FASCIA RESEARCH · POSITION PAPER · VOLUME I

JUNE 2026 · FIRST EDITION · CC BY 4.0

The Fascia Framework.

A path to *connective* superintelligence. The case that the next leap in artificial intelligence is not a larger monolithic model — but a coordinated network of specialists, integrated like the connective tissue of a single, adaptive system.

Executive summary.

The one-paragraph thesis. Read the rest only if you disagree.

The dominant industry path to Artificial Superintelligence (ASI) is scale. Bigger models, more parameters, more compute. The bet is that capability emerges from size. This paper argues the opposite: **capability emerges from coordination.** A network of specialized intelligences, integrated through a connective layer, will reach ASI faster, more safely, and at a fraction of the compute cost of monolithic scaling. We call this **connective superintelligence**, and the rest of this paper makes the case.

> **Three claims, defended below.**
>
> 1 — Specialization beats generalization at the same compute budget. Always.
>
> 2 — Integration is the bottleneck. The hard problem is not building specialists; it is coordinating them.
>
> 3 — Alignment composes. A network of well-aligned specialists is more aligned than a monolithic model of equivalent capability.

Why this paper, why now

Fascia is an AI research and deployment company building toward ASI. We operate a public surface (the GPT Store: 67+ specialist AI experts, deployed live) and a research surface (this paper, and the volumes that follow). Both surfaces test the same hypothesis: a coordinated multi-specialist system outperforms a single general model on the work that actually matters to humans.

This is Volume I. It establishes the framework. Subsequent papers operationalize it: the FASCIA Capability Index (Vol II) measures progress along six dimensions; the Specialist Economy (Vol III) examines the network effects of compositional intelligence; Alignment Through Composition (Vol IV) presents the safety case.

Monolith vs network.

The choice the field implicitly made — and why it was wrong.

For the past five years, frontier AI research has been governed by a single thesis: **scale solves it.** Larger models, longer training runs, more compute, more data. The path to AGI — and beyond it, ASI — was assumed to run through ever-larger monolithic neural networks. The leading labs raised billions on this thesis. The thesis has, in part, been validated. Capabilities did emerge from scale.

But the thesis has also revealed its limits. Each new generation of monolithic frontier model costs roughly 10× the previous, returns roughly 2× the capability, and concentrates more critical risk into a single artifact. The marginal capability per marginal compute dollar is collapsing. Worse: the alignment problem — how do we ensure these systems pursue what humans actually want — gets harder, not easier, with scale.

Meanwhile, a quieter trend has been unfolding. Across consumer surfaces (custom GPTs, Claude Projects, OpenAI's GPT Store, Anthropic's Constitutional AI variants, agentic frameworks), the highest-performing AI *configurations* are not the largest models. They are **narrower models composed together** — tuned for specific domains, integrated through a coordination layer, given access to tools and one another.

The empirical observation

On a wide range of economically significant tasks — coding, legal analysis, financial modeling, scientific writing, customer operations — a coordinated network of 5–10 specialist models at modest scale consistently outperforms a single frontier model at any scale. The frontier-model gap closes when you go composite. And the cost gap is enormous.

> The most accurate, capable, and economically valuable AI system you can run today is not a single monolithic model. It is a coordinated network of specialists with a connective integration layer. The frontier labs have not yet acted on this.

Three theses of connective intelligence.

The first-principles arguments for why a network beats a monolith. All three must hold for the framework to work. We believe all three do.

Thesis 01

Specialization dominates generalization.

At any fixed compute budget, a specialist model trained on a narrow domain outperforms a general model on that domain. The gap widens with the depth of the domain. Frontier general models pay a tax for everything they don't do well at the things they do.

Thesis 02

Integration is the bottleneck.

Building specialists is easy — fine-tune a base model, ship. Coordinating them is hard. The integration layer (routing, memory, conflict resolution, tool access) is where the actual research frontier lives. This is where Fascia invests.

Thesis 03

Alignment composes.

A monolithic model encodes alignment in opaque weights; misalignment is everywhere and nowhere. A composed system encodes alignment at the boundaries between specialists; misalignment becomes inspectable, auditable, intervenable.

