Reshaping State-Space Search: From Dominance to Contrastive Analysis

Álvaro Torralba

Aalborg University, Aalborg, Denmark

Abstract

State-space search is paramount for intelligent decision making when long-term thinking is needed. We introduce dominance and contrastive analysis methods, which enable reasoning about the relative advantages among different courses of action. This re-shapes how agents reason and leads to new families of state-space search algorithms.

Content

Automated Planning deals with sequential decision making (Ghallab, Nau, and Traverso 2004) where agents must find the best strategy in order to achieve their goals and/or maximize their reward, given their knowledge about the environment. In this context, state-space search techniques are widely used due to their ability to plan ahead, considering the consequences of multiple courses of action. The search is guided with evaluation functions, assigning each state a numerical value that estimates how good it is, according to the agent's goals. Multiple algorithms illustrate the success of the state-space search paradigm, such as A* (e.g., for classical planning) or Monte Carlo Tree Search (e.g., for online decision making on adversarial environments). Hence, *redefining how search algorithms reason about the state space can have a huge impact across multiple sub-areas in AI and beyond.*

Despite their variety, most search algorithms share the assumption that search nodes must be evaluated independently of each other. Instead, I propose to guide the search by reasoning about the relative advantages and disadvantages of the already explored states. For example, dominance analysis techniques directly compare states to identify if one is as good as another. This can be used for *dominance pruning* where search nodes are discarded if they correspond to worse paths than others (Torralba and Hoffmann 2015).

More generally, comparing states is a new source of information beyond evaluation/heuristic functions. Search algorithms can use these methods to better understand the structure of the underlying state space being explored. I coin the term *contrastive analysis* to refer to this general setting.

Broadly speaking, there are two questions to consider. First, how states can be compared? We can go beyond identifying states that are strictly better than others. To begin with, we can also quantify how much better/worse a state is (Torralba 2017), or even compare states under any notion of "advancing towards the goal" (Torralba 2018). Next steps are to compare states in terms of their advantages even when one is not as good as another in all aspects.

Second, how search algorithms can utilize this information? Several uses have been suggested beyond dominance pruning, such as identifying irrelevant actions (Torralba and Kissmann 2015), identifying actions that are guaranteed to start an optimal plan on a given state (Torralba 2017), identifying search nodes from which we can safely restart our search (Torralba 2018), or approximating whether the new state has any benefit beyond those explored previously (Groß, Torralba, and Fickert 2020). A more general question, however, is how to make the most out of this information. This may require to introduce completely new search algorithms that are able to reason about the explored region of the search space in new ways. Studying the optimal efficiency of current algorithms was a starting point in this direction (Torralba 2021).

In conclusion, contrastive analysis techniques introduce new ways to reason about the relative outcomes of actions. This leads to new families of search algorithms that can have an impact on AI areas where long-term thinking is important such as planning, game playing, and reinforcement learning.

References

Ghallab, M.; Nau, D.; and Traverso, P. 2004. *Automated Planning: Theory and Practice*. Morgan Kaufmann.

Groß, J.; Torralba, Á.; and Fickert, M. 2020. Novel Is Not Always Better: On the Relation between Novelty and Dominance Pruning. In *Proc. AAAI 2020*, 9875–9882.

Torralba, Á. 2017. From Qualitative to Quantitative Dominance Pruning for Optimal Planning. In *Proc. IJCAI 2017*, 4426–4432.

Torralba, Á. 2018. Completeness-Preserving Dominance Techniques for Satisficing Planning. In *Proc. IJCAI 2018*, 4844–4851.

Torralba, Á. 2021. On the Optimal Efficiency of A* with Dominance Pruning. In *Proc. AAAI 2021*, 12007–12014.

Torralba, Á.; and Hoffmann, J. 2015. Simulation-Based Admissible Dominance Pruning. In *Proc. IJCAI 2015*, 1689–1695.

Torralba, Á.; and Kissmann, P. 2015. Focusing on What Really Matters: Irrelevance Pruning in Merge-and-Shrink. In *Proc. SoCS 2015*, 122–130.