Making goodreads greatreads

Why readers don't rely Goodreads' recommendations and how to rebuild it

Product Case study| Chandra Keerthi Pothina | Long-time Goodreads user

Goodreads tracks reading, doesn't shape it

Goodreads is the world's **largest online book communities with 125M users...**

Owned by **Amazon**, it has over **15 years** of user ratings, **data**,

and THE **default tracking platform** for readers

But readers increasingly **discover new books** through TikTok, Instagram, YouTube,

Friends & book clubs

Goodreads remains a **log** for tracking **history** and hardly shaping **future reads**

Competitors are winning on discovery and recommendations. For now.

	goodreads	The StoryGraph	Hardcover
Users	~125M+ (public estimates)	4M+ users (publicly stated)	~50K–100K community users
Core Positioning	Book tracking + social log	Data-driven, mood-based discovery	Community-built alternative to Goodreads
Recommendation Basis	Shelves + broad activity	Rating-based + mood filters	Ratings + AI-assisted discovery
Personalization Depth	Low–Moderate	High (mood, pacing, tone filters)	Moderate–High
Transparency	Limited explanation	Explains recommendation logic	Shows reasoning + community signals
AI Usage	Minimal	Limited	Active AI assisted recommendations
Social Layer	Strong legacy network	Light	Community-native

Does not leverage its biggest moat, 15+ years of structured ratings at massive scale.

Wins on structured taste matching

Wins on modern UX and transparency

Recommendations are one of the most persistent, high-impact user pain points

Recommendations fail to reflect user's preferences..

- Suggestions are driven by popularity or broad correlation
- Likes and dislikes have limited visible impact
- Even active users receive largely the same generic lists

As a result, users..

- Ignore recommendations altogether
- Rely on external sources (TikTok, Instagram, book tube, friends)
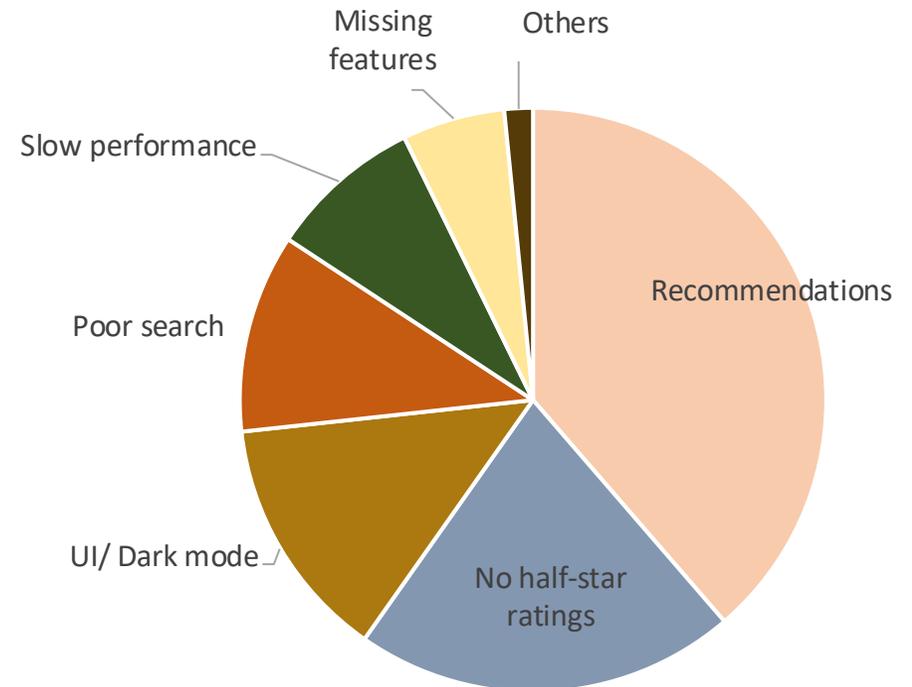- Switch to intelligent book-tracking apps

"Even after rating hundreds of books, I still get the same generic lists."

"I juts want to discover books I may actually enjoy"

"GR still suggests me pre-teen books because that's what I used to read when I first joined. Does not recognize my changed reading habits"

..based on an analysis of 300+ user complaints across

- Reddit discussion threads
- YouTube comments
- Goodreads community posts and reviews

Missing features

Others

Slow performance

Poor search

Recommendations

UI/ Dark mode

No half-star ratings

Recurring themes ranked by frequency
(Approximated based on thematic analysis of complaints)

# Fixing recommendations improves the entire Goodreads ecosystem

When recommendations fail, Goodreads shifts from a **guide for discovery** to a **passive book log**

Discovery expectations have changed

- Readers now expect recommendations to reflect taste, not popularity
- Competing platforms differentiate on personalization, not scale
- Goodreads' size is an advantage only if its data is used intelligently

