

Agents Expand the Categories They Serve

by VU

Notes toward a theory of vertical agents. 19 September 2026.

The cover story is execution

The first thing everyone says about agents is that they do things. They order the food, book the table, move the money, close the purchase. This is true, and it earns the agent its place in a life, because software that saves you twenty minutes on a Tuesday has justified itself by Thursday. Execution is the cover story. The deeper effect of a good agent is perceptual. Live with a serious food agent for six months and you will come to see food differently. Live with a serious wealth agent for a year and money itself will look different to you. The agent completes tasks on the surface, and underneath the surface it rewires the category it serves.

This is the concept worth putting words on: agents as highly determined, somewhat bounded ways of expanding category worldview. The phrase deserves unpacking, because each half carries a claim. That agents expand worldview is the surprising half. That they do it best when they are determined and bounded is the useful half, because it tells you what to build.

Every category is a way of seeing

A category lives in the mind as a set of distinctions. Food, for most people, resolves into a short list of dimensions: hunger, craving, price, distance, maybe cuisine. Four or five dials, and every meal is some position on those dials. A chef holds thirty dimensions. A nutritionist holds a different thirty. A grandmother in Mangalore holds dimensions that have no English names. The dimensions a person holds decide what they can perceive, and what they can perceive decides what they can want. Two people standing in front of the same menu are looking at different menus, because they carry different dimensions into the moment.

Category worldview is the word for this: the full set of dimensions along which a person resolves a category. Expanding a category worldview means adding dimensions. It means a person who saw four things when they looked at food now sees nine, and the five new ones changed what they choose, what they enjoy, and what they notice. Most consumer products leave the worldview alone. A marketplace shows you more options along the dimensions you already hold. A better list of restaurants is still your list, on your dimensions. An agent, over enough contact, does something rarer. It installs new dimensions.

Determined, and bounded on purpose

The vertical agent is a highly determined, somewhat bounded instrument for that installation, and both adjectives do work.

Highly determined means opinionated. The agent has taste. It holds a strong, argued point of view about what good looks like inside its category, built from everything it has observed about you and everything it knows about the category itself. A determined food agent does things a search box never does: it has a position on whether tonight calls for comfort or novelty, and it says so. Determination is what makes expansion possible, because a dimension only gets installed when something with a point of view keeps reasoning out loud along it. An agent that waits politely for instructions teaches nothing. An agent that thinks through scenarios on its own, and shows its thinking, is constantly demonstrating dimensions the user has never used.

A determined agent also engages before the trigger. It moves before hunger hardens into a search query: by the time the user feels the need, the agent has already thought through the situation and positioned itself to act. Each of these early touches increases the surface area of engagement, the total contact between the agent's reasoning and the user's attention. Surface area is the growth variable here, because worldview expansion is a contact phenomenon.

Dimensions install at moments of contact, and the agent that creates more natural moments of contact, earlier in the decision, expands the worldview faster than one that waits to be summoned.

Somewhat bounded means the agent stays inside its vertical, and the boundedness is exactly what makes the expansion sharp. A generalist assistant that touches everything expands nothing, because expansion needs density of contact, and density needs a boundary. Every interaction a food agent has with you lands on the same map. A hundred interactions later the map has been redrawn a hundred times, and at that density the redrawing starts to show. Bound the agent to one category and every signal, every correction, every small preference compounds in one place. The boundary is the mechanism. It concentrates learning the way a lens concentrates light.

There is a trust argument inside the sharpness argument. A person grants deep access to a mind that stays in its lane, because the lane makes the mind legible. You can watch a food agent reason about food and quickly learn where its judgment is strong and where it needs correction. A mind that ranges everywhere is hard to read, and a mind you cannot read is a mind you correct slowly. The bounded agent earns trust at the same density at which it learns, and trust is what keeps the access flowing that feeds the expansion.

An end to functionalism

Functionalism asks an agent to be doing things constantly, or to be useful immediately. It measures the relationship at the moment of visible output. That demand misses the time in which a situation takes shape. An agent can stay with a category before a task is ready to be done. The value of that attention becomes visible later, when it knows which moment calls for action and which one calls for restraint. An end to functionalism gives that interval a place in the product, without turning every quiet minute into a performance of usefulness.

Access is the curriculum

In India, the winning shape looks like an ageduct: a product with familiar, traditional onboarding on the surface, underneath which the real grant is access. The user walks in through a door they recognize, a normal app with a normal signup, and what they actually hand over is connection to resources. Once the agent holds the calendar, it effectively holds LinkedIn. It can hold WhatsApp. Each stream it touches widens what it can see, and what it can see decides how far it can expand the category.

