Insights from Designing Agentic AI for an Improved Natural Language Processing in Civic Innovation Processes

Damir Safin[1,*,†], Dian Balta[1,†]

[1]fortiss GmbH, Research Institute of the Free State of Bavaria for software-intensive systems, Munich, Germany

Abstract

In Germany, a multitude of social projects are initiated each year by individuals and nonprofit organizations. The realization of social projects almost invariably depends on funding. Project initiators, however, often face challenges when searching for suitable funding opportunities. This difficulty stems from the sheer volume of available funding sources, diverse eligibility criteria, and the inherent complexity of matching unique project specifics with varying funder objectives. These multifaceted challenges are often not addressed by German funding search portals successfully. We present a two-stage architectural approach, designed to simplify and improve the search process. In the first stage, the system employs rule-based filtering and NLP techniques for scalable coarse matching of potential funding sources. Subsequently, a second stage utilizes an LLM for fine-grained reranking, based on an assessment of contextual relationships between the project and funding programs. The proposed approach aims to provide users with a prioritized list of relevant funding recommendations. Ultimately, this approach is expected to streamline the funding search and improve the likelihood of successful matches.

Keywords

Project Funding, Citizen-Driven Innovation, Recommender Systems, Conversational Search, Agentic AI, NLP

1. Introduction

While civic innovation processes involving diverse stakeholders [1] can yield sustainable, community-focused results, they are often resource-intensive. A critical bottleneck is securing financial support, as the funding search presents significant hurdles: project initiators navigate numerous sources with varied, extensive eligibility criteria, complicating the alignment of project needs with funder objectives. Consequently, existing German funding search portals [2] and traditional matching methods [3] often struggle with these complexities, particularly in capturing domain-specific subtleties.

To address these shortcomings, this paper introduces an agentic AI system design. This system employs a two-stage architecture aimed at making the funding search more streamlined and impactful, thereby enhancing civic innovation efficiency. The proposed system leverages a combination of techniques: NLP, including the application of LLMs; agent planning rules; and the execution of external tools, such as accessing funding databases. While applying such advanced AI techniques presents known complexities [4], our architecture is designed to harness their strengths for this assessment.

2. System Architecture and Matching Process

The system is designed to match project descriptions—which users provide via natural language queries—to suitable funding opportunities. To achieve this goal effectively, the architecture employs a two-stage matching process intended to balance computational efficiency with a detailed analysis of the relationships between the project and the numerous funding sources within the funding database. The hybrid nature of the architecture stems from its combination of symbolic rule-based filtering with NLP techniques. The decision to implement a hybrid, two-stage architecture is driven by several factors.

Proceedings EGOV-CeDEM-ePart conference, August 31 - September 4, 2025, University for Continuing Education, Krems, Austria

*Corresponding author.

†These authors contributed equally.

safin@fortiss.org (D. Safin); balta@fortiss.org (D. Balta)

0009-0002-7674-9678 (D. Safin); 0000-0001-8311-3227 (D. Balta)

© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073

published 2025-12-12

CEUR-WS.org/Vol-4127/paper51.pdf

A purely rule-based system might lack the flexibility to capture the nuances of project descriptions and funding criteria, while relying solely on neural models to process the entire database could be computationally prohibitive. The initial filtering stage is designed for efficiency, quickly reducing the candidate pool from a potentially vast database. This stage allows the subsequent, more computationally intensive LLM-based reranking to focus its resources on a smaller, more relevant set of funding programs, thereby achieving a practical balance between comprehensive coverage, processing speed, and the depth of contextual analysis.

The matching process begins with a first stage that rapidly filters the extensive funding database to identify an initial set of promising candidate funding programs. This is achieved by concurrently applying rule-based constraints, which leverage structured data like location or funding limits, and by utilizing NLP techniques to assess semantic relevance. Specifically, the NLP techniques employed include keyword-based search methods and cosine similarity of embeddings; these are used to evaluate the degree of relevance between the user's query and the descriptions of the funding programs. This initial filtering is designed to efficiently reduce the number of potential funding programs for subsequent, more detailed analysis.

The second stage performs a deeper evaluation exclusively on the promising candidates identified in the previous step. This stage utilizes an LLM for contextual reranking. In this role, the LLM evaluates the alignment between the specific project details and the funder's objectives, considering aspects such as thematic fit, intended impact, and implicit criteria. The outcome of the matching process is a final, prioritized list of funding recommendations, with the LLM-reranked candidates at the top. Funding sources that were filtered out during the first stage but might still hold some relevance are listed separately at the end of the list. This structured presentation of results is expected to streamline the user's review process and improve the likelihood of finding a successful funding match.

Acknowledgments

This work was partially supported by financial and other means by the following research projects: DUCA (EU grant agreement 101086308), DiProLeA (German Federal Ministry of Education and Research, grant 02J19B120 ff), ROBIN (Grant no: KON-23-039) at the Bavarian Research Institute for Digital Transformation, as well as our industrial partners in the FinComp project.

Declaration on Generative AI

The authors have not employed any Generative AI tools in the preparation of this work.

References

[1] T. Sengewald, A. Roth, Actors' Roles in Citizen-driven Innovation, in: ISPIM Connects Osaka – Connecting and Empowering Society, 2024. URL: https://conferencesubmissions.com/ispim/osaka2024/documents/1415856622_Paper.pdf.

[2] Bundesverband Deutscher Stiftungen e.V., Stiftungssuche, https://stiftungssuche.de/, 2025. Online directory of German foundations by the Association of German Foundations. Accessed: 2025-04-24.

[3] Y. Zhang, S. Qiao, J. Zhang, T.-H. Lin, C. Gao, Y. Li, A survey of large language model empowered agents for recommendation and search: Towards next-generation information retrieval (2025). arXiv:2503.05659.

[4] D. Balta, M. Sellami, P. Kuhn, H. Krcmar, Insights from natural language processing, in: Proceedings of ongoing research, practitioners, posters, workshops, and projects of the international conference egov-cedem-epart, volume 2019, Shefali Virkar, Olivier Glassey, Marijn Janssen, Peter Parycek, Andrea Polini, Barbara Re, Peter Reichstädter, Hans Jochen Scholl, Efthimios Tambouris, 2019, p. 241–243. URL: https://biblio.ugent.be/publication/8626904/file/8626906.pdf.