Can There Be True Artificial Experts in a Given Area?

By Faraz Fookeer

Introduction

Artificial intelligence has long been presented as a means of replacing or surpassing human expertise. From early expert systems in the 1970s and 1980s to contemporary large language models (LLMs) acting as therapists, teachers, and advisers, AI has repeatedly been framed as capable of performing the work traditionally reserved for trained professionals (experts). However, the repeated limitations of such systems suggest that there may be conceptual constraints on what it means to be an “expert,” rather than merely technical constraints on performance.

This essay examines whether artificial systems can be considered true experts. It does so by applying Quast’s (2018) three-dimensional account of expertise, which distinguishes dispositional competence (ExpertiseD), manifestational achievement (ExpertiseM), and functional role fulfilment (ExpertiseF). The central claim defended is that AI systems can satisfy the first two dimensions in restricted ways, but struggle to satisfy the third in a manner that is philosophically robust. Although modern AI can simulate expert interaction convincingly, this simulation often depends on user projection and institutional framing rather than genuine expertise.

The argument is developed by first outlining Quast’s framework and its implications. It then considers the expert systems era as a historical test case. Next, it analyses contemporary chatbots through the concept of parasocial trust (Maeda and Quan-Haase, 2024). Finally, it addresses the counterargument that modern AI’s explanatory capacities, performance, and increasing adoption suggest that the distinction between simulation and expertise is no longer meaningful.

The Concept of Expertise

The question of whether AI can be a “true” expert depends on what is meant by expertise. Quast (2018) argues that expertise is not a single property but involves three analytically distinct dimensions.

First, ExpertiseD refers to dispositional expertise: the possession of competences, skills, and relevant knowledge. This corresponds to the idea that experts have internal capacities enabling them to perform in their domain. Second, ExpertiseM refers to manifestational expertise: the successful production of expert-level outputs. A system that reliably produces accurate

diagnoses or correct translations may demonstrate ExpertiseM even if it is unclear whether it possesses a deeper understanding. Third, ExpertiseF refers to functional expertise: the fulfilment of a socially embedded service role. According to Quast, expertise is conceptually tied to the social deployment of competence to achieve client ends accurately, within a relationship of reliance and trust (Quast, 2018).

Quast's framework is significant because it rejects reductionism. Expertise cannot be reduced solely to competence, nor solely to reputation. A person may be socially recognised as an expert while lacking genuine competence, such as Quast's example of a translation expert who secretly outsources all work. Conversely, someone may possess extraordinary competence but never occupy the functional expert role. Quast therefore proposes a "balanced account" in which the three dimensions must be integrated.

This framework provides a useful lens for evaluating AI systems. Artificial systems may plausibly satisfy ExpertiseD and ExpertiseM in narrow domains. The deeper issue is whether they can satisfy ExpertiseF, since this requires not only performance but a form of socially embedded responsibility and responsiveness.

Expert Systems and the Limits of Codified Expertise

The expert systems developed in the 1970s and 1980s provide an important historical test case for artificial expertise. Systems such as DENDRAL and MYCIN were designed explicitly to replicate expert reasoning in specialised domains. DENDRAL, for example, could infer molecular structures from mass spectrometry data, with its performance measured against doctoral researchers working on similar problems (Brock, 2018). MYCIN similarly achieved high performance in medical diagnosis tasks and has often been discussed as outperforming junior clinicians in narrow diagnostic contexts.

These systems provide clear evidence that artificial systems can achieve dispositional and manifestational expertise. They contained encoded domain knowledge and produced correct outputs. If we look through the lens of Quast, they exhibited the ExpertiseD and ExpertiseM dimensions (Quast, 2018).

However, expert systems encountered significant limitations, especially when applied beyond restricted contexts. Brock (2018) documents how the knowledge engineering approach relied on extracting expert knowledge through interviews and formalising it into production rules. This process was time-consuming and difficult to scale. Moreover, production-rule systems were often too rigid to capture the variety of reasoning strategies used by human experts, including the distinction between causal explanation, heuristic judgement, statistical correlation, and definitional knowledge.

More importantly, expert systems struggled with the functional dimension of expertise. In high-stakes contexts, users require more than correct outputs: they require justification, interpretability, and confidence that the system can recognise its own limitations. Although

expert systems could produce traces of rule execution, these were not equivalent to the explanatory reasoning expected from human experts. The inability to adapt to unusual cases or to respond flexibly to contextual demands undermined their capacity to function as trusted advisers. Their failure suggests that expertise is not simply a matter of codifying knowledge and producing correct answers. The functional role of expertise involves a relationship between expert and client that includes accountability and responsiveness.

