

Cognitive heuristics are generalisation processes that reduce uncertainty

Troy M. Houser
University of Oregon

thouser@uoregon.edu

Copyright. 2017–2023. Psychreg Journal of Psychology
An open access initiative by Psychreg Ltd
ISSN: 2515-138X



Coping with uncertainty is a ubiquitous challenge that all animals constantly face. Uncertainty makes choosing among options particularly difficult because animals are often unsure about the optimal, or “best”, option. This has led some to posit animals as being largely irrational; however, the ubiquitous nature of uncertainty means that computing optimal decisions is simply not possible for animals. Instead, animals evolved the use of heuristics to make rapid decisions that are often useful. Recently, this use of heuristics has been theorised to be a rational response to ecological constraints on computational capacity. In this Opinion, we expand on this view through the lens of the exploration-exploitation dilemma. Specifically, we dissociate between objectively optimal and what we call “platonically optimal” decision-making strategies. Then we describe the computational architecture that can be used to study differences in decision-making, and finally, we speculate on a potential connection between heuristics and generalisation processes. Along the way, we incorporate evolutionary, philosophical, psychological, and neuroscientific findings to substantiate claims made. The goal of this opinion is to weave together disparate literatures to provide a conceptual framework for decision-making in the face of uncertainty.

Keywords: coping; cognitive; decision-making; opinion; uncertainty

Animals need resources from the environment. Consequently, animals are constantly faced with the choice to either exploit known resources, foregoing any unknown resources that may be more rewarding, or forego the known resources to explore the unknown that may be less rewarding. This is known as the exploration-exploitation dilemma. Though there are important differences, the tradeoff between exploitation and exploration can be likened to the speed-accuracy tradeoff (Milli et al., 2021), both of which can be subsumed by the cost-benefit tradeoff (Kahneman, 2003). In order to make adaptive decisions, animals must weigh the costs and benefits for current options and choose whichever option yields a smaller ratio. This is the expected utility theory, which had dominated the philosophy of decision-making for centuries before cognitive scientists realised that people's decisions often deviate from optimality (Bottom et al., 2004; Kahneman & Tversky, 1979). Kahneman (2003) uses the following example to illustrate this latter remark. Imagine there is a bat and a ball, whose sum total costs \$1.10.

The bat costs \$1 more than the ball. How much is the ball? Most people asked this question report (at least) an initial urge to say 10 cents (Kahneman, 2003), and it is reasoned because we use heuristics, or biases in cognition, to facilitate rapid decision-making. There has been extensive study of the use of heuristics in cognition and behaviour, across species; however, far less work has analysed the problem that heuristics are solving. In this Opinion, we propose that heuristics are a solution to uncertainty.

Uncertainty tinges everything that animals do. Given that perception is generated by neural activity induced by sensory signals, the brain never has direct access to the outside world. The outside world is always communicated to it through an intermediary. Therefore, we can never be certain of even our current state, where state can denote a spatial location, a physiological state of being, or an abstract state of mind (Gershman & Uchida, 2019). Moreover, there are more forms of uncertainty that influence our decisions and behaviour, such as how to assign value to states or how to map states to actions. Interestingly, uncertainty maps on quite nicely to the exploration-exploitation dilemma. That is, exploitation is usually equipped with low levels of uncertainty while exploration is equipped with high levels of uncertainty (Krebs et al., 1978; Mehlhorn et al., 2015). Ergo, by studying uncertainty, we may better understand decision-making through the lens of exploration-exploitation. If heuristics are a solution to uncertainty, we would expect to see increased use of heuristics during exploration and decreased use of heuristics during exploitation. We will explore this sub-hypothesis below first.

Also, assuming that heuristics are indeed a solution to uncertainty, then we can ask whether they are an optimal solution or not. If they are optimal, this would provide support for expected utility theory. That is, just because animals make “objectively” suboptimal decisions, they may be making optimal decisions within the current constraints, such as time pressure, resource availability, resource accessibility, and physical capacity. This revised version of expected utility theory has been cast as ecological rationality (Todd & Gigerenzer, 2007), bounded rationality (Todd & Gigerenzer, 2000), and resource-rational decision-making (Milli et al., 2021). Finally, in this Opinion we will draw the connection between heuristics and generalisation. Generalisation is the process of drawing on past experience to respond to novel situations (Taylor et al., 2021). Thus, the use of generalisation to combat uncertainty is already well documented. Here, we make the connection between generalisation and heuristics to (hopefully) unite two disparate literatures as serving a common function with common neurobiological substrates.

Heuristics as solutions to uncertainty

Imagine walking through a forest. Apple trees are sparsely distributed through this forest. You are getting hungry. Your stomach starts rumbling. Finally, you stumble upon one of the apple trees, eat and go home. The next day, you can choose to exploit or explore. Exploitation means you would return to the same tree you visited yesterday because you know that it has delicious apples to offer. Exploration means you ignore the tree you already visited to search for a tree with even better apples. Critical to resolving the exploration-exploitation dilemma is the search process. How long are you searching? What is the terrain like that you must traverse? Are there any dangers lurking? These are

just a few of the uncertainties that are associated with exploration. We propose that heuristics are used to guide this search process through the unknown.

The giving-up time rule (GUT) is one heuristic used by animals to deal with the uncertainty of time. GUT embodies the notion that after a certain amount of time spent foraging in a patch without finding resources, move on. In the example above, you might spend some time searching a patch of forest. If you find an apple tree, you stay in that patch, but if not, you go to a different patch. Formally, GUT corresponds to a Lévy flight through space (Abe, 2020), which is proven to be the most efficient random trajectory through space (Green, 1998). Because Lévy flights are optimal random trajectories, it is likely that evolution provided pressure to adopt Lévy flight search strategies (Viswanathan et al., 1999, 2008). Indeed, numerous species have demonstrated search patterns and exploratory behaviours that approximates Lévy flights, such as great white sharks (Sims et al., 2012), albatrosses (Humphries et al., 2012), bumblebees and deer (Edwards et al., 2007), and even modern-day hunter-gatherers (Raichlen & Polk, 2013).

Further, fossil trace evidence of 50 My-old sea urchins suggest that Lévy flight search strategies evolved a long time ago (Sims et al., 2014). Other search heuristics correspond to sampling strategies. For example, Thompson sampling embodies a form of randomised exploration (Thompson, 1933), where one randomly samples a subset of possibilities and then chooses greedily, or via soft maximisation, which sample to visit (Gershman, 2018, 2019; Tomov et al., 2020). The distance-dependent Chinese Restaurant Process (ddCRP) is a sampling strategy that creates clusters based on the similarity between old clusters and a current observation (Song et al., 2022), and thus, can be used to guide exploration based on observations that are similar to previously-made observations that yielded resources. ddCRP is akin to similarity-based models of categorisation (i.e., the representativeness heuristic; Love et al., 2004; Nosofsky, 1992; Shepard, 1974, 1987; Bowman et al., 2020; Bowman & Zeithamova, 2018; Taylor et al., 2021; Zeithamova et al., 2012).

