

THE RADIAL MODEL™

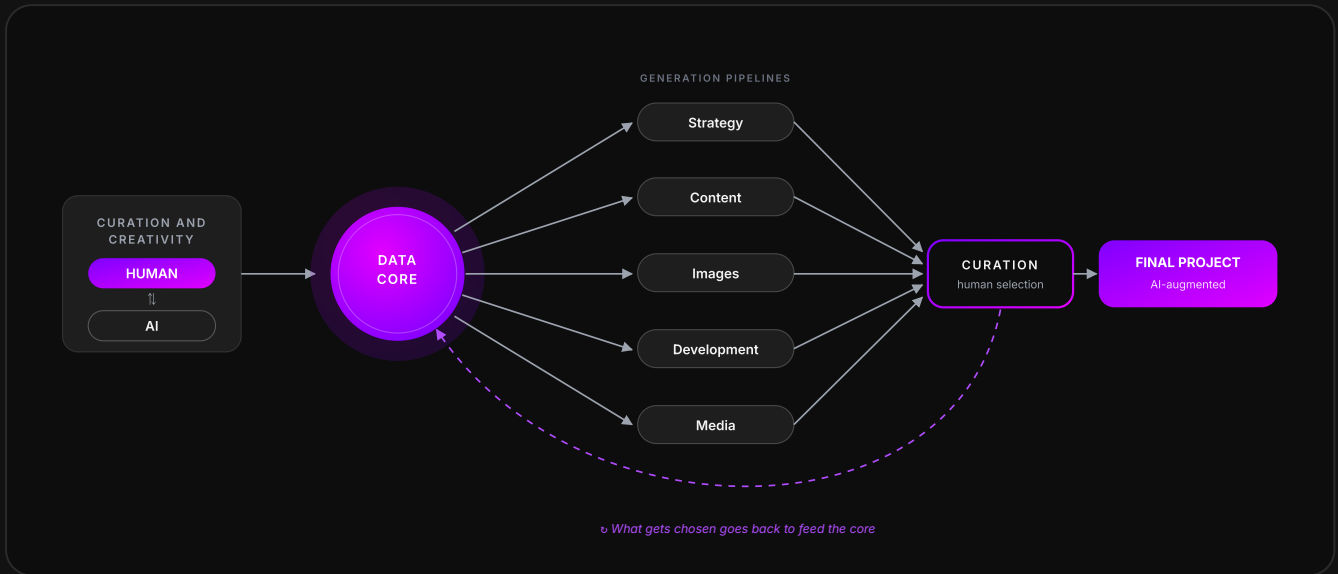
Create with AI, at scale, without losing your soul.

The Radial Model field guide: the full model, the fill-in data core template, typical pipelines, curation checklists, and the traps observed in workshops.

Audric Gagnon

Architect of the model — founder, experimenting.ai

Everything starts from the center. Everything returns to it.



1

The human role: maintain the core and curate. Two functions, both strategic, both irreplaceable.

2

A clean core = coherence at scale. Quality is decided upstream, in the core — not in the prompt.

3

Tools change; the model stays. The radial structure depends on no AI platform.

Three components, one flow

01 The human-AI duo

A closed-loop working relationship. **The human** sets the direction, establishes the constraints, evaluates. **The AI** explores the space of possibilities, generates variants, amplifies. Neither full delegation ("ChatGPT will do the work"), nor a passive tool: a **co-intelligence** where each party does what it does best. Curation is not passively picking from a catalog — it's a creative act in its own right.

02 The data core

The center of the model. The minimal set of **definitive, unambiguous, relevant** data that defines the project: positioning, audiences, message pillars, tone, values, constraints. The single source of creative truth. It's not a technical database — it's a **strategic consensus made explicit**.

Minimalism

the essential, not the exhaustive

Clarity

every term defined

Stability

deliberate revision

Traceability

documented source

Accessibility

reachable by all

Impure core, chaotic outputs; clean core, coherent outputs. In the AI era, the largest share of creative work moves upstream.

03 The generation pipelines

From the core radiate specialized branches of creation — strategy, content, images, development, media — each with its own sequence of steps where human and AI alternate. Pipelines are **modular** (you adjust one step without rebuilding the whole) and they **converge**: their outputs come together in an aligned final project, because they started from the same center.

The loop that closes it all. At the end of the pipelines, curation: the human judgment that keeps what rises above average and rejects the rest. Then the most important gesture in the model: **what gets chosen goes back to feed the core.** The system learns. The next project starts from higher ground.

The data core template

Seven blocks. One to three sentences per element — if a block runs past half a page, it's no longer minimal. Fill it in as a team: the core is a **consensus**, not a memo.

1 · POSITIONING

Who we are, what we offer, to whom, and what sets us apart. One or two sentences, zero slogans.

