Hierarchical Multiple-Instance Data Classification with Costly Features

Jaromír Janisch • **Tomáš Pevný** • **Viliam Lisý**

Artificial Intelligence Center / Department of Computer Science / FEE / Czech Technical University in Prague

Abstract

TL;DR: We apply Deep RL to sequentially select features in tree-like samples to optimize classification accuracy vs. feature cost.

We extend the framework of Classification with Costly Features (CwCF) that works with samples of fixed dimensions to trees of varying depth and breadth (similar to a JSON/XML file). In this setting, the sample is a tree – sets of sets of features. Individually for each sample, the task is to sequentially select informative features that help the classification. Each feature has a real-valued cost, and the objective is to maximize accuracy while minimizing the total cost. The process is modeled as an MDP where the states represent the acquired features, and the actions select unknown features. We present a specialized neural network architecture trained through deep reinforcement learning that naturally fits the data and directly selects features in the tree. We demonstrate our method in seven datasets and compare it to two baselines.

Classification with Costly Features

$$\min_\theta \underbrace{\mathbb{E}\left[\ell(y_\theta(x), y)\right]}_{\text{classification loss}} \quad \text{s.t.} \quad \underbrace{\mathbb{E}\left[z_\theta(x)\right]}_{\text{cost of features}} \leq budget$$

$$\downarrow$$

$$\min_\theta \mathbb{E}\left[\ell(y_\theta(x), y) + \lambda z_\theta(x)\right]$$

Solved as an MDP with *feature-selecting actions* $\mathcal{A}_f$ and a terminal *classification action* $\mathcal{A}_c$. The reward for A_f is $-\lambda z_\theta(x)$ and for A_c it is $-\ell(y_\theta(x), y)$, so the total reward for an episode yields the exact quantity to optimize the main objective.

Method

Object embedding with HMIL: An object processed with Hierarchical Multiple-Instance Learning to create its embedding and internal objects.

Action selection with hierarchical softmax: Starting at the root of the object, stochastic decisions are made to follow the structure to a leaf - a feature. The actions probability is a product of the partial decisions.

Deep RL training: The reward is given as in CwCF. The action probabilities are updated with a policy gradient method (A2C).

Classification of Potentially Malicious Website

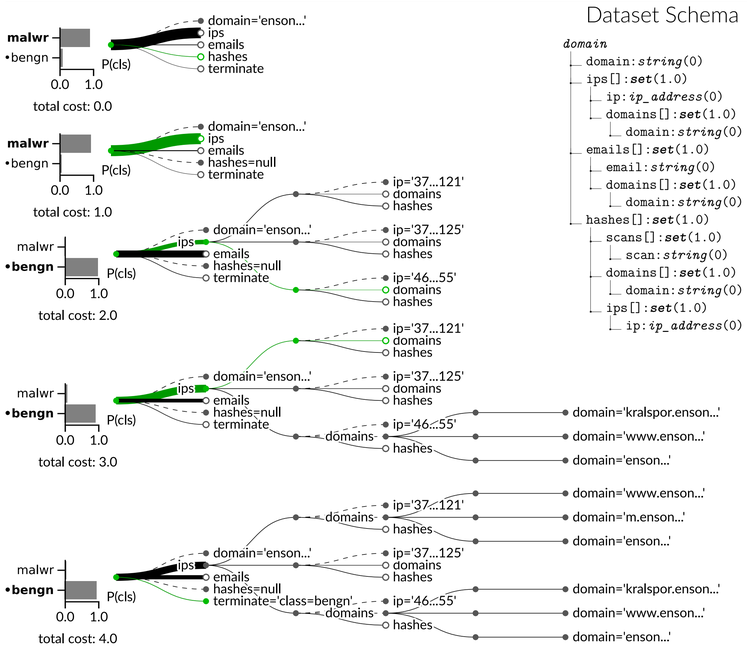

Samples' structure follows the common schema, but the number of objects in sets and their depth vary!

