

ENVIRONMENTAL RESEARCH
LETTERS

PERSPECTIVE

Leveraging large language models to monitor climate technology innovation

OPEN ACCESS

RECEIVED
31 May 2023REVISED
27 July 2023ACCEPTED FOR PUBLICATION
21 August 2023PUBLISHED
7 September 2023Malte Toetzke^{1,*} , Benedict Probst^{1,2} and Stefan Feuerriegel³¹ Group of Sustainability and Technology, ETH Zurich, Zurich, Switzerland² Centre for Environment, Energy and Natural Resource Governance, University of Cambridge, Cambridge, United Kingdom³ Munich Center for Machine Learning, LMU Munich, Munich, Germany

* Author to whom any correspondence should be addressed.

E-mail: mtoetzke@ethz.ch**Keywords:** innovation, large language models, machine learning, climate technologiesOriginal content from
this work may be used
under the terms of the[Creative Commons
Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/).Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.**Abstract**

To achieve net-zero emissions, public policy needs to foster rapid innovation of climate technologies. However, there is a scarcity of comprehensive and up-to-date evidence to guide policymaking by monitoring climate innovation systems. This is notable, especially at the center of the innovation process, where nascent inventions transition into profitable and scalable market solutions. Here, we discuss the potential of large language models (LLMs) to monitor climate technology innovation. By analyzing large pools of unstructured text data sources, such as company reports and social media, LLMs can automate information retrieval processes and thereby improve existing monitoring in terms of cost-effectiveness, timeliness, and comprehensiveness. In this perspective, we show how LLMs can play a crucial role in informing innovation policy for the energy transition by highlighting promising use cases and prevailing challenges for research and policy.

1. Introduction

Accelerating the invention and diffusion of climate technologies that mitigate or remove emissions is crucial for achieving net-zero emissions (IPCC 2023). Around half of the emissions reductions until 2050 will potentially be achieved through technologies that are currently still under development or in demonstration (IEA 2020).

As such, the development of climate technologies and their transition into profitable and scalable market solutions remains a major challenge (Grubb *et al* 2014). While early demonstration and commercialization are crucial for technological ‘learning’, nascent climate technologies are often not competitive with established technologies at market entry. Reasons include higher costs, immature infrastructure, and more pronounced investment risks (Egli *et al* 2018). Therefore, public policy plays a crucial role in promoting the development and deployment of climate technologies and alleviating hurdles along the innovation process by, for example, funding pilot projects, incentivizing industry collaboration, and stimulating

market incentives through subsidies, regulations, and other mechanisms (Doblinger *et al* 2019, Goldstein *et al* 2020, Probst *et al* 2021, Meckling *et al* 2022).

Designing and implementing innovation policies requires evidence of technological developments and surrounding innovation systems, which helps to identify where certain measures could be effective. This includes, for example, evidence of investment landscapes, industry collaborations, or public acceptance. However, collecting such evidence from different sources is highly time-consuming and expensive. Therefore, relevant pieces are often not comprehensive and up-to-date.

Here, we discuss the potential of large language models (LLMs) to improve the monitoring of climate technology innovation. LLMs have gained broad attention for their human-like capabilities of retrieving relevant information from unstructured text at a large scale (Brown *et al* 2020, Liu *et al* 2023). Thereby, LLMs can help researchers and policy-makers retrieve novel and up-to-date evidence from large pools of relevant data sources such as policy documents, social media data, and company reports. In this paper, we

discuss promising use cases and challenges of applying LLMs in climate innovation research and policy.

2. Relevance of monitoring along the innovation process

For the majority of climate technologies, the innovation process can be divided into three main stages: the invention stage, the innovation stage, and the diffusion stage (figure 1; (Grubb *et al* 2014)). This process is iterative and embedded in dynamically evolving technology innovation systems comprising networks of relevant actors (e.g. manufacturers, vendors) and institutions (e.g. regulations, social norms) (Bergek *et al* 2015, Markard 2020). A critical point is the innovation stage where technological solutions need to find profitable markets (Gallagher *et al* 2012, Grubler and Wilson 2014). At this stage, climate innovations often compete with more mature fossil technologies and, thereby, need to overcome financial and techno-economic deficits, such as the range and charging time of electric vehicles (EVs) (Goldstein *et al* 2020).

