Lieder, F., Goodman, ND, & Huys, QJM (2013). Learned helplessness and generalization. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.) Proceedings of the 35th Annual Confer…

Conference Paper · August 2013

Learned helplessness and generalization

Falk Lieder (lieder@biomed.ee.ethz.ch)
Translational Neuromodeling Unit; University of Zurich, & ETH Zurich, Zurich, Switzerland
Noah D Goodman (ngoodman@stanford.edu)
Department of Psychology; Stanford University, Stanford, CA, USA
Quentin JM Huys (qhuys@biomed.ee.ethz.ch)
Translational Neuromodeling Unit; University of Zurich, & ETH Zurich and
Psychiatric University Hospital Zurich, Zurich, Switzerland

Abstract

In learned helplessness experiments, subjects first experience a lack of control in one situation, and then show learning deficits when performing or learning another task in another situation. Generalization, thus, is at the core of the learned helplessness phenomenon. Substantial experimental and theoretical effort has been invested into establishing that a state- and task-independent belief about controllability is necessary. However, to what extent generalization is also sufficient to explain the transfer has not been examined. Here, we show qualitatively and quantitatively that Bayesian learning of action-outcome contingencies at three levels of abstraction is sufficient to account for the key features of learned helplessness, including escape deficits and impairment of appetitive learning after inescapable shocks.

Introduction

Helplessness is a failure to avoid punishment or obtain rewards even though they are under the agent's control. The aetiology and consequences of helplessness have been studied extensively in the animal learning literature using the learned helplessness paradigm [1]. In this paradigm, helplessness is induced in healthy animals by exposure to *in*escapable electric shocks. Helplessness is then measured by the subsequent failure to escape *escapable* shocks in a novel environment. The phenomenon was first demonstrated in dogs in the context of testing the two-factor learning theory using the shuttle-box escape task [2]. In the now classical version of the task, three animals are compared. A master rat experiences electrical shocks. These come on unpredictably, but can be terminated by some action, for instance turning a wheel. A yoked rat is exposed to the exact same sequence of shocks that are delivered to the master rat, but has no action available to terminate the shocks. A third rat, the control, is not exposed to shocks. Compared to the controls, the yoked rats are impaired at acquiring new instrumental responses, but the master rats are either unimpaired or may even show an improvement [1,3]. Effects reminiscent of those in animal learning have been demonstrated in humans. For instance, people who have been exposed to uncontrollable loud noise or insoluble problems are more likely give up on solving anagrams in a subsequent task [4]. Hopelessness theory is a translation of the helplessness concepts to the human and attributional realm [5].

Extensive animal and human experimentation and theoretical work have clarified that the crucial component is a perceived lack of control rather than more specific explanations. First, exposure to one type of uncontrollable reinforcer in one situation can profoundly impair the acquisition of many other types of behaviours (including escape, jumping, immobility, lever pressing, and complex sequential behaviours) in a wide range of *different* situations [1]. Second, the effect of inescapable shocks is not due to shock-induced analgesia, because it also impairs reward seeking in the absence of negative reinforcers [6]; and as uncontrollable rewards can also induce helplessness [7–9]. Third, it is not reducible to an interference between learned motor inactivity as it also affects the ability to acquire behavioural suppression as an escape response [10]. Furthermore, in accordance with the finding that the probability of generalization increases with the variability of the examples (see [11]), learned helplessness lasts longer when it is induced by experiencing multiple mild stressors (chronic mild stress; [12]) than when it is induced by a single severe stressor. In conclusion, there is strong evidence suggesting that helplessness is learned by generalizing from one uncontrollable situation to believing that situations are uncontrollable in general.

Maier and Seligman formalized controllability in terms of the conditional probability of the reinforcer RF (reward or punishment) given whether or not action A is taken (A or $\bar{A}$) [1]. According to their definition, the agent has control if and only if $P(\mathrm{RF}|A) \neq P(\mathrm{RF}|\bar{A})$. Huys and Dayan formalized the essence of this definition, i.e. that helplessness exists when altering behaviour does not alter outcomes, for multiple actions, outcomes, and degrees of desirability [13]. Here we continue [13]'s argument that generalization is central to learned helplessness by investigating two crucial interactions between perceived control and generalization:

1. Learning about the controllability of one situation transfers to novel situations via generalization.

2. Abstract knowledge about control determines how strongly the observation that a particular action had a particular effect will be generalized.

