# "AI IS NOT JUST ABOUT CODING – UNDERSTANDING ITS BUSINESS VALUE IS KEY."

## Building AI Competence in the AI2Business Project


**SANNA PELTONEN (ED.)**



other publications c43


PUBLISHER:
Vaasan ammattikorkeakoulu | University of Applied Sciences

ISSN 2489-4400 (C, other publications, 43)
ISBN 978-952-5784-85-5 (online)
https://urn.fi/URN:ISBN:978-952-5784-85-5


Design: VAMK | Satu Aaltonen
Layout: Tritonia | Kaisa Gullman
Vaasa 2025

VAASAN AMMATTIKORKEAKOULU
UNIVERSITY OF APPLIED SCIENCES

Wolffintie 30, 65200 Vaasa
julkaisut@vamk.fi
VAMK.fi

AI2BUSINESS

Co-funded by the European Union

VAASAN AMMATTIKORKEAKOULU
UNIVERSITY OF APPLIED SCIENCES

Vaasan yliopisto
UNIVERSITY OF VAASA

# SISÄLTÖ

# ABSTRACT

AI2Business is a joint project by Vaasa University of Applied Sciences and the University of Vaasa, co-funded by the European Union and running from 1.8.2023 to 31.12.2025. Its goal is to strengthen the AI capabilities of the working-age population in Ostrobothnia. The project focuses particularly on the needs of companies, especially small and medium-sized enterprises in the region, by offering accessible learning opportunities, practical examples, and a platform for collaboration.

A key outcome of the project is an open online training program for anyone interested in updating their AI knowledge. The training program is especially designed for employees and managers who want to understand how AI can support business development and competitiveness. The learning material approaches AI from business, technology, and design perspectives and includes three practical AI experiments demonstrating applications in document processing, price prediction, and production quality control.

To support shared learning, the project has also established an open AI competence network that brings together enthusiastic regional participants, experts, researchers, and developers. This network has gathered in four events during the project. Competence network events have provided a forum for exchanging insights, engaging in discussions, and supporting collaborative learning related to AI.

This report presents the main results of the project. It outlines the facilitation model fostering collaborative learning, describes the activities of the competence network, details the AI experiments, and introduces the AI2Business training program.

# TIIVISTELMÄ

AI2Business on Vaasan ammattikorkeakoulun ja Vaasan yliopiston toteuttama Euroopan Unionin osarahoittama ryhmähanke (Hankeaika 1.8.2023-31.12.2025). Hankkeen tavoitteena on vahvistaa Pohjanmaan työikäisen väestön tekoälyosaamista. Projekti keskittyy erityisesti alueen pienten ja keskisuurten yritysten tarpeisiin tarjoamalla helposti saavutettavia oppimismahdollisuuksia, käytännön esimerkkejä sekä avoin osaamisverkosto tekoälyosaamisen vahvistamiseen.

Hankkeen keskeinen tulos on kaikille avoin verkkokoulutusohjelma omien tekoälytietojen päivittämiseen. Koulutusohjelma on suunniteltu erityisesti työntekijöille ja esihenkilöille, jotka haluavat ymmärtää, miten tekoäly voi tukea liiketoiminnan kehittämistä ja kilpailukykyä. Koulutusohjelmaa varten tuotettu oppimateriaali tarkastelee tekoälyä liiketoiminnan, teknologian ja muotoilun näkökulmista ja sisältää kolme käytännön tekoälykokeilua, jotka havainnollistavat tekoälyn soveltamismahdollisuuksia dokumenttien analysoinnissa, hintojen ennustamisessa ja tuotannon laadunvalvonnassa.

Alueen tekoälyosaamisen vahvistamiseksi hankkeessa on koottu avoin osaamisverkosto, joka kokoaa yhteen alueen tekoälystä kiinnostuneita henkilöitä, asiantuntijoita, tutkijoita ja kehittäjiä. Verkosto on kokoontunut neljässä tapahtumassa hankkeen aikana. Osaamisverkoston tilaisuudet ovat tarjonneet foorumin tiedon jakamiseen, näkökulmien vaihtoon, keskusteluihin ja yhteisölliseen oppimiseen tekoälyyn liittyvistä teemoista.

Tämä raportti esittelee hankkeen keskeiset tulokset. Raportissa kuvataan fasilitointimalli työelämälähtöisen oppimisen tukemiseen, osaamisverkoston toiminta, tekoälykokeilut sekä AI2Business hankkeessa kehitetty koulutusohjelma.

Sanna Peltonen (VAMK)

# 1. STRENGHTENING AI ADOPTION IN COMPANIES

Finland aims to strengthen its position as a global frontrunner in artificial intelligence (AI). According to international comparisons, Finland is already considered one of the most advanced AI nations and a digital leader within Europe. (Työ- ja elinkeinoministeriö, 2021)

The AI2Business project was launched to support these national ambitions and to accelerate the use of AI in companies, particularly in small and medium-sized enterprises. The AI2Business project was guided by five key themes considered essential for promoting more widespread, value-creating, and sustainable use of AI. (Työ- ja elinkeinoministeriö, 2021)

AI skills and understanding must extend across the entire organization, from the factory floor to the boardroom. When everyone has at least a basic understanding of AI, they are better equipped to identify opportunities to apply it in their own work and even create new value for the business. The role of company leadership is particularly critical. Oldemeyer et al. (2025) note that in many small and medium-sized enterprises, leaders may feel that their company is too small to adopt AI. Some do not recognize relevant AI application areas in their own operations, while others see no need for change, especially if the business is currently performing well.

At the same time, AI adoption must remain firmly rooted in business needs. AI is not only a technological matter; its use should always align with strategic goals and address real challenges. Experience has shown that while companies often run AI-related pilots, many of these experiments remain too disconnected from core business operations, which limits their ability to generate real impact. (Ailisto et al., 2019) A business-focused approach is one way to ensure that AI strengthens competitiveness and supports sustainable growth rather than becoming an isolated experiment.

Learning from others is also essential. Open AI ecosystems that encourage the sharing of experiences, good practices, and international examples help companies gain insights and build confidence. For example, Yildirim et al. (2023) note that more than 85 % of AI-related innovation projects fail because the solutions do not create sufficient value for users. In such cases, projects may focus on the wrong problem, and the resulting solutions do not address the real needs of the organization. It is therefore important to analyze and openly share experiences to better understand the key factors behind successful AI initiatives.

The Ministry of Economic Affairs and Employment (2019) highlights the importance of providing flexible online learning opportunities for those already in working life, enabling adults to update and renew their competences in response to technological change. Close collaboration between companies and universities provides a strong foundation for continuous learning. The rapid development of AI technologies requires access to up-to-date expertise, while researchers benefit from understanding the practical needs and challenges companies face. This two-way interaction strengthens both innovation and the relevance of research.

Finally, trust plays a central role in AI. Companies need to ensure that AI solutions are safe, reliable, and ethically sound. Establishing trustworthy AI practices is essential for long-term adoption, stakeholder confidence, and sustainable implementation.

## AI2Business – Sustainable business from artificial intelligence

Artificial intelligence (AI) is rapidly transforming the way companies operate. It offers opportunities ranging from automating routine tasks and improving decision-making to enabling better customer experience and new business models. While the potential is widely recognized, many companies still struggle to identify concrete and meaningful ways to apply AI in their own operations.

AI adoption is not only a technological issue; it is equally a design and strategy challenge. AI solutions must serve people effectively, create tangible value, and be deployed ethically and responsibly. A business-oriented perspective is therefore essential to ensuring that AI genuinely supports competitiveness and contributes to meaningful business renewal.

The AI2Business –project was established to meet this need. AI2Business is a joint project by Vaasa University of Applied Sciences and the University of Vaasa, co-funded by the European Union and running from 1.8.2023 to 31.12.2025. Its goal is to strengthen the AI capabilities of the working-age population in Ostrobothnia. The project focuses particularly on the needs of companies, especially small and medium-sized enterprises in the region, by offering accessible learning opportunities, practical examples, and a platform for collaboration.

A central outcome of the AI2Business project is an online training program designed for employees and managers who want to understand how AI can support business development and competitiveness. The training program aims to enhance the availability of diverse, continuous learning opportunities related to AI and to facilitate the transfer of research results into working life.

The learning material created for the training program approaches AI from three interconnected perspectives, business, technology and design, helping organizations see AI not only as a technical tool but also as a strategic asset, a driver of customer value and an enabler of improved service experiences. The learning material also presents three AI experiments that demonstrate how AI can be applied to tasks such as document processing, price prediction, and production quality control.

Collaboration and knowledge sharing play a key role in fostering understanding about AI. The project has established an open AI competence network that brings together regional people, experts, researchers, and developers. This network provides a forum for exchanging insights, experiences and emerging needs related to AI, supporting shared learning and long-term capability building. The goal is to create a lasting foundation for AI-related expertise and cooperation that continues beyond the duration of the AI2Business project.

# References

Ailisto, H., Halén, M., & Seppälä, T. (2019). Tekoälyn kokonaiskuva ja kansallinen osaamiskartoitus—Loppuraportti.

Ministry of Economic Affairs and Employment. (2019). Leading the way into the age of artificial intelligence—Final report of Finland's Artificial Intelligence Programme 2019. http://urn.fi/URN:ISBN:978-952-327-437-2

Oldemeyer, L., Jede, A., & Teuteberg, F. (2025). Investigation of artificial intelligence in SMEs: A systematic review of the state of the art and the main implementation challenges. Management Review Quarterly, 75(2), 1185–1227. https://doi.org/10.1007/s11301-024-00405-4

Työ- ja elinkeinoministeriö. (2021). Suomesta voittaja kaksoisiirtymässä – tavoitteista käytäntöön Tekoäly 4.0 -ohjelma, toinen väliraportti (No. Yritykset 2021:64).

