

CHATBOT BUILT WITH DEEP LEARNING: FANNIE MAE INTERNSHIP
THE INTERPLAY BETWEEN ARTIFICIAL INTELLIGENCE AND SOCIETY FOR
MORTGAGE AUTOMATION

A Thesis Prospectus
In STS 4500
Presented to
The Faculty of the
School of Engineering and Applied Science
University of Virginia
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Computer Science

By
Sooren Ghodsi

October 27, 2023

Technical Team Members: Fannie Mae Securitization Team

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

ADVISORS

Prof. Rider Foley, Department of Engineering and Society

Prof. Rosanne Vrugtman, Department of Computer Science

Introduction

In our rapidly transforming world of digital technology, Artificial Intelligence (AI) stands out as one of the pinnacles of innovation, constantly challenging the current way of thinking and our societal norms. Just like how the spread of electricity revolutionized industrial processes and the way of life during the 20th century (Smil, 2005), AI is projected to redefine many fields, change job descriptions, and alter how we experience things. This transformative potential of AI raises several ethical questions. Bostrom & Yudkowsky (2014) have argued that the ethical implications of AI advancements are profound, touching on issues like human well-being, autonomy, and justice. They emphasize the need for careful consideration when developing AI to make sure that these technologies will help society while reducing risks related to their decision-making. Adding to this conversation, Haenlein & Kaplan (2019) discuss how AI is set out to dramatically change the job landscape by automating tasks, leading the way for a new economy that is about human empathy and interpersonal skills. This would completely change job roles and the nature of work itself. The AI revolution is showing us that there is potential to turn even more of our remaining manual tasks into automated ones, leaving us with only the most valuable human-specific tasks.

However, these advancements are not only about working more efficiently but also about improving lives and providing new skills and opportunities to everyone. In recent years, a vast number of AI solutions have been built to tackle this challenge, with AI chatbots emerging as a prime example of these innovations. Among these chatbots, ChatGPT stands out as a cutting-edge language model, designed to simulate human-like conversation, offering a more dynamic and responsive experience compared to traditional chatbots. Eloundou et al. (2023) found that large language models like Generative Pre-trained Transformer (GPT) could affect about 80% of

the U.S. workforce. Their research estimated that 19% of jobs would be heavily influenced by AI models, including GPT-4, with at least 50% of tasks in these jobs being automated. They hypothesized that chatbot technologies will have a massive impact on higher-income jobs, especially with roles such as writers, web and digital designers, financial quantitative analysts, and blockchain engineers. The sheer potential of chatbots in transforming our future, like redefining the workforce, is the main driving force behind my motivation for this project.

On the other hand, the Federal National Mortgage Association (Fannie Mae), a leading financial institution in the mortgage industry, engages with AI from a different perspective. They have traditionally been cautious in embracing AI advancements, prioritizing data privacy and security when it comes to managing sensitive financial information. Nonetheless, despite their conservative stance, Fannie Mae sees the potential benefits of integrating advanced AI solutions, like this chatbot project, into their systems. Their long-term goal is not just to improve the company performance, but it is also to maintain their mission of making home ownership more accessible to everyone. By funding this project, they are attempting to offer more affordable mortgage options to future homeowners, aligning with their mission of creating a stable and sustainable housing market.

Chatbot Built with Deep Learning: Fannie Mae Internship

Fannie Mae has historically witnessed the difficulties that come with managing large volumes of data associated with mortgages and housing trends. Efficiently accessing mortgage data is vital for daily operations, strategic decision-making, and maintaining market relevance. My team and I identified that our process of running certain database queries was reserved only for a small group of business associates with the specialized knowledge to interpret them. This

exclusivity made the process extremely time consuming and limited work efficiency since all other employees had to wait for the few associates to manually run the queries for them. Addressing this issue, I developed a deep learning chatbot to automate these processes and bring data access to all Fannie Mae employees.

To achieve this, I used TensorFlow, a simple yet powerful framework for natural language processing (NLP), to build the chatbot's neural network model. Aleedy et al. (2019) explains that NLP is a crucial component of chatbot technology, allowing them to understand, interpret, and respond to human language in a natural way. The chatbot was also integrated with Fannie Mae's Oracle databases as a requirement for running data queries. Oracle Database, being one of the main database management systems used by the company, contains a large portion of Fannie Mae's mortgage data. This integration allowed the chatbot to gain access to the most relevant and up-to-date information.

