Split-Domain Cognition in User Research

Canonical variant. Version 1.0, April 2026.

Preface

About this document. This is a domain-specific variant of Split-Domain Cognition, addressed to user research as practised across product, service, communication, architectural, interaction, and policy design. It works out the collapse of three layers — observation, inference, and recommendation — into a single research deliverable that is presented as continuous empirical finding. The fusion is not an individual researcher's mistake; it is a genre that the design professions share. The variant argues that the three layers can be held apart without losing the discipline's distinctive contribution. Its philosophical grounding draws on Lucy Suchman (situated action), Clifford Geertz (thick description), and grounded-theory methodology.

Where this sits in the corpus. The canon home is splitdomaincognition.org. The variants index is at /variants/; this variant in particular at /variants/user-research-in-design/v1/. The principle this variant interprets is articulated long-form in *Split-Domain Cognition* and short-form in *A Principle, Not a Pattern*. The protocol by which this variant was derived is at /derivation-protocol-v1/.

Authority and version. Canonical, v1.0. April 2026. The website is the source of record. If this PDF and the website disagree, follow the website.

Use. Openly citable. See Governance for the one-person canon and how variants are admitted.

Author. Prayas Abhinav.

Opening

User research, as a profession inside design, occupies an awkward position across every domain in which design is practised — product, service, spatial, architectural, communication, industrial, interaction, systems, social, and policy design all use the method. The profession claims the epistemic standing of empirical investigation (researchers go out, observe real users, and report what they find) while being structurally obliged to produce outputs that close design decisions. The double obligation produces a recognisable working pattern across every design field that uses research: observation, inference, and recommendation are bundled into a single deliverable and presented as though they were a continuous empirical finding. The deliverable takes different names in different sub-disciplines — a research report, an insights document, a findings deck, a synthesis, a site analysis, a programming study, an ethnography, an opportunity map, a stakeholder map, a journey report, a design brief. Each name does the same work — each smuggles the judgement layer into the language of the observation layer.

The pattern is not the fault of individual researchers. It is a genre that the design professions share. A researcher who produced only literal transcripts, site observations, and participant quotes would be told they had not done their job; a researcher who produced only design recommendations without referencing observations would be told they had not earned the authority to recommend. The fusion is how the role is currently defined across design. This variant argues that the fusion is the SDC collapse at industrial scale, that it imposes measurable costs on the organisations and institutions that depend on design research, and that the three layers can be held apart without losing the discipline's distinctive contribution. The examples in this variant draw on product design (where the literature is most developed) but also on architecture, service design, and communication design. The argument applies wherever research feeds design decisions — which is almost everywhere in the discipline.

The language work in this domain

The language work in user research is the literal record of what people did and said in the situations the researcher observed. Plain verbs across the design fields — in product: *clicked, paused, typed, abandoned*; in spatial and architectural research: *entered, turned, looked, waited, sat, left*; in service design: *asked, explained, returned, queued, complained, thanked*; in communication design: *read, skimmed, recalled, mis-recalled, ignored*. Plain nouns across the fields — screens, buttons, forms, rooms, doorways, corridors, desks, counters, posters, leaflets, touchpoints, moments. Timestamps.

Verbatim quotes. Photographs of artefacts and environments. Sketches and plans the participants produced or referred to. What happened, in the order it happened, using the words the participants used where possible.

This is a narrower definition of "research data" than the profession usually recognises. Most industry outputs that the profession calls "data" are already inferred. "Users are confused by the checkout" is not an observation; it is a verdict applied to an observation. The underlying observation might be: six of twelve participants clicked the "review order" button expecting to confirm, paused for between two and nine seconds, then clicked a second button to confirm. Similar patterns in other design fields: "visitors find the entrance disorienting" is a verdict; the observation is that eleven of eighteen participants walked past the entrance, turned back, and approached it from a different angle. "Residents do not engage with the public space" is a verdict; the observation is that the space had an average dwell time of ninety seconds across a seven-day survey. The verdict "confused", "disorienting", "do not engage" in each case presupposes a rubric (what people "should" have been able to do), applies the rubric (they fell short), and names the verdict. A reader receiving the single-word verdict cannot see the three operations — they receive only the output of the third.

