

Dynamic self-efficacy as a computational mechanism of mania emergence

Jing Li (jing.li@icahn.mssm.edu)

Angela Radulescu (angela.radulescu@mssm.edu)

Center for Computational Psychiatry, Icahn School of Medicine at Mount Sinai, New York, NY, USA

Abstract

Abstract: Bipolar disorder (BD) is a mental health condition characterized by large fluctuations in goal-directed energy and mood. BD is defined by the presence of at least one lifetime episode of mania, a prolonged period of excessive goal-directed behavior, hyperactivity and elevated mood. Previous computational models of BD have primarily focused on explaining mood fluctuations in mania, placing less emphasis on goal-directed symptoms. In this work, we use reinforcement learning (RL), a principled model of goal-directed behavior and learning, to show how augmenting RL agents with *dynamic self-efficacy beliefs* can give rise to goal-directed and mood symptoms characteristic of the mania phase of BD. Our simulations demonstrate that a model-free RL agent that dynamically updates its self-efficacy beliefs learns optimistic overgeneralized value representations. We suggest that these representations may underlie several behaviors associated with mania, such as increased motivational drive and faster initiation of approach behavior (i.e. impatience). We further show that agents with more sensitive self-efficacy beliefs display increased willingness to exert effort in order to achieve higher goals even in the face of costs, a characteristic that is observed in individuals at risk for BD. Finally, unrealistically high self-efficacy beliefs that emerged with learning were accompanied by behaviors such as distractibility and compulsive action selection that have clinical parallels to symptoms of mania.

Keywords: self-efficacy; computational psychiatry; reinforcement learning; bipolar disorder; mania

Introduction

Bipolar disorder (BD) encompasses manic and depressive episodes, interspersed with phases of stable energy levels and mood (Goodwin & Jamison, 2007). The hallmark of BD is the lifetime occurrence of at least one episode of mania, a period marked by increase in activity and energy, an escalation in goal directed behavior, and an unusually intense and persistent elevation in mood (American Psychiatric Association, 2013). Recent updates to diagnostic criteria emphasize a general increase in goal-directed behavior and energy as a key behavioral marker of mania (Mason et al., 2016). And cognitive models of mania suggest a dynamic process by which surges in confidence, along with consistently high ambitions, can lead to excessive goal pursuit, which eventually spirals into mania (Johnson, 2005).

One way to interpret increased goal-directed behavior in mania is through the lens of *self-efficacy*. Previous theoretical work has demonstrated how lowered levels of self-efficacy

can result in impaired goal-pursuit behaviors (Zorowitz et al., 2020). Self-efficacy, defined as one's belief in the capacity to execute actions that achieve desired outcomes, is an adaptive trait that is continuously shaped through performance accomplishments, vicarious experience, social persuasion and physiological signals (Bandura, 1997). For instance, a novice cook begins by making simple dishes. Each culinary success enhances their belief in their cooking abilities. Over time, as they master more complex recipes, their self-efficacy in the kitchen grows, leading to the confidence to experiment and create new dishes. This belief in their capabilities drives them to take on greater culinary challenges, fostering a cycle of continuous learning and skill development. Consistent with Bandura's theory, the relationship between self-efficacy and performance has been empirically supported, showing that higher levels of self-efficacy are associated with better task performance (Themanson & Rosen, 2015).

In the case of mania, the pursuit of goals appears to be dysregulated (Johnson, 2005) and is often coupled with grandiosity – an exaggerated sense of self-importance and personal capabilities (Kendler, 2017). This pattern results in unrealistic self-assessment, and grandiose beliefs that persist even when they are contradicted by reality (APA, 2013). Such beliefs often lead to behaviors that may result in severe consequences for the individual. Johnson (2005) noted that after experiencing initial success, individuals with BD showed larger increases in confidence, and that individuals at risks of mania are more likely to expect more successes after small ones. Moreover, life events, especially those that are perceived as successes, have been shown to predict the onset of manic episodes (Johnson et al., 2012).

These findings led us to hypothesize that when the process of interpreting information from feedback regarding one's self-efficacy is altered, behaviors similar to those observed in mania might emerge through repeated experience with the environment. To test our hypothesis, we turned to reinforcement learning (RL), a computational framework widely used to study goal-directed behavior.

