

A Framework for Long-Term Learning Systems

Diana Benavides-Prado
 Department of Computer Science
 The University of Auckland
 dben652@aucklanduni.ac.nz

Abstract

Increasing amounts of data have made the use of machine learning techniques much more widespread. A lot of research in machine learning has been dedicated to the design and application of effective and efficient algorithms to explain or predict facts. The development of intelligent machines that can learn over extended periods of time, and that improve their abilities as they execute more tasks, is still a pending contribution from computer science to the world. This weakness has been recognised for some decades, and an interest to solve it seems to be increasing, as demonstrated by recent leading work and broader discussions at main events in the field [Chen *et al.*, 2016]. Our research is intended to help fill that gap.

1 Introduction

Metalearning, which embraces the ability of learning systems to learn based on learning experiences, has been one of the fields concerned about the problem of long-term learning systems. A very common form of metalearning can be defined as the ability to transfer information or knowledge from previous tasks. This is the main problem studied in Transfer Learning (TL) and domain adaptation [Pan and Yang, 2010]. The most common setting in TL is to transfer source data, especially when the data for a new task is scarce or difficult to access. Other kinds of information that can be transferred are parameters or feature representations, between homogeneous or heterogeneous domains.

The main difficulty with regular TL is that it still relies on data used for learning previous models. There are scenarios where data is not available for transfer, but instead we can have access to hypotheses learned for related tasks. This has been named as the Hypothesis Transfer Learning (HTL) problem [Kuzborskij and Orabona, 2013], and has gained more interest in recent years. In HTL, it is possible to transfer information directly from previous models or hypotheses. A latent problem is how a long-term learning system can be benefited from transferring along multiple and consecutively learned hypotheses. HTL can act as an enabler for the construction of these kinds of systems, given its reliance in previously learned knowledge rather than in previous data.

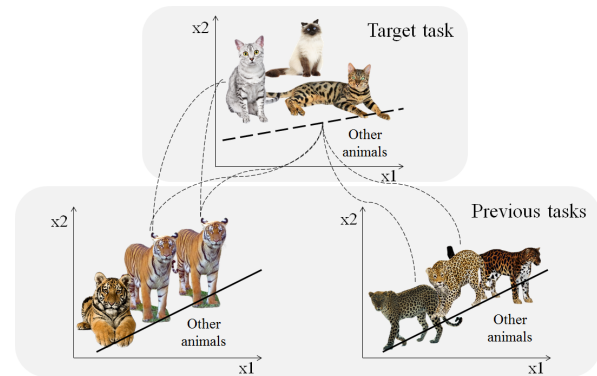


Figure 1: An example of hypothesis transfer learning.

Most research in HTL has concentrated on transferring hypotheses as a whole. The learning of a new function for a target task is biased towards a linear combination of previous hypotheses. These are generally used according to how well they predict the target data. The more accurate the prediction, the more appropriated a hypothesis is for the new task.

Consider the example in Fig. 1. Suppose that we have some training data on animals, including cats, and we have a new learning task to find a function that classifies an animal onto cat or other. We also have access to previous hypotheses. These hypotheses relate characteristics of other Felidae and animals with their corresponding classes (in the example, tiger or leopard, or other animal). We can then use what has been already learned from other Felidae to learn faster or better on our new cats, given that they belong to the same family. This can be understood as the main aim of HTL.

In our cats classification problem, information on some subfamilies of tigers (bengal, for example) might be more useful for learning about some of our new cats (for example, egypt and leopard subfamilies of cats), and some leopards (indian leopard) could better help to learn about other subfamilies of cats (leopard cat). This means that, in some cases, only fragments of source hypotheses are really useful for a new learning task. One of the first opportunities that we have explored concerns the ability to identify and transfer from previous hypotheses in a more selective manner than regular HTL, by using whatever is necessary from source hypotheses, while still learning from the target data.

2 Contributions

Our selective HTL method answers the common questions of what, when and how to transfer, in new ways. For the *what to transfer* question, we propose to transfer coefficients from hypotheses learned with an SVM. These coefficients are potentially useful for our new task, as explained later. We propose to base the decision of *when to transfer* on a two-step evaluation of the similarity between previous hypotheses and their elements, and target data: first, similarity of their distributions as a whole, using a metric like KL-divergence and, second, similarity of source support vectors with target data points. Regarding *how to transfer*, we propose to treat previously learned coefficients as privileged information that is available for a new task at training time [Lapin *et al.*, 2014]. This information serve as upper-bounds on coefficients to be learned on the new task, to reinforce target data points that are similar to previously learned ones.

