Erkenntnis
https://doi.org/10.1007/s10670-019-00174-7

ORIGINAL RESEARCH

Discovering Patterns: On the Norms of Mechanistic Inquiry

Lena Kästner[1,3] · **Philipp Haueis**[2,3]

Received: 1 November 2018 / Accepted: 16 September 2019

Abstract

What kinds of norms constrain mechanistic discovery and explanation? In the mechanistic literature, the norms for good explanations are directly derived from answers to the metaphysical question of what explanations are. Prominent mechanistic accounts thus emphasize either ontic (Craver, in: Kaiser, Scholz, Plenge, Hüttemann (eds) Explanation in the special sciences: the case of biology and history, Springer, Dordrecht, pp 27–52, 2014) or epistemic norms (Bechtel in Mental mechanisms: philosophical perspectives on cognitive neuroscience, Routledge, London, 2008). Still, mechanistic philosophers on both sides agree that there is no sharp distinction between the processes of discovery and explanation (Bechtel and Richardson in Discovering complexity. Decomposition and localization as strategies in scientific research, MIT Press, Cambridge, 2010; Craver and Darden in In search of mechanisms: discoveries across the life sciences, University of Chicago Press, Chicago, 2013). Thus, it seems reasonable to expect that ontic and epistemic accounts of explanation will be accompanied by ontic and epistemic accounts of discovery, respectively. As we will show here, however, recent discovery accounts *implicitly* rely on both ontic and epistemic norms to characterize the discovery process. In this paper, we develop an account that makes explicit that, and how, ontic and epistemic norms work together throughout the discovery process. By describing mechanism discovery as a process of pattern recognition (Haugeland, in: Having thought. Essays in the metaphysics of mind, Harvard University Press, Cambridge, pp 267–290, 1998) we demonstrate that scientists have to develop epistemic activities to distinguish a pattern from its background. Furthermore, they have to determine which epistemic activities successfully describe how the pattern is implemented by identifying the pattern's components. Our approach reveals that ontic and epistemic norms are equally important in mechanism discovery.

Lena Kästner and Philipp Haueis have contributed equally to this work.

✉ Lena Kästner
lena.kaestner@uni-saarland.de

[1] Department of Philosophy, Saarland University, 66123 Saarbrücken, Germany

[2] Department of Philosophy, Bielefeld University, P.O. 100131, 33501 Bielefeld, Germany

[3] Berlin School of Mind and Brain, Humboldt-Universität zu Berlin, 10117 Berlin, Germany

1 Introduction

What kinds of norms guide mechanistic inquiry, viz. the process of discovering mechanisms and developing mechanistic explanations? Philosophical accounts that answer this question should be both descriptively and normatively adequate (Machamer et al. 2000). They should not only describe the actual practice of mechanistic inquiry but also explicate what general principles it follows. While prominent accounts of mechanism discovery (Bechtel and Richardson 2010, Craver and Darden 2013) adequately describe the search for mechanisms, we think that they do not sufficiently explicate the role that different norms play in mechanistic inquiry. In this paper, we develop our own account—the *pattern account*—that both incorporates the descriptive strengths of these accounts and elucidates that, and how, different norms contribute to mechanistic inquiry.

Existing accounts focus on different norms in mechanism discovery. Bechtel and Richardson (2010) emphasize that when researchers develop mechanistic explanations, they aim to increase the *intelligibility* of a phenomenon. Craver and Darden (2013) emphasize that when searching for mechanisms, researchers aim to uncover the *causal structure* relevant to the phenomenon *accurately* and *completely*. Both "Increase intelligibility!" and "Get the causal structure right!" are examples of norms of mechanistic inquiry, i.e. general instructions of how to search for mechanisms. Intelligibility is usually considered an *epistemic* norm, while accuracy and completeness are considered *ontic* norms (Illari 2013). The reason for this terminology is that these norms can be derived from different accounts of what explanations *are*. The norm of intelligibility can be derived from the *epistemic* account, according to which explanations are representations, e.g. texts or diagrams (Bechtel 2008; Wright 2012). According to this account, good explanations render the phenomenon more intelligible than bad ones. The norms of accuracy and completeness can be derived from the *ontic* account, according to which explanations are things in the world. According to this account, "[o]bjective explanations, the causes and mechanisms in the world, are the correct starting point in thinking about the criteria for evaluating explanatory texts" (Craver 2007, p. 27). Hence, good explanatory texts describe the mechanism responsible for the phenomenon as accurately and completely enough for the purposes of explanation (Craver and Kaplan 2018). There is substantive debate about which account captures the nature of explanation best. Discussions on this topic are commonly placed under the heading "ontic-epistemic debate".

While we agree that "ontic" and "epistemic" are important descriptive classifiers which distinguish different norms, our project in this paper is not concerned with the question of what the correct metaphysical account of explanation is. It is not our purpose to take sides in the ontic-epistemic debate. Rather, we are interested in how mechanistic inquiry is *normatively* and *methodologically constrained*. How does the use of scientific methods, tools and skills constrain the search for mechanisms? And how does searching for a particular mechanism constrain which methods, tools and skills researchers use? When answering these questions, we follow philosophers of science who privilege methodological over metaphysical problems when analyzing scientific practice (see Chang 2014; Rouse 2015a; Bursten 2018). Although metaphysical concerns about explanation can affect our answer to methodological questions, they cannot replace these questions altogether. Indeed, mechanistic philosophers of science already embrace a practice-based view as they abandon the logical positivist distinction between the "context of discovery" and the "context of justification" (Reichenbach 1938). Bechtel and Richardson, for example, explicitly motivate their account of mechanistic discovery by attacking such a distinction as artificial (2010, p. 17). Likewise, Craver and Darden (2013, p. 189) explicitly deny that, at least for the life sciences, "distinct norms of explanation and discovery apply". Because we agree that a smooth transition from discovery to explanation adequately describes scientific practice, we subsume mechanism discovery and mechanistic explanation under "mechanistic inquiry". Note that we do not claim that there are no differences between discovery and explanation. Indeed, as both of us have argued elsewhere, there are important philosophical debates to be had about scientific discovery that cannot be answered by consulting philosophical accounts of explanation (e.g. the role that different kinds of experiment play in the description and delineation of the targets of research, see Kästner 2017; Feest 2017; Haueis 2018). We use "mechanistic inquiry" because we are interested in the normative aspects shared by discovery and explanation, not because we think that the search for (explanatory) mechanisms is the only goal of scientific discovery. By fleshing out which norms are operative in mechanistic inquiry, we push the practice-based view further than existing accounts of mechanism discovery.

We believe that Craver and Darden's emphasis on ontic and Bechtel and Richardson's appeal to epistemic norms both pick out important aspects of mechanistic inquiry. Indeed, as we will detail below, their accounts of mechanism discovery implicitly rely on ontic *as well as* epistemic norms. This implicit reliance is a normative weakness, despite the descriptive merit of distinguishing ontic and epistemic norms in the first place. An adequate account of mechanistic inquiry should not only distinguish ontic and epistemic norms but also make explicit how they work together in practice.

In Sect. 2, we introduce the discovery accounts by Bechtel and Richardson (2010) and Craver and Darden (2013). We show that they both implicitly rely on epistemic and ontic norms to characterize the discovery process. Crucially the different norms also interact with one another. To make explicit that, and how, ontic and epistemic norms work together, we develop an account that describes mechanism discovery as a process of *pattern recognition* (Sect. 3). The key ideas of our pattern account are that (1) skilled scientists who use different instruments, and concepts/models form a *pattern recognition practice*, that (2) such practices isolate patterns from their surroundings and experimental noise via different *epistemic activities* such as modeling or experimentation, and that (3) elements within those patterns are tracked via different *epistemic operations*.[1] Our account reveals that ontic and epistemic norms are equally important in mechanism discovery. Section 4 discusses the benefits of our view, illustrating it with concrete examples. Section 5 concludes and discusses the wider-ranging implications of our view for the so-called ontic-epistemic debate about the nature of explanation.

2 Mechanism Discovery Involves Both Epistemic and Ontic Norms

When assessing mechanistic discovery accounts, we think it is useful to distinguish between *norms* and *normative constraints*. Norms are best understood as general instructions of how to search for mechanisms and how to construct good mechanistic explanatory texts. We distinguish between “ontic” and “epistemic” norms according to the *content* of these general instructions. As an example of an epistemic norm for successful mechanistic inquiry we consider the instruction to increase intelligibility of the phenomenon to be explained (Haugeland 1998, ch. 9; Bechtel 2008).[2] As examples of ontic norms we consider the instructions to describe the causal structure of a mechanism accurately and completely enough for the purposes of explanation. We discuss the ontic norms of accuracy and completeness together because they capture complementary aspects of ontic adequacy (Craver and Darden 2013, ch. 6).

