

DYNAMICAL SYSTEMS: ASSISTIVE ROBOTS AND AUTONOMOUS CARS

Interview with Professor Anca Dragan

BY ISABEL CRAIG, NIKHIL CHARI, YANA PETRI, ELENA SLOBODYANYUK

Dr. Anca Dragan is an Assistant Professor in the Department of Electrical Engineering and Computer Science at the University of California, Berkeley. Her lab focuses on developing human-robot interaction algorithms, which not only account for robot function, but also for robot interaction and collaboration with end-users. In this interview, we discuss human-robot collaboration in context of autonomous cars and other dynamical systems.



Professor Anca Dragan
[Source: <https://people.eecs.berkeley.edu>]

BSJ: How did you get involved in the field of Electrical Engineering and Computer Science?

AD: In middle school, I started really liking math and thought about its applications. I loved proving things, but also wanted to see the tangible effects of math on the world. Computer Science combined all of these things for me. Initially, I thought I would do Programming or Software Engineering, but then in high school I came across a book on Artificial Intelligence by Stuart Russell. It was very interesting because you could follow

certain solvable problems and immediately see follow-ups. The notion of creating new algorithms that solve new problems instead of implementing existing algorithms seemed very exciting. I thought I should do research in that field, and I loved the idea of agents that make intelligent decisions on their own, and that's how I got into AI. And then, at the time, I didn't even imagine you could work on AI in industry: it's about problems we have not yet figured out the answer to – that's research! So, here I am.

BSJ: How did you get involved specifically in Robotics?




Figure 1. a) Example of a collaborative table setting scenario, in which the robot needs to adapt to what people plan to do. b) Different timings of motion convey different information about the robot, including the perceived weight of the object being manipulated and how confident the robot is in what it's doing.

AD: I knew that I liked Artificial Intelligence, and Robotics is just the physical manifestation of Artificial Intelligence. With robots, the outcome of your algorithms is right there in front of you, moving! That's what makes them so cool!

BSJ: What can a robot infer from a human's ongoing actions, and how?

AD: Many things. At first, we started looking at how the robot can figure out what a person is reaching for. But since then, we've taken a step back and found that there are many more things that we infer by watching another person. Imagine watching someone perform a normal day-to-day task, like cooking. You can figure out if they're excited about it or bored. You can figure out if they're angry about something because they probably set things down in a much more decisive way. You can figure out if they're an expert or if they're a little hesitant. There's a lot of information (what we call internal state) that is communicated implicitly via actions. How do we infer internal state? Typically, by methods that fall under the category of Bayesian inference. The idea is that there's an underlying state that you can't observe, but you can observe the actions. You can treat those actions as evi-

"How do we infer internal state? Typically, by methods that fall under the category of Bayesian inference"

dence about that underlying state - all you need is what we call the observation model. If this were the correct version of the internal state - if the person were confident - how would they act? That's much easier than the other way around: if I see an action, what's the probability that the person is confident? Luckily, Bayes' rule gives us the way to go from one model to the other. It's this neat little trick that we've been using in Robotics for a long time now, but with the right observation model it becomes applicable to not just localizing a robot in a map, but also to "localizing" the a person's internal state.

BSJ: How do you take advantage of timing to make robot motion more natural?

AD: Well, not only can a robot observe a person and try to make inferences about the person's internal state, but people do the same thing when they observe robots. We perceive robots as agents. Timing plays a big role in the inferences that people make. You can have the same geometric motions (the same path), but if you time it differently - for instance, if you go fast and steady versus pausing repeatedly - people will interpret that in different ways. The same geometric path but different timings lead to different information that you would read into the motion. With my students Allan and Dylan, we've been exploring how to not necessarily make robots more natural, but rather more expressive. How do different timings communicate different information about the robot? It turns out that different timings communicate robot's capability and even the mass of the object that the robot is carrying. If I just carry an object slow and steady, you think that it might be light, but if I all of a sudden slow down and then speed back up, you might think that the object is heavier than it looks.

BSJ: How do you define "confidence" in a robot and how do you assign a quantifiable value to what is typically considered an emotional state?

AD: That's the hard part. This is why goals (what the person is reaching for) are easy – it's trivial to write down an equation for what a goal is, and it's relative easy to learn how people reach for different goals. Confidence is more tricky, because what does it mean mathematically in the first place? What we've done so far is a very simple model. "Confidence" for a robot is just its estimation of the probability of success. But another way to think about it is as a measure of uncertainty. The robot might start moving, and maybe it's uncertain about its location or the location of the object that it's trying to manipulate, so it keeps on gathering observations. Confidence is about the initial precision or uncertainty. Timing is involved in this because to gain the necessary precision at the end, when you're at the goal, the robot needs to slow down to gather more observations if it starts with low precision; whereas if it already has high precision, it can just go for it. That's a way to take something fuzzy like this notion of confidence and try to break it down to some mathematical, tangible model. But it's just a start.

