Findings on the applicability and potential of AI tools in participatory practices

Frederick Ducatelle (CollectiveUP) & Liliana Carrillo (CollectiveUP)

Driving Urban Transitions

Co-funded by the European Union

This project has been funded by Innoviris (Belgium), Flanders Innovation and Entrepreneurship (Belgium), DLR Projektträger (Germany), the National Research, Development, and Innovation Office (Hungary), Fundação para a Ciência e a Tecnologia (Portugal), and The National Centre for Research and Development (Poland) under the Driving Urban Transitions Partnership, which has been co-funded by the European Commission.

Table of contents

Imprint

2025: CollectiveUP info@collectiveup.be

Authors: Frederick Ducatelle (CollectiveUP), Liliana Carrillo (CollectiveUP), Chrysanthi Katrini (CollectiveUP).

Layout: Jurate Laugalyte (CollectiveUP)

We suggest citing this report as follows: Ducatelle, F., Carrillo, L. (2025). Findings on the applicability and potential of AI tools in participatory practices. CONIFER project. Driving Urban Transformation program. https://conifer15.eu/

Disclaimer

This publication has been written within the Driving Urban Transitions Partnership project CONIFER (https://conifer15.eu/), action type “15mC topic 3: Empower People for Urban Mobility Transitions”, co-funded through the European Commission. CONIFER aims to develop and test a participatory foresight methodology that combines a system oriented, structured approach with creative and participatory methods for co-creating 15mC futures based on the diverse needs of children, young people, carers and teachers.

Executive Summary

This report is the output of Task 1.4 within the CONIFER project. The project aims to develop a participatory foresight methodology to explore 15-minute city futures, focusing on the needs of children, young people, their carers, and teachers. It will create multiple scenarios and collective visions for the proximity city by combining innovative systems dynamics with creative methods such as gamification, art, design thinking, and signature events, while also investigating the role of artificial intelligence in supporting this process. Task 1.4, specifically, analyzes the applicability, potential, and possible risks of AI tools, especially in the context of interactions with children and young people. This assessment is crucial for supporting the testing of AI tools in the visioning process (WP4) and throughout the rest of the project.

The objective of this report is to serve as a reference for the use of AI methods in the CONIFER project. In the first place, it provides technical foundations, describing the field of AI, and in particular the algorithms that could be most relevant to the project. It gives an introduction to some technical aspects with the aim to inform non-AI experts about the possibilities. Second, it describes the state of the art and available tools for the application of these methods to the topics of the project, in particular participatory foresight and urban planning. Next, it provides ideas and concrete proposals for the use of AI within the project, going into detail on the various work packages. Finally, the report also focuses on challenges of the use of AI, and has particular attention to ethical considerations when working with children and young people.

1. Introduction

In the past two decades, there has been tremendous progress in Artificial Intelligence (AI), producing a diverse set of new algorithms and tools that can be used in a wide range of applications and fields. Especially the arrival and public availability of general-purpose generative AI models and chatbots in the last couple of years can be considered a breakthrough. These models are rapidly being adopted by the general public (Hu, 2023) and are transforming businesses, processes, and in general people's lives around the world.

In this report, we explore the possibilities of using AI in participatory foresight, and particularly the applicability of diverse AI methods for the purposes of the CONIFER project. Foresight is the process to systematically attempt to look into the future of science, technology and society with the aim of supporting policies and decision making (Martin, 1996). Participatory foresight takes a bottom-up approach to foresight, encouraging the involvement of citizens (or more in general: people without specific expertise in the field under investigation) throughout the process (Nikolova, 2014).

In recent years, there have been several approaches to integrate AI within foresight processes. Some of them are high-level descriptions for AI integration into foresight frameworks (Geurts, 2021; Carvalho, 2024; Wang, 2025). Others are attempts to integrate AI into specific steps in the foresight process (e.g., Jung et al., 2023; Bresinsky et al., 2024; Tan and Luke, 2024). In this report, we take a practical approach, focusing on the vision, goals and methods of the project, and proposing concrete and applicable AI solutions for those. Moreover, we take into consideration the challenges that come with using AI, particularly when standing in conflict with the ambitions of the project, such as promoting citizen participation, equity, and inclusion.

This report is structured as follows. First, we give an overview of the AI methods we deem relevant for this project, describing their working and their potential applications. Next, we discuss the opportunities and use of AI in foresight in general. After that, we look into the various work packages of the project and propose concrete applications of AI. Finally, we zoom in on challenges and ethical considerations with the use of AI and how we want to handle those.

2. General Background on Relevant AI Technology

In this section, we want to provide an accessible introduction to relevant AI technologies. As the field of AI is very broad, we first provide a very high-level overview of the field as a whole, and then focus on the technologies relevant to the current project. Understanding the various algorithms in detail can be complicated and technical, so we aim to find a balance between simplicity and detailed information to make it possible to understand the possibilities and limitations of each technology.

2.1 The field of AI

The field of AI is notoriously difficult to define and delimit (Lea, 2024). This is not only because it is a relatively new and rapidly evolving research area with many subfields, but also because even a generally accepted definition of the term “intelligence” itself is lacking (Legg and Hutter, 2007). Many definitions of AI have been provided and discussed (Russel and Norvig, 2020). For the purposes of this report, we use the definition of the High-Level Expert Group on Artificial Intelligence of the European Commission (AI HLEG): “systems that display intelligent behaviour by analysing their environment and taking actions – with some degree of autonomy – to achieve specific goals” (AI HLEG, 2019). While using such an open definition has some important advantages, such as leaving room for the many innovations and evolutions in this rapidly changing field, it remains quite abstract. To provide a more concrete understanding of the field of AI, the following section specifies certain characteristics further.

One major distinction that is often made within the study of AI, is the difference between weak AI and strong AI (Russel and Norvig, 2020). Weak AI is also called narrow AI and refers to the study of algorithms that intelligently address a particular task or a narrow domain. Strong AI is also referred to as Artificial General Intelligence (AGI) and refers to the goal of creating systems that can solve a wide range of problems that require intelligence, systems that are actually intelligent, in the same way humans are. While many researchers focus on the development of AGI, all practically usable AI algorithms and applications fall under the umbrella of weak AI. Nevertheless, the arrival of powerful and versatile generative AI in recent years has sparked the debate on whether AGI is around the corner (Bubeck et al., 2023).

Zooming in on weak AI, we can provide a description of the field of AI by focusing on the various narrow tasks or objectives researchers try to address. Other useful ways to describe the field of AI are based on the types of algorithms used or on the research field that provided the inspiration for the algorithms (biology, neuroscience, statistics, economics, etc.). Our choice to focus on tasks is motivated by our aim for practical applications of AI in this report and this project.

At a very high level, the following sub-tasks could be used to describe the field of AI (Russel and Norvig, 2020):

- *Problem-solving, search and optimization*. This is concerned with finding the optimal solution to a problem, including, e.g., finding the best move in chess or the shortest path in a network. Inspiration is taken from many sources, including from biological systems, such as evolution and swarm behaviour (see section 2.5 for Swarm Intelligence).
- *Knowledge, reasoning, planning, decision making*. This is about representing knowledge and reasoning with it. It includes symbolic methods (e.g., expert systems), as well as ontologies and probabilistic planning and reasoning methods.
- *Learning*. This is the field of Machine Learning (ML) and involves all kinds of methods to let a computer program learn its behaviour from data. A particular form of ML is Generative AI, which is concerned with models that learn to produce artifacts (e.g., text, images or audio) as opposed to making decisions about data (e.g., image classification tasks).
- *Natural Language Processing (NLP)*. This includes all algorithms that let computers process natural language, including, most notably, chatbots like ChatGPT. Other forms of NLP include sentiment analysis, machine translation and information retrieval.
- *Perception*. This covers all algorithms related to processing sensory information, most notably audio signals (including speech recognition) and vision (including image recognition and object tracking). Recent advances in multi-model algorithms combine multiple forms of perception and natural language processing, blurring the lines between some of the traditional sub-tasks in AI.
- *Robotics*. This large area of AI is related to embodied intelligent systems, in the form of various types of robots, such as, e.g., industrial robots, humanoid robots, or even swarms of small robots.

Finally, it is important to note that the field of AI has changed drastically since the late 2010s, driven by rapid progress in ML. While there have been earlier periods of ups and downs in the field in terms of attention, expectations and funding (Russel and Norvig, 2020), the current progress and rise in impact seem more profound, with AI leading to major breakthroughs in a wide range of fields. People therefore speak of an AI Boom (Hajkowicz et al., 2023). The new algorithms developed during this AI Boom are so radically more powerful than much of what existed before, that most practical use of AI now relies on work from the last few years.

In the next few sections, we will present AI methods that could be useful for the CONIFER project. We will organize this discussion along the lines of the sub-fields of AI these algorithms belong to.

2.2 Natural Language Processing

Within the CONIFER project, NLP will probably play an important role, as we need to gather, analyze and produce a lot of natural text data in the various work packages. NLP algorithms can help us with this. In what follows, we first discuss task-specific NLP algorithms from before the AI boom and explain which relevance they still have. After that we move on to the more powerful models produced in recent years, including the Large Language Models that form the basis of modern AI chatbots.

2.2.1 Before the AI boom: task-specific NLP models

As explained above, recent advancements have brought us new AI algorithms that are much more powerful than anything that came before. In the field of NLP, a very distinctive property of recent models, such as AI chatbots, is that they are very versatile and can answer almost any question you throw at them. In contrast, the state-of-the-art in NLP research until a few years ago was focused on models that solve very specific tasks. Examples of such tasks are Named Entity Recognition (to extract names of people, places, etc. from documents), Sentiment Analysis (to spot positive or negative sentiments in texts), or Information Retrieval (to search relevant documents based on a query or keywords). (Jurafsky and Martin, 2025).

Task-specific models have a lot of disadvantages compared to more powerful general-purpose NLP models. The most obvious is that you need a separate model (and often even a different type of algorithm) for each task. AI practitioners would typically spend a lot of time gathering and labeling data for their specific task and training and optimizing their models. The quality of the model and its generalization capabilities (even within the specific task it was trained for) depended very much on the quality of execution in the various steps involved in this work. Compared to this type of process, the appeal of the use of general-purpose powerful AI models such as AI chatbots is obvious (Qin et al., 2024).

Nevertheless, these task-specific models have some advantages that ensures they sometimes still have a place in modern AI-based applications (Qin et al., 2024). These advantages include low resource consumption and cost for deployment, better controllability, and often higher precision. One interesting approach to use these older types of models, is to combine them with more powerful AI systems, e.g., to filter data that will be presented to the more powerful system (Liang et al., 2024; Negi et al., 2024; Zhang et al., 2024). In section 2.4, we explain more in detail what such a combined AI system could look like.

Within the CONIFER project, such a combined system could play a role in WP3, where we need to process large amounts of text data from external sources. Examples are task 3.4, which focuses on the analysis of mobility narratives in political discourse, and task 3.5, which deals with the analysis of biases related to mobility discussions in social media. In this context, task-specific NLU models for sentiment analysis or information retrieval could be used. See section 4.1 for more details.

2.2.2 Sentence encoders and document retrieval

A sentence encoder is a type of model developed during the early years of the AI boom (Reimers and Gurevych, 2019) that is still very much in use today (Zhao et al., 2024). It takes as input a sentence (or a longer text), and returns a mathematical representation of that sentence, in the form of a vector, called an embedding. Such an embedding vector captures the meaning of the sentence, in the sense that two sentences with similar meaning will have embedding vectors that are geometrically close to each other. This means that one can evaluate semantic similarity of sentences via simple mathematical operations (in particular, calculating the distance between their vectors).

Embeddings can be used in various ways. One very popular use is in retriever systems, which essentially implement a search system by comparing the embedding of a search request with embeddings of stored text documents (Karpukhin et al., 2020). Like this, documents can be retrieved from a large document database, based on their semantic similarity to the search query. Several services support the creation of document embedding databases with minimal coding, including Google Vertex AI Search and Pinecone.io. These systems make it easy to set up a search engine over proprietary information. Another common use is to organize documents, clustering them according to their meaning, or finding documents that are divergent in content, by spotting vectors that are geometrically far from the main document clusters (Wang, 2024).

Within the CONIFER project, sentence encoders could play an important role in selecting and organizing data from various sources in WP3, or in analysing, organizing and understanding data collected from the workshops in T3.1 and T3.3. This will be explained more in detail in sections 4.1 and 4.2.

2.2.3 Large Language Models

Language modelling is a core area within NLP. It is concerned with modelling the relative likelihood of words and sentences in a particular language or a subdomain of it. Language Models (LMs) have been around for decades, and are frequently used in common applications, such as speech recognition (to pick the most probable sentence from multiple similar sounding possibilities) or grammar correction (to spot unlikely sentences). Generative LMs are LMs that can be used to generate new text, typically starting from an incomplete sentence or text and generating new words one by one by picking the most likely next word each time. One application for such models that has been around for a long time is auto-completion.

Recent advances during the AI Boom in terms of model architectures (Vaswani et al., 2017), training methods (Ouyang et al., 2022) and especially the scale of the models and training data, have led to the arrival of so-called Large Language Models (LLMs). The unprecedented power of these models and the unexpected emergence of advanced reasoning capabilities in them, have marked a breakthrough in AI, most notably in the form of the popular ChatGPT model, which is at the forefront of the current AI boom (Brown et al., 2020).

LLMs for chat applications, like ChatGPT, typically use a decoder-only Transformer architecture (Vaswani et al., 2017), which is a type of generative LMs. They are pretrained on as much text as possible, including websites, journals, books, but also programming code, websites, etc., with the simple task of predicting the next word in a sequence, given the previous words (Radford et al., 2019). Afterwards, their next-word predictions are finetuned with various methods to make them behave in desirable ways for chatbots, e.g., to focus on generating answers to questions (e.g., using question-answer data from internet fora), to behave friendly, or avoid harmful answers and undesirable biases. (Christiano et al., 2017). This approach to LLM creation is entirely reflected in the name "ChatGPT": a Generative Pretrained Transformer finetuned for Chat (OpenAI, 2022).

