



Nordic
Innovation

Ethical AI in Practice - Eleven Nordic Use Cases

2024

Foreword

In 2019, the Nordic Prime Ministers decided on a new vision for the Nordic cooperation, Vision2030, which holds that the Nordics is to become the most sustainable and integrated region in the world by 2030. Following this, the Nordic Ministers for Trade, Business and Industry launched eight programs for the period 2021-2024, whereof AI & Data is one. The vision is for the Nordics to be a leading region in digitalization, ethical AI and responsible data use by 2030.

Picture a world where technology not only possesses intelligence but also ensures responsible decision-making in the realms of AI. In this vision, trust, equality, and human rights form the bedrock of our digital future.

The Nordics holds unique strengths that position it to lead in digitalization, ethical AI, and responsible data use. With a rich history of data collection and the world's oldest statistical databases, we have the opportunity to set precedence for the world on how we envision data protection, privacy and human-wellbeing at the core of any AI development and thereby shaping a more integrated, liveable, and green society.

One of the Nordic strongholds is the availability of national datasets that cover many critical areas, such as healthcare, taxes, and employment. Additionally, all Nordic countries have strongly regulated data environments that share common values and understanding of ethics regarding security and responsibility of collecting, sharing, and using data. These common Nordic characteristics are both built on and

complemented by our Nordic identity and values, robust tech ecosystem, and a spirit of collaboration. While there is an excellent opportunity for Nordic countries to forge a competitive edge by collaborating on high-value datasets, there is a need for better overview on practical use cases that empower other organizations to learn, adapt, and innovate.

The purpose of this report is to stimulate and encourage the creation of ethical AI solutions based on Nordic datasets by raising awareness and increasing visibility of successful use cases across Nordics. These use cases also aim to demonstrate best practices in responsible data management and ethical AI by both Nordic businesses that have built innovative ethical AI solutions, but also by public authorities that own and make their datasets available. By inspiring Nordic organizations to innovate within ethical AI, we aim to strengthen the competitive advantage and support sustainable growth in the Nordic region.

We would like to thank Silo AI and all interviewees who contributed with insights into their use cases, but also more broadly in the field of ethical AI and responsible use of data. This report, part of the AI & Data program by Nordic Innovation, aims to serve Nordic businesses, public actors, decision-makers, and other stakeholders in the data value chain, providing valuable information for all involved.

Nordic Innovation, Oslo, June 2024
– Svein Berg
Managing Director

Background

The vision of the AI & Data program is to establish the Nordic region as world leading in digitalisation, ethical AI, and responsible use of data. In more specific terms, the aim is to support the development and deployment of ethical AI and responsible data solutions in Nordic companies through encouraging collaboration, knowledge transfer, competence building and data sharing in the Nordic region.

This report is part of a larger project, Nordic Ethical AI & Data Ecosystem Building, led by Accenture and Silo AI under the AI & Data program at Nordic Innovation. The goal of the project is to explore how Nordic datasets and capabilities in ethical AI and responsible use of data can contribute to the development of new innovative solutions in Nordic businesses.

The report has been written by Silo AI on behalf of Nordic Innovation, with the aim to enhance recognition and understanding of how Nordic organisations in both public and private sectors have responsibly developed AI based products and services, share lessons learned, best practices and capabilities used in these processes, and give visibility to Nordic AI use-cases based on Nordic datasets. The goal is to motivate organizations and companies to responsibly harness AI for stakeholder benefits. The intended audience for this report is Nordic organizations, including businesses, public actors, decision makers, and other stakeholders in the data value chain.



Definitions

This report uses the phrase “ethical and responsible AI” to describe an approach that organizations developing, deploying, and using AI technologies and applications can take. Nordic Innovation and this report define these terms as:

- **AI:** This report references the updated OECD definition¹ that holds a broad and flexible description of AI. Since the AI field is moving at an extremely fast pace, the definition of what constitutes AI moves with the technological development over time.
- **Ethical and responsible AI:** Refers to design, development, deployment and operation of AI aligned with ethical principles and requirements, and the field of questions that accompany how AI impacts its surroundings (e.g. humans, society, animals, fauna, and the planet). In practical terms, ethical and responsible AI can as an example be realized by upholding the seven key requirements for trustworthy AI, set forth by the European Commission and its High-level Expert Group on AI in 2019²:
 - Human agency and oversight
 - Technical robustness and safety
 - Privacy and data governance
 - Transparency
 - Diversity, non-discrimination and fairness
 - Societal and environmental well-being
 - Accountability
- **Ethical and responsible AI practices:** Refers to the organizational structures and steering mechanisms, processes, activities, tools, and methods that organizations apply to achieve ethical and responsible AI, exemplified by the seven key requirements of trustworthy AI defined above. These practices can be multi-disciplinary and cover for example data science, governance, legal, risk management, data security, learning, and development, and company culture.

¹ OECD.AI, Updates to the AI OECD's definition of an AI system explained, 2023

² European Commission, COM(2019)168 Building Trust in Human-Centric Artificial Intelligence, 2019

Executive Summary

The increasing adoption of Artificial Intelligence (AI) across various sectors and regions, including the Nordic countries, is not just about technological advancement but also about the accompanying ethical considerations. The Nordic region, known for its collaboration between public and private sectors, union involvement at workplaces, and awareness of potential adverse effects, is at the forefront of this development.

The use cases described in this report relates to diverse applications such as autonomous driving, micro mobility, money laundering detection, automatic speech recognition, and medical diagnosing. The building blocks span across technologies such as computer vision, large language models, federated learning, data clustering, and predictive modelling.

Despite the large spread in terms of application areas and technologies, the organizations interviewed share similar key factors in their approach to ethical and responsible development and deployment of technologies.

Ethical and responsible AI as a continuous process: ethical AI is not a static goal, but a continuous, iterative process integrated into the project lifecycle. This approach ensures that ethical considerations are an active part of decision-making rather than one-off compliance check.

Beyond binary ethics: the binary perception of projects as simply ethical or unethical is flawed, since ethical complexities are inherent in all projects. This perspective recognizes the diverse ethical challenges and the need to navigate them thoughtfully, often involving trade-offs and competing interests.

Contextual ethics: The ethical considerations presented are highly context-dependent, requiring tailored approaches to each unique use case. Despite this specificity, several overarching themes are emerging:

- **Open communication and culture:** The importance of fostering open dialogue within organizations is highlighted, with cross-disciplinary teams and routine discussions being key to maintaining a focus on non-discrimination and fairness. A diverse workforce helps drive this culture.
- **Transparency:** The balance between openness and protecting intellectual property is explored, with a call for more evidence of transparency in practice, particularly in leadership roles.
- **Accountability and ownership:** The necessity of clear internal accountability and ownership at a high organizational level is underscored to ensure responsible AI deployment. Externally, transparency is vital for building trust with stakeholders.

- **Business alignment:** Ethical AI initiatives must align with business goals and integrate with existing processes to be sustainable and effective. Ethical practices should complement business opportunities and not be isolated from other organizational activities.

In essence, the use cases demonstrate a dedication to harnessing AI's potential responsibly, aiming to inspire further ethical AI initiatives in the Nordic region and reinforcing its standing as a leader in this field.

"The binary perception of projects as simply ethical or unethical is flawed, since ethical complexities are inherent in all projects."

