Perspective-Aware Ai (PAi) for Augmenting Critical Decision Making

1st Marjan Alirezaie *Creative Ai Hub - Toronto Metropolitan University, Canada*
2nd Daniel Platnick *Creative Ai Hub - Toronto Metropolitan University, Canada*
3rd Hossein Rahnama *MIT Media Lab, Cambridge, USA & Toronto Metropolitan University, Canada*
4th Dava J. Newman *MIT Media Lab, Cambridge Massachusetts USA*
5th Alex Paul "Sandy" Pentland *MIT Media Lab, Cambridge Massachusetts USA*

***Abstract*—Perspective-Aware AI (PAi) is a computational innovation in human-AI interaction that allows users to view and interact through each other's perspectives by creating personalized computational models called chronicles. Chronicles capture cognitive and behavioral tendencies from an individual's digital footprint, enabling enhanced decision-making by recognizing and auditing biases. Utilizing federated learning, PAi preserves privacy and ensures data ownership while providing scalability and precision. This approach enhances transparency and clarity in critical decision-making across various domains, including healthcare, education, and business, promoting inclusivity and diverse viewpoints.**

I. Introduction

PERSPECTIVE-Aware Ai (PAi) is a new area of computational innovation in human-Ai interaction that allows users to view and interact through each other's perspectives [1]. PAi involves creating personalized computational models, known as *chronicles*. As shown in Figure 1, chronicles are obtained through a chronological learning process based on an individual's digital footprint. This process captures the cognitive and behavioral tendencies of the individual within various contexts.

PAi is particularly useful for critical decision-making. By leveraging chronicles that represent individual expertise and perspectives, PAi helps identify and consider biases that mirror those in the human mind within given contexts. Recognizing and reasoning about these biases allows PAi to facilitate their auditing which enhances the sharing of expertise and leads to more participatory systems [8][3]. This capability enhances transparency and clarity in critical decision-making in different domains including healthcare, education, politics, journalism, etc., and allows users to discern which opinions or expert perspectives, as represented by their chronicles, should be considered. This approach improves decision-making quality by providing a comprehensive and nuanced understanding of diverse viewpoints and their underlying biases.

A. How to achieve PAi?

Chronicles play an essential role in achieving perspective-aware Ai, serving as reasoning-ready computational models that capture an individual's cognitive as well as behavioral tendencies across multiple contexts and situations. Unlike many user modeling approaches focusing on specific data within limited contexts [4], chronicles aim to be comprehensive digital identity models. They must extend beyond single domain preferences to encompass cognitive processes and behavioral patterns across multiple domains and reflect a user's mentality and personality [5], [6].

The chronicle exists within a pipeline between two key processes: chronicle construction and chronicle utilization (Figure 1). This structure allows for continuous chronological updates through a learning process while maintaining a reasoning-ready structure. An inference engine can then reason about the individual's perspective for other interested (target) users. Assuming the chronicle has a graph-based structure, the construction phase involves graph or structure learning which allows the creation of a dynamic and adaptable model that can evolve with an individual's experiences and behaviors over time [1].

B. What are the challenges?

Structure learning for chronicle construction can be complex and data-intensive by nature [7]. The algorithms must find out relevant variables as well as their relationships; and infer causal structures across diverse domains of human cognition and behavior [1]. As such, complexity grows with more variables and potential interactions, hence necessitating huge amounts of high-quality varied data from numerous contexts to form robust models towards sound chronicles.

While there are several pre-trained multimodal and single-modal models specialized in various types of information that could be employed in our pipeline (Figure 1), we still require lots of data belonging to the same individual to create an accurate digital mental model. However, leveraging individual data in PAi introduces several challenges that are common to traditional computational social systems [2]:

- **Privacy Concerns**: Due to the concerns associated with data misuse, data breaches, or unauthorized access, people usually prefer not to share their data.
- **Data Ownership**: People believe they lose control over their data in centralized systems that are established and maintained by central authorities.
- **Accuracy and Reliability**: The efficiency of the PAi depends on the accuracy and quality of chronicles. If a chronicle generates an inaccurate or wrong result on behalf of its owner, it may raise ethical concerns and

