This is the author's copy of the ACCEPTED MANUSCRIPT. This article should be cited as "Perrykkad, K., & Hohwy, J. (2020). Fidgeting as self-evidencing: A predictive processing account of non-goal-directed action. *New Ideas in Psychology, 56,* 100750. doi:https://doi.org/10.1016/j.newideapsych.2019.100750".
Please see published version available at https://doi.org/10.1016/j.newideapsych.2019.100750

Fidgeting as Self-Evidencing:
A predictive processing account of non-goal-directed action

Kelsey Perrykkad[a] and Jakob Hohwy[a]

[a] Cognition and Philosophy Lab, Philosophy Department, School of Philosophical Historical and International Studies, Monash University, Victoria, Australia

Corresponding Author:
Kelsey Perrykkad
kelsey.perrykkad@monash.edu
20 Chancellors Walk,
Monash University, Clayton, VIC
Australia 3800

Abstract

Non-goal-directed actions have been relatively neglected in cognitive science, but are ubiquitous and related to important cognitive functions. Fidgeting is seemingly one subtype of non-goal-directed action which is ripe for a functional account. What's the point of fidgeting? The predictive processing framework is a parsimonious account of brain function which says the brain aims to minimise the difference between expected and actual states of the world and itself, that is, minimise prediction error. This framework situates action selection in terms of active inference for expected states. However, seemingly aimless, idle actions, such as fidgeting, are a challenge to such theories. When our actions are not obviously goal-achieving, how can a predictive processing framework explain why we regularly do them anyway? Here, we argue that in a predictive processing framework, evidence for the agent's own existence is consolidated by self-stimulation or fidgeting. Endogenous, repetitive actions reduce uncertainty about the system's own states, and thus help continuously maintain expected rates of prediction error minimization. We extend this explanation to clinically distinctive self-stimulation, such as in Autism Spectrum Conditions, in which effective strategies for self-evidencing may be different to the neurotypical case.

Keywords: *fidgeting; self-evidencing; non-goal-directed action; prediction error minimization; autism spectrum conditions*

1. Introduction: Fidgeting and Non-goal Directed Actions

In investigating the psychological and neuroscientific bases of movement, cognitive science has primarily focused on goal-directed action. These actions can be generally defined as those whose function is to achieve some instrumental aim for the individual. For example, reaching for a glass of water to quench one's thirst. However, there is a relatively neglected category of action that has not been well explained under cognitive theories of action. This is non-goal-directed action, which incorporates behaviours such as fidgeting. This common distinction between *goal-directed action* and *non-goal-directed action* is often subtly made, without definition, but rather by attributing a certain cognitive mechanism to goal-directed action specifically. Since goal-directed action is defined in opposition to non-goal-directed action, the latter cannot be about achieving a personal-level end. Other than this negative definition, it is not well explicated (Dayan, 2009). Sometimes, this contrast class is labelled *habits* or *reflexes* (Barfield et al., 2017; Friston et al., 2016).

We label our target class of non-goal-directed actions in this paper, *fidgeting.* Fidgeting is defined partly in opposition to other action types, such as habits. Habits colloquially include automatically performed sequences of productive action (for example brushing your teeth before bed), whereas, superficially at least, fidgeting appears to have no epistemic or pragmatic value. Examples of fidgeting thus excludes brushing one's teeth or always walking the dog along a certain route but includes tapping the table with one's finger or twirling one's hair. The primary outcome of fidgeting, as we define it here, can often be described as reflexive stimulation; caused by the individual, to the individual. Importantly, stimulation of the kind intended here is usually repetitive or patterned and is both self-initiated and self-sustained. Agents may be unaware or aware that they are fidgeting. Typically, agents have not consciously decided to fidget but fidgeting behaviours can be intentionally (and consciously) terminated, resisted or permitted. Note that fidgeting is still distinct from reflexive, goal directed stimulation (such as scratching an itch), because there is no obvious person-level goal that is achieved through the action. Stimulation can occur in multiple sensory domains and is often primarily visual, tactile or auditory (see Table 1 for examples of the target behaviour in different modalities). It can be performed with the body alone, or with a prop, such as a piece of string, a pen and paper or a light source.

There is little consensus among the scientific community about the reason we fidget. There is a possibility that it serves no purpose. Some philosophers use it as the paradigm case of sub-intentional action or mere movement, which, by definition, serves no end for the agent (Hommel, 2015; O'shaughnessy, 1980). However, upon further examination there is growing evidence that fidgeting is importantly and systematically involved in many psychological processes, which suggests a deeper story. Fidgeting is commonly thought to be indicative of a lack of attention, and is increased in individuals with Attention Deficit Hyperactive Disorders (Lis et al., 2010). In the context of lecture based learning, fidgeting was shown to predict recall independently of attention, with increased fidgeting associated with decreased memory (Farley, Risko, & Kingstone, 2013). However, unlike other tasks performed alongside a primary task, doodling has been shown to aid incidental memory (Andrade, 2010). Exaggerated and more frequent fidgeting behaviours are also implicated in psychiatric and neurological conditions such as the stereotypies found in Autism Spectrum Conditions, Rett syndrome, schizophrenia and people who are blind (Barry, Baird, Lascelles, Bunton, & Hedderly, 2011; Morrens, Hulstijn, Lewi, De Hert, & Sabbe, 2006). Together, these findings suggest that fidgeting is not purposeless and, accordingly, several functional accounts of fidgeting have been proposed.

One prominent theory regards fidgeting as *bodily regulation*, such that we fidget to release excess stored energy or reduce fat mass gain (Johannsen & Ravussin, 2008). However, though spontaneous physical activity (including but not limited to fidgeting) is inversely correlated with obesity and was associated with changes in energy expenditure, perturbations in physical activity or weight do not change the amount of spontaneous physical activity (Johannsen & Ravussin, 2008) so the explanation cannot be this simple. Many of the studies under this theory consider fidgeting to be just one of many similar, non-exercise forms of energy expenditure, called non-exercise activity thermogenesis (NEAT), and includes walking and standing along with fidgeting. Thus, the explanandum of these proposals is broader than the current project.