Using a hierarchical Bayesian model of action-outcome-contingencies we show that these interactions are sufficient to explain various deficits observed in the learned helplessness paradigm.

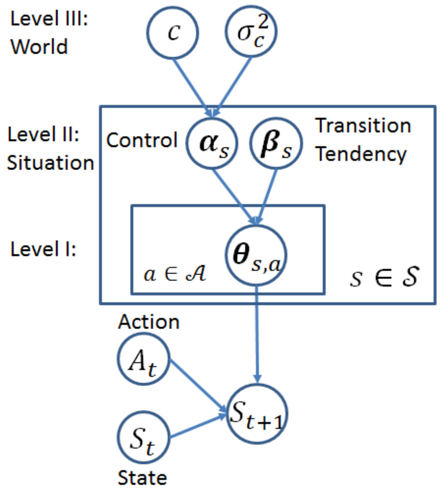

Figure 1: Hierarchical Bayesian model of state-transitions and controllability. $\boldsymbol{\theta}_{s,a}$ are the transition probabilities for taking action a in state s (level I). The second level, abstracts away from particular actions and represents the general outcome tendency $\boldsymbol{\beta}_s$ of situation s and its controllability α_s. The third level abstracts away from any particular state and represents how controllable the world is in general (c) and how much states differ with respect to controllability (σ_c^2).

Methods

In a novel situation s, a rational agent may have to learn how likely each of the available actions $a_1, \cdots, a_m$ is to lead into each of the potential successor states $s_1, \cdots, s_n$. In the absence of knowledge about the particular situation s, the agent can bring experience in other situations s' to bear on the problem, i.e. it can use its knowledge about one part of the transition matrix to inform its belief about others. Hierarchical Bayesian formulations provide a normative framework for such generalizations [11,14]. In this section, we present such a model of state-transition probabilities with three levels of hierarchy (see Figures 1 and 2). At the lowest level are the probabilities that taking action a in state s will lead to state s'. At the second level, the model represents the typical outcome probabilities of actions in any one particular situation s and how different the outcomes of different actions tend to be. The more actions are believed to have similar outcomes, the less control there is. At the third and most abstract level, the model represents knowledge about how controllable situations are in general. In this model beliefs about the world's controllability acts as an over-hypothesis that shapes how the agent learns state-transition probabilities (cf. [14]). Concretely, the agent's belief about the state S_{t+1} resulting from taking action a in state s is a multinomial distribution (Equation 1). The agent assumes that the transition probabilities $\boldsymbol{\theta}_{a,s}$ of the actions $\mathbf{a}$ available in state s are all drawn from the same distribution: a Dirichlet distribution with the state-specific mean vector $\boldsymbol{\beta}_s$ and

$$\begin{aligned}
\forall a, s : S_{t+1} | S_t = s, A_t = a &\sim \text{Multinomial}\left(\boldsymbol{\theta}_{a,s}\right) && (1)\\
\forall a, s : \theta_{\mathbf{a},\mathbf{s}} &\sim \text{Dirichlet}\left(\alpha_s \cdot \boldsymbol{\beta}_s\right) && (2)\\
\forall s : -\log(\alpha_s) &\sim \mathcal{N}\left(c, \sigma_c^2\right) && (3)\\
\forall s : \beta_s &\sim \text{Dirichlet}\left(\mathbf{1}\right) && (4)\\
c &\sim \mathcal{N}(\mu, \sigma_\mu^2) && (5)\\
\sigma_c^2 &\sim \text{InvGamma}(\alpha_\sigma, \beta_\sigma) && (6)
\end{aligned}$$

Figure 2: The functional dependencies of the graphical model in Figure 1.

a second parameter α_s that determines the controllability of situation s (Equation 2). If α_s goes to ∞, then the agent becomes sure that the transition probabilities $\boldsymbol{\theta}_{a_1,s}, \cdots, \boldsymbol{\theta}_{a_N,s}$ are independent of the agent's action a. This means that the situation is uncontrollable (corresponding to the second notion of controllability in [13]). Values of α_s close to zero corresponds to the belief that the transition probabilities $\boldsymbol{\theta}_{a_1,s}, \cdots, \boldsymbol{\theta}_{a_N,s}$ for different actions are uninformative about each other and hence can differ. To allow for the transfer of knowledge between states, a further level is needed: in addition to its belief about the controllability α_s of individual situations s, the agent also has a belief about how controllable situations are in general. This belief is described by a normal distribution on $-\log(\alpha_s)$ (Equation 3). The parameter $c = \mathbb{E}[-\log(\alpha)]$ expresses how controllable situations are on average, and σ_c^2 expresses how much controllability varies from situation to situation.

