

# Simulating Network Paths with Recurrent Buffering Units

Divyam Anshumaan<sup>1\*</sup>, Sriram Balasubramanian<sup>1,2\*†</sup>, Shubham Tiwari<sup>1</sup>  
Nagarajan Natarajan<sup>1</sup>, Sundararajan Sellamanickam<sup>1</sup>, Venkata N. Padmanabhan<sup>1</sup>

<sup>1</sup>Microsoft Research India

<sup>2</sup>University of Maryland, College Park

t-danshumaan@microsoft.com, sriramb@cs.umd.edu, t-shutiwari@microsoft.com, nagarajn@microsoft.com, ssrajan@microsoft.com, padmanab@microsoft.com

## Abstract

Simulating physical network paths (e.g., Internet) is a cornerstone research problem in the emerging sub-field of AI-for-networking. We seek a model that generates end-to-end packet delay values in response to the time-varying load offered by a sender, which is typically a function of the previously output delays. The problem setting is unique, and renders the state-of-the-art text and time-series generative models inapplicable or ineffective. We formulate an ML problem at the intersection of dynamical systems, sequential decision making, and time-series modeling. We propose a novel grey-box approach to network simulation that embeds the semantics of physical network path in a new RNN-style model called Recurrent Buffering Unit, providing the interpretability of standard network simulator tools, the power of neural models, the efficiency of SGD-based techniques for learning, and yielding promising results on synthetic and real-world network traces.

## 1 Introduction

Network simulation provides a cost-effective way of developing and evaluating networking applications (e.g. video-conferencing) and protocols. It is a cornerstone research problem, recognized as such by the networking community, with applications in AI-for-networking (Wei, Gu, and Li 2021).

Network simulation entails delaying or dropping the data packets traversing a sender-receiver network path appropriately. The sender  $S$  typically adapts its sending rate continuously based on the feedback in terms of delays or drops gleaned from the packets sent previously. For the simulation to be realistic, the packet delays and drops produced by the simulation mechanism should reflect the target network conditions faithfully, both at the microscopic and the macroscopic levels, so as to recreate application-level metrics such as throughput distribution.

Widely-used network simulation tools such as ns-3 (NSNAM 2011) require configuring with a certain network topology, link bandwidth, cross-traffic, etc., typically performed manually by networking domain experts. However, it is extremely challenging to ensure realism in such a

manual approach. State-of-the-art (SOTA) data-driven configuration techniques (Yan et al. 2018; Ashok et al. 2020) try to mitigate this challenge, but (a) they do not accommodate real-world network behaviors like packet reordering, and (b) scale poorly as they rely on black-box optimization (Section 2).

In this work, we formulate and study a novel ML problem of simulating a target network path. The goal is to respond to the sending protocol’s actions with *realistic* delay values for every packet, just like the target network would. Note that the sending protocol is provided as input and it could be very different at test vs train. Formally:

**Definition 1** (End-to-end network path simulation). *Let  $(\Pi, \mathcal{N})$  denote the traces collected using a sending protocol  $\Pi$ , over a target network path  $\mathcal{N}$  between a sender  $S$  and a receiver  $R$ , i.e.,  $S \xrightarrow{\mathcal{N}} R$  (e.g., the path between a cloud server and a cellular client in certain locations, during the peak hours of a day). We seek a model  $\hat{\mathcal{N}}$  that simulates  $S \xrightarrow{\hat{\mathcal{N}}} R$  s.t. for a new and previously unseen protocol  $\Pi'$ , the simulated traces  $(\Pi', \hat{\mathcal{N}})$  “closely match” the ground-truth  $(\Pi', \mathcal{N})$ , that would be obtained if we took the trouble of actually running  $\Pi'$  too on the same path  $S \xrightarrow{\mathcal{N}} R$  under identical network conditions. The match is in terms of the metrics that networking applications care about; e.g., the joint delay and throughput distribution.*

This problem poses the following key challenges:

**(A1) Reactive inputs at test time:** At test time, the decisions made by the protocol (e.g., new sending rate), forming the input to the model  $\hat{\mathcal{N}}$ , are a *response* to the delays output by the model. So, we cannot expect the entire input sequence to be available ahead of time, unlike in standard predictive (Rangapuram et al. 2018; Salinas et al. 2020) or generative modeling settings (Esteban, Hyland, and Rättsch 2017; Fu et al. 2019; Smith and Smith 2020).

**(A2) Unseen test protocols:** At test time, the behaviour of inputs to the model (governed by  $\Pi'$ ) can change drastically from that of the training time (governed by  $\Pi$ ), as it depends on *how* the protocol  $\Pi'$  reacts to the (simulated) delays.

**(A3) Non-trivial success metric:** The stochastic nature of the network setting means that the metrics of interest are distributional, e.g., joint delay and throughput distribution. Unlike in sequential decision making, we cannot attach a

\*These authors contributed equally.

†Work partially done as a Research Fellow at Microsoft Research India

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

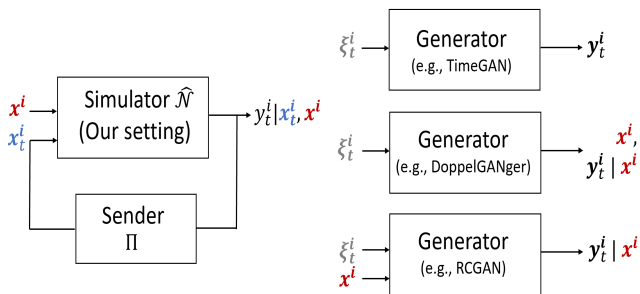


Figure 1: Our simulation (left) vs standard generative modeling settings (right); subscript  $\cdot_t$  denotes time  $t$  (absence denotes static feature), and superscript  $\cdot^i$  denotes series  $i$ .

reward to a given output sequence.

We address this challenging problem (Definition 1) by focusing on modeling the behavior of the target path  $\mathcal{N}$ , using domain-aware neural models, rather than on the actions of the sender protocol that could change drastically at test time. We develop a (conditional) generative modeling technique inspired by network simulation tools like ns-3 that mimic the components of the physical network. A key aspect is explicitly modeling the *unobservable* cross-traffic which competes for resources on the same network path  $\mathcal{N}$  and so critically influences the observed delays.

