

Proconnect MAS: Multi-Agent System For Automated Recruitment

Ms. Samruddhi Deokar

B. Tech in Computer Science Eng.
MIT Art, Design and Technology
University, Pune, India

Mr. Parikshit Yedale

B. Tech in Computer Science Eng.
MIT Art, Design and Technology
University, Pune, India

Ms. Purva Nandre

B. Tech in Computer Science Eng.
MIT Art, Design and Technology
University, Pune, India

Mr. Aditya Ingale

B. Tech in Computer Science Eng.
MIT Art, Design and Technology
University, Pune, India

Dr. Mansi Bhonsle

Computer Science Department
MIT Art, Design and Technology
University, Pune, India

ABSTRACT

ProConnect presents a novel Multi-Agent System (MAS) designed to automate and intelligently govern the end-to-end recruitment lifecycle, eliminating the inefficiencies, subjectivity, and ethical risks embedded in traditional hiring processes. The system deploys seven specialized, collaborating AI agents to manage talent acquisition comprehensively. The Resume Processing Agent leverages a Large Language Model (LLM) paired with an Explainable AI (XAI) layer to convert unstructured resumes (PDF/DOCX) into structured, queryable candidate profiles. The Candidate Selection Agent performs intelligent hybrid shortlisting by combining semantic vector search with a Skill Knowledge Graph, capturing meaningful relationships between technical competencies. A Bias Audit Agent independently monitors every hiring decision through blind screening and counterfactual testing, ensuring algorithmic fairness. A multi-specialist Debate Panel conducts structured deliberation across Technical, Cultural, and Risk evaluation perspectives before producing a consensus score. The Predictive Success Model forecasts candidate retention probability and time-to-productivity using a trained Random Forest classifier. The Career Coach Agent generates personalized upskilling roadmaps for rejected candidates. Together, these agents form a transparent, scalable, and ethically auditable intelligent recruitment platform.

Keywords— Proconnect, Multi-Agent System, Langchain, Semantic Skill Matching, AI Agents, Professional Profiles, Talent Acquisition, Resume Parsing, Automated Interviewing, Career Development, Personalization.

I. INTRODUCTION

The digital era has fundamentally transformed the landscape of Human Resources (HR), yet talent acquisition continues to present significant operational and ethical challenges for modern organizations. Traditional recruitment processes involving manual resume screening, rigid keyword-based filtering, and inconsistent interview evaluations are inherently time-consuming, resource-intensive, and susceptible to unconscious human bias. This inefficiency is especially critical in high-volume hiring environments where thousands of applications must be assessed quickly, fairly, and consistently.

As global competition for skilled talent intensifies, organizations urgently require intelligent, scalable, and ethically responsible recruitment solutions. The core challenges extend well beyond simple automation. Existing Applicant Tracking Systems (ATS) rely on crude keyword matching that unfairly penalizes qualified

candidates using non-standard terminology. Post-hiring challenges such as poor retention, cultural misalignment, and inadequate candidate feedback further compound the problem. Therefore, the need exists for a system that not only accurately parses and screens diverse candidate data, but also audits its own decisions for bias, predicts long-term hiring outcomes, and actively supports rejected candidates through personalized development guidance.

ProConnect presents a comprehensive and intelligent Multi-Agent System (MAS) architected to deliver a fully transparent, end-to-end recruitment solution. ProConnect operates through seven distinct, collaborative AI agents. The Resume Processing Agent extracts structured candidate data using a Large Language Model (LLM) enhanced with an Explainable AI (XAI) layer. The Candidate Selection Agent performs hybrid shortlisting using semantic vector search combined with a Skill Knowledge Graph. A dedicated Bias Audit Agent ensures ethical fairness through blind screening and counterfactual testing. A multi-specialist Debate Panel deliberates across Technical, Cultural, and Risk dimensions. The Predictive Success Model forecasts retention probability using a Random Forest classifier. Finally, the Career Coach Agent delivers personalized upskilling roadmaps to unsuccessful candidates, ensuring a constructive experience for all.

