

— WHITE PAPER

AI Experts: Three Designs for Domain Expertise, from *Retrieval* to *Reasoning*

Current Approach and Future Direction



SECTION 01

Why you need *AI Experts*.

Organizations need access to specialized knowledge faster, and at greater scale, than human experts alone can deliver. AI experts address that problem directly — by encoding a defined body of domain knowledge into a system that can be queried on demand, consulted repeatedly, and deployed across multiple projects without taking up expert time.

The work described here explores three types of AI expert, each designed for a different level of sophistication and a correspondingly different use case. These are not successive versions of the same thing. Two properties set them apart — plurality of perspective and depth of grounding — and only the second and third types sit on the same development axis, where the dynamic expert builds technically on the single-domain expert. Section 2.4 explains how the three relate and which ones remain useful as the others advance.

SECTION 02

A *Taxonomy* of AI Experts.

2.1 The Composite Expert

The composite expert brings together multiple perspectives in a single system. Rather than representing the views of one individual, it behaves more like a committee — combining several voices and viewpoints into a single, consultable entity.

Its knowledge is inherited rather than assembled. The composite draws on whatever the underlying AI model already contains, so it tends to be broadly current without anyone maintaining a dedicated corpus. The trade-off is provenance: because there's no specific body of source material, there's nothing to cite.

The design intent is breadth, not depth. Because this expert holds more than one position on a given question, it's well-suited to eliciting reactions, opinions, and evaluative feedback on concepts, ideas, or project directions — situations where the value lies in the range of responses rather than a single authoritative answer.

2.2 The Single-Domain Expert

The single-domain expert is built from a dedicated body of source material. Its knowledge base is assembled from documents, research, and other inputs supplied specifically for that domain.

Its strength is depth and consistency within a defined area. Unlike the composite expert, it's designed to speak with one voice and remain internally coherent across interactions. At this stage, the knowledge base is static or manually updated — it reflects the material available at the point of construction and changes only through deliberate intervention.

2.3 The Dynamic Expert

The dynamic expert extends the single-domain model by removing the assumption of a fixed knowledge base. Rather than relying solely on material uploaded at the outset, it pulls in new research, patents, articles, and other relevant content on a recurring basis.

The distinguishing property is currency. A single-domain expert is accurate as of its construction date. A dynamic expert is designed to stay accurate as its domain evolves.




Figure 1: The three types positioned by number of perspectives available and the state of the knowledge base behind them. The upper-right quadrant is vacant today — that's where a panel of dynamic experts would sit.

FIGURE 01

2.4 Relationship Between the Three Types

The composite and the dynamic expert aren't competing points on a single scale. The composite offers plurality of perspective; the dynamic expert offers depth of grounding. They overlap in one use case only — feedback on a concept or project direction — and that overlap makes it worth being explicit about where each type is headed.

The defining property of the composite is that it preserves disagreement. Collapsing multiple views into a single synthesized position loses the distribution of opinion. Where a concept is genuinely contested, the most valuable signal is often where the disagreement sits and why — not what the average view happens to be. A panel produces a different shape of output, and the need for that output doesn't disappear as individual experts improve.

What changes is how the panel is constructed. The composite panel is one system holding several viewpoints at once. Once dynamic experts can each form a reasoned view of their own, the panel becomes a group of those experts — each with its own knowledge base — asked the same question. At that point, the composite is no longer a separate type of expert; it's simply a way of using several dynamic experts together. The name already implies as much: a composite is a combination of experts, not a different kind of expert.

Today, where a composite blends voices without a dedicated body of source material behind each one, those voices can speak fluently about their topics but can't point back to the original sources that shaped their answers. A dynamic panel is a different proposition — it's like having three actual specialists, each with their own reference materials open in front of them. The first is smoother to listen to; the second can show you exactly where the answer came from. A panel of grounded dynamic experts is unambiguously stronger on that axis.

Until then, the composite is not merely the easier way to obtain several perspectives — it is the only way.



It is worth stating plainly that today the dynamic panel route is not available at all yet. Asking three dynamic experts the same question currently returns three retrievals from three domains, not three opinions. None of them takes a position yet (see §4). A panel requires each participant to be capable of judgment, and that capability is precisely what §5 describes is still missing in today's models. Until then, the composite is not merely the easier way to obtain several perspectives — it is the only way.

Two cases will keep a lightweight composite in service regardless:

No corpus exists.

Standing up a dynamic expert requires curated material and an update pipeline. Where a fast reaction is needed in a novel or diffuse area, a composite is inexpensive and immediately deployable.

Stance matters more than knowledge.

Questions of the form "how would a skeptical finance director react to this?" are about posture, incentives, and priorities. Deeper domain grounding does little to improve the answer.

