

# Open Research Online

---

The Open University's repository of research publications and other research outputs

## The expressive stance: intentionality, expression, and machine art

### Journal Item

How to cite:

Linson, Adam (2013). The expressive stance: intentionality, expression, and machine art. *International Journal of Machine Consciousness*, 5(2) pp. 195–216.

For guidance on citations see [FAQs](#).

© 2013 World Scientific Publishing Company

Version: Accepted Manuscript

Link(s) to article on publisher's website:  
<http://dx.doi.org/doi:10.1142/S1793843013500066>

---

Copyright and Moral Rights for the articles on this site are retained by the individual authors and/or other copyright owners. For more information on Open Research Online's data [policy](#) on reuse of materials please consult the policies page.

---

[oro.open.ac.uk](http://oro.open.ac.uk)

## THE EXPRESSIVE STANCE: INTENTIONALITY, EXPRESSION, AND MACHINE ART

ADAM LINSON

*Faculty of Mathematics, Computing and Technology,  
Department of Computing, Open University,  
Milton Keynes, MK7 6AA, United Kingdom  
adam.linson@open.ac.uk*

This paper proposes a new interpretive stance for interpreting artistic works and performances that is relevant to artificial intelligence research but also has broader implications. Termed the *expressive stance*, this stance makes intelligible a critical distinction between present-day machine art and human art, but allows for the possibility that future machine art could find a place alongside our own. The expressive stance is elaborated as a response to Daniel Dennett’s notion of the *intentional stance*, which is critically examined with respect to his specialized concept of rationality. The paper also shows that temporal scale implicitly serves to select between different modes of explanation in prominent theories of intentionality. It also considers the implications of the phenomenological background for systems that produce art.

*Keywords:* Intentionality; expression; machine art; machine musicianship; improvisation; time; temporality; aesthetics; phenomenology.

### 1. Introduction

This paper explores the issues of intentionality and expression, as they pertain to an understanding of artistic production — for humans as well as for machines. The relationship between human art and machine-generated art will be explored as an analogue to the relationship between human and machine intentionality. In this context, I will consider the epistemological premise of Daniel Dennett’s intentional stance, the notion that, from a particular interpretive perspective, the question of what is “inside the head” of an agent is irrelevant to a pragmatic account of intentionality. I will then show how Dennett’s notion of rationality, central to his theory of intentionality, is inadequate to account for a certain kind of meaning associated with artistic expression. I will, however, preserve his epistemological premise, according to which I will develop a theory that I will call the *expressive stance*.

Chella & Manzotti [2009] present an overview of some of the philosophical issues pertaining to machine consciousness, including intentionality. One topic not covered in their overview is the topic of art (cf. Chella & Manzotti [2012]). Rather than focus on the broader issue of how a robot might experience art, I will focus on the topic of artistic production by machines. This discussion, however, relates to

a broader discussion concerning the interpretation of art, and several interrelated issues, namely, rationality, the phenomenological background, and temporal scale. These issues will be discussed in order to illustrate the significance of the contrast between Dennett’s account of intentionality and my related account of expression.

The concept of temporal scale, I will argue, is implicit in several prominent philosophical theories of intentionality. Another component common to these theories is the idea of a community of shared meaning. I will examine the relation between temporal scale and community to reveal an important distinction between two dimensions of meaning: an intentional meaning dependent upon shared rationality, and an expressive meaning dependent upon shared experience.

It will be necessary to make a few clarifications, to avoid some potential confusions. The idea of an artist’s (or author’s) intentions — typically assumed to be “inside the head” — and, more specifically, the idea of what an artist intends to express with an artwork, seem to comprise the colloquial senses of *intention* and *expression*. The colloquial versions of these terms are only marginal to the present discussion, which is primarily concerned with their more specialized philosophical counterparts.

Another important point to note in advance is that, while I will not discuss the issue of language at length, I will assume that there are non-linguistic modes of understanding, in everyday life as well as in the experience of art; however, I will also assume that language, a historically fundamental part of human society, is crucial for the elaboration and understanding of certain forms of complex meaning. Indeed, even if we can imagine communities of shared meaning facilitated by non-linguistic practices, some central features of human society, such as literature and law, are inconceivable without language. The implications of this position on language will become clear in my examples, which include painting, improvised music, and installation art. The next section will present a brief account of the relation of language to non-linguistic art before moving on to Dennett’s theory of intentionality.

## 2. Non-linguistic art and the role of language

When art is considered historically, it reveals that all art forms, even non-linguistic ones such as music and painting, are dependent on lived human history, which is thoroughly mediated by language. In Michael Wheeler’s [1996] cognitive account of aesthetic experience, he criticizes Gadamer on the point that “the traditions and heritage which shape understanding are thought to be embedded in language. Skills, crafts, and *works of visual art* — things which are not in themselves linguistic — are held to be intelligible only in so far as they can be brought to articulation in language” [p. 212, original emphasis]. Wheeler’s critical response to Gadamer is to reject language as being at all relevant to certain kinds of aesthetic experience, including gazing at a Rothko painting or “our experience of jazz”, which need not be brought “to articulation in words” [p. 233]. But Wheeler misses something important when he reifies aesthetic experience in this way.

Another possible response to Wheeler’s account of Gadamer’s position would be to say that the traditions and heritage which shape our understanding are part of a sociocultural human history that *includes* language and indeed could not be imagined without it. Language need not be ranked as primary for immediate experience, but it nevertheless permeates the history and development of humanity. It should be clear that even non-linguistic human activities and artifacts, certainly including all works of visual art *as we know them*, could not have surfaced in a human history devoid of language.

Wheeler essentially negates his own point about his non-linguistic aesthetic experience of Rothko’s paintings by relating them to those of Caspar David Friedrich, and thinking about how this canonical context affects his experience. What he leaves out is that the entire construction of a canon of painting that makes such a context possible is built upon an edifice of language that includes written histories and analyses of painting, verbal dialogue among painters and critics, art history and education, training, and so on. A notion of purely abstract sensorial experience does not hold up here, since that would dissolve the context of the history of painting that Wheeler invokes as the basis of his aesthetic experience.

To take another of Wheeler’s examples, Charlie Parker’s saxophone improvisations, it must be pointed out that these improvisations do not simply come out of nowhere; they are made possible by Parker’s personal life experience.<sup>a</sup> This life experience, and the social history of the world into which he was born, were crucially shaped by language. In a manner of speaking, without language, there could not be any Charlie Parker improvisations, nor any audience experience of them.<sup>b</sup>

From the perspective of our role as listeners or viewers, it is only by living in a shared world — itself crucially shaped by language — that makes an aesthetic experience what it is. This remains true even without the immediate intervention of language in the moment. Although Wheeler suggests that a hermeneutics “purged of linguophilia” could yield a non-linguistic phenomenology that would be more harmonious with natural science, it would also lead to an unfortunate reification of aesthetic experience. A version of naturalized aesthetics such as his renders inert many of the socially valued capacities of art. This would ultimately diminish the difference between human art and arbitrary machine output, a difference that will be further explored below.

