

a plan’s goal. For instance, after the goal of recharging a robot is fulfilled, it can act for a longer time. Therefore, the Opportunity Motivator’s fitness biases goal formulation toward goals that have the most actions available per expected state, and leave the agent with the most resources and actions available when the goal is achieved. This function is defined as:

$$f_{\text{opportunity}}(X) = \frac{([\sum_{j=0}^{n-1} N(x_{c+j})] + [w * N(x_{c+n})])}{(n + w)N(s_c) + L(x_{c+n}) - L(s_c) - 1}$$

where $w \geq 1$.

5. Empirical Study

5.1 Rover Domain

Rovers-With-Compass (RWC) is a deterministic navigation domain with hidden obstacles inspired by the difficulties encountered by the Mars Rovers. In this domain, individual locations may be windy, sandy, and/or contain sand pits, but the agent cannot observe these obstacles directly. Sandy locations cause the rover to be covered in sand. While covered in sand, the rover cannot observe its location or recharge its batteries. Sand pits stop the rover from moving, but the rover can dig itself out at a high energy cost. Windy locations clear the sand off of the rover, but due to a malfunction, may confuse the rover’s compass, causing it to move in the wrong direction. Rovers in this domain have a finite amount of energy that is consumed by actions; if a rover runs out of energy, it becomes unable to move. Success in this domain is evaluated based on the agent’s ability to achieve a set of three separate navigation goals using different rovers.

5.2 Experimental Design

Our primary claim is that M-ARTUE should perform comparably to ARTUE, as measured by the percentage of goals successfully achieved, despite its smaller amount of task-specific knowledge. As a secondary claim, we intend to show that both the Opportunity and Exploration Motivators improve the performance of M-ARTUE. In order to measure goal achievement performance, we tested ARTUE and M-ARTUE using three sets of 25 randomly-generated scenarios within the RWC domain. Each scenario consisted of a grid of 36 locations; each location had a fixed probability of being windy, being sandy, being sunny, and containing a sand pit. These probabilities varied across three experimental conditions and correspond to the “danger” of the domain. M-ARTUE was responsible for controlling three rovers in each scenario, and each rover had its own goal location, but the agent could exploit knowledge gained using one rover in planning for another. M-ARTUE’s objective (and social motivation) was to get each rover to its goal location in the scenario, and performance was measured as a percentage of goal locations reached. Each scenario was tested with both M-

ARTUE and ARTUE, which used the same HTN definitions for goal planning. A limit of 80 actions was imposed on both agents to ensure the timely conclusion of all tests. (Allowing more actions could only improve M-ARTUE’s relative performance, as ARTUE completed every scenario in fewer than 80 actions.)

M-ARTUE chose goals using the heuristic-guided goal formulation process described in Section 4. Available goals considered by M-ARTUE included recharging, navigation, and recovery (i.e., removing a rover from a sand-pit); successful completion of scenarios required that all types of goals be used in different situations. ARTUE’s formulation rules asserted recharge goals for partially-discharged rovers, and navigation goals for rovers not at their targets. The priority of the recharge goal varied based on the amount of remaining charge. ARTUE always selected the highest priority goal for which it could find a plan. Both agents planned with SHOP2, using a mapping from goals into tasks and standard hierarchical task decomposition using manually designed hierarchical task methods. In this experiment, C_{social} was set to 1.1, giving a modest growth rate for social urgency. $C_{\text{social-fitness}}$ was set to 5, slightly larger than the average plan length, to keep social fitness largely within the range of $[0,1]$. $C_{\text{exploration}}$ was set to 5, to make exploration most important during the first 5 actions. Finally, w was set to 20, to ensure final states outweighed the rest of longer-than-average plans.

5.3 Results

In our first test, the probability of obstacles was set to 10%, so there was a 10% probability of windy conditions, sandy conditions, and a pit at each location. With 50% likelihood, each location is sunny, meaning the rover can recharge there. Our second test used probabilities of 20% and 40%, respectively, and our third test used 30% and 30%. As a result, successive tests were more dangerous.

Figure 2 compares the performance of M-ARTUE to ARTUE for each of our three test conditions. In the first test, M-ARTUE actually outperformed ARTUE, although the difference is not significant ($p=.07$). For the second and third tests, M-ARTUE and ARTUE did not perform significantly differently, supporting our claim that domain-independent heuristics can perform at the same level as engineered rules for goal formulation.

In Figure 3, we compare the performance of two ablations in our three test conditions. As expected, performance suffers without the Opportunity Motivator. In each test condition, the full M-ARTUE significantly outperforms M-ARTUE without the Opportunity Motivator ($p<.05$). The same is not true for the Exploration Motivator. While there is a significant benefit in the first test ($p=.03$), this decreases as obstacles become more frequent, which matches our intuition that exploration is less useful in more dangerous conditions, supporting our claim that both motivators contribute to superior performance, at least some of the time.

