Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence

Online and Stochastic Learning with a Human Cognitive Bias

Hidekazu Oiwa
The University of Tokyo and JSPS Research Fellow
7-3-1 Hongo, Bunkyo-ku, Tokyo, Japan
hidekazu.oiwa@gmail.com

Hiroshi Nakagawa
The University of Tokyo
7-3-1 Hongo, Bunkyo-ku, Tokyo, Japan
nakagawa@dl.itc.u-tokyo.ac.jp

Abstract

Sequential learning for classification tasks is an effective tool in the machine learning community. In sequential learning settings, algorithms sometimes make incorrect predictions on data that were correctly classified in the past. This paper explicitly deals with such inconsistent prediction behavior. Our main contributions are 1) to experimentally show its effect for user utilities as a human cognitive bias, 2) to formalize a new framework by internalizing this bias into the optimization problem, 3) to develop new algorithms without memorization of the past prediction history, and 4) to show some theoretical guarantees of our derived algorithm for both online and stochastic learning settings. Our experimental results show the superiority of the derived algorithm for problems involving human cognition.

1 Introduction

Online learning and stochastic learning are advantageous for large-scale learning. Sequential processing of data is the key of these methods. For classification tasks, these learning algorithms process a bunch of data one by one and change its classification rule at every round. We call these methods sequential learning in this paper.

Sequential learning algorithms sometimes make wrong predictions on data that were correctly classified in the past. While classical performance evaluation measures for sequential learning, such as the expected loss, do not reflect the history of the past prediction results, previous algorithms have not considered this inconsistent behavior as a crucial factor. The key statement in this paper is that this phenomenon has a crucial impact on the evaluation of algorithms on the condition that humans are evaluators. Humans have a cognitive bias that they attach a higher value to the data that were correctly classified in the past than the other data. This effect originates from the endowment effect that had been widely analyzed in the field of behavior economics. There are motivating examples in which this cognitive bias has important roles:

- **User utility maximization**: Sequential learning has been used in many services such as image object recognition and email filtering (Aberdeen, Pacovsky, and Slater 2010). Many users continuously utilize services whose prediction rules have been changed over time. Furthermore, some users check prediction results of previously seen data. Negative flips may drastically decrease the utilities of these users.
- **Interactive annotation**: There are many human-computer interaction systems based on sequential learning such as active learning-based annotations (Settles 2011). Encouraging people to make annotations is crucial for more data generation and better performance. Some annotators may feel frustrated annotating the data correctly classified in the past as wrong ones.

To maximize the availability of machine learning, algorithms which interact with humans need to adjust update rules to heal the bias derived from the past prediction history. We explicitly deal with this cost as the divestiture loss. We first conducted an experiment to verify whether the endowment effect negatively affects human's evaluations. Next, we set new evaluation measures for sequential learning by incorporating the endowment loss. This measure imposes an additional objective on sequential learning, minimizing the divestiture loss. We note that this new problem setting can be easily dealt with if algorithms could store all previous examples and its prediction results in the memory; however, this memorization is unpractical for large-scale learning setting due to the memory constraint. To solve this problem, we derived new variants of Online Gradient Descent (OGD). Our derived algorithms enable to retain reasonable convergence guarantees for both online learning and stochastic learning settings without data memorization. We lastly conducted experiments and the results showed advantages of our algorithm compared with the conventional ones in the sequential learning framework with a human cognitive bias.

1.1 Notations

Scalars are denoted by lower-case x and vectors are denoted by bold lower-case $\mathbf{x}$. t-th training input vectors and labels are denoted $\mathbf{x}_t$ and y_t. Input vectors are n-dimensional and taken from the input space $\mathcal{X} \subset \mathbb{R}^n$. Output labels are taken from the output space $\mathcal{Y}$. For simplicity, we define $z_t = (\mathbf{x}_t, y_t)$ to describe t-th datum. $\mathbf{x}_{s:t}$ describes a sequence of vectors from s-th to t-th and $\mathbf{x}_{1:0}$ is a empty set. $1_{a=b}$ is a boolean function which becomes 1 only if $a = b$.