VIDEO, MUSIC, AND SOUND METACREATION

ARNE EIGENFELDT
School for the Contemporary Arts, Simon Fraser University
Vancouver, Canada
arne_e@sfu.ca

MILES THOROGOOD
School for the Contemporary Arts, Simon Fraser University
Vancouver, Canada
mthorogo@sfu.ca

JIM BIZZOCCHI
School for the Contemporary Arts, Simon Fraser University
Vancouver, Canada
jimbiz@sfu.ca

PHILIPPE PASQUIER
School for the Contemporary Arts, Simon Fraser University
Vancouver, Canada
pasquier@sfu.ca

TOM CALVERT
School for the Contemporary Arts, Simon Fraser University
Vancouver, Canada
tom@sfu.ca

Keywords: Generative Video, Generative Music, Sound Art, Metacreation

Metacreation is a contemporary approach to generative art, using tools and techniques from artificial intelligence, artificial life, and machine learning (themselves inspired by cognitive and life sciences) to develop software that is creative on its own. We present examples of metacreative art within the fields of music, sound art, the history of generative narrative, and discuss the open potential of the "open-documentary" as an immediate goal of metacreative video.

2014.xCoAx.ORG PORTO, PORTUGAL COMPUTATION, COMUNICATION, AESTHETICS AND X

1. INTRODUCTION

Generative Art has a long tradition, one that dates back to antiquity, according to Galanter (2003). While the potential of codifying artistic decisions may be alluring to many artists, the challenges are many: for example, can the notion of creativity be extended to machines, or can they (should they?) only remain as tools for the creative artist? The nascent field of computational creativity, also known as metacreation, explores these questions, and is populated by psychologists, art theorists, cognitive scientists, artificial intelligence researchers, machine learning specialists, and, perhaps most importantly, artists. As such, it is not merely an academic pursuit: it is, and has been, a fertile creative domain for artists exploring new avenues of production. This paper will explore three such directions – music, sound, video – and provide some examples, many of which are drawn from the authors' work, as proof of its existence. Lastly, we will describe a current project which aims to blend the three formerly disparate media into a single, cohesive artistic medium.

2. METACREATION

Metacreation is the idea of endowing machines with creative behavior (Whitelaw 2004). Metacreation is a contemporary approach to generative art, using tools and techniques from artificial intelligence, artificial life, and machine learning (themselves inspired by cognitive and life sciences) to develop software that is creative on its own. In other words, software is a metacreation if it exhibits behaviors that would be considered creative if performed by humans.

Artists use tools to produce their artwork. Traditionally, the creator of the tool and the artist using the tool has remained distinct; with digital tools, a growing number of tech-savvy artists design and develop the software tools with which they produce their works. By developing tools with *ad hoc* functionalities, these artist/engineers aim to gain more control over their creative processes. What if these artist/scientists could develop tools that do not need a user to create the finished artworks? The notion that the creation of the computer-tool can take precedence over its utilization is at the root of generative arts.

From a research standpoint, the question regarding metacreation is no longer "can a computer exhibit creative behavior?": that question has been answered, in the positive, many times over. AARON's (Cohen 1995) paintings have exhibited at the Tate Gallery in London, and EMI's compositions (Cope 1991) have created music that could not be discerned from human-composed music by the most educated experts (see Section 3 for more examples of metacreations). The question can now be reframed as "how can artists and scientists collaborate to define, structure, and explore the boundaries of this relatively recent multidisciplinary field?".

Two types of approaches are possible for modelling creative behaviour in metacreation research. One can model systems that produce creative behavior in which the system is a "black box", and only its behavior (i.e. its output) matters. This results in processes that explore creativity as it *could be*, rather than model creativity *as it is*. Moreover, although relatively successful, these systems do not mimic humans in the way they operate. It is clear, for example, that a human improviser does not maintain transition probability tables when playing, as is the case with the *Continuator* and its Markov Model (Pachet 2003), and human composers do not evolve populations of musical gestures and simulate their natural evolution, as *GenDash* does (Waschka 2007).

One can also try to model systems that will be creative using the same processes that humans are thought to use. This approach has been relatively unexplored, mostly because these processes are largely unknown. One would have to address more deeply the question of what human creativity is, and produce models that are believable, not only in terms of their output, but in terms of their internal processes. Our group has been pioneering some early attempts in this regard (Maxwell *et al.* 2012), by starting to bridge the gap between the literature in cognitive science, musical perception and cognition, and generative systems.

