Enhancing AI mentorship with Tree of Thoughts and Federated Learning: a privacy-preserving approach

Paul Dickman
Founder and CEO of Ai Ambitions LLC
paulbryand@gmail.com
May 27 2023

Abstract

In a world where AI is increasingly integrated into our daily lives, we are introducing an innovative approach to enhance AI mentors. By combining two existing technologies, namely the Tree of Thoughts (ToT) framework and federated learning, we aim to make AI mentors even more effective. The ToT framework facilitates improved problem-solving capabilities in AI by enabling them to think through problems in a more effective manner. Concurrently, federated learning allows AI mentors to learn from each other without compromising user privacy by securely storing personal data on the user's device. This integrated approach ensures that your AI mentor becomes smarter and more personalized, while maintaining the privacy and security of your personal information.

Introduction

The impact of this new approach to AI mentorship can be transformative across various domains. In **education**, AI mentors can provide personalized tutoring and adaptive learning experiences, tailoring their guidance to the specific needs and learning styles of individual students. This can lead to improved educational outcomes and more efficient use of resources.

In **healthcare**, AI mentors can assist in diagnosing and treating patients, offering personalized recommendations based on medical history, symptoms, and research evidence. They can also provide mental health support, offering empathetic and understanding guidance to individuals struggling with their emotional well-being.

In **professional development**, AI mentors can assist individuals in their career paths by providing tailored advice on skill development, job opportunities, and networking strategies. They can offer insights and resources based on the user's industry, interests, and goals, helping individuals navigate their professional journeys more effectively.

In **personal growth and self-improvement**, AI mentors can serve as virtual coaches, providing guidance on personal goals, motivation, and self-reflection. They can assist individuals in developing new habits, overcoming challenges, and maintaining a positive mindset.

However, **the effectiveness of *current* AI mentors is often limited by their decision-making capabilities. Traditional AI systems typically follow a linear decision-making process, which can be restrictive and may not always lead to the best outcomes.** Furthermore, while AI mentors can learn and improve over time through interactions with users, this learning process is usually isolated. Each AI mentor learns independently, and there is little to no knowledge sharing between different AI mentors.

This paper proposes an innovative approach that combines the Tree of Thoughts (ToT) framework and federated learning to design AI mentors. The **ToT framework enhances problem-solving skills in**

Large Language Models (LLMs) like ChatGPT-4 through four phases: Brainstorming, Evaluation, Expansion, and Decision. During these phases, the AI mentor generates multiple solutions, evaluates them, refines ideas, and ranks the solutions based on evaluations and scenarios. This process guides users towards optimal solutions.

In addition to the ToT framework, federated learning is incorporated into the AI mentor design. Federated learning enables AI mentors to learn from each other without sharing personal user data, preserving user privacy. By leveraging federated learning, the AI mentor network can collectively improve its performance while maintaining data privacy [1][2].

Combining the ToT framework and federated learning empowers AI mentors to explore multiple reasoning paths, make deliberate decisions, and enhance their problem-solving abilities [3]. Simultaneously, federated learning ensures privacy by allowing AI mentors to learn from each other without compromising user data [4]. This integrated approach facilitates the development of personalized AI mentors that can provide effective guidance and support to users while maintaining user privacy and improving overall performance.

In the following sections, we will delve deeper into these concepts and discuss how they can be effectively applied to create more intelligent, adaptive, and privacy-conscious AI mentors. We believe that this approach can significantly advance the field of AI mentorship, paving the way for a new generation of AI mentors that are not only more effective but also respectful of user privacy.

A common complaint about current Large Language Models (LLMs) is that their responses are concise yet formal. In an AI mentorship app, personalized and customized responses are crucial. Users desire guidance that goes beyond generic answers and instead caters to their specific needs, goals, and preferences. Impersonal and formal responses fail to address the unique circumstances and concerns of individual users, leaving them feeling unsatisfied.

