

10 Active Inference as a Unified Theory of Sentient Behavior

In general we are least aware of what our minds do best.

—Marvin Minsky

10.1 Introduction

In this chapter, we wrap up Active Inference’s main theoretical points (from the first part of the book) and its practical implementations (from the second part). Then, we connect the dots: we abstract away from the specific Active Inference models discussed in previous chapters to focus on integrative aspects of the framework. One benefit of Active Inference is that it provides a complete solution to the adaptive problems that sentient organisms have to solve. It therefore offers a unified perspective on problems like perception, action selection, attention, and emotion regulation, which are usually treated in isolation in psychology and neuroscience—and addressed using distinct computational approaches in artificial intelligence. We will discuss the Active Inference perspective on each of these problems (and more) in the context of established theories, such as cybernetics, ideomotor theory of action, reinforcement learning, and optimal control. Finally, we briefly discuss how the scope of Active Inference can be extended to cover other biological, social, and technological topics that are not discussed in depth in this book.

10.2 Wrapping Up

This book offers a systematic account of the theoretical underpinnings and practical implementations of Active Inference. Here, we briefly summarize the discussion of the first nine chapters. This offers an opportunity to

rehearse the key constructs of Active Inference that will be useful in the remainder of this chapter.

In chapter 1, we introduced Active Inference as a normative approach to understanding sentient creatures that form part of action-perception loops with their environment (Fuster 2004). We explained that normative approaches start from first principles to derive and test empirical predictions about the phenomenon of interest—here, the ways living organisms persist while engaging in adaptive exchanges (*action-perception loops*) with their environment. We also considered that one could arrive at Active Inference by following a low road or a high road.

In chapter 2, we illustrated the low road to Active Inference. This road starts from the idea that the brain is a prediction machine, endowed with a *generative model*: a probabilistic representation of how hidden causes in the world generate sensations (e.g., how light reflected off an apple stimulates the retina). By inverting this model, it infers the causes of its sensations (e.g., whether I am seeing an apple, given that my retina is stimulated in a certain way). This view of perception (aka perception-as-inference) has its historical roots in the Helmholtzian notion of unconscious inference and, more recently, in the *Bayesian brain* hypothesis. Active Inference extends this view by bringing action control and planning within the compass of inference (aka control-as-inference, planning-as-inference). Most importantly, it shows that perception and action are not quintessentially separable processes but fulfill the same objective. We first described this objective more informally, as the minimization of a discrepancy between one's model and the world (which generally reduces to surprise or prediction error minimization). Put simply, one can minimize the discrepancy between a model and the world in two ways: by changing one's mind to fit the world (perception) or by changing the world to fit the model (action). These can be described in terms of Bayesian inference. However, exact inference is often intractable, so Active Inference uses a (variational) approximation (noticing that exact inference may be seen as a special case of approximate inference). This leads to the second, more formal description of the common objective of perception and action, as *variational free energy* minimization. This is the core quantity used in Active Inference and may be unpacked in terms of its constituent parts (e.g., energy and entropy, complexity and accuracy, or surprise and divergence). Finally, we introduced a second kind of free energy: *expected free energy*. This is particularly important during planning,

as it affords a way to score alternative policies by considering the future outcome that they are expected to generate. This too may be unpacked in terms of its constituent parts (e.g., information gain and pragmatic value, expected ambiguity and risk).

In chapter 3, we illustrated the high road to Active Inference. This alternative road starts from the deflationary imperative for biological organisms to preserve their integrity and avoid dissipation, which can be described as avoiding surprising states. We then introduced the notion of a *Markov blanket*: a formalization of the statistical separation between the organism's internal states and the world's external states. Crucially, internal and external states can only influence each other vicariously via intermediate (active and sensory) variables, called blanket states. This statistical separation—mediated by the Markov blanket—is crucial to endowing an organism with some degree of autonomy from the external world. To understand why this is a useful perspective, consider the following three consequences.

First, an organism with a Markov blanket appears to model the external environment in a Bayesian sense: its internal states correspond—on average—to an approximate posterior belief about external states of the world. Second, the autonomy is guaranteed by the fact that the organism's model (its internal states) is not unbiased but prescribes some existential preconditions (or prior preferences) that must be maintained—for example, for a fish, being in the water. Third, equipped with this formalism, it is possible to describe optimal behavior (with respect to prior preferences) as the maximization of (Bayesian) *model evidence* by perception and action. By maximizing model evidence (i.e., *self-evidencing*) an organism ensures that it realizes its prior preferences (e.g., a fish stays in the water) and avoids surprising states. In turn, the maximization of model evidence is (approximately) mathematically equivalent to the minimization of *variational free energy*—hence we arrive again (in another way) at the same central construct of Active Inference discussed in chapter 2. Finally, we detailed the relationship between minimizing surprise and *Hamilton's principle of least Action*. This evinces the formal relationship between Active Inference and first principles in statistical physics.

In chapter 4, we outlined the formal aspects of Active Inference. We focused on the passage from Bayesian inference to a tractable approximation—*variational inference*—and the resulting objective for organisms to minimize *variational free energy* via perception and action. The insight from this

treatment is the importance of the generative model that creatures use to make sense of their world. We introduced two kinds of generative models that express our beliefs about how data are generated, using discrete or continuous variables. We explained that both afford the same Active Inference, but they apply when states of affairs are formulated in discrete time (as partially observed Markov decision problems) or continuous time (as stochastic differential equations), respectively.

In chapter 5, we remarked on the difference between the normative principle of free energy minimization and a process theory about how this principle may be implemented by the brain—and explained that the latter generates testable predictions. We then outlined aspects of the process theories accompanying Active Inference, which encompass domains such as neuronal message passing, including neuroanatomical circuitry (e.g., cortico-subcortical loops) and neuromodulation. For example, at an anatomical level, message passing maps nicely to a canonical cortical microcircuit, with predictions that stem from deep cortical layers at one level and target superficial cortical layers at the level below (Bastos et al. 2012). At a more systemic level, we discussed how Bayesian inference, learning, and precision weighting correspond to neuronal dynamics, synaptic plasticity, and neuromodulation, respectively, and how the top-down and bottom-up neural message passing of predictive coding maps to slower (e.g., alpha or beta) and faster (e.g., gamma) brain rhythms. These and other examples illustrate that after designing a specific Active Inference model, one can draw neurobiological implications from the form of its generative model.

In chapter 6, we provided a recipe to design Active Inference models. We saw that while all creatures minimize their variational free energy, they behave in different, sometimes opposite ways because they are endowed with different generative models. Therefore, what distinguishes different (e.g., simpler from more complex) creatures is just their generative model. There is a rich repertoire of possible generative models, which correspond to different biological (e.g., neuronal) implementations and produce different adaptive—or maladaptive—behaviors in different contexts and ecological niches. This renders Active Inference equally appropriate for characterizing simple creatures like bacteria that sense and seek nutrient gradients, complex creatures like us that pursue sophisticated goals and engage in rich cultural practices, or even different individuals—to the extent that ones appropriately characterizes their respective generative models. Evolution

appears to have discovered increasingly sophisticated design structures for brains and bodies that made organisms able to deal with (and shape) rich ecological niches. Modelers can reverse-engineer this process and specify the designs for brains and bodies of creatures of interest, in terms of generative models, based on the kinds of niche they occupy. This corresponds to a series of design choices (e.g., models using discrete or categorical variables, shallow or hierarchical models)—which we unpacked in the chapter.

In chapters 7 and 8, we provided numerous examples of Active Inference models in discrete and continuous time, which address problems of perceptual inference, goal-directed navigation, model learning, action control, and more. These examples were designed to showcase the variety of emergent behaviors under these models and to detail the principles of how they are specified practically.

In chapter 9, we discussed how to use Active Inference for model-based data analysis and to recover the parameters of an individual's generative model, which better explain the subject's behavior in a task. This *computational phenotyping* uses the same form of Bayesian inference discussed in the rest of the book, but in a different way: it helps design and evaluate (objective) models of others' (subjective) models.

