Experiments with generative AI in public administration

Report

ICT in public administration

PUBLICATIONS OF THE MINISTRY OF FINANCE - 2024:48

VALTIOVARAINMINISTERIÖ
FINANSMINISTERIET

Publications of the Ministry of Finance 2024:48

Experiments with generative AI in public administration

Report

Anni Penttilä, Petra Holkko, Niko Mäkilä and Henriikka Tammes-Peters

Ministry of Finance Helsinki 2024

Distribution of publications
Distribution of publications

Government Publications Archives Valto

Publication archives Valto

publications.valtioneuvosto.fi

ISBN pdf: 978-952-367-840-8
ISSN pdf: 1797-9714

Printed by the Government Administrative Unit, Publication

Production Helsinki 2024

Description page
20.9.2024

Experiments with generative AI in public administration Report

Publications of the Ministry of Finance 2024:48		**Subject:**	ICT in public administration
Published by	Ministry of Finance		

Author/s	Anni Penttilä, Petra Holkko, Niko Mäkilä, Henriikka Tammes-Peters		
Language	Finnish	**Number of pages**	64

Summary

This report compares domestic AI experiments in public administration and reviews the international landscape of generative AI. From domestic data on experiments in generative AI were collected through an interview survey. The aim was to identify the current state of play regarding the application of generative AI and its potential to support the work of the government.

The report also aimed to gather lessons learned from the government's generative AI experiments and thus promote cooperation between ministries on generative AI. The sample of the survey is limited to public sector AI experiments launched in 2023 and 2024.

The experiments led to a better understanding of the possibilities and limitations of generative AI. Based on the study, we were able to identify issues that public administrations still need to work on in order to enable sustainable and responsible adoption of generative AI.

The final section of the report contains a selection of AI policies and experiments published by different countries to illustrate global trends. The sample includes the UK, the Netherlands, Germany, France, India, Canada and China. The international landscape is changing rapidly as the adoption of generative AI applications continues to grow in the public sector and government functions.

Keywords	artificial intelligence, public administration, administration, experimentation, public administration ICT		
ISBN PDF	978-952-367-840-8	**ISSN PDF**	1797-9714
Publication address	https://urn.fi/URN:ISBN:978-952-367-840-8		

Presentationblad
20.9.2024

Research on generative artificial intelligence in public administration Report

Finansministeriets publikationer 2024:48		**Her**	Public administration ICT
Utgivare	Finansministeriet		
Authors	Anni Penttilä, Petra Holkko, Niko Mäkilä, Henriikka Tammes-Peters		
Language	finska	**Sidantal**	64

Unit

The report includes a comparison of domestic efforts with artificial intelligence in public administration and an overview of the international situation with regard to generative artificial intelligence. Information on domestic attempts with generative AI was gathered through an interview survey. The aim was to map the current capacity to implement generative AI and the possibilities to support the work within the government.

The purpose of the report is also to gather lessons learned from generative AI projects in public administration and thus to promote cooperation between ministries in this field. The scope of the mapping exercise is small and focuses only on the public sector's proposals for 2023 and 2024.

Tack vare försöken ökade förståsen för möjligheterna och begränsningarna med generativ AI. From the research work, we can put our finger on the areas that need further action within the public administration to enable a sustainable and responsible use of generative AI.

In order to provide an overview of global trends, the second part of the report deals with the efforts of different countries to implement and guide the use of AI. Here are some examples
from Storbritannien, Nederländerna, Tyskland, Frankrike, India, Canada and China. The international situation is changing rapidly, with the use of generative AI programmes increasing throughout the year in the public sector and in public administration.

Nyckelord artificial intelligence, public administration, public administration, public administration ICT, public administration, proposals

ISBN PDF	978-952-367-840-8	**ISSN PDF**	1797-9714
URN address	https://urn.fi/URN:ISBN:978-952-367-840-8		

Description sheet
20 September 2024

Experiments using generative AI in public administration
Report

Publications of the Ministry of Finance 2024:48		**Subject**	Public Sector ICT
Publisher	Ministry of Finance		
Author(s)	Anni Penttilä, Petra Holkko, Niko Mäkilä, Henriikka Tammes-Peters		
Language	English	**Pages**	64

Abstract

This report compares experiments using artificial intelligence in public administration and provides an overview of the international landscape of generative AI. Information on domestic experiments with generative AI were collected through interviews. The aim was to map out current capabilities to apply generative AI and opportunities to use generative AI to support work in central government.

This report also aims to promote cooperation between ministries with respect to generative AI by collecting information on lessons learned through experiments using generative AI in public administration. The survey's sample is narrow, being limited to AI experiments launched in public administration in 2023 and 2024.

The experiments have increased understanding of the possibilities and limitations of generative AI. Through the survey, we were able to identify matters that require further work in public administration in order to enable the use of generative AI in a sustainable and responsible manner.

A collection of AI policies and experiments published by various countries has been included at the end of the report to help outline global trends. The sample includes the UK, the Netherlands, Germany, France, India, Canada and China. The international landscape is changing rapidly as generative AI applications are being increasingly deployed in the public sector and in administrative tasks.

Keywords	artificial intelligence, public administration, administration, experiment, public sector ICT		
ISBN PDF	978-952-367-840-8	**ISSN PDF**	1797-9714
URN address	https://urn.fi/URN:ISBN:978-952-367-840-8		

Content

1 Introduction

Generative or creative AI is described as a disruptive technology with significant societal implications. Generative AI can transform public administration by automating mechanical, repetitive tasks, streamlining decision-making and improving citizen services. Alongside the potential benefits of generative AI, the risks associated with cybersecurity, privacy protection and the proliferation of misinformation also affect the adoption of the technology in government.

Governments around the world are experimenting with generative AI in an effort to make government work more efficiently across different sectors. The greatest potential for applications of generative AI is seen in tasks requiring natural language understanding and production. In Finland, too, there is a complex effort to integrate generative AI in public administration.

For this report, we have gathered information on domestic government AI experiments and created an overview of the international landscape of generative AI by looking at AI policies and experiments in different countries. The aim has been to map the current state of the art of generative AI application capabilities and the potential of the technology to support work in government. The sample of the mapping is limited to the public sector AI experiments launched in 2023 and 2024 with respect to domestic experiments.

Information on domestic GenAI trials was collected through interviews and other means. This included use cases, source material, language models, successes, lessons learned, challenges, and attention to security regulation and ethical considerations. The final part of the report contains extracts from AI policies published by different countries. The section also describes some of the public sector experiments in generative AI in different countries to illustrate global trends.

The report looks at a limited sample of domestic government and public sector AI experiments. For international comparison, a small sample of countries where the development and management of generative AI is being promoted at the whole-of-government level was selected. The field of generative AI is dynamic and rapidly changing. So best practices are in a constant state of flux as the use of generative AI in public administration becomes more common in Finland and around the world.

2 Generative AI tools in general

Generative or creative AI is a broad concept that describes AI technology that can create new text, images, video, sound or code. Large Language Models (LLMs) are AI solutions based on textual learning material that perceive the world through textual interpretation using natural language processing (NLP), which uses advanced machine learning algorithms to analyse natural language and identify grammatical structures, meanings and contexts. Large-scale language models based on NLP technology are used in applications of generative artificial intelligence: they can be used to produce
produce text output based on training data and source material on a probabilistic basis. (Heikniemi, 2023)

OpenAI's ChatGPT and Microsoft Copilot are publicly available online versions of generative AI applications built on top of large language models. Users can communicate in natural language with chatbots and ask the tool to generate text output based on the desired input. Chatbots can be used to simplify many of the tasks that government officials work on every day. For example, like search engines, chatbots can assist in learning about a new topic, drafting or outlining a presentation, summarising a long article or writing a short piece of code for a desired function.

In the case of text generation, these tools use user input to calculate the statistical probability of a word sequence - in other words, each character in the response is based on the system's prediction - with highly variable results. Generative AI applications are trained on huge amounts of data, often collected from the internet. However, there are ethical issues related to the source and use of training data for open language models, including transparency of training data and copyright issues. AI models are said to be only as good as their training data: the information produced by the tools may be based on outdated data, misleading, discriminatory or incorrect. It is important to note that language models are always politically biased in some way: objectivity cannot be assumed (CNAI, 2024).

Users can control applications of generative AI by refining their prompts. The closest to the desired result is when the user's input is well contextualised and contains detailed and relevant information for generating a response. The user is also responsible for what information to feed into the models. Users' public
The prompts given to GenAI applications (such as ChatGPT) can be used for further training of AI systems, i.e. the history of conversations with a bot can be incorporated into other conversations later on. To ensure the safe and ethical development and use of generative AI tools, it is possible to incorporate data management and security protocols on the software and server side, but the responsibility for maintaining privacy and data protection also lies with the user. In particular, a public authority using AI applications in a public administration should be familiar with the principles of the technology it is using and its own responsibilities as a user of the tool.

3 Experiments with generative AI in government 2023-2024

At its best, generative AI can offer new opportunities for improving the efficiency of public sector work and services. The next section of the report presents five recent generative AI trials in different ministries. The sample has been extended to agency and city level for a few of the experiments.
Experiments have used AI to improve the efficiency of law drafting and the processing and analysis of large datasets, for example. The introduction of GenAI in the public sector aims to streamline administrative processes, improve data-driven decision-making and create a flexible citizen-centred service.

However, the development and deployment of generative AI tools involves risks and challenges. Effective and appropriate use of the technology requires an effective information infrastructure, close collaboration between stakeholders and a renewal of the work community's processes. Culture change and user training should also be a priority if generative AI applications are to be deployed in production within an organisation. In the context of this report, the user is primarily the public authority that uses AI tools to support its work.

Table 1. Domestic generative AI experiments in the mapping exercise

The experiment	Responsible organisation	Use case
GoodAI	STM, SM, VM	Creating summaries and comparative analysis
Creative AI to support law-making	LVM	Unpacking the data and searching for commonalities
Generative AI in the consultation process	CCNR, (OM)	Production of a draft summary of the opinions issued in the consultation
Archive Assistant	UK	Create summaries and generated responses based on the data
Hilma's procurement manual assistant	VM, Hansel	Decompilation of the contents of the Procurement Manual and analysis of the data
Support for handling reports of threats and violence	City of Helsinki	Analysing violence incident reports and other related data - from the perspective of the TSA's mission and knowledge management
Helsinki GPT to help with customer feedback	City of Helsinki	Automatic classification and prioritisation of customer feedback and generation of response templates
Increasing customer focus by improving the handling of open feedback	City of Helsinki	Summaries, theming, sentiment analysis and translation work
Increasing customer focus by improving the handling of open feedback	City of Helsinki	Summaries, theming, sentiment analysis and translation work
DVV's Digiturva Chat	DVV	Chatbot to answer security questions - answers are generated based on public digital security guidelines
Specific HR issues	City of Tampere	A chatbot that answers employees' questions securely based on textual data

3.1 Descriptions of AI experiments and key lessons learned and successes

GoodAI

Generative AI was used in the analysis of change programmes in welfare areas, with the aim of creating summaries and comparative analyses based on source data. The experiment created a chat-based graphical user interface.
The source material was only public material: the programmes for the transformation of the welfare regions and the minutes of the decisions of the regional councils' boards and committees. The confidential parts of the service level decisions and the reports on the economic and social situation of the welfare regions were carefully examined to avoid them becoming part of the source material. The source material was in a variety of formats: PDFs of text, PowerPoint presentations with images and tables.

The language model itself was not trained, but the project was implemented as an interface solution. The previous OpenAI language models (3.5 Turbo and 4.0) served as the basis for the solution: they were populated with their own source material and the RAG method was used for data retrieval. The interface included feedback and source reference functions.
The tool provided references to the text it produced, which were intended to indicate the reference books from which the information was quoted. Management was actively involved at the beginning of the experiment. However, interest waned as the trial realised how far away the tool was from becoming an effective daily tool.

The Retrieval-Augmented Generation (RAG) method was recommended for the development of services. RAG allows the dynamic incorporation of relevant data and documents in a language template-based system, allowing the creation of highly customised AI solutions. The method combines the information retrieval component with a text-generating model. Relevant documents and sources are retrieved based on user prompts. The documents are contextualised based on the user's initial input and fed into a text generator that produces the final output. The RAG method is therefore also the basis for the source function implemented in the tool (Choi, 2024).

Lessons learned and successes: work needs to be done on the existing information architecture, semantic interoperability and data mobility. More time should be set aside at the outset to define the use case, for example through self-discovery and study.

The choice of approach should be proactive rather than going in for a trial-and-error approach. Attention should also be paid to how to create a culture change if new AI technologies are introduced at organisational level: many people rely on the old way of doing things.

