
Anticipatory Thinking: A New Frontier for Automated Planning

Adam Amos-Binks

AAMOSBINKS@ARA.COM

Applied Research Associates, Raleigh, NC 27614 USA

Rogelio E. Cardona-Rivera

ROGELIO@CS.UTAH.EDU

School of Computing and the Entertainment Arts and Engineering Program
University of Utah, Salt Lake City, UT 84112 USA

Dustin Dannenhauer

DDANNENHAUER@NAVATEKLLC.COM

Navatek, Arlington, VA 22203 USA

Gene A. Brewer

GENE.BREWER@ASU.EDU

Department of Psychology, Arizona State University, Tempe, AZ 85287 USA

Abstract

Anticipatory thinking (AT) is a key cognitive process in a variety of key societal contexts. Despite the theoretical overlap between planning as a cognitive AT prospection modality and planning as a paradigm of artificial intelligence, there has been little formal work from the latter on how to model, support, or augment the former. This position paper argues that the automated planning community is uniquely poised to investigate this overlap by extending existing techniques to both (i) initiate (in users) and (ii) perform AT. Together, these approaches support our conviction that AT is an exciting frontier that can be readily tackled by convergence, synthesis, and expansion of existing research within automated planning.

1. Introduction

Anticipatory thinking (AT), the deliberate and divergent exploration of relevant possible futures, is a key cognitive process in emergency preparedness, intelligence analysis, and military planning Geden et al. (2019). Cognitive psychologists have long-assessed an individual's capacity for anticipatory thinking, but results are not encouraging Klein et al. (2007): there are cognitive barriers at the individual level and social barriers at the organizational level that preclude effective AT. While there exist methodologies to counteract these barriers – *e.g.* structured analytic techniques as discussed by Heuer (1999) – these have not been subject to rigorous scientific evaluation Iden et al. (2017), and their efficacy remains unknown.

Prospection (future-thinking) is a key component of AT and *planning* is one of its four modalities. We posit there exists a theoretical overlap between planning, the prospection modality, and planning, the artificial intelligence capability. Our position is that the automated planning commu-

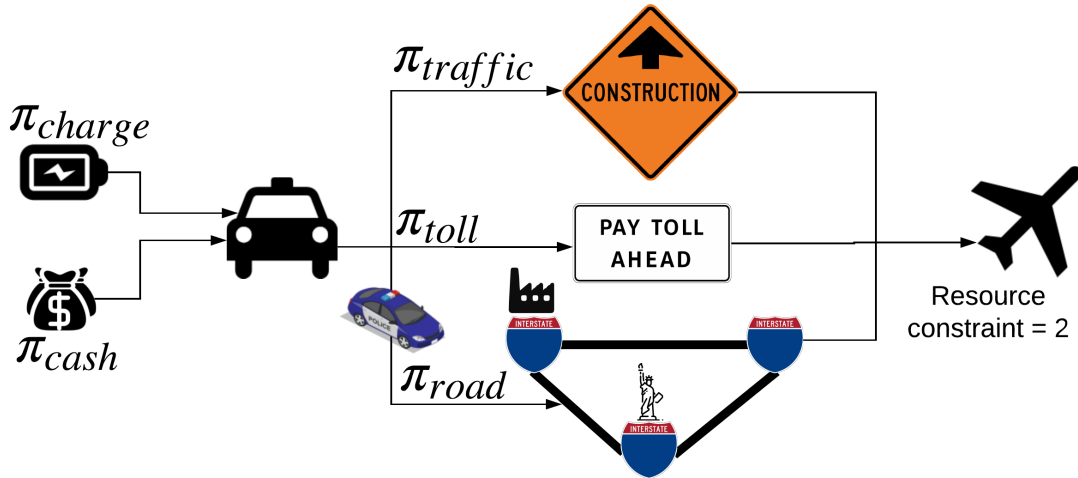


Figure 1. AT of a self-driving car, three plans to minimize travel time and arrive at the airport; plans π_{charge} and π_{cash} are generated as a result of automated (planning-based) AT.

