WANDER · COMPANION DOCUMENT A

Wander: Product Design & Development

How I designed and built a discovery product that starts with one action, and what I learned testing it.

AUTHOR	Shawn Farnum · product design, UX research and design engineering
PROJECT	Independent · September 2026 to present
PRODUCT	wandr.it · publicly available · v1.0 in progress
CASE STUDY	shawnfarnum.com/work/wander
VERSION	1 · November 4, 2026

This document goes deeper than the portfolio case study. It covers Wander as publicly released (v0.7.8, October 27, 2026), and labels anything built since that release as local work that isn't deployed.

The companion Research & Evidence document has the study methods, aggregate results and decision records.

Contents

01 Product framing

02 My role, and how I worked with AI

03 UX and interaction decisions

04 The discovery-selection architecture

05 Discovery-graph evaluation

06 Experiments and rejected alternatives

07 Research interpretation

08 Accessibility and reliability

09 Implementation tradeoffs

10 Current limitations

11 Future product questions

01

Product framing

Search works when I can already name what I want. Wander is for the other moment: when I'm curious but don't have a query yet.

I missed a kind of internet experience that search doesn't offer: stumbling into an unfamiliar archive, an odd interactive site or someone's small obsession, then following a trail I couldn't have planned. StumbleUpon was an early influence. My product question was whether discovery could feel less like a feed and more like opening one door at a time.

Three commitments came out of that question and shaped everything after it:

- **The destination gets the attention.** Wander is a doorway. Time spent inside Wander isn't the goal, so it isn't the measure either.
- **No taste profile up front.** The product can't ask someone to describe their interests before they've experienced it. It has to learn from what they do.
- **Destinations are reviewed, not generated.** Every Discovery is a curated record that a person approved. Nothing is invented by a model at request time.

The core loop

Arrive, press Wander, open the website if it interests you, come back when you're ready, and continue. Each step adds to a trail. The trail is how the product knows where you've been, and it's also what you see of your own path.

02

My role, and how I worked with AI

Wander is an independent project. I set the product direction and its central interaction principle: wandering should teach the experience. I designed the first-visit hierarchy, when each concept and control appears, the path from anonymous use to an optional account, and the visual direction. I wrote the research questions and ran the studies used to challenge those decisions, and I made the scope and release calls.

I also implemented the working web product with AI assistance. AI helped me write and revise code, tests and analysis. I decided what the product should do, reviewed what came back, changed direction when evidence or the implementation contradicted the intent, and controlled what entered a release. The repository began with an imported frontend prototype, so I don't claim that every line or visual originated with me.

> **Where design became system behavior.** "Never show me that as new again" couldn't be a preference stored on one screen. It had to survive new trails, reloads and someone creating an account halfway through. That one product promise decided where history lives (on the server, per identity), how a skipped destination is treated (seen, but not disliked) and what happens when the collection runs out (Wander says so instead of recycling).

03

UX and interaction decisions

Wandering is the onboarding

My starting hypothesis was that the product could teach itself through use. I wanted someone to experience a Discovery before learning any controls or being asked for an account. I think of a street magician who shows you the trick before selling you a ticket: give a reason to care before asking for any commitment. The first screen is a one-line premise and a single Wander button.

The teaching sequence

After the first Discovery, the interface adds one concept at a time, each at the point it becomes useful. The sequence is tracked per identity, so a returning visitor isn't taught the same thing twice.

TEACHING SEQUENCE IN THE V0.7.8 RELEASE

Moment	What appears	Why then
Arrival	Premise and the Wander button	One choice. Nothing to configure, nothing to sign up for.
After Discovery 1	Trail	There's now a path to show.
After Discovery 2	Familiar, Explore, Wild (How Far?)	With a previous step, "how far from here" has a meaning.

Moment	What appears	Why then
After Discovery 3	Reactions: Interesting, More like this, Not for me	Enough of a trail for feedback to steer something.
After Discovery 4	Wander From Here	Starting from your own URL makes sense once the loop is familiar.
After real history	An optional account offer	Only when there's a trail worth preserving.

The timing is implementation evidence. It shows when features enter, not what participants learned from them.

The tradeoff is real: someone can leave before meeting the later features. I accepted that because a first visit doesn't need to be a tour of the interface. The question is whether the central experience makes sense and is worth a person's time.

Accounts preserve; they never unlock

Wandering never requires an account. The offer appears only after meaningful history: eight unique Discoveries with at least three reactions or saves, or, for someone returning a day or more later on a second trail, six Discoveries with one signal. After it has been shown, it won't appear again for 30 days, and a dismissal is respected. An account preserves history across devices. It doesn't change what Wander can do.

