Proceedings Paper:

Jimenez-Rodriguez, A., Prescott, T.J. orcid.org/0000-0003-4927-5390, Schmidt, R. orcid.org/0000-0002-2474-3744 et al. (1 more author) (2020) A framework for resolving motivational conflict via attractor dynamics. In: Vouloutsi, V., Mura, A., Tauber, F., Speck, T., Prescott, T.J. and Verschure, P.F.M.J., (eds.) Biomimetic and Biohybrid Systems : 9th International Conference, Living Machines 2020, Freiburg, Germany, July 28–30, 2020, Proceedings. Living Machines 2020, 28-30 Jul 2020, Freiburg, Germany. Lecture Notes in Computer Science (12413). Springer, Cham , pp. 192-203. ISBN 9783030643126

https://doi.org/10.1007/978-3-030-64313-3_19

A framework for resolving motivational conflict via attractor dynamics

Alejandro Jimenez-Rodriguez$^{1\,[0000-0001-7172-1794]}$, Tony Prescott$^{1\,[0000-0003-4927-5390]}$, Robert Schmidt$^{1\,[0000-0002-2474-3744]}$, and Stuart Wilson$^{1\,[0000-0001-8125-5133]}$

The University of Sheffield, Sheffield, UK
{`a.jimenez-rodriguez, t.j.prescott, robert.schmidt, s.p.wilson`}`@sheffield.ac.uk`

Abstract. Motivation modulates behaviour depending upon contextual and internal cues. Like animals, successful artificial agents must implement different behavioural strategies in order to satisfy dynamical needs. Such causal factors emerge from internal homeostatic or allostatic processes, as well as from external stimuli or threats. However, when two or more needs coalesce, a situation of motivational conflict ensues. In this work we present a four-stage dynamical framework for the resolution of motivational conflict based upon principles from dynamical systems and statistical mechanics. As a central mechanism for the resolution of conflict we propose the use of potentials with multiple wells or minima. This model leads to behavioural switching either by means of a bifurcation or by the stochastic escape from one of the wells. We present analytical and simulation results that reproduce known motivational conflict phenomena observed in the study of animal behaviour, in the case of two conflicting motivations.

Keywords: Motivation · Motivational conflict · dynamical systems · behavioural switching · attractor dynamics

1 Introduction

Dealing with motivational conflict is an important aspect of animal behaviour [19]. Consider, for example, the conflicting need to eat or drink as studied in [15, 16], or the conflict between aggression and mating as exemplified by the response of male stickleback fish to territory invasion by a female conspecific [1]. The former conflict is driven by opposing intrinsic needs or *deficits*, the latter is driven by extrinsic factors, and in both cases conflict resolution requires multiple needs to be satisfied.

Different types of conflict yield different behavioural patterns. Conflict may be resolved via time budgeting, modulated by biological rhythms that occur on different time scales [2]. Or it may be resolved by switching between alternative behaviour systems that may be inhibited or disinhibited by the current behaviour. New patterns of behaviour may emerge (ambivalence), and these may be related (redirection) or unrelated (displacement) to the causes of either motivation [2].

An influential conceptual model of motivation was introduced by Lorenz [10]. Accordingly, independent reservoirs are filled by 'energies' that are action-specific, and when a corresponding energy threshold is exceeded a valve opens to release the corresponding behavioral pattern. Theoretical investigations of motivational conflict often extend this idea to additionally consider direct interactions between motivations via cross-inhibition [12]. Most early investigations incorporated analogies from control theory [18], and in particular used feedback loops to implement Lorenzian energy build-up with competition between motivations driven by internal deficits. Deficits usually inhibit one another directly [11] or bias a decision switch based on 'tendencies' derived from internal state [7]. Such models can be shown to reproduce a wide range of motivational phenomena, but they have been difficult to map directly onto neural systems.

Here we develop a model of how multiple conflicting motivations, each formulated in homeostatic terms, may be resolved to generate appropriate animal (or robot) behaviour. The model is considered first in theoretical terms, and then in terms of simulated robot behaviour.

