International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 05 , May - 2026

Abroad Compass: A Multi-Agent Assistant for International Student Relocation

Arya Mandhare
Savitribai Phule Pune University
Pune, India

Ojvardhan Kale
Savitribai Phule Pune University
Pune, India

Dr. Manju Pawar
Savitribai Phule Pune University
Pune, India

***Abstract*—Relocating to a foreign country for higher education is a multifaceted challenge that involves lack of legal knowledge, organizational inefficiency, and emotional hurdles. Existing digital solutions are often scattered all over the internet, forcing students to manually aggregate data from fragmented sources. This paper presents Abroad Compass— a comprehensive multi-agent AI based system that streamlines the relocation process. Using a LangChain-based routing architecture grounded in zero-shot intent classification router [1], the system integrates specialized agents for legal guidance using Retrieval Augmented Generation (RAG), airfare forecasting using time-series regression [7], [9] with verified real time data, and social matching via SBERT embeddings [11]. Furthermore, the system employs affective computing models [17] to provide contextual emotional support. The implementation leverages a FastAPI backend and an Ionic-based front-end to provide a scalable, cross-platform solution in the form of a chat-bot application. Experimental results indicate that the modular agentic approach significantly reduces user latency and improves the accuracy of domain-specific inquiries compared to monolithic LLM architectures, which reduces hallucinations for sensitive and real-time queries.**

***Index Terms*—Multiagent systems, Retrieval Augmented Generation, SBERT, Flight Fare Forecasting, Legal NLP, Affective Computing, FastAPI, Intent Classification, Hallucination Prevention.**

I. Introduction

The move to a foreign country to study is often more of a logistical marathon than an academic adventure. Millions of students cross borders every year [21] and are overwhelmed by the "information gap." The data available on the internet is rarely personal or integrated. A student must visit a government portal to check visa rules, a commercial aggregator for flights, social media groups for housing, scour through the internet- fact checking every scholarship and that is only the beginning.

This fragmentation creates a "Cognitive Load" [20] that can result in major errors in critical legal documents and financial planning [20]. The psychological burden of navigating these disparate systems simultaneously frequently leads to decision fatigue and suboptimal outcomes for the relocating student.

We developed Abroad Compass with the motive of "a single source of reliable information". Rather than building a generic chat-bot, we designed a modular ecosystem where each agent is a specialist in its domain. This paper details how we leveraged recent advances in NLP, RAG for legal grounding [4], SBERT for semantic matching [11], and sentiment analysis [17] to build a tool that functions less like a search engine and more like a relocation consultant. Inspired by multi-agent conversation frameworks [2] and generative agent architectures [3], these tasks are decomposed into specialized micro-agents orchestrated by a central cognitive router, removing the bottleneck and hallucination hazards of monolithic Large Language Models.

II. Related Work

Abroad Compass represents the intersection of multiple research domains: conversational AI, legal information retrieval, predictive analytics, semantic matching, and affective computing. This section reviews the most relevant prior work in each area. Conversational AI has evolved from simple intent-response pairs to complex, state-aware systems. Chatbots for academic settings have been studied for university FAQ handling [23] and e-learning query resolution [24]. More recently, legal advisory bots [25] and student support frameworks [26] have explored domain-specific deployments. The multi-agent conversation paradigm, formalized by AutoGen [2], enables agent specialization at scale.

A. Legal RAG Agent

The problem of "hallucination" is especially concerning in the legal field. Research on Legal-BERT [4] shows that domain-specific pre-training is necessary to capture the subtlety of immigration-related documents. Retrieval-Augmented Generation (RAG) via dense passage retrieval [5] has become the benchmark approach for grounding model outputs in verified source documents [6].

B. Predictive Analytics in Travel

Airfare pricing is highly volatile, it depends on macroeconomic variables, seasonal trends, and route traffic density and other factors. Prior work has explored GRU-based deep learning models [7], multi-attribute dual-stage attention mechanisms [8], and ensemble machine learning approaches [9] for fare prediction.

C. Semantic Matching and Recommendation

Point-of-interest recommendation has been addressed through context-aware interpretable frameworks such as

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 05 , May - 2026

CAPRI [14]. Social recommendation via Graph Neural Networks [13] enables community-aware suggestions. For roommate matching, personality-aware systems [27] have shown methodologies beyond simple hard-filter approaches.

III. System Architecture

Abroad Compass employs a distributed, multi-layered architecture designed for scalability and low-latency communication. Through integration of these layers, the system undergoes seamless transitions from unstructured user inputs to intent-specific outputs guided by it's dedicated agent.

