Research Statement

Christopher Prodoehl

My research interests are in the philosophy of mind and action, aesthetics, and Kant. I am interested in rational agency, including mental agency, as well as self-consciousness and the epistemology of self-attribution. In each area, I have focused on questions related to creativity. Is having a creative idea an exercise of mental agency? Is acting on the desire to create something an exercise of rational agency? Are intuitions about the direction creative work should take a type of epistemic attitude, e.g., a species of metacognition, or something else? My most recent work has focused on the aesthetics of digital art, and especially digital art that makes use of physics simulations. I also have work in progress on whether generative deep learning models can themselves be creative; I argue there is reason to think they can be. I describe some of this recent work at the end.

My largest ongoing project is a philosophical examination of what Graham Wallas called the "illumination" phase of the creative process.[1] During this phase, the creating person becomes conscious of a new idea that seems promising or worth exploring, and it seems this way to her before she has consciously evaluated the idea. Philosophers have not had very much to say about creative illumination. When they do discuss it, it is often treated as a curiosity, e.g., as involving "a mysterious process whereby ideas simply 'pop' into someone's mind."[2] It is an important feature of the creative process, however, and there are various philosophically interesting questions one can ask about it.

Consider someone who experiences creative illumination, then. They have an idea for something they have never thought of before. It may be an idea for a painting or a theme for a piece of music, or it may be a proposition to figure in a philosophical argument or mathematical proof. In cases like these, the person not only has the idea; they are also motivated to pursue it; and they are motivated to pursue it because they find it promising or worth exploring, and typically, as I noted, before consciously evaluating it in any way.

My work articulates and defends a set of claims about these distinguishable parts of creative illumination. The mental events of having the ideas are not exercises of mental agency. Acting on the motivation to explore the ideas is not an exercise of rational agency. The self-consciousness one has during the mental event is different from the kind one has while, for example, thinking through a simple deductive argument. In other words, the "self" of which one is conscious when having creative ideas is not quite the self of which one is conscious in everyday thinking and reasoning.

Taken together, these claims form what I tentatively call a *selfless* account of creative illumination. At present, the account takes shape over several papers; my plan is eventually to incorporate them into a book.

---

[1] Wallas, Graham. (1926). *The Art of Thought*. J. Cape: p. 86.
[2] Livingston, Paisley. (2005). *Art and Intention*. Oxford University Press: p. 31.

"Creative Feeling" develops ideas at the heart of the account. Creative people in many different disciplines have experiences in which some idea seems promising or worth pursuing, but before any conscious evaluation of the idea has taken place. I call these "creative intuitions." For example, in "Mathematical Creation" Henri Poincaré describes having mathematical ideas and feeling completely certain they are right and can be proved, before attempting to verify them by giving a proof.

One thing that creative intuitions seem to have in common is *feeling* of some kind. Poincaré mentions a feeling of certainty, but there are others: feelings of knowing, of urgency, of fascination with some idea. Call these "creative feelings." What are they? One might think that they are *epistemic feelings*, also sometimes called "metacognitive" feelings, examples of which include feelings of knowing, familiarity, and learning. In "Creative Feeling," however, I argue this isn't right. Many creative feelings are *aesthetic*, i.e., of the same kind as the feeling one may have about an elegant piece of music.

The crux of the paper is its claim about what this aesthetic feeling is a response to. When someone has a creative intuition, I argue, they have an aesthetic feeling in response to the *mental activity* in which they are engaged—something like its form—in addition to whatever response they may have to the contents of particular ideas. By virtue of this response, they ascribe an aesthetic property to their mental activity. One implication of this claim is that the experience of having a new idea and finding it worth pursuing is an intrinsically rewarding one. Another is that the mental activity of having the new idea satisfies norms of aesthetic evaluation; put another way, the mental activity is, aesthetically speaking, as it ought to be. This means that creative mental activity is characterized not only by the (causal) psychological processes that generate new ideas, but also by the normative (aesthetic) properties those processes can have.

