

Final Report

From November 7th-9th I was in DC for the AI-HRI Symposium. On Friday, November 8th my partner from Harvey Mudd College and I presented a 5 minute summary of our paper from the summer and then took part in our poster session. The slides and poster can be seen attached. During the rest of the conference I heard many other presentations and contributed in conversations about artificial intelligence in the human-robot interaction space. Attached you can also find the schedule for the weekend.

Reflection

This summer I attended my first real conference. It was big: there were so many people and different sessions to go to. This one was different, there were maybe 30 people totally and all of the sessions were held in one room. There was no choice on what sessions to go to, you sat in one room for the whole time and heard all of the presentations. This created a very interesting, community type environment. Unfortunately I learned this may not be the topic for me, but I did realize how cool an event like this could be for finding your community.

I think that the most thought provoking even for me was the guest speaker Manuela Veloso. She skyped in from New York and talked to us mostly about her CoBots (collaborative mobile robots) from years ago at MIT. However, the main point of her talk was to talk about how we design robots that are intended to collaborate or interact with humans. Her main point was that we should not be creating robots to make humans feel comfortable. Her argument was that humans adapt and if a robot was built to be it's best with no concerns for humans, humans would get used to them and learn how to work with them since we are smart.

This went totally against what I have ever thought about design. At Olin we focus on human-centered design. And in fact, the reason I thought I might be interested in human-robot interaction is because you think about the humans too, not just the robots. So, to hear this well regarded professor say that humans are smart and can adapt was shocking to me.

I can't say I entirely disagree. I do think humans adapt and learn to not be scared. We adapted to using smart phones, we use Alexa, cars even. But that being said, I don't think any of those were invented without any thought of humans. I think we can adapt to a certain extent, but if, for example, we don't feel safe around a robot, it would take a lot to convince us it is safe. If we don't understand how to interact with it, we would probably give up quickly. Humans are smart and adaptable, but also expect convenience and things to be made for them. I don't think we can completely forget about people when designing humans to be inserted into the real world as Manuela was suggesting.

I was also so interesting and beneficial to see these presentations in person. It is one thing to read papers or articles on line, it is another to hear directly from the researchers. There were countless things I did not understand, but also so many new things I was hearing about. So many robots and projects I had either heard a bit about or never heard anything about and it was so cool to see, in person, where the state of human-robot interaction is. I think I realized the importance of conferences and my desire to go to more because hearing straight from the researcher's mouth is so much more beneficial than reading.

Overall, even though I am not extremely interested in being a part of the human-robot interaction field, it was really interesting to see what is happening in the field. It was also a really great experience to go to this small of a symposium and have a chance to talk to people who devote their entire life to the field.

Artificial Intelligence for Human-Robot Interaction (AI-HRI)

Day 1: Thursday, November 7 (Full Day)

09:00 – 09:15 **Introductions and Announcements**

09:15 – 10:15 **Invited Speaker: Michael Gleicher (45-minute presentation + 15-minute Q&A)**

10:15 – 10:30 **Breakout Session: Topics and Team-ups**

10:30 – 11:00 **Coffee Break**

11:00 – 12:30 **Long Paper Presentations (12-minute presentation + 3-minute Q&A)**

Towards A Robot Explanation System: A Survey and Our Approach to State Summarization, Storage and Querying, and Human Interface

Zhao Han, Jordan Allspaw, Adam Norton, and Holly Yanco

Petri Net Machines for Human-Agent Interaction

Christian Dondrup, Ioannis Papaioannou, and Oliver Lemon

Language-guided Adaptive Perception with Hierarchical Symbolic Representations for Mobile Manipulators

Ethan Fahnstock, Siddharth Patki, and Thomas Howard

Four-Arm Manipulation via Feet Interfaces

Jacob Hernandez, Walid Amanhoud, Anaïs Haget, Hannes Bleuler, Aude Billard, and Mohamed Bouri

Unclogging Our Arteries: Using Human-Inspired Signals to Disambiguate Navigational Intentions

Justin Hart, Reuth Mirsky, Stone Tejada, Bonny Mahajan, Jamin Goo, Kathryn Baldauf, Sydney Owen and Peter Stone

Automated Production of Stylized Animations for Social Robots

Adrian Ball and Ross Mead

12:30 – 14:00 **Lunch**

14:00 – 15:00 **Breakout Session**

15:00 – 15:30 **Lightning Talks (4-minute presentation + 1-minute Q&A)**

A Research Platform for Multi-Robot Dialogue with Humans

Matthew Marge, Stephen Nogar, Cory Hayes, Stephanie Lukin, Jesse Bloecker, Eric Holder, and Clare Voss

