PAIN®

Symptom perception, placebo effects, and the Bayesian brain

Giulio Ongaro[a,*], Ted J. Kaptchuk[b,c]

1. Introduction

The standard and ideal biomedical model of symptom perception treats the brain largely as a passive stimulus-driven organ. It embraces the notion that the brain absorbs sensory signals from the body and converts them, directly, into conscious experience. Accordingly, biomedicine operates under the assumption that symptoms are the direct consequences of physiological dysfunction and improvement is the direct consequence of the restoration of bodily function. Despite its success, the biomedical model has failed to provide an adequate account of 2 well-demonstrated phenomena in medicine: (1) the experience of symptoms without pathophysiological disruption, and (2) the experience of relief after the administration of placebo treatments. This topical review advances the idea that "predictive processing," a Bayesian approach to perception that is rapidly taking hold in neuroscience, significantly helps accommodating these 2 phenomena. It expands on recent high-quality empirical work on predictive processing[1,7,19,24] and outlines, more broadly, how Bayesian models offer an altogether different picture of how the brain perceives symptoms and relief.

2. The Bayesian brain

The nervous system is constantly dealing with a continuous and potentially overwhelming stream of varying signals coming from our body and senses. For the sake of adaptation, the brain must turn this confused play of sensory inputs and neural firings into a reliable perception of the world. Debate in cognitive science has revolved around how exactly the brain accomplishes this task. While previous theories, in line with the current biomedical model of disease, viewed perception mostly as a bottom-up readout of sensory signals, emerging Bayesian models suggest, instead, that perception is cognitively (mostly nonconsciously) modulated,

and might be best viewed as a process of prediction, based on an integration of sensory inputs, prior experience, and contextual cues.[9,10,15,21]

The key suggestion is that to perceive the world, the brain follows a theory of probability known as Bayes rule. In its mathematical form, the rule updates the likelihood of a given hypothesis (or "prior"), given some evidence, by considering the product of the likelihood and the prior probability of the hypothesis.[21] Over rapid time scales, the brain implements Bayes rule by continuously generating a top-down cascade of neurally encoded (mostly nonconscious) hypotheses about the state of the body and world. This top-down flow of hypotheses is met by the bottom-up stream of sensory inputs coming from the senses. Any mismatch between predicted input and actual input results in "prediction error," which prompts the system to revise its hypotheses. Bottom-up perception is thus inseparable from top-down prediction. Some of the hypotheses that account for the most abstract and general features of the world are "built-in" by evolution; others are amenable to progressive refinement through developmental learning. Throughout one's lifespan, the nervous system engages in the continuous updating of these priors to better predict the next incoming sensory inputs and minimize error.

A central implication of the theory is that what we perceive is not the world as it actually is, but the brain's best guess of it, continuously refined by incoming sensory evidence.[10,21] Visual perception, the domain from which much of the evidence for the Bayesian brain has emerged, offers the most intuitive way to grasp its key principle. To mention a simplistic example, sticks in a forest that is infested by snakes might at first be perceived as snakes, until we get a more refined view that updates the hypothesis.

Importantly, the interplay between descending predictions and ascending signals that lie at the heart of predictive processing is flexibly modulated by the "precision" (or "inverse variance," in statistical terms) of hypotheses and sensory evidence. Faced with the task of determining how likely a given set of inputs represents a predicted state, the brain uses prior experience and subtle contextual cues to determine their precision. The example of seeing sticks as snakes represents a case where highly precise hypotheses shaped by previous experience (knowing what a snake looks like, knowing that snakes inhabit the forest) override imprecise visual inputs. Indeed, the Bayesian brain model is able to explain how, in contexts of precise predictions and imprecise inputs, perceptions can deviate from the actual state of the world.[10,33] Conversely, it also elucidates how inferences can be made under conditions of ambiguity that lack precise estimates.[17,28] The model is supported by growing computational and neuroimaging evidence, and advances the notion that the

a Department of Anthropology, London School of Economics and Political Science, London, United Kingdom, b Harvard Medical School, Harvard University, Boston, MA, United States, c Program in Placebo Studies, Beth Israel Deaconess Medical Center, Harvard University, Boston, MA, United States

