Face to Face – Performers and Algorithms in Mutual Dependency

Jan C. Schacher[1], Daniel Bisig[2]

Institute for Computer Music and Sound Technology
Zurich University of the Arts
Zurich, Switzerland
[1] jan.schacher@zhdk.ch
[2] daniel.bisig@zhdk.ch

Abstract. This article explores modes of interaction or 'interfacing' between dancers or musicians and algorithms, and the ways in which inter-dependence and co-performance between human and machine performers arises. Two artistic works, an interactive dance piece and a live-electronic music piece, serve as example cases for the observation and discussion of how algorithms can play a central role in composition processes as well as human-machine interactions during performance. The intended goal of using algorithms in this context is the emergence of idiosyncratic behaviours in interactive systems that reflect and combine aspects of the performance situations, the generative algorithms and the adaptation mechanisms themselves. By comparing the interaction relationships and 'interfaces' of the example works, fundamental differences of the algorithmic systems become visible, and a better understanding of the impact and effect of algorithmic systems in real-life performance can be gained.

Keywords: adaptive mechanisms, algorithmic behaviour, artificial evolution, interactive dance, interactive systems, live-electronic music, machine learning.

Introduction

Human-machine interactions in live-performance represent a broad and common topic in electronic music and interactive dance practices. Less common is the use of adaptive and learning algorithms during composition, systems design, and the performance phases of such works. The use of intelligent and adaptive algorithms enhances, yet potentially also obscures the human performer's role as the source and focus of the expressive force of a piece. Because of this, both technical and conceptual challenges arise when developing strategies for establishing meaningful links and relationships between human and algorithmic agents.

For an improvisation system to respond to a performance situation in an adequate manner, the system's properties need to be adapted to the particular characteristics of the context. Such an adaptation can take at least three different forms: the system can be provided with context-specific source material to operate on, the system's sensory predisposition is specifically selected for perceiving the activity of the performance partners, or the characteristics of the system's behavioural properties are specifically chosen and adjusted to the context.

In this article, we explore the modes of interaction or 'interfacing' between humans and algorithms, and the ways in which a balance between inter-dependence and co-performance of all involved players may be reached. With the help of two exemplary pieces, we attempt to show how two types of intelligent algorithms inform compositional solutions, how principles that are derived from the technical characteristics of chosen interfaces influence behaviours (and vice-versa), and how inherited structures and action patterns can generate a conceptual interface between performers and algorithms.

At the same time, the limits of applying intelligent algorithmic systems to artistic concepts that rely heavily on human expression are exposed. The main constraint is the necessity to control a system's problem and solution space, in other words, to design and train it to consistently produce meaningful output, while remaining adaptive and responsive. An additional characteristic of working in the performing-arts context is the highly specific and idiomatic ways in which performers behave. Dancers, for example, may be repeating a choreography and musicians a given piece of music, but the variability in interpretation due to renewed expressive shaping is considerable. In open-form, improvised performance, without the presence of strictly scored movement- or sound-phrases, the field of possible variations becomes so wide as to prevent the formalisation of generalised solutions that are appropriate for algorithmic systems. Therefore, in both example works, the training and adaptation phases are closely related to the individual performer's phrasing and repertoire as well as the composition's conceptual intentions, and occur during the inception and development phases rather than the performance.

By juxtaposing the central characteristics of each piece, we show the consistent relationships and their 'conceptual interfaces', observe differences in the structure and concept of each piece, and the manner how the generative, algorithmic systems are deployed.

Background

Much of the fascination and potential of generative systems arises from interactive settings and in particular from real-time performance (Lewis 2000). In such settings, the procedural nature of a work is exposed and becomes integrated into a larger web of relationships between algorithm, performer, audience and environment. Through their impact on the characteristics of these relationships, generative methods may substantially alter the creation processes in interactive, technologically mediated dance and music performances (Jones, Brown, and d'Inverno 2012).