Fuzzy Knowledge-Based Architecture for Learning and Interaction in Social Robots

Mehdi Ghayoumi and Maryam Pourebadi

Multimodal Dataset of Human-Robot Hugging Interaction

Kunal Bagewadi, Joseph Campbell, and Heni Ben Amor

Towards Development of Datasets for Human Action Understanding in Human-Robot Interaction

Megan Zimmerman and Shelly Bagchi

Responsive Planning and Recognition for Closed-Loop Interaction

Richard Freedman, Yi Ren Fung, Roman Ganchin and Shlomo Zilberstein

Adaptable Human Intention and Trajectory Prediction for Human-Robot Collaboration

Abulikemu Abuduweili, Siyan Li, and Changliu Liu

15:30 – 16:00 **Coffee Break + Poster Session**

16:00 – 17:00 **Invited Speaker: Manuela Veloso (45-minute presentation + 15-minute Q&A)**

17:00 – 17:30 **Poster Session (cont'd)**

18:00 – 19:00 **Reception**

Artificial Intelligence for Human-Robot Interaction (AI-HRI)

Day 2: Friday, November 8 (Full Day)

09:00 – 09:15 **Announcements**

09:15 – 10:15 **Invited Speaker: Laura Hiatt (45-minute presentation + 15-minute Q&A)**

10:15 – 10:30 **Breakout Session: Topics and Team-ups**

10:30 – 11:00 **Coffee Break**

11:00 – 12:30 **Long Paper Presentations (12-minute presentation + 3-minute Q&A)**

Enabling Intuitive Human-Robot Teaming Using Augmented Reality and Gesture Control
Jason Gregory, Christopher Reardon, Kevin Lee, Geoffrey White, Ki Ng, and Caitlyn Sims

Negotiation-based Human-Robot Collaboration via Augmented Reality
Kishan Chandan, Xiang Li, and Shiqi Zhang

An Alert-Generation Framework for Improving Resiliency in Human-Supervised, Multi-Agent Teams
Sarah Al-Hussaini, Jason M. Gregory, Shaurya Shriyam, and Satyandra K. Gupta

Trust and Cognitive Load During Human-Robot Interaction
Muneeb Ahmad, Jasmin Bernotat, Katrin Lohan, and Friederike Eyszel

Towards an Adaptive Robot for Sports and Rehabilitation Coaching
Martin Ross, Frank Broz, and Lynne Baillie

Selfie Drone Stick: A Natural Interface for Quadcopter Photography
Saif Alabachi, Gita Sukthankar, and Rahul Sukthankar

12:30 – 14:00 **Lunch**

14:00 – 15:00 **Invited Speaker: Matthew Marge (45-minute presentation + 15-minute Q&A)**

15:00 – 15:30 **Lightning Talks (3-minute presentation + 2-minute Q&A)**

MAD-TN: A Tool for Measuring Fluency in Human-Robot Collaboration
Seth Issacson, Gretchen Rice, and James Boerkoel

Building Second-Order Mental Models for Human-Robot Interaction
Connor Brooks and Daniel Szafir

Commitments in Human-Robot Interaction
Victor Fernandez Castro, Aurelie Clodic, Rachid Alami, and Elisabeth Pacherie

Developing Computational Models of Social Assistance to Guide Socially Assistive Robots
Jason Wilson, Seongsik Kim, Ulyana Kurylo, Joseph Cummings and Eshan Tarneja

Towards Effective Human-AI Teams: The Case of Collaborative Packing
Gilwoo Lee, Christoforos Mavrogiannis and Siddhartha Srinivasa

An Automated Vehicle like Me? The Impact of Personality Similarities and Differences between Humans and AVs
Qiaoning Zhang, Connor Esterwood, Xi Jessie Yang and Lionel Robert

15:30 – 16:00 **Coffee Break + Poster Session**

16:00 – 16:30 **Poster Session (cont'd)**

16:30 – 17:30 **Breakout Session**

18:00 – 19:30 **Plenary Session**

Day 3: Saturday, November 9 (Half Day)

09:00 – 09:15 Announcements

09:15 – 10:00 Tech Talks (12-minute presentation + 3-minute Q&A)

Where is My Stuff? An Interactive System for Spatial Relations
Emrah Sisbot and Jonathan Connell