A. Authentication Process

The system implements user access via JSON Web Tokens (JWTs) [22] for authentication via stored credentials (e.g., user email and hashed password using bcrypt) in a SQLite database during the registration process based upon fields collected within the user's profile, including their home country, destination, field of study, and living preference(s). When a user logs in, the system compares the password entered against the stored. If they match, a secure token based upon the user's unique ID will be generated and passed on during following requests for access to secured functionality. FastAPI will then verify the token, retrieve the user's profile from the database, and utilize that profile to provide secure and customized services meeting the user's needs.

B. Presentation Layer (Frontend)

The frontend was developed using React along with Ionic framework for providing cross-platform performance. There are four key areas in the application: authentication, chatting, roommate matching, and searching for scholarships. These pages can be accessed via Ionic React Router. In order to maintain secure connections during the user session, there is an AuthProvider, which deals with the user profiles and JWTs [22]. Whenever an user attempts to open some restricted pages while not being signed in, application will be redirected login page immediately. In terms of the visual component, there is use made of such common Ionic elements as cards, lists, and forms. All requests to the backend were developed using RESTful APIs provided by FastAPI with a security token added to each of them.

C. Application Gateway

The application gateway is constructed using FastAPI, the backend is designed as a gateway that intelligently sends requests to different services. The application is properly configured to talk to the frontend and consists of seven modules, namely, authentication, flights, scholarships, roommates, legal advice (RAG), and chat. Each module acts as a standalone microservice with its own logic. Therefore, when problems arise for specific AI services like ChromaDB and TinyLLaMA, the entire application keeps functioning uninterrupted. Through the utilization of microservices, the application becomes scalable as the gateway would provide security and error handling for the entire application.

D. Cognitive Core: The LangChain Router

All inbound requests are channeled into the Agentic Router once they enter the system. The LangChain Router Agent utilizes a zero-shot ReAct (Reason + Act) approach [1], through which the relevant agent is selected based on the context of the request. The usage of smaller, quantized local models for intent classification reduces the expenses and latency associated with calling a larger LLM model for routing purposes.

Algorithm 1 Zero-Shot ReAct Routing Logic

Require: User Query Q, Conversation History H, Set of Agent Tools $T = \{t_1, t_2, ..., t_n\}$
Ensure: Selected Tool Output O
$Prompt \leftarrow$ FormatContext(Q, H, T)
$Intent \leftarrow$ LocalLLM.Classify($Prompt$)
if $Intent$ matches $t_k \in T$ **then**
    $ExtractedArgs \leftarrow$ ParseArguments(Q, $t_k.schema$)
    $O \leftarrow$ ExecuteTool(t_k, $ExtractedArgs$)
else
    $O \leftarrow$ DefaultConversationalAgent(Q, H)
end if
return O

IV. Detailed Agent Implementation

Abroad Compass relies on five specialized agents, each governed by distinct algorithmic paradigms tailored to their specific domains.

A. Legal Guidance via RAG

The Legal Guidance Agent answers visa and immigration queries across three jurisdictions (UK, Germany, Australia) via a three-stage RAG pipeline [6]:

- **Indexing:** Official immigration PDFs are segmented using LangChain's `RecursiveCharacterTextSplitter` into 1,000-character chunks with 150-character overlap. Each chunk is encoded using `all-MiniLM-L6-v2` [11] and stored in a ChromaDB vector database.
- **Retrieval:** A user query Q is mapped into a dense vector space $E : \mathcal{X} \rightarrow \mathbb{R}^d$ using the same encoder. ChromaDB retrieves the top-k chunks via cosine similarity:

$$\text{sim}(Q, D_k) = \frac{\mathbf{q} \cdot \mathbf{d}_k}{\|\mathbf{q}\| \|\mathbf{d}_k\|} \tag{1}$$

  The jurisdiction is inferred automatically from query context or the user profile.
- **Generation:** Retrieved chunks are passed to a locally-hosted, quantized TinyLLaMA (1.1B parameters) [19] with a system prompt enforcing strict source adherence, preventing hallucination of immigration law details. Responses include source citations and follow-up checklists.

The fully offline architecture ensures sensitive student data is never transmitted to external servers [25].