2 Behaviour under motivational conflict

The proposed model has four stages: Internal physiological state, motivational dynamical system, behavioural selection and pattern expression.

Internal physiological state

We describe the internal physiological state of the agent as a dynamical system. What are referred to in the literature as deficits, are encoded in the state vector $x \in \mathbb{R}^n$. For some motivational systems, e.g., thirst or thermoregulation, the state vector can be associated with a physical quantity such as the amount of water or body heat. For others, such as aggression, it can be related to the accumulation of an action-specific 'energy', in the Lorenzian sense [1]. The internal state vector evolves according to the following dynamical law,

$$\dot{x}(t) = -f(x(t)) + g(u(t)) + h(a(t)). \tag{1}$$

The first term on the right, $-f(x)$, represents the decay of the energy, corresponding to the Lorenzian model, with the function f specifying the nature of the decay (e.g., zero-order, first-order etc.) as well as potential interactions between homeostatic systems. The second term represents an external input $u(t)$, with the function g allowing for a linear or nonlinear transformation of that input (i.e., to represent absorption or thermal conductivity dynamics etc.). The final term represents an autonomic homeostatic process, $a(t)$, with h similarly enabling linear or nonlinear transformations.

The physiological processes we consider are assumed to evolve on slower timescales than that which characterises the behavioural responses, in a close submanifold of $\mathbb{R}^n$, given the existence of physiological limits for all processes, i.e., concentrations can not be negative. To illustrate, consider the following examples.

Simple reservoir. The simplest model is a constant rate of decay of e.g., energy, that can be used to represent *deficits* such as of water or food [7]. Accordingly, the physiological state evolves by $\dot{x}(t) = -\alpha T + u(t)$, where α is the decay rate, T represents some internal autonomic response such as heat generation, and $u(t)$ represents environmental input, e.g., ambient temperature (note that here $f \equiv 0$, $g(u) = u$ and $h(T) = \alpha T$).

Thermoregulation. Consider a recent model of thermoregulatory behaviour proposed by [6], according to which agents are exposed to an ambient temperature x_a and exchange heat upon contact with other agents, $x_c(t)$, resulting in a body temperature, $x(t)$, that evolves according to $\dot{x}(t) = -[k_1 A + k_2(1-A)]x(t) + [k_1 A x_a(t) + k_2(1-A)x_c(t)] + G(t)$. Here, k_1 and k_2 are thermal conductivity constants for the exposed area, A, and the non-exposed area of the body, and $G(t)$ is an autonomic heat generation mechanism that encapsulates different physiological heat sources. Again, the three terms in the right hand side correspond to those of the general model.

Additionally, we propose that at the interface between body and brain i) the state vector first has to be compared to some desired state (or set point) to configure deficits or excesses and respond accordingly. Note that the set point can be variable (allostatic) ii) the corresponding quantities must be normalized and expressed in a *common currency* in order to drive behaviour, and iii) responses to states representing physiological extremes should be differentially weighted to avoid fatal consequences (see [18, 13]). The output of these three transformations is what we call a *drive*, and the component transformations can be expressed as

$$x_d = D(x - x_p), \tag{2}$$

where D is a non-linear map defined as $D = (U \circ N)$, that is, $D(x) = U(N(x))$. The component functions correspond to *normalization*, N, and *urgency*, U. Normalization should limit values to the interval $[0, 1]$, and may be linear (compression of the original domain) or non-linear (e.g., a sigmoidal relationship between physiological state and motivation). Urgency expressions should ensure that extreme values of the physiological range give rise to higher drives than those closer to the set points.

Motivational dynamical system

The motivational state of the agent is modelled as a classical particle undergoing random fluctuations, influenced by the potential $V(x)$. The particle can be thought of as existing in a one-dimensional domain of the *generalized motivational space*, $\mathcal{P}$, which we assume here is equivalent to the real line.

