BRIEFING

Policy Foresight Analysis

European Parliament

Augmented foresight

The transformative power of generative AI for anticipatory governance

SUMMARY

This briefing explores the potential of generative AI in supporting foresight analysis and strategic decision-making. Recent technological developments promise an increased role for large language models (LLMs) in policy research and analysis. From identifying trends and weak signals to fleshing out rich scenario narratives and bringing them to life in experiential and immersive ways, generative AI is empowering foresight analysts in their endeavour to anticipate uncertainties and support policymakers in preparing better for the future. As generative agents powered by LLMs become more adept at mimicking human behaviour, they could offer foresight practitioners and policy analysts new ways to gain additional insights at greater speed and scale, supporting their work.

However, to effectively integrate generative AI and LLMs into foresight practice, it is crucial to critically evaluate their limitations and biases. Human oversight and expertise are essential for ensuring the reliability and validity of AI-generated outputs, as well as the need for transparency, accountability, and other ethical considerations. It is important to note that, while generative AI can augment human capabilities, it should not be seen as a replacement for human involvement and judgment.

By combining human expertise with generative AI capabilities, foresight analysts can uncover new opportunities to enhance strategic planning in policymaking. A proactive and informed approach to adopting generative AI in foresight analysis may lead to more informed, nuanced, and effective strategies when dealing with complex futures.

IN THIS BRIEFING

EPRS | European Parliamentary Research Service

Author: Lucia Vesnic-Alujevic with Salvatore d'Ambrosio

Policy Foresight Unit

PE 774.665 – July 2025

EN

Introduction

The rapid evolution of generative AI (GenAI)[1] and large language models (LLMs)[2] is transforming many jobs. It is also impacting how public policy, research and interconnected science-for-policy endeavours are done. Traditional policymaking, and related scientific and foresight advice, combines quantitative and qualitative insights with expert input. This approach can become more responsive and data-driven with the growing availability of advanced GenAI tools.

LLMs have proven effective in various tasks at the science-policy interface. For instance, models like ChatGPT have assisted policy analysts and researchers in everyday tasks such as organising information, checking grammar and ensuring consistency. As these models improve, their potential expands to include complex analytical tasks, such as data processing, synthesis, and interpretation, as well as conceptual ideation. However, ensuring the responsible, ethical, and transparent use of GenAI, along with its accuracy and accountability, remains a significant challenge for research and similar purposes.

The efficacy of LLMs in supporting policy analysts and, subsequently, policymakers in exploring future developments and informing policy discussions has gained attention in both policy and academic circles. Research suggests that these applications could support foresight practitioners throughout various stages of an entire project.

Within the realm of GenAI and foresight, it is essential to reflect on potential applications of LLMs and assess recent innovations in generative agents' capacities, as they may constitute a valuable complementary approach to conventional methods, offering enhanced efficiency and cost-effectiveness.[3]

The definition of artificial intelligence according to the EU AI Act

According to Article 3 of the AI Act, artificial intelligence is defined as 'a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments'.

Generative AI applications in foresight

Foresight is the discipline of exploring, anticipating and shaping the future. By using participatory exercises and collective intelligence, it can draw useful insights for strategic planning and policymaking. Foresight includes, among other things, the identification of trends and emerging issues, building scenarios and developing strategic visions and related pathways. Its goal is to inform present decisions, and shape the future that we, as a society, want.

The proliferation of data sources and the unprecedented speed and scale at which data are produced today represent a significant opportunity, including for foresight practitioners. As shown further on, the progressive integration of LLM-powered approaches across foresight frameworks demonstrates positive outcomes, enhancing both strategic insights and analytical capacities. Their ability to extract valuable insights from a vast amount of diverse data sources at speed promises the enhanced efficiency and effectiveness that organisations (e.g. political institutions) demand in times of increased uncertainty, complexity, novelty and turbulence. This section focuses on two fundamental foresight methods: horizon scanning and scenario planning. Recent studies and applications, especially in the private sector, have specifically focused on these two methods, using GenAI and LLMs to enhance speed, efficiency, and sometimes even creativity.[4]

Horizon scanning

Horizon scanning is a foresight method used to identify early signals of change ('weak signals') and emerging trends that could potentially become important in the future. It usually consists of several steps, including the collection of signals (via desk research, literature review of academic and grey sources, media analysis, including digital media, as well as participant input), and analysis and sense making of these signals. However, the ever-increasing scale of unstructured online data is driving the need for faster, broader, and more frequent scans.

