

SMA Method

Signal. Meaning. Action.

The missing layer between "make it pop" and hallucination.

Most communication failures in product work are not signal failures. The request was clear enough. They are not action failures. The team knew how to ship. They are meaning failures, the interpretation layer between the two was never designed.

In product, that gap sounds like “make it pop.” In AI systems, it shows up as hallucination and drift. Same failure, two surfaces.

SMA (Signal. Meaning. Action.) makes that layer explicit. It treats every communication as containing three distinct layers: the signal being sent, the meaning that must be decoded, and the action that results. Most systems optimize the first and last. SMA designs the middle.

01 - The Method

What SMA Is

SMA is a structured approach to one of the oldest problems in design and communication: the gap between what someone asks for and what they actually need. It functions as a diagnostic tool, a communication framework, and an execution architecture.

SMA treats every communication as containing three distinct layers:

Signal - the raw input. A piece of feedback, a user behavior, a query, a stakeholder request. The signal is what is expressed, not what is meant. It is the starting point, not the conclusion.

Meaning - the interpretation layer. What does this signal actually represent? What intent, frustration, goal, or need is embedded in it? Meaning is where most systems fail. It requires context, experience, and structured thinking to decode reliably.

Action - the outcome. Once meaning is established, the action becomes specific and defensible. Not a guess. Not a reaction. A decision grounded in decoded intent.

SMA is deceptively simple. Its value is in the discipline it imposes - the refusal to skip from signal directly to action without passing through meaning. That middle layer is where decisions are either grounded or arbitrary.

SMA applies wherever communication produces outcomes. Client relationships. Product design. Organizational strategy. AI interfaces. Digital experiences. Performance marketing. The context changes. The pattern does not.

02 - The Problem

Make It Pop

After 25 years in design, one phrase kept appearing in every client relationship, at every level of seniority, across every industry:

"Can you make it 'pop' more?"

The phrase is vague, subjective, and almost completely uninformative as design feedback. And yet it appears constantly. Not because clients are bad communicators, but because they are accurately reporting an experience they cannot yet name. Something is wrong. They know it. They cannot explain it.

The traditional cycle is predictable:

Stakeholder: "This design doesn't pop."

Designer: "What do you mean by pop?"

Stakeholder: "It just needs more energy."

Designer makes random adjustments hoping to hit the target.

Both parties are stuck. The meeting goes in circles. Revisions pile up. The actual problem remains unsolved. The loop persists because there is no shared language - no common structure for what the signal means or what a correct response looks like.

SMA emerged as the answer to that cycle. Not by dismissing the vague feedback, but by treating it as exactly what it is: a signal. One that, when decoded properly, contains a precise meaning, and points to a clear action.

SMA Applied - "It doesn't pop"

Signal. "It doesn't pop," the raw, unprocessed feedback.

Meaning. The visual hierarchy is not establishing clear priority. The eye has no entry point.

Action. Increase contrast between primary and secondary elements. Establish a dominant focal point. Reduce visual competition.

The signal did not change. The meaning was decoded. The action became obvious. That is SMA working as intended.

03 - The Evolution

From What-How-Why to Signal-Meaning-Action

The first version of the framework was narrower. It emerged from a specific problem: how to structure a SaaS hero section so that a skeptical user, arriving cold, could understand the product in under three seconds.

That began with What-How-Why:

What. What does the company do? The industry, the product, the category.

How. How do they do it? The technology, the approach, the novel differentiator.

Why. Why should I care? The value proposition, the outcome.

What, Why, and How are not stable categories. In practice, they collapse into each other. A "what" is often expressed as a "why." A "how" is used to justify a "what." This ambiguity is exactly the problem SMA resolves. It does not rely on labeling the parts correctly. It enforces the separation between signal, interpretation, and outcome regardless of how they are expressed.

The three elements stayed constant. The hierarchy shifted based on context. Most SaaS companies elevate the How, because technical users are skeptical and need to understand the mechanism before they trust the outcome. A direct-to-consumer brand elevates the Why, because the purchase decision is emotional and the mechanism is often secondary. Same structure. Different priority.

Meaning is the interpretation layer. The hierarchy it generates is what you ship.

The Abstraction

Applied repeatedly across industries and contexts, What-How-Why revealed itself as a specific case of a more general pattern:

What - Signal. The input, the intent, the context.

Why - Meaning. The interpretation, the processing, the structure.

How - Action. The outcome, the response, the reason any of this matters.

The abstraction mattered because it exposed how far the pattern could travel. The same cognitive structure that decoded "make it pop" and made a hero section work also governed database query design, UI personalization, AI system responses, and social platform interactions. The three-part shape was not about websites. It was about how any intelligent system translates signal into outcome.

Meaning remains central. Hierarchy is what you ship. Both properties held across every domain where SMA was applied.

04 - The Decouple

Separating the Interface from the Architecture

Traditional web design treats UI and IA as a single object. The navigation is the architecture. The page structure is the experience. Every user receives the same hierarchy, the same path, the same sequence of decisions. The interface is built once and handed to everyone identically.

That model worked when information was static and audiences were homogeneous. Neither is true anymore. And the arrival of AI has made the problem impossible to ignore, because AI removed the navigation burden of search without restoring the signal that search preserved.

Two Eras of Information Retrieval

Search. You query. You get a list of sources. You triangulate. Effortful, but the disagreement between sources is visible, and that disagreement is the information. Traditional navigation and UI carry the user through that process.

AI. You query. You get one confident answer. Faster, but the disagreement has been smoothed over. The tradeoffs are invisible. Navigation disappears. The user has less work, and less signal.

The shift from search to AI collapsed the interface, but it also collapsed the meaning layer. Users get answers without context. Paths without choice. The decouple is the structural answer.

When UI and IA operate as separate layers, when the architecture remains intact underneath while the surface adapts to the person using it, the same underlying system serves radically different users without requiring two sites, two navigation structures, or two content strategies. The IA becomes the meaning engine. The UI becomes its output surface.

The interface should not make the user navigate the entire architecture to find the one path that was built for them.

SMA's role in this is to design the meaning layer that makes the decouple coherent. To ensure that the signal a user sends, the interpretation the system applies, and the surface it renders are in deliberate alignment. Experience Avatars are the mechanism that makes it usable.

05 - Experience Avatars

Identity and Intent.

SMA began as a communication framework. Experience Avatars extend it into a full UX ontology. A system for dynamically generating interfaces that adapt to the person using them, not the other way around.

