

Machine Learning for End-to-End Congestion Control

Ticao Zhang and Shiwen Mao

This article presents a selected review of the recent applications of ML in the field of end-to-end congestion control. In this survey, the authors start with a brief review of the relationship between congestion control and ML. They then review the recent works that apply ML to congestion control. These works either help the agent to make an intelligent congestion control decision or achieve enhanced performance.

ABSTRACT

End-to-end congestion control has been extensively studied for over 30 years as one of the most important mechanisms to ensure efficient and fair sharing of network resources among users. As future networks are becoming more and more complex, conventional rule-based congestion control approaches tend to become inefficient and even ineffective. Inspired by the great success that machine learning (ML) has achieved in addressing large-scale and complex problems, researchers have begun to shift their attention from the rule-based method to an ML-based approach. This article presents a selected review of the recent applications of ML to the field of end-to-end congestion control. In this survey, we start with a brief review of the relationship between congestion control and ML. We then review the recent works that apply ML to congestion control. These works either help the agent to make an intelligent congestion control decision or achieve enhanced performance. Finally, we highlight a series of realistic challenges and shed light on potential future research directions.

INTRODUCTION

The rapid development of communication technologies has triggered the emergence of new network architectures such as cognitive radio networks, data center networks, ultra dense heterogeneous networks, and millimeter-wave (mmWave) networks. Each network has its own features and performance requirements, which may change dynamically. The increasing capability of the network also enables a variety of new services and applications, such as augmented reality (AR), online gaming, edge computing, and autonomous driving, entailing more stringent requirements on the communication network.

The transport layer plays an important role in the management of end-to-end connections for upper layer services. The performance of emerging new applications depends heavily on the interactions between the underlying network and the transport layer. End-to-end congestion control, as a fundamental part of the transport layer protocol (TCP), ensures network stability and fairness in resource utilization. Today's TCP congestion control mechanism relies on a design that was created in the 1980s for wired networks. It uses a set of pre-defined rules, for example, halving the congestion window (often referred to as CWND) when a packet loss is detected, and adjusting the CWND

according to measured round-trip time (RTT). The same TCP design and its variants have been employed over the past three decades. Although these congestion control mechanisms achieve great success, they may not perform well in today's or future highly dynamic and complex networks where the network performance is affected by a variety of factors. The congestion control problem can be modeled as an optimization problem, while conventional rule-based methods are mostly heuristics with no guarantee to solve the complex problem. They often lead to sub-optimal solutions, which may suffer poor performance.

Recently, machine learning (ML) has made breakthroughs in a variety of application areas, such as speech recognition, computer vision, and robot control. ML can learn from collected data or the environment and build models. Also, with the recent development of computing infrastructures (e.g., GPU, TPU, and ML libraries) and distributed data processing frameworks, there is now an increasing trend in leveraging complex networking problems. For some tasks, such as regression, classification, and decision making, ML performs pretty well. Considering that these tasks play basic but vital roles in networking problems, it is imperative to embrace ML techniques for potential breakthroughs in end-to-end congestion control.

In [1] the connection between artificial intelligence (AI) and network traffic control (NTC) is discussed. The author shows that AI-NTC could be the next frontier of network research, and deep reinforcement learning (DRL) would be a promising model. Reference [2] investigates how ML could benefit network design and optimization. A workflow for applying ML in the network domain is provided. Reference [3] provides an overview of the state of the art in deep learning, a branch of ML, for intelligent network traffic control systems. The above works are focused on network traffic control with AI [1], networking with ML [2], and network traffic control with deep learning [3]. They do not provide a detailed review on the relationship between ML and the specific problem (i.e., end-to-end congestion control) that is the goal of this article.

In the rest of this article, we first survey the state-of-the-art congestion control algorithms and discuss the challenges. We then introduce the concept of ML and the motivations for ML-based congestion control. We next present the recent advances in ML-based congestion control, and

conclude this article with a discussion of potential future directions.

