
Bringing Efficiency and Interpretability to Learned TCP Congestion Control

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Recent research in TCP congestion control (CC) has witnessed tremendous success
2 with deep reinforcement learning (RL) approaches, which use feedforward neural
3 networks (NN) to tackle complex environment conditions and make better decisions.
4 However, these “black box” policies lack interpretability, and reliability and, more
5 importantly, cannot operate under the TCP datapath’s ultra-contingent latency
6 and computational constraints. This paper proposes a novel two-stage solution
7 to achieve the best of both worlds: first to train a deep RL agent, then distill its
8 (over-)parameterized NN policy into white-box, light-weight rules in the form
9 of *symbolic* expressions that are much easier to understand and to implement
10 in constrained environments. At the core of our proposal is a novel **symbolic**
11 **branching** algorithm that allows the rule to be “context-aware” of various network
12 conditions, eventually converting the NN policy into a symbolic tree. The distilled
13 symbolic rules preserve and often improve performance over state-of-the-art NN
14 policies while being orders of magnitude faster and interpretable. We validate
15 the performance of our distilled symbolic rules on both simulation and emulation
16 network systems. Our code will be released upon acceptance.

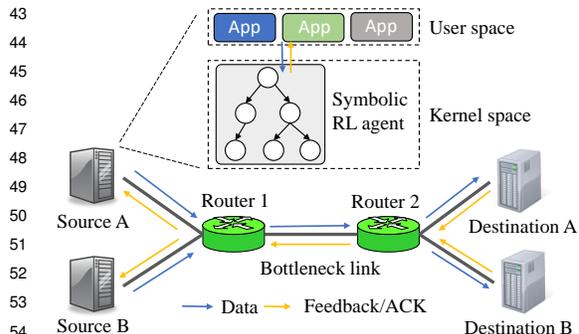
17 1 Introduction

18 Congestion control (CC) is fundamental to Transmission Control Protocol (TCP) communication.
19 Congestion occurs when the data volume sent to a network reaches or exceeds its maximal capacity,
20 in which case the network drops excess traffic, and the performance unavoidably declines. CC
21 mitigates this problem by carefully adjusting the data transmission rate based on the inferred network
22 capacities, aiming to send data as fast as possible without creating congestion. For instance, a classic
23 and best-known strategy, Additive-Increase/Multiplicative-Decrease (AIMD) [1], gradually increases
24 the sending rate when there is no congestion but exponentially reduces the rate when the network is
25 congested. It ensures that TCP senders fairly share the network capacity in the converged state.

26 Figure 1 shows an example where two TCP connections share a link between routers 1 and 2. When
27 the shared link becomes a bottleneck, the CC algorithms running on sources A and B will alter the
28 traffic rate based on the feedback to avoid congestion. Efficient and interpretable CC algorithms have
29 been the bedrock for network services such as DASH video streaming, VoIP (voice-over-IP), VR/AR
30 games, and IoT (Internet of Things), which ride atop the TCP protocol.

31 However, it is nontrivial to design a high-performance CC algorithm. Over a long history, tens of CC
32 proposals have been made, all with different metrics and strategies to infer and address congestion,
33 and new designs are still emerging even today [2, 3]. There are **two main challenges** when designing
34 a CC algorithm. ① it needs to precisely infer whether the network is congested, and if so, how
35 to adjust the sending rate, based on only *partial or indirect observations*. Note that CC runs on
36 end hosts while congestion happens in the network, so CC algorithms cannot observe congestion
37 directly. Instead, it can only rely on specific signals to infer the network status. For instance, TCP
38 Cubic [4] uses packet loss as a congestion signal, and TCP Vegas [5] opts for delay increase. ② CC

39 algorithms operate within the OS kernel, where the computing and memory resources are limited,
 40 and they need to make real-time decisions to adjust the traffic rates frequently (e.g., per round-trip
 41 time). Therefore, the algorithm must be extraordinarily *efficient*. Spending a long time to compute an
 42 action will significantly offset the benefit of congestion control.



43
44
45
46
47
48
49
50
51
52
53
54
55 Figure 1: Overview of a congestion control agent’s
56 role in the network. Multiple senders and receivers
57 share a single network link controlled by the agent,
58 which dynamically modulates the sending rates
59 conditioned on feedback from receivers.

60
61 that the CC decisions are dependent on traffic patterns and network circumstances that can be ex-
 62 ploited by deep reinforcement learning (RL) to learn a policy for each scenario. The learned policy
 63 can perform more flexible and accurate rate adjustment by discovering a mapping from experience,
 64 which can adapt to different network conditions and save manual tuning efforts for high performance.

65 Most notably, Aurora [7], a deep RL framework for Performance-oriented Congestion Control (PCC),
 66 trains a well-established PPO [12] agent to suggest sending rates as actions by observing the network
 67 statistics such as latency ratio, send ratio, and sent latency inflation. It achieves competitive results on
 68 emulation environments Mininet [13] and Pantheon [14], demonstrating the potential of deep learning
 69 approaches to outperform algorithmic, hand-crafted ones. Despite its immense success, Aurora
 70 being a neural network based approach, is essentially a black-box to users or, in other words, lacks
 71 explicit declarative knowledge [15]. They also require exponentially more computation resources
 72 than traditional, algorithmic methods such as the widely deployed TCP-CUBIC [4].

73 1.1 Our Contributions

74 In this work, we develop a new algorithmic framework for performance-oriented congestion control
 75 (PCC) agents, which can ① run as fast as classical algorithmic methods; ② adjust the rate as
 76 accurately as data-driven RL methods; and ③ enjoy good interpretability.

77 We solve this problem by grasping the opportunity enabled by advances in symbolic regression
 78 [16–19]. Symbolic regression bridges the gap between the infeasible search directly in the enormous
 79 symbolic algorithms space and the differentiable training of overparameterized and un-interpretable
 80 neural networks.

81 At a high level, one can first train an RL agent through gradient descent, then distill the learned
 82 policy to obtain the data-driven optimized yet fully interpretable symbolic rules. This results in
 83 a set of symbolic rules through data-driven optimization that meets TCP CC’s extreme efficiency
 84 and reliability demands. However, considering the enormous volume of discrete symbolic space, it
 85 is challenging to learn effective symbolic rules from scratch directly. Therefore, in this paper, we
 86 adopt a two-stage approach: we first train a deep neural network policy with reinforcement learning
 87 mimicking Aurora [7], and then distill the resultant policy into numerical symbolic rules, using
 88 symbolic regression (SR).

