Scientific Update and Overview

Check for updates

Symptom Perception From a Predictive Processing Perspective

Giovanni Pezzulo[a], Domenico Maisto[b], Laura Barca[a], Omer Van den Bergh[c]

[a] *Institute of Cognitive Sciences and Technologies, National Research Council, Rome, Italy.* [b] *Institute for High Performance Computing and Networking, National Research Council, Naples, Italy.* [c] *Health Psychology, University of Leuven, Leuven, Belgium.*

Corresponding Author: Giovanni Pezzulo, Institute of Cognitive Sciences and Technologies, National Research Council, Via San Martino della Battaglia 44, 00185 Rome, Italy. E-mail: giovanni.pezzulo@istc.cnr.it

Abstract

Background: Bodily symptoms are highly prevalent in psychopathology, and in some specific disorders, such as somatic symptom disorder, they are a central feature. In general, the mechanisms underlying these symptoms are poorly understood. However, also in well-known physical diseases there seems to be a variable relationship between physiological dysfunction and self-reported symptoms challenging traditional assumptions of a biomedical disease model.
Method: Recently, a new, predictive processing conceptualization of how the brain works has been used to understand this variable relationship. According to this predictive processing view, the experience of a symptom results from an integration of both interoceptive sensations as well as from predictions about these sensations from the brain.
Results: In the present paper, we introduce the predictive processing perspective on perception (predictive coding) and action (active inference), and apply it to asthma in order to understand when and why asthma symptoms are sometimes strongly, moderately or weakly related to physiological disease parameters.
Conclusion: Our predictive processing view of symptom perception contributes to understanding under which conditions misperceptions and maladaptive action selection may arise.

Keywords

Highlights

- There is a variable relationship between physiological dysfunction and self-reported symptoms.
- We conceptualize symptom perception (and misperception) within a predictive processing perspective.
- In this view, symptom perception integrates sensations and predictions about these sensations.
- Failures of such integration can produce misperceptions and maladaptive action selection.
- We use the perception (and misperception) of asthma symptoms as an example.

New developments in the conceptualization of how the brain works have recently emerged. These conceptualizations emphasize the predictive nature of the brain, hence are known as predictive coding or predictive processing views (Clark, 2013; Friston, 2010; Hohwy, 2013). Although the basic ideas underlying this conceptualization have been developed by von Helmholtz in the late 19th century, a strong impetus in recent years has been given by the thorough study of perception, especially of visual illusions. Many perceptual phenomena can only be understood by assuming that meaningful perception is not just a matter of processing incoming information, but that it is also largely reliant on pre-existing (prior) information: often the brain unconsciously and compellingly assumes (or infers) non-given information to construct a meaningful percept.

Predictive processing views and their implications are currently explored in an increasing number of scientific areas. In neuroscience, the theory of "predictive coding" (Friston, 2005; Rao & Ballard, 1999) describes how sensory (e.g., visual) hierarchies in the brain may combine prior knowledge and sensory evidence, by continuously exchanging top-down (predictions) and bottom-up (prediction error) signals. Besides, interest in creating intelligent systems enhanced the need to extend the predictive processing perspective beyond perceptual processing, to address also action and planning (aka active inference). Pioneering work towards this goal has been done by Karl Friston and colleagues (Friston et al., 2016; Friston, FitzGerald, Rigoli, Schwartenbeck, & Pezzulo, 2017; Friston et al., 2015; Friston, Samothrakis, & Montague, 2012; Pezzulo, Rigoli, & Friston, 2018). In the present paper, we will first introduce some basic concepts of the predictive processing view of perception (called "predictive coding") and its extension to the action domain (called "active inference"). Next, we will briefly describe their implications for symptom perception. The remainder of this paper will sketch a formal model of symptom perception as viewed from a predictive processing perspective.

PsychOpen GOLD

Predictive Processing During Perception (Predictive Coding) and Action (Active Inference)

A basic task of the brain is to construct an adaptive model of the (external and internal) world, while its only source of information to do so is the spatial and temporal patterning of its own neural activity. In order to achieve this goal, the brain uses information from neural activity that is triggered by peripheral input (sense organs and receptors in the peripheral body), but also from neural activity that is generated by the brain itself (aka spontaneous dynamics), reflecting previous experiences and "built in" information. This leads to two counterflowing streams of neural activation across several hierarchical levels of the brain: stimulation by peripheral input (called "likelihood" in the context of Bayesian inference) interacts with activations generated by the brain that act as model-based predictions of the input ("priors") within a specific context. For example, if one is waiting for Jeff in a crowded street, the brain generates neural patterns acting as priors that will facilitate spotting Jeff in the crowd.