CHALLENGES AND OPPORTUNITIES

TCP is designed to provide reliable transmission of packets across an end-to-end connection, where congestion control is incorporated. Typical TCP congestion control works as follows. When started, the endpoint should quickly increase its sending rate to achieve high utilization of network resources. However, when congestion is detected, the endpoints involved should reduce their sending rates; when congestion is gone, the endpoints should increase their rates for high utilization of network bandwidth. Usually, the rate keeps on increasing/decreasing, following the network congestion state. Detection of network congestion is usually performed at the network edge without coordination and communication among users, using loss or delay as indicators.

A summary of the major congestion control mechanisms and their pros and cons are provided in [4, Table 2]. These schemes use some signals to detect congestion and then adjust the CWND or sending rate according to some pre-defined rules. Depending on the type of signals that are used as indicators of congestion, congestion control mechanisms can be categorized into several classes. For example, TCP variants Tahoe, Cubic, Reno, and New Reno are loss-based mechanisms. These mechanisms use packet loss to detect congestion. The additive increase multiplicative decrease (AIMD) algorithm is used to adjust the CWND. Loss may be a signal of congestion, but sometimes congestion happens before the bottleneck buffer is full. Early detection of congestion may help to avoid sharp decrease in throughput. Consequently, delay-based mechanisms such as TCP Vegas and Verus have been developed, which use delay as an indicator of congestion. More recently, some hybrid mechanisms try to combine merits of the existing solutions. For example, Compound TCP, which is available on all Microsoft Windows machines, uses the sum of a delay-based window and a loss-based window as its congestion window. Veno, a hybrid between Vegas and Reno, adopts an explicit model of the bottleneck buffer occupancy. BBR, adopted by Google, estimates both bandwidth and RTT, aiming to keep CWND equal to the bandwidth-delay product (BDP), which is the optimal operating point.

We have to point out that such rule-based mechanisms have several limitations:

- *Adapt to new networks:* Congestion control algorithms that are designed for one specific network may not apply to other types of networks. For example, TCP-NewReno was originally designed for wired links where packet loss is often interpreted as congestion. In wireless networks, packet loss may be caused by either link errors or congestion. The sender will always halve CWND even if the loss is caused by link errors. As a result, the link bandwidth will not be effectively utilized. Furthermore, the CWND update mechanism is not effective to adapt to various network topologies. For example, satellite network links where the RTT is large require a more aggressive CWND increase. Meanwhile, mobile ad hoc networks (MANETs), where

the BDP is low, may demand a more conservative CWND increase.

- *Learn from the past:* The rule-based approach uses a fixed set of rules to handle every situation. It does not leverage past experience and assumes no prior information such as link bandwidth, channel characteristics, and the number of flows on shared links. Suppose that NewReno can learn the link bandwidth information from past experience; it could speed up the sending rate more aggressively in the slow start phase to improve link utilization. Inability to learn from the past prevents endpoints from taking actions proactively.
- *Performance:* Today's network is becoming highly dynamic and complex. Ruled-based congestion control is designed based on humans' understanding of the network. As human knowledge may not always accurately characterize the network features, the resulting congestion control mechanism may not be effective and can only achieve sub-optimal performance.

WHY CONGESTION CONTROL WITH ML

ML is a subset of AI, where the machine (or the agent) has the ability to accomplish a task when facing new data or a new environment after a training process. Existing ML techniques generally fall into three categories: supervised learning, unsupervised learning, and reinforcement learning.