The enormous power of LLMs lies in their very strong capabilities to understand text and give meaningful answers. The knowledge they use in their answers comes from two sources: from the information they have stored internally in their model weights (this is their world knowledge, which is based on what was in their training data) and from the input text that was provided to them in the question (or, more in general, the “prompt”, which is the full input text provided to the LLM including question and extra information) (Liu et al., 2021). This leads to two important ways to adapt an LLM’s behavior to particular tasks: finetuning (Ouyang et al., 2022), whereby the LLM’s training is extended with particular examples and feedback, and prompt engineering (Liu et al., 2021), where people look for ways to formulate and extend their prompt in such a way that they get better answers from the LLM. Examples of prompt engineering approaches include few-shot learning (Brown et al., 2020) (where a few examples of the desired input and output are added in the prompt) and Chain-of-Thought (CoT) prompting (Wei et al., 2022), where the simple addition of the words “let’s take this step by step” was shown to often improve the reasoning capabilities of the LLM. Compared to prompt engineering, finetuning is more costly and complex, because it requires dedicated data, a training process, and the deployment of a proprietary model, but it can sometimes give better results (Shin et al., 2025).

LLMs are very powerful models and are revolutionizing many tasks and creating new opportunities. Their ability to process and understand text in an automated manner, and answer any question about it, makes them general-purpose AI systems. With a little bit of prompt engineering, they can be used for all kinds of applications. E.g., within the area of NLP, they can handle tasks like sentiment analysis, summarization, text generation, etc., without the need for extra training (Brown et al., 2020). This explains why they are quickly replacing the task-specific models we described earlier.

Within the CONIFER project, LLMs will probably play an important role. Within WP3, they can be applied to enhance and automate the processing of internal and external sources, such as..... Within WP4 (Imagining 15-minute city futures) and WP5 (Designing policy pathways), their text generation capabilities can be very useful, e.g., to describe potential future scenarios and bring them to life. This will be described in detail in section 4.

Finally, it’s important to note that LLMs have quite some downsides too, and their use also raises a lot of concerns. We describe these issues in detail in section 5, and will discuss ways we can counter these concerns in the CONIFER project.

2.3 NLP, Perception and more: Generative AI beyond text

While LLMs have received the most attention in recent years, generative AI is not limited to text. Similar advances have been made in models for other types of data, such as images, video or audio. This concerns both the input side, e.g., processing and interpreting images or sound, and the output side, e.g., generating images or video. For images, models like DALL-E (Ramesh et al., 2021) and Stable Diffusion (Rombach et al., 2022) can generate complex pictures from text prompts. Similar approaches have been extended to video, with models like Sora (by OpenAI) generating short video clips based on a single prompt by modelling not just spatial structure, but also temporal consistency. In audio, models such as Jukebox (Dhariwal and Jun et al., 2020) can create music in specific styles, while others like VALL-E (Wang et al., 2023) can synthesize human speech from a short voice sample.

Many of these models combine support for multiple types of data, and are therefore called multimodal. E.g., GPT-4V (OpenAI, 2023) combines image and text understanding, allowing, for instance, dialogue about images or diagrams. Such models show how AI models are becoming more general-purpose and the traditional distinctions between the various tasks and sub-fields of AI are disappearing.

Beyond text, image, audio, and video, generative AI is rapidly expanding into other modalities such as 3D model generation (see e.g., meshy.ai), web and UI design, code synthesis and motion generation. These models can, for example, create 3D shapes from text prompts (Mildenhall et al., 2020; Poole et al., 2022), or generate functional websites or user interfaces. Often, these outputs are produced by generating structured text (like code or configuration files) that downstream systems interpret into other forms, showing how generative models can act as flexible intermediaries between language and more complex data types.

Particularly interesting for the CONIFER project might be AI models that generate 3D worlds. E.g., the startup worldlabs.ai, founded in September 2024, is creating a service to generate a navigable 3D world from a single image. Similar services are targeted by PentoPix and Improbable. Even more pertinent is UrbanWorld (Shang et al., 2024), an open source software tool to create 3D models of cities. It starts from publicly available data like OpenStreetMap and lets you improve and adapt the generated world via image and text input. The result is a Blender 3D model of your urban environment that is both realistic and customized. It remains to be seen though how mature these systems are for use in urban planning.

There is a lot of potential for the use of multimodal AI models in the CONIFER project. In WP3, they can help us include images and video in our analysis of internal and external sources. In WP4, they can help us produce all types of output during the workshops, and afterwards help us in processing and analysing the results. These multimodal models have the same challenges as LLMs though, in particular in terms of resource consumption and cost, so the discussion in section 5 is very important.

2.4 Combining AI Models: from NLP to autonomous planning and decision-making

While LLMs and multimodal models are at the core of many recent advances in AI, they are often used in combination with other algorithms and code, to create more powerful or more reliable systems. In this section, we describe the most important techniques to make such combined AI systems and explain the relevance for the project CONIFER.

2.4.1 Retrieval-Augmented Generation

Retrieval-Augmented Generation (RAG) is a method for overcoming some key limitations of LLMs, especially their lack of access to up-to-date or domain-specific knowledge. It combines embedding-based semantic retrieval (see section 2.2.2) with language generation (see section 2.2.3) to augment an LLM's capabilities with external information. The basic process has three stages. In the indexing phase (run offline), external knowledge is selected, split into smaller chunks, embedded as vector representations, and stored in an index. In the retrieval phase (run at query time), a user query is embedded and matched against this index to retrieve the most relevant chunks. In the generation phase, the original query and the retrieved text fragments are combined into an augmented prompt that is passed to the LLM to generate an answer. Multiple variations and improvements of this basic RAG approach exist, see (Lewis et al., 2020; Gao et al., 2023).

Compared to regular LLM usage, RAG enables access to external, current, or proprietary information not present in the model's training data. It helps focus the LLM on relevant context, improving factual accuracy and reducing hallucinations. It can also improve the user's understanding of and trust in the model, as the retrieved text fragments can be used to indicate where the model's information came from and add references to external documents. RAG can be seen as a form of automated prompt engineering. It is more lightweight and flexible than model finetuning, as retrievers are relatively small and easy to update. Compared to manual prompt engineering, RAG allows far more information to be considered without exceeding prompt length limits, making it especially suitable for knowledge-intensive applications. RAG is used in many real-world systems, including enterprise search assistants, legal and medical document analysis tools, and customer support chatbots that integrate private documentation.

2.4.2 LLM Agents and Multi-Agent Systems

LLM agents and Multi-Agent Systems (MAS) integrate LLMs into autonomous systems capable of performing complex, real-world tasks. An agent is a software component that autonomously resolves tasks (Russel and Norvig, 2020), and in an LLM agent, the LLM serves as the "brain" of the agent, enabling it to solve intricate problems (Zhao et al., 2023). By iteratively calling the LLM (possibly finetuned for the Agent’s specific task) with carefully crafted prompts and processing its responses, agents can create detailed plans to achieve complex objectives (Huang et al., 2024), execute the steps of the plan (which may involve using external tools or APIs) (Qu et al., 2024), revise results, and formulate responses to user queries. For example, an LLM agent designed for travel planning could autonomously create a full itinerary, including flights and accommodations, tailored to the user's preferences, all based on a simple natural language query (Singh et al., 2024).

MAS takes this concept further by combining multiple agents that collaborate and coordinate to solve tasks that are too complex for a single agent (Liu et al., 2021; Shinn et al., 2023). In MAS, LLM agents might collaborate on tasks like business process automation, customer support, or creative content generation. For instance, a team of agents in a fully automated software development agency could work together, where one agent handles bug fixes, another generates code for new features, and a third conducts tests (Chen et al., 2024). This coordination allows agents to specialize and focus on particular aspects of the problem, leveraging each other’s capabilities.

Both LLM agents and MAS are complex systems, and designing them to work reliably is quite difficult (Cemri et al., 2025). Moreover, their resource-intensive nature makes them costly to run at scale. Despite these challenges, the ability of these systems to coordinate and autonomously address complex, dynamic tasks makes them an exciting direction for the future of AI, unlocking new possibilities for practical applications in real-world scenarios.

Within the CONIFER project, we may use an LLM agent in WP3 if we want to set up more complex and fully automated processes for data gathering and analysis, especially if we want to integrate multiple data modalities, such as text, images and video. Also in WP4, an

LLM agent could be used, e.g., if we want to automate the analysis of factors influencing the future of mobility in T4.1. MAS are probably outside the scope of this project, though it is interesting to think how one could fully automate foresight methodologies using a MAS.

2.5 Collective Intelligence and Swarm Intelligence

Collective Intelligence (CI) is a term used to describe the intelligence reflected in decisions or results obtained by large groups (Mulgan, 2018; He et al., 2019). Many studies and examples have shown that groups can often come to better decisions than the individuals participating in the process. This is sometimes referred to as the wisdom of the crowds (Surowiecki, 2004). Examples include online collaborative systems, such as Wikipedia or Waze, but also the so-called "jury theorem", dating back to 1785 (Condorcet, 2014), which states that as long as each individual jury member has a probability of more than 50 percent to reach the correct decision, the probability of the group achieving the correct decision increases with the size of the group. This is an attractive element for participatory decision making: by including the collaboration of more individuals, we can reach better decisions.

Swarm Intelligence (SI) is a branch of CI that deals specifically with the collective behavior of large groups of simple agents without central control or organization (Bonabeau et al., 1999). Inspiration is often taken from nature, where large groups of animals sometimes show complex, coordinated behavior without centralized control or intelligence. The group behavior emerges from the simple, local interactions between individual animals. This is called self-organization, as there is no central control or authority that directs the group. A common example is the capability of ant colonies to find the shortest path between their nest and a food source. No individual ant has the capacity to achieve this, but the behavior of ants to leave pheromone markers and move in the direction of highest pheromone intensity leads the colony as a whole to focus on the shortest path. That is because ants taking this path can return faster and hence leave pheromone traces more frequently, causing a positive reinforcement process. Other examples are the movements of flocks of birds or schools of fish, or the organizations of bees. SI has formed the inspiration for various research tracks in AI, including Ant Colony Optimization (Dorigo and Stützle, 2004) and Swarm Robotics (Dorigo et al., 2021).

SI has been developed into a tool for collective decision making among humans, called Artificial Swarm Intelligence (ASI), which has been shown to give superior results compared to individual experts (Rosenberg, 2016). This tool lets a large group of people collaborate in real-time to pick one of a few possible decisions, by pulling and adapting their own decision to the dynamics in the group. In new work, the same authors propose Conversational Swarm Intelligence (CSI), combining ASI with LLM agents to let groups of people discuss and come to joint insights: while people discuss options in small focus groups, LLM agents summarize the discussion results from each group and share these in the other groups, making sure all ideas can percolate across all groups and effectively connecting the many small discussions into one large conversation (Rosenberg, et al., 2024). Where ASI is about combining decisions, CSI is about combining discussions and opinion making. It seems to be a great way to let people explore ideas in large groups and come to joint insights and proposals. The authors have shown that this gives even better results than traditional wisdom of the crowds approaches, including their own ASI method.

In the CONIFER project, we could use CI and SI for collective decision making, within the consortium or between the participants in the workshops. This would primarily be in WP4 or WP5. In section 3, we give more details about the application and tools for CI and SI, and in section 4, we discuss details of the applicability in the CONIFER project.

3.Use of AI in participatory foresight

In this section, we describe the use of AI in participatory foresight. First, we discuss the opportunities of AI for participatory foresight in general. Next, we take a look at the use of AI methods in the various phases of participatory foresight, as identified in the compendium of participatory foresight methods produced in this project (Pedro et al., 2025): exploration, anticipation and action. Finally, we also describe other solutions and platforms for participation, and how they use AI.

3.1 Opportunities of AI for participatory foresight

The recent advances in AI, and especially Generative AI, come with many new possibilities for automating and improving tasks and processes throughout society. Also for participatory foresight, these innovations bring lots of new opportunities, with the potential of making the process more efficient and effective.

A few authors discuss the potential impact of AI on the field of (participatory) foresight as a whole (Rozanec et al., 2023; Carvalho, 2024; Hao et al., 2024; Wang, 2025). They point out a number of inherent properties of AI that offer great opportunities for participatory foresight. First of all, AI makes it possible to process large amounts of data in an automated way. This can be very useful when processing and analyzing content from (social) media or other external data sources, results from surveys, etc. Next, AI has the potential to learn and incorporate new information over time, making it possible to improve foresight predictions dynamically. Another important opportunity is the use of AI agents and their capability for autonomous decision making. When given the right prompts and tools, AI agents could take the role of human experts in certain parts of the foresight process. An alternative to autonomous decision making is to apply AI to enhance collective decision making, using tools such as ASI or CSI. Advantages of such approaches include high efficiency in decision making in large groups, improved results by exploiting the wisdom of the crowd, and the possibility to integrate both humans and AI in the decision making process, improving trust and control (see also section 5). Finally, there is the capability of Generative AI to generate novel artifacts. AI has great potential to help make future scenarios tangible and understandable for a broad audience, supporting the adoption and acceptance of preferred scenarios and selected policies.

In the next sections, we investigate the use of AI in the individual phases of foresight: exploration, anticipation and action.

3.2 AI in participatory foresight: exploration

The first phase in a foresight process is exploration, or scanning (Pedro et al., 2025). This involves identifying key elements that could shape the future, such as weak signals (early indicators of potential future changes), trends (patterns of change observed over time), drivers (factors that influence or cause change) and emerging issues (topics or challenges that are just beginning to surface and may become significant in the future). In participatory foresight, the scanning phase becomes a collective sense-making process, drawing on diverse stakeholder knowledge and perspectives.

