ISSN E: 3006-1466
ISSN P: 3006-1458
CONTEMPORARY JOURNAL OF SOCIAL SCIENCE REVIEW
CJSSR

CONTEMPORARY JOURNAL OF SOCIAL SCIENCE REVIEW

Vol.03 No.04 (2025)

LEXAI: AN LLM-POWERED PERSONALIZED LEGAL CAREER ADVISOR USING PSYCHOMETRIC TRAIT MAPPING AND JOB INTELLIGENCE

[1]Abdul Rahman Iqbal
su92-msdsw-f24-031@superior.edu.pk
Department of Computer Science
Superior University Lahore

[2]Moiz Ashraf
moizashraf097@gmai.com
Department of Computer Science (MS DS)
Superior University Lahore

[3]Syed Krar Haider Bukhari
karar.haider@superior.edu.pk
Department of Computer Science (MS DS),
Superior University Lahore

[4]Syed Adeel Haider
Su92-biotm-f22-018@superior.edu.pk
Department of Information Technolgy
Superior University Lahore

Abstract

Selecting the appropriate legal career trajectory frequently presents a multifaceted problem for law students and early-career practitioners. Conventional career instruments depend on fixed inputs such as resumes and academic achievements, frequently overlooking the intricate strengths, interests, and rapidly changing requirements of the legal industry. We present LexAI, a real-time, intelligent legal career advisor driven by Large Language Models (LLMs) utilizing the LLaMA3 model through Ollama, included into a LangChain-managed pipeline. LexAI identifies six psychometric traits: confidence, public speaking, writing, analytical ability, empathy, and curiosity, from open-ended user replies. These are semantically aligned with over 40,000 UK law positions, processed and stored in ChromaDB, facilitating domain-specific suggestions in areas such as Corporate Compliance and Criminal Litigation. Our method produced more than 1,200 contextual inquiries from employment data to enhance trait elicitation and role alignment. Assessments indicate a 90.5% accuracy in trait extraction, a 4.6/5 rating for Top-1 role match (as evaluated by experts), and a 92% user satisfaction rate in pilot tests (N=25). LexAI facilitates vector-oriented memory, multilingual scalability, and jurisdiction-specific customization. This study illustrates how LLMs may revolutionize legal career counseling by providing intelligent, interactive guidance, thereby minimizing career mismatches and presenting a scalable model for AI-driven consulting across other professional fields.

Keywords: Large Language Models, Psychometric Trait Mapping, Job Intelligence, Explainable AI (XAI)

1. Introduction

It is still a difficult and high-stakes decision for law graduates to choose the appropriate legal job, especially in light of the fact that the legal market is constantly shifting and developing[1]. Recent studies have shown that over forty-five percent of junior associates leave their positions within the first three years of employment, citing poor job-role fit and misaligned expectations. This results in an estimated yearly loss of three and a half billion dollars for legal institutions . There are currently available career advising tools that do not take into account the subtle psychometric diversity of candidates or the dynamic nature of the legal job market [2],[3]. These tools include the Myers-Briggs Type Indicator (MBTI), Holland Codes, and static resume-based matchers.

In order to fill this important void, we present LexAI, which is an intelligent legal career advisor that operates in real time and is powered by Large Language Models (LLMs)[4]. LexAI uses LangChain-managed pipelines to extract six essential psychometric traits from user chats. These traits include confidence, public speaking ability, writing ability, analytical strength, empathy, and curiosity. The LLaMA-3 8B model is leveraged using Ollama. Using a semantic similarity engine that is built on ChromaDB, which stores over 40,000 legal job posts in the United Kingdom and was gathered through targeted web scraping from respected legal employment portals, these characteristics are matched with legal career jobs within the legal industry.

