A dual-receptor model of serotonergic psychedelics: therapeutic insights from simulated cortical dynamics

Arthur Juliani
Microsoft Research, New York, NY, USA
ajuliani@microsoft.com

Veronica Chelu
McGill University, Montreal, Quebec, Canada

Laura Graesser
Montreal, Quebec, Canada

Adam Safron
Johns Hopkins University, Baltimore, MD, USA

Abstract

Serotonergic psychedelics have been identified as promising next-generation therapeutic agents in the treatment of mood and anxiety disorders. While their efficacy has been increasingly validated, the mechanism by which they exert a therapeutic effect is still debated. A popular theoretical account is that excessive 5-HT2a agonism disrupts cortical dynamics, relaxing the precision of maladaptive high-level beliefs, thus making them more malleable and open to revision. We extend this perspective by developing a theoretical framework and simulations based on predictive processing and an energy-based model of cortical dynamics. We consider the role of both 5-HT2a and 5-HT1a agonism, characterizing 5-HT2a agonism as inducing stochastic perturbations of the energy function underlying cortical dynamics and 5-HT1a agonism as inducing a global smoothing of that function. Within our simulations, we find that while both agonists are able to provide a significant therapeutic effect individually, mixed agonists provide both a more psychologically tolerable acute experience and better therapeutic efficacy than either pure 5-HT2a or 5-HT1a agonists alone. This finding provides a potential theoretical basis for the clinical success of LSD, psilocybin, and DMT, all of which are mixed serotonin agonists. Our results furthermore indicate that exploring the design space of biased 5-HT1a agonist psychedelics such as 5-MeO-DMT may prove fruitful in the development of even more effective and tolerable psychotherapeutic agents in the future.

1 Introduction

Serotonergic psychedelics such as psilocybin, LSD, and DMT have received significant attention from both research scientists and clinicians in recent years for their potential to treat a variety of psychiatric conditions ranging from depression and anxiety to substance use and obsessive compulsive disorders (M. W. Johnson, Hendricks, Barrett, & Griffiths, 2019; Sessa, 2018; Griffiths et al., 2016; M. W. Johnson & Griffiths, 2017; R. Carhart-Harris et al., 2021; Timmermann, Zeifman, Erritzoe, Nutt, & Carhart-Harris, 2024). This transdiagnostic efficacy has led some researchers to hypothesize that there may be a single primary underlying factor of psychopathology which psychedelic therapy is acting to address (R. Carhart-Harris et al., 2022). Despite significant progress in understanding the empirical therapeutic effects of these drugs, theoretical models which describe and predict these effects are less well developed. While the key role of the 5-HT2a receptor system and downstream neurotrophic effects are largely accepted by the scientific community (Shao et al., 2021; Calder & Hasler, 2023; Banushi & Polito, 2023), the role of subjective experience in the therapeutic effects of psychedelics is still a topic of considerable debate (Yaden & Griffiths, 2020; Cameron et al., 2021).

A variety of competing theories have been developed to explain the mechanics behind the subjective effects of psychedelics and their relationship to therapeutic outcomes (Doss et al., 2022). Among these, the *Relaxed Beliefs Under Psychedelics* (REBUS) model (R. L. Carhart-Harris & Friston, 2019) has received additional attention outside the domain of neuroscience for its translational application in guiding protocol development for psychedelic-assisted psychotherapy. The REBUS model hypothesizes that psychedelics exert their therapeutic effect by relaxing the precision of high-level beliefs both acutely and post-acutely, thus making them amenable to modification through introspection and interpersonal therapy. This hypothesis is supported by evidence of an increase in the entropy of cortical transition dynamics during the acute phase of psychedelic use (Singleton et al., 2022), and a relaxation of long-term propositional beliefs after psychedelic therapy (Zeifman et al., 2022).

Despite the existence of some preliminary evidence in support of REBUS, it is not clear whether we should expect the relaxation of beliefs to hold true across the dose-response curve for all substances in the broad class of psychedelics, or whether certain doses and environmental conditions may result in the transient strengthening of beliefs under psychedelics instead. One piece of evidence which would suggest that psychedelics may strengthen beliefs in some cases are reports of insight-experiences (Laukkonen et al., 2023; McGovern et al., 2023), in which individuals arrive a strongly-felt (but potentially erroneous) new beliefs. Another is the experience of so-called “entities” and “alternate realities” under the effects of DMT, which are often described as appearing more real than normal conscious experience (Luan et al., 2023). Furthermore, although the purported mechanism behind the disruption of belief representation is 5-HT2a agonism, there is a non-trivial contribution by other serotonin (5-HT) receptor populations in the changes to neural dynamics which result from psychedelics, with activity at 5-HT1a receptors in particular exerting a significant effect (Deco et al., 2018; Singleton et al., 2022).

The recently introduced *Altered Beliefs Under Psychedelics* (ALBUS) model predicts that psychedelics

alter belief representations in a more non-specific way than REBUS, with the therapeutic belief relaxation seen in clinical studies being a special case which results from specific dose ranges and controlled environmental influences (Safron, 2020). While this alternative hypothesis is based on a study of the theoretical non-linear dynamics in the underlying activity of the relevant circuits in the cortex, it currently lacks empirical validation and is under-specified in a number of ways which limit its practical application as a predictive tool. In this work we begin to address these limitations.

Utilizing a novel simulation paradigm inspired by iterative optimization algorithms and built on the framework provided by ALBUS and hierarchical predictive processing (HPP), we contribute additional theoretical support for the possibility of both belief relaxation and belief strengthening across the dose-response curve of serotonergic psychedelics. Our model and simulation framework has the potential to help answer several important questions about the neuromodulatory effect of serotonin (5-HT) signaling on cortical belief representation. We illustrate that the markers of brain activity characteristic of the psychedelic state and associated with 5-HT2a agonism such as increases in entropy and complexity of cortical states are consistent with both a relaxation as well as a stochastic and transient strengthening of beliefs during an acute drug administration. We demonstrate that within our model this transient strengthening of beliefs is theoretically capable of producing therapeutic outcomes comparable or greater than those of more straightforward belief relaxation.

We also model the neuromodulatory role that 5-HT1a agonism plays in psychedelic effects, both with respect to their acute phenomenology as well as to their long-term clinical efficacy in addressing psychopathology. Psilocybin, LSD, and DMT all have significant affinity for the 5-HT1a receptor in addition to the 5-HT2a receptor (Rickli, Moning, Hoener, & Liechti, 2016). Rather than serving a minor or even completely negligible role, as has been previously assumed, we provide a review of empirical evidence along with simulations under the proposed model which both provide give insight into the key role 5-HT1a agonism may play in psychedelic experiences and therapeutic outcomes. In this work we characterize 5-HT2a agonism as producing disruptions in belief representations through the perturbation of the underlying optimization landscape of the predictively encoded input signal. In contrast, 5-HT1a agonism provides a complementary regularizing effect which prevents the development of overly-precise beliefs during the acute phase of the psychedelic experience. Based on clinical evidence and simulation results, we hypothesize that although 5-HT2a agonism is necessary and sufficient for the psychedelic experience, 5-HT1a agonism also plays a role in modulating the experience such that it is both more tolerable and therapeutic.

Our model allows for an interpretation of the unique phenomenology of the highly biased 5-HT1a agonist psychedelic 5-MeO-DMT, which has been described as producing a “white-out” or “void” state (Gómez-Emilsson, 2020). Within our framework, the 5-MeO-DMT experience corresponding to a state in which belief representation is temporarily regularized to an extreme point where beliefs are represented with a near-zero precision due to the high levels of 5-HT1a agonism. The fact that this drug produces one of the most significant and durable antidepressant effects (Uthaug et al., 2019; Reckweg et al., 2023) is both support for the REBUS model as well as a major complication for the hypothesis that 5-HT2a agonism is the only relevant drug factor in psychedelic-assisted psychotherapy. Properly characterizing the modula-

tory role played by 5-HT1a agonism in the psychedelic experience may be essential to the development of next-generation psychedelic substances which provide both greater psychological tolerability as well as more consistent long-term therapeutic effects.

The rest of this article is organized as follows: We first introduce a guiding framework for the endogenous function of serotonin signalling (Section 1.1) as well as experimental and clinical evidence in support of the modulating roles of 5-HT2a and 5-HT1a agonism in the acute effects and therapeutic efficacy of serotonergic psychedelics (Section 1.2). We then present a review of the REBUS model and predictive processing (Section 1.3), followed by a discussion of the key role of the prefrontal cortex in psychedelic activity (Section 1.4). We next introduce our theoretical model based on the principles of HPP and energy-based models (EBMs) (Section 1.5), within which we propose specific mechanisms for the neuromodulatory effects of 5-HT2a and 5-HT1a agonism on cortical dynamics (Section 2). We follow this with results from idealized simulations of these dynamics, demonstrating their ability to predict potential therapeutic outcomes, markers of cortical entropy, and drug tolerability (Section 3). Finally, we discuss the implications of our model, as well as the experimental evidence which would need to be collected to validate or invalidate our predictions (Section 4).

1.1 Serotonergic neuromodulation as a stress response system

A well-supported model of the serotonin system is that it serves to mediate cognitive and affective responses to stress (Puglisi-Allegra & Andolina, 2015; Murnane, 2019; Zhong, Li, Feng, & Luo, 2017; Hayashi, Nakao, & Nakamura, 2015). This can be seen as an extension of a broader evolutionarily preserved role for serotonin in responding to aversive stimuli (Dayan & Huys, 2009). A prominent model within the paradigm of serotonin-mediated stress response is that relevant serotonin signaling is composed of two unique interdependent systems, corresponding to the two main serotonin receptor populations in the brain. Within the model, these receptor systems instantiate separate “active” and “passive” coping strategies which an organism may engage under stress (R. L. Carhart-Harris & Nutt, 2017). These two systems have independently been identified as “automaticity” and “flux” cognitive states (Shine et al., 2022), which we interpret here as being functionally equivalent to the two coping strategies outlined by R. L. Carhart-Harris and Nutt.

The human nervous system contains over a dozen unique serotonin receptors (Nichols & Nichols, 2008). Of these, 5-HT1a and 5-HT2a are both the most widely distributed within the brain and the most well studied (Beliveau et al., 2017). Of particular interest are the role of the postsynaptically expressed 5-HT1a and 5-HT2a receptors in the cortex. These receptor populations are understood to operate in relatively simple opponency with one another within cortical pyramidal cells, with 5-HT1a receptors inhibiting and 5-HT2a receptors exciting the postsynaptic cell (Aghajanian & Marek, 1997; Puig & Gulledge, 2011). The active and passive coping strategies within the brain are hypothesized to be instantiated by 5-HT2a and 5-HT1a receptor systems, respectively (R. L. Carhart-Harris & Nutt, 2017; Shine et al., 2022).

Within this dual-strategy framework, small to moderate amounts of stress activate the 5-HT1a medi-

ated passive coping system while large amounts of stress activate the 5-HT2a mediated active coping system. In passive coping the animal maintains the current behavioral policy but modulates affect in response to manageable levels of stress. In contrast, active coping is engaged when the level of stress is significant enough to require the instantiation of novel and divergent internal models, beliefs, or behavioral policies (Brouwer & Carhart-Harris, 2021; Safron & Sheikhbahaee, 2021). This bi-modal behavioral distribution can be seen in a simplified form in serotonin-mediated rodent responses to stress from predatory threats, where a low-level threat initiates freezing behavior while a high-level threat initiates fleeing behavior (Seo et al., 2019).

Importantly, both coping strategies are mediated by the same system of serotonergic neuromodulation. Within a stress-response model of serotonin function, the extent to which a given strategy is deployed is a function of 5-HT release in the dorsal raphae nucleus (DRN). A complete account of the computational role of the DRN is still being developed, but evidence exists that DRN neurons may compute unsigned prediction errors (Matias, Lottem, Dugué, & Mainen, 2017), and that the magnitude of this signal corresponds to downstream 5-HT release in the cortex. This prediction error computation is hypothesized to be part of a larger role for the DRN in value prediction complementary to that of the ventral tegmental area (Harkin, Grossman, Cohen, Béïque, & Naud, 2023). The kind of stress which serotonin acts in response to may therefore be mediated by the ability of the organism to form and maintain successfully predictive representations.

Within this framework, 5-HT neuromodulation then serves a role in responding to stress by improving the computational efficiency of the processes required to learn and maintain such successful representations. The ability to efficiently learn and maintain adaptive predictive representations is precisely the construct of cognitive flexibility (Ionescu, 2012)—the characteristic that enables animals or humans to adaptively generate appropriate behavioral responses based on changing sensory stimuli—which psychedelics have been demonstrated to improve (Doss et al., 2021; Torrado Pacheco, Olson, Garza, & Moghaddam, 2023), and is believed to be generally associated with serotonin signalling (Matias et al., 2017).

Given that all classic psychedelics have significant affinity for the 5-HT2a receptor system, they have been characterized as a prototypical example of a substance that induces an “active coping” response in an organism by R. L. Carhart-Harris and Nutt. This manifests in the acute psychedelic state which enables a state of “metaplasticity”—the dynamic regulation of the extent to which synaptic plasticity can be induced (Abraham & Bear, 1996; Nardou et al., 2023), enabling a re-evaluation of previously held beliefs and the spontaneous assumption of novel representations of self or environment (for a review of this effect across various forms of belief, see (Letheby, 2021)). In contrast, R. L. Carhart-Harris and Nutt propose MDMA as an exemplar “passive coping” substance due to its role in releasing endogenous serotonin, which preferentially activates the 5-HT1a receptor system. In humans this is hypothesized to manifest most acutely in the felt sense of equanimity and acceptance that characterizes the unique phenomenology of the MDMA (Greer & Tolbert, 1990; R. L. Carhart-Harris & Nutt, 2017).

A complication to this straightforward account is that psilocybin, LSD, and DMT all bind with significant affinity to 5-HT1a receptors as well 5-HT2a receptors (Rickli et al., 2016) (See Figure 1). Given the

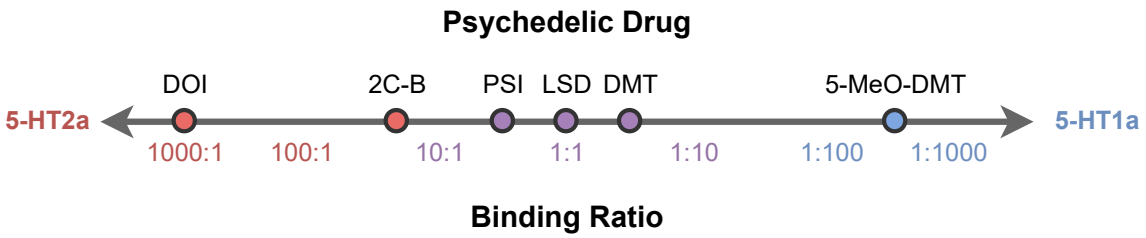

Drug	DOI	2C-B	Psilocybin	LSD	DMT	5-MeO-DMT
5-HT1a affinity	2.355	0.240	0.123	0.003	0.075	0.003
5-HT2a affinity	0.001	0.009	0.049	0.004	0.237	0.907
5HT2a:5-HT1a ratio	3463	27.91	2.510	0.750	0.316	0.003

Figure 1: Relative 5-HT1a and 5-HT2a binding affinities $K_i(\mu M)$ for select psychedelic substances. Lower values correspond to higher binding affinity. DOI has the highest relative 5-HT2a binding affinity, while 5-MeO-DMT has the highest relative 5-HT1a binding affinity. Data sourced from (Halberstadt et al., 2012; Rickli et al., 2015, 2016).

strong positive correlation between receptor binding affinity and efficacy for this class of drugs (Halberstadt, Chatha, Klein, Wallach, & Brandt, 2020), one would expect them to engage both coping systems simultaneously, as opposed to only engaging the active coping system. Below we present a review of existing evidence for the roles of both 5-HT2a and 5-HT1a agonism in the acute effects of psychedelics.

1.2 5-HT2a and 5-HT1a receptor agonism in psychedelics

5-HT2a agonism has been demonstrated to be necessary and sufficient for inducing the "head twitch response" in rodents (Halberstadt, Koedood, Powell, & Geyer, 2011), and this effect is highly correlated with the existence of reportable psychedelic phenomenology in humans (Halberstadt et al., 2020). Despite this, there is also evidence that agonism at the 5-HT1a receptor is responsible for significantly modulating the 5-HT2a mediated effects of serotonergic psychedelics. Here we consider the evidence for this modulating effect across a number of dimensions including behavioral, neural, clinical, and phenomenological.

On the phenomenological level, 5-HT1a agonism has been identified as being the primary driver of stimulus control in rodents trained to discriminate 5-MeO-DMT (Winter, Filipink, Timineri, Helsley, & Rabin, 2000), and there is evidence that it also plays a significant role in the discrimination of LSD and DMT (Reissig, Eckler, Rabin, & Winter, 2005), though not psilocybin (Winter, Rice, Amorosi, & Rabin, 2007) or DOB (a strongly biased 5-HT2a agonist, similar to DOI) (Benneyworth, Smith, Barrett, & Sanders-Bush, 2005). The role of 5-HT1a agonism in stimulus control in these examples is notably correlated with the relative 5-HT2a and 5-HT1a receptor affinities of each drug (Rickli et al., 2016) (see Figure 1), suggesting an underlying pattern whereby psychedelics with greater relative 5-HT1a agonism also manifest phenomenological effects which are more dependent on that agonism.

Despite the fact that the 5-HT1a agonism of psilocybin does not induce stimulus control, there is still

evidence for the role it plays in the drug's behavioral effects, specifically with respect to the modulation of attentional control (Carter et al., 2005), compulsive behavior (Singh et al., 2023), and exploration (Halberstadt et al., 2011). The later study demonstrated that 5-HT1a agonism is likewise primarily implicated in the behavioral and phenomenological effects of 5-MeO-DMT, which is unsurprising, given the high ratio of 5-HT1a to 5-HT2a affinity in this substance.

The behavioral effects in rodents treated with 5-MeO-DMT are accompanied by changes in cortical dynamics in both PFC and visual cortex, both of which are mediated primarily by 5-HT1a agonism (Riga, Soria, Tudela, Artigas, & Celada, 2014; Riga, Bortolozzi, Campa, Artigas, & Celada, 2016; Riga, Lladó-Pelfort, Artigas, & Celada, 2018). One consistent marker of these disruptions is a decrease in fMRI blood-oxygen-level-dependent (BOLD) signal, which is consistent with the inhibitory role of postsynaptic 5-HT1a agonism, particularly in the PFC. 5-HT1a receptor binding maps are also significantly more predictive than any other serotonin binding map except for 5-HT2a in anticipating the changes to brain activity under LSD (Deco et al., 2018), as well as for psilocybin (Singleton et al., 2022).

A prominent theory of the therapeutic efficacy of psychedelics is that they act as powerful psychoplastogens (Olson, 2018), or neurogenesis-inducing agents (Shao et al., 2021; Calder & Hasler, 2023). Although this effect is well documented, there is some debate concerning the mechanisms by which it takes place. Given the critical role 5-HT2a agonism plays in the phenomenological effects of psychedelics, it was historically hypothesized that 5-HT2a agonism may be responsible for the plasticity inducing effects and their corresponding therapeutic benefits as well (Ly et al., 2018). Recent work has refined this hypothesis by implicating intracellular 5-HT2a agonism in particular in neurogenesis (Vargas et al., 2023), thus providing an explanation for why endogenous serotonin does not act as a psychoplastogen as it is only capable of extracellular 5-HT2a agonism.

The hypothesis of 5-HT2a mediated plasticity is complicated by other research implicating 5-HT1a agonism in downstream neurogenesis (Banasr, Hery, Printemps, & Daszuta, 2004; Schreiber & Newman-Tancredi, 2014; Islam, Moriguchi, Tagashira, & Fukunaga, 2014). It has also been shown that a moderate knockdown of 5-HT2A receptor availability disrupts the acute behavioral effect of psilocybin without abolishing its effect on structural plasticity (Shao et al., 2021). Likewise, research into the antidepressant effects of the dissociative drug ketamine have identified a key role of downstream 5-HT1a agonism in PFC neurogenesis (Fukumoto, Iijima, Funakoshi, & Chaki, 2018; Fukumoto et al., 2020), which is a believed to be a key factor in the therapeutic effect of the drug (Deyama & Duman, 2020). Other recent work suggests that the psychoplastogen effects of psychedelics may be largely serotonin-independent and the result of direct binding to the Brain-Derived Neurotrophic Factor (BDNF)—a key neurotrophin involved in synaptic plasticity—receptor TrkB (Tropomyosin Receptor Kinase B) (Moliner et al., 2023), whose activation plays a direct role in inducing neuroplasticity.

Importantly, there is also evidence that 5-HT2a agonism may not play an essential role in the antidepressant effects of psychedelics, even independent of its effects on neuroplasticity. DOI, a highly biased 5-HT2a agonists, has failed to produce robust therapeutic effects in a rodent model of depression (Qu, Chang, Ma, Wan, & Hashimoto, 2023). There is also mixed evidence for the efficacy of DOI as a potential

treatment of anxiety disorders, with a number of studies finding an anxiolytic effect (Dhonnchadha, Hascoët, Jolliet, & Bourin, 2003; Massé, Hascoët, Dailly, & Bourin, 2006), while others finding an anxiogenic effect in rodent models (Masse, Petit-Demouliere, Dubois, Hascoët, & Bourin, 2008; De Paula, Torricelli, Lopreato, Nascimento, & Viana, 2012). In contrast, it has been demonstrated that psilocybin is capable of producing antidepressant effects in rodent models even when the acute effects of 5-HT2a agonism are partially blocked (Hesselgrave, Troppoli, Wulff, Cole, & Thompson, 2021; Shao et al., 2021), potentially implicating 5-HT1a agonism in these effects.

Another proposed mediator of the antidepressant effects of psychedelics is a post-acute increase in cognitive flexibility (Doss et al., 2021). A review of the literature suggests that cognitive flexibility is more frequently associated with preferential 5-HT1a agonism rather than 5-HT2a agonism (Depoortère et al., 2010; Alvarez, Morales, & Amodeo, 2021; Odland, Kristensen, & Andreasen, 2021). Recent evidence also suggests that blocking 5-HT1a agonism in psilocybin impairs the drug's positive effects on cognitive flexibility in a rodent model of anorexia nervosa (Conn et al., 2023), suggesting the important mediating role that 5-HT1a agonism plays in the drug's effects.

Finally, serotonin itself demonstrates a significantly biased agonism for 5-HT1a over 5-HT2a receptors (Barnes & Sharp, 1999), and there are commonalities between the phenomenological effects of serotonin releasing drugs and 5-HT1a agonists. In a study of the effects of the serotonin releasing drug MDMA, 5-HT2a agonism was found to play only a negligible role in the acute phenomenology (Liechti, Saur, Gamma, Hell, & Vollenweider, 2000), suggesting in contrast a more prominent role of 5-HT1a agonism in the subjective effects of the drug. There is also a research literature on the anxiolytic effects of biased 5-HT1a agonists, such as 8-OH-DPAT, in rodent models of psychopathology (Luscombe, Martin, Hutchins, Gosden, & Heal, 1993; Schreiber & De Vry, 1993), and these effects are consistent with those of MDMA and other 5-HT releasing drugs. Reductions in anxiety and avoidance are key mediators in a variety of common psychopathologies and their treatment by psychedelic therapy (Zeifman, Wagner, Monson, & Carhart-Harris, 2023), suggesting a potential link between these effects and 5-HT1a agonism. Notably, outcomes from MDMA assisted therapy are characterized by a significant acute decrease in experiential avoidance (Mitchell et al., 2023).

Taken together, it is clear that at the very least 5-HT1a agonism plays a non-trivial modulating role in the acute phenomenology and long-term efficacy of serotonergic psychedelics. The question that then arises is in what way 5-HT1a agonism may contribute to these effects when the most apparent changes in subjective experience are mediated by 5-HT2a agonism (Preller et al., 2017). Utilizing an HPP framework of inference and learning in the PFC under psychedelic effects as outlined by R. L. Carhart-Harris and Friston and extended by Safron, we hypothesize that 5-HT1a agonism impacts cortical belief representation by regularization toward a flat prior and thus smoothing the optimization landscape underlying predictive coding dynamics. In contrast, we posit the mechanism of action of 5-HT2a agonism is to introduce stochastic perturbations into the neural optimization landscape which produce transient changes resulting in increased plasticity. The predictions of such a model are consistent with the anxiolytic, sedating, and prosocial effects (Luscombe et al., 1993; Thompson, Callaghan, Hunt, Cornish, & McGregor, 2007)

of 5-HT1a agonism as encapsulated in the construct of "passive coping" as well as the stimulating, pareidolic, and insight-inducing effects of 5-HT2a agonism as encapsulated in the construct of "active coping". Only when both sets of effects are taken together can there be a full characterization of the psychedelic experience in humans and other animals (M. W. Johnson et al., 2019; Braun et al., 2024).

1.3 Predictive processing and the REBUS model

The REBUS model is rooted in the idea that the brain uses an internal generative model of the world to predict sensory input based on movement and past sensory experience (R. L. Carhart-Harris & Friston, 2019). This family of theories, collectively referred to as the predictive processing (PP) framework (Clark, 2013, 2016; Keller & Mrsic-Flogel, 2018; Kanai, Komura, Shipp, & Friston, 2015b), includes predictive coding (Huang & Rao, 2011; Chalk, Marre, & Tkavcik, 2018), hierarchical temporal memory and Bayesian inference (K. Friston & Kiebel, 2009; Kanai, Komura, Shipp, & Friston, 2015a). Common to each of these variations of the theory is the idea that an error signal between the predicted and actual sensory input is used to update an internal representation of the world. Within the Bayesian brain perspective in particular (Knill & Pouget, 2004; Doya, Ishii, Pouget, & Rao, 2006), these internal representations are referred to as beliefs.

PP is often described as a processing hierarchy, and referred to as hierarchical predictive processing (HPP) in the context of the human cortex, whereby higher levels of this hierarchy of functional brain areas send top-down signals to a lower areas in the form of predictions of the bottom-up stimulus input to that area. PP need not follow a strict hierarchical arrangement of inter-areal connections (Keller & Mrsic-Flogel, 2018), the connections can also be lateral described by a densely recurrent network (R. P. N. Rao & Ballard, 1999; Buesing, Bill, Nessler, & Maass, 2011) (see Figure 2 for an illustration of a representational layer in the HPP model).

Within this theoretical framework, prediction errors are dynamically adjusted using a gain modulation mechanism which gates neural plasticity in a context-dependent manner (Behrens, Woolrich, Walton, & Rushworth, 2007; Wilmes & Clopath, 2023; Sadeh, Clopath, Palmer, Frank, & Ziv, 2022). Specifically, the neuromodulatory tone may shift the relative contribution of bottom-up and top-down signals and consequently alter the precision or sensitivity by which cortical circuits prioritize and respond to sensory information or report prediction errors. Representations encoded with greater precision are thus able to exert a greater influence on both upstream and downstream computation. The influence of prediction errors can be modulated according to the internal latent cortical representation of the sensory stimuli, which would determine the extent to which bottom-up inputs are used to update the internal model, or how much the internal representation is driven by top-down predictions.