Table of Contents

2	Foreword
3	Background
4	Definitions
5	Executive Summary
7	Table of Contents
8	Introduction
9	Nordic Ethical AI Use cases
10	The Icelandic Centre for Language Technology
12	City of Tampere
14	Saidot
16	Head AI
18	Zenseact
20	AI Sweden
23	Telia Crowd Insights
25	Ruter
26	The CLIC project
28	Finterai
30	DoMore Diagnostics
34	Conclusions
36	About the report
37	Methodology

Introduction

This report presents insights from interviews done with eleven organisations spread around the Nordic countries, each exploring a unique application of ethical AI within the region.

The selection of use cases for the report is intentionally diverse, covering a wide range of industries, company types, and organizations throughout the Nordics. For use cases to be included, they first and foremost needed to demonstrate a tangible and positive effect at various levels; be it on individuals, society at large, or the environment, and on a local, national or global scale. Secondly, the use cases needed to come with a pronounced commitment to tackle key ethical concerns with AI - such as ensuring fairness, mitigating bias, safeguarding data privacy, maintaining transparency, and upholding a focus on human-centric values.

The report aims to enhance recognition and understanding of how Nordic organisations have responsibly developed AI based products and services, share best practices and capabilities used in these processes, and give visibility to Nordic AI use-cases based on Nordic datasets. Through this, the ambition is to provide perspectives proving valuable to the broader conversation on responsible use of AI in the Nordic region and beyond.



Nordic Ethical AI Use cases

The Icelandic Centre for Language Technology (Almannarómur) – Data and infrastructure for Icelandic Language Technology

Almannarómur is the Icelandic Centre for Language Technology and carries out the official Language Technology Programme for Icelandic.

The interview was carried out with Jóhanna Vigdís Guðmundsdóttir, Director at the Icelandic Centre for Language Technology.

Use case description

The Icelandic language technology project is an ambitious initiative aimed at preserving and advancing the Icelandic language. Having products and services supporting Icelandic alongside other languages is believed to be key to preserve and advance the language.

The core of the Icelandic language technology project lies in its focus on building a robust technology infrastructure and creating open data for language technology development. This foundation is crucial for further advancements and solutions in the field. The infrastructure in this context includes support tools for developers to build AI solutions, along with extensive language corpora. These resources have been made available under an open license, encouraging commercialization by third parties. The aim is to develop technology and products that enable the use of Icelandic in various applications, such as speech-to-text, text-to-speech, machine translation, autocorrection, and integration into common IT tools and user interfaces.

A key element of the project has been its collaboration with global technology providers. The Icelandic Centre for Language Technology has collaborated with these technology giants to include Icelandic in the tools from Meta, Amazon, Microsoft, and OpenAI. A tangible result of these efforts is the inclusion of Icelandic in the Microsoft Edge text-to-speech feature, among other potential integrations in Microsoft tools such as MS Teams.

Additionally, the project has successfully established a large open voice database. This was made possible by gathering voice recordings from people reading text aloud, thereby creating a rich resource for AI development.

The Icelandic Centre for Language Technology encourages the commercialization of products and services based on the infrastructure they have built, which they believe is the only way for language technology to become self-sustainable.

Ethical implications

Dealing with the complexities of data privacy has been among the core challenges of this project. Given the regulatory requirements from GDPR, especially concerning voice data, the project has been proactive in securing explicit consent from individuals for using their voice recordings. Additionally, the project has faced questions about the ownership of historical data, such as that originating from university research conducted decades ago.

As for many other AI projects, one of the most critical ethical aspects has been the challenge of mitigating biases, a task that becomes particularly complex when dealing with nuances around gender and ethnicity in AI solutions. The project has collaborated with AI solution providers like OpenAI to tackle these ethical dilemmas. Such partnerships are part of a strategic effort to explore and address bias within AI systems, and the centre is planning to launch a new initiative specifically aimed at this topic.

One of the main work streams of the project has been the collaboration with OpenAI to train the GPT-4 model for the Icelandic language. This collaboration aimed to ensure that the cultural values and nuances of Icelandic are accurately represented. A group of 40 volunteers was recruited to fine-tune GPT-4's Icelandic responses, using the technique known as reinforcement learning from human feedback. This is a process where respondents select their preferred model outputs from several options, and therefore has a significant impact on the output from the system.

Lessons learned and key takeaways

Due to the moderate population sizes in the Nordics, these languages have lower commercial interest than the major languages of the world, and there are also fewer resources available for them. This is particularly true for Iceland and Icelandic. Therefore, the Icelandic Centre for Language Technology sees the role of the public sector as pivotal, and that their responsibility in investing in technological and linguistic infrastructure is vital to support the Icelandic language in software and AI driven tools and services. This is a duty that falls within their purview and interest, and one that private entities is unlikely to undertake.

The use case also underlines the possible benefits that could be achieved with more extensive collaboration between the Nordic countries in the preservation of their languages, aiming to reduce expenses, maximize expertise, and take advantage of synergies due to similarities between languages. Among other advantages, this would build a stronger front in negotiations with global technology companies, and the importance of such collaboration also emphasises the value of nurturing a robust ecosystem within the Nordic region.

City of Tampere - Advanced family analytics project

City of Tampere is Finland's third largest city.

The interview was carried out with Outi Valkama, Program Manager for the City of Tampere.

Use case description

The goal of the Tampere Junior project has been to gain a deeper understanding of the needs of the families living in the city to enhance service provision. This is achieved through data analytics, collating and scrutinizing information on families and children across diverse urban configurations to detect patterns that can inform practical actions. The project was originally launched in Tampere but has expanded to include five additional Finnish cities. The project has previously released an analysis for 2022 and is set to unveil another report in the spring of 2024, offering insights into the developments of 2023.

The project has compiled data from ten different registries, containing information from 80,000 families. This covers aspects such as health, instances of family violence, financial and living conditions, and other descriptive attributes of the families. An initial cluster analysis yielded the categorization of 100 distinct family types, and this analysis lays the foundation for the detection of correlations among various factors.

As one example of the knowledge the analyses brought to light was the identification of a previously unidentified group of approximately 2,000 individuals described as "young people aged 16-18 living independently and encountering new life circumstances." While professionals, such as teachers, might have sporadically noticed students in need, systematic support for this group was absent. The insights from this analysis, however, enabled the city to devise specialized services catering directly to these young people.

On a broader scale, the analysis has resulted in a better overview and understanding of students' diverse backgrounds, which has given input to new strategies to combat school segregation across various neighbourhoods. Among the initiatives launched is the development of digital services for "caring youth", i.e., children who assume familial caregiving roles due to parental illness or other factors limiting their caretaking abilities.

Ethical implications

The project concentrated on discovering data patterns on the population level, rather than making individual inferences. The overarching objective was to create a comprehensive perspective through the combination

of diverse data types at the community level. The project successfully achieved this, constructing a strategic picture of family classifications and their respective needs.

Given the sensitive nature of the project, the City of Tampere decided to add an ethics expert alongside service professionals in a cross-disciplinary team. This person, with expertise in philosophy and theology, played a crucial role in highlighting the ethical questions.

To ensure only community level information was utilized and that individual privacy was preserved, anonymization was imperative. The project therefore adopted two principal measures to maintain the confidentiality of the data.