3. EXAMPLES OF METACREATIVE ART

Metacreative art is the artifact produced by systems, arising from the implementation of specific models of creativity and creative process. These machine-generated artefacts have been used to observe the validity of the model under investigation, and, often, been positioned within cultural contexts such as performance and exhibition venues. The following examples of metacreative art demonstrate the diversity of approaches that researchers have employed in modelling creative behavior in the domains of music, sound art, and moving image.

3.1. METACREATIVE MUSIC

Music has had a long history of applying generative methods to composition, due in large part to the explicit rules involved in its production. A standard early reference is the *Musikalsches Würfelspiel* of 1792, often attributed to Mozart, in which pre-composed musical sections were assembled by the user based upon rolls of the dice (Ihmels and Riedel 2007); however, the "Canonic" compositions of the late 15th century are even earlier examples of procedural composition (Randel 2003).

Exploring generative methods with computers began with some of the first applications of computers in the arts. Hiller's *Illiac Suite* of 1956 utilized Markov chains for the generation of melodic sequences (Hiller and Isaacson 1979). In the next forty years, a wide variety of approaches were investigated – see (Papadopoulos and Wiggins 1999) for a good overview of early uses of computers within algorithmic composition. However, as the authors suggest, "most of these systems deal with algorithmic composition as a problem solving task rather than a creative and meaningful process". Since that time, this separation has continued: with a few exceptions (Cope 1991, Waschka 2007), contemporary algorithmic systems that employ AI methods remain experimental, rather than generating complete and successful musical compositions.

An approach followed by Eigenfeldt in *Kinetic Engine* (Eigenfeldt 2008) was to model the interaction of human improvisors within a drum ensemble through the use of virtual agents. Player agents assume roles and personalities within the ensemble, and communicate with one another to create complex rhythmic interactions. The software was used in performance many times, controlling the robotic percussion instrument *MahaDevibot* (Kapur *et al* 2009), in which the composer acted as a "conductor", directing the virtual agents in response to other live performers (Eigenfeldt and Bahn 2009).

The notion of modelling a software improvising system after human activity was posited by Rowe (1992): "interactive software simulates intelligent behaviour by modeling human hearing, understanding, and response"; however, *Kinetic Engine* is modelled after human *interaction* using the AI paradigm of multi-agents. Intelligent agents are elements of code that operate without direct user interaction (they are autonomous), interact with one another (they are social), interact with their environment (they are reactive), and make decisions as to when they should operate, and what they should do (they are proactive) (Wooldridge 2009). Since these are also attributes required of musicians in improvisational settings, the use of agents to emulate human–performer interaction has proven to be a fertile field of research. Whalley (2009) gives an overview of the recent state of software agents in music and sound art.

Multi-agents were the basis of a series of compositions entitled *Coming Together*. In these systems, agents negotiate musical content within a defined musical environment, with or without direct performer interaction. In each case, agents begin with random musical material, and through the convergence of predefined musical parameters, self-organisation is demonstrated (see Eigenfeldt 2011 for a detailed description of the series). The interaction between virtual agents and humans was explored in *More Than Four* (Eigenfeldt 2012), which also incorporated a curator agent to create complete compositions for performance from a database of pre-generated movements (Eigenfeldt and Pasquier 2012).

3.2. METACREATIVE SOUND ART

Sound art does not have general representation schema, or an established theory like symbolic forms of music. Thus, sound art has been a more difficult field of study due to the shortage of well defined models, and has not been as frequently explored in metacreation research. Sound art is an interdisciplinary practice that is based on acoustics, psychoacoustics, and music principles, but then often contracts knowledge from a diverse range of other fields; acoustic design (Truax 1998), genomics (Fargher and Narushima 2008), or social media (Roma *et al*. 2009), for example. Although sound art may be reified with a physical object (Bandt 2001), or as tape music, it is typical that a work can be positioned along a spectrum of non-symbolic electroacoustic music. This spectrum includes purely electronically generated sound works on one end, whilst on the other we find works of so called found-sound: concrete recordings aimed to evoke in listeners associations of a real time and place.