BSJ: Another area of your research is self-driving cars, or autonomous cars. Why is important to monitor the interactions between human drivers and autonomous cars as a dynamical system?

AD: Autonomous cars work by trying to reach their destination and trying to avoid collisions. Implicitly, other human drivers on the road are simply obstacles that need to be avoided. So the car tries to predict what they might do and get out of the way so that they don't collide. As a result, these robots tend to be very defensive. For example, they would never merge in traffic if there's not a big enough gap. Or at an intersection they might get stuck, because people keep on coming and the autonomous car never gets to go. These cars are physically safe, but they're not necessarily very effective. Thus, to make them a little more effective, we've integrated a model of human response to the robot's actions. This is a dynamical system that incorporates human state as part of its state definition, and for which the dynamics model for how state changes as a function of the robot's actions now incorporates that the person will respond to the robot, which will change their state, which changes the overall system state. Not that it wasn't a dynamical system before: it just was a simpler one, where we were assuming the human state would evolve unaffected by what the robot does. Now, we're trying to add in that coupling, to say, "Well, wait a minute, a person's actions do depend on the robot's decisions." We use inverse reinforcement learning to create a model for how humans drive in response to the robot, and then we use that in our planners to come up with something for the robot to do. Dorsa, the graduate student doing this work, found some interesting behavior being produced by planning in this system. For instance, autonomous cars know that they can sometimes merge in front of someone because the person can slow down to let them in. Or, they decide to inch forward at an intersection to probe the person and figure out whether they're going to let them go. My favorite part is that if the person lets the robot through, it just goes, but if

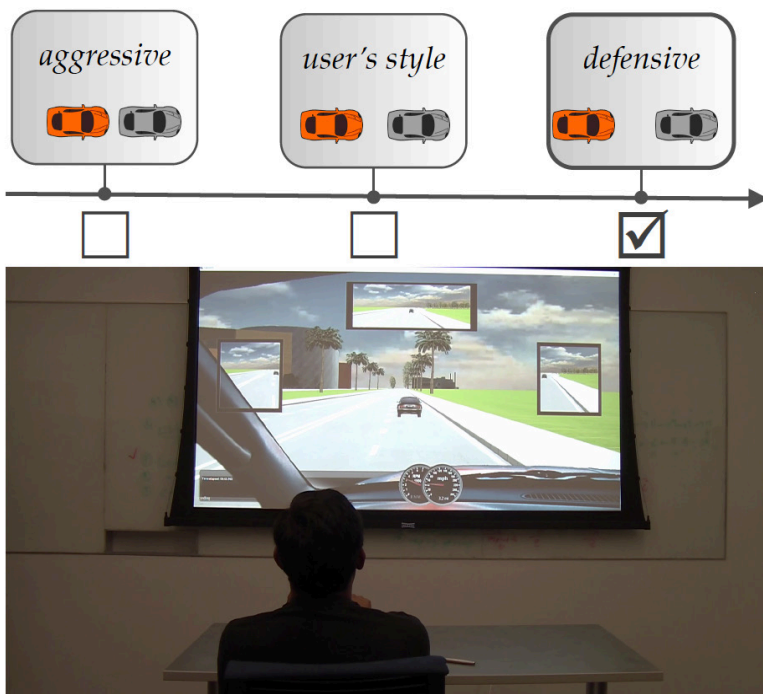


Figure 2. When asked to compare their own style (without knowing it's theirs), a more defensive style, and a more aggressive style, users typically prefer a more defensive style than their own.

“If the person lets the autonomous car through, it goes, but if the person doesn’t let it through, then the robot gently backs up.”

the person doesn’t let it through, then the robot gently backs up to let the person go. I like that a lot.

BSJ:What kind of environment were you testing your robots in?

AD: So far, for a lot of collaborative communication type of work we use our JACO robot, which is actually the first seven-degree of freedom robot that the Kinova company made. We use our JACO robot for a lot of the physical, collaborative interaction tasks. But for driving we’ve been using a simulator. We put the person in front of a steering wheel, and they have pedals, but they’re not driving a car in the real world (that would be a little dangerous at this point), and they’re looking at a monitor where they see their car move, and they’re reacting to the environment that way. We simulate both highway as well as city intersection scenarios.

BSJ:Based on the results of your research, what type of driving style do people prefer? Is it similar to their own driving style, or is it different?

AD: We started with the hypothesis that people would want an autonomous car to drive in the same style as them. So we thought that aggressive drivers would want a more aggressive car, defensive drivers would want a defensive car, and so on. It turns out that’s not the case. Chandrayee, who ran this study, found that people tend to prefer a driving style that’s more defensive than their own. But, interestingly enough, they think that they prefer the car to drive like them. So if you ask them to choose, they will choose not their own driving style, but something more defensive, but they think they are choosing their own style. They thus have a misperception of how they drive.

