Alexander

Your onboarding buddy

Table of Contents

Amazon PR/FAQ

Q: What is Alexander?
A: Alexander is a personal onboarding assistant designed to help new joiners by answering their questions through a chatbot. It can address various inquiries related to training (such as ethics training), specific topics (like the outcome of the last project), facilities (including booking a room or setting up a printer), work administration (such as using the vacation tool for paid time off), and company information (like dress code and company values).

Q: What purposes can Alexander be used for?
A: Alexander can be used to ask questions that enhance the onboarding experience. HR and the relevant department can configure access profiles with confidentiality levels to control the types of content new joiners can see, determining which sensitive documents they can access. HR and the department can upload data to a storage location from which Alexander retrieves information.

Q: How do I use Alexander as a new joiner?
A: Launch the pre-installed application on your desktop, log in with your enterprise credentials, and start entering prompts. The application will provide written answers and, if applicable, include references to internal documents, internal links (such as links to ethics training), and email templates. You have access to a home screen repository where you can check previously and frequently asked questions, as well as the most relevant materials for your current project, offering quick access to essential documents.

How do HR and department heads use Alexander?
A: HR and department heads can access a settings UI to configure profiles. HR can set the confidentiality level for each new joiner on a four-level scale: (1) public, (2) company internal sharing only, (3) sharing only with relevant team members, and (4) highly confidential, which cannot be shared unless there is a justified reason.

Amazon PR/FAQ

Q: What is included in the interface?
A: The interface includes a chat interaction space for entering questions and receiving answers, a home screen repository with the most relevant documents and recently requested documents for new joiners, quick links to common company policies, and a list of important contacts.

Q: What LLM is Alexander using?
A: Alexander uses a third-party large language model (LLM) for initial training with data to enhance accuracy. This approach offers higher-quality outcomes compared to off-the-shelf models.

Q: Who should use Alexander
A: New joiners in a company, department, or team who need answers to their questions—whether simple or complex—should use Alexander.

Q: How does Alexander treat your data?
A: Alexander treats data with high confidentiality. Data is not shared outside of the company's infrastructure, and the model is trained solely on company internal data.

Q: How is Alexander trained?

A: Alexander undergoes supervised initial training involving data labeling with policies, training manuals, and FAQs. Following this, a second training phase utilizes the client company's specific data. A closed development environment is set up for each company, ensuring internal data is used to refine the model without external access, and the data remains solely for internal use.

Problem Definition and Context

Onboarding efforts in the U.S. tend to focus heavily on foundational training content. The top areas covered in onboarding programs include mandatory compliance training (53%), processes and business practices (46%), technical and soft skills development (both 41%), company values and culture (39%), and training on products, services, and internal systems (ranging from 36% to 39%) (Bohne 2024a). While these areas are essential, the data shows that onboarding remains primarily content-driven. However, this process is often fragmented, overwhelming, and inefficient in many organizations, especially large enterprises.

New hires often face the challenge of navigating large volumes of information across fragmented systems like SharePoint or Confluence, frequently without clear structure or role-specific guidance. Although organizations provide documentation, it rarely adapts to the employee's level, business unit, or working model, making it difficult to access the right information at the right time. Global onboarding obstacles include ensuring a positive new hire experience, maintaining engagement, delivering personalized onboarding, and securing leadership buy-in (Bohne 2024b).

Additional challenges include automating workflows, provisioning tools, access, and measuring onboarding program effectiveness (Bohne, 2024b). These findings suggest that many companies struggle not only with delivering content but also with creating structured, measurable, and personalized onboarding processes. Deloitte (2023) found that workers who experience highly effective onboarding are 18 times more likely to feel committed to their company, and 69% are more likely to stay for at least three years.

In response, Deloitte designed onboarding blueprints for distinct employee personas, early career hires, experienced hires, hiring managers, and onboarding buddies, mapping journeys over 12 months with cohort-based milestones tailored to role, hybrid setup, and organizational hierarchy. This approach delivered contextualized, level-specific guidance to enhance retention and accelerate productivity.

Without such frameworks, new employees often lack the clarity and support needed to confidently navigate their roles from the start. Without accessible and well-structured onboarding support, new employees are left to navigate complex processes on their own, significantly impacting their early productivity and overall engagement. In parallel, HR and senior staff are burdened with recurring administrative and explanatory tasks. Instead, time that could be spent on strategic initiatives is consumed by answering routine queries (Bohne 2024b; Deloitte 2023).

Key onboarding challenges faced by HR professionals worldwide include inadequate monitoring of new employees (49%), inconsistencies in how onboarding is applied across the organization (48%), a lack of clarity regarding responsibility for onboarding tasks (43%), and difficulty in measuring the success and effectiveness of onboarding programs (43%) (Bohne 2022a). These issues reflect a broader struggle to establish coherent, accountable, and results-driven onboarding practices across departments and management levels.

The duration of onboarding programs further reflects this challenge. In North America, 38% of onboarding programs are limited to the first day or week, and only 39% extend to 1–3 months. Longer programs are rare, with 16% lasting 3–6 months and fewer than 7% going beyond that (Bohne 2024c). These numbers highlight a trend toward short-term onboarding despite growing expectations that onboarding should support long-term employee engagement, cultural integration, and productivity ramp-up.

Moreover, expectations for what onboarding should achieve are growing. 86% of HR professionals expect onboarding to help new hires feel at ease in the company, 74% want to accelerate time to contribution, 53% focus on improving retention of new employees, and 41% aim to reduce costs (Bohne 2022b).

These expectations reflect a strategic shift from onboarding as a basic administrative necessity toward a holistic tool for cultural integration, engagement, and long-term organizational impact in the shortest amount of time possible. Underlying our solution is a key assumption: in most medium to large enterprises, the most important business information is already documented. The issue is not the absence of knowledge but its discoverability, context, and delivery. Employees often don't know what to search for, where to find it, or how it relates to their specific roles. This lack of context-awareness hinders their ability to perform effectively from the outset.

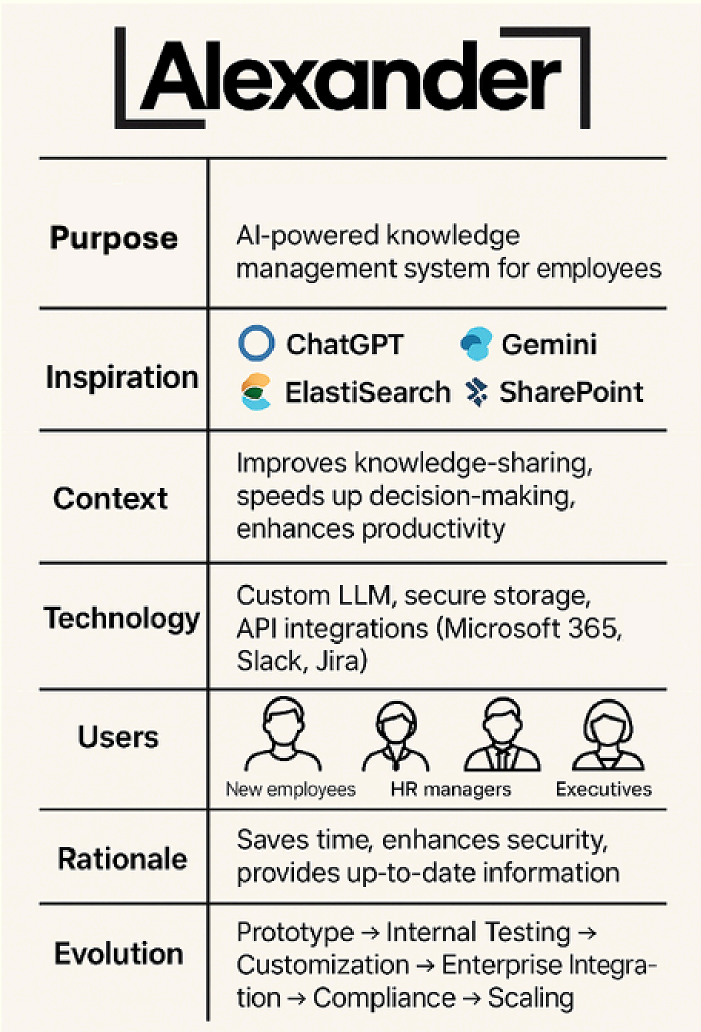
Alexander

Purpose	AI-powered knowledge management system for employees
Inspiration	ChatGPT, Gemini, ElastiSearch, SharePoint
Context	Improves knowledge-sharing, speeds up decision-making, enhances productivity
Technology	Custom LLM, secure storage, API integrations (Microsoft 365, Slack, Jira)
Users	New employees, HR managers, Executives
Rationale	Saves time, enhances security, provides up-to-date information
Evolution	Prototype → Internal Testing → Customization → Enterprise Integration → Compliance → Scaling

FIG 1. PICTURE for Alexander

Our Solution

We propose an AI-based application, “Alexander”, to assist new joiners in all questions related to their onboarding. We intend a standalone application which will process prompts primarily using company internal data, providing the new joiners helpful answers or contact points for further guidance.

Alexander will be connected through an API to other applications of the company such as the knowledge management system to grasp unstructured as well as structured data.