Model

Recent work in computational psychiatry has suggested that a behavioral consequence of *low* self-efficacy is extreme avoidance in the face of threats and withdrawal from approaching rewarding states (Zorowitz et al., 2020). This work provides a computational foundation for modeling the effect of self-efficacy representations on learning and behavior.

Here, we model the *dynamics of self-efficacy* in response to feedback from the external environment. We introduce a learning rule that updates self-efficacy in appraisal of mov-

ing closer to a goal (Bandura, 1997). We propose that self-efficacy — defined as the belief that achieving a goal now is more likely to lead to achieving goals in the future — is a dynamic attribute, continuously shaped by action outcomes.

Q-Learning with dynamic self-efficacy

We implemented a model-free Q-learning agent that learns in sequential grid-world environments (Dulberg et al., 2023; Gagne & Dayan, 2022; Zorowitz et al., 2020). This set-up enabled us to study how changes in self-efficacy affect reward backpropagation across states as the agent interacts with the environment. In our model, reward prediction errors (RPEs) — the discrepancy between expected and actual action outcomes — serve as direct feedback for the appraisal of performance accomplishment. This aligns with the insight that enactive mastery of experiences is a critical source of efficacy information because it is the most direct way for the agent to learn self-efficacy beliefs based on its successes and failures (Bandura, 1997).

The agent selects actions via a softmax policy (inverse temperature = 1), which probabilistically chooses actions in proportion to their value estimates:

$$\pi(a_t) = \frac{e^{Q(s_t, a_t)}}{\sum_a e^{Q(s_t, a)}} \quad (1)$$

After the agent takes an action, the reward prediction error (RPE) is computed as:

$$\delta_t = R_{t+1} + \gamma \cdot w_t \cdot \max Q(s_{t+1}, a_{t+1}) - Q(s_t, a_t) \quad (2)$$

The *self-efficacy belief*, w_t , scales the highest possible future expected reward contingent on action, reflecting the agent’s belief that it can successfully select the best action in the immediate future.

Reward prediction errors serve as the critical signal for updating both the action values and the self-efficacy parameter:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha \cdot \delta_t \quad (3)$$

$$w_{t+1} = w_t + w_{LR+} \cdot \delta_t \quad (4)$$

The *reward learning rate* α indexes how quickly value estimates should be updated based on error.

And the *self-efficacy learning rate*, w_{LR+} , quantifies how sensitive the agent’s beliefs about its own self-efficacy are to positive RPEs, which in a sequential setting can be interpreted as a signal of successfully approaching a goal.

In the next section we detail the structure of the learning environments and the parameters of the agents we trained in each simulation.

Environments

To formalize our hypothesis that the emergence of mania symptoms could be driven by larger updates in self-efficacy, we defined two distinct simulation environments based on a standard deterministic grid-world RL set-up.

Grid-world 1: single goal state In the first environment (Fig. 1 and Fig. 3), the agent starts at the top left location (0,0) and must reach a rewarding terminal state at the bottom right position (9,9). In each state, the agent can either move up, down, left or right. The reward is +1 at the terminal state and 0 everywhere else in the grid-world. To discourage the agent from straying off the grid, a penalty of -0.5 is imposed for such actions. The agent’s training consisted of 200 episodes, each with up to 200 steps (each step represents an action was taken), unless the terminal state was reached sooner, concluding the episode.

Grid-world 2: two goal states with costs In the second environment (Fig. 2), the agent starts from state (4,0), and the two rewarding terminal states are (0,9) and (9,9). The top right terminal state (0,9) results in a reward of 1.5, while the bottom right terminal state (9, 9) yields a reward of 4. Navigating to the latter, more lucrative, terminal state requires the agent to move through two costly preceding states that result in a cost of -2. This set-up allowed to ask under what conditions would the agent exert more effort to get to the more rewarding state, despite the additional cost. To accommodate for learning in this more complex environment, we extended the training to 350 episodes, maintaining the maximum of 200 steps per episode.

Fixed agent parameters All agents have the same rewarding learning rate of $\alpha = 0.6$, and a discount factor of $\gamma = 0.9$. The agent’s action policy is softmax with an inverse temperature of 1, and there is a 0.001 chance that the agent will randomly select from the four actions to encourage exploration.