Following ontic and epistemic norms can be achieved by using specific *normative constraints*. Different such constraints are the *determinates* of the determinable epistemic norm of intelligibility or the ontic norms of accuracy and completeness, respectively. They serve to limit the search space for mechanisms in different ways. Suppose we are searching for an explanation of how a neuron fires action potentials. The heuristic of decomposition is an epistemic constraint that determines the norm of intelligibility in a specific way. Using this heuristic limits the search space to explanations approximating the behavior of the system based on the behavior of independent but interacting components (cf. Bechtel and Richardson 2010, p. xxix). Another epistemic constraint is the role of abstraction from causal detail for the purpose of constructing a model of the system that is mathematically tractable (Boone and Piccinini 2016, p. 5). In contrast, spatial or temporal constraints are ontic constraints that determine the ontic norm of accuracy and/or completeness in a certain way (cf. Craver and Darden 2013, p. 31). They limit the search space to explanations that contain an already identified component of a certain shape, or an already identified activity of a certain duration (e.g. a neurotransmitter being released into

---

[1] We adopt the terms “epistemic activity” and “epistemic operation” from Chang (2014). For explanations see Sect. 3.2.

[2] Sheredos (2016) argues that another epistemic norm of mechanistic inquiry is the instruction to achieve generality, e.g. by categorizing token entities and activities in the same mechanism into types. We do not discuss generality here, because we want to allow for the possibility that describing one-off mechanisms can make a phenomenon intelligible without achieving generality. Besides, Sheredos’ account seems to imply that generality is a variety of the norm of intelligibility, because following this norm makes the scope of an explanatory text intelligible.

the synaptic cleft and being taken up again within 5 ms). Another ontic constraint is the mechanism-to-model-mapping (3M) constraint, which limits the search space to models whose variables map onto components and causal relations within the target mechanism (Kaplan and Craver 2011, p. 611). Importantly, epistemic and ontic constraints can be combined in mechanistic inquiry, e.g. by considering only those among the mathematically tractable models of a phenomenon whose variables map onto some of the causal structure of the mechanism.

The existing accounts of mechanism discovery evaluate the importance of epistemic and ontic constraints differently. Bechtel and Richardson (2010) argue that "while empirical [i.e. ontic] constraints are important to discovery, they are not sufficient" (cf. ibid., p. 5) to decide which of the candidate explanations is the most plausible one. Therefore, they think that mechanism discovery must primarily be driven by epistemic constraints. Cognitive strategies, choice points and the limited capacities of scientific cognizers take center stage in their account. Using heuristics (localization, decomposition) limits the search space and thus makes the task of developing mechanistic explanations epistemically tractable. For instance, cell physiologists and cognitive neuroscientists often divide a system into parts and operations by assuming that the system is nearly decomposable; that is, it consists of independent but interacting components. If the system is in fact nearly decomposable the heuristic assumption will lead to a successful explanation of the system's behavior. If it is not, the heuristic will fail systematically, which tells researchers they have to change their heuristic and search for a different explanation. Given that the success or failure of discovery heuristics implies the success or failure of mechanistic explanation, it appears that the same set of norms governs both mechanistic discovery and explanation. According to Bechtel and Richardson, mechanistic inquiry is primarily guided by epistemic constraints.

By contrast, Craver and Darden (2013) emphasize that "[t]he product of the search for mechanisms shapes the process by which it is discovered" (ibid., p. 7). If they search for a mechanism *producing* a phenomenon, scientists often start from the final product and search for the activities by which the mechanism's entities are transformed into the product. If they search for a mechanism *underlying* the phenomenon, they typically break down a system into its working parts to show how these parts are organized to give rise to the phenomenon to be explained. If scientists search for mechanisms *maintaining* a phenomenon, they search for factors that disturb the phenomenon as well as those correcting for the disturbances. In all of these cases, mechanism discovery iterates through different stages (phenomenon characterization, mechanism sketch, mechanism schema, evaluation, re-characterization,…) until the mechanism has been described completely enough for the purposes of explaining the phenomenon under investigation. At each stage, scientists build on known facts about entities, activities and their organization to construct, evaluate and revise their mechanistic explanations. Their cognitive or epistemic capacities are at best secondary. Craver and Darden's account therefore entails that mechanistic inquiry is primarily guided by ontic constraints.

While Bechtel and Richardson primarily emphasize epistemic constraints and Craver and Darden focus on ontic ones, we argue that both accounts implicitly rely on *both* ontic and epistemic norms to describe mechanism discovery. Bechtel and

Richardson (2010, p. 10) aim to tell a *normative* story about when discovery heuristics succeed or fail. But as an epistemic device that constrains the search space, a heuristic is itself neither successful nor faulty. Normativity only enters the picture when the heuristic is applied to a system in the *world*. Whether it succeeds or fails depends on how the system is *actually* organized. This is clearly an ontic constraint. Craver and Darden (2013, ch. 6) tell a normative story about how researchers use *ontic* constraints to sort complete and/or accurate from incomplete and/or inaccurate mechanistic explanations. Much like Bechtel and Richardson describe mechanism discovery as proceeding along a number of choice points, Craver and Darden conceptualize discovery as multiple successive evaluations and revisions of preliminary mechanism descriptions. Essentially, then, mechanism discovery is a trial and error process. So even in Craver and Darden's account, ontic constraints can only become operative in scientific practice once researchers engage in the epistemic process of trying—and failing—to make the system's behavior intelligible. Thus, without ontic norms, epistemic constraints are normatively inert. Conversely, without epistemic norms, ontic constraints are dormant. To sum up: Despite differences in emphasis, both accounts already *implicitly* rely on both ontic and epistemic norms when describing mechanism discovery.

We are not the first to reach this conclusion. Illari (2013) has already argued for the relevance of ontic and epistemic norms and normative constraints, respectively, for building and evaluating mechanistic explanations. While she focuses on ontic and epistemic conceptions of *explanation*, she also recognizes the importance of ontic and epistemic constraints for mechanism discovery:

> The real achievement of mechanistic (and possibly other forms of) explanation is satisfying both ontic and epistemic constraints simultaneously, to get a story constrained by all the empirical contact with the world that ingenuity can design; a story that we can understand, manipulate and communicate, that we can use, and use collaboratively, to help us manipulate, control and predict the world—and lead science to better knowledge. This is the ongoing challenge of mechanism discovery. (ibid., p. 253)

While we wholeheartedly agree with this statement, we think that more should be said about how this "ongoing challenge" can be met during mechanism discovery. Illari's integrative project primarily seeks to explicate the interplay of ontic and epistemic constraints that already operates implicitly in different accounts of mechanistic *explanation* (cf. ibid., p. 244; p. 248). We continue her project by explicating how ontic and epistemic norms work together in mechanism *discovery*—and thus throughout mechanistic inquiry as a whole. Our account also adds further detail to Illari's integrative project by making explicit that different normative constraints (decomposition, 3M constraint etc.) are determinates of determinable ontic and epistemic norms.[3] The pattern account outlined in Sect. 3 puts this distinction to

---

[3] Illari seems to distinguish between norms and constraints: "Each kind of constraint alone gives us some kind of useful set of *norms* for evaluating, and attempting to build, mechanistic explanations." (ibid., p. 253, emphasis added) While she does not explicitly define "norms", we suspect that Illari means what we call "normative constraints" while her constraints correspond to that we call "norms". However, not much hinges on this terminological difference, since Illari does not put her distinction to work.

work. The examples we discuss below not only highlight *that* researchers consult both ontic and epistemic norms when searching for mechanisms but also *how* they work together to constrain the search space in different ways. Thus, we argue, the pattern account is superior to existing accounts of mechanism discovery, because it shows e*xactly how* ontic and epistemic norms work together throughout mechanistic inquiry, rather than implicitly relying on one of them.

Consider again Bechtel and Richardson's account. They argue that if a heuristic assumption fails to adequately describe the organization of the system under investigation, this mismatch "should be felt in failures of explanation" (2010, p. 32). Beyond this general acknowledgment of the norm of ontic accuracy, however, they do not tell a story about how this norm is used to evaluate discovery heuristics. For example, mechanistic inquiry needs to be constrained from the bottom up, i.e. by evidence about the spatiotemporal characteristics of the system's components. Are these ontic bottom-up constraints simply *objective ontic facts* about the size of an entity or the duration of an activity in the mechanism (rather than *epistemically interpreted findings*)? If they are objective ontic facts, Bechtel and Richardson's account seems to require a *non-inferential* notion of observation; for skillful practitioners need to have direct access to facts about the entities and activities to ontically constrain their search for mechanisms, e.g. by seeing the size of an ion channel through an electron-microscope.[4] Alternatively, bottom-up constraints may refer to *epistemically interpreted findings*, i.e. descriptions of ontic facts whose accuracy is determined by epistemic standards. In this case, Bechtel and Richardson's account requires an *inferential* notion of observation; for researchers must use *theories* of the system to determine whether the data obtained constitutes evidence about components within the system.[5] The distinction between these two notions of observation makes a difference for how bottom-up constraints limit the search for mechanism. On a non-inferential account, it is the ontic facts *themselves*, whereas on an inferential account, it is the facts as they are *made intelligible by theory* which limit the search space for mechanisms. The non-inferential reading implies a direct (i.e. theory-independent), whereas the inferential reading implies an indirect (i.e. theory-dependent) role of ontic constraints in mechanism discovery. Because Bechtel and Richardson only implicitly rely on ontic norms, their account does not disambiguate between the direct and indirect roles of ontic bottom-up constraints in mechanism discovery. Such a disambiguation, however, is needed to explicate why and how ontic and epistemic norms work together in the discovery process (Sect. 4.1).

