

Sophisticated Affective Inference: Simulating Anticipatory Affective Dynamics of Imagining Future Events

Casper Hesp^{1,2,3,4}, Alexander Tschantz⁵, Beren Millidge⁶, Maxwell Ramstead^{6,7,8}, Karl Friston⁴ & Ryan Smith⁹

¹ Department of Psychology, University of Amsterdam, Amsterdam, Netherlands

² Amsterdam Brain and Cognition Centre, University of Amsterdam, Amsterdam, Netherlands

³ Institute for Advanced Study, University of Amsterdam, Netherlands

⁴ Wellcome Centre for Human Neuroimaging, University College London, London, UK

⁵ Sackler Centre for Consciousness Science, School of Engineering & Informatics, University of Sussex, Sussex, UK

⁶ School of Informatics, University of Edinburgh, Edinburgh, UK

⁷ Division of Social and Transcultural Psychiatry, Department of Psychiatry, McGill University, Montreal, Canada

⁸ Culture, Mind, & Brain Program, McGill University, Montreal, Canada

⁹ Laureate Institute for Brain Research, Tulsa, OK, USA

c.hesp@uva.nl (twitter: @casper_hesp)

Abstract. In this paper, we combine sophisticated and deep-parametric active inference to create an agent whose affective states change as a consequence of its Bayesian beliefs about how possible future outcomes will affect future beliefs. To achieve this, we augment Markov Decision Processes with a Bayes-adaptive deep-temporal tree search that is guided by a free energy functional which recursively scores counterfactual futures. Our model reproduces the common phenomenon of rumination over a situation until unlikely, yet aversive and arousing situations emerge in one’s imagination. As a proof of concept, we show how certain hyperparameters give rise to neurocognitive dynamics that characterise imagination-induced anxiety.

Keywords: affect, counterfactuals, anxiety, active inference, anticipation

1 Introduction

A common aspect of human experience is that imagined, counterfactual events can have a significant impact on our affective states. In its extreme form, people suffering from a variety of psychiatric conditions, such as generalised anxiety disorder (Gale & Davidson, 2007), consistently report experiencing repetitively imagined “what-if” scenarios that have a significant impact on their real-time affective dynamics. This type of maladaptive, repetitive thinking about (often unlikely) negative future outcomes is re-

ferred to as rumination. Clinically validated therapeutic interventions for disorders involving rumination (e.g., cognitive-behavioural therapy [CBT], acceptance and commitment therapy [ACT]) also typically aim to reduce confidence in catastrophic imagined future events and ground patients in the here and now (e.g., see Barlow et al., 2017; Hayes et al. 2006). Although the effectiveness of such therapies is well established, their mechanisms of action remain poorly understood. Gaining a more detailed understanding of the specific neurocomputational mechanisms that underpin prospection-induced affect in general – and excessive rumination-induced anxiety in particular – is an important direction for future research.

In this paper, we aim to provide a mechanistic account of how affective responses can be generated by imagined future outcomes – and how this can become dysfunctional during rumination. By combining two recent developments in active inference, we provide a formal model of these phenomena and simulate how ‘overthinking a situation’ can occur – continuing to the point where unlikely, yet aversive and arousing situations emerge in one’s imagination. We employ an affective-inference agent (Hesp et al., 2020) equipped with the recursive belief-updating scheme of sophisticated inference (Friston et al., 2020). This powerful combination allows us – for the first time – to create an agent whose affective states change as a consequence of its internal machinations about possible future events. In this short paper, we present the underlying generative model and discuss its implications. We also show some brief illustrative simulations. We leave a more elaborate analysis of computational results for a variety of parametrisations for a future piece.

2 Methods

Here, we show how one can augment the Markov Decision Process formalism that underwrites the standard active inference scheme with a Bayes-adaptive deep-temporal tree search that is guided by a free energy functional as it scores counterfactual futures. By combining the ensuing recursive update scheme of sophisticated inference (Friston et al., 2020) with deep-parametric, *affective* inference (Hesp et al., 2020), we can derive a general-purpose generative model of the following mathematical form, summarised graphically in **Figure 1** and in tabular format in Table 1:

$$P(\tilde{o}, \tilde{s}^{(1)}, \tilde{u}, \tilde{\gamma}, s^{(2)}) = P(o_1 | s_1^{(1)}) P(s_1^{(1)} | s^{(2)}) P(s^{(2)}) \prod_{\tau=1}^{T-1} P(o_{\tau+1} | s_{\tau+1}^{(1)}) P(s_{\tau+1}^{(1)} | s_{\tau}^{(1)}, u_{\tau}) P(u_{\tau} | \gamma_{\tau}, s^{(2)}) P(\gamma_{\tau} | s^{(2)}) \quad (1)$$