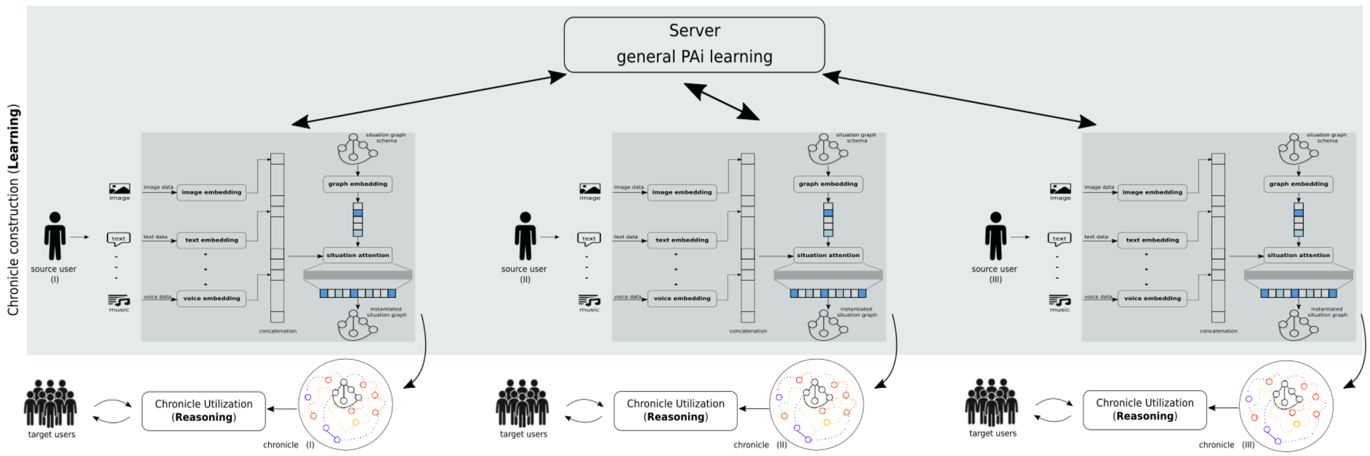

Fig. 1. Chronicle Pipeline Overview: Chronicle construction is achieved through **learning** digital models from multimodal digital footprints of the source individual using federated learning, where data remains decentralized and privacy is preserved. Chronicle utilization is enabled through **reasoning** and interacting with target users (interested in that chronicle).

violate privacy. To make sure that the characteristics of a person are accurately reflected in the system, there should be a thorough assessment and evaluation process.

- **Scalability and Computational Resources**: Building and maintaining chronicles (i.e., as a result of intensive structure learning) for a large number of individuals requires significant resources. Finding a balance between having complex, accurate models and the cost of resources and computing power is an ongoing problem.

C. Distributed Learning as an Enabler for Scalable PAi

To tackle the problems of privacy, data ownership, scalability, and accuracy in PAi systems, we propose a solution that employs distributed and federated learning approaches (Figure 1. This solution aims to create a decentralized ecosystem where individuals, organizations, communities, or any other entity wishing to come up with a sharable interactive digital model can participate while still having control over their data. Our proposed system incorporates data alliance, distributed systems architecture, and federated learning techniques to create a robust, scalable, as well as privacy-preserving framework for PAi. In this architecture, participants form a network where computational resources are shared but data is owned by its owners. Federated learning algorithms allow the creation and maintenance of chronicles without the centralization of sensitive and private information.

This approach addresses the aforementioned issues associated with PAi by:

- Preserving privacy through local data processing & secure model updates.
- Keeping original data with those that created it thus maintaining data ownership.
- Increasing scalability through distributed computation.
- Enhancing accuracy via distributed learning from various sources.

Achieving the proposed distributed architecture will require systematic planning of network architectures, secure communication protocols, and strong model combination methods. However, the distributed and federated learning paradigm presents a potential route in constructing a PAi that preserves the privacy of individuals, promotes data authority, and allows for scalability, and precision. This approach enables the benefits of PAi in critical decision-making while preserving privacy and adhering to data restrictions.

Acknowledgment

The authors would like to thank the team at Flybits, Toronto Metropolitan University, The Creative School, and MIT Media Lab for their support of our ongoing research in this field.

References

[1] M. Alirezaie and H. Rahnama and A. Pentland, *Structural Learning in the design of Perspective-Aware AI Systems using Knowledge Graphs*, AAAI 2024, Digital Human Workshop.

[2] J. Zhang and W. Wang and F. Xia and Y.-R. Lin and H. Tong, *Data-Driven Computational Social Science: A Survey*, Big Data Research, Volume 21, 2020, 100145, ISSN 2214-5796.

[3] H. Rahnama and M. Alirezaie and A. Pentland, *A Neural-Symbolic Approach for User Mental Modeling: A Step Towards Building Exchangeable Identities*, AAAI 2021 Spring Symposium on Combining Machine Learning and Knowledge Engineering, MAKE.

[4] Anders, M.; Obaidi, M.; Paech, B.; and Schneider, K. 2022. A Study on the Mental Models of Users Concerning Existing Software. In Gervasi, V.; and Vogelsang, A., eds., Requirements Engineering: Foundation for Software Quality, 235–250. Cham: Springer International Publishing. ISBN 978-3-030-98464-9.

[5] Nersessian, N. J. 1992. In the Theoretician's Laboratory: Thought Experimenting as Mental Modeling. PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association, 1992: 291–301.

[6] Treur, J.; and van Ments, L., eds. 2022. Mental Models and their Dynamics, Adaptation, and Control: a Self-Modeling Network Modeling Approach. Studies in Systems, Decision and Control (SSDC). Springer Nature Switzerland AG. ISBN 9783030858209.

[7] Zhang, Q., Chang, J., Meng, G., Xu, S., Xiang, S., & Pan, C. (2019). Learning graph structure via graph convolutional networks. Pattern Recognition, 95, 308–318.

[8] Braley, A., Lenz, G.S., Adjodah, D. et al. Why voters who value democracy participate in democratic backsliding. Nat Hum Behav 7, 1282–1293 (2023). https://doi.org/10.1038/s41562-023-01594-w