The connection to other applications enables agentic features allowing the user to further act on the initial prompt. For example, asking Alexander how to write an email in the appropriate tone to a specific client would open a new Microsoft Outlook email with a suitable template. The application consists of a home screen with the most relevant document suggestions (e.g., ethics policies), chat functionality, quick links to key documents and important contacts.

The solution has a role-based access control (RBAC) ensuring that employees can access only permitted documents based on a need-to-know principle. HR-managed user profiles will control access and each department (e.g. marketing) can specify accessible documents (e.g., last year's performance documents).

Governance of Alexander

- Our IT team will collaborate with the client’s IT department for the initial application setup.
- The client's IT department will provide the necessary data sources (with the materials to be processed) for Alexander.
- The client company will also provide initial training data for supervised learning purposes.

UI of Alexander

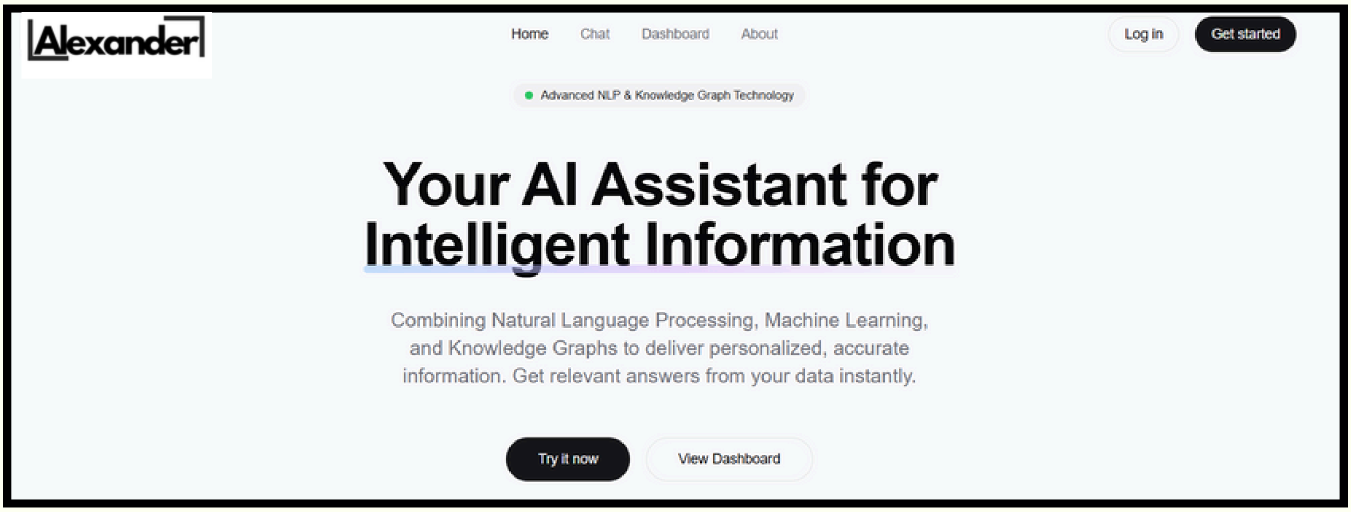

FIG 2. Alexander Welcome Screen

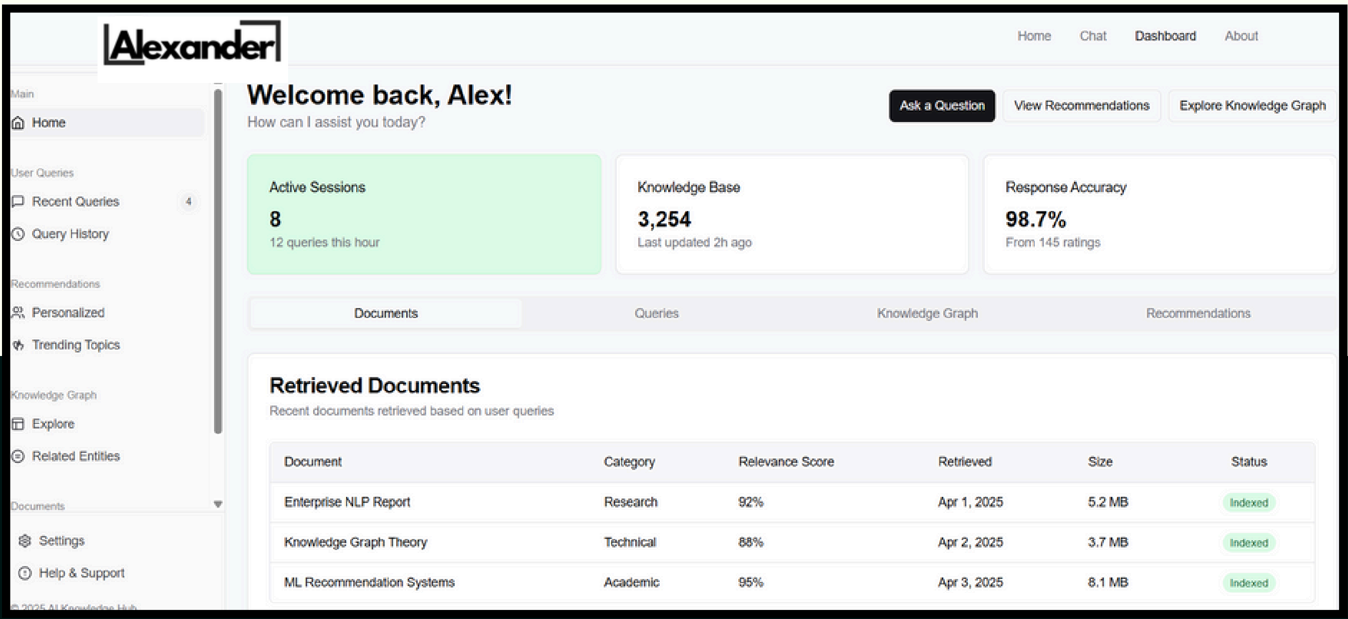

FIG 3. HR Dashboard for Managing Alexander

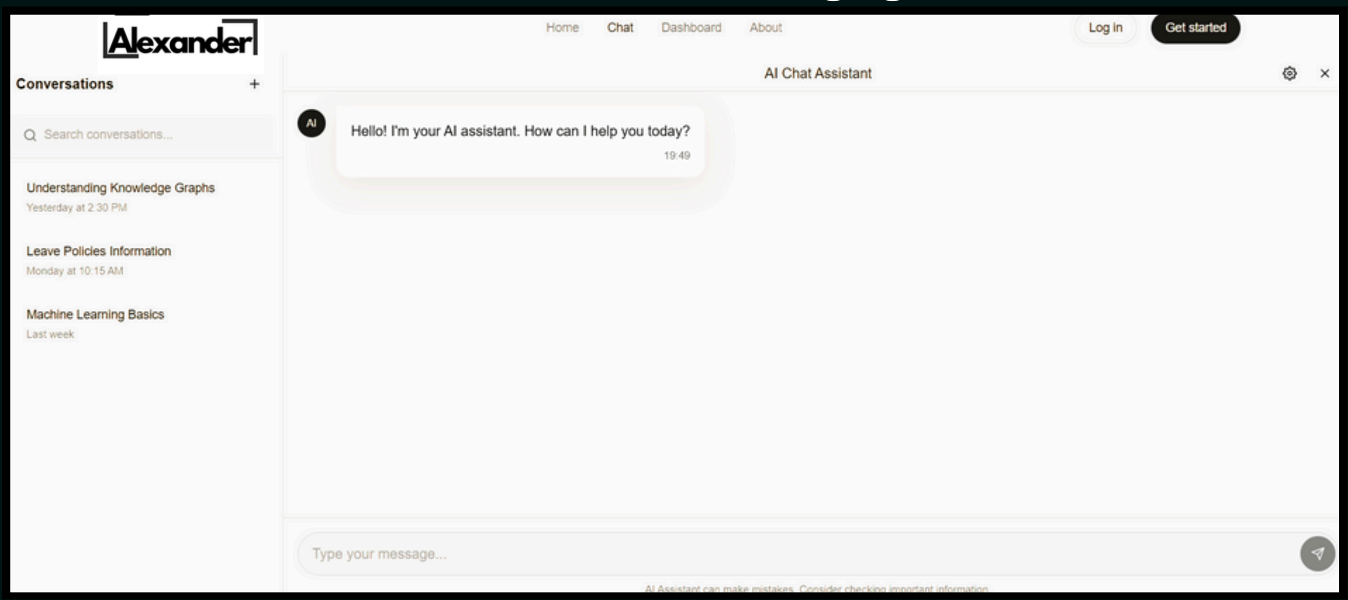

FIG 4. Alexanders Chat Interface

Value proposition

Alexander addresses the problem by providing the new joiner with a virtual onboarding assistant to whom they can pose questions of simple nature or complex nature. The chatbot functionality can answer simple and complex questions, effectively replacing the time senior managers typically spend on these queries. By providing accurate, detailed answers and suggesting the next steps (like drafting emails), the assistant reduces the onboarding workload for managers. The 24/7 availability ensures immediate answers, eliminating waiting times and accelerating the onboarding process. Moreover, new joiners can ask questions they might have previously hesitated to ask, leading to better knowledge acquisition and identification of knowledge gaps.