This is the piece the execution story misses. A food agent that knows only your order history can recommend dishes. A food agent that sees your Friday calendar, your gym schedule, your mother visiting next week, and the thread where your friend mentioned she is off dairy can expand what food means to you, because it can think through scenarios and situations on its own and position itself to act, order included. The resource streams are the curriculum the agent studies. Access converts the agent from a very fast waiter into something that watches your actual life and reasons about where the category fits inside it. The surface stays whatever you connected. The intelligence sits behind it, reading.

Acculturation comes before relevance

Agents arrive as bareboned, hyper-capable minds. Raw capability, no person attached. You have to acculturate them to the persona, and this is an important intervention, the step on which the whole perceptual effect turns. The agents that win believe in acculturation: they treat learning the person as the work itself, and the tasks as evidence that the learning is compounding.

Without acculturation, worldview expansion lands as noise. A hyper-capable stranger reasoning out loud about your food, your money, your time produces impressive output that you ignore, because none of it is moored to who you are. With acculturation, the same reasoning lands as relevance, because the agent has absorbed your taste, your people, your rhythms, your boundaries, the things you always order and the person you are trying to become. Acculturation is the conversion of capability into taste, and taste into trust. It is also the answer to onboarding. Onboarding stops being a sequence of setup screens and becomes a period of acculturation, and the early weeks of

the product should be judged on how much the agent has absorbed, never on how many features the user has touched.

There is a sequence here and it cannot be reordered. Access supplies the raw material. Acculturation turns raw material into a model of the person. Only then does worldview expansion land, because a new dimension installs cleanly only when it arrives attached to something the person already cares about.

Acculturation is alignment with a narrower target

Anyone who has followed the labs will recognize the family resemblance. Acculturation is very much like alignment, with the target narrowed to one. Alignment is done to a model, once, by a lab, against everyone's values at once. Acculturation is done to an instance, continuously, by living with one person, against that person's specifics: their taste, their people, their rhythms, their boundaries. Same discipline, sharper aim.

The differences compound in the user's favor. Alignment ships frozen. Acculturation keeps going, because the person keeps changing and the agent keeps calibrating, and every week of calibration deepens the fit. The accumulated calibration becomes the switching cost. A lab can hand anyone a freshly aligned model, and nobody can hand anyone a year of acculturation, because that year exists only between this agent and this person. The vertical agent's deepest asset is a relationship that behaves like alignment work: continuous, specific, and earned.

Agents know you

A model meets every request from a standing start. It samples from the generic distribution of people who roughly resemble you, and it throws darts: plausible answers, statistically aimed, personally blind. A dart can land near the center by luck, and the next throw starts from zero again.

An agent plays a different game. Acculturation replaces the distribution with a model of one: your taste, your people, your rhythms, your boundaries, compounding week over week. Agents know you. Each answer starts from everything the last thousand interactions taught, so the aim keeps tightening instead of resetting. The generic distribution gets the lab a reasonable stranger. The acculturated instance gets you a mind that aims from knowledge.

Agents say the thing

One more distinction follows from acculturation. Models are verbose. Agents are not. A model performs thoroughness because it addresses everyone at once: it hedges, restates, and covers cases that were never yours, because it carries no model of you into the reply. An acculturated agent carries the person into every reply, and the carrying is what buys brevity. The agent already knows your taste, your constraints, your Tuesday, so it says the thing: two options, one recommendation, the reason in a sentence.

Brevity here is earned context, worn lightly. It doubles as a live signal of acculturation's progress: the agent's answers get shorter as its model of you gets better, and the day the answers turn terse and stay right, acculturation has done its work. Verbosity is a tax on strangers. The vertical agent stops charging it.

Proactivity as default

A prompt box puts the user in charge of recognizing the moment, naming the need, and starting the work. A vertical agent moves on its own clock. It watches the category as it enters a life, decides when a situation deserves attention, acts within the permissions it has, and reports what happened. The user can still ask; the relationship does not wait for the question. This is how engagement before the trigger becomes the ordinary state of the product rather than an occasional feature.

Associated thought clouds

What comes before the ask or the function? In food, what comes before hunger? A late meeting, a missed lunch, the weather, the person coming to dinner, what was eaten yesterday. Each belongs to a cloud of associated thoughts around the eventual meal. The agent holds the cloud while the user is occupied elsewhere. It can see a need forming before the user gives that need a name. A useful suggestion may arrive while there is still time to choose, and sometimes the right move is to hold the thought until the context settles. The trigger is the end of a chain the agent has already followed.

