Anti-intellectualism, instructive representations, and the intentional action argument

Alison Ann Springle[1] · Justin Humphreys[2]

© The Author(s) 2021

Abstract

Intellectualists hold that knowledge-how is a species of knowledge-that, and consequently that the knowledge involved in skill is propositional. In support of this view, the intentional action argument holds that since skills manifest in intentional action and since intentional action necessarily depends on propositional knowledge, skills necessarily depend on propositional knowledge. We challenge this argument, and suggest that instructive representations, as opposed to propositional attitudes, can better account for an agent's reasons for action. While a propositional-causal theory of action, according to which intentional action must be causally produced "in the right way" by an agent's proposition-involving reasons, has long held sway, we draw on Elizabeth Anscombe's insights offer a path toward an alternative theory of action. In so doing, we reject the implicitly Cartesian conception of knowledge at the core of the intentional action argument, while hanging on to the idea that mental states are representations of a certain kind. Our argument provides theoretical support for anti-intellectualism by equipping philosophers with an account of non-propositional, practical content.

> *Can it be that there is something that modern philosophy has blankly misunderstood: namely what ancient and medieval philosophers meant by practical knowledge? Certainly in modern philosophy we have an incorrigibly contem-*

✉ Alison Ann Springle
  alisonspringle@gmail.com

  Justin Humphreys
  Justin.humphreys@villanova.edu

1 University of Tübingen, Baden-Württemberg, Germany

2 Villanova University, Philadelphia, PA, USA

plative conception of knowledge. Knowledge must be something that is judged as such by being in accordance with the facts. The facts, reality, are prior, and dictate what is to be said, if it is knowledge. And this is the explanation of the utter darkness in which we found ourselves. For if there are two knowledges—one by observation, the other by intention—then it looks as if there must be two objects of knowledge; but if one says the objects are the same, one looks hopelessly for the different mode of contemplative knowledge in acting, as if there were a very queer and special sort of seeing eye in the middle of the acting.

(G.E.M. Anscombe, *Intention.* 1957 §62 pp. 57)

1 Introduction

Many philosophers hold that skill requires knowledge-how and some hold the converse, as well (Pavese, 2016a, 2016b). According to intellectualism about knowledge-how, it is just a species of propositional knowledge (Ginet, 1990; Stanley & Williamson, 2001, Snowdon, 2004; Stanley 2011, Pavese, 2015a, 2015b). According to intellectualism about skill, the knowledge involved in skill is propositional (Pavese, 2016a, 2016b; Stanley & Krakauer, 2013).

Intellectualists about skill frequently appeal to what we call the "Intentional Action Argument" (henceforth "IA") (see e.g. Pavese, 2016a, 2016b; Stanley & Krakauer, 2013). According to IA, (P1) skills necessarily manifest in intentional action and (P2) intentional action necessarily depends on propositional knowledge, so (C) skills (at least their manifestation) necessarily depend on propositional knowledge. In what follows, we grant the intellectualist P1 and dispute P2.

Philosophical discussions of intentional action often appeal to Elizabeth Anscombe's seminal treatment of the topic in *Intention* (1957) and the IA for intellectualism about skill is no exception. According to Anscombe, when one is φ-ing intentionally, one has practical knowledge of one's φ-ing: one's practical knowledge of what one is doing is constituted by one's intention in doing it. One way of arriving at P2 is to construe this knowledge propositionally, viz as the knowledge *that* one is φ-ing (Gibbons, 2001). Alternatively, given the common view that an agent's act is intentional iff it is appropriately reason-guided together with the common view that such reasons are propositional attitudes, one might argue that an agent's practical knowledge is propositional by means of an inference to the best explanation (Goldman, 1968; Bratman, 1987; Butterfill & Sinigaglia, 2014; Hawley, 2003; Mylopoulos & Pacherie 2019; Levy, 2017; Pavese, 2015a, 2015b, 2016a, 2016b, 2017a, 2017b, 2019). However exactly one establishes P2, IA concludes that skill essentially depends on propositional knowledge.[1]

---

[1] Differences in terminology present a difficulty of assessing the literature on skill, and indeed that on intentional action more generally. As we employ the terms, skill is necessarily manifested in intentional action, so that accidentally successful action or unintentional action cannot count as skillful. When an agent reliably manifests a skill, we say that she has knowledge-how to bring about the end intended by the action.

Some philosophers take IA to support *strong intellectualism*, according to which the propositional knowledge component exhausts the intelligent aspects of skill. On this view, the epistemic aspects of skill amounts to knowledge-that (e.g. Pavese, 2016a, 2016b; Stanley & Krakauer, 2013; Stanley & Williamson, 2001). Others take it to support only *weak intellectualism*: propositional knowledge is necessary for but does not exhaust the epistemic or intelligent dimensions of skill (e.g. Fridland, 2014, 2015a, b, 2017; Levy, 2017). However, if one accepts IA, it's an uphill battle to defend weak intellectualism. For if skill, understood as a manifestation of knowledge-how, necessarily involves propositional knowledge, parsimony, together with the desire for a unified theory of knowledge speaks against positing a further *epistemic* kind to account for skill and in favor of accounting for the distinguishing features of skill and ability knowledge more generally in terms of a "motor acuity" that's grounded in dispositions or abilities that are not intrinsically epistemic (Stanley & Krakauer, 2013; See Springle, 2019 for response), or in terms of special kinds of "modes of presentation" of propositions (Stanley, 2011; Pavese, 2015b; See Springle, 2019 for response).[2] And in any case, IA certainly threatens the "hardcore" anti-intellectualism we personally prefer, according to which propositional knowledge ultimately depends on non-propositional, practical knowledge.[3]

Our argument against IA will target P2. We'll block the inference to the best explanation defense of the premise by arguing that "embodied instructive representations," as opposed to propositional attitudes, may better account for an agent's reasons for action. We'll block direct defense of the premise by arguing that embodied instructive representations, as opposed to propositional knowledge, may better account for practical knowledge, and the psychological states that constitute intentions more generally.

The indirect version of IA presupposes what we'll call the "Propositional-Causal Theory of Action" (henceforth PCT) according to which intentional action is constitutively an action that is causally produced "in the right way" by an agent's subjective reasons construed as propositional attitudes. Thus, the indirect argument imports into IA substantive assumptions about the nature of reasons for action (henceforth "REASONS"), and about the connection between reasons for action and intentional actions (henceforth "CONNECTION"). These assumptions are summed up in the theses that an agent's subjective reasons for acting are her propositional representations of facts, and that the intentionality of actions derives from their being caused by an agent's propositional attitudes.

These assumptions seem natural to many philosophers, we think, because they encode an implicitly Cartesian conception of *all* knowledge as reflection. Ryle (1949) famously challenged this conception of knowledge but is often understood to have replaced it with behaviorism. To this day, dispositional accounts of the

---

[2] Springle (2019) provides a different but related argument against the IA.

[3] It may not be that every instance of propositional knowledge *reduces* to non-propositional practical knowledge. Rather, our claim is that the possibility of propositional knowledge depends on the existence of non-propositional practical knowledge. This is compatible with there being some species of non-propositional practical knowledge that couldn't exist without propositional knowledge.

mind like Ryle's are widely considered to be anti-representational. Consequently, for philosophers who find it plausible that knowledge essentially involves psychological representations and are thus attracted to P2, it looks like it's PCT or bust. Indeed, despite joining Ryle in launching a neo-Aristotelian counteroffensive on the Cartesian conception of knowledge, Elizabeth Anscombe's views are commonly employed to support IA, and thus to defend PCT. However, we think that PCT and IA fail to respect some of the main insights of Anscombe's account of intentional action, which we can state in five quite general theses[4]:

1. A fruitful distinction can be made between speculative knowledge "derived from the objects known," and practical knowledge, which "is the cause of what it understands." (Anscombe, 1957, p. 87).
2. Some of the causes of an agent's movements can be known by that agent *without observation*, that is, in a direct, first-personal way that requires no gathering of empirical information. (Anscombe, 1957, p. 15).
3. An error in practical judgment is distinct from an error in the performance of an action (Anscombe, 1957, p. 56).
4. Since both propositions and proposals can be affirmed or denied, there are different success states for contemplation and practical judgment. A propositional judgment is true when it is agreement with the facts; in distinction, practical truth is "truth in agreement with right desire" (Anscombe, 1981, p. 76).
5. The "incorrigibly contemplative conception of knowledge" in modern philosophy leads to "utter darkness" in our understanding of action, since it makes the futile effort of attempting to conceive of practical knowledge on the model of the contemplation of facts. (Anscombe, 1957, p. 57).

If we are right, Anscombe's insights offer a path toward an alternative theory of action that denies the problematic aspects of PCT and the Cartesian conception of knowledge it encodes while hanging on to the idea that intentional actions essentially involve psychological representations of a certain kind.[5]

We begin in Sect. 2 by motivating the way the PCT addresses REASONS and CONNECTION. We review PCT's well-known susceptibility to an objection from deviant causal chains and resultant failure to provide sufficient conditions for an action being intentional (Davidson, 1980; Setiya, 2011a, b). We then raise a new objection to how PCT answers REASONS and CONNECTION, and argue that there's reason to desire an alternative to the PCT that does not identify causation via psychological representations as the *constitutive* property of intentional actions.

In Sect. 3 we note that what we call the "Aristotelian-Anscombean Theory of Action" (henceforth AAT), as developed by Michael Thompson (2008), avoids the objection from deviant causal chains by defining intentional action in terms of a

---

[4] Anscombe exegesis can be quite contentious. Unfortunately, there is not space in this paper for a defense of our reading but where possible we refer to relevant primary texts.

[5] Setiya (2008, 2009, 2012, 2016) also defends anti-intellectualism about practical knowledge and knowledge-how on the basis of Anscombe-inspired considerations.

kind of teleological form rather than in terms of an etiological relation to propositional attitudes. However, AAT rejects PCT's identification of REASONS and CONNECTION with propositional attitudes (qua psychological representations). We suggest that the absence of a well worked-out alternative to the propositional attitude model of psychological representation is part of what motivates Thompson to identify reasons with the larger actions for the sake of which an intentional action is done, thereby endorsing a broadly "behaviorist" account of REASONS and CONNECTION.[6] We argue that it would be preferable to develop an alternative to the propositional attitude model of psychological representation that's compatible with AAT. In order to argue that AAT is consistent with an alternative to the propositional attitude model, we begin by sketching Springle's (2021) "problem-solving" analysis of intentional action. Like other versions of AAT, the "problem-solving" analysis avoids defining intentional actions in terms of their etiological relation to psychological representations, but rather defines intentional actions in terms of their form as *solutions to practical problems.*

In Sect. 4, we sketch Springle's (2021) Practical-Epistemic Access (PEA) analysis of representation. According to this account, psychological representations (intentions, beliefs, desires) function to provide PEA to situations (facts), where to have PEA to a situation just is to be in a position to produce an intentional action in response to it. Having PEA is, in other words, having and exercising a practical problem-solving ability. *Direct* PEA grounds perceptions, beliefs, desires, and intentions understood as *embodied instructive representations*, while *indirect* PEA grounds propositions understood as *surrogative representations*. Embodied instructive representations are not essentially attitudes towards propositions, though they can be. They are realized by *problem-solving dispositions*, instantiate a non-propositional "*generative*" species of intentional aboutness or directedness, and have constitutive *appropriateness conditions* rather than truth, veridicality, accuracy, or satisfaction conditions.[7]

In Sect. 5 we connect embodied instructive representations with Anscombe's theses and to a Neo-Aristotelian project of developing an hylomorphic alternative to the Cartesian conception of practical knowledge.

Finally, in Sect. 6 we argue that the account developed in this paper avoids the objection from deviant causal chains as well as the novel objection to PCT, while offering satisfying interpretations of REASONS and CONNECTION. Because embodied instructive representations are themselves essentially abilities to produce intentional actions, they are compatible with AAT, but at the same time preserve the notion that intentional actions are *necessarily* produced by psychological

---

[6] See Setiya (2009) footnote 36.

[7] This view has similarities with Gibson's theory of affordances (1979), Millikan's "pushmi-pullyu representations" (1995); Siegel's "experienced mandates" (2014); Nanay's "pragmatic representations" (2013); Anderson and Rosenberg's "guidance theory of representation" (2008); Cussins's "cognitive trails" (1992); Grush's "skill theory" (2007); Cummins's discussion of "Representations, Targets, and Attitudes" (1996); Pavese's "practical senses" (2015) and "practical representations" (2019); Mandik's "action-oriented representations" (2005); Rowlands's (2006) "deeds"; and to some extent O'Regan and Noe's "sensori-motor contingencies" (2001) and Castañeda's (1975) "practitions."

representations. If successful, our argument provides a significant theoretical boost for hardcore anti-intellectualism. For not only does it block IA, it also equips anti-intellectualism with an account of non-propositional content. We think this upshot offers a powerful basis for the defense of anti-intellectualism.