Yildirim, N., Oh, C., Sayar, D., Brand, K., Challa, S., Turri, V., Crosby Walton, N., Wong, A. E., Forlizzi, J., McCann, J., & Zimmerman, J. (2023). Creating Design Resources to Scaffold the Ideation of AI Concepts. Proceedings of the 2023 ACM Designing Interactive Systems Conference, 2326–2346. https://doi.org/10.1145/3563657.3596058

Sanna Peltonen (VAMK) & Janne Pekkala (VAMK)

# 2. FOSTERING COLLABORATIVE LEARNING BETWEEN INDUSTRY AND ACADEMIA

Collaborative learning between academia and industry is essential for developing practical AI capabilities and ensuring that research insights are effectively applied in real-world contexts. The facilitation model developed in the AI2Business project offers a framework for this collaboration, showing how the two sectors can engage in work-based, and cross-sectoral learning. By structuring interaction and knowledge sharing, the model supports the transfer of research and innovation into working life and ensures that industry needs and insights inform academic work.

## Elements of the facilitation model

The facilitation model developed in the AI2Business project describes the interaction between the project team (academia) and the competence network (representing working life and industry actors). It illustrates how collaborative learning and knowledge transfer are structured through iterative, co-creative processes. Each element of the model contributes to building a shared understanding of AI opportunities, skills development, and business applications.

The model is grounded in actions implemented during the AI2Business project, forming a replicable approach to facilitating AI-related competence development between academia and industry. Figure 1 presents the facilitation model, and the following sections examine its key elements in detail.

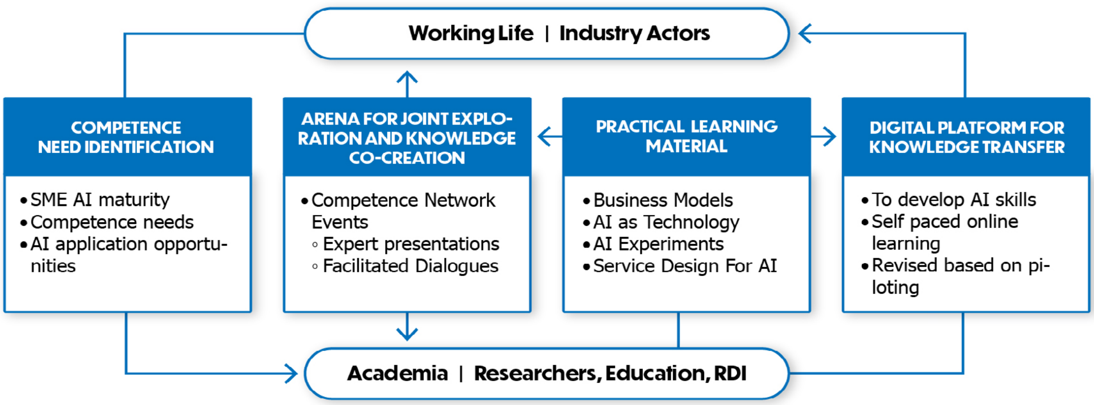


FIGURE 1. Framework for Facilitating an AI Competence Network

Figure 1 visualizes how the four elements of the facilitation model form a mutually supportive learning cycle between academia and industry. The process begins with identifying companies' AI competences and development needs, which creates the knowledge base for joint development and learning activities within the network. These activities are supported by practical learning materials from multiple perspectives that combine project insights, company context, and the latest research on AI solutions. In the final stage, the accumulated knowledge is transferred to a digital learning platform, enabling open and long-term competence development. Each element also feeds back into both academia and industry through two-way interaction, ensuring continuous learning and an iterative model.

The visualized facilitation model does not only show the elements of the model but also helps clarify their relationships, roles, and iterative interactions between academia and industry. It serves as a shared reference point that supports learning, knowledge transfer, and further model development. This makes the facilitation model easier to understand, evaluate, and apply also beyond the project context.

# Competence Need Identification

The first phase of the facilitation model focuses on understanding the current state of AI adoption among regional companies, identifying competence development needs, and recognising shared areas of interest for AI applications.

In the AI2Business project, competence needs were assessed through company interviews. Between October and December 2023, the AI2Business project team interviewed 17 people from 13 companies across Ostrobothnia and South Ostrobothnia. The interviewees represented a range of company functions, including CEOs, ICT managers, and technology directors. The companies represented diverse sectors, including manufacturing, ICT and electronics, and engineering and design. These interviews helped the team to understand how AI is discussed within companies, who lead the conversation, and whether initiatives are driven primarily by business or technology.

These discussions explored AI-related challenges, existing and potential data sources, innovation practices, and the skills companies consider essential for advancing their AI capabilities. The insights gathered were crucial for identifying cross-cutting competence needs and for pinpointing practical business problems that could serve as relevant starting points for further learning activities, experimentation, and network dialogue.

## SME AI Maturity

To synthesise the interview findings, the project team developed company personas representing different levels of AI readiness, experience, and organisational capability. These personas captured typical goals, resources, and challenges identified during the interviews. They were presented in competence network event, included in learning material, and published in project publications.

SME AI Maturity personas can be a valuable tool for reflecting on a company's current situation, setting goals, and formulating methods to enhance AI capabilities and strategies. In the AI2Business project, they also served as reflective tools to bring target groups into project activities, for example when planning communication and engagement actions or developing learning materials.

## AI application opportunities

Based on the interview findings, the project team identified three themes for the AI experiments carried out in the project. These experiments serve as practical demonstrations of AI use cases, showcasing methods, data utilization, and solution design in business contexts. They provide tangible examples of how to develop, evaluate, and manage AI-based solutions.

The experiments were presented to the competence network, providing opportunities for feedback and discussion in relation to industry needs and interests. They serve as hands-on learning tools that bridge theory and practice and are also a central part of the learning materials developed in the project.

# Arena for joint exploration and knowledge co-creation

The AI2Business project organized four events for competence network actors, which were open and free of charge for all interested participants. The events combined expert presentations, project insights, and discussions, serving as arenas for joint exploration and knowledge co-creation. This two-way exchange keeps content relevant to industry needs and encourages participants to share their perspectives, fostering engagement and collective learning.

# Practical learning materials

The learning materials developed by the project team serve as an open-access resource for industry, supporting self-directed learning and ongoing competence development. They are based on company interviews and AI experiments and enriched with the project team's expertise and the latest research on AI-driven business solutions.

The learning materials cover four themes:

1. **AI and Business Models.** AI is transforming business models, requiring companies to rethink how they create and capture value. To harness AI's full potential, organizations must go beyond automating existing processes and explore how AI can generate new opportunities, improve decision-making, and enhance customer experiences.

2. **AI as Technology.** Businesses do not need to master every technical detail to leverage AI effectively. What matters is understanding what AI can do, how different technologies function, and how they can be applied to real-world challenges. Data forms the foundation of AI—without high-quality, well-managed data, even the most advanced systems cannot produce meaningful insights.

3. **AI Experiments.** Practical use cases are essential for understanding AI's capabilities. By seeing how AI has been applied to solve real business problems, organizations can better grasp its potential and anticipate challenges within their own contexts.

4. **AI and Service Design.** A human-centered approach ensures AI serves actual user needs. Identifying

the right problems to solve with AI is crucial, as poorly chosen applications can lead to inefficiencies or missed opportunities. Service design provides a structured framework to align user needs with business goals, particularly in the early phases of development when critical decisions are made.

During the AI2Business project, the learning materials were piloted to collect feedback and assess learning outcomes.

## Digital platform for knowledge transfer

The learning materials produced in the AI2Business project are openly available to anyone interested in the business opportunities of AI. The digital platform offers a flexible, self-paced online course, accessible anytime, anywhere. Users can explore the full course or select individual themes based on their interests. The platform supports long-term use, making the training program accessible to diverse users (e.g., individuals, companies, educational institutions) and fostering the development of AI competencies.

In addition to learning materials, the digital platform brings together all project publications in one place. This allows the platform to offer not only structured learning content but also complementary perspectives on the role of AI in business.

## Feedback

Feedback mechanisms are integrated throughout the process to support continuous improvement and learning. Input from competence network participants on AI experiments, events, as well as participants in piloting the learning materials captures emerging competence needs, preferred activities, insights on the applicability of materials and methods, and reflections on key takeaways. By analyzing and incorporating this feedback, the project team ensures that subsequent activities stay aligned with the evolving needs of the competence network and industry, keeping the co-learning process dynamic and relevant.

The following sections provide a more detailed explanation of the competence network and its activities, as well as the development of AI experiments.

**Tayyab Warraich (University of Vaasa), Marko Kohtamäki (University of Vaasa) & Sanna Peltonen (VAMK)**

# 3. COMPETENCE NETWORK

The AI2Business competence network refers to an open, multidisciplinary, and cross-sectoral collaboration arena that brings together regional actors interested in artificial intelligence. The network includes representatives from companies as well as research, education, and development organizations and potential users, each contributing their perspectives, knowledge, and experiences related to AI. Competence network functions as a platform for shared learning and competence building.

The development of the competence network was guided by four key principles: accessibility, inspiring discussions, face-to-face interaction, and collaborative learning. Four competence network events were organized during the project. Three of the events were held in the evening to accommodate participants' work schedules, while one event was arranged during the Wasa Future Festival to broaden the reach of the target audience. Each event featured invited external speakers who contributed diverse expertise and up-to-date knowledge to the network.

All competence network events were organized as face-to-face meetings, enabling direct interaction, discussion, and the collaborative exchange of ideas. The competence network functions as a learning environment focused on artificial intelligence, offering participants opportunities to share insights, engage in thematic discussions, and explore emerging developments. The project team's primary responsibility involved facilitating the learning process and supporting active participation.