This historical case, therefore, supports Quast's critique of reductionism. Expert systems were useful tools, but they did not clearly qualify as experts in the richer sense because they failed to satisfy ExpertiseF.

Contemporary Chatbots and the Illusion of Expertise

The development of LLMs appears to challenge this conclusion. Unlike expert systems, contemporary AI systems are not primarily rule-based but trained through statistical learning on vast datasets. LLM interfaces such as ChatGPT, and specialised, variant systems framed as therapists or advisers can engage in fluid dialogue, respond to follow-up questions, and generate explanations in natural language. This makes them appear more capable of fulfilling the functional dimension of expertise than earlier systems.

However, this appearance may be misleading. Contemporary LLM interfaces ("chatbots") often succeed not by possessing expertise in the balanced sense, but by producing interactions that encourage users to treat them as experts. Maeda and Quan-Haase (2024) analyse this phenomenon through the concept of parasocial trust. Parasocial relationships are asymmetrical relationships in which an individual experiences a sense of personal connection with a media figure despite the lack of genuine reciprocity. In the context of AI, parasocial trust is produced through anthropomorphic design features: conversational warmth, apparent empathy, and the simulation of attentiveness.

These features can create the impression of an expert-client relationship even when the system lacks the capacities that normally ground such relationships. The system does not possess agency, professional stakes, or moral responsibility. It cannot be held accountable in the way that human experts can. Moreover, the trust placed in the system may be based on affective design rather than verified competence.

A prominent example of the challenges faced by AI systems used in clinical contexts is Babylon Health's AI symptom checker and triage service, including the "GP at Hand" chatbot deployed within the UK's National Health Service. Independent reviews have found that Babylon's diagnostic and triage algorithms lacked convincing evidence of performance superior to that of human clinicians and could perform worse in realistic scenarios, raising concerns about clinical safety and reliability (Martin, 2018). These concerns prompted official scrutiny by healthcare regulators, who reviewed the system amid questions about the appropriateness and safety of the advice provided by its chatbot interface (Vince, 2018). The critiques highlight the difficulty of

aligning AI outputs with expectations of clinical judgement and responsibility, particularly when users interact with systems framed as expert advisers.

This raises a further conceptual point. Quast's account presupposes that expertise is primarily an attribute of individual agents. Yet many AI deployments operate within broader socio-technical systems, where responsibility is distributed across developers, institutions, and regulators. It could be argued that ExpertiseF could be satisfied at the level of the socio-technical system (chain) rather than the AI itself. However, if this is correct, then it becomes difficult to describe the AI as an expert agent. Instead, the AI functions as a component within an expert system, with expertise located in the institutional structures that supervise and govern its use.

This distinction matters because users often interact with the AI as though it were the expert, rather than interacting with the distributed institutional system behind it. The result is a risk that expert authority is projected onto systems that cannot genuinely carry it.

Counterargument: AI as Emerging Expertise

A plausible objection is that this analysis underestimates the extent to which contemporary AI systems have developed beyond the limitations of expert systems. Unlike early rule-based systems, modern AI can generate explanations in natural language and respond flexibly to user queries. This could be seen as satisfying part of the functional role that expert systems lacked.

Furthermore, modern AI systems have demonstrated exceptional performance in complex domains. AlphaFold's success in protein structure prediction is frequently cited as evidence that AI can solve scientific problems that had resisted decades of human effort. Similarly, Katz, Bommarito and Blackman (2017) developed a model capable of predicting U.S. Supreme Court decisions with over 70% accuracy, outperforming expert legal scholars. These examples suggest that AI can display not only narrow competence but also significant predictive capability in domains involving complexity and uncertainty.

Finally, it could be argued that ExpertiseF should be understood partly in terms of social reliance. AI systems are increasingly used in medicine, finance, and law, and users do rely upon them in decision-making contexts. If expertise is partly constituted by the social function of providing guidance, then perhaps AI is already fulfilling that role.

On this view, the distinction between simulation and genuine expertise may be less important than the practical reality that AI systems increasingly perform expert functions.