Another line of work that supports heuristic-guided exploration is entrepreneurship (Holcomb et al., 2009). Entrepreneurs are rarely faced with clear-cut scripts to adhere to (Baron, 1998), meaning their uncertainty levels should be high. Entrepreneurs have to be more exploratory. To cope with this lack of structure, entrepreneurs use heuristics to simplify their decision-making processes (Shepherd & Zacha, 2001; Zacharakis & Shepherd, 2001) a greater degree than normal controls (Nouri et al., 2017). Given this line of work, we might expect to see reduced use of heuristics by people that show more exploitative behavior, such as drug addicts and problem gamblers. On the contrary, multiple studies have implicated addicts in more frequent use of heuristics (Addicott et al., 2017; Dezza et al., 2021; McGuire, 1997); however, we think these findings do not discount the possibility that heuristics are solutions to uncertainty, for this does not preclude their misuse. A process that evolves for one thing can of course be co-opted for something else. Moreover, the increased use of heuristics for exploitation can still be construed as a solution to uncertainty: Perhaps, rather than increase exploitation, addiction increases sensitivity to uncertainty. This construal is supported by studies showing higher intolerance of uncertainty (IU) in addicted individuals than controls (Garami et al., 2017; Radell et al., 2018). IU is a robust measure of how people perceive uncertainty and the amount of information required prior to making decisions. One study showed that people with high IU sampled significantly more marbles from a bag than those with low IU before estimating the number of black versus white marbles in the bag (Ladouceur et al., 1998). For addicts, it is possible that the threshold demarcating sufficient information has become externalised as a drug or slot machines.

In summary, a number of distinct research threads point to the use of heuristics as a solution to uncertainty. Further, heuristics are a solution selected by evolution. In this light, heuristics are not the irrational responses that cognitive scientists assumed them to be (Ariely, 2009; Kahneman & Tversky, 1979). Instead, heuristics are an adaptive solution to something that can never be fully resolved. This entails that decision-making is inherently probabilistic. Moreover, recent advances in computational modeling, specifically reinforcement learning (RL), have demonstrated that there are still optimal solutions to this probabilistic framework. In other words, objective utility might not be possible to maximise but the asymptotic odds of utility are. We will review this idea in more detail next.

Optimal suboptimal decision-making

In a Platonic world, abstract ideas have the same quality as the world. Thus, for every abstract, mathematical notion, there is a corresponding physical phenomenon and vice versa. This is implicit

in maximising “objective” utility. For example, how does one decide on which college to apply to? If we entertain the idea that there is an objectively correct option, then we might begin by ranking options by prestige. But what if I want to go for chemistry and the best school is ranked high in all subjects except for chemistry? Then, I rank the schools based on how good their chemistry department is. But, what if only some of their chemistry professors are the best? Then, I need to factor in different rankings of professors. What if some of the best professors do not like helping new students? Maybe I then factor in student ratings of them professors. What about the rest of the department? Will I be able to get along and collaborate with my colleagues? How difficult was it for alumni to find jobs? This example illustrates the fact that by the time an objectively optimal decision can be computed, we will have died long ago, or at the least, wasted a significant amount of life. As a result, because of the time it would take to compute optimal decisions, (Platonically-) optimal decisions become objectively suboptimal. The distinction between Platonically-optimal and objectively-optimal is important because it means that, if heuristic use is objectively-optimal rather than Platonically-optimal, then expected utility theory is still a valid and plausible framework for the philosophy of decision-making.

We can investigate resource-, bounded-, and ecological-rationality (see the Introduction) through the lens of statistical learning theory (Chowdhary, 2020; Vapnik, 1999), which in turn relies on the efficient coding principle (Attneave, 1954; Barlow, 2013). Both ideas assert that the brain represents information cost-effectively by adapting structural and functional connectivity to statistics of the environment. A Platonically-optimal animal would be able to perfectly assign representations to all their possible outcomes and causes and have accurately assigned weights to each mapping that reflect context-dependent decision-making strategies. On the other hand, biological systems cannot be sensitive to the entire spectrum of sensory information (Bhui et al., 2021). Why is this the case? Quite simply, biological systems do not have a number of sensors equal to the number of possible sensory elements in the world.

Therefore, the brain has to compress the sensory data to make sense of it (Robson, 2001). In other words, statistical learning is a process of information transmission through a limited-capacity communication channel. The efficient encoding scheme is often formalised as a problem of dimensionality reduction (Churchland et al., 2006, 2010; Cunningham & Yu, 2014; Semedo et al., 2019) – the curse of dimensionality – which means that the data being encoded has more dimensions than the structure encoding that data, in turn meaning that some data will be lost. Here, the lost data corresponds to levels of uncertainty. Efficient encoding depends on statistical independence (Bhui et al., 2021; Ebitz & Hayden, 2021; Elsayed et al., 2016; Elsayed & Cunningham, 2017): If two dimensions of data are collinear, then the brain only needs to encode one of those dimensions in order to have knowledge about both of them. Collinearity is not the only requirement for statistical independence. In fact, we can think of efficient coding as a regression problem and potentially inefficient dimensions-to-encode as the dependent variables. Interestingly, heuristics can also be thought of as solutions to a regression problem (Summerfield & Parpart, 2022) by discarding decision-relevant information. Specifically, decision-relevant information is discarded by assigning equal weight to each predictor or by upweighting a single (select few) predictor (s). For example, when making a sandwich, one can pile on a favoured ingredient (i.e., use a single weight heuristic) or require equal contributions from all the available ingredients (i.e., use an equal weight heuristic). Either of these options are easier to decide (i.e., require fewer cognitive resources) than creating a sandwich with carefully constructed portions from all the ingredients.

Statistical learning theory is evolutionarily advantageous as well (Lo, 2013). Consider a population of animals that can live in the valley or on the plateau. If it is sunny outside, building a nest in the valley will lead to offspring while building a nest on the plateau will not because it affords no protection from the sunlight. When it rains, the valley floods, leading to drowned offspring and so the animals should build their nests on the plateau. Now say the probability of sunshine is 75% and the probability of rain is 25%. The Platonically-optimal behaviour would be for all the animals to build their nests in the valley. In doing so, the first rain would wipe out the entire population. Therefore, the objectively-optimal behaviour would be to randomly sample the two locations with 75 probability of sampling the valley and 25% probability of sampling the plateau. This behaviour maximises population growth (Lo, 2013) indicating that statistical learning affords evolutionary advantages (Herrnstein, 1961; Lo, 2013). Similarly, heuristics are evolutionarily advantageous (Marsh, 2002) as demonstrated with bounded-rationality in a number of species (Cohen et al., 2019; Marsh, 2002) and explicit use of heuristics such as GUT (Hills, 2006).