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 05 , May - 2026

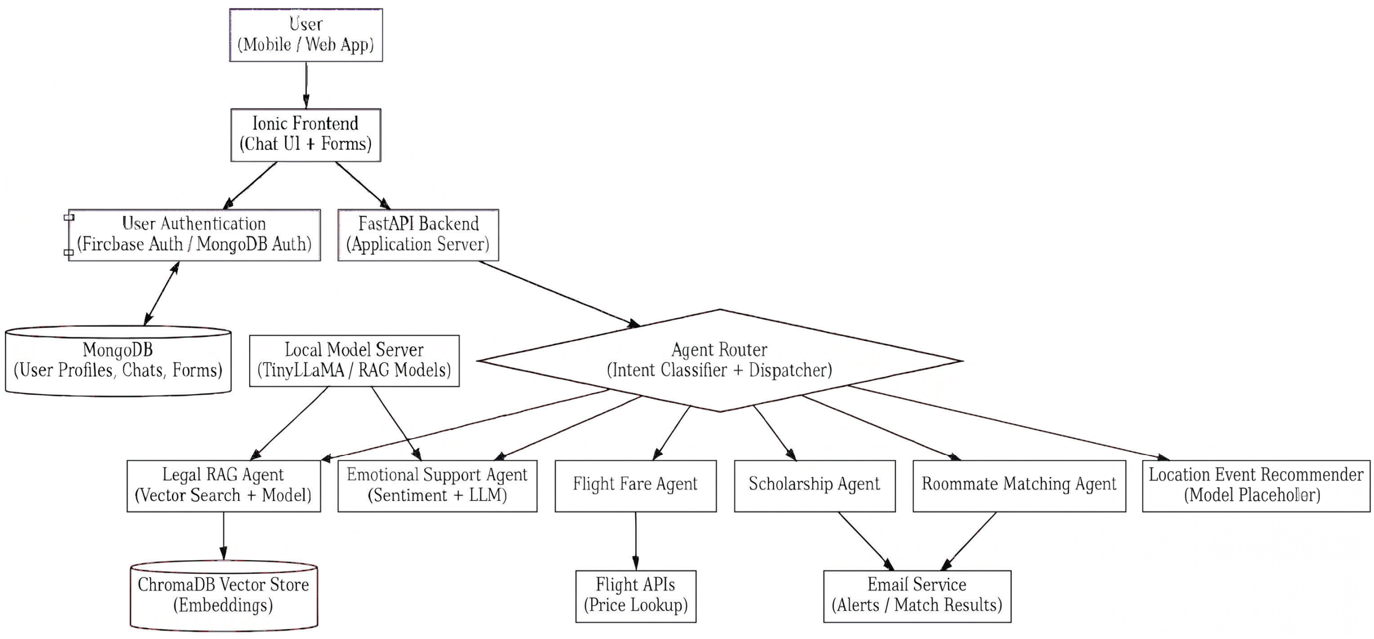

Fig. 1. Detailed System Architecture of Abroad Compass. The figure illustrates the communication between the Ionic Frontend and the FastAPI Backend, highlighting the central role of the LangChain Routing Agent in orchestrating specialized subagents including Legal, Emotional Support, Flight Fare, Scholarship, and Location Recommenders.

B. Flight Fare Forecasting Agent

The Flight Fare Agent employs a hybrid forecasting architecture combining an LSTM-based time-series model [10] for temporal price trend analysis [7], and real-time market validation via the Amadeus API.

- **Feature Engineering:** An 11-dimensional input vector $\mathbf{x}$ is constructed with encoded origin/destination airport codes, lead time (days before departure), base route price, seasonal multiplier, booking window adjustment, departure month, departure day-of-week, booking month, booking day-of-week. The base price incorporates class multipliers and a piecewise booking window function:

$$\alpha(t) = \begin{cases} 1.5 & t < 7 \text{ days} \\ 1.3 & t < 14 \text{ days} \\ 1.1 & t < 30 \text{ days} \\ 0.95 & t < 120 \text{ days} \\ 0.9 & t \geq 120 \text{ days} \end{cases} \tag{2}$$

- **Price Prediction:** A serialized scikit-learn model generates a predicted price $\hat{y}$ from $\mathbf{x}$. An LSTM network [7], [10] independently models temporal fares to capture seasonal trends, with then undergoes ensemble prediction, given by:

$$\hat{y} = \frac{1}{B}\sum_{b=1}^{B} f_b(\mathbf{x}) \tag{3}$$

- **Live Market Reconciliation:** The predicted range is cross-validated against live offers retrieved via the Amadeus Flight Offers Search API (OAuth2). A confidence score c is assigned based on deviation δ between predicted $\hat{y}$ and live price y_{live}:

$$c = \begin{cases} 0.95 & |\delta| \leq 0.15\hat{y} \quad \text{(within range)} \\ \text{flag deal} & y_{\text{live}} < \hat{y} - 0.15\hat{y} \\ \text{wait} & y_{\text{live}} > \hat{y} + 0.15\hat{y} \end{cases} \tag{4}$$

This hybrid approach leverages historical price patterns from the ML model while ensuring users receive up-to-date flight details [8].