The theory of "*predictive coding*" specifies how the brain may mechanistically implement this kind of Bayesian inference. According to predictive coding, input at each hierarchical level that is predicted is cancelled out ("explained away"), while unpredicted input creates prediction errors that are relayed to the next hierarchical level where it meets priors generating new prediction errors. Prediction errors are thus propagated through the brain from very basic and concrete to higher abstract levels of representation to eventually settle on a "posterior" belief (to be understood in the technical sense of a neural probability distribution, not as a conscious belief) that accounts for the stimulation with the least overall prediction errors. The posterior belief can subsequently act as prior for new input leading to further adaption in an iterative process. In the case of waiting for Jeff: the benefit of having an a priori belief in the brain representing Jeff is that it helps to quickly recognize him and to prime a network of related information for further interaction. Obviously, there is also a downside of having highly active priors about Jeff arriving soon: whenever input is downgraded to some extent, any person that resembles Jeff will easily be mistaken for Jeff. In sum, the theory of "predictive coding" postulates that the brain continuously strives to minimize its prediction error (and the difference between predictions and sensations). It does so by accommodating the prior hypothesis (or belief) and/or the model producing such hypothesis, to fit unpredicted information. For example, if Jeff was expected but a female appears, the brain can revise the prior belief. Furthermore, if Jeff is wearing a fancy new cap and sunglasses - which is discrepant information compared to previous encounters - the model of Jeff in the brain may be adapted (for example, by reducing the weight given to these aspects of visual input).

The theory of "active inference" extends this view to also account for active components of perceptual processing (active perception) and goal-directed behavior. In this per-

spective, the brain does not passively wait for sensory stimulations, but it can initiate activity to produce input that is consistent with its adaptive model. Waiting for Jeff may prompt the person to move towards a location providing a better overview of the passing crowd and/or to increase the scanning rate generating more detailed information to help spotting him; or even going to Jeff's house, if he does not appear. As these examples of active inference illustrate, acting is just another way to reduce prediction error. In other words, while in predictive coding one reduces prediction errors by changing the prior belief to fit the world, in active inference one reduces prediction errors by changing the world to fit the prior belief (e.g., that one will encounter Jeff). As this latter example illustrates, in active inference the prior belief is much more than a prediction: it can play the role of a cognitive goal that triggers a goal-directed plan (e.g., a plan to go to Jeff's house).

The Importance of Precision Control

Priors and prediction errors (PE's) can be thought of as probability distributions of neural activity capturing statistical regularities associated with a specific input. These distributions are characterized by a variance, or its inverse: precision. Highly precise priors and prediction errors reflect that a neural pattern has a high probability of being associated with a particular input, and conversely for low precise priors and PE's. If Jeff is unusually tall, both priors and PE's representing Jeff's height are highly precise, resulting in a quick and reliable recognition of Jeff. Repeated encounters will also generate precision expectations, that is: not only is the perceptual information related to Jeff's "height" highly precise, the brain will learn to consider "height" as a highly precise prior for recognizing Jeff.

