NP LAB

Beyond the Search Box: Designing Intelligent Event Discovery

A conceptual case study examining natural-language search, travel-aware recommendations, and responsible AI in event discovery

Concept status: Research exploration. North Parade LLC has not built, tested, or deployed the proposed system. No event provider, ticketing platform, or nonprofit has commissioned or endorsed this work. Examples are fictional and no live ticket availability is represented.

North Parade LLC | October 2026

The problem is not simply a shortage of events

An event platform can display hundreds of listings and still leave a user with few practical choices. Traditional discovery systems tend to prioritize category, date, geographic radius, and inventory availability. These fields are useful, but they do not fully represent how people decide whether to attend an event.

A nearby event may require a ferry, bridge crossing, or difficult drive. A free ticket may still involve parking, tolls, childcare, or several hours of travel. A weekday performance can be technically available but practically inaccessible. A user may travel farther for a favorite team than for an unfamiliar concert.

The resulting problem is not necessarily poor inventory. It is a mismatch between the information stored about events and the constraints that determine whether those events are worthwhile to a particular person.

The proposed concept

This study considers an optional intelligent discovery layer for an established event platform. It would interpret ordinary-language requests, translate them into explicit constraints, retrieve real event records, and rank eligible results according to user preferences and practical travel conditions.

The system would not distribute tickets, invent events, override eligibility rules, or independently promise availability. Existing ticketing and allocation services would remain authoritative.

- Natural-language search alongside ordinary filters.
- Explicit exclusions for unwanted event categories or locations.
- Travel-time estimates based on actual routes rather than distance alone.
- Different travel tolerances for different interests and days.
- Saved, editable preferences and opt-in recommendations.
- Plain-language explanations of why each result was included.

A fictional use scenario

Alex is a fictional resident of coastal Connecticut. Alex enjoys hockey, selected concerts, and occasional baseball games. Weeknight travel is difficult. New York City can be practical by train, while some geographically closer locations on Long Island are inconvenient.

Alex enters: “Find hockey games or concerts within about 90 minutes. Weekends are better. I would travel farther for the Rangers. No soccer, and avoid Long Island.”

The application identifies the preferred categories, explicit exclusions, normal travel limit, weekend preference, and exception for a favorite team. It asks a clarification only if the ambiguity would materially change results, such as whether train travel should be considered.

The result is not an open-ended AI-generated event list. It is a ranked set of actual event records returned by the platform, with eligibility and inventory status verified by the underlying system.

Why traditional filters fall short

A category checkbox expresses a broad interest but cannot distinguish mild interest from a special exception. A radius filter treats roads, bridges, water, traffic, and transit access as though they were irrelevant. An event date filter often fails to represent work schedules or a preference for weekends.

These limitations are not evidence that every search system needs artificial intelligence. Many can be fixed through better interface design, accurate categorization, improved geospatial queries, and explicit user settings.

AI becomes useful when a user describes several interacting preferences in ordinary language and expects the system to translate them into a consistent, inspectable search.

Preferences must remain enforceable

The proposed system would distinguish positive interests, neutral categories, and explicit exclusions. If a user excludes soccer, a language model should not quietly reintroduce soccer because it predicts that the event is popular. A deterministic rule must apply the exclusion after any model-generated interpretation.

Users would be able to see and edit the interpreted search criteria. A model misunderstanding should be correctable without restarting the conversation or navigating a long list of checkboxes.

Proposed technical design

The language model would operate as a constrained intent parser, not as the event database or ticketing authority. It would produce a validated structure containing categories, excluded regions, time windows, transport preferences, maximum travel times, and stated exceptions.

Proposed intelligent discovery architecture

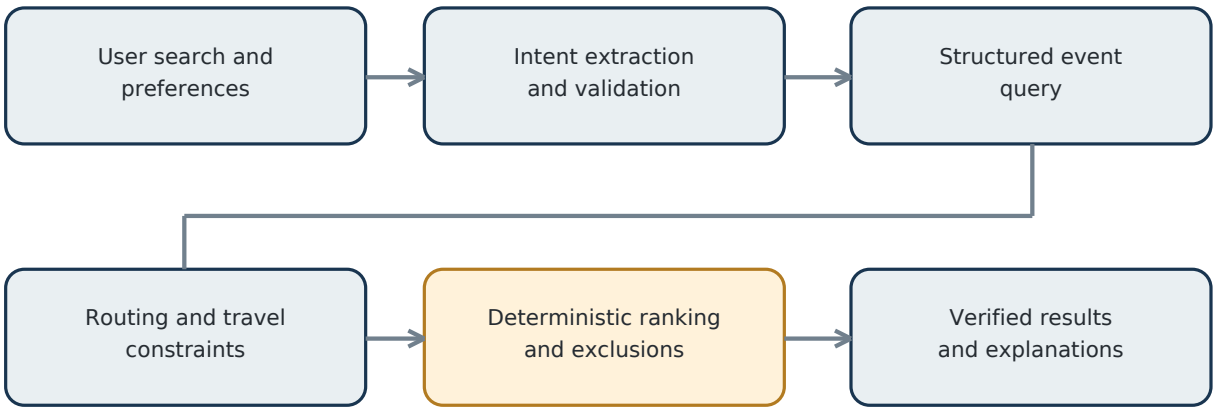

Figure 1. Proposed intelligent discovery architecture

The application would query authoritative event inventory, resolve venue locations, estimate practical routes, and apply deterministic filters. Ranking would consider relevance, timing, travel burden, and user-defined flexibility. Only verified records would be returned.

