Asian Journal of Social Science Studies; Vol. 10, No. 2; 2025
ISSN 2424-8517 E-ISSN 2424-9041
Published by July Press

Architecting Secure and Scalable AI Chatbot for Diversity Data Using Fine Tuned LLMs With RAG Framework and Lambda Functions

Nicole Lee[1], Mary Legrand[1], Sravan Sriram[1] & Rezza Moieni[1]

[1] Diversity Atlas, Collingwood, Australia

Correspondence: Nicole Lee, Diversity Atlas, 2/273-277 Wellington St, Collingwood, VIC 3066, Australia. E-mail: nicole.lee@diversityatlas.io

Received: April 3, 2025 Accepted: May 19, 2025 Online Published: May 30, 2025

doi:10.20849/ajsss.v10i2.1507 URL: https://doi.org/10.20849/ajsss.v10i2.1507

Abstract

In this paper, the authors present a web-based AI chatbot architecture development to understand the challenges of building an intuitive virtual agent capable of extracting insights from diversity data. The developed chatbot combines an information retrieval system with a fine-tuned Large Language Models (LLMs) that is specialised in logic calculations to generate responses. This approach also highlights the importance of balancing diversity's data security with the users' desire for a conversational experience. Unlike modern open-domain AI agents that interpret the data and generate responses with their own agency, the developed chatbot provides more structured guidance in conversations, addressing security and privacy challenges associated with processing personally identifiable information, as well as cost and performance issues.

Keywords: artificial intelligence (AI), chatbot architecture, natural language processing (NLP), large language models (LLMs), human-computer interaction (HCI), diversity data analysis

1. Introduction

Over the last decade, there has been significant research on the role and development of chatbots. Studies have been conducted in various industries, including healthcare (Madhu et al., 2017; Su et al., 2017), travel (Argal et al., 2018), education (Benotti et al., 2014; Lee et al., 2023), and finance (Câmpeanu et al., 2023; Madasamy & Aquilanz, 2023). However, there is a lack of studies on using inclusive language in chatbot applications. Communications, including those of chatbots, are also intercultural and in needs of a robust study of using inclusive language.

Organisations are increasingly using third-party data collectors to gather sensitive Diversity and Inclusion (D&I) data from their employees. This data includes various aspects of both cultural and demographic diversity, such as age, gender, country of birth, language, worldview or religion, cultural identity, ancestry, disability, sexual orientation, and education level. While diversity data has become one of the most valuable resources in D&I industry, it presents a significant opportunity for diversity data collectors due to the need of organisations for access to demographic and cultural data from their employees.

Nonetheless, the challenges with using AI chatbot with diversity data involves a multitude of implications and challenges, including:

- *Data Privacy and Security*: The protection of diversity data such as gender, sexual orientation, country of birth of employees from unauthorised access is essential to ensure data confidentiality and integrity. The data may be unintentionally learned and generated by large language model (LLM) that process large amount of data (Sebastian, 2023). Therefore, legal considerations will play a role in ensuring the system's development complies with regulations. The implementation of AI chatbot should therefore employ information security strategies such as data encryption and segregation as well as audit logs to preserve privacy and data protection.

- *Performance:* In taking intersectionality into consideration, the chatbot should effectively analyse data by intersecting various categories, including ethnicity, gender, and disability and provide insights while ensuring that it does not perpetuate biases against any cultural group. Also, while conversational behaviours such as leading discussions and taking turns are not vital for this chatbot—since it functions more like a data search engine than a conversational agent—it should still be capable of understanding incomplete utterances from users and generating accurate and clear responses for user to make informed sense of data.

- *Cost:* Implement a cost-effective architecture is essential for small-scale companies. To balance performance and cost, rather than adopting a ready-to-go chatbot like Amazon Q, which is more expensive, small businesses may have to develop the architecture from scratch using various tools to reduce costs. In this case study, the authors used couple of AWS tools including Bedrock, Lex, Lambda and S3 Bucket to build the most cost-effective architecture. The chatbot used aggregated data from the dashboard stored in an S3 bucket, with AWS Lambda handling data wrangling, and Bedrock Knowledge Base fine-tuning the response with organisation's terminology and knowledge, while Lambda manages prompt engineering to maintain proper response structure. The Foundation Model (Claude 3 Sonnet) has taken up most of the cost.

To address this concern, this study focuses on building a tailored chatbot for users to access their diversity data in *Diversity Atlas*, a data analytic platform based in Australia that provides insight into cultural and demographic profile within an organisation (Diversity Atlas, 2024). This chatbot was limited to data collected from the dashboard and responded with generative AI and calculation logic. The project aims to present the perspective of diversity data collectors on AI-powered chatbot design and provide a solution that balances user experience and data security.

2. Literature Review

Despite growing public and organisational interest in D&I, early analytical efforts in this field were limited in scope and depth. Scholars such as Moieni and Mousaferiadis (2022) have argued that D&I has been "analytically neglected," with traditional methods failing to capture the inherent complexities and multidimensionality of diversity metrics. Their fractal analysis highlighted that diversity is not a linear or singular construct but rather a complex interplay of overlapping and self-similar patterns that require more sophisticated analytical tools. This literature review synthesises research from both the D&I domain and technological innovation to argue for the integration of advanced computational methods—particularly AI-driven chatbots—in the analysis and dissemination of diversity data.

2.1 Historical Neglect of Diversity and Inclusion Analytics

In the early stages of D&I research, academic attention was predominantly directed toward descriptive statistics and basic categorisations of workforce demographics (Moieni et al., 2022). The literature of that period tended to view diversity as a static attribute, focusing on easily quantifiable dimensions such as gender, age, and ethnicity. Such approaches, however, overlooked the nuanced and intersectional nature of identity, as well as the cultural and contextual factors that influence individuals' lived experiences. As Moieni et al. (2022) demonstrated through fractal analysis, diversity is characterised by recursive patterns and intricate interrelations that cannot be fully understood using traditional linear models. Early studies were also limited by the availability of data and the methods used to analyse it. The lack of comprehensive data collection frameworks and robust analytical tools contributed to an oversimplification of diversity metrics. Researchers noted that many organisations collected data primarily for compliance or superficial reporting, rather than for in-depth analysis that could inform strategic decision-making. As a result, there was a significant gap in the literature regarding how to operationalise the complex, multidimensional nature of D&I in analytical frameworks.

2.2 The Emergence of Technological Interventions in D&I

The advent of digital technologies and big data analytics has prompted a paradigm shift in how organisations approach D&I. With the growing need to comply with regulatory frameworks and to foster inclusive work environments, organisations have increasingly turned to technology to manage and interpret diversity data (Lokman & Ameedeen, 2019). The development of digital platforms for data collection and visualisation has enabled organisations to track a broader range of diversity metrics—from cultural identity and language to more nuanced aspects of intersectionality.