**Contributions:** We make three key contributions:

- (1) Novel ML formulation of end-to-end network path simulation — has significant applications in the development of networking algorithms (Wei, Gu, and Li 2021).
- (2) A grey-box approach to network simulation that embeds the semantics of physical network path in a new RNN-style model called Recurrent Buffering Unit or RBU (Section 3) — provides the interpretability of simulator tools and the expressive power of neural models.
- (3) Efficient and practical solution — scales to sequences of length tens of thousands (leverages domain-specific insights for training in Section 4), orders of magnitude more than what the SOTA time-series GAN techniques (Yoon, Jarrett, and van der Schaar 2019; Jarrett, Bica, and van der Schaar 2021) can handle, yet produces realistic traces in synthetic and real-world network settings (Section 5).

**Related Work:** We highlight the relevant ML work here (and revisit some of these in Section 2).

*Generative models for time-series:* SOTA techniques for generating time-series data use RNNs with a GAN-like objective (Esteban, Hyland, and Rättsch 2017; Lin et al. 2020; Xu et al. 2020; Yoon, Jarrett, and van der Schaar 2019) or imitation learning (Jarrett, Bica, and van der Schaar 2021). While they account for longer-range temporal dynamics, error compounding, conditioning on static meta-data, they do not handle (A1), or scale to very long sequences. Also, evaluation metrics like Maximum Mean Discrepancy, discriminative scores, etc. used in these works are secondary to our domain-specific metrics in (A3).

*Generative models for text:* In the language domain, recent work have used LSTMs (Sutskever, Vinyals, and Le 2014;

Sutskever, Martens, and Hinton 2011) or Transformers (Radford et al. 2018, 2019) to complete or generate sequences, given a context. Indeed the GPT-class models have shown impressive performance in language understanding and text generation tasks, leveraging the self-attention idea in the decoder to capture temporal and positional dependencies while eschewing recurrences of RNNs. However, as with recurrent networks, scaling to very long sequences and obeying domain-specific constraints continue to persist with Transformers, as we observe in our evaluation (Section 5).

*Sequential decision making/RL:* Our problem has the flavor of sequential decision making in (A1). RL formulations applied to such problems (Levine et al. 2020; Ranzato et al. 2016) focus on maximizing expected rewards over multiple trials, which doesn’t apply to our setting as stated in (A3). Frameworks like imitation learning are also infeasible because of the lack of interactive access to the target  $\mathcal{N}$ .

## 2 Problem Setup, Background, & Challenges

A network trace, collected using a sender  $S$  (e.g., file transfer, video call) over a physical network  $\mathcal{N}$ , is packet-level time-series of measurements  $(\mathbf{x}_t, y_t), t = 0, 1, \dots$ , where  $\mathbf{x}_t \in \mathbb{R}^d$  denotes the “input features” for packet  $t$  (e.g., inter-packet spacing  $s_t$ , packet sizes) characterizing the load offered to  $\mathcal{N}$  by  $S$ ; and  $y_t \in \mathbb{R}_{>0}$  the end-to-end delay experienced by packet  $t$ , with the convention  $y_t = \infty$  when packet  $t$  was dropped and so never delivered to the destination. Typically,  $S$  runs a protocol  $\Pi$ , e.g., TCP Cubic (Ha, Rhee, and Xu 2008), for adapting the sending rate based on the feedback it gets in terms of delays and drops experienced by the preceding packets. We denote a set of such traces by  $(\Pi, \mathcal{N})$ . Note that  $\mathcal{N}$  is a complex black-box system, and we treat it as such. In addition to the **packet-level** features  $\mathbf{x}_t$ , we also use 3 **static features**, denoted by  $\mathbf{x}$  (dropping the subscript  $\cdot_t$ ), to model  $\hat{\mathcal{N}}$ : 1) the minimum delay or  $y_{\min}$  (approximating the end-to-end network propagation delay), 2) the maximum delay or  $y_{\max}$ , and 3) the 95th percentile throughput (approximating the bottleneck link bandwidth).

**Setup:** As in Definition 1, we seek a model  $\hat{\mathcal{N}}$ , using  $(\Pi, \mathcal{N})$  for training, that helps produce realistic end-to-end delays, like the actual physical network path  $\mathcal{N}$  between a sender  $S$  and a receiver  $R$  would, even for a new sender protocol,  $\Pi'$ , at test time.  $\hat{\mathcal{N}}$  can be deployed for evaluating new protocols, which may be disruptive or infeasible to perform on the target  $\mathcal{N}$ . The “goodness” of the model  $\hat{\mathcal{N}}$  is determined by how well the application metrics such as the distribution of packet delays and throughput, computed over the simulated traces for the unseen protocol  $\Pi'$ , i.e.,  $(\Pi', \hat{\mathcal{N}})$ , match the ground-truth  $(\Pi', \mathcal{N})$ . While obtaining the ground-truth is challenging in general, it is feasible in controlled settings to enable comparison (described in Section 5).

**How network simulation is done today:** The widely-used solution for network simulation is to use frameworks like ns-3 (NSNAM 2011) that implement the mechanism of physical network components (like links, buffers, end-points) in software. But, it is very challenging to configure them to reflect the target network conditions. Recent efforts (Yan et al. 2018; Ashok et al. 2020) learn the model  $\hat{\mathcal{N}}$  using a simple

abstraction of physical network paths (Figure 3) and domain knowledge-based heuristics. While they show promise for realistic simulation in some settings (Section 5), they (a) are fairly rigid in the type of networks they can model; for instance, they do not accommodate events like link failures, or packets arriving out of order at the receiver, and (b) rely on black-box optimization techniques (because they work with the ns-3 tool directly), e.g., Bayesian Optimization, which makes it challenging to scale. Ashok et al. (2020) also briefly discuss the challenges of using neural formulations (covered by LSTM-based baselines in Section 5) for network simulation, which we tackle in our work.

**Inadequacy of domain-agnostic models:** Consider a typical network trace in Figure 2 obtained using ns-3 tool, configured with a simple topology (in longer version (Anshumaan et al. 2022)), and peak bandwidth of 7.8 Mbps, constituting  $\mathcal{N}$ . Several observations are in order, from left to right in the Figure. First, the sending rate, regulated by the TCP Cubic protocol at  $S$ , stabilizes around the peak bandwidth, after a brief initial “exploration”, characteristic of the protocol. Second, the delays  $y_t$  build up to a peak value of around 0.5 seconds, due to the sender behavior *as well as* (unobserved) cross-traffic along  $\mathcal{N}$ , which together start filling up the bottleneck link buffer. Third, zooming into the start of the trace, the sending rate increases swiftly, and fourth, delay builds up indicating congestion, leading to packet drops once the bottleneck link buffer has been filled up. This example illustrates the *global* (first two plots) and *local* (last two plots) behaviors observed in real-world network traces. These behaviors critically influence the decisions made by the protocol, and in turn the evolution of the network trace, and the application-level metrics. So, learning the model  $\hat{\mathcal{N}}$  entails learning the structure *and* stochasticity in the end-to-end delays  $y_t$  conditioned on the inputs seen until and including  $\mathbf{x}_t$  as well as the previous delays output by  $\hat{\mathcal{N}}$ .