The argument in "Creative Feeling" concerns creativity in any domain in which creativity is possible, including the arts, mathematics, and science. The papers I now describe are focused specifically on artistic creativity.

Having a creative intuition is motivating. More precisely, to have a creative intuition is, in part, to become motivated to pursue some idea. But what exactly is the connection between the intuition and the motivation? How does creative intuition motivate? I pursue this question in "Creative Motivation."

"Creative Motivation" develops one of the core conclusions from "Creative Feeling." According to "Creative Feeling," when artists have a creative intuition, they ascribe aesthetic properties to their mental activity. In doing so, they take aesthetic pleasure in that mental activity. In "Creative Motivation," I argue that this pleasure is *disinterested* in the sense Kant explains in the *Critique of the Power of Judgment*, and, as Kant says there, such pleasures can be motivating. This means, I argue, that when an artist takes disinterested pleasure in her mental activity, she is motivated by that pleasure to engage in creative activity. She has what I call a "creative motive." I argue that acting on creative motives does not amount to an exercise of rational agency. When an artist acts on a creative motive, she remains passive with respect to that motive.

In "Aesthetic Insight and Mental Agency" (*The British Journal of Aesthetics*, 2023), I turn from rational agency, broadly construed, to rational *mental* agency. I argue that artists do not exercise mental agency over their creative intuitions, or as I call them in that paper, "aesthetic insights." I argue that an artist's having an aesthetic insight is not something she does for a reason. The argument turns on general principles about responding to reasons, as well as principles about the particularity of aesthetic evaluation—that is, the absence of general rules for determining whether it is right to ascribe an aesthetic property to something.

One of the quintessential observations about creative experience, and creative illumination, in particular, is that artists sometimes experience their ideas as coming from some source or agency other than themselves. The Muses, for example, were a very early thought about what that source might be. In my view, observations of this kind are best understood, not as evidence of supernatural intervention, but rather as reporting a distinctive kind of self-consciousness. I pursue this idea in "Aesthetic Ideas and Self-Consciousness" (in *Kant's Lasting Legacy: Essays in Honor of Beatrice Longuenesse*, 2025), and "Inspiration and Self-Consciousness."

"Aesthetic Ideas and Self-Consciousness" explores the relationship between what Kant says about *aesthetic ideas* in the *Critique of the Power of Judgment*, and what Kant says about self-consciousness in the *Critique of Pure Reason*. In the latter work, Kant takes it as a datum that we are conscious of ourselves as one and the same subject of different representations. I imagine *this*, think *that*, and remember *something else*, all while experiencing myself as one single thing imagining, thinking, and remembering. According to Kant, this consciousness of ourselves depends, roughly speaking, on a mental activity of combining representations together in distinctive ways.

Aesthetic ideas, according to Kant, are a distinctive kind of mental content that the artistic genius is able to have. They guide her creative work, and, Kant's text suggests, having them involves a mental activity of combining representations together. How does this mental activity of combining representations relate to the one above, which is required for being conscious of ourselves as a single subject of different representations? Kant doesn't discuss the issue, so the paper aims for a rational reconstruction that is consistent with Kant's text. It argues that, while having an aesthetic idea involves combining representations, and so does involve self-consciousness, it is not exactly the same kind of self-consciousness that Kant describes in the first *Critique*. Instead, when an artistic genius has an aesthetic idea, she enjoys a distinctive kind of *aesthetic* self-consciousness.

"Inspiration and Self-Consciousness" pursues closely related ideas. It is also rooted in Kant's discussion of self-consciousness in the *Critique of Pure Reason*, but it focuses more on the creative experiences that artists have described than the paper above does. In "Inspiration and Self-Consciousness," I argue again that the mental activity that takes place during creative illumination is not structured in the way Kant thinks is required for self-consciousness of ourselves as a single subject of different representations. Nonetheless, that mental activity does have enough structure to support a distinct and distinctive kind of

self-consciousness, which I call "alterior self-consciousness." This type of self-consciousness specifies the content of the artist's experience that her ideas come from some source or agency other than herself.