If recommendations remain generic, Goodreads keeps user data but loses influence over discovery.
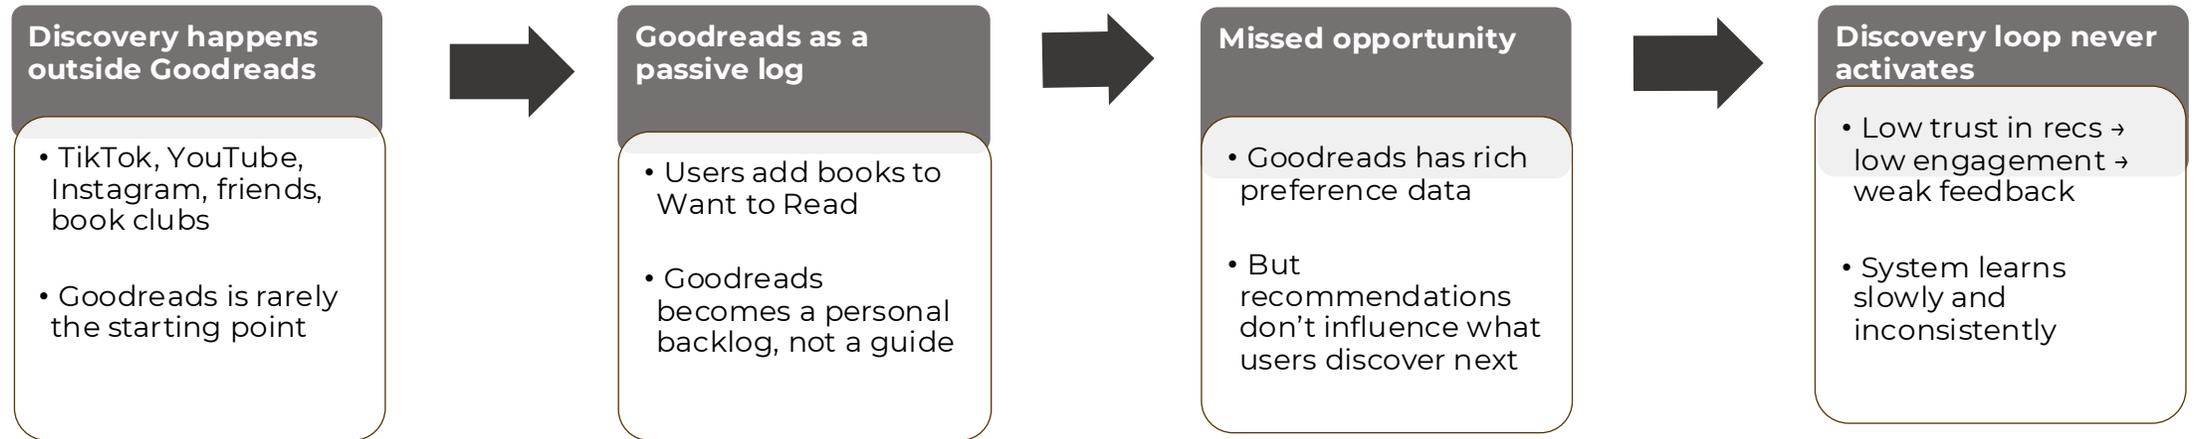
Smart recommendations unlock

Better recommendations

Better signals

Better engagement
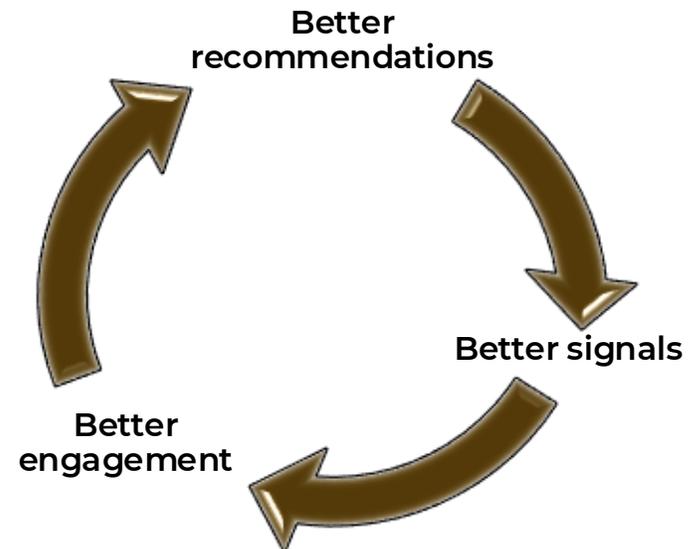
Good recommendations create better feedback, which makes recommendations even better.