Results

Positive overgeneralization

Using the single goal state environment, we first tested how different learning rates for self-efficacy impact the underlying value representation. We repeated the simulation for three different agents with the self-efficacy learning rate w_{LR+} fixed at 0.0, 0.0001, and 0.00015 respectively.

As expected, the three agents acquired different levels of self-efficacy at the end of training, with higher learning rates leading to higher self-efficacy levels (Fig. 1A). To assess how self-efficacy learning rates affect the propagation of rewards in the environment, we examined the learned value representation of the environment after training. We found that for agents that update self-efficacy more quickly, the reward at the terminal states backpropagates to states closer to the start state (Fig. 1B); and the overall value of all states is higher (Fig. 1C). In other words, agents with more sensitive self-efficacy beliefs develop *optimistic overgeneralized future reward expectations*.

This prediction of the model is consistent with the finding that individuals with BD often exhibit a significant boost in confidence that extends beyond specific positive events to wider areas of their life, particularly following initial suc-

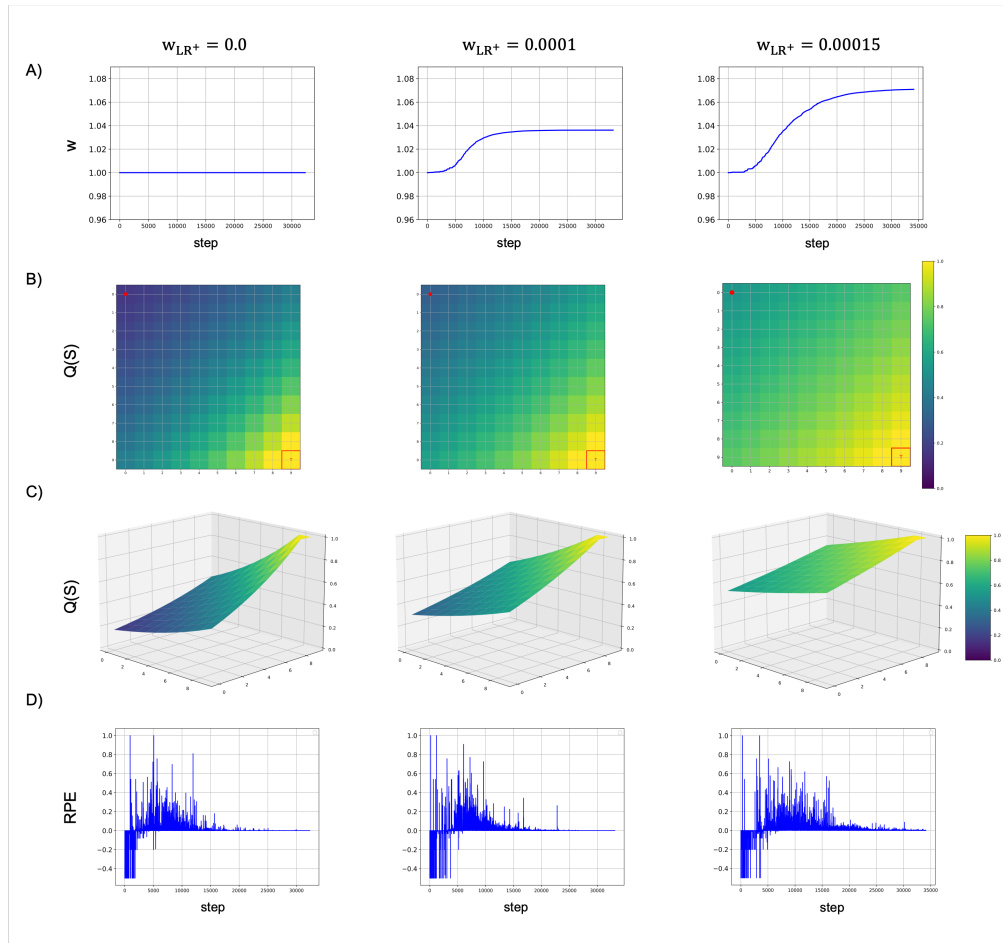


Figure 1: **Positive overgeneralization in agents with dynamic self-efficacy.** A) Agents with different self-efficacy learning rates (left, static; middle, $w_{LR+} = 0.0001$; right, $w_{LR+} = 0.00015$) reach different levels of self-efficacy after learning the grid-world environment. B) Agents with higher self-efficacy learning rates learn to place greater values for states closer to the start position. C) 3-D representations of the learned best action value maps. D) Positive RPEs during the training process are denser and higher for agents with more sensitive self-efficacy learning rates.