Solving Service Robot Tasks: UT Austin Villa@Home 2019 Team Report
Rishi Shah, Yuqian Jiang, Haresh Karnan, Gilberto Briscoe-Martinez, Dominick Mulder, Ryan Gupta, Rachel Schlossman, Marika Murphy, Justin Hart, Luis Sentis, and Peter Stone

MuMMER: Socially Intelligent Human-Robot Interaction in Public Spaces
Mary Ellen Foster and Olivier Canévet

10:00 – 10:30 Tech Discussion: Haves, Needs, and Wants

10:30 – 11:00 Coffee Break

11:00 – 12:00 Future Directions and Discussion

12:00 – 12:30 Wrap-up and Closing Remarks

Overview

We aim to improve human-robot collaboration in order to make the interaction between humans and robots more *fluent*.



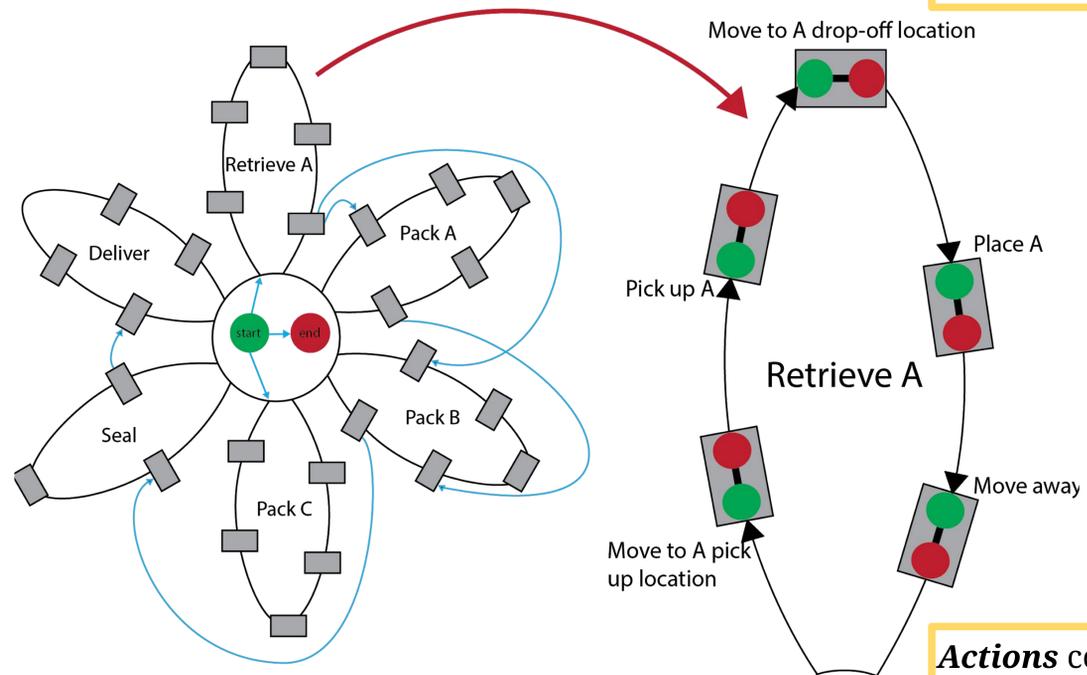
Image: Rethink Robotics

Fluency: “the coordinated meshing of joint activities between members of a well-synchronized team” (Hoffman 2019)

MAD-TN Model (Multi-Agent Daisy Temporal Network)

A **MAD-TN** models collaborative tasks. It consists of petals (subtasks) which may be completed in any order.

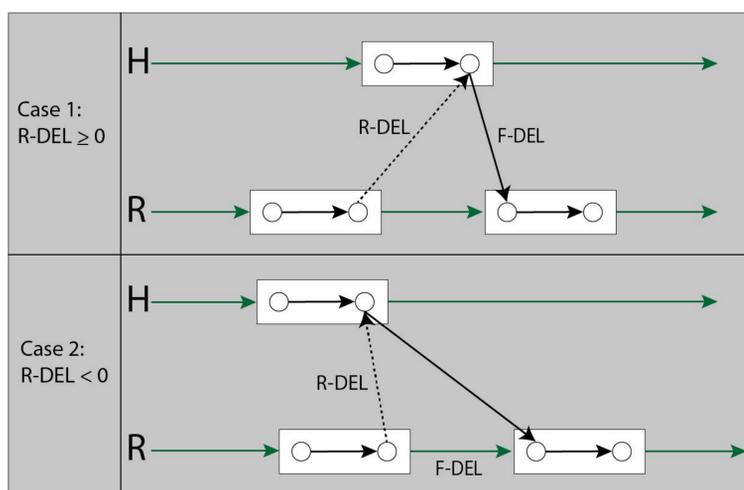
Petals contain atomic actions which must be completed sequentially.