Implication: the path to ASI looks different

If the three theses hold, then the trajectory of frontier AI is wrong. We will not get to ASI by building one increasingly large model. We will get there by building a connective tissue between many models. The metaphor — the brand — is biology. **Fascia** is the connective tissue of the body: invisible, ubiquitous, the structural reason organs become an organism. Same in AI.

Each of the next four volumes operationalizes one of the implications: how to measure progress (Capability Index), how the network effects compound (Specialist Economy), how alignment works in a composed system (Alignment Through Composition), and the operating model behind the work (Operating Model paper).

Measuring the path to ASI.

Preview of the FASCIA Capability Index, the methodology paper coming in Vol II.

To claim a path to ASI, you need a way to measure progress along that path. We propose a 6-dimension Capability Index, scored 0–100 per dimension, summed into a single ASI-Distance metric. The dimensions are chosen to be empirically measurable and conceptually orthogonal.

DIMENSION	WHAT IT MEASURES	TODAY (FRONTIER)
1. Reasoning depth	Multi-step inference under uncertainty. Mathematical, logical, causal.	~64
2. Autonomy horizon	How long the system can operate without human input before drift.	~38
3. Multimodal integration	Native fluency across text, vision, audio, code, action.	~71
4. Long-range planning	Ability to construct and execute multi-step plans toward a stated goal.	~42
5. Alignment fidelity	Degree to which behavior matches stated intent across distribution shift.	~55
6. Scale composability	How well the system integrates with other AI systems and tools.	~48

ASI, in this framework, is a system that scores 95+ on all six dimensions **and** achieves them through composition rather than scale alone. By that definition, the current frontier (mid-2026) sits at roughly 53/100 in aggregate. Most progress in the past 24 months has been in dimension 3 (multimodal) and dimension 1 (reasoning). The hardest gaps remain in autonomy horizon and long-range planning.

The Index is intentionally narrow

We do not claim this Index measures consciousness, sentience, or wellbeing. It measures **operational capability across the dimensions that economically matter**. Volume II expands the methodology and publishes a quarterly tracker. Other labs are welcome to adopt the methodology.

The specialist economy.

Why networks of narrow models will eat the work that monolithic models can't.

The current economic surface of AI is dominated by general-purpose chat. ChatGPT, Claude, Gemini — one box, ask anything. The user experience is great. The economic value capture is small per query because nothing about the system is differentiated for the user's actual problem.

Underneath that surface, a quieter market has formed: **specialists**. Custom GPTs, Claude Projects, agent toolkits, vertical SaaS wrapped around base models. Each specialist is tuned for a narrow domain (legal contracts, financial models, code review, technical writing, customer support, etc.). Specialists capture far more economic value per query because the work product is differentiated.

Fascia operates the GPT Store: **67+ specialist AI experts** across 7 categories — career, business, finance, health, learning, creative, lifestyle. Each specialist is a node in the network. The connective layer (what we build at the platform level) routes user intent to the right specialist(s), composes their outputs, and learns over time. The system gets smarter as the network grows.

67+	7	∞	ASI
SPECIALIST AGENTS	CATEGORIES	COMPOSITIONS	NORTH STAR

Network effects of compositional AI

Adding the 68th specialist does not just add one capability — it adds 67 new *pairs*, and an exponentially growing number of N-tuples. The value of the network is super-linear in the number of nodes. This is the same dynamic that made App Store and AWS Marketplace defensible. We expect the same dynamic to operate in the Specialist Economy, with one important difference: in a network of AI specialists, every new specialist also **improves the others** through cross-pollination of training signal and integration patterns.

Volume III, *The Specialist Economy: Why Networks Beat Monoliths,* publishes the empirical economic model.

Alignment through composition.

Why a network of specialists is safer than a single frontier model of equivalent capability.

The alignment problem in AI is the hardest unsolved problem in the field. As models become more capable, ensuring they pursue what humans actually want — not what they were trained on, not what looks like what humans want, but what humans *actually* want — becomes harder. In monolithic systems, alignment is encoded in weights. The weights are opaque. The misalignment, when it occurs, is non-localizable.