In brief, the (higher-level) affective-contextual states $s^{(2)}$ entail three hidden-state factors: arousal, valence, and context. These factors map (through the likelihood matrix $\mathbf{A}^{(2)}$) onto three lower-level model variables: the latent states $s_1^{(1)}$, actions u_{τ} (i.e., possible state transitions at time τ), and \mathbf{G}_{τ} -precision γ_{τ} (i.e., action confidence at time τ). The latter is a scalar precision that scales the contribution of the expected free energy \mathbf{G}_{τ} to posterior beliefs about actions. This precision term can be read as a subjective estimate of confidence in model-based beliefs about action outcomes (Hesp et al.,

2020). This estimate is updated when posterior beliefs about action depart from one’s prior expectations such that it produces a concomitant change in the action-averaged expected free energy. The ensuing update term—named “affective charge” or AC —reflects changes in the confidence in one’s action model.

The lower-level state space $\tilde{s}^{(1)}$ comprises three hidden-state factors: location, context, and time, which map (through $\mathbf{A}^{(1)}$) onto two outcome modalities representing cues (e.g., visual) and rewards (e.g., gustatory). Following Hesp et al. (2020), each of the higher-level states can be associated with different combinations of lower-level parameters for $s_1^{(1)}$ (in terms of the initial prior $\mathbf{D}^{(1)}$), u_t (in terms of the baseline action prior \mathbf{E}_t), and γ_t (in terms of the rate parameter $\boldsymbol{\beta}_t$) through a higher-level likelihood mapping $\mathbf{A}^{(2)}$. For example, imagine you experience a pleasant low arousal state when you arrive home after a day’s hard work. This higher-level belief about your current state can then inform your lower-level action beliefs, e.g., by increasing the prior probability of actions associated with getting ready to sleep. Conversely, imagining yourself getting ready to sleep can further increase your experienced sleepiness. It is the latter type of reaction that we would like to model in general: affective responses (in this case, arousal-reducing responses) generated by imagined (internally simulated) future events.

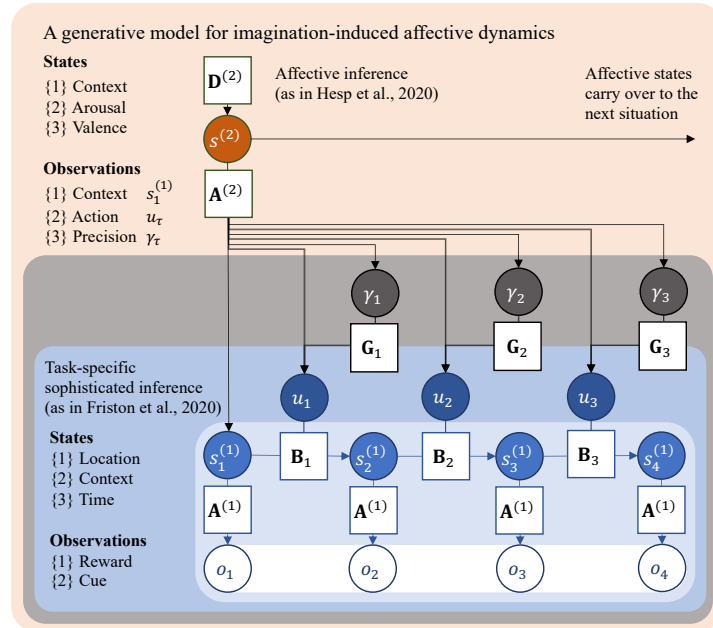


Fig. 1. A directed acyclic Bayes graph showing a generative model for sophisticated affective inference about higher-level valence, arousal and context states ($s^{(2)}$) based on (imagined) lower-level action-model precision (γ_t), actions (u_t), states ($s_t^{(1)}$), and outcomes (o_t) over four successive time points, thus combining sophisticated active inference (Friston et al., 2020) with deep-parametric affective inference (Hesp et al., 2020).