There is thus considerable evidence for the theories of bounded-, resource-, and ecological rationality as variants of expected utility theory. Animal decision-making only appears irrational when compared against the Platonically-optimal strategy. Resource and time constraints put limitations on the computational capacity of biological systems and therefore the objectively-optimal decision-making strategy is more dependent on the current state, which is always equipped with some uncertainty. In sum, the use of heuristics is the objectively-optimal (“optimal suboptimal”) solution to uncertainty. While there is ample evidence for this notion, there has been far less attention paid to how this process occurs in the brain. To resolve this dearth, we will next attempt to explain the process in terms of how the brain generalises information from past experience to novel situations.

Use of heuristics as generalisation

Say you move to a new neighbourhood. The first day you are living there, you see your neighbour leaving. Later that day you see a stranger exiting the same house. From these separate experiences, you can infer that the two people might be partners or relatives. By knowing your neighbour lives in the next-door house, you could infer that someone coming from the same house is in some way related to your neighbour. Now you need to find a good restaurant to eat at because you just moved and have not had time to get groceries. You choose a restaurant from a chain that you previously tried and enjoyed. Both of these examples are forms of similarity-based generalisation: because of the overlapping features from disparate experiences, one can transfer information from one to the other. The other form of generalisation is rule-based generalisation, which explains how we can learn to play games and sports. For example, a child can play hopscotch successfully without having any previous experience with the game. This is possible because the child can generalise the rule across time. The rule “hop on one foot when there is a single square” is applied to the first single square encountered and the *n*th single square. Thus, generalising that rule constrains successful behaviour for every single square.

Rule-based generalisation in particular is quite similar to the use of heuristics. Rather than spending time and energy on learning something that you have experienced before, it is efficient to learn a rule that applies to both situations. Indeed, heuristics function to reduce the cognitive resources we need to use when making decisions (Baker & Nofsinger, 2011; Santos & Relajo-Howell, 2020). However, we can learn numerous rules so knowing which rule to apply in a given situation often requires inferring similarity between states or contexts. Now we can re-cast the use of heuristics as generalising a rule through a similarity-defined space. A similarity-defined space corresponds to Shepard’s (1974, 1987) universal law of generalisation, which defines generalisation as an exponentially decaying function of the similarity between stimuli. Here, similarity is measured as distance in a (multidimensional) psychological space. In graph theoretic terms, stimuli are nodes and edges are possible transitions between stimuli. This is a useful framework, for it affords talk of information diffusion along the structural connectivity of the graph. Generalisation thus corresponds to diffusion from one node through its edges and to adjacent nodes.

Further, graph theory models enable one to adjust edge quality, such as making them descending, ascending, rough, or smooth, consequently shaping the manner in which information diffuses along those edges. Nascent computational models incorporate all of these factors to capture decision-making in environments with spatially-correlated reward distributions, just like foraging in the wild (Garvert et al., 2021; McNamee et al., 2021; Meder et al., 2021; Schulz et al., 2020; Schulz, Konstantinidis et al., 2018a; Schulz et al., 2018b; Wu et al., 2021, 2022). Specifically, these models combine reinforcement learning (RL) and Gaussian Processes (GP) to integrate reward, and rule- and similarity-based generalisation. GPs sample functions from a multivariate distribution with a mean vector and a covariance function (Schulz et al., 2018c).

In this case, making the covariance function a radial basis function kernel means the correlations between stimuli decrease exponentially with distance in psychological space. Thus, a GP can be used to predict reward distributions according to how stimuli correlate. Further, combining the GP with an Upper Confidence Bound (UCB) can predict decisions that account for the uncertainty in the environment. Specifically, a UCB adds the variance of a subjective reward estimate to the subjective reward estimate, making a state’s overall value an additive function of generalised reward and uncertainty. A UCB is also known as a type of exploration bonus (Garivier & Moulines, 2011).

An exploration bonus embodies the idea that it is inherently rewarding to resolve uncertainty. Exploration bonuses are called as such in RL literature, though it also goes by exploratory motivation (Houser et al., 2022; Murty & Adcock, 2017), intrinsic motivation (Leotti & Delgado, 2011), curiosity drive (Loewenstein, 1994), information-seeking (Gottlieb et al., 2013), and epistemic foraging (Kruglanski et al., 2020). Particularly intriguing is the fact that the mesolimbic dopaminergic neural circuitry underpins exploration bonuses (Leotti & Delgado, 2011; Murty et al., 2015, 2018) because this means that evolution co-opted the brain's reward system to motivate resolving uncertainty. One limitation to this view is the neurochemistry behind the explore-exploit tradeoff. That is, dopaminergic release is most often associated with exploitation while noradrenergic release is associated with exploration (Daw et al., 2006;

Dubois et al., 2021; Hills, 2006), suggesting that the brain's reward system should not be active when dealing with uncertainty. Indeed, rodent studies show that dopamine antagonists induce more random decision-making (Cinotti et al., 2019). On the other hand, this conundrum licenses talk of a more complex neurobiological substrate mediating the explore-exploit tradeoff. One possibility is differentiating between phasic and tonic catecholamine release (see (Beeler et al., 2010; Bromberg-Martin et al., 2010; Dubois et al., 2021; Ellwood et al., 2017; Schultz et al., 1997).

In sum, the GP-UCB model discussed above embodies the idea that generalised information guides decision-making in the face of uncertainty. There is a wealth of fear generalisation studies that support this claim (Asok et al., 2019; Dunsmoor et al., 2009, 2011; Dunsmoor & Paz, 2015; Dymond et al., 2015; Onat & Büchel, 2015; Struyf et al., 2017). For example, after threat conditioning, stimuli related to conditioned stimuli will be more likely to evoke a fear response (Starita et al., 2019). This is a highly adaptive mechanism: If I go hiking and encounter a dangerous animal, I will be more cautious the next time hiking that trail. Like substance abuse and problem gambling, however, maladaptive responses can form as well, such as overgeneralisation of a fear response. Regardless, we know that generalisation is a primary response to uncertainty.

The final piece of the puzzle involves how information generalises to particular variables, states, or nodes. According to the GP-UCB model, there needs to be a network graph that determines how information diffuses. One might imagine this network graph as the irrigation channels that guide the flow of water. Here, we suggest that that network graph is implemented in the hippocampus. The hippocampus has a long history of being involved in episodic memory and spatial navigation, however, only recently has it been realised that it plays a perhaps fundamental role in generalisation and abstraction as well (Dusek & Eichenbaum, 1997; Eichenbaum, 2004, 2017c, 2017a, 2017b; Zeithamova et al., 2012, 2019). Specifically, it seems to be the case that the hippocampus cares more about associations in general than particular instances of experience, as predicted by the hippocampal index theory (Goode et al., 2020; Teyler & DiScenna, 1986; Teyler & Rudy, 2007). This is supported by the hippocampus' involvement in spatial and non-spatial foraging (Bellmund et al., 2018; Moser et al., 2008).