Transparency and traceability aspects are important to take into account when outsourcing technical implementation. The co-existence of different AI technologies/applications, such as HyvAI and Copilot, will provide better results in analytical work.

Creative AI to support law-making

Generative AI was used in the preliminary stage of the drafting process to extract the draft legislation; the tool developed in the project was designed to provide interfaces with existing laws. In the experiment, Finnish open language models, Finnish GPT and Reindeer were further trained with legislative texts and related materials. More specifically, the training material was semantic Finlex, the data provided by the Finnish Legal Data Open Linked Data Service.

The challenge in terms of source material was the limitations of using up-to-date legislative material. The experiment therefore had to rely on partly outdated data. The language models were fine-tuned using and relying on data published by Stanford University. A chat-based solution was created for the user interface, with both search results and the source function for the generated text visible. A limited number of testers had access to test the interface via an open link. No specific security measures were implemented, based on the open source material and the open interfaces on which the experiment relied.

Lessons learned and successes: increased understanding of what it takes and requires to build similar services. Increased understanding of the need for better understanding of what it takes to deliver these services.

Generative AI in the consultation process

The experiment sought to learn lessons about the suitability of the common language model as a tool for producing a summary of opinions as part of legislative drafting. The experiment was triggered by the reform of the opinion service, where AI was seen as an opportunity to increase efficiency. The aim of the experiment was also to reflect on strategic autonomy and to explore where language models in our linguistic area could be applied. The use case for the experiment was the production of a draft draft opinion summary based on the comments received in the consultation round. This involved creating summaries, assessing support/opposition and classifying comments.

The aim was also to draw conclusions from the analysis produced by the tool. This objective could not be achieved in practice, as significant shortcomings were identified during the pilot phase: t h e tool produced incorrect content and was dominated by unsatisfactory quality of the source material. The tool did not produce references to the source material and the limitations of the context window of the Poro language model and the limited capabilities of the model were a further challenge. A different user interface was not created in the trial. A generic language model and the lack of a user interface were found to be an unsuitable solution.

Lessons learned and successes: lessons learned from the use of the common language model in the summary of opinions. It was possible to produce a structurally and linguistically satisfactory draft. The demo solution shows potential to facilitate the work of the drafters, but it needs to be
it is human guidance and control. Regular weekly meetings with the editor were seen as useful. Particular attention should be paid to improving the quality of the consultation system, enriching the language model, a wider context window and data harmonisation.

A general language model cannot produce a sufficiently qualitative result; to achieve the desired results, it is necessary to enrich it with specific data. The experiment has led to the realisation that the extraction of productivity gains requires the development of a baseline system and a coherent approach. The need for a language model enriched with legal knowledge was seen as obvious.

Archive Assistant

Generative AI was used to find relevant information from source data, in an effort to efficiently build bounded overviews of large data sets for different needs. The basic use case was therefore to generate summaries and answers to more specific questions based on the data. The experiment also sought to learn lessons on how to integrate own data into the GPT model.

The source material was the complete archives of the Government's publications and the semantic Finlex. The data was public, the experiment was limited, and the agenda was to create a support tool for which the official is responsible for the content produced - thus no specific security measures were taken. Further development of the model into production would require significant additional resources (money, time, skills). The working group assessed the need to extend the source material comprehensively to material from different jurisdictions, for example at EU level.

Lessons learnt and successes: potential identified, provides a case study on integrating own data into the GPT model. Technical feasibility achieved, produces a hallucinated response and returns source references. Includes training capability (match/fail) and learns iteratively. Generative answer provides higher quality results. Open and effective supplier collaboration. The experiment sparked ideas for designing a broader tool for government.

Hilma Procurement Manual Assistant

The pilot created a chat assistant to ask for information on the Procurement Manual for public procurement. The tool is based on the vision of Hilma in terms of increasing user-friendliness; the information to be processed is relatively elusive, fragmented and extensive. The aim of the Procurement Manual Assistant is to provide open data on public procurement in an easily usable format and a single gateway to public procurement information services.

The experiment was carried out using a 3rd party productised AI solution, which allows the underlying language model to be easily changed. Microsoft Azure was chosen as the cloud environment, which provides a platform for leveraging OpenAI language models in proprietary AI solutions. The aim was for the Sourcing Manual Assistant to produce answers using only its own source material, the Sourcing Manual, as a source. The dominance of AI has been and continues to be a challenge for the use of the service. In certain situations, it has been possible to lead the assistant to answer questions from outside the limited source material. Efforts continue to be made to overcome these challenges with the partner. In the context of the trial, the user was responsible for assessing the legality of the results and ensuring quality.

The Hilma procurement manual assistant is already in production and available to all people, but according to the organisation responsible, the service is still in the "pilot phase". The aim is to gather feedback and experience on the usability and reliability of the service before expanding the source material and possibly customising the user interface. Consideration has been given to extending the content of the information, for example for procurement law and public procurement advisory services.

Lessons learned and successes: a chatbot was created with potential for production and use of a wider range of sources. Lessons learned and increased understanding in a limited experiment; future potential for increasing the source material and realising a broader vision; replication of the model for other uses was seen as possible. The approach was cost-effective and different from other experiments; low start-up costs were perceived to be effective for the experiments. Read more about the pilot here: Hilma - Press releases.

Support for handling reports of threats and violence

AI was applied to the analysis of threat and violence reports, risk assessment and used in the process of handling the reports. The pilot produced a function to support the work of the frontline staff, whereby the AI produced suggested actions for the incident reported by the employee at hand. The aim of the pilot was to develop a proactive safety culture, leadership and employee experience through AI.

The technical implementation of the experiment again relied on the customization of pre-trained language models: the Pandas AI Python library and OpenAI GPT models were used. A graphical user interface was created for the tool, where an event description is entered and action suggestions are given in response. The interface contained two text fields: one for security incidents and one for general questions. The user is provided with a printout of the relayed text and the answers, together with a code sample to be explained if necessary.

Lessons learned and successes: the experiment resulted in the development of a discussion bot that provides suggestions for action in the aftercare of a situation when a description of the event is given. The key lessons learned were that the quality of the data (content, utility, usefulness) and the question framing are important factors in achieving the desired outcome. The database, i.e. the source data, did not contain sufficient measure descriptions to allow the tool to produce good enough answers and solutions.

Helsinki GPT in handling customer feedback

Artificial intelligence was used to improve the processing of customer feedback, in order to test how well GPT-based applications can help to process feedback and improve customer orientation in city services. The focus was on the automated classification and prioritisation of customer feedback and the generation of response templates.

The experiment created a GPT-based service environment for which AI was further trained using Python's Pandas AI library. The technical environment was based on Microsoft Azure services. The raw data of the experiment's source data contained personal data, but was anonymised and sote data was removed. The format of the source data was structured data in csv files. The importance of data pre-processing (collection, organisation, cleaning) was highlighted in the experiment.

Lessons learned and successes: the experiment shows that customer feedback can be handled efficiently using AI. The pilot will allow for the efficient use of feedback.
also gained new knowledge about how processes work in Azure and increased understanding of training AI. The experiment increased the pilot team's understanding of the GPT-based AI application and provided insights on how the results of the experiment can be applied to other projects in the city organisation. The importance of internal and external collaboration within the team was highlighted.

Increasing customer focus by improving the handling of open feedback

The aim is to develop a configurable tool based on the GPT model to improve the handling of open customer feedback. The aim of the pilot was to find out what statistically analysable information can be obtained from the feedback and whether it can be used to identify customer needs. The pilot was based on the overall themes of public service development and one of its objectives was to ensure more equal treatment of feedback from foreign language customer groups.

The basic uses of the experiment were summaries, theming, sentiment analysis and translating foreign language feedback into Finnish. In the technical implementation of the experiment, Python's Django web application and GPT language models were used through the interfaces provided by Azure OpenAI. The services are running in the City of Helsinki data centre. The pilot has ended, but further development and production are underway.

Lessons learnt and successes: the experiment showed that GPT language models are useful for analysing open feedback: the processing of customer feedback is more efficient, thus increasing customer orientation in employment services and saving working time for other core tasks. The GPT application allows feedback to be thematised, sentiment to be assessed and foreign language feedback to be translated into Finnish. The tool allows for a better quality treatment of feedback in foreign languages. The lessons learned also included a better understanding of the shortcomings of AI technology: they require a critical review of the results.

Considerations for the future: general skills are needed, e.g. for making invitations (developing internal skills). The need was identified for a generic tool using the language model, with which officials could train. The aim is to develop a prototype of the tool at city level and bring it into production.

Digiturva chat

The pilot created a collection of digital security documents consisting of recommendations from the Information Management Board as a chat service, a chat assistant who can be asked questions about security and whose answers are generated from the public digital security guidelines. The user group was limited to Finnish-speaking citizens. The answers were based on 80 public documents.

In the source data pre-processing, vectorisation, and text generation phases, OpenAI language models were used to formulate the response through interfaces. The LangChain language model development tool was used for the implementation.
The project was carried out in-house, without outsourcing the technical implementation. The larger DVV development project, which included four chatbot experiments, focused on processing and classifying textual data using machine learning models.

Specific HR issues

The pilot developed a chat service for city officials to answer specific HR questions, such as how to record working hours and holidays. The solution was based on Azure's OpenAI technology and an interface implementation using the RAG methodology. The tool was built directly into the city's ICT environment, which allows the security issues of the technical implementation of the solution to be taken into account.

Budget and time resources for experiments

Most of the generative AI experiments and pilots in the mapping have been carried out on a small-scale basis, with budgets ranging from €10 000 to €48 000. Most of the experiments have already been completed, but final reports are still expected in late summer or autumn 2024. In terms of time, the average duration of the experiments has been about half a year. In the case of the Hilma pilot, the demo took a few days' work, both on the part of the editor and on the part of your team. User testing, feedback gathering and possible further development activities have taken up most of the time resources in several trials.

3.2 Language options

As shown in Table 2, the experiments were mainly based on public data. In some experiments, confidential data were filtered out of the original training data, such as the HyvAI data, which were kept confidential.
contributions to service level decisions and well-being area surveys. In some of the experiments, the use case required the use of sensitive data, but in these cases the data were protected by encryption. For example, the City of Helsinki's Threat and
during the implementation phase of the Violence Incident Reporting Support Tool, personal data was anonymised from the source data by manually modifying it.

In terms of user interfaces, the most common solution in the trials has been a graphical chat-based interface. Conversational generative AI tools require appropriate prompts to work as intended. Graphical user interfaces (GUI) provide an easy way to adjust parameters and control the AI. Graphical user interfaces can incorporate functional features (such as automatic follow-up questions) that enable conversational interaction with the bot in a dynamic and intuitive way.

The language model choices have been largely based on the OpenAI GPT models. Reasons given include the fact that the international models were seen to make sense for the scope and purpose of the project. The ChatGPT models were often perceived as flexible due to their good chat features. In the field of AI, the activities of the main developers of language models are overshadowed by a significant lack of transparency, which raises a number of ethical problems as the social impact of these AI models continues to grow (CRFM, 2024).

Table 2. Source material, language choices and user interface

Pilot	Nature of data	Language model	Interface
GoodAI	Public material (text, Powerpoint presentations, pictures, tables)	OpenAI's ChatGPT-3.5 Turbo and ChatGPT-4	Graphical (chatbot-like) user interface
Creative AI to support law-making	Open source, semantic Finlex	English GPT -models, Reindeer language model	Graphical user interface
Generative AI in the consultation process	Data from the consultation rounds, own model examples	Reindeer tongue model	No user interface
Archive Assistant	Government Publications Archives in full and semantic Finlex	OpenAI's ChatGPT	Graphical user interface
Hilma's procurement manual assistant	Handbook on public procurement	OpenAI GPT models	Graphical user interface
Support for handling reports of threats and violence	Occupational safety and health records of primary education workers' reports of threats and violence and recorded measures	OpenAI's ChatGPT	Graphical user interface; two text fields - one for security exceptions and one for general questions
Helsinki GPT to help with customer feedback	Limited amount of customer feedback; 60 000 rows of structured data in csv files.	ChatGPT-3.5	Graphical user interface, user can choose if they want a conversational response
Increasing customer focus by improving the handling of open feedback	Open feedback; varies in quality and language	OpenAI GPT language models	Graphical user interface - able to analyse and summarise base data
DVV's Digiturva Chat	Digital security documents compiled by the Data Management Board as PDF files, public domain material	TurkuNLP GPT3, OpenAI's ADA, GPT-3.5 Turbo and Chat Completion model interface	Web-based access interface
Specific HR issues	Internal documents within the organisation	OpenAI language models in Azure	Graphical user interface

Most of the major AI companies' models are closed-source language models, such as OpenAI's GPT models, Anthropic's Claude models, Google's Gemini, etc. Closed source models are private property, which are made publicly available through licensing. As commercial products, they are often user-friendly and easily accessible, but access to the code and architecture of the model is restricted. Proprietary models are often subject to stringent data protection measures due to their commercial use, aimed at protecting the intellectual property rights of developers and compliance with data protection legislation.