C. Roommate Matching Agent

The Roommate Matching Agent matches international students based on lifestyle compatibility by performing weighted content-based filtering [27], supplemented by learned attribute weights and semantic re-ranking via Sentence-BERT [11].

Each user profile records the following: budget range, housing type, cleanliness level $c \in [1, 5]$, smoking and pet preferences, study habits, sleep schedule, personality tags, and interests. Per-attribute similarity scores $s_i(A, B) \in [0, 1]$ are computed as follows. Cleanliness proximity uses a normalized distance:

$$s_{\text{clean}}(A, B) = \max\left(0,\ 1 - \frac{|c_A - c_B|}{4}\right) \tag{5}$$

Interest overlap is measured via the Jaccard index:

$$s_{\text{int}}(A, B) = \frac{|I_A \cap I_B|}{|I_A \cup I_B|} \tag{6}$$

Attribute weights $\mathbf{w} = [w_1, \ldots, w_6]$ are learned via softmax regression over four ordinal compatibility classes (*Poor,*

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 05 , May - 2026

Neutral, Good, Excellent), trained on synthetically generated profile pairs:

$$P(k \mid \mathbf{s}) = \frac{\exp(\mathbf{w}_k^\top \mathbf{s})}{\sum_{j=1}^{K} \exp(\mathbf{w}_j^\top \mathbf{s})} \tag{7}$$

where $\mathbf{s} = [s_1, \ldots, s_6]$ and $K{=}4$. The final compatibility score is:

$$S(A, B) = \mathbf{w}^\top \mathbf{s} \tag{8}$$

Personality descriptions are encoded into 768-dimensional SBERT vectors, providing a semantic re-ranking signal:

$$\text{sim}_{\text{sem}}(A, B) = \frac{\mathbf{v}_A \cdot \mathbf{v}_B}{\|\mathbf{v}_A\| \, \|\mathbf{v}_B\|} \tag{9}$$

Candidates scoring below $\tau{=}0.5$ are filtered out, with the top-10 results returned in descending order. The content-based design avoids the cold-start problem inherent in collaborative filtering [13].

D. Emotional Support Agent

Relocation is inherently stressful. The Emotional Support Agent provides empathetic assistance to students navigating homesickness, culture shock, academic stress, and loneliness. This agent is built on affective computing principles that focuses on delivering sentimental responses [17].

Intent Detection. Emotional intent is identified by the central LangChain Router via zero-shot intent classification [1], which infers affective intent from query semantics without relying on fixed keyword lists. This ensures generalization to paraphrased or implicitly emotional expressions that rigid lexical triggers would miss.

Sentiment Classification. Queries routed to this agent are passed to a fine-tuned RoBERTa classifier [12], [18] that maps the input into one of four affective states: *neutral, mild distress, moderate distress, severe distress*. The classification confidence score p_e is computed as:

$$p_e = \text{softmax}(\mathbf{W}_e \cdot \mathbf{h}_{\text{CLS}} + \mathbf{b}_e) \tag{10}$$

where $\mathbf{h}_{\text{CLS}}$ is the RoBERTa `[CLS]` token embedding and $\mathbf{W}_e$, $\mathbf{b}_e$ are learned classification parameters.

Context-Aware Response Generation. A rolling window of the ten most recent conversation turns is injected into the prompt as alternating User/Assistant exchanges. When p_e exceeds a distress threshold τ_e, the Routing Agent appends a *compassion directive* to the system prompt, ensuring responses are both factually grounded and emotionally appropriate [17], [18].

Output. The agent returns a structured triple:

$$\mathcal{O}_e = \langle r,\ \texttt{id}_{\text{agent}},\ \mathcal{R} \rangle \tag{11}$$

where r is the generated response, $\texttt{id}_{\text{agent}}$ is the agent identifier, and $\mathcal{R}$ denotes a set of surfaced resource categories (*mental-health-resources, emotional-support-guides*), enabling the frontend to present professional contact information contextually.

E. Scholarship Finding Agent

The Scholarship Agent enables international students to discover relevant funding opportunities based on their academic profile and temporal constraints, querying a structured SQLite database of scholarships indexed by name, location, university, provider, grant type, eligibility criteria, deadline, and grant value.

Profile-Aware Query Construction. The agent exposes a RESTful endpoint accepting optional parameters for country, field of study, grant type, and university. When parameters are absent, the authenticated user's destination country and program field are substituted from their registration profile, ensuring personalized defaults without explicit input.