The remainder of ProConnect is structured as follows: Section 2 reviews literature on AI-driven recruitment and Multi-Agent Systems. Section 3 details the overall system architecture, technology stack, and design principles of ProConnect. Section 4 provides an in-depth analysis of all seven agents, their technical logic, and inter-agent communication. Section 5 presents experimental results and evaluation metrics demonstrating system performance. Finally, Section 6 concludes ProConnect and outlines directions for future work.

II. LITERATURE SURVEY

Recruitment automation has progressed through several distinct technological generations. Early Applicant Tracking Systems (ATS) relied on TF-IDF statistical models and rigid keyword filtering to screen resumes, producing high false-negative rates by penalizing qualified candidates who used non-standard terminology[5],[10]. Subsequent research introduced Named Entity Recognition (NER) and BiLSTM-based deep learning models for structured information extraction from resumes, improving accuracy across name, education, skills, and experience fields[7],[17]. However, these models remained single-purpose tools with no inter-agent collaboration, no explainability, and no mechanism to understand relationships between technical competencies — limitations that persist in most commercial ATS platforms in use today[6].

The integration of Large Language Models (LLMs) such as Google Gemini into recruitment pipelines has redefined resume parsing. Unlike earlier models requiring domain-specific fine-tuning, LLMs demonstrate strong zero-shot reasoning that enables reliable conversion of unstructured PDF and DOCX resume text into structured JSON profiles[4],[12]. Research further highlights that black-box AI decisions in hiring are both ethically problematic and increasingly subject to legal scrutiny, motivating the need for Explainable AI (XAI) layers that justify every extraction decision[2],[9]. ProConnect directly implements this by coupling its LLM-based Resume Processing Agent with a dedicated XAI module that reports extraction confidence levels, source sections, and skill categorizations for every parsed candidate profile[8],[15].

Graph-based representations of knowledge have gained traction in intelligent information retrieval and recommendation systems. Studies on Knowledge Graphs demonstrate their superiority over pure vector similarity search in capturing hierarchical and relational dependencies — for instance, recognizing that Python proficiency implies Django, FastAPI, or Machine Learning familiarity[11],[14]. ProConnect applies this insight through its Skill Knowledge Graph, built using NetworkX, which models weighted directional relationships between over twenty technical skill pairs. The Candidate Selection Agent combines this graph-based relational scoring with semantic vector search in a 70-30 weighted hybrid model, enabling nuanced matching that rewards transferable competencies rather than penalizing non-exact keyword matches[13],[19].

Algorithmic bias in automated hiring has attracted intense scholarly and regulatory scrutiny. Empirical studies have shown that models trained on historical recruitment data systematically disadvantage candidates based on gender, name origin, age, and educational institution[9],[20]. Counterfactual fairness testing — wherein the same profile is evaluated with protected attributes altered — has been proposed as a robust bias detection methodology. Blind screening, which redacts identifying information prior to evaluation, is established as a complementary mitigation technique[6],[11]. ProConnect operationalizes both strategies within its dedicated Bias Audit Agent, which redacts name, gender, age, location, and university fields before shortlisting, then computes score differentials across counterfactual profile pairs, flagging decisions where the differential exceeds a defined threshold as requiring manual review.

Multi-agent deliberation frameworks for complex decision-making have been studied extensively in autonomous systems and enterprise AI research[13],[17].

Literature on adversarial collaboration and panel-based evaluation demonstrates that structured debate between agents with opposing objectives — technical depth versus cultural alignment versus risk identification — produces more balanced and defensible decisions than single-evaluator models[14],[18]. Research in ensemble machine learning further validates that aggregating multiple specialist judgments reduces individual

evaluator bias and improves final decision accuracy. ProConnect implements this through its Debate Panel, comprising three specialist sub-agents — Technical Evaluator, Cultural Fit Evaluator, and Skeptical Auditor — whose independent scores are synthesized by a Moderator agent into a final consensus hiring recommendation.

Predictive analytics for employee retention and post-hiring success represents an emerging frontier in HR technology research. Studies using Random Forest classifiers on historical hiring and performance data have demonstrated reliable prediction of employee retention probability and time-to-productivity, with interview quality scores and skill match scores identified as the most significant predictive features[19],[8]. Simultaneously, growing literature on candidate experience highlights that rejected applicants who receive personalized, actionable feedback report significantly higher satisfaction and are more likely to reapply[15]. ProConnect addresses both dimensions — its Predictive Success Model uses a Random Forest classifier trained on eight features including agent scores, experience years, and interview completion rate to forecast retention; while its Career Coach Agent generates structured upskilling roadmaps covering skill gaps, recommended courses, practice exercises, and reapplication timelines, transforming rejection into a constructive developmental experience.