Goodreads is where discovery gets stored, not where it gets shaped

How readers actually discover books today

Discovery happens outside Goodreads
- TikTok, YouTube, Instagram, friends, book clubs
- Goodreads is rarely the starting point

Goodreads as a passive log
- Users add books to Want to Read
- Goodreads becomes a personal backlog, not a guide

Missed opportunity
- Goodreads has rich preference data
- But recommendations don't influence what users discover next

Discovery loop never activates
- Low trust in recs → low engagement → weak feedback
- System learns slowly and inconsistently

How smart discovery can compound on Goodreads

Personalized Discovery
Personalized recommendations surface books aligned to reader taste and context

Active Engagement
Readers act on suggestions by reading, rating, reviewing, and shelving

Social Proof
Activity creates visibility and community-wide credibility

Commercial Impact
Discovery converts into purchases across print, e-book, and audiobook formats

System Learning
Reader actions generate feedback that improves future recommendations

Rebuilding trust in recommendations requires changing what we optimize for

Preference over activity

- **Principle:** Not all user actions signal liking
- **Implication:** We optimize for ratings (clear preference), not shelves (ambiguous signals)
- **Trade-off:** We sacrifice some recency of activity for clarity of intent

Recency over history

- **Principle:** Reader taste evolves over time
- **Implication:** Recent 5-star ratings matter more than books loved a decade ago
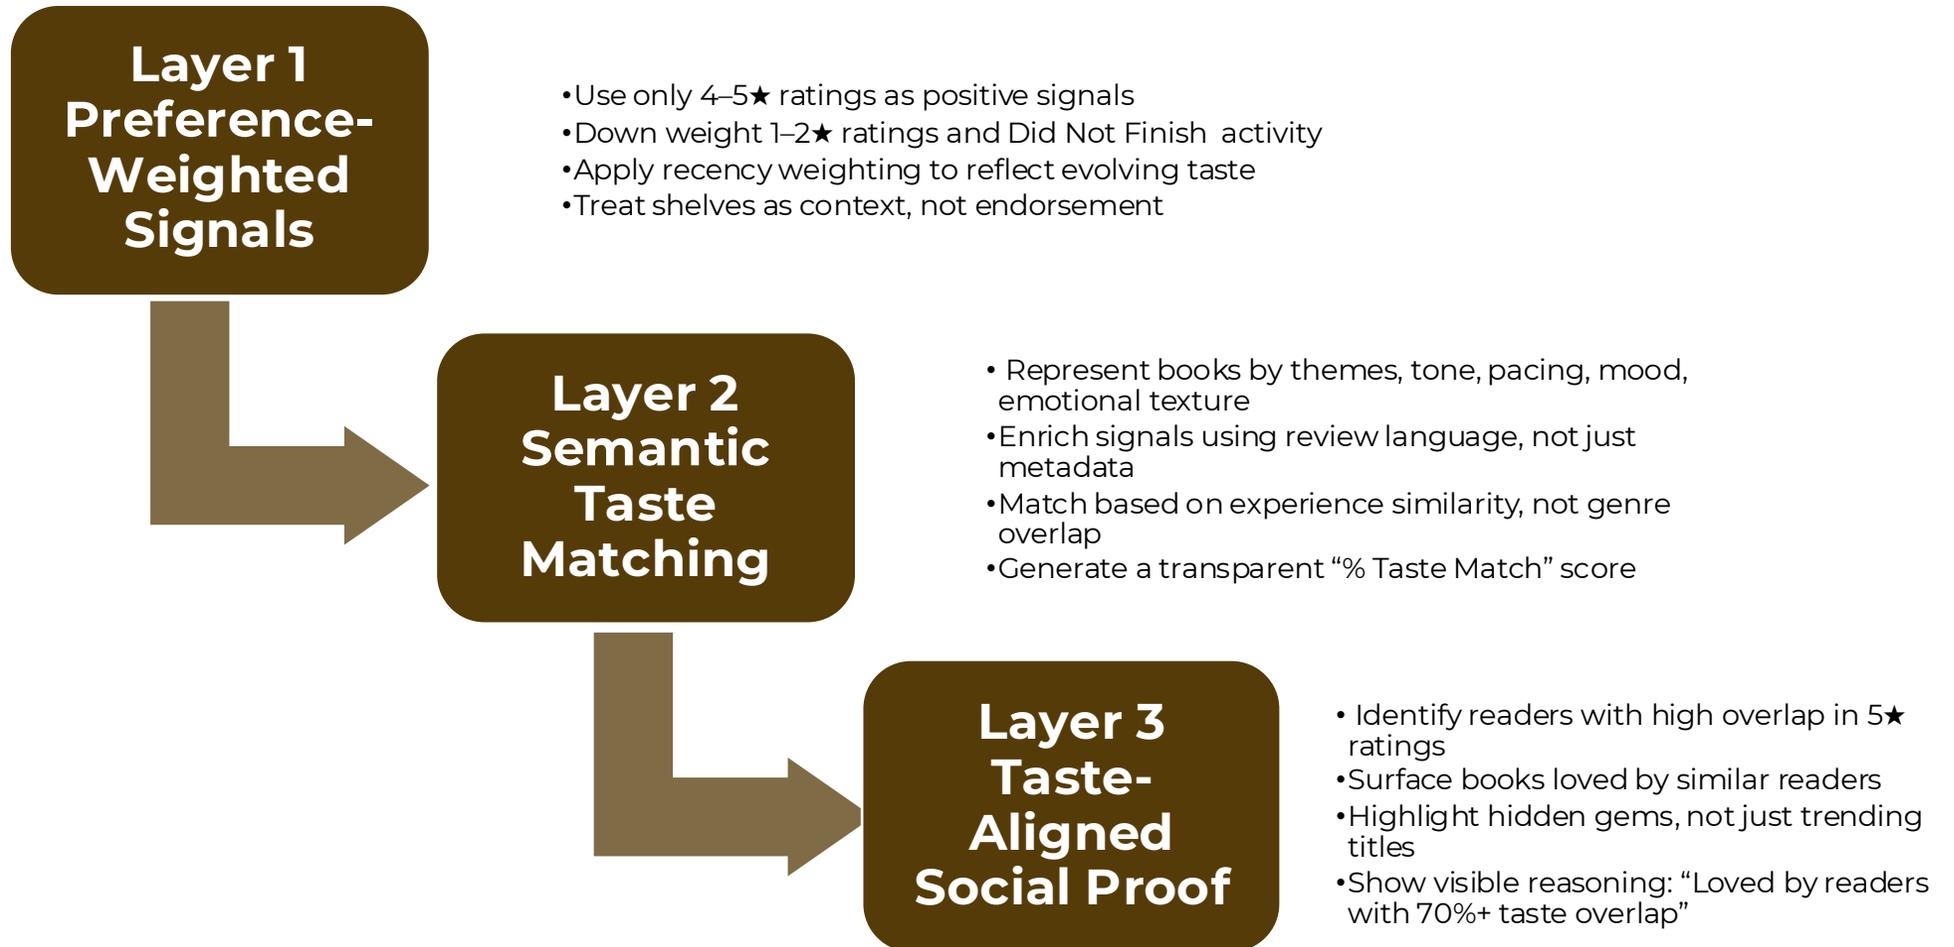- **Trade-off:** New users get less personalized recs initially (cold start is temporary)