This sample daisy represents the packaging task used in our experiment.

Actions comprise start and end times, as well as a duration constraint.

Fluency Metrics



Functional Delay is an existing metric that measures how well an agent predicts the actions of another. We introduce functional delay in the context of actions.

Resource Delay is a *new* metric that measures the delay in handoffs of resources from one agent to another.

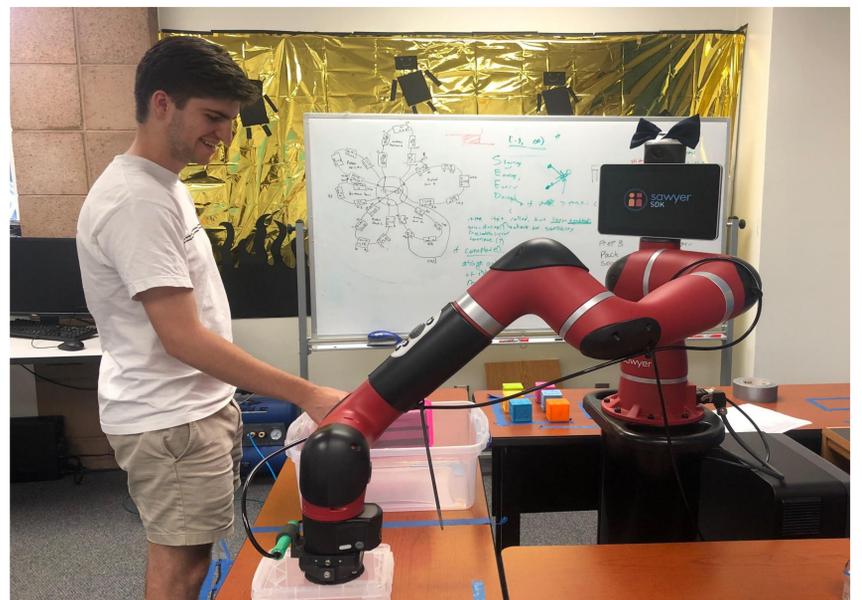
Concurrent Inactivity is a *new* metric which assumes that agents should always be active, and measures how frequently neither agent is active.

Upcoming Experiment

Hypothesis: Resource delay and concurrent inactivity are inversely correlated with subjective measures of fluency.

Human kept safe by collaborative, force-detecting robot.

Tasks dispatched to robot to optimize for fluency.



Clearly labeled zones allow intuitive resource handoffs.

Human's progress tracked to inform task dispatch.

Acknowledgments

Funding for this work was graciously provided by the National Science Foundation under grants IIS-1651822 and CNS-1659805. Thanks to the anonymous reviewers, HMC faculty, staff, and HEATlab members for their support and constructive feedback.

MAD-TN: A Tool for Measuring Fluency in Human-Robot Collaboration

Seth Isaacson*
Harvey Mudd College
sisacson@hmc.edu

Gretchen Rice*
Olin College of Engineering
grice@olin.edu

James C. Boerkoel Jr.
Harvey Mudd College
boerkoel@hmc.edu

Abstract

Fluency is an important metric in Human-Robot Interaction (HRI) that describes the coordination with which humans and robots collaborate on a task. Fluency is inherently linked to the timing of the task, making temporal constraint networks a promising way to model and measure fluency. We show that the Multi-Agent Daisy Temporal Network (MAD-TN) formulation, which expands on an existing concept of daisy-structured networks, is both an effective model of human-robot collaboration and a natural way to measure a number of existing fluency metrics. The MAD-TN model highlights new metrics that we hypothesize will strongly correlate with human teammates' perception of fluency.

Introduction

In the field of HRI, developing a fluent, collaborative schedule requires consideration of a range of parameters: abilities and preferences of teammates, temporal constraints within and between agents' tasks, and team fluidity and efficiency. An ideal solution to this problem adequately addresses each of these concerns and facilitates intuitive measurement of fluency.

Consider a situation in which a human and a robot are collaborating on a packaging task. Each agent may excel at certain tasks, and struggle with or be unable to complete others. In the packaging example, the robot may be better suited to retrieve an object from a dangerous machine. If the robot is stationary, a human would be much better suited to retrieve an object placed across the room. Additionally, some tasks may depend on the completion of other tasks: before the robot can deliver the package, the human must finish sealing it. However, two actions often do not depend on each other, such as retrieving two objects from two different locations. Ideally, the coordination between agents should be sufficiently fluent as to mimic the experience of human-human collaboration.