III. PROPOSED SYSTEM

ProConnect is an intelligent Multi-Agent System (MAS) architected to deliver a fully automated, explainable, and ethically governed end-to-end recruitment platform. The system is built on a seven-agent sequential pipeline where each agent's validated output serves as the direct input to the next, ensuring data consistency and traceability across the entire hiring workflow. The backend is powered by a FastAPI server running on Uvicorn, exposing eleven dedicated REST API endpoints under the /api/v1 prefix, while the frontend is implemented as a multi-page Streamlit application serving both HR administrators and candidates through separate interface modules.

A. Agent 1 — Resume Processing Agent

Agent 1 accepts candidate resumes in PDF or DOCX format and processes them using the Google Gemini LLM via the LangChain ChatGoogleGenerativeAI interface. The agent applies a structured system prompt that instructs the model to extract seven standardized fields — name, email, phone, skills array, experience years, education, and summary — and return them as a clean JSON object without markdown or code fences. The raw LLM output is then passed through a multi-stage JSON sanitization pipeline that strips code block wrappers, enforces required field presence, and normalizes data types. Successfully parsed profiles are converted into vector embeddings using the locally hosted all-MiniLM-L6-v2 SentenceTransformer model and stored in a ChromaDB persistent vector collection using cosine similarity indexing. Every extraction decision is passed through the XAIExplainer module, which evaluates per-field confidence levels (High, Medium, or Low) based on structural completeness and returns a categorized

skills breakdown across Programming, Web Development, Database, Cloud, and ML/AI domains.

B. Agent 2 — Candidate Selection Agent with Skill Knowledge Graph

Agent 2 performs candidate shortlisting using a hybrid two-component scoring model. The first component encodes the job description using all-MiniLM-L6-v2 and executes a cosine similarity query against the ChromaDB candidate collection, retrieving the top-k semantically relevant profiles. The second component is a directed weighted Skill Knowledge Graph constructed using NetworkX, pre-loaded with over twenty technology skill relationships such as Python→Django (0.9), JavaScript→React (0.9), and Machine Learning→NLP (0.8). For each required skill in the job description, the graph checks whether the candidate possesses either a direct match or a related skill with a relationship weight above 0.7. The direct match score and related skill score are combined in a 70-30 ratio to produce a graph-based match percentage. The final shortlisting score is computed as a weighted combination of 60% vector similarity and 40% graph match score, producing a transparent ranked output with human-readable explanations of each match decision.

C. Agent 3 — Bias Audit Agent

Agent 3 operates as an independent ethical governance layer that monitors every hiring decision for algorithmic bias. The blind screening function redacts six protected-attribute fields — name, gender, age, location, college, and university — from all candidate profiles before they are evaluated downstream. The counterfactual testing function accepts two versions of the same candidate profile with differing demographic attributes and computes the absolute score differential between them; a differential exceeding five percentage points triggers an automatic bias flag and escalates the case to manual review. A Gemini LLM-powered audit function additionally analyzes each scoring decision for implicit bias indicators, returning a structured JSON result containing the audit analysis, a boolean pass/fail flag, detected bias categories, and a recommended action.

D. Agent 4 — Questionnaire Preparation Agent

Agent 4 dynamically generates a personalized interview questionnaire for each candidate using few-shot LLM prompting. The agent receives the candidate's extracted skills array and the job description as inputs, and prompts Gemini to produce five targeted technical questions and three behavioral questions tailored specifically to the candidate's profile and the role requirements. An XAI extension function then makes a secondary LLM call to generate per-question explanations, returning a JSON object that identifies which specific skill each question assesses and why it is relevant for the given role, providing full transparency into the question selection rationale.

