

Evaluation of Unsupervised Learning Results: Making the Seemingly Impossible Possible

Ricardo J. G. B. Campello*

May 4, 2019

Abstract of Invited Presentation

When labels are not around, evaluating the final or intermediate results produced by a learning algorithm is usually not simple. In cluster analysis and unsupervised outlier detection, evaluation is important in many different aspects. It is a crucial task, e.g., for model selection, model validation, assessment of ensemble members accuracy and diversity, among others. In cluster analysis this task has been investigated for decades, but it is relatively well understood only under certain oversimplified model assumptions. In outlier detection, unsupervised evaluation is still in its infancy. Even when labels are available in the form of a ground-truth, such as in controlled benchmarking experiments, evaluation can still be challenging because the semantic described by the ground-truth labels may not be properly captured by the data as represented in the given feature space.

In this talk, I intend to discuss some particular aspects regarding outlier and clustering evaluation, focusing on a few recent results as well as some challenges for future research.

*University of Newcastle