*Corresponding author. Address: Department of Anthropology, London School of Economics, Houghton St, London WC2A 2AE, United Kingdom. Tel.: +44 75 40382275. E-mail address: g.ongaro@lse.ac.uk (G. Ongaro).

precision of descending predictions might be "encoded" in the brain by neurotransmitters such as dopamine.[16]

3. Symptom perception and "medically unexplained symptoms"

The idea that what we perceive is not the world as it is but our own best hypothesis of it equally applies to the body[2,27] and subjective bodily states such as medical symptoms. We do not necessarily feel pain—this framework suggests—because we "sense" it directly from the peripheral body. To put it emphatically, we feel pain because we predict that we are in pain, based on an integration of sensory inputs, prior experience, and contextual cues.

The experience of symptoms arises out of the inference that the body has deviated from the physiological constants that define health. From a Bayesian perspective, the experience of health depends on the fact that we maintain a general "healthy body condition" hypothesis (partly determined by evolution, partly by development) that explains away a certain range of normal variations in somatic input (eg, variations in heartbeat frequency, bodily aches, etc.).[30] So long as these variations are kept within the bounds predicted by the "healthy body condition" hypothesis, the brain treats them as "noise" and no symptom is perceived. When, due to a disrupting cause such as disease, the variation of somatic inputs is too large to be successfully predicted by the general hypothesis and prediction error increases, the brain must generate another hypothesis that accounts for the new evidence. According to the theory, we feel symptoms, including pain, when the hypothesis with the lowest prediction error represents an abnormal somatic event.[30] This framing ultimately shows that pathophysiology may be only loosely coupled with symptoms perception because the latter is mediated by internally generated hypotheses about the causes of inputs, not solely determined by inputs themselves. The degree of the correlation between pathophysiology and symptoms will vary according to the relative precision assigned to inputs and hypotheses, respectively.

When a subject unexpectedly encounters a certain painful stimulus for the first time, the ongoing hypothesis that the system is healthy is quickly revised on meeting unambiguous sensory evidence that departs from it. Given that the system has had no previous exposure to the stimulus, sensory signals have higher precision relative to prior hypotheses, and therefore a higher impact on perception. This is why in cases of localized dysfunction and acute pain, we find a high correlation between pathophysiology and symptom perception.

However, for many chronic subjective symptoms, which often involve central sensitization, somatization, aberrant nociceptive amplification, or ambiguous, frequently shifting information, the process can reverse. Here, the perception of symptoms shifts in the direction of the hypotheses generated by the brain, which explains the low correlation that we find between objective pathophysiology and subjective experience across a number of such chronic conditions.[13] From a Bayesian perspective, chronic pain reflects the high precision that is placed on hypotheses vis-a-vis sensory evidence. Slight and harmless variations in interoceptive inputs in certain contexts (which in healthy individuals would be treated as "noise") prompt the brain to mistakenly infer pain as the cause of these inputs,[20] and to feel pain accordingly. Anxiety, fear, threat perception, and catastrophizing, emotional states that often accompany this disorder, have the effect of worsening symptoms by maintaining vigilance to predicted pain.[34]

Furthermore, in the context of chronic pain, the brain does not merely passively perceive pain, but can also play a part in its own intensification. This is so because under the predictive processing framework, another way of minimizing prediction error lies in the generation of bodily action. If, in the context of perception, the brain revises its predictions to match the input, in action, it minimizes prediction error by modifying the inputs so that they can fit the prediction.[6,10] In a condition of chronic pain, the brain may nonconsciously initiate visceral sensations (eg, stomach tension) that match the hypothesis of being in pain. In all this, we see the brain, in a context of precision imbalance, continuing to perform its ordinary Bayesian task of minimizing prediction error to conform inputs to predictions, even if at the detriment of subjective well-being.[20]

