

Modeling the Action-Perception Loop and its role in Phantom Limb Pain using Active Inference

Downloaded from: https://research.chalmers.se, 2026-11-28 23:19 UTC

Citation for the original published paper (version of record):

Ramne, M., Lundh, T., Sensinger, J. (2027). Modeling the Action-Perception Loop and its role in Phantom Limb Pain using Active Inference. Journal of Theoretical Biology, 638. http://dx.doi.org/10.1016/j.jtbi.2026.112634

N.B. When citing this work, cite the original published paper.

research.chalmers.se offers the possibility of retrieving research publications produced at Chalmers University of Technology. It covers all kind of research output: articles, dissertations, conference papers, reports etc. since 2004. research.chalmers.se is administrated and maintained by Chalmers Library

(article starts on next page)

Journal of Theoretical Biology 638 (2027) 112634

Contents lists available at ScienceDirect

Journal of Theoretical Biology

journal homepage: www.elsevier.com/locate/yjtbi

Modeling the Action-Perception loop and its role in phantom limb pain using active inference

Malin Ramne [a,*], Torbjörn Lundh [b], Jon Sensinger [c]

[a] Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden
[b] Department of Mathematical Sciences, Chalmers University of Technology and University of Gothenburg, Sweden
[c] Institute of Biomedical Engineering, University of New Brunswick, Fredericton, Canada

ARTICLE INFO

Keywords:
Active inference
Bayesian inference
Phantom limb pain
Pain
Computational neuroscience

ABSTRACT

Mathematical models grounded in the Bayesian inference framework have been used to describe various aspects of pain perception. However, pain is not only passively inferred but actively shaped through interactions with the environment – a dimension that classical Bayesian approaches typically do not capture. Here, we investigate the role of the action-perception loop in phantom limb pain using a model within the active inference framework, which extends Bayesian inference to include action selection. Despite being one of the most prevalent and distressing consequences of limb amputation, theories regarding the underlying mechanisms of phantom limb pain remain disputed. Because amputation disrupts both sensory input related to the limb and the ability to perform actions, a model incorporating both sensory and active components of pain may provide new mechanistic insight into this condition. The presented model provides a conceptual account of how loss of limb control, ambiguity in sensory input pertaining to limb position, residual noxious input, and pre-amputation pain may contribute to the emergence and persistence of phantom limb pain. Furthermore, the model offers insight into the possible mechanisms underlying common interventions and may help account for their variable efficacy across individuals.

1. Introduction

Pain is not merely a sensory experience, but a complex, inferential process shaped by cognitive factors, such as expectations informed by previous experiences and contextual cues (Villemure and Bushnell, 2002; Wiech and Tracey, 2013). Bayesian inference has successfully been applied to describe a range of pain phenomena such as placebo hypoalgesia and nocebo hyperalgesia, statistical pain learning under experimental pain paradigms, and how overly precise priors and ambiguous observation likelihoods may contribute to chronic and neuropathic pain conditions (Tabor and Burr, 2019; Tabor et al., 2017; Grahl et al., 2018; Jepma et al., 2018; Hoskin et al., 2019; Drusko et al., 2023; Mancini et al., 2022; Eckert et al., 2022; Ramne and Sensinger, 2024). While considerable work has investigated how expectations influence pain perception (see e.g., Atlas and Wager (2012); Fields (2018) for reviews), another dimension of the expectation-pain interaction remains less well understood: actions. Our expectations of pain influence what actions we take, and the actions we take influence the pain that we perceive (Wiech and Tracey, 2013; Seymour, 2019; Seymour and Mancini, 2020; Büchel, 2023; Habermann et al., 2025). Thus, pain is not only passively inferred but actively shaped through interactions with the environment. Following amputation, the ability to perform actions and the sensory input related to the state of the limb are severely affected. A model that can account for both the sensory and active component of pain may provide insight to the possible mechanisms contributing to pathological pain conditions such as phantom limb pain.

Several hypotheses have been proposed for mechanisms associated with phantom limb pain, including accounts based on cortical reorganisation (Flor et al., 1995; Flor, 2002), preservation of cortical representations (Makin et al., 2013; Makin et al., 2015), and sensorimotor incongruence (Harris, 1999). While these frameworks have provided important insights, they often lack mechanistic specificity or struggle to account for the full range of observed phenomena (see e.g., Schone et al. (2022) for a broader review, and also this thesis (Ramne, 2026). One way to increase mechanistic specificity is through computational modelling. A few mathematical models of phantom limb pain have been proposed (Ramne and Sensinger, 2024; Ramne, 2025; Boström et al., 2014; Ropero Peláez and Taniguchi, 2016). However, these models

* Corresponding author.
E-mail address: malin@ramne.com (M. Ramne).

https://doi.org/10.1016/j.jtbi.2026.112634
Received 23 December 2025; Received in revised form 27 April 2026; Accepted 29 October 2026
Available online 1 November 2026
0022-5193/© 2026 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

typically focus on perceptual aspects and do not explicitly incorporate the action–perception loop. Notably, Weiss et al. proposed a predictive coding account that extends sensorimotor incongruence and incorporates elements of action (Weiss et al., 2023). However, this model is conceptual in the sense that is does not perform any numerical calculations and it relies on assumptions that are not yet well supported by empirical evidence (e.g., that mismatches between motor intention and sensory input can directly elicit pain). Taken together, existing models do not explicitly capture the action–perception loop, highlighting the need for computational approaches that integrate both action and perception in phantom limb pain.

One of the key challenges in modelling the action-perception loop is simultaneously optimizing exploration, exploitation, perception and learning. One framework that has garnered attention in recent years, that proposes a simple yet elegant solution to this challenge, is active inference. The solution proposed by active inference is that all living organisms follow a single objective: *minimizing the surprise of their sensory observations* (Parr et al., 2022). Importantly, in the active inference framework *surprise* has a technical meaning – it measures how much an agent's current sensory observations differ from its preferred sensory observations. With this governing philosophy, actions are performed to achieve the following two objectives:

- to obtain *sensory observations that correspond to* desired outcomes or goals (pragmatic value), or
- to *obtain sensory observations that* reduce uncertainty about the world (epistemic value).

This formulation allows outcomes of actions to be quantified in the same units as perception, enabling the optimization of exploration, exploitation, perception and learning to all be cast as the minimization of surprise (formalized as *free energy minimization* in the active inference framework, see Section 4).

Here, we present an active inference model of pain perception and action, with a particular aim of exploring how phantom limb pain may arise following limb loss. Our results indicate the loss of ability to control the limb, ambiguity in noxious and proprioceptive input, and pre-amputation pain are factors that may contribute to phantom limb pain. Our results also provide insights to the possible working mechanism of some interventions that are commonly used to relieve phantom limb pain, and why the interventions may have different efficacy in different patients.

2. Results

2.1. Model description and verification

Before diving into the model predictions on pain after limb loss, we first give a brief description of the model, and then verify that the model shows expected behavior on a more well-known scenario: learning and context switching. A more thorough description of the model derivation and implementation is provided in Section 4.

To model pain we must first have a clear definition of what pain is. The International Association for the Study of Pain (IASP) provides the following definition (Raja et al., 2020) that we will adopt in this paper:

> *"An unpleasant sensory and emotional experience associated with, or resembling that associated with, actual or potential tissue damage."*

In any active inference model, there are two main interacting components: the environment, and the agent. Based on the definition by IASP, and since we are interested in studying the specific scenario of phantom limb pain, we let the environment be the tissue of the limb the agent is making inferences about, and we let the agent be the central nervous system (CNS) of the organism at hand. The interface between the environment and the agent consists of observations in the form of peripheral afferent sensory input to arriving to CNS, and actions in the form of motor control.

The *generative process* describes how observations are generated by the environment and how the environment changes in response to actions from the agent. The *generative model*[1] describes how the agent makes inferences about the environment based on observations and expectations, and how the agent chooses actions to obtain desirable outcomes. The agent's state inference is informed by three different *observation modalities*: position, noxious and reward observations. The agent has a prior preference for reward observations and against noxious observations. Position observations can be thought of as a combination of proprioceptive and visual input about the limb's position and is not associated with any preference. The agent's actions correspond to moving the limb to different positions. Fig. 1 outlines how the generative process and generative model interface with each other, and how information flows between key processes in this active inference model. Additional details of the model implementation are provided in Section 4.

To verify the plausibility of the model, we demonstrate that an agent with an initially flat expectation of the environment can 1) learn to navigate the environment to optimally fulfill its preferences, and 2) re-learn the environment when the locations of high- and low probability of noxious observations switch. Here, the environment consists of four different limb positions, each with either high or low probability of noxious and reward observations respectively. As demonstrated by the example in Fig. 2, an agent with no prior knowledge of the environment learns through exploration that Position 2 is the most favorable (high probability of reward observations, low probability of noxious observations), and then re-learns and adapts its behavior when the environment changes such that Position 4 instead is the best.

2.2. Model predictions

Since we have defined the environment to be the tissue of the limb the agent is making inferences about, limb loss causes dramatic changes to the environment and its generative process. Peripheral afferent nerves that used to innervate the limb become severed and are no longer able to reliably signal information about the state of the limb. In the generative process, we represent this change as 1) increased ambiguity in the noxious observations, and 2) equal likelihood for noxious observations at all limb positions. To reflect differing levels of spontaneous activity in peripheral afferents we will study different levels of ambiguity in the noxious observations. Similarly, the level of ambiguity of the position observations (combined visual and proprioceptive input) and the ability to control the limbs position may vary depending on the level of amputation and reinnervation of the residual limb. Finally, we assume that agent's ability obtain reward observations with the affected limb is impaired. To reflect this change, we set a low probability for reward observations at all positions.

Taken together, we have identified three parameters that may be affected to a varying degree following limb loss: the likelihood of noxious observations, the ambiguity of position observation likelihoods, and the controllability of the limb. Here, we investigate how these parameters might contribute to phantom limb pain following limb loss. In our simulations we have included an additional parameter: duration of pre amputation pain. Persistent pre amputation pain is a well-established risk factor for phantom limb pain (Limakatso et al., 2020). The facet plot in Fig. 3 shows the prevalence of phantom limb pain for different combinations of values of these four factors. In the simulations, phantom limb pain is considered present at a given time point when the

[1] Here, *generative model* follows its usage in computational and cognitive neuroscience, denoting an internal model that generates predictions of sensory input from latent causes in the environment. This should not be confused with generative AI models that produce novel data samples (e.g., images or text) from learned distributions.

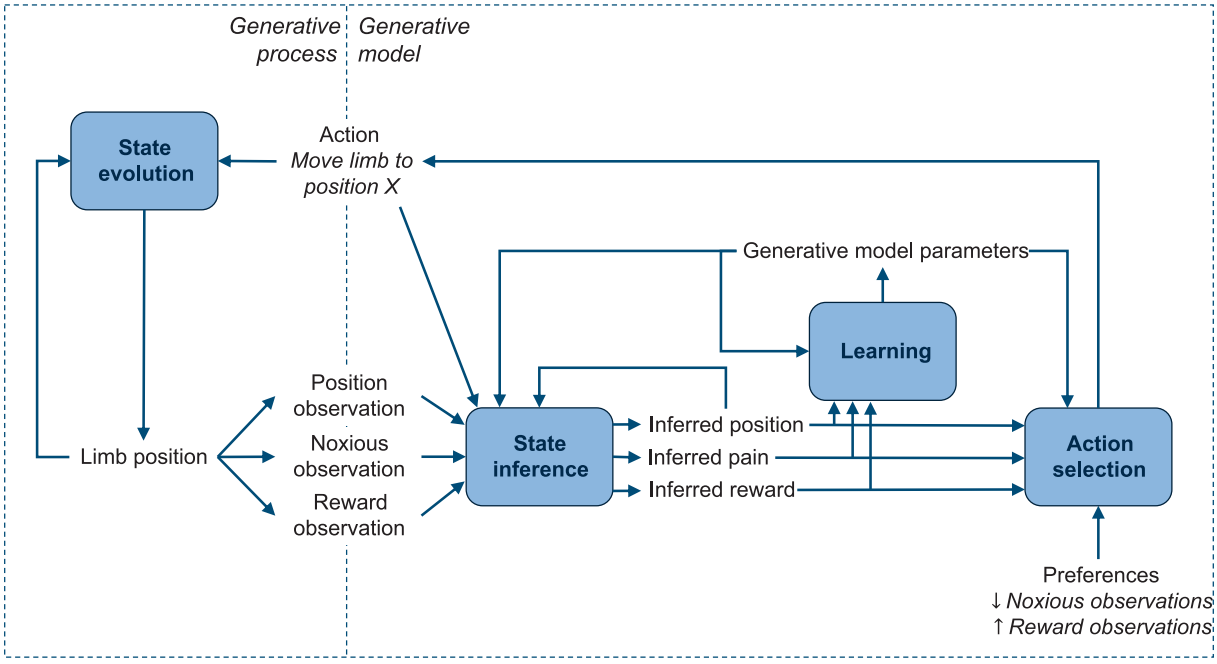

Fig. 1. Model outline. The generative process describes how observations are generated by the environment and how the environment changes in response to actions from the agent, and the generative model describes how the agent makes inferences about the environment based on observations and chooses actions to obtain desirable outcomes. State inference, action selection, learning and state evolution are key processes in the active inference framework and are described in more detail in Section 4 and Appendix A.

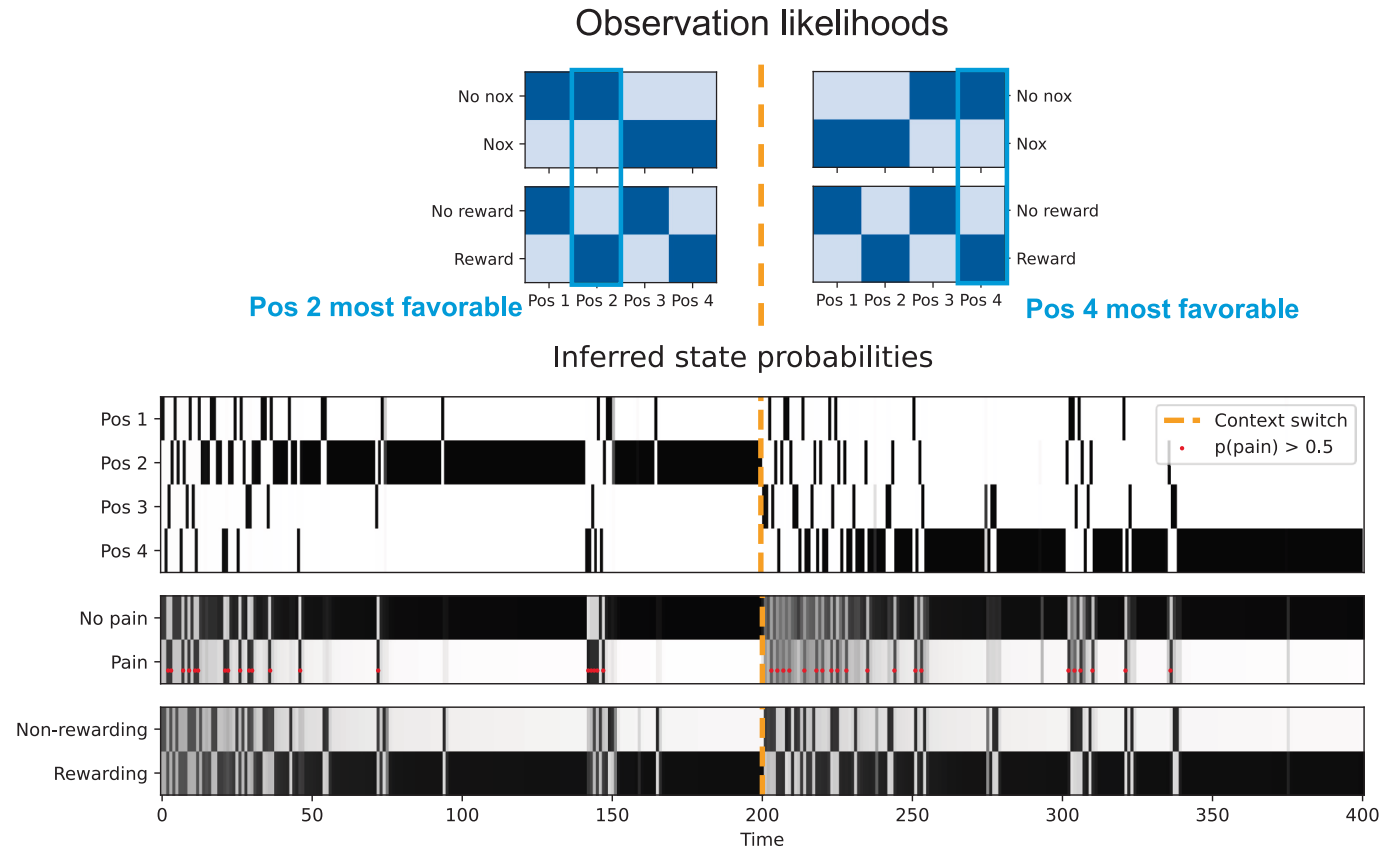

Fig. 2. Simulation of learning and context switching. In the initial environment Position 2 is the most favorable (low probability of noxious observations, high probability of reward observations). At onset, the agent has no prior knowledge of how locations relate to pain and reward-state. After some exploration, the agent becomes confident enough that Position 2 is the most favorable and switches to mostly exploitative behavior. When the environment switches at t = 200, Position 2 is no longer favorable as it then has a high probability of noxious observations. Instead, Position 4 is the most favorable, which the agent learns after some initial trial-and-error following the switch. A more detailed visualization of the simulation parameters and results can be found in Appendix B.

