

diverged twice from the typical analysis pipelines that are used to identify intertumour heterogeneity. First, the authors opted not to perform a procedure called batch correction, which is intended to minimize variability that can arise from the sequencing technique itself, because this correction method can reduce the capture of variability that arises naturally in a biological sample. Second, instead of defining gene-expression patterns in a combined data set involving all tumours, the authors initially defined variable gene-expression programs in each tumour sample and then compared the programs found across the samples.

Among the cancer cells, the authors found more than 5,500 variable gene-expression programs that were robustly observed across tumours. The authors grouped the programs into 41 clusters, which they defined as meta-programs (MPs). For each cluster, Gavish *et al.* identified 50 'consensus' genes, which provided a signature of the MP. Of the 41 MPs, 16 were similar to ones described previously⁵. The authors also applied this analysis to non-cancer cells in the tumour microenvironment and to cancer cell lines grown *in vitro*.

The authors found that MPs in cancer cells recur across a range of cancer types (Fig. 1). To identify the fundamental biological processes that these MPs represent, Gavish and colleagues categorized them into 11 hallmark categories, such as 'cell cycle', 'mesenchymal' and 'senescence'. These groupings indicate that tumours are composed of subpopulations of cancer cells that might drive different aspects of tumour progression, such as its growth, its spread to other sites (metastasis) and resistance to drug treatment. Characterizing these subpopulations and their vulnerabilities for each tumour might therefore be biologically informative and clinically useful.

Gavish *et al.* discovered that the identity of MPs in non-cancer cells helps to explain some of the MPs present in cancer cells. They observed that several MPs in cancer cells were also present in their non-cancer counterparts, suggesting that much of the heterogeneity seen in cancer cells is a combination of fundamental cellular heterogeneity and heterogeneity that arises during tumour formation. Interestingly, this variability was independent of the degree of genetic diversity of the cancer cells. Moreover, the authors identified co-occurrences of various MPs in different cell types, which might indicate that evolutionarily conserved factors drive MPs in both cancer cells and other cells in the tumour microenvironment.

The data also highlight the usefulness of cell lines for studying intratumour heterogeneity. Gavish and colleagues discovered that although cell lines do not recapitulate the full extent of heterogeneity in tumour samples from patients, data from cell lines nevertheless capture a subset of the MPs and certain hallmark-category processes. This result suggests that cell lines can be valuable tools for exploring certain aspects

of tumour heterogeneity and for advancing precision medicine. Moreover, the research provides specific information about which variable features of tumours are captured by the cell lines.

The authors grouped together many gene-expression programs to find recurring MPs. However, a limitation of this approach is that it might miss some biologically meaningful complexity, such as gene-expression programs that involve only a few genes or that exist in rare cells. Furthermore, scRNA-seq data cannot capture sources of heterogeneity beyond the level of messenger RNA abundance, such as variability in the amount of protein produced. Nonetheless, the authors have greatly advanced our understanding of how gene-expression programs are patterned in tumours and across cancers.

Gavish and colleagues open up several exciting prospects for future areas of cancer research. As the size of this scRNA-seq compendium grows, it might be possible to define MPs in more cancer contexts or in rare types of cancer cell. The emergence of data sets that capture RNA-sequencing data together with information about the spatial context of cells in tissues will offer a way to further validate the cell–cell interactions that occur in the tumour microenvironment and that determine the extent to which the heterogeneity

of hallmark processes is organized in specific locations. Future scRNA-seq experiments, paired with more data, might shed light on the determinants of transcriptional heterogeneity (for example, through the use of single-cell 'multi-omics' analysis), and examining clinical data could offer insights into the functional outcomes of this heterogeneity.

The evolutionarily conserved pathways described by Gavish and colleagues present opportunities to identify vulnerabilities across different cancers that could be exploited for therapeutic purposes. Future studies in this area could increase our understanding of how precision medicine should take intratumour heterogeneity into account.

Raymond W. S. Ng and **Sydney M. Shaffer**

are in the Department of Pathology and Laboratory Medicine and the Department of Bioengineering, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA.
e-mail: sydshaffer@gmail.com

1. Gavish, A. *et al.* *Nature* **618**, 598–606 (2023).
2. Couturier, C. P. *et al.* *Nature Commun.* **11**, 3406 (2020).
3. Raghavan, S. *et al.* *Cell* **184**, 6119–6137 (2021).
4. Jerby-Arnon, L. *et al.* *Cell* **175**, 984–997 (2018).
5. Barkley, D. *et al.* *Nature Genet.* **54**, 1192–1201 (2022).