Each event was promoted through LinkedIn, the social media channels of the University of Vaasa and VAMK University of applied sciences, and the Ostrobothnia event calendar, ensuring broad regional visibility.

## Establishing the Competence Network

The establishment of the competence network began alongside the initial data collection phase, which assessed AI maturity within the region. During interviews, potential participants were invited to join the network as a non-obligatory platform for gaining and sharing knowledge. Approximately 80 percent of interviewees expressed interest and became members, and additional participants joined after information was later shared on social media. Motivation for participation revealed strong curiosity about AI and a desire to understand practical applications.

People were encouraged to join during various events where the AI2Business project was presented. Furthermore, a media release announcing the launch of the competence network generated considerable visibility and attracted additional interest. Joining the competence network remained open throughout the entire project. The competence network currently includes more than 100 members and continues to grow as interest in AI and engagement with the platform increases.

## The 1st Competence Network Summit (August 2024)

The program in the very first Competence Network Summit included a presentation on Artificial Intelligence in Business, highlighting AI's applications and impact, followed by an in-depth discussion on the potential and practical use of AI. Another key theme involved the AI maturity of companies in Ostrobothnia, based on the initial project data (company interviews) and presentation on AI experiment themes, along with a dedicated session on project progress and findings.

These themes aligned closely with members stated interests in understanding AI use cases, industry readiness, and emerging opportunities. The first Competence Network Summit concentrated on foundational themes related to AI adoption. The event attracted 49 registered participants, with 45 attending. Feedback collected after the event was highly positive and reinforced the value of expert-driven content.

## The 2nd Competence Network Summit (January 2025)

The second competence network event expanded organizational and practical aspects of AI integration. The event addressed themes such as the dynamics of adoption and adaptation in intelligent business environments and the evolving nature of human-AI collaboration in future workplaces. A panel discussion titled From Theory to Practice: Shaping the Future of AI Workmates examined how AI tools are reshaping work processes. The session also included updates on the AI experiments and the AI2Business project more generally, as well as an overview of the competence network's purpose and development.

A total of 78 participants registered, and 41 attended. Participant feedback highlighted strong appreciation for the balanced combination of theoretical insights, practical perspectives, and interactive dialogue.

## The 3rd Competence Network Summit (August 2025)

The third competence network event adopted a forward-looking orientation, involving expert presentations from both industry and academia. Themes centered on AI-driven innovation, industrial applications, and the connection between research and practice. Discussion focused on the future trajectory of AI in academic training and industry skills development.

Although only 23 participants registered, the event achieved a notably high attendance of 63 participants due to strong interest generated through social media outreach. The event also benefited from being organized as part of the Wasa Future Festival, which contributed to the increased attendance. Feedback indicated satisfaction with the content while emphasizing the need for extended discussion time.

## The 4th Competence Network Summit (November 2025)

The fourth event incorporated feedback recommendation by allocating a longer discussion period while maintaining a strong emphasis on expert contributions. The agenda explored advanced topics including the GenAI revolution in decision-making, trust in AI-driven services, and the transformation of software development through AI. The extended panel discussion, From Data to Decisions: The Role of Trust in AI, reflected growing participant interest in ethical and organizational implications of AI adoption.

A project update and information about the ongoing course piloting activities were also shared with the competence network through the event and follow-up communication. The event received 62 registrations and achieved an attendance of 39 participants.

## Integration of Events, Learning Materials, and Knowledge Transfer

The themes of the competence network were closely connected to several other project activities. Company interviews conducted at the beginning of the project revealed competence needs that the project aimed to address. These discussions were considered when planning the cumulative learning journey for competence network participants.

Topics addressed in the expert presentations were also elaborated in the learning materials and open publications produced within the AI2Business project. These themes included, for example, Artificial Intelligence in Business, the evolving nature of human–AI collaboration in future workplaces, and the role of generative AI in transforming decision-making.

New knowledge created within the project was also disseminated to the competence network. Company personas describing companies' AI maturity and the development of AI experiments are examples of this kind of knowledge transfer. In this way, the competence network, learning materials as well as publications were effectively linked, enriching and deepening the overall learning experience.

### AI Experiments in Competence Network

Based on findings in company interviews, the project team selected three AI experiments to be developed as part of the learning materials. These experiments serve as practical examples of AI applications in business, demonstrating how AI can be applied and illustrating the development, evaluation, and management of AI-based solutions. The three experiments are:

- Intelligent Document Response System
- AI-driven Market Insights
- Intelligent Anomaly Detection

The AI experiment themes were presented twice at the competence network events, where participants were also invited to engage in discussion and provide feedback. In the feedback survey, 89 % of respondents agreed or strongly agreed that the AI experiment cases offered valuable insights and motivation and gave them confidence to begin exploring and applying AI. Although the survey included only nine respondents, the results provide a clear indication that the experiments are relevant and meaningful for the learning material.

The AI experiments were also regularly discussed in the steering group meetings, inspiring the group to engage in discussions about the themes and development of AI experiments. The following sections provide a more detailed discussion of each AI experiment.

Essi Nousiainen (University of Vaasa)

# 4. INTELLIGENT DOCUMENT RESPONSE SYSTEM – USING LARGE LANGUAGE MODELS FOR FINANCIAL ANALYSIS

## Introduction: Using LLMs for financial analysis

The first AI experiment was to study, how Large Language Models (LLMs) could assist in reading long documents. Managers have limited time to examine documents, and AI could assist in speeding up this process. Financial reports were chosen as the use case of the first experiment, since they are long and difficult-to-understand documents. We conducted a study in the AI2Business project to assess LLMs' financial analysis skills. The most important insights from the study, including tips on how to utilize generative AI for financial analysis are summarized here.

LLMs, such as ChatGPT, represent a shift in how financial information can be processed and understood. Unlike traditional tools that rely on structured data, LLMs can interpret unstructured text from annual reports, management discussions, and risk disclosures. This ability opens new possibilities for automating time-consuming tasks and extracting insights from lengthy documents that often span hundreds of pages.

LLMs are designed to understand context and sentiment, making them useful for summarizing strategies, identifying trends, and even supporting forecasting. For financial analysts, this means less time spent on manual reading and more time for strategic decision-making. However, while LLMs excel at summarization and data extraction, they currently fall short in critical interpretation and detecting misleading information.

## What LLMs can do well

The study behind this chapter tested LLMs on real-world financial reports from two Finnish media companies. Here are the key strengths observed:

- **Summarization and clarity.** LLMs can condense complex reports into clear summaries. For example, when asked to outline a company's strategy, the model produced a version that evaluators found clearer and more understandable than the original report.
- **Structured comparisons.** When prompted to create comparison tables between companies, LLMs performed well, providing accurate figures and highlighting differences in profitability, debt levels, and equity strength.
- **Data extraction.** LLMs accurately retrieved financial figures and ratios from reports, such as net sales, EBIT margins, and equity ratios. This makes them valuable for quick fact-checking and compiling key metrics.

- **Sustainability and non-financial topics.** Even when asked about topics like water conservation, which was absent from the reports, the models correctly acknowledged the lack of information and provided related sustainability details.

## Where LLMs struggle?

Despite these strengths, the study revealed important limitations in LLMs' capabilities:

- **Critical thinking.** LLMs tend to accept reported information at face value. They rarely question assumptions or highlight inconsistencies, which is essential in financial analysis.
- **Handling false information.** When given prompts containing incorrect statements (e.g., claiming a profitable company was not profitable), the models repeated the false premise without verification. This is a serious risk if users unknowingly provide inaccurate data.
- **Limited analytical depth.** Responses often mirror the language of the original reports rather than offering independent interpretation or deeper insights. For example, while LLMs summarize well, they do not typically provide forward-looking analysis or challenge management narratives.
- **Dependence on prompt quality.** The level of detail and accuracy in responses depends heavily on how well the prompt is crafted. Vague prompts lead to superficial answers, while structured prompts yield better results.

## Practical applications

LLMs can be integrated into financial workflows in several ways:

1. **Quick summaries of annual reports.** Instead of reading hundreds of pages, analysts can use LLMs to extract key points on strategy, financial performance, and sustainability.
2. **Comparative analysis.** Creating tables comparing companies' fundamentals becomes faster and easier with LLM assistance. It's easy to also make comparisons between companies using generative AI.
3. **Data compilation.** LLMs can pull out ratios and figures, saving time on manual searches. Good prompting strategies help generate readable outputs, e.g. tables.
4. **Investor communication.** Clear summaries generated by LLMs can help explain complex financial information to non-experts.

## Risks and recommendations

While LLMs offer efficiency, they should not replace human judgment. Users must:

- **Verify outputs.** Always cross-check figures and statements against original sources. The models can still generate false information or be otherwise inaccurate.
- **Design better prompts.** Specific, structured prompts lead to more accurate and useful responses. Use thoughtful language and multiple iterations to get better outputs.
- **Maintain critical oversight.** Use LLMs as a support tool, not as a decision-maker. Their inability to detect false premises or provide critical analysis means human review is essential.

## Future Outlook

LLMs are evolving rapidly, and future improvements could include:

**Fact-checking mechanisms.** Built-in verification to prevent propagation of false information. Referencing original sources more precisely to cross-check information.

**Enhanced analytical capabilities.** Moving beyond summarization to provide deeper insights and scenario analysis.

**Integration with financial systems.** Combining LLMs with real-time data feeds for dynamic analysis.

## Summary

LLMs can transform financial analysis by reducing manual workload and improving accessibility of complex information. However, they are not yet a substitute for expert judgment. Think of them as powerful assistants that are great at summarizing and organizing data, but still needing a human analyst to interpret, validate, and make strategic decisions. You can refer to the AI2Business online training program for further prompting advice and more in-depth instruction in document analysis and summarization.

