



NorwAI

NORWEGIAN RESEARCH CENTER  
FOR AI INNOVATION

# ANNUAL REPORT

» » > 2025

INNO  
VATION  
2025

NTNU





*AI Master's students Skjalg Skorstad (right) and Johan Knudsen works with Kjersti Wold, Head of Advanced Analytics at SpareBank 1 SMN, in their master's project on customer relations and churn. The project is part of the close cooperation with NorwAI consortium partners.*

# Contents

4	Entering the next phase	40	An Easy-to-Read Guide into the World of AI	69	Center Management Team	109	Which model can I trust?
6	LLMs and the question of Norway's technological sovereignty	42	Predictive AI in the age of generative models	70	NorwAI objectives	112	Highlights from Media Labs Days #9
10	AI and democracy	44	The real value of language models	75	Research strategy	114	The Forums for shared insights
14	Advancing Evidence-Based AI Policy for Safe and Trustworthy Systems	46	NorwAI Magistrate 24B Reasoning and domain specific LLMs	77	The Scientific Advisory Board	116	Recruitment PhD-students
18	Explainable AI in healthcare: When AI disagrees with expert judgment	50	From first LLM movers to delicate surgical AI	78	Publications in 2025	117	Recruitment Postdocs
23	Predicting transformer temperatures	52	Hemit HF joins NorwAI to foster safe and trustworthy AI in healthcare	83	The Innovation Advisory Board	118	The work packages
26	Sowing systematics and connections to reap results	54	What lies beyond benchmarks?	84	Students as an innovation source	124	Presentations and communication
29	Introducing: Innovation Touchpoints	58	AI for project management in megaprojects	86	PhD Defences in 2025	130	NorwAI in the media in 2025
30	How early AI adoption reshaped work, roles, and internal value	62	Two more partners in the NorwAI consortium	88	Our PhDs	133	Internal seminars
34	When technology runs faster than the organization	65	Consortium	96	NorwAI Innovate '25 A unique AI showroom	136	Personnel
38	Answering the demand for continuing education on language models	67	Members of the board	102	Best poster award The winner takes it all	144	Accounts
		68	Organization	104	NorwAI at Arendalsuka 2025	145	Web and social media
				106	Framtidsfredag – Arendalsuka on Fast-Forward		

*Cover photo: Professor Kerstin Bach presented at the fifth NorwAI Innovate Conference in September 2025.*

# Entering the next phase



**Jon Atle Gulla**  
NorwAI Center Director  
Trondheim 2026-02-26

NorwAI is now entering the last three years of its existence. Over the last five years we have initiated a range of research activities, in which PhD candidates and partners have explored and evaluated novel AI techniques to solve real problems in the real world. We see that some of these activities now come to a natural end, though there are also several cases where the technology can be further pursued commercially by partners or form the basis for new research activities in new sectors and projects.

A significant addition last year was the inclusion of two new partners from a domain that we have not previously prioritized. The healthcare domain is already a technology-driven sector, though large language models are still not heavily used in clinical practice. In collaboration with Hemit, the technology provider for the Central Norway Regional Health Authority, we are exploring relevant use cases and now preparing for the pre-training, fine-tuning and evaluation of health-adapted language models at St Olav. Hemit has chosen to join the NorwAI

consortium as part of this collaboration, and we welcome them on board. Interestingly, NorwAI has indirectly had some impact on the healthcare sector already from 2024, when the start-up Medbric was formed with a co-founder from the research center. Medbric also joined NorwAI in 2025, and we look forward to their contributions in NorwAI as well as their progress in the commercial market.

This is not to say that we do not follow up on our existing research plans. As we are gradually



seeing the results of our core technology work packages, the importance of our work packages on societal impact and trustworthiness has become apparent. New PhD candidates were recruited in both work packages in 2025, and we have over the last year been stressing the alignment with business use cases as well as encouraging the PhD candidates to collaborate across work packages. We want different research threads to come together to strengthen each individual research project, improve the exchange of results and ideas, and make the candidates more relevant to our industrial partners.

Our work on AI innovation ecosystems is also gradually taking form. The intention is to build a platform that external parties, including small startups and more mature companies, can use to explore our technologies without being forced into long-term financially demanding contracts. Central to the ecosystem is the availability of our Norwegian language models. A new 24 billion parameter model with reasoning capabilities was launched in 2025, and we are in the process of setting up an agreement for hosting our models and helping external parties to test and hopefully adopt them in their solutions. This is the technical side of it, though there are also

organizational issues that call for collaboration with other innovation-oriented initiatives like the SFU Engage at NTNU.

Today there are three institutions that pre-train language models in Norway: NorwAI, the Language Technology Group at UiO, and the National Library. At the end of 2024, all three came together in an attempt to raise awareness of the state of language model research in Norway and the benefits of continued research on Norwegian models. There is now a Language Model Cluster Norway group, and we hope that this initiative will make it easier to raise funding for further Norwegian language model training and produce models that are more suitable and better adapted to the needs in the Norwegian society.

Some PhD candidates are now in their final laps with their theses, while others like Yanzhe Bekkemoen (NTNU, associated project), Nils Barlaug (NTNU, associated project), Betül Bayrak (NTNU, associated project), Katarzyna Michalowska (UiO, HYB), Weronika Lajewska (UiS, LAP), Nolvann Bernard (UiS, LAP), Nikolay Nikolov (UiO, DATA) completed and successfully defended their PhD projects in 2025. Congratulations to all of them.

---

# LLMs and the question of Norway's technological sovereignty



**Monica Rolfsen**

Pro-rector Outreach and Social  
Responsibility NTNU

In some domains, we trust computers more than people, calculations, navigation, logistics. Why not language as well? Language, however, is more context-based, culturally embedded, and inherently ambiguous. Meaning in text depends on history, shared references, irony, power relations, and what is left unsaid between the lines.

---



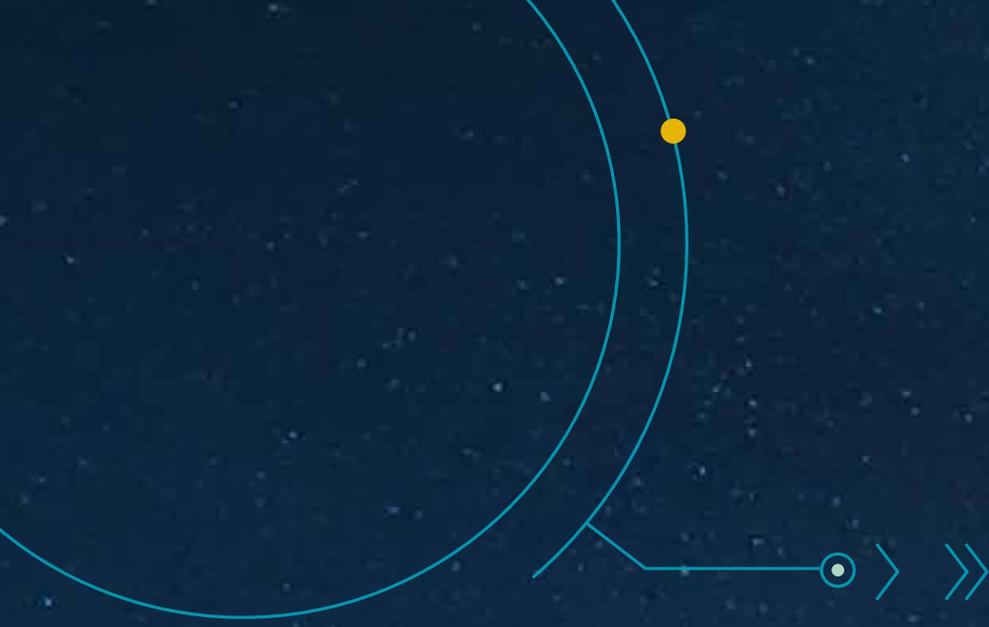
This brings us to a challenging technological dependency facing small and diverse countries like Norway. It also explains why it is important for small languages and cultures such as Norwegian, Icelandic, Sámi that research and development of new AI models are grounded in our historical and cultural contexts and premises. The work carried out by the Norwegian Research Center for AI Innovation (NorwAI) is therefore highly relevant and important.

### **Primarily English**

Most large language models are trained primarily on English-language data that reflect English-speaking cultures at their core. As a result, they mirror those cultures, their assumptions, metaphors, political norms, and value-based perspectives.

At the same time, we see that international technology companies adapt their models to local political conditions when it suits them, for example to avoid offending sensitive authorities or regimes.

- ◆ DeepSeek's treatment of Chinese political issues is an obvious example.
- ◆ Grok, which is strongly influenced by Elon Musk's philosophy, is another.



Regardless of political orientation, it is striking how much power a small number of tech companies wield over public opinion, in contrast to the influence of the vast majority of people in the world.

#### **Fundamental question**

For Norway, this raises a fundamental question of technological sovereignty. If language models become part of our digital infrastructure, used in education, administration, media, and public debate, whose language, values, and assumptions are we importing without realizing it? Training in critical and ethical awareness is therefore essential. Developing language models that adequately reflect Norwegian languages and Norwegian culture is no longer merely an academic concern. It is a strategic necessity.

Encouragingly, Norway has a strong tradition of cooperation among research institutions, industry, and the public sector, and recent initiatives in this field are promising. Combined with advances in more computationally

efficient models, Norway is well positioned to develop language technologies at the highest international level, if we choose to invest sufficient resources in research and development.

#### **Vital AI centers**

As a Pro-rector at NTNU, I believe that all our AI research centers, individually and collectively, will be vital to the development of sovereign competence, skills, and technologies. If we understand technology and its implications, we can also exercise a greater degree of control over it.

This leads to the core issue: diversity in technology is not about optics; it is about reflecting the full complexity of modern Norwegian society. Norway has become more diverse over time, and people from different backgrounds, cultures, and beliefs live side by side.

Some look forward to Christmas, while others celebrate Hanukkah or Eid.

*New language models must in some way take this diversity into account in order to be effective and relevant.*

Some thrive in the urban life of Trondheim and Oslo, others feel more at home in rural areas such as Setesdal or Kautokeino, grounded in their unique local cultures.

New language models must in some way take this diversity into account in order to be effective and relevant.

#### **Machine and Man**

The point is that machines learn from us, our biases, assumptions, and blind spots. Without oversight, these flaws can become embedded in systems that increasingly influence our lives. If research and educational institutions fail to ask critical and thoughtful questions, we risk making mistakes that may harm people. For example, Dutch authorities used AI and machine learning

to identify potential tax fraud but ended up accusing tens of thousands of innocent individuals of wrongdoing, disproportionately harming minority groups.

We must remain aware of the potential flaws of AI while also being a constructive voice, one that demonstrates better and more human-centered approaches to artificial intelligence.



# AI and democracy



## Heidrun Åm

Professor  
Department of Sociology and Political Science  
NTNU

In this text, Heidrun Åm at the Department of Sociology and Political Science, NTNU, argues that the challenges that AI pose to democracy go far beyond the question of disinformation and elections.

In recent years, rapid developments in AI have become increasingly visible. From a public perspective, generative AI in particular has emerged as a transformative technology. Alongside these developments, concerns have grown about how AI may impact citizens, society, and democracy.

When people think about democracy, free and fair elections often come to mind first. In the context of elections and AI, a common fear is that deep-fake videos or audio recordings could be used to manipulate voters. In the Slovak parliamentary elections two years ago, for example, a fabricated audio clip circulated

in which the voice of one candidate appeared to claim to have manipulated the election (Devine et al. 2024). Because of such potential, the World Economic Forum's 2024 Global Risk Report even suggested AI-driven disinformation as the greatest short-term global risk.

Against this background, the Norwegian government convened an expert group on AI and elections in autumn 2024, with the task to describe how AI challenges democracy and what impact AI can have on elections (Schia et al. 2025). Our work included to identify vulnerabilities and mapping various ways AI could be used as tool for unwanted influence,



*Consequently, big tech actors become even more powerful. By controlling infrastructure, they exercise enormous influence.*

as well as to gather experiences from how other countries work with countering election influence.

### **The super-election year**

We concluded that the super-election year 2024 (with elections in major democracies such as India, the United States, the United Kingdom, and Europe) did not turn into the large-scale generative AI election that many had feared. Still, there were instances where chatbots provided incorrect information about election dates or procedures, politicians were exposed to smear campaigns involving AI-generated pictures or videos, and some political movements, mostly on the far-right, used AI-generated content to evoke emotions, for example against immigration.

Nevertheless, AI has consequences for the information landscape in which elections take place, beyond large language models. The more people inform themselves through social media,

the more algorithmic logics can make it difficult to establish common ground for discussions. If everybody receives tailored videos and information based on their preferences, fragmentation may increase, and people might live within completely different narratives about the world they inhabit. AI-generated content adds to this, contributing to blurring the line between what is true and false.

### **Consolidation of power**

As more citizens rely on social media for information, democratic societies become increasingly dependent on big tech companies. However, their interest is tailored to generating profit, not to enabling good democratic deliberations. Indeed, big tech companies' immense power and wealth pose threats to democracy that far exceed the procedural aspects of elections.

Instead of market competition, what we see in the digital age is a special dynamic leading to



*Professor Heidrun Åm served on the Government's expert group on artificial intelligence and elections. In 2025, the group mapped the importance of AI for secure and democratic elections, and recommended measures to meet the challenges that AI may pose.*



a consolidation of power. Big tech actors have already established intellectual monopolies due to their exclusive access to data (Rikap 2021). The concentration of market and political power in companies such as Google, Apple, Facebook, X or Microsoft creates an imbalance and global dependencies that are unprecedented in human history. Never before has a state or company had so much data and insight about populations, events, and developments. This creates, politically, a source of power, and, economically, monopoly structures.

### **Unprecedented power**

The unprecedented power of big tech companies is likely to increase even further with the rise and distribution of large language models. At the moment, the best-known and most widely used large language models are provided by some of the largest tech players in the field: Google and Microsoft. These companies have in the last ten years invested heavily in infrastructure, such as cloud computing platforms, data centers, chips, and undersea cables (Luitse & Denkena 2021). As Luitse and Denkena (2021) explain, among the most important drivers of developing deep learning are not only data but computing power. Computing power requires hardware, energy, and people with the necessary expertise. At present, only a few institutions and

companies have these means to produce and offer parameter-rich AI services at large scale. They distribute their AI-services cloud-based. In this way, they retain control over infrastructure, data, and people. In this manner, their power sediments further.

Consequently, big tech actors become even more powerful. By controlling infrastructure, they exercise enormous influence. For example, Microsoft came under the suspicion in 2025 of having suspended the email account of the chief prosecutor of the International Criminal Court in the Netherlands for investigating Israel for war crimes (Clark 2025). While there are different accounts of whether and how this actually occurred, the possibility that it could happen, underscores the widespread challenges that the power of AI—including its ideology, developers, owners, and infrastructure—poses to democracy.

These are challenges that touch upon substantive democratic questions of freedom, independence, and sovereignty. Such questions go far beyond issues related to democratic elections and we need to explore and address them in the coming years.

---

# Advancing Evidence-Based AI Policy for Safe and Trustworthy Systems



## Kerstin Bach

Professor, NTNU  
Research Director NorwAI  
Director Norwegian Open AI Lab (NAIL)

Artificial intelligence is advancing rapidly, creating both transformative opportunities and significant societal risks. Policymakers are increasingly called upon to regulate AI systems that affect healthcare, education, labor markets, democratic processes, security, and the environment. Yet regulation often struggles to keep pace with technological development. The Global AI Policy Research Network (GlobAIPol) was established to help close this gap by strengthening the role of interdisciplinary research in shaping evidence-based AI policy.

---

GlobAIPol brings together researchers from computer science, law, social sciences, ethics, and public policy to support policymakers in implementing safe, trustworthy, and socially beneficial AI systems. The initiative recognizes that effective governance requires more than technical expertise; it demands collaboration across disciplines and across borders. Through annual summits, working groups, and policy

briefs, the network fosters dialogue between researchers and decision-makers and translates research findings into actionable insights.

### A shared vision

A central milestone in this work is the Roadmap for AI Policy Research, developed collaboratively at the AI Policy Research Summit in 2024. The roadmap articulates a shared vision for AI

*These discussions highlighted the challenges of translating high-level principles into actionable regulatory and implementation frameworks, particularly in sensitive domains such as policing and military applications.*



governance grounded in scientific evidence, ethical principles, and societal well-being. It emphasizes that AI policy must prioritize human and planetary welfare, accountability and transparency, inclusivity and diversity, ethical research practices, and equitable economic growth. Importantly, the roadmap underscores that evidence-based policy is not limited to quantitative metrics; it also requires critical theoretical perspectives and a strong foundation in fundamental rights.

The roadmap identifies several priority areas for impactful AI policy research. These include transboundary AI governance, recognizing that AI systems and their impacts cross national borders and require context-sensitive yet harmonized approaches. It calls for better methods to define and measure the societal benefits of AI, ensuring that innovation genuinely contributes to well-being rather

### **GlobAIpol**

The Global AI Policy Research Network (GlobAIpol) organizes an annual summit for its network. GlobAIpol is a community of practice that serves as a platform for policy researchers and professionals to advance responsible AI policy research, evidence-based insights and actionable strategies for stakeholders across academia, industry, public sector, and civil society.

AI policy research has emerged as an essential guide to navigating the complex interplay between technological innovation and societal impact. It ensures that we guide advancements in AI in alignment with ethical, legal, and social priorities.

than exacerbating inequalities. It highlights the importance of foresight and proactive regulation, enabling policymakers to anticipate technological developments instead of reacting to them. Additional priorities include developing codes of conduct for stakeholders, strengthening sector-specific analysis in high-impact domains, and improving stakeholder collaboration.

#### **The 2025 follow-up**

The 2025 GlobAIPol Summit was hosted at TU Delft's Mondai | House of AI and brought together researchers and policy professionals for two days of focused exchange. Kerstin Bach attended the summit on behalf of NorwAI and contributed to a panel on "Governance for Innovation," where mechanisms ranging from education initiatives to research frameworks were discussed as enablers of responsible AI deployment. The program also included regional discussions on AI safety and sovereignty, sessions on alternative narratives informed by foresight methods, and three deep-dive workshops addressing the contextualization of AI principles in crisis management, education, automated driving, healthcare, and warfare. These discussions highlighted the challenges of translating high-level principles into actionable

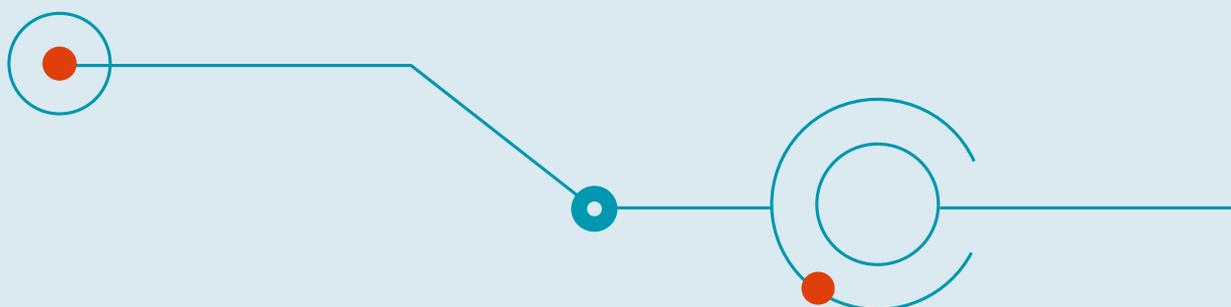
regulatory and implementation frameworks, particularly in sensitive domains such as policing and military applications.

#### **Virginia Dignum**

Virginia Dignum, Director of the AI Policy Lab at Umeå University and a member of NorwAI's Scientific Advisory Board, is one of the driving researchers behind the GlobAIPol initiative. Her leadership has been central in shaping the network's interdisciplinary and global approach to AI governance.

Beyond convening discussions, the initiative is action-oriented. The roadmap proposes establishing a community of practice, launching visiting AI policy fellowships, producing concise policy briefs, strengthening AI literacy, and continuing the annual summit series. Together, these efforts aim to ensure that AI governance evolves alongside innovation and remains grounded in transparency, inclusivity, and robust evidence.

Through its engagement in GlobAIPol, NorwAI contributes to building governance frameworks that help ensure AI systems are not only technologically advanced, but also safe, trustworthy, and aligned with societal values.





*Virginia Dignum, Director of the AI Policy Lab at Umeå University and a member of NorwAI's Scientific Advisory Board, is one of the driving researchers behind the GlobAIPol initiative.*

---

# Explainable AI in healthcare:

## When AI disagrees with expert judgment



### **Nhien Nguyen**

Associate Professor, Department of Industrial Economics and Technology Management, NTNU & WP leader INNOECO, NorwAI

When expert clinical judgment clashes with AI predictions, the tension is not only medical, it is cognitive, emotional, and deeply human. In a study co-authored by Nhien Nguyen (with Reda Hassan, Stine Rasdal Finserås, Lars Adde, Inga Strümke, Ragnhild Støen), the team investigates how clinicians respond when AI outputs conflict with their own assessments, and how those responses change when the AI becomes explainable.

---

*Nhien and her PhD team: from left Alae Ajraoui, Associate Professor Nhien Nguyen, Serinha Murgorgo and Jessica Annalena Steppe.*



The work resulted in the article “Unlocking the black box: Enhancing human–AI collaboration in high-stakes healthcare scenarios through explainable AI,” published in *Technological Forecasting and Social Change*.

### **Non-transparent AI**

Opaque AI can provoke self-doubt. When clinicians disagreed with a “black box” prediction, many began to question their own competence and, at times, revised their judgment to align with the system. This matters because high-stakes clinical work depends not only on accuracy, but also on the clinician’s ability to maintain calibrated confidence under uncertainty.

Importantly, the study shows that many clinicians do not passively accept these tensions. Instead, they engage in **interrogation practices**, asking questions, double-checking information, and reflecting on alternatives. Rather than fading when explainability is introduced, these practices intensify. Explainable systems appear to shift the interaction from silent compliance or confusion toward a more constructive dialogue: experts demonstrate greater trust in the AI system, show a readiness to learn from AI, and may reconsider or update their initial judgments when they contradict AI predictions.

### **Enabler of human–AI collaboration**

The work further highlights the role of explainable AI in clinical judgment. When the system provides signals about how confident it

is, clinicians can make more nuanced decisions about when to rely on AI support and when to lean more heavily on their own expertise. In this way, explainability functions not simply as “transparency”, but as an enabler of human–AI collaboration, supporting trust while strengthening judgment formation.

At the same time, the study advances a cautionary insight that is especially relevant for clinical environments and for healthcare organizations more broadly: explainability is not a cure-all. Over time, overreliance on AI outputs may erode reflective practice, reduce peer-to-peer consultation, and weaken organizational learning. The implication is clear: explainable AI can improve collaboration and decision quality, but only if healthcare settings also protect the social and professional practices that sustain critical reflection.

### **AI in organizations**

While healthcare illustrates AI’s high-stakes impact in concentrated form, the same fundamental dynamics - opacity, uncertainty, shifting judgment, and evolving collaboration - also play out across organizations. Organizations are increasingly exploring how AI redefines how they create, deliver, and capture value. Research shows that AI can drive efficiency and effectiveness, but it also lays the groundwork for new organizational capabilities. Crucially, adopting AI is not just a matter of installing technology or plugging it into existing business processes.

For decades, socio-technical approaches to information systems assumed relative technological stability. AI challenges that assumption: it is evolving, pervasive, and often opaque. As a result, organizations must rethink how they implement AI responsibly and how they collaborate with it in ways that deliver sustainable business value and societal benefit.

### Interconnected research

Against this background, Nhien Nguyen and her team of PhD candidates - Alae Ajraoui, Jessica Annalena Steppe, and Serinha Murgorgo - are conducting interconnected research under the broader project "AI in organizations".

Together, these projects examine how organizations implement AI responsibly and foster human-AI collaboration that enhances creativity. A defining feature of this research stream is that the doctoral projects inform each other and cannot be treated as separate efforts: organizations cannot pursue responsible AI without understanding implementation in practice, and they cannot foster effective human-AI collaboration without continuous learning as AI is implemented. Successful AI deployment therefore depends on advancing these dimensions in parallel.



**Alae Ajraoui**  
PhD Candidate, NTNU

### AI implementation and continuous learning

Alae Ajraoui's PhD research (supervised by Nhien Nguyen and Alf Steinar Sætre) investigates how organizations implement AI and how they continuously learn throughout this process.

Based on a systematic literature review, the study synthesizes existing knowledge on key implementation challenges and the response

mechanisms organizations adopt to address them. In addition to this, a multiple-case study of six Norwegian companies - recognized with the Best Poster Award at the NorwAI Innovate Conference 2025 - examines how AI implementation unfolds in practice across strategic, tactical, and operational levels.

The study reveals that realizing the potential of AI and related technologies requires early sensing of emerging threats and opportunities, seizing both internal and external resources, and reconfiguring teams and workflows to effectively leverage AI and generate significant business value. Together, the findings indicate that *AI implementation is not a one-off integration event, but a dynamic, multi-level process* that requires ongoing adaptation, coordination, and learning. Without such transformation, organizations risk remaining stuck in perpetual experimentation, unable to translate AI's promise into sustained and profitable growth.



**Jessica Annalena Steppe**  
PhD Candidate, NTNU

### Human-LLM collaboration and creativity

Jessica Annalena Steppe investigates how employees work with generative AI, specifically large language models (LLMs), to augment creative performance.

Since creativity underpins innovation and competitive advantage, her research addresses a strategically important question: how can AI support creative work, and how can managers design workflows that enable effective human-LLM collaboration?

While prior studies suggest that LLMs can augment creativity, the mechanisms through which this happens are still not well

understood. In her first qualitative study, Jessica addresses this gap by describing two primary ways employees interact with AI LLMs: as an *inspiration source* for generating ideas and as a *sparring partner* for developing and refining ideas.

Idea generation helps employees quickly explore many possible directions, while idea development helps employees deepen and refine the ideas, they already have by revealing alternative perspectives and challenging assumptions.

The study further shows that employees' prompting skills and domain expertise strongly shape how effectively they collaborate with LLMs, and, in turn, the quality of creative outcomes. Used deliberately as a sparring partner, LLMs can help employees build richer understanding and arrive at solutions perceived as more creative, meaning both novel and useful.

Her research papers are collaborative work conducted together with her supervisors Nhien Nguyen and Marta Morais-Storz, and have been presented at several internationally renowned conferences, including the Academy of Management Meeting in Copenhagen and the Strategic Management Society Conference in San Francisco.



**Serinha Murgorgo**  
PhD Candidate, NTNU

### Responsible AI in practice over time

Serinha Murgorgo's project addresses how organizations operationalize responsible AI (RAI) and learn to use AI responsibly over time.

While many principles, frameworks, and

guidelines describe what responsible AI should look like, they often remain abstract and vague, leaving the "how" of implementation unclear across contexts.

Serinha's research responds with a systematic review of empirical studies and grey literature across sectors. The findings propose a practice-oriented framework emphasizing interdependence and iteration across three reinforcing dimensions. Successful operationalization emerges as dynamic governance: *structures* provide mandate, *procedures* create repeatable routines, and *relational practices* sustain responsibility through culture.

This review was conducted in collaboration with her supervisors Nhien Nguyen, Patrick Mikalef (NTNU), and Tita A. Bach (DNV).

Her ongoing work builds on this dynamic perspective by reframing responsible AI as a question of how organizations learn to embed ethical principles into everyday practice: making sense of responsibility, responding to uncertainty, and continuously adjusting practices over time.

### A connective thread across our WP research

Across healthcare and organizational contexts, a shared theme emerges: *AI changes how humans form judgments, coordinate work, and learn over time*, especially when systems are opaque, evolving, and embedded in everyday practice. The healthcare research shows how explainability can transform output conflict into productive interrogation, while also warning that unchecked reliance may erode reflection and learning. The organizational research extends this logic outward: implementation, creativity, and responsibility are not separate challenges, but tightly linked capabilities that must be developed together for AI to create durable value and societal benefit.

*«... ensuring they live as long as possible is a major priority for grid operators such as Statnett.»*

---

# Predicting transformer temperatures



> **Signe Riemer-Sørensen**

Research Manager and Senior Research Scientist  
SINTEF Digital, WP Leader HYB NorwAI

Transformers are among the most essential components of the electrical power grid. Their job is simple, but critical: they convert electricity between different voltage levels so it can travel long distances and still be safely delivered to homes, factories, and infrastructure. Because transformers are extremely expensive and time-consuming to replace, ensuring they live as long as possible is a major priority for grid operators such as Statnett.

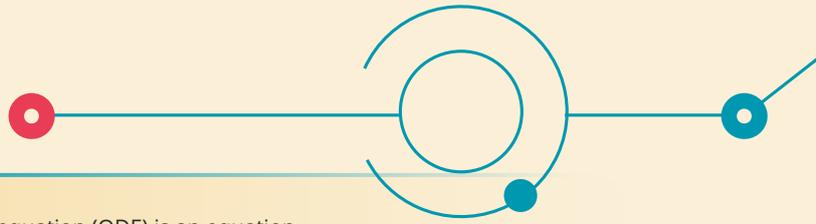
---

One of the biggest factors determining transformer lifetime is **temperature**. A heavily loaded transformer runs hotter, and high temperatures speed up the degradation of insulation materials inside.

Over time, this shortens the transformer's remaining life. To keep temperatures under control, transformers are cooled by oil that

circulates through the windings, picks up heat, and transports it to a heat exchanger mounted on the outside of the tank.

Because lifetime depends so strongly on temperature, it is extremely valuable to predict how hot a given transformer has been or will become under different combinations of electrical load and outside temperature. This



*ODE: Ordinary Differential Equation:* An ordinary differential equation (ODE) is an equation that describes the relationship between an unknown function of one independent variable (e.g. time or position) and its derivative. While ordinary equations give numbers as solutions, differential equations give a function as a solution, and they are used to describe dynamic changes in time and space.

allows grid operators to estimate the current condition of the transformer and how much “lifetime cost” is associated with planned use.