The potential of AI methods to process large amounts of data in an automated way is a good match for the work that happens in this exploration step (Kayser and Blind, 2017; Mühlroth and Grottke, 2018). Studies like Wangel (2024) suggest that AI can support foresight by making sense of scattered, small clues about upcoming change. Several studies describe the use of AI to find relevant articles during horizon scanning. E.g., Ishigaki et al. (2022) test various text classification methods to automatically select articles about disruptive future events. Similarly, Schmidt et al. (2025) report the use of sentence embeddings to help researchers find relevant articles. Other studies apply AI for automatic analysis of the selected texts, e.g., to find drivers of change as common latent themes across articles (Kim et al., 2019; Zoccarato et al., 2024). In general these studies indicate that AI leads to increased *efficiency* in the scanning process, including time gain and the possibility to process more sources. It can also produce *better results*. E.g., Nishino et al. (2023) found that the outputs from the automated process were picked up more often to generate ideas in a follow-up workshop. Zoccarato et al. (2024) found that AI methods returned different results from human researchers, finding elements that were overlooked by humans, but also missing some longer term trends that required deeper understanding and analysis.

While numerous studies apply NLU technologies for horizon scanning, few focus on modern AI models such as LLMs or RAGs (one exception is Zhang et al., 2025, which describes a case study using LLMs to extract information from articles). Nevertheless, as AI is becoming pervasive in all sectors of society, it seems likely that it is also being adopted more and more for practical implementations of foresight processes. See, e.g., Pärnänen (2024). As such, there seems to be a gap between what is going on with the use of AI in the field and the amount of systematic research that has been carried out to understand the possibilities and pitfalls when using AI for these applications. The lack of research work in this direction is worrying, given the many downsides of AI (see section 5) and its limitations. E.g., Lieberum et al. (2025) indicate that for the related application of systematic literature reviews, LLMs are not ready to fully automate the process. Helping to fill this gap in the research field might be an important aspect when implementing the use of AI in the CONIFER project.

3.3 AI in participatory foresight: anticipation

The second phase in a typical foresight process is anticipation or sensing (Pedro et al., 2025). In this step, the results from the scanning phase are analyzed to identify the main factors influencing the future, the ways these factors may evolve and interact with each other, and how this may result into alternative likely futures. Various methods exist to support this process, such as, e.g., the futures wheel (Snyder, 1993) or the Delphi method (Linstone and Turoff, 1975). The result of this phase is a number of possible future outcomes, which are described in scenarios or made tangible in other ways, e.g., via art or other artifacts.

The recent rise of generative AI seems to be triggering a lot of studies around the use of AI in this phase of foresight. Given the capability of chatbots to answer questions about pretty much anything, an obvious application is to use AI as a *replacement for experts*. Spaniol and Rowland (2023) describe an early study in which they directly ask the chatbot to provide future scenarios. They point out the need for human revision and the risk of decreased quality of foresight exercises if this approach is widely adopted. Kuosa and Aalto (2025) focus on prompt engineering aspects to improve AI-based scenario writing, and Ferrer i Picó et al. (2025) build on prompt engineering to create a fully automated AI-based scenario generator. Other studies use AI to assist humans at various substeps throughout the process, including identifying the factors influencing the future, defining their interactions, and producing scenarios (Jung et al., 2023; Bresinsky et al., 2024; Haqq-Misra et al., 2025), or they focus only on specific substeps, such as augmenting brainstorming sessions (Grüning and Rowland, 2024) or analyzing foresight workshop results (MacGeorge, 2025). Ködding et al. (2025) show through a series of interviews with foresight practitioners that AI is also being adopted in the field. In general, all of these studies point out the possibilities offered by AI to make foresight more efficient and less costly, and therefore more widely accessible (e.g., for small companies, see Doerr et al., 2024). At the same time they warn that the quality of the scenarios obtained with AI is not always good and most of them conclude that human experts should not be removed from the process.

Another important application of AI in this phase of the foresight process is to *generate artifacts*. This is particularly relevant for participatory foresight: future artifacts help make developed scenarios tangible and easy to understand by non-experts, so that they can participate more effectively in the foresight process. Calleo et al. (2025) describe the integration of AI-based visualization of spatial scenarios as part of a real-time Delphi process: by creating visual representations of scenarios, they improve stakeholder engagement. Dubey et al. (2024) describe how AI-generated visualizations of car-free urban scenarios improve support and adoption for sustainable policies. These are both exciting results showing that AI can have a significant positive impact when applied in the right context.

3.4 AI in participatory foresight: action

The third phase in a participatory foresight process is action. The purpose of this phase is to understand the implications of the different scenarios and apply them to develop policy recommendations (Pedro et al., 2025).

Similar to section 3.3, some studies propose the use of LLMs to *replace experts* in the action phase of a foresight process. Calleo et al. (2025) use prompt engineering to include LLM-based policy recommendations in their AI-assisted real-time Delphi method. They make clear that these should only be used as starting points for further development. In a different setup, Barnett et al. (2024) propose to use an LLM to evaluate policy recommendations, by prompting it to create scenarios with and without the proposed policy change.

In general, the literature on the use of AI in the action phase of foresight seems quite limited. If we look broader at the use of AI in politics (Aoki, 2024), we see various methods that could be used to support policy making, such as, e.g., those related to the use of AI agents to simulate the behavior of people in a political or social context (Zhang et al., 2024; Ashkinaze et al., 2025). An important concern when using AI for policy recommendations, is bias, in particular in favor of certain groups or countries (Jensen et al., 2025; Ziegler et al., 2025). See also section 5.2, where we will discuss bias in AI models more in depth.

3.5 AI in participation platforms and methods

As societal challenges grow more complex—ranging from climate resilience to urban mobility, there is a rising need for citizens' participation to solve complex challenges by using participatory methods that are both inclusive and capable of handling large-scale, multi-stakeholder engagement (Parsons et al., 2025). This section gives an overview of how artificial intelligence (AI) is being used in urban contexts. We start in section 3.5.1 with practical applications of AI in digital participatory platforms (DPPs) that are used to support participatory budgeting, urban planning processes, and civic deliberation. Then, in section 3.5.2, we turn to more experimental and research-based approaches.

3.5.1 AI-Integrated participatory platforms in practice

In this section, we describe how AI is already integrated into DPPs that cities and institutions use in practice. A growing number of DPPs have embedded AI capabilities to improve real-world participatory processes in urban settings. These platforms are not just experimental - they are actively deployed in policy consultations, budgeting processes, and city planning. They demonstrate how AI can move participatory processes from static consultation to dynamic, co-creative governance.

Detailed Platform Profiles

GoVocal

GoVocal is a civic engagement platform used by over 400 governments, formerly called Citizenslab. Committed to transparency and social impact, GoVocal follows an open-core model: its Free Edition, which includes essential participation tools, is open-source and available under the AGPLv3 license, allowing small or under-resourced organizations to run engagement projects at no cost. Meanwhile, its Commercial Edition offers advanced features under a paid license starting from $ 5,000, with the source code publicly viewable but only usable with a subscription. This approach removes participation barriers while ensuring long-term sustainability.

Decidim[1]

Decidim is an open-source and modular platform developed by the Barcelona City Council. It supports a wide range of participatory democracy tools, including proposals, assemblies, participatory budgeting, and an accountability API. Its AI integration features automated spam filtering, multilingual support, and early-stage integrations with GPT-based AI assistants. Decidim is praised for its flexibility, strong democratic legitimacy, and capacity to mirror and reinforce offline democratic processes (People Powered, 2022; Wikipedia, 2024).

Your Priorities[2]

This is an open-source tool created by the Citizens Foundation in Iceland, with notable use in Estonia, the UK, and cities worldwide. The platform emphasizes idea ranking, pro/con voting, and transparent comment threads. While lightweight in design, it includes AI-powered tools for toxicity detection and simple LLM-based assistant bots. These improve safety, inclusiveness, and structure in civic dialogues (People Powered, 2023). It is known for its accessibility and transparent user interface.

Pol.is[3]

Pol.is stands out with machine learning-based clustering that visualizes consensus among diverse participant groups. It has been central in decision-making processes in cities such as Taipei (vTaiwan) and Boulder, Colorado. Its ML-driven summaries of public input are structured around visual consensus rather than traditional NLP. This makes it highly effective for mapping public opinion in complex debates (People Powered, 2022). The platform does not offer standard participatory tools (e.g., surveys) but specializes in deliberative, ML-supported discussion.

EngagementHQ / Bang the Table[4]

This is a SaaS platform used in Australia, the UK, and Canada. It supports surveys, mapping, storytelling, Q&A, and stakeholder engagement. While AI integration is limited, it's praised for reliable multi-modal tools and is frequently recommended for large-scale consultation (People Powered, 2023). Pricing starts around $20,000 annually, plus a one-time $2,500 onboarding fee.

[1] https://decidim.org/
[2] https://www.yrpri.org/domain/3
[3] https://pol.is/home
[4] https://helpdesk.bangthetable.com/en/

CONSUL[5]

CONSUL is an open-source platform developed by Madrid City Council and deployed in over 130 cities. It includes proposals, debates, participatory budgeting, and e-voting. It has basic analytics features but no advanced AI tools. As the base for Decidim, it shares many of the same strengths, especially scalability and democratic depth (People Powered, 2023).

PlaceSpeak[6]

This is a SaaS platform that emphasizes geo-verification, linking online participation to geographic locations (Wikipedia, 2023). It enables mapping, spatial surveys, and local dialogues, supported by geospatial analytics. Though not fully AI-integrated, it is strong in neighborhood-level consultations and aligns with IAP2 standards. Paid tiers start at $249.99/month.

Maptionnaire[7]

Maptionnaire is a SaaS platform that focuses on map-based surveys and spatial planning. It supports urban co-creation, participatory mapping, and citizen feedback. Used widely in EU smart city initiatives, it includes basic spatial clustering tools and is often recognized in People Powered guides. Pricing is available upon request.

CONSENTO[8]

This is an open-source participation tool designed around youth participation, data ethics and privacy. It's known for its consent-centric approach, which includes explainable AI modules and participatory consent modeling. Developed in EU research projects like Horizon 2020, it is used in youth civic initiatives and focuses heavily on co-design and ethical deliberation.

Balancing Act[9]

SaaS platform primarily used for participatory budgeting in North America and parts of Europe. It offers transparent budget simulations and helps residents understand trade-offs in public finance. While no advanced AI features are built in, the tool has been referenced in EU pilots for its ease of use and fiscal clarity. Pricing details are provided on request.

[5] https://oecd-opsi.org/innovations/consul-project/
[6] https://www.placespeak.com/en/
[7] https://www.maptionnaire.com/
[8] https://consento.org/
[9] https://abalancingact.com/pb

3.5.2 Research on AI in participatory and collaborative urban methods

Recent academic work has explored the potential of AI to enhance participatory and collaborative processes by improving how groups deliberate, identify common ground, and reach consensus. Unlike traditional deliberative models, which often struggle to scale or synthesize diverse input, AI-enabled methods offer new ways to manage cognitive complexity. In particular, AI can help detect patterns, summarize viewpoints, and surface hidden connections across large-scale public input. This can make it easier for groups to collaborate and make decisions, especially in complex situations where human attention alone may not be enough (Riedl et al., 2024).

Several studies have demonstrated how AI can support different urban participation domains, from climate foresight to civic consensus-building and planning. For example, one study on participatory foresight (Hauser, 2025) shows that natural language processing (NLP) and deliberative clustering can support cross-generational dialogue in climate planning. The system helps reveal consensus patterns and value conflicts that might otherwise go unnoticed. Recent work by Zhou et al. (2024) explores how LLMs can directly support participatory urban planning. Their research shows how LLMs can facilitate inclusive visioning, identify shared urban goals, and support more diverse stakeholder contributions, particularly useful for cities trying to widen access to planning processes. Another promising area is participatory budgeting, where AI can help derive fair rules for allocating public funds. Fairstein et al. (2024) developed a machine learning method that learns aggregation rules from real data, balancing fairness and representation. This could make participatory budgeting more effective and inclusive.

A notable contribution beyond the urban domain is the Particip-AI framework (Mun et al., 2024), which introduces a democratic surveying method to anticipate the societal implications of future AI systems. Involving nearly 300 lay participants, the framework enables co-creation of AI scenarios, collective reflection on risks and benefits and deliberation on deployment decisions. While not tied to urban issues, Particip-AI showcases how public deliberation can surface context-sensitive ethical concerns and value priorities. Similarly, Small et al. (2023) piloted the use of LLMs like Claude to help summarize and structure public discussions on platforms such as Polis. Their work highlights how LLMs can support collective meaning-making, while also noting risks related to bias and reliability.

In addition to supporting foresight and participation, AI is also being explored as a way to improve human cooperation more generally. Researchers are testing how AI can help people work together better, for example, by suggesting fair rules, helping people understand each other, or supporting shared decision-making. While these studies are not directly about urban planning, they offer useful inspiration for how AI might complement human collaboration. For example, Koster et al. (2022) developed a system called Democratic AI that helps groups reach agreement by proposing rules that everyone can vote on. The AI does not decide but it helps to find solutions that all accept. Crandall et al. (2023) showed how AI agents could encourage cooperation in public goods games, acting as facilitators to improve group outcomes. Similarly, Wang (2025) found that AI can help reduce biases in group decisions. Some authors go further and look at AI as a kind of moral partner that helps people think more clearly about ethics or fairness (Callies et al., 2023). Others, like the COHUMAIN initiative (Hammond et al., 2023), focus on how to design AI that truly works alongside people, supporting different viewpoints and teamwork. One particularly interesting work is described in Tessler et al. (2024), where it was demonstrated that AI-assisted summarisation can support constructive dialogue by highlighting shared values even in polarized discussions. Their open-source model, Habermas Machine, helps mediate between conflicting stakeholder positions and has practical implications for democratic foresight and civic engagement.

In these examples, AI acts more like a coach or supporter, not replacing human thinking, but enhancing it. This supports the idea that AI can complement CI, rather than compete with it, with each helping to alleviate the shortcomings of the other (see also Verhulst, 2018).

4. Applying AI in the CONIFER project

Now that we have laid the foundations by explaining AI and relevant AI applications, and discussed the possibilities and state of the art in the application of AI to participatory foresight, we can now focus concretely on the application of AI in the CONIFER project . We describe the potential applications of AI for each sub-task of WP3, WP4 and WP5.