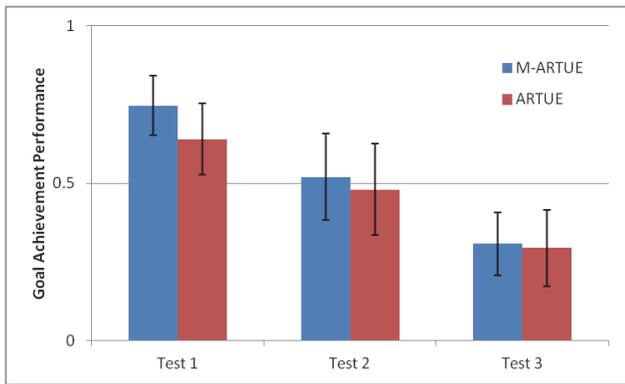


Figure 2: Comparing M-ARTUE vs. ARTUE

6. Conclusions and Future Work

While our experiments support our claim that domain-independent heuristics can, for some tasks and domains, replace hand-coded knowledge in goal formulation, much work remains to be done. First, it's important to identify the generality of these techniques in applications to other domains. Second, our experiments required tuning constants in the motivator functions (see Section 4.1). This use of domain-specific knowledge is undesirable and we plan to instead automatically tune them in our future work. Third, we intend to address the higher overhead of planning for all goals on every GDA cycle, possibly by filtering goals that can be identified *a priori* as non-contributory. Fourth, resources are presently identified as part of the domain description, but could be algorithmically discovered in future work. Finally, we plan to replace the Social Motivator with an interactive system that allows M-ARTUE to learn goal formulation knowledge, as discussed in prior work (Powell *et al.* 2011), to permit long-lived agents to adapt their formulation policies over time.

Acknowledgements

Thanks to ONR 32 for their support of this research.

References

- Bratman, M.E. (1987). *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- Choi, D. (2011). Reactive goal management in a cognitive architecture. *Cognitive Systems Research*, **12**, 293-308.
- Coddington, A.M. (2006). Motivations for MADbot: a motivated and goal directed robot. *Proceedings of the Twenty-Fifth Workshop of the UK Planning and Scheduling Special Interest Group* (pp. 39-46). Nottingham, UK: University of Nottingham.
- Cox, M.T., & Veloso, M.M. (1998). Goal transformations in continuous planning. In M. desJardins (Ed.), *Proceedings of the Fall Symposium on Distributed Continual Planning* (pp. 23-30). Menlo Park, CA: AAAI Press.
- Ghallab, M., Nau, D., & Traverso, P. (2004). *Automated Planning: Theory and Practice*. San Francisco, CA: Morgan Kaufmann Publishers.

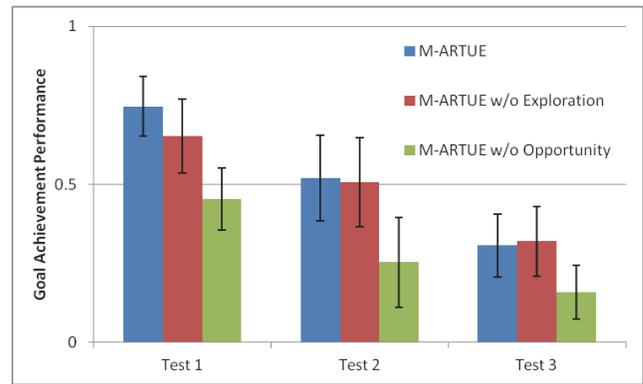


Figure 3: Comparing M-ARTUE vs. ablations

Hawes, N. (2011). A survey of motivation frameworks for intelligent systems. *Artificial Intelligence*, **175**(5-6), 1020-1036.

Hawes, N., Hanheide, M., Hargreaves, J., Page, B., Zender, H., & Jensfelt, P. (2011). Home alone: Autonomous extension and correction of spatial representations. In *ICRA* (pp. 3907-3914). Shanghai, China: IEEE Press.

Jaidee, U., Munoz-Avila, H., & Aha, D.W. (2011). Integrated learning for goal-driven autonomy. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence*. Barcelona, Spain: AAAI Press.

Klenk, M., Molineaux, M., & Aha, D.W. (2012). Goal-driven autonomy for responding to unexpected events in strategy simulations. To appear in *Computational Intelligence*.

Molineaux, M., Klenk, M., & Aha, D.W. (2010a). Goal-driven autonomy in a Navy strategy simulation. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*. Atlanta, GA: AAAI Press.

Molineaux, M., Klenk, M., & Aha, D.W. (2010b). Planning in Dynamic Environments: Extending HTNs with nonlinear continuous effects. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*. Atlanta, GA: AAAI Press.

Molineaux, M., Kuter, U., & Klenk, M. (2012). DiscoverHistory: understanding the past in planning and execution. In *Proceedings of the Eleventh International Conference on Autonomous Agents and Multi-Agent Systems*. Valencia, Spain: ACM Press.

Nau, D.S. (2007). Current trends in automated planning. *AI Magazine*, **28**(4), 43-58.

Nau, D., Au, T.-C., Ilghami, O., Kuter, U., Murdock, J.W., Wu, D., & Yaman, F. (2003). SHOP2: An HTN planning system. *Journal of Artificial Intelligence Research*, **20**, 379-404.

Powell, J., Molineaux, M., & Aha, D.W. (2011). Active and interactive learning of goal selection knowledge. In *Proceedings of the Twenty-Fourth Florida Artificial Intelligence Research Society Conference*. West Palm Beach, FL: AAAI Press.

Sun, R. (2007). The motivational and metacognitive control in CLARION. In W. Gray (Ed.) *Modeling integrated cognitive systems* (pp. 63-75). New York: Oxford University Press.

Sutton, R.S., & Barto, A.G. (1998). *Introduction to Reinforcement Learning*. Cambridge, MA: MIT Press.

Weber, B., Mateas, M., & Jhala, A. (2012). Learning from demonstration for goal-driven autonomy. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*. Toronto, Canada: AAAI Press.