Alexander would consistently provide high-quality answers based on the company’s document knowledge, removing the circumstance that answer quality might differ from employee to employee. Compared to existing competitors, our main differentiation would be the training on proprietary data to ensure highly accurate prediction rates. Firms such as Blackrock have successfully developed models which are more accurate than commonly available generic off-the-shelf models (Blackrock, 2024). Our service offering would be a proprietary model for every company.

Over time, the application can determine the most asked questions and based on that provide the HR department and other departments insights on which questions are frequent. Those departments can then integrate those insights and improve the overall service. Continuous improvement will also be facilitated through reinforcement learning, enabling the application to deliver increasingly better results.

Use case scenarios

Onboarding for employees	New department tasks	Promotion for new roles
Providing new employees with basic information about the company during the first few days/weeks in the company. Helping them to adapt faster.	Helps the HR department to updates tasks and notifications for employees	Support for change of team/department. Providing newly promoted employees new tasks and file access.

FIG 5. Use Case Scenarios Diagram

Target group

The target group users are white-collar workers in professional, managerial, and administrative roles. We assume that these individuals deal with a high load of information during their initial weeks, leading to numerous questions. New joiners might hesitate to ask too many questions due to the new environment and desire to make a strong first impression, potentially fearing that frequent inquiries could reflect negatively on their qualifications.

Alexander's benefits can be fully experienced on a desktop operating system, making new office-based employees an ideal target group due to existing infrastructure (desktops, network access). We are targeting high-growth companies with a higher number of new onboardings as well as companies with a high turnover (e.g. those employing many contractors or consultants) where Alexander can provide added value.

Companies with high personnel costs such as banks would benefit strongly, as it is in their interest to conduct a timely onboarding so that new joiners, particularly senior employees, can become productive quickly.

