A normative theory of approach-avoidance conflicts during dynamic foraging in humans

Arthur Guez[1], Ritwik Niyogi[1], Dominik Bach[2], Marc Guitart-Masip[3], Raymond J. Dolan[4], Peter Dayan[1]

[1] Gatsby Computational Neuroscience Unit, UCL [2] Zurich University Hospital for Psychiatry [3] Karolinska Institutet [4] Wellcome Trust Centre for Neuroimaging, UCL

Introduction

We propose a normative model of the behaviour of human subjects playing a dynamic foraging game containing a time-stochastic threat. The game is intended to capture the essence of the conflict between approach and avoidance. The realistic nature of the task makes planning challenging; we therefore rely on recent innovations in model-based methods to approximate the optimal policy, and on Approximate Bayesian Computation to fit our models.

The Task

Transform an animal paradigm to study approach-avoidance issues in humans.

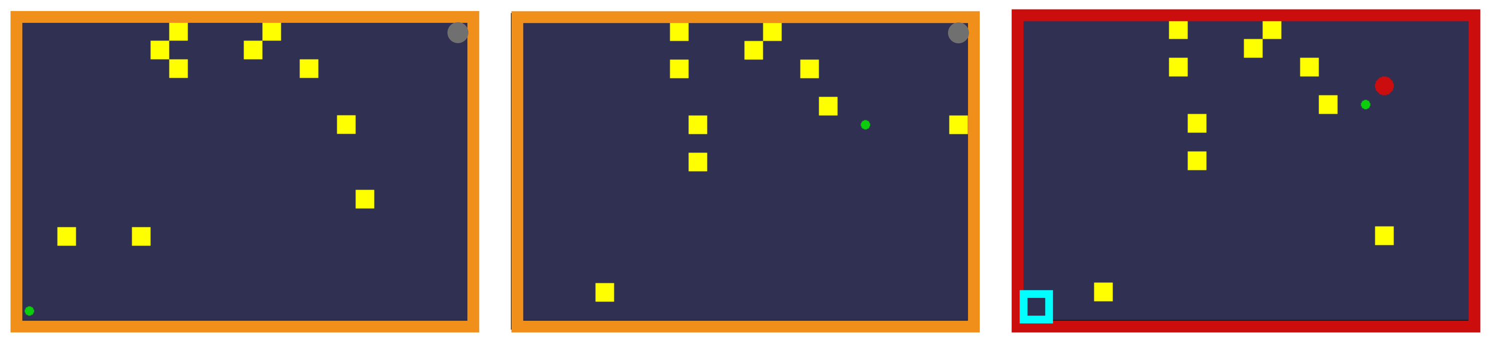

(a) $n = 0$ (b) $n = 3$ (c) $n = 3$

- Human player is foraging for tokens (■) on a 24x16 landscape grid,
- Tokens move randomly to different locations at regular intervals.
- Captured tokens are valuable at the end of the game (**approach motivation**).
- Sleeping robber wakes up at random & chases player (**avoidance motivation**).
- Player can only escape the robber at the safe place. **Loses all tokens if caught.**
- 3 threat level conditions (low, med, high) correspond to prob. of robber waking up.

Subjects: Group1: 25 participants (12 male, 13 female, 23±5yrs); Group 2: 12 controls (7 male, 5 female, 44±7 yrs), 12 hippocampus sclerosis patients (6 male, 6 female, 43±12 yrs).

Methodology: Modeling

Computational Level

- Discrete episodic MDP. Each time-step corresponds to 200ms, episode $\leq$ 15s.
- **State**: position of agent/robber, positions of tokens, wake-up state of robber, token tally, time. Assume transition model known.
- Reward function: 1 for each token, $-$tally$-\beta$ for getting caught.

Algorithmic Level

- Huge combinatorial state space and stochastic transitions.
- Optimal policy not computationally tractable $\Longrightarrow$ look for approximations.
- Many approximations perform well, but they are not all good match for the data!
- Consider heuristic planning and variants of model-based, forward-search, planning algorithms:

Greedy Heuristic
Go to nearest token. If $t \geq \tau$, return to safe place, for some threshold param τ.

