On Human-AI Collaboration in Artistic Performance

Alessandro Saffiotti[1] and **Peter Fogel**[2] and **Peter Knudsen**[3] and **Luis de Miranda**[4] and **Oscar Thörn**[5]

Abstract. Live artistic performance, like music, dance or acting, provides an excellent domain to observe and analyze the mechanisms of human-human collaboration. In this note, we use this domain to study human-AI collaboration. We propose a model for collaborative artistic performance, in which an AI system mediates the interaction between a human and an artificial performer. We then instantiate this model in three case studies involving different combinations of human musicians, human dancers, robot dancers, and a virtual drummer. All case studies have been demonstrated in public live performances involving improvised artistic creation, with audiences of up to 250 people. We speculate that our model can be used to enable human-AI collaboration beyond the domain of artistic performance.

1 Introduction

Should AI systems augment humans, replace humans, or collaborate with humans? This question is being regularly asked by both citizens and policy makers, often boosted by the media, and researchers in all fields are increasingly faced with the many facets of this question.

For the purpose of our discussion, we define the above three models as follows. Consider a task T, traditionally performed by a human. In the *augmentation* model, T is still performed by the human, and this is empowered with new tools and functionalities built through AI. In the *replacement* model, T is instead performed by an artificial agent built using AI technology. In both these models, the task is performed by a single agent. In the *collaboration* model, by contrast, T is performed jointly by two independent but collaborating agents: a human agent, and an AI-based artificial agent.

The topic of *collaborative AI*, or how to make AI systems that collaborate with humans in performing a joint task, is the subject of increasing interest in the AI community. The emphasis on the collaboration model as opposed to the replacement model is also in line with the *Ethics Guidelines* produced by the European High Level Expert Group [14], and later adopted by the European Commission [10], that insists that humans should maintain agency and oversight with respect to AI systems. Aspects of collaborative AI have been studied in several areas, including human-robot teams [18], shared agency [5], hybrid human-AI intelligence [42], mixed-initiative systems [11] and symbiotic systems [8]. In this note, we contribute to the study of human-AI collaboration in a particularly telling domain: collaborative artistic performance.

[1] School of Science and Technology, Örebro University, Örebro, Sweden. Contact email: asaffio@aass.oru.se
[2] University of Skövde, Sweden.
[3] School of Music, Theatre and Art, Örebro University.
[4] School of Humanities, Education and Social Sciences, Örebro University.
[5] School of Science and Technology, Örebro University.

Artistic performance, like music, dance or acting, provides an excellent setting to observe and analyze the mechanisms of human-human collaboration, especially in live or improvised performance. It seems natural, thus, to study human-AI collaboration in this setting. Artistic creativity in general, and in music in particular, has been a topic of interest for computer researchers since the early days of computer science [15] and AI [27]. Ada Lovelace already noted in 1843 that computers "could potentially process not only numbers but any symbolic notations, including musical and artistic ones" [19]. Today, there is a rich literature of computational approaches to music [7, 29], including many AI systems for music composition and improvisation [26]. As pointed out by Thom [38], however, most of these systems focus on the offline creation of music, and not on the online collaborative performance between the human and the AI musicians: the latter is what is usually referred to as *co-creativity* [21, 22]. Notable exceptions in computational music are the early work on jazz improvisation by Walker [41], and the work on a marimba playing robot by Hoffman and Weinberg [17]. Co-creativity has also been studied in other artistic areas, like theater [32], as well as in the more general field of human-computer interaction [25].

In this paper, we study AI systems capable of on-line, collaborative interaction with humans in the context of artistic performance, with a specific focus on live music improvisation. The main contribution of this paper is a general model for collaborative artistic performance, in which an AI system mediates the interaction between a human and an artificial performer. We show how this model has been instantiated in three concrete case studies involving different combinations of human musicians, human dancers, robot dancers, and a virtual drummer. We also briefly discuss the complex problem of how to evaluate human-AI collaborative interaction, especially in an artistic context. We hope that our model can contribute to a better understanding of the general mechanisms that enable successful collaboration between AI systems and humans.