Another way of thinking about the relation between life experience, language, and art is presented in George Lewis’ [1996] account of the “hermeneutics of improvisative music”. Linking the musical to the “extramusical”, he points out that “improvisative musical utterance, like any music, may be interpreted with reference to historical and cultural contexts” [p. 93]. Lewis traces not only different

<sup>a</sup>Lewis [1996] includes this quote from Parker: “If you don’t live it, it won’t come out of your horn” [p. 119].

<sup>b</sup>While this claim is meant in a general sense, it more clearly holds true in a narrower sense, as Parker at times makes musical references imbued with additional layers of meaning that relate to song titles and linguistically conveyed stories.

approaches to improvisative performance, but also critically examines the nuanced ways in which historical and cultural contexts have shaped *writing* about music. He argues that such writing has served to elaborate and solidify ideological positions that in turn inform musical production and reception.

Lewis [1996] refers to an “historically emergent rather than ethnically essential” system of improvisative musicality that he terms “Afrological” [p. 93]. This “logic” refers “metaphorically to musical belief systems and behavior” that evolved in the particular cultural context of African-American traditions of improvisation, with the music of Charlie Parker considered as a key example. A cultural context can also be understood as a community of shared meaning, which provides a reference point for the interpretation of intentionality. This notion of community will be taken up further below in relation to Dennett’s notion of “shared intuitions” [Dennett, 1987, p. 98].

### 3. Dennett’s Intentional Stance

The categories of “belief systems and behavior” have a strong relation to questions of intentionality. Elsewhere, Lewis [1999, 2000] explores his experience with a machine improviser that he developed for human–computer interaction. Recently, I have also undertaken work in this area, having developed my own version of a machine improviser [Linson *et al.*, 2013a]. In empirical studies conducted with this interactive system [Linson *et al.*, 2013b], it was found that a number of participants who performed with the system were inclined to call the machine “he” or “she” rather than “it”, despite their full knowledge that they were interacting with a machine. This situation closely resembles an example given by Dennett, where he found that “loggers in the Maine woods [...] invariably call a tree not ‘it’ but ‘he’, and will say of a young spruce, ‘He wants to spread his limbs, but don’t let him; then he’ll have to stretch up to get his light’” [Dennett, 1981, p. 272]. Dennett writes that it was this experience that led him to consider belief systems and behavior in terms of what he calls the *intentional stance*, and the corresponding notion of *intentional systems*, which together form his theory of intentionality.

Dennett offers a “brief résumé” of his theory of intentionality:

An intentional system is a system whose behavior can be (at least sometimes) explained and predicted by relying on ascriptions to the system of *beliefs* and *desires* (and other intentionally characterized features — what I will call *intentions* here, meaning to include hopes, fears, intentions, perceptions, expectations, etc.). There may *in every case* be other ways of predicting and explaining the behavior of an intentional system — for instance, mechanistic or physical ways — but the intentional stance may be the handiest or most effective or in any case *a* successful stance to adopt, which suffices for the object to be an intentional system [Dennett, 1981, p. 271, original emphasis].<sup>c</sup>

<sup>c</sup>By his own account [Dennett, 1987, p. 3], the “flagship” expression of his theory can be found in

In other words, as long as it makes sense to interpret the behavior of some entity (i.e., a system) as intentional — that is, to take the intentional stance toward it — then it is an intentional system. As he stresses, it is important to bear in mind “how bland this definition of *intentional system* is, and how correspondingly large the class of intentional systems can be” [Dennett, 1981, p. 272].

One of the more drastic epistemological implications of Dennett’s argument is that there is no way of distinguishing how we draw conclusions about human intentionality from how we draw similar conclusions about entities that we recognize as fundamentally different from humans, such as machines, frogs, plants, etc. He handles this potentially counterintuitive implication by maintaining that “even when intentional glosses on (e.g.) tree-activities are of vanishingly small heuristic value, it seems to me wiser to grant that such a tree is a very degenerate, uninteresting, negligible intentional system than to attempt to draw a line above which intentional interpretations are ‘objectively true’” [Dennett, 1981, p. 272]. In other words, if our only means of judging whether anyone or anything exhibits intentional behavior is by external observation, then there is nothing to justify the claim that human intentional systems should count as a difference of kind rather than of degree.

Reflecting on my artificial musical improviser, my performances with it, and my (far more common) collaborative improvisations with other human musicians, I began to consider the fact that with improvised music in general, something like the intentional stance may be brought to bear on the musical situation, and other psychological or physical explanations may serve a role as well. However, there is another dimension of meaning that remains invisible to these interpretive stances. To reveal this dimension of meaning that is relevant to the understanding of art, I will argue, we need the *expressive stance*. In my presentation of the expressive stance below, I will explain how it works, why it is relevant to other art forms as well, and why it is important.

Before explaining the expressive stance, I will first give a more thorough account of several versions of intentionality. The following section will clarify how the question of what is “inside the head” functions in different theories of intentionality. The subsequent three sections will consider intentionality in terms of temporal scale, rationality, and the phenomenological background. These topics will form the basis for my explanation of the expressive stance.

#### 4. Internal and External Explanatory Appeals

Dennett [1991] criticizes John Searle’s [1980] well-known Chinese Room thought experiment, which suggests that a certain kind of biological brain (“neurons with axons and dendrites, and all the rest of it” [Searle, 1980, p. 422]) is presently the only known host for a mind, conceived of as the seat of intentionality. Searle’s argument

“True Believers” [Dennett, 1987, chap. 2].

is specifically aimed against the claim that any computer could have a mind.<sup>d</sup> Searle uses the Chinese Room thought experiment (which I will not recount here) to explain that, for him, it would not be significant if a computer could pass the Turing test, as long as the computer differed internally from a human. For a computer to pass the Turing test (in its most generous formulation), it would have to be indistinguishable from other humans when judged solely on the basis of interaction and observation. Dennett responds with the Systems Reply, “the standard reply of people in AI from the earliest outings” of Searle’s thought experiment [Dennett, 1991, p. 439]. In short, this reply, initially presented by Searle [1980] in a series of encountered critical responses, says that an imagined computer system with tremendous internal complexity could be functionally equivalent to the complex system(s) internal to human beings.

