

# Competence-Aware Systems for Long-Term Autonomy

Doctoral Consortium

Connor Basich

University of Massachusetts Amherst  
Amherst, Massachusetts

## ABSTRACT

Recent years have seen a push towards deploying fully autonomous robots in large, complex domains such as autonomous driving, space exploration, and service robots. However, legal, ethical, or technical constraints have limited the extent of these systems' employable autonomy. In order to successfully achieve their intended goals, these systems must utilize assistance from humans to compensate for their limitations. For such systems to be successful over the course of a long-term deployment, they must both be cognizant of their own competence and have the ability to improve this competence over time in a safe way. Motivated by practical concerns faced in industry, this thesis provides a formal model for such a human-agent system to reason about its own competence and aims in future work to provide effective ways of safely improving the competence of the system over the course of its deployment.

## KEYWORDS

Groups of humans and agents; Single- and multi-agent planning and scheduling; Learning agent capabilities; Long-term (or lifelong) autonomy for robotic systems

### ACM Reference Format:

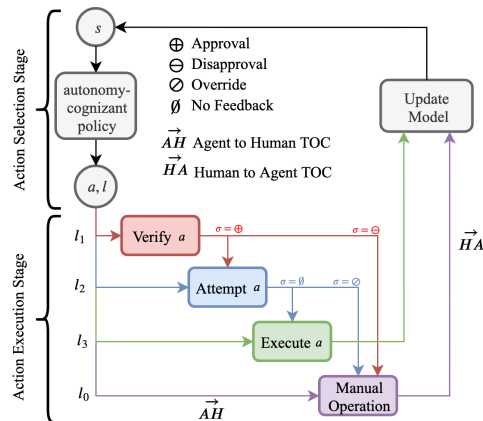
Connor Basich. 2020. Competence-Aware Systems for Long-Term Autonomy. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), Auckland, New Zealand, May 2020, IFAAMAS, 2 pages.

## 1 INTRODUCTION

Recent progress in artificial intelligence and robotics has driven the deployment of increasingly autonomous systems in complex unstructured domains. Such domains include space exploration [4], service robots [3], and, most notably, autonomous vehicles [2]. However, the vast majority of these systems are better characterized as *semi-autonomous systems* (SAS) that can operate autonomously under certain conditions, but may require human intervention or aid in order to achieve their assigned goals [5].

For example, an autonomous vehicle may transfer control to a human when its environment becomes too complex or when lane demarcations are lost. A space exploration robot may be required to power down and wait for explicit commands from the command center when an unexpected obstacle is encountered. This reliance on humans is driven by technical or model-based limitations that restrict the extent of autonomy that the agent can, or should, execute in some situations.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 2020, Auckland, New Zealand. © 2020 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.



**Figure 1: An illustration of a competence-aware system with four levels of autonomy –  $l_0, l_1, l_2, l_3$  – and four type of feedback signals – approval, disapproval, override, and no feedback. The autonomy-cognizant policy can only return levels of autonomy that are allowed for that action.**

More formally, these human interventions are driven by a limited *competence* as defined in our recent work [1], a notion that captures both the ability and robustness of the autonomous agent, as well as factors (legal, ethical, or trust-based) that influence the human's perception of the agent's capabilities but may not be available to the agent itself.

In such complex domains, enumeration of all possible scenarios, and inclusion of all possible environmental features in the agent's model, is intractable to implement and hence impractical to assume. As a result, autonomous systems cannot be expected to perform optimally in all situations encountered, particularly when deployed over months or years.

Consequently, this thesis aims to (1) provide a means for the system to model its own competence in any situation to optimally rely on human assistance and (2) develop effective techniques for competence-aware systems to *safely* improve their competence over the course of their deployment using available human assistance.

## 2 COMPLETED WORK

In an open environment, complete *a priori* determination of the competence of a human-agent system, and the respective capabilities of each actor in the system, is generally either impractical or impossible. As a result, the system's initial policy is likely to be over-reliant or under-reliant on the human in many circumstances; without an ability to adjust the autonomy over time in a *safe* way, the chance of failure and the resources wasted will grow over time. In expensive and safety-critical domains such as

autonomous driving and space exploration, the resultant cost can be very high. Consequently, developing formal mechanisms to explicitly represent, reason about, and optimize the autonomy of a system is an important challenge in artificial intelligence.

To address this issue, we introduced a formal model called *competence-aware systems* (CAS) for optimizing autonomy in semi-autonomous systems over time by learning to optimally leverage the available human assistance [1]. CAS is a decision making framework for semi-autonomous systems where systems can operate at multiple levels of autonomy, each of which corresponds to different constraints on autonomous operation and different ways in which the human can assist the system to compensate for the constraints. Furthermore, each form of assistance is associated with a unique set of feedback signals. By learning from these feedback signals, the CAS can quickly grow to operate at its competence in almost all situations if feasible, effectively optimizing its autonomous operation by minimizing unnecessary reliance on human intervention while relying on the human when optimal to do so. An example of a CAS can be seen in Figure 1.