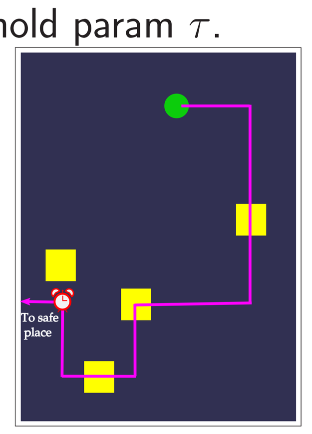

Monte-Carlo Tree Search (MCTS)
Plan using an adaptive **forward-search** tree at the **primitive action level**. Converges to optimal solution but expensive. At leaf nodes, estimate value using Greedy Heuristic (MCTS+GreedyRollouts) or value estimate using function approximation (MCTS+VFA).

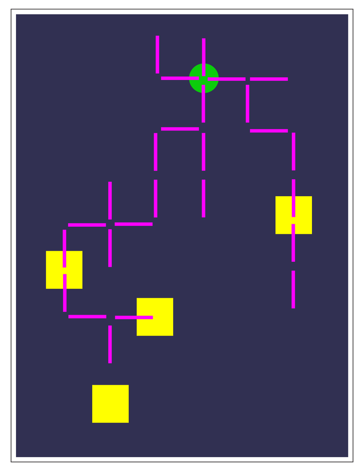

Methodology: Modeling (Cont.)

Monte-Carlo Tree Search with macro steps (MCTS-MS)
Plan using MCTS using **macro actions**. One macro action to go to each of the token and a macro action to return to safe place. A macro action is interrupted if robber wakes up or if target token disappears.

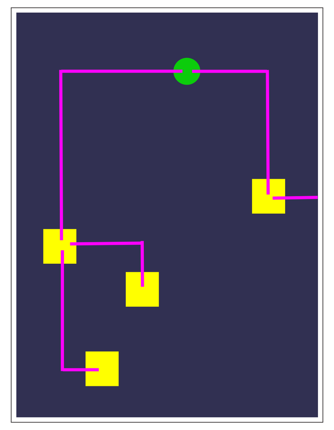

1MS+VFA
Planning 1 macro step ahead followed by a value estimate.

Values are learned using TD(0) in a linear architecture. Features include distances (to robber, safe place, tokens) and timing information.

Methodology: Fitting

With a complex model and task, we cannot directly compute $P(\text{data}|\text{model})$. Instead we rely on a likelihood-free method for model estimation:

Approximate Bayesian Computation (ABC)

- Use form of *approximate* rejection sampling.
- Define m:=model, θ:=params,$\mathcal{D}$:=data.
- Want posterior $P(m, \theta | \mathcal{D})$.
- Use features $\phi = f(\mathcal{D})$ as summary for data.

ABC rejection sampling alg.

1. Sample from prior $m, \theta \sim P(m, \theta)$
2. Simulate with m, θ to obtain $\hat{\mathcal{D}}$.
3. Compute features from simulation: $\hat{\phi} = f(\hat{\mathcal{D}})$.
4. Accept sample if $\epsilon = ||\phi - \hat{\phi}|| < \bar{\epsilon}$.

Features (summary statistics)
Commonly used features in animal literature on risk approach-avoidance plus feature revealing about planning mechanism (see right figure).
Greediness: went for nearest token / collected tokens.

Right: We fit separately the model/parameter for each threat condition (Best model for Low Threat: MCTS+GR, Mid Threat: MCTS-MS+GR, High Threat: MCTS-MS+GR). Data from Group1.