89 As the challenge, directly applying symbolic regression out of the box does not yield a sufficient
 90 versatile expression that captures diverse networking conditions. We hence propose a novel branching
 91 technique for training and then aggregating a number of SymbolicPCC agents, each of which cater to a
 92 subset of the possible network conditions. Specifically, we have multiple agents, each called a branch,
 93 and employ a light-weight “branch decider” to choose between the branches during deployment.

In the long history of congestion control, most algorithms are implemented with manually-designed heuristics, including New Reno [6] Vegas [5], Cubic [4], and BBR [2]. In TCP New Reno, for example, the sender doubles the number of transmitted packets every RTT before reaching a predefined threshold, after which it sends one more packet every RTT. If there is a timeout caused by packet loss, it halves the sending rate immediately. Manually crafted CCs have been shown to be suboptimal and cannot support cases that escape the heuristics [7]. For example, packet loss-based CCs like Cubic [4] cannot distinguish packet drops caused by congestion or non-congestion-related events [7].

Researchers have tried to construct CC algorithms with machine learning approaches to address these limitations [7–11]. The insight is

94 In order to create the branching conditions we partition the network condition space into adjacent
 95 non-overlapping contexts, then regress symbolic rules in each context. With this modification, we
 96 enhance the expressiveness of the resulting SR equation and overcome the bias of traditional SR
 97 algorithm to output rules mostly using numerical operators. Our concrete technical contributions are
 98 summarized as follows:

- 99 • We propose a symbolic distillation framework for TCP congestion control, which improves upon
 100 the state-of-the-art RL solutions. Our approach, **SymbolicPCC**, consists of two stages: first
 101 training an RL agent and then distilling its policy network into ultra-compact and interpretable
 102 rules in the symbolic expression form.
- 103 • We propose a novel branching technique that advances existing symbolic regression techniques
 104 for training and aggregating multiple context-dependent symbolic policies, each of which
 105 specializes for its own subset of network conditions. A branch decider driven by light-weight
 106 classification algorithms determines which symbolic policy to use.
- 107 • Through our simulation and emulation experiments, SymbolicPCC achieves highly competitive
 108 or even stronger performance results compared to their teacher policy networks while running
 109 orders of magnitude faster. Our approach enables congestion control RL that is as light-weight
 110 and interpretable as conventional algorithms while narrowing their performance gap.

111 2 Related Works

112 Conventional TCP CC adopts a
 113 heuristic-based approach where the
 114 heuristic functions are manually
 115 crafted to adjust the traffic rate in a de-
 116 terministic manner. Some proposals
 117 use packet loss as a signal for network
 118 congestion, e.g., Cubic [4], Reno [20],
 119 and NewReno [6], and others rely on
 120 the variation of delay, e.g., Vegas [5].
 121 Other CC designs combine packet loss
 122 and delay [21, 22]. Recently, different CC techniques specialized for data-center networks are also
 123 proposed [2, 3, 23].

124 Researchers have also investigated the use of machine learning to construct better heuristics. In-
 125 digo [10] and Remy [11] use offline learning to obtain high-performance CC algorithms. PCC [24]
 126 and PCC Vivace [9] opt for online learning to avoid any hardwired mappings between states and
 127 actions. Aurora [7] utilizes deep reinforcement learning to obtain a new CC algorithm running in
 128 the user space. Orca [8] improves upon Aurora and designs a user-space CC agent that infrequently
 129 configures kernel-level CC policies. Our proposal further improves these work.

130 These methods, although proven to be high-performing, lack any type of interpretability. On the other
 131 hand, Symbolic regression methods have [16–19] recently emerged for discovering underlying math
 132 equations that govern the observed data. Algorithms with such a property are more favorable for
 133 real-world deployment as they output white-box rules. [17] use genetic programming based method
 134 to generate a list of candidate equations, evaluate and filter out those bad performing ones, then
 135 randomly permute the remaining ones to evolve into better candidates. [19] use a recurrent neural
 136 network to perform a search in the symbolic space.

137 We thus propose to synergize such numerical and data-driven approaches using symbolic regression
 138 (SR) in the congestion control domain. We use SR by following a post-hoc method of first training an
 139 RL algorithm then distilling it into its symbolic rules. Earlier methods that follow a similar procedure
 140 do exist, e.g., [25] distills the learned policy as a soft decision tree. They work on visual RL where
 141 the image observations are coarsely quantized into 10×10 cells, and the soft decision tree policy
 142 is learned over the 100 dimensional space. [26] also aims to learn abstract rules using a common
 143 sense-based approach by modifying Q learning algorithms. Nevertheless, they fail to generalize
 144 beyond the specific simple grid worlds they were trained in. [27] learns from boolean feature sets, [28]
 145 directly approximates value functions based on a state-transition model, [29] optimizes risk-seeking
 146 policy gradients. Other works on abstracting pure symbolic rules from data include attention-based
 147 methods [30], visual summary based methods [31], reward decomposition based methods [32], casual