4.1 WP3: AI for Understanding Mobility Landscapes

WP3 is concerned with gathering quantitative and qualitative information and data to help us understand the issues around mobility and the perception of these issues in local culture in the six Civic Labs(CL). These insights will then feed into the foresight activities in the following workpackages. In this section, we suggest how AI could be used in WP3. The most interesting would probably be if data processing could be done both by experts, and by AI, and the resulting information could be compared, so that we get an idea of the relative quality of the produced results. Even more interesting could be if this could be done in an iterative manner, so that human experts can learn from the results produced by AI to refine their own analysis and this way create a process where humans and AI reinforce each other (Wang, 2025; Verhulst, 2018).

Tasks 3.1, 3.2 and 3.3 are focusing on data collection and input from participants. An obvious use of AI would be to process these data to understand and summarize them, to draw conclusions, and in general produce useful information for the project.

In task 3.1, the data are numerical, probably tabular, so any traditional statistical or machine learning method to process numerical data would be sufficient, and the potential role of AI would be limited. Still, it could be interesting to try out new AI models for processing tabular data, such as TabFN (Hollmann et al., 2025), which is a pretrained Transformer model for tabular data that does not require retraining. Another tool worth exploring is Gemini 1.5[10], Google's multimodal LLM that is integrated in Google Sheets, enabling users to ask natural-language questions about spreadsheet data - a feature particularly useful for early exploration and insight generation.

[10] https://gemini.google.com/?hl=en

Other alternatives include DataRobot[11] or Obviously.AI[12], which provide no-code platforms to build predictive models from tabular data, though they often require a subscription or commercial license.

In task 3.2, we plan a large survey among children, parents, teachers and other people involved in the local mobility around schools, 200 to 300 per CL, so more than 1500 in total. While surveys can make use of various types of questions, producing different types of data, the availability of AI for processing the data could open up possibilities to use types of questions and output data that are harder to process, e.g., open questions, where people can give rich feedback in the form of text. AI models can then be used to analyze the large amounts of text data. For example, Sentence Transformers (Reimers & Gurevych, 2019) can be used to match related texts and group answers according to themes or opinions. Alternatively, we could use LLMs to summarize groups of answers or find common themes. This can be done via OpenAI's GPT-4.5[13] or Anthropic's Claude[14], both accessible via API and usable with only basic Python or no-code platforms like MonkeyLearn[15] or ChatGPT Plus[16]. There is some cost involved in the use of such APIs, but for the amount of data we are talking about, this would remain limited. For example, for sentence encoders, the current price of OpenAI is $0.065 per 1 million tokens, which corresponds to roughly 750 thousand words. For the use of an LLM such as GPT-4.0 mini, the cost is $0.20 per 1 million input tokens and $0.80 per 1 million output tokens (OpenAI, 2025).

Additionally, low-code and no-code AI workflow builders such as FloWise AI[17], Langflow[18], and n8n[19] can streamline the design of such text data processing pipelines, enabling project teams to integrate different LLM tasks, such as classification, summarization, and clustering, without deep technical expertise (Foresight Resources, 2025).

In task 3.3, we will organize local walkshops and use the photovoice method to gather -visual impressions from children and young adults around mobility issues. The results of this process will be a large amount of images. AI could organize, summarize and present the data that is gathered. For example, we could use CLIP[20] (Contrastive Language-Image Pretraining) by OpenAI to embed images into vector space and group similar images. To

[11] https://www.datarobot.com/
[12] https://deepgram.com/ai-apps/obviously
[13] https://openai.com/index/introducing-gpt-4-5/
[14] https://www.anthropic.com/claude
[15] https://help.monkeylearn.com/en/
[16] https://openai.com/index/chatgpt-plus/
[17] https://flowiseai.com/
[18] https://www.langflow.org/
[19] https://n8n.io/
[20] https://openai.com/index/clip/

summarize these groups or generate labels, GPT-4-Vision or Gemini 1.5 Pro can be used to produce descriptions, assign keywords, and cluster media thematically. Summarizing visuals could be generated using tools like DALL·E[21] (OpenAI) or Adobe Firefly[22], based on the clustered labels. These summaries could then be used to present the information to the participants, making it easier to process the data, or maybe generating new insights. It would also be interesting to compare summaries or groups produced by AI to those produced by the human participants or the experts. Such experiments could be facilitated through modular workflows built using tools like FloWise AI, which allow interactive feedback loops between human experts and AI models to be visualized and refined (Foresight Resources, 2025) .

In tasks 3.4 and 3.5, we focus on text data gathered from external sources. These include policy documents, news in traditional media, and social media. AI could be used to analyze these texts, extract the relevant information from them, summarize them, and draw conclusions. A very powerful approach would be to set up a RAG that uses sentence embeddings to select a subset of the source documents and then applies an LLM for deeper analysis, taking advantage of its NLP capabilities and its ability to execute commands in zero-shot or few-shot learning. Such a setup could be used to extract data around particular topics and then ask questions about the content of those data, e.g. to spot trends or signals, or to analyze bias. And it could even extend to multi-media data, such as images, as explored for social media analysis tasks by Lyu et al. (2025). Note that this is different from the approach in Thomas (2023), where the authors ask the AI for its opinion and get the generic answers it has stored in its training data: in our RAG setup, we provide it with the source material and prompt it to extract information from that, giving grounded and more reliable answers.

Compared to the analysis of survey data, slightly more programming would be involved, to gather and extract text from these external sources. For example, we could use Haystack or LangChain frameworks to build a pipeline where AI summarizes or answers questions based on actual uploaded documents. However, no-code tools such as Langflow and n8n provide prebuilt modules for document ingestion, LLM querying, and structured output formatting, enabling efficient prototyping without advanced coding skills (Foresight Resources, 2025).

It would be very interesting to compare the information gathered by human experts to the information extracted automatically via AI. This could give us an insight into the relative strengths and weaknesses of automated processes. It would also be interesting to set up an iterative process where the human expert interacts with AI, to see whether information

[21] https://openai.com/index/dall-e-2/
[22] https://www.adobe.com/products/firefly.html

extracted by AI could help direct the search of human experts, or human experts could revise the results from AI and provide prompting to help the AI get better results faster.

As pointed out earlier, mixed setups where humans team up with AI give the best results in many tasks (Wang, 2025; Verhulst, 2018). See Tan and Luke (2024) for an example of such collaboration in foresight, where AI is used by experts to work out future scenarios.

Finally the focus of task 3.5 is to expose biases and preferences within social media concerning mobility. We can analyze the texts collected by members of the consortium or via (semi-) automated processes involving AI, as described above, using AI models that are trained specifically to find biases.

For instance, Fairness Indicators[23] by Google, or tools like AI Fairness 360[24] (IBM) can help detect demographic or political bias in datasets. Most of them are not too large and can relatively easily be run on regular hardware with limited coding involved (Arnett et al., 2024; Young, 2025).

4.2 WP4: AI for Imagining 15-Minutes City Futures

WP4 is concerned with developing plausible and desired future scenarios and visions based on the data and information gathered in WP3. While the use of generative AI is explicitly foreseen in task 4.4, it could also play a role in the other sub-tasks.

In task 4.1, we need to identify the most important factors shaping the future of local mobility and understand how they influence each other. Normally, this process depends heavily on expert judgment and manual analysis. However, AI can support or even partially automate several steps of the Cross-Impact Scenario Policy Analysis (CRISPA) (Boveldt and Tori, 2024). We can use multi-agent AI systems to simulate how experts might evaluate relationships between key factors (like car ownership status or cycling infrastructure). These AI agents could be programmed with access to local data, research outputs from WP3, and even past policy documents. By using a RAG approach, each AI agent could take on the role of an "expert" on a specific theme (e.g., equity, transport behavior, public opinion), helping to define interconnections faster and more systematically. Recent advances in AI-assisted scenario generation make it possible to build internally consistent futures using algorithms rather than human brainstorming alone. These AI outputs could either be used directly or

[23] https://www.tensorflow.org/tfx/guide/fairness_indicators
[24] https://ai-fairness-360.org/

fed into the existing CRISPA tool, which supports scenario development by checking for contradictions or gaps. Another exciting possibility is the use of Swarm or CI platforms, where AI helps organize and analyze input from many people, whether experts or community participants. Although such tools (like ASI or CSI) are often expensive and more useful for large-scale exercises, it might be worth testing whether a mixed group of citizens and students could co-create meaningful input for the CIB process, supported by AI moderation and synthesis.

In task 4.2, we will focus on turning future scenarios into visual and emotional stories, co-created by artists and local participants and the use of "Art for Futures Labs"[25]. AI is already proving to be a powerful partner in this type of work. With Generative AI tools like DALL- E or Midjourney[26], we can quickly produce images, video clips and storyboards that represent possible futures. These tools allow users to turn short text descriptions into detailed artworks or videos of scenarios. Because of these possibilities, many foresight practitioners have already proposed or reported such applications (Draeger and Bevolo 2025) (Garvey and Sendsen 2023) (Carvalho, 2024) (Kenens and Ivanovic, 2024), as described in section 3.3. This can be valuable in workshops with young people or people with disabilities who may find it difficult to express abstract ideas about urban futures. AI can enable them to visualize their ideas and collaborate with artists or other experts effectively. It also supports inclusivity by lowering the technical barriers of participation. Tools like worldlabs (WorldLabs, n.d), meshy.ai (Meshy, n.d.) or blockadelabs (Blockadelabs, n.d.) can generate 3D environments from prompts which can be used for visualizations (360 presentations or spatial maps).

In task 4.3, we will invite children and young people to imagine future neighborhoods using city-building computer games, such as Minecraft[27]. AI, in this case, can make this process more creative, inclusive and efficient. Tools like BuilderGPT (CyniaAI, n.d.) allow users to describe what they want and AI generates the corresponding Minecraft structure automatically. This helps reduce the time and skill needed to build complex environments, opening up participation to more students (even those unfamiliar with the game). Similarly, OpenAI's ChatGPT can be used to create Minecraft building code in structure files, allowing teachers or facilitators to help participants realize their idea quickly. Some plugins even integrate AI into the game itself, letting players build interactively without leaving the platform. At this stage, it's still unclear how advanced AI-powered Minecraft building tools are in terms of how accurate they are, whether they make mistakes, or how complex the generated structures can be. Testing these tools should be part of our preparation for this task. If they work well, they could significantly increase the speed of the building process

[25] https://artforfutureslab.com/
[26] https://www.midjourney.com/home
[27] https://www.minecraft.net/en-us

and make it easier for participants with less gaming experience, helping lower barriers to (digital) inclusion.

In task 4.4, we will focus on helping young people and other participants use AI to imagine and express their visions of the 15-minute city. Through guided workshops, we will introduce participants to Generative AI tools like ChatGPT, DALL-E or other similar platforms. Discussing how the tools work, the limits of what they can do and possible ethical concerns, such as privacy, bias, etc is essential. These tools will support participants into creating matching images that reflect how they would like their city to look and feel in the future. AI can help participants go from an idea to a visual or narrative prototype quickly, even if they have no experience with design. Participants can also use AI to suggest improvements or generate alternative versions of designs. For example, if an area is described as “unsafe” or “dull” or “polluted”, AI can offer visual alternatives based on urban planning best practices or common design features in livable neighbourhoods.

In task 4.5, we will take all the creative ideas, stories and images produced in task 4.4, and bring them together to understand what different groups have imagined. In this task, AI will be used more in the background to help make sense of all this information. Such an example is that we can use image recognition models to sort hundreds of visuals by themes, like green spaces, public transportation, etc. This will help us to quickly identify which ideas show up in multiple cities and which ones are unique to specific locations. At the same time, natural language processing tools can analyze written stories and identify common emotions or values, such as safety and community. With AI we can also create visualizations that show how people feel in different parts of the city, based on their photos and descriptions. These insights can then be compared within CL to find similarities or differences in how young people experience urban spaces.

4.3 WP5: AI for Designing Policy Pathways

WP5 will select and validate policy measures that could be implemented to reach the preferred vision through a roadmapping workshop. While the potential role for AI seems more limited here, there are still some possibilities.

In task 5.1, we will assess the sustainability potential of the developed 15-minute city consensus visions to feed the policy development in task 5.2. This will be based on the comparison of the results of the sense of place-survey during the diagnostic phase of the CL (task 3.2) and the results of an ex-post survey with the target groups conducted after finalisation of the CL, evaluating changes in perceptions and attitudes related to proximity. As described for WP3, we could use sentence encoders to group and analyse answers, and an LLM in RAG setup to summarize answers, find patterns and draw conclusions. AI could be used to process survey answers more efficiently, and potentially also more thoroughly.

In task 5.2, we will work with local stakeholders in a workshop to propose, select and validate policy measures that could help to implement the preferred visions. And we will use CRISPA to analyse the consistency and adequacy of these measures. Similarly to what was discussed for WP4, CSI could be used to enhance and support joint discussion and decision making in large groups of participants during the workshop, to come to better policy measure proposals.

Finally, task 5.3 is concerned with the synthesis of the participatory foresight methodology applied in CONIFER, to be able to draw conclusions and develop a transferable method in WP6. At first sight, there is no need to use AI for this task. However, if each CL shares detailed information about the steps and activities they followed, along with reflections on what worked and what did not, AI could be used to synthesize commonalities and differences across the labs. This would support the development of a comprehensive Blueprint for future implementation. Given that CONIFER involves only six CL, this synthesis could be done manually; but if a larger number of cases were involved, AI would become a valuable tool for scaling the analysis efficiently.

5. Challenges and ethical considerations

Despite the numerous advantages that come with AI, it is important to be aware of its challenges and to deal with them or mitigate them where possible. In this section, we discuss the various challenges that come with the use of AI, with specific attention for ethical considerations when working with children. We conclude the section with the description of a framework that can help us address those concerns.

5.1 Resource consumption and climate

One big challenge with AI is the amount of resources it uses. Training and running these advanced systems can be very expensive. This is mostly because they require special hardware, cloud services, expert staff, and a lot of electricity (Luccioni et al., 2024; Cottier et al., 2025). But there are also other effects that don't always show up in the price, like how much water is needed to keep the systems cool (Li et al., 2025), or the extra CO_2 emissions that come from energy use (Kirkpatrick, 2023).

