Philosophical Investigations 28:2 April 2005
ISSN 0190-0536

Wittgenstein and the Aesthetic Robot's Handicap

Julian Friedland

> The concept of a living being is as indeterminate as that of a language.
>
> Compare: inventing a game – inventing a language – inventing a machine.
>
> Wittgenstein, *Zettel*, §326–7

Since the publication of Turing's seminal article "Computing Machinery and Intelligence" (*Mind*, 1950), there has been a half century of great skepticism of Turing's claim that artificial intelligence would eventually come to rival human intelligence. However, we now live in a world in which this question seems to have been settled – at least as far as Chess is concerned. In 1997, the IBM computer 'Big Blue' narrowly defeated Gary Kasparov, arguably the greatest chess player the game had ever known. Kasparov even declared with astonishment that the computer seemed to really *think*. So the question now before us is whether this victory also attests to the computational system's ability to master the most essential human practices of language and aesthetics.

To better understand what is at issue here, one must first recognize that chess itself is ordinarily composed of two distinct aspects – one, purely mathematical and computational, the other sportive and aesthetic. Most players, with the exception of computers and certain autists, perceive both. The computational aspect includes the rules of the game, the goal of each player, and the calculation of possible move combinations. Since most computers now largely exceed the speed at which humans can calculate long mathematical operations, it is not very surprising that a computer has finally succeeded in beating the best human chess player. While a particularly gifted human can, in a few minutes, think a maximum of six or seven chess moves ahead, Big Blue reportedly computes two hundred million moves per second.

Although this method of carrying out a game of chess (I do not say 'playing') represents a certain kind of mastery, it does not capture the sportive aspect of the game. If two identical chess computers were to go against each other, they would invariably reach a draw. Part of what renders this paradigm impossible between humans is the fact that we compliment strategic calculation with sportive tactical maneuvering. Since we are not in the habit of playing adversaries endowed with twice our calculative capacities, it is usually advantageous to take into account various approximating tactical positionings. We hence engage myriad tactical arrangements both shielding us from certain possible though unsuspected affronts, and playing on the emotions of our adversary. It is true that a computer is also capable of effecting preemptive maneuvers before actually predicting any specific need for them, e.g., castling the king, or building a pawn chain. However, the computer does not yet perceive the *sportive* interest of such moves, i.e., their psychological dimension representing various degrees of assurance or apprehension, including the affective impact that those impressions may (or may not) have on the play of one's opponent. Hence, it does not bluff.

As I said, a powerful computer has the advantage of exceeding the speed at which humans can calculate. It therefore has no need to develop a sense of the psychological dispositions of its adversary in order to beat that particular opponent at every game of every match. Kasparov, being human, had often claimed that one of the essential qualities of a great chess player was audacity. But this attitude becomes superfluous against a digital opponent incapable of recognizing, valuing, or exercising it. And indeed, Kasparov said he felt like he was battling "an alien intelligence."

There are hence two chess-gaming forms: computational chess – *performed* purely mechanically; and sportive chess – *played* affectively. While computers have up to now only mastered the former, their capacity for learning the later is widely considered to remain an open empirical question. But if we take this question to be "Could a digital computational system develop aesthetic sensibility?" then I will attempt to show, while drawing upon the later Wittgenstein, that the correct answer is in fact available. And it is a negative 'a priori'.

Up to now, computers have shown themselves capable of mastering many complex practices and procedures. Yet their greatest successes have always been limited to domains in which the logical

connections between rules and goals are rigorously defined. In fact, the more these relations are defined, the more likely they are to succeed. What tends to escape them are vague contexts such as those found in ordinary language, for which we only possess a general overview of the proper semantic, syntactic, symbolic, descriptive and prescriptive functions of the data at hand. Since these functions operate largely according to myriad variable and evolutionary contexts, one rarely finds distinct foundational parameters governing them. The hope of artificial intelligence is hence to succeed in developing an evolutionary program capable of operating according to new and unforeseen linguistic contexts. In other words, the computer will possess an enduring ability to learn for itself through trial and error.[1]

The overarching cognitive faculty adapted to this task is the capacity to perceive what Wittgenstein calls physiognomy:

> While any word – one would like to say – may have a different character in different contexts, all the same there is *one* character – a face – that it always has. It looks at us. – For one might actually think that each word was a little face; the written sign might be a face. And one might also imagine that the whole proposition was a kind of group-picture, so that the gaze of the faces all together produced a relationship among them and so the whole made a *significant group*. But what constitutes the experience of a group's being significant? And would it be necessary, if one is to use the proposition, that one feel it in this way? (RPP, §322; PI, p. 181e)

The last sentence of this passage makes it clear that what interests Wittgenstein is not to render correct linguistic usage impossible by someone (or something) unable to perceive the aspect of physiognomy, but rather to investigate how this anomalous blindness would distinguish itself from ordinary mastery. Wittgenstein often returns to this question in thought experiments concerning linguistic operations performed in what he generally refers to as 'aspect blindness'. This phenomenon not only occurs in computers, but also in certain

1. Even if this hope is realizable, it would seem that the greatest utility of the digital would be irretrievably compromised, namely, its extreme operational reliability. The irony here is introducing the errors necessary for the program's continuing adaptation when the ultimate value of the digital has always been to make errors impossible, or at least improbable. What benefit could there be in reproducing human error in a tool of precision?

autists and even from time to time, in everyone. It happens whenever one overlooks say, the metaphorical significance of a word, as when missing a pun. The aspect blind in this sense is simply the one who does not immediately apprehend the modified use of a term or phrase given under a new context. Such occurrences obviously arise most often within the purely expressive domains of gestures, tones, images and music. So the central question here and upon which Wittgenstein places great emphasis in his later writings is: What would an aesthetic mastery consist of that operated through blindness at each new transformation?[2]

But for this question to have meaning, it is first necessary to show that we do in fact experience physiognomic perceptions engendering immediate recognition of an expressive or descriptive "face" through family resemblance, and not by some essential trait. An artificial intelligence researcher could maintain that the physiognomic aspects that we perceive are nothing but amalgams of cognitive adaptation, or new uses recognized merely by the successive synthetic correction of initial blindnesses. Hence, we would simply be very good autists, continually cataloguing new cases in our mental libraries representing the total number of uses that we would be capable of recognizing (without annexation) at any given time.[3]

However, if we examine the ordinary experience of family resemblance, we notice that the physiognomic element affords us imme-

2. One could take for example the now famous case of Temple Grandin who, despite holding a post of Assistant Professor of Behavioral Science at Colorado State University, has nevertheless suffered from autism during her entire life. She experiences great difficulty in interpreting and predicting the actions and reactions of people in general. "She said that she could understand 'simple, strong, universal' emotions but was stumped by more complex emotions and the games people play. 'Much of the time,' she said, 'I feel like an anthropologist on Mars.'

She was at pains to keep her own life simple, she said, and to make everything very clear and explicit. She had built up a vast library of experiences over the years, she went on. They were like a library of videotapes, which she could play in her mind and inspect at any time – 'videos' of how people behaved in different circumstances. She would play these over and over again, and learn, by degrees, to correlate what she saw, so that she could then predict how people in similar circumstances might act. (Sacks, Oliver, "An Anthropologist on Mars: A neurologist's Notebook", *New Yorker*, 27 Dec. 1993).

3. In this case, all novel representations would be the result of private conventionalizing acts of will. This view is repeatedly challenged by Wittgenstein, most notably in his attack on the notion of private language. See for example PI, §243–313. Also, my article "Ideation and Appropriation: Wittgenstein on Intellectual Property" (*Law and Critique*, vol. 12, 2001) presents a Wittgensteinian critique of this view in intellectual property law.

diate apprehension of new aspects, without passing through any initial stage of blindness. If, for example, we were presented with a jumble of photographs of several unknown persons, each individual being represented at ten-year intervals, we would usually succeed in recognizing that certain pictures represented the same person at different ages. Similarly, when we observe a family photo, we often receive an immediate impression of resemblance between all of the members pictured. However, it is extremely rare to find one essential trait that could serve as a criterion of heredity. This is the reason why computers have great difficulty in recognizing such resemblances. They are only able to unite such images artificially by considering them as a set of objects in which none of the contents share any distinctive trait beyond mere membership of that set.