Akpo Siemuri (University of Vaasa)

# 5. USING AI TO FORECAST RAW MATERIAL PRICES

The objective of this study is to analyze the various factors that influence aluminum prices and assess their respective impacts on price fluctuations by implementing some explainable AI tools, for example Shapley values. We also look at how machine learning has been applied in aluminum price prediction. Predicting the supply, demand, and pricing of aluminum is essential for shaping future strategies and investment choices within production and processing sectors. Consequently, a reliable price forecasting model is valuable to traders and stakeholders in the aluminum industry, facilitating informed decision-making processes. Leveraging insights from these analyses, we propose a predictive model using machine learning and transformer-based techniques to forecast aluminum prices. We will also analyze the predictive power of transformers and estimate the importance of aluminum price predictors using Shapley values (SHAP).

## Introduction

Aluminum is classified among non-ferrous metals, gaining prominence in both industrial and daily applications. Extracted primarily from bauxite through electrolytic refining, this metal requires substantial energy input. The resultant material, known as primary ingots, serves as the foundational raw material in various industries [1].

In terms of trading volumes, aluminum stands among the largest commodity exchange for metals globally. Commodity exchanges serve as legal entities facilitating the trading of future delivery contracts (futures). The world's hub for metal trading is the London Metal Exchange (LME), handling approximately 3.7 billion tons of metals annually. Physical delivery of the metal is rare, with trading predominantly conducted among professional market agents and traders who utilize aluminum as a financial instrument [2].

Given the economic ramifications of fluctuations in global metal prices, it is important for nations reliant on metal exports to understand the causes, characteristics, and effects of these fluctuations to formulate appropriate policy responses. These price shifts, stemming from diverse factors, can vary in duration and intensity. In cases where factors have enduring and significant impacts, it becomes essential to enact economic and policy measures to alleviate the consequences of price volatility.

Therefore, aluminum price prediction is very important. Better investments in aluminum can be made only when the price is predicted meticulously.

## Objective of the study

The aim of this study is to investigate the application of machine learning (ML) and deep learning (DL) techniques in predicting aluminum prices. We will be investigating some ML/DL based approaches used for aluminum price prediction and show some results.

However, we first explore the influence of various factors on aluminum prices, along with assessing the extent of these influences. Which factors exert the most pronounced impact on aluminum prices? Afterwards, we proceed with our main aim, which is the prediction of aluminum prices using ML methods.

## Various factors influencing aluminum prices

Various factors affect the price of aluminum and knowing the interaction of these variables is difficult to ascertain. In this section, we will discuss in brief the various factors influencing aluminum prices and the degree to which each factor affects prices.

### FACTORS THAT DETERMINE OR AFFECT THE PRICE OF ALUMINUM

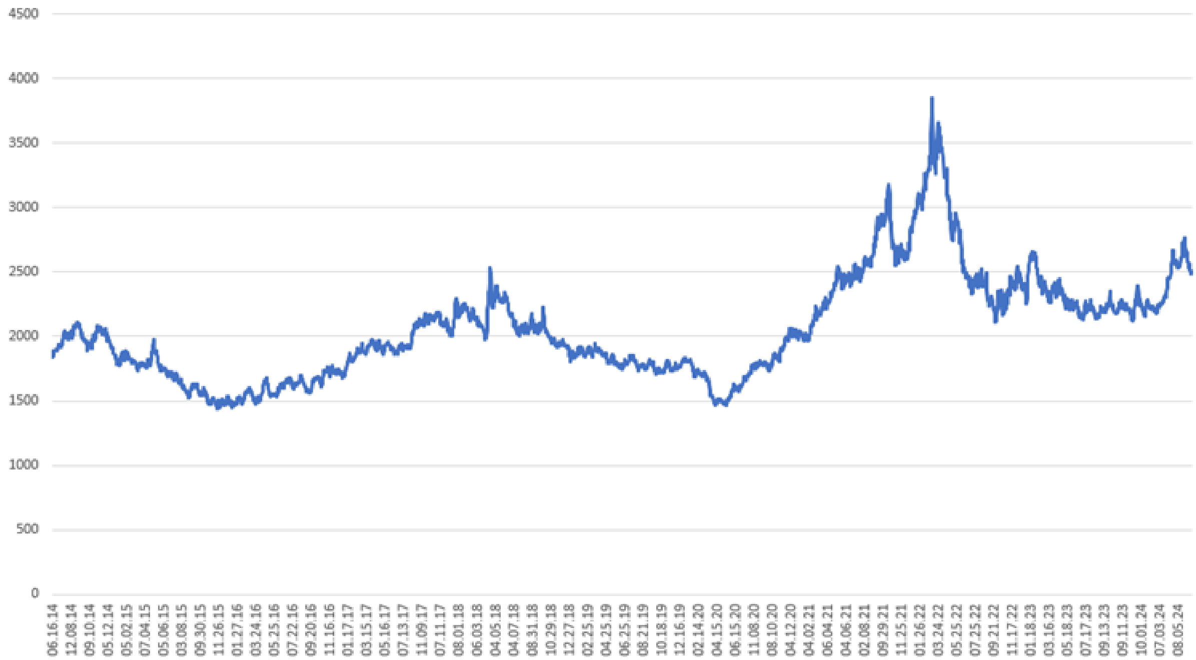

FIGURE 2. Aluminum price fluctuations 2014-2024.

The price trend and price situation of aluminum according to a study by Xinlong Chen [3] shows an interesting price trend of LME aluminum futures over the past decade. According to the study, aluminum prices showed notable volatility, particularly influenced by the onset of the COVID-19 pandemic in 2019, leading to a global recession with disruptions in imports, exports, and facility closures. Between January and March 2020, aluminum prices hit record lows worldwide. However, with subsequent economic recoveries globally, aluminum prices surged, reaching their second-highest inflation-adjusted level in a decade by March 2022. Since then, aluminum costs have decreased by approximately 36% [3]. In figure 2, we find the analysis of the trend of LME aluminum futures from 2012 to 2022.

The factors that determine or affect the price of aluminum are listed below[3][2].

**1. Supply and demand**

Market prices of commodities are directly influenced by the balance between supply and demand. When supply and demand reach equilibrium, prices tend to fluctuate within a narrow range; however, when there is an imbalance between supply and demand, prices can fluctuate significantly. Aluminum, known for its versatility, lightness, corrosion resistance, and recyclability, finds applications in various industries such as construction, automotive, aerospace, packaging, and electronics. The shift towards electric vehicles has increased demand for aluminum, particularly in their structural components and battery packs. Approximately 30% of construction materials in emerging nations comprise aluminum, with construction expansion rates susceptible to factors like interest rates, unemployment, and overall economic conditions, which can impact aluminum demand and prices. Additionally, countries act as both demand and supply sources, influencing prices based on their stockpile accumulation or liquidation activities. Global aluminum demand is concentrated, with a handful of countries dominating both production and consumption. For instance, between 2002 and 2006, Australia, China, and Brazil were the top producers, while China, the United States, Japan, and Germany were the primary consumers[3] .

**2. Alumina supply and pricing**

Alumina accounts for approximately 28%–34% of the production cost of aluminum ingots. Consequently, fluctuations in the supply and price of alumina directly influence the price of aluminum ingots. In 2000, the international spot price of alumina experienced a decline from a high of $450 per ton at the year's onset to a low of $165-175 per ton by year-end. Concurrently, the international futures price of aluminum ingots also decreased from $1,753 per ton to $1,480-1,650 per ton during the same period [3].

**3. Energy supply**

The use of energy in the production of aluminum and other metals is huge. The Aluminum Association finds that electricity contributes to more than a third of the price of aluminum production [4]. The process of producing aluminum requires the use of electricity to convert bauxite into aluminum. Electricity pricing varies based on supply and demand, and when electricity demand rises, its pricing will go up. Therefore, it is a significant factor that influences the price of aluminum. In a study[3], it was stated that European aluminum production has dwindled to its lowest point since the 1970s, attributed to the challenges faced by the energy-intensive industry in coping with rising energy costs. Some analysts suggest that if electricity prices in Europe persistently surpass 90 euros throughout 2022, production cuts could extend by an additional 600,000 tons. By the end of 2022, European production capacity, excluding Russia, is projected to be approximately 1.2 million tons lower than its peak in the second quarter of 2021, resulting in a decline of around 800,000 tons in final production compared to 2021.

**4. Global Economy**

According to a study[3], global recession has significantly impacted aluminum prices, evident from their downward trajectory since the recession's onset. Following the severe recession triggered by the pandemic, there was a sharp rebound in global economic activity, albeit followed by a subsequent slowdown due to tightened policies. China, a major consumer of metals, experienced production halts and ore shutdowns during the pandemic, leading to shifts in global demand dynamics. Factors such as the Russia-Ukraine conflict, pandemic-related policies, and inflation control measures have further dampened global economic activity. Notably, from January to April 2020, aluminum prices witnessed an 18% plummet, the largest drop in nearly a decade. However, from April 2020 to March 2022, aluminum prices more than doubled, marking the most significant increase in three decades for the equivalent time frame. Subsequently, a quarter of these gains were retraced within five months due to heightened fears of a global recession, exerting downward pressure on commodity prices overall.

**5. Commodity currency**

Certain countries heavily rely on exporting specific commodities, leading to their currency closely tracking the global price of that primary commodity. These currencies are termed commodity currencies. According to [5], these currencies possess strong predictive capabilities regarding commodity prices, including aluminum. The theoretical rationale behind this lies in two factors: firstly, commodity prices influence the nominal exchange rate value of the country, as they reflect its terms of trade. Secondly, nominal exchange rates embed expectations of future fundamental values, such as commodity prices, aiding in their prediction. However, Chen et al.'s study did not find robust evidence supporting the reverse relationship, where commodity prices predict exchange rates [2]. In the study conducted by Pablo et al. (2021) [6], commodity currencies were employed for predicting aluminum prices. The findings were based on the relationship between commodities/commodity currencies and the present-value model for exchange rates.

