Crossing the Epistemic Divide

Tracy A. Altman, PhD | Museum of AI

SUMMARY

Now that generative AI is solving math, science, and coding, what is the next frontier? AI creators have an opportunity to benefit more people – but there's a barrier between today's reality and what's possible.

What is standing in the way? There's a divide between the knowledge mastered by today's GenAI and important bodies of interdisciplinary, real-world knowledge. Models need to behave more like human experts on topics where there are no ground-truth answers.

This **epistemic divide** is an obstacle to making human knowledge available to AI models. Crossing this divide will require a fresh mindset:

- We have not exhausted the supply of human knowledge for training.
- AI models can provide definitive responses on problems without ground-truth answers.

The distinct nature of this effort will require new techniques. Getting to the other side won't be as straightforward as sourcing more human experts to annotate more content. Fortunately, several avenues of research and practice can inform choices on the path ahead.

WHAT HAPPENS NEXT

Here, the **Curate** methodology is proposed for capturing post-training data sets on 'fuzzier' interdisciplinary knowledge domains lacking golden answers.

Why should AI creators consider doing this? Deeper model capability and wider user adoption. Rigorous knowledge benefiting human lives, private business, and governments.

THE EPISTEMIC DIVIDE

(abbreviated version)

Problems with ground-truth, unambiguous answers (math, science, engineering)	Problems without ground-truth answers (policy, business strategy, economics, governing)
Coding knowledge	Human expertise & experience
Explicit knowledge: Inform decisions	Implicit knowledge: Make decisions
EXAMPLE TECHNOLOGIES GenAI, RAG, Agentic AI Business intelligence	EXAMPLE TECHNOLOGIES Predictive analytics Decision/risk models

INTRODUCTION

Generative artificial intelligence (AI) models are successfully solving mathematics, physics, and coding. This explosion of AI development has led some to incorrectly conclude that state-of-the-art models have already consumed all human knowledge available for training. That sounds like an unhappy ending.

> "…most models are trained on virtually all human knowledge available today." (Wolf, 2025)

With 'well, this must be it' thinking, numerous business intelligence (BI) vendors failed to innovate beyond shiny dashboards and data visualizations, offering little in terms of explanation or prediction. GenAI creators must avoid falling into a similar trap (Madani, 2025).

No, we're not out of knowledge for model training. AI has not yet mastered interdisciplinary knowledge reflecting human expertise and experience, much of which is implicit and/or lacks ground-truth answers. Vast troves of valuable knowledge are available to those who know where to look. Ingesting it into AI models will require newly designed benchmarks.

> "Benchmarks must expand beyond math and coding to cover commonsense reasoning, causal inference, and ethical decision-making. Real-world performance needs to be the ultimate metric – how well does an AI assist doctors, guide autonomous systems, or navigate complex social interactions?" (Turing Post, 2025)

From certainty to multiple 'right' answers. Producing a trustworthy, comprehensive answer to a 'no ground truth' problem is particularly challenging: More than one response can satisfy the requirement, so judgment is needed to determine which are the most useful answers to each prompt. Rethinking training methods and redesigning model output can add value for the end user and improve their experience with AI products.

More knowledge, more capability. With suitable data annotation and training, AI can assist people grappling with interdisciplinary, real-world problems in business management, policy, and government:

> -What are the critical success factors for implementing a universal basic income (UBI) project? What does UBI success look like?
>
> -What are the impacts of consumer confidence surveys? Do the published findings influence consumer behavior?
>
> -Can governments successfully and economically nudge people toward annual vaccines? What evidence is there to evaluate outcomes of these efforts?

THE EPISTEMIC DIVIDE

To master disciplines in mathematics, science, and coding, AI models are trained on explicit knowledge. That is going quite well by most any measure.

However, tackling other classes of problems requires knowledge of a different nature. Plenty of knowledge is not written down, or is written down obliquely, requiring substantial efforts to extract meaning. But failing to recognize the significance of additional human knowledge would mean failing to apply AI capability to numerous important issues. This puts Gen AI at risk of getting stuck on one side of an epistemic divide.[1]

It's different on the other side. Capturing real-world experience is challenging. Human experts can see what is missing in a given situation and know what the decisions are. It's essential for AI to absorb as much of that knowledge as possible, converting what's unstructured or implicit into structured data.

So how can real-world knowledge be converted into data for training and evaluation? The distinct nature of interdisciplinary, 'fuzzier' domains demands a redesign of traditional methods.

THE EPISTEMIC DIVIDE

Problems with ground-truth, unambiguous answers (math, engineering, science)	Problems without ground-truth answers (policy, economics, business strategy, governing)
Coding knowledge	Human expertise & experience
Explicit knowledge: Inform decisions	Implicit knowledge: Make decisions
Specific domains	Interdisciplinary
Single version of the truth	Multiple true responses
Emphasis on what & operations	Emphasis on why & strategy
EXAMPLE TECHNOLOGIES GenAI, RAG, Agentic AI Business intelligence	EXAMPLE TECHNOLOGIES Predictive analytics Decision/risk models

[1]A recent research paper explores the causes and effects of a 'digital epistemic divide' (Abiri, 2022). However, its focus on equality and societal trust in common epistemic authorities is only tangentially relevant to this paper.