The framework invites us to appreciate the salient role played by the social context in shaping and reinforcing predictions of symptoms. Knowledge of a drug's side effects, the verbal information about imminent pain delivered by a physician, or a culturally specific way of attending to our body heighten the precision of the hypothesis of impending symptoms, leading it to dominate symptom perception. The so called "nocebo effects," in which negative symptoms attributed to medication arise independently of biological activity, fall within the same set of processes.[3,11,29]

Importantly, the approach goes some way towards transcending the artificial but pervasive distinction between "explained" and "unexplained" symptoms in biomedicine (or between "real" and "imaginary" illnesses). One important upshot of the theory is that all symptoms are product of an inferential process that is never strictly reducible to physiological dysfunction and is sometimes only loosely related or unrelated to it. "Explained" and "unexplained" symptoms thus lie on a continuum, differing only in the extent to which they are coupled to an organic disorder. Given that the same inferential process is implicated in both cases, the theory also explains why the so called "real" and "imaginary" symptoms seem to be phenomenologically indistinguishable from the patient's point of view.[30]

4. Symptom relief and placebo effects

A very similar story, if in reverse, applies to the relief of symptoms. From a Bayesian perspective, the experience of recovery is not the direct consequence of the restoration of bodily function, but is itself the process of inferring that certain interoceptive changes are signs that this improvement is taking place. The ongoing hypothesis that we are ill must be revised on meeting evidence that the body is returning to a "healthy body condition." This revision of hypotheses, however, is usually slower or hard to occur if the person is not given any external cues that amelioration is underway. Without receiving this information, the brain might explain away the variation in interoceptive input that follows an effective medical intervention as mere "noise" and adheres to a hypothesis of ongoing pain. Experiments in acute experimental models conducted within the open-hidden paradigm in placebo studies show this very clearly, for they demonstrate that patients who are administered symptom-relieving drugs (eg, analgesics, anxiolytics) in a covert manner tend to experience a much lower relief than patients who are given treatment in full-view.[4,5] The medical ritual prompts the brain to interpret even small interoceptive changes in the body as the consequence of healing, and to experience relief accordingly. Such predictions are self-fulfilling.[1,7,18,19,32] Simultaneously, under precise predictions of incoming health, the brain can also arrive at symptom relief

through processes of active inference. The brain, in short, may initiate healthful visceral sensations (eg, relaxing stomach muscles) that conform to the hypothesis that one is returning to a "healthy body condition"—all this with the purpose of fulfilling the prediction and minimizing error.

Many elements of the therapeutic context can play a role in enhancing predictions of well-being, especially in chronic situations. Experiments have shown that care and supportive verbal suggestions communicated by a trusted clinician are central in eliciting placebo effects,[25] as is the perceived value of the treatment itself. For instance, inert treatments presented as costly to patients tend to be more effective than treatments that patients know are less expensive.[14,23] Social learning–observing first-hand the effects of a treatment on others–also greatly contributes to the formation of placebo responses.[12] Importantly, recent evidence suggests that features of the therapeutic ritual can be effective even when apprehended subliminally.[22] This evidence is fully compatible with the finding that patients can receive benefit in an "open-label placebo" paradigm, in which they are honestly told that they are receiving a placebo.[8,24,26] From a predictive processing perspective, some of this response is probably triggered because of unconscious predictions sparked by the embodied assumption of medication taking[21] and by being in a clinical environment associated with efficacy. The response may also be related to inferences under ambiguity. Two contradictory messages embedded in open-label placebos —"this inert placebo pill may help; this placebo pill cannot work"—may create heightened neurological, cognitive, and embodied dissonance leading to nonconscious inferences that disturb central sensitization.[24]

Obviously, positive predictions of relief are often insufficient to lead to full or even partial recovery. Predictive processing explains the reason why, in the presence of strong physiopathology, placebo effects tend to be difficult to elicit. If a highly weighted prediction of impending relief is met with strong sensory evidence of the contrary, the brain will eventually infer that the body is still in pain. In fact, therapeutic rituals alone tend not to work on physiological conditions that lie outside the reach of the nervous system, and are mostly effective on symptoms of self-appraisal that are uncoupled from pathophysiology.[31]

