



OPEN ACCESS

EDITED BY

Lejian Huang
Northwestern University, United States

REVIEWED BY

William Sedley
Newcastle University, United Kingdom
Daniel S. Barron
Brigham and Women's Hospital and Harvard
Medical School, United States

*CORRESPONDENCE

Anna-Lena Eckert
eckertan@staff.uni-marburg.de

SPECIALTY SECTION

This article was submitted to Pain Research
Methods, a section of the journal Frontiers in
Pain Research

RECEIVED 13 June 2022

ACCEPTED 25 August 2022

PUBLISHED 16 September 2022

CITATION

Eckert A-L, Pabst K and Endres DM (2022) A
Bayesian model for chronic pain.
Front. Pain Res. 3:966034.
doi: 10.3389/fpain.2022.966034

COPYRIGHT

© 2022 Eckert, Pabst and Endres. This is an
open-access article distributed under the terms
of the [Creative Commons Attribution License
\(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or reproduction in
other forums is permitted, provided the original
author(s) and the copyright owner(s) are
credited and that the original publication in this
journal is cited, in accordance with accepted
academic practice. No use, distribution or
reproduction is permitted which does not
comply with these terms.

A Bayesian model for chronic pain

Anna-Lena Eckert*, Kathrin Pabst and Dominik M. Endres

Theoretical Cognitive Science Group, Department of Psychology, Philipps-University Marburg,
Marburg, Germany

The perceiving mind constructs our coherent and embodied experience of the world from noisy, ambiguous and multi-modal sensory information. In this paper, we adopt the perspective that the experience of pain may similarly be the result of a probabilistic, inferential process. Prior beliefs about pain, learned from past experiences, are combined with incoming sensory information in a Bayesian manner to give rise to pain perception. Chronic pain emerges when prior beliefs and likelihoods are biased towards inferring pain from a wide range of sensory data that would otherwise be perceived as harmless. We present a computational model of interoceptive inference and pain experience. It is based on a Bayesian graphical network which comprises a hidden layer, representing the inferred pain state; and an observable layer, representing current sensory information. Within the hidden layer, pain states are inferred from a combination of priors $p(\text{pain})$, transition probabilities between hidden states $p(\text{pain}_{t+1} | \text{pain}_t)$ and likelihoods of certain observations $p(\text{sensation} | \text{pain})$. Using variational inference and free-energy minimization, the model is able to learn from observations over time. By systematically manipulating parameter settings, we demonstrate that the model is capable of reproducing key features of both healthy- and chronic pain experience. Drawing on mathematical concepts, we finally simulate treatment resistant chronic pain and discuss mathematically informed treatment options.

KEYWORDS

chronic pain, interoception, graphical models, computational psychiatry, Bayesian inference, belief-propagation, predictive coding, free energy

1. Introduction

Pain has an undeniable function in daily life. By signalling potential or actual bodily damage and directing both attention and behavior, it ensures the individual's physical integrity in the long run (1). When touching a hot plate, for example, the individual will withdraw their hand immediately to avoid further tissue damage, and seek to soothe the acquired damage.

The same can not be said about chronic pain, where pain persists beyond the presence of acute injury (1). With around 20% of adults and children reporting symptoms of chronic pain, it is a highly prevalent disorder (2, 3). Chronic pain leads to significant losses of quality of life (adults: (4), adolescents: (5)). The disorder is associated with decreased levels of psycho-social functioning and productivity (6) and

frequent utilization of healthcare services (7), causing a substantial economic burden (8–10).

Despite the ever-rising prevalence of chronic pain (11, 12), its underlying mechanisms remain poorly understood. Computational-level mathematical models (cmp. (13)) can help to elucidate the problem that the mind is trying to solve (e.g., detecting harm to the body) and where such solution attempts go wrong. Recent developments in cognitive computational neuroscience and machine learning may offer new opportunities for studying chronic pain from this statistical viewpoint.

A growing body of work suggests interoception, just like perception, follows probabilistic, Bayesian principles (14–18). Perceiving both the external world and the internal milieu requires effectively dealing with noisy, ambiguous and incomplete sensory information, usually from a multitude of sources and modalities. For this, the mind appears to rely on prior knowledge in the shape of a generative model of the environment, which is integrated with the current sensory input (19–21). Following this account, perception is not only shaped by the sensory information reaching the mind or body, but also by the predictions of the generative model this information is met with. This notion is summarized under the free-energy principle (22, 23).

The modulatory effects of prior knowledge on pain perception have been widely demonstrated - from the level of neural processing in the brain stem (24) through to self-reported pain intensity (25–28). For example, Tabor et al. (16) show that pain perception varies in dependence of an exteroceptive cue. When paired with a blue light (signalling safety), a noxious stimulus is perceived as less painful compared to when it is paired with a red light (signalling danger and heat). Further, placebo analgesia describes the phenomenon of experiencing pain relief after taking a placebo pill (15, 29, 30). In line with statistical accounts of pain, this further suggests an integration of expectations, or prior beliefs, with the physical sensory signal. In this context, Hechler and colleagues (2016) have instantiated Bayes' theorem with pain and nociception-related variables to illustrate statistical computations underlying pain perception:

$$p(\text{pain} | \text{sensation}) = \frac{p(\text{sensation} | \text{pain}) * p(\text{pain})}{p(\text{sensation})} \quad (1)$$

where $p(\text{pain} | \text{sensation})$ is the **posterior** probability. The posterior determines the individual's subjective experience of pain. It is proportional to the **likelihood** $p(\text{sensation} | \text{pain})$ times **prior** probability of pain, $p(\text{pain})$. $p(\text{sensation} | \text{pain})$ describes the probability of being exposed to a certain sensation (i.e., nociceptive signals) when the body or a specific body part is in a state of pain. Following this account, chronic pain has been described in terms of aberrant Bayesian

inference (14). For example, aversive life events may lead to a heightened prior expectation of pain, $p(\text{pain})$. This may underlie lowered pain thresholds and hyperalgesia frequently observed in chronic pain patients (31). Long-term learning experiences may alter the associations between sensory information and the inferred experience in patients. In other words, the likelihood term $p(\text{sensation} | \text{pain})$ may be subject to learned biases in patients so that a broader range of stimuli is associated with pain compared to healthy controls (14). This may underlie allodynia, where harmless stimuli are perceived as painful.

We build on these previous lines of work and propose a quantitative Bayesian framework, implementing the above-mentioned considerations within a hierarchical and sequential model. It is built on the assumptions that (i) pain perception emerges from Bayes-optimal combination of sensory input and prior beliefs, (ii) chronic pain is characterized by maladaptive learning processes on longer time-scales, (iii) these learning processes result in generalizing, heightened expectations of pain $p(\text{pain})$ as well as (iv) heightened $p(\text{sensation} | \text{pain})$, i.e. patients erroneously infer being in pain as the most likely cause of a wide range of sensations usually considered harmless.

Our formal implementation is based on the assumption that interoception is an inference problem that requires optimizing the generative models, i.e. the mental representation of causal relations between body state (i.e. a pain-free or pain state) and sensory (e.g. harmless or noxious) inputs. Using concepts from machine learning, such as variational inference, the system's state is approximated (32, 20). Pain is represented in latent state variables and must be inferred from previous model states and sensory information, represented as observable variables. By representing perception and sensation on different computational levels, our model allows the formalization of non-veridical relationships between physical stimulus and pain perception. In other words, pain needs to be inferred using all available information, such as nociceptor firing rates (sensations), general pain expectations and specific pain expectations based on recent percepts. The model incorporates learning, which is implemented using message passing and the minimization of free energy over time (22, 33). We first describe the model's architecture and features, before simulating the effects of the parameters. Here, we focus on the effects of a prior probability of pain, $p(\text{pain})$, a likelihood model $p(\text{sensation} | \text{pain})$, and $p(\text{pain}_{t+1} | \text{pain}_t)$, which describes the probability of remaining in a state of pain between subsequent time steps t and $t + 1$. To illustrate our results, chronic pain can emerge from a high and precise prior belief to be in pain $p(\text{pain})$, which is generally more likely than not being in pain (here over-lined, representing the complement of pain; $p(\overline{\text{pain}})$); $p(\text{pain}) > p(\overline{\text{pain}})$, in patients. The prior is sharpened further with every time-step where pain is inferred as the most likely cause of a given

sensation, leading to a stabilization of the system within this pathological state. The system then is resistant to correctly interpreting harmless sensory information, as represented in an ambiguous $p(\text{sensation} | \text{pain})$.

2. Materials and methods

2.1. Model preliminaries

2.1.1. Sequential data

Chronic pain develops over time, therefore an appropriate computational model has to be sequential in nature. Further, the state of the body a short moment ago is usually the best predictor for its current state. The brain seems to leverage this predictability to process sensory information more efficiently (22, 32). It is hence plausible to introduce a dependency of states over time into the model. Since the mind is assumed to constantly predict future sensory input, the model needs to accommodate longer time-series. The more recent past is usually more relevant to the current situation than the distant past. For example, what you saw three seconds ago is more relevant to your current situation than things you saw 5 weeks ago. The future state of a Markov chain is conditionally independent of all previous states, given the current state (Markov property). A Markov chain model fulfills the requirements of accounting for time-series dependencies while maintaining computational tractability and hence biological plausibility (for more details, see [Appendix A](#)).

2.2.2. Perceptual hierarchies

While a given intensity of noxious stimulation induces intense pain in one subject, another one may barely experience discomfort (34). The non-veridical relationship between sensation and pain prompts the need for a differential representation of the two within a model. This large inter-individual heterogeneity in pain experience further suggests a crucial role of an individual's generative model of pain experience. We use a Hidden Markov model (35, 36) with latent (hidden) variables, representing the pain state, or H , and observable variables, representing sensory input, or S ; see [Figure 1](#) for a graphical representation. Pain is represented in hidden state variables and functions as an indicator of the body's integrity. This information is not directly accessible to the individual and hence has to be inferred from the available sensory information and prior expectations. From a Bayesian perspective, the experience of pain is determined by performing inference on hidden states, given the individual's prior beliefs. Given the non-linear and inter-individually heterogeneous relationships between sensation (i.e. nociceptor firing after physical stimulation) and percept (pain experience), an appropriate model allows the distinction of the two. In our model, this is achieved via

separate hierarchical layers (observable and hidden). Of note, this computational-level model (in the sense of (13)) does not translate directly to implementational, neuroanatomical levels relevant to pain perception (37).

