
VALERIE POULIN, Tutte Institute for Mathematics and Computing
Evaluating Graph Clustering Algorithms: Beware!

An impressive number of graph clustering algorithms have been proposed, studied and compared over the past decades. To identify better graph clustering techniques, one needs a way to score the techniques against one another. A typical method is to compare values of some similarity measure between ground truth partitions of given graphs and the partitions produced by the different algorithms on those graphs. However, the choice of the similarity measure used is crucial and has a huge impact on the conclusions made.

In graph clustering comparison studies, set partition similarities are used as accuracy measures. Typically, a member of the pair-counting family such as Adjusted Rand Index, or of the Shannon information-based family such as Adjusted Mutual Information is used to assess the superiority of a graph clustering algorithm over another. These measures are designed for comparing set partitions and not graph partitions specifically. We call them *graph-agnostic* as they ignore the graph structure.

In this work, we introduce a family of *graph-aware* similarity measures with their adjusted forms for graph partitions and prove that the combination of both graph-agnostic and graph-aware measures is critical for effectively comparing graph partitions.