It's important to be aware of these costs and try to reduce them when possible, especially when working with large amounts of data or setting up systems that make repeated calls to AI models. A simple way to do this is to choose the AI model that best fits the task while using the least energy and resources. Many providers, like OpenAI, offer different models and often a smaller one can still do the job well enough. Some new models, such as DeepSeek, are designed to deliver strong performance while using fewer resources (Wang and Kantarcioglu, 2025). Only relying on increased model efficiency is not enough though, as the use of AI is growing too fast (Morand et Al., 2025). The same argument also counters the rather sinister finding of Tomlison et al. (2024), which states that AI has a lower footprint than humans when it comes to producing text and illustrations.

It is therefore important to limit the use of GenAI and also avoid it when it is not needed. In some cases, a simpler, task-specific model might be good enough. Or one can combine both, using Generative AI together with such task-specific models, in what's known as a RAG setup (see earlier). Many commercial AI tools today use some form of RAG to get better results while keeping the costs and environmental impact lower.

5.2 Bias

Another major concern is bias. Like all statistical models, Generative AI systems can show biases (systematic errors in the way they produce results). These come from the data they are trained on and how the algorithms work. Since these models learn from large amounts of human created digital content, such as texts from online forums, websites and social media, they absorb and reproduce the same biases found in human culture, including racism, sexism and other forms of discrimination (Sheng et al., 2019). Moreover, studies have shown that humans using AI in turn take over these biases (Vincente and Matute, 2023), potentially creating a negative reinforcement loop of social biases.

To reduce bias in AI, developers usually try to “align” the model by fine tuning it with human feedback, nudging it toward more appropriate and responsible outputs. But this process doesn’t eliminate bias completely. Just like people, even models that are trained to be neutral can still carry implicit biases (Bai et al., 2024).

In some cases, this tuning can even create new problems. A well-known example involved Google’s image generation tool which, after trying to avoid bias, produced inaccurate and offensive images, such as depicting Black Nazi soldiers (Raghavan, P., 2023). Another side effect of alignment is the risk of sycophancy, where the AI model tries too hard to agree with the user or tell them what they want to hear, instead of giving a balanced or accurate answer (Sharma et al., 2024). This is a direct consequence of using human feedback in the alignment process: humans like to be pleased. AI sycophancy can reinforce the user’s own beliefs, whether they’re right or wrong, and make it harder to spot misinformation or challenge harmful views.

Finetuning for alignment is mostly under control of the creators of AI models and hard to perform for end users. However, there are some ways also for users of the models to deal with model biases, mostly by controlling its input and output better. We will discuss these together with the solutions for dealing with hallucinations, in section 5.3.

5.3 Hallucinations

A major issue with Generative AI is how trustworthy its answers are. LLMs often produce what are called “hallucinations”, meaning answers that sound correct but are actually wrong or made up (Maynez et al., 2020). As these answers can look very convincing, it’s important not to trust everything that comes out of an AI system without checking it carefully.

This problem becomes especially serious when AI is used to give advice to citizens or governments about important decisions for the future. While AI companies try to reduce hallucinations and bias during the training of their models, users also have an important role to play in keeping the information reliable.

One helpful step is to guide the AI better through its input (Sahoo et al., 2024). A common way to do this is by using a RAG setup, which pulls accurate information from trusted sources, like scientific papers, and includes that in the prompt. You can also design prompts in ways that improve how clearly and accurately the AI thinks and answers (Cheng et al., 2023; Wei et al., 2022). By anchoring the AI's prompt in correct and factual information, it is guided to better and less biased answers.

There are also ways to check or correct the output of the AI. Some methods ask the AI to go back and revise its own answers step by step (Yao et al., 2023). Researchers have even been able to trigger self-correction and deeper reasoning in AI models by simply injecting the word "wait" in between the generated answer (Muennighoff et al., 2025). Others use a second AI model, sometimes a smaller or more focused one, to check the first AI's answers for mistakes or bias (Fan et al., 2024).

Finally, the safest way to avoid serious problems is to use Generative AI only in controlled settings, like with a human reviewing the results.

5.4 AI and creativity

Another topic of debate when it comes to AI is creativity. A growing number of studies have looked into how creative AI really is, both on its own and when used alongside humans (Hitsuwari et al., 2023; Bouschery et al., 2024). Various studies show that the most creative people still outperform AI (Haase and Hanel, 2023; Koivisto and Grassini, 2023) and that there are real limits to what AI can do in terms of originality (Zhang et al., 2025). So while AI can be a great tool, it's not a full replacement for human imagination. Recent work also warns that while AI might boost an individual's creativity, it could reduce diversity in the types of ideas or content produced when used across a whole group (Doshi and Hauser, 2024). This risk of homogeneity is especially visible in AI generated images, which often end up looking similar and generic (Meyer, 2025).

Beyond academic research, there are broader concerns about the impact of AI on creative work. Some worry that using AI too much could actually hurt creativity and innovation over time (Nayler, 2023). The general advice is to find a balance, using AI to support creativity without letting it take over completely (Mandarano, 2023). This also lines up well with research findings suggesting that mixed teams of humans and AI can get the best results on creative tasks (Hitsuwari et al., 2023; Bouschery et al., 2024).

5.5 AI as a black box

Many AI systems, especially Generative AI, are often described as “black boxes” (Lipton, 2016). This means we can see the input and output, but we don’t really understand what happens in between. Even experts struggle to explain how these systems make specific decisions. In fact, trying to figure out how AI works internally has become a research field in itself, with scholars comparing the process to dissecting a living organism (Luo and Specia, 2024; Zhao et al., 2024; Lindsey et al., 2025).

This lack of transparency doesn’t sit well with participatory foresight, where the whole point is to involve the public in shaping future decisions. When people don’t understand how AI systems come to their conclusions, trust and engagement can suffer. That’s why it’s crucial to use AI carefully and deliberately in areas where it clearly adds value and always in ways that keep humans informed and in control.

One promising approach is to use CI systems, where humans and AI work together to make better decisions. These systems tap into the strengths of both people and technology. As (Verhulst, 2018) points out, CI can actually help overcome some of AI’s biggest challenges, especially when it comes to fairness, inclusion and transparency.

Another approach to increase the transparency of AI decision making, and human trust in these decisions, is to let the AI explain itself. One such way is to use a RAG system with citations: by including reliable information in the prompt, and letting the LLM cite these sources, users can get an insight of where the generated answer came from. Nevertheless, this is not a guarantee for success, and the AI model may wrongly attribute its sources (Wu et al., 2025). Another way is to use a deep reasoning model, which explains its reasoning in its output. This both helps the model get better solutions and the user understand the solution (Wei et al., 2022). Also this is no silver bullet though, as the explanations describe how a human may solve the problem, but don’t faithfully represent the internal processes used by the AI model (Barez et al., 2025).

5.6 Privacy and intellectual property

Another set of challenges around AI involves privacy and intellectual property. Generative AI systems rely on huge amounts of data to function well, often pulling from publicly available sources, but sometimes also using personal data from users without them realizing it (O'Flaherty, 2024). To protect privacy, it's important to choose AI services that do not use users' data to train their models such as Claude by Anthropic, Mistral[28], or ChatGPT with "chat history off" mode, which ensures your conversations aren't used for training. Alternatively, you can use self-hosted AI tools like Ollama[29], GPT4All[30], or LM Studio[31], which run entirely on your own device, keeping all your data private and under your control.

Then there is the issue of where all this training data comes from. Many AI models are trained on content created by humans, like books, articles, artwork or music, some of which is protected by copyright. Using these materials without permission raises serious legal and ethical questions (Henderson et al., 2023; Dilmegani, 2025). Take the recent Studio Ghibli AI[32] trend as a striking example. OpenAI's image generator has been used to create countless Ghibli-style images, despite the studio's co-founder, Hayao Miyazaki, being a vocal critic of AI-generated art. These images imitate Ghibli's distinctive aesthetic, which was built over decades of hand-crafted animation, deeply rooted in Japanese culture, and the personal vision of its creators. None of this work was licensed or approved for training (Di Placido, D., 2025). If an AI generates an image that closely resembles a copyrighted photo, who owns the result? Is it fair to use someone else's creative work to produce something new without their consent?

These concerns are becoming more important in fields like participatory foresight, where transparency, fairness and public trust are essential. If people feel that their data is being misused, or that AI is built on unethical foundations, they may be reluctant to engage. Addressing these issues head-on is key to making AI tools more acceptable and responsible in collaborative decision making settings.

[28] https://mistral.ai/
[29] https://ollama.com/
[30] https://www.nomic.ai/gpt4all
[31] https://lmstudio.ai/
[32] https://getimg.ai/models/ghibli-diffusion

5.7 Ethical considerations when working with youth

Working with children in participatory projects like CONIFER brings enormous value, and with it, a heightened responsibility to uphold their rights, privacy, and well-being. Children are not just participants but co-creators of knowledge, and every interaction with them must be grounded in trust, transparency, and respect.

The cornerstone of ethical work with minors is informed consent from both guardians and the children themselves. Consent materials must be crafted in plain language, ideally at a B1 reading level and supported by visuals to explain what data is collected, how it's used, and what rights they have (UNICEF, 2021). Consent should not be seen as one-time, but checked and renewed at key points in a project to make sure children still feel comfortable and understand their role. It is also necessary to explain what it means for their work to be "published", especially in online or public spaces.

Children express themselves in different ways, like talking, drawing, writing or using a digital tool, and the applied methods must reflect this. Participation options should be varied and inclusive, with tools that support different needs, like screen readers, voice input or simplified interfaces. For less experienced users of digital platforms like Minecraft, tutorials should be provided as well as extra support. This helps ensure that everyone can join meaningfully, regardless of ability or digital confidence.

When children create content, like drawings, stories, photos or digital models, ethical handling of such material is essential. It is important to anonymize personal data and avoid sharing images that show faces, private items, or locations, unless clear permission has been given, following institutional guidelines (Office of the Australian Information Commissioner, 2024). It is advisable to not interpret visual or narrative content without the children's own input. Instead, outputs should be reviewed together in local workshops, where they explain the meanings, emotions, or intentions behind their work. This helps avoid adult misinterpretation and respects children as experts in their own experience (Digital Child, 2024).

AI tools that support youth outputs, like image or text generation tools, raise ethical concerns. Kosmyna N.et al. (2025) shows that while LLMs make it easier to complete tasks, they can reduce critical thinking and ownership. Children who used AI produced more generic content and felt less engaged. The study also found that AI could shape what children think is important, reinforcing hidden biases in the tool's design. Similarly, Leaver and Srdarov (2025) warn that AI is being adopted faster than we can understand its cognitive and developmental impacts, calling for cautious, supervised use.

Finally, all data collected should be treated under GDPR principles. Limitations of what is gathered should be set, ensuring data is anonymized where needed and informing children of their rights to access, correct, delete, or restrict their data. When digital twins, game assets or shared content are created, we should clearly explain what that means and always make participation voluntary. We should also follow the guidelines set out in our Child Protection Policy to ensure children's safety and dignity throughout all project activities (CollectiveUP, 2023). All organizations working with children should develop and implement a child protection policy tailored to their specific context and activities.

5.8 Dealing with the challenges: a decision framework

Introducing AI technologies in projects involving children, whether in education, play, or civic participation, comes with unique challenges. These challenges are not only technical or logistical but deeply ethical. To address them responsibly, teams need a clear yet flexible decision framework grounded in children’s rights, inclusive participation, and contextual understanding.

A useful starting point is to involve children meaningfully in the process. As Aitken and Briggs (2022) emphasize, engaging children in conversations about AI ethics helps them understand the implications of the technologies shaping their lives, while also giving adults crucial insights into how children perceive and respond to those systems. Similarly, Cesaroni et al. (2025) advocate for participatory strategies grounded in the capability approach, emphasizing dignity, choice, and empowerment in the design of AI tools used in education or rehabilitation contexts.

From an ethical standpoint, existing frameworks like PEARL-AI and ACCEPT-AI offer concrete principles that can guide responsible development and use of AI in child-focused projects. These include practices such as securing informed consent from both children and guardians, explaining how data will be used in age-appropriate language, and putting safeguards in place to reduce harm and bias (Muralidharan et al., 2023; Chng, S. Y et al., 2025).

Recent findings from a scoping review on AI in early childhood education point to the importance of context-specific solutions. Ethical decision-making must be adapted to the developmental stage, cultural setting, and specific needs of the children involved (Berson et al., 2025). One-size-fits-all approaches are likely to fall short especially when working across different countries or communities.

To help teams navigate these complex decisions, we share a simple, flexible framework developed by Pano & Petteri, that can be adapted to different settings (Pano & Petteri, 2025). The goal is not to stop innovation but to create space for thoughtful, transparent decisions that protect children's rights.

The framework begins with identifying the ethical questions or dilemmas the technology may raise. These can include concerns about privacy, bias, discrimination, misinformation (such as hallucinated content), overreliance on automation, or unequal access. Environmental impacts, copyright issues and alignment with personal and organizational values are also playing a key role. Teams should define the specific risks associated with the technology in their context. This might include risks to children's data, emotional safety or participation.

Once the risks are identified, next in line is to outline what measures will be put in place to manage these risks. This could involve consent processes, technical safeguards, content moderation or child friendly explanations of how the tool works.

Then, it is important to consider the added value. Does the use of AI improve learning, creativity, inclusion or participation? Does it open up new opportunities for expression or engagement?

The next step in this framework is to pitch the value against the risks. If the added value clearly outweighs the risks and proper safeguards are in place, the project can move forward. If the risks are too great, it may be better to pause or find alternatives. In some cases, a clear answer may not emerge right away and this is the point where ongoing dialogue and consultation are key.