One significant technological advancement is the integration of artificial intelligence (AI) into diversity data management systems. AI technologies have been employed to process large volumes of data rapidly, generate insights, and support decision-making. For example, the use of machine learning algorithms in natural language processing (NLP) has allowed for the development of intelligent chatbots capable of interpreting user queries and providing tailored responses (Lokman & Ameedeen, 2019). These systems not only facilitate access to complex data but also democratises data analysis by allowing non-technical users to engage with D&I insights in a conversational manner.

Digital dashboards and interactive platforms have become critical tools for organisations aiming to visualise and interpret diversity data. Setlur et al. (2023) underscore that conversational interfaces embedded in these dashboards can transform static data into dynamic insights. By integrating chatbots into D&I platforms, organisations are able to offer an intuitive interface that bridges the gap between raw data and actionable intelligence. This transformation is particularly important for small and medium enterprises (SMEs), which may lack the resources to deploy advanced data analytics but still require robust mechanisms to assess and improve workplace diversity.

2.3 AI-Driven Chatbots in Diversity Data Analysis

The recent proliferation of AI-driven chatbots marks a turning point in the application of technology to D&I. Modern chatbots leverage advanced NLP, large language models (LLMs), and retrieval augmented generation (RAG) techniques to interpret user queries and generate contextually relevant responses. The architecture of such systems typically includes several interconnected components—knowledge retrieval, text processing, response generation, and machine learning modules—which together facilitate real-time interactions with users (Lokman & Ameedeen, 2019).

A central advantage of these systems is their ability to integrate structured diversity data with unstructured conversational input. By combining retrieval methods with generative models, chatbots are capable of contextualising sensitive diversity information in ways that are both secure and accessible. For instance, a tailored chatbot designed for a diversity analytics platform can extract data from secured sources, apply calculation logic, and then generate human-readable responses that help users navigate complex diversity metrics. This approach ensures that sensitive data—such as cultural and demographic profiles—is handled with the appropriate level of security and precision, mitigating the risk of unauthorised access or data leakage (Giordani, 2024).

Beyond simple data retrieval, AI-driven chatbots are increasingly used to support intersectional analysis by integrating multiple diversity dimensions. As organisational data becomes more complex, incorporating various attributes such as generation, ethnicity, and cultural identity, the ability to process and interpret this information becomes crucial. Vector embedding techniques—rooted in Transformer-based architectures (Vaswani et al., 2017)—are employed to capture the contextual relationships within large datasets. This not only enhances the accuracy of the insights generated but also minimises the risk of perpetuating biases inherent in traditional analytical models (Li et al., 2024).

Another critical facet of chatbot deployment in D&I is the management of privacy and security. With regulations such as the GDPR imposing strict guidelines on the handling of personal and sensitive data, AI systems must be designed with robust security features. Chatbots integrated into diversity dashboards implement multi-layered security protocols, including end-to-end encryption (E2EE), data segregation, and audit logging, to ensure compliance and protect user confidentiality (Blender, 2018; MongoDB, n.d.-a; AWS, 2023). These measures are not only necessary to meet legal requirements but also to maintain trust among users who are increasingly aware of the risks associated with data sharing.

2.4 Contextual and Implementation Challenges

While the integration of AI-driven chatbots into D&I initiatives offers considerable promise, it also presents several contextual and implementation challenges that warrant careful consideration. One of the primary issues relates to the inherent complexity of diversity data. Unlike traditional datasets, diversity information is characterised by its fluidity and multidimensionality. The fractal nature of diversity—as highlighted by Moieni et al. (2020)—implies that diversity patterns exist at multiple scales and are influenced by both visible and latent factors. This complexity necessitates the development of sophisticated analytical models that can adapt to evolving data structures and contextual nuances.

Furthermore, the implementation of AI systems in sensitive domains such as D&I raises ethical and technical concerns. For instance, while chatbots can facilitate rapid data retrieval and interpretation, they must be carefully designed to avoid reinforcing cultural biases. The language models underlying these systems are often trained on large, heterogeneous datasets that may contain biased representations. As a result, there is a risk that AI-generated responses could inadvertently perpetuate stereotypes or marginalise certain groups (Mehra, 2023). Researchers have called for the incorporation of bias mitigation strategies in the training and deployment of these models, emphasising the importance of transparency and accountability in AI systems (Giordani, 2024).

Privacy remains another critical challenge. Although many diversity datasets are aggregated to avoid the exposure of personally identifiable information (PII), the sensitive nature of the underlying data requires stringent privacy measures. AI systems that process diversity information must strike a delicate balance between data utility and confidentiality. As Li et al. (2024) point out, there is a persistent risk that large language models may “leak” sensitive information if proper safeguards are not in place. Therefore, the design of chatbot systems must incorporate data minimisation practices, secure data storage protocols, and regular privacy audits to ensure that user information remains protected.

Another contextual challenge relates to the interoperability of technological systems within diverse organisational settings. Organisations vary widely in their technological readiness and the sophistication of their existing data infrastructures. For smaller organisations or those in emerging markets, implementing a high-tech solution such as an AI-driven chatbot may pose significant resource and integration challenges. Studies have shown that the adoption of such technologies is often contingent on factors such as organisational culture, available expertise, and budget constraints (Setlur et al., 2023). Consequently, the development of cost-effective and scalable solutions is essential to ensure that technological advancements in D&I analytics are accessible to a wide range of organisations.

Moreover, the dynamic nature of language and cultural expression adds another layer of complexity. Chatbots must be able to interpret colloquial expressions, non-standard grammar, and culturally specific references accurately. This requirement is particularly challenging given the global context in which many organisations operate. The linguistic diversity of user populations necessitates the development of multilingual models and culturally adaptive interfaces that can accurately capture the nuances of language. Although current systems show promising results in terms of grammatical accuracy and context sensitivity, further research is needed to address these linguistic challenges comprehensively (Peras, 2018).

In addition to technical challenges, the organisational adoption of AI-driven solutions in D&I is influenced by broader social and political factors. Issues related to trust, ethical use of technology, and the potential for unintended consequences must be addressed to facilitate widespread acceptance. As organisations continue to experiment with AI technologies, there is a growing need for robust governance frameworks that can oversee the ethical implementation of these systems. This includes developing clear policies for data usage, establishing accountability mechanisms for AI decision-making, and ensuring that technological interventions do not inadvertently marginalise vulnerable groups (Hasal et al., 2021).

2.5 Synthesis and Future Directions

The integration of AI and conversational technologies into diversity and inclusion analytics represents a significant advancement over traditional methods. By harnessing the power of AI-driven chatbots, organisations can transform static diversity dashboards into interactive platforms that provide real-time, context-sensitive insights. This not only enhances the accessibility of complex data but also supports more informed decision-making by allowing non-technical users to engage with data in an intuitive manner.

