

Emergence, Agency, and Interaction—Notes from the Field

Simon Penny*

University of California, Irvine

Abstract This article describes the development of several interactive installations and robotic artworks developed through the 1990s and the technological, theoretical, and discursive context in which those works arose. The main works discussed are *Petit Mal* (1989–1995), *Sympathetic Sentience* (1996–1997), *Fugitive I* (1996–1997), *Traces* (1998–1999), and *Fugitive II* (2001–2004)—full documentation at (www.simonpenny.net/works). These works were motivated by a critical analysis of cognitivist computer science, which contrasted with notions of embodied experience arising from the arts. The works address questions of agency and interaction, informed by cybernetics and artificial life.

Keywords

Artificial life, embodied interaction, machine vision, aesthetics of behavior, artificial life art, robotic art, interactive art, autonomous agents, human–computer interaction

I Introduction

Through the 1990s I was an active participant in the development of sensor-based embodied interactive art forms (www.simonpenny.net/works). My practice, like that of many of my peers, involved the development of novel working technological systems from the ground up, directed at functions conceived for cultural purposes. This was partly because of the shortcomings of available computational technology (processor speeds, memory capacity, display limitations, absence of sensors and i/o capability), and partly because the goals identified by artists were (at the time) outside the specifications of research agendas for digital appliances. (See Appendix A (Oops, my art became an appliance) in the online supplementary materials for this article at www.mitpressjournals.org/doi/suppl/10.1162/ARTL_a_00167.)

The challenge for artists working in such contexts was and is to generate culturally significant or meaningful experiences in these new technological environments. At least, that was the challenge I identified and pursued. In my opinion it remains an incomplete project. We lack a great deal in approaching an interdisciplinary aesthetics of automated behavior, which I've referred to previously as a CACA (computationally augmented cultural artefact). In what follows I will recount some technical, aesthetic, and theoretical issues as they emerged in the development of several artworks.

The goals of my own artistic, research, and development practice through the 1990s were informed by my training in art and its implicit commitment to embodiment, materiality, and performativity. These values were not made explicit in theoretical or aesthetic terms in my training, so my engagement with technology demanded an explication of the unspoken values of my field. I had the good fortune to be exposed to high-level development of AI and robotic systems in my capacity as professor of art and robotics at Carnegie Mellon University (international leader for research and

* 3229 Art Culture and Technology, University of California, Irvine, Irvine, CA 92697-2775. E-mail: penny@uci.edu

teaching in these fields). Informed by the intuitions of my art training, my sense of the shortcomings of these technical practices simultaneously forced me to interrogate both my own professional value system and that of AI and robotics.

Contrary to the enthusiasm surrounding graphics, multimedia, and networking, it has been clear to me since the 1980s that the fundamental novelty of computational media and computational art practice is the capacity for *behavior* [19]. I pursued a practice of building systems that conformed to the rigorous requirements of functionality and reliability of robotics but presented working examples of systems that contested assumptions, both implicit and explicit, of these fields.

I began building *Petit Mal*, an autonomous robotic artwork, in 1989. It was designed to explore the possibilities of a physically instantiated autonomous agent contrived to engage the visitor in interaction based in large bodily gesture and movement. It was clear to me that meaningful behavior is dependent not simply on real-time computation, but on the deployment of relevant sensors and sensor-data interpretation. This constitutes such artefacts as quasi-organisms, so biological and ecological metaphors are immediately applicable. From the perspective of the development of cultural practices utilizing these media, a novel territory of aesthetics is thereby opened up—the *aesthetics of behavior* and of behaving systems. Designing automated but contingent behavior is a new design realm and a new genre of aesthetics. The modeling and articulation of behavior itself is a practice to which both cybernetic and artificial life theorization are directly relevant. I referred to *Petit Mal*, *Fugitive I* and *II*, and *Traces* as “autonomous cultural agents” [22]. In *Petit Mal*, *Fugitive*, and *Traces*, part of the aesthetic-theoretical work is the exploration and articulation of embodied interaction, between a single human and an autonomous agent. See Figure 1.

2 Towards a Poetics of Interaction

In designing interactive artworks, one not only addresses aesthetic challenges related to formal attributes familiar in spatial, pictorial, cinematic, and literary forms, but the aesthetic and formal qualities of interaction itself, of situated and dynamical ongoing engagement. In particular, as interaction becomes increasingly embodied, the artist/designer is challenged to build or create a context that is fertile ground for integrating kinesthetic and proprioceptive sensibilities into evocative constellations of subjective sensorimotor experience.

Generally, in the world of computing, either interactivity is taken to be instrumental or it is broken. But in the arts, the challenge of interactivity has been to develop a *poetics of interaction*, a realm of multimodal analogy and provocative association that treads the delicate threshold between predictability and incoherence. This notion of poetics is key to an understanding of the aesthetics of interaction. Interaction can be deployed in an instrumental way, that is, as a device for moving from image to image, scene to scene, paragraph to paragraph, and so on. In this case, the dynamics of interaction itself recede from awareness, and “content” carries the content. In my work I have tried to consciously use the dynamics of interaction in order to engage the user in a way that is reflexively performative. Part of the “content” of the work is conveyed by the sense of “I am doing this now.” This aspect of interaction destroys the comfortable separation between subject and object that characterizes the conventional artwork-viewer relationship.

