Best Paper Track A: Online Learning and Disambiguations of Partial Concept Classes

Wednesday, July 12, 2023 2:50 PM (20 minutes)

Tsun-Ming Cheung, Hamed Hatami, Pooya Hatami and Kaave Hosseini

Best Paper Track A:

Abstract: In many practical learning problems, the learning task is tractable because it is only required to predict the labels of the data points with specific properties. For example, in large-margin classifiers, one seeks to classify only the data points bounded away from the decision boundary by a margin.

In a recent article, Alon, Hanneke, Holzman, and Moran (FOCS '21) introduced a unifying framework to study the learnability of classes of such partial concepts. One of the central questions studied in their work is whether the learnability of a partial concept class is always inherited from the learnability of some "extension" of it to a total concept class.

They showed that this is not the case for PAC learning but left the problem for the stronger notion of online learnability open.

We answer their question in the negative by constructing a class of partial concepts that is online learnable, but no extension of it to a class of total concepts is online learnable (or even PAC learnable).

Presenter: CHEUNG, Tsun-Ming

Session Classification: ICALP 2023 Award-Winning Papers