10.3 Connecting the Dots: The Integrative Perspective of Active Inference

Some decades ago, the philosopher Dennett lamented that cognitive scientists devoted too much effort to modeling isolated subsystems (e.g., perception, language understanding) whose boundaries are often arbitrary. He suggested to try instead modeling “the whole iguana”: a complete cognitive creature (perhaps a simple one) and an environmental niche for it to cope with (Dennett 1978).

One benefit of Active Inference is that it offers a first principle account of the ways in which organisms solve their adaptive problems. The normative approach pursued in this book assumes that it is possible to start from the principle of variational free energy minimization and derive implications about specific cognitive processes, such as perception, action selection, attention and emotion regulation, and their neuronal underpinnings.

Imagine a simple creature that must solve problems like finding food or shelter. When cast as Active Inference, the creature's problems can be

described in enactive terms, as acting to solicit preferred sensations (e.g., food-related sensations). To the extent that these preferred sensations are included (as prior beliefs) in its generative model, the organism is effectively gathering evidence for its model—or, more allegorically, for its existence (i.e., maximizing model evidence or self-evidencing). This simple principle has ramifications for psychological functions traditionally considered in isolation, such as perception, action control, memory, attention, intention, emotion, and more. For example, perception and action are both self-evidencing, in the sense that a creature can align what it expects, given its generative model, with what it senses either by changing its beliefs (about the presence of food) or by changing the world (soliciting food-related sensations). Memory and attention can also be thought of as optimizing the same objective. Long-term memory develops through learning the parameters of a generative model. Working memory is belief updating when beliefs are about external states in the past and future. Attention is the optimization of beliefs about the precision of sensory input. Forms of planning (and intentionality) can be conceptualized by appealing to the capacity of (some) creatures to select among alternative futures, which in turn requires temporally deep generative models. These predict the outcomes that would result from a course of action and are optimistic about these outcomes. This optimism manifests as the belief that future outcomes will lead to preferred outcomes. Deep temporal models can also help us understand sophisticated forms of prospection (where beliefs about the present are used to derive beliefs about the future) and retrospection (where beliefs about the present are used to update beliefs about the past). Forms of interoceptive regulation and emotion can be conceptualized by appealing to generative models of internal physiology that predict the allostatic consequences of future events.

As the above examples illustrate, there is an important consequence of studying cognition and behavior from the perspective of a normative theory of sentient behavior. Such theory does not start by assembling separate cognitive functions, such as perception, decision-making, and planning. Rather, it starts by providing a complete solution to the problems that organisms have to solve and then analyzing the solution to derive implications about cognitive functions. For example, which mechanisms permit a living organism or artificial creature (e.g., a robot) to perceive the world, remember it, or plan (Verschure et al. 2003, 2014; Verschure 2012; Pezzulo, Barsalou et al. 2013; Krakauer et al. 2017)? This is an important

move as the taxonomies of cognitive functions—used in psychology and neuroscience textbooks—largely inherit from early philosophical and psychological theories (sometimes called *Jamesian categories*). Despite their great heuristic value, they may be quite arbitrary—or they may not correspond to separate cognitive and neural processes (Pezzulo and Cisek 2016, Buzsaki 2019, Cisek 2019). Indeed, these Jamesian categories may be candidates for how our generative models explain our engagement with the sensorium—as opposed to explaining that engagement. For example, the solipsistic hypothesis that “I am perceiving” is just my explanation for current states of affairs that include my belief updating.

Adopting a normative perspective may also help in identifying formal analogies between cognitive phenomena studied in different domains. One example is the trade-off between exploration and exploitation, which appears in various guises (Hills et al. 2015). This trade-off is often studied during foraging, when creatures must choose between exploiting previous successful plans and exploring novel (potentially better) ones. However, the same trade-off occurs during memory search and deliberation with limited resources (e.g., time limitations or search effort), when creatures have the choice between exploiting their current best plan versus investing more time and cognitive effort to explore additional possibilities. Characterizing these apparently disconnected phenomena in terms of free energy can potentially reveal deep similarities (Friston, Rigoli et al. 2015; Pezzulo, Cartoni et al. 2016; Gottwald and Braun 2020).

Finally, in addition to a unified perspective on psychological phenomena, Active Inference offers a principled means of understanding the corresponding neural computations. In other words, it offers a process theory that connects cognitive processing to (expected) neuronal dynamics. Active Inference assumes that everything that matters about brains, minds, and behavior can be described in terms of the minimization of variational free energy. In turn, this minimization has specific neural signatures (in terms of, e.g., message passing or brain anatomy) that can be empirically validated.

In the rest of this chapter, we explore some implications of Active Inference for psychological functions—as if we were sketching a psychology textbook. For each of these functions, we also highlight some points of contact (or divergence) between Active Inference and other popular theories in the literature.

10.4 Predictive Brains, Predictive Minds, and Predictive Processing

I have this picture of pure joy
 it's of a child with a gun
 he's aiming straight in front of himself,
 shooting at something that isn't there.

—Afterhours, “Quello che non c'è” (Something that isn't there)

Traditional theories of brain and cognition emphasize feedforward transductions from external stimuli to internal representations and then motor actions. This has been called a “sandwich model,” as everything that is in between stimuli and responses is assigned the label “cognitive” (Hurley 2008). In this perspective, the main function of the brain is to transform incoming stimuli into contextually appropriate responses.

Active Inference departs significantly from this view by emphasizing predictive and goal-directed aspects of brain and cognition. In psychological terms, Active Inference creatures (or their brains) are *probabilistic inference machines*, which continuously generate predictions based on their generative models.

Self-evidencing creatures use their predictions in two fundamental ways. First, they compare predictions with incoming data to validate their hypotheses (predictive coding) and—at a slower timescale—revise their models (learning). Second, they enact predictions to guide the ways they gather data (Active Inference). By doing so, Active Inference creatures fulfill two imperatives: epistemic (e.g., visually exploring places where salient information is present that can resolve uncertainty about hypotheses or models) and pragmatic (e.g., moving to locations where preferred observations such as rewards can be secured). The epistemic imperative renders both perception and learning *active* processes, whereas the pragmatic imperative renders behavior *goal directed*.

10.4.1 Predictive Processing

This predictive- and goal-centric view of brain—and cognition—is closely related to (and provided inspiration for) *predictive processing* (PP): an emerging framework in philosophy of mind and epistemology, which sees prediction as central to brain and cognition and appeals to concepts of “predictive brains” or “predictive minds” (Clark 2013, 2015; Hohwy 2013).

Sometimes PP theories appeal to the specific functioning of Active Inference and some of its constructs, such as generative models, predictive coding, free energy, precision control, and Markov blankets, but they sometimes appeal to other constructs, such as coupled inverse and forward models, which are not part of Active Inference. Therefore, the term *predictive processing* is used in a broader (and less constrained) sense compared to Active Inference.

Predictive processing theories have attracted considerable attention in philosophy, given their potential for unification in many senses: across multiple domains of cognition, including perception, action, learning, and psychopathology; from lower (e.g., sensorimotor) to higher levels of cognitive processing (e.g., psychological constructs); from simple biological organisms to brains, individuals, and social and cultural constructs. Another appeal of PP theories is that they make use of conceptual terms, such as *beliefs* and *surprise*, which speak to a psychological level of analysis familiar to philosophers (with the caveat that sometimes these terms may have technical meanings that differ from common usage).

Yet, as the interest in PP grows, it has become apparent that philosophers have different opinions on its theoretical and epistemological implications. For example, it has been interpreted in internalist (Hohwy 2013), embodied or action-based (Clark 2015), and enactivist and nonrepresentational terms (Bruineberg et al. 2016, Ramstead et al. 2019). The debate around these conceptual interpretations goes beyond the scope of this book.

10.5 Perception

You can't depend on your eyes when your imagination is out of focus.