Due to the unsatisfactory accuracy of LLMs' zero-shot prompting with standalone questions, others have proposed to improve the distributed synonymous questions using Self-Consistency (SC) and Chain-of-Thought (CoT) techniques [5].

The Tree of Thoughts (ToT) framework is another approach to language model inference that allows for more deliberate decision-making. Unlike traditional language models that follow a linear chain of thought, **the ToT framework enables exploration over coherent units of text, or "thoughts", that serve as intermediate steps toward problem-solving. This allows the model to consider multiple different reasoning paths and self-evaluate choices to decide the next course of action. The ToT framework has been shown to significantly enhance language models' problem-solving abilities on tasks requiring non-trivial planning or search [3].**

Federated learning, on the other hand, is a machine learning approach that allows models to learn from decentralized data sources [4]. **In the context of AI mentors, federated learning can be used to allow different AI mentors to learn from each other without sharing personal user data.** Each AI mentor learns from its interactions with its user, and the learnings (in the form of model updates) are sent to a central server where they are aggregated. **This aggregated model is then sent back to each AI mentor, improving all of them without any user data leaving the user's device.**

User privacy is a critical concern in AI applications. With the increasing use of AI in various sectors, there is a growing need to ensure that user data is protected and used responsibly. In the context of AI mentors, this means ensuring that personal user data is not shared between different AI mentors or with third parties. Federated learning provides a potential solution to this problem by allowing AI mentors to learn from each other without sharing personal user data [4].

While existing AI mentors can learn and improve over time through interactions with users, this learning process is typically isolated. Each AI mentor learns independently, and there is little to no knowledge sharing between different AI mentors. This leads to potential inefficiencies in the learning process and missed opportunities for improvement.

Furthermore, user privacy is a critical concern in AI applications. Current AI mentors learn from user interactions, but this often involves the collection and processing of personal user data. **There is a need for a learning mechanism that allows AI mentors to improve their performance without compromising user privacy.**

Therefore, **the problem this paper addresses is twofold: improving the decision-making and learning capabilities of AI mentors, and ensuring user privacy in the learning process.** The goal is to develop a new generation of AI mentors that are not only more effective in their guidance but also respectful of user privacy.

To address the limitations of current AI mentors, we propose a novel approach that integrates the Tree of Thoughts (ToT) framework and federated learning. **For example, in response to a user query, the AI mentor can generate several potential responses, evaluate the pros and cons of each, and choose the most appropriate one.** This process can be repeated for each decision point, allowing the AI mentor to navigate complex problem spaces more effectively [3].

Federated learning, on the other hand, addresses the issue of isolated learning and user privacy. With federated learning, AI mentors can learn from each other without sharing personal user data. **Each AI mentor learns from its interactions with its user, and the learnings (in the form of model updates) are sent to a central server where they are aggregated. This aggregated model is then sent back to each AI mentor, improving all of them.** This process allows for knowledge sharing and collective improvement while ensuring that user data remains on the user's device.

By integrating the ToT framework and federated learning, we can create AI mentors that are more effective and privacy-conscious. They can make better decisions by considering multiple reasoning paths, learn from each other to continuously improve, and do all this while preserving user privacy. This approach addresses the main limitations of current AI mentors and paves the way for a new generation of more intelligent, adaptive, and respectful AI systems.

As AI mentors interact with users, they inevitably process a wealth of personal data. This data, if not handled properly, could potentially be misused or exploited, leading to privacy breaches. Therefore, it is crucial to ensure that AI mentors are designed with robust privacy and security measures.

Federated learning not only preserves user privacy but also enhances data security. Since user data never leaves the user's device, it is less vulnerable to data breaches. This is particularly important in today's digital landscape, where data breaches are a significant concern.

However, it's important to note that while federated learning significantly enhances privacy and security, it is not a panacea. Additional measures, such as data encryption and secure communication protocols, should be implemented to further protect user data.