Searle ultimately disparages the entire series of replies that he presents as being part of an “unashamedly behavioristic” tradition. This tradition follows from Turing’s [1950] point that (external) interaction and observation, rather than mental or physical innards, are the only relevant sources of information in everyday life for our judgments about others, whether human or computer. While Dennett’s intentional stance arguments seem to embrace this “traditional” position, his defense of the Systems Reply is more concerned with imagining the system’s internal construction.

Dennett gives the example of a highly elaborate conversation between a human and a computer to shore up the defense of his response to Searle [Dennett, 1991, pp. 436–437]. Dennett’s point is merely that we could, in principle, *imagine* such a machine, which is enough to combat Searle’s “misdirection”. What is puzzling here is that Dennett’s choice of the Systems Reply invites us to address what specifically is *inside* the computer. Yet, in his example of a “winning” Turing test dialogue, the implication is that we must take an interpretive stance toward the system on the basis external observations and interactions — not unlike the conditions that merit the intentional stance. Here, curiously, for Dennett, drawing conclusions about how our inferences relate to the internal workings of the machine is justifiable.

To be fair, Dennett’s point in this instance is to make a case for why the complexity of human consciousness is, in principle, intelligible as a collection of simpler mechanisms. It is notable, however, that he departs here from the message of another one of his arguments, about *seeing* and visual experience [Dennett, 1987, pp. 106–107] (see also [Dennett, 1991, pp. 338–344]). His point about seeing concerns the fact that we do not have access to the first-person experience of others, but nevertheless, we can make inferences according to interactions with and observations of

<sup>d</sup>More specifically, his argument is aimed against the view that “mental processes are computational processes over formally defined elements” [Searle, 1980, p. 422], which is a restatement of the view that intelligence amounts to “generating and progressively modifying symbol structures until it produces a solution structure” [Newell & Simon, 1976, p. 120]. Thus formulated, I also oppose this view (see [Linson *et al.*, 2012]).

others. In his view — and I am in agreement on this point — there is no way to answer the question about the “essential intrinsic properties of real seeing” [Dennett, 1987, p. 107]. We are faced only with a “tactical and lexical decision about whether it would be misleading” to call certain phenomena “seeing”, such as whether or not a blind person with a prosthetic vision device can “see”. He notes that even though, for example, we have biological evidence that frog vision functions differently from ordinary human vision, we know enough about frog vision to accept that frogs can see; what matters to us about whether or not a frog can see ultimately relates to frog behavior, interpreted through external observation and interaction — for instance, in our assessment that a frog can see that it is being chased and attempt to escape.

Dennett makes a slightly different point regarding intentionality when he argues that we can interpret a chased frog as “*wanting* to escape”, without the need to address the question of how this “want” (or belief or desire) could be “inside the head” of the frog. It seems that, for Dennett, sometimes natural scientific evidence is relevant for interpreting behavioral functions, but such evidence is never relevant for interpreting intentionality. His assumption that consciousness must result from a collection of scientifically identified mechanisms is similar to how he treats seeing, but opposed to how he treats intentionality. This might be considered a tension in his theory.

Dennett [1991] responds to this tension, in a slightly different formulation. Namely, he responds to the charge that the perspective of the intentional stance assumes a *unified* agent, while his theory of consciousness “opposes this central unity” [p. 458]. He argues that there are important theoretical differences that apply to what you are looking at, depending on “how far away you are”. That is, how we should interpret patterns hinges upon what particular patterns we are interpreting, and different perspectives allow us to perceive different patterns. Certain activity patterns can be made sense of by treating their source as a single agent. But, if we are not careful, this treatment can also mistakenly lead us to use the same mode of interpretation for other activity patterns for which there are better modes of interpretation, such as how we understand conscious experience (for a related discussion, see [Polanyi, 1962, chap. 13]).

## 5. Temporal Scale of Interpretation

The question of “how far away you are” is of interest here, and I agree that different explanations (or modes of explanation) are relevant in different circumstances. Yet, while his metaphor of *spatial* distance seems to evoke the right understanding of this point, I find the discussion should be reframed according to what I will call *temporal scale*, given that observation is a fundamentally temporally extended phenomenon. What I mean by temporal scale is not unrelated to Dennett’s notion of scope, but different issues surface when we frame the discussion specifically in terms of time. For example, for the intentional stance to be effective toward a frog that is seen

leaping away, our inferences about the frog’s behavior must be linked to the time scale of what we take to be its salient actions — in this case, leaping.

Such a leap can be, in principle, temporally segmented into much shorter units of time. But, if the segments are too short, they would not be intelligible as frog behavior *per se* (though perhaps as cellular behavior, subatomic behavior, etc.). For the intentional stance to be effective for the frog as an organism, our inferences must be guided by what we could call the *mesoscale* — the time scale related to intelligible, apparently intentional behavior. I wish to introduce a notion of relative temporal scale here, for which the mesoscale would relate to intentional behavior, and for which any further subdivision of activity would be thought of as the *microscale*. The *macroscale*, on this model, would be the life history of the agent. If the agent is biological, the history of the species may also be relevant; we will return to this topic later. (The Latin scale prefixes are merely suggestive, such that each scale has an appropriately flexible span; the prefixes do not necessarily refer to seconds, nor are they meant to imply a picoscale, nanoscale, etc.)

As we have seen in the previous section, with respect to the intentional stance, Dennett has shifted the discussion from the ontological question of “intrinsic” intentionality with which Searle is concerned to one of how we interpret intentionality. Although I will continue to pursue the interpretive line, an ontological critique of Searle does have some relevance to my thematization of time. In particular, Hubert Dreyfus’ [1993] Heideggerian critique of Searle (and Husserl) is relevant here. Dreyfus takes Searle to task on an issue that I relate to temporal scale, although time is only implicit in his discussion:

Searle points out that an expert skier does not have to form a separate intention to shift his weight from one ski to the other or to execute each turn. He just intends to ski down the mountain. This is a safe response since the intentionalist can, indeed, always find a level at which the actor is trying to achieve something [Dreyfus, 1993, p. 30] (see also [Searle, 1983, pp. 150–153]).

In another example, Dreyfus notes that while a tennis player is always trying to win a point, “what he or she is doing seems to be much more fine-grained”. I find that the point Dreyfus is making amounts to the fact that some actions are *temporally* more fine-grained than, for instance, trying to win a point, which is only relevant at a broader time scale.