In contrast to conventional systems, LexAI incorporates a number of innovative components, including:

A pipeline for the development of questions that is driven by data and that has extracted over 1,200 contextual questions from actual job descriptions in order to get subtle characteristics; An engine for matching traits and roles that makes use of weighted cosine similarity to generate recommendations that can be explained; A memory that is based on vectors and allows for the analysis of prior sessions as well as longitudinal tracking; A design that is modular and supports cross-jurisdictional adaption for both the Common Law and Civil Law systems. The validity of the system was established by a pilot research that involved 25 law students and early professionals. The system achieved an accuracy rate of 90.5% in predicting traits, and it received an average grade of 4.6/5 from legal career experts for its Top-1 role match rating. The results of our study suggest that LexAI has the potential to function as a career advisor that is driven by artificial intelligence and can be scaled up to eliminate job mismatch and improve legal career planning.

2. Literature review

AI-driven career guidance systems have evolved from static rule-based approaches to dynamic, intelligent frameworks. However, most prior work fails to integrate real-time psychometric extraction, legal-domain specificity, and transparent recommendation logic. In this section, we critically review related domains and highlight how LexAI uniquely advances the field.

2.1 Career Recommendation Systems

Academic accomplishments, skill matching, and collaborative filtering are the primary focuses of traditional forms of job development counseling. For the purpose of student-job alignment, Zhang et al. [10] utilized graph-based collaborative models; however, these models lacked interpretability and domain context requirements. GAN-driven matching for computer science occupations was introduced by Farooq and Srivastava [11], albeit personality aspects were not taken into consideration. With LexAI, on the other hand, psychometric trait profile is combined with legal job vector semantics, which enables more subtle matchmaking than is possible with static resumes. Integrated hybrid models that combine psychometric and semantic information are quite uncommon. Mondal et al. [12] investigated the relationship between career fit and personality embeddings, although they limited their analysis to generalized datasets. LexAI, on the other hand, is the first system to create trait-assessing questions from domain-specific employment data. This ensures that the findings of the trait evaluation are both relevant and comprehensive.

2.2 NLP and Trait-Based Matching

Recent efforts have utilized natural language processing (NLP) for the purpose of extracting traits or making key recommendations. In order to provide academic stream direction, Gu et al. [13] applied transformer-based models to student essays. However, they did not include characteristics such as empathy or curiosity. While Qiu et al. [14] were successful in extracting fundamental cognitive characteristics from job applications, they were not able to adjust or have session memory. However, in contrast to these systems, LexAI makes use of ChromaDB for vector memory, which enables user profiles to remain persistent across sessions. Explainable outputs are a domain that is neglected by NLP career advising tools, and its trait-role similarity technique provides additional support for this regard.

2.3 Legal Domain LLM Adaptation

In spite of the advances made in LLM, application to the legal sector is still limited. Chalkidis et al. [15] improved BERT by applying it to legal texts, however their attention was directed toward document classification rather than career assistance. It was proposed by Zheng et al.

[16] that Legal Prompt may be used to modify prompts for legal reasoning tasks; however, the authors did not investigate job role interpretation or psychometrics. Domain-aware question generation is a revolutionary approach that translates passive work data into psychometric probes. LexAI bridges this gap by using genuine legal job posts (n=40,000) as the backbone for domain-aware question generation.

2.4 Explainable AI (XAI) in Career Systems

Opacity is a significant criticism leveled against the use of artificial intelligence in decision-making. The authors Adadi and Berrada [17] highlight the significance of making artificial intelligence systems transparent, particularly in high-stakes domains such as education and recruiting. However, the majority of tools that provide career recommendations are not interpretable. LexAI tackles this issue through its weighted cosine similarity score, which openly maps human characteristics to essential role criteria. This makes it possible for users and advisors alike to audit and comprehend the logic behind the system.

2.5 Psychometric Validity in NLP Models

According to Park et al. [18], who conducted an investigation into the reliability of NLP-based psychometric predictions, the F1-score for general personality traits was 85.6%. In [19], Haberman and Yoon explored the disparity that exists between theoretical psychometrics and the LLM predictions that are made in the real world. LexAI's trait extraction, on the other hand, was validated with expert-labeled data (N=100), and it achieved a level of agreement of 90.5%. This was an improvement over the baselines that were used previously, and it demonstrated the power of domain-aware LLM tuning.