The aim of sound art is to transform a concept, devised by the artist, into a set of decisions and processes that will ultimately result in an acoustic work. For example, Philipsz's 2010 soundscape piece *Lowlands* (Philipsz 2010) combines abstracted sounds of the human voice, accompanied with the ambient sounds of modern cities to initiate particular experiential events for the listener. Although the acoustic work is an outcome of sound art, there is, however, no well defined model for achieving it or agreed upon objective function for evaluating its performance.

The want of these constraints has prompted multiple designs of metacreative systems to address the questions of what processes should be used in order to arrive at the sound work, and how to evaluate the work that the system produces. Eigenfeldt and Pasquier (2011) populate a database of sound recordings for the retrieval by autonomous software agents. These artificial agents select recordings based upon semantic tags and spectral audio features, and mix them using a restricted set of digital signal processing techniques. In that work, the concept is established by the domain of tags set by the composer, and the selection criteria employed by agents.

Olofsson (2013) also takes an agent-based approach to generate sound content in his low life series of works. The behaviours of the agents in this work are constituted on rules manifesting from the audio synthesis code they reference. He calls the agents in his "self-referential code" system "audiovisual creatures", which engenders a performative quality to the artificial system. A further example of anthropomorphic agents is demonstrated by Thorogood's artificial life installation *Chatter and Listening* (Thorogood 2007). In this work, behaviour of a bird species is modelled, equipping multiple interacting robots with bird-like characteristics, which then produce a synthesized vocalization based on their behaviour state.

Another approach to metacreative sound art is the use of knowledge representation systems, which aim to model a particular knowledge base in a domain. An example of this type of system is the *ec(h)o* interactive audio museum guide by Hatala, Wakkary, and Kalantari (2005). The authors describe a formal representation for sound objects that address sound content properties, concepts, topics, and themes, including connection to aspects of the exhibition. Their system then updates a visitors display from the input of user data, including physical position of the user, the history of interaction with objects and space, and interests that the user exhibits.

Audio Metaphor (Thorogood and Pasquier 2013) is a system for the generation of sound art from a short text input. The system retrieves labelled audio recordings that have semantic relationships to the input text. These audio recordings are autonomously segmented in response to analysis by a supervised machine-learning algorithm, trained with data from human perceptual classification experiments. The semantic and saliency labelled segments are then processed and combined autonomously based on a composition schema, modelled after production notes from Canadian composer Barry Truax (2008).

3.3. METACREATION, NARRATIVE, AND VIDEO

Some of the first examples of metacreation appeared within the domain of visual art: Romero and Machado (2008) present an overview of many of these systems. While some of the described artworks contain dynamic change, and may border on video – for example, Scott Draves' *Electric Sheep* (2008) – generative video is in a more nascent stage.

Bizzocchi (2011) describes a system entitled *Re:Cycle* which he describes as a generative video engine. The system incorporates a variation on the three aesthetic strategies of his earlier linear video works: strong imagery, manipulation of time base, and carefully designed visual transitions. However, the computationally generative *Re:Cycle* system develops a recombinant aesthetic for the sequencing of shots and transitions drawn from incorporated databases. The system uses meta-tags to nuance randomized selection with a sense of visual flow and coherence.

Traditional video has a long connection with storytelling as a dominant mode. This connection is more problematic with interactive computational video. However, the potential for a sense of "narrativity" rather than traditional "storytelling" is possible within a more open computational approach. Bizzocchi claims that the expressive presentation of character, storyworld, emotion, and narrative theme, as well as a degree of localized "micro-narrative" plot coherence can produce a narrativized experience without the traditional reliance on a full-blown narrative arc (Bizzocchi 2007). We believe that it is also possible to use computationally generative techniques to combine shots, tags, sound and sequencing within a narrativized metacreation aesthetic.

There is a substantial history of writers and artists working across the "narrativity to storytelling" spectrum. Non-digital examples of generative narrativity include a variety of dada and surrealist narrative games from the *Exquisite Corpse* (Gooding 1995) to Burroughs "cut-ups" (Burroughs and Gysin 1978). The most extensive exploration of analog generative narrative is probably found in the Oulipo creators (Wardrip-Fruin and Montfort 2003) and in their digitally-oriented successor groups: Alamo, LAIRE, and Transitoire Observable (Bootz 2012).