Policies can help to overcome such hurdles by addressing infrastructural requirements (e.g. improving charging infrastructure for EVs), providing finance (e.g. through grants or public procurement), or stimulating demand (e.g. through subsidies or regulations, such as carbon taxes) (Gallagher *et al* 2012). To foster the effective implementation of such policies, it is particularly important to monitor climate technology innovations with comprehensive and up-to-date information. In this way, the current functioning of innovation systems can be analyzed, allowing the identification of current hurdles and targeted allocation of resources (Hekkert *et al* 2007). Yet, especially at the center of the innovation process, existing processes for retrieving information on innovation systems are often extremely expensive and time-consuming, which affects the timeliness and comprehensiveness of provided evidence.

3. Existing information retrieval processes for monitoring climate technology innovation

Existing databases that enable monitoring of climate technology innovation are typically based on structured information provided through (1) secondary data sources, (2) surveys, and (3) manual retrieval from unstructured sources. We discuss the three cases in the following.

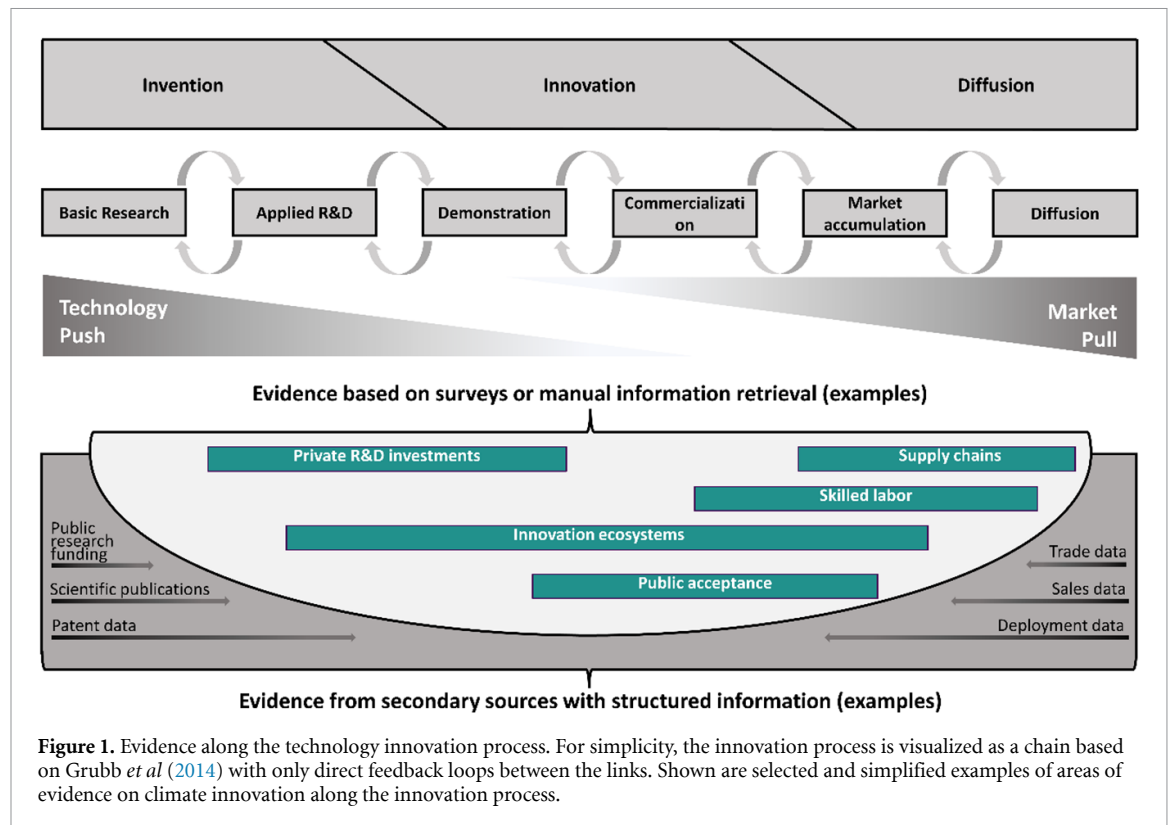
In the first case, databases are assembled from secondary data sources, that is, other databases, where the latter are originally curated for their own distinct purposes like academic publication platforms, patent databases, or trade data. As such, they benefit from individuals and organizations being either legally