Another example that, in my opinion, highlights the contrast between the temporal meso- and microscales of interpreting intentional behavior concerns aircraft pilots. According to Dreyfus, pilots learn to scan their instruments in a particular fixed sequence, one they believe to be following during successful flights. “At one point, however, Air Force psychologists studied the eye movements of the instructors during simulated flight and found, to everyone’s surprise, that [...] their eye movements varied from situation to situation and did not seem to follow any rule at all” [Dreyfus, 1993, p. 29]. He suggests that Searle would respond that scanning

the instruments to maintain the flight is intentional, but that the eye movements themselves are part of a skilled activity that does not need to be represented intentionally (this relates to what Searle calls “Background” [1983, chap. 5], discussed further below). From the perspective of the intentional stance, which makes no appeal to the ontological status of intentionality, we can understand that different time scales of interpretation are at work here: the microscale of eye movements do not lend themselves to intentional interpretation, while the mesoscale of scanning the instruments to maintain the flight readily makes sense according to the intentional stance. The notion of what “makes sense” that arises here will be discussed at length in the next section in terms of rationality.

## 6. Rationality

Part of what makes the intentional stance useful as an interpretive strategy is its pragmatic effectiveness for predicting and explaining behavior. As Dennett states:

For instance, it is a good, indeed the only good, strategy to adopt against a good chess-playing computer. [...] The computer is an intentional system [...] not because it has any particular intrinsic features, and not because it really and truly has beliefs and desires (whatever that would be), but just because it succumbs to a certain *stance* adopted toward it, namely, the intentional stance, [...] the stance that proceeds by considering the computer as a rational practical reasoner [Dennett, 1981, pp. 271–272, original emphasis].

Dennett is careful to give a detailed account of exactly what he means by *rational*, in the sense of the “rational practical reasoner” referred to here to, and elsewhere in his work when he refers more generally to “rational” action: this sense of rationality is not found in the “ideals of Intro Logic” but rather, it is a “pre-theoretical concept of rationality” that can be understood as our “shared intuitions”, i.e., doing what makes sense [Dennett, 1987, p. 98].

He gives a further explanation of this point in another context:

The standard trap is to suppose that the relatively rare cases of conscious practical reasoning are a good model for the rest, the cases in which our intentional actions emerge from processes into which we have no access. Our actions generally satisfy us; we recognize that they are in the main coherent, and that they make appropriate, well-timed contributions to our projects as we understand them. So we safely assume them to be the product of processes that are reliably sensitive to ends and means. That is, they are rational, in *one* sense of that word [Dennett, 1991, p. 252, original emphasis].

There are times, Dennett acknowledges, when we make a consciously deliberated decision and consciously act on it. However, as he points out, such experiences are relatively rare, “and a good thing, too, since there wouldn’t be time” [Dennett, 1991, p. 252]. Here we have a more explicit acknowledgment of the relevance of time, and the fact that what we do “in the moment” (such as a pilot’s microscale

eye movements) could not be consciously and deliberately formulated and carried out.

But, by and large, Dennett argues, we can retrospectively interpret behavior (in order to make sense of it) *as if* it were rationally formulated and carried out. He is careful not to read our reflections after the fact as describing what actually leads to the actions in question. Here, he seems to be in agreement with Dreyfus, who describes the conceptual difficulties that arise from the assumption that “the rules used in the formalization of behavior are *the very same rules* which produce the behavior” [Dreyfus, 1972/1992, p. 190, original emphasis]. Dennett not only *recognizes* this disjunction between formalization and production; his theory of the intentional stance actually *emphasizes* this disjunction: “Intentional system theory is almost literally a black box theory, which makes it *behavioristic* to philosophers like Searle [...] but hardly behavioristic in Skinner’s sense. On the contrary, intentional system theory is an attempt to provide [...] a competence model, in contrast to a performance model” [Dennett, 1987, p. 74, original emphasis].

In other words, an idealized formal description differs from a historical, causal description, and one must be careful to avoid making a category mistake when insights from the former appear relevant to the latter.<sup>e</sup> However, even with an appropriate competence–performance distinction, Dreyfus [1972/1992] raises a further conceptual difficulty in understanding human behavior: there is a critical difference between specialized competencies and the general human competence that makes specialization possible. This difference will be explored in the next section with respect to the notion of the phenomenological background.

## 7. Varieties of Background

Thus far, I have described some of the interrelationships between language, interpretation, rationality, and temporal scale. It is necessary to address one more preliminary topic regarding intentionality before moving on to the expressive stance, namely, the concept of *background*. Searle [1983] defines “Background” as the “non-representational, preintentional capacities” that make possible intentional action [pp. 143–144]; he also subdivides the concept into “deep” and “local” Background (I will capitalize the term as Searle does only when referring to his concept). Deep Background consists of biologically governed capacities, essentially the constraints and facilitating aspects of physical embodiment. Local Background equates to “local cultural practices”, including general socialization, such as opening doors, drinking from bottles, and skiing.

As Dreyfus [1993] points out, however, Searle’s view seems to reduce our sense of being human to having a set of specific skills, which in turn become the only frames of reference for our actions and activities. There is no room in Searle’s model for

<sup>e</sup>For an excellent extended treatment of this theme in relation to behavior and AI, see Hendriks-Jansen [1996].

“everyday coping”, Dreyfus’ gloss on Heidegger’s notion of “being-in-the-world” [Dreyfus, 1991, p. 62] (see also [Heidegger, 1962]). Everyday coping refers to the most basic way of being familiar with the world one experiences, the world one is open to in one’s explorations and traversals of it. This mode of involvement or engagement opposes the traditional dichotomy of an active subject and a passive world of objects. In contrast to Searle’s view, Heidegger’s view has in common with Dennett’s the notion that intentionality “is not something immanent to the subject” [Heidegger, 1982, p. 313] (although this point has a fundamentally different significance in the respective theories of the latter two).

Dreyfus’ Heideggerian critique of AI [1972/1992; 2008], which extends to Searle’s (and Husserl’s) notion of intentionality, in large part relates to the distinction between everyday coping and skilled coping. Skilled coping can be understood as what is taking place in the above example when “an expert skier does not have to form a separate intention to shift his weight from one ski to the other”. For Heidegger, this would be described as someone with extensive experience who is absorbed in a skilled activity. When a specialized activity is considered, such as skiing or chess playing, it is possible to model such activity using computer systems known as *expert systems*, which model expertise in a single domain. These models are based on an “information processing” schema, which, in principle, according to the Heideggerian view, cannot model general everyday coping.

Charles Taylor [1993] gives a concise summary of Heidegger’s critical view of this information processing model, which

offers us the picture of an agent who in perceiving the world takes in “bits” of information from his or her surroundings and then “processes” them in some fashion, in order to emerge with the “picture” of the world he or she has, who then acts on the basis of this picture to fulfill his or her goals, through a “calculus” of means and ends [Taylor, 1993, p. 319].