The construct resolves a long-standing tension in product thinking. Alan Cooper's personas describe who the user is, their stable identity, preferences, and patterns. Clayton Christensen's jobs-to-be-done describe what they are trying to accomplish in a given moment, their situational intent. These have operated as parallel frameworks for over twenty years.

Experience Avatars combine both. Identity and intent are not competing models. They are orthogonal dimensions. One defines the user. The other defines their current need. Meaning emerges from their interaction.

Experience Avatars model users as a composition of identity and intent, rather than a static profile.

The core problem SMA consistently surfaces in digital product design is not just bifurcated audiences, but an industry-wide reliance on static user models. Most systems attempt to define who the user is, but fail to account for why they are here right now.

A developer evaluating a product, implementing it, or debugging it are not three different users. They are the same user operating in different contexts. Traditional navigation and information architecture treat them identically. Same structure. Same hierarchy. Same path. The meaning layer is collapsed.

At the UI, data, and protocol level, systems move from signal to action with minimal structure in between. At best, intent is approximated through metadata passed across APIs without a governing model. MCP enables structured tool access, but it does not define how meaning is represented. Meaning remains implicit.

Experience Avatars make that layer explicit.

What Is an Experience Avatar

An Experience Avatar is not a persona picker, and not a selectable character. It is a structured, computable profile - not a descriptive one. It captures three typed layers simultaneously, each with a distinct persistence model and a distinct role in meaning resolution.

Identity is the stable layer. It defines role, goal, and context - who the user is, what they are ultimately trying to accomplish, and the domain they operate in. An environmentally conscious shopper remains one session to session. A graph database developer remains one. Identity persists.

Environment is the session layer. It captures situational context - the use-case or scenario driving this interaction - and physical context: device, location, time available. These are objective constraints. They shift by session, not by moment.

State is the moment layer. It encodes intent, cognitive load, and experience level. This is the layer most systems ignore because it cannot be reliably inferred from behavior alone. SMA treats it as a first-class input, declared or carried forward from entry signal.

The base avatar stays stable. The environment shifts by session. The internal state shifts by moment.

Meaning is not static. It is computed.

Meaning = $f(\text{Identity, Environment, State})$

The system does not guess intent. It composes it.

How Avatars Actually Attach

Experience Avatars enter the system three ways:

Seeded. Defined upfront by the team, based on product knowledge and observed patterns. The Avatar exists before any user arrives.

Mapped. Assigned probabilistically based on real signal: the ad that was clicked, the entry query, the surface a user landed on, the language they used. Meaning carries forward from the message that brought them in.

Declared. Offered as an optional parallel path alongside traditional navigation. Not a gate. Not a requirement. A shortcut for users who already know what they want.

Self-selection alone fails because users are poor at classifying themselves and any friction at the door kills scale. Experience Avatars do not rely on self-selection. They rely on structured hypotheses that are validated through behavior.

Because Avatars can be mapped from entry signal rather than reconstructed through behavioral surveillance, they sidestep cookies, UTMs, and fingerprinting entirely. The user is not being watched to infer intent. The message that brought them in already carries it. Declaration remains optional. Mapping is silent. Validation is behavioral, but only against the hypothesis already in play.

Context and Vibe

Context is layered. It describes both the external situation and the internal state that drives intent.

What users describe as "vibe" is the system's alignment with that context.

Vibe is not taste. It is not randomness. It is the result of how well the system aligns:

- tone
- intensity
- pacing
- familiarity vs novelty

...with what the user needs in that moment.

A Film Buff might prefer slow, complex, Criterion-level films when alone, but something more accessible and emotionally balanced on date night. The underlying identity does not change. The context does. The meaning the system generates changes with it.

The same applies in enterprise systems. A developer evaluating architecture, implementing a feature, and debugging an issue requires entirely different information. Without context, the system cannot distinguish between them, and the experience feels consistently off.

Context also captures the arc of expertise within a given Avatar. A developer who started on a graph database product as a junior engineer is not the same user they were then. Their base avatar still identifies them as a graph developer, but their internal state has evolved. Their tolerance for assumed context is higher. Their need for foundational examples is lower.

A system with a static user model would either keep serving introductory content or force teams to build multiple versions of the same experience. An Experience Avatar adapts. The same content graph serves both the beginner and the senior. As the user progresses, they either test out of beginner-tagged content or self-identify as experienced and immediately see a different projection.

Same content graph. Same design system. Different presentation.

"It doesn't pop" is often not a failure of design execution. It is a failure of contextual alignment.

Avatars at the Data Layer

Avatars are not a UI construct. They are a data-layer construct with a UI projection.

Traditional user profiles capture surface-level identifiers: name, role, organization. Relational systems model identity. Graph systems extend this by modeling relationships between entities as first-class structures. Nodes and edges reveal how data connects.

But these systems stop at structure. They answer what is connected, they do not answer why it matters right now. Experience Avatars introduce that layer.

At the data layer, this can be represented explicitly:

```
class ContextModifier(BaseModel):
    attribute: str
    multiplier: float
    axis: Literal["outer", "inner"]    # environment vs internal state
    source: Literal["sensed", "declared"] # governance layer
```

Each modifier carries both its origin and its type. The system does not just store values. It stores why those values exist.

The avatar itself follows the same logic. At the schema level:

```
class ExperienceAvatar(BaseModel):
    identity: Identity    \# Role · Goal · Context
```

environment: Environment \# Situational · Physical
state: State \# Intent · Load · Experience

Three typed layers. One computable structure. The diagram and the schema are the same artifact.

A model hallucinates when it is forced to reconstruct meaning from ambiguous signal. It infers the user's identity. It infers their intent. It infers which relationships matter. Each layer adds uncertainty.

When meaning is declared at the source and carried through the data contract, the model stops reconstructing. It reads. Declared context is not just cleaner data, it is less ambiguous data. The model's interpretation space narrows. When context is ambiguous, the model fills gaps with probability. When meaning is declared, those gaps close. The failure mode shifts from inference error - which is invisible and compounds - to execution error, which is addressable.

In graph terms, an Experience Avatar is a first-class node. Nodes and node clusters can share an Avatar, which allows interpretive context to travel with the data rather than being reconstructed at the interface.

A restaurant viewed through a Remote Worker Avatar surfaces wifi, noise level, and outlet availability as its most relevant attributes. The same restaurant viewed through a Date Night Avatar surfaces ambience, lighting, and proximity metrics instead.

