Editorial

WELCOME to the fifth issue of the JOURNAL ON SELECTED AREAS IN INFORMATION THEORY (JSAIT), dedicated to "Sequential, active, and reinforcement learning".

There has been a long history of the interplay between information theory and sequential data analysis in the context of sequential estimation, hypothesis testing, and change-point detection. In recent years, sequential methods have become hugely popular in domains such as reinforcement learning, multi-armed bandits, online convex optimization, and active learning. These methods have assumed tremendous importance due to the wide availability of large amounts of data, including data acquired in real time by advanced sensing technologies, and audio-visual content available on social media platforms. The applications are also aplenty—ranging from real-time recommendation systems, to finding distinct changes in time series, and to next-generation video gaming inspired by reinforcement learning. As an example, users of recommendation systems such as *Yelp*, *Netflix*, and *Amazon* provide real-time feedback (e.g., clicks on items they like) to the system. The system then uses these feedback to refine its recommendations to other users that are assessed to be similar to those who have already provided ratings. Another example that is of concern to humankind pertains to the grave problem of climate change. Climate scientists use real-time data to assess when a certain climate phenomenon, such as the increase in wildfires and heatwaves, started in order to understand and mitigate its adverse effects. The last and perhaps most prominent example involves the use of Artificial Intelligence (AI) to beat human players in games. In March 2016, AlphaGo, which is based on a Monte Carlo tree search algorithm, beat the world champion Lee Sedol in a five-game match in Go; this was heralded by the magazine *Science* as one of the scientific breakthroughs of the year.

Although many practical and efficient algorithms have been developed for making various sequential decisions, there is also a strong need for understanding the *fundamental* performance limits of these algorithms designed for sequential learning tasks. Herein lies an excellent opportunity for information theory to leverage the vast array of techniques in its arsenal to answer these fundamental questions. Indeed, this has already been rather prevalent. For example, impossibility results in multi-armed bandit problems leverage converse techniques in decision theory (i.e., hypothesis testing), the change-of-measure technique (which is also used to derive converse error exponent results in Shannon theory), as well as various information-theoretic inequalities such as the Bretagnolle-Huber inequality and Pinsker's inequality. Simultaneously, sequential learning has already started to motivate new problems and insights in information theory and has led to new perspectives. For example, channel and beam selection in wireless communications can be modelled as multi-armed bandit problems (MAB) and techniques in MAB such as upper confidence bound-based algorithms have been employed for these purposes. Thus, we expect that the design of new algorithms in sequential problems can lead to the development of new communication protocols. This special issue seeks to fertilize new topics at the intersection of information theory and sequential, active, and reinforcement learning, and to promote synergies across these areas of research.

Motivated by the above considerations, around April 2020, we proposed this special issue for JSAIT. Upon the completion of this special issue, we believe it has helped to fertilize further new topics at the intersection of information theory and sequential learning. It has enabled researchers in statistical machine learning to appreciate the role of information theory and promote interactions between these classically disparate areas. This special issue has also served its purpose of sustaining and enhancing the interest in sequential learning from the information-theoretic perspective.

We received 31 papers and eventually accepted 22 of them, including a tutorial on sequential change-point detection. The topics of the accepted papers are diverse. They include classical sequential hypothesis testing and detection with a modern view, including using non-parametric approaches. The topics also include active hypothesis testing, robust detection via information projections, and dynamic changes (moving anomalies). Not unexpectedly, a significant number of papers are also devoted to MABs, including decentralized bandits, bandits in communications setting, correlated arms, with fairness and differential privacy considerations. Along similar lines, the special issue includes a set of papers in active learning and classification with abstention, and also to various aspects of reinforcement learning through the lens of coding and optimization. Finally, optimization-based papers involving stochastic approximations and stochastic gradient descent methods form the final bulk of papers in this special issue.

Many people contributed to the success of the special issue, working in exceptionally difficult times due to the COVID-19 pandemic. We would like to thank the JSAIT Editor-in-Chief Andrea Goldsmith for her constant encouragement, pointing us in the right direction in terms of seeking advice from senior researchers, and making sure that various deadlines for this special issue were met. We would also like to thank the Guest Editorial team, Csaba Szepesvári (University of Alberta and Deepmind), Venugopal Veeravalli (University of Illinois at Urbana–Champaign), Mengdi Wang (Princeton University),

Digital Object Identifier 10.1109/JSAIT.2021.3086220 and Zheng Wen (Deepmind). Their tireless dedication and professionalism in view of the time constraints are to be lauded. We would like to thank the authors for submitting their best work to this special issue and the reviewers for their thorough and meticulous reviews to uphold the quality of the papers. Finally, we would like to thank the JSAIT editorial staff including Alison Larkin and Shannon Campos for their help in keeping us on track.

VINCENT Y. F. TAN
Department of Electrical and Computer Engineering
Department of Mathematics
National University of Singapore
Singapore
(e-mail: vtan@nus.edu.sg)

YAO XIE
H. Milton Stewart School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332 USA
(e-mail: yao.xie@isye.gatech.edu)