2 A model for Human-AI collaborative artistic performance

The model that we propose for Human-AI collaboration in artistic performance is illustrated in Figure 1. In this model, an AI system is used as a *mediator* to coordinate the performance of two autonomous agents: a human performer, and an artificial performer. This model therefore comprises three elements: two performers and a mediator.

For illustration purposes, Figure 1 shows a guitar player as human performer and a dancing robot as artificial performer. We emphasize, however, that we use the term "artificial performer" in a broad sense, to mean any agent that generates physical outcomes: this could be a physical robot producing movements, a virtual instrument producing sounds, or a projector producing artistic visualizations. In the case studies reported below, we use an off-the-shelf commercial virtual drummer and a off-the-shelf humanoid robot as artificial performers.

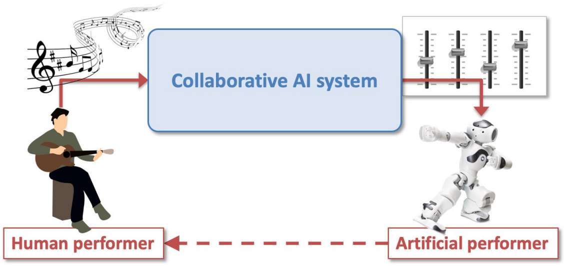

Figure 1. Black box view of our model for Human-AI collaboration. The parameters of the artificial performer are adjusted to adapt to the human performer, and the human performer naturally adapts to the artificial one.

Our model has the following distinctive features.

- **Supervisory**. The AI system does not directly generate the artistic output. Instead, we assume that the artificial performer is capable of autonomous artistic performance, whose modalities are controlled by a fixed number of parameters. The parameters of interest here are expressive parameters, that modulate the behavior of the artificial performer to produce different artistic expressions. For example, a robot may be able to autonomously perform a number of dancing patterns, and parameters may make its motions more aggressive or more subtle.
- **Reactive**. The goal of the AI system is to analyze the artistic expression in the live human performance, and to dynamically adapt the parameters of the artificial performer to match this expression. Thus, the behavior of the artificial performer is influenced by the human performer, but it is not fully determined by this.
- **Proactive**. The AI system may be creative and proactive in setting the performance parameters. The human performer hears or sees what the artificial performer does, and may adjust to it (the right-to-left arrow in Figure 1). Our hypothesis, which we aim to verify trough empirical studies, is that this feedback loop will result in an harmonious joint performance between the human and the autonomous agent.

By combining reactive and proactive behavior, our model implements the two directions in human-AI co-creativity described by Liapis and colleagues [24]: the human guides artificial creativity (reactivity), and the AI system triggers human creativity (proactivity). This also resonates with the idea of "directed improvisation" introduced in the area of multi-agent human-computer interaction [13].

From the three features above, it should be clear that the key move to achieve collaborative artistic performance in our model is to align the artistic expressions of the two performers, each one realized through its specific expressive means. The role of the collaborative AI system is to realize this alignment. This alignment can be seen as an artistic counterpart of inter-modal mapping, that is, the ability of people to associate stimuli received in one modality, e.g., shapes, to stimuli in another modality, e.g., sound [33]. In terms of musical semiology [31], we can also see the function of the collaborative AI system as a mapping from *aesthetics* (perception) to *poietics* (production) that preserves artistic meaning.

3 System architecture

We have implemented the above model in a concrete system for collaborative artistic performance, which has been tested in a number of case studies, three of which are reported below. Figure 2 shows the high-level architecture of this system, in the case where the human performer is a jazz pianist and the artificial performer is a parametric virtual drummer (first case study below).

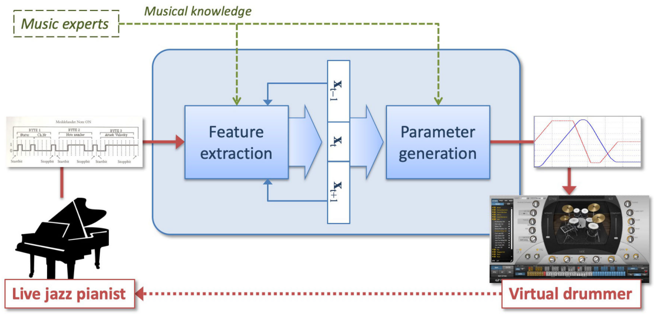

Figure 2. White box view of our model, as implemented in the case studies reported in this paper. The figure refers specifically to the first case study.