Same data. Different meaning projection.

This is what allows systems to scale beyond personalization. Interpretation becomes reusable across domains.

SMA Without Avatars

SMA does not require Experience Avatars.

Experience Avatars are one deployment mode of SMA, not the framework itself. SMA can operate through declared identity, inferred context, predefined semantic contracts, or fully headless interpretation systems.

In a headless deployment, meaning is resolved before execution without requiring a persistent user profile or explicit avatar selection. The system derives relevance from structured data, environmental context, behavioral signals, domain rules, or semantic scoring logic.

In the restaurant demo, meaning can be computed entirely through venue attributes and scoring contracts alone. A Remote Worker deployment may prioritize wifi reliability, outlet access, and noise level through predefined semantic weighting, without knowing who the user is.

Avatar-driven deployments extend this model by attaching meaning to declared or persistent interpretive context. Instead of reconstructing relevance from generalized assumptions, the system can project meaning through a known experiential frame.

The distinction is not whether meaning exists. The distinction is where meaning is resolved.

Headless deployments resolve meaning through context and semantic contracts. Avatar-driven deployments resolve meaning through contextual identity.

Both operate within the same SMA pipeline: Signal → Meaning → Action.

SMA requires a meaning layer. Experience Avatars are one mechanism for stabilizing and projecting that meaning across systems.

06 - Avatars in Practice

Where Avatars Live

Experience Avatars can live in two places, and both can be active at once.

Site-native avatars are defined and managed within a specific product. They arrive one of three ways: seeded by the product team before any user appears, based on existing audience knowledge; mapped probabilistically from entry signal, the ad clicked, the query used, the language in the headline that brought the user in; or declared by the user directly as an optional path alongside standard navigation. A seeded avatar is a hypothesis. A declared avatar is a confirmation. Both are explicit, and both produce more reliable signal than behavioral inference.

Portable avatars travel with the user via a browser extension. The base avatar is user-owned. It carries stable identity, preferences, and declared context from one site to the next. Individual sites can contribute domain-specific refinements within that base - a travel platform might extend it with trip context, a professional tool with role-specific preferences - but the base layer belongs to the user and persists independently of any single site's participation.

When both are active, the portable avatar provides the base and the site-native layer contributes refinements within it. Context does not reset on every visit. The user does not fragment across platforms.

SMA-driven tooling can also be implemented as an overlay on existing sites without requiring the underlying architecture to be rebuilt. The avatar layer sits on top. Sites can adopt it

incrementally at key entry points, using early signal to route each user toward the experience that matches what they came to do. The only infrastructure requirement is a CMS that supports a basic tagging structure. The approach is platform-agnostic.

Transparency Without Surveillance

Most personalization systems are optimized to show users less. Fewer options. Fewer tradeoffs. A narrower surface that maximizes the probability of a specific outcome. The logic is conversion-first: the system decides what the user should see, and the user sees that.

SMA inverts the logic. The goal is not to narrow the surface. It is to make the surface coherent for the person using it, while keeping everything else available.

An environmentally conscious shopper has locally sourced groceries surfaced first, because that is the meaning the system decoded from their avatar. But at checkout, when an item in their cart has a cheaper, non-organic alternative, the system surfaces that option explicitly, with the difference named. It does not suppress the tradeoff. It presents it with context. The user decides. The system does not decide for them.

The common pattern: the avatar narrows the interpretation, not the inventory. What gets prioritized is determined by meaning. What gets hidden is nothing.

Because the avatar is explicit rather than inferred, the system can tell the user exactly what it is responding to. Not a recommendation derived from purchase history. Not a behavioral pattern assembled without their awareness. A declared or mapped intent the system is visibly honoring. That is a different relationship than being watched and approximated, and it produces a different kind of trust, one that compounds across sessions rather than eroding them.

System-Wide Propagation

Once meaning is defined at the source, it propagates consistently across every layer:

- API responses
- UI components
- recommendation systems
- checkout flows
- agent and tool outputs

Each becomes a projection of the same semantic contract.

The UI is not guessing intent. The API is not returning neutral data. The model is not reconstructing context on every request. All layers are aligned because they are derived from the same meaning source.

SMA creates a structured foundation for semantic refinement over time. Behavioral signals, projection outcomes, and interaction patterns can be used to iteratively refine scoring contracts, semantic tags, and avatar mappings. Because the provenance of every scoring weight is explicit and version-controlled, adjustments are auditable rather than opaque.

From Personalization to Interpretation

Traditional systems personalize based on past behavior. SMA systems interpret based on present context. The shift is simple, but fundamental:

From: what does this user like?

To: what does this user need right now?

Most personalization systems operate retrospectively. They observe clicks, purchases, watch history, or engagement patterns, then attempt to predict future behavior from accumulated signals. The system becomes a probability engine trained on behavioral residue.

SMA operates differently. It resolves meaning in the present tense.

A user researching graph databases at midnight from a mobile device while searching fraud detection workflows is not merely a demographic profile or engagement segment. They exist within a situational context that changes what information is relevant, how it should be prioritized, and what form the interface should take.

That distinction matters because behavior alone is often ambiguous. The same click can represent curiosity, urgency, confusion, comparison shopping, or active implementation. Traditional personalization systems collapse those possibilities into statistical inference.

SMA narrows interpretation before execution.

Experience Avatars, semantic contracts, environmental context, and scoring logic work together to reduce interpretive ambiguity before the system decides what to surface, suppress, emphasize, or project.

The result is not simply a more personalized interface. It is a more contextually coherent system.

Personalization optimizes around behavioral prediction. SMA optimizes around semantic alignment.

That is what Experience Avatars make possible.

From Content to Meaning

An avatar defines who the user is and what they need. But meaning also has to be applied to the content they encounter. Raw content, a restaurant listing, a law firm practice area, a SaaS feature page, does not arrive pre-interpreted. It has to be scored.

SMA handles this through scoring contracts. Before any content reaches the interface, it is evaluated against the avatar's defined priorities. A Remote Worker avatar scores venues on wifi reliability, noise level, and outlet access. A Date Night avatar scores the same venues on ambience, lighting, and proximity. The content does not change. The meaning applied to it does.

These contracts are defined during intake, before scoring begins. They specify which attributes matter for this avatar, what evidence counts, and how to handle uncertainty. The result is that meaning is computed before rendering, not inferred from behavior after the fact.

The avatar narrows the interpretation. The scoring contract makes that narrowing explicit and auditable.