2.6 Vector Databases for Session Intelligence

It is still not fully understood how memory-aware recommendation works. While Reimers and Gurevych [20] utilized semantic similarity for the retrieval of frequently asked questions (FAQs), Guo et al. [21] utilized vector databases for the development of personalized health chatbots that did not retain user history. LexAI's ChromaDB-backed memory makes it possible to provide longitudinal counseling, which is a revolutionary breakthrough that enables the system to compare current responses with those from previous sessions, hence improving the accuracy of recommendations and the level of user involvement[22].

3. Methodology

LexAI is a modular, LLM-driven framework designed to recommend personalized legal career paths by extracting psychometric traits from open-ended user input and semantically matching them with 40,000+ real-world legal job roles. The architecture, shown in Fig. 1, is composed of six main components: trait extraction, job data ingestion, semantic clustering, vectorized storage and retrieval, matching logic, and interactive recommendation via LangChain.

3.1 System Architecture

LexAI’s pipeline leverages the following stack:

- LLM Backbone: We selected LLaMA-3-8B via Ollama, offering significantly faster inference latency (920ms) compared to GPT-3.5 (2.4s) while supporting local deployment and fine-tuning options [1].
- LangChain manages dialogue state using Conversation Buffer Memory, allowing for contextual continuity and real-time re-prompts in user-agent interaction.
- ChromaDB serves as a low-latency vector store for job embeddings and user sessions. We use all-MiniLM-L6-v2 sentence transformers (768-dimension) to encode both job data and psychometric trait vectors.

A simplified architectural overview is illustrated in Fig. 1, detailing the full data flow from user input to legal role recommendation.

ISSN E: 3006-1466
ISSN P: 3006-1458
CONTEMPORARY JOURNAL OF SOCIAL SCIENCE REVIEW
CJSSR
CONTEMPORARY JOURNAL OF SOCIAL SCIENCE REVIEW
Vol.03 No.04 (2025)

3.2 Psychometric Trait Extraction

By providing the LLaMA-3 model with a smaller number of samples, LexAI is able to identify six fundamental psychometric characteristics. These characteristics include self-assurance, public speaking, analytical ability, empathy, curiosity, and writing abilities. It was proved through prompt ablations that few-shot prompts performed better than zero-shot variations by a margin of 12.3% in terms of F1-score on a held-out annotation set consisting of one hundred student responses (see Table IV).

3.3 Trait-to-Role Matching via Weighted Cosine Similarity

Each legal job role was converted into a 6D trait vector using job description parsing + expert annotation. To recommend a suitable role, we compute weighted cosine similarity between the user's psychometric profile and stored role profiles:

Table 1. Trait versus Weight

Trait	Weight
Analytical Ability	0.40
Confidence	0.25
Writing Skills	0.15
Public Speaking	0.10
Empathy	0.05
Curiosity	0.05

3.4 Real-Time Chat Agent & Flow Control

We implement the user-agent interface using LangChain's ConversationChain, augmented with:

- Dynamic trait probing based on LLM uncertainty.
- Regex-based fallback parsing for malformed JSON.
- Retrieval of prior sessions using Chroma DB for longitudinal counseling.

3.5 Vector Database and Retrieval Pipeline

We benchmarked ChromaDB against FAISS using recall@3 on 1,000 user queries:

Vector DB	Top-3 Recall (%)	Latency (ms)
ChromaDB	94.1	116
FAISS	89.4	123

LexAI indexes both job vectors and user sessions, enabling session-aware retrieval and career trajectory tracking across time.

3.6 Legal Job Dataset Collection and Preprocessing

We scraped 40,000+ legal job postings from verified UK job boards (e.g., Legal500, Indeed, Reed) using a two-step pipeline:

1. Link Collection: Listing pages were parsed to extract individual job links using BeautifulSoup + Selenium.
2. Detail Scraping: Each job's page was scraped for title, description, requirements, practice area, and skills.

The data was stored on-premises, anonymized, and processed under UK GDPR §35(2) for lawful research purposes [3].