—Mark Twain

Active Inference considers perception as an inferential process based on a generative model of how sensory observations are generated. Bayes' rule essentially inverts the model to compute a belief about the hidden state of the environment, given the observations. This idea of perception-as-inference dates back to Helmholtz (1866) and was often reposed in psychology, computational neuroscience, and machine learning (e.g., analysis-by-synthesis) (Gregory 1980, Dayan et al. 1995, Mesulam 1998, Yuille and Kersten 2006). This generative modeling approach has been demonstrated to be effective in

facing challenging perceptual problems, such as breaking text-based CAPTCHAs (George et al. 2017).

10.5.1 Bayesian Brain Hypothesis

The most prominent contemporary expression of this idea is the Bayesian brain hypothesis, which has been applied to several domains such as decision-making, sensory processing, and learning (Doya 2007). Active Inference provides a normative foundation to these inferential ideas by deriving them from the imperative of minimizing variational free energy. As the same imperative extends to action dynamics, Active Inference naturally models *active perception* and the ways in which organisms actively sample observations to test their hypotheses (Gregory 1980). Under the Bayesian brain agenda, instead, perception and action are modeled in terms of different imperatives (where action requires Bayesian decision theory; see section 10.7.1).

More broadly, the Bayesian brain hypothesis refers to a family of approaches that are not necessarily integrated and often make different empirical predictions. These include, for example, the computational-level proposal that the brain performs Bayes-optimal sensorimotor and multisensory integration (Kording and Wolpert 2006), the algorithmic-level proposal that the brain implements specific approximations of Bayesian inference, such as decision-by-sampling (Stewart et al. 2006), and the neural-level proposals about the specific ways in which neural populations may perform probabilistic computations or encode probability distributions—for example, as samples or probabilistic population codes (Fiser et al. 2010, Pouget et al. 2013). At each level of explanation, there are competing theories on the field. For example, it is common to appeal to approximations of exact Bayesian inference to explain deviations from optimal behavior, but different works consider different (and not always compatible) approximations, such as different sampling approaches. More broadly, the relations between proposals at different levels are not always straightforward. This is because Bayesian computations can be realized (or approximated) in multiple algorithmic ways, even without explicitly representing probability distributions (Aitchison and Lengyel 2017).

Active Inference provides a more integrated perspective that connects normative principles and process theories. At the normative level, its central

assumption is that all processes minimize variational free energy. The corresponding process theory for inference uses a gradient descent on free energy, which has clear neurophysiological implications, explored in chapter 5 (Friston, FitzGerald et al. 2016). More broadly, one can start from the principle of free energy minimization to derive implications about brain architectures.

For example, the canonical process model of perceptual inference (in continuous time) is *predictive coding*. Predictive coding was initially proposed as a theory of hierarchical perceptual processing by Rao and Ballard (1999) to explain a range of documented top-down effects, which were difficult to reconcile with feedforward architectures as well as known physiological facts (e.g., the existence of forward, or bottom-up, and backward, or top-down, connections in sensory hierarchies). However, predictive coding can be derived from the principle of free energy minimization, under some assumptions, such as the Laplace approximation (Friston 2005). Furthermore, Active Inference *in continuous time* can be constructed as a directed extension of predictive coding into the domain of action—by endowing a predictive coding agent with motor reflexes (Shipp et al. 2013). This leads us to the next point.

10.6 Action Control

If you can't fly then run, if you can't run then walk, if you can't walk then crawl, but whatever you do you have to keep moving forward.

—Martin Luther King

In Active Inference, action processing is analogous to perceptual processing, as both are guided by forward predictions—exteroceptive and proprioceptive, respectively. It is the (proprioceptive) prediction that “my hand grasps the cup” that induces a grasping movement. The equivalence between action and perception exists also at the neurobiological level: the architecture of the motor cortex is organized in the same way as the sensory cortex—as a predictive coding architecture, with the exceptions that it can influence motor reflexes in the brain stem and spine (Shipp et al. 2013) and that it receives relatively little ascending input. Motor reflexes permit controlling movement by setting “equilibrium points” along a desired

trajectory—an idea that corresponds to the *equilibrium point hypothesis* (Feldman 2009).

Importantly, initiating an action—like grasping a cup—requires regulation of the precision (inverse variance) of prior beliefs and sensory streams appropriately. This is because the relative values of these precisions determine the way in which a creature manages the conflict between its prior belief (that it holds the cup) and its sensory input (signaling that it does not). An imprecise prior belief about grasping a cup can be easily revised in the light of conflicting sensory evidence—producing a change of mind and no action. Rather, when the prior belief dominates (i.e., has higher precision), it is maintained even in the face of conflicting sensory evidence—and it induces a grasping action to resolve the conflict. To ensure that this is the case, action initiation induces a transient sensory attenuation (or down-weighting sensory prediction errors). Failure of this sensory attenuation can have maladaptive consequences, such as the failure to initiate or control movements (Brown et al. 2013).

10.6.1 Ideomotor Theory

In Active Inference, action stems from (proprioceptive) predictions and not motor commands (Adams, Shipp, and Friston 2013). This idea connects Active Inference to *ideomotor theory* of action: a framework to understand action control that dates back to William James (1890) and the later theories of “event coding” and “anticipatory behavioural control” (Hommel et al. 2001, Hoffmann 2003). Ideomotor theory suggests that *action-effect links* (similar to forward models) are key mechanisms in the architecture of cognition. Importantly, these links can be used bidirectionally. When they are used in the *action-effect* direction, they permit generating sensory predictions; when they are used in the *effect-action* direction, they permit selecting actions that achieve desired perceptual consequences—implying that actions are selected and controlled on the basis of their predicted consequences (hence the term *ideo+motor*). This anticipatory view of action control is supported by a body of literature that documents the effects of (anticipated) action consequences on action selection and execution (Kunde et al. 2004). Active Inference provides a mathematical characterization of this idea that also includes additional mechanisms, such as the importance of precision control and sensory attenuation, which are not fully investigated in (but are compatible with) ideomotor theory.

10.6.2 Cybernetics

Active Inference is closely related to cybernetic ideas about the purposeful, goal-directed nature of behavior and the importance of (feedback-based) agent-environment interactions, as exemplified by the TOTE (Test, Operate, Test, Exit) and related models (Miller et al. 1960; Pezzulo, Baldassarre et al. 2006). In both TOTE and Active Inference, the selection of actions is determined by the discrepancy between a preferred (goal) state and the current state. These approaches diverge from simple stimulus-response relationships, as more commonly assumed in behaviorist theory and computational frameworks like reinforcement learning (Sutton and Barto 1998).

The notion of action control in Active Inference is particularly akin to *perceptual control theory* (Powers 1973). Central to perceptual control theory was the notion that what is controlled is a perceptual state, not a motor output or action. For example, while driving, what we control—and keep stable over time in the face of disturbances—is our reference or desired velocity (e.g., 90 mph), as signaled by the speedometer, whereas the actions we select for this (e.g., accelerating or decelerating) are more variable and context dependent. For example, depending on the disturbance (e.g., wind, a steep road, or other cars), we would need to either accelerate or decelerate to maintain the reference velocity. This view implements William James's (1890) suggestion that “humans achieve stable goals via flexible means.”

While in both Active Inference and perceptual control theory it is a perceptual (and specifically a proprioceptive) prediction that controls action, the two theories differ in how control is operated. In Active Inference but not perceptual control theory, action control has anticipatory or feedforward aspects, based on generative models. In contrast, perceptual control theory assumes that feedback mechanisms are largely sufficient to control behavior, whereas trying to predict a disturbance, or exerting feedforward (or open-loop) control, is worthless. However, this objection was mainly intended to address the limitations of control theories that use inverse-forward models (see next section). Under Active Inference, generative or forward models are not used to predict a disturbance but to predict future (desired) states and trajectories to be fulfilled by acting—and to infer the latent cause of perceptual events.

