Steve: Your Personal AI Career Coach

Balaji Rao, Naveen Mathews Renji, Elena Korshakova, Carlo Lipizzi
{brao, n3, ekorshak, clipizzi}@stevens.edu

[1]Stevens Institute of Technology
1 Castle Point Terrace, Hoboken
NJ 07030 USA

Abstract

Steve is an AI career coaching platform that turns a resume and insights from an AI-enabled chat with a user into a personalized skill gap report and upskilling roadmap. The platform suggests a personalized course plan, supports continuous learning, and helps shape the user's career trajectory. Steve is built around schema-constrained JSON artifacts and a configurable career-tree ontology. The system compares confirmed skills against role-specific requirements, prioritizes gaps (critical/important/beneficial), and translates the analysis into embedding-based queries over the course index. Steve has three personas (Interview Coach, Resume Evaluator, and Career Coach) that provide concise feedback tailored to the user's goals and context. Steve also supports speech Input/Output (I/O) via Whisper-based speech-to-text and a dual-voice text-to-speech layer, enabling users to talk to Steve. The platform offers flexible adaptability across institutions, enabling them to configure deployments by substituting their own ontologies and course catalogs. Our demo uses STEM trajectories as a case study, but the pipeline is domain-agnostic by design. Users can edit inputs, check speech recognition accuracy, and observe consistent updates, illustrating a reproducible, human-in-the-loop pattern for deploying LLMs in career guidance. Steve is currently in its alpha stage and available for demonstration at: https://calm-pond-070924d0f.4.azurestaticapps.net/

Introduction

Career progression requires a combination of consistent efforts, continuous skill-building, and an understanding of rapidly evolving industry demands. Career guidance typically relies on scarce and costly human coaching, as well as on static data that fails to reflect a candidate's evolving context. Our work addresses this gap by introducing Steve, a platform that combines AI-driven skill assessment with an AI-enabled chat to deliver detailed, actionable recommendations. It provides guidance on which certifications or courses to pursue for career progression.

Contributions

- **Schema-first orchestration:** Every stage (resume parsing, AI-enabled chat with a user, gap analysis, course query) emits compact JSON for auditing and re-runs.
- **Ontology-guided analysis:** Gap detection is grounded in a configurable career-tree (requirements), avoiding invented criteria.
- **Retrieval bridge:** Analysis becomes a structured query (domains/keywords/must-have) used for embedding-based course retrieval in Qdrant database.
- **Demo-ready interaction:** Dual-mode personas (concise user interface output + JSON), optional speech I/O, and fast refresh under live edits.

Steve is designed for both technical educational institutions, whose students are predominantly in finance, technology, and other STEM-related fields, as well as large companies seeking to help employees progress in their careers by acquiring specialized technical skill sets.

Related Work

Prior efforts combine NLP for resume parsing, conversational agents for guidance, and recommendations for learning pathways. NLP-based parsing reliably extracts structured signals from unstructured resumes (Sinha, Amir Khusru Akhtar, and Kumar 2021). Chatbots support career planning and feedback (Lee et al. 2019; Shilaskar et al. 2024); and recent LLM-driven approaches assess skill gaps directly from profiles (Kothari et al. 2024). Career recommendation has leveraged hierarchical job/skill structures and ontologies (Xu et al. 2018; Chen 2009) as well as transformer-based models to align candidates with role requirements (Guan et al. 2024). Commercial systems emphasize internal mobility, document generation, or habit coaching, but often rely on static signals and opaque pipelines. In contrast, *Steve* centers (i) schema-constrained JSON artifacts for every stage, (ii) an explicit, configurable career-tree that grounds comparisons to role-specific requirements, and (iii) an end-to-end, reproducible workflow that allows users to edit inputs, inspect artifacts, and observe consistent updates.

Methodology

Steve converts a PDF resume and a brief intake session into an ontology-grounded analysis that maps the user to roles in a career tree, prioritizes skill gaps, and assembles a course plan.

Steve's workflow:

Resume parsing Users upload their resume in a PDF format. The text is then extracted, tokenized into manageable chunks, and normalized. Function-calling is used to produce outputs that conform to a predefined JSON schema; Named Entity Recognition within the GPT model extracts structured fields (contact, education, experience, skills). Duplicate fragments are merged to preserve unique entries.

Career path analysis Steve evaluates parsed experience, education, and skills to propose short- and mid-term career titles using function-calling. With the additional information provided from an AI-enabled chat with a user, a second function call re-evaluates interests and competencies and suggests career titles for each step. Both functions embed the career-tree logic and emphasize recent skills to provide accurate recommendations across various domains.

Skill assessment This module consolidates the resume and the insights from an AI-enabled chat with a user into a skill profile and compares it against predefined role requirements to flag gaps, prioritized as critical/important/beneficial with short justifications. The result is a structured JSON report summarizing strengths and targeted gaps.

Course retrieval The gap report is converted to a structured query (domains, keywords, must-have and good-to-have qualifications). Courses (title, description, link, outcomes) are embedded and stored in Qdrant database; we use cosine similarity to return top-k matches aligned with the identified gaps (Qdrant 2023; Thakur et al. 2021).

Platform Overview

A public web interface allows users to set up a profile (using Single Sign-On-based onboarding), upload a resume, and engage in a brief AI-enabled chat with a user. Steve extracts technical and soft skills from the resume and project descriptions, maps profiles to a configurable career tree via semantic similarity, and proposes next-possible roles. Through a conversational workflow, users receive their skill report and recommended courses, and can update responses or upload a revised resume to refresh the analysis. Users can also interact with the chatbot to get further clarification on their questions.