BSJ:In the future when you buy a car you’ll be able to select among different options?

AD: Even better, an option will be made tailored just to you! It’s more like in the future when you buy a car, you’re at the dealership, and you sit down in a nice simulator and there’s this virtual agent that says, “Hey, if I’m your car, and this is the environment, do you want me to do this,” and it plays out a trajectory, “or, do you want me to do this other thing.” Maybe one of them is more aggressive and one of them is more defensive. So you say, “I want this one,” and then you repeat this a few times, and the algorithm that we’re looking at is one that will actually enable this agent to, as quickly as possible, converge on the right style. This is an active learning type of approach, where you search for the most informative queries and comparisons that you can ask the person. And after a few such queries, hopefully the car will have converged to the driving style that you want.

BSJ:Are such dynamical systems specific to driving?

AD: No, they’re not, that’s a very good question. We’ve applied this to driving, but, if you think about it, robot actions influence human actions in all sorts of tasks. In fact, we’ve looked at a problem that’s a collaboration. So it’s no longer “I drive and I have my own objective, you drive and you have your own objective.” It’s an actual collaboration where a human and a robot work together to do a task, but the person is actually responding to what the robot is doing. And in particular, if the person is not perfect at optimizing and responding to the robot, meaning they can’t think many steps ahead, then a smart robot can guide the person to a better overall plan and compensate for a person’s myopathy. To make this concrete, we were looking at a handover example where a robot gives the person an object, and the way the robot decides to hold an object influences how the person grabs the object. If I give you a bottle upright, you can maybe grab it from the top, but if I tilt it you have a whole different set of choices for grasping it. So what’s interesting is that people grab the bottle in the most natural way. However, if they have to do something with it, like put it in a cupboard, they’re not very good at thinking ahead and grabbing it so that they can quickly put it upside down. So they end up having to re-grasp the object or twist their arms. But if the robot actually accounts for the fact that the person is a bit myopic and doesn’t think many steps ahead, then the robot essentially influences the person’s choices. For example, a robot can give a person a bottle in a way that would be most convenient for them to grasp in a

natural way and then put in a cupboard. I think it's a beautiful example of robots using the influence they have on people's actions to gently guide the person to perform better at this collaborative task.

BSJ: What has inspired you to become a co-PI for the Center for Human Compatible Artificial Intelligence (AI)?

AD: The Center for Human Capable AI is trying to develop an AI that is not just functional, but that can be beneficial to and compatible with humans. What got me most interested in the problem of human compatibility is that it's difficult to specify objective functions for robots. We specify some objective function, but inevitably, we don't think about corner cases or new situations. It's very common for AI systems, as they become better at optimizing the objectives that we give them, to end up not doing what we actually want and produce unintended consequences. In the center, we're working on the value alignment problem: how a robot or AI agent would be able to, over time, arrive at the correct objective function. One of the key ideas here is for the agent to not look at the objective function that's been given initially as set in stone but to be able to cooperate with the person and figure out the true objective function. This is cooperative inverse reinforcement learning: a collaboration with Stuart, Pieter, and our student Dylan.

BSJ: Would you say that your research has a lot of intersections with psychological research?

AD: Yes, in that our mission is to formally incorporate human state and human action into robotics, and account for the fact that human state is different and much more interesting than physical state. In order to figure out a model for human state, including human decisions and beliefs, and how these change based on the action of the robot, that's where psychology comes in. In particular, a branch of psychology called computational cognitive science that is all about developing computational models for how people reason and make decisions. AI agents can then use these mental models for how a person works so that they know how to interact with them.

BSJ: What are the future applications of your work?

AD: I think the applications are fairly broad. We recently focused on autonomous driving. That's probably going to continue because driving is a present and exciting problem. But we've also focused on manipulation: on interacting with robot arms and

more humanoid-like robots, which are harder problems because those robots don't work well right now. We do research on things that are applicable across multiple situations because we're trying to study the fundamental theory and algorithms that enable robots to correctly reason about people and their actions. This applies to anything from manufacturing, where you have robots moving out of the cages and working side by side with human workers, to robot-wheelchairs that can help people with disabilities live more independently, or to robots in the home helping you clean up the dining room table. There too, we want interaction, because we don't want all the people on the couch watching TV and the robots in the kitchen doing everything. Even in that domain, interaction is important.

REFERENCES:

1. Basu, C., et al. Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction, 2017.
2. Liu, C., et al. Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, 2016.
3. Sadigh, D., et al. Proceedings of the Robotics: Science and Systems Conference (RSS), 2016.