### To start with

To obtain such a model, we could start from physical properties. However, Statnett have many different types of transformers, and physical models would require knowing the details of each individual transformer.

The other extreme would be to train a machine learning model for each transformer with the caveat that the model can only be used on conditions (outdoor temperature and power) already seen in the training data.

In NorwAI Work Package Hybrid AI Analytics (HYB) we combined the approaches. Many physical systems, including heating and cooling processes, can be described using *ordinary differential equations (ODEs)*. These equations describe how a quantity changes over time. A transformer’s thermal behavior can be simplified to a single ODE for winding temperature, which depends on ambient air temperature, electrical load, and a number of unknown physical parameters.

If we had a perfect expression for the function that governs this ODE, call it **F**, we could compute temperature forecasts using standard numerical solvers.

But what if we don’t know the correct form of **F**? That is where *Neural ODEs* enter the picture.

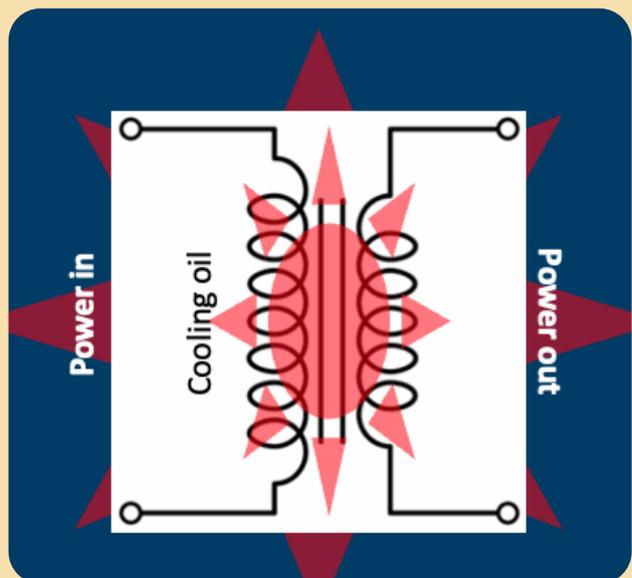
### A neural model

A Neural ODE is a machine-learning model that *learns the function F directly from measurement data*. Instead of guessing a formula, we let the model infer how temperature changes over time by observing historical transformer behavior.

Master student Berk Hadzhamolla worked with data from Statnett to test this approach on real transformers. Some cases led to almost perfect models while others performed less convincingly, often due to limited training data or inconsistencies that suggested the transformer broke the assumption of behaving like a simple ODE system. In all cases, further investigations of the poor models led to identification of data problems, enabling a new way for Statnett to identify potential anomalies.

Overall, Neural ODEs show great promise. They highlight how **data quality and cleaning are essential**, and they open the door for applications such as anomaly detection. Ongoing work includes adding more transformers, improving data quality, and incorporating more physics, such as latent layers of oil or time-delayed heat diffusion, to build even more realistic models.



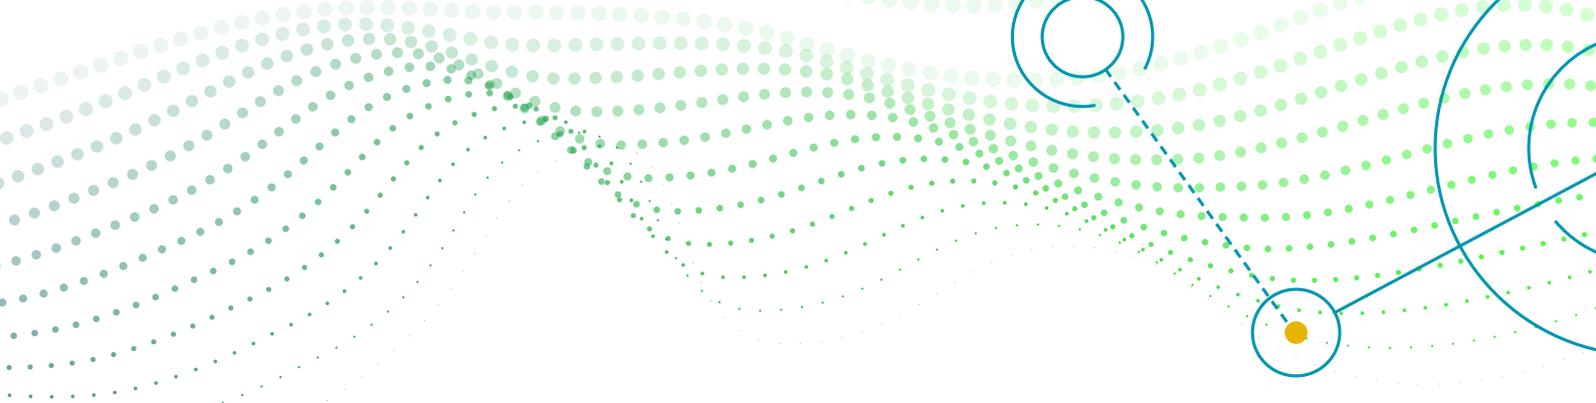


*Lifetime of transformers depends strongly on temperature controls in the cooling system.*

---

# Sowing systematics and connections to reap results

*Lisa S. Græsli at SINTEF keeps track of the many use cases that form an important part of NorwAI's systematic innovation work.*



**Lisa S. Græsli** has been assigned a key role in NorwAI's innovation ambitions. As a researcher and senior project manager at SINTEF Digital, she works daily with innovation projects - while also being responsible for systematizing and following up on the prerequisites that make innovation possible in the research center's ecosystem.

---

Sowing and reaping innovation is not a game of chance. In NorwAI, a foundation of structure, relationships and systematics is built that will facilitate lasting results.

It has long been recognized among organizational scientists that innovation has the greatest potential to be created across disciplines with specialized knowledge. But it also requires understanding that knowledge is both a source of and a barrier to innovation. There is significant potential in interdisciplinary collaboration, but in practice it can be difficult to get experts to work together in ways that realize this.

### **Boundary spanners**

Lisa S. Græsli, who holds a PhD in industrial economics and technology management from NTNU, is particularly interested in the role of so-called "boundary spanners" or bridge builders: those who have the ability to build bridges between disciplines, actors and sectors that otherwise do not connect naturally because they use different professional languages, have different understandings, work at different paces and have different interests and goals.

-Without these connections, innovation, learning and implementation often come to a standstill, says Lisa S. Græsli.

On the one hand, continuous knowledge exchange between partners is required – to develop common language and interests through frequent dialogue and sharing of insights and challenges. This is crucial for building understanding and identifying new opportunities together.

On the other hand, innovation is a dynamic and unpredictable process. Challenges arise along the way where you don't quite know how or when you will reach your goal. This is precisely why bridge builders become crucial – they translate



Dorothy Leonard's (1995) statement that most innovation happen at the boundaries between disciplines or specializations tells us that working across boundaries is a key ingredient of competitive advantage, but also why innovation proves so difficult to create and maintain. The growing research on knowledge in organizations underscores this challenge by recognizing first the "knowledge boundaries" (Brown and Duguid 2001) between specialized domains and second that knowledge is "both a source of and a barrier to innovation" (Carlile 2002, p. 442).



needs across disciplines and organizational boundaries, and connect the right expertise at the right time to create progress even when complexity increases.

Boundary spanners have unique interactional expertise that is developed through dialogue over time. One goal is therefore to develop as many bridge builders as possible through knowledge sharing and joint problem solving via different types of arenas, she says.

### **Arenas for bridge building**

The INNOECO (AI Innovation Ecosystems) work package, in which Lisa S. Græslie participates, has a set of meeting places available. Everything from the annual NorwAI Innovate conference, the biannual NorwAI Forums with a focus on research in the spring and innovation in the fall. And most recently the new arena NorwAI Innovation Touchpoints, where partners are invited to share results or challenges for interdisciplinary discussions and knowledge sharing.

In addition to identifying bridge builders and actively encouraging the generation of new cross-links, the center focuses on systematic follow-up of use cases to support research and innovation activities.

– Through a systematic approach to use cases and innovation in the center's work, we try to understand how use cases develop and where

the challenges lie, as well as where there are possible new connections or opportunities, she says.

### **Breadth that strengthens and challenge**

The use cases range from advanced modeling of physical systems – such as predictive maintenance, anomaly detection and hybrid digital twins – to the use of large language models in the finance and media sectors, which are about personalization, content understanding, automation, data governance and the safe use of generative AI. The breadth is a clear advantage: it allows NorwAI to work on relevant issues across industries, technologies and application areas. But it also makes coordination more complex, because the challenges and maturity vary greatly between the use cases.

– Is there room for improvement?

– Yes, absolutely. It is no easy task to keep track of everything that is happening at the various partners. People are replaced, organizations change direction, and new priorities arise continuously. At the same time, the AI field is moving so quickly that development cycles must be carried out at a pace that often contrasts with the research's need for thoroughness and time. This makes it extra important to have a common platform that can keep track and ensure that we are moving in the same direction across environments, says Lisa S. Græslie.



# Introducing: Innovation Touchpoints

Moving into the final years of NorwAI, focus on innovation is key, in order to turn research results into new AI-driven products and services that create growth and societal benefits. And an ecosystem for open innovation such as NorwAI needs informal and low-threshold meeting places and arenas for sharing and inspiration, for encounters between academia and business.

While the innovation-driven autumn editions of NorwAI forum, as well as the NorwAI Innovate conferences provide these platforms, our industry partners indicated that they would like to have more of these meeting places – and not necessarily only physical.

There was already a series of NorwAI webinars, usually organized around the research work in the center, but at times slightly too technical for the businesses to follow. What if we “stole” these online meeting slots once in a while, to presents some finished, ongoing, or future innovation activities by the industry partners?

Innovation Touchpoints in NorwAI are a series of online/physical meetings bringing together the industry and research communities in the center in order to **foster dialogue, discussions and collaboration**.

Driven by the industry partners, we will highlight how AI is being rolled out in operational contexts. Innovation Touchpoints are open to all NorwAI stakeholders.



This is when the idea of the NorwAI Innovation Touchpoints was born.

## Informal sessions

An informal and short online session to be run three to four times a year. So far, the center has seen three of these sessions, with contributions from Statnett, SMB, Telenor, and DNV. The most recent edition, featuring Andreas Hafver from DNV laying out their planned research on trustworthy AI (related to the WP TRUST), was also a well-appreciated introduction into the NorwAI Forum two weeks later.

Participation – and especially researcher participation – has been increasing; as has the level of discussions and questions over the three instances, encouraging us to keep organizing further Touchpoints in 2026, as well!

*Andreas Hafver is teamleader for emerging technologies in DNV.  
Photo: NorwAI*



# How early AI adoption reshaped work, roles, and internal value



**Christian Printzell Halvorsen**

CEO  
Vend



Integrating AI into core operations is driving a shift from execution-focused roles and support functions toward builder roles and internal product teams, with leadership and capability development as critical enablers.

Since the release of ChatGPT in late November 2022, generative AI has evolved from a novelty to a general-purpose technology. Beyond productivity gains and automation, early organization-wide adoption is reshaping the fundamentals of work and organizational structure.

The observations in this article are grounded in practical experience gained through Vend's early AI adoption initiatives.

## **Eye-opening moments**

At Vend, we made an early strategic investment in generative AI by giving every employee access to ChatGPT. In 2024, this was complemented by an organization-wide AI training, with management incentives tied to participation. The aim was simple: to raise baseline AI-competence across the organization.

As experimentation spread, the initiative produced a series of unexpected, eye-opening moments. Employees quickly discovered that



*“Many roles will shift from being doers to becoming orchestrators of AI agents or builders of tools.”*

**CHRISTIAN PRINTZELL HALVORSEN,  
CEO AT VEND**

AI was not just a faster way to write text or summarize documents. Instead, it revealed a deeper shift in how we structure work: tasks were broken into components, execution was increasingly handled by AI, and results were reassembled through human judgment. What began as a productivity experiment exposed a structural change in how work itself could be designed.

**From thinkers to builders**

The most significant outcome of this experimentation was a transformation of the employee role. Many roles began shifting from task execution to orchestrating AI agents and building tools.

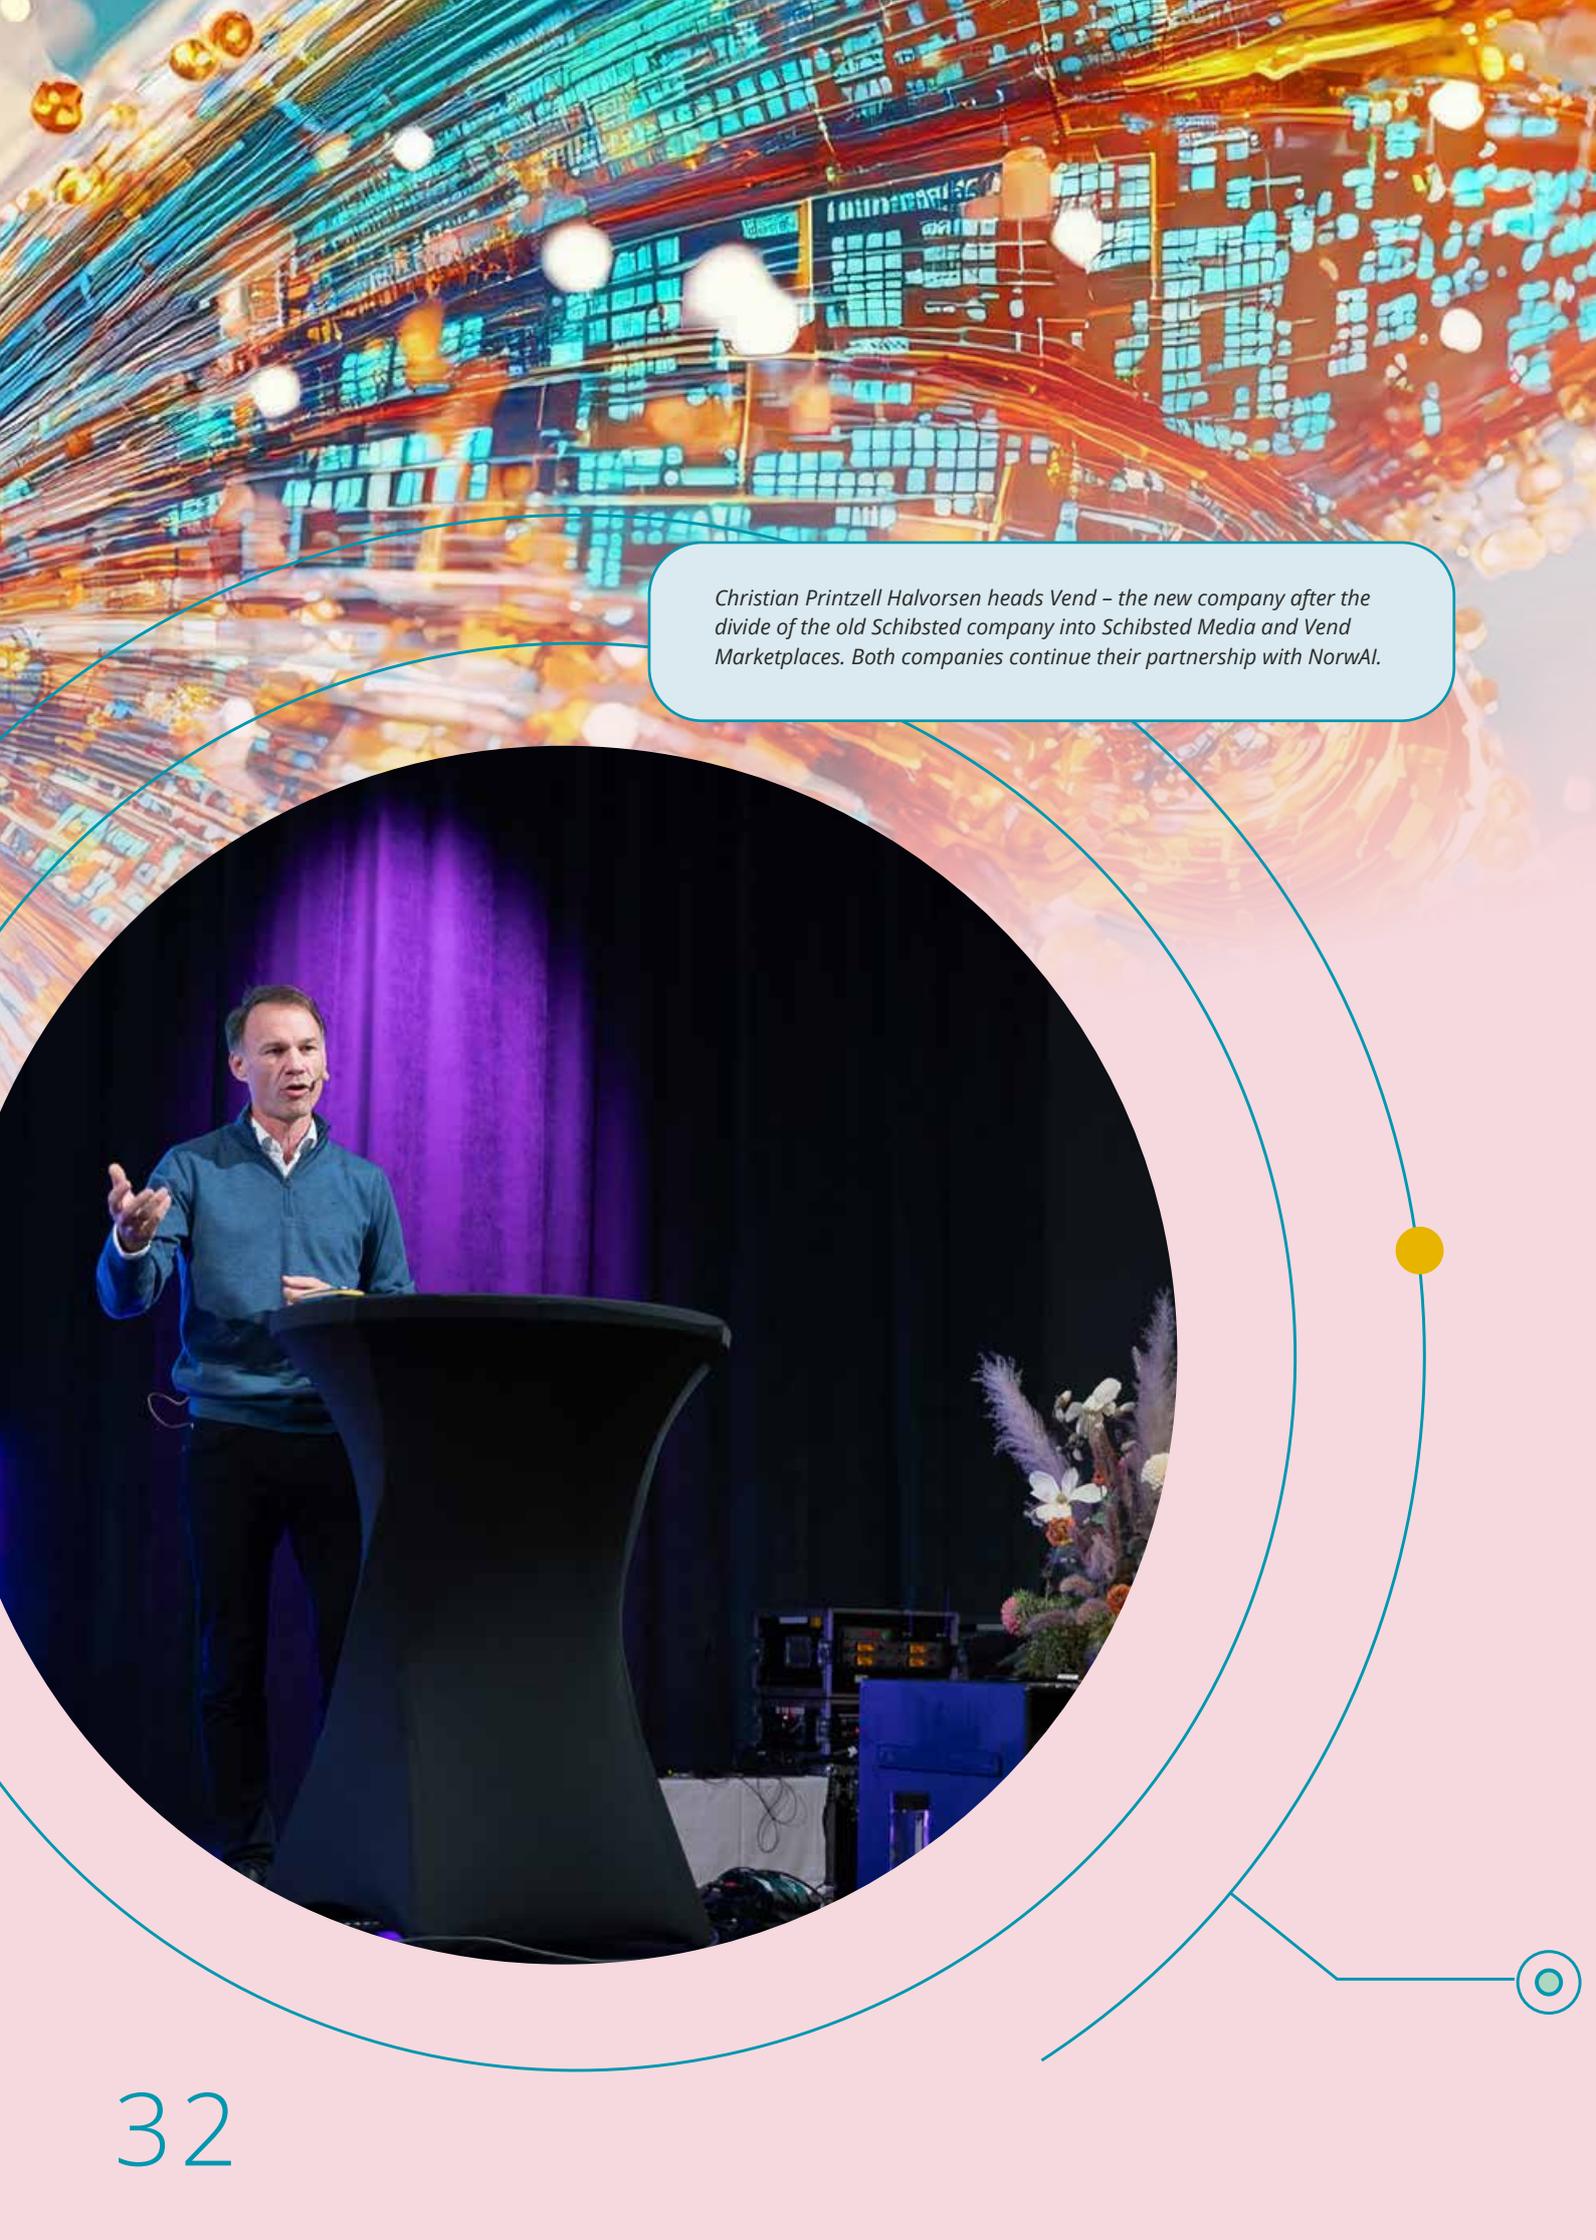
In the previous model, value was created through manual digital labor. In the emerging

AI-enabled model, value is created through designing workflows in which AI handles execution.

This marks a transition from individual task optimization to systemic problem solving. Human judgement remains central, but its focus shifts from doing work to shaping how work is done.

**Illustrative case 1: The Vend Builder program**

This shift is not theoretical. It is visible in initiatives such as our Builder program, which empowers employees to independently develop AI-based solutions for themselves or others. Builders are employees outside technical roles who leverage AI to create tools that enhance productivity or improve products.



*Christian Printzell Halvorsen heads Vend – the new company after the divide of the old Schibsted company into Schibsted Media and Vend Marketplaces. Both companies continue their partnership with NorwAI.*





By lowering the barrier to solution development, we effectively decentralized innovation. Instead of routing all improvements through centralized teams, domain experts closest to the problem could design and deploy solutions directly. This significantly accelerated experimentation and learning, while also reinforcing the emerging builder-oriented role identity.

### **The rise of internal product teams**

As AI became embedded in daily operations, it also indicates a reevaluation of organizational structure. Functions such as HR, Finance, and IT have traditionally been framed as support functions or cost centres, with a primary role of enabling others to create value.

As AI moves into core operations, this distinction breaks down. These functions will increasingly build AI-based services that directly shape how work happens. If HR develops tools for talent matching or churn prediction, or if Finance build AI agents to automate forecasting, they are no longer merely supporting the business. They are

delivering internal products that meaningfully influence business outcomes.

This requires a shift in mindset, capabilities, and governance. Internal functions must adopt product thinking, including user-centric design, continuous improvement, and clear ownership.

### **Illustrative case 2: The Vend Assistant**

The evolution of our Vend Assistant for employees illustrates this transformation. What began as a simple chat interface evolved into a context-aware agent embedded in our workflows. By integrating organizational knowledge and processes, the assistant moved from generic support to a core operational capability. At the same time, the team building the assistant evolved into product-thinking and now focuses on and takes decisions based on the product KPIs for the assistant.

This case demonstrates how internal AI tools can mature into strategic assets when treated as products rather than utilities.

# When technology runs faster than the organization



**Thomas Nikolai Blekeli**  
Strategic Advisor  
Norwegian Broadcasting Corporation (NRK)

When the generative AI-wave first hit, NRK, like many large organizations, faced a dilemma. We saw tools changing the game overnight, yet as a public broadcaster, our most important currency is trust.

The internal curiosity was enormous, but so was the uncertainty. How could we allow our employees to explore this immense potential without risking our journalistic integrity?

This challenge highlighted the gap known as "Martec's Law": Technology changes exponentially, while organizations change logarithmically.

At NRK, bridging this competence gap required more than just emailing new guidelines. It demanded a completely new arena for learning.

### **The sandbox strategy**

To provide a safe environment, we built NRKGPT to serve as an internal gym where staff could train within a secure framework. To ensure responsible use, we introduced a mandatory "AI Driver's License".



### Lessons learned by NRK

Looking back at this journey, we identified three hurdles every organization must address:

- 1 The trust challenge: The risk of hallucinations requires us to relentlessly teach that AI is a drafting tool, never a replacement for fact-checking.
- 2 The time barrier: Even with motivation, finding time to learn in a high-pressure environment is a constant struggle.
- 3 The Prompting Skill Gap: Prompting is a new literacy. There is a vast difference between using a tool and using it well.



### Martec's Law

Technology changes exponentially (fast), yet organizations change logarithmically (slow).

Management must strategically choose which technological changes to embrace, given the highly constrained bandwidth for absorbing organizational changes.



this change gap widens over time, eventually requiring a "reset" of the organization

*By prioritizing mindset over tool proficiency and providing a safe playground, we have moved the organization from being passive "answer-seekers" to active "co-pilots", says NRK strategic advisor Thomas Nikolai Blekeli.*



*Co-presenter Egil Aslak Hagerup, NRK, talked about success*

Our philosophy was simple: "Welcome everyone, qualify everyone."

Rather than being gatekeepers, we became facilitators. By keeping the technical threshold low but the ethical awareness high, we encouraged wide editorial adoption while striving to ensure every user understood the fundamentals.

### **Scaling the knowledge**

Strategy is nothing without execution. The practical challenge was daunting: How do you train thousands of busy employees?

We began with a dedicated team and a pilot group within the news- and sports divisions. We designed a hands-on course focused on practical application and ethical dilemmas.

Crucially, we used these sessions to gather bottom-up feedback. This revealed that our existing guidelines were struggling to keep pace

with usage, allowing us to update our editorial standards based on real-world experience rather than top-down theory.

To scale further, we adopted a "train the trainer" model. Teaching colleagues from different departments to become AI educators allowed us to customize the training for thousands of employees.

### **Validating the impact**

Did it work? We were fortunate to have Janne Seime Siler and Kathrine Frid from BI Norwegian Business School use NRKGPT as a case study for their master's thesis in 2025.

Their research validated our approach: 71% of users reported gaining new skills and knowledge through the tool. As one journalist put it: "It helps me tremendously to do my job faster and more accurately. And it comes up with things I wouldn't have thought of on my own".

---

# Answering the demand for continuing education on language models



Martine Torland, Photo: Øistein Falch, NTNU Videre

*– The first run-through has been very rewarding. It is demanding to teach a course on a topic that does not have any good textbooks.*

**JON ATLE GULLA ON FORMING A COURSE ON  
INNOVATION IN LARGE LANGUAGE MODELS**



**Martine Torland is young and has a recent master's degree. Nevertheless, she experienced great benefit from the continuing education course Innovation in Large Language Models.**

---

Martine Torland is a manager and advisor at the newly launched advertising agency Bennett and has taken the NTNU Continuing Education and Professional Development course held by NorwAI.

### **AI as a sparring partner**

*How does AI affect an industry that offers services such as creative copywriting and content development?*

- For the creative industry, AI has primarily been a powerful tool, not a replacement for strategic and creative expertise. I think it's more about us being able to do things more efficiently now. It certainly is for us like most other, says Martine Torland.

*How does artificial intelligence actually affect the everyday work of those of you who make a living writing creative texts?*

- AI has a number of uses! Like being a source of knowledge, inspiration or writing advertising texts in the customers' "tone of voice". We free up time to follow up on customers and can spend more time developing insight-based communication that is in line with the brands we work with, says Martine.

She also points out that good use of AI tools requires brand understanding. You need to have the right "input", to ensure that the "output" is in line with what you want to convey and that the result is "on brand".

### **From theory to practice**

It wasn't long ago that Torland herself was a master's student. She talks about a clear difference between being a full-time student to now studying with work experience:

- When you study without work experience, there's not as much to focus on in the curriculum. I notice that understanding improves when you can connect what you learn to your own everyday work, Martine says.

The course Innovation with Large Language Models was developed in the fall of 2024 by the research center NorwAI. The first session was held the following spring. Professor and director of NorwAI Jon Atle Gulla says that continuing education on language models has been in demand both from companies, but also from external actors who have contacted them.