The average controllability c and the variability of control σ_c^2 are unknown properties of the world that have to be learned from experience. We describe the agent's prior beliefs about these quantities by a normal distribution on c (Equation 5) and an Inverse-Gamma distribution on σ_c^2 (Equation 6). In this model helplessness results from a probabilistic belief that one's control over the world is low on average (low c) and varies very little across situations (low σ_c^2).

Assuming that this hierarchical Bayesian model captures the subjects' internal representation of transition probabilities and control, we can examine how they infer the controllability c of the world in general from the observations $\mathbf{o} = \{(s_1, a_1, s_2), \cdots, (s_{t-1}, a_{t-1}, s_t)\}$ of the state transitions $(s_1, s_2), \cdots, (s_{t-1}, s_t)$ and their actions $a_1, \cdots, a_t$. In addition, we can simulate the weaker generalization on the situation-specific controllability α_s and transition-tendency $\boldsymbol{\beta}_s$ by computing $P(\alpha_s, \boldsymbol{\beta}_s | \mathbf{o})$. Finally, we can investigate how this generalization is shaped by abstract beliefs about control ($P(c, \sigma_c^2)$).

Simulations The effects of controllable and uncontrollable shocks were modelled by simulating the learning process taking place during the shocks as Bayesian inference on c and σ_c^2. The naive subjects' belief about controllability was modelled by a probability distribution with $\mathbb{E}[\alpha] = 10$ and $\text{Var}[\alpha] = 100$; the variance of the prior beliefs was 100 for c and 0.1 for σ_c^2. To model the observations resulting from controllable ($\mathbf{o}_\text{c}$) versus uncontrollable shocks ($\mathbf{o}_{\neg\text{c}}$), we assumed one observation per second during 64 shocks lasting 60 seconds each. For controllable shocks there was one action (a_1) that would always terminate the shock ($s_1 \to s_2$) and four actions that did not ($s_1 \to s_1$), whereas there was no such action for uncontrollable shocks.

Learning after exposure to shocks was modelled as Bayesian inference on $\boldsymbol{\alpha}, \boldsymbol{\beta}$ given $c = \mathbb{E}[c|o_{\neg c}]$ and $\sigma_c^2 = \mathbb{E}[\sigma_c^2|o_{\neg c}]$ for uncontrollable shocks versus $c = \mathbb{E}[c|o_c]$ and $\sigma_c^2 = \mathbb{E}[\sigma_c^2|o_c]$ for controllable shocks.

Inference was performed using Markov chain Monte-Carlo (MCMC) methods. To sample from $P(\boldsymbol{\alpha}, \boldsymbol{\beta}, c, \sigma_c^2|\mathbf{o})$, we used a Metropolis-Hastings algorithm with Gaussian random-walk proposals on c and $-\log(\alpha)$. The proposal for $\boldsymbol{\beta}_{t+1}$ was Dirichlet$(10n \cdot \boldsymbol{\beta}_t)$ where n is the number of states, and the proposal for $\sigma_{c,t+1}^2$ was an Inverse-Gamma distribution with mean $\sigma_{c,t}^2$ and variance 1. 50 Markov chains were run for 51000 iterations with a burn-in period of 1000 iterations. $P(\boldsymbol{\alpha}, \boldsymbol{\beta}|c, \sigma_c^2, \mathbf{o})$ was computed in the same way. The posterior expectation of $\boldsymbol{\theta}$ was computed using Monte-Carlo integration: $\mathbb{E}[\boldsymbol{\theta}|\mathbf{o}, c, \sigma_c^2] = \int \mathbb{E}[\boldsymbol{\theta}|\alpha, \beta, \mathbf{o}] \cdot p(\boldsymbol{\alpha}, \boldsymbol{\beta}|\mathbf{o}, c, \sigma_c^2) d\boldsymbol{\alpha} d\boldsymbol{\beta} \approx \frac{1}{m}\Sigma_{i=1}^m \mathbb{E}[\boldsymbol{\theta}|\boldsymbol{\alpha}_i, \boldsymbol{\beta}_i, \mathbf{o}]$ with $\boldsymbol{\alpha}_i, \boldsymbol{\beta}_i \sim P(\boldsymbol{\alpha}, \boldsymbol{\beta}|\mathbf{o}, c, \sigma_c^2)$.