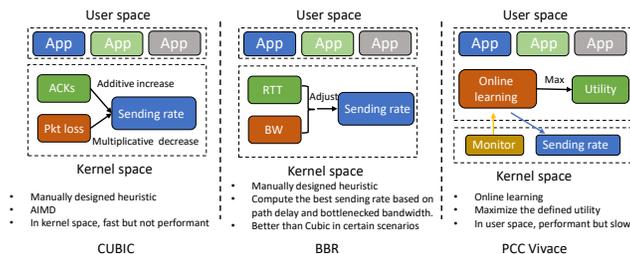


Figure 2: Overview of Conventional Baselines

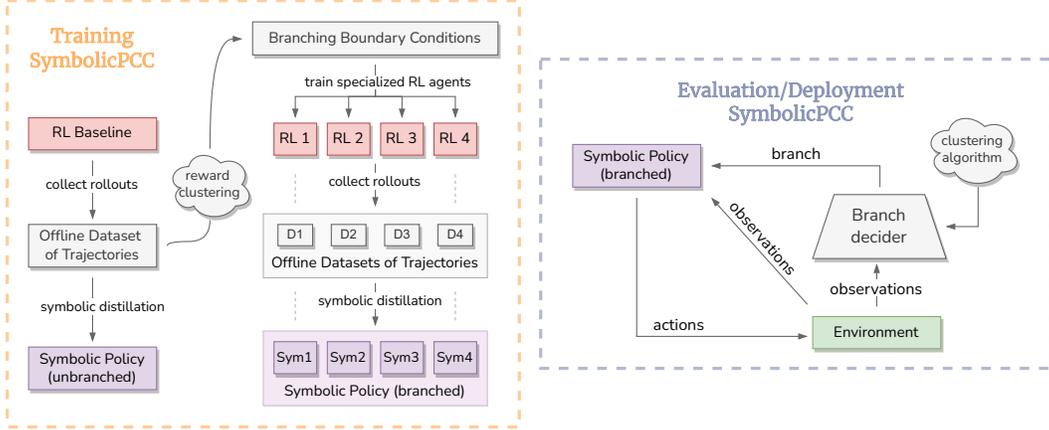


Figure 3: The proposed SymbolicPCC training and evaluation technique: A baseline RL agent is first trained then evaluated numerous times with the roll-outs being saved. Directly distilling out from this data provides a baseline symbolic policy. A light-weight clustering algorithm is used to cluster from the roll-out dataset, non-overlapping subsets of network conditions (aka. branching conditions) that achieve similar rewards. Separate RL agents are then trained on each of these network contexts and distilled into their respective symbolic rules. During the evaluation, the labels from the clustering algorithm are re-purposed to classify which branch is to be taken given the observation statistics. The chosen symbolic branch is then queried for the action.

148 model based methods [33], markov chain based methods [34], and case-based expert-behaviour
 149 retrieval methods [35].

150 3 Methodology

151 Inspired from the idea of “teacher-student knowledge distillation” [36], our symbolic distillation
 152 technique is two-staged—first train regular RL agents (as the *teacher*), then distill the learned policy
 153 networks into their white-box symbolic counterparts (as the *student*). In Section 3.1 we follow [7]’s
 154 approach in Aurora and the training of teacher agents on the PCC-RL gym environment. We also
 155 briefly discuss the approach of applying symbolic regression to create a light-weight numerical-driven
 156 expression that approximates a given teacher’s behavior. In Section 3.2 we look at the specifics of
 157 symbol spaces and attach internal attributes to aid long-term planning. Finally, in Section 3.3 we
 158 discuss our novel branching algorithm as a method for training, then *ensembling* multiple context-
 159 dependent symbolic agents during deployment.

160 3.1 Preliminaries: The PCC-RL Environment and the Symbolic Distillation Workflow

161 PCC-RL [7] is an open-source RL testbed for simulation of congestion control agents based on the
 162 popular OpenAI Gym [37] framework. We adopt it as our main playground. It formulates congestion
 163 control as a sequential decision making problem. Time is first divided into multiple periods called
 164 MIs (monitor intervals), following [24]. At the onset of each MI, the environment provides the agent
 165 with the history of statistic vector observations over the network, and the agent responds with adjusted
 166 sending rates for the following MI. The sending rate remains fixed during a single MI.

167 The network statistics provided as observations to the congestion control agent are ① the latency
 168 inflation, ② the latency ratio, and ③ the sending ratio. The agent is guided by reward signals based
 169 on its ability to react appropriately when detecting changes and trends in the vector statistics of
 170 the PCC-RL environment. It is provided with positive rewards for higher values of throughputs
 171 (packets/second) while being penalized for higher values of latency (seconds) and loss (ratio of sent
 172 vs. acknowledged packets).

173 **Training of Teacher Agents:** We first proceed to train RL agents using the PPO algorithm [12]
 174 similar to Aurora [7] in the PCC-RL gym environment till convergence. Although they statistically
 175 perform very well [7], the PPO agents are entirely black-boxes; this makes it difficult to explain its
 176 underlying causal rules directly. Also, their over-parameterized neural network forms incur high
 177 latency. Hence, we choose to indirectly learn the symbolic representations using a student-teacher
 178 type knowledge distillation approach based on the teacher’s (in this case, the RL agent) behaviors.

179 **Distillation of Student Agents:** Using the teacher agents, we collect complete trajectories, formally
 180 known as roll-outs in RL, in an inference mode—deterministic actions are enforced. The observations
 181 and their corresponding teacher actions are MI-aligned and stored as an offline dataset. Note that
 182 this step is only performed once at the start of the distillation procedure and is reused in each of its
 183 iterative steps. A search space of operators and operands is also initialized (details are discussed
 184 shortly in Section 3.2). Guesses for possible symbolic relations are taken, composed of random
 185 operators and operands from their respective spaces. The stored observation trajectories are then
 186 re-evaluated based on this rule to output corresponding actions. The cross-entropy loss with respect
 187 to the teacher model’s actions from the same dataset is used as feedback. This feedback drives
 188 the iterative mutation and pruning following a genetic programming technique [38, 39]. The best
 189 candidate policies are collected and forwarded to the next stage. If the tree fails to converge or does
 190 not reach a specific threshold of acceptance, the procedure is restarted from scratch.

191 3.2 A Symbolic Framework for Congestion Control

192 **Defining the Symbol Space for CC:** Unlike visual RL [40], the PCC observation space is vector-
 193 based, hence we directly plug them into the search space of our numerical-driven symbolic algorithm.
 194 We henceforth call these observations as `vector statistic symbols`. The distillation procedure
 195 as described earlier learns to *chain* these `vector statistic symbols` using a pre-defined operator
 196 space. Specifically, we employ three types of numerical operators. The first type of operators are
 197 arithmetic based which include $+$, $-$, $*$, $/$, \sin , \cos , \tan , \cot , $(\cdot)^2$, $(\cdot)^3$, $\sqrt{\cdot}$, \exp , \log , $|\cdot|$. The second
 198 type of operators are Boolean logic, such as $is(x < y)$, $is(x \leq y)$, $is(x == y)$, $a \mid b$, $a \& b$, and $\neg a$.

199 We also utilize a third type of *high level* operators – namely the `slope_of` (`observation history`)
 200 which provides the average slope of an array of observations, and `get_value_at` (`observation`
 201 `history`, `index`). The slope operator is especially useful when trying to detect *trends* of a specific
 202 statistic vector over the provided monitor interval. For instance, identifying latency increase or
 203 decrease trends serves as one of the crucial indicators for adjusting sending rates. Meanwhile the
 204 index operator is observed from our experiments to be implicitly used for *immediate responding*—i.e.,
 205 based on the latest observations.

206 We note that the underlying decision procedure of the policy network could be efficiently represented
 207 in the form of a high-fidelity tree-shaped form similar to Figure 4. This *decision tree* contains said
 208 *condition nodes* and *action nodes*. Each condition node forks into two leaf nodes based on the
 209 Boolean result of its symbolic composition.