In certain clinical cases in particular, fidgeting may help the body regulate its own function. Fidgeting in patients with autonomic failures has been shown to counteract symptomatic drops in blood pressure (Cheshire, 2000), and may also prevent endothelial dysfunction by improving blood flow during prolonged sitting (Morishima et al., 2016). Perhaps then it is natural to think that the usual instance of fidgeting is associated with autonomic self-regulation where movements would always increase bloodflow and blood pressure. However, fidgeting often occurs in situations of high stress, where the self-regulatory autonomic goal should be to decrease these cardiac parameters (assuming one is regulating automatic fight or flight responses). Further, these purely homeostatic and bodily explanations do not explain the cognitive effects associated with fidgeting, or why stereotypies occur in psychiatric conditions without clear autonomic dysfunction. As such, these explanations do not capture the whole story. While it should be noted that there may be cases of fidgeting that are restricted to specific mechanistic failures of the body (as, perhaps, in patients with specific autonomic failures), we suggest that an explanation of fidgeting in the general case should account for both autonomic regulation and cognitive elements of the phenomenon. In this paper, we shall appeal to active inference as the best framework for explaining fidgeting, and this notion subsumes homeostatic self-regulation in a broader (allostatic) framework, addressing these misgivings about the homeostatic account of fidgeting specifically.

Another theory regards fidgeting as *cognitive regulation*. As we briefly reviewed earlier, evidence shows that fidgeting is associated with cognitive states such as attention and memory, and varied psychiatric conditions. If fidgeting is associated with so many cognitive functions, it is plausible to think the best explanation will be a cognitive or neurological one. One theory is that fidgeting is bodily *mind wandering* because the tendency to fidget and mind wander are correlated between individuals (Carriere, Seli, & Smilek, 2013). However, this fails to illuminate either notion. Further, the usual approach of correlating fidgeting behaviour with contextual differences fails to explain which features are significant to these contexts, and ignores explanatorily fruitful individual differences in behaviour. A comprehensive account of fidgeting behaviour should specify both why and when we do it. More empirical research on the causal relations between fidgeting and cognitive processes such as attention, memory and mind-wandering is needed.

Lovaas, Newsom, and Hickman (1987), argue that self-stimulatory activity occurs as a form of operant conditioning, in which the reinforcer is the resulting perceptual stimuli. This is the most similar theory to the one presented in this paper, however, it does not account for *why* self-stimulation and particular perceptual consequences should be rewarding, which our proposal aims to explain. In general, an account that relies on reinforcement learning would be attractive, because it integrates fidgeting with other forms of action.

Our explanation will rather employ the *active inference* account of action from the predictive processing framework. Briefly, this account situates action as a form of prediction error minimisation. It is a contemporary alternative to reinforcement learning that would also unify the account of fidgeting with an account of action in general. Active inference arguably has computational advantages over reinforcement learning explanations because the former does not rely on an independent value system (Friston, Daunizeau, & Kiebel, 2009; Friston, FitzGerald, Rigoli, Schwartenbeck, & Pezzulo, 2017). Further, since reinforcement learning requires an action (or series of actions) to be fully completed and yield an outcome before one can learn from them, it raises an issue in explaining how and why we might choose to switch to fidgeting in the middle of another task (as the initial policy is not yet completed) (see also Friston (2017b)). The active inference account allows learning at all stages of the policy execution, and so the experience of changing to fidgeting mid-task is more amenable to our account.

The problem for this strategy is, however, that fidgeting does not obviously seem to specify any expected low prediction error state. Conversely put, fidgeting becomes a challenge for those who believe active inference is a unifying account of action. This sets the challenge for this paper, and we will argue that active inference does in fact have the resources to accommodate fidgeting, in a way that interestingly reveals fidgeting's function and throws light on the nature of self-representation.

2. Predictive Processing Framework

In this section, we will introduce the specific elements of the predictive processing framework relevant to fidgeting behaviour including active inference. We will explain why fidgeting poses a challenge to this framework, and will lay the groundwork for our solution.

The predictive processing account of brain function asserts that the brain hierarchically compares incoming information from the senses with descending prior expectations. Where these signals mismatch, a prediction error is propagated up the neural hierarchy. The brain's overall function is to minimize this error by making contextually appropriate adjustments of its internal expectations and thus creating a dynamic model of the hidden causes of sensory input (both environmental and bodily). This theoretical framework parsimoniously explains both perception and action as methods of prediction error minimisation (Clark, 2015; Friston, 2010; Hohwy, 2013).

2.1 Policy selection.

Each individual's particular internal model of the world influences their perception and action selection enabling prediction error minimisation over their lifespan. A *policy* is a set of possible actions (or individual control states) that have been grouped together by the individual for its history of success as a strategy to reduce prediction error when faced with situations with learned commonalities, which cue success for that policy.

Policies are considered under uncertainty, in the sense that each policy has a certain probability that it will lead to a desired outcome, defined in terms of the expected sensory outcomes. Since there is not a one-to-one relation between actions and outcomes, policy selection presents an inference problem, hence the label *active inference.* This implies that, in some circumstances, agents can misrepresent causal relations between actions and outcomes when there are none. For example, Skinner (1948) describes pigeons who learn to perform "superstitious" actions for rewards given at random intervals. Policies influence the dynamics of states of the world, each of which is associated with a likelihood whose *precision* describes the fidelity of the mapping from those states to their sensory consequences. Many extremely precise policies are performed automatically without any conscious awareness, such as the movements associated with walking for an ambulatory adult. The precision of a likelihood mapping is inversely related to its *ambiguity*, where high ambiguity means it is unclear which sensory outcome should be predicted. As such, alternative policies are associated with different levels of ambiguity in virtue of the alternative states (and therefore likelihoods) whose dynamics they prescribe. Active inference is therefore sensitive to precision and ambiguity.