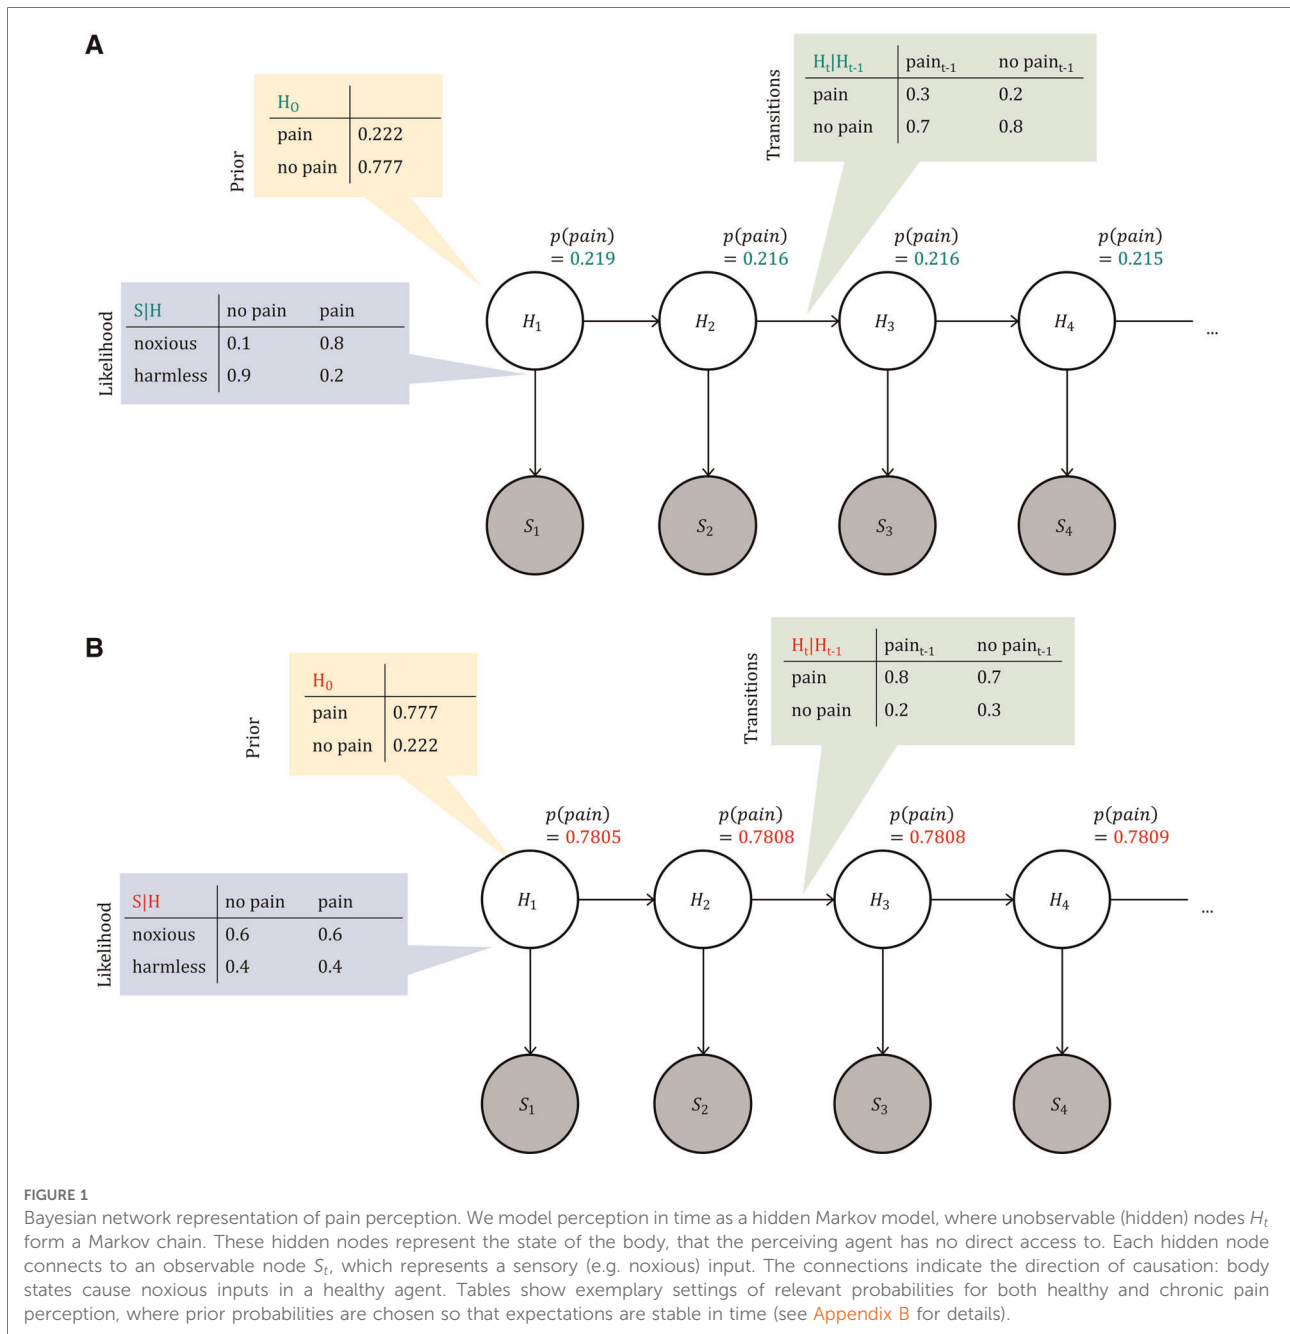
2.3.3. Inference and learning

The perceiving individual does not have direct access to their surroundings. It has to rely on information from its sensory organs, such as signals from sensory receptors, to obtain a representation of its surroundings. The accuracy of this information is crucial to the individual's survival in a given environment. Since the information arriving at the sensory organs is noisy and ambiguous, the mind needs to infer the real environmental causes for the sensation. Both prior beliefs and sensory information appear to be represented in a probabilistic manner (38). Perception then is the Bayes-optimal combination of prior beliefs and said sensory information. Conscious experience is determined by the hidden state that is assigned the highest posterior probability. In the context of pain perception, inference means determining the extent to which an observation $S = s$ allows the perceiver to update their knowledge about a hidden state $H = h$. For this purpose, it is necessary to infer the marginal probabilities of the latent variables from observations. We use the sum-product algorithm, also known as belief-propagation, as an efficient means of deriving marginal probabilities in singly-connected graphical models (36, 39). Inference and learning over time is governed by the imperative to minimize free energy, a principle which has been proposed as a general theory of brain function (22, 20, 40). Variational message passing, as proposed in this framework, has the advantage of both computational tractability and biological plausibility (22). The interested reader is referred to [Appendix A](#) for a detailed description of the approach to inference applied here. While inference in this framework refers to short-term conclusions about marginal probabilities of pain, learning refers to longer-term updates of model parameters. We here implemented batch-updates, where model parameters are updated once after the observation of a full time series. Psychologically, this would correspond to retrospective memory consolidation, for example during sleep.

3. Simulations

All models and simulations were implemented in Python (version 3.9.7), in particular the numpy package (version 1.21.2, (41)) and a package for variational inference with exponential family models by author D.E. (33). All plots were created using the matplotlib.pyplot library for Python (version 3.4.3, (42)).

A Hidden Markov Model with $N = 20$ time-steps and free-energy learning over time is implemented (see [Appendix A](#) for



details). Its hidden state variables, or nodes, H_t , represent the system's internal state at time point t . There are two possible states in each state, ($H \in \{\text{pain}, \overline{\text{pain}}\}$) which are inferred from the combination of previous nodes' states and incoming sensory information. Sensory information is represented within the observable nodes S_t . There are two possible observations: noxious or harmless sensory information, or $S \in \{\text{noxious}, \text{harmless}\}$. Observable nodes can remain unobserved. The top-down factor contains the likelihood $p(S_t | H_t)$ which allows explicit modelling of the association between sensory information and internal model

state. The hidden variable nodes form a Markov chain, connected through additional free-energy factor nodes necessary for message passing. Here, the transition probabilities $p(H_t | H_{t-1})$ represent the implicit expectations of pain, or the pain prior, over time. In other words, transition probabilities translate into the individual's expectation regarding the persistence of pain. A fixed and precise transition probability can constrain the model to arrive at a steady state, where conflicting sensory information does not alter the model's inferred hidden state, which we will discuss in detail below. We systematically vary

the above-mentioned parameters to capture healthy pain perception as well as one that is biased towards chronically inferring pain. While the underlying inference engine is the same in both cases, there are four starting points from which to manipulate the model's behavior, also summarized in **Figure 1**.

1. The **prior** of the first hidden variable H_1 ;
2. The **likelihood function**, or the top-down factor connecting observable and hidden nodes, $p(S_t | H_t)$;
3. The **transition probabilities** $p(H_t | H_{t+1})$ between the hidden variables;
4. The **sensory information** at observable nodes S_t .

Parameters are characterized by both the probabilities for specific states or observations, represented by their sufficient statistics λ ; and their precision, formalized by their pseudocounts ν (see **Appendix A** for details). ν is referred to as the *pseudocount* as it keeps count of the number of observed data points.

Healthy interoceptive inference (“healthy observer,” from here onward) is characterized by a low, but precise prior expectation of pain ($p(\text{pain}) = 0.2$, $\nu = 100$) and a precise likelihood function that allows accurate inferences based on incoming sensory information (e.g., $p(\text{noxious}_t | \text{pain}_t) = 0.8$, $\nu = 100$). Transitions towards pain-free states are more likely than transitions to painful states, as implemented in the transition probabilities between hidden states (e.g. $p(H_t = \overline{\text{pain}} | H_{t-1} = \text{pain}) = 0.7$).

In contrast, *chronic pain inference* (“chronic pain observer”) is characterized by an increased prior expectation of pain ($p(\text{pain}) = 0.9$, $\nu = 100$). The likelihood functions are imprecise and do not allow accurate inference from incoming sensory data, with $p(\text{noxious}_t | \text{pain}_t) = 0.6$, $\nu = 20$. Transitions between hidden states are biased so that transitions into painful states are more likely than transitions into pain-free states, e.g. $p(H_t = \text{pain} | H_{t-1} = \overline{\text{pain}}) = 0.7$.

3.1. Mixed observations

In a first simulation, both observers were exposed to the same sensory information: noxious input for time steps $S_1 - S_5$, and harmless input at $S_{12} - S_{18}$ for 40 learning trials. Results are shown in **Figure 2**. In the healthy observer (**Figure 2A**), the exposure to noxious stimuli between time steps 1 and 5 leads to an acute and highly precise inference to be in pain. In contrast, the presentation of harmless sensory information (time steps 12–18) decreases the probability of inferring to be in pain drastically. In the chronic pain model, however, the response to acute noxious sensations follows a different pattern (**Figure 2B**)—the baseline probability of inferring pain is at ceiling levels already. In contrast, harmless information barely has any

significant effects on the probability of inferring pain, which remains high.

3.2. Prolonged exposure

In another simulation, the two observers are exposed to only-noxious vs. only-harmless sensory observations for all 20 time steps and 40 learning trials. Results are shown in **Figure 3**. In individuals with chronic pain (3A,B), the type of sensory information the system is exposed to does not result in significant changes of $p(\text{pain})$. Especially compared to the healthy inference model (3C,D), exposure to harmless information does not lead to decreased probabilities of inferring pain. In healthy inference, prolonged exposure to noxious information leads to acute heightened probabilities of experiencing pain. In contrast, harmless stimulation results in very low probabilities of inferring pain. Of note, the inferred probabilities of pain at the first- and final time step differ from intermediate time steps. This is an artifact caused by the model architecture with a definite end node, and our batch-updates (learning) that are performed at the end of one time series (see figure caption for details).

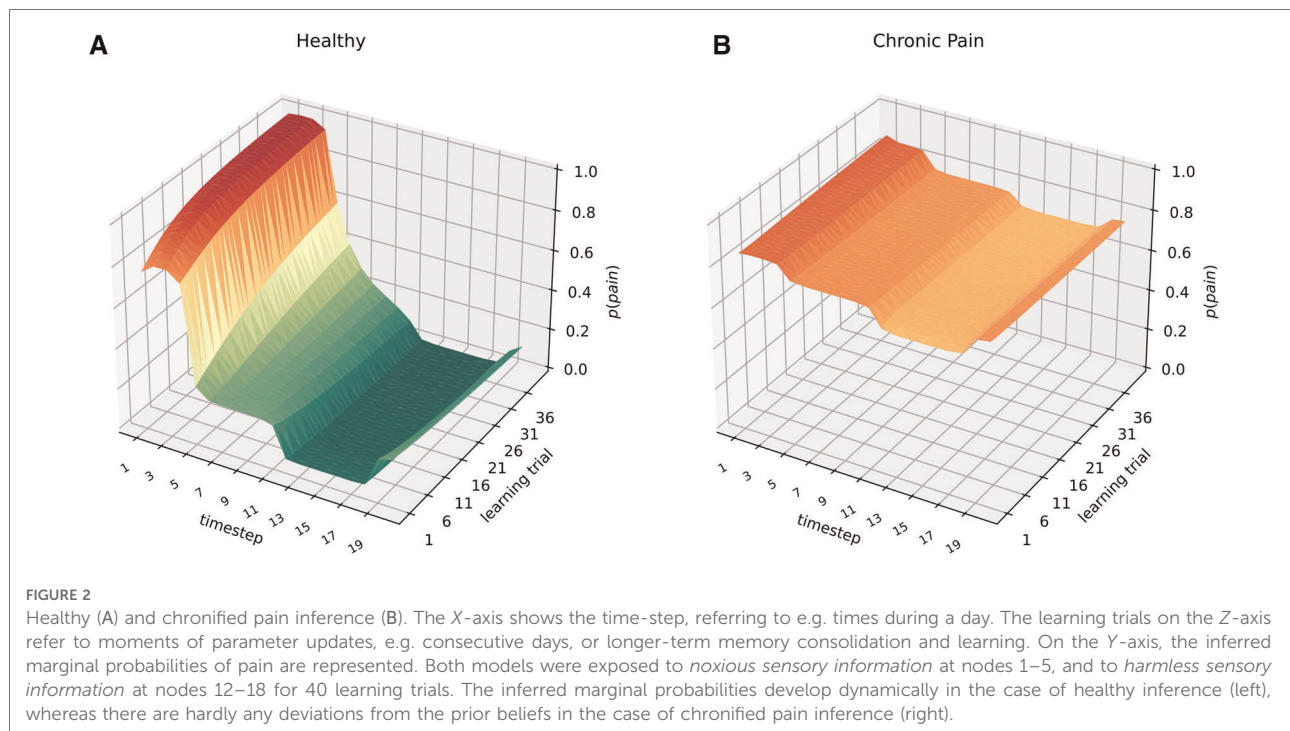
4. Treatment-resistant chronic pain: a null space problem?

Some forms of chronic pain are remarkably resistant to psychotherapeutic (43) or pharmacological (44) intervention, which has led to calls for ever more radical treatments such as transcranial magnetic stimulation, or TMS, (45), electroconvulsive therapy (46) or ketamine infusions (47).

A computational model of chronic pain therefore needs to be able to demonstrate a rigid, “treatment-resistant” pain experience over time. In the present model, this translates to a high temporal stability of the inferred, latent pain-state, regardless of the varying sensory input at the observable nodes (i.e., psychotherapy) or modulation of the sensations (i.e. pharmacotherapy). To derive the parameter settings necessary to arrive in a fully stable model state, we borrow the concept of the null space from linear algebra (36).

The null space of a matrix A contains all vectors \vec{x} such that $A\vec{x} = 0$ (48). We are interested in the null space of the matrix that constrains the transition probability distributions $p(H_{t+1} | H_t)$ to yield a $p(H_t = \text{pain})$ after waiting for a sufficiently long time, i.e. we are looking for the stable states of the Markov chain. In other words, we are interested in changes to the transition probabilities that have no effect in the long run – i.e. a resistance of the system to treatment.