We should recall that interactivity itself, as well as machine systems with agentic behavior, was prototyped by the British cyberneticians Gordon Pask, Ross Ashby, and Grey Walter and in the Art-and-Technology movement of the 1960s [31]. It is important to note that the prevailing cybernetic-systems-theoretic approach of the time provided a different and more flexible paradigm for experimenting with the aesthetics of behaving systems than that of the more recent digital context. Two of many reasons for that are that the distinction between hardware and software was not established, and nor was the narrow idea of intelligence as reasoning on symbols. One might make the argument that, contra the rhetoric of “multimedia” in the 1990s, the desktop computer *closed down* the range of creative exploration in interactivity. This was for a combination of reasons—the limited i/o capacity of desktop computers and the individualist paradigm of “instrumental” interaction where the computer

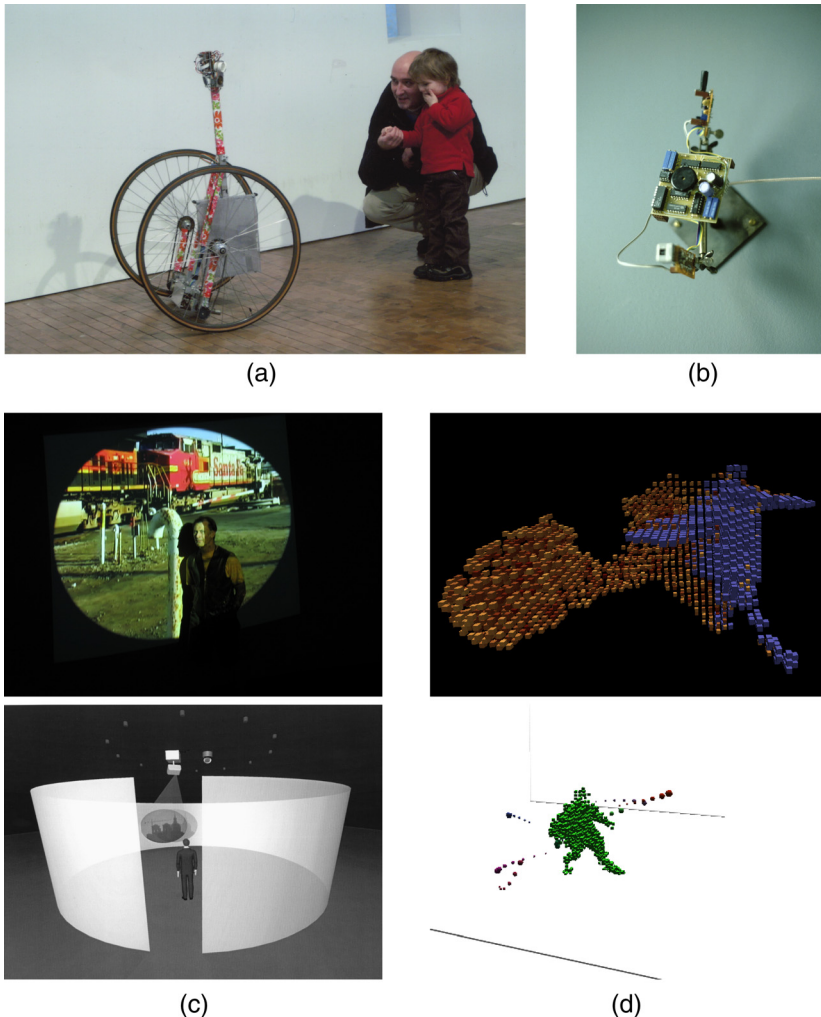


Figure 1. (a) *Petit Mal*, 1989–1995. Shown here at Smile Machines exhibition curated by Anne Marie Duguet at Transmediale 2006, Berlin. (b) *Sympathetic Sentience*, 1993–1997. One of 12 units of the second version (*Sympathetic Sentience II*, 1997). (c) *Fugitive*, 1996–2004. Shown here is an image from *Fugitive II*. Australian Center for the Moving Image, Melbourne, Australia, 2004. (d) *Traces*. Volumetric-machine-vision-based embodied interaction in the CAVE. Ars Electronica, Linz, Austria, 1999.

figured as an efficient work tool for single-user productivity. In comparison, the cybernetic notion was more broad and ecologically inspired—a paradigm of an agent adapting to an environment through iterative feedback loops. This conception is both temporally enactive and externalist, in opposition to the computationalist conception of both brain and computer functioning as symbol processors, “communicating” in a Shannonesque mode. In comparison with cybernetic thinking, internalist computationalism has minimal ability to deal with interaction per se, or embodied interaction in particular.

Embodied interaction presents radical challenges to conventional aesthetics, as it confounds the conventional subject-object divide so integral to the humanist worldview. The artwork is no longer at some conveniently remote distance from the “audience,” nor is it conveniently static. Not just kinetic, it is responsive. The experience of the artwork involves subjective bodily experience in concert with awareness of changing environment. As such, it collapses the conventional subject-object

binary that characterizes art consumption as much as it characterizes scientific practice. It is contingent, processual, and relational—one moves through the experience, and the behavior of the system is spatiotemporally coupled with the user. Its manifestations can be almost of the order of a bodily reflex—more like jumping out of the way of an oncoming car than like contemplating the beauty of a sunset. The mode of experience of embodied interaction is closer to the mode of normal bodily engagement with the world than is the conventional mode of engagement of artworks. Installation art, of which interactive installation art can be seen as a subset, occupies a middle ground.

From the perspective of artificial life, “interactivity” in the sense of human-computer interaction is a subset of agentic interaction behavior. Systems and subsystems can be agentic if they exhibit behaviors that are contingent upon the sensing or measurement of aspects of their world, and interpretation of those measurements. Such action may be with respect to a digital world or with respect to a physical world, or a combination; may or may not be “real time”; and may involve interaction with other similarly behaving entities or with qualitatively different entities or phenomena. In the case of interactivity, those other entities are often people. Agents exist in a hybrid ecology, usually comprising in-silico and in-vivo components.

3 Cognitivism and Embodied Experience

My experience making kinetic sculpture, installation, and performance provided me rich perspectives concerning the nature of interaction, particularly bodily interaction, and the importance of attention to all the sensorial, material, and spatial details of a work. Among technical and scientific workers, the understanding of what artists do tends to be naive and dated (undoubtedly the reverse is also true)—the usefulness of artists is often seen as to “make things pretty.” I had no interest in serving a decorative function.