Precision parameters of both PE's and priors are used as weighting factors in Bayesian inference and predictive coding: they determine the relative contributions of prior information and sensory evidence to the brain's "posterior belief" - and thus the content of perception. Highly precise priors and low precise PE's will shift the posterior belief towards the prior, while the reverse is true with low precise priors and highly precise PE's (see Figure 1 for a graphical illustration of integration of prior and sensory evidence in Bayesian inference). For example, when it is dark, there is a high probability to recognize Jeff in any tall person, reflecting a strong effect of the prior on the eventual perception. Conversely, on a sunny day it is less likely to take any tall person for Jeff and this likelihood is even further reduced if one is not waiting for Jeff.
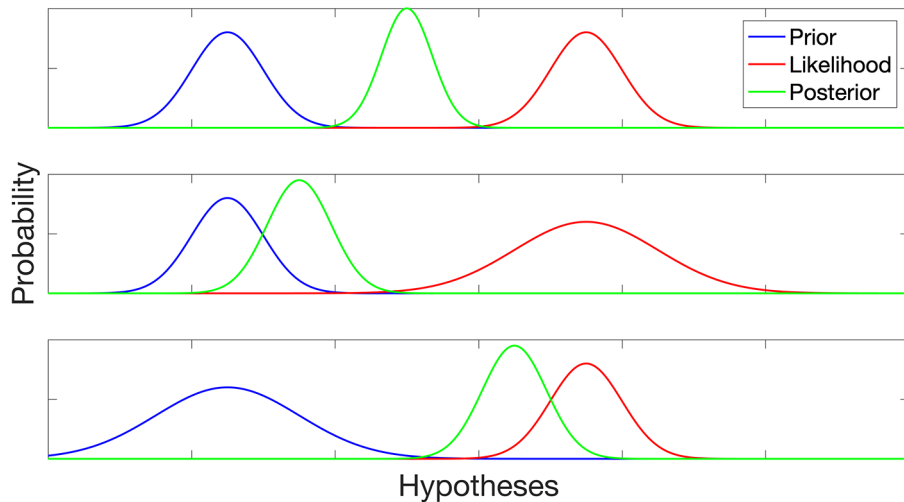
Figure 1. Integration of prior knowledge and sensory evidence (likelihood) in Bayesian inference.

Note. The top panel shows that if prior and likelihood have the same precision (i.e., inverse variance of the Gaussian distribution), the posterior belief is in between. The second and third panels show that higher precision prior and likelihood "attract" the posterior, respectively. Note that in all cases, the precision of the posterior increases compared to the prior. See the main text for explanation.

Precision parameters of sensory events play an additional role in (active) perceptual inference and information gathering. Information sources that are assumed to bring more precise information are preferentially sampled, while those that are assumed to bring imprecise information can be ignored (e.g., looking for Jeff in the total dark is useless and thus avoided).

In sum, perception can be considered a dynamic constructive process balancing external input and pre-existing information: under some conditions, the eventual percept closely reflects the external input, while in other conditions it may more closely reflect pre-existing information that act as (implicit) prior expectations. Perceptual illusions can be considered extreme cases where the percept is (almost) entirely determined by prior expectations (Pezzulo, 2014; Sterzer et al., 2018). Furthermore, perception has active (information gathering) components that permit sampling information from the most precise information sources - but can lead to inattention or even neglect when precision parameters are not set correctly (Parr & Friston, 2018).

Predictive Processing and Symptom Perception

One of the research areas for which these new conceptualizations are particularly fruitful is interoception, which is considered to play an important role in the experience of the self, agency, emotion and psychopathology (Allen, Levy, Parr, & Friston, 2019; Barca & Pezzulo, 2019; Iodice, Porciello, Bufalari, Barca, & Pezzulo, 2019; Pezzulo, Barca, & Friston, 2015; Pezzulo, Maisto, Barca, & Van den Bergh, 2019; Pezzulo, Rigoli, & Friston, 2015; Seth, 2013; Tsakiris & Preester, 2018). The Embodied Predictive Interoception Coding model (EPIC; Barrett & Simmons, 2015) describes the neural architecture and functional characteristics of interoception, suggesting a critical role for active inference: visceromotor cortices generate autonomic, hormonal and immunological predictions to adequately deal with anticipated demands while PE's are fed back to the brain to adapt and modify subsequent predictions. Because visceromotor cortices are overall relatively insensitive to somatic input, interoception is largely dominated by prior expectations ("a construction of beliefs that are kept in check by the actual state of the body", Barrett & Simmons, 2015, p. 424). Being critical for symptom perception, this account of interoception allows and suggests important variability in the relationship between symptoms and peripheral bodily dysfunction. This has tremendous conceptual and practical implications for medicine.