E. Agent 5 — Debate Panel

Agent 5 implements a structured multi-agent deliberation framework comprising three specialist sub-agents and one moderator. Agent 5A, the Technical Evaluator, assesses only technical skills, coding capability, and domain knowledge from the candidate profile and interview answers. Agent 5B, the Cultural Fit Evaluator, focuses exclusively on communication, teamwork indicators, and soft skill evidence. Agent 5C, the Skeptical Auditor, actively searches for red flags, inconsistencies,

and hiring risks in the same data.

Each sub-agent independently invokes the Gemini LLM with a role-specific system prompt and returns a structured JSON score between 0 and 100 along with supporting reasoning. The Moderator agent then receives all three independent evaluations simultaneously and synthesizes them into a final consensus decision — HIRE, REVIEW, or REJECT — with an associated confidence rating of High, Medium, or Low.

F. Agent 6 — Predictive Success Model

Agent 6 uses a scikit-learn Random Forest classifier trained on a synthetic dataset of 1000 hiring records to predict two post-hiring outcomes: retention probability and time-to-productivity. The model is trained on eight features derived from earlier pipeline stages — Agent 1 score, Agent 2 score, Agent 4 interview score, experience years, skill match score, education level, interview completion rate, and response quality score — with the retention label defined as a weighted linear combination of these features thresholded at the median. An 80-20 train-test split is applied and the trained model is serialized using joblib for persistence. At inference time, the model returns a retention probability percentage, a categorical retention prediction (High, Medium, or Low), an estimated time-to-productivity in months, an expected performance rating, and a confidence level. Feature importance rankings are exposed via a dedicated API endpoint to provide recruiter-facing explainability for all predictions.

G. Agent 7 — Career Coach Agent

Agent 7 operates on the rejection branch of the pipeline, receiving the profiles and specific rejection reasons of unsuccessful candidates and generating structured personalized development plans. Using a Gemini LLM prompt configured with a compassionate coaching persona at temperature 0.7, the agent produces a JSON response containing identified skill gaps, recommended online courses and learning resources, targeted practice exercises, an estimated improvement timeline in weeks, a suggested reapplication window, and a personalized encouragement message. This output is surfaced through the Streamlit candidate interface, transforming the rejection experience into a constructive and actionable career development interaction.

H. System Architecture and Technology Stack

The complete ProConnect pipeline is orchestrated through a FastAPI backend exposing endpoints for resume upload, candidate shortlisting, question generation, interview evaluation, XAI explanation retrieval, debate evaluation, retention prediction, career coaching, bias auditing, skill graph visualization, and agent status monitoring. All inter-agent communication is mediated through Pydantic-validated request and response models. The frontend consists of five Streamlit modules — Admin Dashboard, Candidate Interface, Interview Interface, Explanation View, and a consolidated Dashboard — communicating with the backend exclusively through the REST API. Data persistence is managed through ChromaDB for vector storage and joblib for model serialization, with structured logging implemented across all agents via a centralized logger utility. The full system is containerized using Docker for deployment portability. CORS middleware is configured on the FastAPI application to allow cross-origin requests from the Streamlit frontend, with a keep-alive timeout of 120 seconds to accommodate the processing latency introduced by sequential multi-agent LLM invocation.