The **features extraction** module analyzes input from the human performer, in this case in the form of MIDI signals, and estimates the value of a set of variables that represent the *musical expression* of the performance, like the keypress velocity and the rhythmic density. This module computes an expressive state represented by variables that depend on the past and current input values, as well as variables that predict future states. Examples of the former are the instantaneous velocity $v(t)$, the average velocity $\bar{v}(t)$ over the last bar, and the velocity slope $\Delta_v(t)$; an example of the latter is a predicted climax $\hat{c}(t+1)$ at the end of an ongoing crescendo. Variables referring to past, current and predicted states are represented in the figure by $\mathbf{x}_{t-1}$, $\mathbf{x}_t$ and $\mathbf{x}_{t+1}$, respectively.

The **parameter generation** module uses the above variables to decide the values of the execution parameters of the artificial agent, so as to continuously adapt its performance to match the current *musical expression* of the human performer. In the case of a virtual drummer shown in the picture, these parameters include the intensity $I(t)$ and complexity $C(t)$ of the drumming, the drumming pattern $P(t)$, and the selective muting $M(t)$ of some of the drums.

The above architecture can be interpreted in terms of the semantic perception-production mapping mentioned above: in this view, feature extraction would correspond to aesthetics, parameter generation to poietics, and expressive state variables represent artistic meaning. The architecture also reminds of the listener-player schema for interactive music systems originally proposed by Rowe [35], and later used in several works [28]. However, the crucial difference is that what are generated in our case are not performance contents (music or movements) but performance parameters.

To implement both feature extraction and parameter generation, we relied on a knowledge-based approach where knowledge from the music experts was manually encoded into the system (the top arrows in Figure 2). Our team includes both computer scientists and musicians: discussions among these revealed that musicians possess heuristic knowledge of how the drummer's parameters depend on the pianist's play, and that this knowledge can be expressed in terms of approximate rules using vague linguistic terms, like:

If rhythmic complexity on the lower register is *high*,
Then rhythmic complexity of drums should *increase strongly*.

This type of knowledge is suitably encoded in fuzzy logic [23], and consequently we implemented both feature extraction and parame-

Case	Input device	Extracted expressive features	Output device	Generated expressive parameters
4.1	MIDI piano	velocity, rhythmic density, velocity slope, density slope, step change	Strike drummer	intensity, complexity, pattern, fill, selective drum mute
4.2	Motion tracker	count, distance	Strike drummer	pattern, selective drum mute
4.3	MIDI piano	velocity, rhythmic density, velocity slope, density slope, step change	Pepper robot	motion type, selective joint mute

Table 1. Configurations used in the three case studies

ter generation using multiple-input multiple-output Fuzzy Inference Systems (FIS). Each FIS is based on the usual fuzzify-inference-defuzzify pipeline found in classical fuzzy controllers [9]. To take the temporal aspect into account in the feature extraction FIS, we use a recurrent fuzzy system [1] that takes the current estimated state and predictions as input. This solution allows us to capture the knowledge of the musician about temporal patterns, e.g., about what counts as a "sudden drop in intensity", in a way that is both explicit and easy to modify.

The same implementation has been used in all the case studies reported below, with only minor changes to the fuzzy rules and the membership functions. Further details of this implementation can be found in [39]. For the purpose of this paper, we shall not discuss the technical realization of our model in any depth; rather, we want to demonstrate its applicability in breadth across different types of artistic collaboration, and different types of human and artificial players.

4 Case studies

We now give concrete examples of how the above model was used in three different case studies involving human-robot collaborative artistic performance. Each case was implemented using fuzzy rule-based systems as discussed in the previous section. In each case, the features extracted from the input characterize the detected musical expression of the human performer; and the parameters sent as output represent the desired artistic expression of the artificial performer. The input device, the output device, the extracted features and the generated parameters are different for each case study, and they are summarized in Table 1.