While predictive processing focuses on the brain's generation and adjustment of predictions to minimize prediction errors, the free energy principle (FEP) (K. Friston, 2010) which inspired the REBUS model, provides a broader theoretical perspective, framing these processes as part of a larger imperative for biological systems to minimize "free energy" and, by extension, surprise or unpredictability in their interactions

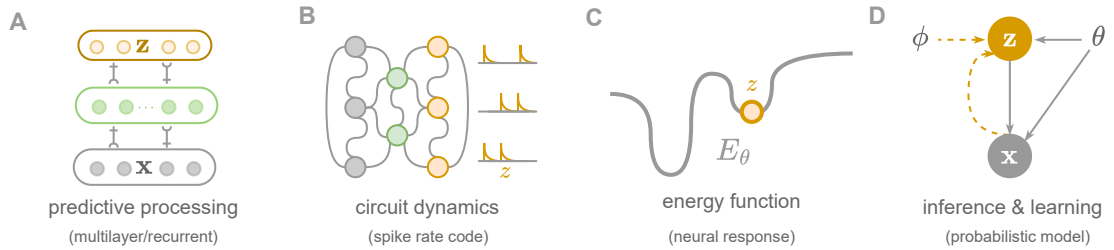

Figure 2: **(A) Inter-Areal Communication in the HPP Framework**. The synaptic connections between the neural response vector $\mathbf{z}$ (a vector of firing rates) and the input stimulus $\mathbf{x}$ is updated based on a comparison of top-down prediction and bottom-up input. Prediction errors are sent forward in the hierarchy, while predictions are sent backward. **(B) Population dynamics in the PFC.** Illustrative representation of neural population dynamics in the PFC and associated rate spike code. **(C) Energy function.** A prototypical representation of a 1D energy function induced by synaptic connectivity structure θ, along with a sampled neural response z. **(D) The directed graphical model under consideration.** Solid lines denote the generative model $p_\theta(\mathbf{z})p_\theta(\mathbf{x}|\mathbf{z})$, dashed lines denote the variational approximation $p_\phi(\mathbf{z}|\mathbf{x})$ to the intractable posterior $p_\theta(\mathbf{z}|\mathbf{x})$, which uses the variational parameters ϕ.

with the environment. When there is a mismatch between predictions and sensory input, a discrepancy between top-down and bottom-up inputs generates prediction errors (cf. Figure 2-A). Minimizing these prediction errors aligns with the "free energy" principle's goal of reducing surprise or uncertainty.

According to the REBUS model, the spontaneous activation of neurons due to increased postsynaptic 5-HT2a agonism results in a desynchronization of those neuronal populations, and thus a disruption in their ability to robustly encode high-level beliefs or internal representations of the world. This disruption is equivalent to a decrease in the precision of the beliefs which these representations encode, thus making them also more malleable. Under the REBUS formulation, psychopathology pertaining cognitive inflexibility can be interpreted as inference under aberrant precision with respect to strong, rigid priors (R. Carhart-Harris et al., 2022; Juliani, Safron, & Kanai, 2024). 5-HT2a modulation by psychedelics is thought to relax overly precise prior precision and enable the revision of maladaptive priors through experience-dependent learning.

At the neural level, the disruption of high-level functional connectivity induced by psychedelics is accompanied by a decrease in activity and correlation within and between these networks, particularly those involved in self-referential processing. This includes regions of the default-mode network (DMN), a brain network that shows increased endogenous fluctuations when an individual is awake but not engaged in a task requiring attention (Kometer, Pokorny, Seifritz, & Volleinweider, 2015; M. W. Johnson et al., 2019), a reconfiguration of communication in the brain characterized by increased brain activity, increased diversity of neural activation patterns compared to normal waking consciousness, elevated Shannon-Boltzmann entropy (Shannon, 1948) of intra-brain-network synchrony (R. L. Carhart-Harris et al., 2014), induced "metaplasticity"—thought to be one of the mechanisms underlying the establishment of critical periods

(Abraham & Bear, 1996; Nardou et al., 2023).

We can draw a link between an organism's ability to minimize prediction errors, which can be interpreted as a definition of free energy within HPP, and its long-term behavioral success (Kirchhoff, Parr, Palacios, Friston, & Kiverstein, 2018). Theoretical work has further linked this success in minimizing prediction errors to not only successful behavioral outcomes, but also positive mental health outcomes (Smith, Badcock, & Friston, 2021). Large prediction errors in cortical representation corresponds to the agent's inability to resolve the uncertainty in an environment.

There is both theoretical and empirical research connecting a failure of predictive error minimization to negatively valenced affect in an individual's subjective experience (Smith et al., 2022). Greater felt uncertainty for example is often characteristic of anxiety and mood disorders (Carleton et al., 2012; Smith, Kirlic, et al., 2021). This is particularly the case when the inability to resolve prediction errors takes place at higher levels of the cortical hierarchy, as this suggests a fundamental maladaptivity in the foundational beliefs or attentional policies of the organism. Within this work we use this connection between prediction error and negatively valenced affect to derive a measure of the potential therapeutic efficacy of a given pharmacological intervention.

A critical characteristic of a successful learning system is robustness to changing tasks, contexts, and environments (Thrun, 1998; Parisi, Kemker, Part, Kanan, & Wermter, 2019). A maladaptive neural representation is one which is locally optimal with respect to the current environmental context, characterized by a particular history-dependent frame of reference, but sub-optimal with respect to future contextual frames of reference (DeYoung, 2015). The canalization model of psychopathology characterizes the tendency to be caught in high-precision local optima of the cortical optimization landscape as playing a causal role in a variety of psychopathologies (R. Carhart-Harris et al., 2022). A similar "plasticity loss" phenomenon has been observed in adaptive learning algorithms in the context of continual learning in non-stationary environments (Lyle et al., 2023).

An implication of connecting prediction error minimization with behavioral fitness is that confident neural responses are not inherently pathological if they are both relevant to the current environmental context and also adaptable to future contexts (Juliani et al., 2024). Indeed, such high-precision, stable, and adapted representations may provide a buffer against various forms of psychopathology (DeYoung & Krueger, 2018). The inverse is also true that even if certain encoded predictions are amenable to change (more adaptable), they can still be maladaptive if they are unable to instantiate a robust enough behavioral policy to ensure success in a given environmental context. Crucially, the ideal situation involves learning models which are both adapted to the current environmental context and malleable enough to make possible future adaptation when the environmental context changes. This highlights the stability-plasticity dilemma (Mermillod, Bugaiska, & Bonin, 2013), which refers to the extent a system must be plastic to integrate novel information and stable in order not to catastrophically interfere with consolidated knowledge.

1.4 Identifying a key role for prefrontal cortex

In this work we focus in particular on the role of the prefrontal cortex (PFC) as a target site of both endogenous serotonergic neuromodulation as well as exogenous neuromodulation by psychedelics. The PFC is involved in higher cognitive functions such as decision-making, working memory, and executive control (Friedman & Robbins, 2022), and is a strong candidate region to focus our attention for a number of reasons. It both contains a large number of neurons which express both 5-HT2a and 5-HT1a postsynaptic receptors and is a primary target of 5-HT neuromodulation arising from the dorsal raphae nucleus (Puig & Gulledge, 2011). These facts alone would make it likely that the PFC plays a significant role in the action of psychedelic drugs, and there is clear neuroimaging evidence that this is the case (Wood, Kim, & Moghaddam, 2012; Luppi et al., 2024).

The PFC is also implicated in a variety of psychopathology ranging from depression to schizophrenia (Callicott et al., 2000; Pizzagalli & Roberts, 2022), making it an ideal target region for a pharmacological intervention which may be addressing an underlying causal factor of psychopathology (R. Carhart-Harris et al., 2022). It is also the case that the PFC represents high-level beliefs and goals of the organism concerning self, others, and world (Miller, Freedman, & Wallis, 2002), as well as the high-level behavioral and attentional policies required to achieve those goals (Domenech & Koechlin, 2015). These collectively make up the primary psychological material which is involved in psychedelic-assisted psychotherapy, as well as psychotherapy more broadly. Indeed, within the REBUS model the PFC is positioned as one of the primary regions both responsible for the acute effects of psychedelics as well as their more lasting therapeutic effects, given its position at the top of the predictive processing hierarchy.

Given its role both in the maintenance of critical representations of self, world, and goals, as well as its causal influence on downstream cortical dynamics and ultimately behavior, we consider the prefrontal cortex to be a prime candidate to model within our theoretical framework and simulations. The PFC also expresses strong recurrent dynamics (S. Wang, Falcone, Richmond, & Averbeck, 2023), which are necessary for the formation and maintenance of stable attractors responsible for the representation of high-level beliefs (Mante, Sussillo, Shenoy, & Newsome, 2013). In our context, these stable attractors corresponding to minima in the optimization landscape of the PFC represent high-level beliefs concerning the organism's overall fitness in a given environmental context, its high level goals, as well as mental and physical behavioral policies. It is precisely this set of beliefs which psychedelic-assisted psychotherapy has been shown to be capable of altering (Letheby, 2021; Zeifman et al., 2022).

1.5 Energy-based models of neural dynamics

Dynamical systems approaches have become increasing popular in computation neuroscience, being used to model both single neurons as well as population responses of large numbers of neurons (Favela, 2020). The dynamical systems approach explicitly describes neural population responses as time-varying trajectories in a high-dimensional state space and views the dynamics as acting to shape these trajectories (Khona & Fiete, 2022).

The brain operates in the face of substantial uncertainty due to ambiguity in the inputs, and inherent unpredictability in the environment. Probabilistic inference offers a principled framework for understanding both behaviour and cortical computation. Based on the idea that cortical circuits implement Bayesian inference in latent variable models (Knill & Richards, 1996; R. P. Rao, Olshausen, & Lewicki, 2002), it has been proposed that neural firing rates might be viewed as representing Monte Carlo (MC) (Metropolis, 1987) samples from the posterior distribution over the latent variables capturing the neural response, given the observed input or stimuli (Hoyer & Hyvärinen, 2002; Orbán, Berkes, Fiser, & Lengyel, 2016). In this view, the neural response variability is related to the uncertainty, about world parameters, which is inherent in any stimulus.

The PFC forms a recurrent network whose spiking activity encodes multiple types of learning-relevant information (Mante et al., 2013; J. X. Wang et al., 2018). A number of plausible mechanisms have been proposed for the persistence of neural activity in a population of neurons for the time a stimulus is perceived, and there is significant evidence of this being the case for many clinically relevant sub-networks of the cortex, and for the PFC in particular (Li, Daie, Svoboda, & Druckmann, 2016; Khona & Fiete, 2022; J. X. Wang et al., 2018; S. Wang et al., 2023). These stable and persistent patterns of neural activity are thought to be the result of attractor dynamics within the cortex. This assumption of the existence of cortical attractor dynamics has formed the basis for multiple recent theoretical models of psychedelic's effect on brain activity (Girn et al., 2023; Hipólito, Mago, Rosas, & Carhart-Harris, 2023; Ruffini, Lopez-Sola, Vohryzek, & Sanchez-Todo, 2024).

Energy functions can represent various aspects of neural activity, including factors such as synaptic weights, membrane potentials, and firing rates. One common approach is to define an energy function analogous to the Hamiltonian in classical mechanics (Hamilton, 2000). This energy function captures the total "energy" of the neural network. Changes in neural activity and synaptic connectivity can be modeled as gradients in this energy landscape, driving the system towards stable states (attractors) or trajectories corresponding to specific cognitive processes or behaviors.

Here we study the abstract properties of PFC neural circuits by examining the dynamics of neural activity, their evolution over time, and whether they tend to converge to stable attractor states (Keller & Mrsic-Flogel, 2018). Energy functions can be used to analyze the stability of these states by assessing how changes in neural activity affect the total energy of the system. In this perspective, attractors are minima in the energy landscape in which a stable and persistent configuration of neural activity in reached. From a Bayesian perspective, this persistent activity corresponds to a internal belief state, and the difficulty of the system to escape this configuration corresponds to the precision with which the belief is encoded.

In complex systems, the energy landscape represents all possible configurations and their associated energies, with the system evolving towards lower energy configurations. Understanding the energy landscape's structure helps predict the behavior and stability of dynamical systems. In optimal control problems, manipulating inputs or synaptic connections aims to minimize an optimization objective, which determines the energy function and the associated probabilistic model it generates (Hamilton, 2000). Within the hierarchical predictive processing (HPP) framework, this objective functional calibrates the energy function and

shapes the landscape governing cortical dynamics based on the predictive error between decoded and incoming stimuli (Keller & Mrsic-Flogel, 2018).

Following these premises, we propose an optimization algorithm within the predictive processing framework which utilizes an energy-based model (EBM) of neural dynamics (Ackley, Hinton, & Sejnowski, 1985; Hinton, 1999; LeCun, Chopra, Hadsell, Ranzato, & Huang, 2006), a.k.a. energy networks (Blondel, Llinares-López, Dadashi, Hussenot, & Geist, 2022). We consequently assume that the attractor dynamics operate within the HPP framework (cf. Figure 2), whereby higher or downstream functional networks attempt to predict incoming information from lower or upstream networks via top-down predictions cf. Figure 2 (A)). In this context, the formation of attractors represents the persistence of stable predictions, or patterns of activity in upstream or high-level functional networks within HPP. These patterns of stable activity then correspond to decodable neural responses which are equivalent to internal representations of beliefs from a cognitive perspective.

2 Methods

Our theoretical model is based on existing work using EBMs to model neural responses, but differs in a number of important ways. (i) We first describe a general algorithm for learning and inference in EBMs of neural responses to stimuli in the context of HPP; (ii) Then we present our algorithm with added serotonin modulation; (iii) Lastly, we introduce the metrics of interest considered in our simulation results presented in Section 3. For additional details on the theoretical model and relevant connections to similar models, see Appendix B. For additional details on the specific parameters used in the simulations, see Appendix C.

2.1 Energy-based model preliminaries

Taking HPP as a starting point, we assume the brain is modeling the distribution of possible interoceptive and exteroceptive sensory signals which can be encountered in a given environmental context. Concretely, at any moment in time an organism is receiving some sensory stimuli $\mathbf{x}$ which is distributed according to some probability distribution $p(\mathbf{x})$. We model a single layer of a predictive processing hierarchy (Figure 2(a)), and thus our framework is generic enough to remain agnostic with respect to the signal's origin. In this work however we assume $\mathbf{x}$ is internal to the organism and corresponds to a representation upstream of the PFC. The goal of the model is to produce a downstream neural response $\mathbf{z}$ which is capable of predicting $\mathbf{x}$. To do so, the model must learn to correctly represent the distribution $p(\mathbf{x})$.

The distribution of stimuli is a function of some environmental context (Mante et al., 2013; C. Chen et al., 2024), which may persist for minutes, hours, days or weeks (Muysers et al., 2024). This context may correspond to a particular configuration of social relationships, the physical environment, or goals which govern the individual's mental and physical behavior. In psychedelic-assisted therapy, an environmental context would correspond to the "set and setting" for a particular drug-induced experience. In order to simplify the learning procedure, we consider a fixed environmental context, and correspondingly a stationary

distribution over stimuli.

In order to learn a response function which can model the stimuli distribution $p(\mathbf{x})$, we take inspiration from control theory and machine learning (LeCun et al., 2006; Lotter, Kreiman, & Cox, 2016; Ranzato, Poultney, Chopra, & Cun, 2006; Ngiam, Chen, Koh, & Ng, 2011; Bond-Taylor, Leach, Long, & Willcocks, 2021), and employ EBMs, a.k.a. energy networks. EBMs are models which perform inference and learning implicitly through an energy function, typically parameterized by a multi-layer or recurrent neural network. We follow prior approaches from the literature (Hoyer & Hyvärinen, 2002; Hennequin, Aitchison, & Lengyel, 2014; Aitchison & Lengyel, 2016; Orbán et al., 2016; Echeveste, Aitchison, Hennequin, & Lengyel, 2020; Qi & Gong, 2022; Dong & Wu, 2023) and consider inference and learning in EBMs where the energy function is defined over possible neural responses, or patterns of neural activity.

To begin with, any probability density function $p(\mathbf{x})$ can be expressed via a Boltzmann distribution as

$$p(\mathbf{x}) \doteq \frac{\exp(-E(\mathbf{x}))}{Z}, \quad Z \doteq \int_{\mathbf{x}} \exp(-E(\mathbf{x})) \tag{1}$$

where $E(\mathbf{x})$ is an energy function that maps input $\mathbf{x} \in \mathcal{X}$ to a scalar, and Z is the normalizing constant with respect to $\mathbf{x}$ (also known as the partition function). Ideally, an energy function should assign low energy values to the samples drawn from data distribution, and high values otherwise.

In computational models of learning in the brain using EBMs, inference is carried out through neural sampling (Ackley et al., 1985; Hoyer & Hyvärinen, 2002), while learning happens via long-term synaptic plasticity (Bliss & Lomo, 1973; Miles, Poncer, Fricker, & Leinekugel, 2005). The key challenges are the following: (i) The intractable partition function Z makes learning the energy function via a maximum likelihood estimate is not always straightforward, particularly under normative constraints, such as locality. (ii) The inference procedure via neural sampling should be able to explore complex multi-modal distributions and be fast enough to not violate normative neural constraints.

2.2 Predictive processing using EBMs

The structure of the EBM is such that synaptic connectivity, captured with parameters θ, can be functionally represented by means of a mapping between stimuli $\mathbf{x}$ and neural responses $\mathbf{z}$, which defines a joint probabilistic model $p_\theta(\mathbf{x}, \mathbf{z})$ (Figure 2 (D)). By manipulating the parameters or synaptic connections θ, we may influence the dynamics of neural activity and guide the system towards desired states or behaviors.

EBMs can provide a framework for inference and learning using spatially localized rules (Dong & Wu, 2023), which fit normative theories that postulate plasticity can only be affected by locally-available information (Crick, 1989; Grossberg, 1987; Lillicrap, Santoro, Marris, Akerman, & Hinton, 2020; Bredenberg & Savin, 2023), and thus refrain from violating neural processing constraints.

Following prior research emphasizing the role of gradients in neural computation (Richards et al., 2019a), we use a gradient-based learning algorithm, within the predictive processing framework, which corresponds to learning the parameters θ defining the synaptic connections such that the distribution of the stimuli under the model $p_\theta(\mathbf{x})$ matches as accurately as possible the perceived distribution of stimuli $p(\mathbf{x})$

an organism may encounter in a given environmental context. This is carried out by maximum likelihood estimation with respect to θ

$$\min_{\theta} D_{\mathrm{KL}}(p(\mathbf{x}), p_\theta(\mathbf{x})) = \min_{\theta} \mathbb{E}_{\mathbf{x}\sim p(\mathbf{x})}\left[\log \frac{p(\mathbf{x})}{p_\theta(\mathbf{x})}\right] = \min_{\theta} -\mathbb{E}_{\mathbf{x}\sim p(\mathbf{x})}\left[\log p_\theta(\mathbf{x})\right]$$

with $D_{\mathrm{KL}}(p, q)$ the Kullback–Leibler (KL) divergence between p and q. The marginal likelihood $p_\theta(\mathbf{x}) = \int_z p_\theta(\mathbf{z}) p_\theta(\mathbf{x}|\mathbf{z}) d\mathbf{z}$ however is computationally intractable in most cases of interest. One approach to solve this challenge is to approximate the posterior distribution $p_\theta(\mathbf{z}|\mathbf{x})$, with another parameterized density model $p_\phi(\mathbf{z}|\mathbf{x})$, using synaptic connection weights ϕ. Then, minimize the evidence lower bound (ELBO) (Beal, 2003; Kingma & Welling, 2022) on the marginal likelihood $-\mathbb{E}_{\mathbf{x}\sim p(\mathbf{x})}\left[\log p_\theta(\mathbf{x})\right] \leq \mathcal{L}(\theta, \phi)$, with respect to both the generative parameters θ and the variational parameters ϕ, where the ELBO is

$$\begin{aligned}\mathcal{L}(\theta, \phi) &\doteq D_{\mathrm{KL}}(p_\phi(\mathbf{x}, \mathbf{z}), p_\theta(\mathbf{x}, \mathbf{z})) = \int \frac{\log p_\phi(\mathbf{z}|\mathbf{x})}{\log p_\theta(\mathbf{x}, \mathbf{z})} p_\phi(\mathbf{x}, \mathbf{z}) d\mathbf{x} d\mathbf{z} \\ &= \mathbb{E}_{\mathbf{x}\sim p(\mathbf{x})}\left[D_{\mathrm{KL}}(p_\phi(\mathbf{z}|\mathbf{x}), p_\theta(\mathbf{z})) - \mathbb{E}_{p_\phi(\mathbf{z}|\mathbf{x})}\left[p_\theta(\mathbf{x}|\mathbf{z})\right]\right]\end{aligned}$$

To approximate $\nabla_\theta \mathcal{L}(\theta, \phi) = -\mathbb{E}_{\mathbf{x}\sim p(\mathbf{x})}\mathbb{E}_{\mathbf{z}\sim p_\phi(\mathbf{z}|\mathbf{x})}[\nabla_\theta \log p_\theta(\mathbf{x}, \mathbf{z})]$, the loss gradient with respect to the generative parameters θ, given individual stimuli sampled from $\mathbf{x}^{(i)} \sim p(\mathbf{x}), \forall i \geq 0$ and generated samples from the posterior $\mathbf{z}^{(i)} \sim p_\phi(\mathbf{z}|\mathbf{x})$, we may use an online estimate $\widehat{\nabla}_\theta \mathcal{L}(\theta, \phi) \doteq \nabla_\theta \log p_\theta(\mathbf{x}^{(i)}, \mathbf{z}^{(i)})$, or an empirical average over a temporal window $\widehat{\nabla}_\theta \mathcal{L}(\theta, \phi) \doteq \sum_{\mathbf{x}^{(i)}} \sum_{\mathbf{z}^{(i)}} \nabla_\theta \log p_\theta(\mathbf{x}^{(i)}, \mathbf{z}^{(i)})$ and then update the parameters θ with gradient descent $\theta' \leftarrow \theta + \Delta\theta$, where $\Delta\theta \doteq -\alpha \widehat{\nabla}_\theta \mathcal{L}(\theta, \phi)$, and α is a small positive learning rate.

Similarly, to estimate the gradient with respect to the parameters of the recognition model ϕ, given individual stimuli sampled from $\mathbf{x}^{(i)} \sim p(\mathbf{x}), \forall i \geq 0$ and generated samples from the posterior $\mathbf{z}^{(i)} \sim p_\phi(\mathbf{z}|\mathbf{x})$, we may use an online estimate $\widehat{\nabla}_\phi \mathcal{L}(\theta, \phi) \doteq \nabla_\phi \log p_\phi(\mathbf{z}^{(i)}|\mathbf{x}^{(i)})$, or an empirical average over a temporal window $\widehat{\nabla}_\phi \mathcal{L}(\theta, \phi) \doteq \sum_{\mathbf{x}^{(i)}} \sum_{\mathbf{z}^{(i)}} \mathbf{r}(\mathbf{x}^{(i)}, \mathbf{z}^{(i)}) \nabla_\phi \log p_\phi(\mathbf{z}^{(i)}|\mathbf{x}^{(i)})$, where $\mathbf{r}(\mathbf{x}, \mathbf{z}) = \log \frac{p_\theta(\mathbf{x}, \mathbf{z})}{p_\phi(\mathbf{z}|\mathbf{x})}$ can be viewed as a reward in the context of the "weight transport" or credit assignment problem (Williams, 1992; Bellec et al., 2020; Richards et al., 2019b). We may then update the variational parameters ϕ with gradient descent $\phi' \leftarrow \phi + \Delta\phi$, where $\Delta\phi \approx -\alpha \widehat{\nabla}_\phi \mathcal{L}(\theta, \phi)$.

In EBMs we may represent the generative and variational parameters jointly $p_\theta(\mathbf{x}, \mathbf{z})$, without explicit modelling of the posterior $p_\theta(\mathbf{z}|\mathbf{x})$ if we can still draw samples from it (Dong & Wu, 2023). Samples from the posterior or the joint model may be obtained using stochastic sampling by means of a Monte Carlo gradient estimator.

2.3 Sampling from EBMs

The neural sampling perspective of probabilistic inference in the cortex posits that the brain infers a posterior distribution over neural responses consistent with Bayesian inference (Knill & Richards, 1996; Ernst & Banks, 2002; Körding & Wolpert, 2004; Knill & Pouget, 2004; Tenenbaum, Griffiths, & Kemp, 2006; Doya et al., 2006). To model the variability in neural responses, it has been proposed (Hoyer & Hyvärinen, 2002) to interpret the neural responses as Monte Carlo sampling (Metropolis, 1987) of the posterior distribution

$p_\phi(\mathbf{z}|\mathbf{x})$ or the joint model $p_\theta(\mathbf{x}, \mathbf{z})$. To work around the need to represent the posterior explicitly, we may also use the relationship $\nabla_\mathbf{z} \log p_\theta(\mathbf{z}|\mathbf{x}) = \nabla_\mathbf{z} \log p_\theta(\mathbf{x}, \mathbf{z}) = -\nabla_\mathbf{z} E_\theta(\mathbf{x}, \mathbf{z}) - \nabla_\mathbf{z} Z_\theta = -\nabla_\mathbf{z} E_\theta(\mathbf{x}, \mathbf{z})$, since $\nabla_\mathbf{z} Z_\theta = 0$. The sampling process itself involves iteratively descending the energy landscape in order to arrive at a representation where the energy is locally minimal.

Prior works (Hoyer & Hyvärinen, 2002; Welling & Teh, 2011; Buesing et al., 2011; T. Chen, Fox, & Guestrin, 2014; Hennequin et al., 2014; Aitchison & Lengyel, 2016; Orbán et al., 2016; Echeveste et al., 2020; Qi & Gong, 2022) have developed and applied a number of approaches to sample from the posterior or the joint efficiently by defining the inference dynamics as performing walks in the latent space of neural responses, e.g., Gibbs sampling (Ackley et al., 1985), Langevin sampling (Neal, 2012; Welling & Teh, 2011), Hamiltonian sampling (Neal, 2012; T. Chen et al., 2014).

2.4 Operationalizing serotonergic neuromodulation

Within the framework we lay forward, pathologies in optimization correspond to failures to find stable and predictive configurations of neural activity, and likewise a failure to encode adaptive beliefs or internal representations. If these pathologies are severe and prolonged enough, they may lead to downstream failures in behavioral fitness and the development of psychopathology associated with either over- or under-canalization (R. Carhart-Harris et al., 2022; Juliani et al., 2024; Ruffini et al., 2024). These optimization issues can include the inability to find global optima (J. Chen & Luss, 2018; Ajalloeian & Stich, 2020), ill-conditioning of the energy landscape manifesting in high sensitivity in response to small changes in the input, causing a deceleration in convergence to a stable attractor, overfitting and plasticity loss (Ahmed, Roux, Norouzi, & Schuurmans, 2019; Mei et al., 2020; Bolland, Louppe, & Ernst, 2023; Lyle et al., 2023).

Using Markov Chain Monte Carlo (MCMC) (Metropolis, 1987) we can sample from the posterior distribution over neural responses $p_\theta(\mathbf{z}|\mathbf{x})$ by translating the density function for this distribution to a "potential" energy function $E_\theta(\mathbf{x}, \mathbf{z})$ using Eq 1. We then simulate a Markov chain from a stochastic process designed to explore the posterior distribution of neural responses, given observed data. Langevin dynamics (LD) has been applied to sample from a posterior distribution, as Bayesian inference, by performing stochastic gradient descent on the manifold of the energy function defined as the negative log-posterior distribution $z_{t+1} \leftarrow z_t - \beta \nabla_\mathbf{z} E_\theta(\mathbf{x}, \mathbf{z}) + \alpha\xi$, for $E_\theta(\mathbf{x}, \mathbf{z}) = -\log p_\theta(\mathbf{z}|\mathbf{x})$.

Alone however, this response sampling procedure is slower than the processing speed of the cortex (Dong & Wu, 2023), and is furthermore susceptible to the optimization pathologies described above. Introducing a surrogate objective via an auxiliary function (a "kinetic" energy function) that presents additional neuromodulation mechanics, the resulting stochastic dynamical system may be better able to explore the posterior density and escape from local minima of the "potential" energy landscape. We model the neuromodulatory effects of serotonin using an inference process presenting such neuromodulatory mechanisms which indirectly impact the learning process for the synaptic connectivity of the energy function over neural responses inducing the posterior distribution. We operationalize the effects of 5-HT2a and 5-HT1a agonism as modulators of the energy function landscape, and simulate their effect following existing empirical

evidence, particularly implementing the following mechanisms.