We consider distinct motivations to correspond with specific locations of the phase space $\mathcal{P}$ (Figure 1). The energy landscape provided by the *motivational* *potential* specifies regions of minimal energy that trap the particle for a period in its evolution, with unstable regions serving as barriers. This energy landscape is changed dynamically as a function of the motivational factors, i.e., the drives. More precisely, if $\rho(t)$ represents the position of the particle in the motivational coordinate, its time evolution is given by

$$\dot{\rho} = -\frac{1}{\tau}\frac{\partial V}{\partial \rho} + \sigma dW(t). \tag{3}$$

The first term in the right, $V(\rho, a_m)$, is a potential field that depends upon the drives, and other motivational factors described shortly, and τ is the time constant for the evolution. We formalize the concept of a motivation by restricting $V : \mathcal{P} \times \mathbb{R}^n \to \mathbb{R}$ to be a motivational potential if and only if there exist elements $\bar{\rho}_1, \bar{\rho}_2, \dots, \bar{\rho}_k$, $k > 0$ such that $\frac{\partial V}{\partial \rho}(\rho_k, a_1, \dots, a_k) = 0$ when $a_m = 0$ for all $m = 1, \dots, k$. This characterizes the initial, undisturbed shape with k motivations. The second term to the right in Equation 3 is a noise term with variance, σ^2. We refer to the variance of the fluctuations, analogous to the influence of heat on Brownian motion, as *arousal.* Note that when $\sigma = 0$, $\frac{dV}{dt} = \frac{\partial V}{\partial \rho}\frac{d\rho}{dt} = -\left[\frac{\partial V}{\partial \rho}\right]^2 < 0$. Therefore the dynamics will always tend to select one motivation. Accordingly, the motivational conflict problem is recast in terms of the *escape* or *Kramers* problems that are familiar in classical statistical mechanics [5]. For fixed drives the motivational state, once trapped at point (a) in Figure 1, will eventually escape with a probability determined by the height of the barrier at point (b). The higher the drive, the less likely escape it to occur.

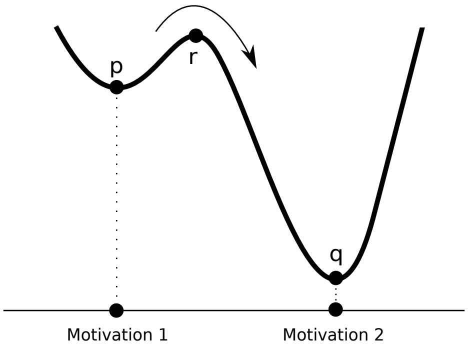

Fig. 1. Concept of motivational potential

Behavioral selection

We conceptualize each motivation as a state. Different readouts of such states produce *tendencies* towards different actions, which we relate, conceptually, to motivation and behavior via *motivational tendency kernels.* Given a motivation

$\bar{\rho}_i$ in the generalized motivational space, a tendency kernel is a function $\xi : \mathcal{P} \to [0,1]$ with finite support, such that $\xi(\bar{\rho}_i) = 1$.

Using, as an example, a *Gaussian tendency kernel*, $\xi(\phi) = G(\rho|\bar{\rho}_i, \sigma^2)$ with a certain width σ (unrelated to the arousal), the incentive will decay with the distance of the system from a given motivation in the phase space. As a second example, for a *step tendency kernel* the incentive will be maximal within a region $a \leq \bar{\rho}_i \leq b$ around the given motivation.

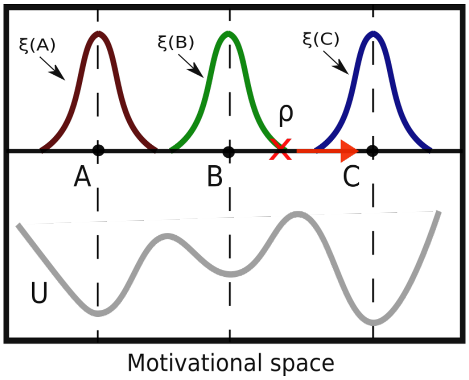

Fig. 2. Illustration of the tendency kernels (top), and the potential U (bottom).

Behavioral pattern expression

The tendencies from the previous section are used to modulate behaviours by biasing action selection and/or unfolding Fixed Action Patterns (FAPs), or by modulating Taxes (i.e. chemotaxis).

