Online Learning for Traffic Routing under Unknown Preferences

Devansh Jalota[1], Karthik Gopalakrishnan[2], Navid Azizan[3], Ramesh Johari[4], Marco Pavone[2]

[1] Institute for Computational and Mathematical Engineering, Stanford University
[2] Department of Aeronautics and Astronautics, Stanford University
[3] Department of Mechanical Engineering, Massachusetts Institute of Technology
[4] Department of Management Science and Engineering, Stanford University

Introduction and Motivation

- Road **tolls** help cope with efficiency losses due to **selfish** user routing by users to minimize individual travel costs
- However, the efficacy of road tolling schemes relies on having access to **complete information** on users' trip attributes, e.g., their values of times and O-D pairs
- We propose an **online learning** approach to set tolls that preserves **user privacy**

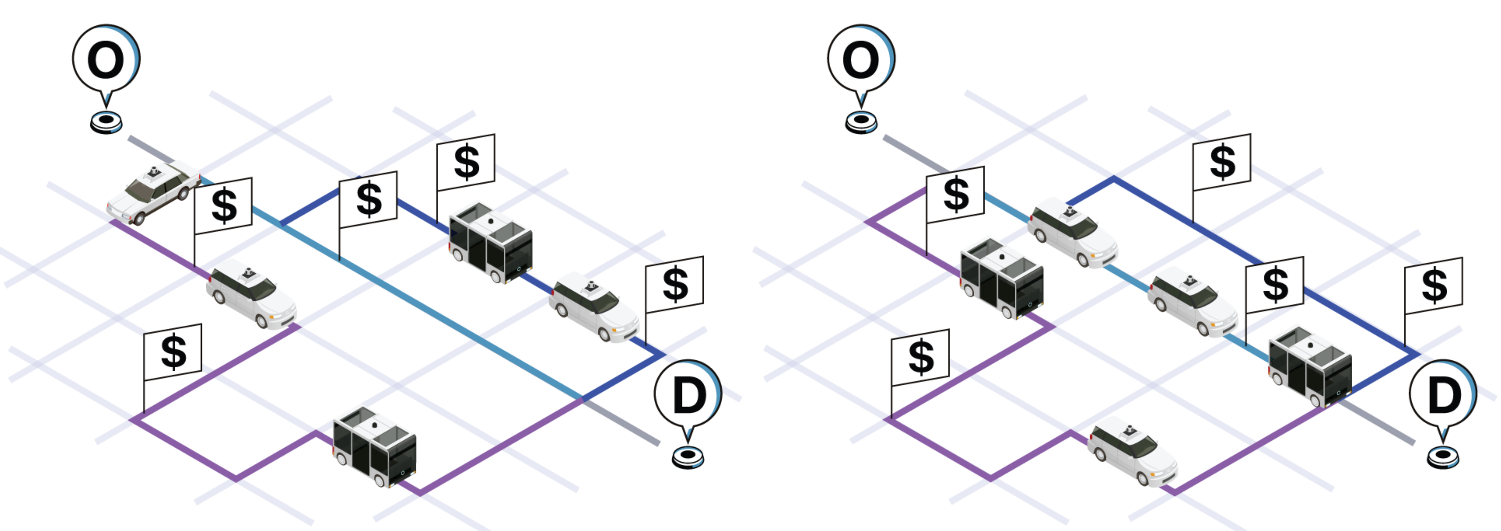

VoT Distribution 1 VoT Distribution 2

For the same tolls, the resulting traffic flows will be different if the value of time distribution of users is different

System Optimization

Minimize Total Travel Cost while satisfying Road capacity constraints

$$\min_{\boldsymbol{f}, \boldsymbol{f}_o} \quad U^* = \sum_{u \in \mathcal{U}} \left(v_u \sum_{P \in \mathcal{P}_u} l_P f_{P,u} + \lambda_u f_{o,u} \right), \leftarrow \textbf{Total Travel Cost}$$

$$\text{s.t.} \quad \sum_{P \in \mathcal{P}_u} f_{P,u} + f_{o,u} = 1, \quad \forall u \in \mathcal{U}, \leftarrow \textbf{Allocation Constraints}$$

$$\boldsymbol{f}_o \geq \boldsymbol{0}, f_{P,u} \geq 0, \quad \forall P \in \mathcal{P}_u, u \in \mathcal{U}$$

$$\sum_{u \in \mathcal{U}} \sum_{P \in \mathcal{P}_u : e \in P} f_{P,u} \leq c_e, \quad \forall e \in E, \leftarrow \textbf{Capacity Constraints}$$