Finally, another important point of contact between Active Inference and perceptual control theory is the way they conceptualize control hierarchies. Perceptual control theory proposes that higher hierarchical levels control

lower hierarchical levels by setting their reference points or set-points (i.e., *what* they have to achieve) by leaving them free to select the means to achieve them rather than by setting or biasing the actions that the lower levels have to perform (i.e., *how* to operate). This stands in contrast with most theories of hierarchical and top-down control, in which higher levels either directly select plans (Botvinick 2008) or bias the selection of actions or motor commands at lower hierarchical levels (Miller and Cohen 2001). Similar to perceptual control theory, in Active Inference one can decompose hierarchical control in terms of a (top-down) cascade of goals and subgoals, which can be autonomously achieved at the appropriate (lower) levels. Furthermore, in Active Inference, the contribution of goals represented at different levels of the control hierarchy can be modulated (precision weighted) by motivational processes, in such a way that the more salient or urgent goals are prioritized (Pezzulo, Rigoli, and Friston 2015, 2018).

10.6.3 Optimal Control Theory

The way Active Inference accounts for action control is significantly different from other models of control in neuroscience, such as *optimal control theory* (Todorov 2004, Shadmehr et al. 2010). This framework assumes that the brain's motor cortex selects actions using a (reactive) control policy that maps stimuli to responses. Active Inference, instead, assumes that the motor cortex conveys predictions, not commands.

Furthermore, while both optimal control theory and Active Inference appeal to internal models, they describe internal modeling in different ways (Friston 2011). In optimal control, there is a distinction between two kinds of internal models: *inverse models* encode stimulus-response contingencies and select motor commands (according to some cost function), whereas *forward models* encode action-outcome contingencies and provide inverse models with simulated inputs to replace noisy or delayed feedback, hence going beyond a pure feedback control scheme. Inverse and forward models can also operate in a loop that is detached from external action-perception (i.e., when inputs and outputs are suppressed) to support internal, "what if" simulations of action sequences. Such internal simulations of action have been linked to various cognitive functions, such as planning, action perception, and imitation in social domains (Jeannerod 2001, Wolpert et al. 2003) as well as various disorders of movement and psychopathologies (Frith et al. 2000).

In contrast to the forward-inverse modeling scheme, in Active Inference forward (generative) models do the heavy lifting of action control, whereas inverse models are minimalistic and often reduce to simple reflexes resolved at the peripheral level (i.e., in the brain stem or spinal cord). Action is initiated when there is a difference between anticipated and observed states (e.g., desired, current arm positions)—that is, a sensory prediction error. This means a motor command is equivalent to a prediction made by the forward model as opposed to something computed by an inverse model as in optimal control. The sensory (more precisely, proprioceptive) prediction error is resolved by an action (i.e., arm movement). The gap to be filled by action is considered so small that it does not require a sophisticated inverse model but a much simpler motor reflex (Adams, Shipp, and Friston 2013).¹ What renders a motor reflex simpler than an inverse model is that it does not encode a mapping from inferred states of the world to action but a much simpler mapping between action and sensory consequences. See Friston, Daunizeau et al. (2010) for further discussion.

Another crucial difference between optimal motor control and Active Inference is that the former uses a notion of *cost* or *value function* to motivate action, whereas the latter replaces it with the Bayesian notion of prior (or prior preference, implicit in expected free energy)—as we discuss in the next section.

10.7 Utility and Decision-Making

Action expresses priorities.

—Mahatma Gandhi

The notion of a cost or value function of states is central in many fields, such as optimal motor control, economic theories of utility maximization, and reinforcement learning. For example, in optimal control theory, the optimal control policy for a reaching task is often defined as the one that minimizes a specific cost function (e.g., is smoother or has minimum jerk). In reinforcement learning problems, such as navigating in a maze that includes one or more rewards, the optimal policy is the one that permits maximizing (discounted) reward while also minimizing movement costs. These problems are often solved using the Bellman equation (or the Hamilton-Jacobi-Bellman equation in continuous time), whose general idea is that the value

of a decision can be decomposed in two parts: the immediate reward and the value of the remaining part of the decision problem. This decomposition affords the iterative procedure of *dynamic programming*, which is at the core of control theory and reinforcement learning (RL) (Bellman 1954).

Active Inference differs from the above approach in two main ways. First, Active Inference does not consider utility maximization alone but the broader objective of (expected) free energy minimization, which also includes additional (epistemic) imperatives, such as the disambiguation of current state and novelty seeking (see figure 2.5). These additional objectives are sometimes added on to classical rewards—for example, as a “novelty bonus” (Kakade and Dayan 2002) or “intrinsic reward” (Schmidhuber 1991, Oudeyer et al. 2007, Baldassarre and Mirolli 2013, Gottlieb et al. 2013)—but they arise automatically in Active Inference, enabling it to resolve exploration-exploitation trade-offs implicit in many decisions. The reason for this is that free energies are functionals of beliefs, which means we are in the realm of belief optimization as opposed to external reward functions. This is essential in explorative problems, wherein success depends on resolving as much uncertainty as possible.

Second, in Active Inference, the notion of cost is absorbed into the prior. The prior (or prior preference) specifies an objective for control—for example, a trajectory to follow or an endpoint to reach. Using priors to encode preferred observations (or sequences) may be more expressive than using utilities (Friston, Daunizeau, and Kiebel 2009). Using this method, finding the optimal policy is recast as a problem of inference (of a sequence of control states that realize the preferred trajectory) and does not require a value function or the Bellman equation—although can appeal to a similar recursive logic (Friston, Da Costa et al. 2020). There are at least two fundamental differences between the ways priors and value functions are normally used in Active Inference and RL, respectively. First, RL methods use value functions of states or of state-action pairs—whereas Active Inference uses priors over observations. Second, value functions are defined in terms of the expected return of being in a state (or performing an action in a state) following a specific policy—that is, the sum of future (discounted) rewards obtained by starting in the state and then executing the policy. In contrast, in Active Inference, priors do not usually sum future rewards, nor do they discount them. Rather, something analogous to the expected return only emerges in Active Inference when the expected free energy

of a policy is calculated. The implication is that expected free energy is the closest analogue to the value function. However, even this differs in the sense that expected free energy is a functional of beliefs about states, not a function of states. Having said this, it is possible to construct priors that resemble value functions of states in RL—for example, by caching expected free energy calculations in these states (Friston, FitzGerald et al. 2016; Maisto, Friston, and Pezzulo 2019).

Furthermore, absorbing the notion of utility into the prior has an important theoretical consequence: priors play the role of goals and render the generative model biased—or optimistic, in the sense that the creature believes it will encounter preferred outcomes. It is this optimism that underwrites inferred plans that achieve desired outcomes in Active Inference; a failure of this sort of optimism may correspond to apathy (Hezemans et al. 2020). This stands in contrast with other formal approaches to decision-making, such as Bayesian decision theory, which separate the probability of events from their utility. Having said this, this distinction is somewhat superficial, as a utility function can always be rewritten as encoding a prior belief, consistent with the fact that behaviors that maximize a utility function are a priori (and by design) more probable. From one (slightly tautological) deflationary perspective, this is the definition of utility.

10.7.1 Bayesian Decision Theory

Bayesian decision theory is a mathematical framework that extends the ideas of the Bayesian brain (discussed above) to the domains of decision-making, sensorimotor control, and learning (Kording and Wolpert 2006, Shadmehr et al. 2010, Wolpert and Landy 2012). Bayesian decision theory describes decision-making in terms of two distinct processes. The first process uses Bayesian computations to predict the probability of future (action- or policy-dependent) outcomes, and the second process defines the preference over plans, using a (fixed or learned) utility or cost function. The final decision (or action selection) process integrates both streams, thus selecting (with higher probability) the action plan that has the higher probability of yielding the higher reward. This stands in contrast to Active Inference, in which the prior distribution directly signals what is valuable for the organism (or what has been valuable during evolutionary history). However, parallels could be drawn between the two streams of Bayesian decision theory and the optimization of variational and expected free energy, respectively. Under

Active Inference, the minimization of variational free energy affords accurate (and simple) beliefs about the state of the world and its likely evolution. The prior belief that expected free energy will be minimized through policy selection incorporates the notion of preferences.