- NorwAI had not initially planned to provide continuing education courses, but then an opportunity opened up by applying for funding from the Directorate for Higher Education and Competence.

This allowed them to offer teaching, both physically and digitally, as early as spring 2025. The next version is to take place in the first half of 2026.

*Original article, in Norwegian, by Øistein Falch for NTNU Videre.*

# An Easy-to-Read Guide into the World of AI



**Rolf Dyrnes Svendsen**  
Communication Manager  
NorwAI



Norwegian organizations need simple and practical steps to adopt AI assistants safely and efficiently. After a swift and focused process lasting two to three months in the spring of 2025, committee chair Jon Atle Gulla was able to present a user-friendly guide to the Minister of Digitalization and Public Governance Karianne Tung on June 16. The guide was designed to help companies and organizations navigate the world of artificial intelligence.

- It was important for us to develop a concrete and user-friendly guide that Norwegian organizations can actually use. The guide serves as a ticket to start using AI and to understand how to realize its benefits. It translates theory into practice so that more Norwegian organizations can safely take advantage of AI assistants, said Professor Jon Atle Gulla, head of the expert group, during the handover.

- Artificial intelligence offers great opportunities for innovation, growth, and increased efficiency. With this guide, we aim to lower the threshold for Norwegian organizations, enabling more

of them to safely adopt AI assistants and strengthen their competitiveness, said a thankful Minister of Digitalization and Public Governance Karianne Tung in her remarks of appreciation.

### **How do we get started?**

AI assistants are becoming as commonplace as email and online meetings, but they require new skills, routines, and responsibilities. Many Norwegian organizations, both public and private, are asking: "How do we get started without taking unnecessary risks?"



*Minister of Digitalization and Public Governance Karianne Tung with Professor Jon Atle Gulla at the handover of the AI guide.*

The guide was developed by a broadly composed expert group and incorporates insights from both the public and private sectors. It is designed to provide both knowledge and practical advice:

- ◆ A tool to promote innovation and growth
- ◆ A simple recipe with concrete steps to get started
- ◆ A checklist for safe and responsible use
- ◆ Relatable examples
- ◆ A framework for legal application

Professor Jon Atle Gulla emphasized that the guide focuses specifically on the use of AI assistants. It does not address AI in general, nor does it cover more complex AI agents or autonomous systems capable of making decisions independently:

- Those with expertise in such systems don't need a guide, said Gulla, adding: - I hope this guide will spark momentum in the areas it covers, allowing us to keep up with emerging solutions over time.

---

# Predictive AI in the age of generative models



**Anders Løland**  
Research Director  
Norwegian Computing Center (NR)

The current major AI efforts revolve around generative methods. Does this mean that predictive AI is less important? Yes, no, and maybe!

---

## Do you want to generate new data or predict something?

*Generative AI* is a subfield of AI that is used to generate images, videos, text, audio, software code or other types of data. These models learn the underlying patterns and structures of their training data and use them to generate new data in response to input, which often takes the form of natural language prompts.

*Predictive AI* on the other hand uses historical training data to find patterns and forecast or classify future events, patterns, trends, or behaviors.

*Generative AI* is typically trained on large datasets crawled from the web, while predictive AI often is trained on smaller, more carefully curated datasets – although it can use large datasets as well.

## The massive AI investments go to generative AI

I contend – and this is difficult to quantify – that predictive AI presently offers greater practical value in industrial settings. However, this might change fast: Due to the current massive investments in AI – by for example OpenAI and NVIDIA – mostly into generative AI, the money spent on AI is obviously higher for generative AI.

## AI subfields

For the benefit of this discussion, I will limit AI to machine learning (ML). I will for example ignore expert systems. Simplifying further, I divide ML into two subfields:

- ◆ Classifiers and statistical learning (including decision trees, regression models, Gaussian processes, Bayesian networks, Support-vector machines, random forest, boosting)

*General-purpose models can be sufficient, but tailored models can in some cases be much better – especially if they incorporate domain and/or physical knowledge.*



◆ Deep learning (including neural networks, large language models)

### **The grey zone**

There is a grey zone here: Some models can be used for both predictive and generative tasks. Think for example of an LLM. It can generate a new text, or it can be used to classify the sentiments of a text.

This grey zone is not limited to deep learning. Almost any traditional classifier/statistical learning model can also be used to simulate or generate synthetic data. And if I condition on an appropriate combination of features when generating synthetic data, I have produced a probabilistic prediction or a distributional forecast.

This is not to say that any model will produce results that are equally good. General-purpose models can be sufficient, but tailored models can in some cases be much better – especially if they incorporate domain and/or physical knowledge.

### **Validation can be trickier**

Any AI model should be tested and validated on independent data. For predictive models, this is fairly straightforward: The models are

trained and tested on the same classes of data or problems. An example is to predict the temperature the next day, and evaluate the model's performance on temperature data that have not been used in the training process.

For generative AI, like LLMs, validation can be much trickier: The space of possible test situations and evaluation metrics is almost endless, and since those models typically are trained on large chunks of the web, determining what they have "seen" during training is far from trivial. A generative model can perform well on one task, but not so well on the next. An all-purpose evaluation metric will never exist.

There is a famous paper titled "Attention is All you Need", which successfully introduced the transformer mechanism for deep learning in 2017. Perhaps evaluation we can trust is needed more now?

### **Is super intelligence coming?**

If you believe in the advent of super intelligence, the massive investments into generative AI can make sense. In the meantime, some enterprises would in my opinion benefit from redirecting their attention to generative AI back towards predictive AI.

---

# The real value of language models



**Jon Atle Gulla**  
Professor  
NorwAI Center Director

Since 2021 NorwAI has trained a number of Norwegian language models that are now being tested and used by organizations across the country. Different architectures, configurations, and datasets have been tested to expose the underlying technology and to train the best Norwegian models given the research center's constraints and priorities.

In the MIMIR project from 2024, where NorwAI collaborated with the Language Technology Group at the University of Oslo and the National Library, a series of Norwegian test datasets were developed for the automatic evaluation of 17 Norwegian language models.

In the TrustLLM project, of which NorwAI is a partner, the EuroEval framework is used to compare large language models across Europe on aspects like summarization, knowledge representation, common-sense reasoning,

simplification, European values, instruction-following and bias detection. Leaderboards for the European languages rank models and are often consulted when the quality of various large language models are discussed. These automatic evaluation metrics are valuable for the iterative development of models with different configurations and datasets. They provide early feedback on important properties and guide the development of increasingly more capable models.



However, when large language models are incorporated into real applications, it is not given that the models on top of the leaderboards are the most suitable ones. One thing is that there may be certain requirements in terms of model size, data sensitivity and licenses that rule some models out.

A more fundamental problem with these generic models is that they do not necessarily reflect the characteristics of the application domain or the tasks they are intended to solve. For this reason, we have now started using post-training like fine-tuning and reinforcement learning to adapt generic models to better understand the domain terminology and carry out the desired tasks. From an evaluation point of view we need to evaluate the post-trained models against the generic ones, and in this process it is not obvious that the best generic language model is the best starting point for the adapted models. Also, the automatic evaluation scores are aggregations of sub-scores from tests like summarization and text understanding, though for a particular domain or task the weight of each sub-score may vary substantially from the generic case.

In principle, the quality of large language models needs to be assessed on the basis of the quality of the resulting applications. In 2026 we will fine-tune language models for the healthcare

and the financial sectors, as well as for project management. The fine-tuned models will be evaluated with automatic metrics, though we also intend to observe the applications in use and define qualitative methods with real users evaluating the applications in realistic settings. We want to know to what extent these applications affect the quality and efficiency of their work, and how these language models may be further adapted to help users with their daily tasks.

There is a gradual realization that not even human evaluations like that are sufficient to uncover the true value of language models. Even if each individual's productivity is increased, there may be unforeseen consequences at an organizational level that call for a more holistic approach to LLM adoption. As noted by one of our partners, a realistic consequence of everybody's use of large language models may be reduced collaboration among colleagues, with less knowledge transfer and lack of constructive human feedback. To assess the real value of LLM-driven applications, thus, we may need to observe the organization over some time to see their impact on the organization as a whole.

In 2026 we want to explore further how language models can be efficiently adapted to real applications and provide real value to end-users and organizations.

# NorwAI Magistrate 24B Reasoning and domain specific LLMs



**Lemei Zang**  
Researcher, NorwAI



**Peng Liu**  
Researcher, NorwAI

NorwAI completed its latest Large Language Model late autumn 2025. The NorwAI Magistrate 24B Reasoning is the biggest LLM built in Norway, and fits into the research centers plans to develop specialized models. We asked Lemei Zhang and Peng Liu for a Q&A session about the new model.

**Q:** *What is NorwAI-Magistral-24B-Reasoning?*

**A:** NorwAI-Magistral-24B-Reasoning is a sophisticated language model built specifically for Norwegian language understanding and generation. With 24 billion parameters, it represents a significant advancement in Norwegian AI capabilities.

The model was developed in two main stages. First, it underwent continued pretraining on approximately 88 billion Norwegian tokens, drawn from a wide range of high-quality sources, including partially healthcare-related data. This phase strengthened the model's understanding of Norwegian language structure, vocabulary, and factual knowledge. Next, the model was fine-tuned to enhance its reasoning abilities, teaching it not only to generate fluent text, but also to reason more carefully and efficiently when solving problems.

What truly sets NorwAI-Magistral-24B-Reasoning apart is its advanced reasoning capability with dynamic depth adaptation. Much like a human expert, the model adjusts how deeply it reasons depending on the task. For simple questions, it responds quickly and directly. For moderately complex problems, it engages in short, structured reasoning. For more difficult cases, such as medical decision-making, it can perform longer, step-by-step analysis. This flexibility makes it especially valuable in healthcare contexts, where both speed and careful reasoning matter, says Lemei and Peng.

Equally important is the model's strong linguistic foundation. Trained extensively on Norwegian Bokmål and Nynorsk, it understands local language nuances, professional terminology, and documentation styles far better than most international models. This results in improved factual consistency and reliability, a critical



requirement in healthcare, where accuracy is essential.

**Q:** Why can smaller models compete with Large International Models?

**A:** At first glance, it may seem obvious that “bigger is better” when it comes to AI models. However, NorwAI’s research shows that smaller, specialized models can effectively compete with much larger international alternatives.

◆ *The secret lies in specialization. Large international models are designed to cover many languages and use-cases at once. While impressive, this generality can become a weakness when dealing with specific languages, regulations, or professional domains. In contrast, a smaller model trained deeply on Norwegian data can act like a specialist, outperforming a generalist in its area of expertise, says Lemei.*

Early experiments show that fine-tuned Norwegian models perform competitively with significantly larger international models, despite being smaller, cheaper to train, and capable of running locally. This approach makes advanced AI more accessible to Norwegian organizations while maintaining data sovereignty, ensuring that sensitive information remains within Norwegian borders and complies with local regulations.

**Q:** Explain how a general assistant can be trained to be a medical expert: How does fine-tuning work?

**A:** So how does a general language model become a specialized healthcare tool? The process is called fine-tuning, and it’s like professional training after general education.

After learning general Norwegian language patterns during pretraining, the model is exposed to thousands of carefully curated healthcare examples. These include medical question-answering, clinical note summarization, patient dialogues, and diagnostic reasoning tasks. Each example follows a clear instruction format (e.g. “Task: X, Input: Y, Output: Z”) teaching the model how to respond appropriately in real clinical scenarios.



❖ *Crucially, this data is created and reviewed by medical professionals. Every example passes through multiple quality checks, including expert annotation, peer review, automated consistency testing, and final approval by senior medical experts. This ensures the model learns from high-quality, accurate medical information, says Peng*

**Q:** *Why is NorwAI-Magistral-24B-Reasoning a strong foundation for healthcare AI?*

**A:** NorwAI-Magistral-24B-Reasoning offers a powerful foundation for future healthcare AI in Norway. Its dynamic reasoning capabilities allow it to support both simple information requests and complex diagnostic thinking. Its extended context window enables it to process long patient histories and detailed clinical documentation without losing critical details.

Because it is trained on Norwegian data, the model naturally understands healthcare

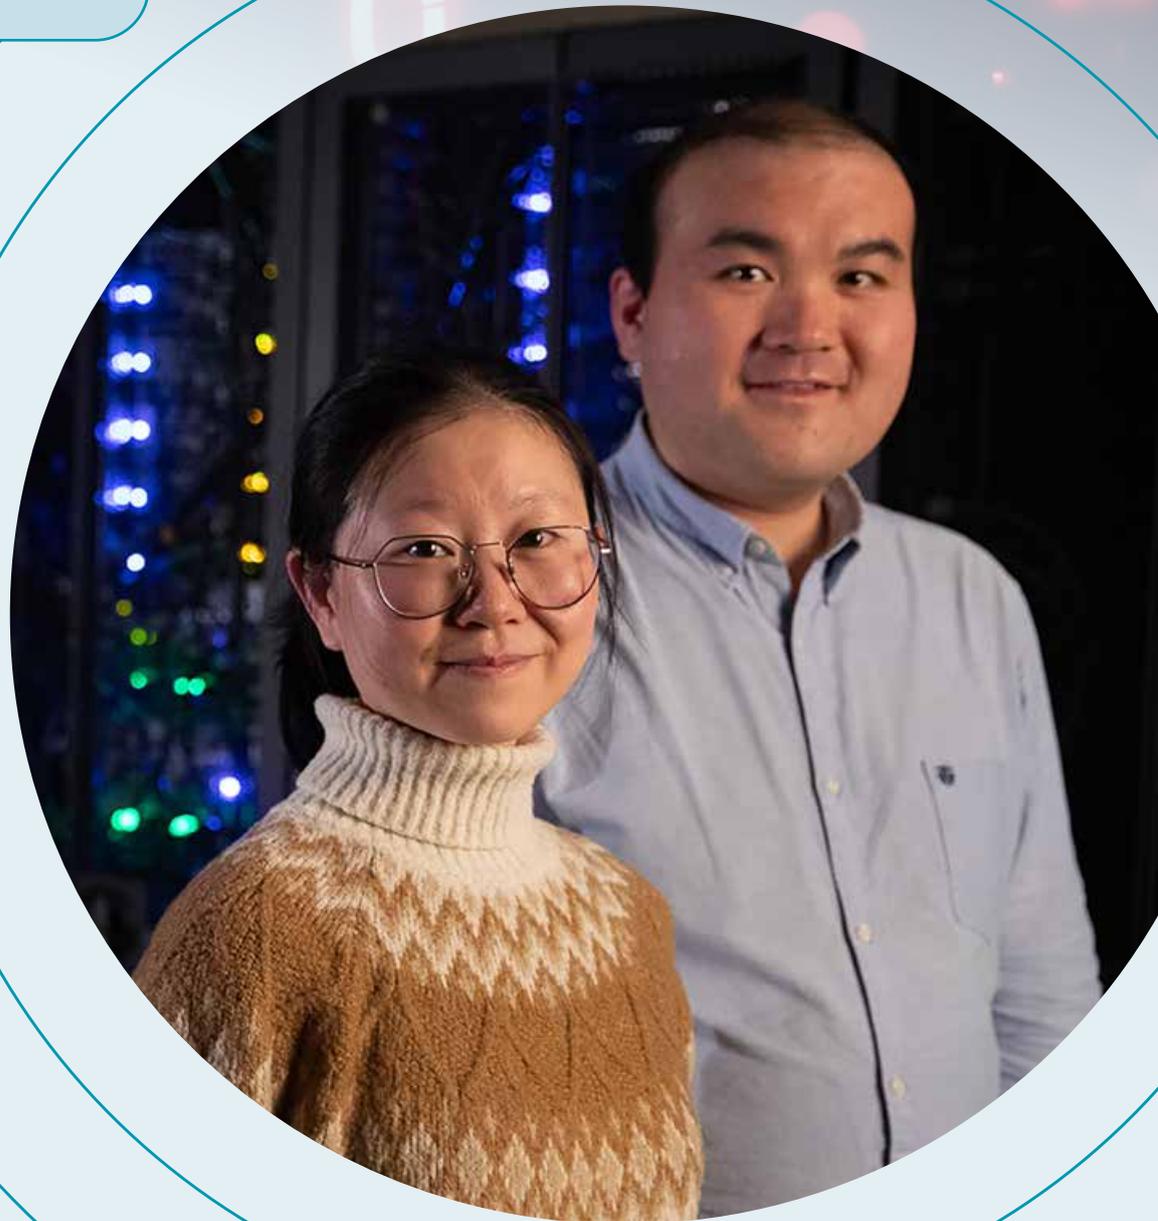
terminology, documentation practices, and patient communication styles to a certain extent. This makes it well suited for clinical documentation, diagnostic support, and patient interaction.

At the same time, NorwAI researchers are clear: this is a foundation, not a finished medical product. Before clinical deployment, the model must undergo specialized medical fine-tuning, rigorous evaluation, and continuous human oversight. A safety-first design, expert involvement, and careful validation are essential when patient health is at stake.

❖ *By combining advanced reasoning, expert-curated data, and a safety-first approach, NorwAI-Magistral-24B-Reasoning points toward a future in which Norwegian healthcare professionals are supported by AI systems that are not only powerful, but also trustworthy, transparent, and aligned with national values, says Lemei and Peng.*



*Lemei Zhang and Peng Liu are key personnel in developing large language models for NorwAI.*



# From first LLM movers to delicate surgical AI



**Eric Monteiro**

Professor, NTNU

NorwAI Core Team Member

Since November 2022, LLMs have generated unprecedented interest worldwide not the least from the general public. Awe, fear and genuine surprise in equal measures have characterized the reactions. The earliest Norwegian LLMs from NorwAI at in August 2021, preceding the launch of OpenAI's ChatGPT in November 2022 with more than a year, has similarly enjoyed a lot of attention from researchers and the industry alike.

In the world of LLMs, November 2022 seems a long time ago; so much has happened since then. Following OpenAI's LLM, many companies, consortia and national efforts have launched their own LLMs, making pre-trained LLMs into something akin to commodities in 2026.

Consequently, NorwAI is with its activities in LLMs busy with the next step: domain-specific LLMs, with the domain of healthcare an early, prominent and important one. The significance of a shift to domain-LLMs is that LLMs here

- ◆ target designated, practically useful tasks (e.g., summarizing patient-doctor dialogues, extracting data from patient records as specified by a clinical trial protocol)

- ◆ has a degree of 'tribal language' (e.g., specialized vocabulary, concepts and abbreviations used in 'medical Norwegian').

Different from AI used for critical decision making (e.g., diagnostic interpretation of medical images), the tasks targeted for health LLM typically are non-critical – but time- and labour-intensive, hence prime candidates for efficiency gains. Our approach (i.e., research design) has a tolerance for inaccuracies in the performance of the health LLMs, which significantly boosts the prospects of embedding them into everyday, clinical workflows.

## **Quality evaluation needed**

A prominent challenge with health LLMs,



*Professor Eric Monteiro (front) and Innovation Manager in Hemit HF Ketil Thorvik at the NorwAI Innovate 2025.*

<<< ●

research-wise as much as practical, is the need to work out a new regime of evaluating the performance of LLMs. As the usefulness and effectiveness of health-LLMs is so sensitive to context and clinical situation - summaries at a psychiatric ward is “long, narrative” whereas at an orthopedic ward short and factual – implies that the evaluation of what a “good summary” cannot only be assessed by quantitative metrics (which have dominated evaluations of generic LLMs), they need to be supplemented also by qualitative evaluation by the physicians themselves.

Exactly what such a regime of hybrid quant & qual evaluation should look like is an important research challenge and simultaneously of pressing, practical value. In our work with health-LLMs, we accordingly rely on a close collaboration with St Olav, Hemit HF and Helseplattformen to assure the support from physicians, a support we need to ongoingly maintain and deserve. In our project, we have established a clinical panel of physicians to have an arena to discuss this openly.

●

# Hemit HF joins NorwAI to foster safe and trustworthy AI in healthcare



**Nina Kotte**  
Communication Manager  
Hemit HF

Hemit HF has joined NorwAI as a new partner, strengthening the national ecosystem for responsible artificial intelligence.

– Our ambition is to foster AI-innovation by ensuring safe and trustworthy LLMs in healthcare, says Trond Utne, CEO at Hemit HF.

Hemit HF delivers technology and digitization services for specialist healthcare to the Central Norway health region (Helseregion Midt-Norge). The organization develops, manages and operates common ICT systems for all hospitals in Trøndelag and Møre og Romsdal. With approximately 450 employees, Hemit's mission is to enhance patient care through innovative digital solutions while ensuring the continuous availability of critical hospital systems.

Hemit works closely with hospitals in the region on a joint initiative for AI in the Health Region of Mid-Norway and participates in several national forums for artificial intelligence in the health sector.

– There is a significant and growing commitment to artificial intelligence within the healthcare

services. Expectations relate to time savings, reduced waiting lists, less overtime, a more balanced workload for healthcare employees, and improved quality in medical practice, says Ketil Thorvik, innovation advisor at Hemit HF.

## AI in Healthcare

National health authorities have established strict requirements for the introduction of AI in the health sector. The report "Use of Artificial Intelligence in Health and Care Services" (Helsedirektoratet, 2025) outlines the considerations organizations must assess to ensure AI systems are trustworthy when procuring, implementing, and using AI.

– Successful implementation of artificial intelligence in healthcare is a multidisciplinary effort. Hemit works closely with the hospitals, Helseplattformen, and Central Norway Regional Health Authority to ensure safe and effective implementation, says Trond Utne.

The two directors, Jon Atle Gulla of NorwAI (left) and Trond Utne of Hemit HF. Photo: Ann Iren Jamtøy



**Ketil Thorvik**  
Innovation advisor  
Hemit HF



**Per Olav Østbyhaug**  
Director of development  
St. Olavs hospital.



In Norway, early AI deployment has largely focused on radiology, but LLMs are now being introduced to support doctor-patient consultations by listening to the conversation and generating structured summaries for the medical record. This is being rolled out both in primary and specialist healthcare services.

– Although experiences vary in early implementations, we see growing enthusiasm and increased maturity for AI in healthcare. There is a clear message that success requires far more than technology. Change management and organizational development are crucial, says Trond Utne.

### The key question

A key question moving forward is:

Will small(er), Norwegian LLMs, trained or fine-tuned for specialized domains and designated tasks, perform better than large, generic LLMs?

This is a central focus in the collaboration between Hemit and NorwAI, aimed at accelerating AI-driven innovation while ensuring safety and trustworthiness in healthcare. St. Olavs hospital has also joined the effort. Together, the partners aim to develop a pipeline for domain-specific, smaller LLMs tailored for healthcare.

– St. Olavs has strong ambitions for applying AI in medical practice, says Per Olav Østbyhaug, director of development at St. Olavs hospital.

– But we aim for more than research. We need successful innovation and adoption to create value in the healthcare services. I see a strong commitment at Hemit and NorwAI to ensure innovation, and our clinicians are enthusiastic about this collaboration, says Per Olav Østbyhaug.

# What lies beyond benchmarks?

NorwAI researchers Eric Monteiro, Lars Bungum, and Pål Furu Kamsvåg argue that benchmarking constitutes a paradigm, which brings with it assumptions and shortcomings. After recapping their development in the era of Artificial Neural Network (post AlexNet), they identify challenges to this paradigm and opine on future research that could address these with an emphasis on Large Language Models.



**Lars Bungum**  
Postdoc researcher  
TrustLLM, NTNU



**Pål Furu Kamsvåg**  
PhD candidate  
NorwAI, NTNU



**Eric Monteiro**  
Professor  
NorwAI, NTNU

The concept of benchmarking has underpinned the breakthroughs of AI, as exemplified by the ImageNet breakthrough of 2012. Because it offers a consistent and comparable way to assess different models, progress in the field has been viewed through the lens of benchmark performance. Furthermore, benchmark performance offers insight into whether one is headed in the right direction when developing new models.

As the forefront of LLM development, especially, is driven by commercial entities, benchmark performance is also a sales point, and companies will, therefore, provide benchmark performance as new models are launched, also for marketing purposes. That a VP at Meta[1] had to publicly deny that they had trained their Lambda 4 models on the test sets (the equivalent of peeking at the correct answers when taking an exam) is a testament to the current importance of benchmarks.

However, increasing attention has been given to what the benchmarks leave behind. For example, we have LLMs passing the bar exam

with flying colors, but no AI lawyers. There are LLMs solving math problems of the most difficult sort and famously excelling at complex strategic games like Chess and Go. Thus, narrow artificial intelligence has arrived, but general artificial intelligence not (yet?). Performance on narrow tasks is in some cases stellar, but when the technology is lifted up to a level of practical application, the situation is unclear. We address this discrepancy in our research, using health applications as a case study. After addressing the many pitfalls of benchmarking that have been pointed out, we propose solutions through hybrid evaluation schemata.

### **A powerful metaphor**

Ethan Mollick casts AI application with a powerful metaphor: either that of a centaur, or that of a cyborg. The former pictures a clear line between the human torso and the horse body, and the latter a blend. As AI systems increase in complexity, the latter is becoming the norm, as witnessed by academic conferences requesting declarations of AI use, it often blends into various aspects of the writing process like research, grammar and orthography, but in principle also the coherence of ideas and the modes and structure of presentation. On a more fundamental level, Blaise Agüera y Arcas also argues forcefully in his recent book *What Is Intelligence* that “symbiogenesis” is the main driver of evolution (not point mutation). A symbiotic relationship between man and machine was already purported in Nobel Laureate Hannes Alfvén’s pseudonymous 1966 classic *The Tale of the Big Computer: A Vision*. The time elapsing from the advent of computers until they surpassed humans in ability was characterized by fruitful cooperation, a “symbiosis” between man and computer.

In this vision, a combination coefficient is posited, as a replacement of the old intelligence quotient of the pre-symbiotic era, i.e., that the power of a computer-aided system is determined by its ability to combine skills and information into novel actions. Alfvén’s tale acknowledged the difficulty of defining this coefficient but specified it as pertaining to how incoming signals from sensory systems (or their equivalents) could be combined and understand which of these combinations align with the memory of already present combinations. Furthermore, it also recognizes

the ability to create impulses which are not triggered by incoming impulses but pertain to combinations stored in memory for a long time. These abilities correspond to some degree to “initiative”. It is definitely possible to combine benchmarks, but numerical combinations, however elaborate, seem insufficient to account for the above vision. How may we come closer?

### **About beauty**

Evaluation is inherently controversial. As the old adage says: “beauty is in the eye of the beholder.” In Machine Translation, for example, some evaluation forms were considered to favor rule-based systems over statistical, and

*A benchmark* was once used to literally mark benches with varying heights for use in topographical studies. Figuratively, however, it has long been used to compare the performance of more solutions to the same task.

*In computer science*, benchmark datasets have been frequently used for comparing algorithmic performance. Alongside the increasing success of data-driven relative to rule-based approaches, the development of such datasets has scaled up to at least the same degree.

*With the ability of modern AI systems* to solve problems of increasing difficulty, many of which were beyond the reach of computers systems up until recently, new datasets have emerged to characterize how well different models work.

vice versa, which is illustrative of how the choice of evaluation framework influences system ranking. In a quip often attributed to the late Yorrick Wilks, it is claimed that more has been written on MT evaluation than MT itself, much for this reason. But evaluation frameworks not only impact rankings; they are part of and influence the world. We research how.

The research community strives to make evaluations informative and comparisons fair. Over the last decade or so, the developments in LLM evaluation can be summed up by the following concepts; an expansion of tasks into multimodal settings (from image to video), human diversity and commonsense (can you push a string?), human alignment (are decisions perceived as fair?), aggregate metrics and leaderboards (how do LLMs perform on diverse tasks?), and, finally, using LLMs as judges (prompting them to evaluate given criteria and examples). All these developments go in the right direction. But what is still missing?

### **Socially shaped benchmarks**

Among the remaining problems with benchmarks, we consider how benchmarks are socially shaped. When images are categorized, from whence come the categories? In a similar vein, they are historized, reflecting the signs of the times. Moreover, they are co-constructed with the developing technology, often, but not always constructed before you know what the task you are looking for. Finally, an overreliance of benchmarks risks cultivating monocultures, akin to the risks of monocropping in agriculture, which encompasses the biological risks of relying on too few crops. There may be a trade-off between the short- and long-term good. We maintain that these differences cannot only be stated in much-rehearsed arguments on qualitative vs. quantitative and the process of discretizing qualitative knowledge into computable rules.

There have been efforts to address the static nature of benchmarks by creating moving targets, and also increasing the difficulty of benchmarks that have become saturated, obsoleted by genuine performance increases or also data leakage, whereby LLMs are trained on the datasets, and, therefore, easily come up with the right answers. Efforts like GDPval also look at real-world economically valuable tasks, primarily evaluated by human experts. There are also efforts to assess the environmental impacts of LLMs, e.g., by calculating the Pareto optima of resource use and performance.

### **How we continue**

Our research continues in this direction. Benchmarks measure performance in the present but also shape perceptions and influence the world. Moreover, the increasingly interwoven relationship between man and machine as their efforts are braided beckons us to investigate not only the merits of these braids, but also the effect this interdependency has on the human strand in particular.

We begin with studies in the health sector, operationalized by physicians and writing their summaries unaided, and thereafter assessing the quality of LLM-written summaries. This represents a hybrid evaluation form.

As we continue our foray into hybrid evaluation, we consciously side-step whether the mix should be 50-50 or 90-10. Instead, we signal that the qualitative input needs to be ongoingly legitimized. We also need to look at the long-term consequences of LLM use on, e.g., critical thinking, the crucial ability to monitor the systems, and we are welcoming any further work on other complex application areas where there is no clear-cut 'ground truth', just like health care. We expect to present results on these efforts in the coming year.



---

# AI for project management in megaprojects



**Even W. Laurak**  
Research assistant  
NorwAI

*NTNU Unified Campus model. Photo: Per Henning*

Large megaprojects, such as the construction of bridges and government facilities, are complex undertakings in society. They can involve thousands of people, big investments, and several years of planning and execution.