2 The propositional-causal theory of action

According to PCT, an agent's reasons for action are (facts as they are represented by) her propositional attitudes, in particular, her means-end beliefs, desires, or intentions. This is how PCT addresses REASONS. Similarly, according to PCT, intentional actions are constitutively caused by an agent's reasons, understood as propositional attitudes (propositional representations of the facts that are her reasons). This is how PCT addresses CONNECTION.

PCT's interpretation of REASONS and CONNECTION are intuitive. You grind coffee beans because you want to make a cup of coffee and you believe that grinding coffee beans is a means to this end. This is one way in which an agent's subjective reasons—the ends an agent's desires represent together with what her perceptual and means-ends beliefs represent—explain her intentional actions. And if they explain intentional actions, plausibly that's because they are constitutive of them. At the very least, an agent's subjective reasons don't appear to be accidentally connected to her actions.

Anscombe herself suggests one reason to think that they are not accidentally connected: intentional actions are those to which the question "why?" has application. We typically answer such questions by citing how we take the world to be (our beliefs and perceptions) and how we want it to be (our desires). What's more, intentional actions can be modeled by an Aristotelian practical syllogism, the premises of which map onto desires and means-end beliefs, while the conclusion appears to map onto the intention (Anscombe, 1957; Davidson, 1963, 1970; Goldman, 1970; Audi, 1986). Since they have the form of premises in an Aristotelian practical syllogism, propositional attitudes seem to account for the ways in which an agent's mental representations combine to rationalize an intentional action. This is the beauty of the modal separability of the propositional and attitude components of propositional attitudes. Attitudes are defined in terms of causal roles, while propositional contents are defined in terms of truth or satisfaction conditions and can be shared by different attitudes. And if an agent's reasons for action are her beliefs and desires, and these are propositional attitudes, then so are intentions and practical knowledge.

Alas, PCT can only provide necessary but not sufficient conditions for an action's counting as intentional because it is vulnerable to deviant causal chains. On PCT, the intentionality of action is derived from the intentionality of an agent's subjective reasons, conceived as propositional attitudes. And because propositional attitudes and contents are modally separable, it is possible to devise scenarios in which:

(i)   The agent intends to φ,
(ii)   What happens is in conformity with the agent's representation of what is to happen, and

(iii) That event was caused by the intention to φ, but where what happens is accidental rather than intentional.

Consider, for instance, the following example from Alfred Mele:

> A philosopher intends to knock over his glass of water in order to distract his commentator. However, his intention so upsets him that his hand shakes uncontrollably, striking the glass and knocking it to the floor. (1992, p. 182).

As a result of deviant causal chain examples like this, PCT must say that intentional actions are constitutively caused by an agent's propositional attitudes "in the right way." But what does "in the right way" amount to? Some philosophers believe it can be cashed out in terms of a causal account of the way in which a propositional attitude controls an action.[8] Others, like Anscombe,[9] believe that PCT's susceptibility to such counterexamples is a symptom of treating the intentionality of action as causally derivative.[10] Some of these latter philosophers endorse what we're calling the AAT according to which actions are *intrinsically* intentional by virtue of their teleological form.[11] We'll return to AAT in the next section.

PCT is also susceptible to an additional objection, which as far as we are aware, has not been noticed in the literature on intentional action.[12] In order to launch this new objection to PCT, we assume that the property in virtue of which anything represents is functional in nature. This "Functional Analysis Assumption" is motivated by the fact that most philosophers hold that representations have constitutive success conditions, so that a necessary condition on something's being a representation is that it be in principle capable of *mis*representing. If this is right, it's plausible that the constitutive property of representations is functional.[13] As we have seen, according to the PCT, intentional actions must be analyzed in terms of mental representations causing them in the right way. Thus, if an account of mental representation must appeal to intentional actions, then not only is PCT incomplete, since it cannot specify what the "right way" is, but also, we hold, viciously circular.

Our argument turns, first, on a general distinction between constitutive and necessary properties. We take it that an analysis of a kind F lists the constitutive properties

---

[8] See e.g. Peacocke (1979), Bishop (1989); Setiya (2007, pp. 32–3); Frankfurt (1978).

[9] See Anscombe (1989, pp. 110–1).

[10] See e.g. Wald and Tenenbaum (2018).

[11] There are causalist teleological accounts of action. See e.g. Wilson (1989), Ginet (1990) and see Sehon (1997, 2005) for criticism. We are sympathetic with Roth (2000) who argues that reasons explanations may be irreducibly teleological *and* at once cite reasons as efficient causes. Causalists views do not obviously avoid the problem of deviant causal chains. For an example of an anti-causalist, see McLaughlin (2012). We focus on a version of AAT developed by Thompson (2008) as we consider his view, more than other teleological views, an exemplar of the Aristotelian-Anscombean theory. Thompson holds that the debate between causalists and anti-causalists is misconceived. As we read Thompson, this is because he has an hylomorphic account of agency and intentional action. The view we'll develop (based on Springle 2021) may be understood as attempting to combine insights from both Thompson and Roth. Springle's view is also indebted to Stout (1996, 2014).

[12] Though see Stout (2014) for a similar worry about circularity. See Springle (2021) for a longer version of the objection developed here.

[13] See also Burge (2010), Millikan (1984) and Dretske (1988).

that individuate Fs. Other properties may be necessary to Fs but not constitutive. Such properties need not show up in an analysis of Fs, but an adequate theory of Fs should ultimately illuminate the connection between the constitutive properties of Fs and those that are necessary but not constitutive. Second, observe that PCT is incompatible with theories of what constitutes mental representation that reference intentional actions. This is because PCT is itself an analysis of what constitutes intentional action that makes reference to mental representations (beliefs, desires, intentions). PCT thus renders analyses of what constitutes mental representation that reference intentional action viciously circular. What's more, if an analysis of what constitutes mental representation must make essential reference to intentional actions, then PCT will itself turn out to be *viciously* circular, since, in referring to mental representations, it will indirectly refer to intentional actions.

As we've seen, PCT moves from the observation that when an agent acts intentionally she does so on the basis of reasons, where these are commonly taken to be an agent's mental representations, to an analysis of intentional actions that makes essential appeal to mental representations that cause the actions that are intentional "in the right way." In so doing, PCT generates a constraint on theories of mental states:

> **The CT Constraint:** An account of mental states cannot make (ineliminable) reference to intentional actions, on pain of circularity.

If our argument is correct, PCT is problematic because the CT Constraint it generates is problematic.

Why think it's problematic? For one thing, there's good reason to think that an account of mental states must appeal to intentional actions. Behaviorist accounts make mental states entirely a matter of dispositions for behavioral actions. But not everything that might be considered a behavior (sweating, reflexes, etc.) is plausibly a mental state, so behaviorist theories will need to appeal to intentional action to distinguish between mental states and other dispositions, violating the CT Constraint.

That said, most philosophers these days are representationalists. So perhaps the CT Constraint just gives representational theories another advantage over dispositional theories. But that depends on whether philosophers can analyze mental representations in a way that doesn't make essential appeal to intentional actions, and there's reason to think this isn't possible. Consider one prominent account of mental representation—the teleosemantic account—which appeals to causal or information ("tracking") relations together with biological functions. For teleosemanticists, something functions to represent iff it functions to produce certain effects. But just as the behaviorist needed to restrict the class of behaviors, teleosemanticists must restrict the class of effects. Not just any old biologically fit effects will do if the analysis of representational function is to exclude, for instance, the cellular activity in one's gut. Thus, in distinguishing effects of representations from non-representational, causal activity, the teleosemanticist will also plausibly have to appeal to intentional actions. If so, then the CT constraint will again generate circularity. To wit, in defining mental representations, we believe that the teleosemanticist must distinguish non-representational information states from robustly representational psychological states and to do that they must appeal to the fact that the robust

variety causes intentional actions. PCT defines intentional actions as actions that are caused by psychological representations ("in the right way"). So the definition of psychological representation points us towards intentional action and the definition of intentional action points back towards psychological representations.

Now, it might be thought that what are sometimes called "phenomenal intentionalist" accounts of mental states, i.e. accounts that ground the representational powers of mental states first and foremost in phenomenal properties, are consistent with the CT constraint.[14] Such views treat phenomenal properties and a subject's awareness of them as a primitive and explain a subject's intentional relation to, i.e. her conscious awareness of, things in the world in terms of her awareness of phenomenal properties. Such views plausibly avoid the vicious circularity in relation to the CT constraint just described. Still, we think they do encounter a nearby tension when it comes to how such intentionalist views account for intentional relations and thus reference to mind-independent particulars.

Consider, for instance, David Chalmers's "two-stage" account of perceptual content. According to Chalmers, the "*phenomenal content* of a perceptual experience is a representational content that is determined by the experience's phenomenal character," i.e. its "presentational phenomenology." (3) And, according to Chalmers, "the presentational phenomenology of an experience immediately grounds an Edenic content" (55) i.e. content that refers to a merely possible "Edenic world." (27) Edenic contents are "Fregean" in that they refer to the actual world via a "matching relation." (28):

> The phenomenal character of an experience determines an Edenic phenomenal content, and it determines an unsaturated Fregean phenomenal content. According to the unsaturated phenomenal content, an experience is veridical iff the relevant object has properties that match the relevant Edenic properties. Once combined with a standard of matching, this unsaturated content determines a saturated Fregean content. (60–61)

But where does the relevant matching relation come from? According to Chalmers, "the answer is clear: it comes from the *inferential role* of visual experience." (54) Moreover, Chalmers suggests that there is "good reason to believe that quite generally, mental content is tied to inferential role." He claims that "the core inferential role of a perceptual experience is reflected in the pattern of judgments about veridicality and nonveridicality that the subject of such an experience makes, or more strictly, in the pattern of judgments that they should rationally make." (54).

Our worry, in a nutshell, is that inference is essentially a kind of intentional action and, arguably, one essentially grounded in social practice (Brandom, 1994, 2000).[15] Indeed, inferentialist theories of linguistic meaning have pragmatist roots and are a

---

[14] See also Mendelovici (2018), Pautz (2007).

[15] One might object that inferences can be carried out by computers, so are not intentional actions. However, these are only inferences by virtue of being treated as such by human users, or by virtue of the intentional activities of the humans that designed them. Computer inferences presuppose human intentional activities.

species of "use theories" of meaning. And if "edenic contents" only become actual-world-referring contents by virtue of intentional actions (i.e. the activities and practices in which inference has its home), then actual-world-referring contents cannot explain the intentionality of actions. Moreover, it isn't clear how purely edenic contents can "present" the actual world in a way that can explain the intentionality of actions insofar as actions are actual-world-directed. They might somehow explain "edenic actions," but until the content of mental representations is actual-world-directed, it's hard to see how they could account for the actual-world-directedness of intentional actions. And we worry that it is precisely the actual-world-directedness of intentional actions that allows appealing to them (or activities or practices that essentially involve them) to account for the world-directedness of intentional contents.[16] So we think there's good reason to worry that the CT constraint also generates tension in relation to phenomenal intentionalist accounts of mental representation.

In light of the inter-theoretic tensions the CT-Constraint gives rise to, we consider it a desideratum on a theory of intentional action that it avoids generating the CT-Constraint.

3 An alternative to the propositional-causal theory

A natural response to the puzzles that result from adopting PCT is to resist the temptation to define intentional actions in terms of their causal etiology in psychological representations. As an alternative to PCT, the AAT demarcates intentional action as having a particular kind of teleological form. On its own, AAT diverges from PCT with respect to how it answers REASONS and CONNECTION: AAT denies that intentional actions are constitutively caused by propositional attitudes. AAT can thereby sidestep objections involving deviant causal chains. And on its own, the AAT account avoids generating the CT-constraint.