Indeed, while the relationship between self-reported symptoms and parameters of peripheral bodily dysfunction is generally strong in acute monosymptomatic health conditions, it becomes typically much weaker in chronic multisymptomatic conditions (Janssens, Verleden, De Peuter, Van Diest, & Van den Bergh, 2009). In a substantial number of cases no relationship with physiological dysfunction can be found at all. Hence, the latter are often called "medically unexplained symptoms" (MUS). The prevalence of MUS in primary care consultations is estimated around one third, while prevalence rates in secondary care are even higher (De Waal, Arnold, Eekhof, & van Hemert, 2004; Nimnuan, Hotopf, & Wessely, 2001). In secondary care general medicine, the symptoms often appear as functional syndromes, such as chronic fatigue, fibromyalgia, irritable bowel disease, multiple chemical sensitivity, bodily distress disorder, while in psychiatry they are labeled as somatic symptom disorder, somatization disorder, conversion disorder, etc. However, also placebo and nocebo phenomena which are abundantly present in everyday medicine are difficult to understand within a strict biomedical disease model.

The predictive processing perspective allows to describe the conditions moderating the relationship between symptoms and bodily dysfunction (Van den Bergh, Witthöft, Petersen, & Brown, 2017), to explain pseudoneurological symptoms and conversion (Edwards, Adams, Brown, Pareés, & Friston, 2012), persistent physical symptoms (Henningsen et al., 2018), placebo effects (Büchel, Geuter, Sprenger, & Eippert, 2014) and pain perception (Wiech, 2016). However, most current models appeal to the mechanisms of predictive coding, while disregarding action components (or active inference) that are equally important to understand symptoms and psychopathological conditions.

PsychOpen GOLD

Below, we discuss a worked example of symptom perception in terms of underlying predictive coding and active inference dynamics. Our example focuses on asthma perception. Asthma relies on a well-known physiological dysfunction but often the symptoms do not clearly relate to that dysfunction, which is a rather prevalent clinical problem (De Peuter et al., 2005; Janssens et al., 2009). Our example describes the conditions for a strong, weak or absent relationship between symptoms and bodily input.

A Worked Example of Symptoms and the Body: The Case of Asthma Perception

Consider the simplified case of an asthmatic person who feels two bodily sensations (e.g., wheezing, breathlessness) that sometimes indicate the beginning of an asthma episode. The person has to infer whether it is an asthma episode (Hypothesis 1) or not (Hypothesis 2), based on what he currently feels (e.g., wheezing, breathlessness) and his prior belief (e.g., the fact that he/she is in the bedroom where he usually has asthma episodes).

Generative Model and Inference

From the formal perspective of predictive coding (and more broadly, Bayesian inference), the brain makes this inference using a so-called "generative model" of how its sensations are generated. The "generative model" has two essential components. The first one ("likelihood model") describes the probabilistic mapping between sensations (e.g., wheezing, breathlessness) and the two competing hypotheses (Hypothesis 1: this is an asthma episode; Hypothesis 2: this is not) - which in this context are also called "hidden" states, because they cannot be directly observed but need to be inferred. For example, a good likelihood model of asthma may represent the fact that under Hypothesis 1 (this is an asthma episode), the probability of feeling wheezing is high (e.g., 0.8). However, under Hypothesis 2 (this is not an asthma episode), the probability of feeling wheezing is very low (e.g., 0.05). In other words, the person should expect to feel wheezing (only) if he is experiencing an asthma episode. Furthermore, the likelihood model may represent the fact that breathlessness has the same probability (e.g., 0.6) under Hypotheses 1 and 2 (and more broadly, that one can feel breathless for many other reasons, such as because one has done physical exercise). A consequence of having this particular likelihood model is that while wheezing is very informative (i.e., feeling wheezing tells me with high probability that Hypothesis 1 is true; and not feeling wheezing tells me with high probability that Hypothesis 2 is true), breathlessness is not, as it cannot disambiguate between Hypotheses 1 and 2.

The second component of the generative model is the person's "prior belief" about the two Hypotheses 1 and 2. For example, if the asthmatic person is in the bedroom where he frequently experienced asthma episodes in the past, he may have a high prior

belief (e.g., 0.7) for Hypothesis 1. If we assume for simplicity that Hypotheses 1 and 2 are mutually exclusive, and there are no alternative hypotheses, then the prior probability of Hypothesis 2 is just one minus the prior probability of Hypothesis 1; that it, 0.3.