Service providers are usually responsible f o r maintaining and updating private models, which means that security audits are carried out on a regular basis. On the other hand, the models are dependent on updates from vendors, which can limit the scope for customisation and further development of the models. Closed models are sometimes compared to black boxes. The challenge in placing them is that the user does not know what is happening inside the model. When the source code cannot be checked, the company may include code that performs functions that are hidden from the user. The client can only rely on the fact that sensitive data, such as that used to train the model, and user input are protected.

In **open source** language models, the source code is publicly available. Anyone can freely explore, modify and share it. The main advantages of open source models are seen in their transparency and free usability: developers can build on top of what already exists, understanding what their own solution is based on. Models are agile and can be customised and fine-tuned to meet specific needs. Security auditing of the models is done externally: the developer community can detect and (if desired) fix security flaws. The choice of open source requires specific expertise and a dedicated team. In the AI community, advocates of ethical AI argue for open source models. (Davis, 2023) When evaluating security when choosing a language model, the advantages and disadvantages of open and closed models must be evaluated on a case-by-case basis.

There are big differences between language models in **the context windows**, i.e. the amount of textual information that the AI can take into account when processing a single input. The context window is an essential variable in determining a model's ability to produce coherent and contextual responses or analyses. When the number of characters in a conversation exceeds the context window, the model loses access to the earliest stages of the history of the conversation. The size of context windows is measured in tokens (words or parts of words). A context window of 16,000 tokens corresponds to the capacity to handle about 20 pages of text with a single prompt. OpenAI's GPT-3.5 Turbo has 16,000 and GPT-4 has 32,000 or 128,000 token windows, depending on the model. Meta's Llama and Anthropic's Claude models have 4,096 and 100,000 or 200,000 token windows respectively.

The larger the context window, the more computing power is required to run the model. A very large context window, in addition to its advantages, can also overload the language model: the probability of the model getting lost in irrelevant information or losing essential information is higher as the context window increases.
In addition, the performance of the model can slow down as the computational power requirement increases (Aryani, 2023; Rahul, 2024).

In the trials, comparisons with Finnish language models were not often made at the supplier tendering stage, but in many cases the choice of language model came as part of the package deal that was procured from the supplier. In the context of **strategic autonomy**, i.e. generative AI, the aspect of developing domestic AI capabilities emerged in the justification of language model choices in a few of the experiments.

3.3 Challenges and problems

Challenges of source material

Problems with source data were one of the recurring challenges, and one of the first to be addressed in a number of experiments. The quality of the training data or source material was found to have a significant impact on the results of the experiments. Not all data is ready for use in generative AI: the applicability of AI technology depends on the quality and quantity of data available. To produce relevant and meaningful results, high quality data in a usable format is needed. Inadequate, outdated or otherwise poor quality source/training data is more likely to lead to inaccuracies and erroneous outputs. Source data needs to be curated, transformed and organised in a structured format so that language models can be effectively trained or pre-trained models fine-tuned to perform as desired. The suitability of ambiguous source material as training data needs to be assessed on a case-by-case basis.

Data security and privacy issues concern the selection and handling of source material and the handling of content generated by the models. In the research experiments, the source material consisted mainly of public data, exploited in a closed environment. In some of the City of Helsinki experiments, the source material contained personal data that required anonymisation. The encryption of personal data was done manually, as there was no tool developed for anonymisation. Manual work was not considered sustainable in the long term. In the limited time window of the trial, manual work was not a priority. In this case it was felt that in order to streamline the experiment it would be good to have a ready-made environment that meets data protection requirements or, alternatively, to deal with different, non-confidential material in a less regulated environment.

When processing personal data in AI systems, systems should map all flows of personal data during the development, testing and deployment phases to minimise risks, and reduce the risk of personal data identifiability by processing it using various privacy-enhancing techniques. Pseudonymisation means processing personal data in such a way that there is no risk of identifiability without additional data. The processing of pseudonymised data should still be subject to data protection rules. Anonymisation is the processing of personal data in such a way that the individual can no longer be identified. Anonymised data are no longer considered personal data and are therefore not subject to data protection rules (Office of the Data Protection Ombudsman, n.d.).

Most of the experiments were carried out **as interface solutions**, where large language models were tailored to suit the specific use case. This is a custom language model solution, where the existing language model is customised with its own source material to improve the performance of the model for a particular use case. In several experiments, this process relied on Azure OpenAI services, a collaboration between Azure and OpenAI that provides developers with tools for AI development.
the use of intelligence and machine learning in the development of applications. Other language model application development tools, such as LangChain, were also used. In order to prevent the storage or retention of information provided by **the user as prompts**, some of the experiments chose a development environment that included privacy controls to prevent the storage and retention of user information.

Validating and sanitising source data and the data used in prompts is one way to prevent security attacks. In the context of Generative Artificial Intelligence (GenAI), data sanitisation is the process of removing or modifying sensitive data from source data in an orderly manner before it is used to build AI models. Sanitisation can also be referred to as user input processing.
in the context of removing potentially harmful content from the body, which is also a critical action to prevent various security breaches (Haltu, 2023).

Intellectual property issues specifically concern the selection and treatment of source material, and the further use and distribution of content produced from source material. Data used to train AI systems is often protected by copyright. The copyright of creative professionals is threatened when works and content available on the Internet are exploited by AI systems.

as educational data. There is an ongoing debate about the IPR infringements associated with the mining, downloading and processing of copyrighted data and its use in the training and output of AI models.

Another topical issue concerns the ownership and copyright of AI-generated material. The question of AI as a copyright holder is a globally challenging issue. In the US and the EU, it is thought that copyright protection can only be granted to a natural person. However, the use of AI tools as an auxiliary medium is not in itself a barrier to copyright protection: the registration of AI-assisted content is subject to case-by-case discretion. Discussions about copyright ownership have also extended to the rights of the developer and owner of the AI system. (Karus, 2024) Traditional intellectual property laws are challenging to apply to AI-generated content. Public sector actors need to define how openly AI-generated content is to be distributed, or whether it should be protected by special licences, for example in the case of a model trained on classified material. For more information on the IPR issues and risk management of generative AI, see the World Intellectual Property Rights report.

Governance: in the trials, the main challenges related to the reliability of the tools concerned the governance of the models and the distribution of responsibilities for quality control. AI models sometimes produced false, misleading or completely fictitious results. Systematic manipulation occurred in almost all experiments. AI researchers are divided on whether they see dominance as a fundamental challenge for applications of genetic AI. However, there is research evidence of increasingly effective methods and algorithms for tracking hallucinations (Farquhar et al., 2024). The better the instructional material matches its intended use, the less the model hallucinates. In addition, effective frame design can reduce hallucination and the harm it causes.

Prompt engineering is the so-called interaction with AI to guide applications of generative AI to produce high-quality and relevant results. A prompt for a text-based AI model can take the form of a question or command, or a longer request to feed the model with relevant additional material, instructions or a reference to a discussion list. In practice, framework design is the strategic use of language to formulate clear prompts, questions and commands to guide or train the AI or conversational bot to perform in the most optimal way. There are numerous practical guides to body design.

Quality control: Since the results produced by generative AI are not deterministic but probabilistic, the teams maintaining the AI tools must ensure continuous quality control and management of the systems. A reliable GenAI tool will provide high quality answers and clear explanations of how it reached these conclusions. When AI is used to support decision making, its explanatory power is crucial for decision makers, who need to understand the rationale behind the recommendations.

In some of the experiments, transparency and traceability had been achieved by having the AI bot also provide search results and source references when responding. In most of the experiments, quality control issues were left to a later stage, based on the scale of the solutions and the demo phase. However, in the interviews we found that there was sometimes little open discussion with the suppliers about the technical implementation of the trials and the explanatory power of the tools. The lack of visibility into the operation of machine learning algorithms and the unpredictable results are a challenge to maintaining human control. They also make it difficult to allocate responsibility and liability for errors made by AI tools.

The challenge of technology integration

It is important to consider whether generative AI offers a standalone solution in itself, or whether it was more sensible to integrate it with other technologies such as vector databases, semantic search and other natural language processing tools to produce higher quality results. Seamless integration will improve the performance of the tool and the quality of the results. Complex and time-consuming acquisition and implementation processes can be streamlined by proactively identifying the requirements for AI integration.

The experiments also identified the need to make more effective and planned use of existing tools and other (AI) technologies that have been sidelined by the rise of generative AI. Generative AI should only be introduced when it is seen to offer concrete benefits over and above other tools. Table 3 illustrates the suitability of generative AI for different use cases. Taking sufficient time to carefully define the use case and assess the suitability of generative AI is a key element in the early stages of successful AI experiments.

Table 3. Use case family and relative generative model utility

Family of sites	Usability of the generative model (for now)	Case studies
Forecasting	Weak	Risk forecasting, sales/demand forecasting
Design	Weak	Operational research, optimisation, route planning
Decision making, decision enquiry	Weak	Decision support, automation
Autonomous systems	Weak	self-driving cars, advanced robotics, drones
Segmentation/ classification	Reasonable	Clustering, customer segmentation, object classification
Referral systems	Reasonable	Personalised recommendations/guidelines, NBA (next best action)
Observation	Reasonable	Object recognition, speech recognition
Intelligent automation	Reasonable	Intelligent document processing, OCR (object character recognition), hyper automation
Anomaly identification/measurement	Reasonable	Detection of abnormal transactions, outlier detection, measurement
Content extraction	Good for	Text generation, image and video creation, synthetic data
Conversational access-interfaces	Good for	Virtual assistants, chatbots, 'digital' workers
Finding information	Good for	Data storage, retrieval, extraction

Source: Gartner, 2024

The challenges of basic information architecture

Data fragmentation and lack of interoperability of systems challenge the management of data within and between organisations. Semantic interoperability refers to the ability of information systems to communicate meaning, i.e. to understand and exchange data in an unambiguous way. Data must have a common standard and protocol, such as a common syntax and data format (e.g. as in XML). Data from multiple databases can be used to train generative AI, which can lead to semantic inconsistencies if the data models are not compatible across systems. This situation requires extracting, collecting and transforming the data so that it can be used in a new system.

The training and maintenance of AI tools requires smooth data mobility within the organisation's data infrastructure. However, data mobility can be hampered by privacy concerns, data silos and regulatory restrictions. Anonymisation techniques and secure data sharing policies can facilitate data mobility. Public administrations often deal with fragmented data stored in different systems. The challenge is to bridge silos and enable seamless data transfer. Ensuring data security and compliance with data protection regulations (including GDPR) when moving data between different databases and cloud environments is a key challenge that the basic information architecture must be able to address.

Scaling up generative AI may therefore require investments in upgrading data management platforms and modernising legacy technology environments. A well-organised and common data platform in public administrations can gradually improve data mobility, facilitating the transition from few and limited use cases and pilots to the integration of generative AI at the organisational level. Enabling the use of siloed data, improving data management, adding AI capabilities to the data management platform itself, and integrating software and hardware capabilities associated with AI systems are all issues related to scaling generative AI, which will increase the cost of bringing new AI tools into production (BCG Global, 2024).

Cloud computing solutions are key to data management. Cloud computing provides computing resources such as remote servers, storage, databases, networking and analytics as a service over the internet; companies using cloud computing do not need to manage the underlying hardware infrastructure running the applications. The most prominent cloud providers include Amazon Web Services, Google Cloud and Microsoft Azure. Cloud services are centrally managed, making them easier to maintain and update. The rapid development of GenAI systems has created a huge amount of new data to the internet, posing new challenges to existing computing and communication frameworks. The ever-increasing volume of data transfers and access requests means that cloud computing is increasingly prone to delays. For low latency use cases, it is worth considering other computing approaches to GenAI (Ashtari, 2022).

---

edgecomputing is a type of distributed computing architecture where data processing and analysis takes place closer to the data source, at the 'edge of the network' rather than in a central data centre. In edge computing, data is processed locally, reducing the need for data transfer and resulting in lower latency and faster response times. Low latency is critical for real-time decision making. Local data processing reduces network traffic and, in terms of data protection, sensitive data is preserved locally.

