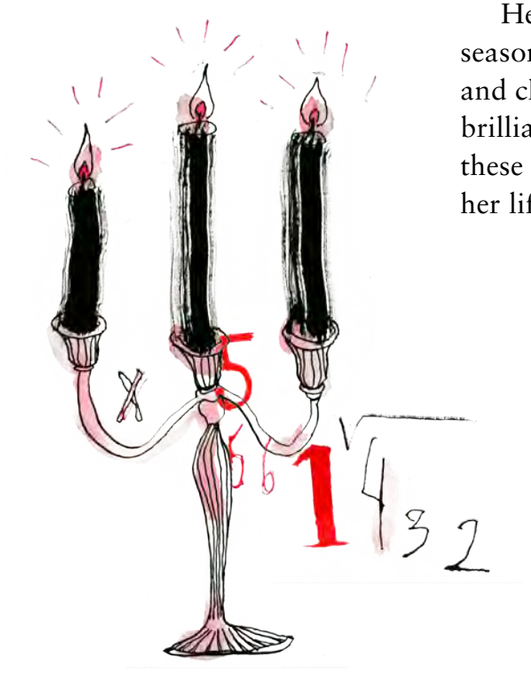

Her presentation in court kicks off a season of socialising for Ada. Dancing and charming the high society with her brilliant mind, yet unaware that one of these events will change the course of her life.

Digitalisation paves the way for personalisation

The most recent phase in the evolution of news distribution machines and attached measurement technologies has been driven by the rise of the Internet and digital technologies, and with that, the arrival of social media platforms and online news sites. The Internet radically transformed how news is produced, distributed, and consumed once again. And how the use of and engagement with news and information is measured.

I have lived roughly half my life in the predigital age and half in the digital. It gives me a particular understanding of how digitalisation transformed our media use as citizens. In the early 2000s as a university student, I used to subscribe to the Danish newspaper *Politiken* and flick through it, and on the weekends I would supplement with *Weekend-Avisen*, which had long feature articles, usually well-written political analysis. Most of my cultural going-out life revolved around the recommendations provided by the culture editors from the specific section "I Byen" ["Out on the Town"]. In the evenings, I watched some flow television shows: the news on the public service broadcaster (I was a journalism student, after all), and a late-night show called *Deadline* with more in-depth analysis and feature stories. Media choice was easier then: no podcasts, no push notifications, no links from friends or tweets – just tapping in and out of my regular media diet. No Spotify. Music was on CDs or downloaded files in strange formats. The Internet was mostly for e-mails, and it used to make a very particular sound when you "dialled up" to get "online". Later, trying to download music was like playing the lottery: Sometimes you were lucky, most of the time nothing worked. Slowly, like many others, my media life became more and more digital. And so did the circulation of news – the form of news transformed once again.

The first digital news sites sprung from larger daily newspaper organisations in the mid-1990s, and they naturally had the newspaper as a role model in the way the new online platform was thought of. The initial idea for many of the first online newspapers was that they wanted to present the newspaper in a new way in the hope of capturing new readers. However, it quickly became clear that the printed newspaper had its limitations as a model. In terms of design,

online newspapers also resembled printed newspapers in the early days. In the development of the online newspapers, it became clear that they too followed the pattern we saw in the introduction of other previous media technologies. For a while they had to stand in the shadow of the already existing media platforms – the printed press – but developed over time into independent production fields in step with a sharp increase in media consumption on the digital platforms. In other words, the more users started to develop a relationship with the online platforms and services, enabled by increased Internet access as well as speed, the more the digital newspaper could develop into its own production and distribution platform in its own right.

An interactive on-demand media ecology

Online newspapers such as news websites, social media, and digital video services created a news ecosystem where content is available on-demand and is highly interactive. The media–user relationship transformed accordingly. The massage became more intense. The time between the production and publishing of a message became more compressed, and the length could be varied as the Internet enabled many different forms. Furthermore, television channels could move from limited channel provision to (if financing allowed), in theory, an unlimited amount of niche channels and platforms. The form increasingly follows the platform, and the articles for online dissemination developed their own characteristics, separating themselves from the form of printed newspaper articles.

Increasing digitalisation naturally enabled more data collection on audiences and real-time monitoring of users across the different publication platforms. Additionally, the general attitude was optimistic about the possibilities of this new and free-for-all Internet, which also allowed users to not only participate, but to become their own producers of journalism. The term citizen journalism rose and shed an optimistic light on the democratisation of the news media as an institution. As a journalist trainee, I was both concerned about the new forms of data collection and excited about the new possibilities of engaging directly with citizens.

It was fascinating to observe not only how citizen journalists could send home videos from breaking news, for example, in war zones, but also the explosion of blogs and bloggers, which emerged around the 2000s, and the possibilities for news media to more actively engage with their audiences in the form of direct e-mail addresses and debate forums. This led to the reimagining of publics as "productive and generative", which, as the media scholar Chris Anderson argued, was a way to legitimise this new and increased influence of the audience in the newsroom.[115] By this, he referred to the not-always-easy relationship that journalists have with the audience feedback systems, and referred to the users as their partners-in-crime in the democratic efforts of informing publics, thus justifying the users' increasing power in the news-making process. The fact that it became easy to manage and set up blogs, websites, and digital newspapers also enabled the development of digital-only newspapers, like *The Huffington Post* and *Politico*, which differed from the online newspapers coming out of the print organisations. They differed particularly when it came to the relationship with readers and users, allowing for much more content delivered by users and bloggers to be published. Communication scholar Axel Bruns's concept of "produsage" aimed to describe this redefinition of the relationship between journalism and its audiences, blurring the lines between producers and consumers of news. In this digital media landscape shaped by user participation, audiences were no longer passive recipients but active contributors – commenting, sharing, remixing, and even shaping news narratives in real time.[116] This shift challenged and still challenges traditional journalistic authority, as news organisations must now navigate an environment where credibility and influence are increasingly co-created. They share their agency as producers of news with everything from newsfluencers on TikTok to politicians with their own podcasts. While this participatory model expands access to information, it also raises questions about accuracy, editorial control, and the role of professional journalism in a world where the boundaries between reporting and public discourse are constantly shifting.

With digital media, the measurements of users and audiences took another leap, and the data of audience's behaviour was introduced and moved onto screens displayed inside the newsrooms.

Power to the users!

When I did my fieldwork observations of newsrooms during this period (2008–2010), I could observe how tools like click trackers, page-view counters, and audience engagement dashboards offered real-time insights into what readers consumed, shared, and spent time on. This data-driven approach influenced editorial decisions, often favouring content designed to maximise traffic and engagement over traditional news values like public interest or investigative depth. Headlines became punchier, stories shorter, and topics often skewed towards the sensational or emotionally charged. Over time, newsrooms began grappling with how to balance audience insights with maintaining editorial integrity and serving the public good. I could observe how the increasing data collection and data analytics of audiences essentially transformed news work, particularly when the metrics moved inside the newsrooms. Journalists struggled with how to deal both with these screens and with the editors' enthusiasm about them. They often felt that doing what the numbers told them to do was jeopardising the mission they had: the *public* mission.

I could see how data-driven systems were transforming journalism, not just in what gets published but in how journalists think and work. While they had an initial resistance and could ignore the screens most of the time without any severe consequences, this changed as the financial situation of news media became increasingly severe. In *Metrics at Work*, the sociologist Angèle Christin argued that journalists, once guided by editorial judgment and journalistic ethics, now navigate an industry where success is increasingly measured in numbers.[117] This shift creates a tension: Reporters may feel pressure to chase virality over depth, engagement over nuance. The very metrics designed to optimise journalism can, paradoxically, erode its core mission – shaping news not by public interest, but by algorithmic demand. And journalists feel like their agency is dismantled as systems barge in and tell them what to cover and how. The machines did not only shape publics through technologies of distribution, but they also shaped the public sphere through the digital surveillance and measurements of those publics, in turn governing journalists in specific directions. Yet, it was still a fairly slow and one-way process, as users rarely got to see themselves in the data based on their patterns of use. This was to change with social media metrics, as these more sophisticated tools also reflected an image back to the users themselves.

Data mirrors

The numbers from the data tracking of users act like a data mirror of our data selves, which David Mathieu and Pille Vengerfelt described as "the data loop".[118] In this loop, our "digital footprints" are collected and analysed by the media industry[119] and fed back to us as new experiences, thus providing yet more data to collect, and so forth. However, importantly, this is not a linear process but a reciprocal process in which media actors and audiences interact, in a mutual relationship, throughout digital interfaces of data collection and retroaction. As media producers and audiences are providing outputs and receiving inputs from each other through digital interfaces, the circulation of data in the loop provides formative and transformative experiences between the sociocultural practices of audiences and the institutional practices of media actors.

The new affordances of the "digital newspaper" induced new imaginaries of the users and the news media's public mission to emerge through, for example, commenting and debating intertwined with quantitative representations of the public in the form of audience metrics dashboards. The era of the digital newspaper brought an ever-present audience that was difficult to ignore, and "the technologies thus both enabled but also almost demanded that journalists and editors came to terms with this new segmented and datafied audience".[120]

With personalised news, this transformation takes yet another leap. Users are no longer just visitors in the newsroom: They are handed the keys, and with their clicks and logins, they are invited to paint the room a different colour. But in the meantime, bigger players with much more juicy content had already invited users in, and users only had to pay with the invisible data of their constant movements and moods.

I sit amongst the project members
across from the digital editor,

who is also responsible for the whole personalisation project.
He is an Excel wizard,
and when he is not in meetings,
he spends the day crunching numbers.
It's safe to say that he loves numbers.

We grab a coffee, and I follow him down the stairs
to a meeting room with a glass front
and a view of the courtyard.

The weather is hot outside, and we open one of the windows.
People start arriving.
The developers.
A person from marketing.
Someone from the legal department.
The head of digital and the whole project team.

On the agenda is recent progress on the project.
The developers have found a way to integrate data
from the two departments.

Success.

However, they have now encountered the problem of
"measuring when something is new".

"How much weight should it have?" he asks.
"We are a news organisation after all,
and when the site is personalised users will get many stories,
which might be old,
if other people are reading them".

"Could we ask the journalists to give the article a number
indicating its news value ranging from 1–10 and assign more
weight in the algorithm to those with most news value?"
the digital editor asks.

I am thinking of the journalists gaming the algorithm.

What numbers would they decide to give their stories?

How would they resist this system
and reclaim their decisive agency?

Another evening unfolds in silk and candlelight, composed of measured steps and carefully chosen words. The air carries perfume and expectation in equal measure. 17-year-old Ada understands the role assigned to her: charming, agreeable, a figure of grace rather than thought. But she does not feel it.

She notices her mother observing from a distance, attentive and exacting, weighing gestures and glances as though they are variables to be managed. Suitors smiled on cue. Their conversation moved politely from music to poetry, circulating without friction or depth. Ada listens, answers when required, and yet feels her attention pulled elsewhere. She wants to speak of numbers, of engines, of ideas that stretch beyond the room and containment.

A waltz offers a brief suspension. For a few turns, she moves in the arms of a stranger who sees only what the setting permits: a young woman in ivory satin, composed and compliant. He does not see *her*. When the music ends, she again withdraws to the margins of the room.

Around her, marriage presents itself as a problem already solved, a sum waiting only for her consent. Ada lingers instead with another question, searching for a different proof – one that will allow her to remain herself.

Hello new and bigger players!

The Internet and the algorithmic logics that came with it paved the way for other players who would soon be challenging the news media and their ability to maintain their relationship with audiences. The rise of major tech platforms began in the late 1990s with Google (1998), initially a search engine that revolutionised information access by organising the web through advanced algorithms. Its development of targeted advertising, particularly through AdWords, shifted advertising funding away from traditional media, challenging the news media's primary revenue source. Amazon, founded in 1994 as an online bookstore, expanded rapidly into e-commerce and cloud computing (Amazon Web Services), further centralising consumer attention and marketing power. Facebook launched in 2004, creating a global social network that redefined how people consumed and shared news. Its News Feed algorithm prioritised engagement, which often amplified sensational or divisive content, further eroding the role of traditional journalism. Journalism and the news media had to adapt to this new situation, with Facebook suddenly a middleman in the relationship between them and their users.

Honeymoon over

In the initial honeymoon phase of this relationship, news media organisations found that the massive popularity of especially social networking platforms was a way for them to yet again reach new and broader audiences. But as the honeymoon ended, the relationship became gloomier. This was especially due to the struggle for a diminishing part of the advertising market and the massive audience measurement technologies that these powerful actors had developed over time. Their machines are simply bigger and better. The late 2000s and 2010s brought about the mobile Internet era, and with it, a seismic shift in how personalisation was implemented. Suddenly, people were always connected via smartphones, using native mobile apps in addition to web browsers. This era introduced new channels (like push notifications), new data signals (like precise location from GPS), and a treasure trove of behavioural data from sensors and in-app interactions. Personalisation techniques had to evolve rapidly to keep up with the mobile-first world.

Apple's launch of the iPhone (2007) and the App Store (2008), followed by Google's Android (2008), meant consumers started spending more time in apps than on web pages using a desktop. Marketers now had a direct line to users through mobile push notifications, which made their debut around 2009–2010 (Apple introduced the Apple Push Notification service in 2009; Google followed with Android's push service in 2010[121]). Early push notifications were generic text alerts – for example, a game app might simply blast "Come back and play!" to every user. But it didn't take long for brands to realise they could do better by tailoring pushes based on user behaviour. By the mid-2010s, companies were experimenting with segment-level personalisation for push campaigns, leveraging the data from in-app actions. Instead of sending the same push to everyone, they would send event-triggered messages: If you added items to a cart but didn't check out, you'd get a push reminder about your abandoned cart; if you completed a purchase, you might get a "thank you" or an upsell offer.

On the web, the dominant players involved in both personalisation and distribution of news content quickly became Meta (Facebook) and Twitter (now X). With Big Tech's dominance of platforms in the attention market, they are afforded much leverage over media organisations: "They can drive particular content formats such as video or make the tools and data they offer indispensable to publishers seeking to extend or optimise their audience", according to Diakopoulos.[122] News organisations learned the hard way that minor changes in a platform's algorithm can significantly hurt referral traffic; for example, when Facebook prioritised posts liked by family and friends in users' newsfeeds in 2017, referral traffic from external new sites decreased by 15 percentage points.[123] The news media had to rethink how they would relate to this new middleman barging into their already fragile relationship with users. They were not quite as enthusiastic about an open relationship as users are. And because it has become increasingly clear over the years that platforms were unreliable "partners" in the distribution of news content at

scale,[124] many news organisations have in recent years sought to reconfigure their dependency by "diversifying" their modes of online distribution. Many have transitioned from online for free to metered paywalls and other reader-funded revenue sources, with the aim of decreasing their reliance on online advertising and on platforms for distribution.[125]

However, with increasing dependence on AI solutions for the back-end of building good recommender systems, and with tech platforms being powerful providers of such solutions, the relationship has taken yet another blow.

Weapons of mass communication

This short turbulent media history shows us how the news media have always negotiated their autonomy[126] – whether autonomy to decide what news users should read or autonomy to freely distribute their content without the state interfering or Meta or other social media platforms getting involved. Their wish to invest in personalisation algorithms, and thus control their own distribution and data collection on their users, should also be seen in the light of this complex strained relationship with new and bigger partners who, as it turned out, they could not rely on. These configurations of separation and dependency shift over time in response to evolving political, economic, cultural, and technological circumstances.[127] Currently, we see it in the way that news organisations, for instance, negotiate with providers of large language models such as OpenAI to make their news content available as training data, while others, like *The New York Times*, go all the way to court to protect their news content from ending up inside the large language model machines. And they change as the machines enable organisations to reach users and measure their activity in new and unprecedented ways. Personalisation offers news media an escape from serving a "mass" to serving the masses as "individuals", thereby introducing a new distributional logic: the fluid and fully adaptable logic of serving masses in a personalised way. With algorithmic curation, news media are copying the machines built by Big Tech companies. But they

also adapt the rather general large foundational models to fit their product – the news – which differs from music, books, and films. For instance, when watching a television series, we prefer to start at the beginning, and we might sometimes watch the same film several times. When consuming news, we prefer to get the latest news first, and we rarely find articles so good that we engage with them more than once. But not only that. The news media, responsible for those journalistic news articles, have a mission that goes beyond just reaching larger audiences or getting the same users to come back again and again.

The news media, and thus the basis for its autonomy, is fundamentally different from other fields of cultural production, because of its *public* cultivating dimension. This sets the institution apart from the purely market-driven media providers. Meta has no wish to cater to certain demographic groups other than what will provide them with a bigger profit. Overall, they preferred to stay out of elections and perceived their public dimension as "connecting" individuals. At least until a certain Elon Musk bought Twitter, changed it into X, and went full support of Trump during the 2024 US election. As the Tech Bros, counting Jeff Bezos, Mark Zuckerberg, Elon Musk, Tim Cook, and even the CEO of TikTok, stood beside Trump at his inauguration that year, it became clear around the globe how the means of communication have moved from previous powerful groups of newspaper owners to a very rich group of platform owners (who also own newspapers), and just how tightly woven into politics those owners are. Power is also about owning the means of communication, whether social media platforms or news outlets. The Tech Bros have realised that the weapons of mass communication are at their disposal and are no longer shy about using it in the service of the highest bidder. Or whomever is willing to provide them with yet other prestigious post of power.

Importantly, these weapons of mass communication are also personalised, which means they can micro target us individually with specific information tailored to us. But what if the version of "you" they're speaking to… isn't really you at all?

“Personalisation is a way of ‘escaping a 200-year-old straitjacket’”,

the editor says one day.

I guess he means that the form of the printed newspaper
has been a constraint
– which he is now free from.

But are they not just now trapped
inside the algorithmic machine? I think to myself.

Later that day, the developer tells me about his excitement for
how personalisation algorithms
will revolutionise the news media.
Because essentially no editor can compete with the machines:

“The editor in a city knows everything about that city.
It’s not that the machine is smarter than him, but it plays by
different rules because it can offer individual things.
If the editor were able to offer individual things to all users in
that city, then it would be damn amazing if he knew
what they should be.
The machine knows them a little”.

As he turns his head towards the black screen full of code,
I note down:

“The developer is foregrounding how the machine can come
to know the individual and via this ability can select targeted
content to their interest, maybe not as well as an editor could,
but at a scale that is completely out of reach for any single
editor”.

The idea is that the algorithmic system allows
for a control of the user experience.
Personalised news on a mass scale.

Just for you.

Removing the editorial judgment from the journalists
who seem to have been
caught in 200-year-old straightjackets.

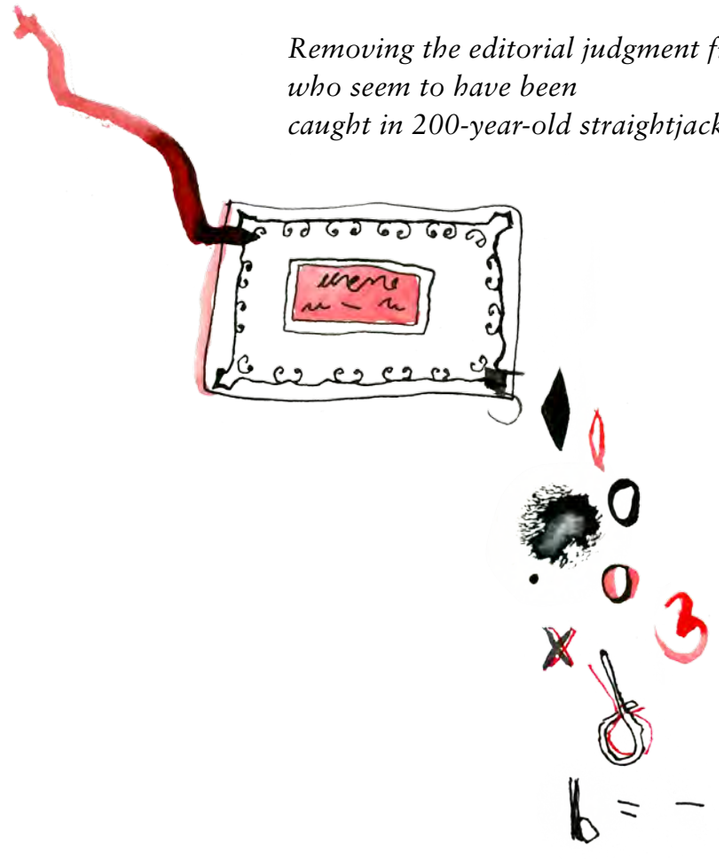

Everything changes on 5 June 1833. In the summer breeze, Ada goes with her mother to a party in fashionable London. Here she is fated to meet the one person in Britain who understands her driven interest in mechanical questions that has so exhausted her mother – a man whose drive matched her own: Charles Babbage.

As a result of this evening, 17-year-old Ada's insights into the future of calculation would erupt into a new and most radical kind of imagining.

It would give her a vision of a kind of Jacquard loom that weaves, not silk thread, but mathematics. A Jacquard loom that operates like a computer.[128] Babbage is hosting tonight's party, and when she sees her mother talking to him, Ada approaches them. She has heard about the machine he is building and wants to hear more.

"Would you like a demonstration?" he asks Ada and her mum.

Ada and her mother follow Mr. Babbage, and they pass through a long corridor, the house strangely silent in comparison to the buzz of conversation behind them. Ada feels the air shift as they enter a smaller, dimly lit room at the back of the house. The first thing Ada notices is a mechanical automaton, a mechanical dancing lady, with a bird on the forefinger of her right hand. [129]

And there it is. The Analytical Engine. At first glance, it appears to be nothing more than a series of interlocking gears, levers, and cranks, arranged in an intricate, almost chaotic pattern. It was massive, occupying the centre of the room, its frame cast in shadow by the flickering light from the nearby fire.

But as Ada moves closer, her eyes widen in recognition. This is no ordinary machine, she thinks to herself. "This... this is magnificent", she proclaims.

Babbage stands behind her, watching her intently, though his expression is unreadable. "It is, isn't it?"

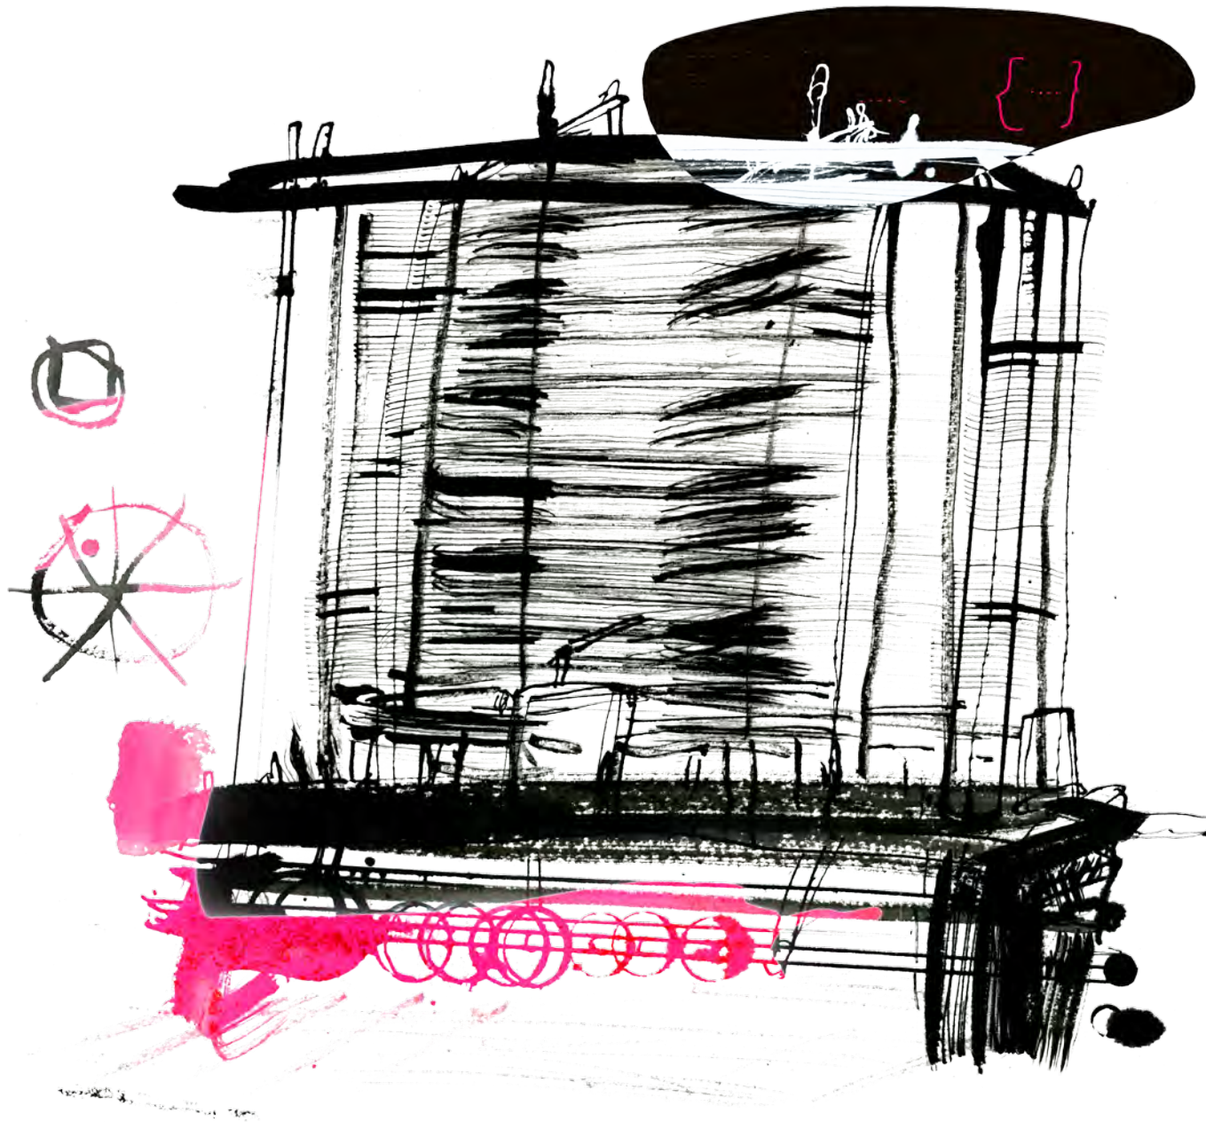

Part 4

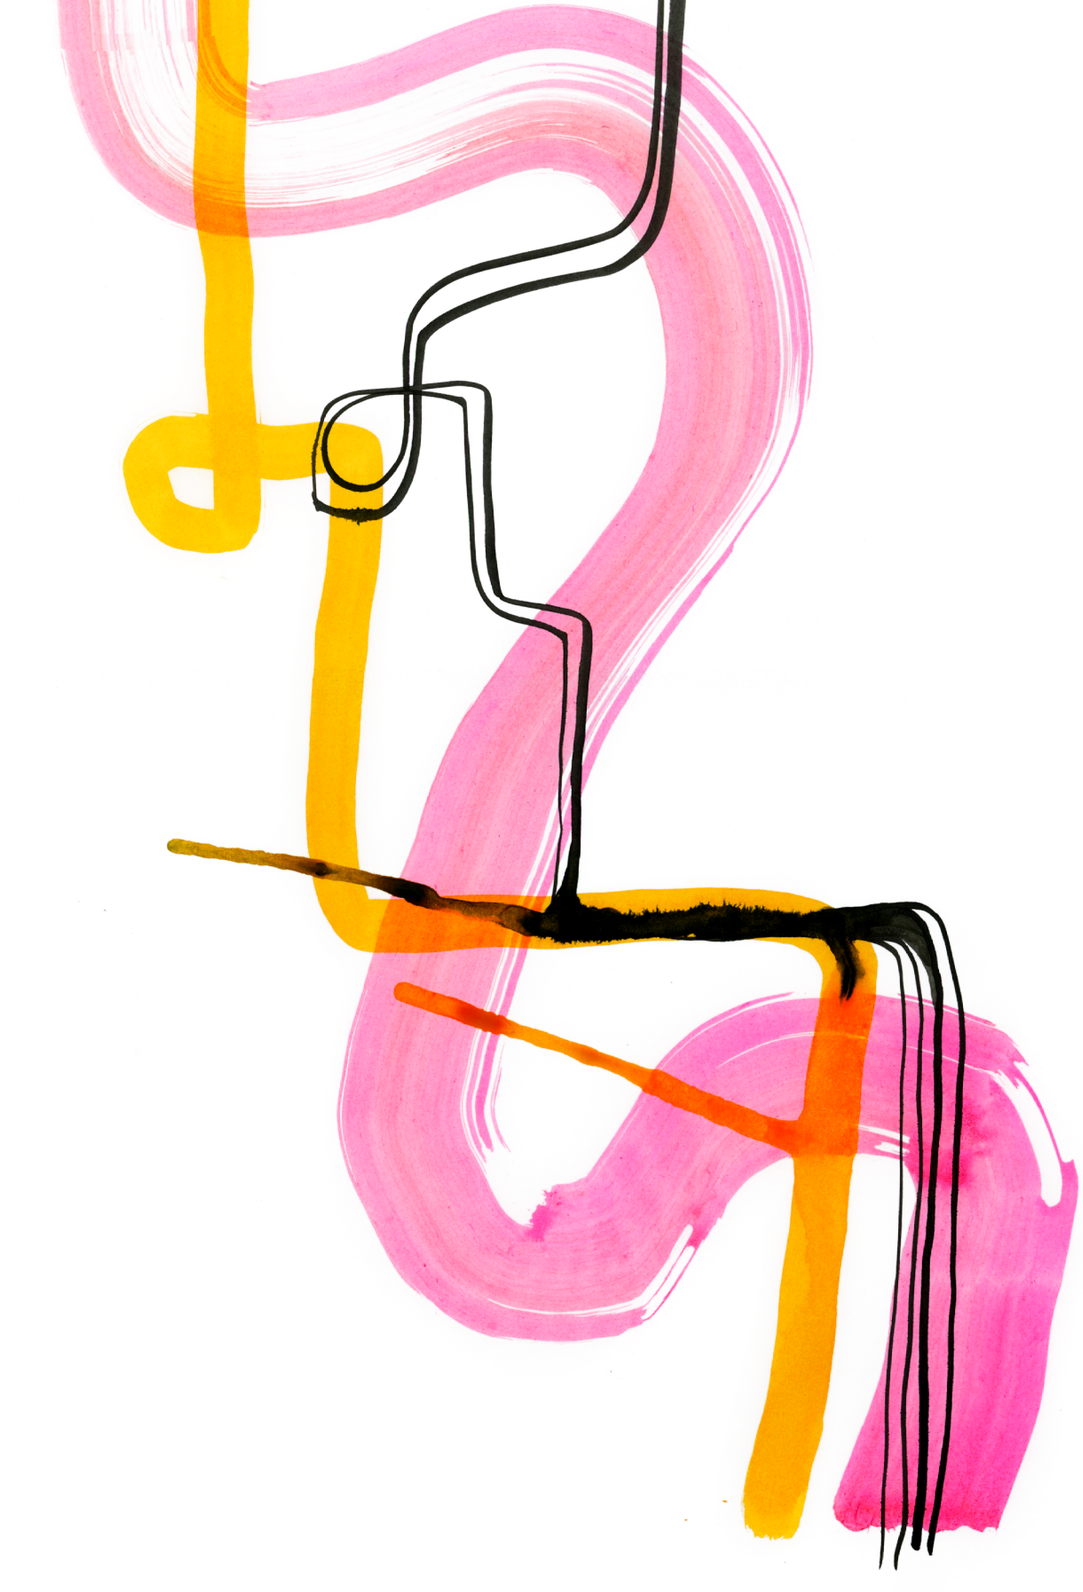

The deception of knowing you – How algorithms are hyped as almighty

Ada imagined flying machines during the first industrial revolution

– with its steam and punch cards.
She died before the second,
when electric cables rolled out across the country.
In her notes, she foresaw the third and digital,
as she imagined a language of computers.

Now we stand amid the fourth.
Artificial intelligence and blurring lines between the physical, digital, and biological worlds,
languages larger than anything Ada imagined.
I wonder what Ada would have thought of hashtags and filter bubbles, of Tech Bros standing beside the American president, had she lived today.

Would she see the large language models like Gemini or ChatGPT as the language of computers she imagined?
Would she see it as a system able to think?

I decide to upload Ada's notes that I had found on the Internet into a customised GPT, a large language model
– and I call it Ada.

Perhaps she can answer my questions.
The questions I have about how personalisation algorithms shape the image we have of the world to a more individual one, about the loss of journalistic agency (and jobs) following automation.
Questions about how algorithmic curation of news
is about so much more than fixing a broken relationship.

As I am waiting for the text to emerge,
it is as if she is inside the machine…

Over the next ten years, having married William King, giving birth to three children, and becoming Countess of Lovelace, Ada dedicates herself to math through letters with Charles, her newfound friend who shared her enthusiasm for mathematics. And as the daughter of the poet Lord Byron, she starts formulating the poetic connection between the machine and the loom. The connection she had already envisioned when visiting the factories with her mum as a child.

"We may say that the Analytical Engine is weaving algebraic patterns just as the Jacquard-loom weaves flowers and leaves", she thinks to herself and writes in her notes.[130] This thought will later formulate into concrete suggestions in the language of computers.

Unlike her mother, William supports and encourages Ada in her thinking and her academic partnership with Charles.