An example of data-driven horizon scanning that has been successfully used for several years is the JRC TIM tool. This tool uses text mining to collect signals from academic publications (Scopus) and patents (European Patent Office PATSTAT) that can show novel developments. GenAI can further enhance this large-scale process, by auto-curating, clustering and classifying the content while updating it in real time. As a result, weak signals and trends that would potentially be missed can be more easily identified.

There is a growing body of research exploring the potential of machine learning (including LLMs) to efficiently manage the scanning of large volumes of unstructured data and streamline literature reviews and evidence synthesis. For example, a study developed a model trained to identify news articles likely to contain 'future signals', using a dataset of articles previously labelled as relevant by foresight experts. Another study focused on automating the retrieval and ranking of horizon-scanning articles. This approach significantly reduced the need for manual review, leading to the system successfully retrieving most of the articles (95 %) that a human analyst would have considered important during a scan.

Integrating LLMs into horizon scanning processes can enhance efficiency. This is achieved by accelerating data processing and enabling real-time trend monitoring through automated signal detection and analytical interpretation. By leveraging LLMs, foresight practitioners can focus on sense-making activities and more creative parts of the process.

Initial implementations of AI-automated horizon scanning within private sector and foresight consultancies also demonstrate encouraging preliminary results (e.g. the Foresight Strategy Cockpit and the Futures Platform).

Automating horizon scanning with LLMs: Step-by-step guide

1. **Design smart prompts**: Ask the LLM to search for *weak signals* – early signs of change with future impact.
2. **Feed fresh info**: Overcome the model's knowledge cut-off by uploading news, reports, or real-time web content.
3. **Use online data sources smartly**: Connect LLMs to live sources while browsing to keep insights up to date.
4. **Focus scanning**: Guide the model towards future-relevant domains (e.g. climate, tech shifts).
5. **Iterate often**: Refine prompts and data to sharpen foresight over time.

Scenario planning

Scenario planning is a foresight method used to support long-term planning by helping decision-makers understand trends and uncertainties and act upon them. Through the development of different plausible futures, scenario planning contributes to identifying strategic options across these futures, as well as assessing potential risks and opportunities that different futures might bring. By doing so, more resilient strategies and more informed and adaptable choices can be made.

The process usually consists of desk research and participatory exercises (including workshops, interviews, and surveys). Developing a robust set of scenarios can be both time- and resource-intensive, as well as methodologically challenging. Several software solutions (e.g. ScMI, Parmenides Eidos) already help practitioners overcome some of these limitations. However, GenAI is now able to support them throughout the entire scenario planning process in a more sophisticated way and at greater speed and scale. It offers improvements in areas such as:

- The identification of key trends and drivers.
- The development of scenario narratives. This includes assisting with refining scenarios and checking for consistency, considering complex interactions among factors and the need for scenario plausibility.
- Providing guidance for creating strategy and policy options.

Scenario building with LLMs: Step-by-step guide[5]

1. **Frame the question**: Set the scope and purpose – what future challenge or topic are you exploring? Identify user, purpose and main applications.
2. **Generate drivers**: Use an LLM to surface relevant trends, uncertainties, and change forces from the research corpus. Refine prompts to steer the model. More specific input at this stage leads to more relevant and useful scenarios.
3. **Mix and match**: Identify most relevant drivers and combine them into distinct, plausible scenarios.
4. **Draft narratives**: LLMs can rapidly draft narratives based on your input and ideas.
5. **Validate and iterate**: Refine scenarios with feedback, real-world context, or new data. Check for consistency.
6. **Develop policy implications and options**: Based on a developed set of scenarios inserted as part of a prompt, use LLM to assist in developing policy implications and options.
7. **Human-led shaping**: Choose at each step whether to co-create, steer, or fully author the outputs.

Using multiple LLMs together (an AI swarm) can strengthen and enrich the machine outputs. Even so-called 'hallucinations', errors generated by GenAI models, can positively impact the process of scenario building, sparking creativity and leading to unexpected hypotheses. This is important in a domain such as foresight, where imagination is essential. Hallucinations can produce creatively disruptive scenarios that stretch the limits of plausible futures. However, the mechanisms by which they are triggered and the ways their use can be formalised are still not well understood.

Creativity in GenAI is also shown through its power to create images and videos that bring future scenarios to life. Tools like Midjourney, Sora, and Runway can turn foresight narratives into immersive visuals, helping stakeholders to better grasp abstract or long-term changes. Paired with technologies such as Extended Reality (XR), including Augmented Reality (AR) and Virtual Reality (VR), these tools could also potentially simulate future worlds, thereby turning strategic planning into an experiential journey.