Crossing the Divide

Once a decision had been made to cross over the epistemic divide, some issues require attention. When training to respond to complex, real-world prompts, AI creators should:

PURSUE

- Comprehensive, truthful responses
- Rigorous human data annotation
- Innovative formats and structures
- Benchmarks on real-world performance

AVOID

- Overly long responses
- Traditional presentation
- Fixed training data sets

Taking design cues from these disciplines can inform how to overcome these challenges.

Evidence-based medicine. Complex knowledge is synthesized to provide clinical decision support. AI models guide professionals when answering questions that lack clear-cut, ground-truth answers. In both research and practice, the relative values of knowledge sources are recognized and a hierarchy of evidence is followed (Wikipedia/Evidence-based medicine).

Knowledge representation and reasoning. KRR structures information to formally represent it as knowledge, then understand it and interpret it (Wikipedia/Knowledge representation and reasoning).

Graph theory. This discipline models pairwise relations between objects. A graph is made up of vertices (nodes) which are connected by edges (Wikipedia/Graph theory).

Decision analysis. A popular program at Stanford, DA is interdisciplinary, comprising the activities 'necessary to address important decisions in a formal manner'. Practitioners use visuals to explicitly connect specific actions with outcomes (Wikipedia/Decision analysis).

Data and information visualization. This term refers to several methods of summarizing and presenting complex text or data using visuals. (Wikipedia/Data and information visualization).

Argumentation theory. The interdisciplinary study of how 'conclusions can be supported or undermined by premises through logical reasoning… [including] the arts and sciences of civil debate, dialogue, conversation, and persuasion' (Wikipedia/Argumentation theory).

CURATE METHODOLOGY

Proposed is a methodology for curating data artifacts for model post-training and using them to structure human expert annotation projects. **Curate** is intended primarily for interdisciplinary, real-world knowledge – but it's also applicable to translational math and science research. The purpose is to capture rigorous, transparent knowledge for training data sets.

STEP ONE: CURATING DATA ARTIFACTS

Each artifact is a visual anchor that explicitly states a claim connecting an action with an outcome. These are created in two different ways.

1. Human subject-matter experts are guided by a custom user interface as they curate artifacts. They then annotate evidence supporting the claim being made, creating structured model training data.

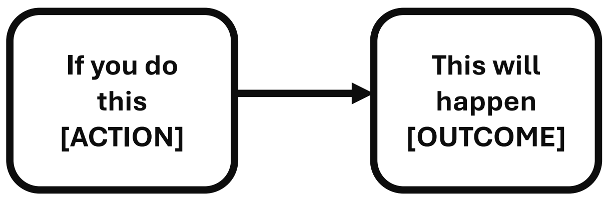

Data artifact design based on Curate methodology

2. Alternatively, an AI model draws a conclusion by reviewing a collection of available documents and then creates artifacts representing that knowledge. Human experts review and annotate these.

Process is critical. In all human annotation projects, effective guidelines and rubrics are essential to collecting reliable training data (MacDonald, 2025).

STEP TWO: PRESENTING MODEL RESPONSES

Here, **Curate** augments the traditional presentation formats employed by today's models. The emphasis is creating a positive experience for professionals in business, policy, and government. Data artifacts relevant to answering a prompt are used to structure a portion of a response, presenting a concise analysis in a simple visual style.

Multiple responses can satisfy a prompt on a no-ground-truth problem, so answers are aggregated. Think of this as guiding a model toward an evidence-based 'golden path' rather than toward a single, concrete 'golden answer'.

Potential Benefits

Higher-quality, rigorous human expert annotations. Dynamic training data sets that continuously adapt based on human and AI knowledge. Deeper model capability and wider product adoption. Valuable assistance for people grappling with important problems in social science, business management, public policy, or government.

REFERENCES

Abiri, G., and Buchheim, J. *Beyond True and False: Fake News and the Digital Epistemic Divide*, J. 7-Apr-2022, SSRN (Social Science Research Network).

Alfariz et al., *Why Does It Fail? A Systematic Review of End-Users' Challenges in Business Intelligence*, 27-Nov-2024, International Journal of Research and Innovation in Social Science.

Bouchard , L. and Iusztin, P. *Our blueprint to cut LLM costs by 88%: Domain-specific AI: RAG + fine-tuning*, 27-Mar-2025, Decoding ML/Substack.

MacDonald, J., Oliver, J., Scowcroft, J. *The Human Touch: Enhance Gen AI with Quality PostTraining and Evaluation Data* (Session S73235), 19-Mar-2025, Nvidia GTC 2025.

Madani, A. *The Hidden Gaps, Challenges, and Failures in Business Intelligence (BI) and Data Engineering (Part 1)*, 21-Jan-2025, Medium.

Turing Post, #90: *Why AI's Reasoning Tests Keep Failing Us*, 3-Mar-2025, turingpost.com.

Wikipedia/Argumentation theory, retrieved March 26, 2025.

Wikipedia/Data and information visualization, retrieved March 25, 2025.

Wikipedia/Decision analysis, retrieved March 3, 2025.

Wikipedia/Evidence-based medicine, retrieved March 25, 2025.

Wikipeda/Graph theory, retrieved March 26, 2025.

Wikipedia/Knowledge representation and reasoning, retrieved March 30, 2025.

Wolf, T. *The Einstein AI model*, 26-Feb-2025, thomwolf.io.

Contact: Tracy Altman
tracy@museumof.ai | 1+ 3036184479
https://museumof.ai
Linkedin: /in/tracyaltman/
Twitter: @UglyResearch