The authors declare no competing interests.
This article was published online on 31 May 2023.

Artificial intelligence

Online tools help language models to solve problems

Aleksandra Piktus

The large language models popularized by chatbots are being taught to alternate reasoning with calls to external tools, such as Wikipedia, to boost their accuracy. The strategy could improve fact-finding outcomes, as well as online shopping.

As the success of ChatGPT shows, large language models (LLMs) are becoming ever easier to use. The power of these systems lies in their ability to respond to textual prompts by generating natural-sounding language. But what makes them even more effective is their ability to learn by example: by taking in just a few demonstrations, LLMs can markedly improve their performance on complex tasks¹. Precisely how to formulate these examples to elicit accurate answers is an open problem, but researchers have some ideas. Reporting at the Eleventh International Conference on Learning Representations (go.nature.com/42qwbwg), Yao *et al.*² propose ReAct – a prompting strategy that improves on existing methods by breaking down a multifaceted reasoning

problem into bite-sized tasks and outsourcing them to external tools.

The name ReAct refers to Yao and colleagues' integration of an 'acting' step into the process of solving a problem through reasoning. This step gives LLMs a way to interact with other tools, such as Wikipedia, by executing programming requests known as API (application programming interface) calls. The authors incorporated this step into an existing strategy called chain-of-thought prompting³, in which tasks requiring complex reasoning are decomposed into more-granular steps.

ReAct prompts LLMs to work through a trajectory of reasoning tasks and 'actions' to solve a given problem (Fig. 1). A reasoning step generates text, which Yao and colleagues call

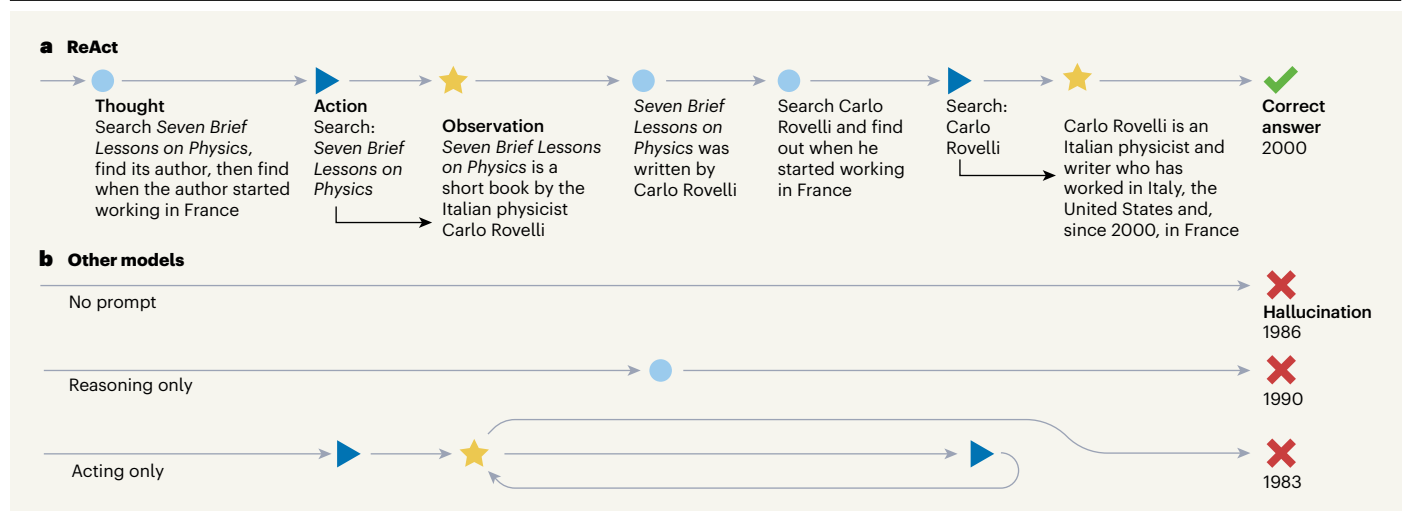


Figure 1 | A strategy to improve the output of large language models (LLMs). Yao *et al.*² designed a method called ReAct to help LLMs to solve problems, such as finding out when the author of the book *Seven Brief Lessons on Physics* first started working in France. **a**, ReAct prompts LLMs to work through a trajectory of reasoning tasks (called ‘thoughts’) interwoven with queries to external tools (called ‘actions’), yielding ‘observations’ that inform further thoughts. **b**, Other models produce incorrect answers (known as

hallucinations) more often than does ReAct, particularly when no prompts are used. Methods that use reasoning only generate thoughts using the language accessible to the LLM without any external input. ‘Acting’ strategies that use only actions can fail to make sense of the observations they acquire. Yao and colleagues’ combination of reasoning and acting reduces the number of hallucinations produced by LLMs. (Adapted from go.nature.com/3nbkxcv.)