My skills and sensibilities led me to diagnose a fundamental rift in implicit notions of intelligence and embodiment. It became clear to me that in the computer sciences, the separation of matter and information was an unquestioned assumption which grounded almost all practice in the field. As an artist, steeped in making practices, this idea was not nearly as self-evident to me as it appeared to be to my CS colleagues. I recall in the mid 1990s raising such issues with a cognitive scientist at Carnegie Mellon. I inquired about what seemed to be a general assumption that the brain was a computer. His response, which alarmed me, was “we don’t even question that, we know it’s the case.” Yet the physically grounded and performative nature of robotics, and the pragmatics of building robots that could navigate abandoned mines or the surface of Mars, implicitly challenges ideas of disembodied minds.

In my view, the values that underlie the discipline of computer science are Platonist and Cartesian, and paradigmatically enforce a hierarchical split between matter and information in the hardware-software duality (a direct analogy to the mind-body duality), from the computationalist paradigm in cognitive science and related disciplines to the conception of bodiless existence in virtual reality (VR) and cyberspace so clearly imagined by William Gibson and by Hans Moravec [16].

I was troubled by the way these ideas insinuated themselves, Trojan-horse-like, in the technology, leading to an implicit imposition of this dualism in many aspects of culture newly touched by computers. I rejected the imposition of this hierarchical dualism because it denied the significance of *bodily intelligences* upon which the making and enjoyment of art depend, and because it reflects a transcendentalist dream of liberation from the body, an idea I reject. I had the distinct sense that these were deep and serious issues, and was surprised that most artists and computer scientists seemed unaware of the great philosophical battle that is occurring beneath the everyday production of code.

4 Critical Technical Practice

Contrary to accepted rhetoric, computing is neither neutral nor absolute. Like any other human practice, it is historically contingent. I sensed the ever-present danger of being blindsided by the

incorporation of techniques that were based in ideas antithetical to those I was pursuing in my theoretically activist work, so I was concerned to explore and expose the assumptions that undergird computing and the culture surrounding it, and are built into hardware and software. Such a reflexive approach is what Philip Agre called *critical technical practice* [2].

The cognitivist conception of intelligence is as a process of logico-mathematical reasoning, located inside a black box separate from the world. This notion of thinking as symbol manipulation is expressed in the *physical symbol system hypothesis* of Newell and Simon (1976) [18], based in Fodor's functionalism (1978) [10]. The idea is axiomatic in what John Haugeland called "good old-fashioned AI" (GOFAI) (1995) [12]. The von Neumann machine automates logical operations, but there is no reason to imagine that (human) perception and action can be modeled appropriately as a von Neumann machine, with a linear process of input, processing, and output stages (analogous to industrial automation). According to this analogy, perception and action are reduced to the status of analogue-to-digital conversion "peripherals," as if perception were nothing but a routine translation from voltages to binary numbers.

Partly based on my experience as an artist, I found this idea simply wrong. Formed by such assumptions, standard computational practices often precluded, or made difficult to implement in principled ways, ideas of cognition that contest these assumptions. I wanted to implement systems that see cognition as closed-loop-embodied sensorimotor engagement with the world—a way of thinking about being in the world that, in my view, is a key aspect of practices in the arts. In the language of enactivist cognitive science, this is called "structural coupling," or "codependent arising," to use Varela's autopoietic terminology [35].

Over the last 25 years, embodied and situated cognitive science has emerged and undergone evolution, regrettably independent from the development of embodied interactive art, although the fields have much to share. The silo effect not only prevents researchers from sharing research, but generates disciplinary hubris. In the case of media art, the modes of presentation were physically located in institutional contexts with minimal crossovers with the academy—international media art events such as Ars Electronica, or international competitions such as VIDA. More importantly, this kind of work is ontologically different—performative as opposed to representational, with no aspiration to proof, truth, scientific fact, or original contributions to knowledge.

5 The Virtual

Through the 1990s, the notion of virtuality was central to emerging digital arts discourses. Immersive environments (meaning usually interactive stereoscopic multi-screen computer graphic environments) were an exciting topic in technical research and marketing rhetoric. The "virtual" meant several different things, including the realm of online or in-computer abstract information, as well as immersive environments. The general confusion about such matters is captured in various popular films of the time, for instance *The Lawnmower Man* (Leonard, 1992).

My previous art practice (sculpture, installation, etc.) brought to my design of robotic and interactive systems an understanding of the role of embodiment and bodily movement in the engagement of kinetic and spatialized aesthetic projects and an equally visceral awareness of the ways that experiences are built from the summation and interaction of diverse and multimodal components: the physical placement of components with respect to the scale of the body and range of movement; the effect of ambient light and ambient acoustics; the sensorial and associative qualities of specific materials. This sculptor's sense of the task of crafting the totality of a sensorial and sensorimotor experience permitted me rather different insights into the task of interaction design.

Perhaps as a result of my background in situated and embodied practices of sculpture, performance, and installation, I was never excited by the rhetoric of fantastical immersion in Virtual Reality nor by the fetishization of the stereo-visual spectacle and other manias for display systems of

increasing technological sophistication and presumed visual verisimilitude, which characterized the intersection of cinema and VR. One of the arguments for immersive interaction was that it was “embodied.” To me, this claim was asymmetrical. While as user you felt you were “in” a world, this feeling of being “in” went no further than the visual, generated by the integration of a stereoscopic viewpoint with head angle data. The system had no data about the rest of you. You might as well have been a giant slug. As I observed at the time, you checked your body at the door. *Fugitive* was in part a response to this situation, and in *Traces*, I set about to build a truly embodied volumetric machine vision system as a counterargument to this hollow rhetoric of embodied interaction. (See Appendix B (Seeing clearly - a critique of machine vision) in the online supplementary materials.)