At execution time, the scoring engine applies those weighted semantic contracts directly to the content layer:

```
def calculate_focus_score(venue):
    return weighted_average([
        venue.wifi_reliability,
        venue.noise_level,
        venue.outlet_access
    ])
```

A scoring contract for a Remote Worker avatar, applied to venue discovery, looks like this:

```
class RemoteWorkerContract(BaseModel):

    wifi_reliability = {
        "weight": 0.30,
        "scale": "0 = absent or unreliable | 10 = strong, consistent, confirmed",
        "evidence": ["reviews", "venue_tags", "photos"]
    }

    noise_level = {
        "weight": 0.25,
        "scale": "0 = quiet, sustained focus possible | 10 = loud, nightlife",
        "evidence": ["reviews", "time_of_day_mentions"]
    }
```

```
}  
  
outlet_access = {  
  "weight": 0.20,  
  "scale": "0 = no outlets | 10 = accessible power, long-session ready",  
  "evidence": ["photos", "review_patterns"]  
}
```

Each venue is evaluated against this contract before it is surfaced. The ranking is not algorithmic preference. It is decoded meaning.

View Schemas

An Avatar does not only determine which content is relevant. It determines the shape the experience takes. A developer who declares they are new to graph databases requires a curated progression - definitions, use cases, architecture examples, starter tutorials, case studies, next steps. That is a schema-driven landing page where the heading hierarchy and content sequence are generated from the avatar, not from a template. The same developer, after applying a fraud detection filter, may prefer a different mode entirely: a binge-style stream of every relevant article, video, documentation page, and product resource within that overlap. Three-column channelized views separate that content by modality - video, editorial, product - letting the user self-select how they consume without losing the avatar-scoped filter.

The underlying content graph does not change. The view schema changes.

This is the decoupling at the content layer. The IA holds. The UI becomes a projection of the avatar's resolved context, rendered in the form that best matches how the user needs to absorb the information.

The Avatar resolves relevance. The view schema resolves presentation.

07 - The UX Layer

The Meaning Gap at Every Altitude

The meaning gap does not only appear at the strategy level, where a product ships without a defined audience, or at the communication level, where a stakeholder says "it doesn't pop" and no one can locate the failure. It appears at every layer of UI execution. The same structural problem that produces vague feedback in a kickoff meeting produces layout instability in a hover state.

This is not a metaphor. It is the same gap, operating at a smaller scale.

A design token carries a value. It does not carry a rule for how that value is allowed to behave. The token arrives. The meaning rule does not. The execution either honors the intent or violates it, and the system has no mechanism to tell the difference. But treating it as a token problem or a CSS problem misframes the cause. The execution rule is missing because the meaning was never declared at the implementation layer.

SMA addresses this at every altitude through the same mechanism: making the meaning explicit before the action executes. In UI systems, this requires a full application layer with its own internal hierarchy. It is what SMA looks like when it goes all the way down.

The Execution Stack

The UX layer has four distinct levels. Each is a surface where the meaning gap can open, and each failure it produces looks different from the others.

The rendered layer is what users see. This is where most people assume the design work ends, and where most UX failures become visible. It is not where they originate.

The execution layer is where values are translated into implementation. This is where token values arrive without governing constraints. The system implements them literally. The result is technically correct and experientially wrong. The meaning of the interaction was never encoded as a rule the implementation was required to follow.

The behavioral layer is where the interface responds over time. Motion, timing, transitions, scroll behavior, feedback tempo. This is not about what the interface shows. It is about how it moves and when. Most interfaces treat this layer as a stylistic afterthought, a set of animation values attached to components without a governing model. SMA treats it as a first-class meaning expression, because how something moves communicates intent as precisely as what it displays.

The integrity layer is the floor. It defines what the system is forbidden to do regardless of input, avatar, or mode. No layout shift. No dimension change on interaction. No spatial disruption that breaks orientation. These constraints hold across all modes and all avatars. They are not preferences. They are the non-negotiable minimum of a coherent interface.

Each level is a place the meaning gap can open. Each level requires an explicit Interaction Pattern Constraint to close it.

Avatars, UX Modes, and UI Behavior

The behavioral layer is where Experience Avatars operate within the UX layer, and the connection is not direct.

Avatars do not control CSS. They do not generate styles, modify tokens, or alter component properties. The chain of influence is longer and more deliberate than that.

An avatar maps to a UX mode. A UX mode is a semantic description of experience intent, a human-readable statement of how an interaction should feel given who the user is and what they need. "Fluid, ambient, unhurried." "Task-oriented, responsive, immediate." "Exploratory, open, discovery-paced." These are not visual specifications. They are experience intentions. They live at the UX layer, above implementation.

A UX mode then maps to a UI behavior pattern, a constrained parameter set that translates experience intent into allowed execution ranges. Motion duration. Scale limits. Easing type. Permitted and forbidden animated properties. This is where the semantic becomes mechanical. Not as a style decision, but as a bounded parameter set the rendering layer must operate within.

The system runs on three behavior patterns. Focus, for task-oriented, stability-first experience patterns. Explore, for discovery-oriented, more expressive patterns. Balanced, for neutral-state interactions that require neither extreme. Three behavior patterns, not infinite tuning. The constraint is intentional. Infinite tuning produces inconsistency, unpredictability, and systems that cannot be debugged or explained. Discrete modes are deterministic. They are inspectable, demoable, and extensible without breaking.

The full chain is:

Avatar resolves UX Mode. UX Mode constrains UI Behavior Patterns. UI Behavior Patterns governs IPC. Interaction Pattern Constraints (IPC) instruct the behavior engine. The behavior engine executes against CSS and animation primitives.

At no point does an avatar directly produce a style value. At every point, the connection between intent and rendering is traceable.

Interaction Pattern Constraints (IPC)

Between the declared meaning of an interaction and its permitted implementation sit the IPCs. These are the mechanisms that ensure execution honors intent without unintended side effects.

An IPC is not a style rule and not a design token. A token describes what a value is. An IPC describes what a value is allowed to do, what it must not do, and under what conditions it is considered to have succeeded.

Applied to a ranked result surfacing in the restaurant demo:

Signal: the top-ranked venue should be visually emphasized when the avatar resolves.

Meaning: emphasis communicates priority, not instability. The card should feel elevated, not repositioned.

IPC: visual emphasis must be achieved without layout shift. The card may receive increased shadow and a border treatment. It must not change dimensions. Adjacent cards must not reflow. Scroll position must be preserved.