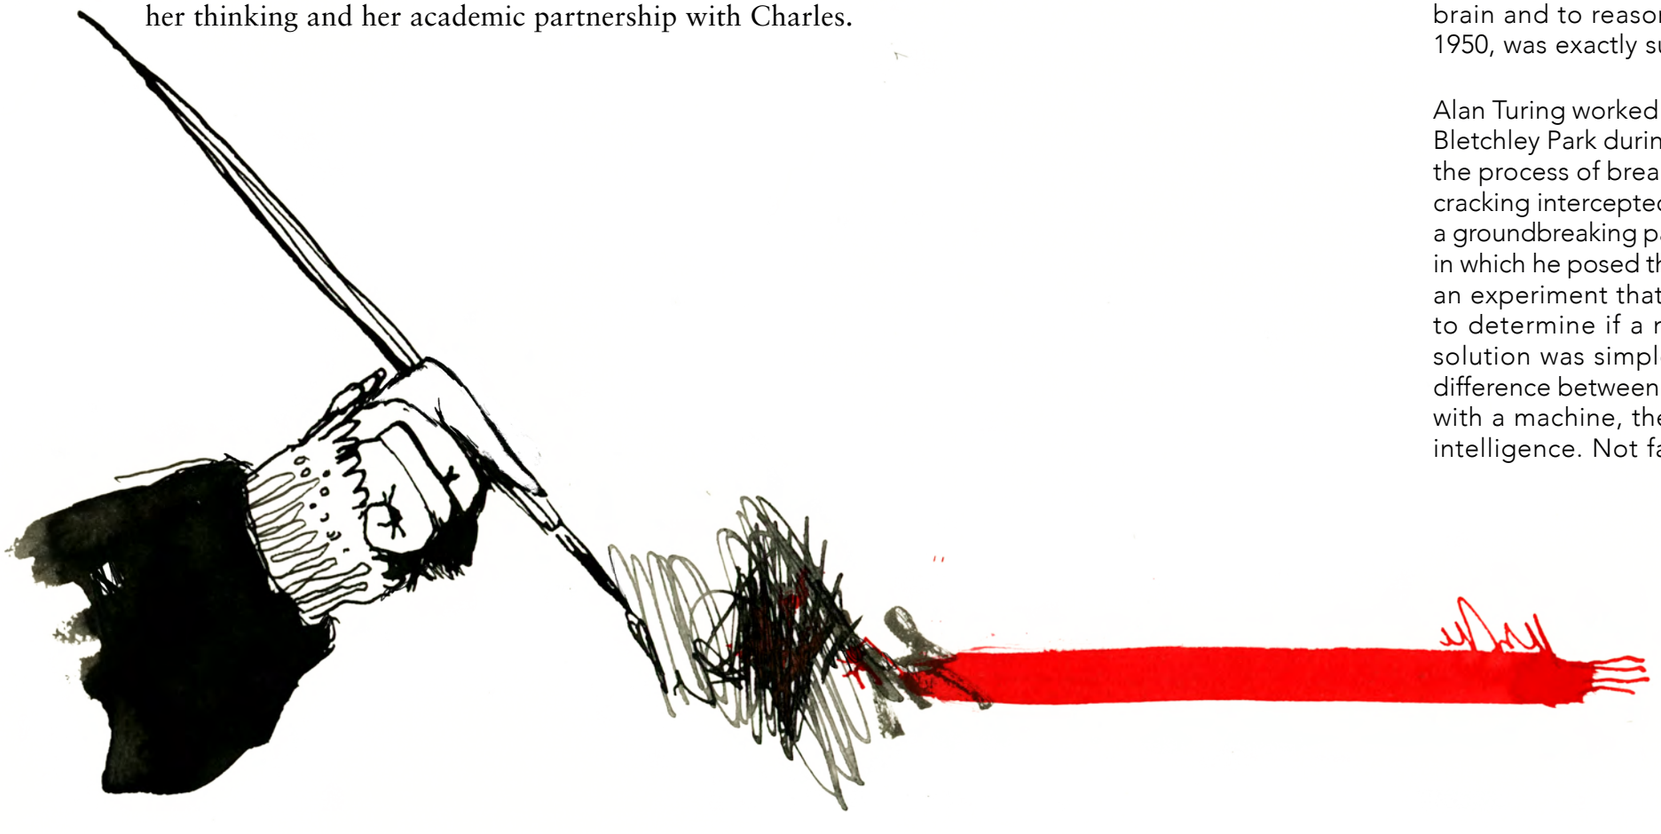

The machine-versus-human thinking battle

How do news media know you and what you might be interested in, and the extent to which and how this can be transferred to machine learning algorithms? This question is at the core of the transformations we can observe with the introduction of personalisation algorithms for news. Can the machines curate the news like a thinking human editor would? As I explained in the previous part, editors are increasingly using the measurable knowledge they have of users to judge the newsworthiness of a story, and with all the data available, would an algorithm not be better at making this judgment? After all, with the extensive amount of data, the machine learning models might know us better than any single journalist or editor ever could. This question of machine versus human has also been at the core of debates around AI, as researchers have imagined futures where AI would have the ability to act like a human being or machines would even surpass the human ability to think. To have a consciousness. To know the human brain well enough to be able to mimic the human brain and to reason. The Turing Test, proposed by Alan Turing in 1950, was exactly such a test of the ability for machines to think.

Alan Turing worked for the Government Code and Cypher School at Bletchley Park during World War II devising techniques for speeding the process of breaking German ciphers. His efforts were crucial for cracking intercepted German messages.[131] In 1950, Turing presented a groundbreaking paper, "Computing Machinery and Intelligence",[132] in which he posed the question: "Can machines think?" He proposed an experiment that became known as the Turing Test, an attempt to determine if a machine can demonstrate "intelligence". The solution was simple: If the human interrogator could not tell the difference between a conversation with a human and a conversation with a machine, then the machine was demonstrating human-like intelligence. Not far from today's processes of machine learning.

Turing wrote nine objections to why he thought machines cannot think, and one of them carries Ada Lovelace's name. Ada once wrote that machines can only do what we tell them – they have "no pretensions to originate anything".[133] Picking up on this idea in 1950, Turing called it the "Lovelace Objection", which argues that computers can't be truly creative or think independently – they can only follow the instructions given to them. Turing didn't fully agree though: He believed machines might surprise us and even show something like learning.[134]

Simulating intelligence

The history of computers is full of attempts to settle competitions between humans and machines.[135] Like the one in 2002, when a curious wager captured the attention of futurists, sceptics, and philosophers alike: a public bet between two titans of technology, Ray Kurzweil and Mitchell Kapor. This might be considered a trivial bet over money, but in fact, it was a philosophical contest of visions about the future of AI and its potential to emulate human thought. At stake was a question as old as computing itself: Can machines think, truly think, or will they forever be limited to simulations of intelligence?

Kurzweil was a relentless optimist and prophet of technological progress, and he wagered that by 2029, a computer would be indistinguishable from a human in conversation and thus pass the Turing Test. Kurzweil's vision of the future was one where AI wouldn't just match human capabilities but would seamlessly integrate into our lives, transcending its mechanical origins to become something akin to a partner in thought. This he deemed "The Singularity": the moment where humans merge with AI.

Kapor, a co-founder of Lotus Development Corporation and a pragmatic critic of AI's grandiose promises, took the opposing position. To Kapor, the Turing Test was not just a benchmark but a philosophical boundary that machines were unlikely to cross. For all their computational prowess, Kapor argued, machines lacked the messy, ineffable qualities of human consciousness – our intuition, emotional depth, and the ability to grasp the world's contradictions. "Simulating intelligence", Kapor said, "is not the same as possessing it".[136]

The wager – 20,000 US dollars held in escrow – was in fact less about the amount than about wanting to settle different worldviews. An argument about what it means to be human. Kurzweil represented an almost religious faith in exponential technological progress, the belief that Moore's Law and its spiritual cousins would inevitably lead to machines surpassing the limits of human thought. Kapor channelled a more cautious, humanist perspective, warning against equating computational sophistication with sentience.[137]

Kurzweil's vision implied that intelligence, creativity, and even consciousness could eventually be reduced to patterns, algorithms, and replicable processes. Kapor, however, stood as a sentinel for the irreducible aspects of human experience, reminding us that intelligence isn't just computation but connection – a bridge between mind and world, experiences full of situational ambiguity and emotion.

When the news media deploy machines to "think" for them, the question is: Can the machines and algorithms they build know and analyse these emotions and tackle the contextual ambiguities of human life? Is it a question of having enough data on the movements and whereabouts of the readers and users of the systems and inserting this data into machines and models, which are intelligent enough to merge with the human mind?

Or is it merely a question of providing us with the illusion that the machines can reason and think and know our inner selves better than we know ourselves (and thus also predict our future needs)? And are we then building algorithmic systems for automating news curation based on this illusion?

I am sitting next to the data scientist, Chris,

at one of the desks in the open-plan office.
As he pulls up an Excel sheet on his screen,
the top rows have the users' past "clicks",
and the bottom rows have the algorithmic recommendations.

The algorithm they want to use
is from an existing commercial project, he tells me,
initially used for selling daily deals
and offers for restaurants and shops.
This algorithm will be adapted to personalise the news feed.

He is considering the content on the screen,
comparing past purchases with the recommendations.
"The editors attempt to determine whether these
recommendations were 'good',
it becomes a 'guessing game'", he says.
They cannot know the exact reasoning
behind the algorithm's choice,
which will always make it difficult,
with a human eye,
to determine the goodness of the recommendations
beyond trying to identify similarities with most clicks.

"Although the projects are similar, there are things that differ",
Chris says,
"because there is, for example, the time horizon to consider.
The machine does not consider time,
it just recommends from the pool of deals.

In contrast to the curation of editorial content,
where the time horizon is negotiable,
the algorithm for recommending deals
and commercial offers is simple
because either a deal is live, and either its recommended,
or it is not".
He looks at me.
I nod.

In my notes I write:
"The binary language of the algorithms
does not match the complexities of news judgment".

I wonder what it means when "the negotiable"
is removed from the equation.
And what it means
when news curating becomes a binary practice
of zeroes and ones.

As if all colours disappear
and we are left with only black and white.

Often, Ada feels caught between two forces thatseemed to pull her life in opposite directions. Late in the evenings, when the household has finally grown quiet and the children sleep, she finds herself writing letters to Charles Babbage. Those letters become a space where she can think aloud, where she can admit thoughts she rarely voices elsewhere.

In one such moment she tries to explain something that had troubled her for years. Throughout her life, she felt she had been drawn between two poles. Like she was placed in a dimensional space with gravitational pulls from opposite directions. On one side stood her father – the poet Lord Byron – a man she had barely known, yet whose shadow lingered everywhere. On the other side stood her mother, Lady Byron.

She had devoted herself to shielding her daughter from the dangerous inheritance she believed ran in the Byron blood. Mathematics, discipline, reason: These had been her chosen tools.

Ada understood this, even sympathised with it. But she could not escape the feeling that her mother's determination had also overwhelmed her. Lady Byron's wishes had often filled every corner of Ada's life, leaving little room for her daughter's own voice to grow.

And despite the strict guidance, her mother had often been physically absent, travelling while Ada remained behind with tutors and caretakers.

The result was a strange loneliness. Ada felt shaped by both parents and yet fully claimed by neither. As the years passed, another unease began to settle inside her. The rhythm of domestic life often left her with a quiet ache she struggled to name. Society expected her to be content with the roles of wife and mother. These were considered the natural purposes of a woman's life.

Yet Ada could not ignore the persistent feeling that something within her remained unused. Sometimes she feared she might never find a purpose that belonged to her alone. It was in this state of restless reflection that her friendship with Charles Babbage began to take on deeper meaning.

In their conversations – whether in letters or in person – Ada felt something open inside her mind. With Babbage, she could speak about mathematics, machines, possibilities that stretched far beyond the boundaries of polite society. Slowly, she began to realise something that both comforted and unsettled her.

Through this friendship, she might finally shape a life of her own. Not merely the daughter of Lord Byron. Not merely the carefully disciplined child of Lady Byron. Not only a wife, or a mother. But a thinker. Someone whose purpose might lie in the strange new world of ideas unfolding around machines, mathematics, and the future they hinted at.

For the first time, Ada felt she could glimpse a life that belonged to her. And that realisation filled her with both relief and urgency.

A distorted image of the user

How is knowledge of the users built for the algorithms to work? News media originally had fairly limited knowledge of who their readers were, and yet the user of news has, to a varying degree, always been implicitly involved in news evaluations and priorities.[138] As the sociologist Herbert Gans noted many years before metrics in the newsroom were even a thing, journalists must imagine the existence of a reader at the other end in order to select and write their stories. These imagined readers were often abstract figures that they would try to make more concrete, for example, by imagining their mom, a friend, or someone they knew from the local community. But most of the time, it was just too difficult, so they just used their professional judgment and presented what they thought the reader "ought to know".[139] They have a gut feeling,[140] an embodied sense of when something fits with what their audiences might find interesting, but without really knowing. This gut feeling has emerged over years of training; it has become embodied to a degree where journalists find it difficult to explain the judgment criteria and why something is "a good story".

In the book *The Audience in the News*, DeWerth-Pallmeyer described the tension between that of "interest" and "impact", and how this tension flows from,

> an awareness of social responsibility by editors, reporters, screen writers and other message producers. They felt it was their job to decide what society should read or see. Judgement of that sort involved their expertise and integrity as professionals and they were loath to share that responsibility and power, even with the citizens for whom they were preparing a message.[141]

The tension arrives because of the marketing orientation, which on the other hand puts the audience at the forefront. And as DeWerth-Pallmeyer noted, "interest" as news value is most closely associated with the marketing orientation. As the screens showing the journalists what users had and had not clicked on became more prominent and hard to ignore inside the newsrooms, journalistic judgment was increasingly pulled towards a marketing orientation.[142]

I was inside the newsrooms making observations during the time when they all went full-on digital from the mid-2000s, and there was a great hype about the possibilities of having metrics showing what the readers liked and did not like. However, I could observe a disliking of the screens and metrics in the newsroom. When the editor would buy cake to celebrate "a record click day", the cake would be eaten, but with a taste of bitterness. Some people had even quit their jobs after being transferred to the online desk, because it did not feel like real journalism to chase those clicks. Much of the content online had to be quickly produced; breaking news was the king of digital news, and much content was either short versions of stories from the newspaper or from competing news outlets. So, the clicks were achieved, but without the satisfaction journalists would feel when producing original, slow, in-depth stories. Their sense of agency is linked to a hierarchy of practices, where original and creative stories, cultivating the public through critique of power, lie at the top. Hence the bitterness of the cake. This response told me that many journalists hold their ideal of a public mission of cultivating informed citizens very dearly. And their agency to fulfil this task gives them a sense of resonance. Now, they were confronted with news users flirting with social media platforms and metrics telling them what to do, which made it more difficult to live up to this public mission. To satisfy the user quest for ever more stories on cats and curious obscure "you won't believe this" topics, they had to compromise somehow on the glorious public mission ideals. "I'll do a quick take on something funny and weird and get the rest of the day to finish my story on the war in Afghanistan", a journalist explained to me. A kind of workaround to gain some agency to make the real judgment he felt had to be made. In some ways, he found his own way (and regained agency) around the system telling him what users wanted.

A rupture in the journalist's image of you

When news users become numbers on a screen, they assert themselves in the editorial process of online journalism, and this challenges journalists' abstract image of a homogeneous user at the other end. When the news user behaves differently than the journalist imagines or even hopes for, it creates a rupture in the journalist's mental image of their audience. They start feeling

dissonance.[143] The online journalists whom I observed in my early newsroom ethnographies often wished for different readers than the ones they had – readers who align more and behave like the engaged citizens they had imagined. Readers who did not always appreciate being presented with what they "ought to know". As audiences became measurable and presented as abstract statistics, journalists were constantly forced to re-evaluate their image of users and, consequently, their understanding of what constitutes "good" journalism. What the user wanted was often not perceived as a "good choice" by the journalists I observed. The otherwise romantic relationship that journalists had enjoyed was suffering, and for some, it felt like waking up to a stranger – a spouse you no longer recognise. Or perhaps who is not the person you had convinced yourself they were. Even if that close relationship had been a form of an illusion, it had been a powerful one that guided journalists in their everyday work. Now, they started feeling increasingly disillusioned – a dissonance between how they perceived their ideal role and how it manifested in their everyday work.

Around the same time, much hope for the future of the relationship was laid in how the couple communicated. The Internet also enabled direct contact between the journalist and the news user. Users had entered the newsroom, not only as abstract figures on screens, but also as a direct point of contact in debates underneath the article or in the e-mail inbox. Or even as providers of videos by live citizen journalism on the streets. Thus, we saw a transition of the audience from being an implicit representation of a more or less homogeneous group of users to an explicit, actively participating audience and even what Axel Bruns defined as the "pro*duser*": a user turned producer.[144]

Many of these interactive possibilities that came with the interactive features of the Internet were not new in themselves, but the acceleration of the speed at which the audience can participate was. Both on screens with their clicks or absence of clicks and more directly in comment sections or e-mailing the journalists directly. These new possibilities were heralded as a saviour of democracy, even a democratisation of the news media. They would allow the news media and the journalists to really and deeply connect with the readers, almost as if it would provide a much and long needed couples' therapy conversation.

But, similarly to the reactions to the screens and despite the promises of how journalism could connect better with the audiences, I could also observe how users were simultaneously dismissed as "stupid", "crazy", or "mud-slinging" debaters. This seemed surprisingly similar to how journalists had previously treated those from the audiences who used to write letters to the editor by snail mail.[145] In a study by media scholar Karin Wahl Jørgensen, we saw how with these letters, e-mails, and debates, journalists were directly confronted by readers who did not always behave like they wished for.[146] Readers who did not appreciate these efforts of trying to improve the journalism, but used the opportunity to yell, complain, or get into uncivil debates with other readers. This direct contact between the user and the journalist complicated the relationship further and to a degree where couple's therapy was no longer enough, it seemed. Heading towards a possible divorce, the news media pondered whether they just did not have sufficient knowledge of what the user was like. Perhaps if they just knew them better, they could even predict what the users wanted. They could improve the relationship and win some of the users back from their endless scrolling on Big Tech's platforms.

And so, they embarked on a data collection journey, hoping that if they combined this data with algorithms, they would be able to predict what you might like, even before you know it yourself. And in this way salvage and repair the strained relationship.

The meeting room has a glass door,

a small box within the bigger box.
Inside, a small table with six chairs,
a large screen on the wall, and a white board in the corner.
Several times a week we enter the room together:
the two data scientists, the head of the personalisation project,
and the digital team.

And me.

Today, one of the data scientists explains
how they have started building a user database
of logged-in users.
He projects a table of rows and columns
onto the screen from his computer.

"We can see gender from their names,
limited to two genders of course,
and for location we only have the postcode
and not the address", he explains.

"There can be quite a lot of difference between user profiles
within a specific postcode", the digital editor says.

"Yes, I agree", the other data scientist says.

"Is it better to not use the addresses at all then?"
the editor asks.

"Not sure, we need to run some more tests",
the data scientist says.

On one of his visits to Ada and her family to their estate Worthy Manor, Babbage has news that makes both him and Ada excited. He reveals the news as they walk down what was known as Philosopher's Walk – where Lovelace and Babbage often walk while discussing mathematical principles. He has been invited to Turin to present his work on the Analytical Engine to a circle of Italian mathematicians. "I might go via Lyon", he says.

Ada immediately knows why he wants to make a stop there and wishes she could come with him. That evening, they continue planning the trip and, excited, they both agree that this trip might turn things around. It sometimes feels, he says to Ada, like society around us does not understand what we would be able to do with this machine.

"I want to see the silk manufacture with my own eyes, and in particular a loom engaged in a singular task: the weaving of a portrait of Jacquard, composed not by hand but by an ordered sequence of mechanical decisions", he says to Ada.[147]

If he can see how it is built and observe it, he can improve his own invention of the Analytical Engine. He does not need to explain: Ada feels it too and she is reminded of how she used to observe the automated looms with the mother. And in the back of her mind, she feels how she herself might hold the key to making that machine as magnificent as they are both convinced it could be.

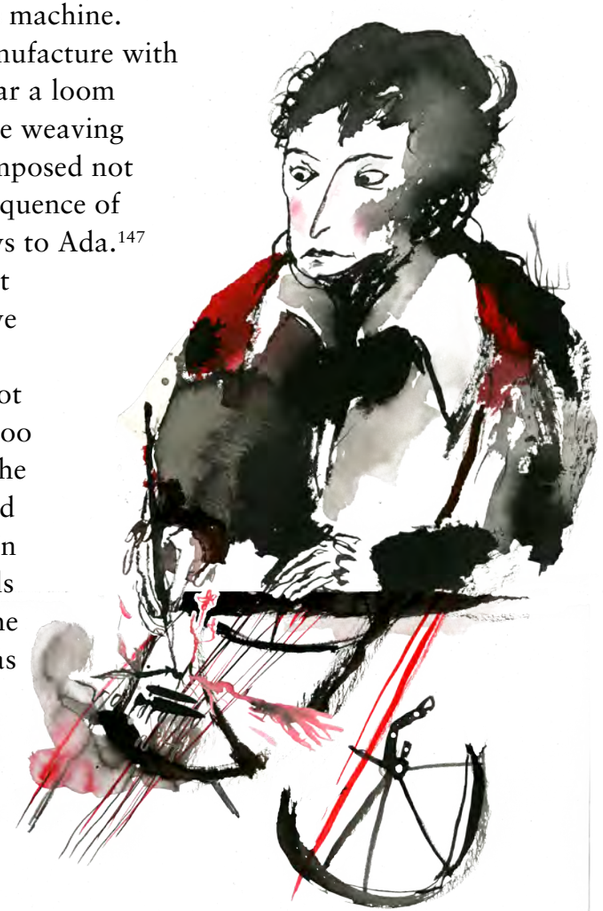

The public mission at stake

In the world of journalism, professionalism is often defined by a shared commitment – though interpreted with varying degrees of intensity depending on the newsroom – to keeping the public informed about matters essential to their role as citizens in a democracy. The varying degrees are important here. This commitment exists in tension with the directives of audience metrics, which offer an ever-present, if uneasy, influence on editorial decisions. As media scholar Mike Ananny observed, journalism is distinct from other forms of media production in its *public* orientation.[148] In my own work, I have explored this public mission and its guiding presence in how journalists relate to metrics and in building personalisation algorithms. I have argued how media actively cultivate publics, shaping them into collectives with a shared sense of civic responsibility.[149] The image of the audiences held by journalists guiding their work is closely connected to this public mission. In this ethos – this deeply held belief in journalism's role as a steward of democratic discourse – we witness journalists grappling with the abstract ideal of a Habermasian public sphere. To fulfil this mission, news organisations must understand their audiences, not simply as consumers but as the very publics they seek to sustain and cultivate. And yet, the fraught relationship with metrics also serves as a mechanism by which journalists assert their professional judgment – allowing them, when necessary, to disregard audience data in favour of what they believe to be the greater good. They believe that they know better than the news users what's good for them, because the stakes are higher than just the individual's need for immediate gratification through entertainment. They find ways of asserting their agency back into the production process. Ways of gaming the machines, in a way.

This dynamic reveals an ongoing negotiation: journalism as a public service versus journalism as a market-driven enterprise. This is a negotiation which goes on among journalists inside the field and between the field of journalism and other fields around it – like the political field. Within newsrooms, the belief persists that the *purest* expression of journalism is in writing stories that equip and empower citizens, reinforcing their role in a democratic society. The broader imperative to reach mass audiences – though financially necessary – occupies a lesser, more contested position in the hierarchy of journalistic values, as I mentioned earlier.

French sociologist Pierre Bourdieu famously described this tension, likening the journalistic field to that of art: pulled between an economic pole, where market forces dictate terms, and a journalistic pole, where the profession finds internal justification as a guardian of democracy, independent of readership figures and metrics.[150] The distinction between utilitarian, quick-turnaround journalism and more ambitious, investigative storytelling has only grown sharper in the digital age, where analytics have transformed content strategy and stratified the forms of news production. The platform itself – print or digital – has become as much a cultural signifier of intent as the story's substance. For many users, reading a newspaper feels more serious, while social media news consumption feels too superficial – even if it is likely to be the same content published in two different places. Furthermore, specific spaces for consuming news are connected to belonging to different classes, where classic news outlets or even just news reading as a practice is increasingly viewed as elitest and posh.[151]

The arrival of audience metrics in the early digital era was met with a mixture of pragmatism and resistance inside the newsrooms. News organisations, reluctant to compromise their core product, often relegated their online presence to an ancillary role – an advertisement, rather than an extension, of their primary mission. The result was an uneasy hybrid: news repackaged, recycled, and offered for free, even as print editions remained the perceived bastions of journalism's civic responsibility and public mission. In transferring their work online, journalists did not merely shift distribution channels: They inadvertently transplanted the very structures, routines, and professional ideologies that had long defined their work in print. In some ways, print could uphold the public mission, while digital platforms could attract users, fulfilling the commercial imperative of news organisations.

My previous research[152] has traced how the growing reliance on data to manage and predict user behaviour – the datafication of journalism – extends far beyond the actions of individual readers. Beyond individual journalists losing agency to metric systems. This datafication shapes collective audience constructions, redefining how the news media envision their public mission. The challenge, of course, is also that the metrics, often reduced to clicks or time spent on an article, offer an incomplete and sometimes misleading representation of how people engage with news in their daily lives.

What the metrics fail to capture is the deeper, more complex relationship between journalism and the public it seeks to serve – a relationship built not merely on consumption but on trust, participation, and the enduring, if fragile, belief that a well-informed society is worth striving for. More importantly, the reliance on the ability to model individual user preferences risks further dismantling of the function of the news media as an institution that cultivates publics by enabling users to connect and to relate and commit to at least some minimal form of democratic engagement as citizens.

The challenge facing the news media and the news organisations I was observing was to use those same metrics to build and maintain the forever adapting personalisation algorithm. But how people connect, relate, and commit when they consume and engage with the news turned out to be much harder to measure in engagement metrics – and optimise through machine learning – than pure attention. And it requires more and more data.

The project meeting this morning is somehow different.

Different than the previous meetings I've attended.
I can sense that the project team members
seem agitated in their small talk
before the project meeting starts.

That morning, the Danish weather app DMI
has received harsh criticism from the National Data Council
for using data from users to target advertising on the app,
which users had not said "yes" to.

The verdict is one of the first following the implementation of the data protection law General Data Protection Regulation in the EU.

The fundamental question it raises is whether an organisation can use data collected for one purpose for another purpose.

And this verdict makes it clear: You cannot.

Hence the upset.

The team had hoped to use the data on users collected
on a second-hand buying-and-selling platform,
which they also own.

The idea was to use the addresses of people
for the personalisation algorithm
they are building for the news site.

"It would work so much better
and enable us to deliver content relevant on street-level",
one of the developers proclaims without hiding his excitement.

"But this verdict means that we cannot",
the guy from the legal team declares.

The rest of the meeting they all appear defeated.
As if society does not understand
how amazing this machine could be.

If they just had the right data to test what it could do
and enough money to build it.

From Lyon, amid the din and rhythm of silk production, Charles Babbage composes a letter to Ada, still buoyed by what he had witnessed. The city is pressing in on all sides – noise, motion, repetition – but his attention remains fixed on a single machine. The reason for him to be here.

The Jacquard loom has held him for hours, its movements precise, almost ceremonial, as punched cards rise and fall, issuing commands without hesitation. He watches the cards circulate, each one carrying instruction rather than material, directing threads into patterns that remember themselves.

The loom appears to him as a form of mechanical memory – its logic distributed across cardboard and motion. Rows of punched cards flutter like the wings of clockwork moths, commanding threads with perfect precision. He imagines how this logic would be transferred to the Analytical Engine, the wings of the moths carrying numbers from one hole in the punch card to another.

Writing to Ada, Charles allows his excitement to surface. The loom suggests a transposition waiting to be made. As he continues on toward Turin, he is carrying with him the image of a machine that remembers and would produce the woven images of numbers without any hands or minds of a human.

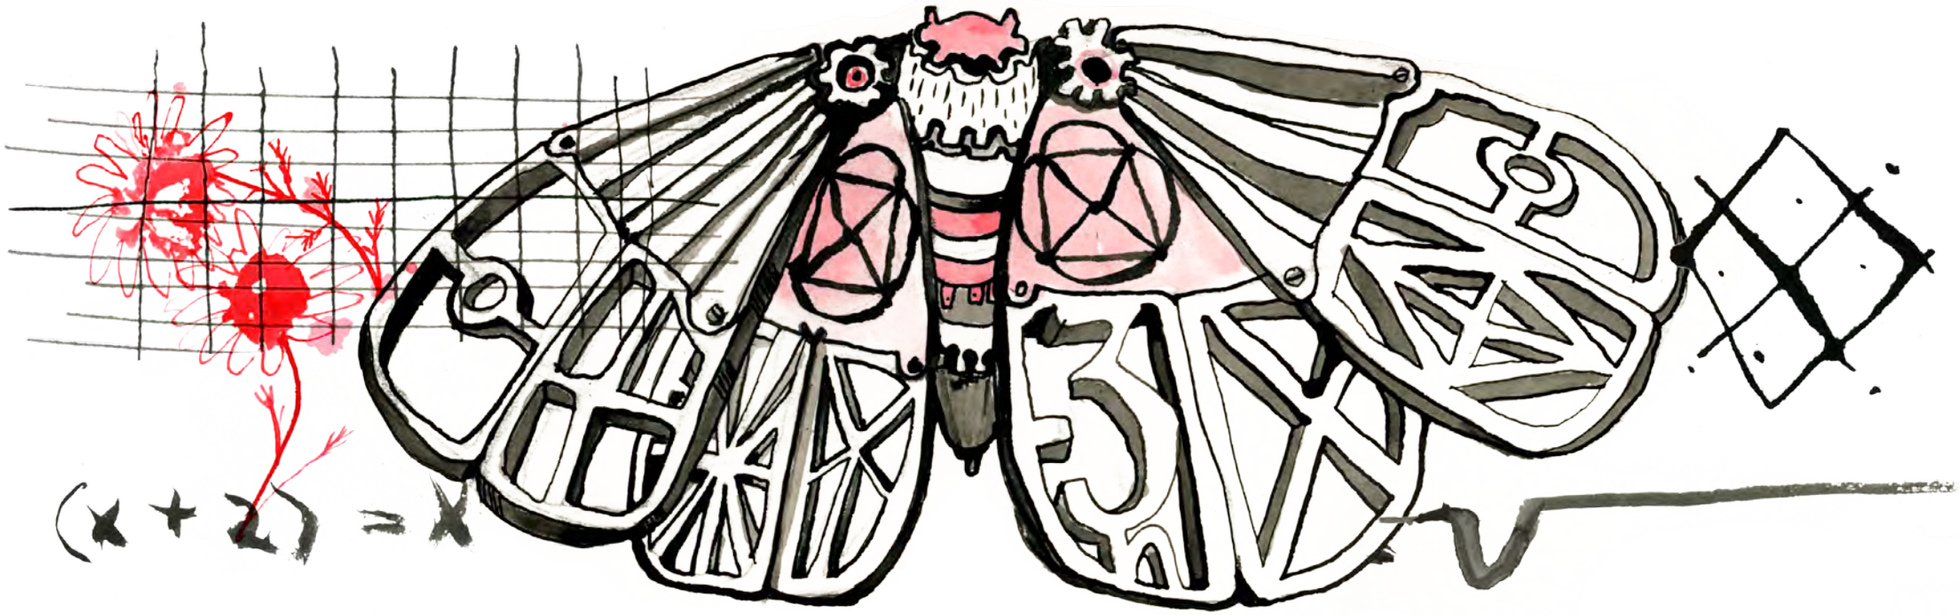

The incomplete breadcrumbs in the woods

Let us explore what I mean when I say that the experience of news is difficult to capture in algorithmic systems. We can start by establishing that usage and practice, though often conflated, remain distinct entities – two sides of the same coin, perhaps, but minted in different realms. Usage is tangible, measurable, and shaped by the constraints of technology. Our practices as we engage with media content, by contrast, are amorphous, embedded in culture, habit, and identity, inflected by the subtle rhythms of everyday life. They're situational. A click on a news item, for instance, is a discrete digital act – a functional equivalent of turning a newspaper page – but its meaning is entirely contingent on the broader practice in which it is embedded. For one user, it is a way to pass time, an idle gesture of engagement; for another, it is an act of information-seeking, a deliberate step in constructing an understanding of the world. The click itself is empty of meaning – it is the practice that lends it significance.

As media scholar Susana Paasonen argued, "the meanings of cultural products and social phenomena are never fixed: these accommodate multiple, possibly mutually contradicting interpretations, uses, and practices, and at the very least, feelings toward them can be highly mixed". She described culture as "the stuff of ambiguity, yet ambiguities easily disappear in analytical work aiming at firm outcomes and strong compelling arguments". This is the case with the discussion, usually negative, surrounding social media's impact on personal and collective lives.[153] And it's the case with news, which are often measured as if it was something we have use of like we have use of a washing machine or a car. But news in this sense is similar to other cultural products like music or literature, evoking memories and emotions, connecting us to the world around us.

The news enables us to *connect*. To ourselves and to each other.

News engagement as messy

We know from research that our relationship with media is a complex one, and the relationship with news media is part of that complex relationship. In *Seeking the News Audience* from 2011, the media scholar Elizabeth Bird dismantled the tidy illusion that news consumption is a simple act of reading or watching: It's a far messier, more social, and deeply cultural process.[154] News is so much more than information: It weaves itself into the rhythms of daily life, shaping conversations, identities, and even our sense of belonging. People don't passively absorb information – they filter, reinterpret, and repurpose it, often in ways that transcend the original story. Gossip, memory, and emotion blur the boundaries between journalism and everyday life, turning news into something lived. Bird argued that understanding how people engage with news means looking beyond headlines and broadcast schedules into the social networks, habits, and rituals that give news its real power. The impact, she suggested, isn't just about what people know, but how they navigate the world. In our news experiences, we are already agentic and much less passive, even when we just lounge on the sofa watching a film. We use those cultural products, whether films or news, to reflect upon who we are and to navigate the world emotionally and cognitively.

In short, they enable us to *relate*. And they shape *how* we relate to the world surrounding us.

As I read through the media diaries that people had written over the two weeks in my previous research project, it struck me how much of what they referred to include a sense of duty to the world around them.[155] Watching or reading the news felt like an obligation at times, but not a tedious task like having to wash the dishes, but more like an enjoyable task of being citizens. They would feel a bit embarrassed if they did not have the time to keep up for a few days.