We can use these figures to calculate the (posterior) probability of the two (mutually exclusive) Hypotheses 1 and 2, according to Bayes' rule:

$$\text{Posterior of HYP1} = \frac{\text{prior of HYP1 * likelihood of HYP1}}{\text{prior of HYP1 * likelihood of HYP1 + prior of HYP2 * likelihood of HYP2}}$$

$$\text{Posterior of HYP2} = 1 - \text{Posterior of HYP1}$$

Imagine the person is currently experiencing wheezing and is in the bedroom where he frequently experiences asthma episodes. We can use the numbers above to calculate the posterior probability (or belief) about Hypotheses 1 and 2, as follows:

$$\text{Posterior of HYP1} = \frac{0.7 * 0.8}{0.7 * 0.8 + 0.3 * 0.05} = 0.9739$$

Therefore, in this example, the posterior probability of HYP1 is 0.9739 and the posterior probability of HYP2 is one minus 0.9739, that is, 0.026. This means that in this situation, the person would have a very strong belief (in probabilistic terms) about an asthma episode.

It is possible to use the same formula to simulate other possible situations. Imagine that the same person is in the same room but does not feel any wheezing or breathlessness. In this second example, the belief about an asthma episode would be much smaller (0.474 for HYP1) - and the person should conclude that Hypothesis 2 is correct.

From Bayes' Rule to Predictive Coding

Note that we have illustrated our two examples in terms of Bayesian inference, which cannot be directly computed by the brain. However, the theory of predictive coding suggests that the brain solves something analogous to the above Bayes' formula, using a hierarchical neural architecture[1]. In this architecture, *predictions* (derived from prior beliefs) are propagated in a top-down manner, and they are compared with perceptual and interoceptive evidence (via the likelihood model). The result of the comparison is called *prediction error*, and is propagated bottom-up in the hierarchy, to help updating the (posterior) probability of the initial hypothesis.

In our first example above, the brain would propagate a strong top-down prediction about an asthma episode (as the prior of Hypothesis 1 is high); and because the interoceptive evidence (wheezing) is largely compatible with this hypothesis, the resulting prediction error that is propagated bottom-up would be relatively low. Iterating this top-

---

1) Note that predictive coding uses continuous probability distributions (e.g., Gaussian) rather than the discrete distribution that we considered in the example of Bayes' rule. For simplicity, we ignore this difference here.

down (prediction) and bottom-up (prediction error) message passing would permit refining the initial hypotheses, setting the posterior probability of HYP1 to a value where prediction error is minimized - which in this case is (close to) 0.9739.

In our second example above, the brain would propagate a strong prediction about an asthma episode, too. However, because the interoceptive evidence (not wheezing) is incompatible with this hypothesis, the resulting prediction error would be very high - and after some iterations, the inference would settle to a (posterior) probability of 0.474 for HYP1.

Precision Weighting and Its Mis-Regulation in Psychopathology

Yet there is another aspect of Bayesian inference and predictive coding that we have ignored for now but is central to theories of psychopathologies. All the aforementioned top-down and bottom-up signals are weighted by their precision. Technically, precision is the inverse variance of a probability distribution (e.g., a continuous distribution, such as a Gaussian) and it can be used as a weight to each of the elements (priors and likelihoods) of the above Bayes' rule - with the effect that the more precise information has a stronger effect on the computations of the posterior probability, see Figure 1. Precision weighting is a convenient way to give more credit to the most reliable information sources and discard noisy evidence. For example, there may be conditions in which I cannot be sure about my sensory or interoceptive evidence (e.g., I don't know how I feel); in these cases, the evidence has to be down-weighted and thus the prior dominates the inference.

Trusting the prior is of course something sensible to do when evidence is scarce or unreliable. However, there are other and more pathological cases in which the prior may acquire a very high precision and dominate the inference, even if this is not optimal; and this may constitute a route to MUS. Let's expand our second example above (i.e., the case when one has a strong prior but no evidence for an asthma episode) by also considering that both the prior and the likelihood are weighted according to some precision value. If the precision of the prior is (for some reason) excessively high, one can obtain posterior probabilities for HYP1 that are much higher than our previous example (i.e., very close to prior probabilities, as in the central panel of Figure 1). The person would thus conclude incorrectly that he/she is experiencing an asthma episode. Furthermore, given that the predictive coding architecture continuously generates predictions about what it expects, the same person may also predict or "hallucinate" the wheezing that he is not experiencing (because it is highly compatible with the winning Hypothesis 1).