Semantic Retrieval via TF-IDF and NER. Rather than relying on rigid pattern matching, the agent employs Named Entity Recognition (NER) to extract structured entities — field of study, institution type, funding category — from the user's academic profile. These entities are matched against the scholarship corpus using TF-IDF weighted retrieval:

$$\text{TF-IDF}(t, d, D) = \text{tf}(t, d) \cdot \log\left(\frac{|D|}{|\{d \in D : t \in d\}|}\right) \tag{12}$$

where t is an NER-extracted term from the user profile, d is a scholarship document, and D is the full corpus. This ensures retrieval is driven by semantic relevance rather than exact string overlap.

Temporal Filtering. A two-stage deadline filter retains only scholarships whose deadlines δ satisfy:

$$t_{\text{now}} \le \delta \le t_{\text{start}} + 365 \tag{13}$$

where t_{now} is the current date and t_{start} is the user's program start date, ensuring all returned opportunities are temporally actionable. Grant values undergo ISO 4217 currency parsing to extract numeric amounts for consistent downstream comparison.

Output. Results are mapped to a structured schema containing title, country, field of study, deadline, amount, currency, eligibility type, grant type, university, provider, and an external link — surfaced directly within the chat interface for immediate access..

F. Location and Event Recommender

The Location Agent integrates spatial APIs with the CAPRI context-aware recommendation framework [14] and neural collaborative filtering [15]. Rather than returning random suggestions, the agent clusters university campus coordinates and student venues using DBSCAN [16]. A data point p qualifies as a core cluster point when:

$$|N_\epsilon(p)| \ge \text{MinPts} \tag{14}$$

The agent then recommends safe neighborhoods and culturally relevant events within these clusters based on the inferred user preference profile. This agent is in the active development phase.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 05 , May - 2026

V. Implementation and Results

Abroad Compass is developed as a fully workable cross-platform application and evaluated by the system demonstration on all five specialized agents. The main interface components are presented in the following subsections.

A. Authentication and User Onboarding

At sign-up, the authentication layer collects detailed profile metadata including destination country, field of study, program start date, budget range, and housing preferences. This profile is then used by the Scholarship and Roommate agents as default query context which eliminates redundant user input across sessions.

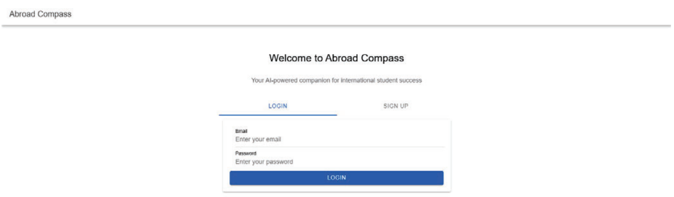

Fig. 2. Login page with password encryption.

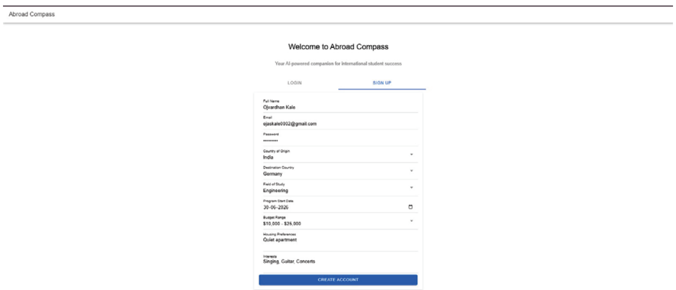

Fig. 3. User registration form for capturing user profile attributes for personalized agent responses across the platform.

B. Central Chat Interface and Agent Routing

All user interaction happens through a single chat interface. The LangChain Router transparently directs queries to the correct specialized agent, and the active agent identifier is surfaced in the top-right of the interface (e.g., *Flight Finder*, *Legal Help*, *Scholarships*). The sidebar navigation allows direct access to each agent module and displays the destination context of the authenticated user (destination country and field of study).

C. Legal Guidance Agent

Figure 6 shows the Legal RAG Agent answering a query about documentation for a German study visa. The answer is based on chunks retrieved from the official immigration PDFs (`gerenal_info_germany_visa.pdf`,

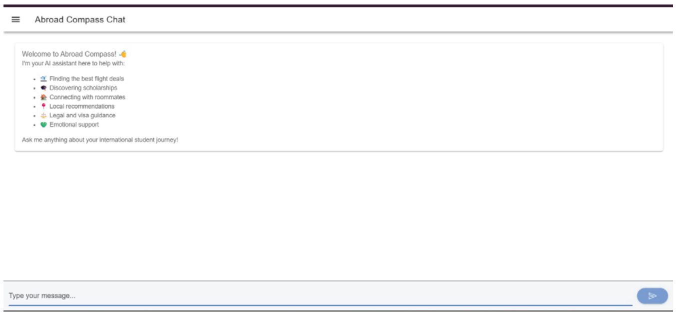

Fig. 4. Abroad Compass Chat interface, showing the system welcome message and enumerating all available agent capabilities.