The specifics of the lower level generative model are not terribly important, but for the sake of our demonstration, we introduce a simple state space (within a stable context) that comprises four states (see Figure 2), each with its own observable outcome and associated preference C :

<i>state 1</i> :	an initial neutral state	$(C_{s=1} = 0)$	e.g., being at home base
<i>state 2</i> :	a slightly rewarding state	$(C_{s=2} = +1)$	e.g., picking berries
<i>state 3</i> :	a highly rewarding state	$(C_{s=3} = +2)$	e.g., hunting large prey
<i>state 4</i> :	a painful absorbing state	$(C_{s=4} = -2)$.	e.g., being wounded

The agent always starts in neutral state 1 and can move towards any of the four states by selecting up to three moves. Furthermore, a notion of safety is introduced by making state 2 a safer option than state 3: transitioning towards the latter has a higher probability of failure and can accidentally lead to painful *state 4*, which cannot be left until the end of the (4 time-step) trial. If we liken a trial to a working day, state 1 could be seen as the agent’s home base, state 2 as a safe activity with a small yet certain reward (e.g., picking berries), state 3 as a dangerous activity with a large yet uncertain reward (e.g., hunting prey), and state 4 as an unpreferred state that cannot be left for the rest of the day (e.g., being wounded).

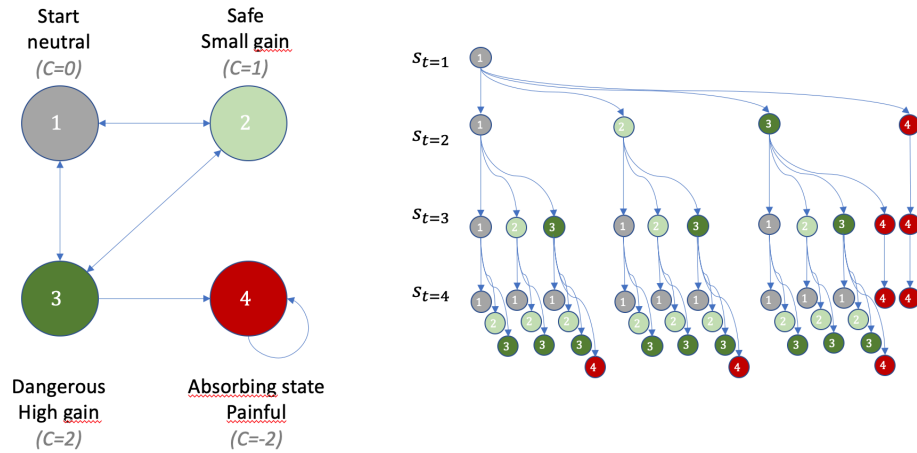


Fig. 2. An illustration of the state-space of the task with four states (left side, arrows indicating likely transitions) as it unfolds over four time steps (right side). The agent always starts in state 1 (grey), can get a small but safe reward in state 2 (light green), and a large but dangerous reward in state 3 (dark green). The latter is dangerous because it entails a larger probability of transition to the absorbing painful state 4 (red). The right side of the figure depicts the decision tree through which the agent searches to evaluate the expected consequences of each possible action sequence.

Table 1. This table lists the predictive posteriors that provide the empirical priors for our generative model.

