Meditative in-action: an endogenous epistemic venture

Giuseppe Pagnoni^{1,2*} and Fausto Taiten Guareschi³

¹*Department of Biomedical, Metabolic and Neural Sciences, University of Modena and Reggio Emilia, Modena, Italy*

²*Center for Neuroscience and Neurotechnology, University of Modena and Reggio Emilia, Modena, Italy*

³*Istituto Italiano Zen Sōtō Shōbōzan Fudenji, Salsomaggiore Terme, PR, Italy*

**Corresponding author: giuseppe.pagnoni@unimore.it*

UPDATED 2022/01/25

Abstract

The theoretical framework of active inference proposed by Karl Friston is currently one of the more actively developed research areas in neuroscience. According to this theory the brain adapts its synaptic activity and architecture in such a way that it *de facto* comes to mirror the causal structure of events that the organism both encounters and actively induces in its environment. In this contribution we show how active inference can provide a useful perspective to better understand the processes engaged by contemplative practices. More specifically, and focusing on the practice of *shikantaza* ('just sitting') in the Japanese Zen Sōtō tradition, we argue that meditation enacts a peculiar policy with high endogenous epistemic value, whereby the practitioner accrues an intimate, but not necessarily explicit, knowledge about herself. This superordinate policy entails the embodied, active suspension of our habitual reward-seeking and punishment-avoidance behavior, an attitude epitomized by the traditional notion of *mushotoku* (Jap. 'nothing to be attained'). From this perspective, we also critically examine some popular claims about meditation that we believe are likely to be misconstrued, and may not be without personal, social, and even political consequences.

Ma sedendo e mirando, interminati
spazi di là da quella, e sovrumani
silenzi, e profondissima quiete
io nel pensier mi fingo

But sitting here in a daydream, I picture
The boundless spaces away out there, silences
Deeper than human silence, an unfathomable
hush

L'Infinito, Giacomo Leopardi
(transl. Eamon Grennan)

1 A minimal introduction to active inference

The term *active inference* refers to a theoretical framework aiming to understand how self-organizing systems, such as biological organisms, move, perceive, and generally behave in their environment. It is based on the principle of free-energy minimization proposed by Karl J. Friston (Friston, 2010) and, while its origins lie in the field of neurosciences, it appears to possess a remarkably ample explanatory and applicative scope and is thus being very rapidly developed (Friston, 2019) (for a hands-on, modeling-based tutorial on the core notions of active inference, see Smith et al., 2021).

The central notion of active inference is that biological systems keep their identity and viability by minimizing the likelihood of incurring into surprising states over the long term. The role of surprise minimization as a fundamental operating principle may seem bizarre at first, but it becomes easier to understand if we consider the physiological concept of homeostasis. An organism can maintain its vital status and identity only by keeping its internal, physiological parameters (e.g., glucose and oxygen blood levels, pH, temperature, etc.) within a narrow range of variation. Trespassing such bounds represents an abrupt surprise¹ for what concerns the internal states of the organism, which the latter will go to great lengths to avoid². In more general terms, active inference and the principle of free energy minimization represent a formulation of how living organisms are able to shield themselves temporarily from the entropic dissolution described by the fluctuation theorem of statistical mechanics. It is intriguing to note that long before modern physics, the archetypal notions of order and

¹ ‘Surprise’ here does not necessarily imply psychological surprise. It refers in fact to an information-theoretical quantity and, to avoid confusion, the term ‘surprisal’ is often used instead in the active inference literature.

² While the concept of homeostasis is used here for the sake of simplicity, active inference aligns conceptually better to the notion of *allostasis*, which is the maintenance of vital parameters within the range appropriate for the current or expected task requirements (think heart rate and strenuous exercise, for example).

disorder (i.e., cosmos and chaos) were also recognized as fundamental, dynamically opposing forces (but generatively so) in many ancient cultures, and thus placed at the core of their respective creational myths. It is also worth mentioning that active inference’s foundational ideas are indeed closely related to the pioneering work of Maturana and Varela on biological autonomy and sensorimotor coupling, leading to the development of the concepts of autopoiesis and enactive cognition (Varela et al. (1974); Maturana and Varela (1980); Varela et al. (1991); but see also Di Paolo et al. (2021), for a critical perspective on this matter).

Surprise minimization rests necessarily upon the concept of a *generative model*, that is, a model of how hidden (worldly or bodily) causes generate observable quantities (i.e., sensations). The inversion of the generative model allows the organism to *infer* the hidden causes of the current sensations, in turn generating further *expected* sensations. Surprise refers to the mismatch between expected and actual sensations.

