

Invited Talk

Shield Synthesis for Safe Reinforcement Learning

Roderick Bloem
TU Graz, Austria

You have a reinforcement learning system?
Sure, it works great, but does it give you any guarantees?
I thought not.

We will describe methods to use reactive synthesis to construct runtime enforcement modules (shields) that can ensure that a system works correctly, even if the system has bugs. If the system doesn't have too many bugs, the behavior of the shielded system will stay close to the behavior of the original system. We will show extensions in probabilistic and timed settings.