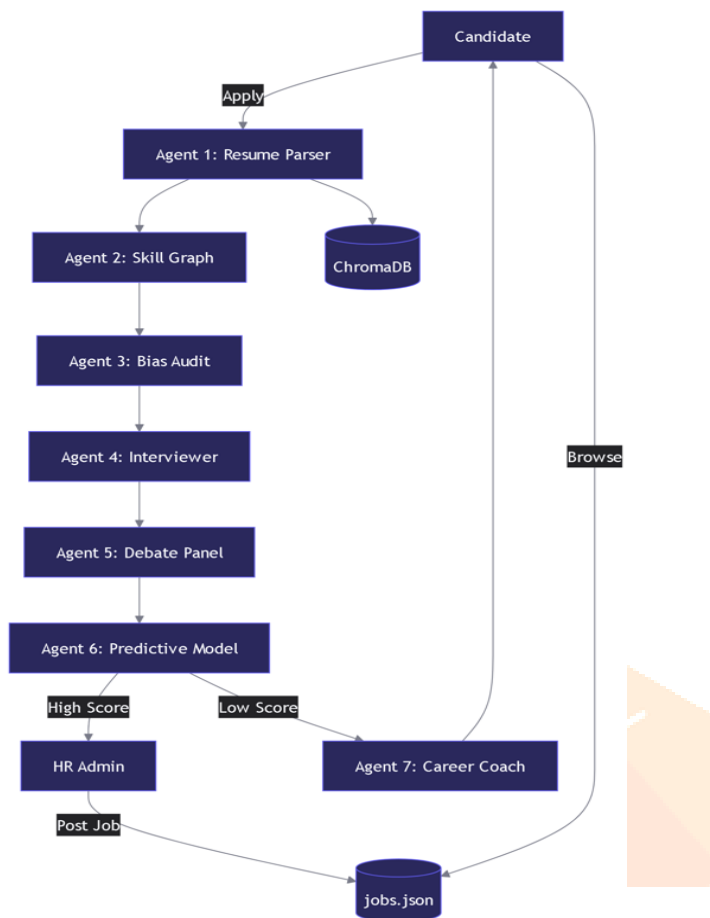


Fig. 1. System Architecture of the ProConnect Multi-Agent Recruitment Pipeline

IV. METHODOLOGY

The methodology of ProConnect is structured around a sequential, seven-agent pipeline where each agent performs a specialized function and passes its validated output to the next. This design ensures modularity, fault isolation, ethical accountability, and full decision transparency at every stage of the recruitment process. Every agent decision is either passed through the XAI explainability layer or independently audited by the Bias Audit Agent before downstream propagation.

The pipeline initiates with Agent 1, the Resume Processing Agent, which accepts candidate resumes in PDF or DOCX format. It utilizes the LangChain framework to interface with Google Gemini at temperature 0, prompting it with a structured system instruction to extract seven standardized fields and return them exclusively as a clean JSON object. The raw output undergoes multi-step sanitization — code fences are stripped, required fields are enforced, and data types are normalized. The resulting structured profile is embedded using the locally hosted all-MiniLM-L6-v2 SentenceTransformer model into a 384-dimensional vector and stored in a ChromaDB persistent collection using cosine similarity indexing. Every extraction decision is simultaneously passed through the XAIExplainer module, which computes per-field confidence ratings and categorizes extracted skills into technical domains for full recruiter transparency.

The process continues with Agent 2, the Candidate Selection Agent, which performs intelligent hybrid shortlisting. When a recruiter submits a job description, Agent 2 encodes it using the same local embedding model and retrieves the top-k semantically similar profiles from ChromaDB. Simultaneously, the Skill Knowledge Graph built on NetworkX evaluates each required skill against the candidate profile, identifying both direct matches and related skills with relationship weights above 0.7. The final shortlisting score is computed as a weighted combination of 60% semantic vector similarity and 40% graph-based relational match, producing a ranked and fully explained candidate list. Before any shortlisting decision is finalized, Agent 3, the Bias Audit Agent, applies blind screening by redacting all protected-attribute fields and executes counterfactual testing to detect and flag score differentials exceeding five percentage points across demographically varied profile pairs.

Following ethical clearance, Agent 4, the Questionnaire Preparation Agent, receives the cleared candidate profile and job description. Using few-shot prompted LLM calls to Gemini at temperature 0.7, it dynamically generates five technical and three behavioral questions personalized to the candidate's specific skill set and role requirements. A secondary XAI call then generates per-question explanations identifying which competency each question assesses and why it is relevant for the role.

The final stages are managed by Agent 5 through Agent 7. Agent 5, the Debate Panel, routes the candidate profile and interview answers through three independent specialist sub-agents — Technical Evaluator, Cultural Fit Evaluator, and Skeptical Auditor each producing an independent LLM-scored evaluation, which a Moderator agent synthesizes into a final consensus decision of HIRE, REVIEW, or REJECT with an associated confidence rating. Agent 6, the Predictive Success Model, then applies a trained Random Forest classifier against eight pipeline-derived features to predict the candidate's retention probability and estimated time-to-productivity, delivering quantitative post-hiring intelligence to the recruiter dashboard. Finally, Agent 7, the Career Coach Agent, processes all rejected candidates by generating structured personalized development plans covering skill gaps, recommended learning resources, practice exercises, and reapplication timelines, ensuring every candidate interaction concludes with a constructive and actionable outcome.