First, to prevent the possibility of tracing back to individual families, a threshold was established wherein each cluster included for analysis had to include at least 200 families. If any cluster fell short of this number, clusters would be merged to form larger analytical segments, maintaining the anonymity threshold.

Second, due to the sensitive nature of the data, which included health information and other personal indicators, each piece of data had to be meticulously anonymized. The Finnish Institute for Health and Welfare (THL) was commissioned to apply their robust anonymization procedures to the data, procedures originally designed for financial data.

To prevent any potential data breaches, the process was very carefully planned. At the funding application stage, the project was required to justify the specific data usage and combination, verifying that only relevant and lawful data was collected. The project had to substantiate the inclusion or exclusion of sensitive variables, like family structure, religion, and language, navigating issues related to stigmas and societal taboos. In addition, strict adherence to a minimum cluster size was enforced to ensure anonymity throughout the project.

Lessons learned and key takeaways

The analyses have proven to be an effective tool to discover strategic patterns at the community level, allowing the City of Tampere to identify groups of the population that can be supported in a better way. The project is seen as a good example that data analytics can contribute to building better societies.

As described above, ethical expertise was included as part of the project team. This was seen as a success and a critical support to the technology experts in surfacing the ethical questions. Additionally, the project engaged with various personnel from job-seeking services across different cities.

The project analyses were intentionally directed at community-level patterns, rather than focusing on individual or family-specific data. Ethical considerations followed the project's life cycle, and ethical judgments and dilemmas were actively deliberated upon as they emerged. It was acknowledged that ethics often involves navigating a landscape of divergent, sometimes opposing, constraints and interests; rarely does one encounter a straightforward choice. This required dialogue within the project team, with some topics revisited as an iterative process as the project evolved. Dedicated meetings were convened to discuss ethics, but ethical considerations were also included in the routine project team gatherings, communications planning, and announcements.

Saidot – AI governance solutions

Saidot is a Finnish based company focusing on products to systematically govern and monitor AI driven systems. Their goal is to create tools that help their clients apply ethical principles and share information with stakeholders.

The interview was carried out with Meeri Haataja, CEO.

Use case description

Saidot's platform is designed with the aim to simplify access to compliance and AI ethics resources, and to provide a standardised and easily searchable method to document all relevant aspects of algorithmic processes. These functionalities should be accessible for AI developers and product teams without the assistance of lawyers and other external experts.

The platform provides a library of best practice templates and a curated list of up-to-date regulations relevant for the AI field. It can give users step-by-step guidance through compliance processes and it offers tools to identify and mitigate risks, support in selection of foundation models (also taking risk into account), as well as methods to monitor models as external factors are evolving. Finally, it provides tools to enhance AI systems' transparency towards users and other stakeholders through functionality such as stakeholder reports.

Saidot works with large organisations adopting AI technologies, and the public sector is seen as a key area for them due to the critical need for responsible development practices when AI enabled systems potentially impacts large populations. They have been working with capital cities such as Helsinki and Amsterdam. The other primary customer segment for Saidot's platform is AI product companies who are building AI enabled solutions for their clients.

Ethical implications

AI governance needs to be a process intrinsic to an organisation or enterprise, beginning with a comprehensive understanding of the policies and requirements related to the product or service being developed. Saidot highlights the need that the requirements should stem directly from the policy framework, and AI policy templates can assist in ensuring such coherence while also taking the context of the product or service into account.

Transparency, accountability, and ownership are vital to successful development and deployment of AI based solutions. Transparency, in particular, is critical as it enables the verification of the presence and adherence to ethical principles. Ideally, there should be a designated owner for each system or product within a company.

At the same time, ethics should be approached as a collective endeavour, engaging all relevant stakeholders, and not undertaken in isolation. By opening the process to broader scrutiny, the feedback received can lead to adaptations and improvements, fostering a culture of continuous learning and responsiveness. This is the reason behind the emphasis on giving developers, users, and other stakeholders access to all relevant information regarding AI based solutions through Saidot's offerings.

Lessons learned and key takeaways

Based on the company's experience from working with systems for systematic governance for and monitoring of AI enabled products and services, three key success factors were discussed during the interview. First, establishing ownership is paramount. It is essential to clearly identify a sponsor or owner for the initiative to develop and deploy AI based systems in a responsible manner. Such ownership must be at a high enough level in the organisation to ensure that decisions being made will get the traction needed.

The second success factor relates to the integration with established processes, such as those related to sustainability or privacy. AI related projects should not live in a vacuum but be tightly knit together with other initiatives. Many of the principles governing other kinds of projects will be equally important and applicable in AI related projects, and one should avoid creating artificial barriers between AI projects and other activities. This leads to the third success factor, which is to leverage resources that are already available. This includes utilizing existing policies and platforms, as well as taking advantage of the expertise of competent individuals.

Head AI – Skills data & synthetic data

Head AI is a Finnish company delivering technology to make text-based data interoperable between individuals, companies, education, and labour markets. Their AI engine Graphmind is capable to process large volumes of textual data, make the data interoperable between different services and analyses, detect patterns, and give users actionable insights to make informed decisions.

The interview was carried out with Anu Passi-Rauste, Chief Marketing Officer and Harri Ketamo, Founder/CEO.

Use case description

At the time of this project, the Tampere region in Finland had 40,000 job vacancies and 20,000 unemployed individuals. Obvious questions in this situation are whether a larger proportion of the population can be employed through better matching between available job skills and opportunities, and whether a further increase could be achieved by putting reasonable efforts into upskilling the available work force. If data science could assist with this, it would benefit many stakeholders like job applicants, employers in search of talent, providers of training and education, and the community as such.

The concrete use case was a simulation study conducted to explore the feasibility of job and talent matching using synthetic data. 58 individuals were first surveyed about their job skills and competencies, and these data entries constituted the labelled real world data set used as a starting point. The data was anonymized, and an additional 1,500 data entries were synthetically generated in a manner reflecting the original dataset's distribution.

In the simulation study, the synthetic data set with skill profiles was matched with real world job openings. For such purposes in a wide context, Head AI harnesses a mixture of publicly funded data (e.g., EU investments, green initiatives) and open data sources (e.g., job portals, academic data repositories, Wikipedia, CNN, Business Finland, Vinnova). One of the points made during the interview was the significant potential in these resources, the need for data openness for ease of access and the necessity of curation for machine readability.

The simulations unveiled that 10% of the jobless population could be directly placed into employment, while an additional 20-60% could be engaged via reasonable upskilling initiatives. The upsides of this would be very clear, both to the individuals being employed and to the society at large: by successfully employing 10% of the unemployed through talent matching, the city could realize savings of 14 million euros. This simulation study gives the City of Tampere information needed to discuss strategies for reducing unemployment and better supporting local employers.



Ethical implications

AI based recommendation systems must navigate ethical considerations carefully. In the interview the significance of framing talent matching outcomes in a constructive manner was emphasized. As an example, such a system would in certain cases tell a user that he or she has no competencies or skills of any value, unless the system is designed with sensitivity to such issues.

Traditional career counselling practices also present ethical challenges, such as gender biased guidance being offered to students. This highlights the opportunity for a system that not only avoids ingrained biases like gender stereotypes but also considers how individuals perceive the recommendations made to them. A potential solution lies in crafting user interfaces and data presentations that are interpretable. For instance, an AI system could provide more nuanced suggestions, indicating for example "65% fit between your skills and this job", rather than a strict binary outcome. This enables job seekers to reflect on and seek secondary opinions on the system's recommendations.