Without the IPC, the system knows to emphasize the top result. With the IPC, it knows what emphasis is not allowed to do. That is the difference between a token library and a governed system.

IPC applies across all four levels of the execution stack. At the rendered layer, they define visual output constraints. At the execution layer, they define permitted and forbidden implementation patterns. At the behavioral layer, they define motion parameter ranges by mode. At the integrity layer, they define the absolute prohibitions that hold regardless of context.

The Separation That Holds

The UX layer is what Action looks like when it operates above the implementation level of a UI system. This separation is not semantic housekeeping. It is the structural property that allows SMA to remain domain-agnostic at the core while producing a fully specified execution system at the UI layer. A strategist applying SMA to a product decision and a design systems engineer applying it to a component interaction are using the same method at different altitudes.

The Agility Fallacy describes what happens when meaning is skipped at the process level. The Protocol Level describes what happens when meaning is skipped in AI systems. The UX Layer describes what happens when meaning is skipped at the execution level of an interface, and it names the mechanism that prevents it.

Tokens describe values. SMA defines meaning. Interaction Pattern Constraints govern execution. All three are required. Any one of them, operating without the others, produces a system that works in isolation and fails in practice.

08 - The Agility Fallacy

The Hybrid Stream

The evolution of product development is usually framed as a clean transition from Waterfall to Agile. The reality is messier. Most organizations now operate in a hybrid friction zone where the

velocity of development clashes with the depth required for design. Not Waterfall's cascading drops. Not Agile's closed, reflective loops. Something in between. A stream.

The stream flows constantly. It picks up speed. It lacks the structural integrity of Waterfall and the reflective pauses of true Agile. Because it never stops, there is a bias toward motion. If something is not moving toward a release, it is perceived as stagnant.

That perception is the trap. Thinking looks like stagnation, or worse, slippage. Defining meaning looks like delay. The only legible progress is output.

At the protocol level, code must be written and deployed for a product to exist. That makes Dev the hard requirement. Because design is perceived as soft or visual, it becomes the elastic variable, the first thing compressed when timelines tighten.

When time is stripped from UX to feed a Dev-hungry sprint, the Meaning layer is deleted. The team moves straight from a noisy Signal to a rushed Action.

That is the Agility Fallacy in one sentence: Agile optimizes for action cadence while leaving the meaning layer unmanaged.

SMA Breaks the Cycle

When UX is bypassed, the cycle does not just continue. It decays.

Features ship that do not solve the core problem, because the semantic grounding was never done. That creates Technical Design Debt on top of Technical Debt.

Design systems are meant to be a source of truth. A protocol for consistency. When design is rushed, teams stop contributing to the system and start spinning up one-offs just to clear the ticket. The design system becomes a graveyard of inconsistent components. The connective tissue of the product breaks.

By the time anyone notices, the cost of repair exceeds the cost of having done it right the first time.

To stop the ruin, SMA re-legitimizes the design phase at the system level:

Signal. A Dev ticket is a signal, not a command.

Meaning. This is the non-negotiable step where the UX logic and semantic structure are defined. If this layer is skipped, the Action will be inherently flawed.

Action. Development begins when the Meaning has been mapped to the protocol, not before.

By treating Meaning as a functional requirement, as critical as a database schema, the work moves from a crushed final touchpoint to an integrated architecture. The stream still flows. The meaning layer is just no longer optional.

09 - The Protocol Level

The Meaning Layer in AI Systems

The Agility Fallacy does not stop at the human layer. It extends down into how machines communicate with each other, and with us.

Today's Model Context Protocol servers expose tools that move directly from Signal (the model's interpreted request) to Action (a tool call and its raw result). MCP's contract is structural - it defines what tools exist and how to call them. It does not define what the task means, who is asking, or what outcome is declared. That is a different contract, and an absent one. The meaning layer is implicit, reconstructed by the model on every turn, from surface signals alone.

This is the Agility Fallacy at the protocol level. Fast action cadence. Unmanaged meaning. The model ends up guessing intent, structure, and relationships simultaneously, which is exactly where hallucination and drift originate.

MCP is the current mechanism, not the requirement. SMA operates at the semantic contract layer - the definition of what the entity is, what role it plays, and what the user's declared intent is. That contract can be carried by MCP, by a custom API integration layer, or by any protocol that supports structured meaning alongside raw data. What matters is that meaning travels with the signal, not which pipe it travels through.

SMA requires a single shared meaning layer that sits above the protocol and produces three synchronized outputs from one semantic source:

Human UI. What the user experiences and sees.

Structured Data (JSON / Schema). What machines index and reason on.

Agent Response. What the agent consumes.

All three are projections of the same declared meaning. They are not independent renderings that happen to describe the same thing. They are rendered from one source and therefore synchronized by design. Accidental drift becomes impossible. Deliberate change becomes visible, version-controlled, and auditable. That is the difference between a system that drifts and a system that breaks contracts explicitly.

Why This Removes Guesswork

A model hallucinates when it is forced to reconstruct meaning from ambiguous input. It infers intent. It infers structure. It infers relationships. Three layers of guesswork, compounding.

When the agent response carries explicit meaning - what the entity is, what role it plays, what outcome the user declared they want - the model is no longer reassembling context from surface signals. It is reading a contract. The UI and the agent response share the same semantic root, so the agent cannot drift away from what the user is seeing.

This does not eliminate hallucination. It removes one of the primary conditions that amplifies it - the need to simultaneously reconstruct identity, intent, and relationships from surface signals. The model stops guessing context because context is declared. What remains is execution, which is auditable.

Alignment is a data structure problem, not just a training problem.

RAG and GraphRAG improve retrieval. Relational databases store facts. GraphRAG maps relationships between them. SMA declares the meaning behind those relationships - the intent, the context, and the semantic contract that governs how they should be interpreted and projected across interfaces, schemas, and agents. They are complementary layers, not competing approaches.

"It doesn't pop" has a machine-layer equivalent. "This feels off." "The agent is confused." "The output is wrong but I can't say why."

Vague feedback exists because there is no shared object to point at. A meaning layer gives evaluators, agents, and teams three addressable failure points:

- The meaning definition is wrong.
- The UI projection is wrong.
- The agent response is wrong.

Each becomes its own fixable artifact. Teams stop arguing about output and start correcting the source, which is the only move that compounds.