Results

As a first step in assessing whether generalization can account for the differential effects of controllable versus uncontrollable stress, we simulated Bayesian learning from these experiences according to the model shown in Figure 1. Figure 3A shows the simulated changes in perceived controllability induced by the escapable and inescapable shocks administered in the learned helplessness paradigm [15]. After inescapable shocks, the subjects' perceived control c was reduced, whereas controllable shocks increased it. Furthermore, controllable shocks increased the estimated variability of control across situations, whereas no such change was observed after inescapable shocks (Figure 3B). Thus, the two kinds of shocks have opposite effects on the subjects' high-level beliefs about controllability.

We next asked whether the beliefs induced by uncontrollable shocks were sufficient to impair escape learning in a different task, and whether the beliefs induced by controllable shocks have the opposite (mastery) effect. We modelled these beliefs by the inferred mean and variance of c for escapable and inescapable shocks and simulated learning in the shuttle-box escape task. As a first step, we simulated learning from given observations with one action that did (a_1) and four actions that did not cancel the shock ($a_2, \cdots, a_5$). Concretely, we simulated how strongly naive subjects, subjects who had experienced inescapable shocks (yoked), and subjects who had experienced escapable shocks (masters) would believe that action a_1 cancels the shock after having taken action a_1 for $0, 1, 2, 4, 8, 16$ or 32 times and each of the four other actions 8 times. Figure 4 shows that yoked subjects (red) acquired the escape response more slowly than naive subjects (blue): more evidence was required before they believed that the action was efficient in terminating the shock. Furthermore, the model

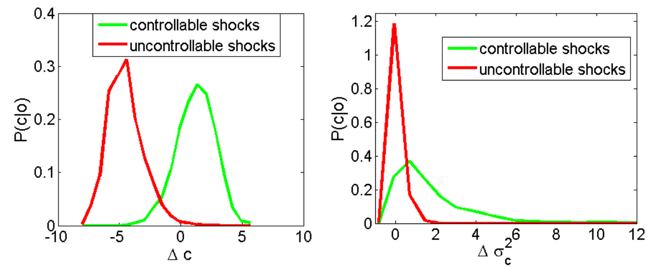

Figure 3: **A:** Expected controllability learned from controllable or uncontrollable electric shocks. The values on the x-axis are the change Δc relative to the controllability expected by naive subjects and height of the bars shows how strongly the simulated agent beliefs in the corresponding value of c. **B:** Variance of controllability learned from controllable and uncontrollable electric shocks. The values on the x-axis are the change Δc relative to the variability expected by naive subjects and height of the bars shows how strongly the simulated agent beliefs in the corresponding value of σ_c^2.

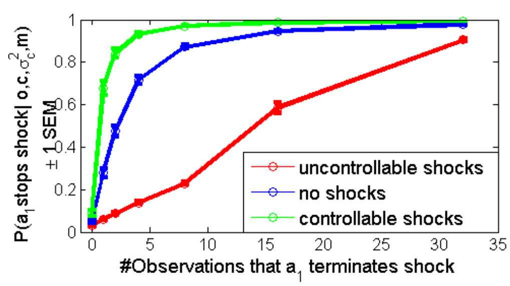

Figure 4: Simulated effects of controllable and uncontrollable shocks on the speed of learning that action 1 terminates the shock.

also captured mastery effects, whereby prior exposure to controllable shocks leads to faster learning (green; [16]).

