

Complete Explanations
Position Paper
PTHG-18

Eugene C. Freuder
Insight Centre for Data Analytics
University College Cork
Cork, Ireland
eugene.freuder@insight-centre.org

As AI becomes more ubiquitous there is a renewed interest in computers being able to provide explanations, and the European GDPR provides special impetus. In constraint satisfaction, on the one hand we are fortunate in that an explanation for a successful solution is very straightforward: “see, the constraints are satisfied”. However, when a constraint satisfaction problem is unsolvable, explanations are difficult, as there can be an exponential number of reasons for failure, corresponding to every way that the constraints cannot be satisfied, and there can be many different routes to arriving at the conclusion that satisfaction is impossible. To misquote Tolstoy: Solvable CSPs are all alike; every unsolvable CSP is unsolvable in its own way.

A number of approaches have been taken to providing explanations for constraint satisfaction failure. This position paper is restricted to efforts to provide explanations to users as opposed to explanations intended to make algorithms more efficient or to aid programmers. Given the exponential threat, and to address specific needs, these efforts generally start with an abstracted or higher-level form of explanation, e.g. sets of unsatisfiable constraints, and then quickly limit their focus, e.g. to minimal sets of unsatisfiable constraints.

The position taken here is that it can be worthwhile to start with truly complete explanations and abstract and limit from there. The goal is to provide a high-level “big picture” of the structure of the problem, in a form readily meaningful to a human user. The hope is that this may, as well, lead to general insights into constraint satisfaction problem structure.

Even though it is straightforward to explain any one solution, one may also want a more complete explanation of solvability that encompasses all the solutions. Again, we can look for abstract or compact representations. Of course, for most problems some options will be failures and some successes, and we can combine complete representations of both, and abstractions thereof. (One may also be interested in how solutions or failures are arrived at.)

Abstractions and limitations derived from complete explanations can still address scaling concerns. Also as AI becomes increasingly pervasive in everyday life, attention will shift to some degree to smaller scale problems where explanation size is not as great an issue. In any case, the key issue is not so much how to minimize explanation size, as how to maximize explanation utility. We need not just scalable algorithms but effective human-computer interfaces, including visualization tools, that help users grasp the big picture and explore their options.