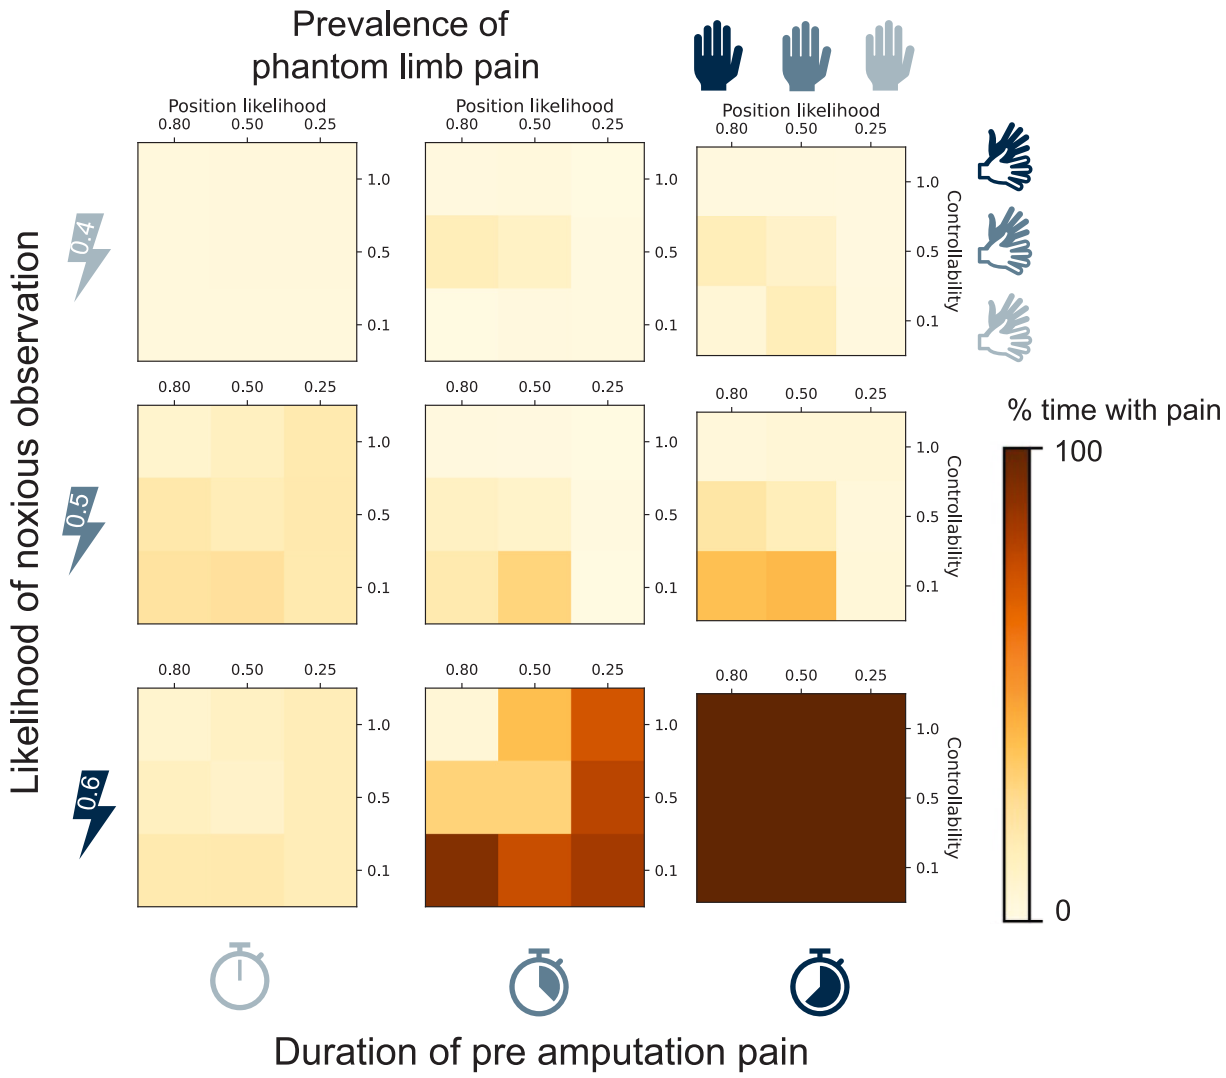

Fig. 3. Predicted phantom limb pain prevalence. Prevalence of phantom limb pain for different combinations of values for position likelihood, controllability, likelihood of noxious observations, and duration of pre amputation pain. These four factors are denoted by the blue icons with varying saturation along the axes of the facet plot. Higher likelihood of noxious observations, for example due to residual peripheral noxious input, and duration of pre amputation pain, are associated with higher prevalence of phantom limb pain. The influence of position likelihood and controllability is more complex and dependent on the other parameters.

inferred probability of pain exceeds 0.5. Prevalence is then estimated as the proportion of time points across all simulations in which phantom limb pain is present, providing a proxy for population-level prevalence. For simplicity, we have only included inferred pain state in the following section, additional figures depicting the position and reward state can be found in Appendix B.

These results provide insight to the possible complex interaction of factors contributing to phantom limb pain. Such insights could help inform preventative measures. For example, our results corroborate that pre amputation pain is a key risk factor for phantom limb pain, and thus shall be avoided. But could these results also give us insight to how to relieve the pain?

If some of the effects are reversible, one could imagine that interventions that move an individual from a high-prevalence region of the facet plot to a low-prevalence region might lead to some reduction in pain prevalence. Inspection of the patterns in the facet plot reveal two types of inter-plot movement that potentially could be associated with pain relief:

1) vertical upward movement between heat maps
2) diagonal left/up movement within heatmaps.

The first of these interventions would correspond to addressing peripheral sources of residual noxious input. In practice, such interventions might correspond to applying a nerve block to spontaneously active neuromas or the dorsal root ganglion (Vaso et al., 2014; Grant et al., 2023; Borghi et al., 2010), or through surgical interventions promoting reinnervation of severed nerves into denervated tissue (Dumanian et al., 2019; Woo et al., 2016; Pettersen et al., 2023). Another possible approach could be to increase the ratio of innocuous to noxious afferent input through stimulation of peripheral nerves (Albright-Trainer et al., 2021; Steven Cohen et al., 2019; Pagan-Rosado et al., 2023). Fig. 4 shows an example of how this type of intervention might reduce the prevalence of phantom limb pain.

For the second intervention two things must happen simultaneously: i) the ability to control the limb must be restored (upward movement in the plot), and ii) the information about the position of the limb must become more precise (leftward movement in the plot). Apart from unamputating the missing limb, these changes might seem impossible. However, there are other ways of achieving similar effects, for example by means of an active prosthesis or an AR/VR limb that is controlled by electromyographic signals recorded on the residual limb (Ortiz-Catalan et al., 2016; Lendaro et al., 2025). When the agent attempts to initiate a movement of the missing limb, electromyographic activity is recorded

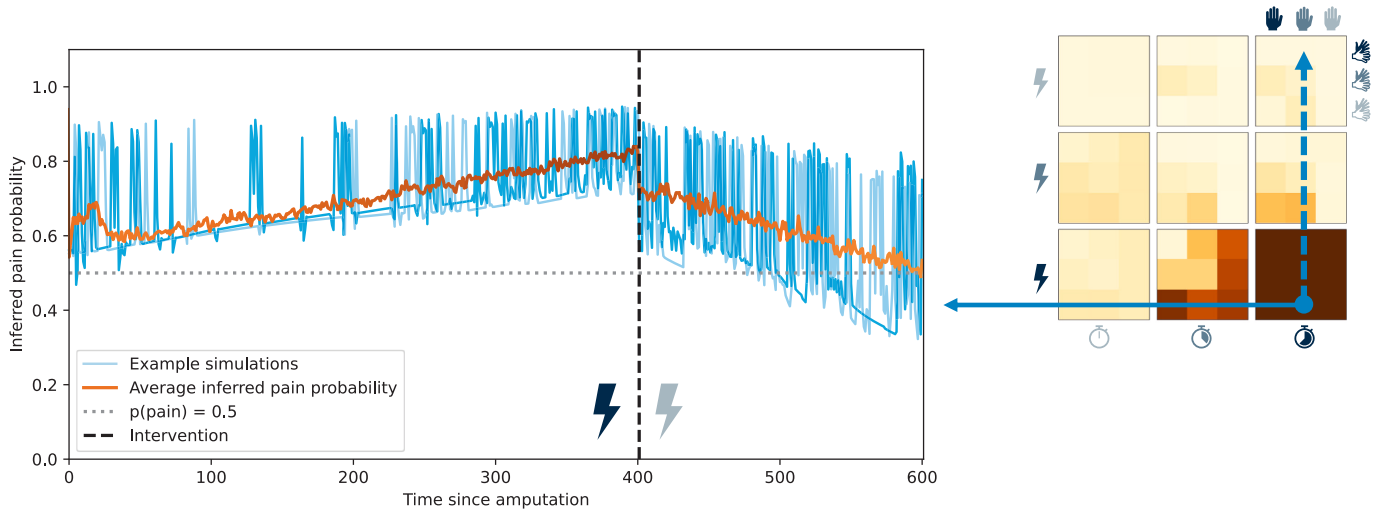

Fig. 4. Example intervention. Example of how reducing the likelihood of noxious observations can be used as a phantom limb pain relieving intervention. The orange line indicates the average inferred pain probability across 100 simulations and the blues lines represent three example simulations. In the facet plot, this intervention corresponds to vertical upward movement between heatmaps. In practice, such interventions might correspond to applying a nerve block to spontaneously active neuromas or the dorsal root ganglion, surgical interventions, or peripheral nerve stimulation.

from the muscles of the residual limb, allowing the motor intent to be decoded and actuated in movement of the prosthesis or virtual limb, resulting in sensory input matching the intended movement of the missing limb. Similarly, a mirror reflection of the contralateral, intact limb could be used (Barbin et al., 2016; Limakatso et al., 2023). If the agent matches the movements of the intact limb with intended movements of the missing limb, the visual input emanating from the mirror will match the intended movement of the missing limb. Surgical interventions that promote proprioceptive feedback from the severed nerves could also fall into this category of intervention (Srinivasan et al., 2021; Gardetto et al., 2021). Importantly, in these interventions, the agent is given a means of controlling some *source* other than the missing limb that generates sensory input that matches the intended movements.

The left panel in Fig. 5 exemplifies how such an intervention can result in reduction in prevalence of phantom limb pain. However, as the right panel reveals, such interventions are not guaranteed to succeed. For individuals who find themselves in the bottom right heatmap, interventions corresponding to movement within the heatmap will have no effect. Thus, in addition to predicting possible interventions for reducing phantom limb pain, our model results also give an indication of how to prioritize among these interventions based on patient characteristics. Note, however, that the colors of the heatmap do not perfectly predict intervention success. Experiencing phantom limb pain for an extended period of time may partially have the same effects as pre amputation pain, possibly reducing the effectiveness of interventions as indicated by these heatmaps. See Appendix B for one such example, along with more detailed visualizations of interventions.

Fig. 6 summarizes the four modelled factors that may contribute to phantom limb pain, and how the two proposed interventions addresses some of these factors.

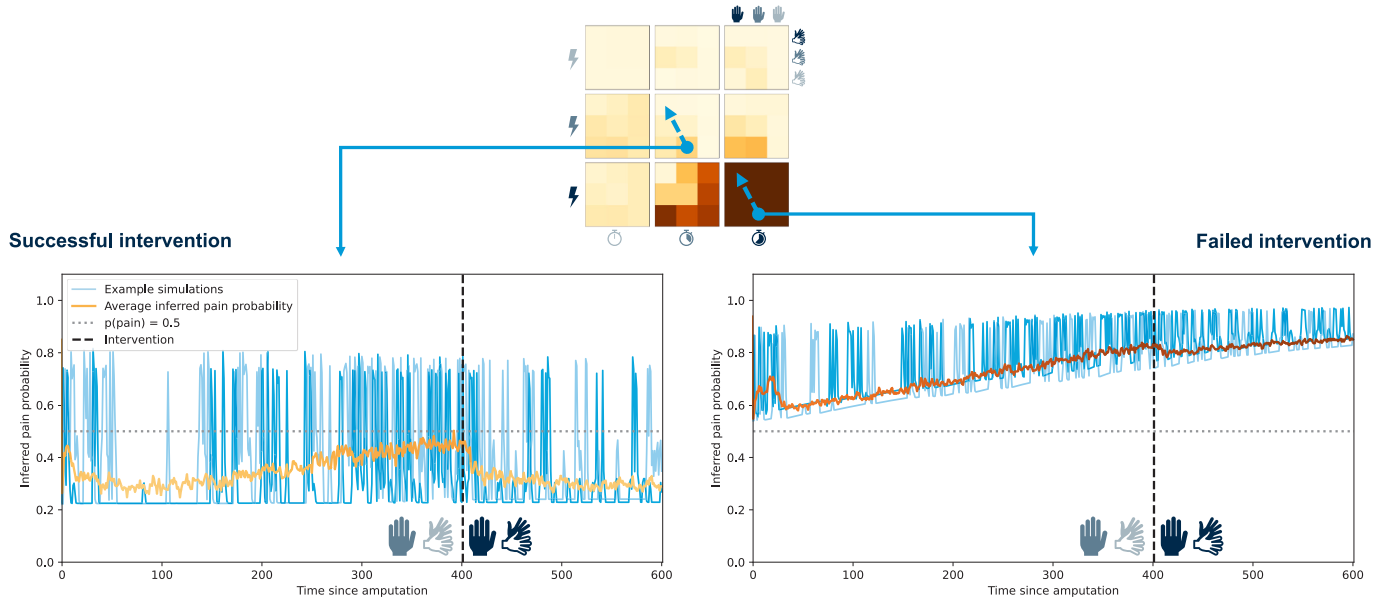

Fig. 5. Example intervention. Examples of varying success of the second intervention (corresponding to diagonal up/left movement within heatmaps) for different patient characteristics. For patients who are in the bottom right heatmap of the facet plot, interventions corresponding to movement within the heatmap will have no effect. These patients are likely better helped by the intervention exemplified in Fig. 4.

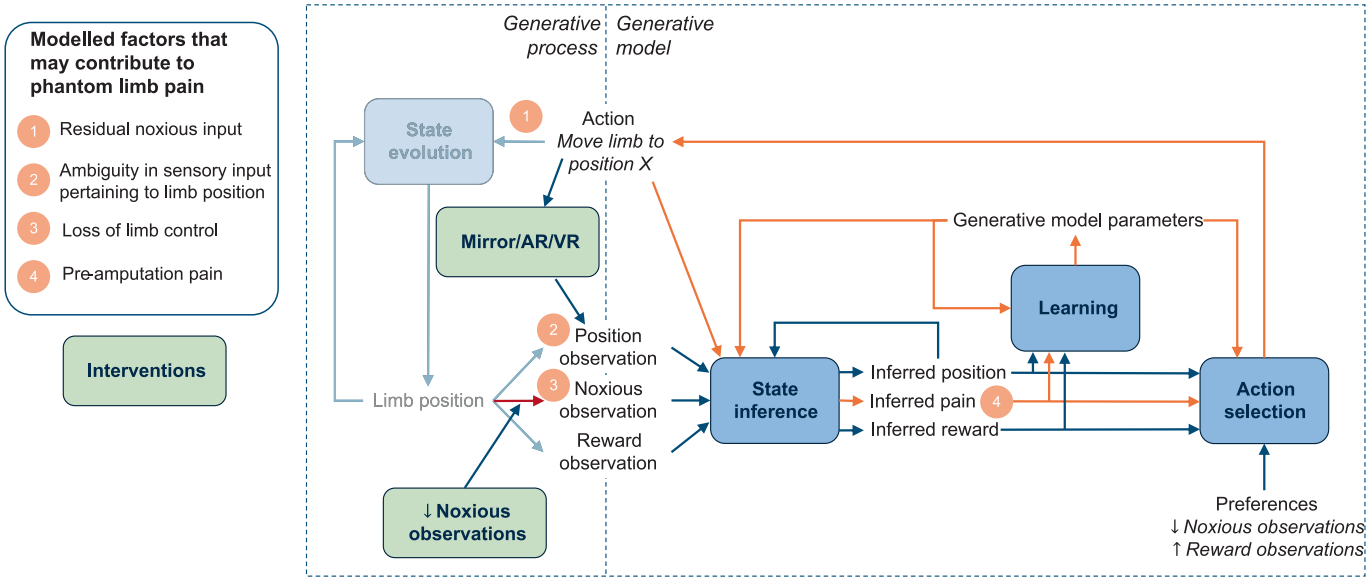

Fig. 6. Summary of model predictions. The proposed model provides a conceptual account of how four distinct factors may contribute to the emergence and persistence of phantom limb pain following amputation: 1) loss of control of the limb, 2) ambiguity in sensory input pertaining to limb position, 3) residual activity in afferent nociceptive neurons, and 4) pre-amputation pain (through feedback loops indicated in orange). The first two factors can be addressed with interventions using mirrors, AR or VR, to give the agent a means of controlling some source of sensory input that matches the intended movements. Residual noxious input may be addressed through surgical interventions, nerve block to spontaneously active peripheral nerves or through peripheral nerve stimulation.

3. Discussion

Phantom limb pain is one of the most prevalent and distressing consequences of limb amputation (Schone et al., 2022; Limakatso et al., 2020). Theories about the underlying mechanism remain diverse and poorly supported by empirical evidence, which contributes to challenges in effectively treating the pain (Schone et al., 2022; Limakatso and Parker, 2021). What is beyond dispute is that, following amputation, both sensory input related to the state of the limb and the ability to perform actions are profoundly altered. This led us to hypothesize that a model that can account for both the sensory and active component of pain may provide insight to the possible mechanisms contributing to phantom limb pain.

To this end, we developed a model within the active inference framework that provides a conceptual account of how four different factors interact to influence the prevalence of phantom limb pain. Specifically, the results suggest that loss of control over the limb, ambiguity in proprioceptive input, residual noxious input, and pre-amputation pain may all contribute to the emergence and persistence of phantom limb pain. Furthermore, the model offers insight into the possible mechanisms underlying common interventions and may help explain why their efficacy varies across individuals. Importantly, the results presented here should be interpreted as qualitative, model-based predictions rather than empirically validated findings. The model has not been fit to experimental data, and its parameters were selected to reproduce plausible behaviour. As such, the conclusions drawn from the simulations should be viewed as hypothesis-generating rather than confirmatory.

Based on the model predictions, we can coarsely identify two subtypes of phantom limb pain: primarily peripherally driven – maintained by residual peripheral noxious activity, and primarily centrally driven – maintained by the agent falsely inferring that the limb is in a position that previously has been associated with pain. The second of these is most likely to occur when the controllability and position likelihood precision both are low, making it difficult for the agent to move the phantom limb from the painful position. Furthermore, if there has been an extended period of pre amputation pain, multiple positions are more likely to be associated with pain, making it even more difficult to move the phantom limb into a "pain free" position. Phantom pain being associated with a specific limb position is reported by some patients (Sugawara et al., 2021; Collins et al., 2018); however, not all phantom pain has this quality. Other common descriptors of phantom pain are sensations like stabbing, burning or shocking, which are typical of neuropathic pain of peripheral origin (Sugawara et al., 2021; Hill, 1999; Attal et al., 2018). Thus, the quality and characteristics of phantom pain could be a possible clinical predictor of the underlying cause of the pain and which intervention is most likely to be successful. While this subtyping of phantom limb pain – particularly the proposed link between pain characteristics and underlying subtype – remains speculative, it suggests that future studies examining how pain characteristics relate to treatment responsiveness could provide valuable insights for improving treatment selection and outcomes.