In discussions of data bias, it should also be acknowledged that bias is inherent in all data. The imperative is not just to recognize the presence of bias but to understand its nature. The AI system must be designed to be cognizant of its data biases and capable of conveying this information transparently to job seekers through the user interface.

For the success of the use case, establishing unambiguous protocols for all participants in the ecosystem is critical. For example, if a person consents to share his or her data with the City of Tampere, the city must

ensure that the data is not further shared with private corporations. Explicit regulations governing data storage, processing, and sharing both across and within the ecosystem are essential. The objective is to adopt an approach anchored in user consent and ensuring data traceability.

Lessons learned and key takeaways

The company identified three important insights from their experience with the use case, which are aligned with their learnings from tens of similar kinds of use cases. First, they recognized that the discourse and theoretical underpinnings of ethics as a field have a long history, and that there is a lot to learn from the science of ethics. A multidisciplinary approach is needed though, as the field of ethical AI extends beyond the scientific domain of ethics alone.

Second, the essence of adopting an ethical AI methodology lies in the comprehensive understanding of one's actions. It is important to grasp the underlying mechanics of the technology in use, such as the dependencies within a deep learning model. For this, it is needed to learn also from the methodologies implemented by other parties.

Lastly, it is crucial to acknowledge that ethics is more than a mere preset checklist; it is an ongoing operational mode and strongly related to fundamentals of technology in use and the attitude towards problem solving in general. Ethics must be interwoven into the very fabric of the company's culture and integrated into the tools and practices it employs.

Zenseact – Autonomous driving and advanced driver assistance

Zenseact is a Swedish based company building AI powered software for advanced driver assistance systems (ADAS) and autonomous driving (AD). The company is fully owned by Volvo Cars and develops the automated driver support software stack for cars from Volvo and Polestar.

The interview was carried out with Jenny Widahl, Chief IP Officer, Legal and IP team lead and Veronika Nihlén, Head of Brand & Communications.

Use case description

Every year, traffic kills more than one million people worldwide, with the vast majority of accidents being caused by driver error. Zenseact sees automation as a key to drastically reducing this number. ADAS assist drivers with the safe operation of the vehicle and could include features such as adaptive cruise control and collision avoidance. With full AD, once in place, the car is capable of operating with no human input.

The company is working on technology encompassing most aspects of

automation, from how the car perceives and understands the world to how it navigates through it. This is a technically complex field where machine learning and deep learning is fundamental to solving the associated challenges. Adding to this, the technology will be used in cars globally, meaning that there is a need to adhere to regulations all around the world.

Ethical implications

Zenseact's goal is to make every mile safer. By steadily improving ADAS technology on the road towards AD, the benefit is that new technology will for a long time to come be tested with humans still being in control of the vehicle.

The overarching goal of improved safety is clearly for the greater good, but it is also evident that there is a range of challenges of an ethical nature that needs to be solved. Many of these are related to data and data curation. The data used for development is collected by human drivers, and data collection needs to be designed in a way that is not biased (meaning collection needs to happen in culturally diverse locations), preserves privacy (as data will contain person identifiable elements like faces and license plates), and adheres to regulations around the world. As an example, in certain parts of the world anonymization will be required to happen before the data is leaving the vehicle. At the same time, as the purpose of the data collection efforts is to fuel training and verification of new AI models, it is crucial that the processing being done to protect the privacy of people does not alter the properties of the data in a way that invalidates its use.

Going forward, Zenseact is considering the opportunity to complement some of the recorded material with synthetic data created through use of generative models. While still experimental, synthetic data offers benefits like enhanced image quality, reduced privacy concerns, lower computational and annotation demands, and the generation of rare, real-world scenarios. In the company's view more research is needed to prove how using synthetic data for training would impact AI model accuracy. Their work with synthetic data is therefore brought to use in a step-wise approach, where research guides the product development.

Finally, the topic of explainability is also crucial in the domain where Zenseact is operating. There is a need to be able to convey why using ADAS and eventually AD technologies is safe. If the team is unable to explain the actions taken by their software stack, the company is not in a position to state clearly that the technology is indeed safe. Explainability is therefore an important topic in interactions with users, the market, the company owners, and the authorities.

Lessons learned and key takeaways

Zenseact operates in a field where AI ethics plays a crucial role. Emphasis is given to the diversity of their team and the harnessing of an open, inclusive, and flat organization. The arguments given for this are that they

need everyone to contribute at their highest level to be able to solve problems as complex as autonomous driving, but also that a diverse workforce will be better able to identify different perspectives and detect bias. It is a belief in the company that diversity translates to their data used for training and validation.

One of the steps taken in Zenseact is to establish an AI ethics board. The board members primarily come from inside the company, although some of the members also could have affiliations with outside organisations. The members come from different roles such as development, research, legal compliance and communications. The need to bring people in with sound backgrounds from humanistic sciences was emphasised. The board is tasked with contributing to Zenseact's technology strategy, but also to constantly question the direction taken.

Finally, the discussion centred around publishing research and data sets as a way of progressing the field and build trust in the solution. The Zenseact team publishes research papers in well regarded scientific conferences and journals. They have also published an open data set and see this being used by their peers in the industry and by the research community.

AI Sweden – AI Ethics Lab and Natural Language Understanding

AI Sweden is the Swedish national centre for applied AI. Their mission is to accelerate the use of AI for the benefit of Swedish society, Swedish competitiveness, and everyone living in Sweden.

The interview was carried out with Francisca Hoyer, Head of Responsible AI and Operations NLU and Karin Vajta Engström, Head of Communications.

Use case description

In this interview, two use cases were explored, AI Sweden's AI Ethics Lab as well as their Natural Language Understanding (NLU) initiative.

The AI Ethics Lab was launched in 2021, initiated by the AI Sweden's Steering Board in 2021. The aim of this initiative has been to progress from broad ethics guidelines to more actionable, tangible recommendations. The Lab created a network with a diverse set of organizations, ranging from startups to public sector entities. A core part of this initiative was to identify relevant use cases within the network and construct a knowledge base from the gathered insights.

The Natural Language Understanding project was set up by AI Sweden in

collaboration with RISE Research Institute of Sweden. It originated from the launch of the Google BERT language model, which due to its large data requirements, was primarily beneficial for major languages like English, leaving smaller languages and communities underserved. Recognizing this as an ethical concern, RISE began developing a Swedish language model.

The state of art of large language models were later pushed forward with the introduction of Generative Pretrained Transformer (GPT) models. This led AI Sweden to develop a Swedish GPT model, focusing on creating the foundational technology. In the fall of 2023, GPT-SW3 was released openly, with contributions also from WASP WARA Media & Language. GPT-SW3 is a collection of large decoder-only pretrained transformer language models built on massive amounts of Swedish, Norwegian, Danish, Icelandic, and English text data, with the explicit purpose of being able to generate Swedish and Nordic languages texts. GPT-SW3 is not an off-the-shelf product or service ready for use, but rather it must be adapted by developers to build chatbots or other specific applications.