In terms of generative open-endedness, automated forms of real-time adaptation are important since they influence the diversity of a generative system's behaviours (Galanter 2009) and may be considered to exhibit a high level of meta-creativity (Eigenfeldt et al. 2013). However, in order for generative, algorithmic systems to become a productive part of creation processes, it is essential that their unpredictability be balanced by reproducible outcomes. While consistency is essential for understanding generative behaviour, it is only through unpredictable behaviours that generative, algorithmic systems can provoke creative discoveries (Beilharz and Ferguson 2007).

Open-form and improvisational practices both in music and dance provide the opportunity for algorithms to act as artificial performers and to engage in a playful and exploratory manner with their human counterparts. These improvisation situations constitute examples of "synergistic hybrid human-computer systems" in which a generative system assumes the role of a colleague (Lubart 2005). But these algorithms can only usefully contribute to a performance, if they are sensitive to context and environment rather than merely autonomous and independent (Bown and Martin 2012).

Example Works

Two artistic projects shall serve as real-life use-cases in order to observe and discuss how algorithms can play a central role in composition processes as well as human-machine interactions during performance. The two examples come from the complementary fields of interactive dance and open-form music with live-electronics. The concepts and premises of each work are different, yet both exhibit an overlapping, blended space (Fauconnier and Turner 2003) between performer and algorithmic system. A comparative analysis of these examples can generate a better understanding of the role and import of adaptive algorithms in human-machine interactions in general and in performing arts with algorithmic systems in particular.

Phantom Limb: Virtual Body Structures for Interactive Dance

The interactive dance performance "Phantom Limb" was realised in collaboration with composer Pablo Palacio and choreographer Muriel Romero. The performance employs a generative approach to allow a dancer to alter and extend her bodily presence and movement possibilities. This approach is based on the simulation of artificial body structures whose morphological and behavioural properties are tightly interrelated with the dancer's body and movements. The virtual body structures take on the role of extensions of the dancer's body that are capable of responding to her activities through a combination of reactive and pro-active behaviours.

The morphology of both the dancer and the body structures are represented by the simulation. The morphology consists of a mass-spring system (MSS) that is organised into a branching tree-like structure. The simulation models spring tension forces according to Hooke's law and simulates directional restitution forces that push springs towards relative rest directions. In addition, the simulation implements an Artificial Neural Network (ANN). This network can possess recurrent connections and signals propagate with time delays. The activity of the ANN affects the properties of the MSS and vice versa. This functionality is realised via the implementation of sensing and actuation elements. Sensing elements read the property of a spring and modify the activity of a neural node. Actuation elements change the property of a spring based on the activity of a neural node. For a thorough discussion of the simulation principles please refer to (Palacio and Bisig 2014).