The literature reviewed here underscores both the transformative potential and the multifaceted challenges of applying advanced technological solutions to D&I analytics. On the one hand, AI technologies—through the use of advanced NLP, LLMs, and secure data processing techniques—offer a promising pathway to address the longstanding analytical neglect of diversity (Moieni et al., 2020). On the other hand, significant challenges remain regarding bias mitigation, privacy protection, interoperability, and ethical governance. Addressing these challenges will require continued interdisciplinary collaboration among computer scientists, data analysts, organisational theorists, and ethicists.

Future research should focus on refining the algorithms that underpin AI-driven chatbots to ensure that they can more accurately interpret and analyse diversity data without compromising ethical standards. In particular, there is a need for longitudinal studies that assess the impact of these systems on organisational decision-making and employee perceptions of inclusion. Researchers are also encouraged to explore the integration of multilingual and culturally adaptive models, which could further enhance the relevance and utility of chatbot interfaces in global organisations.

Moreover, as regulatory frameworks around data privacy continue to evolve, the development of AI systems must remain agile and compliant with emerging standards. This may involve the incorporation of real-time privacy audits, the use of decentralised data storage solutions, or the deployment of federated learning approaches that reduce the risk of centralised data breaches. In sum, while AI-driven chatbots offer a compelling solution to many of the challenges inherent in D&I analytics, their successful implementation hinges on a careful balance between technological innovation, ethical responsibility, and organisational readiness.

3. Case Study: Build AI Chatbot (demo) for Diversity Atlas

Figure 1 illustrates the overall architecture of the developed chatbot.

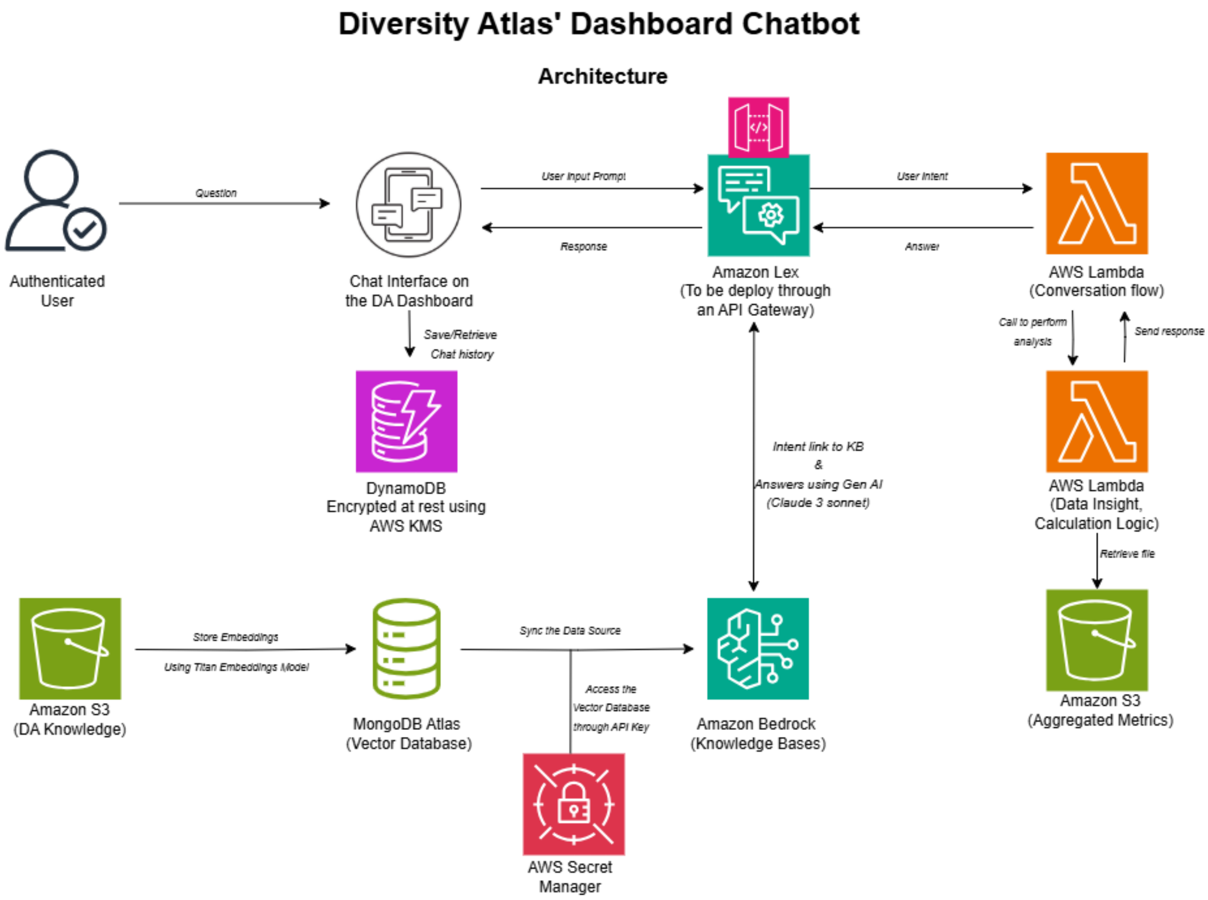

Figure 1. Architecture of the developed chatbot

Following a AI-empowered Business Intelligence framework proposed by Azmi et al. (2023), below is a flowchart of the chatbot-user interactions in the developed chatbot:

1. *User Interaction:* User interacts with the Chatbot through a pop-up chat window in a web browser. User can select options and input text to ask questions. The chatbot provides two options for user to start with: 1) General questions, and 2) Data insights.
2. *NLP*: User's input is passed through the NLP component of the Chatbot. AWS Lex understands and interprets the text.
3. *Query Processing*: Chatbot processes the interpreted text and formulates a data query to retrieve the data insights
4. *Data retrieval*: Depending on what the user selects, the chatbot connects to different AWS S3 buckets for data. If a user chooses "General questions", the chatbot reads information from the knowledge base in a AWS S3 bucket. If the user selects "Data Insights", the chatbot accesses another AWS S3 bucket that contains the diversity data of the user and retrieves the data.
5. *Insights generation*: For queries related to "General questions", the retrieved data is processed within AWS Bedrock. For the "Data insights" query, calculation logics in AWS Lambda are applied to perform calculations and generate insights.

6. *Response generation*: Based on the processed data and insights, AWS Lex generates a response to the user in text.
7. *Dialogue management*: The Chatbot handles follow up questions or further queries from the user.
8. *User feedback and learning*: Chat history would be logged to enhance its understanding and response capabilities.

Regarding architectural design, the rest of this section will be organised around the four main elements (Baird, 2024; Lokman & Ameedeen, 2019):

3.1 Knowledge Domain

The developed chatbot is a closed-domain chatbot with specialised knowledge about Diversity Atlas (DA). The knowledge was built in Amazon Bedrock Knowledge Base, retrieving knowledge data from Amazon S3 bucket that stores knowledges such as FAQs, Cultural Calendar data, and DA's website pages.

