

Infinity Learning: Learning Markov Chains from Aggregate Steady-State Observations

Jianfei Gao,¹ Mohamed A. Zahran,² Amit Sheoran,³ Sonia Fahmy,⁴ Bruno Ribeiro⁵

Department of Computer Science, Purdue University

305 N. University St, West Lafayette, IN 47907

{gao462¹, mzahran², asheoran³}@purdue.edu, {fahmy⁴, ribeiro⁵}@cs.purdue.edu

Abstract

We consider the task of learning a parametric Continuous Time Markov Chain (CTMC) sequence model without examples of sequences, where the training data consists entirely of aggregate steady-state statistics. Making the problem harder, we assume that the states we wish to predict are unobserved in the training data. Specifically, given a parametric model over the transition rates of a CTMC and some known transition rates, we wish to extrapolate its steady state distribution to states that are unobserved. A technical roadblock to learn a CTMC from its steady state has been that the chain rule to compute gradients will not work over the arbitrarily long sequences necessary to reach steady state—from where the aggregate statistics are sampled. To overcome this optimization challenge, we propose ∞ -SGD, a principled stochastic gradient descent method that uses randomly-stopped estimators to avoid infinite sums required by the steady state computation, while learning even when only a subset of the CTMC states can be observed. We apply ∞ -SGD to a real-world testbed and synthetic experiments showcasing its accuracy, ability to extrapolate the steady state distribution to unobserved states under unobserved conditions (heavy loads, when training under light loads), and succeeding in difficult scenarios where even a tailor-made extension of existing methods fails.

Introduction

Can we learn a parametric sequence model given only aggregate statistics as training data? As machine learning expands into new applications, new learning paradigms emerge, such as learning a sequence model from a set of observations without any clear time order between them.

Traditional supervised and unsupervised learning methods are essentially tasked with problems that can be learned from examples (*interpolation*). In a host of key applications of parametric sequence models, we want to *extrapolate*, i.e., take these aggregate observations and extrapolate them to a scenario *not observed* in the training data.

For instance, servers in the cloud collect system logs—aggregate statistics such as response-time distribution, queue length distribution—under light-load conditions. Under high-loads, however, these servers may disable statistics

collection (logs) due to the potential performance penalty of logging (Newman 2017). Capacity planning requires knowing how the servers perform under medium to high load conditions, which requires extrapolated predictions of request loss probability and server response times from the collected light-load data.

Hence, in this work we consider the task of learning a parametric Continuous Time Markov Chain (CTMC) sequence model—with transition rate matrix $Q(x, \theta)$ where x are known parameters but parameters θ must be learned—without examples of sequences, where the training data consists entirely of aggregate steady-state statistics. Making the problem harder, we further assume that the states we wish to predict are unobserved in the training data. More specifically, given an inductive bias over the transition rates of a CTMC and some known transition rates, we wish to extrapolate its steady state distribution to states that are unobserved. We focus on the application of predicting failures in queuing systems under heavy loads—e.g., predicting request loss rates in overloaded cloud services—with training data that contains only aggregate statistics of the system under light loads, and no observed losses. Traditionally, CTMCs are learned from observations of their transient (sequences given by transitions between states) not from observations of their steady state, even less so if only a subset of the state space is observable.

Remark 1. *Extrapolation v.s. generalization error: In our task we must make a distinction between generalization error—which is the error on unseen data that reduces with more training examples even without inductive biases—and extrapolation error (Marcus 1998)—which is a type of generalization error over unseen states and domains that does not reduce with more training data without the help of a modeling assumption. Our task is to learn a parametric model that is capable of extrapolation.*

Contributions. Our work introduces the general problem of learning a *parametric* CTMC from aggregate steady-state observations (frequencies) of part of the CTMC states, focusing on queueing systems as our application. We also introduce a novel method (∞ -SGD) to learn parametric CTMCs from aggregate steady-state observations, which

work even if the observations are over a restricted set of states. Our approach, ∞ -SGD, is a novel, theoretically principled, optimization approach that, among other things, uses randomly-stopped estimators (McLeish 2011). In our experiments ∞ -SGD finds significantly better maximum likelihood estimates than the baselines in real testbed and synthetic scenarios, both for the training and test data. We also see that ∞ -SGD can successfully extrapolate from training data under light queuing loads to predictions under heavy loads. We expect ∞ -SGD to be a useful tool in applications that collect aggregate statistics but need to learn parametric CTMCs.

Preliminaries

Consider a stationary and ergodic Continuous-Time Markov Chain (CTMC) $\mathcal{Y} = (\mathbf{Y}_\tau)_{\tau \geq 0}$ over a finite state space \mathbb{S} , where \mathbf{Y}_τ is the state of the Markov chain at time τ . The CTMC is governed by Kolmogorov’s Forward Equation

$$\frac{\partial}{\partial \tau} \mathbf{p}_{\mathbf{x}, \boldsymbol{\theta}}(\tau)^\top = (\mathbf{p}_{\mathbf{x}, \boldsymbol{\theta}}(\tau))^\top \mathbf{Q}(\mathbf{x}, \boldsymbol{\theta}), \quad (1)$$

where $\mathbf{Q}(\mathbf{x}; \boldsymbol{\theta})$ is a transition rate matrix parameterized by both \mathbf{x} (a vector of observed parameters, *e.g.*, request rate) and $\boldsymbol{\theta}$ (a vector of hidden parameters), $\mathbf{p}_{\mathbf{x}, \boldsymbol{\theta}}(\tau)$ is a column vector of dimension $|\mathbb{S}|$, with $\Pr[\mathbf{Y}_\tau = i] = \mathbf{p}_{\mathbf{x}, \boldsymbol{\theta}}(\tau)_i$ as the probability of being at state $i \in \mathbb{S}$ at time $\tau \geq 0$, given that \mathcal{Y} starts at state $j \in \mathbb{S}$ with probability $\Pr[\mathbf{Y}_0 = j] = \mathbf{p}(0)_j$.

The transition rate matrix $\mathbf{Q}(\mathbf{x}, \boldsymbol{\theta})$ is such that for $i \neq j$, $(\mathbf{Q}(\mathbf{x}, \boldsymbol{\theta}))_{ij} \geq 0$ describes the rate of the process transitions from state i to state j . The diagonal $(\mathbf{Q}(\mathbf{x}, \boldsymbol{\theta}))_{ii}$ is such that each row of $\mathbf{Q}(\mathbf{x}, \boldsymbol{\theta})$ sums to zero, irrespective of the values of \mathbf{x} and $\boldsymbol{\theta}$. Because \mathcal{Y} is stationary and ergodic, the solution to Equation (1) implies a unique steady state distribution $\boldsymbol{\pi}(\mathbf{x}, \boldsymbol{\theta}) = \lim_{\tau \rightarrow \infty} \mathbf{p}_{\mathbf{x}, \boldsymbol{\theta}}(\tau)$.