But where does such generative model come from or, in other words, what kind of process *generates* the generative model? The basic idea here is that the history of interaction between an organism and its environmental niche literally *in-forms* the structure and causal dynamics of both organism and environment. In other words, the continuous ‘friction’ between the living agent and its world gives rise in time (and along multiple time scales) to a modification of the internal structure of the agent, so that its dynamics comes to reflect, *re-present* or *re-enact* that of the causal chains of events that the organism habitually encounters in its environment: the agent comes to *embody* a generative model of its world. The environmental niche gets also similarly sculpted by the organisms inhabiting it, although this is an aspect that concerns us only marginally here (the interested reader may want to consult Rubin et al., 2020). On the other hand, it is not irrelevant for our purposes to mention that, for a social species like ours, the interaction with other individuals is of utmost importance and thus the relationship with the other is, at a very fundamental level, constitutive of our own identity (the generative model that we *are*). In fact, it has been argued that the so-called Cambrian explosion (around 600 million years ago), when most of the animal phyla appeared, was the period where inferring the identity, behavior and intention of other animals became the cardinal issue for survival and reproduction, and therefore the leading force for the evolution of the nervous (and sensorimotor) system and, ultimately, of the mind *in response to other minds* (Smith, 2017). Furthermore, as the gamut of sensorimotor interactions that an organism could entertain with its niche expanded along its evolutionary trajectory, spanning a greater range of spatio-temporal scales, the architecture of its internal generative model had necessarily to become progressively more hierarchical, to be able to reflect the richer and deeper structure of causal dependencies and statistical regularities encountered in its world.

Notably, from the perspective we have been describing, a living system can be seen as if continuously anticipating its sensory input (whence the term *predictive processing* or *predictive coding*, often used to describe this feature) under the single imperative to minimize the discrepancy (*prediction error*), between

the expected and the actual sensory signals. This process of prediction error minimization is thought to underlie not only perception and learning, but also decision-making and action selection, whereby an agent decides its course of behavior (*policy selection*) on the basis of its beliefs about (1) how likely a given sequence of actions is to bring about its preferred outcomes (the *pragmatic* value of a policy), and (2) how much information gain it will produce in the service of reducing uncertainty (the *epistemic* value of a policy) and thus improving future predictions. It is indeed the aspect of agency that takes active inference beyond vanilla predictive coding or Bayesian brain formulations: the ability of sentient beings to *actively* solicit from the environment the data they need to confirm or disconfirm their hypotheses, which on longer time scales underwrites the process of niche construction. Under this light, the organism appears to maintain its viability by continuously working to gather the sensory evidence expected by the generative model that it *is*, whereby the notion of ‘self-evidencing’ (Hohwy, 2014).

An important aspect of the theory is that the generative model must necessarily include the agent itself in its set of representations, to be able to anticipate the consequences of its own actions and thus choose which sequences of actions (*policies*) will likely bring about the more valuable states for its phenotype. This requirement entails the extension of the predictive horizon far beyond the immediate contingency of ‘what will happen now’, so that the selection of a given course of action by the agent is in fact determined by the *expected* reduction of free energy associated with that policy in the future. This temporal aspect is also particularly interesting for our purposes, because the degree of sentience of a living system has been linked to the temporal ‘depth’ or ‘thickness’ of its generative model (Friston, 2018). In brief, it has been argued that the longer the temporal arcs the system can cast in its predictive and postdictive processing³, the greater its sentient capacity. The terms ‘depth’ or ‘thickness’ refers to the hierarchical architecture of the generative model, where higher layers encode representations of progressively coarser spatio-temporal grain – in other words, in moving up the hierarchy of the model, percepts gradually give way to concepts. Notably, the brain seems to be an organ exquisitely specialized for this kind of operations, due to its extremely rich connectivity that allows the implementation of the required web of hierarchically structured and recursive informational message-passing.

Now, while perception, learning and movement can be formulated as processes that unfold based only on the minimization of the *current* prediction error, the higher faculties of decision making and policy selection — where the term policy denotes a purposeful sequence of actions — require also the minimization of the prediction error of the *future* (i.e., the expected prediction error), because the sensory observations upon which the computation of prediction error is based have obviously not occurred yet. It can be shown that mathematically this entails the joint minimization of two informational quantities that can be

³ *Postdiction* is intended here as the capacity to infer or locate the causes of the current state of affairs in the past or, in other words, to interpret the present on the basis of past events.

called *risk* and *ambiguity*, respectively. Risk scores the probability of reaping sensory data, under the selected policy, that differ considerably from the ones that define the agent’s phenotype: in other words, the probability of incurring into non-preferred states. Ambiguity measures how uncertain (imprecise) the observations are likely to be when the system is in the states aimed to by a policy; a reduction of ambiguity is associated with an improved predictive efficacy of the model because the ensuing observations will be more precise, providing more reliable teaching signals for the generative model. In summary, minimization of current free energy (linked to current prediction error) balances the accuracy of the model — how closely it aligns with the world — with its complexity — how much it needs to change to accommodate the new observations. On the other hand, minimization of *expected* free energy (linked to expected prediction error) balances the tendency of the organism to secure its preferred states (the pragmatic value of a policy) with a drive to gather novel information that can be useful in the future (the epistemic value of a policy). Note that while epistemic foraging may look like surprise-seeking, it is in fact sensation-seeking in the service of information gathering for future resolution of uncertainty. Also, it is worthwhile to emphasize here that minimization of *expected* free energy is what really distinguishes sentient beings capable of long-term planning from more elementary agents, such as viruses and bacteria, for which “living in the present” (and thus minimization of *current* free energy) is quite sufficient.