Active inference can minimize prediction error in two ways, for epistemic value or for utility. Epistemic value refers to actions that implicate causally intervening on modelled causal processes (e.g., shaking a present to find out what is inside) and can provide insight into the accuracy of that model in the same way as experimenting in science can provide insight into causal processes. Policies can then be inferred based on their expected epistemic utility, in the set of actions expected to maximise model updating (i.e., how much the existing model improves upon acting), and thereby minimise uncertainty (Friston et al., 2016; Friston et al., 2015). This reveals a prior expectation of agents, namely that they expect to act to occupy states where they minimize uncertainty. Action in this framework can also just be about changing the world such that it conforms with the agent's expected states, and thereby reduces prediction error (e.g., opening a present to obtain the reward inside) (Friston et al., 2015). As such, polices are *selected* by being estimated to be the most probable course of action to induce state transitions which minimize prediction error in one of these two ways. Notice that, for both types of action, the inferred policies provide information to the agent about what kind of agent they are; policies reveal how an agent is likely to act in different contexts, and thereby help the agent infer their own traits and develop a sense of identity and selfhood.

Active inference concretely relates to prediction error minimization in the sense that the inferred policy predicts sensory input that is not actually occurring, which induces prediction error that will be minimized by moving the body. That is, bodily movement arises in the minimization of prediction error generated from an inferred policy that specified some expected sensory input (servicing either epistemic value or utility). Under the predictive processing framework, this is the *only* explanation for action, and is proposed as an alternative to traditional motor command accounts (Adams, Shipp, & Friston, 2013).

All this implies that fidgeting arises as an inferred policy. However, under the active inference account, it is a challenge to understand why our inferred policy would be to fidget given that it doesn't seem likely to reduce prediction error in any straightforward, meaningful way.

2.2 Volatility.

Agents subsist in the actual, noisy, changeable and uncertain world, which is characterized by dynamic changes in sensory input as causal factors at various spatiotemporal levels (including causes associated with the agent's own actions) interact with each other creating non-linear and unexpected fluctuations in the sensory input. Overall, this means that the expected uncertainty changes dependent on context. Under a predictive processing framework, human brains represent such change as an estimate of *volatility* – the expected variability (or variance) in the dynamics of the states causing sensory input (i.e., changes to the standard deviation across contexts).

Volatility reflects an expectation about the rate at which our models change in perceptual inference, and how much they can change when acting for epistemic value. This expectation can be implicitly represented in the brain in terms of the overall uncertainty of what is being inferred together with the precision (or fidelity) with which states or actions generate outcomes. This would provide a dynamic estimate of how much information would be gained by a certain action. Intuitively, the distance with which the agent's beliefs move following an observation is the information gain, and this reduction in uncertainty corresponds to the degree to which prediction errors are minimized.

Some types of agents might have evolved explicit representation of this rate of prediction error minimization, represented in the brain as a hyperprior – a deep hierarchical prior about the dynamic nature of the agent's own model. Technically, the rate at which a prediction error is minimised is partially determined by the precision (or confidence) ascribed to that prediction. Intuitively, we update our beliefs faster and more dramatically when we believe our data to be more reliable and when we suspect the world is frequently liable to change (for a more formal account of the dynamic effect of variability and volatility on learning rate, see Mathys, Daunizeau, Friston, and Stephan (2011)). This means that those creatures whose internal models allow them to make predictions about this precision implicitly hold (sub-personal) beliefs about the expected rate of error-minimisation, conditioned on the data they choose to sample. Given the inverse relationship between precision and ambiguity, we can associate beliefs about the ambiguity expected under a given policy with beliefs about the expected rate of uncertainty (or error) minimisation.

Expectations for volatility and rate of prediction error minimization will impact on policy selection. Policies that will give the expected prediction error minimisation should be selected but this inference is sensitive to beliefs about the extent to which there will be underlying change in the world during execution of the policy. Actions which are effective in one context will likely be rendered less optimal when the statistics of the environment change, as various hidden causes both enter and leave the causal chain.

2.3 Hypothesis decay.

The expectation for change and volatility will exert pressure on the current best evidenced hypothesis about the causes of sensory input, decreasing its strength with time. The longer one sticks to a winning hypothesis, in the presence of higher-level information about causal interactions, the less likely it is that this hypothesis will remain efficient at minimising prediction error (Hohwy, Paton, & Palmer, 2016, p. 320). In short, the strength of evidence for a particular hypothesis decays over time, contingent on the inferred volatility. A given prior, which was efficient to guide inference at some time may soon be less efficient as hidden causes in the environment change. This carries over to policy selection via the fidelity of the mapping from states and actions to outcomes, which is sensitive to underlying change in the causal landscape – as time passes the agent must begin to anticipate that ambiguity will increase, that is, that uncertainty will increase and rate of prediction error minimization will decrease. As the mapping between actions and outcomes is anticipated to change under volatility and thus becomes more ambiguous with time, the precision of the estimation for the belief that "I can effectively minimise prediction error" decreases. This leads us to our discussion of self-evidencing.

2.4 Self-evidencing and self-models.

Biological systems, under predictive processing, are embodied models of the statistical regularities in the world which resist entropy by occupying a limited number of possible homeostatic states, relative to which they minimise prediction error by active and perceptual inference (Friston, 2013; Friston & Ao, 2012). Minimising prediction error under a model is equivalent to maximising the evidence for that model. If the model is embodied as the agent, then the reduction in prediction error becomes evidence for the existence of the agent; this is called *self-evidencing* (Friston, 2010, 2013; Hohwy, 2016). This means that an agent will model its own existence implicitly in terms of a belief that it is self-evidencing. Meeting an expected rate of prediction error minimisation (or information gain) in perceptual and active inference essentially reassures the agent of its own existence.

In modelling the causes of sensory input, one relatively stable cause of changes to the sensed environment is oneself. As the agent's own body, via its actions, causes endogenously generated changes to sensory information, the model that best explains all of the agent's sensory experiences will include a model of the agent itself (Apps & Tsakiris, 2014); this is the *self-model*.

The self-model is hierarchical, spanning basic body parameters (limb size and reach), occurrent control states (beliefs about how desire for sip of coffee leads to reaching movements), habits (having coffee every morning), and character traits (introvert coffee connoisseur). Interactions among the self-models' elements help minimize prediction error caused by the agent's own actions (Hohwy & Michael, 2017).

The agent represents itself as a (complex) cause that is cyclically perturbed as the system interacts with the world and deals with the uncertainty inherent in spatially or temporally changing environments. This representation then is also implicated in policy selection, as the scope of possible intervention by the agent and the threshold for irreducible uncertainty is inferred through this model.