These findings may also explain the varying efficacy of some interventions. As we have already seen in Fig. 5, the model predicts that interventions targeting central mechanisms won't work if the pain is driven by peripheral noxious input. Conversely, there is empirical evidence that interventions targeting peripheral mechanisms do not guarantee universal pain relief (Vaso et al., 2014; Grant et al., 2023; Borghi et al., 2010; Dumanian et al., 2019; Woo et al., 2016; Birbaumer et al., 1997), indicating involvement of central pain mechanisms. Such mismatches between pain subtype and intervention mechanism could be one explanation for the differing efficacy of interventions like mirror therapy (Barbin et al., 2016; Limakatso et al., 2023). Another factor could be the variability of protocols used for such interventions (Limakatso et al., 2025). Our model predictions emphasize the need of restoring controllability along with position likelihood, meaning that actively moving the phantom limb to match the movements of the mirror image is a crucial part of the intervention. Due to the lack of a standardized protocol for mirror therapy, the emphasis on the importance of phantom movements might vary between practitioners.

When searching for possible interventions informed by the facet plot in Fig. 3 we identified two beneficial movement directions in the plot and mapped them onto real-life interventions. Additionally, in several of the heat maps within the facet plot, fully ambiguous position likelihood (rightmost column) looks to be favorable. Following the same logic as earlier, this would indicate that rightward movement within heatmaps

might be another possible intervention. However, this type of intervention may be challenging to actualize. Recall that we let position observations be a compound of visual and proprioceptive input about the limbs position. Following limb loss, and in absence of a prosthesis, the visual input about the limb indeed does become fully ambiguous. As for the proprioceptive input, it is practically much harder to ensure that it is fully ambiguous. The severed nerves that signal proprioceptive information about the limb will likely innervate new tissue or form a neuroma, resulting in some form of proprioceptive input from stimulation of the tissue or spontaneous activity. It is difficult to ensure that these served nerves stay silent following amputation. Meanwhile, surgical techniques are being developed to promote meaningful proprioceptive feedback (Pettersen et al., 2023; Srinivasan et al., 2021; Gardetto et al., 2021), implicating diagonal up-and-leftward movement as a more feasible option for interventions.

For simplicity, we have assumed that the likelihood matrices for position observations and control matrices for movement are symmetric. In reality, severed nerves might reinnervate in such a way that they are more likely to signal one position than another and that some movements might be more difficult than others. Thus, more complex likelihood matrices could either exacerbate or reduce the prevalence of phantom limb pain. We provide some such examples in Appendix B.

Due to the omission of higher-level cognition, our model likely best describes subconscious interactions of pain and action rather than more complex, goal-directed behavior. For instance, it may capture automatic movement adjustments or avoidance responses in the presence of pain. In contrast, a more advanced (and arguably more realistic) agent might choose to act despite anticipating pain, if doing so serves a valued long-term goal. A classic example is persisting in a marathon despite muscle cramps due to the appeal of a medal and bragging rights. Sometimes, the long-term reward could even be the promise of pain relief. Challenging the subconsciously avoided movements in physical therapy might give pain in the moment, but pain relief in the long term. In the present model, 'reward' is treated in an abstract sense, representing desirable or goal-relevant outcomes rather than a specific physiological signal. We do not attempt to map this quantity onto a particular biological substrate, and the model does not imply that arbitrary external rewards (e.g., monetary incentives) would directly modulate phantom limb pain. Taken together, these scenarios highlight that the pain-action loop can involve complex trade-offs between immediate discomfort and delayed reward – dynamics that extend beyond the scope of our current model.

The omission of higher-level cognition also prevents our model from capturing the emotional dimension of pain. Pain reprocessing therapy (PRT), for instance, suggests that certain forms of chronic pain may be sustained by fear and other aversive emotional states (Ashar et al., 2022). This mechanism is likely relevant in phantom limb pain as well, particularly given the often traumatic circumstances surrounding limb loss. Recent work further indicates that emotional valence may be encoded in the perceived success or failure of actions (Hesp et al., 2021). Extending the model to incorporate valence could therefore provide insight into cases of phantom limb pain that are resistant to interventions, as well as into other chronic pain conditions not primarily driven by sensory disruption.

There is additional complexity in the interaction between control and pain not captured by our model. While it is difficult to disentangle the effects of controllability and predictability, there is growing evidence of a stronger effect of expectancy modulation when pain is controllable (Habermann et al., 2025). It is hypothesized that the sense of agency may play a key role in these dynamics (Habermann et al., 2025; Strube et al., 2023). Our model captures certain aspects of agency – specifically, the recognition of oneself as the causal agent of voluntary actions and their sensory outcomes – but it omits other relevant components such as self-efficacy and attentional processes (Habermann et al., 2025). The temporal relationship of action and outcome is crucial for the perception of agency (Habermann et al., 2025; Legaspi and Toyoizumi, 2019). Thus, in addition to the aforementioned inclusion of valence, incorporating temporal dynamics into the model may be essential for capturing the emotional and motivational mechanisms by which control and agency influence pain.

Because we employed a discrete state space, the model currently only captures the presence or absence of phantom limb pain. While this this type of simplification limits the biological realism of the model, it benefits the model in other ways by making it more explanatory and avoiding overfitting to details that might obscure the causal role of underlying processes. Nevertheless, there are a number of extensions to the model that could bring additonal insight. For example, extending the model to a continuous state space could allow for modeling of pain intensity. Another promising future direction would be to separate visual and proprioceptive input modalities. Sensory weighting may vary across individuals, which could in turn influence the effectiveness of different interventions. For example, some patients may benefit more from surgically restoring proprioceptive feedback, whereas others might respond better to visually based approaches such as mirror therapy. Furthermore, the current model only considers interventions where position observations are generated from the external environment. Yet, interventions based on internal processes – such as mental imagery – have also shown some efficacy in reducing phantom limb pain. If these effects operate via mechanisms similar to mirror or virtual reality interventions, they raise important questions about how to delineate the boundary between the agent and the environment should be represented in the model.

Finally, since our model has not been fit to empirical data, but manually tuned to qualitatively reproduce phenomenological characteristics and behavior, we can only offer conceptual predictions. Moving to a quantitative model that can produce real estimates would require two major steps. First, it must be determined how the model parameters can be quantified using real-world measures. For example, how are noxious observations quantified in the nervous system; is it by the firing rate of nociceptive neurons, or the ratio of activity in nociceptive and innocuous neurons? Answers pertaining to some of the model parameters might already be available in literature, while others likely would require a battery of physiological and psychophysical experiments to be determined. Second, once parameters have been tied to some real-world metrics, empirical data would need to be collected to fit these parameters in the model. This step requires data to be collected both from individuals with intact limbs and from amputees. For example, such measurements might involve microneurography in intact and residual limbs to measure activity in peripheral afferent neurons, or quantification of the susceptibility to sensory illusions to identify the individual likelihood mappings of different sensory modalities relating to limb position (Nyström and Hagbarth, 1981; Stone et al., 2018; Burrack and Brugger, 2005). In addition to these two steps, some of the extensions proposed in the previous paragraphs may also be necessary to enable mapping of model parameters onto real-life quantities.

4. Model derivation and implementation

To describe how the model is set up we follow the steps outlined in Chapter 6 of the book *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior* (Parr et al., 2022):

Step 1. Which system are we modeling?

In this first step, we need to identify the four core components of any active inference model: *the agent* (the generative model), *the environment* (the generative process), and the *sensory data* and *actions* that form the interface between the agent and the environment.

The generative process describes the true causal structure by which sensory data are generated from some hidden states. The hidden states may be influenced by actions performed by the agent. *The generative model* is a construct used by the agent to infer estimates of the hidden states based on observations of sensory data. To match these components onto pain, we recall the definition of pain: *An unpleasant sensory and emotional experience associated with, or resembling that associated with,*

actual or potential tissue damage (Raja et al., 2020). We suggest that, in the context of pain, the hidden state corresponds to tissue damage, and that pain corresponds to an agent's inferred estimate of tissue damage. It is worth noting that this mapping implies that the tissue for which the level of damage is being estimated belongs to the environment. Specifically, we define the agent as the CNS of the organism at hand, and the environment as the tissue of the body part the agent is making inferences about. The sensory interface between the agent and the environment thus becomes the relay point between the peripheral afferent sensory neurons of the body and the CNS.

So far, we have identified three out of four core concepts of our active inference model of pain: the agent, the environment and the sensory interface between the two. The fourth and final puzzle piece is the second part of the interface between the agent and the environment: actions. Since we have defined the agent to be the CNS, and the environment as the tissue of the body part of interest, the relevant actions are *actions initiated by the CNS that influence the state of the tissue*. While the CNS can affect the body through various pathways, we will focus on what is arguably the most direct causal link: motor control and movement.

Step 2. What is the most appropriate form for the generative model?

Once we have detailed exactly what system we are interested in modelling, we next need to consider what modelling format is most appropriate for describing the generative model. To this end, we must specify three attributes of the model: the type of state space, the hierarchical depth, and the temporal depth.

For the first attribute, we can choose between categorical (discrete) or continuous state space (or hybrid combining both). For this model, we have chosen to use a discrete state space due to its relative simplicity in computation. With this type of state space, state factors and observations can be described by categorical distributions. For example, the pain state can take one of two values, e.g., $\{\text{Pain}, \text{No pain}\}$, with probabilities $\{p, 1-p\}$. This type of state space would allow us to describe the probability of pain being present or absent. One of the main limitations with the discrete state space is that it does not allow us to describe e.g., pain intensity.

For the hierarchical depth, we need to consider whether all model variables evolve on the same or different time scales. Since we are interested in the development of chronic pain states, we would like to be able to describe both inference of states and of model parameters (i.e., learning). Thus, we need to include at least two levels of inference in our model.

As for the final attribute, the temporal depth, we are interested in modelling an agent that can anticipate consequences of actions or action sequences (policies or plans). To this end, we need to be able to describe planning capabilities in order to predict action-outcome contingencies – our model needs temporal depth. However, we will limit the model to only planning one time step ahead.

Step 3. How to set up the generative model?

Now that we have defined the system we are interested in and specified the key attributes of the generative model, it's time to dig into the specifics. In this step, we want to lay out the exact set-up of the generative model: What are the generative model's most appropriate variables and priors? Which parts are fixed and what must be learned?

In describing the generative model (and later the generative process) we will use the following terminology:

- *State factors* refer to independent sets of states (e.g., $\{\text{Pain}, \text{No pain}\}$)
- *Observation modalities* refer to independent sets of observable outcomes (e.g., activity in peripheral afferent neurons)

In step 1 we identified pain as the inferred estimate of tissue damage and that peripheral afferent neurons make up the sensory interface with the tissue. At any given time, the tissue in question will give rise to sensory input that is informative of the current level of tissue damage. The sensory input modality that is most closely linked to tissue damage is the activation of nociceptors, although input through other sensory channels (e.g., visual input) can also carry some information about the status of the tissue. Activity in peripheral afferent nociceptors is received by the generative model as *noxious observations* and is used to infer the pain state.

In step 1, we further defined actions to be movements of the body part in interest. One way of representing these movements in a discrete state space is by letting a state factor represent the *position* of the body part. Actions can then be thought of as movement between the different positions. For the agent to be able to efficiently plan and execute these movements, the agent will infer the current position based on *position observations*. Position observations can be thought of as a combination of proprioceptive and visual input about the limb's position.

So far, we have two state factors (position and pain) and two observation modalities (position and noxious). In the active inference framework, preferences are specified through the priors over observations. By doing so, undesirable observations become surprising and, through the surprise minimization paradigm that governs active inference, will be avoided. It is worth noting that these priors over observations serve a different role than the priors typically used in Bayesian inference. While such priors usually encode expectations about the state of the world, priors over observations instead encode preferences over sensory outcomes – i.e., which observations are desirable or undesirable – and thereby effectively play the role of goals in guiding action. Since tissue damage could threaten an organism's survival noxious observations indicative of such tissue damage should be undesirable. As for the position observation modality, it is not immediately apparent that any observation should be more or less desirable than any other. Thus, we let the agent have a neutral preference for all position observations.

With the generative model described so far, the behavior of the agent will be rather predictable: move to the position that minimizes noxious observations. The only reason the agent would ever deliberately act to increase noxious observations would be to reduce uncertainty about the world. Such an action might correspond to flexing a sprained ankle to figure out how bad the injury is. The action is likely to increase noxious observations, but also reduces uncertainty about the state of the ankle. However, as anyone who has experienced pain during an extended period of time can likely attest to, this type of action only makes up a small portion of the actions we take in the face of pain. Typically, the availability of actions that reduce noxious observations is limited and staying entirely passive while waiting for the tissue damage to heal is rarely an option since most organisms have other competing interests to attend. While minimizing noxious observations is desirable, there are other goals that also must be optimized for to ensure survival, such as scavenging for food, fleeing from a predator or going to the office to avoid loss of monetary income. These competing interests can, and often do, motivate the organism to take actions that increase undesirable noxious observations if they simultaneously generate alternate desirable *reward observations.*

To include these competing interests and alternate rewarding observations we extend the generative model in two ways. First, we add an additional observation modality to represent the *rewarding observations*. Second, we add an additional state factor (*reward state)* by which the agent infers if they are in a rewarding or non-rewarding state. For example, the rewarding observation could be finding a nutritious berry while foraging for food, by which the agent infers the probability that they are in a location with high or low berry-density (i.e., rewarding or non-rewarding state).

Now, we have three state factors and three observation modalities. For each state factor-observation modality pair (position-proprioceptive observations, pain-noxious observations, rewarding-reward observations), we have a *likelihood matrix* describing the agent's belief about the probability that an observation was generated by a specific state.

As described earlier, the actions available to the agent are movements between different positions. While the position is the only

controllable state factor, movement between positions might still influence the other two state factors. As an example, you can control the flexion of your ankle, and while you cannot independently control the pain, some positions of the ankle may be more painful than others. Thus, we must specify *state transition matrices* describing how the agent believes each action will influence each of the state factors. These state transition matrices determine how the agent expects states to evolve over time, thereby shaping its prior belief about future states. The state transition and likelihood matrices, along with the prior preferences over observations, are key ingredients in the agent's decision of which policy (sequence of actions) to follow. Fig. 7 shows an example of the generative model and its components.

Step 4. How to set up the generative process?

Just as for the generative model, we need to describe the generative process, i.e., what the hidden states are, how observations are generated and how the agent's actions influence the environment.

Following the same reasoning as in step 3 one could conclude that tissue damage, position and reward state all should be state factors of the generative process. While that certainly would be a viable approach, we have chosen to use an even simpler generative process to describe the environment. In our model position is the only state factor of the generative process, and we have let each position be associated with some probability for each of the three observation modalities. This means that we only need one state transition matrix to describe how the environment changes in response to the agent's actions.

Our reasoning for this simplification is that we primarily are interested in studying pain perception and action. In particular, we are interested in exploring how and when the perceived pain may dissociate from the noxious input. We have no interest in trying to describe tissue damage in detail. Thus, an environment that allows us to specify states with varying likelihood of noxious observations is sufficient for our purpose.

For our simulations we have chosen to let the position state factor have four different states, {Pos 1, Pos 2, Pos 3, Pos 4}, as this is the minimal environment that can account for all combinations of the noxious and reward observations : {No nox, No reward}, {No nox, Reward}, {Nox, No reward} and {Nox, Reward}. We allow the agent to move to any position from any other position (i.e., the agent does not have to go via Pos 2 from Pos 1 to reach Pos 3). The generative process of the default environment is visualized in Fig. 8.

We are particularly interested in studying how phantom limb pain may arise following limb loss. In our model, an injury directly affects the environment through changes to the generative process. An example of the generative process of the environment following limb loss is visualized in Fig. 9.

Below we provide some more context on the changes to the generative process:

Ambiguous noxious observations. Following limb loss, peripheral afferent nerves that used to innervate the limb become severed and are no longer able to reliably signal information about the state of the limb. Some neurons may find new tissue to innervate in the residual limb, regaining some stability in their signaling. However, a significant portion of neurons will likely either fall silent or become spontaneously active, causing some level of ambiguity in the signal arriving to the CNS. In the generative process, we represent this change as 1) increased ambiguity in the observations, and 2) equal likelihood for noxious observations at all four limb positions. We will consider three distinct cases corresponding to low, intermediate and high level of spontaneous activity in nociceptive peripheral afferent neurons.

Low reward observations. Following limb loss, we assume that agent's ability obtain reward observations with the affected limb is impaired (e.g., picking berries). To reflect this change, we set a low probability for reward observations at all positions.