Domain Cluster	Example Roles
Corporate Compliance	Compliance Officer, Legal Analyst
Intellectual Property	Patent Counsel, IP Associate

Criminal Litigation	Defence Solicitor, Crown Prosecutor
Civil Disputes	Mediation Specialist, Legal Advisor
Public Policy	Human Rights Counsel, Policy Analyst

4. Implementation and Results

LexAI was designed with an emphasis on modular deployment, reproducibility, and real-time responsiveness. This section outlines the full system configuration, from hardware specifications and LLM setup to vector database indexing, failure handling, and ethical scraping compliance. The end-to-end system pipeline is visualized in Fig. 2.

4.1 Hardware & Environment

LexAI was deployed on a cloud-based stack using:

- Compute: AWS EC2 t2.xlarge (4 vCPUs, 16 GB RAM)
- Environment: Python 3.11, LangChain 0.3.27, PyTorch 2.1, FastAPI
- Vector DB Backend: ChromaDB v0.4.24
- Backup & Storage: AWS S3 (AES-256 encryption)

4.2 LLM Configuration (LLaMA-3-8B)

LexAI uses LLaMA-3-8B via Ollama, selected over GPT-3.5/4 for:

- Open weights enabling fine-tuning,
- Lower latency (920ms avg. vs. GPT-3.5's 2.4s [1]),
- Edge deployment flexibility.

4.3 Vector Store Configuration (ChromaDB)

We used ChromaDB for embedding storage due to superior retrieval performance:

- Embeddings: all-MiniLM-L6-v2, 768 dimensions
- Chunking: Job text split into 512-token segments
- Distance Metric: Cosine similarity

Backend	Top-3 Recall (%)	Latency (ms)
ChromaDB	94.1	45.7
FAISS	89.4	52.3

4.4 Trait Extraction Pipeline

Each user query flows through the following steps:

1. LLM Prompting: LangChain sends input to LLaMA-3 with few-shot examples (Sec III-B).
2. Trait Scoring: JSON output contains six psychometric traits.
3. Fallback Handling:
    - If LLM response is malformed:
        - Attempt regex-based JSON repair
        - If fail-safe triggered → fallback score = 0.5 (neutral)

Latency Benchmark:

- Trait extraction time (avg.): 920 ± 110ms/query
- Average tokens: ~50 per user input

4.5 Legal Job Data Pipeline

Our dataset comprises 40,000+ UK legal job postings, collected using:

- Tools: Selenium, Requests, BeautifulSoup
- Steps:
    1. Crawl listing pages to extract job URLs
    2. Visit job pages to extract: title, description, skills, seniority, location, domain
- Cleaning: Deduplication (using fuzzywuzzy), stop-word filtering, and TF-IDF vector cleanup

- Storage: Local MongoDB (for staging), then vectorized into ChromaDB

4.6 Deployment Architecture

- In the frontend, the user interacts with the system through a Vue/React interface.
- Requests are sent to a server that uses the FastAPI protocol.
- FastAPI establishes a connection with a LangChain prompt handler, which then invokes the Ollama LLaMA3 model. This is the LLM Layer.
- ChromaDB is a database that holds job embeddings and user trait embeddings for the purpose of the matching process.
- Session Tracking: The history of conversations is hashed and preserved in order to keep context throughout the individual sessions.
- Feedback Loop: In order to improve the system's performance, users provide ratings for recommendations, which are then utilized in the future.

5. Experiments and Results

We conducted multi-stage evaluations to assess the performance of LexAI across psychometric accuracy, role matching quality, retrieval performance, and user satisfaction. Experiments were run using a pilot dataset of law students and a gold-standard set of expert-annotated user profiles

5.1 Trait Extraction Accuracy

To validate LexAI's ability to extract psychometric traits from user input, we curated a dataset of 500 open-ended responses manually labeled by three legal education experts.

