On Teammate-Pattern-Aware Autonomy

JAAMAS Track

Edmund H. Dufee
University of Michigan
Ann Arbor, MI, USA
durfee@umich.edu

Abhishek Thakur
BRINC Drones
Las Vegas, NV, USA
thakabhi@umich.edu

Eli Goldweber
University of Michigan
Ann Arbor, MI, USA
edgoldwe@umich.edu

ABSTRACT

We describe an approach for constraining robot autonomy based on the robot's awareness of *patterns* of its human teammates' behaviors, rather than either ignoring its teammates (which is fast but dangerous) or inferring their plans (which is safer but slow). We evaluate this approach in a series of simulated problems where an unmanned ground vehicle and its human teammates must rapidly respond to a sudden context shift, and identify conditions that should be (purposely) met such that a pattern-aware approach is particularly effective compared to the alternatives.

KEYWORDS

Human-Robot Coordination, Autonomy, Organizational Design

ACM Reference Format:

Edmund H. Dufee, Abhishek Thakur, and Eli Goldweber. 2021. On Teammate-Pattern-Aware Autonomy: JAAMAS Track. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 3 pages.

1 MOTIVATION

We are concerned with mixed manned-unmanned teams (specifically involving unmanned ground vehicles (UGVs)) undergoing sudden context shifts, such as when a natural disaster strikes or an adversary attacks. Under nominal mission conditions, a human operator can closely oversee a UGV, but during sudden context shifts the operator must often focus on personally-important tasks like self preservation, and thus at such times can least afford to provide guidance/supervision to the UGV. As a consequence, it is at precisely these times of change and uncertainty that the UGV should shoulder greater autonomy for controlling its own responses. The danger, though, is that a UGV's autonomous behavior, based on algorithmic assessments balancing what it knows and has learned about the risks and rewards of different courses of action in its perceived environment, might deviate from its human teammates' immediate expectations, compounding their confusion, and threatening team goals and even teammate safety [5]. Hence, robotic agents that can flexibly come into contact with human teammates [7] may need to trade away some degree of task-execution optimality in order to satisfy teammates' preferences [6]. This problem, of how a UGV can rapidly exercise its autonomy in ways that support, not confound, its human teammates, is the focus of the *JAAMAS* article [4] that we summarize here.

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online.

2 CONCEPTUAL CONTRIBUTION

We characterize *teammate-aware autonomy* as when a UGV's autonomous behaviors are informed and constrained by awareness of its teammates' plans and expectations. Teammate-aware autonomy can be (and has been) realized in various ways. For example, in response to a sudden context shift, the UGV could be **treammate-runtime-aware** by observing all it can about the teamwide-situation, infer what its teammates will do, and plan its actions around theirs; however, all the sensing, inference, and planning could catastrophically delay its own responses. The UGV could instead approximate the teammate-aware approach with **teammate-*un*aware** planning (ignoring its teammates), followed by runtime sensing to avoid negative interactions with them (e.g., halting before a collision); however, as our empirical results confirm, reactively-repaired teammate-unaware plans can be quite inefficient. To speed up its response while still planning in a teammate-aware way, the UGV could try to pre-plan for every conceivable teamwide-situation offline, and then just retrieve the right plan at runtime; however, the space of possible situations to examine and the memory required to store all of the contingent responses make such an approach infeasible in interesting settings.

Our conceptual contribution is a new **teammate-pattern-aware autonomy** approach that combines aspects of the preceding offline and runtime planning approaches. Our approach's offline component samples many possible team situations, and for each predicts how the teammates will respond so as to contribute best to coordinated team behavior. It then uses the results across the sampled situations to discern *patterns* in the behaviors of teammates, using techniques inspired by research into automating the organizational self-design process [1–3, 8–10] When computing the UGV's autonomous response to a sudden context shift, the approach's online component incorporates the teammates' abstract behavior patterns into the UGV's local planning process, leading quickly to local plans that account for teammates' abstract plans and behaviors, fulfilling expected roles.