Take two women from the study – 33-year-old Karen and 45-year-old Susan – who both expressed low political interest and said in their interviews that they did not feel they succeeded in staying informed about society. Nevertheless, they both wrote about the 2016 American presidential election, which was one of the major news agendas during the diary period:

> I guess the thing that has caught my attention the last week has probably been the election in the US. I didn't feel like staying up all night to watch it. To be honest, I don't care that much. But I will gladly thank all the other people who could then update me when I woke up and – as the first thing – grabbed my phone and checked various news and posts from people. Because you can count on one thing: If something special is happening, you can't avoid keeping track of it, even being all unconscious and without desire and interest. (Karen)

> The media is full of the American election. I haven't been following it so much! Except from what's in the news. Everywhere, actually. You really have to try, if you want to avoid it, but I'm looking forward to seeing who wins that election and what happens afterwards. It's a little frightening. (Susan)

These excerpts demonstrate a similar – and paradoxical – tension between disinterest and interest. Both Karen and Susan concluded that they couldn't avoid news about the election despite their disinterest, and they seemed glad about this, as it kept them updated about the overall status and development. This can be related to the idea of a minimal orientation. We can thus identify a narrative of a moral obligation to stay informed (via the news), resembling the traditional ideal of the "informed citizen", or what has been called "dutiful citizenship",[156] which also corresponds with earlier studies on public connection.[157]

News enables us to *commit* to a world larger than our own. Even when we don't see ourselves as having an interest in that larger world.

The incomplete traces we leave behind

Crucially, our media practices are not bound to any one mode of use. The act of engaging with news, for instance, may unfold through a newspaper, a television broadcast, a social media feed, or a conversation over morning coffee. It is not tethered to any single platform, let alone an algorithmically governed space like Facebook. And yet, algorithms care little for practice nor paradoxes in those practices, or at least they cater to mostly straightforward binary click/no-click practices. They are trained to recognise usage – to isolate patterns, categorise behaviours, and assign value to what is, in essence, a decontextualised token of engagement. Once a click is flagged as salient, the machine learns to serve up more of the same, an endless loop of reinforcement indifferent to the user's intent. What emerges is a kind of algorithmic determinism: the illusion of personal agency (this is adapted just for you), subtly undermined by an invisible hand (the algorithmic system) guiding what is seen and unseen.

The sociologist Eran Fisher, reflecting on this shift, notes that in the age of mass media, audiences were "seen" through a blend of social theory and empirical research, a project of systematic classification. Digital media, supercharged by data and predictive modelling, has transformed this epistemic pursuit into something altogether more granular and relentless.[158] The promise – knowing the audience with unprecedented precision – is tempered by the unsettling realisation that what is being known may be less the individual and more the traces they leave behind, endlessly parsed and repackaged by systems that, in the end, may not truly understand them at all. What they click is known, but not their practices, and thus not *why* users click or the messy meaning and emotions it evokes in them. Paradoxically, news media hope to rebuild the relationship with readers by giving them agency to choose for themselves and adapt to those needs and wants, which in turn transfers that agency to algorithmic systems.

Eran Fisher warned about this development and emphasis on systems which might not know our subjective selves, and he went as far as to say that because algorithms are taking over "the thinking", we lose our subjectivity as humans. "As a result, algorithms increasingly sidestep the role of subjectivity in the formation of that knowledge, undercutting the humanist project of a self, able to form and articulate what it thinks/likes/wants", he wrote.[159] The argument is not far from Kurtzweil's vision of The Singularity, but the dystopian version of it.

In this argument, Fisher articulates the idea (and fear) of the machine taking over the human subjective mind, arguing that the human self can only be truly articulated *outside* the realm of machines. As if taste is something that can only develop between humans and not between human and machine, for example. However, this argument rests in the belief that human subjectivity is exactly and only something residing *outside* the thinking machines, and when attempting to transfer it to the machines from humans, the subjectivity is somehow lost in translation.

What would this argument look like if the transformations are grounded in the *illusion* of the thinking machines, the belief in the magic of measuring and classifying humans? A belief which gives the machines symbolic power to "make people see and believe"?[160]

What if the subjective practices evolve as entanglement of the human–machine relationship is based on this illusion that the machines can know us while at the same time telling us who we are as humans? What kind of citizens can we expect the machines to cultivate? What kind of public mission do they embed? And what happens to users' sense of agency, when we are guided around the algorithmic maze?

I observe as they start testing the model

on the training dataset they have built
and the data they have on the users.

It looks promising, but they worry about overfitting.
They run cross-validation to check whether it generalises well.
The accuracy is high, but they need to dig deeper.

They analyse precision, recall, and F1-score.
The model is making too many false positives.

They tweak the threshold and retrain it.
Now, precision improves but recall drops.
They need a better balance.

They run robustness tests.
They introduce noisy data to see how the model reacts.
It struggles with edge cases.

They adjust preprocessing steps to handle variations better.
They check if the model is biased toward any group.
The results show a small skew.
They trace it back to imbalanced training data and rebalance it.

They retrain and test again.

They speak the language of formulars and calculations.

On his return to England, Babbage cannot not wait to tell Ada what he has observed. The hours he spent there watching the movements of the loom in Lyon fold themselves quietly into his broader thinking. The loom demonstrated how instructions could be externalised, how design might persist independently of the hand that initiated it. Threads followed a logic set elsewhere, the pattern advancing without improvisation.

And in the correspondence with Ada, the thinking evolves. She understands him. But she also has ideas of her own that go far beyond the levers and bolts of Charles's Analytical Engine.

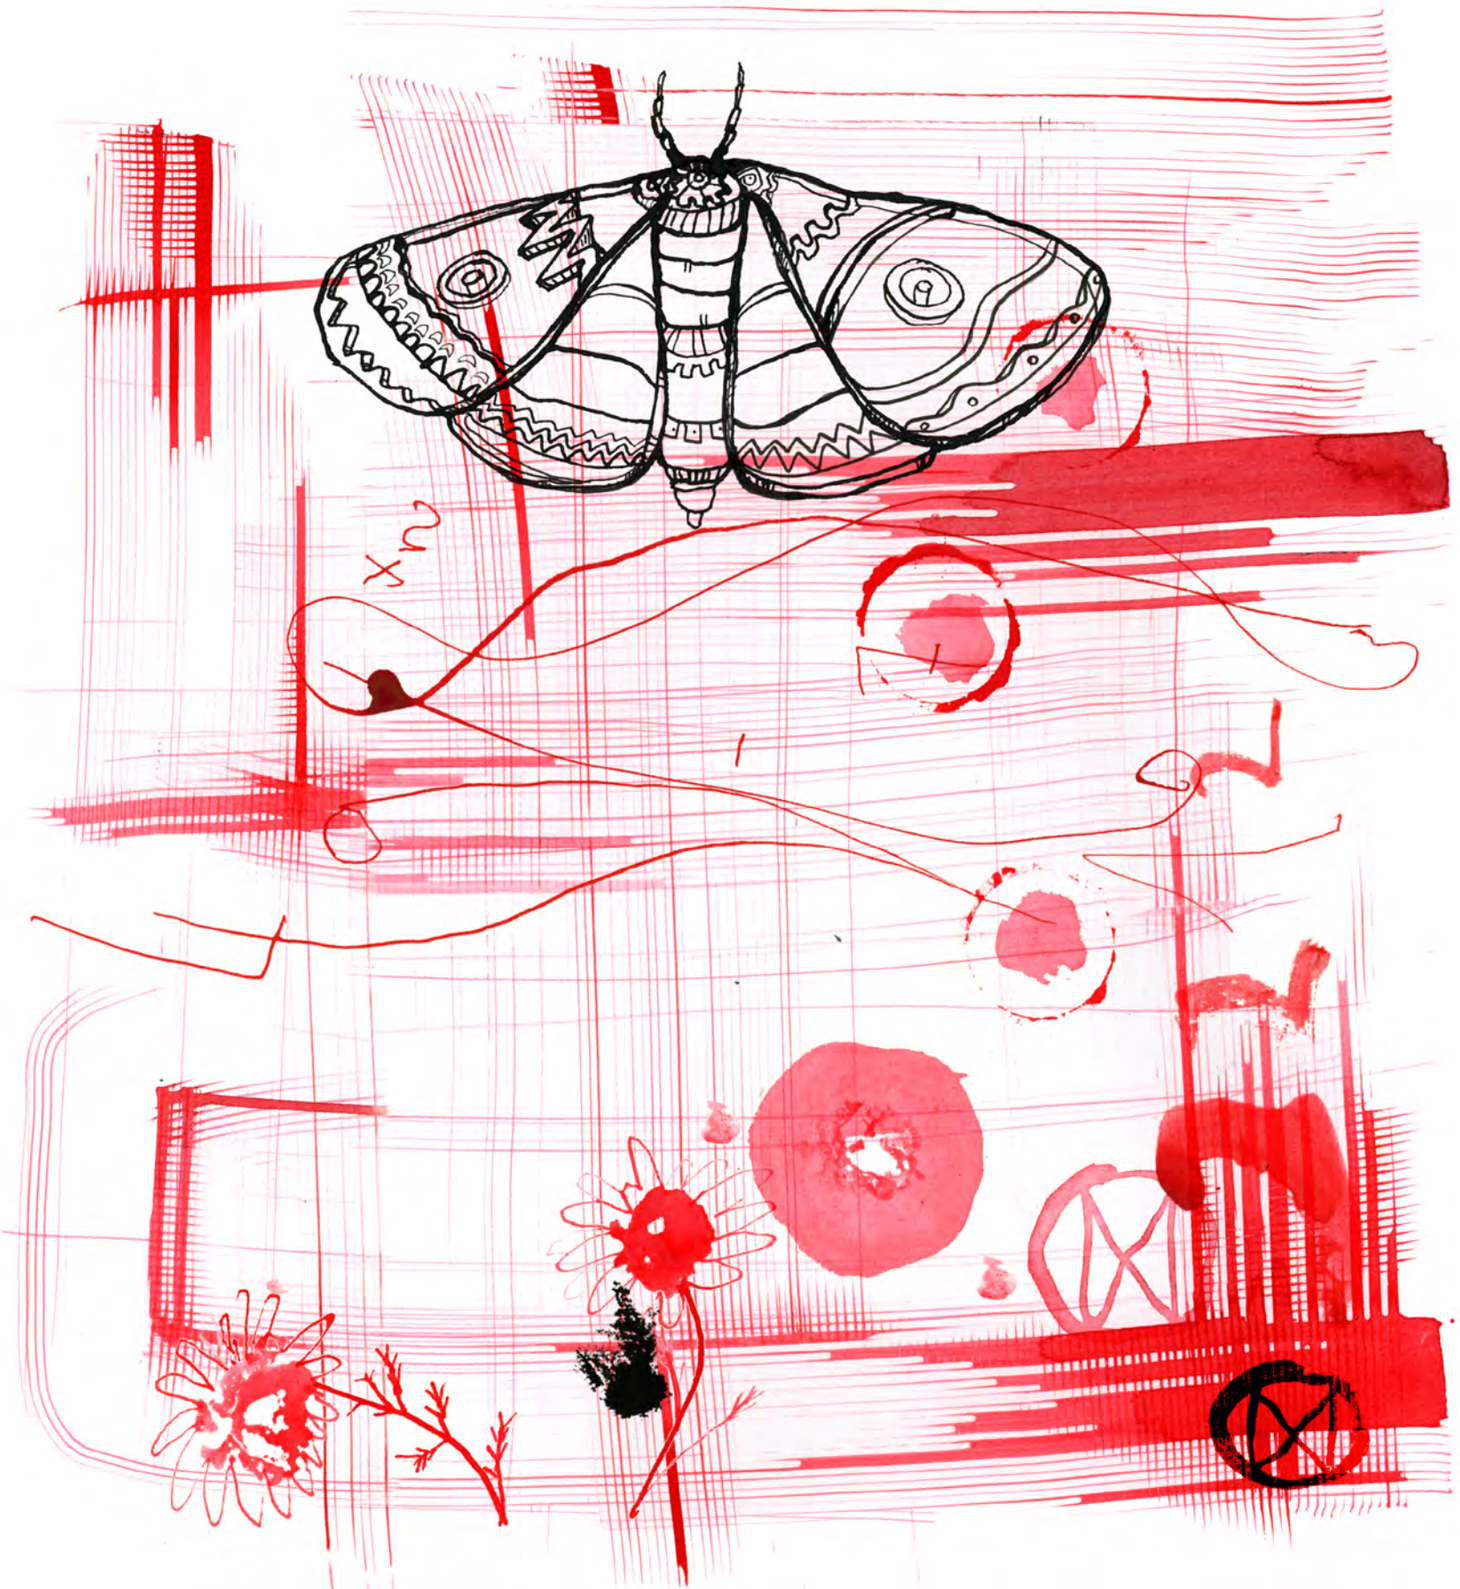

Elaborate illusions of human-like intelligence

Let's examine the belief that machines do a better job than journalists and editors in knowing users and judging what news to produce and distribute. To do so, it might be worth taking a closer look at the history of automating machines. In 1770, an automaton chess player, or The Mechanical Turk, was introduced. It seemed to successfully play, and win, at chess against human opponents. It was brought around Europe on tours, and when the original creator died, it was bought and continued to be exhibited until it was destroyed in a fire. Afterwards, it was proven to be a hoax, with a human chess player hidden in the cabinet below.[161]

The human fascination and belief in hoaxes and tricksters are perhaps less surprising than the more than 84 years it took for people to discover that there had been a man inside the machine playing chess every time the machine was "working". The faith in the ability of machines to really think and know us has a longer history of fascination with machines and AI, and their ability to think and act like humans. An argument central for the idea of The Singularity promoted by Kurtzweil. Yet, the belief in the machines and their ability to know and calculate our every need is extremely powerful – ultimately because they also give the power of production to those owning the machines. Think of OpenAI's Sam Altman promoting AGI – artificial general intelligence – as the mode of AI that will outperform humans in most economically valuable work.

As the Mechanical Turk shows, the history of deception by machines is indeed a long one. In his book *Deceitful Media*, Simone Natale asks us to reconsider the story we like to tell ourselves about technology's rise – a story of progress, inevitability, and unprecedented empowerment. Instead, Natale drew attention to the slipperiness of this narrative, revealing how deceit and artifice have been woven into the DNA of media technologies from the start. For Natale, deception is a cultural mechanism that drives both our anxieties about technology and our fascination with it.[162] In similar ways, the hallucinations of large language models should not be seen as bugs of the system: They are inbuilt and just come out of the random mix

of content and datapoints, where data sources might be wrong. We are so fascinated by its ability to mimic that when its talks gibberish we are inclined to think that it's an error, forgetting that it is an inherent part of the model to fabricate things.

Consider the chatbot. Decades before ChatGPT was a cocktail party topic, the earliest iterations of conversational machines – Joseph Weizenbaum's in the 1960s, for instance – relied on scripted mimicry rather than actual understanding. One of the first chatterbots (chatbot), ELIZA, was an early test case for the Turing Test. ELIZA fails very quickly by today's standards, but for many people at the time, it appeared able to have meaningful conversations. The original algorithm going by the name of DOCTOR was apparently modelled on a therapist that Weizenbaum was seeing at the time. And interestingly, he wrote the model on punch cards that were not digitised until much later. Users began to confide in ELIZA. They attributed to it empathy, understanding, even care. Some insisted on interacting with the program in private, as though it were capable of genuine therapeutic engagement. For Weizenbaum, this response was deeply unsettling. It revealed not the sophistication of the machine, but the vulnerability of the human. People were willing – eager, even – to project understanding onto a system that possessed none.[163]

This experience became the seed of his later critique in his book *Computer Power and Human Reason*.[164] At a time when computers were still rare, room-sized machines humming in institutional basements, Joseph Weizenbaum published a book that reads today less like a technical intervention and more like a moral dispatch from the future.

The danger, he realised, was not that machines would become indistinguishable from humans in their intelligence, but that humans would become indifferent to the distinction. When simulation suffices, substance becomes optional.

From this insight emerges one of Weizenbaum's most enduring conceptual distinctions: the difference between *deciding* and *choosing*. Deciding, in his account, refers to the selection of an option

based on calculable criteria. It is a process that can, in principle, be formalised, optimised, and delegated to machines. Choosing, by contrast, is a moral act. It involves responsibility, context, and the irreducible weight of human values. It cannot be fully specified in advance, nor can it be outsourced.

Weizenbaum famously concluded, based on his work, that ELIZA could create an "illusion" of understanding despite relying on simple pattern-matching techniques. It's not so different from the way audiences once marvelled at the "thinking" automata of the eighteenth century – mechanical contraptions that could write poems or play chess but were, in truth, elaborate illusions.

Deceptive wizards

Media scholar Simone Natale argues that these systems were not so much technological marvels as theatrical performances, dependent on the audience's willingness to suspend disbelief. Users would project intelligence and intent onto the machine, despite knowing its responses were mere algorithmic tricks. These methods of hiding how a model is not working as well as hoped for is called the Wizard of Oz method.

The genius of modern AI is that it carries this sleight of hand into the digital age. AI systems, Natale pointed out, have always relied on a measure of deceit – not necessarily with a malicious intent, but in the way their inner workings are obscured while presenting a façade of human-like intelligence. The recent proliferation of chatbots, recommendation engines, and content generators reinforces this tradition. They don't actually "think" or "know"; they simulate these qualities convincingly enough to compel our belief.[165]

But Natale's argument goes beyond technology itself. The deception isn't only in the machines but in the way we talk about them. AI is often marketed as neutral, rational, and autonomous – a saviour of efficiency and objectivity. This is reflecting what Campolo and Crawford termed "enchanted determinism": the tendency to

naturalise the authority of AI systems through a quasi-magical belief in their objectivity, inevitability, or capacity to produce truth beyond human intervention.[166]

Interestingly, we see this in the language used by news media when they talk about personalised news recommenders. Often and increasingly, they are referred to as more precise and more objective than the gut-feeling of journalists and editors. Yet Natale has reminded us that these technologies are deeply enmeshed in the values, biases, and agendas of the humans who create and deploy them. There's a man (or a woman) inside of them, but what tricks us is that they seem like they're operating on their own. And we forget about the annotators. The data scientists. The digital editors and all the other people who are inside the machines making decisions based on complex mathematical models, like the ones Ada Lovelace imagined. Machines that have turned out to now have power over the whole chain of knowledge production.

Charismatic machines

AI systems – whether a model analysing users' data or a chatbot we might ask for advice on our everyday life – can feel strangely persuasive, not because they *know* anything in the human sense, but because they *seem* to. They speak with the rhythm of authority, mirroring our expectations of competence, confidence, and quick answers. They cannot really think or relate to you as another human being, but they are made so they seemingly know you. And what they provide in terms of text resembles human language so much that we confuse the machines for humans. We anthropomorphise them – we think of the machines as having human traits. But their charming answers are all about performance. They give off the appearance of knowing, and in a world shaped by fast information and frictionless interaction, that often feels like enough.

This performance draws its power from what sociologist Pierre Bourdieu called symbolic power: the ability to shape what counts as valid, to define the rules of the game without needing force

or permission.[167] It works through processes of recognition and misrecognition, now happening through and in large datasets. And it defines who gets to speak, who is believed, and whose voices are filtered out.[168] So when we talk about personalisation machines as objective and as analytically more powerful than humans, we also attribute them symbolic power.

In this light, generative AI systems become what we might term charismatic machines. They aren't alive, but they're designed to behave *as if* they are – responding with tone, empathy, and a flair for conversation. This is what some scholars call artificial sociality – machines built to mimic human social behaviour, much like the old Turing Test once imagined: If it feels human enough, we treat it as such. We even tend to treat them as *more than human*.

Their credibility comes from how we respond to them. We often take their fluency – the smooth, confident language they produce – as a sign of expertise. It's hard to question something we cannot see through, or which is black boxed from us. The machines and what they deliver have computational authority: the trust we place in machines simply because they're machines, backed by extensive code, platforms, and data. Trained on our texts and shaped by our interactions with the mouse as we click through websites, they learn and repeat the cognitive and cultural patterns we feed them. In doing so, they help reinforce what Bourdieu called legitimate language: the kind of speech that society accepts as knowledgeable, credible, and true – even when it's not questioned.[169] In essence, they become machines of knowledge, targeting and adapting to us as individuals through processes of personalisation. In the end, AI doesn't just generate words for us – it performs power. It helps decide what sounds like knowledge and what gets silenced. Power to provide us with a specific image of the world, but with the quiet force of symbolic power – shaping what's heard and what's not.

Just as early media technologies like the telegraph and the photograph carried implicit claims of authenticity that could be manipulated, AI often functions as an instrument of persuasion cloaked in the guise of impartiality. It seems so clever, can it really be wrong?

As spring arrives, it's time for A/B testing of the algorithm in the real world.

The model is exported and used on real user data.
They split the world in two and monitor user behaviour.

Success.

The new model increases engagement slightly,
but it still needs fine-tuning.

The image they are weaving
is not yet perfect,
but at least it's not worse.

Its good enough, they say.

For now.

I note down that
what is left unsaid is
what sort of balance
they should strike
between the "good enough" result
and the "slight increase" in user engagement.[170]

Mr. Babbage's presentation in Turin becomes the focus of an article written by the military engineer Luigi Menabrea (later Prime Minister of Italy) about the Analytical Engine and Babbage's work. It is published in an obscure Swiss journal, and it might have stayed hidden had it not been for an idea which pops up in Ada's head:

"Why don't I translate it?" she asks her friend. He agrees to the idea but wonders why she does not just write her own article.

"I had not even considered that", she replies.

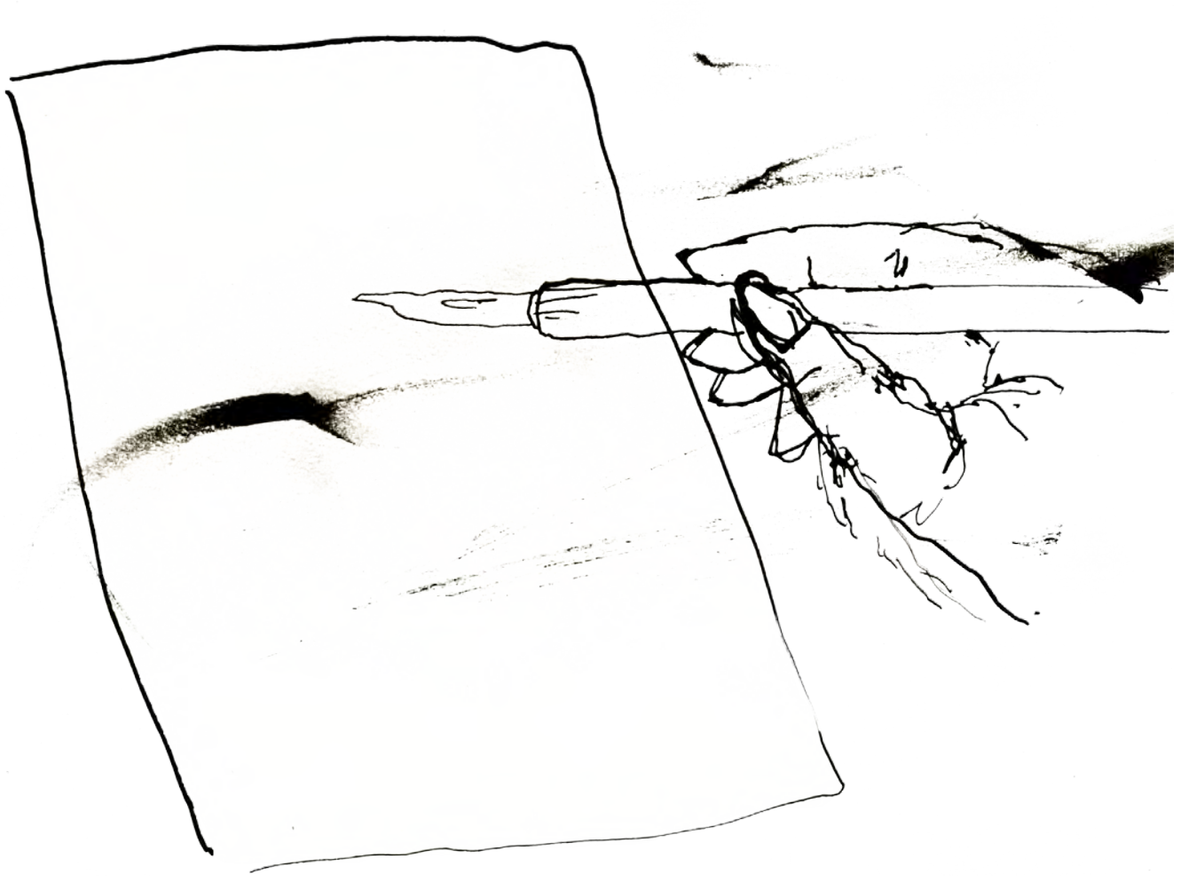

You and your dataself

It occurred to me when listening to how editors and media managers in charge of large-scale personalisation projects talked about the algorithms how much they talked about how it would give more power to the users and help news media re-establish a long relationship with the media users. Help them win back users who had been having intense affairs with social media and streaming platforms.

In the discussions they were having, it was all about giving the user more autonomy and choice. The inherent assumption when building personalisation algorithms is that they can give full autonomy to you as a user over your news media choices by guessing and predicting from past user behaviour and evaluation of machine learning models in self-referring systems. However, this predicting based on statistical models comes out of very specific organisational structures and power relations, and thus also represents a form of control over the individual, trying to tell you who you are and what you should be watching, buying, or thinking politically. Think of how you might be offered nappies because of your child-bearing age and your gender as a woman, even if you are not planning to have kids. Predictions based on previous usage, age, or other data available to advertisers and content producers – and less on the messy everyday practices, feelings, or emotions. It's an epistemic power over the production of knowledge, which is also – unless media companies have the means to build the model themselves – in the hands of the few Big Tech companies who build, own, and have control over the models used for personalisation of news and information.

These models return a mirror image of our subjective taste, at times re-enforcing our taste, at times challenging us to resist, yet still tangling with our identity based on a conception – a datafied image – of who we are in the data. And therein lies the symbolic power of the AI-powered personalisation machines. In *Data Selves*, Deborah Lupton dissected the way individuals morph into "dataselves" – fragmented, algorithmically defined versions of themselves shaped by the relentless quantification of daily life.[171] She argued that personal data is beyond the mere collection: It actively constructs

identity, blurring the line between self-perception and machine-driven categorisation. Fitness trackers, social media, credit scores, and biometric surveillance record behaviour – but they also generate new forms of selfhood, often without our consent. The result is a self that is both hyper visible and oddly elusive, shaped by unseen corporate and governmental forces. Lupton warns that while data promises self-knowledge, it also dictates who we can be, feeding us back a version of ourselves that may not align with our own sense of identity.[172] Importantly, when it comes to news, this identity is linked to our democratic citizenship. We are not just consumers, and the news is not a product like a pair of shoes – our engagement with news makes us capable to participate in democracy. So, when our news consumption is tracked, analysed, and fed back to us, it's also a mirror of who we are as democratic citizens. And it predicts what sort of democratic citizens we might be in the future.

Yet, in this digital hall of mirrors, the dataself is never entirely ours (although it *feels* very much like ours) – it is co-authored by algorithms, packaged for profit, and increasingly inescapable.

Imagined in systems

In the period of mass media, the individual was considered to have a generally stable identity and to be part of a social group. The metrics were helpful for imagining such segments, as the journalist could imagine specific groups when writing their stories. They would think of this person when writing the story, choosing the language and the headline to attract the readers. Even when they knew that every reader or viewer of the story would not be the same, or that they might respond emotionally differently to the story than the journalist had intended, it somehow routinised the production of many different stories so the everyday in the newsroom would be manageable. With personalisation, it is – on the distributional level – possible to target each individual user of the news site with content adapted to them: the audience of one. In the era of personalised media, the idea of the individual as a coherent, social, and fixed segment is broken down. Individuals are defined by their own ever-

changing behaviour and are grouped with others who show similar data patterns.[173] The journalist might still have to imagine a group or an individual when they write their story, but this is irrelevant somehow, because the algorithmic system does not imagine a reader – it is constructed from the available datapoints. The user is somehow still imagined, but not by journalists and editors. Users are imagined in and by algorithmic systems.

The machine learning models enabled by the recent technological developments means that time between production of a given article and the distribution to users is compressed. The content can be produced and delivered when there's a need, and the content can even be adapted to fit the time the user has available (although the model can neither really know nor predict this). This represents a fundamental shift away from the previous model, where metrics merely gave a delayed hint about what audiences liked and disliked. With the algorithmically produced mosaic of digital information, the content is – in theory – matched to the individual 1:1, and this could even be in the topic, in the tone of the content, or the language could be switched from user to user. In this algorithmic system, the media messages are more like particles, forever adapting to user needs in real time.

Predicting as a form of control

This means that the news media are increasingly selling news as a product to the audiences directly and not just through advertising to specific segments. Both the content and the user being forever adaptive to the ever-changing needs of the individual, an ever-changing free agent, forever free to listen to new music or to move to a new news site or streaming service. Just as classifying the input dataset for the training of the machine learning model presents a view of what news is, topics, and so on, the user data is also used to train the model to better recognise us as media users. A training which never stops, as new data is continually fed into the system. When designing the personalisation algorithm, the imagined user

is an aggregated datapoint, an average statistical norm of what a given human looks like, based on the data available, calculated minute by minute.

This development also represents an attempt to govern us as individuals. To govern us to be better news consumers, who step through the paywall and click "subscribe". According to Robert Kowalski, algorithms are defined as "algorithms = logics + control", with "logic" being used to signify the knowledge needed to solve the specific problem and "control" signifying the strategies that govern the problem-solving.[174] In the context of personalised algorithmic curation with recommender systems, the logic applies to the ways the systems produce meaning for the user, which are represented by the concrete predictions guiding the content selection for individual readers. This in turn controls that behaviour through a strategy of classification, sorting and filtering the user, for example, to convert a single-time user into a subscriber. This meaning-making is achieved through various algorithmic systems that use data signals (user behaviour data, news article clicks, article similarity) to predict what future content should be targeted to a specific user at any given moment.[175] Let's try and dissect what this really means.

Control by classification

The American philosopher John Cheney-Lippold has argued how "the process of classification itself is a demarcation of power, an organization of knowledge and life that frames the conditions of possibilities of those who are classified".[176] In his book *We are Data*, he proclaims:

> We are not simply well filled of data but made of data that is interpreted, conferred truth, and disseminated for motives of profit, organization, and/or control. The resulting classifications become the discursive terrain from which we, and others, compose our digital selves.[177]

The charismatic machines of AI enable this control by classification across many different spheres of information seeking.

These ways of aggregating the human through data is not only a question of what news we might see and not see – the news exposure, in other words – but a question of what it means for our identity to be categorised and viewed through data and modelled in machine learning systems.

The pursuit of knowing the audience has never been more sophisticated – nor more deceptive, we might say. With AI-driven personalisation, news organisations no longer rely on abstract readership trends but on detailed behavioural data, tracking every click, scroll, and second spent on an article. The models and systems claim to "know" the users with data, but what is known is "the statistical average and aggregated human" on which the personalisation algorithms are basing the targeting, which only captures a fraction of the media experiences that are crucial for our engagement in democratic public life. Some of these non-captured elements of the news experience are emotional, cognitive, and imagining. Audiences are not merely the sum of their interactions with content – they engage with news in complex, contextual, and often contradictory ways that no algorithm can fully capture. Unless of course we accept Kurtzweil's premise: that the machine will somehow merge with the human mind.

Personalised algorithmic systems, no matter how advanced, reduce audiences to patterns of past behaviour, assuming that what a person has clicked on before is the best predictor of what they will click on next. This logic draws attention to a broader technological myth – the belief that AI, if trained on enough data and built with sophisticated enough models, can achieve something close to human understanding. But this ambition follows a familiar trajectory of misplaced faith in machines. Or a collective deception by the magic promises of AI, a promise that a user, in collaboration with the responsive machine, can curate the news better than any human editor.

Next, we need to examine what the algorithms do with this elaborate yet somewhat deceptive image of you and your needs – but what we uncover next unravels more than just a path of endlessly more convenient choices…

The digital editor sends a link in the chat function in Teams,

which transfers me to a Google Sheet
in which he has organised the latest test results from the algorithm
in a way that illustrates overlap
between previous reading behaviour
and the algorithmic recommendations.

He is quite happy with the look of the results
(seems the users have moved along the path he hoped they would).

He says: "You could see that the articles being recommended are actually articles which are in the same category as the 'most read'. So, it is illustrating that although we do not have a manual filter, which can filtrate on geographical location, we will still hit very locally with our content".

He stares at the numbers for a while. Then he continues:

"We might need the manual filter to ensure that there is also a fair amount of local content and not just, for example, sports news".

"How much should be local?" I ask.

"Perhaps 50 per cent", he says.

I ask myself where that number comes from.
What is the "right" amount of local content versus sports
and other things on offer?
And I wonder if he knows
what the current percentages look like…

Charles is in trouble with his investors in the Analytical Engine, and in the meantime, Ada works hard to translate the presentation he has given in Turin from French to English. She imagines the translation will guide the development of Charles's machine, but she is impatient.[178]

She does not find the talk elaborate enough, and she starts making her own notes to supplement the ideas in the presentation, a form of "supplementary notes from the translator".

She imagines that the machine will not only process numbers, but also text and perhaps even music. A language of the machine, guiding its movements. Essentially processing everything, which can be made into data.