4.1 A human pianist and a virtual drummer

The first case study involves collaboration in live jazz performance. The human performer was a pianist performing improvised jazz music, while the robot performer was the commercial virtual drummer Strike 2.0.7 [2]. The tempo and style were agreed before the performance starts, as is commonly done among musicians, and manually set into the virtual drummer. The other parameters of the virtual drummer were decided in real time by the AI system based on the musical expressive features of the piano performance using the architecture in Figure 2, as described in the previous section.

The architecture was implemented in Python 3.6.8 with the MIDO library (1.2.9). The input comes from a MIDI piano, or from a MIDI file for debugging purposes. The output was a MIDI signal, encoding the parameters to be sent to the Strike drummer.

The resulting system was tested in two public concerts given at the Music School of Örebro University, Sweden, in Spring 2019, attended by about 60 and about 100 people, respectively. Video recordings from these concerts are available online at [3]. Figure 3 is a snapshot from video of the second concert: the background screen shows a visualization of the Strike drummer; the monitor on the right shows the output membership functions generated by our system, from which the system extracts the control parameters sent to Strike. Although we did not collect structural feedback (e.g., questionnaires), informal comments by the audience were very positive, with many people remarking that the drummer appeared to follow (and sometime anticipate) the pianist in a natural way.

In addition to feedback from the audience, we also collected informal feedback from the artists. The pianist at the concerts commented that the AI-controlled drummer was perceived as collaborative and "human like". He also remarked that it was often "surprising" in a way that he did not expect a machine to be, and sometimes more "proactive" than a human drummer might be, leading him to follow what the drummer seemed to suggest, as per the feedback arrow in Figure 1.

It is worth speculating a bit on the last point. We believe that this feeling of proactivity is partly due to the use of expectations ($\mathbf{x}_{t+1}$) in parameter generation, leading to an *anticipatory* behavior in the drummer [34]. For example, when a step change is predicted, e.g., expecting to go from a *forte* to a *piano*, the system first mutes the kick, and if the change is confirmed it then also mutes the snare. The pianist may perceive the absence of the kick as a suggestion for a change in mood, and either follow the suggestion and go to a *piano*, or not follow it and persist with the *forte*. In the first case, the drummer will also mute the snare; in the second case, it will unmute the kick. An example of this proactive interaction can be observed in the video recording at [3] around time 22:34.

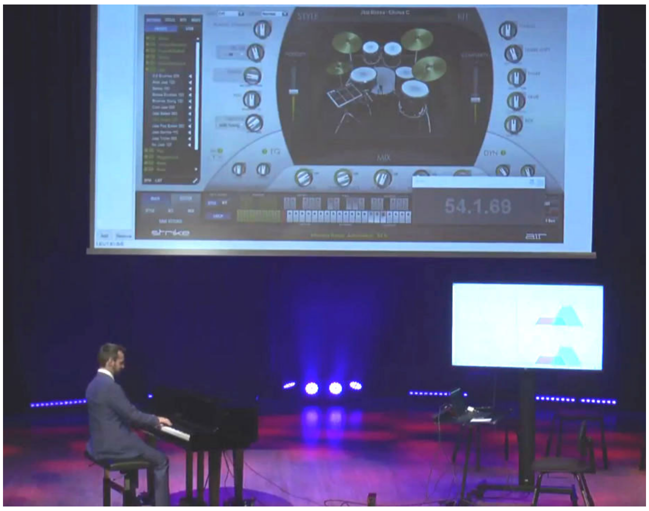

Figure 3. A snapshot from a public concert given on June 12, 2019, during the international symposium "Humans Meet Artificial Intelligence".

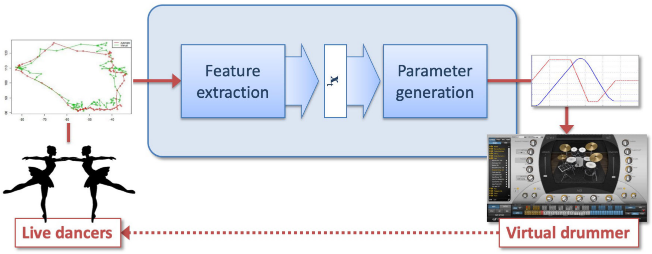

Figure 4. The variation of our system used for the second case study. The variables x_t depend on the current position and distance of the dancers.