At a first glance, this resembles auto-regressive time-series formulations studied in predictive (Borovykh, Bohte, and Oosterlee 2017; Rangapuram et al. 2018) and generative settings (Sutskever, Martens, and Hinton 2011; Graves 2013). These techniques factorize the joint  $P(y_{1:T}|\cdot)$  into a product of conditionals  $\Pi_t P(y_t|\cdot)$ , and use RNN-based or Transformer-based (Radford et al. 2019) models to learn a parameterized distribution for the conditional (e.g., multinomial or Gaussian). MLE-based training of these models helps learn  $P(y_t|\cdot)$  with a low step-wise loss *in expectation*, but there is no guarantee that *samples* satisfy nuanced dynamics observed in network traces. While such modeling could capture the high-level structure in sequences if carefully trained (Bengio et al. 2015; Ranzato et al. 2016), they fail at capturing the micro-level characteristics (see Section 5), that we articulated using Figure 2.

On the other hand, GAN techniques that directly yield samples (Yoon, Jarrett, and van der Schaar 2019; Xu et al. 2020; Jarrett, Bica, and van der Schaar 2021) are meant for generating synthetic data, which is different from our setting (see Figure 1). They do not explicitly model the conditional dynamics (A1), and scale poorly with the sequence length — training TimeGAN (Yoon, Jarrett, and van der Schaar 2019)

to synthesize sequences of length 600 takes over a week on a V100 GPU, with their TensorFlow code. In contrast, network traces are 50x – 100x longer.

It is quite unclear whether domain-agnostic neural models can capture the fine-grained behaviors in network traces. In the next section, we show how we address the challenges by priming the neural model with domain knowledge in the form of queuing dynamics of real network paths.

### 3 Proposed Model: Recurrent Buffering Unit

It seems unlikely, a priori, that the problem posed in Definition 1, in the face of challenges (A1)–(A3), can be solved satisfactorily, even under some assumptions on the sender protocols. The promise comes from a growing body of research underscoring the importance of incorporating the knowledge of physical systems and processes in neural models (Li et al. 2020; Xu, Pradhan, and Duraisamy 2021; Beucler et al. 2021). Especially, to tackle (A2), it is imperative that we model the behavior of the target path  $\mathcal{N}$ , rather than the network responses to the (observed) actions of the sender protocol — which could be very different at test time. Also, any acceptable model for simulating  $\mathcal{N}$ , in terms of domain-specific metrics (A3), should preserve path dynamics at the level of *consecutive* packets. For instance, we want the delays,  $y_t$  and  $y_{t+1}$ , imposed on packets,  $t$  and  $t + 1$ , to ensure that these packets delivered at the receiver  $R$  are spaced apart in accordance with the bottleneck bandwidth, i.e., a higher (lower) bandwidth would mean a shorter (longer) inter-packet spacing at  $R$ —otherwise, packets  $t$  and  $t + 1$  being delivered arbitrarily close to each other in time would imply impossibly high available bandwidth for  $S$ .

We appeal to how network simulation tools preserve path behaviors and physical constraints *by construction*, i.e., by implementing, in code, the semantics of physical network path composed of links, buffers, and nodes. As we saw in Section 2, the key difficulty in working with such tools is to appropriately configure them. Our first technical contribution is, in essence, to turn the (discrete) simulator tool into a *learnable* model via deriving an end-to-end differentiable formulation.

Consider an abstraction of the physical path  $S \overset{\mathcal{N}}{\rightsquigarrow} R$  in Figure 3. For clarity, we consider a single bottleneck link (where the path is most constrained in terms of bandwidth) of unknown bandwidth  $B$  and a FIFO (First-in First-out) queue of unknown buffer size  $\tau$ , as in (Yan et al. 2018; Ashok et al. 2020). Later in the section, we extend the ideas to multi-path networks, which, among other things, allows us to accommodate phenomena such as packet re-ordering.

The end-to-end delay,  $y_t$ , suffered by the packet  $t$  along the network path in Figure 3 admits a nice structure, comprising (1) the end-to-end propagation delay of  $S \overset{\mathcal{N}}{\rightsquigarrow} R$ ,  $d_{\text{prop}}$ , arising from the speed of light, (2) the “transmission delay”,  $d_{\text{trans}} \propto 1/B$ , or the time taken to transmit a packet onto the network link, and (3) the “queuing delay”  $d_{\text{queue}}$ , or the time spent by packet  $t$  in the buffer, waiting for its turn to be transmitted. In other words,  $y_t = d_{\text{prop}} + (d_{\text{trans}} + d_{\text{queue}})$ .

In the rest of the discussion, we define  $d_{\text{trans}} + d_{\text{queue}}$  to be the “bottleneck delay”,  $d_t$ , for packet  $t$ , with the subscript  $\cdot_t$ ,




Figure 2: Global (left two) and local (right two, zoomed into the first 10% packets) characteristics of a typical network trace.




Figure 3: Abstraction of a physical network path.

as the transmission delay and the queuing delay would, in general, vary from packet to packet depending on the packet size and the length of the queue encountered by a packet.

The key component of stochasticity affecting  $d_t$  are the (unobserved) competing cross-traffic packets  $c_t$  also filling the buffer, marked by shaded regions in the figure.

So, the bottleneck link abstraction of network paths is specified by  $d_{\text{prop}}$ ,  $d_{\text{trans}}$ ,  $\tau$ , and the dynamic cross-traffic  $c_t$ .

**Modeling parameters**  $d_{\text{prop}}$ ,  $d_{\text{trans}}$ ,  $\tau$ : While it is possible to estimate these parameters from offline traces using simple heuristics in some cases, such estimates can be grossly inaccurate for real-world traces. So, we model the 3 parameters, with a bounded sigmoid function, in terms of static features and the packet size (details in Section 4).

**Modeling cross-traffic  $c_t$** : Cross-traffic typically arises from other senders whose traffic flows via the same bottleneck link buffer. We model the (random) cross-traffic  $c_t \in [0, 1]$  as a fraction of the *remaining* buffer space occupied when packet  $t$  from  $S$  arrives. Noting that cross-traffic could react to changes in the network state just as the sender  $S$ , we model  $c_t$  via a non-linear dynamical system:

$$c_t = \sigma(\langle \mathbf{w}_c, \mathbf{h}_{t-1} \rangle), \quad \mathbf{h}_t = \sigma(\langle W_h, \cdot \rangle + \langle U_h, \mathbf{h}_{t-1} \rangle), \quad (1)$$

where  $\mathbf{h}_t$  that models the local state of the network path at time  $t$  is a standard RNN with weight matrices  $U_h$  and  $W_h$ ; the input  $\cdot$  to this RNN comprises packet features  $\mathbf{x}_t$  and the global state of the path as we will see in Section 4.