We demonstrate in **Appendix B** that for each marginal $p(\text{pain})$, there are transition probabilities $p(H_{t+1} | H_t)$ that

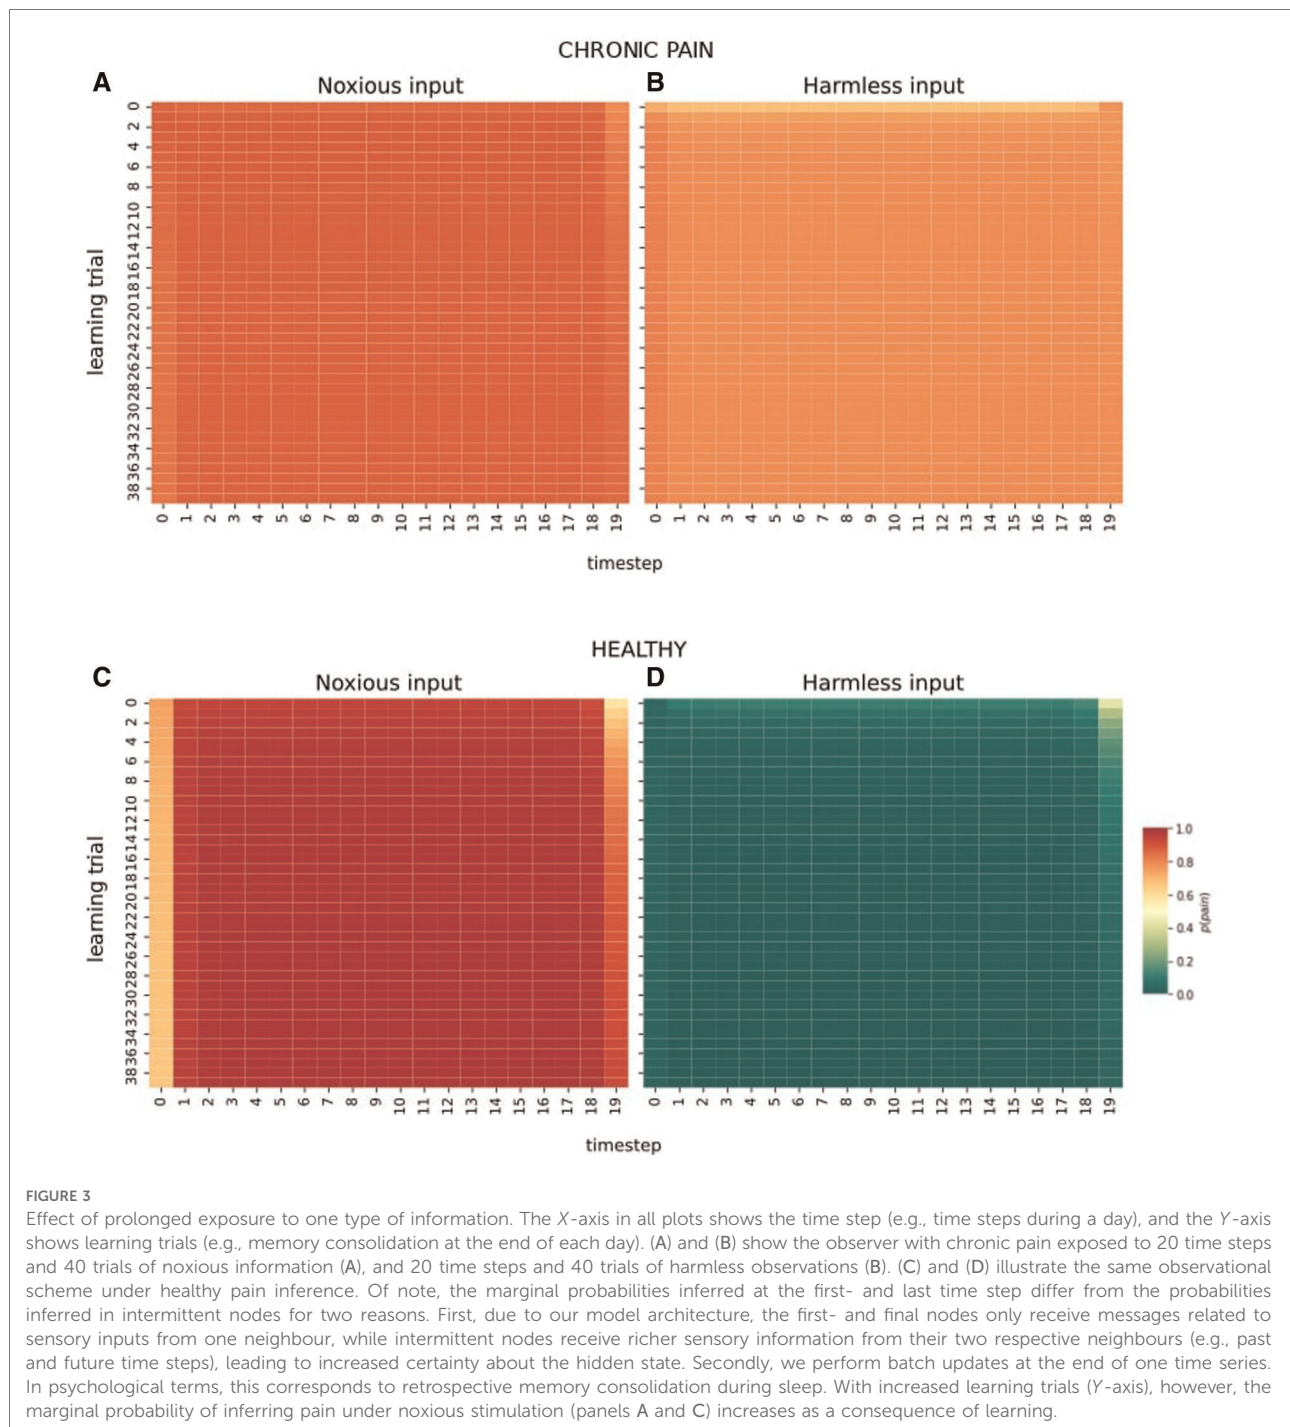


satisfy this condition. Resulting is a line, the one-dimensional null space, containing matching transition probabilities for a given marginal probability of pain. The combination of marginal and derived transition probabilities has a marked effect on the inferred state: it remains stable within the range of the preset marginal probability of pain, irrespective of the quality of observations made, or changes to the prior. In **Figure 4**, we illustrate the example of the marginal $p(\text{pain}) = 0.7$. First, as above, a model with 20 time steps and 40 learning trials is created. The prior $p(H_0 = \text{pain})$ on the first time step is chosen randomly. Crucially, we sample transition probabilities from the null space and then sample the inferred probability of pain at the last time step. It becomes evident that within a very short period, the inferred probability of pain stabilizes at the pre-defined marginal $p(\text{pain})$ of 0.7—regardless of any changes to the prior information.

4.1. Therapy for chronic pain

Exposure therapy is an effective treatment for chronic pain, where patients with high levels of fear-avoidance are gradually exposed to stimuli or movements that elicit fear (49). Relating the idea of exposure therapy to our model, its goal would translate to changing the interpretation of harmless sensory information so that the inference becomes less biased towards pain. As an example, a person with chronic low back pain may avoid lifting a basket with groceries out of their car. In graded exposure

therapy, the therapist would guide the patient through lifting baskets with increasing weight, all while closely monitoring the patient's fearful assumptions and physical sensations. Over the course of the exposure treatment, we would expect a change in $p(\text{lift} - \text{basket} | \text{pain})$ in this patient. Altered likelihoods may lead to less guarding behavior in patients, an increase in corrective sensory information and, ultimately, remission. In the specific case of fixed transition probabilities, however, this approach may not be sufficient. This picture would be in line with treatment-resistant chronic pain. The maintenance of $p(\text{pain})$ in these patients is mostly determined by the transition probabilities, while incoming sensory information is mostly disregarded. We simulated two observers who receive ten sessions of exposure therapy (presentation of harmless sensory information for one time step) over the course of fifty time steps. In both cases, the pre-set marginal probability of inferring pain is $p(\text{pain}) = 0.7$, and the transition probabilities were sampled from the respective null space of a marginal $p(\text{pain}) = 0.7$. Results are illustrated in **Figure 5**. In the case of the first patient, the likelihood functions, $p(\text{sensation} | \text{pain})$ were ambiguous ($p(\text{noxious} | \text{pain}) = 0.6$, $p(\text{harmless} | \text{pain}) = 0.4$, $p(\text{noxious} | \overline{\text{pain}}) = 0.6$, $p(\text{harmless} | \overline{\text{pain}}) = 0.4$, and imprecise ($v = 20$). In the simulation, the marginal probability of inferring pain dips slightly, before moving towards the initial, high $p(\text{pain}) = 0.7$ again. In a second case, the likelihood functions were intact, that is, accurate ($p(\text{noxious} | \text{pain}) = 0.8$, $p(\text{harmless} | \text{pain}) = 0.2$, $p(\text{noxious} | \overline{\text{pain}}) = 0.1$, $p(\text{harmless} | \overline{\text{pain}}) = 0.9$) and precise ($v = 100$). Here, the inference moves away from the high prior faster and remains relatively pain-free for longer.



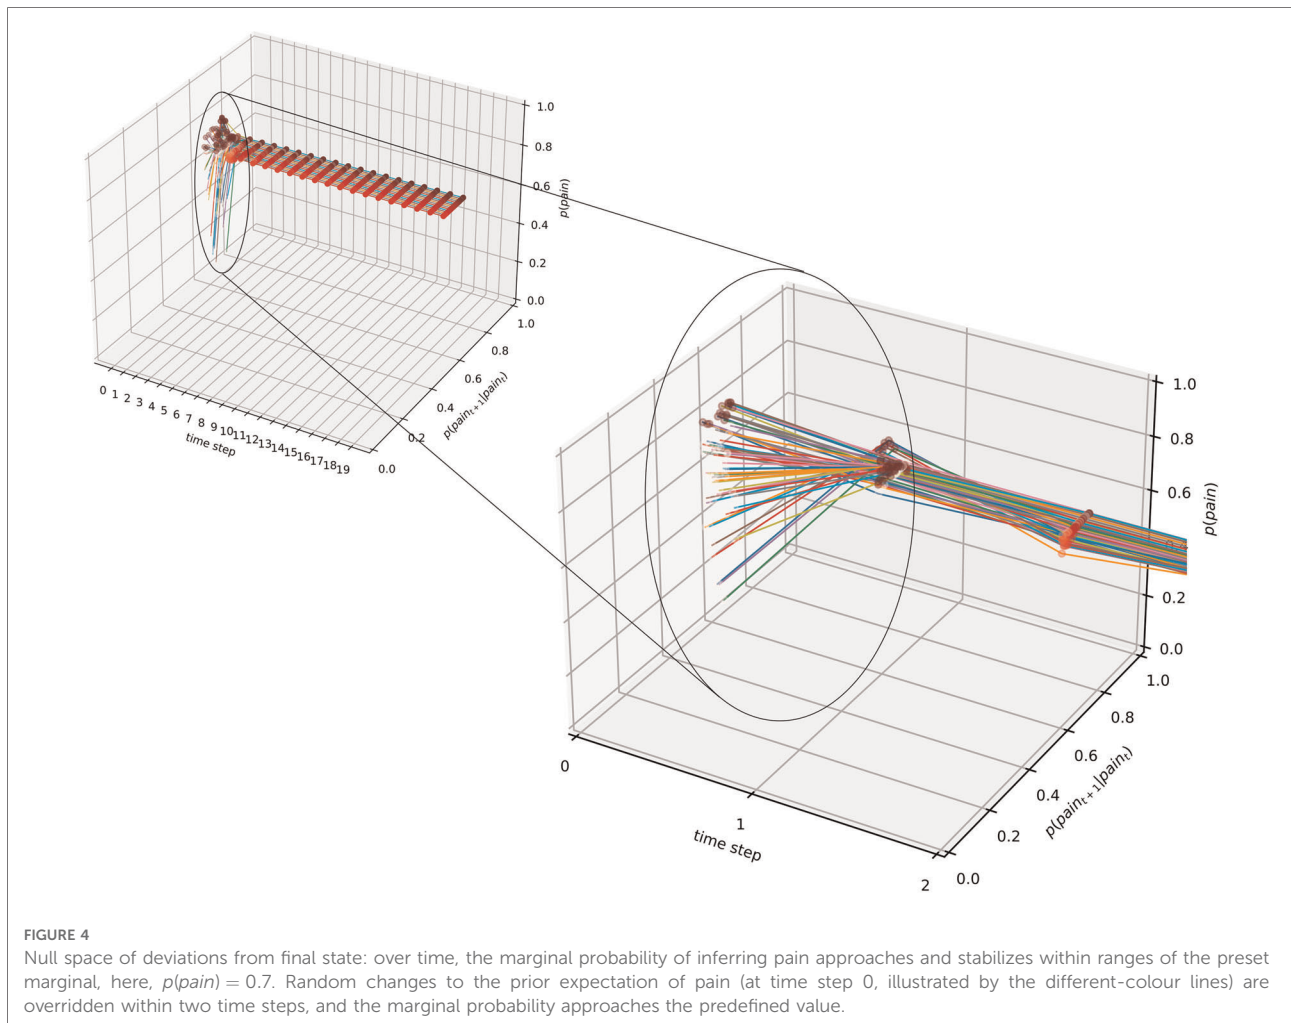
5. Discussion

5.1. A Bayesian model for chronic pain

The experience of pain seems to be the result of an inferential process, where prior expectations are integrated with current sensory information over time (14, 16). This process is probabilistic and non-linear, with large inter-individual

differences in the relationship between physical stimulation and pain experience. While statistical accounts of pain have gained momentum over the past decade, insights into computational underpinnings of the chronification of pain have been sparse (50, 18, 14).