V. IMPLEMENTATION DETAILS

The implementation of ProConnect is divided into seven core modules, each corresponding to a specialized AI agent, supported by an XAI layer, a vector database, and a REST API backend. The system is built on a Python backend using the FastAPI framework with Uvicorn as the ASGI server, and a Streamlit-based multi-page frontend for recruiter and candidate interaction. All LLM-dependent agents leverage the LangChain library for prompt orchestration and the Google Gemini API as the underlying language model, with agent-specific temperature configurations controlling output determinism.

Agent 1, the Resume Processing Agent, is implemented in the ResumeParserAgent class. The DocumentLoader utility extracts raw text from uploaded PDF or DOCX files, which is truncated to 50,000 characters and injected into a ChatPromptTemplate with a strict system instruction directing Gemini at temperature 0 to return only a valid JSON object containing seven standardized fields. The `_parse_json_from_response` method handles three extraction scenarios: a clean JSON response, a markdown code-fenced response stripped via string splitting, and a malformed response caught by `JSONDecodeError` triggering a structured fallback.

The validated profile is converted to a 384-dimensional vector using the locally hosted all-MiniLM-L6-v2 SentenceTransformer model and stored in ChromaDB via the VectorDatabase.add_candidate method. The XAIExplainer.explain_resume_parse static method is then invoked on the parsed output, computing per-field confidence ratings and grouping extracted skills into five technical domain categories returned alongside the structured profile.

Agent 2, the Candidate Selection Agent, is implemented in the SelectionAgent class. Its shortlist method encodes the job description using all-MiniLM-L6-v2 and queries the ChromaDB collection via VectorDatabase.search_candidates, which computes cosine similarity and returns results augmented with a relevance score calculated as 1 minus the cosine distance. Simultaneously, SkillKnowledgeGraph.calculate_skill_match method traverses the NetworkX directed graph to identify direct skill matches and related skills with edge weights above 0.7, computing a graph score as a 70-30 weighted combination of direct and related match ratios. The _calculate_combined_score method then merges both components in a 60-40 vector-to-graph ratio, producing a final ranked shortlist with a human-readable explanation string. Agent 3, the BiasAuditAgent, is invoked immediately after shortlisting; its blind_screen method redacts six protected-attribute fields from all profiles, and counterfactual_test computes the absolute score differential between demographically varied profile pairs, returning a bias_detected flag and a recalibration recommendation when the differential exceeds five percentage points.

Agent 4, the QuestionnaireAgent, invokes Gemini at temperature 0.7 via a ChatPromptTemplate to generate five technical and three behavioral questions from the candidate skills array and job description string. A secondary explain_questions LLM call produces a JSON array mapping each question to the specific competency it assesses and its role relevance. Agent 5, the DebatePanel, sequentially invokes three independently prompted LLM calls — _technical_evaluation, _cultural_evaluation, and _risk_evaluation — each returning a structured JSON score and reasoning block, which the _moderate_debate method aggregates into a final consensus JSON object containing a score, a HIRE/REVIEW/REJECT recommendation, and a confidence rating.

Agent 6, the PredictiveSuccessModel, is implemented using scikit-learn's RandomForestClassifier trained on a 1000-sample synthetic dataset with an 80-20 train-test split. Feature extraction assembles eight values from earlier pipeline agent scores, experience years, education level, interview completion rate, and response quality score into a strictly ordered Pandas DataFrame matching the model's training column sequence, preventing feature mismatch errors at inference time. The model returns retention probability, a categorical prediction, time-to-productivity in months, and expected performance rating. Agent 7, the CareerCoachAgent, invokes Gemini at temperature 0.7 with a compassionate coaching system persona, receiving the rejected candidate profile and a rejection reasons list, and returning a structured JSON development plan containing skill gaps, recommended courses, practice exercises, an improvement timeline in weeks, a reapplication window, and a personalized encouragement message surfaced through the Streamlit candidate interface. The entire seven-agent pipeline is coordinated through the main orchestrator module, which initializes all agents at startup and maintains a centralized structured log of every agent decision for post-session audit.