obliged (e.g. companies that are required to report trade activities to national customs authorities) or strongly incentivized to contribute specific information (e.g. inventors that file a patent to gain protection for an invention). While, in general, secondary data sources are relatively comprehensive, the underlying self-reporting process may also cause specific selection biases. For example, the database PATSTAT (European Patent Office 2023) collects patent data from national patent offices of leading industrialized and developing countries. While PATSTAT contains more than 300 thousand patent documents on clean energy technologies, some inventors might refrain from using patents to protect climate inventions because of time, costs, or secrecy, which likely leads to a lower representation of small companies, lower income countries, or software-based inventions.

In the second case, surveys are used to collect specific information directly from primary sources. However, only a few organizations can convince companies and other organizations to directly contribute information without legal obligations or substantial value delivered in return. For example, the International Energy Agency (IEA) requests information on technology costs directly from companies, and the Organizations of Economic Cooperation and Development (OECD) collects policy-related information from national ministries. While acquiring such information directly from the source organizations usually leads to accurate information, the underlying bureaucratic workload only allows for a limited number of sources. In some cases, such as the estimation of technology costs across countries, generalizability can be affected if the number of contributing organizations only represents a small share of the total population.

In the third case, databases are manually collected, so that information on innovation systems from various unstructured sources is captured. For example, databases, such as Bloomberg New Energy Finance (BNEF) (BloombergNEF 2023) or i3 (Cleantech Group 2023), provide information on clean energy companies, but are manually compiled by screening news announcements, company reports, and company websites. This is an extremely expensive and time-consuming process. As a result, such databases usually charge high subscription fees. Furthermore, the data-collection process is often not transparent, revealing little information about potential biases and other limitations. For example, although BNEF is considered the most comprehensive database for renewable energy project finance, there are still major concerns about the accurate representation of projects across different countries (Lilliestam *et al* 2020).

Overall, secondary sources with structured information on climate innovation are commonly located at the early invention and late diffusion stages of the innovation process. Monitoring climate



innovation around the center of the innovation process thus often requires information retrieval via surveys or manual data collection. Both involve extensive bureaucratic and manual workload. This, for example, applies to information retrieval of private R&D investments, innovation ecosystems, or skilled labor (see figure 1). As we argue below, LLMs can help to automate information retrieval from unstructured sources, speeding up the processes while, at the same time, reducing costs. Thereby, LLMs can enable information retrieval from much larger pools of data sources, such as newspapers, company websites, or social media, and produce more timely evidence on innovation systems. Furthermore, LLMs can be used to enrich evidence from patents or academic publications with additional context from corresponding raw texts.

