

EPISTEMOLOGY OF MACHINE LEARNING
ASSIGNMENT #1

Deadline: **Thursday May 17**, before class. Strive to answer concisely yet informatively. Aim for about half a page per question; certainly not more than one page.

- (1) Lipton distinguishes two problems about induction: the problem of justification and the problem of description. Explain both problems for the specific case of the ERM algorithm (“enumerative induction” in the words of Harman & Kulkarni).
- (2) Vapnik in the introduction of his book contrasts the theoretical approach of SLT with that of the “AI hardliners” (those “who declared that ‘complex theories do not work; simple algorithms [on large amounts of data] do,’” p. 13). Discuss whether and in what respects Wheeler draws a similar contrast between traditional epistemology and his proposed “machine epistemology.”
- (3) Suppose that we have an hypothesis class \mathcal{H} that is PAC learnable with the ERM algorithm, and that the realizability assumption holds. We fix some confidence parameter δ and error parameter ϵ , and we draw more than $m_{\mathcal{H}}(\epsilon, \delta)$ examples i.i.d. from the unknown distribution \mathcal{D} , giving a sample S . At this point we select, in accordance with the ERM procedure, an hypothesis $h_S \in \mathcal{H}$ with empirical error $L_S(h_S) = 0$. Can we now proclaim that with probability at least $1 - \delta$ the true risk $L_{(\mathcal{D}, f)}(h_S)$ of this particular hypothesis h_S is less than ϵ ?

(Hint: point out the difference between objective-frequentist and epistemic probability. It might also be helpful to have a look at section 9.2 of the paper by von Luxburg & Schölkopf, 2011.)