We propose a Bayesian model for interoceptive perception and pain experience in chronic pain. Drawing on machine learning concepts such as belief-propagation and free-energy



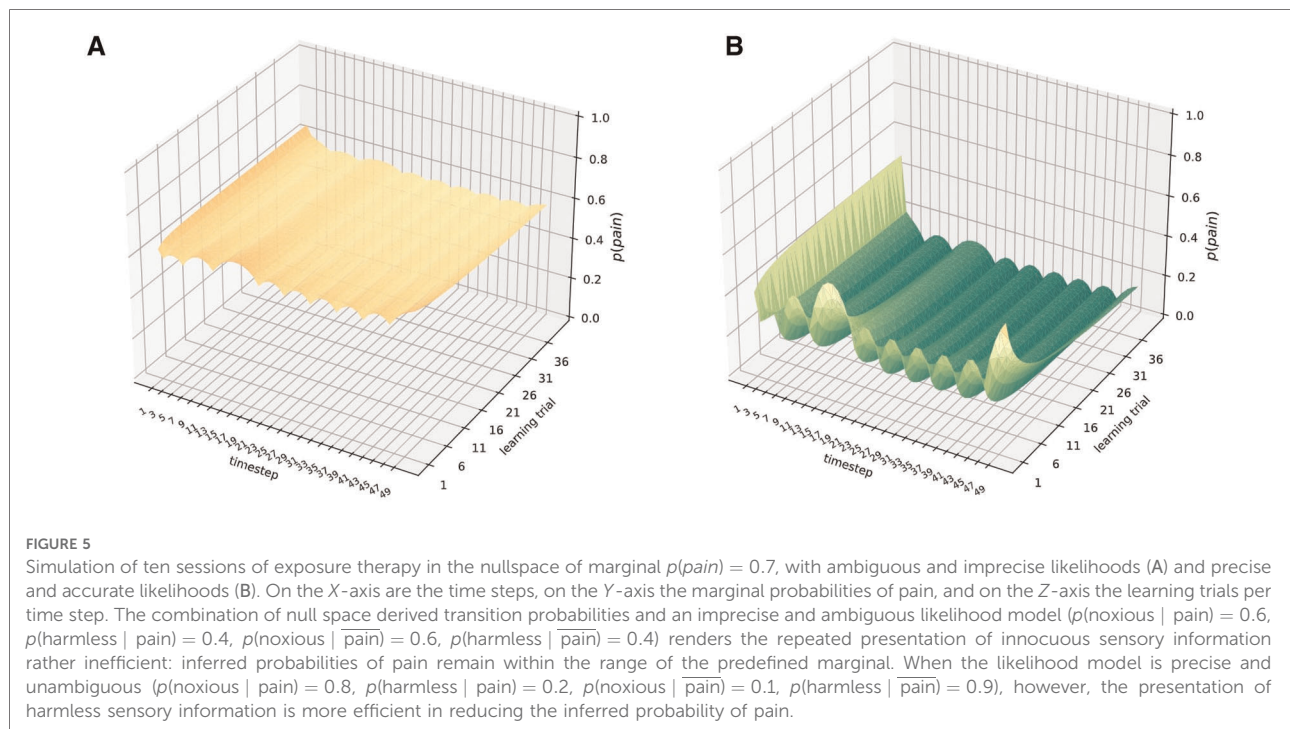
minimization, we demonstrate that the model captures well-studied phenomena such as decreased pain thresholds and allodynia (i.e. perceiving innocuous stimuli as painful) in individuals with chronic pain.

Our “chronic pain” observer is characterized by heightened prior expectations of pain, aberrant associations between sensory input and pain perception and heightened assumptions about the persistence of pain. In our model, this refers to overly precise priors for being in a state of pain, $p(\text{pain})$, ambiguous top-down or likelihood functions $p(\text{sensation} | \text{pain})$ and more rigid transition probabilities between hidden states $p(\text{pain}_t | \text{pain}_{t-1})$, respectively.

5.2. Chronic pain as biased inference

It has been suggested that chronic pain is associated with a broad and heightened expectation of pain (51, 14, 24). Heightened expectations of pain may stem from early aversive life events or injuries (52), or social factors such as an

overprotective, fearful parental style (14). In the patient model, a heightened pain prediction led to inferring to be in pain more readily compared to the healthy model. Conflicting innocuous sensory stimulation hardly influenced the inferred state. In accordance with this, empirical studies suggest chronic pain patients may not attend to sensory information but rather rely on their prior expectations of pain (53, 54). Our observations fit into a larger discussion about expectations as core features of a multitude of mental disorders (51, 14). A further aspect of our model is that harmless sensory information is associated with pain over time. This phenomenon is commonly observed in patients and referred to as allodynia (55, 14). The Bayesian Brain hypothesis postulates that the perceiving individual perpetually tries to infer the real-world causes for incoming sensory information (22). We here present a perspective on pain as an abstract percept of the body’s state. Following this view, pain is an inference on a real-world “cause” (cmp. (56)), bodily damage, based on all available prior- and current sensory information. The chronic pain patient infers pain



from a wide range of sensory stimuli that would not be perceived as painful by healthy controls. Further, the associations between sensory information and hidden state are noisy and unclear in the chronic pain observer. While a healthy observer maintains precise beliefs about the external causes of harmless sensory input, $p(\text{harmless} | \text{pain})$, this probability is completely ambiguous in our patient model. Over time, the chronic pain patient is less and less likely to infer a pain-free state from any type of incoming sensation.

Contrarily, a flexible and highly dynamic pattern can be observed under healthy interoceptive inference. Here, noxious stimuli are associated with pain, and a pronounced and acute state of pain is inferred from incoming harmful stimuli. However, once the stream of harmful sensory input stops, the probability of pain normalizes very quickly. This dynamic pattern is in line with the function of pain as a warning signal—motivating quick and adaptive behavior change following a threat to the individual’s physical integrity (57).

Note that our results also hold under the active inference framework (23, 58). Canonical versions of active inference contain a generative model that is very similar to ours (Hidden Markov model, HMM), with the difference of an additional layer that models an agent’s actions, or choices (59). Actions then influence the transitions between states so that more likely, or expected states under the current observations, are reached. In other words, in active inference, an agent infers the hidden state from the combination of prior beliefs, sensory information and their own actions. As an example, consider the case of a child with chronic

abdominal pain who involuntarily keeps picking at her abdomen (checking behavior, example from (14)), so that increased pain is experienced. Under active inference, this can be interpreted as causing observations (noxious sensations) that are expected under the most likely state (pain). Acting directly towards bringing about a certain observation changes the causal relationship between observation and hidden state (i.e., “do-operations”, (60)). However, it is impossible to act directly on one’s nociceptors, so that pain still needs to be inferred from sensations, even when the noxious stimulation is self-induced.

5.3. Exposure and the null space of psychotherapy

If the individual performs implicit Bayesian inference over hidden states, then we here demonstrate that this inference can be biased towards inferring pain. This has important implications for the treatment of chronic pain patients. Cognitive behavioral therapy (CBT) shows statistically significant, but small effects on certain aspects of chronic pain (pain, mood, disability, pain catastrophizing; (61)). Techniques usually include relaxation training, behavioral activation (i.e. engaging in behaviors that were previously avoided due to fear of pain), setting behavioral goals, problem-solving training and cognitive re-structuring (61, 62). Further, exposure therapy has shown promising results in the treatment of chronic pain (63, 64, 49). Exposure therapy

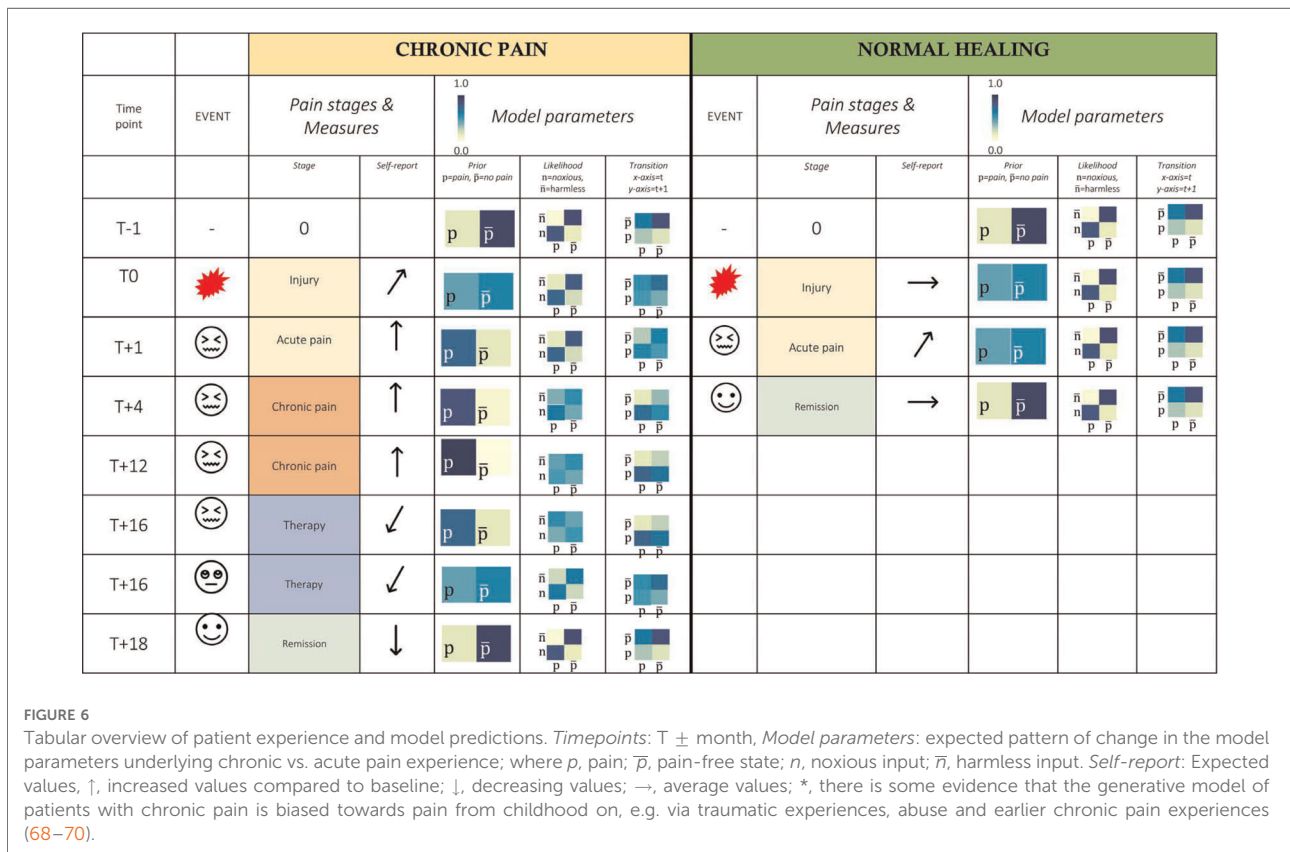
seems to reduce disability more effectively than CBT in patients with chronic lower back pain (49). It is designed to specifically target emotional responses to pain- or expectations of pain, which lead to excessive guarding behaviors. Guarding behaviors, in turn, may limit the opportunity for receiving corrective, conflicting sensory information, leading to the solidification of fearful pain beliefs (65, 66). However, we show that inference can stabilize towards pain, even in light of conflicting sensory information. Specifically, when the transition probabilities between hidden states have a stable state with high marginal pain probability, a simulated patient model returns to the inference of being in pain while harmless sensory information is provided. In other words, the initial deviation from the high pain state, which might have been reached by therapeutic means, has vanished. This simulation could be a model for treatment-resistant chronic pain. We further show that these patients may be best served by targeting their beliefs about the associations between sensory stimuli and pain expectations (likelihoods). When likelihoods allow the precise and accurate association between sensory stimuli and pain, exposure to harmless sensory information leads to more promising results. However, when the patient does not have accurate likelihoods, exposure to certain sensory stimuli may not be a promising approach. It can hence be derived that patients may need to re-learn precise and accurate inferences about different types of

sensory stimuli, ideally before being guided through exposure therapy.

One can speculate about the role of pharmacotherapy in the context of our model. Potentially, analgesics target the patient's likelihood model by inhibiting the transmission of noxious information from pain receptors (67). This assumption requires empirical testing with appropriate data of medicated chronic pain patients. In this case, combining pharmacological and psycho-therapeutical strategies may be promising. More data is necessary to assess which treatments target what model parameter. For an overview of expected self-report scores and hypothesized underlying model parameters, see Figure 6.

In Figure 6, we illustrate two hypothetical patients, one developing chronic pain after an acute injury, whereas the other patient does not develop chronic pain beyond acute injury.

