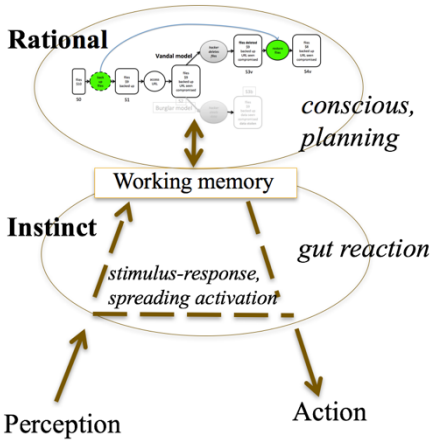
In previous work we have developed the DASH agent platform aimed at simulating human behavior in order to understand and mitigate the negative consequences of human security decision-making for security protocols. The DASH architecture, shown below, includes a dual-process approach with modules for rational and instinctive behavior, helping account for differing human responses to situations depending on factors such as level of attention, fatigue, stress or affect. It includes a reactive planning module that allows the agent to exhibit goal-directed behavior that is responsive to changes in the environment and support for different mental models that humans may use to reason about a situation.



The DASH agent architecture includes a dual-process model, with support for (1) reactive planning and (potentially flawed) mental models within the rational module and (2) spreading activation within the instinctive module.

This model has been successfully applied in a number of situations, including humans managing a portfolio of passwords on different sites under bounded rationality and planning to complete time-sensitive work under contention for shared resources, and launching cyber attacks with various goals and levels of expertise. The agents have been informed by surveys and behavioral experiments run by our team or others and have also been used to provide a realistic context for behavioral experiments.

Here we describe FARM, a toolkit that builds on the agent platform to (1) help developers describe the precise relationship between parameters in agents and the experimental work or data that supports them and (2) partially automate the exploration of the effect that parameter assumptions may have on the overall result of a simulation experiment. FARM includes an explicit representation of the underlying experiments and the confidence levels in agent parameters they may give rise to. Developers building on these agents can have the tool select the most plausible parameter values and can easily see what experimental data supports these values. The tool can also test the dependence of the outcomes of an experiment on these values and can in some cases suggest follow-on experiments or analysis that would improve confidence in the outcomes. The same approach can also help the experimenter simplify parts of the agent model, improving computational aspects with a known tradeoff on accuracy, hence the name, which stands for 'Finding the Appropriate level of Realism for Modeling'.

We demonstrate FARM retroactively on an experiment on human defenders that used DASH attacker agents, and to show confidence intervals in an agent modeling password behavior.