Parameterized transition rate matrix $\mathbf{Q}(\mathbf{x}, \boldsymbol{\theta})$. We exemplify $\mathbf{Q}(\mathbf{x}; \boldsymbol{\theta})$ with one of the simplest CTMCs: the birth-death process (BD). BD has two parameters: the request (birth) rate $\mathbf{x} = (\lambda)$ and the service (death) rate $\boldsymbol{\theta} = (\mu)$. The transition rate matrix is

$$\mathbf{Q}(\mathbf{x}; \boldsymbol{\theta}) = \begin{bmatrix} -\lambda & \lambda & 0 & \cdots & 0 \\ \mu & -(\mu + \lambda) & \lambda & \cdots & 0 \\ 0 & \mu & -(\mu + \lambda) & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \cdots & \cdots & \cdots & \cdots & -\mu \end{bmatrix},$$

where request rate λ is known but the service rate μ needs to be learned. In our work, $\mathbf{Q}(\mathbf{x}, \boldsymbol{\theta})$ can be significantly more complex, as we only assume $\mathbf{Q}(\mathbf{x}, \boldsymbol{\theta})$ is differentiable w.r.t. $\boldsymbol{\theta}$. More generally, we can have an $n \times n$ matrix

$$\mathbf{Q}(\mathbf{x}; \boldsymbol{\theta}) = \begin{bmatrix} -\sum_{i \neq 1} f_{1,i}(\mathbf{x}, \boldsymbol{\theta}) & f_{1,2}(\mathbf{x}, \boldsymbol{\theta}) & \cdots & f_{1,n}(\mathbf{x}, \boldsymbol{\theta}) \\ f_{2,1}(\mathbf{x}, \boldsymbol{\theta}) & -\sum_{i \neq 2} f_{2,i}(\mathbf{x}, \boldsymbol{\theta}) & \cdots & f_{2,n}(\mathbf{x}, \boldsymbol{\theta}) \\ \vdots & \vdots & \ddots & \vdots \\ f_{n,1}(\mathbf{x}, \boldsymbol{\theta}) & f_{n,2}(\mathbf{x}, \boldsymbol{\theta}) & \cdots & -\sum_{i \neq n} f_{n,i}(\mathbf{x}, \boldsymbol{\theta}) \end{bmatrix},$$

for some appropriate set of functions $\{f_{i,j}\}_{i,j}$ of \mathbf{x} and $\boldsymbol{\theta}$ (whose image must be in $[0, \infty)$).

Learning task. Consider learning $\boldsymbol{\theta}$ from a set of steady-state observations from a subset $\mathbb{S}' \subseteq \mathbb{S}$ of the states of the CTMC \mathcal{Y} . That is, even though \mathcal{Y} evolves over \mathbb{S} , the observations from states in $\mathbb{S}' = \mathbb{S} \setminus \mathbb{S}'$ are unavailable to us —*e.g.*, consider a system that disables statistics collection (logs) when it reaches a set of *system overload* states \mathbb{S}' .

Training data: Our training data consists of M time windows from which we have observed aggregate steady state data: $\mathcal{D} = \{(\mathbf{x}_m, \mathbf{y}_m)\}_{m=1}^M$, where $y_{m,j} \equiv (\mathbf{y}_m)_j$ is the number of steady state observations of state $j \in \mathbb{S}'$ at the m -th time window.

Loss function: The minimum negative log-likelihood of the model must be conditioned on only observing states of \mathbb{S}' in steady state (*i.e.*, $\tau \rightarrow \infty$),

$$\boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta}} \sum_{m=1}^M \mathcal{L}(\mathbf{y}_m, \lim_{\tau \rightarrow \infty} \mathbf{p}_{\mathbf{x}, \boldsymbol{\theta}}(\tau)), \quad (2)$$

where

$$\mathcal{L}(\mathbf{y}, \boldsymbol{\pi}) = - \sum_{j \in \mathbb{S}'} y_j \log \left(\frac{\pi_j}{\sum_{j' \in \mathbb{S}'} \pi_{j'}} \right), \quad (3)$$

such that the denominator ensures the observations are conditioned on only observing states in \mathbb{S}' —a detailed description of the math behind this conditional can be found in Meyer (1989).

In theory, we could optimize $\boldsymbol{\theta}$ in Equation (2) via gradient descent but the derivative of Equation (2) w.r.t. $\boldsymbol{\theta}$ requires computing the derivative of the steady state $\lim_{\tau \rightarrow \infty} \mathbf{p}_{\mathbf{x}, \boldsymbol{\theta}}(\tau)$, which is challenging as our steady state distribution does not have a closed-form expression.

The identifiability of \mathbf{Q} is irrelevant to our task: In our task, we wish to predict the steady-state distribution of unobserved states from samples from the steady state of observed states. Specifically, we wish to extrapolate those predictions such that we can predict these steady state distributions even when the observed parameters, \mathbf{x} of $\mathbf{Q}(\mathbf{x}, \boldsymbol{\theta})$ change. Because there are infinitely many \mathbf{Q} that can give the same correct steady state distribution predictions (see Supplementary Material A1), it is irrelevant to us knowing whether we recovered the “true” \mathbf{Q} . In fact, in our formulation there is no notion that we can ever learn a “true” \mathbf{Q} . We only care if it gives the correct steady state distribution.

Next, we review the related work.

Related Work

Inverting an MC steady state. Bernstein and Sheldon (2016) is one of the most closely related works, showing an estimator for an existing optimization approach from econometrics, Conditional Least Squares (CLS) (Miller 1952; Van Der Plas and others 1983; Kalbfleisch and Lawless 1984), which can be used to learn a Markov chain from aggregate statistics. This approach, however, is not designed to learn a parametric model (our $\mathbf{Q}(\mathbf{x}, \boldsymbol{\theta})$ needs derivatives w.r.t. $\boldsymbol{\theta}$) and thus, cannot extrapolate to unobserved states

in the training data. Moreover, our Markov chain is not homogeneous across observation time windows, requiring \mathbf{x} to also change, which conflicts with the assumptions in CLS.

Maystre and Grossglauser (2015) and Ragain and Ugander (2016) are the also closely related works, which learn the transition rates of a Plackett–Luce-type model CTMC from samples of its stationary distribution. In an earlier work, Kumar et al. (2015) learns a discrete-time Markov chain model similar to the Plackett–Luce’s model in the context of Web navigation. These earlier works, however, make domain-specific assumptions on \mathbf{Q} that make computing $\boldsymbol{\pi}$ from \mathbf{Q} trivial. We consider a general parametric $\mathbf{Q}(\mathbf{x}, \boldsymbol{\theta})$ that may have no trivial solution.

Hashimoto, Gifford, and Jaakkola (2016) uses far-apart observations of a MC to learn transition probabilities, and Pierson et al. (2018) uses cross-sectional data to learn a temporal model; these works are focused on specific diffusion processes. Our problem is also related to the more general problem of learning over distributions Szabó et al. (2016), which in our scenario requires a solution designed for the task.