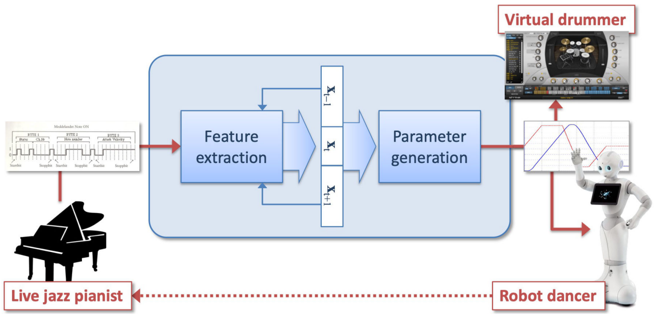

Figure 6. The system used for the third case study. The robot has been enriched with control software to perform classical ballet movements.

4.2 Two human dancers and a virtual drummer

Our second case study happened by serendipity. In October 2019, the Music Tech Fest (MTF) art festival [30] was hosted at Örebro University. There, our team met the *Accents in Motion* team, who had previously researched the use of body movements to control sound. We decided to join forces and explore if and how the performance of the virtual drummer could follow two human dancer improvisers, using our model.

To do this, we used the simplified version of our architecture shown in Figure 4. The input to the system was taken from a Vicon tracking system mounted in a large laboratory space. Together with the artists, we decided the features to extract and how the drummer should react to those. We drew an area on the floor to act as a "black box" where dancers would be invisible to the tracking system. We decided to extract two single features from the tracking system data: the number of dancers that are visible (i.e., outside the "black box"), and their mutual distance. For the parameter control part, we used two simple principles. The distance among dancers would influence the pattern of the drummer: the closer the dancers, the more complex the pattern. The number of visible dancers would influence which instruments are played: none with no dancer, only cymbals with one dancer, all cymbals and drums with two dancers.

The above system was realized in collaboration with the *Accents in Motion* team within an MTF lab, during two hectic days of work. A performance was recorded on October 19, 2019, and shown at the MTF closing night. A clip from that recording is available at [3]. Figure 5 shows a snapshot from the clip. The plot at the bottom shows the temporal evolution of the 'number' and 'distance' variables: at the time of the snapshot, both dancers have just jumped out of the "black box" and became visible to the system.

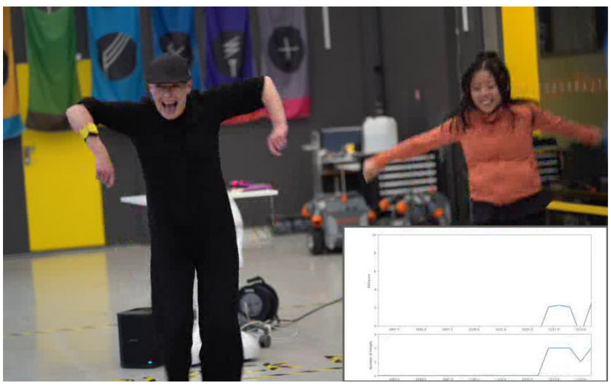

Figure 5. A snapshot from the "Music Tech Fest" lab on October 19, 2019

4.3 A human pianist and a robot dancer

In our last case study, we used our model to realize a collaborative live performance of a jazz pianist and a robot dancer. Like in the first case study the jazz pianist improvises, but this time the AI system controls the execution parameters of a humanoid robot.

For this experiment, we have used the commercial robot Pepper, produced by Softbank Robotics, as artificial performer. We used the system shown in Figure 6. The collaborative AI system is the same one used in the first case study, but now the performance parameters are sent both to the virtual drummer and to the robot.

The robot has been enriched with a control software to continuously perform dancing motions, synchronized with the pre-defined beat, and generated from a library of basic motions inspired by classical ballet [12]. Movements may involve one or both arms, the head, or the base, in any combination. The dancing motions are selected, modified and chained depending on the parameters received from the collaborative AI system. Some parameter values are mapped to different combinations of motions, which are decided randomly in order to produce a more lively performance. Selective muting disables some of the degrees of freedom, like the head or the base, and is typically used in response to more quiet passages by the piano. Further details on the implementation of this test case are given in [40].