The self-model represents the internal states of the agent from which action outcomes are generated. Therefore, what the self-model represents is centrally implicated in the fidelity of the mapping from actions to outcomes, and thus for the expected reduction of uncertainty or rate of prediction error minimization. This creates a conceptual link between the self-model ("what kind of cause in the world am I?") to existence ("are my actions self-evidencing?"). This link is captured in the ambiguity of the action-outcome mapping – a highly ambiguous mapping predicts poor self-evidencing. As already briefly discussed, some agents may be able to represent the rate of self-evidencing explicitly, as a belief capturing the expected rate of prediction error minimization. Further, sensitivity to dynamic changes in the rate of prediction error minimisation may function in a similar way to explicit estimations of volatility, improving stability following volatile changes in statistical contexts and allowing for quicker and more differentiated changes in response to volatility (Joffily & Coricelli, 2013).

In this section, we have reviewed relatively subtle aspects of the predictive processing framework, which we will recruit in order to meet the challenge of explaining why fidgeting policies should be inferred.

3. Fidgeting as a Solution to Self-Model Decay

Since the representation of self is central to the agent's model of the world, it too is sensitive to dynamic estimates of environmental volatility. The hypotheses generated from the self-model will then also have decaying evidence given the right (or wrong) circumstances. Self-hypothesis decay implies that the self-model is deprived of evidence, undermining the agent's confidence that it is minimizing error at the expected rate – implicitly, its confidence that it is self-evidencing and will continue to exist. When there is hypothesis decay (contingent on expectations for volatility) about the very mapping between our actions and their consequences we will increasingly violate our expectation for the rate at which we can minimise prediction error. In that case, our current policy for interacting with the world no longer accurately predicts how much or how little prediction error it can minimise. This raises the question of what we should do when there are unexpected rates of prediction error minimization, that is, when belief in self-evidencing is threatened.

In practice, it is likely that the actual rate of prediction error minimisation deviates from expected rates of prediction error minimisation (for the given context) during the extremes of psychological arousal. When the agent is bored, the available policies are not expected to have much epistemic value, even though the ambiguity is low – the model is already largely optimised for the environment. This can be interpreted as the model performing better than expected in terms of achieving states with high utility, eliminating prediction error faster than expected in the context, in spite of relative sparsity of sensory perturbations and enacted policies. Consider boredom during on-task behaviour such as when reading dull administrative emails. In this situation, the individual may have an expectation that opening a new email should produce some level of prediction error (i.e., verify a model with somewhat ambiguous mapping between action and sensory effect), however, on finding that they perfectly predict the content, they will be bored – reducing prediction error faster than expected based on the modelled ambiguity. The modelled ambiguity has failed to adhere with the expected result (greater prediction error) of the most precise policy. Note that in overwhelming and overly complex situations one might also become bored, such as in a very complicated lecture. In these cases, the most precise policy (which is often to disengage from the inordinate task) reduces the total set of presumed model-able hidden causes, making the bound on estimated irreducible prediction error higher, but enabling the rate of prediction error minimisation to be more precisely estimated.

At the other extreme, when the agent is stressed, the model reduces uncertainty at a slower rate than expected, leaving the system with much unaccounted-for prediction error (see Peters, McEwen, and Friston (2017) for an account of stress in terms of irreducible uncertainty). In this scenario, the administrative email which was expected to have a small range of content may have contained disturbing news requiring a whole new set of policies and goal states for the day with increased complexity and ambiguity.

These deviations in expected rate of prediction error minimisation could come about for different reasons, such as in the case of boredom, when the world may be less complex (in terms of quantity of interacting hidden causes), less volatile and/or the agent is more efficacious than expected. In the case of stress, the world may be more complex, more volatile and the agent may be have less effective causal powers than expected (for reasons such as tiredness). It is in these situations that the self-model is most susceptible to the processes of hypothesis decay. They represent instances where the fidelity of the action-outcome mapping deteriorates in different ways. This is because there is increasing ambiguity in policies learned over a long history of perception and action which have reciprocally constructed the inferred self-model. If your actions no longer do what they have always done, then this will motivate the inference that your self-model is itself subject to volatility. Because your understanding of what kind of agent you are in the world is based on beliefs about how you act in general and how those actions should affect prediction error minimisation, you begin to lose evidence that you exist (at least as the same kind of agent that you always have been).

This is the context in which, we argue, we should understand fidgeting. The problem, described above, was to explain why we might infer a fidgeting policy when such actions seem to provide no utility – agents obtain no obvious utility for fidgeting. The notion of active inference also allows action for epistemic value but it is also difficult to see how fidgeting could provide epistemic value, since it is not obviously a matter of exploring the world to reduce uncertainty about the model. These actions are precisely ones that the agent has performed many times before with unchanging results (compare eye movement to look more closely at something to determine what it is – a clear case of acting for epistemic value). Nevertheless, we propose that, in a setting where there is self-hypothesis decay, fidgeting can in fact be conceived as action for epistemic value.

The idea is that fidgeting reinforces the fidelity of the action-outcome mapping and thereby solidifies the agent's belief that it is efficiently self-evidencing. Our actions more or less reliably have certain expected consequences. This reliability is the precision of our inferred policy, which determines how efficiently sensory prediction errors are minimised. At the most basic level, even aimless moving will reliably increase the signal in proprioceptive inputs relative to the noise. So fidgeting should always reduce ambiguity, that is, speeding up the error minimisation process in relation to, for this particular example, the skeletomotor system. Perhaps, because this will always the case, this is the type of behaviour we default to if other behaviours cannot reduce uncertainty (either because there is no other relatively informative policy available in the context (i.e. boredom), or because the environment is so ambiguous that efficient policy selection is impossible (i.e. stress)). In this way, fidgeting brings the rate of prediction error minimisation back to what we expect it to be, while also changing the expected sensory input to include this repetitive self-stimulation. The proposal is thus that fidgeting is action for epistemic value in the foundational sense of self-evidencing. In other words, when I can't resolve uncertainty about anything else, I can still resolve uncertainty about myself.