Familiar, Explore and Wild change the next request

The How Far? control expresses intended distance from where you are: Familiar stays near, Wild leaps. Explore is the default. Choosing a mode changes the weighting of the next Discovery's score (section 4). It doesn't navigate anywhere by itself, so it's a steering choice, not a second button that does the same thing as Wander. These labels express intent; they aren't three separately validated experiences.

Reactions, saves and skips mean different things

Reactions shape direction. Save is a separate bookmark and doesn't change ranking. A destination you skip is remembered as seen, so it won't come back as new, but it isn't counted as a dislike. Keeping those three apart avoids punishing a destination just because someone moved on quickly.

Visual direction

Editorial typography, generous space and warm light and dark themes keep the interface quiet so the destination carries the attention. These are my design intentions. I haven't measured emotional outcomes and don't claim them.

The mobile hierarchy

A local layout investigation found Keep wandering, the primary action, below the first screen on short phones, after the destination preview and description. I compared tighter spacing, a shorter preview and moving the action earlier. I kept the preview and moved Keep wandering closer to the destination link. The evidence was a measured comparison of matched local layouts, not a participant study.

I also tested putting How Far? above Keep wandering. It made the link between modes and the next Discovery explicit, but the extra height pushed the primary action lower in three of four matched short journeys and added three keyboard stops, so I rejected it. A disclosure menu would have hidden the choices behind another action. A compact variant that places the modes after Keep wandering is frozen as a research comparison candidate. No participant comparison result exists yet, so it isn't a winner.

04

The discovery-selection architecture

The runtime is topic-based scoring with a small random draw. It isn't graph traversal, and it doesn't use embeddings.

The Discovery record

Each Discovery is a curated record with a canonical URL, title, description, content type, topic labels, an editorial quality score, an obscurity score, its domain and a publication status. The corpus is a versioned, reviewed index that the application loads. It isn't generated live.

One request, end to end

The client sends an action (wander, wander from here, start a new trail, react, save). A single server-side operation handles it, independent of the web framework, so every client and the database tests exercise the same path. The visitor's trail, history, reactions and saves are stored in Postgres per identity, so the rules hold across devices and return visits. The selection code stays behind the API.

STEP 1 · CURATED DISCOVERY INDEX

All records, loaded from the reviewed corpus.

STEP 2 · ELIGIBILITY AND SEEN FILTERING

Keep only published records the identity has never been shown, counting every trail and the current one. A Wander From Here starting point that matches an indexed Discovery is excluded too. If nothing remains, the request reports that the visitor has wandered everywhere Wander can currently take them.

STEP 3 · TOPIC AND CONTEXT SCORING

Score every remaining candidate on the factors below, weighted by the current mode.

STEP 4 · CANDIDATE SELECTION

Sort by score and pick uniformly at random from the top three.

STEP 5 · NEXT DISCOVERY

Return it, append it to the trail and record it as seen. The new trail position and history feed the next request.

Scoring factors

Topic overlap between two sets is the number of shared topics divided by the size of the larger set. When either side has no topics, overlap falls back to a neutral 0.35.

FACTORS IN THE V0.7.8 RELEASE

Factor	What it measures
Quality	The editorial quality score on the record.
Relevance	75% topic overlap with the previous Discovery (or the Wander From Here starting point), plus 25% from the visitor's reaction affinities for the candidate's topics.
Novelty	One minus overlap with every topic already visited on this trail.
Variety	Zero for the same domain as the previous Discovery; otherwise one minus topic overlap with it.
Obscurity	The editorial obscurity score on the record.
Exploration	A random draw per candidate, so near-ties don't always resolve the same way.
Domain cooldown	Not weighted: the whole score is reduced 18% for each time the candidate's domain appears among the last three Discoveries shown, up to 45%.

MODE WEIGHTS IN THE V0.7.8 RELEASE

Factor	Familiar	Explore (default)	Wild
Quality	30%	30%	25%
Relevance	35%	20%	8%
Novelty	15%	20%	25%
Variety	10%	15%	20%
Obscurity	7%	10%	14%
Exploration	3%	5%	8%

Read across a row and the intent shows: Familiar gives the most weight to relevance, Wild trades it for novelty, variety and obscurity, and quality stays high in all three.

Reactions as fading topic affinities

Each reaction adds to an affinity for every topic on that Discovery: Interesting +0.35, More like this +0.8, Not for me −0.7. The contribution shrinks by a factor of 0.82 for each step the trail has moved since, and each topic's affinity is kept between −1 and 1. Starting a new trail keeps reactions and history; it changes the path, not the memory.