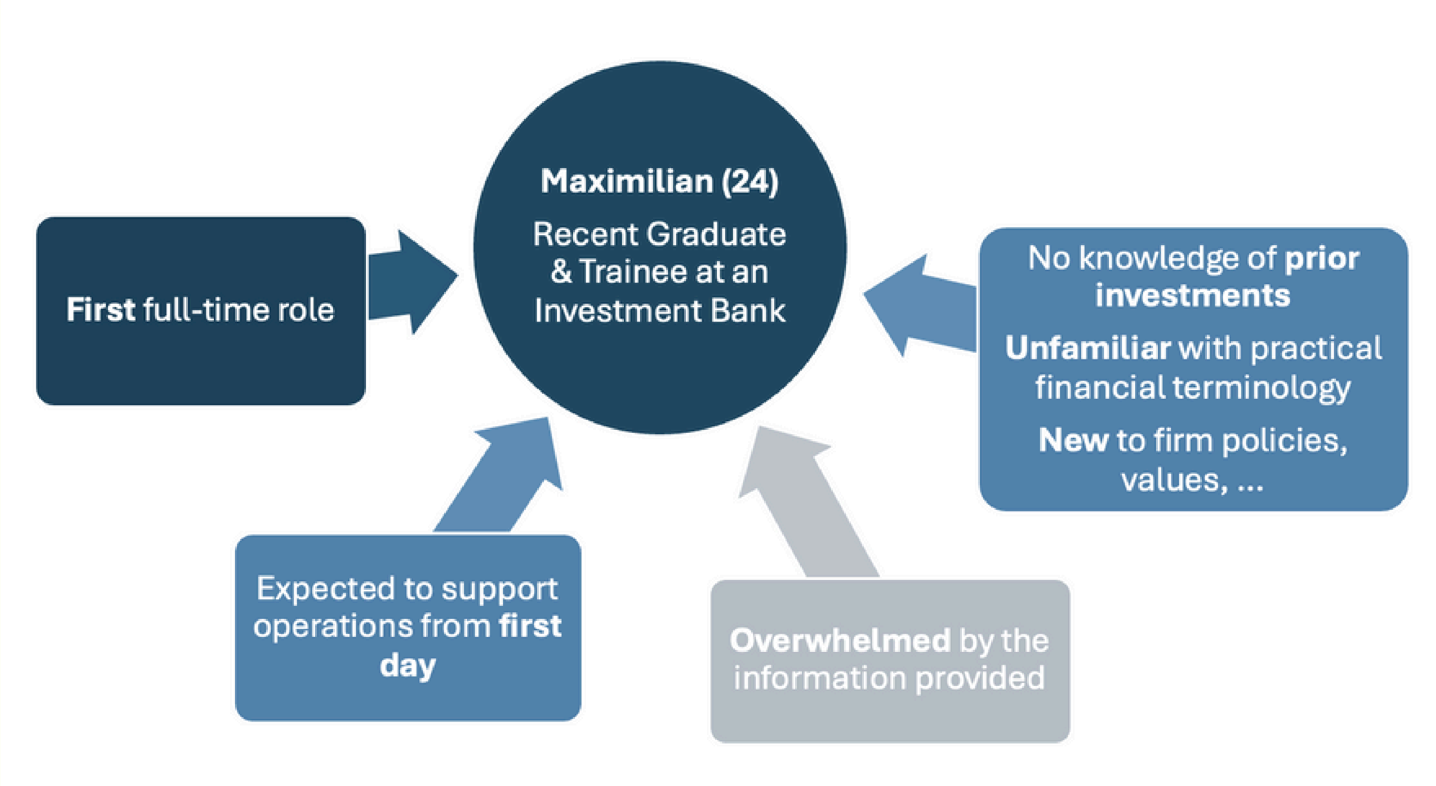

FIG 6. Situation and Concerns of the example user Maximilian

User Journey

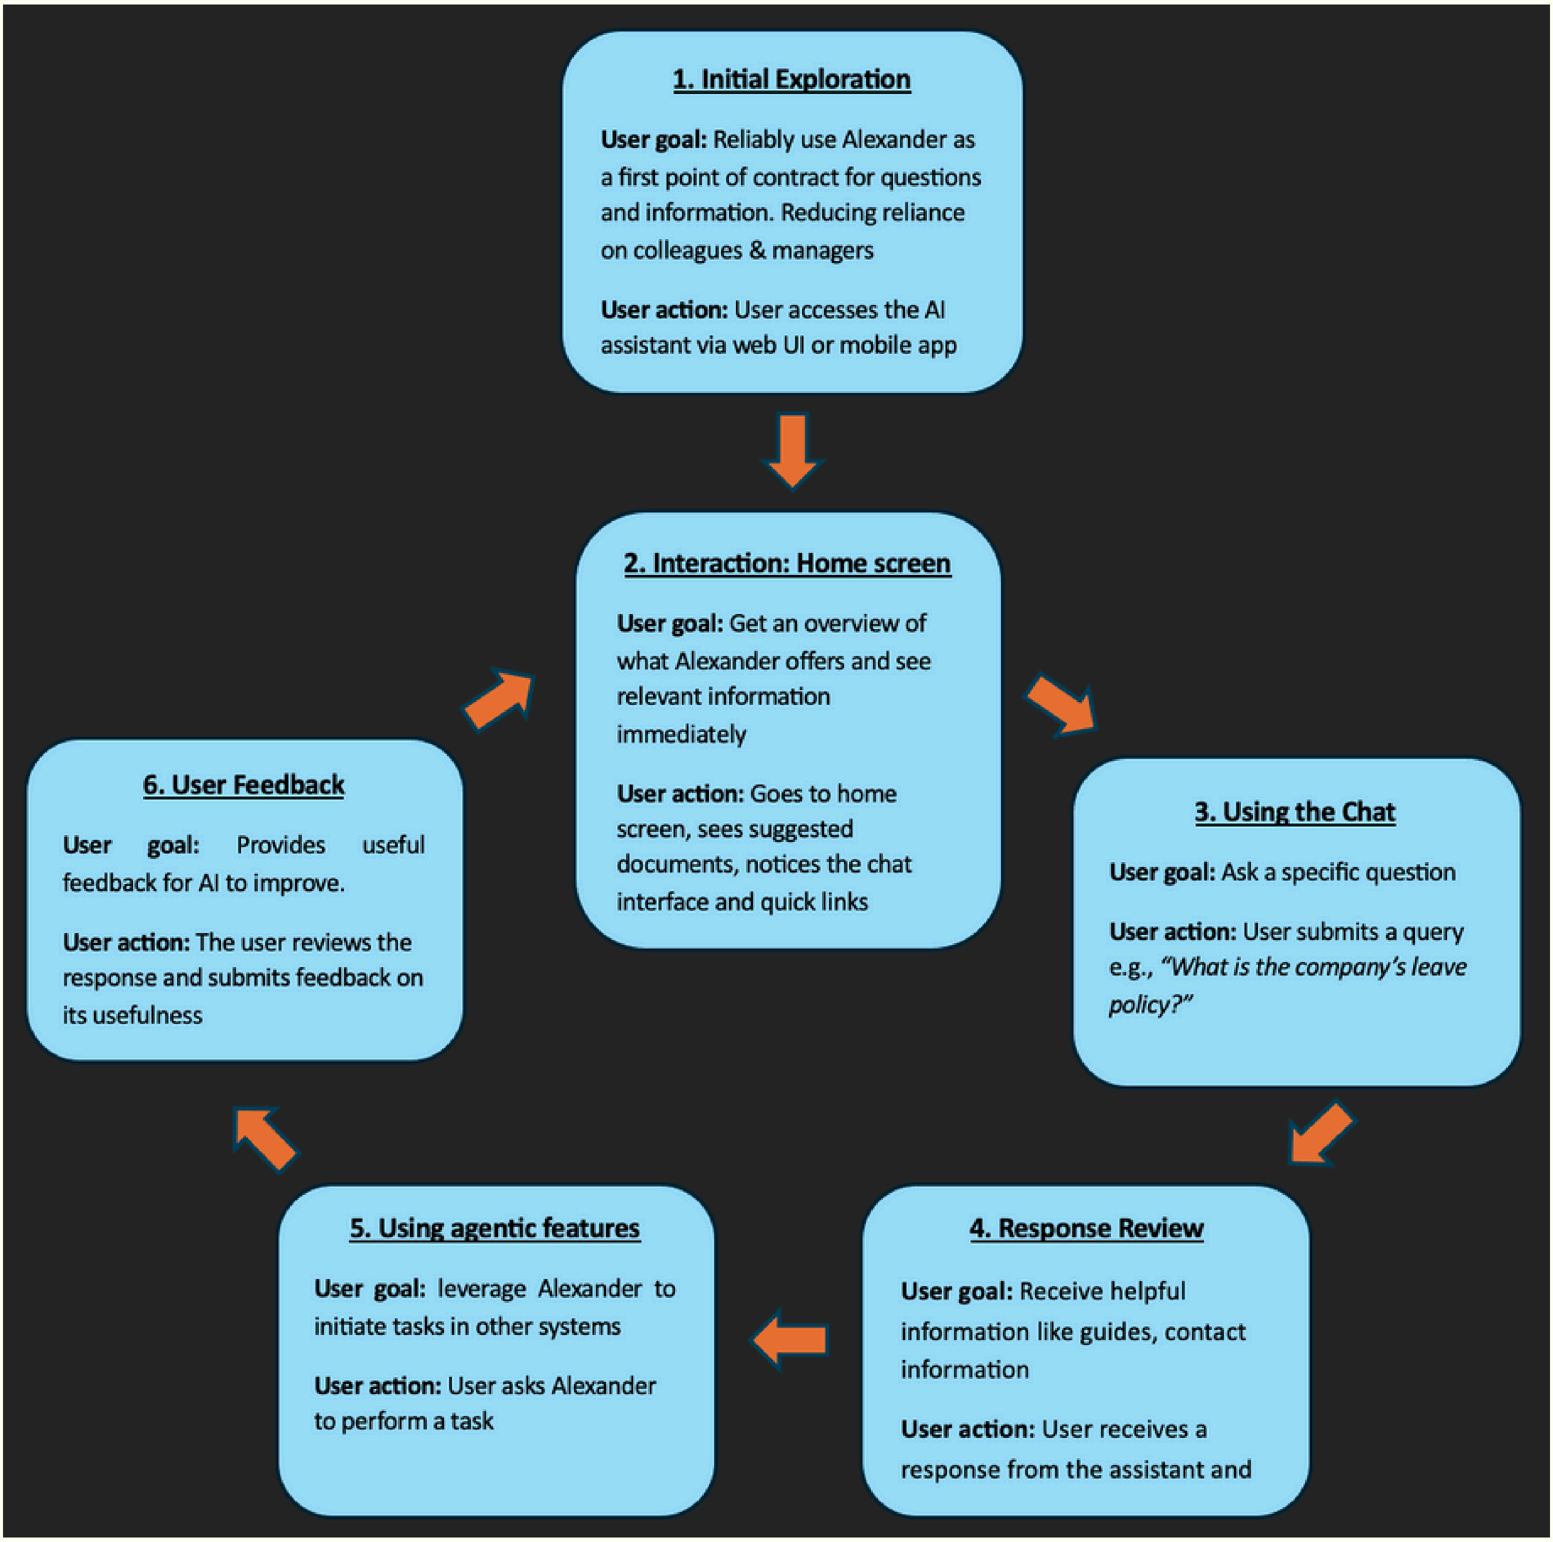

FIG 7. User journey diagram for Alexander

Tech Overview

Overview of the Technical Solution

This technical solution involves creating a chatbot that helps new employees learn about company policies by leveraging Natural Language Processing (NLP), Machine Learning for Recommendations, and Knowledge Graphs. The company's documentation is ingested into the system via an API and stored on our servers for training. The model absorbs the data and self-annotates the company information based on its initial training. It then compiles this data and accesses it via input data from the central database.

The main interaction component is the chatbot interface managed by the NLP, which connects to the knowledge graphs and is continuously improved by machine learning tools integrated into its overall digital structure. Each component interacts seamlessly to provide accurate, relevant, and personalized information to users. Below is a detailed analysis of each section and how they contribute to the overall system.