We read Anscombe as endorsing a version of AAT in an important passage of *Intention*. From the observation that one can meaningfully ascribe intentions not only to human language-users, but also to non-linguistic animals, Anscombe suggests that possession of a propositional (descriptive) belief or desire is neither necessary nor sufficient for an agent's action to count as intentional. Rather,

> ...we describe what [non-linguistic animals] do in a manner perfectly characteristic of the use of intention concepts: we describe what *further* they are doing *in* doing something (the latter description being more immediate, nearer to the merely physical): the cat is stalking a bird *in* crouching and slinking along with its eye fixed on the bird and its whiskers twitching. The enlarged description of what the cat is doing is not all that characterizes it as an intention (for enlarged descriptions are possible of any event that has describable

---

[16] Nor do we think that the gap can be filled by appealing to merely causal relations for more or less the same reasons Chalmers doesn't want to ground the matching relation in merely causal relations.

effects), but to this is added the cat's perception of the bird, and what it does if it catches it. The two features, knowledge and enlarged description, are quite characteristic of description of intention in acting. Just as we naturally say 'The cat thinks there is a mouse coming,' so we also naturally ask: why is the cat crouching and slinking like that? and give the answer: it's staling that bird; see, its eye is fixed on it. We do this, though the cat can utter no thoughts, and cannot give expression to any knowledge of its own action, or to any intentions either. (*Intention* 1957, §47; 86)

While this passage is somewhat enigmatic, Anscombe appears to suggest that there are two features that make the cat's action count as intentional. The first feature is that the action can be explained in an enlarged description of that action. That is, we understand that the cat is crouching and slinking along *intentionally* when we understand that it's crouching and slinking along *in order to* stalk a bird. Indeed, insofar as the cat's crouching and slinking along are constitutive of its stalking the bird, "stalking a bird" is just an expanded description *of the same action*. But it is this expanded description that allows us to understand the crouching and slinking as intentional. However, as Anscombe notes, the enlarged description is not sufficient to secure intention since "enlarged descriptions are possible of any event that has describable effects." In addition to the enlarged description, what's needed is knowledge that is manifested in the action. In Anscombe's example, the cat's knowledge is perceptual and, insofar as cats cannot speak, non-propositional and non-descriptive.[17]

We think that the notion of practical knowledge Anscombe invokes—the non-descriptive (and non-observational) knowledge the cat has in perceiving the bird it's hunting and that Anscombe also described in terms of knowledge-how and exercises of general capacities in particular fields—is best understood as involving a form of psychological representation. For at root, we take it that psychological representations constitute an agent's *grasp* or *understanding* of her situation. A fact can be a reason for which an agent acts only insofar as an agent has a psychological grasp of or hold on that reason. How an agent grasps, understands, or perceives a fact is essentially related to the ways in which she's prepared to respond intentionally to that fact. Where the response is non-accidentally successful—in the way in which an action can be successful independently of whether it's actually performed—t's because the representation—grasp, understanding—was successful, i.e. knowledge.

The problem is this: the standard view of psychological representation analyzes them in terms of propositional attitudes. But propositional attitudes are causally related to intentional actions. So, if practical knowledge (along with enlarged description) accounts for the intentionality of actions, and if practical knowledge is a propositional attitude, we seem to have found our way back to PCT. In short, an essential part of what makes AAT a genuine alternative to PCT is that it rejects the notion that propositional attitudes are constitutive of practical knowledge.

---

[17] See also Anscombe's "Under a Description" (1979) pp. 221. In the later sections of *Intention* (§87–90), she clearly connects practical knowledge with knowledge-how and skill.

We venture that it is for this reason, together with the fact that there's no well worked-out alternative to the propositional attitude model, that Thompson (2008) focuses his attention on Anscombe's (Aristotle-inspired) "expanded description" criterion and develops a way of making sense of reasons for action that leaves psychological representations out of the story. Thompson develops a version of AAT that connects the teleological form of an intentional action with an Aristotelian notion of a life-form. So, in accordance with the AAT, Thompson does not define intentional action in terms of causation by an agent's reasons understood as her psychological grasp of facts. Indeed, an agent's psychological grasp of facts appears to play no essential role in Thompson's version of AAT. Thompson argues that an action is ultimately "rationalized" not by a description of the agent's psychological states but by locating the action in the "developing process" or context of her other actions (ultimately, the activities that are characteristic of her form of life). On this view, the question, "why are you cracking eggs?" is sufficiently rationalized with the answer "because I'm baking a cake" which makes no reference to the agent's intentions or desires but rather places the action of egg-breaking in a sense-making context. Thompson allows for a sense in which propositional attitudes—in particular, "intentions" and "desires"—can be invoked in some rationalizations: the question, "why did you go to the store?" can sensibly be answered by "because I wanted to buy eggs." But on Thompson's view, such ascriptions do not refer to psychological states in which agents are related to propositions or to the world by means of propositions. Rather, they are "sophisticated" descriptions of the agent's larger action or developing process (plus additional information).

We are sympathetic to Thompson's version of AAT, but we worry that his way of understanding REASONS and CONNECTION generates a different version of the CT-Constraint insofar as it appears to be compatible only with a behaviorist account of psychological states.[18] In so doing, it risks losing track of the second feature Anscombe highlights in the passage. In a word, we worry that it threatens to throw the baby of psychological representation out with the bathwater of propositional attitudes. But this would be unnecessary if we could find an account of psychological representations that allows us to maintain a version of AAT and thus to sidestep the deviant causal chain objection and avoid generating the CT-constraint. This is the positive project of the rest of the paper.

3.1 The problem-solving account of intentional action

We begin by explicating Springle's account of intentional action which, like Thompson's version of the AAT, defines intentional actions in terms of a kind of form that's distinctive of the activities of living organisms or life-forms. Springle cashes-out this form via her notion of a "solution" to a "practical problem." Specifically, Springle defines intentional action in terms of the notion of an "application-accountable local

---

[18] See Setiya (2009) footnote 36.

solution." In order to unpack this definition, we consider the notion of a "solution" more generally.

For Springle, substances that are life-forms are distinguished by their capacity to flourish. Different forms of life (or species of life form) have different types of needs or conditions for flourishing. Springle uses the notion of needs for flourishing to define one of the central notions in her account of intentional action: "practical problems" ("problems" for short). A type of practical problem arises out of a life-form's needs (basic or derived) for flourishing and is defined in terms of types of "ends"—particular forms of interaction between agents and "targets," in the world such as berries, oceans, and stop signs. For example, a life-form needs the metabolic resources to maintain and energize its body.[19] That is, she needs certain targets—some amount of water, minerals, carbohydrates, etc. But a life-form doesn't merely need these targets to exist out in the world somewhere, she needs to acquire and incorporate them into herself: she needs to *interact* with them. So "eating" is a basic type of problem or need, in that it is a type of end, involving a type of interaction (acquisition and incorporation), with a type of target (edibles). "Eating" is a quite general type of problem in that it is defined in terms of a general type of end itself defined in terms of a general type of target and interaction. One may define more specific subtypes in terms of more specific types of targets (more specific types of sources of specific sorts of nutrients) and interactions (the types of processes of acquisition required for different types of sources, e.g. grazing, hunting, or shopping). Nutrition and protection from threats are basic types of needs for biological flourishing, but Springle does not attempt to reduce all other types of needs or problems to some subset of basic needs. Accordingly, friendship or intellectual needs might be basic for some forms of life, e.g. for humans.

Problems can be local or global. A local problem is an instance of a type of problem, for example, a dehydrated animal faces the local problem of needing to drink some water. Likewise, an animal being chased by a predator faces the local problem of needing to escape from that predator. Indeed, these needs can coincide in the same animal, who faces the local problem of needing to drink some water as well as the local problem of needing to escape a predator. But suppose that if the animal were to stop for a drink of water before escaping the predator, the predator would most likely kill the animal. The animal now faces a global problem. Global problems are defined in terms of the end of organizing (ordering, prioritizing) solutions to local problems in a way that promotes an animal's overall flourishing. Since, in the example just given, we are supposing that the animal's overall flourishing depends more on evading the predator than getting a drink at this moment, the animal faces the global problem of prioritizing evading the predator over stopping for a drink.

Solutions are defined in terms of the problems they solve. A solution is a type of interaction between a life-form or subject and a type of target[20] that is generated by a life-form that *non-accidentally* satisfies a type of need for flourishing. Examples of

---

[19] According to Springle, practical problems apply to the whole organism/animal; the "biological problems" an animal's organs solve are derivative.

[20] Including the vehicles of surrogative representations—see below.

solutions include anything from drinking and evading predators to making friends to producing objects such as paintings and shoes. A type of solution, e.g. eating, may be realized (tokened) in a number of different manners and may involve different sequences of sub-solutions some of which correspond to what are sometimes called "basic actions" like walking, raising an arm, etc.

A type of "local solution" solves a type of local problem. A type of local problem has the potential to define a type of global problem in relation to other types of local problems. For instance, a type of local solution that involves staying in one place is not compatible with a type of local solution that involves locomoting to a different place. They therefore mutually define a type of global problem. Global problems need not necessarily be defined in terms of spatio-temporally incompatible local solutions. For instance, swimming immediately after one eats might not typically promote one's overall flourishing. Nor are global problems necessarily defined in terms of incompatible local solutions. In some cases, different local solutions (in particular with respect to their sub-solutions) may be complementary. For example, the trash can in my bedroom is full. If I need to run up to my room to grab something before I go out, it makes sense to take the trash, so I can drop it off in the outdoor can as I leave the building. Global problems are solved by selecting the right sequences of local solutions at the right times. In short, a local solution answers the question "how do I solve this local problem?" while global solutions answer the question "which of the local problems I'm facing should I try to solve now?" by way of answering the question "which of these local solutions should I prioritize, i.e. perform now?"

Springle proposes that types of intentional action are grounded in types of local solutions which, in relation to other types of local solutions, are types of global solutions. A token intentional action is thus a token of a type of local solution. Some local solutions instanced (potentiated) in a context are global solutions (whether a local solution instance is a global solution may change as the context changes), so some intentional actions (namely those we'd deem rational) are global solutions. But all local solutions are potentially global solutions by virtue of standing in particular relations vis-à-vis other types of local solutions with respect to a subject's global flourishing.

According to Springle, agents have abilities for types of intentional actions (local solutions). When an agent intentionally acts, she applies a token of a type of local solution to a situation. When a type of local solution is applied, a token solution of that type is potentiated and a subject is thereby in a position to execute that solution. That token is "appropriate" iff, were it performed to completion, it would non-accidentally actualize an instance of the type of solution it is. This actualization occurs when the token is applied to, i.e. potentiated in response to or for, a token of the problem type, which is of the type the token solution solves. For example, eating is a local solution that solves the need for food iff it is directed at food. So it is appropriate when one needs food and only when what one eats is food. In other words, it only when one needs food that eating solves this problem, and only when the eating is directed at food that it solves this problem. So a token intentional action is always a local solution, but it may be an *inappropriate* local solution. It is an inappropriate local solution if it fails to "practically fit" the context in which it is applied

(for which it is potentiated). A global solution may likewise be inappropriate if it is applied to (selects) a local solution that would not, were it performed to completion in the relevant context, non-accidentally promote the subject's overall flourishing. An intentional action is only an appropriate global solution if it is an appropriate local solution, but it may be an appropriate local solution even if it is not an appropriate global solution. For instance, eating a bag of chips may be an appropriate solution in that it solves one's need for certain nutrients, but not globally appropriate because it isn't compatible with solving the problem of lowering one's cholesterol.

To *apply* a solution is not yet to *perform* it: a type of local solution is applied when a token of that type is potentiated. To wit, a type of solution is applied when an agent is *in a position to perform* one of its instances, where an agent is in a position to perform one of its instances iff an instance has been determined, i.e. potentiated. Accordingly, Springle stresses the importance of distinguishing between (in) appropriateness of a solution as applied on the one hand and the success or failure of its performance or execution on the other. Performance failures can result from circumstances that require one to abandon the action, for example, when it is no longer globally the appropriate thing to do, though it still may be locally appropriate. They may also result when a token of an appropriate type of action involves executional errors, in the sub-actions that compose a token action's manner of execution. Walking to the kitchen is an instance of a solution to your needing to make coffee, but you might trip and injure yourself on your way to the kitchen, and consequently never finish your act of making coffee. In that case, the solution *walking to the kitchen* was never fully actualized. But it was still a potentiated solution since had you not fallen, you would have made coffee. So, the appropriateness of an intentional action is a matter of its potential, is robust to performance errors, and does not depend on being actualized.