An LLM-augmented process could, therefore, accelerate strategic foresight and scenario planning (although it is not yet capable of fully automating existing approaches), making it more accessible and reducing costs. This is especially important for smaller entities that do not have sufficient capacity to engage with large-scale foresight projects. However, the efficacy and success of AI-generated scenarios rely on human expertise in guiding and interpreting the outputs, underscoring the necessity of a hybrid human-AI approach. A dynamic interplay between human insights and AI-generated content, in which they continuously learn from and refine each other, can indeed lead to more comprehensive and innovative future scenarios and to more dynamic foresight processes. In scenario generation using current LLMs, well-crafted, detailed prompts are key to guiding the AI to produce diverse and relevant futures.

Overview of generative agents

Research shows that LLMs can simulate expert perspectives through generative agents. These AI-powered tools can produce methodologically informed and relatively low-cost responses that may be similar to those of human participants in certain contexts. Generative agents may signal a potential shift in approaches to social science research, including foresight. However, while opening up new opportunities, they also pose important questions about their validity and appropriate use.

Case study 1 – A hypothetical use of GenAI agents in a project

Imagine the following hypothetical situation often faced by policy analysts and foresight practitioners:

As part of your project, you have compiled a list of 100 subject-matter experts with diverse demographic backgrounds, in order to gather 20 expert insights through a participatory workshop, focus group or survey. Yet, two weeks later, you have received only 15 replies from experts agreeing to take part in your study. With deadlines looming and resources depleted, traditional remedies offer no clear path forward.

You have heard that emerging technologies offer promising alternatives. LLMs can simulate expert perspectives through generative agents. These AI-powered tools can produce methodologically informed, low-cost responses that research suggests may approximate those of human experts in certain contexts.

To complement your sample, you decide to use five generative agents tailored to the specific requirements of your study. Armed with your LLM of choice, refined prompts, and ideal expert bio inputs, your synthetic agent is just a click away – so is your interview.

Eventually, you complete your project on time.

***

However, the following questions emerge: How robust is such a foresight process? How useful could such results be for policymakers, for whom you are doing the study? How accurate are they? And what are the ethical implications of such an approach?

The potential to use these agents as substitutes for human participants has been studied so far in, for example, human-computer interaction, marketing, economics, and human behaviour. In these fields, researchers have examined the potential of generative agents to simulate human responses to various prompts and stimuli (e.g. images), demonstrating that these agents can meet key accuracy and validity criteria.[6]

An increasing amount of research supports the development and deployment of generative agents. Studies suggest that since LLMs can encode a broad range of human behaviours from their training data, LLM-powered agents can create realistic responses. When provided with a narrowly defined context, these agents can mimic persons and their behaviour, generating responses that are remarkably realistic. So far, agents have demonstrated their ability to replicate existing economics experiments, store and retrieve factual knowledge and generate synthetic data surveys. To achieve this, structured prompts including examples to demonstrate the desired response patterns are used to encourage step-by-step reasoning and guide the model's response.[7] These approaches are particularly effective for simulating context-specific behaviours, such as how a fictional 'troll' might respond to a social media post or how a robot might navigate a room. Researchers are working to increase agents' sophistication and multi-dimensional characteristics, providing them with personal memory banks they can access at any time.

Notable applications include doctor-patient interactions in hospital settings, macroeconomic activity modelling, and classroom-based educational simulations. A recent newspaper commentary states that the UK Prime Minister used synthetic users to poll the electorate. This indicates that the technology has already found its way into the political arena.

However, human behaviour is multi-faceted and guided by goals and social cues that are difficult to model by rigid simulations. Only agents able to recall past experiences, reflect on their memories to form generalisations, and generate responses that are coherent in the moment and long-term could successfully capture the full complexity, including unpredictability, of human behaviour. In our foresight work, this would also mean being able to replicate fears, hopes and expectations of the future, as well as creativity and intuition that humans might have in connection to a specific topic.