Ethical implications

The AI Ethics project revealed that while gaining general insights is relatively straightforward, applying these insights specifically and practically to applications is challenging. It was also noted that an ethics-focused team within an organization should include multidisciplinary expertise, involving case owners, developers, and individuals with diverse academic backgrounds, such as in social science. However, challenges are often highly specific to the context, making it hard for external groups to draw conclusions about projects they are not directly involved in. Hence, integrating multidisciplinary expertise into the project itself is critical, with an emphasis on weaving ethics into the process throughout the project, rather than as a peripheral consideration at the start or end.

Another common issue identified, was the communication gap between the leadership and the practitioners developing AI based solutions. To achieve successful outcomes, it is essential to bridge this gap and facilitate interaction between these two levels.

In terms of project ethics, the observation was that labelling projects as ethical or unethical is an oversimplification. Instead, ethical challenges are present in all projects, often manifesting themselves as dilemmas that require careful deliberation, such as assessing the risks associated with not pursuing a certain project. There is a balance to be struck, as being overly cautious can lead to missed opportunities due to the fear of ethical missteps. Ethics is just one facet of responsible AI. There are other critical, sometimes conflicting factors that need to be considered to ensure overall responsibility.

In the Natural Language Understanding project, transparency was a key ethical consideration. This involved ensuring open access to all aspects of the models, including data and source code, and highlighted challeng-

es like the collection and availability of data and the dilemma of making these models accessible while preventing misuse by malicious actors.

The project also emphasized the need to avoid perpetuating data biases. This was approached by involving a team of experts from the start to select and curate the training data, aiming to create a representative dataset. The interdisciplinary team of 34 experts, including journalists, gender studies experts, historians, anthropologists, and sociologists, brought a wealth of contextual understanding to the data. This resulted in better-documented training data, enabling model developers to make more informed decisions.

The multidisciplinary composition of the team, including experts from physics, computer science, humanities, and social sciences, has proven highly beneficial. The subsequent section discusses the advantages of more systematic expert involvement.

Lessons learned and key takeaways

A key insight from the interview was the value of fostering a community within organizations to engage in ethics discussions. It is crucial to have multidisciplinary expert groups, comprising individuals from diverse backgrounds, to oversee and document processes such as system development and data collection. These groups should cultivate an inclusive atmosphere where open dialogue and inquiry are encouraged.

For sustainability and scalability, especially in larger organizations, establishing permanent expert pools is advisable. These pools can form the basis for assembling temporary project teams, minimizing the need for repetitive onboarding and preserving knowledge continuity. This approach also assists the reapplication of insights gained from previous teams.

Another important takeaway was the need to integrate ethics into the entire AI solution lifecycle. From the outset of each project, ethical considerations should be a guiding principle, with product development influenced by both expert group feedback and user input, possibly through participatory design methods. However, companies face a dilemma: whether to launch products early for user feedback or wait until a certain level of product maturity is achieved.

The interview also highlighted the significance of transparency and open science in large language model development. AI Sweden advocates for European open-source alternatives to proprietary large language models, not only for trust-building, but to facilitate community and interdisciplinary communication among developers.

Lastly, for a more ethically-informed AI landscape in the long term, enhancing education at all levels, from elementary school to university, will be essential. Robust technology education ensures a comprehensive understanding of AI's capabilities and limitations. Simultaneously, reinforcing humanities and social sciences education is crucial, equipping future technology developers and users with the wide-ranging perspectives required to navigate complex ethical challenges.



Telia Crowd Insights – Societal mobility patterns 🇸🇪

Telia Company AB is a Swedish multinational telecommunications company and mobile network operator present in Sweden, Finland, Norway, Denmark, Estonia, Latvia, and Lithuania.

The interview was carried out with Kristofer Ågren, Head of Division X Product at Telia.

Use case description

Telia Crowd Insights is a service that offers data and insights on societal mobility patterns, helping to comprehend the dynamics of population movement. Actual data derived from people's movements is essential in this context.

Crowd insights have a range of applications. For instance, in the retail sector, it can aid in identifying the most strategic locations for stores by analysing residential and workplace data. For the fields of public transportation and urban development, this service is instrumental in understanding travel patterns, thus facilitating more informed decisions about resource allocation and infrastructure investments. For environmental analysis, crowd insights enable the quantification of CO2 emissions originating from public transportation. This methodology involves transforming data on people's movement, gathered from buses and trains, into carbon emission figures, integrating this with travel survey data. Given that 290 municipalities in Sweden are mandated to report CO2 emissions uniformly, the advantages of this service are substantial. As a final example, crowd insights were used extensively during the COVID-19 pandemic, where it provided critical statistical data to support decision-making processes.

Ethical implications

Telia has opted for a transparent approach in their data collection and analysis processes. During the pandemic, debates were going in the public as well as internally in Telia. Through these processes, it became clear that the only way to create trust around the use of such data was through transparency and openness. As one example, Telia chose to publish a research paper together with The Public Health Agency of Sweden describing how the service works, hence also making it open for the scrutiny of peers in this field. The most critical concern about such services is clearly the potential for individuals to be tracked. Telia Crowd Insights safeguards the privacy through a multi-step process. First anonymization is done on the network side using k-anonymity, a technique designed to protect individual privacy in datasets. This is however not sufficient, and hence multiple aggregation steps are implemented as safeguards to prevent individual identification. The data is reported at a group level and the minimum geographical area in urban settings is restricted to 500 x 500 meters.

The Telia product team continuously assesses new data combinations for their impact on anonymity. For instance, one evaluation considering potential real-time uses, such as in terrorism scenarios, led to the implementation of a 24-hour delay in data updates.

Telia has integrated ethical principles deeply into its policy and processes, viewing ethics as a core component of their value proposition, not an add-on. This stems from the telecommunication industry's stringent regulations, which encompass broader ethical principles. AI ethics at Telia aligns with their compliance processes, with issues like bias evaluated case-by-case.

In terms of AI model explainability, customer demand for clarity has prompted a preference for simpler deterministic models over complex stochastic ones.

Lessons learned and key takeaways

A key insight from this use case is the recognition of ethical AI as not only a moral imperative, but also a business advantage. Ethical AI practices shape product development and are vital for sustainable business. Without this alignment, maintaining ethical practices can become challenging. The Telia team has observed that the public communication around the ethical aspects of their services is helping them in sales presentations, as their potential customers are well versed with these topics before the meetings, and they can hence get faster to discussing the potential value and benefits to the customers.

The second lesson is the value of utilizing anonymized data. Leveraging such data simplifies processes and speeds up operations. Care must be taken that the anonymization techniques being used are sufficient to protect individual's privacy, and in the context of crowd insights this means paying close attention to the geographical resolution.

The third point emphasizes the importance of considering ethical aspects from the early stages of development. It is crucial to have a team that includes experts in sustainability, legal matters, and ethics, rather than relying solely on data scientists and developers. The teams need to integrate ethical AI into their training and professional development. This has become increasingly important compared to a few years ago, and professions such as data scientists are now expected to be versed not only in technology and natural sciences, but also in the principles of ethical AI.

Ruter – Responsible use of photos from public places

Ruter is a public transport company operating in Oslo and parts of Akershus county in Norway. It is owned by Oslo municipality and Akershus county. The company provides a variety of services, including buses, trams, trains, and ferries. Ruter's goal is to provide a sustainable, accessible, and affordable transportation system for the region.

The interview was conducted with Umair M. Imam, Head of Data & AI at Ruter.