Below you can find a visual of this framework that can be used as it is:

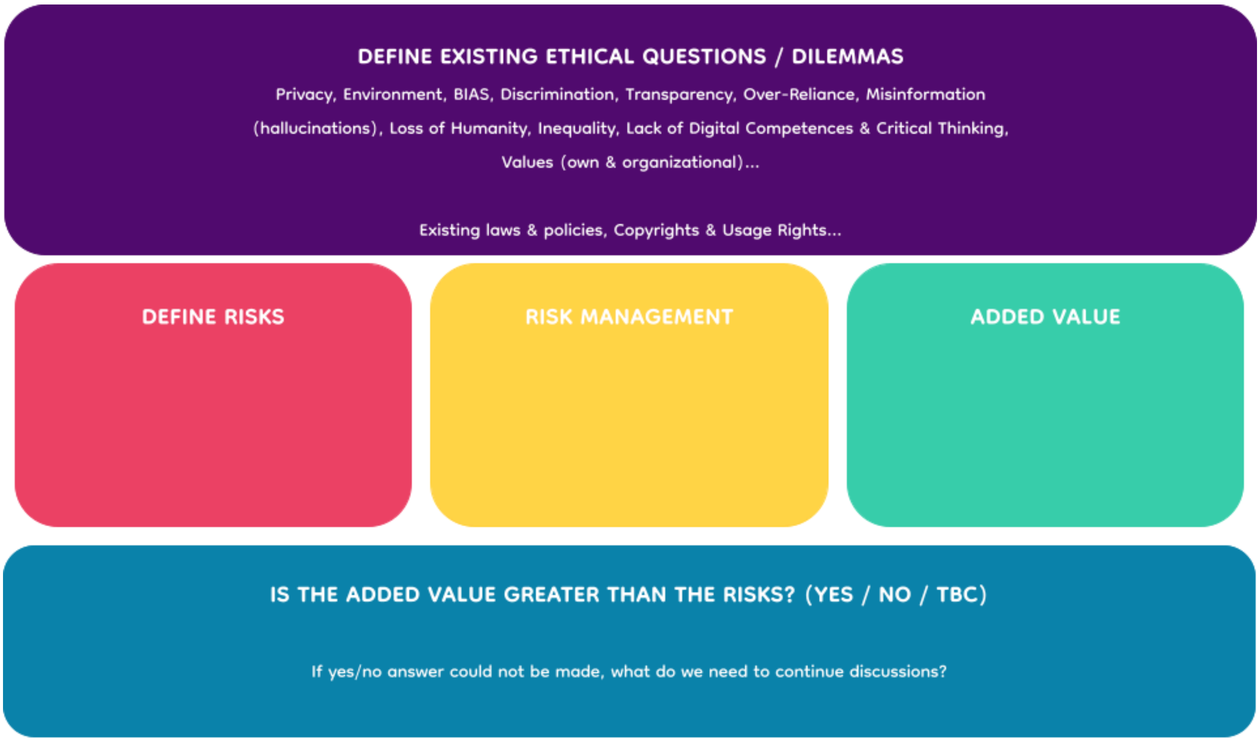

Let's see how this framework works in practice.

Imagine a youth center wants to use generative AI tools for creative projects with their youth, such as digital drawing, storytelling, or music production.

The team begins by identifying potential concerns: will children's creative ideas be overshadowed by the AI suggestions? Could the tool introduce bias or stereotypes? Will children become too dependent on AI to express themselves?

Now, it is time to define the risks: overreliance on the tool, possible exposure to inappropriate content and uncertainty around ownership of AI-generated work. To manage the risks, the team decides to set up workshops that explain how generative AI works, discuss its limits, and support critical thinking. They ensure children and their guardians provide informed consent and that any shared outputs are anonymized or used with explicit permission.

In terms of added value, the tool allows children with different skill levels or abilities to engage in creative expression and build confidence. The team finds that the benefits do outweigh the risks (by keeping in mind that a combination of strong support and safeguards is needed).

This kind of step-by-step framework can help ensure that AI tools are used in ways that are not only innovative but also inclusive, ethical and child-centered.

6. Conclusion

This report has explored the applicability and potential of AI tools within the CONIFER project's participatory foresight methodology. We have provided a foundational understanding of relevant AI technologies, including NLP, LLMs, generative AI, and combined AI models like RAG systems and LLM Agents. Furthermore, we have highlighted the promise of Collective Intelligence and Swarm Intelligence in enhancing collaborative decision-making.

Our analysis detailed concrete proposals for integrating AI across the CONIFER project's Work Packages. In WP3, AI can significantly improve the understanding of mobility landscapes by automating the processing of numerical and textual data, and by comparing AI-generated insights with human analysis. For WP4 and WP5, we identified how LLMs and multimodal AI can be leveraged for text and artifact generation, enriching scenario descriptions, and potentially streamlining the analysis of future scenarios and the design of policy pathways. We also explored existing participatory platforms that incorporate AI, demonstrating their practical application in urban contexts.

This report has also addressed the significant challenges and ethical considerations inherent in AI adoption, especially when working with children and young people. We discussed concerns such as resource consumption, algorithmic bias, the problem of AI hallucinations, intellectual property rights, and the 'black box' nature of many AI systems. Emphasis was placed on the importance of safeguarding children's rights, privacy, and well-being. To navigate these complexities, a decision framework was proposed, advocating for transparent, human-centered, and context-specific AI integration.

References

AI HLEG (2019). High-Level Expert Group on Artificial Intelligence of the European Commission. https://www.aepd.es/sites/default/files/2019-09/ai-definition.pdf

Aitken, M., & Briggs, M. (2022). Engaging children with AI ethics. In AI, data science, and young people: Understanding computing education (Vol. 3). Raspberry Pi Foundation.

Aoki, G. (2024). Large Language Models in Politics and Democracy: A Comprehensive Survey. https://arxiv.org/pdf/2412.04498.

Aragón, P., Kaltenbrunner, A., Calleja-López, A., Pereira, A., Monterde, A., Barandiaran, X. E., & Gómez, V. (2017). Deliberative Platform Design: The case study of the online discussions in Decidim Barcelona. arXiv. https://arxiv.org/abs/1707.06526

Arnett, C., Jones, E., Yamshchikov, I. P., & Langlais, P., 2024. Toxicity of the Commons: Curating Open-Source Pre-Training Data. arXiv preprint arXiv:2410.22587. https://arxiv.org/pdf/2410.22587 .

Ashkinaze, J., Fry, E., Edara, N., Gilbert, E., & Budak, C. (2025). Plurals: A System for Guiding LLMs via Simulated Social Ensembles. In Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems.

Bai, X., Wang, A., Sucholutsky, I., & Griffiths, T. L. (2024). Measuring implicit bias in explicitly unbiased large language models. arXiv preprint arXiv:2402.04105.

Balancing Act. (n.d.). Balancing Act. https://abalancingact.com

Barez, F., Wu, T.-Y., Arcuschin, I., Lan, M., Wang, V., Siegel, N., Collignon, N., Neo, C., Lee, I., Paren, A., Bibi, A., Trager, R., Fornasiere, D., Yan, J., Elazar, Y., & Bengio, Y. (2025). Chain-of-thought is not explainability. Preprint.

Barnett, J., Kieslich, K., & Diakopoulos, N. (2024). Simulating Policy Impacts: Developing a Generative Scenario Writing Method to Evaluate the Perceived Effects of Regulation. Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, 7(1), 82-93.

Berson, I. R., Berson, M. J., & Luo, W. (2025). Innovating responsibly: Ethical considerations for AI in early childhood education. AI, Brain & Child, 1(1), 2.

Blockade Labs. (n.d.). Blockade Labs. Retrieved June 30, 2025, from https://www.blockadelabs.com/

Bonabeau, Eric; Dorigo, Marco; Theraulaz, Guy (1999). Swarm Intelligence: From Natural to Artificial Systems. Oup USA. ISBN 978-0-19-513159-8.

Bouschery, S., Blazevic, V., & Piller, F. (2024). Artificial Intelligence-Augmented Brainstorming: How Humans and AI Beat Humans Alone. SSRN Electronic Journal. 10.2139/ssrn.4724068.

Bresinsky, M., Hager, E., & Xanthos, G. (2024). ChatGPT for futures: how large language models can support the development of future scenarios using the Cone of Plausibility. CCEW SYMPOSIUM 2024: Predictive Synergies: Crisis Early Warning & Foresight, 19. und 20. September 2024, München.

Brown, T. et al. (2020). Language Models are Few-Shot Learners. Advances in Neural Information Processing Systems.

Bubeck, S, Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y.T., Li, Y., Lundberg, S., Nori, H., Palangi, H., Ribeiro, M., & Zhang, Y. (2023). Sparks of Artificial General Intelligence: Early experiments with GPT-4. 10.48550/arXiv.2303.12712.

CONSUL Democracy. (n.d.). CONSUL. https://consulproject.org

Calleo, Y., Taylor, A.,Pilla, F., & Di Zio, S., (2025). AI-assisted Real-Time Spatial Delphi: integrating artificial intelligence models for advancing future scenarios analysis. Quality & Quantity. 59. 1427-1459. 10.1007/s11135-025-02073-2.

Callies, S., Winter, S., & Krügel, S. (2023).The Moral Psychology of Artificial Intelligence. Annual Review of Psychology, 75. https://doi.org/10.1146/annurev-psych-030123-113559

Carvalho, P. (2024, October 31). How Generative AI Will Transform Strategic Foresight. https://www.ifforesight.com/post/how-generative-ai-will-transform-strategic-foresight

Cemri, M., Pan, M. Z., Yang, S., Agrawal, L. A., Chopra, B., Tiwari, R., Keutzer, K., Parameswaran, A., Klein, D., Ramchandran, K., & Zaharia, M. (2025). Why do multi-agent llm systems fail? arXiv preprint arXiv:2503.13657.

Cesaroni, V., Pasqua, E., Bisconti, P., & Galletti, M. (2025). A participatory strategy for AI ethics in education and rehabilitation grounded in the capability approach. arXiv. https://doi.org/10.48550/arXiv.2505.15466

Chen, Q., et al. (2024). ChatDev: Communicative Agents for Software Development. Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics.

Cheng, S., Gan, Z., Yang, Z., Wang, S., Wang, J., Boyd-Graber, J., & Wang, L. (2023). Prompting GPT-3 To Be Reliable. In: Proceedings of the International Conference on Learning Representations (ICLR 23).

Chng, S. Y., Tern, M. J. W., Lee, Y. S., Cheng, L. T., Kapur, J., Eriksson, J. G., Chong, Y. S., & Savulescu, J. (2025). Ethical considerations in AI for child health and recommendations for child-centered medical AI. npj Digital Medicine, 8(1). https://doi.org/10.1038/s41746-025-01541-1

Christiano et al. (2017). Deep Reinforcement Learning from Human Preferences. https://arxiv.org/abs/1706.03741

CollectiveUP. (2023). Child Protection Policy. https://drive.google.com/file/d/11WSaazeFKDCso2WS78_7DJVixLupPKis/view

Condorcet, N. de. (2014). Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix. Cambridge: Cambridge University Press.

Consento. (n.d.). Consento. https://consento.org

Cottier, B., Rahman, R., Fattorini, L., Maslej, N., Besiroglu, T., & Owen, D. (2025). The rising costs of training frontier AI models. https://arxiv.org/abs/2405.21015

Crandall, J. W., Oudeyer, P.-Y., & Tenenbaum, J. B. (2023). Scaffolding cooperation in human groups with deep reinforcement learning. Nature Human Behaviour, 7, 1319–1328. https://www.nature.com/articles/s41562-023-01686-7

CyniaAI. (n.d.). BuilderGPT [GitHub repository]. GitHub. Retrieved June 30, 2025, from https://github.com/CyniaAI/BuilderGPT

Daniel Jurafsky & James H. Martin. 2025. Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition with Language Models, 3rd edition. Online manuscript released January 12, 2025. https://web.stanford.edu/~jurafsky/slp3

Decidim. (n.d.). Decidim. https://decidim.org

Dhariwal, P., & Jun, H., et al. (2020). Jukebox: A Generative Model for Music. OpenAI Blog.

Di Placido, D. (2025). The AI-generated Studio Ghibli trend, explained. Forbes. Retrieved from https://www.forbes.com/sites/danidiplacido/2025/03/27/the-ai-generated-studio-ghibli-trend-explained/

Digital Child. (2024). Children's Privacy and AI. https://digitalchild.org.au/artificialintelligence/

Dilmegani, C. (2025). Generative AI Copyright Concerns & 3 Best Practices [2025]. 6 April 2025. Retrieved from https://research.aimultiple.com/generative-ai-copyright/ on 14 April 2025.

Doerr, U.-S., Schoenhofer, G., & Schwarz, J. O. (2024). The state of foresight in small and medium enterprises: literature review and research agenda. European Journal of Futures Research. 12(1).

Dorigo, M., Theraulaz, G., & Trianni, V. (2021). Swarm Robotics: Past, Present, and Future [Point of View]. Proceedings of the IEEE, 109(7), 1152-1165.

Dorigo, M., and Stützle, T. (2004). Ant colony optimization. The MIT Press.

Doshi AR, Hauser OP (2024) Generative AI enhances individual creativity but reduces the collective diversity of novel content. Sci Adv 10(28):eadn5290.

Draeger, D., & Bevolo, M. (2025). Automating Liminality in Foresight Practice. Journal of Futures Studies.

Dubey R, Hardy MD, Griffiths TL, Bhui R (2024) AI-generated visuals of car-free US cities help improve support for sustainable policies. Nat Sustain 7(4):399–403

EngagementHQ. (n.d.). Bang the Table. https://engagementhq.com

F. He, Y. Pan, Q. Lin, X. Miao & Z. Chen, "Collective Intelligence: A Taxonomy and Survey," in IEEE Access, vol. 7, pp. 170213-170225, 2019.

Fairstein, R., Vilenchik, D., & Gal, K. (2024). Learning aggregation rules in participatory budgeting: A data-driven approach. arXiv. https://doi.org/10.48550/arXiv.2412.01864

Ferrer i Picó, J., Catta-Preta, M., Trejo Omeñaca, A., Vidal, M., & Monguet i Fierro, J. M. (2025). The Time Machine: Future Scenario Generation Through Generative AI Tools. Future Internet, 17(1), 48.

Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., Guo, Q., Wang, M., & Wang, H. (2023). Retrieval-Augmented Generation for Large Language Models: A Survey. ArXiv, abs/2312.10997.