Currently, GenAI services rely heavily on traditional cloud architecture due to the need for large computing resources. An efficient centralised computing infrastructure is needed to run the huge number of language models and handle user requests. The heavy computation of cloud computing consumes a significant amount of energy. A centralised computing system under heavy demand is unecological and cost inefficient.

The development of edge-cloud computing has gained momentum in recent years with the exponential growth of data-intensive applications. Hybrid edge and cloud computing exploits the powerful computing resources of cloud servers and the fast data management and communication capabilities of edge servers. Combining edge and cloud servers is seen as a potential option for building scalable GenAI systems. (Wang et al., 2023)

Resource challenges and skills

Among the resources allocated to the experiments, the time allocated to the project in particular created challenges for several working groups. In about half of the trials, the time resources were found to be insufficient to adequately train, test and develop the model based on user feedback.

Another challenge that was often raised in the debate was the resources available **for training civil servants**. At the organisational level, changing the culture and paradigms of work requires considerable resources, more than was anticipated in the experiments. The technical and operational aspects of AI tools are no guarantee of productivity gains if staff lack the capacity and willingness to adopt the technology and new ways of working. The organisation's management needs to be involved in setting up new working models. For example, common methods must be developed to identify and deal with incorrect or misleading information. Sufficient time should be allocated for user testing and for sharing and unpacking experiences between the working and testing teams. While the technology is still a work in progress, users are expected to be persistent and accountable, as they will always be responsible for how the AI-generated services are used.

In-house **IT skills** were also mentioned by interviewees as one of the areas for development. The lack of skills slows down the adoption of AI in public administrations and the public sector. Technical expertise is needed when considering the overall architecture of organisations and the development of models from the perspective of scalability and security. Knowledge sharing between different disciplines needs to be promoted. For AI development, governance and regulatory enforcement, it is important to seek closer cooperation between private sector actors, AI developers and legislators. Training, human resources and collaboration with the private sector will help to bridge the knowledge gap.

The mapping exercise revealed that in many of the experiments**, the target setting** was limited to a very short timeframe. The experiments were conducted with limited resources, which was reflected in the final results. Experience and lessons were learnt about the functioning and applicability of generative intelligence, but the quality of the results of the tools was still far from ready for productive application in several experiments. The potential for scaling up applications was, however, identified, but **the transition to production** was seen in several experiments as very distant, if not impossible on the current basis.

A common challenge in the trials was the lack of clear plans to bring the demo solutions to scalability. Generic AI is still in its infancy and the development and maintenance of AI solutions requires specific expertise that may not be sufficiently available in the public sector. Setting clear and realistic targets also requires an understanding of the potential and limitations of the technology. The ability to apply generative AI in experiments has therefore provided an understanding for future AI projects.

The findings of the mapping exercise raise the question: is there a need for more trials in the future, or should good use cases and demo solutions be identified and existing resources be directed towards their further development and production instead of new trials? **Strengthening** in-house **IT expertise** is a better guarantee that solutions can be designed for production use from the outset, taking into account aspects of technology integration, secure software development and application maintenance in addition to the administrative side. Being proactive in identifying the challenges of the transition to production allows projects to identify in advance the areas that are prone to failure, thus providing a better basis for scaling up solutions that have been proven to work.

The role of private IT companies and the challenges of outsourcing

In the interviews about the pilots, the question of the **role of private IT companies in** the development, production and especially maintenance of AI solutions in public administration was also raised. As AI applications are brought into production in the public sector and the aim is to deploy the solution at organisational level, the issues of maintaining the technical implementation become more relevant. What are the challenges of outsourcing technical implementation to private IT companies? If a public service relies on a technical solution that is maintained by a private company, liability issues can be complex.

Firstly, if a private IT company is responsible for the technical maintenance of a public service, it is important to clearly define who is responsible for any errors or system malfunctions. The quality control of the functioning of the applications is then divided between the server side, which is operated by the private operator, and the subscriber, which controls the content and use of the applications. This may require precise contracts and continuous monitoring. Secondly, security issues also come to the fore. Public sector services often handle sensitive data, and the service provider needs to ensure that their security practices are strong enough to protect this data and also interoperable with government information systems.

Dependence on private operators can lead to a situation where the public sector loses some of its own technical expertise and independence. This can challenge future development projects and make the public sector more vulnerable to market changes and business decisions by private companies. For private companies
The services of IT companies can be expensive, and budget constraints can limit the public sector's ability to maintain and update its systems as needed. In addition, the cost of standalone solutions relying on private vendors can become significantly high in a situation where multiple AI tools are ordered from different IT houses, but

they rely on similar basic solutions in many respects. The trials showed that the applications are largely similar, especially at ministerial level, and the variable element is mainly the source material depending on the application environment.

In the case of this survey, the technical implementation was outsourced to private IT companies in nine out of ten trials. In this case, **transparency of the development process of** the product or service commissioned between the commissioner and the supplier must be a primary requirement, since reliable and explainable AI tools are based on transparency throughout the product life cycle. Some of the experiments raised the desire for more open communication between the subscriber and the supplier. Direct and open communication about the development of the model and the technical implementation phases also prevents 'black boxes'.
problem for the end-user. The customer must be kept informed of the solutions on which the technical implementation is based and of the causes of any shortcomings or defects detected in the tools. The experience of the City of Helsinki working group was that the service design could have been used at the beginning of the project to support the alignment of the supplier's and the user's perceptions. In some of the experiments, service design was part of the package deal.

Financial resources and cost estimation challenges

The small **budgets of** the micro-projects limited the scope for experimentation and further development of alternative approaches in some of the experiments. There were few estimates of the ongoing costs of operating and maintaining AI tools, as implementations were carried out with very limited resources for very short-term use to accumulate lessons. For possible future scaling-up and production development, the assessment of the ongoing costs needs to be done accurately to provide at least an indicative estimate of the real cost impact of the systems. However, there are uncertainties in the estimation of costs.

Scaling up generative AI models can lead to unexpectedly **high costs**, making budget management challenging. **There are ongoing costs** associated with the use of tools, search functionality and the generation of responses. Ongoing costs depend, among other things, on the choice of language model, the length of the user's body, and the number of times the tools are used. Thus, only hypothetical cost estimates can be made for the ongoing costs of using the services. In addition, the rapid development of artificial intelligence is constantly influencing the cost of using generative AI tools. The final report of the STM's HyvAI experiment presents indicative cost estimates for the regular use of a discussion bot.

The ability to understand the technical implementation and the ongoing costs of the server side is also required to assess the **costs associated with maintaining** and **operating the** systems. For example, interface solutions using OpenAI language models can insidiously create unexpected costs as the length and detail of conversations increase. For example, in interfaces that store conversation history, a multi-step conversation costs cumulatively, not per round, as the entire conversation is sent through the interface to be processed by the language model on each input.

Costs are also generated by data management: data storage, processing and management costs.

costs are high when it comes to scalable AI systems. Data cleaning and preparation require computing resources and time. Identifying, formatting and removing data bias when needed can also be laborious and costly. In addition, regular training of staff on the effective and appropriate use and management of AI systems increases operational capacity. Various tools and guidelines are available for estimating the costs of AI deployment. The following frameworks provide guidance for assessing ongoing costs: the Granica AI Cost Analysis Guide and the QAT Global report on the economic impact of AI.

If language models are to be trained from scratch, it must be taken into account that the process requires significant computational resources and a considerable amount of electricity. Stable hardware and infrastructure is the basis for LLM development and is often the most important part of the overall budget. Powerful servers and AI chips such as graphical processing units (GPUs) are essential for training LLM models. Since the release of ChatGPT, sales of GPU chips have grown at a phenomenal pace. Supply chains have struggled to keep up with demand, challenging the operations and short-term plans of large technology companies, such as the delayed release of OpenAI's multimodal models (Patel, 2023).

Since the advent of generative AI, there has been a heated debate about what will be the limiting factor for the technology. In the medium to long term, the biggest limitation is thought to be the lack of computing power. The cornerstone of the AI infrastructure is AI data centres, and energy constraints are the main bottleneck to their construction. In addition, Goldman Sachs semiconductor analysts are of the opinion that the shortage of two critical chip components, HBM memory type and silicon-based CoWoS packages, will also become a key limiting factor for the adoption of generative AI in the coming years (Nathan et al., 2024).

Global competition for AI chips has intensified in recent years and tensions between China and the US have risen. The interest rate pandemic and geopolitical tensions have highlighted the fragility of global supply chains for semiconductors. Chips are produced in a complex network where no single country, region or company controls production. China has invested in GPU chips to stockpile them for its own national interest. In addition, China spends as much money each year to support its domestic chip industry as the rest of the world combined. The US has tightened its export controls on AI chips exported to China and is seeking to increase its domestic production (Forbes India, 2024).

Taiwan and Korea are the largest chip manufacturers, TSMC and Samsung Electronics. Most AI chips are designed in the US and manufactured in Asia. Several nation states, such as Saudi Arabia and the UAE, are seeking to procure hundreds of millions of dollars worth of GPU chips to develop their own AI capabilities. Several economies, including Japan and Europe, are seeking to promote their own chip industries through incentive programmes. (Patel, 2023) The EU has identified advanced semiconductors as critical technologies in its economic security framework. The Finnish government is also committed to repatriating funding from the EU chip legislation (Mikkilä, 2024).

3.4 Ethical issues

> ***Ethical issues in AI are an integral part of the development and deployment of the technology. The use of generative AI in public administration requires an impact assessment not only in terms of the cost-benefits and value of the investment, but also in terms of the social and environmental consequences. As AI tools are deployed in public sector offices, different aspects of sustainability should be taken into account throughout the life cycle of the technologies.***

Legislation and regulation

The legal issues of generative AI are not new, even if it is a n e w technology. They include privacy and data protection, sourcing and licensing agreements, intellectual property rights, equality and equity, and human rights. Finnish and EU laws provide, inter alia, for

including the already accepted use of artificial intelligence, digital services, equality and equity, privacy and data protection (GDPR) and data use. In spring 2023, the law on automated decision-making in public administration will also enter into force, but this does not apply to learning systems such as AI.

Intelligent systems that make predictions and act at the individual and societal level can violate international agreements, unintentionally or intentionally. Such agreements include, in particular, UN-level agreements on human rights, non-discrimination and, increasingly, environmental impact (Suomi.fi, 2024).

The safe and responsible use of AI technologies requires compliance with existing legislation and ethical guidelines and agreements. The European Union's AI Act, the first internationally binding AI legislation, entered into force in August 2024. The full implementation of its provisions will take place over the next 36 months, with a series of deadlines for implementation. The AI Regulation will become the main regulatory framework for the requirements and use of AI systems in EU Member States. With the implementation of the Regulation, the Commission expects a more harmonised governance of AI technologies to harmonise the AI market across the European Union (Mikkilä, 2024).

After the breakthrough of ChatGPT and content-generating AI models, the scope of the regulation was extended to include so-called general-purpose AI models. According to the Regulation, a general-purpose AI model is the basis of an AI system that determines how the system processes information for specific tasks, such as language translation or data analysis. For example, language models with a wide range of tasks are considered to be general-purpose AI models. These models can be integrated into different systems or applications.

In several governmental experiments, the basic solution relied on just such a general-purpose AI model, such as OpenAI's GPT models. All general-purpose models are subject to transparency and information requirements. If general-purpose AI systems involve systemic risk, offering them on the EU market is not a viable option.
in these countries, more stringent requirements must be met. The European AI Office, to be set up within the EU Commission, will be responsible for monitoring the compliance of general-purpose AI models. The AI Office will play a key role in the implementation of the AI Regulation by supporting Member State administrations in their tasks of overseeing AI systems in Member States.

If the general-purpose AI model is further developed, such as fine-tuned or tailored with its own source data to better meet a limited use case, the requirements on the system will change. General-purpose AI models can thus be embedded in AI systems whose regulation is based on the original four risk levels of the AI system. The risk level of AI systems determines the requirements to which the system operator/provider is subject to be able to operate in the EU market. At national level, compliance of AI systems is monitored by the competent authorities designated by the Member State (Technology Industry, 2024).

In the context of public administration, legal liability also challenges AI systems. When problems arise, such as the handling of erroneous outputs, the question of ultimate responsibility is challenging, as the system involves multiple actors and dimensions, such as data generated by other parties, technical providers, programmers, algorithms and end-users. The creation and distribution of responsibility requires a link between the official and the decision-making process. To ensure that algorithmic decision-making in public authorities does not endanger citizens' legal security, the operating principles of AI-based systems must be traceable and transparent (Suomi.fi, 2024).