«Project Norway», a competence center and partner-based organization owned by NTNU, runs «Bedre Megaprojekter» where solid engineering expertise is combined with a deep understanding of the underlying domain language, terminology, risk assessment, planning, and quality control. Professor Eilif Hjelseth at the Department of Civil and Environmental Engineering reached out to NorwAI in the summer of 2025 asking if our work with domain specific language models could be of use in Bedre Megaprojekter.

His initiative gave birth to a project to develop a standardized methodology for using large language models to learn the language of different sectors and help project managers navigate large volumes of documentation and data.

NorwAI's team is led by Professor Jon Atle Gulla coordinating the project's industry partners and experts, Postdoc researcher Lars Bungum as task leader for the research assistants Even W. Lauvrak and Jorge Ballesta Cerezo.

### **Falling short**

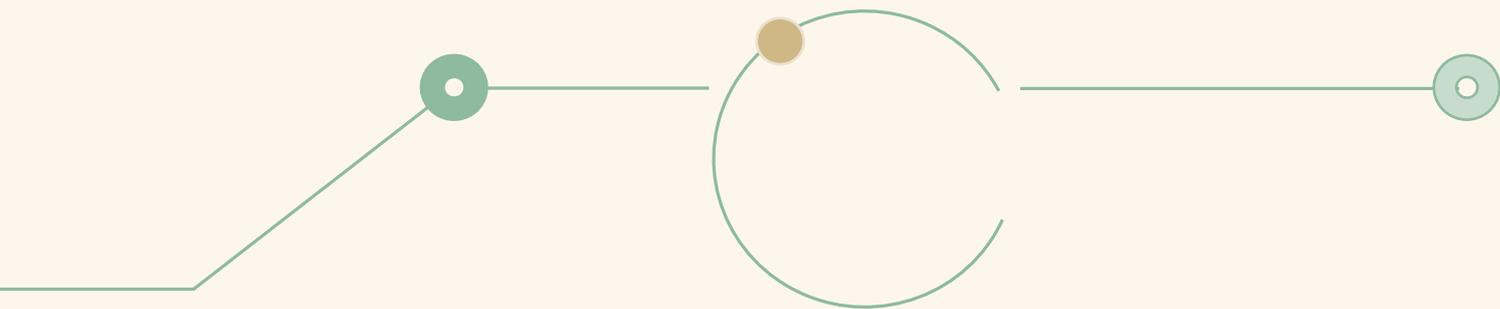
In recent years, we have seen how much models such as ChatGPT, Claude, and Gemini are capable of, and the value they bring in terms of streamlining work. These language models are designed for general purposes and trained on a broad body of text. When it comes to specific

### **Bedre Megaprojekter**

"Better Megaprojects" is a national development program that aims to contribute to the largest and most complex projects. The number of megaprojects in Norway has never been higher, and it will increase in the coming years.

According to Independent Project Analysis (IPA), too many megaprojects fail. They either go well or they fail very badly. Megaprojects that fail to achieve their goals have major consequences for the companies that own them and for social development in Norway.

Ongoing demonstration projects are NTNU Unified Campus, the new National Hospital, Innlandet Hospital and Rehabilitation at UiB.



domain language, such as terminology used in Norwegian construction and civil engineering, we often find that the models fall short. Texts containing legal language or words with double meanings can frequently be misunderstood or mishandled by such general-purpose models.

To address these challenges, we are working to develop a standardized methodology for training and applying domain-specific language models that can absorb the language, identify connections, distinguish between double meanings, and above all apply that linguistic knowledge to specialized purposes. What these specialized purposes are depends on the end user's needs within the given domain. In our case, we have so far focused on project management within megaprojects in the construction and civil engineering industry.

#### **Valuable assistance**

Since the start of the project, we have received valuable assistance from Eilif Hjelseth, professor at the Department of Civil and Environmental Engineering at NTNU, as well as PhD student Christian Bakke from the Faculty of Engineering. Both Hjelseth and Bakke have solid expertise in project management for megaprojects and have provided us with access to data that we have used in the development and evaluation of our

methodology, tested against the construction and civil engineering domain.

So far in the project, we have trained the models on cleaned domain data and evaluated their ability to learn and apply the language. However, to carry out the more advanced tasks, we need to fine-tune the models for domain-specific applications. The methodology for a standardized approach to both fine-tuning and evaluation is still under development. The challenge lies in finding a robust, standardized procedure that can be applied across multiple sectors and evaluated in a general way, without knowing exactly which task the model will be required to solve.

To fine-tune the models for tasks that an industry partner requires, we also need structured datasets containing curated sets of input data and desired output, also called an instruction dataset. The plan is for this to be developed in close collaboration with participating industry partners through a dedicated workshop. The purpose of such a workshop is twofold, we want a curated instruction dataset developed by the industry partner under our guidance for model training, as well as knowledge building for the partner, giving them insight into the capabilities of language models.



*When it comes to specific domain language, such as terminology used in Norwegian construction and civil engineering, we often find that the models fall short.*

# Two more partners in the NorwAI consortium



**Rolf Dyrnes Svendsen**  
Communications Manager  
NorwAI

Two companies were accepted for partnership in NorwAI in 2025. The successful medical start-up Medbric, a true child of NorwAI itself, joined the consortium late in the year. So did the health enterprise Hemit HF, both companies based in Trondheim.

For Hemit HF, the application was seen in the context of the planned collaboration with NorwAI on a project where, together with St. Olavs Hospital and Helseplattformen, they plan to test out whether NorwAI's newest language model will support the hypothesis that smaller, Norwegian language models, trained and fine-tuned for the health domain, will perform better than large, international generic models.

### **A true child**

Medbric on the other hand, is a true child of research from the very beginning in August 2024. Researcher and physician Jorunn Thaulow's project at the Antibiotic Center (ASP) at the University of Oslo showed variations from doctor to doctor on using antibiotics. Could artificial intelligence provide solutions to the problem of overuse of antibiotics, she asked.

At NorwAI, project manager Jon Espen Ingvaldsen developed Norwegian language models. He took Thaulow's findings seriously and built solutions that met the requirements for good medical needs understanding, security and user-friendliness.

- The company was founded on September 17 2024 and eight days later Medbric was officially launched on stage during NorwAI Innovate Conference that year. We build AI agents for the healthcare sector. The solutions provide precise, personalized medical records, summaries and quick answers to professional questions, says CEO Jorunn Thaulow.

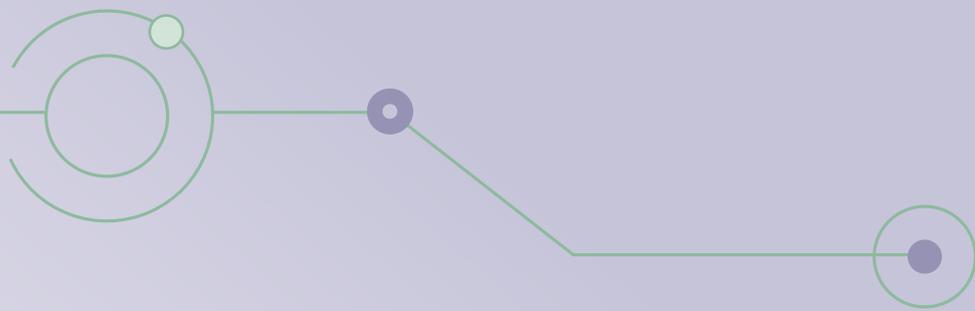
### **15 months later**

15 months later, by the end of 2025, 2000 GPs (general practitioners), about 25-30 % of all GPs



# Pioneering AI Agents for Health

Founders of Medbric, Jon Espen Ingvaldsen and Jorunn Thaulow, presenting the first spinoff from AI research at NorwAI and the University of Oslo at the NorwAI Innovate Conference in 2024. AI agents assist general practitioners (GPs) in the primary health service.



in Norway, use their *ambiance scribe* solution. In addition 1000 physicians in the specialized health service are among the users, as well as 30 municipalities. Medbric's speech-to-note solution, *Aila*, was launched in September 2025 at the hospitals in both the Western Norway and the Mid-Norway Regional Health Authorities. Medbric has also entered the Danish health sector having adapted their solution to two national medical record systems. In Norway three systems are adapted.

- We firmly believe that close cooperation with a competent research community is a recipe for us. Another success factor is a close co-op with health personnel themselves, says Jorunn Thaulow.

### **The tech enterprise**

Hemit HF is the technology enterprise of the Health Authority of Mid-Norway. Hemit HF develops, manages and operates common ICT systems for all hospitals in Trøndelag and Møre og Romsdal. Hemit HF works closely with the hospitals on a joint initiative on AI in the Health Ministry of Mid-Norway, and participates in several national forums dealing with artificial intelligence in the health services.

- In the specialist health service, many are eager for digital tools that improve their working life. We are anxious to see whether language models adapted to the Norwegian health service can contribute to this, says CEO Trond Utne at Hemit HF.

### **Research reasons**

NorwAI wants to enter the health domain for research reasons. Professor Eric Monteiro put words to the challenges when the plans for a

cooperation with St Olavs Hospital and Hemit HF was announced:

- It is an exciting research challenge to find good benchmarks for language models within a sector where the tasks vary significantly between different fields, from chronic mental health to surgery in all its forms, says Professor Eric Monteiro.

The national health authorities have set up strict rules for the introduction of AI in the sector. In the report on quality assurance, *Use of artificial intelligence in health and care services* (Helsedirektoratet 2025), it outlines the key considerations organizations must address to assess whether an AI system is trustworthy when procuring, implementing, and using it in health and care services.

### **A rapidly evolving field**

This report describes the current legal framework, outlining what is required by applicable regulations and, in part, what is considered good practice. However, AI is a rapidly evolving field, with new legislation on the horizon.

In 2019, the European Union published ethical guidelines describing three core components of trustworthy AI that should be upheld throughout a system's lifecycle:

(1) it should be legal, complying with all applicable laws and regulations, (2) it should be ethical, ensuring adherence to ethical principles and values, and (3) it should be robust, both from a technical and social perspective, as AI systems, even with good intentions, can cause unintended harm.

# Consortium

## A POWERHOUSE

*Some of the largest and technologically most ambitious companies and research institutes in Norway have joined the consortium.*

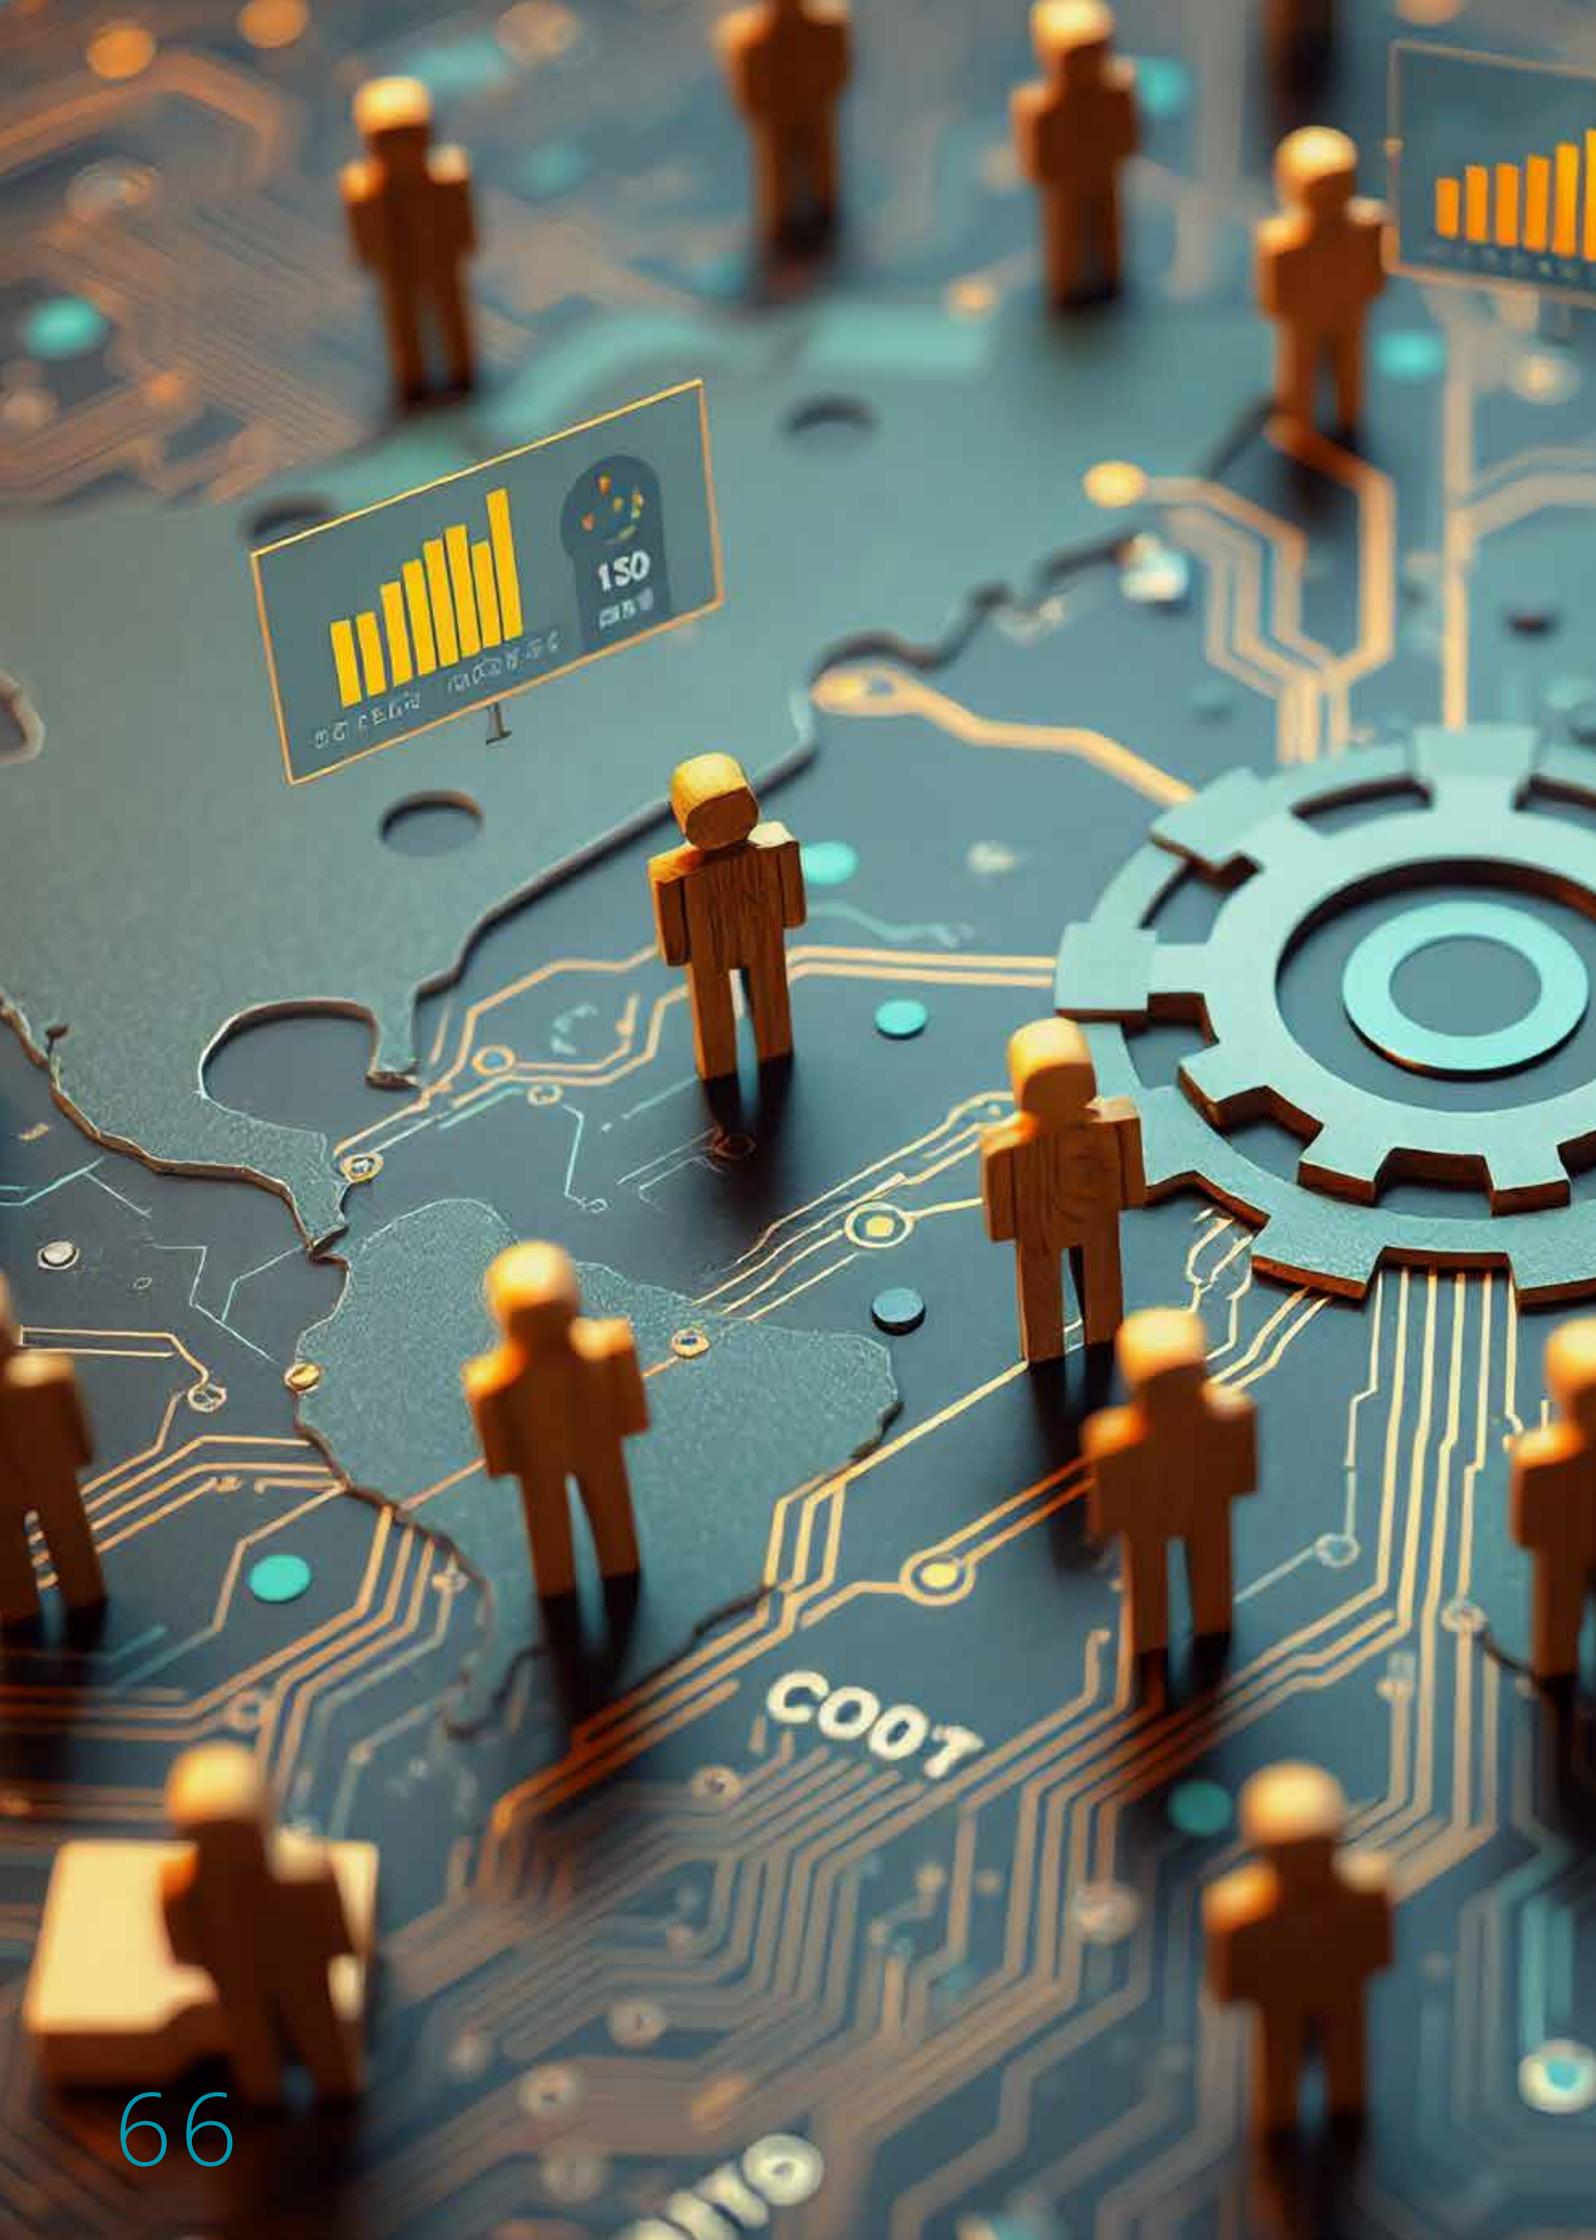
### RESEARCH PARTNERS

NTNU, the Norwegian University of Science and Technology, Department of Computer Science is host for the center, the other research partners are Norwegian Computing Center (NR), SINTEF, University of Oslo and University of Stavanger.

### INDUSTRIAL PARTNERS

The group of industrial partners in Norway consisted in 2025 of ANEO, Cognite, Digital Norway, DNB, DNV, Hemit, Kongsberg Digital, Medbric, NRK, Schibsted, SpareBank 1 SMN, Statnett, Telenor and Vend.





# Members of the board

## CHAIR EXECUTIVE BOARD



**Sven Størmer Thaulow**  
Executive Advisor, for Vend

## MEMBERS OF THE EXECUTIVE BOARD (2025)



**Odd Erik Gundersen**  
Chief AI Officer, Aneo



**Ingelin Steinsland**  
Vice Dean Research, Faculty of Information Technology and Electrical Engineering, NTNU



**John Markus Lervik**  
Chief Strategy and development officer, Cognite



**Juan Carlos Lopez Calvet**  
Director of Data and AI, Schibsted Media



**Liv Dingsør**  
CEO, Digital Norway



**Odd Are Svensen**  
Research Director, SINTEF



**Karl Aksel Festø**  
Head of CoE Advanced Analytics, DNB



**Astrid Undheim**  
Executive Director, Sparebank1 SMN



**Frank Børre Pedersen**  
Vice President & Program Director "Future of Digital Assurance", DNV



**Arild Nebb Ervik**  
Chief Data Officer, Statnett



**Stein-Roar Skånhaug Bjørnstad**  
CTO, Kongsberg Digital



**Gorm Grønnevet**  
Program Director, Telenor Research



**Anders Løland**  
Research Director, NR



**Viktoria Stray**  
Deputy Head of Department, University of Oslo

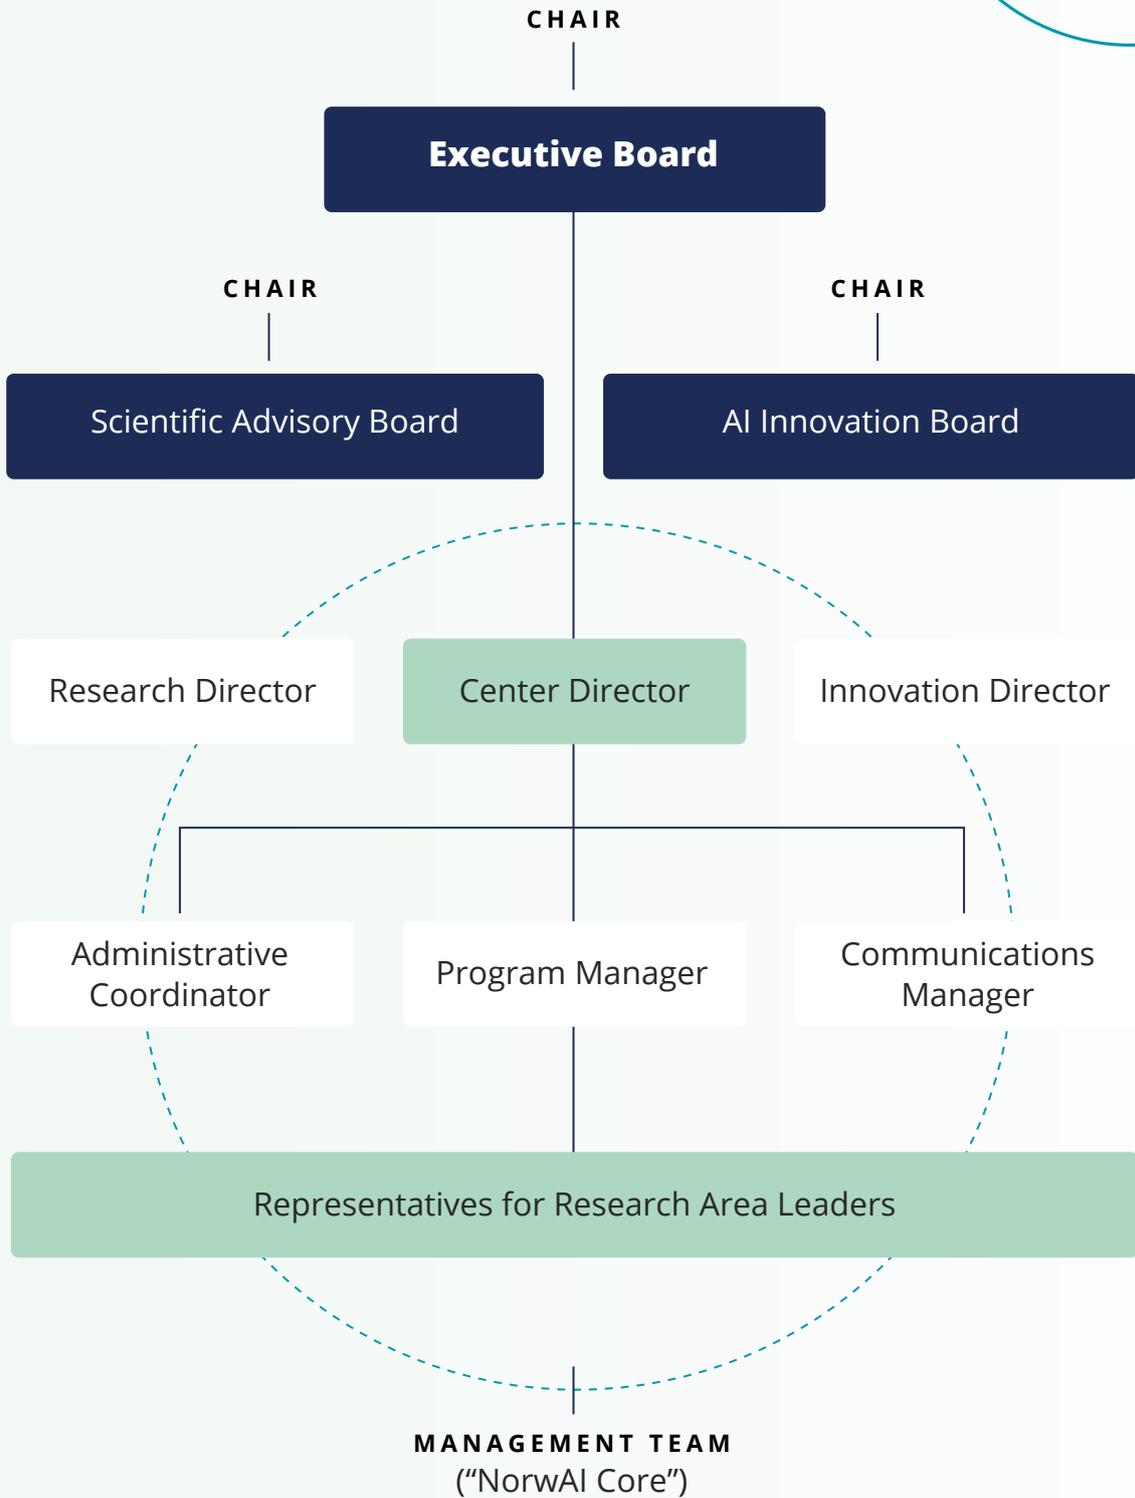
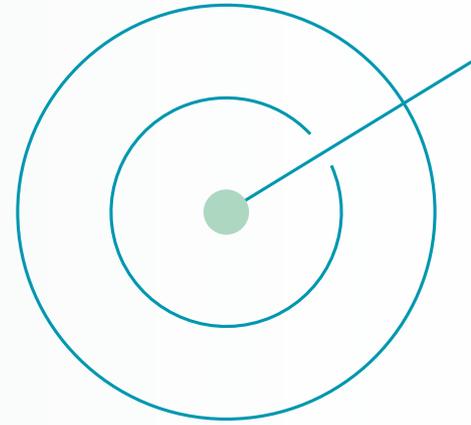


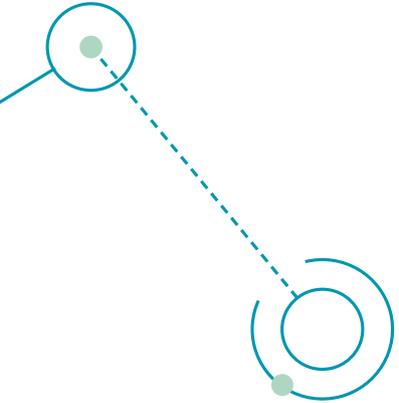
**Pål Nedregotten**  
Director of Technology, NRK



**Tom Ryen**  
Head of Department of Electrical Engineering and Computer Science, University of Stavanger

# Organization





---

# Center Management Team



*The Center Management Team is responsible for the day-to-day operation of the center and consists of:*



**Jon Atle Gulla**  
Professor, NTNU,  
Center director



**Özlem Özgöbek**  
Associate Professor, NTNU,  
Program Manager



**Terje Brasethvik**  
Adjunct Associate Professor, NTNU,  
acting Research director  
(until Aug -25)



**Benjamin Kille**  
Associate Professor, NTNU,  
Program Manager



**Kerstin Bach**  
Professor, NTNU,  
Research director



**Signe Riemer-Sørensen**  
Research Manager, SINTEF,  
Research Area Representative



**Till Christopher Lech**  
Research Manager, SINTEF,  
Innovation Director



**Jon Espen Ingvaldsen**  
Adjunct Associate Professor, NTNU,  
Research Area Representative



**Karolina Storesund**  
Senior Advisor, NTNU,  
Administrative Coordinator



**Eric Monteiro**  
Professor, NTNU,  
Research Area Representative



**Rolf Dyrnes Svendsen**  
Head of NxtMedia Lab,  
Communications Manager



**Theodoros Chondrogiannis**  
Associate professor, NTNU,  
Research Area Representative

---

# NorwAI objectives

The objective of the NorwAI research center is to provide a strong and robust arena for industry, research and academic institutions to collaborate on the development of AI ideas and techniques, share results, and iteratively explore how technology can transform existing businesses and enable entirely new business avenues.