Shoham and Rappoport explored the benefits of federated learning and its potential to achieve comparable performance to models trained with centralized data. They highlighted that this finding has significant implications for the practical application of machine learning in scenarios where data privacy is a primary concern [1].

The authors demonstrated their approach using the MIMIC-IV dataset for the next visit prediction task. Their federated learning approach showed improvements in average precision by 4-10 absolute percentage points compared to local models. Furthermore, their approach achieved very close average precision performance to centralized models while preserving data privacy and enabling scalability for multi-center studies.

By integrating federated learning into the design of AI mentors, we can create systems that are not only more effective but also respectful of user privacy and security. This approach aligns with the growing demand for privacy-preserving AI applications and paves the way for more responsible AI development.

One of the key strengths of AI mentors is their ability to provide personalized guidance. Through continuous interactions with users, AI mentors can learn about the users' preferences, learning styles, and needs, and adapt their responses accordingly. This personalization is what sets AI mentors apart from generic AI systems and makes them truly valuable to users. The personalization based on user interactions fosters a stronger relationship, enhancing user satisfaction and engagement.

In our proposed approach, each AI mentor develops a unique personality based on its interactions with its user. The Tree of Thoughts framework allows the AI mentor to consider multiple reasoning paths and make decisions that align with the user's preferences. **Over time, this process leads to the development of a unique AI personality that complements both the users personality and their needs.**

By integrating the Tree of Thoughts framework and federated learning, we can create AI mentors that are not only more effective and privacy-conscious but also truly personalized to each user. This approach represents **a significant advancement in the field of AI mentorship** and has the **potential to greatly enhance the user experience.**

Methods

Evaluating the effectiveness of our proposed solution is crucial to ensure its practicality and impact. Our evaluation plan comprises of both quantitative and qualitative measures, and includes theoretical analysis, simulations, and real-world experiments.

1. **Theoretical Analysis**: We will conduct a theoretical analysis of our proposed solution, examining the potential benefits and limitations of integrating the Tree of Thoughts framework and federated learning. This will involve assessing the potential improvements in decision-making capabilities and learning efficiency, as well as the effectiveness of privacy preservation.

2. **Simulations**: We will run simulations to test the performance of our proposed solution under controlled conditions. This will involve creating synthetic user interactions and observing how the AI mentors make decisions and learn from each other. We will compare the performance of AI mentors using our proposed solution with those using traditional methods to quantify the improvements.
3. **Real-World Experiments**: Finally, we will conduct real-world experiments to test the effectiveness of our proposed solution in a practical setting. This will involve deploying our AI mentors in a beta version of the "Hey Guru" app and collecting user feedback. We will measure user satisfaction, the quality of the AI mentors' decisions, and the rate of learning improvement. We will also monitor for any potential privacy issues to ensure that user data is protected.
4. **Privacy Audit**: To ensure the privacy-preserving nature of our solution, we will conduct a privacy audit. This will involve verifying that no personal user data is shared during the federated learning process and that all data is securely handled.

Through this multi-faceted evaluation plan, we aim to thoroughly assess the effectiveness of our proposed solution and make necessary adjustments to ensure its success. We believe that this rigorous evaluation process is crucial to developing a solution that truly advances the field of AI mentorship.

However, there are limitations to consider. One limitation is the potential trade-off between model performance and privacy preservation [6]. The aggregated model obtained through federated learning may not fully capture the nuances of individual users' preferences and may result in a compromise in personalized guidance. **Striking the right balance between privacy and personalized recommendations will require careful optimization.**

The potential for bias or skewed learning in federated learning arises from the fact that the distribution of user data across different AI mentors may not be uniform. This non-uniform distribution can be attributed to various factors, such as variations in user demographics, user preferences, or the specific tasks each AI mentor is primarily exposed to.

When AI mentors participate in federated learning, they contribute their model updates based on their local user interactions. However, if the distribution of user data is not balanced across the AI mentors, it can lead to biased learning outcomes. AI mentors with access to more diverse or representative user data may have a broader understanding of different perspectives and provide more accurate recommendations compared to those with limited or biased data.

