MACHINE LEARNING

Priyanka Sharma

Raj Kumar Goel Institute of Technology for Women, Ghaziabad, (India)

ABSTRACT

Much of current machine learning (ML) research has lost its connection to problems of import to the larger world of science and society. From this perspective, there exist glaring limitations in the data sets we investigate, the metrics we employ for evaluation, and the degree to which results are communicated back to their originating domains. What changes are needed to how we con- duct research to increase the impact that ML has? We present six Impact Challenges to explicitly focus the field's energy and attention, and we discuss existing obstacles that must be addressed. We aim to inspire ongoing discussion and focus on ML that matters.

Keywords: ML (Machine Learning), CALO (Cognitive Assistant That Learns and Organizes), ICML (International Council of Machine Learning, ROC (Receiver Operating Characteristic)

I. INTRODUCTION

At one time or another, we all encounter a friend, spouse, parent, child, or concerned citizen who, upon learning that we work in machine learning, wonders "What's it good for?" The question may be phrased more subtly or elegantly, but no matter its form, it gets at the motivational underpinnings of the work that we do. Why do we invest years of our professional lives in machine learning research? What differEnce does it make, to ourselves and to the world at large?

Much of machine learning (ML) research is inspired by weighty problems from biology, medicine, finance , astronomy, etc. The growing area of computational sustainability (Gomes, 2009) seeks to connect ML advances to real-world challenges in the environment, economy, and society. The CALO (Cognitive Assistant that Learns and Organizes) project aimed to integrate learning and reasoning into a desktop assistant, potentially impacting everyone who uses a computer (SRI International, 2003-2009) translation (Koehnetal , 2003), two problems of global import. And so on.

And yet we still observe a proliferation of published ML papers that evaluate new algorithms on a handful of isolated benchmark data sets. Their "real world" experiments may operate on data that originated in the real world, but the results are rarely communicated back to the origin. Quantitative improvements in performance are rarely accompanied by an assessment of whether those gains matter to the world outside of machine learning research.This phenomenon occurs because there is no widespread emphasis, in the training of graduate stu- dent researchers or in the review process for submitted papers, on connecting ML advances back to the larger world. Even the rich assortment of applications-driven ML research often fails to take the final step to trans- late results into impact.

Many machine learning problems are phrased in terms of an objective function to be optimized. It is time for us to ask a question of larger scope: what is the field's objective function? Do we seek to maximize performance on isolated data sets? Or can we characterize progress in a more meaningful way that measures the concrete impact of machine learning innovations?

This short position paper argues for a change in how we view the relationship between machine learning and science (and the rest of society). This paper does not contain any algorithms, theorems, experiments, or results. Instead it seeks to stimulate creative thought and research into a large but relatively unaddressed issue that underlies much of the machine learning field. The contributions of this work are 1) the clear identification and description of a fundamental problem: the frequent lack of connection between machine learning research and the larger world of scientific inquiry and humanity, 2) suggested first steps towards addressing this gap, 3) the issuance of relevant Impact Challenges to the machine learning community, and 4) the identification of several key obstacles to machine leaearning impact, as an aid for focusing future research efforts. Wether or not the reader agrees with all statements in this paper, if it inspires thought and discussion, then its purpose has been achieved.

II. MACHINE LEARNING FOR MACHINE LEARNING'S SAKE

This section highlights aspects of the way ML research is conducted today that limit its impact on the larger world. Our goal is not to point fingers or critique individuals, but instead to initiate a critical self-inspection and constructive, creative changes. These problems do not trouble all ML work, but they are common enough to merit our effort in eliminating them.The argument here is also about "theory versus applications." Theoretical work can be as inspired by real problems as applied work can. The criticisms here focus instead on the limitations of work that lies be- tween theory and meaningful applications: algorithmic advances accompanied by empirical studies that are divorced from true impact.

2.1 Hyper-Focus on Benchmark Data Sets

Increasingly, ML papers that describe a new algorithm follow a standard evaluation template. After presenting results on synthetic data sets to illustrate certain aspects of the algorithm's behavior, the paper reports results on a collection of standard data sets, such as those available in the UCI archive (Frank & Asuncion, 2010). A survey of the 152 non-cross-conference papers published at ICML 2011 reveals:

148/152 (93%) include experiments of some sort 57/148 (39%) use synthetic data 55/148 (37%) use UCI data 34/148 (23%) use ONLY UCI and/or synthetic data 1/148 (1%) interpret results in domain context

The possible advantages of using familiar data sets include 1) enabling direct empirical comparisons with other methods and 2) greater ease of interpreting the results since (presumably) the data set properties have been widely studied and understood. However, in practice direct comparisons fail because we have no standard for reproducibility. Experiments vary in methodology (train/test splits, evaluation metrics, parameter settings), implementations, or reporting. Interpretations are almost never made. Why is this?

First, meaningful interpretations are hard. Virtually none of the ML researchers who work with these data sets happen to also be experts in the relevant scientific disciplines. Second, and more insidiously, the ML field neither motivates nor requires such interpretation. Reviewers do not inquire as to which classes were well classified and which were not, what the common error types were, or even why the particular data sets were chosen. There is no expectation that the authors report whether an observed x% improvement in performance promises any real impact for the original domain. Even when the authors have forged a collaboration with qualified experts, little paper space is devoted to interpretation, because we (as a field) do not require it.

III. CONCLUSION

Machine learning offers a cornucopia of useful ways to approach problems that otherwise defy manual solution. However, much current ML research suffers from a growing detachment from those real problems. Many investigators withdraw into their private studies with a copy of the data set and work in isolation to perfect algorithmic performance. Publishing results to the ML community is the end of the process. Successes usually are not communicated back to the original problem setting, or not in a form that can be used.

Yet these opportunities for real impact are widespread. The worlds of law, finance, politics, medicine, education, and more stand to benefit from systems that can analyze, adapt, and take (or at least recommend) action. This paper identifies six examples of Impact Challenges and several real obstacles in the hope of inspiring a lively discussion of how ML can best make a difference. Aiming for real impact does not just increase our job satisfaction (though it may well do that); it is the only way to get the rest of the world to notice, recognize, value, and adopt ML solutions.

IV. ACKNOWLEDGMENTS

I thank Mrs. Kadambri Aggrwal, Ms. Charu Gupta, Mrs. Deepali Singhal & insightful anonymous reviewers for suggestions on this paper.

REFERENCES

[1] http..//www.coursera.org/course/ml
[2] cs229.stanford.edu/
[3] azure.mmicrosoft.com/en-us/services/machinelearning/