A number of digital works link knowingly to this literary tradition of generative and recombinant narrativity. Hayles claims that "Generative art… is currently one of the most innovative and robust categories of electronic literature" (Hayles 2007). Andrews and Wardrip-Fruin explicitly recognize their own extensions of Burroughs's cut-up aesthetic in the works *On Lionel Kearns*, *Regime Change*, and *News Reader* (Wardrip-Fruin 2005). Bill Seaman's installation work *The World Generator* (2002) uses "images, sound, and spoken text to create a recombinant poetics that created emergent and synergistic combinations of all these modalities" (Hayles 2007).

Many contemporary works that rely on generative computation include an explicit commitment to more traditional storytelling. Expressive developments in generative digital narrative works can be seen in contemporary projects such as *Curveship* (Montfort 2009), *Mexica-impro* (Perez *et al*. 2011), *Soft Cinema* (Manovich and Kratky 2002), or the series of works by Harrell (2007). Montfort's *Curveship* systematically modifies storytelling modalities (such as voice, style, focalization) in narrative constructions. Perez's *Mexica* uses a computational cycle of "story generation" and "reflection" to systematically move a narrative to its conclusion. Manovich's *Soft Cinema* video artwork uses database and algorithm to build a recombinant cinema aesthetic. Harrell has designed generative systems based on shuffling text and image to build a series of expressive and emotionally evocative narrative systems: GRIOT, GENIE, and Renku (Harrell and Chow 2008). Montfort, Perez, Harrell, and Campana are currently developing *Slant*, an integrated system capable of generative storytelling (Montfort *et al*. 2013).

4. FUTURE DIRECTIONS

From what we have seen in the literature and practices outlined here, research in regard to metacreative art is concentrated within individual domains rather than across media forms. We propose to integrate our work to explore its metacreative multi-mediated potential. We will do this through the development of generative video systems that are fully integrated with sound and music metacreation.

Clearly, there are several difficulties with which we are faced, perhaps the greatest being that many of the tools used in music and sound analysis do not translate easily into video. Techniques such as recombinance work well in these domains when there is some understanding of the material; within audio, this can be derived from tools found within music information retrieval (Tzanetakis and Cook 2000). While methods of meta-tags, already used by Bizzocchi, could be extended with further associative descriptions, the machine learning described in Thorogood's *Audio Metaphor*, has as of yet to be written for video analysis.

We propose to explore the potential of the "open documentary" through extension of *Re:Cycle* through enhanced "narrativity", relying on variables identified in analysis of interactive narrative (Bizzocchi and Tanenbaum 2012). Work has begun in combining Bizzocchi's generative video system with Thorogood's *Audio Metaphor*, in which each video clip is provided with a descriptive, metaphorical, and contrapuntal commentary which can be used by the audio system to provide a complementary soundworld.

Further, we have begun research toward a system of generative sound design. Leveraging the successes from *Audio Metaphor*, this new system analyzes sentences and systematically selects and segments sound files. Using a state of the art planning algorithm, composition plans are generated and evaluated based on existing principles of sound design (Murch 2007). This research has already shown encouraging directions for generative sound design. We see that the ambient video generation and this new development as a promising avenue for further investigation.

Popular forms of narrative – such as mainstream cinema and novels – typically rely on the complete commitment to the narrative arc as the backbone and the engine for the storytelling experience. Other narrative forms, however, such as video games, song lyrics, television commercials, and the long history of generative narrative art show that narrative can follow other paths. The potential for "narrativity" exists in the design and presentation of character, storyworld, emotional tenor, and thematic sequencing. The ongoing development of micro-narratives and associated moments of narrative coherence in a generative system can do the work of the unitary narrative arc of a more traditional form.

Our work may ultimately approach a more complete commitment to unitary storytelling and the metacreation of a tight narrative progression. This is a much higher order problem to solve – one that may or not be attainable in the context of our current project. A generative and recombinant storytelling system implies significant control over the details of plot sequencing, narrative arc, and narrative closure. This, in turn, will require much higher standards for computation, metadata tagging, and shot selection. The commitment to a loose "narrativity" in an "open documentary" context is a far more reachable intermediate goal. Progressive development of the documentary system will inch towards ever-increasing narrativity. It will be interesting to see how closely it approaches - and whether it ultimately realizes - the narrative coherence of a true storytelling system.