3 Recasting motivational conflict

Here we recast some concepts from classic ideas about motivation in the light of the model outlined.

Switching between two motivations

Some of the early research in motivational conflict was related to the switching between two motivated behaviours (i.e. eating and drinking)[7, 14, 16, 17]. The switching from eating to drinking can be thought of as the escape of the particle from the well corresponding to (p) to that corresponding to (q) in Figure 1, effectively jumping the barrier (r). As such, the expected escape time will be the latency of the motivation 2 given that the agent is in p, and the bout duration is the time spent in each well.

For the following analyses we make the simplifying assumption that the drives change on a time scale much slower than that governing the motivational particle ρ, that motivations are read out in terms of non-overlapping step tendency kernels, and that the most salient action will always be performed.

Stationary distribution It is well know from the theory of Stochastic processes that the dynamics of ρ satisfies the Fokker-Plank equation [5]

$$\frac{\partial p(\rho,t)}{\partial t} = \frac{\partial}{\partial \rho}\left[V'(\rho)p(\rho,t)\right] + \sigma^2 \frac{\partial^2 p(\rho,t)}{\partial \rho^2}, \tag{4}$$

where $V' = dV/dx$, and $p(\rho,t)$ is the probability of finding the particle at position ρ at time t. For long enough times, and slowly varying drives (and, therefore, a fixed potential shape), such a probability distribution will reach a stationary value given by

$$p_\infty(\rho) = K \exp\left(-V(\rho)/\sigma^2\right). \tag{5}$$

This captures the intuition that an agent should spend more time in the deepest well. Note that for very small arousal (σ), the agent will remain trapped in that motivation for as long as the shape of the potential is unchanged.

Motivational transition By analogy with the analysis of Kramers for chemical reaction systems [9], the expected exit time of the agent from motivation 1 (p) to motivation 2 (q in figure 1) is given by

$$T(p \to q) = \pi \left[|V''(q)|V''(p)\right]^{1/2} \exp\left\{[V(q) - V(p)]/\sigma^2\right\}. \tag{6}$$

In the absence of arousal, the exit time will be infinity, and two forms of motivational switching emerge. When $\sigma^2 \to 0$, escape becomes improbable and transitions occurs by *competition* [8], defined here as the transition from motivation 1 due to changes in the *causal factors* of motivation 2 (i.e., equivalent to drives). Such a transition occurs when one of the minima is lost due to the increase in the drives for the second motivation. To illustrate, we need to study a specific form of the potential V that is differentiable, for which we choose

$$V(\rho,a,b) = (\bar{\rho}_1 - \rho)^2(\bar{\rho}_2 - \rho)^2 + a(\bar{\rho}_1 - \rho)^2 + b(\bar{\rho}_2 - \rho)^2, \tag{7}$$

where a and b are the corresponding drives for motivation 1 and 2 respectively. For this potential, it can be shown by differentiation that in order for the minima to exist, the following *well relation* must be satisfied:

$$-(\bar{\rho}_1 - \rho)(\bar{\rho}_2 - \rho)\left(\frac{\bar{\rho}_1 + \bar{\rho}_2}{2} - \rho\right) = -\frac{a+b}{2}\rho + \frac{a\bar{\rho}_1 + b\bar{\rho}_2}{2}. \tag{8}$$

As such, wells will exist as long as the line specified by the right hand side intersects with the third-degree polynomial specified on the left hand side (see Figure 3). If we increase the drive (b), the well in the vicinity of the first motivation disappears and the system will tend inexorably to $\bar{\rho}_2$. In the case when drive a is zero and drive b increases, the corresponding minimum will reach the height of the barrier exactly when $b = \frac{1}{2}\left(\frac{\bar{\rho}_1 - \bar{\rho}_2}{2}\right)^2$. In the well relationship, this change is equivalent to changing the slope and intercept of the line (Figure 3).

Note that for $a = b \gg 0$, two wells can merge into one. Indeed, whenever $a = b = \frac{\bar{\rho}_1 + \bar{\rho}_2}{2} - \bar{\rho}_1 \bar{\rho}_2$, the two wells will converge at the mid-way point, leaving only one stable state. In this case the two motivations will have equal tendencies, depending on the form of the tendency kernels, and as such the agent would end up dithering.