In some circles, there are concerns about the status of Bayesian decision theory. This follows from the complete class theorems (Wald 1947, Brown 1981) that say for any given pair of decisions and cost functions, there exist some prior beliefs that render the decisions Bayes optimal. This means that there is an implicit duality or degeneracy when dealing separately with prior beliefs and cost functions. In one sense, Active Inference resolves this degeneracy by absorbing utility or cost functions into prior beliefs in the form of preferences.

10.7.2 Reinforcement Learning

Reinforcement learning (RL) is an approach to solving Markov decision problems that is popular in both artificial intelligence and the cognitive sciences (Sutton and Barto 1998). It focuses on how agents learn a policy (e.g., pole balancing strategy) by trial and error: by trying out actions (e.g., move to the left) and receiving positive or negative reinforcements, depending on action success (e.g., pole balanced) or failure (e.g., pole fallen).

Active Inference and RL address overlapping sets of problems but differ in many respects mathematically and conceptually. As noted above, Active Inference dispenses with the notions of reward, value functions, and Bellman optimality that are key to reinforcement learning approaches. Furthermore, the notion of *policy* is used differently in the two frameworks. In RL a policy denotes a set of stimulus-response mappings that need to be learned. In Active Inference, a policy is part of the generative model: it denotes a sequence of control states that need to be inferred.

Reinforcement learning approaches are plentiful, but they can be subdivided into three main families. The first two methods try to learn good (state or state-action) value functions, albeit in two different ways.

Model-free methods of RL learn value functions directly from experience: they perform actions, collect rewards, update their value functions, and use them to update their policies. The reason they are called *model-free* is because they do not use a (transition) model that permits predicting future states—of the sort used in Active Inference. Instead, they implicitly appeal to simpler kinds of models (e.g., state-action mappings). Learning value functions in model-free RL often involves computing *reward* prediction

errors, as in the popular temporal-difference rule. While Active Inference often appeals to prediction errors, these are *state* prediction errors (as there is no notion of reward in Active Inference).

Model-based methods of RL do not learn value functions or policies directly from experience. Rather, they learn a model of the task from experience, use the model to plan (simulate possible experiences), and update value functions and policies from these simulated experiences. While both Active Inference and reinforcement learning appeal to model-based planning, they use it differently. In Active Inference, planning is the computation of the expected free energy for each policy, not a means to update value functions. Arguably, if the expected free energy is seen as a value functional, it could be said that inferences drawn using the generative model are used to update this functional—offering a point of analogy between these approaches.

The third family of RL approaches, *policy gradient methods*, tries to optimize policies directly, without intermediate value functions, which are central to both model-based and model-free RL. These methods start from parameterized policies, able to generate (for example) movement trajectories, and then optimizes them by changing the parameters to increase (decrease) the likelihood of a policy if the trajectory results in a high (low) positive reward. This approach relates policy gradient methods to Active Inference, which also dispenses with value functions (Millidge 2019). However, the general objective of policy gradients (maximizing long-term cumulative reward) differs from Active Inference.

Besides the formal differences between Active Inference and RL, there are also several important conceptual differences. One difference regards how the two approaches interpret goal-directed and habitual behavior. In the animal learning literature, goal-directed choices are mediated by the (prospective) knowledge of the contingency between an action and its outcome (Dickinson and Balleine 1990), whereas habitual choices are not prospective and depend on simpler (e.g., stimulus-response) mechanisms. A popular idea in RL is that goal-directed and habitual choices correspond to *model-based* and *model-free* RL, respectively, and that these are acquired in parallel and continuously compete to control behavior (Daw et al. 2005).

Active Inference instead maps goal-directed and habitual choices to different mechanisms. In Active Inference (in discrete time), policy selection is quintessentially model-based and hence fits the definition of goal-directed, deliberative choices. This is similar to what happens in model-based RL, but with a difference. In model-based RL, actions are *selected* in a prospective

manner (using a model) but are *controlled* in a reactive way (using stimulus-response policies); in Active Inference, actions can be controlled in a proactive way—through fulfilling proprioceptive predictions (on action control, see section 10.6).

In Active Inference, habits can be acquired by executing goal-directed policies and then caching information about which policies are successful in which contexts. The cached information can be incorporated as a prior value of policies (Friston, FitzGerald et al. 2016; Maisto, Friston, and Pezzulo 2019). This mechanism permits executing policies that have a high prior value (in a given context) without deliberation. This can be thought of simply as observing “what I do” and learning that “I am the sort of creature that tends to do this” over multiple exposures to a task. In contrast to model-free RL, where habits are acquired independently of goal-directed policy selection, in Active Inference habits are acquired by repeatedly pursuing goal-directed policies (e.g., by caching their results).

In Active Inference, goal-directed and habitual mechanisms can cooperate rather than only compete. This is because the prior belief over policies depends on both a habitual term (a prior value of policies) and a deliberative term (expected free energy). Hierarchical elaborations of Active Inference suggest that reactive and goal-directed mechanisms could be arranged in a hierarchy rather than as parallel pathways (Pezzulo, Rigoli, and Friston 2015).

Finally, it is worth noting that Active Inference and RL differ subtly in how they conceive behavior and its causes. RL originates from behaviorist theory and the idea that behavior results from trial-and-error learning mediated by reinforcement. Active Inference assumes instead that behavior is the result of an inference. This leads us to the next point.

10.7.3 Planning as Inference

In the same way that it is possible to cast perceptual problems as problems of inference, it is also possible to cast control problems in terms of (approximate) Bayesian inference (Todorov 2008). In keeping with this, in Active Inference, planning is seen as an inferential process: the inference of a sequence of control states of the generative model.

This idea is closely related to other approaches, which include *control-as-inference* (Rawlik et al. 2013, Levine 2018), *planning-as-inference* (Attias 2003, Botvinick and Toussaint 2012), and *risk-sensitive* and *KL control* (Kappen et al. 2012). In these approaches, planning proceeds through inferring a posterior

distribution over actions, or sequences of actions, using a dynamic generative model that encodes probabilistic contingencies between states, actions, and future (expected) states. The best action or plan can be inferred by conditioning the generative model on observing future rewards (Pezzulo and Rigoli 2011, Solway and Botvinick 2012) or optimal future trajectories (Levine 2018). For example, it is possible to clamp (i.e., fix the value of) the future desired state in the model and then infer the sequence of actions that is more likely to fill the gap from the current state to the future desired state.

Active Inference, planning-as-inference, and other related schemes use a prospective form of control, which starts from an explicit representation of future, to-be-observed states rather than from a set of stimulus-response rules or policies, as is more typically done in optimal control theory and RL. However, the specific implementations of control- and planning-as-inference vary along at least three dimensions—namely, what form of inference they use (e.g., sampling or variational inference), what they infer (e.g., a posterior distribution over actions or action sequences), and the goal of inference (e.g., maximizing the marginal likelihood of an optimality condition or the probability of getting reward).

Active Inference takes a unique perspective on each of these dimensions. First, it uses a scalable approximate scheme—variational inference—to solve the challenging computational problems that arise during planning-as-inference. Second, it affords model-based planning, or the inference of a posterior over control states—which correspond to action sequences or policies, not single actions.² Third, to infer action sequences, Active Inference considers the *expected free energy* functional, which mathematically subsumes other widely used planning-as-inference schemes (e.g., KL control) and can handle ambiguous situations (Friston, Rigoli et al. 2015).

10.8 Behavior and Bounded Rationality

The wise are instructed by reason, average minds by experience, the stupid by necessity and the brute by instinct.