One challenge in promoting regulation is the different approaches of stakeholders in different jurisdictions and legal systems. The AI community does not have the same level of experience with the regulatory implications of technological developments and the implementation of AI-specific rules as regulators. At the same time, legislators and DPAs, despite their deep regulatory experience, may lack the necessary technical expertise to understand how, for example, personal data is used in the design, development and deployment of AI systems. This technical gap may lead to a misdirection of regulatory focus. Balancing risk management and innovation requires cooperation and information sharing between stakeholders.

Data security and privacy

Generative AI systems may process personal data during their training and testing phases, and potentially create outputs containing personal data, references to personal data or sensitive data. When using AI technology, consideration must be given to the protection of personal data, compliance with data protection legislation and minimising the risk of privacy breaches from the outset.

In October 2023, the Global Privacy Assembly (GPA) published a resolution advocating a number of data protection and privacy principles as key elements in the development, operation and deployment of generative AI systems (GPA, 2023), including: lawfulness of data processing, purpose limitation and restriction of use, data minimisation, accuracy, data protection, data protection and privacy. six, transparency, security, privacy by design and by default, the rights of data subjects and accountability.

When personal data is processed in the development or use of AI applications, you must clearly define why the personal data is being processed and then use the data only for that purpose. When using personal data, it is important to know the minimum amount of personal data that will allow a quality implementation. In this context, mechanisms for protecting sensitive data should be ensured and confidential data should be processed in systems only for as long as justified by the purpose of the use.

AI systems must be protected in their ICT environment and comply with security and privacy regulations, such as the General Data Protection Regulation. Regular security auditing and vulnerability testing are essential. Recommended solutions to mitigate identified security risks include DLP *(*Data Loss Prevention) practices to protect sensitive data from accidental or intentional disclosure.

Accountability and quality control

Monitoring the legality and quality of the operation of AI technologies applies to the whole life cycle of applications. At the organisational level, a risk management model must be established, based on which measures to reduce risks and their impact assessment will be carried out. Automated testing and quality control methods can be implemented on the server side to ensure continuous evaluation of the model's performance. Human resources should also be allocated to monitoring the performance of the system and quality control to ensure human control. The user of the AI tool is always responsible for the further use of the content produced by the AI.

Human control must be integrated into automated decision-making systems at all stages of their life cycle, from training to testing and production. In this way, the proactive and ex-post control of algorithmic systems can be strengthened and unforeseen consequences in terms of system transparency can be avoided. Human and generative AI technologies have different strengths and weaknesses in ensuring fair results.

Generative AI cannot use emotional intelligence or an understanding of the wider context. On the other hand, humans have unconscious biases and beliefs that influence cognitive processes. Thus, strong governance structures are needed to enable generative AI, as well as the application and alignment of existing organisational practices, such as risk management, with new technologies.

The trials reviewed report emphasised that it is a support tool, so the person is always in charge. This responsibility ranges from the formulation of the input to the correctness of the response provided by the AI. The AI's response must always be verified, as it is only a draft. The very formulation of the input is a kind of risk.
The trials showed that during the production phase, special attention needs to be paid to these issues by increasing user understanding through training

Transparency and traceability

Transparency and traceability are the cornerstones of trustworthy AI. Trustworthy AI is an umbrella concept that, according to the EU definition, encompasses legality, ethics, technical and social sustainability. In addition to existing data protection legislation, the development of future AI tools should proactively take into account the upcoming AI regulation (AI Regulation).

AI systems need to be transparent and understandable in how they work and how they make decisions. Citizens have the right to know how decisions based on AI are made. This is the concept of Explainable Artificial Intelligence (XAI). Extensive language models can be explained using techniques such as decision trees, rule-based systems and model-agnostic methods that allow the underlying reasoning chains of the models to be unlocked (Zhao et al., 2023). Especially in the transition to production, openness and transparency become more important as algorithmic bias is rapidly propagated in scalable systems.

Users and society need to be aware of how model developers share information about their activities and training data. OpenAI has been criticised for its lack of transparency in the training of its language models. Developers closely guard the underlying training details of all high-performance models, such as GPT-4. In addition, AI companies hold significant non-public information about the capabilities and limitations of their systems, the adequacy of protection protocols, and the risk levels of various harms (Heikkilä, 2023).

The fundamental problem of a lack of transparency in the AI industry is highlighted in the *Foundation Model Transparency Index* report by Stanford University's CRFM research centre (CRFM, 2024).The lack of transparency in the industry is not just about education data: the index scores the transparency of companies' supply chains by one hundred indicator. The 14 key platform developers (OpenAI, Google, Meta, etc.) have been selected for the mapping exercise. On average, enough data is provided to cover 58 out of 100 indicators.

Developers are the least likely to share information on the upstream resources needed to build the founding models (46%), covering areas such as data, manpower, data availability and computation. The wage data for the data work required to develop models, the wider environmental impact of models and measures to reduce copyright infringement were reported by a maximum of four out of fourteen developers. At OpenAI, for example, in addition to data on training data, data on labour, computation and externalities are often closed to the public (Bommasani et al, 2024).

The lack of transparency is being addressed through regulation: for example, the EU AI Regulation and the *AI Foundation Model Transparency Act* proposed by the US are taking steps forward by imposing disclosure requirements. By establishing similar transparency reporting practices for founders, stakeholders (e.g. developers, customers, investors, policy makers) can make informed decisions and seek to build a common understanding. Promoting transparency has been shown to lead to better social outcomes through increased accountability, better science and better policy (Bommasani et al, 2024).

Inclusiveness, neutrality and non-discrimination

Fairness as a concept is embedded in several areas of law and regulation, equality and human rights, data protection, consumer and competition law and public law. Fairness in the context of generative AI means ensuring that the outputs of models are unbiased and do not exacerbate existing social, demographic or cultural differences. Thus, fair AI is inclusive, as unbiased as possible and non-discriminatory.

AI systems are designed, developed and deployed by people who are bound by the constraints and biases of their contexts. This human bias can be reflected in educational data environments, for example in problems of representation of minority groups. The dominance of educational data must take into account the use case, the needs of the users, and aim to minimise algorithmic bias.

In generative AI applications, harmful biases can occur in various formats, text, image, sound and video, which perpetuate stereotypical or unfair treatment based on race, gender, ethnicity or other individual characteristics. Generative AI models are trained with data that reproduce current and past biases in social structures. These biases can occur throughout the life cycle of generative AI, from data generation to the writing of prompts. The complexity of systems and lack of transparency can make it difficult to find accurate information about where and how bias arises and occurs.

The international text-based templates are trained on internet data, i.e. mainly data-rich, major languages (e.g. English, Spanish and Mandarin). Building inclusive digital ecosystems requires that you are able to have more multilingual training data to train language models, in an easily usable format. This can help to bridge the digital language gap. The drive for multilingualism is also linked to the objectives of developing domestic AI capabilities. A language model trained on Finnish data can provide more suitable answers in our own language area than international language models that sometimes make rather clumsy translations.

Environmental impact

AI systems are operated from data centres, which require huge computing capacity. The high power consumption of supercomputers has significant environmental impacts, the assessment of which is important for policy-making. In particular, large-scale language models increase carbon emissions due to their high electricity consumption, both during the training phase and when they are used. For example, compared to a single Google search, a query using ChatGPT consumes almost 10 times as much electricity (Goldman Sachs, 2024).

According to a July report by Goldman Sachs, traditional data centres, or data halls, have expanded to meet growing demand from consumers, service providers and large technology companies such as Google, Amazon, Meta and Microsoft. According to Sachs' conservative baseline scenario, the expansion of traditional machine shops could increase electricity consumption in Europe by around 10-15% in the coming years. In contrast, AI data centres could consume up to ten times as much energy as traditional data centres, especially during their training phase. Sachs estimates that the rapid expansion of AI data centres amidst the increasing adoption of generative AI and the gradual digitalisation of the data centre environment will have a significant impact on the cost of training. acceleration in the pace of global warming could boost Europe's electricity demand by around 40-50% over the next decade (Goldman Sachs, 2024).

data centres need a lot of water to cool their data centres and the manufacturing process of key components such as graphics processing units (GPUs) increases the extraction of rare metals (Bashir et al., 2024).

While awareness of the environmental impacts of AI development is growing, the accurate and timely externalities of AI technologies are challenging to measure and model. This is partly because most companies developing AI do not openly share energy consumption statistics or other data on externalities. In addition, the environmental footprint of generative AI comes from multiple sources.

In the case of large language models such as ChatGPT, the environmental footprint includes first of all the training of huge foundational models, which provide the basis for many products and solutions. Secondly, the maintenance and use of models and services has its own environmental impact. As users or deployers, we have less opportunity to influence the first phase of the life cycle of AI technologies. However, we do make an active choice as to whether to use generative AI or another AI technology and, if so, which model or product to choose.

From an environmental point of view, it rarely makes sense to train your own model if suitable pre-trained models are available. The environmental impact, operating costs and cost benefits of using pre-trained models should be balanced when selecting the appropriate model size. Generative AI applications in general consume significant energy and are expensive to train, and should not be deployed for tasks where other machine learning methods can be used.

The responsible development of generative AI requires not only efficiency improvements, but also a shift to new frameworks for assessing benefits and costs. Since applications of generative AI are significantly energy consuming and expensive to operate, a comprehensive impact assessment should precede the selection of technologies, vendors and language models. MIT's The Climate and Sustainability Implications of Generative AI report and the publications Evaluating the Social Impact of Generative AI Systems in Systems and Society and Evaluating Large Language Models Trained on Code provide perspectives on the wider impact assessment beyond environmental impacts.

The development of AI capabilities can be carried out in a more fair and open way: the Finnish supercomputer LUMI is internationally recognised for its ecological and energy-efficient design. The LUMI supercomputer has been used, for example, to compute the largest purely Finnish language model, FinGPT3. The supercomputer, located in Kajaani, is powered by entirely carbon-neutral hydropower. Unlike OpenAI's GPT models, TurkuNLP's FinGPT3 is fully open and accessible to all (CSC, 2023).

3.5 Developing domestic AI capabilities

The government's experiments Creative AI in support of law-making, and Generative AI in the consultation process were carried out using domestic language models. In the Sitra-funded experiments, one of the priorities was to take strategic autonomy into account. The debate on the national development of large-scale language models (LLMs) will be an issue if the aim is to strengthen national and European AI capabilities and reduce external dependencies on generative AI technologies. Strategic autonomy in Europe is currently challenged by a strong dependence on services provided by US technology giants, the so-called gatekeeper companies. This dependence on these players creates uncertainty for the future prospects of the European data economy. For example, future changes in licensing fees and service conditions can only be predicted.

A report published by Sitra in January 2024 makes recommendations on how to strengthen strategic autonomy in the context of artificial intelligence. The ability to act autonomously in the development of generative AI is a goal that requires competitive and reliable European large-scale language models that can be used to develop
new, sector-specific applications and services in a self-sufficient way. The development of European language models in small, less dominant European languages is also seen as a step towards a more inclusive and robust European AI system. The Sitra report also points to the need to develop the high performance computing capabilities and skills required to develop foundational models such as large-scale language models (Sitra, 2024).

The development of domestic AI capabilities also involves other aspects than strengthening self-sufficiency and national and sovereignty. Developing technology in partnership with local communities, businesses and administrations will ensure that solutions support local needs and employment. The development and selection of national or European language models and AI capabilities can reduce dependence on large international technology giants on a small scale and prevent further consolidation of the monopoly position of global AI giants.

4 Comparison of AI policies at international level

I n order to provide a broader picture of the GenAI debate and landscape, this report has also mapped the content and promise of AI policies and experiments in generative AI at the international level. The identification of new innovative approaches and use cases can broaden the perspective on what kind of projects it is possible and reasonable to undertake in Finland with existing resources and information systems. In this chapter, we present the strategic priorities and policies for AI in different countries and also touch on generative AI use cases in the public sector.

Since 2017, a number of national AI strategies have been published across the country to promote the development and adoption of AI at national level. The strategies focus on different aspects of AI policy and reflect different geopolitical goals and values. Priorities include scientific research, skills development, education, AI adoption in the public and private sectors, ethics and inclusiveness, standards and regulation, data and digital infrastructure (Dutton, 2018).

At the level of public administration, a strong international trend is the proliferation of AI bots that collect information and enhance information retrieval to support the work of civil servants and help citizens. Limited information is available on the deployment of generative AI around the world, but based on the analysis in this report, the applications appear to be similar to chatbot variations, as far as generative AI has been applied to ICT environments. More advanced integration is certainly already happening in many countries and the landscape is constantly changing. The sample in this list provides some examples from around the world.