- **FAQs:** FAQs are the frequently asked questions and answers provided in DA public website https://diversityatlas.io/faq/. It builds the knowledge of the terminology used within the dashboard. See Table 1 for FAQs example.

Table 1. Extract of FAQs. This data explains the definition of diversity and how it relates to the DA dashboard

Question	Answer
What Is Diversity?	Diversity can be any attribute that makes us different from one another. Diversity is both attributive and cognitive, and can refer to differences based on cultural heritage, gender, age, language, religion, disability and sexual orientation, and to many more dimensions such as education, occupation, tenure, personality, socioeconomic status, marital or parental status, and so on. The list is almost endless! We have given great consideration to the metrics that best capture salient diversity for any organisation. We regularly review and conduct research in order to ensure Diversity Atlas captures key diversity data. We emphasise cultural diversity, demographic diversity and intersectionality in order to provide nuanced, considered and useful insights in diversity results.

- **Cultural Calendar data:** Cultural Calendar is a feature of DA dashboard that display culturally significant days and events around the globe. It builds the knowledge of the DA dashboard and enable users to search for events relevant to their cohort based on their own data. See Table 2 for Cultural Calendar data example.

Table 2. Extract of Cultural Calendar data. This data presents the date range and name of the East Timor Independence Restoration Day

Start Date	End Date	Event Name
20-05-2024	21-05-2024	East Timor Independence Restoration Day

3.2 Response Generation

A hybrid approach was used to produce responses with both retrieval and generative methods. The developed chatbot, written in Python, used Amazon Lex, AWS Lambda, and Large Language Models (LLMs) hosted on Amazon Bedrock to build a Retrieval Augmented Generation (RAG) system. RAG enables the chatbot to obtain information from Knowledge Base to serve as context before creating a response and then supply it to the LLMs to generate a response based on that information (Baird, 2024). This ensures the response generated is contextually appropriate and precise.

Figure 2 illustrates the workflow of Amazon Lex in collecting user input and producing responses.

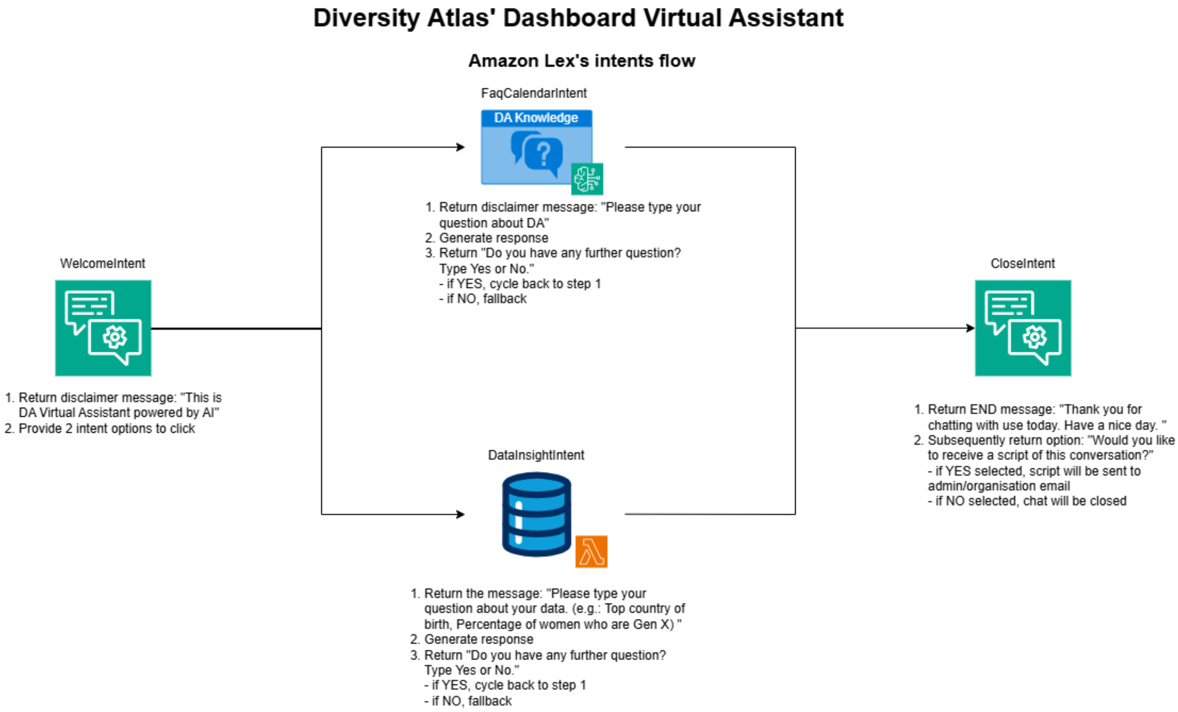

Figure 2. Architecture of Amazon Lex intents flow

There are four intents in total:

1. *Welcomeintent:* This is an intent to activate welcome message when user start the conversation in the chatbot.
2. *FaqCalendarIntent:* This is an intent for “General questions” where user can ask questions about Diversity Atlas.
3. *DataInsightIntent:* This is an intent for “Data insights” where user can ask questions about their diversity data.
4. *CloseIntent*: This is an intent to end the conversation in the chatbot.

Regarding “DataInsightIntent”, when the user selects "Data insights" in the chatbot, it will read the diversity data of the user. For demonstration purposes, sample diversity data was used, which does not contain any real client data.

The diversity data is a set of 29 questions and answers collected from Diversity Atlas (DA), a platform that gathers and visualises D&I data for organisations. The sample data has been aggregated. Data frames for each set of questions and answers were created to pre-process the data, which were uploaded to Amazon S3 bucket. See Table 3 for sample diversity data example.

Table 3. Sample diversity data (extract). It shows the data on country of birth by gender

Your country of birth							
Continent	**Country**	**State**	**Woman**	**Man**	**Total**	**Prefer not to answer**	**Other**
Europe	France	Ile De France	3	6	9	0	0
Europe	France	Provence Alpes Cote D'Azur	1	4	5	0	0
Europe	France	Hauts De France	1	2	3	0	0

Calculation logic is written and stored in the AWS Lambda to serve as the context to define how the calculation should be performed for Data insights.

The tailored calculation logic adapts to the data structure of the diversity data to generate more nuanced insight. For example, defining the birth years for Generation X helps in accurately categorising and analysing the data to reduce errors. It also ensures that every query is processed using the same logic for generating consistent, reliable and comparable insights for the clients (see Appendix B for the code).

The chatbot's data processing pipeline in AWS lambda function incorporates multiple data wrangling techniques including Pandas' data frames aggregation, normalising demographic data across comprehensive intersectional categories. These categorises encompass generation (i.e. birth year), gender, ancestry, education level, sexuality, disability status and language preference. Through the implementation of Amazon Bedrock's guardrails and content filtering in Amazon Lex, the system ensures culturally sensitive, accessible, and bias-free responses across diverse user populations and contexts.

