

Georgia State University

ScholarWorks @ Georgia State University

Philosophy Theses

Department of Philosophy

8-13-2019

Experience-Based Intuitions

Tiffany Zhu

Follow this and additional works at: https://scholarworks.gsu.edu/philosophy_theses

Recommended Citation

Zhu, Tiffany, "Experience-Based Intuitions." Thesis, Georgia State University, 2019.
doi: <https://doi.org/10.57709/14939798>

This Thesis is brought to you for free and open access by the Department of Philosophy at ScholarWorks @ Georgia State University. It has been accepted for inclusion in Philosophy Theses by an authorized administrator of ScholarWorks @ Georgia State University. For more information, please contact scholarworks@gsu.edu.

EXPERIENCE-BASED INTUITIONS

by

TIFFANY ZHU

Under the Direction of Neil Van Leeuwen, PhD

ABSTRACT

In this thesis, I argue that many identification intuitions, such as one that helps you identify the authorship of a painting you are seeing for the first time, fall under the class of experience-based intuitions. Such identification intuitions cannot arise without intuition generating systems (IGSs) that are shaped by experiences accumulated during one's life. On my view, experience-based intuitions are produced by domain-general learning systems of hierarchical abstraction which may be modeled by deep convolutional neural networks. Owing to the mechanism of such IGSs, the reliability of experience-based intuition X depends on the quality of the experiences underlying the IGS which produces X. Lastly, I suggest that insofar as some philosophical thought experiments elicit experience-based identification intuitions, we can use the case method to glean information about our experiences as well as uncover certain conceptual commitments.

INDEX WORDS: Intuition, Philosophy of mind, Epistemology, Deep learning, Reliabilism, Perception, Philosophical methodology, Thought experiments

EXPERIENCE-BASED INTUITIONS

by

TIFFANY ZHU

A Thesis Submitted in Partial Fulfillment of the Requirements for the Degree of

Master of Arts

in the College of Arts and Sciences

Georgia State University

2019

Copyright by
Tiffany Ameline Zhu
2019

EXPERIENCE-BASED INTUITIONS

by

TIFFANY ZHU

Committee Chair: Neil Van Leeuwen

Committee: Eddy Nahmias

Electronic Version Approved:

Office of Graduate Studies

College of Arts and Sciences

Georgia State University

May 2019

ACKNOWLEDGEMENTS

I would like to thank Neil for his mentorship and for meeting with me on an almost weekly basis since summer of 2017. I would like to thank Eddy for teaching me, from scratch, to write a philosophy paper. Thank you both for not giving up on me. Dr. Andrea Scarantino has helped me improve my thesis and made me a more careful writer. Dr. Sandy Dwyer has not only read early drafts but also tirelessly supported me over the past years and helped me become a better teacher. Lastly, I thank my friends who help me understand and see the world. You continue to make me a kinder and happier person.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	V
LIST OF FIGURES	VII
1 INTRODUCTION.....	1
2 FEATURES OF INTUITION.....	6
3 EXPERIENCE-BASED INTUITIONS.....	9
4 EXPERIENCE-DEPENDENT INTUITION GENERATING SYSTEMS	13
4.1 Three characteristics of experience-dependent IGSs	13
4.2 Value of a computational model	14
4.3 How DCNNs explain experience-dependent IGSs	15
4.4 Five features of intuition revisited.....	18
5 RELIABILITY (AND NECESSITY) OF EXPERIENCE-BASED INTUITIONS OF IDENTIFICATION.....	20
6 EXPERIENCE-BASED INTUITIONS IN THOUGHT EXPERIMENTS	26
6.1 Do some thought experiments elicit experience-based identification intuitions?	26
6.2 What can experience-based thought experiment intuitions reveal?	31
7 CONCLUSIONS	34
REFERENCES.....	35

LIST OF FIGURES

Figure 1.1 The Mulberry Tree, Vincent van Gogh (c. October 1889).....	2
---	---

1 INTRODUCTION

It seems to me that existing philosophical theories cannot adequately account for a broad subset of intuitions. To get a sense of the sort of intuitions I have in mind, imagine the following scenario. You are visiting a new art museum for the first time. While you have some familiarity with canonical works of major artists, you are no art historian. Upon entering the first gallery of the museum, one painting catches your eye. The subject of the painting – a single tree – is portrayed with bold brushstrokes and vibrant blues and yellows. Although this painting is neither prominently displayed nor surrounded by museum visitors, you get a sudden sense that it is a van Gogh. But you can't quite articulate why. You've never seen this painting before, so you are not merely recalling its image from memory. You have seen reproductions of *Starry Night*, *Sunflowers* and a few of van Gogh's self-portraits, but those paintings look very different from the painting in front of you. Even more puzzlingly, many other landscapes hanging in the same gallery – the Impressionist wing of the museum – seem similar in style to the painting that caught your eye. When asked how you know the painting is a van Gogh, you might point out its bold colors or abstract manner of depiction. But a dozen other paintings in the gallery can be described in the same way. Nevertheless, your hunch turns out to be right! A label confirms that the painting is *The Mulberry Tree* (1889) by Vincent van Gogh.



Figure 1.1 The Mulberry Tree, Vincent van Gogh (c. October 1889)
Norton Simon Art Foundation

Your spontaneous and surprisingly accurate identification of the painting is prompted, I suggest, by an experience-based intuition.

A survey of literatures outside of philosophy reveals a consensus on people's reliance on similar sorts of identification intuitions across a panoply of domains of experience, including chess, sports, medicine, and military operations.¹ For instance, a seasoned trauma nurse making quick triage decisions relies not on referencing medical textbooks but on looking at patients. Yet existing philosophical accounts do not explain such identification intuitions very well. There seem to be two major camps of philosophical accounts of intuition. Rationalist accounts generally posit that intuitions are direct and *a priori* apprehension or insight (Bealer 1998,

¹ See Cokely and Feltz (2014) for an overview.

Bengson 2015, BonJour 1998). In contrast, a second type of account claims that intuitions largely result from systems which evolved to solve common problems in our evolutionary environment (Bargh 2011, Gigerenzer 2007, Klein 1998). Conceiving of intuitions as the outputs of evolved systems may be helpful for demonstrating how intuitions can be calibrated to a limited extent (Nagel 2012). However, as I discuss in more detail in section two, neither rationalist nor evolved systems accounts are particularly well-suited to explain the identification intuition in cases like the van Gogh example, which is *made possible*, not merely improved, by first-person experience. Moreover, thinking of intuitions as not rooted in experience has also effected a trend toward doubting their evidentiary role in judgment-making. Peter Railton is one of few contemporary philosophers optimistic about the reliability of intuitions, perhaps owing in part to his emphasis on the role of experience in shaping intuitions diachronically. For Railton, that intuitions are “direct, immediate, non-analytic” is just a user illusion – probabilistic computations are hard at work behind the scenes (2016: 38). But his account focuses on how representations of probability are updated through experience, not on how we come to have intuitions based on representations of concepts in the first place.² Thus, to better explain the sorts of identification intuitions that interest me, it seems that a radically empiricist account of intuition is in order.