Arising from my sculpture, installation, and performance sensibilities, *Petit Mal—an Autonomous Robotic Artwork* (1989–1995)—sought to move interaction off the desktop, out of the shutter glasses, and into the physically embodied and social world. The focus of my artistic practice in that period was on the bodily experience of the user in the context of behaving installations and on the construction of a fluid relation between bodily dynamics and technological effects. Theoretically, *Petit Mal* sat at a confluence of embodied art practice, artificial life, and the cognitivist crisis.

Petit Mal engaged visitors in large-scale bodily interaction—a dance. The goals of *Petit Mal*, apart from the obvious one of building an autonomous mobile robot that was an artwork, were:

- To build an autonomous human-scaled machine that was perceived as an active intelligence, but did not resort to anthropomorphism or zoomorphism (at least not in its form, though its *behavior* is zoomorphic. Leafing through an Edwards Scientific catalog recently, I saw any number of relatively simple mechanical toys called “robots” solely on account of the application of self-adhesive plastic googly eyes. This was precisely what I wanted to avoid.)
- To build a computational machine for which the interface was entirely gestural, bodily, and kinesthetic, in which there was no textual or iconic interface, no buttons or menus, no keyboards or mice, no screens or codes of flashing lights.
- To build a behaving machine that elicits play behavior among people. *Petit Mal* implements a non-instrumental kind of play that is quite incommensurable with the conventional computer-game logic of competition, numerical scoring, and “levels,” which has more to do with rationalized industrial labor than with play [25].
- To provide a working example of a situated and reactive robot, providing a physical and performative critique of conventional AI approaches to robot control and navigation.

Midway through this project I became aware that my research agenda, arising substantially out of art interests, was consistent with progressive thinking in robotics, cognitive science, and AI. I found that my intuition about behavior programming was consonant with the bottom-up and reactive robotics work of Brooks, Steels, and others [4–7]. I saw *Petit Mal*, technically, as a vindication of a *reactive* robotics strategy and a critique of conventional AI-based robotics, as well as an experiment in artificial sociality.

Motivation to interact seemed driven by curiosity. It was interesting to observe that people willingly and quickly adjusted their behavior and pacing to extract as much behavior from the device as possible, motivated entirely by pleasure and curiosity. (The only demographic who were unwilling to interact were adolescents.) Like *ELIZA* (Weizenbaum, 1966), *Petit Mal* often elicited assumptions that the thing was more clever than it really was. My emphasis on engagement of the user in a situated and embodied way was consistent with emerging critiques of AI. These critiques included Hubert Dreyfus’ *What Computers Can’t Do* [8, 9], John Searle’s *Chinese room* [33], and Stevan Harnad’s article on the “Symbol Grounding Problem” [11]. New ideas about embodied and situated cognition were coming to light in work such as Lucy Suchman’s *Plans and Situated Actions* [34], Varela, Thomson, and Rosch’s *Embodied Mind* [35], and the work of Edwin Hutchins [13] and David Kirsh [14].

The context in which this work developed is significant. I had already begun the project when I took up a cross-disciplinary position at Carnegie Mellon University as professor of art and robotics in 1993. I brought to that context my experience in installation, performance, and machine sculpture, along with substantial experience in designing performative technologies and persuasive sensorial experience—and more subtly, with predictions regarding the cloud of cultural associations that might be elicited by a particular set of cues, materials, gestures, and references.

The period of development of *Petit Mal* was crucial to the development of my understanding of the hardcore engineering realities of robotics and the development of my critique of cognitivism. I was fortunate to have had the opportunity to move in circles with leading roboticists and to come to terms first hand with the technical realities and motivations of robotics. I began to recognize that my experience in creating materially instantiated sensorially affective (art)work provided me with a different leverage on robotics from that of many in the Robotics Institute whose backgrounds were in computer science and engineering. When the term “socially intelligent agents” was abroad in AI circles in the late 1990s, I coined the term “culturally intelligent agents” and was bemused when “affective computing” became a buzz-phrase in that world.

Given the available technology of the time, and the unusual nature of the project, I had to engage with mechanics, electromechanics, and computational hardware and software at a low level. *Petit Mal* used a combination of ultrasonic and pyroelectric sensors to locate people. I designed and built my own sonar drive circuitry, pyroelectric sensor array, motor drive circuitry, brake system, and rotary encoders, each of which took weeks or months to design, as well as source components, prototype, and tests. I managed mechanical reliability, power budget, and charging techniques so that the device could function robustly with the public in a large environment for 10–12 hours a day (using low-cost components). This was a significant achievement for any robot at the time. Most research robots ran for shorter durations between downtime. (See Appendix C (A critique of plug-n-play), Appendix D (Auto-Pedagogic interface), and Appendix E (On migration, reconstruction, and obsolescence) in the online supplementary materials.)

6 *Sympathetic Sentience: Stigmergy Art*

The motivation for *Sympathetic Sentience* (1996) and *Sympathetic Sentience II* (1997) arose from an interest in emergent complex behavior and specifically in stigmergy, a key concept in the discourses around artificial life in the early 1990s [3]. My goal was to build a minimal physically instantiated system of multiple autonomous but communicating units that demonstrated emergent complex behavior, which I defined at the time as a manifestation of behavior more complex than that explicitly coded in the system. I conceived of a group of twelve autonomous communicating electronic entities that generated complex patterns of rhythmic sound out of, and only because of, the process of communication.

In the process of development, we (I and my collaborator Jamieson Schulte) found that a one-way serial loop of communication was sufficient to generate the complexity we sought. As in a game of Chinese whispers (also known as “telephone”), the units communicated in a serial loop, via a custom infrared signaling system. Each unit received, processed, and forwarded a continuous stream of data. At the outset, some random unit makes a chirp, which passes to the next via an infrared signal. At some point another unit chirps, and the data stream being passed consists of two chirps. This process continues, the rhythms slowly cycle around the group, increasing in complexity. Each device voiced the stream of data as it passed through the unit, occasionally adding or subtracting a chirp of its own. This generated a constantly changing but ordered rhythm and melody, and one could hear that melody repeated across the room, delayed and slightly changed by other units in the loop.