Without engaging with the ontological question, we can note here that in Dennett’s alternative approach, the means–ends calculus is simply a way of making sense of agents *post hoc*, without claiming this logic has anything to do with what drives an agent’s actions. This point is reinforced when Dennett [2001] states that:

In the wider world of human activity [...] [w]e don’t find the tale of Jane going to the supermarket on her way home from work interesting precisely because it all unfolds so predictably from the intentional stance; today she never encountered any interesting options, given her circumstances. Other times, however, the most rational thing for an agent to do is far from obvious, and may be practically incalculable [p. 306].

In other words, at the temporal mesoscale we can make sense of certain activity by interpreting it as rational, intentional action.

But, arguably, activity such as going to the supermarket is not *only* constrained (and facilitated) by deep and local Background, that is, by biology and the enculturation and socialization that makes such action intelligible. Indeed, even while we

can in this way make sense of a person performing actions, there is a further sense, unaccounted for by Background, that a person’s entire life history manifests itself in a unique extended general engagement with the world. This *individuality* gives a broader context to any specific skilled engagement. That is to say that going to the supermarket certainly makes use of various skills and capacities, but the abstraction that it is simply a performance that fulfills certain conditions of satisfaction is too limited, and risks discarding other dimensions of experience as trivial. This claim will be further explored below.

One final point before moving on: when I introduced the the temporal macroscale, I mentioned that with a biological agent, the history of the species may be a relevant addition to the individual history of the agent. To follow up on this, we can note that humans, and perhaps some other animals, are born into a world that is crucially shaped by historically contingent social and cultural customs of everyday life, in addition to the traditions of specific skilled activities. For example, the practice of sailing, clearly a specialized skilled activity, itself developed in relation to a wide number of much broader everyday life practices at the nexus of sociocultural and material forces (see Law [1987]). It is true that an individual sailor today may learn a set of sailing skills that can be analytically abstracted from the individual’s life history and the sociocultural history of the world. But, ultimately, the totality that made possible the individual’s engagement with sailing includes the human history of relationships to bodies, materials, natural and social forces, and so on. This totality not only extends beyond, but also fundamentally shapes the Background described by Searle.

On the basis of this totality, the life history of any human may be said to be influenced by the whole of human history in a way that does not necessarily hold with respect to the life of a frog, for example. For a human, the macroscale of temporal interpretation must be understood not only in terms of one person’s cumulative perceptions and behaviors, but also in terms of how his or her perceptions and behaviors are relative to an even broader time scale, a “supermacro” time scale. To keep matters conceptually manageable, however, we can understand this broader scale in terms of an individual human life history by keeping in mind the fact that everyday experience — especially when understood as everyday coping — can always be shown to have been shaped by forces that are only intelligible at a broader human time scale. Thus, the macroscale for a human agent, even when confined to a life history, always implies a historical, sociocultural totality.

## 8. The Expressive Stance

My argument thus far has been that the intentional stance, and the accompanying retrospective ascription of rationality, can be effective at the mesoscale of temporality. But, at the same time, adopting this stance means that microscale actions (like those of the pilot’s eye movements) can only be understood as being in the service of (or as analytical decompositions of) mesoscale activities. If we grant that

the microscale eye movements are necessary for the mesoscale activities of scanning the instruments, maintaining the flight, and so on, we have no way to understand the significance of specific microscale actions for which other available and valid microscale actions could be substituted. Even if, at times, one possible microscale action edges out another because it is somehow (in retrospect, at least) better, or *more* rational, we can certainly imagine a case in which, given all the relevant information, two or more options would be *equally* rational.

Dennett [1981] gives an example of just such a case, in choosing between identical soup cans from a supermarket shelf, “all roughly equidistant from your hands” [p. 291]. He is willing to grant that microscale differences may result in “slight variations in timing, style and skeletal implementation”, but notes that these are irrelevant to the intelligibility of the grabbing of a soup can at the mesoscale. In this case, we must either interpret the microscale actions that lead to the grabbing of one particular soup can according to some variety of hard determinism, or simply as random.<sup>f</sup> But, we might hope for more in the case of some performer’s choice of a hauntingly beautiful musical note, played in the moment, without conscious deliberation. Even when Dennett suggests that some choices more important than which soup can to grab may be the result of an intelligent, inventive, creative selection from a swift generate-and-test heuristic process, he maintains that the result of this process must admit to his model of rationality.

This leaves us, in my opinion, with an unsatisfactory version of a unique individual person. If microscale actions are deterministic or random, then the only significance we can ascribe to them is based on how they fit within broader mesoscale actions. I shall argue that we *do* take microscale actions as significant, even when abstracted from mesoscale actions, and even when they cannot be served by an appeal to rationality called for by the intentional stance. An example to illustrate this point can be found in the musical arts. Consider a performance by a group of jazz musicians adhering to a particular style, such as bebop. When they perform a piece, there are a number of constraints that, although violable, nevertheless affect what is played at any given moment during the piece. There are formal musical constraints, such as rules of harmony, perhaps a certain agreed upon turn-taking structure, and so on. There are also informal musical constraints, such as generally agreed upon practices about how to structure a solo or accompaniment. In addition to these sociocultural constraints, clearly, there are also the physical constraints of both human physiology, including lungs, muscles, nerves, and ears, and the physics of musical instruments, including solid materials and air.

All of these constraints are enough to limit the total possibilities of what could be played at a given moment — yet the constraints allow for more than one possible

<sup>f</sup>As Dennett points out, a further distinction applies between “random” as “patternless” (what computer scientists usually mean) and as “physically undetermined” (more likely in discussions of metaphysics) [Dennett, 1981, p. 298]. For the purposes of this discussion, either version has the same consequences.

action. So if a performer’s mesolevel actions could be explained by a rational, intentional appeal (observing that it was one player’s turn for a solo, the solo was reaching a denouement, certain melodic, harmonic, and rhythmic elements were given, and so on), then how should we understand the significance of the soloist playing one particular note as opposed to another note that would have done equally well? And what is the significance of *that* soloist playing that note within the same constraints another player might encounter? If everything is deterministic, or, if players are following random paths within the constraints, in either case, no notes or players could be significant. Let us assume that of two available musical notes within a set of constraints, neither could be rationalized as “better” than or “more preferable” to the other. The rational, intentional appeal would mean that, as long as the available notes were interchangeable (like cans of soup), the actual note played would be meaningless, except as it relates to the broader appeal. I do not find this conclusion satisfactory. It seems to me a better strategy to work backward from the sense that a particular note by a particular player is in fact meaningful.