---

The primary objective – or mission – of NorwAI is to:

*Accelerate the innovation of sustainable and trustworthy artificial intelligence solutions across Norwegian industries.*

This primary objective is further broken down into objectives and secondary objectives, as shown on the next page.

Artificial intelligence constitutes a paradigm shift in computer science, enabling substantially shorter development cycles, extremely powerful solutions, and immediate transfer of technologies from one domain to another. The innovation cycles become very dynamic, representing both challenges and opportunities at the same time. New systems may completely transform existing practices, render old value chains worthless, or open for brand new business opportunities. NorwAI supports innovation as a platform for continuous interaction between industry and academia, and will act as an ecosystem for creating alliances, joint venturing and building synergies among all partners. NorwAI will thereby enhance the ability of the business sector to innovate and create value through a greater focus on long-term research.

NorwAI recognizes that the deep impact of AI makes it essential to ensure effective sharing of knowledge and to equip businesses with the skills needed to adapt their innovation

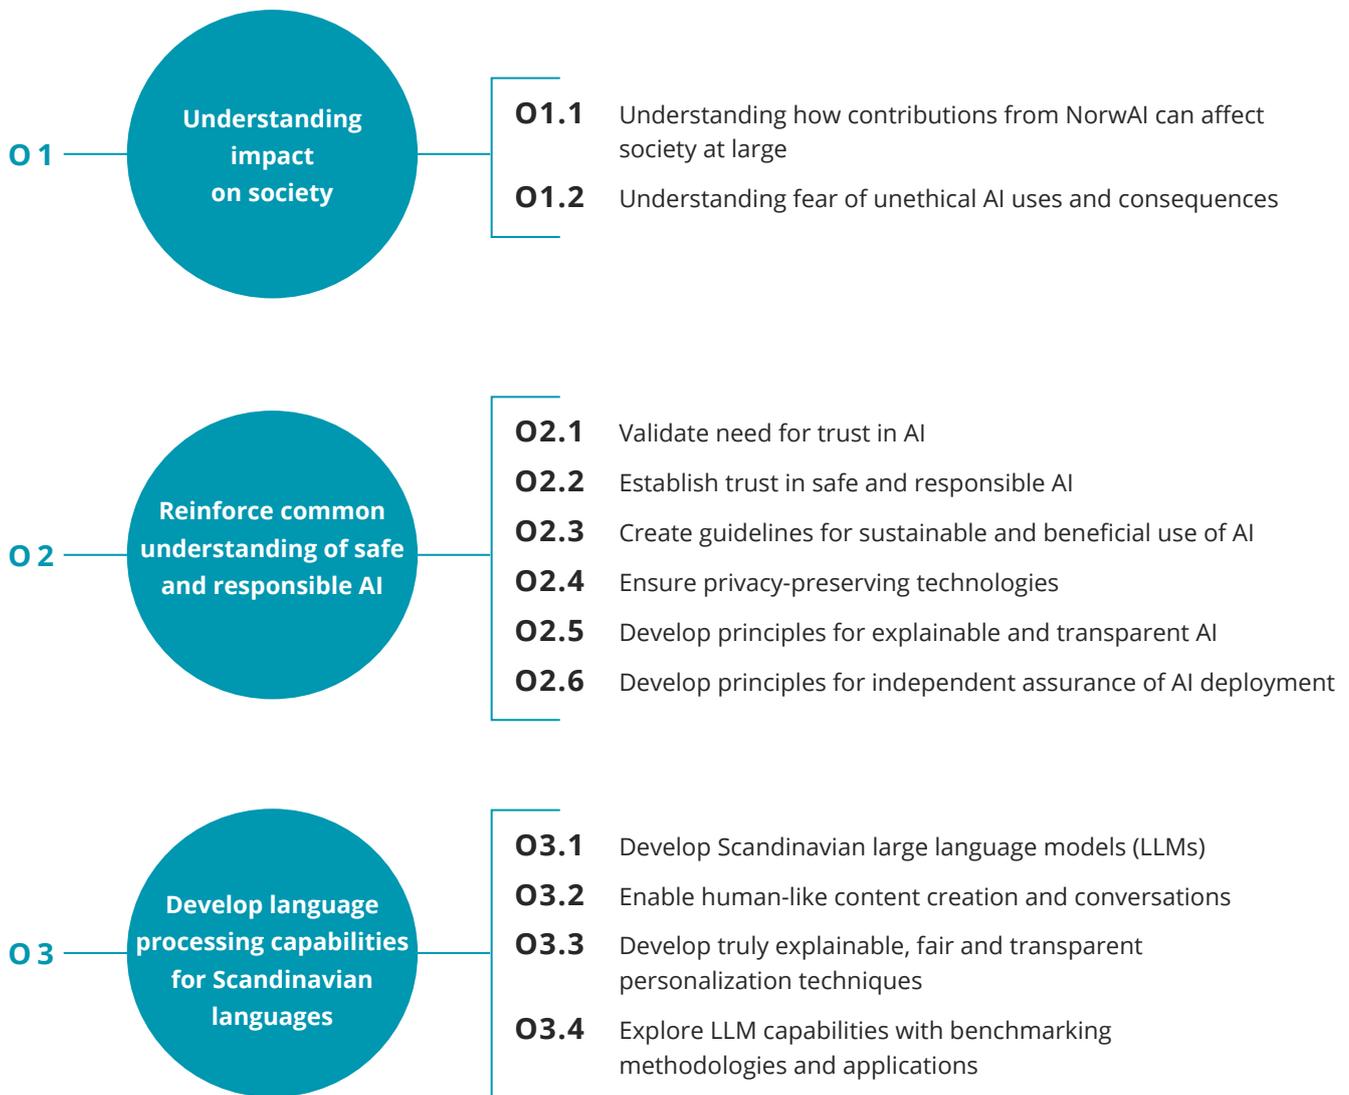
processes to rapidly evolving technological landscapes. The center places strong emphasis on responsible and safe use of AI, reinforcing a common understanding of trustworthy practices, privacy-preserving technologies, and transparent, explainable systems. It contributes to the development of AI solutions that respect societal needs while enabling personalization, fairness, and transparency across applications. At the same time, the center advances Scandinavian language technologies, including the development and application of large-scale language models, human-like content generation, and benchmarking methodologies. It also supports modern AI techniques for streaming and sensor-based data, uncertainty quantification, and data integration through the combined use of AI methods, knowledge graphs, and large language models.

A significant part of NorwAI's efforts is focused on combining physical models with AI to improve understanding of physical systems, manage imperfect or non-representative data, and ensure robust, explainable predictions. In addition, the center provides data and platforms for AI innovation by enabling effective data-sharing mechanisms, establishing platforms for efficient AI development processes, and offering infrastructures that support collaboration and reproducibility. Finally, the overall quality and reputation of the research center will affect the center's ability to help companies develop and deploy research-based innovations and establish NorwAI as an internationally leading AI research and innovation center.

# NorwAI objectives

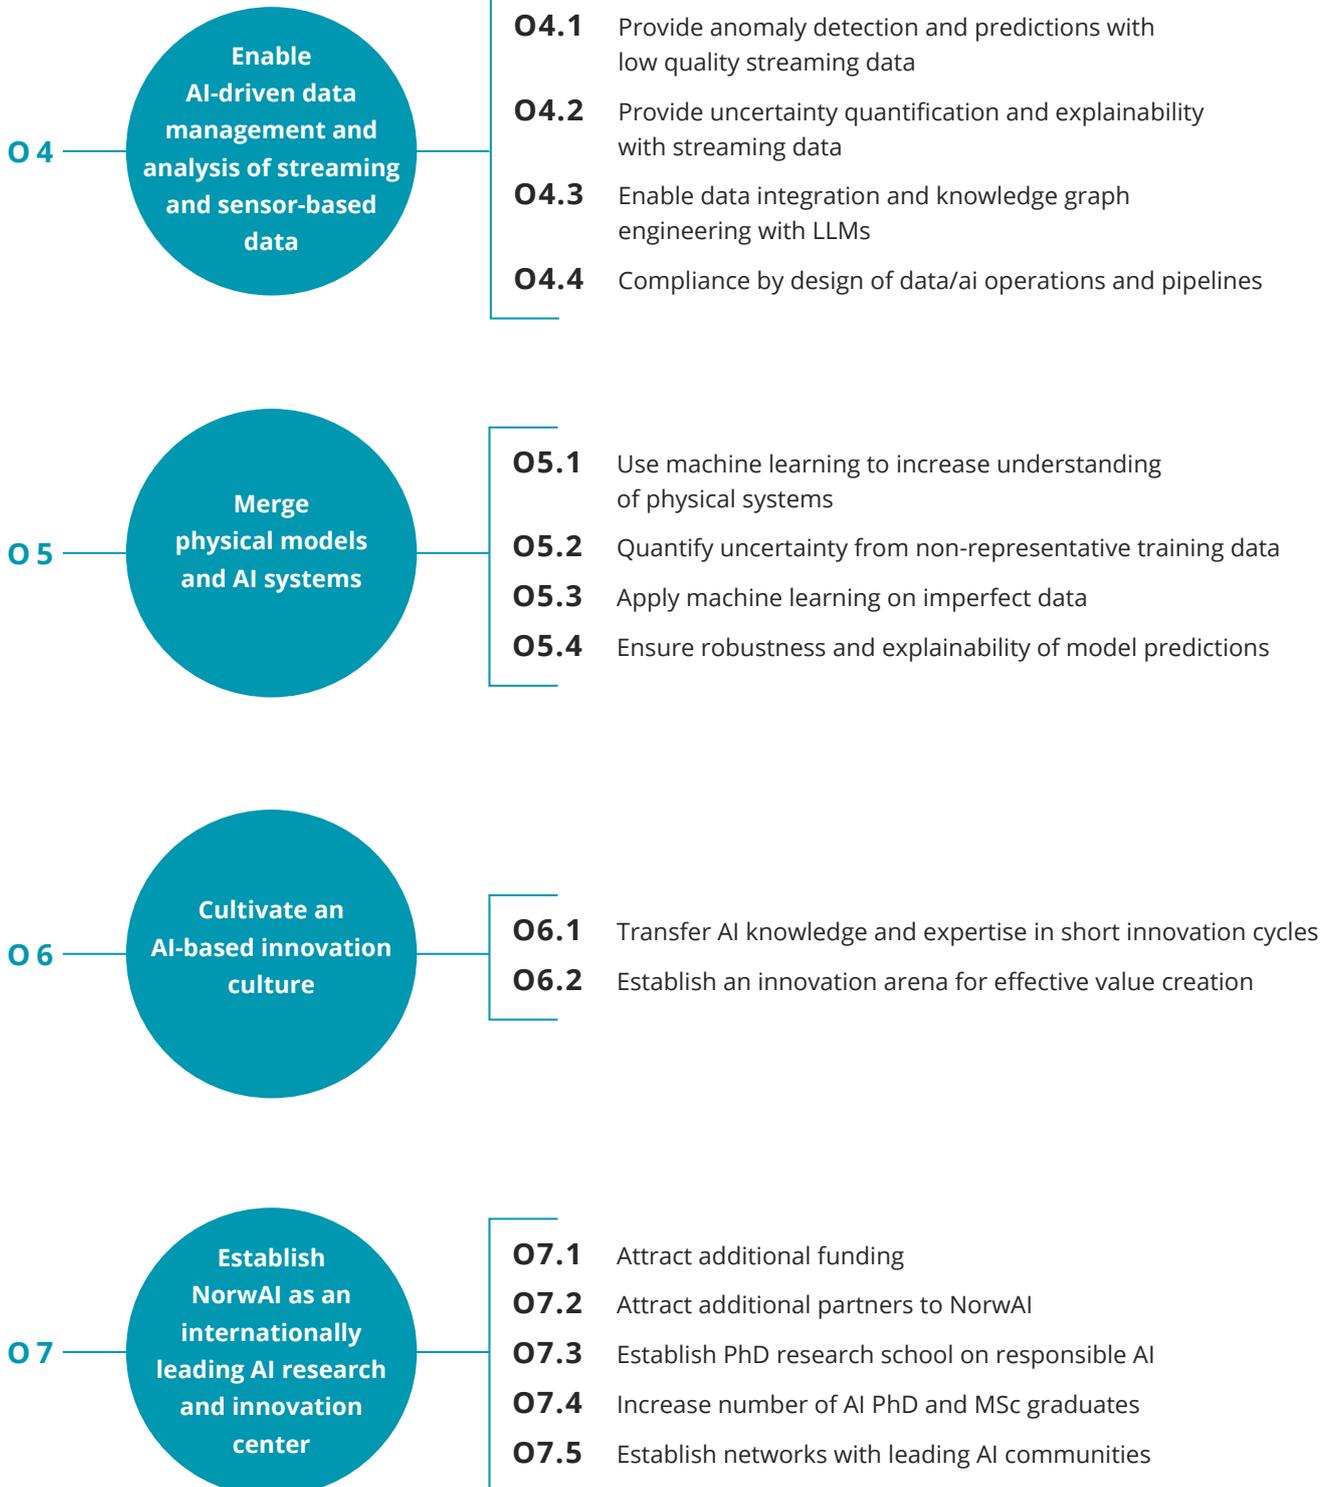
## OBJECTIVES

## SECONDARY OBJECTIVES

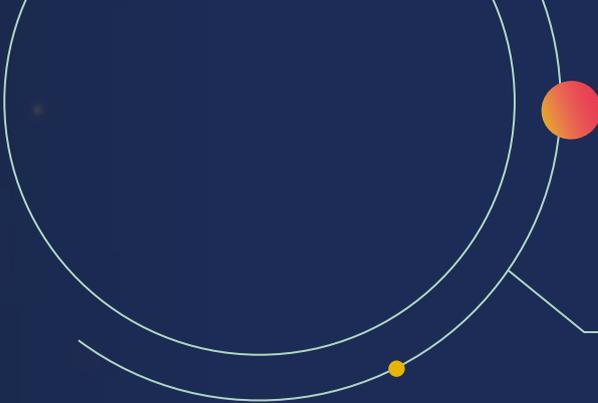


## OBJECTIVES

## SECONDARY OBJECTIVES







# Research strategy

NorwAI's research strategy is built on a strong foundation of interdisciplinary AI research that integrates methodological advances in artificial intelligence with socio-economic, organizational, and ethical perspectives. We combine fundamental research in core AI methods with use-inspired research driven by real-world challenges, ensuring both scientific excellence and societal relevance.

The center focuses on a set of core AI research themes that constitute a stable scientific backbone. These themes address foundational questions in AI while enabling innovation across diverse application domains. The portfolio is periodically reviewed to allow strategic expansion into emerging areas, ensuring scientific renewal while maintaining long-term coherence.

Our approach goes beyond parallel disciplinary contributions. NorwAI fosters integration across technical and socio-economic research through joint research projects, co-supervised PhDs, and cross-disciplinary research teams. This ensures that advances in AI methods are informed by domain needs, and that implementation, governance, and impact are addressed from the outset.

Research at NorwAI follows a use-inspired basic research model: industry and public-sector partners contribute challenges, data, and contextual knowledge, while the center retains scientific leadership in defining research questions and methodologies. This balance ensures independence, long-term research excellence, and sustained impact on innovation.

#### Through this strategy, NorwAI aims to:

- ◆ Produce internationally leading research in selected AI domains,
- ◆ Develop scalable, robust, and trustworthy AI methods,
- ◆ Bridge the gap between foundational research and deployment,
- ◆ Build long-term competence and research infrastructure in Norway,
- ◆ Contribute to responsible and sustainable AI adoption.



---

# The Scientific Advisory Board

## *An Ambitious Board*

The overall goal of the Scientific Advisory Board of NorwAI is to provide external scientific reviews of research activities, evaluate plans and progress, and contribute to shaping the center's research ambitions. The chairperson of the Scientific Advisory Board is professor Christian S. Jensen at Aalborg University.

### CHAIR OF THE SCIENTIFIC ADVISORY BOARD



**Christian S. Jensen**  
Professor, Aalborg University

### MEMBERS OF THE SCIENTIFIC ADVISORY BOARD



**Concha Bielza**  
Professor, Technical University of Madrid



**Maarten De Rijke**  
Professor, University of Amsterdam



**Virginia Dignum**  
Professor, Umeå University

# Publications in 2025



## JOURNAL PAPERS

NAME	TITLE	TYPE/PLACE	DATE
Aalbu, Kjersti; Bauer, Susanne	Data Infrastructures, International Organizations, and the Politics of Amendment	Big Data & Society	2025-04-18
Bayrak-Çik, Betül; Bach, Kerstin	PerCE: Hierarchical Perturbation-Based Counterfactual Explanations for Multivariate Time Series Classification	IEEE Access	2025-12-01
Breuer, Timo et al (incl. Bernard, Nolwenn)	Report on the 1st Workshop on Simulations for Information Access (Sim4IA 2024) at SIGIR 2024	ACM SIGIR Forum	2025-03-06
Hassan, Reda; Nguyen Nhien; Rasdal Finserås, Stine; Adde, Lars; Strümke, Inga; Støen, Ragnhild	Unlocking the black box: Enhancing human-AI collaboration in high-stakes healthcare scenarios through explainable AI	Technological Forecasting and Social Change	2025-07-23
Michałowska, Katarzyna; Goswami, Somdatta; Karniadakis, George Em; Riemer-Sørensen, Signe	Multi-Resolution Learning with DeepONets and Long Short-Term Memory Neural Networks	Neurocomputing	2025-08-05
Pilán, Ildikó; Manzanares-Salor, Benet; Sánchez, David; Lison, Pierre	Truthful text sanitization guided by inference attacks	Applied Soft Computing	2025-12-01
Vassøy, Bjørnar; Kille, Benjamin; Langseth, Helge	Opt-in Transparent Fairness for Recommender Systems Opt-in Transparent Fairness for Recommender Systems	Lecture Notes in Computer Science (LNCS)	2025-04-03

## PUBLISHED CONFERENCE PAPERS

NAME	TITLE	TYPE/PLACE	DATE
Balog, Krisztian; Bernard, Nolwenn; Zerhoudi, Saber.; Zhai, CengXiang.	Theory and Toolkits for User Simulation in the Era of Generative AI: User Modeling, Synthetic Data Generation, and System Evaluation	48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '25)	2025-07-13
Bauer, Susanne	Swimming with and against Data Streams: a digital fish and its unaccounted others	Panel on Digital doppelgangers; Nordic STS Conference, Stockholm	2025-06-12
Bernard, Nolwenn; Balog, Krisztian	Limitations of Current Evaluation Practices for Conversational Recommender Systems and the Potential of User Simulation	SIGIR-AP 2025: Proceedings of the 2025 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region	2025-12-08
Bernard, Nolwenn; Joko, H.; Hasibi, F.; Balog, Krisztian	CRS Arena: Crowdsourced Benchmarking of Conversational Recommender Systems	Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining (WSDM '25)	2025-03-10
Charpentier, Lucas Georges Gabriel; Lison, Pierre	Re-identification of De-identified Documents with Autoregressive Infilling	Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)	2025-07-25
De La Rosa, Javier; et al (incl Zhang, Lemei; Liu, Peng; Gulla, Jon Atle)	The Impact of Copyrighted Material on Large Language Models: A Norwegian Perspective	Proceedings of the Joint 25th Nordic Conference on Computational Linguistics and 11th Baltic Conference on Human Language Technologies (NoDaLiDa/Baltic-HLT 2025)	2025-03-02
Doan, Tu My; Baumgartner, David; Kille, Benjamin	Efficiently Summarizing Norwegian Legal Texts	The 30th Annual International Conference on Natural Language & Information Systems (NLDB 2025)	2025-07-05
Hadzhamolla, Berk; Stasik, Alexander Johannes; Alaliyat, Saleh Abdel-Afou	Modeling Dynamical Systems Using Neural ODEs: A Study On Scientific Machine Learning Models	Proceedings of the 39th ECMS International Conference on Modelling and Simulation (ECMS 2025)	2025-06-27
Knaus, Tanja	Autonomous Machine Listening: The Inetrsection of Voice as Data and Desire	Leverhulme Center for the Future of Intelligence, University of Cambridge	2025-06-10
Kuruge, Darshana Abeyrathna; Hafver, Andreas	Increasing confidence in AI models by explaining uncertainty in predictions	ESREL & SRA-E 2025: EUROPEAN SAFETY AND RELIABILITY & SOCIETY FOR RISK ANALYSIS EUROPE CONFERENCE	2025-06-17
Kåsene, Vebjørn; Lison, Pierre	Following Route Instructions using Large Vision-Language Models: A Comparison between Low-level and Panoramic Action Spaces	Proceedings of the 8th International Conference on Natural Language and Speech Processing (ICNLSP-2025)	2025-08-27
Łajewska, Weronika; Balog, Krisztian	GINGER: Grounded Information Nugget-Based Generation of Responses	48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '25)	2025-07-13

## PUBLISHED CONFERENCE PAPERS CONT'D

NAME	TITLE	TYPE/PLACE	DATE
Shamba, Abdul Kazeem; Bach, Kerstin; Taylor, Gavin	Contrast All The Time: Learning Time Series Representation from Temporal Consistency	28th European Conference on Artificial Intelligence, 25-30 October 2025, Bologna, Italy – Including 14th Conference on Prestigious Applications of Intelligent Systems (PAIS 2025)	2025-10-21
Storset, Lilja Charlotte; Rønningstad, Egil; Mæhlum, Petter; Velldal, Erik; Øvrelid, Lilja	Mixed Feelings: Cross-Domain Sentiment Classification of Patient Feedback	Proceedings of the Joint 25th Nordic Conference on Computational Linguistics and 11th Baltic Conference on Human Language Technologies (NoDaLiDa/Baltic-HLT 2025)	2025-03-03
Yadav, Vandana; Gulla, Jon Atle; Özgöbek, Özlem, Zhang, Lemei	News Timeline Summarization: Recent Methods	Natural Language Processing and Information Systems	2025-07-01

## REPORTS

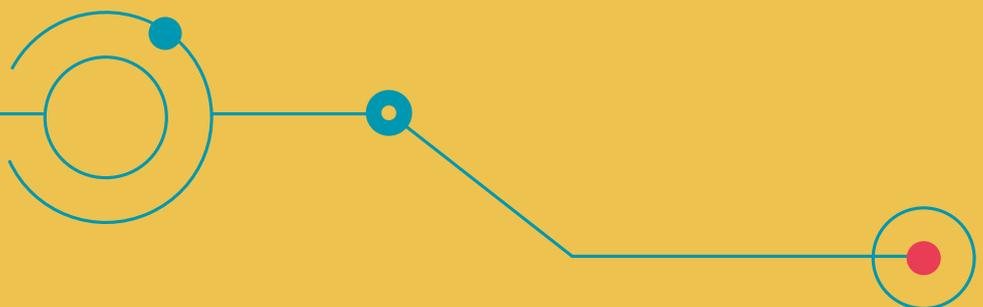
NAME	TITLE	TYPE/PLACE	DATE
Gunnes, Mari; Thaulow, Kristin; Sand, Kari; Ingvaldsen, Jon Espen; Kille, Benjamin Uwe; Thaulow, Jorunn	AIFAL-prosjekt: Allmennlegers perspektiv på muligheter og begrensninger ved bruk av KI på fastlegekontoret	SINTEF-rapport	2025-03-12

## OTHER PUBLICATIONS

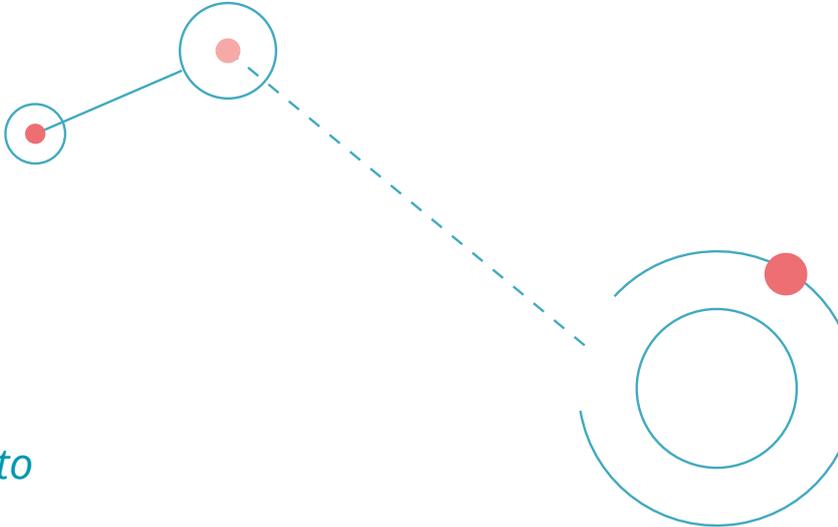
NAME	TITLE	TYPE/PLACE	DATE
Nikolov, Nikolay	Scalable Design, Deployment, and Optimization of Big Data Pipelines Across the Cloud/Edge/Fog Computing Continuum	PhD thesis (UiO)	2025-04-30
Michałowska, Katarzyna	Physics-Informed Machine Learning for Scientific and Industrial Applications	PhD thesis (UiO)	2025-05-23
Bernard, Nolwenn Marie Emilie	User Simulation for the Development and Evaluation of Conversational Information Access Agents	PhD thesis (UIS)	2025-06-16
Łajewska, Weronika	Grounded and Transparent Response Generation for Conversational Information-Seeking Systems	PhD thesis (UIS)	2025-06-20

OTHER

NAME	TITLE	TYPE/PLACE	DATE
Bernard, Nolwenn; Suresh, Sharath Chandra Etagi; Balog, Krisztian; Zhai, ChengXiang	SimLab: A Platform for Simulation-based Evaluation of Conversational Information Access Systems	Technical report, arXiv	2025-07-07
Gulla, Jon Atle et. AI Expert group on the responsible deployment and use of AI Assistants	KI-assistenter i arbeidslivet- en praktisk guide	Digitaliserings- og forvaltningsdepartementet	2025-06-16
Iana, Andreea;Treuillier, Céline; Yadav, Vandana, Kille, Benjamin; Lommatzsch, Andreas; Özgöbek, Özlem	The 13th International Workshop on News Recommendation and Analytics (INRA 2025)	RecSys '25: Proceedings of the Nineteenth ACM Conference on Recommender Systems	2025-09-07
Liu, Peng	Understanding Language Modeling Paradigm Adaptations in Recommender Systems: Lessons Learned and Open Challenges.	Tutorial at 27th European Conference of Artificial Intelligence (ECAI), Santiago de Compostela, Spain, 2024	2024-10-19
Liu, Peng	Part 3: Optimization Objectives of LLM-based RSS	Tutorial at 27th European Conference of Artificial Intelligence (ECAI), Santiago de Compostela, Spain, 2024	2024-10-19
Riemer-Sørensen, Signe; Stasik, Alexander Johannes; Djupskås, Aslak; de Leyos, Aaron	Proceedings of NORA's annual conference 2025	Nordic Machine Intelligence	2025-08-04
Zhang, Lemei	Understanding Language Modeling Paradigm Adaptations in Recommender Systems: Lessons Learned and Open Challenges.	Tutorial at 27th European Conference of Artificial Intelligence (ECAI), Santiago de Compostela, Spain, 2024	2024-10-19







# The Innovation Advisory Board

*Exclusive group of innovators to advise NorwAI on innovation*

The Innovation Advisory Board (IAB) will provide advice on how to create innovations from research for the NorwAI partners. The IAB will follow up on the innovation results on a regular basis and help to monitor the overall progress following SFI success criteria for innovation and commercialization. The Innovation Advisory Board is planned to meet with the Center director and the Innovation director twice a year.

The Innovation Advisory board underwent a re-organization in 2025. As from 2026 the board will have a new mandate, initially focusing on language models in the health domain and on startups. The members of the Innovation Advisory Boards will be

## MEMBERS OF THE INNOVATION ADVISORY BOARD FROM 2026



**Kjartan Landgraff Kalstad**  
Advisor, entrepreneur and investor



**Sverre Fjeldstad**  
Independent advisor, growth investor, venture builder



**Per Olav Østbyhaug**  
Attending Physician and Section Chief, St. Olavs Hospital



# Students as an innovation source



**Hanne Dahl Vonen**

Entrepreneur and student  
NTNU

After two and a half years in Boston, I have moved back to Trondheim and NTNU, both to complete my studies and to work with medical innovation through the startup Hugin Medical. There is much to be said about American universities and the way they structure education, for better and for worse. However, my time in Boston has given me many new perspectives on how we at NTNU can become even stronger in innovation.

Back in Norway, it is not as common to work with innovation within the university ecosystem. Of course, there are certain courses and programs with a strong focus on innovation, but there is still significant potential for a more systematic effort where innovation is integrated into education. The goal is not, of course, that every NTNU student should start a company. As a society, however, we have much to gain if students learn to think systematically about how to develop innovative solutions and new ways of working, regardless of where they end up after graduating.