**Modeling bottleneck delay  $d_t$** : We derive  $d_t$  using the FIFO buffer dynamics. Let  $s_t$  denote the delta between the sending timestamps of packets  $t$  and  $t-1$ , i.e., the inter-packet spacing at  $S$ . We exploit the following mutually exclusive conditions, when packet  $t$  arrives in the buffer. If there are no other packets from  $S$  ahead in the queue, i.e.,  $s_t \geq d_{t-1}$ , then  $d_t$  is proportional to the cross-traffic in the queue; else, the

packet  $t-1$  has not yet drained, and the additional delay  $a_t$  accrued by packet  $t$  is  $d_{t-1} - s_t$ . With  $c_t$  modeled as in (1) and  $\text{ReLU}(z) = \max(z, 0)$ , we have:

$$a_t = d_{\text{trans}} + \text{ReLU}(d_{t-1} - s_t), \quad d_t = a_t + c_t(\tau - a_t). \quad (2)$$

Note that when  $s_t \geq d_{t-1}$ ,  $d_t$  increases with the fraction  $c_t$  of cross-traffic occupying the buffer of capacity  $\tau$ ; when  $s_t < d_{t-1}$ ,  $d_t$  increases with the fraction  $c_t$  of cross-traffic occupying the buffer of (shrunk) capacity  $\tau - a_t$ .

**Modeling the output:** The end-to-end delay  $y_t$  and the packet drop probability  $p_t$  are given by:

$$y_t = d_t + d_{\text{prop}}, \quad \text{and} \quad p_t = \sigma(d_t - \tau), \quad (3)$$

where  $d_t$  is given by (2) and  $\sigma(\cdot)$  is the sigmoid function.

**Recurrent Buffering Unit:** The proposed Recurrent Buffering Unit (RBU) for modeling end-to-end network delays and packet drops (Figure 4, right) is given by recurrences (1), (2), (3).

**Proposition 1.** RBU preserves the semantics of single-bottleneck link network path in Figure 3, when  $\sigma$  in (3) is the step function. That is, for any two packets originating at  $S$  at timestamps  $t$  and  $t'$ , with  $t < t'$ , they are delivered in order at  $R$ , i.e., their output delays satisfy  $t + y_t < t' + y_{t'}$ , or packet  $t$  is dropped.

**Proof:** It suffices to show that  $t' + d_{t'} > t + d_t$ . Note that  $t' \geq t + s_{t'}$ . If we show the inequality for  $t' = t + s_{t'}$ , when  $t'$  is indeed the immediate next packet, then it holds for all future packets. In this case, it reduces to showing  $s_{t'} + d_{t'} > d_t$  or  $d_{t'} > d_t - s_{t'}$ . Now, from (2), we have either  $d_{t'} \geq a_{t'}$  or packet  $t$  is dropped, in which case we are done, since  $a_t > \tau$  and  $c_t \in [0, 1]$  together imply  $d_t > \tau$ . So,  $d_{t'} \geq a_{t'} = \max(d_t - s_{t'}, 0) + d_{\text{trans}} > \max(d_t - s_{t'}, 0) \geq d_t - s_{t'}$ , which proves the claim.

**Multi-path networks:** In reality, links can fail momentarily over the course of a call, or packets may be randomly routed via different paths and arrive out of order at the receiver. To this end, we consider a generalized multi-path abstraction with a bottleneck link along each path  $k$  parameterized by  $d_{\text{trans}}^{(k)}$  and  $\tau^{(k)}$ . Consider a packet that enters queue  $k$ . To model the dynamics here, we need to keep track of the time elapsed between the immediately preceding packet that entered the same queue  $k$  and the current packet. We

denote this quantity by  $s_t^{(k)}$ , and update it using the recurrence:  $s_0^{(k)} = 0$  and  $s_t^{(k)} = s_{t-1}^{(k)} + s_t$ . With  $c_t$  as in (1), the bottleneck delay for the queue  $k$  that packet  $t$  enters is given by recurrences analogous to (2):

$$a_t^{(k)} = d_{\text{trans}}^{(k)} + \text{ReLU}(d_{t-1}^{(k)} - s_t^{(k)}), \quad (4)$$

$$d_t^{(k)} = a_t^{(k)} + c_t(\tau^{(k)} - a_t^{(k)}) \quad (5)$$

For all other queues  $k' \neq k$  at time  $t$ :  $a_t^{(k')} = a_{t-1}^{(k')}$ , and  $d_t^{(k')} = d_{t-1}^{(k')}$ . Then, we obtain  $y_t$  and  $p_t$  from (3) with  $d_t = d_t^{(k)}$  and  $\tau = \tau^{(k)}$ . Finally, we reset the accumulative elapsed time of the queue  $k$ , i.e.,  $s_t^{(k)} = 0$ , as the current packet has entered queue  $k$ .

## 4 Training and Inference

There are three key challenges in learning the RBU model, using the traces  $(\Pi, \mathcal{N})$ .

**(C1) Very long traces.** Traces are extremely long in general (tens of thousands of packets). Trying to jointly learn all the model parameters in the recurrence relations, even in the single bottleneck link case, (1), (2) and (3), using only the observed end-to-end delays in the traces, may be ill-posed. Working with aggregated or sub-sampled traces is out of question in the simulation setting, unlike in synthetic data generation setting of GANs, because we need to respond at a packet-level to the sender protocol at test time. On the other hand, we could divide the (packet-level) traces into independent chunks that are sufficiently small amenable to efficient training. However, the independence of the chunks means that the global temporal structure of traces (e.g., slow build-up of congestion and surges in cross-traffic) is not captured. We address this challenge using a two-level architecture that preserves both the global and fine-grained characteristics of traces, yet being computationally and sample-efficient. We train a packet-level model within the confines of individual chunks, but the global structure of the traces is integrated via a coarser, window-level model.

**(C2) Cross-traffic  $c_t$  estimation.** Despite the structure RBU model imposes on the delays unlike a vanilla RNN, training the model using standard MLE techniques (Graves 2013; Borovykh, Bohte, and Oosterlee 2017; Salinas et al. 2020) does not perform well for our problem, owing to the absence of direct feedback, especially  $c_t$  (1). Using domain knowledge and the global trace structure via the window-level model, we estimate packet-level  $c_t$  robustly.