### 8.1. *Bach and Dolphy*

Although Dennett does not spend a lot of time addressing music or the arts in general, he does make the following claim:

We honor Bach, the artistic genius, but he was no “natural” doodler, an intuitive genius just “playing by ear”. He was the master musical technologist of his day, the inheritor of musical instruments that had had their designs honed over several millennia, as well as some relatively recent additions to the music-maker’s toolbox — a fine system of musical notation, keyboard instruments that permitted the musician to play many notes at once, and an explicit, codified, rationalized *theory* of counterpoint. These mind-tools were revolutionary in the way they opened up musical design space for Bach and his successors [Dennett, 2001, p. 319, original emphasis].

Here, Dennett opposes artistic genius and technological mastery to intuitive genius and playing by ear. He argues that being a “master” of musical technologies allowed Bach to find the right points within the “musical design space” that would become the most popular. He fails to acknowledge that one of these musical “technologies” was improvisation, for which Bach was famous in his own time, and instead focuses on Bach’s chorales. He describes Bach’s compositional process as one of “breeding” popular and memorable works. By focusing on the level of a whole, consciously formulated work, Dennett is, as with the intentional stance, operating on the temporal mesoscale. Specific note choices in this context can be rationalized as being in the service of a number of broader concerns, such as how the note functions within the melodic and harmonic framework, and how it functions according to the broader intention to make the piece popular and memorable.

But, from a different perspective, we might be interested in a specific “in the moment” note choice made during an improvisation. In this case, we may consider

a more recent improviser, Eric Dolphy, another well-regarded, though less widely known “artistic genius”. In an interview between Dolphy and Leonard Feather [ca. 1964], excerpted below,<sup>§</sup> Dolphy highlights aspects of his approach to improvisation. His description resists the rational framework proposed by Feather, who appears frustrated by this fact:

- LF: “[There are] certain things I don’t understand about what you are doing- [...] What people like you are doing, harmonically, it’s very hard to explain, very hard to analyze. Can you put it into words?”
- ED: “Well, [...] the lines are not held to no chord pattern, harmonically.”
- LF: “No, but what I don’t understand is what *are* they held to? I mean what is the- what is the difference between the limitations? There must be some limitations otherwise it- you would be arbitrary, you could just play any notes you like.”
- ED: “Well that’s the idea. You *can* play every note that you like. Of course, you only can play what you can hear. [...]”
- LF: “Well, if your foundation is not a chord sequence, which is what the traditional basis of jazz was, then what is the foundation?”

Feather’s traditional musicological rationalization that could also be applied, with some minor adjustments, to Bach’s music clearly does not apply here. It is interesting to see how Dolphy identifies the relationship between being a “master musical technologist” (which he undoubtedly was) and being “intuitive” and “playing by ear” in the following few excerpts. As Dolphy explains:

You see, some things you play are not based on chords [...]. You start with one line and you keep inventing as you go along, linewise, and you keep creating until you state a phrase. And quite naturally you — what is the word, intuitive? — whatever’s around you, or with you, [you are] working together [with it]. [...] So harmonically, it’s not held down to the old thing of where you have a seventh chord and you keep running the chord [playing within its harmonic constraints]. Quite naturally, you run the chord, but you use other notes in the chord to give you other- to express another kind of sort — or you’d be playing the same thing everybody else is playing.

At the end of this passage, he seems to be referring to other expressive possibilities that extend beyond a superficial adherence to formal constraints, while, at the same time, pointing out that he maintains an awareness of formal constraints. The intuitive sense he describes sounds more like an ecological approach (see Clarke [2005]) than a generate-and-test procedure. He also seems to be describing a process of real-time development of musical material that takes place in relation to the

<sup>§</sup>Some filler words have been omitted for readability.

unique individual perspective that a player has, rather than a process of just-in-time selection of the nearest available note.

Later in the interview, Dolphy states:

A musician coming up now, he has had more training. Not to mean that he's better, I'm not saying that he's much better. I'm saying he has training, [is] better equipped, that he has a little more technique, so he doesn't know what to do with it. In the case of myself, I had to find something- what to do. Not to say in the sense of finding something to do just to exhibit my technique, but to find something to do to enhance some kind of musical- [to] make some kind of musical sense.

And I found that within my playing that I could play notes- not at first, because at first I couldn't hear these notes, so I wouldn't play them. But as I play more and more, I hear more notes to play against the more common chord progressions, and a lot of people say they're wrong. Well, I can't say they're right, and I can't say they're wrong. [...] To my hearing, I'm right.

This clearly reveals that his efforts are not focused on crafting a popular, memorable tune. Moreover, as Dolphy said earlier on, simply following the formal constraints could potentially lead one to be “playing the same thing everybody else is playing”, lacking the significance of a unique contribution. But in his further elaboration, he points out that, rather than having no constraints, the skill of listening and working together with what is around you functions in relation to technical mastery. As he describes it elsewhere in the interview:

[This approach to improvising] opens up a whole different type of hearing. [...] I knew about this [approach] quite a while ago, but I couldn't do it because I couldn't hear, you know, and it's not a case of just like, going out and then saying, “Well, I'm just gonna play anything”, because I couldn't make any musical sense out of it — you know what I mean? — to build a line against what I hear. [...] It's not a question of just running notes. And so, a lot of people say- they hear guys playing [and] they say, “he's just running notes”, just to be running notes at random. But that isn't true [...].

So for Dolphy, the notes are clearly not random, which would also render them meaningless.

Finally, it is interesting to point out that when Dolphy endeavors to explain himself through a retrospective reconstruction of his approach, he finds his own explanation inadequate: “It's hard to say at the moment, as I'm sitting here, because you know improvisation- the thing only happens at the moment of when you're doing it.” This sentiment fits nicely with Dennett's recognition (cited above) that not all action can be consciously deliberated: “and a good thing, too, since there wouldn't be time”. Yet, here, it would not seem right to equate note choice with soup can choice. Moreover, Dolphy's statements resonate with Dreyfus' point that in-the-moment production may extend beyond the narrow conception of a skill as

a specialized set of habits.

## 8.2. *Informed interpretation*

The question of how a particular note could be meaningful is thus far unanswered. A potential philosophical move at this point in the argument could be to say that a note expresses a particular intention of the performer. It could be claimed that this intention is “inside the head” of the performer, or that the execution of the note itself is a manifestation of a performer intention that need not have been formed in advance (on the latter, see Taylor [1979]). I do not wish to follow either of these routes. Here, I find myself partially in agreement with Dennett in that, with respect to interpreting individual agent action, the intentional stance appears to provide a sufficient account of what is relevant to our understanding.