The term intuition is used in many ways in philosophical literature. Here I will use the term *intuition* to refer to the conscious mental state that is a token of any given intuition type. For instance, one may experience many instances of moral intuitions, of possibility intuitions, or of art identification intuitions. *Intuition generating system* (IGS) refers to the cognitive structures that generate intuitions. *IGS-shaping experiences* are those experiences that shape an IGS by

² Railton has also tended to focus on normative intuitions, especially moral intuitions.

enriching it with content and refining its structure prior to the moment of intuiting. Lastly, *intuition-prompting experience* refers to any token experience that jumpstarts the process to engender a token intuition in real time. To clarify, there is a widespread consensus that intuitions are *prompted* by experience. What I take issue with is the extent to which experience accumulated prior to the moment of intuiting makes certain intuitions possible. In the art identification scenario, then, prior encounters with artworks are the IGS shaping experiences that have cultivated in you a relatively robust IGS with content pertaining to artworks. Seeing the painting is the intuition-prompting experience. That you possess such an IGS makes it possible for you to have an art identification intuition at the museum that turns out to be accurate.

In this thesis, I argue for the following two theses:

Thesis 1: There exists a class of intuitions which cannot arise without experience-dependent *intuition generating systems* (IGSs) that are influenced by *IGS-shaping experiences* encountered during one's life. I call these *experience-based intuitions*.³

Thesis 2: The reliability of an experience-based intuition X depends on the quality of the IGS-shaping experiences underlying the IGS which produces X.⁴

The plan is as follows. In section one, I outline five features of intuition. In section two, I argue that the conscious mental state which prompts your identification of the van Gogh painting is an experience-based intuition. The first two sections constitute my defense of Thesis 1. In section three, I argue that experience-dependent IGSs are structures which perform hierarchical abstraction from experiential data. I briefly consider a computational model for such IGSs. In section four, I argue that the structure of these IGSs explains the varying degrees of reliability of experience-based intuitions and address some skeptical challenges. I also suggest that even if

³ The existence claim in Thesis 1 does not preclude the possibility that some IGSs are *a priori* or shaped primarily by evolutionary forces.

⁴ Joshua Greene (2017) has suggested that the acquisition of good moral intuitions require representative data and value-aligned training, which appears to agree with Thesis 2.

intuitions cannot always guide us rightly, identification intuitions rooted in experience may be a necessary component of many judgments. In the last section, I address some implications of my account to philosophical methodology. In particular, I suggest that some philosophical thought experiments produce experience-based intuitions of the same sort as those prompted in everyday situations. I end by suggesting that even if these thought experiment intuitions do not reveal a priori truths, they could uncover veiled conceptual commitments and tell us about the experiences that underpin our intuitive judgments.

2 FEATURES OF INTUITION

Some theorists have noted that the kind *intuition* is heterogeneous. Jennifer Nado (2014), for instance, points out that intuitions across domains have diverse etiologies and content. Even so, there remains some consensus about features that a mental state should have in order to be called an intuition. Here I enumerate five commonly accepted features of intuition.

There is a consensus that intuitions:

(1) Are conscious

Intuitions are experienced consciously – it *feels* a certain way to have an intuition. The terms many philosophers use to describe having an intuition unveil a shared assumption that intuitions possess a subjective quality of experience. George Bealer (1998), for instance, describes an intuition as a “conscious episode” (208). Railton (2014) describes the experience of intuiting as finding “ourselves with a spontaneous ‘sense’ (815). And, as Panaccio (2014) puts it, an intuition “is itself a distinct reality within the mind” (65).

(2) Are induced by sensory or imagined input

Intuitions are most often triggered by an experience, as when a trial attorney gets the sense from the expressions of jurors that her argument fails to convince (Railton 2014: 818). Figments of our own imagination can also prompt intuitions.⁵

(3) Are induced non-voluntarily

We may choose to engage in activities that end up giving rise to an intuition, but having this choice does not grant us control over the resulting intuition. When we describe an intuition as being triggered by an experience, we mean that the intuition *happens* to us. For instance, though consciously entertaining a thought experiment may trigger an intuition, that intuition is

⁵ Chandra Sripada (2016) suggests that future-oriented affects and intuitions can result from imagined scenarios.

not formed spontaneously or in a way we can control (Gopnik & Schwitzgebel 1998).

Spontaneity also doesn't imply habitual recalling of information learned by rote (Ohlsson 2011).

(4) Are produced by processes that are introspectively opaque to the intuitor

Though intuitions are consciously experienced, *how* we come to experience them is opaque to us (Kornblith 2002). When we attempt to explain how we have arrived at an intuition, we might come up with something that sounds convincing. But such explanations amount to *post hoc* justifications, since we lack direct access to the processes generating our intuitions. This inaccessibility can be especially glaring in cases where our explanations fail to support the intuitions in question. We might even be surprised by our intuitions because we have little grasp on how they come about.

(5) Are disposed to guide deliberation, judgment and/or action

Theories of intuition abound because there is a shared assumption that intuitions do guide actions to varying degrees, regardless of whether we think they *should*.⁶ Many think intuitions amount to judgments, while others, including myself, disagree.⁷ Either way, intuitions have a strong disposition to guide further reasoning and action. However, intuitions can at times be recalcitrant in the face of contrary evidence or conscious reasoning, leading to judgments that may conflict with other attitudes held by the intuitor.

In sum, intuitions generally (1) are conscious, (2) are induced non-voluntarily, (3) are induced by sensory or imagined input, (4) are produced by processes that are introspectively opaque, and (5) are disposed to guide deliberation and action. But does intuition really play a

⁶ Some have suggested that philosophers, for instance, may not crucially rely on intuitions when philosophizing (Cappelan 2012, Deutsch 2010, Ichikawa 2014, Williamson 2007).

⁷ To proceed from having an intuition to having a propositionally articulated judgment, interpretation in light of existing beliefs may be an intermediary step (McGahhey & Van Leeuwen, 2018). Intuitions may thus be necessary though insufficient conditions for some judgments.

role in helping you identify a van Gogh painting?

3 EXPERIENCE-BASED INTUITIONS

The experience that you have in response to the van Gogh painting is an intuition in the relevant sense, since it has all five features of intuition outlined above.

(Feature 1) You are suddenly overcome with a conscious, if murky, sense of familiarity when you encounter the painting.⁸ It just seems to you like a van Gogh.

(Feature 2) This sudden sense is prompted by the perceptual input you receive from your viewing experience.

(Feature 3) Although a series of deliberate actions – such as those involved in planning a museum visit – make your encounter with the painting possible, a sense about the painting’s authorship comes to you spontaneously. You need not have the explicit goal of identifying any artwork to in fact identify the van Gogh painting. Further, no effortful reasoning is employed in the moment of intuiting. You simply don’t compare the artwork you are viewing, feature by feature, against a repertoire of potential matches recalled from memory.

(Feature 4) The processes leading to the identification are opaque to you. You just get the sense that you’re looking at a van Gogh without being able to explain how you came to that insight. When asked to justify the identification, you might appeal to certain visual features of the painting. But each explanation you articulate fails to sufficiently support the identification, since you would describe many other paintings in the museum in the same ways.

(Feature 5) Lastly, this hazy sense, rather than other forms of reasoning, is what leads you to identify the painting. Moreover, even if a well-informed fellow visitor claims that the museum houses no artworks by van Gogh, your sense might persist.