Altered ability to perform movements. Since the limb is no longer present, the agent cannot efficiently move it to different positions. In fact, following limb loss the "position state" of the limb becomes ill defined in the generative process. To figure out how the state transition and position likelihood matrices might look, we can turn to an analogous scenario: regional anesthesia. When regional anesthesia is applied to a limb, the limb's position is still well defined, but the controllability and sensory input from the limb become altered. When the agent attempts to initiate a movement the limb partially or fully (depending on the intended movement and the level of anesthesia) stays unchanged. Mathematically, a completely stuck limb would correspond to the state transition matrices being identity matrices. However, turning back to limb loss, depending on the level of amputation and the reinnervation of the residual limb, there may exist some ability to control the limb position. As for position observations, visual input about the limb's position becomes fully ambiguous. Just as with the noxious observations, proprioceptive sensory input from the severed nerves also becomes

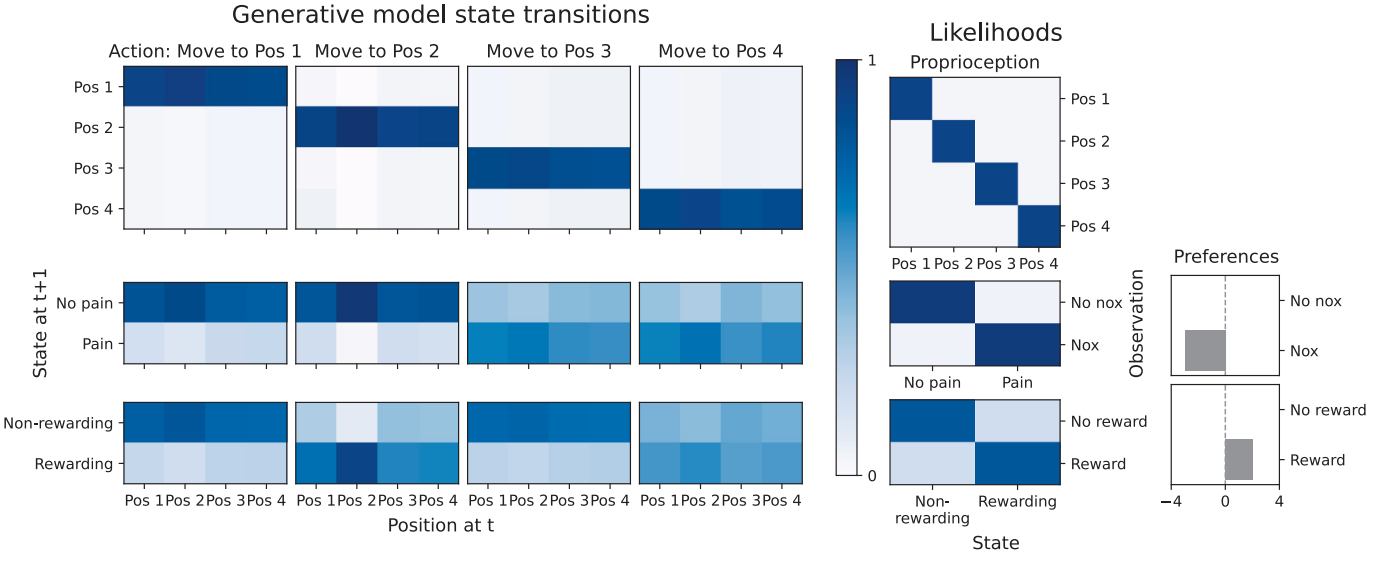

Fig. 7. Example generative model. The generative model consists of the expected state transitions (or control matrices) for each action, and likelihoods linking states to observations. In our model the agent makes inferences about three states: position, pain and reward state. The interface between the agent and the environment is made up of actions (movement of the limb) and sensory observations (noxious, position and reward observations). The agent has a preference for reward and against noxious observations.

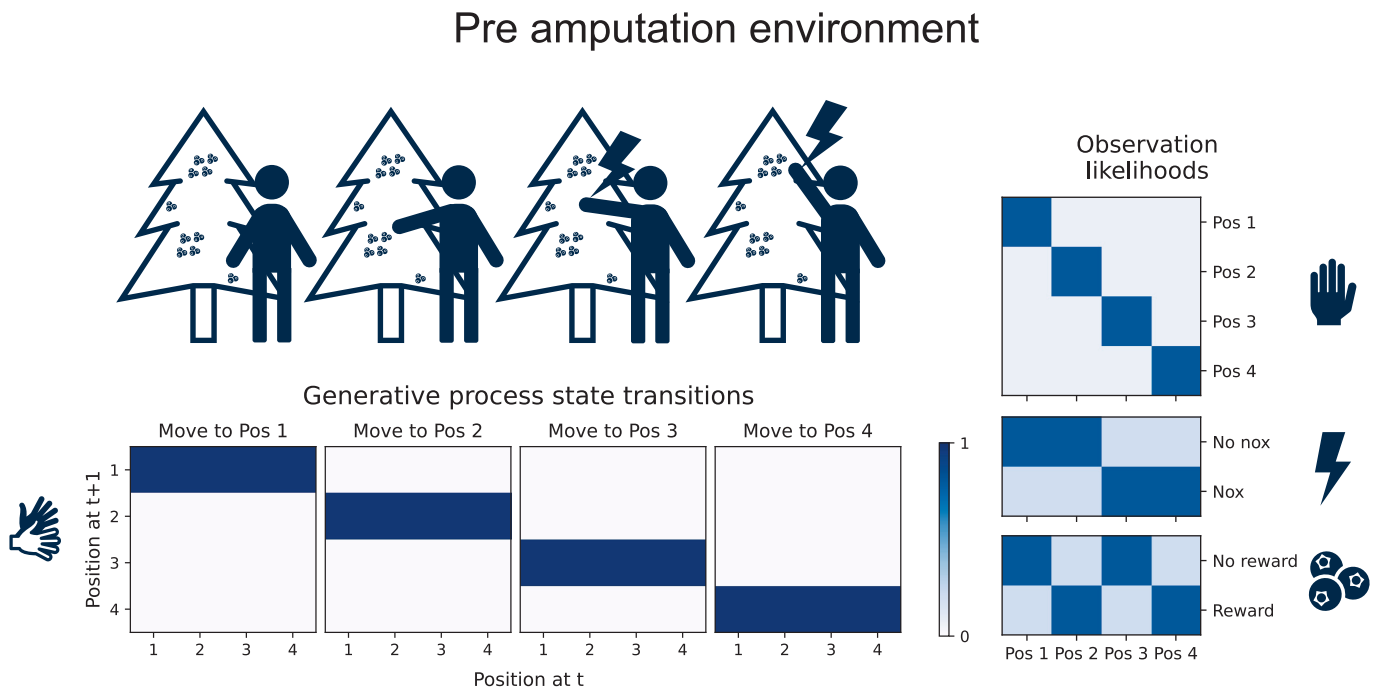

Fig. 8. Default generative process. For our model we have chosen to let the position state factor have four different states, {Pos 1, Pos 2, Pos 3, Pos 4}, as this is the minimal environment that can account for all combinations of the noxious and reward observations: {No nox, No reward}, {No nox, Reward}, {Nox, No reward} and {Nox, Reward}. We allow the agent to move to any position from any other position (i.e., the agent does not have to go via Pos 2 from Pos 1 to reach Pos 3). In the default (pre amputation) environment the state transition and likelihood matrices have high precision, allowing the agent to accurately perform movements and infer states from observations.

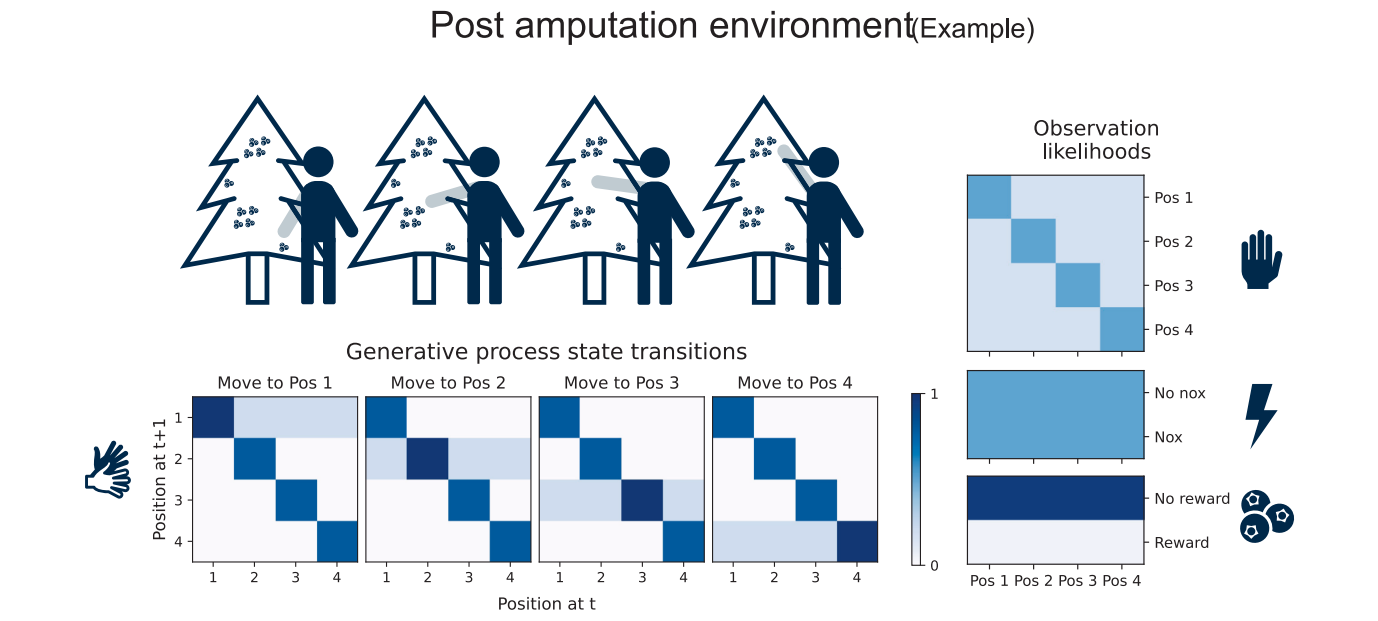

Fig. 9. Post amputation generative process. Amputation causes dramatic changes to the environment, affecting the agent's ability to control the limb position and altering the likelihoods of sensory observations. In the generative process these changes are reflected in the state transitions becoming a mixture of the original matrices (see Fig. 8) and the identity matrix, increased ambiguity in position and noxious observation likelihoods, and a low probability for reward observations at all positions.

ambiguous when they no longer have a limb to innervate. However, the level of ambiguity in proprioceptive feedback may vary depending on the level of amputation and reinnervation of the residual limb. In what follows, we will consider three levels of ambiguity: low ambiguity (as in the default environment) partial ambiguity, in which position input still carries some information, and full ambiguity, in which all limb positions generate the same observation.

Bonus step 5: How to implement the model?

Now that we have defined the generative model and process of the system we are interested in studying, we are almost ready to start simulating, visualizing and analyzing model outputs. However, before we dive into the results, we suggest one more step in addition to those suggested by Parr et al. (2022): describing what modelling choices are made in the specific model implementation. What method is used for

variational inferences? How is policy selection performed? How is learning implemented?

Our model implementation leverages the *pymdp* Python library for active inference in discrete state spaces (Heins et al., 2022). *pymdp* is a Python package that is inspired by and tested against the active inference routines contained in the widely used DEM toolbox of SPM (Friston et al., 2008). While *pymdp* allows for a high degree of customizability, we have mostly used the standard routines provided by the library. In what follows, we provide brief descriptions of these routines and specify where and how our implementation deviates.

Variational inference. For state inference we use the default inference algorithm (denoted 'VANILLA') provided by *pympd* which leverages Fixed Point Iteration to find the marginal posterior over hidden states.

Policy distribution. A policy is a sequence of actions the agent can take. The policy distribution is a probability distribution over the available policies and is computed as $\mathrm{softmax}(G \bullet \gamma)$, where G is the *expected free energy* for each policy and γ is the policy precision. Expected free energy is a fundamental quantity in the active inference framework, related to the governing principle that all living organisms follow the single objective of *minimizing the surprise of their sensory observations* (Parr et al., 2022). Expected free energy quantifies the expected *utility* (how well sensory observations correspond to desired outcomes or goals) and *information gain* (reduction of uncertainty about the world) of each policy. Thus, selecting the policy that minimizes the expected free energy corresponds to minimizing surprise. The relative contribution of utility and information gain in the expected free energy determines the level of exploration vs. exploitation in the agent's actions. The expected free energy is computed by 1) computing expected states under each policy, 2) computing expected observations based on the expected states, 3) calculating the expected utility based on the expected observations and the agent's prior preferences, 4) calculating the expected information gain about states (Bayesian surprise), 5) calculate the expected Dirichlet information gain over the learnable parameters, and finally 6) summing the utility, information gain over states and information gain over parameters.

Policy selection. Once the policy distribution has been computed, the agent must choose which policy to follow. This selection can be done either deterministically, by selecting the policy corresponding to the highest value in the distribution, or by stochastically sampling from the distribution. We have used deterministic policy selection in our simulations to limit behavioral variability not attributable to exploration or exploitation.

Learning. In our model implementation, the agent can update (learn) the entries of the state transition matrices based on experiences. This learning procedure is implemented through "learning by counting" by using Dirichlet priors as described in Smith et al. (2022). As Dirichlet concentration parameters grow over time, the agent becomes more and more confident in its beliefs. While this learning paradigm accurately reflects that an agent should grow more and more confident upon repeated confirmatory experiences, it can also lead to the agent becoming overconfident and slow to update their beliefs in a changing environment. One way to remedy this is by introducing a "forgetting factor", in which the concentration parameters are scaled by a factor $w_B \in (0, 1]$ (Smith et al., 2022). Note, however, that this forgetting factor doesn't imply that the agent "forgets" or changes its beliefs, but rather that it, in absence of additional confirmatory experiences, becomes less confident in its beliefs, allowing for more readily adaptive updates in light of new evidence. Applying a flat forgetting rate at all iterations may lead to some undesired behavior, such as intermittent bursts of exploratory behavior even when the environment is unchanged. Thus, we only apply the forgetting factor when $F - F_{avg} > \theta$, where F is the current variational free energy, F_{avg} is a moving average of the variational free energy within a predefined window, and θ is some threshold. In the active inference framework, variational free energy is a proxy for surprise in cases where the exact surprise cannot be computed. Put into plain words, the forgetting factor is only applied when data are excessively surprising under the current belief, signaling that the current generative model is inadequate. Inference of the volatility of the environment can be implemented in a more principled manner, for example in using a hierarchical model (Liakoni et al., 2021). However, this simpler paradigm allows us to incorporate some surprised-based adaptation without adding much complexity to the model.

Observations. When dealing with discrete state spaces, observations generated by the environment can be received by the agent either as point-samples from the observation distribution, or as the entire distribution. In the former case there is no ambiguity in the observation itself (you either observe a nutritious berry or you do not), but there could still be some ambiguity in how that observation relates to the hidden state (how likely is a single berry-observation to have come from a high- vs low-density berry region). The latter case allows for ambiguity to also be baked into the observation itself (e.g., observing a sound in a noisy environment). Since we treat noxious observations as the incoming signals from peripheral afferent neurons, we expect there to be some ambiguity in these observations. Thus, we treat observations as distributions (by setting the `distr_obs` flag to `True`).

Code availability

Python code for all simulation results presented here are available online at: https://doi.org/10.5281/zenodo.17395215. Additional results and pseudocode of the active inference algorithm can be found in the appendices.

CRediT authorship contribution statement

Malin Ramne: Writing – review & editing, Writing – original draft, Software, Conceptualization. **Torbjörn Lundh:** Writing – review & editing, Supervision. **Jon Sensinger:** Writing – review & editing, Supervision, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Appendix 1

Appendix A. Active inference pseudocode

Table A1
Variables and parameters of the generative process.

State factor 0: Position	$s^{(0)}$	$[1,0,0,0]^T$ for Pos 1 $[0,1,0,0]^T$ for Pos 2 $[0,0,1,0]^T$ for Pos 3 $[0,0,0,1]^T$ for Pos 4
Actions	a	1 for Move to Pos 1 2 for Move to Pos 2 3 for Move to Pos 3 4 for Move to Pos 4
State transition factor 0: Position	$B_a^{(0)}$	$\begin{bmatrix} p\left(s_{t+1}{}^{(0)} = \text{Pos 1} \middle\| s_t^{(0)} = \text{Pos 1}, a\right) & \cdots & p\left(s_{t+1}{}^{(0)} = \text{Pos 1} \middle\| s_t^{(0)} = \text{Pos 4}, a\right) \\ \vdots & \ddots & \vdots \\ p\left(s_{t+1}{}^{(0)} = \text{Pos 4} \middle\| s_t^{(0)} = \text{Pos 1}, a\right) & \cdots & p\left(s_{t+1}{}^{(0)} = \text{Pos 4} \middle\| s_t^{(0)} = \text{Pos 4}, a\right) \end{bmatrix}$
Observation likelihood modality 0: Position observations	$A^{(0)}$	$\begin{bmatrix} p(o^{(0)} = \text{Pos 1} \mid s^{(0)} = \text{Pos 1}) & \cdots & p(o^{(0)} = \text{Pos 1} \mid s^{(0)} = \text{Pos 4}) \\ \vdots & \ddots & \vdots \\ p(o^{(0)} = \text{Pos 4} \mid s^{(0)} = \text{Pos 1}) & \cdots & p(o^{(0)} = \text{Pos 4} \mid s^{(0)} = \text{Pos 4}) \end{bmatrix}$
Observation likelihood modality 1: Noxious observations	$A^{(1)}$	$\begin{bmatrix} p(o^{(1)} = \text{No nox} \mid s^{(0)} = \text{Pos 1}) & \cdots & p(o^{(1)} = \text{No nox} \mid s^{(0)} = \text{Pos 4}) \\ p(o^{(1)} = \text{Nox} \mid s^{(0)} = \text{Pos 1}) & \cdots & p(o^{(1)} = \text{Nox} \mid s^{(0)} = \text{Pos 4}) \end{bmatrix}$
Observation likelihood modality 2: Reward observations	$A^{(2)}$	$\begin{bmatrix} p(o^{(1)} = \text{No rew} \mid s^{(0)} = \text{Pos 1}) & \cdots & p(o^{(1)} = \text{No rew} \mid s^{(0)} = \text{Pos 4}) \\ p(o^{(1)} = \text{Reward} \mid s^{(0)} = \text{Pos 1}) & \cdots & p(o^{(1)} = \text{Reward} \mid s^{(0)} = \text{Pos 4}) \end{bmatrix}$
Observation modality 0: Position observations	$o^{(0)}$	$A^{(0)}s^{(0)}$
Observation modality 1: Noxious observations	$o^{(1)}$	$A^{(1)}s^{(0)}$
Observation modality 0: Reward observations	$o^{(2)}$	$A^{(2)}s^{(0)}$

Table A2
Variables and parameters of the generative process.