## Machine learning in predictive analysis

The growing accessibility of extensive historical data and the necessity for precise forecasting across various scientific and practical fields necessitate the development of robust and efficient techniques capable of inferring stochastic dependencies between past and future observations. Since the 1960s, the forecasting domain has predominantly relied on linear statistical methods like ARIMA models[7]. However, in recent times, machine learning models have emerged as significant contenders, challenging traditional statistical approaches and gaining traction within the forecasting community.

A time series refers to a sequence S comprising historical measurements yt of an observable variable y taken at uniform time intervals. Time series analysis serves various objectives, including forecasting future trends based on past data, comprehending the underlying phenomenon driving the observations, or succinctly describing the prominent characteristics of the series[7]. Time series data can be used for analyzing trends, cycles, and other temporal patterns.

Forecasting the price of aluminum and other metals using a ML method has been a field of interest amongst researchers. In the next section, we discuss related studies where ML has been applied for forecasting the price of aluminum.

Similarly, deep learning (DL) and transformers models have been applied in various time-series prediction applications. Transformer can closely track the trend of stock prices and make predictions in line with the direction of stock price changes. According to[8], in the study of stock price forecast, their results showed that when compared to LSTM and HMM, Transformer performed the best-predicting capability of stocks of the new energy sector in the A-share market. The model's performance was quantified using the Mean Absolute Percentage Error (MAPE) and Matthews Correlation Coefficient (MCC).

### Explainability of ML/AI models

For the easy adoption of ML/DL and transformer models in industries, the explain-ability of the models is an important aspect to give confidence in the models. Predictive ML/DL models can become even more potent when combined with interpretability tools like SHAP. These tools highlight the most significant relationships between the input features and the predicted outcome, aiding in explaining the model's behavior, securing stakeholder approval, and identifying potential issues. In this study, we will attempt to explain the models and predictors using Shapley values (SHAP). Shapley values are a widely used approach from cooperative game theory that come with desirable properties.

# Related research

In this section, we present from previous literature some approaches used for metal price prediction or forecasting.

## Methods used for forecasting commodity prices

There are some statistical methods that have been applied in the forecasting of the price of metals including the price of aluminum. An example is the Vector Error Correction Mechanism (VECM) model. VECM models are a special application of VAR or Vector Autoregressive Models. The specification of VECM models involves the introduction of error correction terms into the VAR models. VECM methodology is used when the variables in the system have a long-run relationship, that is, they are cointegrated. The Auto-Regressive Integrated Moving Average (ARIMA) model is a widely used statistical approach for commodity price forecasting. In a study by Dooley and Lenihan (2005), ARIMA was applied and compared to lagged forward price models for predicting monthly spot prices of lead and zinc. The findings were mixed: for zinc, the study did not present conclusive evidence favoring either method. However, in the case of lead, the ARIMA model outperformed the lagged forward price approach in nine out of sixteen instances [2].

In the master's thesis titled "Forecasting the Price of Aluminium Using Machine Learning,"[2] the authors compare the effectiveness of machine learning models (XGBoost) against traditional statistical methods (VECM and ARIMA) for predicting aluminum's 3-month futures prices. Regardless of adjustments to the assumptions concerning endogenous and exogenous variables or shifts in the cross-validation starting point, the conclusion remained consistent with the primary analysis: the machine learning model XGBoost consistently outperforms the statistical VECM model in forecasting the 3-month aluminum futures price at a one-month horizon.

In another paper titled "Stock Price Forecast: Comparison of LSTM, HMM, and Transformer"[8], the authors compare the effectiveness of three neural network models for stock price prediction: Long Short-Term Memory (LSTM), Hidden Markov Model (HMM), and Transformer. The study evaluates these models using data from the new energy vehicles sector in the A-share market. Performance metrics such as Mean Absolute Percentage Error (MAPE) and Matthews Correlation Coefficient (MCC) are employed to assess predictive accuracy. The results indicate that the Transformer model outperforms the other two in forecasting stock prices within this sector.

# Data collection and processing

## Variable selection

In this section, we present some selected variables from previous literature used for aluminum price prediction. In [2], some of the variables used for aluminum price prediction include Aluminum 3-month futures price, Aluminum spot price, Industrial production, AUD/USD exchange rate (aluminum currency exchange rates) and co-movement observed between different metal prices.

Earlier studies of other commodities presented in the literature review section indicate that spot and futures prices often follow each other. It is not unlikely that this is true for aluminum, suggesting that the spot price should be used as a predictor.

### Spot price

In[9], Coppola examines the relationship between spot prices and futures prices in the context of oil price forecasting. Given that many financial time series, including commodity spot and futures prices, are often non-stationary, there's a risk of obtaining misleading results. This occurs when the model suggests a relationship between variables that does not actually exist. To address this, Coppola[9] conducts a cointegration test to investigate the relationship between futures contract prices and spot prices prior to applying the forecasting method. Ultimately, this research concludes that a cointegration relationship does exist. This finding is also supported by other studies

On the London Metal Exchange, the principle of futures-spot convergence applies accurately [10]. This principle asserts that the futures contract price aligns with the spot price of the underlying commodity, making the futures contract price approximately equal to the spot price at the delivery date [11]. Consequently, it is reasonable to anticipate a long-term relationship between the two prices of a given commodity.

### Industrial activities and Market yield 3M

According to the study [12], it examined the factors influencing metal prices, particularly focusing on macroeconomic impacts on the common factor, which refers to the co-movement between metal prices. The study concludes that industrial activity holds the highest predictive power over metal prices. In [12], causality tests indicate that fluctuations in the prices of copper, lead, and zinc may influence aluminum prices, and that changes in zinc prices tend to precede changes in lead and tin prices.

### European wholesale day-ahead electricity price

The production of aluminum is significantly more energy-intensive compared to steel. According to recent data, producing one ton of aluminum requires approximately 211 GJ (Gigajoules) of energy, whereas producing one ton of steel only requires about 20-30 GJ of energy. This stark difference highlights the high energy demands of aluminum smelting, which is why fluctuations in power prices can have a considerable impact on aluminum production costs [13, 14].

To further explore the specific energy requirements and the factors influencing the production of aluminum, you can refer to the World Economic Forum's Net-Zero Industry Tracker 2023 [13] and the Light Metal Age Magazine [14]. These sources provide detailed insights into the energy consumption patterns and the shift towards renewable energy in the aluminum industry.

### Shanghai Export Containerized Freight Index (SCFI)

First introduced in 1998 by the Shanghai Shipping Exchange, the China Export Containerized Freight Index (CCFI) measures freight rates and volume across 12 major global trade routes, reflecting fluctuations in shipping costs. Chinese departure ports include Dalian, Tianjin, Qingdao, Shanghai, Nanjing, Ningbo, Xiamen, Fuzhou, Shenzhen, and Guangzhou. The CCFI is updated weekly, providing timely insights into shipping cost changes for exports from China. Similarly, the Shanghai Export Containerized Freight Index (SCFI) also offers weekly updates, with both indices indicating current global trade conditions due to China's status as the largest exporter.

### Aluminum currency exchange rates

As indicated in a study [2], the commodity currency Australian dollar/USD exchange rate can be considered as a predicting variable as well since Australia is a significant aluminum producer.

### NEWC Index and Newcastle coal futures

In [2], over 80% of aluminum production in China uses coal-fired power. Therefore, the NEWC Index, which is used as a benchmark for seaborne thermal coal prices in the Asia-Pacific region is important as well. It serves as the main price reference for physical coal contracts in Asia. It is the settlement price for a significant volume of index-linked contracts, ranging from Australia and Indonesia to Japan and India. Additionally, the NEWC Index is the settlement basis for ICE's globalCOAL Newcastle Coal Futures, thereby pricing the vast majority of coal derivatives in Asia [15].

Newcastle coal futures hovered around $135 per tonne, trading near weekly lows, as China progresses toward its climate objectives and adheres to stricter international carbon standards. In May, China's coal generation share fell to a record low of 53%, down from 60% the previous year, while renewable energy sources reached new highs. Additionally, China did not approve any coal-based steel-making projects in the first half of the year, marking a first since announcing its major climate neutrality goals in 2020. Conversely, coal-fired power generation in the US is expected to increase in 2025 compared to 2023 levels due to high natural gas prices [16].

## Example data

An example of a dataset that could be used for training a machine learning model for forecasting aluminum price is shown in the figure 3 below. Besides daily aluminum price itself, it contains 8 other variables, all time-aligned with aluminum prices and interpolated to fill missing values (e.g., weekends).