To more quantitatively relate the learning dynamics shown in Figure 4 to empirical data, we simulated learning and decision making in the fixed-ratio operant conditioning task of [17]. In this task, rats have to learn to press a lever, but only every third lever press terminates the shock. This task was modelled as sequential-decision making. To do so, we partitioned the 60 seconds of each trial in [17]'s experiment into 30 bins, each 2 seconds long, and simulated one decision, one observation, and one belief update for every bin. The simulated choices were to stay still (a_0), to push the lever (a_1), or to perform a different action (a_2). The reward for staying still and receiving the shock was modelled as -1 ($r(s_1, a_0, s_1) = -1$). Moving and receiving a shock was assumed to incur a small additional cost ($r(s_1, a_2, s_1) = -1.2$). If the action stopped the shock, it was assumed to incur only the cost of movement ($r(s_1, a_0, s_2) = 0$ and $r(s_1, a_i, s_2) = -0.2$ for $i \in \{1, 2\}$). We assumed that rats

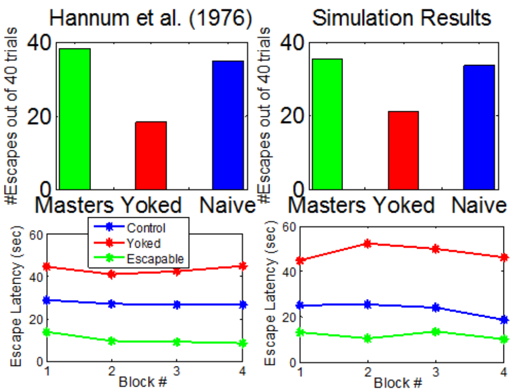

Figure 5: The left panel shows empirical data acquired by [17] as shown in [1]. The plots on right show our simulations of the experiment.

learn the probability $P(S_{t+1} = s_2|S_t = s_1, A_t = a_i)$ that an action a_i terminates the shock (an alternative model might consider treating different numbers of lever presses as separate actions). The subjects' internal representation of transition probabilities was modelled as the hierarchical Bayesian model shown in Figure 1. The rat's decision making was simulated by a sampling algorithm to produce behaviour akin to probability-matching [18]. Specifically, we assumed that the rat simulates five outcomes $u_{i,j} = r(s, a_i, s'_j), s'_j \sim P(S_{t+1}|S_t = s, A_t = a_i)$ of each action a_i and chooses the action a_i for which the average utility $\frac{1}{5}\sum_{j=1}^{5} u_{i,j}$ was largest, and ties were broken at random. Under these assumptions, the learning dynamics shown in Figure 4 capture the qualitative effects of uncontrollable shocks on the probability to escape shock and the time required to do so [17]: yoked subjects failed to escape more often than naive subjects (Figure 5, left panel), and when they succeeded to escape it took them longer (Figure 5, right panel). Furthermore, our model accounts for the mastery effect that rats who had been exposed to controllable shocks prior to the task, escaped faster than rats with no prior exposure to shock.

As outlined in the introduction, learned helplessness impairs not only the ability to learn from punishments but also from rewards. To assess whether our model captures this effect, we simulated the experiment by [6]. In the experiment's appetitive choice task, rats were rewarded with food for going into one of two chambers after they had been trained to prefer the other chamber. We modelled this task as a sequence of decisions, observations, and belief updates as described above. As Figure 6 shows our model captures that uncontrollable shocks reduced the probability that a rat would first seek out the chamber in which a reward would be delivered. Thus this apparently anhedonic behaviour can be explained purely in terms of impaired associative learning due the generalization that the world is uncontrollable.

Next, we asked whether our model can account for the finding that the effect of learned helplessness is most pronounced in tasks that are complex and require persistence. To answer this question, we simulated decision

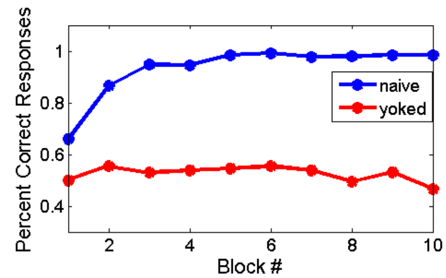

Figure 6: Simulation of the appetitive choice distinction task by [6]. Our simulation captures that yoked rats performed worse than naive rats across all 10 blocks of the experiment.