cesses (Johnson, 2005; Johnson et al., 2012). Moreover, people diagnosed with BD often demonstrate a bias towards overly optimistic self-assessment and evaluation of their emotional states (Jones et al., 2006). Notably, there are some data to suggest that among individuals at risk for BD, higher positive overgeneralization predicts elevated hypomanic symptoms (Stange et al., 2012).

Our simulations also show that as self-efficacy learning rates increase, actions leading towards the goal from equivalent states have higher value for agents that more quickly update self-efficacy beliefs. Behaviorally, this implies that agents with higher self-efficacy learning rates will consistently show faster response times (i.e. be more impatient) during goal pursuit. This feature of mania has been noted in clinical research emphasizing that people with BD display elevated drive to accomplish goals (Johnson, 2005) and complete task with rewards faster than healthy controls (Hayden et al., 2008).

The higher overall value across all states also predicts that larger self-efficacy updates should result in more effort in approaching a goal (Frömer et al., 2021). Previous theoretical work has considered weighing the cost of exerting effort against the payoffs of completing a task as a cost-benefit analysis (Shenhav et al., 2013), suggesting that people will expend more effort if the potential payoff is high. Thus, the overgeneralized optimistic value representations in our simulations imply that higher self-efficacy learning rates should result in expending more effort to approach a goal, a prediction consistent with the observation that people prone to mania will expend more effort toward reward pursuit (Johnson et al., 2012).

Finally, we examined the reward prediction errors during the learning process and observed that agents with higher learning rates on self-efficacy “experienced” denser positive RPEs at larger values (Fig. 1D). Previous work modeling mood has suggested recency-weighted RPEs as a driver of

mood (Eldar & Niv, 2015). In the context of our simulations, this would indicate more prolonged and higher positive mood increases for agents with more sensitive self-efficacy beliefs.

More willingness to pursue difficult-to-obtain goals

Using the two goal state environment with costs, we next examined how different levels of self-efficacy learning rates would impact the agent's willingness to approach goals with low reward/no cost versus goals with high reward/high cost. We repeated the simulation for three different agents with the self-efficacy learning rate w_{LR+} fixed at 0.0, 0.00005, and 0.00008 respectively. The self-efficacy learning rates were adjusted to accommodate the enriched rewards in the more complex environment.

As expected, agents with higher self-efficacy learning rates acquired higher levels of self-efficacy at the end of training (Fig. 2A). At a self-efficacy learning rate of 0, the agent learned to move towards either of the two rewarding terminal states, indicated by the learned best actions in each state (Fig. 2C). However, as the self-efficacy learning rate increased, the agent showed a stronger tendency to approach the bottom right terminal state with a higher reward, despite the cost to-be-incurred in one of the preceding adjacent states. Upon examining the representation of the best action values at two states equidistant to both goals, (0,6) and (9,6), we found that the action leading towards the bottom right terminal state acquired higher value when the self-efficacy learning rate is higher (Fig. 2B). This result implies that with a higher self-efficacy learning rate, the agent is more willing to pursue goals with higher rewards, despite increased cost.

These results align with the finding that individuals with vulnerability to hypomania may be more likely to pursue more difficult goals (Johnson, 2005) and people with BD remain engaged in pursuing rewards longer as the tasks became more difficult (Harmon-Jones et al., 2008). Moreover, the increased difference between the value representations of the two goal states also suggests increased willingness to expend effort towards the more difficult goal, a trait consistently related to BD (Johnson et al., 2012).