—Marcus Tullius Cicero

Behavior in Active Inference automatically combines multiple components: deliberative, perseverative, and habitual (Parr 2020). Imagine a person who is walking to a shop close to her house. If she predicts the consequences

of her actions (e.g., turning left or right), she can elaborate a good plan to reach the shop. This deliberative aspect of behavior is provided by *expected free energy*, which is minimized when one acts in a way to achieve preferred observations (e.g., being in the shop). Note that *expected free energy* also includes a drive to reduce uncertainty, which can manifest in deliberation. For example, if the person is unsure about the best direction, she can move to an appropriate vantage point, from which she can find the way to the shop easily, even if this implies a longer route. In short, her plans acquire epistemic affordance.

If the person is less able to engage in deliberation (e.g., because she is distracted), she may continue walking after reaching the shop. This perseverative aspect of behavior is provided by *variational free energy*, which is minimized when one gathers observations that are compatible with current beliefs, including beliefs about the current course of actions. The sensory and proprioceptive observations that the person gathers provide evidence for “walking” and hence may determine perseveration in the absence of deliberation.

Finally, another thing the person could do—when she is less able to deliberate—is select the usual plan to go home, without thinking about it. This habitual component is provided by the prior value of policies. This could allocate high probability to a plan to go home—a plan she has observed herself enacting multiple times in the past—and can become dominant if not superseded by deliberation.

Note that deliberative, perseverative, and habitual aspects of behavior coexist and can be combined in Active Inference. In other words, one can infer that, in this situation, a habit is the most likely course of action. This is different from “dual theories,” which assume that we are driven by two separate systems, one rational and one intuitive (Kahneman 2017). The mixture of deliberative, perseverative, and habitual aspects of behavior plausibly depends on contextual conditions, such as the amount of experience and the amount of cognitive resources one can invest in deliberative processes that may have a high complexity cost.³

The impact of cognitive resources on decision-making has been widely studied under the rubric of bounded rationality (Simon 1990). The core idea is that while an ideal rational agent should always fully consider the outcomes of its actions, a bounded rational agent has to balance the costs, effort, and timeliness of computation—for example, the information-processing costs of deliberating the best plan (Todorov 2009, Gershman et al. 2015).

10.8.1 Free Energy Theory of Bounded Rationality

Bounded rationality has been cast in terms of *Helmholtz free energy minimization*: a thermodynamic construct that is strictly related to the notion of variational free energy as used in Active Inference; see Gottwald and Braun (2020) for details. The “free energy theory of bounded rationality” formulates the trade-offs of action selection with limited information-processing capabilities in terms of two components of free energy: energy and entropy (see chapter 2). The former represents the expected value of a choice (an *accuracy* term), and the latter represents the costs of deliberation (a *complexity* term). What is costly during deliberation is decreasing the entropy (or complexity) of one’s beliefs before a choice to render them more precise (Ortega and Braun 2013, Zénon et al. 2019). Intuitively, the choice would be more accurate (and potentially entail higher utility) with a more precise posterior belief, but because increasing the precision of beliefs has a cost, a bounded decision-maker has to find a compromise—by minimizing free energy. The same trade-offs emerge in Active Inference, thus producing forms of bounded rationality. The notion of bounded rationality also resonates with the use of a variational bound on evidence (or marginal likelihood) that is a definitive aspect of Active Inference. In sum, Active Inference provides a model of (bounded) rationality and optimality, where the best solution to a given problem results from the compromise between complementary objectives: accuracy and complexity. These objectives stem from a normative (free energy minimization) imperative that is richer than classical objectives (e.g., utility maximization) usually considered in economic theory.

10.9 Valence, Emotion, and Motivation

Consider your origins: you were not made to live as brutes, but to follow virtue and knowledge.

—Dante Alighieri

Active Inference focuses on (negative) free energy as a measure of fitness and the capacity of an organism to realize its goals. While Active Inference proposes that creatures act to minimize their free energy, this does not mean that they ever have to compute it. Generally, it is sufficient to deal with the gradients of the free energy. By analogy, we do not need to

know our altitude to find the top of a hill but can simply follow the slope upward. However, some have suggested creatures may model how their free energy changes over time. Proponents of this hypothesis suggest that it might permit characterizations of phenomena like valence, emotion, and motivation.

On this view, it has been proposed that emotional valence, or the positive or negative character of emotions, can be conceived as the rate of change (first time-derivative) of free energy over time (Joffily and Coricelli 2013). Specifically, when a creature experiences an increase in its free energy over time, it may assign a negative valence to the situation; whereas when it experiences a decrease of its free energy over time, it may assign it a positive valence. Extending this line of thought to long-term dynamics of free energy (and second time-derivatives), it may be possible to characterize sophisticated emotional states; for example, the *relief* of passing from a phase of low valence to a phase of high valence, or the *disappointment* of passing from a phase of high valence to a phase of low valence. Monitoring free energy dynamics (and the emotional states they elicit) may permit adapting the behavioral strategies or learning rates to long-term environmental statistics.

It may seem a bit of a leap to assume a second generative model whose role is to monitor the free energy of the first. However, there is another way in which these ideas can be interpreted. An interesting formalization of these perspectives rests on thinking about what causes rapid changes in free energy. As it is a functional of beliefs, a rapid change in free energy must be due to fast belief updating. The key determinant of this speed is precision, which acts as a time-constant in the dynamics of predictive coding. Interestingly, this ties in with the notion of higher derivatives of the free energy, as precision is the negative of the second derivative (i.e., the curvature of a free energy landscape). However, this begs the question as to why we should associate precision with valence. The answer comes from noticing that precision is inversely related to ambiguity. The more precise something is, the less ambiguous its interpretation. Choosing a course of action that minimizes expected free energy also means minimizing ambiguity and therefore maximizing precision. Here we see a direct association between high order derivatives of the free energy, its rate of change, and motivated behavior.

Expectations about (increases or decreases of) free energy may play motivational roles and incentivize behavior, too. In Active Inference, a surrogate

expectation about changes (increases or decreases) of free energy is the precision of beliefs about policies. This again highlights the importance of this second order statistic. For example, a highly precise belief signals that one has found a good policy—that is, a policy that can be confidently expected to minimize free energy. Interestingly, the precision of (beliefs about) policies has been linked to dopamine signaling (FitzGerald, Dolan, and Friston 2015). From this perspective, stimuli that increase the precision of beliefs about policies trigger dopamine bursts—which may indicate their incentive salience (Berridge 2007). This perspective may help shed light on the neurophysiological mechanisms linking expectations of goal or reward achievement to increases in attention (Anderson et al. 2011) and motivation (Berridge and Kringelbach 2011).

10.10 Homeostasis, Allostasis, and Interoceptive Processing

There is more wisdom in your body than in your deepest philosophy.

—Friedrich Nietzsche

A creature's generative model is not just about the external world but also—and perhaps even more importantly—about the internal milieu. A generative model of a body's inside (or interoceptive schema) has a dual role: to explain how interoceptive (bodily) sensations are generated and to ensure the correct regulation of physiological parameters (Iodice et al. 2019), like body temperature or sugar levels in the blood. Cybernetic theories (touched on in section 10.6.2) assume that a central objective of living organisms is maintaining homeostasis (Cannon 1929)—ensuring that physiological parameters remain within viable ranges (e.g., body temperature never becomes too high)—and that homeostasis can only be achieved by exerting a successful control over the environment (Ashby 1952).

This form of homeostatic regulation can be achieved in Active Inference by specifying the viable ranges of physiological parameters as priors over interoceptive observations. Interestingly, homeostatic regulation can be achieved in multiple, nested ways. The simplest regulatory loop is the engagement of autonomic reflexes (e.g., vasodilation), when certain parameters are (expected to be) out of range—for example, when body temperature is too high. This autonomic control can be constructed as *interoceptive inference*: an Active Inference process that operates on interoceptive streams

rather than proprioceptive streams, as in the case of externally directed actions (Seth et al. 2012, Seth and Friston 2016, Allen et al. 2019). For this, the brain may use a generative model that predicts interoceptive and physiological streams and triggers autonomic reflexes to correct interoceptive prediction errors (e.g., a surprisingly high body temperature). This is analogous to the way motor reflexes are activated to correct proprioceptive prediction errors and steer externally directed actions.