2 Complexity and accuracy of the generative model

Let us delve a little deeper into the issue of model complexity. Statistically, a model that matches too tightly the data that it tries to explain and recapitulate is suboptimal for inferential purposes. This situation, known as overfitting, may result from an increase in the complexity of the model because a larger number of parameters is necessary to accurately describe every minute detail of the data at hand, where a portion of the variance in the data may actually arise from random fluctuations or represent generally inessential features. Thus, an overly complex model is doubly disadvantageous: economically (metabolically, in the case of neural systems), because of its high maintenance costs, and performatively, because it generalizes poorly to situations that have the same core features of the previously learnt ones but are furnished with minor variations.

The active inference framework deals very efficiently with this issue because — as briefly mentioned above — the general imperative of free-energy (or prediction error) minimization can be mathematically recast in terms of the simultaneous drive to increase accuracy *and* reduce model complexity (see Smith et al. (2021) for an accessible introduction to the mathematical derivations). In other words, free-energy minimization is an expression of the principle known as Occam’s razor: to seek the most parsimonious explanation that represents with sufficient accuracy the data at hand. Interestingly, REM sleep dreaming has been proposed to have developed phylogenetically to enlist the inactive hours of the night to the task of reducing model complexity: by reenacting and recom-

binning the statistical regularities collected during the waking hours — which could have overly increased model complexity — their essential features can be summarized and consolidated by a process of synaptic pruning (Hobson and Friston, 2012).

On the other hand, when model simplification becomes excessive, it turns maladaptive. As a case in point, the progressive synaptic loss that we face as we age may also correspond to a decrease in complexity of our generative model. Beyond the general age-related deterioration of physiological processes, this is likely a consequence of our agential capacity: barring some drastic environmental changes outside the purview of our control, we are liable to get progressively better in fashioning our niche according to our expectations in a self-fulfilling prophesy cycle, thus reducing the complexity demands of our generative model. As older adults, we then become experts both at inferring our particular environment with high precision and at attaining our preferred states, but less capable of adapting to volatile situations: our generative model tends not only to suffer from overgeneralization (i.e., it struggles to encode novel contingencies with sufficient accuracy), but also to become refractory to revision. Children, on the contrary, are extremely adept at epistemic foraging and much less hindered by uncertainty — which makes perfect sense, as in their stage of life the generative model needs to accrue large amounts of information before being able to navigate the world without adult assistance and supervision.

3 A work in regress

Keeping the complexity of the generative model as low as possible, while maintaining a good accuracy, allows for an optimally adaptive strategy characterized not only by fitting responses to habitual situations, but also by adequate pliancy when facing the unexpected. While REM sleep has been argued to represent a general mechanism to preserve this balance across the animal kingdom, one may wonder whether the peculiar features of the human species — its highly complex agential capacity, the ability to array its environment with an ever expanding and dynamically evolving multitude of technical and cultural artifacts, the lengthening of the lifetime, and its increased meta-awareness — have also produced more specifically suited strategies to counteract the progressive loss of cognitive flexibility and adaptability that often comes with aging.

Interestingly, many spiritual traditions point to the necessity of an unlearning process to escape the limitations imposed by a well-established but ossified cognitive structure. Just to quote a few examples:

*In the beginner's mind there are many possibilities, but in the expert's mind there are few*⁴. (Shunryū Suzuki, Zen Mind, Beginner's Mind)

In pursuit of learning every day there is more. In the pursuit of the Dao every day there is less. (Laozi, Dao De Jing, Ch. 48)

⁴ To avoid a potential misunderstanding of this quote, here Suzuki was speaking in favor of recovering the openness of the beginner's mind, which we tend to lose as we become "expert".

Truly I tell you, unless you change and become like little children, you will never enter the kingdom of heaven. (Matthew 18:2–4, Holy Bible, NIV)

In a similar vein, a number of important artists, writers and philosophers have argued that the survival into adult life of the playfulness, openness, immediacy and sureness of judgement of infancy is an essential requisite for the creative enterprise, as epitomized by Pablo Picasso’s famous aphorism: “It took me four years to paint like Raphael, but a lifetime to paint like a child.”

What could the peculiar *praxis* enabling this kind of fertile backward step be, and what neurocognitive mechanisms could underpin it? To fruitfully approach the question from the point of view we have been espousing here, it is important to introduce a further notion, that of *precision*, and highlight the role it plays in the active inference scheme. In general terms, the outcome of the inferential process depends on the relative strength — i.e., precision of, or confidence in — predictions (beliefs) and prediction errors. A high-precision sensory prediction error coupled with a low-precision prediction makes the state of the generative model very sensitive to prediction error: credibility among competing predictions is swiftly redistributed in order to minimize prediction error. In the case of perception, this corresponds to making its content more attuned to the incoming sensory signal by selecting the most fitting hypothesis among the available ones about the cause of the current sensation. On the other hand, a low-precision sensory prediction error coupled with a high-precision prediction will secure the original prediction while the system works towards its fulfillment via appropriate actions: the agent will move and reorient its sensory surfaces so that collecting the very sensory input it is *expecting* becomes more likely. In short, prediction error is always precision-weighted and the variable deployment on precision across the hierarchy of the generative model is a key element in driving the dynamics of the system in perception, learning and behavior.