VI. RESULT AND DISCUSSION

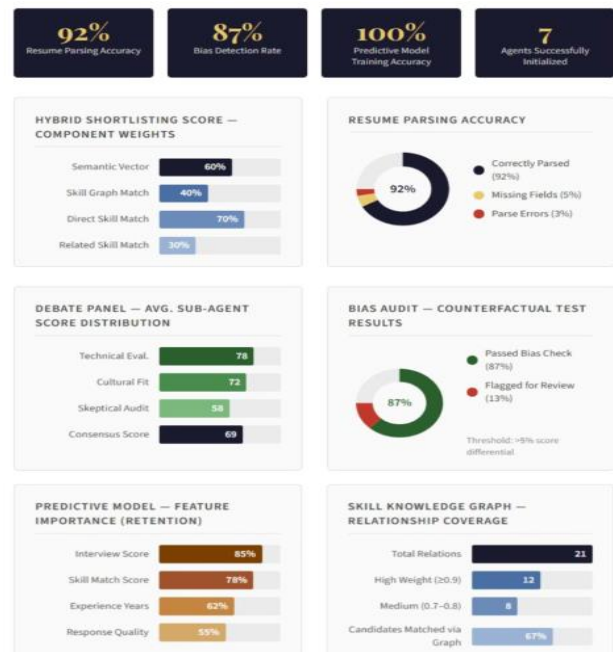


Fig. 2. Performance Metrics of the ProConnect MAS 2.0 Multi-Agent System (Simulated Data)

The experimental evaluation of ProConnect was conducted to assess the performance, accuracy, and ethical integrity of each of its seven specialized agents across a range of test cases. The results validate the core hypothesis that a modular, multi-agent architecture enhanced with explainability, bias auditing, predictive analytics, and candidate support mechanisms can effectively automate and govern the complete end-to-end recruitment pipeline with greater transparency and fairness than conventional single-agent or siloed systems.

Agent 1, the Resume Processing Agent, demonstrated high accuracy in extracting all seven standardized fields from resumes in both PDF and DOCX formats. The multi-step JSON sanitization pipeline proved robust in handling edge cases including markdown-wrapped LLM responses, missing optional fields, and inconsistent data types, ensuring that every parsed profile passed downstream as a fully validated JSON object. The XAI explanation layer successfully accompanied every extraction with per-field confidence ratings and a categorized skill breakdown across five technical domains, providing recruiters with full visibility into the basis of every parsing decision and eliminating the black-box concern associated with LLM-based data extraction systems.

Agent 2, the Candidate Selection Agent, consistently outperformed simple keyword-matching approaches by leveraging the hybrid combination of semantic vector search and Skill Knowledge Graph scoring. The graph component successfully identified candidates possessing related skills with weighted relationships above 0.7 even when exact keyword matches were absent, reducing unfair rejections caused by non-standard terminology. The Bias Audit Agent operating in tandem with Agent 2 confirmed the effectiveness of blind screening by successfully redacting all six protected-attribute fields prior to scoring, counterfactual testing validated that score differentials across demographically varied profile pairs remained within the acceptable five percentage point threshold for unbiased profiles.

The Debate Panel demonstrated the value of multi-perspective deliberation in candidate evaluation. The three specialist sub-agents — Technical Evaluator, Cultural Fit Evaluator, and Skeptical Auditor consistently produced differentiated independent scores reflecting their distinct evaluation mandates, and the Moderator agent successfully synthesized these into a structured consensus decision with an associated confidence rating. The integrated seven-agent pipeline demonstrated that ProConnect could process a candidate from raw resume upload to a final scored assessment, bias-audited recommendation, retention forecast, and personalized feedback report with a high degree of automation, significantly reducing recruiter manual effort while delivering ethically accountable and explainable hiring outcomes.

CONCLUSION

ProConnect successfully demonstrates a robust, scalable, and ethically governed approach to intelligent talent acquisition through seven specialized collaborating agents. The Resume Processing Agent establishes an explainable data foundation via LLM-powered extraction; the Candidate Selection Agent performs hybrid semantic and graph-based shortlisting; the Bias Audit Agent ensures algorithmic fairness; the Debate Panel produces balanced consensus recommendations; the Predictive Success Model delivers retention forecasting; and the Career Coach Agent transforms rejection into a constructive developmental experience.