Use case description

As an addition to the traditional modes of public transportation, Ruter now also offers their users access to micromobility vehicles, such as e-scooters, e-bikes and city bikes, through their own app. The vehicles are operated by separate vendors, and not by Ruter themselves. E-scooter services delivered by companies like VOI and Tier were first introduced in Oslo in 2019, and soon after issues arose due to users' lack of good judgment when parking the vehicles. Common problems seen were parking on private property and parking in conflict with local traffic rules and other regulations. To mitigate these problems, the vendors now have rules that users need to adhere to related to where and how to park. In order to prove that they have parked responsibly, including that the e-scooter stands and is not lying on the ground, users need to take a photo of the e-scooter through the app after parking.

In implementing access to these services through their own app, Ruter needed to introduce functionality to take pictures of parked e-scooters and transfer these to the e-scooter service provider. The challenge arising from this are the pictures that may include person sensitive data when combined with time and location. To mitigate this problem, Ruter applied AI/ML methods to blur parts of the image containing sensitive data, in this case defined as faces and license plates from cars and motor bikes.

Ethical implications

Photos taken in public places may incidentally identify people despite this not being the purpose. In this particular use case, the intention is not to

publish these photos, but Ruter has only little control over the use of the material once transferred to the e-scooter vendor. Rather than making assumptions about this, Ruter has chosen to implement own procedures to blur such sensitive information irrelevant to the use case.

Object classification in visual data is a relatively mature technology, and for face and license plate recognition Ruter chose to deploy models available to them from their cloud service vendor. This approach avoids issues with the large data efforts that would otherwise be needed for development, and the accuracy would also be expected to be very high. Nevertheless, having sound validation strategies is seen by the company as a key part of their ethical AI strategy. Therefore, the algorithm was systematically tested with images showing people of different genders and age, with different skin colours, and with people wearing glasses and face masks. No specific issues were detected during the validation of the algorithms.

Lessons learned and key takeaways

The primary takeaway from this use case was the need to hypothesize around potentially discriminatory behaviour of algorithms (such as face recognition accuracy being dependent on gender or skin colour), and then designing validation strategies targeting such factors. Having a strong focus on designing fair validation strategies is seen as key to ethical deployment of AI.

Ruter has deployed AI in various parts of their organisation. As a public entity, Ruter embraces openness and transparency, fostering an environment for open dialogues. Discussing the societal impact, business impact, and legal implications of potential new solutions is seen as a natural part of this culture and is done with a wide involvement of stakeholders in the organisation.

The organization does not have separate guidelines for developing and deploying AI solutions but instead relies on their existing standards for responsible

conduct to govern their technology related initiatives. The responsibility for both the technological and ethical aspects of the solutions developed in the end lie with the AI team leader, while other functions in the organization contribute to these discussions. The importance of good support from the legal team is seen as particularly useful in these situations.

The CLIC project – Data platform for AI based image diagnostics

The Clinical Image Consortium (CLIC) is an ongoing project funded by the Danish Innovation Fund, aiming to facilitate an accelerated development and implementation of imaging technology in healthcare services.

The interview was carried out with Troels Bierman Mortensen, project manager for CLIC and founder / CEO of DataFair.

Use case description

The objective of the CLIC Project is to accelerate the development and implementation of image-based AI solutions in the Danish health care sector, against a backdrop of radiologist shortages and an increasing demand for radiological examinations.

The purpose of the project is to develop an agile data platform to grant commercial as well as non-commercial entities secure access to data and services supporting development, validation and implementation of image-based AI solutions for radiology. This will help facilitate collaboration between companies offering AI solutions and the public health care system wanting to integrate AI solutions into their clinical workflow.

One of the advantages of working with radiological data lies in the fact that image data formats are typically standardized, simplifying data sharing from a technical standpoint. This is countered, however, by the legal complexities. The high costs associated with developing a quality dataset currently pose significant hurdles in addressing the challenges of AI based image diagnostics in clinical practices.

The ultimate aim for CLIC is to make anonymized datasets accessible, ensuring full compliance with relevant laws. Transparency is critical in how companies handle this data, including clear communication about the data's use, the validity of results, and specifics about the data contributing to these results. Gaining trust from both physicians and patients will be crucial in this process.

Ethical implications

Due to the ethically sensitive nature of healthcare data, the use of such data has for many years been governed by laws and regulations. Initially, GDPR was perceived as a hindrance for innovation, but the interviewee highlighted that as GDPR-like regulations are being adopted globally, the Nordic region's stringent standards and culture for data handling might provide a competitive edge.

In healthcare, AI based solutions typically are considered medical devices and need to be certified as such. This means that the process of updating a model comes with a considerable complexity. Another challenge arises because the population used for training and validating the model often differs from the population encountered in clinical settings, which in essence implies that the data is biased. The model accuracy cannot be guaranteed across different datasets and equipment. Therefore, it is crucial for physicians to be able to scrutinize inferences made by AI models. Additionally, monitoring for model drift, making biases explicit, and implementing mitigation strategies are essential.

Requirements might mandate companies to demonstrate that their solutions are as effective, for example, on a Danish population as they are on the original population used for training. This highlights the need for companies to access necessary data for such validations.

Lessons learned and key takeaways

This interview outlines three key learnings. Firstly, one might argue that ethics play a more significant role in healthcare compared to many other sectors. This is reflected in the higher level of regulations that govern this domain. Secondly, there is a pronounced need for functional ecosystems and technical platforms. Such infrastructures should facilitate interaction among suppliers, users such as clinicians, and patients, allowing each group to address their ethical concerns. They should also tackle the unique characteristics of machine learning systems, particularly concerning the precision of trained models. Finally, questions about data ownership, especially in cases where data has been enhanced or cleansed through expensive procedures, present complex challenges in this context.

Finterai – Federated learning for the financial sector

Finterai is a Norwegian start-up, developing and selling infrastructure for federated learning to the financial sector.

The interview was carried out with Alexander Skage, CEO.

Use case description

The use case centres around the application of federated learning within the financial sector. The primary challenge lies in enabling banks to collaboratively improve their AI models for financial transaction monitoring and money laundering detection, while not wanting to share data with their competitors in the traditional sense. A significant issue in banking is the high rate of false positives in money laundering detection, with the majority share of alerts typically being unfounded. Better money laundering detection models are therefore needed to avoid both false negatives and false positives.

Since money laundering is relatively speaking a rare phenomenon, individual banks possess limited numbers of data points representing these cases. The rare occurrence of money laundering also leads to highly imbalanced data sets, further complicating matters for AI based detection methods. It is therefore expected to bring substantial value if models can be improved based on data from multiple institutions in the sector.

Another potential application of federated learning is in facilitating cooper-

ation between different branches of a bank without the need for direct data sharing. Finterai's platform combines self-hosted infrastructure components with Infrastructure-as-a-Service (IaaS). Part of the solution is hosted in their data center, while another segment operates on the premises of the customer, adhering to the principles of federated learning. Finterai aims to establish a consortium of insurance companies, where each member develops its own model, but benefits from data from other organizations for training. In this situation, each consortium will set its own rules regarding model sharing.

Ethical implications

From an ethical standpoint, the primary concern in federated learning is data privacy. Federated learning offers a privacy-preserving solution as each participating company retains its data within its own database, rather than pooling it into a central repository. This non-sharing of data upholds user privacy. Federated learning is also beneficial for data processors handling data for data controllers, and it simplifies the deletion of an individual's data since it is not widely disseminated.