2 · AUDIENCES

Two to four real segments. For each: context, need, resistance.

3 · MESSAGE PILLARS

Three to five ideas repeated in everything we create. If everything is a priority, nothing is.

4 · TONE AND VOICE

How we speak — and how we never speak. Vocabulary to use, register, examples.

5 · VALUES AND INTENTIONS

What the project stands for. What every creation should make people feel.

6 · CONSTRAINTS

Regulatory, legal, visual, budgetary, technical. The walls of the playing field.

7 · FORBIDDEN

What we never say, what we never show, the taboo subjects. As important as everything else.

GOLDEN RULE

Every element of the core must be definitive, unambiguous and relevant. When in doubt about an element: take it out of the core and settle it as a team before putting it back.

The clean core checklist

Put your core to the test before launching a single pipeline. **One "no" = something to fix first.** Cleaning the core is cheaper than revising fifty misaligned creations.

- Does it fit on one page?** A core that overflows isn't a core — it's documentation.
- Is every term defined?** "Innovative," "premium," "accessible" — defined, or deleted.
- Are there contradictions?** "Bold" and "cautious" in the same core = chaotic outputs guaranteed.
- Are priorities settled?** If two pillars conflict, which one wins? The core must say so.
- Has the team validated it?** A core written by one person is an opinion. Validated by the team, it's a consensus.
- Does every element have a source?** Team decision, research, client brief — you know where it comes from.
- Is it accessible to all?** Humans and AI. A core inside the creative director's head doesn't count.
- Are the forbiddens written down?** What you never do is as structuring as what you do.
- Would it survive your absence?** The ultimate test: a new person + the core + an AI = a result in the spirit of the project?
- Does it have a revision date?** Stable doesn't mean frozen. Plan revision by cycle — not on the fly.

THE TRANSFER TEST

Give your core to someone who doesn't know the project, with access to an AI. If what comes out respects the spirit of the project without further explanation, your core is clean. If not, whatever that person was missing — that's exactly what your core is missing.

Three typical pipelines

Pipelines change with the domain; the structure never does: core → steps (human ≠ AI) → curation → loop. At every step, the AI generates and the human decides.

Marketing / campaign

CORE → Campaign strategy → Positioning & messaging → Visual identity → Campaign concepts → Content creation → Media plan → CURATION → ∞

The AI explores message and aesthetic variants; the human validates strategic alignment and the brand at every node.

Editorial / content

CORE → Detailed brief → Directed brainstorm → Editorial angle → Intentions & audiences → Development & assets → Visuals → CURATION → ∞

The human checkpoints carry the editorial responsibility: verification, tone, mission. The AI accelerates; the human answers for the content.

Podcast / narrative

CORE → Concept → Development → Synopsis → Narration → Emotions & audience → Moodboard → Storyboard → CURATION → ∞

The visual branch (moodboard, prompt extraction, image iteration) runs in parallel with the content branch — and reconverges before final curation.

BUILD YOUR OWN

- List your current creation steps — as they actually happen.
- For each step: who decides? What can the AI generate? What must the human settle?
- Mark the mandatory human checkpoints — the places where you never delegate.
- Define each step's output: what moves on to the next one?
- Five to eight steps are enough. Too short = no structure; too long = bureaucracy.

The curation checklist

Curation is the anti-average filter: reject what is merely "fine" and keep only what beats the algorithmic consensus. Five questions, in order, for every generated creation.

1 Is it faithful to the core? Positioning, tone, pillars, forbiddens. If not — reject, whatever the quality. A beautiful off-core creation is more dangerous than an average one inside it.

2 Does it beat the average? Would anyone with the same AI have gotten this? If yes, it's the statistical average — iterate or reject.

3 Is it true? Facts, figures, quotes, references: verified. Responsibility for what ships is human — always.

4 Does it serve the intent? Not "is it good" — does it accomplish what this piece was meant to accomplish, for this audience, at this moment?

5 What does it teach us? If a choice repeats from project to project, it may belong in the core. That's the loop: curation feeds the source.