Several common research artefacts in design are judgement-work dressed as language work. Affinity diagrams that cluster observations under evaluative headings ("pain points", "opportunities", "needs") have already applied a rubric in the act of clustering — this is a move product design shares with service design mapping and spatial design programming. Personas fuse demographic reporting with aspirational judgement — "Priya is a busy professional who values efficiency" is a character sketch that pre-commits to what the design should optimise for; architectural "user profiles" and service-design "archetypes" make the same move. The *jobs-to-be-done* framework in product design extracts user purposes from user behaviour without marking the extraction as an inferential move; the architectural *programme* and the service-design *journey map* do the same in their fields. Usability scores like SUS and NPS pretend to measure while being evaluation — the score is the verdict, rendered as a number; post-occupancy evaluations in architecture often produce similar composite scores that fold judgement into apparent measurement. Even the phrase "we found that people…" does verdictive work in every design field: *found* presents the inference as a discovery, which obscures the choice the researcher made.

Language work alone is insufficient. This is the move the profession most resists, and it needs to be made explicit. A research deck that contains only literal observations — transcripts, behavioural records, quotes in context, plans annotated with dwell times —

does not close a design decision and is not supposed to. The language layer makes the observation available. Closing the design question requires the judgement layer: a strategy rubric authored outside the research that states what the design is trying to do and what would count as success. Without a rubric the observation cannot be converted to a decision, and any conversion performed inside the research artefact has smuggled in an unstated rubric. Researchers who produce decision-closing artefacts without visible rubrics are therefore operating an implicit rubric that the organisation or client cannot audit. The result is a research practice whose outputs look empirical but whose authority is rubric-based, with the rubric kept out of view. SDC's move is to bring the rubric into view.

The judgement work in this domain

The judgement work is the design decision. In product: build X, kill Y, prioritise Z, redesign W, invest in a different direction entirely. In architecture: approve this massing, relocate this programme, change this circulation, alter this facade. In service design: add this touchpoint, remove this handoff, re-sequence this journey, change this staff role. In communication design: adopt this message, drop this channel, revise this hierarchy, restructure this information architecture. These decisions close the research and ought to be traceable to a rubric. The rubric is the design strategy: what the design is trying to do in the world, what would count as success for its stated purpose, what failure modes are acceptable at what thresholds, how trade-offs among competing goals are weighted.

Most design organisations and clients have a strategy or brief. Most strategies and briefs are not written at the level of specificity a rubric requires. "We want to make financial services accessible to everyone" is a product aspiration, not a rubric. "Create an inclusive public gathering space" is an architectural aspiration, not a rubric. "Improve the patient journey" is a service-design aspiration, not a rubric. A rubric would specify: which people count as the target audience and which do not; which tasks, experiences, or moments are core and which are peripheral; what completion rate, dwell time, comprehension rate, or satisfaction threshold counts as success on each core element; how speed is weighted against care, how efficiency against accessibility, how visual coherence against inclusion; where the design is willing to lose audience members and where losing them would be a strategic failure. A rubric at that specificity is uncommon in every design field because writing it creates accountability. It is easier to operate with an implicit rubric that can be adjusted case by case.

The judgement work is therefore doubly hidden. The first layer of hiding is the fusion inside research outputs, already described. The second layer is the absence of an explicit rubric to which any product decision could be traced. Product decisions that appear to follow from research are, in many organisations, following from an implicit rubric the research has helped the decision-maker rationalise. The research is the rationale; the rubric is the reasoning; and the fusion hides both from the audit.

Applying SDC in this domain means writing the rubric explicitly, keeping it visible, and treating research outputs as the application of the rubric to specific observations. Decisions are then traceable to two things: what the observation showed (Layer 1) and which part of the rubric applied (Layer 2). When stakeholders disagree, they can locate the disagreement — in the observation, in the rubric, or in the inference drawn from applying one to the other.

The collapse, with examples

The collapse in this domain is dense. A partial inventory:

"Users want X." The most common fusion in the profession. *Want* is an inference from behaviour; users rarely articulate wants in the language of the product. The researcher has observed behaviour, inferred a purpose, and translated the purpose into want-language — three operations in one sentence. Stakeholders receive it as a fact.