In the chronic pain patient, pain persists for prolonged periods post injury. Empirically, we expect to see pain and disability scores that are persistently heightened in the months following injury. In experiments (i.e. fear conditioning, pain ratings following noxious stimuli), our model predicts increased levels of prior expectations of pain (note the shift in probabilities over time), as well as decreased precision of sensory information, corresponding to an ambiguous and imprecise likelihood model. Finally, for pain to persist in a very rigid way, transitions towards a pain-free state become



more unlikely until therapy can reverse this effect. We can speculate that likelihood models may be treated via medication-supported exposure therapy, whereas prior beliefs could be best targeted using CBT.

In the acute pain patient, the injury is followed by a quick recovery and remission into a pain-free state. The noxious stimulus leads to a transient state of experiencing pain, including an elevated prior expectation of pain, which is a consequence of learning about an acutely noxious stimulus. Disability and pain experience as assessed with self-report measures are elevated in times of acute injury. With the injury healing and the noxious information subsiding, the pain-free state is re-established within 4 months. Note that the likelihood and transition models are assumed to stay relatively intact in this patient.

5.4. Expectations as primary therapy target

Following our model, there are three factors that could be promising targets for patients with chronic pain. Firstly, therapy could target the heightened and generalized expectations of pain. This is in line with Rief et al. (51), Panitz et al. (71), who discuss the importance of psychological interventions that target expectations. Secondly, our model suggests a need for re-learning the associations between sensory information and the state of pain. This could be achieved via gradual, guided exposure to harmless stimuli in therapy. Lastly, a patient's belief about the persistence of pain over time could be targeted specifically with cognitive interventions (e.g. retrospective evaluations, re-learning to differentiate pain intensities), which may be especially relevant in cases of treatment-resistant chronic pain.

5.5. Limitations

The presented model makes several simplifying assumptions and can be extended in numerous ways. Currently, latent- and observable variables could only take one of two possible states for the sake of simplicity. Continuous random variables could allow for more complex conclusions. It is generally feasible to perform inference on graphical models with continuous random variable nodes (72, 73). However, this added complexity might detract from our main conclusions, which we expect to hold in the continuous case, too. In future work, we will extend our model to continuous sensory signals and internal states.

Further, future efforts to extend this framework should account for the vast heterogeneity that is found among chronic pain patients. Assuming a multitude of underlying etiological factors, they might need differential representation in quantitative terms as well.

With this model, we focus on computational and psychological aspects of pain inference. We have formalized a hypothesis about the generative process underlying chronic pain. While several salient features of acute and chronic pain perception are captured by our model, we do not have access to the true generative process and can hence not draw any definite conclusions about it.

Besides psychological factors (e.g. beliefs), a multitude of additional factors could be considered when extending this model. As Kiverstein and colleagues (2022) note, a model of chronic pain needs to take on a bio-psycho-social perspective, integrating a multitude of empirical findings under one approach. For example, pain experience seems to be alleviated by the presence of a close friend or partner, i.e. social support (74). In a recent review, Mohr and Fotopoulou (2018) argue that social support may serve as a security signal which renders prediction errors caused by threats to the physical integrity less precise. Future iterations of this model may be extended to incorporate these findings, e.g. by means of a context-layer. Here, however, we show that important characteristics of chronic pain experience can be simulated from a Bayesian perspective, by assuming that a) sensory information relates to pain experience through (biased) parameters of a generative model and b) the inferred states are marked by correlations, or stability, over time.

5.6. Future directions

Rigid expectations that are not sufficiently constrained by incoming sensory information seem to lie at the heart of several mental disorders. Psychosis, for instance, can be regarded as resulting from an exaggerated and inflexible reliance on prior expectations (75, 76). Computational accounts of mental disorders can provide novel mechanistic insights into their etiology, exacerbation and maintenance over time (77).

The computational model of chronic pain presented here can be applied to empirical data. Below, we outline possible avenues towards testing the model presented here.

Firstly, data from learning experiments in chronic pain patients and controls could shed light on the real-world significance of the model's parameters. In fear conditioning experiments, participants learn the association between a painful unconditioned stimulus (UCS, such as e.g. heat stimulation with a thermode or electric shocks) and unrelated, conditioned stimuli (CS, such as a visual cue). In a first step, significant group differences in model parameters would need to be investigated. Then, given their heightened pain prediction, our model would predict decreased pain thresholds in patients with chronic pain. Biased likelihood terms may lead to altered learning about the UCS-CS contingencies in chronic pain patients. Similarly, other

assumptions that were used to furnish this model could be tested empirically (e.g. biased and imprecise likelihood models, heightened prior predictions of pain, or rigid and exaggerated beliefs about the persistence of pain in chronic pain patients).

A key challenge to empirical tests of our model in humans is the need to approximate observable sensations, i.e., nociceptor firing. Sensations can be approximated by measuring physical properties of experimental stimuli, such as the temperature of a heat thermode, the current used for electric stimuli, or the intensity of mechanical pressure. To empirically test our model, future studies could assess the relationship between physical sensations and the pain percept. For this, concurrent measurements of sensations (or proxies thereof, such as physical stimulus properties) and reports on the pain percept (e.g. pain scores) are necessary. Pain scores in response to stimuli with different intensities could be collected from healthy individuals and chronic pain patients. We would expect to find significant group differences with respect to prior, transition- and likelihood model parameters. Such data could further shed light on an important detail of patient experience: what model parameter is most relevant to the chronification of pain and the experience of patients? In all cases, the full Bayesian model would need to be compared to both simpler models, such as e.g. single-trial vs. sequential models, and more complex models, such as deeper predictive hierarchies, in a model comparison.

5.7. Conclusion

A novel Bayesian model is able to reproduce common features of healthy- and chronified pain perception. Machine learning approaches, such as a hidden Markov model and variational inference, allowed the exploration of parameters that may underlie the development of chronic pain. We have demonstrated that this model is able to capture many phenomena observed in individuals with chronic pain, and during acute, healthy pain perception. Additionally, the model was able to simulate rigid, treatment resistant chronic pain. Further research could significantly advance our mechanistic understanding of chronic pain, which may help to inform treatment selection.

References

1. IASP. Iasp terminology. (2017). Available at: <https://www.iasp-pain.org/Education/Content.aspx?ItemNumber=1698navItemNumb>
2. King S, Chambers CT, Huguet A, MacNevin RC, McGrath PJ, Parker L, et al. The epidemiology of chronic pain in children, adolescents revisited: a systematic review. *Pain* (2011) 152:2729–38. doi: 10.1016/j.pain.2011.07.016
3. Breivik H, Collett B, Ventafridda V, Cohen R, Gallacher D. Survey of chronic pain in Europe: prevalence, impact on daily life, and treatment. *Eur J Pain* (2006) 10:287–333. doi: 10.1016/j.ejpain.2005.06.009
4. Lamé IE, Peters ML, Vlaeyen JW, Kleef MV, Patijn J. Quality of life in chronic pain is more associated with beliefs about pain, than

Data availability statement

All code used in this manuscript is publicly available here: <https://data.uni-marburg.de/handle/dataumr/163>.

Ethics Statement

Ethical review and approval was not required for this study in accordance with the local legislation and institutional requirements.

Author contributions

ALE and DE conducted the first set of analyses, further developed by KP, DE and ALE. ALE, KP and DE produced the final results presented here. ALE wrote a first draft of this manuscript. All authors reviewed and edited the manuscript and approved the submitted version.

Funding

This work was supported by the cluster project “The Adaptive Mind”, funded by the Excellence Program of the Hessian Ministry for Science and the Arts and the DFG SFB-TRR289 “Treatment Expectation”, project number 422744262.

Conflict of interest statement

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher’s note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- with pain intensity. *Eur J Pain* (2005) 9:15–24. doi: 10.1016/j.ejpain.2004.02.006
5. Hunfeld JA, Perquin CW, Duivenvoorden HJ, Hazebroek-Kampschreur AA, Passchier J, van Suijlekom-Smit LW, et al. Chronic pain and its impact on quality of life in adolescents and their families. *J Pediatr Psychol* (2001) 26:145–53. doi: 10.1093/jpepsy/26.3.145
 6. Tunks ER, Crook J, Weir R. Epidemiology of chronic pain with psychological comorbidity: prevalence, risk, course, and prognosis. *Can J Psychiatry* (2008) 53:224–34. doi: 10.1177/070674370805300403
 7. Mäntyselkä P, Kumpusalo E, Ahonen R, Kumpusalo A, Kauhanen J, Viinamäki H, et al. Pain as a reason to visit the doctor: a study in Finnish primary health care. *Pain* (2001) 89:175–80. doi: 10.1016/s0304-3959(00)00361-4
 8. Gaskin DJ, Richard P. The economic costs of pain in the united states. *J Pain* (2012) 13:715–24. doi: 10.1016/j.jpain.2012.03.009
 9. Phillips CJ. Economic burden of chronic pain. *Expert Rev Pharmacoecon Outcomes Res* (2006) 6:591–601. doi: 10.1586/14737167.6.5.591
 10. Groenewald CB, Essner BS, Wright D, Fesinmeyer MD, Palermo TM. The economic costs of chronic pain among a cohort of treatment-seeking adolescents in the United States. *J Pain* (2014) 15:925–33. doi: 10.1016/j.jpain.2014.06.002
 11. Frebarger JK, Holmes GM, Agans RP, Jackman AM, Darter JD, Wallace AS, et al. The rising prevalence of chronic low back pain. *Arch Intern Med* (2009) 169:251–8. doi: 10.1001/archinternmed.2008.543
 12. Sá KN, Moreira L, Baptista AF, Yeng LT, Teixeira MJ, Galhardoni R, et al. Prevalence of chronic pain in developing countries: systematic review and meta-analysis. *Pain Rep* (2019) 4:1–8. doi: 10.1097/PR9.0000000000000779
 13. Marr D. *Vision: a computational investigation into the human representation and processing of visual information*, Cambridge, Mass.: MIT Press (1982).
 14. Hechler T, Endres DM, Thorwart A. Why harmless sensations might hurt in individuals with chronic pain: about heightened prediction and perception of pain in the mind. *Front Psychol* (2016) 7:1638. doi: 10.3389/fpsyg.2016.01638
 15. Büchel C, Geuter S, Sprenger C, Eippert F. Placebo analgesia: a predictive coding perspective. *Neuron* (2014) 81:1223–39. doi: 10.1016/j.neuron.2014.02.042
 16. Tabor A, Thacker MA, Moseley GL, Körding KP. Pain: a statistical account. *PLoS Comput Biol* (2017) 13:e1005142. doi: 10.1371/journal.pcbi.1005142
 17. Wall P, McMahon S. The relationship of perceived pain to afferent nerve impulses. *Trends Neurosci* (1986) 9(6):254–5. doi: 10.1016/0166-2236(86)90070-6
 18. Kiverstein J, Kirchhoff MD, Thacker M. An embodied predictive processing theory of pain experience. *Rev Philos Psychol* (2022):11–26. doi: 10.1007/s13164-022-00616-2
 19. Dayan P, Hinton GE, Neal RM, Zemel RS. The helmholtz machine. *Neural Comput* (1995) 7:889–904. doi: 10.1162/neco.1995.7.5.889
 20. Friston K, Kiebel S. Predictive coding under the free-energy principle. *Philos Trans R Soc Lond B Biol Sci* (2009) 364:1211–21. doi: 10.1098/rstb.2008.0300
 21. Barrett LF, Simmons WK. Interoceptive predictions in the brain. *Nat Rev Neurosci* (2015) 16:419–29. doi: 10.1038/nrn3950
 22. Friston K. The free-energy principle: a unified brain theory?. *Nat Rev Neurosci* (2010) 11:127–38. doi: 10.1038/nrn2787
 23. Friston K, Daunizeau J, Kiebel S. Active inference or reinforcement learning. *PLoS ONE* (2009) 4:e6421. doi: 10.1371/journal.pone.0006421
 24. Fairhurst M, Wiech K, Dunckley P, Tracey I. Anticipatory brainstem activity predicts neural processing of pain in humans. *Pain* (2007) 128:101–10. doi: 10.1016/j.pain.2006.09.001
 25. Wiech K, Vandekerckhove J, Zaman J, Tuerlinckx F, Vlaeyen JW, Tracey I. Influence of prior information on pain involves biased perceptual decision-making. *Curr Biol* (2014) 24:R679–81. doi: 10.1016/j.cub.2014.06.022
 26. Wiech K. Deconstructing the sensation of pain: the influence of cognitive processes on pain perception. *Science* (2016) 354:584–7. doi: 10.1126/science.aaf8934
 27. Hird EJ, Jones A, Talmi D, El-Deredy W. A comparison between the neural correlates of laser, electric pain stimulation, their modulation by expectation. *J Neurosci Methods* (2018) 293:117–27. doi: 10.1016/j.jneumeth.2017.09.011
 28. Gligorov N. Telling the truth about pain: informed consent and the role of expectation in pain intensity. *AJOB Neurosci* (2018) 9:173–82. doi: 10.1080/21507740.2018.1496163
 29. Anchisi D, Zanon M. A Bayesian perspective on sensory and cognitive integration in pain perception and placebo analgesia. *PLoS ONE* (2015) 10:e0117270. doi: 10.1371/journal.pone.0117270
 30. Ongaro G, Kaptchuk TJ. Symptom perception, placebo effects, and the Bayesian brain. *Pain* (2019) 160:1. doi: 10.1097/j.pain.0000000000001367
 31. Imamura M, Chen J, Matsubayashi SR, Targino RA, Alfieri FM, Bueno DK, et al. Changes in pressure pain threshold in patients with chronic nonspecific low back pain. *Spine* (2013) 38:2098–107. doi: 10.1097/01.brs.0000435027.50317.d7
 32. Friston K, Kilner J, Harrison L. A free energy principle for the brain. *J Physiol Paris* (2006) 100:70–87. doi: 10.1016/j.jphysparis.2006.10.001
 33. Endres DM, Pabst K, Eckert A-L, Schween R. A collection of identities for variational inference with exponential-family models. (2022). doi: 10.17192/ldr/112
 34. Fillingim RB. Individual differences in pain: understanding the mosaic that makes pain personal. *Pain* (2017) 158:S11. doi: 10.1097/j.pain.0000000000000775
 35. Fine S, Singer Y, Tishby N. The hierarchical hidden Markov model: analysis, applications. *Mach Learn* (1998) 32:41–62. doi: 10.1023/A:1007469218079
 36. Bishop CM. *Pattern recognition, machine learning*, New York: Springer (2006).
 37. Schnitzler A, Ploner M. Neurophysiology and functional neuroanatomy of pain perception. *J Clin Neurophysiol* (2000) 17:592–603. doi: 10.1097/00004691-200011000-00005
 38. Hohwy J. *The predictive mind*. Oxford: Oxford University Press (2013).
 39. Pearl J. *Reverend Bayes on inference engines: a distributed hierarchical approach*. Menlo Park: Cognitive Systems Laboratory, School of Engineering, Applied Science (1982).
 40. Friston K. A theory of cortical responses. *Philos Trans R Soc Lond B Biol Sci* (2005) 360:815–36. doi: 10.1098/rstb.2005.1622
 41. Harris CR, Millman KJ, van der Walt SJ, Gommers R, Virtanen P, Cournapeau D, et al. Array programming with NumPy. *Nature* (2020) 585:357–62. doi: 10.1038/s41586-020-2649-2
 42. Hunter JD. Matplotlib: a 2D graphics environment. *Comput Sci Eng* (2007) 9:90–5. doi: 10.1109/MCSE.2007.55
 43. Borsook D, Youssef AM, Simons L, Elman I, Eccleston C. When pain gets stuck: the evolution of pain chronification, treatment resistance. *Pain* (2018) 159:2421. doi: 10.1097/j.pain.0000000000001401
 44. Morrone L, Scuteri D, Rombola L, Mizoguchi H, Bagetta G. Opioids resistance in chronic pain management. *Curr Neuropharmacol* (2017) 15:444–56. doi: 10.2174/1570159X14666161101092822
 45. Lefaucheur J-P, Nguyen J-P. A practical algorithm for using rTMS to treat patients with chronic pain. *Neurophysiol Clin* (2019) 49:301–7. doi: 10.1016/j.neucli.2019.07.014
 46. Suda S, Takagai S, Inoshima-Takahashi K, Sugihara G, Mori N, Takei N. Electroconvulsive therapy for burning mouth syndrome. *Acta Psychiatr Scand* (2008) 118:503–4. doi: 10.1111/j.1600-0447.2008.01261.x
 47. Pickering G, Pereira B, Morel V, Corriger A, Giron F, Marcaillou F, et al. Ketamine and magnesium for refractory neuropathic pain: a randomized, double-blind, crossover trial. *Anesthesiology* (2020) 133:154–64. doi: 10.1097/ALN.0000000000003345
 48. Greub WH. *Linear algebra*. Vol. 23. New York: Springer Science & Business Media (2012).
 49. Glombiewski JA, Holzapfel S, Riecke J, Vlaeyen JW, de Jong J, Lemmer G, et al. Exposure and CBT for chronic back pain: an RCT on differential efficacy and optimal length of treatment. *J Consult Clin Psychol* (2018) 86:533. doi: 10.1037/ccp0000298
 50. Tabor A, Burr C. Bayesian learning models of pain: a call to action. *Curr Opin Behav Sci* (2019) 26:54–61. doi: 10.1016/j.cobeha.2018.10.006
 51. Rief W, Glombiewski JA, Gollwitzer M, Schubö A, Schwarting R, Thorwart A. Expectancies as core features of mental disorders. *Curr Opin Psychiatry* (2015) 28:378–85. doi: 10.1097/YCO.0000000000000184
 52. Victoria NC, Murphy AZ. Exposure to early life pain: long term consequences and contributing mechanisms. *Curr Opin Behav Sci* (2016) 7:61–8. doi: 10.1016/j.cobeha.2015.11.015
 53. Tracey I. Getting the pain you expect: mechanisms of placebo, nocebo and reappraisal effects in humans. *Nat Med* (2010) 16:1277–83. doi: 10.1038/nm.2229
 54. Edwards MJ, Adams RA, Brown H, Parees I, Friston KJ. A Bayesian account of “hysteria”. *Brain* (2012) 135:3495–512. doi: 10.1093/brain/aww129
 55. Flor H, Turk D. *Chronic pain: an integrated biobehavioral approach*. Washington: IASP Press (2015).
 56. Hohwy J. *The predictive mind*. Aalphen an de Rijn: Oxford University Press (2013).
 57. Wiech K, Tracey I. Pain, decisions, and actions: a motivational perspective. *Front Neurosci* (2013) 7:46. doi: 10.3389/fnins.2013.00046