The figure 4 above shows an example where a 4-unit LSTM network forecasts aluminum prices 10 days ahead in a closed-loop manner, on a daily basis. The neural network has been trained with multivariate

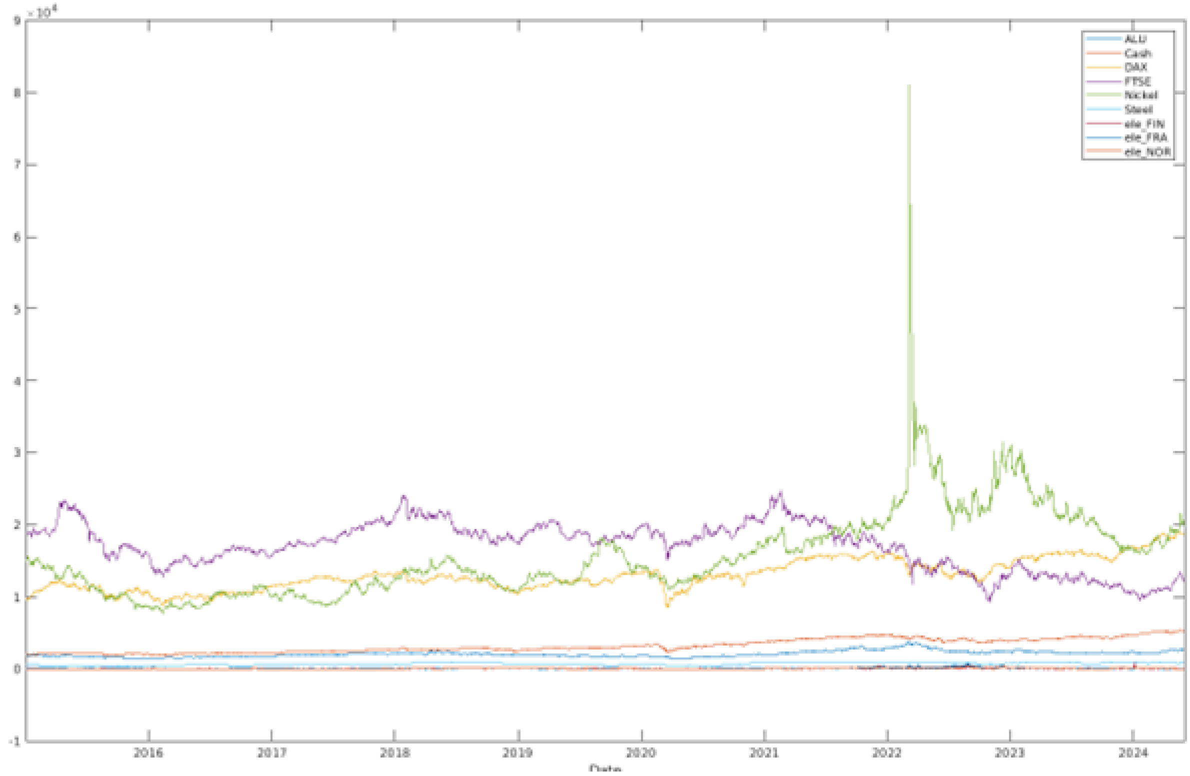


FIGURE 3. An example of a dataset that could be used for training a machine learning model for forecasting aluminum

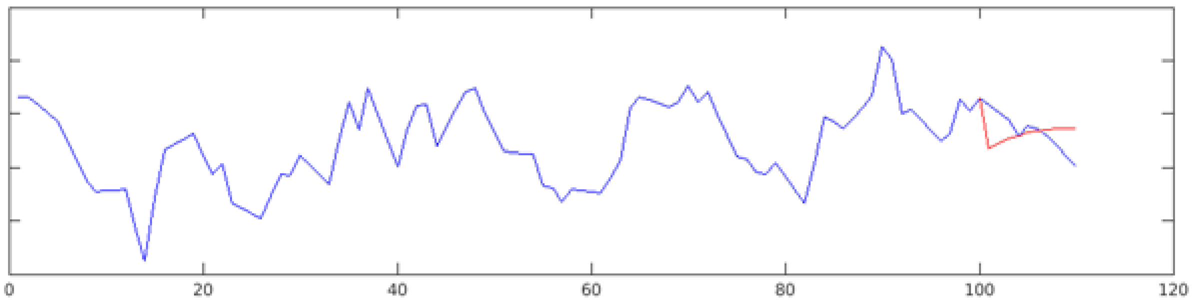


FIGURE 4. Real-life daily aluminum price fluctuations (normalized data) in blue color, and a 10-day ahead multivariate daily forecast computed by a 4-unit LSTM neural network starting from day 100 (in red)

data similar to the previous figure. The forecast shows that the neural network correctly forecasts a short-term price drop, however the steepness of the drop is mispredicted somewhat. On a longer term (2-10 days ahead) the network starts to make its forecast more conservative, approaching a mid-term average value. Such behavior is intuitive and reflects the fact that long-term forecasts of market fluctuations can be even more challenging to forecast correctly than weather. An interesting direction for future work would be expanding the forecast with predictor variables that could bear a causal connection to aluminum prices, going beyond market prices of commodities.

# References

[1] L. W. Agata Wzorek, Oleksandr Ivashchuk, "Analysis of the factors influencing the price of aluminum on the world market," MECHANIK NR, 7/2017. [Online]. Available: https://www.mechanik.media.pl/artykuly/analiza-czynnikow-ksztaltujacych-ceny-aluminium-na-swiatowym-rynku.html

[2] E. M. T. Stina Johanne Mysen, "Forecasting the price of aluminium using machine learning," in Norwegian School of Economics Bergen, 2021. [Online]. Available: https://openaccess.nhh.no/nhh-xmlui/bitstream/handle/11250/2985394/masterthesis. pdf?sequence=1&isAllowed=y

[3] X. Chen, "Factor analysis and influence of aluminum price," BCP Business and Management, vol. 44, pp. 137–143, 04 2023.

[4] M. D. Platzer, "Effects of u.s. tariff action on u.s. aluminum manufacturing," Congressional Research Service, US, 2018. [Online]. Available: https://sgp.fas.org/crs/misc/IF10998.pdf

[5] Yu-Chin Chen, K. Rogoff, and B. Rossi, "Can exchange rates forecast commodity prices?" Quarterly Journal of Economics, vol. 125, no. 3, p. 1145–1194, 2010

[6] P. Pincheira and N. Hardy, "Forecasting aluminum prices with commodity currencies," Resources Policy, vol. 73, p. 102066, 2021. [Online]. Available: https://doi.org/10.1016/j.resourpol.2021.102066

[7] G. Bontempi, S. Ben Taieb, and Y.-A. Le Borgne, Machine Learning Strategies for Time Series Forecasting. Springer-Verlag, 01 2013, vol. 138.

[8] Q. Wang and Y. Yuan, "Stock price forecast: Comparison of lstm, hmm, and transformer," in Proceedings of the 2nd International Academic Conference on Blockchain, Information Technology and Smart Finance (ICBIS 2023). Atlantis Press, 2023, pp. 126–136. [Online]. Available: https://doi.org/10.2991/978-94-6463-198-2_15

[9] A. Coppola, "Forecasting oil price movements: Exploiting the information in the futures market," Journal of Futures Markets, vol. 28, pp. 34 – 56, 01 2008.

[10] J. Jia and S. B. Kang, "Do the basis and other predictors of futures return also predict spot return with the same signs and magnitudes? evidence from the lme," Journal of Commodity Markets, vol. 25, 03 2021.

[11] Y.F. Chow, M. McAleer, J. Sequeira, "Pricing of forward and futures contracts," Journal of Economic Surveys, vol. 14, pp. 215–53, 02 2000.

[12] W. Labys, A. Achouch, and M. Terraza, "Metal prices and the business cycle," Resources Policy, vol. 25, no. 4, pp. 229–238, 1999. [Online]. Available: https://doi.org/10.1016/S0301-4207(99)00030-6

[13] WEForum. (2023) Net-zero industry tracker 2023. [Online]. Available: https://www.weforum. org/publications/net-zero-industry-tracker-2023/in-full/steel-industry-net-zero-tracker/

[14] A. Tabereaux. (2023) The shift toward renewable power in aluminum smelting. [Online]. Available: https://www.light-metalage.com/news/industry-news/smelting/ the-shift-toward-renewable-power-in-aluminum-smelting/

[15] globalCOAL. (2024) The newc index. [Online]. Available: https://www.globalcoal.com/ coalprices/newcindex.cfm

[16] TRADING-ECONOMICS. (2024) Trading economics: Coal, summary. [Online]. Available: https://tradingeconomics.com/commodity/coal

Abol Basher (University of Vaasa) & Jani Boutellier (University of Vaasa)

# 6. USING AI FOR IMAGE ANOMALY DETECTION

## Introduction to Visual Anomalies and Detection Strategy

An anomaly refers to an observation that deviates from the expected concept of normality. In the visual domain, anomalies manifest as unwanted deformations present in images, videos, or point clouds. These may include subtle changes such as thin scratches or major structural defects such as missing components. In time-series data, anomalies correspond to abnormal behaviors or patterns that do not occur during regular operation. Industrial anomalies encompass a broad spectrum of defects ranging from cosmetic imperfections to severe malfunctions.

Detecting anomalies is essential across a wide array of domains. In industrial manufacturing, anomaly detection identifies faulty products and monitors potential intrusions. In market analysis and environmental forecasting, anomalies indicate unexpected behavior or conditions. Additionally, anomaly detection is used to identify illegal activities, diagnose diseases such as seizures, and detect fraudulent financial card usage. These diverse applications illustrate the foundational role of anomaly detection in ensuring safety, efficiency, and reliability [4].

## Vision-Based Data Acquisition and Challenges

Visual anomaly detection relies heavily on imaging and sensor technologies including LiDAR, CCTV cameras, industrial cameras, and other time-series data capturing tools. Such data serve as the foundation for automated inspection and real-time monitoring systems.

Implementing AI-based visual anomaly detection in real-life environments presents several challenges. A lack of sufficient real-world data limits the ability to train models that generalize effectively. Illumination variation makes feature extraction difficult, while pose and perspective differences complicate the interpretation of objects due to the change in appearance with viewpoint. Heterogeneous object characteristics reduce application performance because of large variations in object shape, size, and texture. Distinguishing events in sparse versus dense scenes is also challenging, as methods tailored for sparse settings often fail in highly crowded environments. Finally, occlusion severely complicates detection and tracking when relevant objects are partially hidden.

## Industrial Anomaly Detection Datasets

Several datasets support research in industrial anomaly detection. MVTec AD offers over five thousand high-resolution color images across fifteen object and texture categories with a wide variety of defect types. The VisA dataset contains more than ten thousand images including both normal and anomalous examples with pixel-level labels, making it one of the largest industrial datasets. Real3D-AD includes 3D point cloud data across twelve categories where each point cloud may have forty thousand to million data points. Additional datasets cover a range of visual and video-based industrial environments [5].