**(C3) Discrete path selection.** In the multi-path scenario, recurrences involve a discrete step of selecting a queue  $k$  for packet  $t$ . This introduces discontinuity in the model at training. In our implementation, we use a smoothed version of recurrences.

**Leveraging global structure:** Consider the step-wise loss (i.e., packet-level), in the spirit of auto-regressive formulations (Graves 2013; Borovykh, Bohte, and Oosterlee 2017; Salinas et al. 2020), to learn the RBU model  $\Theta_{\text{RBU}}$ :

$$J_{\text{pkt}}(\cdot; \Theta_{\text{RBU}}) := \sum_{i=1}^N \sum_{t=1}^{T_i} \ell_{\text{pkt}}((\hat{y}_t^{(i)}, \hat{p}_t^{(i)}), y_t^{(i)}), \quad (6)$$

where  $(\hat{y}_t^{(i)}, \hat{p}_t^{(i)})$  denote the delay and drop probability for packet  $t$  in trace  $i$  predicted by RBU, and  $\ell_{\text{pkt}}$  is the loss in (11). We can apply SGD to minimize (6), with standard tricks like chunking, mini-batching, and (truncated) back-propagation through time (Sutskever 2013). But, this performs poorly in our evaluation (Section 5) given extremely long traces.

We devise a two-level architecture (Figure 4) that helps mitigate the issues. We use a coarser window-level model to obtain an embedding of the global state of the network path, which then provides crucial feedback on cross-traffic  $c_t$  for the packet-level RBU model. For the window-level model, we use (2-layer) LSTM, parameterized by  $\Theta_{\text{window}}$  and static trace features  $\mathbf{x}$  as input, to compute an embedding  $\mathbf{h}_w$  of the path  $S \xrightarrow{\mathcal{N}} R$  state. The model operates over fixed-length (100 ms, in experiments, corresponding to the round-trip time on typical network paths when the global state could change), non-overlapping windows:

$$\mathbf{h}_w = \text{LSTM}(\mathbf{h}_{w-1}, \mathbf{x}; \Theta_{\text{window}}). \quad (7)$$

**Estimating cross-traffic:** We estimate the *expected* fraction of cross-traffic filling the bottleneck buffer, denoted  $c_w \in [0, 1]$ , using a linear layer over the global path state  $\mathbf{h}_w$  with sigmoid activation. We modify the packet-level  $c_t$  in (1) to incorporate this information with a hyper-parameter  $\gamma \in [0, 1]$ :

$$c_t = (1 - \gamma) c_w + \gamma \sigma(\langle \mathbf{w}_c, \mathbf{h}_{t-1} \rangle). \quad (8)$$

Furthermore, we obtain a crude estimate of  $c_w$  from the training data, by inverting the RBU recurrences to approximate  $c_t$ ; to do so, we use  $\gamma = 0$  in (8) and simple heuristic estimates (Ashok et al. 2020) for  $d_{\text{trans}}$ ,  $d_{\text{prop}}$  and  $\tau$  (discussed in Section 5). We then use the distribution of (discretized)  $c_t$  values for packets  $t \in \text{window } w$ , denoted  $\tilde{c}_w$ , as the “ground-truth” for  $c_w$ , to compute the cross-entropy term:

$$J_{\text{win}}(\cdot; \Theta_{\text{window}}) := \sum_{i=1}^N \sum_{\text{window } w} \ell_{\text{CE}}(c_w^{(i)}, \tilde{c}_w^{(i)}). \quad (9)$$

We show in Section 5 that even the crude estimate  $\tilde{c}_w$  helps improve the performance of RBU significantly, by providing an effective “domain-specific regularization” while training.

**RBU training:** The input to the RNN  $\mathbf{h}_t$  of the cross-traffic model (8) comprises  $\mathbf{x}_t$  and  $\mathbf{h}_w$  obtained from the window model (7). The parameters  $d_{\text{prop}}$ ,  $d_{\text{trans}}$ ,  $\tau$  must satisfy certain physical constraints; in particular, for a trace with static features  $y_{\text{min}}$  and  $y_{\text{max}}$ ,  $0 < d_{\text{prop}}, d_{\text{trans}} \leq y_{\text{min}}$  and  $0 < \tau \leq y_{\text{max}}$  by definition. So, we use the bounded sigmoid function for estimating each of the 3 parameters:

$$g(\mathbf{x}) = (b_{\mathbf{x}} - a_{\mathbf{x}})\sigma(\langle \mathbf{w}_g, \mathbf{x} \rangle) + a_{\mathbf{x}}, \quad (10)$$

where  $a_{\mathbf{x}}$  and  $b_{\mathbf{x}}$  are the lower and upper bounds respectively for the parameter, given  $\mathbf{x}$ . Let  $p = \mathbb{1}_{\{y=\infty\}}$  denote the packet drop status in the training data, i.e.,  $p = 1$  when a packet is dropped (i.e.,  $y = \infty$ ), else  $p = 0$ . We use squared loss for delays and cross-entropy loss  $\ell_{\text{CE}}$  for drops:

$$\ell_{\text{pkt}}((\hat{y}, \hat{p}), y) = p \ell_{\text{CE}}(\hat{p}, 1) + (1 - p)(\ell_{\text{CE}}(\hat{p}, 0) + (\hat{y} - y)^2). \quad (11)$$




Figure 4: Left: Window-level (LSTM) and packet-level (RBU) models unrolled across time/packets. Right: RBU cell.

We set up the optimization problem (Figure 4):

$$\min_{\Theta_{\text{RBU}}, \Theta_{\text{window}}} J_{\text{pkt}}(\cdot; \Theta_{\text{RBU}}) + \lambda J_{\text{win}}(\cdot; \Theta_{\text{window}}), \quad (12)$$

where  $\Theta_{\text{RBU}}$  is the set of RNN weights in (1), and the weights for  $g$  in (10) needed to compute  $d_{\text{prop}}$ ,  $d_{\text{trans}}$ ,  $\tau$ . We use stochastic gradient-descent to learn the model parameters jointly, with mini-batching, and weight decay on the model parameters.