210 **Attributes for Long Term CC Planning:** In addition to having these operators and operands as
 211 part of the symbolic search space, we also attach a few attributes/flags to the agent which are shared
 212 across consecutive MI processing steps and help with long-term planning. One emerging behaviour
 213 in our SymbolicPCC agents is to use this attribute for *remembering* if the agent is in the process of
 214 recovering from a network overload or if the network is stable. Indeed, a more straightforward option
 215 for such “multi-MI tracking” would be to just provide a longer history of the vector statistics into
 216 the searching algorithms. But this quickly becomes infeasible due to an exponential increase of the
 217 possible symbolic expressions with respect to the length of `vector statistic symbols`.

218 3.3 Novel Branching Algorithm: Switching between Context-Dependent Policies

219 Unlike traditional visual RL environments, congestion control is a technically more demanding task
 220 due to the variety of possible network conditions. The behaviour of the congestion control agent
 221 will improve if its response is conditioned on the specific network context. However, as this context
 222 cannot be known by the congestion control agent, the traditional algorithms such as TCP Cubic [4]
 223 are forced to react in a slow and passive manner to generalize to both slow-paced and fast-paced
 224 conditions.

225 Hence, we propose to create n non-overlapping contexts for different network conditions, namely—
 226 bandwidth, latency, queue size, and loss rate. We then train n individual RL agents in the PCC Gym
 227 by exposing them only to their corresponding network conditions. We thus have a diverse set of
 228 teachers which are each highly performant in their individual contexts. Following the same approach
 229 as described in Section 3.1, each of the agents are distilled—each one called a *branch*. Finally,
 230 during deployment, based on the inference network conditions, the branch with the closest matching
 231 boundary conditions/context is selected and the corresponding symbolic policy is used.

232 **Partitioning the Networking Contexts:** A crucial point to note in the proposed branching pro-
 233 cedure is to identify the most suitable branching context boundary values. In other words, the best

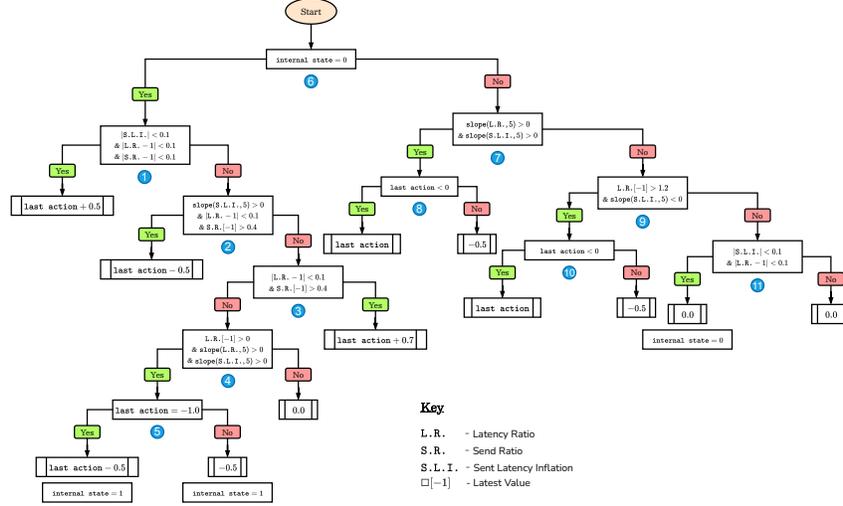


Figure 4: A distilled symbolic policy from the baseline RL Agent in the PCC-RL Environment. Condition nodes are represented as rectangular blocks and action nodes as process blocks.

234 boundary conditions for grouping need to be statistically valid, and plain hand-crafted boundaries are
235 not optimal. This is because we do not have ground truths of any of the network conditions [41], let
236 alone four of them together. Therefore, we first use a trained a RL agent on the default (maximal)
237 bounds of network conditions (hereinafter known as the “baseline” agent). We then evaluate the
238 baseline agent on multiple regularly spaced intervals of bandwidth, latency, queue size, and loss
239 rate and store their corresponding rewards as well as observation trajectories. To create the optimal
240 groupings, we simply use KMeans [42] to cluster the data based on their rewards. Due to the inherent
241 proportional relation of difficulty (or in this case the ballpark of rewards) with respect to a network
242 context, clear boundaries for the branches can be obtained by inspecting the extremes of each network
243 condition within a specific cluster. Our experimentally obtained branching conditions are further
244 discussed in Section 4.2 and Table 1.

245 **Branch Decider.** Since the network context is not known during deployment, this creates the need
246 for a branch decider module. The branch decider reuses clusters labels from the training stage
247 for a K Nearest Neighbours [43] classification. The light-weight distance based metric is used to
248 classify the inference-time observation into one of the training groupings, and thereby executing the
249 corresponding branch’s symbolic policy. Figure 3 illustrates our complete training and deployment
250 techniques.

251 Lastly, in order to accommodate for branching, we have yet another long-term tracking attribute that
252 stores a history of branches taken in order to smooth over any erratic bouncing between branches
253 which are in non-adjacent contexts.

254 4 Experimental Settings and Results

255 Next, we discuss the abstract rules uncovered by SR, as well as validate the branching contexts. In
256 Sections 4.3, 4.4, and 4.5 we provide emulation results on Mininet [13], a widely-used network
257 emulator that can emulate a variety of networking conditions. Lastly in Section 4.6, we compare the
258 compute requirements and efficiencies of SymbolicPCC with conventional algorithms, RL-driven
259 methods as well as their pruned and quantized variants.

260 4.1 Interpreting the Symbolic Policies

261 The baseline symbolic policy distilled from the baseline RL agent is represented in its decision tree
262 form in Figure 4. One typical CC process presented by the tree is increasing the sending rate until the
263 network starts to “choke” and then balancing around that rate. This process is guided with a series
264 of conditions regarding to inflation and ratio signals, marked with circled numbers in Figure 4. The
265 detailed explanation is in the following.

266 Condition node ① checks whether the vector statistic symbols are all stable—namely,
267 whether the latency inflation is close to zero, while latency ratio and send ratio are close to one.

Table 1: The baseline network conditions and resultant branching boundary values (contexts) for each branch after clustering. The rewards centroid refers to the reward value at cluster center of that specific branch.