3.3 Text Processing and Machine Learning Models

The use of machine learning is mainly on converting client's diversity data into vector embedding through MongoDB Atlas by Titan Embeddings that was hosted on Amazon Bedrock. Vector embeddings are numeric representations of the text data that capture the contextual relationships within the dataset (Lokman & Ameedeen, 2019).

The model used is based on Transformer-based architecture (Vaswani et al., 2017). It is different from Recurrent Neural Networks (RNNs) such as Long Short-Term Memory Network (LSTMs) or Gated Recurrent Units (GRUs), Transformers could capture complex relationship within aggregated diversity data that is not sequential in faster training times.

4. Evaluation

The success of this case study is defined from four perspectives: Security, Privacy, Performance, and cost, which are the challenges aimed to be addressed at the beginning.

4.1 Security

There are multiple potential security threats to user messages when transferred to the server hosting the chatbot, as well as to the user's data processed, stored, and shared on the server backend (Hasa et al., 2021).

On the Diversity Atlas dashboard, users must verify their login credentials in the chatbot again after logging into the dashboard, prior to requesting any data (see Appendix A). Although this chatbot could be implemented on other platforms, such as WordPress, it is important to remove the confidential diversity data and adapt the chatbot to meet the specific requirements of those websites.

Article 32 (a) of the GDPR specifically requires that organisations pseudonymise and encrypt personal data (GDPR-Text.com, n.d.). The developed chatbot responds data queries with personal data, so End-to-End Encryption (E2EE) is required to ensure only the organisation admins have access to the conversation. By using Amazon Lex and Amazon DynamoDB, all conversations between user and the chatbot, as well as the chat history, are encrypted with encryption keys and Transport Layer Security (TLS) to ensure the data in transit over the network and data at rest is encrypted (Blender, 2018). MongoDB Atlas employs strong encryption measures by using 256-bit Advanced Encryption Standard (AES-256) encryption algorithm for data at rest and (TLS) for data in transit (MongoDB, n.d.-b). To enhance data security, database-level encryption using AWS Key

Management Service (KMS) will be implemented, which provides centralized control of encryption keys (AWS, 2023).

The overall security assessment will involve the following components:

- *Data protection:* This includes advanced encryption for data at rest and in transit, secure solutions to protect confidential data, and robust data transfer protocols using secure channels like TLS/SSL.
- *Authentication and authorisation mechanism:* To further safeguard the DA dashboard, we will implement an additional multi-factor authentication system for accessing the chatbot, requiring users to verify their identity.
- *Vulnerability test:* This includes conducting security audits to identify weaknesses, continuous monitoring for real-time threat detection, and regular penetration testing to evaluate our measures.
- *Incident Response:* A comprehensive incident response strategy that includes ongoing threat monitoring and clear protocols for addressing security breaches, ensuring quick resolution while minimising impacts.

4.2 Privacy

The diversity dataset itself is not Personal Identifiable Information (PII) as it has been aggregated by DA at the first place. However, these data are confidential information that require protection. While Li et al. (2024) highlights potential risk of data privacy leakage in the LLMs when these models could access and learn from the dataset for training and fine-tuning, it is a critical concern to ensure that the models employed in the chatbot adhere to data privacy protection standards.

By building the Retrieval Augmented Generation (RAG) system, the client's diversity data is provided to the machine learning models only when Amazon Lex is asking for generated content, which is not stored in the models themselves as knowledge (Proser, 2024). The knowledge base of the developed chatbot is built from the training data provided to the LLMs, which is sample data from a demo account that does not contain real customer data.

The evaluation of the chatbot privacy centres on three critical components:

- *Data minimisation:* Conduct a systematic audit of all data collection points, conversation logs and RAG implementations
- *Data retention and deletion practices:* Reviewing retention policies by testing automatic deletion triggers and complete data purge across all storage systems including backups and cached data
- *Compliance with regulations:* Perform comprehensive compliance assessment against privacy regulations (GDPR) documentations

4.3 Performance

To evaluate the performance of the chatbot, its user experience, information retrieval and linguistic aspects would be discussed based on Peras's chatbot evaluation metrics (2018).

Regarding user experience, before developing this chatbot, organisation administrators had to navigate the platform to obtain diversity data. This involved clicking on the correct panel and dropdown menu to find specific information. For example, to check employee responses to the "Identity Priority" question, users had to first click on the "Culture" panel in the menu, then select "Identity Priority" from the dropdown in a table (see Figure 3).

With the introduction of the chatbot, users no longer need to click through various panels and dropdown menus after logging into the dashboard. Instead, the chatbot opens a pop-up window on the home page to start the conversation. This new approach has significantly reduced the time required to obtain specific data, proofing the effectiveness of the chatbot.

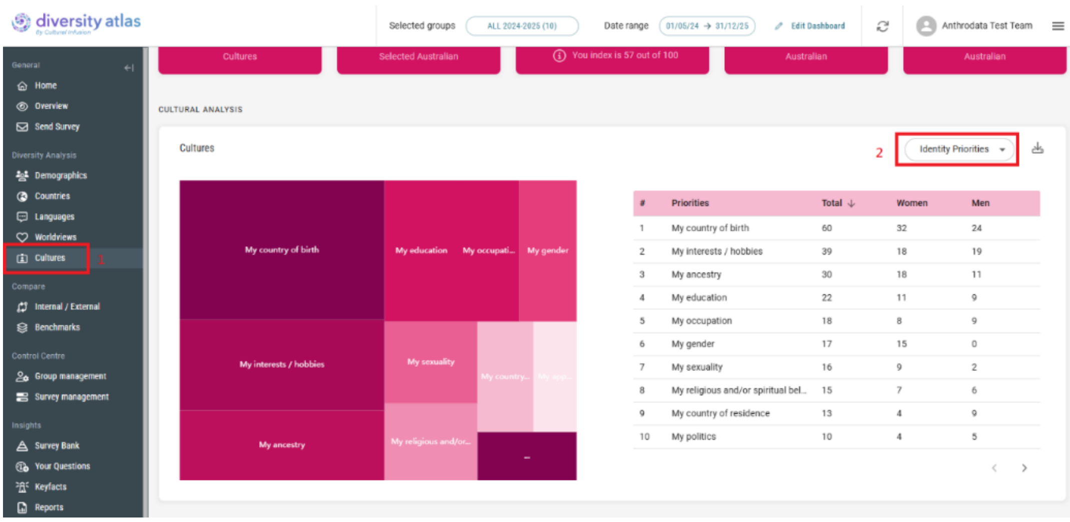

Figure 3. Dashboard screenshot in Culture panel

Regarding information retrieval, although the chatbot is unable to retrieve all information available on the dashboard—such as comparisons between internal groups and third-party census data and internal benchmarking—it is capable of providing accurate intersectional information stored in its knowledge base.