Similar findings have been reported for the entorhinal cortex (Bao et al., 2019; Bellmund et al., 2016, 2018; Constantinescu et al., 2016; Doeller et al., 2010; Julian & Doeller, 2021; Moon et al., 2020; Nau et al., 2018; Park et al., 2021). A particularly revealing behavioural manipulation conveys this point beautifully. Suzuki et al. (1980) ran rodents in a radial arm maze task (Olton, 1987; Olton & Samuelson, 1976) with reward at the end of each arm and tested them for efficiency in completing the maze. The maze was completed when all rewards were collected, and thus, the most efficient path through the maze is to visit each arm once and none repeatedly. It had previously been shown that rodents could learn this maze, though, most believed that it was because rodents were storing individual arms in memory and thereby keeping track of which ones were already visited (Olton, 1987; Olton & Samuelson, 1976).

To test if this is what was really happening, Suzuki and colleagues attached a visual stimulus to the end of each arm. If rodents were tracking each arm individually, the visual cues should help facilitate this process. That is, they merely need to remember whether they visited the cue at the end of the arm. After learning the maze, rodents were assigned to one of two other mazes: (1) a rotated maze, or (2) a transposed maze. The rotated maze simply rotated each arm 180 degrees such that arm 1 became arm 5, arm 2 became arm 6, and so on. Importantly, the global configuration of cues remained the same: Every arm was related to every other arm in the same way as before. The transposed maze featured a random shuffling of the arms such that the original global configuration was destroyed. Incredibly, rodents had no trouble with the rotated maze, but their performance was seriously

impaired in the transposed maze, indicating that they were not remembering each individual arm but rather the global configuration (i.e., the graph network) of the task space (Nadel & Willner, 1980).

In sum, computational models have provided us with the rationale for how generalisation guides decision-making in the face of uncertainty and a number of empirical studies across multiple species have found neural signatures of heuristics and generalisation that overlap heavily (Abivardi et al., 2020; Goel et al., 2004; Korn & Bach, 2018, 2019; Krause & Drugowitsch, 2022; Wimmer & Shohamy, 2012). This idea—to our knowledge—has never been explicitly put forth so we hope to unite future research forces in verifying some of the broad implications laid out in this article.

Concluding remarks

For thousands of years, cognitive science has sought after Platonic ideals that have little bearing on the real world. One of those ideals is (Platonic) optimality in decision-making. Platonically-optimal decision-making assumes that one can make decisions with full knowledge of all possible outcomes. While this may indeed be possible theoretically, it does not offer an objectively-optimal decision-making strategy because the time it would take to compute Platonic optimality is far too costly for finite, resource-dependent beings. Thus, deviations from Platonic optimality are in fact a rational decision-making strategy. This stands in contrast to popular theories of decision-making that cast animals as highly irrational due to their failure to make decisions that coincide with Platonic optimality. The use of heuristics, which—ironically—was once the evidence for irrational decision-making, has been theorised to be a rational solution to finite resources (Bhui et al., 2021; Milli et al., 2021).

In this paper, we attempted to extend this theory by claiming that heuristics are a rational solution to uncertainty in general. We used work based on the exploration-exploitation dilemma to bolster this claim. Specifically, exploration is a venture into the unknown, and is thus always equipped with uncertainty while exploitation encompasses perseveration on known information. Consequently, the use of heuristics may be a hallmark of (epistemic) exploration. We also suggested that irregular sensitivity to uncertainty can lead to their misuse which can manifest in maladaptive exploitation such as substance abuse and problem gambling.

Next, we went into more detail about the difference between Platonic and objective optimality, using robust computational principles of neural encoding to substantiate this claim. Overall, objective optimality offers a more ecologically valid framework for biological-based decision-making. Finally, we introduced a novel take on the use of heuristics, specifically arguing that they can be guided by an abstract structural representation of task space instantiated in the hippocampus and surrounding cortices. We ground this claim by relating the characteristics of heuristics to generalisation processes that are already known to be a solution to uncertainty.

Moreover, we laid the foundations for the computational architecture underpinning generalisation-guided decision-making in states of uncertainty. This formalisation provides a tool with which we can evaluate explore-exploit tradeoffs in individuals. For example, using maximum likelihood techniques to find the best fitting parameters for the GP kernel, the exploration bonus, and the sampling strategy (e.g., inverse temperature for soft maximisation) can reveal how much generalisation, exploration, and exploitation, respectively, are weighted during decision-making (Wu et al., 2018). Generalisation-guided decision-making is an evolutionarily-advantageous strategy for searching, planning, and navigating that has endured across the ages. Moreover, it is a single phenomenon that requires both decision-making and generalisation, and thus, has the potential to bridge the gap that exists between these two literatures.

Acknowledgements: None declared

Conflict of interests: The author declare no competing interests.

Ethical approval: Not applicable

Funding: No funding was received.