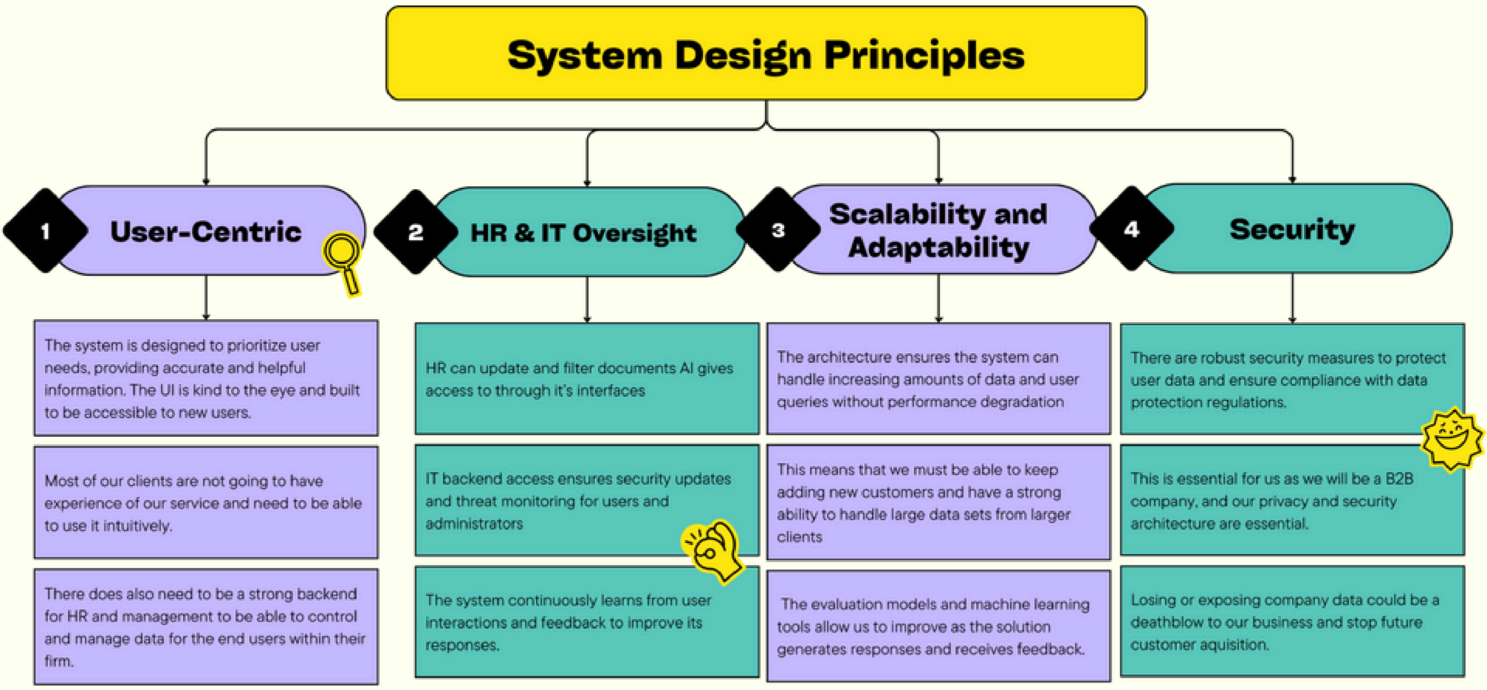

FIG 8. Alexanders System Design Principles

Models Interactions and Output Generation

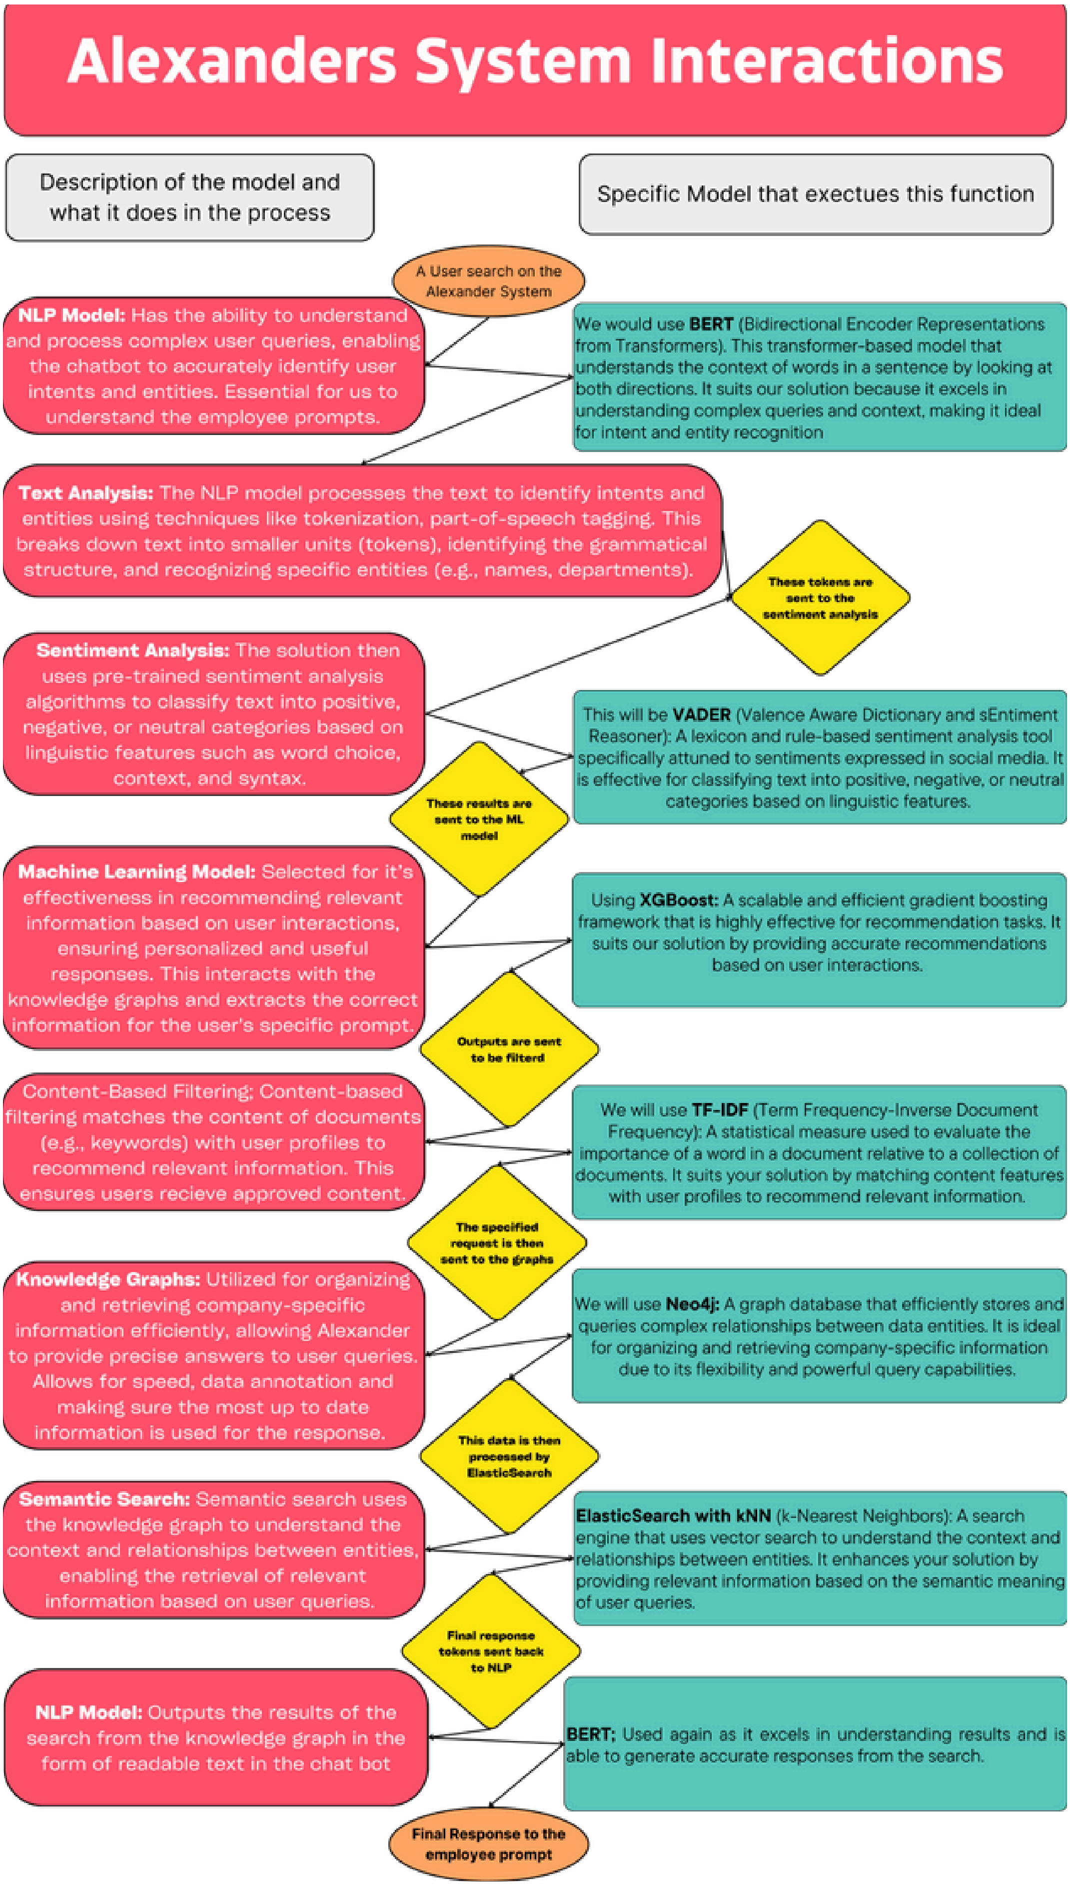

The data is delivered through either an API, Manual interface in the UI, or through uploading and image or Manifest/Excel to the techinical solution

FIG 9. Alexanders System and Models interaction Diagram) (Chandra, 2024), (ElasticSearch, 2025) (Google Dev, 2025) (Kumar, 2024), (Murel, 2024), (Neo4j, 2025).

Knowledge Graphs

A knowledge graph are structured forms of information that connects data points through relationships. It organizes data into nodes (representing entities like people, places, or things) and edges (representing the relationships between these entities). A good way of understanding this concept is of an advanced labelling technique that allows for quick searching of large data sets. This structure allows for more intuitive and efficient querying and analysis of complex data.