We come, finally, to the concept of an "application-accountable" solution. As we just saw, whether a type of solution (intentional action) is appropriate depends on the conditions in which it is applied. The conditions in which a type of solution (local and global) is applied depend on its "application algorithms." An application-accountable solution is indexed to individual life-forms rather than a species of life-form and is "accountable" to the appropriateness or inappropriateness of its applications iff the execution of inappropriate solutions leads to modifications of its *own* application algorithms. In other words, an intentional action is a type of solution that is not indifferent to the appropriateness of its instances. For instance, call the type of interaction that solves the problem of needing nutrients "Nutrient Consumption," of which "Eating" is a subspecies. Suppose that "Eating", but not "Nutrient Consumption" is an application-accountable solution. If an individual agent has the ability to eat, as opposed to merely having the ability to "consume nutrients," then the agent has the ability to apply a type of solution whose application algorithms are at least capable of being modified any time the agent eats inappropriately, i.e. when the "Eating" solution is applied inappropriately *and* executed. In such cases, the solution of which that particular agent is in possession is itself disposed to modifications in respect of its application algorithms as they are realized in that particular agent. In contrast, a creature that merely has the ability to consume nutrients does not have an ability to apply a type of solution whose application algorithms are capable of being

modified any time the agent consumes nutrients inappropriately. Acts of consuming nutrients are *almost* intentional actions, but not quite, because they are instances of a type of solution that is not application-accountable.[21]

With Springle's problem-solving analysis of intentional action in hand, we turn in the next section to an explication of her analysis of representation. Since Springle does not depend on psychological representations in her definition of intentional action, she may define psychological representations in terms of intentional actions without circularity. This way of analyzing psychological representations in terms of the problem-solving account of intentional action has repercussions for how we think about psychological representation: rather than building intentions out of sui generis psychological states, we can understand beliefs and desires as modes of intentions. Moreover, psychological representations are understood not as propositional attitudes with constitutive satisfaction conditions but rather as "embodied instructive representations" with constitutive appropriateness conditions.

4 Embodied instructive representation

The main concept of Springle's analysis of representation is the notion of "Practical-Epistemic-Access" (PEA):

Def. PEA: An agent has PEA to X iff an agent is in a position to produce an intentional action in response to X and that is supposed to practically-fit X.
Def. Representation: R is a representation of (about, directed at) a situation X iff R constitutes a form of PEA to X.

According to Springle, a representation is constitutively something that provides PEA to situations (facts, needs and targets). PEA is defined in terms of an agent's being in a position to produce an intentional action in response to a situation. The relevant intentional action is supposed to "practically fit" the situation in that it is supposed to be an appropriate response to it. So an agent represents a situation insofar as an agent is in a position to produce an intentional action as a response to it—i.e. that is itself directed at the situation. As we saw in the previous section, Springle defines intentional actions as application-accountable solutions. So, PEA is grounded in the exercise of an ability for a type of application-accountable solution where an exercise constitutes the application of that solution to a situation. Again, to apply a solution is not yet to perform it: a type of local solution is applied when a token of that type is potentiated. So on this view, PEA, and thus representation, consists in the potentiation of a solution (intentional action) that's directed at a situation. In other words, to be representationally related to a situation X just is to be in

---

[21] One might worry that sweating would count as an intentional action. But sweating is not a candidate global solution; it does not compete with other, potentially incompatible, local solutions for performance. Indeed, there seems to be no space between the activation of dispositions to sweat and the actualization of sweating. In contrast, intentional actions begin as potentiations (the potentiality of a determinate token of a type of solution) and whether they are performed depends on whether they are selected as global solutions. We thank an anonymous reviewer for alerting us to the need to address this sort of objection.

possession of a potentiated solution that's directed at X. In this way, the PEA analysis inverts the standard order of explanation: rather than thinking that the intentional directedness of action depends on the intentional directedness of representations, the intentional directness of representations just is the intentional directness of potentiated intentional actions, i.e. application-accountable local solutions.

Springle distinguishes "direct" from "indirect" species of PEA:

Direct PEA: R provides an agent with direct PEA to X iff in having R, an agent is positioned to produce an intentional action that is a direct response to X, where a response is direct when its immediate target is X.

Indirect PEA: R provides an agent with indirect PEA to X iff in having R, an agent is positioned to produce an intentional action that is an indirect response to X, where a response is indirect when its immediate target is something that takes the place of X—something one acts on directly so as to act on X indirectly.[22]

Direct PEA grounds "embodied instructive representations" while indirect PEA grounds "surrogative representations":

Embodied Instructive Representations generatively (see below) stand-for ("instruct" in a special first-personal sense) intentional actions by virtue of being potentiations of them, i.e. applications of application-accountable local solutions. They have constitutive appropriateness conditions that are identical to the local appropriateness conditions of the intentional actions they instruct. They have the schematic form "ϕ-THAT-THUS!", where ϕ is a type of solution (e.g. eating), THAT is target (e.g. a particular apple), THUS is a manner of realizing ϕ (which may involves various sub-acts, e.g. climbing an apple tree), and "!" indicates the "force" or urgency of the instruction.

Surrogative Representations surrogatively stand-in-for targets by virtue of being (or functioning to be) solutions to proximity problems involving those targets.[23] They have constitutive alethic appropriateness conditions: they are appropriate to the extent that they are "true" or "veridical" to their targets, i.e. to the extent that they have features they need to have to be solutions to the

---

[22] Indirect PEA presupposes direct PEA: one must have direct PEA to the R that provides indirect PEA.

[23] Under some conditions, targets are "unavailable" for certain kinds of direct interactions. A direct interaction may, for instance, be too dangerous or otherwise costly or difficult, or it may not be possible because the relevant targets are facts that obtained in the past but no longer, or facts that may come to obtain in the future but not yet, or facts that are mere logical possibilities, or facts that consist not of particulars but of generalities, patterns, universals, etc. if such things can exist. When an agent has a practical problem that requires a type of interaction with an unavailable target for its solution, she faces a ***proximity problem***. When an agent faces a proximity problem, she needs a go proxy (stand-in) for the unavailable target, with which she can interact and thereby satisfy her overarching need. So, something that solves a proximity problem functions as a phase in a larger solution to a practical problem and it functions to contribute to the larger solution by enabling a type of indirect-interaction with a target when direct-interaction is not a (good) option. Indirect PEA is just such a solution (and it may be more or less appropriate). When appropriate, indirect PEA extends agents' practical-problem-solving abilities by enabling agents to intentionally respond to targets that would otherwise either not be accessible for the relevant form of practical engagement at all or not in a way that's convenient or safe.

proximity problems those targets partially define. At least paradigmatically, they have the schematic form "S is F" where S is a target they stand-in for and F is the way they describe, picture, or model S to be (the way they stand-in for S).

Springle argues that psychological representations—intentions, beliefs, and desires—are embodied instructive representations. Embodied instructive representations are realized by abilities for local solutions, i.e. "problem-solving dispositions." Local solutions are realized by exercises of "instructive dispositions" while global solutions are realized by exercises of "organizing dispositions." Before we turn to explaining Springle's account of problem-solving dispositions and embodied instructive representations, a few remarks are in order.

First, Springle proposes that instructive dispositions are an agent's "embodied practical concepts," a.k.a. "instructive concepts."[24] Second, according to Springle, there's an important distinction between representing an action and "acting on it" or, in the case of an embodied instructive representation, "acting it out." This is just the difference discussed above under the rubric of the "potentiation" as opposed to the "performance" of an intentional action. Third, paradigmatic surrogative representations are non-psychological or public representations and acquire their indirect PEA from the intentional actions of agents who produce and utilize them; they refer to their targets by virtue of being used by agents in order to solve proximity problems. However, Springle acknowledges that it is possible that there are also surrogative psychological representations and that representations that are grounded in indirect PEA could be non-artifactual or "naturally occurring."

For our purposes here, Springle's central thesis is that the reasons that motivate actions are psychological representations, but psychological representations are embodied instructive representations that are realized by instructive dispositions. One of the ways Springle motivates her account is by asking what picture of mental representational content we would have if we took intentions rather than beliefs as our theoretical starting place. Accordingly, rather than following PCT in supposing that beliefs and desires are more basic because they correspond to the premises in the practical syllogism that lead up to the conclusion that corresponds to the intention, Springle's account of mental representation takes intention as basic.

We believe that embodied instructive representations offer an alternative to the propositional attitude model of mental representation. In order to show this, in Sect. 4.1 we sketch Springle's account of "problem-solving dispositions."[25] As noted, these are of two types: "instructive dispositions," which realize solutions to local problems, and "organizing dispositions," which, in connection with instructive dispositions, realize solutions to global problems. In Sect. 4.2 we provide a detailed

---

[24] We recognize that some philosophers mean something importantly different by "concept" and we ask that the reader resist importing anything into our notion beyond what we've specified. A standard sort of example of a concept as we understand it is "Red". If an agent possesses the concept "Red," it is true of that agent that, if she were to see a red item and someone asked her what color it is, she would (ceteris paribus) answer "red.

[25] These – in particular their input sub-components – are implementations of application-algorithms.

description of the embodied instructive representations problem-solving dispositions realize. There are different modes of activating instructive dispositions and that constitute I-instructions, N-instructions, and T-Instructions, which contribute in different ways to determining a potential intentional action and thus its appropriateness conditions and whether they're satisfied. Finally, in Sect. 4.3, we explicate the "Generative" species of aboutness Springle's identifies with embodied instructive representations.

4.1 Problem-solving dispositions

Instructive dispositions are a kind of active causal capacity: they function to enable agents to produce or generate intentional actions. According to Springle, they essentially involve an active causal component, the "φ-output sub-capacity." Since intentional actions constitutively function to solve practical problems, the φ-output sub-capacity of any given instructive disposition must be modulated by two passive causal sub-capacities: a "need-input sub-capacity" that responds to instances of a type of need (local problem), and a "target-input sub-capacity" that responds to instances of the types of targets that satisfy those needs (Fig. 1).

A type of need (problem) is for a type of "terminal target" and for a type of interaction with that terminal target. For instance, if the need is nutrition, the terminal target is food (something edible), and the type of interaction is eating. So an agent's "Eat" instructive disposition functions to respond to the agent's hunger and to prepare and thus enable the agent to eat something edible.

The N-input sub-capacity functions to respond to agent's hunger, i.e. to activate the "Eat" instructive disposition when and only when the agent needs food. It may fail to fulfill this function by activating the "Eat" instructive disposition when the agent doesn't need food. As we'll see below, when this happens, the agent has an inappropriate N-instruction.

In order to fulfill the function of preparing and thus enabling the agent to eat something edible, the T-input sub-capacity functions to respond to terminal targets. But sometimes, an agent will be hungry, and will need to do something before there's a terminal target to which it can respond. In these contexts, the interaction with a terminal target that's required to solve a problem is "extended": it involves sequences of sub-actions and sub-targets by means of which an agent "secures" a terminal target. In order to secure something edible, an agent may need to check in certain places where food might be: a lion will scan the savanna for the movement of potential prey, a person will look in the fridge. But an agent may need to travel (a lion chases its prey, a person goes shopping) or engage in other sorts of activities (crying, asking for food); and sometimes must do all of these things.

Agents must also be sensitive to the absence of terminal targets. If an agent is hungry and food isn't immediately available, she will need to deploy a sequence of sub-actions that interact with non-terminal targets in order to secure some food. These are sub-solutions; they solve the problem of eating when food must be secured by solving the problem of securing food. Such sub-solutions are embedded in token acts of eating, since they are parts of the intentional action of eating.

They are realized by "sub-instructive dispositions." Sub-instructive dispositions are activated by instructive dispositions; they are abilities for solutions that indirectly serve an agent's flourishing by contributing to solutions that directly serve an agent's flourishing (Fig. 2).[26]

Sub-instructive dispositions are N-activated by (super-) instructive dispositions that are themselves N-activated. Sub-instructive dispositions may be T-activated by targets that are present in the context. Their φ-outputs are tokens of types of sub-actions, e.g. searching for food by going to particular places. There may be sub-sub-instructive dispositions that are N-activated by N-activated sub-instructive dispositions, and so on. Sub-sub instructions are T-activated by fine-grained features of contexts and their φ-output is just the modulating effect they exert on the φ-outputs of the (super) instructive dispositions that (ultimately) N-activate them.

Normally, (super-) instructive dispositions will deploy sub-instructive dispositions even when a terminal target is present. Since securing a terminal target typically involves some sub-actions that need to be responsive to the particular properties of a terminal target relevant to the interaction as well as to the properties particular to the context in which it is secured. Moreover, a single particular can be an N-input and a T-input at the same time. Consider, for instance, a rabbit's protective "Avoid" instructive disposition. If a rabbit sees a predator, e.g. a fox, it will at once N-activate and T-activate her "Avoid" instructive disposition. Avoiding the fox e.g. by escaping into a hole into which the fox can't fit, depends on the rabbit registering both a need (escape) and a target (the rabbit hole).

Instructive dispositions tell us what to do. Organizing dispositions function to modulate instructive dispositions, determining which of the things we're told to do are to be performed. In this way, they enable our solutions to local problems to be solutions to global problems, as well. Organizing dispositions function to enable short-term overall flourishing by taking as input the strength of two or more activated instructive dispositions and modifying instructive dispositions as well as their connections to one another through learning. In short, organizing dispositions enable an agent's instructive dispositions to promote an agent's overall flourishing (Figs. 3 and 4).