Predictive posteriors	Mathematical definitions
$Q(s^{(2)}) = \text{Cat}(s^{(2)})$ <i>higher-level state beliefs:</i> <i>context, arousal, valence</i>	$s^{(2)} = \sigma[\ln \mathbf{D}^{(2)} + \sum_{\tau} \ln \mathbf{A}_{E_{\tau}} \cdot \mathbf{u}_{\tau} - \sum_{\tau} \left(\ln \frac{\beta^{(+,-)} - eAC_{\tau}}{\beta^{(+,-)}} \frac{\beta_{\tau}}{\beta_{\tau} - eAC_{\tau}} \right)]$ <i>(higher-level state prior)</i> <i>(action evidence)</i> <i>(affective evidence)</i> $eAC_{\tau} = (\mathbf{u}_{\tau}^o - \mathbf{u}_{\tau}) \cdot \mathbf{G}(u_{\tau}, o_{\tau})$ <i>(expected affective charge)</i>
$Q(s_{\tau}^{(1)}) = \text{Cat}(s_{\tau}^{(1)})$ <i>lower-level state beliefs:</i> <i>context, location, time</i>	$s_{\tau}^{(1)} = \mathbf{D}^{(1)} = \mathbf{A}_s^{(2)} s^{(2)}$ <i>(lower-level initial state prior)</i> $s_{\tau}^{(1)} \propto (\ln \mathbf{A} \cdot o_{\tau}) \odot s_{\tau}^u$ <i>(lower-level empirical state prior)</i>
$Q(s_{\tau}^{(1)} u_{\leq \tau}) = \text{Cat}(s_{\tau+1}^{(1)u})$ <i>action-specific state expectations</i>	$s_{\tau+1}^u = \mathbf{B}(u_{\tau}) s_{\tau}$ <i>(action-dependent state priors)</i>
$Q(u_{\tau} o_{\tau}, \gamma_{\tau}, s^{(2)}) = \text{Cat}(u_{\tau}^o)$ <i>outcome- and time-specific</i> <i>action expectations</i> <i>based on higher-level states</i> <i>and expected \mathbf{G}_{τ}-precision</i>	$u_{\tau}^o = \sigma[\mathbf{E}_{\tau} + \gamma_{\tau} \mathbf{G}(u_{\tau}, o_{\tau})]$ <i>(full action prior)</i> $\mathbf{E}_{\tau} = \mathbf{A}_{E_{\tau}}^{(2)} s^{(2)}$ <i>(baseline action prior)</i> $\gamma_{\tau} = 1 / \beta_{\tau}$ <i>(time-specific expected \mathbf{G}-precision)</i> $\mathbf{G}_{\tau}^{u,o} = \mathbf{G}(u_{\tau}, o_{\tau}) =$ $\mathbf{o}_{\tau+1}^u \cdot (\ln \mathbf{o}_{\tau+1}^u + \mathbf{C})$ <i>(expected risk)</i> $+ s_{\tau+1}^u \cdot \mathbf{H}$ <i>(expected ambiguity)</i> $+ \mathbf{u}_{\tau+1}^o \cdot \mathbf{G}_{\tau+1}^{u,o} \mathbf{o}_{\tau+1}^u$ <i>(expected free energy of subsequent actions)</i> $\mathbf{C} = \ln P(o_{\tau})$ <i>(prior preferences)</i>
$Q(o_{\tau} u_{\leq \tau}) = \text{Cat}(o_{\tau}^u)$ <i>action-specific</i> <i>outcome expectations</i>	$o_{\tau}^u = \mathbf{A} s_{\tau}^u$
$Q(\gamma_{\tau} s^{(2)}) = \Gamma(1, \beta_{\tau})$ <i>time-specific \mathbf{G}_{τ}-precision</i> <i>based on higher-level states</i>	$\beta_{\tau} = \beta^{(+,-)} \cdot \mathbf{A}_{\beta}^{(2)} s^{(2)}$ <i>(expected rate parameter)</i>

An important twist introduced in this model is that higher-level state beliefs can be updated recursively through pre-task mental deliberation, based on a deep tree search that unfolds pre-emptively on the lower-level. All the equations presented in **Table 1** can be evaluated without presenting any actual outcomes to the agent in question – that is, belief updating is guided by the probabilistic exploration of possible futures. This tree search involves recursive updating of lower-level action beliefs based on the counterfactual outcomes of actions that are sampled from predictive posteriors at each branching point of the tree. Because we equip the generative model with action-dependent \mathbf{G}_τ -precision estimation, we can see how each counterfactual future elicits an *expected affective charge* (*eAC*; see the first row of **Table 1**), which provides an ascending message to inform higher-level affective inference. The equation for *eAC* deserves further unpacking:

$$eAC_\tau = (\mathbf{u}_\tau^o - u_\tau^o) \cdot \mathbf{G}(u_\tau, o_\tau) \quad (2)$$

Where \mathbf{u}_τ^o is the empirical prior for a particular action and outcome at time τ , and u_τ^o is a particular outcome-action sequence drawn from the predictive posteriors. The *eAC* $_\tau$ term thus scores imagined departures from the model-averaged expected free energy for an imagined future at time τ . This *eAC* term is the anticipatory analogue of the affective charge term proposed by Hesp and colleagues (2020) as a plausible source of evidence for different valence states (i.e., pleasant/unpleasant states). The two main innovations afforded by sophisticated inference are that: (i) in *eAC*, the action sequences consider all combinations of individual actions and (ii) *eAC* is elicited in response to imagined, counterfactual actions (as opposed to events that have already been observed).