Crucially, the predictive processing approach shows that therapeutic ritual and active ingredients of the intervention, albeit through different routes, act on the same inferential process whereby we experience symptoms relief. The first strengthens predictions of impending health by offering external evidence that recovery is taking place (through the ritual drama, verbal interaction, etc.). The second strengthens predictions of impending health by removing the source of nociceptive inputs, or, in the case of symptom-relieving drugs, by stimulating neurotransmitters that encode for the precision of top-down predictions.[7] Given that, whichever the route, the same basic inferential process lies at the heart of symptom relief, the framework explains why healing that is primarily related to a medical intervention (so called "treatment effect") and healing that is related to the therapeutic ritual (so called "placebo effect") seem to be as equally real from the patient's point of view.

5. Conclusions

Symptoms without a physical cause and relief through placebo intervention are anomalies for the biomedical model of disease. The Bayesian approach to perception explains and accommodates these 2 phenomena. It exposes placebo and nocebo effects, not as aberrant events, but as facets of the overall modus operandi of the nervous system. It shows, also, that these act on the same inferential processes as "real" disease and "real" treatments do. The implication of this approach is that, to be truly patient-focused, medicine must attend to the predictive process that lies at the basis of symptom perception, and thereupon evaluate what efficient courses of action can lead the brain to predict the body's health.

Conflict of interest statement

The authors have no conflict of interest to declare.

References

[1] Anchisi D, Zanon MA. Bayesian perspective on sensory and cognitive integration in pain perception and placebo analgesia. PLoS One 2015;10: e0117270.
[2] Barrett LF, Simmons WK. Interoceptive predictions in the brain. Nat Rev Neurosci 2015;16:419–29.
[3] Benedetti F, Lanotte M, Lopiano L, Colloca L. When words are painful: unraveling the mechanisms of the nocebo effect. Neuroscience 2007; 147:260–71.
[4] Benedetti F, Maggi G, Lopiano L, Lanotte M, Rainero I, Vighetti S, Pollo A. Open versus hidden medical treatments: the patient's knowledge about a therapy affects the therapy outcome. Prev Treat 2003;6:1a.
[5] Benedetti F, Carlino E, Pollo A. Hidden administration of drugs. Clin Pharmacol Ther 2011;90:651–61.
[6] Brown H, Adams RA, Parees I, Edwards M, Friston K. Active inference, sensory attenuation and illusions. Cogn Process 2013;14:411–27.
[7] Büchel C, Geuter S, Sprenger C, Eippert F. Placebo analgesia: a predictive coding perspective. Neuron 2014;81:1223–39.
[8] Carvalho C, Caetano JM, Cunha L, Rebouta P, Kaptchuk TJ, Kirsch I. Open-label placebo treatment in chronic low back pain: a randomized controlled trial. PAIN 2016;157:2766.
[9] Clark A. Whatever next? Predictive brains, situated agents, and the future of cognitive science. Behav Brain Sci 2013;36:181–204.
[10] Clark A. Surfing uncertainty: prediction, action, and the embodied mind. Oxford: Oxford University Press, 2015.
[11] Colagiuri B, Quinn VF, Colloca L. Nocebo hyperalgesia, partial reinforcement, and extinction. J Pain 2015;16:995–1004.
[12] Colloca L, Benedetti F. Placebo analgesia induced by social observational learning. PAIN 2009;144:28–34.
[13] Edwards MJ, Adams RA, Brown H, Parees I, Friston KJ. A Bayesian account of "hysteria". Brain 2012;135:3495–512.
[14] Espay AJ, Norris MM, Eliassen JC, Dwivedi A, Smith MS, Banks C, Allendorfer JB, Lang AE, Fleck DE, Linke MJ, Szaflarski JP. Placebo effect of medication cost in Parkinson disease a randomized double-blind study. Neurology 2015;84:794–802.
[15] Friston K. A theory of cortical responses. Philos Trans R Soc Lond B Biol Sci 2005;360:815–36.
[16] Friston K, Shiner T, Fitzgerald T, Galea JM, Adams R, Brown H, Dolan RJ, Moran R, Stephan KE, Bestmann S. Dopamine, affordance and active inference. PLoS Comput Biol 2012;8:e1002327.
[17] Gersham S. A unifying probabilistic view of associative learning. PLoS Comput Biol 2015;11:e1004567.