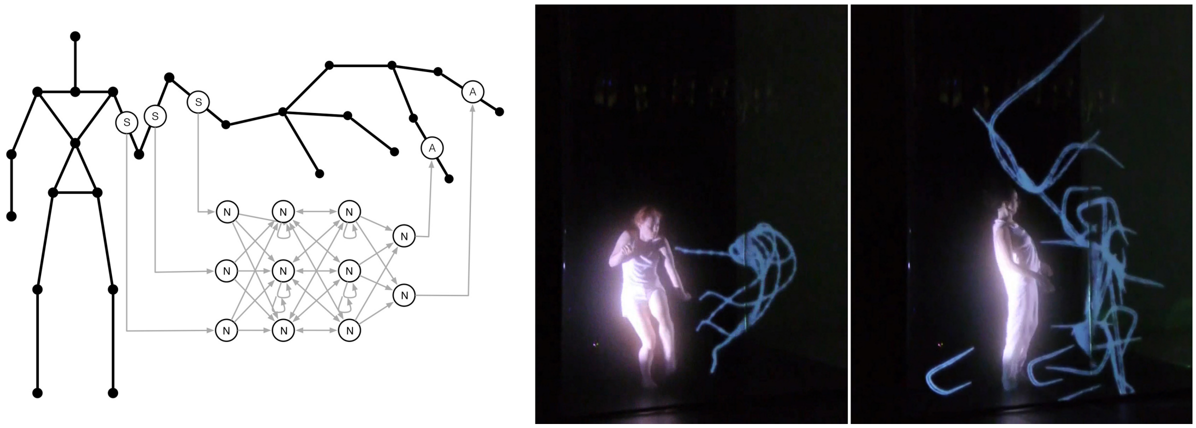

Figure 1: Phantom Limb: Simulation and Presence of a Body Structure. Left Side: Schematic depiction of a morphological and neural coupling between a skeletal representation of a dancer and a virtual body structure (black circles: mass points, black lines: springs, circular outlines: neurons (N), sensors (S), actuators (A), grey lines: neural connections). Right Side: Video projection of body structures on a transparent screen that hangs in front of a dancer.

Behavioural coupling between dancers and virtual body structures is based on their common abstraction and operational use as MSS and ANN (see left side of Figure 1). The MSS representing the dancer's body is acquired by a Kinect-based skeleton tracking mechanism. This skeletal MSS is extended with a simple ANN that serves as a proprioceptive sensing mechanism which translates the dancer's body postures into neural activity patterns. By connecting elements of the dancer's skeletal structure with those of a virtual body, the latter becomes a physically coupled mechanical system that propagates the dancer's movements. A neural coupling is established by axonal connections between the proprioceptive sensing elements located on the dancer's skeletal joints and the ANN of a body structure. This allows the dancer's movements to affect the neural activity in the ANN, which in turn initiates and modulates the active behaviours of the body structure.

For most body structures, the MSS was designed by hand whereas the ANN was generated automatically by a process of artificial evolutionary adaptation. This adaptation process took place during the development phases of the piece and served to create a repertory of behaviours that are related to specific movements of a particular dancer. Adaptation is based on a genetic algorithm that affects neurons, sensors and actuators in the model. The fitness function evaluates the quality of behavioural synchronisation between dancer and artificial body structure and can be manually overwritten in order to favour certain traits.

The visual rendering highlights the morphology of the body structure by rendering it as three dimensional tube-like structures that represent the branching topology of the underlying MSS (see right side of Figure 1). This rendering is projected on a transparent screen that hangs in front of the dancer and places the resulting image at a position and scale that aligns with the dancer's own body position and size. This settings fulfils two purposes: the correspondence between the dancer's body and movement and the virtual body structures is clearly visible for the dancer, and the audience perceives the appearance of the dancer and that of the virtual body structures in a visual superposition.

The sonification of the body structures complements the visualisation and foregrounds the sensorimotor coupling between the dancer and the artificial body structures. It does so by rendering audible the internal structural relations and dynamic processes that underlie the simulated behaviour. The sonification employs a combination of different sound synthesis and mapping techniques. Additive synthesis renders movements of the body joints audible as coordinated glissandi. Alternately, the shape of the body structures is used to control an extended dynamic stochastic concatenative synthesis model (Luque 2009). A third approach maps activity bursts within the ANN to the parameters of a synthetic percussion model, thereby generating polyrhythmic musical structures. For an in-depth discussion of the sonification principles please refer again to (Palacio and Bisig 2014).

Double Vortex: Supervised Learning in Live-Electronics

The investigation into the potential of machine learning (ML) algorithms as generative tools in electronic music is carried out with the piece "Double Vortex" for trombone and live-electronics. The interaction of the instrumentalist with the algorithmic system reflects varying degrees of inter-dependence. The connections range from purely analogous linkages to interactive and quasi-independent decision taking by the ML algorithms. The core question in the compositional process is that of agency and inter-subjective interaction, or simply the interplay between trombone player and algorithmic system. The intent is to generate a conversational interaction model that may also exhibit abstract autonomous behaviours, which are not necessarily perceptually connected to the sounds or actions performed by the musician.

Live-electronic sound processing and motion sensing represent the compositional domain that problematises the relationship between musician, instrument, movement, sound, and the algorithmic system. For the audience the declaration and subsequent recognition and reading of these techniques during performance generates expectations: they want to see and recognise the linkages and dependencies that are 'at play'. A way of playing with these expectations through the composition is by letting the system sometimes fulfil them and sometimes propose an alternate modality of interplay, the most abstract of which are autonomous, algorithmically generated musical decisions. Thus, the technology sits at the conceptual nexus between the instrumentalist's actions and the natural or electronically extended sounds.