Randomly stopped estimators have been used in unrelated machine learning tasks (Xu, Srivastava, and Sutton 2019) and (Filippone and Engler 2015), with significantly different tasks and estimators than ours. Applying randomly stopped estimators is mostly about proving that a specific estimator gives finite-variance estimates.

Queueing systems. Cloud computing has transformed IT operations and management by deploying services on commodity hardware in public or private data centers, saving millions of dollars in both capital and operational expenses (Armbrust et al. 2010; ETSI 2014). The savings in operational expenses can only be attained if the allocation of compute, memory, networking and storage resources scales based on the workload (Cao et al. 2018; Amazon 2019; Google Cloud 2019). A key problem in this elastic scaling is anticipating overload and failures in order to proactively allocate and initialize additional resources. This prediction needs to be done without sufficient data on overload and failures (Weiss and Hirsh 1998). Fortunately, several novel cloud computing services can be modeled by queueing systems (Yang et al. 2009; Khazaee, Misić, and Misić 2012). Existing approaches, however, require knowing the transition rate matrix rather than learning it from aggregate observations (as we do).

Learning Transition Rates from Aggregate Steady State Metrics

In this section, we will describe why a good parametric model of \mathbf{Q} is key to learn a $\boldsymbol{\theta}^*$ that can predict the steady state distribution of the states $\bar{\mathbb{S}} = \mathbb{S} \setminus \mathbb{S}'$ that are not observed from the states that are observed \mathbb{S}' . We will then introduce a few naïve methods to learn $\boldsymbol{\theta}^*$ from Equation (2) and show they are unsuitable for learning accurate CTMC transition rates, including an extension of BPTT. Finally, we will propose a novel approach to learn $\boldsymbol{\theta}^*$ that is significantly more accurate and more computationally efficient than the naïve approaches.

For the ease of notation, sometimes we abbreviate transition rate matrix that $\mathbf{Q} \equiv \mathbf{Q}(\mathbf{x}, \boldsymbol{\theta})$ and we may denote $(\mathbf{Q}(\mathbf{x}, \boldsymbol{\theta}))_{ij}$ by q_{ij} .

The need for a good parametric model of \mathbf{Q} : Without tied parameters in \mathbf{Q} through $\boldsymbol{\theta}$, the steady state distribution would be flexible enough to make $\pi_i, \forall i \in \mathbb{S}'$, and $\pi_j, \forall j \in \bar{\mathbb{S}}$, have arbitrarily different probabilities. This would make it impossible to correctly extrapolate the observed data and predict π_j for observations of states in $\bar{\mathbb{S}}$. An example is provided in Supplementary Material A1.

Lagrangian multipliers (e.g., Conditional Least Squares) are undesirable. To solve Equation (2), $t \gg 1$, we can add the condition $\boldsymbol{\pi}^T \mathbf{Q}(\mathbf{x}, \boldsymbol{\theta}) = 0$ as a Lagrangian multiplier as if $\boldsymbol{\pi}$ are extra learnable parameters. Then, the loss function is redefined as

$$-\sum_{j \in \mathbb{S}'} y_j \log \left(\lim_{t \rightarrow \infty} \frac{\pi_j}{\sum_{j' \in \mathbb{S}'} \pi_{j'}} \right) + \lambda \|\boldsymbol{\pi}^T \mathbf{Q}(\mathbf{x}, \boldsymbol{\theta})\|, \quad (4)$$

$\lambda > 0$, which is the Conditional Least Squares (Miller 1952) for a CTMC, a regularization applied the definition of a steady-state distribution. We found, however, that this approach is very challenging by design, since $\boldsymbol{\pi}$ is a function of \mathbf{x} and $\boldsymbol{\theta}$, and \mathbf{x} varies in the training data. Hence, the Lagrangian multiplier λ depends on the loss function (which is conditional) and on \mathbf{x} and $\boldsymbol{\theta}$, a challenging task.

Moreover, if we assume a constant λ , the resulting approach needs to work as as a bi-level optimization procedure (Bhatnagar and Borkar 1998; Colson, Marcotte, and Savard 2007). In computing the derivatives of the loss w.r.t. $\boldsymbol{\theta}$, there is essentially no connection between the data $(\{y_j\}_{j \in \mathbb{S}'})$ and $\boldsymbol{\theta}$, which is the reason why the approach fails. Fixing these Lagrangian multiplier issues is future work.

Solution through Uniformization and Chain Rule

Our first step to a solution is to *uniformize* the Markov chain \mathbf{Q} in order to transform the CTMC of Equation (1) into a discrete-time Markov chain (DTMC) with probability matrix $\mathbf{P}(\mathbf{Q}(\mathbf{x}, \boldsymbol{\theta}))$ as described below (Jensen 1953). Hence, we will see that an approximation of the steady-state distribution can be obtained by recursively applying $\mathbf{P}(\mathbf{Q}(\mathbf{x}, \boldsymbol{\theta}))$, and the derivative of this recursive application of the transition probability can be obtained via chain rule.

Definition 1 (Uniformized Markov Chain). *Let $\mathbf{Q}(\mathbf{x}, \boldsymbol{\theta})$ be a stationary and ergodic CTMC. We define a set of Chapman-Kolmogorov equations representing the CTMC at the events (arrivals) of a Poisson process with rate $\gamma(\mathbf{x}, \boldsymbol{\theta}) > \max(-\text{diag}(\mathbf{Q}(\mathbf{x}, \boldsymbol{\theta})))$. The distribution after $t \geq 0$ of these events is*

$$\mathbf{p}^{(\text{events})}(t; \mathbf{x}, \boldsymbol{\theta})^\top = \mathbf{p}^{(\text{events})}(0)^\top \mathbf{P}(\mathbf{Q}(\mathbf{x}, \boldsymbol{\theta}))^t, \quad (5)$$

where $\mathbf{p}^{(\text{events})}(0)$ is some initial distribution and $\mathbf{P}(\mathbf{Q}(\mathbf{x}, \boldsymbol{\theta})) = \mathbf{I} + \mathbf{Q}(\mathbf{x}, \boldsymbol{\theta})/\gamma(\mathbf{x}, \boldsymbol{\theta})$, where \mathbf{I} is the identity matrix.

By construction, since \mathcal{Y} is ergodic and has a steady state π , the Markov chain described by $P(Q(x, \theta))$ has the same steady state as the CTMC described by $Q(x, \theta)$ (Jensen 1953), i.e., for $\mathbf{p}_{x, \theta}(\tau)$ as described in Equation (1),

$$\begin{aligned} \pi(x, \theta) &= \lim_{\tau \rightarrow \infty} \mathbf{p}_{x, \theta}(\tau) = \lim_{t \rightarrow \infty} \mathbf{p}^{(\text{events})}(t; x, \theta) \\ &= \lim_{t \rightarrow \infty} (\mathbf{p}^{(\text{events})}(0)^\top P(Q(x, \theta))^t)^\top. \end{aligned} \quad (6)$$

Note that $(P(Q(x, \theta))^t)_{ij}$ is the probability that the CTMC starts at state i and reaches state j after $t \geq 0$ events of the Poisson process with rate $\gamma(x, \theta)$ given by Definition 1.