Excitatory postsynaptic effect of 5-HT2a agonism Given the excitatory postsynaptic effect of 5-HT2a agonism on cortical pyramidal cells, we model its role in the neural optimization process as a stochastic perturbation of the energy landscape used to infer the neuronal response. This stochasticity injection has been hypothesized to disrupt neural responses through the desynchronization of relevant neural populations in a manner that is capable of both strengthening or relaxing the neural representations of populations of neurons in a context-dependent manner (Safron, 2020).

Inhibitory postsynaptic effect of 5-HT1a agonism We model the inhibitory postsynaptic effect of 5-HT1a agonism on cortical pyramidal cells via a local smoothing or regularization of the energy landscape. This serves to destabilize locally optimal, but potentially globally sub-optimal, attractor points of the optimization surface. By regularizing the energy function, the optimization algorithm is encouraged to explore the solution space effectively while still exploiting promising regions for optimization. This relaxation prevents the encoding of overly-precise beliefs, which may become maladaptive in future contexts. It also helps prevent overfitting by promoting a broader and less confident exploration of the solution space—an exploration-exploitation tradeoff (Kukavcka, Golkov, & Cremers, 2017).

2.5 Connection to Hamiltonian dynamics

A connection can be drawn between our proposed mechanisms of 5-HT neuromodulation and Hamiltonian Monte-Carlo (HMC), which relies on Hamiltonian dynamics to sample from posterior probabilities (Neal, 2012). HMC can be implemented by the interactions of recurrently coupled excitatory and inhibitory populations in cortical circuits, and has been previously used to simulate neural responses (T. Chen et al., 2014; Aitchison & Lengyel, 2016; Dong & Wu, 2023). In Hamiltonian dynamics, the state of the system behaves as a particle moving on a (high dimensional) surface, frictionless but with momentum. The algorithm leverages an auxiliary variable, which can be interpreted as momentum in the context of neural adaptation, and using which it is able to adapt the inference dynamics. The surface determines the potential energy of the particle, corresponding to the negative logarithm of the probability distribution that needs to be sampled from $E(\mathbf{x}, \mathbf{z}) \propto -\log p(\mathbf{z}|\mathbf{x})$ (such that high probability states correspond to low potential energy).

The momentum variable, denoted $\mathbf{e}$ corresponds to an auxiliary random variable capturing a "kinetic" energy function $\mathbf{e}$, defined using a normal distribution centered around the neural responses $\mathbf{z}$, corresponding to smoothing the probability density of $\mathbf{z}$ (5-HT1a effect), with standard deviation σ, $p(\mathbf{e}|\mathbf{x}, \mathbf{z}) = \mathcal{N}(\mathbf{z}, \sigma)$. The momentum term accumulates smoothed stochastic gradients of the potential energy function, which are subject to a Gaussian convolution and are applied a stochastic perturbation using independent Gaussian-distributed random variables ξ_t representing the stochastic noise terms (5-HT2a effect). As a consequence, the particle will accelerate as it heads towards the minimum of the "potential" energy land-

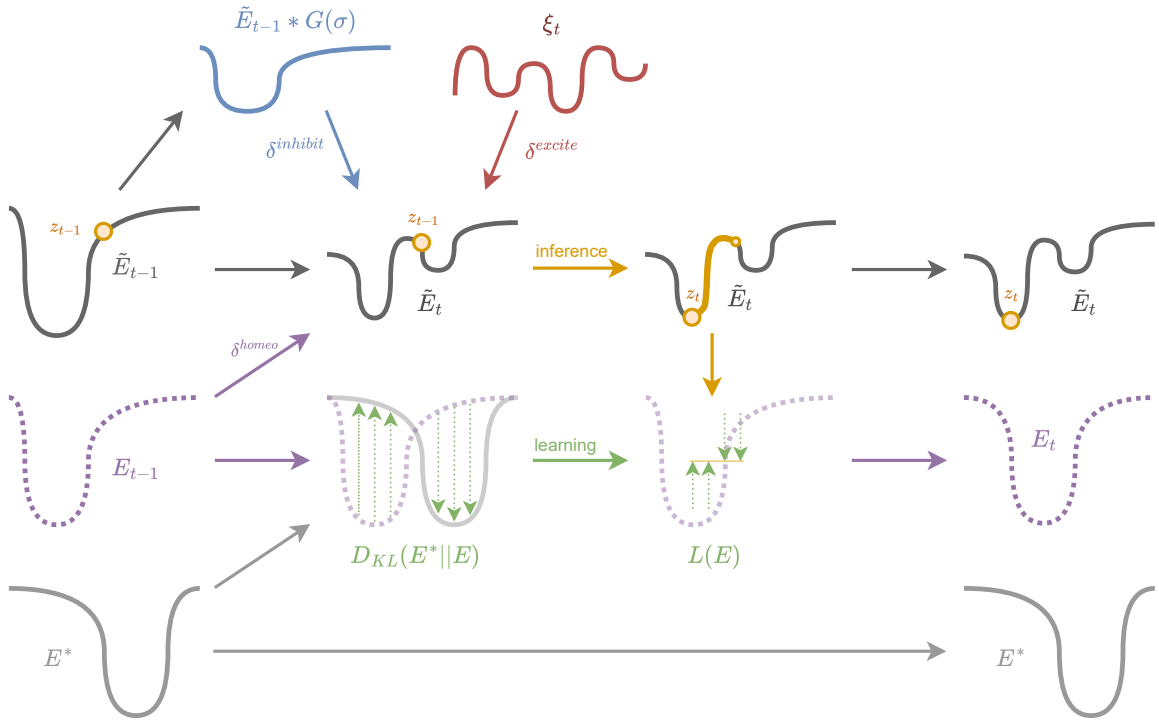

Figure 3: Simplified diagram of the inference and learning procedure used to update the internal representation $\mathbf{z}$, the energy function $E_\theta(\mathbf{z})$, and the neuromodulated energy function $\tilde{E}_\theta(\mathbf{z})$ at each time-step. A one-dimensional energy function is used here for illustrative purposes. See Algorithm 1 for a step-by-step account of the process.

scape

$$\mathbf{e}_{t+1} \leftarrow \mathbf{e}_t + \Delta\mathbf{e}_t, \qquad \Delta\mathbf{e}_t \doteq -\beta\nabla_{\mathbf{z}_t} E_\theta(\mathbf{z}_t, \mathbf{x}_t, \mathbf{e}_t) + \alpha\xi_t, \qquad p(\mathbf{e}_t|\mathbf{z}_t, \mathbf{x}_t) \doteq \mathcal{N}(\mathbf{z}_t, \sigma_t)$$
$$\mathbf{z}_{t+1} \leftarrow \mathbf{z}_t + \Delta\mathbf{z}_t, \qquad \Delta\mathbf{z}_t \doteq \mathbf{e}_{t+1}$$

The above equation indicates the instantaneous change in neural response is driven by a momentum term $\mathbf{e}$ which accumulates stochastic smoothed gradients of an energy function $\nabla_{\mathbf{z}} E_\theta(\mathbf{z}, \mathbf{x})$ subject to injected stochasticity and smoothness. Thus, the neural adaptation and sampling process explores the space of neural responses by combining gradient-driven dynamics towards lower energy regions with stochastic noise and smoothing, allowing the system to cover a wider range of neural responses. After a sufficient number of time steps, the final neural response of the system can be considered as a sample from the distribution $p_\theta(\mathbf{x}, \mathbf{z})$ of interest.

2.6 Concrete algorithm for 5-HT neuromodulated EBMs

Above we presented idealized learning and inference algorithms in EBMs with neuromodulation mechanisms within an HPP framework. In this work we are primarily interested in studying the effect of exogenous 5-HT2a and 5-HT1a neuromodulation on the inference procedure of neural responses in EBMs. In order to keep our analysis clear and tractable, we make a set of simplifications to the above procedures.

Algorithm 1 Neuromodulated optimization algorithm

Input: Homeostatic plasticity strength α^{homeo}, Hebbian plasticity step-sizes α_t^{plast}
Input: Neuromodulation level sampling function $\texttt{drug}(t)$, 5-HT1a gaussian kernel size σ_t
Input: target energy function $E^*(\mathbf{z})$, initial energy function $E_0(\mathbf{z})$, prior over neural responses $p_0(\mathbf{z})$
$\tilde{E}_0(\mathbf{z}) = E_0(\mathbf{z})$ ▷ Initialize surrogate neuromodulated energy function
$\mathbf{z}_0 \sim p_0(\mathbf{z})$ ▷ Sample initial neural response
for $t = 0, 1 \ldots T-1$ and $\forall \mathbf{z}$ **do**
  $\alpha^{\text{excite}}, \alpha^{\text{inhibit}} \sim \texttt{drug}(t)$ ▷ Sample neuromodulator level
  $\delta^{\text{excite}} = \alpha^{\text{excite}} \xi_t$ ▷ 5-HT2a agonism effect
  $\delta^{\text{inhibit}} = \alpha^{\text{inhibit}} [\tilde{E}_t(\mathbf{z}) * G(\sigma_t) - \tilde{E}_t(\mathbf{z})]$ ▷ 5-HT1a agonism effect
  $\delta^{\text{homeo}} = \alpha^{\text{homeo}} [E_t(\mathbf{z}) - \tilde{E}_t(\mathbf{z})]$ ▷ Homeostatic plasticity effect
  $\tilde{E}_{t+1}(\mathbf{z}) \leftarrow \tilde{E}_t(\mathbf{z}) + \delta^{\text{inhibit}} + \delta^{\text{excite}} + \delta^{\text{homeo}}$ ▷ Apply neuromodulatory effect
  $\mathbf{z}_t^{(0)} \doteq \mathbf{z}_t$
  for $k = 0, 1 \ldots K-1$ **do**
    $\xi^{(k)} \sim \mathcal{N}(0, 1)$ ▷ Sample optimization noise
    $\mathbf{z}_t^{(k+1)} \leftarrow \mathbf{z}_t^{(k)} - \nabla_{\mathbf{z}} \tilde{E}_t(\mathbf{z}_t^{(k)}) + \xi^{(k)}$ ▷ Recurrently infer neural responses
  end for
  $\mathbf{z}_t \doteq \mathbf{z}_t^{(K-1)}$
  $\mathcal{L}(E_t) \doteq 1/2 \|E^*(\mathbf{z}_t) - E_t(\mathbf{z}_t)\|_2^2$ ▷ Objective for approximate optimization of EBM
  $E_{t+1}(\mathbf{z}_t) \leftarrow E_t(\mathbf{z}_t) + \alpha^{plast} [E^*(\mathbf{z}_t) - E_t(\mathbf{z}_t)]$ ▷ Update EBM with gradient-based update on $\mathcal{L}(E_t)$
end for

First, we omit directly representing $\mathbf{x}$, and consider only $\mathbf{z}$ in the learning procedure. Second, we use a fixed target $E^*(\mathbf{z})$ to provide an estimate of the fitness of a given neural response while avoiding the need to explicitly model the generation of stimuli. As such, $p^*(\mathbf{z})$ corresponds to a distribution over neural responses which is able to optimally predict $p(\mathbf{x})$. Third, $E_\theta(\mathbf{z})$ and $E^*(\mathbf{z})$ are both modeled non-parametrically, as a tabular representation, with $N \times N$ set of points, each with an associated value that can be changed during learning. Thus, there are no learned parameters θ (or ϕ), rather we update $E(\mathbf{z}) \doteq E_\theta(\mathbf{z})$ directly, henceforth dropping θ from the notation. A consequence of the non-parametric representation of $E(\mathbf{z})$ is that all gradients in the simulator are approximated using finite differences.

The non-parametric learned and target energy functions $E(\mathbf{z})$ and $E^*(\mathbf{z})$ are sampled from a distribution of possible energy function landscapes $E \sim \mathcal{E}(N, \sigma)$. Specifically, $\mathcal{E}(N, \sigma) = M * G(\sigma)$, where $M \in \mathbb{R}^{N \times N}$ with $M_{ij} \sim \mathcal{N}(0, 1), \forall i, j \in \{1, \ldots, N\}$, N is the size of the non-parametric representation, $G(\sigma)$ is a Gaussian function which is convolved with M, and σ is the size of the Gaussian function kernel. The values of these representations are then normalized to be within the range $[0, 1]$ at the start of the optimization process.

Finally, we employ a simplified learning objective for $E(\mathbf{z}) \propto -\log p(\mathbf{z})$ to describe changes in synaptic connectivity, namely via $\min_E \mathcal{L}(E)$, with $\mathcal{L}(E) = D_{\text{KL}}(p, p^*) = \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} \left[\log \frac{p(\mathbf{z})}{p^*(\mathbf{z})} \right] \approx \sum_{\mathbf{z}^{(i)}} [E^*(\mathbf{z}^{(i)}) -$

$E(\mathbf{z}^{(i)})]$, capturing the difference in energy function values at sampled points. Finally, we introduce use an auxiliary energy function $\tilde{E}(\mathbf{z})$, a surrogate corresponding to the energy function under neuromodulation, used for sampling neural responses $\mathbf{z}$.

Neuromodulatory signals We consider two unique neuromodulatory signals and one additional constraint:

- α^{excite}: modulates the 5-HT2a excitatory drug effect
- α^{inhibit}: modulates the 5-HT1a inhibitory drug effect
- α^{homeo}: modulates the strength of homeostatic constraints

Sampling neuromodulation level The ranges for α^{excite} and α^{inhibit} can either be set constant for the duration of the optimization process, or be determined dynamically. Given that we are interested in exogenous neuromodulation effects from drugs, we use a function $\texttt{drug}(t)$ to determine the neuromodulation levels as a function of time. This function uses a beta distribution to provide a response curve, where the probability density function at a given point corresponds to the drug strength at that point in time. This function is used to provide a simple and generic representation of prototypical change in blood plasma concentration levels of a drug over time.

5-HT2a drug effect At each iteration t, a stochastic perturbation is applied to the auxiliary energy function $\tilde{E}(\mathbf{z})$

$$\delta_t^{\text{excite}} = \alpha_t^{\text{excite}}[\xi_t^{\text{excite}}]$$

where $\xi_t^{\text{excite}} \sim \mathcal{E}(N, \sigma)$ is structured noise sampled from the energy function distribution and then scaled to be within the range $[-0.5, 0.5]$, and α_t^{excite} the relative excitatory strength of 5-HT2a signaling.

5-HT1a drug effect We simulate the inhibitory effect of 5-HT1a agonism by applying a Gaussian filter kernel centered around $\tilde{E}(\mathbf{z})$ and apply this effect to the auxiliary energy function $\tilde{E}(\mathbf{z})$

$$\delta_t^{\text{inhibit}} = \alpha_t^{\text{inhibit}}[\tilde{E}(\mathbf{z}) * G(\sigma_t) - \tilde{E}(\mathbf{z})]$$

We use a local Gaussian filter kernel $G(\sigma)$ which is convolved with $\tilde{E}(\mathbf{z})$, where σ_t denotes the size of the Gaussian kernel and $\alpha_t^{\text{inhibit}}$ the relative inhibition strength of 5-HT1a signaling.

Homeostatic plasticity effect Homeostatic plasticity is a mechanism for stabilizing neural activity levels across a network, ensuring that neurons maintain an appropriate level of excitability. This regulatory mechanism restores stability and mitigates imbalances in the activity levels of neurons by regularizing the neuromodulated energy landscape towards the baseline

$$\delta_t^{\text{homeo}} = \alpha_t^{\text{homeo}}[E(\mathbf{z}) - \tilde{E}(\mathbf{z})]$$

with α_t^{homeo} the relative strength of the update toward the previous value of the unmodulated energy function $E(\mathbf{z})$.

Applying neuromodulatory effects The three neuromodulatory effects are then additively combined into a single term δ_t and accumulated into $\tilde{E}(\mathbf{z})$ before the sampling process takes place each iteration. The optimization procedure including the aforementioned neuromodulation mechanisms is described in Algorithm 1, and visualized in Figure 3.

2.7 Metrics of interest

We consider six primary metrics of interest which we collect and analyze during each simulation experiment. Figure 4 (1st row) contains a simple graphical representation of each metric under consideration, with the final two being represented by the rightmost graphic.

Energy function value The energy function value corresponds to the $\mathbb{E}_{\mathbf{z}_t}[\tilde{E}(\mathbf{z}_t)]$ for a set of sampled $\mathbf{z}_t$ values at a given iteration t of the simulation. This metric provides insight into how a given intervention on the neuromodulatory system changes the energy level of sampled $\mathbf{z}_t$ values relative to a non-modulated baseline over time. The EBM represents the probability of different responses implicitly, by the frequency with which it visits their representations via its dynamics. Higher energy of some neural response $\mathbf{z}$ means neural activity will likely visit $\mathbf{z}$ with less frequency, whereas lower energy corresponds to a larger proportion of the time neural activity representing $\mathbf{z}$.

Gradient magnitude We consider the average gradient magnitude of the energy function at a given time-step. This is calculated by uniformly evaluating $||\nabla\tilde{E}(\mathbf{z})||_2$ for every possible value of $\mathbf{z}$ in our non-parametric representation. Functions with bounded gradient magnitudes are smooth, ensuring that changes in input variables result in proportionally limited changes in the function's output. This is crucial for maintaining stability and convergence of optimization algorithms towards optimal solutions. The average gradient magnitude value also corresponds to the expected precision assigned to the set of internal representations capable of being encoded by $\mathbf{z}$. Larger values are desirable to the extent that the representations encoded by $\mathbf{z}$ are both able to reduce prediction errors in the current and future environmental contexts. This is more likely to be the case the more stationary the environmental context is over time. In non-stationary environments, a smaller average gradient magnitude over time will correlate with smoothness, and may be desirable as it will produce a more diverse sampling of $\mathbf{z}$.

Local minima count We consider the number of local minima present in the energy function at a given time-step. This is calculated by uniformly evaluating our non-parametric representation of $\tilde{E}(\mathbf{z})$ and determining the local gradient at each point. The local minima count reflects the complexity of the energy landscape (the more local minima, the more complex the landscape) and corresponds to possible stable attractor states in which the stochastic sampling process may enter. Within the context of our theoretical

model, the number of local minima is predictive of how successful both the learning and inference processes will be. Fewer local minima increase the likelihood that the stochastic sampling process will infer a neural response closer to global minima of the energy function.

State visitation count We calculate state visitation count by tracking the number of unique values which $\mathbf{z}$ takes over the course of the optimization process in a given experiment. Within the context of our theoretical model, this metric illustrates the ability to cover a broad range of plausible interpretations of its input stimuli by exploring the entire posterior space, where better coverage (i.e. more visited states) leads to more accurate estimation of the energy function. As such, this metric can be seen as a proxy for diversity over neural states. The number of states $\mathbf{z}$ is correlated with diversity of neural activity patterns which might be recorded using neuroimaging techniques. Given the inherent low-dimensional nature of our $\mathbf{z}$ values, we simply count unique states rather than compute typical complexity measures such as Lempel–Ziv, which assumes the compressibility of states.

Energy function divergence We also track the surrogate objective $\mathcal{L}(\tilde{E}) = D_{KL}(\tilde{p}, p^*)$, where $\tilde{p} \propto -\exp \tilde{E}(\mathbf{z})$, defined as the KL-divergence between the posterior probability distribution induced by the auxiliary variable—a surrogate energy function $\tilde{E}(\mathbf{z})$ resulting from the acute drug-infused neuromodulation, and the target posterior induced by E^*. This value is a proxy for the magnitude of the HPP prediction error that can be expected by sampling from $\tilde{p}(\mathbf{z})$ at a given iteration, and thus smaller values correspond to reduced expected error. Within the context of our theoretical model, an optimal posterior over neural responses (i.e. $\mathcal{L}(\tilde{E}) = 0$) implies that an organism is able to infer stimuli correctly and reduce uncertainty relative to its environment. We conjecture this correlates with a reduction in stress and negative affect in an individual, and is consequently a clinically desirable outcome. Consequently, we use this value relative to a placebo baseline as a proxy for the therapeutic efficacy of a given neuromodulatory intervention.

Divergence trend monotonicity We finally consider a measure that quantifies the evolution of the KL-divergence during optimization. In an ideal scenario where temporal monotonicity is achieved, the convergence curve of the optimization process would show a consistently decreasing trend over iterations. However, in practice, there exist fluctuations or plateaus due to system noise and the influence of the neuromodulatory signal. Consequently, we resort to quantitatively capturing the overall trend in temporal monotonicity via the cumulative increase in KL-divergence values over the course of an experiment. This corresponds to the cumulative sum of positive consecutive changes in optimality gap, measured under the KL-divergence $\sum_{t\geq 0} \max\{D_{\mathrm{KL}}(\tilde{p}_{t+1}, p^*) - D_{\mathrm{KL}}(\tilde{p}_t, p^*), 0\}$ evaluated between each pair of consecutive iterations. As a result of the learning process, the optimality gap captured by the divergence should decrease over time, so the change in divergence should be negative $D_{\mathrm{KL}}(\tilde{p}_{t+1}, p^*) - D_{\mathrm{KL}}(\tilde{p}_t, p^*) \leq 0$. This metric quantifies fluctuations and transitory changes, thereby measuring the disruption in the learning process caused by neuromodulation. Within the context of our theoretical model, this corresponds to the extent to which uncertainty or misaligned internal representations are introduced by the pharmacological in-

tervention. We therefore use this value as a measure of the tolerability of the acute experience. A perfectly tolerable intervention would strictly decrease the divergence through a robust monotonically-converging optimization processes.

3 Results

In order to evaluate our proposed model of serotonergic neuromodulation on cortical dynamics, we run a series of idealized simulations of the algorithm presented above. We first examine the individual impacts of 5-HT2a and 5-HT1a neuromodulation on cortical optimization, finding that they respectively induce transient overfitting and underfitting in the optimization process of the posterior probability over neural responses. We correlate these phenomena with our metrics of therapeutic efficacy and tolerability, respectively. We next consider the effect of mixed agonism, and find that this maintains the desirable properties of both individual agonism types to provide both greater therapeutic efficacy as well as greater tolerability. Finally, we explore the full space of biased agonists and report the relative level of 5-HT2a to 5-HT1a agonism which optimally trades off between our proxy metrics of therapeutic efficacy and tolerability, finding it to be a biased 5-HT1a agonist.

3.1 5-HT2a agonism induces transient overfitting

We first consider the effects of pure 5-HT2a agonism on the final KL-divergence values across a range of dose levels (see Figure 4, 2nd row). We find that only two levels of 5-HT2a agonism decrease the KL-divergence relative to the baseline: `heavy` [two-tailed t-test: $t(99) = 3.1, p < 0.001$], and `max` [$t(99) = 4.83, p < 0.001$] (column 5). Given that we use the KL-divergence objective function as a proxy for therapeutic efficacy, we predict that only `heavy` and `max` levels of 5-HT2a agonism may be capable of inducing a therapeutically significant effect.

When we examine the energy function values over the course of the experiment, we observe that across all dose ranges there is a transient acute reduction of the energy value below the baseline. Furthermore, in cases of `heavy` and `max` dose levels, this acute reduction falls below the final values in the post-acute stage of the simulation (column 1). We also find that the gradient magnitude increases over baseline during the acute drug effect, and that this effect is likewise correlated with dose strength, with `heavy` [$t(99) = -3.249, p < 0.001$] and `max` [$t(99) = -11.391, p < 0.001$] doses inducing significantly larger gradient magnitudes during the acute effect (column 2).

The transient increase in gradient magnitude paired with the energy value reaching values lower than would be possible in the absence of neuromodulation can be interpreted as the induction of transient overfitting in belief representation during the acute phase of 5-HT2a agonism when a dose strength of `heavy` or `max` is applied. Significantly, this overfitting is not to the target distribution $p^*(\mathbf{z})$, but rather to the distribution induced by the current neuromodulated energy function $\tilde{p}(\mathbf{z})$. This transient overfitting may serve as a basis for the induction of transient (and sometimes false) insight experiences and supernatural

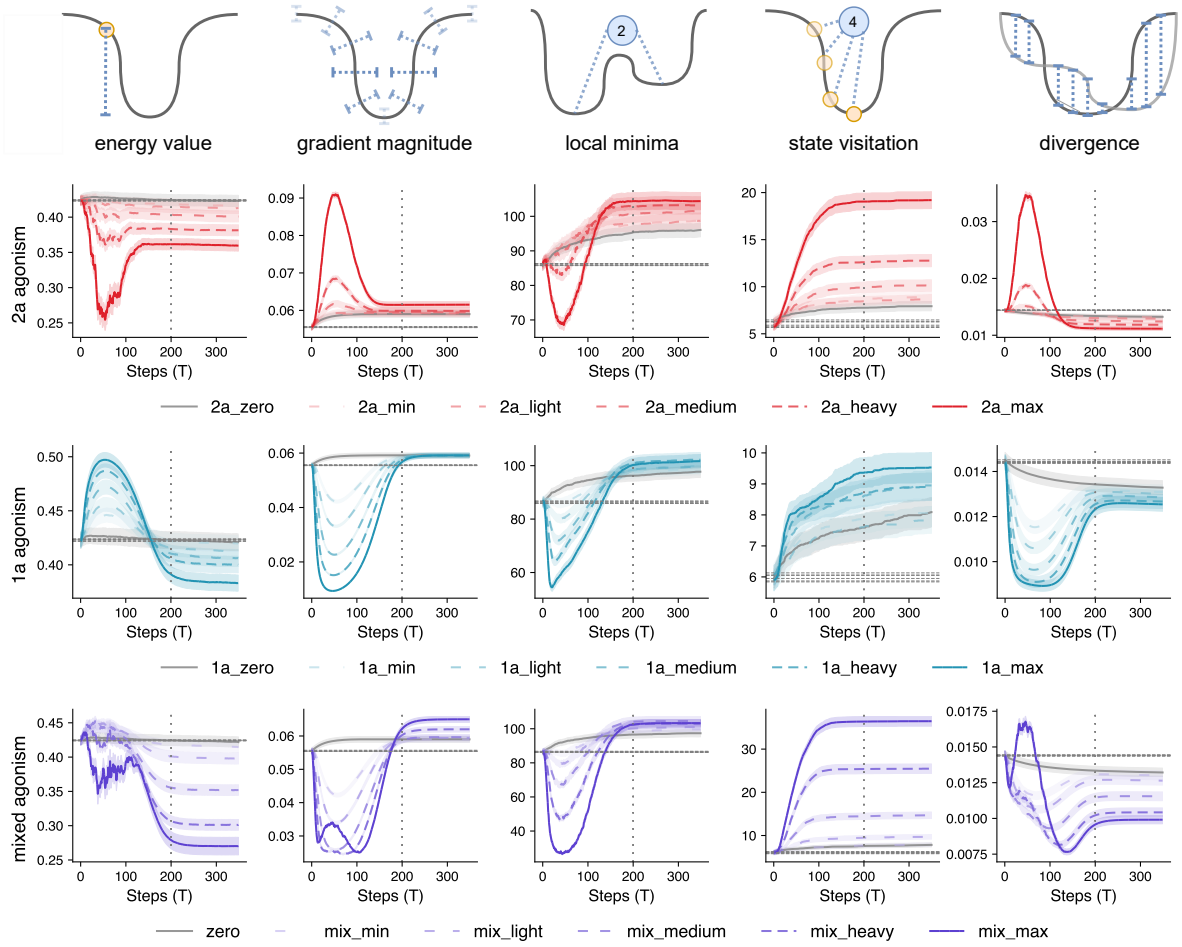

Figure 4: **(First row)** Graphical representations of the five primary metrics considered in the simulations. The representation in the header of each column is simplified to use a one-dimensional energy function for visualization purposes whereas the energy function used in the simulation is two-dimensional. The representation corresponds to the neuromodulated energy function $\tilde{E}(\mathbf{z})$. Orange circles correspond to samples from the EBM $\tilde{p}(\mathbf{z})$. For each metric we convey in blue an intuitive depiction of its measurement. The x-axis shows the number of steps T in Alg.1. **(Second row)** Shows the average metric values for 5-HT2a pure agonist experiments. **(Third row)** Shows the average metric values for 5-HT1a pure agonist experiments. **(Fourth row)** Shows the average metric values for balanced mixed 5-HT2a and 5-HT1a agonist experiments. In all time series plots the solid grey line corresponds to the baseline without neuromodulation. Horizontal dashed line represents value of metric at start of drug intervention. Vertical dotted line represents point at which drug effect ends. Semi-transparent bands around lines correspond to standard error.

percepts during the acute effects of psychedelic drugs (Griffiths, Hurwitz, Davis, Johnson, & Jesse, 2019; McGovern et al., 2023).