Wander From Here

Someone can start from any public website address. Wander validates it as a public HTTP or HTTPS URL and strips credentials, query and fragment. If the URL matches an indexed Discovery, that Discovery's topics are used. Otherwise the server derives topic hints from words in the host and path, plus a small hand-written map of related topics (for example, "architecture" also suggests art, places and history). It never fetches the submitted page during the recommendation request. It's a starting-point heuristic, not an understanding of arbitrary web content.

What the system is not

- No stored semantic edge graph, embeddings or shortest-path search at runtime.
- No model generating or summarizing destinations when someone wanders.
- No research data in ranking: study participation is consented separately and doesn't feed recommendations.

Built to be evaluated offline

The selection engine accepts optional injected random sources. Production passes none. An offline harness can supply them so two engine variants are compared on the same random draws,, which keeps luck from deciding a comparison like those in section 6.

05

Discovery-graph evaluation

I think of Wander as a landscape of possible journeys rather than a list of websites. To see how well the metadata supports that idea, I built an offline graph: each Discovery is a node, and two Discoveries are linked if they share a topic. It's a diagnostic representation, not something the product runs.

There are two versions of the graph. The **any-topic** graph counts all shared topics, including nine broad family topics (such as art, history and play) that the ingestion process adds by pattern. That's what the relevance factor sees. The **specific-topic** graph leaves the family topics out. It's the stricter, more trustworthy approximation, because family topics are added automatically, not by a reviewer.

OFFLINE SNAPSHOT OF 118 PUBLISHED DISCOVERIES (HISTORICAL, NOT THE CURRENT LIVE COUNT)

Measure	Any shared topic	Specific topic only
Edges (of 6,903 possible pairs)	3,153	239
Density	0.46	0.035
Connected components	2	39
Largest component	117 (99.2%)	51 (43.2%)
Isolated Discoveries	1	31

Put another way: 239 pairs share a specific topic, 2,914 are connected only through broad family topics, and 3,750 share no topic at all. Thirty of the 31 specifically isolated Discoveries are connected only through family labels. The broad topics hold the graph together: art alone is the only link for 902 pairs, and history for 822.

> **What this does and doesn't mean.** These counts describe how the metadata represents the corpus. They aren't measures of user value. An isolated Discovery can be excellent; a missing path can mean the real relationship just isn't represented. I didn't remove, retag or rerank anything because of this audit.

Inspecting a trail

I also built an offline inspector that explains a single simulated trail: the path between each step in both graphs, whether a move crosses components, and flags such as a move connected only through family topics. On one simulated trail, no selected move had a specific-topic path and all 11 moves changed component. That explains one simulated trail. It's not a finding about people.

Asking people

Because a tidy graph isn't evidence that a transition feels good, I ran a small registered review in which people rated constrained sets of transitions without knowing how they were produced. Specific-topic transitions were rated as more related in a limited set of cases. Other findings were mixed or unsupported, and none of it establishes that topic proximity should decide ranking. The reviewers assessed constrained transitions, not open-ended journeys.

06

Experiments and rejected alternatives

Rather than assume new weights would fix the experience, I tested alternatives against quality gates I set before running them. None entered the app.

SELECTION EXPERIMENTS AND THEIR DISPOSITIONS

Approach	What it tried	Result	Disposition

Approach	What it tried	Result	Disposition
Context-aware Explore shortlist	Make Explore's candidates depend on the current trail context.	Larger candidate pool, but below its quality comparison, and it favored sparsely tagged destinations.	Not adopted
One-over-rank draw	Weight candidates by the inverse of their rank instead of a uniform top-three pick.	Failed quality gates across modes.	Not adopted
Eligibility filter	An added filtering step on which candidates are eligible.	Passed its hard gates, but by design doesn't improve variety.	Not integrated
Later challenger (E3)	A preregistered challenger, gated on a technical check before its later phases.	Overall source-balanced quality missed the required bound; ten of twelve crossed test conditions failed.	Stopped; remaining phases not run

Why I stopped instead of tuning

Each experiment had its gate written down before it ran. When one failed, I closed it rather than adjusting it until a preferred result appeared. The lesson isn't that the current engine is validated. It's that interesting destinations don't automatically make an interesting journey. The current system still relies on a small curated corpus, topic labels that can be broad, and a top-three draw, and whether its paths feel coherent and worth following remains a human question.

07

Research interpretation

I ran two studies and kept their findings separate. The Research & Evidence companion has the methods and full aggregate tables. Here's how I read them.