Optimal Tolls can be set through the dual variables of the capacity constraints of the above problem, but this requires complete information on users' values of time and O-D pairs

Online Learning Algorithm

Privacy-Preserving Tolling Algorithm: Only relies on observed edge flows and not on any user attributes

Algorithm 1: Efficient Routing via Privacy-Preserving Tolls

```
Input   : Time Period T, Road Capacities c
Set the Toll τ^(1) ← 0
for t = 1, ..., T do
    Phase I: User Equilibrium for Toll τ^(t)
    Observe Edge Flows x ;
    Phase II: Toll Update
    τ^(t+1) ← (τ^(t) − γ(c − x^t))_+
end
```

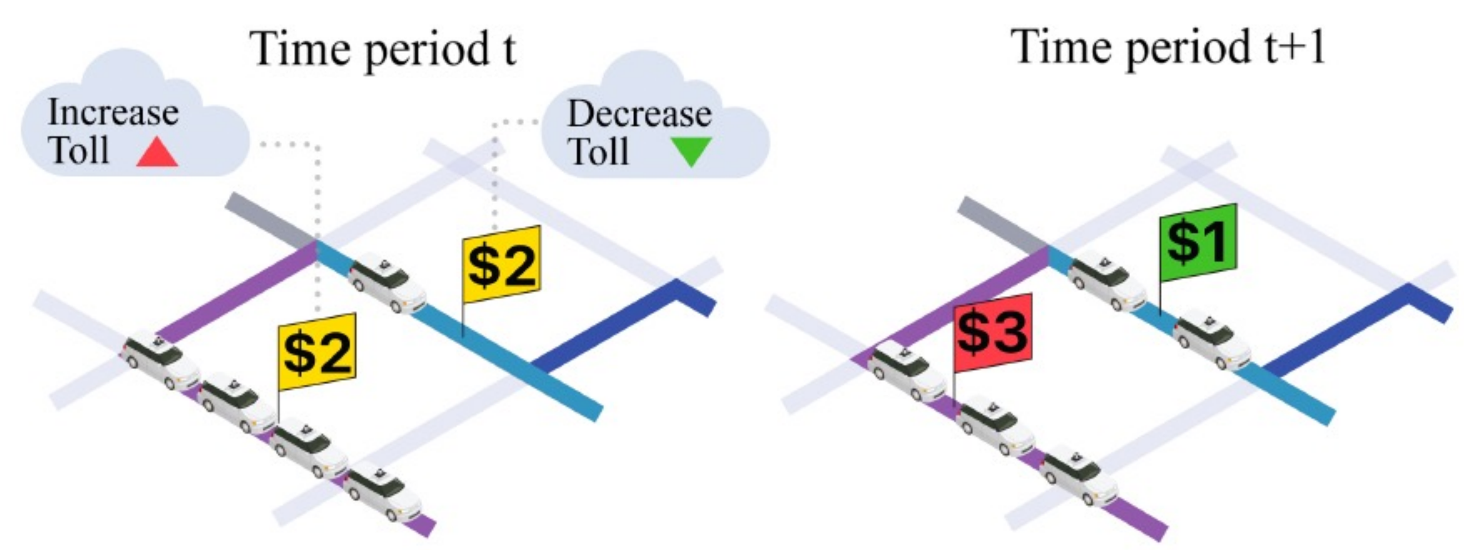

Between two time periods, the toll on each edge is increased (decreased) if there is more (less) flow than the capacity of the edge

Performance Metrics

Regret

Optimality gap between an online allocation and an offline oracle with complete information on users' preferences

