Conversing Learning: Actively looking for human assistance to improve Machine Learning tasks

Saulo D. S. Pedro
Federal University of Sao Carlos
saulods.pedro@gmail.com

Estevam R. Hruschka Jr.
Federal University of Sao Carlos
estevam@dc.ufscar.br

ABSTRACT

In Machine Learning (ML) systems we often apply techniques to learn about real world problems behavior from data. Traditionally, such data comes from instances of datasets (that represents the target problem) which we want to learn from. This approach has been broadly used in many different application domains such as recommending systems, meteorological prediction, medical diagnosis, etc. The recent years of quick development of communications technology, that made the Internet faster and available, made possible the acquisition of information to feed a growing number of Machine Learning applications and, in addition, brought light to the use of human computation and crowdsourcing approaches commonly applied to problems that are easy for human but difficult for computers. Thus, the Social Web has been the focus of many research in Artificial Intelligence (AI) and Machine Learning. In this work we want to show how we can take advantage from the Social Web to add value to Machine Learning systems which can actively and autonomously ask for web users help to improve learning performance.

Categories and Subject Descriptors

I.2.6 [**Artificial Intelligence**]: Learning—*knowledge acquisition*

General Terms

Machine Learning

Keywords

self-supervision, self-revision, crowdsourcing

1. CONVERSING LEARNING

This approach is inspired by other Machine Learning techniques such as Active Learning (AL) [2, 6] and Interactive Learning (IL) [7]. In a nutshell, AL allows an oracle to select of a subset of a dataset to improve the learning task that depends on this data. Moreover, in IL we can explore the oracle to continuously improve a learning system in each cycle of consulting. In the IL approach, the learning system analyzes results acquired from the last consult and deploy an improved knowledge base as a new iteration. The main goal of this paper is to present and discuss the main ideas and principles of a new approach designed to take advantage of the union of Machine learning and crowdsourcing, called Conversing Learning (CL)[3]. Conversing Learning is based on both, Active Learning and Interactive Learning, and is intended to allow machines to convert their knowledge base into human understandable content and then, actively and autonomously ask people (Web users) to take part into the knowledge acquisition (and labeling) process[3]. The idea of Conversing Learning was first introduced in [5, 3] and has been applied mainly to NELL (as in [4]), a Never-Ending Language Learner that is running since January 2010 [1]. This paper is intended to briefly overview CL main characteristics and principles, as well as to show that it can be applied to a broad set of different applications and domains.

The application of Conversing Learning to resolve Machine Learning problems depends on the addressing of a few core questions. First it's needed to decide what is the problem that is going to be put into human consideration. These problems can include for an example validation and revision. Another feature that needs definition is the method of conversation with these web users. We have already connected with users through social media applications but there could be several other ways to perform this task such as surveys and forums. The method of conversation also includes addressing how to convert the machine knowledge into a human understandable form which might require resolution of Natural Language Processing issues, since we need to use human language. Finally the CL system must include a way to analyze the web users contributions and feedback the Machine Learning system.

2. SS-CROWD

In [5], we designed an algorithm to take advantage of the wisdom of crowds and put to test the idea of using human collaboration in autonomous learning tasks. The system was named SS-Crowd (Self-Supervisor Agent Based on the Wisdom of Crowds). SS-Crowd can be seen as one of the methods to connect with web users through web communities. The algorithm was first experimented with a set of rules created by NELL (Never-Ending Language Learner) system, which is a computer system that runs 24 hours per day, 7 days per week focusing on gathering knowledge from web pages in English and using its acquired knowledge to become a better learner each day [1]. The rules created by NELL were put into web users from *Yahoo! Answers* and *Twitter* attention through SS-Crowd and its basic work flow is as follows:

- Convert the information from knowledge base into understandable questions.
- Input the question on Twitter and Yahoo! Answers
- Gather opinion from users and combine them into a single opinion