nity is uniquely poised to investigate this theoretical overlap. We argue that our collective capacity for AT will be improved by extending existing automated planning paradigms to model, support, and augment AT. We focus on two case challenges. The first is initiating AT in non-domain expert users, who exhibit systematic failures of imagination that preclude their effective AT; current work on diverse planning and goal-recognition communities is well-positioned to tackle – but has not yet solved – this challenge. The second is performing AT as an automated process; existing work cannot yet generate what we characterize as an *AT Plan* in a principled manner. Thus, applying planning to AT is not straightforward: off-the-shelf solutions do not wholly solve AT problems. The convergence, synthesis, and expansion of existing work within automated planning has the potential to address (at least two) AT challenges; this paper elucidates how.

2. Motivating Example

To illustrate how planning initiates and performs AT we will build on a running example of a self-driving car. Suppose the car’s machine learning systems have learned, through many miles of driving, generally how to respond to road conditions. However, machine learning fails to learn infrequent but highly impact events, such as running out of gas in remote areas. These events remain, at present, elusive to learn in advance and prevent widespread adoption of self-driving cars Gilpin (2019). AT’s role is to identify and mitigate against these low-probability but high-impact events.

Consider the scenario in Figure 1. The self-driving car’s goal is to minimize travel time (to increase opportunities for revenue) and arrive at the airport in time for a passenger to catch their flight. Within the solution space, there are five plans. The plan with highest probability of success is π_{traffic} , but it requires driving the modestly-charged car through a construction zone where there are

no charge stations. Though very low probability, a traffic jam could occur, depleting the battery and making on-time arrival impossible. To ensure arriving on time even with a traffic jam, the car must be charged more in advance; this is plan π_{charge} . The fastest plan, π_{toll} , passes by a toll road that reduces travel time but is rarely open. To take advantage of the cash-only toll road, exact change must be acquired in advance; this is plan π_{cash} . A final plan, π_{road} , involves driving on a road network with access to other part sof the city where both travel time and probability of success in between those of π_{traffic} and π_{toll} , but with no potential reduced travel time or breakdown. Importantly, π_{traffic} and π_{toll} are viable options on their own and do not strictly require charging the car or obtaining exact change; π_{traffic} is desirable because it is the most likely to succeed while π_{toll} minimizes travel time. We expect that planning systems would normatively generate π_{traffic} or π_{toll} as solutions, but not π_{charge} and π_{cash} as they contain non-essential actions. In what follows, we discuss how we might generate and reason-with this space of plans.

3. Initiating Anticipatory Thinking

Human predictions of future events are subject to common biases in imagination; *e.g.* asking people to “imagine if you had a car accident last week” leads to more details, more realistic scenarios, and a less-stereotyped imagining Anderson & Dewhurst (2009) than asking people to “imagine getting into a car accident next week” Kane et al. (2008). Here, we are interested in using planners as decision-support tools that elicit more-numerous, more-detailed, and less-prototypical predictions of future events.

3.1 Divergent Possible Futures via Diverse Planning

The further in the future we ask people to predict events, the more stereotyped the predictions are D’Argembeau & Van der Linden (2004). We might imagine helping people escape the “functional fixedness” Duncker & Lees (1945) that results from people settling on the first reasonable idea by relying upon diverse planning techniques. Instead of finding a single solution that maximizes a single quality metric (typically cost), diverse planners generate a set of plans with some operational measure of diversity amongst the set (*e.g.* Nguyen et al., 2012). Modern approaches specify a desired k amount of diverse plans to generate in advance of searching for them (*e.g.* Sohrabi et al., 2016), where diversity of the plan-set is measured as the average pair-wise distance between plans (*cf.* Coman & Muñoz-Ávila, 2011).