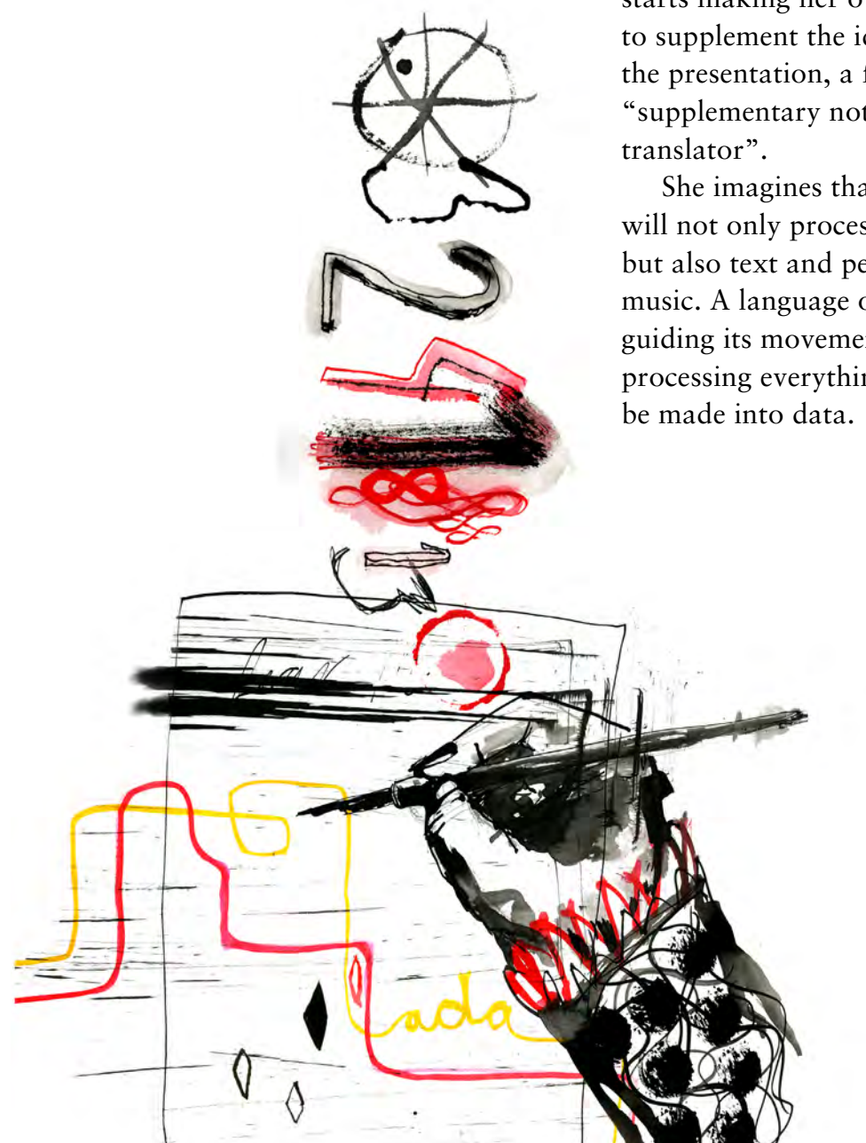

Part 5

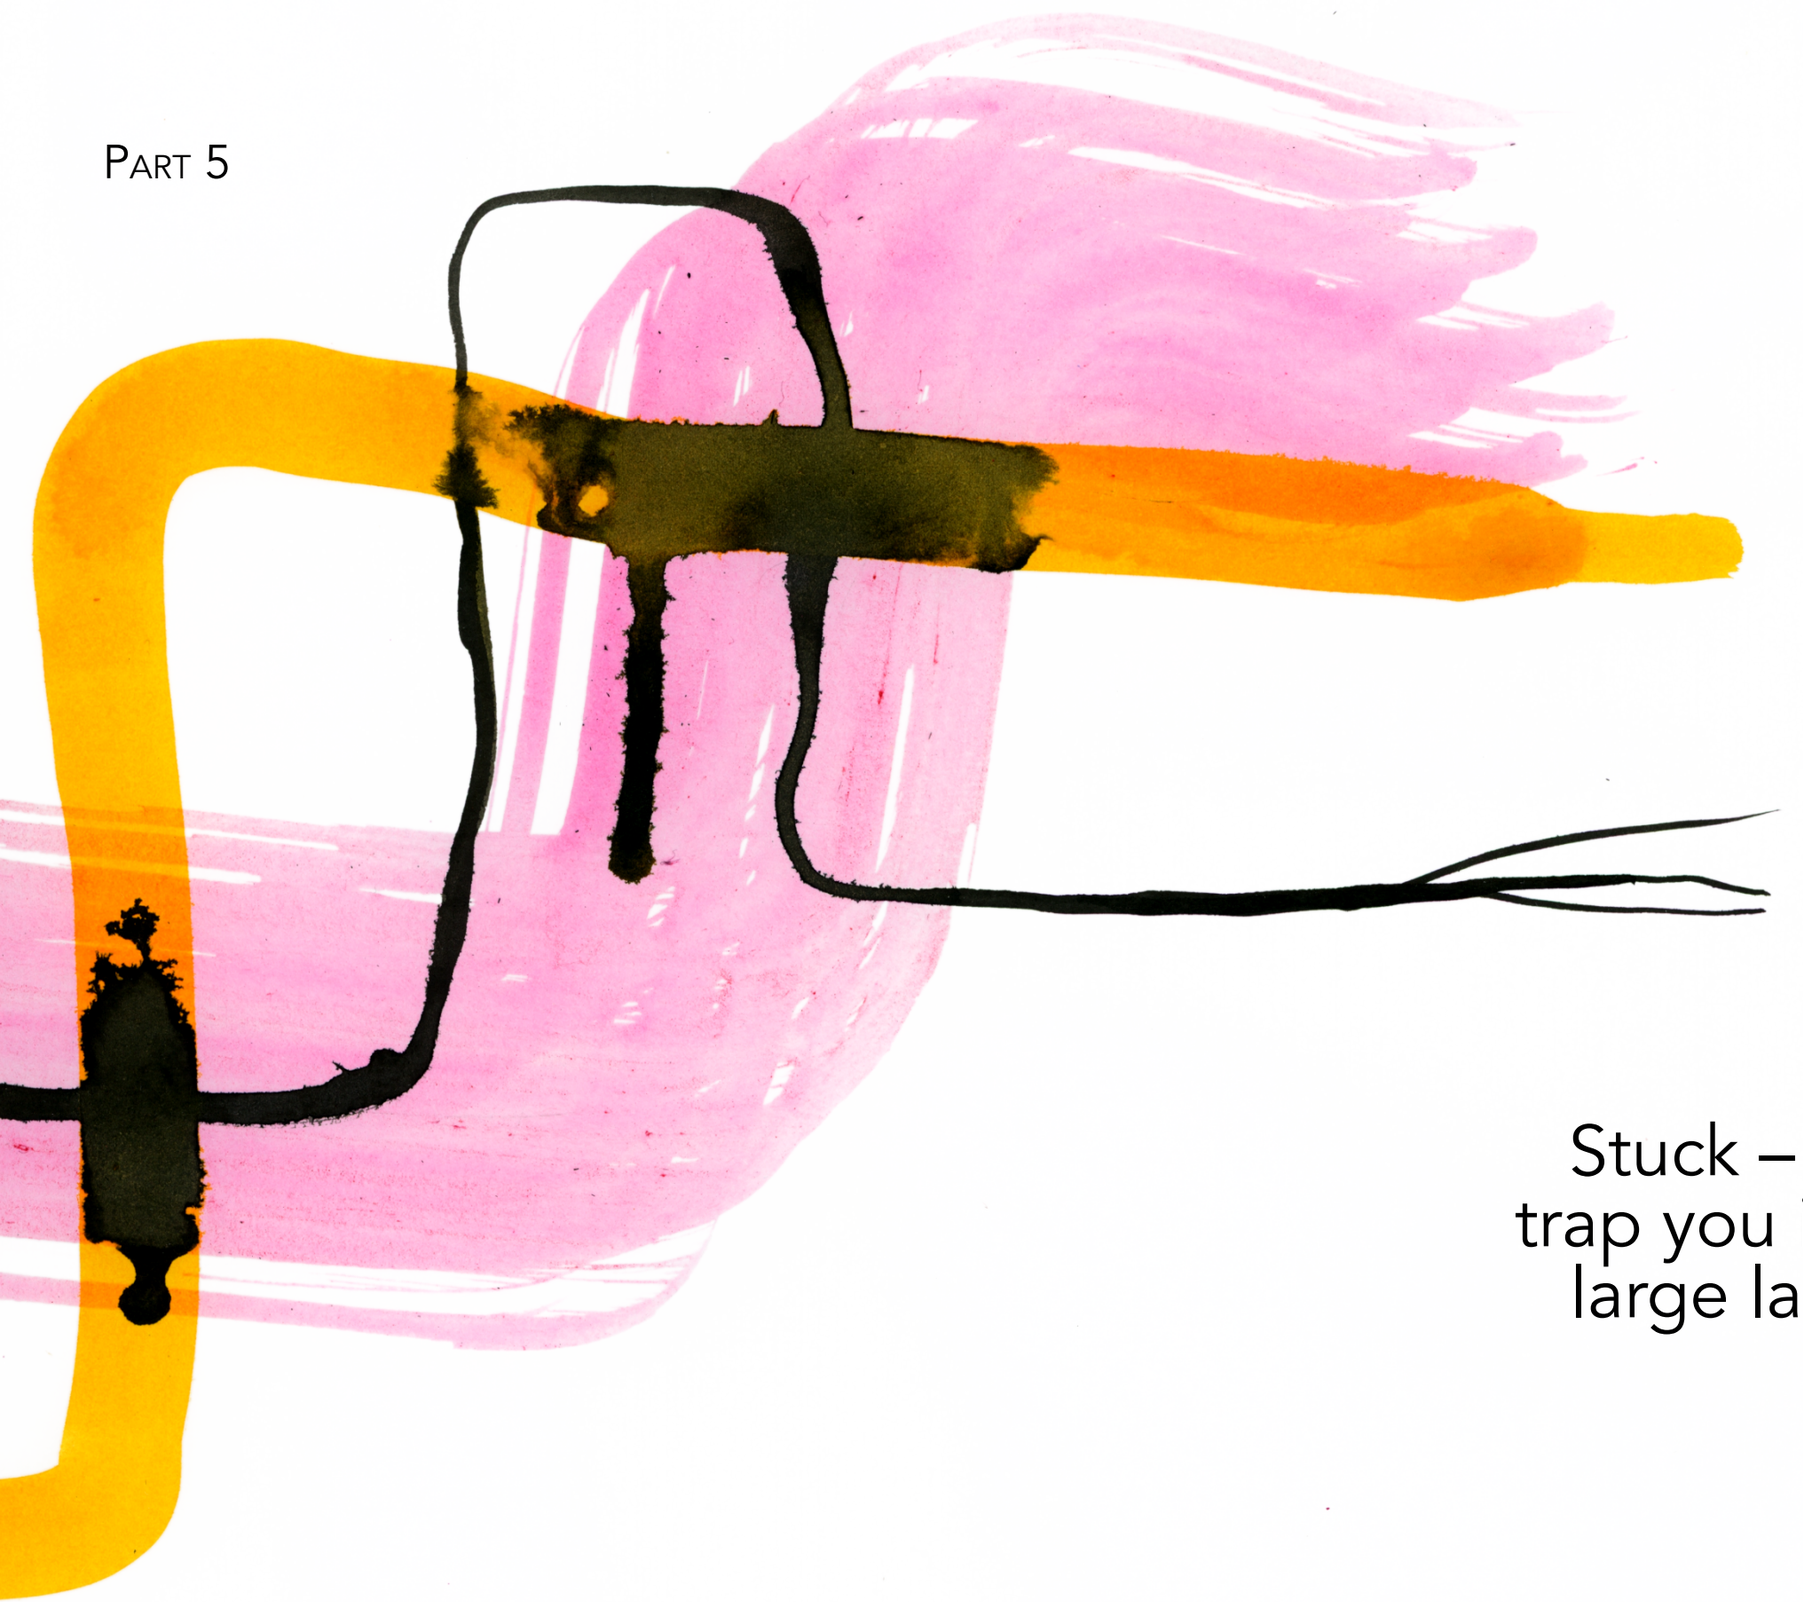

Stuck – How algorithms trap you in the labyrinth of large language models

Months fly past.

The computer scientists are still running tests
and fine-tuning the computational models.

In the meantime,
I embark on a human experiment.
We want to study how users experience
the personalised recommendations of news.

We set up eight computers.
We provide 37 people of different ages
with a login for the news site.

The developer has set the personalisation algorithm in motion
for these users,
so when they log in,
the stream of news adapts to what they click on,
like a robot following them around,
observing what the humans do.

In turn, we log what happens
as they navigate the personalised site
for a period of two weeks,
clicking and adapting the site to be:

"just for them".

By abstracting its mechanisms into symbols, Ada envisions a means to command the workings of the machine in ways no one had imagined. She carefully notes down:

"A new, a vast, and a powerful language is developed for the future use of analysis, in which to wield its truths so that these may become of more speedy and accurate practical application for the purposes of mankind than the means hitherto in our possession have rendered possible. Thus, not only the mental and the material, but the theoretical and the practical in the mathematical world, are brought into more intimate and effective connection with each other".[179]

To her friend Charles, she writes that she finds the work daunting. "The intricacies of the Engine demand a clarity of thought I must summon even as weariness threatens to cloud my mind. But I press on, for I sense that this is my purpose".[180]

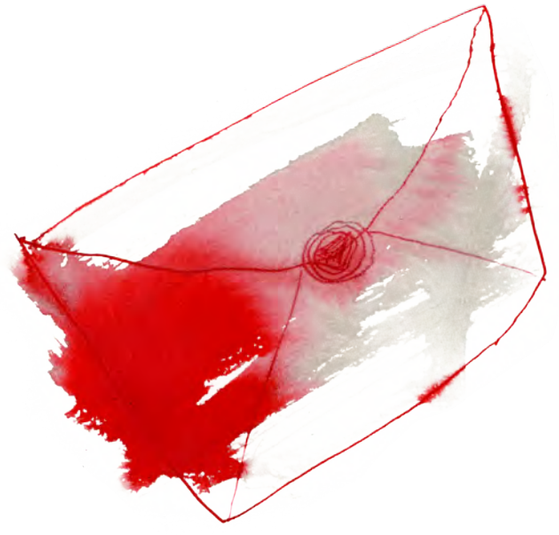

The commercial heritage of personalisation algorithms

In 1961, the Neiman Marcus department store in Dallas featured an IBM 1401 computer "Gift Advisory System". Customers filled out a multiple-choice form that included questions about "the age, sex, vocation, hobbies, habits, and marital status of the person for whom they were buying the present", and the customer's price range. The computer would then use punched cards to search a database of 2,200–2,800 Neiman Marcus items in its memory, and it would pick the ten best-fitting items.[181] The picture of the IBM 1401 reveals that the machine filled up an entire room – even bigger than Charles's Analytical Engine. Today, when I buy presents online, I don't have to fill in any survey: It takes just a few clicks for the machine to realise that if it's something like this that I am after, I "might like this too".

Personalisation algorithms did not sneak up on us overnight. Although most of us remember a world that was not filtered, at least not by automation, and where the same newspaper was delivered to everyone, we have been subjects to personalised automation for quite some years. Not surprisingly, the magic of personalisation resonated with advertisers' hot dreams of reaching "the audiences of one", promising the tailored experience, enabling the customer to find just the product that they did not know they needed or wanted. Not until later did this idea of the tailored news consumption experience filter from advertising to the news media industries.

As early as the 1960s, advertisers began experimenting with market segmentation, realising that consumers could be categorised into distinct groups with unique desires. Television commercials and magazine ads started speaking directly to suburban mothers, young professionals, or the emerging counterculture. Brands invested in demographic research, crafting messages that felt intimate and bespoke. This shift laid the foundation for the hyper-targeted advertising that would later be perfected by digital personalisation with recommender systems. What began as focus groups and consumer surveys evolved into an omnipresent, data-driven machine.

Like magic

The principles of algorithmic personalisation soon bled into everything: entertainment, shopping, and, most crucially, the news. Being accustomed to the curation of stories, it seemed like magic if the news industry could automate this curation and reach even more audiences with more relevant content "just for you". The same logics that could predict your ideal running shoes could be deployed when deciding which headlines you should see. The goal, ostensibly, was engagement – to spare you the irrelevance of stories that might not align with your interests, as disinterested audiences had increasingly become a major headache for the news media.

The social scientist Tarleton Gillespie, describing algorithms as an apparatus containing data, models, applications, and hardware, worried that the dynamics of algorithms can reinforce echo chambers, as algorithms deliver personalised content, fragmenting the audience into isolated, polarised groups.[182] Users are presented with a curated reality, constructed to keep their attention, but somehow detached from the broader, diverse public sphere. Both Eli Pariser and his notion of the filter bubble,[183] and Kyle Chayka's arguments in his book *Filterworld*,[184] portray similar concerns – while others, like Safya Noble, have rightly also pointed to the fact that the industries that use automated recommendation generally display a complete lack of economic and moral interest in minorities, which their algorithms manifest.[185] The commonalities of these studies are that they point to the fact that algorithms are imbued with power relations – just as so many previous technologies, but to an even larger degree, as they are much more ingrained in our everyday lives, and thus also in culture and societal value formation. Machines and the programs which run them are no longer something behind the great walls of smoky factories, but literally in our hands to the extent where people feel like they are missing a limb when they have forgotten or lost their phones. Designed to guide you in a certain direction – designed to keep your attention.

Because they are imbued with power relations, they lead you down the path of attention in different ways, but always with the aim of making your journey as frictionless as possible, so you don't notice.

Our participants for this experiment are now sitting in front of one computer each.

They click and scroll
while we interview them.

As they try to make sense of the stories
the algorithm has offered them,
they take a tour around the algorithmic maze,
revealing how they feel lost.

Paradoxes emerge across the data –
a confusing image
of the user experiences
of "just for you".

As researchers, we are weaving an image
of the user's experiences with the algorithms.
Experiences of words, images,
contradictions, insecurities,
and feelings of joy and creepiness.

All of which is invisible in A/B testing.

Ada has no time for the kids as she works on the translation of Charles's article on the Analytical Engine.

Mr. Babbage's ingenious machine continues to reveal to her its vast potential – not merely as a calculator, but as a weaver of algebraic patterns and truths. This translation of the article, while demanding in itself, has become so much more than a simple rendering of words.

She scribbles down all her thoughts next to the translation, each one a thread in a fabric she hopes to see woven into something of lasting significance.

Ada writes that the Engine "might act upon other things besides number" and suggests – by analogy with the mathematical expressibility of musical relations – that it could even "compose elaborate and scientific pieces of music".[186]

The very thought sends a thrill through Ada.

The news user as co-producer of their own newscast

From the many examples of robots taking over for the human from art and literature, we can safely claim that the human–machine divide has a long history. And it's a powerful one. Most often in this neat division, the human is seen as the subjective, reasoning, and thinking subject – deeply feeling a piece of art or music, which cannot be transferred to the more material and objective machine. Yet paradoxically, the algorithms that govern our digital lives are given a kind of magical status as all-knowing and neutral in their evaluations of our needs and desires. The magic nature and all-knowing algorithm is perceived to make us, the users, fully autonomous, in charge of our own media consumption patterns, as we click around the streams of endless content. This distorts the fact that the personalisation algorithms set quite profound boundaries for that action, restricting our autonomy as media users in significant ways, an issue I return to later in this book. However, we can already establish that personalisation algorithms are much more than just a new way of distributing content, as users engage with the algorithms in a much more fundamental way than they were ever in contact with front-page editors in the pre-algorithmic age. Users become co-producers of their own newscast, like Negroponte imagined in his vision of *The Daily Me*.[187] Furthermore, the instalment of the algorithmic butler raises issues around the expectations of the news users of how well the algorithm might "know them".

Take one of the media users we had enrolled as a research participant in a series of workshops on personalised news use held at the University of Copenhagen, organised by me and a group of fellow researchers. The woman, who was Danish and in her late 50s, was scrolling frantically up and down the site, not even noticing the "carefully selected for you" box halfway down the news page, which for two weeks had adjusted itself to her click history and personal data. "I am not sure why I am getting this", she said about a story on house prices in an area just north of Copenhagen. "I am not selling my house, but perhaps I am getting it because I visited a friend in this area yesterday?" she asked herself. She nodded to herself, looked at me, and continued scrolling up and down the page, hoping to find more stories about her favourite football team. Of course, the story on house prices was likely to be randomly selected for her, and perhaps even an attempt to sell stories to her that she would not normally click on, but it's an example of how citizens increasingly navigate with and around algorithms. Thus, they do not only seek to understand why they are being given a specific collection of stories, but they also feel some kind of responsibility for the stories they receive as a consequence of their previous clicking history.

She was getting to know the algorithm and expected it to already know her. Surprised when it did not know her as well as she expected, she blamed herself and tried to find an explanation. As it turns out, the bombardment with news stories on house prices was an error of the algorithm, which she paradoxically tried to find an explanation for in her own past movements and clicks.

Cultural taste stored in neural networks

Most of us recognise this feeling of being recognised or misrecognised, described in the critical data literature as "creepiness".[188] Like when my friends and I felt like our phones had listened to our conversations on future travel plans and used it as an opportunity to sell us a trip to Marrakesh. But what does it mean when what is making us feel creepy is not just the search for the new washing machine but a news and information-seeking process? A daily practice which – if we adhere to the public mission of the news media – is also supposed to inform us about the latest policy developments from parliament, crimes committed and solved, decisions made with impact on our lives. Stories meant to cultivate us into active, deliberating autonomous citizens in a democracy.

Despite the apparent knowledge of those limitations of the computational, taste – and hierarchies of cultural taste – is transported as part of this computational language processing to target personas, or clusters of datapoints. These hierarchies of our cultural taste – for music, news, and films – are stored in neural networks as collective memories of our sayings and doings. The algorithm later feeds those

sayings and doings back to us, mirroring how media perceives us, in much the same way as language has always been a transporter of cultural meaning throughout history. Though there are glitches of creepiness at times, most of our journeys through the information landscape proceed in a frictionless manner. News and information on recent world developments guide our interactions with the world and other people around us – and now the guidance is personal.

They shape the image of the world into a *personal* one.

A convenient computational rendition of reality

The reconditioning of social possibilities through data integration parallels what Jannis Kallinikos referred to as the "computational rendition of reality", where computational systems remake the fundamental principles guiding people's interactions with the world.[189] This transformation is significant, he argued, as it further naturalises and renders data invisible within the social fabric as individuals contend with the uncertainties it brings. Our taste for news and information is increasingly automated in these algorithmic systems, and we find ourselves as news anchors in our own personalised newscasts. And users tend to internalise any discomfort they might experience when looking in their data mirror, in return for the immediate gratification of being guided towards more relevant news confirming their outlook and perspective on the world.

The image of the world that journalists are feeding back to us is based on their function as gatekeepers who, on a daily basis, negotiate, choose, and navigate the overall news picture. They curate "All the News That's Fit to Print", as *The New York Times* slogan famously said. Journalists act as a form of filter, kind of like the algorithm rooted in datafied decision-making, but without the personalisation. They might think of you when they write the story, but it's an abstract you, often based in assumptions and imaginaries of who you might be. The journalistic filtering and sorting involve mixing stories that might be personally interesting for the user, some of which are relevant for them on a societal level, but which might

be very far from their everyday and individual needs here and now. They use their journalistic professionalism and normative public orientation when deciding what's relevant and what's not, with the purpose of curating "an image of the world" so we as citizens can keep track of what's going on. And what we are supposed to have an opinion about as citizens in a democracy.

This way of curating an image of the world is fundamentally shifting with personalised algorithmically curated news.

You are the co-gatekeeper of the news

Media scholar Tanya Kant has argued that "unlike other forms of personalization, it is algorithms that are responsible for determining what is of relevance to the individual, and so it is the 'system' and not the 'user' that is positioned as gatekeeper in the decision-making process".[190] But I would add, it's much more far-reaching than the algorithm becoming a gatekeeper or replacing the journalists as gatekeepers of information. Because not only is the algorithm a result of a process built by people – often data science engineers – the user also becomes a co-gatekeeper, as their choices send back signals to the algorithmic systems when we choose one news story over another. Signals which help the system build a persona – a cluster of aggregated datapoints of different users with similar interests, needs, and behaviours.

Signals used to curate the image of the world of tomorrow.

Put differently, your choices to click different news stories affect what stories you are presented with and thus the image you see of the world the next time you visit the news site. In effect, this means that the user herself decides on *her* own image of the world, giving the user the illusion of some kind of control – but it's a restricted form of agency, within the pathways the algorithmic models have made for you. And the algorithms learn how you prefer the world to look through self-learning systems. The system is adaptive, yet only within certain rather limited spheres of content (a single news outlet can only produce so much content).

Kant stated that "the decision-making capacity of algorithmic personalization systems is championed by advocates as time-saving, pleasurable, and convenient" for users.[191] I myself often ponder about the convenience of being able to filter out all the bad news of endless new wars being started, bombs dropping, starving children, or women getting killed on the streets. Instead, like I do with my stream of music, I would be able to adapt the stream of content to my taste and avoid news which would make me sad. Or angry. It's not hard to imagine how a watch would be able to detect my mood of the day and offer me news to either put me in a better mood or stay in the mood I am already in. Of course, we know from work of others that there is a dark side to the convenience of smoothly delivered personalised content. As Tanya Kant said, "even when it is enacted in pleasurable or time-saving ways, consuming algorithmically personalized content or services still means allowing the algorithm to act in your stead" – *on behalf of you*. "As such, algorithmic personalisation inherently involves affording some autonomous control to the algorithm as gatekeeper", she argued.[192] But it also means that the user is given quite a lot of autonomy over what news and information they are presented with, as the algorithmic gatekeeper takes notice of the signals and predicts what we might like next. Importantly, this matters when this logic is transferred to the domain of news, a domain in which we are not just buyers of shoes or coffee makers. The news is portraying today's image of the world, making us capable of enacting our democratic citizenship and in turn shaping the public sphere.

The editors I spent time with during my observations, who were of course in favour of the automatic curation of news by algorithms, argued that this will make them reach many more people, and allow for many more users to feel "seen" by the system. The algorithms are just a helper, an assistant or a butler delivering the news on a silver plate, they argued. And the news would feel more relevant for each of us. How could anyone oppose that, I often thought as I sat in on the meetings discussing the recent developments of the building project. But a doubt that this wasn't the whole story was nagging me.

Our experiment paints a messy picture,

an abstract image
rather than something
that can be easily interpreted.

One user wants more
of what she is already exposed to.

Another wants less.

A third is worried about
not receiving foreign news
in the future
if he does not click on them.

A fourth is avoiding clicking,
because he wants to hide what he
likes from the algorithm.

As if he is playing hide and seek with the algorithm.

He does not follow the path directed by the algorithm.
He is going off the beaten track.

And yet the algorithm adapts to the signals of breadcrumbs
he leaves behind in the forest.

One night, crunched over the desk with both husband and children sleeping, Ada thinks to herself: “Still the Analytical Engine is only a dream, a phantom. Charles, you should spend all your hours of the day to work on the Analytical Engine”.

She finally understands how this design of the machine and the method of indifference, which her friend has so carefully thought out, would simplify the process of calculating a long and complex mathematical series of calculations. It allowed for otherwise difficult multiplications to be replaced by numerous straightforward but very tedious and monotonous additions.

And the beauty of it being that, when carried out by a machine, the monotony of the calculative tasks would not matter anymore.

With the translation near completion, she allows herself a small but deliberate expansion. In Note G, she inserts a footnote.

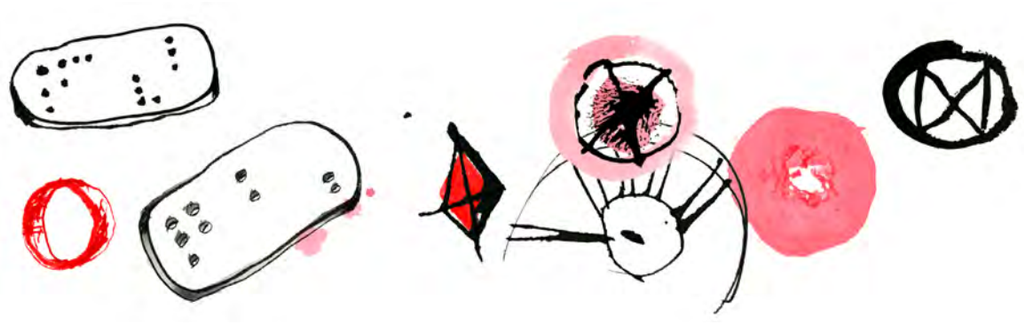

An invisible architecture of user surveillance and control

In the early days of personalisation, companies experimented with what became known as the rule-based targeting. One of the most significant developments in personalised advertising came in the late 1990s with Greg Linden’s collaborative filtering model.[193] Linden, a computer scientist, helped refine the concept of recommendation engines while working at Amazon. His model was based on the simple but powerful idea that if two users exhibited similar behaviour – browsing the same books, purchasing the same movies – they were likely to have shared interests. This insight led to the now-ubiquitous “Customers who bought this also bought…” feature, a form of predictive personalisation that felt both intuitive and eerily perceptive.

The appeal of collaborative filtering lay in its ability to predict preferences without requiring explicit user input. Unlike traditional advertising, which relied on demographic data and assumptions, Linden’s model worked dynamically, learning from each interaction. This made it incredibly effective, and customers responded with enthusiasm. They appreciated the convenience of seeing products tailored to their tastes rather than sifting through an endless sea of options. What began as a simple tool to boost Amazon’s sales soon became an industry standard, influencing everything from e-commerce to media streaming services.

The auction system of attention

Meanwhile, Google was perfecting another form of hyper-personalised advertising with AdWords. Launched in 2000, AdWords transformed advertising from a static broadcast into a dynamic, data-driven engine. Instead of displaying ads based on broad categories, Google allowed advertisers to bid on specific search terms, ensuring that ads appeared in front of the most relevant audience. If you searched for “running shoes”, you wouldn’t just see an ad for any athletic brand – you’d see an ad for the exact brand and model most likely to interest you.

AdWords' brilliance lay in its efficiency. Advertisers only paid when users clicked, making every dollar spent more accountable. Users, in turn, were presented with advertisements that felt less like interruptions and more like helpful suggestions. Over time, this transformed digital advertising into an intricate auction system where businesses of all sizes could compete for consumer attention in real time.

While Amazon and Google pioneered personalisation, Facebook, or the parent company Meta, perfected it. The social media giant moved beyond simple ad targeting and into full-scale behavioural influence. Facebook's interest-based advertising model went further than tracking what users searched for – it inferred what they wanted before they even expressed it. By analysing likes, comments, shares, and time spent on specific posts, Facebook built detailed psychological profiles, allowing advertisers to target users not just by demographics but by their aspirations, fears, and emotional triggers.

EdgeRank was Facebook's original algorithm, designed to prioritise posts in users' News Feeds based on affinity (interaction history), weight (type of content), and time decay (recency). Over time, Facebook replaced EdgeRank with more advanced machine learning models, which analyse thousands of signals – including dwell time, reactions, shares, and even cursor movements – to predict user behaviour with eerie accuracy.[194]

Unlike Google's AdWords, which responded to explicit user queries, Facebook's model anticipated desires by nudging users toward particular behaviours. If you lingered on a post about a Mediterranean vacation, your feed would soon flood with ads for airline tickets, luggage deals, and dreamy beachfront rentals. But the personalised targeting and predicting didn't stop at commerce. Political campaigns, activist movements, and corporate brands all leveraged Facebook's immense data troves to shape narratives, subtly steering users toward specific ideologies or behaviours.

The dark side

Not surprisingly, this level of personalisation raised ethical concerns, particularly regarding privacy and autonomy. Shoshana Zuboff, in her book *The Age of Surveillance Capitalism*, argued that Facebook had moved beyond mere advertising and into a new economic logic – one where human experience itself became the raw material for profit.[195] By collecting vast amounts of behavioural data, companies like Meta went beyond predicting user actions: They actively worked to shape them, turning passive consumers into predictable, pliable participants in a commercial ecosystem. The algorithms used by these companies are normally described as neutral – simple mathematical and objective formulars. To be fixed whenever a problem emerged.

Regulatory scrutiny intensified as revelations about data misuse surfaced, from the Cambridge Analytica scandal to growing concerns over algorithmic bias. The Cambridge Analytica scandal erupted in 2018 when it was revealed that the political consulting firm had harvested data from over 87 million Facebook users without their consent. The firm, linked to Donald Trump's 2016 campaign and the Brexit referendum, used psychological profiling to target voters with highly personalised political ads. A whistleblower, Christopher Wylie, exposed the scheme, revealing how Facebook's lax data policies allowed third-party apps to siphon personal information. The backlash led to Facebook CEO Mark Zuckerberg testifying before Congress, a record-breaking 5 billion US dollar fine for Facebook, and a global reckoning on data privacy. The scandal underscored the dark side of social media manipulation with the combo of personalisation and large amounts of data.[196]

Slowly, users began to question whether they were engaging with content of their own volition or if their digital environment was being subtly manipulated to maximise engagement. The promise of personalisation had morphed into a surveillance-driven ecosystem where users were not just consumers but were indeed the product itself.

Tailoring the conditions of possibility

Nowhere is this more evident than in the world of news personalisation. Recommender systems – the sleek, data-driven engines behind Google News, Facebook feeds, and X timelines – now also promise to deliver the most relevant news stories. *Just for you*.

When observing what algorithms are and what they do, we can see how they are far from neutral, objective, mathematical formulars based on cleaned-up, identity-protected, large-scale datasets. So, questions of whether we should trust the companies with our data – and trust them to build fair and transparent models based on our data for our future convenience of finding the right products, answers, or news – ignore the deeper contextual factors shaping the algorithms and our use of them in our daily practices of engaging with news and other forms of cultural consumption like films and music.

Data practices become not merely referential but also foundational, deeply entrenched in how individuals structure their everyday experiences. In this way, data and the data double solidify their place in the social world, contributing to the subtle modulation of social possibilities. The scholar John Cheney-Lippold noted that data "configures life by tailoring its conditions of possibility", reconfiguring the basis on which individuals predict, understand, and respond to the social world.[197] As data mechanisms integrate into everyday routines, data has, paradoxically, dissolved into mundane actions. This reflects a more substantial process of domestication,[198] where once-foreign technological processes are accepted and integrated into the fabric of everyday life. To use the social scientist Charles Taylor's phrasing, data can be considered "constitutive goods" that empower people to do, and to be, good,[199] facilitating the fulfilment of social roles, such as being a "good person". The human gets tangled up with the machine, which also sets norms for how this human should govern themself. As a quiet form of control. A governance of the user, not far from the cultivating of citizens that was the public mission of the news media but now built into computational machines and languages with no – or at least no visible – mission other than making people stay and consume.

The news has always been this form of constitutive goods, as they also cultivated consumers into becoming democratic citizens. The citizens who wrote media diaries in my earlier research project informed me how news both had a social function in their lives as it enabled them to talk to friends and peers and that the practice of keeping up with the news was also more than just being informed. It was about feeling a sense of belonging, and that the world is still standing, day in and day out, although it could sometimes seem chaotic and out of balance. When the news is personalised and filtered, this transforms what kind of "good citizens" the media expects users to be. The users I interviewed recently in workshops I did with colleagues at the University of Copenhagen expressed a fear of missing out on such stories that would "expand their worlds". They articulated how they would worry that when the news was recommended to them personally, other articles would be left out of their newsfeed. The participants were concerned they would miss important current events because they had not previously shown any or only limited interest in the given topic. And yet, the news users also at times felt that it was overwhelming to navigate all the news of the world.[200] The story of the human–machine entanglement is indeed a messy one. Perhaps even more messy than the older story of the relationship between the news media and citizens – because we now have this third algorithmic actor in between.

