

Machine Learning Techniques for MultiAgent Systems

Yoad Lewenberg

The Hebrew University of Jerusalem

yoadlew@cs.huji.ac.il

1 Introduction

Research in artificial intelligence ranges over many subdisciplines, such as Natural Language Processing, Computer Vision, Machine Learning, and MultiAgent Systems. Recently, AI techniques have become increasingly robust and complex, and there has been enhanced interest in research at the intersection of seemingly disparate research areas. Such work is motivated by the observation that there is actually a great deal of commonality among areas, that can be exploited within subfields. One example of a successful combination is the intersection of machine learning and multiagent systems. For example, [Kearns *et al.*, 2001] proposed an efficient graphical model-based algorithm for calculating Nash equilibria. Going in the other direction, [Datta *et al.*, 2015] showed that solution concepts from cooperative game theory can be used to uniquely characterize the influence measure of classifiers.

1.1 Machine Learning

The recent surge in research on machine learning has shown an extremely diverse set of domains where machine learning approaches have excellent performance. In particular, approaches based on *Deep Convolutional Neural Networks* achieve very good performance on various domains such as image classification, speech recognition, and control tasks. In many cases, CNN-based approaches outperform systems developed over many years.

Another machine learning approach that has earned significant and growing interest is *Probabilistic Graphical Models*. Probabilistic graphical models are powerful tools for compactly representing complex probability distributions and for expressing conditional dependencies among random variables. Graphical models are widely studied in artificial intelligence, and were used, for example, for Bayesian probabilistic matrix factorization, and for large-scale recommender systems.

1.2 MultiAgent Systems

Multiagent systems is a subfield of AI that aims to provide both principles for the construction of complex systems involving multiple agents, as well as mechanisms for coordination of independent agents' behaviors. In recent years, especially since the rapid growth in the number of people using the World Wide Web, research in the field has dramatically

expanded. In particular, there is increased interest in developing autonomous agents that can act effectively on our behalf in interactions that are commonly seen on the web and on social media, such as digital currencies, voting, and cooperation among selfish agents.

1.3 Integration of Machine Learning and MultiAgent Systems

As machine (and human-machine) interactions typically involve several self-interested entities, those interactions have significant inherent complexity. Therefore, there is much interest in using machine learning techniques to help deal with this complexity. Indeed, the combination of machine learning and multiagent systems opens up a potentially wide field of research. Examples include multiagent learning algorithms, and the application of machine learning techniques to multiagent environments.

The use of multiagent reinforcement learning algorithms has attracted attention because of its generality and robustness. Those techniques have been applied in both cooperative environments and non-cooperative environments.

Another line of work is the application of machine learning tools for uncertain environments in multiagent systems. This research direction stretches from voting to cooperative games.

In my own research, I strive to extend this approach, and combine topics from multiagent systems with those from machine learning

2 Ongoing Research

The research I have carried out during my master's and PhD studies so far has been focused primarily on multiagent systems and machine learning.

2.1 MultiAgent Systems

In multiagent systems we have studied issues in an all-pay auction, digital currencies [Lewenberg *et al.*, 2015a; 2015b], and district voting.

All-pay auctions All-pay auctions, a common mechanism for various human and agent interactions, suffers, like many other mechanisms, from the possibility of bidders' failure to participate in the auction. We model such failures and show how they affect the equilibrium state [Lewenberg *et al.*, 2013].

District voting Voting systems in which voters are partitioned to districts encourage accountability by providing voters an easily identifiable district representative, in some cases, a party may have a majority of the popular vote, but lose the elections due to districting effects. We studied this phenomenon [Bachrach *et al.*, 2016a], and provided bounds for “cost of districting” under various voting rules. Further, we showed [Lewenberg *et al.*, 2017b] that a geographically-based manipulation in district-based elections is NP-complete in the worst case. However, a greedy algorithm can be used to solve that problem.

2.2 Machine Learning

In machine learning we have studied applications of graphical models [Bachrach *et al.*, 2016b] and deep convolutional neural networks.

Graphical models Most systems for classifying political ideologies use a Left-Right classification. However, most political parties take an ideological stand regarding many issues, and one dimension does not seem sufficient for expressing the political spectrum. Using graphical models we showed [Lewenberg *et al.*, 2016a] that although the political spectrum is richer than a simple left-right structure, fewer than 10 dimensions are sufficient to represent peoples’s entire political opinion.