4. Tracking climate technology innovation with LLMs

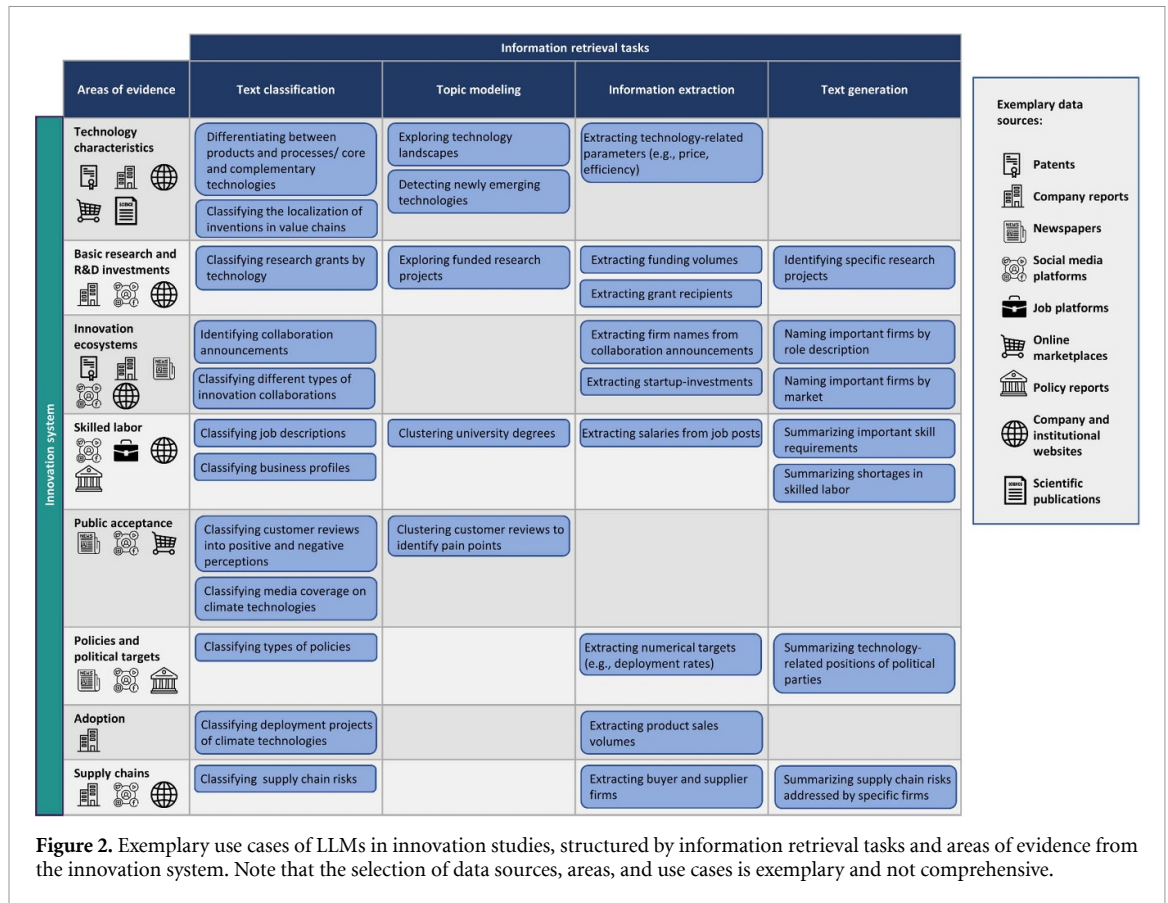
4.1. Background on information retrieval with LLMs

LLMs are deep learning models for processing and generating human language. Compared to previous approaches in natural language processing, LLMs are much more powerful. For example, the parameter size of pretrained LLMs typically ranges from 20 million to more than 100 billion parameters (Devlin *et al* 2018, Lewis *et al* 2019, Floridi and Chiriatti 2020, Touvron *et al* 2023).

Three main abilities distinguish LLMs from previous approaches in natural language processing. First, instead of treating text as a bag of words without ordering, LLMs take the sequence of words into account and thereby understand hierarchies and relationships between entities. As a result, LLMs can capture complex semantics and even generate meaningful content as output. Second, LLMs allow users to retrieve information via so-called ‘prompts’ which are commands in natural language. As such, they can perform information retrieval tasks through in-context learning, where, examples of the task can be directly provided in the prompts (Liu *et al* 2023). Third, LLMs are pre-trained on vast corpora of text data. This enables few-shot learning, where high performance can be achieved with little additional training data, potentially avoiding large sets of manually-annotated data (Brown *et al* 2020).

LLMs can perform a wide range of information retrieval tasks on texts, and each task has particular advantages and limitations. We discuss these information retrieval tasks in the context of climate innovation monitoring in the following:

Through *text classification*, the model learns to categorize documents along specified dimensions (e.g. the sustainable development goals; climate technology classes (Toetzke *et al* 2022c)). Text classification offers customized, consistent, and replicable outcomes and reports clear performance metrics (e.g. accuracy, F1-score). By nature, the task is restricted to a set of predefined labels (e.g. solar energy, wind energy, etc., for classifying climate technologies).