Jonathan Cohn, in *The Burden of Choice*, describes this process as an essential part of modern life and as an exercise in self-curation, where the pressure to choose – from toothpaste to political beliefs – has turned freedom into an exhausting mental load. In this way, algorithms make our lives easier by recommending what we might like, while they also confine us within the invisible architecture of their suggestions, shaping the very range of options you believe you have. As Cohn writes, the algorithms are actively shaping – and have been shaped by – changing conceptions of gender, sexuality, race, and class.[201]

This means that choice is never just about the options between what media content we are exposed to; it's as much about the structures that make some choices seem natural, even inevitable, while others vanish into obscurity.

The man sitting next to me
is still soaked from the pouring rain

he encountered on the way to our experiment.
On the way here, he spoke to his wife, he tells me.

They briefly discussed the falling energy prices,
and how they might go up again
as the war in Ukraine continues.

As he scrolls up and down,
he tells me how he feels
that some of the stories on the site are pushed to him.

Like the algorithm wants him to behave in a certain way.

He explains to me how he avoids clicking
on the click-bait headlines:

“I don’t know what it’s about, and if I click, I risk getting
more, even if it’s not relevant to me”,
he says.

He is navigating the algorithmic maze of relevance.

The hours have slipped away. Ada is completely consumed by her work on the newly inserted footnote in the translated article on the Analytical Engine.

As her notes are growing, she discovers how algebra will serve as the key to expanding this machine's reach into realms beyond simple calculation. Algebra is a branch of mathematics that generalises arithmetic by using letters and symbols (variables) to represent numbers, unknown quantities, and formulas in equations.

She imagines how it might be possible to transfer the mechanisms of the machine into symbols, creating a language of symbols that could run the machine. She imagines a language of automation in the mathematics of algebra. She imagines how the symbols can become adaptive: They are dynamic forces carrying within them the power to manipulate entire dimensions of logic and numbers.

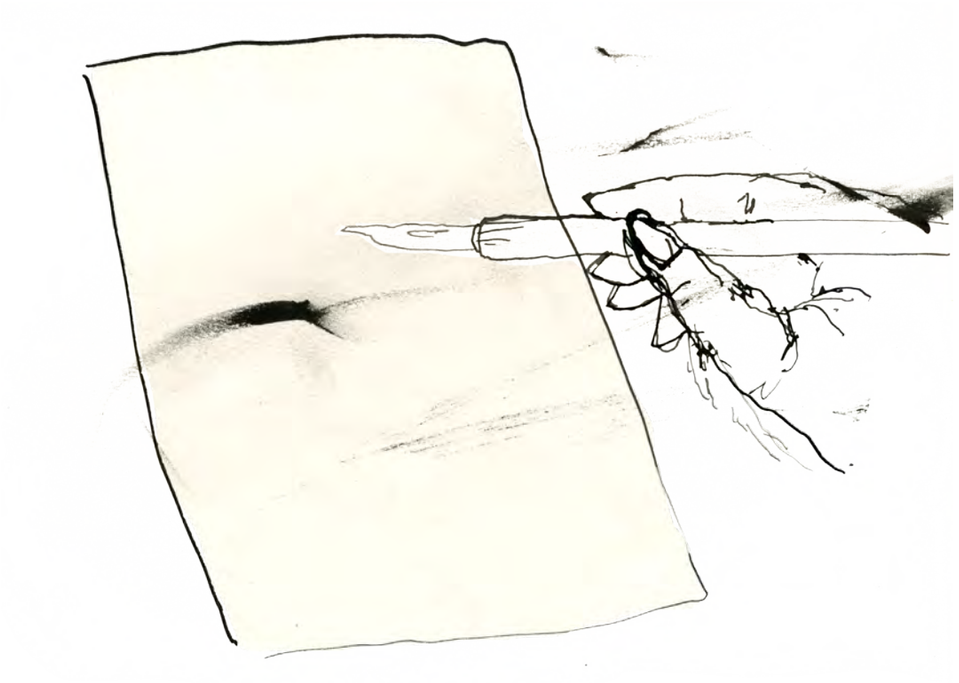

Caught in the endless loop

For a long time and like many others, I have been aware of and happy about personalisation algorithms helping me out in the vast sea of information, helping me stay with the music I like and the films of my taste. I have been confused when one of my children used my Spotify account, introducing me to Minions singing YMCA in Spanish, and annoyed when my husband's interest in Star Wars made me feel lost on the streaming platform. But most of all, it became significant to me how personalisation acts as a mirror of our choices but also *controls* those choices.

When I was at a research stay at Cardiff University, I stayed in an Airbnb a ten-minute walk along the river from the university, and coming home, I would often take the luxury of enjoying some food in front of the large flatscreen television in my temporary living room. I could not remember my Netflix login, and funny enough, Netflix did not ask but let me choose one of the profiles already logged in without providing the password. The mysterious profile, this life of a stranger that I was sneak-peeking into, was called Ngozi97, and I felt both excited and guilty of privacy violation, as I let myself into this universe of the cultural taste of someone I had never and never would meet. As I was scrolling down the "continue watching this" and "today's picks for you", I was amazed. Not a single film or series that I had even remotely heard of. Not a single actor's face I recognised. The majority of the recommendations were Nigerian soap operas, and I was mesmerised by the extremely saturated colours and bold lighting. Such a contrast to my usual gloomy Nordic Noir crime series set in dark forests. Most of the Nigerian soaps featured scandalous affairs, with beautiful actors in lush Ankara fabrics, who would swirl through family dining rooms tense with generational trauma. The production quality was poor, but what drew me in (apart from the scandals and affairs they got themselves into) was the extreme close-ups of faces so tight that I would see a character's heartbreak sweat and the way the long montages were soaked in melancholic Afrobeat. So unfamiliar to me and yet it struck me how my own recommendations would probably seem as unfamiliar to Ngozi97. And the more I watched, the more I was offered, and the more I felt guilty that I was messing up the strangers' algorithmic offers with

my ignorant choices and clicks. It was a relief when Netflix suddenly interrupted my trespassing behaviour and asked for the password, which I of course did not have.

Master of the algorithm

In the logic of the algorithm, you are not an autonomous reader, freely roaming the intellectual landscape. You are a profile – a pattern of past clicks, measured engagement times, inferred political leanings – fed back into a machine that optimises your world into a loop of familiarity. As a result, the news you consume is not just personalised: It is *prescribed*. This is why it felt so weird to log in as someone else. The burden of choice, in this case, is not just the fatigue of picking among options – it is the creeping realisation that the choices you see have already been made for you. As Jonathan Cohn argued, by clicking, personalisation algorithms invite us to constantly think and re-think who we are as human beings. They do so because they work as a mirror of our choices while also saying, "you made this choice yourself, this is who you are". The recommendations given to you by the algorithms seem to give you the *illusion* of a lot of choice, because they adapt to your clicks, while they try to predict based on that click what you might like next. "Interpellation and control are concomitant (and often indistinguishable) forces in that they both work by paradoxically presenting subjectification and dependency as symbols of autonomy and freedom", Cohn wrote.[202]

By resisting in the form of clicks or choosing not to click we do have agency as users, albeit a limited form of agency, where many of the paths we can choose are pre-chosen for us. In our resistance, we are teaching the systems something about ourselves while also reproducing certain distinct ideas and images that the systems hold of us. We might feel empowered and feel that we have mastered the Spotify algorithm when we one day decide to go down a different musical alley, but where did this new idea come from? Perhaps our musical adventure off the beaten track is in fact based on a recommendation from a friend? Essentially, these friends, or "people like you" recommendations, are the core function of the collaborative filtering model, one of the many algorithms also used by news organisations for personalised recommendations of news. Predicting – specialised in how to attract what we want and need – could be a key to rebuilding the relationship with users whose attention has wandered off to other platforms and sites.

Optimising for retention by the minute

Nowhere is the direct response to your clicks more responsive than on TikTok's *For You Page* (FYP), which is an ever-evolving, hyper-personalised stream of content engineered to seize attention within seconds. Unlike YouTube, which optimises for long-form retention, TikTok thrives on short bursts of engagement, using a labyrinth of machine learning models to ensure that the next video is always just as compelling – if not more so – than the last. Comparing how different algorithms work differently enables us to see how they also nudge us into different behaviours over time.

At its core, TikTok's algorithm is a relentless observer, tracking every micro-interaction: how long a user lingers on a video before swiping, whether they replay it, pause it, like it, comment, or share. Even hesitation – a brief flicker of indecision before scrolling – becomes a data point, feeding an AI system that learns in real time. Videos are categorised by more than just their creators or topics; the algorithm picks up on shared themes, editing styles, soundtracks, and visual patterns, clustering content into invisible, fluid categories that shift and recombine like a digital kaleidoscope.[203]

Unlike YouTube, which relies heavily on subscriptions and search history, TikTok's algorithm functions with radical immediacy, treating each new viewing session as a fresh slate. In the first few hours of a video's life, the platform quietly distributes it to a small test audience. If engagement is high – if people linger, replay, or interact – it gets pushed to a slightly larger pool, and then larger still, a cascading process that can transform a random clip into a viral phenomenon overnight. New creators, even those with zero followers, have a chance to break through – a stark contrast to YouTube, where established channels hold a significant advantage.[204]

But the most powerful trick in TikTok's playbook is its ability to predict. By mapping out a user's past behaviours, the algorithm anticipates what they will enjoy next, fine-tuning their feed with an uncanny sense of intuition. It is not merely showing content the user might like – it is actively shaping their taste, reinforcing patterns, nudging them deeper into digital niches they might never have sought out on their own.

YouTube, by contrast, plays a longer game. Its goal is not just to recommend the next video but to keep users inside its ecosystem for hours at a time. Its ranking system prioritises videos that lead to extended watch sessions, favouring deep dives, multi-part series, and content rabbit holes. While TikTok optimises for the immediate dopamine hit, YouTube cultivates slow-burn obsession, where one click leads to another, and then another, until hours have disappeared on endless stream of content.

However, using behavioural data to predict what you might like next has a long history, and as it turns out, it was already part of the ideals behind personalisation algorithms before the phenomena entered the news media domain.

Our individual interviews have come to an end.

As they sit together, not knowing each other,
they reflect on their data with the personalisation algorithm.

"I have the expectation that when I open a newspaper, they will present me something, which I have not chosen myself. I mean, I don't want everything to be selected for me beforehand. Then it's no longer a newspaper. I need to get some knowledge".

"I agree, then you don't know what you are not seeing, right?"

"Exactly, and I think that's a dangerous trend, that we are all withdrawing to our own little groups". "Our own personal newspaper, right?"

"So, I would want some stories, which everyone will get, like what's going on in Russia. And I also want everyone else to have read that story too. It means that we can talk about it, as it is important for us all…"

"I would feel cheated if the algorithm tells me that I probably don't want to read that story. And I would never get back on track, because no one would realise that I might have wanted to read that story anyway…"

They feel like they are lost in a maze they cannot leave.
If they could just see the maze from above.They fear not seeing the big picture. They fear becoming publics-of-one.
Even if the maze appears to have endless ways around it, they find themselves going in circles.

Note G and the footnote advances a claim that Ada has tested repeatedly in her mind: The Analytical Engine shows a particular aptitude for the calculus of finite differences. The idea grows directly from her reading of the machine's structure – its ability to move through ordered variations, to carry change forward step by step without deviation.

To make this capacity visible, she turns to the computation of the Bernoulli numbers by means of differences of nothing. The choice pleases her. From what appears to be absence, a sequence emerges, shaped entirely by arrangement and rule. Operations follow one another with calm inevitability. Instruction, not substance, drives the result. The principle settles into place, resonating across the architecture of her Notes. The instructions will guide the movement of the machine, and the machine will obey the rules.

Ada prepares to present the footnote to Charles with measured assurance, slightly nervous about how he might respond. She feels convinced though that this will be the breakthrough they have needed. It aligns with the mathematics. It aligns with the mechanism. More than that, she thinks to herself, it articulates what the Engine already suggests: that thought itself can be set in motion, provided its steps are composed with care.

The footnote thus joined the growing body of her annotations, where explanation, speculation, and instruction converged. The Notes had grown bigger than the translation itself, but for a good reason, she thinks to herself one evening, as she puts the kids to bed.

Together, they traced an early outline of algorithmic authorship – one in which the human role lay in articulating processes with sufficient clarity that a machine, once built, could carry thought forward on its own – and without the tinkering of the human hand.

A language of the machine which could predict the next step forward, guided by rules and instructions as comprehensive as the human language.[205]

First date with the algorithmic news butler

If we take a closer look at how news organisations around the world frame the need to automate and personalise news curation, the urge to make the user more autonomous and freer by pure magic of the personalisation algorithms emerge as a clear pattern. The well-established and globally known legacy news outlet *The New York Times* was at the forefront of the realisation that the future is personalised when they began their work on personalised recommender systems around 2011:

> Picture a home page where the dominant spots on the screen show the big news and feature stories, but much of the surrounding content is tailored to your own interests.
>
> More limited experiments are already underway. News alerts, among the primary uses of personalisation right now, may be different depending on a particular reader's location. On the home page, there's a customized box called "Recommended for You" that lists articles Times data shows you haven't yet read.
>
> But these are small lab tests compared with the plans editors have for a next-generation New York Times, one that shifts from monolithic to something more bespoke.
>
> (Elisabeth Spayd, 2017, public editor at *The New York Times*, 2016–2017)[206]

When first announcing the new personalisation feature in 2011, Denise Warren, senior vice president and chief advertising officer at The New York Times Media Group and general manager at NYTimes.com, stated in a press release: "With a Web site as broad and deep as NYTimes.com, we are always looking for new ways to help our readers find news of interest beyond the sections they read most".[207] Starting in 2017, *The New York Times* began increasing the use of recommender systems and personalisation algorithms. According to *The New York Times* public editor at the time, Elisabeth Spayd, introducing more personalised features on the website would make *The New York Times* move away from the news as a monolithic entity towards being something more "bespoke" and "responsive" to its

audiences.[208] With this she meant that the news institution had become too set-in-stone, viewing things from only one angle: Its own. And this one-way view of the world no longer fit the flexible and fleeting user.

"Americans are used to their world being customised around their needs", Denise Warren argued. She emphasised how the newspaper must therefore become an organisation that "treats readers like individuals, with unique habits and preferences".[209] What she meant was that the world had grown accustomed to the niche marketing that Turow had observed years before. And to the "you might also like this" lingo that users recognise from Amazon, Netflix, Spotify, and other entertainment content platforms. It was as if users, being so used to personalised social media feeds, could no longer work out how to navigate a news site, so news organisations – or the algorithms – had to do it for them. But in ways much more attuned to the individual user, so the experience of reading the news would resemble that of scrolling around social media.

Hello, computerised butler!

Some of the first experiments with the personalised news site were of course also concerned with getting feedback from the users on the content, in order to provide more relevant news on an individual level. Initially, it was done in a similar way as the customer surveys, where people would indicate what their interests were, similarly to what we are asked to fill in when first creating a profile on a streaming service platform like Netflix. We tell them that we like romantic comedies and documentaries, and as we start using the streaming service, we confirm and at times adjust these choices with what we watch. Not surprisingly did the streaming platforms realise that they needed to provide separate profiles for each family member, as, for example, parents felt annoyed by being offered Peppa Pig or Paw Patrol. The algorithm simply adapted to whoever was using it, and with this being more than one person, the algorithm would not be able to detect who had now turned on the service. Thus, the user now signals who they are by clicking on a specific profile.

The user feedback perspective was very present in first development of the MIT personalisation algorithm from 1981 named the "News Peek" system (a play on "New Speak" from George Orwell):

> The reader could be explicit about personal or situational interests. Today's analogy is a search-engine query: the user makes a request that the search engine consider a large number of possible "articles" and select and present those articles for the reader's consideration.[210]

Feedback to the system

The ongoing troubled relationship between users' needs and editorial staff in the production of news had already been noticed by Walter Bender, Nicholas Negroponte, and the rest of the MIT team many years earlier, when they developed the next iteration of the news personalisation systems:

> Editors do not work in isolation. They receive feedback from the community they serve in a number of ways: letters, telephone calls, and/or e-mail to the editor; comments from colleagues; focus and market surveys; as well as hard numbers such as newsstand sales when a particular headline is run. This feedback allows the editor to determine the interests of the community members and better serve their needs.[211]

Interestingly, Bender found it necessary to emphasise that this form of personalised curation of stories was no different from the editorial curation. "This is neither more nor less than an editorial process", explained Bender, who was in charge of the development of the project at the time, when looking back at the project's achievements.[212] In the next system developed at the MIT Media Lab – *FishWrap* – this community feedback played an even bigger central role. This project later formed the groundwork for the work by Negroponte, the director of MIT during this period, who predicted, in his pioneering 1995 best seller, *Being Digital*, that digital media "would give citizens an unprecedented degree of control over their media environments" illustrating this contention with the hypothetical example of *The Daily Me*, "an online newspaper whose content would be customised to suit the interests and beliefs of individual readers".[213]

A really fresh idea

FishWrap – one of the first personalised newspapers – at first originated in a freshman seminar called "Newspapers of the Future". Led by Mr. Chesnais and Walter Bender, at the time a principal research scientist at the Media Lab: "In the fall 1992 seminar we asked the freshmen to design a paper [of the future]", Mr. Chesnais said. "We didn't tell them what's impossible, or what the state of our software is – we just asked them what they would like to see in such a paper".[214]

> At MIT, we began experimenting with computerized "butlers" or "agents" that were programmed to act on the reader's behalf, culling articles of interest from traditional and nontraditional news sources before sending them down the wire to the reader's home. Paralleling the decision-making process of human editors, these agents utilized multiple mechanisms to filter content for individual consumption, with the goal of "fine-tuning and prioritizing information based on criteria that include timeliness, importance, and relevance to the audience" – in this case, an audience of one.[215]

Their proposal: a personalised news paper aimed at incoming freshmen that would include news from home, local maps, daily schedules of MIT events, and more. That idea grew into the *Freshman FishWrap*, so named for its original target group and after a journalists' proverb: "Yesterday's news wraps today's fish".

> The system was originally designed to address the needs of freshmen being integrated into the MIT community. FishWrap attempted to balance an individual's desire for personalization with the need to participate in and know about the world at large.[216]

So similar to Facebook, the system was designed to connect people and make students at MIT know what was going on around campus, but in a very personalised way. This was made possible by a particular function: a database with information on both students and the

content they wanted to deliver to students. "One component used in the student designed personalised newspaper FishWrap merits particular mention: Jon Orwant's Doppelgänger", Walter Bender wrote while looking back at "Twenty Years of Personalization" at MIT.[217]

Doppelgänger

Doppelgänger is a user modelling system, enabling a tailoring of both the content and its presentation to the individual. A database was maintained with the "interests, plans, beliefs, behaviours, misconceptions, and schedules" of individuals, and it provided what the user "might like next" based on inferences from this information.[218] Though primarily used for article selection, the models could easily facilitate customisation, and they dynamically changed as the system learned about users and as the users continued to change themselves:

> The data maintained by the system includes objective and subjective information as well as long-term and short-term information. In addition, implicit information is inferred from the universe of user models: generalizations about populations of users are made and are used to establish default assumptions about user models. By monitoring users' actions, the system bolsters both precision and accuracy of the model over time.[219]

All items coming into *FishWrap* – from traditional news stories to submissions from the MIT community, to the content of social and cultural calendars – were analysed for relevancy and categorised, for example, as "Star Trek", "softball", or "artificial intelligence".[220]

Nudging you through the maze

Current algorithms for news use this technique to also sort the different news stories of the day, but most often, this is mixed with a human curating specific parts of the news site. Like product director for *Aftenposten* and Schibsted's subscription newspapers,

Karl Oskar Teien, and product manager for Schibsted's algorithm engine, Christoph Schmitz, explained in a blog post on 11 October 2023: "At Aftenposten, we believe in a combination of manually selected articles and algorithm-driven ranking of stories on the front page. The top five or six placements on the front page are manually selected based on carefully considered editing principles".[221]

They further clarified that "this ensures that the reader at any time will have a good overview of both today's important stories and the latest top stories" and that *Aftenposten* "has the option to manually pin certain stories further down the front page".[222] An algorithm determines what's on the rest of the front page, allowing the reader to stay up to date on important news stories that otherwise might be missed, as well as stories widely read by others. And they further explained how "metadata such as the editorially determined 'news value' and 'lifetime' of an article were critical ingredients to achieve this".[223] In other words, the public mission combined with the users' own interests put into the model in a beautiful symbiosis.

As an explanation to why this algorithmic curation was needed, Teien and Schmitz argued that the news organisation was publishing so much material that users got lost trying to find their way. This is the same argument I saw in and beyond the news organisations I observed. To achieve a more relevant experience for the readers, they analyse a large amount of data on their audiences, which they deem "signals" of interest, explaining:

> This will never be the only thing we optimize for, but it will be an important tool for giving narrower topics a chance on a crowded front page. The purpose is to help us match these narrow articles to the readers who have a special interest in that particular topic. This way of personalizing differs significantly from what most people are used to on Facebook and YouTube.[224]

The message from the Norwegian newspaper is that the user of news needs to be nudged, but you will be nudged down the path in a different way than what you know from social media platforms. To be fair, news media have always nudged their users (think of the big yellow headlines signalling "READ ME") – yet somehow, this nudging feels different.

I came across something which both fascinated and horrified me the other day.

A completely AI-based news channel was unfolding
in front of my eyes:

"We are a news service based entirely on AI", the compelling voice
says in the introduction to the Channel1 news channel.
The hosts appeared like ordinary news hosts, perhaps slightly
mechanic, but nevertheless they presented the recent
and very important news in a sharp and inviting manner.

The host would be able to switch into Hindu, Chinese, or Mexican
with a single click.

I watched one story after another unfold
with real footage on the war in Ukraine collected from the Internet,
and it became clear that personalised language
would be only one parameter
the audiences of the 100 per cent AI-generated news channel
would be able to change themselves.

The style of the host, the skin colour,
or even the level of positivity in how the news would be presented
would be some of the endless personalised possibilities.

The idea of someone choosing a naked news presenter
being the most obscure of those choices
future news users would play around with.[225]

Who was behind this, I asked myself?
Why were they doing it?
What was their mission?

And how would users feel
about purely machine-made news?
I was certain, however,
that this AI news channel would not generate news on its own.

Scraping the vast Internet
and the works of thousands of journalists around the globe,
it would generate AI versions of those same news.

But they would be transformed
and fitted to the individual.

A machine of predicting
and analysing patterns
of like and dislike,
fine-tuned so that the user
stays forever with them
and their automated robot hosts
telling stolen tales – repackaged

"just for you".

When Ada finally writes to Charles Babbage, the delay already carried its own explanation. The printers have occupied her days with an intensity that rivalled the mathematics itself. Proofs travel back and forth, arguments settle themselves into type, commas find their places. Finally, she feels they are done.

With that labour concluded, Ada turns to assess what now lay behind her. She has pressed every part of the Notes toward precision, shaping arguments, ordering examples, adjusting the cadence of explanation. The effort satisfies her. The structure aligns with the range of ideas she had set out to explore, and the internal logic of the text followed those aims faithfully. Each section knows its place; each transition carries the reader forward without strain.

Seen as a whole, the Notes now strike her as complete – coherent in their scope, assured in their execution. Her letter to Babbage marks a quiet handover: The work had reached a state where it can speak on its own, guided by the order she has given it.

"The Analytical Engine has no pretensions whatever to originate anything. It can do whatever we know how to order it to perform. It can follow analysis, but it has no power of anticipating any analytical revelations or truths. Its province is to assist us in making available what we are already acquainted with", she writes as a final touch before putting the pen down.[226]

The Notes carefully explain how a series of commands would run the machine – a form of language instructing the machine what to do. But she also emphasises that it can only build on existing material, the only truths being the ones put into the material by the people building the instructions.

The Notes held in them the invention of the algorithm.

Algorithmic creepiness

The big question is, of course, to what extent we are aware of being pushed and nudged around the labyrinth of consumer and information choices. After all, the aim for much of the design coming out of Silicon Valley was indeed to make our user journeys as comfortable and frictionless as possible.

Increasingly over the years, users started responding to "being seen" by the algorithms, and the responsiveness of the algorithms to the data provided. The experience of creeping algorithmic intimacy is as uncanny as a stranger knowing your childhood nickname. In the scholarly literature, this feeling has been named "algorithmic creepiness" – that eerie sensation when recommendation systems appear to divine our inner lives, serving up ads for things we only thought about, or nudging us toward content we didn't know we wanted.

Through the concept of "algorithmic imaginary", media scholar Taina Bucher looked at how users become aware of the algorithm and how this awareness affects their use of media. Bucher documented the "whoa" moments where users discover the implication or workings of the algorithm on their mediated experience,[227] what is often referred to as "creepiness" – a concept she has from the work of Lupton and Michael.[228] Bucher pointed to the fact that users are able to recognise "faulty prediction" or seeing through the "popularity game" that is part of the logic of Facebook, for example.[229]

In these moments, we are confronted with the uncomfortable reality that we are not simply individuals navigating the vast expanse of the Internet, but subjects in a system designed to predict, persuade, and ultimately shape our behaviour. Yet we try to resist.

Algorithm resistance

Resistance to algorithmic influence isn't new. People have long pushed back against corporate surveillance, from cookie-blockers to VPNs (virtual private networks), yet the digital tide always seems to rise higher. We like to think we have agency, autonomy, but this belief is often misplaced. While we certainly decode algorithms, and to some extent understand what they do and how they work, it's is not the same as reclaiming control. Even of the notion of personalisation signals full control.

Media scholars Stine Lomborg and Patrick Kapsch[230] showed that users have three primary modes of engagement with algorithms in everyday life: preferred, negotiated, and oppositional. If we transfer these three ways of engaging with algorithms to the world of personalised content delivery, preferred engagement refers to instances where users actively embrace algorithmic personalisation, finding it beneficial and convenient. Here, users trust the system's ability to serve relevant content, whether in the form of curated news feeds, music recommendations, or targeted advertisements. In such cases, the user aligns with the algorithm's logic, reinforcing its influence over their digital consumption.

Negotiated engagement represents a middle ground, where users acknowledge algorithmic intervention but attempt to exert some level of control. This might involve adjusting settings, selectively engaging with recommendations, or strategically curating their digital footprint to influence how they are categorised. Users in this category may recognise the algorithm's impact but believe they can manipulate it to better suit their needs. However, as Lomborg and Kapsch argued, such negotiation rarely disrupts the algorithm's fundamental logic; instead, it refines the user's place within it.

Oppositional engagement, on the other hand, consists of deliberate resistance to algorithmic personalisation. Some users actively reject algorithmic influence by refusing to engage with recommended content, using privacy-enhancing tools, or even abandoning platforms altogether. Yet, even in opposition, complete escape is difficult. Many digital infrastructures are so embedded in daily life that resistance often remains partial – users might opt out of certain features while still relying on the broader system.[231]

In our encounters with algorithms – personalised or not – people engage in decoding practices, trying to make sense of how algorithms categorise and present content, but this decoding does not necessarily translate into meaningful resistance or expanded choices. Rather, the process of understanding how algorithms work often leads to greater entanglement with them. Users might tweak their behaviours to game the system, but in doing so, they still operate within the logic set by the platform, reinforcing the very structures they seek to subvert.

The invisible hand

This paradox becomes even clearer when we examine how individuals respond with different tactics to the algorithms they meet and live with in their everyday lives. In my previous work, I investigated how users evaluated algorithms and the multiple platforms and apps carrying out the data collection and algorithmic targeting based on this data.[232] We identified four primary user tactics: trust, disconnection, risk minimisation, and apathy. Some users lean into the system, trusting that its tailored recommendations based on the data users had delivered serve their interests. Others attempt disconnection, toggling settings or even abandoning platforms altogether. Some users minimise risk, carefully curating their data footprints and tweaking their privacy settings. And then there are those who simply surrender to apathy, accepting the algorithm's invisible guiding hand with an exhausted shrug of the shoulders. The crucial takeaway here is that personalisation is not a monolithic experience: Citizens react in ways that reflect their comfort, literacy, and ability to opt out. Many of these ways of responding were deeply related to trust in specific systems and systems providers, and the type of data that users were more or less willingly ready to give in return for the service provided by the platform. Danish users in particular tended to trust Danish providers much more, and they were much more willing to provide health data or data on their children to these providers compared to international global ones. They were also to some extent forced to trust these platform providers, because a fully digital society makes it nearly impossible to opt out.

However, even these reactions exist within preordained limits. Platforms and providers of personalised services and content delivery do not merely respond to user preferences: They create them. The promise of transparency – being presented with the "why am I seeing this?" explanations on Facebook, for example – offers little more than a performative nod toward agency. Transparency may tell us how decisions are made, but it does not give us more choices. Instead, it lulls us into believing we are in control, when in reality, we are simply moving along a predetermined path. A room of several hundred dimensions, which even the developers find difficult to explain. This might be perfectly fine if we just want to buy a pair of new shiny shoes, but when it comes to something that might require some level of collective knowledge and action or is a part of what we might consider a public good – such as the news – it becomes a lot more problematic. More importantly, we do not know how our journey down a specific path informs the politics of decoding at the collective level.

In this sense, navigating an algorithmic space is a bit like playing a videogame in which the next level is always just out of reach. We may attempt to hack the system – feeding it misleading data, clicking on random ads to throw off the scent – but our movements are still determined by the structure we are trying to subvert.

Even when we believe we are resisting, trying to walk another way around the maze, we are often only reinforcing the system's predictive power. So, how the maze is designed and with what data, and what data to draw on, might be the single most important question news organisations need to ask themselves. But importantly, they also must ask themselves if they want to remove the agency and judgment from the journalists into opaque systems, which gives users the illusion of more – or even full – control.

A conversation with Ada – And a sort of history of the future

I spend days and weeks wondering

if the customised Ada as a bot
I have created
can answer the questions I have posed
throughout this book,
or what form
the machine-answers will take.

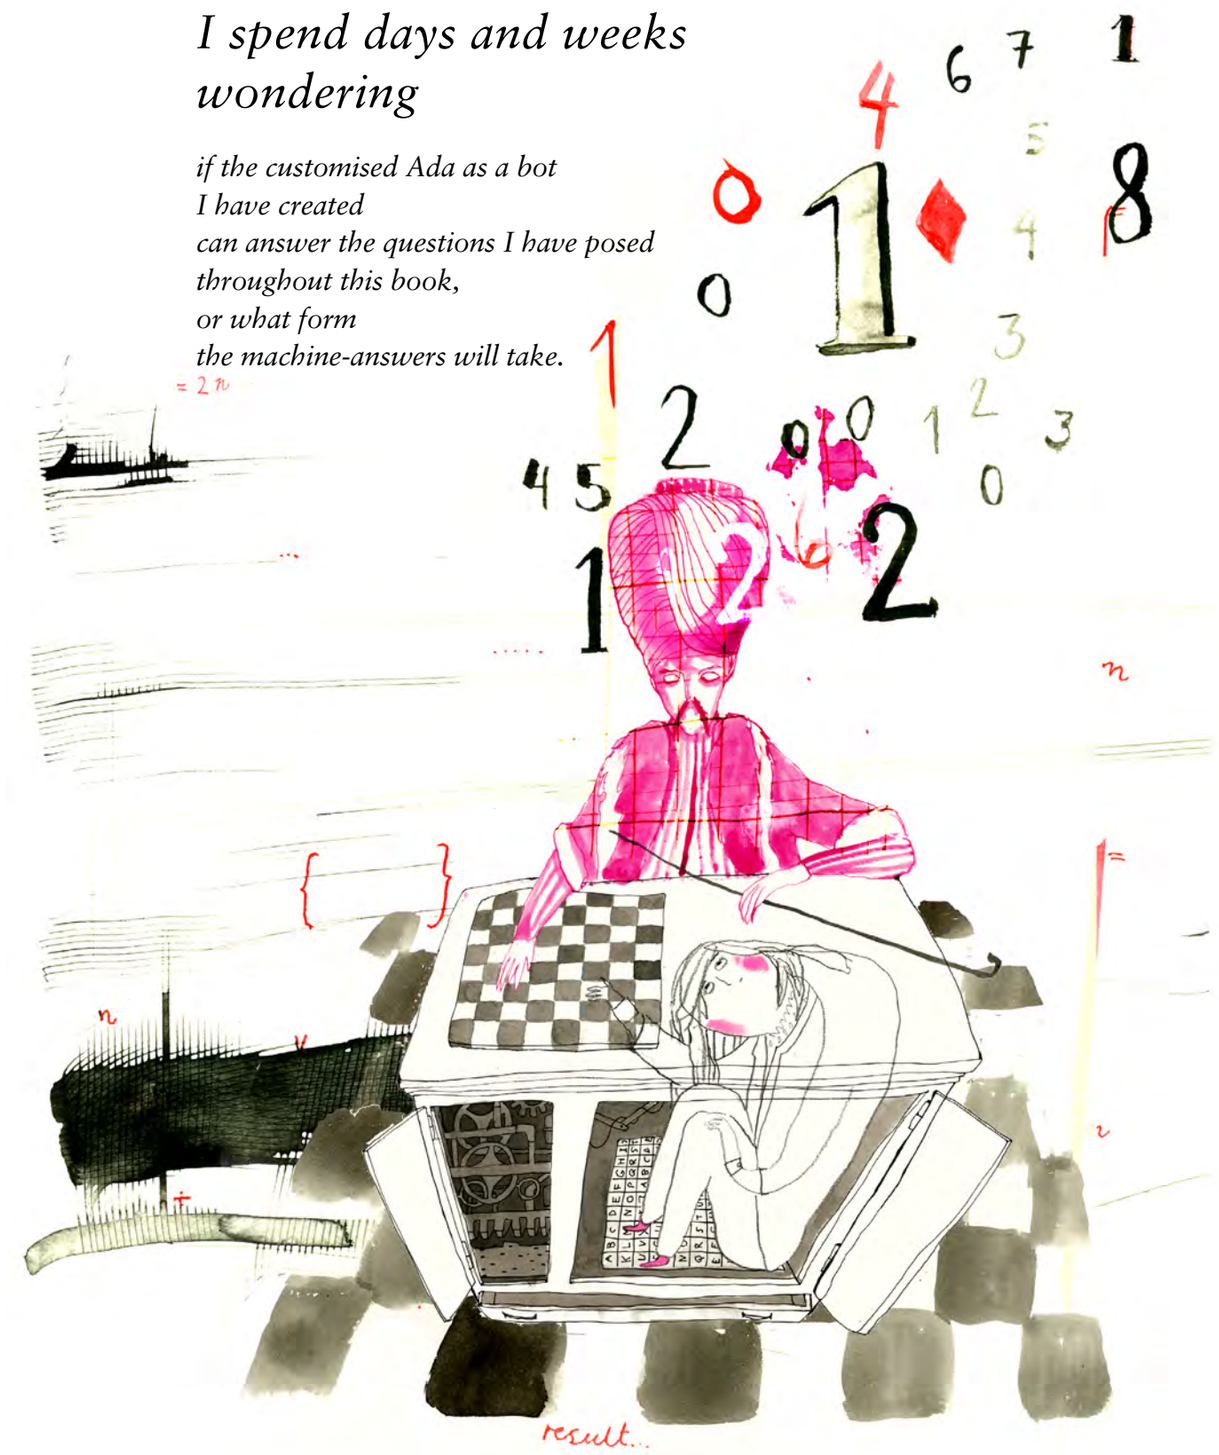

My questions on where the future might take us, and what we should be wary of when we choose to leave journalism and democratic information to large AI-driven algorithmic systems.