For example, a basic query like “What is the most common non-heterosexual selection?” can be processed and responded to within two seconds (see Figure 4). The response meets users' information needs with high accuracy and efficiency, without bias towards people who are not heterosexual.

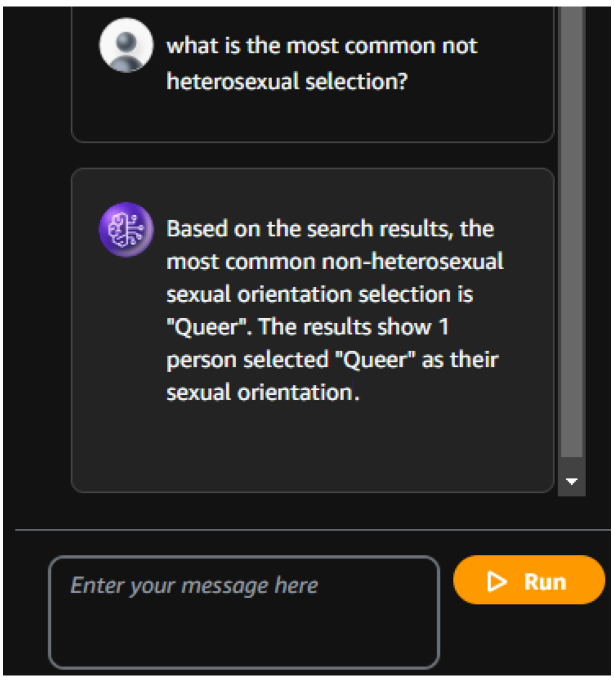

Figure 4. Screenshot of the chatbot (in testing mode) in answering “What is the most common not heterosexual selection?”

Regarding linguistic, Figure 5 illustrates an example where the user inquired about the five most common countries of birth but did so using incorrect grammar. The response, however, provided a grammatically accurate output, listing five country names clearly and without ambiguity. The number of people associated with each country is relevant to the context of the question, demonstrating a high level of linguistic accuracy. Currently, the chatbot does not support multilingual features, which could be considered for future development.

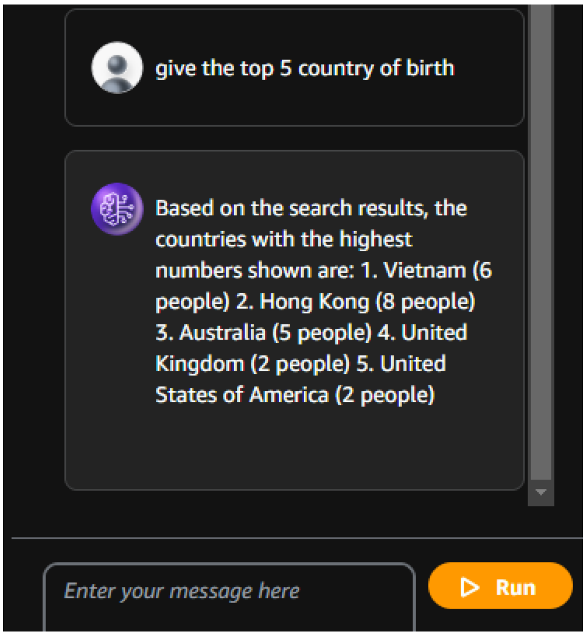

Figure 5. Screenshot of the chatbot (in testing mode) in answering "Give the top 5 country of birth"

The performance framework will evaluate the chatbot operational efficiency through these metrics:

- *Response time:* Perform latency analysis by measuring the speed of initial response speed and total conversation completion times, including average response times and variability in delays.
- *Accuracy of response:* Evaluate the accuracy of responses based on the actual data and the semantic relevance, which determines if they meet user expectations.
- *Reliability:* Review the reliability of the system by monitoring uptime and error rates, including the frequency of data inaccuracy during interaction.

4.4 Cost

The cost associated with developing and operating a chatbot can be unpredictable, influenced by factors such as usage and performance. One of the key objectives of this case study is to create a customised architecture that uses the right tools to help small organisations like Diversity Atlas reduce their operational costs. Throughout this study, various AWS tools were employed to achieve a balance between performance and cost. Table 4 below presents the cost of each tool and the total cost.

Table 4. Cost of the AI chatbot architecture in this case study

Tools	**Cost (USD)**
1. AWS Lex	$0.004 per speech & $0.00075 per text request
2. AWS S3	$0.023 per GB / month
3. AWS Bedrock	$0.008 / 1000 input tokens & $0.024 / 1000 output tokens
4. AWS Lambda	• $0.0000166667 per every GB-second & Requests • $0.20 per 1M requests for the first 6 billion GB-seconds per month
5, Anthropic Claude 3 Sonnet	$0.003 per 1000 input tokens & $0.015 per 1000 output tokens
6. MongoDB Atlas	$9.00 (shared cluster)

7. DynamoDB	$1.25 per million writes & $0.25 per million reads
8. API Gateway (Rest APIs)	$3.5 / million requests per month (for the first 333 million) + data transfer charges

The cost of this chatbot would be greatly reduced by paying per usage rate, instead of number of users, which could be an affordable solution in the case of smaller organisations.

The economic viability of the chatbot depends on factors such as

- *Development costs:* This includes costs related to software development, design coding, testing, deployment, and integration with existing systems.
- *Operational costs:* Ongoing expenses such as cloud infrastructure usage, maintenance support, regular updates, and licensing fees for AWS services.
- *Cost efficiency:* Comparing chatbot performance against manual operation in the dashboard
- *Resourceful usage:* Evaluating how effectively the chatbot optimises resources includes analysing metrics for system performance, scalability, and its capacity to manage increasing conversations without sacrificing quality.

5. Conclusion

In this paper, a review of AI-driven chatbot concept and challenges with using the chatbot for reporting diversity data are introduced. The case study explored the possible chatbot architecture and various AWS tools that balance the user experience and expectation of interpretating diversity data while addressing security and privacy concerns. This paper proposed that a customised architecture, using the right tools, can be both well utilised and cost-effective for small organisations. This approach can mitigate data security issues and provide a compelling alternative to more advanced solutions available in the market.

6. Limitations and Future Work

This chatbot was developed using the Claude 3 Sonnet model, which requires a greater amount of pre-processing time for data wrangling compared to the Claude 3.5 Sonnet model. The limitation is due to the unavailability of the Claude 3.5 Sonnet model in the AWS Asia Pacific (Sydney) region at the time of development. Developing the chatbot in the Sydney region was, however, necessary because the Diversity Atlas database is hosted there.

Future development efforts should consider incorporating the Claude 3.5 model if it becomes available. Besides, the Claude 3 model has already been retired in many other regions and may soon be retired in Sydney as well. Therefore, the architecture should be updated in the future based on the availability of the model.