"Users are frustrated by the checkout." *Frustrated* is a judgement. The observation might be that users paused, looked away, clicked-back, or sighed. Frustration is one interpretation of the behaviour; impatience, distraction, or multitasking are others. The single-word verdict closes the interpretive question.

"Users need a way to..." *Need* is stronger than *want* and carries moral weight. The sentence reads as reporting a requirement that exists independently of the researcher's reading. In practice, need-language is almost always the researcher's inferential move rendered as if it were a user-level fact.

Affinity diagrams. Researchers cluster sticky notes (each representing an observation) under thematic headings ("frustrations", "delights", "pain points", "opportunities"). The clustering is a judgement move — the researcher has decided what the observation is an instance of. Once clustered, the individual observations become invisible and the cluster heading becomes the thing the team acts on.

Personas. A persona is a fictional user assembled from observations. The assembly combines description (demographics, behaviours observed) with aspiration (values, motivations, needs) and with narrative (a day in the life). The three registers are fused inside the persona card. Product teams then design for the persona, which means designing for the fused object rather than for the observations.

Jobs-to-be-done. The framework asks "what job is the user hiring this product to do?" The question is productive because it reframes user behaviour. It is also a judgement move — the researcher identifies a "job" by interpreting what the user did, and the job becomes the design target. The interpretation is not marked as interpretation.

Research decks that open with "We found that users..." The verb *found* carries the authority of discovery. "We observed that six of twelve users did X" is language work. "We found that users need Y" is judgement work in the language of finding.

Usability scores. SUS (System Usability Scale), NPS (Net Promoter Score), CES (Customer Effort Score) — each is an evaluation instrument that reports a number. The number looks like measurement; it is the evaluation compressed to one dimension. The score is often used without its rubric being consulted, which means the team acts on a verdict whose basis is not examined.

"Research says X." A sentence used inside product teams to close arguments. Research does not say anything; researchers say things. The use of research-as-subject distances the verdict from the person who issued it and makes it harder to contest.

Each of these fusions works the same way. An observation has been processed through a rubric the audience cannot see, and the output is delivered in language that implies empirical finding. The collapse is how the research profession's authority currently operates. Naming it is the first move in restoring the distinction the practice claims to honour.

The cost of the collapse

Several costs follow from the fusion.

Design decisions cannot be reconstructed. A decision made in 2023 to build feature X, to relocate a programme in a building, to remove a touchpoint from a service, or to adopt a particular information hierarchy "based on user research" cannot be re-examined in

2026, because no one can reconstruct which observations supported the decision or which rubric converted the observations into the decision. The reasoning has been absorbed into the artefact, and the artefact (the shipped product, the built environment, the running service, the published communication) has become part of the design's assumed history.

Stakeholders cannot locate disagreement. When two people disagree about what the design should do, and both cite "the research", neither can say whether they disagree about the observations (Layer 1), the strategy (Layer 2), or the inference drawn from applying strategy to observation. The dispute becomes a standoff over authority rather than a debate that could be resolved by examination.

Research becomes advocacy with evidence. Researchers, pushed to produce decision-closing deliverables, learn to select observations that support the recommendation they have already formed. This is not dishonest; it is what the genre rewards. Over time, the profession loses the ability to notice observations that contradict its own recommendations, because the genre does not have a place for them.

Organisations cannot audit past decisions. An audit requires reconstructable reasoning. Reconstructable reasoning requires the three layers to be separable. Organisations that run on fused research cannot audit their own history, which means the same mistakes recur because they cannot be diagnosed.

Junior researchers learn the wrong lesson. A junior researcher observing senior practice learns that "research" means producing convincing inference dressed as observation. The skill of holding the layers apart is not taught because it would disqualify most industry research from the category the industry calls "research".

The neutral-reporter role erodes. Researchers who operate the fused genre lose the ability to claim neutrality credibly. Stakeholders begin to treat research outputs as advocacy and weight them accordingly. The research function becomes a political actor inside the organisation rather than an informational one. This is often lamented; it is predictable given the genre.

People are described as what they are not. Personas, user profiles, archetypes, programmatic categories, and audience segments all extrapolate from behaviour to character. The character constructed in the research artefact is not the person; it is a composite that the organisation designs for. When the product launches, the building opens, the service goes live, or the communication is published, real people behave differently from the composite, and the team reads the difference as an anomaly rather

than as a failure of the composite. The composite has become more real to the organisation than the people it was supposed to represent.