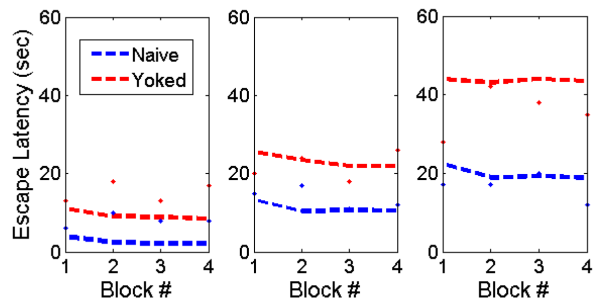

Figure 7: Simulation of the experiment by [19]. Dashed lines are model predictions; diamonds are data points. The three columns correspond to the experimental conditions requiring 1, 2, or 3 lever presses.

making and learning in the experiment by [19]. In this experiment, yoked rats did learn to escape response when one or two, but not when three lever presses were required. In Figure 7, we show that the model can quantitatively capture the increasing penetrance of inescapable shock exposure with increasing escape response requirements.

Discussion

Our results indicate that a normative account of generalization of action-outcome contingencies is sufficient to produce a wide range of phenomena observed in learned helplessness experiments. The account captures (i) how helplessness is induced by uncontrollable stressors and why it transfers to novel situations, (ii) why controllable stress fails to induce helplessness, (iii) that helplessness results from impaired learning that different actions have different effects, (iv) mastery effects, (v) impaired reward seeking, and (vi) the interaction between helplessness and task requirements. This suggests that the generalization of experienced control may be sufficient to account for many learned helplessness effects.

Note that our model explains helplessness as the consequence of rational learning and generalization (cf. [11, 14]) from uncontrollable stress. Mirroring the fact that learned helplessness induces depression-like states in healthy animals and affects healthy humans, this suggests one pathway by which learned helplessness may arise as a rational adaptation to an uncontrollable environment rather than from negative biases or dysfunctional information processing. In the present account, the generalization that leads to helplessness is that different actions tend to have the same effect. This generalization can occur at two levels of abstraction: (i) from the outcomes of some actions in a given situation to the outcomes of all actions available in that situation, and (ii) from the controllability of one situation to the controllability of situations in general. Our model predicts that after uncontrollable stress generalizations of the first type will be unusually strong. This may capture overgeneralization – a frequent feature of depressive thought in humans [20] – and the model explains how this learning style increases the risk for helplessness by fostering the belief that all actions are equal.

According to the classical notion [1], control requires that taking an action or not alters the probability of outcomes. Our model generalizes this notion by allowing for multiple actions, multiple outcomes, and two additional levels of abstraction, and it expands it from a binary distinction to a graded, quantitative measure of controllability. As a result, our model instantiates [13]'s second type of controllability which captured on the achievability of different outcomes. While the notion of control presented here does not take into account the outcomes' desirability (type 3 in [13]), it refines [13]'s proposal of how helplessness might be learned by generalization in two regards. First, [13] juxtaposed two extremes of generalization: the controllabilities of different environments are (i) independent (no generalization) or (ii) identical (full generalization). Arguably, both extremes correspond to pathologically inaccurate models of the world. By contrast, the hierarchical generative model proposed here formalizes the more realistic intermediate assumption that although some situations are more controllable than others, knowing about the controllability of one situation is informative about the controllability of other situations. Second, inference in this model captures an important aspect of attribution: Was the outcome due to the action I took or due to the situation I was in? In the model, a perceived lack of control induces misattributions that impair learning: the over-hypothesis that the world is uncontrollable renders implausible any interpretation according to which different actions have different effects. Therefore the perceived lack of control biases people to attribute the outcome of taking action a in state s to the state s rather than to the action that they have taken. For extreme helplessness the situation's action-independent outcome tendency $\boldsymbol{\beta}_s$ will be updated just as much as the outcome probabilities $\boldsymbol{\theta}_{s,a}$ of taking action a in this situation. Conversely, the action-specific outcome probabilities $\boldsymbol{\theta}_{s,a}$ will be updated no more than $\boldsymbol{\beta}_s$, and the outcomes of actions $b, c, d, \cdots$ will influence the belief about $\boldsymbol{\theta}_{s,a}$ almost as much as the outcomes of action a itself. This increases the amount of evidence required to discover that there is an action that achieves the goal while most other actions do not, and this is how the perceived lack of control impairs learning. Thus, according to our model, helpless behaviour in simple tasks results from slowed learning of transition probabilities. This complements a recent model of how the perceived lack of control impairs planning in complex, sequential decision problems [21].

