Needwise explained: conversational recommendation without text generation

25 November 2026 · 7 min read · blogs.iictp.org/needwise-conversational-recommendation-typed-decisions-explained

Short answer: Needwise is a conversational recommender, described in an arXiv preprint submitted on 23 November 2026, that ranks items from typed decisions (named questions answered with option probabilities) instead of asking a language model to re-read the whole dialogue each turn. The authors report it ranks Google Local requests better than GPT-6 Sol (0.383 against 0.343 nDCG@10) at a fraction of the cost and in 2.5 seconds at the median against 5.6. It matters for search and local marketing because it shows how needs such as price, place and type of item can be tracked explicitly. This is an arXiv preprint. It has not been peer reviewed, and its numbers are the authors' own claims.

TL;DR The short version

- Needwise is a conversational recommender that tracks needs as typed decisions with option probabilities instead of having an LLM re-read the whole dialogue.
- The authors report results comparable to the strongest LLM ranker on Google Local and Yelp, with auditable reasons.
- It recognises only 19 to 27% of withdrawn needs; for local marketers, the text you publish is the evidence such systems read.

Read online, listen to it, or change the text size: https://blogs.iictp.org/needwise-conversational-recommendation-typed-decisions-explained

What is Needwise?

The abstract frames a problem: users add, withdraw and reverse needs across a conversation, while catalogs record only some attributes. LLM rankers re-read the whole dialogue each turn, which is slow, and their ordering does not show which needs an item meets. Needwise keeps a typed state of needs instead.

How it works

- **Typed requests.** Named questions over fixed options, answered with option probabilities rather than generated text.
- **Per-turn mapping.** One request maps the user's message onto the catalog's fields, keeps what no field expresses in the user's own words, and decides how the turn changes each earlier need.
- **Must-first ranking.** Needs on the kind of item, its place and its price form a tier; inside it a holistic decision and need-level verdicts (supported, contradicted or unknown) are combined with equal weight.

What is new compared with LLM rankers

An LLM ranker returns an order without an auditable reason. Needwise's verdicts per need let you see why an item ranked where it did, and the typed state is meant to survive a user switching requests and coming back.

Results reported by the source

- With graded attributes withheld, so needs must be inferred from names, descriptions and reviews, Needwise-HF ranks Google Local requests significantly better than GPT-6 Sol: 0.383 against 0.343 nDCG@10.
- It is reported as not significantly different from the strongest LLM ranker on Google Local and Yelp.
- Cost is reported as a fraction of GPT-6 Sol's API price, with a median of 2.5 s against 5.6 s on Google Local.
- In conversations, the typed state keeps the ranking across switches and returns, but recognises only 19 to 27% of withdrawn needs, the authors say.

Limitations and open questions

- The weak point the authors report is withdrawn needs: 19 to 27% recognised.
- It is a preprint, and the benchmarks are Google Local and Yelp, so results may not carry to other catalogs.
- The paper's abstract does not say how it handles reviews written in different languages. Sources were read on 25 November 2026. Products, rules and standards change, so check the linked originals before you rely on any detail.

How to apply the idea

1. List the needs your customers state in a request: type, place, price, and soft preferences such as "quiet".
2. Check which of them your own product or listing data actually records.
3. Fill gaps in descriptions, attributes and reviews, since the system infers needs from that text.
4. Log each need as supported, contradicted or unknown for your top items and see where you are unknown.

What it means for local and search marketers

The idea reinforces something practical: when a system infers whether a business meets a need from names, descriptions and reviews, the text you publish is the evidence. Clear, specific descriptions and review responses make "supported" verdicts more likely. For Google's own AI search behaviour, see query fan-out in Google Search and Ask Maps.

Frequently asked questions

- **What is Needwise?** A conversational recommender built on typed decisions rather than free-text generation.
- **What does it report on Google Local?** 0.383 nDCG@10 against 0.343 for GPT-6 Sol, in a median of 2.5 s against 5.6 s.
- **What is its main weakness?** It recognises only 19 to 27% of withdrawn needs.
- **Is it peer reviewed?** No, it is a preprint.

Related reading

- Ask Maps: agentic appointments and local visibility
- Query fan-out in Google Search
- AI Practitioner certification

Related credential: AI Practitioner. Explore the exam facts and syllabus.

Sources and further reading

1. Kim S, Huang W, Reddy MRS, Srinivasan A. Needwise: Generation-Free Conversational Recommendation from Need-Level Typed Decisions. arXiv 2611.31175, submitted 23 Nov 2026 (preprint). arxiv.org/abs/2611.31175

Published by IICTP, an independent technical credential body, on 25 November 2026. The current version of this post is at https://blogs.iictp.org/needwise-conversational-recommendation-typed-decisions-explained. More guides: blogs.iictp.org.