‘thoughts’, using all the language to which the LLM currently has access. This process doesn’t involve any external input, but the thoughts are interwoven with actions that result in ‘observations’ that augment the language accessible to the LLM.

This combination of reasoning and acting proves to be an effective one. It lends itself to a diverse array of tasks – from knowledge-intensive tasks that require questions to be answered and facts checked, to decision-making problems, such as online shopping, that involve interacting with an external environment through various tools.

Knowledge-intensive tasks are complicated by the phenomenon of ‘hallucinations’, a term used to describe instances of plausible but factually incorrect text generated by an LLM. Yao *et al.* tackled this problem by giving LLMs access to Wikipedia through an API call. The ability to retrieve relevant passages is not essential for solving the task at hand, because LLMs can easily attempt to answer questions or fact-check statements without it. However, as the authors found, this step helped to reduce the number of hallucinations of standard chain-of-thought models in the samples that they studied.

Yao *et al.* borrowed from the terminology of reinforcement learning when working with decision-making problems. Specifically, they cast the LLM as an ‘agent’ that interacts with an environment through a set of well-defined actions that are formulated in everyday language. In one of Yao and colleagues’ examples, ReAct played a conversation-based game; in another, it interacted with a fictitious online shop to pick a product on the basis of criteria provided by the user. In both cases, the ability

to interact with the environment was an essential aspect of the task itself. Gaining access to the APIs of different environments is therefore key to expanding the set of problems to which LLMs can be successfully applied.

It’s worth noting that smaller, dedicated models tailored to specific problems typically perform better than LLMs on tasks for which they were trained. Question answering is a case in point: fine-tuning a model with ReAct-inspired examples works much better than simply prompting the model (Fig. 1), and both techniques lag behind state-of-the-art methods that are dedicated to the task. The trade-off is between efficiency and generality: specialized task-specific solutions are effective but have limited applicability, whereas models that can solve many problems often do so in a suboptimal and cost-inefficient way.

Although Yao and colleagues’ approach shows promise, a word of caution is warranted. Motivated by the need to make LLM predictions easier to interpret and more trustworthy, the authors embraced the idea of designing their models to solve problems as they assume a person would. This tendency leads to an anthropomorphism that manifests itself, for example, by referring to reasoning steps as ‘thoughts’. Such metaphors can be misleading because they suggest that LLMs actually carry out reasoning that they then explain to us if prompted in the right way. This obfuscates the fact that these elaborate prompting techniques are being developed precisely because LLMs fail to predict the right answer in the first place. By contrast, other authors⁴ recognize that the optimal way for a machine to incorporate API calls could differ from the approach of a human.

Prompting methods such as ReAct seem to be focused on improving the accuracy of LLMs and ensuring that their outputs are grounded in facts. Achieving this goal by increasing the complexity of the prompts is bound to boost the credibility of future generations of LLMs. But this might turn out to be a double-edged sword: improved prompting is likely to fix some hallucinations, but in doing so, it might make the new ones even more insidious, precisely because they look so credible. The cost of spotting mistakes will inevitably increase as prompting methods become more sophisticated.

At the same time, the idea that LLMs could develop into a general interface for interacting with various online tools is highly attractive – particularly in relation to tools that use natural language. ReAct is a tentative step in this direction: it solves some problems while creating others. Yet it is not hard to imagine a near future in which the ability to use tools such as ReAct will become crucial in the next big breakthrough in artificial intelligence.

Aleksandra Piktus is at Hugging Face, New York City, New York 11201, USA. e-mail: ola.piktus@gmail.com

1. Brown, T. B. *et al.* Preprint at <https://arxiv.org/abs/2005.14165> (2020).
2. Yao, S. *et al.* Preprint at <https://arxiv.org/abs/2210.03629> (2022).
3. Wei, J. *et al.* Preprint at <https://arxiv.org/abs/2201.11903> (2022).
4. Schick, T. *et al.* Preprint at <https://arxiv.org/abs/2302.04761> (2023).

The author declares no competing interests. This article was published online on 1 May 2023.