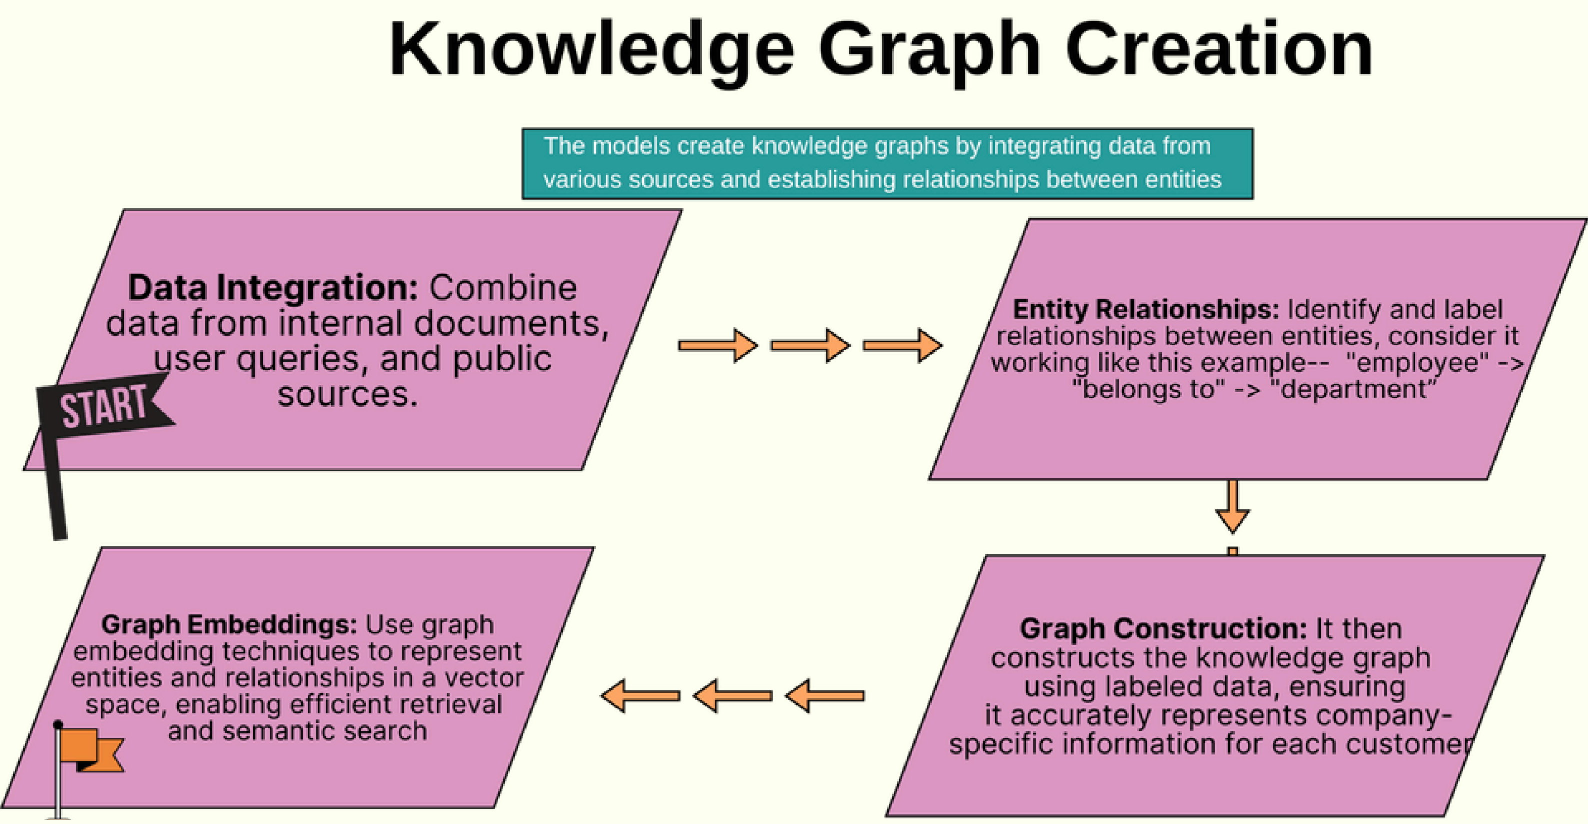

FIG 10. - Data Sources and Data Preprocessing

Training

Data Annotation

Data annotation will be crucial for training the models effectively in a two-step process. First manual annotation using human annotators label data based on predefined categories to create a high-quality training set. Then on new datasets the model can do automated annotation: As the model learns, it begins to self-annotate new data, reducing the need for manual intervention.

The Ingression Models

This process will be done via supervised learning on this initial Data Set we use labelled datasets to train the model. This includes company-specific data such as policies, training manuals, and FAQs. The initial data set helps the model understand the structure and content of company information. Buy this or from our first partner, probably easier to just buy it outright from a bankrupt company! Then the model is trained to self-annotate new data by recognizing patterns and relationships within the initial data set. This involves labelling intents and entities automatically based on learned patterns. Human in the loop training will be required, and we will get a third party to do this for us.

Transfer learning for the main line models will use pre-trained models leveraging pre-trained models for the NLP like BERT or GPT that have been trained on large datasets. These models are then fine-tuned on company-specific data to adapt them to the unique requirements of the company. Over time we will have to do fine-tuning to adjust the pre-trained models using the company's data to improve accuracy and relevance in responses.

Onto the Training Timeline

The Training Timeline

Learning paths guide the progression of training, ensuring the system learns effectively and adapts to new information: we will have the initial training focusing on understanding company data and self-annotating information. This phase involves intensive training with the initial data set to establish a strong foundation. Then ongoing training with regular updates and fine-tuning based on new data and user interactions. This phase ensures the model remains accurate and relevant. Finally, we have a state of advanced training incorporating feedback mechanisms and continuous evaluation to refine the model further. This phase focuses on iterative improvement and adaptation.

FIG 11. The Solutions Training Timeline

Data Types & Transfers

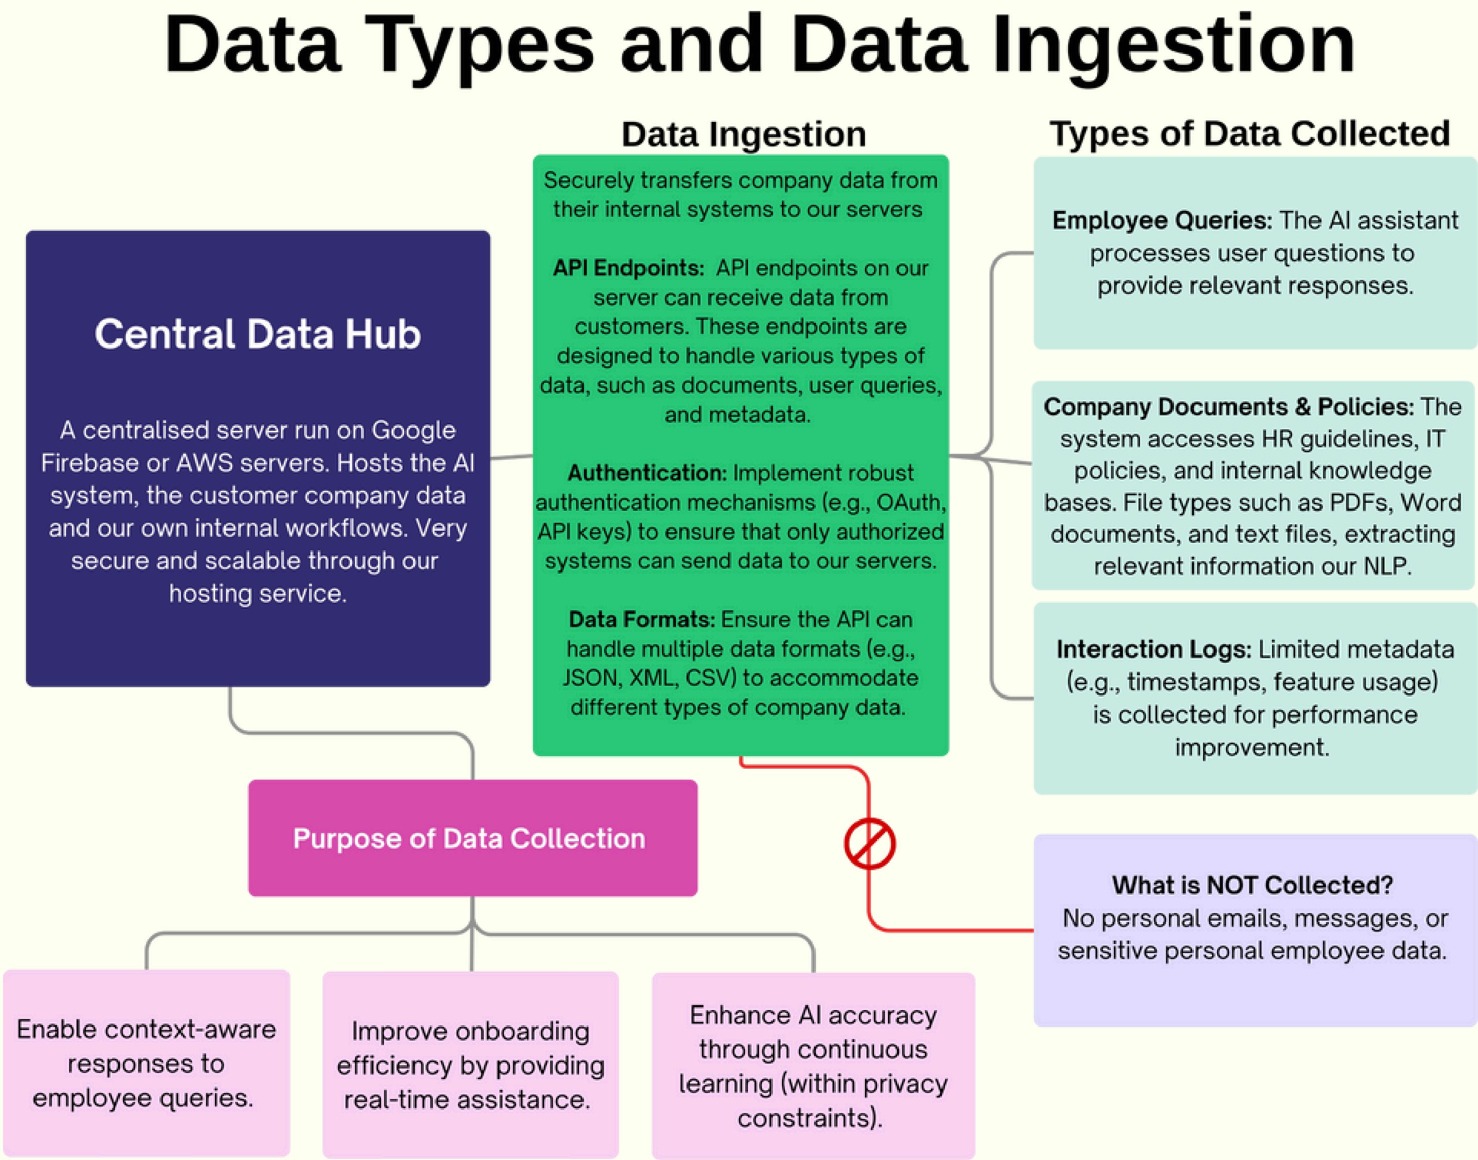

FIG 12. Data Types and Data Ingestions

The Data Flow

The model collects data from muliple sources to get better accurancy in the prompt response outputs. This data is collected and delivered to the central proccessing hub, where the seqeunce of algorithms process inputs from the API's and generate the ouput throught the NLP BERT into the chat.