Through *topic modeling*, the model identifies prominent topics from a large corpus of documents (e.g. policy reports). As such, topic modeling analyzes text documents for exploratory purposes (e.g. ‘What are most prominent topics in company sustainability reports?’). Topic modeling is an unsupervised task, which means that it does not require predefined labels. Instead, it explores potential clusterings of the text documents at hand. This is beneficial for monitoring tasks where no categorization scheme yet exists. However, in turn, topic modeling does not allow users to target their analyses on specific questions of interest.

Through *information extraction*, the model extracts numbers (e.g. financials, technical parameters) and named entities (e.g. names of regions or companies) from texts. LLMs allow such information extraction with few- or even zero-shot learning where examples can be included in the prompts (e.g. ‘Extract names of companies, such as Siemens AG or BlackRock, Inc, from the following text’). By nature, information extraction only returns information that is explicitly mentioned in the text (e.g. revenue numbers, legal paragraphs).

Through *text generation*, the model generates customized summaries of large text corpora based on specific user questions. As such, text generation can answer semantically complex questions about large texts without the need for downstream training (e.g.

‘Name the most promising technologies for energy storage in the IPCC reports’). On the one hand, text generation is a very versatile task with highly flexible outputs. On the other hand, it is comparatively difficult to evaluate, difficult to control algorithmically, and prone to hallucination, where plausible but factually wrong answers can be returned.

4.2. Promising use cases for LLMs in climate innovation research

Through the above-described tasks, LLMs can automate information retrieval and, thereby, help to address various evidence gaps along the innovation process. In the context of climate innovation monitoring, the selection of information retrieval tasks is essential as different tasks enable different use cases. Figure 2 provides a schematic overview of exemplary use cases, attributed to different areas of evidence, information retrieval tasks, and potential data sources.

4.2.1. Technology characteristics

Applications of LLMs can help to provide more context to existing monitoring of technology characteristics by analyzing scientific publications, patents, or company reports at scale. For example, text classification allows for distinctions in patents between product and process inventions or core versus complementary inventions. As such, one could analyze

whether inventions of specific climate technologies, such as e-fuels, focus on technology production or downstream technology integration. Topic models, conversely, are more suitable for exploratory tasks, such as clustering product descriptions to identify technology variations (e.g. in carbon removal technologies). Through information extraction, LLMs can extract technology parameters, such as input materials for EV batteries from patents or battery capacities from product descriptions.

4.2.2. Basic research and applied R&D investments

Public research funding is usually reported in detail by public financiers (e.g. research foundations) and recipients (e.g. universities), including structured information (e.g. research disciplines) and additional textual descriptions on the project level (Meckling *et al* 2022). Here, LLMs can be used to further contextualize existing monitoring via topic models, text classification (e.g. regarding the inclusion of economical, technological, or social aspects), or chats with text generation models to identify projects of interest (e.g. ‘Which research projects investigate the potential of hydrogen production in sub-Saharan Africa?’). In contrast, companies report only basic information on their R&D spending, such as overall R&D budgets. While further details are often not disclosed (e.g. spending by technology), LLMs can still be used to systematically peruse millions of company reports and extract relevant information where existing.

4.2.3. Innovation ecosystems

Databases for monitoring innovation ecosystems mostly focus on specific types of innovation networks, such as venture capital investments in startups or project finance for renewable energy. Such databases are usually compiled manually from different sources, including news announcements and company websites. However, often the underlying data lacks evidence on emerging technologies or developing countries and is biased toward specific sources. Here, LLMs can automate the information retrieval process and systematically review much larger pools of secondary data. For example, a recent study uses LLMs to analyze partnership announcements of climate technology companies from several million social media posts, classifying them by the type of interaction (e.g. equity funding, R&D collaboration) and targeted technology, and extracting the names of collaborating actors and with corresponding roles (e.g. investor, developer) (Toetzke *et al* 2023).