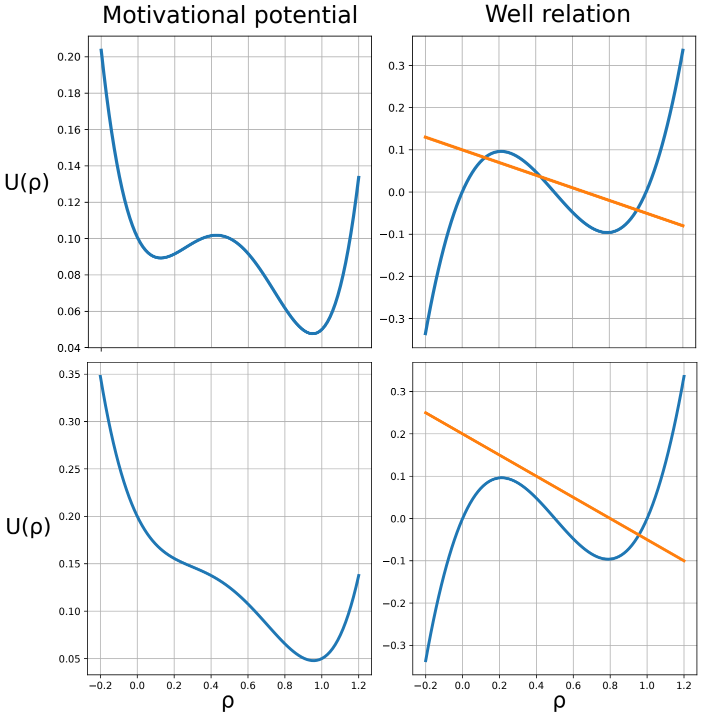

Fig. 3. Top. Existence two wells for $a = 0.05$, $b = 0.1$, $\bar{\rho}_1 = 0$ and $\bar{\rho}_2 = 1$. Bottom. The potential loses one of its minima when the difference between the drives is high enough. The number of minima, and therefore the shape, of the motivational potential is illustrated by the number of points in which the orange line crosses the blue cubic curve.

4 Simulation results

Time-sharing

When the “arousal” term in Equation 6 is non-zero, there is a chance of random transitions between motivations without a corresponding change in the drives. In terms of the motivational literature, these changes can be thought of as time-sharing, where two motivations are observed to switch in bouts of a consistent (average) duration [8], independently from the underlying causal factors.

This can be observed in terms of the potential (Equation 7), when the two wells have the same depth. Figure 4, left, shows that the stationary distribution predicts an equal time spent by the agent in the two motivational states. The corresponding motivational dynamics display the expected switching that occurs more or less frequently depending on the level of arousal, σ. When the heights of the two wells are different, the escape time (and thus the motivational switch) will differ (Equation 6). In Figure 5 a simulated experiment for the transition between two motivations shows how the latency and bout duration change as a function of the shape of the potential. A simulated agent starts in the first well corresponding to e.g., hunger, the escape time for the well marks the switch to e.g., thirst, and the bout duration is the time taken to return to hunger. It can be seen from the simulation that the escape time decreases as the relative height of the two wells increases (as the agent becomes thirsty), as predicted by Equation 6. On the other hand, as the drinking well becomes deeper, it takes a longer time to return to the former motivation. This corresponds to similar phenomena reported in the animal behaviour literature [16].

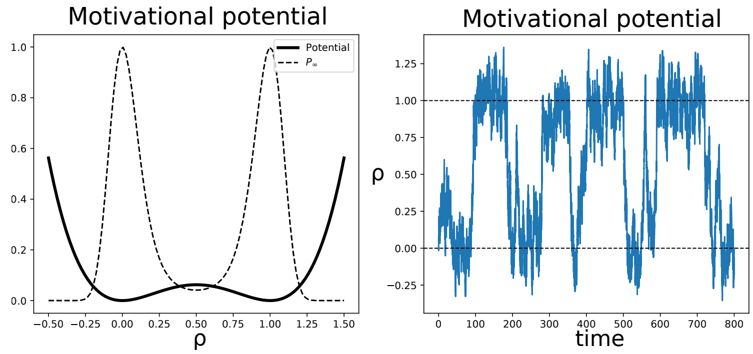

Fig. 4. Time sharing phenomena in the model, $\sigma = 0.12$. Left. The motivational potential (solid) and stationary distribution (dotted) from equation (5). Right. Simulation of the behaviour of the motivational particle switching between the two motivations.