The evidence provided by fidgeting is particularly strong and precise because fidgeting policies are relatively simple, involving few interacting hidden causes between the movement and the sensory effects of the movement. In this way, fidgeting polices are robust to volatile changes in the external environment. In the clearest examples, such as bouncing one's leg or playing with one's hair, the only hidden causes are part of the agent's own body. When other causes are involved, they are reliable and familiar, such as the sounds generated from clicking a pen. The particular ways an individual chooses to fidget will depend on their learned history about what fidgeting policies are most precise in each context, which will also be individually delimited based on learned cues to relevant similarities across environments.

The proposal captures several important features of fidgeting. Fidgeting is often self-stimulatory and reflexive, reflected in the tight causal loop, and involvement of few hidden causes, with good robustness across environmental contexts. The stimulation is patterned because this increases the predictability within the current context and provides temporal rhythm that also increases the dimensions across which the stimuli are predictable. As a precise, expertly completed policy, fidgeting can be performed consciously or unconsciously but is usually initiated non-deliberatively because it does not fulfil person-level goals, but rather is an automatic way to reaffirm the existence of the entity to which person-level goals are attributed. The sensory modality over which fidgeting is completed does not matter for its self-evidencing role, so policies may span these domains within an individual (Table 1). It can involve external objects, such as clicking a pen,visually tracking a rotating fan or rocking on a chair; or be limited to touching or moving parts of ones' own body, for example bouncing one's leg, playing with one's hair or biting one's nails.

In some ways, this may seem similar to the bodily-regulation accounts of fidgeting reviewed earlier, in that fidgeting policies are a way to stay within a relatively small range of expected states. However, our proposal differs for two primary reasons. First, the states that cause the deviation from expected states need not be primarily bodily (homeostatic) states, but are related to cognitive states. So the prediction error being eliminated by fidgeting is not necessarily directly to do with states like heart rate and blood pressure. Second, the account allows for fidgeting in *anticipation* of prediction error, inferred from a deviation from the expected trend. In this way, where a classic homeostatic regulation account would place fidgeting as a response to prediction error, our account is more allostatic, in that fidgeting may occur before being in the unexpected bodily state (Corcoran & Hohwy, 2019). The prediction error it addresses is related to unexpected changes in the modelled transitions between states rather than the result of being in a particular state.

Even though fidget policies are very reliable and provide strong self-evidencing when executed, (neurotypical) agents do not over-indulge in fidgeting. This is because in the real, volatile world, a fidget-only strategy would accumulate prediction error in the long run, that is, there is hypothesis-decay even for the hypotheses leading to fidgeting. Concretely, feeding, exploring, and paying the bills are precluded by fidgeting, meaning that the agent eventually strays from its expected states. Self-evidencing agents will therefore dynamically recruit from a varied repertoire of policies. We have argued that fidgeting belongs in this repertoire.

The extent to which fidgeting policies should be inferred raises the question of maladaptive fidgeting, especially in psychiatric conditions. It is far from a trivial inference problem to continuously vary the inference of different types of policies from one's repertoire. Active inference requires a good sense for which kinds of expected states can actually be achieved – the world does not always co-operate with one's desires. Confident policy inference depends on quite high-level, abstract representations of the underlying statistics, at various time-scales, that might have bearing on the current context. Learning of this type may be compromised in some mental and developmental conditions, which could in some cases lead to increased inference of fidgeting policies. We turn to this topic next.

4. Fidgeting in Psychiatric Conditions: The Case of Autism

One of the diagnostic criteria for Autism Spectrum Conditions in the DSM-5 is "Stereotyped or repetitive motor movements, use of objects, or speech" (American Psychiatric Association, 2013). Autistic individuals' fidgeting is pathologised, as it is often done in a socially inappropriate way (for example larger, more noticeable, more unique actions), and seemingly happens more frequently. It is commonly called *stimming*, which is short for self-stimulation.

Our proposal can make sense of stimming in the autistic case. Previous research has shown that self-cognition may be different in autism (Frith & Happé, 1999; Hobson, 2011; Huang et al., 2017; Lombardo & Baron-Cohen, 2011; Molnar-Szakacs & Uddin, 2016; Perrykkad & Hohwy, 2019; Uddin, 2011; Williams, 2010). From the predictive processing perspective, research suggests that the autistic self-model has less hierarchical depth (Perrykkad & Hohwy, 2019), and autistic people have a high expectation for volatility (Lawson, Mathys, & Rees, 2017), meaning autistic people would accumulate uncertainty faster than neurotypicals. Relatively shallow models with high expectations for volatility would consider much sensory input as unexpected surprise, leading to high uncertainty overall. Autistic individuals could then be expected to recruit fidgeting policies more and for longer, in order to establish and maintain some fidelity of their action-outcome mapping and reduce uncertainty quickly. With a relatively shallow self-model, they will have fewer cognitive resources to model volatile changes in the sensory input, which means that fidgeting policies may remain attractive for longer, even as prediction error arises.

It might be revealing to note how stereotypies, as in the case of autism, differ from tics, as in the case of Tourette syndrome. The characteristics of stereotypies are more similar to the characteristics of fidgeting as we have defined it, than are the characteristics of tics. These differences are outlined nicely by Mills and Hedderly (2014) (in a table based on Barry et al. (2011)), such that stereotypies, concordantly with our approach to fidgeting, are “fixed, identical, foreseeable… rhythmic”, whereas tics are “variable”.

Additionally, while tics are often vehemently avoided, stereotypies are commonly described by autistic people as enjoyable. To our knowledge, the only other explicit functional account of the rate of prediction error minimisation is in determining emotional valence, such that, in general, a higher rate of prediction error minimisation is associated with positive affect (Joffily & Coricelli, 2013; Van de Cruys, 2017; Wilkinson, Deane, Nave, & Clark, 2019). In this way, switching to a policy with a reliably steep prediction error minimisation rate, as in fidgeting, should be enjoyable under a predictive processing explanation. In a similar vein, increased prevalence of anxiety in autism (van Steensel, Bögels, & Perrin, 2011) might also be explained by higher ambiguity for overall expected states which leads to “‘faster and faster’ increase in the violation of the expectations about … existential causes of sensations, eliciting fear at these levels” (Joffily & Coricelli, 2013, p. 12).