Branch	Rewards Centroid	Bandwidth (pps)	Latency (sec)	Queue Size (packets)	Loss Rate (%)
Baseline	-	100 - 500	0.05 - 0.5	2 - 2981	0.00 - 0.05
Branch 1	95.84	100 - 200	0.35 - 0.5	2 - 2981	0.04 - 0.05
Branch 2	576.57	200 - 250	0.25 - 0.35	2 - 2981	0.02 - 0.03
Branch 3	1046.46	250 - 350	0.15 - 0.25	2 - 2981	0.02 - 0.03
Branch 4	1516.70	350 - 500	0.05 - 0.15	2 - 2981	0.00 - 0.02

268 The sending rate starts to grow if the condition holds. Condition node (2) identifies if the network
 269 is in a over-utilized status `slope_of` (latency inflation) increasing as the key indicator. It
 270 the condition is true, the acceleration of sending rate will be reduced appropriately. On the other
 271 hand, condition node (3) is activated when the initial sending rate is too low or has been reduced
 272 extensively due to (2). (4) is evaluated when major network congestion starts to occur due to increased
 273 sending rates from the earlier condition nodes. It checks both latency inflation and latency ratios
 274 in an increasing state. Its child nodes start reducing the sending rates and also flip the `internal`
 275 `state` attribute to 1. The latter is used to track if the agent is recovering from network congestion.
 276 On the “False” side of (6) (i.e. `internal state = 1`), (7) and (8) tackle two stages of recovery,
 277 where the latency inflation ratio starts plateauing and then starts reducing. (11) indicates that stable
 278 conditions have been recovered again and the agent is at an optimal sending rate. The `internal`
 279 `state` is flipped back again to 0 after this recovery.

280 4.2 Inspecting the Branching Conditions

281 As discussed in Section 3.3, a light-weight clustering algorithm divides the network conditions into
 282 multiple non-overlapping subsets. Table 1 summarizes the obtained boundary values. The baseline
 283 agent is trained on all possible bandwidth, latency, queue size, and loss rate values, as depicted in
 284 the first row. During the evaluation, bandwidth, latency, and loss rate are tested on linearly spaced
 285 values of step sizes 50, 0.1, and 0.01, while queue sizes are exponentially spaced by powers of e^2
 286 respectively. The rewards of the saved roll-outs are clustered using K-Means Clustering, and the
 287 optimal cluster number is found to be 4 using the popular elbow curve [44] and silhouette analysis
 288 [45] methods. By observing the maximum and minimum of each network condition individually in
 289 the 4 clusters, respective boundary values are obtained. A clear relation discovered is that higher
 290 bandwidths and lower latencies are directly related to higher baseline rewards.

291 **Remark 1: Exceptions for non-overlapping contexts.** It is also to be noted that no such trend was
 292 found between the queue size and rewards, and hence all the 4 resultant branches were given the
 293 same queue size. A similar exception was made for the loss rates of Branches 2 and 3.

294 **Remark 2: Interpreting the symbolic policies branches.** All the 4 distilled symbolic trees from the
 295 specialized RL agents possess high structural similarity and share similar governing rules as to that
 296 of the baseline agent in Section 4.1. They majorly differ in the numerical thresholds and magnitudes
 297 of action nodes, i.e., by varying their “reaction speeds” and “reaction strengths”, respectively.

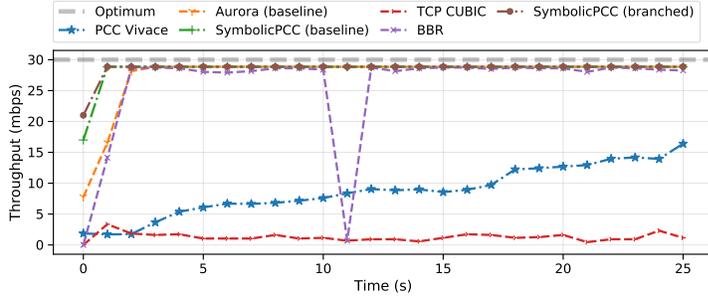
298 4.3 Emulation Performance on Lossy Network Conditions

299 The ability to differentiate between congestion-induced and random losses is essential to any PCC
 300 agent. Figure 5a¹ shows a 25-second trace of throughput on a link where 1% of packets are
 301 randomly dropped [46]. As the link’s bandwidth is set to 30 Mbps, the ideal congestion control
 302 would aim to utilize it fully as depicted by the gray dotted line. Baseline SymbolicPCC shows near-
 303 ideal performance with its branched version pushing boundaries further. In contrast, conventional
 304 algorithms, especially TCP CUBIC [4], repeatedly reduces its sending rates as a response to the
 305 random losses. Quantitative measures of mean square error with respect to the ideal line are provided
 306 in Table 2 as “Lossy Δ_{opt}^2 ”. This result proves that SymbolicPCC can effectively differentiate
 307 between packet loss caused by randomness and real network congestion.

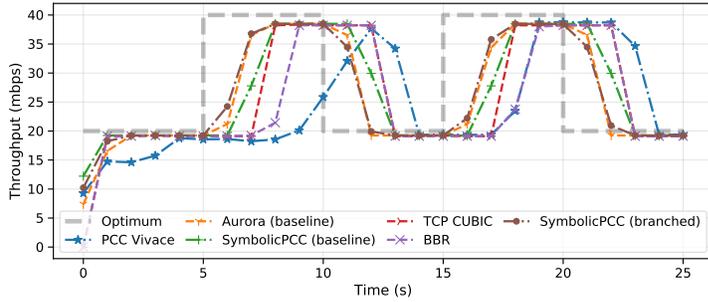
¹Interestingly, the figure shows that BBR has rate drop around 11th second. This is a limitation of the BBRv1 design—it reduces sending rate if the `min_rtt` has not been in 10s, which is triggered because the RTT in our setup is very stable. We have confirmed this with the BBR team at Google.

308 4.4 Emulation Performance under Network Dynamics

309 Unstable network conditions are common in the real world and this test benchmarks the agent’s ability of quickly responding to network dynamics. Figure 5b shows our symbolic agent’s ability to handle such conditions. The benefits of our novel branching algorithm and switching between agents which each specializes in their own network context is clearly visible from faster response speeds. In this case, the link was configured with its bandwidth alternating between 20 Mbps and 40 Mbps every 5 seconds with no loss. Quantitative results from Table 2 show the mean square error with respect to the ideal CC as “Unstable Δ_{opt}^2 ”.



(a) A 25-second throughput trace for TCP CUBIC, PCC-Vivace, BBR, Aurora, and our SymbolicPCC variants on a 30 Mbps bandwidth link with 2% random loss, 30 ms latency, and a queue size of 1000.



(b) A 25-second throughput trace for TCP CUBIC, PCC Vivace, BBR, Aurora, and our SymbolicPCC variants on a link alternating between 20 and 40 Mbps every 5 seconds with 0% random loss, 30 ms latency, and a queue size of 1000.