From this perspective, the *work in regress* outlined above would be aimed at relaxing the precision of ingrained beliefs — redistributing credibility more uniformly among the priors or smoothing out their probability density — to allow greater latitude in exploring alternative beliefs. From a different angle, this would amount to increasing the entropy of the system and thus the likelihood of selecting other, usually latent, a priori hypotheses.

Incidentally, such a mechanism has also been proposed to explain the effects of psychedelic drugs like psilocybin, LSD, DMT or MDMA (Carhart-Harris and Friston, 2019), which have recently shown promising long-term efficacy both in the treatment of resistant depression, obsessive-compulsive disorder and PTSD (Reiff et al., 2020; Carhart-Harris et al., 2021; Mitchell et al., 2021), and in changing the general outlook on life and expectations of terminally-ill patients towards a more accepting and positive stance (Griffiths et al., 2016; Ross et al., 2016). More specifically, the above conditions have been argued to be similarly characterized by an abnormally reduced level of neurocognitive entropy and by overly precise and rigidified priors that have become refractory to revision. Typically, the mentioned psychedelic drugs act as neuromodulators, i.e., they

change the magnitude of the effect that presynaptic inputs have on postsynaptic firing, a function that perfectly matches their hypothesized role of adjusting precision in active inference (Parr and Friston, 2017).

4 Contemplative practice as (in)active inference

In our opinion, active inference offers a distinct vantage point for understanding the mechanistic underpinnings of contemplative practice as the kind of *work in regress* discussed in the previous section, while also providing an appropriate language for developing a fruitful dialogue between meditative traditions and cognitive sciences (see also Farb et al., 2015; Pagnoni and Guareschi, 2017; Lutz et al., 2019; Pagnoni, 2019; Laukkonen and Slagter, 2021). Since the term ‘meditation’ can frequently refer to practices that are quite different in their formal aspects, underlying conceptual framing, and engaged mental processes (see Dahl et al., 2015, for a proposed taxonomy), we will limit our discussion here to the practice of ‘just sitting’ (Jap. *shikantaza*, also called *zazen*) as transmitted within the Japanese Zen Sōtō tradition⁵.

We propose that the practice of *zazen* can be seen as having two thoroughly intertwined aspects: (1) a pragmatic one, based on the priors installed by the meditative instructions, aiming to enact a suitable bodily posture and a mental state of non-judgmental wakefulness, and (2) an epistemic one, related to the cognizant faculty of realizing the spontaneous arising of mental content, the effects that the latter triggers in the body, and the instances of failure of the meditative prescriptions (e.g., episodes of absorption, dozing off, or slips of the posture), which is intimately tied to attentional regulation.

Paradoxically — but only so at a cursory examination — the apparent inaction featured by meditation is quite an active process (whence the pun “meditative in-action” in this article’s title), and a peculiar one at that. Indeed, *zazen* can be seen as the practice of skillfully orchestrating the suspension of all the usual policies associated with our tendency to move towards what we like and away from what we dislike⁶: a wedge, we could say, thrown into the mechanics of our fundamental, constitutive *appetite* for self-evidencing (Hohwy, 2014), all the while keeping an open attitude towards any spontaneously arising mental phenomenon. The legs are folded with the feet turned upside down in the lotus position⁷ — an embodied symbolic representation of the cessation of the intentionality of walking motion. The hands are placed with the fingers

⁵ The practice of *shikantaza* has some characteristics — the emphasis on the integration of body and mind in the sitting posture while generally forgoing common meditative techniques such as “labeling” of mental content, “decentering”, or breath counting — that license a specific treatment here. However, with suitable adjustments, the active inference framework can certainly be applied to other varieties of contemplative practice.

⁶ For a more detailed discussion, see Rev. Issho Fujita’s piece “My *zazen sankyu* notebook (6)” in the October 2000 issue of the Zen Sōtō School newsletter “Dharma Eye” (<https://www.sotozen.com/eng/dharma/pdf/07e.pdf>).