Along with what is traditionally considered autistic stimming, some further diagnostic criteria are also relevant here. Specifically, “insistence on sameness, inflexible adherence to routines, or ritualised patterns of verbal or nonverbal behaviour… highly restricted, fixated interests that are abnormal in intensity or focus… hyper- or hyporeactivity to sensory input or unusual interests in sensory aspects of the environment.” (American Psychiatric Association, 2013). These activities can be conceived as self-evidencing behaviours that proactively (if unconsciously) construct an environmental niche in which modelling hidden causes is less complex (Constant, Bervoets, Hens, & Cruys, 2018). This is a longer-term way of building an environment in which the individual is more likely to have an unambiguous action-outcome mapping, meet their expected rate of prediction error minimisation and avoid the need to engage in short term policies for prediction error minimisation like stimming.

In some cases, fidgeting involves self-harm, including some of the more worrisome and clinically problematic versions of stimming in autism, such as head-banging or cutting. This may seem inconsistent with our proposal: if the fidgeting literally destroys the boundaries of the agent’s own body, or damages the vehicle of this self-model, how can it be self-evidencing? We speculate that pain, in this case, may be functioning as a precise source of information about the self because the presence of acute pain is more certain than the presence of a touch or movement of a limb. In spite of their averse nature, actions that involve self-inflicting pain spring from policies with precise action-outcome mappings, which could explain why they are sometimes inferred. In a clinical context, this suggests that instead of focusing on just stopping the self-harm outright, a good strategy may be to replace it with a less harmful stim. The struggle for clinicians will be finding a stimming policy that yields equally precise self-evidencing, given the hierarchical structure of the self-model.

Using autism as a case study, we can then see how impaired fidgeting might arise. Fidgeting maximises reliability of the action at the expense of learning changing and complex environmental regularities. In this sense, it is like pushing Occam’s razor too far. While fidgeting begets precise prediction error, it is too simplistic to capture much of the causal structure of the world.

For the case of autism, we think it is likely that some individuals have quite profound experiences relating to self-evidencing, as illustrated in this conversation:

> Mukhopadhyay: *Rules are formed by an Autistic person to simplify the ongoing uncertainty which is taking place around him. The uncertainty may lead the Autistic person to lose his identity. And because that would be a total chaotic situation, he tends to take the shelter of his rules, which he has created, choosing certain phenomenon from the greater uncertainty surrounding him. …*
>
> Biklen: *I would think that if a rule is known only to you, this could cause difficulty to those around you?*
>
> Mukhopadhyay: *Rules are somewhat the very proof to an Autistic person that he exists. He would have guidelines about these rules, which rule would be performed by him to the extremities of forming a rigid system of ritual. I am no exception and I get a sort of self existing sense when I have followed a routine set of activities.*
>
> *(Tito Rajarshi Mukhopadhyay, who is autistic, speaks with ethnographer Biklen (2005, p. 126)).*

5. Mental fidgeting?

Given that self-evidencing and minimising decay of the self-model relies on effective active inference, this raises a question about what happens in cases where an individual cannot move, such as in locked-in syndrome, paralysis, or imprisonment. In these cases, we would expect that the sense of self transforms or even fades. However, we also must consider mental actions such as imagining and thinking, subtle movements such as eye movements, and homeostatic regulation of bodily states for which the brain receives feedback, all of which might provide some self-evidencing. Mental fidgeting may include rumination or repetitive thoughts, or lapses of concentration and frequent sojourns into mindwandering. We would expect such processes to be more frequent in conditions of constricted or absent agency.

Temporary fading of the belief in bodily existence can reportedly be achieved by the practices of meditation or sensory deprivation. One interpretation of these cases is that they still involve self-hypothesis decay but that for some time period it is possible to maintain a belief that the body is dispersing, rather than choosing to act to re-confirm bodily existence. This may lead to common experiences of the dissolution of the ego boundary under these conditions (see also Letheby and Gerrans (2017)), as self-model decay is not actively resisted. These states can be pleasurably and convincingly upheld until bodily states create endogenous volatility outside of the agent's control – for example, the feeling of hunger. Conversely, in some disorders of the self where self-evidencing is compromised, a tranquil absence of action may be too hard to maintain, and we should see increased mental and bodily fidgeting.

Similarly, there may be cases where repetitive actions induce a trance-like state which has analogous ego-dissolving characteristics. In these cases, individuals deliberatively engage in fidget-like behaviours in order to *create* self-model decay. Here, these individuals are not naturally in a context of growing uncertainty, but are intentionally ignoring the complexities in the world and purposefully inducing longer-term prediction errors. While initially these actions might serve the usual fidgeting function of self-evidencing, when the usual drive to switch policies in order to occupy the most probable states for the agent (as discussed earlier, such as feeding, exploring, paying bills) is *intentionally* delayed, these actions begin to violate the expected rate of prediction error, and serve the opposite function. They make the world too simple/predictable for too long, which leads to ego-dissolution due to (in this case) intentional violation of the expected volatility.

6. Fidgeting in animals and babies?

On the predictive processing framework, self-evidencing is a necessary characteristic of any biological, active inference system. This means that any biological system will implicitly or explicitly represent the expected reduction in uncertainty from its mapping of actions to outcomes, or its expected prediction error minimization. Some creatures may even have explicit representations of themselves, that is, have a sense of self. This implies that many if not all organisms will experience occasional deviations from their expected rate of self-evidencing. This in turn predicts that fidgeting should be observed in many species.

This prediction is modulated by details about the timescale over which such organisms expect to self-evidence and the depth of their hierarchical model such that the overall expected rate of prediction error minimisation will interact with how often a policy like fidgeting will be learned as likely to reduce uncertainty. Organisms with a higher threshold for expected irreducible prediction error and a slower expected rate of prediction error minimisation (e.g., due to a less complex model) may find fidgeting not a very useful policy, as it is too short term. However, it is conceivable that examples of animal fidgeting include a dog chasing its tail, an elephant rocking in a zoo and a Tasmanian devil pacing along circular tracks in its enclosure. Note too that the latter two cases are situations of reduced agency as discussed at the beginning of section 5.