Alone, each unit could only generate a single chirp at regular but semi-random intervals, but in collaboration, this community of simple electronic organisms manifested emergent behavior. While

the serial loop persisted, the system and each unit demonstrated unpredictable, constantly changing but melodic and rhythmic behavior and was inherently self-stabilizing. After an initial startup period, the ongoing sound maintained about 50% saturation. That is, over any time period, the total duration of silence was always roughly equal to the total duration of sound. After an initial buildup period, the system was never silent, nor was it ever fully saturated. Due to the use of an XOR gate in each circuit, the varying density of the bits in the stream meant that a bit injected into the stream might be additive, or might erase or cancel out a bit, adding a space. The density of the sound was thus self-governing. In cybernetic terms, it was a homostat of sorts. The rhythm cycles were never constant but continually in development. This quasi-stability was a surprisingly complex (emergent?) consequence of a comparatively simple circuit comprising little but a simple astable pulse generator, an XOR gate, and a short time delay in each unit.

The system did have minimal, suppressive interactivity—a user could unknowingly block this chain of communication by moving through the space and thereby interrupting an IR beam, thus imposing a gap or silence in the melodic passage, which would then slowly fill up again. A longer interruption resulted in complete silencing of the whole group. When the blockage was removed, slowly a new rhythm would build up. The focal concern was not with human agency, but with the agency, autonomy, and creativity of a minimally complex system.

7 Machine Vision Works: *Fugitive* and *Traces*

My commitment to physically instantiated systems was tempered by the rigors of hardware and machine development, which necessitated the construction of new modules for each new process. Instituting a new behavior for *Petit Mal* might involve a month of cutting metal and soldering. The pull towards materiality is strong, but the flexibility of software is an undeniable pragmatic attraction. As a result I moved back and forth between physically instantiated works and image-based works. *Fugitive I* and *II* and *Traces* involve digital imagery and substantial software development.

These works pursued the creation of an aesthetically rich, fully embodied experience that users could and would interact with utterly intuitively. There were no hand-held input devices, no strap-on sensors, no suiting up with reflective baubles, no codified interaction procedures, and no pre-use training sessions. The systems were crafted to respond instantly to the normal bodily behavior of users within the interaction space. They presented an experience of technological immanence: Users had no immediate awareness that any computer technology was involved. There was no obvious task or goal. The works encouraged and rewarded exploratory play.

Much of the conventional VR/media art of the 1990s “represented” or “depicted” in a manner consistent with conventional (pre-computational) static and linear narrative forms. They told a story. Contrarily, *Fugitive* and *Traces* were centrally concerned with the *performative*, with users’ awareness of their ongoing bodily experience in the context of an aesthetic environment contrived to provoke certain kinds of explorations. The subject matter in these works was the users’ experience of their own ongoing bodily engagement with the system. It was not unusual that users emerged from their experience of such works sweating and panting, but smiling and laughing: a testament to their bodily engagement and desire to explore the full gamut of possibilities of interaction.

They were single-user systems. I aspired to the goal of multi-user sensing and response, but there were limitations at each interface. At the sensing interface, it was computationally challenging to reliably distinguish one user from another. But we were usually able to extract more information about the users’ behavior than we were able to represent back in ways that were comprehensible to the user, while at the same time permitting the kind of poetics that made these systems more than simply instrumental. Producing output that each user could readily identify as related to their own input was also a challenge without the relatively easy identification with avatars in some story space.

7.1 *Fugitive*

In 1996–1997 I developed *Fugitive*—a machine-vision-driven, motion-controlled digital video installation—at ZKM with my software engineer collaborator Andre Bernhardt. It was provoked in part by my critique of rhetoric around virtual reality, specifically the themes of embodiment, immersion, and navigation [20]. I set about to build an interactive, immersive, and navigable environment, (perhaps perversely) without the use of synthetic imagery. This deployment of projected video was entirely opposed to the cinematic conception of passive viewing of visual spectacle. My interest was on the developing dynamic between agent and user. Thus it was necessary to build a sensing system that responded to whole-body gestures, rather than the *xy* position of a single point, on the top of the head or on the dataglove.

A work by another interactive media artist, Miroslav Rogala, helped to focus my critique. In his *Lovers Leap* (1995), the user moved on a checkerboard floor flanked by two projection screens. Video sequences from an interactive CD ROM were triggered by the user's position on the grid. From a users' perspective, their precise grid location was far less significant than whether they jumped, sidled, or crawled there. Experiences with such interactive installations—whose modality of interaction was premised on quasi-objective spatial frameworks made tractable by hardware and software framed by such conceptions—confirmed my commitment to build sensor systems that responded to dynamics of bodily movement that were meaningful to the user. This work predated similar work in the AI community, such as Agre and Chapman's development of deictic programming [1].

In a later work (*Voices of Bedlam, 2001–2003*) I implemented a body-centric polar coordinate system for gesture response. A user's hand position or arm movement was not interpreted in terms of an *xy* world-space grid, but in terms of radii with respect to the vertical axis of the body. The data generated was consistent with the subject's proprioceptive distinction between expansive, fast or violent gestures that perturbed the equilibrium of the body, and slow or gentle gestures that did not.

Fugitive detected bodily movement and large gestures of a user in a 10-m circular space deploying a custom machine vision system (running on a pre-Pentium PC). It analyzed and responded to user behavior as process, not as raw position. I called this set of procedures the “mood analysis engine” in a wry comment on the tendency in the field to deploy titles and names that made dubious claims (“artificial intelligence,” “knowledge engineering,” etc.). This behavior drove the selection of video from a structured database of video and sent it to a motion-controlled video projector, which displayed the images in varying locations on the wall of the cylindrical room. As with all my work, all code and most of the hardware was custom.