Similarity over popularity

- **Principle:** Readers with aligned taste are the strongest signal
- **Implication:** We recommend books loved by similar readers, not just bestsellers
- **Trade-off:** May surface niche books over safe mainstream choices (this is a feature, not a bug)
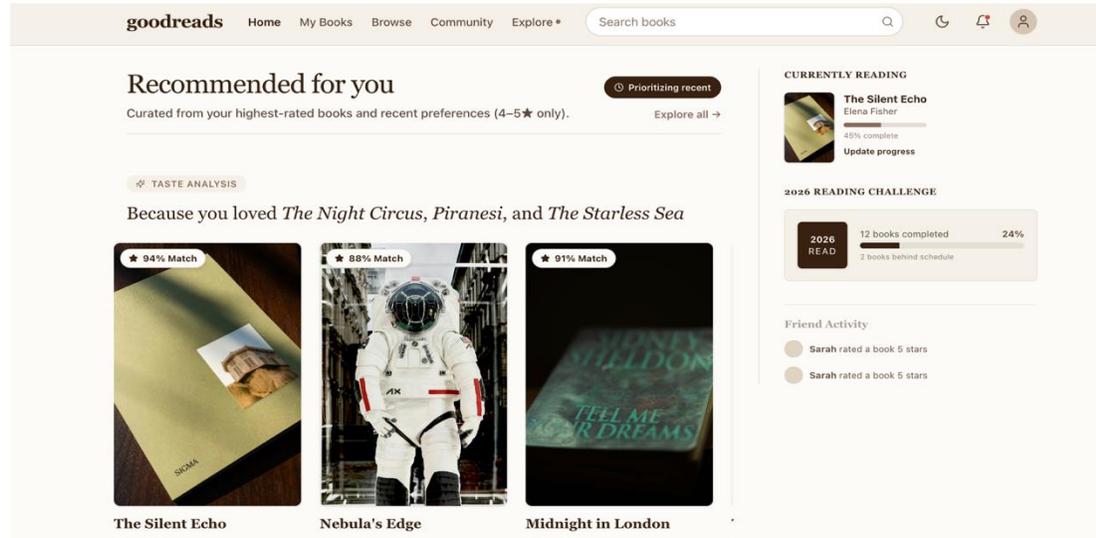
Transparency builds trust