Different ways of relating musical playing techniques and movement characteristics are explored either perceptually, or with the aid of movement and orientation sensing. The juxtaposition of movement- and sound-instructions for the musician lead to sections of the piece where simultaneous playing and moving of body and instrument produce a perceptual shift between eye and ear (see right half of Figure 2). With this strategy, complex movement patterns are overlaid to musical elements and influence the instrumental sound through a physiological impact, affecting the breath and destabilising the air-column by disturbing the player's posture.

To clarify the technical system a brief description of the systems structure of Double Vortex follows (see also left half of Figure 2). For an in-depth description of the technical tools used in this piece, please refer to (Schacher, Miyama, and Bisig 2015).

The machine learning tool is configured with three pipelines that simultaneously observe the trombone player's movements (see left half of Figure 2). The two flavours of supervised ML algorithms, Dynamic Time Warping (DTW) from the Gesture Recognition Toolkit GRT (Gillian and Paradiso 2014) and the Gesture Variation Follower (GFV) (Caramiaux et al. 2014), need to be trained, i.e. provided with templates of the movement patterns to look for. The configuration of sensors on the instrument, along with the use of patterns that are specific to the performer, demands that training be done by the instrumentalist himself, in circumstances as close to the performance as possible. This points to the fact that training should eventually become an integral part of the performance, since training by demonstration can be done in real-time with the appropriate algorithms.

For this work, and in relation to the trombonist's movement repertoire, six archetypal movement sequences are trained that can be recognised easily during performance. They are achieved with more or less ease during playing, either fully formed in the proper tempo or as elements of other gestures. It turns out that training the system with very precise movement patterns leads to a well trained ML system, but also to a narrow window of recognition. Providing a more broadly varying set of templates renders recognition more tolerant and lowers the threshold for obtaining meaningful results.

The output of these algorithms is used to control live-electronics processing for capturing and a subsequently deconstructed rendering of the sound materials performed by the trombone player himself. In order to achieve this, the boundaries of the algorithm's capabilities are explored. In addition, the differences in reaction time between the two types of algorithms are leveraged. GFV, on the one hand, is relatively fragile but provides a continuous answer about the position within a template. DTW on the other hand is quite robust, but only provides an answer after recognising a segment. In this system the two algorithms cooperate. The former takes the decisions for capturing sounds and the latter triggers a response using these deconstructed sounds-elements.

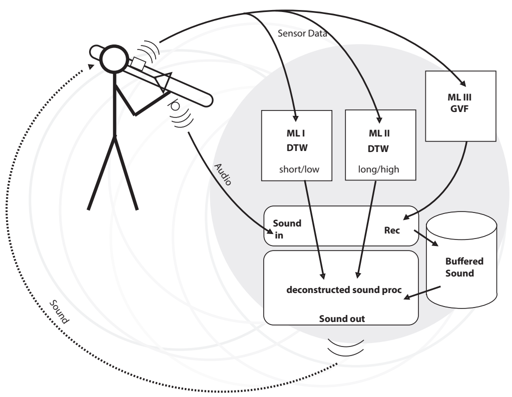

Figure 2: Double Vortex: Left: Relationships between trombone player and software agent. Of the three machine learning pipelines (ML), the first two (DTW) control sound generation, the third pipeline (GVF) controls the capture of audio. A central element is the sound feedback given to the musician. Right: Stage configuration with the moving performer and a single speaker.

Conceptual Interfacing

The two artistic examples presented here provide an opportunity to reflect on common characteristics in order to better understand the role and impact in the live-performance context of algorithmic systems endowed with some form of autonomy. In addition to the technical modes of linking the performers to the machine and vice-versa, conceptual couplings become visible that represent the central 'interfaces' for both works. The following overview juxtaposes the two tiers of interfacing: on the one hand the **technical** linkage and on the other hand the **conceptual** role of the algorithmic system in relation to the human performer.