**RBU inference:** During simulation, the sender  $S$ , configured with protocol  $\Pi'$ , transmits data for the duration of the run (1 minute, in our evaluation). Unbeknownst to  $S$ , we replace the real network  $\mathcal{N}$  with the trained RBU model  $\hat{\mathcal{N}}$ . We sample the static features  $\mathbf{x}$  uniformly from the training data  $(\Pi, \mathcal{N})$ . We sample  $c_w$ , needed in (8), for  $1 \leq w \leq N_w$  ( $= 600$ , corresponding to 100ms windows over 1 minute), by unrolling the window-level LSTM (7) at once with input  $\mathbf{x}$ . For each packet  $t$  that  $S$  sends out, we form the features  $\mathbf{x}_t$  needed as input, together with  $\mathbf{x}$ , for  $\hat{\mathcal{N}}$ . We do a forward pass of  $\hat{\mathcal{N}}$  on the input (which takes around 2 ms on a standard GPU for RBU). The output (delay value or packet drop) from  $\hat{\mathcal{N}}$  is provided as feedback to  $S$ . Then,  $S$  sends out the next packet  $t+1$ , with inter-packet spacing of  $s_{t+1}$ , as determined by  $\Pi'$  acting on the model feedback, and so on.

**Multipath RBU training:** In the multi-path scenario, first note that we can rewrite the set of recurrences (4) for the links when packet  $t$  arrives, using the indicator function  $\mathbf{1}_t(k) = 1$ , if packet  $t$  enters queue  $k$  and  $\mathbf{1}_t(k') = 0$ , for  $k' \neq k$ . To mitigate (C3), we relax this indicator function, during training, with the probability of traversing queue  $k$  for packet  $t$ , denoted by  $q_t(k)$ . In some cases, we may be able to model  $q_t(k)$  using indirect observations in the training traces. For instance, in the two-path case, the fraction  $\tilde{q}_w$  of packets sent by  $S$  that arrived out-of-order at  $R$  in a given time window  $w$  (which can be easily computed for any trace) gives an approximation to  $q_t(2)$  (or  $q_t(1)$  which is  $1 - q_t(2)$ ) for  $w$ . We estimate  $q_w$  from the window-level embedding (7), just as  $c_w$  above, via incorporating a loss term  $\ell_{\text{CE}}(q_w, \tilde{q}_w)$  corresponding to (9). Details of the training and inference procedure for the multi-path scenario are given in the longer version (Anshuman et al. 2022). In Section 5, we demonstrate how this technique helps model multi-path behaviors like packet reordering in real-world network traces.

## 5 Experiments

**Compared methods:** We compare the RBU model with: (1) iBoxNet (Ashok et al. 2020): a SOTA network simulation approach that uses network domain knowledge to infer parameters from network packet traces, (2)  $\text{LSTM}_{\text{win}}$ : Autoregressive modeling of delays (Sutskever, Martens, and Hinton 2011; Graves 2013) (referred to as ‘‘T-forcing’’ in (Yoon, Jarrett, and van der Schaar 2019; Jarrett, Bica, and van der Schaar 2021; Xu et al. 2020)) as the factorized conditional  $\Pi_t P(y_t | \cdot)$ , implemented with LSTMs trained on windowed traces using  $\mathbf{x}_t$  as input and (discretized)  $y_t$  as output; at inference, we use the arg max of the output distribution as the delay value for all the packets in the window, (3)  $\text{LSTM}_{\text{pkt}}$ : Same as  $\text{LSTM}_{\text{win}}$  during training, but we sample a value from the output distribution independently for each packet in the window at inference, (4)  $\text{LSTM}_{\text{pkt, FIFO}}$ : same as  $\text{LSTM}_{\text{pkt}}$ , but enforces no-packet-reordering constraint (Proposition 1) while sampling delays, (5) Transformer: We use a GPT (decoder) model (Radford et al. 2018).

**Datasets:** We (1) design a synthetic benchmark using ns-3 as in (Ashok et al. 2020), consisting of 4200 traces for 4 different TCP protocols, on a variety of cross-traffic patterns and network configurations; and (2) use a subset of traces from a real physical network testbed Pantheon (Yan et al. 2018) for 2 TCP protocols. The ns-3 data corresponds to single-path configuration (hence, no packet reordering), while the Pantheon data includes naturally occurring reordering from real networks.

In all our experiments, we use only the TCP Cubic protocol (dominant on the internet) traces for training, and the other TCP protocols (Vegas, NewReno, and LEDBAT) for testing.

**Implementation:** We implement all the models in PyTorch. The static trace features are normalized to  $[0, 1]$ . For  $\text{LSTM}_{\text{win}}$  and  $\text{LSTM}_{\text{pkt}}$ , we (a) normalize the delays and the sending rates, and (b) use a 2-layer LSTM with 256 hidden units and a fully connected layer with discretized  $y_t$  as output (100-dimensional), tuned to maximize mean delay and throughput distribution match, on the training protocol. For RBU, we (a) use the same LSTM architecture, to be consistent, for the window-level model in (7), with discretized  $c_w$  in (8) as output, (b) set  $\gamma = 0.1$  in (8) and size of  $\mathbf{h}_t$  in (1) to 1, which works well across datasets, and (c) use single-bottleneck buffer RBU model (just as the ground-

Protocol	Model	WD (Tput, Delay)	WD (P95 Delay)
Cubic (Train)	iBoxNet	<b>0.015</b> ± 0.000	<b>0.000</b> ± 0.000
	LSTM <sub>pkt</sub>	0.271 ± 0.011	0.155 ± 0.001
	LSTM <sub>win</sub>	0.164 ± 0.002	0.150 ± 0.001
	LSTM <sub>pkt,FIFO</sub>	0.214 ± 0.009	0.119 ± 0.000
	Transformer	0.224 ± 0.015	0.030 ± 0.003
	RBU	0.032 ± 0.004	0.007 ± 0.000
Vegas (Test)	iBoxNet	0.054 ± 0.000	0.088 ± 0.000
	LSTM <sub>pkt</sub>	0.084 ± 0.000	0.112 ± 0.001
	LSTM <sub>win</sub>	0.108 ± 0.003	0.110 ± 0.001
	LSTM <sub>pkt,FIFO</sub>	0.061 ± 0.000	0.091 ± 0.000
	Transformer	0.079 ± 0.003	<b>0.007</b> ± 0.001
	RBU	<b>0.041</b> ± 0.004	0.057 ± 0.007
LEDBAT (Test)	iBoxNet	0.056 ± 0.000	<b>0.005</b> ± 0.000
	LSTM <sub>pkt</sub>	0.156 ± 0.002	0.154 ± 0.001
	LSTM <sub>win</sub>	0.118 ± 0.019	0.149 ± 0.000
	LSTM <sub>pkt,FIFO</sub>	0.106 ± 0.000	0.122 ± 0.000
	Transformer	0.103 ± 0.001	0.043 ± 0.005
	RBU	<b>0.049</b> ± 0.000	0.008 ± 0.000
NewReno (Test)	iBoxNet	<b>0.053</b> ± 0.000	<b>0.007</b> ± 0.000
	LSTM <sub>pkt</sub>	0.207 ± 0.009	0.152 ± 0.001
	LSTM <sub>win</sub>	0.133 ± 0.012	0.147 ± 0.001
	LSTM <sub>pkt,FIFO</sub>	0.166 ± 0.007	0.124 ± 0.000
	Transformer	0.165 ± 0.013	0.038 ± 0.002
	RBU	0.094 ± 0.018	0.024 ± 0.011

Table 1: (mean ± std. dev) Wasserstein distances (WD), the lower the better, for traces obtained via different models and protocols, on the ns-3 data. The best numbers are in bold.

truth) for ns-3, and 2-path RBU model for Pantheon. For iBoxNet, we use their official code. Training RBU on the largest dataset (ns-3) takes only about 3 minutes per epoch on V100 GPU. We report mean and std. dev., over 3 independent simulations, for all the metrics.