Craver and Darden (2013) also do not make the interaction of ontic and epistemic norms explicit. They emphasize that throughout discovery researchers repeatedly encounter *anomalies*, viz. empirical findings that seemingly conflict with the currently most plausible model of a mechanism (cf. ibid., p. 144). While anomaly resolution is a crucial motor for scientific discovery, anomalous findings can indicate anything from a mundane measurement error to a significant mismatch between the current model and the causal structure of the world.

---

[4] For a non-inferential account of microscopic observation see Hacking (1981).

[5] The *locus classicus* of an inferential account is van Fraasen (1980).

Depending on the source of the anomaly, resolving anomalies can shape the discovery process in various ways. To resolve an anomalous finding successfully, one needs to know which type of anomaly one is facing. Suppose we are trying to measure an activity in a mechanism (say, using a patch clamp to measure a sodium channel's opening contributing to the action potential), but the measurement took shorter than expected. Does this deviation indicate (a) experimental or measurement errors (e.g. due to a broken patch clamp or a missing experimental control), (b) that the actual duration of the activity deviates from the duration postulated in our model or (c) that we accidently measured a different, previously unknown activity (e.g. the opening of a different ion channel)? This is the classical Duhemian underdetermination problem (Duhem 1906/1954), i.e. the shortened measurement alone does not tell us where the error lies.

Craver and Darden think that (a)-type anomalies can often be resolved through experimental replication (ibid., p. 155), but they do not explain how scientists distinguish between (b)-type and (c)-type anomalies. Since (a)-type anomalies are rooted in methodological problems, their resolution requires consulting epistemic constraints (e.g. checking for errors in measurement or experimental design). By contrast, (b)-type and (c)-type anomalies are *both* rooted in a mismatch between ontic structure and the current mechanistic model. The question remains, however, *how* this mismatch should be resolved. Consider what Craver and Darden write about dealing with an unexpectedly short measurement:

> If an hypothesized stage would be expected to take longer than a time-course experiment shows that it does, then the resulting anomaly indicates the need for a change in the hypothesized mechanism (Darden and Craver 2002, p. 15).

Hence, resolving (b)-type or (c)-type anomalies requires researchers to consult ontic (temporal) constraints. But Craver and Darden's account does not tell researchers how to revise their model to resolve this anomaly—should they posit (b) a different actual duration of sodium channel openings or should they posit (c) a missing ion channel? In other words: should researchers follow the ontic norm of accuracy and revise the duration of sodium channel openings or should they follow the ontic norm of completeness and posit a novel type of ion channel? Since such a decision is undetermined by the experimental evidence, scientists have to *epistemically* constrain—at least—their decision whether an unexpected finding indicates a (b)-type or (c)-type anomaly. However, because Craver and Darden focus predominantly on ontic constraints and rely on epistemic constraints only implicitly, they do not explain when and how the latter guide the evaluation and resolution of anomalies during discovery (see Sect. 4.2 for an extended example).

Taken together, accounts of mechanism discovery that only implicitly rely on either ontic or epistemic constraints leave central methodological questions about mechanistic inquiry unanswered. As such, they cannot paint a clear picture of how ontic and epistemic norms work together in practice. To remedy these issues, we introduce a novel approach that makes explicit how both kinds of norms interact. In what follows we therefore describe mechanism discovery as a *pattern recognition* process.

3 Mechanism Discovery as Pattern Recognition

We think that a one-sided emphasis on either ontic or epistemic norms in mechanism discovery is reinforced by an ontic or epistemic answer the question of what explanations are (Bechtel 2008; Craver 2014). In this paper, however, we want to bracket this metaphysical question. The reason is that methodological questions about ontic constraints on heuristics and epistemic constraints on anomaly resolution cannot be replaced altogether by metaphysical concerns. Since we are interested in the interaction of ontic and epistemic norms, we seek an impartial framework that is (1) not laden with either an ontic or epistemic conception of explanation and (2) closely tied to scientific practice. We therefore describe mechanistic inquiry as a process of *pattern recognition* (cf. Haugeland 1998). Our pattern account allows us to see how ontic and epistemic norms actually work together in mechanistic inquiry.

3.1 Pattern Recognition in Science

Haugeland's (1998) account of patterns and their recognition provides an impartial framework for analyzing mechanistic inquiry we shall draw on. Haugeland emphasizes from the get-go that there are two equally important aspects of scientific normativity. First, science is a *collective achievement* supplying researchers with models and instruments to study the phenomena in their target domain as well as *social and epistemic norms* that regulate how these instruments and models are to be used. Second, scientific norms must apply *beyond our own epistemic situation*; otherwise our explanations will not be accountable to the world—a claim reminiscent of ontic norms. Importantly, Haugeland's account also tells us how these two aspects of normativity work together: The norms of the epistemic collective help scientists to decide which experimental results count as *correct* or *incorrect*. Sorting correct from incorrect results is a precondition for making phenomena in a scientific domain intelligible; for only correct results can enhance intelligibility. "Correctness" here means that the results accord with the epistemic standards of the domain, and not whether the results accurately account for the world. To achieve ontic accountability, scientists need to find out what (entity or activity) the results are evidence *for*. This is where ontic norms enter the picture: "In the relation between entities and ordinary ontical claims about them, the entities are always in the driver's seat." (Haugeland 2013, p. 58) That is to say the truth of a statement depends—in an asymmetric way—on the entity or activity which it is a statement *about*: the entity or activity *determines* the truth of the statement but not vice versa (ibid.). This asymmetric dependence constrains the epistemic standards appropriate for scientific discovery; they are bound to the entities investigated. In other words: scientists must not adopt standards which make it impossible to formulate true statements about the entities and activities to be discovered.[6] Likewise, no ontic structures can be discovered

---

[6] To make this more concrete: What we have in mind here is that scientists should not make assumptions that are highly untenable given their explanatory interests, like, say postulating that the moon is made of cheese when trying to explain its surface structure.

without scientists following epistemic standards. As such, Haugeland's view is *not biased* to either ontic or epistemic norms and the dependence between ontic and epistemic norms is no longer implicit.

The key concept to unlock the inextricable link between ontic and epistemic norms is that of *pattern*. Generally, patterns are discernible regularities. Descriptions of patterns can be used to represent, classify, and process information efficiently. In philosophy, the perhaps best-known piece on patterns is Daniel Dennett's "Real Patterns" (1991). Dennett makes the general point that higher-order patterns are real even though materially they are nothing over and above the units that compose them.[7] Haugeland (1998, ch. 11) develops Dennett's account of patterns by disambiguating two senses of "pattern" at work in the 1991 paper. On the one hand, "pattern" refers to an "orderly or non-random arrangement—the opposite of chaos" (Haugeland 1998, p. 274). On the other hand, "patterns are 'by definition' candidates for discernment or *recognition*" (ibid., p. 273). These two senses of pattern as *orderly arrangements* and *candidates for recognition* are tightly coupled. On the one hand, picking out arrangements that are orderly presupposes that scientists can recognize the elements which are orderly arranged. On the other hand, recognizing individual elements of a pattern presupposes an orderly arrangement in which the elements are assembled.

As orderly arrangements, patterns *persist from below* and as candidates for recognition they are *salient from above* (cf. Haugeland 1998, p. 270f.). This is already reminiscent of the mechanistic view: To say patterns persist from below means they exist in virtue of their implementing lower-level structures. To say they are salient from above means they are recognizable by their higher-level features such as their functional role.[8] Note again that we do not aim to use Haugeland's pattern concept to elucidate the metaphysics of mechanisms (Kaiser and Krickel 2017); or contribute to the debate on the metaphysics of patterns (Ladyman and Ross 2007; Wallace 2003). While these could be interesting projects in their own right, we are here focusing on the question of how scientists' search for mechanisms is constrained by ontic and epistemic norms, respectively. By drawing on the two senses of pattern, our metaphysical commitments are nothing more than what is already inherent in most of the mechanistic literature. These commitments are that (1) there is causal structure in the world and that (2) some of this structure exhibits regularities or phenomena. Mechanistic inquiry aims to identify what brings these phenomena about, i.e. "to carve mechanisms out of the busy and buzzing confusion that constitutes

---

[7] For illustration, Dennett discusses Conway's "Game of Life", a cellular automaton based on a 2D-grid of ON and OFF cells. Once the initial configuration is set, one can start the game and watch how the grid evolves. The evolution is governed by an algorithm which specifies, at each step, for any individual cell whether it will be on or off next time around depending on the cell's current status as well as that of its eight neighbors. As a result, players can see "figures" move across the grid. Strikingly, observers will soon be able to recognize certain "species" and predict their "behavior" without knowing the rules in the algorithm. For Dennett, this illustrates that for higher-level causal generalizations to hold, we do not need to know what lower-level principles govern higher-level regularities. .

[8] Our talk of "higher" and "lower" levels here is compatible with the mechanistic commitment that mechanisms form local nested hierarchies (see also Craver 2007, p. 191f.).

the causal structure of the world." (Craver 2014, p. 140) The mechanisms scientists carve out are patterns in the sense that they are *orderly arrangements* within this causal structure that are *candidates for recognition* during the discovery process that will eventually be described or modeled in mechanistic explanations. Against this background, our focus on the norms operative in mechanistic inquiry highlights the inextricable link between epistemology and ontology: On the one hand, scientific practice with its methods and tools epistemically constrains what patterns in the causal structure of the world can be recognized as mechanisms. On the other hand, patterns in the causal structure ontically constrain which scientific tools will serve to recognize them as orderly arrangements persistent from below and salient from above, respectively. The pattern account explicates this feature of mechanistic inquiry (we elaborate on this in Sect. 3.2).