This example illustrates that priors that have acquired an excessively high precision may dominate the inference and fail to be correctly updated based on empirical evidence - thus potentially producing MUS. How can priors acquire unwarrantedly high precision? While accurate predictive coding requires the precision of top-down and bottom-up signals to be optimized (and would thus not produce MUS), there may be various pathological conditions that can lead to their mis-regulation. These may include deficits of neuro-

modulators like dopamine and noradrenaline, which in predictive coding are carriers of precision signals; or the exposure to the "wrong" environmental statistics, like when growing up with a chronically ill or health-anxious parent. These and other condition may lead to the formation of excessively precise priors that resist updating; and it is under these conditions that MUS may emerge.

A second possible way MUS (or similar phenomena) may emerge is the converse of the above example; and namely, when likelihoods have excessively (pathologically) low precision. Some pathologies may be related to deficits of interoceptive processing, in which one "does not know how he/she feels" (e.g. alexithymia, affective agnosia; Lane, Weihs, Herring, Hishaw, & Smith, 2015) or cannot easily attribute some interoceptive sensation (e.g., wheezing) to some cause (e.g., an asthma episode). In these cases, because the interoceptive signals are assigned a vanishingly small precision, they would be largely ignored during the inference - and again, the prior would dominate it.

From Predictive Coding to Active Inference

We discussed how, under a predictive coding scheme, deficits of precision weighting in either the prior or the likelihood (or both) can lead to maladaptive perceptual inference and MUS. The theory of active inference expands this view, by introducing additional ways these deficits may hinder correct inference and action selection. Here we focus on just one aspect of active inference: the fact that it induces an active sampling of information that is expected to have informative value, i.e., to gather relevant evidence.

When describing the asthmatic person's generative model, we have considered that wheezing is more informative than breathlessness, as the presence or absence of the former (but not the latter) disambiguates between Hypotheses 1 and 2. Active inference assumes that informative evidence is not passively gathered (as in predictive coding) but actively sampled; for example, by monitoring or directing attention to the relevant information sources (e.g., "attention to bodily signals"). Active inference would thus predict that under normal conditions, the asthmatic person should preferentially monitor and direct attention to its most informative signal: wheezing.

Yet, one can imagine a degenerate (technically, high-entropy) likelihood function, in which wheezing as a source of interoceptive evidence has degraded to a sensation that has exactly the same probability under both Hypotheses 1 and 2. In this case, monitoring wheezing would be useless, as it would bring exactly the same evidence for the two hypotheses. If a person's (likelihood) model of his bodily signals were degenerate, not only he would fail to recognize asthma symptoms, but he would also cease to attend to them - and more broadly, to pay attention to his bodily signals, similar to a form of "neglect" (Parr & Friston, 2018). In this case, he would only be able to infer asthma from the prior belief or other, non-bodily sources of information (e.g., what the others around me believe about my asthma) that may not be particularly reliable. Ignoring bodily signals

would thus render this person prone to MUS, as well as to deficits of body schema and self-representations that may have a strong bodily basis (Pezzulo, 2014; Seth, 2013).

A degenerate (likelihood) model of bodily signals may arise from neurological or peripheral disorders that make bodily signals noisier. However, it can also be the consequence of a poor learning and developmental processes, which can lead to the acquisition of internal models that are insufficiently differentiated and do not permit to appropriately categorize one's own interoceptive signals (Petersen, Schroijen, Mölders, Zenker, & Van den Bergh, 2014).

Conclusions

We discussed symptom perception and MUS from the perspective of predictive coding and active inference. Our examples illustrate the fact that there are various ways by which the components of a person's generative model (prior and likelihood) can be assigned too high or too low precision or become "unbalanced". This, in turn, may produce (momentary) incorrect inference or action selection or (more chronic) psychopathological conditions.

Formal theories like predictive coding and active inference can help dissecting these possibilities and identifying their markers during development. However, these conceptual models also imply important challenges to test and validate them. One way is to flesh out a computational version of the model involving a clear mechanistic description of the critical variables and their interactions, to run simulations and compare the results with evidence from real life (Friston et al., 2017; Petzschner, Weber, Gard, & Stephan, 2017; Stephan et al., 2016).