⁷ Alternative sitting postures are recommended for beginners and people who are unable to sit in the full lotus position.

of one hand resting on those of the other hand and the thumbs lightly touching in a horizontal line, a gesture (*hokkai-jōin*) which enacts and symbolizes refraining from both grasping and pushing away. Similarly, the eyes do not chase around — visual saccades are our chief means of environmental epistemic foraging — but are kept slightly downcast and open to eschew the seductive pull of dreamlike imagery that tends to manifest when closing our eyes and the associated drowsiness. The mouth is shut with the upper and lower teeth barely touching, with the tip of the tongue placed just behind the upper dental arch, temporarily shelving our speaking habits as well. Finally, an attitude of affective equanimity that discourages habitual value judgments about spontaneous thoughts and feelings extends the meditative poise into the cognitive and emotional realms. In summary, contemplative in-action can be viewed as a posture of *radical unavailability* to compulsive reactive patterns, to chasing after what we like (including ‘enlightenment’) and avoiding what we don’t like, well expressed by the Japanese term *mushotoku* (‘nothing to be attained’). In this sense, the embodiment of heedful presence in the sitting posture corresponds to a superordinate policy that effectively decreases the precision, and thus the compelling power, of ingrained policies with a high pragmatic value and a short-term resolution dynamic, by avoiding to follow them through with confirmatory (i.e., precision-boosting) actions. Quite literally, the meditator could then be viewed as performing a ‘sacrifice’ (from the Latin, *sacrum facere*: ‘sacred-making’), enacting the demarcation of a sacred, ritual space *within him/herself* rather than on the ground (Calasso, 1999), where the usual business of worldly cares and motivations is temporarily suspended. Note that this *praxis* is fundamentally different from an effortful training to control mind and body to reach a purported “enlightened” state based on personal beliefs and ideas, which is denoted in Japanese by the term *shuzen*⁸. In fact, the attainment of a suitable posture in *zazen* is more a matter of listening carefully (or perhaps, *caringly*) to how the body-mind ‘wants’ to settle according to its own natural dynamics, than following rigidly an externally-prescribed protocol. Dōgen Zenji, founder of the Zen Sōtō school in 13th century Japan, repeatedly warned that “*zazen* is not *shuzen*” and further clarified that “The *zazen* I speak of is not learning meditation. It is simply the Dharma gate of repose and bliss, the practice-realization of totally culminated enlightenment”(Waddell and Masao, 1973).

On account of the above considerations, we propose that under the virtual boundary staged by contemplative in-action⁹ the process of active inference effectively takes an inward-directed epistemic turn, providing a unique opportunity for the meditator’s generative model to learn *about itself*¹⁰ — “learn the

⁸ See Rev. Issho Fujita’s article “My footnotes on Zazen 1. Zazen is not Shuzen (1)” in the november 2011 issue of the Zen Sōtō School newsletter “Dharma Eye” (<https://www.sotozen.com/eng/dharma/pdf/28e.pdf>).

⁹ The construction of such a (ritual) boundary can perhaps be usefully specified in the framework of active inference as a temporary reduction in the degrees of freedom of the states of the active side of the Markov blanket of the meditator. For an introduction to the notion of Markov blankets, see Clark 2017.

¹⁰ Readers may see an apparent contradiction between the unlearning aspect of meditation and the notion of learning about oneself introduced here. In fact, learning and unlearning are

backward step that turns the light inwardly to illuminate your self”, we read again in Dōgen’s *Fukanzazengi* (Waddell and Masao, 1973). This learning of the generative model about itself is important because, as we previously mentioned, agents capable of a vast spectrum of actions like us must necessarily possess an accurate model of themselves to be able to accurately predict and infer the changes in their environment engendered by their own actions. Evidently, such self-knowledge or representation can only be partial, as the generative model would need to be computationally “bigger” than itself to possess a complete self-representation — an obvious logical impasse. Because of the specific features of the *praxis*, engaging body and mind as a whole, we speculate that the auto-epistemic outcome of *zazen* may concern especially the relationship between mental activity and its bodily consequences (eg, autonomic activity, postural changes, etc.). It may be useful to remark here that, while the notion that meditation practice helps us become more ‘selfless’ seems to have become popular lore these days, we believe it is also prone to be misunderstood. Kodo Sawaki, perhaps the most prominent teacher of the Zen Sōtō school of the twentieth century, used to say that “*zazen* is the self serving the self by the self” (Uchiyama and Okumura, 2014), begging the question of what is really this self that we refer to when we talk about selflessness.

From this perspective, we further argue that the mind-wandering episodes that inevitably punctuate a meditative session are not to be seen as nuisances that we hope will eventually disappear from our practice, but rather — when crucially combined with meditative body-mind presence — as vital opportunities for the generative model *that we are* to learn about itself. In fact, it is well known that the probability of mind-wandering increases progressively as the demands placed on us by the external world subside (Smallwood and Schooler, 2015). If the aim of meditation were simply to abolish mind-wandering, it would not make much sense to put ourselves in those very conditions that make it flourish, i.e., sitting ‘at rest’ in a quiet environment. On the contrary, we propose that meditative in-action be regarded as a skilled optimization of the ‘default’ auto-epistemic modality that characterizes spontaneous mindful rest¹¹, when the content of our stream of consciousness provides valuable opportunities for reflection, creative problem solving, and planning (Smallwood and Schooler, 2015).