58. Friston K, FitzGerald T, Rigoli F, Schwartenbeck P, Pezzulo G. Active inference: a process theory. *Neural Comput* (2017) 29:1–49. doi: 10.1162/NECO_a_00912
59. Heins C, Millidge B, Demekas D, Klein B, Friston K, Couzin I, et al. pymdp: a Python library for active inference in discrete state spaces. [Preprint] (2022). Available at: arXiv:2201.03904.
60. Pearl J. *Causality*. Cambridge: Cambridge University Press (2009).
61. Ehde DM, Dillworth TM, Turner JA. Cognitive-behavioral therapy for individuals with chronic pain: efficacy, innovations, and directions for research. *Am Psychol* (2014) 69:153. doi: 10.1037/a0035747
62. Beehler G, Murphy JL, King PR, Dollar KM. Brief cognitive behavioral therapy for chronic pain: therapist manual. US Department of Veterans Affairs.
63. Vlaeyen JW, Morley S, Linton SJ, Boersma K, de Jong J. *Pain-related fear: exposure-based treatment of chronic pain*. Washington, D.C.: IASP Press (2012).
64. Trost Z, Parsons TD. Beyond distraction: virtual reality graded exposure therapy as treatment for pain-related fear and disability in chronic pain. *J Appl Biobehav Res* (2014) 19:106–26. doi: 10.1111/jabr.12021
65. Vlaeyen JW, Kole-Snijders AM, Rotteveel AM, Ruesink R, Heuts PH. The role of fear of movement/(re) injury in pain disability. *J Occup Rehabil* (1995) 5:235–52. doi: 10.1007/BF02109988
66. Vlaeyen JW, Crombez G, Linton SJ. The fear-avoidance model of pain. *Pain* (2016) 157:1588–89. doi: 10.1097/j.pain.0000000000000574
67. Finnerup NB, Sindrup SH, Jensen TS. Chronic neuropathic pain: mechanisms, drug targets and measurement. *Fundam Clin Pharmacol* (2007) 21:129–36. doi: 10.1111/j.1472-8206.2007.00474.x
68. Davis DA, Luecken LJ, Zautra AJ. Are reports of childhood abuse related to the experience of chronic pain in adulthood?: a meta-analytic review of the literature. *Clin J Pain* (2005) 21:398–405. doi: 10.1097/01.ajp.0000149795.08746.31
69. Lampe A, Doering S, Rumpold G, Sölder E, Krismser M, Kantner-Rumplmair W, et al. Chronic pain syndromes and their relation to childhood abuse and stressful life events. *J Psychosom Res* (2003) 54:361–7. doi: 10.1016/S0022-3999(02)00399-9
70. Goldberg RT, Pachasoe WN, Keith D. Relationship between traumatic events in childhood and chronic pain. *Disabil Rehabil* (1999) 21:23–30. doi: 10.1080/096382899298061
71. Panitz C, Endres DM, Buchholz M, Khosrowtaj Z, Sperl MF, Mueller EM, et al. A revised framework for the investigation of expectation update versus maintenance in the context of expectation violations: the violex 2.0 model. *Front Psychol* (2021) 13(11):5237. doi: 10.3389/fpsyg.2021.726432
72. Noorshams N, Wainwright MJ. Belief propagation for continuous state spaces: stochastic message-passing with quantitative guarantees. *J Mach Learn Res* (2013) 14:2799–835.
73. Friston KJ, Parr T, de Vries B. The graphical brain: belief propagation and active inference. *Netw Neurosci* (2017) 1:381–414. doi: 10.1162/NETN_a_00018
74. von Mohr M, Krahé C, Beck B, Fotopoulou A. The social buffering of pain by affective touch: a laser-evoked potential study in romantic couples. *Soc Cogn Affect Neurosci* (2018) 13:1121–30. doi: 10.1093/scan/nsy085
75. Fletcher PC, Frith CD. Perceiving is believing: a Bayesian approach to explaining the positive symptoms of schizophrenia. *Nat Rev Neurosci* (2009) 10:48–58. doi: 10.1038/nrn2536
76. Sterzer P, Adams RA, Fletcher P, Frith C, Lawrie SM, Muckli L, et al. The predictive coding account of psychosis. *Biol Psychiatry* (2018) 84:634–43. doi: 10.1016/j.biopsych.2018.05.015
77. Huys QJ, Browning M, Paulus MP, Frank MJ. Advances in the computational understanding of mental illness. *Neuropsychopharmacology* (2021) 46:3–19. doi: 10.1038/s41386-020-0746-4
78. Friston K, Mattout J, Trujillo-Barreto N, Ashburner J, Penny W. Variational free energy and the laplace approximation. *Neuroimage* (2007) 34:220–34. doi: 10.1016/j.neuroimage.2006.08.035
79. Parr T, Friston KJ. Generalised free energy and active inference. *Biol Cybern* (2019) 113:495–513. doi: 10.1007/s00422-019-00805-w
80. Yang X. Understanding the variational lower bound. In: *variational lower bound, ELBO, hard attention*. (2017) p. 1–4.
81. Kullback S, Leibler RA. On information and sufficiency. *Ann Math Stat* (1951) 22:79–86. doi: 10.1214/aoms/1177729694
82. Jensen JLWV. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Math* (1906) 30:175–93. doi: 10.1007/BF02418571
83. Kolmogorov AN, Bharucha-Reid AT. *Foundations of the theory of probability*. 2nd English ed. Mineola, N.Y.: Courier Dover Publications (2018).

Appendix A. Variational inference

Graphical models are a useful tool when representing and manipulating joint probability distributions. In the present context, we are interested in characterizing the mind’s probability computations of being in pain, given a series of noxious or harmless observations, i.e. *inference* (36). For this purpose, we use a modified state-space model architecture, i.e. a Hidden Markov Model (36). It features a chain of unobservable nodes, representing the inferred pain state at different time points. Each unobservable node is connected to one observable node representing sensations at the same time point (see **Figure A1**).

Assuming that hidden state inference evolves in an auto-correlated manner over time seems plausible because pain states tend to persist at least for some time. For this reason, a Markov chain of hidden states was selected as the basis of the model. The Markov property states that inference at a given future node, S_{t+1} , is conditionally independent (\perp) from all but the previous states S_t :

$$S_{t+1} \perp S_{t-1} \mid S_t \tag{A1}$$

In other words, the future state S_{t+1} does not depend on the distribution of the past state S_{t-1} given the current state S_t .

