Letter from the Special Issue Editor

Privacy-preserving data management is a critical and timely topic in our increasingly data-driven world. As organizations across various sectors, ranging from healthcare and finance to social media and government, collect and process vast amounts of data, the need to protect individual privacy has never been more crucial. Privacy-preserving data management encompasses a range of techniques and practices designed to ensure that sensitive information remains secure while still allowing data to be used for beneficial purposes such as research and analytics.

The importance of privacy-preserving data management cannot be overstated. On one hand, data drives innovation and growth, enabling breakthroughs in medical research, enhancing customer experiences, and optimizing operations across industries. On the other hand, mishandling or exposing sensitive data can lead to significant harm, including identity theft, financial loss, and erosion of trust. Regulatory frameworks such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA) highlight the need for stringent data protection measures, but compliance alone is not enough. Effective privacy-preserving data management requires a proactive approach, integrating advanced technical solutions with sound governance practices.

This special issue features a collection of six papers from expert researchers that push the boundaries of our current understanding and capabilities, reflecting diverse perspectives on privacy protection in the management of data. We start with a paper by Sohn *et al.* that surveys the latest advancements in secure and private database systems. With the surge in data collection and cloud computing, the authors highlight critical privacy challenges and solutions. They delve into technologies like differential privacy, secure multiparty computation, and zero-knowledge proofs, explaining how these methods protect sensitive data while allowing for meaningful analytics. The paper serves as a practical guide for navigating the complexities of implementing privacy-preserving techniques in database systems. The second article, by He and Zhang, explores the application of differential privacy to provenance data, which are records that describe the history of data, including how it was collected, processed, and used. The authors provide a comprehensive framework that ensures privacy while preserving the utility of the data for various applications. The authors present detailed methodologies and case studies that highlight the effectiveness of their approach in real-world scenarios. Next, we feature a paper by Gavidia-Calderon *et al.* that introduces SQLSynthGen, a method for generating synthetic relational datasets, particularly focusing on healthcare data from National Health Service (NHS) hospitals in the UK. The tool addresses the dual need for data fidelity and privacy, providing a white-box approach to synthetic data generation that includes differential privacy mechanisms for enhanced security. The fourth paper, by Mao *et al.*, provides a comprehensive survey of differential privacy applied to time series data. The paper discusses the unique challenges of protecting privacy in time series due to their volume and temporal correlations. The survey covers various techniques to balance privacy and utility, and it highlights the open challenges and future directions in this evolving field. The fifth contribution, by Monir *et al.*, reviews adaptive techniques in Differentially-Private Stochastic Gradient Descent (DP-SGD), focusing on improving the privacy-accuracy trade-off. It also discusses the data management challenges associated with the deployment of DP-SGD, providing insights into enhancing computational and memory efficiency. Finally, the last article by Islam *et al.* focuses on text data, and examines the intersection of bias, fairness, and differential privacy in NLP models. The article provides an in-depth analysis of how differential privacy impacts the fairness of NLP models and offers strategies to mitigate potential biases, ensuring that privacy-preserving techniques do not inadvertently harm model performance or fairness.

Overall, these works offer new insights, methodologies, and tools that address the pressing challenges that we face in the filed of privacy-preserving data management. We would like to thank all authors for their valuable contributions.

Xiaokui Xiao
National University of Singapore.