However, I object when Dennett [1981] discards as trivial any in-the-moment selection between equally valid choices [pp. 290ff.]. Returning to the example of the performing musician, I wish to explain the meaningful choice of note another way. Taking a cue from Heidegger, we could say that if the skilled activity of a performance is situated within the broader world of the performer — within the performer’s everyday coping that extends far beyond the time frame of a single performance — then this everyday coping must also have some bearing on what happens *during* a skilled activity. While this may be less clear with respect to pragmatic activities, it is apparent when we consider artistic activity. I will argue that we make sense of why one note was played rather than another on the basis of who a performer is *as a person*. This argument is based on the premise that all life experience affects everyday coping, which in turn affects skilled coping. But would this mean that we need to know who a performer is as a person to understand the significance of what notes he or she plays? My answer is no, or at least not entirely, and this is where the role of the expressive stance will be made clear.

The expressive stance parallels the intentional one in the sense that it is an interpretive stance, dependent upon having something to interpret. With the intentional stance, we use “a sort of hermeneutical process that tells the best, most rational story that can be told” [Dennett, 1987, p. 92]. When things do not add up, when we are faced with an agent’s fallings-short when demonstrating “imperfect and inappropriate proclivities and inactivities”, these diminish our grounds for ascribing intentionality [Dennett, 1987, p. 92]. Simply put, information has the capacity to affect our judgments; new information allows us to revise a previous judgment, or, what might amount to the same thing, to invalidate a previous judgment and make a new one. This seems uncontroversial; it is certainly the case in a court of law, and seems to be the case for most of our everyday considerations.

This process of information gathering and judgment should not be considered any different in the evaluation of an artwork or artistic performance. The role of assumption is particularly relevant here. When we have an aesthetic experience of an artwork, the experience is guided by a large number of assumptions. To see a

painting, for example, and have an aesthetic experience, goes beyond the merely sensorial. Our everyday sense impression of a work does not yet make it art. Any situation in which we encounter, for example, a smear of oil paint on a canvas may be enough to bring our assumptions to bear on it. Further information may in some cases bring with it more assumptions rather than less, such as learning that the canvas is usually on display in a gallery or museum. We also have implicit assumptions about what it would mean for a human to become a painter and produce a particular painting. If we later discover a machine produced the painting, our original sensory experience has not changed, but the surrounding assumptions, interpretations, and judgments have changed with the new information.

Or consider a reversal of this scenario, in which something goes unrecognized as art until more information has been gleaned. For example, at the 1997 art exhibition *documenta X* in Kassel, Germany, at an abandoned train platform, an

unused track was overgrown with weeds, among which the Viennese artist Lois Weinberger had planted more weeds, indistinguishable from those already there. One learned later, however, that they were of a fast-growing introduced kind, resistant to herbicides and capable of overrunning the indigenous species. [...] Once made aware, the subtlest irony and humor tuned in visitors to the complex layering of unsuspected meaning that lurked in the details of the most innocent-looking ‘installation’ [Miyoshi, 1998, p. 152].

In this case, various assumptions led to an initial failure to recognize the weeds as an artwork. But, upon discovering their context, the additional information allowed for a new interpretation and aesthetic judgment. Further information about the artist’s life work — for example, that he developed a relation between gardening and contemporary art over several decades — allows for further interpretation and judgment, and so on. As a result of this process, one could have an extended aesthetic experience that is irreducible to the sensory impression of the artwork.

It is important to point out that this argument does not fall into the trap of elevating the artist’s authority. The familiar phrase “death of the author” stands in for the (by now) widely recognized claim that artists do not have privileged access to the meaning of their works. A spectator or critic may provide an interpretation of a work or performance that is equally valid to that of the artist, or perhaps even more valid, depending on how the interpretation is grounded. While there can, in principle, be no definitive interpretation, nonetheless, the context of an artist’s life may provide crucial information to help judge, interpret, and thus experience specific works. Even when a work is compelling in isolation, our implicit assumptions guide us, and further information can potentially alter our initial assessment, for better or worse.

### 8.3. *Expressive timing*

The expressive stance works by treating the human situation implied by an artwork as the background against which we interpret the significance of the microscale. This sense of “expressive” is ultimately consistent with the notion of “expressive timing” that has been identified in studies of various musical performance practices; such studies acknowledge that expressive timing variation (even of well-known composed works) is not solely determined by musical structural constraints (e.g., Repp [1990]). In considering the sociocultural role of local traditions of dance, or rituals related to the history of slavery, Vijay Iyer [2002] connects a broader (macrolevel) context with manifestations of physical embodiment and situatedness in the arts, in particular, concerning expressive microtiming in musical performance. But Iyer does not go far enough in his criticism of how we construe intentionality in improvised music. He points out that it does not make sense in improvisation to consider questions about whether someone “intended” to play something specific, because this would presume that something specific was “supposed” to happen (quotation marks in original). In contrast, he explains, “from the perspective of an improviser, the notion of a mistake is supplanted by the concept of interaction with the structure suggested by the sonic, physical, and temporal environment” [p. 408]. While this characterization is at one level accurate, there is nevertheless a problem with how it is framed.

On his model, there is still a (*post hoc*) rational story we can tell about why a specific course of interaction was undertaken with respect to the stated constraints of the sonic, physical, and temporal environment. That is, one could say that this multidimensional environment points to something like the embodied situation captured by Searle’s deep Background, and the sociocultural contingencies of the musical tradition form something like the local Background. On this basis, even as Iyer attempts to resist an intentional characterization, he winds up with a hierarchy in which a specific microscale action (e.g., a single note) derives its meaning from a mesoscale activity such as an expressive performance (e.g., one that makes use of a certain kind of rhythm), while the performance derives its parameters of meaning from the macroscale of a performance tradition (as practiced by an individual). I do not in principle object to this strategy of intentional interpretation, which I find compelling for certain modes of analysis. This perspective does, however, render invisible the other dimension of meaning I wish to describe here.

With the expressive stance, the specific moments of an improvisation are not only relative to the constraints leading up to these moments. A specific note played at a given moment is different than what someone else under the same constraints might have played at that moment, not merely because they selected a different note among the available options, as an exercise of skilled activity, and not merely because they have a different physical body. Rather, the expressive stance reveals that a person brings their life experience to bear on an aesthetic situation in a way that allows for a new perspective on what is possible within a set of constraints. Even if two performers, under formally identical conditions, would have played the

same note — seemingly dictated by the constraints — the difference is the person. The expressive stance makes this relationship between microscale action and the (macroscale) life of a person intelligible.