United Kingdom

The UK's National AI Strategy 2021 sets out a ten-year vision. The strategy emphasises long-term action from a science and forward-looking perspective, and targeting the benefits of AI technology across all sectors. The strategy encourages investment and innovation, but also recognises aspects of security and fundamental values. The strategy stresses the need for an innovation-friendly regulatory and governance framework to protect citizens.

In January 2024, the UK government published a comprehensive report on generative AI. framework for AI in government (GOV.UK, 2024b), which sets out 10 principles for the use of GenAI, and provides a framework for how to start building solutions based on generative AI technology in a safe, secure and efficient way. These guidelines encourage the use of generative AI, but also identify risks such as misinformation, hallucinations and data privacy. The guidelines will be updated every six months. The guidelines for the use of generative AI are: avoid it in automated decision making, in high-risk cases, in contexts requiring high explanatory power and accuracy, and in contexts of limited data.

The UK's Algorithmic Transparency Recording Standard (ARTS) helps public sector organisations to proactively and transparently provide information about the algorithmic tools they use and the reasons for their use. The standard requires that information about algorithmic tools and the decisions they assist in making is transparent, understandable, easily accessible and free (OECD.AI, 2021).

One example of a use case is **the "red box" AI tools** for improving the efficiency of ministries. By combining ChatGPT versions with securely self-hosted open source AI models, tools have been created to facilitate administrative work in preparing draft responses to parliamentary questions and requests for information from public authorities. The tool can extract and summarise information from official sources, parliamentary minutes or speeches of the Council of Ministers (Incubator for Artificial Intelligence n.d.).

Netherlands

The Netherlands has developed an administration-wide vision for generative AI in 2024, describing the opportunities and risks of this disruptive but promising new technology. The government wants to create the conditions for the responsible development and use of generative AI, maintaining both digital and strategic autonomy. This will include close monitoring of developments, stakeholder collaboration and policy development (Government of the Nether- lands, 2024).

The policy on generative AI for civil servants contains four principles:

1. the development and deployment of AI is done safely and securely
2. using artificial intelligence for consistency and equality
3. AI technology serves human well-being and protects human autonomy
4. the AI revolution must be sustainable and contribute to shared prosperity.

The Dutch central bank has an internal **ChatDNB, a chatbot** capable of answering questions related to banking supervision. The questions are based on more than 10 000 pages of source material from the central bank's open database of laws, regulations and supervisory information. ChatDNB is implemented by
with the support of Microsoft's Azure OpenAI service. The pilot, using selected source material, received positive feedback from managers and decision-makers. The positive feedback and the interest generated by ChatDNB in the supervised banks encouraged the DNB team to further develop its products (Central Banking, 2024).

The system has safeguards to ensure its reliability and security: references outside the Bank's documents are prevented and references are included in the answers. Further development will also make use of confidential data. It is worth noting that the EU AI legislation that will enter into force will impose requirements on the use of ChatDNB, as the application is in a supervisory risk category.

France

The March 2024 report of the French AI Commission (*Our AI: Our Ambition for France*) highlights the need to invest heavily in European and French AI capabilities in order to overcome decades of European technological
and the downward trend in economic development could be reversed to ensure the continent's prosperity and independence. This structural change will be achieved through the creation of a "France & AI" fund, to which €10 billion will be allocated. (Artificial Intelligence Commission, 2024).

The programme proposes a strategy of six policies:

1. In order to educate the nation and raise awareness, a plan must be launched immediately, taking into account the management of public discourse, the training offered by universities, the integration of AI into social interaction (as an object and as a tool).
2. 10bn. creation of a €1.5 billion AI fund.
3. Setting up data centres in France, increasing computing capacity and promoting and supporting the development of national AI capabilities.
4. Facilitating access to data: abolition of certain prior authorisation procedures for personal data, development of technical infrastructure with European and French cultural data while respecting intellectual property rights.
5. Experimenting with an "AI exemption" for public research: freeing researchers from administrative constraints, increasing fees, doubling resources for public research specialised in AI.
6. Promoting global governance of AI: to create a World AI Organization to assess and monitor AI systems, an International Fund for AI to serve the public interest and a "1% AI" solidarity mechanism for developing countries.

AI is being used to improve public services: to simplify education, personalise patient care, anticipate professional transitions and reduce bureaucracy. The solution is seen as institutional change: public administrations need to strengthen their infrastructure and pilot AI projects. The aim is to accelerate, deepen and scale up the connection to AI across all public services.

Germany

In 2018, Germany was among the first countries to adopt a national AI strategy. The strategy set out to promote growth and competitiveness and the responsible and reliable development of AI. In December 2020, the German Federal Government adopted an updated AI strategy. The update of the strategy can be characterised as a shift towards developing the overall capabilities of the public sector. The updated report focuses on the following policy areas: research, knowledge and expertise, transfer and application, regulatory framework and society. In addition, new initiatives focused on sustainability, environmental/climate protection, managing the interest rate virus pandemic and international/European cooperation, among others (OECD, 2024).

Germany's strong computing capacity supports the research sector's ability to study and develop AI models. With emerging AI and foundational models, AI research is becoming increasingly computationally intensive. Germany has numerous supercomputers to train AI models and meet the growing computational needs of the research and private sectors in AI development. While the widespread availability of computing resources remains to be assessed, computing capacity is a major competitive advantage for Germany.

Use cases of AI in government: in Germany, however, **the level of digitalisation of the public sector lags behind** many other OECD countries, due to a lack of interoperability between government databases. This is partly due to Germany's federal structure, which limits the ability to use public sector data at the national level. Nevertheless, public sector AI initiatives have already been widely publicised at both national and more local level. The Oxford Insights Government AI Readiness Index 2023 ranks Germany eighth out of 193 countries in terms of its ability to integrate AI into public administration for the public good (Oxford Insights, 2023).

Use case: at local level, several cities use LM-based tools to develop public services. The cities of Heidenheim and Heidelberg use chatbots, where citizens and residents can find information in a conversational way. The state of Baden-Württemberg also has a chatbot for citizen tax surveys and as a support tool for civil servants in their daily tasks. The Land of Bavaria also uses chatbots to provide citizens with information on 2 800 different services.

Baden-Württemberg's F13 language model assists civil servants in their daily tasks, freeing up their time and capacity to work on important projects by streamlining their work processes. F13 was designed to condense long texts, streamline the handling of government presentations and invoices, assist in research, create text from human input (users can download documents and they are synthesised into a coherent text, adjustable parameters are length and thematic focus of the text). The tool is characterised by a clearly defined scope: staff are not allowed to use the model for any task, only for predefined sets of tasks. This makes it possible to streamline long administrative processes and to invest in personal interaction with citizens or politicians, for example.

An experiment using natural language processing (NLP) to answer citizens' questions on the interest rate pandemic and the restrictions in place has also been successful. The tool reduced people's dependence on the opening and response times of the relevant authorities' offices. Several ministries worked together to develop this service, which was completed in just 40 days. The pilot can be used as a blueprint for chatbots in other sectors (OECD, 2024).

Canada

Canada was the first country to publish a national AI strategy (Pan-Canadian AI Strategy 2017) and has invested more than $2 billion in AI, digital research and innovation. Canada's AI strategy has focused primarily on research and talent promotion. Initiatives taken under this strategy aim to strengthen Canada's position internationally in the field of AI research and education. The 2024 budget will allocate
2.4 billion package of measures to maintain Canada's competitive advantage in AI. The investments are expected to accelerate job growth in Canada's AI and other sectors, boosting productivity by helping researchers and companies develop and deploy AI while ensuring technology accountability. (PM of Canada, 2024)

The Government of Canada has also published a guide to the use of generative AI in public administration institutions, which aims to set out principles for the responsible use of AI, identify challenges and provide good practices. The guide also aims to promote coordination between federal institutions. The guide also emphasises the importance of involving key stakeholders and stakeholder group collaboration when introducing AI tools in the public sector. These stakeholders include lawyers, privacy and security experts, the Office of the Chief Information Officer of the Government of Canada, bargaining agents, advisory groups, and customers of government services (Secretariat, T.B. of C., 2024).

There is a desire to increase the use of AI in Canadian public services, but the government does not want AI tools such as ChatGPT to be used by civil servants when dealing with confidential information such as cabinet documents (CBC News, 2024).

China

In its 2017 *Next Generation Artificial Intelligence* development plan, China has announced its ambition to become a world leader in AI research, technologies and applications. The development plan expresses China's ambition to become the world's largest AI

(State Council of China, 2017) In addition to development, China has long been active in the development of artificial intelligence and algorithms. With the wave of generative AI, a proactive approach to regulation has been pursued to ensure that AI models embody the socialist core values of the state and do not threaten national security as generative AI spreads into the corporate world. At the same time, China seeks to promote innovation and technological development.

Beijing sees AI as a key driver of economic and military competitiveness. Government support for generative AI is expressed, but the AI measures set out include comprehensive obligations on content management, privacy and security, and transparency of generative AI. The Generative AI Measures build on previous regulations on deep synthesis and algorithm management and represent the pattern of China's broader data management regulatory architecture. (Gong & Qu, 2023)

In 2023, the government will publish a regulatory framework for generative AI, which will require AI service providers to register language models in the government's algorithm registry in order to control their legality. In addition, the Department of Cyber Security has since published its own approved LLM training material. In particular, unregistered templates cannot be used
further development of the models, and as a result the use of foreign open source models is prohibited. Another notable policy is that the government recently blocked access to the Hugging Face repository, one of the main open databases where developers can collaborate and work on many generative AI models (Chang, 2024).

In May 2024, China published a new draft regulation on generative artificial intelligence, entitled "Cybersecurity Technology - Basic Security Requirements for Generative Artificial Intelligence (AI) Service". The draft will serve as guidance for both service providers and regulators. It provides guidance for conducting security assessments and drafting regulations (China Briefing News, 2024).

The draft defines critical safety requirements for generative AI services, which include the following subsections:

Security of training data: ensuring the security and integrity of the data used to train AI models.

Model security: protecting AI models from potential threats and ensuring their integrity throughout their lifecycle.

Safety measures: identify the essential safety measures to be taken to reduce risks effectively.

India

In 2018, the Indian government, through the public policy think tank NITI Ayog, published the National AI Strategy #AIForAll, which included AI research and development guidelines for various sectors: healthcare, agriculture, education, smart cities, infrastructure, smart mobility and disruption. The strategy focuses on how India can leverage transformative technologies to enable social and inclusive growth, in line with the government's development philosophy.

In addition, India aims to replicate these solutions in other similarly positioned emerging countries. From an applications perspective, the aim is to identify the sectors with the highest potential to benefit from the adoption of AI solutions, which will require the government to take the lead in developing an AI adoption plan (including the agriculture sector, which constitutes the economic backbone of India) (NITI Ayog, 2018).

In 2021, NITI Ayog published two documents on the principles of responsible AI. The first examines the various ethical aspects of AI deployment in India, divided into systems and societal perspectives. The systems perspectives address the general principles underlying decision-making, while the social perspectives focus on the impact of automation and job creation and employment. The second part focuses on the operationalisation and implementation of the principles of responsible AI. The report identifies the actions that both public and private sectors need to take, in collaboration with research institutions, to implement regulatory measures, develop policies and create a framework for compliance with relevant AI standards (Express Computer, 2024).

The government's efforts to regulate AI have been pro-innovation in nature, but at the same time policies and guidelines have been developed that identify the ethical concerns and risks associated with the use and deployment of AI. Specific guidelines on the use of AI in the provision of public services have not yet entered into force. Niti Ayog's policies and recommendations on responsible AI are not binding in nature and are not specifically designed for public administrations.

India recognises the transformative power of generative AI and is working to deploy the technology across sectors throughout society. Use cases already exist, such as **LawBot Pro**, a tool to help citizens apply for legal aid. In India, the proportion of lawyers in the population is very low, making legal aid inaccessible to many. LawBot Pro was created by lawyer Mandaar Mukesh Giri, who launched the project on a pro bono basis. Other sectors in which Genai tools can be seen
The major potentials for the future include agriculture, insurance, education and entertainment (OECD, INDIAai, 2023).

India has a culturally and linguistically diverse population, but of the 22 official languages spoken in the country, only Hindi and English are used by the central government for communication, posing numerous administrative challenges. To overcome these challenges, the Indian government is using various AI tools to improve administrative efficiency and narrow the language gap. For example, Jugalbandi, a free and open platform, combines ChatGPT and Indian translation models to bring conversational AI solutions to all sectors. The platform is an example of the potential of generative AI to address the diverse challenges of India's large population: enabling AI-based conversations in local languages, breaking down language barriers and aiming for greater accessibility (OECD, INDIAai, 2023).