The tested accuracy of generative agents based on interview

A recent study tested the accuracy of generative agents in simulating human behaviour. They sampled 1 052 real individuals, stratified across multiple demographic dimensions, and conducted two-hour audio interviews with them. The interviews served as input to train the models in a way that reflected each participant's unique perspective and life experience. After the model training, the agents' ability to predict behaviour was evaluated by comparing human and generative agents' answers in an opinion and attitudes questionnaire, a personality questionnaire, and behavioural games (e.g. the prisoner's dilemma). The accuracy between responses given by humans and agents varied between 66 % and 85 %. The same study compared model training based on qualitative interviews with two other ways of model training (summarising demographic characteristics of human participants or a paragraph summarising a person's profile) and outperformed them in all aspects.

Agents can also contribute to and enhance engagement with scenarios by creating more vivid and experiential contexts. For example, generative AI can bring traditional scenario narratives to life via 'persona' narratives, stories of fictitious but realistic characters. When prompted with the scenario story, generative agents can respond to questions in real time with answers grounded in it. This helps scenario users 'immerse' themselves in and better understand scenario narratives.

Moreover, generative agents have already found applications in the private sector. For example, some platforms are offering *synthetic consumers* tailored to specific product categories, for marketing purposes (e.g. Synthetic Users, Evidenza). Here, generative agents are used to run customisable surveys and interviews delivering results quickly. Some studies, such as the one conducted by the synthetic research start-up Evidenza in collaboration with EY, report up to 95 % alignment in insights between synthetic respondents and their human counterparts. Also, big tech companies' CEOs, such as Sam Altman and Mark Zuckerberg, predict that generative agents will have a transformative impact across many industries in the near future.

Case study 2 – Bringing scenarios to 'life' with generative agents

The authors used GPT-4 to create simple AI agents, drawing on the EPRS publication 'EU Capabilities in Space: Scenarios for Space Security by 2050' and its **Inertia** scenario.

By 2050, the EU has fallen behind in the global space race due to division, underinvestment, and slow decision-making. Major powers dominate space militarily and commercially, while the EU lacks counter space capabilities and struggles to regulate the booming private sector. This decline has led to a brain drain, social unrest, and growing risks from unregulated satellite proliferation and space debris.

The scenario was input into the AI model, which was prompted to generate a detailed description of a *persona* living in that future, including detailed demographic characteristics. This persona was then 'interviewed' about future space legislation, aligning with the study's core topic.

ATTRIBUTE	DETAILS
Name	Matteo Lefevre
Age	44
Gender	Male
Nationality	French
Location	Dordogne region, France (off-grid agri-commune)
Occupation	Agroecological systems steward (manages AI-assisted permaculture and soil regeneration)
Education	Degree in Environmental Science; Certified in BioTech-AI Synergy for Agriculture
Family	In a long-term co-housing partnership with two adult companions; raising one child (shared guardianship through community cooperative)
Hobbies	Stargazing with modular observatory pod; drone-based landscape art; retro-tech restoration
Personality	Stoic, community-oriented, wary of centralized tech governance but values applied science

POLICY INTERVIEW PROMPT

"Matteo, what would you ask your MEP to prioritise in future space legislation?"

"I just want them to remember us", Matteo says, looking out at the darkened countryside. "We're affected by global space events, yet we're pushed to the side. Rural areas need a voice in space policy, not just cities. I would ask them to prioritise real regional representation and protection from both debris and foreign dominance. European space policy must serve all Europeans, not just the digital elite living in glass-and-metal bubbles. I'd ask the EU to focus first on strengthening regulations regarding space debris mitigation. As a farmer, I've seen first-hand the damage that space junk can cause. Every time a piece burns up overhead, I can't help but worry it might land in my field. I'd push for the EU to invest in a rural space debris monitoring network. The current system isn't equipped to properly handle these issues."

AUTHORS' COMMENT:

This simple experiment demonstrates the promising potential of using LLM-powered agents to 'materialise' foresight scenario narratives. The model can rapidly generate rich, multidimensional personas with unique backstories, enabling researchers and policymakers to 'interview' them and create human-like dialogues. This approach makes abstract future scenarios more relatable and concrete, helping policymakers understand and engage with issues that they may not personally face. However, it also raises questions about the role and impact of generative AI in scenario development and policy analysis more broadly. For instance, how does this dialogue impact our perception of policy issues and the development of possible options? How could we ensure that different perspectives and voices are represented? Careful consideration and critical evaluation of its output are needed.