The above case study was demonstrated in a public performance at the official yearly celebration of Örebro University, attended by about 250 people. A video recording is available at [3], Figure 7 shows

Figure 7. A snapshot from the public concert on January 31, 2020

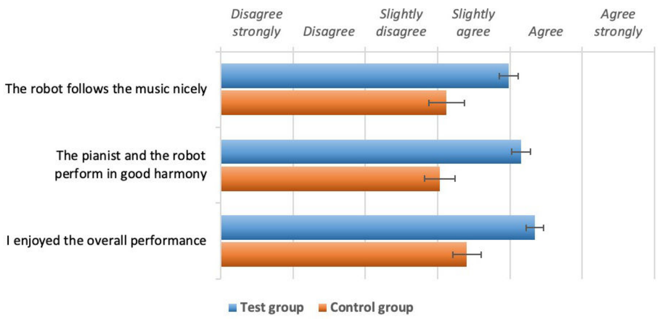

Figure 8. Results from the online user study. Error bars represent ±1 standard error of the mean.

a snapshot from that video. The reaction from the audience to the performance was overwhelmingly positive.

5 Evaluating collaborative performance

An open question for a human-AI collaborative system is how to evaluate the effectiveness and added value of the collaboration. This question is even more complex in the case of artistic collaboration, where we are faced with the double problem of evaluating the collaborative aspect and the artistic aspect. Recently, some works have been reported on the evaluation of collaborative human-robot systems [16] and of artificial artistic creations [6, 20], but much still needs to be understood in both domains and in their combination.

Bown [4] has suggested that the evaluation of artificial creative systems should look at the subjective experiences of humans. In the case of our model, the informal feedback received at the live events indicated that both the audience and the musicians experienced a feeling of harmonious collaboration between the performers. We then decided to evaluate this feeling in quantitative terms, and we run an online user study aimed at measuring the subject's perception of collaborative artistic performance in the third case study above.

The experimental setup was designed to highlight the collaboration aspect rather than the quality of the robot's performance. We created two versions of the system based on Figure 6, a *test* one and a *control* one. Both versions used the same artificial performers: the Strike 2 virtual drummer, and the Pepper robot performing dancing movements synchronized with the music beats. However, while the test case used our collaborative AI system to decide the parameters of the robot's performance, in the control case those parameters were selected randomly. (The parameters for the virtual drummer were generated by our system in both cases.)

We recruited 90 subjects using the Amazon Mechanical Turk. Subjects were randomly assigned to a test group (58), that were shown videos of performances using the test version of the system; and to a control group (32), that were shown videos of performances using the control version. These videos can be seen at `https://tinyurl.com/yyg67eco`. Subjects were asked to rate a few statements about the performance, using a 6-step Likert scale. The survey was created and run using PsyToolkit [37].

The results of the experiment are visualized in Figure 8. Subjects in test group consistently rated the statement "The robot follows the music nicely" higher than those in the control group, showing that our system successfully aligns the artistic performance of the robot to the one of the pianist. The test group also gave higher rates to the statement "The pianist and the robot perform in good harmony", supporting our hypothesis that the loop in Figure 1 leads to a perceived sense of collaborative performance. Finally, the test group consistently rated the statement "I enjoyed the overall performance" higher than the control group, suggesting that our system may result in increased perceived artistic value. Full details of this user study are reported in [40].

6 Conclusions

We have proposed a model for collaboration between human agents and artificial agents in artistic performance. Our model does not focus on the production of behavior in the artificial agent. Instead, we assume that the artificial agent is already capable of autonomous performance, and we focus on how an AI system can be used to modulate this performance through the manipulation of its parameters, to harmoniously adapt to the performance of the human.

Co-existence, co-operation and co-creation between humans and AI systems are today extremely important areas of investigation. Collaborative artistic performance among humans is one of the domains where these phenomena are most visible, and it is therefore an ideal domain where to study the foundations of human-AI collaboration. We hope that the model, the case studies and the evaluation presented in this note contribute to this study.