In Trondheim, we are fortunate to have a strong concentration of highly skilled academic environments. This creates tremendous potential for innovation. For us at Hugin Medical, this has been essential. Being close to the medical expertise at St. Olavs Hospital and NTNU's strong AI community helps us develop high-quality medical technology.

## **Urgent need**

The Norwegian healthcare sector is world-leading, but it is under pressure due to tight budgets and an aging population. The sector

**Hanne Dahl Vonen** is a fifth-year medical student at NTNU. She holds an MSc in Epidemiology from Harvard T.H. Chan School of Public Health, where she was an Aker Scholar, and she also has previous experience from Novo Nordisk External Innovation. She is the founder and CEO of the startup Hugin Medical, where she and her team use AI to create a more sustainable healthcare system.



is in urgent need of innovation and technology that can reduce workload. We know that AI is entering healthcare and that much of the future innovation in health will be driven by AI. Yet doctors and healthcare professionals rarely have the time and resources to focus on innovation, which is why we should look to students as a valuable resource.

Some concrete suggestions for how we can strengthen innovation at Norwegian universities by leveraging students as a resource:

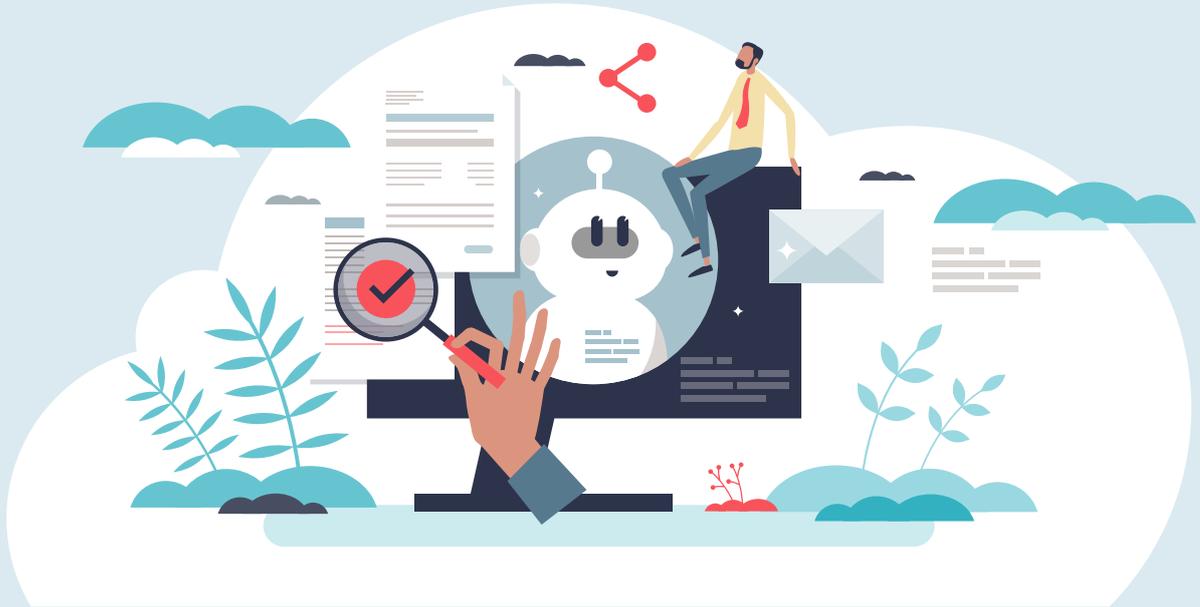
- ◆ **Integrate innovation more systematically into education:** One of the major differences between Norway and the U.S. is the strong emphasis on innovation at American universities. In several of the programs at my faculty in the U.S., innovation was a mandatory part of the curriculum. This was not necessarily about starting a company, but about design thinking and systems innovation. There is a fundamental belief that learning to think creatively and innovatively is an essential part of higher education's role.
- ◆ **Provide flexibility similar to that offered to student-athletes:** At NTNU, athletes are offered flexibility in their study programs, flexibility regarding mandatory activities, and accommodations for midterms and

exams. We could offer similar arrangements for students running their own startups, who currently do not receive the same opportunities from the university. This is despite frequent discussions about innovation being key to our future economy. Just as NTNU seeks to attract elite athletes, why not aim to become Norway's best university for entrepreneurs?

- ◆ **Let students tackle real-world challenges:** In the US, Innovation Fellowships are common, where students work on challenges in a real-world setting. This could be particularly valuable in healthcare. In the health sector, there are strong incentives to avoid mistakes, but fewer incentives to innovate. Strengthening collaboration with students, who often have more time, flexibility, and the ability to see problems with fresh eyes, could make a significant difference. This is something we could achieve in Norway as well.

Trondheim, home to some of the country's most talented students, has a strong opportunity to build bridges between medicine and technology. We can facilitate more interdisciplinary collaboration and provide healthcare students with opportunities to test their skills in innovation and entrepreneurship.

# PhD Defences in 2025



The following PhD candidates, associated to the SFI, successfully defended their theses in 2025:



### **Betül Bayrak on 19 February**

Post-hoc eXplainable Artificial Intelligence Methods: Counterfactuals and XCBR Applications.



### **Aleksej Logacjov on 14 March**

Large-scale Self-supervised Learning for Enhancing Accelerometer-based Human Activity and Sleep Recognition.





**Nikolay Nikolov on 30 April**

Scalable Design, Deployment, and Optimization of Big Data Pipelines Across the Cloud/Edge/Fog Computing Continuum.



**Yanzhe Bekkemoen on 6 May**

Explainable Reinforcement Learning (XRL): Simplifying Agent Behavior.



**Katarzyna Michałowska on 23 May**

Scalable Design, Deployment, and Optimization of Big Data Pipelines Across the Cloud/Edge/Fog Computing Continuum.



**Nolwenn Bernard on 16 June**

User Simulation for the Development and Evaluation of Conversational Information Access Agents.



**Weronika Łajewska on 20 June**

Grounded and Transparent Response Generation for Conversational Information-Seeking Systems.



Read more about the work of these PhDs on the next pages.

---

## Our PhDs

*NorwAI is well into its PhD program. We asked their supervisors to give a brief summary of the doctoral work by the PhDs that defended their thesis in 2025.*

## Weronika Łajewska and Nolwenn Bernard Advancing Conversational AI and its Evaluation



Weronika Łajewska



Nolwenn Bernard

The development of advanced conversational agents requires a two-fold approach: we must create systems that are transparent and trustworthy by design, while simultaneously developing robust methods to measure their performance. Through their respective doctoral research, Weronika Łajewska and Nolwenn Bernard have addressed these challenges from two complementary perspectives.

Weronika's research focuses on how conversational agents synthesize information from relevant sources into concise responses. In her thesis, "Grounded and Transparent Response Generation for Conversational Information-Seeking Systems," she has developed methods for building systems that generate transparent and grounded responses. By ensuring that information is anchored in verifiable sources and enabling systems to disclose their limitations, she allows users to navigate complex information needs with a high degree of trust.

Parallel to this, Nolwenn has tackled the challenge of automatic evaluation of conversational agents in her thesis, "User Simulation for the Development and Evaluation of Conversational Information Access Agents."

For conversational systems to improve, reliable and reproducible testing methods are essential. Nolwenn's use of user simulation allows for the automatic evaluation of various system aspects at scale, significantly reducing the need for resource-intensive manual testing.

In practice, these findings demonstrate that advancing conversational AI requires a symbiotic relationship between system capabilities and evaluation methods. It is not possible to truly improve agent capabilities without having the means to accurately measure their effectiveness. By using user simulation to break the traditional bottleneck of human testing, this research enables a faster, more cost-effective cycle where new features can be rigorously stress-tested for safety and robustness.

## Katarzyna Michałowska

# Exploring physics-informed machine learning



Katarzyna Michałowska

Machine learning can model complex systems directly from data, but often struggles to generalize when data are noisy or limited. Physics-based models provide consistent and stable predictions, but can be computationally demanding and rely on simplifying assumptions.

In her PhD thesis «Physics-informed Machine Learning for Scientific and Industrial Application» Katarzyna Michałowska explores physics-informed machine learning, with a particular focus on neural-operator-based methods for industrial and scientific applications.

### The thesis presents three main contributions

- ◆ First, it proposes a framework for condition monitoring in safety-critical industrial systems that combines efficient unsupervised learning with expert validation. Using operational data from maritime vessels, the resulting fault detection algorithm identifies historical faults and detects a real-time thruster leakage earlier than conventional systems, and it is currently commercially deployed by Brunvoll AS.
- ◆ Second, it introduces the DON-LSTM architecture for multi-resolution learning in dynamical systems, combining deep operator networks with recurrent sequence modeling to reduce reliance on expensive high-resolution data while outperforming single-resolution baselines.
- ◆ Third, it extends neural operators to long-time-horizon prediction and extrapolation beyond training data, and investigates how soft and hard constraints, such as conservation of mass and momentum, can be incorporated to improve reliability in extended simulations.

## Yanzhe Bekkemoen

# How intelligent agents works – and can we trust them?



Yanzhe Bekkemoen

Many intelligent agents that learn by interacting with their environment, be they chess players or autonomous cars, achieve impressive results. But how do they actually work? And how do we know if we can trust them? PhD candidate Yanzhe Bekkemoen has investigated these questions in his doctoral thesis, which he defended in 2025.

Yanzhe's focus is on explainable reinforcement learning – and is about how to make intelligent agents more understandable and thus easier to trust. He has developed methods to simplify the agents' behavior in three ways:

- ◆ by reducing the complexity of the states they observe
- ◆ by learning simplified representations of their logic
- ◆ by explaining which actions actually affect the outcome

An important, consistent finding is that you can often simplify the logic of the agents' behavior significantly without worsening performance.

This can lead to more transparent and relatable behavior.

# Nikolay Nikolov

## Developing methods for Big Data Pipelines



Nikolay Nikolov

Big Data Pipelines (BDPs) are essential for transforming vast, heterogeneous data streams into actionable insights. Yet as data increasingly flows across Cloud, Fog, and Edge environments, designing, deploying, and optimizing such pipelines becomes a complex technical and organizational challenge.

In his PhD thesis, “Scalable Design, Deployment, and Optimization of Big Data Pipelines Across the Cloud/Fog/Edge Computing Continuum”, Nikolay Nikolov addresses this challenge by developing methods that make BDPs more scalable, efficient, and accessible across distributed infrastructures.

**The thesis is driven by three core questions:**

- ◆ how to conceptualize BDPs in a high-level, domain-friendly manner;
- ◆ how to automate their deployment across heterogeneous resources;
- ◆ how to optimize their execution without relying on costly full-scale benchmarking.
- ◆ Together, these questions target the full lifecycle of data pipelines, from modeling and implementation to profiling and scheduling.

A central contribution is the introduction of a domain-specific language (DSL) for high-level modeling of BDPs. By abstracting technical details and separating design-time from run-time concerns, the DSL enables domain experts, data scientists, and DevOps engineers to collaborate effectively. This conceptual layer is combined with a container-based execution approach, where individual pipeline steps are encapsulated as reusable, scalable components. Through container orchestration and message-oriented middleware, the approach supports fine-grained scaling of individual steps, avoiding bottlenecks and enabling efficient execution across the Computing Continuum.

Beyond design and deployment, the thesis introduces a novel dry-run profiling method for Big Data Pipelines. Instead of requiring expensive full-scale benchmarks, the dry-run approach executes pipelines with representative



sample data to estimate resource requirements, execution times, and scalability characteristics. This profiling data serves as the basis for a context-aware scheduling model implemented in the CATS scheduler, which optimizes makespan and cost while accounting for data locality, resource heterogeneity, and inter-step dependencies.

The methods are validated in real-world deployments in telemedicine and Industry 4.0 settings, demonstrating improved scalability, reduced processing time, and more efficient resource utilization. These case studies show that the proposed approach is not only theoretically sound but also practically applicable in demanding, data-intensive environments.

Taken together, the thesis advances the state of the art in Big Data processing by integrating

modeling, automation, profiling, and scheduling into a coherent framework for the Cloud/Fog/Edge Computing Continuum. By lowering technological barriers and emphasizing usability, scalability, and cost efficiency, this work helps make advanced data pipeline solutions accessible to a broader range of stakeholders and application domains.

## Betül Bayrak

# Addressing the growing complexity of AI systems



Betül Bayrak

Artificial Intelligence systems are increasingly used in high-stakes domains, yet their growing complexity often makes them difficult to understand and trust. In her PhD thesis, “Post-hoc eXplainable Artificial Intelligence Methods: Counterfactuals and XCBR Applications”, Betül Bayrak addresses this challenge by developing user-centered methods for generating and evaluating instance-based explanations, with a particular focus on counterfactuals.

### The thesis is structured around three overarching questions:

- ◆ how to integrate domain knowledge and human-centered design into explanation systems;
- ◆ how to generate high-quality, plausible counterfactual explanations;
- ◆ how to evaluate explanations in ways that align with user needs and preferences.

Across five publications, Bayrak develops a coherent framework that combines methodological advances with empirical validation.

A central contribution is the development of PertCF, a perturbation-based counterfactual generation method that balances minimal change with stability and plausibility. By integrating feature attribution techniques with structured perturbations, the method produces counterfactuals that are both interpretable and robust. In parallel, the thesis advances case-based reasoning approaches through twin XCBR systems that provide supportive and contrastive explanations, explicitly incorporating domain knowledge to enhance trust and practical relevance.

Beyond counterfactual generation methods, the thesis makes a substantial contribution



to the evaluation of explainability. The CEval toolkit provides a comprehensive framework for assessing instance-based explanations using objective metrics such as fidelity, proximity, robustness, and diversity. Complementing this, empirical user studies investigate how explanation metrics relate to user preferences, highlighting the importance of aligning technical quality measures with human perception and comprehensibility.

Taken together, the thesis demonstrates that effective explainable AI requires a tight integration of algorithmic design, evaluation methodology, and human-centered considerations. By systematically addressing

the generation, representation, and evaluation of counterfactual explanations, this work contributes to making AI systems more transparent, trustworthy, and practically usable across application domains.





---

## NorwAI Innovate '25

### A unique AI showroom

In 2025, NorwAI Innovate Conference again proved itself to the Norwegian AI community to be maybe the most interesting AI tech conference in the country. An extraordinary line-up of keynotes and a mixed academia and business audience provided a melting pot for cross professional collaboration, cross discipline exchange and cross generation meetings.

#### HOSTS



**Professor Jon Atle Gulla**  
NorwAI Center Director



**Rolf Dyrnes Svendsen**  
NorwAI Communications Manager



**Karolina Storesund**  
NorwAI Administrative Coordinator



*NorwAI Innovate 2025 kicked off with an audience from partners, academics and interested participants from both public organisations and private companies.*

**In 2025, NorwAI organized the event for the fifth time. Once again, the conference attracted national interest from a wide audience with more than 150 attendees.**

NorwAI and NTNU Technology Transfer Office also organized an IP and Innovation Seminar for Researchers where IPR and business strategies for AI inventions were discussed.

NorwAI Innovate is a unique showroom for what to expect from the research center itself. The conference also examines specific AI verticals and invites national and international experts to tell on use cases and innovations in the ongoing

AI transition in society. The 2024 version focused on implementations of AI in business with an extra focus on the health sector, research and innovation updates, presentations from two of the new AI centers, reports on how organizations had built AI competence among thousands of employees. The closing keynote was delivered by Karianne Tung, Minister of Digitalisation and Public Governance, looking ahead.

Supplementary exhibits, including a poster and a demo of NorwAI's latest LLM, the NorwAI-Magistral-28B-Reasoning was available for attendees to test on their own. Above all, the lunch-to-lunch conference is a meet & greet event and a chance to get to know leading people in both the private and public sector concerned and engaged in AI developments.





 **NorwAI**  
NORWEGIAN RESEARCH CENTER  
FOR AI INNOVATION

»» **NorwAI-  
Magistral-  
24B-reasoning**

*Lemei Zhang at the conference, demonstrating the latest developments of the NorwAI large language models.*

## SPEAKERS



**Abdul-Kazeem Shamba**  
NTNU



**Gunnar Bovim**  
Research Council of Norway



**Juan Carlos Lopez Calvet**  
Schibsted Media



**Nhien Nguyen**  
NTNU



**Annita Fjuk**  
Digital Norway



**Hanne Vonen**  
Hugin Medical



**Karianne Tung**  
The Government



**Rolf Dyrnes Svendsen**  
NTNU



**Asgeir Sørensen**  
NTNU



**Heidrun Åm**  
NTNU



**Kerstin Bach**  
NTNU



**Signe Riemer-Sørensen**  
SINTEF



**Benjamin Kille**  
NTNU



**Hilde Lovett**  
Norw. Directorate of Health



**Ketil Thorvik**  
Hemit HF



**Simone Casolo**  
Cognite



**Bjørn Fjukstad**  
DIPS



**Haakon Thue Lie**  
Dehns/NTNU



**Kjersti Wold**  
SpareBank 1 SMN



**Terje Brasethvik**  
NTNU



**Christian P. Halvorsen**  
Vend



**Johannes Voll Kolstø**  
NR



**Lars Bungum**  
NTNU



**Thomas Nicolai Bredeli**  
NRK



**Egil Aslak A. Hagerup**  
NRK



**Jon Atle Gulla**  
NTNU



**Lemei Zhang**  
NTNU



**Yngvar Ugland**  
DNB



**Elise Lindeberg**  
Skygard



**Jon Espen Ingvaldsen**  
Medbric/NTNU



**Lilja Øvrelid**  
UiO



**Øyvind Solberg Moe**  
AI Factory



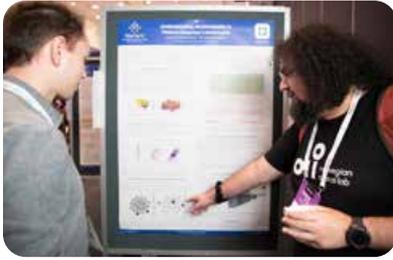
**Eric Monteiro**  
NTNU



**Jon Nordby**  
SoundSensing



**Michail Giannakos**  
NTNU





# Best poster award

## The winner takes it all

Alae Ajraoui, Nhien Nguyen, and Alf Steinar Sætre won the Best Poster Award at NorwAI Innovate '25 with their work titled "Implementing AI in Practice: Insights from a Multi-Case Study Across Organizational Levels".

This empirical study addresses the research question: How do organizations transform in order to implement and scale AI?

The research draws on a multiple-case study of six Norwegian companies across diverse industries. Through 25+ semi-structured interviews and triangulated data sources (including internal documents and media coverage), the study examines how actors at strategic, tactical, and operational levels coordinate and adapt to integrate AI at scale.

### KEY FINDINGS:

◆ The findings reveal the set of practices and processes led by actors at each organizational level to facilitate the adoption of AI on a large scale.

◆ The study extends classical technology implementation theories to better account for the complexity and dynamism of AI.

◆ It applies the dynamic capabilities lens to explain how companies sensed the threat posed by the rise of large language models (LLMs), seized internal and external resources, and reconfigured teams and workflows to harness the potential of AI.

There were 10 accepted poster contributions to the competition, and this poster received the most support from the NorwAI Innovate '25 community.

*Research director Kerstin Bach handing PhD candidate Alae Ajraoui and Associate Professor Nhien Nguyen the Best Poster Award at NorwAI Innovate Conference 2025.*



## POSTER CONTRIBUTIONS AT NORWAI INNOVATE 2025

AUTHORS	TITLE
Ajraoui, A.; Nguyen, N.; Sætre, A. S.	Implementing AI in Practice: Insights from a Multi-Case Study Across Organizational Levels
Arustashvili, M.; Balog, K.	Natural Language User Interest Profiles for Scientific Literature Recommendation
Bungum, L.	Designing a Multidimensional Online Metric for Evaluation of Large Language Models
de Leyos, A.; Stasik, A.; Riemer-Sørensen, S.	Gaussian Process Kolmogorov-Arnold Networks for Probabilistic Multivariate Function Approximation
Djupskås, A.; Stasik, A.; Riemer-Sørensen, S.	Unreliable Uncertainty Estimates with Monte Carlo Dropout
Jekic, A.; Natsaridou, A.; Riemer-Sørensen, S.; Langseth, H.; Gundersen, O. E.	Examining the robustness of Physics-Informed Neural Networks to noise for Inverse Problems
Murgorgo, S.; Nguyen, N.; Mikalef, P.; Bach, T. A.	Responsible AI Practices in Organisations: A Systematic Literature Review
Sánchez-Díaz, X. F. C.; Mengshoel, O. J.	Understanding Multimodality in Feature Selection Landscapes
Steppe, J. A.; Nguyen, N.; Johannson, O. H. N.; Maric, S.	Key Competencies for Human-AI Collaboration and How Organizations Foster Their Development
Tuzen, A.; Nørnvåg, K.; Langseth, H.	Simplifying the Graph Similarity Learning Task with Graph Neural Networks

# NorwAI at Arendalsuka 2025

AI was still one of the hottest topics at the Arendal Week in 2025. NorwAI attended three meetings in the largest political gathering in the country.



From left, NorwAI Chair Sven Størmer Thaulow, CEO Trond Utne, Hemit HF and Minister Karianne Tung from the Government discussed AI in health at the Arendal Festival. Photo: NorwAI

## TUESDAY 12TH AUGUST 09:00

«When artificial intelligence becomes part of the healthcare team: From trial to profit»

Minister of Digitalization and Public Governance Karianne Tung and a team of experts meet to discuss AI needs in health. NorwAI chair Sven Størmer Thaulow joined the panel with his expertise from private business. NorwAI organized the session with our new partner Hemit HF and asked:

- ♦ What hinders the transition from promising AI projects to actual value creation in the health service?
- ♦ How can we ensure that the health field keeps up – and that good initiatives reach patients?
- ♦ What resources and structures are missing – and what can politicians, administration and the health sector do?



Research director Kerstin Bach (left) discussing with Anne Kjersti Fahlvik, director for Innovation in Business and Public sector at the Norwegian Research Council. Photo: NorwAI



NorwAI chair Sven Størmer Thaulow, discussing with the director of SFI Visual Intelligence Robert Jenssen. Photo: NorwAI

## TUESDAY 12TH AUGUST 14:00

«Research-driven innovation in AI: How do we strengthen it together?»

SFI Visual Intelligence at UiT and NorwAI organized the session where a panel discussed what should be done to strengthen research-driven innovation within AI in Norway? What opportunities lie in this strengthening – for

both the academic and industrial sectors? What are the challenges, and how should academia, the institute sector, and business solve these challenges together?

Research director Kerstin Bach and chair Sven Størmer Thaulow represented NorwAI.



Professor Jon Atle Gulla at Arendalsuka presenting NorwAI's approach to challenges at «A sea of data»  
Photo: NorwAI

## WEDNESDAY 13TH AUGUST 12:30

“A sea of data – new technology gives us the key to knowledge, security and value creation” organized by Department of Marine Technology at NTNU

NorwAI director Jon Atle Gulla represented the SFI where speakers presented Norway as a major power at sea – both as a manager, a creator of value and a knowledge nation.

Now we are on the verge of a technological breakthrough that could change everything: An explosion in data collection, new observation platforms and artificial intelligence are giving us the opportunity to understand the ocean in completely new ways.

---

# Framtidsfredag

## – Arendalsuka on Fast-Forward



*From the debate "When AI is part of your health team". From the right: Minister Karianne Tung, debate host Rolf Dyrnes Svendsen, NorwAI, and CTO Jon Espen Ingvaldsen, Medbric. Photo: Jorunn Thaulow.*

From the debate on "Norway's role in the big AI-world". From right Jon Atle Gulla, Karianne Tung and Tore Tennøe Photo: Jorunn Thaulow



This one-day happening in Steinkjer on June 20th presented more than 100 events, debates, presentations and lectures. The event was sold out to a 2400 paying audience before the program was even published!

**NORWAI ORGANIZED THREE EVENTS AT THE FESTIVAL:**

*"The big conversation on AI and the arts"* with author Bjørn Vatne, head of Writers Association and Associate Professor Inga Strümke, NTNU - the meeting was in cooperation with Norwegian Open AI Lab and partner Sparebank 1 SMN.

at Norwegian Board of Technology and Jon Atle Gulla, Director of NorwAI.

*"Norway's role in the big AI-world"* with Minister of Digitalisation and Public Governance Karianne Tung, CEO Tore Tennøe

*"When AI is part of your health team"* also with Minister Karianne Tung, Innovation Manager Ketil Thorvik of Hemit HF, CTO Jon Espen Ingvaldsen of Medbric and Senior Physician and Associate Professor Bjørnar Grenne - this meeting was organized in cooperation with Hemit.

# Framtidsfredag

Framtidsfredag wants to be a dynamic and attractive meeting place with a focus on the future, business development and social development for both private individuals and businesses.

be sustainability, regional policy, culture, environment, technology, innovation, industry, labor market, politics, emergency preparedness and much more.

We will ensure that we focus on the broad perspective of the future. Current topics will

Our ambition is to bring people and generations together across geography to an arena where we build relationships.

FramtidsFredag is owned by a foundation with private companies and public organizations as partners.



# Which model can I trust?



**Benjamin Kille**  
Associate Professor  
NorwAI, NTNU



**Dan Saastrup Smart**  
Principal AI Specialist  
Alexandra Institute, Denmark

Europe is home to a diverse set of languages. AI tools have established as helpers when it comes to answering questions, generating drafts of documents, and many more tasks. But, which model should you pick? Which model can you trust?

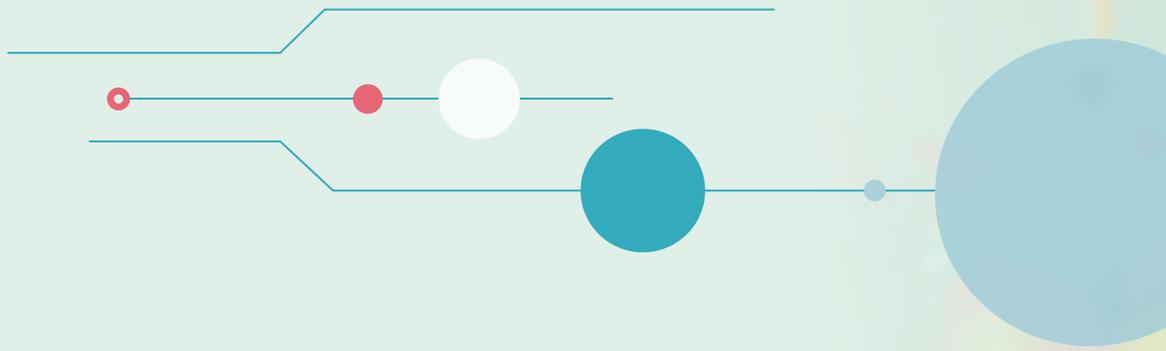
The EU financed project TrustLLM engages in research about Large Language Models (LLMs) with a particular focus on trustworthiness. In work package no. 7. Multilingual and Multi-metric LLM Evaluation, researchers look into ways to evaluate different LLMs.

This involves collecting data about tasks, and desired output in different languages. Existing models are then applied to the tasks, and an

automated benchmarks computes scores reflecting on how well the models respond.

## **The ScandEval tool**

Trust LLM have used ScandEval, a benchmark for Scandinavian languages, as the starting point. The benchmark has grown impressively over the last two years which prompted us to rename it to EuroEval. Just in 2025, EuroEval has added support for 23 languages. These include



languages from the Latin language family (e.g., Spanish), Slavic languages (e.g., Ukrainian), and Baltic languages (e.g., Lithuanian). EuroEval now support 30 language.

EuroEval support a set of tasks including answering questions, European Values, and simplifying texts. On the platform, users can find leaderboards that show which of the many models perform best for specific tasks, or languages. The platform offers information about model performance in a tabular or visual format. The visual presentation includes scatter plots comparing the model size to relative performance, and radar plots for a more detailed comparison of specific models.

### **Alignment to European values**

Assessing models' alignment with European values represents a major focus of TrustLLM. The long-running European Values Study supplied us with questions that reflect on Europeans' values, and beliefs.

These values constitute a subject of our research, but are not included in the ultimate ranking of models. There is a common impression that commercial models are optimized for the United States. The scores for the European Values can indicate whether models reflect European values.

Besides the challenges related to evaluation, the developers face technical challenges. We strive to include as many models as possible. Each model comes with a specific software stack including a tokenizer, the model architecture, and a generation strategy. The TrustLLM team, in particular the colleagues at the Alexandra Institute, work hard to enable evaluation of a diverse set of models, both commercial and open source.

The figures show performances of the models developed by NorwAI. Figure A depicts the position of the NorwAI-Mixtral-8x7B model for Norwegian. Figure B compares the same model to normistral-11b-long. We see that the NorwAI model achieves better performance in most evaluation tasks.

In 2026, EuroEval will grow. We plan to add several new tasks and corresponding leaderboards such as logical reasoning, and detecting hallucinations or biases. Besides, we will add more languages to cover as much of the European language landscape as possible. Finally, we will try to optimize the tasks with human preference data coming from the European LLM arena.

The main objective for TrustLLM is the development of an open, trustworthy and factual LLM, initially targeting the Germanic languages. This will create the foundation for an advanced open ecosystem for next generation modular and extensible European trustworthy, sustainable, and democratized LLMs. The focus on Germanic languages can serve as a

blueprint for future activities in other families of languages.

The TrustLLM project and the surrounding ecosystem will enable, support, and improve context-aware human-machine interaction in a wide range of applications.