4.2 Instructive dispositions as embodied instructive representations

Springle proposes that the instructive dispositions an agent possesses are her embodied practical concepts, and that their activations constitute subjective instructive representings. For example, Raquel the Racoon's "Eat" instructive disposition constitutes her practical "Eat" or "Eat-[THAT]-[THUS]" concept (Fig. 5).

---

[26] Note that this does not mean that they have indirect PEA-Functions. Sub-instructive dispositions have direct PEA-Functions in that they function to enable an agent to produce a solution (a type of sub-action). They are indirect in the sense that their PEA-Functions depend on (super-) instructive dispositions, for indirect instructive dispositions only function to enable an agent to solve derivative problems, i.e. problems that arise in the context of solving more basic problems.
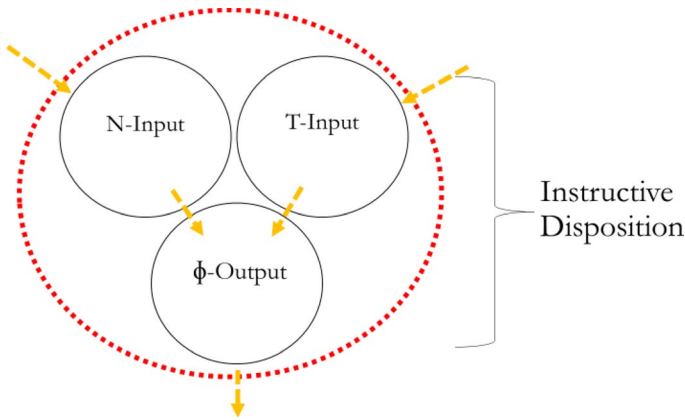
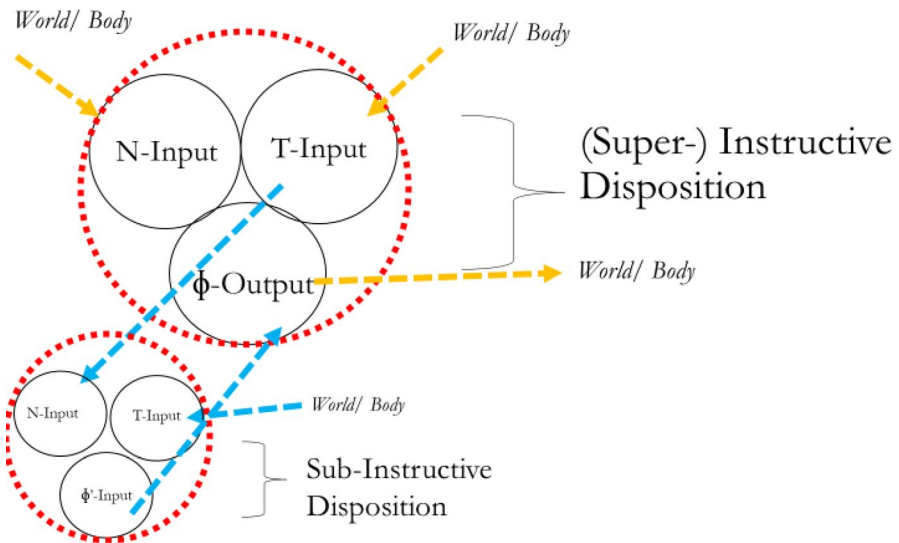**Fig. 1** Instructive disposition

Fig. 2 Sub-instructive dispositions

This practical concept stands for a determinable type of action, namely, eating. The concept is applied when the instructive disposition that realizes it is activated. When the concept is applied, an act of eating is instructed. In other words, when the instructive disposition is inactive, it stands for eating where that action is entirely determinable. The conditions of application of the concept determine a potential act of eating in that (1) the N-input that N-activates the disposition determines a particular instance of the type of practical problem (nutrition), (2) the T-input and its context determine a particular terminal target ("THAT") or manner of securing a terminal target ("THUS") which involves a sequence of sub-acts and their targets. The conditions of activation and thus application of
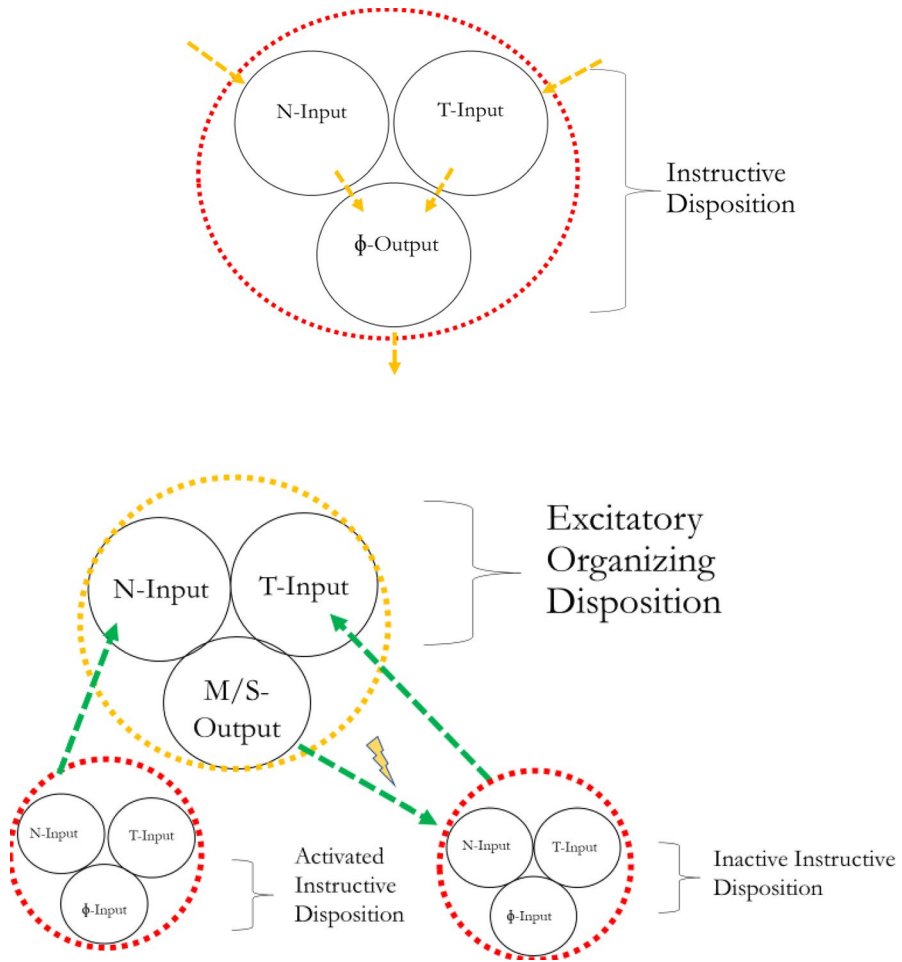
Fig. 3 Organizing disposition (excitatory)

the concept thus determine the appropriateness conditions for the instructed act of eating.

There are four modes of instructively representing which correspond to four modes of activating instructive dispositions and four degrees of determinateness of an instructed action: T-instructions, N-instructions, i-instructions, and I-instructions.

T-instructions correspond to the T-input mode of activation. In the image below, Raquel's seeing an apple consists in it serving as an activating (visual) T-input to her "Eat" instructive disposition. This constitutes an application of her practical "Eat-[_ THAT_]-[_THUS_]!" concept that instructs Raquel to "Eat-[APPLE]-[THUS]." In other words, it instructs an act of eating that is determinate to the extent that it has a particular terminal target together with a particular realizing manner ("THUS")— probably just a matter of a sequence of sub-sub-acts—that's determined by the
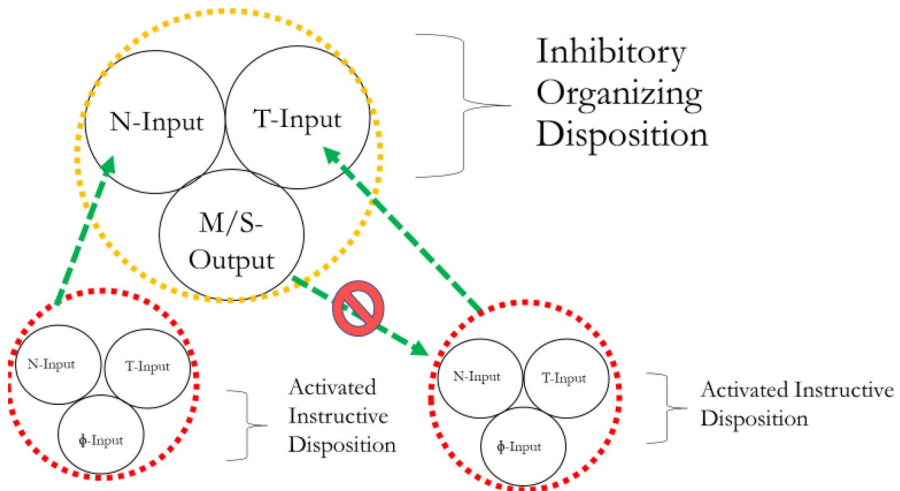
Fig. 4 Organizing disposition (inhibitory)

terminal target and its context. However, Raquel's "Eat" instructive disposition is *only* activated by the T-input; she only has a T-instruction. In the absence of a corresponding N-instruction, Raquel won't eat (since she isn't hungry!). While Raquel's visual experience T-instructs her "Eat-[APPLE]-[THUS]," she isn't motivated by that instruction.

The activation of the N-input sub-capacity of the Eat instructive disposition constitutes a representation (a T-instruction) only by virtue of the sub-capacity being a constituent of the instructive disposition such that activating the sub-capacity constitutes an activation (if only partial) of the instructive disposition. So no sub-capacity represents independently of its role in one or more instructive dispositions, and when a sub-capacity is activated it constitutes an instructive representation.[27] Some instructive dispositions have a jointly sufficient set of T-inputs. For instance, a visual apple T-input might not be sufficient to activate the T-input component of her Eat disposition. For there might be contexts in which apples are unlikely to be edible to her, e.g. when they are rotten, so that apples won't activate the N-input of Raquel's "Eat" instructive disposition unless they make some particular impact on both Raquel's visual system and her olfactory system in a certain way.

---

[27] This is compatible with one T-input sub-capacity figuring in multiple instructive dispositions. Say Raquel is an inventive and industrious racoon who suffers from back pain. Seeing the apple may thus activate the N-input component of both her "Eat" instructive disposition and a sub-instructive disposition of her "Sooth" instructive disposition—an instructive disposition that functions to instance solutions to injury-related problems. Her already N-activated "Sooth-with-[THAT]-[THUS]" may thus instruct her to "Sooth-with-[Apple]-[THUS]" where "THUS" involves rolling her back against the apple. (Alternatively, the instruction might be "Sooth-[Back]-[THUS]" and sub-instructive dispositions that specifies the "THUS" in terms of "Roll against-[apple]-[THUS].").
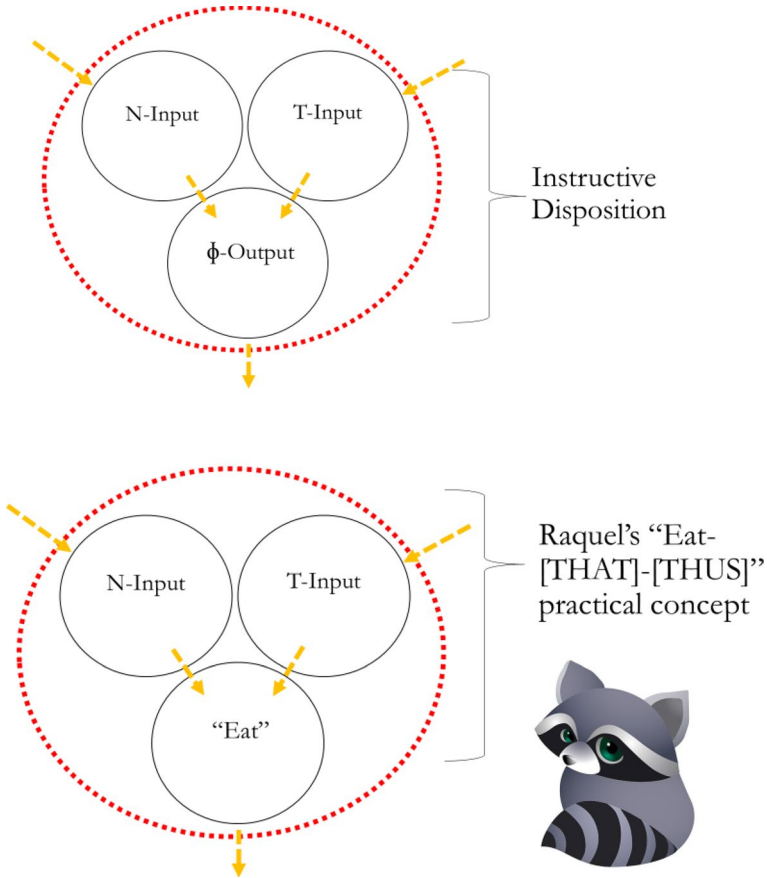
Fig. 5 Instructive dispositions as embodied practical concepts

This scenario illustrates how T-instructions are comparable to means-end beliefs. Raquel has a T-instruction but no corresponding N-instruction. So, T-instructions are like means-end beliefs in that they can occur independently of N-(and thus i-) instructions. By considering a case in which a T-instruction occurs in the absence of a corresponding N-instruction, we can also isolate the contribution of the T-input to preparing an agent to produce, and thereby instructing, a particular instance of a type of action. In this case, Raquel's T-instruction prepares her for an instance of eating that is individuated by the particular apple that is the potential action's target, which in its context determines the manner of the potential action.