Interfaces		**Phantom Limb**	**Double Vortex**
technical	sensing input	camera/skeleton tracking	motion (IMU), microphone
	media output	image, sound	sound
conceptual	semantic	behavioural coupling w/algorithm	behavioural coupling w/algorithm
	inheritance	morphology	sound material
	adaptation	evolutionary adaptation	machine learning

Table 1: Comparing interfacing relationships between the two works.

In both cases the **technical interfacing** between system and dancer/musician is informed by the necessities of stage-worthiness and portability.

Phantom Limb employs camera-based tracking to update the dancer's representation as an MSS and to generate new sensory inputs for the ANN. The stationary camera provides an absolute spatial reference that allows to position the dancer's body in relation to the virtual body structures.

Double Vortex connects the performer with the system with the aid of a motion-sensing inertial system, that provides data about the absolute orientation of the instrument in space. This information enables the ML system to observe the performer, and controls by traditional mapping aspects of processing of the instrumentalist's sound.

A microphone on the bell of the instrument is the source of all sound treated by the live-electronics, and also generates an acoustic feedback path between the instrument and loudspeaker.

In *Phantom Limb*, the creation of synthetic audio and video is based on a sonification and visualisation approach and fulfils the functional requirements of highlighting the kinaesthetic correspondences between dancer and body structures.

In *Double Vortex* the algorithmic system as well as the traditional live-electronics processing produce their results as sound. This is played back on a single speaker situated next to the performer, thus endowing the system with the character of an independent 'object' if not 'subject' that is present on stage.

The **conceptual interface** exists in the manner that performers and algorithmic systems relate to each-other in the domains of perception, measurement, and particularly behaviour.

The *semantic* relationship describes how significance or meaning is generated by the presence of both human performer and algorithmic system.

The design of the simulation in *Phantom Limb* is informed by concepts from the field of embodied artificial intelligence (Pfeifer and Bongard 2006). Phantom Limb connects to the notion that fundamental aspects of agency as well as cognitive capabilities are grounded by mutual inter-dependencies between a system's physical, morphological, perceptual and neural characteristics (Sørensenn and Ziemke 2007). This approach enables the integration of an interactive system with embodied improvisation techniques in dance and emphasises the relationship between bodily predispositions and kinaesthetic imagination and creativity.

Double Vortex is designed to endow the algorithmic system with a semblance of autonomy. This perceived autonomy challenges the notion of control by the musician and creates a slight ambiguity with regard to authorship and the origins of the perceived sounds. At the same time it demands of the musician to act and interact with the system during the performance, thereby enforcing a shared focus between inner and outer perception (Marcel 2003).

The *inheritance* relationship describes how behaviours and structures are inherited from one side to the other.

In *Phantom Limb*, the visual and acoustic rendering of a body structure possesses little similarity with the dancer's appearance. Rather, inheritance originates from a correlation of simulation-based abstractions. A direct form of inheritance is established through a mechanical coupling between dancer and artificial body structure. A less direct form of inheritance is created by giving the ANN a sensitivity to particular types of dance movements. This increases a dancer's influence on the behaviours of an artificial body structure.

In *Double Vortex* both sides operate in the movement as well as sound domain. By linking the technical observation to movement and not sound, a gestural rather than sonic link is established between musician and algorithm. Since the resulting responses of the system occur in the sonic domain, by reusing, deconstructing, and therefore abstracting the instrumentalists sound contribution, a first type of inheritance is established. As a consequence, the reaction to the decoupled sound-events that originate from the system puts the human player into an inter-dependence, thereby generating a 'dialogical' and truly interactive situation.

The *adaptation relationship* describes how the presence and actions of one side influence the behaviour and output of the other.

The evolutionary adaptation in *Phantom Limb* serves to provide the ANN with sensitivity to particular dance movements. This results in a continuous mapping between input data, simulation behaviour and audiovisual rendering rather than static relationships between input categories and musical events. The criteria specified loosely by the evolutionary fitness function enable a large solution space. This enables the automated adaptation process to come up with a wide variety of potentially interesting solutions, later to be narrowed down according to additional quality criteria.