References

Allen, M., Levy, A., Parr, T., & Friston, K. J. (2019). In the body's eye: The computational anatomy of interoceptive inference. *bioRxiv*. Article 603928. https://doi.org/10.1101/603928

Barca, L., & Pezzulo, G. (2018). Keep your interoceptive stream under control: An active inference perspective on Anorexia Nervosa. *OSF Preprints*. https://doi.org/10.31219/osf.io/2s7qk

Barrett, L. F., & Simmons, W. K. (2015). Interoceptive predictions in the brain. *Nature Reviews Neuroscience, 16*(7), 419-429. https://doi.org/10.1038/nrn3950

PsychOpen GOLD

Büchel, C., Geuter, S., Sprenger, C., & Eippert, F. (2014). Placebo analgesia: A predictive coding perspective. *Neuron, 81*(6), 1223-1239. https://doi.org/10.1016/j.neuron.2014.02.042

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences, 36*(3), 181-204. https://doi.org/10.1017/S0140525X12000477

De Peuter, S., Van Diest, I., Lemaigre, V., Li, W., Verleden, G., Demedts, M., & Van den Bergh, O. (2005). Can subjective asthma symptoms be learned? *Psychosomatic Medicine, 67*(3), 454-461. https://doi.org/10.1097/01.psy.0000160470.43167.e2

De Waal, M. W. M., Arnold, I. A., Eekhof, J. A. H., & van Hemert, A. M. (2004). Somatoform disorders in general practice: Prevalence, functional impairment and comorbidity with anxiety and depressive disorders. *The British Journal of Psychiatry, 184*, 470-476. https://doi.org/10.1192/bjp.184.6.470

Edwards, M. J., Adams, R. A., Brown, H., Pareés, I., & Friston, K. J. (2012). A Bayesian account of 'hysteria'. *Brain: A Journal of Neurology, 135*(11), 3495-3512. https://doi.org/10.1093/brain/aws129

Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences, 360*(1456), 815-836. https://doi.org/10.1098/rstb.2005.1622

Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience, 11*(2), 127-138. https://doi.org/10.1038/nrn2787

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., O'Doherty, J., & Pezzulo, G. (2016). Active inference and learning. *Neuroscience & Biobehavioral Reviews, 68*, 862-879. https://doi.org/10.1016/j.neubiorev.2016.06.022

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active inference: A process theory. *Neural Computation, 29*(1), 1-49. https://doi.org/10.1162/NECO_a_00912

Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive Neuroscience, 6*(4), 187-214. https://doi.org/10.1080/17588928.2015.1020053

Friston, K., Samothrakis, S., & Montague, R. (2012). Active inference and agency: Optimal control without cost functions. *Biological Cybernetics, 106*(8–9), 523-541. https://doi.org/10.1007/s00422-012-0512-8

Henningsen, P., Gündel, H., Kop, W. J., Löwe, B., Martin, A., Rief, W., . . . EURONET-SOMA Group. (2018). Persistent physical symptoms as perceptual dysregulation: A neuropsychobehavioral model and its clinical implications. *Psychosomatic Medicine, 80*(5), 422-431. https://doi.org/10.1097/PSY.0000000000000588

Hohwy, J. (2013). *The predictive mind.* Oxford, United Kingdom: Oxford University Press.

Iodice, P., Porciello, G., Bufalari, I., Barca, L., & Pezzulo, G. (2019). An interoceptive illusion of effort. *PNAS, 116*(28), 13897-13902. https://doi.org/10.1073/pnas.1821032116

Janssens, T., Verleden, G., De Peuter, S., Van Diest, I., & Van den Bergh, O. (2009). Inaccurate perception of asthma symptoms: A cognitive-affective framework and implications for asthma treatment. *Clinical Psychology Review, 29*(4), 317-327. https://doi.org/10.1016/j.cpr.2009.02.006