- **Principle:** Users should understand WHY they're seeing a recommendation
- **Implication:** Always show the reasoning ("Because you loved X, Y, and Z...")
- **Trade-off:** Limits black-box ML approaches. Choosing explainability as it matters more than marginal accuracy gains
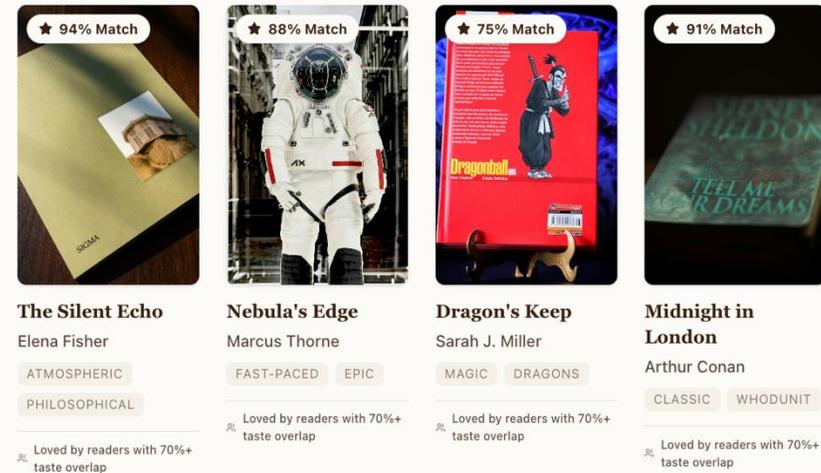
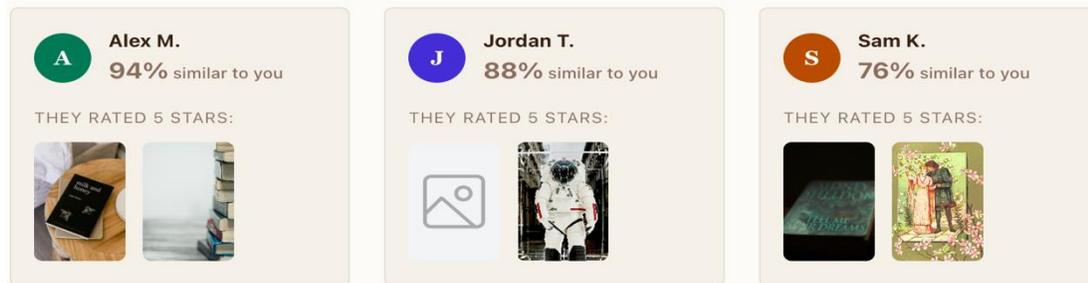# A 3 layered approach to move from activity to preference-based recommendations

Layer 1 Preference-Weighted Signals

- Use only 4–5★ ratings as positive signals
- Down weight 1–2★ ratings and Did Not Finish activity
- Apply recency weighting to reflect evolving taste
- Treat shelves as context, not endorsement

Layer 2 Semantic Taste Matching

- Represent books by themes, tone, pacing, mood, emotional texture
- Enrich signals using review language, not just metadata
- Match based on experience similarity, not genre overlap
- Generate a transparent "% Taste Match" score

Layer 3 Taste-Aligned Social Proof

- Identify readers with high overlap in 5★ ratings
- Surface books loved by similar readers
- Highlight hidden gems, not just trending titles
- Show visible reasoning: "Loved by readers with 70%+ taste overlap"

How preference-first discovery looks on Goodreads

Interactive prototype (Figma link) [click here](#)

Each recommendation is designed...

Each book shows a % match based on semantic similarity and aligned readers.

Midnight in London
Arthur Conan
CLASSIC WHODUNIT

Loved by readers with 70%+ taste overlap

Why this recommendation?
Because you rated 'Sherlock Holmes' 5 stars.

Only 4–5 ⭐ and recent ratings power these recommendations.

Continue exploring your taste ›

94% Match — **The Silent Echo** — Elena Fisher — ATMOSPHERIC PHILOSOPHICAL — Loved by readers with 70%+ taste overlap

88% Match — **Nebula's Edge** — Marcus Thorne — FAST-PACED EPIC — Loved by readers with 70%+ taste overlap

75% Match — **Dragon's Keep** — Sarah J. Miller — MAGIC DRAGONS — Loved by readers with 70%+ taste overlap

91% Match — **Midnight in London** — Arthur Conan — CLASSIC WHODUNIT — Loved by readers with 70%+ taste overlap

The Silent Echo
Elena Fisher
ATMOSPHERIC PHILOSOPHICAL

Loved by readers with 70%+ taste overlap

Why this recommendation?
Because you enjoy dystopian themes with philosophical depth.