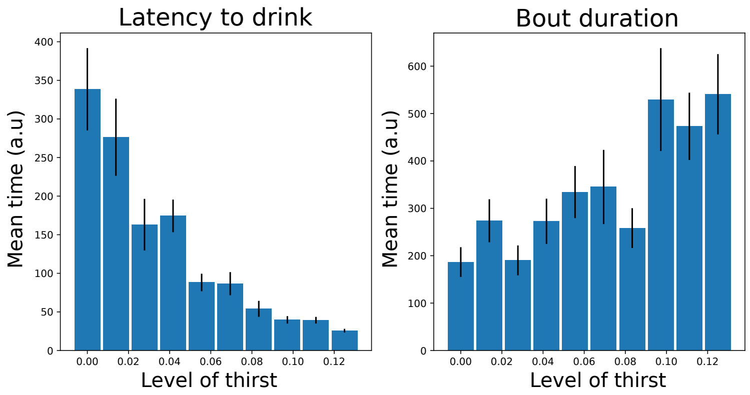

Fig. 5. Left. Latency to motivational switch. Right. Bout duration. $\bar{\rho}_1 = 0$, $\bar{\rho}_2 = 1$, $a = 0$. The value of the drive b assumes values in the interval $[0, 1/8]$, where the upper bound is the theoretical limit at which the first well completely disappears. The escape time is defines as the first time point such $\rho > 0.5$. The bout is assumed to properly start when $\rho > 0.7$. The bout duration is the time it takes to return to $\rho < 0.5$ after the bout has started.

Thermoregulation versus feeding

Here we consider a simulated agent in a two-dimensional (x,y) environment with a chemical (odour) gradient radiating from a food source with a two-dimensional Gaussian profile, and a temperature gradient that is linear in the x coordinate (Figure 6). The physiological state of the agent consists of two homeostatic variables, energy and temperature, that evolve according to the laws presented in section 2. The normalization and urgency maps in Equation 2 are given by sigmoidal and cubic functions. The motivational state is given by the potential in Equation 7 and the parameters a and b are the absolute value of the drives. Motivational kernels for behavioural selection are Gaussian.

We implement two taxis behaviours (and no fixed action patterns) and consider the agent to receive a ‘shot’ of energy when it enters the vicinity of the food source. The agent is modelled as a Braitenberg vehicle [3] with bilateral sensors for the chemical and temperature signals.

The agent and its environment are shown in Figure 6. Initially, the agent is driven towards a region in which its body temperature is maintained in the thermoneutral zone, which creates a conflict because the food is located in a colder region. Two phenomena are observed. First, the agent hesitates to approach the food, and returns multiple times, illustrating a trade-off between the *cost* of obtaining the food and the level of hunger; only when it is hungry enough is the excursion for the food complete. Secondly, between 40 ms and 60 ms the two minima of the motivational potential merge, and the agent is in an *ambiguous* motivational state (comparable with animal behaviour, see [1]).

Spontaneous transitions are observed when the level of arousal is increased (Figure 7 at 20ms), with the motivational state changing in the absence of a

