



PiTree: Practical Implementations of ABR Algorithms Using Decision Trees



Zili Meng* Jing Chen* Yaning Guo* Chen Sun* Hongxin Hu† Mingwei Xu*
*Tsinghua University †Clemson University

Session 5C 14:15 Thu Oct 24 Rhodes 9 zilim@ieee.org transys.io/pitree

Motivation

ABR algorithms are increasingly heavyweight

Server-side Implementation


- High operating expenses.
- Up to millions of concurrent viewers.

Client-side Implementation

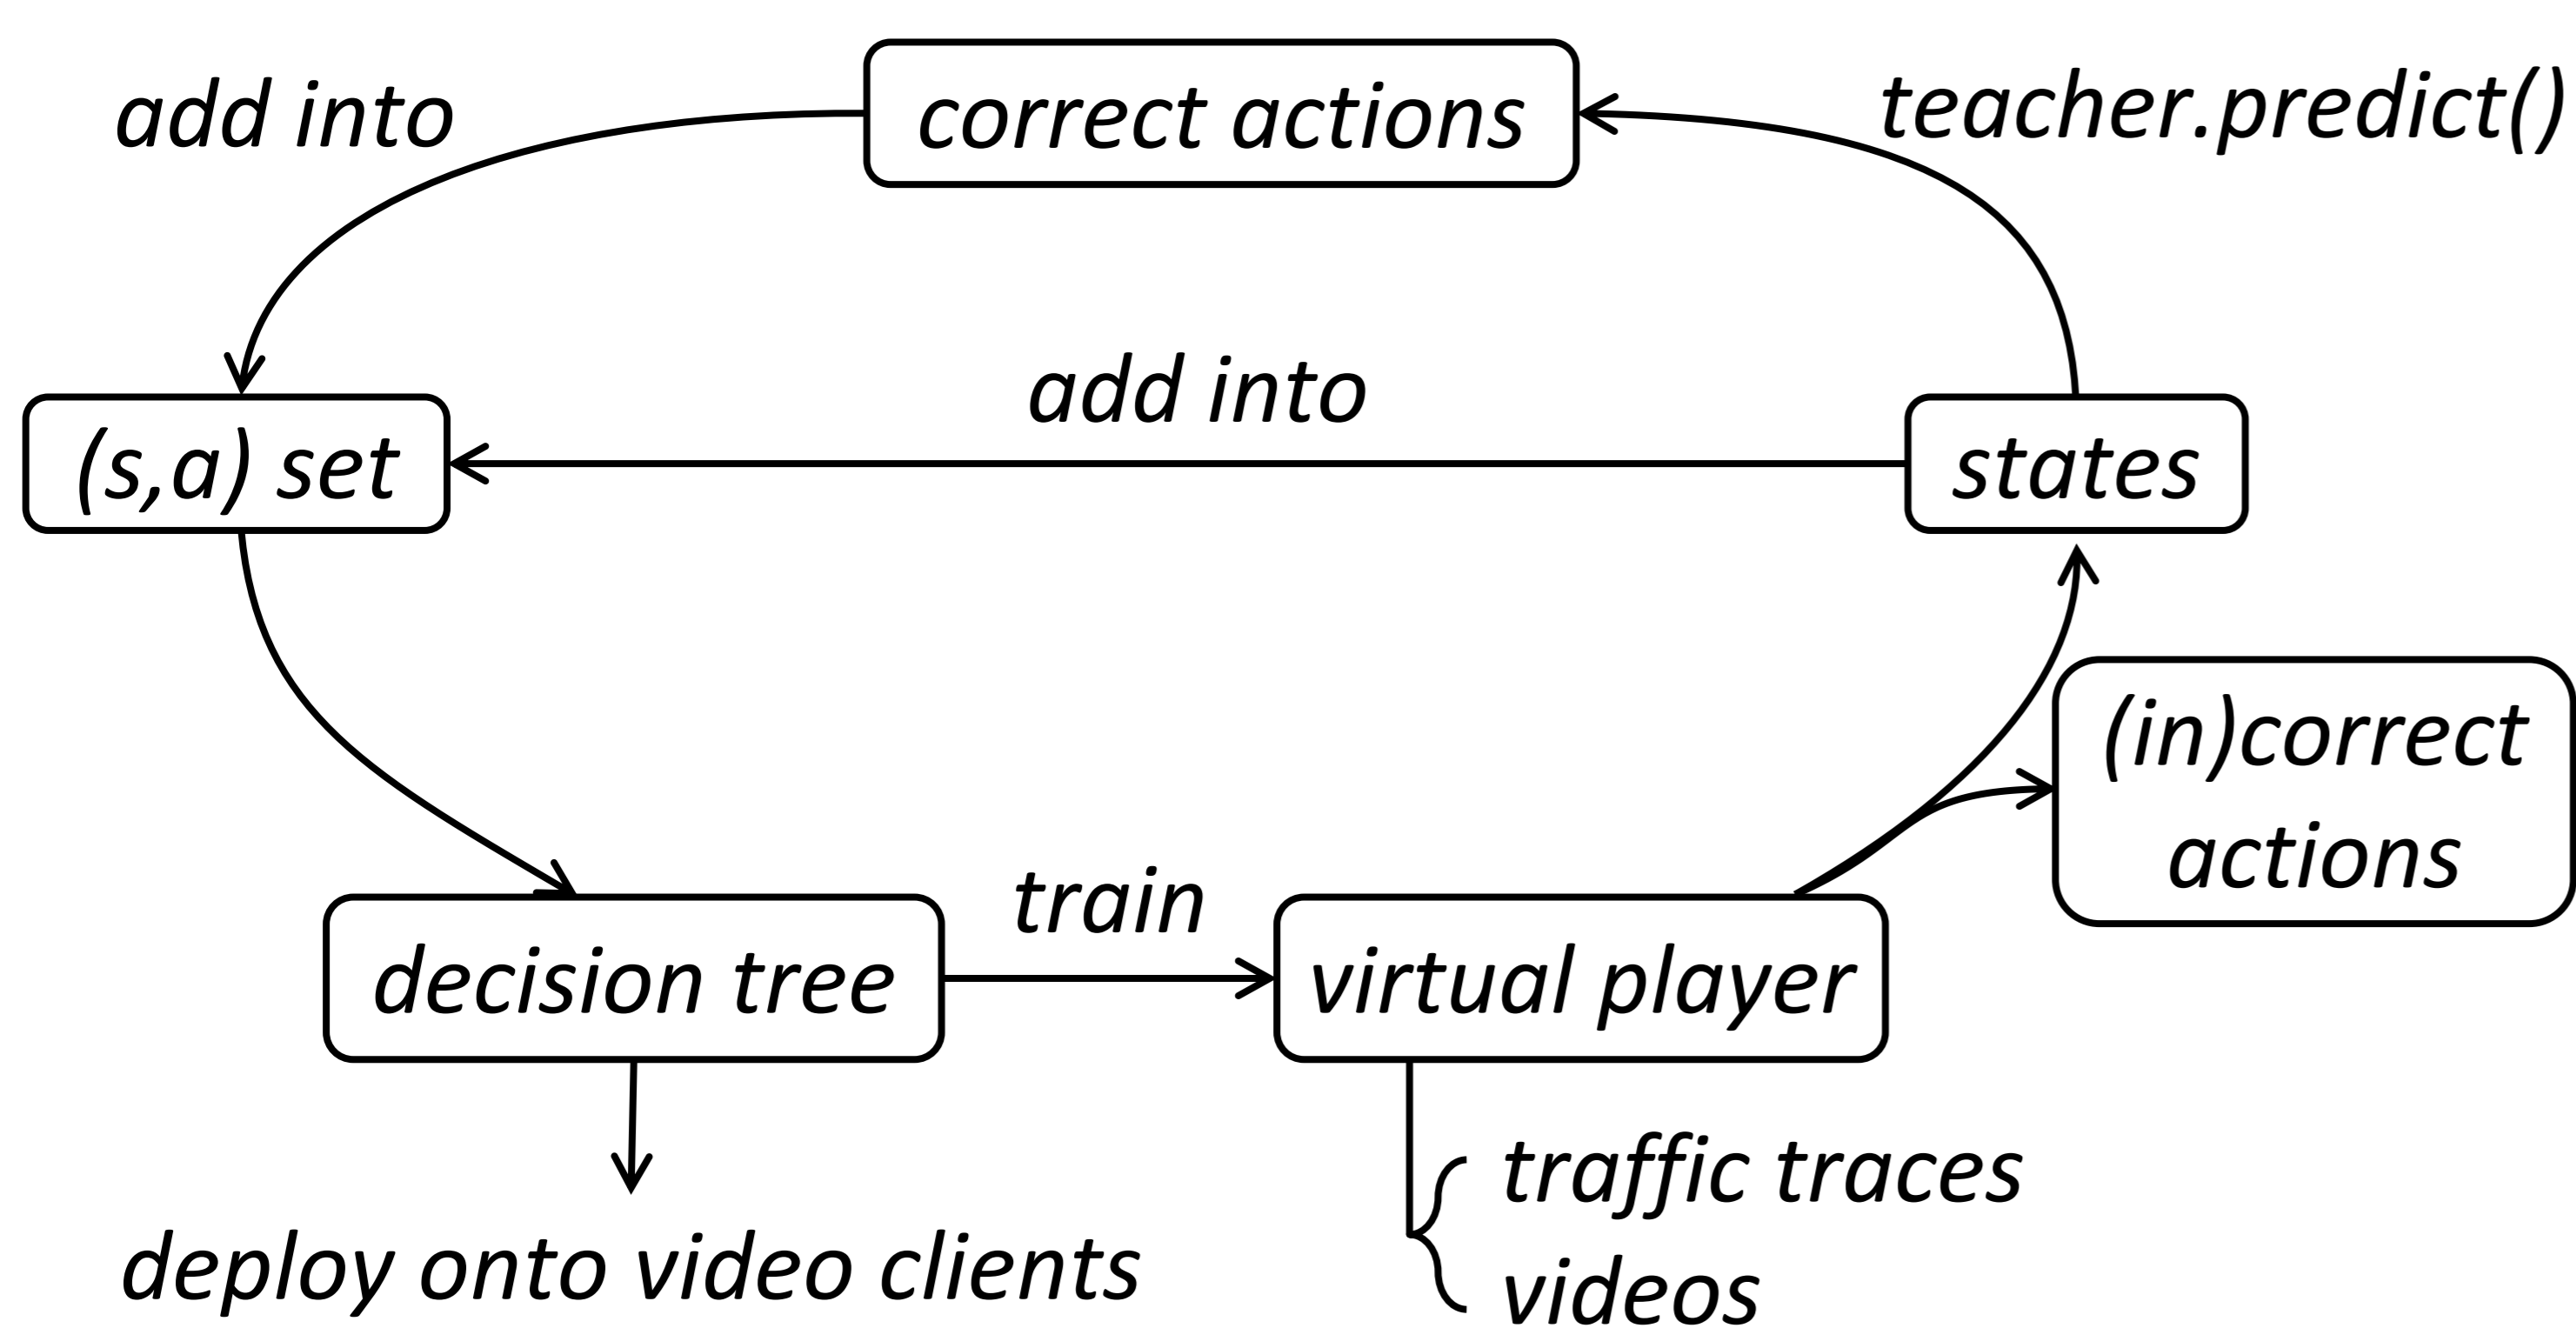
- Large page size.
- Page load time of Pensieve is increased by ~10s.
- Long decision latency.
- Decision latency of RobustMPC > chunk length.