REFERENCES

- Abivardi, A., Khemka, S., & Bach, D. R. (2020). Hippocampal representation of threat features and behaviour in a human approach–avoidance conflict anxiety task. *Journal of Neuroscience*, 40(35). <https://doi.org/10.1523/JNEUROSCI.2732-19.2020>
- Addicott, M. A., Pearson, J. M., Sweitzer, M. M., Barack, D. L., & Platt, M. L. (2017). A primer on foraging and the explore/exploit trade-off for psychiatry research. *Neuropsychopharmacology*, 42(10). <https://doi.org/10.1038/npp.2017.108>
- Ariely, D. (2009). The end of rational economics. *Harvard Business Review*, 87(7).
- Asok, A., Kandel, E. R., & Rayman, J. B. (2019). The neurobiology of fear generalisation. *Frontiers in Behavioural Neuroscience*, 12. <https://doi.org/10.3389/fnbeh.2018.00329>
- Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review*, 61(3). <https://doi.org/10.1037/h0054663>
- Baker, H. K., & Nofsinger, J. R. (2011). Behavioural finance: An overview. Behavioural finance: investors, corporations, and markets. <https://doi.org/10.1002/9781118258415.ch1>
- Bao, X., Gjorgieva, E., Shanahan, L. K., Howard, J. D., Kahnt, T., & Gottfried, J. A. (2019). Grid-like neural representations support olfactory navigation of a two-dimensional odour space. *Neuron*, 102(5), 1066–1075.e5. <https://doi.org/10.1016/j.neuron.2019.03.034>
- Barlow, H. B. (2013). Possible principles underlying the transformations of sensory messages. Sensory Communication. <https://doi.org/10.7551/mitpress/9780262518420.003.0013>
- Baron, R. A. (1998). Cognitive mechanisms in entrepreneurship. *Journal of Business Venturing*, 13(4). [https://doi.org/10.1016/s0883-9026\(97\)00031-1](https://doi.org/10.1016/s0883-9026(97)00031-1)
- Beeler, J. A., Daw, N., Frazier, C. R. M., & Zhuang, X. (2010). Tonic dopamine modulates exploitation of reward learning. *Frontiers in Behavioural Neuroscience*. <https://doi.org/10.3389/fnbeh.2010.00170>
- Bellmund, J. L. S., Deuker, L., Schröder, T. N., & Doeller, C. F. (2016). Grid-cell representations in mental simulation. *ELife*. <https://doi.org/10.7554/eLife.17089>
- Bellmund, J. L. S., Gärdenfors, P., Moser, E. I., & Doeller, C. F. (2018). Navigating cognition: Spatial codes for human thinking. *Science*, 362(6415). <https://doi.org/10.1126/science.aat6766>
- Bhui, R., Lai, L., & Gershman, S. J. (2021). Resource–rational decision making. *Current Opinion in Behavioral Sciences*, 41. <https://doi.org/10.1016/j.cobeha.2021.02.015>
- Bottom, W. P., Gilovich, T., Griffin, D., & Kahneman, D. (2004). Heuristics and biases: The Psychology of intuitive judgment. *Academy of Management Review*, 29(4). <https://doi.org/10.2307/20159081>
- Bowman, C. R., Iwashita, T., & Zeithamova, D. (2020). Tracking prototype and exemplar representations in the brain across learning. *ELife*, 9, 1–47. <https://doi.org/10.7554/eLife.59360>
- Bowman, C. R., & Zeithamova, D. (2018). Abstract memory representations in the ventromedial prefrontal cortex and hippocampus support concept generalisation. *Journal of Neuroscience*, 38(10), 2605–2614. <https://doi.org/10.1523/JNEUROSCI.2811-17.2018>
- Bromberg–Martin, E. S., Matsumoto, M., & Hikosaka, O. (2010). Dopamine in motivational control: Rewarding, aversive, and alerting. *Neuron*, 68(5). <https://doi.org/10.1016/j.neuron.2010.11.022>
- Chowdhary, K. R. (2020). Statistical learning theory. Fundamentals of artificial intelligence. https://doi.org/10.1007/978-81-322-3972-7_14
- Churchland, M. M., Santhanam, G., & Shenoy, K. v. (2006). Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. *Journal of Neurophysiology*, 96(6). <https://doi.org/10.1152/jn.00307.2006>
- Churchland, M. M., Yu, B. M., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., Newsome, W. T., Clark, A. M., Hosseini, P., Scott, B. B., Bradley, D. C., Smith, M. A., Kohn, A., Movshon, J. A., Armstrong, K. M., Moore, T., Chang, S. W., Snyder, L. H., Lisberger, S. G., Shenoy, K. V. (2010). Stimulus onset quenches neural variability: A widespread cortical phenomenon. *Nature Neuroscience*, 13(3). <https://doi.org/10.1038/nn.2501>
- Cinotti, F., Fresno, V., Aklil, N., Coutureau, E., Girard, B., Marchand, A. R., & Khamassi, M. (2019). Dopamine blockade impairs the exploration–exploitation trade-off in rats. *Scientific Reports*, 9(1). <https://doi.org/10.1038/s41598-019-43245-z>
- Cohen, D., Teichman, G., Volovich, M., Zeevi, Y., Elbaum, L., Madar, A., Louie, K., Levy, D. J., & Rechavi, O. (2019). Bounded rationality in *C. elegans* is explained by circuit-specific normalisation in chemosensory pathways. *Nature Communications*, 10(1). <https://doi.org/10.1038/s41467-019-11715-7>
- Constantinescu, A. O., O'Reilly, J. X., & Behrens, T. E. J. (2016). Organising conceptual knowledge in

- humans with a grid like code. *Science*, 352(6292). 1464–1468.
<https://doi.org/10.1126/science.aaf0941>
- Cunningham, J. P., & Yu, B. M. (2014). Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, 17(11). <https://doi.org/10.1038/nn.3776>
- Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095). <https://doi.org/10.1038/nature04766>
- Dezza, I. C., Noel, X., Cleeremans, A., & Yu, A. J. (2021). Distinct motivations to seek out information in healthy individuals and problem gamblers. *Translational Psychiatry*, 11(1).
<https://doi.org/10.1038/s41398-021-01523-3>
- Doeller, C. F., Barry, C., & Burgess, N. (2010). Evidence for grid cells in a human memory network. *Nature*, 463(7281), 657–661. <https://doi.org/10.1038/nature08704>
- Dubois, M., Habicht, J., Michely, J., Moran, R., Dolan, R. J., & Hauser, T. U. (2021). Human complex exploration strategies are enriched by noradrenaline-modulated heuristics. *ELife*, 10.
<https://doi.org/10.7554/eLife.59907>
- Dunsmoor, J. E., Mitroff, S. R., & LaBar, K. S. (2009). Generalisation of conditioned fear along a dimension of increasing fear intensity. *Learning and Memory*, 16(7).
<https://doi.org/10.1101/lm.1431609>
- Dunsmoor, J. E., & Paz, R. (2015). Fear generalisation and anxiety: Behavioural and neural mechanisms. *Biological Psychiatry*, 78(5). <https://doi.org/10.1016/j.biopsych.2015.04.010>
- Dunsmoor, J. E., Prince, S. E., Murty, V. P., Kragel, P. A., & LaBar, K. S. (2011). Neuro behavioural mechanisms of human fear generalisation. *NeuroImage*, 55(4).
<https://doi.org/10.1016/j.neuroimage.2011.01.041>
- Dusek, J. A., & Eichenbaum, H. (1997). The hippocampus and memory for orderly stimulus relations. *Proceedings of the National Academy of Sciences of the United States of America*, 94(13), 7109–7114. <https://doi.org/10.1073/pnas.94.13.7109>
- Dymond, S., Dunsmoor, J. E., Vervliet, B., Roche, B., & Hermans, D. (2015). Fear generalisation in humans: Systematic review and implications for anxiety disorder research. *Behaviour Therapy*, 46(5). <https://doi.org/10.1016/j.beth.2014.10.001>
- Ebitz, R. B., & Hayden, B. Y. (2021). The population doctrine in cognitive neuroscience. *Neuron*, 109(19), 3055–3068. <https://doi.org/10.1016/j.neuron.2021.07.011>
- Edwards, A. M., Phillips, R. A., Watkins, N. W., Freeman, M. P., Murphy, E. J., Afanasyev, V., Buldyrev, S. V., da Luz, M. G. E., Raposo, E. P., Stanley, H. E., & Viswanathan, G. M. (2007). Revisiting Lévy flight search patterns of wandering albatrosses, bumblebees and deer. *Nature*, 449(7165). <https://doi.org/10.1038/nature06199>
- Eichenbaum, H. (2004). Hippocampus: Cognitive processes and neural representations that underlie declarative memory. *Neuron*, 44(1), 109–120. <https://doi.org/10.1016/j.neuron.2004.08.028>
- Eichenbaum, H. (2017a). On the integration of space, time, and memory. *Neuron*, 95(5), 1007–1018.
<https://doi.org/10.1016/j.neuron.2017.06.036>
- Eichenbaum, H. (2017b). The role of the hippocampus in navigation is memory. *Journal of Neurophysiology*, 117(4), 1785–1796. <https://doi.org/10.1152/jn.00005.2017>
- Eichenbaum, H. (2017c). Time (and space) in the hippocampus. *Current Opinion in Behavioural Sciences*, 17, 65–70. <https://doi.org/10.1016/j.cobeha.2017.06.010>
- Ellwood, I. T., Patel, T., Wadia, V., Lee, A. T., Liptak, A. T., Bender, K. J., & Sohal, V. S. (2017). Tonic or phasic stimulation of dopaminergic projections to prefrontal cortex causes mice to maintain or deviate from previously learned behavioural strategies. *Journal of Neuroscience*, 37(35).
<https://doi.org/10.1523/JNEUROSCI.1221-17.2017>
- Elsayed, G. F., & Cunningham, J. P. (2017). Structure in neural population recordings: An expected by-product of simpler phenomena? *Nature Neuroscience*, 20(9).
<https://doi.org/10.1038/nn.4617>
- Elsayed, G. F., Lara, A. H., Kaufman, M. T., Churchland, M. M., & Cunningham, J. P. (2016). Reorganisation between preparatory and movement population responses in motor cortex. *Nature Communications*, 7(1). <https://doi.org/10.1038/ncomms13239>
- Garami, J., Haber, P., Myers, C. E., Allen, M. T., Misiak, B., Frydecka, D., & Moustafa, A. A. (2017). Intolerance of uncertainty in opioid dependency – Relationship with trait anxiety and impulsivity. *PLoS ONE*, 12(7). <https://doi.org/10.1371/journal.pone.0181955>
- Garivier, A., & Moulines, E. (2011). On upper-confidence bound policies for switching bandit problems. Lecture notes in computer science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 6925 LNAI. https://doi.org/10.1007/978-3-642-24412-4_16
- Garvert, M. M., Saanum, T., Schulz, E., Schuck, N. W., & Doeller, C. F. (2021). Hippocampal spatio-temporal cognitive maps adaptively guide reward generalisation. *BioRxiv*.