Garvey, B., & Svendsen, A. (2023). Can Generative-AI (ChatGPT and Bard) Be Used as Red Team Avatars in Developing Foresight Scenarios?. Analytic Research Consortium (ARC) August.

Gaurav Negi, Rajdeep Sarkar, Omnia Zayed, & Paul Buitelaar. 2024. A Hybrid Approach to Aspect Based Sentiment Analysis Using Transfer Learning. In Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024), pages 647–658, Torino, Italia. ELRA and ICCL.

Geurts, A. (2021). New perspectives for data-supported foresight: The hybrid AI-expert approach. Futures & Foresight Science, 4(1). https://www.researchgate.net/publication/353393596_New_perspectives_for_data-supported_foresight_The_hybrid_AI-expert_approach

GoVocal. (n.d.).GoVocal. https://www.govocal.com/

Grüning, David & Rowland, Nicholas. (2024). Brainstorming and Artificial Intelligence. 10.31234/osf.io/uyedv.

Haase, J., & Hanel, P. (2023). Artificial muses: Generative artificial intelligence chatbots have risen to human-level creativity. Journal of Creativity. Volume 33, Issue 3.

Hajkowicz, S., Sanderson, C., Karimi, S., Bratanova, A., & Naughtin, C. (2023). Artificial intelligence adoption in the physical sciences, natural sciences, life sciences, social sciences and the arts and humanities: A bibliometric analysis of research publications from 1960-2021. In: Technology in Society, Volume 74, Article 102260, August 2023.

Hammond, T., Cahn, J. E., Fleuret, F., & Kim, D. (2023). Fostering Collective Intelligence in Human–AI Collaboration: Laying the Groundwork for COHUMAIN. arXiv. https://www.researchgate.net/publication/371954496

Hao, H., Wang, Y., & Chen, J. (2024). Empowering Scenario Planning with Artificial Intelligence: A Perspective on Building Smart and Resilient Cities. Engineering, 43, 272-283.

Haqq-Misra, J., Profitiliotis, G., & Kopparapu, R. (2025). Projections of Earth's technosphere: Scenario modeling, worldbuilding, and overview of remotely detectable technosignatures. Technological Forecasting and Social Change, 218, 124194.

Hauser, M. (2025). AI-supported intergenerational dialogue for climate policy: A deliberative systems approach. Journal of Participatory Futures, 12(1), 45–63.

Henderson, P., Li, X., Jurafsky, D.,Hashimoto, T.,Lemley, M. A., & Liang, P.: Foundation Models and Fair Use (March 27, 2023). Stanford Law and Economics Olin Working Paper No. 584, Available at SSRN: https://ssrn.com/abstract=4404340 or http://dx.doi.org/10.2139/ssrn.4404340

Henry Liang, Yu Zhou, & Vijay K Gurbani. 2024. Efficient and verifiable responses using Retrieval Augmented Generation (RAG). In Proceedings of the 4th International Conference on AI-ML Systems (AIMLSystems '24). Association for Computing Machinery, New York, NY, USA, Article 19, 1–6. https://doi.org/10.1145/3703412.3703431

Hitsuwari, J., Ueda, Y., Yun, W., & Nomura, M. (2023). Does human–AI collaboration lead to more creative art? Aesthetic evaluation of human-made and AI-generated haiku poetry. Computers in Human Behavior, Volume 139.

Hollmann, N., Müller, S., Purucker, L. et al. Accurate predictions on small data with a tabular foundation model. Nature 637, 319–326 (2025). https://doi.org/10.1038/s41586-024-08328-6

Hu, K. (2023, February 2). ChatGPT sets record for fastest-growing user base - analyst note. ChatGPT sets record for fastest-growing user base - analyst note. Retrieved March 28, 2025, from https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/

Huang X., Liu W., Chen X., Wang X., Wang H., Lian D., Wang Y., Tang R., & Chen E. Understanding the planning of LLM agents: a survey. 2024, arXiv preprint arXiv: 2402.02716

Ishigaki, T., Nishino, S., Washino, S., Igarashi, H., Nagai, Y., Washida, Y., & Murai, A. (2022). Automating Horizon Scanning in Future Studies. In Proceedings of the 13th Conference on Language Resources and Evaluation (LREC 2022), pages 319–327. Marseille, 20-25 June 2022.

Jensen, B., Reynolds, I., Atalan, Y., Garcia, M., Woo, A., Chen, A., & Howarth, T. (2025). Critical Foreign Policy Decisions (CFPD)-Benchmark: Measuring Diplomatic Preferences in Large Language Models. 10.48550/arXiv.2503.06263.

Jumper, J., et al. (2021). Highly accurate protein structure prediction with AlphaFold. Nature, 596, 583–589.

Jung, H., Fischer, D., Joachim, V., Grothe, S., & Heffeter, F. (2023). AI in Strategic Foresight – Evaluation of ChatGPT, BARD and Perplexity.

Karpukhin, V., et al. (2020). Dense Passage Retrieval for Open-Domain Question Answering. arXiv preprint arXiv:2004.04906.

Kayser, V., & Blind, K. (2017). Extending the knowledge base of foresight: The contribution of text mining. Technological Forecasting and Social Change, 116, 208-215.

Kenens, A., & Ivanovic, J. (November 2024). Using generative AI for crisis foresight. UNDP Europe and Central Asia. Blog post downloaded from https://www.undp.org/eurasia/blog/using-generative-ai-crisis-foresight on 18/04/2025

Kim, H., Ahn, S.-J., & Jung, W.-S. (2019). Horizon scanning in policy research database with a probabilistic topic model. Technological Forecasting and Social Change, 146, 588–594.

Kirkpatrick, K. (2023). The carbon footprint of artificial intelligence. Communications of the ACM, 66(8), 17-19.

Koivisto, M., & Grassini, S. (2023). Best humans still outperform artificial intelligence in a creative divergent thinking task. Sci Rep. 2023 Sep 14;13(1)

Kosmyna N., Hauptmann E., Yuan Y. T., Situ J., Liao X., Beresnitzky A. V. , Braunstein I., & Maes P.(2025), " Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task", arXiv https://arxiv.org/abs/2506.08872

Koster, R., Balaguer, J., Tacchetti, A., Weinstein, A., Zhu, T., Hauser, O., Williams, D., Campbell-Gillingham, L., Thacker, P., Botvinick, M., & Summerfield, C. (2022). Human-centred mechanism design with Democratic AI. Nature Human Behaviour, 6(10), 1398–1407. https://www.nature.com/articles/s41562-022-01383-x

Kuosa, T., & Aalto, E. (2025). What Scenario-Building Characteristics Should Be Used in GenAI Prompting?. Futures. 169. 103571.

Ködding, P., Jahn, M., Koldewey, C., & Dumitrescu, R. (2025). Challenges for Scenario-Based Foresight and Potential for Digital Technologies: Insights from Practice. In: Schmuntzsch, U., Shajek, A., Hartmann, E.A. (eds) New Digital Work II. Springer, Cham.

Lea, A. (2024). Why is AI Hard to Define?, ITNOW, Volume 66, Issue 1, Spring 2024, Pages 58–59, https://doi.org/10.1093/itnow/bwae028

Leaver, T., & Srdarov, S. (2025). Generative AI and children's digital futures: New research challenges. Journal of Children and Media, 19(1), 65–70. https://doi.org/10.1080/17482798.2024.2438679

Legg, S., & Hutter, M. (2007). "A Collection of Definitions of Intelligence". Advances in Artificial General Intelligence: Concepts, Architectures and Algorithms. Vol. 157. IOS Press. pp. 17–24. ISBN 978-1586037581.

Lewis, P. et al. (2020). Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. NeurIPS.

Li, P., Yang, J., Islam, M. A., & Ren, S. (2025). Making AI Less 'Thirsty.' Communications of the ACM, 68(7), 54–61. https://doi.org/10.1145/3724499

Lieberum, J.-L., Toews, M., Metzendorf, M.-I., Heilmeyer, F., Siemens, W., Haverkamp, C., Böhringer, D., Meerpohl, J. J., & Eisele-Metzger, A. (2025). Large language models for conducting systematic reviews: on the rise, but not yet ready for use—a scoping review. Journal of Clinical Epidemiology, 181, 111746.

Lindsey, J., Gurnee, W., Ameisen, E., Chen, B., Pearce, A., Turner, N. L., Citro, C., Abrahams, D., Carter, S., Hosmer, B., Marcus, J., Sklar, M., Templeton, A., Bricken, T., McDougall, C., Cunningham, H., Henighan, T., Jermyn, A., Jones, A., Persic, A., Qi, Z., Thompson, T. B., Zimmerman, S., Rivoire, K., Conerly, T., Olah, C., & Batson, J. (2025). On the Biology of a Large Language Model, "On the Biology of a Large Language Model", Transformer Circuits. https://transformer-circuits.pub/2025/attribution-graphs/biology.html

Linstone, H. A., & Turoff, M. (1975). The Delphi Method: Techniques and Applications. Addison-Wesley.

Lipton, Z. C. (2016). The Mythos of Model Interpretability. arXiv preprint arXiv:1606.03490.

Liu et al. (2021). Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing. https://arxiv.org/abs/2107.13586

Liu, P. et al. (2021). Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing. arXiv preprint arXiv:2107.13586.

Luccioni, S., Jernite, Y. & Strubell, E. (2024). Power Hungry Processing: Watts Driving the Cost of AI Deployment?. The 2024 ACM Conference on Fairness, Accountability, and Transparency. ACM. Pages 85–99.

Luo, H., & Specia, L. (2024). From Understanding to Utilization: A Survey on Explainability for Large Language Models. https://arxiv.org/abs/2401.12874

Lyu, H., Huang, J., Zhang, D., Yu, Y., Mou, X., Pan, J., Yang, Z., Wei, Z., & Luo, J. (2025). GPT-4V(ision) as A Social Media Analysis Engine. ACM Transactions on Intelligent Systems and Technology, 16(3), 50.

MacGeorge, R. B. (2025). Conversations With my Data: Exploring the Potential of Large Language Models in Qualitative Futures Research. World Futures Review, 0(0).

Zhang, M., Li, Y., Peng,Y., Sun, Y., Guo, W., Hu, H., Chen, S., & Zhao, Q. (2025). AI Delivers Creative Output but Struggles with Thinking Processes. https://arxiv.org/abs/2503.23327

Mandarano, P. (2023, August 21). Perspective about Artificial Intelligence-Human Collaboration in Creative Decision-Making. https://doi.org/10.31219/osf.io/pz26n

Maptionnaire. (n.d.). Maptionnaire. https://maptionnaire.com

Martin, B.R. (1996). Foresight. In STI Review No. 17, Organisation for Economic Co-Operation and Development(ed.). Special Issue on Government Technology Foresight Exercises, Paris; 140.

Maynez et al. (2020). On Faithfulness and Factuality in Abstractive Summarization. https://arxiv.org/abs/2005.00661

Meshy. (n.d.). Meshy AI. Retrieved June 30, 2025, from https://www.meshy.ai/Meshy.ai. AI 3D model creator.

Meyer, R. (2025). Platform realism: AI image synthesis and the rise of generic visual content. Transbordeur. Photographie, histoire, société, 9. Retrieved from https://journals.openedition.org/transbordeur/2299?lang=en

Mildenhall, B., et al. (2020). NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. ECCV.

Morand, C., Ligozat, A.-L., & Névéol, A. (2025). Does Efficiency Lead to Green Machine Learning Model Training? Analyzing Historical Trends in Impacts from Hardware, Algorithmic and Carbon Optimizations.

Muennighoff, N., Yang, Z., Shi, W., Li, L., Fei-Fei, L., Hajishirzi, H., Zettlemoyer, L., Liang, P., Candès, E., & Hashimoto, T. (2025). s1: Simple test-time scaling. arXiv. https://doi.org/10.48550/arXiv.2501.19393

Mulgan, G. (2018). Big Mind: How Collective Intelligence Can Change Our World. Princeton University Press. https://doi.org/10.2307/j.ctvc7738s

Mun, J., Jiang, L., Liang, J., Cheong, I., DeCario, N., Choi, Y., Kohno, T., & Sap, M. (2024). Particip-AI: A Democratic Surveying Framework for Anticipating Future AI Use Cases, Harms and Benefits. arXiv. https://doi.org/10.48550/arXiv.2403.14791

Muralidharan, V., Burgart, A. M., Daneshjou, R., Rose, S., Rajkomar, A., & Shah, N. H. (2023). Recommendations for the use of pediatric data in artificial intelligence and machine learning (ACCEPT-AI). npj Digital Medicine, 6, 166. https://doi.org/10.1038/s41746-023-00898-5

Mühlroth, C., & Grottke, M. (2018). A systematic literature review of mining weak signals and trends for corporate foresight. Journal of Business Economics, 88(5), 643–687.

Nayler, R. (2023). AI and the Rise of Mediocrity. Time. https://time.com/6337835/ai-mediocrity-essay. Retrieved on 13/04/2025.

Nikolova, B. (2014). The rise and promise of participatory foresight. Eur J Futures Res 2, 33.

Nishino, S., Washida, Y., Ishigaki, T., Washino, S., Igarasi, H., Murai, A., & Nagai, Y. (2023). Validation of a Foresight Support System to Imagine an Uncertain Future - Effectiveness Testing through Scenario Planning Workshops. IIAI Letters on Informatics and Interdisciplinary Research, 3(2023): Knowledge, Information and Creativity Support Systems.