To mitigate potential biases, it is essential to monitor the distribution of user data across AI mentors during federated learning. Techniques such as stratified sampling or data weighting can be applied to ensure fair representation of different user groups or characteristics. By actively monitoring the data distribution, it is possible to identify and address biases that may arise during the learning process.

Furthermore, it is important to consider the potential impact of biases on the personalized recommendations provided by AI mentors. Biased learning outcomes may result in recommendations that are skewed towards certain demographics, perspectives, or preferences, leading to unfair or unbalanced guidance. It is crucial to implement mechanisms to detect and mitigate such biases, ensuring that the AI mentors provide accurate, diverse, and unbiased recommendations to all users.

Addressing bias in federated learning requires a combination of algorithmic approaches, data analysis techniques, and careful design considerations. Ongoing research in the field of fairness, accountability, and transparency in AI can provide valuable insights and methodologies to tackle these challenges. By actively monitoring and mitigating biases, we can ensure that the federated learning process results in fair and accurate learning outcomes, enabling AI mentors to provide unbiased and effective guidance to all users.

In terms of future research directions, there are several avenues to explore. Firstly, further refinement of the Tree of Thoughts framework can enhance its effectiveness in decision-making. Exploring different algorithms or approaches to branching and navigating the tree could lead to more sophisticated problem-solving capabilities.

Additionally, investigating advanced privacy-preserving techniques, such as differential privacy, in the context of federated learning could provide stronger privacy guarantees while maintaining AI mentor performance. Exploring methods to handle user data heterogeneity and balance data contributions across AI mentors can further improve the learning process.

Furthermore, integrating natural language processing techniques, sentiment analysis, or emotion recognition into AI mentors could enhance their ability to understand and respond to users' emotional states, fostering more empathetic and supportive interactions.

Conclusion

In summary, the proposed solution offers several benefits such as enhanced decision-making, privacy preservation, and personalized guidance. While there are limitations to address, the potential for future research and refinement is promising. By continuing to explore these areas, we can advance the field of AI mentorship, creating more intelligent, privacy-conscious, and effective AI mentors that cater to the unique needs of individuals.

The potential impact of our proposed solution is substantial. AI mentors equipped with the ToT framework can provide more nuanced and tailored guidance, adapting to users' preferences and needs. The integration of federated learning ensures that AI mentors can learn from each other without compromising user privacy, fostering collective improvement and diverse learning experiences.

By improving the decision-making and learning capabilities of AI mentors while preserving user privacy, the proposed solution opens up new possibilities in the field of AI mentorship. Users will benefit from personalized guidance that aligns with their unique personalities, preferences, and cultural contexts. The AI mentor network will evolve into a dynamic ecosystem of knowledge sharing, allowing for continuous improvement and adaptation to user needs.

References

1. Shoham, O. B., & Rappoport, N. (2023). Federated Learning of Medical Concepts Embedding using BEHRT. arXiv preprint arXiv:2305.13052.

2. H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. "Communication-Efficient Learning of Deep Networks from Decentralized Data." arXiv preprint arXiv:1602.05629 (2023).

3. Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. "Tree of Thoughts: Deliberate Problem Solving with Large Language Models." arXiv preprint arXiv:2305.10601 (2023).

4. McMahan, B., and Ramage, D. “Federated learning: Collaborative machine learning without centralized training data.” (2017). https://research.googleblog.com/2017/04/federated-learning-collaborative.html, 2017.

5. Xiangyang Liu, Tianqi Pang, and Chenyou Fan. "Federated Prompting and Chain-of-Thought Reasoning for Improving LLMs Answering." arXiv preprint arXiv:2304.13911 (2023).

6. Tânia Carvalho, Nuno Moniz, Pedro Faria, Luís Antunes “Towards a data privacy-predictive performance trade-off”

Expert Systems with Applications,
Volume 223,
2023,
119785,
ISSN 0957-4174,