In what follows, we sometimes denote the probability matrix as $P \equiv P(Q(x, \theta))$ and steady state distribution as $\pi \equiv \pi(x, \theta) \equiv \pi(P(Q(x, \theta)))$.

Chain rule to learn θ from a steady-state approximation.

Learning θ^* in Equation (2) through gradient descent can be approximated for a large enough value of $t^* \gg 1$ through the chain rule (Werbos 1990). The derivative of the loss in Equation (2) is

$$\begin{aligned} \frac{\partial \mathcal{L}(y, \mathbf{p}^{(\text{events})}(t^*; x, \theta))}{\partial \theta_k} &= \\ \sum_{\substack{i \neq j, \\ i, j \in \mathbb{S}}} \left(\frac{\partial \mathcal{L}(y, \mathbf{p}^{(\text{events})}(0)^\top P^{t^*})}{\partial (P^{t^*})_{ij}} \frac{\partial (P(Q(x, \theta))^{t^*})_{ij}}{\partial \theta_k} \right), \end{aligned} \quad (7)$$

In order to compute $\partial(P(Q(x, \theta))^{t^*})/\partial \theta_k$, we have to recursively apply the chain rule, which leads to a backpropagation through time (BPTT)-style method.

BPTT challenges. Directly using BPTT, however, has both theoretical and practical barriers. The theoretical challenge is finding a large-enough value of t^* that allows $P(Q(x, \theta))^{t^*}$ to approximate the steady-state distribution $\pi(x, \theta)$ for any assignment of $Q(x, \theta)$ that our optimization might find. The computational challenge is both of computational resources and of numerical precision.

The following definition gives a divide-and-conquer aid to the computational challenge of calculating BPTT over $P(Q(x, \theta))^{t^*}$:

Definition 2 (Divide-and-Conquer BPTT (DC-BPTT)). *Assume $t^* = 2^T$ for some $T > 1$. Rather than backpropagating over t^* time steps—which is difficult if t^* is large due to vanishing and exploding gradients—, we will use a divide-and-conquer approach to reduce the backpropagation steps to $\log_2 t^* = T$, by noting that*

$$P^{2^T} = ((P^{2^{T-2}})^2)^2 = (\dots (P^2)^2 \dots)^2. \quad (8)$$

That is, rather than multiplying an intermediate P^t by P to obtain P^{t+1} , we multiply P^t by itself to obtain P^{2t} .

Computing P^{t^*} , $t^* = 2^T$ with $T > 1$, from Definition 2 is more computationally efficient than the naïve t^* multiplications $P \cdot \dots \cdot P$ because the computation graph is a tree whose backpropagation paths from the root to the leaves give the same derivatives at the same tree height. Unfortunately, as we see in our experiments, DC-BPTT still fails in the most challenging tasks.

Solution via Infinity Learning

An alternative to BPTT is to dive deeper into the chain rule equations and look for mathematical equivalences. Rather than using BPTT to compute the gradient $\partial P(Q(x, \theta))^{t^*}/\partial \theta_k$, $k \in \mathbb{S}$, in Equation (7), we can make use of the following observation.

Lemma 1. *Let $Q(x, \theta)$ be a K -state transition matrix and $P(Q(x, \theta))$ be its uniformized Markov Chain. $P(Q(x, \theta))^t$ is the Markov chain after t steps where $t > 0$, then the gradients of P^t w.r.t. θ_k is*

$$\begin{aligned} \nabla_{\theta_k}^{(t)} P(Q(x, \theta)) &\equiv \frac{\partial P(Q(x, \theta))^t}{\partial \theta_k} \\ &= \sum_{l=1}^t P(Q(x, \theta))^{t-l} \frac{\partial P(Q(x, \theta))}{\partial \theta_k} P(Q(x, \theta))^{l-1}, \end{aligned} \quad (9)$$

where

$$\frac{\partial P(Q(x, \theta))}{\partial \theta_k} = \sum_{ij} \frac{\partial P(Q)}{\partial q_{ij}} \frac{\partial q_{ij}(x, \theta)}{\partial \theta_k}.$$

The proof is in the Supplementary Material B1. Because $\sum_{l=1}^{\infty} P(Q(x, \theta))^{l-1}$ diverges for any valid x and θ , it is not obvious that Equation (9) converges to a unique fixed point for $t \rightarrow \infty$. In what follows we show that the gradient in Equation (9) exists and is unique as $t \rightarrow \infty$:

Proposition 1 (Infinite Gradient Series Simplification). *Let Q be a K -state transition rate matrix of a stationary and ergodic MC. Equation (9) for $t \rightarrow \infty$, henceforth denoted $\nabla_Q^{(\infty)} P(Q) \equiv \lim_{t \rightarrow \infty} \nabla_Q^{(t)} P(Q)$, exists and is unique and can be redefined as*

$$\begin{aligned} (\nabla_Q^{(\infty)} P(Q))_{ij} &\equiv \lim_{t \rightarrow \infty} \sum_{l=1}^t \left(P^{t-l} \frac{\partial P(Q)}{\partial q_{ij}} P^{l-1} \right) \\ &= \Pi \sum_{l=0}^{\infty} \frac{\partial P(Q)}{\partial q_{ij}} P^l, \end{aligned} \quad (10)$$

where Π is a matrix whose rows are the steady state distribution π . Note that the diagonal $i = j$ is trivial to compute but should be treated as a special case.

The proof in the Supplementary Material B2 shows that $\nabla_Q^{(\infty)} P(Q)$ converges because the term inside the sum converges to a matrix of zeros as $l \rightarrow \infty$. Using Proposition 1 it is easy to prove that Equation (9) converges as $t \rightarrow \infty$.

While Proposition 1 shows that $\nabla_Q^{(\infty)} P(Q)$ converges, evaluating the infinite sum in Equation (10) is challenging. Truncating the sum would make the gradient biased, deviating the fixed point solution of Equation (2). To circumvent the infinite sum in Equation (10), we propose ∞ -SGD, a numerically stable stochastic gradient descent method that can optimize gradients with infinite sums—as long as the sum is a weakly convergent series. Our experiments show that ∞ -SGD consistently outperforms BPTT in stability to hyperparameters in convergence rate, and in estimation accuracy.

