Undirected Task Switching as Optimization: A Quantum Probabilistic Model

Lena Rosendahl

Advised by: Jonathan Cohen and Naomi Leonard

Princeton University MAE

Introduction and Background:

Cognitive control has been defined as the "ability (of the human cognitive system) to configure itself for the performance of specific tasks through appropriate adjustments in perceptual selection, response biasing, and the on-line maintenance of contextual information" (Botvinick et. al, '04). The application of cognitive control is understood to be subject to limitation, which manifests in the form of boredom and fatigue, but for which the mechanism is not well-understood.

The Expected Value of Control theory suggests that the brain "specifies how much control to exert according to a rational cost-benefit analysis" (Lieder et al. '18), estimating an optimal amount of control according to, "the expected payoff from a (task), the amount of control that must be invested (…), and the cost in terms of cognitive effort." (Shenhav, Botvinick, Cohen, '13). In a similar approach Kurzban et al., "argue that the phenomenology of effort can be understood as the felt output of these cost/benefit computations", and that the build-up of these costs over time "motivates reduced (control) in the service of the present task", leading to "performance reductions". However, in the absence of explicit, deliberate cost-benefit calculations, the nature of these cost-benefit assessments is unclear.

One source of insight is the study of task-switching and the exploration of the cognitive stability-flexibility dilemma. A person doing nearly anything is a person performing a task, and a person switches tasks every time they shift their attention. A person, for example, reading a book on the bus must switch tasks to monitor the environment to assess the proximity of their stop. The problem of how and when we choose to switch tasks can be framed as an explore-exploit tradeoff and, thereby, as an optimization problem that humans solve frequently and, often, without explicit calculation. Work by Kreuger, Wilson, and Cohen indicates that "explore-exploit decisions are driven by three independent processes: directed and random exploration, and a baseline uncertainty seeking that is driven by a prior". They explain that directed exploration is an information seeking behavior while random exploration is driven by "decision noise". In deciding whether to explore, agents must also consider a loss in the form of "switch cost", measured by slowed reaction time and, consummately reduced reward, and attributed to the physical process of changing tasks (Kreuger, Wilson, Cohen '17).

Recent work by Musslick et al.suggests that limiting control is a means of optimizing in the cognitive stability-flexibility dilemma, which describes the need to focus enough to succeed in one task, but not so much as to incur excessively high switch costs to other rewarding tasks, as in the case of hyper-focusing in ADHD patients.

Busemeyer has provided strong arguments for applying quantum probability in assessing cognitive states, citing order effects, interference, and observer-dependency in measurement.

I will present a model building on EVC, the motivational model, and quantum cognition work to create a model of task-switching and fatigue. This model offers hope of a unifying description of multiple well-known but previously disparate phenomena.

I describe the task landscape as a dynamic, n-dimensional set of quadratic potentials, where initial depth is determined by reward and wells update in discrete time as agents gather information about the tasks. As the task landscape is internal to the agent, it depends on the agent's perception of the tasks. The potentials (also labeled as "attractor states") become more or less attractive as the agent gains information about their potential for reward.

In the context of experiment, wherein an agent is in a closed, controlled environment with a fixed number of tasks with known, static rewards dispensed upon successful completion of a task, we expect that the agent's assessment of task attractiveness will be a function only of the reward and the agent's assessment of their likelihood of succeeding at the task. As the agent succeeds or fails at the task, they collect evidence to either support or contradict their expectations of reward and will update the landscape accordingly.

We expect the depth of the wells, then, to update as a result of the agent's success or failure. I will explore multiple plausible update algorithms, to be fitted to experimental data.

An agent's state is described analogously to a quantum particle within the task landscape: the agent has "kinetic energy", a sum of a control constant and gaussian random noise, and is acted upon by the task potentials. The control term is determined by the agent and represents the amount of focus the agent chooses to apply to the tasks. The noise arises from several sources: noise internal to the agent (commonly seen in cognitive psychology), and an uncertainty bias arising from all tasks not explicitly assigned a reward term. The agent's kinetic energy determines whether they are found in a bound or scattering state. In a bound state, the agent has a higher probability of being found in the task wells than elsewhere. In a scattering state, the agent has equal probability of being found anywhere in the task landscape.

In a physical QM system, a particle is said to be found within a potential well when it is measured to be in the space occupied by that potential. In this model, an agent is said to be found in the well of the task they perform. We are required to make a distinction between the task the agent *chooses* to *attempt*, and the task that the agent *performs.*

In the case of two tasks, an agent will choose to attempt one task but, in addition to successfully completing that task, they may also fail by either succeeding at the other task (as commonly occurs in tasks that exhibit interference) or by failing both tasks.

In order for the wells to update appropriately, it is important to be able to classify both forms of failure, as an agent will not experience "accidental success" in the same way as "deliberate success".

In the single task landscape, failure only results in cognitive fatigue as the single potential diminishes. In a two-task landscape, an agent also experiences fatigue through diminishing task potentials as a result of failure. However, the agent may switch to another rewarding task. There is always a non-zero probability of a spontaneous switch.

I will also discuss experimental verification of and future directions for this model.