- Gershman, S. J. (2018). Deconstructing the human algorithms for exploration. *Cognition*, 173. <https://doi.org/10.1016/j.cognition.2017.12.014>
- Gershman, S. J. (2019). Uncertainty and exploration. *Decision*, 6(3). <https://doi.org/10.1037/dec0000101>
- Gershman, S. J., & Uchida, N. (2019). Believing in dopamine. *Nature Reviews Neuroscience*, 20(11). <https://doi.org/10.1038/s41583-019-0220-7>
- Goel, V., Makale, M., & Grafman, J. (2004). The hippocampal system mediates logical reasoning about familiar spatial environments. *Journal of Cognitive Neuroscience*, 16(4). <https://doi.org/10.1162/089892904323057362>
- Goode, T. D., Tanaka, K. Z., Sahay, A., & McHugh, T. J. (2020). An integrated index: Engrams, place cells, and hippocampal memory. *Neuron*, 107(5)805–820. <https://doi.org/10.1016/j.neuron.2020.07.011>
- Gottlieb, J., Oudeyer, P. Y., Lopes, M., & Baranes, A. (2013). Information-seeking, curiosity, and attention: Computational and neural mechanisms. *Trends in Cognitive Sciences*, 17(11). <https://doi.org/10.1016/j.tics.2013.09.001>
- Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behaviour*, 4(3). <https://doi.org/10.1901/jeab.1961.4-267>
- Hills, T. T. (2006). Animal foraging and the evolution of goal-directed cognition. *Cognitive Science*, 30(1). https://doi.org/10.1207/s15516709cog0000_50
- Holcomb, T. R., Ireland, R. D., Holmes, R. M., & Hitt, M. A. (2009). Architecture of entrepreneurial learning: Exploring the link among heuristics, knowledge, and action. *Entrepreneurship: Theory and Practice*, 33(1). <https://doi.org/10.1111/j.1540-6520.2008.00285.x>
- Houser, T. M., Tompary, A., & Murty, V. (2022). Agency enhances temporal order memory in an interactive exploration game. *PsyArXiv*.
- Humphries, N. E., Weimerskirch, H., Queiroz, N., Southall, E. J., & Sims, D. W. (2012). Foraging success of biological Lévy flights recorded in situ. *Proceedings of the National Academy of Sciences of the United States of America*, 109(19). <https://doi.org/10.1073/pnas.1121201109>
- Julian, J. B., & Doeller, C. F. (2021). Remapping and realignment in the human hippocampal formation predict context-dependent spatial behaviour. *Nature Neuroscience*, 24(6), 863–872. <https://doi.org/10.1038/s41593-021-00835-3>
- Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioural economics. *American Economic Review*, 93(5). <https://doi.org/10.1257/00028280322655392>
- Kahneman, D., & Tversky, A. (1979). Prospect theory - An analysis of decision under risk. *Econometrica*, 47.
- Korn, C. W., & Bach, D. R. (2018). Heuristic and optimal policy computations in the human brain during sequential decision-making. *Nature Communications*, 9(1). <https://doi.org/10.1038/s41467-017-02750-3>
- Korn, C. W., & Bach, D. R. (2019). Minimising threat via heuristic and optimal policies recruits hippocampus and medial prefrontal cortex. *Nature Human Behaviour*, 3(7). <https://doi.org/10.1038/s41562-019-0603-9>
- Krause, E. L., & Drugowitsch, J. (2022). A large majority of awake hippocampal sharp-wave ripples feature spatial trajectories with momentum. *Neuron*, 110(4). <https://doi.org/10.1016/j.neuron.2021.11.014>
- Krebs, J. R., Kacelnik, A., & Taylor, P. (1978). Test of optimal sampling by foraging great tits. *Nature*, 275(5675). <https://doi.org/10.1038/275027a0>
- Kruglanski, A. W., Jasko, K., & Friston, K. (2020). All thinking is ‘Wishful’ thinking. *Trends in Cognitive Sciences*, 24(6). <https://doi.org/10.1016/j.tics.2020.03.004>
- Ladouceur, R., Blais, F., Freeston, M. H., & Dugas, M. J. (1998). Problem solving and problem orientation in generalised anxiety disorder. *Journal of Anxiety Disorders*, 12(2). [https://doi.org/10.1016/S0887-6185\(98\)00002-4](https://doi.org/10.1016/S0887-6185(98)00002-4)
- Leotti, L. A., & Delgado, M. R. (2011). The inherent reward of choice. *Psychological Science*, 22(10). <https://doi.org/10.1177/0956797611417005>
- Lo, A. W. (2013). The origin of bounded rationality and intelligence. *Proceedings of the American Philosophical Society*, 157(3).
- Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*, 116(1). <https://doi.org/10.1037//0033-2909.116.1.75>
- Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A Network Model of Category Learning. *Psychological Review*, 111(2). <https://doi.org/10.1037/0033-295X.111.2.309>
- Marsh, B. (2002). Do animals use heuristics? *Journal of Bioeconomics*, 4(1). <https://doi.org/10.1023/A:1020655022163>