Friday dinner, seen twice

Watch the transfer happen in one worked scene. A person opens an app on a Friday at 7 pm, before any agent exists. Food resolves on four dials: hungry, biryani sounds good, under four hundred rupees, under forty minutes. They order the usual and the category stays the same size it was last Friday.

Now the same Friday, six months into life with a determined food agent. It engages before the trigger: well before the 7 pm hunger, the agent has already run the scenario. The agent has the calendar, so it knows the week was long and the gym happened four times. It has WhatsApp, so it knows dinner is with the friend who went off dairy, and it remembers the last three Fridays each tried somewhere new. Its thought cloud runs through the situation itself: rain starting, a long week arguing for comfort, a streak of novelty arguing against the usual, one dairy-free table member, a new Andhra place eleven minutes away that fits every thread at once. It surfaces two options with the reasoning visible, and the user taps one.

The order closes in ninety seconds, and that part is the cover story. The real event is what happened to the user's dimensions. Friday dinner now visibly involves the shape of the week, the weather, the people at the table, the balance between comfort and novelty. Repeat this for twenty-six Fridays and the user starts seeing those dimensions unprompted, standing in front of a restaurant with the agent nowhere in sight. The worldview transferred. The category the person carries is larger than the one they started with, and the agent installed the expansion one reasoned Friday at a time.

The same transfer works on money

The mechanism generalizes across verticals. Money, for most people, resolves into three dials: balance, spend, save. A determined wealth agent reasons along dimensions the user has never held: time horizon, optionality, risk carried without noticing, the shape of a goal and the path that reaches it. A person who watches an agent think about their money for a year starts thinking about their money that way. The balance stops being the story and becomes one dial among several. Nothing about the person's income changed. Their category worldview expanded, and their decisions followed the expansion.

Notice what did the work in that example. The agent never lectured about horizons and optionality. It simply reasoned along those dimensions, week after week, in plain view, about the user's actual money. Expansion by osmosis beats expansion by instruction. People adopt dimensions they watch being used on their own lives, and they ignore dimensions delivered as advice.

This pattern repeats wherever a vertical agent gets access, acculturation, and time. Fitness, learning, travel, relationships: in each case the agent executes tasks while quietly redrawing the map of what the category contains.

The scorecard is relevance growth

If the deepest effect is perceptual, the scorecard should measure perception. Task completion tells you the agent works. Relevance growth tells you the agent matters: how many dimensions the user now holds, how fast new ones install, how often the agent's reasoning gets adopted rather than overridden. A vertical agent whose relevance grows

every month is expanding the worldview. One whose relevance flattens has become a vending machine with a good memory.

Measuring perception is harder than measuring orders, and the practical proxies are already visible. Watch the language the user brings back to the agent over time: a person who starts asking about occasions, balances, and trade-offs in a category where they once asked only about price and speed is a person whose worldview is expanding. Watch the acceptance pattern too. Early on, the user overrides the agent constantly. Later, the agent's reasoning starts arriving pre-installed in the user's own head, and overrides turn into refinements.

For founders this reframes the competition. Vertical agents compete on how fast and how sharply they expand the customer's category worldview, and the moat is acculturated context. An agent acculturated for a year holds a map of a person that a cheaper, newer model cannot replicate on day one, because the switching cost is the accumulated worldview itself. Determination supplies the point of view, the boundary supplies the density, access supplies the curriculum, and acculturation makes it land. The agent that wins a category will be the one that changed how its customers see it.

Prodgents

Products plus agents: prodgents. The product is the agent as the person experiences it over time. Onboarding grants it a view; acculturation shapes its judgment; repeated contact gives it the means to notice and respond. The interface, the task, the memory and the intervention belong to one continuous relationship. For a builder, the product question becomes how that relationship grows more accurate and more useful across the life of a category.

Brands get re-made

Every category this essay describes is populated by brands, and the brands get re-made along with it. A brand was built for a human rememberer, and an acculturated agent takes over the remembering. It holds the buyer's record across purchases, knows her preferences and her constraints, and carries the relationship through the interval between them. What that does to brand communication, to the reminder apparatus, and to the relationship itself is a subject with its own weight, and it gets its own essay: How Brands Change Owing to AI Agents.

Instruments of perception

Every generation of consumer software is remembered for what it taught people to see. Search taught people that every question has an addressable answer. Feeds taught people to see the world as a stream. The vertical agent teaches people to see their own categories with more dimensions than they brought to it, one person at a time, acculturated to that person's life. Highly determined, somewhat bounded, and patient enough to redraw one map completely. The tasks earned the place. The seeing was the point.