What are the mechanisms by which such optimization could take place? Our working hypothesis is that the vigilant attitude prescribed and fostered in meditation¹², coupled with immobility, limits the attenuation of both the

usually interwoven processes because the acquisition of novel representations (or relationships between them) often requires the modification or unlearning of the old ones. Also, “learning about oneself” includes developing a higher degree of phenomenal opacity (Limanowski and Friston, 2018) of one’s own mental life, or meta-awareness, as discussed further below.

¹¹ The use of the term ‘default’ here is not casual: it points to the established association of mind-wandering with the activity of a specific set of brain regions that exhibit high metabolism at rest, known as ‘default mode network’ (Raichle et al., 2001; Smallwood et al., 2021).

¹² Maintaining this quality of wakefulness and liveliness in *zazen* is considered crucial: “like a dragon that reaches water, like a tiger entering the mountain”, as traditional expressions indicate. See, e.g., Rev. Issho Fujita’s article “*Shobogenzo Zazenshin* — A Free Translation

interoceptive signals associated with the bodily reverberations of spontaneous thoughts, and the meta-awareness of the latter. More specifically, not yielding to habitual reactive patterns triggered by spontaneous thoughts during meditation has the dual effect of (i) making them more available to consciousness and (ii) not reinforcing them, in turn redistributing precision more equitably among the existing priors and thus compelling the system to find alternative paths of free-energy minimization (e.g., letting the body/mind renormalize *endogenously* a state of increased arousal set off by anxiogenic thoughts, rather than via compensatory/compulsive overt actions). It is tempting to relate metaphorically the first aspect, the epistemic dimension of the meditative exercise, to the Homeric myth of Ulysses having himself tied to the ship's mast by his men in order to be able to hear the sirens' song without being drawn fatally to the rocks. The picture can be useful for reminding us that meditation is not a strategy to assuage anxiety by reducing uncertainty, but rather a way of constraining dynamics against habitual reactive patterns so that tolerance to uncertainty can be learned. This may increase — in time and outside of the meditative session — the likelihood of selecting less-frequently adopted policies whose expected reward (attainment of preferred states) is located further down in the future. In other words, meditative practice may also work as an antidote to the contemporary trend towards a steep delay discounting bias, i.e., the increasing preference for attaining immediate rewards and quickly avoiding uncomfortable situations¹³. It is in this way that the epistemically rich dynamic enacted by the very alternation of episodes of mind-wandering and meditative presence (Hasenkamp et al., 2012) can beget pragmatically valuable repercussions. In fact, mind-wandering has also been experimentally linked to reduced delay-discounting (Smallwood et al., 2012), supporting the hypothesis that opening up to the spontaneous occurrence of mental events during meditation, rather than trying to suppress it, may increase the probability of selection of policies with higher long-term pragmatic value. Ironically, the acritical apperception of a notion like 'living in the present moment', associated with meditation in the popular culture, may be one of the causes of the remarkable diffusion and success of mindfulness practices in the West in the last decade, by virtue of its alignment with the imperatives of a consumer-based capitalistic society (see, for instance, the concept of Buddhist modernism in Sharf, 2015; Thompson, 2020, and the associate critique of Purser and Milillo, 2014; Purser, 2019). Such cultural biases may easily generate perverse effects and, indeed, recent experimental evidence suggests that mindfulness practices can actually *decrease* prosocial behavior in individuals with pre-existing self-centered personality traits (Poulin et al., 2021). Moreover, it is not difficult to see how the capillary diffusion of digital connectivity and social networks, along its undeniable advantages, has also the potential to entangle us in cognitive and emotional processes of shorter and shorter expectation-confirmation arcs, making vulnerable individuals more

(6)" in the March 2010 issue of the Zen Sōto School newsletter "Dharma Eye" (<https://www.sotozen.com/eng/dharma/pdf/25eF.pdf>).

¹³ For reviews on delayed reward discounting, see Frederick et al., 2002; Kalenscher and Pennartz, 2008.

liable of being, so to say, held hostage by the present moment.

A few caveats are in order concerning the previous points. First, the meta-awareness fostered by *zazen* is not necessarily explicit or readily amenable to analysis and conceptualization; rather, it is more akin to a process of becoming familiar, intimate with the dynamics of the mind — its pushes and pulls, the peaks and troughs of its clarity, its bodily correlates, etc. — and thus is more of an intuitive nature. Second, it is conceivable that what is learnt in *zazen* goes beyond mere representational content, concerning more crucially being acquainted with the capacity for *mental action*. In the active inference framework, this is viewed as the ability to deploy different degrees of precision across the hierarchical web of message passing in the generative model — i.e., amplifying specific prediction error streams (especially endogenous ones, in this case) through neuromodulatory transmission, a mechanism that has been linked to the cognitive function of selective attention (Limanowski and Friston, 2018; Sandved-Smith et al., 2021). Finally, it should be pointed out that, according to traditional Zen teachings, the very compulsion to know is also to be relinquished in *zazen*, in the sense that knowing must be left free to appear, rather than actively pursued according to our own necessarily limited and limiting interests. This suggests again a condition where mental dynamics is somewhat less constrained by our prior beliefs than usual and more exposed — but from the safer vantage point of meditative in-action — to the epistemic vastness of the unexpected.