To implement this simple yet effective selective HTL method, we only need to update a constraint on the dual optimization function of an SVM learning problem:

$$\begin{aligned} & \max_{\alpha} \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i,j=1}^n \alpha_i \alpha_j y_i y_j K(x_i, x_j) \\ \text{s.t.} & \sum_{i=1}^n y_i \alpha_i = 0, \forall i \ 0 \leq \alpha_i \leq C + c_i, c_i = \frac{s_i}{|\overline{H}|} \sum_{k=1}^{sv} \overline{\alpha}_k \end{aligned} \quad (1)$$

For a specific coefficient α_i that must be learned during a new task, a modified constraint, c_i , is learned from previous hypotheses, as represented by coefficients $\overline{\alpha}$ from the corresponding subset of similar source support vectors of size sv . The new constraint can be learned from source hypotheses, and allows some target data points to contribute more to the maximization problem, thus implying their importance. s_i corresponds to the number of previous hypotheses contributing to c_i , as found by our method, from the set of available source hypotheses of size $|\overline{H}|$. This method has shown faster convergence rates for classification tasks, with a decrease on the number of iterations by up to 56% on average compared to learning with no transfer and up to 92% compared to regular HTL methods, while maintaining similar accuracy levels.

3 Future Work

Apart from selective HTL, there are additional challenges that remain to be solved for long-term learning systems that operate on the basis of transferring information between hypotheses. One of them is how to deal with the non-stationary nature of data and tasks within a learning system. This is similar to the concept drift problem, a common issue that can arise for systems or models that operate over time. Several research works have tried to provide solutions for learning in the presence of drifting concepts, from drift detection to model adaptation. It is also a recognised problem in transfer learning research [Wang *et al.*, 2014]. HTL for drifting concepts is a relevant research avenue for long-term learning systems that operate on the basis of hypothesis transfer, since these transfer mechanisms should be able to detect, react and adapt to new concepts that arise or old concepts that evolve.

An additional research avenue in long-term learning systems based on HTL is how to transfer among changing feature representations. In regular TL, this problem has been extensively treated under the umbrella of transfer between heterogeneous domains [Pan and Yang, 2010]. However, the general assumption in HTL has been that both source hypotheses and target data have the same feature representation.

Last but not least, our framework should be able to provide answers to common questions that arise for what has been named as the Lifelong Learning problem [Chen *et al.*, 2016]. High level Lifelong Learning frameworks have been proposed [Silver *et al.*, 2013], along with recent works (eg. [Ruvolo and Eaton, 2013; Fei *et al.*, 2016]) that denote an increasing interest. Challenges in this area are related, for example, to learning of concepts in the long-term or to evaluation metrics for Lifelong Learning. To the best of our knowledge, this problem has not been explored from the point of view of hypotheses that can work collaboratively over time by selectively transferring information between learning tasks.

The development of a learning machine that retains models or hypotheses over time and that use them effectively to improve the learning ability of an overall system, remains as an opportunity. Novel methods to transfer from previous hypotheses, that provide answers to the main challenges, and that can be embedded onto a framework as a basis for the development of similar systems, are still to be explored. To the best of our knowledge, a framework for long-term learning systems based on selective transfer of information would be the first of this kind.

References

- [Chen *et al.*, 2016] Zhiyuan Chen, Estevam Hruschka, and Bing Liu. Lifelong machine learning and computer reading the web. SIGKDD Tutorial, 2016.
- [Fei *et al.*, 2016] Geli Fei, Shuai Wang, and Bing Liu. Learning cumulatively to become more knowledgeable. In *SIGKDD*, pages 1565–1574. ACM, 2016.
- [Kuzborskij and Orabona, 2013] Ilja Kuzborskij and Francesco Orabona. Stability and hypothesis transfer learning. In *ICML*, pages 942–950, 2013.
- [Lapin *et al.*, 2014] Maksim Lapin, Matthias Hein, and Bernt Schiele. Learning using privileged information: SVM+ and weighted SVM. *Neural Networks*, 53:95–108, 2014.
- [Pan and Yang, 2010] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE TKDE*, 22(10):1345–1359, 2010.
- [Ruvolo and Eaton, 2013] Paul Ruvolo and Eric Eaton. Ella: An efficient lifelong learning algorithm. *ICML*, 28:507–515, 2013.
- [Silver *et al.*, 2013] Daniel L Silver, Qiang Yang, and Lianghao Li. Lifelong machine learning systems: Beyond learning algorithms. In *AAAI Spring Symposium: Lifelong Machine Learning*, volume 13, pages 49–53. Citeseer, 2013.
- [Wang *et al.*, 2014] Xuezhi Wang, Tzu-Kuo Huang, and Jeff G Schneider. Active transfer learning under model shift. In *ICML*, pages 1305–1313, 2014.