Mach Learn
DOI 10.1007/s10994-013-5335-x

From machine learning to machine reasoning

An essay

Léon Bottou

Received: 27 April 2012 / Accepted: 13 February 2013

Abstract A plausible definition of "*reasoning*" could be "*algebraically manipulating previously acquired knowledge in order to answer a new question*". This definition covers first-order logical inference or probabilistic inference. It also includes much simpler manipulations commonly used to build large learning systems. For instance, we can build an optical character recognition system by first training a character segmenter, an isolated character recognizer, and a language model, using appropriate labelled training sets. Adequately concatenating these modules and fine tuning the resulting system can be viewed as an algebraic operation in a space of models. The resulting model answers a new question, that is, converting the image of a text page into a computer readable text.

This observation suggests a conceptual continuity between algebraically rich inference systems, such as logical or probabilistic inference, and simple manipulations, such as the mere concatenation of trainable learning systems. Therefore, instead of trying to bridge the gap between machine learning systems and sophisticated "all-purpose" inference mechanisms, we can instead algebraically enrich the set of manipulations applicable to training systems, and build reasoning capabilities from the ground up.

Keywords Machine learning · Reasoning · Recursive networks

1 Introduction

Since learning and reasoning are two essential abilities associated with intelligence, machine learning and machine reasoning have both received much attention during the short history of computer science. The statistical nature of learning is now well understood (e.g., Vapnik 1995). Statistical machine learning methods are now commonplace (NIPS 1987–2010). An internet search for "*support vector machines*" returns more than two million web pages.

Editors: Ronan Collobert and Luke Zettlemoyer.

This essay is an updated version of the unpublished report (Bottou 2011).

L. Bottou (✉)
Microsoft Research, Redmond, WA, USA
e-mail: leon@bottou.org

Published online: 10 April 2013