In our example, the car might opt to pursue the highest success-rate plan π_{traffic} . A breakdown here would mean late arrival, resulting in the passenger missing their flight. A diverse planner might generate a plan-set of similar success-likelihood but that would not be so disastrous upon breakdown (*e.g.* π_{charge}). This diverse plan-set may then be presented to the passenger, to elicit their AT and override the car’s plan. However, at least 2 limitations preclude the immediate application of diverse planners to initiating AT.

First, plan-set generation methods cannot yet provide guarantees on solution space coverage. That is, asking for k -many diverse plans does not direct the search algorithm to *maximize* the diversity of the output plan-set with respect to k ; diversity is only analytically identified post-search. There may exist relaxed planning graph-based heuristics Bryce & Kambhampati (2007) that would

guide the search toward diverse regions of the solution space; the work by Amos-Binks *et al.* (2017) is a first domain-specific step.

Second, we need a more nuanced notion of plan-set diversity. Measuring diversity on the basis of *average* pair-wise distance may obscure plans that are outliers or “events in the long-tail,” precisely the sequences of events we care to characterize in AT Geden *et al.* (2019). A straightforward step toward solving this is to measure diversity on the basis of other descriptive statistics (*e.g.* inter-quartile range) and/or higher-order statistical moments (*e.g.* variance, skewness, kurtosis), but it is presently unclear how these might be useful for expanding the descriptiveness of plan-set diversity measures.

3.2 Deliberate Possible Futures via Intent Elicitation

People are conservative in creative tasks: McCaffrey & Spector (2012) found that in order to elicit more creative designs for candles, they had to specifically *ask* participants to change features that tend to not vary much (*e.g.* changing motion or sound) to avoid having participants rely on features that normatively vary a great deal (*e.g.* changing color or scent). We might imagine deliberately asking a human to imagine what features of an environment have to change in order to better predict the activities of others, a task similar to goal recognition design.

In plan recognition, an observer agent aims to predict the future behavior of an actor agent given a sequence of observed steps from the actor. The observer is external to the environment and assumes the actor is pursuing one goal from a finite set of possible goals and must predict: (a) which goal explains the actor’s actions and (b) which next actions the actor will take to achieve the goal. There are many techniques for solving this task (*cf.* Sukthankar *et al.*, 2014); of interest is a technique that depends on a *domain model* Ramírez & Geffner (2009). This technique assumes that the actor is behaving optimally with respect to plan cost and are behaving honestly Masters & Sardina (2019), thus transforming the actor’s observed steps into goals that a planning agent (the observer) must satisfy.

Within domain model-based plan recognition, goal recognition design (grd) is the task of designing a plan recognition problem that minimizes the *worst-case distinctiveness* (wcd). The wcd is the maximal number of actions an actor can take before they *must* take an action that reveals their intent to an observer.

In the original formulation of grd Keren *et al.* (2019), the wcd is minimized via edits to the set of all ground actions in the domain. In contrast, the *intent-elicitation problem* (iep) is a novel reformulation of grd that keeps the set of possible actions fixed and seeks to plan to achieve a state of the world that minimizes the wcd. Framed differently, within the iep, the observer is no longer external to the environment or reacting to an actor’s activity; instead, the observer is proactively deducing the actor’s intent (goal and corresponding plan) by establishing conditions in the world that prompt the actor to act in a way that reveals their intention.

To expand upon our example, imagine a police officer who suspects the self-driving car to be transporting a criminal. The officer knows there are many paths a criminal might take to continue their criminal activity, but (for the sake of this example) none of them involve the airport: how might the officer act upon their suspicion in a productive manner? We might apply an iep-solver to deliberately initiate possible-futures reasoning in the officer and thereby help them come up with

a plan. The officer has *not* observed the car charging nor obtaining the cash toll, so the solver can rule out π_{charge} and π_{cash} . It prompts the officer to set up a blockade on the π_{road} since pruning it from the solution space removes access to other parts of the city and leaving only two direct routes to the airport; π_{traffic} and π_{toll} . If the car chooses either of these options it indirectly demonstrates the passenger’s innocence but if the car chooses neither, they will remain under suspicion.