In this paper, we amend the Daisy model proposed by Maniadakis, Hourdakakis, and Trahanias (2017) so that it more generally and accurately captures the activities and temporal constraints of human-robot teams while

facilitating measures of fluency. The resulting Multi-Agent Daisy Temporal Network (MAD-TN) formalizes daises in the vocabulary of temporal networks. We show that this model is well-suited for intuitively representing collaborative, real-world tasks while also enhancing the monitoring and management of team fluency. Finally, we propose two new metrics that we hypothesize will correlate with human perception of team fluency. We further provide a new lens for understanding an existing fluency metric—functional delay—on two scales.

MAD-TN

We introduce the Multi-Agent Daisy Temporal Network (MAD-TN)—which we also call a Daisy—as a scheduling problem formulation that supports fluent collaboration in multi-agent tasks. At a high level, a daisy models multi-agent systems by breaking large tasks into smaller tasks that can be completed independently. Our work builds on the work of Maniadakis, Hourdakakis, and Trahanias, who first proposed daisy-structured temporal networks for use in the context of multi-agent collaboration (2017). We modify their formulation to achieve a more general model while maintaining the helpful traits they identified. Our model more precisely characterizes the daisy framework in the language of temporal networks. We also formalized actions to have start and end times, which in turn, are useful for better understanding the timing of interactions between agents

Background: Temporal Constraint Networks

A Temporal Constraint Network (TCN) is generally a set T of timepoints where $t_i \in \mathbb{R}$ with a set C of constraints on those timepoints. These networks are often encoded as directed graphs where nodes are timepoints and edges represent the constraints. An assignment of times to each $t_i \in T$ that satisfies all the constraints is called a *schedule*. A network with at least one valid schedule is called *consistent*. The Simple Temporal Network (STN) is an example of a TCN with constraints of the form $t_j - t_i \in [-c_{ji}, c_{ij}]$, where $c_{ij}, c_{ji} \in \mathbb{R}$ (Dechter, Meiri, and Pearl 1991). A Multi-agent STN generalizes STNs by designating which agents are responsible for the scheduling and execution of each timepoint (Boerkoel. and Durfee 2013). Fi-

* Authors listed alphabetically but contributed equally.

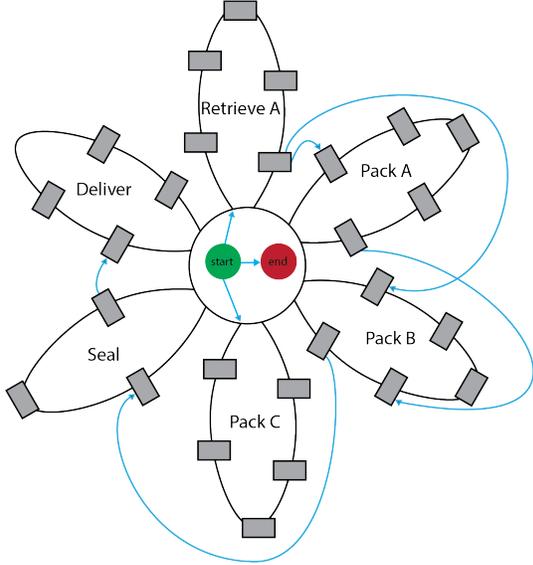


Figure 1: A daisy model of the experiment, adapted from Maniatakis, Hourdakakis, and Trahanias (2017). Each petal is labeled as the sub-task it represents and includes minimal actions, represented by the grey boxes. The blue arrows represent external constraints between petals.

nally, a Disjunctive Temporal Network (DTN) is similar to an STN, but allows for disjunctive constraints of the form $\bigvee_k (t_j^k - t_i^k \in [-c_{ji}^k, c_{ij}^k])$, which represent a disjunctive choice among different simple temporal constraint options. While the DTN model is a more powerful representation that allows deciding how to order events, finding a solution for a disjunctive network is NP-complete, whereas STNs are polynomial time solvable (Bhargava and Williams 2019).