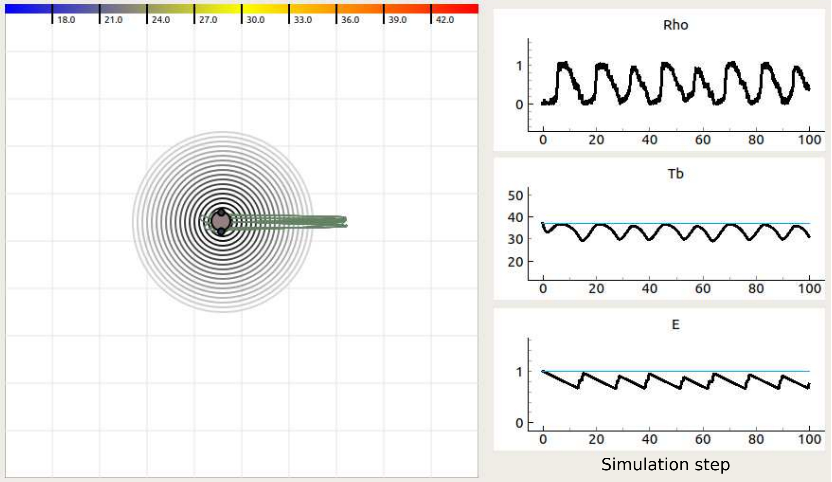

Fig. 6. Simulation environment for the motivated agent. The square to the left represents the 2D arena with a temperature (colored bar at the top) and a chemical (concentric circles in the middle) gradient. The panels to the right display the time course of the motivational dynamical variable ρ (top) and the two homeostatic variables (bottom). The agent moves autonomously back and forth the food source and its preferred spot in the thermal gradient.

corresponding behavioural output. This is comparable with *displacement* phenomena observed in animals, whereby an external cue is shown to drive a novel behaviour that is otherwise unrelated to the current behaviour.

5 Discussion

We developed a four-stage dynamical model to study how behaviours are affected by motivational conflict. The first stage keeps track of the internal state of the agent, the second represents the motivational state, the third contains multiple readouts or measurements in the motivational phase space that give rise to tendencies for certain behaviours, and the fourth stage controls behaviour expression.

A theoretical analysis revealed how the motivational state of the agent can depend on its internal state, allowing us to recast some motivational phenomena described in the animal behaviour literature. For example, we showed how the fixed points of the motivational dynamical system can disappear, appear, or merge, as a function of internal state. We showed how merging of the fixed points can generate the kinds of ambiguous motivational states that have been observed in some animals [1]. And by incorporating stochasticity, we showed how the model can account for phenomena such as time-sharing [7, 14] in terms of spontaneous transitions and disinhibition [15].

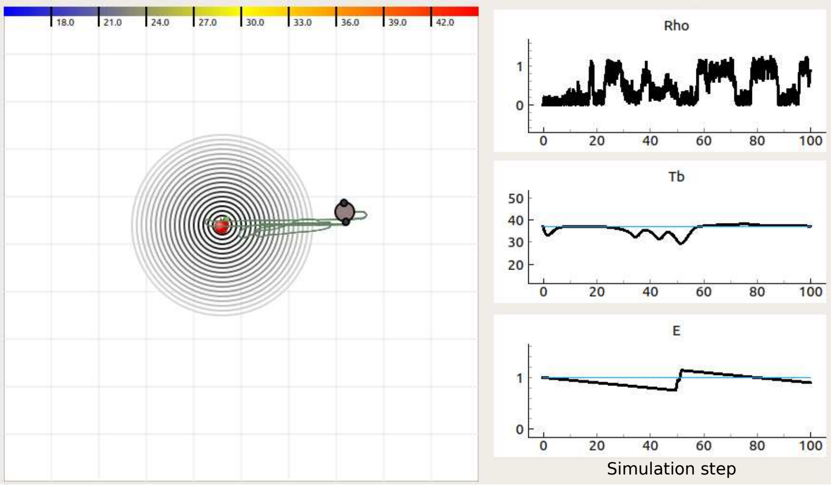

Fig. 7. An increase in arousal increases the rates of spontaneous switching without behavioural output

Our model resembles several classical approaches [19] to this problem. For example, the changes in the underlying physiological dynamics that are described correspond to Lorenzian energy accumulation, but with the valve that releases behaviour recast as a bistability in the motivational state space. However, early models of motivational conflict typically involve positive feedback loops according to a control theory framework [17, 18, 7], and our model does not include explicit feedback loops, with the control instead directly implemented through the effect of the environment on the internal state of the agent.