Fugitive pursued the notion of a spatial embodied non-symbolic and autopedagogic interface. My goals in the interaction scheme of *Fugitive* were precisely to resist the tendency toward scopophilic focus on the image, and rather to draw the attention of the user to the temporal continuity of their own embodiment. This was in part motivated by a critique of VR. In conventional VR, the disembodied gaze had the ability to travel within a prestructured architectonic environment. In *Fugitive*, the subject is the subject. This presented novel design challenges, as one had to construct the images and the interaction to counter the normal assumption that when looking at an image, it is the image, rather than the looking, that is important. As I made explicit at the time, in *Fugitive* the continuity that structured the experience was the subjective temporal continuity of the users' embodiment, or more correctly their proprioceptive and kinesthetic awareness. In conventional VR environments, the user was disciplined by the interactive and architectonic order. In *Fugitive*, virtual worlds arose and collapsed on the basis of bodily continuity.

Fugitive was conceived as an agent, and the user-system dynamics were not unlike those in *Petit Mal*. The user was engaged in a one-to-one interaction. *Fugitive* read movements and responded with changes in images, and its behavior was manifested in the changes and movement of the image. It remains a curious but undeniable fact that interaction with *Petit Mal* is far more affective for most users; they speak of empathy. While users actively engaged and enjoyed *Fugitive*, I never heard of anyone identifying with the agent in an affectionate way, or even identifying it as something separate

from the environment. After much consideration, I can only explain this effect in terms of the discrete embodiment of *Petit Mal* and its animal-like quality of physical movement.

7.2 Spatiality and Temporality in Interaction

tis Cinna, I do know him by his gait

—William Shakespeare, *Julius Caesar*, Act 1, Scene 3.

As discussed above, conventional machine vision has been preoccupied with (static) images. But people (and animals) are extremely sensitive to temporal patterns, rhythms, and such, as McCulloch, Lettvin, Maturana, and Pitts made clear in their ground-breaking article “What the Frog’s Eye Tells the Frog’s Brain” (1959) [17]. My instinct with respect to machine vision for embodied interaction was similar to that of the Canadian artist David Rokeby with his *VNS* (*Very Nervous System*, 1988). Images and frames were simply the wrong structures to impose. Shape and gesture recognition were not only difficult and costly, but unnecessary. With limited processing power, we found it was more important to keep the temporal resolution high, even at the cost of spatial resolution.

This data was made more useful (and computation was simplified) when some short-term history was utilized. While developing *Fugitive*, I was faced with the question of how much user history to base the system response on. A basic premise of *Fugitive* was that instantaneous position is of minimal significance in orchestrating bodily interaction. Much more important is the recent history of embodied action, or muscular memory. *Fugitive* based its actions on an ongoing history of bodily behavior. But the question was: How much history makes sense to the user? If the machine responds to something the user has forgotten she’s done, the behavior seems random to the user. Empirically, I determined that around 2.75 seconds was an optimum period, that people seem to maintain relatively accurate memory of bodily movement for only that long (unless they are, for instance, a trained dancer). Basing the immediate action of the system at every time step on the past 2.5 seconds of user behavior seemed intuitive to the user; integration of data over longer periods did not.

David Rokeby notes that when he reduced the latency in his down to almost nothing, the *Very Nervous System* sense was that the actions of the system were integrated bodily with the user:

I once programmed “Very Nervous System” to respond very clearly as soon as it saw the slightest movement. In every instance, the system responded before I realized that I’d started moving. In fact, the system seemed to respond at the moment that I decided I would move. This delay in consciousness makes it possible for systems with high sampling rates and response speeds to slip under the user’s consciousness. At this point, the system and its responses are experienced in the same way that we experience our own body. The interactive system becomes integrated into our proprioceptive system—the same internal sensing system that defines our sense of being in our body, and establishes the relative position of our arms and legs to our “point of consciousness.” [32]

(See Appendix F (The CAVE and rhetorics of embodiment in VR) in the online supplementary materials.)

7.3 Traces

I have tried to imagine and develop systems that would performatively refute mind-body dualism. I wanted to build systems that responded to bodily behavior and that required no textual or iconic abstraction, so that the system came to meet the user in the user’s own world, a world of kinesthetic behavior such as those required to ride a bicycle or to walk on an icy footpath, and bodily intelligences such as the way we unconsciously interpret body language. The dualistic bias of the technology elided

these complex intelligences, because of the underlying commitment to dualism in computer science and because such intelligences resist codification in deterministic terms.

The goal in *Traces* was to build an immersive system in which a user—unencumbered by wires or sensors of any kind, or by a symbolic interface with keys or buttons or scroll wheels—could build virtual sculpture in real time, that is, the user, with large bodily gestures, could *dance a sculpture*. People had to move to interact with the system. They jumped, kicked, danced, and crawled on the floor. This in itself was proof of the eccentricity of *Traces*. As a result of my experience with my single-camera overhead vision system, I determined to develop a multi-camera volumetric vision system compatible with live video display. *Traces* involved the development of a custom real-time volumetric machine vision system for deployment in the CAVE virtual reality immersive environment (Ars Electronica, 1999). (A second iteration of *Fugitive* (*Fugitive II*, 2000–2004) deployed the *Traces* volumetric machine vision system.)