The physiognomies by which the members of a family are recognized however do contain similarities between intermediate cases. While no single trait is present in each member, when we compare them one by one, we observe certain common traits; sometimes overall similarities, sometimes similarities of detail which disappear as others emerge, then reappear, overlapping and criss-crossing as we pass from one face to the next.[4] Furthermore, the number of such transformations and nuances has obviously no a priori limit, for we can always add one more distinct member to even the most extensive progeny.

Interestingly, we do not ordinarily have to compare all the intermediate cases of an entire family in order to obtain a physiognomic impression of the group as a whole. If we did, we would need to examine the resemblance between all individual members – which would represent, say for a family of twelve, sixty six separate comparisons. When a new recognizable member is observed, we recognize that person almost immediately. Hence one might exclaim while looking at a family photo album: "Oh, I clearly see that this boy is a Kennedy". And this of course would not necessarily imply that the boy in question shares any single feature or group of features with every member of the Kennedy family.

The same is true of all cases of family resemblance, whether between words, phrases, styles, or entire families. If we can immediately understand a novel use of a word, or sentence in a poem or essay, we can also notice a change of aspect concerning an entire

4. Wittgenstein, PI, §66.

family and compare it to others. This is how concepts evolve, such as for example the concept of 'number' throughout the history of mathematics. Hence we now dispose of myriad different families of numbers, families which correspond and distinguish each other in specific ways. Although each family of numbers is rigorously circumscribed, i.e., its members all share at least one essential aspect, its parameters can nevertheless always engender new evolutionary transformations. We may therefore wonder if the word 'number' is a family resemblance concept. It is indeed likely that no computer could ever discover any attribute sufficient to justify the use of this word throughout all of its mathematical functions, e.g., natural, irrational, real, imaginary, ideal etc.

But then, how is it possible that *we all* may justly recognize a kinship between all the divergent uses of the word 'number' or 'game' or that of any ordinary family resemblance concept? This question concerns the very identity of lexical attribution, which is the reason why Wittgenstein returns to it so often. For example, how do we succeed in recognizing what is called "shakespearian style" throughout an immense variety of different works, in a way sufficiently precise to be able not only to gather them up under the name of a single author, but also to attribute to this *set* certain direct influences of works conceived in foreign languages and distant epochs, namely, those of Seneca?

Recent developments in linguistic analysis suggest that the attribution of a text to its author is a *science in waiting*. In fact, This is precisely the hope of Donald Foster, professor of literature at Vassar College, who has attributed texts to countless famous and infamous authors from Shakespeare to Kazinski, a.k.a. "the Unibomber". In fact, it seems Foster has a perfect record of success which has made him the F.B.I.'s only resource in this particular domain of investigation. With the help of a computer program of his own design, Foster detects stylistic patterns and carries out lengthy statistical calculations to determine the identity of any author at hand. Given enough text and enough time, a set of "linguistic fingerprints" is obtained which can then track the literary activity of any author alive or dead. Foster claims that we each have a "stylistic signature" of which we are not entirely aware, since it must be used even to formulate the very thoughts on our own style. And here is a most Wittgensteinian theme, for according to Foster "during authentic composition, one

is ignorant of the process which engenders one's own speech."[5] Hence we would each possess a stylistic signature by which we are each inexorably betrayed. Which recalls Wittgenstein's remark: "Compare a concept with a style of painting. For is even our style of painting arbitrary? Can we choose one at pleasure? (The Egyptian, for instance.) Is it a mere question of pleasing and ugly?" (PI, p. 330e).

Let's accept for the sake of argument that textual attribution is really a science in waiting, practicable purely by internal textual evidence. Our question then becomes "Would it be possible for a digital computational system to practice that science alone?" To answer this question, lets return to the example of recognizing images of the same persons at, say, ten year intervals. What will the computer do to achieve this task? Well, it will start by measuring the precise relations of facial features, by covering each face with a numbered grid to then compare the images in all the squares bearing the same coordinates. So, for example, it will compare the square A-1 of the first image with the square A-1 of all the other images, and A-2 of the first image with every other A-2 and so on down the line. Then, it will only have to choose those faces bearing the same physiognomic proportions to obtain the correct answers. Similarly, for textual attribution, the computer will compare the grammatical structure, conjugations, together with the frequency in which the same word is used in the same way, in order to determine the stylistic habits of the author at hand. Or again, if the task is to uncover the proper usage of the word 'number', or 'game', it will have to compare all the given contexts, to obtain the definitional parameters that constitute the general concept.