I have other papers on closely related issues. In "Lucky Artists" (*Analytic Philosophy*, 2023), I argue that artists, and performers especially, exercise a distinctively non-practical type of control over their bodily movements; I call this "receptive control." For an example, think of a pianist who doesn't need to carefully monitor her bodily activity because of her training and skill. I argue that, to affect fine-grained, aesthetically-relevant details of her performance, such as tone and dynamics, she listens to the results of her playing, and she imagines the way she wants her performance to sound. Neither listening nor imagining in this way, I argue, involves practical, control-oriented attitudes toward her playing. Neither is an intention to play a certain sequence of notes, for example, or even an intention to play the passage in some particular way. Each nonetheless recruits bodily activity, and so counts as a kind of control. Both listening and imagining in this way are, I argue, instances of exercising receptive control.

"Personal and Impersonal Expression" (*The British Journal of Aesthetics*, 2025) examines the role of self-expression in creating art. While some artists encourage self-expression in the creative process, others caution against it. What are they disagreeing about? What alternative do the cautious have in mind? After critically discussing some potential answers, I argue that the artists who caution against self-expression are recommending a different way—an *impersonal* way—of expressing feeling in the process of creating art. Expressing feeling in this impersonal way does not lead to self-knowledge in the same way that self-expression does.

In future work, I plan to develop these lines of argument into an account of creativity in all its phases. Much of the important groundwork is in place. The bodily activity by which artists create works of art is an exercise of receptive control, not intentional control. Acting on the motivation produced by creative illumination is not an exercise of rational agency. Events of creative illumination themselves are not exercises of rational *mental* agency, but they do involve a distinctive kind of self-consciousness. But what about a positive characterization? What is creative illumination if it is none of the above? Is it essentially passive, akin to perception? Maybe a species of passive imagination? These are questions for future work.

My most recent work has focused on the aesthetics of computation and the question of whether generative deep learning models can be creative. In "Computer Animation and the Aesthetics of Simulations" (presented at the American Society for Aesthetics Annual Meeting, 2024), I explore the aesthetics of computer animation that makes uses of physics simulations. Think, for example, of the Disney film *Moana*, which uses fluid simulations to animate its oceans. I argue that the correct way to aesthetically appreciate such work is *as a simulation*, or as computational—rather than, for example, simply as an artistic representation of nature, akin to landscape painting (or hand-drawn animation of an ocean). I work through a few basic suggestions about what this means in the paper. For instance,

part of what it is to appreciate some work (or some aspect of a work) as simulated is to appreciate it as distinctively independent of the attitudes of the person who made it, something such work shares with photography. The basic argument in this paper will generalize, I hope, to images and other artifacts produced by, e.g., deep learning models, and may suggest ways of thinking more generally about aesthetic properties related to computational processes.

"AI, Creative Generation, and Being Creative" focuses on whether generative deep learning models, such as large language models, are capable of *creative conception*. Can an LLM that outputs a witty couplet, for instance, be said to have produced that output creatively? There are reasons to think not. Among them is the fact that LLMs are trained to learn statistical features of their training data. The process of producing statistically appropriate next tokens, one might think, is not a creative process. I argue, however, that this and other considerations are not decisive. I do so by drawing on recent work on *mechanistic interpretability*, which attempts to explain the output of deep neural networks (like LLMs) by identifying networks of nodes—sometimes called "circuits"—responsible for producing that output. By drawing on this work, I argue there is a distinction between standard and non-standard routes by which a model can generate its output. This distinction, together with other considerations, suggests it is possible for deep learning models to produce output in both routine and creative ways. It is possible, then, for deep learning models to do something akin to creative conception.