- Supervised learning uses a set of labeled samples to learn a mapping between the input and output spaces. Depending on whether the output is continuous or not, supervised learning can be categorized as *regression* and *classification*. Deep neural network (DNN), convolutional neural network (CNN), recurrent neural network (RNN), naive Bayesian (NB), decision tree (DT), and support vector regression are typical techniques in supervised learning.
- Unsupervised learning focuses on classifying unlabeled samples into different clusters. It is mainly used for data dimension reduction in continuous cases or clustering in discrete cases. Restricted Boltzmann machine (RBM), autoencoder (AE), Gaussian mixture model (GMM), principal component analysis (PCA), and *k*-means clustering are usually used in unsupervised learning.
- Reinforcement learning (RL) is an environment-based approach where the agent is trained to solve decision making problems through interactions with the environment. Typical RL techniques include Q-learning, the actor-critic algorithm, and the deep deterministic policy gradient (DDPG) algorithm.

A comparison between supervised and unsupervised learning is given in Table 1. The latest breakthroughs, such as deep learning, generative adversarial networks (GANs), and transfer learning, also have great potential to deal with many complex problems such as end-to-end congestion control.

Conventional congestion control only considers several measurements such as packet loss and/or RTT as indicators of congestion. The entire decision making process relies on these measure-

ML is a subset of AI, where the machine (or the agent) has the ability to accomplish a task when facing new data or a new environment after a training process. Existing ML techniques generally fall into three categories: supervised learning, unsupervised learning, and reinforcement learning.




Figure 1. An illustration of Remy congestion control (figure courtesy of [5]).

	Supervised learning	Unsupervised learning
Training data	Labeled	Unlabeled
Discrete case	Classification	Clustering
Continuous case	Regression	Dimensionality reduction
Accuracy of results	High	Less accurate
Number of classes	Known	Not known

Table 1. Supervised vs. unsupervised learning.

ments and the pre-defined rule based on human understanding of the network. The rule-based mechanism is more susceptible to many unpredictable factors, resulting in poor performance. ML, on the other hand, aims to construct algorithms or models that can learn to make decisions directly from past experience or the network environment. It does not need accurate network models. Hence, it has the potential to outperform the rule-based mechanism.

Nowadays intensive computational resources are provided by both the central controller and edge servers. Cutting-edge technologies such as edge computing, network slicing, and software-defined networking are changing the way network traffic is managed. Moreover, the emergence of several dedicated libraries, such as TensorFlow, Caffe, and PyTorch, has greatly simplified the process of building an ML model. By leveraging these computational resources, ML-based congestion control is becoming feasible. In practice, a model can be trained with global information over a long time interval with the computation resources. The model parameters can be updated one at a time. In the online deployment stage, the decision can be inferred using the trained ML model with a few computations.

RECENT ADVANCES IN ML-BASED CONGESTION CONTROL

Over the past few years, there have been several works on ML-based congestion control. TCP Remy [5] is the first example, where the authors model congestion control as a decision making problem under uncertainty. Each endpoint can make decisions on whether to send packets or not. As shown in Fig. 1, the network model is assumed to be Markovian. The network state consists of three variables: an exponentially weighted moving average (EWMA) of the inter-arrival times of acknowledgments (ACKs), an EWMA of

the sending times of those same packets, and the ratio between the most recent RTT and the lowest measured RTT. The traffic is modeled as a stochastic process that switches unicast flows on or off between sender-receiver pairs. After observing the network state, the agent adjusts the CWND to achieve a balance between high throughput and low latency. This ML-based algorithm outperforms the human-designed end-to-end algorithms, including Cubic, Compound, and Vegas [5]. However, TCP Remy works well at the price of a stringent assumption on the network and traffic models.

Unlike the offline training approach adopted in TCP Remy, Performance-Oriented Congestion Control (PCC) [6] uses online training and does not make similar assumptions on the network model. In each micro-experiment, PCC chooses a sending rate and observes the selective ACK (SACK), which is used to measure the utility of an action (delivery ACK, loss, and latency). PCC runs several micro-experiments continuously, and an empirically optimal rate control policy is learned in an online manner such that the utility can be maximized. PCC achieves 10 times performance improvement, with better fairness and stability, over TCP CUBIC. However, its performance over wireless networks is affected by bufferbloat, and it has not been tested in real-world network environments.