## 9. Conclusion

We need language to equip ourselves with certain ways of relating to the world and to one another — to give us evidence and understanding of shared experience and shared intuition. Much of this language may be relegated to scaffolding that plays no direct role in a wide variety of non-linguistic experience. Yet, without language, we cannot go beyond intentional attributions, as we may do with a frog or specialized machine, and we cannot go beyond a reified, apparently unmediated series of moments. With language, we can obtain relevant information about shared experience — or reflect upon our assumptions about shared experience — to make interpretations from the expressive stance. The expressive stance allows us to attribute significance to the specific elements of a given artwork or artistic performance judged in relation to its wide-ranging connections to shared everyday experience, not merely in terms of an imagined rationally justifiable motivation for each of its elements.

Through observation and interaction, especially linguistic interaction, we may one day conclude that a machine shares our intuitions and sensibilities about everyday experience. At such point, whether or not we would count the machine as having consciousness is reduced to a “tactical and lexical decision about whether it would be misleading”. But, if we were to be convinced that a machine shared our everyday experience, its artistic production would take on a different meaning than the skilled activities performed by today’s expert systems; it could no longer be thought of as simply an aesthetic extension of its designers. If that point were ever reached, only then would we be free to interpret a machine’s art as an expression akin to our own.

## Acknowledgments

Many thanks to Antonio Chella and Iain D. Thomson for their encouragement and helpful suggestions on earlier versions of this paper, and to Andrea R. English for helpful discussions at various stages of its development.

## References

- Chella, A. and Manzotti, R. [2009] Machine consciousness: a manifesto for robotics, *International Journal of Machine Consciousness* **1**(1), 33–51.
- Chella, A. and Manzotti, R. [2012] “Jazz and machine consciousness: Towards a new Turing test,” in V. Müller & A. Ayesch (eds.), *Revisiting Turing and his test: Comprehensiveness, qualia, and the real world (AISB/IACAP symposium, Alan Turing Year 2012)*, pp. 49–53.

- Clarke, E. [2005] *Ways Of listening: An ecological approach to the perception of musical meaning* (Oxford University Press, Oxford).
- Dennett, D. [1981] *Brainstorms: Philosophical essays on mind and psychology* (MIT Press, Cambridge).
- Dennett, D. [1987] *The Intentional Stance* (MIT Press, Cambridge).
- Dennett, D. [1991] *Consciousness explained* (Little, Brown, and Co., New York).
- Dennett, D. [2001] The evolution of culture, *The Monist* **84**(3), 305–324.
- Dolphy, E. and Feather, L. [ca. 1964] Leonard Feather interviews Eric Dolphy, <http://adale.org/Discographies/Feather.html>, transcription modified for accuracy using original audio source.
- Dreyfus, H. L. [1972/1992] *What computers still can't do: A critique of artificial reason* (MIT Press, Cambridge).
- Dreyfus, H. L. [1991] *Being-in-the-world: A Commentary on Heidegger's Being and Time, Division I* (MIT Press, Cambridge).
- Dreyfus, H. L. [1993] Heidegger's critique of the Husserl/Searle account of intentionality, *Social Research* **60**(1), 17–38.
- Dreyfus, H. L. [2008] Why Heideggerian AI failed and how fixing it would require making it more Heideggerian, in P. Husbands, O. Holland & M. Wheeler (eds.), *The mechanical mind in history* (MIT Press, Cambridge), pp. 331–371.
- Heidegger, M. [1962] *Being and Time* (Harper and Row), trans. J. Macquarrie and E. Robinson.
- Heidegger, M. [1982] *The basic problems of phenomenology* (Indiana University Press, Bloomington), trans. A. Hofstadter.
- Hendriks-Jansen, H. [1996] *Catching ourselves in the act: Situated activity, interactive emergence, evolution, and human thought* (MIT Press, Cambridge).
- Iyer, V. [2002] Embodied mind, situated cognition, and expressive microtiming in African-American music, *Music Perception* **19**(3), 387–414.
- Law, J. [1987] On the social explanation of technical change: The case of the Portuguese maritime expansion, *Technology and Culture* **28**(2), 227–252.
- Lewis, G. [1999] Interacting with latter-day musical automata, *Contemporary Music Review* **18**(3), 99–112.
- Lewis, G. [2000] Too many notes: Computers, complexity and culture in voyager, *Leonardo Music Journal* **10**, 33–39.
- Lewis, G. E. [1996] Improvised music after 1950: Afrological and Eurological perspectives, *Black Music Research Journal* **16**(1), 91–122.
- Linson, A., Dobbyn, C. and Laney, R. [2012] “Critical issues in evaluating freely improvising interactive music systems,” in M. Maher, K. Hammond, A. Pease, R. Pérez y Pérez, D. Ventura & G. Wiggins (eds.), *Proceedings of the Third International Conference on Computational Creativity*, pp. 145–149.
- Linson, A., Dobbyn, C. and Laney, R. [2013a] A parsimonious cognitive architecture for human-computer interactive musical free improvisation, in *Advances in Intelligent Systems and Computing 196: Biologically Inspired Cognitive Architectures*

22 *References*

- 2012 (Springer, Berlin), pp. 219–224.
- Linson, A., Dobbyn, C., Lewis, G. and Laney, R. [2013b] An artificial agent for collaborative free improvisation, (*submitted*).
- Miyoshi, M. [1998] Radical art at documenta X, *New Left Review* **I/228**, 151–160.
- Newell, A. and Simon, H. A. [1976] Computer science as empirical inquiry: Symbols and search, *Communications of the ACM* **19**(3), 113–126.
- Polanyi, M. [1962] *Personal Knowledge: Towards a post-critical philosophy* (Routledge and Kegan Paul Ltd., London).
- Repp, B. [1990] Patterns of expressive timing in performances of a Beethoven minuet by nineteen famous pianists, *The Journal of the Acoustical Society of America* **88**(2), 622–641.
- Searle, J. [1980] Minds, brains, and programs, *Behavioral and brain sciences* **3**(3), 417–457.
- Searle, J. [1983] *Intentionality: An Essay in the Philosophy of Mind* (Cambridge University Press, Cambridge).
- Taylor, C. [1979] Action as expression, in C. Diamon & J. Teichman (eds.), *Intention and Intentionality: Essays in Honor of G.E.M. Anscombe* (Cornell University Press, Ithaca), pp. 73–89.
- Taylor, C. [1993] Engaged agency and background in Heidegger, in C. Guignon (ed.), *The Cambridge Companion to Heidegger* (Cambridge University Press, Cambridge), pp. 317–336.
- Turing, A. M. [1950] Computing machinery and intelligence, *Mind* **59**(236), 433–460.
- Wheeler, M. [1996] From robots to Rothko: The bringing forth of worlds, in M. Boden (ed.), *The philosophy of artificial life* (Oxford University Press, Oxford), pp. 209–236.