This problem formulation is readily afforded by combining several existing well-defined problems and techniques; future work should elucidate potential pitfalls and solutions.

3.3 Evaluation

As discussed, initiating AT is an instance of human-in-the-loop planning (HILP, Kambhampati & Talamadupula, 2015): a planning system accounts for a human operator, and the interaction elicits the deliberate and divergent exploration of relevant possible futures. There is a dearth of studies within HILP (historically, “mixed-initiative planning”) that assess the impact of augmenting human cognitive ability with automated planning capability. While identifying the best interface for such a loop is an open problem Amershi et al. (2019), we note that human-subjects evaluations will be key for discerning the utility of AT initiating techniques. We propose to investigate multi-modal physiological data (*e.g.* eye-tracking) as a way to assess a planning system’s efficacy in initiating AT. Specifically, de Rooij et al. (2018) have linked pupillary dilation to divergent thinking in naturalistic creativity tasks, Roca et al. (2011) have linked number, locations, and durations of fixations to anticipatory thinking and decision making, and Hintz et al. (2017) have linked anticipatory eye-movements to interactions between language and the visual world. Taken together, these diverse studies highlight that a variety of signals are present in eye-tracking data that can be used in real-time to assess AT.

4. Performing Anticipatory Thinking

We also find compelling the case of performing AT without reference to a human operator that might interpret the results. Several planning and goal reasoning paradigms take into account the uncertainties during plan execution and flexible goal achievement. Their respective outputs can be stitched together to reflect the *result* of performing AT, but do not search the space of (what we would characterize as) AT solution plans in a principled manner. These paradigms include: (a) contingent planning, which generates conditional plans for use in plan execution when the initial state or action effects are inconsistent with initial beliefs Hoffmann & Brafman (2005), (b) oversubscription planning Van Den Briel et al. (2004) has the capability to manage goal achievement by identifying a maximal subset of the defined goals, (c) agent-based preservation of maintenance goals while pursuing achievement goals Duff et al. (2006), (d) metacognitive planning, in which agents formulate their own goals and plans in response to their environment Cox et al. (2016), (e) learning-driven goal generation, which anticipates new goals from experience Pozanco et al. (2018), (f) agent-based modeling of future state belief for diagnosing plan failures and discrepancies in execution Muñoz-Avila et al. (2019), among others.

Current paradigms fail to address AT because they do not critically examine opportunities for events that can foil or advance the pursuit of goals in a manner independent over a sensing-and-

execution context (as exemplified by contingent planning). Further, approaches that mitigate for points of potential failure exhaustively reason over *all* such points, failing to reason over which failure points are worth reasoning over and why. In AT, the determination of *relevant* possible futures (in terms of potential risk or reward) is key; put simply, in AT, not all threat/opportunity points are relevant. Finally, in considering threat points, current paradigms explore *mitigation* strategies; in AT, we are interested in *preventative* or *preparatory* strategies.

4.1 AT Plans

Amos-Binks & Dannenhauer (2019) claim that AT is metacognitive process and proposes an approach to perform edits upon existing plans that introduce anticipatory actions, with no commitment on the planning approach that produced the input plan in the first place. Here, we go beyond that work by proposing a (preliminary) definition of an AT plan relative to a planning domain model \mathcal{D} , with the goal of spurring the development of planning systems that deliberately generate them.

Definition 1 (AT plan) A tuple, $\tau = \langle A, AA \rangle$, where A is a set of actions a_1, \dots, a_n and AA is a set of anticipatory actions taken to create an advantage.

In the example, τ could be either π_{cash} or π_{charge} , because these include anticipatory actions.