Preliminary Evaluation

During our initial evaluations, skill extraction and immediate next-role suggestions were stable and accurate for most resumes; however, second-jump trajectories were sometimes over-optimistic relative to demonstrated skills, as reflected in course scores. We observed consistent outputs under repeated runs, suggesting robustness to conversational variability.

Conclusion

We present *Steve*, a schema-first, ontology-guided career coach platform that turns resumes and AI-enabled chats with users into skill gap reports and course plans. The modular pipeline (resume parsing, role mapping, gap analysis, retrieval) yields audit-able JSON artifacts and supports controlled re-runs.

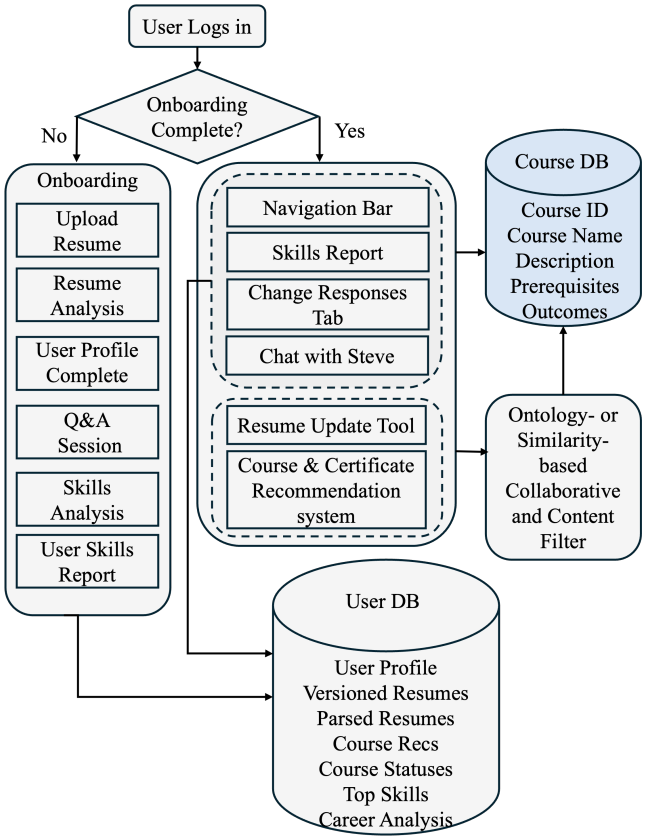

Figure 1: AI Powered Career Coach

The platform is available for demonstration purposes, and we welcome suggestions for improvements.

Ethical Considerations and Limitations

AI-based career assessment can reflect societal biases if not grounded. Steve mitigates this by anchoring comparisons in role-specific requirements and incorporating insights from AI-enabled chats with users. The prototype relies on Subject Matter Expert validation of the career hierarchy and role skill requirements. Our future work will focus on ensuring proper data protection and access control, as resumes contain sensitive personal information. We also aim to incorporate objective skill evaluations for more accurate assessments and provide regular updates on the skill requirements for each role. Additionally, we plan to conduct thorough user testing and evaluations by career experts to guide our next steps.

Acknowledgments

We gratefully acknowledge Accenture's financial support for advancing career-coaching platforms powered by Generative AI, large language models (LLMs), and machine learning.

References

Chen, C.-M. 2009. Ontology-based concept map for planning a personalised learning path. *British Journal of Educational Technology*, 40(6): 1028–1058.

Guan, Z.; Yang, J.-Q.; Yang, Y.; Zhu, H.; Li, W.; and Xiong, H. 2024. JobFormer: Skill-Aware Job Recommendation with Semantic-Enhanced Transformer. *ACM Transactions on Knowledge Discovery from Data*.

Kothari, P.; Mehta, P.; Patil, S.; and Hole, V. 2024. InterviewEase: AI-powered interview assistance. *Research Square*. Preprint.

Lee, T.; Jagannath, K.; Aggarwal, N.; Sridar, R.; Wilde, S.; Hill, T.; and Chen, Y. 2019. Intelligent Career Advisers in Your Pocket? A Need Assessment Study of Chatbots for Student Career Advising. In *Americas Conference on Information Systems (AMCIS) 2019*, AMCIS '19. ISBN 978-0-9966831-8-0.

Qdrant. 2023. High-Performance Vector Search at Scale. https://qdrant.tech/.

Shilaskar, S.; Bhatlawande, S.; Sawle, P.; Gupta, S.; and Buche, R. 2024. Conversational AI for Career Counseling. In *2024 MIT Art, Design and Technology School of Computing International Conference (MITADTSoCiCon)*, 1–5. IEEE.

Sinha, A. K.; Amir Khusru Akhtar, M.; and Kumar, A. 2021. Resume screening using natural language processing and machine learning: A systematic review. *Machine Learning and Information Processing: Proceedings of ICMLIP 2020*, 207–214.

Thakur, N.; Reimers, N.; Daxenberger, J.; and Gurevych, I. 2021. Augmented SBERT: Data Augmentation Method for Improving Bi-Encoders for Pairwise Sentence Scoring Tasks. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 296–310. Online: Association for Computational Linguistics.

Xu, H.; Yu, Z.; Guo, B.; Teng, M.; and Xiong, H. 2018. Extracting Job Title Hierarchy from Career Trajectories: A Bayesian Perspective. In *IJCAI*, 3599–3605.