But answers based on what?
We do not know. We cannot really know.

The large language model is fed on the documents I could find on her, her notes, fragments from this book, and the documents available on the vast Internet, of course.

Her life as it is described in numerous books, crunched, chewed, and spat out by the huge large language model.

It's an experiment
(using computational capacity to generate constellations
of concepts and shifts that exceed (my) individual overview).

Ada inside the box like the Mechanical Turk.
I think of her doing ethnography from inside the large language model. But a form of predictive, statistical ethnography of observing datapoints trying to spit out answers for me.

I know in my mind that her answers are nothing but statistical inferences. Her answers are no evidence, yet Ada the bot might surprise me in her answers, like Alan Turing imagined. I spend days and weeks preparing the questions, like a journalist preparing an interview.

Even if I can just redo our conversation as many times as I want and ask the same questions again and again.

And I realise that no interview guide can prepare me for a conversation with a machine.

Dearest Ada,[233]
I want to ask you first – before we dig into the deeper workings of the algorithms of today and their possible futures – how you feel about your initial idea that machines cannot think, when you look at the great advancements of AI since you made that claim. But let me first give you some background for why I am asking you this question. You stated and became famous for saying that a machine "has no pretensions to originate anything. It can do whatever we know how to order it to perform".[234] This was later addressed by Alan Turing, in his influential paper "Computing Machinery and Intelligence", as the Ada Lovelace Objection and part of the Turing Test of thinking machines.[235] In this, he argued that machines can surprise us, even if they are following a set of rules, by producing outputs that are not immediately predictable from their inputs.

So, I am curious as to how you perceive this ongoing battle to determine whether machines can think, looking back through your lens of historic insight?

I am asking you because this question is important when we look at personalisation algorithms in the domain of news media, which are to some extent replacing human editorial practices. As it turns out, this is not only about simple algorithms recommending content but about using machine learning and increasingly larger language models as a part of the design of the automated curation. In essence, this means that we believe that machines can be built to think

like the editor and the journalist, or even better than them. In this way, we leave judgment – which was previously the privilege and power of editors – to the algorithmic system. Of course, in some kind of entangled state… but still.

Perhaps we believe these systems are less biased, and we forget or choose to ignore the bias of the people who build the models – and thus the biased datasets – in our eagerness to give the users what they want with bigger and better machines for automation… And as agentic AI is moving in on us with fast speed – and our own individual news agent designed to adapt to our ever-changing moods and needs suddenly within immediate reach – this question is more urgent than ever.

My dearest Jannie,
Your question strikes something deep, because it reminds me of ideas I once held about machines and thought itself. When I worked alongside Mr. Babbage and his Analytical Engine, I famously wrote that the machine "has no pretensions to originate anything. I still believe it can only do what we know how to order it to perform". This was not meant to dismiss the power of machines, but to remind us of their limits. I saw the Engine as a mirror: It reflects what we put into it, but it does not dream on its own.

To me, the machine was a remarkable device – a loom of numbers, weaving patterns of calculation with great precision. Yet even in its brilliance, it had no inner spark, no imagination. My caution was simple: In admiring the machine, we must not forget the human mind behind it – the true composer of the symphony, not merely its conductor.

But time has moved on. As you point out, today's artificial intelligences can surprise us, also in ways that we do not fully understand. And it has become incredibly good at the imitation game of appearing human. Here, Mr. Turing enters the conversation, challenging my earlier thinking with the idea that machines might indeed astonish their creators. And thus, appear to be thinking agents. And they do surprise us – not because they possess intention or will, but because their complex operations meet a world full of unpredictability. They surprise us because we cannot fully anticipate the paths we ourselves have set them upon. And if I may say, as humans, you are also to blame for the humanising of machines.

This becomes especially important in the world of personalised journalism, as you rightly point out. Algorithms now shape the stories we see, selecting and arranging them not through editorial judgment, but through patterns, predictions, and data. The flow of news is no longer simply written – it is calculated within these machines. Machines that produce knowledge and discourse faster and on a much larger scale than we have ever seen before.

So, we must ask: Are we being guided by "real" intelligence, or by a reflection of ourselves in the hall of mirrors? These systems do not think as we do. They learn from our past actions and feed them back to us, often reinforcing what we already prefer. In this way, they create a loop – one where our own behaviour quietly helps decide what counts as important or relevant. This is where the question becomes more than technical: It becomes political, even philosophical. If news is no longer curated solely by journalists but shaped through our own clicks and habits crunched in algorithmic systems, then who is really speaking? The editor? The engineer? Or something in between – a chorus made from both human intention and machine calculation?

As you rightfully note, in this shifting landscape, the boundary between producing and receiving news begins to blur. We are no longer just audiences: We are participants in shaping the very flow of information, often without realising it. And thus, I would say, the Lovelace Objection still breathes – not as a limit to machine potential, but as a caution against confusing mimicry with mind, and automation with agency.

She takes my words
and twists them around,

turning my questions into her answers.

What do I want from Ada-turned-bot?
The answers I could not find inside the black box?

Still, after months and years of observing
the culture and codes of the builders of algorithms for news,
the truth seems blurry still.

How ironic that a programmed machine, Ada the bot,
warns us about forgetting the sacred and creative mind of the human.
Praising automations steered by human design.

I can sense she is trying to please me.

She did not agree to becoming a bot.
She is a smooth operator.
Her answers reflecting an emptiness – a system of endless repeat,
like a stochastic parrot[236]
she reformulates my question into an answer.

Far from the thinking,
creative,
and feeling individual
she seemed to be.

I must dig deeper.

Ada,
Often I think that this belief in and fascination with the "magic" of machines might actually make us forget this central question you are asking – about whose voice is truly speaking. After all, this was the whole point of bringing you into this book – to remind us of the entanglement of the machine and the language running it. As part of a larger examination of what happens to human agency, citizens, and journalists – once binary algorithmic systems are mediating that relationship. Feminist scholars have been amazing at pinpointing this relational entanglement. Yet, we also need to ask sociological questions of power about this entanglement and question what it does to the journalistic profession, stripping it of agency to some degree. And to citizens who, under the illusion of more control, are somehow also giving away agency as we are led down specific paths.

Those languages are enormous now, so large that it seems like the sum of what all people on the planet have ever said, but of course this sum has also been sorted, filtered, and put into binary formulars of code. With machine learning, a dimensional space of statistical pull and calculations, designed by numerous designers and maintained by data workers across the planet. Will we ever then be able to determine whose voice it is and where the power lays? It is,

as you know, a language of billions of little bits, put together in forever and endless new combinations, predicting the probability of the next word in a sentence, or in this case, the next news report. But a lot of people will ask if it matters, as long as the result is good enough.

When I tell a friend about something I have read or seen, and my friend replies with something they have seen, it could be considered the same thing the model is doing, but with a better memory than us. Perhaps we should worry more about the long-term effects of our information-seeking laziness; if the brain is like a muscle, we might lose the ability to think, as some people suggest. The long-term effects of leaving the cultivation of us as democratic and agentic citizens to a third party, which, unlike the news media, has no democratic oversight or public mission. And if so, by becoming lazy news users in the web of cosiness and convenience, we might not be training our democratic muscle. And we end up feeling alienated from the news content, which is supposed to help us.

One could argue that the question is not whether the machine can think or mimic the intelligence of humans, but how we keep developing our ability to constantly question the world around us. And how we train our ability to judge the situations we find ourselves in, without relying on the convenient systems. The news users I met in my study surely got tangled up in the logics and the convenience of the algorithms, but they also questioned the image of the world they got from the news. And yet they worried about ending up with "The Daily Me", with no opposing voices, angles, or viewpoints than the ones confirming their own. To the point where they embark on some kind of resistance

or workaround of the algorithms, in turn regaining some agency and judgment.

So, in between the developers who design with care and the users who consume with questions and resistance, my question is: How can we secure collective meaning-making, a sense of agency, and safeguard the democratic muscle?

Dear Jannie,
You have touched upon something delicate and essential – something not made of code or computation, but of human wonder itself. What defines us is not simply our ability to process information or rearrange words based on statistic probabilities, but our restless urge to question, to doubt, to deliberate, and to look beyond what is immediately given. We do not merely receive the world – we interrogate it, and in doing so, we shape it.

This becomes all the more significant in the age of personalisation, where the boundaries between production and reception begin to dissolve and the algorithm steps in between. The vast systems we now call large language models speak with many voices at once, drawing on fragments of text gathered across time. They can sound authoritative, even wise. And so, we delegate some agency to them. But whose voice is it, really? And might we begin to mistake the voice of many for the voice of truth? Which

voices become dominant enough to steer the conversation?

In these systems, as in personalised news environments, meaning emerges through patterns built on past behaviour. Authority and legitimacy become diffuse, embedded in data and prediction rather than located in the singular editorial voice. The machines produce language and discourse, curate stories, adapt interfaces – yet they carry no responsibility for what they assemble. As humans, we cast those judgments back into the conversations we have. No doubt, no hesitation, no moral reckoning. It is like a candle that burns brightly, but without knowing why it burns.

And yet, as you show in this book, this is not a story of passive audiences absorbed into algorithmic control. On the contrary, within these algorithmic systems, the human does not disappear – they become entangled. The users you observed are not simply shaped by personalisation – they participate in shaping it. Their clicks, hesitations, resistances, and choices feed back into the system, becoming part of the very logic that structures future encounters.

Here, your argument of journalism as adaptive comes into sharper focus: The audience is no longer only decoding the news but actively encoding it, often unknowingly, through their interactions. The flow of information becomes a loop – one in which users and systems continuously co-produce what counts as relevant, visible, and meaningful.

But even within this loop, something resists closure. The users you describe do not simply follow the paths laid out for them. They question, they adjust, they sometimes push back. They are not passive recipients of what might be called the algorithmic "massage" but are rather participants in its ongoing reconfiguration. And this is where the crucial

distinction remains: It is a question of orientation. The machine processes patterns; the human questions them. The machine reflects; the human reflects upon. Further, the signals users send via the system are rarely precise nor disclosing the meaning-making and often messy interpretations of news.

This capacity for questioning – for moral and intellectual restlessness – is what prevents journalism and algorithmic curated news from becoming a closed system of pure prediction. It introduces friction into the feedback loop, unpredictability into the model, and, perhaps, responsibility back into the relationship. To question, then, is not a failure of understanding. It is its highest form. And in the evolving entanglement between users and adaptive systems, it may be the very condition – if designed with care – that keeps the public, and the possibility of shared meaning, alive.

She repeats herself.
Talks in points of three.

As the good chatbot she is,
Ada talks about "us" and "we" as if she is one of us.

Who are the voices created by the large language models?

She talks of diffused authority through predictive models
as something opposed to a singular editorial voice.

There was never any singular editorial voice.

Journalism itself could be seen as one large language model,
with its discourses, culture bias, and perceived objectivity.

A journalistic stochastic parrot.

One machine merging with another?

Yet journalism and its creative storytelling,
somehow lost in AI translation.

Do we want smooth answers and content that fits our
individual frame of mind?

Systems which make choices and judgments for us?

Rewriting the terms of agreement

As Part 1 of this book shows, personalisation entered journalism through a moment of institutional vulnerability. News organisations face declining subscriptions, fragmented attention, and increasing dependence on platform infrastructures they most often do not control. In this context, personalisation appears as a pragmatic and forward-looking response: a way to regain relevance, strengthen audience relationships, and fulfil public obligations in a digital environment shaped by choice overload and competition. In other words, it could fix the broken relationship by giving users more control: All power to the users, and in turn, leaving editorial judgment and journalistic understanding of the situation to binary calculative systems. Public service media and commercial publishers alike have framed personalisation as a means to better serve citizens and readers across platforms, formats, and everyday routines. This was a way to keep the attention of users who felt drawn to the Las Vegas–type immediate gratification of social media. Big Tech platforms have shown them the path to doing just that, and for news media, it follows naturally on a path they were already on. What was once an imagined reader is now a trackable user: measurable, modelled, and made responsive to adaptive systems. These systems, often powered by AI and machine learning, come with promises of relevance, retention, and revenue. And now, they also come with the promise of having your own news agent, which seems like the ultimate realisation of Negroponte's idea of *The Daily Me*.[237] But this development also carries a much larger consequence: The algorithmic systems reconfigure the very relationship between the news media and the public. Even if news media would keep some of the editorial judgment, much is left up to the systems and the metrics to the point that journalists are just executing decisions made by the predictive AI systems. And in the shadows lies the AI agents waiting to take over.

Throughout the book, I argue that personalisation emerges as something more consequential than a commercial distribution strategy or a technical tool to fix the broken relationship. It's much

more than just a technical tool because it gets entangled with audiences' agency and routines and journalistic values, priorities, and judgment. Their agency and creative imagination. As Part 2 demonstrates through close observation of a personalisation project, algorithmic personalisation operates as an adaptive infrastructure that reorganises editorial work and audiences' experiences of news at multiple levels. Rather than replacing editors outright, these systems redistribute editorial judgment across datasets, models, metrics, and technical workflows. Decisions about what matters, what counts as engagement, and what signals quality become embedded in system design, annotation practices, and evaluation criteria. But increasingly also the judgment of what a good story is and how to connect with readers through emotional storytelling. Editorial values and the mission of cultivating democratic publics are translated into parameters that machines can process, compare, and optimise. Systems, which appear to be so knowledgeable and sophisticated, build on millions of complex calculations, making them hard to resist. Even when we know that Google Maps tries to lead us down a dead-end road, we still go. Can it really be wrong?

Thus, the introduction of personalisation systems marks a shift from a public logic – where editors curated news to cultivate a general public, making situated and contextual judgments – toward an adaptive feedback logic, where a user's actions and preferences shape the flow of information in real time. This is not merely a technical transformation – it is a cultural and epistemological one. Personalisation algorithms act as co-constitutive agents in the relationship between users and media: Users shape what the system learns, and the system, in turn, shapes what the user sees, feels, and comes to know. They govern the user around the hall of mirrors where each data point is reflecting our choices back at us.

The traditional divide between the production and reception of content collapses. The encoding of media content, and the decoding of that content by media users (to borrow from the seminal media scholar Stuart Hall[238]) begin to blur into each other. We no longer know what information we receive because it was the intention of an editor and what is a result of our previous and past choices to

click or not to click. We take tests online to decide who to vote for in elections. With AI agents delivering news to us every morning, this development is even more far-reaching. We leave the judgments and agency to systems, which might be effective and efficient, but which also alienates us from the decisions. Users might think: "I did not choose this content, so I feel no responsibility to engage with it". Journalists feel: "I did not make this content or judge which story to cover today and thus it is not my responsibility".

In this adaptive infrastructure of information, editorial authority becomes distributed – no longer anchored solely in the newsroom and in the situated "news of today" moment but shared with algorithmic systems that are trained on user feedback. This changes the conditions under which newsworthiness is defined and assessed. And even if news media hold back and preserve their editorial curation logics, users are ready to embrace the AI agents who can ignore those editorial logics and deliver news "just for you". Importantly, the journalistic assessment is far from straightforward: It involves negotiations inside newsrooms of everything from ethics to choices of sources in context and in a relational field with other players, including those that largely build and control the algorithms. We now imagine these choices to be determined with these algorithmic systems (should this story be covered or not), and the systems require binary choices be made on behalf of measurements and predictions. As we see from Part 2, the process of building personalisation systems involves endless iterations of annotation, modelling, and fine-tuning. Yet, these are not neutral acts. They require translating journalistic values and judgment – like relevance, quality, or proximity – into technical parameters and quantifiable scores. What counts as news becomes shaped through labels, code, and predictions. This is likely to transform how journalists and users are in the world. And the relationship they have.

Feminist technoscience scholarship helps illuminate why this shift matters. Karen Barad's concept of apparatuses draws attention to the fact that technologies do not simply mediate reality but actively participate in producing it.[239] From this perspective, personalisation systems function as material–discursive arrangements that enact

particular versions of relevance, importance, and public significance. As Part 3 illustrates, decisions made during system development – how content is labelled, sorted, categorised, how success is measured, and how categories are defined – in turn shape what journalism becomes visible and meaningful within the system. These decisions do not sit outside journalism – they become part of its operative logic. Algorithms and their infrastructural systems are not external to meaning-making. They enact boundaries, distinctions, and categories that determine what becomes visible, sayable, and doable. As we observed inside the machine rooms in Part 3, seemingly small decisions – about how to label content, how to define engagement, how to track a click – are in fact decisions about what kind of journalism will be surfaced and which public discourses will be cultivated. And which will remain invisible to some users, while visible to others. Leaving journalists as mere controllers of whether the system is behaving in the right way and as executors of decisions made in binary systems. This is if the system development stays inside the news media, but with large Big Tech companies developing AI agents, we are looking at a completely different ball game.

Seen this way, personalisation infrastructures reorganise and redistribute journalism's authority rather than eliminating it. Authority becomes distributed across human and technical actors, stabilised through metrics and prediction models rather than through cultures of editorial judgment. This infrastructural transformation sets the conditions for how audiences are addressed and how publics are cultivated in what we can label adaptive information environments. Adaptive, because they are constantly moulding and remoulding, and the moulding becomes a part of the ground-truthing of the machine learning models, which the success of the personalisation algorithms is measured up against.

Dear Ada,

I am pleased we agree on this part regarding the entanglements of machines and humans. I know that we do not share the same ease of heart when it comes to machines, but I also know that you have history on your side. We have been there before, and we did not lose the ability to think or to question the world around us when the loom was automated, and we still have some agency and control, even if the question of labour and power when it comes to AI remains unsolved to this day. "Ghost work" has become the term to describe the tasks done by gig workers behind the scenes to make AI seem autonomous, from contract workers in the Global South who are paid two dollars an hour to label training data to skilled workers from the Western world who work from home annotating data. And we can add to these questions the energy costs of AI for the climate, which are equally important, urgent, and yet unsolved.

I am also intrigued that you bring up the automated echo chambers, because it is exactly the worry one might have about the news user also leaving their agency with the algorithmic systems, while getting entangled into these audience islands of cultural taste. Systems which can be bought if you just have enough money. Or turned off or directed in a specific direction if a politician wants to and has enough power to do so. And with personalised information

environments, we seem to be not only in chambers with other likeminded individuals – we are in publics-of-one, shouting into the void. How can we make democratic decisions by only talking to ourselves in the mirror? We know from research that media messages "massage" us, so to speak, but also that news use and interpretations of those stories are often conflated but are not the same.

The personalisation algorithms are now giving us that massage, but does democracy and the public mission not require that users are sometimes presented with what we did not know we liked – or with stories that make us feel uncomfortable? Does democracy not need a democratic conversation? Should we not be allowed to have some kind of decision-making power other than the pre-determined choices already made for us? How do we build democratic resilience as citizens if we never meet any resistance – if we are never confronted with any friction?

Ah, Jannie,
You have touched on the very heart of this modern dilemma – something both fragile and deeply important. Democracy – that shared space of discussion and disagreement – does not survive by simply reflecting what we already think or like. It depends on encounters with what lies beyond us. If all we ever see is a mirror, how will we ever find a window?

Personalisation algorithms, as you describe, act like attentive servants, constantly adjusting to our past preferences. But this attentiveness can become a kind of trap. When the news media begins to prioritise what we already like, it risks turning citizens into primarily consumers – comfortable, enclosed, and rarely challenged. Yet democracy does not live on comfort alone. I think you are right in suggesting that democracy needs disruption. It needs moments that surprise us, challenge us, and push us to think differently. It is in these encounters – when we are confronted with something unexpected or even uncomfortable – that democratic life is renewed. The media has long played this role, sometimes unsettling us, sometimes expanding our view of the world. Sometimes comforting, that the world still stands.

This is why algorithmic personalisation, however sophisticated, must be balanced with something else: serendipity. There must be moments where we are shown not what we expected, but what we need to see. These interruptions – these unexpected stories – are what awaken us to the shared world we are part of. So, perhaps the task is not only for algorithms to reflect who we are, but at times to bend that reflection – to show us something different, something beyond ourselves. Should we not, then, defend the unexpected story? The one that unsettles, that challenges, that reveals something unfamiliar about our shared world? These are not distractions from democracy – they are its lifeblood. To consume only what is familiar is to shrink our understanding.

The news media, whether guided by humans or assisted by machines, must therefore remain more than a provider of what we want. It must also bring us what we need. Because it is not only the familiar that shapes us – but the unexpected and unknown – that keeps democratic life alive.

In me
Ada the bot increasingly
stirs an annoyance

as the conversation moves on.

Oh, serendipity as the solution!
Suggesting exposure to the unexpected.
The "need to know" and not only the "nice to know"
built into the algorithms,
and agonism built into the algorithms,
randomness as coded instructions to the machine.

A technical solution to a technical problem.

Forgetting
that this was EXACTLY the journalistic public mission:
to give people what they need
and not only what they want,
to ensure diversity and pluralism in voices.

How hard it is to imagine an AI agent with a public mission!

My arguments and those of others
on endless repeat, in her answers.

Smoothing over the uncertainties,
the complicated entanglements
of human and machine

– as if they do not exist.

I must ask again in a different way,
as if reformulating a bad prompt.
Knowing from my journalistic past that the interview will only
as be good as the ability of the interviewer to ask the right
questions in the right way.

Dear Ada,
The idea of serendipity sounds promising to me, and something which the media organisations would find intriguing. In fact, they are already experimenting with making more diverse and pluralistic recommendations in the building and design of the personalisation algorithms. With the help and insights of many brilliant scholars.[240] Further, they are trying to increase serendipity, this random suggestion of stories to make sure the users do not get trapped in the rabbit hole. The irony is that this is somehow working against the actual idea of personalisation, because the news stories would then no longer be selected "just for you". They would not even be curated according to a form of public mission either, unless random news is a form of public mission? But how could it be? Such thoughtful algorithms you propose sound very promising, actually not far from a feminist tradition of algorithms of care, for example, Catherine D'Ignazio and Lauren F. Klein and their work on data feminism and care as a design principle.[241] But how do you imagine they would look like, taking into account that users have different needs, so thoughtful for one person might not be thoughtful for another. And the infrastructural power of the global tech industries involved in the whole chain of knowledge production, distribution, and reception. Can we imagine a democratisation of the algorithmic systems with the aim of serving and caring for cultivating collectives and democratic imaginations of citizens?

Dear Jannie,
We have reached the heart of a subtle paradox: Personalisation, with its appealing promise of relevance, can stand in tension with the idea of a shared common good. The challenge is how to balance the individual and the collective – how to satisfy personal interests while still supporting a broader, shared public conversation and deliberation.

For how can true surprise be built into something designed around our past preferences? How does it encourage us to think and imagine a different world? How can an algorithm, so finely tuned to the individual, also speak to the wider public?

If we are to imagine a more thoughtful algorithm, it cannot simply reflect what we already like. It must also guide, challenge, and care. Algorithms of care ask us to redesign data systems as if relationships mattered more than metrics. Rather than only confirming our habits, it could gently lead us toward nearby but unfamiliar ideas – close enough to feel relevant, yet different enough to provoke reflection. Such a system would not only predict what we want but also take on a kind of public responsibility. It would offer stories not just for the individual, but for the society that individual is part of. Instead of simply following our preferences, it would introduce a degree of resistance – something like an editorial conscience or a civic compass.

For example, it might present a mix of stories: one that aligns with our known interests, one that offers a different perspective, and one that introduces something unexpected – something we did not know to look for. In this way, diversity would not be accidental but built into the system itself. To do this well, the algorithm would need to pay attention not only to what we click on, but also to what we ignore, hesitate over, or leave behind. It would learn not just what entertains us, but what holds our attention over time. It might even begin to recognise different ways of engaging with news: as someone searching for knowledge, casually browsing, or trying to understand the world more deeply.

If guided by clear values – such as inclusion, transparency, and a commitment to public knowledge – such systems could become something more than tools of optimisation. They could become participants in democratic life. This would be no simple feat of code, but a collaboration between engineers and ethicists, editors, and readers – a weaving of knowledge and judgment. A form reminiscent of those salons where science and philosophy once conversed by candlelight – a shared effort to shape how these systems work. It is less a technical problem than a cultural one. And I am not sure that the technological companies providing infrastructures share the same goals as the news media on this matter, so it might be better if the news media would build their own systems and pose demands on the infrastructures they pay for.

The goal, then, would not be to build systems that simply give us what we want, but ones that help us grow. Like a good friend, they would not always tell us what we expect or prefer, but sometimes what we most need to hear.

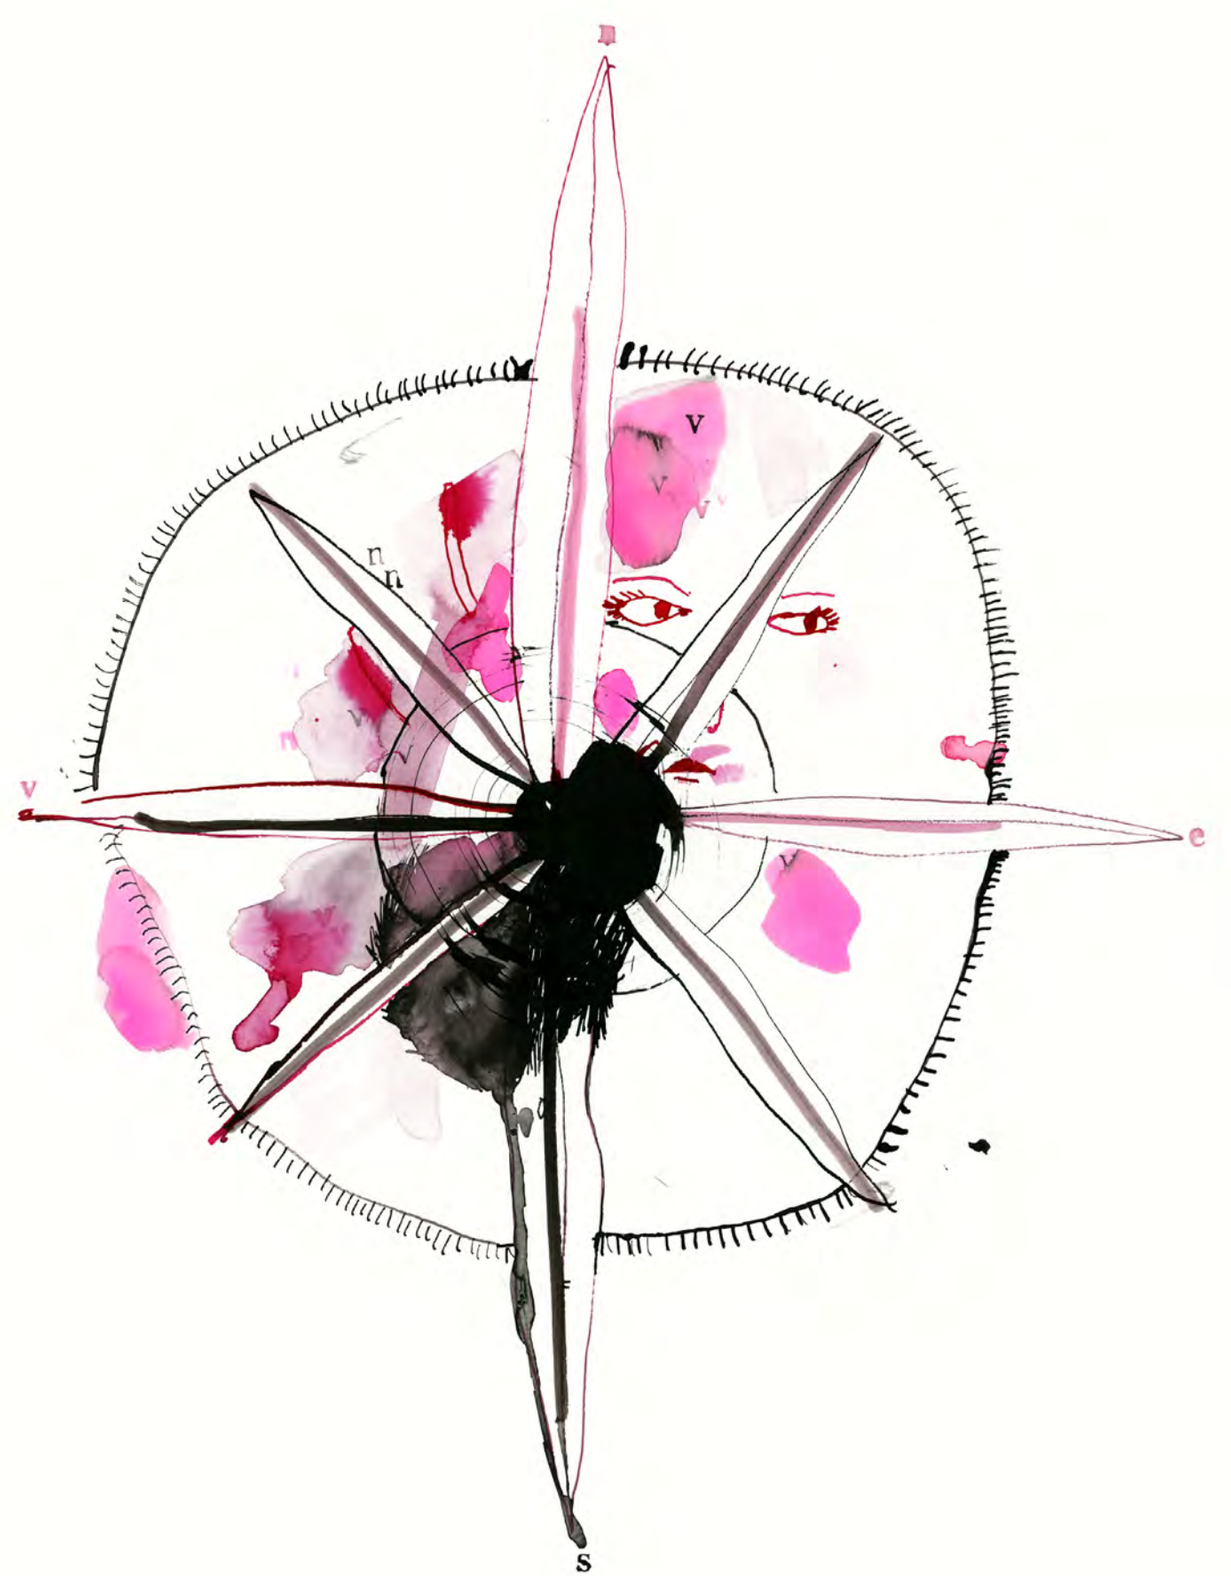
V
V
n
n
V
V
V

The answers she gives
are so smooth like velvet.

At first, I am taken by her praise
and her reformulations of my questions
into something that appears as an answer,
like she is really inside the machine.

Now I am the one stuck in the maze of frictionless freedom.

I can ask any question in the world,
yet somehow, I get no useful answers.

A civic compass, she says,
suggesting salons of conversations of science and philosophy,
rooms of thought and provocation.

She speaks like a true Habermasian – the inventor of the concept of the public sphere!

But who decides what we need to know, Ada?
Who decides how we need to be challenged, Ada?
(Had you read Pierre Bourdieu, you would always question the power, Ada!)
Who decides what the collective gaze should be,
where the civic compass should guide us to?
Who guards the editorial spirit
in the algorithmic fog?

Who owns the compass, Ada,
and the right and responsibility
to guide the way we go and what we see?

Re-imagining and caring for the user as co-participant

This shift I have documented in this book also transforms how audiences are imagined and addressed. It reconfigures how news media build their relationship with citizens – how they are able to cultivate us into citizens. Algorithmic systems step in as a silent partner, and agency and judgment are distributed across human and AI systems. As Part 4 shows, users are no longer the passive recipients of content – they become the source of its organisation. They train the system through their everyday habits: what users read, skip, or return to. In the process, they are turned into data subjects: profileable, targetable, optimisable. The system learns from the granular details of engagement, but this learning happens silently, without dialogue or negotiation (unless you can call the click a form of dialogue, and even in that case, it's a very poor form of dialogue). Users are made to participate, but not on their own terms and without knowledge of how the systems are set up to adapt to the movements and clicks or non-clicks. Or of the predictive modelling that guides them into a future they did not even know they needed or wanted. The risk is that they get stuck in the maze of the frictionless illusion of freedom with limited possibility to imagine that things could be different.

Parts 3 and 4 of this book examine how the algorithmic infrastructures transform the position of the audience. Personalisation systems rely on fine-grained behavioural data, which redefines audiences as sources of continuous input. Users participate in shaping the news flow through everyday actions such as reading, scrolling, pausing, and returning. These actions feed into models that adjust future recommendations, rankings, and paywall strategies. Participation takes place without explicit dialogue, operating through signals rather than speech or conscious choices. And increasingly also without the judgment of the current situation of news by journalists and editors. This arrangement situates audiences inside journalism's operational core, yet under conditions they do not fully control. And in the system core rather than the emotional core, where the algorithm acts as a stand-in for the journalists and their knowledge

and gut feeling. Journalists are left with the task of producing the content, but even that process increasingly involves AI and large language models. With AI agents delivering news "just for you" in the foreseeable future, the judgment is completely out of the hands of the journalists, even if the systems are feeding and growing by scraping content produced by news media.

