

# Distributional Metareasoning for Heuristic Search

Tianyi Gu

Department of Computer Science, University of New Hampshire, USA  
 gu@cs.unh.edu

## Abstract

Heuristic search methods are widely used in many real-world autonomous systems. Yet, people always want to solve search problems that are larger than time allows. To address these challenging problems, even suboptimally, a planning agent should be smart enough to intelligently allocate its computational resources, to think carefully about where in the state space it should spend time searching. For finding optimal solutions, we must examine every node that is not provably too expensive. In contrast, to find suboptimal solutions when under time pressure, we need to be very selective about which nodes to examine. In this work, we will demonstrate that estimates of uncertainty, represented as belief distributions, can be used to drive search effectively. This type of algorithmic approach is known as metareasoning, which refers to reasoning about which reasoning to do. We will provide examples of improved algorithms for real-time search, bounded-cost search, and situated planning.

## Introduction

Heuristic search methods are widely used in many real-world autonomous systems (Mudgal et al. 2005; Ferguson, Howard, and Likhachev 2008; Mudgal et al. 2005; Karpas and Magazzeni 2020). However, people always want to solve search problems that are larger than time allows. To solve the challenging problems, even suboptimally, a planning agent better be smart enough to intelligently allocate its computational resources, to think carefully about where in the state space it should spend time to search. In contrast to finding optimal solutions, in which we must examine every node that is not provably too expensive, we need to be very selective about which nodes to examine when under time pressure. This type of algorithmic approach was named metareasoning, which refers to reasoning about what to reason about. An agent equipped with a metareasoning component would solve a meta-level-reasoning problem in addition to the conventional object-level-reasoning problem (Horvitz 1990). These two problems typically differ in their utility functions. The meta-level utility is the expected utility associated with inference-related cost (i.e., deliberation cost), while the object-level utility is the expected utility associated with the state of the world (i.e., solution cost) with-

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

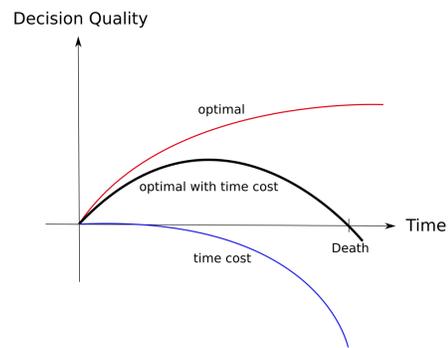


Figure 1: Illustration of optimal action selection with time cost.

out regard to the cost of reasoning. The relation of different types of cost is illustrated in Figure 1. The red line shows the so-called ‘intrinsic utility’ obtained by the agent as time progresses. The solution quality would increase given more and more computational time. The blue line is the time cost which would be getting worse and worse as time progresses. The black line shows the ‘net utility’ of taking an action at a particular time, which is the intrinsic utility minus the time cost. In this work, we propose practical metareasoning methods that would carefully allocate effort due to time pressure and thus optimize combined utility that takes into account both reasoning cost and object-level cost.

Historically, the metareasoning problem has been discussed for a long time since it was proposed (Good 1971; Russell and Wefald 1991; Zilberstein 2008). However, there only have been a few practical search algorithms that actually do metareasoning. It ought to be beneficial to design such an agent that explicitly optimizes both object-level utility and meta-level utility. For an agent to plan under time pressure, obviously, it should not only optimize the solution cost but also the planning time to achieve the best performance. Therefore, We pursue practical metareasoning components that can enhance various families of traditional search algorithms.

Conventional metareasoning approaches follow decision theory that tells the agent to select the action that maximizes the expected utility. However, when modeling an uncertain

value, scalar expected value often is not as powerful as a belief distribution. Intuitively, distribution helps because it quantifies uncertainty, which is what search resolves. Recently, distributional methods have been proposed in the RL community (Bellemare et al. 2017; Dabney et al. 2020), and it has been shown that distributional-informed methods can often outperform scalar-expected-value-informed methods by taking advantage of reasoning on value uncertainty. Therefore, taking inspiration from this prior work from RL, our work pursues an alternative class of rational metareasoning that takes advantage of distributional methods to have a better estimate model instead of only relying on the expected utility. The distributions can be constructed through offline learning or online learning. The thesis of this work is that distributional metareasoning can be a viable component for search and planning. We aim to improve algorithms for the following three problem settings: (1) real-time search (2) bounded-cost search, and (3) situated planning.

### Enhance Search by Metareasoning

Firstly, We will address real-time decision-making. In this setting, when the time bound is reached, the agent has to commit to the best action on hand even if it only has a partial solution plan. Because the time bound tightly limits the computation that an agent can perform, metareasoning could play an important role in this setting. Traditional real-time search methods were adapted directly from off-line search methods like A\* and make online action selection decisions based on a lower bound rather than an expected cost, which is not appropriate as a basis for rational action selection. To do a rational real-time search, it might be worth it for the agent to gather information about the value uncertainty due to the bounded rationality and make online decisions based on the value uncertainty as well as the expected utility. The traditional real-time search approaches are lacking this kind of metareasoning component. The first part of this work makes contributions to designing a rational real-time search approach (Fickert et al. 2020b,a).