Turning to the remaining metrics, we observe that local minima count decreases relative to the baseline during the acute phase of neuromodulation for the `heavy` and `max` dose levels [both $p < 0.001$] (column 3). State visitation count in contrast increases according to the dose strength, with `medium`, `heavy`, and `max`

doses producing significantly higher state counts than baseline [all $p < 0.005$] (column 4). Finally, we find that as the dose strength increases the KL-divergence metric increases as well, with significant differences found for all levels [all $p < 0.001$] (column 5).

As a result of 5-HT2a agonism, both the number of attractor states and the location and energy value of these states change significantly during the acute drug phase and this effect is correlated with drug strength. These results suggest the number of attractor states decreases and the optimization procedure exhibits oscillatory or diverging dynamics during the acute drug phase, an effect correlated with drug strength. This finding is consistent with observed increases in the entropy of cortical dynamics as measured by fMRI in individuals experiencing the acute effects of serotonergic psychedelics (R. L. Carhart-Harris et al., 2014). Furthermore, the acute increase in divergence at `heavy` and `max` dose levels suggests that 5-HT2a agonism produces a state in which there is an expected increase in divergence, and thus may correlated with a more challenging subjective experience.

3.2 5-HT1a agonism induces transient underfitting

We next consider the effects of pure 5-HT1a agonism on the metrics of interest across the same range of doses (see Figure 4, 3rd row). We first find that 5-HT1a agonism is unable to produce a lasting decrease in the KL-divergence between the learning and target posterior probability distributions, even at the `max` dose level (column 5). This suggests that within the assumptions of our theoretical model, 5-HT1a agonism alone may have limited long-term therapeutic efficacy. On the other hand, during the acute phase of the simulation we find that 5-HT1a agonism transiently, but significantly decreases the KL-divergence values for all levels of agonism [all $p < 0.001$] (column 5). This suggests that 5-HT1a agonists may provide acute therapeutic relief, even if these effects do not persist once the drug effect has ended.

Similar to what we observe with 5-HT2a agonism, we find that 5-HT1a agonism is able to induce decreases in the post-acute energy function values. We find however that only the `max` dose level is able to produce a significant decrease compared to the baseline [$t(99) = 3.498, p = 0.001$] (column 1). Unlike 5-HT2a agonism however, 5-HT1a agonism induces an opposite effect during the acute phase of the drug by increasing the energy value above baseline in a manner correlated with drug strength, with `medium`, `heavy`, and `max` doses producing significant acute increases [all $p < 0.001$]. Examining the change in gradient magnitude, we find an inverse effect to that of 5-HT2a agonism, with 5-HT1a agonism inducing significant transient decreases in gradient magnitude as a function of drug strength for all dose levels [all $p < 0.001$] (column 2).

Transient increases in energy values along with decreases in gradient magnitude together can be interpreted as the induction of transient underfitting in neural response inference procedure during the acute effects of 5-HT1a agonism. This transient underfitting is equivalent to belief relaxation, where the precision of encoded beliefs is significantly decreased. These transient effects are also correlated with the acute decrease in the KL-divergence objective metric. We can connect this finding to the clinically observed phenomenon of temporary relief in individuals from maladaptive beliefs during the acute administration of

drugs such as MDMA (Greer & Tolbert, 1990).

Examining the remaining metrics, we find that 5-HT1a agonism produces a significant acute decrease in the number of local minima within the energy function which are greater than those of 5-HT2a agonism (column 3), and that this effect is correlated with dose strength [all dose levels above min $p < 0.001$]. We find that the number of visited states is also significantly increased by 5-HT1a agonism (column 4), though at a level which is considerably less than that of 5-HT2a agonism, with only `max` [$t(99) = -2.104, p = 0.037$] dose levels produce significantly greater state counts than baseline. These results suggest that 5-HT1a agonism alone is unlikely to produce the increases in neural response diversity comparable to those which result from administration of serotonergic psychedelics (R. L. Carhart-Harris et al., 2014).

3.3 Effect of balanced serotonin agonism

We next report the effect of applying evenly balanced 5-HT2a and 5-HT1a agonism on the metrics of interest across a range of doses (see Figure 4, 4th row). We find that these neuromodulation patterns are able to produce final KL-divergence values which are significantly lower than the considered baseline for the `medium`, `heavy`, and `max` dose levels [all $p < 0.001$] (column 5). This suggests that these dose levels may provide long-term therapeutic benefit for individuals under our conjectured therapeutic correlation. We also observe that for all doses except for `max`, there is a transient decrease in KL-divergence values during the acute phase of the drug. We correlate this observation with possible short-term therapeutic benefits to balanced 5-HT1a/5-HT2a agonism, similarly to 5-HT1a agonism.

We observe that balanced agonism also produces significant decreases in post-acute energy function values at the `medium`, `heavy`, and `max` dose levels [all $p < 0.001$] (column 1). Furthermore, these values are lower than either pure 5-HT2a or 5-HT1a agonism alone at the equivalent dose level. The acute energy function values for balanced agonists are also neither above or below pre- or post-acute levels. We also observe that mixed agonism produces significant transient decreases in gradient magnitude at all dose levels [all $p < 0.001$], but that these are less extreme than those of equivalent pure 5-HT1a agonism.

Taken together, these results suggest that balanced agonism is able to reduce post-acute energy function values without the introduction of either predominant underfitting or overfitting of the neural response function during the acute phase. Given the transient but unstable decrease in gradient magnitude, we can still interpret these interventions as inducing a form a belief relaxation, although one which is not as consistent as that of pure 5-HT1a agonism. It also differs from pure 5-HT1a agonism in that the states being sampled have a higher likelihood than they would under the baseline (as a result of the lower energy value), but a lower precision (as a result of the lower average gradient magnitude).

Turning to the remaining metrics, we find that balanced agonism induces a decrease in the number of local minima during the acute drug phase which is greater than that seen for either 5-HT1a or 5-HT2a agonism alone [all except `min` dose level $p < 0.001$] (column 3). Correspondingly, state visitation count is also increased in balanced agonism over 5-HT2a and 5-HT1a agonism, with all dose levels above `min` producing significantly more visited states [$p < 0.01$] (column 4). This would likewise correspond to a

measurable increase in neural response diversity, which is one of the signatures of mixed serotonergic psychedelics such as LSD and psilocybin (R. L. Carhart-Harris et al., 2014).

Relative to monotonic decrease of the optimization objective over the course of the simulation observed in the baseline, here we observe that there are significant increases in KL-divergence [all $p < 0.001$] (column 5). This suggests under our therapeutic formulation that balanced agonists may also present questions of tolerability. Notably however we see more monotonicity in the KL-divergence trend compared to that of 5-HT2a agonists, suggesting an overall better potential tolerability profile. Collectively, these results suggest that balanced agonists may provide both greater short- and long-term therapeutic efficacy than either 5-HT2a or 5-HT1a agonists in addition to providing a better tolerability profile than pure 5-HT2a agonists.

3.4 Exploring the optimal serotonin agonism bias

Above we found a synergistic effect of balanced agonism on post-acute values of the KL-divergence as well as the state visitation and acute local minima counts. Balanced agonism also produced a moderating effect on the non-monotonicity of the KL-divergence trend. Given these improvements as compared to pure agonists, we hypothesized that a neuromodulation profile which ideally balances our conjectured proxy metrics for therapeutic efficacy and tolerability might involve a form of biased rather than balanced mixed agonism. In order to determine this, we simulated all possible permutations of 5-HT2a and 5-HT1a dose strengths. These results are presented in Figure 5. We consider final values of the KL-divergence and the non-monotonicity of the KL-divergence trend over the optimization procedure as the primary two metrics of interest, given their interpretation within our model as proxies of therapeutic efficacy and tolerability, respectively.

We find that a `max` dose of a balanced 5-HT2a and 5-HT1a agonist produces the greatest post-acute reduction in optimization objective. Likewise, the general trend is that increases across both agonist ranges produces improvements in this metric, with a greater portion of the improvement coming from increases in 5-HT2a rather than 5-HT1a agonism. This result alone would imply that a maximal dose of a balanced agonist may provide the greatest long-term therapeutic benefit. In contrast, we find that the KL-divergence trend is most monotonic at low levels of both 5-HT2a and 5-HT1a agonism, but is most strongly affected by increases in 5-HT2a agonism. Accordingly, smaller doses, and a bias towards 5-HT1a agonism may improve the expected tolerability of the drug intervention.

In order to assess the trade-off between these two metrics, we normalized values of both final divergence and divergence trend non-monotonicity and equally weighted them to produce a new score metric which balances the two desirable drug properties. Here we find that a `max` level of 5-HT1a agonism paired with a `medium` or `heavy` level of 5-HT2a agonism provides the optimal trade-off between the two metrics of interest. In particular, this level of biased agonism produces post-acute KL-divergence values which are comparable to that of pure max 5-HT2a agonism, but does so with a trend which is significantly more monotonic than that produced by max 5-HT2a agonism. For cases in which tolerability and short-term effi-

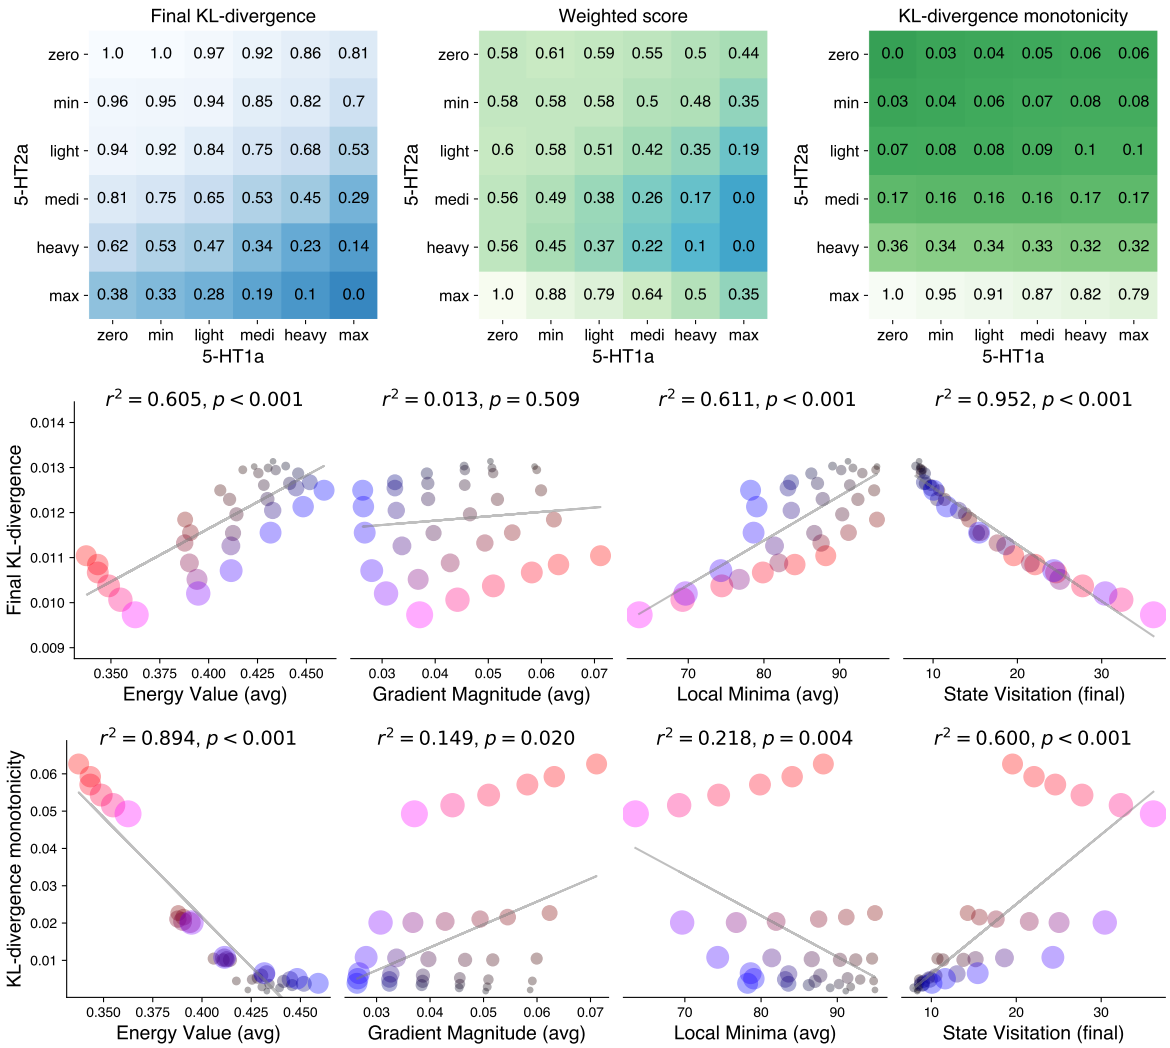

Figure 5: **(Top-left)** Matrix reporting normalized final energy values for set of possible 5-HT2a and 5-HT1a agonist combinations. Lower values correspond to lower energy. **(Top-right)** Matrix reporting normalized final non-monotonicity values for set of possible 5-HT2a and 5-HT1a agonist combinations. Lower values correspond to higher monotonicity. **(Top-middle)** Matrix reporting normalized weighted energy values for set of possible 5-HT2a and 5-HT1a agonist combinations. Lower values correspond to lower weighted score. **(Middle)** Correlations between metrics and final KL-divergence values. **(Bottom)** Correlations between metrics and KL-divergence monotonicity. Color of data points correspond to relative 5-HT2a (red) and 5-HT1a (blue) agonism. Size of data point corresponds to average dose level.

cacy is more important, a greater 5-HT1a bias would be desirable. In contrast, for cases in which long-term therapeutic efficacy is most important, a balanced agonism profile is desirable.

We finally consider the relationship between the metrics describing the acute drug effects and our two proxy outcome metrics. We find that while average gradient magnitude is uncorrelated with the final KL-divergence, final energy value, average local minima count, and final state visitation count are all significantly correlated with final KL-divergence value [linear regression: all $p < 0.001$] (see Figure 5). Among

these, state visitation count is able to account for nearly all of the variance in final optimization objective, with an $r^2 = 0.952$. This finding aligns with the hypothesis that there is a causal relationship between increases in neural response diversity and therapeutic outcomes (R. L. Carhart-Harris et al., 2014). From an optimization perspective this is natural since increased exploration of the posterior probability landscape can improve the sampling procedure, consequently leading to better learning of the target.

Average energy value, average gradient magnitude, and final state visitation count are also linearly predictive of monotonicity in the KL-divergence trend [all $p < 0.01$]. Of these, we find that average energy value accounts for the greatest variance in the proxy metric, with an $r^2 = 0.894$. Although this relationship is not causal, it does imply that overfitting in the energy function during the acute drug experience is associated with greater divergence between the current and target energy distributions. Given that this is our hypothesized proxy for tolerability, it suggests that acute overfitting may be associated with decreases in drug tolerability.

4 Discussion

Operating within the preexisting theoretical models of predictive processing and energy-based models of neural response dynamics, we proposed concrete mechanisms whereby 5-HT2a and 5-HT1a agonism drugs such as psilocybin, LSD, and DMT may exert their effects, and simulated these effects across a range of dose strength and relative agonism biases. This model and simulations results have a number of potential implications for theoretical debates concerning belief representation, and motivate further study concerning clinical application of psychedelic-assisted therapy, as well as next-generation psychedelic drug discovery. Below we attempt to address these considerations.

4.1 Alterations in belief representation under psychedelics

Overall we found that both 5-HT2a and 5-HT1a agonists are able to produce meaningful improvements in KL-divergence, the metric which serves as our proxy metric for therapeutic efficacy. In particular, 5-HT2a agonism provides greater long-term benefits across the dose response curve, whereas 5-HT1a agonism confers benefit only during the acute phase of neuromodulation. Our findings predict that 5-HT2a agonists can be characterized as acutely inducing a transient overfitting into the cortical dynamics, resulting in the generation of a sequence of ever-shifting high-precision beliefs which may or may not be veridical. In contrast, 5-HT1a agonists can be characterized as acutely inducing transient underfitting into cortical dynamics, resulting in the generation of a more stable sequence of lower-precision beliefs than either a placebo baseline or 5-HT2a agonism.

We can interpret these observations in light of the debate around whether psychedelics are belief-strengthening, belief-relaxing, or some third category (R. L. Carhart-Harris & Friston, 2019; Safron, 2020). Strictly speaking, our findings predict that 5-HT2a agonists such as DOI and DOB will be belief-strengthening. This is apparent in the sharp increase in gradient magnitude produced by higher doses of 5-HT2a agonism

in our simulations as well as the acute decrease in energy function values relative to baseline. This must be qualified however by the fact that this strengthening of beliefs does not refer to post-acute outcomes. Rather than a fixed set of beliefs being strengthened for the duration of the drug effect, the increase of belief precision is stochastic and transient, thus resulting in a higher state visitation count as attractors in the energy landscape are constantly disrupted during the acute drug effect. This account is able to reconcile the phenomenological evidence for pareidolia and insight experiences, which are consistent with belief strengthening, during acute psychedelic experiences with the finding that psychedelics increase cortical entropy entirely within a pure 5-HT2a agonist model of psychedelics.

Our understanding of the belief-altering properties of psychedelics becomes more nuanced however when we consider the significant modulating role that 5-HT1a agonism also plays in the acute effects of psychedelics. Our simulations predict that 5-HT1a agonism produces a belief-relaxation effect, which can be measured both by the decrease in gradient magnitude as well as energy function values during the acute phase of the experiment. Both of these together serve as our proxies for belief precision, suggesting that 5-HT1a agonists relax beliefs acutely. Furthermore, we find that 5-HT1a agonists decrease the number of local minima present in the energy function surface, which corresponds to the number of stable beliefs a network can encode. This effect is consistent with the phenomenological reports of 5-MeO-DMT use, a highly biased 5-HT1a agonism which at heavy doses produces experiences in which users become unable to represent objects, ideas, space, or even time (Reckweg et al., 2022; Dourron et al., 2023; Gómez-Emilsson, 2020).

We also consider relatively balanced serotonin agonists such as psilocybin, LSD, or DMT, which have also been the subject of the greatest scientific focus in recent years. Here we predict both belief-relaxation and belief-strengthening operating simultaneously, albeit in different ways. In this case we find both an increase in state visitation count as well as an overall decrease in gradient magnitude and number of local minima. We likewise find that balanced agonists produce a more monotonic trend in KL-divergence compared to pure 5-HT2a agonists, suggesting greater potential tolerability, while also decreasing the post-acute divergence values beyond what the placebo or either class of pure agonists are capable of. Taken together, these findings are both consistent with the ALBUS and REBUS models of psychedelic action. Whether psychedelics are viewed as strengthening or relaxing beliefs is then a question of the particular drug, timescale, and proxy metric of interest under consideration.

This opponency effect between 5-HT2a and 5-HT1a agonism on belief representation has previously been identified in the context of a model contrasting DMT and 5-MeO-DMT (Gómez-Emilsson, 2020). It is also is consistent with the passive vs active coping model introduced by R. L. Carhart-Harris and Nutt, and the automaticity vs flux model of Shine et al.. In all cases, 5-HT1a agonism can be understood to acutely relax beliefs whereas 5-HT2a agonism acutely and stochastically strengthens beliefs. The effect of the former is to produce a state which may provide short-term relief from stress, whereas latter provides the potential for long-term benefit. Significantly, the class of serotonergic psychedelics currently being clinically evaluated are capable of inducing both effects simultaneously due to their mixed binding profiles, thus providing both benefits in a combined manner.

4.2 Connection to endogenous neuromodulation

Our work also has implications for endogenous neuromodulation and is consistent with current theoretical work on the role of postsynaptic 5-HT agonism in the cortex as a stress-response system (R. L. Carhart-Harris & Nutt, 2017; Shine et al., 2022). Here we extend this perspective by proposing a set of concrete mechanisms by which cognitive responses to stress can be deployed to improve the adaptedness of an organism within a given environmental context. Although we did not explore it here, a natural extension of our model is to consider the mechanism by which endogenous neuromodulation takes place in response to stress.

While critically important to affect and cognition, the receptor populations we consider here are only two of a dozen 5-HT receptor types expressed throughout the central nervous system. There is evidence that other 5-HT receptor populations including 5-HT1B and 5-HT2C may also play important roles in the etiology of certain mood disorders (Ruf & Bhagwagar, 2009; Millan, 2005; Carr & Lucki, 2011). These two populations in particular are also common targets of classic psychedelics including LSD and psilocybin, though they do not possess the same level level of binding affinity or causal influence as either 5-HT2a or 5-HT1a (Rickli et al., 2016; Deco et al., 2018). The action of these receptor populations on cortical activity is also less straightforward than the apparent opponency between 5-HT2a postsynaptic excitation and 5-HT1a postsynaptic inhibition. Further basic neuroscientific research will be required to fully understand the complex role that the family of 5-HT receptor populations plays throughout the cortex and the central nervous system more broadly.

There is also a complex interplay between serotonin and other neuromodulatory systems such as dopamine (Giacomelli, Palmery, Romanelli, Cheng, & Silvestrini, 1998; Sakashita et al., 2015; Grandjean et al., 2021; den Ouden et al., 2013). While previously it has been believed to exist an opponency between the two neuromodulators (Daw, Kakade, & Dayan, 2002; Boureau & Dayan, 2011), this hypothesis has more recently been refined to suggest the pair of neurotransmitters having doubly dissociable effects, particularly in the context of cognitive flexibility (den Ouden et al., 2013). The dual-receptor nature of the serotonin coping systems can also be compared to the dopamine system's balance between reflective, goal-oriented decision making, linked to the dorsomedial striatum (caudate) and involving future planning, and reflexive, habitual behavior, linked to the dorsolateral striatum (putamen) and based on past actions or previous experience (Dolan & Dayan, 2013).

4.3 Connection to neural annealing theory

A popular metaphor for the therapeutic action of psychedelics on brain dynamics is that of annealing in metallurgy or simulated annealing in probability theory (Bollweg, Maurer, & Kroll, 1997). This analogy has been formalized by various theorists into a neural annealing theory of psychedelic action (M. Johnson, 2019; R. Carhart-Harris et al., 2022). According to neural annealing theory, the dynamics of the brain are "heated," "manipulated," and then "cooled" in a manner analogous to that of the process of annealing in metallurgy. Rather than being literally heated or cooled, the neural annealing process is meant to

correspond to the phase of transient acute increases in the entropy of neural patterns of activity during psychedelic use (R. L. Carhart-Harris et al., 2014), followed by the slow process within which the cortical dynamics return to baseline—subsequently also referred to as "metaplasticity" (Nardou et al., 2023). The neural annealing theory has also been extended beyond the particulars of psychedelics to encompass any intervention which may be capable of increasing the entropy of brain dynamics, such as specific breathing or meditation techniques (Gómez-Emilsson, 2021).

The theoretical model and set of simulations which we present here serve as a possible concrete computational instantiation of the neural annealing process. In this work we demonstrate first that both 5-HT1a and 5-HT2a agonism have the capacity to increase the neural response variability, a measure of cortical entropy (Keshmiri, 2020), during the acute phase of drug effects, and that these dynamics return to baseline when the drug effect ends, thus satisfying the basic criteria of a neural annealing process. We also demonstrate that this increase in cortical entropy corresponds to a period in which the synaptic connectivity structure becomes more malleable, thus resulting in lasting changes to cortical dynamics in the post-acute phase.

A therapeutically successful neural annealing process involves two primary outcomes. The first is that cortical dynamics settle into a regime in which they are better able to reduce prediction errors for the current environmental context, which we refer to as adaptation. This is reflected at two temporal scales. The first is in the current neural dynamics settling into an attractor which is closer to the global minima than would result from a placebo in which the stochastic sampling process is local, insufficiently exploratory and thereby slow. The second is the KL-divergence between the current and target energy functions decreasing significantly. We find that in our simulations 5-HT2a agonism is able to achieve this goal at suitable dose levels, with balanced agonism producing an even stronger therapeutic effect.

The second desirable outcome of a neural annealing procedure is that the synaptic connectivity of the energy function enables fewer local minima and a smaller overall gradient magnitude. Together these are expected to prevent neural dynamics from getting trapped in overly-precise attractors in the future, which we refer to as adaptability. This increase in adaptability may be what is typically associated with post-acute increases in markers of plasticity after psychedelic drug administration (Olson, 2018). Within the context of our simulations, we find that 5-HT1a agonism is capable of producing this effect acutely. As we do not model the potential effects that psychoplastogens may have on the underlying energy function, we cannot speak to the long-term changes in local minima. Regardless, an increase in future adaptability will likely be at the expense of decreasing current adaptedness, highlighting a trade-off inherent in learning systems (Mermillod et al., 2013), including the human brain, between adaptation to a given context and adaptability to potential future contexts.

4.4 Empirically validating our model and predictions

Our theoretical model is based on a number of assumptions which are still in need of experimental validation. Our model operates on the basis of the Bayesian brain hypothesis, and assumes that cortical

dynamics are governed by a process of hierarchical predictive processing (Keller & Mrsic-Flogel, 2018). More specifically, we make the assumption that HPP can be modeled as learning and stochastic sampling in an energy-based model using gradient descent (Hoyer & Hyvärinen, 2002; Aitchison & Lengyel, 2016; Dong & Wu, 2023). Human behaviour is consistent with Bayesian inference in many sensory (Ernst & Banks, 2002), motor (Körding & Wolpert, 2004) and cognitive tasks (Tenenbaum et al., 2006). While there is accumulating evidence for predictive processing in various aspects of cortical computation, it is still not universally accepted among theorists (Walsh, McGovern, Clark, & O'Connell, 2020). A similar state of evidence exists for explicit attractor dynamics in cortical computation, which would be amenable to modeling using energy functions (Khona & Fiete, 2022). While research into the existence of such dynamics is still nascent, there is accumulating evidence that within the context of high-level decision making PFC cortical dynamics can be modeled using energy functions (Mante et al., 2013; S. Wang et al., 2023).

In our model, we assume that both 5-HT1a and 5-HT2a agonists rapidly and continuously modify the entire energy function governing cortical dynamics. Validating the particular nature of this modification would pose some difficulty given the current capabilities of neuroimaging. Current recording techniques are only able to build model of this function by sampling from a number of sequential trials over an extended period of time (S. Wang et al., 2023). Novel experimental design or data analysis methods may serve to provide insight into the effects of these receptor populations on the induced energy function in simple contexts such as two choice decision making tasks or basic perceptual tasks.