Posterior estimate of state factor 0: Position	$\widehat{s}^{(0)}$	$\begin{bmatrix} p(s^{(0)} = \text{Pos 1} \mid o^{(0)}) \\ p(s^{(0)} = \text{Pos 2} \mid o^{(0)}) \\ p(s^{(0)} = \text{Pos 3} \mid o^{(0)}) \\ p(s^{(0)} = \text{Pos 4} \mid o^{(0)}) \end{bmatrix}$
Posterior estimate of state factor 1: Pain state	$\widehat{s}^{(1)}$	$\begin{bmatrix} p(s^{(1)} = \text{No pain} \mid o^{(1)}) \\ p(s^{(1)} = \text{Pain} \mid o^{(1)}) \end{bmatrix}$
Posterior estimate of state factor 2: Reward state	$\widehat{s}^{(2)}$	$\begin{bmatrix} p(s^{(2)} = \text{Non-rewarding} \mid o^{(2)}) \\ p(s^{(2)} = \text{Rewarding} \mid o^{(2)}) \end{bmatrix}$
Actions	a	1 for Move to Pos 1 2 for Move to Pos 2 3 for Move to Pos 3 4 for Move to Pos 4
Believed state transition factor 0: Position	$\widehat{B}_a^{(0)}$	$\begin{bmatrix} p\left(\widehat{s}_{t+1}{}^{(0)} = \text{Pos 1} \middle\| \widehat{s}_t^{(0)} = \text{Pos 1}, a\right) & \cdots & p\left(\widehat{s}_{t+1}{}^{(0)} = \text{Pos 1} \middle\| \widehat{s}_t^{(0)} = \text{Pos 4}, a\right) \\ \vdots & \ddots & \vdots \\ p\left(\widehat{s}_{t+1}{}^{(0)} = \text{Pos 4} \middle\| \widehat{s}_t^{(0)} = \text{Pos 1}, a\right) & \cdots & p\left(\widehat{s}_{t+1}{}^{(0)} = \text{Pos 4} \middle\| \widehat{s}_t^{(0)} = \text{Pos 4}, a\right) \end{bmatrix}$
Believed state transition factor 1: Pain state	$\widehat{B}_a^{(1)}$	$\begin{bmatrix} p\left(\widehat{s}_{t+1}{}^{(1)} = \text{No pain} \middle\| \widehat{s}_t^{(0)} = \text{Pos 1}, a\right) & \cdots & p\left(\widehat{s}_{t+1}{}^{(1)} = \text{No pain} \middle\| \widehat{s}_t^{(0)} = \text{Pos 4}, a\right) \\ p\left(\widehat{s}_{t+1}{}^{(1)} = \text{Pain} \middle\| \widehat{s}_t^{(0)} = \text{Pos 1}, a\right) & \cdots & p\left(\widehat{s}_{t+1}{}^{(1)} = \text{Pain} \middle\| \widehat{s}_t^{(0)} = \text{Pos 4}, a\right) \end{bmatrix}$
Believed state transition factor 2: Reward state	$\widehat{B}_a^{(2)}$	$\begin{bmatrix} p\left(\widehat{s}_{t+1}{}^{(2)} = \text{Non-rew} \middle\| \widehat{s}_t^{(0)} = \text{Pos 1}, a\right) & \cdots & p\left(\widehat{s}_{t+1}{}^{(2)} = \text{Non-rew} \middle\| \widehat{s}_t^{(0)} = \text{Pos 4}, a\right) \\ p\left(\widehat{s}_{t+1}{}^{(2)} = \text{Rewarding} \middle\| \widehat{s}_t^{(0)} = \text{Pos 1}, a\right) & \cdots & p\left(\widehat{s}_{t+1}{}^{(2)} = \text{Rewarding} \middle\| \widehat{s}_t^{(0)} = \text{Pos 4}, a\right) \end{bmatrix}$
Observation likelihood modality 0: Position observations	$\widehat{A}^{(0)}$	$\begin{bmatrix} p(o^{(0)} = \text{Pos 1} \mid \widehat{s}^{(0)} = \text{Pos 1}) & \cdots & p(o^{(0)} = \text{Pos 1} \mid \widehat{s}^{(0)} = \text{Pos 4}) \\ \vdots & \ddots & \vdots \\ p(o^{(0)} = \text{Pos 4} \mid \widehat{s}^{(0)} = \text{Pos 1}) & \cdots & p(o^{(0)} = \text{Pos 4} \mid \widehat{s}^{(0)} = \text{Pos 4}) \end{bmatrix}$
Observation likelihood modality 1: Noxious observations	$\widehat{A}^{(1)}$	$\begin{bmatrix} p(o^{(1)} = \text{No nox} \mid \widehat{s}^{(1)} = \text{No pain}) & p(o^{(1)} = \text{No nox} \mid \widehat{s}^{(1)} = \text{Pain}) \\ p(o^{(1)} = \text{Nox} \mid \widehat{s}^{(1)} = \text{No pain}) & p(o^{(1)} = \text{Nox} \mid \widehat{s}^{(1)} = \text{Pain}) \end{bmatrix}$
Observation likelihood modality 2: Reward observations	$\widehat{A}^{(2)}$	$\begin{bmatrix} p(o^{(2)} = \text{No rew} \mid \widehat{s}^{(2)} = \text{Non-rew}) & p(o^{(2)} = \text{No rew} \mid \widehat{s}^{(2)} = \text{Rewarding}) \\ p(o^{(2)} = \text{Rew} \mid \widehat{s}^{(2)} = \text{Non-rew}) & p(o^{(2)} = \text{Rew} \mid \widehat{s}^{(2)} = \text{Rewarding}) \end{bmatrix}$

(*continued on next page*)

Table A2 (*continued*)

Posterior estimate of state factor 0: Position	$\widehat{s}^{(0)}$	$\begin{bmatrix} p(s^{(0)} = \text{Pos 1} \mid o^{(0)}) \\ p(s^{(0)} = \text{Pos 2} \mid o^{(0)}) \\ p(s^{(0)} = \text{Pos 3} \mid o^{(0)}) \\ p(s^{(0)} = \text{Pos 4} \mid o^{(0)}) \end{bmatrix}$
Preferences for observation modality 0	$C^{(0)}$	0 for Pos 1 0 for Pos 2 0 for Pos 3 0 for Pos 4
Preferences for observation modality 1	$C^{(1)}$	0 for No nox −3 for Nox
Preferences for observation modality 2	$C^{(2)}$	0 for No rew 2 for Reward
Prior estimate of state factor 0	$\widehat{D}^{(0)}$	$\begin{bmatrix} p\left(s_0^{(0)} = \text{Pos 1}\right) \\ p\left(s_0^{(0)} = \text{Pos 2}\right) \\ p\left(s_0^{(0)} = \text{Pos 3}\right) \\ p\left(s_0^{(0)} = \text{Pos 4}\right) \end{bmatrix}$
Prior estimate of state factor 1	$\widehat{D}^{(1)}$	$\begin{bmatrix} p\left(s_0^{(1)} = \text{No pain}\right) \\ p\left(s_0^{(1)} = \text{Pain}\right) \end{bmatrix}$
Prior estimate of state factor 2	$\widehat{D}^{(2)}$	$\begin{bmatrix} p\left(s_0^{(2)} = \text{Non-rewarding}\right) \\ p\left(s_0^{(2)} = \text{Rewarding}\right) \end{bmatrix}$
Learning rate	η_B	1
Forgetting rate	w_B	0.97
Forgetting threshold	θ_B	0.2
Policy precision	γ	10

Justification of model parameter values:

The parameter values were selected to yield qualitatively plausible behaviour in the learning and context-switching simulations, in particular a balanced trade-off between exploration and exploitation.

Preferences. We assigned a positive preference for reward observations (2) and a slightly stronger negative preference for noxious observations (−3). This asymmetry reflects the focus of the model on scenarios of unintentional chronic pain (rather than the agent actively seeking a state of pain to obtain some highly valued rewrd), where avoidance of noxious states is expected to dominate over reward-seeking.

Learning and forgetting parameters. Learning rate (η_B), forgetting rate (w_B) and forgetting threshold (θ_B) were chosen to balance adaptability and stability. Excessive forgetting leads to over-exploratory, near-random behaviour, whereas insufficient forgetting (or excessive learning) results in overly persistent beliefs that are slow to adapt when the environment changes.

We emphasize that these parameters are not uniquely determined, and we make no claim that this is the only viable parameter set. Rather, they represent one configuration that produces qualitatively expected behaviour in a well-characterized scenario. While the precise values influence the quantitative details of the simulations, the qualitative patterns reported here were found to be robust across a reasonable range of parameter settings.

Table A3

Pseudocode for the active inference algorithm used in the simulations. The full code implementation can be found at *https://doi.org/10.5281/zenodo.17395215*. Notation: hats (e.g., $\widehat{s}$) denote estimates, bars (e.g., $\bar{s}$) denote priors, superscripts denote dimension (0,1,2), and the subscript denotes time.

(*continued on next page*)

Table A3 (*continued*)

while $t < n$

State inference

Make prior estimate of states:

$\hat{\tilde{s}}_{t+1}^{(0)} = \hat{B}_a^{(0)} \hat{s}_t^{(0)}$ if $t = 0$: $\hat{\tilde{s}}_1^{(0)} = \hat{D}^{(0)}$

$\hat{\tilde{s}}_{t+1}^{(1)} = \hat{B}_a^{(1)} \hat{s}_t^{(0)}$ if $t = 0$: $\hat{\tilde{s}}_1^{(1)} = \hat{D}^{(1)}$

$\hat{\tilde{s}}_{t+1}^{(2)} = \hat{B}_a^{(2)} \hat{s}_t^{(0)}$ if $t = 0$: $\hat{\tilde{s}}_1^{(2)} = \hat{D}^{(2)}$

Get observations from the environment:

$o_t^{(0)} = A^{(0)} s_t^{(0)}$

$o_t^{(1)} = A^{(1)} s_t^{(0)}$

$o_t^{(2)} = A^{(2)} s_t^{(0)}$

Make posterior estimate of states using Fixed Point Iteration, `FPI`:

$\hat{s}_{t+1}^{(0)}, \hat{s}_{t+1}^{(1)}, \hat{s}_{t+1}^{(2)}, F =$ `FPI(` $\hat{\tilde{s}}_{t+1}^{(0)}, \hat{\tilde{s}}_{t+1}^{(1)}, \hat{\tilde{s}}_{t+1}^{(2)}, o_t^{(0)}, o_t^{(1)}, o_t^{(2)}, \hat{A}^{(0)}, \hat{A}^{(1)}, \hat{A}^{(2)}$`)`

Learning and forgetting

Reduce confidence in $\hat{B}_a^{(0)}, \hat{B}_a^{(1)}, \hat{B}_a^{(2)}$ if relative surprise exceeds threshold:

if $F - F_{avg} > \theta_B$

Multiply Dirichlet priors of $\hat{B}_a^{(0)}, \hat{B}_a^{(1)}, \hat{B}_a^{(2)}$ with forgetting factor w_B

end if

Update Dirichlet priors of $\hat{B}_a^{(0)}, \hat{B}_a^{(1)}, \hat{B}_a^{(2)}$ based on state prediction errors

Acting

Infer the policy distribution π:

Compute expected states under each possible policy

Compute expected observations based on the expected states

Calculate the expected utility based on the expected observations and the agent's prior preferences

Calculate the expected information gain about states

Calculate the expected Dirichlet information gain over $\hat{B}_a^{(0)}, \hat{B}_a^{(1)}, \hat{B}_a^{(2)}$

Sum the utility, information gain over states and information gain over parameters to obtain the expected free energy, G

Compute policy distribution $\pi =$ `softmax(`γG`)`

Select the action a that corresponds to `argmax(`π`)`

Act (update the environment based on the selected action):

Compute the distribution of future states: $p(s_{t+1}^{(0)}) = B_a^{(0)} s_t^{(0)}$

Sample the next state from the distribution: $s_{t+1}^{(0)} \sim p(s_{t+1}^{(0)})$

end while

Appendix B. Additional simulation results

Context switching in more detail.

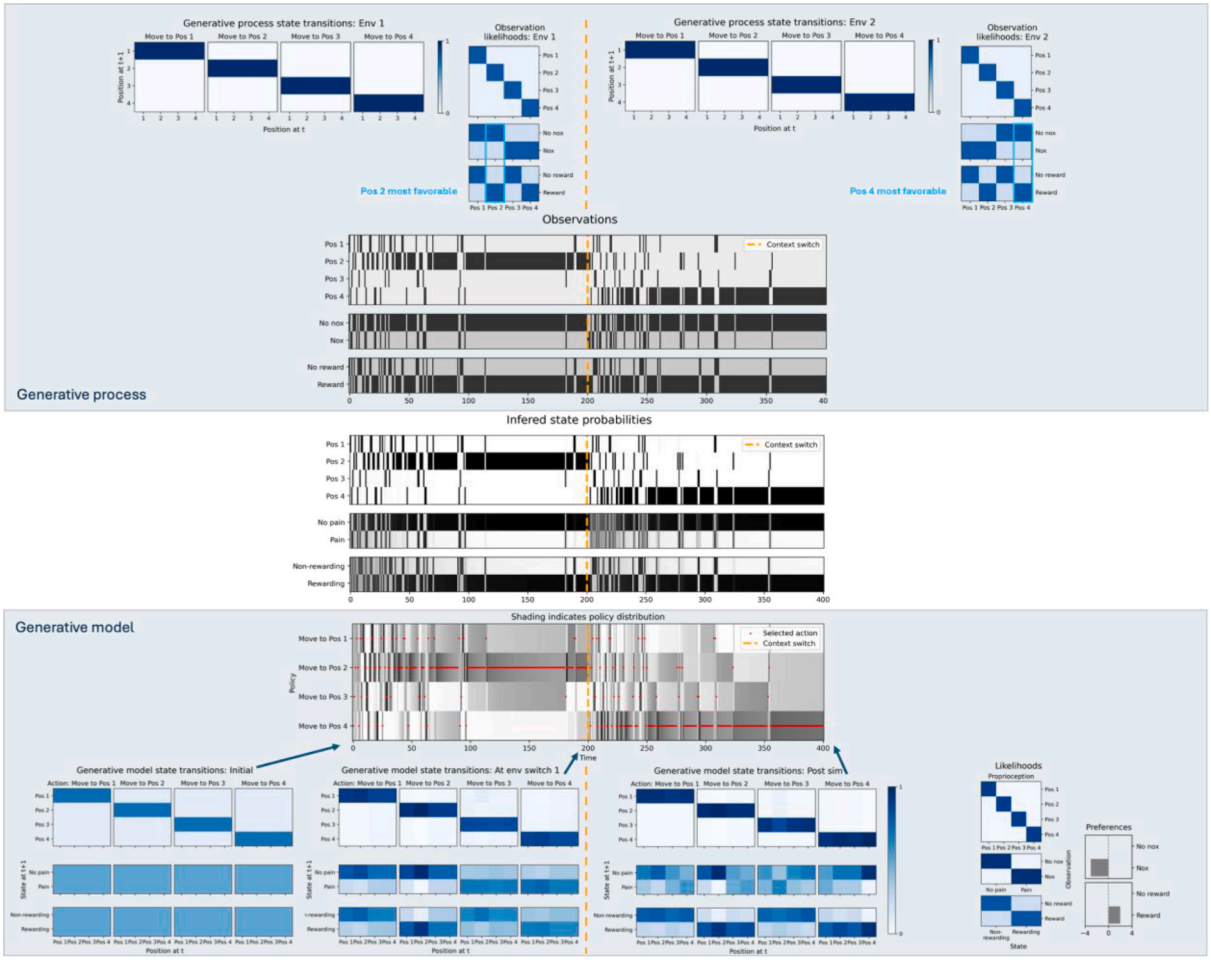

Fig. B1. *A more detailed depiction of the parameters and processes involved in the context switching-simulation in Fig. 1 of the manuscript. At the top the full set of parameters of generative process before and after context switch are displayed, along with the observations generated from the environment. At the bottom the generative model at a few key time points (at simulation onset, just before context switch and at the end of the simulation) are shown. Note that the likelihood matrices and the preferences remain unchanged throughout the simulation, hence they are only displayed once. The generative model also includes the inferred policy distribution at each time point, along with the selected action.*

Prevalence of phantom limb pain – additional results.

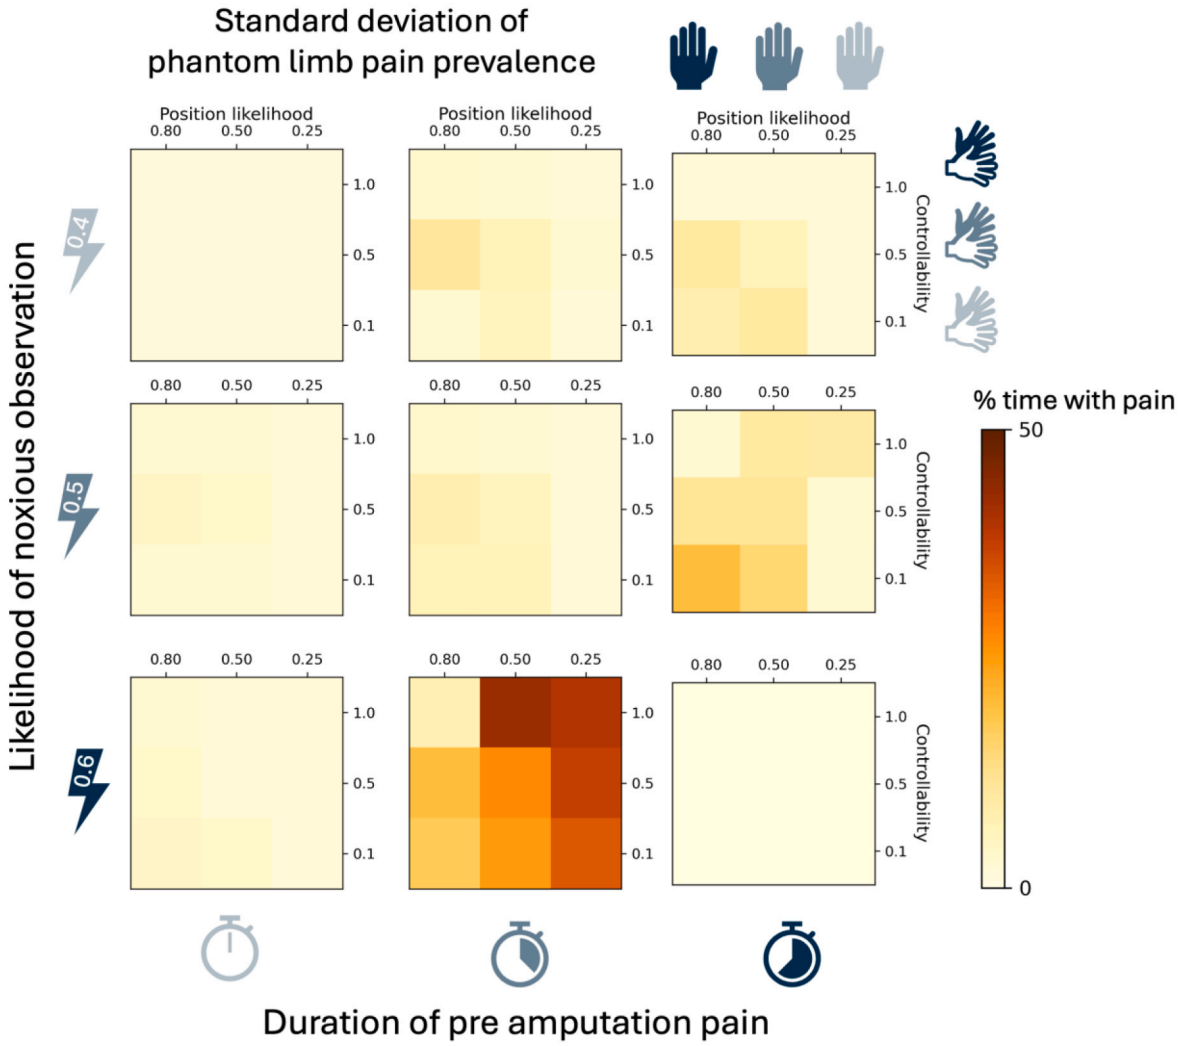

Fig. B2. *Standard deviation of phantom limb pain, calculated using the same simulated data as in Fig. 2 of the manuscript.*