The implementation of the model used in our test cases is purely knowledge-based. In this initial stage, this approach was chosen because it afforded us a quick bootstrap using existing music knowledge. The knowledge-based approach also allowed us to go through an open, modular and incremental development loop. Interestingly, the music experts found that the process of eliciting knowledge was rewarding for them. For example, they found that the need to describe music performance in logical terms led them to develop a new analytical perspective on how, when and why different styles are being chosen and used. Notwithstanding the advantages of the knowledge-based approach, we plan in the near future to integrate this approach with a data-driven approach for the feature extraction part, the parameter generation part, or both. This might help to complete the handwritten rules, or to adapt them to the artistic taste of a given musician. It might also allow us to use sub-symbolic input, like sound, rather than symbolic one, like MIDI [36].

ACKNOWLEDGEMENTS

We are grateful to the *Accents in Motion* team (Andreas Bergsland, Lilian Jap, Kirsis Mustalahti, Joseph Wilk) and to MD Furkanul Islam for their collaboration in the second case study. Madelene Joelsson kindly acted as pianist in the third case study, and Vipul Vijayan Nair helped us with the survey. Work by A. Saffiotti and L. de Miranda was supported by the European Union's Horizon 2020 research and innovation programme under grant agreement 825619 (AI4EU). The cross-disciplinary perspective was made possible thanks to the CREA initiative (`crea.oru.se`) of Örebro University.

REFERENCES

[1] Jürgen Adamy and Roland Kempf, 'Regularity and chaos in recurrent fuzzy systems', *Fuzzy Sets and Systems*, **140**(2), 259–284, (2003).

[2] AIR music technology. `https://www.airmusictech.com/product/strike-2/`.

[3] Alessandro Saffiotti (maintainer). Human meets AI in music. Website: `http://crea.oru.se/Music`, 2020.

[4] Oliver Bown, 'Empirically grounding the evaluation of creative systems: Incorporating interaction design', in *Int Conf on Computational Creativity (ICCC)*, pp. 112–119, (2014).

[5] Michael E Bratman, *Shared agency: A planning theory of acting together*, Oxford University Press, 2013.

[6] Rebecca Chamberlain, Caitlin Mullin, Bram Scheerlinck, and Johan Wagemans, 'Putting the art in artificial: Aesthetic responses to computer-generated art.', *Psychology of Aesthetics, Creativity, and the Arts*, **12**(2), 177, (2018).

[7] David Cope, *Computer models of musical creativity*, MIT Press Cambridge, 2005.

[8] Silvia Coradeschi and Alessandro Saffiotti, 'Symbiotic robotic systems: Humans, robots, and smart environments', *IEEE Intelligent Systems*, **21**(3), 82–84, (2006).

[9] Dimiter Driankov, 'A reminder on fuzzy logic', in *Fuzzy Logic Techniques for Autonomous Vehicle Navigation*, eds., D Driankov and A Saffiotti, chapter 2, Springer, (2001).

[10] European Commission. White paper on artificial intelligence: A european approach to excellence and trust. `https://ec.europa.eu/commission/presscorner/detail/en/ip_20_273`, 2020.

[11] George Ferguson and James Allen, 'Mixed-initiative systems for collaborative problem solving', *AI magazine*, **28**(2), 23–23, (2007).

[12] Ann Hutchinson Guest, *Choreo-graphics: a comparison of dance notation systems from the fifteenth century to the present*, Psychology Press, 1998.

[13] Barbara Hayes-Roth, Lee Brownston, and Robert van Gent, 'Multiagent collaboration in directed improvisation.', in *Proc of the First Int Conf on Multiagent Systems*, pp. 148–154, (1995).

[14] High Level Expert Group on AI. Ethics guidelines for trustworthy AI. `https://ec.europa.eu/digital-single-market/en/high-level-expert-group-artificial-intelligence`, 2019.

[15] Lejaren A Hiller Jr and Leonard M Isaacson, 'Musical composition with a high-speed digital computer', *J. of the Audio Engineering Society*, **6**(3), 154–160, (1958).

[16] Guy Hoffman, 'Evaluating fluency in human-robot collaboration', *IEEE Transactions on Human-Machine Systems*, **49**(3), 209–218, (2019).