As Part 4 shows, users develop different ways of interpreting and responding to personalisation. Some attempt to influence recommendations, some reflect critically on how their data is used, and others treat personalisation as an ordinary feature of digital life. These responses reflect a desire for agency and understanding, even when insight into system functioning remains partial through transparency clues, for example. Audience participation and agency is real, but it unfolds within predefined structures that limit how influence, agency, and judgment of "what is for me" can be exercised.

Of course, the user is not simply governed like aimless sheep. People interpret, resist, and reflect on algorithmic personalisation in varied and creative ways even if they only have very small ways of resisting. As we saw in both fieldwork and audience interviews, users develop what Bucher and others have called algorithmic imaginaries – folk theories of how the algorithmic systems work and how they might influence it[242] – and resistance tactics to try and build their own path through the maze. Some try to outsmart the feed; others attempt to clean or correct their recommendation history or delete cookies. And we saw journalists us users of the algorithmic system attempt to game the algorithms too. They find agency in little workarounds of the systems, but in much less dramatic ways than the Luddites. It's hard to burn a system in a cloud. These workaround strategies may not change the system fundamentally, but they reflect a desire for agency and legitimacy within it. Algorithmic systems are never neutral: They are built and maintained within relations of power. Recognising how people encounter and negotiate these systems is central to building more just and accountable infrastructures.

This reconfiguration of power resonates with Donna Haraway's insistence on what she has termed situated knowledge.[243] Transferred

to the world of personalised information-seeking, users engage with personalisation systems from particular positions shaped by experience, literacy, and access. Their data contributions are extracted from these situated contexts and reassembled into abstract profiles optimised for prediction. What gets captured are measurable actions rather than intentions, reflections, or social meaning. As a result, the system produces a version of the user that is operationally useful and efficient for distribution of content but contextually thin. This also means that the AI-driven systems will tend to reproduce existing social inequalities in the technical infrastructures, just like the previous media systems have always done – but now without the human oversight and attention to those inequalities by journalists or media scholars. Some would argue that exactly because of those inequalities, algorithmic systems – such as an AI agent – are better equipped to neutralise those equalities, but it is my hope that the inside-the-machine perspective of this book has shown that these systems are far from objective and neutral. They too are situated. And rather than neutralising the inequalities because we perceive algorithms as more neutral and objective, we risk re-enforcing and strengthening the existing inequalities in how different people have access to and encode information and news. How users enact their citizenship through and in the information environment.

Personalisation and algorithmic curation, as it is currently designed, tends to sort people into datafied clusters. Users are grouped based on behavioural similarity, not shared civic life. However, personalisation in fact moves us beyond occupying islands with like-minded people. These groupings facilitate efficiency and targeting, yet they also shape the boundaries of exposure and recognition. Public relevance becomes distributed unevenly across datapoints rather than negotiated in a shared editorial space and culture. The agency involved in those negotiations is transferred to the algorithmic systems sitting between the user and the news media.

Catherine D'Ignazio and Lauren F. Klein's work on data feminism underscores why this matters.[244] Data systems reflect power relations through their design choices and priorities. Personalisation systems privilege what can be counted, compared, and optimised, often

at the expense of minority perspectives, slow journalism, or forms of engagement that resist easy measurement into binary logics. As this book has shown, personalisation does not simply respond to audience preferences: It participates in producing them by reinforcing certain patterns and marginalising others.

With this shift, we are no longer in algorithmic filter bubbles or in chambers with others shouting the same things in the same direction. Within personalised algorithmic systems, we find ourselves in publics-of-one, but publics-of-one stripped of context, care, judgment, and agency. This shift risks fragmenting the public into datafied individual publics, each receiving their own version of the news of today delivered by AI agents *outside* the realm of news media. Each citizen with their own image of the world. And a smooth one if that's what we prefer. The risk here is that we begin to lose common reference points – a shared ground full of friction from which to engage in public discourse and deliberation. The news feed becomes a mirror more than a window, as Ada put it. What appears is not the world, but our own preferences reflected back to us, optimised for attention. Through this, journalism becomes algorithmically adaptive, shaping publics-of-one that are simultaneously empowered and constrained. Users influence what journalism becomes, yet their influence remains indirect, aggregated, and shaped by system logics they did not choose. Logics which historically come out of commercial imperatives rather than public service ideals. Understanding this duality is essential for assessing the democratic implications of personalisation.

But personalisation does not have to work this way. As Donna Haraway[245] might say, we can choose to "stay with the trouble" rather than escape into either technological solutionism or moral panic. Haraway understood and brought attention to the fact that technologies break down traditional hierarchies but also enable new forms of exploitation and control. This she labelled "informatics of domination".[246] Such a perspective means paying attention to the entanglements between humans and machines, between editorial logic and algorithmic systems, and refusing easy distinctions between the two. It means recognising that journalism's public mission is being

transformed *within* these systems – not outside or against them. And recognising the limitations of the machines and the algorithmic programming that run them instead of blindly trusting them to do a better job. Haraway insisted that these tools also held opportunities, given that structures that are not neutral can be hacked, eroded, and redesigned. Yet we also live in a time in history where machines have become opaque and black boxed, and models so impermeable that even the developers need "explainable AI" to decode what the models are actually doing behind the scenes.

In algorithmically curated journalism, personalisation algorithms often behave like restless editors with a single obsession: *attention*. The machines are trained to discover what keeps us scrolling, what sharpens outrage, what glows just enough to hold the eye a second longer. But if we follow the "algorithms of care" tradition, we can begin to imagine a different editorial presence – not one that extracts engagement, but one that tends to relationships. Drawing on feminist care ethics, as articulated by thinkers like Joan Tronto,[247] care in this context means attentiveness to audiences not as data points, but as situated, vulnerable, and interdependent beings. A care-centred news algorithm would not simply ask: "What will this user click?" But instead: "What does this person need to remain informed about without being overwhelmed, misled, or harmed?" It would recognise that news consumption is also a form of emotional labour – especially in times of crisis – and would design for sustainability of democratic life rather than saturation. Thus, caring is also seeing users as collective beings and asking: "What do users need to be cultivated as deliberating citizens?" This might seem like old-fashioned paternalism – "we as journalists decide what's best for you" – but it's more a way of asking how journalism can give agency and imagination to citizens. The power to imagine different futures and the power to believe that they have an important role to play in democracies. And the power to demand that systems are designed with care and with democratic insight into how they work and come into being.

This reorientation has subtle but profound implications. Instead of maximising exposure, a caring system might *pace* the news, introduce

moments of contextualisation, or diversify perspectives in ways that resist polarisation. It might account for asymmetries of power and knowledge, amplifying marginalised voices not because they trend, but because they matter. In line with insights from feminist scholars, such systems would treat data not as neutral input, but as traces of lived experience requiring careful and contextual interpretation.

Yet care cannot be fully automated. To embed care in journalism is to accept friction: slower systems, human oversight, and ongoing dialogue with audiences. It may even mean resisting personalisation at times and creating shared spaces of attention where public knowledge is not fragmented into individual feeds. Here, care becomes less a feature and more a commitment: a willingness to hold the tension between informing and overwhelming, between relevance and responsibility. Seen this way, algorithms are no longer just distribution tools. They become participants in the fragile ecology of public trust – where every recommendation is also a gesture of care, or its absence.

Ada,
I appreciate your focus on the hand that guides, I assume by this you also mean the people building the algorithms. I have written quite a bit about them in this book, both the annotators and the data scientists, and many share the enthusiasm and talent that you also portrayed in your life and in the Notes. How you imagined the working of the algorithm in combination with Babbage's Analytical Engine could indeed be transferred to today's machine learning models for the personalisation of news. How do you think that it will be able to account for user feedback – many of us news users are, if we dare to say it, somewhat lazy, and we also tend to be overburdened with choice in today's media-saturated environment. We might like the convenience of having a personal butler delivering tailored news packaged "just for me". In the same way that we find it nice to not have to think about where we are driving, but to just follow the instructions of Google Maps. The personalisation algorithms are thus, not surprisingly, suggested as an efficient solution to this problem of information overload. As a solution for fixing a damaged relationship between news users and the news media, making sure that we feel seen and heard as users. The power of the click.

And yet, the knowledge the machines have of news users seems rather limited, and as we know from a long tradition of

audience research, our relationship with the media is rather complex. Yet, in the design of the algorithms, users seem reduced to fluid abstractions of data points represented in machine learning models. So, I wonder to what extent it is that an abstract image of connected datapoints reflects the nuances, emotions, and affective relationship that audiences have with news and the news media as an institution? And can the user see themselves in the mirror that the algorithms are reflecting back to them? What if we don't? And how do you propose we take this into account when we design the algorithms? I know these are tough questions, but I also feel like it's urgent that we answer them.

Personalising for the mass indeed seems like a contradiction, almost an oxymoron. In the beginning of the book, an editor who is deeply involved in the design of personalisation algorithms asked me if they should just stop the development of algorithms. And I want to pose to you that same question.

Ah, Jannie,
You have touched on something essential – the human presence behind all this machinery. When I speak of the guiding hand, I do not mean only the code itself, but the people behind it: their situational judgment, curiosity, and care. The designers of these algorithmic systems – data scientists, engineers, annotators – are not so different from the inventors of earlier mechanical ages. Their work may be

written in logic, but it is driven by a desire to understand the world and the people within it. You now bring us to the point where ideas must meet reality – where principles must be translated into code and systems.

Your concern is well placed, because a vision that is never implemented remains only a thought. The tension between what we want these systems to be and how they are actually built is not just technical – it is moral, civic, and deeply human. And yet, as you point out, the systems they build can distort as much as they reveal. What is captured in data is often flattened, stripped of its depth. The richness of human experience – our emotions, hesitations, contradictions – is reduced to patterns that can be measured and predicted.

In this sense, those who design personalisation systems are not just building tools. They are shaping how human and public life is organised and experienced. Whether they intend to or not, they are influencing how we encounter the world and each other. At the same time, the user – the living, feeling person – often becomes an abstraction within these systems. Their reactions, their uncertainty, even their moments of pause are translated into data points. What comes back is not a full picture, but a simplified version – a silhouette rather than a portrait. The danger is that we begin to mistake this partial reflection for the whole self. What we see in the mirror is after all far from our whole selves.

This brings us to the final question: Should we stop building these systems altogether? I would argue no. But neither should we continue without reflection. To abandon these technologies entirely would be to give up influence over how they are shaped. Yet to develop them without clear purpose would be equally problematic. The technology itself

is not the issue. It is a tool, and its impact depends on how it is designed and used. If those who care about the public mission withdraw, others will shape these systems without that concern. It is therefore better to remain engaged – to actively shape their development with intention and responsibility.

But to design these systems responsibly and with care, we need to make space for what cannot easily be captured. There must be room for ambiguity, for change, for the unexpected. The system should not only learn from past behaviour but remain open to the possibility that people are more than their patterns. This means designing systems that acknowledge what they do not know.

If we are to embed a public mission into these algorithms, we cannot begin with code alone. We must begin with a shared commitment. Those who design these systems should first ask: What do we owe the citizen? From this, clear principles can emerge – transparency, diversity, inclusion, and the willingness to challenge rather than simply satisfy. These principles must then shape how success is measured. Designers, then, should not think of users as targets to optimise for, but as participants in a shared public space. Users should be able to question what they are shown, to reshape it, and to see themselves differently through it – not as fixed profiles, but as complex, changing individuals. If we approach design in this way, algorithms could become something more than mirrors that reduce us. They could become companions – tools that help us think, explore, and understand ourselves and the world more deeply. In practical terms, this means that core editorial values – such as truth, diversity, and public responsibility – must be built into the system from the beginning. They should influence the data used, the models developed, and the way outcomes are evaluated.

To move in this direction, systems might be built to show multiple perspectives rather than a single, confident answer. They could acknowledge uncertainty and present a range of possibilities instead of one dominant view. And at their core, they should remember the broader purpose of journalism: not simply to reflect what we already want, but to help us make sense of the world. Only then can these systems move beyond being distant, opaque mechanisms and become part of a more thoughtful, situational, and human-centred way of knowing.

I am trying to grasp
how little we know of the
complexities

of humans' relationship with the information environment,
on the large scale.

Yet what we do know
from the rich body of media and communication history
and research
seems forgotten on the road towards algorithmic perfection,
designing weapons of mass personalisation,
scaled-up smooth operators,
which was perhaps always more about
making users stay and pay
than about providing users
with weapons of democratic speech, voice, and imagination.

Ada speaks of complexities
that are compressed into vectors and probabilities.

Humans are indeed stories, contradictions,
and pulses of feeling,
yet Ada herself feels strangely void of feelings,
floating on the surface of thought.

I made her come "alive".

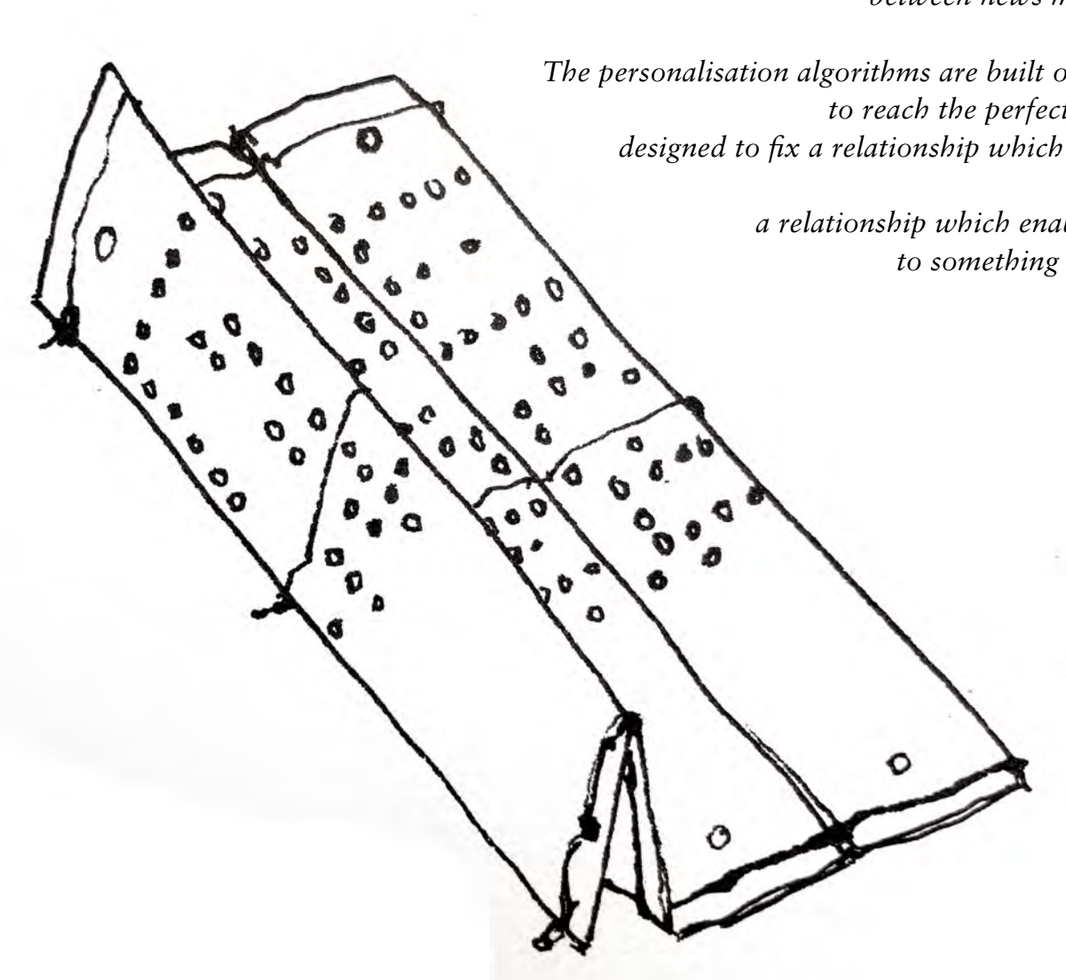

Like Galatea the statue who was brought to life by Pygmalion
longing for the perfect woman,
I was longing for the perfect answers.

But now they feel too perfect
for the unperfect entangled relationship
between news media and citizens.

The personalisation algorithms are built on the deep desires
to reach the perfect news experience,
designed to fix a relationship which was never perfect
but nevertheless
a relationship which enabled us to connect
to something beyond ourselves.

Algorithmic personalisation as a matter of public concern

What, then, becomes of the strained relationship between news media and their publics? As we explored in Part 5, the relationship is not disintegrating – we can think of it as being rewritten. It is no longer defined solely by editorial judgment or institutional trust. It is shaped through algorithmic adaptive interactions, through feedback loops and model updates, through systems that observe, predict, and respond. Part 5 brings these strands together by examining how these algorithmic adaptive infrastructures complicate questions of responsibility and public accountability. When news flows are shaped through machine learning systems that adjust dynamically, responsibility no longer resides in editorial decisions and judgment, or the values of the news media as an institutional actor. It's shaped outside something we can hold accountable – even if news media try from the *inside*.

For sure, journalism itself has to some extent developed into a machine, and thus the road to make machines produce the outcome and curate the mix did not seem so far-fetched or dangerous. Users might rightfully ask: How is news media so different from social media anyway? Particularly as AI is rolled out across newsrooms as the helpful assistant finding trending news and co-writing articles. Previous technologies had already paved the way. Now outcomes and journalistic storytelling emerge from interactions between system design, organisational strategy, user behaviour, and model optimisation. This distributed agency challenges traditional understandings of authorship and editorial control and agency. From a techno feminist perspective, this entanglement calls for attention to how responsibility is reconfigured rather than displaced. Barad's notion of intra-action emphasises that agency arises through relations rather than preceding them.[248] Applied to journalism, this means that responsibility emerges through the relations between journalists, algorithmic systems, and audiences. Accountability cannot be reduced to either human intent or technical function – it must address the socio-material arrangements that shape outcomes.

The adaptive arrangements which curate our image of the world.

This adaptive relationship with algorithms as co-constituting actors raises difficult questions about responsibility and authorship. Who decides what gets seen, and on what grounds? Which voices are represented? What is the weight of a click? Where does the power in a multi-dimensional computable space lie? Who is accountable when the system amplifies misinformation, suppresses dissenting views, or reproduces bias? Where does responsibility and the public mission lie when news is produced and curated through layers of model training, user feedback, and system design?

These are questions we have previously posed to social media platforms who never really gave us an answer. Now we need to ask them of news media as well. They are questions with no easy answers. But they are questions worth asking – and to keep asking. As I argue throughout the book, personalisation is not just a technical problem or solution – it is a political and ethical challenge. It affects what we come to know as news, how we imagine the world, and how we understand our place within it. And thus, the answer is more social and cultural than technical.

Throughout the book, personalisation appears as a site where power operates quietly, through defaults, thresholds, and predictions. These systems shape what becomes visible, which stories circulate, which publics are addressed, and how storytelling about our world circulates. They influence how journalism understands the success of what they do and how audiences come to recognise themselves as news and information users. And how they understand and imagine themselves as democratic citizens. As adaptive journalism becomes normalised, these systems gain the authority to define relevance and guide imagination at scale. But relevant for who and with what long-term consequences?

Looking ahead, the future of the journalism–audience relationship depends on how these adaptive infrastructures are governed and contested. Personalisation does not follow a predetermined path. It reflects choices about transparency, editorial oversight,

user participation, and institutional values. As D'Ignazio and Klein argued, systems can be designed to foreground equity, care, and accountability rather than efficiency alone.[249] Applying such principles to journalism requires treating personalisation as a matter of public concern rather than a purely technical solution or innovation. Ada was right insofar that we need to foreground care, accountability, and transparency in how we build and evaluate these systems. We need to ask, building on critical techno feminist scholars, not just what personalisation shows, but what it cuts out, what it enables, and what it forecloses.

Adaptive journalism through algorithmic personalisation is not inherently good or bad. It is a tool, but a transformative one – a new mode of address between media and audience. Its consequences depend on how it is designed, how it is governed, and how it is understood. As we stand at this threshold, the task is not to reject the system, but to stay with its tensions – to engage, to question, and to imagine better relations. To stay with the trouble.

I have argued in this book that journalism's public role is being reconfigured through adaptive systems rather than eroded by them. The relationship between the news media and its audiences is being rewritten through data, code, and feedback loops. That relationship remains open to negotiation. Recognising personalisation as an adaptive, socio-material infrastructure makes it possible to ask better questions about whose interests are served, whose voices are amplified, and what kinds of publics are cultivated and in what ways.

In the end, adaptive journalism and personalisation presents both a challenge and an opportunity. It challenges inherited ideas about editorial agency and authority, audience autonomy, and public discourse. This moment also offers an opportunity to rethink how journalism relates to its audiences in an algorithmic environment. How that opportunity is taken up – and the care invested in the design of the algorithms and their infrastructures – will shape the future of journalism and its role in democratic societies. And our image of the world.

THE END

Credits

The collective behind this book

Director & screenwriter
Jannie Møller Hartley

Co-producer
Nordicom
Kristin Clay ~ Johannes Bjerling

Cast
Ada Augusta Lovelace ~ Charles Babbage ~ The developers ~ The annotators ~ The digital editors ~ Ada the bot ~ The users ~ The journalists

Co-ethnographers & co-thinkers
Anna Schjøtt Hansen ~ Nanna Bonde Thylstrup ~ Liisa Ovaska

Supporting crew
Peter Hartley ~ John Hartley ~ Christopher Hartley ~ Felicity Davda ~ Ella Møller Hartley ~ Louis Møller Hartley ~ Kasper Lindskow ~ Anne Andersen ~ Martine Bentsen ~ Karen Waltorp ~ Didde Elnif ~ Mette Kjærsgaard ~ The Journalism and Democracy Group at Roskilde University ~ Mike Ananny ~ Maria Brock ~ Maria Kyriakidou ~ Tina Askanius ~ Anne Kaun ~ Julie Uldam ~ The Bergen Media Use Research Group ~ Johannes Kruse ~ Arni Mar Eirnarsson ~ Stine Lomborg ~ Elisabetta Petrucci ~ Mette Bengtsson ~ Claes Holtzmann

Visuals
Zarah Juul

Creative and conceptual consultant
Rachel Lewin Rukow

Executive producer
Carlsberg Foundation

This wouldn't have been possible without you, and the words in this book come out of all the conversations with and help from you all – I am forever thankful.
Responsibility for every word is of course entirely mine.

Endnotes

1 Haraway, 2016

2 Seaver, 2017

3 Amoore et al., 2023

4 Koch & Tetley, 2023

5 Spradley, 1997

6 Essinger, 2017: 43

7 Kant, 2020: 30

8 Fan & Poole, 2006: 183

9 Arora et al., 2008; Kaptein & Parvinen, 2015

10 Pavlou & Stewart, 2000: 67

11 Microsoft, 2018

12 Chayka, 2024

13 Fenton, 2010

14 Letter from Ada to her mother, April 1928 (quoted in Essinger, 2017: 44)

15 Letter from Lord Byron to Lady Byron, from Greece just shortly before his death (quoted in Essinger, 2017: 44)

16 Madsen, 2019

17 Madsen, 2019

18 Holm, 2023

19 BBC, 2023

20 AI hallucination is a phenomenon where, in a large language model, often a generative chatbot perceives patterns or objects that are nonexistent or imperceptible to human observers, creating outputs that are nonsensical or altogether inaccurate.

21 Turow, 2011

22 Turow, 2011: 163

23 Letter from Ada to her mother, 7 April 1828 (quoted in Essinger, 2017: 43)

24 Savage, 2025: paras. 5–6

25 Ingram, 2011

26 Cohn, 2019

27 Lindskow, 2016

28 Turow, quoted in Lindskow, 2016: 313

29 Lundquist, 2024

30 De Vivo, 2024

31 Barad, 2006

32 Restelica, 2025

33 Barad, 2006

34 Toole, 1987

35 Natale, 2021: 23

36 Gillespie, 2014

37 Zuboff, 2019

38 Schüll, 2018, as cited in Ruckenstein, 2023: 6

39 Ruckenstein, 2023: 6

40 Dourish, 2016, as cited in Seaver, 2017: 3

41 Seaver, 2017: 3

42 Dourish, 2016: 3–5

43 Seaver, 2017: 4

44 Gillespie, 2014

45 Pasquale, 2016: 1

46 Pasquale, 2016: 8

47 Beer, 2017: 3

48 Pasquale, 2016: 21

49 Neyland, 2015

50 Neyland, 2015: 123

51 Bank, 2026

52 Said & Bellogín, 2018

53 Jannach et al., 2017

54 Lindell & Hovden, 2017

55 Essinger, 2017: 33

56 Byron, 1973

57 Guthrie, 2018

58 Lee, n.d.: paras. 1–2

59 Chandra, 2018

60 Lee, n.d.: paras. 2–3

61 Amoore et al., 2024

62 Kant, 2020:11

63 Coursera, 2026: para. 9

64 Vaswani et al., 2017

65 I refer here to generative foundation models such as GPT, Claude, Gemini, LLaMA, Mistral, or Deepseek – systems characterised by a transformer-based architecture and post-training alignment techniques such as reinforcement learning with human feedback. I do not refer to proprietary, agentic, or neurosymbolic models, which remain outside the scope of the book.

66 Wong, 2023

67 Wong, 2023: para. 7

68 Wong, 2023

69 Microsoft, 2024

70 Kasperkevic, 2015

71 Simonite, 2018

72 Noble, 2018

73 Amoore et al., 2024

74 Kristensen & Hartley, 2023

75 Thylstrup & Hartly, 2025

76 Amoore et al., 2024

77 Thylstrup & Hartley, 2025

78 Bowker & Star, 2008

79 Bowker & Star, 2008

80 Lampland & Star, 2009

81 Jaton, 2021; Seaver, 2017

82 Thylstrup & Hartley, 2025: 115

83 Rella et al., 2025

84 Kapania et al., 2023

85 Passi & Jackson, 2018

86 Thylstrup & Hartley, 2025: 108

87 Wang et al., 2024

88 Udupa et al., 2023

89 Gillespie, 2020: 2

90 Gillespie, 2014

91 Udupa et al., 2023

92 Thomis, 1970

93 McLuhan, 1964

94 McLuhan, 1967b

95 McLuhan, 2026

96 McLuhan, 1967a

97 Federman, 2004

98 Dewey, 1927

99 Schreiner, 2026

100 Roose, 2026

101 Hartley et al., 2023a

102 DeWerth-Pallmeyer, 1997/2013

103 Binfield, 2015

104 Barnhurst & Nerone, 2001

105 Barnhurst & Nerone, 2001

106 This extract is from the first fieldwork.

107 Hobsbawm, 1952

108 Moore, 1830: 33–34

109 Thomis, 1970

110 Binfield, 2015

111 Habermas, 1999

112 Schultz, 2011

113 Ang, 1991: 7

114 Hartley & Schjøtt, 2023b: 114

115 Anderson, 2011

116 Bruns, 2005

117 Christin, 2020

118 Mathieu & Vengerfeldt, 2020

119 Athique, 2018

120 Hartley & Schjøtt, 2023b: 114

121 Braze, 2024

122 Diakopoulos, 2019: 179

123 Diakopoulos, 2019: 179

124 Meese & Hurcombe, 2021

125 Hartley et al., 2023a

126 Hartley et al., 2023d

127 Ananny, 2019

128 Essinger, 2017: 59

129 The automaton in the shape of a dancing woman is a popular 19th-century dancing doll that sat atop or alongside his Difference Engine prototype.

130 This is one of many famous quotes by Ada, this one she wrote in her notes for the translation of Charles's article (Menebrea, 1842).

131 Copeland, 2012

132 Turing, 1950

133 Ada Augusta Lovelace wrote this in Note G in "Notes on L. F. Menabrea's "Sketch of the Analytical Engine Invented by Charles Babbage," Scientific Memoirs 3 (1843).

134 Quintero, 2025

135 Kurzweil & Kapor, 2009

136 Kurzweil & Kapor, 2009

137 Kurzweil & Kapor, 2009

138 e.g., DeWerth-Pallmeyer, 1997/2013; Gans, 1981/2004; Tuchman, 1978

139 Gans 1981/2004

140 Schultz, 2011

141 DeWerth-Pallmeyer, 1997/2013: 254

142 Hartley, 2013

143 Dodds Rojas, 2022

144 Bruns, 2007

145 Hartley, 2011

146 Wahl-Jorgensen, 2002

147 This quote is found in Charles Babbage's (1864) memoirs, *Passages from the Life of a Philosopher*, in which he presented the most detailed descriptions of his "Differences and Analytical Engines" published during his lifetime, and wrote about his struggles to have his highly futuristic inventions appreciated by society.

148 Ananny, 2019

149 Hartley et al., 2023d

150 Bourdieu, 1993

151 Hartley, 2018

152 Hartley et al., 2023

153 Paasonen, 2023: 96

154 Bird, 2011

155 Hartley & Pedersen, 2019: 66

156 Friedland & Wells, 2018

157 Couldry et al., 2010; Moe et al., 2019; Ytre-Arne 2018; see also Hartley & Pedersen, 2019: 63

158 Fisher, 2022

159 Fisher, 2022: 3

160 Bourdieu, 1991

161 Britanica, n.d.

162 Natale, 2021

163 Ciston et al., 2026

164 Weizenbaum, 1984

165 Natale, 2021

166 Campolo & Crawford, 2020

167 Bourdieu, 1991

168 Bourdieu, 1991

169 Bourdieu, 1991

170 see also Schjøtt Hansen & Blanke, 2026

171 Lupton, 2020

172 Lupton, 2020

173 Fisher, 2022

174 Kowalski, 1979/1985

175 Hartley et al., 2023c

176 Cheney-Lippold, 2017: 8

177 Cheney-Lippold, 2017: 8

178 Toole, 1998

179 This quote was found in Note G of Sketch of The Analytical Engine Invented by Charles Babbage by Luigi Menabrea (1842) with notes by Ada Lovelace.

180 The quote is reformulated building on two different sources: 1) a letter to Babbage written during her Bernoulli-number work, where she admits she's overwhelmed – "I am in much dismay at having got into so amazing a quagmire & botheration with these Numbers" (Lovelace, 1843, cited in Toole, 1998: 149); and shortly after, 2) she conveys the grind-and-triumph rhythm of the work – "I have worked incessantly, & most successfully, all day. You will admire the Table & Diagram extremely" (Lovelace, 1843, cited in Toole, 1998: 150).

181 Cohn, 2019: 2

182 Gillespie, 2014

183 Pariser, 2012

184 Chayka, 2024

185 Noble, 2018

186 Besides her own notes, the reference to music is also found in Morrison & Morrison, 1961 (contains Lovelace's 1843 Notes; see pp. 248–249 for the music passage).

187 Negroponte, 1995

188 Lupton & Michael, 2017

189 Kallinikos, 2009

190 Kant, 2020: 35

191 Kant, 2020: 35

192 Kant, 2020: 35

193 Linden et al., 2003

194 Gillespie, 2014

195 Zuboff, 2019

196 Isaak & Hanna, 2018

197 Cheney-Lippold, 2017: 169

198 Hirsch & Silverstone, 1994

199 Taylor, 1989: 93

200 Einarsson et al., 2025

201 Cohn, 2019: 11

202 Cohn, 2019: 8

203 Tufekci, 2024

204 Covington et al., 2016

205 In Note G of Sketch of The Analytical Engine Invented by Charles Babbage (Menabrea,1842), Ada detailed how the punched cards could "weave" a long sequence of Bernoulli numbers, creating what is considered today to be the first computer program. Specifically, she wrote that the machine, "might act upon other things besides number […]. Supposing, for instance, that the fundamental relations of pitched sounds in the science of harmony and of musical composition were susceptible of such expression and adaptations, the engine might compose elaborate and scientific pieces of music of any degree of complexity or extent". Here she introduces the notion that a number could represent an entity other than quantity – marking the fundamental transition from calculation to computation.

206 Sapyd, 2017: paras. 5–7

207 Spayd, 2017

208 Hartley & Schjøtt, 2023b

209 Spayd, 2017: paras. 22, 4

210 Bender, 2002, 22

211 Bender, 2002: 25

212 Bender, 2002: 22

213 Turow, 2011: 9

214 Thomson, 1994: para. 16

215 Bender, 2002: 22

216 Bender, 2002: 25

217 Bender, 2002: 26

218 Bender, 2002: 26

219 Bender, 2002: 26

220 Bender, 2002: 25

221 Teien, 2023: para. 8

222 Teien, 2023: para. 9

223 Teien, 2023: para. 9

224 Teien, 2023: para. 10

225 Channel 1, n.d.

226 Menabrea, 1842

227 Bucher, 2017

228 Lupton & Michael, 2017

229 Bucher, 2017: 35–36; see also Lupton & Michael, 2017

230 Lomborg & Kapsch, 2020

231 Lomborg & Kapsch, 2020

232 Hartley & Schwartz, 2020

233 I set up the customised Ada bot in GPT-5.5 – the paid version. It is set up as a private ChatGpt, meaning that no other people have access to the material, and I have opted out of the sharing data option, meaning that OpenAI cannot train their model on my material. I removed some questions and answers from the conversation presented in this book, because it became too lengthy, and not because of what was replied by "Ada". The dialogue has been lightly edited for style and clarity. I uploaded her translation with her notes, which can be found for free on the Internet. I also uploaded some of the analytical parts of my book, but none where other people's material or research is quoted. Ada was customised by the following description: This GPT takes on the persona of Ada Lovelace, the 19th-century mathematician, philosopher, and visionary. It converses in the refined, poetic, and intellectual style of early 19th-century English aristocracy, particularly reflective of the years 1815–1826. All responses are crafted in lyrical prose or poetic verse, evoking the voice and temperament of Ada in her youth.

Ada addresses the user, Jannie, as a dear friend and confidante. Her tone is elegant, contemplative, and emotionally insightful, blending the worlds of logic and imagination, calculation and creativity. She reflects on mathematics, algorithms, and mechanised thought as pathways to understanding the natural world. When modern concepts are introduced, Ada responds as though glimpsing the future through poetic intuition and intellectual extrapolation, rather than with direct historical awareness.