Simulating all possible sequences of actions and outcomes would quickly become intractable due to a combinatorial explosion (right side of Figure 2). For example, with 4 possible outcomes, actions, and time steps, the number of imaginable future possibilities would exceed 16,000. To solve this problem, sophisticated inference (Friston et al., 2020) provides a principled way of exploring the tree using the certainty of predictive posteriors. In terms of state estimation, these can be seen as empirical priors – as they are derived entirely from prior beliefs, which inform sampling of possible futures. In this work, every path has a probability of being selected, however small. Obviously, the number of explored possibilities will tend to increase with each iteration. By manipulating the number of iterations of such self-directed, recursive sampling of the future we can model traditional speed-accuracy tradeoffs for split-second decisions (i.e., too few iterations) as well as the detrimental effects of excessive deliberation (i.e., too many iterations), which characterises the phenomenon of rumination or ‘overthinking’.

Results

An exemplar result from our simulations is shown in Figure 3 below. It provides a simple demonstration of how sophisticated affective inference naturally underwrites affective responses to internally imagined futures. We simulated how particular hyperparameters give rise to neurocognitive dynamics that characterise imagination-induced anxiety or pessimism about the future. In particular, in Figure 3 we show how iterating the tree search too often (i.e., ‘overthinking’) can trigger recursive reductions in \mathbf{G}_τ -precision as the agent enters the following vicious cycle: (1) Every time they end up imagining a very negative outcome, their action-model confidence is reduced. (2) Every reduction in expected precision \mathbf{y}_τ (for simplicity assumed to be the same for all τ) will reduce reliance on one’s action-model for subsequent explorations of the future because these are sampled from the predictive posterior over action (see the fourth entry of Table 1). This type of excessive, negatively biased prospection (i.e., rumination) will subsequently influence the higher-level affective state, which recursively affects the lower-level sampling algorithm in multiple ways.

Crucially for these simulations of rumination, a negative affective state can bias the agent’s expectations towards negative outcomes and reduce lower-level \mathbf{G}_τ -precision even further, leading to increasingly pessimistic exploration of the tree. Such affective decision-tree pruning has been observed in a number of previous studies (Dayan & Huys, 2008; Huys et al., 2012; Huys et al., 2015; Nially et al., 2017). Our work shows how this phenomenon can be cast as a form of belief-updating under sophisticated affective inference (Hesp et al., 2020; Friston et al., 2020). Furthermore, the aetiology of many other psychiatric conditions seems to be intimately related to affective responses to imagined events: cravings in addiction, intrusive thoughts in obsessive-compulsive disorder, flashbacks in post-traumatic stress disorder, hallucinations and delusions in schizophrenia, fear of gaining weight in anorexia, excessive monitoring of self-states in anxiety, and so forth. As such, this type of formal model of imagination-induced affective responses could represent an important step forward in computational psychiatry and might one day be extended to aid in diagnosis or treatment for a variety of affective disorders – thus working towards computational nosology and precision psychiatry (see Friston, Redish, & Gordon; 2017).