The technique used in *Double Vortex* to create adaptive responses from both sides is to deliberately blur the precision of pattern-recognition and to modulate the sensitivity of the algorithmic responses in such a way as to generate a self-reinforcing adaptive loop between performer and system. In addition, the juxtaposition of decisions by the musician with semi-autonomous musical decisions by the algorithmic system creates a balance. Both agents need to perceive the other in order to meaningfully react. The performer does this by listening to the algorithmic sound responses, the machine does it by watching the player's movements.

Discussion and Conclusion

By juxtaposing the two case studies several arguments can be developed which may inform a common heuristic for integrating human performers and algorithmic systems into shared improvisation contexts.

A core concern when using generative algorithms in an interactive setting is the issue of comprehensibility of a perceivable, albeit complex, correlation between human and machine activities that becomes visible both from the performer and the audience perspectives. Three elements affect this: human and machine possess comparable perceptual capabilities, human and machine exhibit a shared musical and/or visual presence, and human–machine interaction is at least partially reproducible.

The machine possesses perceptual capabilities that enable it to read and respond to aspects of the performing partner's presence and activities, which in turn are readily observable by the human player. In both *Phantom Limb* and *Double Vortex*, the machine's perceptual capabilities are connected to the performer's own kinaesthetic body awareness.

The visual and/or sonic presence of both interactive systems is sufficiently related to the presence of the human performers for comparing them within a perceptual as well as aesthetic domain. In *Phantom Limb*, this shared presence is apparent in the visual overlap between the dancer and the simulated bodies but also exists in the metaphorical domain as a sonic translation of kinaesthetic activities. In *Double Vortex*, the shared presence arises from the combination of the physicality of both the human musician and the loudspeaker. The speaker-object acts as a physical representation or 'place-holder' for the interactive system whose musical co-presence is based on live-input and transformation of the instrumental sounds performed during the piece.

In both cases, the interaction between human performer and algorithmic system only becomes traceable and comprehensible once clear relationships between their respective behaviours are established. While the characteristics of these relationships are discovered, developed and reinforced through the use of potentially open-ended adaptive mechanisms, these mechanisms themselves are no longer operational during performance. They occur during the composition, design and rehearsal phases. In Phantom Limb the evolutionary adaptation relates to idiomatic movement phrases of a single performer, whereas in Double Vortex the training of the machine learning serves to capture gesture templates that are recurrent or demanded of the musician during performance. This approach preserves the capability of generative systems to create surprising interactive relationships, but at the same time stabilises the relationships in order to make them reproducible.

Employing machine learning and evolutionary adaptation facilitates the integration of generative, algorithmic approaches with compositional and choreographic techniques. While it is still possible and fruitful to integrate musical or movement concepts directly and explicitly on an algorithmic level, the generative system can be primed with particular performance ideas through an iterative process of exposition to representative elements of the performance. By doing so, the generative system is treated as an artificial performance partner throughout the creation and rehearsal phases of a new work. This establishes a more fluid and immediate mode of exchange between algorithmic design decisions, choreographic and compositional goals, and experimentation and experience in performance.

A core concept that is shared by both case studies relates to the automated adaptation mechanisms. They are used to transfer some of the stylistic or subjective qualities of a particular performer and performance situation to the interactive system. This transfer reduces the generic aspects of human machine interaction in favour of a relationship that conveys elements of an inter-personal exchange. Thus the interactive system becomes sensitive towards higher-level expressive qualities that are unique to a performance situation. This enables the emergence of idiosyncratic behaviours in the interactive system that reflect and combine aspects of the performance situation, the generative algorithms and the adaptation mechanism itself.

In closing, we wish to emphasise the complexity of the relationships generated when creating work in this manner. These relationships are by no means limited to the performer and the algorithmic system. Rather, the performer's presence is preceded by the composer's and choreographer's intentions and efforts, and is completed by the audience's partaking of the performance. A particular aspect of the distributed authorship between human and algorithmic performers is the inherently unstable relationship and the necessity to continuously negotiate the position of each player in this game. After all, the effectiveness of these strategies can only provide an aesthetic gain to a person who manages to perceives all the intertwined and relational elements of a technological performance as one.