RED FLAGS

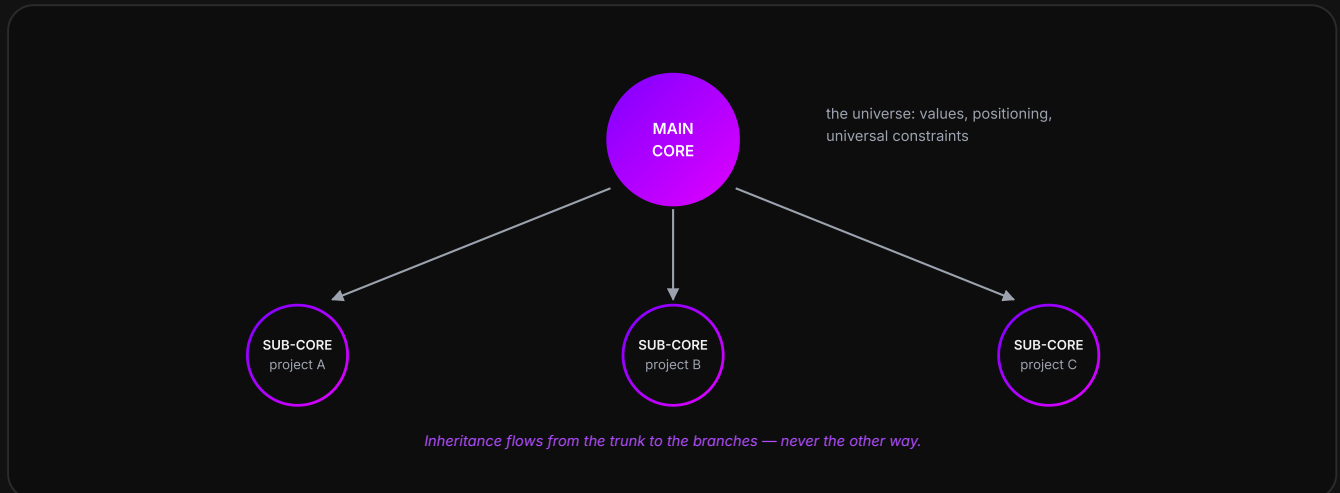
- You approve everything — curation has become rubber-stamping.
- You rewrite everything — the core or the pipeline is broken, not the generation.
- No one can say why a piece was kept — the criteria aren't explicit.

TO MAKE IT SCALE

- Written criteria (the 5 questions) — curation no longer depends on a single person.
- Sample when volume explodes: not everything deserves the same scrutiny.
- Tier it: junior curation first, senior arbitration on edge cases.

The core tree

The model's known limit: **a core that grows too large overflows the context window of AI tools.** A universe documented across dozens of pages can't be swallowed whole at every generation — it consumes, it slows down, it dilutes the direction. The answer: grow the core **as a tree.**



WHEN TO SPLIT

- The core exceeds what the tool can absorb without diluting the direction.
- A recurring project (series, podcast, annual campaign) develops its own constants.
- Two projects pull the same core in different directions.

INHERITANCE RULES

- A sub-core **inherits** from the main core and carries only what its project needs.
- A sub-core **never** contradicts the main core — if it must, it's the main core you revise.
- The connection between cores isn't automatic yet: the human maintains the lineage. That's the model's current frontier.

The traps observed in workshops

Six ways to make the model fail — all observed in the field, several times.

THE TRAP	WHAT HAPPENS	THE FIX
"Let's just start creating" core ignored or fuzzy	Incoherent outputs, endless revisions, "AI doesn't understand our vision."	Stop. The core first — even a single page changes everything.
"ChatGPT will do the work" full delegation	Lost intent, generic content, statistical average.	Mandatory human checkpoints at the pipeline nodes.
"We take what comes out" curation missing	Mediocre work ships, the brand erodes, errors accumulate.	The 5 curation questions, written down and applied.
"Every project is unique" ad hoc pipeline	Everything gets rebuilt every time; no velocity ever sets in.	One documented pipeline per domain, adjusted — never improvised.
The prompt-as-core confusing tactics and strategy	Everyone uses the same prompt: convergence toward the same result (the groupthink documented by Wharton).	The core gives the direction; everyone pilots their own prompts in their pipeline.
The encyclopedia-core documenting everything "just in case"	Overflows the tools' context, dilutes the direction, nobody reads it.	Minimalism — and the core tree when the project grows.

THE SYMPTOM THAT SUMS IT ALL UP

If your team spends more time revising the AI's creations than defining what it expects from it, the problem isn't downstream — it's in the core. **We automate execution before clarifying intent: that's the root trap all the others flow from.**

Start tomorrow morning

No transformation program needed. One project, one team, five steps.

- 1 Pick a real, recurring project.** Not an artificial test — something you already produce regularly.
- 2 Fill in the data core template as a team** (page 4). One working session. The disagreements that surface are the real deliverable: settle them.
- 3 Run the clean core checklist** (page 5). The transfer test is non-negotiable.
- 4 Draw your first pipeline** (page 6) and produce one full cycle, core → curation, with the feedback loop.
- 5 Compare with your old way of working.** Coherence, velocity, revision load. That's your internal proof — and your argument for extending the model.

The model's canonical reference lives at experimenting.ai/radial — versioned and dated. To cite the model:

Gagnon, A. (2026). The Radial Model: a model for AI-augmented creation (v1.0). experimenting.ai.

THE MODEL IS BEST LEARNED BY PRACTICING IT

Workshops, adoption pathways and guidance: experimenting.ai · hello@experimenting.ai