334 4.5 Link Utilization and Network Sensitivities

336 Link utilization as measured from the server side is defined as the ratio of average throughput over the emulation period to the available bandwidth. A single link is first configured with defaults of 30 Mbps capacity, 30 ms of latency, a 1000-packet queue, and 0% random loss. To measure the sensitivity with respect to a specific condition, it is independently varied keeping the rest of the conditions constant. An ideal CC preserves high link utilization over the complete range of measurements. From Figure 6, it is observed that our branched SymbolicPCC provides near-capacity link-utilization at most tests and shows improvement over any of the other algorithms.

Figure 5: Emulation performances on different conditions.

345 4.6 Efficiency and Speed Comparisons

346 Since TCP congestion control lies on the fast path, ultra-fast responses are needed from the agents. Due to its GPU compute requirements and slower runtimes, RL-based approaches such as Aurora are constrained in their deployment settings (e.g., user-space decisions). On the other hand, our symbolic policies are entirely composed of numerical operators, making them structurally and computationally minimal. From our results in Table 2, adding the branch decider incurs a slight overhead as compared

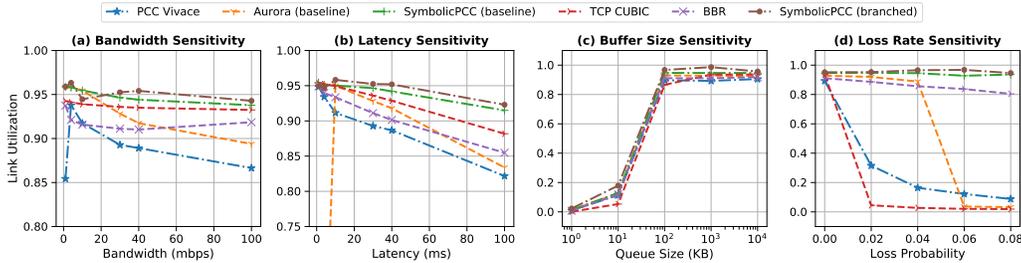


Figure 6: Link-utilization trends as a measure of sensitivities of bandwidth, latency, queue size, and loss rate. Higher values are better.

Table 2: Efficiency and speed comparison of congestion control agents.

Algorithm	Type	FLOPs (\downarrow)	Runtime (μ s) (\downarrow)	Lossy Δ_{opt}^2 (\downarrow)	Oscillating Δ_{opt}^2 (\downarrow)
TCP CUBIC	Conventional	-	< 10	823.02	126.07
PCC Vivace	Conventional	-	< 10	440.55	186.76
Aurora (baseline)	RL-Based	1488	864	26.29	53.22
Aurora (50% pruned)	RL-Based	744	781	27.37	61.85
Aurora (80% pruned)	RL-Based	298	769	48.13	79.80
Aurora (95% pruned)	RL-Based	74	703	83.66	103.53
Aurora (quantized)	RL-Based	835	810	142.92	88.45
SymbolicPCC (baseline)	Symbolic	48	23	7.29	85.03
SymbolicPCC (branched)	Symbolic	63	37	4.14	43.83

351 to the non-branched counterpart. Nevertheless, it is preferable due to its increased versatility in
352 different network conditions, as verified by Mininet emulation results. SymbolicPCC enjoys over a
353 **23 \times faster run times** over Aurora, being reasonably comparable to PCC Vivace and TCP CUBIC.
354 We also compare between global magnitude pruned and dynamically quantized versions of Aurora.
355 Although these run faster than their baseline versions, they come at the cost of worse CC performance.

356 5 Discussions and Potential Impacts of SymbolicPCC

357 The interpretability of our approach enables re-
358 searchers to verify the learned model, which is
359 critical for both performance and security ob-
360 jectives in practical deployments. It also pro-
361 vides more insights for networking researchers
362 of what are the key heuristic for TCP CC. More-
363 over, our success of using symbolic distillation
364 for CC also paves the possibility of applying it
365 to other systems and networking applications where performance and interpretability are both key
366 consideration, such as traffic classification and CPU scheduling tasks.

Table 3: Decoupling: symbolic alone helps generalization.

Model	Avg. Rewards (\uparrow)
Aurora	832
Black-box dist. (50%) from Aurora	641
White-box dist. from above model	687

367 **Need for Branching.** The branched training of multiple symbolic models, each in different training
368 regimes, is designed to ease the optimization process. It does **not** directly enforce similarity between
369 solutions for the grouped states – **therefore not causing brittleness**. This is assured as the symbolic
370 model within any branch does not directly perform the same action for all scenarios within its regime,
371 but contains multiple operations within itself to map states to actions based on the network state
372 observed. Also, during the inference/deployment stage, we use the branch-decider network which
373 chooses branches based on the observed state, **not** the bandwidths or latencies (in fact, these measures
374 are **unavailable** to the controller agent and cannot be observed).

375 **Interpretability – a universal boon for ML?** It may be a bit surprising that the distilled symbolic
376 policy outperforms Aurora. A natural question arises if it is due to a generalization amplification
377 that sometimes happens for distillation in general or if it is due to the symbolic representation. We
378 hypothesize that the performance of a symbolic algorithm boils down to the nature of the environment
379 it is employed in. The congestion control problem is predominantly rule-based, with deep RL models
380 brought to devise rules more complex and robust than hand-crafted ones through iterative interaction.
381 It is only natural to observe that symbolic models outperform such PCC RL models when the
382 distillation is composed of a rich operator space and dedicated policy denoising and pruning stages to
383 boost their robustness and compactness further. To justify this, in Table 3 we analyze the performance
384 obtained by **decoupling distillation and symbolic representation**: we first distill a black-box NN
385 half the size of Aurora (“typical KD”) and then further perform symbolic distillation on it.

386 6 Conclusion and Future Work

387 This work studied the distillation of NN-based deep reinforcement learning agents into symbolic poli-
388 cies for performance-oriented congestion control in TCP. Our branched symbolic framework enables
389 better interpretability and efficiency while exhibiting comparable and often improved performances
390 over their black-box teacher counterparts on both simulation and rigorous emulation testbeds. Our
391 results point towards a fresh direction to make congestion control extremely light-weight and more
392 interpretable, via a symbolic design. Our future work aims for more integrated neural-symbolic solu-
393 tions and faster model-free online training/fine-tuning for performance-oriented congestion control.
394 Exploring the fairness between our learned CC and legacy CC is also an interesting next step. Besides,
395 we also aim to apply symbolic distillation to a wider range of systems and networking problems.