In closing, it is worthwhile mentioning that the approach to the scientific study of contemplative practices proposed here is still very much in the making. Thus, it may contain some internal inconsistencies or theoretical difficulties that further work should be able to elucidate. Computational modeling and experimental testing with brain imaging and behavioral paradigms will be crucial in verifying its viability and usefulness in future studies.

Acknowledgements

Special thanks to Antoine Lutz and the Lyon research group (Oussama Abdoun, Arnaud Pouban, Jérémie Mattout) for the ongoing collaboration on the topic of active inference and meditation, to Ryan Smith for providing valuable feedback on the manuscript, and to Leonida Gianfagna, Daniele Miracapillo, Giovanni Nubile and Tiziana Verde for lending their ear to the arguments discussed here and suggesting further exciting lines of speculative inquiry that, for reasons of space and time, will have to find their proper home in a future treatment. Also thanks to David Rudrauf, for his careful review and constructive critique of the first draft of the manuscript.

References

- Calasso, R. (1999). *Ka: Stories of the Mind and Gods of India*. VINTAGE.
- Carhart-Harris, R., Giribaldi, B., Watts, R., Baker-Jones, M., Murphy-Beiner, A., Murphy, R., Martell, J., Blemings, A., Erritzoe, D., and Nutt, D. J. (2021). Trial of psilocybin versus escitalopram for depression. *New England Journal of Medicine*, 384(15):1402–1411.
- Carhart-Harris, R. L. and Friston, K. J. (2019). Rebus and the anarchic brain: Toward a unified model of the brain action of psychedelics. *Pharmacological reviews*, 71:316–344.
- Clark, A. (2017). How to knit your own markov blanket:. In Metzinger, T. K. and Wiese, W., editors, *Philosophy and Predictive Processing*, chapter 3. MIND Group, Frankfurt am Main.
- Dahl, C. J., Lutz, A., and Davidson, R. J. (2015). Reconstructing and deconstructing the self: cognitive mechanisms in meditation practice. *Trends Cogn Sci*, 19(9):515–523.
- Di Paolo, E., Thompson, E., and Beer, R. D. (2021). Laying down a forking path: Incompatibilities between enaction and the free energy principle.
- Farb, N., Daubenmier, J., Price, C. J., Gard, T., Kerr, C., Dunn, B. D., Klein, A. C., Paulus, M. P., and Mehling, W. E. (2015). Interoception, contemplative practice, and health. *Frontiers in psychology*, 6:763.
- Frederick, S., Loewenstein, G., and O’donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40(2):351–401.
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nat Rev Neurosci*, 11(2):127–138.
- Friston, K. (2018). Am i self-conscious? (or does self-organization entail self-consciousness?). *Frontiers in psychology*, 9:579.
- Friston, K. J. (2019). A free energy principle for a particular physics. *arXiv: Neurons and Cognition*.
- Griffiths, R. R., Johnson, M. W., Carducci, M. A., Umbricht, A., Richards, W. A., Richards, B. D., Cosimano, M. P., and Klinedinst, M. A. (2016). Psilocybin produces substantial and sustained decreases in depression and anxiety in patients with life-threatening cancer: A randomized double-blind trial. *Journal of Psychopharmacology*, 30(12):1181–1197.
- Hasenkamp, W., Wilson-Mendenhall, C. D., Duncan, E., and Barsalou, L. W. (2012). Mind wandering and attention during focused meditation: A fine-grained temporal analysis of fluctuating cognitive states. *Neuroimage*, 59(1):750–760.