To illustrate the two senses of "pattern", consider the following (non-mechanistic) example. Picking out walls (orderly arrangements) from a range of material configurations requires us to know that the pattern "wall" (candidate for recognition) is a solid structure made of cuboid elements staggered on top of each other. At the same time, picking out cuboid elements as "bricks" in the pattern "wall" (candidate for recognition) presupposes we are looking at a wall (orderly arrangement) rather than, say, a heap of pebbles. In other words, for the wall pattern to be an orderly arrangement presupposes that it is a candidate for recognition—and vice versa. Therefore, neither sense of "pattern" is conceptually prior.[9]

It is important to note that the orderliness of a pattern as a whole is recognizable without knowing all the parts—even though it is constituted by those parts (cf. Haugeland 1998, p. 277). To continue our non-mechanistic example: a wall is implemented by staggered bricks and can be discriminated from other material configurations (like a heap of pebbles) by its role to support a building. To recognize a wall, we do not need to identify every single brick, neither do we need to specify the specific bond structure. Mechanistic discovery accounts are driven by the same principled idea: recognizing the phenomenon without already knowing the mechanism is the very starting point of mechanism discovery (Sect. 3.2).

In Haugeland's account, recognition refers to an observer's ability to discriminate a pre-specified pattern from its surroundings and from noise (cf. ibid., p. 272). Whereas recognition is a skill of an observer, *recognizability* is a feature of a pattern so recognized. This distinction is crucial for scientific discovery to be possible at all: there are recognizable patterns that scientists cannot yet recognize. To determine which recognizable (salient from above) patterns are orderly arrangements (persisting from below), scientists must determine which patterns are *stably* recognizable. Stability stems from the fact that orderly arrangements exhibit discernable regularities (Sect. 3.2). A pattern is *real* if and only if it is stably recognizable (ibid.).

Recognizing patterns is a normative affair because recognition skills can fail. For instance, we may misrecognize a fence as a wall. To distinguish different

---

[9] A stronger way to read Haugeland is to claim that neither sense of "pattern" is metaphysically prior. For current purposes, however, we bracket the metaphysics of patterns and instead focus on the role that skills play for pattern recognition in scientific practice.

cases of failure, Haugeland usefully introduces two kinds of recognition: inner and outer. *Outer recognition* is "telling whether something (a pattern) *is there*" (Haugeland 1998, p. 285). A mason's recognizing a bond type in a wall is a case in point. By contrast, *inner recognition* amounts to "recognizing whether what is present, the current element, fits the pattern" (ibid.). For our mason that means to recognize whether a brick's position does or does not fit the bond type of the wall. Inner recognition can fail while outer recognition succeeds: a mason may misrecognize brick positions that implement a "Sussex bond" as fitting the pattern of a "Monk bond". Conversely, outer recognition can fail while inner recognition succeeds: a mason may see a wallpaper from a distance and recognize a "Sussex bond" while there is not actually a brick wall behind the paper. These different kinds of failures illustrate that inner and outer recognition can conflict with one another. Resolving these conflicts is an important part of the pattern recognition process, and consequently, of anomaly resolution in mechanism discovery (as we will detail in Sect. 4.2).

In scientific discovery, pattern recognition always happens within an epistemic collective of which an individual scientist is a part. Members in the collective share an *epistemic perspective* (cf. Kästner 2018). This expresses what Dennett calls "taking a stance" and Haugeland describes as "applying certain rules of recognition" when recognizing patterns. An epistemic perspective not only allows individual scientists to recognize specific patterns, but also provides normative standards to evaluate recognition performances. For example: from the perspective of electrophysiology, a researcher can recognize the overall pattern of the action potential by analyzing spike trains in her electrode recordings. For the analysis to be correct, she should follow the current standards of analysis set by the electrophysiology community. A given perspective, that is, provides both skills for recognition and normative constraints on the execution of these skills. If we shift our perspective (take a different stance, apply different rules) and apply a separate set of recognition skills, different elements of a pattern may be revealed. From the perspective of molecular neurobiology, for instance, we can recognize distinct phases in the overall pattern of the action potential, and investigate which molecules get released during each phase. Which perspective an agent will choose largely depends on the tools available to her, her knowledge of the phenomenon and research question, as well as which type of experiment is possible given limited time and resources. Still, successful pattern recognition will also depend on the behavior of the entities that make up the orderly arrangement being investigated (e.g., the molecules and their charges in the action potential example above).

Thus far, we framed mechanistic inquiry as pattern recognition. In our view, scientists adopt epistemic perspectives to constrain their search for orderly arrangements that are stably recognizable. The subsequent discovery of the elements of such patterns involves recognizing that the pattern is there (outer recognition) as well as recognizing how a given element fits that pattern (inner recognition). In the next section, we show that mechanistic inquiry naturally lends itself to be analyzed as such a pattern recognition process. Because this framework is not biased towards either epistemic or ontic norms it allows us to explicate how epistemic and ontic norms work together in mechanism discovery.

3.2 How Ontic and Epistemic Norms Guide Mechanistic Inquiry: The Pattern Account

We conceive of mechanistic inquiry as a process during which scientists apply recognition skills to search for patterns in the world. As mentioned above, the two senses of "pattern" nicely map onto central features of mechanistic discovery accounts. Patterns as orderly arrangements persisting from below correspond to those portions of the causal structure of the world that produces, maintains or underlies the phenomenon to be explained (Craver and Darden 2013). Patterns as candidates for recognition that are salient from above correspond to the phenomena that scientists seek to explain. They isolate orderly arrangements which exhibit *discernable regularities*—recognized as phenomena to be explained—from causal background factors and experimental noise. This isolation process usually invokes heuristic assumptions. An example is the assumption that causal interactions within a mechanism are denser/stronger than the interactions between the mechanism and its surroundings (Bechtel and Richardson 2010). Following isolation researchers start searching for the pattern's elements, which requires an interplay between inner and outer recognition. Discovering parts requires researchers to inner recognize how those parts fit into the pattern as a whole (i.e. how they contribute to the mechanism's behavior). But this presupposes that researchers can outer recognize that a pattern exhibits a discernable regularity (a phenomenon) and is not just an isolated curiosity (cf. Rouse 2015a, p. 235). The interplay between inner and outer recognition is present in different discovery strategies: mutual manipulation experiments (Craver 2007, ch. 4), establishing how a component studied *in vitro* contributes to the behavior of a mechanism (Rheinberger 1997; Craver and Darden 2013, ch. 7), or uncovering the organization of the pattern through visualization techniques (Kästner 2015).

In experimental practice, the joint application of inner and outer recognition skills generates data. Acquiring and analyzing those data requires instruments, basic methodological skills to properly use those instruments, and concepts/models to adequately describe the data. Together, instruments, scientists with their methodological skills and concepts/models form what we call a *pattern recognition practice*. In mechanism discovery, these elements of a pattern recognition practice have a common reference point: a mechanism in the world.[10] We thus highlight both the importance of ontic structures and epistemic practices in mechanism discovery. Our account also provides a clear criterion to individuate practices: one pattern recognition practice is distinguished from another one by the mechanism it investigates.[11] Pattern recognition practices do not merely refer to a mechanism "out there"; the

---

[10] What we mean here is not that scientists involved in the pattern recognition practice already agree on the details of the mechanism which will be the product of the discovery process. Rather, we suggest they have a shared agenda to explain a specific phenomenon by identifying the entities and activities responsible for it. This does still allow for disagreement as to which entities and activities are involved (see also below).

[11] As we will detail below, because there can be considerable uncertainty regarding the mechanism during discovery, clearly individuating the corresponding pattern recognition practice is often only possible in retrospect.

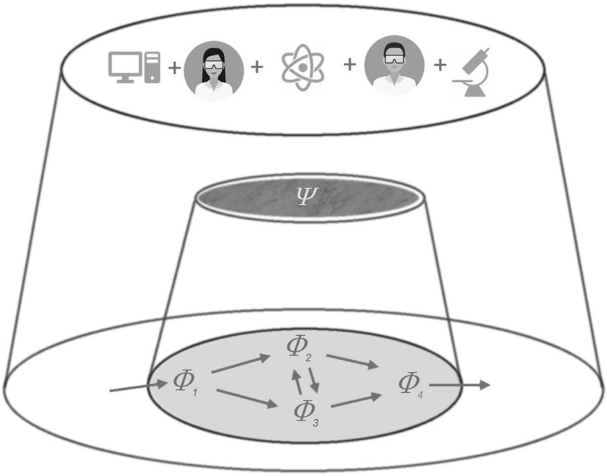

Fig. 1 Ontic mechanisms as parts of pattern recognition practices

mechanism is actually *part* of the recognition practice (Fig. 1). Because the elements of a recognition practice repeatedly occur together, the practice forms itself a pattern in the world, i.e. an orderly arrangement of instruments, scientists/skills, and scientific concepts/models. Rouse (2015b) presents a similar idea: "The real patterns that comprise the functioning of 'ontic' mechanisms in the world are *components* of a larger pattern that also incorporates scientific capacities for pattern-recognition and articulation." (ibid., p. 14). The idea that mechanisms (as causal structures) are parts of their recognition practices squares well with mechanistic accounts of discovery, insofar as they emphasize that practitioners work on carving out the mechanism, e.g. by manipulating it to reveal its causal structure. This is not to impose strong metaphysical assumptions. As we will discuss below, on the pattern account, the overall organization of the mechanisms carved out is always practice-relative.

