# Experimental Evaluation of Subgraph Isomorphism Solvers

Chritine Solnon

March 26, 2019

Subgraph Isomorphism (SI) is an NP-complete problem which is at the heart of many structural pattern recognition tasks as it involves finding a copy of a pattern graph into a target graph. In the pattern recognition community, the most well-known SI solvers are VF2, VF3, and RI. SI is also widely studied in the constraint programming community, and many constraint-based SI solvers have been proposed since Ullman, such as LAD and Glasgow, for example. All these SI solvers can solve very quickly some large SI instances, that involve graphs with thousands of nodes. However, McCreesh et al. have recently shown how to randomly generate SI instances the hardness of which can be controlled and predicted. In particular, they have shown how to generate small instances (with thirty pattern nodes and 200 target nodes, for example) which are computationally challenging for all solvers. This study also showed that some small instances which are easily solved by constraint-based solvers, appear to be challenging for VF2 and VF3. In this talk, we will widen this study by considering a large test suite of 14,621 instances coming from eight different benchmarks. We will show that, as expected for an NP-complete problem, the solving time of an instance does not depend on its size, and that some small instances coming from real applications are not solved by any of the considered solvers. We will also show that, if RI and VF3 are able to solve very quickly a large number of easy instances, for which Glasgow or LAD need more time, they fail at solving some other instances that are quickly solved by Glasgow or LAD, and they are clearly outperformed by Glasgow on hard instances. Finally, we will show that we can easily combine solvers to take benefit of their complementarity.