Panel on Research Challenges for Agent Verification

The first speaker was Dr. Kerstin Eder from the University of Bristol and Bristol Robotics Lab. She began by highlighting a quote from Yaron Kashai at Cadence, "Model Checking works best for well-defined models that are not too huge. Most of the world is thus not covered," which she felt encapsulated the major challenge faced by Agent Verification - namely the need to explore a wide range of techniques beyond model-checking, in order to provide coverage of all aspects of verification in robotic systems. She also promoted the need to design robotic systems "for verification".

She said that she, herself, would be interested in pursuing further research in Runtime Verification of robotic systems. At this point Dr. Dejanira Araiza-Illan pointed out that the new ISO Robotics standard (ISO 13482 2014) mentioned the need for runtime verification.

Dr. Brian Logan from the University of Nottingham asked what people really wanted to verify. He put forward the challenge of moving up through the levels of abstraction above temporal logic in order to identify the truly high level requirements. This needed to grapple with questions like "what is autonomy?" and what it meant for a robot to behave ethically. This also had ramifications for how we explain what we do to policy makers in order to make sure they had a realistic understanding of what verification could offer. He felt lack of understanding from policy makers presented a key risk to verification.

Dr. Dave Parker from the University of Birmingham mentioned scalability and learning.

In ensuing discussion Brian Logan highlighted the need to move from verification of existing designs to systems that were correct by construction, though he observed that this could cause problems with engineers and regulators. In particular, he had encountered an issue in his own work where engineers would not trust synthesized controllers. This led to a general discussion of what engineers might actually be doing, beyond simply producing a controller that satisfied the requirements, when they designed something. In particular it was generally agreed that requirements were often incorrect, and the design process served as a back and forth between design and requirements. The lack of transparency affected the trust system designers have in synthesis tools. In this context Kerstin Eder described a WHY button offered in a state-of-the-art formal verification tool for microelectronic design verification; this button could be used to drill down into the logic of a design in order to understand verification results. Something similar may be needed to increase the confidence and trust of engineers and developers in using correct by construction systems for autonomous robotics.

Dr. Michael Rovatsos from the University of Edinburgh highlighted existing social media systems that use algorithms from artificial intelligence often in obscure ways and asked about the challenges these presented to verification. It was far from clear that these systems, as they currently existed, adequately met the requirements of their users (or even their developers) and there were challenges in understanding

and defining requirements which encompassed multi-criteria, multi-stakeholder viewpoints. This related to questions of how we talk about properties of adaptive/ learning based systems and also drew on the question of what was needed to make people operate and trust autonomous systems, beyond simple regulation and verification.

Several of the speakers touched on other challenges: scalability, systems that could learn or otherwise self-adapt, and modelling the environment.

In the general discussion that followed questions of learning, ethically-aware autonomous systems, scalability of formal verification, interaction with regulators, and the consequences of verification failing were all discussed. There was some enthusiasm for the development of a relatively realistic challenge example, particularly one that could incorporate aspects of machine learning.