Theorem 1 (Infinity Stochastic Gradient Descent (∞ -SGD)). *Let \mathbf{Q} be the transition rate matrix of a stationary and ergodic CTMC. Assume strictly positive values for the learnable parameters $\theta^{(h)}$ at the h -th step of the optimization. Let $\mathbf{P}(\mathbf{Q}(\mathbf{x}, \theta^{(h)}))$ be its unnormalized transition probability matrix per Definition 1. Let $\mathcal{L}(\mathbf{y}, \boldsymbol{\pi})$ be as in Equation (3). Reparameterize $\tilde{\mathcal{L}}(\mathbf{y}, \mathbf{x}, \theta^{(h)}) = \mathcal{L}(\mathbf{y}, \boldsymbol{\pi}(\mathbf{x}, \theta^{(h)}))$ as the loss function with respect to $\theta^{(h)}$. Let $X^{(h)} \sim \text{Geometric}(p^{(h)})$, $X^{(h)} \in \mathbb{Z}^+$, be an independent sample of a Geometric distribution with $p^{(h)} < \delta^{(h)}$, where $\delta^{(h)}$ is the spectral gap of $\mathbf{P}(\mathbf{Q}(\mathbf{x}, \theta^{(h)}))$. Then, for $0 < \epsilon \ll 1$ and for all learnable parameters θ ,*

$$\theta_k^{(h+1)} = \max \left(\theta_k^{(h)} - \eta^{(h)} \nabla_{\theta_k} \tilde{\mathcal{L}}(\mathbf{y}, \mathbf{x}, \theta) \Big|_{\theta=\theta^{(h)}}, \epsilon \right),$$

where

$$\begin{aligned} \nabla_{\theta_k} \tilde{\mathcal{L}}(\mathbf{y}, \mathbf{x}, \theta) &= \sum_{ij} \sum_{mn} (\mathbf{p}^{(\text{events})}(0))_m \frac{\partial \mathcal{L}(\mathbf{y}, \boldsymbol{\pi})}{\partial \pi_n} \Big|_{\boldsymbol{\pi}=\boldsymbol{\pi}(\mathbf{x}, \theta)} \\ &\quad \times \pi(\mathbf{x}, \theta)_n \Gamma_{ijmn}(\mathbf{x}, \theta) \frac{\partial \mathbf{Q}(\mathbf{x}, \theta)_{ij}}{\partial \theta_k} \end{aligned}$$

with $h = 0, 1, \dots$, where $\boldsymbol{\pi}(\mathbf{x}, \theta)$ is the steady state distribution defined in Equation (6), $\eta^{(h)}$ is the learning rate with $\sum_{h=0}^{\infty} \eta^{(h)} = \infty$, $\sum_{h=0}^{\infty} (\eta^{(h)})^2 < \infty$, and

$$\Gamma_{ijmn}(\mathbf{x}, \theta) = \sum_{t=0}^{X^{(h)}} \left[\frac{\partial \mathbf{P}(\mathbf{Q}(\mathbf{x}, \theta))}{\partial q_{ij}} \frac{\mathbf{P}(\mathbf{Q}(\mathbf{x}, \theta))^t}{\mathbb{P}[X^{(h)} > t]} \right]_{mn}, \quad (11)$$

is a stochastic gradient descent method that minimizes Equation (2).

The proof of the theorem is in the Supplementary Material B3. The main insight is the use of Proposition 1 to produce a randomly-stopped unbiased estimator. The requirement in Theorem 1 that $p^{(h)} < \delta^{(h)}$ comes from a loose bound, *i.e.*, in practice $p^{(h)}$ can be relatively large (larger than the spectral gap) as our empirical results show —*e.g.*, all of our empirical results use the constant $p^{(h)} = 0.1, \forall h$. We have also tested some experiments with $p^{(h)} = 0.01$, which works as well as $p^{(h)} = 0.1$ (see Supplementary Material C4). As it is application-dependent, the value of $p^{(h)}$ should be seen as a hyperparameter. In what follows we introduce our empirical results.

Results

In this section, we contrast the accuracy and convergence of ∞ -SGD (Theorem 1) against DC-BPTT (Definition 2) and find that ∞ -SGD is more stable and consistently learns more accurate models. The primary application of our experiments is predicting *request loss rates in a queueing system* from data that has no observed losses, under the following conditions: (a) we learn θ^* of Equation (2) as a function of known request rate $\mathbf{x}^{\text{light}} \in \Lambda^{\text{light}}$ under light load (no losses) in the training data, and predict $\boldsymbol{\pi}^{\text{heavy}}$, the steady state request loss rates under heavy loads in the test data (out-of-sample extrapolation), where $\boldsymbol{\pi}^{\text{heavy}}$ is such that

$(\boldsymbol{\pi}^{\text{heavy}})^\top \mathbf{Q}(\mathbf{x}^{\text{heavy}}; \theta^*) = 0$ with $\mathbf{x}^{\text{heavy}} > \max(\Lambda^{\text{light}})$; moreover, (b) only part of the state space is observed in the training data, $\mathbb{S}' \subset \mathbb{S}$, and we wish to predict the steady state probability of the unobservable states $\bar{\mathbb{S}}' = \mathbb{S} \setminus \mathbb{S}'$.

Baseline method. Due to the absence of methods on parametric inference of CTMCs from steady-state observations, our main baseline is the DC-BPTT of Definition 2. In most of our simulations, we set $t^* = 128 = 2^7$ and $\mathbf{p}^{(\text{events})}(0) = \mathbf{1}^\top / |\mathbb{S}|$ throughout all our experiments. We also tested BPTT without divide and conquer but find the optimization unstable due to the long backpropagation paths. We tested $t^* \in \{16, 128\}$ and found that smaller values of t^* are easier to optimize but —as expected— generally produce worse approximations of the steady state for heavy loads.

Infinity learning. Our experiments also test our proposed approach, ∞ -SGD, with $X^{(h)} \sim \text{Geometric}(p)$ of Theorem 1, where p is a constant success probability, *i.e.*, $\mathbb{E}[X] = 1/p$. In most of our experiments, $p = 0.1$, that is, on average we consider only the first ten terms in the sum of Equation (10). Contrast, ∞ -SGD's 10 summation terms with matrix powers that need no chain rule, with the baseline DC-BPTT approach (Definition 2) where $\mathbf{P}(\mathbf{Q}(\mathbf{x}, \theta))^{128}$ needs to be computed together with a chain rule to compute gradients over the matrix multiplications. It is no surprise that ∞ -SGD is a more stable optimization method (no vanishing or exploding gradients); interestingly, ∞ -SGD also works well on the tested slow-mixing CTMCs, while baseline methods like DC-BPTT fail in these scenarios (see Supplementary Material C5).