As we have seen, all five features obtain for the art identification intuition. Thus, the

⁸ Though I focus my discussion on the identification component of the van Gogh encounter, other phenomenological content is also likely present, such as vague like or dislike of the painting.

general conception of intuitions as a mental state should include similar identification intuitions. And since intuitions of identification co-exist amongst other intuitions under the same psychological class, they are of comparable theoretical interest.

A thread unifying many intuitions across divergent areas of expertise is that prior experience is *necessary* for generating such intuitions. For the sake of clarity, in this paper I focus on the identification of artistic style as a case study. But similar identification intuitions seem to be relevant in medicine (Cork et al. 2014, Hogarth 2001, Elstein et al. 1990), chicken sexing (Horsey 2002), chess (de Groot 1965 and 1996, Chase & Simon 1973), business (Provis 2010) and military operations (Banks & Dhami 2014). Just as a chess novice is guided less by an intuitive grasp of board positions than by consciously recalling rules of the game, someone who is unfamiliar with any work by van Gogh simply cannot have the intuition you had at the museum.⁹ Such identification intuitions are marked by a crucial and distinguishing characteristic – they cannot be *a priori*. Thus, rationalist accounts fail to explain at least some intuitions.

Further, the cognitive mechanism underlying art identification intuitions would be mischaracterized as evolved systems. Here one might object to my claim by suggesting that such an intuition could be the byproduct of an evolved cognitive system responsible for identification of novel objects in the environment, which system had aided our ancestors in the detection of potential threats. To be sure, evolved systems may help to explain *all* intuitions to a limited extent. After all, our visual apparatus – a system necessary for visually identifying an artwork – has itself been shaped by evolutionary forces. However, an evolved system within the context of intuition production is usually characterized by domain specificity,¹⁰ ontogenetic stability, and minimal reliance on and improvement by training. Consider intuitions about biological kinds,

⁹ Such a person would also be hard-pressed to identify a van Gogh through non-intuitive means.

¹⁰ Though there are dissenters – see Barrett and Kurzban (2006).

which in part help us understand that animals within a species share certain features. Intuitive biological thinking bears the mark of being the product of an evolved system of intuition generation. It applies specifically to intuitions about species. More importantly, essentialist biological thinking persists at least from adolescence into early adulthood and is influenced minimally by biology education (Coley et al. 2017). Not only are we endowed with rigid “templates” for biological kinds (Boyer 2001), these templates, even in cases where they become unhelpful, remain largely unsusceptible to revision. In contrast, art identification intuitions only *appear* to be domain specific insofar as they pertain to one area of experience, but the mechanisms underlying visual identification applies across the board to all images. Identification intuitions also depend on experience rather than on ontogenetic development. That is, when it comes to artistic style, the accumulation of art viewing experiences, rather than developmental stage, is what determines the possibility of having an intuition. The same holds true for recognizing board positions in chess, for triaging trauma patients, and for identifying the composer of a piece of music you’re hearing for the first time. Moreover, even if it is possible to construct templates helpful for artistic style identification – e.g., blue and yellow swirls = Vincent van Gogh – those very templates must be devised by individuals who are familiar with works by van Gogh and other artists. Therefore, the van Gogh intuition, and others like it, are not adequately explained by evolved systems accounts of intuition.

To be sure, not all expert intuitions are experience-based intuitions. While a trauma nurse’s triage decisions are experience-based intuitions of identification, not all medical decisions qualify as such intuitions. For instance, a physician may come, through repeated practice, to perform biopsies in a way that appears second nature. Robert McCauley (2011) suggests that *practiced naturalness* may arise from extensive practice and may explain some

expert intuitions. On my account, practiced naturalness is distinct from experience-based intuitions, since, according to McCauley, *effortful practice* is essential in the acquisition of practiced naturalness. While effortful practice such as reading a medical textbook is *sometimes* part of the prior experience enabling one to have an experience-based intuition, such as might be the case in the triage nurse example, it is certainly not *necessary*. After all, the process by which we come to be able to intuitively tell chairs apart from non-chairs involves little to no conscious effort or instruction. Further, McCauley's account of practiced naturalness also focuses on learned *physical* skills such as swinging a golf club, rather than on intuitions which enable conscious decisions. To illustrate the difference, consider the sort of practice that world-class soccer players go through versus the sort of experience that could allow someone to identify that a pass was successful or not. Consider also the differences between being a comedian who excels at comedic timing and being able to identify good comedic timing. In each case, the latter ability requires prior familiarity, while the former ability constitutes a skill requiring extensive *practice*.

In sum, cases such as the van Gogh example involve an identification intuition made possible by relevant prior experience. By definition, intuitions are produced by intuition generating systems. Specifically, an IGS responsible for producing an intuition such as an art identification intuition (1) is not the product of any specific evolutionary pressure and (2) requires extensive experience, rendering it different in kind from some other IGSs. Thesis 1 has thus been established. Next I turn to the structure of experience-dependent IGSs.

4 EXPERIENCE-DEPENDENT INTUITION GENERATING SYSTEMS

Everything which we observe imprints itself uncomprehended and unanalyzed on our percepts and ideas, which then, in their turn, mimic the process of nature in their most general and striking features. In these accumulated experiences, we possess a treasure-store, which is ever close at hand, and of which only the smallest portion is embodied in clear articulate thought. (Mach 1883/1960: 36)

In the passage, Ernst Mach suggests that we obtain much unthematized or unanalyzed information from the environment. Of that information, a small portion is encoded into articulable knowledge while a larger portion is organized non-theoretically into “percepts and ideas” not subject to conscious control. The resulting percepts and ideas in turn help us process information and navigate through the environment. Following Mach,¹¹ I suggest that mental categories derived from experience inform identification intuitions. To imbue my account of experience-based intuitions with theoretical richness, in this section, I argue that: (1) IGSs responsible for producing identification intuitions operate by generating abstract categories from experience, (2) deep convolutional neural networks (DCNNs) provide a helpful and biologically plausible model for experience-dependent IGSs, and (3) all five features of intuition are well-explained by such a model.

4.1 Three characteristics of experience-dependent IGSs

To begin, an IGS capable of producing experience-based intuitions of identification should have three characteristics. First, as I have argued in section two, it should be a general-purpose mechanism that is highly receptive to training through experience (*general-purpose*). Second, it can generate representations of categories by grouping similar objects through detecting irrelevant information. This process of abstraction should be hierarchical, such that increasingly fine-grained categories can be formed from one sufficiently large repertoire of

¹¹ According to Panaccio (2014), Ockham has a similar account of concept acquisition which posits the human mind as being endowed with a mechanism for generating representations from exemplars.

experience (*hierarchical category grouping*). An example of categories ranked from most to least general might be: artifacts, artworks, paintings, Impressionist paintings, paintings by van Gogh. Relatedly, the intricacy of hierarchical categories represented by an IGS is positively correlated with the amount of experience that has gone into shaping it. Third, an experience-dependent IGS should have the flexibility to use one experience to inform multiple categories, not just within one hierarchical lineage (*cross-category abstraction*). For instance, seeing the *Mulberry Tree* might modify artistic style categories in addition to other categories such as *tree*.