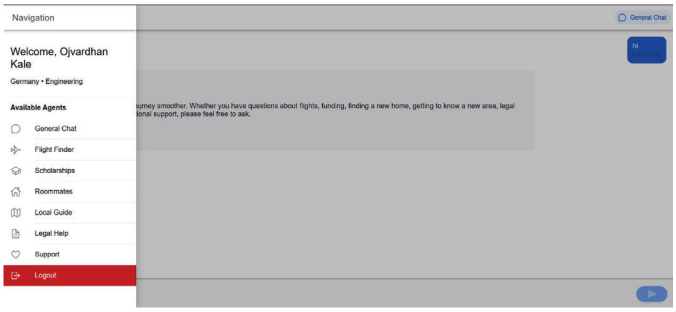

Fig. 5. Personalized agent modules and user context displayed in the navigation drawer (destination and field of study).

`infostudents-data-germany-india.pdf`) with the source citations directly below the response.

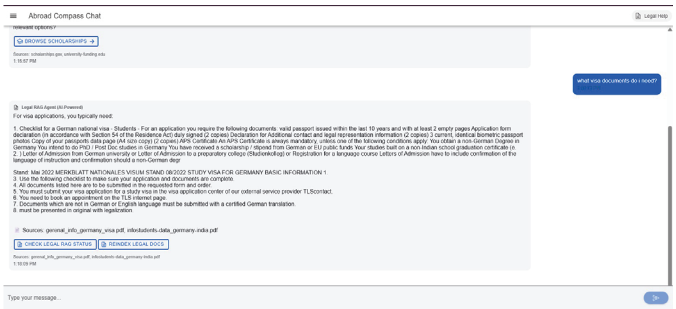

Fig. 6. Legal RAG Agent answering a query about German study visa documentation, citing retrieved source PDFs inline.

D. Flight Fare Forecasting Agent

Figure 7 shows the Flight Fare Agent making a price prediction for the DEL $\rightarrow$ FRA route. The correct cities are selected based on the user preferences filled in the sign up form, if no user specifications are given. The response contains a predicted price range ($591-$722), a confidence score (70%), a recommended booking window (2-4 weeks in advance), and live offers retrieved via the Amadeus API cross-validated against the model prediction. The transition from an emotional support query to a flight query in the same session demonstrates the router's ability to switch agent context dynamically without losing conversational continuity.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 05 , May - 2026

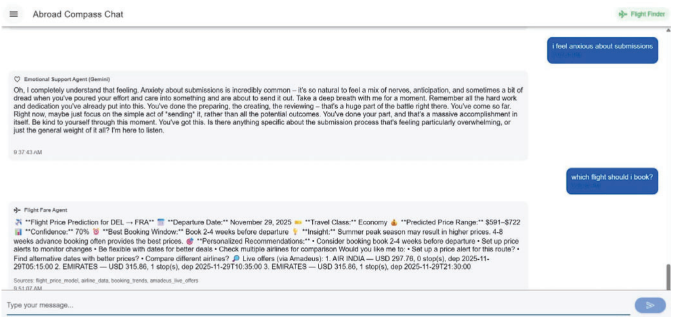

Fig. 7. Flight Fare Agent delivering a DEL→FRA price prediction with Amadeus live offer reconciliation. The previous turn depicts dynamic agent switching from Emotional Support in the same session.

E. Emotional Support Agent

As shown in Figure 8, the Emotional Support Agent responds in a caring way to the query, ”I feel stressed.” It correctly identifies relocation-related anxiety based on the earlier conversation about legal issues. The response acknowledges the user’s feelings, normalizes the experience, and encourages further discussion. This approach aligns with the compassion guideline added by the router when it detects emotional intent. Resource categories, such as `mental-health-resources` and `emotional-support-guides`, are included in the response metadata.