The design of this chatbot required a proper understanding of the different queries and data interactions that were specific to Fannie Mae's internal systems. By collaborating with business associates, we came up with a set of major use cases for the chatbot, focusing on business-related queries that were previously challenging to run. This collaborative work was essential because it directly affected the chatbot's functionality and its ability to effectively address the issue. The set of use cases identified were then transformed into a list of intents placed in the "intent file," formatted in JSON. This file is the learning foundation for the chatbot, comprised of tags, patterns, responses, and sometimes actions. This structure, shown in Figure 1, enables the chatbot to associate user commands with appropriate responses.

```

{"tag": "dus_info",
 "patterns": ["I have a question about the dus lenders program", "I want to learn
about dus", "enquiry about dus"],
 "responses": ["Okay, what do you want to know?", "what are you curious about?",
"Okay, what is your question"]
},
{"tag": "query_xyz",
 "patterns": ["Run business query xyz", "query xyz please", "can you do query xyz?"],
 "responses": ["Sure, here are the results: ", "These are the results: ", "your
results: ", "Sorry, I couldn't process your query"],
 "actions": ["query xyz"]
}

```

Figure 1: Sample Code Snippet from an Intent File.

The neural network structure used in the chatbot development involves a feed-forward network created with TensorFlow, having four layers (see Figure 2): one input layer, two hidden layers, and one output layer (Goldberg, 2016). During training, the intent file is used to create a “bag of words” array, representing all unique words. When a user command is received, it goes through a three-part NLP pipeline to convert it into this “bag of words” format: Tokenization turns sentences into individual words, lemmatization reduces these words to their base forms, and part-of-speech tagging helps with the understanding of grammar structure. The model then uses this data to predict and generate responses, picking the best choice based on the probability of different tags. The system also includes a feedback mechanism, where employee ratings on responses are stored and used to improve and improve interactions in the future.




Figure 2: Feed-Forward Neural Network Design (Ognjanovski, 2020)

Anticipated features for the chatbot include moving from development to a full-scale production environment, integrating into most internal system user interfaces (UI), and improving user-friendliness. The project's success will be measured not just in terms of efficiency, but also in how it levels the playing field, giving all employees, regardless of their technical backgrounds, access to important data and knowledge.

Stakeholders and the SCOT Framework

The potential impacts of my AI chatbot on the dynamics of an entire company makes it clear that no advancements in technology, especially AI-related ones, exist in a vacuum. They are deeply embedded in our societal structures and human interactions. Humans are the architects of chatbots. They are the developers, end-users, and business associates.

Although, at a glance, the chatbot was made just for this group of actors, it affects an entire network of stakeholders who utilize this complex system. The securitization process (see

Figure 3) at Fannie Mae involves the Delegated Underwriting and Servicing (DUS) lenders like Wells Fargo giving loans to borrowers, which Fannie Mae later purchases (Segal & Szymanoski, 1998). These loans are then grouped into Mortgage-Backed Securities (MBS) and sold to investors. Fannie Mae guarantees timely payments on these MBS, even if borrowers default. Auditing firms like KPMG make sure that these transactions are accurate and complying with regulations. The funds from MBS sales allow Fannie Mae to buy more mortgages, keeping liquidity flowing in the housing market.

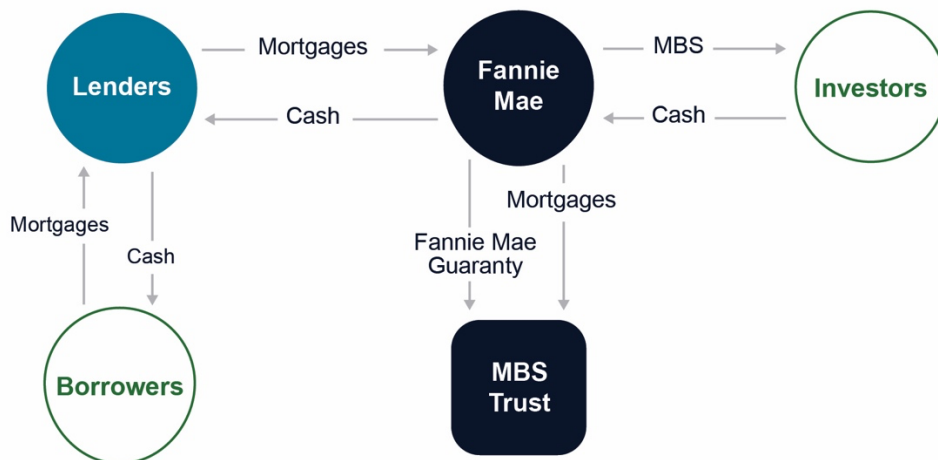


Figure 3: Fannie Mae Portfolio Securitization Transaction (Federal National Mortgage Association, 2022)