4.2 Value of a computational model

Judging from these three characteristics of experience-dependent IGSs, deep convolutional neural networks (DCNNs) seem to have the potential to be a helpful model. Before delving into their mechanism, I want to briefly discuss the virtues of using (primarily) a computational model to explain a human cognitive structure. First, artificial neural networks are often modeled on mammalian neural networks. While artificial neural networks solve machine learning problems, they are nonetheless biologically realistic models. Second, when it comes to visual perception, there is evidence suggesting that the 6-layer “deep” structure of the mammalian neocortex and the hierarchical processing in the ventral stream have analogues in artificial neural networks like the DCNN.¹² From these similarities, Cameron Buckner (2018) has recently argued that DCNNs model core aspects of abstraction – what I argue to be a key capacity of experience-dependent IGSs – in the mammalian brain. Third, the strategy of using computational models to explain human abilities has a demonstrable lineage. For example, O’Loughlin and Thagard (2000) use a connectionist model to enrich and test Uta Frith’s (1989) weak central coherence theory of autism. Chandra Sripada (2016) proposes a deep learning

¹² Buckner 2018. See also Khaligh-Razavi & Kriegeskorte 2014, Yamins & DiCarlo 2016.

account of the role of mind-wandering in prospection. Bayesian models of a plethora of human capacities also abound.¹³ Lastly, one might object that AI researchers are sometimes at a loss as to how artificial networks that they themselves created are able to produce certain results, suggesting that philosophers are ill-advised to look to AI for answers about intuition production. However, there is a rich body of machine learning literature on how artificial neural networks operate, what sorts of tasks they can perform, how their performance can be improved, etc. Not to mention, a computational model offers exciting opportunities for testing hypotheses, since artificial networks are more easily manipulated than are human subjects. Thus, even acknowledging the possibility that DCNNs might later turn out to be an imperfect model for experience-dependent IGSs, using such a model is still valuable at present.

4.3 How DCNNs explain experience-dependent IGSs

In contrast to other types of artificial neural networks, DCNNs are “deep” because they have multiple intermediate layers between input and output layers. According to Buckner (2018), convolutional filters pass outputs up the processing hierarchy by detecting and amplifying the presence of desired features and minimizing other information (19-20). At present, DCNNs primarily perform recognition tasks. Impressive practical applications for DCNNs include detection of tumors in scans and natural language processing. Here I focus on how they perform in detecting and naming visual content, since such tasks are most comparable to the van Gogh example.

DCNNs possess the three characteristics of experience-dependent IGSs, since they are general-purpose learning mechanisms that perform hierarchical category grouping and cross-category abstraction. First, DCNNs are general-purpose systems that are highly responsive to

¹³ Battaglia et al. 2012, Oliver et al. 2000, Tenenbaum 1999

training. Even the most intricately-designed algorithm cannot produce accurate results if it lacks sufficient training input. DCNNs are also general-purpose in that an algorithm equipped to process images can process images of any kind. Moreover, DCNNs can be trained without supervision. Human IGSs are not solely shaped by labeled information – a minority of the information gleaned through the environment is labeled, mostly during our formative years and in formal education. A DCNN can similarly learn from unlabeled datasets without explicit instructions (Silver et al. 2017).¹⁴ In short, not only are DCNNs trained by experience, they are trained in ways that are analogous to how experience-dependent IGSs are trained.

Second, DCNNs are unique amongst artificial neural networks in their ability to perform hierarchical abstractions. The presence of nuisance variables, including size, position, and angular rotation in visual identification tasks has long stumped rule-based algorithms. For instance, a chair viewed from the bottom can appear vastly different from the same chair viewed from the front, even after controlling for size and proportion. Low-level visual features of images can exhibit such variety to render it impossible to compose a rule comprehensive enough to capture all images of chairs. But DCNNs are not rule-based, direct input-output systems. They are multi-layered systems that perform abstractions. DCNNs are remarkably adept at accentuating task-relevant features while controlling for nuisance variation (Patel et al. 2015). At each convolutional layer, idiosyncratic presentations of images, such as lines at different positions, sizes, or orientation get filtered out to produce more abstract presentations that are fed to the next layer (Buckner 2018: 23). Moreover, DCNNs' performance of hierarchical category grouping from exemplars has an analogue in human memory consolidation and learning.

¹⁴ Promising new developments in AI combine reinforcement learning with an “episodic buffer” that replays sessions in “offline training,” simulating memory consolidation during sleep and daydreams in mammals. See also Blundell et al., 2016; Hassabis et al., 2017; Kumaran, Hassabis, & McClelland 2016; Mnih et al., 2015.

According to an influential account of complementary learning systems by McClelland et al. (1995), hippocampal synaptic changes associated with new memories trigger changes in the neocortical system responsible for the consolidation and continued revision of remote memories, suggesting that one process of learning from experience involves extracting underlying commonalities from accumulated episodic memories.

Turning now to the third feature of experience-dependent IGSs, DCNNs need to have the ability to generate across categories from a given selection of data. The intuition that interests me in the van Gogh example is not one that tells you the painting depicts a tree, or one that alerts you to the fact that you're looking at a painting. Rather, the intuition in question concerns the *style* of the painting. But DCNNs are tested primarily on identification of images of everyday objects such as chairs. To be successful at such a task, an algorithm needs to treat style of depiction as a nuisance factor. Can a DCNN trained on a selection of paintings generate categories based on style *in addition* to identifying the objects depicted? As it turns out, they can. In fact, modelers of DeepArt¹⁵ rely on the ability for an algorithm to extract input images' style to render any image uploaded by users in the style of famous artists.

There are a couple of ways to account for this capability, both of which lend further support to using DCNNs to model experience-dependent IGSs. The first way is through considering task relevance. With respect to any perceptual inference task, the target to be identified is distinguished from nuisance factors according to the goal. For a museum visitor, the frame in which a painting is displayed and the color of the wall surrounding the painting are appropriately considered nuisance. But if the task is to distinguish a framed painting from an unframed one, then the frame itself is relevant. What a DCNN takes to be nuisance is similarly

¹⁵ <https://deepart.io>

task relevant. Second and more intriguingly, one can also appeal to Buckner’s notion of *transformational abstraction*. In short, DCNNs must generate information about the nuisance that has been subtracted from exemplars when performing abstractions. An algorithm trained on paintings featuring chairs does not only generalize about the category *chair* but also artistic styles, since an input image’s style needs to be well-detected by the algorithm or it cannot be filtered out.¹⁶ How DCNNs perform abstractions across categories without explicit direction to do so mirrors how we sometimes extract information from the environment.

Finally, though I’ve focused my discussion of DCNNs on visual categorization tasks, they may perform similarly in the generation of many other types of categories.¹⁷ Insofar as I take experience-based intuitions to encompass far more than the identification of images, the possibility that artificial systems like DCNNs can generate non-perceptual abstractions is another point in their favor as a helpful model.