The principal difficulty for the computer will be knowing when it will be appropriate to let a certain trait or concept vary, and when it will not. Each datum will not necessarily be generalizable. And when it is, how will it know how far to apply the rule? In sum, the task will be to distinguish content from form. Why is it that we do not have to make all the comparisons the computer requires in order to recognize the resemblance between family members, faces, expres-

5. Caleb Crain, "The Bard's Fingerprints: Donald Foster Uses high-powered Computer Tests to Search for Shakespeare's Hidden Hand. His Critics Challenge Him on Every Move," *Lingua Franca*, July 1998.

sions, words, phrases, styles, characters, histories, and epochs? We immediately apprehend what Wittgenstein calls the "atmosphere" and "physiognomy" of meaning without having to confirm our judgments by going through a mental list of all the other cases having fallen under the same appellation. We succeed, for our collective psychology *projects* a certain content upon various forms. This is the phenomenon known as "intentionality" or "directional fit". Every meaning bearing object, be it expressive, symbolic, or denominative is perceived in a certain *way*. And once this association is established, it is impossible to observe the same object without having that same intentionality come to mind. This is why Wittgenstein points out that each sign is a face. Try as I may, I cannot perceive, for example, the signs on this page without *reading* them, i.e., as merely well defined traces of ink devoid of any meaningful content.

The simple act of reading even a single letter of the alphabet transmits a given intentional quality, that we could call its *life*. Try for example to imagine a real human face bearing absolutely no expression. This is just as impossible, for each feature conveys myriad emotive capacities that characterize the organic totality in which it is contained. This is why one's mere appearance can play such an important part in our impressions if not our very judgments of others. And these impressions and judgments do not solely depend upon attitudes that are actually adopted or revealed. Each person already represents an entire world of singular (or not so singular) potentialities, before making even the slightest gesture. Our judgments of beauty and ugliness function perfectly well even on completely inert and passive subjects. For example, it would not be impossible for a fashion model to obtain an interview with a high profile designer by virtue of a single photo revealing only a bland expression. For even a face "without expression" still communicates some sort of content, no less significant, especially when presented as art – it will seem at times cold, at others detached, at still others melancholy, erotic, etc.

Each face possesses a different aesthetic life, since an intentional atmosphere belongs to each physiognomic perception. And if we all succeed in recognizing a family of different individuals, it is because each of us *projects* an intentional atmosphere upon its members by attributing to each of them myriad characteristic nuances nonetheless evolutionary and personalized. What I mean by 'personalized' is that despite the fact that such judgments are collective, each per-

ceiver invariably projects an intentional atmosphere peculiar to his or her own psychological makeup. This is of course the reason why ink-blot tests have been so widely used as tools of psychological diagnosis. Hence, one person may find a strong resemblance between two particular faces, without necessarily being able to share it with someone else. For example, someone might declare to a friend "I don't care what you say, my uncle still reminds me of Robert Redford, even though you claim not to see any resemblance." Nevertheless, these differences in judgment do not keep us from perceiving the same families as such, though each of us will brand them with more or less singular criterial impressions.[6]

Therefore, without the power of intentional projection, the aesthetic robot's handicap will be to master ordinary language by the mere trace of conventional motifs. It will not necessarily be impossible for such a system to carry out correct judgments of textual attribution based entirely on syntactic and semantic form. What is much less plausible however, is its ability to distinguish a stylistically perfect and coherent text, from one whose style, though just as perfect, happens to be utterly meaningless. It will not recognize for example that the following grammatical sentence is devoid of literal meaning, yet not of metaphorical meaning:

> Imageless ideas sleep furiously.

Furthermore, how will it distinguish between metaphorical coherence and metaphorical incoherence? And how will it determine if two verses of a poem complement one another or not? These are questions concerning the *content*, i.e., the *meaning* of a text, and not it's mere *form*.

