On Neural Networks and Decision Support

Optimising Human and Artificial Decision Making

Author: Sam Verboven
Promotors: prof. Vincent Ginis, Ph.D (VUB) and prof. Wouter Verbeke, PhD (KU Leuven)

Thesis outline :

The objective of this thesis was to advance the state of the art of deep learning techniques and applications in decision support literature. To this end, we introduced a novel framework for strategic decision support based on unsupervised learning, and three new adaptive multi-task learning algorithms.

In chapter 1, we describe the broader streams in the literature relevant to this work. Mainly, we highlight how; (i) algorithmically generated representations can be informative to human decision making (CH2), (ii) multi-task learning can improve performance and algorithms and hence further aid decision making and (CH3,4,5), (iii) accurate modeling of causal effects is imperative to make the right decisions (CH5).

In chapter 2, we first review the literature on strategic decision support and human decision making. Consequently, a first experiment highlights the inconsistency of human decision making in strategic decisions. Next, we propose unsupervised algorithmic decision support as a solution. Specifically, we propose the Autoencoder as a generator of strategically informative representations. We present results on a three-point experimental setup to validate the performance of the Autoencoder. Unsupervised learning and in particular the Autoencoder are shown to be valuable tools to support strategic decision-making.

In the following three chapters, our work focuses on the idea of using look-ahead main-task gain to inform multi-task loss weighting. Based on this general idea, we formulate three novel algorithms, each of which provide a solution in a different domain.

In chapter 3, we introduce HydaLearn, an intelligent multi-task weighting algorithm that uses main-task gain to determine dynamic loss weights at the mini-batch level. To calculate main-task gain, we introduce the concept of ‘look-ahead updates’. We demonstrate the performance increase of HydaLearn compared to state-of-the-art baselines for two two-task supervised learning settings on synthetic data and two real datasets for default prediction and mortality prediction, respectively.

In chapter 4, we propose a first specialized extension to HydaLearn that generalizes to any number of auxiliary tasks and is robust to label noise. This novel algorithm calculates and performs parameter updates based on each task separately and sequentially. Furthermore, a novel ‘look-ahead’ calculation scheme is used to guarantee that obtained ‘look-aheads’ are always representative of the final update. When training with very small mini-batches, the algorithm effectively ignores task-specific noise and capitalizes on noiseless examples. Using synthetic regression and classification data, and semi-synthetic data on an inverse dynamics regression problem, we investigate the behaviour of our algorithm and confirm its robustness to label noise.

In chapter 5, we propose a third new algorithm, extending the idea of using ‘look-aheads’ to inform task weighting to uplift modeling. Previous algorithms focused either on ranking or on accurate point-estimation of the uplift. Our novel uplift algorithm combines the best of both worlds; use of ranking information during training, and point estimate outputs. In our experiments, we show appropriate weight initialisation is highly important for MTL uplift models. To this end, we use a simple meta-loop to find good initialisation parameters. Furthermore, the superior performance of our algorithm is empirically validated on precision marketing and student dropout minimization problems.

In chapter 6, we conclude with a summary of the main contributions, critical reflections and a discussion of future avenues. From a technical perspective, three new highly effective multi-task learning algorithms were introduced using the novel concept of look-ahead metric gain to inform task weighting. From a practical perspective, we have pushed the state-of-the-art in several existing high-impact decision support settings and introduced a novel approach to strategic decision support. Furthermore, the proposed algorithms require less tuning compared to conventional multi-task learning algorithms, thus decreasing the burden of finding optimal hyperparameters on the practitioners.