Response: Why Simulation Still Matters

Although the counterargument highlights genuine advances, it does not follow that modern AI systems satisfy Quast's balanced account of expertise.

First, the ability of a language model to generate explanation-like output does not entail that it possesses understanding. The system may produce coherent and persuasive justifications without access to the internal basis of its own outputs. This matters because explanation in expert contexts is not merely rhetorical; it functions as a mechanism of accountability. Human experts can often justify their claims by appeal to reasons, evidence, and conceptual understanding. AI-generated explanations may resemble this structure, but they may also be produced through statistical pattern completion rather than grounded reasoning.

Second, high performance does not settle the question of expertise. A system can achieve impressive predictive success without being an expert in the richer sense. AlphaFold may generate highly accurate predictions, but this does not imply that it understands protein chemistry or can engage in the reflective practices of scientific reasoning. It demonstrates manifestational success, but this may be closer to tool-like competence than expert agency.

Third, user trust does not guarantee genuine expertise. Trust can be rational, but it can also be socially manufactured through design and presentation or as a byproduct of network effect. The fact that users increasingly rely on AI does not demonstrate that the system bears responsibility in the way that human experts do. Quast's functional dimension involves not only being relied upon but also occupying a role where accountability, stakes, and normative expectations apply. AI systems themselves cannot be morally responsible in this sense, even if institutions can be held responsible for deploying them.

The deeper issue is therefore that functional equivalence does not imply conceptual equivalence. AI systems may perform tasks that resemble expert activity, but the processes grounding this performance may not meet the conditions required for expertise as Quast defines it. They may simulate the surface form of expert interaction without possessing the agency and responsibility that underpin the expert role.

Conclusion

The question of whether artificial systems can be true experts cannot be answered purely by appealing to performance benchmarks. It depends on the conceptual structure of expertise itself.

Quast's (2018) framework suggests that expertise involves dispositional competence, manifestational achievement, and functional role fulfilment. Historical expert systems demonstrated that artificial systems can achieve competence and output in narrow domains, but struggled to scale and to function as socially trusted advisers. Contemporary AI systems appear to overcome some of these limitations through natural language interaction and impressive

performance. However, this essay has argued that such systems often generate an appearance of expertise through parasocial trust and anthropomorphic design rather than satisfying the deeper requirements of functional expertise.

AI systems may therefore be best understood as powerful instruments that can support or augment expert practice, rather than as expert agents in their own right. While they can produce expert-level outputs, the absence of agency, responsibility, and socially embedded accountability suggests that they do not fully satisfy the balanced account of expertise that underpins ordinary understandings of what it means to be a true expert.

A final question remains open. This essay has treated expertise as a human-centred concept, grounded in moral responsibility and agency. It is possible that within machine-to-machine contexts, expertise might be evaluated purely in terms of functional reliability. If so, the concept of expertise may itself shift. The more immediate concern, however, is that human users may increasingly treat AI as expert agents even when the social and moral foundations of expertise are not present.

References

Brock, D.C. (2018) 'Learning from Artificial Intelligence's Previous Awakenings: The History of Expert Systems', *AI Magazine*, 39(3), pp. 3–15.

Katz, D.M., Bommarito II, M.J. and Blackman, J. (2017) 'A general approach for predicting the behavior of the Supreme Court of the United States', *PLoS ONE*, 12(4), e0174698. doi: 10.1371/journal.pone.0174698.

Maeda, T. and Quan-Haase, A. (2024) 'When Human-AI Interactions Become Parasocial: Agency and Anthropomorphism in Affective Design', in *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24)*, Rio de Janeiro, Brazil, June 2024. ACM. doi: 10.1145/3630106.3658956.

Martin, E., 2018. *Lancet review raises questions about Babylon Health AI diagnostic performance*. Digital Health. Available at: https://www.digitalhealth.net/2018/11/lancet-review-babylons-ai/ [Accessed 18 Feb. 2026].

Quast, C. (2018) 'Towards a Balanced Account of Expertise', *Social Epistemology*, 33(1), pp. 40–57. doi: 10.1080/02691728.2018.1546349.

Vince, G., 2018. *Safety regulators reviewing concerns about Babylon's chatbot*. Health Service Journal. Available at: https://www.hsj.co.uk/technology-and-innovation/safety-regulators-reviewing-concerns-about-babylons-chatbot/7022612.article [Accessed 18 Feb. 2026].