In a follow-up study, we tackled the surveying problem [Lewenberg *et al.*, 2017a]. In this domain, we attempted to predict how a target user is likely to respond to questions by iteratively querying that user, collaboratively-based on the responses of a sample set of users. We proposed a method for solving the problem based on a Bayesian dimensionality reduction technique. We empirically evaluated our method, contrasting it to benchmark approaches based on augmented linear regression, and showed that it both achieves much better predictive performance, and is much more robust when there is missing data.

Deep convolutional neural networks We considered the task of predicting various traits of a person given an image of their face [Lewenberg *et al.*, 2016b]. We attempted to estimate both objective traits, such as gender, ethnicity, and age, as well as subjective traits, such as the emotion a person expresses or whether they are humorous or attractive. We began by using a CNN architecture, and corroborated that CNNs are promising for facial attribute prediction. To further improve performance, we proposed a novel approach that incorporates facial landmark information for input images as an additional channel, helping the CNN learn face-specific features so that the landmarks across various training images hold a correspondence. We empirically analyzed the performance of our proposed method, showing consistent improvement over the baselines across traits.

3 Future Work

We now briefly discuss how approaches and techniques from machine learning can be integrated within multiagent systems tasks to provide representations of agents’ beliefs and for control tasks in multiagent environments.

In multiagent group decision problems, the agents are usually asked to report their complete ranking over the set of

alternatives. However, in some cases, the ranking may be incomplete or even inconsistent. In those cases, probabilistic graphical models can be used for representing agents’ beliefs. The graphical model can observe the incomplete ranking, infer the ballot and after applying the voting rule reach a decision.

In some domains, agents are situated in a cooperative environment, where the agents need to negotiate in order to form teams and solve a joint problem. Agents for such tasks can be designed using deep q-learning methods. One of the advantages of deep neural networks over feature learning methods and deep q-learning methods over q-learning is that there is no need for designing features by hand, as deep learning techniques can automatically learn feature representations. After an agent has been trained, the learned features can be compared to solution concepts in cooperative game theory, and we may observe whether the natural representation of the environment that the neural network has learned is consistent with game-theoretic solution concepts.

References

- [Bachrach *et al.*, 2016a] Yoram Bachrach, Omer Lev, Yoad Lewenberg, and Yair Zick. Misrepresentation in district voting. In *IJCAI*, pages 81–81, 2016.
- [Bachrach *et al.*, 2016b] Yoram Bachrach, Yoad Lewenberg, Jeffrey S Rosenschein, and Yair Zick. Tracking performance and forming study groups for prep courses using probabilistic graphical models. In *AAMAS*, pages 1359–1360, 2016.
- [Datta *et al.*, 2015] Amit Datta, Anupam Datta, Ariel D Procaccia, and Yair Zick. Influence in classification via cooperative game theory. *arXiv preprint arXiv:1505.00036*, 2015.
- [Kearns *et al.*, 2001] Michael Kearns, Michael L Littman, and Satinder Singh. Graphical models for game theory. In *UAI*, pages 253–260, 2001.
- [Lewenberg *et al.*, 2013] Yoad Lewenberg, Omer Lev, Yoram Bachrach, and Jeffrey S Rosenschein. Agent failures in all-pay auctions. In *IJCAI*, 2013.
- [Lewenberg *et al.*, 2015a] Yoad Lewenberg, Yoram Bachrach, Yonatan Sompolinsky, Aviv Zohar, and Jeffrey S Rosenschein. Bitcoin mining pools: A cooperative game theoretic analysis. In *AAMAS*, pages 919–927, 2015.
- [Lewenberg *et al.*, 2015b] Yoad Lewenberg, Yonatan Sompolinsky, and Aviv Zohar. Inclusive block chain protocols. In *FC*, pages 528–547, 2015.
- [Lewenberg *et al.*, 2016a] Yoad Lewenberg, Yoram Bachrach, Lucas Bordeaux, and Pushmeet Kohli. Political dimensionality estimation using a probabilistic graphical model. *UAI*, 2016.
- [Lewenberg *et al.*, 2016b] Yoad Lewenberg, Yoram Bachrach, Sukrit Shankar, and Antonio Criminisi. Predicting personal traits from facial images using convolutional neural networks augmented with facial landmark information. *IJCAI*, pages 1676–1682, 2016.
- [Lewenberg *et al.*, 2017a] Yoad Lewenberg, Yoram Bachrach, Ulrich Paquet, and Jeffrey S. Rosenschein. Knowing what to ask: A bayesian active learning approach to the surveying problem. *AAAI*, 2017.
- [Lewenberg *et al.*, 2017b] Yoad Lewenberg, Omer Lev, and Jeffrey S. Rosenschein. Divide and conquer: using geographic manipulation to win district-based elections. *AAMAS*, 2017.