Emergence of extreme value attractors

We conducted an exploratory analysis and further increased the learning rate of self-efficacy in the single goal-state environment. This adjustment was informed by clinical literature, particularly Johnson et al. (2005), which demonstrates that individuals at risk for mania tend to experience a significant boost in confidence and set higher goals following initial successes. We asked if a larger increase in self-efficacy in response to successes will lead to other aspects of mania symptoms not observed at lower levels of self-efficacy learning rates. We found that self-efficacy beliefs display stepwise dynamics. When the self-efficacy learning rate was increased to 0.0002, a tipping point in the self-efficacy level was observed, such that late in training self-efficacy rapidly escalated to its artificial limit (Fig. 3A).

This stepwise increase in self-efficacy was related to the emergence of attractors in the value function. Initially, the reward backpropagation from the terminal state followed the expected Q-learning convergence pattern, with action values being highest near the terminal state and gradually becoming lower closer to the start state (Fig. 3B). However, as the training continued, an unexpected shift occurred in the peak value of the learned value representation. This peak gradually moved from the vicinity of the terminal state to more distant locations (Fig. 3C-D).

This shift continued (Fig. 3E-F) until the learned best actions formed a closed loop around the states with peak values of the learned value map (Fig. 3G). This loop entrapped the agent moving within a cycle of states with extremely high values, resulting in a rapid escalation of self-efficacy and extremely high action value representations in states far from the goal state.

Taken together, the rapid escalation of self-efficacy levels, the emergence of attractor-like value representations, and the consequent behavioral patterns have interesting parallels to certain aspects of mania. The extreme values of the self-efficacy beliefs can be interpreted as grandiosity, an exaggerated representation of personal capabilities (Kendler, 2017). The quick shifts in peak value representations during training can be likened to distractibility, a common clinical feature of mania (APA, 2013). Finally, the simulation revealed a representation of overly high values in a selective small number of states, leading to a limited set of repetitive actions that are qualitatively similar to compulsive behaviors (Kesebir et al., 2012).

Discussion

This study aimed to formalize the hypothesis that a higher sensitivity of self-efficacy beliefs to goal-directed feedback could provide a mechanism for the emergence of mania in bipolar disorders. We proposed a computational model based on reinforcement learning that augments model-free RL agents with dynamic self-efficacy beliefs. In our model, self-efficacy is continuously updated based on reward prediction errors (RPEs), which provide a direct signal of goal attainment. Informed by clinical literature showing that individuals at risk for mania demonstrate greater increase in confidence (Johnson, 2005), we varied the level of self-efficacy learning rates in response to successes and showed that higher sensitivity of self-efficacy beliefs can give rise to several cognitive and behavioral features observed in mania, including: optimistic overgeneralization of reward expectations; consequent increased drive and motivation for pursuing goals; impatience in pursuing goals; and an increased tendency to exert effort in pursuit of higher and more difficult to attain goals (APA, 2013; Stange et al., 2012; Johnson et al., 2017).

Exploratory simulations showed that an extremely high level of self-efficacy akin to grandiosity can emerge as a result of stepwise dynamics in self-efficacy beliefs. This sudden increase in self-efficacy was accompanied by rapid shifts




Figure 2: **Increased willingness to approach the more difficult goal with higher self-efficacy levels.** A) Agents with different levels of self-efficacy learning rates acquired different levels of self-efficacy after training (left, static; middle, $w_{LR+} = 0.00005$; right, $w_{LR+} = 0.00008$). B) The difference between the best action value in two states, (0,6) and (9,6), which are equidistant from the two terminal states, increases as self-efficacy learning rates increase (states (0,6) and (9,6) are highlighted in blue boxes in Fig. 2C). C) The agent with static self-efficacy ($w = 1$) learns to approach both terminal states with different reward levels (top right terminal state has a reward value of 1.5; bottom right terminal state has a reward value of 4 and is surrounded by states with a cost of -2). Each arrow represents the learned best action in that state. The red arrow denotes the general tendency of the agent to approach either the top or bottom state.

in peak value representations, as well as a qualitative change in the agent’s policy. The resulting policy was characterized by a narrow set of closed-cycle repetitive actions that mirror compulsive behavior observed in mania.

One open question is under what conditions dynamic self-efficacy could be adaptive. The self-efficacy parameter in our model acts as a multiplier over the valuation of future action values in addition to the discount factor. As the agent experiences positive outcomes, increased self-efficacy leads to higher expectations for future rewards, effectively accelerating and amplifying the backpropagation of the value of future events.