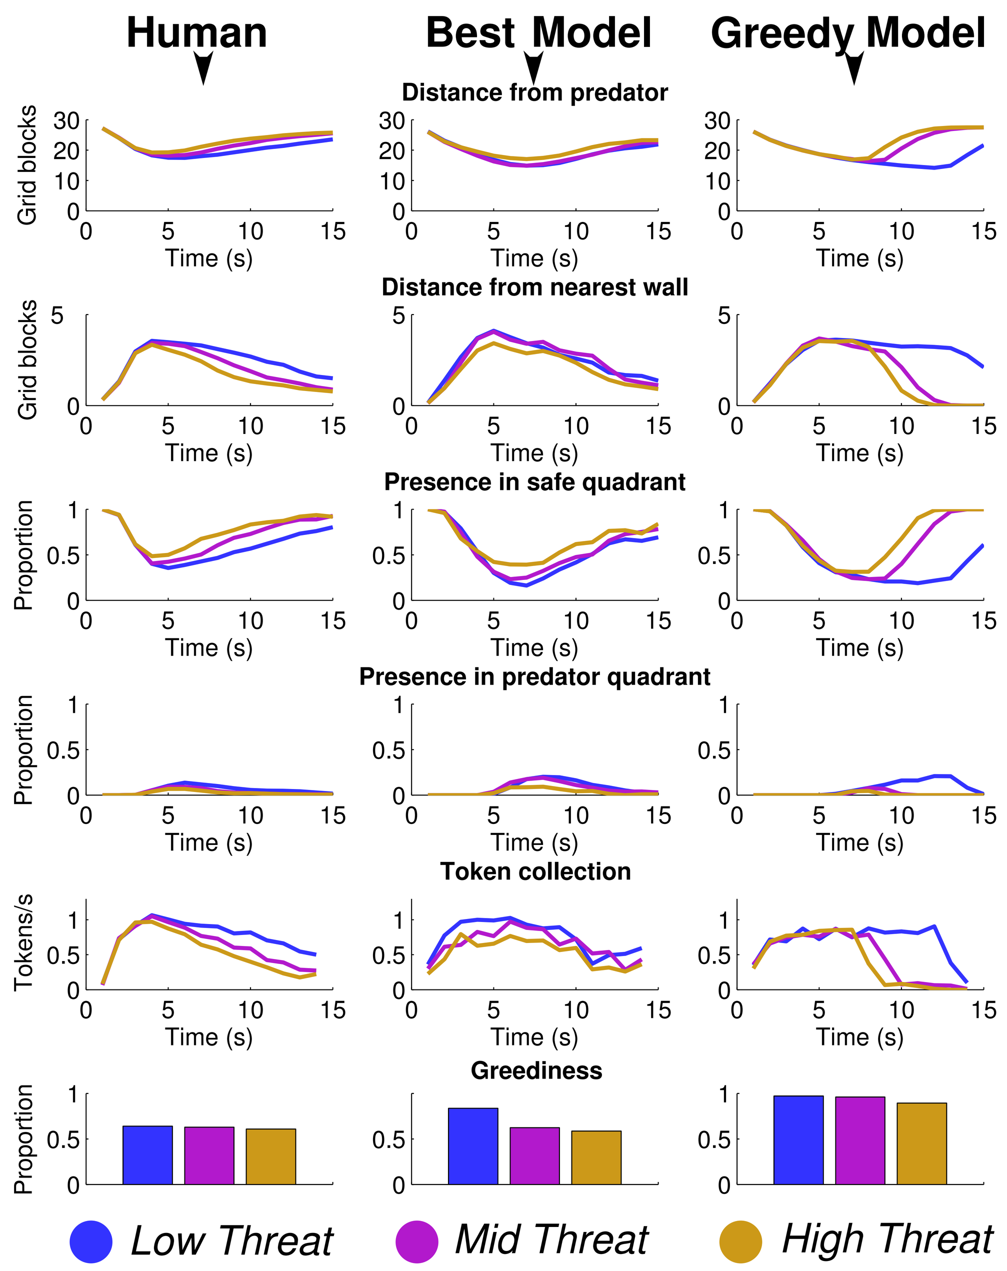

Sanity check

Testing explanatory power of the features.

1. Generate data $\mathcal{D}$ from model 2 (MCTS-MS+GreedyRollouts)
2. Run ABC with six different models.
3. Recover model 2 as most likely in resulting posterior.

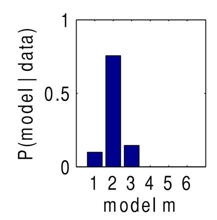

Risk and Behavior

Risk affects behavior in different ways. Multiple systems at play:

- Modify/bias loss function (e.g. be sensitive to variance).
- Pavlovian responding, pre-encoded behaviors (PAG).
- Hippocampal lesions in rodents have some anxiolytic characteristics. Associated with threat level in this task.