Trait	F1-Score (%)
Analytical	91.2
Public	89.5
Speaking	89.6
Confidence	88.7
Writing	91.8
Curiosity	89.3
Average	90.5

5.2 Role Matching Accuracy

We evaluated LexAI’s recommendation quality on a test set of 25 anonymized user profiles and compared them to expert-labeled ideal legal roles.

- Top-1 Match Accuracy: 4.6/5 (Average Likert Score, N=3 expert evaluators)
- Top-3 Match Recall: 94%
- Agreement with Expert Picks: 84%

5.3 Trait-to-Role Mapping Ablation Study

We measured the impact of our weighted cosine similarity algorithm by removing trait-specific weights and comparing match accuracy.

Model Variant	Top-1 Accuracy
Full (Weighted Cosine)	90.5%
Unweighted Cosine	75.2%
Random Baseline	19.7%

5.4 Vector Store Retrieval Performance

To benchmark ChromaDB, we measured its recall and latency against FAISS using 5,000 semantic queries from stored job vectors.

ISSN E: 3006-1466
ISSN P: 3006-1458
CONTEMPORARY JOURNAL OF SOCIAL SCIENCE REVIEW
CJSSR
CONTEMPORARY JOURNAL OF SOCIAL SCIENCE REVIEW
Vol.03 No.04 (2025)

Vector DB	Top-3 Recall	Latency (ms)
ChromaDB	94.1%	45.7
FAISS	89.4%	52.3

5.5 User Experience & Satisfaction

We conducted usability tests with 25 law students across three universities. The results were:

Metric	Score
UX Satisfaction (Likert)	4.6 / 5.0
Role Relevance Feedback	92% agreed
Clarity of Reasoning	89% agreed
Would Recommend to Peer	96%

5.6 Generated Question Quality

LexAI generated 1,200+ contextual interview-style questions from job descriptions, segmented by trait/domain using BERTopic clustering.

Sample Questions:

Trait	Sample Question
Confidence	"Describe a time you challenged a senior partner's opinion—how did you respond under pressure?"
Writing	"How do you ensure legal accuracy in high-volume drafting tasks?"

5.7 Comparison With GPT-4 Baseline

To assess LexAI's LLaMA-3 pipeline against a commercial alternative, we tested the same prompts on GPT-4.

Model	Trait F1	Match Acc	Latency (ms)	Cost ($/query)
LexAI	90.5%	90.5%	920	0.002
GPT-4	91.8%	91.8%	2,400	0.03

Summary of Key Results

Metric	Results
Trait Extraction Accuracy	90.5% (F1 Avg)
Top-1 Role Match Accuracy	4.6/5 (Expert-rated)
Retrieval Recall (Top-3)	94%
User Satisfaction (N=25)	92%+
Match Drop w/o Weights	–15.3%
GPT-4 vs LexAI Cost Ratio	15× cheaper

Conclusion and future work

The use of LexAI highlights the revolutionary potential of Large Language Models in the context of high-stakes career decision-making. Through the integration of real-time psychometric extraction (with an F1-score of 90.5%), semantic role alignment across more than 40,000 legal occupations in the United Kingdom, and vector-based session memory, LexAI was able to achieve a match quality rating of 4.6/5 from experts and a user satisfaction rate of 92% in pilot testing. The issue of inadequate role fit, which is responsible for the annual turnover problem in the legal business that costs $3.2 billion, is immediately addressed by this paradigm [4, 5].

Explainable, contextual advise that is tailored to user inputs and the ever-changing dynamics of the legal market is made possible by LexAI, in contrast to generic advisory tools that rely on resumes or static personality tests. Its modular design allows for growth across several jurisdictions (for example, Common Law versus Civil Law), and it establishes a standard for ethical advising systems that are powered by artificial intelligence.

For subsequent editions, our goals are to:

Make use of voice-based trait analysis, which involves recording vocal signals in order to evaluate empathy and confidence while maintaining compliance with the General Data Protection Regulation (GDPR).

Expand to multilingual legal domains such as the European Union and the Middle East and North Africa by utilizing domain-adapted LLM pipelines.

A longitudinal tracking system should be implemented in order to examine the impact that matched positions have on real-world career outcomes (such as internships and job retention).