Capacity Violation

Norm of cumulative excess flow beyond each edge's capacity

Theorem (Lower Bound): ***Under i.i.d. user arrival the regret of any algorithm is $\Omega(\sqrt{T})$, where T is the number of time periods***

Theoretical Upper Bound

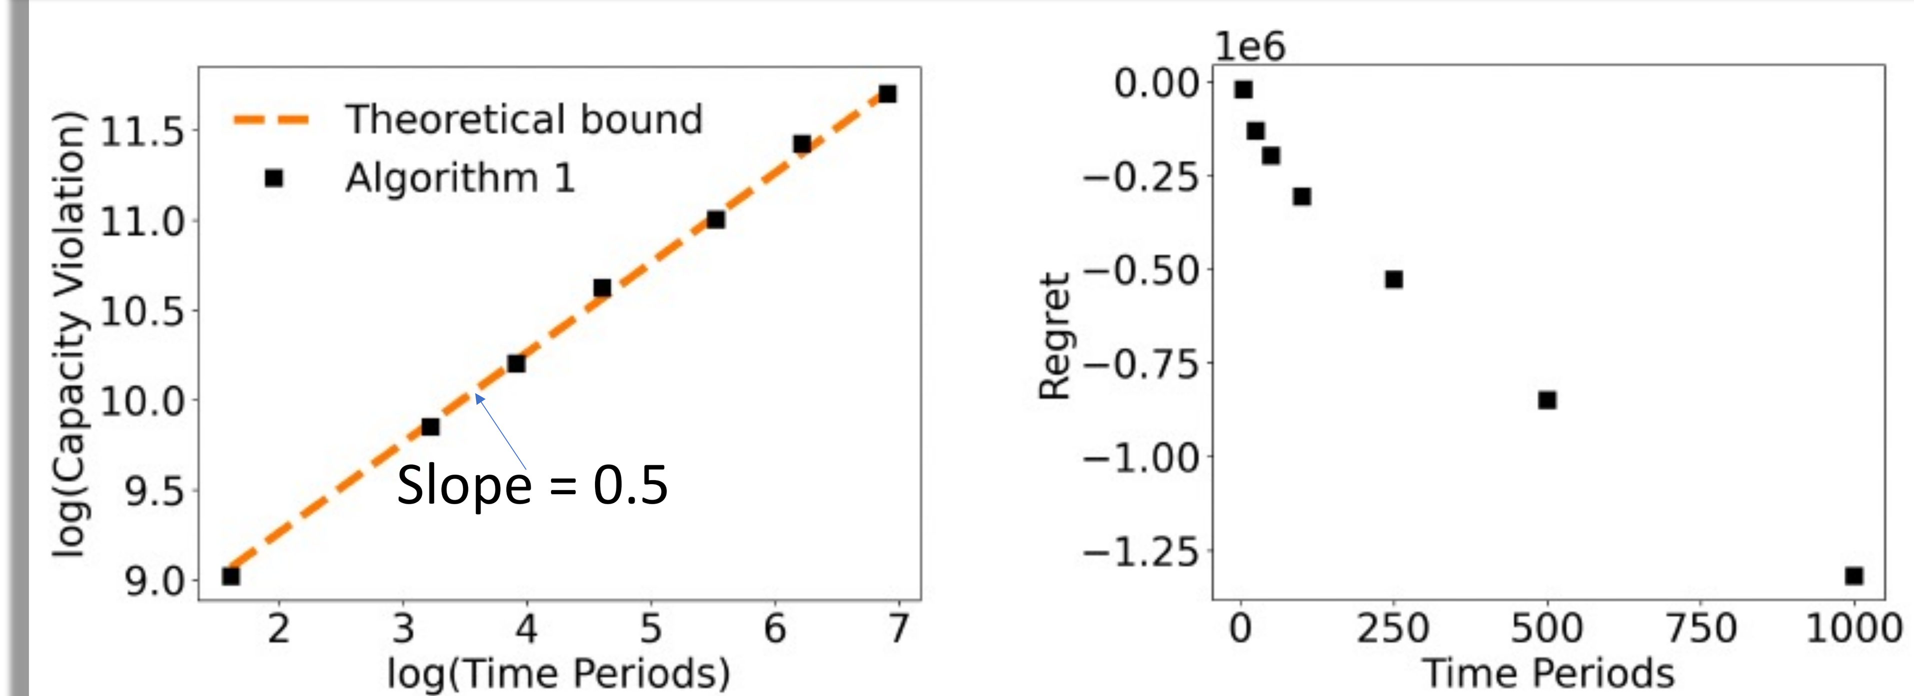

Theorem (Upper Bound):