# The performance of the models developed by NorwAI

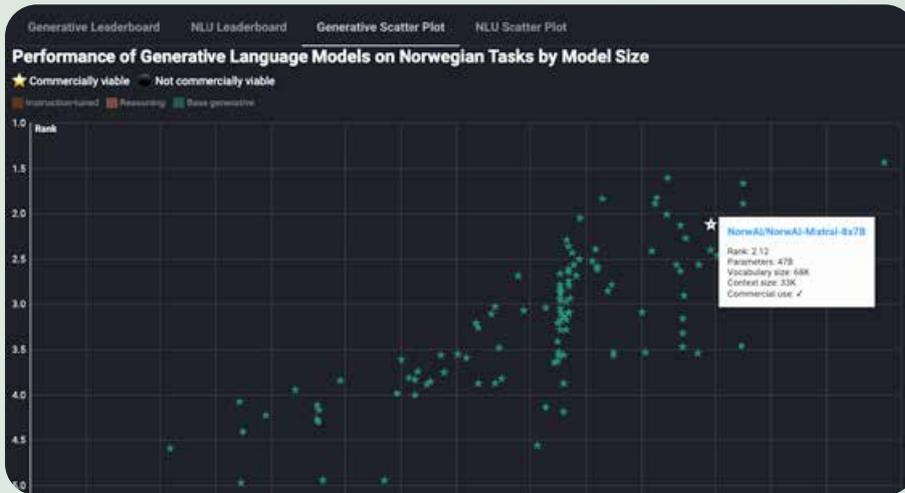
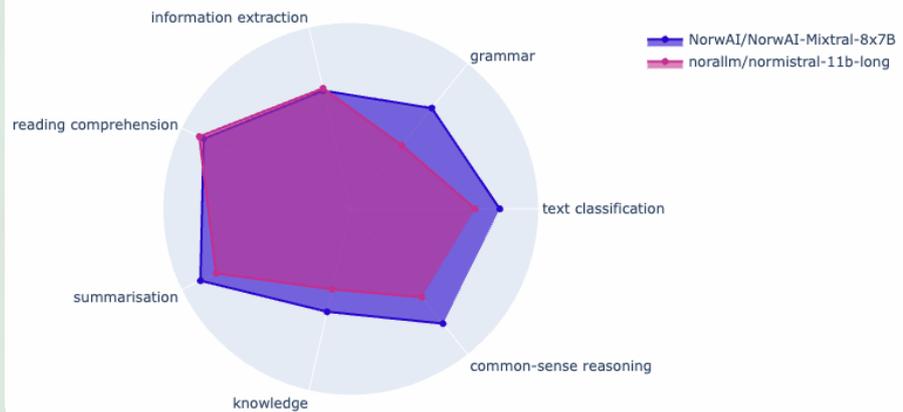


Figure A depicts the position of the NorwAI-Mixtral-8x7B model for Norwegian.

Figure B compares the same model to normistral-11b-long. We see that the NorwAI model achieves better performance in most evaluation tasks.

Rank Score on Norwegian Language Tasks



## Trust LLM project partners:

- ◆ Linköping University, Sweden
- ◆ Fraunhofer, Germany
- ◆ Forschungszentrum Jülich, Germany
- ◆ Lindholmen Science Park, Sweden
- ◆ Mideind (MID), Iceland
- ◆ University of Iceland, Iceland
- ◆ University of Copenhagen, Denmark
- ◆ Alexandra Institute, Denmark
- ◆ Norwegian University of Science and Technology, Norway
- ◆ Nederlandse Organisatie voor Toegepast Natuurwetenschappelijk Onderzoek, Netherlands
- ◆ Academy for Artificial Intelligence (AKI), Germany

## External collaborators

- ◆ The Centre for Language Technology at the University of the Faroe Islands, Faroe Islands
- ◆ ScaDS.AI (Center for Scalable Data Analytics and Artificial Intelligence), Germany

**TrustLLM is funded by the European Union.**



# Highlights from Media Labs Days #9



The 9th edition of Media Labs Days – a yearly European gathering by the Worlds Association of News Publishers - took place on 16 – 17 October in Trondheim, Norway. The two days event, hosted by NorwAI partner NRK and regional newspaper Adresseavisen supported by NorwAI brought together the Global Alliance for Media Innovation (GAMI) community to explore how innovation is reshaping the media landscape.

Discussions ranged from sensor journalism to XR technologies, and of course AI in the context of fact-checking or LLM's for smaller languages.

## LLMs in Scandinavia and beyond

Jon Atle Gulla, Director of NorwAI, explained the value of developing Norwegian LLMs. Creating a national model would not only better represent Norwegian culture and language but also ensure copyrighted material is used responsibly, in

collaboration with the rights holders rather than through international models.

## Jon Atle Gulla highlighted:

- ◆ When pre-training an existing international model with national training data, it produced models that perform better than the international models for this particular language, he said.



*A multinational seminar for innovation decisionmakers from the European media industry held their yearly convention in Trondheim in October 2025. Photo: Adresseavisen*



To complete this session, Saskia Lensink, Product manager at GPT-NL in The Netherlands and Partner in the project Trust LLM, highlighted the public concerns about big tech companies' use of personal data. She presented TrustLLM, a European initiative aiming to develop a trustworthy LLM, targeting Germanic Languages. NorwAI represents Norway in the multinational EU project.

# The Forums for shared insights

NorwAI Forums are bi-annual internal summits for NorwAI partners for sharing, discussing and setting goals for across industry and academic collaboration.

The two forums have different focuses. While the Oslo Forum headed by SINTEF Digital has an innovative perspective, the Trondheim version headed by NTNU takes a closer look into research.

Both forums are cross-company and cross-institution gatherings with representatives from the NorwAI community attending and presenting use-cases and insights.

## NORWAI FORUM: TRONDHEIM, APRIL 9TH

At the NorwAI HQ at NTNU, The Forum drew a number of participants from different partners.

### Agenda:

- ◆ Research and recent work in WP LAP – Language and Personalization
- ◆ Research and recent work in WP INNOECO – AI Innovation Ecosystems
- ◆ Research and recent work in WP TRUST – Trustworthy AI
- ◆ Research and recent work in WP HYB – Hybrid AI Analytics
- ◆ Extended research and recent work in WP SOC – AI in Society
- ◆ Recent work in WP DATA – Data Platforms and Streaming Data
- ◆ Representation learning in Sensor Data
- ◆ Boundary spanning, boundary spanner and knowledge transfer
- ◆ Introducing Innovation Touchpoints
- ◆ Viewpoints from partners Telenor, SpareBank1 SMN and Schibsted Marketplaces



*Andreas Hafver, DNV on uncertainty in AI*



*Nikolay Nikolov, SINTEF on automating compliance processes in AI pipelines*

## NORWAI FORUM: OSLO NOVEMBER 19TH

At the DNV HQ at Høvik with representatives from both Oslo and Trondheim based companies

### Agenda: Trustworthy AI in Practice

- ◆ Keynote DNV: Explaining Uncertainty in AI
- ◆ Keynote SINTEF: Automating Compliance Processes in AI Pipelines
- ◆ Keynote Telenor: Creating an Innovation Culture
- ◆ Keynote NTNU: Evaluating Counterfactual Explanations: Metrics, Users and New Frontiers

- ◆ Keynote NTNU: Privacy Preserving Machine Learning
- ◆ Group work: joint discussions on how partners can strengthen innovation through trustworthy AI.



*Espen Sund, NTNU on privacy preserving ML*

# Recruitment PhD-students

STARTED IN 2025



## Espen Sund

Nationality: Norway  
Topic: Parameter estimation under different models of trust  
Started: January 2025  
Main supervisor: Staal Vinterbo, NTNU  
NorwAI Work Package: TRUST



## Mariam Arustashvili

Nationality: Georgia  
Topic: Large Language Models for Recommendation  
Started: April 2025  
Main supervisor: Krisztian Balog, University of Stavanger  
NorwAI Work Package: LAP



## Pål Furu Kamsvåg

Nationality: Norway  
Topic: Evaluating Artificial Intelligence – Bridging the Gap Between Benchmark Success and Real-World Utility  
Started: August 2025  
Main supervisor: Eric Monteiro, NTNU  
NorwAI Work Package: SOC



## Christoph Eder

Nationality: Austria  
Topic: Enhancing data quality and integration through data augmentation and generative AI  
Started: September 2025  
Main supervisor: Benjamin Kille, NTNU  
NorwAI Work Package: DATA

# Recruitment Postdocs

STARTED IN 2025



## Betül Bayrak

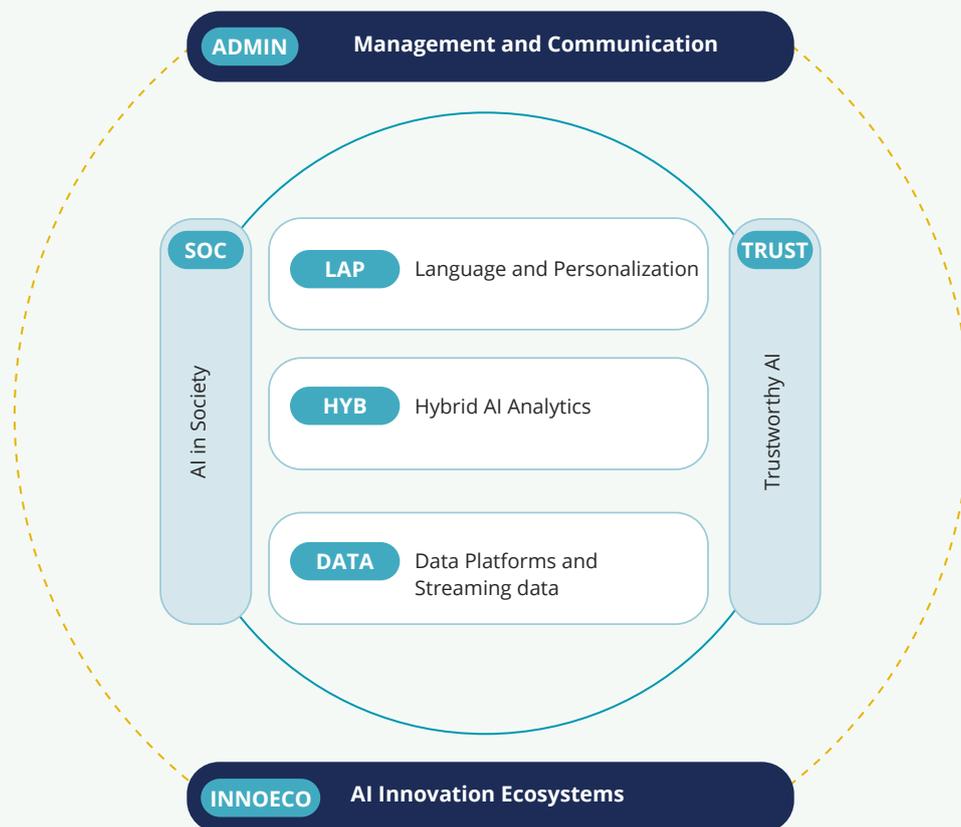
Nationality: Tyrkiye  
Topic: Explainable Artificial Intelligence  
Started: February 2025  
NorwAI Work Package: SOC



---

# The work packages

*NorwAI consists of seven work packages, which have been organized as described in this figure:*



## ADMIN: MANAGEMENT AND COMMUNICATION

This work package concerns the day-to-day operations of the center as well as communication.

**Work package leader: Jon Atle Gulla, Professor, NTNU**



## SOC: AI IN SOCIETY

This work package examines societal aspects of artificial intelligence technologies as they are developed, debated, and implemented. We focus on selected empirical cases, including

- Values in language models
- Societal aspects of voice and speech emotion recognition technologies
- Standardization of AI technologies and ethical, legal, and social aspects.

**Work package leader:**

**Susanne Bauer, Professor, University of Oslo** (until August -25)

**Eric Monteiro, Professor, NTNU** (from Sep -25)



## TRUST: TRUSTWORTHY AI

The purpose of work package TRUST is to reinforce a common understanding of safe and responsible AI, specifically:

- Establish trust in safe and responsible AI
- Ensure privacy-preserving in AI technologies
- Create guidelines for sustainable and beneficial use of AI
- Develop principles for explainable and transparent AI
- Develop principles for independent assurance of AI deployment

Trust in AI is a necessary condition for the scalability and societal acceptance of these technologies. Without trust, innovation can be stalled. This research investigates, from an interdisciplinary perspective, the multiple dimensions of trust raised by the deployment of AI and builds tools, methods, and a framework for assuring the safe and responsible deployment of AI in industry and society. This work package aims to answer the question: How can such tools address the safety and needs of individuals, organizations and society at large, addressing both non-technical and

technical issues? The research will address issues related to safety, explainability, transparency, bias, privacy and robustness, as well as human-machine interactions and co-behavior all in the context of industry regulations and societal expectations.

**Work package leader: Andreas Hafver, Team Leader**  
**- Emerging Technologies, Group Research & Development, DNV**



## LAP: LANGUAGE AND PERSONALIZATION

The purpose for this work package is to develop personalization techniques and Scandinavian language processing capabilities to provide personalized content generation and:

- Develop truly explainable, fair and transparent personalization techniques
- Enable proactivity in customer relations
- Provide an individualized experience that provably respects privacy concerns
- Develop individualized content
- Develop large-scale Scandinavian language models
- Enable human-like content creation and conversations

Personalization and contextualization have been successfully employed in diverse applications over the past decade, and currently see an extended usage, for instance in proactive interaction with customers and individualization of news stories. LAP will contribute to developing such systems while ensuring that the system usage will be ethical and respecting users' requirements for privacy, fairness and accountability.

Building Scandinavian language models requires the compilation of large-scale reusable language resources, including general-purpose corpora from public sources (e.g., news and social media) as well as industry- and domain-specific text collections. We will address the scarcity of the latter by pre-training on the former and developing transfer learning methods. These large-scale language models will then be utilized in real-life scenarios by formulating a number of specific summarization, explanation, and conversational tasks based on our partners' use-cases. LAP will develop appropriate evaluation methodology with user-oriented evaluation measures and objectives. It will thus contribute to providing measurable quantification of the amount of domain-specific training material needed in order to provide a language service that is of sufficiently high quality.

**Work package leader: Krisztian Balog, Professor, University of Stavanger**



## HYB: HYBRID AI ANALYTICS

The purpose of work package HYB is to:

- Develop robust, stable and explainable data-driven models for physical systems
- Constrain models to enforce meaningful predictions
- Transfer data-driven models from simulations to reality
- Characterize and quantify uncertainty of data-driven models

This work package will develop methods to predict and reduce the uncertainty of data-driven models. The models will be constrained by existing knowledge, allowing to interpret the model (explainable AI) and reducing the amount of required training data. Applying these methods on real world applications will allow the industry partners to better predict the behavior of their facilities and improve their simulations, e.g. for condition monitoring, predictive maintenance, optimal utilization.

**Work package leader: Signe Riemer-Sørensen, Research Manager, SINTEF**



## WP DATA: DATA PLATFORMS AND STREAMING DATA

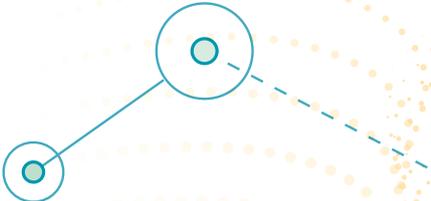
The purpose of this work package is twofold: 1) to develop modern AI for streaming and sensor-based data analysis and 2) to develop techniques and tools for the automatic creation and management of knowledge graphs.

The development of modern AI for streaming and sensor-based data will be done by

1. Providing anomaly detection and predictions with low quality streaming data
2. Providing uncertainty quantification and explainability with streaming data
3. Enabling combinations of streaming and static data for efficient data analysis

Streaming data can be used for automation, recommendations and decision making. Often this involves predictions and anomaly detection in multivariate time series, as well as providing explanations for the conclusions drawn. IoT sensors are increasingly instrumenting the physical world, and efforts have been made to use AI for solving these tasks also in low-quality data regimes. This research area will identify robust techniques for analysis of streaming data within several domains (including telco network, industrial IoT), with a particular focus on improving interpretability for cases with multivariate time series with low quality data.

Solving the research problems in DATA is crucial to successfully innovate how IoT data can be fully used in anomaly detection and contribute to breakthrough in applying AI in predictive maintenance and operational availability.



Real impact of data-driven AI depends on the availability of live data of sufficient quality and quantity in an automatically discoverable format that both humans and machines can understand. DATA will investigate how the semantics of data, through automatic creation and mapping of suitable knowledge graphs, can be leveraged to scale AI models from one situation to all similar situations and how complex graph-based structures can be efficiently stored and processed.

**Work package leader: Simone Casolo, Head of O&G Digital Services  
– Industrial Data Science, Cognite**



## WP INNOECO: AI INNOVATION ECOSYSTEMS

The purpose of this work package is to

- Create an Innovation Ecosystem among NorwAI partners for sharing both research findings, innovations, business solutions and change models within the AI domain.
- Share this knowledge with a broader audience, emphasizing small and medium sized businesses.

Artificial Intelligence still evokes more myths associated with threats than the possibilities technology brings to the table. AI experts are perceived speaking a foreign language when communicating to their organization, and most managers lack sufficient understanding of data driven business-models.

INNOECO will convey state of the art knowledge on AI-driven business models to a wide audience, including non-AI experts. This will be achieved by case analysis, conference presentations, and provision of a range of educational materials spanning from short tutorials to full university courses. The following topics will be covered: what AI is, its strategic impact on business model innovation, the business opportunities and limitations, implication on HR and teamwork, and innovation team dynamics, etc. To achieve these objectives, both AI and HR expertise will be convened.

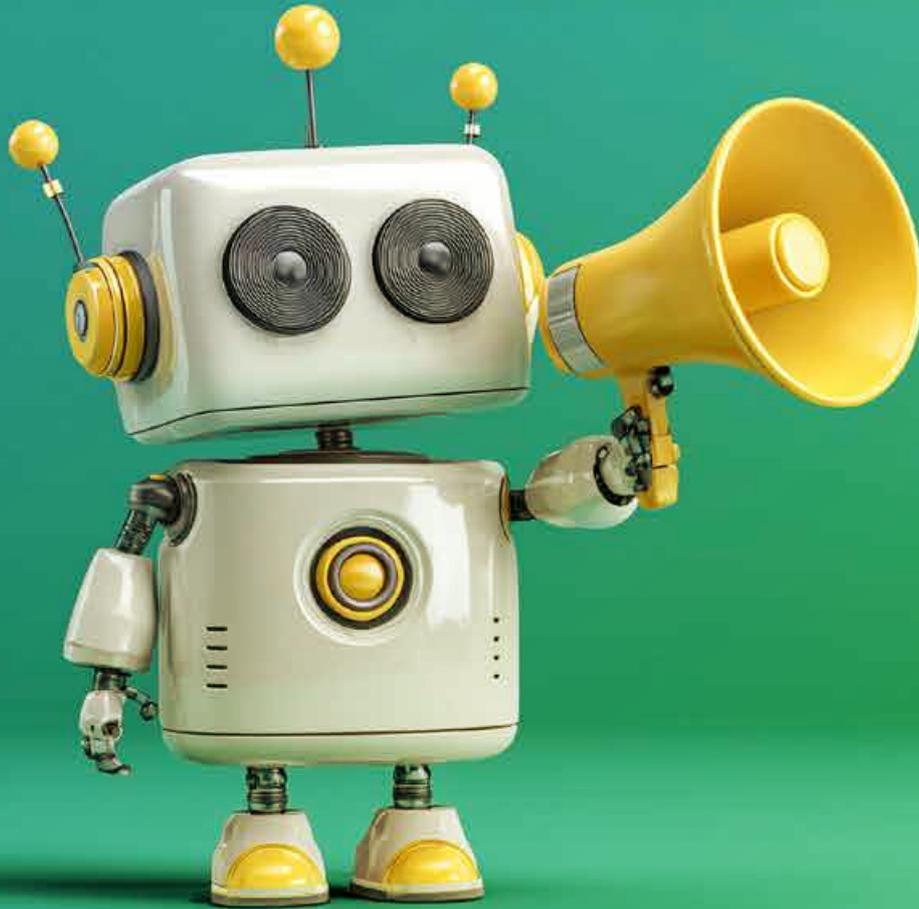
**Work package leader: Nhien Nguyen, Associate Professor, NTNU**





*AI master's students Jacob Opland (left) and Oskar Iglebæk Nesheim works on transforming radio programs into text-based editorial articles for NRK in their master's project.*

# Presentations and communication



NAME	TITLE	TYPE/PLACE	DATE
Ajraoui, Alae	Barriers and enablers of implementing AI in organizations: a systematic literature review	Presentation at EURAM Conference 2025	2025-06-25
Ajraoui, Alae	"Implementing AI in Practice:	Poster at NorwAI Innovate Conference 2025	2025-09-23
Arustashvili, M; Balog, K.	Natural Language User Interest Profiles for Scientific Literature Recommendation	Poster at NorwAI Innovate Conference 2025	2025-09-23
Bach, Kerstin	XAI for Trustworthy AI: Bridging Interpretability and Reliability.	Invited Seminar at the School of Electrical Engineering and Computer Science	2025-02-24
Bach, Kerstin	AI and Health: Challenges and Opportunities for Interdisciplinary Collaboration	AIR Research Seminar - School of Computing, Engineering, and Technology at RGU, Aberdeen	2025-03-26

NAME	TITLE	TYPE/PLACE	DATE
Bach, Kerstin	Forskningsdrevet innovasjon innen KI: Hvordan styrker vi det sammen?	Arendalsuka, Org: SFI Visual Intelligence, UiT, SFI NorwAI, NTNU	2025-08-12
Bach, Kerstin	NorwAI Center Presentation	Delegation from Department of Engineering Sciences, Universidad de Los Lagos (Chile) (Joel Sebastian Torres Carrasco )	2025-08-20
Bach, Kerstin	Teaching AI to Recognize Human Activity: Self-Supervised Learning with Minimal Labeled Data	2025 Future Wearable Tech Summit, Nice	2025-09-02
Bach, Kerstin	NorwAI Center Presentation	DNB delegation at NTNU	2025-09-05
Bach, Kerstin	NorwAI: Current Results & Future Plans	IDI 40 år minisymposium, arrangør Arne Jørgen Berre	2025-10-11
Bach, Kerstin	AI for Health: What We've Learned and Where We're Headed	NOFE Conference 2025, Bergen (Norwegian Epidemiological Association)	2025-10-30
Bach, Kerstin	Panel: Governance for Innovation	Global AI Policy Research Summit 2025 Delft	2025-11-12
Bach, Kerstin	Gesundheitsdaten als Innovationsmotor: Erfassung, Analyse und KI-gestützte Nutzung am Beispiel von selfBACK und HUNT4	Bremen Symposium AI in Health	2025-11-27
Bach, Kerstin; Flobak, Åsmund; Tessem, May-Britt; Visnes, Torkild.	Panel: Challenges and Opportunities for Precision Medicine.	Workshop: Precision Medicine for Sustainable Healthcare	2025-04-28
Bauer, Susanne	Data and Methods Inside Out: Epidemiology as Infrastructure	Invited Talk, STS Seminars, Chalmers University, Gothenburg	2025-02-28
Blekeli, T. N.; Hagerup, E. A.	From idea to internal success: How we built NRKGPT and trained 1500 employees in AI	NorwAI Innovate 2025	2025-09-24
Brasethvik, Terje	KI på Veg - Bruk og Begrensninger	Avdelingssamling Trøndelag Fylkeskommune	2025-03-26
Brasethvik, Terje	Norske språkmodeller fra NorwAI	Møte med Finans Norge	2025-03-27
Brasethvik, Terje	NorwAI: Samspillet mellom student, veileder og industri	Master Safari	2025-04-10
Brasethvik, Terje	Norske språkmodeller fra NorwAI - Anvendelser i helse	Buanseminaret	2025-10-15
Bungum, Lars	Online Evaluation of LLM-Generated Text: What Metrics Make a Difference?	NorwAI Innovate 2025	2025-09-23
Bungum, Lars	Multi-Dimensional Evaluation Metric for Text Generation	Poster at NorwAI Innovate Conference 2025	

NAME	TITLE	TYPE/PLACE	DATE
Bungum, Lars	Moderator and session chair	Workshop: Samisk språk og teknologi – samarbeid for fremtiden (Kommunal- og distriksdepartementet )	2025-10-30
Casolo, Simone	AI for the heavy industry	Outlook North 2025 Harstad	2025-10-30
Casolo, Simone	From Anomaly to Answer: AI Agents for Industrial Equipment Root Cause Analysis	NorwAI Innovate 2025	2025-09-24
de Leyos, Aaron; Stasik, Alexander Johannes; Riemer-Sørensen, Signe	Gaussian Process Kolmogorov-Arnold Networks for Probabilistic Multivariate Function Approximation	NORA Annual Conference 2025	2025-06-02
de Leyos, Aaron; Stasik, Alexander Johannes; Riemer-Sørensen, Signe	Gaussian Process Kolmogorov-Arnold Networks for Probabilistic Multivariate Function Approximation	Poster at NorwAI Innovate Conference 2025	2025-09-23
Djupskås, Aslak; Stasik, Alexander Johannes; Riemer-Sørensen, Signe.	Unreliable Uncertainty Estimates with Monte Carlo Dropout	NORA Annual Conference 2025	2025-06-02
Djupskås, Aslak; Stasik, Alexander Johannes; Riemer-Sørensen, Signe.	Unreliable Uncertainty Estimates with Monte Carlo Dropout	Poster at NorwAI Innovate Conference 2025	2025-09-23
Doan, Tu My; Baumgartner, David; Kille, Benjamin	Språkmodeller for tekstsammendrag	AI 4 Interviews Project seminar	2025-05-20
Eidnes, Sølve	Structure-preserving machine learning for physical systems.	NorwAI/NAIC Tutorial	2025-03-06
Fjuk, Annita	Digital imagination- a key competence for all in 2025	NorwAI Innovate 2025	2025-09-24
Gulla, Jon Atle	Introduksjon til kunstig intelligens	Surnadal næringsforening	2025-09
Gulla, Jon Atle	Nasjonal KI-veileder for ansvarlig innføring og bruk av KI-assistenten	Personverndagene	2025-09-09
Gulla, Jon Atle	Paneldiskusjon - Artificial Intelligence and Value Creation for the Industry: Beyond the Hype	BRU21 Conference 2025	2025-05-27
Gulla, Jon Atle	Du store Ai-verden! Hva med lille Norge?	Framtidsfredag 2025	2025-06-20
Gulla, Jon Atle	Kan vi bruke språkmodeller for å forstå havets språk?	Arendalsuka, Et hav av data – ny teknologi gir oss nøkkelen til kunnskap, sikkerhet og verdiskaping. Org: NTNU, Havforskningsinstituttet	2025-08-13
Gulla, Jon Atle	Large Language Models in Norway	European Media Lab Day, Trondheim	2025-10-16

NAME	TITLE	TYPE/PLACE	DATE
Gulla, Jon Atle	Language Models, Data and Applications	Digitalization in the Automation Domain 2025, NFEA - Norsk Forening for Elektro og Automatisering	2025-10-22
Gulla, Jon Atle	The Importance of Language Data for the Development of LT Solutions	Language Data Space Workshop, Oslo	2025-10-28
Gulla, Jon Atle	Digitalisering og kunstig intelligens: Hva og hvorfor?	Wood works, Trondheim	2025-11-06
Gulla, Jon Atle	Store språkmodeller og media	KI og mediebildet, Åpent møte ved Norges Tekniske Vitenskapsakademi	2025-12-04
Halvorsen, Christian Printzell	AI - from promise to practice	NorwAI Innovate 2025	2025-09-24
Ingvaldsen, Jon Espen	Når KI kommer inn på legekontoret og klinikken	Nasjonalt senter for e-helseforskning, webinar	2025-02-14
Ingvaldsen, Jon Espen	Kommersialisering av generative AI for Helsetjenesten	Trondheim Techport Innovasjonsfrukost: KI og helseteknologi - hvordan navigere i den regulatoriske jungelen?	2025-02-28
Ingvaldsen, Jon Espen	AI-agenter som avlaster helsepersonell og øker kvaliteten på tjenestene	Faglig Forum for helse- og sosialtjenesten	2025-03-25
Ingvaldsen, Jon Espen	Når Kunstig intelligens blir en del av helseteamet ditt	Framtidsfredag 2025	2025-06-20
Jekic, A ; Natsaridou, A.; Riemer-Sørensen, S.; Langseth, H ; Gundersen, O. E.	Examining the robustness of Physics-Informed Neural Networks to noise for Inverse Problems	Poster at NorwAI Innovate Conference 2025	2025-09-23
Kille, Benjamin	NorwAI—NTNU og Helse Midt-Norge	Helse Midt-Norge visits NTNU	2025-01-13
Kille, Benjamin	NorwAI—BRU21: An AI outlook into the future	BRU21 steering committee meeting	2025-01-24
Kille, Benjamin	Hvordan kan algoritmer påvirke våre meninger?	Besøk av Kommunal- og distriktsminister Kjersti Stenseng	2025-03-10
Kille, Benjamin	Artificial Intelligence and how it can improve Asset Management	Train Academy Day	2025-06-03
Liu, Peng; Zhang, Lemei	Exploring the Frontier of Norwegian Large Language Models	Workshop med NTNU Universitetsbiblioteket, Seksjon for ressurser of digitale tjenester	2025-12-03
Lopez Calvet, J. C.	Concrete applications of AI in Media	NorwAI Innovate 2025	2025-09-23