4.4 Five features of intuition revisited

As I have shown, DCNNs provide a biologically plausible model for experience-dependent IGSs. They are general-purpose systems of hierarchical abstraction that can create multiple categories from a given set of data. Let’s return to the five features of intuition in the van Gogh case study to see how an experience-based identification intuition could be produced by such a system. First, your intuition is triggered by viewing the painting. Your past knowledge of art has contributed to shaping the IGS that produced this token intuition you are experiencing in the gallery. Second, viewing the painting constituted the right sort of intuition-prompting experience such that the corresponding IGS is triggered and spits out an intuition. Third, you

¹⁶ Buckner 2018, Gatys, Ecker, & Bethge 2016

¹⁷ Evidence suggests that AlphaGo (an algorithm that has defeated expert Go players at a notoriously complex game) performs transformations beyond visual modalities, encompassing abstract notions including “influence,” “connection,” and “stability” which are cited by human players as key notions of the game (Buckner 2018: 28).

don't have much if any control over how the intuition comes about in the moment of intuiting, since the generation of categories is sub-personal and not available to introspection. Fourth, the purpose of the IGS is to allow unthematized processing to feed information to consciousness. Thus, it makes sense that the output is an intuition that is felt. Lastly, this output is interpreted in light of background beliefs to result in the realization that you are looking at a van Gogh painting.

5 RELIABILITY (AND NECESSITY) OF EXPERIENCE-BASED INTUITIONS OF IDENTIFICATION

With the structure of experience-dependent IGSs in mind, I now turn to Thesis 2: the reliability of an experience-based intuition depends on the quality of the experiences that shape the corresponding IGS. Based on Thesis 1, I argue that while the accuracy of token intuitions varies, the mechanism which produces experience-based intuitions is generally reliable. I review two lines of skeptical challenges and respond to them. I highlight the importance of distinguishing the *reliability* of an intuition-generating system and the accuracy of token intuitions that it produces from the *stability* of categories of objects we identify intuitively. I end by suggesting that though experience-based intuitions of identification are fallible just like other sources of justification, they may be necessary components of decision making.

Experience-dependent IGSs have the potential to generate remarkably accurate identifications. But the capacity for having good intuitions does not guarantee that you will in fact have them. Having a sufficiently large repertoire of varied and representative experiences is vital in the generation of accurate intuitions. The more abundant and varied the exemplars, the more sophisticated the corresponding IGS. Consider an art historian, a college freshman taking an introductory course, and a person who has never seen a painting. The naïve person would be hard-pressed to experience *any* intuition of the authorship of any painting, much less an accurate one. As you become familiar with more artworks, the IGS starts spitting out intuitions when prompted by viewing experiences, but these intuitions are only sometimes accurate. With increased exposure to works across mediums and styles, provided that labeled information is labeled correctly, the identification intuitions of the art historian become increasingly accurate, owing to a much more robust IGS that has been shaped by many exemplars. Experience thus

explains the heightened reliability of expert intuitions.¹⁸ Simply put, good experience is necessary for good identification intuitions.

But is experience sufficient for producing intuitions that we can rely on? I now consider two lines of challenges to the reliability of intuition as a source of justification. To begin, reliability is usually defined in terms of the extent to which intuitions lead to accurate judgments. That is, an art identification intuition is reliable insofar as it leads to the correct identification of a painting, and a medical diagnostic intuition is reliable to the extent that the patient in question has the illness with which she has been diagnosed. One line of skeptical challenge claims that token intuitions can lead us astray if we rely on them to make decisions. Thus, intuitions in general are claimed to be unreliable. A second line of objection claims that even if intuitions at times lead to accurate judgments, they cannot be said to be reliable if they merely track consensus instead of truth (Koriat 2008).

The first line of objection is often framed in terms of moral intuitions being inflexible in comparison to conscious reasoning. For instance, Greene (2017) has insisted that even if experience can improve intuitions generally, any token intuition would be inflexible at the time of deployment. In response, it should first be noted that such objections conflate intuitions with intuitive judgments. As mentioned earlier, intuitions motivate judgment but require interpretation. When some skeptics examine token intuitive judgments that they deem to be subpar, they should take care to locate the source of the faulty judgment. The problem may have arisen on the way from intuition to intuitive judgment – one may have interpreted a good intuition in light of false beliefs.

Further, while it is true that in the moment of intuiting, an intuition can only be as good

¹⁸ It might be easy to point out a Seurat even in a room full of Impressionist works. But identifying a Raphael in a room full of Renaissance paintings might prove to be a challenge to the average art enthusiast.

as prior conditioning, other acceptable sources of justification are susceptible to similar inflexibility. Nagel (2012, 2014) has argued that the existence of perceptual illusions does not preclude vision from having justificatory power. Analogously, the possibility of illusions in epistemic intuitions does not preclude epistemic intuitions from being generally trustworthy indicators of what counts as knowledge, especially when said intuitions are prompted by meticulously crafted scenarios. In general, it seems unfruitful to determine the reliability of a broad class of psychological states on token instances. After all, even conscious reasoning is fallible, not least when the information used to reason with is incorrect. Instead, it may be more helpful to focus not on the quality of certain intuitions but rather on investigating the robustness of intuition generating systems.¹⁹ Moreover, the influence of contextual or framing effects on intuitive judgments may be more easily investigated in non-moral intuitions of identification. Whereas experiments testing people's moral intuitions are often susceptible to the criticism that the designated "right" answer is in fact wrong on certain interpretations, experiments involving non-moral intuitions of identification may be much less prone to such experimental design flaws.

My account might also help to address some concerns of inaccurate token intuitions argued from interpersonal inconsistency (Goldman 2007). Conflicting intuitions across individuals might be explained by appealing to differences in cognitive abilities, such as those involving memory consolidation and retrieval, as well as the quality and depth of prior experience.²⁰ After accounting for the differences in experience, we may find that each person's intuition, though they conflict interpersonally, are well-justified. In short, inaccurate intuitive judgments do not entail that intuitions are less reliable than other means of justification.

In response to the second challenge on the capacity for intuitions to track truth, I

¹⁹ This is in line with what process reliabilism advocates (Goldman 2006).

²⁰ I explore the implications of my account for social epistemology in my thesis.

recognize that experience-based intuitions are rooted in personal categories rather than natural kind categories.²¹ However, the general trend remains that with experience, an IGS forms categories that are increasingly fine-grained and resistant to nuisance. Further, experience does not merely come from idiosyncratic encounters. Some high-quality experience may come from testimony from reliable persons, not to mention formal education. More importantly, while some take consensus to be a lackluster measure of accuracy, many targets of investigation are artifacts the definitions of which are moving targets. Such concepts as chairs, Impressionism, knowledge, and moral salience may be best defined by consensus. When art authenticators exhaust means of chemical analysis and provenance research, consensus among experts – all of whom may be relying on intuition – may be the only way to determine the authorship of a newly discovered painting. Sometimes consensus is all we could hope to track when truth is not easily defined.

Underlying my responses to both challenges to the reliability of intuitions is the important distinction between the *reliability* intuition-generating systems (and the accuracy of their resulting token intuitions) versus the *stability* of the environment in which we experience intuitive judgments. Take moral intuitions as an example. While intuitive judgments may seem unreliable for difficult moral cases, including edge cases for in which people do not generally have adequate training, the cause of such unreliability is often the inherent instability or inconsistency in morality itself rather than the unreliability of intuition generating systems. Groups of people who are epistemic peers may stably arrive at similar moral intuitions given their common experiences. But whether those intuitions lead to moral judgments that all moral agents would deem to be accurate token intuitions is a separate issue. In other words, groups of epistemic peers who share similar experiences in the moral realm may arrive at comparable

²¹ Goldman (2007) construes the targets of philosophical analysis as concepts in the psychological and personal senses (6), not as concepts tracking natural kinds.

moral intuitions owing to the reliability of experience-based intuition generation systems, but the accuracy of those moral intuitions are not always evaluated by their epistemic peers. Moreover, according to Kahneman and Klein (2009), “skilled intuitions will only develop in an environment of sufficient regularity which provides valid cues to the situation” (520). Morality may be one such unstable environment, along with Kahneman and Klein’s example of the stock market. In contrast, highly stable environments include chess, medicine, and art identification, among others.