T-instructions also contribute to the appropriateness conditions of instructed actions in a way that corresponds to their distinctive contribution to the determination of the potentiated token action. For an instructed action with an inappropriate terminal target will not succeed. For instance, if the particular apple Raquel is looking at is actually a wax apple, then the instruction to eat that apple is inappropriate. After all, if she were to eat it, it would not in fact nourish her. The act of eating that

particular apple would not constitute a solution, and thus it would not be appropriate. Therefore, the T-instruction to do that act is inappropriate.

The way T-instructions contribute to the appropriateness conditions of instructed actions provides another point of similarity between T-instructions and means-end beliefs. Just as the success (truth) of means-end beliefs depends on whether the relevant X is really a means to the relevant Y, so the appropriateness conditions of T-instructions depend on whether the targets they incorporate into the means of the actions they instruct render the actions they instruct means (solutions) to their ends (the problem the action functions to solve). Raquel's T-instruction "Eat-[apple]-[THUS]" is appropriate iff eating the apple (thus) would, if executed, solve a nutritional problem (if she had one). It's possible for an instructive representation to involve a chronically defective T-input sub-capacity. If a T-input sub-capacity is hyper-sensitive, it will frequently instruct actions with inappropriate targets. If it's hypo-sensitive, it will frequently fail to be activated, and thus frequently fail to supply N-instructed actions with targets, or it will frequently fail to produce N-instructions. Chronic defects are either mechanical in nature or they may stem from the influence of defective organizing dispositions.

T-instructions and means-end beliefs seem to support the same sorts of counterfactual conditionals. If Raquel believes that the apple before her is edible—would be a means to her end of satisfying her hunger—and she's hungry, she'll eat it. If Raquel's belief is false and she eats the apple, she'll need to eat something else, since the apple won't satisfy her hunger. Likewise, if Raquel's experience of the apple instructs her to eat it, then, if she's hungry, she'll eat it. She's already prepared to eat it should she become hungry. If the instruction is inappropriate, Raquel will need to eat something later, since the apple won't solve her nutritional problem (Fig. 6).

N-instructions are constituted by N-input activations of instructive dispositions. As with T-inputs, as N-input activations constitute N-instructions by virtue of constituting partial activations of the full instructive disposition. In the image below, Raquel's "Eat" instructive disposition is activated via its N-input component by her hunger. But unlike T-instructions, N-instructions motivate: the N-instruction Raquel receives from her stomach "moves" her to eat. N-instructions are thus comparable to desires.

As noted in the discussion of instructive dispositions above, instructed actions are often extended, meaning that they require a sequence of sub-acts to secure a terminal target. The context of the time of the N-instruction determines the initial manner of the instructed action. This is the situation pictured in the image below: Raquel receives the "Eat" N-instruction from her stomach when there's no food in sight. This instructed act of eating is individuated not only by the instance of the problem it functions to solve, but also by the context in which that problem is instanced, as contributes to determining the manner of the instructed action. The manner is a matter of the sub-actions [THUS] involved in securing a target for eating where those sub-acts will be guided by the targets available to those sub-acts in her initial context. For instance, perceptible features of Raquel's current context may instruct her to locomote to a nearby location where she's likely to find food. The context in which the N-instruction is instanced will not, however, fully determine

the manner of the action, which depends on how the action is actualized. In short, then, N-instructions token-individuate the actions they instruct by determining the token problem the potential action constitutively functions to solve, and by determining the initial conditions for the action.

The primary way N-instructions contribute to the appropriateness conditions of the actions they instruct is by determining the problem instance. It is possible for an instructive disposition's N-input sub-capacity to be activated by a "pseudo-problem"- something that activates the N-input sub-capacity but isn't an instance of the type of problem the instructive disposition functions to enable the agent to solve. When this happens, the N-instruction is inappropriate, since if there's no problem to be solved, then the intentional action can't be an instance of a solution. It's also possible for a chronic defect in an N-input sub-capacity to cause an agent to respond to instances of one type of problem as if they were an instance of a different type of problem. Thus, although N-instructions are desire-like, they are also belief-like in that their appropriateness turns on whether they are sensitive to genuine problems of the right type.

Finally, N-instructions support the same sorts of counterfactual conditionals desires support. Generally, if an agent desires X, she's likely to take the means to achieving X. Whether she does so will depend on what else she desires at the time, what means are available, and the strength of her desire. Likewise, if an agent is N-instructed to ϕ, she's likely to ϕ. Whether she does so will depend on whether the N-instruction is an i-instruction, and on the relative strength of the corresponding I-instruction, which is a function of the strength (itself a function of the intensity and weight) of the N-instructions together with the strength of the corresponding T-instructions, as modulated by the agent's organizing dispositions (Fig. 7).

I-instructions are just i-instructions that have been selected by organizing dispositions. And i-instructions are just instructive dispositions that are N-activated and T-activated (by a terminal target or by its absence when its absence N-activates sub-instructions). So, typically, an N-instruction will also be an i-instruction. An i-instruction is locally appropriate iff its corresponding N- and T-instructions are appropriate. Since an I-instruction reflects the influence of organizing dispositions, it is evaluable not only in terms of the local appropriateness but also global appropriateness conditions (Figs. 8, 9 and 10).

Thanks to the similarities between these different modes of embodied instructive representations and propositional attitudes, the former can replace the latter. However, embodied instructive representations do not altogether eliminate propositional attitudes. Rather, they absorb them: when embodied instructive representations involve surrogative representations as targets, i.e. when they are responses to surrogative representations as such, or when the actions they instruct essentially involve the production of surrogative representations (e.g. the production of a descriptive judgement), embodied instructive representations function as (intrinsically contentful) attitudes towards propositions. After all, embodied instructive concepts and representations represent (local application-accountable) solutions and, as noted in Sect. 3.1 above, problems can be basic or derived and they can consist in ends that are rather sophisticated, such as responding to or producing a propositional judgment, i.e. a surrogative representation. The appropriateness of an embodied
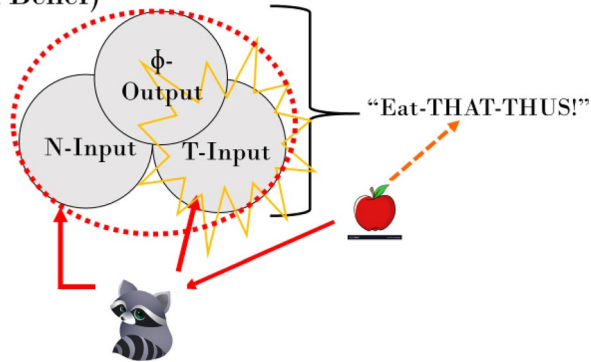
Fig. 6 T-instruction

instructive representation in which the production of a proposition (surrogative representation) figures as a constituent of the action depends on whether the surrogative representation is appropriate. And Springle argues that truth (accuracy, veridicality, satisfaction) may be understood as what it is for a surrogative representation to be appropriate; it may be the surrogative species of appropriateness, or at least an essential aspect of this species of appropriateness. Therefore, embodied instructive representations that consist in or essentially involve the production of surrogative representations like judgments are appropriate iff the surrogative representations one is instructed to produce are true. Likewise, the appropriateness conditions of embodied instructive representations that instruct responses to surrogative representations as such depend on the contents, and thus the truth-conditions, of those surrogative representations.

4.3 Generative directedness

On Springle's view, an intentional action is successful—i.e. appropriate—just when it's appropriately guided by embodied instructions.[28] Since Springle construes subjective reasons not as propositional attitudes but as appropriate N-, T-, and I-instructions, we believe her view provides the anti-intellectualist with the resources to reject P2 of IA. However, the intellectualist may object that embodied instructive representations are either propositional representations in disguise—e.g. imperatives—or else they are not really representations at all. We've already discussed a few reasons for distinguishing the embodied instructive species of representation

---

[28] Springle argues that surrogative representations may be understood as "inter-subjective reasons" while embodied intrusive representations may be understood as "intra-subjective reasons," where intra-subjective reasons may be responses to inter-subjective reasons. Intra-subjective reasons play a distinguishing role in Human intentional action.
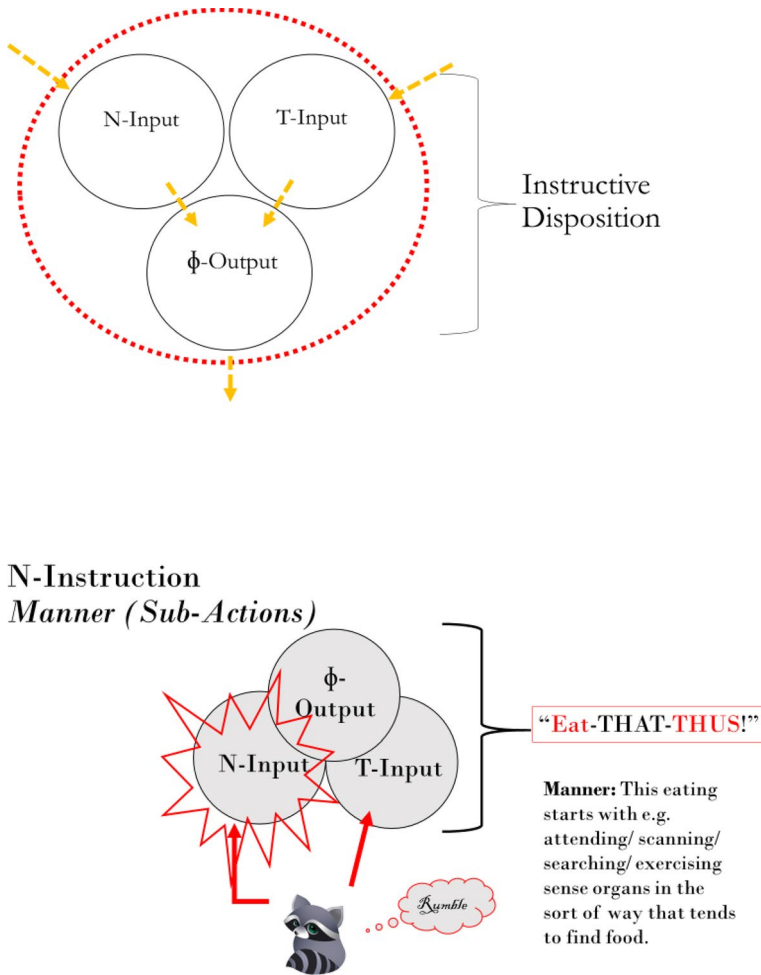
Fig. 7 N-instruction

from the surrogative, propositional species. They're grounded in different species of PEA-functions (direct and indirect, respectively). The former have appropriateness conditions but not truth (veridicality, accuracy, or satisfaction) conditions, while the latter have truth conditions understood as a species of appropriateness conditions. And the fact that both propositional (surrogative) and embodied instructive representations are grounded in PEA functions is a reason to count both as genuine species of representation. But embodied instructive representations and propositional representations can also be distinguished by their respective modes of aboutness, directedness, presentation or reference, where being a mode of aboutness captures what it is to be a representation.