Acknowledgements. This article is based on work carried out in the contexts of the 'Motion Gesture Music' Project at the Institute for Computer Music and Sound Technology of the Zurich University of the Arts that is funded by the Swiss National Science Foundation Grant No. 100016_149345, and the 'METABODY' project that is coordinated by the Reverso Association and supported by the Culture Programme of the European Union.

References

Beilharz, Kirsty, and Sam Ferguson. 2007, June. "Gestural hyper instrument collaboration with generative com- putation for real time creativity." *Proceedings of the 6th Conference on Creativity & Cognition*. Washington, DC, USA, 213–222.

Bown, Oliver, and Aengus Martin. 2012. "Autonomy in Music-Generating Systems." *Proceedings of the International Workshop on Musical Metacreation at AIIDE*. 8–13.

Caramiaux, Baptiste, Nicola Montecchio, Atau Tanaka, and Frédéric Bevilacqua. 2014. "Adaptive Gesture Recognition with Variation Estimation for Interactive Systems." ACM Transact. on *Interactive Intelligent Systems* 4, no. 18.

Eigenfeldt, Arne, Oliver Bown, Philippe Pasquier, and Aengus Martin. 2013. "Towards a Taxonomy of Musical Metacreation: Reflections on the First Musical Metacreation Weekend." *Proceedings of the International Workshop on Musical Metacreation*.

Fauconnier, Gilles, and Mark Turner. 2003. *The way we Think: Conceptual Blending and the Mind's Hidden Complexities.* Reprint 2008. New York, USA: Basic Books.

Galanter, Philip. 2009. "Truth to Process – Evolutionary Art and the Aesthetics of Dynamism." *Proc. of the Generative Art Conference*. Milan.

Gillian, Nicholas, and Joseph A Paradiso. 2014. "The Gesture Recognition Toolkit." *The Journal of Machine Learning Research* 15 (1): 3483–3487.

Jones, Daniel, Andrew R. Brown, and Mark d'Inverno. 2012. "The Extended Composer, Creative reflection and extension with generative tools." In *Computers and Creativity*, edited by Jon McCormack and Mark d'Inverno, 175–203. Springer Berlin Heidelberg.

Lewis, George E. 2000. "Too many notes: Computers, complexity and culture in voyager." *Leonardo Music Journal* 10:33–39.

Lubart, Todd. 2005. "How can computers be partners in the creative process: classification and commentary on the special issue." *International Journal of Human-Computer Studies - Special issue: Computer support for creativity* 63 (4-5): 365 – 369.

Luque, Sergio. 2009. "The Stochastic Music of Iannis Xenakis." *Leonardo Music Journal*, vol. 19.

Marcel, Anthony. 2003. "The Sense of Agency: Awareness and Ownership of Action." In *Agency and Self-Awareness*, edited by Johannes Roessler and Naomi Eilan, 48–93. Oxford University Press Oxford, England.

Palacio, Pablo, and Daniel Bisig. 2014. "Phantom Limb. Connecting Cognitive Neurosciences, Sound Synthesis, Generative Video and Dance." *Proceedings of the Congreso Internacional Espacios Sonoros y Audiovisuales*. Madrid.

Pfeifer, Rolf, and Josh C. Bongard. 2006. *How the Body Shapes the Way We Think: A New View of Intelligence* (Bradford Books). The MIT Press.

Schacher, Jan C., Chikashi Miyama, and Daniel Bisig. 2015, May 31 – June 3. "Gestural Electronic Music using Machine Learning as Generative Device." *Proceedings of the International Conference on New Interfaces for Musical Expression, NIME'15,*. Baton Rouge, USA.

Sørensenn, Mikkel Holm, and Tom Ziemke. 2007. "Agents without Agency?" *Cognitive Semiotics* 0 (1): 102–124.