Figure 1 visualizes the idea that mechanisms are part of the pattern recognition practices investigating them. The inner figure shows the schema of constitutive mechanisms adapted from Craver (2007). ψ is the phenomenon constituted by the entities ($\phi_{1–4}$) that causally interact with one another via activities (arrows). The outer figure schematically depicts a pattern recognition practice. It consists of skilled scientists (avatars) who use instruments (computer on the left, microscope on the right) and concepts/models (middle) to describe, manipulate or predict the behavior of a mechanism according to the methodological standards that regulate inquiry in their domain of inquiry. The mechanism's causal structure is part of this practice because the scientists' activities are directed towards it. When electrophysiologists record from an ion channel within a neuron, for example, their measurements are directed towards the spatiotemporal organization of the action potential.

By incorporating a full mechanism as part of a pattern recognition practice, Fig. 1 captures the outcome of a successful mechanism discovery episode. During discovery, however, the mechanism under investigation is at least partially unknown to the researchers. So how do researchers begin the discovery process, i.e. how do they identify the mechanism around which their recognition practice is centered? And how do the elements of their practice develop as they discover more relevant details about the mechanism?

Following mechanistic discovery accounts, we assume that a discovery episode typically starts with characterizing a phenomenon, i.e. distinguishing a pattern from its background (Bechtel and Richardson 2010, ch. 3; Craver and Darden 2013, ch. 4). Skilled scientists can achieve such a characterization via different *epistemic activities,* i.e. actions they undertake to produce or improve knowledge about a pattern (cf. Chang 2014, p. 72).[12] Examples of epistemic activities include modeling how a given input is transformed into a certain output (think of box-and-arrow diagrams in cognitive psychology), or measuring the behavior of the phenomenon experimentally with the help of operational definitions (Feest 2011). Such epistemic activities do not occur in isolation. Rather, they form a coherent set that constitutes a pattern recognition practice.[13] Each epistemic activity involves all three elements of a pattern recognition practice (skills, tools and concepts/models). Yet, various epistemic activities emphasize them differently (e.g., experiments lean more heavily on instruments than theoretical modeling).

It is not enough for a pattern recognition practice to simply characterize a phenomenon, i.e. what is salient from above. Mechanistic explanation also requires researchers to specify the elements of a pattern, i.e. what makes the pattern persist from below. To achieve this goal, researchers must introduce various *epistemic operations* that track the entities and activities constituting the pattern. The selection of such epistemic operations is ontically constrained: scientists must tailor them to the particular spatiotemporal characteristics of the entities and activities they are supposed to track. Collections of epistemic operations jointly constitute an epistemic activity. For instance, experimental operations on ion channels, membrane potentials or neurotransmitter release etc. jointly constitute the epistemic activity of empirically investigating the action potential. Operations belonging to different epistemic activities can track the same entity or activity. For instance, (1) injecting a current in a neuron and measuring the response of an ion channel with a patch clamp, and (2) mathematically modeling the response of that ion channel to voltage changes belong to separate epistemic activities. Yet they track the same entity in the mechanism of the action potential.

Figure 2 visualizes our conception of a pattern recognition practice. It consists of two different epistemic activities [modeling (left), experimenting (right)] in which scientists perform different epistemic operations to track various entities and activities in the mechanism they investigate. To integrate the findings obtained by these operations, researchers must *coordinate* their epistemic activities. Coordination involves both ontic and epistemic norms. Different epistemic operations may share

---

[12] Epistemic activities are typically governed by rules (e.g. standards of a discipline), but as Chang points out, these rules need not be articulated.

[13] Individual epistemic activities are carried out from a particular epistemic perspective (see Sect. 3.1). Yet, a pattern recognition practice may combine epistemic activities that adopt different epistemic perspectives.

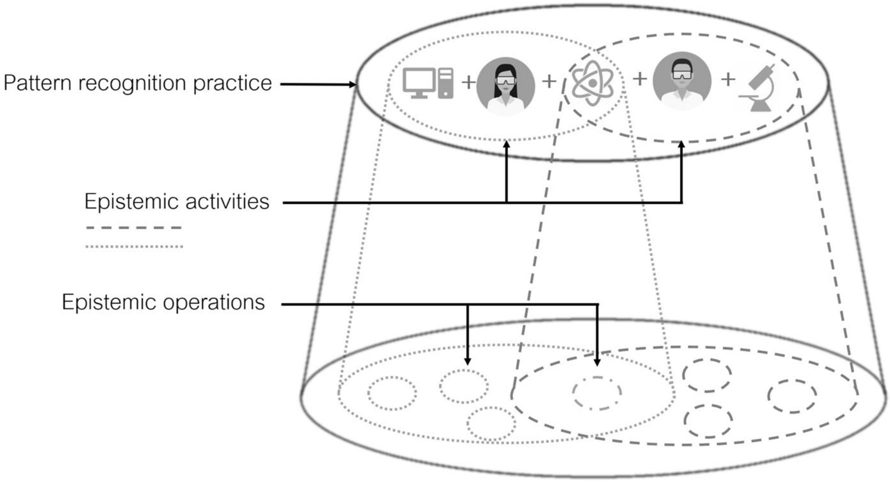

Fig. 2 The structure of a pattern recognition practice

the same ontic referent (epistemic operation in the middle of Fig. 2). An ontic constraint in this case is that findings obtained by the two epistemic operations cannot remain mutually incompatible (cf. Haugeland 1998, p. 335). An epistemic constraint is that comparing findings obtained through the operations may require researchers to adopt a shared conceptual framework (indicated by the atom icon above).

Note that our talk about "shared ontic referents" only requires a metaphysically undemanding form of realism about entities and activities. There are two realist positions which are congenial with mechanist accounts of discovery. One is Hacking's (1983) entity realism, which asserts that if scientists can manipulate an entity and use it to manipulate other entities then they have grounds to assume that this entity actually exists. The other is Wimsatt's robustness analysis (1981) which asserts that if an entity or activity can be reliably observed or measured by independent means, then there is evidence for its reality. Both these positions go hand in hand with mechanists' emphasis on experimental manipulations and multiple measurement techniques as a means to identify entities and activities contributing to a phenomenon (e.g., Craver 2007, ch. 7; Craver and Darden 2013, ch. 9).

Realism about entities and activities in this sense does not stand in opposition to the *practice-relativity* of mechanisms inherent in the pattern account: a mechanism as a whole is always bound to a specific pattern recognition practice, which can contain multiple epistemic activities each containing different models, skills, and instruments. Within a single pattern recognition practice, different epistemic activities can robustly track the same entities and/or activities. But that does not imply that all practitioners share the same *beliefs* about a given entity or activity (e.g. with respect to its role for the operation of a mechanism), neither does it establish that the detected entities and activities are in fact parts of the pattern to be discovered. To figure out the boundaries of the mechanism, and how its components are organized together, researchers have to further investigate how the entities and activities they identified work together to bring about the phenomenon that is the candidate for recognition.

While it is acknowledged in the literature that drawing boundaries and uncovering organization are important to achieve integrated mechanistic explanations (cf. Craver 2007, ch. 7; Bechtel 2008), it is not usually discussed how such integration actually proceeds. We think that one motor for integration is resolving anomalies and *disagreements* about the boundaries, components and organization of the mechanism. Because the pattern account leaves room for such disagreement, it can specify how the integration of different epistemic activities works within a pattern recognition practice (see Sect. 4.2 for an example). Despite the presence of disagreement, a pattern recognition practice remains coordinated as long different researchers share the commitment to investigate the entities and activities responsible for the phenomenon (i.e. candidate for recognition) in question (cf. Haugeland 1998, ch. 13). This coordination involves sharing both ontic and epistemic norms within a given pattern recognition practice. Researchers must mutually acknowledge that (1) the different epistemic activities they perform produce knowledge about the pattern (epistemic norm), and (2) the success or failure of their epistemic activities depends on the causal structure of that pattern (ontic norm). Note that, at least in the context of mechanism discovery, we cannot have one norm without the other. Without ontic norms, it is—impossible evaluate the success or failure of the various *epistemic* activities such as modeling or experimentation.[14] And without epistemic norms, it is impossible to identify and communicate about *ontic* structures, and to make intelligible how they contribute to the phenomenon (Sect. 2).

Our pattern account thus clarifies the relation between epistemic practices and ontic structures (Fig. 1) and the mutual reliance of ontic and epistemic norms (Fig. 2). Moreover, it usefully explicates how researchers proceed from establishing a pattern recognition practice to successfully discovering the mechanism responsible for the phenomenon of interest (Fig. 3a–c). The starting point is usually a characterization of this phenomenon through one or more epistemic activities (Fig. 3a).