Relaxing the parametric model. In some of our experiments, we will construct transition rate matrix $\mathbf{Q}'(\mathbf{x}, \theta, \tilde{\mathbf{Q}}) = \mathbf{Q}(\mathbf{x}, \theta) + \tilde{\mathbf{Q}}$ with an $\alpha \|\tilde{\mathbf{Q}}\|_2^2$ regularization penalty, $\alpha > 0$, where $\tilde{\mathbf{Q}}$ is an additional non-parametric learnable matrix s.t. $\tilde{\mathbf{Q}}_{ij} = 0$ whenever $(\mathbf{Q}(\mathbf{x}, \theta))_{ij} \neq 0$, otherwise $\tilde{\mathbf{Q}}_{ij}$ is a learnable parameter of our model. This allows some uncertainty on the form of our parametric models. It also allows us to learn the parameters through an interpolation between parametric and non-parametric CTMC models.

The regularization term $\alpha \|\tilde{\mathbf{Q}}\|_2^2$ is added to the negative log-likelihood loss in Equation (2) to ensure that we can control how much flexibility we want. With small values of α , we are testing how overparameterization, *i.e.*, having too many extra parameters in \mathbf{Q}' , affects learning and generalization. Our experiments show that $\alpha \gg 1$ gives the best results, *i.e.*, the correct parametric model works best. We also see that $\alpha \approx 1$ still gives competitive results (refer to Supplementary Material C3), showing that some model flexibility is tolerable. In contrast, we see that $\alpha = 0.1$ tends to significantly hurt our ability to extrapolate queue losses in the test data.

Table 1: [MAPE] Simulation results showing MAPE/100 ($\langle \text{Mean Absolute Error} \rangle / \langle \text{true value} \rangle$) errors between predicted steady state and ground-truth for failure states in test data (heavy load). Training data collected under light loads and restricted observed states (queues zero and one). Mixing rates are determined by the spectral gaps δ_n observed in training data over multiple time windows ($n = 1, \dots, 50$). With 95% confidence intervals.

	δ_n (spectral gap)	DC-BPTT $t^* = 16$	DC-BPTT $t^* = 128$	∞ -SGD ($p = 0.1$)
Testbed Emulation (Upper Trig.)	N/A	$1.43 \times 10^1 \pm 0.00$	$1.88 \times 10^1 \pm 0.00$	$9.33 \times 10^{-1} \pm 8.91 \times 10^{-2}$
M/M/1/K (fast-mix)	[0.022, 0.043]	$2.04 \times 10^{-1} \pm 2.86 \times 10^{-4}$	$1.32 \pm 4.94 \times 10^{-3}$	$1.18 \times 10^{-2} \pm 8.01 \times 10^{-3}$
M/M/1/K (slow-mix)	[0.005, 0.008]	$8.91 \times 10^3 \pm 2.02 \times 10^2$	$6.68 \times 10^1 \pm 7.61$	$8.88 \times 10^{-1} \pm 1.48$
M/M/m/m + r	[0.013, 0.024]	$4.11 \times 10^{-1} \pm 4.47 \times 10^{-2}$	$4.01 \times 10^{-1} \pm 8.86 \times 10^{-2}$	$1.52 \times 10^{-1} \pm 8.50 \times 10^{-2}$
M/M/Multiple/K	[0.068, 0.096]	$9.09 \times 10^{-1} \pm 1.20 \times 10^{-2}$	$4.03 \times 10^1 \pm 6.74 \times 10^{-2}$	$2.27 \times 10^{-1} \pm 1.47 \times 10^{-2}$

Table 2: [MSE] Simulation results showing MSE errors between predicted steady state and ground-truth for failure states in test data (heavy load). Training data collected under light loads and restricted observed states (queues zero and one). Mixing rates are determined by the spectral gaps δ_n observed in training data over multiple time windows ($n = 1, \dots, 50$). With 95% confidence intervals.

	δ_n (spectral gap)	DC-BPTT $t^* = 16$	DC-BPTT $t^* = 128$	∞ -SGD ($p = 0.1$)
Testbed Emulation (Upper Trig.)	N/A	$4.80 \times 10^{-1} \pm 0.00$	$8.45 \times 10^{-1} \pm 0.00$	$2.41 \times 10^{-3} \pm 5.20 \times 10^{-4}$
M/M/1/K (fast-mix)	[0.022, 0.043]	$6.81 \times 10^{-3} \pm 1.99 \times 10^{-5}$	$2.45 \times 10^{-1} \pm 1.84 \times 10^{-3}$	$4.98 \times 10^{-5} \pm 5.37 \times 10^{-5}$
M/M/1/K (slow-mix)	[0.005, 0.008]	$1.14 \times 10^{-2} \pm 1.54 \times 10^{-4}$	$4.84 \times 10^{-4} \pm 1.99 \times 10^{-4}$	$9.36 \times 10^{-4} \pm 1.58 \times 10^{-3}$
M/M/m/m + r	[0.013, 0.024]	$4.25 \times 10^{-2} \pm 7.25 \times 10^{-3}$	$2.87 \times 10^{-2} \pm 1.09 \times 10^{-2}$	$6.65 \times 10^{-3} \pm 5.61 \times 10^{-3}$
M/M/Multiple/K	[0.068, 0.096]	$8.97 \times 10^{-3} \pm 1.24 \times 10^{-4}$	$6.84 \times 10^{-1} \pm 1.99 \times 10^{-3}$	$5.34 \times 10^{-4} \pm 8.23 \times 10^{-5}$

Testbed Experiments

We now contrast DC-BPTT against ∞ -SGD in a real-world testbed emulating a Voice-over-LTE (VoLTE) system in a wireless cellular network. The testbed is configured as a single server with a waiting queue of size $K = 20$. The training data (86 time windows) is generated under light loads (with mean 7.7 and median 3 call losses) and the test data (137 time windows) under heavy loads (with mean 135.2 and median 254 call losses). *Moreover, we also restrict the observations in the training data, $\mathbb{S}' \subset \mathbb{S}$, to queue sizes one and two, estimated from the request processing delays collected at the clients.* We define $\mathbf{Q}(x; \theta)$ symbolically using Pytorch’s autodiff function (Paszke et al. 2017). The Supplementary Material C1 contains the details of our experimental setup and methodology. We study two parametric $\mathbf{Q}(x, \theta)$ models:


(A) *M/M/1/K model:* We start with arguably the most fundamental CTMC parametric model of a queueing system, the M/M/1/K queue with a single server and a single queue of size $K = 21$, which can hold up to 20 requests waiting for service. The queue can be described by a CTMC $\mathbf{Q}(x_m, \theta^*)$, where service requests arrive according to a Poisson process with request rate x_m assumed constant over time window m —the length of a time window is the minimum time resolution of our logs (one second). If a new request arrives when the queue is full, it is dropped by the server. Service times are exponentially distributed with rate θ^* , assumed constant—the service rate capacity of the system. We assume K is known. The request call rate x_m at time window n is known while θ is the only parameter that needs to be learned. The system is assumed in steady state even over short time windows.