Secondly, We propose algorithms for bounded-cost search settings where the agent is given a specific cost bound along with the search problem. The goal is to find a complete solution within the cost bound as quickly as possible. Bounded-cost search is also very useful since its users can have control over the solution quality. Traditional methods for bounded-cost search are focused on designing inadmissible heuristics that could guide the search toward the search nodes that have a high chance for finding a solution within the bound (Stern, Puzis, and Felner 2011). BEES (Thayer et al. 2012) explicitly tries to find a solution within the bound as quickly as possible, which is a meta-level problem. However, the performance of BEES can be very sensitive to the error of its estimate. In the second part of this work, We propose a distributional method to not only explicitly optimize the time to find a solution within bound but also take advantage of knowing the uncertainty of the estimate and thus better guide the search. (This work is currently under review.)

Thirdly, We also propose a metareasoning algorithm for online planning, specifically, answering the question of when to commit an action. When the planner commits to

an action, it re-roots its search at the node representing the outcome of that action. We assume that the system cannot be uncontrolled, so the planner must commit to a new action (perhaps a no-op) before the previously chosen action completes. In this setting, it can be beneficial to commit early, in order to devote more lookahead search focused below an upcoming state. In the third part of this work, we propose a principled method for making this commitment decision. (This work is still in progress.)

### Conclusion

Metareasoning can be a viable component for search and planning. This work has shown that one can adapt the distributional perspective from Bayesian RL to various deterministic single-agent planning settings and how deterministic planning can provide useful testbeds for methods that metareason about uncertainty during planning.

The first contribution of this work provides a metareasoning-based approach for rational real-time search. Instead of relying only on single heuristic value guidance, our framework, DDNancy, can learn a data-driven cost-to-go distribution to provide more rational guidance. We also adapt other RL methods (Interval Estimation and Monte-Carlo Tree Search) to real-time search, considering value uncertainty but not explicitly construct belief distribution. The results suggest that it can often achieve better performance than traditional approaches by considering value uncertainty.

The second contribution of this work introduces a distributional-enhanced approach for bounded-cost search. Our approach, XES, uses not just point estimates but belief distributions to estimate the probability that a node will lead to a plan within the bound. Our methods can often outperform current state-of-the-art bounded-cost search methods with the help of deliberate thinking on the value uncertainty.

The third contribution of this work design a metareasoning method to deal with the complex online problem setting of concurrent planning and execution. We propose a rational metric for the decision-making problem of when the agent should choose to commit action and think more about the future.

This work hopes to encourage further efforts in widening the applicability of metareasoning for suboptimal search and planning under time pressure. The results in this dissertation demonstrate that the methods proposed here bring rational decision-making to its day. Metareasoning can be very powerful when an agent has to deal with complex problems, especially when resources are tight. Many kinds of suboptimal search and planning methods can benefit from adding a metareasoning enhancement. Just like we did for real-time search, bounded-cost search, and situated planning.

### References

- Bellemare, M. G.; Dabney, W.; Munos; and Remi. 2017. A Distributional Perspective on Reinforcement Learning. In *the International Conference on Machine Learning*.
- Dabney, W.; Kurth-Nelson, Z.; Uchida, N.; Starkweather, C. K.; Hassabis, D.; Munos, R.; and Botvinick, M. 2020. A

distributional code for value in dopamine-based reinforcement learning. *Nature* 671–675.

Ferguson, D.; Howard, T. M.; and Likhachev, M. 2008. Motion planning in urban environments. *Journal of Field Robotics* 25(11-12): 939–960.

Fickert, M.; Gu, T.; Staut, L.; Lekyang, S.; Ruml, W.; Hoffmann, J.; and Petrik, M. 2020a. Real-time Planning as Data-driven Decision-making. In *the ICAPS-20 Workshop on Bridging the Gap Between AI Planning and Reinforcement Learning (PRL-20)*.

Fickert, M.; Gu, T.; Staut, L.; Ruml, W.; Hoffmann, J.; and Petrik, M. 2020b. Beliefs We Can Believe In: Replacing Assumptions with Data in Real-Time Search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 9827–9834.

Good, I. J. 1971. Twenty-seven principles of rationality. In *Godambe, V.P. & Sprott, D.A. (Eds), Foundations of Statistical Inference*, 108–141.

Horvitz, E. J. 1990. *Rational metareasoning and compilation for optimizing decisions under bounded resources*. Knowledge Systems Laboratory, Medical Computer Science, Stanford University.

Karpas, E.; and Magazzeni, D. 2020. Automated planning for robotics. *Annual Review of Control, Robotics, and Autonomous Systems* 3: 417–439.

Mudgal, A.; Tovey, C.; Greenberg, S.; and Koenig, S. 2005. Bounds on the travel cost of a mars rover prototype search heuristic. *SIAM Journal on Discrete Mathematics* 19(2): 431–447.

Russell, S. J.; and Wefald, E. 1991. *Do the right thing: studies in limited rationality*. MIT Press.

Stern, R. T.; Puzis, R.; and Felner, A. 2011. Potential search: A bounded-cost search algorithm. In *Twenty-First International Conference on Automated Planning and Scheduling*.

Thayer, J. T.; Stern, R.; Felner, A.; and Ruml, W. 2012. Faster bounded-cost search using inadmissible estimates. In *Twenty-Second International Conference on Automated Planning and Scheduling*.

Zilberstein, S. 2008. Metareasoning and Bounded Rationality. In *AAAI Workshop on Metareasoning: Thinking about Thinking*. Chicago, Illinois.