"Why this recommendation?" expands with reasoning reflecting transparency

...for better book discovery and experience
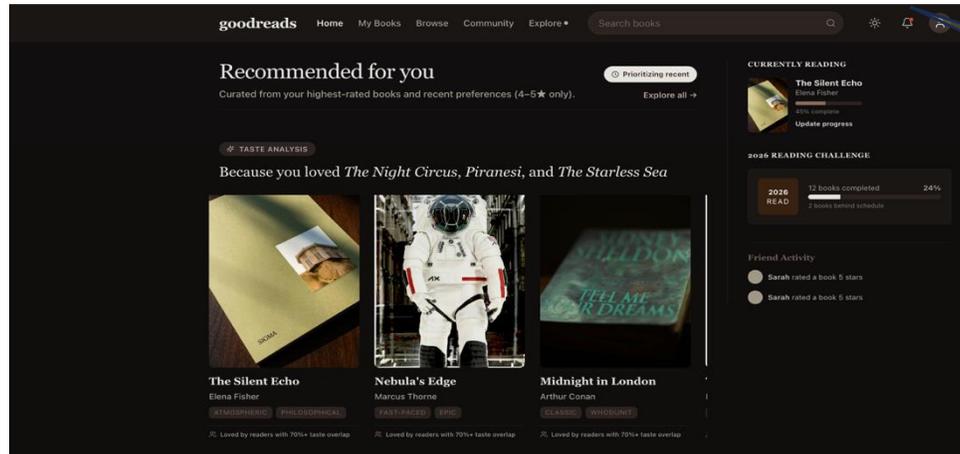
Discovery stays consistent across a much-demanded for 'Dark mode'
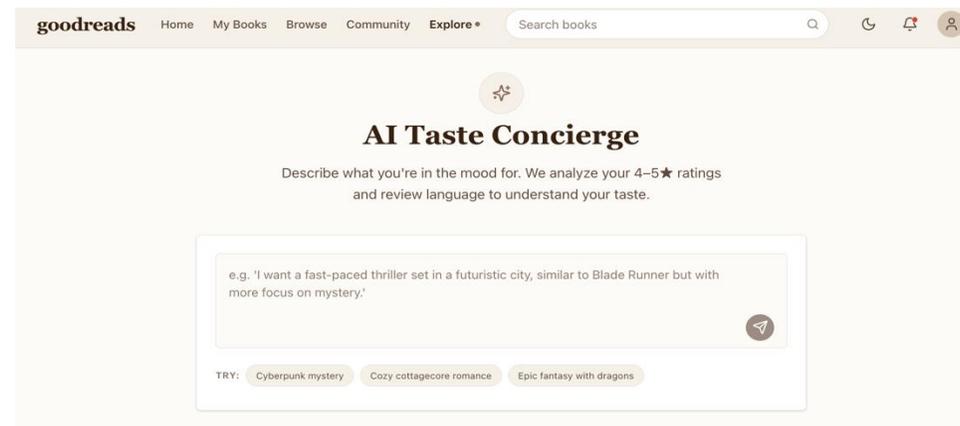
AI Concierge is an optional exploration tool. Core recommendations are driven by our preference-first matching engine, not generative AI.

High-match books prioritized even if they are not mainstream, based on similar readers' preferences

The system design powering preference-first discovery

USER INTERACTION LAYER
Desktop Interface

- 📖 Desktop Web Interface
- ☆ Book Ratings (1–5 Stars)
- 💬 Review Text Analysis
- ▽ 'Not Interested' / DNF
- ⬤ Dark Mode Toggle
- 💡 AI Concierge Input
  Optional natural language query