Example Human-Robot Packaging Task

Next, we describe the components of the daisy and how to measure fluency within it using an example packaging scenario as shown in Figure 1. In this packaging scenario, we pair one robot with one human. This collaborative packing task contains sub-tasks such as: “Retrieve Object A”, “Prepare and Pack Object B”, “Pack Object A”, “Pack Object C”, “Seal Package”, and “Deliver Package.” Each sub-task will be assigned to the human or the robot. This task was designed to have distinct hand-off points where a resource (such as Object A) is transferred from one agent to another. We believe these interactions are important for fluency.

Actions

An action contains a start vertex, end vertex, and a makespan constraint between the two. The makespan constraint should always constrain the durations of actions to be non-negative. To provide the most flexibility and control when scheduling, actions should represent

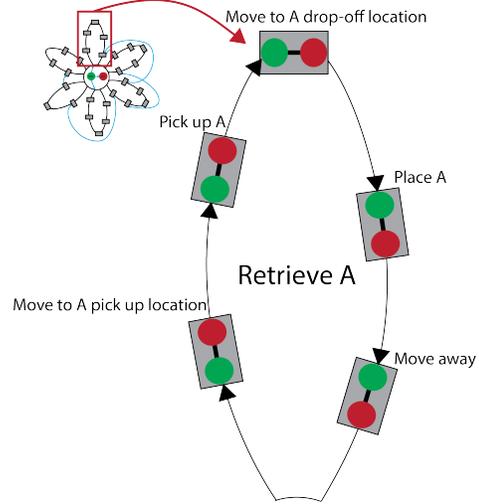


Figure 2: The “retrieve A” petal containing five actions represented as labeled grey boxes.

atomic tasks, such as moving to a location or picking up an object. In the packaging example, one action in the “Retrieve Object A” subtask is “Pick Up Object A.” This action contains a “Start Pick” vertex and “Finish Pick” vertex and might have a constraint between the nodes with bounds $[0.5, 3]$ meaning the pick may take anywhere from 0.5 to 3 seconds.

Petals

A petal serves to break an overall collaborative task into sub-tasks that can be done by a single agent. A petal contains a sequence of actions, connected by a set of constraints that imputes the order in which the actions are completed. A petal should generally deal with the handling of a single resource or set of resources. We assume that transition times between actions within a petal are negligible unless an agent must wait due to an external constraint (explained next). Figure 2 illustrates the “Retrieve Object A” petal from the packaging task.

Daisy Structure

A daisy consists of a set of petals, and any constraints between them. More formally, $D = \langle P, C, V_s, V_e \rangle$ where $P = \{P_1, \dots, P_k\}$ for a daisy D of k petals. V_s is a start vertex representing the start time of the daisy, and V_e is an end vertex representing the end time of the daisy. C is the set of constraints that are not contained within one petal. A constraint in C between V_s and V_e restricts the overall makespan of the daisy.

A daisy’s petals may be completed in *any* order, subject to the constraints between them. For example, since Object A has to be retrieved before it can be packaged, there would be a handoff constraint between the “Retrieve A” and “Pack A” petals, as illustrated in Figure 1. This mandates that the “Retrieve A” petal cannot follow the “Pack A” petal. So while petals can generally

be done in any order, C often enforces a partial ordering of petals. For our discussion, we will assume that a consistent total ordering of petals has already been determined.

Each petal in the daisy is completed by an agent. Generally, petals can be assigned to any agent capable of completing the subtask and once assigned, that petal becomes part of their local subproblem. The way agents are assigned to petals is an interesting planning problem that likely influences the overall fluency of the collaborative task. For instance, Maniadakis et al. (2016) developed a greedy planner that assigns petals based on each agent’s ability. However, for the purposes of this discussion, we assume that each petal has been preassigned to either the human or robot.

Fluency

According to Hoffman, fluency is a “coordinated meshing of joint activities between members of a well-synchronized team” (2019). Hoffman discusses human idle time, robot idle time, concurrent activity, and functional delay as quantitative metrics of fluency. The MAD-TN provides a natural way to measure each of these metrics. Further, the MAD-TN structure allows us to define two *new* fluency metrics: *concurrent inactivity* and *resource delay*, both of which we hypothesize will impact human teammates’ perceptions of fluency.

Agent Idle Time

Agent idle time refers to the amount of time an agent is not executing an action. In the daisy, an agent’s idle time is the sum across two sources of idle time. The first is rest time, which is the amount of time an agent is within a daisy but waiting to begin work on a petal (i.e., the time spent resting between subtasks). The second is time spent between actions. For most actions, this source of idle time will be negligible, since we assume agents will complete actions within their subtasks/petals without pause. However, an agent may be forced to wait due to an external constraint involving another petal.