Over time, researchers introduce more and more epistemic operations tracking different entities and activities that potentially contribute to the phenomenon (Fig. 3b). At this stage, boundaries of a mechanism may remain unclear. Researchers may have developed epistemic operations that track entities and activities without knowing whether these actually belong to the mechanism (ϕ with question marks in Fig. 3b). They may also be uncertain about what their epistemic operations actually track or mistakenly assume that there is anything to track at all (question marks in Fig. 3b). Researchers may only partially know the temporal order of activities and interactions between entities (e.g. in Fig. 3b they know the activity

---

[14] We agree with Halina (2018) that the norms of accuracy and intelligibility can sometimes pull in different directions: "intelligibility may take priority in pedagogical contexts; while conveying information about the target mechanisms may become more important in those contexts where advanced researchers are attempting to understand and intervene on a target system." (ibid., p. 221, see also Kaplan and Craver 2011, 609f. for a similar point). However, here we are only concerned with the latter contexts, i.e. contexts in which researchers perform epistemic activities to generate *novel* knowledge about entities and activities. Thus our point about the need for both ontic and epistemic norms still holds for these contexts.

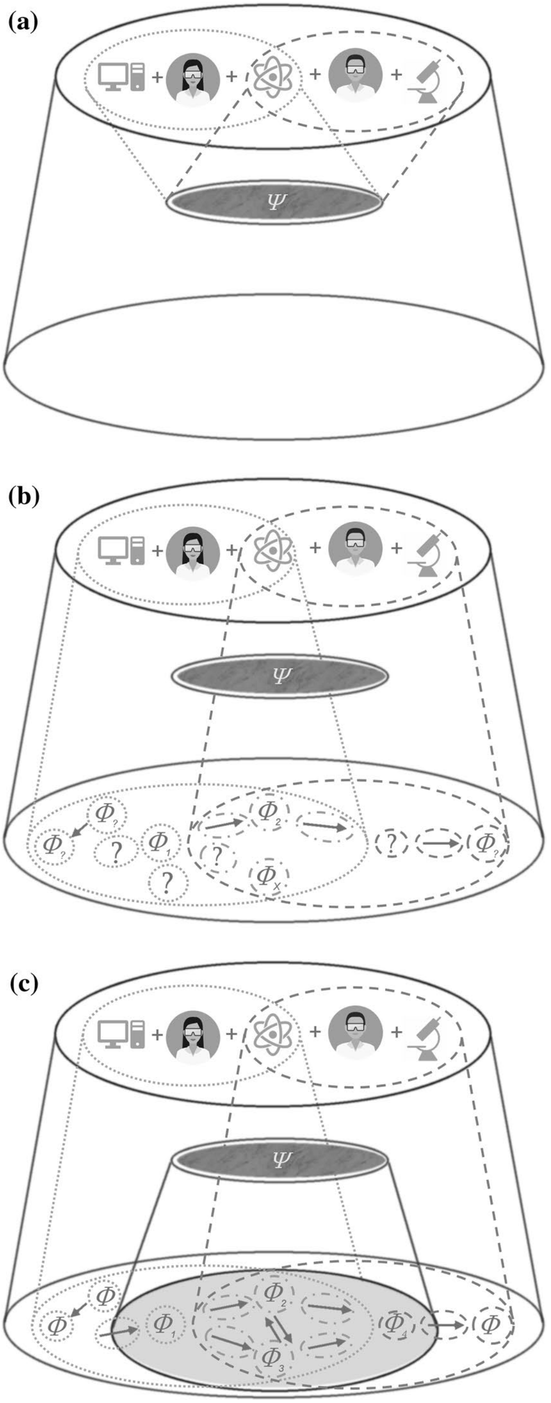

Fig. 3 **a** Establishing a pattern recognition practice by characterizing the phenomenon. **b** Tracking entities and activities with epistemic operations. **c** Successful discovery: pattern recognition practice as a whole has the full mechanism in view

linking ϕ_1 and ϕ_2, but not that linking ϕ_2 and ϕ_x). Depending on how these issues are resolved, researchers take different trajectories through the search space of possible mechanisms.

If successful, the pattern recognition practice as a whole will eventually have in view the mechanism responsible for the phenomenon from which the discovery episode started. That is, each entity and activity in the mechanism is being accurately tracked by at least one epistemic operation (Fig. 3c). Which exact mechanism researchers end up with depends on the trajectory they take through search space. Trajectories are characterized by choice points such as where to draw the boundaries of the mechanism; they are, like the mechanism as a whole, practice-relative. The upshot of combining realism about entities and activities with the practice-relativity of mechanisms as a whole is that the possible trajectories through the search space must be both ontically and epistemically constrained to lead to the discovery of an explanatory mechanism. It is ontically constrained because researchers must choose which epistemic operations accurately track entities and activities. And it is epistemically constrained because researchers aim to arrive at a coherent understanding of the mechanism, which requires e.g. resolving anomalous findings and reducing errors throughout the discovery process (see also Sect. 4.2).

Figures 3a–c of course present an idealized picture of how pattern recognition practices can successfully discover mechanisms. We emphasize that this process is often not as straightforward as the above description suggests. Researchers typically have to recognize patterns in the face of noise in experimental systems, assess causal roles based on imperfect manipulations, and build mathematical models by approximating (incomplete) empirical results. Oftentimes, phenomena that are mechanistically decomposed later are initially investigated in different contexts of inquiry (see the discussion of long-term-potentiation in epilepsy research in Craver 2003, pp. 167–170). Similarly, new entities are often unexpectedly discovered under experimental circumstances that were initially geared towards a different aspect of a mechanism (see the discussion of discovering transfer RNA in protein synthesis research in Rheinberger 1997, p. 154f., p. 189). Because of these practical complications, actual discovery episodes do not typically proceed straight from characterizing the phenomenon to discovering all relevant aspects of a mechanism. They often take unexpected turns, encounter dead ends, or deviate towards other phenomena.

Despite such practical complications, it is important to recognize that error reduction and anomaly resolution are an integral part of scientific discovery. And they are only possible if the reference to a pattern as orderly arrangement remains constant over time. In ongoing research, it may be unclear what exactly the overall mechanism is as it is not yet in full view. Also, the description of both the mechanism and the phenomenon may change over time. Uncertainty and dynamic re-descriptions are thus integral parts of investigating the "same mechanism" (Bechtel and Richardson 2010; Feest 2011; Craver and Darden 2013; Rouse 2015b). But if a pattern recognition practice fails to refer to an orderly arrangement altogether, then its epistemic activities will cease to contribute to the mechanistic discovery process.[15] In order to avoid such failure, researchers need to continuously try to align their epistemic activities of pattern recognition with the gradually emerging boundaries of the mechanism(s) they are investigating.

In sum, we have couched mechanism discovery in terms of a pattern recognition practice centered around an orderly arrangement (a mechanism situated in the causal structure of the world). A pattern recognition practice consists in various epistemic activities whose epistemic operations track different entities and activities in the mechanism (the elements of the pattern). By describing how ontic and epistemic norms guide each step throughout the discovery process, we have made explicit what remained implicit in previous mechanistic discovery accounts: *how* ontic and epistemic norms work together. With the pattern account in place, we can now address the methodological questions left open by existing discovery accounts: how ontic constraints contribute to the evaluation of discovery heuristics (Sect. 4.1) and how epistemic constraints contribute to the resolution of anomalies (Sect. 4.2).

4 Benefits of the Pattern Account of Mechanistic Inquiry

In this section, we show why it is beneficial to conceive of mechanism discovery as a process of pattern recognition. To do this, we return to the two challenges for existing discovery accounts outlined in Sect. 2: Bechtel and Richardson (2010) leave unanswered the question whether ontic bottom-up constraints play a direct or indirect role in the heuristically guided search for mechanisms. Craver and Darden (2013) do not conclusively explain how researchers can identify the ontic source which gave rise to the anomaly, and how exactly discovery should proceed in the face of certain kinds of anomalies. We now show that the pattern account can answer both these questions.

4.1 Bechtel and Richardson: How Do Ontic Constraints Guide Mechanism Discovery from the Bottom Up?

Let us first consider Bechtel and Richardson's (2010) account. Though they do acknowledge that ontic accuracy may serve as a bottom-up constraint on discovery, we argued in Sect. 2 that their exposition of bottom-up constraints remains ambiguous. Such constraints can function as *direct* (i.e. theory-independent) constraints if they are *objective ontic facts*, or they can function as *indirect* (i.e. theory-dependent) constraints if they are *epistemically interpreted findings*. We now show that whether ontic constraints play such direct or indirect roles depends on the exact kind of problem practitioners are facing.

Ontic constraints play a *direct* role when researchers aim to evaluate whether a given epistemic *operation* is appropriate to characterize a given element of the

---

[15] For examples of such failed systems in genetics and molecular biology see Rouse (2015a, p. 312) and Rheinberger (1997, p. 50, p. 196).

pattern (an entity or activity in the mechanism). This follows from our realism about entities and activities, which implies a non-inferential notion of observation. Thus, on the pattern account, successful epistemic *operations* (e.g. seeing ion channels with a microscope) provide researchers with direct access to facts about the entities and activities they track (see Sect. 2). This direct access is granted because which epistemic operations are *appropriate* depends on the spatio-temporal characteristics of the entities and activities they are supposed to track (Sect. 3.2). For example: the duration of the action potential directly constrains which temporal resolution is appropriate to measure it. That is, researchers must match the precision of their instruments to the *scale* of the entities or activities they are investigating (Potochnik and McGill 2012; Haueis 2014, 2018. ch. 2).