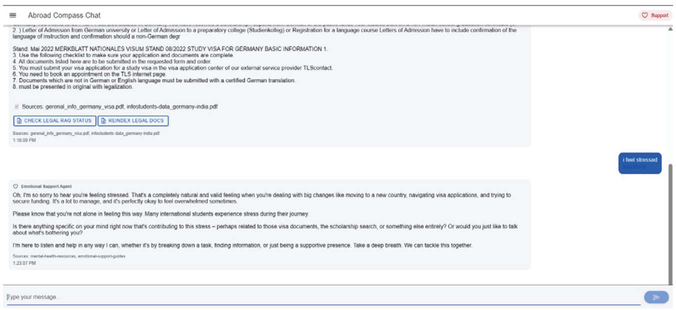

Fig. 8. Emotional Support Agent’s response to a query with contextually grounded empathetic output, following a legal guidance exchange in the same session.

F. Scholarship Finding Agent

Figure 9 shows the Scholarship Recommender filtering out suitable scholarship sources for an Engineering student traveling to Germany, employing defaults based on user profile information without specifying any explicit filter settings. Figure 11 illustrates the specialized scholarship browser that supports multiple attribute filters, including country, field, type of grant and university, resulting in output such as the *Deutschlandstipendium*

G. Roommate Matching Agent

Figures 12 and 13 illustrate the Roommate Matching Agent workflow. Users fill a structured lifestyle profile that includes budget range, housing type, cleanliness level, smoking and pet

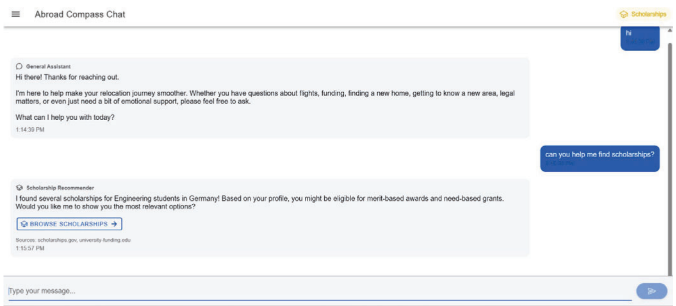

Fig. 9. Scholarship Recommender agent identifying profile-matched funding opportunities via the chat interface, with a direct link to the scholarship browser.

Fig. 10. Scholarship browser displaying filtered results for Engineering students in Germany, including grant value, deadline, institution, and eligibility type.

preferences, study habits, sleep schedule, personality traits, and interests. The SBERT-augmented matching pipeline returns ranked candidates with weighted compatibility scores; as per the given tests, the top match achieved a 97% compatibility score across five shared interest categories, with subsequent matches at 82% and 80%. This reflects meaningful differentiation while scoring the matches.

VI. Conclusion and Future Work

Abroad Compass provides an altogether different agent infrastructure in comparison to monolithic chatbots. Through the utilization of RAG for legal correctness [4], [5], time-series regression for fare estimation [7], semantic matching for social compatibility [11], and affective computing for emotional intelligence [17], we provide a holistic relocation assistant experience. The cloud native architecture ensures that there is scalability and high availability.

Modularization ensures that specific agents can be upgraded or swapped out without affecting the core router functionality. Some potential improvements that could be done include: (1) Increasing the volume of documents related to other destinations within the legal document corpus; (2) Automate the intake process for scholarship and immigration policies through live updates; (3) Utilizing multiple modalities for input within the agent pipeline; (4) Complete implementation of location and event finder agent; and (5) Utilizing cloud infrastructure to increase availability and dynamic scaling as user traffic increases.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 05 , May - 2026

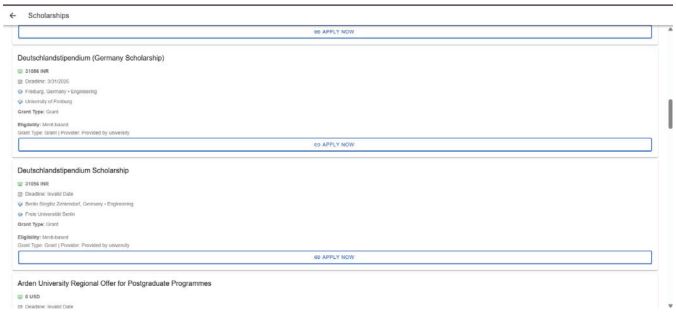

Fig. 11. (Extended) Scholarship browser displaying filtered results for Engineering students in Germany, including grant value, deadline, institution, and eligibility type.

Fig. 13. Roommate match results ranked by weighted compatibility score (97%, 82%, 80%), displaying shared interests, budget range, and housing type for each candidate.