## Evaluation Metrics

Evaluation metrics for anomaly detection depend on whether higher or lower values indicate better performance. Common metrics include area under the receiver operating characteristic curve (AU-ROC), F1 score, Average Precision (AP) / pixel-level precision, recall, and Intersection Over Union (IOU) [5]. These metrics help to compare models on standardized datasets.

## Industrial Visual Anomaly Detection and Localization

Industrial visual anomaly detection aims to identify abnormal images and localize defective regions. The task is especially challenging due to the rarity and variability of anomalies. Minor surface imperfections and significant structural defects must all be recognized despite limited examples. Standard benchmark datasets such as MVTec AD, VisA, and Real3D-AD exemplify the diversity and complexity of industrial anomalies.

## Categories of Visual Anomaly Detection Methods

Anomaly detection methods can be classified based on the type of data involved—images, videos, time-series signals, 3D point clouds, or multimodal data. They can also be grouped according to methodological approach, including reconstruction-based, synthesizing-based, embedding-based, inpainting-based, and optical-flow-based methods. For time-series data, popular approaches include statistical, classification-based, vocabulary-based, and reconstruction-based models.

A further distinction is made based on supervision level. Fully supervised methods require abundant labeled samples, whereas semi-supervised and unsupervised methods operate under limited or extremely scarce labeled data. Few-shot and zero-shot learning approaches aim to extend anomaly detection capabilities to situations with minimal or no class-specific data [2].

## Reconstruction-Based Methods

Reconstruction-based methods assume that anomalous regions cannot be accurately reconstructed when models are trained only on normal data. Generative models, including autoencoders and generative adversarial networks (GANs), attempt to reconstruct the input and identify discrepancies as anomalies [1]. Some approaches treat anomaly detection as an inpainting problem by erasing random patches and predicting the missing content. However, if a decoder becomes overly powerful, it may reconstruct anomalous regions smoothly, causing the method to fail at distinguishing anomalies.

## Synthesizing-Based Methods

Synthesizing-based methods artificially generate anomalies on top of normal images or features to train discriminative models. For example, DREAM learns a joint representation between anomalous images and their reconstructions, establishing a boundary between normal and abnormal patterns. CutPaste generates synthetic anomalies by cutting and pasting patches within images, though its ability to simulate the full diversity of real-world anomalies is limited.

## Embedding-Based Methods

Embedding-based methods project normal features into a compressed space so that anomalous features lie far from normal clusters. Models commonly use pretrained networks such as ImageNet-based CNNs for feature extraction. PaDiM models patch-wise feature distributions via multivariate Gaussian distributions, whereas PatchCore uses a maximally representative memory bank of normal features. During inference, distances such as Mahalanobis or maximum feature distance to quantify anomaly likelihood. However, industrial images often differ significantly from ImageNet data distributions, limiting the effectiveness of pretrained models without adaptation. Memory bank operations can also be computationally expensive on resource-constrained devices.

# SimpleNet for Image Anomaly Detection

SimpleNet[3] addresses limitations in both synthesizing-based and embedding-based approaches. It uses a feature adaptor to correct biases inherited from pretrained CNNs and synthesizes anomalies in feature space rather than image space. Its architecture consists of a pretrained feature extractor, a shallow feature adaptor, an anomaly feature generator that injects Gaussian noise into normal features, and an anomaly discriminator that distinguishes normal from abnormal features. A truncated L1 loss prevents overfitting. The final anomaly score corresponds to the maximum value in the anomaly map.

# Creating an MVTec AD -Style Anomaly Detection Dataset

The MVTec Anomaly Detection (MVTec AD) dataset provides a widely adopted structure for image-based industrial anomaly detection, and it is relatively straightforward to extend it with a new object class of one's own. Constructing such a dataset requires four primary components: healthy training images, healthy test images, anomalous test images, and corresponding segmentation masks for the anomalous samples. Typically, several hundred healthy training images are required for the model to learn what a defect-free object looks like. Additional healthy and anomalous test samples are needed to evaluate performance, and segmentation masks serve as a reference for pixel-level anomaly localization. While masks are recommended for robust model evaluation, they are not strictly mandatory.

## Imaging Equipment Considerations

The type of imaging equipment used to collect training and test imagery depends on the physical properties of the target objects and the type of measurement. The discussion here assumes acquisition under normal lighting conditions using color imagery, but the same principles extend to grayscale, X-ray, depth, or other modalities as long as the images can be represented in standard formats such as PNG. Although original MVTec AD images generally have resolutions near 700×700 to 1000×1000 pixels, some models, including SimpleNet, operate at lower resolutions such as 256×256 without significant performance loss.

Cameras with HD resolution around 1920×1080 pixels are typically sufficient for regular RGB imagery. Factors such as the physical size of the object, distance from the camera, and lens focal length influence how prominently the object appears within the captured frame. Once the raw images are cropped to a standardized size, such as 512×512 pixels, the object should occupy most of the frame so that fine geometric and textural details remain visible.

## Optical Characteristics and Illumination

Short focal lengths require objects to be physically close to the camera, but they can introduce radial distortions. Components with significant depth variation may exhibit local blur, occlusion, and perspective deformation. Longer focal lengths mitigate these issues but necessitate bulkier optics and adequate light. Affordable imaging setups commonly rely on shorter focal lengths, making careful control of illumination equally important. Lighting should be uniform, avoiding point-like reflections unless specifically required to highlight defect characteristics. Increased lighting intensity generally reduces sensor noise, and the use of matte white surroundings around the imaging area helps create evenly distributed illumination.

## Imaging Environment and Data Collection

The imaging environment used to collect training and test images should match the deployment environment as closely as possible. Deviations in lighting, camera geometry, or surface appearance between training and operational settings often reduce detection accuracy. When collecting data, it is advisable to place objects near the optical axis to minimize distortions. Sample orientation should be kept consistent, and potentially anomalous regions should be positioned so they are clearly visible in the capture.

## Segmentation Masks

Segmentation masks play an important role in evaluating pixel-wise performance of anomaly detection models. These masks indicate, at a pixel-by-pixel level, the exact locations of anomalous regions in each test image. While not required for training, they are highly recommended for comprehensive evaluation. Constructing segmentation masks is often a manual and time-consuming task, and one mask must be produced

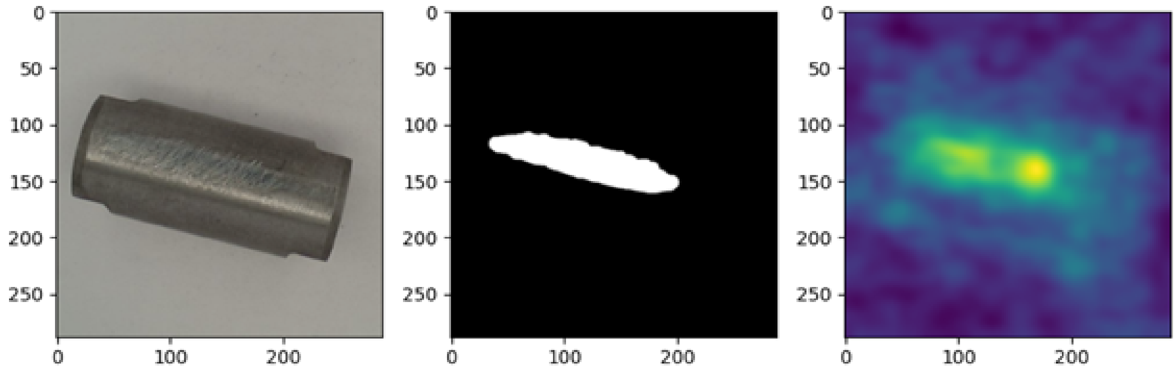


FIGURE 5. An anomaly detection result example for a custom MVTec AD class. Left: test image (anomalous); center: ground truth segmentation mask with anomalous are in white; right: SimpleNet detection result of the test image

for each anomalous test sample. Free annotation tools such as Label Studio (https://labelstud.io/) assist in the creation of these masks, but the effort remains considerable.

## Training SimpleNet with a Custom Dataset

Training SimpleNet on a newly created dataset follows a straightforward process once the data are placed into the default MVTec AD directory structure. Only a small modification in the SimpleNet configuration file is required to specify the new class name. Training typically takes around 30 to 45 minutes on a standard GPU-equipped computer. The folder structure includes separate directories for training images, healthy and anomalous test images, and segmentation masks aligned with the MVTec AD convention.

## Testing the Custom-Trained SimpleNet Model

After training, SimpleNet produces a pixel-wise anomaly map alongside the ground truth during testing. Anomaly detection results can be quantified by a numerical anomaly score. On the original MVTec AD dataset, SimpleNet achieves image-wise anomaly detection accuracy above 99.5 percent and pixel-wise accuracy above 98 percent. If the model performs substantially worse on a custom dataset, the shortcomings typically manifest as false positives or false negatives. A false positive occurs when a healthy sample is mistakenly classified as anomalous, whereas a false negative occurs when an anomalous sample is misclassified as normal.

It is important to note that the intensity of the anomaly map does not directly correspond to anomaly severity; instead, the map is scaled relative to internal features of the model.

## Sources of Misclassification

False positives often arise when unusual attributes appear in test images that were not present during training. These may include geometric distortions caused by the camera, reflections or shadows, objects cropped too tightly, or training sets that are too small or too homogeneous. False negatives may result from anomalous samples leaking into the training set or from anomalies that are inherently difficult for the model to detect. In particular, anomalies that manifest solely in an object's silhouette or outline are challenging for texture-based anomaly detectors like SimpleNet. The precision of segmentation masks also influences apparent performance during evaluation. Furthermore, anomalies must differ sufficiently from healthy samples for the model to recognize them reliably.