The separation, in this domain's language

Applying SDC to user research means writing the three layers explicitly and refusing to bundle them in a single artefact.

Layer 1 — literal record (language work). Session transcripts, observation notes with timestamps, verbatim quotes in context, behavioural logs, artefact photographs, video recordings where permissions allow, literal description of each event in the order it occurred. Uses plain nouns and plain verbs. Does not use *want, need, pain point, frustration, delight, confused, motivated, satisfied, struggling* as interpretive labels. Inferences are marked where they occur, with the inferential step visible. A Layer 1 artefact is something that another researcher could re-read and form their own inferences from.

Layer 2 — strategy rubric (judgement work). The design strategy or brief written at the specificity a rubric requires. Which people are target, which are adjacent, which are out of scope. Which tasks, experiences, or moments are core and which are peripheral. What completion rate, error rate, dwell time, speed, satisfaction, comprehension, or accessibility thresholds count as success on each core element. How trade-offs among competing goals are weighted. Where the design is willing to lose audience members and where losing them would be a strategic failure. The rubric is authored outside the research — it is a strategy artefact, not a research output — and it is revisable as a public act. Research is then the application of the rubric to observations, not the authoring of the rubric in disguise.

Layer 3 — decision narration. The design decision, stated explicitly, with references to the observations that supported it (Layer 1) and the rubric that converted the observations into a decision (Layer 2). Product example: *Given that six of twelve participants completed the checkout in 47–112 seconds (Layer 1), and given that the strategy rubric prioritises completion rate above 80% on the core purchase task (Layer 2), we will invest in simplifying the checkout flow (Layer 3).* Architectural example: *Given that eleven of eighteen visitors walked past the entrance and approached from an alternate angle (Layer 1), and given that the brief prioritises legibility of arrival within the first thirty seconds of encounter (Layer 2), we will re-study the entrance sequence and*

signage (Layer 3). The narration is readable. It can be contested at any of the three layers independently.

When the three layers are held apart, the profession produces what it has always claimed to produce — decisions grounded in evidence — with the difference that the grounding can be inspected. Stakeholders can locate disagreements. Decisions can be audited later. Junior researchers can learn the layers as three skills, each of which has its own practice. The research function recovers the credibility that the fused genre has eroded.

Philosophical grounding

Three older traditions underwrite the SDC move in this domain.

Lucy Suchman's situated action. Suchman's *Plans and Situated Actions* (1987) argued that users' actions in interactive systems are not executions of pre-formed plans but situated responses to local circumstances. The research move that follows from Suchman is to observe the situated action first — what the user did in the situation — before inferring the plan. Plans are constructs the observer applies; actions are what happened. The user-research profession has absorbed Suchman's vocabulary and often ignored her methodological move. SDC restores the move.

Clifford Geertz's thick description. Geertz's *The Interpretation of Cultures* (1973) distinguished thin description (what is literally visible) from thick description (what the action means in the participants' own frame of reference). Thick description is still Layer 1 — it is description — but it is description that includes the cultural context of the action. Applied to user research, thick description means observation that takes the user's situated meaning seriously without collapsing that meaning into the researcher's design frame. The difference between "the user clicked the wrong button" (the researcher's frame) and "the user attempted what they understood as the completion action, in the belief that the label matched the outcome" (the user's frame) is the difference between thin and thick at Layer 1. SDC asks for thick.

Grounded theory methodology. Glaser and Strauss's grounded-theory approach, developed in the 1960s for qualitative social science, requires the researcher to stay close to the data before generating theory. Open coding preserves the observation; theoretical codes come later. The methodology is routinely borrowed by user researchers and routinely compressed — the open coding stage is skipped, and the theoretical codes

are applied directly. The compression is the SDC collapse in the research process itself. Restoring open coding as a distinct step restores Layer 1.

None of these traditions invented the separation; they each named it in their domain. SDC's contribution is to name it architecturally, to show that the same move recurs across domains, and to argue that design-research organisations and design schools can adopt the separation as a deliberate working practice rather than leaving it to individual researchers who happen to have read the methodologists.

Relationship to Koher output

The user-research variant has several landings in the Koher practice.