Two Methods, Two Modes

Deployment method defines where the semantic contract lives architecturally. Runtime mode defines how meaning is resolved during interaction. Experience Avatars define how identity and intent enter the system. These layers operate independently but combine into a unified semantic contract.

SMA is not a single implementation pattern. The semantic contract can enter a system at different points depending on what infrastructure already exists and where meaning is most efficiently declared. Two questions decide how it enters: where meaning resolves, and whether a declared identity drives it. The first answer is the method. The second is the mode. They are chosen independently, which is why they combine.

The Methods: Where Meaning Resolves

Semantic Interface. Meaning is resolved through a stateful semantic layer operating between the underlying system and the rendered surface. Identity, scoring contracts, permissions, behavioral state, and structured context can all participate in meaning resolution at runtime. This method suits enterprise platforms, authenticated SaaS systems, agent ecosystems, marketplaces, and any environment where meaning must persist, adapt, or coordinate across sessions and systems.

Embedded. Meaning is pre-declared and baked into the content. The semantic projections are authored in advance and resolved locally, without requiring a centralized semantic runtime layer. The semantic logic travels with the page. JSON and JavaScript resolve it in the browser, with no database, no enrichment pipeline, and no external calls. This method suits content-first sites, publisher platforms, and any context where the infrastructure requirement must be zero. If the platform supports basic content tagging, SMA can be layered on top.

The Modes: Whether Identity Drives It

Headless. No avatar is declared. Meaning is carried by the content and resolved against the user's filter state. A user weighting one attribute high and another low is not selecting an avatar; they are assembling one through behavior. The system resolves meaning from the intersection of enriched content and declared weights. This mode suits any system where user identity is unknown or irrelevant but content meaning can be structured in advance.

Avatar. An identity is attached at input, seeded from product knowledge, mapped from entry signal, or declared directly. Meaning is grounded at the source. The system does not infer who the user is, because the avatar already answers that. The flow is Signal - (Avatar + Meaning) - Action. This mode suits SaaS platforms, personalized content systems, and enterprise tools where identity drives the experience.

A deployment is one method and one mode, chosen independently. The four demos that follow show how the pairs play out in practice. Whichever pair is chosen, the output structure is the same: a human surface, a structured data surface, and an agent surface, all projections of the same declared meaning. The method and the mode determine how meaning enters the system. The framework governs what happens once it does.

The Demos

These deployment combinations are not theoretical. Each is implemented as a working demo, built against a different industry vertical to prove that the framework is architecture, not analogy. The industry changes. The semantic contract does not.

Venue Discovery - Semantic Interface + Headless.

No avatar is declared. Weight what matters, and the meaning assembles itself. Google Places API pulls raw venue data as signal - name, address, hours, aggregate rating. The enrichment engine applies scoring contracts across noise, intimacy, service formality, kitchen pacing, and group fit, transforming commodity JSON into meaning-enriched venue objects. The user applies weighted filters. A user prioritizing intimacy and low noise is not selecting a category; they are assembling a meaning declaration through behavior. The system resolves which venues surface and in what order based on the intersection of enriched content and declared filter weights. Signal enters one end. A meaning-governed, ranked result exits the other. The avatar is never named. It is constructed from what the user does.

Fintech Dashboard - Semantic Interface + Avatar.

A deal-evaluation dashboard presents one commercial real estate credit deal to four views: Default, Analyst, Partner, and Risk. The raw extracted signal is identical for every avatar: NOI, DSCR, IRR, the comp set, the anomalies. Selecting an avatar resolves a different meaning and a different action surface in place. Risk assumes the deal has a problem and leads with the flags, the audit trail, and escalation. Partner gets a compressed decision view: the key metrics, the comparables, and a single action. Analyst gets source detail and validation methodology. The avatar also selects the behavior mode, Focus, Balanced, or Explore, which governs how the interface responds, while an Interaction Pattern Constraint keeps the flags elevating in place without displacing a single metric card. Same extracted signal. Different meaning. Different execution.

SaaS Marketing Site - Embedded + Avatar.

A marketing site for a data-infrastructure product opens in a default state that speaks to everyone at once, the full content mix. Setting an experience avatar reorders the same fixed content around a declared audience: the developer, the economic buyer, the operator. The hero, the value propositions, and a fourteen-item resource library all reweight in place. The developer sees the SDK references, the quickstart, and the API docs surface first. The buyer sees the ROI calculator, the customer result, and the executive brief. The operator sees the security, compliance, and hardening material. The content is pre-tagged by audience and focus, so no runtime scoring is required; declaring the identity flips the page to the projection already written for it. Same site, same content, three readers, one semantic source.

Law Firm Site - Embedded + Avatar.

One firm site serves several client contexts from a single IA: a default view, a personal injury claimant, a business in a dispute, and in-house counsel. The content reorders without the architecture changing. A business-dispute visitor sees the page lead with litigation proof, case outcomes, and direct language, the litigation partner surfaced first and emotional framing removed. A personal injury claimant sees the same firm reframed around empathy, eligibility, and a clear path to contact. In-house counsel sees the advisory and ongoing-counsel framing. Same firm, same attorneys, same practice areas, reordered and reframed by who declared themselves at the door.

These four demos share one framework and one output contract. The method and the mode are product decisions. The semantic layer is the constant.

Protocol Agnosticism Is the Point

MCP is the tool layer. A2A is the coordination layer. Neither is the meaning layer.

A2A lets agents coordinate. MCP lets agents use tools. Neither defines what the task means for this user - the identity, the declared intent, the semantic contract that makes the result interpretable. The current agent stack assumes meaning is a precondition. In practice it is reconstructed on every hop, from surface signals alone.

A lingua franca for agents that has no word for meaning is a faster way to coordinate incorrect answers at scale.

SMA is the contract that fills that gap. A semantic contract is not a feature of the pipe. It is what makes the pipe worth running. MCP can transport it. A2A can route tasks on it. A custom API can wrap it. The protocol changes. The meaning is the same source.

When the pipes change, and they will, the contract persists. That is not a hedge. It is the structural property that determines SMA's reach.

This is why SMA does not become obsolete when the protocol landscape shifts, and the protocol landscape is shifting faster than any framework can track. The only layer that remains stable across that churn is the one that operates above protocol: the semantic contract that defines what must be present before any system acts on a signal.

Without the contract, faster pipes carry wrong answers faster.

The pipe is not the point. The meaning is the point. You can always change the pipe.

10 - The Paid Media Objection

Replacing Discovery with Definition

Meta, Google, and every performance platform built on probabilistic targeting operate on a single premise: throw variations into a large population, watch behavior, converge on what works. The system learns by brute force.