Through repeated practice and experience, humans enhance their skills, allowing for more efficient task performance with improved outcomes. An optimal update of self-efficacy as skills improve may be advantageous because: 1) it sets higher expectations of future rewards, thus fueling motivation to achieve more, and 2) it aids in prioritizing accom-

plishment of high-reward options, allowing the allocation of more resources towards endeavors that are more rewarding, albeit more effortful and costly.

Intriguingly, Johnson (2005) noted that individuals with a history of manic or hypomanic episodes often achieve higher socioeconomic status and report greater creative, educational, and occupational attainments despite the disruptive symptoms of the condition. This is consistent with the possibility that while an optimal level of self-efficacy update could be adaptive, individuals at risk for mania may have more sensitive self-efficacy beliefs. Whether such heightened sensitivity in self-efficacy beliefs could emerge through early-life experience is an open question (Harhen & Bornstein, 2024). Finally, we draw connections to the literature on intrinsic rewards. Chew, Blain, Dolan, and Rutledge (2021) showed that skill mastery is intrinsically rewarding and significantly contributes to affective dynamics. The design of our model, which integrates self-efficacy and value functions,

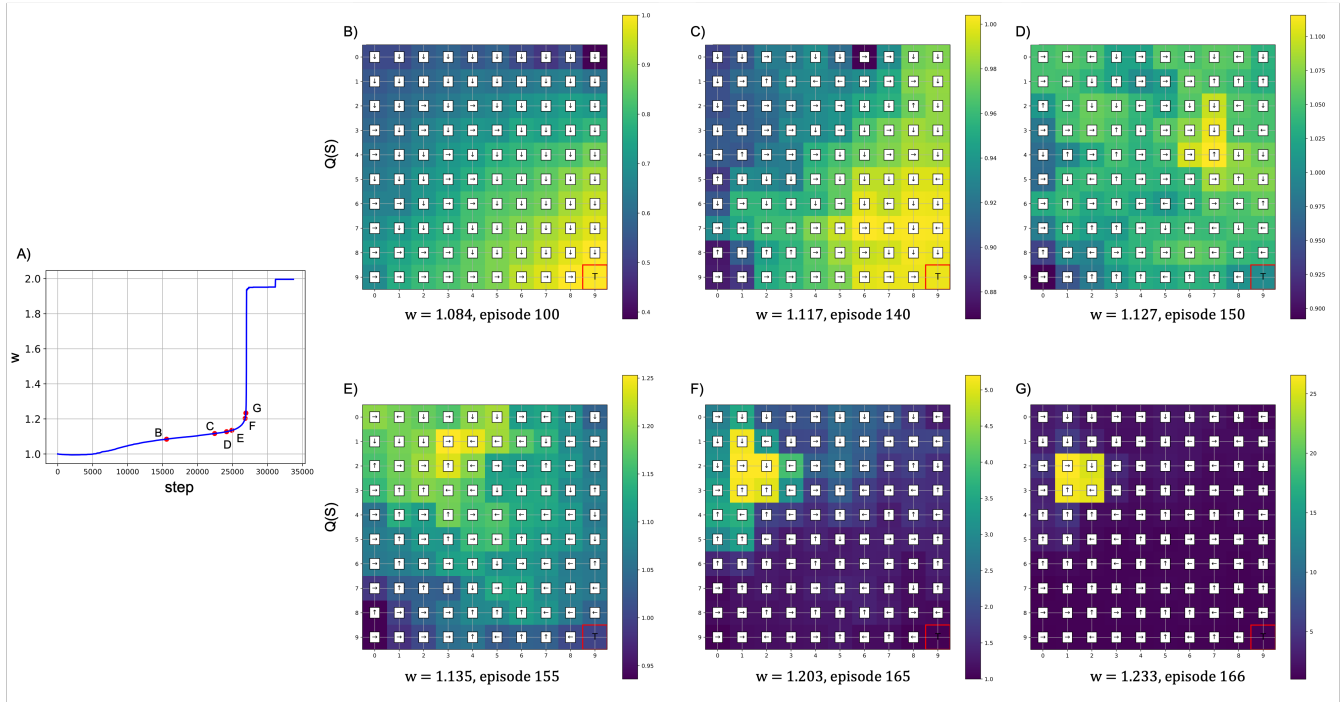


Figure 3: **Attractor-like states emerge as the self-efficacy learning rate increases** ($w_{LR+} = 0.0002$). A) Self-efficacy levels steadily increase until reaching an inflection point and then exponentially grow to the artificial limit. Each red dot represents a time point when value maps shown in B-G emerge. B-G) At a high self-efficacy learning rate, higher state values around the rewarding terminal state transfer to other non-terminal states (C-D) and then shift again to new locations (E-F) which eventually collapse to an attractor-like state (G) where the agent is 'trapped', exponentially increasing the self-efficacy level and state values as it repeats its actions within attractor-like neighboring states. The arrows in the gridworld represent the learned best action in each state.