5. CONCLUSION

In general, there is a continuum between traditional praxis or performance tools, and metacreations. At one end, the software simply acts as a tool to be manipulated by the creator: the artist or composer has to do most, if not all, of the creative work, by manipulating the various functionalities of the tool. On the other extreme, pure metacreations are autonomous and proactive in their creative choices, and require no human intervention once running (although human intervention is still needed at design time). Interactive systems that allow for a constructive dialogue, or interaction between the system and its user, are situated in the middle.

We have described several successful metacreations within music and sound art, and noted the dearth of metacreative video; however, the history of generative narrative demonstrates a potential for a true metacreation in this medium. The first step will be the exploration of the "open-documentary".

ACKNOWLEDGEMENTS

This research was funded by grants from the Social Sciences and Humanities Research Council of Canada and by the GRAND Network of Centres of Excellence.

REFERENCES

Bandt, R. *Sound Sculpture: Intersections in Sound and Sculpture in Australian Artworks.* Craftsman House, Fine Art Publishing Pty Limited, 2001.

Bizzocchi, J., *Games and Narrative: An Analytical Framework.* Loading: the Journal of the Canadian Game Studies Association, 1(1), 2007

———. *Re: Cycle – A Generative Ambient Video Engine.* Entertainment Computing – ICEC, Springer Berlin Heidelberg, 2011.

Bizzocchi, J., and Tanenbaum, J. *Mass Effect 2 A Case Study in the Design of Game Narrative.* Bulletin of Science, Technology & Society 32.5, 2012.

Bootz, P. *From OULIPO to Transitoire Observable The Evolution of French Digital Poetry.* http://www.dichtung-digital.de/en/journal/archiv/?postID=248, 2012, retrieved January 14, 2014.

Burroughs, W. S., and Gysin, B. *The Third Mind.* Viking Press, 1978.

Cohen, H. *The further exploits of AARON, painter.* Stanford Humanities Review 4:2, 1995.

Cope, D. *Computers and Musical Style.* Madison, WI: A-R Editions, 1991.

Draves, S. *Evolution and Collective Intelligence of the Electric Sheep.* The Art of Artificial Evolution. Springer Berlin Heidelberg, 2008.

Eigenfeldt, A. *Emergent rhythms through multi-agency in Max/MSP.* Computer Music Modeling and Retrieval. Sense of Sounds. Springer Berlin Heidelberg, 2008.

———. *Real-time composition as performance ecosystem.* Organised Sound 16.2, 2011.

———. *More Than Four.* Performed at Goldcorp Centre for the Arts, Vancouver, May 25 2012. http://youtu.be/ao_B0eW3gDI

Eigenfeldt, A., and Bahn, C. *Bhatik.* Performed at Western Front, Vancouver, April 13 2009. http://youtu.be/R-DQ5WPd4ao

Eigenfeldt, A,. and Pasquier, P. *Negotiated Content: Generative Soundscape Composition by Autonomous Musical Agents in Coming Together: Freesound.* Proceedings of the Second International Conference on Computational Creativity, 2011.

———. *Creative Agents, Curatorial Agents, and Human-Agent Interaction in Coming Together.* Proceedings of Sound and Music Computing, Copenhagen, 2012.

Fargher, C, and Narushima, T. *Knit Two Together: Art/Science Collaborations in BioHome. The Chromosome Knitting Project.* The Scientific Imaginary in Visual Culture 5: 181, 2008.

Galanter, P. *What is Generative Art? Complexity theory as a context for art theory.* GA2003–6th Generative Art Conference, 2003.

Gooding, M. *A book of surrealist games.* Boston, MA: Shambhala Redstone Editions, 1995.

Harrell, F. *GRIOT's Tales of Haints and Seraphs: A Computational Narrative Generation System. Second Person: Role-Playing and Story in Games and Playable Media.* MIT Press, Cambridge, 2007.

Harrell, F., and Chow, K. *Generative visual renku: linked poetry generation with the griot system.* Proceedings, electronic literature organization conference, Vancouver, 2008.

Hatala, M., Wakkary, R., and Kalantari L. *Rules and Ontologies in Support of Real-Time Ubiquitous Application.* Journal of Web Semantics 3, 2005.

Hayles, K. *Electronic Literature: What is it?* http://eliterature.org/pad/elp.html, 2007, retreived January 14, 2014.

Hiller, L., and Isaacson, L. *Experimental Music; Composition with an Electronic Computer.* Greenwood Publishing Group Inc. Westport, CT, USA, 1979.