Verifying the predictions concerning the clinical effects of serotonergic psychedelics is potentially more tractable within pre-existing methodological frameworks. The simulation results presented in Section 3 provide the basis for a number of predictions which may guide neuroimaging research, clinical protocols, and drug development. The first of these concerns the modulating role of 5-HT1a agonism in the psychedelic experience. Our model predicts that unlike 5-HT2a agonist psychedelics, 5-H1a agonists should not significantly increase acute cortical entropy, as measured by the number of unique cortical states visited as well as the transition dynamics between states (R. L. Carhart-Harris et al., 2014). Our model also predicts that balanced mixed agonists will produce levels of cortical entropy which are significantly greater than either 5-HT1a or 5-HT2a agonists alone, suggesting a complex synergistic effect between the two receptor systems.

As has been theorized by R. Carhart-Harris et al., we expect that increases in cortical entropy will be significantly correlated with greater post-treatment therapeutic efficacy. Our model also predicts that 5-HT2a agonists will provide greater long-term therapeutic efficacy than 5-HT1a agonists, and that mixed agonists will produce the greatest therapeutic effect, as they also produce the greatest levels of cortical entropy. Our model also predicts that mixed agonists will have greater therapeutic efficacy at lower doses than either pure 5-HT2a or 5-HT1a agonists alone, which has implications for both safety and tolerability of treatment.

From a clinical perspective, it is also desirable to maximize the tolerability and short-term efficacy of the acute experience during psychedelic-assisted therapy. Our model predicts that pure 5-HT1a agonists will be able to produce measurable therapeutic improvements during the acute phase of the drug effect

as well as provide greater acute tolerability. In contrast, we predict that a pure 5-HT2a agonist may have the oppose acute effect. This is consistent with pre-existing findings that 5-HT1a agonists produce acute anxiolytic effects, while 5-HT2a agonists can be anxiogenic. There already exists some evidence that increasing 5-HT1a receptor activation during the acute psychedelic experience decreases the intensity of the reported acute experience (Strassman, 1995; Pokorny, Preller, Kraehenmann, & Vollenweider, 2016), but this should be examined in greater depth by explicitly modulating the relative levels of 5-HT2a and 5-HT1a agonism across a range of doses.

Our model also has implications for research into the experience of insight, a key aspect of problem solving (Topolinski & Reber, 2010; Laukkonen et al., 2023). Psychedelic use is known to transiently induce insight experiences, some of which may be so-called "false insights," due to their non-veridical nature (McGovern et al., 2023). We can operationalize the insight experience as the cortical dynamics for a given internal representation arriving in a novel attractor within the energy function landscape that is able to reduce prediction error significantly better than previous attractor states. The intensity of the insight experience then corresponds to the value of the energy function at that point in combination with the local gradient magnitude around the point, the latter of which is a measure of the precision with which the novel belief is encoded.

Our model predicts that 5-HT2a agonists induce insight experiences through the generation of transient attractor states with unusually small energy values and large gradient magnitudes around them. This can potentially account for the intensity of the reported insights during psychedelic-assisted therapy, along with the potential for these insights to be false or spurious, as they are generated by an energy function which has been perturbed by exogenous neuromodulation from the drug and may not align with the target. Our model also predicts that 5-HT1a agonism will modulate the number and intensity of these insight experiences, potentially resulting in fewer, but more stable insights during the acute treatment experience. These insights may then be better positioned to serve as a useful basis for the post-acute process of psychotherapeutic integration.

4.5 Potential implications for drug design

Current efforts in psychedelic drug design are largely focused on the development of non-hallucinogenic 5-HT2a agonists (Cameron et al., 2021; Qu et al., 2023; Lewis et al., 2023). Although progress has been made in this area, there is evidence to suggest that the acute phenomenology of the psychedelic experience plays a non-trivial causal role in the long-term therapeutic efficacy of psychedelic-assisted therapy (Griffiths et al., 2011; Yaden & Griffiths, 2020). If this is indeed the case, then there is value in understanding the pharmacological mechanisms responsible for different aspects of the acute phenomenology and their relationship to clinical outcomes.

Early clinical evidence suggests that 5-MeO-DMT, a highly biased 5-HT1a agonist psychedelic, has significant potential as a treatment for major depressive disorder (Uthaug et al., 2019; Reckweg et al., 2023). Some concern exists however regarding the safety profile of 5-MeO-DMT both in the acute and

post-acute phases of the drug effect. One particular concern is the experience of "reactivations" in which the drug phenomenology of the effect spontaneously returns days or weeks after acutely consuming the drug (Dourron et al., 2023). Other concerns include the potential for 5-MeO-DMT to induce seizures in individuals with a predisposition, or for it to increase body temperature to unsafe levels.

Our model predicts that a highly-biased 5-HT1a agonist such as 5-MeO-DMT will only have clinically relevant effects at very high doses, and this seems to be the case from early empirical studies (Reckweg et al., 2023). Increases in dose strength however produce decreases in psychological tolerability as well as corresponding undesirable side-effects resulting from high levels of 5-HT agonism at non-clinically relevant sites throughout the nervous system, potentially leading to serotonin syndrome in extreme cases (Volpi-Abadie, Kaye, & Kaye, 2013). An ideal balance between 5-HT2a and 5-HT1a agonism may involve a relatively smaller but pharmacologically meaningful bias towards 5-HT1 agonism over 5-HT2a agonism.

Notably, 5-HT itself is biased in this way with a three to five fold preference for 5-HT1a over 5-HT2a agonism (Barnes & Sharp, 1999), which is consistent with the optimal binding profile reported in our results. It may prove worthwhile to explore the development of drugs with a similar profile to 5-HT itself, but which are acutely psychoactive and amenable for use in psychedelic-assisted therapy. 5-MeO-DMT is part of a larger class of "5-methoxy" psychedelic substances which all posses greater binding affinity for 5-HT1a over 5-HT2a receptors (Puigseslloses et al., 2024). In particular, drugs such as 5-MeO-MiPT or 5-MeO-DiPT have been reported to have binding affinity ratios within the range highlighted here as potentially optimal, and should be given greater potential clinical attention.

The existence of modern drug discovery methods utilizing deep learning make the process of discovering a molecule with the desirable binding properties considerably easier than in previous decades (Dara, Dhamercherla, Jadav, Babu, & Ahsan, 2022; King-Smith et al., 2024). Recent work has demonstrated the ability to utilize state-of-the-art protein folding models to develop candidate molecules which possess strong 5-HT2a binding profiles (Lyu et al., 2023). The extension of this technology to discovering 5-HT1a binding molecules is well within the capabilities of current models, and may yield the development of promising new psychotherapeutic agents with fewer off-target effects than those of currently known psychedelics.

4.6 Limitations and future directions

There are a number of limitations to both our theoretical model and the simulations presented here. Many of these were made in order to avoid unnecessary obfuscating complexity into a novel paradigm, but they should nonetheless be addressed in the future. First, we limit our consideration to the dynamics of only a single fixed target energy function and a single associated learned energy function. This is a significant simplification of real neural dynamics, where predictive processing is hypothesized to take place at many levels of a spatiotemporal hierarchy. We are also unable to capture the potential for complex non-linear interactions between relevant functional networks within the cortex that may be induced by 5-HT neuromodulation.

Second, we make the simplifying assumption that the environmental context which the organism finds

itself in is held fixed for the duration of the simulation. Consequently, the generative model's parameters, from latent neural responses to reconstructed stimuli, presumably context-dependent, are assumed fixed, and omitted throughout the drug experience. This allowed us to avoid considerations of how the external dynamics of the world may alter the sensitivity of the neural optimization process, particularly its response to sensory information, or attentional modulation of prediction errors over time. At the same time, it has prevented us from studying how the drug effect may impact lifelong or continual learning (Parisi, Kemker, Part, Kanan, & Wermter, 2018). A number of clinical protocols for psychedelic-assisted therapy have also utilized repeated doses over the course of multiple weeks (R. Carhart-Harris et al., 2021; Mitchell et al., 2023), whereas here we focus on simulating the effects of a single dose, which may be less effective.

Third, we resorted to a non-parametric representation of the optimization landscape, which is modeled via a two-dimensional neural response $\mathbf{z}$, which we update directly by gradient descent on the energy function. While this is a significant simplification of neural dynamics in living organisms, there is evidence to suggest that for a given functional network, the underlying optimization dynamics may often lay on a lower-dimensional manifold (Khona & Fiete, 2022). Indeed, multiple studies have found that cortical attractor dynamics can be projected down to as simple as a 1D space while retaining significant predictive power in the case of manifolds for high-level decision making and representation (Inagaki, Fontolan, Romani, & Svoboda, 2019; Okazawa, Hatch, Mancoo, Machens, & Kiani, 2021; S. Wang et al., 2023).

Fourth, our model also does not take into account the role that action plays in the perceptual process. Given that perception and belief formation is ultimately in the service of adapted and goal-directed action, this would serve as a worthwhile extension to the model. The PFC in particular has been heavily studied for not only its role in high-level belief representation, but also for high-level action selection (Domenech & Koechlin, 2015; S. Wang et al., 2023). Numerous theoretical models currently exist which attempt to provide normative descriptions of this role of the PFC, often from the perspective of reinforcement learning, meta-learning, or both (J. X. Wang et al., 2018; Muller et al., 2024).

Fifth, it is the case that the simulations utilized in this work are abstracted from real neural data. Each of the starting energy functions are sampled from a set of structured noise functions, rather than being naturally induced by a specific set of parameters and perceptual or behavioral task. A promising future direction is to extend this work by considering datasets and excitatory-inhibitory models represented via recurrent neural networks which would themselves induce more ecologically valid energy functions. This would allow for a more thorough examination of the neuromodulatory properties of 5-HT1a and 5-HT2a agonists under conditions which are more similar to the actual set of tasks which an organism may face in its daily life. The currently expanding literature on work utilizing deep neural networks as models of brain function may be of particular interest towards realizing this goal (Aitchison & Lengyel, 2016; Barak, 2017; Richards et al., 2019a; Juliani et al., 2024; Dong & Wu, 2023).

Finally, in order to realize the potential benefits of next-generation psychotherapeutic agents, we emphasize the importance of a measured and careful approach to any research in the mental health and psychedelic drugs space. It is also important that due attention be paid to issues of health equity and ethics to ensure that the consequences of such research are beneficial and equitably distributed.

4.7 Conclusion

Psychedelic-assisted psychotherapy has the potential to improve the lives of many individuals suffering from cognitive, affective, and behavioral disorders which are not amenable to current treatment modalities (M. W. Johnson & Griffiths, 2017; M. W. Johnson et al., 2019). A critical step towards realizing this potential is to develop sophisticated models which can be used to predict the therapeutic efficacy of candidate psychedelic substances and doses. While the amount of clinical and neuroscientific evidence which can be used to develop such models continues to accumulate at a rapid pace (R. Carhart-Harris et al., 2021; Doss, de Wit, & Gallo, 2023), actionable models are still in their infancy.

One popular theoretical framework understands psychedelics primarily as pharmacological agents which engender greater neuroplasticity across various temporal scales (Olson, 2018; Calder & Hasler, 2023). While appealing for its simplicity, this model cannot account for the mediating role that the acute phenomenological experience and its correlated neural activity play in driving positive clinical outcomes (Yaden & Griffiths, 2020). In contrast, frameworks privileging changes to subjective experience have focused on the role of belief alteration (Safron, 2020; Letheby, 2021), typically in the form of belief relaxation as a result of psychedelic use both acutely and post-acutely (R. L. Carhart-Harris & Friston, 2019). Here we presented a theoretical model based on predictive processing and a set of simulations in which this hypothesized process of belief alteration can be formally studied. We hope that it will be able to guide further clinical and neuroscientific research on this topic.

Within our framework, we modeled the acute action of psychedelics on belief representation as the result of two unique sets of effects, generated by the populations of 5-HT2a and 5-HT1a receptors present in the prefrontal cortex. We operationalize these effects as two different forms of modulation of the process which encodes and maintains belief representations. At a high level, we find that both the effects of stochastic perturbation through postsynaptic excitation (5-HT2a) as well as uniform smoothing through postsynaptic inhibition (5-HT1a) result in improvements in this optimization process across various metrics, thus pointing to a potential therapeutic effect for both. We characterize these effects in optimization terminology as 5-HT2a agonism introducing a transient overfitting of beliefs and 5-HT1a agonism introducing a transient underfitting of beliefs.

We predict that mixed agonists which activate both 5-HT2a and 5-HT1a populations will have the most desirable properties, as they may balance both long-term therapeutic efficacy as well as acute drug tolerability. Furthermore, a bias towards preferential 5-HT1a agonism may provide an even greater relative benefit to tolerability with only a minimal decrease in long-term therapeutic efficacy. Interestingly, the psychedelics currently being most thoroughly investigated for their potential as clinical tools all have various levels of mixed 5-HT agonism. Rather than being an undesirable property to be engineered away, this mixed agonism may be key to the unique profile of subjective tolerability and therapeutic efficacy which psilocybin, LSD, and DMT seem to possess, and pure 5-HT2a agonists such as DOI and DOB may not.

Looking ahead, increasing clinical interest is being placed on the biased 5-HT1a agonist psychedelic 5-MeO-DMT (Dourron et al., 2023), a substance which is capable of inducing powerful and long-lasting

antidepressant effects in minutes (Uthaug et al., 2019). Clinical attention is also beginning to be paid to the positive effects of co-administration of MDMA along with LSD or psilocybin (Straumann et al., 2023; Zeifman, Kettner, et al., 2023), a combination which may also result in a net 5-HT1a agonism bias and greater tolerability (Schechter, 1998). Rather than focusing on the development of non-hallucinogenic psychedelics, which may prove to have limited efficacy outside of animal models (Yaden & Griffiths, 2020), significant evidence exists in support of exploring the pharmacological design space of psychedelic substances with a biased 5-HT1a binding profile similar to that of 5-MeO-DMT or LSD co-administered with MDMA. Doing so may yield substances which are psychologically safer and more effective for clinical use than those currently under investigation. Such novel substances would enable the more mainstream adoption of psychedelic-assisted therapy, creating life-changing opportunities for the large population of individuals whose needs have been unmet by current front-line psychiatric treatments.

5 Acknowledgements

The authors thank Andrés Gómez Emilsson, Alex Meurice, Blake Richards, Colin Bredenberg, Michael Johnson, and Yad Konrad for the valuable insights they shared during preliminary discussions of the ideas eventually presented here. VC is grateful for support from Fonds de recherche du Québec – Nature et technologies (FRQNT), B2X (dossier 321675) and IVADO, Fonds d'excellence en recherche Apogée Canada, Bourse d'excellence au doctorat 2021.

References

Abraham, W. C., & Bear, M. F. (1996, Apr). Metaplasticity: the plasticity of synaptic plasticity. *Trends Neurosci*, *19*(4), 126–130.

Ackley, D. H., Hinton, G. E., & Sejnowski, T. J. (1985). A learning algorithm for boltzmann machines. *Cognitive science*, *9*(1), 147–169.

Aghajanian, G., & Marek, G. (1997). Serotonin induces excitatory postsynaptic potentials in apical dendrites of neocortical pyramidal cells. *Neuropharmacology*, *36*(4-5), 589–599.

Ahmed, Z., Roux, N. L., Norouzi, M., & Schuurmans, D. (2019). Understanding the impact of entropy on policy optimization. In K. Chaudhuri & R. Salakhutdinov (Eds.), *Proceedings of the 36th international conference on machine learning, ICML 2019, 9-15 june 2019, long beach, california, USA* (Vol. 97, pp. 151–160). PMLR. Retrieved from http://proceedings.mlr.press/v97/ahmed19a.html

Aitchison, L., & Lengyel, M. (2016, Dec). The hamiltonian brain: Efficient probabilistic inference with excitatory-inhibitory neural circuit dynamics. *PLoS Comput Biol*, *12*(12), e1005186.

Ajalloeian, A., & Stich, S. U. (2020). Analysis of SGD with biased gradient estimators. *CoRR*, *abs/2008.00051*. Retrieved from https://arxiv.org/abs/2008.00051

Alvarez, B. D., Morales, C. A., & Amodeo, D. A. (2021). Impact of specific serotonin receptor modulation on behavioral flexibility. *Pharmacology Biochemistry and Behavior*, *209*, 173243.

Banasr, M., Hery, M., Printemps, R., & Daszuta, A. (2004). Serotonin-induced increases in adult cell proliferation and neurogenesis are mediated through different and common 5-ht receptor subtypes in the dentate gyrus and the subventricular zone. *Neuropsychopharmacology*, *29*(3), 450–460.

Banerjee, A., Merugu, S., Dhillon, I. S., & Ghosh, J. (2005). Clustering with bregman divergences. *J. Mach. Learn. Res.*, *6*, 1705-1749. Retrieved from https://api.semanticscholar.org/CorpusID:8197416

Banushi, B., & Polito, V. (2023). A comprehensive review of the current status of the cellular neurobiology of psychedelics. *Biology*, *12*(11), 1380.

Barak, O. (2017). Recurrent neural networks as versatile tools of neuroscience research. *Current opinion in neurobiology*, *46*, 1–6.

Barnes, N. M., & Sharp, T. (1999). A review of central 5-ht receptors and their function. *Neuropharmacology*, *38*(8), 1083–1152.

Beal, M. J. (2003). Variational algorithms for approximate bayesian inference.. Retrieved from https://api.semanticscholar.org/CorpusID:11861569

Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, *10*(9), 1214–1221.

Beliveau, V., Ganz, M., Feng, L., Ozenne, B., Højgaard, L., Fisher, P. M., ... Knudsen, G. M. (2017). A high-resolution in vivo atlas of the human brain's serotonin system. *Journal of Neuroscience*, *37*(1), 120–128.

Bellec, G., Scherr, F., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., & Maass, W. (2020). A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature Communications*, *11*(1), 3625.

Benneyworth, M. A., Smith, R. L., Barrett, R. J., & Sanders-Bush, E. (2005). Complex discriminative stimulus properties of (+) lysergic acid diethylamide (lsd) in c57bl/6j mice. *Psychopharmacology*, *179*, 854–862.

Bliss, T. V., & Lomo, T. (1973, Jul). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J Physiol*, *232*(2), 331–356.

Blondel, M., Llinares-López, F., Dadashi, R., Hussenot, L., & Geist, M. (2022). *Learning energy networks with generalized fenchel-young losses.*

Bolland, A., Louppe, G., & Ernst, D. (2023). *Policy gradient algorithms implicitly optimize by continuation.*

Bollweg, W., Maurer, H., & Kroll, H. (1997). Numerical prediction of crystal structures by simulated annealing. In I. M. Bomze, T. Csendes, R. Horst, & P. M. Pardalos (Eds.), (pp. 253–288). Boston, MA: Springer US.

Bond-Taylor, S., Leach, A., Long, Y., & Willcocks, C. G. (2021). Deep generative modelling: A comparative review of vaes, gans, normalizing flows, energy-based and autoregressive models. *IEEE transactions on pattern analysis and machine intelligence*.

Boureau, Y.-L., & Dayan, P. (2011). Opponency revisited: competition and cooperation between dopamine and serotonin. *Neuropsychopharmacology*, *36*(1), 74–97.

Braun, D., Rosenberg, A. M., Rabaniam, E., Haruvi, R., Malamud, D., Barbara, R., ... Kawashima, T. (2024). High-resolution tracking of unconfined zebrafish behavior reveals stimulatory and anxiolytic effects of psilocybin. *Molecular Psychiatry*, 1–17.

Bredenberg, C., & Savin, C. (2023). *Desiderata for normative models of synaptic plasticity.*

Brouwer, A., & Carhart-Harris, R. L. (2021). Pivotal mental states. *Journal of Psychopharmacology*, *35*(4), 319–352.

Buesing, L., Bill, J., Nessler, B., & Maass, W. (2011, 11). Neural dynamics as sampling: A model for stochastic computation in recurrent networks of spiking neurons. *PLOS Computational Biology*, *7*(11), 1-22.

Calder, A. E., & Hasler, G. (2023). Towards an understanding of psychedelic-induced neuroplasticity. *Neuropsychopharmacology*, *48*(1), 104–112.

Callicott, J. H., Bertolino, A., Mattay, V. S., Langheim, F. J., Duyn, J., Coppola, R., ... Weinberger, D. R. (2000). Physiological dysfunction of the dorsolateral prefrontal cortex in schizophrenia revisited. *Cerebral cortex*, *10*(11), 1078–1092.

Cameron, L. P., Tombari, R. J., Lu, J., Pell, A. J., Hurley, Z. Q., Ehinger, Y., ... others (2021). A non-hallucinogenic psychedelic analogue with therapeutic potential. *Nature*, *589*(7842), 474–479.

Carhart-Harris, R., Chandaria, S., Erritzoe, D., Gazzaley, A., Girn, M., Kettner, H., ... others (2022). Canalization and plasticity in psychopathology. *Neuropharmacology*, 109398.

Carhart-Harris, R., Giribaldi, B., Watts, R., Baker-Jones, M., Murphy-Beiner, A., Murphy, R., ... Nutt, D. J. (2021). Trial of psilocybin versus escitalopram for depression. *New England Journal of Medicine*, *384*(15), 1402–1411.

Carhart-Harris, R. L., & Friston, K. J. (2019). Rebus and the anarchic brain: toward a unified model of the brain action of psychedelics. *Pharmacological reviews*, *71*(3), 316–344.

Carhart-Harris, R. L., Leech, R., Hellyer, P. J., Shanahan, M., Feilding, A., Tagliazucchi, E., ... Nutt, D. (2014). The entropic brain: a theory of conscious states informed by neuroimaging research with psychedelic drugs. *Frontiers in human neuroscience*, 20.

Carhart-Harris, R. L., & Nutt, D. (2017). Serotonin and brain function: a tale of two receptors. *Journal of Psychopharmacology*, *31*(9), 1091–1120.

Carleton, R. N., Mulvogue, M. K., Thibodeau, M. A., McCabe, R. E., Antony, M. M., & Asmundson, G. J. (2012). Increasingly certain about uncertainty: Intolerance of uncertainty across anxiety and depression. *Journal of anxiety disorders*, *26*(3), 468–479.

Carr, G. V., & Lucki, I. (2011). The role of serotonin receptor subtypes in treating depression: a review of animal studies. *Psychopharmacology*, *213*, 265–287.

Carter, O. L., Burr, D. C., Pettigrew, J. D., Wallis, G. M., Hasler, F., & Vollenweider, F. X. (2005). Using psilocybin to investigate the relationship between attention, working memory, and the serotonin 1a and 2a receptors. *Journal of Cognitive neuroscience*, *17*(10), 1497–1508.

Chalk, M., Marre, O., & Tkavcik, G. (2018). Toward a unified theory of efficient, predictive, and sparse coding. *Proceedings of the National Academy of Sciences*, *115*(1), 186–191.

Chen, C., Altafi, M., Corbu, M.-A., Trenk, A., van den Munkhof, H., Weineck, K., ... Ponomarenko, A. (2024). The dynamic state of a prefrontal–hypothalamic–midbrain circuit commands behavioral transitions. *Nature Neuroscience*.

Chen, J., & Luss, R. (2018). Stochastic gradient descent with biased but consistent gradient estimators. *CoRR*, *abs/1807.11880*. Retrieved from http://arxiv.org/abs/1807.11880

Chen, T., Fox, E., & Guestrin, C. (2014). Stochastic gradient hamiltonian monte carlo. In *International conference on machine learning* (pp. 1683–1691).

Clark, A. (2013, Jun). Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behav Brain Sci*, *36*(3), 181–204.

Clark, A. (2016). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press. Retrieved from https://doi.org/10.1093/acprof:oso/9780190217013.001.0001 doi: 10.1093/acprof:oso/9780190217013.001.0001

Conn, K., Milton, L. K., Huang, K., Munguba, H., Ruuska, J., Lemus, M. B., ... Foldi, C. J. (2023). Psilocybin prevents activity-based anorexia in female rats by enhancing cognitive flexibility: contributions from 5-ht1a and 5-ht2a receptor mechanisms. *bioRxiv*, 2023–12.

Crick, F. (1989). The recent excitement about neural networks. *Nature*, *337*(6203), 129–132.

Dara, S., Dhamercherla, S., Jadav, S. S., Babu, C. M., & Ahsan, M. J. (2022). Machine learning in drug discovery: a review. *Artificial Intelligence Review*, *55*(3), 1947–1999.

Daw, N. D., Kakade, S., & Dayan, P. (2002). Opponent interactions between serotonin and dopamine. *Neural networks*, *15*(4-6), 603–616.

Dayan, P., & Huys, Q. J. (2009). Serotonin in affective control. *Annual review of neuroscience*, *32*, 95–126.

Deco, G., Cruzat, J., Cabral, J., Knudsen, G. M., Carhart-Harris, R. L., Whybrow, P. C., ... Kringelbach, M. L. (2018). Whole-brain multimodal neuroimaging model using serotonin receptor maps explains non-linear functional effects of lsd. *Current biology*, *28*(19), 3065–3074.

den Ouden, H. E. M., Daw, N. D., Fernandez, G., Elshout, J. A., Rijpkema, M., Hoogman, M., ... Cools, R. (2013, Nov). Dissociable effects of dopamine and serotonin on reversal learning. *Neuron*, *80*(4), 1090–1100.

De Paula, D., Torricelli, A. S., Lopreato, M. R., Nascimento, J. O. G., & Viana, M. d. B. (2012). 5-ht2a receptor activation in the dorsolateral septum facilitates inhibitory avoidance in the elevated t-maze. *Behavioural brain research*, *226*(1), 50–55.

Depoortère, R., Auclair, A. L., Bardin, L., Colpaert, F. C., Vacher, B., & Newman-Tancredi, A. (2010). F15599, a preferential post-synaptic 5-ht1a receptor agonist: activity in models of cognition in comparison with reference 5-ht1a receptor agonists. *European Neuropsychopharmacology*, *20*(9), 641–654.

Deyama, S., & Duman, R. S. (2020). Neurotrophic mechanisms underlying the rapid and sustained antidepressant actions of ketamine. *Pharmacology Biochemistry and Behavior*, *188*, 172837.

DeYoung, C. G. (2015). Cybernetic big five theory. *Journal of research in personality*, *56*, 33–58.

DeYoung, C. G., & Krueger, R. F. (2018). A cybernetic theory of psychopathology. *Psychological Inquiry*,

29(3), 117–138.

Dhonnchadha, B. Á. N., Hascoët, M., Jolliet, P., & Bourin, M. (2003). Evidence for a 5-ht2a receptor mode of action in the anxiolytic-like properties of doi in mice. *Behavioural brain research*, *147*(1-2), 175–184.

Dolan, R. J., & Dayan, P. (2013). Goals and habits in the brain. *Neuron*, *80*, 312 - 325. Retrieved from https://api.semanticscholar.org/CorpusID:7199332

Domenech, P., & Koechlin, E. (2015). Executive control and decision-making in the prefrontal cortex. *Current opinion in behavioral sciences*, *1*, 101–106.

Dong, X., & Wu, S. (2023). *Neural sampling in hierarchical exponential-family energy-based models.*

Doss, M. K., de Wit, H., & Gallo, D. A. (2023). The acute effects of psychoactive drugs on emotional episodic memory encoding, consolidation, and retrieval. *Neuroscience & Biobehavioral Reviews*, 105188.

Doss, M. K., Madden, M. B., Gaddis, A., Nebel, M. B., Griffiths, R. R., Mathur, B. N., & Barrett, F. S. (2022). Models of psychedelic drug action: modulation of cortical-subcortical circuits. *Brain*, *145*(2), 441–456.

Doss, M. K., Povavzan, M., Rosenberg, M. D., Sepeda, N. D., Davis, A. K., Finan, P. H., ... others (2021). Psilocybin therapy increases cognitive and neural flexibility in patients with major depressive disorder. *Translational psychiatry*, *11*(1), 574.

Dourron, H. M., Nichols, C. D., Simonsson, O., Bradley, M., Carhart-Harris, R., & Hendricks, P. S. (2023). 5-meo-dmt: An atypical psychedelic with unique pharmacology, phenomenology & risk? *Psychopharmacology*, 1–23.