The flaw is not the optimization. It is the starting point. These systems are reverse-engineering meaning from clicks, because meaning was never defined upfront.

The wasted spend is not in testing. Testing is necessary. The waste is in using paid distribution to discover who the audience is and what resonates with them, signals that should have been defined before any ad went live.

That is the monotonous, low-return phase. Ten headline swaps with no hypothesis. Broad audience blasts to see what sticks. Redundant variations that are basically the same idea. All optimizing toward shallow engagement signals.

Most ad spend is wasted compensating for unclear thinking.

The platform is not wrong to optimize this way. It is the only tool available when the meaning layer is missing. The problem is that teams are paying distribution networks to do the interpretive work that should happen upstream.

An Avatar is a defined hypothesis of who responds and why. It is not a demographic. It is a semantic profile: fears, motivations, language tolerance, skepticism level, desired outcome framing.

With an Avatar defined, the starting point shifts:

Traditional. Ad Variants - Large Audience - Behavior Signals - Inferred Persona - Optimization

SMA. Avatar - Meaning Layer - Ad Creative - Targeted Delivery - Validation

Testing does not go away. It changes role. Instead of testing ten guesses to discover what resonates, you test three expressions of the same defined meaning to refine it. Direct. Metaphorical. Confrontational. Same Avatar, same meaning, different surfaces.

The platform's optimization engine still runs. It is just being fed better inputs.

Meaning Goes Upstream

Meta is excellent at optimizing distribution. It is not designed to define meaning. These are different jobs, and conflating them is where most paid media budgets quietly disappear.

SMA handles meaning before the campaign launches. The avatar is defined, the message is derived from it, and the platform receives a higher-signal input than behavioral inference can produce. The optimization engine still runs. It is just no longer being asked to do interpretive work it was never built for.

The framing is not anti-targeting. It is pro-clarity. SMA does not replace what the platform does. It upgrades what the platform has to work with.

When an ad is generated from a defined Avatar and meaning, the landing page does not have to dump the user into a generic homepage. The Avatar that carried them in also determines what they see when they arrive. Message and surface stay aligned. The user does not have to re-orient. Steps between intent and action compress.

This is the same "one meaning, three surfaces" logic applied to the ad stack. The ad, the landing, and the downstream experience are all projections of the same semantic source.

The instinctive pushback from paid media is predictable and valid: users are bad at self-selecting, and any friction kills scale.

The answer is that Avatars do not require self-selection. They are seeded from product knowledge, mapped from entry signal - the ad clicked, the query used, the language in the headline - and validated through performance. The declared path, when offered, is an optional shortcut, not a gate.

11 - The Unified Theory

A Single Layer, Across the Lifecycle

SMA did not begin as a theory of AI product design. It began as a way to decode vague client feedback. But every time it was applied to a new problem, the same pattern held. Every time the pattern held, the surface area of SMA expanded. At some point it stopped being a diagnostic tool and became something else: a single meaning layer that moves through every phase of building an intelligent product.

What follows is the actual through-line, from the first moment a product idea is named to the last moment it is measured.

Research and Strategy

Every product begins with signal. User interviews, market observation, stakeholder ambition. All of it raw, unfiltered input. Most teams rush to turn that input into roadmap items, feature lists, and brief documents. They skip to Action, and the product is already compromised before a screen has been designed.

SMA forces the meaning layer in at the start. Before anything is built, the team has to name what the product means, to whom, and which interpretive contexts it is actually for. Seeded Avatars, semantic contracts, and contextual hypotheses get defined here, upfront, not reconstructed later from ad performance. Positioning is not a marketing deliverable that gets bolted on at the end. It is the semantic source every downstream surface will project from.

Architecture and UX

The architecture is where meaning becomes structure. This is where the decouple matters most. The IA is not the navigation. It is the meaning engine, the graph of relationships that describes how content connects under different interpretive contexts. The UI is one of its output surfaces, not the thing itself.

When the meaning layer is designed first, the architecture does not have to be rebuilt for every audience. The same graph serves the developer and the decision maker. The same graph serves the exploratory browser and the intent-driven buyer. Resolved context shifts which nodes are elevated. The underlying structure holds. Two sites do not need to be built. One site renders multiple experiences from the same source.

But the architecture's relationship to the UI does not end at content selection. How content is selected is one output. How the interface responds when it delivers that content is another. The UX Layer governs the second.

When resolved context activates a UX mode, the architecture is not just deciding what to surface. It is activating a set of experience constraints that determine how the surface behaves: the tempo of transitions, the energy of interactions, the permitted motion range. The meaning engine and the behavior engine operate from the same declared source. Content and behavior are both projections of the same meaning layer, the same contract, the same resolved context.

Content and Messaging

Content is where meaning becomes voice. Most organizations treat content as an isolated discipline, generated somewhere between strategy and launch, owned by a team that was not

in the room when meaning was defined. The result is copy that contradicts the architecture, messaging that contradicts the product, and voice that drifts from page to page.

SMA collapses that isolation. Every headline, every paragraph, every caption becomes a projection of resolved meaning for a defined interpretive context. The work is not writing. The work is translation.

And this is where “it doesn't pop” stops appearing as feedback. When meaning is explicit and context is declared, vague feedback has somewhere to land. Revisions become targeted. Cycles compress.

Engineering and Protocol

Engineering is where meaning becomes contract. This is the protocol-level application. APIs, schemas, agent responses, tool outputs. Every one of them is a surface where meaning either gets carried forward or gets reconstructed from scratch. In AI-native products, the second option is where hallucination lives.

SMA treats the semantic contract as a functional requirement, alongside the database schema and the API spec. The three surfaces, Human UI, Structured Data, and Agent Response, get generated from one source rather than invented independently. Drift becomes structurally impossible because there is nothing to drift from. The model reads a contract instead of reassembling one.

Launch and Paid Media

Launch is where meaning meets distribution, and where the Agility Fallacy usually wins. Speed is legible. Meaning is not. Teams rush to buy impressions. The platform reverse-engineers the audience from clicks. The meaning layer was skipped upstream, so the system now spends money trying to reconstruct it downstream, one blind test at a time.

SMA replaces that scramble with a defined contextual hypothesis. The semantic target seeded in strategy is the same target the ad is written for. The message, the landing page, and the follow-up sequence are all projections of the same source. Distribution platforms get better inputs. The optimization engines do what they were built to do, and stop doing what they were never good at.