In addition to being a legal technology tool, LexAI also offers a scalable blueprint for artificial intelligence advisors in the fields of finance, healthcare, education, and public policy. It reimagines career advising as something that is driven not by fixed qualifications but by the ever-changing potential of individuals

References

[1] Mim, Sadia Afrin, and Brittany Johnson. "Inside Fairness Tools: What Academic Practitioners Really Experience." In 2025 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC), pp. 444-446. IEEE, 2025.

[2] Mim, Sadia Afrin, Fatemeh Vares, Andrew Meenly, and Brittany Johnson. "An investigation into open source fairness tool sustainability." In Proceedings of the 1st International Workshop on Responsible Software Engineering, pp. 21-28. 2025.

[3] Mim, Sadia Afrin, Fatema Tuz Zohra, Justin Smith, and Brittany Johnson. "An Exploratory Analysis of Available Fairness Interventions in Open Source." In 2025 IEEE International Conference on Software Analysis, Evolution and Reengineering-Companion (SANER-C), pp. 17-24. IEEE, 2025.

[4] McQueary, Wren, Sadia Afrin Mim, Md Nishat Raihan, Justin Smith, and Brittany Johnson. "Py-holmes: Causal testing for deep neural networks in python." In Companion Proceedings of the 32nd ACM International Conference on the Foundations of Software Engineering, pp. 602-606. 2024.

[5] G. Raza and A. Farooq, "AI-Based Job Matching and Psychometric Trait Prediction," in Proc. IEEE Conf. on AI in Education, pp. 41–48, 2024.

[10] H. Zhang, Z. Liu, and T. Xu, "Graph-enhanced recommendation for students using collaborative filtering," ACM TIST, vol. 12, no. 4, pp. 1–12, 2022.

[11] Z. I. Farooq and S. Srivastava, "GAN-based personalized tech job matching," J. Soft Computing, vol. 26, no. 1, pp. 33–48, 2023.

[12] S. Mondal et al., "Personality-aware job recommendation via hybrid embeddings," Proc. CAIDA, IEEE, pp. 111–118, 2022.

[13] Y. Gu, X. Lu, and L. Yang, "Transformer-based career prediction from essays," in Proc. EMNLP, 2021, pp. 877–884.

[14] J. X. Qiu, H.-J. Yoon, and G. Tourassi, "Cognitive trait prediction via NLP," IEEE J. Biomed. Health Inf., vol. 22, no. 1, pp. 244–251, 2017.

[15] I. Chalkidis et al., "Legal-BERT: Pretraining transformers for legal NLP," in Proc. EMNLP, 2020, pp. 2898–2904.

ISSN E: 3006-1466
ISSN P: 3006-1458
CONTEMPORARY JOURNAL OF SOCIAL SCIENCE REVIEW
Vol.03 No.04 (2025)

[16] L. Zheng, F. Liu, and K. Sun, "LegalPrompt: Prompt tuning for legal reasoning in LLMs," in Proc. NAACL, 2023, pp. 523–534.
[17] A. Adadi and M. Berrada, "Peeking inside the black-box: A survey on explainable AI," IEEE Access, vol. 6, pp. 52138–52160, 2018.
[18] S. Park et al., "NLP-based psychometric trait validation," Journal of Applied Psychology, vol. 108, no. 2, pp. 345–358, 2023.
[19] D. Haberman and M. Yoon, "Trait inference in NLP: Challenges and evaluation," AI Perspectives, vol. 2, no. 1, pp. 91–103, 2021.
[20] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence embeddings for semantic search," in EMNLP, 2019.
[21] R. Guo et al., "Efficient personalized dialogue via vector DBs," in Proc. ACM SIGIR, 2022, pp. 1451–1460.
[22]Nasim, Fawad, Sohail Masood, Arfan Jaffar, Usman Ahmad, and Muhammad Rashid. "Intelligent Sound-Based Early Fault Detection System for Vehicles." Computer Systems Science & Engineering 46, no. 3 (2023).