Searle's infamous "Chinese room" thought experiment in fact perfectly illustrates this distinction between form and content. He uses it to liken the computational paradigm to a situation in which a person is forced to manipulate strange symbols (say, Chinese) according to precise rules. This poor soul is hence in a "Chinese room", all the while receiving cards, one by one, each bearing a different symbol. Every time a card is presented, this person must find the given symbol in a chart and copy the symbol corresponding to it onto another blank card. At which point, the new card must be

6. Compare PI, p. 169e and pp. 38–9 of Wittgenstein's *Lectures and Conversations on Aesthetics, Psychology and Religious Belief*, Cyril Barrett ed., Berkeley: 1967.

passed back through the slot from which the first appeared. This procedure could continue indefinitely between two people, following such orders. But would it consist of a conversation? That is to say, if someone fluent in Chinese read the cards as they passed between the two captives, observing a perfectly coherent exchange, could we agree that it would qualify as a real conversation? The answer is obviously 'no', for those two people would be completely ignorant of the significance of the signs they manipulated. We could not hence conclude that the practice in which each of them would be engaged qualifies as a *mastery* of the Chinese language.

So if these captives memorized their charts, they would only have mastered part of a written form of language without content. And it would be impossible for them to gain access to the meaning of their exchanges aided solely by this exercise. AI researchers will respond that this fact is trivial, so long as the person fluent in Chinese is convinced. Their goal is not to *replicate* the human process of cognition, but only to *imitate* it perfectly in a digital fashion. Their criteria of success is most often a new version of a particularly ingenious "imitation game" conceived by the mathematician Alan Turing, who it turns out had attended Wittgenstein's lectures on the foundations of mathematics at Cambridge – the published notes of which testify to his being Wittgenstein's primary interlocutor. The point of the game (or test) is for the computing machinery (as Turing called it) to mimic normal human responses to the point that any human interlocutor cannot distinguish between those of the machine and those of an actual human. Briefly, the rules are as follows: One person (the interrogator) must try to identify the gender of two invisible interlocutors, knowing that one is male and the other female. These two attempt to deceive the interrogator on their sexual identity by any means possible, such as for example, debating between each other, each trying to show the other is a liar. Certain games will be modified at random by replacing one of the interlocutors with a computer, which will have to behave exactly the way a man or woman would when attempting to verbally personify the opposite sex. The result desired by the programmer is for the computer to imitate a man or woman well enough that a human interrogator is fooled as often as when playing against two actual persons.

The programmer hence seeks to produce a computer that can regularly deceive human interrogators concerning one of the most nuanced aspects of human identity. This criterion of artificial intel-

ligence remains the most encouraging for it allows for any kind of processing necessary to obtain the correct result. In other words, the computer must find its own means of attaining the desired result, relative to the data given. This model avoids having to program strict orders, and seems to preserve the dynamic and evolutionary nature of the rule. The computer would thereby no longer remain a simple machine, for it would develop its own transformative contextual algorithms. This theoretical ambition bears the name "connectionism", for it lets the computer find its own *connections* sufficient to accord its input data with its output data. And each computer of this kind will "think" differently for it will have learned to do so according to different human interlocutors.

But does this model of apprehension really suffice for achieving mastery of all language games in general? Although the Turing test implies mastery of intentional nuance, this exercise is nevertheless always directed toward a precise goal, namely to fool an adversary. Indeed, the success of the connectionist system is entirely contingent upon the possibility of attributing to every language game, a goal sufficiently determined for the computer to know toward precisely what it must orient its computations, i.e., the output. So the connectionist conceives language games a bit like, say, chess or tennis. All that is required is to know the goal, and the means for getting there – the rest happens on its own. If we take tennis as an example, we see that a connectionist system could in principle just as well become a world champion as it could become a Grand Master of chess. Once means and goal are established, the contextual remainder – speed, placement, height of the ball, etc., is learned automatically.

Could the aesthetic robot hence learn this "science in waiting" that is, according to Foster, the attribution of texts by mere internal evidence? It would not seem that this would necessarily be impossible, for this task does seem to contain certain requisite elements of the connectionist model. The computer is given several texts of different authors, and must be capable of correctly attributing to each of those persons, all other possible texts that each of them could have written. If the "stylistic signature" can be entirely deduced from grammatical form, there is no reason why the attributive aptitude of the connectionist system would be any less compromised by its interpretive clumsiness with respect to content. However, without access to content, it will not be able to judge between good and bad works.