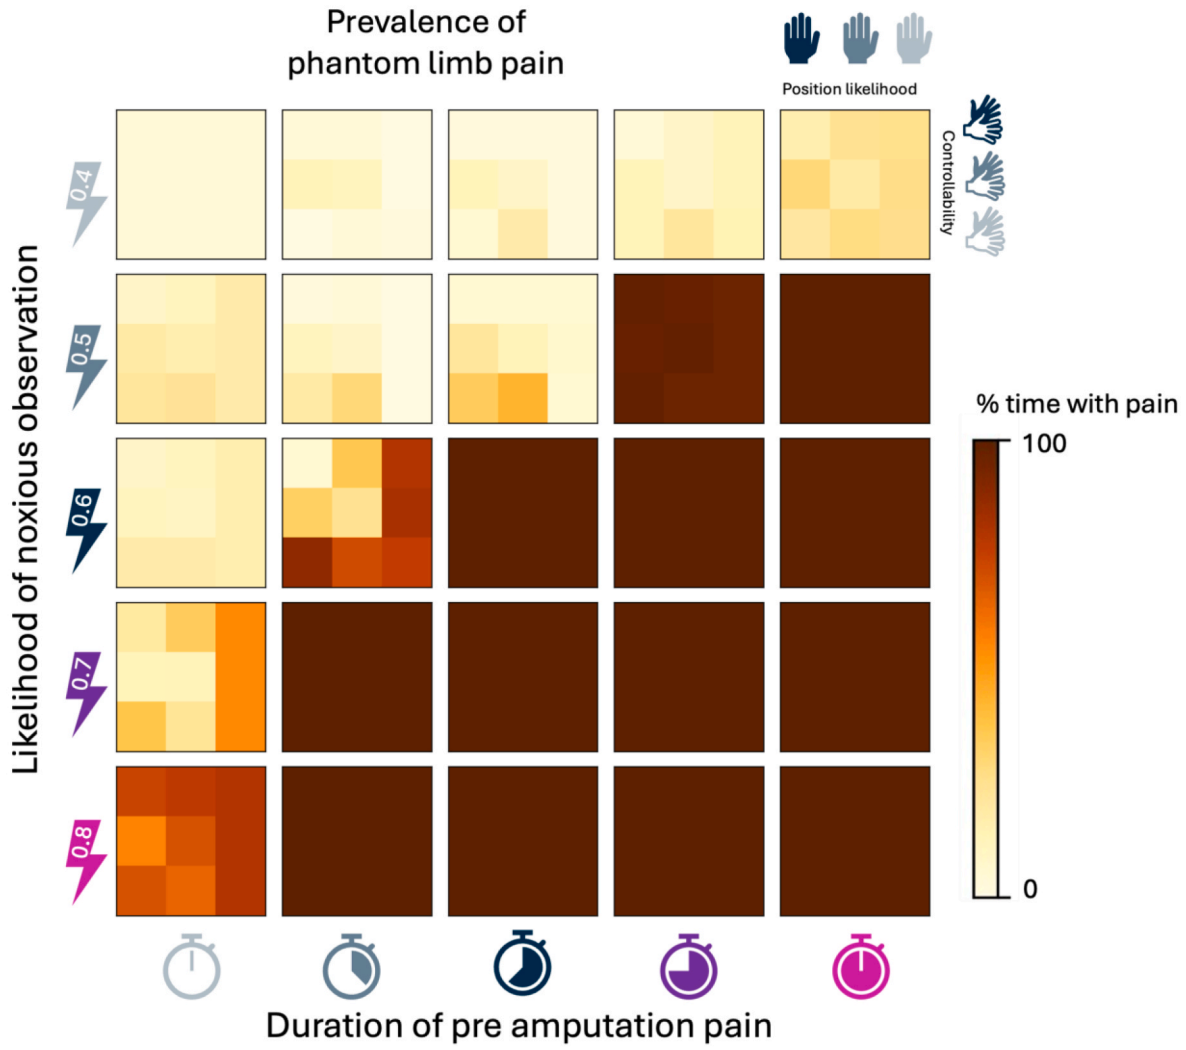

Fig. B3. *Similar results as presented in Fig. 2 of the main manuscript, but with additional simulations for higher likelihoods of noxious observations (0.7 and 0.8) and longer times of pre amputation pain. The added parameter values are indicated with purple and pink colors of the icons.* ***These results further strengthen the finding that residual noxious input and duration of pre amputation pain are risk factors for phantom limb pain.*** *The values for position likelihood and controllability are unchanged as they already cover the full range of the corresponding parameters.*

Examples of individual simulations of phantom limb pain

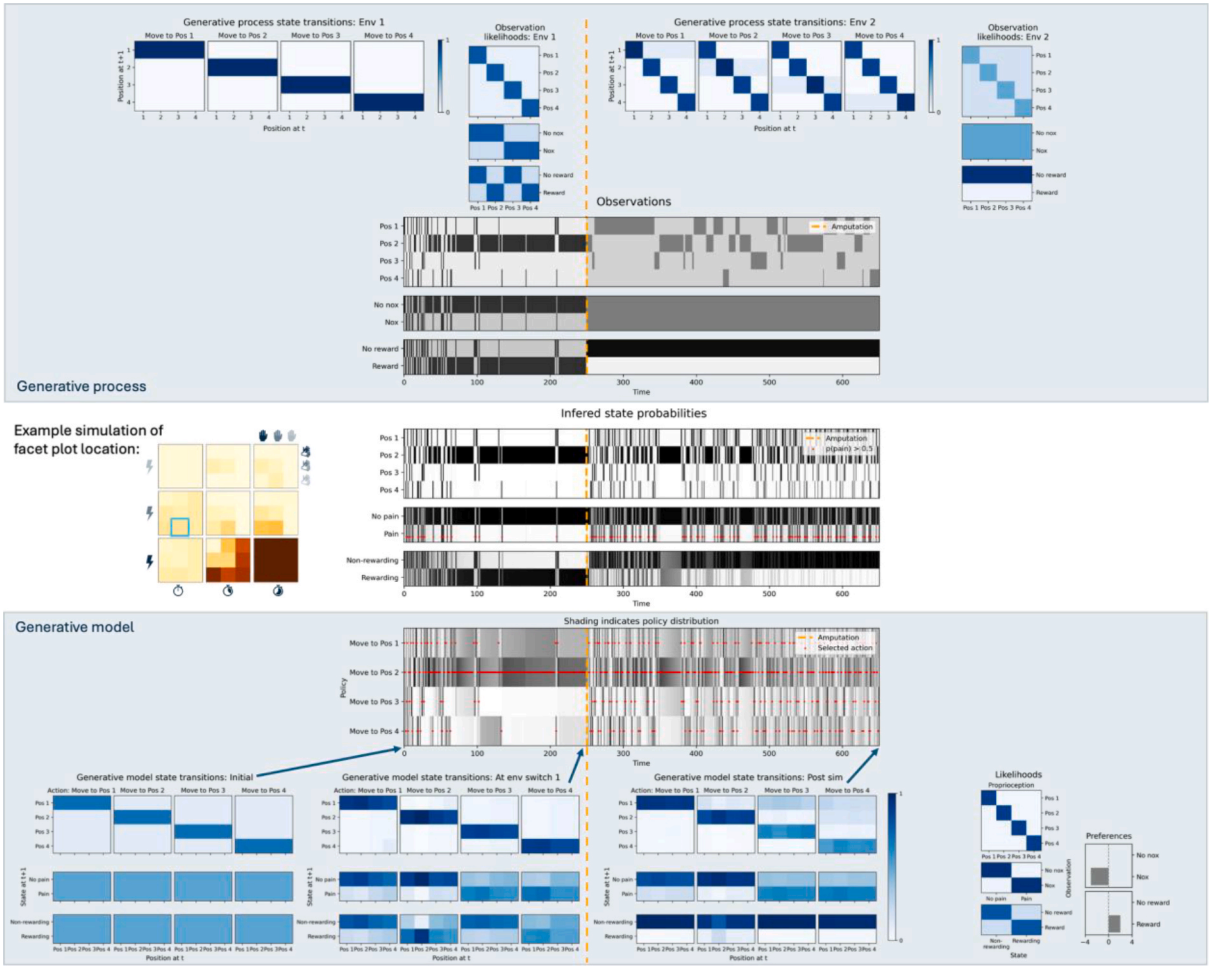

Fig. B4. *Example simulation of the indicated facet plot location, demonstrating that* ***even though the overall prevalence of phantom limb pain may be relatively low, individuals without pre amputation pain may also experience episodes of phantom limb pain*** *(indicated by the red dots in the middle-most figure of Inferred state probabilities). At the top the full set of parameters of generative process before and after amputation are displayed, along with the observations generated from the environment. At the bottom the generative model at a few key time points (at simulation onset, just before amputation and at the end of the simulation) are shown. Note that the likelihood matrices and the preferences remain unchanged throughout the simulation, hence they are only displayed once. The generative model also includes the inferred policy distribution at each time point, along with the selected action.*

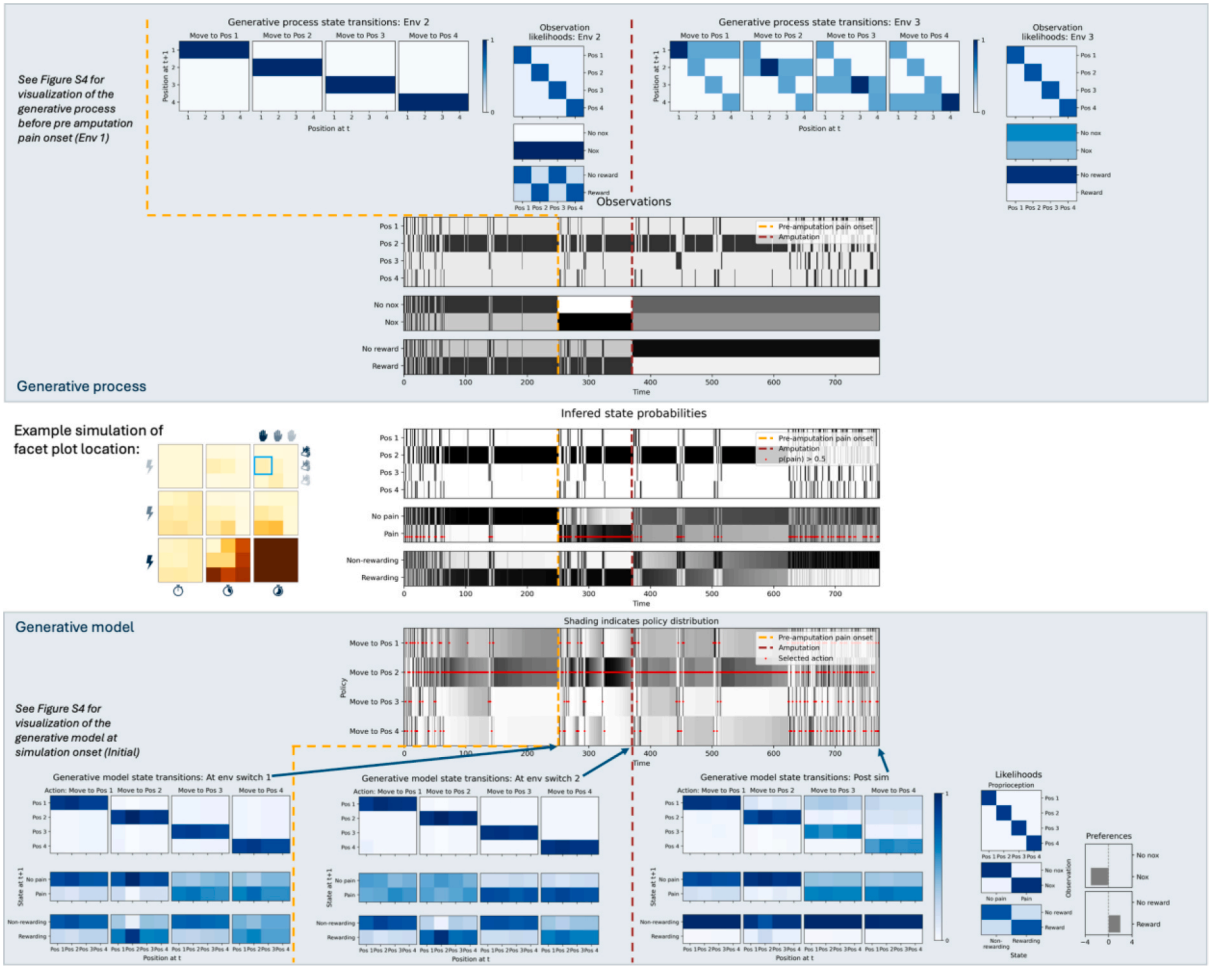

Fig. B5. *Example simulation of the indicated facet plot location, demonstrating that* ***even though the overall prevalence of phantom limb pain may be relatively low, individuals without low likelihood of noxious observations may also experience episodes of phantom limb pain*** *(indicated by the red dots in the middle-most figure of Inferred state probabilities). At the top the generative process during the period with pre amputation pain before amputation and after amputation are displayed, along with the observations generated from the environment. The reader is referred to for a visualization of the generative process parameters during the period before pre amputation pain onset. At the bottom the generative model at a few key time points (at pre amputation pain onset, just before amputation and at the end of the simulation) are shown. The reader is referred to for a visualization of the generative model parameters at simulation onset. Note that the likelihood matrices and the preferences remain unchanged throughout the simulation, hence they are only displayed once. The generative model also includes the inferred policy distribution at each time point, along with the selected action.*

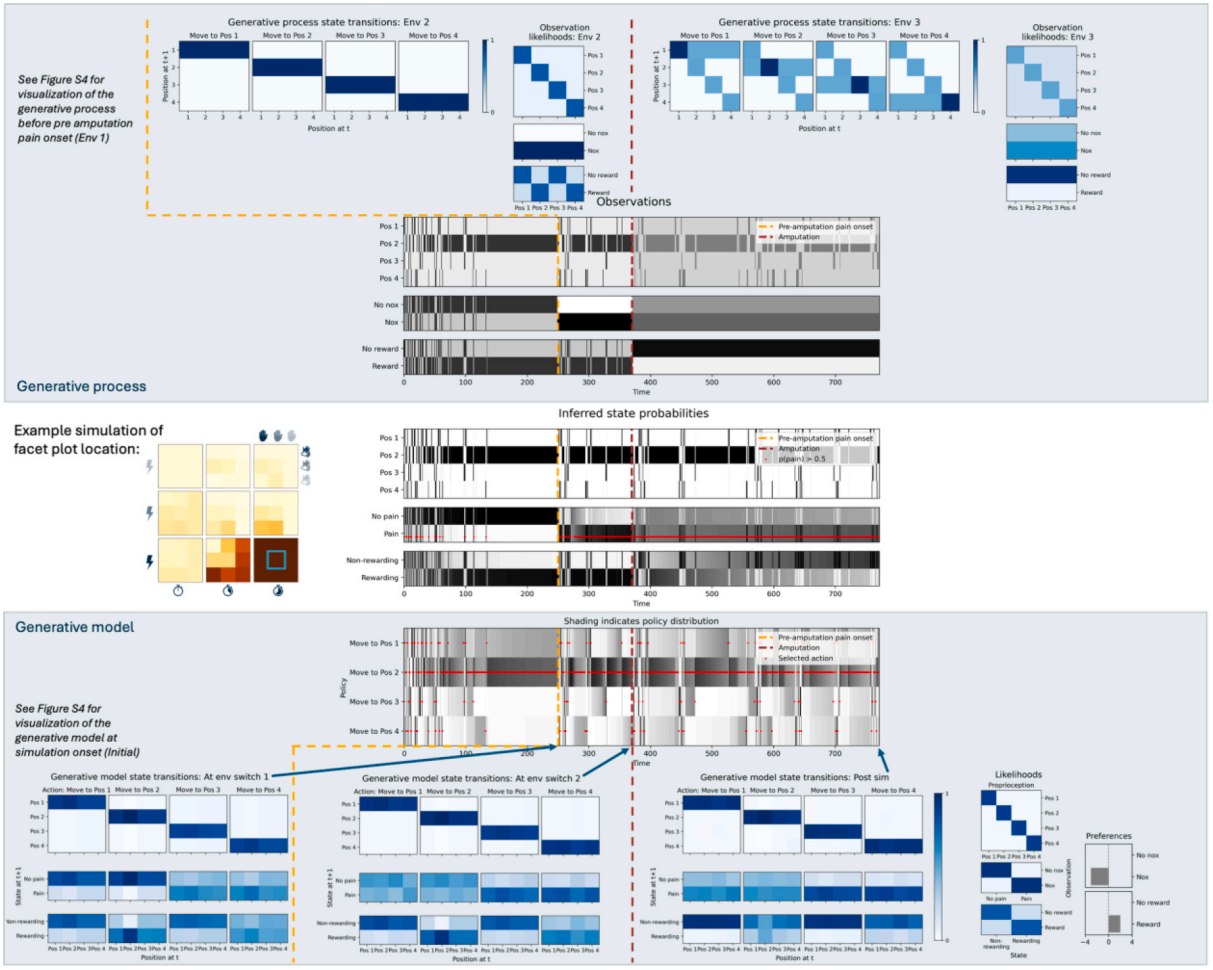

Fig. B6. ***Example simulation of the indicated facet plot location where prevalence of phantom limb pain is high****. At the top the generative process during the period with pre amputation pain before amputation and after amputation are displayed, along with the observations generated from the environment. The reader is referred to for a visualization of the generative process parameters during the period before pre amputation pain onset. At the bottom the generative model at a few key time points (at pre amputation pain onset, just before amputation and at the end of the simulation) are shown. The reader is referred to for a visualization of the generative model parameters at simulation onset. Note that the likelihood matrices and the preferences remain unchanged throughout the simulation, hence they are only displayed once. The generative model also includes the inferred policy distribution at each time point, along with the selected action.*

Additional details and examples of interventions.