Lastly, I want to highlight the necessity of experience-based identification intuitions. Some skeptics conclude that since intuitions are fallible due to various reasons, we ought not to rely on them. However, they neglect to recognize that, setting aside the question of whether we are justified in using such intuitions, we may not be able to avoid using them. Consider again the popular contention that since moral intuitions are less flexible than conscious reasoning, we should favor reasoning over intuition. It seems to me that this sort of claim makes the mistake of unduly circumscribing intuitions to the realm of only some subsets of intuitions. Even proponents of intuitions are susceptible to making this mistake. For instance, Nagel (2012) suggests that reflective thinking is triggered in novel cases pertaining to the identification of knowledge (500). However, the very recognition that a given case is unusual involves an identification intuition that is experience-based. We don’t come to the realization that a scenario is unusual by comparing it, feature by feature, to a “usual” case. We might not even be able to articulate which features are relevant to defining a scenario as belonging to one kind versus another. Instead, some things just strike us as unusual when we encounter them. This realization of something as being different than expected at least involves discerning relevant features of the case at hand that distinguishes it from other similar cases. Thus, even if we prefer to carefully

reason through a novel case before making a judgment, we still rely on an identification intuition rooted in experience to start the process.

Thus, it seems that good empirical experiences give people better intuitions across a wide range of intuition categories that are experience-based. Some common challenges against the reliability of intuitions do not successfully demonstrate that intuitions are in general less reliable, or more dispensable, than other sources of justification.

6 EXPERIENCE-BASED INTUITIONS IN THOUGHT EXPERIMENTS

I would like to devote the last section to a discussion of some implications of my account of experience-based intuitions to debates in philosophical methodology, since the nature and reliability of intuitions might contribute to greater understanding of the role of intuitions in thought experiments. I suggest that the insofar as the mental states elicited by philosophical thought experiments are intuitions, what we get from the case method is likely neither the objective nature of the target of investigation nor a set of its necessary and sufficient qualifying conditions. On the other hand, we might also be overstating the case if we conclude based on discrepant case method judgments across individuals that philosophers ought not to appeal to intuitions, whether their own or that of non-philosophers. I suggest that some intuitive judgments elicited by thought experiments, insofar as they are prompted by the sort of experience-based descriptive intuitions of identification I've thus described – and I argue that many thought experimental intuitions are of this kind – tell us about the quality of experience we have in domains relevant to the subject under investigation. Our intuitions, however divergent across individuals, also have the potential to reveal conceptual commitments that are not directly accessible.

6.1 Do some thought experiments elicit experience-based identification intuitions?

The goal of this section is to get to the nature of what is revealed by thought experiment intuitions. But first it needs to be established that thought experiments in fact elicit experience-based identification intuitions that of the same kind as intuitions we have in response in everyday situations. I argue that not only do thought experiments elicit intuitions, such intuitions as a whole are also not importantly distinct from intuitions we experience in everyday situations. I also suggest that many thought experiments in philosophical literature seem to fall into the

category of experience-based intuitions of identification which can be explained by the account I have thus put forth.

To begin, there is a tradition of employing thought experiments as a means of philosophizing. Many, including Edouard Machery (2017), use the term “method of cases” to refer to the methodology that philosophers use when they present a real or imagined vignette and then reason from his or her judgment of the relevant philosophical issue at hand from the vignette. Machery writes, “cases are descriptions of actual or hypothetical situations, and philosophical cases are cases put forward by philosophers” (11). Thus, I will use thought experiments and the method of cases interchangeably. It’s also generally accepted that philosophical thought experiments are “almost always meant to elicit a judgment or some other mental state about the situations they describe” (Ibid.). Though Machery’s phrasing leaves room for one to argue that thought experiments can elicit any number of mental states, including intuitions, he later goes on to dismiss the need to describe the relevant mental state in the case method as intuitions, citing a multitude of reasons, not least of which is a plethora of ways in which the notion of intuition is cashed out (36).

Yet not only is it common to describe the method of cases as involving intuitions, it again seems important here to reiterate the distinction between intuitions and judgments which are prompted by intuitions.²² While it may be true that philosophers rely on explicit and explicable judgments when making arguments – after all, the content of the judgment should be made clear to the reader who is following along attempting to glean the point that is made by the use of the thought experiment – it’s a separate question whether the mental state immediately elicited when the thought experiment is considered amounts to a judgment. In other words, just because

²² Goldman 2007, etc.

philosophers take the propositional attitude that amounts to a *judgment* which resulted from considering a thought experiment to make their argumentation, it does not follow that an intuition plays no role in the process. In fact, we have seen that intuitions appear to not only have a distinctive phenomenology and are triggered in certain conditions but also play a role in influencing judgment and action. The judgment that a certain mental state is or is not a justified true belief in Gettier cases is informed by an intuition, even though the intuition itself, because it lacks content, plays much less visible of a role in argumentation.

Further, not only are intuitions triggered by thought experiments, these intuitions are also not distinct in kind from intuitions that we experience in everyday situations, such as the intuition in the van Gogh case. Though there are differences we could draw between them, there doesn't seem to be good reasons for why those differences provide enough justification for drawing a hard line between them. According to Machery, there are three ways in which the mental states elicited by thought experiments can be characterized:

1. The exceptionalist posits that there is something importantly distinct from about the intuitions elicited by philosophical thought experiments that sets these intuitions apart from everyday intuitions in various ways (17). It might be posited that thought experiment intuitions have distinct phenomenologies – e.g., a sense of heightened necessity, as though your intuition must be right about the thought experiment. They could also have a distinct epistemic role, for instance, in being *a priori* justified. And they could have a distinct etiology, such as expressing one's conceptual competence.²³
2. The particularist holds that the mental states prompted by thought experiments are a particular type of everyday judgment (19). As such, they have certain properties that some, but not all, of everyday judgments possess. For instance, some particularist characterizations may identify thought experiment intuitions by means of their content, phenomenology, epistemic status, or etiology, while maintaining that these properties are par for the course for intuitions in general.
3. The minimalist holds that philosophical thought experiments do not elicit attitudes that are distinct in kind from those elicited by non-thought-experiment cases (20). While thought experiment intuitions are often prompted by descriptions of situations one should not count on encountering in real life (e.g., Swampman), the attitudes themselves are of the same kind as everyday intuitions.

²³ Though it is not clear to me what this is in opposition to.