**Note:** We give all the key results in this section. For additional details on datasets, implementation, metrics, and for more comprehensive qualitative and quantitative results, we defer the reader to the longer version of our paper (Anshumaan et al. 2022).

### Qualitative Evaluation of Traces

In the top row of Figure 5, we show 4 randomly picked ground-truth (GT) traces for Cubic (train) and Vegas (test) protocols from the ns-3 dataset. Each trace, shown in a different color for a protocol, consists of a sending rate series and a delay series, trimmed to the first few seconds to show the local behaviors, in separate plots. The bottom row shows (a) for Cubic, traces obtained by running the RBU model with static features obtained from the same 4 GT Cubic traces (to enable direct comparison), and (b) for Vegas, example RBU traces that give similar throughput as the 4 GT traces.

Note how the RBU traces reflect the much lower delays with Vegas vs Cubic, just as in GT. RBU is able to achieve such accurate recreation, even though it was trained only on Cubic data.

### Quantitative Evaluation on Unseen Protocols

We first look at application-level metrics obtained via different models on the ns-3 data. We quantify the distributional match using the standard Wasserstein distance (WD). In Table 1, we present the 2-dimensional WD for the joint throughput, mean delay distribution, and WD for the P95 delay distribution. RBU is competitive w.r.t. the SOTA iBoxNet across all the protocols. LSTM-based baselines, on the other hand, fare poorly (as hypothesized in Section 2). The LSTM<sub>pkt,FIFO</sub> method, where we explicitly constrain the sampled delays to satisfy no re-ordering (as in the GT traces), does improve over the standard variants, but is often worse compared to RBU. Notably, we find that the Transformer (GPT) model outperforms the LSTM-based models in many cases (as one would expect), but is not as good as RBU. As noted in Section 2, our intuition for this is that (a) it is important to model the delay dynamics carefully, which the Transformers do not, and (b) Transformer models continue to suffer from the pitfalls of LSTM models at inference time. In Table 2, we see that RBU performs significantly better than the SOTA iBoxNet on the real-world Pantheon data on both Cubic (train) and Vegas (test) protocols.

Next, we quantify how well RBU preserves fine-grained temporal patterns. We divide the traces (sending rates, delays) into small chunks (of 15 packets) and compute the maximum mean discrepancy (MMD) between the simulated and the GT chunks, using RBF kernel. We show the MMD (the lower the better) of chunks over time in Figure 5 for the Vegas (test) protocol. RBU has much lower MMD in general, and especially relative to the baseline LSTM models. This suggests that RBU captures local temporal patterns over very long traces. Also, consistent with Table 1, LSTM<sub>pkt,FIFO</sub> performs better than the baselines, and iBoxNet is competitive. Details on MMD computation and results for different chunk lengths, protocols, and more baselines are in the longer version (Anshumaan et al. 2022).

**GAN techniques:** As we mentioned in Section 2, GAN techniques (Yoon, Jarrett, and van der Schaar 2019; Jarrett, Bica, and van der Schaar 2021) are unable to scale to sequences of lengths even in the order of hundreds, and for modest model sizes. Setting the output sequence length of the generator to a manageable size, and adversarially training it with aggregate traces, is also not meaningful in our setting — the output samples from the generator cannot be used to drive simulation as the sender requires continuous packet-level feedback.

### Simulating Real-World Network Phenomena

We now demonstrate RBU’s ability to simulate real-world network behaviors using packet re-ordering phenomenon, i.e., packets sent by  $S$  arriving out of order at receiver  $R$ , observed in the Pantheon traces. This is an important behavior from




Figure 5: First two plots: (Top row) Ground-truth sending rates, delays for 4 sample TCP Cubic & Vegas traces; (Bottom row) traces from the RBU model trained on Cubic, tested on Cubic and Vegas. Last plot:  $MMD^2$  vs chunks for TCP Vegas (ns-3 data).




Figure 6: Fraction of packets reordered in calls (first two), windows (last two) for TCP Cubic, Vegas (Pantheon)

Protocol	Model	Wasserstein Distances			
		2D (Tput, Mean Delay)	2D (Tput, P95 Delay)	1D Mean Delay	1D P95 Delay
Cubic (Train)	iBoxNet	0.150 ± 0.000	0.125 ± 0.000	0.142 ± 0.000	0.116 ± 0.000
	LSTM <sub>pkt</sub>	0.225 ± 0.018	0.245 ± 0.015	0.065 ± 0.007	0.101 ± 0.001
	LSTM <sub>win</sub>	0.288 ± 0.023	0.279 ± 0.024	0.078 ± 0.004	0.087 ± 0.009
	RBU	<b>0.098 ± 0.002</b>	<b>0.084 ± 0.002</b>	<b>0.038 ± 0.001</b>	<b>0.034 ± 0.001</b>
Vegas (Test)	iBoxNet	0.098 ± 0.000	0.184 ± 0.000	0.082 ± 0.000	0.211 ± 0.000
	LSTM <sub>pkt</sub>	0.254 ± 0.029	0.153 ± 0.038	0.234 ± 0.012	0.125 ± 0.005
	LSTM <sub>win</sub>	0.265 ± 0.020	0.145 ± 0.030	0.270 ± 0.019	0.135 ± 0.006
	RBU	<b>0.091 ± 0.029</b>	<b>0.089 ± 0.025</b>	<b>0.036 ± 0.005</b>	<b>0.043 ± 0.010</b>

Table 2: Wasserstein distances (WD) for Pantheon data for different models and protocols. The lower the better.

the application’s perspective as reordered packets could be treated as lost if they don’t arrive before a certain timeout (depending on the protocol). The metric of interest is the fraction of packets re-ordered in the calls.