UML Diagrams

Use Case Diagram

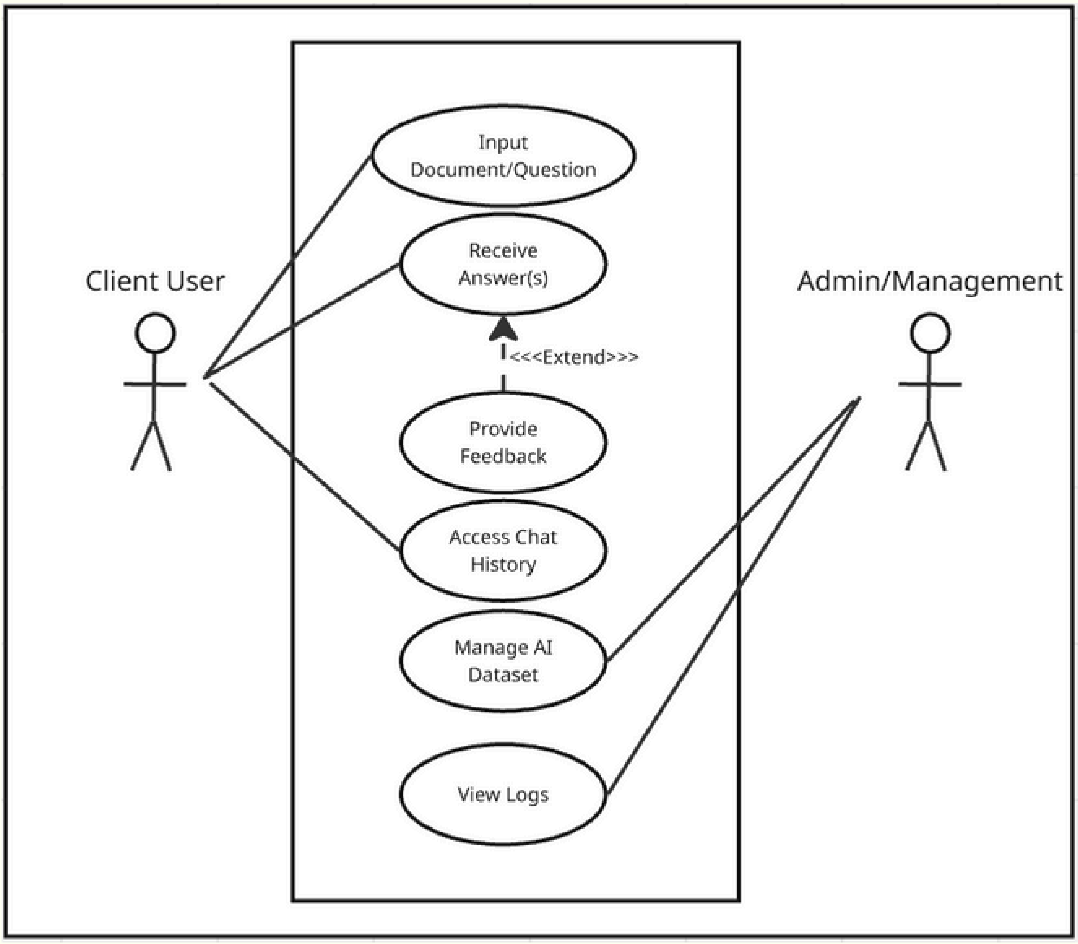

FIG 13. Use Case Diagram

In this use case diagram above illustrates the key interactions between the two main actors of the onboarding assistant system, the Client User and the Administrator/Management. The client user represents a new employee or any staff member using the chat-based assistant.

Their primary interactions include inputting documentation or questions into the chatbot, receiving answers, providing feedback and accessing chat history. On the other hand, the Administrator/Management represents HR personnel or system administrators. Their interactions are focused towards maintaining and optimizing the assistant's backend by managing the AI dataset and monitoring interactions.

Class Diagrams

Class Model

The class diagram above outlines the key data entities and their relationships within the onboarding assistant system. It focuses on how different components interact to deliver a streamlined, AI-powered onboarding experience for new employees in various firms. There are users (new employees), documents/questions, summaries, and feedback. Document/Question represents the input provided by users, summaries are the AI-generated responses based on the input, and feedback is the users' opinions about the generated responses.

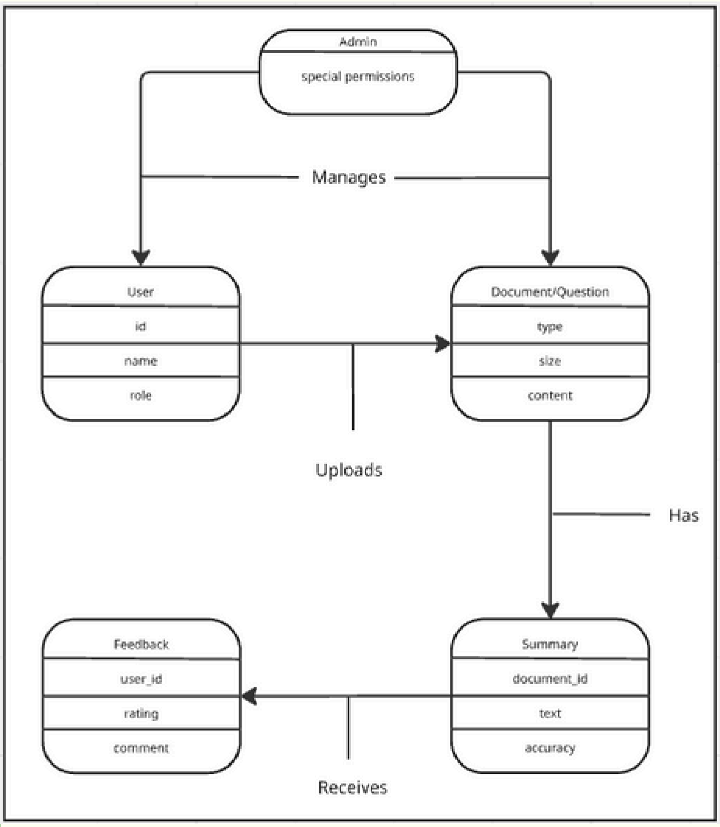

FIG 14. Class Diagram

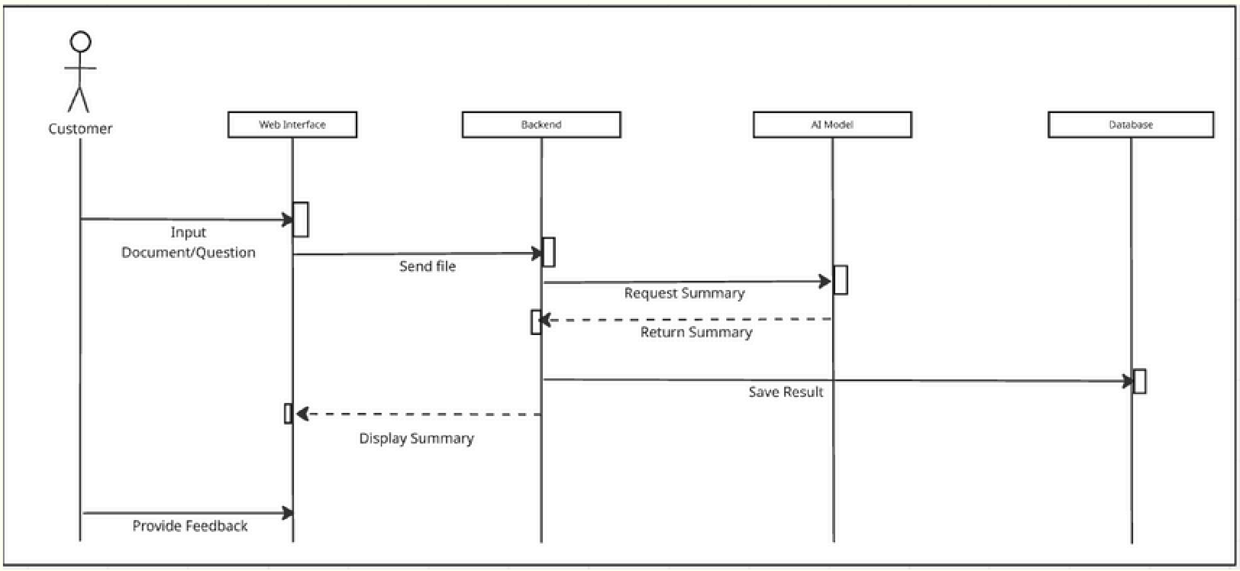

FIG 15. Sequence Diagram

Sequence Diagram

This sequence diagram captures the end-to-end interaction flow between a new employee and the onboarding assistant system, showing how a user query or document is processed and responded to through the system's layered architecture. It is a diagram composed of the actors & components within the system, such as customers and the web interface, among others, while illustrating the process flow.

Implementation - Structure and Platform Interaction

Firstly, addressing the system architecture of the client and how Alexander fits into this issue. The client's system will involve local document storage and cloud-based operations. Alexander will act as a SaaS layer. Clients will upload their files or questions via web interface or API; the data will go to a cloud-based backend, such as AWS S3 + Lambda or EC2, it will be processed by AI, and then the results will be returned.

The platform selection will be a primary interface that is fully web based and accessible via browsers with a responsive design for both desktop and mobile. The tech stacks mention using Langchain, OpenAI API, AWS, Python/Flask or something similar. Meanwhile, the future app integrations can be discussed as a scalable vision.