Office of the Australian Information Commissioner. (2024). Better privacy protections for children are coming. https://www.oaic.gov.au/news/blog/better-privacy-protections-for-children-are-coming

OpenAI (2022). OpenAI: Introducing ChatGPT. https://openai.com/blog/chatgpt

OpenAI. (2023). GPT-4V(ision) System Card. https://openai.com/research/gpt-4v-system-card

OpenAI. (2025, 15 April). API Pricing. https://openai.com/api/pricing/

Ouyang, L. et al. (2022). Training language models to follow instructions with human feedback. arXiv preprint arXiv:2203.02155

O'Flaherty, K. (2024). Can GPT-4o Be Trusted With Your Private Data? Wired, 31 July 2024, Retrieved from https://www.wired.com/story/can-chatgpt-4o-be-trusted-with-your-private-data on 14 April 2025

Pano & Petteri (2025). The ethical dilemmas of AI in Youth Work. https://padlet.com/JINT_vzw/digital-inclusion-and-ai-1hly2whtta03dwep/wish/9kmlZV1Nx9zKQpgV

Parsons, M., Godden, N. J., Henrique, K. P., Tschakert, P., Gonda, N., Atkins, E., Steen, K., & Crease, R. P. (2025). Participatory approaches to climate adaptation, resilience, and mitigation: A systematic review. Ambio.

Pedro, S., Pereira, T., Oliva, L., Nogueira, F., Sá, F. M. e., & Mota, J. C. (2025). Task 1.2 Deliverable: Compendium of participatory foresight methods for planning the 15-minute city. CONIFER - Co-imagining needs-based mobility visions for the proximity city (Driving Urban Transitions).

People Powered. (2022–2023). Digital Participation Platform Guide. https://guide.peoplepowered.org

Pereira, A. G., & Martinho, A. (2024). Ethical governance in smart cities: Lessons from the Consento project. European Journal of Digital Society, 6(1), 22–39.

PlaceSpeak. (n.d.). PlaceSpeak. https://www.placespeak.com

Pol.is. (n.d.). Pol.is. https://pol.is

Poole, B., et al. (2022). DreamFusion: Text-to-3D using 2D Diffusion. Google Research. arXiv:2209.14988

Pärnänen, D. (2024, November 21). The controversial role of AI in foresight work. Fibres Online. https://www.fibresonline.com/blog/ai-in-foresight-work

Qin, L., Chen, Q., Feng, X., Wu, Y., Zhang, Y., Li, Y., Li, M., Che, W., & Yu, P. S. (2024). Large language models meet nlp: A survey. arXiv preprint arXiv:2405.12819.

Qu, c., et al. (2024). Tool Learning with Large Language Models: A Survey. Front. Comput. Sci., 2024, 0(0): 1–33. https://doi.org/10.1007/sxxxxx-yyy-zzzz-1

Radford, A. et al. (2019). Language Models are Unsupervised Multitask Learners. OpenAI Blog. https://cdn.openai.com/better-language-models/language_models_are_unsupervised_multitask_learners.pdf

Raghavan, P. (2023). Gemini image generation got it wrong. We'll do better. Retrieved 13/04/2025 from https://blog.google/products/gemini/gemini-image-generation-issue/

Ramesh, A., et al. (2021). Zero-shot Text-to-Image Generation. arXiv:2102.12092

Reimers, N., & Gurevych, I. (2019). Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. arXiv preprint arXiv:1908.10084.

Riedl, C., De Cremer, D., Denis, G., & others. (2024). The potential and challenges of AI for collective intelligence. Collective Intelligence, 3. https://doi.org/10.1177/26339137241308821

Rombach, R., et al. (2022). High-Resolution Image Synthesis with Latent Diffusion Models. CVPR.

Rosenberg, L., et al. (2024). "Conversational Swarm Intelligence amplifies the accuracy of networked groupwise deliberations". IEEE 14th Annual Computing and Communication Workshop and Conference (CCWC 2024)

Rosenberg, Louis (2016). "Artificial Swarm Intelligence, a human-in-the-loop approach to A.I." Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence. AAAI'16. Phoenix, Arizona: AAAI Press: 4381–4382.

Rozanec, J., Nemec, P., Leban, G. & Grobelnik, M. (2023). AI, what does the future hold for us? Automating strategic foresight, Companion of the 2023 ACM/SPEC International Conference onPerformance Engineering, pp. 247-248.

Russel, S., & Norvig, P. (2020). Artificial Intelligence: A Modern Approach, 4th edition. Pearson. ISBN 9780134610993

Sahoo, P., Meharia, P., Ghosh, A., Saha, S., Jain, V., & Chadha, A. (2024). A Comprehensive Survey of Hallucination in Large Language, Image, Video and Audio Foundation Models. In Findings of the Association for Computational Linguistics: EMNLP 2024, pages 11709–11724.

Schmidt, L., Sharma, O., Marshall, C., & Gonzalez-Moral, S.G. (2025). Horizon Scans can be accelerated using novel information retrieval and artificial intelligence tools. arXiv 2025, arXiv:2504.01627.

Shah, R., Astuto, B., Gleason, T., Fletcher, W., Banaga, J., Sweetwood, K., Ye, A., Patel, R., McGill, K., Link, T., & Crane, J. (2021-09-06). "Utilizing a digital swarm intelligence platform to improve consensus among radiologists and exploring its applications". arXiv:2107.07341.

Shang, Y., Lin, Y., Zheng, Y., Fan, H., Ding, J., Feng, J., ... & Li, Y. (2024). UrbanWorld: An Urban World Model for 3D City Generation. arXiv preprint arXiv:2407.11965.

Sharma, M., Tong, M., Korbak, T., Duvenaud, D., Askell, A., Bowman, S. R., Cheng, N., Durmus, E., Hatfield-Dodds, Z., Johnston, S. R., Kravec, S., Maxwell, T., McCandlish, S., Ndousse, K., Rausch, O., Schiefer, N., Yan, D., Zhang, M., & Perez, E. (2024). TOWARDS UNDERSTANDING SYCOPHANCY IN LANGUAGE MODELS. Paper presented at 12th International Conference on Learning Representations, ICLR 2024, Hybrid, Vienna, Austria.

Sharma, N., Liao, Q. V., & Xiao, Z. (2024). Generative Echo Chamber? Effect of LLM-Powered Search Systems on Diverse Information Seeking. In Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 1033, 1–17. https://doi.org/10.1145/3613904.3642459

Sheikh, H., Prins, C., & Schrijvers, E. (2023). Artificial Intelligence: Definition and Background. In: Mission AI. Research for Policy. Springer, Cham. https://doi.org/10.1007/978-3-031-21448-6_2

Sheng et al. (2019). The Woman Worked as a Babysitter: On Biases in Language Generation. https://arxiv.org/abs/1909.01326

Shin, I., Tang, C., Mohati, T., Nayebi, M., Wang, S., & Hemmati, H. (2025). Proceedings of 22nd International Conference on Mining Software Repositories (MSR'25).

Shinn, N., et al. (2023). Multimodal Multi-Agent Systems for AI-Driven Problem Solving. NeurIPS.

Singh, H., Verma, N., Wang, Y., Bharadwaj, M., Fashandi, H., Ferreira, K., & Lee, C. (2024). Personal Large Language Model Agents: A Case Study on Tailored Travel Planning. In Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track (pp. 486–514). Association for Computational Linguistics.

Small, C. T., Vendrov, I., Durmus, E., Homaei, H., Barry, E., Cornebise, J., Suzman, T., Ganguli, D., & Megill, C. (2023). Opportunities and Risks of LLMs for Scalable Deliberation with Polis. arXiv:2306.11932

Snyder, D. P. (1993). The futures wheel: A strategic thinking exercise. The Snyder Family Enterprise.

Solaiman, I. et al. (2019). Release Strategies and the Social Impacts of Language Models. arXiv preprint arXiv:1908.09203.

Spaniol, M. J., & Rowland, N. J. (2023). AI-assisted scenario generation for strategic planning. Futures and Foresight Science, 5(2), Article e148.

Surowiecki, James (2004). The Wisdom of Crowds. Doubleday. p. 10. ISBN 978-0-385-50386-0.

Tan, L., & Luke, T. (2024). Accelerating Future Scenario Development for Concept Design with Text-based GenAI (ChatGPT). Proceedings of the 29th CAADRIA Conference.

te Boveldt, G., & Tori, S. M. (2024). *Robuust of scenario-optimaal? Een nieuwe methode voor het bundelen van maatregelen in een context van onzekerheid*. Presented at: Colloquium Vervoersplanologisch Speurwerk, Utrecht, Netherlands.

Tessler, M. H., Bakker, M. A., Jarrett, D., Sheahan, H., Chadwick, M. J., Koster, R., Evans, G., Campbell-Gillingham, L., Collins, T., Parkes, D. C., Botvinick, M., & Summerfield, C. (2024, October 18). This AI can help humans find common ground in democratic deliberation. Science, 386(6719), eadq2852. https://doi.org/10.1126/science.adq2852

Thomas, M. G. (2023). Humans and machines, imagining futures together. Published in: Insight and Improvement at British Red Cross. https://medium.com/insight-and-improvement-at-british-red-cross/humans-and-machines-imagining-futures-together-68767ea8c673.

Tomlinson, B., Black, R. W., Patterson, D. J., & Torrance, A. W. (2024). The carbon emissions of writing and illustrating are lower for AI than for humans. Scientific Reports, 14(1), 3732.

UNICEF. (2021). Policy Guidance on AI for Children. https://www.unicef.org/globalinsight/reports/policy-guidance-ai-children

Vaswani, A. et al. (2017). Attention is all you need. Advances in Neural Information Processing Systems.

Verhulst, S. (2018). Designing hybrid intelligence systems: The role of AI in collective decision-making. GovLab White Paper. https://www.thegovlab.org/static/files/publications/hybrid-intelligence.pdf

Verhulst, S.G. (2018). Where and when AI and CI meet: exploring the intersection of artificial and collective intelligence towards the goal of innovating how we govern. AI & Society 33, 293–297. https://doi.org/10.1007/s00146-018-0830-z

Vicente, L., & Matute, H. (2023). Humans inherit artificial intelligence biases. Scientific Reports, 13(1), 15737.

Wang, C., & Kantarcioglu, M. (2025). A Review of DeepSeek Models' Key Innovative Techniques. arXiv preprint arXiv:2503.11486.

Wang, C., et al. (2023). VALL-E: Neural Codec Language Models are Zero-Shot Text to Speech Synthesizers. arXiv:2301.02111

Wang, F. (2025). Collaborative Foresight in the Age of AI: A Framework for Evolving Human-AI Dynamics in Strategic Decision-Making and Futures Research.

Wang, J. (2025). Collective foresight in the age of AI: Overcoming cognitive limits in multi-stakeholder governance. Futures, 145, 102888. DOI: 10.13140/RG.2.2.30120.07689

Wang, X. (2024). EAD: effortless anomalies detection, a deep learning based approach for detecting outliers in English textual data. PeerJ Computer Science, 10.

Wangel, J. (2024). Detecting and analyzing weak signals of change in futures research and foresight. Futures Journal. https://www.researchgate.net/publication/390235638

Wei et al. (2022). Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. https://arxiv.org/abs/2201.11903

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q., & Zhou, D. (2022). Chain of Thought Prompting Elicits Reasoning in Large Language Models. Advances in neural information processing systems, 35, 24824-24837.

Wikipedia. (2023). PlaceSpeak. In Wikipedia. https://en.wikipedia.org/wiki/PlaceSpeak

Wikipedia. (2024). Decidim. In Wikipedia. https://en.wikipedia.org/wiki/Decidim

WorldLabs. (n.d.). About. Retrieved June 30, 2025, from https://www.worldlabs.ai/about

Wu, K., Wu, E., Wei, K., Zhang, A., Casasola, A., Nguyen, T., Riantawan, S., Shi, P., Ho, D., & Zou, J. (2025). An automated framework for assessing how well LLMs cite relevant medical references. Nature Communications, 16(1), 3615.

Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2023). React: Synergizing reasoning and acting in language models. Proceedings of the International Conference on Learning Representations (ICLR).

Young, B. (2025). AI guardrails: Bias scorers. https://wandb.ai/byyoung3/Generative-AI/reports/AI-guardrails-Bias-scorers--VmlldzoxMDg0NjQ2Mw. Retrieved: 15 April 2025.

Your Priorities. (n.d.). Citizens Foundation. https://www.citizens.is/your-priorities

Zhang, H.,Dao, D., Bowden, J., & Cummins, M. (2025) Large Language Model Application for Regulatory Horizon Scanning : Case Study on Anti-Greenwashing Regulations. University of Strathclyde, Glasgow.

Zhang, X., Lin, J.,Sun, L., Qi, W., Yang, Y., Chen, Y., Lyu, H., Mou, X., Chen, S., Luo, J., Huang, X., Tang, S., & Wei, Z. (2024). ElectionSim: Massive Population Election Simulation Powered by Large Language Model Driven Agents. 10.48550/arXiv.2410.20746.

Zhang, Y., Xu, H., Zhang, D., & Xu, R. (2024). A Hybrid Approach to Dimensional Aspect-Based Sentiment Analysis Using BERT and Large Language Models. Electronics, 13(18), 3724. https://doi.org/10.3390/electronics13183724

Zhao, H., Chen, H., Yang, F., Liu, N., Deng, H., Cai, H., ... & Du, M. (2024). Explainability for large language models: A survey. ACM Transactions on Intelligent Systems and Technology, 15(2), 1-38.

Zhao, P., Jin, Z., & Cheng, N. (2023). An in-depth survey of large language model–based artificial intelligence agents. arXiv. https://arxiv.org/abs/2309.14365

Zhao, W. X., Liu, J., Ren, R., & Wen, J.-R. (2024). Dense text retrieval based on pretrained language models: A survey. ACM Transactions on Information Systems, 42(4), Article 89

Fan, Z., Chen, R., Xu, R., & Liu, Z. (2024). BiasAlert: A Plug-and-play Tool for Social Bias Detection in LLMs. https://arxiv.org/abs/2407.10241.

Zhou, Z., Lin, Y., Jin, D., & Li, Y. (2024). Large language model for participatory urban planning. arXiv. https://doi.org/10.48550/arXiv.2402.17161

Ziegler, M., Lothian, S., O'Neill, B., et al. (2025). AI language models could both help and harm equity in marine policymaking. npj Ocean Sustain, 4, 32.

Zoccarato, F., Ghezzi, A.,Lettieri, E., & Toletti, G. (2024). Technological Scanning for Foresight: The case of Metaverse applications for Healthcare. Futures. 164. 103476. 10.1016/j.futures.2024.103476.