By contrast, ontic constraints play an *indirect* role when researchers aim to put a variety of facts into a coherent representation (e.g. when building a mathematical model of the action potential). This maps onto the practice-relativity of mechanisms as a whole on the pattern account. Because a given set of facts is compatible with different models, researchers need to epistemically constrain which of these models to adopt. Ontic facts play only an indirect role because they need to be epistemically interpreted to constrain model selection; that is, they need to be made intelligible through a conceptual framework and are thus theory-dependent. Direct constraints on the epistemic activity of model selection are primarily epistemic. For example, scientists can favor a model which can be used to derive testable hypotheses at the current stage of the discovery process. Such a selection marks a choice point along the trajectory the discovery process takes through search space. It puts aside other empirically equivalent models that contain hypotheses that are currently untestable, e.g., because of technological, practical or ethical limitations. Another epistemic constraint is to favor those models which are mathematically tractable. Following this constraint puts aside other models which include more empirical details but which cannot be solved by the computational means currently available to the scientists. When choosing a computationally tractable model, constraints which determine the norm of ontic accuracy (e.g. the 3M constraint) only play an *indirect* role in discovery (cf. Boone and Piccinini 2016, p. 5).

Because Bechtel and Richardson only implicitly rely on ontic bottom-up constraints they do not distinguish between cases in which ontic constraints play direct and indirect roles, respectively. The pattern account can do so because it distinguishes between (1) epistemic operations that are constrained by the spatiotemporal properties of entities and activities, and (2) the epistemic activity of model selection that is constrained directly by epistemic criteria (overall intelligibility, computational tractability, generating testable hypotheses, etc.) and only indirectly constrained by ontic bottom-up constraints.

Our response to the challenge that Bechtel and Richardson's account cannot meet is thus rooted in how the pattern account relates practice-relativity and realism about entities and activities. It highlights that without ontic norms, epistemic constraints are normatively inert. To normatively evaluate an explanation (assess whether it is good or bad) we must consult ontic constraints that—directly or indirectly—constrain our search for mechanisms; and the pattern account elaborates how this works.

4.2 Craver and Darden: How Do Epistemic Constraints Help with Anomaly Resolution?

Let us now turn to Craver and Darden's (2013) account. Though they emphasize that epistemic constraints contribute to the resolution of experimental errors, i.e. (a)-type anomalies, they do not explicate how researchers epistemically constrain their choice between the ontic norm of accuracy, i.e. (b)-type anomalies, and the ontic norm of completeness—i.e. (c)-type anomalies (see Sect. 2). With the pattern account in place, we can now approach this question.

Let us consider another case of a *temporal anomaly* discussed by Craver and Darden: the unexpectedly short measurement of a protein synthesis initiation in the so-called PaJaMo experiment. Pardee et al. (1959) observed that that protein synthesis starts quickly after a functional gene is inserted in *E. coli* bacteria which lack that gene. This finding presented an anomaly to biochemists and molecular biologists, who operated from different epistemic perspectives when researching protein synthesis. Both groups, however, assumed that DNA synthesis happens in the ribosome, and which works at slower rates than observed in the PaJaMo experiment (cf. Darden and Craver 2002, p. 15). The crucial question is whether this short duration of the measurement arises from (a) flaws in the measurement device or experimental design (b) deviations between the actual rate of ribosomal DNA synthesis and the rate postulated by the model, or (c) accidently measuring a previously unknown type of RNA (so-called *messenger* RNA, or mRNA).

Craver and Darden rule out that the PaJaMo experiment presents an (a)-type anomaly because the same results were replicated in subsequent experiments (Craver and Darden 2013, p. 155). Since in (a)-type anomalies the error lies in the experiment and not the mechanistic model, the resolution of such anomalies does not require researchers to consult ontic norms of accuracy and completeness. With regard to (a)-type anomalies, the pattern account agrees with Craver and Darden's account: it assumes that researchers can identify (a)-type anomalies by checking whether *basic methodological skills* (e.g. how to measure an action potential, how to run a genetic cross-breeding experiment etc.) were properly applied. If they were not, the anomalous finding should be counted as an incorrect result; it stems from problems with data acquisition and analysis. If, by contrast, the current epistemic norms of the pattern recognition practice have been followed, the anomalous finding should be counted as a correct result; it does not stem from methodological problems. Hence, scientists should consult ontic norms for anomaly resolution. But having ruled out experimental errors, should researchers (b) follow the ontic norm of accuracy and adjust the duration of ribosomal DNA synthesis in the current model or should they (c) follow the ontic norms of completeness and posit a new *type* of RNA which could achieve faster synthesis rates? And how should they epistemically constrain this choice?

Biochemists and molecular biologists actually disagreed on how to answer these questions. Some prominent biochemists (e.g., Sol Spiegelman) wanted to (b) adjust the established model in which ribosomal RNA acted as the template for DNA synthesis, while some prominent molecular biologists (e.g. Jim Watson and Sydney Brenner) wanted to (c) propose a new model which includes mRNA as novel

RNA type and in which ribosomal RNA acts not as a template but as an unspecific "reading head" (cf. Judson 1996, p. 417, 422f.). Craver and Darden acknowledge that "interfield competition between biochemists and molecular biologists" existed (2002, p. 16). However, their discovery account does not capture that protein synthesis researchers actually disagreed about *which* ontic norm to follow when resolving this temporal anomaly. Because the pattern account makes the disagreement about which ontic norm to follow a crucial ingredient of anomaly resolution, it presents an improvement over Craver and Darden's account.

The pattern account draws on inner and outer recognition skills (Sect. 3.1) as specific epistemic constraints that researchers use to decide whether to follow the ontic norms of accuracy or completeness, respectively. Outer recognition tracks whether we are looking at the same pattern across different experiments. Inner recognition tracks how an obtained measurement of a purported element fits that pattern. The crucial feature here is that both are *monitoring skills* (cf. Haugeland 1998, p. 335), which are importantly different from basic methodological skills: rather than sorting correct experimental results from incorrect ones, monitoring skills help scientists determine what a given result is evidence *for*. Thus, inner and outer recognition skills help researchers to epistemically constrain which ontic norm to consult when resolving the anomaly at hand.

Let us now use inner and outer recognition skills to describe how researchers resolve the temporal anomalies like the one of the PaJaMo experiment. If they inner recognize a shorter duration for the synthesis of new ribosomal RNA while stably outer recognizing the overall pattern of protein synthesis, researchers face a (b)-type anomaly. Thus, they need to follow the ontic norm of accuracy and adjust their assumptions about the duration of ribosomal DNA synthesis. For example, Noumura et al. (1960) conducted a centrifuge experiment indicating that after bacteriophage infection, *E. coli* bacteria synthesized novel ribosomal RNA faster than previously observed. This result led them to follow the ontic norm of accuracy: they suggested to adjust the synthesis rate of ribosomal RNA in the established model. By contrast, if researchers can inner recognize different *kinds* of RNA, while outer recognizing how these different kinds contribute to the overall pattern of protein synthesis, researchers face a (c)-type anomaly. Thus, they need to follow the ontic norm of completeness and posit a new type of RNA. For example, Brenner et al. (1961) were able to inner recognize existing ribosomal RNA and newly synthesized mRNA, which led them to follow the ontic norm of completeness: they regarded the mRNA as an additional element in the pattern. They also disagreed that Noumura et al. (1960) successfully inner recognized enough newly synthesized ribosomal RNA to account for the temporal anomaly by revising the established model (cf. Judson 1996, p. 422).[16]

---

[16] While this peer disagreement continued for several years, the pattern recognition practice eventually converged a shared conceptual framework ("genetic code", "information transfer" etc.) as well as experimental systems (e.g., *E. coli* in-vitro system) to investigate protein synthesis (cf. Rheinberger 1997, ch. 12). The now shared epistemic perspective is evident Watson's (1965) textbook, which presents a new model of protein synthesis including mRNA, together with biochemical and molecular biological details about the mechanism (cf. Darden and Craver 2002, p. 17).

In sum, our treatment of anomaly resolution highlights that without epistemic norms, ontic constraints are dormant. To decide which kind of ontic constraint a given experimental finding imposes, we must carefully exercise and assess appropriate methodological along with monitoring skills. Inner and outer recognition act as epistemic constraints on the choice of which ontic norm to follow when resolving anomalies. The pattern account thus demonstrates not only *that*, but also *how* ontic and epistemic norms work together in anomaly resolution. While the case we discussed here demonstrates only one way in which ontic and epistemic norms interact, we believe there are certainly others, too. But that is a matter for future research. Still, we have illustrated that the pattern account can address the challenges that existing discovery accounts cannot answer satisfactorily.

5 Conclusions

We argued that while existing accounts of discovery emphasize ontic or epistemic norms respectively, they implicitly recruit the other set of norms (Sect. 2). Bechtel and Richardson (2010) have to rely on ontic norms to evaluate when it is normatively appropriate to apply a certain discovery heuristic to a particular system. Without ontic norms, epistemic constraints are normatively inert. Conversely, Craver and Darden (2013) have to rely on epistemic norms to outline how scientists decide which kind of anomaly they are facing in the discovery process. Without epistemic norms, ontic constraints are dormant. Both resolving anomalies and evaluating heuristics is crucial to determine which trajectory scientists should take trough search space to discover mechanisms. A satisfying account of mechanism discovery should therefore make both ontic and epistemic constraints on this search process explicit. To achieve this, we introduced the pattern account (Sect. 3). It shows that mechanistic inquiry, as a process of pattern recognition, is both ontically and epistemically constrained. Scientists perform epistemic activities to make an orderly arrangement in the world a candidate for recognition. To investigate the elements of this arrangement they need to employ specific epistemic operations tailored to the spatio-temporal properties of the entities and activities being tracked. Eventually, a successful pattern recognition practice will have in view the overall mechanism responsible for the phenomenon characterized at the outset of the discovery episode. To reach this goal, practitioners must share the commitment that (1) they search for the entities and activities responsible for the phenomenon in question *and* (2) that their epistemic activities make intelligible how the mechanism is responsible for the phenomenon to be explained.