Doya, K., Ishii, S., Pouget, A., & Rao, R. P. (2006). *Bayesian brain: Probabilistic approaches to neural coding*. The MIT Press. Retrieved from https://doi.org/10.7551/mitpress/9780262042383.001.0001 doi: 10.7551/mitpress/9780262042383.001.0001

Echeveste, R., Aitchison, L., Hennequin, G., & Lengyel, M. (2020). Cortical-like dynamics in recurrent circuits optimized for sampling-based probabilistic inference. *Nature Neuroscience*, *23*(9), 1138–1149.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*(6870), 429–433.

Favela, L. H. (2020). Dynamical systems theory in cognitive science and neuroscience. *Philosophy Compass*, *15*(8), e12695.

Friedman, N. P., & Robbins, T. W. (2022). The role of prefrontal cortex in cognitive control and executive function. *Neuropsychopharmacology*, *47*(1), 72–89.

Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature reviews neuroscience*, *11*(2), 127–138.

Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical transactions of the Royal Society B: Biological sciences*, *364*(1521), 1211–1221.

Friston, K. J., & Stephan, K. E. (2007). Free-energy and the brain. *Synthese*, *159*, 417-458. Retrieved

from https://api.semanticscholar.org/CorpusID:264460

Fukumoto, K., Fogaça, M. V., Liu, R.-J., Duman, C. H., Li, X.-Y., Chaki, S., & Duman, R. S. (2020). Medial pfc ampa receptor and bdnf signaling are required for the rapid and sustained antidepressant-like effects of 5-ht1a receptor stimulation. *Neuropsychopharmacology*, *45*(10), 1725–1734.

Fukumoto, K., Iijima, M., Funakoshi, T., & Chaki, S. (2018). Role of 5-ht1a receptor stimulation in the medial prefrontal cortex in the sustained antidepressant effects of ketamine. *International Journal of Neuropsychopharmacology*, *21*(4), 371–381.

Giacomelli, S., Palmery, M., Romanelli, L., Cheng, C. Y., & Silvestrini, B. (1998). Lysergic acid diethylamide (lsd) is a partial agonist of d2 dopaminergic receptors and it potentiates dopamine-mediated prolactin secretion in lactotrophs in vitro. *Life Sci*, *63*(3), 215–222.

Girn, M., Rosas, F. E., Daws, R. E., Gallen, C. L., Gazzaley, A., & Carhart-Harris, R. L. (2023). A complex systems perspective on psychedelic brain action. *Trends in Cognitive Sciences*.

Grandjean, J., Buehlmann, D., Buerge, M., Sigrist, H., Seifritz, E., Vollenweider, F. X., ... Rudin, M. (2021, Jan). Psilocybin exerts distinct effects on resting state networks associated with serotonin and dopamine in mice. *Neuroimage*, *225*, 117456.

Greer, G. R., & Tolbert, R. (1990). The therapeutic use of mdma. In *Ecstasy: The clinical, pharmacological and neurotoxicological effects of the drug mdma* (pp. 21–35). Springer.

Griffiths, R. R., Hurwitz, E. S., Davis, A. K., Johnson, M. W., & Jesse, R. (2019). Survey of subjective" god encounter experiences": Comparisons among naturally occurring experiences and those occasioned by the classic psychedelics psilocybin, lsd, ayahuasca, or dmt. *PloS one*, *14*(4), e0214377.

Griffiths, R. R., Johnson, M. W., Carducci, M. A., Umbricht, A., Richards, W. A., Richards, B. D., ... Klinedinst, M. A. (2016). Psilocybin produces substantial and sustained decreases in depression and anxiety in patients with life-threatening cancer: A randomized double-blind trial. *Journal of psychopharmacology*, *30*(12), 1181–1197.

Griffiths, R. R., Johnson, M. W., Richards, W. A., Richards, B. D., McCann, U., & Jesse, R. (2011). Psilocybin occasioned mystical-type experiences: immediate and persisting dose-related effects. *Psychopharmacology*, *218*, 649–665.

Grossberg, S. (1987). Competitive learning: From interactive activation to adaptive resonance. *Cognitive Science*, *11*(1), 23-63.

Gómez-Emilsson, A. (2020). *5-meo-dmt vs. n,n-dmt: The 9 lenses.* Retrieved from https://www.qri.org/blog/5meo-vs-dmt

Gómez-Emilsson, A. (2021). *Healing trauma with neural annealing.* Retrieved from https://www.qri.org/blog/Neural-Annealing

Halberstadt, A. L., Chatha, M., Klein, A. K., Wallach, J., & Brandt, S. D. (2020). Correlation between the potency of hallucinogens in the mouse head-twitch response assay and their behavioral and subjective effects in other species. *Neuropharmacology*, *167*, 107933.

Halberstadt, A. L., Koedood, L., Powell, S. B., & Geyer, M. A. (2011). Differential contributions of serotonin receptors to the behavioral effects of indoleamine hallucinogens in mice. *Journal of psychopharma-*

cology, *25*(11), 1548–1561.

Halberstadt, A. L., Nichols, D. E., & Geyer, M. A. (2012). Behavioral effects of α, α, β, β-tetradeutero-5-meo-dmt in rats: comparison with 5-meo-dmt administered in combination with a monoamine oxidase inhibitor. *Psychopharmacology*, *221*, 709–718.

Hamilton, W. R. S. (2000). On a general method of expressing the paths of light, and of the planets, by the coefficients of a characteristic function by..

Harkin, E. F., Grossman, C. D., Cohen, J. Y., Béïque, J.-C., & Naud, R. (2023). Serotonin predictively encodes value. *bioRxiv*.

Hayashi, K., Nakao, K., & Nakamura, K. (2015, Apr). Appetitive and aversive information coding in the primate dorsal raphénucleus. *J Neurosci*, *35*(15), 6195–6208.

Hennequin, G., Aitchison, L., & Lengyel, M. (2014). Fast sampling-based inference in balanced neuronal networks. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, & K. Weinberger (Eds.), *Advances in neural information processing systems* (Vol. 27). Curran Associates, Inc.

Hesselgrave, N., Troppoli, T. A., Wulff, A. B., Cole, A. B., & Thompson, S. M. (2021). Harnessing psilocybin: antidepressant-like behavioral and synaptic actions of psilocybin are independent of 5-ht2r activation in mice. *Proceedings of the National Academy of Sciences*, *118*(17), e2022489118.

Hinton, G. E. (1999). Products of experts. In *Artificial neural networks, 1999. icann 99. ninth international conference on (conf. publ. no. 470)* (Vol. 1, pp. 1–6).

Hipólito, I., Mago, J., Rosas, F. E., & Carhart-Harris, R. (2023). Pattern breaking: a complex systems approach to psychedelic medicine. *Neuroscience of Consciousness*, *2023*(1), niad017.

Hoyer, P. O., & Hyvärinen, A. (2002). Interpreting neural response variability as monte carlo sampling of the posterior. In *Proceedings of the 15th international conference on neural information processing systems* (p. 293–300). Cambridge, MA, USA: MIT Press.

Huang, Y., & Rao, R. P. N. (2011). Predictive coding. *WIREs Cognitive Science*, *2*(5), 580-593. Retrieved from https://wires.onlinelibrary.wiley.com/doi/abs/10.1002/wcs.142 doi: https://doi.org/10.1002/wcs.142

Inagaki, H. K., Fontolan, L., Romani, S., & Svoboda, K. (2019). Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature*, *566*(7743), 212–217.

Ionescu, T. (2012). Exploring the nature of cognitive flexibility. *New ideas in psychology*, *30*(2), 190–200.

Islam, M., Moriguchi, S., Tagashira, H., & Fukunaga, K. (2014). Rivastigmine improves hippocampal neurogenesis and depression-like behaviors via 5-ht1a receptor stimulation in olfactory bulbectomized mice. *Neuroscience*, *272*, 116–130.

Johnson, M. (2019). Neural annealing: Toward a neural theory of everything. *URL https://opentheory.net/2019/11/neural-annealing-toward-a-neural-theory-of-ev erything*.

Johnson, M. W., & Griffiths, R. R. (2017). Potential therapeutic effects of psilocybin. *Neurotherapeutics*, *14*, 734–740.

Johnson, M. W., Hendricks, P. S., Barrett, F. S., & Griffiths, R. R. (2019). Classic psychedelics: An integrative review of epidemiology, therapeutics, mystical experience, and brain network function.

Pharmacology & therapeutics, *197*, 83–102.

Juliani, A., Safron, A., & Kanai, R. (2024). Deep canals: A deep learning approach to refining the canalization theory of psychopathology. *Neuroscience of Consciousness*.

Kanai, R., Komura, Y., Shipp, S., & Friston, K. (2015a). Cerebral hierarchies: predictive processing, precision and the pulvinar. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *370*(1668), 20140169.

Kanai, R., Komura, Y., Shipp, S., & Friston, K. (2015b, May). Cerebral hierarchies: predictive processing, precision and the pulvinar. *Philos Trans R Soc Lond B Biol Sci*, *370*(1668).

Keller, G. B., & Mrsic-Flogel, T. D. (2018). Predictive processing: a canonical cortical computation. *Neuron*, *100*(2), 424–435.

Keshmiri, S. (2020, Aug). Entropy and the brain: An overview. *Entropy (Basel)*, *22*(9).

Khona, M., & Fiete, I. R. (2022). Attractor and integrator networks in the brain. *Nature Reviews Neuroscience*, *23*(12), 744–766.

Kingma, D. P., & Welling, M. (2022). *Auto-encoding variational bayes.*

King-Smith, E., Berritt, S., Bernier, L., Hou, X., Klug-McLeod, J. L., Mustakis, J., ... others (2024). Probing the chemical 'reactome'with high-throughput experimentation data. *Nature Chemistry*, 1–11.

Kirchhoff, M., Parr, T., Palacios, E., Friston, K., & Kiverstein, J. (2018). The markov blankets of life: autonomy, active inference and the free energy principle. *Journal of The royal society interface*, *15*(138), 20170792.

Knill, D. C., & Pouget, A. (2004, Dec). The bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci*, *27*(12), 712–719.

Knill, D. C., & Richards, W. (Eds.). (1996). *Perception as bayesian inference*. USA: Cambridge University Press.

Kometer, M., Pokorny, T., Seifritz, E., & Volleinweider, F. X. (2015). Psilocybin-induced spiritual experiences and insightfulness are associated with synchronization of neuronal oscillations. *Psychopharmacology*, *232*, 3663–3676.

Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, *427*(6971), 244–247.

Kukavcka, J., Golkov, V., & Cremers, D. (2017). Regularization for deep learning: A taxonomy. *arXiv preprint arXiv:1710.10686*.

Laukkonen, R. E., Webb, M., Salvi, C., Tangen, J. M., Slagter, H. A., & Schooler, J. W. (2023). Insight and the selection of ideas. *Neuroscience & Biobehavioral Reviews*, 105363.

LeCun, Y., Chopra, S., Hadsell, R., Ranzato, A., & Huang, F. J. (2006). A tutorial on energy-based learning.. Retrieved from https://api.semanticscholar.org/CorpusID:8531544

Letheby, C. (2021). *Philosophy of psychedelics*. Oxford University Press.

Lewis, V., Bonniwell, E. M., Lanham, J. K., Ghaffari, A., Sheshbaradaran, H., Cao, A. B., ... others (2023). A non-hallucinogenic lsd analog with therapeutic potential for mood disorders. *Cell reports*, *42*(3).

Li, N., Daie, K., Svoboda, K., & Druckmann, S. (2016, Apr). Robust neuronal dynamics in premotor cortex during motor planning. *Nature*, *532*(7600), 459–464.

Liechti, M. E., Saur, M. R., Gamma, A., Hell, D., & Vollenweider, F. X. (2000). Psychological and physiological effects of mdma ("ecstasy") after pretreatment with the 5-ht2 antagonist ketanserin in healthy humans. *Neuropsychopharmacology*, *23*(4), 396–404.

Lillicrap, T. P., Santoro, A., Marris, L., Akerman, C. J., & Hinton, G. (2020). Backpropagation and the brain. *Nature Reviews Neuroscience*, *21*(6), 335–346.

Lotter, W., Kreiman, G., & Cox, D. (2016). Deep predictive coding networks for video prediction and unsupervised learning. *arXiv preprint arXiv:1605.08104*.

Luan, L. X., Eckernäs, E., Ashton, M., Rosas, F. E., Uthaug, M. V., Bartha, A., ... others (2023). Psychological and physiological effects of extended dmt. *Journal of Psychopharmacology*, 02698811231196877.

Luppi, A. I., Girn, M., Rosas, F. E., Timmermann, C., Roseman, L., Erritzoe, D., ... others (2024). A role for the serotonin 2a receptor in the expansion and functioning of human transmodal cortex. *Brain*, *147*(1), 56–80.

Luscombe, G. P., Martin, K. F., Hutchins, L. J., Gosden, J., & Heal, D. J. (1993). Mediation of the antidepressant-like effect of 8-oh-dpat in mice by postsynaptic 5-ht1a receptors. *British journal of pharmacology*, *108*(3), 669–677.

Ly, C., Greb, A. C., Cameron, L. P., Wong, J. M., Barragan, E. V., Wilson, P. C., ... others (2018). Psychedelics promote structural and functional neural plasticity. *Cell reports*, *23*(11), 3170–3182.

Lyle, C., Zheng, Z., Nikishin, E., Pires, B. A., Pascanu, R., & Dabney, W. (2023). Understanding plasticity in neural networks. *arXiv preprint arXiv:2303.01486*.

Lyu, J., Kapolka, N., Gumpper, R., Alon, A., Wang, L., Jain, M. K., ... others (2023). Alphafold2 structures template ligand discovery. *bioRxiv*, 2023–12.

Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *nature*, *503*(7474), 78–84.

Massé, F., Hascoët, M., Dailly, E., & Bourin, M. (2006). Effect of noradrenergic system on the anxiolytic-like effect of doi (5-ht2a/2c agonists) in the four-plate test. *Psychopharmacology*, *183*, 471–481.

Masse, F., Petit-Demouliere, B., Dubois, I., Hascoët, M., & Bourin, M. (2008). Anxiolytic-like effects of doi microinjections into the hippocampus (but not the amygdala nor the pag) in the mice four plates test. *Behavioural brain research*, *188*(2), 291–297.

Matias, S., Lottem, E., Dugué, G. P., & Mainen, Z. F. (2017). Activity patterns of serotonin neurons underlying cognitive flexibility. *Elife*, *6*, e20552.

McGovern, H., Grimmer, H., Doss, M., Hutchinson, B., Timmermann, C., Lyon, A., ... Laukkonen, R. (2023). The power of insight: How psychedelics solicit false beliefs. *PsyArXiv preprint: https://psyarxiv. com/97gjw*.

Mei, J., Xiao, C., Dai, B., Li, L., Szepesvari, C., & Schuurmans, D. (2020). Escaping the gravitational pull of softmax. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, & H. Lin (Eds.), *Advances in neural*

information processing systems (Vol. 33, pp. 21130–21140). Curran Associates, Inc.

Mermillod, M., Bugaiska, A., & Bonin, P. (2013). *The stability-plasticity dilemma: Investigating the continuum from catastrophic forgetting to age-limited learning effects* (Vol. 4). Frontiers Media SA.

Metropolis, N. (1987). The beginning of the monte carlo method. *Los Alamos Science*, *Special Issue 15*, 125–130.

Miles, R., Poncer, J.-C., Fricker, D., & Leinekugel, X. (2005, Oct). The birth (and adolescence) of ltp. *J Physiol*, *568*(Pt 1), 1–2.

Millan, M. J. (2005). Serotonin 5-ht2c receptors as a target for the treatment of depressive and anxious states: focus on novel therapeutic strategies. *Therapies*, *60*(5), 441–460.

Miller, E. K., Freedman, D. J., & Wallis, J. D. (2002). The prefrontal cortex: categories, concepts and cognition. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *357*(1424), 1123–1136.

Mitchell, J. M., Ot'alora G, M., van der Kolk, B., Shannon, S., Bogenschutz, M., Gelfand, Y., ... others (2023). Mdma-assisted therapy for moderate to severe ptsd: a randomized, placebo-controlled phase 3 trial. *Nature Medicine*, 1–8.

Moliner, R., Girych, M., Brunello, C. A., Kovaleva, V., Biojone, C., Enkavi, G., ... others (2023). Psychedelics promote plasticity by directly binding to bdnf receptor trkb. *Nature Neuroscience*, *26*(6), 1032–1041.

Muller, T. H., Butler, J. L., Veselic, S., Miranda, B., Wallis, J. D., Dayan, P., ... Kennerley, S. W. (2024). Distributional reinforcement learning in prefrontal cortex. *Nature Neuroscience*, 1–6.

Murnane, K. S. (2019). Serotonin 2a receptors are a stress response system: implications for post-traumatic stress disorder. *Behavioural pharmacology*, *30*(2-), 151.

Muysers, H., Chen, H.-L., Hahn, J., Folschweiller, S., Sigurdsson, T., Sauer, J.-F., & Bartos, M. (2024). A persistent prefrontal reference frame across time and task rules. *Nature Communications*, *15*(1), 2115.

Nardou, R., Sawyer, E., Song, Y. J., Wilkinson, M., Padovan-Hernandez, Y., de Deus, J. L., ... Dölen, G. (2023). Psychedelics reopen the social reward learning critical period. *Nature*, *618*(7966), 790–798.

Neal, R. (2012, 06). Mcmc using hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*.

Ngiam, J., Chen, Z., Koh, P. W., & Ng, A. Y. (2011). Learning deep energy models. In *Proceedings of the 28th international conference on machine learning (icml-11)* (pp. 1105–1112).

Nichols, D. E., & Nichols, C. D. (2008). Serotonin receptors. *Chemical reviews*, *108*(5), 1614–1641.

Odland, A. U., Kristensen, J. L., & Andreasen, J. T. (2021). The selective 5-ht2a receptor agonist 25cn-nboh does not affect reversal learning in mice. *Behavioural Pharmacology*, *32*(5), 448–452.

Okazawa, G., Hatch, C. E., Mancoo, A., Machens, C. K., & Kiani, R. (2021). Representational geometry of perceptual decisions in the monkey parietal cortex. *Cell*, *184*(14), 3748–3761.

Olson, D. E. (2018). Psychoplastogens: a promising class of plasticity-promoting neurotherapeutics. *Journal of experimental neuroscience*, *12*, 1179069518800508.

Orbán, G., Berkes, P., Fiser, J., & Lengyel, M. (2016, Oct). Neural variability and sampling-based proba-

bilistic representations in the visual cortex. *Neuron*, *92*(2), 530–543.

Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. (2018). Continual lifelong learning with neural networks: A review. *Neural networks : the official journal of the International Neural Network Society*, *113*, 54-71. Retrieved from https://api.semanticscholar.org/CorpusID:73497737

Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. (2019). Continual lifelong learning with neural networks: A review. *Neural networks*, *113*, 54–71.

Parr, T., Pezzulo, G., & Friston, K. J. (2022). *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. The MIT Press. Retrieved from https://doi.org/10.7551/mitpress/12441.001.0001 doi: 10.7551/mitpress/12441.001.0001

Pizzagalli, D. A., & Roberts, A. C. (2022). Prefrontal cortex and depression. *Neuropsychopharmacology*, *47*(1), 225–246.

Pokorny, T., Preller, K. H., Kraehenmann, R., & Vollenweider, F. X. (2016). Modulatory effect of the 5-ht1a agonist buspirone and the mixed non-hallucinogenic 5-ht1a/2a agonist ergotamine on psilocybin-induced psychedelic experience. *European Neuropsychopharmacology*, *26*(4), 756–766.

Preller, K. H., Herdener, M., Pokorny, T., Planzer, A., Kraehenmann, R., Stämpfli, P., ... Vollenweider, F. X. (2017). The fabric of meaning and subjective effects in lsd-induced states depend on serotonin 2a receptor activation. *Current Biology*, *27*(3), 451–457.

Puglisi-Allegra, S., & Andolina, D. (2015). Serotonin and stress coping. *Behavioural brain research*, *277*, 58–67.

Puig, M. V., & Gulledge, A. T. (2011). Serotonin and prefrontal cortex function: neurons, networks, and circuits. *Molecular neurobiology*, *44*, 449–464.

Puigseslloses, P., Nadal-Gratacós, N., Ketsela, G., Weiss, N., Berzosa, X., Estrada-Tejedor, R., ... others (2024). Structure-activity relationships of serotonergic 5-meo-dmt derivatives: insights into psychoactive and thermoregulatory properties. *Molecular Psychiatry*, 1–13.

Qi, Y., & Gong, P. (2022). Fractional neural sampling as a theory of spatiotemporal probabilistic computations in neural circuits. *Nature Communications*, *13*(1), 4572.

Qu, Y., Chang, L., Ma, L., Wan, X., & Hashimoto, K. (2023). Rapid antidepressant-like effect of non-hallucinogenic psychedelic analog lisuride, but not hallucinogenic psychedelic doi, in lipopolysaccharide-treated mice. *Pharmacology Biochemistry and Behavior*, *222*, 173500.

Ranzato, M., Poultney, C., Chopra, S., & Cun, Y. (2006). Efficient learning of sparse representations with an energy-based model. *Advances in neural information processing systems*, *19*.

Rao, R. P., Olshausen, B. A., & Lewicki, M. S. (2002). *Probabilistic Models of the Brain: Perception and Neural Function*. The MIT Press.

Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79–87.

Reckweg, J. T., Uthaug, M. V., Szabo, A., Davis, A. K., Lancelotta, R., Mason, N. L., & Ramaekers, J. G. (2022). The clinical pharmacology and potential therapeutic applications of 5-methoxy-n, n-dimethyltryptamine (5-meo-dmt). *Journal of Neurochemistry*, *162*(1), 128–146.

Reckweg, J. T., van Leeuwen, C. J., Henquet, C., Van Amelsvoort, T., Theunissen, E. L., Mason, N. L., ... Ramaekers, J. G. (2023). A phase 1/2 trial to assess safety and efficacy of a vaporized 5-methoxy-n, n-dimethyltryptamine formulation (gh001) in patients with treatment-resistant depression. *Frontiers in Psychiatry*, *14*, 1133414.

Reissig, C., Eckler, J., Rabin, R., & Winter, J. (2005). The 5-ht 1a receptor and the stimulus effects of lsd in the rat. *Psychopharmacology*, *182*, 197–204.

Rezende, D. J., Mohamed, S., & Wierstra, D. (2014). *Stochastic backpropagation and approximate inference in deep generative models.*

Richards, B. A., Lillicrap, T. P., Beaudoin, P., Bengio, Y., Bogacz, R., Christensen, A., ... others (2019a). A deep learning framework for neuroscience. *Nature neuroscience*, *22*(11), 1761–1770.

Richards, B. A., Lillicrap, T. P., Beaudoin, P., Bengio, Y., Bogacz, R., Christensen, A., ... Kording, K. P. (2019b). A deep learning framework for neuroscience. *Nature Neuroscience*, *22*(11), 1761–1770.

Rickli, A., Luethi, D., Reinisch, J., Buchy, D., Hoener, M. C., & Liechti, M. E. (2015). Receptor interaction profiles of novel n-2-methoxybenzyl (nbome) derivatives of 2, 5-dimethoxy-substituted phenethylamines (2c drugs). *Neuropharmacology*, *99*, 546–553.

Rickli, A., Moning, O. D., Hoener, M. C., & Liechti, M. E. (2016). Receptor interaction profiles of novel psychoactive tryptamines compared with classic hallucinogens. *European Neuropsychopharmacology*, *26*(8), 1327–1337.

Riga, M. S., Bortolozzi, A., Campa, L., Artigas, F., & Celada, P. (2016). The serotonergic hallucinogen 5-methoxy-n, n-dimethyltryptamine disrupts cortical activity in a regionally-selective manner via 5-ht1a and 5-ht2a receptors. *Neuropharmacology*, *101*, 370–378.

Riga, M. S., Lladó-Pelfort, L., Artigas, F., & Celada, P. (2018). The serotonin hallucinogen 5-meo-dmt alters cortico-thalamic activity in freely moving mice: Regionally-selective involvement of 5-ht1a and 5-ht2a receptors. *Neuropharmacology*, *142*, 219–230.

Riga, M. S., Soria, G., Tudela, R., Artigas, F., & Celada, P. (2014). The natural hallucinogen 5-meo-dmt, component of ayahuasca, disrupts cortical function in rats: reversal by antipsychotic drugs. *International Journal of Neuropsychopharmacology*, *17*(8), 1269–1282.

Ruf, B., & Bhagwagar, Z. (2009). The 5-ht1b receptor: a novel target for the pathophysiology of depression. *Current drug targets*, *10*(11), 1118–1138.

Ruffini, G., Lopez-Sola, E., Vohryzek, J., & Sanchez-Todo, R. (2024). Neural geometrodynamics, complexity, and plasticity: a psychedelics perspective. *Entropy*, *26*(1), 90.

Sadeh, S., Clopath, C., Palmer, S. E., Frank, M. J., & Ziv, Y. (2022). Contribution of behavioural variability to representational drift. *eLife*, *11*, e77907.

Safron, A. (2020). On the varieties of conscious experiences: Altered beliefs under psychedelics (albus).

Safron, A., & Sheikhbahaee, Z. (2021). Dream to explore: 5-ht2a as adaptive temperature parameter for sophisticated affective inference. In *Joint european conference on machine learning and knowledge discovery in databases* (pp. 799–809).

Sakashita, Y., Abe, K., Katagiri, N., Kambe, T., Saitoh, T., Utsunomiya, I., ... Taguchi, K. (2015). Effect

of psilocin on extracellular dopamine and serotonin levels in the mesoaccumbens and mesocortical pathway in awake rats. *Biol Pharm Bull*, *38*(1), 134–138.

Schechter, M. D. (1998). Candyflipping': synergistic discriminative effect of lsd and mdma. *European journal of pharmacology*, *341*(2-3), 131–134.

Schreiber, R., & De Vry, J. (1993). Neuronal circuits involved in the anxiolytic effects of the 5-ht1a receptor agonists 8-oh-dpat, ipsapirone and buspirone in the rat. *European journal of pharmacology*, *249*(3), 341–351.

Schreiber, R., & Newman-Tancredi, A. (2014). Improving cognition in schizophrenia with antipsychotics that elicit neurogenesis through 5-ht1a receptor activation. *Neurobiology of learning and memory*, *110*, 72–80.

Seo, C., Guru, A., Jin, M., Ito, B., Sleezer, B. J., Ho, Y.-Y., ... others (2019). Intense threat switches dorsal raphe serotonin neurons to a paradoxical operational mode. *Science*, *363*(6426), 538–542.

Sessa, B. (2018). The 21st century psychedelic renaissance: heroic steps forward on the back of an elephant. *Psychopharmacology*, *235*(2), 551–560.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell Syst. Tech. J.*, *27*, 623-656. Retrieved from https://api.semanticscholar.org/CorpusID:55379485

Shao, L.-X., Liao, C., Gregg, I., Davoudian, P. A., Savalia, N. K., Delagarza, K., & Kwan, A. C. (2021, 2024/01/10). Psilocybin induces rapid and persistent growth of dendritic spines in frontal cortex in vivo. *Neuron*, *109*(16), 2535–2544.e4.

Shine, J. M., O'Callaghan, C., Walpola, I. C., Wainstein, G., Taylor, N., Aru, J., ... John, Y. J. (2022). Understanding the effects of serotonin in the brain through its role in the gastrointestinal tract. *Brain*, *145*(9), 2967–2981.

Singh, S., Botvinnik, A., Shahar, O., Wolf, G., Yakobi, C., Saban, M., ... Lifschytz, T. (2023). Effect of psilocybin on marble burying in icr mice: role of 5-ht1a receptors and implications for the treatment of obsessive-compulsive disorder. *Translational Psychiatry*, *13*(1), 164.