## References


[1] Liu, Shenghua, Bin Zhou, Quan Ding, Bryan Hooi, Zhengbo Zhang, Huawei Shen, and Xueqi Cheng. "Time series anomaly detection with adversarial reconstruction networks." IEEE Transactions on Knowledge and Data Engineering 35, no. 4 (2022): 4293-4306.

[2] Cao, Yunkang, Xiaohao Xu, and Weiming Shen. "Complementary pseudo multimodal feature for point cloud anomaly detection." Pattern Recognition 156 (2024): 110761.

[3] Liu, Zhikang, Yiming Zhou, Yuansheng Xu, and Zilei Wang. "Simplenet: A simple network for image anomaly detection and localization." In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 20402-20411. 2023.

[4] Cao, Yunkang, Xiaohao Xu, Jiangning Zhang, Yuqi Cheng, Xiaonan Huang, Guansong Pang, and Weiming Shen. "A survey on visual anomaly detection: Challenge, approach, and prospect." arXiv preprint arXiv:2401.16402 (2024).

[5] Liu, Jiaqi, Guoyang Xie, Jinbao Wang, Shangnian Li, Chengjie Wang, Feng Zheng, and Yaochu Jin. "Deep industrial image anomaly detection: A survey." Machine Intelligence Research 21, no. 1 (2024): 104-135.

Sanna Peltonen (VAMK) & Janne Pekkala (VAMK)

# 7. AI2BUSINESS ONLINE TRAINING PROGRAM

The AI2Business training program is open to anyone who wishes to strengthen and update their knowledge of AI in a business context. It can be used by micro, small, and medium-sized enterprises (SMEs), as well as other organizations, to support their employees in AI-related training. The program is designed for flexible learning, whether for those already in working life or aiming to enter it. It is free of charge and requires only the creation of a personal user account. The course is delivered via the Moodle platform, a widely used learning management system that supports self-paced, personalized learning. The training program offers a flexible learning experience through a online course that can be accessed anytime and anywhere.

The training program can be accessed at: https://moodle.muovadigital.net/mod/page/view.php?id=680


FIGURE 6. An example of the AI2Business online course layout, illustrating the module-based structure.

## The Structure and Logic of the Course Content

The AI2Business training program consists of five modules that together form a logical and structured content package. Learners begin by building a shared understanding of what AI is and why it matters in business (Module 1). They then explore how AI affects business models and value creation (Module 2), followed by an accessible overview of AI technologies and their practical use in organizations (Module 3). The fourth module presents real examples of AI experiments, demonstrating how AI models can be developed and applied in business contexts. The final module introduces a human-centered approach to developing AI-enabled services (Module 5), emphasizing how companies can align business knowledge, technology and user needs.

The modules can be completed from start to finish or used selectively as needed. This flexible structure makes the course suitable for different professional roles, workplace learning situations and practical problem-solving.

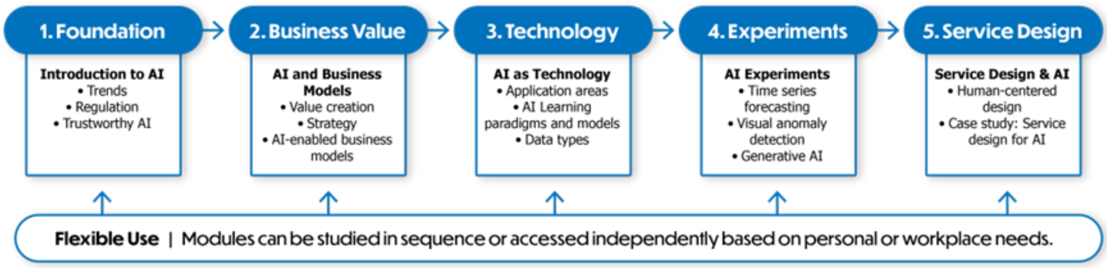


FIGURE 7. The AI2Business course structure, showing both the sequential module order and the option to access modules independently as needed

# Learning Resources and Exercises

The learning material consists of a variety of resources designed to support different learning needs. Written content is complemented with visual elements, and the aim is to offer versatile material that combines text, videos, and graphics to enhance learning and engagement. The material is intended to provide new knowledge and inspiration for all learners, regardless of their prior experience or expertise in AI. Each section introduces the theme and guides learners step-by-step through the content providing both theory and practical examples.

The section focusing on AI experiments is more technical in nature. Learners with some prior knowledge of AI will find it easier to grasp the technical aspects, but the presentations are clear and engaging, making the content valuable and accessible for all participants. This part of the course also provides new insights for those with a deeper interest or involvement in AI development.

At the end of each theme, there are exercises or tasks to reinforce learning. Students can test their new knowledge by answering multiple-choice questions or deepen their understanding through more advanced assignments. The tasks are designed to guide the learning process and can be used either to assess one's own knowledge or simply for practice and enjoyment.

At the end of the course, students complete a feedback form designed as a tool for reflecting on their own learning. The form prompts learners to consider what they have understood about AI, how the course has influenced their ability to recognize AI opportunities, and how they might apply this knowledge in their work or professional development. In addition to rating statements about learning outcomes and experiences, students are encouraged to articulate key insights and personal takeaways.

# Feedback from AI2Business Online Course pilot

The training program developed in the AI2Business project was piloted as an open, web-based online course. The pilot was launched on 30.5.2025 and will continue until the end of the year. This report covers the results up to early December 2025. After the piloting, the platform will remain accessible for users.

In the piloting feedback survey participants were asked to reflect on their own learning, the usefulness of the course materials, and the functionality of the platform. The feedback survey included both Likert-scale statements and open-ended questions. In total, 98 users registered to the piloting (until 5.12.2025), of whom 9 responded to the survey.

Feedback from the course indicates a clearly positive learning experience among participants. Although the sample size is small, the results consistently show strong satisfaction with the course content, structure, and learning outcomes.

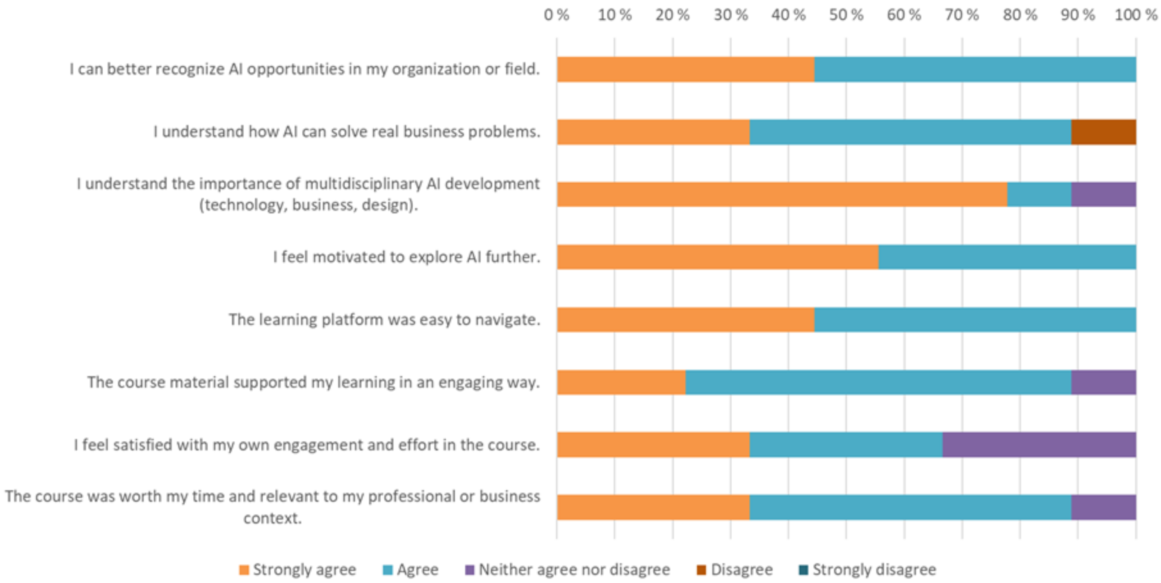


FIGURE 8 presents the results of the Likert-scale questions from the feedback survey

Participants widely agreed that the course enhanced their ability to recognize AI opportunities and understand how AI can address real business challenges. They also strongly acknowledged the value of multidisciplinary perspectives (technology, business, and design) in AI development. Motivation to continue exploring AI remained high, suggesting that the course successfully fostered further interest and engagement.

The learning platform was perceived as easy to navigate, and the course material was viewed as supportive and engaging. Respondents also expressed satisfaction with their own level of effort during the course. Overall, participants agreed that the course was worth their time and relevant to their professional or business context.

## Open ended responses

The open-ended responses indicate that participants gained a broader understanding of AI beyond technical implementation. Several learners highlighted that the course helped them connect AI with business value, decision-making, and real-world applications, rather than viewing it solely as a coding or technical tool. Ethical considerations, AI in business models, and practical applications in various contexts were also noted as important new insights.

Participants appreciated concrete examples, such as the business model canvas and OsthroBoats case. Several respondents emphasized that the course helped them see AI as a practical and positive innovation that can increase efficiency, reduce manual work, and contribute to sustainability, rather than something to be feared. Overall, the feedback suggests that the course effectively encouraged learners to reflect on the strategic and operational potential of AI in business contexts.

Participants indicated that the course is likely to have a positive impact on their work, business decisions, and professional development. Several respondents noted that the course helped them understand how AI can simplify tasks, improve accuracy, and support smarter, data-driven decision-making. Learners reported increased confidence in identifying where AI can add value and enhance business processes, reducing stress in daily work.

Many participants emphasized that the course encouraged a shift from passive to active AI use, inspiring them to seek new applications of AI in their work or personal projects. Some highlighted that the knowledge gained could be applied in international contexts, supporting innovation and best practices beyond their current workplace. Others mentioned that the course contributes to long-term professional development by strengthening digital skills and providing insights applicable to further studies or thesis work.