(B) *Upper Triangular model:* Since it is difficult to fit a real system with an exact CTMC queueing model, we consider an embedded birth-death process $\mathbf{Q}(x_m; \theta^*)$, called *Upper Triangular model*, which sets the upper triangular

portion of \mathbf{Q} to x_m in the upper diagonal and zeros everywhere else. This model (setting the upper triangular to zeros) indicates that only a client request can increase the queue size. The lower triangular part of \mathbf{Q} is populated with different parameters that we will learn, which implies learning $|\theta^*| = |\mathbb{S}|(|\mathbb{S}| - 1)/2$ parameters for a CTMC with $|\mathbb{S}|$ states. The learnable parameters θ are initialized with zeros and learned with either DC-BPTT or ∞ -SGD. We also know the maximum queue size $K = 21$ as in the M/M/1/K scenario. Since we want to predict overload, we add a new state to represent that case. Any arrival after reaching the last state, when the current queue size is K , will force the system to transit to this overload state. Thus, we have $|\mathbb{S}| = K + 2$.

Results. Figure 1a shows the mean absolute percentage error (MAPE) between the predicted call drop probability (given the call request rate) and the true call drop probability in the test data under heavy loads, for the upper triangular and M/M/1/K parametric model learned with ∞ -SGD. The *Upper Triangular* model achieves much lower test mean squared error (MSE) ($2.57 \times 10^{-3} \pm 5.75 \times 10^{-4}$) corresponding to MAPE of about 80% to 100% over test call drop probabilities in the range [0.0045, 0.1629], *i.e.*, it more accurately extrapolates the training data (light loads, just observing queues of size one and two) to the test (heavy loads, full queue). The M/M/1/K parametric model is too simple and performs poorly with test MAPE of 449% (test MSE is reasonable at 7.78×10^{-2}).

We also investigate the transition rate matrix learned by upper triangular model in Figure 1c. We note that the learned queue is quite similar to an M/M/1/K, but the service rate is decreasing as the queue size is increasing. Surprisingly, this is a real phenomenon when real systems start to become overloaded (Jain and Ramakrishnan 1988). We also see some *reset* transitions, where the system goes from a full queue to a nearly empty queue. Finally, Figure 1b shows that



(a) (Training curves) Effect of parametric models on test MAPE (learned with ∞ -SGD).

(b) (Training curves) Effect of learning methods on test MAPE of upper triangular model.

(c) $Q(x, \theta^*)$ learned by upper triangular model with ∞ -SGD. Upper triangular (x) is removed or shows zeros, lower triangular shows θ^* .

Figure 1: Real-world experiment results on VoLTE testbed. (a-b) Test error (MAPE) of unseeing failure state —here, *ground-truth* of dropped call probability— under heavy load (while training under light loads), as a function of training epochs. In these plots we verify the better generalization and stability of ∞ -SGD. (a) Shows that more flexible *Upper Triangular* parametric model has much smaller (near-zero) test error than the more strict *M/M/1/K* parametric model. (b) Shows that ∞ -SGD significantly outperforms DC-BPTT (which fails to learn). (c) Learned $Q(x, \theta^*)$ by the upper triangular parametric model with ∞ -SGD, showing an emergent block structure.

only ∞ -SGD can learn the Q of the upper triangular model —which has $|\mathbb{S}|(|\mathbb{S}| - 1)/2 = 210$ parameters—, while DC-BPTT has vanishing gradients for both $t^* \in \{16, 128\}$ —we note that $t^* = 16$ has a smaller loss than $t^* = 128$. Finally, Tables 1 and 2 reaches the obvious conclusion that ∞ -SGD learns significantly better models for extrapolation over the test data than DC-BPTT.

In what follows we explore the differences between DC-BPTT and ∞ -SGD in synthetic experiments.

Synthetic Experiments

We now turn our attention to simulations. Due to space limitations, we give a succinct description of the experiments, relegating details and additional results to Supplementary Materials C2, C3, C4, and C5.

Birth-death queues: We start with arguably the most fundamental parametric CTMC queueing system, the diagonal structure of the *birth-death process*. The birth-death process approximates a number of queueing systems, such as the *M/M/1/K* queue with a single queue of size K and a single server, and the *M/M/m/K* queue with m servers that has been used to approximate cloud services (Yang et al. 2009; Khazaei, Mistic, and Mistic 2012). Since queue size K must be larger than m to support all servers, the *M/M/m/K* queue is typically denoted as *M/M/m/m + r*, where $r \geq 0$.

Training and testing data. Structures used to simulate data are provided in Supplementary Material A3. We predefine the service rates ($\theta^* = 25$ for *M/M/1/K* (slow-mix and fast-mix), $\theta^* = 5$ for *M/M/5/5+r*, $\theta^* = (15, 10, 5)$ for *M/M/Multiple/K*) and queue sizes $K = 20$ ($r = 15$), then at each time window in the training data we sample a request rate uniformly in the interval $x \in [11, 15]$ (light load), except for *M/M/1/K* slow-mix $x \in [21, 30]$ (to decrease the spectral gap). At test time, in the test data, we sample a request rate uniformly in the interval $x \in [31, 60]$ (heavy load), except for *M/M/1/K* slow-mix $x \in [11, 40]$ (to decrease the spectral gap). We assume we *only* observe the

queue size if it is empty or it has exactly one request, *i.e.*, $\mathbb{S}' = \{0, 1\}$. This emulates a common trend in logging critical infrastructure systems, where logging stops as soon as the server load is non-trivial (Newman 2017). We also have extra results with different transition rates (in an easier task where ∞ -SGD does even better) and more details on our training and test data generation in Supplementary Materials C2 and C5.

Results. Our goal is to predict the request loss probability against ground-truth under a range of both heavy and lighter loads, while training under a narrow range of light loads. In the *M/M/1/K* and *M/M/m/m+r* simulations, the training data consists of the aggregate frequencies observed for queue sizes zero and one during one second, along with the request rate. For *M/M/Multiple/K*, aggregate frequencies for queue sizes zero to three are observed.

Tables 1 and 2 compare the extrapolation error of DC-BPTT and ∞ -SGD in our synthetic experiment using MAPE and MSE errors, respectively. Our approach, ∞ -SGD, is consistently better than DC-BPTT over all simulations and on both error metrics (MAPE and MSE). For a slow-mixing *M/M/1/K*, ∞ -SGD extrapolation MAPE error is 1/100-th of DC-BPTT MAPE error, considering the confidence interval. In some of the scenarios, DC-BPTT finds gradient vanishing problems (failing to learn) giving very large errors (see training curves in Supplementary Material C5), while ∞ -SGD never fails to obtain gradients that can reduce the loss during the optimization.

MAPE result shows that $t^* \leq 128$ is not enough to see the slow-mixing chain in steady state. Success in MSE for slow mixing while failing in MAPE shows that DC-BPTT has trouble learning parametric CTMCs well enough to predict out-of-sample (extrapolated) target states that have small probabilities. Moreover, DC-BPTT $t^* = 16$ tends to achieve lower errors (both MAPE and MSE) than DC-BPTT $t^* = 128$ in the *M/M/1/K* fast mixing scenarios.