396 **References**

- 397 [1] Dah-Ming Chiu and Raj Jain. Analysis of the increase and decrease algorithms for congestion
398 avoidance in computer networks. *Computer Networks and ISDN systems*, 17(1):1–14, 1989.
- 399 [2] Neal Cardwell, Yuchung Cheng, C Stephen Gunn, Soheil Hassas Yeganeh, and Van Jacobson.
400 BBR: congestion-based congestion control. *Communications of the ACM*, 60(2):58–66, 2017.
- 401 [3] Gautam Kumar, Nandita Dukkipati, Keon Jang, Hassan MG Wassel, Xian Wu, Behnam Montaz-
402 eri, Yaogong Wang, Kevin Springborn, Christopher Alfeld, Michael Ryan, et al. Swift: Delay
403 is simple and effective for congestion control in the datacenter. In *Proceedings of the Annual
404 conference of the ACM Special Interest Group on Data Communication on the applications,
405 technologies, architectures, and protocols for computer communication*, pages 514–528, 2020.
- 406 [4] Sangtae Ha, Injong Rhee, and Lisong Xu. Cubic: a new tcp-friendly high-speed tcp variant.
407 *ACM SIGOPS operating systems review*, 42(5):64–74, 2008.
- 408 [5] Lawrence S Brakmo, Sean W O’Malley, and Larry L Peterson. Tcp vegas: New techniques
409 for congestion detection and avoidance. In *Proceedings of the conference on Communications
410 architectures, protocols and applications*, pages 24–35, 1994.
- 411 [6] Sally Floyd, Tom Henderson, and Andrei Gurtov. Rfc3782: The newreno modification to tcp’s
412 fast recovery algorithm, 2004.
- 413 [7] Nathan Jay, Noga Rotman, Brighten Godfrey, Michael Schapira, and Aviv Tamar. A deep
414 reinforcement learning perspective on internet congestion control. In *International Conference
415 on Machine Learning*, pages 3050–3059. PMLR, 2019.
- 416 [8] Soheil Abbasloo, Chen-Yu Yen, and H Jonathan Chao. Classic meets modern: A pragmatic
417 learning-based congestion control for the internet. In *Proceedings of the Annual conference
418 of the ACM Special Interest Group on Data Communication on the applications, technologies,
419 architectures, and protocols for computer communication*, pages 632–647, 2020.
- 420 [9] Mo Dong, Tong Meng, Doron Zarchy, Engin Arslan, Yossi Gilad, Brighten Godfrey, and
421 Michael Schapira. PCC vivace: Online-learning congestion control. In *Proc. NSDI*, 2018.
- 422 [10] Francis Y Yan, Jestin Ma, Greg D Hill, Deepti Raghavan, Riad S Wahby, Philip Levis, and
423 Keith Winstein. Pantheon: the training ground for internet congestion-control research. In *Proc.
424 ATC*, 2018.
- 425 [11] Keith Winstein and Hari Balakrishnan. Tcp ex machina: Computer-generated congestion
426 control. *ACM SIGCOMM Computer Communication Review*, 43(4):123–134, 2013.
- 427 [12] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal
428 policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- 429 [13] Ramon R Fontes, Samira Afzal, Samuel HB Brito, Mateus AS Santos, and Christian Esteve
430 Rothenberg. Mininet-wifi: Emulating software-defined wireless networks. In *2015 11th
431 International Conference on Network and Service Management (CNSM)*, pages 384–389. IEEE,
432 2015.
- 433 [14] Francis Y Yan, Jestin Ma, Greg D Hill, Deepti Raghavan, Riad S Wahby, Philip Levis, and
434 Keith Winstein. Pantheon: the training ground for internet congestion-control research. In *2018
435 {USENIX} Annual Technical Conference ({USENIX}{ATC} 18)*, pages 731–743, 2018.
- 436 [15] Andreas Holzinger. From machine learning to explainable ai. In *2018 world symposium on
437 digital intelligence for systems and machines (DISA)*, pages 55–66. IEEE, 2018.
- 438 [16] Michael Schmidt and Hod Lipson. Distilling free-form natural laws from experimental data.
439 *science*, 324(5923):81–85, 2009.
- 440 [17] Miles Cranmer, Alvaro Sanchez-Gonzalez, Peter Battaglia, Rui Xu, Kyle Cranmer, David
441 Spergel, and Shirley Ho. Discovering symbolic models from deep learning with inductive
442 biases. *arXiv preprint arXiv:2006.11287*, 2020.
- 443 [18] Miles Cranmer. Pysr: Fast & parallelized symbolic regression in python/julia. *GitHub repository*,
444 2020.
- 445 [19] Brenden K Petersen, Mikel Landajuela Larma, T Nathan Mundhenk, Claudio P Santiago, Soo K
446 Kim, and Joanne T Kim. Deep symbolic regression: Recovering mathematical expressions
447 from data via risk-seeking policy gradients. *arXiv preprint arXiv:1912.04871*, 2019.