CONGESTION CONTROL WITH DRL

More recently, advanced learning algorithms such as RL and DRL have been incorporated into congestion control design. Just as a human learns a skill from past experience, with RL the endpoint can learn the optimal congestion control policy from its observations of the network environment and past experiences. These mechanisms do not rely on pre-defined rules and exhibit stronger ability to intelligently adapt to the changing environment. DRL embraces the advantages of DNN to train the learning process and to achieve scalability (over RL); thereby, learning speed and learning ability can both be improved.

In congestion control, each endpoint adjusts its sending rate or CWND based on its observations of measured RTT, ACK, and so on. This can be viewed as a decision making problem. Under uncertain and stochastic environments, the decision making problem is usually modeled as a Markov decision process (MDP). The goal is to find an optimal policy for the MDP so that the expected

cumulative reward can be maximized. As shown in Fig. 2, a DRL agent interacts with the network environment by adjusting its actions (e.g., the sending rate or CWND) based on observed network state (e.g., RTT, CWND, and the inter-arrival time of ACKs). A DNN is trained to map state to action so that the reward metric (e.g., throughput and/or latency) can be optimized. Recent works on DRL-based congestion control are summarized in Table 2. QTCP is the first work that leverages RL to design congestion control algorithms [7]. It helps senders gradually learn the optimal congestion control policy in an online manner, without requiring prior knowledge of the network model. Reference [8] presents a DRL-based scheme called Aurora, which uses a fully connected DNN to learn the state-action pairs from stored historical data. This method is surprisingly robust to environments outside the scope of training. TCP-RL, proposed in [9], uses RL to dynamically configure the congestion control parameters for both short and long TCP flows when the environment changes. The performance evaluation is based on a real implementation.

References [10, 11] investigate the multi-path TCP (MPTCP) congestion control problem. A single agent is used in [10] to dynamically and jointly optimize congestion control for all active flows. Apart from incorporating a long short-term memory (LSTM) model in the DRL framework to better capture the network dynamics, [10] for the first time integrates the actor-critic mechanism into DRL for continuous congestion control. In [11], the authors focus on the congestion control performance degradation problem of MPTCP caused by path diversity. For practical deployment, the authors in [11] propose an asynchronous RL-based framework to decouple model training and execution. The congestion control rules are generated offline and applied for real-time online decision making of window adjustment. Reference [12] considers the initial congestion window (IW) selection problem in an mmWave network. A DRL-based online decision making approach is proposed to adjust the IW such that the flow completion time (FCT) is optimized.

PERFORMANCE ENHANCEMENTS

The congestion control problem has its own characteristics and is impacted by many factors. ML techniques, especially supervised learning and unsupervised learning, have high potential to classify packet losses and predict congestion-related parameters for improved performance.

Loss Classification: As mentioned before, the inability to distinguish packet loss caused by congestion and degraded channel quality in wireless networks will result in poor throughput. For next generation cellular networks, where wireless links are more easily blocked, accurately identifying the cause of packet loss will be of vital importance to improve the congestion control performance.

Various ML models have been used to infer the cause of packet loss at the network edge. For example, [13] uses expectation-maximization (EM) to identify the packet losses caused by contention and congestion in optical networks. The classification helps to improve the TCP throughput. Reference [14] develops a loss-predictor-based congestion control mechanism with supervised




Figure 2. Framework of DRL for congestion control.

learning for wireline networks. It achieves a better trade-off of throughput and delay compared to NewReno. Such classifiers usually achieve higher classification accuracy than non-ML approaches. TCP variants built upon these ML-based classifiers have been shown to outperform the standard rule-based solutions.