We now look at ground-truth θ^* parameters and their

estimates $\hat{\theta}$ from ∞ -SGD. We note that the estimates are very close to the true values. In the M/M/1/K model (true $\theta^* = 25$), the (slow mixing) scenario gives $\hat{\theta} = 25.003$, and (fast mixing) gives $\hat{\theta} = 25.083$. For M/M/m/m+r (with true $\theta^* = 5$) obtains $\hat{\theta} = 5.15$. For M/M/Multiple/K, ∞ -SGD obtains $\hat{\theta} = (13.5, 8.3, 5.4)$, close to the ground truth $\theta^* = (15, 10, 5)$. This conclusively shows ∞ -SGD to be a reliable optimization method.

Conclusions

This work introduces ∞ -SGD, the first theoretically principled optimization approach that can accurately learn general *parametric* Continuous Time Markov Chains (CTMCs) from aggregate steady-state observations. Our approach, ∞ -SGD, works even when the observations are over a restricted set of states. We have shown that ∞ -SGD finds significantly better maximum likelihood estimates than the baseline (DC-BPTT) in both a real testbed and synthetic scenarios. Moreover, in the context of queueing systems, ∞ -SGD consistently better extrapolates from training data in light loads to heavy loads in test data. We expect ∞ -SGD to be a useful tool in other tasks where parametric models are needed and sequence data is only available as aggregate frequencies.

Acknowledgement

This work has been sponsored in part by the ARO, under the U.S. Army Research Laboratory contract number W911NF-09-2-0053, the Purdue Integrative Data Science Initiative, and the National Science Foundation grants CNS-1717493, OAC-1738981, and CCF-1918483.

References

- Amazon. 2019. Amazon EC2 Auto Scaling. <https://aws.amazon.com/ec2/autoscaling/>.
- Armbrust, M.; Fox, A.; Griffith, R.; Joseph, A. D.; Katz, R.; Konwinski, A.; Lee, G.; Patterson, D.; Rabkin, A.; Stoica, I.; and Zaharia, M. 2010. A view of cloud computing. *Communications of the ACM* 53(4):50–58.
- Bernstein, G., and Sheldon, D. 2016. Consistently estimating Markov chains with noisy aggregate data. In *Artificial Intelligence and Statistics*, 1142–1150.
- Bhatnagar, S., and Borkar, V. S. 1998. A two timescale stochastic approximation scheme for simulation-based parametric optimization. *Probability in the Engineering and Informational Sciences* 12(4):519–531.
- Cao, L.; Fahmy, S.; Sharma, P.; and Zhe, S. 2018. Data-driven resource flexing for network functions virtualization. In *Proc. ANCS*.
- Colson, B.; Marcotte, P.; and Savard, G. 2007. An overview of bilevel optimization. *Annals of operations research* 153(1):235–256.
- ETSI. 2014. ETSI Network Functions Virtualisation (NFV) Architectural Framework. http://www.etsi.org/deliver/etsi_gs/NFV/001_099/002/01.02.01_60/gs_NFV002v010201p.pdf.
- Filippone, M., and Engler, R. 2015. Enabling scalable stochastic gradient-based inference for gaussian processes by employing the unbiased linear system solver (ulisse). *arXiv preprint arXiv:1501.05427*.
- Google Cloud. 2019. Autoscaling Groups of Instances. <https://cloud.google.com/compute/docs/autoscaler/>.
- Hashimoto, T.; Gifford, D.; and Jaakkola, T. 2016. Learning population-level diffusions with generative rnns. In *International Conference on Machine Learning*, 2417–2426.
- Jain, R., and Ramakrishnan, K. K. 1988. Congestion avoidance in computer networks with a connectionless network layer, part i: Concepts, goals and methodology. In *Proc. Computer Networking Symposium, Washington, D.C.*, 134–143.
- Jensen, A. 1953. Markoff chains as an aid in the study of markoff processes. *Skand. Aktuarietidskrift* 36:87–91.
- Kalbfleisch, J. D., and Lawless, J. F. 1984. Least-squares estimation of transition probabilities from aggregate data. *Canadian Journal of Statistics* 12(3):169–182.
- Khazaei, H.; Mistic, J.; and Mistic, V. B. 2012. Performance analysis of cloud computing centers using m/g/m/m+r queueing systems. *IEEE Transactions on Parallel and Distributed Systems* 23(5).
- Kumar, R.; Tomkins, A.; Vassilvitskii, S.; and Vee, E. 2015. Inverting a Steady-State. In *WSDM*.
- Marcus, G. F. 1998. Rethinking eliminative connectionism. *Cognitive psychology* 37(3):243–282.
- Maystre, L., and Grossglauser, M. 2015. Fast and Accurate Inference of Plackett–Luce Models. In *NIPS*.
- McLeish, D. 2011. A general method for debiasing a monte carlo estimator. *Monte Carlo Methods Appl.*
- Meyer, C. D. 1989. Stochastic complementation, uncoupling Markov chains, and the theory of nearly reducible systems. *SIAM Rev.* 31(2):240–272.
- Miller, G. A. 1952. Finite markov processes in psychology. *Psychometrika* 17(2):149–167.
- Newman, A. 2017. *Benchmarking Java logging frameworks*. <https://www.loggly.com/blog/benchmarking-java-logging-frameworks/>.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in pytorch. <https://pytorch.org>.
- Pierson, E.; Koh, P. W.; Hashimoto, T.; Koller, D.; Leskovec, J.; Eriksson, N.; and Liang, P. 2018. Inferring multi-dimensional rates of aging from cross-sectional data. *arXiv preprint arXiv:1807.04709*.
- Ragain, S., and Ugander, J. 2016. Pairwise Choice Markov Chains. In *NIPS*.
- Szabó, Z.; Sriperumbudur, B. K.; Póczos, B.; and Gretton, A. 2016. Learning theory for distribution regression. *The Journal of Machine Learning Research* 17(1):5272–5311.
- Van Der Plas, A. P., et al. 1983. On the estimation of the parameters of markov probability models using macro data. *The Annals of Statistics* 11(1):78–85.
- Weiss, G. M., and Hirsh, H. 1998. Learning to predict rare events in event sequences. In *Proceedings of the 4th International Conference on Knowledge Discovery and Data Mining*.
- Werbos, P. J. 1990. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE* 78(10):1550–1560.
- Xu, K.; Srivastava, A.; and Sutton, C. 2019. Variational russian roulette for deep bayesian nonparametrics. In *International Conference on Machine Learning*, 6963–6972.
- Yang, B.; Tan, F.; Dai, Y.; and Guo, S. 2009. Performance evaluation of cloud service considering fault recovery. In *Proc. of 1st International Conference on Cloud Computing*.