According to Springle, surrogative aboutness is the philosophically familiar mode of aboutness; it's the aboutness of a representation that descriptively
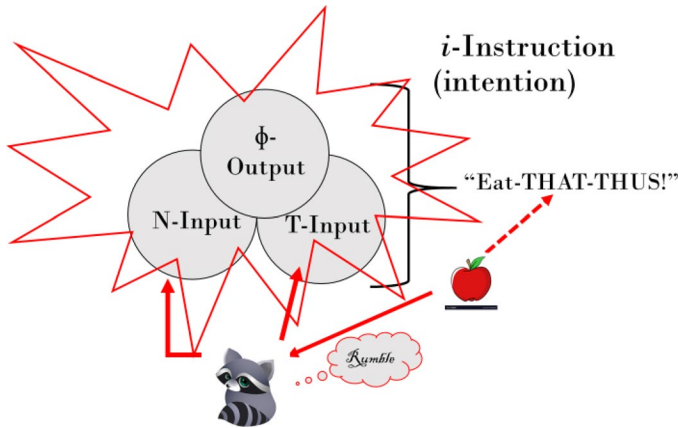
Fig. 8 i-instruction

stands-in-for and attributes some features to a subject. Springle gives the kind of aboutness embodied instructive representations instantiate the name "generative directedness". An embodied instructive representation does not stand-in-for what it represents. Rather, it stands-for and thus represents an intentional action by virtue of constituting the potentiality for it. An activation of an instructive disposition potentiates a token of the type of action (solution) the disposition stands-for qua embodied practical concept. The token intentional action is itself directed at the particular need and target(s) that activated the disposition; the particular need and target(s) are the individuating components of that token.

As noted, embodied instructive representational contents have the logical form φ-THAT-THUS! Following Springle, we'll use this form to illuminate the nature of generative aboutness and its connection with the nature of the instructive representations that realize the embodied instructive representations that instantiate this mode of aboutness.

When an agent has an embodied instructive representation, she is generatively directed at φ and THUS and generatively directed at THAT.[29] The agent is generatively directed at φ and THUS because they token a type of action that the agent is prepared to produce by virtue of the activation of the relevant instructive disposition. The agent is generatively directed at THAT because the target that individuates the token of φ contributes to determining the individuating THUS, and the agent is generatively directed at φ and THUS. In other words, when an agent has an embodied instructive representation, she is generatively directed at an intentional action and, ipso facto generatively directed to the target at which that action is directed (likewise for the need). In other words, according to Springle,

---

[29] Springle suggests that an agent is generatively directed at her action and through that action she may be "generatively-demonstratively" directed at the targets and needs at which her action is causally directed. For simplicity we omit the generative/ generative-demonstrative distinction.
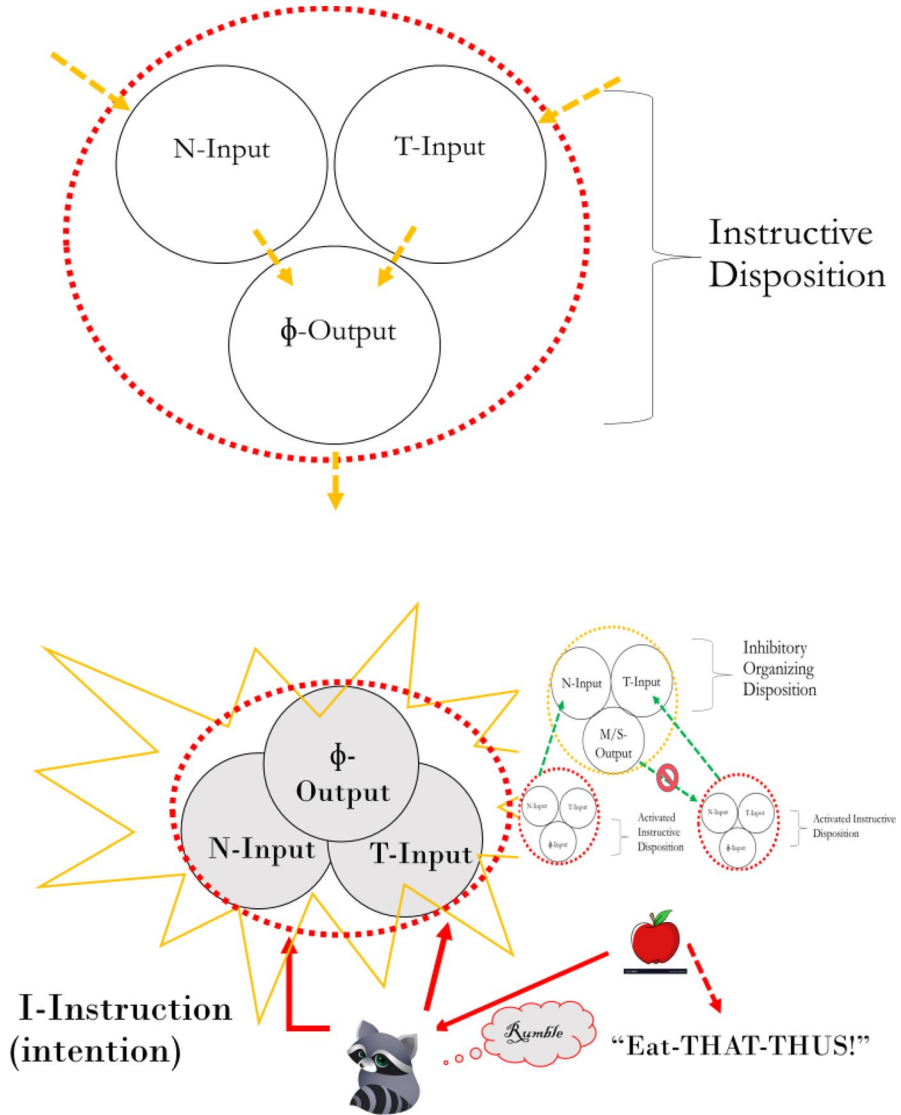
Fig. 9 I-instruction

fundamental intentional directedness to things in the world (e.g. targets) is generative directedness. One is not, fundamentally, intentionally directed at one's own actions by observation, or by virtue of prior intentional states that represent the action, or by its outcome. Rather, a subject is fundamentally intentionally directed at the world in exactly the same way, and at exactly the same time, as she is fundamentally intentionally directed at her own actions, as being generatively directed at an action is at once to be generatively directed at the world (targets and needs) at which that action is directed.

Fig. 10 Instructive attitudes

Generative directedness is essentially sensory and subjective. It's sensory insofar as T- and N-instructions correspond to perceptions and emotions: T-and N-input sub-capacities are activated via causal impacts on sensory organs; when one T- or N-instructs, one is responding to the causes of such impacts. Embodied instructive representations are, in this way, perceptual experiences of targets and needs; perceptual experience is essentially a matter of sensory capacities potentiating and thus instructing action that's directed at the targets and needs to which they causally respond. While sensory, generative aboutness is also essentially subjective: one does not observe one's potential actions, rather one is in a position to produce them. And, of course, one can only be generativity directed at the very actions one is *oneself* in a position to produce.

5 Anscombe's theses

We believe that embodied instructive representations may help illuminate the five Anscombean theses we mentioned in our introduction. First, the distinction between embodied instructive representations and surrogative representations provides a way of interpreting Anscombe's distinction between speculative knowledge "derived from the objects known," and practical knowledge, which "is the cause of what it understands." (Anscombe, 1957, §87). Indeed, on Springle's account, an agent's generative grasp of an intentional action is constituted by and therefore inseparable from the potentiation of that action. Moreover, embodied instructive representations are generative modes of presentation of actions; an action is not presented as an object but rather an agent is presented with the world insofar as there's an embodied instructive representation (potentiated action) that's directed at it. In this way, we think generative aboutness provides a potential interpretation of the first two

Anscombean theses: the idea that practical knowledge is a kind of knowledge of an agent's action that is not by observation (Anscombe, 1957, §29, §51, §57) and that it is the cause of what it understands (§87–88).

Second, Springle's distinction between the appropriateness of an intentional action qua applied solution and the successful performance of an intentional action (or successful actualization or execution of an applied solution) helps capture Anscombe's distinction between errors in practical judgment and errors in the performance of an action. Consider Anscombe's example of the shopping husband. Anscombe observes that the husband's list expresses his intention, but the intention (list) may also be expressed as a kind of order (in our terms, an instruction) if his wife gave it to him. Anscombe contrasts the list qua intention to a list produced by a detective following the man, which records each item the husband buys. Anscombe asks, "What then is the identical relation to what happens, in the order and the intention, which is not shared by the record?" (Anscombe, 1957, p. 56). Anscombe suggests that while the detective would make a mistake if his record failed to record the facts, if the man with the shopping list makes a mistake by, say, buying margarine when butter is on the list, this is an error in performance. Only if he had put something on the list that is inappropriate to his ends would there be an error in his list *qua* an expression of his intention (or his wife's order), i.e. *qua* practical judgment.[30]

The standard way of understanding this example takes it that the detective's list is an indicative propositional representation (belief) while the man's list is an imperative propositional representation (desire). But if an intention or practical knowledge is a propositional attitude, then the success of an intention or practical knowledge is a matter of the truth or satisfaction of its propositional content. So if the man buys margarine when butter is on the list, this would presumably count as a failure of intention or practical knowledge. But Anscombe says it's not. It's therefore not obvious that propositional attitudes can comfortably account for the difference between the two kinds of error countenanced in Anscombe's example. In contrast, Springle's account suggests that the list expresses an embodied instruction in a surrogative form, while the record is the detective's surrogative representation of the man's actions. And Springle's account makes a clean distinction between these two different kinds of error in terms of the distinction between the appropriateness of an embodied instructive representation and the success of the performance of the intentional action it generatively represents. In this respect, her account picks up on Anscombe's insight that there are different success states accorded to an intentional action qua practical judgment and its performance.

While a propositional judgment is true just in case it represents facts correctly, Anscombe notes that ancient and medieval philosophers distinguished between such contemplative truth, and appropriateness or "practical truth" that is not reducible to contemplative knowledge but is to be construed as "truth in agreement with right desire" (Anscombe, 1981, p. 76). Anscombe finds such a view in Aristotle, who explicitly distinguishes between contemplative thought and the "truth" of actions:

---

[30] For one compelling account of this distinction, see Frost (2014).

As affirmation and denial are in thinking, so pursuit and avoidance in desiring. Hence as much as (ὥστ' ἐπειδὴ) as character virtue is a dispositional state of one's choice (ἕξις προαιρετική) and choice is deliberate desire (προαίρεσις ὄρεξις βουλεθτική), on account of this it is necessary that the reason be true and the desire right, if the choice is to be good, and that desire must pursue the same things that the reason affirms. Indeed, this is thought, and practical truth (ἡ ἀλήθεια πρακτική). While in theoretical thought, which is not concerned with action or production, doing "well" is the true and "badly" is false (for this is the function of all thinking); yet the function of practical thinking is the attainment of truth in agreement with right desire. (*Nicomachean Ethics*, 1139a21-31).

Though the ultimate, correct interpretation of this passage is the matter of debate, for our purposes it is sufficient to observe that Aristotle here divides practical and theoretical thinking by means of a distinction between their success states.[31] It is the function (ἔργον) of theoretical or speculative thinking to hit on factual truth. When one reaches a true judgment, one's theoretical reason is considered "good" in virtue of this function being satisfied. But contemplative truth is not the success state of practical thinking. Rather, practical thinking is right only when it corresponds to right desire, that is, desire formed on the basis of character virtues like courage, moderation, and justice. Given Aristotle's view that possession of these virtues is necessary for a flourishing human life, it is reasonable to conceive of the actualization of these virtues in producing practical truth as contributing to one's flourishing.

This Aristotelian doctrine does not, on its own, answer the modern question of how best to characterize intentional action. But we think that Springle offers a framework in which it can be brought to bear on that question. Indeed, we hope to have offered the beginning of a remedy to the syndrome of modern action theory that Anscombe referred to as the "incorrigibly contemplative conception of knowledge" (Anscombe, 1957, p. 57). Because PCT depends on fitting together the heterogeneous elements of psychological representations and embodied actions, it reproduces the Cartesian problem of how to account for bodily effects by means of mental causes. But replacing propositional attitudes with larger actions or mere dispositions for actions seems to omit any interesting explanatory role for an agent's psychological representations understood as her first-personal perspective, grasp, or understanding of her situation.

We think Springle's account may bypass this conceptual quagmire because it is, in the spirit of Aristotle, Anscombe, and arguably Ryle, essentially hylomorphic: intentional actions are material processes with the form of (local, application-accountable) solutions.[32] And mental representations are identical to the potentiations of processes with said form. On this view, an agent's psychological representations are explanatory of what she does because they constitute her being in a position to do

---

[31] For a comprehensive discussion of the various interpretations that have been offered, see Olfert (2017).