Ethical challenges and limitations

Numerous ethical challenges related to generative AI have been discussed in policy and academic spheres. Among other documents, in 2019, the Commission's high-level expert group on AI published 'Ethical guidelines on trustworthy AI'. Despite their constant improvements, LLMs still face limitations that constrain their utility in foresight, policy analysis, and research. Some of the most relevant limitations for foresight are the following:

1. Biased data

LLMs show biases because they are trained on data that are neither representative nor demographically balanced. Generative agents inherit the same biases as well. This means that we need to be cautious when using and interpreting LLMs' outputs. For example, when using LLMs to create scenarios, if marginalised communities or unconventional ideas are underrepresented in the training data, the automated scenarios will fail to capture their views. This could ultimately lead to producing one-dimensional or exclusionary policy advice, rather than bringing in balanced and nuanced views of the plausible futures.

Bias in LLMs can be reduced by diversifying training data, constant monitoring, de-biasing techniques and using bias detection tools, prompt engineering and more transparency overall.

2. Dependence on historical datasets and knowledge boundaries

LLMs are constrained by the temporal and thematic limits of their training data. This can limit their ability to address emerging issues or detect weak signals, especially if they cannot access real-time data from the web. Therefore, their foresight outputs might just be projections of past trends and patterns rather than anticipations of disruptive or novel developments. This creates a risk of reinforcing the *status quo*, while transformative approaches might be needed when looking ahead.

However, preliminary evidence suggests that LLMs may be capable of extrapolating causal relationships beyond their training data and applying their learned strategies to completely new variables. Different techniques, such as data augmentation using qualitative interviews (as seen above), and new approaches can allow models to adapt to evolving trends and weak signals and reduce reliance on statistic datasets.[8]

Knowledge boundaries can also be a constraint when thinking about the future, trying to identify future uncertainties or imagining plausible future scenarios. However, in combination with human input and participatory activities, such as workshops, this obstacle can be overcome. Finally, more research to understand how LLMs represent cause-and-effect relationships could lead to their broader use in foresight projects.

3. Overreliance

Overreliance on AI represents a significant challenge because it risks diminishing original human thought, collective intelligence based on collaborative engagement and the social capital that is built through such interactions.

Foresight practitioners should explore the optimal balance in implementing the use of LLMs and generative agents in their foresight processes. Determining how to best integrate agents within larger project frameworks will be key to preventing over-reliance on AI and preserving the creative and critical thinking that could serve as an input to human-led foresight advice.

4. Questionable reliability

The performance of LLMs and generative agents needs to be tested further. Sensitivity to prompt phrasing and inconsistencies between models complicate the reliability of their outputs, because small changes in prompts can lead to completely different results. Similarly, prompt manipulation can generate fabricated or misleading content. New methods that incorporate lessons from past outcomes into model tuning show promise for improving both model realism and institutional learning (i.e. systematically learning from past actions and outcomes).[9]

Also, LLMs have a limited understanding of the context of the prompt. This can lead to misinterpretations and the identification of unreliable patterns or trends. They may also struggle to handle the uncertainties and ambiguity that foresight deals with. Their poor performance has also been noted when handling out-of-distribution data (i.e. input or scenarios that significantly differ from the data LLMs were trained on), leading to often unpredictable outcomes. Similar challenges arise when dealing with complex systems and their interactions, including non-linear relationships and unpredictable outcomes, where expert knowledge or intuition still outperforms LLMs.

Finally, while hallucinations can serve as a source of creative inspiration, they also pose risks to aspects of foresight projects that depend on desk research that could potentially introduce inaccurate or misleading information. Human oversight remains critical to ensure the reliability and integrity of the process.

Conclusions

Just a few years ago, AI was seen as having a modest, complementary role in foresight, used to support traditional methods. Since then, the rapid evolution of LLMs and the emergence of generative agents have significantly broadened the horizon of possibilities. The use of these technologies has been reshaping the research process itself, providing faster and more cost-effective pathways to knowledge generation. LLMs and generative agents could support the exploration of complex futures, provided they are used critically and with care, consideration, and continuous human oversight.

Research shows that LLMs enhance foresight capabilities and strategic decision-making, but cannot replace human expertise in critical analysis and strategic judgment. The integration of these technologies should only serve to augment and strengthen human decision-making processes, creating more effective synergies between technological capabilities and expert oversight. Ensuring that AI continues to serve as a complement to, rather than a substitute for, human intelligence remains essential for achieving robust, informed, and contextually nuanced foresight outcomes and related advice to policy.

Moreover, more transparency about the data on which a model was trained is needed, together with the reasoning process, traceability and explainability (Recital 27 of the AI Act). Identifying key data points or analogies that informed AI-made decisions is key to trustworthy AI-assisted analysis.