Definition 2 (Anticipatory action) A tuple $\langle AAS, CE \rangle$, where AAS is an action sequence, aa_1, \dots, aa_n added to an AT plan τ such that at least one effect reduces the impact of a conditioning event $CE \in CE(\mathcal{D})$.

Definition 3 (Conditioning event) an action in the domain model \mathcal{D} executed by an agent or occurring in environment that prunes desirable areas of the solution space. The set of all such events is denoted by $CE(\mathcal{D})$,

In our example in Figure 1, the car has the option to charge. If a traffic jam occurs in the construction zone, the car will have enough energy to sustain the delay and arrive on-time at the airport. Generating a plan with the *drive* and *charge* anticipatory actions for conditioning event *jam* is an anticipatory plan with actions $AA(\tau) = \{\{drive, charge\}, jam\}$.

Of course there are many conditioning events to account for in a domain and preparing for them all is not realistic in practical sense. Budgets rarely account for preparing for all possible catastrophes, and also could be intractable computation wise. To this end, we introduce a resource constraint that ensures the anticipatory actions do not cross an allocated threshold of resources. Individual planners may decide if it is the number of conditioning events they wish to act upon (e.g. oversubscription planning Van Den Briel et al. (2004)) or the total impact of a few events.

Definition 4 (Resource constraint) A resource constraint RC is an additional goal condition that provides a planner a balance of anticipatory actions and efficiency.

Figure 1 specifies the resource constraint at two. Assuming uniform action cost, this means solution plans must contain two or less anticipatory actions to prepare for one of the traffic jam or toll road events.

Table 1. Support for AT solution plan desiderata within IPC-2018’s probabilistic track domains; ● is full, ◐ is partial, and ○ is no support for Conditioning Events (CE), Anticipatory Actions (AA), and Resource Constraints (RC).

IPC-2018 Domain	CE?	AA?	RC?
Academic advising	○	○	○
Chromatic dice	○	○	○
Cooperative recon	●	◐	●
Earth observation	◐	○	○
Manufacturer	○	○	○
Push your luck	○	○	○
Red-finned blue-eye	◐	○	○
Wildlife preserve	◐	○	○

4.2 Evaluation

The maturity of the planning community is perhaps best epitomized by the international planning competition (IPC). Benchmark domains across a number of tracks measure the continued progress of state-of-the-art planners. Of particular interest to our position is the probabilistic track, where partially-observed conditioning events (such as traffic jams or open toll roads) are supported in RDDDL and require consideration by a planner.

Table 1 contains an analysis of the IPC-2018 domains and the ability of a planner to generate AT plans from planning problems within them. Both the *Academic advising* and *Manufacturer* both require traversing a directed acyclic graph (DAG). While there are some strategic ways to traverse the DAG, the DAG is static and not subject to conditioning events such as a curriculum change. The *Push your luck* and *Chromatic dice* are similar in that solutions are both dependent on a single random variable. Again, the lack of conditioning events affords no opportunity to anticipate. *Earth observation*, *Red-finned*, *blue-eyed*, and *Wildlife preserve* all have changing environmental conditions (weather, water level, and poachers, respectively) that could qualify as conditioning events but are essential to the progression of the domain. These three domains also have no actions to prepare for these events. Lastly, *Cooperative recon* has conditioning events by way of probabilities in cells that damage sensing equipment. While no anticipatory actions can be taken in advance, travelling or co-operating with other rovers for redundant equipment prepares for hazards that may emerge. Resource constraints can be represented as action costs.

Our qualitative analysis leads us to believe that many of the probabilistic benchmark domains, save perhaps *Cooperative recon*, lack the representation to support planning problems for which AT plans are solutions. However, this can be overcome by extending the domains to include both conditioning events that are not essential to the progression of the domain and actions that reduce the impact or take advantage of conditioning events.