India has a huge amount of digital data due to its large population, which gives the government a significant advantage in developing generative AI models due to the abundance of educational data. The international consultancy EY reports that the public sector
At national and government level, around 50% of organisations have set a target to deploy the first generative AI solutions. The EC-India statement notes that the potential of generative AI lies in particular in the efficiency and accessibility of public service delivery. According to the report, the main challenges organisations face in adopting GenAI are data protection, skills shortages and finding the right use cases (The Economic Times, 2024).

5 Conclusions of the mapping exercise

The international landscape of generative AI is dynamic and rapidly changing. National AI strategies emphasise different priorities depending on the country: key transnational themes include ethics and regulation, investment and innovation, education and skills, and international cooperation and standards. Around the world, efforts are being made to develop AI-oriented legislation and more binding regulatory frameworks. Generative AI has been in the spotlight since the breakthrough of ChatGPT and efforts are being made to promote its adoption in many countries across all sectors of society. Despite this momentum, estimates of the potential productivity benefits of generative AI are still blurred and divide the scientific community (Nathan et al., 2024).

The work on generative AI in different national administrations focuses on policy and guidance measures. Guidelines for the use of reliable and safe generative AI by public authorities are becoming more common in publications by the departments coordinating national AI development. In addition to guidance measures, the uptake of generative AI technologies in administrative tasks is increasing, but is still at an early stage in many countries.

Some experiments and pilots have been carried out in public administrations, such as the UK government's "red box" AI tools, which can be used to prepare draft responses to parliamentary questions and requests for information from public authorities. At the local level, generative AI tools are already widely used to support the work of public authorities and to improve digital services for citizens. In Germany, for example, much is being done in the field of generative AI at state level. Several cities are using LM-based tools to develop public services.

Our survey of domestic trials shows that there are many similarities between several of the ministries' trials, both in terms of use cases and lessons learned and challenges faced. Generative AI has mostly been used in the trials of the mapping exercise for tasks where the processing and analysis of large amounts of data is to be accelerated and made more efficient.

Potential uses included the pre-processing of drafting material for the legislative drafting process, the production of summaries of opinions, and more efficient information retrieval and the creation of limited overviews of large data sets.

Of the lessons learned, the working groups most emphasised the growing understanding of the possibilities and limitations of AI. The importance of data quality and interoperability was highlighted in several of the experiments. The value of source data is different if it i s in multiple formats or if there are inconsistencies in its structure (missing values, fields, duplicates, inconsistent formatting, or bias). High quality data is not only not sufficient but also a prerequisite for the successful application of AI technology.

The potential of generative AI as a support tool for mechanical, repetitive tasks was recognised, but at the same time the perception that AI requires human guidance and control was reinforced. Control bias was present in all experiments. Factors contributing to its reduction include the quality of the training data and its suitability for the application, as well as high quality body design.

The accuracy and legality of the results produced by AI systems must always be critically assessed and there must be defined verification mechanisms for data validation. The importance of quality control is highlighted when AI-generated outputs are used, for example, to support decision-making. As users of AI, public authorities must be aware of its limitations, risks and their own responsibilities both at the stage of writing requests and in assessing the further use of AI-generated content.

There was also an increased understanding of what is required and needed to move to the production phase: more basic work on data architecture, data mobility and semantic interconnectivity should be done before applying higher technologies. The choice of approaches should focus on being proactive, rather than reacting to challenges as they arise. Attention should also be focused on creating a culture change if new technologies are introduced in the organisation.

A key challenge in working with the private sector is the outsourcing of the technical implementation of AI projects. The role of private IT companies in the development, production and maintenance of AI solutions for public administrations raises questions about the independence of the public sector in AI procurement, especially when solutions are moved into production. This dependence on private actors may lead to a loss of technological competence and autonomy in the public sector, making the public sector more vulnerable to market changes.

The technical implementation of AI solutions should also be planned across ministerial boundaries, otherwise there is a risk of creating a fragmented development and production environment in the state administration. Several experiments were based on similar basic solutions. The source material varied according to the subject matter, but the use cases were very much focused on data retrieval, analysis, processing and summarisation. The interviews introduced the idea of a cross-agency baseline that could be further developed on a ministry- and agency-specific basis to suit the specific use case.

Many of the experiments relied on limited resources and a fast implementation schedule, which was reflected in the final results. Although valuable lessons were learned, the quality of the results was still far from production-ready. The lack of clear and realistic long-term objectives and the need for specific skills were highlighted. For the future, the need was identified to focus on the implementation of lessons learned, i.e. the identification and production of solutions that have been proven to work. Moving from the pilot phase to production of a demo solution requires consistent planning and goal setting, supported in particular by the strengthening of in-house IT expertise.

Building and maintaining scalable AI systems requires capital-intensive investments, whose impact and potential cost benefits must also be assessed on other parameters in addition to productivity gains. The pro-active consideration of AI legislation is becoming increasingly topical with the entry into force of the EU AI Regulation. The requirements of the GDPR and the AI Regulation must be taken into account in the design and development phase of AI systems and also in the selection of suppliers and service providers if the basic solution uses language models from companies outside the EU.

In the integration of generative AI, reliable and ethical technology development, deployment and continuous quality control are key to the success of AI systems in production.
the basis for solutions. Ethical challenges related to the technology, such as the proliferation of misinformation, the bias of source data and potential privacy violations, must be proactively taken into account when moving forward with new AI projects.

When selecting language models, the diversity, representativeness and inclusiveness of the education data directly affect the reliability of the model results. Embedding verification systems on the software side as part of AI solutions reduces the risk of spreading erroneous information produced by AI. Open discussions with third parties and stakeholders on the ethics and limitations of generative AI are needed.

Each organisation's journey towards adopting generative AI is different and must be designed around the unique needs that arise from its role, structure and services. An organisation's priorities in adopting AI should be guided by use case-specific goals and individual needs for AI capability development. Ideally, generative AI can be deployed and developed through an iterative approach, whereby Gen- AI tools can be gradually added to the existing technology and organisational environment as the need for further development, upgrades or new tools arises.

While generative AI promises to streamline mechanical, repetitive tasks and decision-making, in reality, issues related to transparency, model interpretation and stakeholder engagement remain major challenges on the AI path of public administration. The rapid adoption of generative AI may look promising, but there is also a risk that the wider AI space will be overshadowed by generative AI. Once risks are taken into account and guardrails are put in place, generative AI can be seen as having significant potential to be an effective support tool for the work of public authorities.

L AHTEE T

Generative AI tools in general:

CNAI (2024), Fact sheet on the use of generative AI tools in the Federal Administration. Available: https://cnai.swiss/wp-content/uploads/2024/01/Fact-sheet-on-the-use-of-generative-AI-tools-in-the-Federal-Administration_V1.2_EN_clear.pdf [Accessed 19 June 2024].

Heikniemi, J. (2023). Artificial intelligence based on language models opens up the productivity race. [online] Devisioona. Available at: https://www.devisioona.fi/2023/03/kielimallei- hin-based-technology-opens-the-productivity-race/ [Visited 23.7.2024].

Ramos, L., Yan, B., Khandabattu, H. & Rigon, G. (2024). When Not to Use Generative AI. Gartner.

Generative AI tools and guidance for officials:

CNAI (2023), Instruction sheets for the use of AI within the federal administration, https://cnai.swiss/en/products-other-services-instruction-sheets/ [Accessed 19.6.2024].

GOV.UK. (2024a). Guidance to civil servants on use of generative AI. [online] Available: https://www.gov.uk/government/publications/guidance-to-ci-vil-servants-on-use-of-generative-ai/guidance-to-civil-servants-on-use-of-gene-rative-ai [Accessed 20.6.2024].

CT policies and experiments:

Artificial Intelligence Commission. (2024). OUR AI: OUR AMBITION FOR FRANCE COMMISSION DE L'INTELLIGENCE ARTIFICIELLE. Available: https://www.info.gouv.fr/upload/media/content/0001/09/02cbcb40c3541390be391feb-3d963a4126b12598.pdf [Accessed 11 July 2024].

Central Banking. (2024). Initiative of the year: the Netherlands Bank's ChatDNB - Central Banking. [online] Available: https://www.centralbanking.com/awards/7960892/initiative-of-the-year-the-netherlands-banks-chatdnb [Accessed 9.7.2024].

Chang, W. (2024). China keeps generative AI on simmer | Merics. [online] Available: https://www.merics.org/en/comment/china-keeps-generative-ai-simmer [Exported 11.7.2024].

Chawdhry, Mt. (2024). Governing The Government's Use Of Artificial Intelligence In India. [online] Available: https://thesecretariat.in/article/governing-the-government-s-use-of-artificial-intelligence-in-india [Accessed 15 July 2024].

CBC News. Thompson, E. (2024). Federal government plans to increase its use of AI - with some big exceptions. [online] CBC. Available: https://www.cbc.ca/news/politics/public-service-artificial-intelligence-1.7216222.

China Briefing News. (2024). China Releases New Draft Regulations for Generative AI. [online] Available: https://www.china-briefing.com/news/china-releases-new-draft-regulations-on-generative-ai/.

CII Blog. (2023). Artificial Intelligence in Governance. Available at: https://ciiblog.in/ artificial-intelligence-in-governance/.

Congress.gov. H.R.6881 - 118th Congress (2023-2024): "AI Foundation Model Transparency Act of 2023". Available: https://www.congress.gov/bill/118th-congress/house-bill/6881/text. [Accessed 9 July 2024].

Department of International Cooperation, Ministry of Science and Technology, P.R. China. (2017). Next Generation Artificial Intelligence Plan. Saatavissa: http://fi.china-embassy.gov.cn/eng/kxjs/201710/P020210628714286134479.pdf.

Dutton, T. (2018). An Overview of National AI Strategies [online] Medium. Available: https://medium.com/politics-ai/ an-overview-of-national-ai-strategies-2a70ec6edfd.

Express Computer, Sharma, A. (2024). Unlocking potential: GenAI regulation and India's economic growth [online] Express Computer. Available: https://www.expresscomputer.in/guest-blogs/unlocking-potential-genai-regulation-and-indias-economic-growth/112408/ [Accessed 15 July 2024].

Gong, J. and Qu, H. (2023). What You Need to Know About China's New Generative AI Measures. [online] Available: https://www.twobirds.com/en/insights/2023/ china/what-you-need-to-know-about%20china%E2%80%99s-new-generati- ve-ai-measures [Accessed 22 July 2024].

GOV.UK. (2023). Algorithmic Transparency Recording Standard Hub. [online] Available: https://www.gov.uk/government/collections/algorithmic-transparency-recording-standard-hub.

GOV.UK. (2024b) Generative AI Framework for HMG (HTML) [online]. Available at: https://www.gov.uk/government/publications/generative-ai-framework-for-hmg/generative-ai-framework-for-hmg-html.

Government of the Netherlands. (2024). Dutch government presents vision on generative AI - News item - Government.nl. [online] www.government.nl. Available: https://www.government.nl/latest/news/2024/01/18/dutch-government-dutch-government-presents-vision-on-generative-ai.

Incubator for Artificial Intelligence - GOV.UK. (n.d.). Redbox. [online] Available: https://ai.gov.uk/projects/redbox/ [Accessed 22.7.2024].

NITI Ayog. (2024). National Strategy on Artificial Intelligence. Available at: https://www.niti.gov.in/sites/default/files/2023-03/National-Strategy-for-Artificial-Intelligence.pdf

OECD.ai. (2021). OECD AI Policy Observatory Portal. [online] Available: https://oecd.ai/en/dashboards/policy-initiatives/http:%2F%2Faipo.oecd.org%2F2021-data-policyInitiatives-27669 [Accessed 22.7.2024].

OECD (2024). OECD Artificial Intelligence Review of Germany [online] OECD iLibrary. Paris: Organisation for Economic Co-operation and Development. Saatavissa: https://www.oecd-ilibrary.org/docserver/609808d6-en.pdf?expires=1720698685&id=id&accname=guest&checksum=40F77CB2499C04D068AC21F3D97F8326 [Vierailtu 11.7.2024].

OECD.ai, INDIAai. (2023). OECD AI Policy Observatory Portal. [online] Available: https://oecd.ai/en/wonk/india-generative-ai [Accessed 22.7.2024].

Oxford Insights. (2023). Government AI Readiness Index 2023. Available: https://oxfordinsights.com/wp-content/uploads/2023/12/2023-Government-AI-Readiness-Index-1.pdf.

Prime Minister of Canada. (2024). Securing Canada's AI advantage. [online] Available: https://www.pm.gc.ca/en/news/news-releases/2024/04/07/securing-canadas-ai.