can be interpreted as a reward function which couples extrinsic and intrinsic rewards. Higher self-efficacy leads to greater valuation of actions, while successful action outcomes further enhance self-efficacy levels. The agent in our dynamic self-efficacy model maximizes a reward signal that encompasses both intrinsic rewards such as satisfaction from mastery and extrinsic rewards such as tangible outcomes. Similar designs have been proposed in machine learning, where intrinsically-motivated RL systems have been designed to create autonomous agents capable of solving wide ranges of complex problems (Aubret et al., 2019; Barto, 2013).

Related Work

Several researchers have examined the relationship between efficacy and exertion of effort in pursuit of desired outcomes using the Expected Value of Control (EVC) framework (Shenhav et al., 2021). EVC defines efficacy through the lens of control. In this view, one's efficacy is composed of control efficacy – the extent to which higher levels of control translate into better performance; and performance efficacy – the extent to which better performance can translate into better outcomes. Experimental tests of EVC theory showed that people dynamically update expectations of their performance efficacy based on feedback (Grahek et al., 2023).

In our model, self-efficacy is an independent construct, such that an agent can represent how well they can accomplish a task without first considering how much effort they need to put in for control. Rather than using past efficacy estimates to update beliefs, the agent uses reward prediction errors (RPEs) as the main input for updating self-efficacy. In this way, self-efficacy beliefs arise directly through the appraisal of goal attainment, allowing us to test how beliefs about self-efficacy are shaped through the interplay of the agent's actions with feedback from the environment.

Limitations and future directions

Although our simulations provide meaningful insights, they significantly differ from human contexts. Additionally, while our model can explain some core symptoms of mania as discussed above, it does not account for all symptoms of mania, such as pressured speech, decreased need for sleep or racing thoughts. Future work will seek to expand the model to encompass other symptoms associated with mania and test the model predictions using empirical data from human studies.

Acknowledgments

The authors would like to thank Sheri L. Johnson for valuable discussions.