Despite the encouraging results reported here, more research is needed to establish the validity of this modelling approach. Our simulation of the experiment reported in [19] did not fully capture the rats' learning dynamics and overestimated their performance in the simplest condition. These discrepancies could be due to the simplistic assumption that rats do not associate shock termination with the number of lever presses but only with their most recent action. As a result, the fit achieved by the reported simulation is not substantially better than the fit obtained by reducing the subjective intensity of the shock (data not shown). Therefore the results of our simulations (Figure 7) could be mimicked by analgesia [19, 22]. Furthermore, our model failed to capture the precise pattern of the data in [6]. However, the two hypotheses might be discernible using data that reveal the dynamics of learning, or by directly probing subjects' beliefs about outcome probabilities in the absence of rewards.

There are important aspects of helplessness that our model does not capture yet. It needs to assume surprisingly weak priors to explain why helplessness can be learned so rapidly—an assumption that is difficult to extend to mature animals with a lifetime of experience. Conversely, it would be challenging to explain with our model why even severe shock-induced helplessness tends to fade within 48 hours. To capture these two suggestions of temporal (in)stability of helplessness, our model could be extended by replacing the constant c by a Gaussian random walk $C(t)$. This would take into account that how much control a person has, changes throughout his or her life. It could be used to explain why a very brief period of uncontrollable stress can induce helplessness despite a lifetime of experience in a controllable environment. This extension may also explain why the speed of learning about controllability increases with the environment's perceived volatility (cf. [23, 24]). Beyond these cognitive signatures of helplessness, our model has yet to engage with the affective and subjective aspects of helplessness.

Our model suggests a number of avenues for future work. As normative inference, exposure to sufficient controllability will lead to the correct inference, suggesting that experience of control should heal helplessness, and thus that the escape deficit should be temporary. Why and how the experience of shocks that an animal did not attempt to escape may nurture helpless beliefs may necessitate adaptations in the inference. Nevertheless, the model does capture that controllable stress can immunize subjects against the effects of uncontrollable stress, and suggests two computational mechanisms: first a higher expected controllability, and second a higher estimate of the variability of control across situations. If the initial expected controllability is higher, then a given reduction in perceived control may no longer be sufficient to induce learned helplessness. Alternatively, a higher estimate of the variability of control across situations would reduce the strength of the generalization from uncontrollable stress in one situation to the controllability of the world in general. This theoretical work does, however, also echo the likely limited benefits of exposure because a subjects' helpless choices in some situations, paired with strong tendencies to generalise, can produce sufficient evidence of lack of control to drown any islands of controllability.

Our model may be able to illuminate why and how the effects of chronic-mild stress [12] differ from the effects of severe-acute stress in their scope, severity, and duration. The results suggests that the underlying mechanisms can be understood in terms of well-studied general generalization phenomena [11]. For instance, since the strength of a generalization increases with the variability of the examples, the experience of multiple stressors should render the effects of chronic mild stress more general than the equivalent amount of stress experienced in a single situation.

Learned helplessness is a behavioural paradigm with parallels in humans and animal models, and with established validity in research on depression [25]. Using methods from reinforcement and machine learning, this work has shown that abstract learning and generalization about controllability explain many of the key features of learned helplessness in animals.

References

[1] S. F. Maier and M. E. Seligman, "Learned helplessness: Theory and evidence.," *J EXP PSYCHOL GEN*, vol. 105, pp. 3–46, 1976.

[2] B. Overmier and M. Seligman, "Effects of inescapable shock upon subsequent escape and avoidance responding.," *J COMP PHYSIOL PSYCH*, vol. 63, pp. 28–33, 1967.

[3] S. F. Maier and L. R. Watkins, "Stressor controllability and learned helplessness: the roles of the dorsal raphe nucleus, serotonin, and corticotropin-releasing factor.," *NEUROSCI BIOBEHAV R*, vol. 29, pp. 829–841, 2005.

[4] C. Peterson, S. F. Maier, and M. E. P. Seligman, *Learned Helplessness: A theory for the age of personal control.* OUP, 1993.