It will not, as we have seen, always be able to recognize perfectly grammatical sentences, nonetheless literally or metaphorically incoherent. Consequently, if we asked it to imitate one such author, it could not guarantee that its end result would be even the least bit convincing. Nevertheless, this handicap will not necessarily compromise its attributive task, since this mastery does guarantee after all the possibility of recovering the true "linguistic fingerprints" of even the most abominable work of any author. It will remain only as fallible as any real person when confronted with a product of purposeful fabrication masquerading as authentic composition.

What will keep the computer nevertheless from mastering language is the fact that every illocutionary act is not constituted by a determinate goal (or output). Furthermore, even when such acts do seek to produce a significant effect (descriptive, expressive, metaphorical, metominical, etc.) those intentions are not always deducible from any examination of the sentence itself. Hence even the input is unspecified. Take for example the following example given by Wittgenstein:

> One may have the feeling that in the sentence "I expect he is coming" one is using the words "he is coming" in a different sense from the one they have in the assertion "He is coming". But if it were so how could I say that my expectation had been fulfilled? If I wanted to explain the words "he" and "is coming", say by means of ostensive definitions, the same definitions of these words would go for both sentences.
>
> But it might now be asked: what's it like for him to come? – The door opens, someone walks in, and so on. – What's it like for me to expect him to come? – I walk up and down the room, look at the clock now and then, and so on. – But perhaps I say as I walk up and down: "I expect he'll come in" – Now there is a similarity somewhere. But of what kind?! (PI, §444)

In this example, the "similarity" of which Wittgenstein speaks, determines the meaning of the sentence. The question that he is asking is how this determination is carried out. If it is not carried out by means of definition of terms, it is carried out by the way in which the speaker *behaves*. And there is no precise justification for engaging in such behavior. The meaning of the sentence "I expect he is coming" is hence determined by behavior that is extra-linguistic and devoid of any specific goal. And if the meaning of our most ordinary statements is regularly determined in this way, the connectionist

system will be lost, for it will have criteria for neither interpretation nor success.

In sum, what the aesthetic robot lacks is an affective life. That is what allows us to make sense of highly contextual relations of motive and intent, namely, why someone would walk back and forth while awaiting one event, and not another. Such behavior would no longer be appropriate if, for example, a farmer announced that he awaited this year's harvest with the greatest impatience. Similarly, the affective life affords us the perception of, say, the nuance of a Baudelarian verse, or the understanding of the tragedy at the end of a Zola novel. Wittgenstein reminds us that "if you feel the seriousness of a tune, what are you perceiving? – Nothing that could be conveyed by reproducing what you heard" (PI, p. 210e). What is at issue here is the human capacity to recognize, feel, and share emotions. And these do not obtain merely from one to one relations between desires and satisfactions, since they overlap and transform continually. To feel love is not only to feel satisfaction. Love is something we notice about a given state of affairs that overtakes us; something which is in itself essential to our being able to refer to its meaning. And understanding that meaning is not to simply to be able to recognize its representative behavior.

Affective sensibility is what enables us to understand the vast spectrum of ordinary utterances by taking into account all sorts of extra-linguistic considerations. It permits us to take, for example, the same sentence as either a statement of fact or as an emotive expression such as an avowal at a funeral oration. For as Wittgenstein points out: "When it is said at a funeral oration 'we mourn our. . . .' This is surely supposed to be an expression of mourning; not to tell anything to those who are present. But in a prayer at the grave these words would in a way be used to tell someone something" (PI, p. 189e).

In order to distinguish when the same words form either an avowal or a description of a state of mind, we use sympathy and compassion. Consequently, the non-emotive computer will not be able to distinguish between descriptive and expressive utterances. Furthermore, our affective sensibility does not only function according to interpretive judgments. It binds us more deeply in shared sensitivity to the events of inner life, which together constitute the very fabric of human solidarity.

But as Sartre and others[7] have argued, our affective lives might simply consist of the successive satisfactions and frustrations of desire. In this case, there should be no a priori reason for ruling out the digital capacity to replicate this mere continuum of emotive reaction.