Congestion Prediction: Congestion prediction plays an important role in dynamic routing, congestion control, congestion avoidance, and proactive network management. In practice, when congestion actually happens, it has already affected the throughput performance significantly, and it may be too late to adopt further actions. If congestion or congestion-related parameters, such as TCP throughput and RTT, can be accurately predicted, the sender can proactively respond to congestion. Currently, there are mainly two directions of research on estimation of congestion-related parameters: formula-based and history-based. Formula-based approaches integrate the sender's measurements, such as RTT, packet loss rate, and CWND, into a formula to generate predictions. However, timely gathering of such information is not easy, especially in today's highly dynamic and complex networks. Moreover, the ever-evolving feature of TCP makes it hard to maintain an accurate formula-based model. History-based approaches refer to some time series analysis techniques, such as the exponential weighted moving average (EWMA) algorithm used by TCP for estimating RTT. However, time series analysis may also be inaccurate in some cases.

To deal with the limitations of conventional mechanisms, ML-based approaches (mainly supervised and unsupervised learning) have been adopted for congestion prediction. Reference [15] applies support vector regression (SVR), a useful ML tool for multivariate regression, to predict TCP throughput. Recently, [14] has developed a loss predictor using random forest, a supervised learning technique, to predict the probability of packet loss caused by congestion. This method can predict and reduce packet loss events, lower the frequency of sending rate reduction, and achieve higher throughput. These works all predict congestion-related parameters from passive measurements with ML approaches, which have great potential in parameter prediction.

Algorithm	States	Actions	Rewards	Pros	Cons
QTCP [7] (2018)	Sending interval, ACK interval, RTT	Select CWND with RL	Large throughput, low latency	1. Generalization to different network topologies, 2. Higher throughput than NewReno	Limited performance evaluation
Aurora [8] (2019)	Latency gradient, latency ratio, sending ratio	Adjust sending rate with DRL	Large throughput, low latency, low packet loss rate	1. Robust to environment, 2. Outperforms BBR, PCC-Vivace and Remy	Limited test on network changes
TCP-RL [9] (2019)	Throughput, RTT, loss rate	Adapt the initial window of connection and the congestion control algorithm	Large throughput, low RTT	1. Dynamical parameter configuration, 2. Improved TCP transmission performance	Untested in highly dynamic network conditions
DRL-CC [10] (2019)	Sending rate, RTT, RTT deviation, goodput, CWND	Select CWND with single agent DRL in MPTCP	Large goodput, fairness	1. Performance improvement in terms of goodput, 2. Flexible and robust, 3. TCP friendly	Large state space increases the complexity
SmartCC [11] (2019)	Sending rate, ACK interval, (subflow)	Select congestion control rules in MPTCP	Large throughput, low RTT, low jitter fairness	1. Addresses path diversity in HetNet, 2. Asynchronous learning framework to reduce overhead, 3. Performance improvement in throughput, RTT, e.t.c.	1. Does not consider TCP friendliness 2. Lack real deployment test
IW-DRL [12] (2019)	FCT, RTT, throughput	Select the initial CWND with DRL	Low FCT	1. Algorithms converges stable and fast, 2. Adaptive on-line decision making, 3. Software-defined networking (SDN) implementation, 4. Flow completion time(FCT) reduction	Does not optimize congestion directly

Table 2. A Summary of ML-based congestion control schemes.

ML-BASED CONGESTION CONTROL WORKFLOW

As real network data is generally hard to obtain and label, we have to point out that RL, which does not require network data, will be a dominant model to implement ML-based congestion control. It has the potential to find the best decision based on *trial and error* and help the endpoints to quickly react to environment changes. Supervised learning and unsupervised learning can leverage historical data for potential performance improvement. For example, [12] combines the merit of both DRL and supervised learning. The system is implemented with DRL for online decision making. A supervised learning approach is utilized to extract features from the collected data during online learning for performance improvement.

Figure 3 illustrates the generic workflow for implementing congestion control with ML. The problem is first formulated as a decision making problem. Depending on the measured signals and the goal, the corresponding state, action, and reward should be properly defined. In the model training process, various training methods can be used to help the model learn the best control policy through interactions with the environment. The training data can be collected, based on which features can be extracted with supervised/unsupervised learning for potential performance improvement. Finally, the model is ready to be deployed in a real environment.