Ihmels, T., and Riedel, J. *Die Methodik der generativen Kunst.* 2007.

Kapur, A., Eigenfeldt, A., Bahn, C., and Schloss, W. *Collaborative composition for musical robots.* Journal of Science and Technology of the Arts, 1(1), 2009.

Manovich, L., and Kratky, A. *Soft cinema.* ZKM Center for Art and Media Karlsruhe, 2002.

Maxwell, J., Eigenfeldt, A., Pasquier, P., Gonzalez, N. *MusiCOG: A Cognitive Architecture for Music Learning and Generation.* Proceedings of the 9th Sound and Music Computing Conference, 2012.

Montfort, N. *Curveship: an interactive fiction system for interactive narrating.* Proceedings of the Workshop on Computational Approaches to Linguistic Creativity. Association for Computational Linguistics, 2009.
Montfort, N., y Pérez, R. P., Harrell, D. F., and Campana, A. *Slant: A Blackboard System to Generate Plot, Figuration, and Narrative Discourse Aspects of Stories.* Proceedings of the Fourth International Conference on Computational Creativity, Sydney, 2013.
Murch, W. *Dense clarity–clear density.* Ps1. org (Museum of Modern Art) 10, 2007.
Olofsson, F. *low life*. Musical Metacreation Weekend, Sydney (Old School Building, University of Sydney). Presented in association with ISEA 2013, curated by Oliver Bown. Sydney, Australia, 15-16 June 2013.
Pachet, F. *The continuator: Musical interaction with style.* Journal of New Music Research 32:3, 2003.
Papadopoulos, G., and Wiggins, G. *AI methods for algorithmic composition: A survey, a critical view and future prospects.* AISB Symposium on Musical Creativity, Edinburgh, UK, 1999.
Perez y Perez, R., Ortiz, O., Luna, W. A., Negrete, S., Peñaloza, E., Castellanos, V., and Ávila, R. *A System for Evaluating Novelty in Computer Generated Narratives.* Proceedings of the Second International Conference on Computational Creativity, México City, 2011.
Philipsz, S. *Lowlands.* Tate, Britain (2010 Turner Prize winner). Curated by Helen Little and Katharine Stout, London, England, Oct, 2010
Randel, D. *The Harvard Dictionary of Music.* Belknap Press, 2003.
Roma G., Herrera, P., and Serra, X. Freesound Radio: Supporting Music Creation by Exploration of a Sound Database. Computational Creativity Support Workshop CHI09, 2009.
Romero, J., and Machado, P. eds. The art of artificial evolution: a handbook on evolutionary art and music. Springer, 2008.
Rowe, R. *Interactive music systems: machine listening and composing.* MIT press, 1992.
Seaman, B. *Recombinant Poetics: Emergent Explorations of Digital Video in Virtual Space.* New Screen Media: Cinema/art/narrative, London: BFI Pub, 2002.
Thorogood, M. *Chatter and Listening: A Wireless Multi-node Sound Sculpture.* Proceedings of the Australasian Computer Music Conference, 2007.
Thorogood, M., and Pasquier, P. *Computationally Generated Soundscapes with Audio Metaphor.* Proceedings of the Fourth International Conference on Computational Creativity, 2013.
Truax, B. *Models and Strategies for Acoustic Design*, Acoustic Ecology Conference, 1998.
———. *Soundscape composition as global music: electroacoustic music as soundscape*. Organised Sound 13.02, 2008.
Tzanetakis, G., and Cook, P. *Marsyas: A framework for audio analysis.* Organised sound 4.3, 2000.
Wardrip-Fruin, N., and Montfort, N. eds. *The New Media Reader*. The MIT Press, 2003.
Wardrip-Fruin, N. *Playable media and textual instruments.* Dichtung Digital 1, 2005.
Whalley, I. *Software Agents in Music and Sound Art Research/Creative Work: Current State and a Possible Direction.* Organised Sound 14, 2009.
Waschka, R. *Composing with Genetic Algorithms: GenDash.* Evolutionary Computer Music. Springer London, 2007.
Whitelaw, M. *Metacreation. Art and Artificial Life.* Cambridge: MIT Press, 2004.
Wooldridge, M. An Introduction to Multiagent Systems. Chichester: Wiley, 2009.