While *Traces* was built for the CAVE, it did not exploit any of the standard modalities of virtual reality interaction. In particular, there was no architectonic space clad with texture-mapped panels, and no weightless, bodiless travel at infinite speed to the virtual horizon. As a result, many of the standard VR authoring tools, or major aspects of them, were irrelevant to us. There was no navigation and no “world,” so there were no walls to texture-map. There was just the system responding immediately to the bodily behavior of the user. Ironically, these design choices meant the graphics in *Traces* had to be very basic. We were generating a lot of voxels, and even with the superfast Silicon Graphics Reality Engine, we were limited to graphical routines instantiated in hardware, which limited real-time graphics to simple geometries with simple color and light effects. Not only the authoring tools, but also the hardware had been optimized for the standard style of VR, based on navigation through texture-mapped geometries.

When it comes to making systems and objects, one of the key differences between scientists and artists is that for artists, the immediate experience is primary—it is persuasive or it is a failure. In science, experimental demonstrations stand as a proof of concept, and function as a pointer toward an abstract conclusion—the demonstration is not an end in itself. I approached the *Traces* project (and the works before it) as an installation artist, for whom the direct immediate bodily sensorial experience was central.

7.4 *Traces*—Behaviors

In *Traces*, the bodily behavior of the user generated real-time graphics in a static virtual space. The real-time body model of the user was a virtual brush that seeded algorithmic behaviors of voxels. The goal of *Traces* is precisely not to present a panoptic spectacle for the user, but, as in *Fugitive*, to turn the attention of the users back onto their own sense of embodiment through time. The movement of the user through the space left *traces*: volumetric and spatial-acoustic residues of user movement, which slowly decayed.

Users experienced three behaviors of increasing sophistication. This had the effect of teaching the user the dynamics of the system. The first behavior, the *drifting trace*, was a 3D analogue of time lapse photograph. A virtual wind blew through the space, and voxels of the body model drifted into the distance, becoming increasingly transparent. As time progressed, the traces became more active, and in the later stage of the user experience, autonomous entities that had complex behaviors of their own were spawned by the user.

The second behavior was a 3D cellular automaton. In this behavior, masses of voxels (the body model) propagated, percolated, and expired, like some sort of bubbling fermenting floating mass. The third behavior, *Chinese dragons*, involved a flock of autonomous agents (deploying extended Reynoldsian flocking behaviors). Entities of the flock were generated by the gestures of the user. Once generated, they wheeled around in a flock, and like pigeons in a park, were sometimes attracted to the user and sometimes repelled. To my knowledge, this part of the project was the first case in which 3D flocking agentic behavior was juxtaposed with embodied and immersive interaction, where the ALife agents interacted with the bodily gestures of the users.

After development periods at the GMD Bonn and at Carnegie Mellon, *Traces* and the *Traces Vision System* were shown at Ars Electronica in 1999 and received a Prix Ars Electronica honorable mention. I remain indebted to my four collaborators—Andre Bernhardt, Jeffrey Smith, Phoebe Sengers, and Jamie Schulte—for their intelligence and expertise and their conscientious dedication to the *Traces* project. (See Appendix G (*The Traces Vision System (TVS)*: the case for volumetric machine vision) and Appendix H ((Failed) Immersive Telematics) in the online supplementary materials.)

8 Conclusion

The period from the late 1980s to the early 2000s was a period of bewilderingly fast technological change. During this period, researchers of all stripes, including artists, explored emerging technologies. I grappled with the available technologies (and my limited technical education) to build systems for embodied interaction that were both intuitive and aesthetically rich, and that, by their existence and functioning, proposed viable alternatives to conventional ideas about computing and cognition. My writing at the time emerged out of this combination of at-the-coal-face practice in this exciting discursive environment (see, for instance, [19–25]). Central to my thinking was a notion of agency, based in artificial life conceptions. The deployment of artificial life ideas in my work was not casual. My occupation with artificial life techniques and post-computationalist conceptions of cognition emerged as reactions to my critique of conventional GOFAI ideas. The motivations for making these works were in large part informed by a skepticism regarding the system of values that informed the cultures of computing and computer use, and to a large extent still do (see [26–29]).¹

As some anonymous sage observed, “the difference between theory and practice is greater in practice than in theory.” It is one thing to proclaim ideas regarding agency, autonomy, embodiment, or an aesthetics of behavior. It is quite another to build working systems, especially if the technologies you are working with imply or inhere contrary value systems. The excavation of these value systems hidden deep in the architectures is challenging, not least because experts in the field may be unaware of those qualities (conforming as they do to the worldview of the discipline). This intellectual work, in parallel with technical R&D, is what Philip Agre called *critical technical practice* [2]. In my opinion, such practice is often missing from technical education and research, and yet pursuing research “against the grain” keeps a discipline lively and opens up new realms of research, ultimately contributing to the substantiality and coherence of the field.

In this essay I have reported the experience of such a process. In the same way that the works themselves have to be experienced, in the process of making them, the hands-on negotiation of the tools and procedures cannot be elided. The resolution of the larger vision only occurs through the manipulation of technical minutiae, and the larger vision informs that process, creating criteria for utilizing one process or tool as opposed to another.

The negotiation and combination of emerging technologies, emerging techno-scientific ideas, and emerging cultural practices is an interdisciplinary exercise of some complexity. It seems important to emphasize that in addition to the obvious tangible achievement of such exercises—namely, making a new thing or new knowledge that was not possible within either discipline—there is a less tangible but more profound result: Such interdisciplinary projects expose and interrogate the assumption and

¹ It is gratifying that one of the anonymous reviewers of this article had this to say: “This is a historical survey of interactive installations and robotic artworks by the author, inspired and informed by issues within Cybernetics and Artificial Life. It focuses on the 1990s. Tension between different ideologies, between disembodied cognitive notions widespread in many cognitive science communities and alternative viewpoints that stress the centrality of embodied experience, forms a central theme to this history. The picture presented rings absolutely true to those, like this reviewer, who saw those tensions at that time. And exactly the same issues are still very live through to the present day, so this historical survey has up-to-the-moment relevance.”

values of disciplines in unexpected ways. This, in my opinion, has thoroughgoing epistemological importance.

