

Letter from the Special Issue Editors

How cool would it be if we can record and recall every moment in our lives? Then we will never worry about remembering the answers to questions such as “Where did I put my keys?”, “I had some very good Korean hotpot a while back but which restaurant was that?”, or “Does this skirt have higher quality and lower price than the similar-looking one I saw at Macy’s yesterday?” Asking such questions has been a dream for decades and dates back to Vannevar Bush’s MEMEX (MEMory & EXpansion) vision in 1945. Now with the rapid advancements in AI technology, especially in 2023, the Year of AI, are we getting closer to the dream of capturing and organizing personal information smoothly and accessing them effortlessly?

To realize this dream, we face four challenges. First and foremost, we need to be able to **capture personal information**. Scopewise, do we capture mainly digital footprints of a user (messages, browsing and purchase history, and app usage logs), or also physical aspects of a user (what she sees and experiences, and the time and location of those) through special devices? For the latter, what should be the frequency of capturing of data so it answers our needs? Second, we need to be able to **integrate personal data**, often from different sources, organize them in a meaningful way and find a way to store the potentially large volume of captured data. Understanding of the semantics and effective aggregation are the keys, but achieving them is non-trivial. Third, we shall be able to effectively **leverage personal data** to make life easier or memorable for the owner in future; afterall, the data is worthless unless we can effectively use it. The usage includes helping users easily retrieve past memory, and offering recommendations of books, articles, and products that they may enjoy. Solution to this challenge also decides the solutions to previous challenges, such as how we would like to present and store the data, how much data we really need to capture, and whether we shall have an aging mechanism to ensure efficient usage of the past data and limit its effect on future recommendations. Last but not least, **privacy** guarantees are a must-have for trustworthy personal information management, to use it for good, not for harm.

This issue collects a set of papers around personal information management shedding lights on how to address the aforementioned challenges. We start with a paper by **Tran et al.**, describing lifelogging of user activities from both online and offline in the real world, and sharing the learnings from the annual Lifelog Search Challenge Benchmarking Workshops. The logged data include both our digital footprints, and information collected by devices and sensors to track heart rate, sleep quality, locations, and so on. In the second paper, **Kalokyri et al.** focused on digital data and described the YourDigitalSelf project. The project collects users’ digital traces from different applications, integrates them, reconstructs meaningful episodes from them, and supports memory recall using the episodes. The next two papers focus on particular domains and discuss how to build a personal knowledge graph to support applications: **Chakraborty et al.** described Personal Research Knowledge Graph to organize important information for scientists, such as grants, lab equipments, papers, journal and conferences; **Nidhi et al.** described G-VARS for gun violence susceptibility analysis, through information collected on behavior, mental health, access to firearms, social networks and online activities. We next come to how to apply personal information in applications. **Sun et al.** described QALinkPlus, a novel graph-based method that constructs an entity co-occurrence graph derived from QA datasets, describes context QA-specific subgraphs, and leverages this to improve QA, especially personalized QA. The last paper in this issue is an insightful survey User Modeling in the Era of LLMs, where **Tan and Meng** surveyed user modeling leveraging LLM techniques, and pointed our new research directions for personalized search and recommendation.

Overall, the above papers represent an interesting sample of the ongoing work on the new trend of personal information management. We hope that this special issue will further help and inspire the research community in its quest to solve this challenging problem. We would like to thank all the authors for their valuable contributions, as well as Haixun Wang for giving us the opportunity to put together this special issue, and Nurendra Choudhary for his help in its publication.

Xin Luna Dong*, Alon Halevy[†], Shane Moon*
*Meta (Facebook),[†]Amazon