While we purposefully bracketed metaphysical issues from our current discussion, we would like to close by bringing the following issue to the reader's attention: Claiming that ontic and epistemic norms are equally important in mechanism discovery, and that the same set of norms applies to both mechanistic discovery and explanation, seems inconsistent with prioritizing either ontic or epistemic norms in mechanistic explanations. Given that the norms of mechanistic explanations are immediately tied to the metaphysical question of what explanations are, we wonder what the implications of the pattern account for the so-called ontic-epistemic debate about the nature of explanation might be. Although we did not develop such an account explicitly, we think that the pattern account of mechanism discovery provides an important step in this direction. It suggests that neither ontic nor epistemic norms are primary when sorting good mechanistic explanations from bad ones (see Illari 2013 for a similar conclusion). As long as we do not want to resurrect the logical positivist distinction between discovery and justification, we should embrace the conclusion that ontic and epistemic norms jointly constrain the entire process of mechanistic inquiry.

Acknowledgements We are indebted to Carl Craver, Uljana Feest, Joseph Rouse, and Beate Krickel for ample discussion on the subject and helpful feedback on earlier versions of this paper. We would also like to thank the participants of the Workshop "Patterns in Science" held at Berlin School of Mind and Brain in 2015, the SPSP conference 2016 at Rowan University, the GAP conference 2018 at University of Cologne, and discussants at the philosophy of science colloquium at Australian National University in 2019 as well as three anonymous reviewers.

References

Bechtel, W. (2008). *Mental mechanisms: Philosophical perspectives on cognitive neuroscience*. London: Routledge.

Bechtel, W., & Richardson, R. (2010). *Discovering complexity. Decomposition and localization as strategies in scientific research*. Cambridge, MA: MIT Press.

Boone, T. W., & Piccinini, G. (2016). Mechanistic abstraction. *Philosophy of Science, 83,* 1–13.

Brenner, S., Jacob, F., & Meselson, M. (1961). An unstable intermediate carrying information from genes to ribosomes for protein synthesis. *Nature, 190,* 576–581.

Bursten, J. (2018). Smaller than a breadbox. Scale and natural kinds. *British Journal of Philosophy of Science, 9*(1), 1–23.

Chang, H. (2014). Epistemic activities and systems of practice: Units of analysis in philosophy of science after the practice turn. In L. Soler, S. Zwart, M. Lynch, & V. Isreal-Jost (Eds.), *Science after the practice turn in philosophy, history and social studies of science* (pp. 67–79). London: Routledge.

Craver, C. (2003). The making of a memory mechanism. *Journal of the History of Biology, 36,* 153–195.

Craver, C. (2007). *Explaining the brain*. Oxford: Oxford University Press.

Craver, C. (2014). The ontic account of scientific explanation. In M. Kaiser, O. Scholz, D. Plenge, & A. Hüttemann (Eds.), *Explanation in the special sciences: The case of biology and history* (pp. 27–52). Dordrecht: Springer.

Craver, C., & Darden, L. (2013). *In search of mechanisms: discoveries across the life sciences*. Chicago: University of Chicago Press.

Craver, C., & Kaplan, D. (2018). Are more details better? On the norms of completeness for mechanistic explanations. *The British Journal for the Philosophy of Science*. https://doi.org/10.1093/bjps/axy015.

Darden, L., & Craver, C. (2002). Strategies in the interfield discovery of the mechanism of protein synthesis. *Studies in the History and Philosophy of the Biological and Biomedical Sciences, 33,* 1–28.

Dennett, D. (1991). Real patterns. *Journal of Philosophy, 88*(1), 27–51.

Duhem, P. (1906/1954). *The aim and structure of physical theory*, trans. from *La Théorie Physique: Son Objet et sa Structure* (Paris: Marcel Riviera & Cie.). Princeton, NJ: Princeton University Press.

Feest, U. (2011). Re-membering (short term) memory: Oscillations of an epistemic thing. *Erkenntnis, 75*(3), 391–411.

Feest, U. (2017). Phenomena and objects of research in the cognitive and behavioral sciences. *Philosophy of Science, 84,* 1165–1176.

Hacking, I. (1981). Do we see through a microscope? *Pacific Philosophical Quarterly, 62*(4), 305–322.

Hacking, I. (1983). *Representing and intervening. Introductory topics in the philosophy of the natural sciences*. Cambridge: Cambridge University Press.

Halina, M. (2018). Mechanistic explanation and its limits. In S. Glennan & P. Illari (Eds.), *The Routledge handbook of mechanisms and mechanical philosophy* (pp. 213–224). London: Routledge.

Haueis, P. (2014). Meeting the brain on its own terms. *Frontiers in Human Neuroscience*. https://doi.org/10.3389/fnhum.2014.00815.

Haueis, P. (2018). *Meeting the brain on its own terms. Exploratory concept formation and noncognitive functions in neuroscience*. Dissertation, Otto-von-Guericke Universität Magdeburg.

Haugeland, J. (1998). Pattern and being. In *Having thought. Essays in the metaphysics of mind* (pp. 267–290). Cambridge, MA: Harvard University Press.

Haugeland, J. (2013). Dasein disclosed. In J. Rouse (Ed.), *John Haugeland's Heidegger*. Cambridge, MA: Harvard University Press.

Illari, P. (2013). Integrating the ontic and epistemic. *Erkenntnis, 78,* 237–255.

Judson, H. F. (1996). *The eighth day of creation: The makers of the revolution in biology (expanded edition)*. Cold Spring Harbour, NY: Cold Spring Harbor Laboratory Press.

Kaiser, M., & Krickel, B. (2017). The metaphysics of constitutive mechanistic phenomena. *British Journal of Philosophy of Science*, *68*(3), 745–779.

Kaplan, D. M., & Craver, C. F. (2011). The explanatory force of dynamical and mathematical models in neuroscience: A mechanistic perspective. *Philosophy of Science, 78*(4), 601–627.

Kästner, L. (2015). Learning about constitutive relations. In U. Mäki, I. Votsis, S. Ruphy, & G. Schurz (Eds.), *Recent developments in the philosophy of science: EPSA13 Helsinki. European studies in philosophy of science* (pp. 155–167).

Kästner, L. (2017). *Philosophy of cognitive neuroscience: Causal explanations, mechanisms & empirical manipulations*. Berlin: Ontos/DeGruyter.

Kästner, L. (2018). Integrating mechanistic explanations through epistemic perspectives. *Studies in the History and Philosophy of Science*, *68*, 68–79.

Ladyman, J., & Ross, D. (2007). *Every thing must go. Metaphysics naturalized*. Oxford: Oxford University Press.

Machamer, P., Craver, C., & Darden, L. (2000). Thinking about mechanisms. *Philosophy of Science*, *67*(1), 1–25.

Noumura, M., Hall, B. D., & Spiegelman, S. (1960). Characterization of RNA synthesized in Escherichia coli after bacteriophage T2 infection. *Journal of Molecular Biology, 2,* 306–326.

Pardee, A. B., Jacob, F., & Monod, J. (1959). The genetic control and cytoplasmic expression of 'inducibility' in the synthesis of beta-galatosidase. *Journal of Molecular Biology, 1,* 165–178.

Potochnik, A., & McGill, B. (2012). The limitations of hierarchical organization. *Philosophy of Science, 79,* 120–140.

Reichenbach, H. (1938). *Experience and prediction*. Chicago, IL: University of Chicago Press.

Rheinberger, H.-J. (1997). *Toward a history of epistemic things*. Stanford, CA: Stanford University Press.

Rouse, J. (2015a). *Articulating the world. Conceptual understanding and the scientific image*. Chicago, IL: University of Chicago Press.

Rouse, J. (2015b). Mechanisms as modal patterns. In *Manuscript presented at "patterns in science" workshop, Dec 3–4, 2015*. Berlin School of Mind and Brain.

Sheredos, B. (2016). Re-reconciling the epistemic and ontic views of explanation (or, why the ontic view cannot support norms of generality). *Erkenntnis, 81,* 919–949.

Van Fraasen, B. C. (1980). *The scientific image*. Oxford: Clarendon Press.

Wallace, D. (2003). Everett and structure. *Studies in History and Philosophy of Modern Physics, 34,* 87–105.

Watson, J. D. (1965). *Molecular biology of the gene*. New York: W. A. Benjamin.

Wimsatt, W. (1981). Robustness, reliability and overdetermination. In M. Brewer & B. Collins (Eds.), *Scientific inquiry in the social sciences (a festschrift for Donald T. Campbell)* (pp. 123–162). San Francisco: Jossey-Bass.

Wright, C. (2012). Mechanistic explanation without the ontic conception. *European Journal for Philosophy of Science*, *2*(3), 375–394.