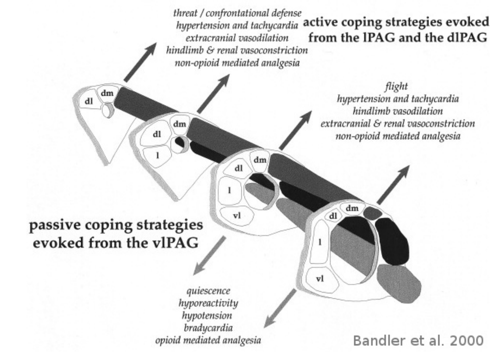

Columnar organisation of Periaqueductal gray (PAG).

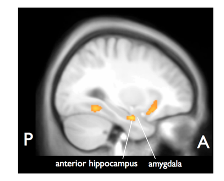

Brain area associated with threat level.

Preliminary Results

Different planning models/parameters supported by data for the different group/threat levels.

Left: **Control group** Right: **Hipp. lesion group**

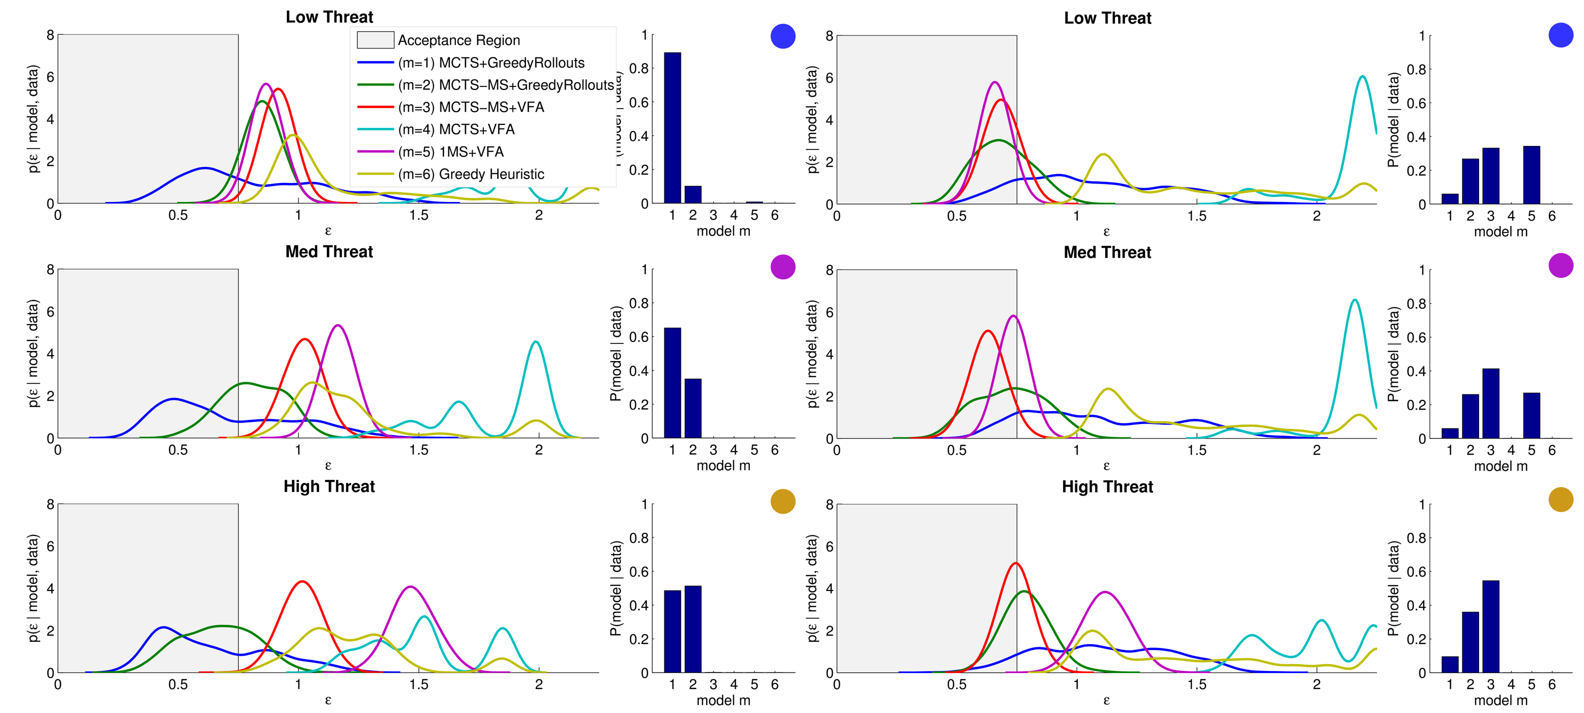

ABC posterior over planning models. *(x-axis: error ϵ)*

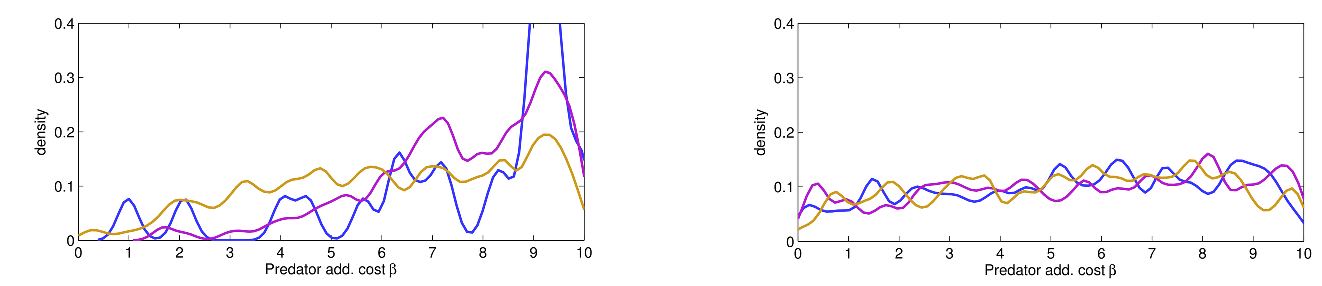

ABC posterior over param β (Added cost) for model MCTS-MS (m=2).

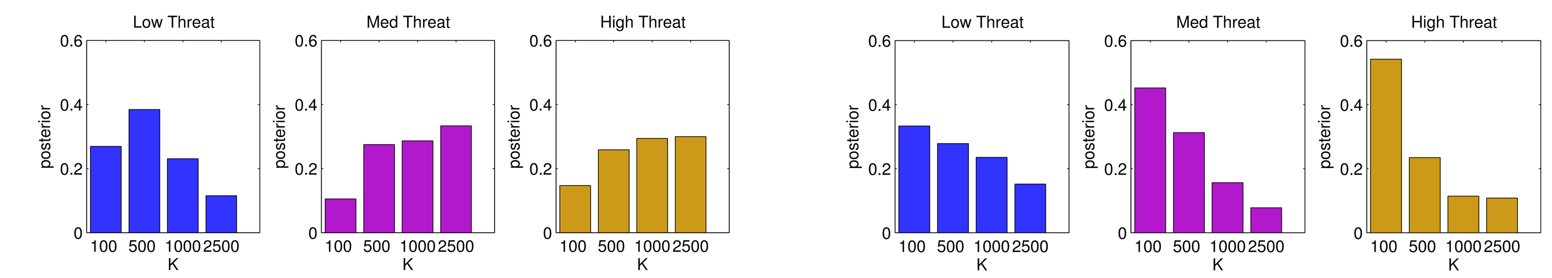

ABC posterior over param K (# sims) for model MCTS-MS (m=2).

Conclusion

- Study behavior models for risk approach-avoidance in humans.
- Complex task can be more revealing about subject's planning mechanisms,
- But also requires more intricate modeling and fitting.
- Preliminary results suggests possible causes for discrepancies between control and patient behavior.

References

- Bach et al. Characterising the role of the human hippocampus in approach-avoidance conflict. Submitted.
- Gray & McNaughton (2000) The Neuropsychology of anxiety: an enquiry into the functions of the septo-hippocampal system.
- Kocsis & Szepesvàri (2006) Bandit based Monte-Carlo planning.

This work was funded by the Gatsby Charitable Foundation and the Natural Sciences and Engineering Research Council of Canada.

Contact: aguez@gatsby.ucl.ac.uk