- McGuire, W. J. (1997). Creative hypothesis generating in psychology: Some useful heuristics. *Annual Review of Psychology*, 48. <https://doi.org/10.1146/annurev.psych.48.1.1>
- McNamee, D. C., Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2021). Flexible modulation of sequence generation in the entorhinal–hippocampal system. *Nature Neuroscience*, 24(6). <https://doi.org/10.1038/s41593-021-00831-7>
- Meder, B., Wu, C. M., Schulz, E., & Ruggeri, A. (2021). Development of directed and random exploration in children. *Developmental Science*, 24(4). <https://doi.org/10.1111/desc.13095>
- Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., Hausmann, D., Fiedler, K., & Gonzalez, C. (2015). Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, 2(3). <https://doi.org/10.1037/dec0000033>
- Milli, S., Lieder, F., & Griffiths, T. L. (2021). A rational reinterpretation of dual–process theories. *Cognition*, 217. <https://doi.org/10.1016/j.cognition.2021.104881>
- Moon, H. J., Gauthier, B., Park, H. D., Faivre, N., & Blanke, O. (2020). Sense of self impacts spatial navigation and hexadirectional coding in human entorhinal cortex. *BioRxiv*, 41, 1–39. <https://hal.archives-ouvertes.fr/hal-03215424> <https://hal.archives-ouvertes.fr/hal-03215424/document>
- Moser, E. I., Kropff, E., & Moser, M. B. (2008). Place cells, grid cells, and the brain’s spatial representation system. *Annual Review of Neuroscience*, 31. <https://doi.org/10.1146/annurev.neuro.31.061307.090723>
- Murty, V. P., & Adcock, R. A. (2017). Distinct medial temporal lobe network states as neural contexts for motivated memory formation. The hippocampus from cells to systems: Structure, connectivity, and functional contributions to memory and flexible cognition. https://doi.org/10.1007/978-3-319-50406-3_15
- Murty, V. P., DuBrow, S., & Davachi, L. (2015). The simple act of choosing influences declarative memory. *Journal of Neuroscience*, 35(16). <https://doi.org/10.1523/JNEUROSCI.4181-14.2015>
- Murty, V. P., Dubrow, S., & Davachi, L. (2018). Decision–making increases episodic memory via post encoding consolidation. *Journal of Cognitive Neuroscience*, 31(9). https://doi.org/10.1162/jocn_a_01321
- Nadel, L., & Willner, J. (1980). Context and conditioning: A place for space. *Physiological Psychology*, 8(2), 218–228. <https://doi.org/10.3758/BF03332853>
- Nau, M., Navarro Schröder, T., Bellmund, J. L. S., & Doeller, C. F. (2018). Hexadirectional coding of visual space in human entorhinal cortex. *Nature Neuroscience*, 21(2), 188–190. <https://doi.org/10.1038/s41593-017-0050-8>
- Nosofsky, R. M. (1992). Exemplar, prototypes, and similarity rules. From Learning Theory to Connectionist Theory: Essays in honour of William K. Estes.
- Nouri, P., Imanipour, N., Talebi, K., & Zali, M. (2017). Heuristics and biases in entrepreneurial marketing: Some new insights. *International Journal of Entrepreneurship*, 21(3).
- Olton, D. S. (1987). The radial arm maze as a tool in behavioural pharmacology. *Physiology and Behavior*, 40(6), 793–797. [https://doi.org/10.1016/0031-9384\(87\)90286-1](https://doi.org/10.1016/0031-9384(87)90286-1)
- Olton, D. S., & Samuelson, R. J. (1976). Remembrance of places passed: Spatial memory in rats. *Journal of Experimental Psychology: Animal Behaviour Processes*, 2(2), 97–116. <https://doi.org/10.1037/0097-7403.2.2.97>
- Onat, S., & Büchel, C. (2015). The neuronal basis of fear generalisation in humans. *Nature Neuroscience*, 18(12). <https://doi.org/10.1038/nn.4166>
- Optimal Strategies and Heuristics for Ecological Search Problems. (2019). *Cognitive Search*. <https://doi.org/10.7551/mitpress/9529.003.0026>
- Park, S. A., Miller, D. S., & Boorman, E. D. (2021). Inferences on a multidimensional social hierarchy use a grid-like code. *Nature Neuroscience*, 24(9), 1292–1301. <https://doi.org/10.1038/s41593-021-00916-3>
- Radell, M. L., Allen, M. T., Favaloro, B., Myers, C. E., Haber, P., Morley, K., & Moustafa, A. A. (2018). Intolerance of uncertainty and conditioned place preference in opioid addiction. *Peer Journal*, 5. <https://doi.org/10.7717/peerj.4775>
- Raichlen, D. A., & Polk, J. D. (2013). Linking brains and brawn: Exercise and the evolution of human neurobiology. *Royal Society Biological Sciences*, 280(1750). <https://doi.org/10.1098/rspb.2012.2250>
- Santos, A. & Relajo-Howell, D. (2020). Lifestyle and cognitive functioning of Filipino older adults as basis for cognitive enhancement programme. *Psychology & Society*, 4(82), 97–105. <https://doi.org/10.35774/pis2020.04.097>
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306). <https://doi.org/10.1126/science.275.5306.1593>
- Schulz, E., Franklin, N. T., & Gershman, S. J. (2020). Finding structure in multi-armed bandits.