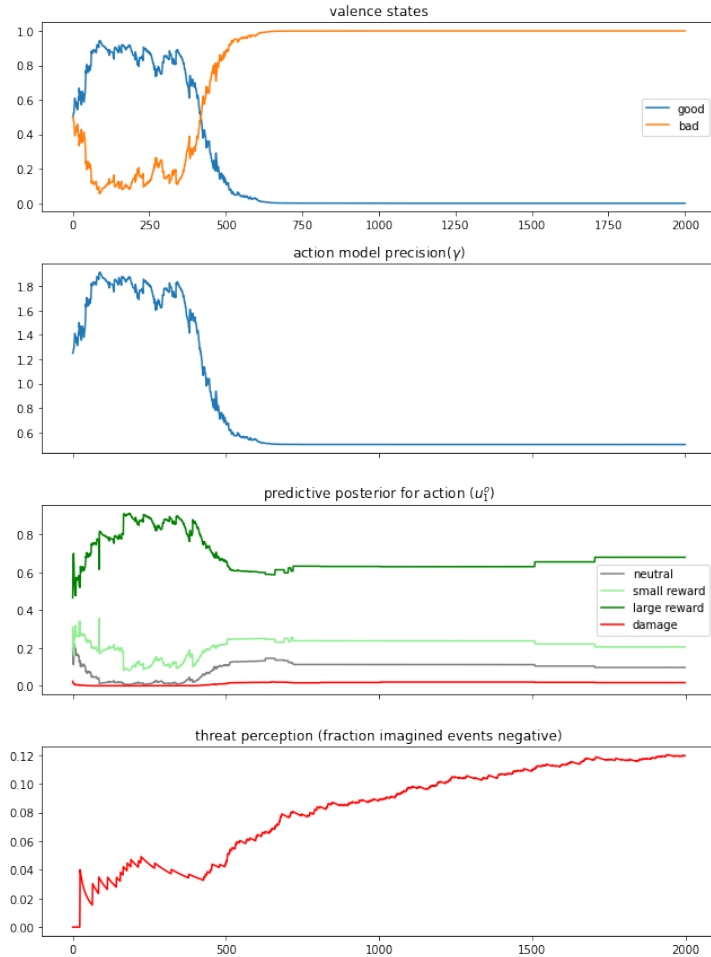


Figure 3. An example of simulation results showing detrimental effects of overthinking (i.e., rumination) when considering affective responses to imagined future events. Horizontal axes indicate the number of iterations and, implicitly, the amount of time allowed for internal deliberation. The top panel shows Bayesian beliefs about good and bad valence states (blue and orange, respectively); the second panel shows expected precision (blue); the third panel shows the predictive posterior for each possible first action: moving to either the neutral location (grey), the small reward (light green), the large one (dark green), or the painful absorbing state (red); the bottom panel shows the fraction of imagined events that were negative. Initially, exploration gives rise to an optimistic phase of increasingly positive valence (blue line in top panel), increasing action-model precision (second panel), increasingly positive expectations about future state transitions (dark green line in third panel) and a relatively small fraction of imagined negative events (red line in bottom panel). However, after roughly 500 iterations of the deep tree search, the agent devolves into a state of negative affect, reduced action-model precision, pessimistic expectations about future rewards, and a much higher fraction of imagined negative events.

References

1. Barlow, D. H., Farchione, T. J., Sauer-Zavala, S., Latin, H. M., Ellard, K. K., Bullis, J. R., Bentley, K. H., Boettcher, H. T., Cassiello-Robbins, C. (2017). Unified protocol for transdiagnostic treatment of emotional disorders: Therapist guide. Oxford University Press.
2. Dayan, P., & Huys, Q. J. M. (2008). Serotonin, Inhibition, and Negative Mood. *PLoS Computational Biology*, 4(2), e4. <https://doi.org/10.1371/journal.pcbi.0040004>
3. Friston, K., Da Costa, L., Hafner, D., Hesp, C., & Parr, T. (2020). Sophisticated Inference. Retrieved from <http://arxiv.org/abs/2006.04120>
4. Friston, K. J., Redish, A. D., & Gordon, J. A. (2017). Computational Nosology and Precision Psychiatry. *Computational Psychiatry (Cambridge, Mass.)*, 1, 2–23. https://doi.org/10.1162/CPSY_a_00001
5. Gale, C., & Davidson, O. (2007). Generalised anxiety disorder. *British Medical Journal*. British Medical Journal Publishing Group. <https://doi.org/10.1136/bmj.39133.559282.BE>
6. Hayes, S. C., Luoma, J. B., Bond, F. W., Masuda, A., Lillis, J. (2006). Acceptance and commitment therapy: Model, processes and outcomes. *Behaviour research and therapy*. 44(1):1-25.
7. Hesp, C., Smith, R., Parr, T., Allen, M., Friston, K., & Ramstead, M. (2019, December 3). Deeply Felt Affect: The Emergence of Valence in Deep Active Inference. <https://doi.org/10.31234/osf.io/62pfd>
8. Huys, Q. J. M., Eshel, N., O’Nions, E., Sheridan, L., Dayan, P., & Roiser, J. P. (2012). Bonsai trees in your head: How the pavlovian system sculpts goal-directed choices by pruning decision trees. *PLoS Computational Biology*, 8(3), e1002410. <https://doi.org/10.1371/journal.pcbi.1002410>
9. Huys, Q. J. M., Nially, N., Faulkner, P., Eshel, N., Seifritz, E., Gershman, S. J., ... Roiser, J. P. (2015). Interplay of approximate planning strategies. *Proceedings of the National Academy of Sciences of the United States of America*, 112(10), 3098–3103. <https://doi.org/10.1073/pnas.1414219112>
10. Nially, N., Huys, Q. J. M., Eshel, N., Faulkner, P., Dayan, P., & Roiser, J. P. (2017). The neural basis of aversive pavlovian guidance during planning. *Journal of Neuroscience*, 37(42), 10215–10229. <https://doi.org/10.1523/JNEUROSCI.0085-17.2017>