We now turn to assessing the quality of an AT plan so we might measure how effective a planner is at generating them. Amos-Binks and Dannenhauer 2019 measure how successful an AT plan is with the product of i) the ratio of conditioning events identified to mitigated ones and ii) the ratio of

action costs to reduced conditioning event impact. This measure is an important first step, however, it neglects two factors. First, it only considers negative valence conditioning events. In our example, it would only be a relevant calculation for π_{charge} whereas conditioning events can also have a positive valence, such as the toll road in π_{cash} . Second, it does not consider the resource constraints allocated. Searching to find a set of high impact conditioning events with low cost mitigations is a behavior that merits rewarding. Addressing these two challenges enables meaningful comparisons between planners for their ability to perform AT.

5. Conclusion

Learning without actually experiencing is a hallmark of human intelligence. AT is cognitive process that exemplifies this hallmark and is in widespread use across many real-world domains. One of its defining features is its reliance on symbolic reasoning to identify possible futures, not just likely ones, and take action to create an advantage. Our position is that a theoretical overlap exists between the planning modality of prospection, one of the underlying mechanisms of AT, and the automated planning research community.

A fledgling research community has only just begun to measure and investigate the phenomena of AT, focusing on assessing an individual's AT Geden et al. (2019), methodologies that initiate it Heuer (1999); Iden et al. (2017), and the role of expertise in performing it Klein et al. (2011). This focus on individual AT limits widespread adoption at an organizational level that is especially problematic for organizations where AT is germane to its core function. As a result, AT remains confined to a small group of narrow domain and methodology experts. Addressing the challenges we have articulated in this position paper is tractable and can overcome this limitation, scaling AT to an organizational level and contributing to one of AI's original goals, the replication of human intelligence, along the way.

References

- Amershi, S., et al. (2019). Guidelines for human-ai interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*.
- Amos-Binks, A., & Dannenhauer, D. (2019). Anticipatory thinking: A metacognitive capability. *Working Notes of the AAAI Fall Symposium on Cognitive Systems for Anticipatory Thinking*.
- Amos-Binks, A., Potts, C., & Young, R. M. (2017). Planning graphs for efficient generation of desirable narrative trajectories. *Proceedings of the 10th Workshop on Intelligent Narrative Technologies at the 13th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*.
- Anderson, R. J., & Dewhurst, S. A. (2009). Remembering the past and imagining the future: Differences in event specificity of spontaneously generated thought. *Memory, 17*, 367–373.
- Bryce, D., & Kambhampati, S. (2007). A tutorial on planning graph based reachability heuristics. *AI Magazine, 28*, 47–47.
- Coman, A., & Muñoz-Ávila, H. (2011). Generating Diverse Plans Using Quantitative and Qualitative Plan Distance Metrics. *Proceedings of the 25th AAAI Conference on Artificial Intelligence* (pp. 946–951).