4.2.4. Skilled labor

While employment data helps monitor different occupations across industries, scant evidence exists on labor demand and available skills. Due to the energy transition, labor markets are changing dramatically, leading to both extensive job loss and the creation of new jobs (Zaussinger *et al* 2023). Here,

LLMs can help to monitor changes in the job market structure, on the demand side, by classifying job descriptions (e.g. required qualifications, job linkage to climate or fossil technologies), and on the supply side, by analyzing the availability of skills (e.g. through clustering university degrees). Furthermore, LLMs can extract information on job-specific salaries. This can help to compare whether job losses in fossil industries can be compensated through climate technology jobs in the same region and requiring equivalent skills, thereby supporting a just energy transition.

4.2.5. Public acceptance

For consumer products (e.g. EVs or rooftop photovoltaic systems), public acceptance can be proxied through product sales. Here, using LLMs to analyze product reviews on, for example, marketplaces can help to understand consumer behavior and identify decisive factors in purchase decisions (e.g. charging time of EVs; costs of electricity for solar power). For infrastructure-heavy climate technologies, such as wind parks or nuclear plants, public acceptance is challenging to monitor with traditional methods. Opinion polls are expensive and only capture small samples (Cox *et al* 2020). Here, LLMs can be used to monitor public acceptance across large populations by classifying social media posts or newspaper coverage along relevant aspects (e.g. safety, sustainability, costs), while distinguishing between positive, negative, and neutral stances (Toetzke *et al* 2022b).

4.2.6. Policies and political targets

Monitoring policies and political targets traditionally requires substantial resources. For example, the Net Zero Tracker (Lang *et al* 2023) relies on large numbers of research assistants and volunteers to collect information on emission targets across countries, cities, and industries. Other organizations, such as the OECD, use self-reporting by national ministries to distribute the excessive workload of retrieving information from relevant sources (Toetzke *et al* 2022a). Here, LLMs can automate the information retrieval process through classification (e.g. types of innovation policies), information extraction (e.g. emissions reduction targets), or even interactions with chatbots pre-trained on relevant policy reports and databases (e.g. ‘Which countries currently have a carbon price above USD 100?’). While the compilation of relevant sources remains a necessary prerequisite, using LLMs allows organizations to speed up the information retrieval process and scale it to much larger pools of data sources and varieties of parameters, at low costs.

4.2.7. Adoption

The adoption of climate technologies is often not monitored comprehensively, especially for developing countries and emerging technologies. Here, LLMs can help identify demonstration projects through

project announcements via different channels, such as company reports, newspapers, or social media. However, smaller installations, such as residential solar PV systems, are likely to not be captured by these secondary data sources. In this case, sales figures for climate technologies could be extracted from company reports and websites, if disclosed properly.

4.2.8. Supply chains

LLMs can help to map climate technology supply chains by extracting buyer-supplier ties from social media announcements, company reports, or websites. A major bottleneck here is that, in many cases, buyer-supplier information is not fully disclosed, which could lead to incomprehensive and potentially biased evidence. LLMs can also be used to monitor risks of climate technology supply chains (e.g. scarce input materials) and analyze how companies address such risks in their reporting.

5. Discussion

LLMs show great potential for informing climate innovation research and policy. Across the entire innovation process, they can help to retrieve new evidence through classification, topic modeling, information extraction, and text generation from large pools of relevant text data. Thereby, they can provide more comprehensive and timely insights on climate innovation systems through automated, cost-efficient, continuous, and scalable monitoring. As such, monitoring can support targeted innovation policies, identifying emerging innovation clusters, suitable deployment locations, and promising markets. Especially at the center of the innovation process, LLMs can improve information retrieval processes that are currently cumbersome and limited in scope, by mapping innovation ecosystems, analyzing public acceptance, and comparing innovation policies across countries at large scale. Thereby, innovation policy can address prevalent bottlenecks (e.g. investment shortages or public safety concerns) and accelerate the translation of climate inventions into scalable solutions.

However, using LLMs for monitoring climate innovation has limitations in terms of model applicability and model outcomes. In terms of model applicability, data availability and computational resources represent major bottlenecks. While, in many contexts, relevant data is made publicly available (e.g. patents by inventors, research funding by public institutions, or public opinions via social media), in some contexts (e.g. corporate R&D investments or supply chain information), available data can be limited due to non-disclosure. Here, LLMs still facilitate the search for relevant information by screening documents from different sources at a scale that would be impractical to conduct manually.