2.5 What a Panel of Dynamic Experts Would Add

A panel assembled from dedicated dynamic experts produces the same kind of output as a composite — a range of views rather than a single verdict — but improves on it in six meaningful ways:

- 01 Each view is backed by evidence.** In a composite, a perspective is a stance the system adopts. In a panel, the regulatory view comes from an expert whose knowledge base contains the actual regulations, and the basis for any position can be requested and inspected.
- 02 The disagreement is substantive.** Voices generated within one system tend to converge, because they draw on the same underlying material and differ mainly in tone. Separate experts with separate corpora disagree because they're examining different evidence — which makes the location of the split genuinely informative.
- 03 Every panel member stays current.** Each participant inherits the dynamic expert's update mechanism, so the panel as a whole doesn't quietly fall out of date.
- 04 The panel is composed per question.** Regulatory, consumer, and technical experts can be convened for one concept; a commercial and supply-chain pairing for the next. In a composite, the mix of voices is fixed when the expert is built.
- 05 The experts are reusable.** Each one serves individual consultation, panel work, and any project within its domain. The organization accumulates assets rather than configurations.
- 06 Coverage scales by addition.** Extending the system to a new domain means adding one expert — not rebuilding a composite so the new voice sits correctly alongside the existing ones.

These advantages carry real costs. A dynamic panel is slower and more expensive per question than a single composite. Every participant requires a curated corpus and a working update pipeline. And something must convene the panel and present its divergence coherently — that last component is the principal unbuilt piece. Five dynamic experts each returning a long answer is a worse outcome than a composite unless the points of genuine disagreement are surfaced and explained.

Put simply: the advantage is the same panel output, but with each opinion traceable, current, and independently sourced — and with the panel composition chosen per question rather than fixed at the point of construction.



SECTION 03

Architecture of the Dynamic Expert.

The dynamic expert is built from three components:

01

A domain-specific knowledge base

The curated foundation of source material that defines the expert's area of competence.

02

Ongoing updates from relevant external sources

A mechanism for regular ingestion of new research, patents, publications, and related material.

03

Memory of prior interactions within a project

Retained context that lets the expert build on earlier exchanges rather than treating each query in isolation.

Together, these components allow the expert to remain both current and context-aware — not frozen at the moment of its creation.

SECTION 04

Current Limitation and Future Development.

Right now, the dynamic expert is best understood as a highly capable knowledge resource. It retrieves relevant information reliably and responds well to specific factual questions within its domain.

What it cannot yet do reliably is reason across the full body of its knowledge the way a human expert would in order to arrive at a holistic judgment, opinion, or recommendation. That distinction matters: retrieving the relevant evidence and forming a considered position on the basis of that evidence are different operations, and only the first is currently dependable.

The practical consequence is that the dynamic expert today is stronger as a source of knowledge than as a substitute for evaluative expert feedback. Where a project requires the latter, the system should be treated as an input to human judgment — not a replacement for it.



That said, our intended direction of development is to extend the expert from retrieval to reasoning, and we believe this is achievable with targeted adjustments.

One path we are considering is an additional reasoning layer positioned above retrieval, whose function is to synthesize retrieved material into an expert-style point of view rather than returning it as an unresolved body of evidence.

This layer changes how panel-style feedback is produced — it doesn't remove the need for it. Several reasoning-capable experts can be convened on the same question, with their divergence preserved rather than averaged into a single position. That's the function currently served by the composite expert, reconstituted as a mode of use rather than a separate type (§2.4).

In practical terms, this defines a two-stage trajectory:

— TODAY

Today, the dynamic expert is strongest as a continuously updated domain knowledge source.

— OVER TIME

Over time, it could evolve into a system that also provides expert-like opinions and evaluative feedback.

Our objective is a system that can not only surface relevant knowledge but also help interpret it and provide structured feedback on ideas, concepts, and project decisions.

— SECTION 05

Strategic *Value.*

If our development path succeeds, the result is an expert system that is not static but grows with its domain — one that stays current, accumulates project context, and becomes progressively more valuable the longer it's used and the more projects it serves.

The value isn't confined to any single application. Because the architecture is domain-agnostic — the same structure applies regardless of which body of knowledge populates it — it constitutes a reusable, scalable product capability that can be instantiated across categories rather than rebuilt from scratch for each engagement.

Conclusion

All three expert types address the same underlying problem from two directions: getting specialized knowledge to many people at once, and making that knowledge useful when someone has a decision to make.

The composite expert offers several points of view. The single-domain expert offers one reliable view with real depth behind it. The dynamic expert adds the ability to stay current and to remember the project context it's operating within.

Once dynamic experts can reason properly, the composite stops being something you build and becomes something you do: you convene several dynamic experts around the same question rather than commissioning a special-purpose expert for it. The need for a panel doesn't go away as individual experts improve. A single expert has to give one answer — and where views genuinely conflict, one answer either picks a side or averages them. The average hides the disagreement, which is usually the most useful thing to know.

The principal open question remains the transition from retrieval to reasoning.

We are are taking steps to make it happen.