I take the minimalist approach, since thought experiments possess the same defining features that other intuitions possess. They neither possess properties that distinguish them from everyday judgments (except for their unusual subject matter) nor can they be identified with a particular type of everyday judgment (20). Thought experiment intuitions, according to Machery, do not possess a phenomenology that everyday intuitions do not have (e.g., they do not carry a heightened sense of urgency), do not have distinctive epistemic status (e.g. they are not justified *a priori*), do not have a distinctive semantic status (e.g., they are not analytic), they do not have a distinctive etiology, and so on (20). Further, while one could make the argument that the subject matter of thought experiment intuitions often being unusual could qualify them as a distinctive sort of mental states, as Machery also notes, everyday judgments can also be about unusual cases. He writes, “when we judged that the first iPhone was a phone, our judgment had a novel subject matter, but it was of the same kind as other application of the concept PHONE. Generally, everyday judgments can be made about esoteric subject matters” (21). I agree with Machery on this point and contend that thought experiment intuitions seem to play the same epistemic role as other intuitions, even though they tend to be prompted more by imagined and unusual input. Furthermore, thought experiments can indeed involve descriptions of actual rather than hypothetical scenarios. In cases where the case is actual rather than hypothetical, it seems a far reach to contend that just because a philosopher is describing the scenario that the resulting mental state in the intuitor needs to be categorized differently.

There remains the question of whether some thought experiment intuitions fall under the class of experience-based intuitions of identification. It seems to me that many philosophical intuitions are indeed of this kind. For starters, I think that when we consider the goal of many

thought experiments used in the areas of philosophy of language, epistemology, and action theory, for instance, the aim of the cases is to draw out intuitions that help with classification. Goldman (2007) has suggested that philosophers consider actual and hypothetical examples and ask whether these examples provide instances of the target category or concept – e.g., knowledge, reference, causation, etc. (1). The mental responses that people have to these cases, which he happily calls intuitions, are treated as evidence for category membership of the case. Classification intuitions, according to Goldman, are intuitions about how cases are to be classified, or whether various categories or concepts apply to selected cases (4). Indeed it seems that thought experiments such as Gettier cases or Twin Earth are put forth with the explicit aim of prompting identification intuitions that help us make judgments about what sort of concept is represented in the case. As Goldman puts it, the “discovery” that knowledge isn’t equivalent to justified true belief was made not through Edmund Gettier’s declaration about the category membership of his examples, but rather by the agreement in the intuitive judgments that Gettier and his readers shared about the examples. In other words, intuitions in response to many philosophical thought experiments, specifically those intuitions that prompt us to make judgments about what sort of concept we are dealing with.

Further, these identification intuitions seem to rely on experience in the sense that it seems implausible that one could have intuitions about edge cases of what could be considered knowledge if one does not have a good foundational grasp of more mundane cases of knowledge might look like. It seems far-fetched to say that the concept of knowledge, for instance, has been endowed through what I have called evolved systems. The concept of knowledge might not be something that we need to acquire through *formal* learning, but it seems to require everyday experience having to do with knowing things and forming implicit notions of what knowledge

might constitute based on those experiences. Goldman posits that the process of generating classification intuitions has more in common with memory retrieval than with purely intellectual thought, the core of the *a priori* (20). I agree and take that the identification of personal category membership in response to philosophical cases relies significantly on prior experience.

6.2 What can experience-based thought experiment intuitions reveal?

What, then, is gleaned through the identification intuitions we experience when we consider philosophical thought experiments? Some might contend that through our intuitions, we glean rational, *a priori* truths about the objective categories we are attempting to draw. Skeptics, including Machery (2017), have argued that since there is such vast interpersonal variation in case method intuitive judgments that our intuitions really are not revealing much about their intended, philosophically relevant targets, but perhaps our own biases and prejudices. Some skeptics might go further and suggest that as result of the various biasing effects correlated with interpersonal variations in intuitive judgments, appeals to intuition as a philosophical methodology is ill-advised in general.

On my account, intuitions in response to thought experiments are poised to convey a lot of useful information. First, they reveal information about our experience in the relevant domains pertaining to the subject matter in the thought experiment. Just as the accuracy of my van Gogh intuition tells me about the quality and depth of my prior art-viewing experiences, so too does my intuition leading to the identification of philosophical categories reveal experience with these categories. This is especially apparent, for instance, when the case presented is meant to probe intuitions about object category membership. Consider, for instance, being presented with a description of an edge case of an object that might be used for sitting, and being asked to consider whether that object might qualify as a chair. One's prior experience and knowledge of

chairs would be crucial in informing the intuition, even though the intuition is here prompted by a description, perhaps coupled with an imagined object from the description, rather than by viewing the actual object. Furthermore, the way in which we imagine this description of the object used for sitting may itself be informed by prior experience. That is, the sorts of chairs I've come across prior to reading the description of a hypothetical chair affects the object that I imagine as I'm reading the description. In this way, the category-identifying intuitive judgment that I get from the description is deeply intertwined with experience, and, upon reflection, has the potential to reveal that experience and how that experience might have shaped the current intuition in question. This also means that biases may also be revealed, if we consider the sorts of experiences and how our position in the world makes certain experiences possible and more likely. Not to mention some philosophers have constructed thought experiments using highly atypical cases for illustrative purposes. Consider Donald Davidson's Swampman case. Rather than genuinely testing the confines of personal identity, such a case may be designed foremost to be illustrative and used to reveal the reader's theoretical commitments. So Machery's conclusion is helpful as long as we don't use it to discount thought experiment intuitions altogether.

Furthermore, thought experiment intuitions may help uncover the stability of a given environment as well as the typicality of cases described in the thought experiment. In other words, two factors predict the accuracy of token thought experiment intuitions. First, a thought experiment is more likely to elicit accurate intuitions if the case described is typical. That is, if intuitors can be realistically expected to have come across cases in the past that are comparable to the case described in a thought experiment, then that thought experiment has the capacity to elicit accurate intuitions by triggering a robustly trained IGS. Second, a thought experiment that describes events which occur in more stable environments is likely to elicit more accurate

intuitions than one that describes events which occur in highly unstable environments. For instance, cases involving the identification of species of trees are much more likely to produce reliable intuitive judgments than are cases involving the identification of moral action, owing to higher consistency in the taxonomy of trees in comparison to right versus wrong action.

But can descriptive thought experiment intuitions provide any genuine new knowledge of the external world, especially when it comes to category memberships that are not solely determined subjectively? That is, granting that we can glean, piecemeal, the contours of our psychological categories through considering edge cases in thought experiments, it still seems important to ask the question of what, if anything, could be learned through the case method about the external world. Michael Strevens (2019) has recently defended the case method and argued that intuitive judgments about cases come from ordinary beliefs that we form inductively through our interactions with the environment. Thus, according to Strevens, “fresh knowledge about the ultimate basis of category membership” is within grasp through the case method (138).

While I’m skeptical of Strevens’ contention that intuitions can lead us to know the objective structure of the world or natural kind membership, I think he’s right that since categories are inductively derived – that is, representations of categories emerge as we experience the world – thought experiment intuitions provide some data to work with for unveiling concepts. This is an important task, since some of our conceptual commitments are introspectively opaque. When first-person direct reports of our conceptual commitments, or the content of our concepts – whether they correspond to natural kind categories or not – generating intuitions by the method of cases helps to uncover our commitments and our positions in the world which have afforded us certain sets of experiences rather than others.