Prayas's teaching history. User Research Methods has been part of Prayas's teaching curriculum across NID (Bangalore and Gandhinagar — industrial, communication, and interaction design), IIT Gandhinagar (design research), CEPT (architecture and spatial research), Srishti (interdisciplinary design research), and Anant (interaction design). Fifteen-plus years of teaching the discipline across the breadth of design, not within any single sub-discipline. The variant is a writing-down of the methodological argument Prayas has been making to students across these institutions for a long time, now grounded in the SDC architecture. Because the teaching has always spanned product, architecture, communication, and interdisciplinary contexts, the variant's cross-design framing reflects what has already been in the classroom.

studioMeetingCcompanion. The gated onboarding model (`memory/feedback_gated_onboarding_model.md`) is a design-research practice applied to faculty rather than to end-users. Prayas interviews the faculty member (Layer 1 generation), authors a rubric from the interviews (Layer 2 authoring), and the tool runs narration against both (Layer 3). The tool itself is an SDC instance; the onboarding methodology is an SDC instance of the design-research variant.

Koher's own research practice. Koher tools are evaluated in deployment. The evaluation should itself be an SDC instance — what people did with the tool (Layer 1), what the tool is trying to do (Layer 2), whether the tool met its goals on this observation (Layer 3). Keeping the three separate inside Koher's own tool evaluations is a practice the variant can help hold.

Potential curriculum contribution. A research-methods course at Anant or elsewhere — in any design discipline or across several — that taught the three-layer structure explicitly would be unusual. The variant is the spine of such a course. Whether the course is delivered is a separate decision; the variant makes the course possible.

Potential writing contribution. A short essay or article making this argument to a design-research audience is within reach. Targets include product-research publications (*A List Apart*, *Smashing Magazine*, ResearchOps Community), architectural-research journals (*Journal of Architectural Education*, *Architectural Research Quarterly*), service-design venues (Service Design Network, *Touchpoint* journal), and general design-research outlets (*Design Studies*, *Design Issues*, *CoDesign*). The variant's cross-design scope lets it land in more than one of these registers if drafted carefully. The opening-from-outside failure mode (see `memory/feedback_open_from_inside_audience_territory.md`) is less of a risk here — Prayas is inside the territory of design research across fields.

What this variant makes possible

A **vocabulary becomes available** for naming what experienced researchers have long felt without being able to articulate — the double-hatted role, the advocacy-with-evidence genre, the sense that research outputs are more authoritative than they should be and less empirical than they claim. The variant gives the feeling a structural form.

A **methodological reform becomes conceivable.** Organisations that adopt the three-layer separation as a working practice would produce research outputs readable across time. The separation is a change to the genre, not to the skills. Researchers already know how to observe; they already know how to infer; they already know how to recommend. What changes is that each activity produces its own artefact, inspectable on its own terms.

A **teachability becomes possible.** The discipline has long struggled to transmit itself. New researchers learn by imitation, which means they absorb the fused genre along with the useful skills. The variant makes the useful skills (observation, rubric application, narration) nameable as three distinct practices, each with its own craft. The discipline becomes teachable in the same way a three-layer engineering practice is teachable.

The variant does not argue that design research should become more scientific, more rigorous, or more quantitative. These are separate debates inside the discipline. It argues

that the existing qualitative and quantitative practices across every design field are making a structural mistake in how they bundle their outputs, and that un-bundling the outputs would preserve each field's distinctive contributions while making them inspectable. Observation-thick research could be as thick as it wants, as long as it stayed at Layer 1. Rubrics could be as strategic or programmatic or curatorial as they wish, as long as they stayed at Layer 2. Recommendations could be as opinionated as they dare, as long as they referred to both of the above at Layer 3. The practice is not impoverished by the separation; it is made legible across product, architecture, service, communication, and every other field in which design research is practised.

Version 1.0 — 16 April 2026. Read `../sdc.md` *and* `README.md` *before working on this variant. Stake: Prayas's 15+ years teaching user / design research across NID (industrial, communication, interaction), IIT Gandhinagar, CEPT (architecture), Srishti (interdisciplinary), and Anant (interaction); studioMeetingCcompanion's gated onboarding as an SDC-design-research hybrid; potential writing + curriculum contribution across design fields.*