Measurement and Iteration

The final phase is where most organizations fail silently. Metrics come back. Some things performed, some did not. The team reads the numbers and guesses at why. Vague feedback reappears, this time at the system level. The campaign underperformed. The feature did not stick. Users are confused.

With a defined meaning layer, measurement finally has something to compare against. The hypothesis was named. The interpretive context was declared. The projections were explicit. When something fails, the failure is addressable. The meaning was wrong, or the UI projection was wrong, or the agent response was wrong. Each one is its own fixable artifact.

The organization learns. SMA compounds. The next cycle starts from a higher floor instead of a blank one.

Why It Is Unified

Every phase above is usually handled by a different team, a different discipline, a different budget line. Research is one function. UX is another. Engineering is its own empire. Marketing lives somewhere else entirely. The handoffs between them are where products lose coherence, and where intelligent systems built on those products inherit that incoherence.

SMA is the theory that says those phases are not separate. They are projections of one meaning layer onto different surfaces. The same discipline that decodes “it doesn't pop” decodes what a model is hallucinating. The same resolved context that defines a paid media campaign defines the landing page composition the user arrives at. The same semantic contract that governs the agent response governs the human UI.

That is what makes it unified. SMA is not a design method adapted to AI. It is a semantic coordination layer that operates across the full lifecycle of an intelligent system.

A product is not a collection of deliverables. It is a coordinated projection of meaning across every surface it touches.

Most products fail not because the design was wrong or the engineering flawed, but because meaning was never treated as the through-line between them.

SMA is the discipline that makes it one.

12 - The Standard

A Method, Not a Style Guide

Brad Frost's Atomic Design gave the industry a shared language for building scalable interface systems. Atoms, molecules, organisms, templates, pages. It described how interfaces are structurally composed and how those compositions scale across products and teams.

SMA operates on a different but complementary layer. Where Atomic Design focuses on structure, SMA focuses on interpretation. Atomic Design answers: how is this built? SMA

answers: what does this mean in context, how should it behave, and what should it produce? One governs composition. The other governs semantic coordination.

Neither replaces the other. A design system without semantic coordination produces interfaces that are visually and structurally consistent but directionless in behavior and interpretation. SMA without scalable structural systems produces clarity without operational consistency. Together, they create systems that are coherent both in construction and in effect.

If Atomic Design functions as the grammar of a design system, SMA functions as the semantic layer that coordinates meaning across interfaces, behaviors, context, system responses, and adaptive outputs. But grammar and semantics alone still leave one important layer unresolved: execution governance.

A sentence can be grammatically correct and semantically clear while still violating the rules of the environment it operates within. Interface systems suffer from the same problem. A component can be atomically composed and semantically aligned while still being implemented in a way that violates the experience it was intended to create. A hover state defined correctly in a token library can still produce layout shift if implemented without governing constraints. The structure was correct. The meaning was declared. The execution behavior drifted.

Within SMA, IPCs (Interaction Pattern Constraints) govern that execution layer. Tokens define what something is. SMA defines what it means. IPCs constrain how that meaning is permitted to resolve during execution. Atomic Design governs structure. SMA governs semantic coordination. IPCs operate within SMA as behavioral governance mechanisms that preserve intended meaning through implementation.

SMA is not a style guide or visual trend system. It is a structured discipline for coordinating meaning across people, interfaces, systems, and increasingly, intelligent agents. Once internalized, it changes how teams interpret feedback, how systems adapt to context, and how execution decisions are governed across both human and machine environments.

Because SMA operates above the protocol layer, it does not become obsolete when infrastructure changes. Frameworks evolve. Protocols shift. Models improve. The need for semantic coordination remains.

13 - The Adoption Gap

Constructing the Layer

SMA assumes a meaning layer exists. Most systems today do not carry that layer.

This is the starting condition, not a flaw in the model. Meaning has to be resolved before it can be inherited.

In early implementations, SMA operates by defining semantic contracts and applying them to raw content. A restaurant listing, a product page, a law firm practice area, or an API response does not arrive pre-interpreted. It is evaluated against a resolved interpretive context, whether avatar-driven, inferred, environmental, or contract-defined, and meaning is computed before it is rendered. This is the constructed state.

In headless deployments, that resolution may come entirely from scoring logic, environmental context, or predefined semantic weighting. In avatar-driven deployments, the system resolves meaning through declared or persistent contextual identity. The mechanism differs. The pipeline does not.

As systems adopt SMA, that work does not repeat indefinitely. The meaning layer becomes portable and reusable. Entities begin to carry structured interpretation alongside their raw attributes. A venue is no longer just a name and location. It includes the attributes that matter to different contexts, already evaluated, weighted, and stored.

At that point, systems stop constructing meaning and begin inheriting it.

Inheriting the Layer

When two SMA-enriched systems interact, resolution becomes mechanical rather than inferential. The request carries resolved context, whether declared, inferred, environmental, or avatar-driven. The response carries structured meaning aligned to that context. The system selects or composes the projection that best matches the resolved interpretive frame. Reconstruction is minimized.

Meaning is not recomputed at every step. It is read, selected, and, when necessary, combined from compatible structures.

This is where the model shifts from method to layer. The same semantic contract can move across interfaces, APIs, structured metadata, and agent systems without being reinterpreted each time.

Large platforms that already aggregate content at scale are positioned to accelerate this transition. Reviews, metadata, behavioral signals, and domain-specific scoring systems can generate the initial semantic layer. Once attached, that meaning becomes portable. Other systems can consume it directly instead of reconstructing it from scratch.

SMA does not require a team to adopt a new stack, migrate to a new CMS, or rebuild their architecture. The meaning layer can be constructed on top of existing systems, using any protocol, starting now.

SMA begins as a way to interpret signal. At scale, it becomes a shared layer for carrying meaning.

14 - About the Author

25 Years of Application

SMA was not developed in a research context. It emerged from practice. Client meetings, stakeholder reviews, design critiques, site builds and crashes, product decisions made under pressure and rarely with the luxury of being wrong. The technical foundation came from Berkeley, where computer science and fine art fought for my attention without ever fully resolving the tension between them. That tension is probably why the method exists.

The companies and client interactions where it was forged, slowly and without a name for what it was becoming, include Neo4j, Anthem, Callaway Golf, MetLife, Motorola, AT&T, Bristol Myers Squibb, Fortress Investments, and 1-800-Dentist, among others.