CHALLENGES AND OUTLOOK

ML exhibits great potential for end-to-end congestion control. However, most ML algorithms so far are optimized for simple network topologies. There is no fully implemented and well-tested ML-based congestion control mechanism in a real environment. We outline some potential future research directions here.

REAL DATA COLLECTION

Collecting a large amount of high-quality data along with network profiles is challenging. Most of the existing works rely on simulated datasets generated based on a specific network type. Hence, the resulting ML model may not be compatible with a practical system. For supervised learning, the collection of labeled data may be labor-intensive and costly. It could save a lot of repeated experiments if there are unified open source datasets and a shared standard platform for researchers to compare their developed mechanisms. Moreover, due to the high test cost of large-scale networks and the difficulty in accessing these networks, simulators with high speed and high fidelity are in great demand.

FAIRNESS, ROBUSTNESS, AND GENERALIZATION

ML-based congestion control might be unfair to existing solutions. If ML-based congestion control is trained in an environment where it competes with existing mechanisms, it might learn to occasionally cause packet loss to force other TCP protocols to back off so that it can occupy more network resources. How to coexist properly with other existing protocols is challenging. Moreover, due to specific requirements of network systems, network protocols often require a worst performance guarantee. ML-based congestion control needs to be robust to the rapidly changing environment. Finally, the generalization ability of the ML algorithm is also required to ensure that the model can adapt to network dynamics without retraining the model every time the network condition changes.

CROSS-LAYER OPTIMIZATION

In ML, it is common to use the reward function as the optimization objective. Common reward functions can maximize throughput, minimize latency, reduce loss or fairness, or a combination of these




Figure 3. Workflow of ML-based congestion control.

functions. Since the reward function can be chosen by the designer, some higher-layer application metrics, such as users' quality of experience (QoE), can also be incorporated in the reward function design. For example, it is reported that in 2020, more than 80 percent of Internet traffic is video related, and much of it is carried by TCP. Congestion control algorithms that aim to maximize users' QoE could be of great importance. Cross-layer optimization is a promising research direction.

DEPLOYMENT IN REAL ENVIRONMENTS

Most of the existing DRL-based congestion control algorithms are tested in simulated environments but not in real-world environments. In the future, the rule table needs to be embedded in the operating system kernel, and the ML model needs to be implemented as a system service. Then a prototype can be developed, based on which the adaptation ability of the congestion control algorithms can be tested under dynamic network conditions.

CONCLUSIONS

Today's network is becoming more and more complex, and it is imperative to embrace ML techniques to design effective congestion control algorithms. Despite the advances made in recent works, there is still considerable room for improving network performance by redesigning smart congestion control protocols with ML. There is still a long way to go toward putting ML-based congestion control into practice for some practical issues. In this article, we provide a selected review of the recent advances on ML-based end-to-end congestion control. We also discuss open problems that need to be further investigated from both the networking and ML perspectives.

ACKNOWLEDGMENTS

This work is supported in part by the NSF under Grant ECCS-1923717.

REFERENCES

- [1] J. Xu and K. Wu, "Living with Artificial Intelligence: A Paradigm Shift Toward Future Network Traffic Control," *IEEE Network*, vol. 32, no. 6, Nov. 2018, pp. 92–99.
- [2] M. Wang et al., "Machine Learning for Networking: Workflow, Advances and Opportunities," *IEEE Network*, vol. 32, no. 2, Mar./Apr. 2018, pp. 92–99.
- [3] Z. M. Fadlullah et al., "State-of-the-Art Deep Learning: Evolving Machine Intelligence Toward Tomorrow's Intelligent Network Traffic Control Systems," *IEEE Commun. Surveys & Tutorials*, vol. 19, no. 4, 4th qtr. 2017, pp. 2432–55.