Only 4–5★ and recent ratings are treated as strong signals.
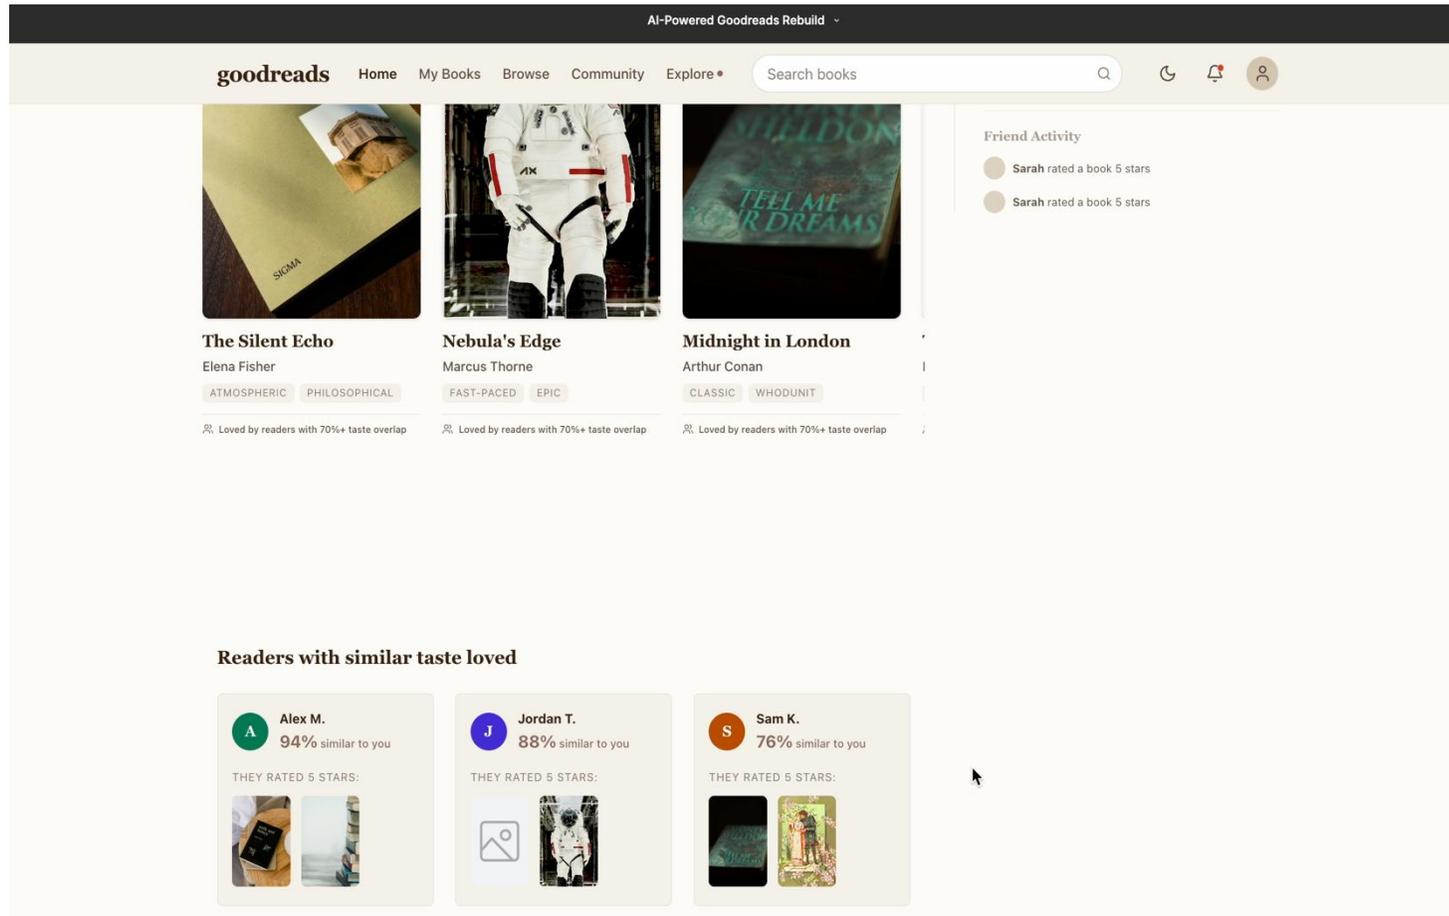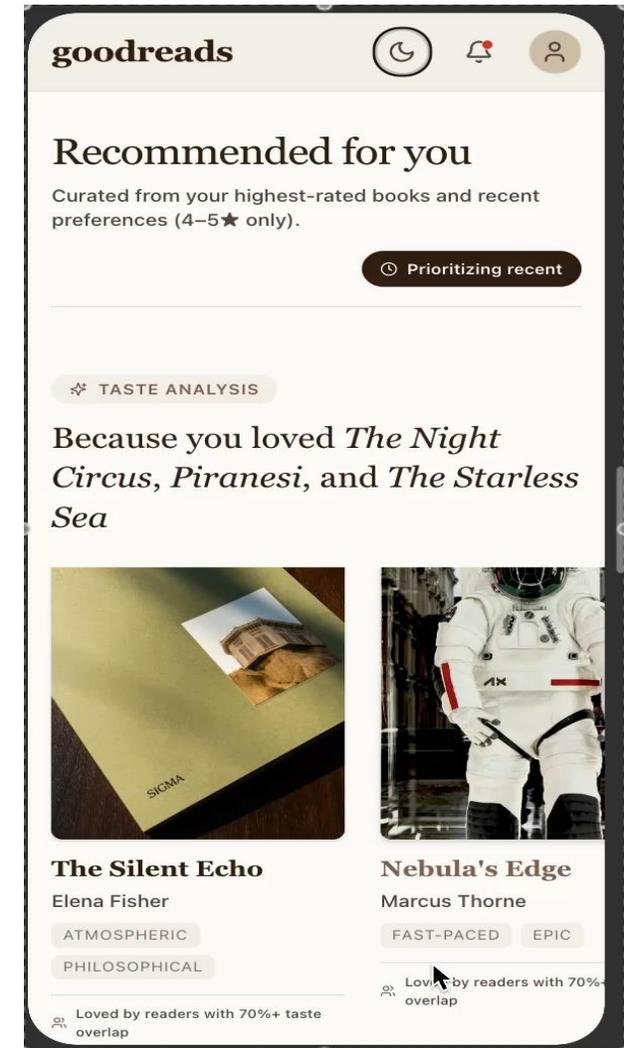