References

- Aubret, A., Matignon, L., & Hassas, S. (2019). A survey on intrinsic motivation in reinforcement learning. (Publisher: arXiv Version Number: 2)
- Bandura, A. (1997). *Self efficacy: the exercise of control*. New York (N.Y.): W. H. Freeman.
- Barto, A. G. (2013). Intrinsic Motivation and Reinforcement Learning. In G. Baldassarre & M. Mirolli (Eds.), *Intrinsically Motivated Learning in Natural and Artificial Systems* (pp. 17–47). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Chew, B., Blain, B., Dolan, R. J., & Rutledge, R. B. (2021). A Neurocomputational Model for Intrinsic Reward. *The Journal of Neuroscience*, 41(43), 8963–8971.
- Diagnostic and statistical manual of mental disorders: DSM-5* (5th ed ed.). (2013). Washington, D.C: American Psychiatric Association.
- Dulberg, Z., Dubey, R., Berwian, I. M., & Cohen, J. D. (2023). Having multiple selves helps learning agents explore and adapt in complex changing worlds. *Proceedings of the National Academy of Sciences*, 120(28), e2221180120.
- Eldar, E., & Niv, Y. (2015). Interaction between emotional state and learning underlies mood instability. *Nature Communications*, 6(1), 6149.
- Frömer, R., Lin, H., Dean Wolf, C. K., Inzlicht, M., & Shenhav, A. (2021). Expectations of reward and efficacy guide cognitive control allocation. *Nature Communications*, 12(1), 1030.
- Gagne, C., & Dayan, P. (2022). Peril, prudence and planning as risk, avoidance and worry. *Journal of Mathematical Psychology*, 106, 102617.
- Goodwin, F. K., & Jamison, K. R. (2007). *Manic-depressive illness: bipolar disorders and recurrent depression* (2nd ed ed.). New York: Oxford University Press.
- Grahek, I., Frömer, R., Prater Fahey, M., & Shenhav, A. (2023). Learning when effort matters: neural dynamics underlying updating and adaptation to changes in performance efficacy. *Cerebral Cortex*, 33(5), 2395–2411.
- Harhen, N. C., & Bornstein, A. M. (2024). Interval Timing as a Computational Pathway From Early Life Adversity to Affective Disorders. *Topics in Cognitive Science*, 16(1), 92–112.
- Harmon-Jones, E., Abramson, L. Y., Nusslock, R., Sigelman, J. D., Urosevic, S., Turonie, L. D., ... Fearn, M. (2008). Effect of Bipolar Disorder on Left Frontal Cortical Responses to Goals Differing in Valence and Task Difficulty. *Biological Psychiatry*, 63(7), 693–698.
- Hayden, E. P., Bodkins, M., Brenner, C., Shekhar, A., Nurnberger, J. I., O'Donnell, B., & Hetrick, W. P. (2008). A multimethod investigation of the Behavioral Activation System in bipolar disorder. *Journal of Abnormal Psychology*, 117(1), 164–170.
- Johnson, S. (2005). Mania and dysregulation in goal pursuit: a review. *Clinical Psychology Review*, 25(2), 241–262.
- Johnson, S., Edge, M. D., Holmes, M. K., & Carver, C. S. (2012). The Behavioral Activation System and Mania. *Annual Review of Clinical Psychology*, 8(1), 243–267.
- Johnson, S., Swerdlow, B. A., Treadway, M., Tharp, J. A., & Carver, C. S. (2017). Willingness to Expend Effort Toward Reward and Extreme Ambitions in Bipolar I Disorder. *Clinical Psychological Science*, 5(6), 943–951.
- Jones, S., Mansell, W., & Waller, L. (2006). Appraisal of hypomania-relevant experiences: Development of a questionnaire to assess positive self-dispositional appraisals in bipolar and behavioural high risk samples. *Journal of Affective Disorders*, 93(1-3), 19–28.
- Kendler, K. S. (2017). The clinical features of mania and their representation in modern diagnostic criteria. *Psychological Medicine*, 47(6), 1013–1029.
- Kesebir, S., İřitmez, S., & Gündoğar, D. (2012). Compulsive buying in bipolar disorder: Is it a comorbidity or a complication? *Journal of Affective Disorders*, 136(3), 797–802.
- Mason, B., Brown, E., & Croarkin, P. (2016). Historical underpinnings of bipolar disorder diagnostic criteria. *Behavioral Sciences*, 6(3), 14.
- Shenhav, A., Botvinick, M., & Cohen, J. (2013). The Expected Value of Control: An Integrative Theory of Anterior Cingulate Cortex Function. *Neuron*, 79(2), 217–240.
- Shenhav, A., Prater Fahey, M., & Grahek, I. (2021). Decomposing the Motivation to Exert Mental Effort. *Current Directions in Psychological Science*, 30(4).
- Stange, J. P., Molz, A. R., Black, C. L., Shapero, B. G., Baccelli, J. M., Abramson, L. Y., & Alloy, L. B. (2012). Positive overgeneralization and Behavioral Approach System (BAS) sensitivity interact to predict prospective increases in hypomanic symptoms: A behavioral high-risk design. *Behaviour Research and Therapy*, 50(4), 231–239.
- Themanson, J. R., & Rosen, P. J. (2015). Examining the relationships between self-efficacy, task-relevant attentional control, and task performance: Evidence from event-related brain potentials. *British Journal of Psychology*, 106(2), 253–271.
- Zorowitz, S., Momennejad, I., & Daw, N. D. (2020). Anxiety, Avoidance, and Sequential Evaluation. *Computational Psychiatry*, 4(0), 1.