- [4] M. Polese et al., "A Survey on Recent Advances in Transport Layer Protocols," *IEEE Commun. Surveys & Tutorials*, vol. 21, no. 4, 4th qtr. 2019, pp. 3584–3608.
- [5] K. Winstein and H. Balakrishnan, "TCP Ex Machina: Computer-Generated Congestion Control," *Proc. ACM SIGCOMM '13*, Hong Kong, China, Aug. 2013, pp. 123–34.
- [6] M. Dong et al., "PCC: Re-Architecting Congestion Control for Consistent High Performance," *Proc. USENIX NSDI '15*, Oakland, CA, May 2015, pp. 395–408.
- [7] W. Li et al., "QTCP: Adaptive Congestion Control with Reinforcement Learning," *IEEE Trans. Net. Sci. Eng.*, vol. 6, no. 3, July 2019, pp. 445–58.
- [8] N. Jay et al., "A Deep Reinforcement Learning Perspective on Internet Congestion Control," *Proc. 2019 Int'l. Conf. Machine Learning*, Long Beach, CA, June 2019, pp. 3050–59.
- [9] X. Nie et al., "Dynamic TCP Initial Windows and Congestion Control Schemes Through Reinforcement Learning," *IEEE JSAC*, vol. 37, no. 6, June 2019, pp. 1231–47.
- [10] Z. Xu et al., "Experience-Driven Congestion Control: When Multi-Path TCP Meets Deep Reinforcement Learning," *IEEE JSAC*, vol. 37, no. 6, June 2019, pp. 1325–36.
- [11] W. Li et al., "SmartCC: A Reinforcement Learning Approach for Multipath TCP Congestion Control in Heterogeneous Networks," *IEEE JSAC*, vol. 37, no. 11, Nov. 2019, pp. 2621–33.
- [12] R. Xie, X. Jia, and K. Wu, "Adaptive Online Decision Method for Initial Congestion Window in 5G Mobile Edge Computing Using Deep Reinforcement Learning," *IEEE JSAC*, 2019.
- [13] A. Jayaraj, T. Venkatesh, and C. S. R. Murthy, "Loss Classification in Optical Burst Switching Networks Using Machine Learning Techniques: Improving the Performance of TCP," *IEEE JSAC*, vol. 26, no. 6, Aug. 2008, pp. 45–54.
- [14] Y. Kong, H. Zang, and X. Ma, "Improving TCP Congestion Control with Machine Intelligence," *Proc. 2018 Wksp. Network Meets AI & ML*, New York, NY, Aug. 2018, pp. 60–66.
- [15] M. Mirza et al., "A Machine Learning Approach to TCP Throughput Prediction," *IEEE/ACM Trans. Net.*, vol. 18, no. 4, Aug. 2010, pp. 1026–39.

BIOGRAPHIES

TICAO ZHANG received his B.E. degree in 2014 and his M.S. degree in 2017 from the School of Electronic Information and Communications, Huazhong University of Science and Technology, Wuhan, China. He is currently pursuing a Ph.D. degree in electrical and computer engineering at Auburn University, Alabama. His research interests include video coding and communications, and optimization and design of wireless multimedia networks.

SHIWEN MAO [S'99, M'04, SM'09, F'19] received a Ph.D. in electrical and computer engineering from Polytechnic University, Brooklyn, New York, in 2004. He is the Samuel Ginn Professor and Director of Wireless Engineering Research and Education Center at Auburn University. His research interests include wireless networks, multimedia communications, and smart grid. He is a co-recipient of several best paper awards and the best demo award from IEEE conferences. He received the IEEE ComSoc TC-CSR Distinguished Technical Achievement Award in 2019, the NSF CAREER Award in 2010, and the 2004 IEEE Communications Society Leonard G. Abraham Prize in the Field of Communications Systems.

Today's network is becoming more and more complex, and it is imperative to embrace ML techniques to design effective congestion control algorithms. Despite the advances made in recent works, there is still considerable room for improving the network performance by redesigning smart congestion control protocols with ML.