The stochastic motivational system resembles accumulator models [12]. Indeed, given two stochastic motivations, v_1 and v_2, as in [12], the motivational particle, ρ, could be associated with the difference $v_1 - v_2$ of such accumulators. However, interaction between motivations is assumed not to happen explicitly in our model (i.e. no explicit cross-inhibition), but it is instead assumed to be mediated via an external field.

Our decision to embed motivational attractors in a metric space corresponds to an assumption that the underlying space in which behaviours reside has a definite topology, i.e., that behaviours can be meaningfully ordered and that the ordering determines the interactions between behaviours that may be observed in various conflict scenarios. As such, it should be possible to devise experiments to determine a definite pattern of *displacement* phenomena for a given animal (or species).

Finally, we note that the interaction between the motivational state and the readouts resembles observations of the interactions between the lateral hypothalamus and the ventral tegmental area of the mammalian brain, and that the *tendency* readout could be potentially associated with the dopamine signals that relate to value in these areas [4]. Further connections with the neurobiology of decision-making in animal brains, with a focus on possible relationships with reinforcers and reward, will be explored in future work.

Simulation code is available at https://github.com/ABRG-Models/MammalBot/

Acknowledgments. This work was supported by the EU H2020 Programme as part of the Human Brain Project (HBP-SGA2, 785907).

References

1. Barnard, C.J.: Animal behaviour: mechanism, development, function and evolution. Pearson Education (2004)
2. Bolhuis, J.J., Giraldeau, L.A.E.: The behavior of animals: Mechanisms, function, and evolution. Blackwell Publishing (2005)
3. Braitenberg, V.: Vehicles: Experiments in synthetic psychology. MIT press (1986)
4. Dabney, W., Kurth-Nelson, Z., Uchida, N., Starkweather, C.K., Hassabis, D., Munos, R., Botvinick, M.: A distributional code for value in dopamine-based reinforcement learning. Nature pp. 1–5 (2020)
5. Gardiner, C.W.: Handbook of stochastic methods: for physics, chemistry and the natural sciences (2004)
6. Glancy, J., Groß, R., Stone, J.V., Wilson, S.P.: A self-organising model of thermoregulatory huddling. PLoS computational biology **11**(9) (2015)
7. Houston, A., Sumida, B.: A positive feedback model for switching between two activities. Animal Behaviour **33**(1), 315–325 (1985)
8. Houston, A.I.: Transitions and time-sharing. Animal Behaviour **30**(2), 615–625 (1982)
9. Kramers, H.A.: Brownian motion in a field of force and the diffusion model of chemical reactions. Physica **7**(4), 284–304 (1940)
10. Lorenz, K.Z.: The comparative method in studying innate behavior patterns. (1950)
11. Ludlow, A.: Towards a theory of thresholds. Animal Behaviour **30**(1), 253–267 (1982)
12. Marshall, J.A., Favreau-Peigne, A., Fromhage, L., Mcnamara, J.M., Meah, L.F., Houston, A.I.: Cross inhibition improves activity selection when switching incurs time costs. Current zoology **61**(2), 242–250 (2015)
13. McFarland, D., Bösser, T., Bosser, T.: Intelligent behavior in animals and robots. Mit Press (1993)
14. McFarland, D.: Mechanisms of behavioural disinhibition. Animal Behaviour **17**, 238–242 (1969)
15. McFarland, D., McFarland, F.: Dynamic analysis of an avian drinking response. Medical and biological engineering **6**(6), 659–668 (1968)
16. Roper, T., Crossland, G.: Mechanisms underlying eating-drinking transitions in rats. Animal Behaviour **30**(2), 602–614 (1982)
17. Slater, P.: A simple model for competition between behaviour patterns. Behaviour pp. 236–258 (1978)
18. Toates, F., Archer, J.: A comparative review of motivational systems using classical control theory. Animal Behaviour **26**, 368–380 (1978)
19. Toates, F., Jensen, P.: Ethological and psychological models of motivation-towards a synthesis. In: Proceedings of the first international conference on simulation of adaptive behavior on From animals to animats. pp. 194–205 (1991)