2.1 Experiments with SS-Crowd

With that work we intended to show how we can take advantage of collected intelligence in ML tasks as well as determining and addressing issues raised when working with human generated content. We also started investigating how the use of Conversing Learning in different communities could be explored. When working with human generated content by an stimulus, we might have to resolve important matters such as the determination of the channels and method of communication. We conducted experiments based on *Yahoo!Answers* as well as *Twitter* focusing on understanding how different communities would react to the same stimulus. So we used SS-Crowd to perform a validation task of NELL's KB by asking users about the rules created by NELL. From the total set of rules we took those that would have greater impact in NELL's KB, thus performing an AL task. This approach was possible because NELL has metadata indicating how much beliefs would be affected by an specific rule. We put the users of both communities to the same set of rules that were converted into 62 questions and received 350 answers for *Yahoo! Answers* and 72 answers for *Twitter*. First we determined that the opinion of the users differ 45% of their answers which increases our belief to be working with independent sources instead of redundant information. Since we had previously annotated opinion of NELL's developers about the validity of each rule, we built a naive-bayes classifier to help us decide which is the better community to perform this specific test. The classifier accuracy (user's answer matching NELL's developer answers) was 77% for *Twitter* and 70% for *Yahoo! Answers*. Based on the obtained results, we concluded that the motivation as well as people that takes part in web communities can vary depending on the different communities targeted. Therefore, the choice of the community can influence the success of this approach. The same variation happens when we think about how to automatically formulate the questions that will bring more useful answers. In our experiments we performed keyword based method to capture the user's opinion about NELL's knowledge. In the beginning we had issues to extract the user opinion from the answer, but the accuracy of this task increased when we driven the user for a more simple answer, which are easier to process. This was achieved by changing the questions to Yes/No questions. Besides the simple approach, the total of questions left with undetermined opinion went from 24% to 5%[3]

3. DISCUSSION

Considering the amount of available data and the potential of Machine Learning algorithms and techniques, Knowledge bases can have a tendency to be infinite, and NELL's KB is an example of it. Therefore, it would be difficult to use the wisdom of the crowds (in a CL approach) to validate every bit of knowledge in such huge knowledge bases. Although we have millions of people to ask, focusing on more critical piece of knowledge might bring more accurate collaboration faster. Communities on the web may change focus quickly and it might be difficult to reach human collaboration unless we are working with a community specially designed for Machine Learning purposes. One way to approach this matter could be the application of Active Learning tasks to select the subset of knowledge that, when put into human consideration would make greater impact on the knowledge base. This subset can be considered critical because of its potential to improve the ML system accuracy substantially. The same idea can be applied when choosing an specific subset of channels. Those are examples of the principles behind CL. Also, it is important to consider that Web users are sensitive to the awareness of the intentions of the query they are being prompted. Users tend to behave differently when they know a ML research is going to use the data generated from them. We noticed that the awareness of users influences the plainness and detailing of their answers and this difference could be used in the favor of CL systems while being developed.

In previous work we have compared the opinion of Yahoo! Answers and Twitter with opinions previously annotated by NELL's specialists during the human supervision procedure. The results we had, from this comparison, were not enough accurate to use Conversing Learning methods as a surrogate of NELL's current verification procedure. It could be used, however, to point out subsets of knowledge that needs attention, thus performing a self-supervision task. Such characteristic could also be extended to a self-revision task. Considering, for instance, adding a new step to CL, in which Web users are asked to give their help to correct knowledge flagged as wrong, in this sense, the wisdom of the crowds could be used as a review process.

Moreover, when we work with different sources of human generated content, if we understand the channel of communication and its limits and features, then we can drive specific issues to specific communities. Following along these lines, we have noticed that *Yahoo!Answers* users give more attention to details and provide more complex answers. Twitter users, on the other hand, are more straightforward (mainly because they are NELL followers, thus, they are motivated to help NELL to learning better). With a deep understanding of these capabilities, we could design CL systems that know where to look for human assistance depending on the problem being assessed, thus making possible a future raise of self-reflection tasks in the CL system itself. With this analysis, we came to the conclusion that applying Conversing Learning to improve learning tasks can bring more challenges to crowdsourcing research. The CL capability of actively asking for human collaboration allows the creation of autonomous tasks such as validation and knowledge acquisition which are critical for systems that learns continuously.