Federated learning can also contribute to reducing bias and promoting fairness in AI models. By enabling clients to incorporate a more diverse range of data, the resulting datasets are more representative and balanced.

However, broader collaboration and alignment among organizations are needed to fully leverage federated learning. The potential is significant, especially considering challenges such as money laundering that require cross-organizational cooperation. In the Nordic context, the culture of collaboration is a strength, contrasting with the competitive nature seen among banks in some other parts of the world.

Yet, there are adoption challenges, akin to a "chicken and egg" scenario. Federated learning, being a relatively new technology, faces integration difficulties with existing legacy systems in sectors like banking and healthcare. There is hesitancy among clients to commit to new consortiums. Another technical challenge is standardizing the features or data inputs for AI models. Successful implementation of federated learning requires a higher maturity in AI practices, including well-defined AI workflows, standardized ETL (Extract, Transform, Load) processes, consistent use of databases or feature stores for data extraction, and the use of a declarative machine learning framework.

Finally, it should be noted that federated learning as a technology is still maturing, and it comes with its own set of challenges like data heterogeneity issues, increased model and system complexity, and in particular vulnerabilities to cybersecurity threats, including adversarial attacks from malicious participants that could in given circumstances lead to sensitive data being leaked. Partly as a consequence of such issues, relatively few market players are currently actively pursuing this technology.

Lessons learned and key takeaways

Federated learning presents a promising approach for upholding data privacy, as it allows each participating company to retain its data within its own database, rather than consolidating it into a central repository. This model ensures that data privacy is maintained since individual datasets are not shared.

Additionally, federated learning aids in reducing bias and fostering fairness in AI models. This is achieved by enabling companies to utilize a broader and more diverse range of data, resulting in AI models that are built on more representative datasets.

In the context of the Nordics, the region's strong culture of collaboration offers a conducive environment for the successful implementation of federated learning. This collaborative ethos enhances the potential for federated learning to be effectively utilized in various sectors within the Nordic countries.

It should still be noted that the technology comes with its own set of challenges, as discussed above.

DoMore Diagnostics – Cancer diagnostics solutions

DoMore Diagnostics is a Norwegian start-up with a goal to transform cancer diagnostics with AI to improve patient care and make drug development more effective.

The interview was carried out with Torbjørn Furuseth, co-founder/CEO.

Use case description

DoMore Diagnostics is aiming to change cancer treatment by tailoring it to the individual needs of patients. The current approach in cancer treatment often involves chemotherapy. While this treatment is life saving for many patients, it is also well known to cause significant side effects. The company's goal therefore is to apply AI to predict more precisely the outcome of a patient's cancer and to guide personalized treatment.

To be able to work in this domain, DoMore Diagnostics has focused on building a rich data set with 4,500 patient samples and 90 million image tiles from more than 200 sites in 8 countries. The samples were mounted on glass slides by different laboratories and scanned at two sites with two different scanners, in order to create a robust data set for training and verification of AI models. The interviewee emphasizes that while



anonymization is typically very complex for health data, in this case only very specific data is needed, i.e., histology images accompanied by clinical outcomes.

AI is believed to represent a significant potential for improvements in cancer treatment, particularly in speeding up diagnoses and analyzing tumour tissues, but also to select personalized treatment strategies. DoMore Diagnostics has trained AI models based on histology images, and these models are now capable of identifying features indicative of high-risk tumours and categorizing patients into high or low risk groups, accordingly. The solution has been validated with 1,100 patients and is currently undergoing further clinical evaluation.

Ethical implications

One challenge expressed in the interview is the stereotypical expectation that AI-based solutions need to be close to perfect to be used in health-care industry, and therefore deemed often not be sufficiently good. This could be the situation despite such solutions surpassing the quality of human-only based decision-making. In the interview, this was described as a conservative attitude, with a reluctance to embrace AI advancements.

Pathologists traditionally examine tumours manually, inadvertently leading to variations in assessments done by different pathologists, or in assessments done by the same pathologist at different times. AI methods do not suffer from such issues but have their own set of challenges. One of these is the explainability of results from complex AI models. Several

approaches exist, such that pathologists could examine samples taking advantage of the findings from the AI based analysis. This way the AI system would become an aid to the doctors, rather than having the purpose to replace human expertise.

Another significant technical challenge is the variability in laboratory environments, such as different scanners from various manufacturers. Developing an universal AI solution compatible with all scanners is a complex challenge, requiring iterative development and a strong focus on the verification and validation procedures. It also requires collecting comprehensive data from hospitals for long-term patient follow-up. This is difficult but crucial for understanding treatment impacts over extended periods.

High-quality data and its processing are vital in this context. The Nordics have a competitive advantage due to their robust data sets, unified health systems, centralized registries, and a history of meticulous data collection. Due to the sensitive nature of health data, some entities employ federated learning across hospitals without data sharing. However, federated learning also has its complexities and shortcomings as described for a previous use case in this report.

Despite the size of DoMore Diagnostics' data set, certain issues still remain. While being designed with gender balance in mind, the data contains a higher proportion of men than women, reflecting the higher incidence of cancer in males. Regarding ethnic diversity, the dataset primarily comprises UK and Norwegian patient material, showing limited ethnic variation. The model trained from this data is currently being tested in the Netherlands, Japan and the US to validate its effectiveness across different populations.

DoMore Diagnostics has licensed their deep learning AI models from renowned universities. These research groups have published findings in well-respected journals, hence contributing to the research community and at the same time being open to the scrutiny from peers in the field. They are also addressing how to interpret and respond to the findings based on AI models. For instance, the algorithm suggests that patients identified as low risk may not require aggressive treatment. There are further validation studies ongoing to document this approach in different geographies and treatment settings, underscoring the need for careful risk assessment in determining treatment options.

Lessons learned and key takeaways

Several interesting points were raised during the interview. One of these was the perception that people may typically require higher standards from automated solutions than what they do for existing solutions and humans solving the same task, and that acceptance for erroneous diagnosis is lower when committed by technology compared to through a manual process. Hence, automated solutions need to show significantly



improved accuracy to be accepted, particularly for safety critical tasks like diagnosing. Similar discussions can be found for instance in relation to autonomous driving.

AI based cancer diagnostic solutions are not meant to replace pathologists, but rather to support in their work, making them more effective and help them making diagnoses more precise. There is a scarcity of pathologists in most countries, and improving the efficiency in this line of work will therefore be an advantage to the society. The approval processes for medical technology are slow, and it therefore takes a long time before these benefits can be realized. In the interview this was described as a balance between "misuse" and "missed use".

DoMore Diagnostics has a strong focus on high quality data sets, needed for training and validation of AI based methods. Contrary to many other applications in the healthcare domain, this particular use case is not the most challenging when it comes to privacy issues, as the data can be anonymized with simple means.

Finally, the company has chosen to in-license their technology from academic institutions, which also ensures that methods and results are published. This transparency is seen as a way to build trust in the technology, although it was acknowledged during the interview that complete transparency is not always feasible due to the risk of competitors duplicating their findings, which would in turn jeopardize their competitive situation.

