

Summary of Neural Adaptive Video Streaming with Pensieve

Zhang Shiwei | May 2018

In *Neural Adaptive Video Streaming with Pensieve*[1], Mao et al. presented a novel system called *Pensieve* as an alternative to current adaptive bitrate (ABR) algorithms in video streaming. The authors also made extensive experiments in both simulated and real environments to show the effectiveness and robustness of it. As declared in the paper, *Pensieve* improves the quality of experience (QoE) for 12% - 25% than current state-of-the-art algorithms, and generalize well to various types of networks.

The key idea of *Pensieve* is employing a reinforcement learning (RL) algorithm to train a neural network model as a replacement of current hand-crafted ABR models. RL considers an agent interacts with an environment, where at each time step the agent chooses an action and get a reward. By setting network observations and past actions like bandwidth, bitrate, and buffer size as the environment, bitrate of the next chunk as the action, and QoE as the reward, existing RL algorithms can be used to learn a model that select bitrates automatically to maximize QoE under certain observations.

According to the paper, the primary challenges of ABR algorithms are

- the variability of network throughput;
- the conflicting QoE requirements (high bitrate, minimal rebuffering, smoothness, etc.);
- the cascading effects of bitrate decisions (e.g., selecting a high bitrate may drain the playback buffer to a dangerous level and cause rebuffering in the future);
- and the coarse-grained nature of ABR decisions;

Current algorithms, as stated in the paper, cannot address these problems well because they rely on fixed heuristics or simplified (and thus inaccurate) system models. By contrast, RL-generated algorithms learn from actual performance resulting from different decisions. By incorporating this information into a flexible neural network policy, RL-generated ABR algorithms can automatically optimize for different network characteristics and QoE objectives.

The architecture of the neural network is basically a 3-layer MLP, with an additional 1-D convolution layer for sequential data like past throughput or next chunk sizes. The authors came up with this architecture by experiments. However, for sequential data, recurrent neural networks are often considered performing better than convolutional neural networks[2], which is used in this paper.

I was skeptical about the work because the authors train the model in a simulated environment, while claim it performs better than other models that may be based on the statistics of even larger corpus. However, the paper gives a lot of experiment details including some in the real environment. The only explanation is that *Pensieve* did design a better generic ABR algorithm from the network trace data than the previous experts. We can try this method in similar tasks like TCP congestion control algorithms.

References

- [1] Hongzi Mao, Ravi Netravali, and Mohammad Alizadeh. Neural Adaptive Video Streaming with Pensieve. pages 197–210, Los Angeles, CA, USA, 2017. ACM Press.
- [2] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, May 2015.