As aforementioned, CL has been mainly applied to NELL. NELL's architecture allows the system to generate constraints that will guide the learning process, thus, avoiding wrong concepts to be inserted in the KB. To be even more precise, and also, to review and verify whether the facts stored in its KB are correct or not, currently the system can take advantage of four different approaches: i) ***Human Supervision***: RTW[1] group members can spend 5 minutes per day validating NELL's extractions; ii) ***Web Querying***: NELL can query the Web on specific facts to verify correctness, or to predict the validity of a new fact; iii) ***Conversing Learning***[5]: NELL can autonomously talk to people (have conversations) in web communities and ask for help on validating specific facts, rules or meta-data (features, mutually exclusiveness relationships, etc.); and iv) ***Hiring Labelers***: NELL can autonomously hire people (using web services such as Mechanical Turk, by Amazon) to label data and help the system to validate acquired knowledge.

4. CONCLUSION AND FUTURE WORK

All those possibilities of reviewing and verifying the truth value of facts stored in the knowledge base can be used in order to help the system's self-supervision and self-reflection. In this sense, suppose for example that, NELL can use the Mechanical Turk labeling component to infer categories (or relations) which have a low accuracy and low volume of learned facts. Then, the system can use Conversing Learning, in the sequel, and also, its Web-Querying component, to identify new features for the learners. NELL can perform this "Mechanical Turk-Conversing Learning" cycle until identifying that the accuracy and the volume of the learned facts have improved. This *Mechanical Turk-Conversing Learning* cycle reveals that CL is not the only way to explore the wisdom of the crowds in a Machine Learning environment, it also shows that there are specific situations where CL is more suitable and situations where other approaches will fit better.

In summary, it is possible to say that CL is not suitable in situations where too many queries should be performed. In such a scenario, where it is important to validate every bit of knowledge in huge knowledge bases, a Mechanical Turk task, for instance, might be better. In scenarios where the wisdom of the crowds will be used focusing only in specific aspects, or in critical portions of knowledge base, then, CL might bring more accurate results faster.

In the future, we want to apply Conversing Learning in different problems to understand better how broad this approach can be and which tasks could be unleashed by it. We are also interested in exploring the web users to provide information that will be be used in the metadata (features, mutually exclusiveness constraints, etc.) of learning algorithms. That is, use human generated opinion to make changes to the core of learning algorithms aiming to enhance its learning task.

5. REFERENCES

[1] A. Carlson, J. Betteridge, B. Kisiel, B. Settles, E. R. Hruschka Jr., and T. M. Mitchell. Toward an architecture for never-ending language learning. In *Proceedings of AAAI2010*, volume 2, pages 1306–1313, 2010.

[2] D. Cohn, Z. Ghahramani, and M. Jordan. Active learning with statistical models. *Journal of Artificial Intelligence Research*, 4:129–145, 1996.

[3] S. D. Pedro and E. R. Hruschka Jr. Conversing learning: Active learning and active social interaction for human supervision in never-ending learning systems. In *Advances in Artificial Intelligence–IBERAMIA 2012*, pages 231–240. Springer, 2012.

[4] S. D. S. Pedro, A. P. Appel, and E. R. Hruschka,Jr. Autonomously reviewing and validating the knowledge base of a never-ending learning system. In *Proceedings of the 5th Int. Workshop on Web Intelligence & Communities*, WI&C13. ACM, 2013.

[5] S. D. S. Pedro and E. R. Hruschka,Jr. Collective intelligence as a source for machine learning self-supervision. In *Proceedings of the 4th Int. Workshop on Web Intelligence & Communities*, WI&C12, pages 5:1–5:9. ACM, 2012.

[6] B. Settles. Active learning literature survey. *SciencesNew York*, 15(2):201–221, 2010.

[7] B. Settles. Closing the loop: Fast, interactive semi-supervised annotation with queries on features and instances. In *Proc. of EMNLP2011*, pages 1467–1478, Edinburgh, Scotland, UK., July 2011.

[1] http://rtw.ml.cmu.edu