Hommel (2017) suggests that newborn babies do not perform goal-directed actions, but rather start out with many *reflex* behaviours, which are slowly replaced by goal directed actions in development, during which time babies develop a sense of self (Verschoor & Hommel, 2017). In our framework, we conceive of these babies as self-evidencing from their first active inference, and with a marginal self-model at this point representing how they can reduce prediction error over time (see also Friston (2017a)). Their behaviours can then be likened to proto-fidgeting, that is, the first steps in setting up and exploring mappings between actions and outcomes.

7. Conclusion

We have provided a novel understanding of fidgeting as action for self-oriented epistemic value. Fidgeting leads to uncertainty reduction using a policy for action that involves only few hidden causes and which therefore furnishes a highly precise mapping of actions to outcomes. These policies are therefore rationally inferred in some contexts, despite being unfit for bringing about reward in a traditional sense. Fidgeting thus conforms with a form of self-evidencing, and helps the agent confirm its own existence in situations where evidence for its self-model might be waning. Impaired fidgeting can then be seen to arise for individuals who have compromised learning of expected levels of self-evidencing.

References

Adams, R. A., Shipp, S., & Friston, K. J. (2013). Predictions not commands: active inference in the motor system. *Brain Structure and Function, 218*(3), 611-643. doi:10.1007/s00429-012-0475-5

American Psychiatric Association. (2013). *Diagnostic and statistical manual of mental disorders (DSM-5®)*: American Psychiatric Pub.

Andrade, J. (2010). What does doodling do? *Applied Cognitive Psychology: The Official Journal of the Society for Applied Research in Memory and Cognition, 24*(1), 100-106.

Apps, M. A., & Tsakiris, M. (2014). The free-energy self: a predictive coding account of self-recognition. *Neurosci Biobehav Rev, 41*, 85-97. doi:10.1016/j.neubiorev.2013.01.029

Barfield, E. T., Gerber, K. J., Zimmermann, K. S., Ressler, K. J., Parsons, R. G., & Gourley, S. L. (2017). Regulation of actions and habits by ventral hippocampal trkB and adolescent corticosteroid exposure. *PLoS biology, 15*(11), e2003000.

Barry, S., Baird, G., Lascelles, K., Bunton, P., & Hedderly, T. (2011). Neurodevelopmental movement disorders–an update on childhood motor stereotypies. *Developmental Medicine & Child Neurology, 53*(11), 979-985.

Biklen, D. (2005). *Autism and the myth of the person alone*: NYU Press.

Carriere, J. S., Seli, P., & Smilek, D. (2013). Wandering in both mind and body: Individual differences in mind wandering and inattention predict fidgeting. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale, 67*(1), 19.

Cheshire, W. P. (2000). Hypotensive akathisia: autonomic failure associated with leg fidgeting while sitting. *Neurology, 55*(12), 1923-1926.

Clark, A. (2015). *Surfing uncertainty: Prediction, action, and the embodied mind*: Oxford University Press.

Constant, A., Bervoets, J., Hens, K., & Cruys, S. V. d. (2018). Precise Worlds for Certain Minds: An ecological perspective on the relational self in autism. *Topoi*(The Relational Self - Basic Forms of Self-Awareness). doi:10.1007/s11245-018-9546-4

Corcoran, A. W., & Hohwy, J. (2019). Allostasis, interoception, and the free energy principle: Feeling our way forward. In M. Tsakiris & H. De Preester (Eds.), *The Interoceptive Mind* (pp. 272-292). Oxford, UK: Oxford University Press.

Dayan, P. (2009). Goal-directed control and its antipodes. *Neural Networks, 22*(3), 213-219. doi:https://doi.org/10.1016/j.neunet.2009.03.004

Farley, J., Risko, E. F., & Kingstone, A. (2013). Everyday attention and lecture retention: the effects of time, fidgeting, and mind wandering. *Front Psychol, 4*, 619. doi:10.3389/fpsyg.2013.00619

Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience, 11*(2), 127-138.

Friston, K. (2013). Life as we know it. *Journal of The Royal Society Interface, 10*(86).

Friston, K. (2017a). Self-evidencing babies: Commentary on “Mentalizing homeostasis: The social origins of interoceptive inference” by Fotopoulou & Tsakiris. *Neuropsychoanalysis, 19*(1), 43-47. doi:10.1080/15294145.2017.1295216

The Variational Principles of Action 207-235 (Springer International Publishing 2017b).

Friston, K., & Ao, P. (2012). Free energy, value, and attractors. *Computational and mathematical methods in medicine, 2012*.

Friston, K., Daunizeau, J., & Kiebel, S. J. (2009). Reinforcement Learning or Active Inference? *PloS one, 4*(7), e6421. doi:10.1371/journal.pone.0006421

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., J, O. D., & Pezzulo, G. (2016). Active inference and learning. *Neurosci Biobehav Rev, 68*, 862-879. doi:10.1016/j.neubiorev.2016.06.022

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active Inference: A Process Theory. *Neural Comput, 29*(1), 1-49. doi: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.

Frith, U., & Happé, F. (1999). Theory of mind and self-consciousness: What is it like to be autistic? *Mind & Language, 14*(1), 82-89.

Hobson, P. R. (2011). Autism and the self. In S. Gallagher (Ed.), *The Oxford Handbook of the Self* (pp. 571-591). Oxford, UK: Oxford University Press.

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

Hohwy, J. (2016). The Self-Evidencing Brain. *Noûs, 50*(2), 259-285.

Hohwy, J., & Michael, J. (2017). Why should any body have a self? In F. de Vignemont & A. Alsmith (Eds.), *The Body and the Self, Revisited*: MIT Press.

Hohwy, J., Paton, B., & Palmer, C. (2016). Distrusting the present. *Phenomenology and the Cognitive Sciences, 15*(3), 315-335. doi:10.1007/s11097-015-9439-6

Hommel, B. (2015). Action control and the sense of agency. *The sense of agency*, 307-326.

Hommel, B. (2017). Goal-directed actions. *Handbook of Causal Reasoning*, 265-278.