However, to improve the overall comprehensiveness of monitoring, it is essential that policy-makers foster company reporting and open data. Furthermore, LLMs demand significant computational resources at both training and inference time (e.g. high-performance GPUs and ample storage capacity), which pose financial and technical barriers.

In terms of model outcomes, applying LLMs for information retrieval raises important questions concerning validity and transparency. This is especially relevant for prompt-based information retrieval or in-context learning where task-specific training on large sets of self-annotated training data is avoided. This enforces the black-box character of LLMs as inferences become less transparent. Because the underlying reasoning of LLMs is based on excessive primary training data, which also includes factually doubtful sources, outcomes can potentially be biased. Furthermore, LLMs are prone to hallucination where incorrect answers are generated. To address these shortcomings, we expect the following procedures to be helpful: First, user studies or annotated test sets should validate the model performance, using data that have not been used for prompt engineering, context-provision, or fine-tuning. Second, codebooks should be manually compiled by experts to provide important definitions and categorizations as a basis for performance validations and, potentially, as additional context for the model task. Third, where available, fine-tuned models could be used that have been specifically trained on factually proven sources related to the required task, such as ChatClimate, which has been trained on the latest IPCC report (Vaghefi *et al* 2023).

The application of LLMs has several practical benefits compared to traditional approaches from natural language processing. As such, LLMs allow users to avoid several labor- and cost-intensive steps in text preprocessing and model training. For example, LLMs enable few-shot learning, where only a small number of data samples are needed for training. Currently, we observe rapid and continuous improvements in model efficiency, which allows the deployment of LLMs on traditional hardware (Rasley *et al* 2020, Hu *et al* 2021). In the future, we expect a steep growth in the availability of user-friendly interfaces for customized training and deployment of LLMs as well as improvements in the verifiability of outputs through generating citations to credible references.

In summary, we expect substantial advances in innovation research and policy driven by the application of LLMs. As the field continues to evolve, further work should be dedicated to addressing major challenges, such as the availability of open data, algorithmic interpretability, and cross-disciplinary collaboration to maximize the potential of LLMs in monitoring climate innovation.

Data availability statement

No new data were created or analysed in this study.

Acknowledgments

MT acknowledges funding from the Swiss National Science Foundation (SNSF) as part of the Eccellenza grant 186932 as well as from the Swiss Federal Office of Energy's "SWEET" programme as part of the PATHFNDR consortium.