Under i.i.d. user arrival, the regret and constraint violation of Algorithm 1 are $O(\sqrt{T})$, where T is the number of time periods

Comparison to Benchmarks

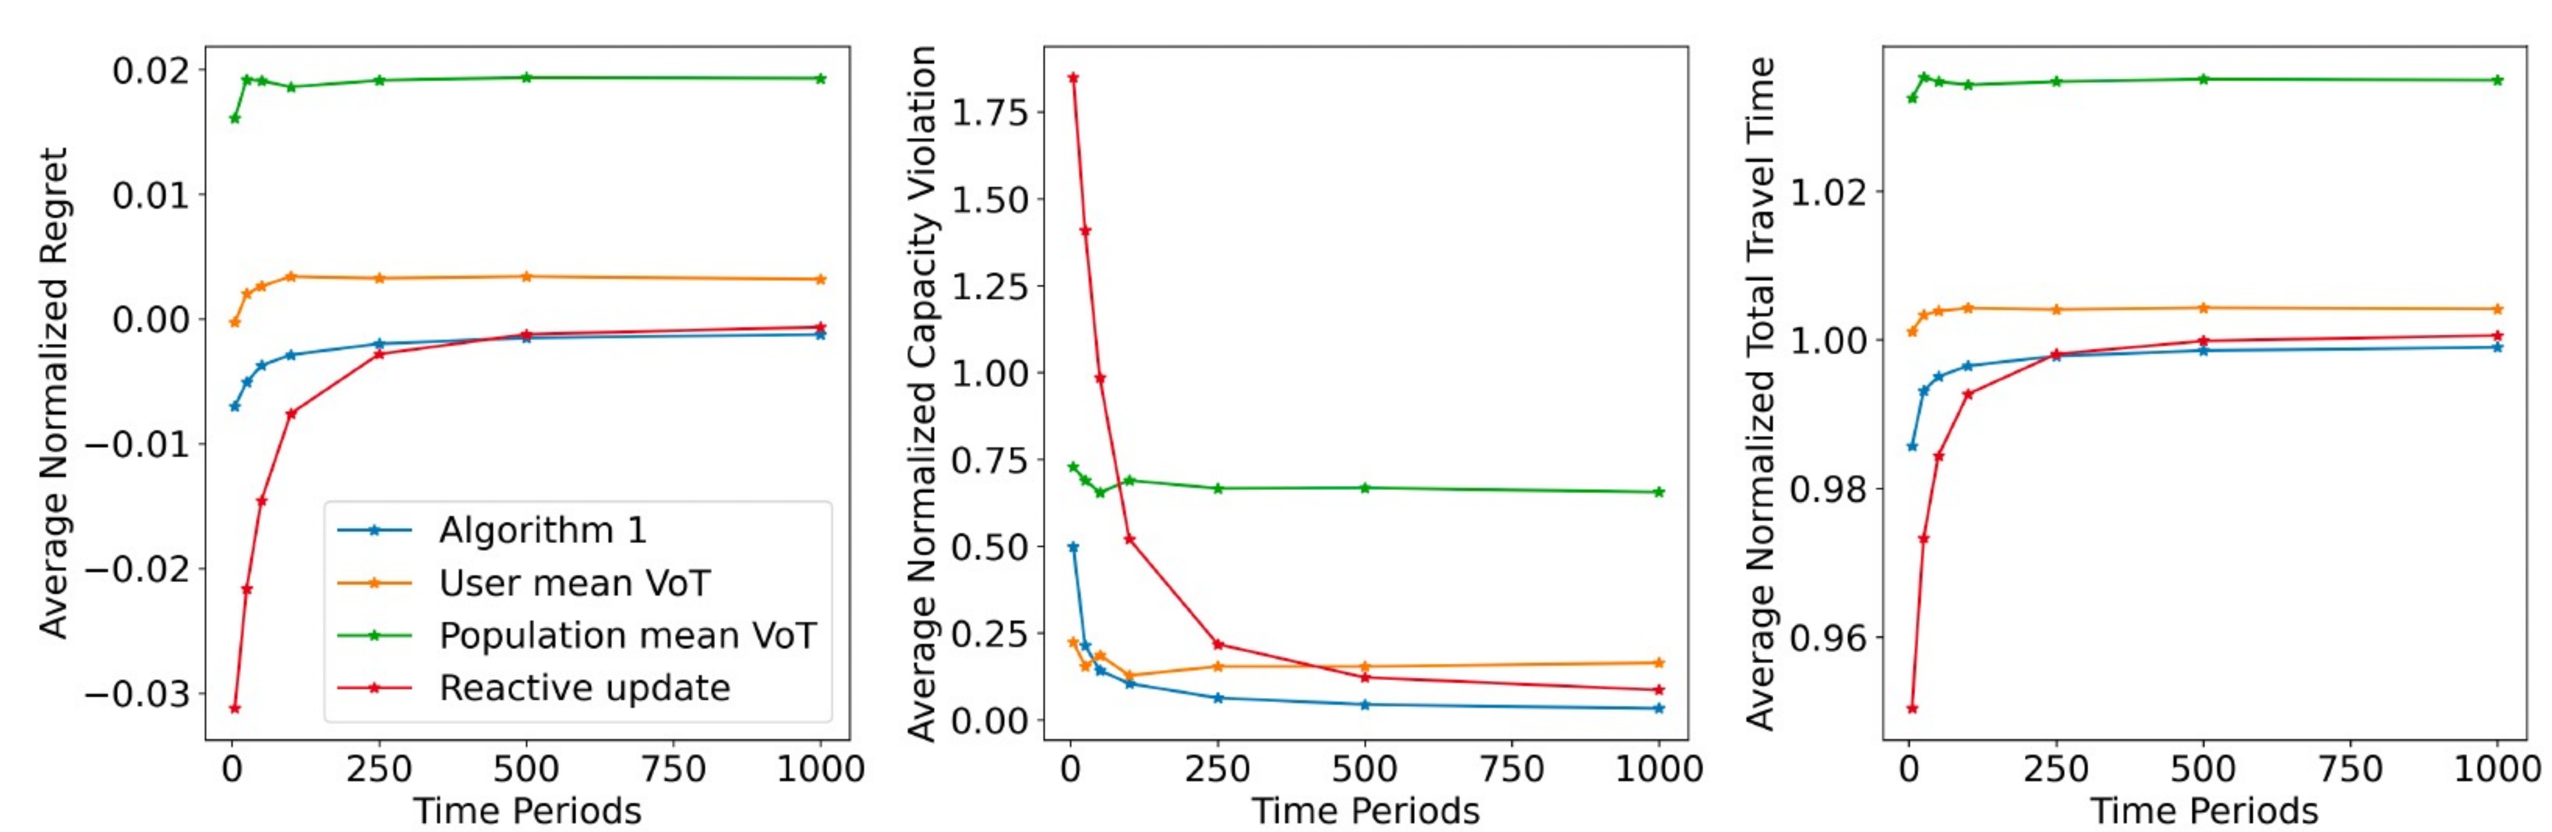

Algorithm 1 performs better than two static tolling benchmarks that have information on the mean values of time of users and the entire population. It also performs better than a reactive dynamic tolling benchmark on both regret and capacity violation metrics

Future Work

1. Developing regret and constraint violation guarantees when users' trip attributes are not i.i.d.
2. Extending results to setting when travel times vary as a function of the flows
3. Extending to objectives beyond system efficiency, e.g., fairness or revenue