Model Architecture

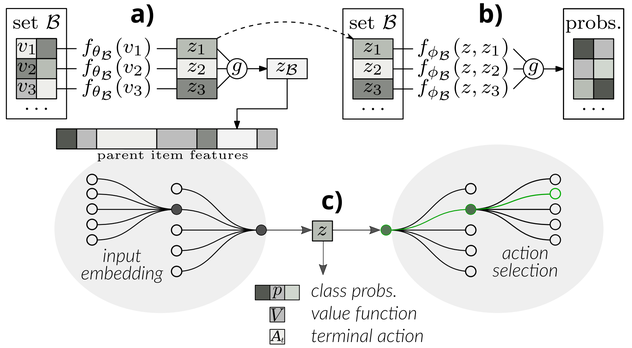

Experimental Results

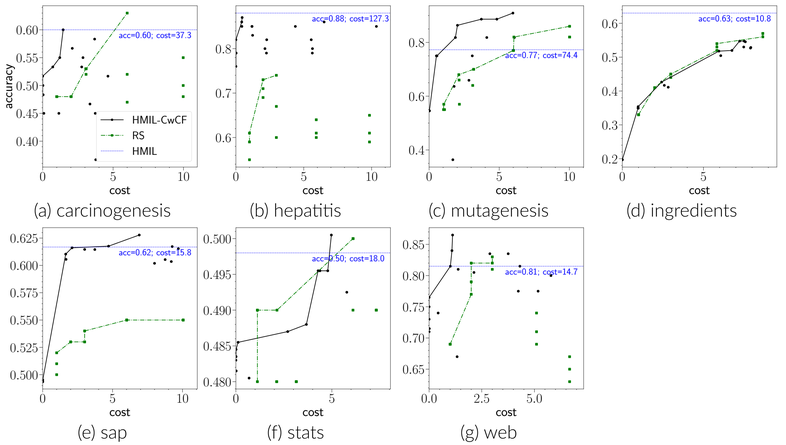

(a) carcinogenesis (b) hepatitis (c) mutagenesis (d) ingredients

(e) sap (f) stats (g) web

Age Classification of StackExchange User

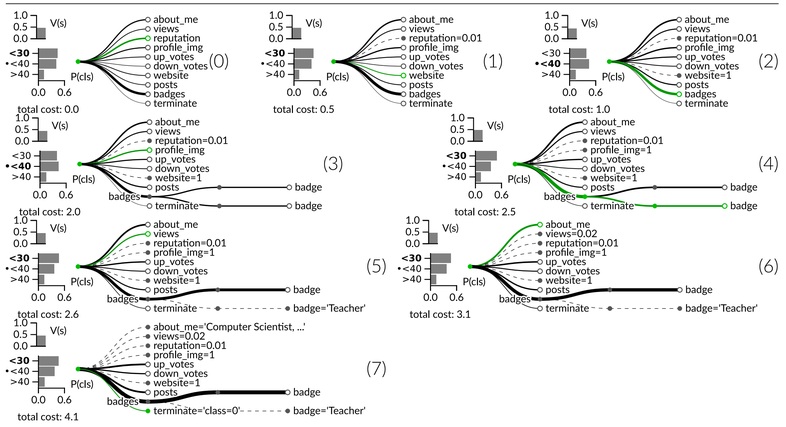

Acknowledgments: This research was supported by the European Office of Aerospace Research and Development (grant no. FA9550-18-1-7008) and by The Czech Science Foundation (grants no. 18-21409S and 18-27483Y). The GPU used for this research was donated by the NVIDIA Corporation. Computational resources were supplied by the project "e-Infrastruktura CZ" (e-INFRA LM2018140) provided within the program Projects of Large Research, Development and Innovations Infrastructures. The authors acknowledge the support of the OP VVV funded project CZ.02.1.01/0.0/0.0/16_019/0000765 "Research Center for Informatics".