Active Inference extends beyond simple autonomic loops: it can correct the same interoceptive prediction error (high body temperature) in increasingly sophisticated ways (Pezzulo, Rigoli, and Friston 2015). It can use predictive, *allostatic* strategies (Sterling 2012, Barrett and Simmons 2015, Corcoran et al. 2020) that go beyond homeostasis and preemptively control physiology in an allostatic fashion before interoceptive prediction errors are triggered—for example, finding shade before overheating. Another predictive strategy entails mobilizing resources before expected excursions from physiological setpoints—for example, increasing cardiac output before a long run in anticipation of increased oxygen demands. That requires modifying the priors over interoceptive observations dynamically, going beyond homeostasis (Tschantz et al. 2021). Eventually, predictive brains can develop sophisticated goal-directed strategies, such as ensuring that one brings cold water to the beach, meeting the same imperative (controlling body temperature) in richer and more effective ways.

Biological and interoceptive regulation may be crucial for affect and emotional processing (Barrett 2017). During situated interactions, the brain's generative model constantly predicts not just what will happen next but also what the consequences for interoception and allostasis are. Interoceptive streams—elicited during the perception of external objects and events—imbue them with an affective dimension, which signals how good or bad they are for the creature's allostasis and survival, hence making them “meaningful.” If this view is correct, then disorders of this interoceptive and allostatic processing may engender emotional dysregulation and various psychopathological conditions (Pezzulo 2013; Barrett et al. 2016; Maisto, Barca et al. 2019; Pezzulo, Maisto et al. 2019).

There is an emerging bedfellow for *interoceptive inference*—namely, *emotional inference*. In this application of Active Inference, emotions are considered part of the generative model: they are just another construct or hypothesis that the brain employs to deploy precision in deep generative

models. From the perspective of belief updating, this means anxiety is just a commitment to the Bayesian belief “I am anxious” that best explains the prevailing sensory and interoceptive queues. From the perspective of acting, the ensuing (interoceptive) predictions augment or attenuate various precisions (i.e., covert action) or enslave autonomic responses (i.e., overt action). This may look much like arousal, which confirm the hypothesis that “I am anxious.” Usually, emotional inference entails belief updating that is domain general, assimilating information from both interoceptive and exteroceptive sensory streams—hence the intimate relationship between emotion, interoception, and attention in health (Seth and Friston 2016; Smith, Lane et al. 2019; Smith, Parr, and Friston 2019) and disease (Peters et al. 2017, J. E. Clark et al. 2018).

10.11 Attention, Salience, and Epistemic Dynamics

True ignorance is not the absence of knowledge, but the refusal to acquire it.
—Karl Popper

Given the number of times we have referred to *precision* and *expected free energy* in this chapter alone, it would be negligent not to devote a little space to *attention* and *salience*. These concepts recur throughout psychology, having been subject to numerous redefinitions and classifications. Sometimes these terms are used to refer to synaptic gain control mechanisms (Hillyard et al. 1998), which preferentially select some sensory modality or subset of channels within a modality. Sometimes they refer to how we orient ourselves, through overt or covert action, to gain more information about the world (Rizzolatti et al. 1987; Sheliga et al. 1994, 1995).

Although the uncertainty afforded by the many meanings of *attention* underwrites some of the epistemic attractiveness of this field of study, there is also value in resolving the attendant ambiguity. One of the things offered by a formal perspective on psychology is that we do not need to worry about this ambiguity. We can operationally define *attention* as the precision associated with some sensory input. This neatly maps to the concept of gain control, as sensations we infer to be more precise will have greater influence over belief updating than those inferred to be imprecise. The construct validity of this association has been demonstrated in relation to psychological paradigms, including the famous Posner paradigm (Feldman and Friston

2010). Specifically, responding to a stimulus at a location in visual space that is afforded a higher precision is faster than responding to stimuli in other locations.

This leaves the term *salience* in want of a similar formal definition. Typically, in Active Inference, we associate salience with expected information gain (or epistemic value): a component of the expected free energy. Intuitively, something is more salient when we expect it to yield more information. However, this defines *salience* in terms of an action or policy, while *attention* is an attribute of beliefs about sensory input. This fits with the notion of salience as overt or covert orienting. We saw in chapter 7 that we could further subdivide expected information gain into salience and novelty. The former is the potential to infer, while the latter is the potential to learn. An analogy that expresses the difference between attention and salience (or novelty) is the design and analysis of a scientific experiment. Attention is the process of selecting the highest quality data from what we have already measured and using these to inform our hypothesis testing. Salience is the design of the next experiment to ensure the highest quality data.

We do not dwell on this issue to simply add another reclassification of attentional phenomena to the literature but to highlight an important advantage in committing to a formal psychology. Under Active Inference, it does not matter if others define *attention* (or any other construct) differently—as we can simply refer to the mathematical constructs in question and preclude any confusion. A final point of consideration is that these definitions offer a simple explanation for why attention and salience are so often conflated. Highly precise data are minimally ambiguous. This means that they should be afforded attention and that actions to acquire these data are highly salient (Parr and Friston 2019a).

10.12 Rule Learning, Causal Inference, and Fast Generalization

Yesterday I was clever, so I wanted to change the world. Today I am wise, so I am changing myself.

—Rumi

Humans and other animals excel at making sophisticated causal inferences, learning abstract concepts and the causal relationships between objects, and generalizing from limited experience—in contrast to current machine

learning paradigms, which require a large number of examples to attain similar performance. This difference suggests that current machine learning approaches, which are largely based on sophisticated pattern recognition, may not fully capture the ways humans learn and think (Lake et al. 2017).

The learning paradigm of Active Inference is based on the development of generative models that capture the causal relations between actions, events, and observations. In this book, we have considered relatively simple tasks (e.g., the T-maze example of chapter 7) that require unsophisticated generative models. In contrast, understanding and reasoning about complex situations require deep generative models that capture the *latent structure* of the environment—such as hidden regularities that permit generalizing across a number of apparently dissimilar situations (Tervo et al. 2016; Friston, Lin et al. 2017).

One simple example of a hidden rule that governs sophisticated social interactions is a traffic intersection. Imagine a naïve person who observes a busy crossroad and has to predict (or explain) on which occasions pedestrians or cars cross the road. The person can accumulate statistics about the co-occurrence of events (e.g., a red car stopping and a tall man crossing; an old woman stopping and a big car passing), but most are ultimately useless. The person can eventually discover some recurrent statistical patterns, such as that pedestrians cross the road soon after all cars stop at a certain point on the road. This determination would be deemed sufficient in a machine learning setting if the task were just to predict when pedestrians are about to walk, but it would not entail any understanding of the situation. In fact, it may even lead to the erroneous conclusion that the stopping of cars explains the movement of pedestrians. This sort of error is typical in machine learning applications that do not appeal to (causal) models—and cannot distinguish whether the rain explains the wet grass or the wet grass explains the rain (Pearl and Mackenzie 2018).

On the other hand, inferring the correct hidden (e.g., traffic light) rule provides a deeper understanding of the causal structure of the situation (e.g., it is the traffic light that causes the cars to stop and the pedestrians to walk). The hidden rule not only affords better predictive power but also renders inference more parsimonious, as it can abstract away from most sensory details (e.g., the color of cars). In turn, this permits generalizing to other situations, such as different crossroads or cities, where most sensory details differ significantly—with the caveat that facing crossroads in some

cities, like Rome, may require more than looking at traffic lights. Finally, learning about traffic light rules may also enable more efficient learning in novel situations—or to develop what is called a “learning set” in psychology or a *learning-to-learn* ability in machine learning (Harlow 1949). When facing a crossroad where the traffic light is off, one cannot use the learned rule but may nevertheless have the expectation that there is another, similar hidden rule in play—and this could help understanding what the traffic police officer is doing.