Concurrent Activity

Concurrent activity is the amount of time both agents are active at the same time. Concurrent activity is simple to measure in the MAD-TN structure as the sum of the overlapping action times.

Functional Delay

Functional delay (F-DEL) is the delay (positive or negative) between when one agent stops working on a task and another agent begins work on a dependent task. Functional delay can be separated into the delay caused by a transition from a robot to a human or *vice-versa*.

Functional delay can be measured on daisies on two scales—at the petal/sub-task level and at the action level. Here we discuss it on the petal scale. If two petals

owned by differing agents share a constraint, a dependence is implied between the petals. Take, for example, when the human must seal the box (“Seal” petal in Figure 1) before the robot delivers it (“Deliver” petal in Figure 1). The robot’s delivery depends on the human’s sealing. This introduces the possibility of robot functional delay. That delay is measured as the difference between when the robot begins work on its petal minus the time when the human ends work on their petal. If that time is positive, it indicates the robot introduced positive functional delay. Negative functional delay indicates the robot anticipated the human’s action and began its petal before the human finished theirs. We discuss action-level functional delay below.

Concurrent Inactivity

We introduce a new metric, **concurrent inactivity**, that measures the amount time both agents are simultaneously inactive. Since the daisy structure naturally allows parallel task execution, agents should mostly always be active. Thus, both agents being inactive is a worst-case-scenario for efficiency. We expect concurrent inactivity to be inversely correlated with perceived fluency.

Resource Delay

When agents collaborate over a shared resource, a fluent handoff of that resource would be one where one agent finishes its use of the resources just in time for the other agent to begin using it. Thus, a fluent handoff is characterized by **Resource Delay (R-DEL)** — a *new* metric we introduce — and action-level functional delay. Resource delay captures the time between when an agent is ready to use a resource and when that resource becomes available. Action-level functional delay captures the time between when a resource is available and when an agent begins to use it.

Consider the situation in which the human is moving Object A to an accessible location for the robot to pack. If the robot arrives before the box has been placed down, it is being *blocked* by the human. This would be a positive resource delay. On the other hand, if the robot arrives after the human places Object A, the handoff becomes *stale*, which is a negative resource delay. In summary, R-DEL is the amount of time (positive or negative) between when one agent is ready to start an action and when the that action is enabled by the other agent involved in the handoff.

In our precise definitions of resource and action-level functional delay below, we focus on typical handoffs where one agent must complete the use of one resource prior to yielding it to the other agent. As such, the definitions that follow assume these external handoff constraints will have a lower bound of 0. In the future, we hope to extend our definitions of functional delay and resource delay to other types of external constraints, including e.g., synchronization constraints or handoffs with positive wait times associated them.

When there is a positive resource delay, action-level functional delay is the delay between when one agent

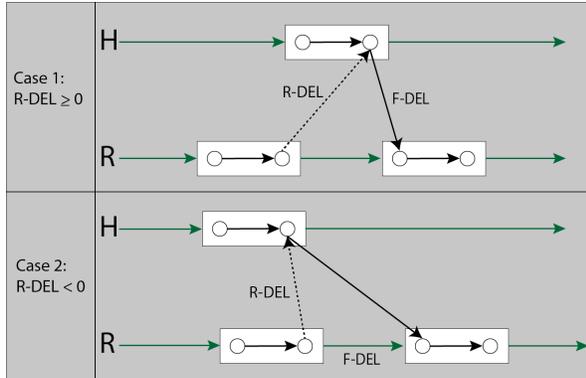


Figure 3: Resource delay (dashed line) and functional delay (solid line) for a scenario where the human blocks robot progress (Case 1) and one where the resource becomes stale (Case 2). Here, white boxes are actions while the contained circles are start and end nodes.

enables the others’ action and the time that agent begins, e.g., the delay between the human placing Object A and the robot beginning to pack Object A. A positive functional delay represents the transition time between agents due to the handoff, e.g., the time it takes the robot to process that Object A has been placed as shown as Case 1 in Figure 3. A negative functional delay represents anticipatory action, e.g., the robot began picking up object A before the human finished placing it down.

In the second case of a negative resource delay, we instead define functional delay to be the time it takes the agent to process and act on the fact that the resource is already available, and thus is measured as the time between the start of the activity and the end of the preceding one. In our running example, this would be the time between when the robot finishes the “Move to Object A” action and when it begins the “Pick Object A” action. Since the resource delay was stale, functional delay in this situation can only be positive (since the agent has arrived late, it is now impossible to begin work early). A situation with staleness and positive resource delay can be seen in Figure 3 in Case 2.