- Cognitive Psychology*, 119. <https://doi.org/10.1016/j.cogpsych.2019.101261>
- Schulz, E., Konstantinidis, E., & Speekenbrink, M. (2018a). Putting bandits into context: How function learning supports decision making. *Journal of Experimental Psychology: Learning Memory and Cognition*, 44(6). <https://doi.org/10.1037/xlm0000463>
- Schulz, E., Speekenbrink, M., & Krause, A. (2018c). A tutorial on Gaussian process regression: Modelling, exploring, and exploiting functions. *Journal of Mathematical Psychology*, 85. <https://doi.org/10.1016/j.jmp.2018.03.001>
- Schulz, E., Wu, C. M., Huys, Q. J. M., Krause, A., & Speekenbrink, M. (2018b). Generalisation and search in risky environments. *Cognitive Science*, 42(8). <https://doi.org/10.1111/cogs.12695>
- Semedo, J. D., Zandvakili, A., Machens, C. K., Yu, B. M., & Kohn, A. (2019). Cortical Areas Interact through a Communication Subspace. *Neuron*, 102(1). <https://doi.org/10.1016/j.neuron.2019.01.026>
- Shepard, R. N. (1974). Representation of structure in similarity data: Problems and prospects. *Psychometrika*, 39(4). <https://doi.org/10.1007/BF02291665>
- Shepard, R. N. (1987). Toward a universal law of generalisation for psychological science. *Science*, 237(4820). <https://doi.org/10.1126/science.3629243>
- Shepherd, D. A., & Zacha, A. (2001). The venture capitalist-entrepreneur relationship: Control, trust and confidence in co-operative behaviour? *Venture Capital*, 3(2). <https://doi.org/10.1080/13691060110042763>
- Sims, D. W., Humphries, N. E., Bradford, R. W., & Bruce, B. D. (2012). Lévy flight and Brownian search patterns of a free-ranging predator reflect different prey field characteristics. *Journal of Animal Ecology*, 81(2). <https://doi.org/10.1111/j.1365-2656.2011.01914.x>
- Sims, D. W., Reynolds, A. M., Humphries, N. E., Southall, E. J., Wearmouth, V. J., Metcalfe, B., & Twitchett, R. J. (2014). Hierarchical random walks in trace fossils and the origin of optimal search behaviour. *Proceedings of the National Academy of Sciences of the United States of America*, 111(30). <https://doi.org/10.1073/pnas.1405966111>
- Song, M., Jones, C. E., Monfils, M.-H., & Niv, Y. (2022). Explaining the effectiveness of fear extinction through latent-cause inference. *ArXiv*.
- Starita, F., Kroes, M. C. W., Davachi, L., Phelps, E. A., & Dunsmoor, J. E. (2019). Threat learning promotes generalisation of episodic memory. *Journal of Experimental Psychology: General*, 148(8). <https://doi.org/10.1037/xge0000551>
- Struyf, D., Zaman, J., Hermans, D., & Vervliet, B. (2017). Gradients of fear: How perception influences fear generalisation. *Behaviour Research and Therapy*, 93. <https://doi.org/10.1016/j.brat.2017.04.001>
- Summerfield, C., & Parpart, P. (2022). Normative principles for decision-making in natural environments. *Annual Review of Psychology*, 73. <https://doi.org/10.1146/annurev-psych-020821-104057>
- Suzuki, S., Augerinos, G., & Black, A. H. (1980). Stimulus control of spatial behaviour on the eight-arm maze in rats. *Learning and Motivation*, 11(1), 1–18. [https://doi.org/10.1016/0023-9690\(80\)90018-1](https://doi.org/10.1016/0023-9690(80)90018-1)
- Taylor, J. E., Cortese, A., Barron, H. C., Pan, X., Sakagami, M., & Zeithamova, D. (2021). How do we generalise? *Neurons, Behaviour, Data Analysis, and Theory*. <https://doi.org/10.51628/001c.27687>
- Teyler, T. J., & DiScenna, P. (1986). The hippocampal memory indexing theory. *Behavioural Neuroscience*, 100(2), 147–154. <https://doi.org/10.1037/0735-7044.100.2.147>
- Teyler, T. J., & Rudy, J. W. (2007). The hippocampal indexing theory and episodic memory: Updating the index. *Hippocampus*, 17(12), 1158–1169. <https://doi.org/10.1002/hipo.20350>
- Thompson, W. R. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4). <https://doi.org/10.2307/2332286>
- Todd, P. M., & Gigerenzer, G. (2000). Précis of Simple heuristics that make us smart. *Behavioral and Brain Sciences*, 23(5). <https://doi.org/10.1017/S0140525X00003447>
- Todd, P. M., & Gigerenzer, G. (2007). Environments that make us smart. *Current Directions in Psychological Science*, 16(3). <https://doi.org/10.1111/j.1467-8721.2007.00497.x>
- Tomov, M. S., Yagati, S., Kumar, A., Yang, W., & Gershman, S. J. (2020). Discovery of hierarchical representations for efficient planning. *PLoS Computational Biology*, 16(4). <https://doi.org/10.1371/journal.pcbi.1007594>
- Vapnik, V. N. (1999). An overview of statistical learning theory. *IEEE Transactions on Neural Networks*, 10(5). <https://doi.org/10.1109/72.788640>
- Viswanathan, G. M., Buldyrev, S. v., Havlin, S., da Luz, M. G. E., Raposo, E. P., & Stanley, H. E. (1999). Optimising the success of random searches. *Nature*, 401(6756). <https://doi.org/10.1038/44831>

- Viswanathan, G. M., Raposo, E. P., & da Luz, M. G. E. (2008). Lévy flights and super diffusion in the context of biological encounters and random searches. *Physics of Life Reviews*, 5(3).
<https://doi.org/10.1016/j.plrev.2008.03.002>
- Wimmer, G. E., & Shohamy, D. (2012). Preference by association: How memory mechanisms in the hippocampus bias decisions. *Science*, 338(6104). <https://doi.org/10.1126/science.1223252>
- Wu, C. M., Schulz, E., & Gershman, S. J. (2021). Inference and search on graph-structured spaces. *Computational Brain and Behaviour*, 4(2). <https://doi.org/10.1007/s42113-020-00091-x>
- Wu, C. M., Schulz, E., Pleskac, T. J., & Speekenbrink, M. (2022). Time pressure changes how people explore and respond to uncertainty. *Scientific Reports*, 12(1). <https://doi.org/10.1038/s41598-022-07901-1>
- Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2018). Generalisation guides human exploration in vast decision spaces. *Nature Human Behaviour*, 2(12).
<https://doi.org/10.1038/s41562-018-0467-4>
- Zacharakis, A. L., & Shepherd, D. A. (2001). The nature of information and overconfidence on venture capitalists' decision making. *Journal of Business Venturing*, 16(4).
[https://doi.org/10.1016/S0883-9026\(99\)00052-X](https://doi.org/10.1016/S0883-9026(99)00052-X)
- Zeithamova, D., Mack, M. L., Braunlich, K., Davis, T., Seger, C. A., Kesteren, M. T. R. V., & Wutz, A. (2019). Brain mechanisms of concept learning. *Journal of Neuroscience*, 39(42).
<https://doi.org/10.1523/JNEUROSCI.1166-19.2019>
- Zeithamova, D., Schlichting, M. L., & Preston, A. R. (2012). The hippocampus and inferential reasoning: Building memories to navigate future decisions. *Frontiers in Human Neuroscience*. <https://doi.org/10.3389/fnhum.2012.00070>