In our recent work [1], we provide the following results:

- We define the *competence* of a CAS as the optimal level of autonomy for every state-action pair conditioned on perfect knowledge of the human’s feedback distribution.
- We define a CAS to be *level-optimal* in a state if the optimal policy is executed at the CAS’s competence in that state.
- We prove that under standard convergence assumptions, **given any initial estimate of the human’s feedback distribution, the competence-aware system will converge to be level-optimal in the limit.**
- We validate the theory experimentally on multiple domains, including an autonomous vehicle domain motivated by real-world problems as part of a collaboration with *Alliance Innovation Lab Silicon Valley*, where we demonstrate that the number of episodes and feedback signals needed to reach level-optimality in all reachable states is quite reasonable (~30 and 50 respectively), leading to a significant decrease in the expected cost from ~25 to ~10 in the same span.
- We have deployed and tested an initial version of our approach on a fully operational AV prototype.

### 3 DIRECTIONS FOR FUTURE WORK

In general, it is difficult to determine if a planning model is missing important features in its representation of the environment based only on observed performance. The CAS model, however, provides an immediate ability to do so using existing information in the model. Under the assumption that humans are consistent and not ‘random’ in their feedback, by identifying state-action pairs that have received *seemingly* random feedback above an expected level of noise, the system can identify situations for which a human is likely making decisions based on information that is unused by the autonomous system. We are currently working on a method for autonomously recognizing potential candidates, determining the features most likely to be discriminators, and then using human feedback to determine which, if any, to include in the model.

When a modification is made to the system, it is not easy *a priori* to determine the extent of scenarios that the change will affect. As a result, what the agent learned prior to the modification

can no longer be strictly adhered to. However, gaining feedback information can be expensive and so it is wasteful to simply throw it out and restart. As a result, it is important to develop a well-defined means of bootstrapping the modified system with the old information to avoid wasting it, while ensuring that the system safely acts in a manner that is appropriate given the modification. One possible approach is to reset the level of autonomy in all states affected by the update for all actions to a lower, more conservative, level, effectively forcing the agent to reacquire feedback information from the human to move back to a higher level of autonomy.

Finally, in many industrial settings there is not one but many similar systems all deployed simultaneously. An obvious example of this is a fleet of autonomous vehicles. A natural extension of the CAS model is to learn the competence over a much larger space by using the information from a *fleet of systems*. However, the assumption that feedback comes from a stationary distribution no longer holds when each system may have their own human authority providing feedback from a unique distribution. One approach to this problem would be to have standardized feedback criteria, or to consider a training phase prior to data collection that would enable the human to gain a consistent view of the system’s capabilities.

### 4 CONCLUSION

Autonomous systems intended for deployment in highly complex domains for extended periods of time cannot be expected to capture all possible scenarios at deployment time. To avoid wasting operational productivity, and to minimize the chance of failure, it is important for such systems to have a well-defined means of identifying their competence for any situation, and subsequently be able to improve that competence over time. To this end, we have developed the *competence-aware system* model, and are investigating the following topics for future research:

- Introspectively identifying unused features in the current model using apparent randomness in human feedback to improve the robustness of the system.
- Developing a well-defined approach to bootstrapping a modified system using existing feedback acquired prior to the modification, while ensuring that the new system operates in a way that safely adheres to the modifications.
- Extending CAS to a multi-agent setting with multiple humans providing feedback from different distributions.

### REFERENCES

- [1] Connor Basich, Justin Svegliato, Kyle Wray, Stefan Witwicki, Joydeep Biswas, and Shlomo Zilberstein. 2020. Learning to Optimize Autonomy in Competence-Aware Systems. In *International Joint Conference on Autonomous Agents and Multiagent Systems*. forthcoming.
- [2] Alberto Broggi, Pietro Cerri, Mirko Felisa, Maria Chiara Laghi, Luca Mazzei, and Pier Paolo Porta. 2012. The VisLab Intercontinental Autonomous Challenge: An extensive test for a platoon of intelligent vehicles. *International Journal of Vehicle Autonomous Systems* 10, 3 (2012), 147–164.
- [3] Nick Hawes, Christopher Burbridge, Ferdian Jovan, Lars Kunze, Bruno Lacerda, Lenka Mudrova, Jay Young, Jeremy Wyatt, Denise Hebesberger, Tobias Kortner, et al. 2017. The STRANDS Project: Long-Term Autonomy in Everyday Environments. *IEEE Robotics & Automation Magazine* 24, 3 (2017), 146–156.
- [4] John F. Mustard, D. Beaty, and D. Bass. 2013. Mars 2020 science rover: science goals and mission concept. In *AAS Division for Planetary Sciences Meeting Abstracts*, Vol. 45.
- [5] Shlomo Zilberstein. 2015. Building Strong Semi-Autonomous Systems. In *AAAI Conference on Artificial Intelligence*.