While the impact of generative AI and human-AI interactions will grow in foresight as well as in related domains, we propose that, for now, its role is primarily complementary to participatory foresight processes. From automating desk research tasks and testing ideas on non-human agents before engaging people, to visually representing or prototyping foresight outcomes, GenAI can already add substantial value to many foresight activities with little to no friction.

The broader sociotechnical implications of integrating LLMs and generative agents into scientific advice to policy, including foresight advice, must also be further considered. AI literacy and skills are necessary for foresight practitioners to learn how to work with AI, understand its limitations and engage with the output, potentially adopting new working modalities. Similar to other AI-based tools, LLM-powered foresight needs to become both technically sound and societally acceptable.

Therefore, the promise of AI in foresight will only be realised if we actively address technical limitations, embed ethical safeguards, and ensure that these systems complement – rather than substitute – human judgment and experience. If done right, we might look forward to a future where LLMs help us better understand uncertainties and generate insights and innovations that truly serve society. Continuous validation, expert oversight, and transparent methodologies are essential to ensure that these tools enhance, rather than undermine, informed policymaking.

The key for further GenAI application in foresight is to start experimenting, because while artificial intelligence itself will not replace policy and foresight analysts, analysts who use AI might have an advantage over those who do not.

MAIN REFERENCES

Argyle, L., Busby, E., Fulda, N., Gubler, J., Rytting, C. and Wingate D., *Out of One, Many: Using Language Models to Simulate Human Samples*, Cambridge University Press, 2023.

Park, J., Zou, C., Shaw, A., Mako Hill, B., Cai, C., Ringel Morris, M., Willer, R., Liang, P. and Bernstein, M., *Generative Agent Simulations of 1,000 People*, arXiv, 2024.

Pérez-Ortiz, M., *From Prediction to Foresight: The Role of AI in Designing Responsible Futures*, Journal of Artificial Intelligence for Sustainable Development 1 (1), pp. 1–9, 2024.

Safaei, M. and Longo, J., *The End of the Policy Analyst? Testing the Capability of Artificial Intelligence to Generate Plausible, Persuasive, and Useful Policy Analysis*, Digital Government: Research and Practice, Volume 5, Issue 1, Article No: 4, p. 1, 2024.

Si, C., Yang, D. and Hashimoto, T., *Can LLMs Generate Novel Research Ideas? A Large-Scale Human Study with 100+ NLP Researchers*, arXiv, 2024.

Spaniol M. and Rowland N., *AI-assisted scenario generation for strategic planning*, Futures & Foresight Science, Volume 5, Issue 2, 2023.

Su, J., Jiang, C., Jin, X., Qiao, Y., Xiao, T., Ma, H., Wei, R., Jing, Z., Xu, J. and Lin, J., *Large Language Models for Forecasting and Anomaly Detection: A Systematic Literature Review*, arXiv, 2024.

ENDNOTES

[1] Generative AI uses deep-learning models to generate different types of content (e.g. text, images, music). It encompasses LLMs and other AI models.

[2] Large language models are a type of AI that 'understands' and generates human language text, based on patterns learned from vast amounts of text datasets.

[3] *Generative agents* are defined as 'computational entities designed to simulate authentic human behavioural patterns'. Similar concepts used in academic and policy discourse are synthetic people/users, AI replicas, and simulation agents.

[4] The impact of LLMs on creativity remains a key concern: while they can enhance creativity during assisted tasks, they may undermine independent creative abilities with repeated use, emphasising the importance of designing AI systems that support long-term autonomous thinking.

[5] Many foresight organisations are already using LLMs in scenario creation. The guide below draws on our own experiments with LLMs, as well as insights from others in the field.

[6] The concepts of 'algorithmic fidelity' and 'silicon sampling' have been used to test for accuracy and validity of LLMs to simulate human behaviour.

[7] This includes few-shot prompting (providing the model with a few examples to demonstrate the desired response pattern) or chain-of-thought prompting (encouraging the model to reason step-by-step before reaching a conclusion).

[8] These include ensemble modelling (multiple models working together to improve accuracy and robustness), continual learning (to allow models to adapt over time), and transfer learning (knowledge from one task used to enhance performance of another) .

[9] These methods include reinforcement learning from human feedback (RLHF) and reinforcement learning from experience feedback (RLXF).

DISCLAIMER AND COPYRIGHT

eprs@ep.europa.eu (contact)

www.eprs.ep.parl.union.eu (intranet)

www.europarl.europa.eu/thinktank (internet)

http://epthinktank.eu (blog)