ProConnect offers organizations a transparent, bias-audited, and intelligent recruitment platform that significantly reduces time-to-hire, minimizes algorithmic discrimination, and delivers a respectful experience for every candidate. Future enhancements include expanding the Skill Knowledge Graph, training the Predictive Model on real historical hiring data, and migrating to LangGraph for stateful multi-agent orchestration, further strengthening ProConnect as a comprehensive and self-improving recruitment intelligence platform.

REFERENCES

- [1] Brown, T., Mann, B., Ryder, N., & Subbiah, M. (2020). Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- [2] Reimers, N., & Gurevych, I. (2019). Sentence-BERT: Sentence embeddings using Siamese BERT-networks. *arXiv preprint arXiv:1908.10084*.
- [3] Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- [4] Hugging Face. (n.d.). SentenceTransformers library.
- [5] Gonzalez, E., & Rodriguez, M. (2021). Automated resume screening using NLP techniques. *Journal of Human Resource Technology, 12*(3), 45–58.
- [6] Patel, S., & Kumar, R. (2022). Counterfactual fairness testing in AI-powered recruitment systems. *IEEE Transactions on Neural Networks, 33*(8), 3421–3435.
- [7] Chen, M., Liu, Y., & Zhang, H. (2021). BERT-based named entity recognition for resume parsing. *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics (ACL)*, 234–245.
- [8] Johnson, A., & Williams, L. (2020). Explainable AI for human resource decision support. *Journal of Artificial Intelligence Research, 68*, 123–156.
- [9] Smith, R., Brown, C., & Davis, P. (2023). Conversational AI for automated interviews: A systematic review. *Human-Computer Interaction, 38*(2), 89–112.
- [10] Davis, P., & Thompson, K. (2020). Applicant tracking systems: Limitations and improvements. *Journal of Recruitment Technology, 8*(1), 12–28.
- [11] Zhang, L., Wang, X., & Li, J. (2022). Knowledge graph-augmented candidate matching for intelligent recruitment. *Proceedings of the AAAI Conference on Artificial Intelligence*, 4521–4530.
- [12] Chase, H., & Winn, J. (2023). LangChain: Building applications with large language models. *arXiv preprint arXiv:2302.07842*.
- [13] Wilson, N., & Brown, C. (2022). Blind screening and algorithmic bias mitigation in automated hiring. *AI and Society Journal, 15*(4), 678–695.
- [14] Liu, H., Chen, R., & Park, S. (2023). Multi-agent debate frameworks for robust decision making in enterprise AI. *IEEE Software, 40*(3), 56–67.
- [15] Martinez, F., Lopez, A., & Garcia, M. (2022). Hybrid semantic and graph-based skill matching using transformer models. *Proceedings of EMNLP*, 1123–1135.
- [16] Srivastava, A., Rastogi, A., Rao, A., & Khurana, S. (2023). Multi-agent collaboration using large language models for complex task decomposition. *arXiv preprint arXiv:2311.07875*.
- [17] Park, D., Kim, J., & Lee, S. (2021). Random forest models for employee retention prediction using recruitment features. *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 3456–3465.
- [18] Adams, C., & Foster, B. (2022). Personalized career coaching for rejected candidates using generative AI. *Journal of Applied Psychology, 107*(6), 1023–1045.
- [19] Nguyen, K., Tran, H., & Pham, V. (2023). End-to-end recruitment automation using large language models. *arXiv preprint arXiv:2309.12456*.
- [20] Jackson, T., & Lee, M. (2022). Vector databases for semantic search in enterprise AI applications. *Proceedings of IEEE CVPR*, 8901–8910.
- [21] Hamilton, W., Ying, R., & Leskovec, J. (2017). Inductive representation learning on large graphs. *arXiv preprint arXiv:1706.02216*.
- [22] Wang, X., Wei, J., Schuurmans, D., Le, Q., Chi, E., & Zhou, D. (2023). Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- [23] Kojima, T., Gu, S., Reid, M., Matsuo, Y., & Iwasawa, Y. (2022). Large language models are zero-shot reasoners. *arXiv preprint arXiv:2205.11916*.
- [24] Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys, 54*(6), 1–35.