[5] L. Y. Abramson, G. I. Metalsky, and L. B. Alloy, "Hopelessness depression: A theory-based subtype of depression," *Psychol. Rev.*, vol. 96, pp. 358–372, 1989.

[6] R. A. Rosellini, J. P. DeCola, and N. R. Shapiro, "Cross-motivational effects of inescapable shock are associative in nature.," *J EXP PSYCHOL ANIM B*, vol. 8, pp. 376–388, 1982.

[7] F. Goodkin, "Rats learn the relationship between responding and environmental events: An expansion of the learned helplessness hypothesis," *LEARN MOTIV*, vol. 7, 1976.

[8] R. L. Welker, "Acquisition of a free-operant-appetitive response in pigeons as a function of prior experience with response-independent food," *LEARN MOTIV*, vol. 7, pp. 394–405, 1976.

[9] Q. J. M. Huys, J. Vogelstein, and P. Dayan, "Psychiatry: Insights into depression through normative decision-making models," in *Adv. Neural Inf. Process. Syst. 21* (D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, eds.), pp. 729–736, MIT Press, 2009.

[10] R. L. Jackson, S. F. Maier, and P. M. Rapaport, "Exposure to inescapable shock produces both activity and associative deficits in the rat," *Learn. Motiv.*, vol. 9, pp. 69–98, 1978.

[11] J. B. Tenenbaum and T. L. Griffiths, "Generalization, similarity, and bayesian inference.," *Behav Brain Sci*, vol. 24, 2001.

[12] P. Willner, "Validity, reliability and utility of the chronic mild stress model of depression: a 10-year review and evaluation," *Psychopharm*, vol. 134, pp. 319–29, 1997.

[13] Q. J. M. Huys and P. Dayan, "A bayesian formulation of behavioral control.," *Cognition*, vol. 113, pp. 314–328, 2009.

[14] C. Kemp, A. Perfors, and J. B. Tenenbaum, "Learning overhypotheses with hierarchical bayesian models," *Dev. Sci.*, vol. 10, pp. 307–321, 2007.

[15] M. Seligman and S. Maier, "Failure to escape traumatic shock.," *J EXP PSYCHOL*, vol. 74, pp. 1–9, 1967.

[16] J. R. Volpicelli, R. R. Ulm, A. Altenor, and M. E. P. Seligman, "Learned mastery in the rat," *LEARN MOTIV*, vol. 14, pp. 204–222, 1983.

[17] R. D. Hannum, R. A. Rosellini, and M. E. Seligman, "Learned helplessness in the rat: Retention and immunization.," *DEV PSYCHOL*, vol. 12, pp. 449–454, 1976.

[18] R. J. Herrnstein and D. H. Loveland, "Maximizing and matching on concurrent ratio schedules.," *J Exp Anal Behav*, vol. 24, pp. 107–116, 1975.

[19] M. Seligman and G. Beagley, "Learned helplessness in the rat.," *J COMP PHYSIOL PSYCH*, vol. 88, pp. 534–541, 1975.

[20] A. T. Beck, A. J. Rush, B. F. Shaw, and G. Emery, *Cognitive therapy of depression.* Guilford Press, 1979.

[21] F. Lieder, N. D. Goodman, and Q. J. M. Huys, "Controllability and resource-rational planning.," *Cosyne Abstracts 2013*, 2013.

[22] R. L. Jackson, S. F. Maier, and D. J. Coon, "Long-term analgesic effects of inescapable shock and learned helplessness," *Science*, vol. 206, no. 4414, pp. 91–93, 1979.

[23] T. E. J. Behrens, M. W. Woolrich, M. E. Walton, and M. F. S. Rushworth, "Learning the value of information in an uncertain world.," *Nat Neurosci*, vol. 10, pp. 1214–1221, 2007.

[24] C. Mathys, J. Daunizeau, K. J. Friston, and K. E. Stephan, "A bayesian foundation for individual learning under uncertainty.," *Front Hum Neurosci*, vol. 5, 2011.

[25] L. Y. Abramson, L. B. Alloy, M. E. Hogan, W. G. Whitehouse, M. Cornette, S. Akhavan, and A. Chiara, "Suicidality and cognitive vulnerability to depression among college students: a prospective study.," *J Adolesc*, vol. 21, pp. 473–487, 1998.