Making artworks, at least of the kind I have made, is an exercise in performative technology and more generally an exercise in what some would describe as a non-modern ontology [15]. The general mode of scientific evidence and demonstration is representational—it seeks to expose facts as document or record. The works presented here are not conceived as didactic or authoritative, but convey what they have to offer through bodily experience, through a “dance of agency,” as Andy Pickering would say [30]. And this is, in large part, what the works are about.

References

1. Agre, P. E., & Chapman, D. (1987). Pengi: An implementation of a theory of activity. In *Proceedings AAAI'87* (pp. 268–272). Seattle: AAAI Press.
2. Agre, P. E. (1997). Towards a critical technical practice: Lessons learned in trying to reform AI. In G. C. Bowker, S. L. Star, W. Turner, & L. Gasser (Eds.), *Social science, technical systems, and cooperative work: Beyond the great divide* (pp. 131–158). Hillsdale, NJ: Lawrence Erlbaum Associates.
3. Beckers, R., Holland, O. E., & Deneubourg, J. L. (1994). From local actions to global tasks: Stigmergy and collective robotics. In R. A. Brooks & P. Maes (Eds.), *Artificial life IV, Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*. Cambridge, MA: MIT Press.
4. Brooks, R. A. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1), 14–23.
5. Brooks, R. A. (1990). Elephants don't play chess. *Robotics and Autonomous Systems*, 6, 3–15.
6. Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence Journal*, 47, 139–159.
7. Brooks, R. A. (1991). *Intelligence without reason* (A.I. Memo No. 1293). Massachusetts Institute of Technology, Artificial Intelligence Laboratory.
8. Dreyfus, H. L. (1972). *What computers can't do: A critique of artificial reason*. New York: Harper & Row.
9. Dreyfus, H. L. (1979). *What computers still can't do: A critique of artificial reason*. Cambridge, MA: MIT Press.
10. Fodor, J. (1978). *RePresentations. Philosophical essays on the foundations of cognitive science*. Cambridge, MA: MIT Press.
11. Harnad, S. (1990). The symbol grounding problem. *Physica D*, 42, 335–346.
12. Haugeland, J. (1985). *Artificial intelligence: The very idea*. Cambridge, MA: Bradford Books, MIT Press.
13. Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.
14. Kirsh, D., & Maglio, P. (1995). On distinguishing epistemic from pragmatic actions. *Cognitive Science*, 18, 513–549.
15. Latour, B. (1993). *We have never been modern*. Cambridge, MA: Harvard University Press.
16. Moravec, H. (1988). *Mind children*. Cambridge, MA: Harvard University Press.
17. McCulloch, W., Lettvin, J., Maturana, H., & Pitts, W. (1959). What the frog's eye tells the frog's brain. *Proceedings of the IRE*, 47(11), 1940–1951.
18. Newell, A., & Simon, H. A. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of the ACM*, 19(3), 113–126.
19. Penny, S. (1987). Simulation, digitization, interaction: The impact of computing on the arts. *Artlink* (Adelaide, South Australia) 7(3,4), Art+Tech issue.
20. Penny, S. (1994). VR as the end of the Enlightenment project. In G. Bender & T. Druckrey (Eds.), *Culture on the brink: Ideologies of technology*. Seattle: Bay Press.
21. Penny, S. (1996). The Darwin machine: Artificial life and interactive art. *New Formations UK*, 29, 59–68.
22. Penny, S. (1998). Embodied cultural agents and synthetic sociology. In *AGENTS '98, proceedings of the Second International Conference on Autonomous Agents* (pp. 463–464). New York: ACM Press, Addison-Wesley.

23. Penny, S. (1999). Agents as artworks and agent design as artistic practice. In K. Dautenhahn (Ed.), *Human cognition and social agent technology* (pp. 395–414). Amsterdam: John Benjamins.
24. Penny, S., Smith, J., Sengers, P., Bernhardt, A., & Schulte, J. (2001). Traces: Embodied immersive interaction with semi autonomous avatars. *Convergence: The Journal of Research into New Media Technologies*, 7(2), 47–65.
25. Penny, S. (2004). Representation, enaction and the ethics of simulation. In P. Harrigan & N. Wardrip-Fruin (Eds.), *First Person*. Cambridge, MA: MIT Press.
26. Penny, S. (2010). Artificial life art: A primer. In *Proceedings, DAC09 (Digital Art and Culture 2009)*. California Digital Library. Available at http://escholarship.org/uc/ace_dac09 (accessed December 2014).
27. Penny, S. (2010). Twenty years of artificial life. *Digital Creativity*, 21(3), 197–204.
28. Penny, S. (2011). Sixty years of robotic art. *AI and Society*, 28. Available at <http://www.springerlink.com/openurl.asp?genre=article&id=doi:10.1007/s00146-012-0404-4> (accessed December 2014).
29. Penny, S. (2012). Art and artificial life, performativity and process: An intellectual genealogy of a heterogeneous field. In *VIDA 13*. Madrid: Telefonica Foundation.
30. Pickering, A. (1995). *The mangle of practice: Time, agency, and science*. Chicago: University of Chicago Press.
31. Pickering, A. (2010). *The cybernetic brain: Sketches of another future*. Chicago: University of Chicago Press.
32. Rokeby, D. (1998). The construction of experience: Interface as content. In C. Dodsworth, Jr. (Ed.), *Digital illusion: Entertaining the future with high technology* (pp. 27–47). New York: ACM Press, Addison-Wesley.
33. Searle, J. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3), 417–457.
34. Suchman, L. A. (1987). *Plans and situated actions: The problem of human-machine communication* (2nd ed.). Cambridge, UK: Cambridge University Press.
35. Varela, F., Thompson, E., & Rosch, E. (1992). *The embodied mind: Cognitive science and human experience*. Cambridge, MA: MIT Press.