- **Alpha01 (five participants) was encouraging.** All five reported finding something they probably wouldn't have found alone, and all five continued from a tenth Discovery to an eleventh. It's a small, self-selected formative sample with uneven activity, not a population estimate or evidence of retention.
- **It also showed mental models varied.** People described Trail and the distance controls differently. Nobody recorded a Save or a Wander From Here. Zero use can't tell me whether a feature is hidden, unclear or simply unnecessary.
- **Firstvisit01 tested the premise without me.** It was remote and unmoderated. Several people began without guidance, but short and uneven exploration (a median of three Discoveries) and only three questionnaires left comprehension and perceived value unresolved.

What I did with that: I didn't redesign Trail, How Far?, Save or Wander From Here from a five-person pilot, and I didn't respond to the first-visit study by adding interface instructions. I want to understand what participants meant before changing the experience.

08

Accessibility and reliability decisions

- **Memory is enforced where it can't be bypassed.** The API and database enforce no-repeat across trails and returning visits, not the browser.
- **The corpus is reviewed and versioned.** It's loaded by the application, not generated by a model, so what someone can land on is known in advance.
- **Destinations are treated as untrusted.** A separate ingestion process verifies candidate URLs before editorial approval. That doesn't guarantee every destination stays safe or available forever.
- **No third-party page loads inside Wander.** Previews are static images, so a destination's scripts never run in Wander's page.
- **Identity is optional.** People begin anonymously. Accounts preserve history across devices. Export and deletion are available under Your data. Research participation is consented separately and doesn't feed ranking.

An accessibility correction (local, not deployed)

A local accessibility review caught a failure a screenshot wouldn't show. The sign-in and Wander From Here dialogs clipped their controls on short viewports, and a recoverable account error could leave keyboard focus in the wrong place. I corrected the dialog scrolling and made focus restoration conditional on what happened. Four of the nine original reproduction probes had failed. All nine passed afterward, along with focused automated browser checks.

This correction isn't deployed yet and isn't an accessibility certification. Physical-device, assistive-technology and hosted verification remain open.

09

Implementation tradeoffs

CHOICES, WHAT THEY BOUGHT AND WHAT THEY COST

Choice	What it gives	What it costs
Progressive disclosure	A calm first visit that leads with the experience.	Some people will leave before meeting the later features.
A curated, reviewed corpus	Every destination was approved by a person; nothing is invented.	A small collection, so running out is a real state the product has to handle honestly.
Topic labels as the relationship signal	Simple, explainable scoring I can inspect factor by factor.	Broad family topics dominate, and the specific-topic graph is sparse.

Choice	What it gives	What it costs
Top-three random draw	Results aren't fully deterministic, so the same state doesn't always give the same next step.	Harder to reason about a single trail; experiments need shared random draws to compare fairly.
Server-side selection and history	No-repeat holds across trails, reloads and devices.	Every Wander is a server round trip, and anonymous identity needs protecting from abuse.
Anonymous first, optional account	No gate in front of the core experience.	History stays tied to one browser until someone chooses an account.
Static preview images	No third-party page runs inside Wander.	A preview can drift from what the live site looks like now.

10

Current limitations

- The recommendation approach has known limits, and no experiment established a better production system.
- The corpus is small, and its topic labels are often broad.
- Participant evidence is small and formative. Comprehension of the first visit and perceived value are unresolved.
- There's no evidence yet of sustained return, broad appeal, revenue or a winning How Far? design.
- The accessibility correction and further reliability work are local and not deployed. Assistive-technology and physical-device testing haven't happened.

My release standard for v1.0 isn't a perfect algorithm or a feature-complete roadmap. It's a trustworthy destination, an understandable first action, dependable memory, essential accessibility, and evidence that people find something worth their time.

11

Future product questions

The longer-term idea is whether relationships between discoveries could reach beyond websites, to creative work, music, events, places or real-world experiences. That's an exploration. It isn't a roadmap or a validated need, and none of it is built. Before extending the model, I'd need answers to questions like these:

- What counts as a meaningful connection across different kinds of things, when shared topic labels are already too thin for websites alone?
- How can people steer without turning wandering into search?

- How should place and time be represented responsibly, and how is privacy protected when discovery leaves the web?
- Closer in: does a first visit make sense without explanation, and where should How Far? sit on a small screen?

I started with a belief that the internet could still surprise us. Building Wander has made me more careful about what counts as evidence for that belief. The product is real and available. Whether it consistently earns someone's curiosity is what I'm still working to learn.

Selection details in this document come from the v0.7.8 release source. Study figures are aggregates; the companion document explains their limits.