Huang, A. X., Hughes, T. L., Sutton, L. R., Lawrence, M., Chen, X., Ji, Z., & Zeleke, W. (2017). Understanding the Self in Individuals with Autism Spectrum Disorders (ASD): A Review of Literature. *Frontiers in psychology, 8*(1422). doi:10.3389/fpsyg.2017.01422

Joffily, M., & Coricelli, G. (2013). Emotional valence and the free-energy principle. *PLoS computational biology, 9*(6), e1003094.

Johannsen, D. L., & Ravussin, E. (2008). Spontaneous physical activity: relationship between fidgeting and body weight control. *Current Opinion in Endocrinology, Diabetes and Obesity, 15*(5), 409-415.

Lawson, R. P., Mathys, C., & Rees, G. (2017). Adults with autism overestimate the volatility of the sensory environment. *Nature neuroscience*. doi:10.1038/nn.4615

Letheby, C., & Gerrans, P. (2017). Self unbound: ego dissolution in psychedelic experience. *Neuroscience of Consciousness, 3*(1), nix016-nix016. doi:10.1093/nc/nix016

Lis, S., Baer, N., Stein-en-Nosse, C., Gallhofer, B., Sammer, G., & Kirsch, P. (2010). Objective measurement of motor activity during cognitive performance in adults with attention-deficit/hyperactivity disorder. *Acta Psychiatr Scand, 122*(4), 285-294. doi:10.1111/j.1600-0447.2010.01549.x

Lombardo, M. V., & Baron-Cohen, S. (2011). The role of the self in mindblindness in autism. *Consciousness and Cognition, 20*(1), 130-140.

Lovaas, I., Newsom, C., & Hickman, C. (1987). Self-Stimulatory Behaviour and Perceptual Reinforcement. *Journal of Applied Behavior Analysis, 20*(1), 45-68. doi:10.1901/jaba.1987.20-45

Mathys, C., Daunizeau, J., Friston, K. J., & Stephan, K. E. (2011). A Bayesian Foundation for Individual Learning Under Uncertainty. *Frontiers in human neuroscience, 5*, 39. doi:10.3389/fnhum.2011.00039

Mills, S., & Hedderly, T. (2014). A guide to childhood motor stereotypies, tic disorders and the tourette spectrum for the primary care practitioner. *The Ulster medical journal, 83*(1), 22.

Molnar-Szakacs, I., & Uddin, L. Q. (2016). The Self in Autism. In M. Kyrios, R. Moulding, G. Doron, S. S. Bhar, M. Nedeljkovic, & M. Mikulincer (Eds.), *The Self in Understanding and Treating Psychological Disorders* (pp. 144-157). Cambridge, UK: Cambridge University Press.

Morishima, T., Restaino, R. M., Walsh, L. K., Kanaley, J. A., Fadel, P. J., & Padilla, J. (2016). Prolonged sitting-induced leg endothelial dysfunction is prevented by fidgeting. *American Journal of Physiology-Heart and Circulatory Physiology, 311*(1), H177-H182.

Morrens, M., Hulstijn, W., Lewi, P. J., De Hert, M., & Sabbe, B. G. C. (2006). Stereotypy in schizophrenia. *Schizophrenia Research, 84*(2), 397-404. doi:https://doi.org/10.1016/j.schres.2006.01.024

O'shaughnessy, B. (1980). *The will* (Vol. 2): Cambridge University Press Cambridge.

Perrykkad, K., & Hohwy, J. (2019). Modelling Me, Modelling You: the Autistic Self. *Review Journal of Autism and Developmental Disorders*. doi:10.1007/s40489-019-00173-y

Peters, A., McEwen, B. S., & Friston, K. (2017). Uncertainty and stress: Why it causes diseases and how it is mastered by the brain. *Progress in Neurobiology, 156*, 164-188. doi:https://doi.org/10.1016/j.pneurobio.2017.05.004

Skinner, B. F. (1948). 'Superstition' in the pigeon. *Journal of Experimental Psychology, 38*(2), 168-172. doi:10.1037/h0055873

Uddin, L. Q. (2011). The self in autism: An emerging view from neuroimaging. *Neurocase, 17*(3), 201-208. doi:10.1080/13554794.2010.509320

Van de Cruys, S. (2017). Affective Value in the Predictive Mind. In T. K. Metzinger & W. Wiese (Eds.), *Philosophy and Predictive Processing*. Frankfurt am Main: MIND Group.

van Steensel, F. J. A., Bögels, S. M., & Perrin, S. (2011). Anxiety Disorders in Children and Adolescents with Autistic Spectrum Disorders: A Meta-Analysis. *Clinical Child and Family Psychology Review, 14*(3), 302. doi:10.1007/s10567-011-0097-0

Verschoor, S. A., & Hommel, B. (2017). Self-By-Doing: The Role of Action for Self-Acquisition. *Social Cognition, 35*(2), 127-145. doi:10.1521/soco.2017.35.2.127

Wilkinson, S., Deane, G., Nave, K., & Clark, A. (2019). Getting Warmer: Predictive Processing and the Nature of Emotion. In L. Candiotto (Ed.), *The Value of Emotions for Knowledge* (pp. 101-119). Cham: Springer International Publishing.

Williams, D. (2010). Theory of own mind in autism: Evidence of a specific deficit in self-awareness? *autism, 14*(5), 474-494. doi:10.1177/1362361310366314

Acknowledgements:

The authors would like to thank Thomas Parr for useful discussion and generous comments on this paper.

Author’s Note:

On behalf of all authors, the corresponding author states that there is no conflict of interest. Kelsey Perrykkad is supported by an Australian Government Research Training Program (RTP) Scholarship. Jakob Hohwy is supported by the Australian Research Council (DP160102770, DP190101805).

Table 1
Examples of fidgeting in different modalities for typically developing adults.

Primary Modality	*Examples*
Visual	Doodling Visually tracking a rotating fan Absentmindedly arranging objects on desk
Vestibular	Rocking on a chair Absentminded head nodding
Tactile	Playing with own hair Touching own face Rubbing a soft sweater
Auditory	Clicking a pen Tapping foot Humming
Taste	Chewing gum Sucking on a toothpick
Proprioception	Bouncing one’s leg