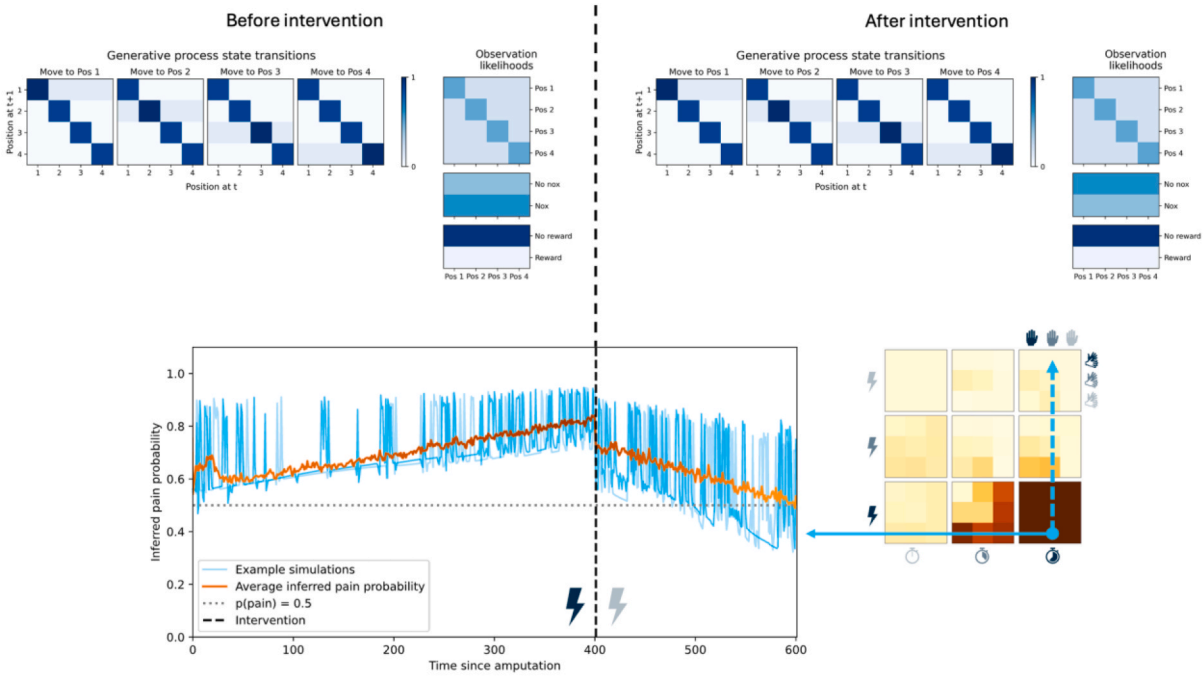

Fig. B7. ***Details on the generative process before and after the intervention shown in*** Fig. 3 *of the main manuscript. Such interventions might correspond to applying a nerve block to spontaneously active neuromas or the dorsal root ganglion, surgical interventions, or peripheral nerve stimulation*

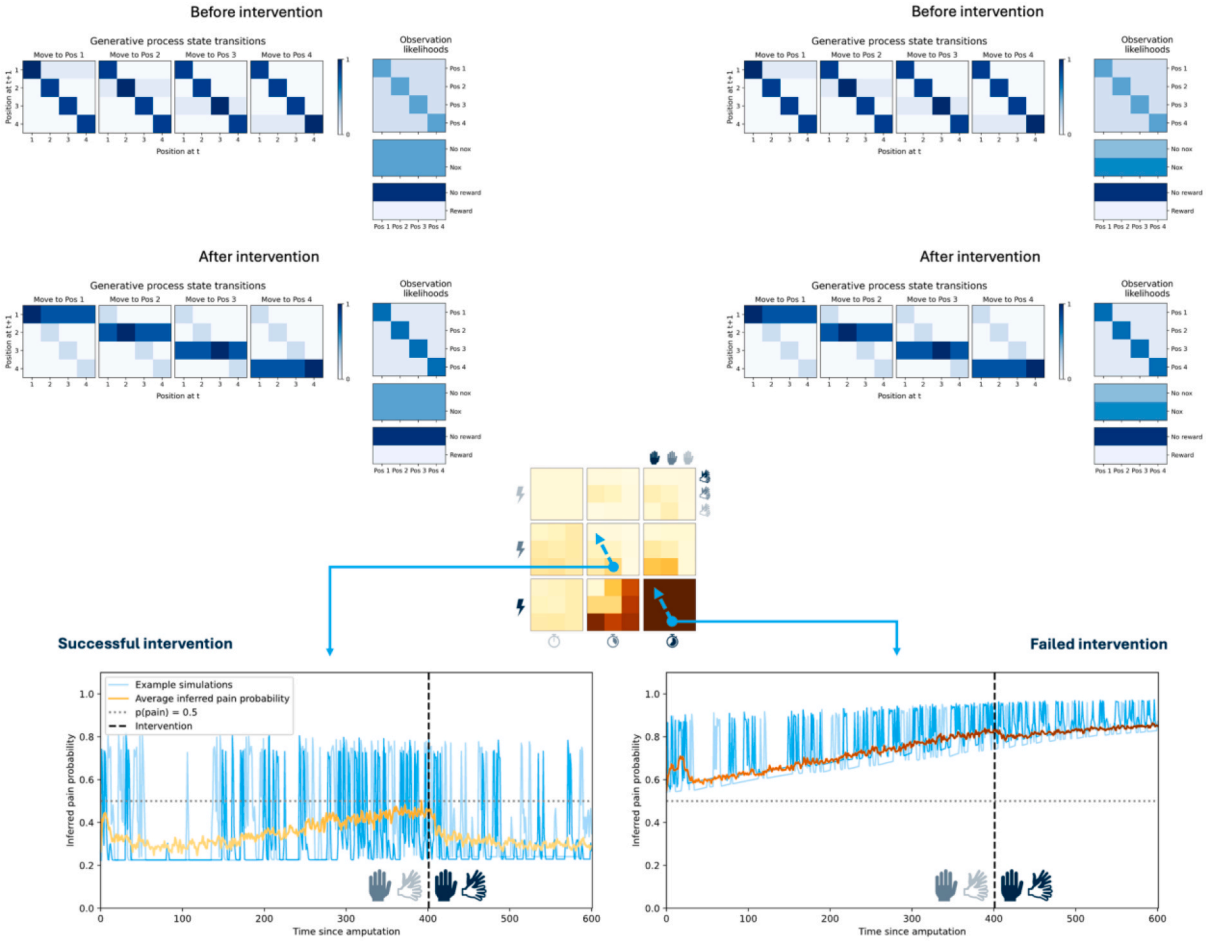

Fig. B8. ***Details on the generative process before and after the intervention shown in*** Fig. 4 *of the main manuscript. In these interventions, the agent is given a means of controlling some source other than the missing limb that generates sensory input that matches the intended movements, such as an active prosthesis or AR/VR limb.*

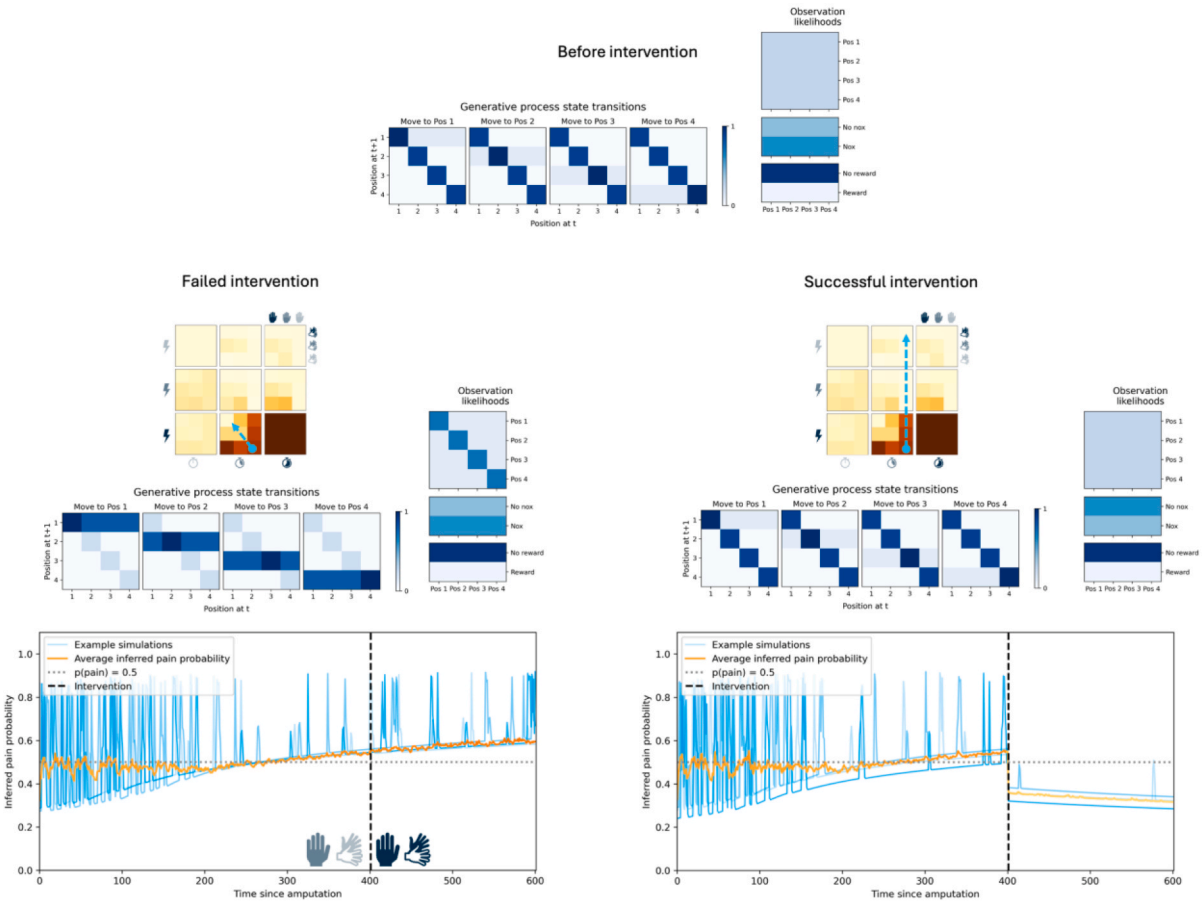

Fig. B9. ***Left: Example where the colors of the heatmat do not predict intervention success.*** *The time with phantom limb pain may partially have the same effect as pre amputation pain, thus reducing the effectiveness of interventions that look to be effective based on the colors of the heatmap.* ***Right:*** *Addressing residual sources of noxious input does relieve the pain.* ***These results further strengthen the finding that sources of residual noxious input must be addressed as the first line of action.***

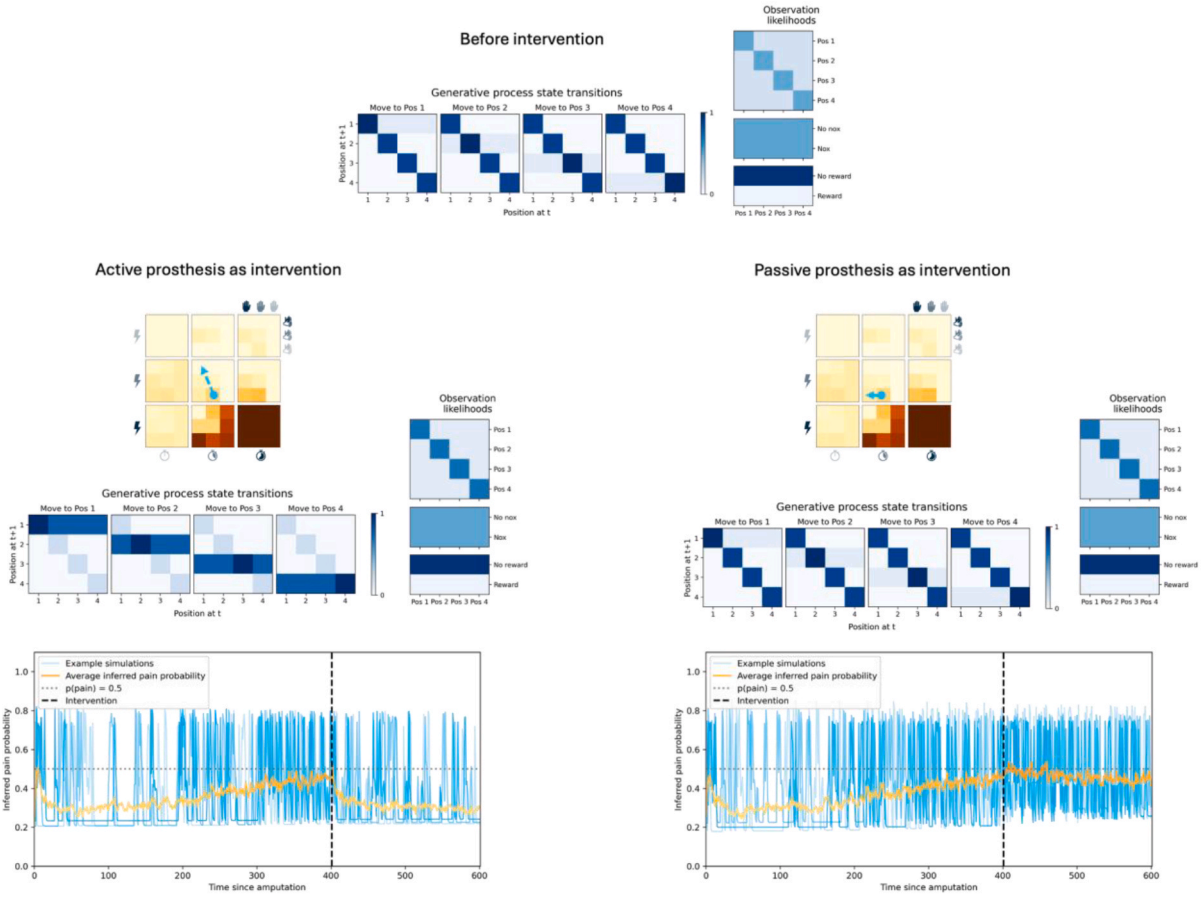

Intervention	Time with pain after intervention	Avg inferred probability of pain after intervention
Active prosthesis	10 %	.31
Passive prosthesis	29 %	.46

Fig. B10. ***Example of passive* vs *active prosethesis as intervention***, *assuming that a passive prosthesis corresponds (partially) restoring position observations but not controllability (horizontal leftward movement within heatmaps). In this example, patients are helped by the active but not the passive prosthesis.*

Examples of asymmetrical likelihood and control matrices.

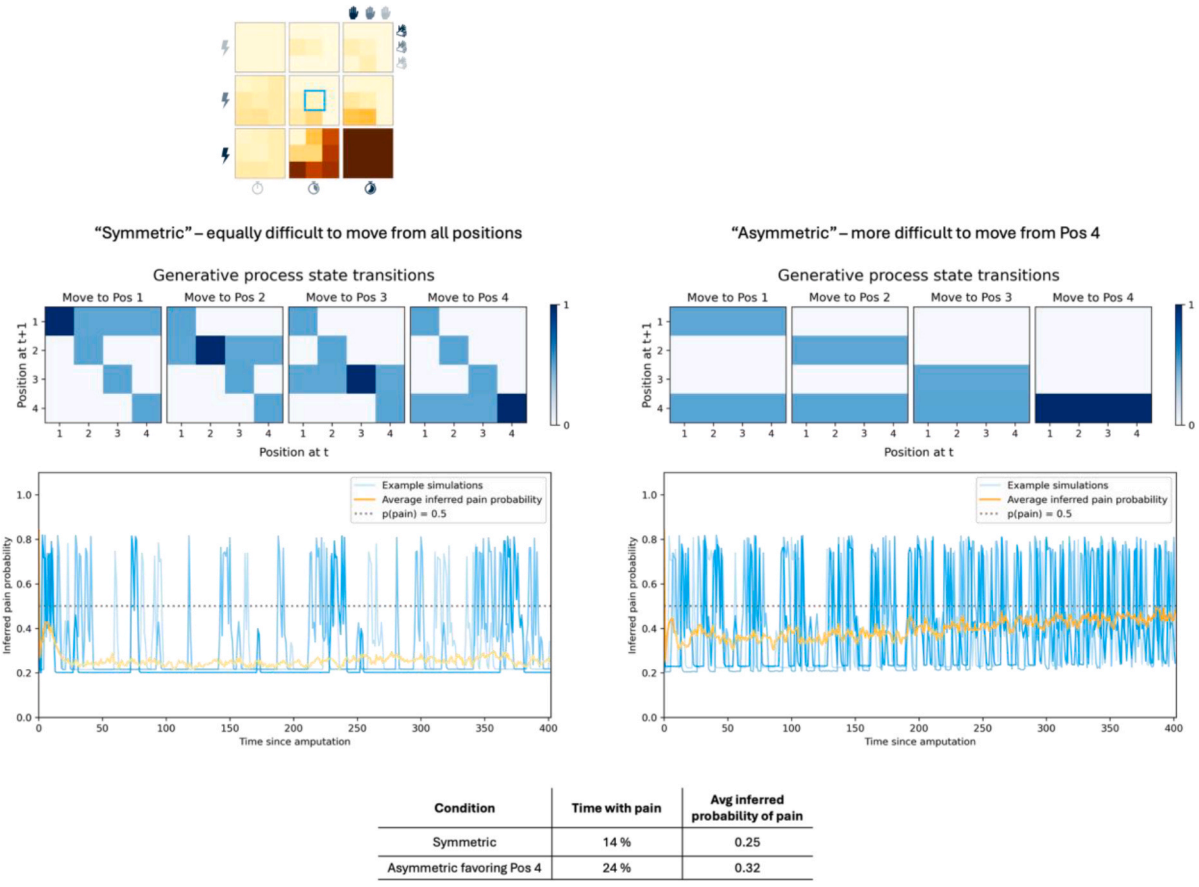

Condition	Time with pain	Avg inferred probability of pain
Symmetric	14 %	0.25
Asymmetric favoring Pos 4	24 %	0.32