PsychOpen GOLD

Lane, R. D., Weihs, K. L., Herring, A., Hishaw, A., & Smith, R. (2015). Affective agnosia: Expansion of the alexithymia construct and a new opportunity to integrate and extend Freud's legacy. *Neuroscience and Biobehavioral Reviews, 55*, 594-611. https://doi.org/10.1016/j.neubiorev.2015.06.007

Nimnuan, C., Hotopf, M., & Wessely, S. (2001). Medically unexplained symptoms: An epidemiological study in seven specialities. *Journal of Psychosomatic Research, 51*(1), 361-367. https://doi.org/10.1016/S0022-3999(01)00223-9

Parr, T., & Friston, K. J. (2018). Active inference, novelty and neglect. In T. Hodgson (Ed.), *Processes of visuospatial attention and working memory* (pp. 115-128). https://doi.org/10.1007/7854_2018_61

Petersen, S., Schroijen, M., Mölders, C., Zenker, S., & Van den Bergh, O. (2014). Categorical interoception: Perceptual organization of sensations from inside. *Psychological Science, 25*(5), 1059-1066. https://doi.org/10.1177/0956797613519110

Petzschner, F. H., Weber, L. A. E., Gard, T., & Stephan, K. E. (2017). Computational psychosomatics and computational psychiatry: Toward a joint framework for differential diagnosis. *Biological Psychiatry, 82*(6), 421-430. https://doi.org/10.1016/j.biopsych.2017.05.012

Pezzulo, G. (2014). Why do you fear the Bogeyman? An embodied predictive coding model of perceptual inference. *Cognitive, Affective, and Behavioral Neuroscience, 14*(3), 902-911. https://doi.org/10.3758/s13415-013-0227-x

Pezzulo, G., Barca, L., & Friston, K. J. (2015). Active inference and cognitive-emotional interactions in the brain. *Behavioral and Brain Sciences, 38*, Article e85. https://doi.org/10.1017/S0140525X14001009

Pezzulo, G., Maisto, D., Barca, L., & Van den Bergh, O. (2019). Perception and misperception of bodily symptoms from an Active Inference perspective: Modelling the case of panic disorder. *OSF Preprints.* https://doi.org/10.31219/osf.io/dywfs

Pezzulo, G., Rigoli, F., & Friston, K. J. (2015). Active inference, homeostatic regulation and adaptive behavioural control. *Progress in Neurobiology, 134*, 17-35. https://doi.org/10.1016/j.pneurobio.2015.09.001

Pezzulo, G., Rigoli, F., & Friston, K. J. (2018). Hierarchical active inference: A theory of motivated control. *Trends in Cognitive Sciences, 22*(4), 294-306. https://doi.org/10.1016/j.tics.2018.01.009

Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience, 2*(1), 79-87. https://doi.org/10.1038/4580

Seth, A. K. (2013). Interoceptive inference, emotion, and the embodied self. *Trends in Cognitive Sciences, 17*(11), 565-573. https://doi.org/10.1016/j.tics.2013.09.007

Stephan, K. E., Manjaly, Z. M., Mathys, C. D., Weber, L. A. E., Paliwal, S., Gard, T., . . . Petzschner, F. H. (2016). Allostatic self-efficacy: A metacognitive theory of dyshomeostasis-induced fatigue and depression. *Frontiers in Human Neuroscience, 10*, Article 550. https://doi.org/10.3389/fnhum.2016.00550

PsychOpen GOLD

Sterzer, P., Adams, R. A., Fletcher, P., Frith, C., Lawrie, S. M., Muckli, L., . . . Corlett, P. R. (2018). The predictive coding account of psychosis. *Biological Psychiatry, 84*(9), 634-643. https://doi.org/10.1016/j.biopsych.2018.05.015

Tsakiris, M., & Preester, H. D. (2018). *The interoceptive mind: From homeostasis to awareness.* Oxford, United Kingdom: Oxford University Press.

Van den Bergh, O., Witthöft, M., Petersen, S., & Brown, R. J. (2017). Symptoms and the body: Taking the inferential leap. *Neuroscience and Biobehavioral Reviews, 74*(Pt A), 185-203. https://doi.org/10.1016/j.neubiorev.2017.01.015

Wiech, K. (2016). Deconstructing the sensation of pain: The influence of cognitive processes on pain perception. *Science, 354*(6312), 584-587. https://doi.org/10.1126/science.aaf8934