Besides, regarding “General questions” intent, currently there is no calculation logic applied on it. In the future, the chatbot may suggest cultural events based on the diversity data of the user. For example, a Buddhism-related event could be suggested if Buddhism is the most common worldview selected in the organisation. This would require more resources from using both Bedrock and Lambda functions to access the knowledge base and diversity dataset.

Acknowledgement

This research was supported by Diversity Atlas and their provision of data has been instrumental in shaping the findings of this study. We thank the entire Diversity Atlas team for their support, although they may not agree with all the interpretations/conclusions of the results.

References

Argal, A., Gupta, S., Modi, A., Pandey, P., Shim, S., & Choo, C. (2018). Intelligent travel chatbot for predictive recommendation in echo platform. In *2018 IEEE 8th annual computing and communication workshop and conference (CCWC)*, 176-183. IEEE. https://doi.org/10.1109/CCWC.2018.8301732

AWS. (2023). *AWS Key Management Service features*. Retrieved January 21, 2025, from https://aws.amazon.com/kms/features/

AWS. (2023). *AWS Key Management Service*. Retrieved January 21, 2025, from https://aws.amazon.com/kms/

Azmi, M., Mansour, A., & Azmi, C. (2023). A context-aware empowering business with AI: Case of chatbots in business intelligence systems. *Procedia Computer Science, 224*, 479-484. https://doi.org/10.1016/j.procs.2023.03.073

Baird, C. (2024). Prototype a RAG chatbot with Amazon Bedrock, Kendra, and Lex. *AWS Community*. Retrieved from https://community.aws/content/2eBsTWhvFFPUtzh2secQlNRBgta/prototype-a-rag-chatbot-with-amazon-bedrock-kendra-and-lex?lang=en

Benotti, L., Martínez, M. C., & Schapachnik, F. (2014). Engaging high school students using chatbots. In *Proceedings of the 2014 conference on Innovation & technology in computer science education*, 63-68. Retrieved from https://dl.acm.org/doi/proceedings/10.1145/2591708

Blender, J. (2018). How to encrypt and sign DynamoDB data in your application. *AWS*. Retrieved from https://aws.amazon.com/blogs/security/how-to-encrypt-and-sign-dynamodb-data-in-your-application/

Blender, J. (2018). How to encrypt and sign DynamoDB data in your application. AWS Security Blog. Retrieved January 21, 2025, from https://aws.amazon.com/blogs/security/how-to-encrypt-and-sign-dynamodb-data-in-your-application/

Câmpeanu, I. A., Dragomir, D. A., & Zota, R. D. (2023). Banking chatbots: How artificial intelligence helps the banks. In *Proceedings of the International Conference on Business Excellence* (Vol. 17, No. 1, pp. 1716-1727). https://doi.org/10.2478/picbe-2023-0153

Diversity Atlas. (2024). How it works. Retrieved January 21, 2025, from https://diversityatlas.io/how-it-works/

GDPR-Text.com. (n.d.). Article 32 of the General Data Protection Regulation (GDPR) – Security of processing. *GDPR-Text.com*. Retrieved January 21, 2025, from https://gdpr-text.com/read/article-32/

Giordani, J. (2024). Mitigating chatbots AI data privacy violations in the banking sector: A qualitative grounded theory study. *European Journal of Applied Science, Engineering and Technology*, *2*(4), 14-65. https://doi.org/10.59324/ejaset.2024.2(4).02

Hasal, M., Nowaková, J., Ahmed Saghair, K., Abdulla, H., Snášel, V., & Ogiela, L. (2021). Chatbots: Security, privacy, data protection, and social aspects. *Concurrency and Computation: Practice and Experience*, *33*(19), e6426. https://doi.org/10.1002/cpe.6426

Hasal, M., Nowaková, J., Ahmed Saghair, K., Abdulla, H., Snášel, V., & Ogiela, L. (2021). Chatbots: Security, privacy, data protection, and social aspects. *Concurrency and Computation: Practice and Experience, 33*(19), e6426. https://doi.org/10.1002/cpe.6426

Lee, J., An, T., Chu, H. E., Hong, H. G., & Martin, S. N. (2023). Improving science conceptual understanding and attitudes in elementary science classes through the development and application of a rule-based AI chatbot. *Asia-Pacific Science Education, 9*(2), 365-412. https://doi.org/10.1163/23641177-bja10070

Li, Q., Hong, J., Xie, C., Tan, J., Xin, R., Hou, J., … Li, B. (2024). LLM-PBE: Assessing data privacy in large language models. *arXiv preprint arXiv:2408.12787*. https://doi.org/10.48550/arXiv.2408.12787

Li, Q., Hong, J., Xie, C., Tan, J., Xin, R., Hou, J., … Song, D. (2024). Llm-pbe: Assessing data privacy in large language models. *arXiv preprint arXiv:2408.12787*. https://doi.org/10.48550/arXiv.2408.12787

Lokman, A. S., & Ameedeen, M. A. (2019). Modern chatbot systems: A technical review. In *Proceedings of the Future Technologies Conference (FTC) 2018* (Volume 2, pp. 1012-1023). Springer International Publishing. https://doi.org/10.1007/978-3-030-02683-7_75

Madasamy, S., & Aquilanz, L. L. C. (2023). The evolution of chatbots: Cloud and AI synergy in banking customer interactions. *Journal of Emerging Technologies and Innovative Research*, *10*, 1-9. Retrieved from https://www.jetir.org/papers/JETIR2303312.pdf

Madhu, D., Jain, C. N., Sebastain, E., Shaji, S., & Ajayakumar, A. (2017). A novel approach for medical assistance using trained chatbot. In *2017 International conference on inventive communication and computational technologies (ICICCT)* (pp. 243-246). IEEE. https://doi.org/10.1109/ICICCT.2017.7975195

Mehra, P. (2023). OpenAI temporarily shutdown ChatGPT to fix bug exposing user chat titles.