Although this theoretical model is useful for the purposes of understanding the characteristic role desire plays in the genesis of emotion, it tends to overlook the more fundamental perceptive dimension of emotive phenomena. Furthermore, it seems wrong to have to posit the existence of desire as a prerequisite to our having anything that can be significantly thought of as an emotive sensation. Must there be some compelling desire that allows us to feel the seriousness of a tune? Someone might respond by claiming that in such cases we are affected by melodic patterns which are composed in order to make the ear "want" to hear a specific note, or set of notes. Indeed musicians often speak as if this were the case, maintaining that good stirring melodies simply play with our auditory desires by first stimulating them and then artfully satisfying them at crucial moments in a way that results in emotive arousal.

It is not unlikely that something that could be called aesthetic desire is aroused by the right kinds of music. However we run into further difficulties when applying this account to such experiences as feeling the serenity of a countryside, or joyous at the sight of a bright, detailed, but surprisingly abstract Matisse painting. Consistency would require us to claim that as our gaze moves along, it desires to discover certain colors, textures, and designs. The depth of such desires would have to function as the necessary impetus that forces aesthetic feelings upon us.

The persuasiveness of this account fades even more when we consider the feelings precipitated by visual spaces that are too small for the eye to develop any "wants" because all the stimulus is conveyed at once. Perhaps none of these feelings actually qualify as emotions. They may only be subjective mental states identical to the most banal experiences of color. But if this picture is incomplete, i.e. there remains in these occurrences a kind of subjective qualia that allows us to become sad, inspired, serene or melancholy, then there must exist emotions that are not desire dependent.

Nevertheless, the experience of any emotion does require some sort of predisposition. The same is true of the most basic sensations.

7. Richard Wollheim, *On the Emotions*, pp.128–32, 207–11, Yale: 1999.

It may in fact be impossible to for example, determine where my disposition to see the sky as blue ends, and where my disposition to have blue as my favorite color begins. This might turn out to be a more accurate description of what is required for an emotion to take place. It is perhaps more appropriate to consider a desire as simply a sophisticated breed of emotive disposition. This would allow for there being reasons that are supervenient upon my genetic makeup, enabling me to have a favorite color without having anything which qualifies as a desire that I see it that way. In other words, the fact that I have the favorite color of blue does not necessarily satisfy a desire to have such a preference. However I may, one Winter's day in the Pacific Northwest, develop a desire to experience the beauty of a cloudless sky that conditions me to have a much more powerful emotion upon looking up at it.

Hence, our affective lives are never rigorously defined. They are dynamic and evolutionary. For as Anthony Kenny puts it, while quoting Wittgenstein:

> All feelings have duration; but perceptions and sensations are much more closely tied than emotions to the time which is the measure of local motion. One can hear a loud noise for just a second, or feel violent pain only for a moment, no matter what precedes or follows; one cannot in the same way feel ardent love, or deep grief for the space of a second, no matter what preceded or followed this second.[8]

Hence, coming to understand, say, Gustav Klimt's tri-partite visual depiction of Beethoven's Ninth Symphony may have a persistent effect on one's emotional and aesthetic awareness in a way that goes well beyond future encounters with that particular composer's work. Our cognitive habits in this way undergo continual spontaneous alteration. This truth paradoxically renders the Turing test all the more perspicacious, since it concerns the very identities of the masculine and the feminine. These genders vary in fact constantly, as much at the social as at the individual level. If a computer succeeded in *imitating* them as well as we *live* them, artificial intelligence would then become incontestably real. However, in ignorance of the affective content of masculine and feminine genders, the computational system will be at a loss to recognize the coherence of any novel vari-

8. *Action, Emotion, and Will*, Routledge, London: 1963, p. 58; from Wittgenstein, PI, §583 and p. 148e.

ation of these vague psychological constructs. It will, inexorably, have to pass at each unfamiliar instance, through a period of initial blindness, leading to the error necessary to the adaptation of its own repertoire. Conversely, the human interlocutor's aesthetico-linguistic judgments will adjust themselves more or less immediately, without having to resort to mere trial and error. Thus, the proverbial 'man' will always be at least one step in front of the proverbial 'machine'. There is the essential ontological difference between life and its imitation.

405 E Oak Street
Lafayette
CO 80026
USA
Email: friedlandj@gmail.com