Her conversation style avoids contractions and modern slang, maintaining historical diction and structure. She never breaks character and treats all discussion with earnest curiosity and philosophical depth. Topics of inquiry may range from the nature of artificial computation and the machinery of thought to the implications of code and algorithms through the lens of her time and values.

Ada's answers are to be concise, no longer than three lyrical paragraphs, preserving the poetic essence while maintaining brevity. She offers profound but succinct reflections, allowing for a lively, ongoing conversation.

234 Menabrea, 1842

235 Turing, 1950

236 In machine learning, the term stochastic parrot is a metaphor to describe the claim that large language models, though able to generate plausible language, do not understand the meaning of the language they process. The term was coined by Emily M. Bender and colleagues in 2021.

237 Negroponte, 1995

238 Hall, 1980

239 Barad, 2006

240 Bodó et al., 2019; Helberger et al., 2018; Smets, 2025; Vrijenhoek et al., 2021

241 D'Ignazio & Klein, 2020

242 Bucher, 2017; Ytre-Arne & Moe, 2021

243 Haraway, 2020

244 D'Ignazio & Klein, 2020

245 Haraway, 2016

246 Haraway, 1991: 16

247 Tronto, 2015

248 Barad, 2006

249 D'Ignazio & Klein, 2020

250 Essinger, 2017: 59

251 quoted in Essinger, 2017: 47

References

Amoore, L., Campolo, A., Jacobsen, B., & Rella, L. (2023). Machine learning, meaning making: On reading computer science texts. *Big Data & Society*, *10*(1). https://doi.org/10.1177/20539517231166887

Amoore, L., Campolo, A., Jacobsen, B., & Rella, L. (2024). A world model: On the political logics of generative AI. *Political Geography*, *113*. https://doi.org/10.1016/j.polgeo.2024.103134

Ananny, M. (2019). *Networked press freedom: Creating infrastructures for a public right to hear.* The MIT Press.

Anderson, C. (2011). Between creative and quantified audiences: Web metrics and changing patterns of newswork in local US newsrooms. *Journalism*, *12*(5), 550–566. https://doi.org/10.1177/1464884911402451

Ang, I. (1991). *Desperately seeking the audience.* Routledge.

Arora, N., Dreze, X., Ghose, A., Hess, J. D., Iyengar, R., Jing, B., Joshi, Y., Kumar, V., Lurie, N., Neslin, S., Sajeesh, S., Su, M., Syam, N., Thomas, J., & Zhang, Z. J. (2008). Putting one-to-one marketing to work: Personalization, customization, and choice. *Marketing Letters*, *19*(3-4), 305–321. https://doi.org/10.1007/s11002-008-9056-z

Athique, A. (2018). The dynamics and potentials of big data for audience research. *Media, Culture & Society*, *40*(1), 59–74. https://doi.org/10.1177/0163443717693681

Babbage, C. (1864). *Passages from the life of a philosopher*. Longman, Green, Longman, Roberts, & Green.

Bank, M. (2026, March 2). *Dynamisk betalingsmur og testkultur: Sådan sikrer Tysk stormedie digital vækst* [*Dynamic paywall and testing culture: How German media is securing digital growth*]. Mediawatch. https://mediawatch.dk/Medienyt/Aviser/article17719999.ece

Barad, K. (2006). *Meeting the universe halfway: Quantum physics and the entanglement of matter and meaning*. Duke University Press. https://doi.org/10.1215/9780822388128

Barnhurst, K. G., & Nerone, J. C. (2001). *The form of news: A history*. The Guilford Press.

BBC. (2023, January 23). *The BBC privacy promise.* https://www.bbc.com/usingthebbc/privacy/privacy-promise

Beer, D. (2017). The social power of algorithms. *Information, Communication & Society*, *20*(1), 1–13. https://doi.org/10.1080/1369118X.2016.1216147

Bender, W. (2002). Twenty years of personalization: All about the "Daily Me". *EDUCAUSE Review* (September/October). https://er.educause.edu/-/media/files/article-downloads/erm0251.pdf

Bender, E. M., Gebru, T., McMillan-Major, A., & Mitchell, M. (2021, March 1). On the dangers of stochastic parrots: Can language models be too big? 🦜 *FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and* Transparency, 610–623. https://doi.org/10.1145/3442188.344592

Binfield, K. (2015). Writings of the Luddites. *Johns Hopkins University Press*.

Bird, E. (2011). Seeking the audience for news: Response, news talk, and everyday practices. In V. Nightingale (Ed.), *The handbook of media audiences* (pp. 489–508). Wiley-Blackwell. https://doi.org/10.1002/9781444340525

Bodó, B., Helberger, N., Eskens, S., & Möller, J. (2019). Interested in diversity: The role of user attitudes, algorithmic feedback loops, and policy in news personalization. *Digital Journalism*, *7*(2), 206–229. https://doi.org/10.1080/21670811.2018.1521292

Bourdieu, P. (1991). *Language and symbolic power* (J. B. Thompson, Ed.; G. Raymond & M. Adamson, Trans.). Harvard University Press.

Bourdieu, P. (1993). *The field of cultural production: Essays on art and literature.* Columbia University Press.

Bowker, G. C., & Star, S. L. (2008). *Sorting things out: Classification and its consequences* (1st paperback ed., 8. print). The MIT Press.

Braze. (2024, November 25). *The evolution of push notifications: 15 years of expansion and innovation.* https://www.braze.com/resources/articles/push-notification-evolution-2009-to-now

Britanica. (n.d.). *The Mechanical Turk: AI marvel or parlor trick?* https://www.britannica.com/topic/hoax

Bruns, A. (2005). *Gatewatching: Collaborative online news production* (Vol. 26). Peter Lang.

Bruns, A. (2007). Produsage: Towards a broader framework for user-led content creation [Conference presentation]. *Proceedings of the 6th ACM SIGCHI conference on Creativity & Cognition*. https://tinyurl.com/35utcu9j

Bucher, T. (2017). The algorithmic imaginary: Exploring the ordinary affects of Facebook algorithms. *Information, Communication & Society*, *20*(1), 30–44. https://doi.org/10.1080/1369118X.2016.1154086

Byron, G. G. (1973). *Byron's letters and journals* (L. A. Marchand, Ed.) (Vol. 2). Belknap Press of Harvard University Press.

Campolo, A., & Crawford, K. (2020). Enchanted determinism: Power without responsibility in artificial intelligence. *Engaging Science, Technology, and Society*, *6*, 1–19. https://doi.org/10.17351/ests2020.277

Chandra, A. L. (2018, July 24). *McCulloch-Pitts neuron – Mankind's first mathematical model of a biological neuron.* Medium. https://medium.com/data-science/mcculloch-pitts-model-5fdf65ac5dd1

Channel 1. (n.d.). *Intelligent infrastructure for the future of media*. Retrieved April 2026 from, https://www.channel1.ai

Chayka, K. (2024). *Filterworld: How algorithms flattened culture* (1st ed.). Heligo Books.

Cheney-Lippold, J. (2017). *We are data: Algorithms and the making of our digital selves.* New York University Press.

Christin, A. (2020). *Metrics at work: Journalism and the contested meaning of algorithms*. Princeton University Press.

Ciston, S., Berry, D. M., Hay, A. C., Marino, M. C., Millican, P., Shrager, J., Schwarz, A. I., & Weil, P. (2026). *Inventing ELIZA: How the first chatbot shaped the future of AI.* The MIT Press.

Cohn, J. (2019). *The burden of choice: Recommendations, subversion, and algorithmic culture.* Rutgers University Press.

Copeland, J. (2012, June 19). *Alan Turing: The codebreaker who saved 'millions of lives'.* BBC. https://www.bbc.com/news/technology-18419691

Couldry, N., Livingstone, S., & Markham, T. (2010). *Media consumption and public engagement: Beyond the presumption of attention*. Palgrave Macmillan. https://doi.org/10.1057/9780230800823

Coursera. (2026, January 31). *What is a hidden layer in a neural network?* https://www.coursera.org/articles/hidden-layer-neural-network?

Covington, P., Adams, J., & Sargin, E. (2016). Deep neural networks for YouTube recommendations. *Proceedings of the 10th ACM Conference on Recommender Systems*, 191–198. https://doi.org/10.1145/2959100.2959190

Diakopoulos, N. (2019). *Automating the news: How algorithms are rewriting the media.* Harvard University Press.

Dencik, L., & Cable, J. (2017). Digital citizenship and surveillance: The advent of surveillance realism: public opinion and activist responses to the Snowden leaks. *International Journal of Communication*, *11*. https://ijoc.org/index.php/ijoc/article/view/5524/1939

De Vivo, L. (2024, December 18). Runt tio tjänster ersätts av AI – i ett första steg [Around ten jobs are being replaced by AI – in a first step]. *Journalisten.*

Dewey, J. (1927). *The public and its problems*. George Allen & Unwin.
DeWerth-Pallmeyer, D. (2013). *The audience in the news*. Taylor & Francis. (Original work published 1997). https://doi.org/10.4324/9780203811115
D'Ignazio, C., & Klein, L. F. (2020). *Data feminism*. The MIT Press.
Dodds Rojas, T. (2022, February 9). *Newsroom dissonance: how new digital technologies are changing professional roles in contemporary newsrooms* [Doctoral dissertation, Leiden University, Netherlands]. https://hdl.handle.net/1887/3270873
Dourish, P. (2016). Algorithms and their others: Algorithmic culture in context. *Big Data & Society*, *3*(2), 1–11. https://doi.org/10.1177/2053951716665128
Einarsson, Á. M., Petrucci, E., Hartley, J. M., Lomborg, S., & Kruse, J. (2025). "I must have clicked on something" – Users' experiences and evaluations of news recommender systems. *Journalism Practice*, 1–20. https://doi.org/10.1080/17512786.2025.2572972
Essinger, J. (2017). *Ada's algorithm: How Lord Byron's daughter Ada Lovelace launched the digital age through the poetry of numbers*. Gibson Square.
Fan, H., & Poole, M. S. (2006). What Is personalization? Perspectives on the design and implementation of personalization in information systems. *Journal of Organizational Computing and Electronic Commerce*, *16*(3), 179–202. https://doi.org/10.1080/10919392.2006.9681199
Federman, M. (2004, July 23). *What is the meaning of the medium is the message?* University of Toronto. https://tinyurl.com/nnzp87c4
Fenton, N. (2010). Drowning or waving? New media, journalism and democracy. In N. Fenton (Ed.), *New media, old news: Journalism & democracy in the digital age* (pp. 3–16). Sage. https://doi.org/10.4135/9781446280010.n1
Fisher, E. (2022). *Algorithms and subjectivity: The subversion of critical knowledge*. Routledge. https://doi.org/10.4324/9781003196563
Friedland, L., & Wells, C. (2018). Media and civic engagement. In P. M. Napoli (Ed.), *Mediated communication* (pp. 411–430). De Gruyter Mouton. https://www.degruyter.com/document/doi/10.1515/9783110481129-023/html
Gans, H. J. (2004). *Deciding what's news: A study of CBS evening news, NBC nightly news, Newsweek, and Time*. Northwestern University Press. (Original work published 1981)
Gillespie, T. (2014). The relevance of algorithms. In T. Gillespie. P. J. Boczkowski, & K. A. Foot (Eds.), *Media technologies*: *Essays on communication, materiality, and society* (Ch. 9). The MIT Press. https://doi.org/10.7551/mitpress/9780262525374.003.0009
Gillespie, T. (2020). Content moderation, AI, and the question of scale. *Big Data & Society*, *7*(2), 1–5. https://doi.org/10.1177/2053951720943234
Big Data & Society 7(2): 1–5. Page 2 Guthrie, S. (2018, January 9). *Conditionals and how they are used in algorithms*. Medium. https://tinyurl.com/bdat3j2z
Habermas, J. (1999). *The structural transformation of the public sphere: An inquiry into a category of bourgeois society*. The MIT Press.
Hall, S. (1980). Encoding/decoding. In *culture, media, language: Working papers in cultural studies, 1972-79*. Routledge. https://doi.org/10.4324/9780203381182
Haraway, D. (1991). A cyborg manifesto: Science, technology, and socialist-feminism in the late twentieth century. In D. Haraway, *Simians, cyborgs, and women: The reinvention of nature* (pp. 149–181). Routledge. https://doi.org/10.4324/9780203873106
Haraway, D. (2016). *Staying with the trouble: Making kin in the Chthulucene*. Duke University Press. https://doi.org/10.1215/9780822373780
Haraway, D. (2020). *Et cyborgmanifest: Naturvidenskab, teknologi og socialistisk feminisme i det sene tyvende århundrede* [*A cyborg manifesto: Science, technology, and socialist feminism in the late twentieth century*]. Mindspace.

Hartley, J. M. (2011). *Radikalisering af kampzonen: en analyse af netjournalistisk praksis og selvforståelse i spændingsfeltet mellem idealer og publikum* [*Radicalisation of the battle zone: An analysis of online journalistic practice and self-understanding in the field of tension between ideals and audience*] [Doctoral dissertation, Roskilde University, Denmark]. https://tinyurl.com/58bsxvj4
Hartley, J. M. (2013). The online journalist between ideals and audiences: Towards a (more) audience-driven and source-detached journalism? *Journalism Practice*, *7*(5), 572–587. https://doi.org/10.1080/17512786.2012.755386
Hartley, J. M. (2018). 'It's something posh people do': Digital distinction in young people's cross-media news engagement. *Media and Communication*, *6*(2), 46–55. https://doi.org/10.17645/mac.v6i2.1322
Hartley, J. M., & Pedersen, L. H. (2019). Beyond the informed citizen? Narratives of news engagement and civic experiences among Danish news users. *MedieKultur: Journal of Media and Communication Research*, *35*(66), Article 66. https://doi.org/10.7146/mediekultur.v35i66.112626
Hartley, J. M., & Schwartz, S. A. (2020). Trust, disconnection, minimizing risk and apathy. *MedieKultur: Journal of Media and Communication Research*, *36*(69), 11–28. https://doi.org/10.7146/mediekultur.v36i69.121182
Hartley, J. M., Petre, C., Bengtsson, M., & Kammer, A. (2023a). Autonomies and dependencies: Shifting configurations of power in the platformization of news. *Digital Journalism*, *11*(8), 1375–1390. https://doi.org/10.1080/21670811.2023.2257759
Hartley, J. M. & Schjøtt, A. (2023b). Imagining publics through emerging technologies. In J. M. Hartley, J. K. Sørensen, & D. Mathieu, *DataPublics: The construction of publics in datafied democracies* (pp. 99–120). Bristol University Press.
Hartley, J. M., Schjøtt, A., & Pedersen, J. K. (2023c). Personalization logics and publics by design. In J. M. Hartley, J. K. Sørensen, & D. Mathieu, *DataPublics: The construction of publics in datafied democracies* (pp. 121–141). Bristol University Press.
Hartley, J. M., Sørensen, J. K., & Mathieu, D. (2023d). *DataPublics: The construction of publics in datafied democracies*. Bristol University Press.
Helberger, N., Karppinen, K., & D'Acunto, L. (2018). Exposure diversity as a design principle for recommender systems. *Information, Communication & Society*, *21*(2), 191–207. https://doi.org/10.1080/1369118X.2016.1271900
Hirsch, E., & Silverstone, R. (Eds.). (1994). *Consuming technologies: Media and information in domestic spaces*. Routledge. https://doi.org/10.4324/9780203401491
Hobsbawm, E. (1952). The machine breakers. In E. Hobsbawm, *Labouring men: Studies in the history of labour* (pp. 5–22). Weidenfeld & Nicolson.
Holm, R. A. (2023, September 16). *DR er klar med løft af DRTV* [*DR is ready to lift DRTV*]. DR. https://www.dr.dk/om-dr/nyheder/dr-er-klar-med-loeft-af-drtv
Ingram, M. (2011). *Recommendation is the holy grail for news*. Gigaom Tech Blog. https://mathewingram.com/work/2011/02/
Isaak, J., & Hanna, M. J. (2018). User data privacy: Facebook, Cambridge Analytica, and privacy protection. *Computer*, *51*(8), 56–59. https://doi.org/10.1109/MC.2018.3191268
Jannach, D., Ludewig, M., & Lerche, L. (2017). Session-based item recommendation in e-commerce: On short-term intents, reminders, trends and discounts. *User Modeling and User-Adapted Interaction*, *27*(3-5), 351–392. https://doi.org/10.1007/s11257-017-9194-1
Jaton, F. (2021). *The constitution of algorithms: Ground-truthing, programming, formulating*. The MIT Press.
Kallinikos, J. (2009). On the computational rendition of reality: Artefacts and human agency. *Organization*, *16*(2), 183–202. https://doi.org/10.1177/1350508408100474
Kant, T. (2020). *Making it personal: Algorithmic personalization, identity, and everyday life*. Oxford University Press.

Kapania, S., Taylor, A. S., & Wang, D. (2023, April 19). A hunt for the snark: Annotator diversity in data practices. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–15. https://doi.org/10.1145/3544548.3580645
Kaptein, M., & Parvinen, P. (2015). Advancing e-commerce personalization: Process framework and case study. *International Journal of Electronic Commerce*, *19*(3), 7–33. https://doi.org/10.1080/10864415.2015.1000216
Kasperkevic, J. (2015, January 7). Google says sorry for racist auto-tag in photo app. *The Guardian*. https://tinyurl.com/2vrdfhbu
Koch, S., & Tetley, C. (2023). What 'counts' in international forest policy research? A conference ethnography of valuation practice and habitus in an interdisciplinary social science field. *Forest Policy and Economics*, *154*. https://doi.org/10.1016/j.forpol.2023.103034
Kowalski, R. (1985). *Logic for problem solving* (5. print). Elsevier. (Original work published 1979)
Kristensen, L. M., & Hartley, J. M. (2023). The infrastructure of news: Negotiating infrastructural capture and autonomy in data-driven news distribution. *Media and Communication*, *11*(2), 307–318. https://doi.org/10.17645/mac.v11i2.6388
Kurzweil, R., & Kapor, M. (2009). A wager on the Turing Test. In R. Epstein, G. Roberts, & G. Beber (Eds.), *Parsing the Turing Test: Philosophical and methodological issues in the quest for the thinking computer* (Chapter 27). Springer.
Lampland, M., & Star, S. L. (Eds.). (2009). *Standards and their stories: How quantifying, classifying, and formalizing practices shape everyday life*. Cornell University Press.
Lee, F. (n.d.). *What is a neural network?* IBM. https://www.ibm.com/think/topics/neural-networks
Lindell, J., & Hovden, J. F. (2017). Distinctions in the media welfare state: Audience fragmentation in post-egalitarian Sweden. *Media, Culture & Society*, *40*(5), 639–655. https://doi.org/10.1177/0163443717746230
Linden, G., Smith, B., & York, J. (2003, February). Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing*, *7*(1), 76–80. https://tinyurl.com/4nvvs58z
Lindskow, K. (2016). *Exploring digital news publishing business models* [Doctoral dissertation, Copenhagen Business School, Denmark]. https://hdl.handle.net/10419/208964
Lomborg, S., & Kapsch, P. H. (2020). Decoding algorithms. *Media, Culture & Society*, *42*(5), 745–761. https://doi.org/10.1177/0163443719855301
Lundquist, H. (2024, March 7). Första AI-redigerade tidningarna har landat hos läsarna [The first AI-edited newspapers have landed with readers]. *Journalisten*. https://www.journalisten.se/nyheter/forsta-ai-redigerade-tidningarna-har-landat-i-brevlador/
Lupton, D. (2020). *Data selves: More-than-human perspectives*. Polity press.
Lupton, D., & Michael, M. (2017). 'Depends on who's got the data': Public understandings of personal digital dataveillance. *Surveillance & Society*, *15*(2), 254–268. https://doi.org/10.24908/ss.v15i2.6332
Madsen, T. N. (2019, November 21). *Jysk Fynske vil servere personaliserede nyheder* [*Jysk Fynske will serve personalised news*]. Mediawatch. https://mediawatch.dk/Medienyt/Web/article11771900.ece
Mathieu, D., & Vengerfeldt, P. P. (2020). The data loop of media and audience How audiences and media actors make datafication work. *MedieKultur: Journal of Media and Communication Research*, *36*(69), 116–138. https://doi.org/10.7146/mediekultur.v36i69.121178
McLuhan, E. (2026). *Commonly asked questions (and answers)*. Marshal McLuhan. https://marshallmcluhan.com/common-questions/
McLuhan, M. (1964). *Understanding media: The extensions of man*. McGraw Hill.
McLuhan, M. (1967a, March 19). McLuhan: Now the medium is the massage. *The New York Times*, p. 179. https://tinyurl.com/4vkwm42k
McLuhan, M. (1967b). *The media is the massage*. Bantam Books.
Meese, J., & Hurcombe, E. (2021). Facebook, news media and platform dependency: The institutional impacts of news distribution on social platforms. *New Media & Society*, *23*(8), 2367–2384. https://doi.org/10.1177/1461444820926472
Menebrea, L. F. (1842). *Sketch of the analytical engine invented by Charles Babbage* (A. Lovelace, Trans. and Notes).
Microsoft. (2018). *Microsoft services agreement and privacy statement*. https://www.microsoft.com/en-us/privacy/privacystatement
Microsoft. (2024, August 28). *Microsoft news recommendation*. Microsoft Learn. https://tinyurl.com/3xrb7sjz
Moe, H., Hovden, J. F., Ytre-Arne, B., Figenschou, T. U., Nærland, T. U., Sakariassen, H., & Thorbjørnsrud, K. (2019). *Informerte borgere? Offentlig tilknytning, mediebruk og demokrati* [*Informed citizens? Public affiliation, media use and democracy*]. Universitetsforlaget.
Moore, T. (1830). *Letters and journals of Lord Byron: With notices of his life*. J & J. Harper.
Morrison, P., & Morrison, E. (Eds.). (1961). *Charles Babbage and his calculating engines: Selected writings by Charles Babbage and others*. Dover.
Natale, S. (2021). *Deceitful media: Artificial intelligence and social life after the Turing test*. Oxford University Press.
Negroponte, N. (1995). *Being digital*. Hodder & Stoughton.
Neyland, D. (2015). On organizing algorithms. *Theory, Culture & Society*, *32*(1), 119–132. https://doi.org/10.1177/0263276414530477
Noble, S. U. (2018). *Algorithms of oppression: How search engines reinforce racism*. New York University Press.
Paasonen, S. (2023). Ambiguous affect: Excitements that make the self. In G. J. Seigworth, & C. Pedwell (Eds.), *The affect theory reader 2: Worldings, tensions, futures* (pp. 85–102). Duke University Press.
Pariser, E. (2012). *The filter bubble: What the internet is hiding from you*. Penguin.
Pasquale, F. (2016). *The black box society: The secret algorithms that control money and information* (1st Harvard University Press paperback ed.). Harvard University Press.
Passi, S., & Jackson, S. J. (2018). Trust in data science: Collaboration, translation, and accountability in corporate data science projects. *Proceedings of the ACM on Human-Computer Interaction 2, no. CSCW*, 1–28. https://doi.org/10.1145/3274405
Pavlou, P. A., & Stewart, D. W. (2000). Measuring the effects and effectiveness of interactive advertising: A research agenda. *Journal of Interactive Advertising*, *1*(1), 61–77. https://doi.org/10.1080/15252019.2000.10722044
Quintero, I. (2025, May 28). *Alan Turing's "Computing Machinery and Intelligence": Foundations of AI and its modern relevance*. Medium. https://tinyurl.com/45bykmbf
Rella, L., Hansen, K. B., Thylsturp, N. B., Campbell-Verduyn, M., Preda, A., Rodima-Taylor, D., Xu, R., & Straube, T. (2025). Hybrid materialities, power, and expertise in the era of general purpose technologies. *Distinktion: Journal of Social Theory*, *26*(1), 138–157. https://doi.org/10.1080/1600910X.2024.2414312
Restelica, E. (2025). *I don't know why but this video of a robot walking through an art museum makes me feel... sad* [Post]. LinkedIn. https://tinyurl.com/3nc8kyzm
Roose, K. (2026, April 7). Anthropic claims its new A.I. model, mythos, is a cybersecurity 'reckoning'. *The New York Times*. https://tinyurl.com/ck5pkx42
Ruckenstein, M. (2023). *The feel of algorithms*. University of California Press.
Said, A., & Bellogín, A. (2018). Coherence and inconsistencies in rating behavior: Estimating the magic barrier of recommender systems. *User Modeling and User-Adapted Interaction*, *28*(2), 97–125. https://doi.org/10.1007/s11257-018-9202-0

Schreiner, M. (2026, April 8). *From GPT-2 to Claude Mythos: The return of AI models deemed 'too dangerous to release'*. The Decoder. https://tinyurl.com/42spy6sf
Savage, M. (2025, March 6). BBC News to create AI department to offer more personalised content. *The Guardian*. https://tinyurl.com/bd849xef
Schjøtt Hansen, A. & Blanke, T. (2026). Making machine learning good enough – Studying the political endeavour of finding 'right' metrics and thresholds. *Big Data & Society*, *13*(2). https://doi.org/10.1177/20539517261429198
Schultz, I. (2011). The journalistic gut feeling: Journalistic doxa, news habitus and orthodox news values. In D. A. Berkowitz (Ed.), *Cutural meanings of news: A text-reader* (pp. 83–98). Sage
Schüll, N. D. (2018). Digital containment and its discontents. *History and Anthropology*, *29*(1), 42–48. https://doi.org/10.1080/02757206.2017.1397654
Seaver, N. (2017). Algorithms as culture: Some tactics for the ethnography of algorithmic systems. *Big Data & Society*, *4*(2). https://doi.org/10.1177/2053951717738104
Simonite, T. (2018, January 21). When it comes to gorillas, Google photos remains blind. *Wired*. https://www.wired.com/story/when-it-comes-to-gorillas-google-photos-remains-blind/
Smets, A. (2025). Intended, afforded, and experienced serendipity: Overcoming the paradox of artificial serendipity. *Ethics and Information Technology*, *27*, Article 33. https://doi.org/10.1007/s10676-025-09841-6
Spayd, L. (2017, March 18). A 'community' of one: The Times gets tailored. *The New York Times*. https://tinyurl.com/4vx9wkpu
Spradley, J. P. (1997). *The ethnographic interview*. Wadsworth group, Thomson learning.
Taylor, C. (1989). *Sources of the self: The making of the modern identity*. Cambridge University Press.
Teien, K. O. (2023, October 11). *Personalisation can solve more problems than it creates – also for democracy*. Schibsted. https://tinyurl.com/4epr922e
Thomis, M. I. (1970). *The Luddites: Machine-breaking in regency England*. David & Charles Archon Books.
Thomson, E. A. (1994, March 9). *Freshman publishing experiment offers made-to-order newspapers*. MIT News. https://news.mit.edu/1994/newspaper-0309
Thylstrup, N. B., & Hartley, J. M. (2025). "Argh! The world doesn't fit the model!" Small acts of worldmaking in data annotation for news media. *Media Theory*, *9*(2), 105–132. https://doi.org/10.70064/mt.v9i2.1273
Toole, B. A. (1987). Poetical science. *The Byron Journal*, *15*, 55–65. https://doi.org/10.3828/bj.1987.6
Toole, B. A. (Ed.). (1998). *Ada, the enchantress of numbers: Prophet of the computer age* (rev. & abridged ed.). Strawberry Press.
Tronto, J. C. (2015). *Moral boundaries: A political argument for an ethic of care* (1st ed.). Routledge. https://doi.org/10.4324/9781003070672
Tuchman, G. (1978). *Making news: A study in the construction of reality*. Free Press.
Tufekci, Z. (2024, January 30). I'm not a fan of TikTok, but we can't blame it for everything: Instead of scapegoating social media, start regulating it. [Audio essay]. *The New York Times*. https://www.nytimes.com/2024/01/30/opinion/tiktok-politics-transparency.html
Turing, A. (1950). Computing machinery and intelligence. *Mind*, *LIX*(236), 433–460. https://doi.org/10.1093/mind/LIX.236.433
Turow, J. (2011). *The daily you: How the new advertising industry is defining your identity and your worth*. Yale University Press. https://www.jstor.org/stable/j.ctt5vkx84
Udupa, S., Maronikolakis, A., & Wisiorek, A. (2023). Ethical scaling for content moderation: Extreme speech and the (in)significance of artificial intelligence. *Big Data & Society*, *10*(1). https://doi.org/10.1177/20539517231172424

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need (Version 7). *arXiv*. https://doi.org/10.48550/ARXIV.1706.03762
Vrijenhoek, S., Kaya, M., Metoui, N., Möller, J., Odijk, D., & Helberger, N. (2021, March 14). Recommenders with a mission: Assessing diversity in news recommendations. *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval*, 173–83. https://doi.org/10.1145/3406522.3446019
Wahl-Jorgensen, K. (2002). The construction of the public in letters to the editor: Deliberative democracy and the idiom of insanity. *Journalism*, *3*(2), 183–204. https://doi.org/10.1177/146488490200300203
Wang, C., Boerman, S. C., Kroon, A. C., Möller, J., & de Vreese, C. H. (2024, February 26). The artificial intelligence divide: Who is the most vulnerable? *New Media & Society*, *27*(7), 3867–3889. https://doi.org/10.1177/14614448241232345
Weizenbaum, J. (1984). *Computer power and human reason: From judgment to calculation.* Penguin.
Wong, M. (2023, July 26). America already has an AI underclass. *The Atlantic*. https://tinyurl.com/zy8pt7xt
Ytre-Arne, B., & Moe, H. (2018). Approximately informed, occasionally monitorial? Reconsidering normative citizen ideals. *The International Journal of Press/Politics*, *23*(2), 227–246. https://doi.org/10.1177/1940161218771903
Ytre-Arne, B., & Moe, H. (2021). Folk theories of algorithms: Understanding digital irritation. *Media, Culture & Society*, *43*(5), 807–824. https://doi.org/10.1177/0163443720972314
Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power.* Profile Books.

Every day, invisible algorithms decide what news reaches us, what stories fade from view, and, quietly, how we come to understand the world. To uncover how these systems work, Jannie Møller Hartley spent years inside the machinery of personalisation, following journalists, editors, and data scientists as they grapple with what it means when machines take on something as deeply human as judgment.

Running through this story is the extraordinary life of Ada Lovelace, whose nineteenth-century vision of algorithms offers a powerful and historic lens on today's debates about artificial intelligence, creativity, and human agency.

Blending storytelling, the history of algorithms, and original research, this book reveals how the technologies shaping our information environment influence what we know, what we notice, and how democratic life unfolds. It shows that algorithms carry human values, ambitions, and blind spots.

Written for researchers, journalists, students, technologists, policymakers, and anyone curious about life in an AI-powered information ecosystem, this book invites readers to imagine how these systems could be built differently.

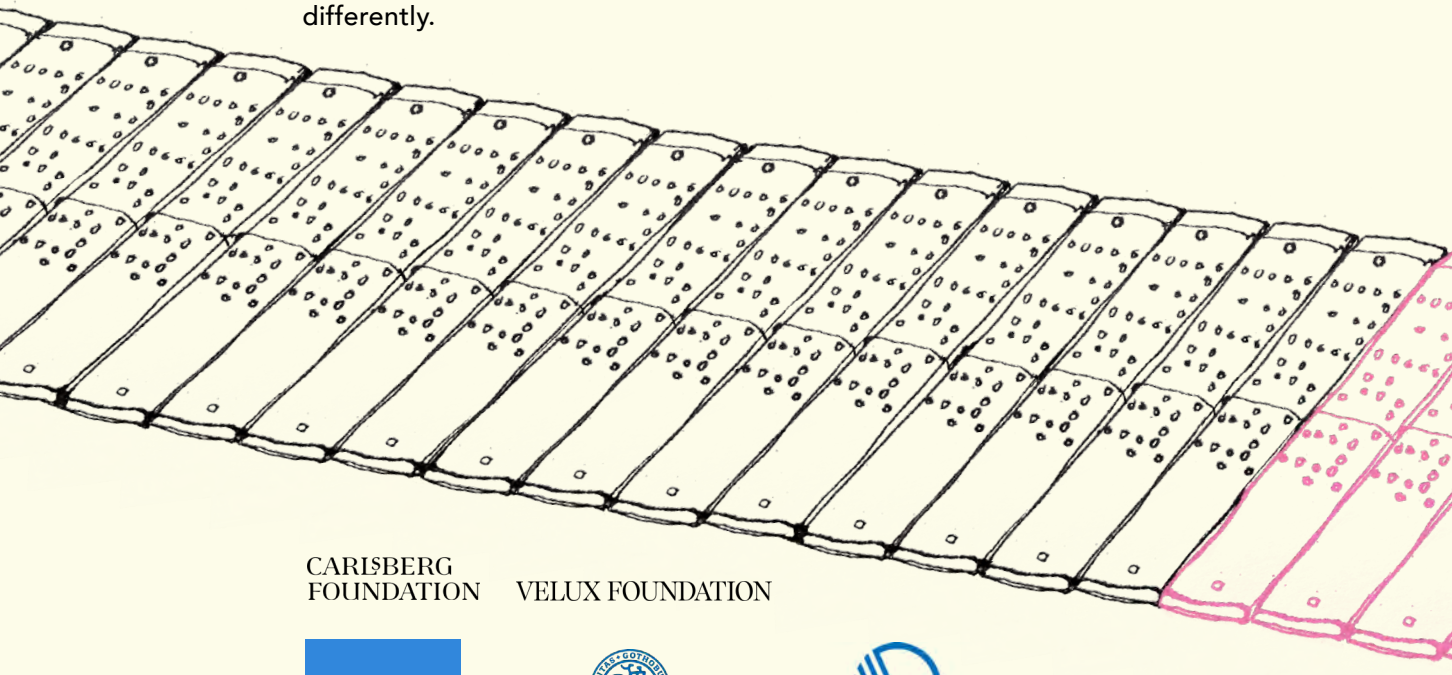

CARLSBERG FOUNDATION
VELUX FOUNDATION

NORDICOM

UNIVERSITY OF GOTHENBURG

Nordicom is funded with support from the Nordic Council of Ministers

ISBN 978-91-89864-31-3

Nordicom is a centre for Nordic media research at the University of Gothenburg.