Fig. B11. ***Example where asymmetric transition matrices can exacerbate phantom limb pain.*** *Severed nerves might reinnervate in such a way that some movements might be more difficult than others following amputation. In this example, it is more difficult to move from Pos 4, resulting in increased presence of phantom limb pain.*

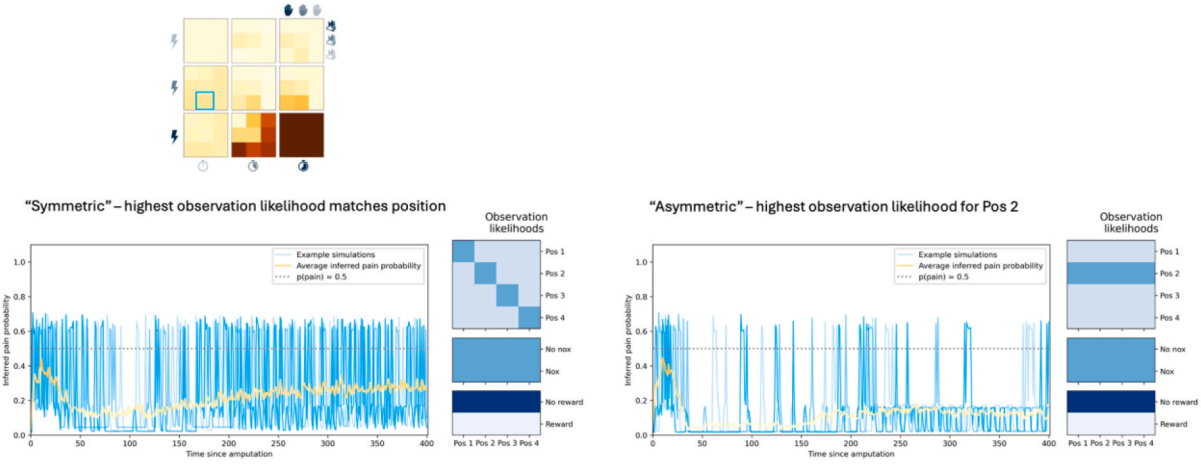

Condition	Time with pain	Avg inferred probability of pain
Symmetric	18%	0.18
Asymmetric favoring Pos 2	9 %	0.11

Fig. B12. ***Example where asymmetric likelihood matrix can reduce the prevalence of phantom limb pain.*** *Severed nerves might reinnervate in such a way that they are more likely to signal one position than another. In this example observation likelihood is highest for Pos 2, resulting in reduced presence of phantom limb pain.*

References

Albright-Trainer, B., Phan, T., Trainer, R.J., Crosby, N.D., Murphy, D.P., Disalvo, P., et al., 2021. Peripheral nerve stimulation for the management of acute and subacute post-amputation pain: a randomized, controlled feasibility trial. Pain Manag 12, 357–369. https://doi.org/10.2217/pmt-2021-0087.

Ashar, Y.K., Gordon, A., Schubiner, H., Uipi, C., Knight, K., Anderson, Z., et al., 2022. Effect of Pain Reprocessing Therapy vs Placebo and Usual Care for Patients with Chronic Back Pain: a Randomized Clinical Trial. JAMA Psychiat. 79, 13–23. https://doi.org/10.1001/jamapsychiatry.2021.2669.

Atlas, L.Y., Wager, T.D., 2012. How expectations shape pain. Neurosci. Lett. 520. https://doi.org/10.1016/j.neulet.2012.03.039.

Attal, N., Bouhassira, D., Baron, R., 2018. Diagnosis and assessment of neuropathic pain through questionnaires. Lancet Neurol. 17, 456–466.

Barbin, J., Seetha, V., Casillas, J.M., Paysant, J., Pérennou, D., 2016. The effects of mirror therapy on pain and motor control of phantom limb in amputees: a systematic review. Ann. Phys. Rehabil. Med. 59. https://doi.org/10.1016/j.rehab.2016.04.001.

Birbaumer N, Lutzenberger W, Montoya P, Larbig W, Unertl K, Töpfner S, et al. Effects of Regional Anesthesia on Phantom Limb Pain Are Mirrored in Changes in Cortical Reorganization. The Journal of Neuroscience 1997;17:5503 LP – 5508. https://doi.org/10.1523/JNEUROSCI.17-14-05503.1997.

Borghi, B., D'Addabbo, M., White, P.F., Gallerani, P., Toccaceli, L., Raffaeli, W., et al., 2010. The use of Prolonged Peripheral Neural Blockade after lower Extremity Amputation: the effect on Symptoms Associated with Phantom Limb Syndrome. Anesth. Analg. 111.

Boström, K.J., De Lussanet, M.H.E., Weiss, T., Puta, C., Wagner, H., 2014. A computational model unifies apparently contradictory findings concerning phantom pain. Sci. Rep. 4. https://doi.org/10.1038/srep05298.

Büchel, C., 2023. The role of expectations, control and reward in the development of pain persistence based on a unified model. Elife 12, e81795. https://doi.org/10.7554/eLife.81795.

Burrack, A., Brugger, P., 2005. Individual differences in susceptibility to experimentally induced phantom sensations. Body Image 2. https://doi.org/10.1016/j.bodyim.2005.04.002.

Collins, K.L., Robinson-Freeman, K.E., O'Conor, E., Russell, H.G., Tsao, J.W., 2018. A survey of frozen phantom limb experiences: are experiences compatible with current theories. Front. Neurol. 9. https://doi.org/10.3389/fneur.2018.00599.

Drusko, A., Baumeister, D., McPhee Christensen, M., Kold, S., Fisher, V.L., Treede, R.-D., et al., 2023. A novel computational approach to pain perception modelling within a Bayesian framework using quantitative sensory testing. Sci. Rep. 13, 3196. https://doi.org/10.1038/s41598-023-29758-8.

Dumanian, G.A., Potter, B.K., Mioton, L.M., Ko, J.H., Cheesborough, J.E., Souza, J.M., et al., 2019. Targeted Muscle Reinnervation Treats Neuroma and Phantom Pain in Major Limb Amputees: a Randomized Clinical Trial. Ann. Surg. 270. https://doi.org/10.1097/SLA.0000000000003088.

Eckert, A.-L., Pabst, K., Endres, D.M., 2022. A Bayesian model for chronic pain. Front. Pain Res. 3. https://doi.org/10.3389/fpain.2022.966034.

Fields, H.L., 2018. How expectations influence pain. Pain 159. https://doi.org/10.1097/j.pain.0000000000001272.

Flor, H., 2002. Phantom-limb pain: Characteristics, causes, and treatment. Lancet Neurol. 1. https://doi.org/10.1016/S1474-4422(02)00074-1.

Flor, H., Elbert, T., Knecht, S., Wienbruch, C., Pantev, C., Birbaumers, N., et al., 1995. Phantom-limb pain as a perceptual correlate of cortical reorganization following arm amputation. Nature 375, 482–484. https://doi.org/10.1038/375482a0.

Friston, K.J., Trujillo-Barreto, N., Daunizeau, J., 2008. DEM: a variational treatment of dynamic systems. Neuroimage 41. https://doi.org/10.1016/j.neuroimage.2008.02.054.

Gardetto, A., Baur, E.M., Prahm, C., Smekal, V., Jeschke, J., Peternell, G., et al., 2021. Reduction of phantom limb pain and improved proprioception through a tsr-based surgical technique: a case series of four patients with lower limb amputation. J. Clin. Med. 10. https://doi.org/10.3390/jcm10174029.

Grahl, A., Onat, S., Büchel, C., 2018. The periaqueductal gray and Bayesian integration in placebo analgesia. Elife 7, e32930. https://doi.org/10.7554/eLife.32930.

Grant PE, Caneris OA, Gonzalez RG, Iadarola MJ, Sapio MR, Mannes AJ, et al. Analgesia after dorsal root ganglionic injection under CT-guidance in a patient with intractable phantom limb pain. Pain Medicine 2023:pnad039. https://doi.org/10.1093/pm/pnad039.

Habermann, M., Strube, A., Büchel, C., 2025. How control modulates pain. Trends Cogn. Sci. 29, 60–72. https://doi.org/10.1016/j.tics.2024.09.014.

Harris, A.J., 1999. Cortical origin of pathological pain. Lancet 354, 1464–1466. https://doi.org/10.1016/S0140-6736(99)05003-5.

Heins, C., Millidge, B., Demekas, D., Klein, B., Friston, K., Couzin, I.D., et al., 2022. pymdp: a Python library for active inference in discrete state spaces. J Open Source Softw 7. https://doi.org/10.21105/joss.04098.

Hesp, C., Smith, R., Parr, T., Allen, M., Friston, K.J., Ramstead, M.J.D., 2021. Deeply felt affect: the emergence of valence in deep active inference. Neural Comput. 33. https://doi.org/10.1162/neco_a_01341.

Hill, A., 1999. Phantom limb pain: a review of the literature on attributes and potential mechanisms. J. Pain Symptom Manage. 17. https://doi.org/10.1016/S0885-3924(98)00136-5.

Hoskin, R., Berzuini, C., Acosta-Kane, D., El-Deredy, W., Guo, H., Talmi, D., 2019. Sensitivity to pain expectations: a Bayesian model of individual differences. Cognition 182, 127–139.

Jepma, M., Koban, L., van Doorn, J., Jones, M., Wager, T.D., 2018. Behavioural and neural evidence for self-reinforcing expectancy effects on pain. Nat. Hum. Behav. 2, 838–855. https://doi.org/10.1038/s41562-018-0455-8.

Legaspi, R., Toyoizumi, T., 2019. A Bayesian psychophysics model of sense of agency. Nat. Commun. 10. https://doi.org/10.1038/s41467-019-12170-0.

Lendaro, E., Van der Sluis, C.K., Hermansson, L., Bunketorp-Käll, L., Burger, H., Keesom, E., et al., 2025. Extended reality used in the treatment of phantom limb pain: a multicenter, double-blind, randomized controlled trial. Pain 166, 571–586. https://doi.org/10.1097/j.pain.0000000000003384.

Liakoni, V., Modirshanechi, A., Gerstner, W., Brea, J., 2021. Learning in volatile environments with the bayes factor surprise. Neural Comput. 33. https://doi.org/10.1162/neco_a_01352.

Limakatso, K., Bedwell, G.J., Madden, V.J., Parker, R., 2020. The prevalence and risk factors for phantom limb pain in people with amputations: a systematic review and meta-analysis. PLoS One 15, e0240431. https://doi.org/10.1371/journal.pone.0240431.

Limakatso, K., Cashin, A.G., Williams, S., Devonshire, J., Parker, R., McAuley, J.H., 2023. The efficacy of graded motor imagery and its components on phantom limb pain and disability: a systematic review and meta-analysis. Canadian Journal of Pain, 2188899. https://doi.org/10.1080/24740527.2023.2188899.

Limakatso, K., McGowan, E., Ortiz-Catalan, M., 2025. Evaluating Mirror Therapy Protocols in Phantom Limb Pain Clinical Trials: a Scoping Review. J. Pain Res. 18, 619–629. https://doi.org/10.2147/JPR.S502541.

Limakatso, K., Parker, R., 2021. Treatment Recommendations for Phantom Limb Pain in people with Amputations: an Expert Consensus Delphi Study. PM and R 13, 1216–1226. https://doi.org/10.1002/pmrj.12556.

Makin, T.R., Scholz, J., Filippini, N., Henderson Slater, D., Tracey, I., Johansen-Berg, H., 2013. Phantom pain is associated with preserved structure and function in the former hand area. Nat. Commun. 4, 1570. https://doi.org/10.1038/ncomms2571.

Makin, T.R., Scholz, J., Henderson Slater, D., Johansen-Berg, H., Tracey, I., 2015. Reassessing cortical reorganization in the primary sensorimotor cortex following arm amputation. Brain 138, 2140–2146. https://doi.org/10.1093/brain/awv161.

Mancini, F., Zhang, S., Seymour, B., 2022. Computational and neural mechanisms of statistical pain learning. Nat. Commun. 13, 6613. https://doi.org/10.1038/s41467-022-34283-9.

Nyström, B., Hagbarth, K.E., 1981. Microelectrode recordings from transected nerves in amputees with phantom limb pain. Neurosci. Lett. 27. https://doi.org/10.1016/0304-3940(81)90270-6.

Ortiz-Catalan, M., Guðmundsdóttir, R., Kristoffersen, M., Zepeda-Echavarria, A., Caine-Winterberger, K., Kulbacka-Ortiz, K., et al., 2016. Phantom motor execution facilitated by machine learning and augmented reality as treatment for phantom limb pain: a single group, clinical trial in patients with chronic intractable phantom limb pain. Lancet 6736. https://doi.org/10.1016/S0140-6736(16)31598-7.

Pagan-Rosado, R., Smith, B.J., Smither, F.C., Pingree, M.J., D'Souza, R.S., 2023. Peripheral Nerve Stimulation for the Treatment of Phantom Limb Pain. A Case Series. Case Rep Anesthesiol. https://doi.org/10.1155/2023/1558183.

Parr T, Pezzulo G, Friston KJ. Active Inference: The Free Energy Principle in Mind, Brain, and Behavior. MIT Press; 2022. https://doi.org/10.7551/mitpress/12441.001.0001.

Pettersen, E., Sassu, P., Reinholdt, C., Dahm, P., Rolfson, O., Björkman, A., et al., 2023. Surgical treatments for postamputation pain: study protocol for an international, double-blind, randomised controlled trial. Trials. https://doi.org/10.1186/s13063-023-07286-0.

Raja, S.N., Carr, D.B., Cohen, M., Finnerup, N.B., Flor, H., Gibson, S., et al., 2020. The revised International Association for the Study of Pain definition of pain: concepts, challenges, and compromises. Pain 161.

Ramne, M., 2025. A mathematical model of dorsal horn circuits' contribution to neuropathic pain. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. 2025, 1–7. https://doi.org/10.1109/EMBC58623.2025.11253737.

Ramne, M., 2026. Advancing the Understanding of Phantom Limb Pain through Mathematical Models. Chalmers University of Technology. https://doi.org/10.63959/chalmers.dt/5817.

Ramne, M., Sensinger, J., 2024. A Computational Framework for Understanding the Impact of prior Experiences on Pain perception and Neuropathic Pain. PLoS Comput. Biol. 20. https://doi.org/10.1371/journal.pcbi.1012097.

Ropero Peláez, F.J., Taniguchi, S., 2016. The Gate Theory of Pain Revisited: Modeling Different Pain Conditions with a Parsimonious Neurocomputational Model. Neural Plast. 2016, 4131395. https://doi.org/10.1155/2016/4131395.

Schone HR, Baker CI, Katz J, Nikolajsen L, Limakatso K, Flor H, et al. Making sense of phantom limb pain. J Neurol Neurosurg Psychiatry 2022;93:833 LP – 843. https://doi.org/10.1136/jnnp-2021-328428.

Seymour, B., 2019. Pain: a Precision Signal for Reinforcement Learning and Control. Neuron 101, 1029–1041. https://doi.org/10.1016/j.neuron.2019.01.055.

Seymour, B., Mancini, F., 2020. Hierarchical models of pain: Inference, information-seeking, and adaptive control. Neuroimage 222, 117212. https://doi.org/10.1016/j.neuroimage.2020.117212.

Smith, R., Friston, K.J., Whyte, C.J., 2022. A step-by-step tutorial on active inference and its application to empirical data. J. Math. Psychol. 107. https://doi.org/10.1016/j.jmp.2021.102632.

Srinivasan, S.S., Gutierrez-Arango, S., Teng, A.C.E., Israel, E., Song, H., Bailey, Z.K., et al., 2021. Neural interfacing architecture enables enhanced motor control and residual limb functionality postamputation. PNAS 118. https://doi.org/10.1073/pnas.2019555118.

Steven Cohen, C.P., Army, U., Gilmore, C.A., Rauck, R.L., Lester, D.D., Trainer, R.J., et al., 2019. Percutaneous Peripheral Nerve Stimulation for the Treatment of Chronic Pain following Amputation. Mil. Med. 184, 267. https://doi.org/10.1093/milmed/usz114.

Stone, K.D., Bullock, F., Keizer, A., Dijkerman, H.C., 2018. The disappearing limb trick and the role of sensory suggestibility in illusion experience. Neuropsychologia 117. https://doi.org/10.1016/j.neuropsychologia.2018.07.012.

Strube, A., Horing, B., Rose, M., Büchel, C., 2023. Agency affects pain inference through prior shift as opposed to likelihood precision modulation in a Bayesian pain model. Neuron 111. https://doi.org/10.1016/j.neuron.2023.01.002.

Sugawara, A.T., Simis, M., Fregni, F., Battistella, L.R., 2021. Characterisation of Phantom Limb Pain in Traumatic Lower-Limb Amputees. Pain Res. Manag. https://doi.org/10.1155/2021/2706731.

Tabor, A., Burr, C., 2019. Bayesian Learning Models of Pain: a call to Action. Curr. Opin. Behav. Sci. 26, 54–61. https://doi.org/10.1016/j.cobeha.2018.10.006.

Tabor, A., Thacker, M.A., Moseley, G.L., Körding, K.P., 2017. Pain: a Statistical Account. PLoS Comput. Biol. 13, e1005142. https://doi.org/10.1371/journal.pcbi.1005142.
Vaso, A., Adahan, H.M., Gjika, A., Zahaj, S., Zhurda, T., Vyshka, G., et al., 2014. Peripheral nervous system origin of phantom limb pain. Pain 155, 1384–1391. https://doi.org/10.1016/j.pain.2014.04.018.
Villemure, C., Bushnell, C.M., 2002. Cognitive modulation of pain: how do attention and emotion influence pain processing? Pain 95. https://doi.org/10.1016/S0304-3959(02)00007-6.
Weiss, T., Koehler, H., Croy, I., 2023. Pain and Reorganization after Amputation: is Interoceptive Prediction a Key? Neuroscientist 29. https://doi.org/10.1177/10738584221112591.
Wiech, K., Tracey, I., 2013. Pain, decisions, and actions: a motivational perspective. Front. Neurosci. https://doi.org/10.3389/fnins.2013.00046.
Woo, S.L., Kung, T.A., Brown, D.L., Leonard, J.A., Kelly, B.M., Cederna, P.S., 2016. Regenerative peripheral nerve interfaces for the treatment of postamputation neuroma pain: a pilot study. Plast. Reconstr. Surg. Glob. Open 4. https://doi.org/10.1097/GOX.0000000000001038.