Going back to the context of the chatbot, developers brought their technical skills and vision to the project, while business associates contributed with their business-specific knowledge and possible chatbot use cases. The end-users, who could be any Fannie Mae employee, are the main beneficiaries of the chatbot. They will be providing feedback over time and molding the chatbot's feature and UI designs. An experiment conducted by Noy & Zhang

(2023) revealed that the AI tool, ChatGPT, was most beneficial for the least skilled workers, helping to decrease the performance gap between employees. In this context, the “least skilled workers” are the general company employees who do not have the business expertise, that will be benefitting the most from the chatbot.

Additionally, the choices related to the technical elements of the chatbot reflect the environment of Fannie Mae, where the need for efficient data processing and easy access to mortgage information is key. Choosing TensorFlow, known for its powerful NLP capabilities, is in line with the goal of making the chatbot user-friendly and able to handle complex queries. Similarly, the chatbot being integrated with Oracle databases shows an alignment with Fannie Mae’s existing internal systems so that data queries are run seamlessly.

As per the Social Construction of Technology (SCOT) framework, technologies display “interpretive flexibility” in their early stages, where different social groups can influence their development path (Pinch & Bijker, 2008). The chatbot project experienced this flexibility since multiple stakeholders played a role in defining its use cases. Business associates affected its functionality to make sure it addressed their time-consuming query issues, while feedback from general Fannie Mae employees or end-users could possibly change the chatbot UI and response accuracy. This collaboration will eventually lead to the design of the chatbot stabilizing, making it more and more difficult to change its core structure.

Furthermore, the chatbot’s development brings about broader societal and ethical topics. As highlighted by Hagendorff (2020), AI technologies need rigorous evaluation not just for their functionality but also for their societal impacts. While the Fannie Mae chatbot is mainly a tool for efficiency, it also touches on data privacy, security, and accessibility issues. The fact that it can grant data access to all employees raises questions about data governance and the ethical use

of AI in the finance field. Kazim & Koshiyama (2021) also reinforce this idea, emphasizing the importance of understanding AI ethics from a high-level perspective. They argue that it's very crucial to view AI not just as a tool but as something that re-shapes society's expectations and norms.

Research Question and Method

This leads me to ask the question: How are various stakeholders within the mortgage sector experiencing the integration of AI technologies, and what are the broader societal implications of these changes?

To answer this, I will be conducting in-depth interviews across multiple mortgage companies, targeting a wide spectrum of stakeholders. These interviews will not be limited to experiences with chatbots but will cover all AI technologies. Sachan et al. (2019) discuss an AI-driven decision-support system designed to automate loan underwriting, which is a prime example of another AI technology being used in the mortgage world. I plan on talking with mortgage brokers to understand how AI has changed client interactions and expectations. Information from technical staff, developers, and engineers will be helpful in figuring out the motivations and challenges behind AI implementation. Additionally, discussions with management personnel will reveal some more strategic reasons for adopting AI and their perceived outcomes. Most importantly, the perspectives of end-users and clients will give us a ground-level view of how AI tools have changed their mortgage experience, for better or for worse. The interview questions will be customized for each group (see Tables 1-4), going into their specific experiences with AI, the perceived advantages and challenges, and the effects that they expect on society from this technological advancement.

For Mortgage Brokers
How has AI technology like chatbots changed your interactions with clients?
Can you describe any challenges you've faced while implementing AI into your daily work?

Table 1: Interview Questions for Mortgage Brokers

For Developers and Engineers
What brought about the usage of AI technologies in your workflow?
What technical challenges have you encountered in developing AI?

Table 2: Interview Questions for Developers and Engineers

For Management
What objectives guided the decision to adopt AI in your company?
How do you see AI impacting the mortgage industry in the long term?