Moieni, R., & Mousaferiadis, P. (2022). Analysis of Cultural Diversity Concept in Different Countries Using Fractal Analysis. *The International Journal of Organizational Diversity, 22*(1), 43. https://doi.org/10.18848/2328-6261/CGP/v22i01/43-62

MongoDB. (n.d.-a). *MongoDB Atlas security*. Retrieved from https://www.mongodb.https//www.mongodb.com/resources/products/capabilities/mongodb-atlas-securitycom/resources/products/capabilities/mongodb-atlas-security

MongoDB. (n.d.-b). *Security – Encryption at rest and in-transit*. https://www.mongodb.com/products/capabilities/security/encryption

Peras, D. (2018). Chatbot evaluation metrics. *Economic and Social Development: Book of Proceedings*, 89-97. Retrieved from https://www.proquest.com/conference-papers-proceedings/chatbot-evaluation-metrics-review-paper/docview/2176212638/se-2

Proser, Z. (2024). Retrieval augmented generation (RAG). *Pinecone*. Retrieved January 21, 2025, from https://www.pinecone.io/learn/retrieval-augmented-generation/

Sebastian, G. (2023). Privacy and data protection in ChatGPT and other AI chatbots: Strategies for securing user information. *International Journal of Security and Privacy in Pervasive Computing (IJSPPC), 15*(1), 1-14. https://doi.org/10.4018/IJSPPC.325475

Setlur, V., Correll, M., Satyanarayan, A., & Tory, M. (2023). Heuristics for supporting cooperative dashboard design. *IEEE transactions on visualization and Computer Graphics*, *30*(1), 370-380. https://doi.org/10.1109/TVCG.2023.3327158

Su, M. H., Wu, C. H., Huang, K. Y., Hong, Q. B., & Wang, H. M. (2017). A chatbot using LSTM-based multi-layer embedding for elderly care. In *2017 International conference on orange technologies (ICOT)* (pp. 70-74). IEEE. https://doi.org/10.1109/ICOT.2017.8336091

TECHciRCLE. Retrieved January 21, 2025, from https://www.techcircle.in/2023/03/21/openai-temporarily-shutdown-chatgpt-to-fix-bug-exposing-user-chat-titles

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., … Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems, 30*, 31st Conference on Neural Information Processing Systems, Long Beach, CA, USA. Retrieved from https://papers.nips.cc/paper/7181-attention-is-all-you-need

Appendix A. Login Credential in the Chatbot

See Figure 6 below for the design of login credential in the chatbot after login on the dashboard.

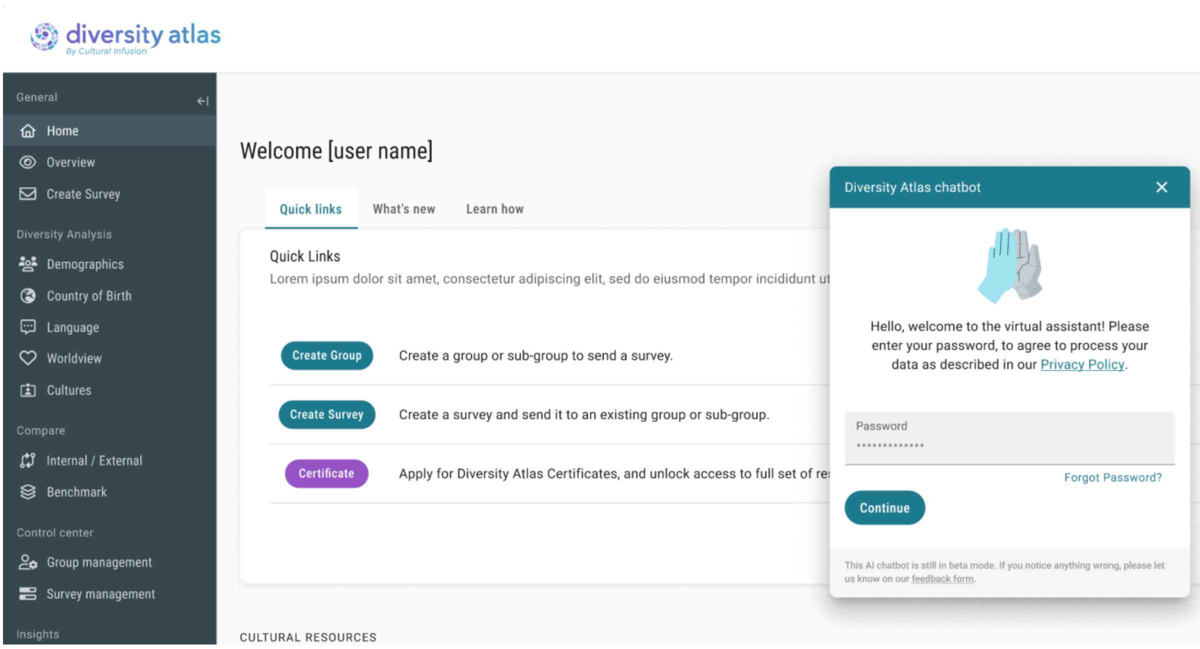

Figure 6. Screenshot of chatbot on the dashboard with login credentials required

Appendix B. Codes of the Calculation Logics in Lambda

See Code A and Code B below for the example codes of the calculation logics used for different analysis.

```
elif analysis_type == "DEMOGRAPHIC_FILTER":
        generation = None
        generation_keywords = {
            'silent generation': 'Silent Generation',
            'baby boomer': 'Baby Boomers',
            'gen x': 'Gen X',
            'gen y': 'Gen Y',
            'gen z': 'Gen Z'
        }
        question_lower = question.lower()
        for keyword, gen_name in generation_keywords.items():
            if keyword in question_lower:
                generation = gen_name
                break

        if not generation:
            logging.error("Could not determine generation from question")
            raise ValueError("Generation not specified in question")
```

```
gen_filtered_df = df[df['Generation'] == generation]

total_count = df[metric].sum()
gen_count = gen_filtered_df[metric].sum()

value_counts = {
    'generation': generation,
    'count': float(gen_count),
    'total': float(total_count),
    'percentage': round((gen_count / total_count * 100) if total_count > 0 else 0, 2)
}

logging.info(f"Demographic analysis results: {value_counts}")
```

Code A. The analysis "DEMOGRAPHIC_FILTER" will be triggered if the question asked has the terms "generation" or any of the generation name. Example: The percentage of women who are generation X.

```
elif analysis_type == "COUNT":
        filtered_df = df.copy()

        if 'Level' in df.columns:
            logging.info("Before filtering:")
            logging.info(f"Available levels: {filtered_df['Level'].unique()}")
            logging.info(f"Total rows: {len(filtered_df)}")

            level_filter = None
            if 'advanced' in question.lower():
                level_filter = 'Language 3'
            elif 'intermediate' in question.lower():
                level_filter = 'Language 2'
            elif 'basic' in question.lower():
                level_filter = 'Language 1'

            if level_filter:
                filtered_df = filtered_df[filtered_df['Level'] == level_filter]
                logging.info(f"After filtering for {level_filter}:")
                logging.info(f"Remaining rows: {len(filtered_df)}")
                logging.info(f"Sample of filtered data:\n{filtered_df.head()}")

        if 'Total' in filtered_df.columns:
            filtered_df['Total'] = pd.to_numeric(filtered_df['Total'], errors='coerce')
```

```
value_counts = filtered_df.groupby('Language')['Total'].sum().fillna(0).to_dict()

value_counts = {k: v for k, v in value_counts.items() if v > 0}
value_counts = dict(sorted(value_counts.items(), key=lambda x: x[1], reverse=True))

logging.info(f"Final counts for analysis: {value_counts}")
```

Code B. The analysis "COUNT" is triggered for questions like "Language spoken at advanced level".