Conclusions

The use cases presented in this report have covered different industries, solving different problems with different use of AI technologies. The emphasis on the different aspects of ethical and responsible AI has varied across the use cases, but together they span all the seven key requirements for trustworthy AI as set forth by the European Commission and its High-level Expert Group on AI in 2019. In some cases, the focus has primarily been on data and privacy preservation, whereas in other cases the focus has been on creating trust and transparency. However, all the use cases have demonstrated ways of using AI for good by solving real-life problems and developing better products, services, and insights for Nordic citizens and society.

One of the common denominators across the use cases is the need to consider the focus on ethical and responsible behaviours as an ongoing and iterative process where the ethical questions are discussed throughout the project lifecycle, as opposed to being a one-time checkbox exercise.

There is also a shared perception that thinking of projects as being either ethical or unethical is a misconception. Rather, there are ethical challenges in all projects, often with multiple possible routes with different advantages and disadvantages, and sometimes in the form of divergent or even opposing constraints and interests. Rarely does one encounter a straightforward choice.

As demonstrated by cases presented in this report, the ethical questions are highly context specific and need to be addressed in relation to a particular use case. However, a number of general conclusions can be drawn from these eleven use cases:

- **Open communication / Culture:** It is important to enable open discussions in the organisation, e.g. by ensuring cross-disciplinary teams. It is also crucial to make room for discussions in the teams and make these part of the daily routines. For example, in one of the use cases the ethical questions were discussed every time the project teams met. In another case a cross-functional AI ethics board had been established as an instrument to make sure that the company constantly questions the direction taken and that they stay true to developing ethical and responsible AI. The importance of a diverse and open working culture in order to uphold a focus on non-discrimination and fairness was a point made in most of the interviews held with the use case stakeholders.

- **Transparency:** It is beneficial to share data, methods and results as openly as possible, e.g. by publishing research papers, talking to the media, and answering questions. This is of course a balancing act with protecting the intellectual property and value of a company. However, in order to create trust with customers, users, and other stakeholders, and to demonstrate a commitment to accountability, transparency is of utmost importance. Even though transparency is at the heart of Nordic society, there is still some work to do. As stated by Meeri Haataja from Saidot: "I would love to say that transparency is at the heart of our Nordic democratic society. However, I miss evidence for actually living it out. The narrative is not really translated into leadership. I was hoping to see much more happening now."
- **Accountability and ownership:** These are crucial factors, both internally and externally. Internally, it is essential to clearly identify a sponsor or owner for the initiative to develop and deploy AI based systems in a responsible manner. Such ownership must be at a high enough level in the organisation to ensure that decisions being made will get the traction needed.
- **Alignment with business goals and existing business processes:** Initiatives towards ethical and responsible AI need to be aligned with the business goals of the organisation to make them a priority and sustainable for the business. One key takeaway from the interviews is that potential initiatives must be assessed in terms of the business opportunities that are opened as a result. An ethical practice that is not good for the business will be very hard to sustain. Another point made was the integration with established processes, such as those related to sustainability or privacy. AI related projects should not live in a vacuum but be tightly knit together with other initiatives. Many of the principles governing other kinds of projects will be equally important and applicable in AI related projects, and one should avoid creating artificial barriers between AI projects and other activities.

In summary, the use cases presented in this report showcase along their different dimensions their commitment to harnessing the potential of AI for the benefit of society while upholding ethical standards, data privacy, and transparency. It is the goal of this report that these examples serve as inspiration for new initiatives in the Nordic region that can contribute to the region's reputation as a leader in ethical and responsible AI development and implementation.

About the Report

This publication is a part the Nordic Ethical AI and Data Ecosystem project, in the AI and Data program by Nordic Innovation, a project led by Accenture and Silo AI.

Acknowledgements

We would like to express our sincere appreciation to all contributors that provided input to this report. This includes The Icelandic Language Technology Center, City of Tampere, Saidot, HeadAI, Zenseact, AI Sweden, Telia Crowd Insights, Ruter, The CLIC project, Finterai, and DoMore Diagnostics.

Nordic Innovation

Nordic Innovation is an organization under the Nordic Council of Ministers. Nordic Innovation aims to make the Nordics a pioneering region for sustainable growth and works to promote entrepreneurship, innovation and competitiveness in Nordic businesses. We support the Nordic Prime Ministers' vision that the Nordic region will become the most sustainable and integrated region in the world in 2030.

Silo AI

Silo AI is Europe's largest private AI lab – a trusted AI partner that brings competitive advantage to product R&D. We build AI-driven solutions and products to enable smart devices, autonomous vehicles, industry 4.0, and smart cities. Silo AI provides its customers unique access to world-class AI expertise, as well as the Silo AI product offerings to speed up AI development and deployment. Established in 2017, Silo AI is on a mission to build a European flagship AI company, with presence currently in Finland, Sweden, Denmark, Norway, the Netherlands, Germany, United States and Canada. Visit us at www.silo.ai.

Accenture

Accenture is a global professional services company with leading capabilities in digital, cloud and security. Combining unmatched experience and specialized skills across more than 40 industries, we offer Strategy and Consulting, Interactive, Technology and Operations services—all powered by the world's largest network of Advanced Technology and Intelligent Operations centers. Our 710,000 people deliver on the promise of technology and human ingenuity every day, serving clients in more than 120 countries. We embrace the power of change to create value and shared success for our clients, people, shareholders, partners and communities.

Disclaimer: This publication is part of the work facilitated by Nordic Innovation. The authors are responsible for its content. The use case descriptions are based on views expressed in interviews with representatives from the organizations, and not all of these have been independently verified.

Methodology

Identification of companies and organisations

The project team first compiled a list of approximately twenty initial prospective use cases based on searching media sources, event announcements, and networks. The process aimed at identifying use cases from different industries and business areas, consumers and providers of technology, with organisations spread across the Nordic countries, and with a good variation with respect to the key ethical questions being asked for the respective use cases. Thirteen of the prospective organisations were successfully contacted and agreed to take part in interviews. The precise use case definition was part of the interview methodology, and in two cases it was decided that a suitable use case for this report could not be described. Hence this report presents eleven use cases related to ethical AI across the Nordic countries.

Interviews

Interviews were carried out during May to August in 2023 with the selected Nordic organisations, later also supplemented in April 2024. The interviews had a semi-structured format, going first into the business challenge and the background of the use case, followed by questions focused on the ethical considerations and the methodologies employed to address them. All interviews were based on the same interview questions, but new ideas and themes were allowed to be brought up during the interview.

The organisations taking part in the interviews and their use cases were:

- Icelandic Language Technology Center (Iceland): Data and infrastructure for Icelandic Language Technology
- City of Tampere (Finland): Advanced family analytics project
- Saidot (Finland): AI governance solutions
- HeadAI (Finland): Skills data and synthetic data
- Zenseact (Sweden): Autonomous driving and advanced driver assistance
- AI Sweden (Sweden): AI Ethics Lab and Natural Language Understanding
- Telia Crowd Insights (Sweden): Societal mobility patterns
- Ruter (Norway): Responsible use of photos from public places
- The CLIC project (Denmark): Data platform for AI based image diagnostics
- Finterai (Norway): Federated learning for the financial sector
- DoMore Diagnostics (Norway): Cancer diagnostics solutions

Images : Cover pages and page elements: Joni Tuohimaa / Pexels: <https://www.pexels.com/photo/low-angle-shot-of-conifer-trees-against-northern-lights-16171185/>

P. 3, 8, 17, 23, 31 Freepik