It follows that the conditional distribution for the current state S_t given *all* previous S is given by only its predecessor’s state S_{t-1} :

$$p(S_t \mid S_1, \dots, S_{t-1}) = p(S_t \mid S_{t-1}) \tag{A2}$$

The main inference goal is the estimation of the hidden random variables’ H_t state, given a series of observations. This is expressed in the random variable’s respective marginal probabilities. Computing marginal probabilities is challenging, because it usually entails computing the sums of integrals over high-dimensional spaces. There are a number of widely used algorithms for efficient inference on graphical models. However, to perform inference with message-passing algorithms, we first need to transform the Bayesian network to a factor graph. A graphical representation of the factor graph is illustrated in **Figure A1** (36). A factor graph represents the factorization of a function, or here, the probability distributions of interest. In other words, it provides a graphical representation of how probabilities of certain events depend on adjacent events in the observable and hidden layers. A factor graph consists of variable- and factor nodes, denoted by circles and squares, respectively. In our model, variable nodes can take on a binary set of values

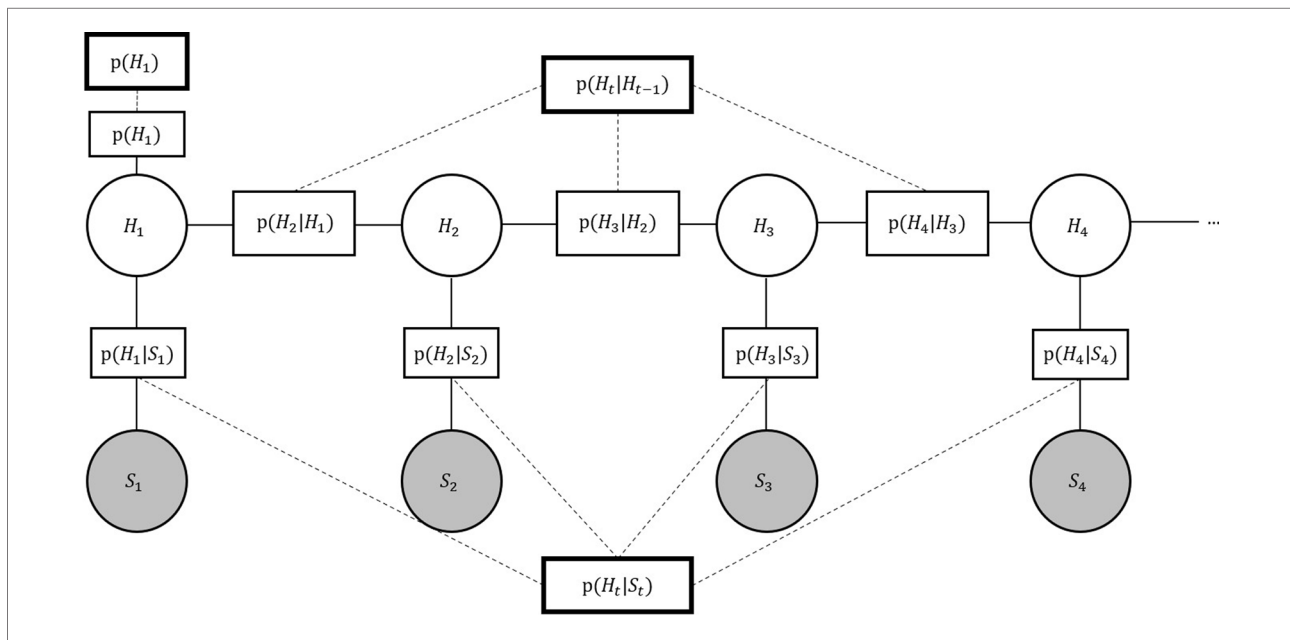


FIGURE A1
 Graphical model (fragment). Graphical model depicting the inter-dependencies between hidden H_t and observable S_t variables. Connected to each variable node, factor nodes are necessary to perform variational inference message passing. This approach is computationally efficient because it allows message passing without re-instantiating the model at each time-step. The observable nodes S_t represent incoming sensory information, $S \in \{\text{noxious}(1), \text{harmless}(0)\}$. The hidden states, representing the system’s internal model, can in turn take on the values 0 (no pain) or 1 (pain), or $H \in \{\text{pain}, \overline{\text{pain}}\}$. The likelihood $p(S_t \mid H_t)$ is the top-down factor in this model and allows quantifying the association between the model’s state H_t and the observed sensory input at S_t . The hidden variable nodes form a Markov chain and contain transition probabilities $p(H_t \mid H_{t+1})$, which can be interpreted as the development of the pain prior over time. In individuals with chronic pain, we here assume that $p(H_t = \text{pain} \mid H_{t-1} = \text{pain}) > p(H_t = \overline{\text{pain}} \mid H_{t-1} = \text{pain})$.

(here; *pain*, $\overline{\text{pain}}$ for the hidden nodes, and *noxious*, *harmless* input in the observable nodes). They can be observed or unobserved. Factor nodes contain exponential family distributions, parameterized by their natural parameters. These parameters represent probabilities for either transitioning between different hidden states or for observing certain types of sensory information, given a specific latent model state. We here further introduce hypernodes, denoted by thick squares, that govern the learning of factor-inherent natural parameters over time. All nodes are connected via edges to their respective neighbour nodes.

Once transformed, it is now appropriate to apply inference algorithms to the factorized graphical model. Sum-product message passing, also known as belief propagation, is a well-known algorithm that allows the efficient computation of marginal probability distributions on graphical networks without loops (39). In sum-product message passing, all nodes compute sums and integrals locally, before passing the results on towards the set of their neighbours in the shape of messages (72). In the present use-case, all variables are binary. In sum-product message-passing, the marginal probabilities for each state of a variable node V are given by the product of all incoming messages from factor nodes F :

$$p(V) = \prod_{l \in ne(V)} \mu_{F_l \rightarrow V}(V) \tag{A3}$$

where $ne(V)$ are the are the factor nodes' neighbouring the variable node V , and conversely for factor nodes. $\mu_{F_l \rightarrow V}$ is the message received from factor node F_l . Variable nodes can only have factor nodes as neighbours. Messages from a factor node F to a variable node V are given by

$$\mu_{F \rightarrow V}(V) = \sum_{V_1} \cdots \sum_{V_M} F(V, V_1, \dots, V_M) \prod_{m \in ne(F) \setminus V} \mu_{V_m \rightarrow F}(V_m) \tag{A4}$$

Here, the product of the factor and all messages the factor node received from its neighbouring variables nodes except V , $ne(F) \setminus V$ (a set of variable nodes without the recipient node), $\mu_{V_m \rightarrow F}(V_m)$ is computed and summed over all variables except V . When $ne(F) = \text{None}$, that is, F 's only neighbour is V , the factor node sends its factor information:

$$\mu_{F \rightarrow V}(V) = F(V). \tag{A5}$$

Messages from variable node V to a neighbouring factor node are given by

$$\mu_{V \rightarrow F}(V) = \prod_{l \in ne(V) \setminus F} \mu_{F_l \rightarrow V} \tag{A6}$$

that is, V computes the product of all received messages from neighbouring factor nodes $\mu_{F_l \rightarrow V}$ except F . $ne(V) \setminus F$ is empty when F is the only neighbour of V , in which case the message reduces to

$$\mu_{V \rightarrow F}(V) = 1 \tag{A7}$$

where the value 1 is sent from unobserved nodes. Otherwise, the message value corresponds to the observation made at V . Note that variable nodes can also remain unobserved.

We here adapt this scheme and transform messages into log-space to avoid over- or underflow of the usual floating point computer arithmetic. With this, the marginal of a variable V is now computed via

$$p(V) = \exp \left(\sum_{l \in ne(V)} \log(\mu_{F_l \rightarrow V}(V)) \right), \tag{A8}$$

i.e., the marginal of a variable node is given by the sum of all incoming, log-transformed messages $\log(\mu_{F_l \rightarrow V}(V))$ from factor nodes F_l .

Exact inference can become computationally inefficient or intractable in case of continuous variables. This is the case in our model due to the natural parameters of the probability distributions. To avoid this, we use variational inference and the lower bound approximation as implemented in the free-energy framework (78, 79, 36). Instead of estimating the (potentially intractable) posterior $p(X | D)$, with X representing model parameters and data D , an approximate distribution $q(X | D)$ is estimated. This further sidesteps the increasingly costly computation of sums over all variables and their values by introducing more simple update rules in order to compute the posterior distribution. The optimal $q(X | D)$ is estimated by maximizing the lower bound on the marginal log-likelihood of the data D , $\log(p(D))$, also referred to as the evidence lower bound (ELBO, (80)). In the present case, the marginal log-likelihood of D is given by

$$\begin{aligned} \log(p(D)) &= \log \left(\sum_x p(D, X) \right) \\ &= \log \left(\sum_x q(X) p(D | X) \frac{p(X)}{q(X)} \right) \end{aligned} \tag{A9}$$

Meaning that the marginal log-likelihood of the data is given by the sum over all joint probability distributions $p(D, X)$, or the sum of an approximate distribution of X , $q(X)$, and conditional probabilities of the data given X , $p(D | X)$ and a normalization factor $p(X)/q(X)$. The last term of Equation A9 represents a factorization of $\log(p(D))$.

When trying to approximate $\log(p(D))$, there are two essential ingredients: Jensen’s inequality and the Kullback-Leibler divergence (81, 82, 36). Applying Jensen’s inequality for convex functions, it follows that

$$\begin{aligned} \log\left(\sum_x q(X)P(D | X)\frac{p(X)}{q(X)}\right) \\ \geq \sum_x q(X)\log\left(p(D | X)\frac{p(X)}{q(X)}\right). \end{aligned} \tag{A10}$$

The log sum of the factorized marginal probability distribution of $p(D)$ is larger or equal to the sum of all approximate distributions $q(X)$ times the log-normalized conditional probability distribution $p(D | X)$.

In a second step, we decompose the right-hand side of Equation A10, using $\log(x * y) = \log(x) + \log(y)$ into the expected likelihood and the Kullback-Leibler (KL) divergence between the approximate posterior and prior. The KL divergence can be interpreted as the dissimilarity of two functions. In this context, we are interested in the dissimilarity between our prior and the approximated marginal probability $p(D)$.

$$\begin{aligned} \log\left(\sum_x q(X)p(D | X)\frac{p(X)}{q(X)}\right) \\ \geq \langle \log(p(D | X)) \rangle_{q(X)} - KL(q(X)||p(X)) \end{aligned} \tag{A11}$$

where $KL(q(X)||p(X))$ is the non-negative KL divergence (81, 36) between prior $p(X)$ and approximate distribution $q(X)$. It follows that a lower-bound $L(q(X), D)$ on the marginal log-likelihood $\log(P(D))$ is given by

$$\begin{aligned} L(q(X), D) = \langle \log(p(D | X)) \rangle_{q(X)} - KL(q(X)||p(X)) \\ \leq \log(p(D)). \end{aligned} \tag{A12}$$

Maximising the lower bound L w.r.t. $q(X)$ is equivalent to optimising the approximation between $p(X)$ and $q(X)$. When these conditions are met, updated model parameters will improve the approximation of $p(X)$ over time.

We now turn to describing the parameter updates in more detail. Variational inference schemes with identities vastly facilitate parameter updates by re-parameterizing the messages that are sent within the graphical model (36, 33). We specifically base parameter updates on exponential family distributions. The main reason for this choice is the mathematically convenient property of conjugate priors. A conjugate prior ensures that the prior probability distribution and the posterior probability distribution come from the same

distribution family (36), which vastly facilitates parameter updates.

In the present model, we use the multinomial distribution, which in its standard form is given by (see (33))

$$P(x | q) = \prod_{k=1}^{K-1} q_k^{x_k} \left(1 - \sum_{k=1}^{K-1} q_k\right)^{x_K}. \tag{A13}$$

The multinomial distribution is a generalization of the Bernoulli distribution to K possible outcomes (36). In it, multinomial random variables x are represented by vectors with K components. Each component is either 0 or 1, $x_k \in \{0, 1\}$ and all components sum up to 1, $\sum_{k=1}^K x_k = 1$. $q = q_1, \dots, q_K$ are the respective probabilities that $x_k = 1$, i.e. we use a 1-out-of- K or one-hot representation.

The conjugate prior on the multinomial distribution is given by the Dirichlet distribution. The density of the Dirichlet distribution is given by Endres (33):

$$p(q | \alpha) = M \prod_{k=1}^{K-1} q_k^{\alpha_k - 1} \left(1 - \sum_{k=1}^{K-1} q_k\right)^{\alpha_K - 1}. \tag{A14}$$

where M is

$$M = \frac{\Gamma(\sum_{k=1}^K \alpha_k)}{\prod_{k=1}^K \Gamma(\alpha_k)}, \tag{A15}$$

and $\Gamma(x)$ is the gamma function.

Variational inference allows for learning over time and avoids computationally costly new instantiations. For this, we need to introduce free energy nodes with natural parameters η (given by sufficient statistics λ and pseudocounts ν) into the graphical model. Free energy factors (in the following: factors/ factor nodes) are described by conjugate prior and exponential family distributions introduced above. Parameter updates in this context lead to improved model predictions so that here, we can interpret the updating process as learning over time.

Further efficiency is introduced into the model by incorporating hypernodes, e.g. $p(H_t | H_{t-1})$ and $p(S_t | H_t)$ in **Figure A1**. Hypernodes ensure equal prior and posterior distributions of $P(H_t | H_{t-1})$ and $p(S_t | H_t)$ at each time step t . They are connected to the more local, time point specific factor nodes implicitly via the lower bound, so that no loops occur and the application of belief-propagation algorithms remains feasible. In this design, each factor node’s parameters ν and λ are set to the specific parameter values of their associated hypernode. Hypernodes are not involved in sum-product message passing. Also, local factor parameters are not updated individually, but rather, the hypernodes’ parameters

TABLE A1. Detailed overview of exponential family parameter settings for the simulation of healthy- and chronic pain inference.