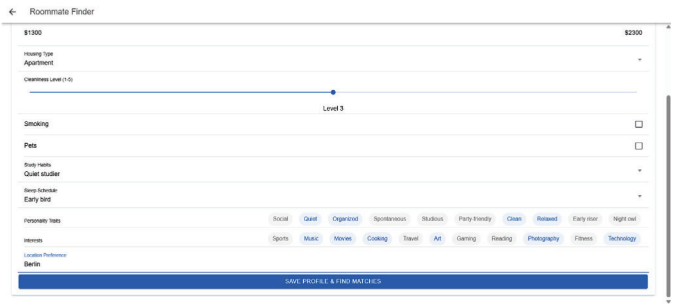

Fig. 12. Roommate profile form capturing lifestyle attributes including cleanliness level, study habits, sleep schedule, personality traits, and interests used for SBERT-augmented compatibility scoring.

References

[1] S. Yao et al., "ReAct: Synergizing Reasoning and Acting in Language Models," *arXiv preprint arXiv:2210.03629*, 2022.

[2] Q. Wu et al., "AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation," *arXiv preprint arXiv:2308.08155*, 2023.

[3] J. S. Park et al., "Generative Agents: Interactive Simulacra of Human Behavior," in *Proc. ACM UIST*, 2023.

[4] I. Chalkidis et al., "LEGAL-BERT: The Muppets Straight Out of Law School," in *Findings of EMNLP*, 2020.

[5] V. Karpukhin et al., "Dense Passage Retrieval for Open-Domain Question Answering," in *Proc. EMNLP*, 2020.

[6] P. Lewis et al., "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, pp. 9459–9474, 2020.

[7] H. Shu et al., "Deep-Learning-Powered GRU Model for Flight Ticket Fare Forecasting," *Applied Sciences*, vol. 13, no. 10, p. 6032, 2023.

[8] X. Sun et al., "Civil Airline Fare Prediction with a Multi-Attribute Dual-Stage Attention Mechanism," *Applied Intelligence*, 2021.

[9] K. Tziridis et al., "Airfare Prices Prediction Using Machine Learning Techniques," in *Proc. IEEE EAIS*, 2020.

[10] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[11] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence Embeddings Using Siamese BERT-Networks," in *Proc. EMNLP*, 2019.

[12] Y. Liu et al., "RoBERTa: A Robustly Optimized BERT Pretraining Approach," *arXiv preprint arXiv:1907.11692*, 2019.

[13] W. Fan et al., "Graph Neural Networks for Social Recommendation," in *Proc. WWW*, 2019.

[14] K. Yao et al., "CAPRI: Context-Aware Interpretable Point-of-Interest Recommendation Framework," *arXiv preprint arXiv:2306.11395*, 2022.

[15] X. He et al., "Neural Collaborative Filtering," in *Proc. WWW*, 2017.

[16] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise," in *Proc. 2nd Int. Conf. on Knowledge Discovery and Data Mining (KDD-96)*, pp. 226–231, 1996.

[17] H. Rashkin et al., "Towards Empathetic Open-domain Conversation Models: A New Benchmark and Dataset," in *Proc. ACL*, 2019.

[18] W. Shen et al., "TED: Turn Emphasis with Dialogue Feature Attention for Emotion Recognition in Conversation," *arXiv preprint arXiv:2501.01123*, 2020.

[19] P. Zhang, G. Zeng, T. Wang, and W. Lu, "TinyLlama: An Open-Source Small Language Model," *arXiv preprint arXiv:2401.02385*, 2024.

[20] J. Sweller, "Cognitive Load During Problem Solving: Effects on Learning," *Cognitive Science*, vol. 12, no. 2, pp. 257–285, 1988.

[21] International Organization for Migration, "World Migration Report 2024," IOM, Geneva, 2024. [Online]. Available: https://worldmigrationreport.iom.int

[22] M. Jones, J. Bradley, and N. Sakimura, "JSON Web Token (JWT)," IETF RFC 7519, May 2015. [Online]. Available: https://tools.ietf.org/html/rfc7519

[23] G. Molnár and Z. Szüts, "The Role of Chatbots in Formal Education," in *Proc. IEEE SACI*, 2018.

[24] P. Smutny and P. Schreiberova, "Chatbots for Learning: A Review of Educational Chatbots for the Facebook Messenger," *Computers & Education*, vol. 151, 2020.

[25] R. Santos et al., "LAWBOTS: Utilization of AI Chatbots for Legal Advising in the Philippines," in *Proc. IEEE*, 2024.

[26] N. Haidari et al., "A Proposed Cognitive Framework Model for a Student Support Chatbot in a Higher Education Institution," *Academia*, 2022.

[27] A. Dixit et al., "Personality-Aware Roommate Recommendation System," 2018.