[32] See also Springle and Grush (2018).

them. Moreover, where an intentional action is locally inappropriate, the embodied instructive representation that generatively represents it—that is its potentiation—is locally inappropriate, owing to a problem in the realizing instructive disposition's T- or N-input sub-capacities or both. We think, then, that Springle's view may offer a structure in which to characterize Aristotelian-Anscombean intentions in the terms of contemporary philosophy.

6 Conclusion

IA depends on the PCT, which, we argued, faces two serious problems. Springle's account provides an alternative to PCT that avoids these problems and thereby provides a superior account of intentional action, REASONS and CONNECTION to boot.

First, because Springle's problem-solving analysis of intentional action is a version of AAT, it avoids generating the CT constraint. Springle is therefore in a position to analyze representation in terms of PEA and PEA in terms of problem-solving abilities that position agents to produce intentional actions. Indeed, Springle's PEA analysis of representation provides a way of unifying different species of representation, including propositional (surrogative) representations and her problem-solving account of intentional action enables her to unify her analysis of intentional action with her analysis of representation.

Second, as a result of the PEA analysis, mental representations are necessarily connected with intentional actions. Indeed, embodied instructive representations provide an account of intentions (I-instructions), desires (N-Instructions) and means-end beliefs (T-instructions), without construing these psychological states as propositional attitudes. On Springle's account, psychological representations are identical to (full or partial) potentiations of intentional actions: an agent's situation activates her instructive dispositions where this contributes to the determination of a potentiated intentional action that is directed at that situation. That situation figures into whether the potentiated intentional action and thus psychological representation is appropriate, and the potentiation of an intentional action that is directed at a situation constitutes an agent's generative awareness of, i.e. her direct PEA to, a situation. We think, then, that Springle's notion of generative aboutness can account for the sense in which an agent's psychological representations constitute her first-personal perspective or grasp of her situation.[33]

Third, embodied instructive representations can account for the sense in which an agent's psychological representations—what she believes, desires, and intends—are reflected in her intentional actions in a way that's compatible with AAT (in Springle's case, the problem-solving account) and therefore in a way that does not render her account of action susceptible to counterexamples involving causal deviance. Embodied instructive representations are protected from deviance because they are

---

[33] An agent's individual history is reflected in, because it contributes to, the nature of her problem-solving dispositions.

not composed out of more basic beliefs and desires with modally separable contents. Instead, beliefs and desires (T-instructions and N-instructions) are modes of instructively intending. And instructive intentions potentiations of intentional actions, or intentional actions qua potentiated, so it is not possible for an embodied instructive intention to manifest in a non-intentional action.

Finally, while embodied instructive representations can explain sophisticated activities such as complex, planned actions, or acts of judgment, they are also plausibly attributed to human infants and non-human animals. For those find it prima facie plausible that infants and non-human animals act intentionally and can possess skills, but doubt that they are capable of propositional thought, this provides yet another reason to prefer Springle's account over PCT and in particular, embodied instructive representations over propositional attitudes.[34]

All that is required for the anti-intellectualist to resist the intellectualist conclusion that intentional action and thus skill always necessarily depend on propositional knowledge is the existence of an alternative account. Yet if we are right to think that Springle's account of action and psychological representation is superior to that given by PCT, then the anti-intellectualist is in a strong position to reject P2 of the Intentional Action Argument. Moreover, Springle's account provides prima facie support for hardcore anti-intellectualism. Since paradigmatic surrogative representations (e.g. sentences in a public language) are the products or targets of intentional actions, agents must have embodied instructive representations that are the potentiations of the intentional activities in order to produce and/or exploit them.

We've argued that embodied instructive representations, as opposed to propositional attitudes, may better account for an agent's reasons for action, practical knowledge, and the psychological states that constitute intentions more generally. To the extent that we've succeeded, we've blocked the defense of P2. But more generally, and perhaps more importantly, we hope we've persuaded the reader that anti-intellectualists have a promising, indeed potentially quite powerful, framework at their disposal.

---

[34] See Springle (2019) for a related argument.

References

Anderson, M., & Rosenberg, G. (2008). Content and action: The guidance theory of representation. *The Journal of Mind and Behavior*, *29*(1/2), 55–86.

Anscombe, G. E. M. (1957). *Intention*. Harvard University Press.

Anscombe, G. E. M. (1979). Under a description. *Nous*, *3*(2), 219–233.

Anscombe, G. E. M. (1989). Practical inference. In M. Geach & L. Gormally (Eds.), *Human life, action and ethics.* (Vol. 2005, pp. 109–147). Imprint Academic.

Anscombe, G. E. M. (1981). Thought and action in aristotle: What is 'practical truth'? In G. E. M. Anscombe (Ed.), *From parmenides to wittgenstein.* Basil Blackwell.

Audi, R. (1986). Acting for reasons. *Philosophical Review, 95*(4), 511–546.

Bishop, J. (1989). *Natural agency*. Cambridge: Cambridge University Press.

Brandom, R. (1994). *Making it explicit: Reasoning, representing, and discursive commitment*. Harvard University Press.

Brandom, R. (2000). *Articulating reasons: An introduction to inferentialism*. Harvard University Press.

Bratman, M. (1987). *Intentions, plans, and practical reasoning*. Harvard University Press.

Butterfill, S. A., & Sinigaglia, C. (2014). Intention and motor representation in purposive action. *Philosophy and Phenomenological Research, 88*(1). https://doi.org/10.1111/j.1933-1592.2012.00604.x.

Burge, T. (2010). *Origins of objectivity*. Oxford University Press.

Castañeda, H. N. (1975). *Thinking and doing. The philosophical foundations of institutions*. Reidel.

Cummins, R. (1996). *Representations, targets, and attitudes*. MIT Press.

Cussins, A. (1992). Content, embodiment and objectivity: The theory of cognitive trails. *Mind*, *101*(404), 651–688.

Davidson, D. (1963). Actions, reasons and causes. *Journal of Philosophy*, *60*, 685–700.

Davidson, D. (1970). Mental events. In L. Foster & J. W. Swanson (Eds.), *Experience and theory*. London: Duckworth.

Davidson, D. (1980). *Essays on Actions and Events*. Oxford: Oxford University Press.

Dretske, F. (1988). *Explaining behavior: Reasons in a world of causes*. MIT Press.

Frankfurt, H. (1978). The problem of action. *American Philosophical Quarterly, 15*, 157–162.

Fridland, E. (2014). They've lost control: Reflections on skill. *Synthese, 191*(12), 2729–2750.

Fridland, E. (2015a). Knowing-how: Problems and considerations. *European Journal of Philosophy, 23*, 703–727.

Fridland, E. (2015b). Automatically minded. *Synthese*. https://doi.org/10.1007/s11229-014-0617-9.

Fridland, E. (2017). Skill and motor control: Intelligence all the way down. *Synthese, 174*(6), 1539–1560.

Frost, K. (2014). On the very idea of direction of fit. *Philosophical Review, 123*(4), 429–484.

Ginet, C. (1990). *On action*. Cambridge University Press.

Gibbons, J. (2001). Knowledge in action. *Philosophy and Phenomenological Research, 62*(3), 579–600.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.

Goldman, A. (1968). Actions, predictions, and books of life. *American Philosophical Quarterly, F5*(3), 135–151.

Goldman, A. (1970). *A theory of human action*. Prentice Hall.

Grush, R. (2007). Skill Theory v.2.0. *Synthese*, *159*(3), 389–416.

Hawley, K. (2003). Success and knowledge-how. *American Philosophical Quarterly, 40*, 19–31.

Levy, N. (2017). Embodied savoir faire: Knowledge-how requires motor representations. *Synthese, 194*(2), 511–530.

Mandik, P. (2005). Action-oriented representation. In K. Akins & A. Brook (Eds.), *Cognition and the brain: The philosophy and neuroscience movement*. Cambridge: Cambridge University Press.

McLaughlin, B. (2012). Why Rationalization is not a species of causal explanation. In J. D'Oro (Ed.), *Reasons and causes: Causalism and anti-causalism in the philosophy of action.* Palgrave McMillan.

Mele, A. R. (1992). *Springs of action: Understanding intentional behavior*. Oxford University Press.

Mendelovici, A. (2018). *The phenomenal basis of intentionality*. OUP.

Millikan, R. (1984). *Language, thought and other biological categories*. Cambridge: MIT Press.

Millikan, R. (1995). Pushmi-pullyu representations. *Philosophical Perspectives, 9*, 185–200.

Mylopoulos, M., & Pacherie, E. (2019). Intentions: The dynamic hierarchical model revisited. *Wiley Interdisciplinary Reviews: Cognitive Science, 10*, e1481. https://doi.org/10.1002/wcs.1481.

Nanay, B. (2013). *Between perception and action*. Oxford University Press.

Olfert, C. M. M. (2017). *Aristotle on practical truth*. Oxford University Press.

O'Regan, K., & Noe, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, *24*, 939–1031.

Pautz, A. (2007). Intentionalism and perceptual presence. *Philosophy of Mind, 21*(1), 495–591.

Pavese, C. (2015a). Knowing a rule. *Philosophical Issues, 25*(1), 165–188.

Pavese, C. (2015b). Practical senses. *Imprint, 15*(29), 1–25.

Pavese, C. (2016a). Skill in epistemology I: Skill and knowledge. *Philosophy Compass, 11*(11), 642–649.

Pavese, C. (2016b). Skill in epistemology II: Skill and know how. *Philosophy Compass, 11*(11), 650–660.

Pavese, C. (2017a). Know-how and gradability. *Philosophical Review, 126*(3), 345–383.

Pavese, C. (2017b). A theory of practical meaning. *Philosophical Topics, 2*, 65–96.

Pavese, C. (2019). The psychological reality of practical representation. *Philosophical Psychology, 32*(5), 784–821.

Peacocke, C. (1979). *Holistic explanation: Action, space, and interpretation*. Oxford University Press.

Roth, A. (2000). Reasons explanation of actions: Causal, singular, and situational. *Philosophy and Phenomenological Research, 59*, 839–874.

Rowlands, M. (2006). *Body language*. MIT Press.

Ryle, G. (1949). *The concept of mind*. Routledge.

Sehon, S. (1997). Deviant causal chains and the irreducibility of teleological explanation. *Pacific Philosophical Quarterly, 78*, 195–213.

Sehon, S. (2005). *Teleological realism: Mind, agency, and explanation*. Cambridge MA: MIT Press.

Setiya, K. (2007). *Reasons without rationalism*. Princeton University Press.

Setiya, K. (2008). Practical knowledge. *Ethics*, *118*(3), 388–409.

Setiya, K. (2009). Reasons and causes. *European Journal of Philosophy, 19*(1), 129–157.

Setiya, K. (2011a). Knowing how. *Proceedings of the Aristotelian Society, 112*(3), 285–307.

Setiya, K. (2011b). Knowledge of intention. In A. Ford, J. Hornsby, & F. Stoutland (Eds.), *Essays on Anscombe's intention.* (pp. 170–197). Harvard University Press.

Setiya, K. (2016). Anscombe on practical knowledge. In *Practical knowledge: Selected essays*. Oxford University Press.

Siegel, S. (2014). Affordances and the contents of perception. In B. Brogaard (Ed.), *Does perception have content?* (pp. 51–75). Oxford University Press.

Snowdon, P. (2004). Knowing how and knowing that: A distinction reconsidered. *Proceedings of the Aristotelian Society, 104*(1), 1–29.

Springle, A. (2019). Methods, minds, memory, and kinds. *Philosophical Psychology, 32*(5), 634–660.

Springle, A. (2021). Dissertation manuscript.

Springle, A., & Grush, R. (2018). Agency, perception, space & subjectivity. *Phenomenology and the Cognitive Sciences, 18*(5), 799–818.

Stalnaker, R. (2012). Intellectualism and the objects of knowledge. *Philosophy and Phenomenological Research, LXXXV*(3), 754–761.

Stanley, J. (2011). *Know how*. Oxford University Press.

Stanley, J., & Krakauer, J. W. (2013). Motor skill depends on knowledge of facts. *Frontiers in Human Neuroscience, 7*, 503. https://doi.org/10.3389/fnhum.2013.00503.

Stanley, J., & Williamson, T. (2001). Knowing how. *The Journal of Philosophy, 98*(8), 411–444.

Stout, R. (1996). *Things that happen because they should: A teleological approach to action*. Oxford University Press.

Stout, R. (2014). *Action*. Routledge.

Thompson, M. (2008). *Life and action*. Harvard University Press.

Wald, B., & Tenenbaum, S. (2018). Reasons and action explanation. In D. Star (Ed.), *The Oxford handbook of reasons and normativity*. Oxford, UK.

Wilson, G. (1989). *The intentionality of human action*. Stanford University Press.