Parameter	Healthy inference	Pathological inference
Prior on 1st node	$\lambda = p(\text{pain}) = 0.222,$ $\nu = 100$	$\lambda = p(\text{pain}) = 0.777,$ $\nu = 100$
Top-down prior	$\lambda = p(\text{noxious}_t \text{pain}_t) = 0.8,$ $\lambda = p(\text{harmless}_t \overline{\text{pain}}_t) = 0.9,$ $\nu = 100$	$\lambda = p(\text{noxious}_t \overline{\text{pain}}_t) = 0.6,$ $\lambda = p(\text{harmless}_t \text{pain}_t) = 0.4,$ $\nu = 20$
Transition prior	$\lambda = p(H_t = \text{pain} H_{t-1} = \overline{\text{pain}}) = 0.2,$ $\lambda = p(H_t = \overline{\text{pain}} H_{t-1} = \text{pain}) = 0.7,$ $\nu = 100$	$\lambda = p(H_t = \text{pain} H_{t-1} = \overline{\text{pain}}) = 0.7,$ $\lambda = p(H_t = \overline{\text{pain}} H_{t-1} = \text{pain}) = 0.2,$ $\nu = 100$
Observations	$S_1 - S_5 = \text{noxious input}, S_{12} - S_{18} = \text{harmless input for both cases}$	

Further, the first observation scheme to produce Figure 2 is shown below. Probabilities and their respective counter-probabilities sum up to 1 (83).

are updated after message-passing and then passed on to the local factor nodes. It can be shown that updates of hypernode parameters are done via:

$$\tilde{\nu} := \nu + r \tag{A16}$$

$$\tilde{\lambda} := \frac{\nu\lambda + r}{\tilde{\nu}} \tag{A17}$$

where $\tilde{\nu}$ are the updated ν , or pseudocounts and $\tilde{\lambda}$ denotes the updated λ , or natural parameters. β is an inverse temperature parameter. In case of Bayes-optimal updates, $\beta = 1$. ν is referred to as the *pseudocount* as it keeps count of the number of observed data points. $\tilde{\lambda}$ can be interpreted as a weighted mean of the prior λ and the observed data. The responsibilities r , i.e. the accumulated posterior probabilities of observations, can be computed from all messages a factor node F received:

$$r = \frac{\exp(\log(F(X)) + \sum_{m \in ne(F)} \log(\mu_{v_m \rightarrow F}(V_m)))}{\sum_{m \in ne(F)} \exp(\log(F(X)) + \sum_{m \in ne(F)} \log(\mu_{V_m \rightarrow F}(V_m)))} \tag{A18}$$

Hypernodes collect responsibilities from all neighbouring factor nodes:

$$r_H := \sum_{n \in ne(H)}^N r_n \tag{A19}$$

This way, all observations made at the observable variable nodes in each time step are taken into account. Before message-passing can be instantiated anew, the posterior parameters of all local factor nodes are set to those of their neighbouring hypernode. The exponential family parameterization used in our simulations of healthy vs. chronic pain inference are summarized in Table A1.

Appendix B. Derivation of null space and stable pain states

Chronic pain is marked by very tenacious pain perception in the absence of acute tissue damage. This translates to stable pain states in our model. To achieve this stability of hidden states, we need to derive transition probabilities $P(\text{pain}_{t+1} | \text{pain}_t)$ and $P(\overline{\text{pain}}_{t+1} | \overline{\text{pain}}_t)$ such that the inferred probability of pain in some next time step $t + 1$ is equal to the probability of pain in the current time step t , which in turn is equal to some stable marginal probability of pain, $P(\text{pain}_\infty)$:

$$P(\text{pain}_{t+1}) = P(\text{pain}_t) = P(\text{pain}_\infty) \tag{B1}$$

The probability of pain in a future time step $t + 1$ is given by the sum over the possible previous states $\in \{\text{pain}, \overline{\text{pain}}\}$ and the transition probabilities:

$$P(\text{pain}_{t+1}) = P(\text{pain}_{t+1} | \text{pain}_t) \cdot P(\text{pain}_t) + (1 - P(\overline{\text{pain}}_{t+1} | \overline{\text{pain}}_t)) \cdot (1 - P(\text{pain}_t)) \tag{B2}$$

Thus, for a stable pain percept after a long time, we require:

$$P(\text{pain}_\infty) = P(\text{pain}_{t+1} | \text{pain}_t) \cdot P(\text{pain}_\infty) + (1 - P(\overline{\text{pain}}_{t+1} | \overline{\text{pain}}_t)) \cdot (1 - P(\text{pain}_\infty)) \tag{B3}$$

Conversely, the condition for a stable no-pain percept is:

$$(1 - P(\text{pain}_\infty)) = (1 - P(\text{pain}_{t+1} | \text{pain}_t)) \cdot P(\text{pain}_\infty) + P(\overline{\text{pain}}_{t+1} | \overline{\text{pain}}_t) \cdot (1 - P(\text{pain}_\infty)) \tag{B4}$$

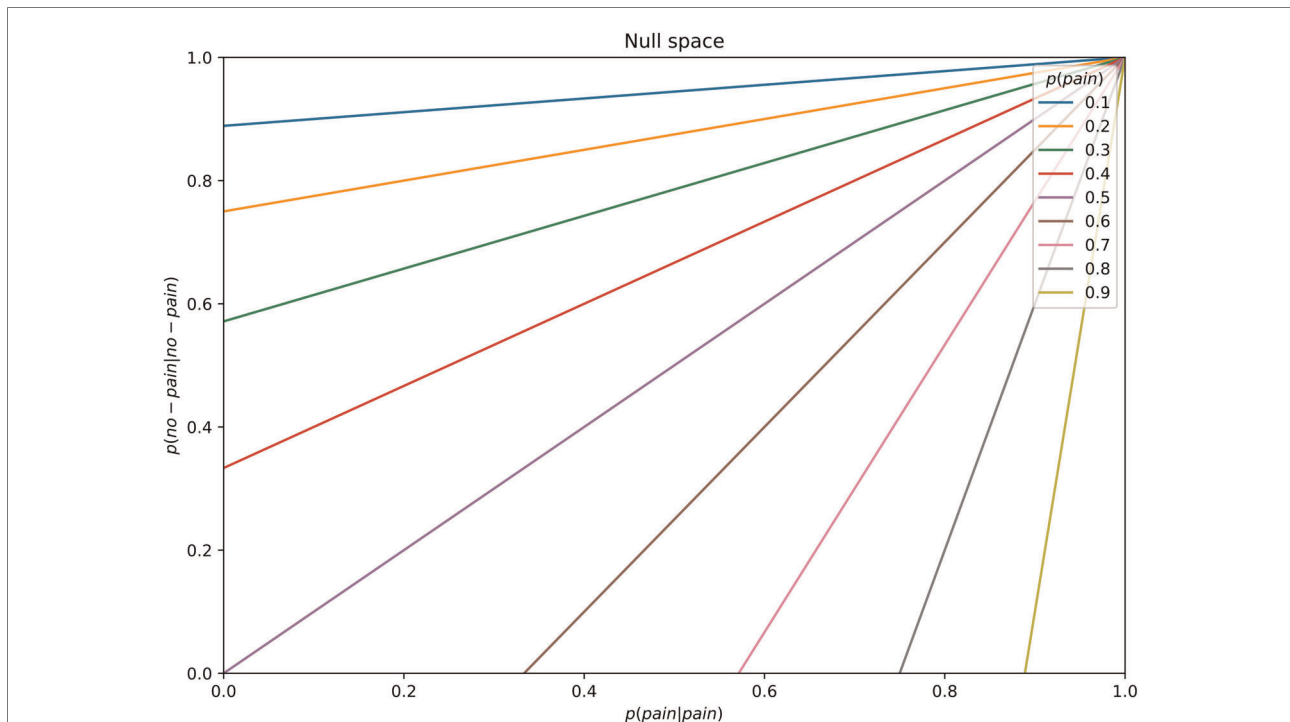


FIGURE B1
 Illustration of null space of transition probability deviations for different marginal $p(\text{pain}_\infty)$. We show that when the transition probabilities $p(\text{pain}_{t+1} | \text{pain}_t)$ and $p(\overline{\text{pain}}_{t+1} | \overline{\text{pain}}_t)$ are sampled from the one-dimensional null space of a given marginal probability of pain $p(\text{pain}_\infty)$, the marginal probability $p(\text{pain}_t)$ will approach and stabilize within the range of $p(\text{pain}_\infty)$. Here, we illustrate the transition probabilities for $p(\text{pain}_\infty) \in [0.1, 0.9]$ (legend).

Solving equation B3 for the transition probabilities, we find

$$\begin{aligned}
 &P(\text{pain}_{t+1} | \text{pain}_t) \cdot P(\text{pain}_\infty) - P(\overline{\text{pain}}_{t+1} | \overline{\text{pain}}_t) \cdot (1 - P(\text{pain}_\infty)) \\
 &= P(\text{pain}_\infty) + P(\text{pain}_\infty) - 1 = 2P(\text{pain}_\infty) - 1
 \end{aligned}
 \tag{B5}$$

Similarly, for $(1 - P(\text{pain}_\infty))$ (Equation B4):

$$\begin{aligned}
 &-P(\text{pain}_{t+1} | \text{pain}_t) \cdot P(\text{pain}_\infty) + P(\overline{\text{pain}}_{t+1} | \overline{\text{pain}}_t) \cdot (1 - P(\text{pain}_\infty)) \\
 &= 1 - P(\text{pain}_\infty) - P(\text{pain}_\infty) = 1 - 2P(\text{pain}_\infty)
 \end{aligned}
 \tag{B6}$$

In matrix-vector form, the last two equations can be written as

$$\begin{pmatrix} P(\text{pain}_\infty) & -(1 - P(\text{pain}_\infty)) \\ -P(\text{pain}_\infty) & (1 - P(\text{pain}_\infty)) \end{pmatrix} \begin{pmatrix} P(\text{pain}_{t+1} | \text{pain}_t) \\ P(\overline{\text{pain}}_{t+1} | \overline{\text{pain}}_t) \end{pmatrix} = \begin{pmatrix} 2P(\text{pain}_\infty) - 1 \\ 1 - 2P(\text{pain}_\infty) \end{pmatrix}
 \tag{B7}$$

To solve for the transition probabilities, we decompose the

transition probability vector into the sum of q and r :

$$(q + r) = \begin{pmatrix} P(\text{pain}_{t+1} | \text{pain}_t) \\ P(\overline{\text{pain}}_{t+1} | \overline{\text{pain}}_t) \end{pmatrix}
 \tag{B8}$$

where q and r are constrained such that

$$\begin{pmatrix} P(\text{pain}_\infty) & -(1 - P(\text{pain}_\infty)) \\ -P(\text{pain}_\infty) & (1 - P(\text{pain}_\infty)) \end{pmatrix} \cdot q = \begin{pmatrix} 2P(\text{pain}_\infty) - 1 \\ 1 - 2P(\text{pain}_\infty) \end{pmatrix}
 \tag{B9}$$

in other words, q is that summand of the transition probability decomposition which produces the desired stable marginal probability of pain.

In contrast, for r we require that

$$\begin{pmatrix} P(\text{pain}_\infty) & -(1 - P(\text{pain}_\infty)) \\ -P(\text{pain}_\infty) & (1 - P(\text{pain}_\infty)) \end{pmatrix} \cdot r = 0
 \tag{B10}$$

i.e. r is a vector in the null space of the matrix in Equation B7.

Solving for the components of $q = (q_1, q_2)^T$ and $r = (r_1, r_2)^T$ yields

$$\begin{pmatrix} P(\text{pain}_\infty) & -(1 - P(\text{pain}_\infty)) \\ -P(\text{pain}_\infty) & (1 - P(\text{pain}_\infty)) \end{pmatrix} \begin{pmatrix} q_1 \\ q_2 \end{pmatrix} = \begin{pmatrix} 2P(\text{pain}_\infty) - 1 \\ 1 - 2P(\text{pain}_\infty) \end{pmatrix} \tag{B11}$$

i.e.

$$q_1 = 2P(\text{pain}_\infty) - 1 = -q_2 \tag{B12}$$

Similarly, for r ,

$$\begin{pmatrix} P(\text{pain}_\infty) & -(1 - P(\text{pain}_\infty)) \\ -P(\text{pain}_\infty) & (1 - P(\text{pain}_\infty)) \end{pmatrix} \begin{pmatrix} r_1 \\ r_2 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \tag{B13}$$

we find

$$\frac{P(\text{pain}_\infty)}{1 - P(\text{pain}_\infty)} r_1 = r_2 \tag{B14}$$

Thus, the transition probabilities for a given $P(\text{pain}_\infty)$ can be written as

$$\begin{pmatrix} P(\overline{\text{pain}}_{t+1} | \overline{\text{pain}}_t) \\ P(\text{pain}_{t+1} | \text{pain}_t) \end{pmatrix} = \begin{pmatrix} 2P(\text{pain}_\infty) - 1 \\ 1 - 2P(\text{pain}_\infty) \end{pmatrix} + \beta \begin{pmatrix} 1 \\ \frac{P(\text{pain}_\infty)}{1 - P(\text{pain}_\infty)} \end{pmatrix} \tag{B15}$$

for any $\beta \in \mathbb{R}$ such that $P(\text{pain}_{t+1} | \text{pain}_t), P(\overline{\text{pain}}_{t+1} | \overline{\text{pain}}_t) \in [0, 1]$. This implies a one-dimensional null space, where pain percepts remain stable even if the transition probabilities change within this null space (further illustrated in **Figure B1**). Alternatively, we can eliminate β by solving the first component equation for β and substituting the result into the second component equation:

$$\begin{aligned} P(\overline{\text{pain}}_{t+1} | \overline{\text{pain}}_t) &= \frac{1 - 2P(\text{pain}_\infty)}{1 - P(\text{pain}_\infty)} \\ &+ \frac{P(\text{pain}_\infty)}{1 - P(\text{pain}_\infty)} \cdot P(\text{pain}_{t+1} | \text{pain}_t) \end{aligned} \tag{B16}$$