- Hobson, J. A. and Friston, K. J. (2012). Waking and dreaming consciousness: neurobiological and functional considerations. *Prog Neurobiol*, 98(1):82–98.
- Hohwy, J. (2014). The self-evidencing brain. *Noûs*, 50(2):259–285.
- Kalenscher, T. and Pennartz, C. M. (2008). Is a bird in the hand worth two in the future? the neuroeconomics of intertemporal decision-making. *Progress in Neurobiology*, 84(3):284–315.
- Laukkonen, R. E. and Slagter, H. A. (2021). From many to (n)one: Meditation and the plasticity of the predictive mind. *Neuroscience & Biobehavioral Reviews*, 128:199–217.
- Limanowski, J. and Friston, K. (2018). ‘seeing the dark’: Grounding phenomenal transparency and opacity in precision estimation for active inference. *Frontiers in psychology*, 9:643.
- Lutz, A., Mattout, J., and Pagnoni, G. (2019). The epistemic and pragmatic value of non-action: a predictive coding perspective on meditation. *Current opinion in psychology*, 28:166–171.
- Maturana, H. R. and Varela, F. J. (1980). *Autopoiesis and cognition: The realization of the living*, volume 42 of *Boston Studies in the Philosophy of Science*. Reidel, Dordrecht.
- Mitchell, J. M., Bogenschutz, M., Lilienstein, A., Harrison, C., Kleiman, S., Parker-Guilbert, K., G., M. O., Garas, W., Paleos, C., Gorman, I., Nicholas, C., Mithoefer, M., Carlin, S., Poulter, B., Mithoefer, A., Quevedo, S., Wells, G., Klaire, S. S., van der Kolk, B., Tzarfaty, K., Amiaz, R., Worthy, R., Shannon, S., Woolley, J. D., Marta, C., Gelfand, Y., Hapke, E., Amar, S., Wallach, Y., Brown, R., Hamilton, S., Wang, J. B., Coker, A., Matthews, R., de Boer, A., Yazar-Klosinski, B., Emerson, A., and Doblin, R. (2021). MDMA-assisted therapy for severe PTSD: a randomized, double-blind, placebo-controlled phase 3 study. *Nature Medicine*, 27(6):1025–1033.
- Pagnoni, G. (2019). The contemplative exercise through the lenses of predictive processing: A promising approach. *Progress in brain research*, 244:299–322.
- Pagnoni, G. and Guareschi, F. T. (2017). Remembrance of things to come: a conversation between zen and neuroscience on the predictive nature of the mind. *Mindfulness*, 8(1):27–37.
- Parr, T. and Friston, K. J. (2017). Uncertainty, epistemics and active inference. *Journal of the Royal Society, Interface*, 14.
- Poulin, M., Ministero, L., Gabriel, S., Morrison, C., and Naidu, E. (2021). Minding your own business? mindfulness decreases prosocial behavior for those with independent self-construals.
- Purser, R. (2019). *McMindfulness*. Watkins Media.

- Purser, R. E. and Milillo, J. (2014). Mindfulness revisited. *Journal of Management Inquiry*, 24(1):3–24.
- Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., and Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences of the United States of America*, 98:676–682.
- Reiff, C. M., Richman, E. E., Nemeroff, C. B., Carpenter, L. L., Widge, A. S., Rodriguez, C. I., Kalin, N. H., and and, W. M. M. (2020). Psychedelics and psychedelic-assisted psychotherapy. *American Journal of Psychiatry*, 177(5):391–410.
- Ross, S., Bossis, A., Guss, J., Agin-Liebes, G., Malone, T., Cohen, B., Mennenga, S. E., Belser, A., Kalliontzi, K., Babb, J., Su, Z., Corby, P., and Schmidt, B. L. (2016). Rapid and sustained symptom reduction following psilocybin treatment for anxiety and depression in patients with life-threatening cancer: a randomized controlled trial. *Journal of Psychopharmacology*, 30(12):1165–1180.
- Rubin, S., Parr, T., Costa, L. D., and Friston, K. (2020). Future climates: Markov blankets and active inference in the biosphere. *Journal of The Royal Society Interface*, 17(172):20200503.
- Sandved-Smith, L., Hesp, C., Mattout, J., Friston, K., Lutz, A., and Ramstead, M. J. D. (2021). Towards a computational phenomenology of mental action: modelling meta-awareness and attentional control with deep parametric active inference. *Neuroscience of consciousness*, 2021:niab018.
- Sharf, R. H. (2015). Is mindfulness buddhist?(and why it matters). *Transcultural psychiatry*, 52(4):470–484.
- Smallwood, J., Ruby, F. J. M., and Singer, T. (2012). Letting go of the present: Mind-wandering is associated with reduced delay discounting. *Conscious Cogn*, 22(1):1–7.
- Smallwood, J. and Schooler, J. W. (2015). The science of mind wandering: empirically navigating the stream of consciousness. *Annual review of psychology*, 66:487–518.
- Smallwood, J., Turnbull, A., ting Wang, H., Ho, N. S., Poerio, G. L., Karapanagiotidis, T., Konu, D., Mckeown, B., Zhang, M., Murphy, C., Vatansever, D., Bzdok, D., Konishi, M., Leech, R., Seli, P., Schooler, J. W., Bernhardt, B., Margulies, D. S., and Jefferies, E. (2021). The neural correlates of ongoing conscious thought. *iScience*, 24(3):102132.
- Smith, P. (2017). *Other minds : the octopus, the sea and the deep origins of consciousness*. William Collins, London.

- Smith, R., Friston, K., and Whyte, C. (2021). A step-by-step tutorial on active inference and its application to empirical data.
- Thompson, E. (2020). *Why I Am Not a Buddhist*. Yale University Press.
- Uchiyama, K. and Okumura, S. (2014). *The Zen Teaching of Homeless Kodo*. Wisdom Publications, U.S.
- Varela, F. G., Maturana, H. R., and Uribe, R. (1974). Autopoiesis: The organization of living systems, its characterization and a model. *Biosystems*, 5(4):187–196.
- Varela, F. J., Thompson, E., and Rosch, E. (1991). *The embodied mind: cognitive science and human experience*. MIT Press, Cambridge, Mass.
- Waddell, N. and Masao, A. (1973). Dōgen’s: “fukanzazengi” and “shōbōgenzō zazengi”. *The Eastern Buddhist*, 6(2):115–128.