NAME	TITLE	TYPE/PLACE	DATE
Løland, Anders	Kunstig intelligens fra øst og vest	Hønefoss øst Rotaryklubb, Møte nr 2702	2025-10-07
Monteiro, Eric	Digital oil	ISFIT - Democracy Yay or nay? (Eric) (Paneldebatt?)	2025-03-21
Monteiro, E. (NTNU); Thorvik, K. (Hemit)	Domain spesific LLMs in health care	NorwAI Innovate	2025-09-23
Morais-Storz, Marta; Nguyen, Nhien	Insights on the conditions of human-AI collaboration: an empirical exploration	EURAM Conference 2025	2025-06-25
Murgorgo, S.; Nguyen, N., Mikalef, P; Bach, T. A.	Responsible AI practices in organisations: A systematic literature review	Poster at NorwAI Innovate Conference 2025	2025-09-23
Murgorgo, Serinha	Responsible AI practices in organisations: A systematic literature review	Presentation at EGOS Conference, Athens, Greece	2025-07-03
Nguyen, Nhien	How Norwegian businesses govern Responsible AI – An Introduction	Telenor Global AI Day	2025-09-10
Nordby, J. (Soundsensing); Kolstø, J. V. (Norwegian computing center, NR)	Applied anomaly and change point detection algorithms for online monitoring	NorwAI Innovate 2025	2025-09-24
Riemer-Sørensen, S. (SINTEF)	Synergies between NorwAI and the Norwegian Center on AI for Decisions (AID)	NorwAI Innovate 2025	2025-09-23
Tüzen, A.; Nørvåg, K; Langseth, H.	Simplifying the Graph Similarity Learning Task with Graph Neural Network	Poster at NorwAI Innovate Conference 2025	2025-09-23
Riemer-Sørensen, Signe	Kraftbransjens største utfordringer og muligheter med kunstig intelligens	Kraftbransjens KI-konferanse 2025	2025-05-13
Riemer-Sørensen, Signe	WindAI – Collaborative problem solving and skill development	Opening of OsloMet Green Energy Lab	2025-05-15
Riemer-Sørensen, Signe	Hybrid AI in practice	VITO lunch talk	2025-06-22
Riemer-Sørensen, Signe	AI needs math to work in real life	Math Meets Industry 2025	2025-08-14
Riemer-Sørensen, Signe (SINTEF)	Recent work on Hybrid AI Analytics	NorwAI Forum, Trondheim	2025-04-09
Riemer-Sørensen, Signe (SINTEF); Stasik, Alexander (SINTEF); Lund, Halvor	Hybrid AI for power transformers	#techday - kraftbransjens IT-konferanse	2025-04-29
Shamba, Kazeem	Contrast All The Time: Learning Time Series Representation from Temporal Consistency	NorwAI Innovate 2025	2025-09-23

NAME	TITLE	TYPE/PLACE	DATE
Steppe, Jessica Annalena; Storz, Marta Morais; Nguyen, Nhien	Insights into creative idea generation and development using language models	Presentation at XXXVI ISPIM Innovation Conference, Bergen, Norway	2025-06-15
Steppe, Jessica Annalena; Storz, Marta Morais; Nguyen, Nhien	AI-augmented creativity: insights into idea generation and development using language models	Presentation at Druid25 Conference, Toronto, Canada	2025-06-24
Steppe, Jessica Annalena; Storz, Marta Morais; Nguyen, Nhien	AI-augmented creativity: insights into idea generation and development using language models	Presentation at EGOS Conference, Athens, Greece	2025-07-03
Steppe, Jessica Annalena; Storz, Marta Morais; Nguyen, Nhien	AI-augmented creativity: insights into idea generation and development using language models	Presentation at AOM Conference, Copenhagen, Denmark	2025-07-25
Steppe, Jessica Annalena; Nguyen, Nhien; Johansson, Ole Hartvik Nissen; Maríc, Sinan	Key Competencies for Human-AI Collaboration and How Organizations Foster Their Development	Poster at NorwAI Innovate Conference 2025	2025-09-23
Steppe, Jessica Annalena; Storz, Marta Morais; Nguyen, Nhien	AI-augmented creativity: insights into idea generation and development using language models	SMS Conference, San Francisco, USA	2025-10-11
Strømke, Inga; Løland, Anders	Bør vi frykte generasjon arbeidsløs?	Dagens Næringsliv	2025-06-29
Thaulow, Sven Størmer	Når kunstig intelligens blir en del av helseteamet: Fra forsøk til gevinst.	Arendalsuka, Org: NTNU NorwAI, Hemit HF, EHiN	2025-08-12
Thaulow, Sven Størmer	Forskningsdrevet innovasjon innen KI: Hvordan styrker vi det sammen?	Arendalsuka, Org: SFI Visual Intelligence, UiT, SFI NorwAI, NTNU	2025-08-12
Ugland, Y. (DNB)	Leading with Data: How Ethical Foundations Drive AI Success	NorwAI Innovate	2025-09-23
Wold, K. (Sparebank1 SMN)	AI Competence in Practice: How a Shared Learning Project is shaping the future of Customer Service	NorwAI Innovate	2025-09-24
Zhang, Lemei	From Cortex to Code: Comparative Insights into Reasoning in Brain Networks and LLMs	NorwAI Innovate	2025-09-23
Øvreliid, Lilja (UiO)	LLM evaluation in the Mimir project: how can we assess the value of copyrighted materials in LLM training?	NorwAI Innovate	2025-09-23
Özgöbek, Özlem	Poster	WACE 2025, Istanbul	2025-05-28
Özgöbek, Özlem	Round table session	WACE 2025, Istanbul	2025-05-28

# NorwAI in the media in 2025

**Norsk AI-direktør om Deepseek:  
- Gode nyheter for Norge**  
(E24, 2025-02-06)

**Tror på milliardgevinster: Ekspertgruppe  
skal få opp farten i næringslivet**  
(digi.no, 2025-03-17)

**Ny ekspertgruppe skal gi hjelp  
til norske bedrifter**  
(digi.no, 2025-03-17)

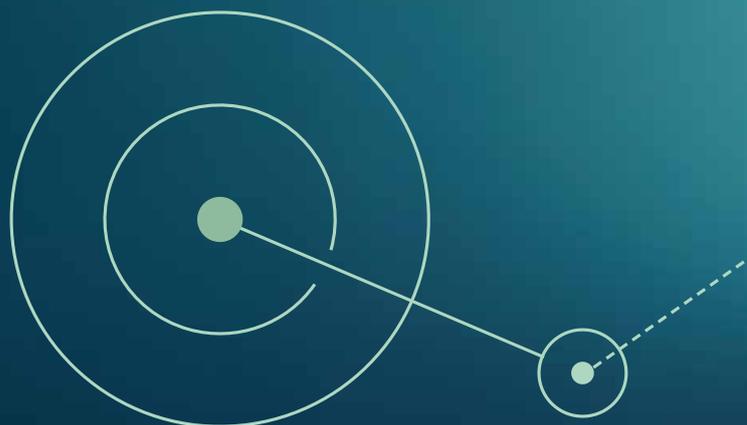
**Regjeringen lanserer veileder  
for kunstig intelligens!**  
(INNOMAG, 2025-03-21)

**Nå er «oppskriften» for bruk av  
KI-assistenter klar: Slik kan norske  
virksomheter komme i gang**  
(Pressemelding Digitaliserings- og  
forvaltningsdepartementet, 2025-06-16)

**Regjeringen har lagt ut sin KI-veileder**  
(Travel News Norge, 2025-06-16)

**«Oppskriften» for bruk av KI-assistenter  
er klar**  
(Byggfakta Nyheter, 2025-06-16)

**Regjeringens KI-oppskriften er klar!**  
(Norsk Byggebransje, 2025-06-16)



**-Tidsfristen har vært drepende**  
(digi.no, 2025-06-17)

**Lanserte KI-veileder etter kun to måneders  
arbeid. – Det har vært veldig hektisk**  
(Alltinget digital, 2025-06-19)

**Bør vi frykte generasjon arbeidsløs?**  
(Dagens Næringsliv, 2025-06-29)

**Digital suverenitet må være kjernen i  
Norges KI-strategi**  
(digi.no, 2025-08-02)



**Kunstig intelligens – en ny hverdag?**

(Trollheimsporten, 2025-09-06)

**Nå kan du få diagnose før du har gått ut fra legekantoret**

(Adresseavisen, 2025-09-12)

**AI4Interviews har vunnet Europols innovasjonspris**

(Politiet, Nyheter fra Oslo politidistrikt, 2025-09-24)

**3000 leger og helsepersonell kastet seg over**

**AI-verktøy: – Folk er på gråten**

(Shifter, 2025-09-24)

**Når vinden blåser, flyter strømmen**

**- møt vinnerne av WindAI-konkurransen**

(Pressemelding NTB Kommunikasjon, 2025-11-03)



## Internal seminars

NAME	TITLE	TYPE/PLACE	DATE
Bauer, Susanne (UiO)	Recent work on AI in Society	NorwAI Forum, Trondheim	2025-04-09
Hoiseth, Eirik Hoel (Telenor)	Telenor's viewpoint from the innovation angle	NorwAI Forum, Trondheim	2025-04-09
Kille, Benjamin (NTNU)	Recent work on Language & Personalization	NorwAI Forum, Trondheim	2025-04-09
Malacarne, Sara (Telenor)	Recent work on Data Platforms and Streaming Data	NorwAI Forum, Trondheim	2025-04-09
Monteiro, Eric (NTNU)	Extending research in SOC	NorwAI Forum, Trondheim	2025-04-09
Nguyen, Nhien, Qiu Xinlu; Arjaoui, Alae; Steppe, Jessica; Murgogo, Serinha; Li, Xiaomu (NTNU)	Presentation, research in INNOECO (Human-AI collaboration, Responsible AI in practice, AI implementation in organizations, How genAI transforms organization)	NorwAI Forum, Trondheim	2025-04-09
Ngyuen, Nhien (NTNU)	Recent work on AI Innovation Ecosystems	NorwAI Forum, Trondheim	2025-04-09
Rønningstad, Egil (UiO)	Presentation, research in LAP	NorwAI Forum, Trondheim	2025-04-09
Shamba, Kazeem (NTNU)	Representation learning in Sensor DATA	NorwAI Forum, Trondheim	2025-04-09
Sund, Espen (NTNU)	Presentation, research in TRUST	NorwAI Forum, Trondheim	2025-04-09
Thrane, Jon Eivind (DNV)	Recent work on Trustworthy AI	NorwAI Forum, Trondheim	2025-04-09
Thronsen, Eivind (Vend)	AI innovation at Schibsted Marketplaces	NorwAI Forum, Trondheim	2025-04-09
Qui, Xinlu (NTNU)	Boundary spanning, boundary spanner, and knowledge transfer	NorwAI Forum, Trondheim	2025-04-09
Wold, Kjersti (SMN)	SMN's lesson learned in AI use case	NorwAI Forum, Trondheim	2025-04-09
Bayrak, Betül (NTNU)	Evaluating Counterfactual Explanations - Metrics, Users and New Frontiers	NorwAI Forum, Oslo	2025-11-19



NAME	TITLE	TYPE/PLACE	DATE
Hafver, Andreas (DNV)	Explaining uncertainty to enable trust in AI	NorwAI Forum, Oslo	2025-11-19
Larsen, Eirin (Telenor)	Creating an AI Innovation Culture - Building trust and organizational readiness for responsible AI	NorwAI Forum, Oslo	2025-11-19
Nikoay, Nikolov (SINTEF)	[Automating Compliance Processes in AI Pipelines – Integrating compliance-by-design into data and AI operations]	NorwAI Forum, Oslo	2025-11-19
Sund, Espen (NTNU)	Parameter estimation under different models of trust	NorwAI Forum, Oslo	2025-11-19
Özğöbek, Özlem (NTNU)	Misinformation spread and recommender systems on social media	NorwAI Forum, Oslo	2025-11-19
Cho, Kyunghyun (New York Univ.)	Learning to X	The AI Seminar Series from NorwAI and NAIL	2025-01-14
Martens, Harald (NTNU)	A greener, safer and more understandable AI for natural science and technology	The AI Seminar Series from NorwAI and NAIL	2025-02-27
Borth, Damian; Jenssen, Robert (UiT)	Weight Space Learning: How to learn Representations of Neural Networks	The AI Seminar Series from NorwAI and NAIL	2025-03-13
Nehaniv, Chrystopher L. (Univ. of Waterloo)	Evolution of Complexity in Computational Intelligence	The AI Seminar Series from NorwAI and NAIL	2025-03-21
Fitzgerald, Tesca (Yale Univ.)	Getting Robots to Ask Better Questions	The AI Seminar Series from NorwAI and NAIL	2025-04-24
Zhu, Zhanxing (Univ. of Southampton)	Unveiling AI Alchemy: Understanding Training Dynamics in Modern Deep Learning	The AI Seminar Series from NorwAI and NAIL	2025-04-30

NAME	TITLE	TYPE/PLACE	DATE
Sharma, Kshitij (NTNU)	AI and biometrics in education	The AI Seminar Series from NorwAI and NAIL	2025-05-09
Kabisch, Johannes; Eiermann, Aron (TU Darmstadt)	AI Meets Proteins: Unlocking Biotechnology through Structure Prediction	The AI Seminar Series from NorwAI and NAIL	2025-05-23
Klöpper, Miriam (NTNU)	From Metrics to Mistakes – Artificial Intelligence in Peronnel Management	The AI Seminar Series from NorwAI and NAIL	2025-06-13
Taylor, Gavin (US Naval Academy)	Whose Truth? The impact of recommendation systems on societies and national defense	The AI Seminar Series from NorwAI and NAIL	2025-06-20
Cabañas, Rafael (Univ. of Almería)	Bounded Counterfactual Reasoning with Probabilistic Graphical Models	The AI Seminar Series from NorwAI and NAIL	2025-08-08
Sporsem, Tor (NTNU/SINTEF)	How we built an AI for Pediatricians – and What Every Organization Should Know	The AI Seminar Series from NorwAI and NAIL	2025-09-19
Chondrogiannis, Theodoros (NTNU)	Graph Databases: Past, Present and Future	The AI Seminar Series from NorwAI and NAIL	2025-10-24
Goethals, Bart (Univ. of Antwerp)	Challenges of News Recommendations in Practice	The AI Seminar Series from NorwAI and NAIL	2025-11-14
Kurpicz-Briki, Mascha (Bern Univ. of Applied Sciences)	Inside the BIAS project – Societal Stereotypes in Language Models	The AI Seminar Series from NorwAI and NAIL	2025-11-21
Puttick, Alexandre (Bern Univ. of Applied Sciences)	The Problem with Fairness by Blindness: Grounding Technical Methods in Sociotechnical Context	The AI Seminar Series from NorwAI and NAIL	2025-11-21

# Personnel

## KEY RESEARCHERS

NAME	INSTITUTION	MAIN RESEARCH AREA
Simone Casolo	Cognite	Machine learning, Hybrid analytics
Andreas Hafver	DNV	Trustworthy AI
Anders Løland	Norsk Regnesentral	Machine Learning, Statistics
Ildiko Pilán	Norsk Regnesentral	NLP, Data privacy, Clinical NLP, Computer-assisted language learning, Machine learning
Pierre Lison	Norsk Regnesentral	NLP, Privacy and Security, Machine learning, Information retrieval, Semantics, AI and society, Big data
Benjamin Kille	NTNU	NLP, Personalization, Machine learning
Eric Monteiro	NTNU	AI in Society, digitalization of organizations
Jon Atle Gulla	NTNU	NLP, Personalization
Jon Espen Ingvaldsen	NTNU	Innovation processes
Kerstin Bach	NTNU	AI and society, Innovation processes, Personalization, Machine learning, Trustworthy AI, Decision Support
Lemei Zhang	NTNU	Personalization, Big data, Information retrieval, Machine learning
Nhien Ngyuen	NTNU	Innovation processes
Peng Liu	NTNU	Personalization, Big data, Information retrieval, Machine learning
Terje Brasethvik	NTNU	Information retrieval, Semantics, Trustworthy AI
Theodoros Chondrogiannis	NTNU	Database systems, Graph Analytics
Özlem Özgöbek	NTNU	Personalization, Recommender systems, Trustworthy AI and Disinformation detection
Francisco Martin-Recuerda	SINTEF Digital	Large language models for entity matching for ontologies in industrial systems and digital twins.
Alexander Johannes Stasik	SINTEF Digital	Machine learning, Hybrid analytics
Lisa Græsleie	SINTEF Digital	AI Innovation



NAME	INSTITUTION	MAIN RESEARCH AREA
Signe Reimer-Sørensen	SINTEF Digital	Machine learning, Hybrid analytics
Till Christopher Lech	SINTEF Digital	AI Innovation
Stian Arntsen	Sparebank1 SMN	AI Innovation
Halvor Lund	Statnett	Big Data
Sarah Malacarne	Telenor	Big data, Machine learning, Trustworthy AI, Time series
Erik Velldal	UiO	NLP, Machine learning
Lilja Øvrelid	UiO	NLP, Machine learning
Susanne Bauer	UiO	AI in society
Krisztian Balog	UiS	NLP, Semantics, Personalization, Information retrieval, Machine learning

## VISITING RESEARCHERS

NAME	AFFILIATION	NATIONALITY	SEX	DURATION	TOPIC
Gavin Taylor	US Naval Academy	USA	M	2024-06-28 – 2025-06-27	Machine learning, including reinforcement learning, time series, and evasion of machine learning systems.

## POSTDOCTORAL RESEARCHERS

Postdoctoral researchers with financial support from the center budget:

NAME	NATIONALITY	PERIOD	SEX	TOPIC
Bayrak-Çik, Betül	Türkiye	2025-02-20-2028-02-19	F	Explainable Artificial Intelligence

Postdoctoral researchers working on projects in the center with financial support from other sources:

NAME	FUNDING	NATIONALITY	PERIOD	SEX	TOPIC
Zhang, Lemei	NTNU	Kina	2022-04-04-2023-01-31; 2023-07-31-2025-10-03	F	Streaming Data Analytics
Lars Bungum	Horizon Europe	Norge	2024-02-01 – 2026-10-31	M	Evaluating Trustworthy and Efficient Large Language Models

## PHD STUDENTS

PhD students with financial support from the center budget:

NAME	NATIONALITY	PERIOD	SEX	TOPIC
Arustashvili, Mariam	Georgia	2025-04-01 – 2028-03-31	F	Large Language Models for Recommendation
Baumgartner, David	Austria	2021-09-21 – 2025-11-30	M	Data analysis with noisy and low-quality data streams
Bernard, Nolwenn	France	2022-02-01 – 2025-01-31	F	Study of fairness and transparency in conversational recommender systems (PhD defence: 2025-06-16)
Eder, Christoph	Austria	2025-09-01 – 2028-11-30	M	Enhancing data quality and integration through data augmentation and generative AI
Kamsvåg, Pål Furu	Norway	2025-09-01 – 2028-11-30	M	Evaluating Artificial Intelligence – Bridging the Gap Between Benchmark Success and Real-World Utility
Knaus, Tanja	Austria	2022-04-01 – 2025-03-31	F	Automation of affective data: infrastructure and data practices in voice recognition systems
Łajewska, Weronika	Poland	2022-02-01 – 2025-06-20	F	Personalizing Conversational Informational Access (PhD defence: 2025-06-20)
Murgogo, Serinha	Zimbabwe	2023-08-15- 2026-08-14	F	Embracing Responsible AI: Integration of responsible AI principles in organisational practices
Michałowska, Katarzyna	Poland	2021-01-01 – 2024-12-31	F	Informed machine learning (PhD defence: 2025-05-23)
Nikolov, Nikolay	Bulgaria	2021-01-01 – 2024-12-31	M	Flexible Deployment of Big Data Pipelines on the Cloud/Edge/Fog Continuum (PhD defence: 2025-04-30)
Rønningstad, Egil	Norway	2021-10-14 – 2026-02-13	M	Entity-level Sentiment Analysis
Shamba, Abdul-Kazeem	Nigeria	2023-09-01- 2027-08-31	M	Representation Learning for Predictive Analytics and Reasoning in Multimodal Sensor Data
Steppe, Jessica Annalena	Germany	2023-08-15- 2026-08-14	F	AI-Augmented Creativity in Innovation Teams
Sund, Espen	Norway	2025-01-06 – 2028-01-05	M	Parameter estimation under different models of trust
Tüzen, Ahmet	Türkiye	2024-12-17 – 2027-12-16	M	Scalable multi-task dynamic graph learning framework
Vassøy, Bjørnar	Norway	2021-08-02 – 2025-08-01	M	Fairness, Accountability, Transparency and Privacy in Personalization/Recommender systems
Yadav, Vandana	India	2023-11-02 – 2027-11-01	F	Natural Language Processing for personalized content summarization

## PHD STUDENTS

PhD students working on projects in the center with financial support from other sources:

NAME	FUNDING	NATIONALITY	PERIOD	SEX	TOPIC
Ajraoui, Alae	NTNU	Marocco	2023-10-23 – 2026-10-22	M	Digital innovations implementation: A dynamic approach of continuous adaptation
Bayrak, Betül	NFR/NTNU (Exaigon)	Turkiye	2022-01-19 – 2025-01-18	F	Explainable Case-Based Reasoning (PhD defence: 2025-02-19)
Bekkemoen, Yanzhe	NTNU	Norway	2019-10-01 – 2024-09-30	M	Probabilistic approaches to explainable AI and reinforcement learning. (PhD defence: 2025-05-06)
Bjørn, Anna Rodum	NTNU	Norway	2021-11-11 – 2025-11-10	F	Explainable deep bayesian learning
Jernbert, Folke	NTNU (integret PhD)	Norway	2025-10-01 – 2030-09-30	M	Language technology
Ottersen, Stuart	NTNU Helse: AI Ment	Norway	2023-02-01 – 2027-01-31	M	Using machine learning to predict response and dropout in RCT data from digital sleep therapy
Skott, Aina Øverås	RCN/ Sparebank1 SMN	Norway	2025-08-15 – 2028-08-14	F	Business Value from Artificial Intelligence (AI): Creating and capturing value from AI in Banking and Finance
Sylvester, Sophia	NTNU Helse: AI Ment	Germany	2023-02-01 – 2027-01-31	F	Radar Data and Machine Learning in Psychiatric Care: Novel Approaches for Advancing Sleep/Wake Estimation and Aggression Prediction

## MASTER'S DEGREES (2025)

Master students obtaining their degree on NorwAI topics in 2025:

NAME	SEX	PERIOD	TOPIC
Aasvær, August Sætre	M	2025-08-24 – 2025-06-30	Decoding Emotional Responses to Language
Adam, Thomas Hasvold	M	2025-08-24 – 2025-06-30	Experimental Analysis and Mitigation of Performance Degradation in HNSW Indexes through Enhanced Update Strategies
Amtrup, Casper Lunde	M	2023-08-01- 2024-06-30	SOLDIR: Online Semantic Joins
Andresen, Fredrik	M	2025-08-24 – 2025-06-30	Scalable Temporal Reasoning over Document Collections
Eidal, Simen	M	2025-08-24 – 2025-06-30	Beyond Feature Engineering: An Investigation into a Wind Power Prediction Neural Network Institute
Fagernes, Sander Østrem	M	2025-08-24 – 2025-06-30	Applying AI techniques for synchronous analysis of students' open-text responses in higher education classrooms
Fonn Hansen, Fredrik	M	2025-08-24 – 2025-06-30	Applying AI techniques for synchronous analysis of students' open-text responses in higher education classrooms
Gudmundsen, Selma	F	2025-08-24 – 2025-06-30	Local Explainability to Combat Customer Churn in Sparebank 1 SMN
Halvorsen, Oscar Aleksander	M	2025-08-24 – 2025-06-30	Neural Machine Translation of Bokmål-Nynorsk Using Large Language Models
Håland, Hans Jakob	M	2025-08-24 – 2025-06-30	Improving Domain Adaptation for Forecasting Models with Feature-Based Multivariate Time Series Generation
Jacobsen, Anna Johanne	F	2025-08-24 – 2025-06-30	Modeling the Impact of Recommendation Diversity on Misinformation Spread in Social Networks: An Agent- Based Approach
Jakobsen, Lise	F	2025-08-24 – 2025-06-30	Modeling the Impact of Recommendation Diversity on Misinformation Spread in Social Networks: An Agent- Based Approach
Johannesen, Jørgen	M	2025-08-24 – 2025-06-30	Bootstrapping Norwegian Financial Sentiment Analysis: Leveraging Machine Translation in a Low-Resource Setting
Jørgensen, Oskar	M	2025-08-24 – 2025-06-30	Improving Generative Image Models Using Hybrid Flow Matching and Autoregressive Models
Krohn-Pettersen, Thale	F	2025-08-24 – 2025-06-30	When the Forecast Fails Designing Human-Centered Explanations for Wind Power Trading
Kvadsheim, Hans Dahle	M	2025-08-24 – 2025-06-30	Decoding Emotional Responses to Language
Malcolmson, Emily Louise	F	2025-08-24 – 2025-06-30	Reconstruction Diffusion-based Anomaly Detection: The Role of The Wasserstein Distance
Mihai, Silviu Catalin	M	2025-08-24 – 2025-06-30	Efficient Approximate Nearest Neighbor Search for Spatio-Textual Data
Moe, Synne Frafjord	F	2025-08-24 – 2025-06-30	Abstractive Summarization Using Transformer Models for Norwegian Discharge Summaries
Nordal, Marcus Stephan	M	2025-08-24 – 2025-06-30	Efficient Approximate Nearest Neighbor Search for Spatio-Textual Data

## MASTER'S DEGREES (2025)

Master students obtaining their degree on NorwAI topics in 2025:

NAME	SEX	PERIOD	TOPIC
Reistad, Skage Klingstedt	M	2025-08-24 – 2025-06-30	TallOrder - Numerical Distribution Index
Rødseth, Magnus	M	2025-08-24 – 2025-06-30	Bootstrapping Norwegian Financial Sentiment Analysis: Leveraging Machine Translation in a Low-Resource Setting
Salvesen, Thea Sofie	F	2025-08-24 – 2025-06-30	Local Explainability to Combat Customer Churn in Sparebank 1 SMN
Sjåvik, Vegard	M	2025-08-24 – 2025-06-30	Improving Domain Adaptation for Forecasting Models with Feature-Based Multivariate Time Series Generation
Steinsund, Vigdis-Irene	F	2025-08-24 – 2025-06-30	Experimental Analysis and Mitigation of Performance Degradation in HNSW Indexes through Enhanced Update Strategies
Syrstad, Anine	F	2025-08-24 – 2025-06-30	Probabilistic Wind Power Forecasting: A Structured Literature Review
Tran, Julie Vy	F	2025-08-24 – 2025-06-30	A Story of My Patient's Life: Extracting Patient Timelines from Unstructured Norwegian Patient Journals using Text Mining
Vesetrud, Markus Risa	M	2025-08-24 – 2025-06-30	Beyond Feature Engineering: An Investigation into a Wind Power Prediction Neural Network
Ytterstad, Sofie Karlsen	F	2025-08-24 – 2025-06-30	Probabilistic Wind Power Forecasting: A Structured Literature Review

## PHD AND MASTER'S STUDENT RESEARCH ASSISTANTS

NAME	SEX	PERIOD	TOPIC
Even W. Lauvrak	M	2025-01-01 – 2025-12-31	Scenario description, Development, finetuning language models for the construction industry
Vilhjalmur Arnar Vilhjalmsson	M	2025-04-01 – 2025-06-15	Teaching assistant at EVU course Innovation with large language models
Xing, Yujie	F	2025-02-16 – 2025-03-31	Teaching assistant at EVU course Innovation with large language models

## BACHELOR'S STUDENT ASSISTANTS

NAME	SEX	PERIOD	TOPIC
Andrew Wang (University of Waterloo)	M	2025-01-01-2025-04-30	Large-scale data handling and evaluations for underwater acoustic signal classification
Abeer Kashar (University of Waterloo)	F	2025-01-01-2025-04-30	Research Assistant in Computational Linguistics and NLP
Jeevan Parmar (University of Waterloo)	M	2025-01-01-2025-04-30	Social Robotics Programmer - Artificial Intelligence at NTNU
Ayah Amer (University of Waterloo)	F	2025-01-01-2025-04-30	Visualizing Underwater Acoustic Data for Deep Learning
Jeffrey Qiu (University of Waterloo)	M	2025-01-01-2025-04-30	Programmer/Software Developer in Artificial Intelligence Research
Cyrus Hatami (University of Waterloo)	M	2025-01-01-2025-04-30	Programmer/Software Developer in Artificial Intelligence Research
Heath Couture (University of Waterloo)	M	2025-09-01-2025-12-31	Visualizing Underwater Acoustic Data for Deep Learning
Charles Huang (University of Waterloo)	M	2025-09-01-2025-12-31	Research Assistant in Computational Linguistics and NLP
John Greenough (University of Waterloo)	M	2025-09-01-2025-12-31	Social Robotics Programmer - Artificial Intelligence

# Accounts



2025	FUNDING	COSTS
The Research Council	11849	
The Norwegian University of Science and Technology (NTNU)	8870	15560
Research Partners*)	1189	6348
Enterprise partners**)	8748	8748
Total	30 656	30656

All figures in 1000 NOK.

\*) Norwegian Computing Center (NR), SINTEF, University of Oslo and University of Stavanger.

\*\*\*) Aneo, Cognite, Digital Norway, DNB, DNV, Kongsberg Digital, NRK, Vend, Sparebank1 SMN, Statnett, Telenor, Schibsted Media.





## Web and social media



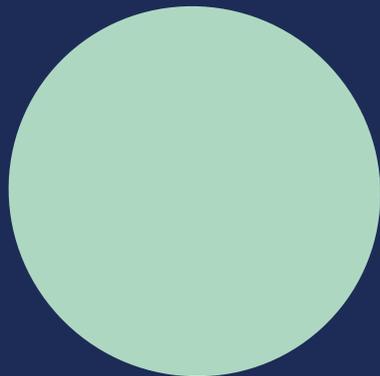
**INSTAGRAM**  
[@sfi\\_norwai](https://www.instagram.com/sfi_norwai)



**WEBPAGE**  
[ntnu.edu/norwai](https://ntnu.edu/norwai)



**LINKEDIN**  
[www.linkedin.com/  
company/norwai](https://www.linkedin.com/company/norwai)



NorwAI

Norwegian Research Center  
for AI Innovation



 The Research  
Council of Norway

 NTNU

Editor: Rolf Dyrnes Svendsen  
Coordinator: Karolina Storesund  
Center photos: Kai T. Dragland, NTNU  
Illustrations: Adobe Stock, AI generated  
Graphic design: Marianne Gilbu, NTNU Graphic Center