Additionally, a clean dashboard UI for employees with drag-and-drop upload, with a chat interface, the summary generation or reply includes feedback buttons that determine whether the answer was helpful or not. The backend can be accessed by management or the administrative department, where they can access data logs, user activity, and model feedback through an admin portal.

Finally, detailing the integration strategies. An initial onboarding process will be conducted through demos and firm outreach. To encourage further adoption, trials or API access could be offered. For firms operating within an existing cloud infrastructure, a seamless plug-in style could be implemented. There will be an ongoing use through retraining the AI with firm-specific documentation.

Definition of Requirements

Considering modern offices, it is important to address the desktop requirements to be able to run the program. Hence, a desktop that can run a modern browser (Chrome, Firefox, Edge), a stable internet connection and no heavy local resources will be required to run the program.

Storage and Safety

Regarding storage, on-premises is challenging to implement; therefore, Alexander opts for the Cloud, since the use of AWS services will be instrumental. Such as using S3 for storage, Lambda/EC2 for processing and DynamoDB/RDS for metadata storage. There is the advantage and benefit of easy scalability through serverless architecture. Data encryption (AWS KMS) and role-based access control are the central points of the program.

Data Access and Processing

The data will be stored in AWS as the files are uploaded to the program, while relying on real-time processing on the cloud between the shared devices. Hence, the model, fine-tuned locally, processes and returns a reply in seconds. The program will support dynamic scaling based on the file volume.

Manual Entries and Learning Inputs

Concerning manual data entry, users can manually upload non-digital or annotated notes, and the interface will further support this by allowing tagging and categorization of non-documented information. A feedback loop will be in place, allowing users to provide insight via feedback. Additionally, offering an option to input case-specific notes or comments will help enrich context.

On the other hand, for data validation, using version control and timestamping to avoid outdated documents and keep data as current and reliable as possible. For retraining purposes, only authorized documents will be allowed to be used. Finally, feedback will be flagged and verified before being used in the model's learning process.

Evaluation of the app in the real world

To assess the real-world efficacy of Alexander, a comprehensive evaluation strategy has been developed to examine its usability, functional performance, and broader organisational impact. Designed to support new employees as they navigate complex information systems and social environments in the early stages of their employment, Alexander must be tested for its technical accuracy and ability to address human and organisational needs.

The overarching goal of the evaluation is to determine how effectively Alexander reduces onboarding friction, enhances access to critical knowledge, and alleviates the burden placed on HR teams and senior staff. Specifically, we aim to measure whether the assistant is intuitive and engaging for users, whether it meaningfully reduces the time required for the new employees to reach productivity and whether it delivers relevant, accurate responses that are personalized to users' roles and departments. Importantly, we also seek to understand the social dynamics of adoption. Whether users trust the system, feel more confident asking questions through AI than human channels, and whether HR and management perceive it as an asset or an added complexity.

To achieve these aims, we adopt a mixed-methods approach that balances quantitative rigour with qualitative insight. One strand of the evaluation involves a comparative study, where two groups of new employees, one using Alexander and one onboarded through conventional methods, are assessed across metrics such as onboarding duration, number of HR interactions, and time-to-productivity, as judged by their direct supervisors. These forms of A/B testing are widely recognized for their reliability in isolating the impact of new technologies in workplace settings and will offer a baseline understanding of Alexander's measurable benefits.n parallel, we will conduct task-based usability testing with a selected group of participants. These individuals will be asked to complete daily onboarding tasks using Alexander, such as locating internal policies or finding departmental contact points.

Evaluation of the app in the real world

ITheir interaction will be monitored for task success rate, efficiency, and user satisfaction, with follow-up surveys using to provide a standardized perceive usability assessment. This will help identify design strengths and usability bottlenecks, ensuring the assistant is accessible to users with varying levels of technical proficiency.

Complementing these direct interactions, we will also analyse anonymized system usage logs to uncover behavioural patterns. This includes monitoring the frequency and timing of queries, the type of documents accessed and the success rate of information retrieval. The assistant's feedback loop, where users can rate the relevance or helpfulness of the answer, will be critical for identifying response quality and informing future model refinement. To enrich our understanding of Alexander's organizational value, we will conduct semi-structured interviews with HR professionals and team managers.

These stakeholders can offer a critical perspective on how the tool integrates into existing workflow to reduce repetitive tasks. Their feedback will be essential in refining the role-based access model and ensuring compliance with internal policy standards. Finally, all new joiners participating in the trial deployment will be invited to complete a survey reflecting on their onboarding experience with Alexander. This will gauge their satisfaction, perceived confidence and willingness to rely on the assistant over traditional channels.

Evaluation responsibilities will be shared across the project team, HR, and departmental leadership. The technical team will lead the analytics and usability testing. HR will coordinate participant recruitment and data collection; managers will provide performance-based feedback on employee progression. This collaborative approach ensures that Alexander is functional, secure and genuinely valuable for addressing modern organizations' onboarding challenges.

Future use cases and design reflections

While Alexander is designed with the purpose of helping employees to better adapt to their new work environment, potential future expansions can be made for the app to better optimise in different work settings. Research conducted by the Office for National Statistics (ONS) found that after the covid pandemic, more than 50% of people living in London are working remotely outside of their office (ONS, 2020), indicating that remote work has now become a popular work culture.

Similarly, while Alexander is intended to help white collar workers to adapt faster during their first few weeks working at a traditional office environment, non-white collar workers could utilise AI assistance during their onboarding phase. For instance, employees who have a peripatetic work schedule (e.g. traveling salespeople, care workers, plumbers...) are required to travel frequently from various locations for their work, this makes the onboarding process more complicated and obscure for new joiners as they are required to be mobile without any colleagues or superiors present at scene to guide them.

In addition to the design of Alexander's chatbot interface, the current prototype consists of text-based communication between the machine and the user. While this may be beneficial for users that are familiar with texting via a keyboard or touchscreen interface, it is a challenging medium for disabled users to navigate. To provide a more user-friendly option for users with physical disabilities, audio-based communication can be implemented as an alternative communication method between a human agent and the AI chatbot, allowing the system to register via speech recognition application and respond with audio messages in response to the user's queries.

Lastly, as an AI assistant tool designed for improving workplace efficiency and organisation management, it is imperative to keep in mind how cultures may be an important factor affecting organisational practices. The current version of Alexander targets users in the Western hemisphere with regions that are predominantly English language dominated, this gives the opportunity for a potential examination into how Alexander could adapt to various other cultures.

References

- BlackRock. (2024). How AI is transforming investing. Retrieved from https://www.blackrock.com/us/individual/insights/ai-investing
- Bohne, R., (2022a). Leading challenges with onboarding according to HR worldwide in 2018. Statista. https://www.statista.com/statistics/950216/onboarding-challenges-hr-worldwide/
- Bohne, R., (2022b). Expectations of HR practitioners worldwide from people onboarding in 2018. Statista. https://www.statista.com/statistics/950156/hr-expectations-onboarding-worldwide/
- Bohne, R., (2024a). Key areas of onboarding training in the U.S. 2018. Statista. https://www.statista.com/statistics/1026265/onboarding-delivery-areas-united-states/
- Bohne, R., (2024b). Biggest obstacles for onboarding programs 2018, by region. Statista. https://www.statista.com/statistics/951349/onboarding-program-obstacles-worldwide-region/
- Bohne, R., (2024c). Typical length of onboarding programs worldwide in 2018, by region. Statista. https://www.statista.com/statistics/951359/onboarding-program-typical-length-worldwide-region/
- Chandra, R. (2024). 9 Machine Learning Algorithms For Recommendation Engines. https://insights.daffodilsw.com/blog/machine-learning-algorithms-for-recommendation-engines
- ElasticSearch. (2025). What is Semantic Search? | A Comprehensive Semantic Search Guide. https://www.elastic.co/what-is/semantic-search
- Google Dev. (2025). Content-based filtering | Machine Learning. Google for Developers. https://developers.google.com/machine-learning/recommendation/content-based/basics
- Kumar, N. (2024). Top 10 AI Tools for NLP: Enhancing Text Analysis. Analytics Insight. https://www.analyticsinsight.net/artificial-intelligence/top-10-ai-tools-for-nlp-enhancing-text-analysis
- Murel, J. (2024). What is collaborative filtering? | IBM. https://www.ibm.com/think/topics/collaborative-filtering
- Neo4j. (2025). Knowledge Graph. Graph Database & Analytics. https://neo4j.com/use-cases/knowledge-graph/
- Office for National Statistics. (2020). Coronavirus and Homeworking in the UK: April 2020. https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/bulletins/coronavirusandhomeworkingintheuk/april2020
- Ramlochan, S. (2024). Generate Knowledge Graphs for Complex Interactions. Prompt Engineering Institute. https://promptengineering.org/knowledge-graphs-in-ai-conversational-models/
- Reyes, V., Bordeaux, C., Sivak, M., Mathur, M. (2023). Transforming the new hire onboarding experience. Deloitte. https://www2.deloitte.com/us/en/blog/human-capital-blog/2023/new-hire-onboarding.html
- Smith, G. (2025). Top 10 NLP Models (Natural Language Processing) You Should Know. https://www.theknowledgeacademy.com/blog/nlp-models/