Challenges

ABR Control is a sequential decision-making process. One wrong prediction may drive the student off the teacher's trajectory.



Design



Algorithm pseudocodes:

```

(S, A) ← VirtualPlay(π*)
For i from 1 to M:
  π_i ← TrainDT(S, A)
  (S_i, A_i) ← VirtualPlay(π_i)
  A_i* ← Predict(π*, S_i)
  Aggregate S ← S ∪ S_i, A ← A ∪ A_i*

```

Loss in decision tree training: $\ell(r; r_0) = \frac{(r-r_0)^2}{(R_{max}-R_{min})^2}$


Theoretical bound

For any $\delta > 0$, with training loss ϵ_M , there exists a policy $\hat{\pi} \in \{\pi_1, \dots, \pi_M\}$ s.t. the average optimization loss satisfies:


$$\mathbb{E}_{s \sim d_{\hat{\pi}}}[\ell(\hat{\pi}(s); \pi^*(s))] \leq \epsilon_M + \Theta(1/T)$$

Solution

Offline training/design






Online deployment



Evaluation

RobustMPC [SIGCOMM'15]	HSDPA	Lin
ABR: Pensieve [SIGCOMM'17]	Traces: FCC	QoE: Log
HotDASH [ICNP'18]	Oboe	HD



Our experiments: HotDASH: PC: Intel Core i7-8550 Mobile: Qualcomm Snapdragon 821
<0.1s @ 1Mbps 14.5s @ 1Mbps

Coming soon

Explaining Complex Networked Systems

Decision trees are not only lightweight, but also explainable. Partial results are online at [arXiv:1910.03835](https://arxiv.org/abs/1910.03835).