As this simple example illustrates, learning rich generative models—of the latent structure of the environment (aka structure learning)—affords sophisticated forms of causal reasoning and generalization. Scaling up generative models to address these sophisticated situations is an ongoing objective in computational modeling and cognitive science (Tenenbaum et al. 2006, Kemp and Tenenbaum 2008). Interestingly, there is a tension between current machine learning trends—wherein the general idea is “the bigger, the better”—and the statistical approach of Active Inference—which suggests the importance of balancing the *accuracy* of a model with its *complexity* and to favor simpler models. Model reduction (and the pruning of unnecessary parameters) is not simply a way to avoid wasting resources—it is also an effective way to learn hidden rules, including during offline periods like sleep (Friston, Lin et al. 2017), perhaps manifesting in resting state activity (Pezzulo, Zorzi, and Corbetta 2020).

10.13 Active Inference and Other Fields: Open Directions

It has to start somewhere, it has to start sometime,
what better place than here? What better time than now?

—Rage Against the Machine, “Guerrilla Radio”

In this book, we mainly focus on Active Inference models that address biological problems of survival and adaptation. Yet Active Inference can be applied in many other domains. In this last section, we briefly discuss two such domains: social and cultural dynamics and machine learning and robotics. Addressing the former requires thinking about the ways in which multiple Active Inference agents interact and the emergent effects of such interaction. Addressing the latter requires understanding how Active Inference can be endowed with more effective learning (and inference) mechanisms to scale

up to more complex problems—but in a way that is compatible with the basic assumptions of the theory. Both are interesting open directions for research.

10.13.1 Social and Cultural Dynamics

Many interesting aspects of our (human) cognition relate to social and cultural dynamics rather than individualistic perceptions, decisions, and actions (Veissière et al. 2020). By definition, social dynamics require multiple Active Inference creatures that engage in physical interactions (e.g., joint actions, such as playing team sports) or more abstract interactions (e.g., elections or social networking). Simple demonstrations of inter-Active Inference between identical organisms already produced interesting emergent phenomena, such as the self-organization of simple life forms that resist dispersion, the possibility to engage in morphogenetic processes to acquire and restore a body form, and mutual coordinated prediction and turn taking (Friston 2013; Friston and Frith 2015a; Friston, Levin et al. 2015). Other simulations have addressed the ways in which creatures can extend their cognition to material artifacts and shape their cognitive niches (Bruineberg et al. 2018).

These simulations capture only a fraction of the complexity of our social and cultural dynamics, but they illustrate the potential of Active Inference to expand from a science of individuals to a science of societies—and how cognition extends beyond our skulls (Nave et al. 2020).

10.13.2 Machine Learning and Robotics

The generative modeling and variational inference methods discussed in this book are widely used in machine learning and robotics. In these fields, the emphasis is often on how to learn (connectionist) generative models—as opposed to how to use them for Active Inference, the focus of this book. This is interesting as machine learning approaches are potentially useful to scale up the complexity of the generative models and of the problems considered in this book—with the caveat that they may call on very different process theories of Active Inference.

While it is impossible to review here the vast literature on generative modeling in machine learning, we briefly mention some of the most popular models, from which many variants have been developed. Two early connectionist generative models, the *Helmholtz machine* and the *Boltzmann machine* (Ackley et al. 1985, Dayan et al. 1995), provided paradigmatic

examples of how to learn the internal representations of a neural network in an *unsupervised* way. The Helmholtz machine is especially related to the variational approach of Active Inference, as it uses separate *recognition* and *generative* networks to infer a distribution over hidden variables and sample from them to obtain fictive data. The early practical success of these methods was limited. But afterward, the possibility to stack multiple (restricted) Boltzmann machines enabled learning of multiple layers of internal representations and was one of the early successes of unsupervised deep neural networks (Hinton 2007).

Two recent examples of connectionist generative models, *variational auto-encoders* or VAEs (Kingma and Welling 2014) and *generative adversarial networks* or GANs (Goodfellow et al. 2014), are widely used in machine learning applications, such as recognizing or generating pictures and videos. VAEs exemplify an elegant application of variational methods to learning in generative networks. Their learning objective, the evidence lower bound (ELBO), is mathematically equivalent to variational free energy. This objective enables learning of an accurate description of the data (i.e., maximizes accuracy) but also favors internal representations that do not differ too much from their priors (i.e., minimizes complexity). The latter objective acts as a so-called regularizer, which helps to generalize and avoid overfitting.

GANs follow a different approach: they combine two networks, a *generative network* and a *discriminative network*, which continuously compete during learning. The discriminative network learns to distinguish which example data produced by the generative network are *real* or *fictive*. The generative network tries to generate fictive data that fool (i.e., are misclassified by) the discriminative network. The race between these two networks forces the generative network to improve its generative capabilities and produce high fidelity fictive data—an ability that has been widely exploited to generate, for example, realistic images.

The above generative models (and others) can be used for control tasks. For example, Ha and Eck (2017) have used a (sequence-to-sequence) VAE to learn to predict pencil strokes. By sampling from the internal representation of the VAE, the model can construct novel stroke-based drawings. Generative modeling approaches have been used to control robot movements, too. Some of these approaches use Active Inference (Pio-Lopez et al. 2016, Sancaktar et al. 2020, Ciria et al. 2021) or closely related ideas, but in a connectionist setting (Ahmadi and Tani 2019, Tani and White 2020).

One of the main challenges in this domain is that robot movements are high dimensional and require (learning) sophisticated generative models. One interesting aspect of Active Inference and related approaches is that the most important thing to be learned is a forward mapping between actions and sensory (e.g., visual and proprioceptive) feedback at the next time step. This forward mapping can be learned in various ways: by autonomous exploration, by demonstration, or even by direct interaction with a human—for example, a teacher (the experimenter) who guides the hands of the robot along a trajectory to the goal, hence scaffolding the acquisition of effective goal-directed actions (Yamashita and Tani 2008). The possibility to learn generative models in various ways greatly expands the scope of robot skills that can be eventually achieved. In turn, the possibility to develop more advanced (neuro-) robots using Active Inference could be important not just for technological but also for theoretical reasons. Indeed, some key aspects of Active Inference, such as the adaptive agent-environment interactions, the integration of cognitive functions, and the importance of embodiment, are naturally addressed in robotic settings.

10.14 Summary

Home is behind, the world ahead,
and there are many paths to tread
through shadows to the edge of night,
until the stars are all alight.

—J. R. R. Tolkien, *The Lord of the Rings*

We started this book by asking whether it is possible to understand brain and behavior from first principles. We then introduced Active Inference as a candidate theory to meet this challenge. We hope that the reader has been convinced that the answer to our original question is yes. In this chapter, we considered the unified perspective that Active Inference offers on sentient behavior and what implications this theory has for familiar psychological constructs, such as perception, action selection, and emotion. This gave us the opportunity to revisit the concepts introduced throughout the book and to remind ourselves of the fascinating questions still open for future research. We hope this book provides a useful complement to related works on Active Inference, including on the one hand the philosophy (Hohwy 2013, Clark 2015) and on the other hand the physics (Friston 2019a).

We are now at the end of our journey. Our aim has been to offer an introduction to those interested in using these methods—both at conceptual and formal levels. However, it is important to emphasize that Active Inference is not something that can be learned purely in theory. We encourage anyone who has enjoyed this book to think about pursuing it in practice. Important rites of passage in theoretical neurobiology are trying to write down a generative model, experiencing the frustration when simulations misbehave, and learning from violations of your prior beliefs when something unexpected happens. Whether or not you choose to pursue this practice at a computational level, we hope that you will reflect on it as you engage in Active Inference in day-to-day life. This may manifest in the compulsion to direct your eyes to resolve uncertainty about something in your peripheral vision. It may be in choosing to eat at a favorite restaurant to fulfill prior (gustatory) preferences. It may be in reducing the heat when the shower is too hot to ensure the temperature conforms to your model of how the world should be. Ultimately, we are confident that you will continue to pursue Active Inference in some form.