ORCID iD

Malte Toetzke  <https://orcid.org/0000-0002-1153-2702>

References

- Bergek A, Hekkert M, Jacobsson S, Markard J, Sandén B and Truffer B 2015 Technological innovation systems in contexts: conceptualizing contextual structures and interaction dynamics *Environ. Innov. Soc. Transit.* **16** 51–64
- BloombergNEF 2023 BloombergNEF (available at: www.bnef.com/) (Accessed 31 January 2023)
- Brown T et al 2020 Language models are few-shot learners *Advances in Neural Information Processing Systems* vol **33** pp 1877–901
- Cleantech Group 2023 Discover i3 Market Intelligence (available at: <https://i3connect.com/>)
- Cox E, Spence E and Pidgeon N 2020 Public perceptions of carbon dioxide removal in the United States and the United Kingdom *Nat. Clim. Change* **10** 744–9
- Devlin J, Chang M-W, Lee K and Toutanova K 2018 Bert: pre-training of deep bidirectional transformers for language understanding (arXiv:1810.04805)
- Doblinger C, Surana K and Anadon L D 2019 Governments as partners: the role of alliances in US cleantech startup innovation *Res. Policy* **48** 1458–75
- Egli F, Steffen B and Schmidt T S 2018 A dynamic analysis of financing conditions for renewable energy technologies *Nat. Energy* **3** 1084–92
- European Patent Office 2023 PATSTAT (available at: www.epo.org/searching-for-patents/business/patstat.html)
- Floridi L and Chiriatti M 2020 GPT-3: its nature, scope, limits, and consequences *Minds Mach.* **30** 681–94
- Gallagher K S, Grübler A, Kuhl L, Nemet G and Wilson C 2012 The energy technology innovation system *Annu. Rev. Environ. Resour.* **37** 137–62
- Goldstein A, Doblinger C, Baker E and Anadon L D 2020 Patenting and business outcomes for cleantech startups funded by the advanced research projects agency-energy *Nat. Energy* **5** 803–10
- Grubb M, Hourcade J-C and Neuhoff K 2014 *Planetary Economics* (Routledge)
- Grubler A and Wilson C 2014 *Energy Technology Innovation* (Cambridge University Press)
- Hekkert M P, Suurs R A A, Negro S O, Kuhlmann S and Smits R E H M 2007 Functions of innovation systems: a new approach for analysing technological change *Technol. Forecast. Soc. Change* **74** 413–32
- Hu E J, Shen Y, Wallis P, Allen-Zhu Z, Yuanzhi L, Wang S, Wang L and Chen W 2021 Lora: low-rank adaptation of large language models (arXiv:2106.09685)
- IEA 2020 Energy technology perspectives 2020 (available at: www.iea.org/reports/energy-technology-perspectives-2020)
- IPCC 2023 Climate change 2023: synthesis report (available at: www.ipcc.ch/report/ar6/syr/)
- Lang J, Hyslop C, Yeo Z Y, Black R, Chalkley P and Hale T 2023 Net zero tracker (available at: zerotracker.net)
- Lewis M, Liu Y, Goyal N, Ghazvininejad M, Mohamed A, Levy O, Stoyanov V and Zettlemoyer L 2019 Bart: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension (arXiv:1910.13461)
- Lilliestam J, Melliger M, Ollier L, Schmidt T S and Steffen B 2020 Understanding and accounting for the effect of exchange rate fluctuations on global learning rates *Nat. Energy* **5** 71–78
- Liu P, Yuan W, Jinlan F, Jiang Z, Hayashi H and Neubig G 2023 Pre-train, prompt, and predict: a systematic survey of prompting methods in natural language processing *ACM Comput. Surv.* **55** 1–35
- Markard J 2020 The life cycle of technological innovation systems *Technol. Forecast. Soc. Change* **153** 119407
- Meckling J, Galeazzi C, Shears E, Xu T and Anadon L D 2022 Energy innovation funding and institutions in major economies *Nat. Energy* **7** 876–85
- Probst B, Touboul S, Glachant M and Dechezleprêtre A 2021 Global trends in the invention and diffusion of climate change mitigation technologies *Nat. Energy* **6** 1077–86
- Rasley J, Rajbhandari S, Ruwase O and Yuxiong H 2020 DeepSpeed: system optimizations enable training deep learning models with over 100 billion parameters *Proc. 26th ACM SIGKDD Int. Conf. on Knowledge Discovery & Data Mining* pp 3505–6
- Toetzke M, Banholzer N and Feuerriegel S 2022a Monitoring global development aid with machine learning *Nat. Sustain.* **5** 533–41
- Toetzke M, Probst B, Tatar Y, Feuerriegel S and Hoffmann V 2022b Analyzing the global energy discourse with machine learning *NeurIPS 2022 Workshop on Tackling Climate Change with Machine Learning* (available at: www.climatechange.ai/papers/neurips2022/96)
- Toetzke M, Re F, Probst B, Feuerriegel S, Diaz Anadon L and Hoffmann V 2023 Mapping global innovation networks around clean energy technologies *ICLR 2023 Workshop on Tackling Climate Change with Machine Learning* (available at: www.climatechange.ai/papers/iclr2023/32)
- Toetzke M, Stünzi A and Egli F 2022c Consistent and replicable estimation of bilateral climate finance *Nat. Clim. Change* **12** 897–900
- Touvron H et al 2023 Llama: open and efficient foundation language models (arXiv:2302.13971)
- Vaghefi S A et al 2023 Chatclimate: grounding conversational AI in climate science
- Zaussinger F, Egli F and Schmidt T S 2023 The impact of the low-carbon transition on employment and skills in Europe