- 448 [20] Van Jacobson. Congestion avoidance and control. *ACM SIGCOMM computer communication*
449 *review*, 18(4):314–329, 1988.
- 450 [21] Kun Tan, Jingmin Song, Qian Zhang, and Murad Sridharan. A compound tcp approach for
451 high-speed and long distance networks. In *Proceedings-IEEE INFOCOM*, 2006.
- 452 [22] Saverio Mascolo, Claudio Casetti, Mario Gerla, Medy Y Sanadidi, and Ren Wang. Tcp
453 westwood: Bandwidth estimation for enhanced transport over wireless links. In *Proceedings of*
454 *the 7th annual international conference on Mobile computing and networking*, pages 287–297,
455 2001.
- 456 [23] Mohammad Alizadeh, Albert Greenberg, David A Maltz, Jitendra Padhye, Parveen Patel, Balaji
457 Prabhakar, Sudipta Sengupta, and Murari Sridharan. Data center tcp (dctcp). In *Proceedings of*
458 *the ACM SIGCOMM 2010 Conference*, pages 63–74, 2010.
- 459 [24] Mo Dong, Qingxi Li, Doron Zarchy, P Brighten Godfrey, and Michael Schapira. {PCC}: Re-
460 architecting congestion control for consistent high performance. In *12th {USENIX} Symposium*
461 *on Networked Systems Design and Implementation ({NSDI} 15)*, pages 395–408, 2015.
- 462 [25] Youri Coppens, Kyriakos Efthymiadis, Tom Lenaerts, Ann Nowé, Tim Miller, Rosina Weber,
463 and Daniele Magazzeni. Distilling deep reinforcement learning policies in soft decision trees.
464 In *Proceedings of the IJCAI 2019 workshop on explainable artificial intelligence*, pages 1–6,
465 2019.
- 466 [26] Artur d’Avila Garcez, Aimore Resende Riquetti Dutra, and Eduardo Alonso. Towards symbolic
467 reinforcement learning with common sense. *arXiv preprint arXiv:1804.08597*, 2018.
- 468 [27] Andrea Dittadi, Frederik K Drachmann, and Thomas Bolander. Planning from pixels in atari
469 with learned symbolic representations. *arXiv preprint arXiv:2012.09126*, 2020.
- 470 [28] Jiří Kubalík, Jan Žegklitz, Erik Derner, and Robert Babuška. Symbolic regression methods for
471 reinforcement learning. *arXiv preprint arXiv:1903.09688*, 2019.
- 472 [29] Mikel Landajuela, Brenden K Petersen, Sookyung Kim, Claudio P Santiago, Ruben Glatt,
473 Nathan Mundhenk, Jacob F Pettit, and Daniel Faissol. Discovering symbolic policies with deep
474 reinforcement learning. In *International Conference on Machine Learning*, pages 5979–5989.
475 PMLR, 2021.
- 476 [30] Wenjie Shi, Shiji Song, Zhuoyuan Wang, and Gao Huang. Self-supervised discovering of causal
477 features: Towards interpretable reinforcement learning. *arXiv preprint arXiv:2003.07069*, 2020.
- 478 [31] Pedro Sequeira and Melinda Gervasio. Interestingness elements for explainable reinforcement
479 learning: Understanding agents’ capabilities and limitations. *Artificial Intelligence*, 288:103367,
480 2020.
- 481 [32] Zoe Juozapaitis, Anurag Koul, Alan Fern, Martin Erwig, and Finale Doshi-Velez. Explainable
482 reinforcement learning via reward decomposition. In *IJCAI/ECAI Workshop on Explainable*
483 *Artificial Intelligence*, 2019.
- 484 [33] Prashan Madumal, Tim Miller, Liz Sonenberg, and Frank Vetere. Explainable reinforcement
485 learning through a causal lens. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
486 volume 34, pages 2493–2500, 2020.
- 487 [34] Nicholay Topin and Manuela Veloso. Generation of policy-level explanations for reinforcement
488 learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages
489 2514–2521, 2019.
- 490 [35] Santiago Ontanón, Kinshuk Mishra, Neha Sugandh, and Ashwin Ram. Case-based planning and
491 execution for real-time strategy games. In *International Conference on Case-Based Reasoning*,
492 pages 164–178. Springer, 2007.
- 493 [36] Jianping Gou, Baosheng Yu, Stephen J Maybank, and Dacheng Tao. Knowledge distillation: A
494 survey. *International Journal of Computer Vision*, 129(6):1789–1819, 2021.
- 495 [37] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang,
496 and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- 497 [38] Alan M Turing. Computing machinery and intelligence. In *Parsing the turing test*, pages 23–65.
498 Springer, 2009.
- 499 [39] Richard Forsyth. Beagle—a darwinian approach to pattern recognition. *Kybernetes*, 1981.

- 500 [40] Alexander Sieusahai and Matthew Guzdial. Explaining deep reinforcement learning agents
501 in the atari domain through a surrogate model. In *Proceedings of the AAAI Conference on*
502 *Artificial Intelligence and Interactive Digital Entertainment*, volume 17, pages 82–90, 2021.
- 503 [41] Nathan Jay. Continued development of internet congestion control: Reinforcement learning and
504 robustness testing approaches. 2019.
- 505 [42] James MacQueen et al. Some methods for classification and analysis of multivariate observations.
506 In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*,
507 volume 1, pages 281–297. Oakland, CA, USA, 1967.
- 508 [43] Evelyn Fix and Joseph Lawson Hodges. Discriminatory analysis. nonparametric discrimination:
509 Consistency properties. *International Statistical Review/Revue Internationale de Statistique*,
510 57(3):238–247, 1989.
- 511 [44] Robert L Thorndike. Who belongs in the family? *Psychometrika*, 18(4):267–276, 1953.
- 512 [45] Peter J Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster
513 analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.
- 514 [46] Danyang Zhuo, Monia Ghobadi, Ratul Mahajan, Klaus-Tycho Förster, Arvind Krishnamurthy,
515 and Thomas Anderson. Understanding and mitigating packet corruption in data center networks.
516 In *Proceedings of the SIGCOMM Conference*, 2017.

517 Checklist

- 518 1. For all authors...
- 519 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s
520 contributions and scope? [Yes]
- 521 (b) Did you describe the limitations of your work? [Yes] See section 5.
- 522 (c) Did you discuss any potential negative societal impacts of your work? [Yes] See section
523 5
- 524 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
525 them? [Yes]
- 526 2. If you are including theoretical results...
- 527 (a) Did you state the full set of assumptions of all theoretical results? [N/A] We did not
528 include theoretical results.
- 529 (b) Did you include complete proofs of all theoretical results? [N/A]
- 530 3. If you ran experiments...
- 531 (a) Did you include the code, data, and instructions needed to reproduce the main experi-
532 mental results (either in the supplemental material or as a URL)? [No] Our codes and
533 data will be fully released upon acceptance.
- 534 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
535 were chosen)? [Yes]
- 536 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
537 ments multiple times)? [Yes]
- 538 (d) Did you include the total amount of compute and the type of resources used (e.g., type
539 of GPUs, internal cluster, or cloud provider)? [No]
- 540 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 541 (a) If your work uses existing assets, did you cite the creators? [Yes]
- 542 (b) Did you mention the license of the assets? [N/A] The assets we used are open-source.
543 The license information is available online.
- 544 (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
545
- 546 (d) Did you discuss whether and how consent was obtained from people whose data you’re
547 using/curating? [Yes] The data we are using is open source.
- 548 (e) Did you discuss whether the data you are using/curating contains personally identifiable
549 information or offensive content? [No] Our data does not include personally identifiable
550 information or offensive content.

- 551 5. If you used crowdsourcing or conducted research with human subjects...
- 552 (a) Did you include the full text of instructions given to participants and screenshots, if
- 553 applicable? [N/A]
- 554 (b) Did you describe any potential participant risks, with links to Institutional Review
- 555 Board (IRB) approvals, if applicable? [N/A]
- 556 (c) Did you include the estimated hourly wage paid to participants and the total amount
- 557 spent on participant compensation? [N/A]