3 ILLUSTRATION

We illustrate this process using the simplified grid environment in Figure 1, where the UGV and its 10 human teammates are sweeping northward, side-by-side, with the UGV at the west end of the line, when suddenly an attack comes from the northeast. The standard operating procedure for this situation is for the humans to move to safety at the west end, while the UGV should move eastward to a target position from which it can observe the enemy. If the UGV is unaware of its teammates and blindly moves eastward, however, it can collide with and injure them. For safety, it could reactively stop when encountering a teammate (humans can safely

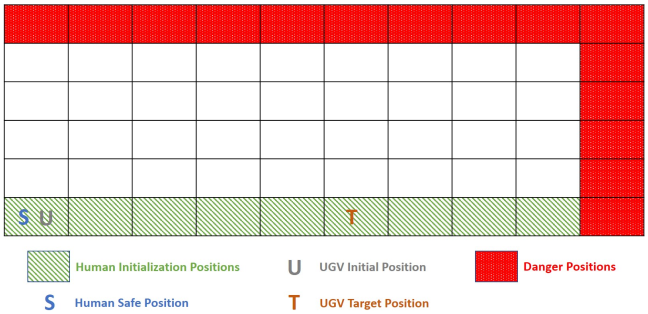

Figure 1: Environment with Constrained Movements

move around a *stationary* UGV), but repeatedly pausing for each of the 10 teammates delays it in reaching its target location. A teammate-runtime-aware UGV instead first collects information about the positions of all 10 teammates, then infers all of their movement plans, and finally optimizes its own plan around theirs. Unfortunately, all the collecting and inferring takes time, delaying when the UGV begins executing its plan.

Our teammate-pattern-aware UGV strikes a different balance. Offline (before the mission), it samples from possible runtime placements of its teammates and infers their plans. It then finds patterns over these plans, identifying the most likely places teammates might be at various times. In this simple scenario, for example, it concludes that teammates become increasingly concentrated in the western cells of this row. At runtime, then, the UGV uses this information when planning its own path: in this case, it takes a longer distance path by moving northward (to get out of the flow of westbound teammates), then eastward, and then southward into its target location after all of the teammates would have gone past. Even though the UGV travels farther, in the JAAMAS paper we show empirically that it arrives at its target location sooner because it never needs to pause to avoid collisions, and it doesn't spend time inferring the situation-specific plans of its 10 teammates.

4 FEATURES AFFECTING PERFORMANCE

This illustrative setting was particularly apropos for the teammate-pattern-aware approach for several reasons, including (1) the pattern was restricted (rational teammates could only be in a relatively small number of places over time), (2) there were many teammates (so someone probably *would* be in a place where someone *could* be), and (3) a plan that could entirely avoid teammates involved only a slightly longer route to the target location. We thus empirically tested our pattern-aware approach to better understand under which conditions it would be expected to perform well or poorly.

Effect of population size. For the setting in Figure 1, our teammate-pattern-aware approach devises the same (longer) path regardless of the number of teammates, and thus can be outperformed by other approaches when there are few teammates. With 1 teammate, for example, runtime planning is cheap and the best plan is to directly move to the target location with a pause as the teammate goes by.

Sensitivity to the accuracy of teammate model. Both the runtime-aware and our pattern-aware approaches depend on inferring the teammates' plans. As one would expect, as the teammate model becomes less accurate, the less useful the inferences are (whether at runtime or offline) and the more relatively effective a teammate-unaware approach becomes.

Sensitivity to the accuracy of initial conditions. Our approach finds patterns based on sampling from the space of possible problems. Unsurprisingly, if its samples are not representative (for example, its probability distributions over where teammates might be are wrong), then the resulting patterns will be less useful. The runtime-aware and unaware approaches are not affected by this form of inaccuracy.

Sensitivity to availability of (good) alternative plans. Awareness of its teammates is most helpful when the UGV has more choices over possible plans/actions. For example, if the UGV in Figure 1 were limited to only moving east-west, then its pattern-aware behavior is simply to wait until even the easternmost teammates would reach its position before it itself starts moving eastward. As with its roots in organizational self-design, our experiments show that a pattern-aware approach works best when one agent (the UGV) can act (adopt a role) so as to complement the patterns (roles) of its teammates.