Model output would be validated against an explicit schema. Unknown categories, unsupported venue claims, fabricated ticket counts, and unauthorized actions would be rejected. A failure of the AI service should leave conventional event search fully usable.

Travel time is contextual, not absolute

Travel calculations depend on mode, departure time, traffic, rail schedules, bridge access, and the reliability of third-party routing data. An estimated 75-minute drive on a Sunday morning may be much longer on a weekday evening.

The platform should display estimates and assumptions, not guarantees. Users could choose driving, transit, or both, and specify whether tolls, ferries, or certain regions are unacceptable. More expensive routing calculations could be limited to promising candidate events.

Recommendations without unnecessary surveillance

A useful system needs some memory of preferences, but it does not need an unrestricted personal profile. Interests, exclusions, travel tolerances, and notification settings should be explicit, editable, and removable.

Sensitive identity-verification records should remain separate from discovery preferences. A recommendation engine should not receive identity documents, government identifiers, or eligibility evidence merely to rank events.

The system should avoid inferring sensitive traits from attendance behavior. It should also provide meaningful discovery without requiring users to disclose personal histories or enable continuous location tracking.

Accessibility and practical usability

Natural-language search should be an option rather than a requirement. Clear category filters, readable text, large touch targets, keyboard support, screen-reader labels, and straightforward error messages remain essential.

An older or less technically confident user should be able to search without understanding prompts or AI. The system should display its interpretation in plain language, allow one-tap corrections, and avoid blocking the application with unnecessary announcements or acknowledgments.

Search relevance and respectful account support are different concerns. An AI search feature cannot compensate for inaccessible verification workflows, poor customer support, or inconsistent interface behavior.

A responsible first demonstration

A prototype should begin with a synthetic event dataset representing concerts, hockey, baseball, theater, and other categories across a realistic region. It should include inconvenient travel corridors, weekday events, closed inventory, and deliberately ambiguous event metadata.

The prototype could compare three approaches: ordinary category-and-radius filters, improved deterministic filters, and natural-language intent extraction layered on the same underlying search.

- Can users find relevant events in fewer steps?
- Are explicit exclusions honored every time?
- Does the system distinguish distance from practical travel time?
- How often does the model misunderstand a preference?
- Does the explanation accurately reflect the ranking decision?
- What are latency, routing expense, and model operating costs?
- Can users complete the task without AI when the model is unavailable?

A demonstration should not claim improved ticket availability or attendance without measured evidence. Success would mean better discovery of existing inventory, not creation of additional donations.

Operational risks and boundaries

The principal risks are fabricated event details, misleading availability, inconsistent interpretation of exclusions, privacy leakage, poor routing estimates, algorithmic bias toward popular events, and excessive complexity for a task that simpler filters could solve.

Hard constraints should be enforced outside the model. Eligibility, allocation, and ticket delivery must remain in existing trusted services. User-facing explanations should cite the actual event data and disclose uncertainty where appropriate.

The concept should be reconsidered if testing shows that conventional filtering performs just as well, model interpretation repeatedly violates user intent, or third-party routing and inference costs outweigh the benefit.

Potential applications beyond event ticketing

The underlying approach is not limited to entertainment. Similar preference-aware discovery could help people locate community programs, public resources, service appointments, museum activities, volunteer opportunities, or other location-dependent services.

In each case, the same boundary applies: language models can help interpret human requests, while authoritative systems must supply the facts and enforce the rules.

Conclusion

Better event discovery is not primarily a question of making an interface more conversational. It is a question of accurately representing the constraints people use when making decisions.

Some improvements require no AI at all: clearer navigation, respected exclusions, readable mobile layouts, and better regional filters. A controlled language-model layer may add value by interpreting nuanced requests and making the results easier to refine.

The proposed design deliberately separates understanding from authority. AI interprets the request. Structured services retrieve real events. Deterministic rules enforce constraints. The user remains in control.

That distinction makes the concept reasonable to explore without implying that it has been implemented or proven effective. The objective is not to add AI to a search box. It is to help people find events they can realistically attend.

Research and implementation references

The following public technical resources are relevant to a future implementation. They are references for design exploration, not evidence that the proposed system has been tested.

- PostgreSQL Documentation. Full-text search and geographic query support through extensions. https://www.postgresql.org/docs/
- PostGIS Documentation. Spatial indexing, geometry, and geographic queries. https://postgis.net/documentation/
- OpenStreetMap. Community-maintained geographic data and associated usage policies. https://www.openstreetmap.org/
- OSRM Project. Open-source routing engine and route computation. https://project-osrm.org/
- W3C Web Content Accessibility Guidelines (WCAG) 2.2. https://www.w3.org/TR/WCAG22/
- NIST AI Risk Management Framework. Governance and evaluation of AI systems. https://www.nist.gov/itl/ai-risk-management-framework