Singleton, S. P., Luppi, A. I., Carhart-Harris, R. L., Cruzat, J., Roseman, L., Nutt, D. J., ... Kuceyeski, A. (2022). Receptor-informed network control theory links lsd and psilocybin to a flattening of the brain's control energy landscape. *Nature communications*, *13*(1), 5812.

Smith, R., Badcock, P., & Friston, K. J. (2021). Recent advances in the application of predictive coding and active inference models within clinical neuroscience. *Psychiatry and Clinical Neurosciences*, *75*(1), 3–13.

Smith, R., Kirlic, N., Stewart, J. L., Touthang, J., Kuplicki, R., Khalsa, S. S., ... Aupperle, R. L. (2021). Greater decision uncertainty characterizes a transdiagnostic patient sample during approach-avoidance conflict: a computational modelling approach. *Journal of Psychiatry and Neuroscience*, *46*(1), E74–E87.

Smith, R., Varshney, L. R., Nagayama, S., Kazama, M., Kitagawa, T., & Ishikawa, Y. (2022). A computational neuroscience perspective on subjective wellbeing within the active inference framework. *International Journal of Wellbeing*, *12*(4).

Strassman, R. J. (1995). Human psychopharmacology of n, n-dimethyltryptamine. *Behavioural brain research*, *73*(1-2), 121–124.

Straumann, I., Ley, L., Holze, F., Becker, A. M., Klaiber, A., Wey, K., ... Liechti, M. E. (2023). Acute effects of mdma and lsd co-administration in a double-blind placebo-controlled study in healthy participants. *Neuropsychopharmacology*, *48*(13), 1840–1848.

Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences*, *10*(7), 309–318.

Thompson, M., Callaghan, P. D., Hunt, G. E., Cornish, J. L., & McGregor, I. S. (2007). A role for oxytocin and 5-ht1a receptors in the prosocial effects of 3, 4 methylenedioxymethamphetamine ("ecstasy"). *Neuroscience*, *146*(2), 509–514.

Thrun, S. (1998). Lifelong learning algorithms. In S. Thrun & L. Pratt (Eds.), *Learning to learn* (pp. 181–209). Boston, MA: Springer US. Retrieved from https://doi.org/10.1007/978-1-4615-5529-2_8 doi: 10.1007/978-1-4615-5529-2_8

Timmermann, C., Zeifman, R. J., Erritzoe, D., Nutt, D. J., & Carhart-Harris, R. L. (2024). Effects of dmt on mental health outcomes in healthy volunteers. *Scientific Reports*, *14*(1), 3097.

Topolinski, S., & Reber, R. (2010). Gaining insight into the "aha" experience. *Current Directions in Psychological Science*, *19*(6), 402–405.

Torrado Pacheco, A., Olson, R. J., Garza, G., & Moghaddam, B. (2023). Acute psilocybin enhances cognitive flexibility in rats. *bioRxiv*, 2023–01.

Uthaug, M., Lancelotta, R., Van Oorsouw, K., Kuypers, K., Mason, N., Rak, J., ... others (2019). A single inhalation of vapor from dried toad secretion containing 5-methoxy-n, n-dimethyltryptamine (5-meo-dmt) in a naturalistic setting is related to sustained enhancement of satisfaction with life, mindfulness-related capacities, and a decrement of psychopathological symptoms. *Psychopharmacology*, *236*, 2653–2666.

Vargas, M. V., Dunlap, L. E., Dong, C., Carter, S. J., Tombari, R. J., Jami, S. A., ... others (2023). Psychedelics promote neuroplasticity through the activation of intracellular 5-ht2a receptors. *Science*, *379*(6633), 700–706.

Volpi-Abadie, J., Kaye, A. M., & Kaye, A. D. (2013). Serotonin syndrome. *Ochsner Journal*, *13*(4), 533–540.

Walsh, K. S., McGovern, D. P., Clark, A., & O'Connell, R. G. (2020). Evaluating the neurophysiological evidence for predictive processing as a model of perception. *Annals of the new York Academy of Sciences*, *1464*(1), 242–268.

Wang, J. X., Kurth-Nelson, Z., Kumaran, D., Tirumala, D., Soyer, H., Leibo, J. Z., ... Botvinick, M. (2018). Prefrontal cortex as a meta-reinforcement learning system. *Nature Neuroscience*, *21*(6), 860–868.

Wang, S., Falcone, R., Richmond, B., & Averbeck, B. B. (2023). Attractor dynamics reflect decision confidence in macaque prefrontal cortex. *Nature Neuroscience*, 1–11.

Welling, M., & Teh, Y. W. (2011). Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th international conference on machine learning (icml-11)* (pp. 681–688).

Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, *8*(3), 229–256.

Wilmes, K. A., & Clopath, C. (2023). Dendrites help mitigate the plasticity-stability dilemma. *Scientific Reports*, *13*(1), 6543.

Winter, J., Filipink, R., Timineri, D., Helsley, S., & Rabin, R. (2000). The paradox of 5-methoxy-n, n-dimethyltryptamine: an indoleamine hallucinogen that induces stimulus control via 5-ht1a receptors. *Pharmacology Biochemistry and Behavior*, *65*(1), 75–82.

Winter, J., Rice, K., Amorosi, D., & Rabin, R. (2007). Psilocybin-induced stimulus control in the rat. *Pharmacology Biochemistry and Behavior*, *87*(4), 472–480.

Wood, J., Kim, Y., & Moghaddam, B. (2012). Disruption of prefrontal cortex large scale neuronal activity by different classes of psychotomimetic drugs. *Journal of Neuroscience*, *32*(9), 3022–3031.

Yaden, D. B., & Griffiths, R. R. (2020). The subjective effects of psychedelics are necessary for their enduring therapeutic effects. *ACS Pharmacology & Translational Science*, *4*(2), 568–572.

Zeifman, R. J., Kettner, H., Pagni, B. A., Mallard, A., Roberts, D. E., Erritzoe, D., . . . Carhart-Harris, R. L. (2023). Co-use of mdma with psilocybin/lsd may buffer against challenging experiences and enhance positive experiences. *Scientific Reports*, *13*(1), 13645.

Zeifman, R. J., Spriggs, M. J., Kettner, H., Lyons, T., Rosas, F., Mediano, P. A., . . . Carhart-Harris, R. (2022). From relaxed beliefs under psychedelics (rebus) to revised beliefs after psychedelics (rebas): Preliminary development of the relaxed beliefs questionnaire (reb-q).

Zeifman, R. J., Wagner, A. C., Monson, C. M., & Carhart-Harris, R. L. (2023). How does psilocybin therapy work? an exploration of experiential avoidance as a putative mechanism of change. *Journal of Affective Disorders*.

Zhong, W., Li, Y., Feng, Q., & Luo, M. (2017, Sep). Learning and stress shape the reward response patterns of serotonin neurons. *J Neurosci*, *37*(37), 8863–8875.

A Notation

Table 1: Notation

Indices	
t	outer-loop iteration index $t \in \{0, ..., T-1\}$ (neuromodulation)
T	total number of outer-loop iterations.
k	inner-loop iteration index $k \in \{0, ..., k-1\}$ (sampling)
K	total number of inner-loop iterations (sampling).
Parameters	
$\theta \in \Theta \equiv \mathbb{R}^d$	parameters of the (generative) model $p_\theta(\mathbf{x}, \mathbf{z}) = p_\theta(\mathbf{x}\|\mathbf{z})p_\theta(\mathbf{z})$
$\phi \in \Phi \equiv \mathbb{R}^d$	parameters of the posterior recognition model $p_\phi(\mathbf{z}\|\mathbf{x})$
$\mathbf{z} \in \mathcal{Z} \equiv \mathbb{R}^n$	neural response
$\mathbf{x} \in \mathcal{X} \equiv \mathbb{R}^N$	stimuli
α	hyper-parameters of the neuromodulatory map/neuromodulatory strength/ learning-rates
Maps	
$E^* : \mathcal{Z} \to \mathbb{R}$	target energy function (ground truth)
$E : \mathcal{Z} \to \mathbb{R}$	learning energy function
$\tilde{E} : \mathcal{Z} \to \mathbb{R}$	auxiliary energy function for neuromodulation (exploration)
$\mathcal{L} \equiv D_{\mathrm{KL}} : \mathcal{P}(\mathcal{Z}) \times \mathcal{P}(\mathcal{Z}) \to \mathbb{R}$	KL-divergence between two distributions

B Theoretical and computational model

In this section, we give additional details concerning the theoretical framework in which the computation model operates and the theory surrounding it.

B.1 Inference and learning in latent-variable models

Given stimulus $\mathbf{x}$ we seek a good neural representation, denoted with $\mathbf{z}$, such that the neural response elicited encodes maximal information about the stimulus. We formalize inference and learning in latent variable models as an optimization problem using the variational inference framework (Kingma & Welling, 2022). We assume the stimuli should be decodable from neural responses, i.e. there exists a generative distribution $\mathbf{x} \sim p_\theta(\mathbf{x}|\mathbf{z})$ (typically obtained from top-down processing) represented by means of synaptic connection weights θ (subject to learning under Hebbian plasticity) that takes a neural response and decodes the stimuli. We also assume there exists a posterior recognition distribution over latent neural responses $\mathbf{z}$, which is used to encode the sensory stimuli $p_\phi(\mathbf{z}|\mathbf{x})$, and it has a synaptic representation captured by ϕ.

Problem setting To fit a generative distribution to incoming sensory stimuli using maximum likelihood, we arrive at the minimization objective over the synaptic weights θ

$$D_{\mathrm{KL}}\left(p(\mathbf{x}), p_\theta(\mathbf{x})\right) = \int \log\left(\frac{p(\mathbf{x})}{p_\theta(\mathbf{x})}\right) p(\mathbf{x}) d\mathbf{x} \tag{2}$$

The objective is minimized when $p_\theta(\mathbf{x})$ matches $p(\mathbf{x})$, in which case the data distribution is perfectly modelled. Using Bayes' Theorem, we may obtain the distribution over latent neural responses $\mathbf{z}$ when a stimulus $\mathbf{x}$ is received

$$p_\theta(\mathbf{z}|\mathbf{x}) = \frac{p_\theta(\mathbf{z}) p_\theta(\mathbf{x}|\mathbf{z})}{\int p_\theta(\mathbf{z}) p_\theta(\mathbf{x}|\mathbf{z}) d\mathbf{x} d\mathbf{z}}$$

For complex, nonlinear models, the high-dimensional integral in the denominator is generally computationally or analytically intractable.

Evidence lower bound (ELBO) One approach is estimate the recognition map $p_\theta(\mathbf{z}|\mathbf{x})$ with a parametric model $p_\phi(\mathbf{z}|\mathbf{x})$ and use it to encode the stimulus by minimizing the following KL-divergence over the parameters ϕ

$$\mathbb{E}_{p(\mathbf{x})}[D_{\mathrm{KL}}\left(p_\phi(\mathbf{z}|\mathbf{x}), p_\theta(\mathbf{z}|\mathbf{x})\right)] = \int \log\left(\frac{p_\phi(\mathbf{z}|\mathbf{x})}{p_\theta(\mathbf{z}|\mathbf{x})}\right) p_\phi(\mathbf{z}|\mathbf{x}) p(\mathbf{x}) d\mathbf{z} d\mathbf{x} \tag{3}$$

Adding Eq.2 and Eq.3, we obtain an upper bound on the original objective of interest

$$
\begin{aligned}
D_{\mathrm{KL}}\left(p(\mathbf{x}), p_\theta(\mathbf{x})\right) &\leq D_{\mathrm{KL}}\left(p(\mathbf{x}), p_\theta(\mathbf{x})\right) + \mathbb{E}_{p(\mathbf{x})}[D_{\mathrm{KL}}\left(p_\phi(\mathbf{z}|\mathbf{x}), p_\theta(\mathbf{z}|\mathbf{x})\right)] \\
&= \int \log p(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} - \int \log p_\theta(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} + \int \log p_\phi(\mathbf{z}|\mathbf{x}) p_\phi(\mathbf{z}|\mathbf{x}) p(\mathbf{x}) d\mathbf{z} d\mathbf{x} \\
&\quad - \int \log p_\theta(\mathbf{z}|\mathbf{x}) p_\phi(\mathbf{z}|\mathbf{x}) p(\mathbf{x}) d\mathbf{z} d\mathbf{x} \\
&= \int \log p(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} - \int \log p_\theta(\mathbf{x}) p(\mathbf{x}, \mathbf{z}) d\mathbf{x} d\mathbf{z} + \int \log p_\phi(\mathbf{z}|\mathbf{x}) p_\phi(\mathbf{x}, \mathbf{z}) d\mathbf{z} d\mathbf{x} \\
&\quad - \int \log p_\theta(\mathbf{z}|\mathbf{x}) p_\phi(\mathbf{x}, \mathbf{z}) d\mathbf{z} d\mathbf{x} \\
&= \int \log p(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} + \int \log p_\phi(\mathbf{z}|\mathbf{x}) p_\phi(\mathbf{x}, \mathbf{z}) d\mathbf{z} d\mathbf{x} - \int \log p_\theta(\mathbf{x}, \mathbf{z}) p_\phi(\mathbf{x}, \mathbf{z}) d\mathbf{z} d\mathbf{x} \\
&= \int \log p(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} + \int \log \left(\frac{p_\phi(\mathbf{z}|\mathbf{x})}{p_\theta(\mathbf{x}, \mathbf{z})} \right) p_\phi(\mathbf{x}, \mathbf{z}) d\mathbf{z} d\mathbf{x}
\end{aligned}
\tag{4}
$$

We only care about the negative log likelihood part of the original objective, i.e. the only part that is parameterized

$$
D_{\mathrm{KL}}\left(p(\mathbf{x}), p_\theta(\mathbf{x})\right) = \int \log p(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} - \int \log p_\theta(\mathbf{x}) p(\mathbf{x}) d\mathbf{x}
$$

Plugging back in the lower bound in Eq 4 we arrive at the negative variational free energy or ELBO (left), the evidence lower bound of the log likelihood (right)

$$
\begin{aligned}
-\mathcal{L}(\theta, \phi) &\doteq -D_{\mathrm{KL}}(p_\phi(\mathbf{z}|\mathbf{x}), p_\theta(\mathbf{x}, \mathbf{z})) \\
&= -\int \log \left(\frac{p_\phi(\mathbf{z}|\mathbf{x})}{p_\theta(\mathbf{x}, \mathbf{z})} \right) p_\phi(\mathbf{x}, \mathbf{z}) d\mathbf{z} d\mathbf{x} \\
&\leq \int \log p_\theta(\mathbf{x}) p(\mathbf{x}) d\mathbf{x}
\end{aligned}
$$

The ELBO has been fundamental in the Free-Energy Principle algorithms (K. J. Friston & Stephan, 2007; Parr, Pezzulo, & Friston, 2022), variational autoencoders (Rezende, Mohamed, & Wierstra, 2014) and Expectation-Maximization algorithms (Banerjee, Merugu, Dhillon, & Ghosh, 2005). It may further be expressed as

$$
\mathcal{L}(\theta, \phi) \doteq -\mathbb{E}_{p(\mathbf{x})}[D_{\mathrm{KL}}\left(p_\phi(\mathbf{z}|\mathbf{x}), p_\theta(\mathbf{z})\right) + \mathbb{E}_{p_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})]
$$

Maximizing the ELBO on the log-likelihood requires the variational distribution $p_\phi(\mathbf{z}|\mathbf{x})$ over the latent neural response $\mathbf{z}$, which can be interpreted, in the context of information theory as representing the cost of coding the errors and parameters. This implies finding neural responses $\mathbf{z}$ that maximize the ELBO of the likelihood under internal constraints, such as metabolic costs, bounds on capacity/bandwidth and/or noise.

Approximating the loss gradient We want to differentiate and optimize the lower bound $\mathcal{L}(\theta, \phi)$ with respect to the generative parameters θ and the variational parameters ϕ. Taking the gradient with respect to the generative parameters θ yields

$$
\nabla_\theta \mathcal{L}(\theta, \phi) = -\int \log p_\theta(\mathbf{x}, \mathbf{z}) p_\phi(\mathbf{z}|\mathbf{x}) p(\mathbf{x}) d\mathbf{x} d\mathbf{z} \tag{5}
$$

whereas the gradient with respect to recognition parameters ϕ yields

$$\nabla_\phi \mathcal{L}(\theta,\phi) = \int \log \frac{p_\theta(\mathbf{x},\mathbf{z})}{p_\phi(\mathbf{z}|\mathbf{x})} \nabla_\phi p_\phi(\mathbf{z}|\mathbf{x}) p(\mathbf{x}) d\mathbf{x} d\mathbf{z} = \int \underbrace{[\log p_\theta(\mathbf{x},\mathbf{z}) - \log p_\phi(\mathbf{z}|\mathbf{x})]}_{\mathbf{r}(\mathbf{x},\mathbf{z})} \nabla_\phi p_\phi(\mathbf{z}|\mathbf{x}) p(\mathbf{x}) d\mathbf{x} d\mathbf{z}$$

where $\mathbf{r}(\mathbf{x},\mathbf{z}) = \log \frac{p_\theta(\mathbf{x},\mathbf{z})}{p_\phi(\mathbf{z}|\mathbf{x})}$ can be viewed as a reward for a credit assignment problem.

Given individual stimuli sampled from $\mathbf{x}^{(i)} \sim p(\mathbf{x})$, and generated samples from the posterior $\mathbf{z}^{(i)} \sim p_\phi(\mathbf{z}|\mathbf{x})$ we may approximate the gradients with respect to the generative parameters θ using an empirical mean, and update them using gradient-descent

$$\theta' \leftarrow \theta + \Delta\theta \qquad \Delta\theta \approx -\alpha \nabla_\theta \mathcal{L}(\theta,\phi) \propto \textstyle\sum_{\mathbf{x}^{(i)}} \sum_{\mathbf{z}^{(i)}} \nabla_\theta \log p_\theta(\mathbf{x}^{(i)},\mathbf{z}^{(i)})$$

where α is a small positive learning rate. Similarly, to estimate the gradient with respect to the recognition parameters ϕ, given individual stimuli sampled from $\mathbf{x}^{(i)} \sim p(\mathbf{x})$, and generated samples from the posterior $\mathbf{z}^{(i)} \sim p_\phi(\mathbf{z}|\mathbf{x})$ we may use

$$\phi' \leftarrow \phi + \Delta\phi, \qquad \Delta\phi \approx -\alpha \nabla_\phi \mathcal{L}(\theta,\phi) \propto \textstyle\sum_{\mathbf{x}^{(i)}} \sum_{\mathbf{z}^{(i)}} \mathbf{r}(\mathbf{x}^{(i)},\mathbf{z}^{(i)}) \nabla_\phi \log p_\phi(\mathbf{z}^{(i)}|\mathbf{x}^{(i)})$$

with $\mathbf{r}(\mathbf{x}^{(i)},\mathbf{z}^{(i)}) = \log \frac{p_\theta(\mathbf{x}^{(i)},\mathbf{z}^{(i)})}{p_\phi(\mathbf{z}^{(i)}|\mathbf{x}^{(i)})}$.

B.1.1 Joint generation and neural response sampling

EBMs adopt the sampling method to generate neural responses from the posterior distribution or the joint model $p_\theta(\mathbf{x},\mathbf{z})$. The latter is an option due to the relationship $\nabla_\mathbf{z} \log p_\theta(\mathbf{z}|\mathbf{x}) \propto \nabla_\mathbf{z} \log p_\theta(\mathbf{x},\mathbf{z}) \propto \nabla_\mathbf{z} E_\theta(\mathbf{x},\mathbf{z})$. This means we can use Monte Carlo sampling (MCMC) to obtain the samples $\mathbf{z}^{(i)}$. There exist several MCMC samplers that can perform the procedure.

Langeving sampler The simplest of them is Langevin sampling, which uses

$$\mathbf{z}' \leftarrow \mathbf{z} + \Delta\mathbf{z}, \qquad \Delta\mathbf{z} \doteq -\beta \nabla_\mathbf{z} \log p_\theta(\mathbf{x},\mathbf{z}) + \alpha\xi$$

A real concern with MCMC methods is that the Markov chain move through all areas of significant probability. If the target distribution has many modes or "islands" of high density, then it will take a long time to move from one island to another. When the initial proposal distribution used when sampling has a very large variance, then the chance of landing on a high-density island is small. Auxiliary variable MCMC methods such as stochastic Hamiltonian Monte-Carlo (HMC) and other neural adaptation procedures have been developed to address these concerns, making non-local jumps possible so that we can more easily jump from one mode to another.

Neural adaptation with a Hamiltonian sampler HMC adds a random momentum variable and then simulates a particle moving on an energy surface. The momentum term keeps the network moving while reaching the local minima. This kind of neural adaptation accumulates noise terms ξ coming, e.g., from ion concentrations, release of neural transmitters, activation/inactivation of ion channels, drug-induced activation of receptors, etc. The momentum variable typically also aggregates an independent component,

generally a random normal distribution $\mathcal{N}(0, \sigma)$, which may be independent of $\mathbf{z}$ or centered around it. Together these dynamics act as momentum to accelerate the sampling process and enhance the exploration of the posterior. Consider the simplified momentum-based dynamics

$$\mathbf{e}' \leftarrow \mathbf{e} + \Delta\mathbf{e}, \qquad \Delta\mathbf{e} \doteq -\beta\nabla_{\mathbf{z}} \log p_\theta(\mathbf{z}, \mathbf{e}, \mathbf{x}) + \alpha\xi$$
$$\mathbf{z}' \leftarrow \mathbf{z} + \Delta\mathbf{z}, \qquad \Delta\mathbf{z} \doteq \mathbf{v}'$$

with $p(\mathbf{e}|\mathbf{z}, \mathbf{x}) = \mathcal{N}(\mathbf{z}, \sigma)$ Gaussian noise dependent on $\mathbf{z}$ capturing the ability of inhibitory cells to respond to input, and ξ the perturbation added to the system by adding noise.

C Simulation Methods

The cortical dynamics simulator code was written in python, and is freely available to use with an open-source license at: `https://github.com/awjuliani/serotonin-ebm`. Table 2 presents the hyperparameters used in the experiments presented in this work. Section C.1 presents a set of studies which assess the sensitivity of our simulated results to values of a number of key hyper-parameters.

Parameter	**Value**
surface size (N x N)	50x50
surface pattern	"noise"
energy function σ	5
replications	100
total time-steps (T)	400
warm-up time-steps	50
drug time-steps	200
cooldown time-steps	150
beta distribution (α, β)	(2, 5)
grad steps per time-step (k)	10
homeostatic α	0.05
plasticity α	0.5
dose level descriptors	[`zero`, `min`, `light`, `medium`, `heavy`, `max`]
5-HT1a α	[0.0, 0.03, 0.06, 0.13, 0.25, 0.5]
5-HT2a α	[0.0, 0.03, 0.06, 0.13, 0.25, 0.5]
5-HT1a σ	2.5

Table 2: Parameters used for simulation experiments.

C.1 Simulator parameter sensitivity studies

We assessed the sensitivity of the simulated results to four key hyper-parameters; (1) Hebbian plasticity, (2) homeostatic strength, (3) the number of gradient steps taken when inferring the neural response (k in Algorithm 1), and (4) the energy function initializations.

Hebbian plasticity Figure 6 presents results for a range of plasticity values. We find that our simulated results and overall conclusions remain consistent when plasticity α is varied from 0.1 to 1.0 (the default is 0.5). We do observe that when $\alpha = 1.0$ the post acute number of local minima for mixed agonists falls relative to the baseline plasticity value, however this does not meaningfully affect our conclusions as the difference is small relative to the acute changes in local minima count. However, when plasticity is very large ($\alpha = 2.0$), our conclusions no longer hold. In this case divergence increases for all agonist mixes and strengths, including the baseline (zero). Here, we observe that the post acute gradient magnitude is substantially larger compared to the baseline plasticity, as is the number of local minima. This indicates the energy function has become more complex and it is easier to get stuck in sub-optimal local minima, likely leading to the observed increase in post-acute divergence. Note that since we observe the the baseline divergence increasing we do not think that 2.0 to be a very realistic value for plasticity α. We also present results for $\alpha = 0$ to show the simulator dynamics in the absence of plasticity, however we note this is also not a realistic value for this parameter as it implies no learning is taking place.

Homeostatic strength Figure 7 presents results varying strengths of the homeostatic constraint. Overall we find that our simulated results and overall conclusions remain consistent when homeostatic α is varied from 0.025 to 0.1 (the default is 0.05). The weaker the homeostatic constraint, the more the acute drug phase effects are amplified. For example, compare the acute phase reduction in energy value or increase in gradient magnitude and divergence for `2a_max` when $\alpha = 0.05$ and $\alpha = 0.01$. The stronger the homeostatic constraint, the faster the dynamics stabilize after the drug effect ends (vertical dotted line). When homeostatic α is very low, at 0.01, the energy value, gradient magnitude, local minima, and divergence all continue to change substantially for many time steps after the drug effect wears off. Further, the energy value and gradient magnitude have not stabilized by the end of the simulated experiment. This suggests that 0.01 is not a realistic value for this parameter. We also present results for $\alpha = 0$ to show the simulator dynamics in the absence of any homeostatic constraint, however we note this is not a biologically plausible value for this parameter.

Gradient steps We find that the simulated results are not particularly sensitive to the number of inference gradient steps each iteration, and that our overall conclusions would remain the same across a range of values. Figure 8 indicates that results are broadly consistent between 1 and 100 steps (the default is 10). Taking 100 or 10 gradient steps leads to almost identical results. When 1 gradient step is taken, the divergence and gradient magnitude is very similar to those when taking 10 or 100 steps. State visitation counts increase slightly while the number of local minima decreases marginally for the baseline and lower

strength doses. The acute phase energy value for mixed agonists is somewhat sensitive to changes in the number of gradient steps. As gradient steps increase from 1 to 10, energy decreases during the acute phase, especially for the `mix_max` agonist. This is intuitive, since more gradient steps creates more opportunities for the neural response to descend the modulated energy surface.

Energy function initialization We considered six different functions — `noise` (default), `log`, `sine`, `swirl`, `tangent`, `twisted sine` — for generating the initial energy function and the results are presented in Figure 9. Details of the functions are given in Table 3. We find that our main conclusions hold across a range different methods for initializing the energy functions $E^*(\mathbf{z})$ and $E_(\mathbf{z})$. Additionally, in most cases the results and rank ordering of different agonists are consistent. There are two exceptions, the tangent landscape, and the energy value metric.

When the tangent landscape is used, final KL-divergence values are higher and the error bars wider. However, consistent with our default generating function, we find that the mixed agonists `max-max`, `medium-medium` and `1a-max-2a-heavy` still result in the lowest final KL-divergence value and that the values are statistically significant [all $p < 0.001$] when compared with the baseline. Pure 2a agonism at `max` strength is also significantly lower than the baseline [$p < 0.001$]. Additionally, when we simulate all possible permutations of 5-HT2a and 5-HT1a dose strengths (see Figure 10), the results are consistent with those presented in Figure 5. Interestingly, in this case a biased agonist of `1a-max-2a-heavy` has a relatively lower final KL-divergence value compared with the default noise landscape, and this value is close to the max-max combination. Finally, we note that the baseline divergence increases for this generating function, suggesting that the other choices of generating function are more suitable.