5 IMPROVING PATTERNS

Among the conclusions from our empirical studies is that pattern-aware autonomy is the approach of choice when teammates behaviors are more bounded (patterns are "tighter") and separable from those of the UGV. At first glance, it might seem that such conditions are rare. However, a crucial observation is that the human teammates benefit from the UGV's success, and thus have an incentive to themselves exhibit activity patterns that are helpful to the UGV. For example, in simulated gridworld settings, a teammate's diagonal path might be randomly chosen from several equally-good (Manhattan distance) alternatives, leading to a diffuse pattern. But if its teammates are purposely biasing their motions (for example, completing all north-south moves before east-west moves), the UGV can discover and exploit the resulting tighter pattern.

We have examined the promise of bias that is exogenously provided (like above) and that also can emerge more organically (such as a teammate preferring its shortest path that visits the fewest previously-unvisited locations). We have shown that biasing indeed can pay great dividends, and that emergent biases can perform competitively with hand-crafted biases. Taking a step back, moreover, our results suggest potential lessons for the broader organizational self-design community. Specifically, instead of organizational roles arising based on patterns of actions that agents are independently performing, agents that *know* that their actions are being scrutinized for patterns might purposely behave so as to make such patterns easier to discover and/or more helpful to use.

ACKNOWLEDGMENTS

Dr. Jonathon Smereka helped formulate this problem, and our implementation repurposed some code written by Dr. Jason Sleight. This work was supported, in part, by the Automotive Research Center (ARC), via Cooperative Agreements W56HZV-14-2-0001 (U.S. Army TARDEC) and W56HZV-19-2-0001 (U.S. Army GVSC).

REFERENCES

[1] Gideon Dadik Bibu. 2019. Modelling Organisations As Electronic Institutions. *International Journal of Scientific and Technology Research* 8, 6 (2019), 37–43.

[2] Daniel D. Corkill. 1983. *A Framework for Organizational Self-Design in Distributed Problem-Solving Networks*. Ph.D. Dissertation. University of Massachusetts.

[3] Stephen Cranefield, Samhar Mahmoud, Julian A. Padget, and Ana Paula Rocha (Eds.). 2017. *Coordination, Organizations, Institutions, and Norms in Agent Systems XII - COIN 2016 Revised Selected Papers*. Lecture Notes in Computer Science, Vol. 10315. Springer.

[4] Edmund H. Durfee, Abhishek Thakur, and Eli Goldweber. 2020. Teammate-pattern-aware autonomy based on organizational self-design principles. *Autonomous Agents and Multi-Agent Systems* 34, 39 (2020), 1–30.

[5] Victoria Groom and Clifford Nass. 2007. Can robots be teammates? Benchmarks in human-robot teams. *Interaction Studies* 8, 3 (2007), 483–500.

[6] Bradley Hayes and Brian Scassellati. 2013. Challenges in Share-Environment Human-Robot Collaboration. In *Proceedings of the Collaborative Manipulation Workshop at the 2013 Human-Robot Interaction Conference*.

[7] Vladimir Murashov, Frank Hearl, and John Howard. 2016. Working Safely with Robot Workers: Recommendations for the New Workplace. *Journal of Occupational and Environmental Hygiene* 13, 3 (2016), 61–71.

[8] Mark Sims, Daniel D. Corkill, and Victor R. Lesser. 2008. Automated organization design for multi-agent systems. *Autonomous Agents and Multi-Agent Systems* 16, 2 (2008), 151–185.

[9] J. Sleight and E. H. Durfee. 2013. Organizational design principles and techniques for decision-theoretic agents. In *International Conference on Autonomous Agents and Multiagent Systems, (AAMAS)*. 463–470.

[10] J. Sleight and E. H. Durfee. 2015. Effective Influence Abstractions for Organizational Design. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 1267–1274.