In Figure 6, the metric CDFs for the RBU model (with 2 bottleneck links) traces align significantly better with GT, compared to LSTM<sub>win</sub>, for both train and test protocols; iBoxNet is not even shown here since its rigid single FIFO queue model precludes the recreation of reordering. We also show a baseline where we fix the length of the second queue  $\tau^{(2)}$  to be twice the first queue  $\tau^{(1)}$ ; this tends to reorder packets flowing through the second (longer) queue; while it performs reasonably well on the train protocol, the match is relatively poor on the test protocol, which underscores the effectiveness of our technique, and the joint learning of the RBU parameters.

**Limitations:** To capture and recreate real-world network behaviors such as reordering, we would need domain-specific insights on the new behaviors of interest. It is unlikely that the full expressive power of RBU can be exploited otherwise.

## 6 Conclusions

We formulate a novel ML problem at the intersection of sequential decision making, dynamical systems, and time-series generative modeling. We present the RBU construct that combines domain knowledge with the expressive power of neural models, yielding significantly better match for application-level metrics for network simulation than existing neural techniques and pure domain-knowledge based techniques. We also demonstrate that RBU is flexible enough to model real-world network phenomena like packet reordering accurately, which is currently not possible using domain-knowledge based techniques like iBoxNet.



## References

- Anshumaan, D.; Balasubramanian, S.; Tiwari, S.; Natarajan, N.; Sellamanickam, S.; and Padmanabhan, V. N. 2022. Simulating Network Paths with Recurrent Buffering Units. *arXiv preprint arXiv:2202.13870*.
- Ashok, S.; Duvvuri, S. S.; Natarajan, N.; Padmanabhan, V. N.; Sellamanickam, S.; and Gehrke, J. 2020. iBox: Internet in a Box. In *Proceedings of the 19th ACM Workshop on Hot Topics in Networks*, 23–29.
- Bengio, S.; Vinyals, O.; Jaitly, N.; and Shazeer, N. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems*, 1171–1179.
- Beucler, T.; Pritchard, M.; Rasp, S.; Ott, J.; Baldi, P.; and Gentine, P. 2021. Enforcing analytic constraints in neural networks emulating physical systems. *Physical Review Letters*, 126(9): 098302.
- Borovykh, A.; Bohte, S.; and Oosterlee, C. W. 2017. Conditional Time Series Forecasting with Convolutional Neural Networks. *stat*, 1050: 16.
- Esteban, C.; Hyland, S. L.; and Rätsch, G. 2017. Real-valued (Medical) Time Series Generation with Recurrent Conditional GANs. *arXiv:1706.02633*.
- Fu, R.; Chen, J.; Zeng, S.; Zhuang, Y.; and Sudjianto, A. 2019. Time series simulation by conditional generative adversarial net. *arXiv preprint arXiv:1904.11419*.
- Graves, A. 2013. Generating Sequences With Recurrent Neural Networks. *CoRR*, abs/1308.0850.
- Ha, S.; Rhee, I.; and Xu, L. 2008. CUBIC: a new TCP-friendly high-speed TCP variant. *ACM SIGOPS Operating Systems Review*, 42(5).
- Jarrett, D.; Bica, I.; and van der Schaar, M. 2021. Time-series Generation by Contrastive Imitation. *Advances in Neural Information Processing Systems*, 34.
- Levine, S.; Kumar, A.; Tucker, G.; and Fu, J. 2020. Offline Reinforcement Learning: Tutorial, Review, and Perspectives on Open Problems. *arXiv e-prints*, arXiv–2005.
- Li, Z.; Kovachki, N. B.; Azizzadenesheli, K.; Bhattacharya, K.; Stuart, A.; Anandkumar, A.; et al. 2020. Fourier Neural Operator for Parametric Partial Differential Equations. In *International Conference on Learning Representations*.
- Lin, Z.; Jain, A.; Wang, C.; Fanti, G.; and Sekar, V. 2020. Using GANs for Sharing Networked Time Series Data: Challenges, Initial Promise, and Open Questions. In *Proceedings of the ACM Internet Measurement Conference*, 464–483.
- NSNAM. 2011. ns-3 Network Simulator. <https://www.nsnam.org/>. Accessed: 2023-03-18.
- Radford, A.; Narasimhan, K.; Salimans, T.; and Sutskever, I. 2018. Improving language understanding by generative pre-training. [https://cdn.openai.com/research-covers/language-unsupervised/language\\_understanding\\_paper.pdf](https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf). Accessed: 2023-03-18.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Rangapuram, S. S.; Seeger, M. W.; Gasthaus, J.; Stella, L.; Wang, Y.; and Januschowski, T. 2018. Deep State Space Models for Time Series Forecasting. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems* 31, 7785–7794. Curran Associates, Inc.
- Ranzato, M.; Chopra, S.; Auli, M.; and Zaremba, W. 2016. SEQUENCE LEVEL TRAINING WITH RECURRENT NEURAL NETWORKS. *International Conference on Learning Representations*.
- Salinas, D.; Flunkert, V.; Gasthaus, J.; and Januschowski, T. 2020. DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*, 36(3): 1181–1191.
- Smith, K. E.; and Smith, A. O. 2020. Conditional GAN for timeseries generation. *arXiv preprint arXiv:2006.16477*.
- Sutskever, I. 2013. *Training recurrent neural networks*. University of Toronto Toronto, Canada.
- Sutskever, I.; Martens, J.; and Hinton, G. E. 2011. Generating text with recurrent neural networks. In *ICML*.
- Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27.
- Wei, W.; Gu, H.; and Li, B. 2021. Congestion Control: A Renaissance with Machine Learning. *IEEE Network*, 35(4): 262–269.
- Xu, J.; Pradhan, A.; and Duraisamy, K. 2021. Conditionally Parameterized, Discretization-Aware Neural Networks for Mesh-Based Modeling of Physical Systems. *Advances in Neural Information Processing Systems*, 34.
- Xu, T.; Wenliang, L. K.; Munn, M.; and Acciaio, B. 2020. COT-GAN: Generating Sequential Data via Causal Optimal Transport. *Advances in Neural Information Processing Systems*, 33.
- Yan, F. Y.; Ma, J.; Hill, G. D.; Raghavan, D.; Wahby, R. S.; Levis, P.; and Winstein, K. 2018. Pantheon: the training ground for Internet congestion-control research. In *2018 {USENIX} Annual Technical Conference ({USENIX}{ATC} 18)*, 731–743.
- Yoon, J.; Jarrett, D.; and van der Schaar, M. 2019. Time-series Generative Adversarial Networks. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems* 32, 5508–5518. Curran Associates, Inc.