Human Perceptions of Fluency

Hoffman (2019) validated each of the existing fluency metrics against human perceptions of fluency using a simulated human-robot interaction in which a human and robot completed alternating tasks involving manipulating objects in a common workspace. He used an on-line platform to have participants watch videos of the simulation and answer questions about how fluent the interactions appeared. He found a statistically significant correlation with human idle time and the viewers perception of fluency, but not with robot idle time or concurrent activity. The strongest correlation reported was between function delay and fluency.

Overall, we believe the MAD-TN will provide a useful tool for empirically exploring the efficacy of these

five metrics at predicting human’s perceptions of fluency. Indeed, our ongoing investigation includes plans to design, implement, and analyze a physically embodied human-robot experiment that tests correlations of each of these with human’s perception. We posit that the degree to which these fluency metrics correlate with human perceptions of team fluency may be influenced by the experimental setup. For instance, a human’s teammate perception of fluency might be much more impacted by an idle physically-embodied robot teammate than an observer of a simulation. Further, we believe that the concurrent nature of MAD-TNs will highlight the importance of some metrics such as concurrent activity and concurrent inactivity. Finally, we are particularly motivated to further investigate the role that resource delay plays in human teammates’ perceptions of fluency. We suspect a human whose progress is impeded by their robot teammate might cause undue frustration or dissatisfaction. We believe our ongoing explorations will be key in understanding, designing, and planning more fluent human-robot interactions.

Conclusion

In this paper we introduced the Multi-Agent Daisy Temporal Network, or MAD-TN, which generalizes an existing model (Maniadakis, Hourdakis, and Trahanias 2017) by more precisely characterizing it in the language of temporal networks. The MAD-TN provides a convenient framework for measuring the fluidity of human-robot interactions using both existing measures and our new fluency metrics, Resource Delay and Concurrent Inactivity.

As we continue this work, we plan to further formalize the MAD-TN and empirically assess it’s usefulness in supporting fluent human-robot interactions. We also plan to design an experiment that validates our new and existing fluidity metrics. In particular, we hypothesize that collaborative tasks with significant resource delay will hurt overall perceptions of fluidity by human teammates. Eventually, we hope to explore how to order and assign petals to agents to maximize team efficiency and fluency.

While beyond the scope of this paper, we believe there are many interesting extensions of the MAD-TN that would be useful in human-robot teaming. We believe the disjunctive nature of petals will be a critical feature in scenarios that are more complex than the planned packaging task. Further, we believe MAD-TN could be extended to accommodate recursively defined daisy networks, where, e.g., actions within petals could themselves be daisies and petals of these sub-daisies could be assigned to sub-teams of agents. We also believe our action-level resource and functional delay measures will be critical for actions that require explicit synchronization such as when a robot physically hands an object over to a human teammate (Cakmak et al. 2011). Finally, the petal structure supports interesting questions of efficient resource allocation and issues of the privacy and autonomy of agents’ tasks.

Acknowledgements

Funding for this work was graciously provided by the National Science Foundation under grants IIS-1651822 and CNS-1659805. Thanks to the anonymous reviewers, HMC faculty, staff and HEATlab members for their support and constructive feedback.

References

- Bhargava, N., and Williams, B. C. 2019. Complexity bounds for the controllability of temporal networks with conditions, disjunctions, and uncertainty. *Artificial Intelligence* 271:1 – 17.
- Boerkoel, J. C., and Durfee, E. H. 2013. Distributed reasoning for multiagent simple temporal problems. *Journal of Artificial Intelligence Research* 47:95–156.
- Cakmak, M.; Srinivasa, S.; Lee, M. K.; Kiesler, S.; and Forlizzi, J. 2011. Using spatial and temporal contrast for fluent robot-human hand-overs. 489–496.
- Dechter, R.; Meiri, I.; and Pearl, J. 1991. Temporal constraint networks. *Artificial Intelligence* 49(1):61 – 95.
- Hoffman, G. 2019. Evaluating fluency in human–robot collaboration. *IEEE Transactions on Human-Machine Systems* 49(3):209–218.
- Maniadakis, M.; Aksoy, E. E.; Asfour, T.; and Trahanias, P. 2016. Collaboration of heterogeneous agents in time constrained tasks. In *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, 448–453.
- Maniadakis, M.; Hourdakis, E.; and Trahanias, P. 2017. Time-informed task planning in multi-agent collaboration. *Cognitive Systems Research* 43:291 – 300.