Considering the energy value metric, we observe that the acute and post acute values for pure agonist (whether 5-HT2a or 5-HT1a) are consistent across different generating functions. We consistently observe the acute phase transient over-fitting for 5-HT2a agonists and acute phase transient under-fitting for 5-HT1a agonists. However for mixed agonists, both the acute and post acute values appear sensitive to the choice of generating function. The acute phase may be characterized by an increase, decrease, or mixed change in the energy value. Additionally, the relative difference between the final energy value for mixed agonists compared to pure agonists, as well as the ordering between `mix-max` and `1a-max-2a-heavy` varies. What remains consistent across all landscapes is (1) that the final energy value for mixed agonists are lower than any acute phase value for the same agonist and (2) the final energy value of both the `mix-max` and `1a-max-2a-heavy` is lower than the baseline and any of the pure agonists. In 60% of landscapes (sine, tangent, twisted sine) the relative difference is similar to the baseline, in 20% of cases (swirl) the difference is moderately less but still substantial, and in 20% of cases (log) the gap is much smaller.

Name	**Function**
`noise` (default)	n
`log`	$n \log(\sqrt{x^2 + y^2})$
`sine`	$n \sin(\sqrt{x^2 + y^2})$
`swirl`	$n \sin(x^2 - y^2)$
`tangent`	$n \tan(\sqrt{x^2 + y^2})$
`twisted sine`	$n \sin x \cos y$

Table 3: Initialization functions for $E^*(\mathbf{z})$ and $E_(\mathbf{z})$. $n \sim \mathcal{N}(0, 1)$, (x, y) are coordinates of an $N \times N$ grid of points. Additional processing is performed by convolving the matrix with a guassian kernel $G(\sigma)$ and normalizing the values to be within the range $[0, 1]$.

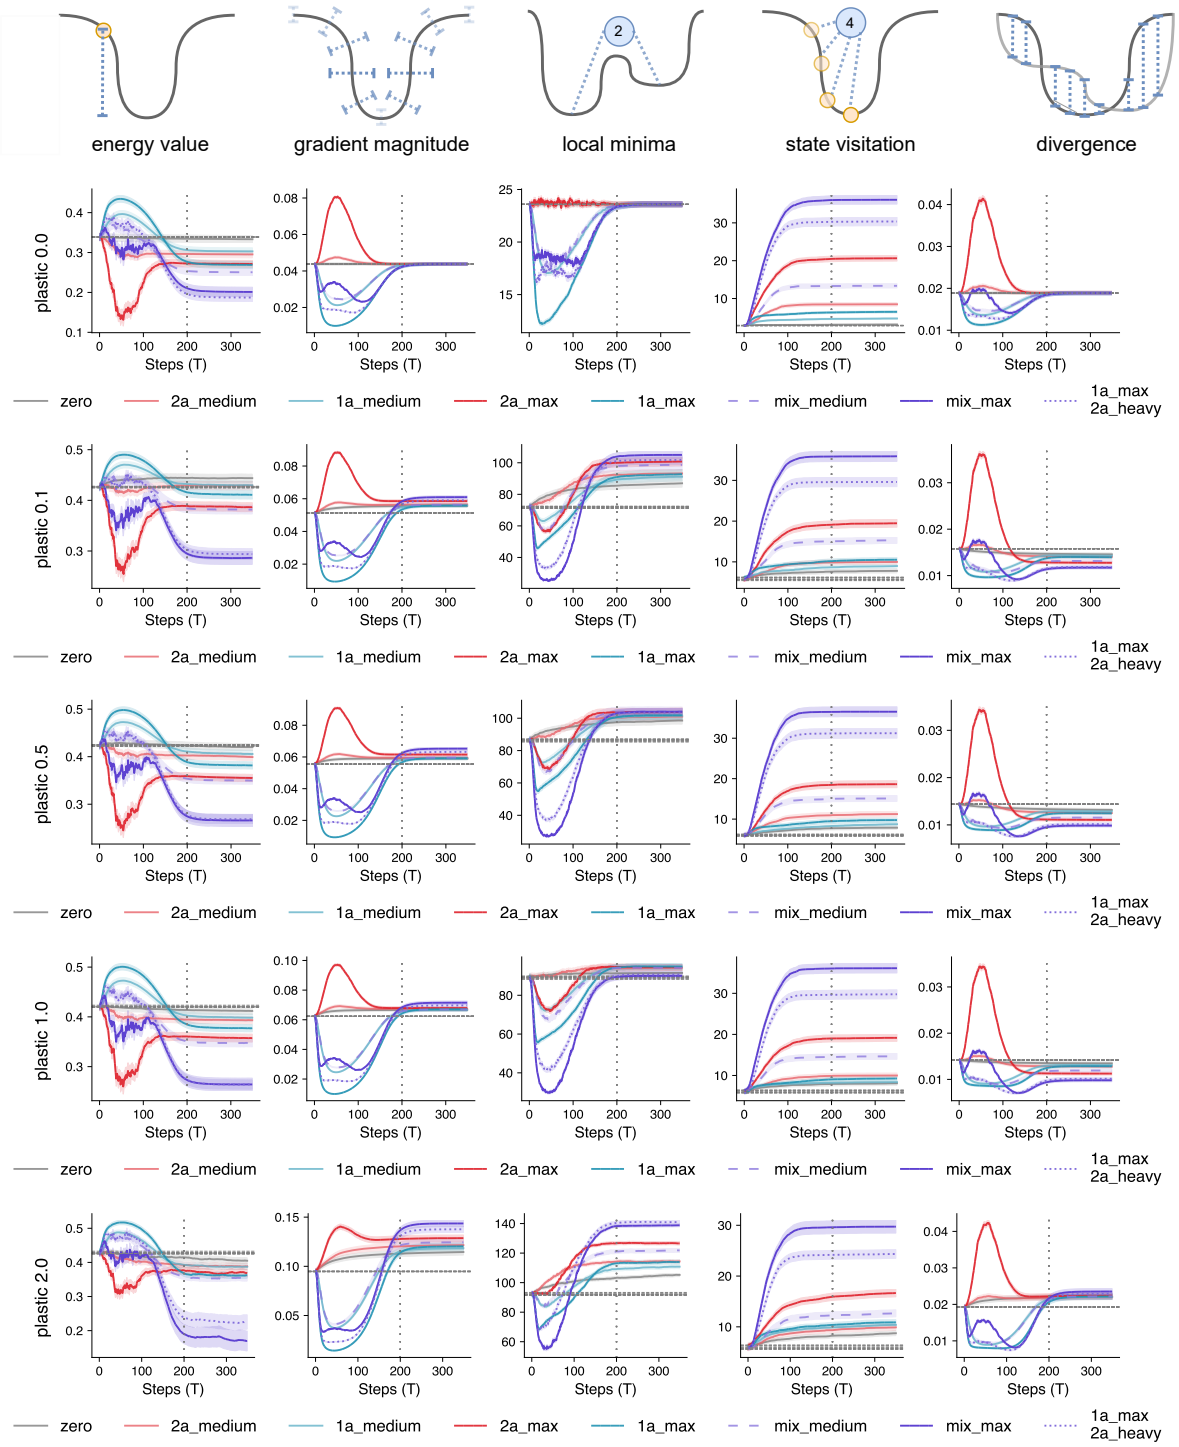

Figure 6: Hebbian plasticity hyper-parameter ($plastic_str$) sweep.

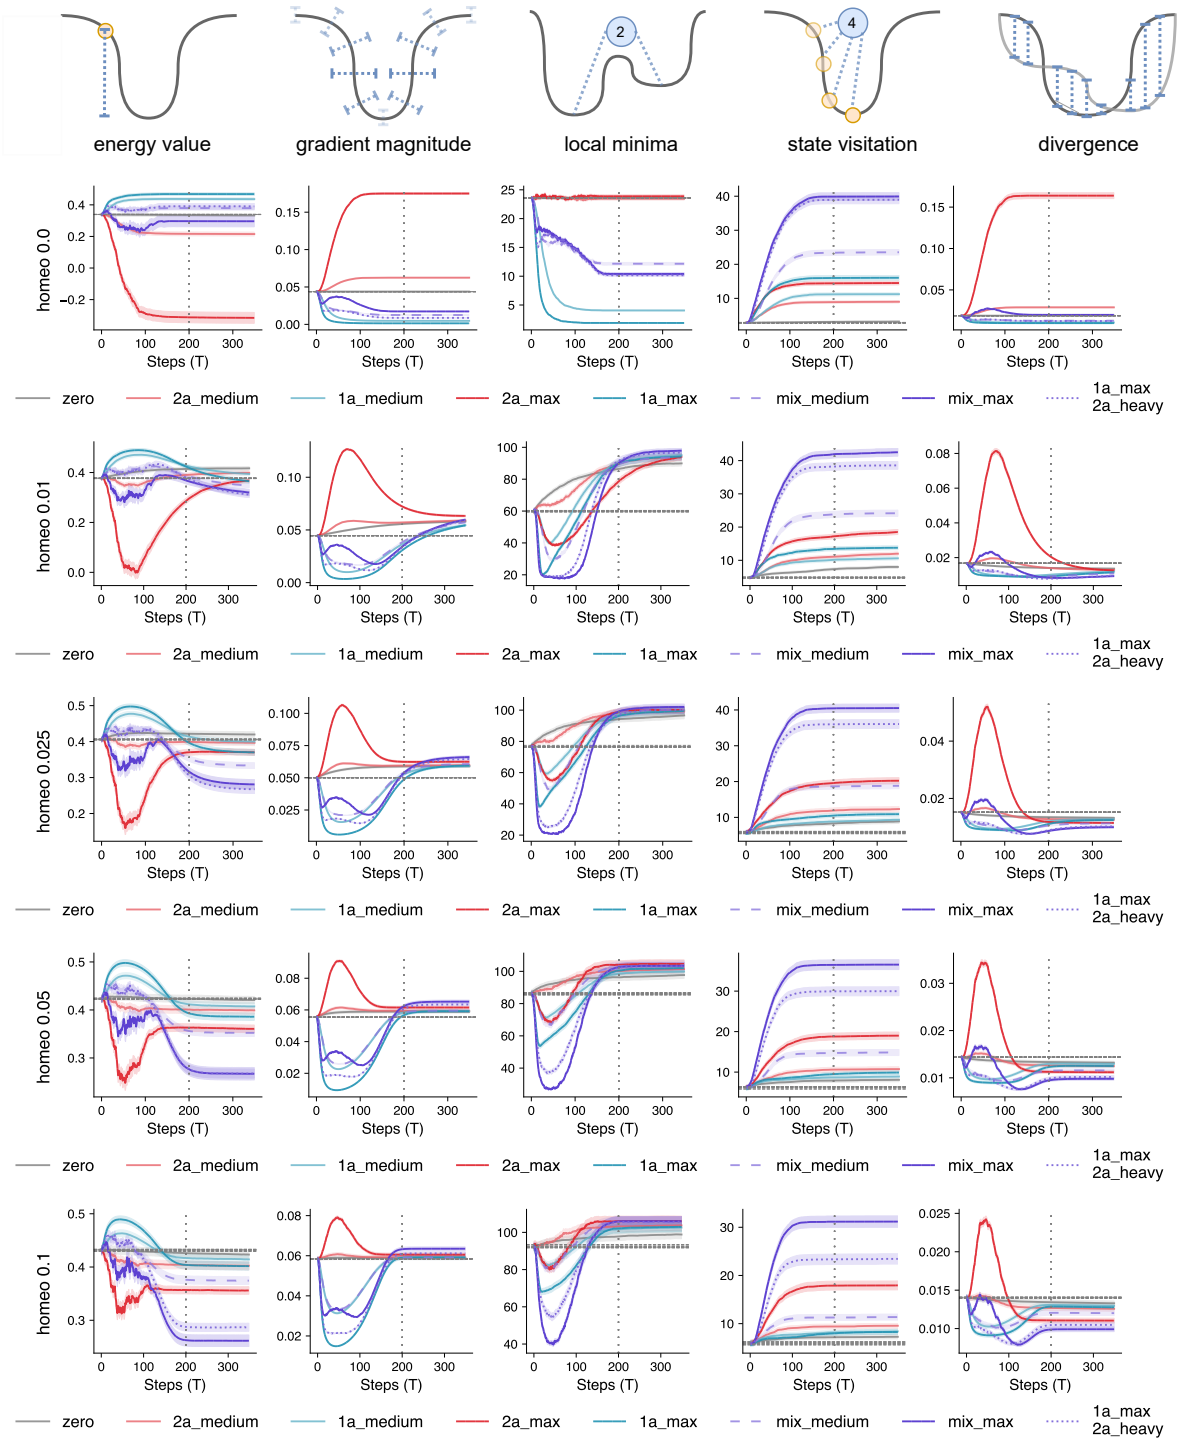

Figure 7: Homeostatic constraint hyper-parameter ($constraint_str$) sweep.

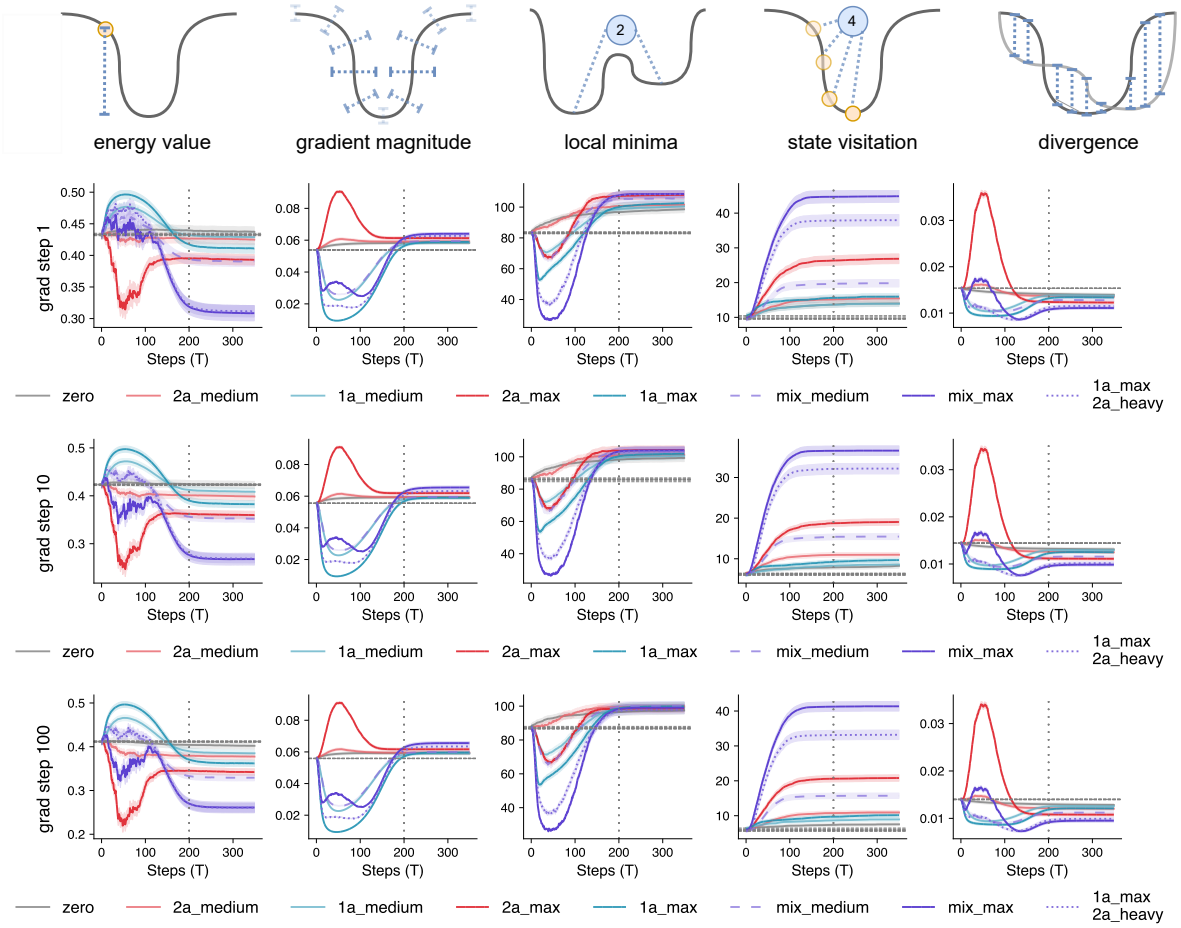

Figure 8: Number of gradient steps on z ($grad_step_num$) sweep.

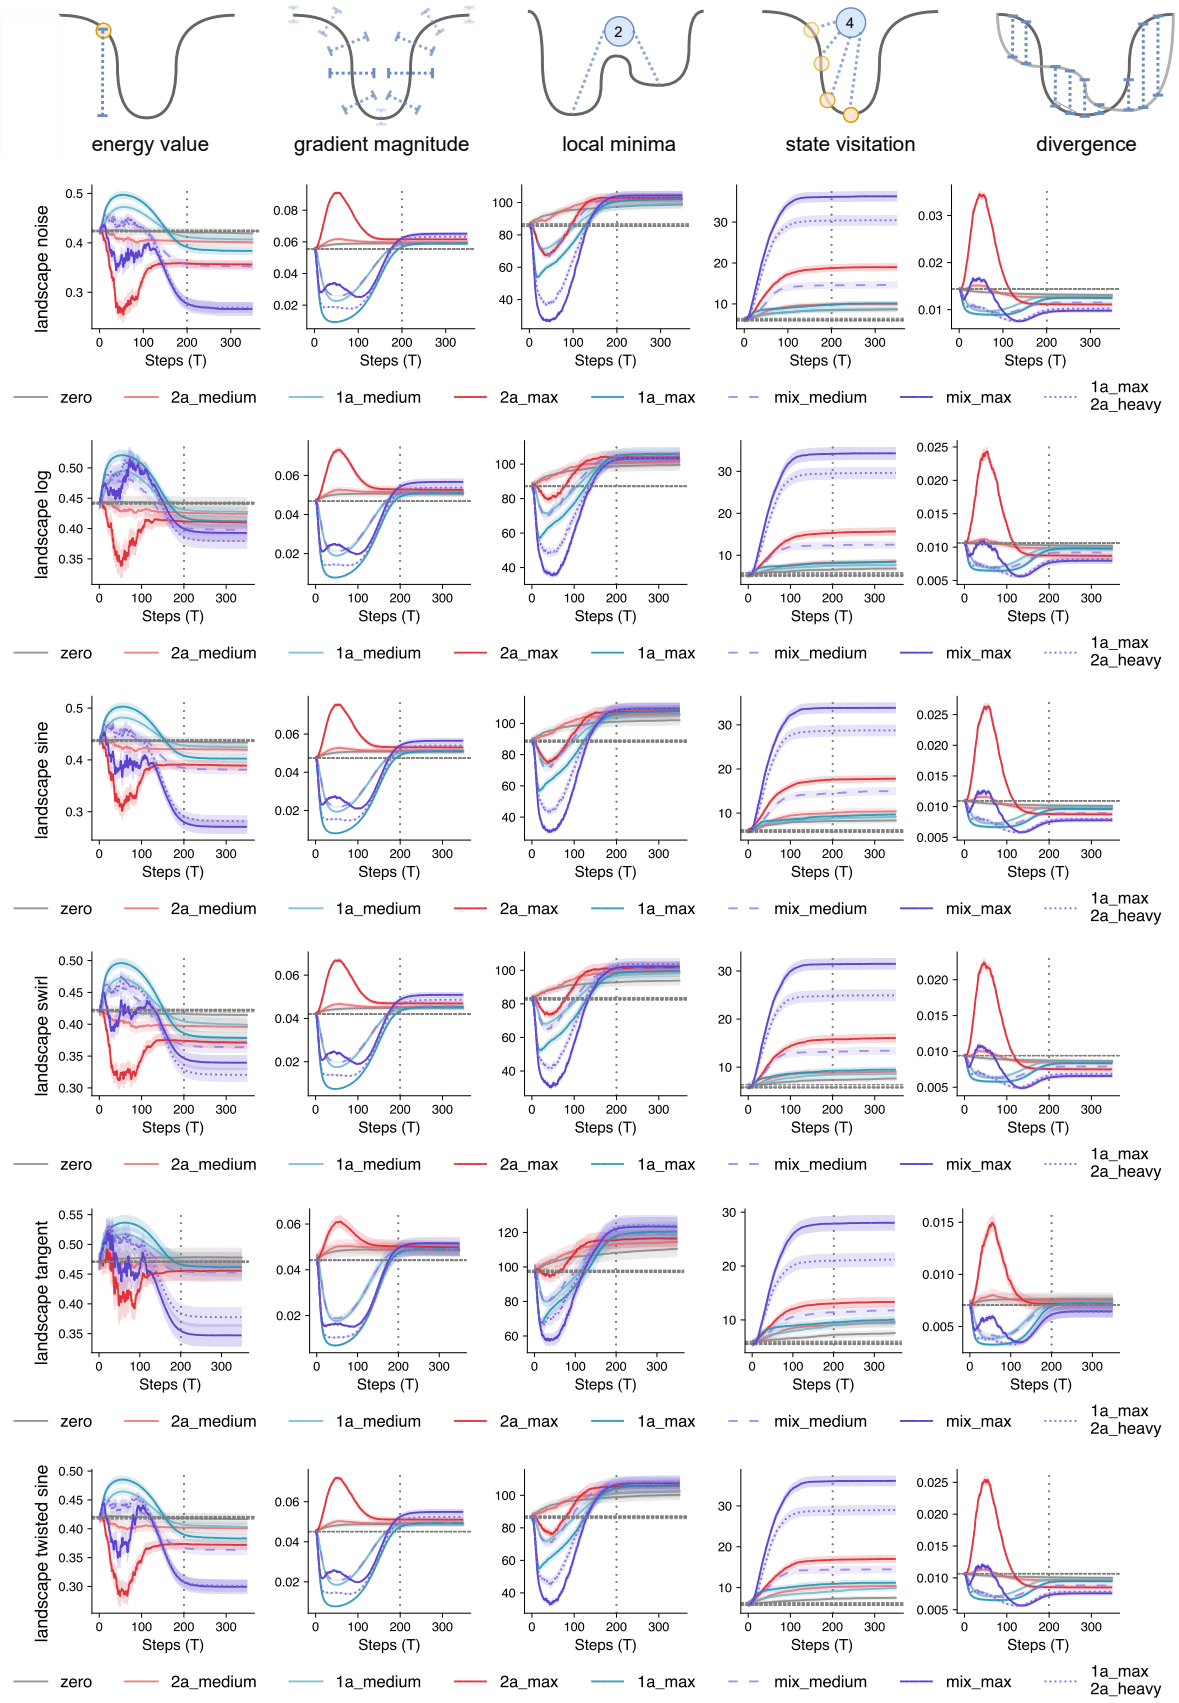

Figure 9: Energy function generator ($surface_pattern$) sweep.

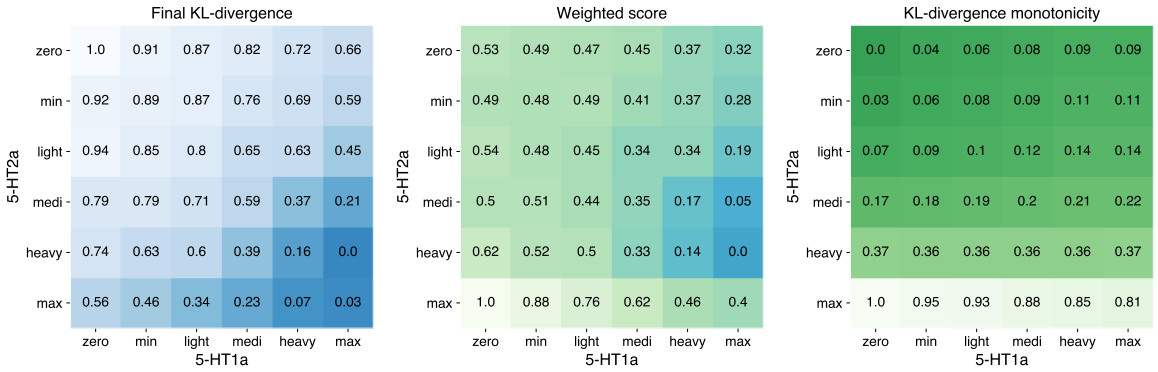

Figure 10: Matrix reporting normalized final energy values for set of possible 5-HT2a and 5-HT1a agonist combinations using a tangent energy function initialization. The results and conclusions are consistent with those reported in Figure 5.

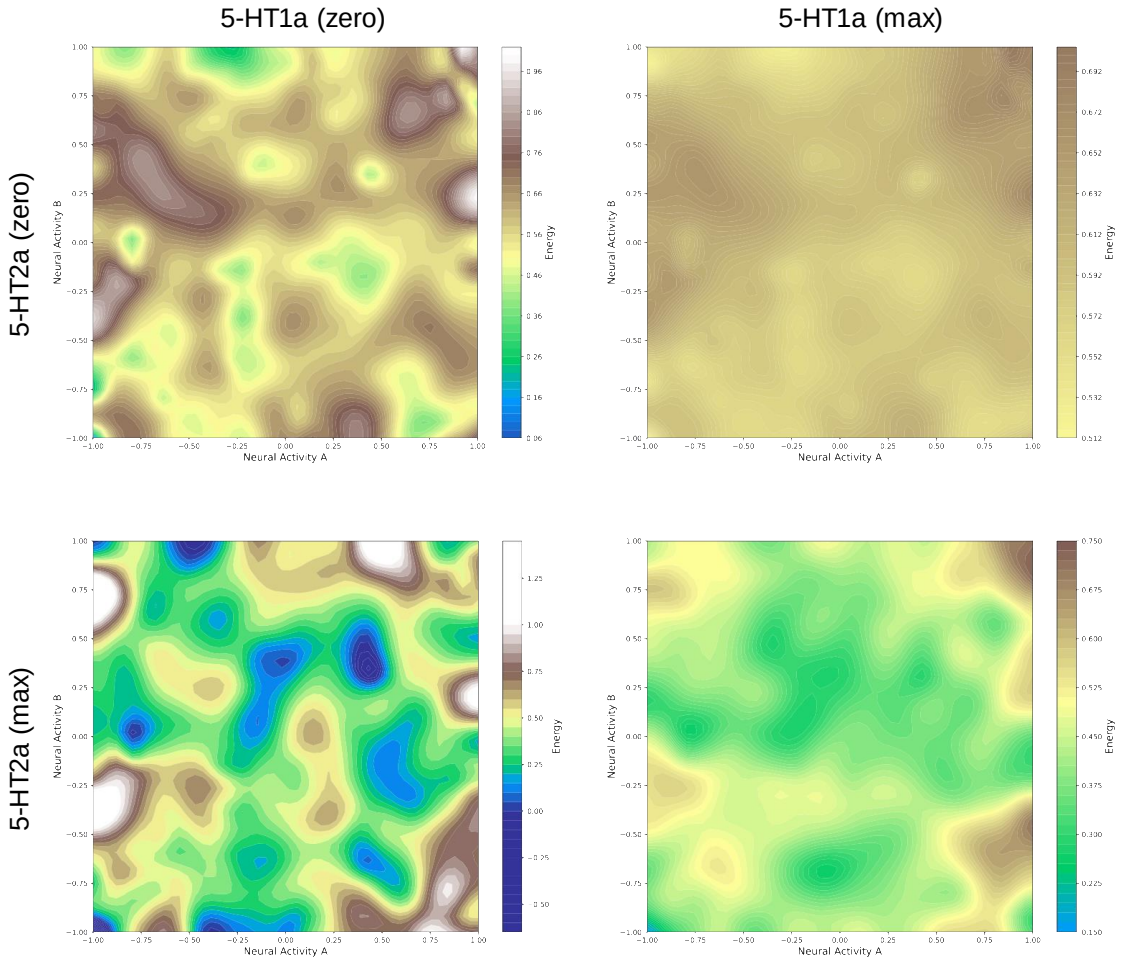

Figure 11: Representative examples of the energy function $\tilde{E}$ at the peak of the drug effect. **(Top Left)** No neuromodulation. **(Top Right)** Max level of 5-HT1a neuromodulation. **(Bottom Left)** Max level of 5-HT2a neuromodulation. **(Bottom Right)** Max levels of both 5-HT2a and 5-HT1a agonism.