Preference-First Logic

PREFERENCE ENGINE
Signal Processing

- ▽ Filter 4–5★ Ratings
- ▽ Exclude 1–2★ + DNF
- 🗄 Recency Decay Function
  Recent ratings weighted 3x vs historical
- ☰ Ignore Shelves
  Reduces noise
- ✴ Theme Extraction
  NLP on review language

Only 4–5★ ratings, adjusted by recency, are treated as strong signals.

TASTE MATCHING ENGINE
Core Intelligence

🗄 SEMANTIC EMBEDDINGS
Books mapped by themes, tone, pacing, and emotional depth vectors.

👥 COLLABORATIVE FILTERING
Identify readers with >60% overlap in 5★ ratings history.

✴ HIDDEN GEM BOOST
Surface high-match books with low popularity scores.

Hybrid System: Semantic + Collaborative (Not Generative AI)

RANKING + TRANSPARENCY
Scoring Model

Weighted Scoring Model

% Match Calculation

```
% Match = (Semantic
Similarity × 0.5) + (Taste
Overlap × 0.4) + (Recency
Weight × 0.1)
```

Similar Reader Validation

Explanation Generator
"Because you rated 'The Secret History' 5 stars..."

Trust via Visibility

USER EXPERIENCE
Frontend Delivery

98% Match Badge on covers

Social Proof
'Loved by readers with 70% taste overlap'

Expandable Context
'Why this recommendation?'

Dark Mode Support

OPTIONAL TOOL
🔍 **AI Taste Concierge** Natural language discovery, distinct from core recs.

← CONTINUOUS LEARNING LOOP

DATA LAYER (FOUNDATION)

UserRatings	**ReviewText**	**BookEmbeddings**	**TasteGraph**
UserID, BookID, Rating, Timestamp	Theme extraction input	Vector representations	User-to-user similarity scores

Great reads recommendations in action

Desktop version
(hover to click play)

Mobile app version

Preference-first discovery delivers impact but also has trade-offs

Smart discovery has clear benefits for both users and goodreads...

For Readers: Discovery becomes intentional

- Recommendations reflect books they rated highly, not everything they touched
- Clear "Why this book?" explanations build confidence
- Less time spent filtering irrelevant suggestions
- Reduced dependence on external discovery platforms

For Goodreads: Stronger engagement and feedback loops

- Higher click-through on recommendation modules
- More 4–5★ ratings (high-confidence signals increase)
- Improved quality of collaborative similarity graph
- Greater influence over downstream book purchasing

...But comes with a few potential risks which can be mitigated

Risks including

- New users may get weak recommendations if they haven't rated enough books
- System may over-personalize and keep users in narrow taste bubbles
- Niche clusters could limit discovery diversity
- Matching and similarity scoring increases infrastructure load

Mitigations and Design Controls

- Guided onboarding to collect data on 5 strong preference signals upfront
- Inject 10–20% "exploration" titles outside core cluster
- Minimum similarity thresholds before boosting books
- Batch processing and incremental updates to manage compute cost

How to measure if preference-first discovery actually works

North Star Metric
% of Books Added From Recommendations
Currently <5% (assumption)→ Target: 25%+

L1 Metrics: Trust & Engagement Signals

Are users acting on recommendations?
- Recommendation click-through rate
- % of recommendations added to Want to Read
- % of recommended books later rated 4–5★
- Repeat visits to Explore section
- "Not Interested" / hide actions ↓

L2 Metrics: Business & Ecosystem Impact

Is discovery creating business value?
- Amazon click-out conversion rate
- Hidden gem discovery rate (<10K ratings books)
- Retention of active readers
- Cross-surface engagement (Explore → My Books → Reviews)
- Time spent in discovery flows

Guardrail Metrics: Monitoring for risk

- Diversity score – Number of unique genres in recommendations
- Over-personalization rate – Recommendations from 2 genres or less
- Cold start quality -% of new users who add ≥1 recommendation in first week

The future of Goodreads is not social. It's intelligent.

Preference-first discovery makes Goodreads

Personal.
Trusted.
Indispensable.