7 CONCLUSIONS

In this thesis, I discussed the nature of and mechanism underlying intuitions of identification. I argued that many identification intuitions, such as the one that helps you identify the authorship of a painting you are seeing for the first time, fall under the class of experience-based intuitions. On my view, experience-based identification intuitions are produced by domain-general learning systems of hierarchical abstraction that are capable of extracting similarities and forming categories from experience. I suggested that deep convolutional neural networks may be a promising model for understanding how such IGSs function.²⁴ Owing to the mechanism of experience-dependent IGSs, the reliability of experience-based intuition X depends on the quality of the experiences that have shaped the IGS which produced X. I further argued that not only are experience-based intuitions no more fallible than other sources of justification, they may be indispensable. Lastly, even if intuitions turn out not to reveal a prior truths or knowledge about natural kind categories when employed in philosophical inquiry, those intuitions that rely on prior experience might be helpful in unveiling conceptual commitments.

²⁴ One particularly intriguing implication of using artificial neural networks to explain how human beings come to have certain intuitions is the possibility that such networks are themselves capable of having intuitions!

REFERENCES

- Bargh, J. A. (2011). Unconscious Thought Theory and Its Discontents: A critique of the critiques. *Social Cognition*, 29(6): 629-647.
- Barrett, H. C. and Kurzban, R. (2006). Modularity in Cognition: Framing the debate. *Psychological Review*, 113(3): 628-647.
- Battaglia, P., Ullman, T., Tenenbaum, J., Sanborn, A., Forbus, K., Gerstenberg, T., and Lagnado, D. (2012). Computational Models of Intuitive Physics. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 34(34): 32-33.
- Bealer, G. (1998). Intuition and the Autonomy of Philosophy. *Rethinking Intuitions: The psychology of intuition and its role in philosophical inquiry*. Eds. M. DePaul and W. Ramsey. Rowman & Littlefield, 201-239.
- Bengson, J. (2015). The Intellectual Given. *Mind*, 124(495): 707-760.
- BonJour, L. (1998). *In Defense of Pure Reason*. Cambridge University Press.
- Boyer, P. (2001). *Religion Explained: The Evolutionary Origins of Religious Thought*. Basic Books.
- Buckner, C. (2018). Empiricism without Magic: Transformational abstraction in deep convolutional neural networks. *Synthese*, 12: 1-34.
- Cokely, E. T. and Feltz, A. (2014). Expert Intuition. *Rational Intuition: Philosophical Roots, Scientific Investigations*. Eds. Lisa M. Osbeck and Barbara S. Held. Cambridge University Press.
- Coley, J. D., Arenson, M., Xu, Y., and Tanner, K.D. (2017). Intuitive Biological Thought: Developmental changes and effects of biology education in late adolescence. *Cognitive Psychology*, 92: 1-21.

- Cork, L. L. (2014). Nursing Intuition as an Assessment Tool in Predicting Severity of Injury in Trauma Patients. *Journal of Trauma Nursing*, 21(5): 244-252.
- Greene, J. D. (2017). The Rat-A-Gorical Imperative: Moral intuition and the limits of affective learning. *Cognition*, 167: 66-77.
- Greene, J. D. (2014). Beyond Point-and-Shoot Morality: Why cognitive (neuro)science matters for ethics. *Ethics*, 124(4): 695-726.
- Gigerenzer, G. (2008). *Gut Feelings: The Intelligence of the Unconscious*. Penguin Books.
- Goldman, A. (2007). Philosophical intuitions: Their target, their source, and their epistemic status. *Grazer Philosophische Studien*. 74: 1-26.
- Goldman, A. (2006). Immediate Justification and Process Reliabilism. *Epistemology: New Essays*. Ed. Quentin Smith. Oxford University Press.
- Gopnik, A. and Schwitzgebel, E. (1998). Whose Concepts Are They, Anyway? The role of philosophical intuition in empirical psychology. *Intuitions: The psychology of intuition and its role in philosophical inquiry*. Eds. M. DePaul and W. Ramsey. Rowman & Littlefield, 76-91.
- Ichikawa, J. J. (2014). Intuition in Contemporary Philosophy. *Rational Intuition: Philosophical Roots, Scientific Investigations*. Eds. Lisa M. Osbeck and Barbara S. Held. Cambridge University Press.
- Kahneman, D. and Klein, G. (2009). Conditions for Intuitive Expertise: A failure to disagree. *American Psychologist*, 64(6): 515-526.
- Koriat, A. (2008). Subjective Confidence in One's Answers: The consensuality principle. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(4): 945-959.
- Kornblith, H. (2002). *Knowledge and Its Place in Nature*. Oxford University Press.

- Mach, E. (1883/1960). On Thought Experiments. *Knowledge and Error*, sixth edition. Trans. T. McCormack and P. Foulkes. Reidel, 134-147.
- Machery, E. (2017). *Philosophy Within Its Proper Bounds*. Oxford University Press.
- McCauley, R. N. (2011). *Why Religion is Natural and Science is Not*. Oxford University Press.
- McClelland, J. L., McNaughton, B. L., and O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3): 419-457.
- McGahhey, M. and Van Leeuwen, N. (2018). Interpreting Intuitions. *Third-Person Self-Knowledge, Self-Interpretation, and Narrative*. Eds. Julie Kirsch and Patrizia Pedrini. Springer Verlag.
- Nado, J. (2014). Why Intuition? *Philosophy and Phenomenological Research*. 86(1): 15-41.
- Nagel, J. (2012). Intuitions and Experiments: A defense of the case method in epistemology. *Philosophy and Phenomenological Research*, 85(3): 495-527.
- Nagel, J. (2007). Epistemic Intuitions. *Philosophy Compass*, 2(6) 792-819.
- O'Loughlin, C. and Thagard, P. (2000). Autism and Coherence: A computational model. *Mind & Language*, 15(4): 375-392.
- Ohlsson, S. (2011). *Deep Learning: How the mind overrides experience*. Oxford University Press.
- Oliver, N., Rosario, B. and Pentland, A. (2000). A Bayesian Computer Vision System for Modeling Human Interactions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8): 831-843.
- Osbeck, L. M. and Held, B. S. (2014). Introduction. *Rational Intuition: Philosophical Roots*,

- Scientific Investigations*. Eds. Lisa M. Osbeck and Barbara S. Held. Cambridge University Press.
- Panaccio, C. (2014). Ockham: Intuition and Knowledge. *Rational Intuition: Philosophical Roots, Scientific Investigations*. Eds. Lisa M. Osbeck and Barbara S. Held. Cambridge University Press.
- Patel, A., Nguyen, T. and Baraniuk R. G. (2015). *A Probabilistic Theory of Deep Learning*.
- Railton, P. (2016). Intuitive Guidance: Emotion, information, and experience. *Homo Prospectus*. Eds. Martin E.P. Seligman, Peter Railton, Roy F. Baumeister, and Chandra Sripada. Oxford University Press.
- Railton, P. (2014). The Affective Dog and Its Rational Tale: Intuition and attunement. *Ethics*, 124:813-859.
- Sripada, C. (2016). Imaginative Guidance: A mind forever wandering. *Homo Prospectus*. Eds. Martin E.P. Seligman, Peter Railton, Roy F. Baumeister, and Chandra Sripada. Oxford University Press.
- Strevens, M. (2019). *Thinking off Your Feet: How empirical psychology vindicates armchair philosophy*. The Belknap Press of Harvard University Press.
- Tenenbaum, J. B. (1999). Bayesian Modeling of Human Concept Learning. *Advances in Neural Information Processing Systems*. Eds. M.S. Kearns, S.A. Solla and D.A. Cohn. MIT Press.