[18] Geuter S, Koban L, Wager TD. The cognitive neuroscience of Placebo effects: concepts, predictions, and physiology. Annu Rev Neurosci 2017; 40:167–88.

[19] Grahl A, Onat S, Büchel C. The periaqueductal gray and Bayesian integration in placebo analgesia. Elife 2018;7:e32930.

[20] Hechler T, Endres D, Thorwart A. Why harmless sensations might hurt in individuals with chronic pain: about heightened prediction and perception of pain in the mind. Front Psychol 2016;7:1638.

[21] Hohwy J. The predictive mind. Oxford: Oxford University Press, 2013.

[22] Jensen KB, Kaptchuk TJ, Kirsch I, Raicek J, Lindstrom KM, Berna C, Gollub RL, Ingvar M, Kong J. Nonconscious activation of placebo and nocebo pain responses. Proc Natl Acad Sci 2012;109:15959–64.

[23] Kam-Hansen S, Jakubowski M, Kelley JM, Kirsch I, Hoaglin DC, Kaptchuk TJ, Burstein R. Altered placebo and drug labeling changes the outcome of episodic migraine attacks. Sci Transl Med 2014;6:218ra5.

[24] Kaptchuk TJ. Open-label placebos: reflections on a research agenda. Perspect Biol Med 2018;61:311–34.

[25] Kaptchuk TJ, Kelley JM, Conboy LA, Davis RB, Kerr CE, Jacobson EE, Kirsch I, Schyner RN, Nam BH, Nguyen LT, Park M. Components of placebo effect: randomised controlled trial in patients with irritable bowel syndrome. BMJ 2008;336:999–1003.

[26] Kaptchuk TJ, Friedlander E, Kelley JM, Sanchez MN, Kokkotou E, Singer JP, Kowalczykowski M, Miller FG, Kirsch I, Lembo AJ. Placebos without deception: a randomized controlled trial in irritable bowel syndrome. PLoS One 2010;5:e15591.

[27] Seth AK, Suzuki K, Critchley HD. An interoceptive predictive coding model of conscious presence. Front Psychol 2012;2:395.

[28] Tenenbaum JB, Kemp C, Griffiths TL, Goodman ND. How to grow a mind: statistics, structure, and abstraction. Science 2011;331:1279–85.

[29] Van den Bergh O, Brown RJ, Petersen S, Witthöft M. Idiopathic environmental intolerance: a comprehensive model. Clin Psychol Sci 2017;5:551–67.

[30] Van den Bergh O, Witthöft M, Petersen S, Brown RJ. Symptoms and the body: taking the inferential leap. Neurosci Biobehav Rev 2017;74: 185–203.

[31] Wechsler ME, Kelley JM, Boyd IO, Dutile S, Marigowda G, Kirsch I, Israel E, Kaptchuk TJ. Active albuterol or placebo, sham acupuncture, or no intervention in asthma. N Engl J Med 2011;365:119–26.

[32] Wiech K. Deconstructing the sensation of pain: the influence of cognitive processes on pain perception. Science 2016;354:584–7.

[33] Wilkinson S. Accounting for the phenomenology and varieties of auditory verbal hallucination within a predictive processing framework. Conscious Cogn 2014;30:142–55.

[34] Zaman J, Vlaeyen JW, Van Oudenhove L, Wiech K, Van Diest I. Associative fear learning and perceptual discrimination: a perceptual pathway in the development of chronic pain. Neurosci Biobehav Rev 2015;51:118–25.