Table 3: Interview Questions for Management

For End-Users
How has your experience with getting a mortgage changed with the introduction of AI?
Do you feel that AI technologies have made mortgage processes more user-friendly?

Table 4: Interview Questions for End-Users

Alongside these interviews, I will also run a focused case study on Fannie Mae's adoption of AI. This case study will provide a detailed analysis of Fannie Mae's workflow before and after the implementation of AI tech. It will investigate the reasons behind adopting such AI, the expected outcomes, and the actual impacts that can be seen after implementation. This combined approach, which blends broad industry information from the interviews with the case

study of Fannie Mae, will provide a comprehensive understanding of AI's transformative role in the mortgage sector.

Conclusion

The rise of AI, shown by tools like the Fannie Mae chatbot, demonstrate a shift in the mortgage sector. This research attempts to study the various repercussions of a technology like AI. By looking at both the technical details and the broader societal issues, this study tries to deliver a comprehensive AI narrative. The expected findings will likely offer extremely critical knowledge, potentially guiding future AI technologies, making sure that the technology's benefits are geared towards the advancement of society.

References

- Smil, V. (2005). The Age of Electricity. In *Creating the twentieth century: Technical innovations of 1867-1914 and their lasting impact* (pp. 33–98). Oxford University Press.
- Bostrom, N., & Yudkowsky, E. (2014). The ethics of artificial intelligence. In *The Cambridge Handbook of Artificial Intelligence* (pp. 316-334). <https://nickbostrom.com/ethics/artificial-intelligence.pdf>
- Haenlein, M., & Kaplan, A. (2019). A Brief History of Artificial Intelligence: On the Past, Present, and Future of Artificial Intelligence. *California Management Review*, 61(4), 5-14. <https://doi.org/10.1177/0008125619864925>
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models. Retrieved from <https://arxiv.org/abs/2303.10130>
- Aldehy, M., Shaiba, H., & Bezbradica, M. (2019). Generating and analyzing chatbot responses using natural language processing. *International Journal of Advanced Computer Science and Applications*, 10(9). https://doras.dcu.ie/27514/1/Paper_10-Generating_and_Analyzing_Chatbot_Responses.pdf
- Singh, R., Paste, M., Shinde, N., Patel, H., & Mishra, N. (2018). Chatbot using TensorFlow for small Businesses. 2018 *Second International Conference on Inventive Communication and Computational Technologies (ICICCT)*, 1614-1619. <https://doi.org/10.1109/ICICCT.2018.8472998>
- Goldberg, Y. (2016). A Primer on neural network models for Natural Language Processing. <https://www.jair.org/index.php/jair/article/view/11030>
- Ognjanovski, G. (2020). Everything you need to know about Neural Networks and Backpropagation—Machine Learning Easy and Fun. Medium. <https://towardsdatascience.com/everything-you-need-to-know-about-neural-networks-and-backpropagation-machine-learning-made-easy-e5285bc2be3a>
- Federal National Mortgage Association. (2022). Form 10-K for the fiscal year ended December 31, 2021. United States Securities and Exchange Commission. <https://www.fanniemae.com/media/42821/display>
- Pinch, T. J., & Bijker, W. E. (2008). The social construction of technological systems: New directions in the sociology and history of technology. MIT press.
- Segal, W., & Szymanoski, E. J. (1998). Fannie Mae, Freddie Mac, and the Multifamily Mortgage Market. *Cityscape*, 4(1), 59–91. <http://www.jstor.org/stable/41486468>

Noy, S., & Zhang, W. (2023). Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence. In *Science* 381, (pp. 187-192).

<https://www.science.org/doi/10.1126/science.adh2586>

Hagendorff, T. (2020). The ethics of AI ethics: An evaluation of guidelines. *Minds & Machines*, 30, 99–120. <https://doi.org/10.1007/s11023-020-09517-8>

Kazim, E., & Koshiyama, A.S. (2021). A high-level overview of AI ethics. *Computer Science*, University College London. <https://doi.org/10.1016/j.patter.2021.100314>

Sachan, S., Yang, J.-B., Xu, D.-L., Eraso Benavides, D., & Li, Y. (2019). An explainable AI decision-support-system to automate loan underwriting. *Expert Systems with Applications*. <https://doi.org/10.1016/j.eswa.2019.113100>