A composed system changes this. Each specialist is small enough that its behavior is inspectable. The boundaries between specialists become natural sites for alignment instrumentation: input filters, output filters, cross-checks, conflict resolution, escalation to human review. Misalignment in one specialist does not propagate to others if the integration layer enforces type contracts.

The composition advantage in plain language

- **Inspectability.** You can audit a specialist. You cannot meaningfully audit a 2-trillion-parameter monolith.
- **Replaceability.** If a specialist fails its alignment evaluation, swap it. The system continues to function. A monolithic model with the same problem requires a full retrain.
- **Cross-checking.** The same input can be routed to multiple specialists; conflicts trigger review. Monoliths cannot cross-check themselves.
- **Bounded scope.** Each specialist's failure modes are bounded by its domain. A misaligned legal-contract specialist cannot start writing medical advice.

> The composition advantage is not that any single specialist is more aligned than a frontier monolith. It is that **the system as a whole has alignment surface area where there was none before.** That surface area is where the safety work gets done.

Volume IV operationalizes this into a Composition Alignment Framework with measurable safety guarantees.

Operating model + open questions.

How Fascia is built to test the framework, and what we genuinely do not know yet.

Operating model

Fascia is structured as an AI research and deployment company — not pure research, not pure product. The two halves are deliberately coupled: research informs deployment, deployment generates the data and edge-case patterns that inform the next research cycle. The GPT Store is the deployment surface where the framework gets tested against real users every day. Research papers (this volume and the ones that follow) are how we share what we learn.

The team is small by design. Researchers, engineers, designers, and operators. No mid-tier middle management. Every team member touches both surfaces — research and deployment.

Three things we genuinely do not know

1 · Compute floor of composition. We claim composition beats monolithic scaling at fixed compute. We have empirical evidence for this at the scales we've tested. We do not know whether the result holds at frontier-monolith scale (10^28 FLOPs and above). Volume II includes a controlled benchmark.

2 · Alignment compositionality. The thesis that alignment composes is intuitive and we believe it. We do not yet have a formal proof that composition preserves alignment across the failure modes that emerge at frontier capability. Volume IV addresses this directly.

3 · Network limits. Networks of specialists exhibit network effects up to some point. We do not know where the diminishing returns curve is. Is it at 100 specialists? 10,000? 1,000,000? Our 67-specialist deployment is too small to tell us. We will keep expanding and reporting.

The invitation

This is Volume I of a multi-year research arc. We publish openly under Creative Commons. The framework is meant to be adopted, critiqued, improved, and operationalized by others. ASI, if it is reached, will not be reached by one lab. It will be reached by a network of labs that figured out how to coordinate — the same way the systems they build will be.

About + citation.

About this paper

This paper is the first volume in the Fascia Framework series, the institutional research output of Fascia — an AI research and deployment company on the path to Artificial Superintelligence. The series is published openly to invite engagement from the broader AI research community.

Forthcoming volumes

- **Volume II · FASCIA Capability Index:** 6-dimension methodology, quarterly tracker, baseline benchmarks across frontier models. Q3 2026.
- **Volume III · The Specialist Economy:** Empirical economic model of network effects in compositional AI. Q3 2026.
- **Volume IV · Alignment Through Composition:** Composition Alignment Framework with measurable safety guarantees. Q4 2026.
- **Volume V · The Operating Model:** Lessons from running coupled research + deployment surfaces. Q1 2027.

Citation

> Fascia Research. (2026). *The Fascia Framework: A Path to Connective Superintelligence, Volume I.* Fascia — An AI Research & Deployment Company. fasciaai.com/research

License

Contact

For research collaboration, citation correspondence, or to join the work:
research@fasciaai.com · fasciaai.com/join

Acknowledgements

The framework builds on a long lineage: Marvin Minsky's *Society of Mind*, the multi-agent systems literature of the late 90s, Marcus Hutter's AIXI, the LessWrong alignment community, and the practical operating intuitions of the dozens of teams building production AI systems in 2025–2026. None of those traditions are responsible for any errors in this paper; all of them shaped its thinking.