Reynolds, M. (2024). Could generative AI help break down language barriers plaguing the Indian justice system? [online] ABA Journal. Available: https://www.abajournal.com/web/article/could-generative-ai-help-break-down-language-barriers-plaguing-the-indian-justice-system [Accessed 22 July 2024].

Secretariat, T.B. of C. (2024). Guide on the use of generative artificial intelligence. [online] www.canada.ca. Available: https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/guide-use-generative-ai.html#toc-5.

The Economic Times. (2024). Half of government, public service organisations in India to apply Gen AI solutions in one year: EY survey (2024) [online] 8 Apr. Available: https://economictimes.indiatimes.com/tech/technology/government-and-public-service-organisations-embracing-gen-ai-solutions-ey-report/articleshow/109127091.cms [Accessed 15 July 2024].

National experiments:

Aryani, A. (2023) What is your Context Window? - Amir Aryani - Medium [online] Medium. Available: https://medium.com/@amiraryani/what-is-your-context-window-772b1c65b881 [Accessed 23 July 2024].

Ashtari, H. Spiceworks. (2022). Edge Computing vs. Cloud Computing: 10 Key Comparisons. [online] Available: https://www.spiceworks.com/tech/cloud/articles/edge-vs-cloud-computing/.

Bashir, N., Donti, P., Cuff, J., Sroka, S., Ilic, M., Sze, V., Delimitrou, C. and Olivetti, E. (2024). The Climate and Sustainability Implications of Generative AI. An MIT Exploration of Generative AI. [online] Available: https://mit-genai.pubpub.org/pub/8ulgrckc/release/2.

BCG Global. (2024). Generative AI for the Public Sector: the Journey to Scale. [online] Available: https://www.bcg.com/publications/2024/gen-ai-journey-to-scale-in-government.

Choi, N. (2024). What is retrieval-augmented generation, and what does it do for generative AI? [online] The GitHub Blog. Available: https://github.blog/2024-04-04-what-is-retrieval-augmented-generati-.on-and-what-does-it-do-for-generative-ai/.

Davis, G. (2023). A developer's guide to open source LLMs and generative AI. [online] The GitHub Blog. Available: https://github.blog/ai-and-ml/llms/a-developers-guide-to-open-source-llms-and-generative-ai/ [Accessed 13 Aug 2024].

Farquhar, S., Kossen, J., Kuhn, L. and Gal, Y. (2024). Detecting hallucinations in large language models using semantic entropy. Nature, [online] 630(8017), pp.625-630. Available at: https://doi.org/10.1038/s41586-024-07421-0.

Forbes India. (2024). The Global Race For AI Chips Intensifies. Where Does India Stand? [online] Available: https://www.forbesindia.com/article/take-one-big-story-of-the-day/the-global-race-for-ai-chips-intensifies-where-does-india-stand/92517/1.

granica.ai. (2024). AI Cost Analysis Guide. [online] Available: https://granica.ai/blog/ai-cost-analysis-grc [Accessed 19.7.2024].

Gunashekar, S., van Soest, H., Qu, M., Politi, C., Aquilino, M.C. and Smith, G. (2024). Examining the landscape of tools for trustworthy AI in the UK and the US: Current trends, future possibilities, and potential avenues for collaboration. [online] www.rand.org. Available: https://www.rand.org/pubs/research_reports/RRA3194-1.html [Accessed 9.7.2024].

Haltu Ltd. (2023). Secure software development - best practices. [online] Available: https://www.haltu.fi/blogi/tietoturvallinen-ohjelmistokehitys [Exported 26.7.2024].

Mittal, A. (2023). The black box problem in LLM: challenges and new solutions [online] Unite.AI. Available: https://www.unite.ai/fi/the-black-box-problem-in-llms-challenges-and-emerging-solutions/ [Accessed 9 Jul 2024].

Murthy, S. Abu Bakar, A., Abdul Rahim, F., and Ramli, R. "A Comparative Study of Data Anonymization Techniques," *2019 IEEE 5th Intl Conference on Big Data Security on Cloud (BigDataSecurity), IEEE Intl Conference on High Performance and Smart Computing, (HPSC) and IEEE Intl Conference on Intelligent Data and Security (IDS)*, Washington, DC, USA, 2019, pp. 306-309, doi: 10.1109/BigDataSecurity-HPSC-IDS.2019.00063.

Nathan, A. et al. (2024). Goldman Sachs Research Newsletter. Gen AI: Too much spend, too little benefit?. [online] Available: https://www.goldmansachs.com/intelligence/pages/gs-research/gen-ai-too-much-spend-too-little-benefit/report.pdf.

Patel, D. (2023). AI Capacity Constraints - CoWoS and HBM Supply Chain. [online] www.semianalysis.com. Available: https://www.semianalysis.com/p/ ai-capacity-constraints-cowos-and.

QAT Global (2024). Financial Implications of AI: Costs & Benefits of Developing GenAI [online] QAT Global. Available: https://qat.com/financial-implications-of-ai/ [Accessed 19 July 2024].

Rahul, S (2024). Context Window of Language Models - Rahul S - Medium. [online] Medium. Available: https://ogre51.medium.com/context-window-of-language-models-a530ffa49989 [Accessed 30 Jul 2024].

Sitra. (2024). Ensuring European leadership in the data economy. [online] Available: https://www.sitra.fi/en/publications/ensuring-european-leadership-in-the-data-economy/ [Visited 3.7.2024].

Teknologiateollisuus. (2024) EU AI Regulation - Safety regulation for AI systems and general-purpose AI models* Available: https://teknologiateollisuus.fi/sites/default/files/inline-files/EU_teko%C3%A4lyasetus_Teknologiateollisuus.pdf [Visited 6.8.2024].

Mikkilä, J. (Technology Industry) (2024). Microchips are Finland's sustainable growth accelerators. [online] Available: https://teknologiateollisuus.fi/mikrosi- rut-are-suomen-sustainable-growth-accelerators/ [Visited 6.8.2024].

www.promptingguide.ai. (n.d.). Body design guide - Nextra. [online] Available: https://www.promptingguide.ai/fi [Accessed 30.7.2024].

Y. -C. Wang, J. Xue, C. Wei and C. -C. J. Kuo, "An Overview on Generative AI at Scale With Edge-Cloud Computing," in *IEEE Open Journal of the Communications Society*, vol. 4, pp. 2952-2971, 2023, doi: https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=10268594.

Zhao, H., Chen, H., Yang, F., Liu, N., Deng, H., Cai, H., Wang, S., Yin, D. and Du, M. (2023). Explainability for Large Language Models: A Survey. ACM Transactions on Intelligent Systems and Technology, 15(2), pp.1-38. doi: https://doi.org/10.1145/3639372. url: https://arxiv.org/pdf/2309.01029

Data security and competition law:

OECD. (2024). The intersection between competition and data privacy. [online] Available: https://www.oecd.org/en/publications/the-intersection-between-competition-and-data-privacy_0dd065a3-en.html?utm_campaign=Copy%20of%20STI%20News%209%20July&utm_content=Read%20the%20report&utm_term=sti&utm_medium=email&utm_source=Adestra

Intellectual property rights:

Karus, R. (2024). Generative AI and copyright in 2024. [online] IPRinfo. Available: https://iprinfo.fi/uutiset/ generative-artificial-intelligence-and-copyright-in-2024/.

World Intellectual Property Organization.(2024) Generative AI Navigating Intellectual Property IP and Frontier Technologies.(2024). Available at: https://www.wipo.int/export/sites/www/about-ip/en/frontier_technologies/pdf/generative-ai-factsheet.pdf.

The environmental impact of artificial intelligence:

Bashir, Noman, Priya Donti, James Cuff, Sydney Sroka, Marija Ilic, Vivienne Sze, Christina Delimitrou, and Elsa Olivetti. 2024. "The Climate and Sustainability Implications of Generative AI." An MIT Exploration of Generative AI, March. Available at: https://mit-genai.pubpub.org/pub/8ulgrckc/release/2

digital-skills-jobs.europa.eu. (n.d.). LMI, Europe's newest supercomputer | Digital Skills & Jobs Platform. [online] Available at: https://digital-skills-jobs.europa.eu/en/ latest/news/lumi-new-european-supercomputer .

Intelligence, G.T. (2024). The environmental cost of AI. [online] Verdict. Available: https://www.verdict.co.uk/the-environmental-cost-of-ai/.

/thoughtworks. (2023). How can we respond to the environmental impact of generative AI? [online] Thoughtworks. Available: https://www.thoughtworks.com/insights/articles/how-respond-to-environmental-impact-generative-ai

Fjeld, J., Achten, N., Hilligoss, H., Nagy, A., Srikumar, M.: Principled artifcial intelligence: mapping consensus in ethical and rights-based approaches to principles for AI. SSRN Electron. J. (2020). https://doi.org/10.2139/ssrn.3518482.

Jobin, A., Ienca, M., Vayena, E.: The global landscape of AI ethics guidelines. Nat. Mach. Intell. 1(9), 389-399 (2019). https://doi.org/10.1038/s42256-019-0088-2.

Roche, C., Wall, P.J. and Lewis, D. (2022). Ethics and Diversity in Artificial Intelligence Policies, Strategies and Initiatives. AI and Ethics, 3. doi: https://doi.org/10.1007/s43681-022-00218-9

Solaiman, I. et al (2024), Evaluating the Social Impact of Generative AI Systems in Systems and Society. [online] Available: https://arxiv.org/pdf/2306.05949.

Suomi.fi. (2024). Consider the laws and recommendations - Responsible use of AI - Suomi.fi for developers. [online] Available: https://kehittajille.suomi.fi/oppaat/responsible-tech/ethics/consider-the-laws-and-recommendations

Ethical challenges:

Antal, S. (2023). Artificial Intelligence in the Global South: Will AI Advancement Deepen Digital Divides and Inequalities? [online] Portulans Institute. Available: https://portulansinstitute.org/artificial-intelligence-in-the-global-south/

Bommasani, R., Klyman, K., Kapoor, S., Longpre, S., Xiong, B., Maslej, N. and Liang, P. (n.d.). The Foundation Model Transparency Index v1.1 May 2024. [online] Available at: https://crfm.stanford.edu/fmti/paper.pdf.

crfm.stanford.edu. (2024). Foundation Model Transparency Index. [online] Available: https://crfm.stanford.edu/fmti/May-2024/index.html.

CSC. (2023). Research team created the largest-ever Finnish language model using the LUMI supercomputer [online] Available: https://csc.fi/uutinen/ research-team-created-the-largest-ever-suomi language-model-using-the-UMI-supercomputer/.

Global Privacy Assembly. (2023). 45th Closed Session of the Global Privacy Assembly. (2023). Available at: https://www.edps.europa.eu/system/files/2023-10/ edps-gpa-resolution-on-generative-ai-systems_en.pdf.

Heikkilä. M. (2023). MIT Technology Review. It's high time for more AI transparency. [online] Available: https://www.technologyreview.com/2023/07/25/1076698/its-high-time-for-more-ai-transparency/.

Mikkilä, J. (2024). The EU's historic AI regulation has now moved forward - read the Technology Industry's compilation | Teknologiateollisuus. [online] Available: https:// teknologiateollisuus.fi/en/ajankohtaista/eun-historallinen-tekoalyasetus-nytkah- ti-eteenpain-tutustu-teknologiateollisuuden

Office of the Data Protection Ombudsman (n.d.). Pseudonymisation and anonymisation of personal data. [online] Available at: https://tietosuoja.fi/ pseudonymisation-anonymisation

Other references:

Chen, M. et al (2021). Evaluating Large Language Models Trained on Code [online] Available: https://arxiv.org/pdf/2107.03374.

European Union. (2022). EU chip act. [online] Available: https://www.consilium.europa.eu/en/infographics/eu-chips-act/.

Mittal, A. (2023). The role of vector databases in modern generative artificial intelligence applications [online] Unite.AI. Available: https://www.unite.ai/fi/the-role-of-vector-databases-in-modern-generative-ai-applications/.

The Personal Data Protection Commission (2018). Guide to basic data anonymisation techniques. [online] Available at: https://www.pdpc.gov.sg/-/media/Files/ PDPC/PDF-Files/Other-Guides/Guide-to-Anonymisation_v1-(250118).pdf.

VALTIOVARAINMINISTERIÖ
FINANSMINISTERIET

MINISTRY OF FINANCE
Snellmaninkatu 1 A
PL 28, 00023 COUNCIL OF STATE
Phone 0295 160 01 vm.fi

ISSN 1797-9714 (pdf)
ISBN 978-952-367-840-8 (pdf)

September 2024