- Cox, M. T., Alavi, Z., Dannenhauer, D., Eyorokon, V., Munoz-Avila, H., & Perlis, D. (2016). Midca: A metacognitive, integrated dual-cycle architecture for self-regulated autonomy. *Proceedings of the 30th AAAI Conference on Artificial Intelligence*.
- D'Argembeau, A., & Van der Linden, M. (2004). Phenomenal characteristics associated with projecting oneself back into the past and forward into the future: Influence of valence and temporal distance. *Consciousness and cognition*, *13*, 844–858.
- Duff, S., Harland, J., & Thangarajah, J. (2006). On proactivity and maintenance goals. *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems* (pp. 1033–1040).
- Duncker, K., & Lees, L. S. (1945). On problem-solving. *Psychological monographs*, *58*, i–113.
- Geden, M., Smith, A., Campbell, J., Spain, R., Amos-Binks, A., Mott, B., Feng, J., & Lester, J. (2019). Construction and Validation of an Anticipatory Thinking Assessment. *PsyArxiv:9reby*.
- Gilpin, L. (2019). Explaining possible futures for robust autonomous decision-making. *Working Notes of the AAAI Fall Symposium on Cognitive Systems for Anticipatory Thinking*.
- Heuer, R. J. (1999). *Psychology of intelligence analysis*. Center for the Study of Intelligence, Central Intelligence Agency.
- Hintz, F., Meyer, A. S., & Huettig, F. (2017). Predictors of verb-mediated anticipatory eye movements in the visual world. *Journal of Experimental Psychology: Learning Memory and Cognition*, *43*, 1352–1374.
- Hoffmann, J., & Brafman, R. (2005). Contingent planning via heuristic forward search with implicit belief states. *Proceedings of the 15th International Conference on Automated Planning and Scheduling*.
- Iden, J., Methlie, L. B., & Christensen, G. E. (2017). The nature of strategic foresight research: A systematic literature review. *Technological Forecasting and Social Change*, *116*, 87–97.
- Kambhampati, S., & Talamadupula, K. (2015). Human-in-the-loop planning and decision support. AAAI Tutorial.
- Kane, J., McGraw, A. P., & Van Boven, L. (2008). Temporally asymmetric constraints on mental simulation: Retrospection is more constrained than prospection. *The handbook of imagination and mental simulation*, (pp. 131–149).
- Keren, S., Gal, A., & Karpas, E. (2019). Goal recognition design in deterministic environments. *Journal of Artificial Intelligence Research*, *65*, 209–269.
- Klein, G., Snowden, D., & Chew, L. P. (2007). Anticipatory thinking. *Proceedings of the 8th International Conference on Naturalistic Decision Making*.
- Klein, G., Snowden, D., & Chew, L. P. (2011). Anticipatory thinking. *Informed by knowledge: Expert performance in complex situations*, (pp. 235–245).
- Masters, P., & Sardina, S. (2019). Goal recognition for rational and irrational agents. *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems* (pp. 440–448).
- McCaffrey, T., & Spector, L. (2012). Behind every innovative solution lies an obscure feature. *Knowledge Management & E-Learning: An International Journal*, *4*, 146–156.

- Muñoz-Avila, H., Dannenhauer, D., & Reifsnnyder, N. (2019). Is everything going according to plan? expectations in goal reasoning agents. *Proceedings of the AAAI Conference on Artificial Intelligence* (pp. 9823–9829).
- Nguyen, T. A., Do, M., Gerevini, A. E., Serina, I., Srivastava, B., & Kambhampati, S. (2012). Generating Diverse Plans to Handle Unknown and Partially Known User Preferences. *Artificial Intelligence*, *190*, 1–31. From <http://dx.doi.org/10.1016/j.artint.2012.05.005>.
- Pozanco, A., Fernández, S., & Borrajo, D. (2018). Learning-driven goal generation. *AI Communications*, *31*, 137–150.
- Ramírez, M., & Geffner, H. (2009). Plan recognition as planning. *Proceedings of the 21st International Joint Conference on Artificial Intelligence*.
- Roca, A., Ford, P. R., McRobert, A. P., & Williams, A. M. (2011). Identifying the processes underpinning anticipation and decision-making in a dynamic time-constrained task. *Cognitive Processing*, *12*, 301–310.
- de Rooij, A., Vromans, R. D., & Dekker, M. (2018). Noradrenergic modulation of creativity: evidence from pupillometry. *Creativity Research Journal*, *30*, 339–351.
- Sohrabi, S., Riabov, A. V., Udrea, O., & Hassanzadeh, O. (2016). Finding diverse high-quality plans for hypothesis generation. *Proceedings of the 22nd European Conference on Artificial Intelligence* (pp. 1581–1582).
- Sukthankar, G., Geib, C., Bui, H. H., Pynadath, D., & Goldman, R. P. (2014). *Plan, Activity, and Intent Recognition: Theory and Practice*. Elsevier.
- Van Den Briel, M., Sanchez, R., Do, M. B., & Kambhampati, S. (2004). Effective approaches for partial satisfaction (over-subscription) planning. *AAAI* (pp. 562–569).