[17] Guy Hoffman and Gil Weinberg, 'Interactive improvisation with a robotic marimba player', *Autonomous Robots*, **31**(2-3), 133–153, (2011).

[18] Tariq Iqbal and Laurel D Riek, 'Human-robot teaming: Approaches from joint action and dynamical systems', in *Humanoid robotics: A reference*, eds., A. Goswami and P. Vadakkepat, 2293–2312, Springer, (2017).

[19] Walter Isaacson, *The Innovators: How a Group of Hackers, Geniuses, and Geeks Created the Digital Revolution*, Simon & Schuster, 2014.

[20] Anna Jordanous, 'Evaluating evaluation: Assessing progress and practices in computational creativity research', in *Computational Creativity: The Philosophy and Engineering of Autonomously Creative Systems*, eds., Tony Veale and Amílcar F. Cardoso, 211–236, Springer, (2019).

[21] Anna Kantosalo and Hannu Toivonen, 'Modes for creative human-computer collaboration: Alternating and task-divided co-creativity', in *Proc of the Int Conference on Computational Creativity*, pp. 77–84, (2016).

[22] Pegah Karimi, Jeba Rezwana, Safat Siddiqui, Mary Lou Maher, and Nasrin Dehbozorgi, 'Creative sketching partner: an analysis of human-ai co-creativity', in *Proc of the Int Conf on Intelligent User Interfaces*, pp. 221–230, (2020).

[23] George J. Klir and Tina A. Folger, *Fuzzy sets, uncertainty, and information*, Prentice Hall, 1988.

[24] Antonios Liapis, Georgios N. Yannakakis, Constantine Alexopoulos, and Phil Lopes, 'Can computers foster human users'creativity? theory and praxis of mixed-initiative co-creativity', *Digital Culture & Education*, **8**(2), 136–153, (2016).

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

[26] Jon McCormack, Toby Gifford, Patrick Hutchings, Maria Teresa Llano Rodriguez, Matthew Yee-King, and Mark d'Inverno, 'In a silent way: Communication between ai and improvising musicians beyond sound', in *Proc of the CHI Conf on Human Factors in Computing Systems*, pp. 38:1–38:11, (2019).

[27] Marvin Minsky, 'Music, mind, and meaning', in *Music, mind, and brain: The neuropsychology of music*, ed., Manfred Clynes, 1–19, Springer, (1982).

[28] René Mogensen, 'Swarm algorithm as an improvising accompanist: an experiment in using transformed analysis of george e. lewis's "voyager"', in *Proc of the 1st Conf on Computer Simulation of Musical Creativity*, (2016).

[29] Bhavya Mor, Sunita Garhwal, and Ajay Kumar, 'A systematic literature review on computational musicology', *Archives of Comp. Methods in Engineering*, 1–15, (2019).

[30] Music Tech Fest. `https://musictechfest.net`, 2019.

[31] Jean-Jacques Nattiez, *Music and discourse: Toward a semiology of music*, Princeton University Press, 1990.

[32] Brian O'Neill, Andreya Piplica, Daniel Fuller, and Brian Magerko, 'A knowledge-based framework for the collaborative improvisation of scene introductions', in *Int Conf on Interactive Digital Storytelling*, pp. 85–96, (2011).

[33] Vilayanur S. Ramachandran and Edward M. Hubbard, 'Hearing colors, tasting shapes', *Scientific American*, **288**(5), 53–59, (2003).

[34] Robert Rosen, *Anticipatory Systems: Philosophical, Mathematical & Methodological Foundations*, Pergamon Press, 1985.

[35] Robert Rowe, *Interactive Music Systems: Machine Listening and Composing*, The MIT Press, Cambridge, MA, 1993.

[36] Robert Rowe, 'Split levels: Symbolic to sub-symbolic interactive music systems', *Contemporary Music Rev.*, **28**(1), 31–42, (2009).

[37] Gijsbert Stoet, 'PsyToolkit: A software package for programming psychological experiments using linux', *Behavior research methods*, **42**(4), 1096–1104, (2010).

[38] Belinda Thom, 'Interactive improvisational music companionship: A user-modeling approach', *User Modeling and User-Adapted Interaction*, **13**(1-2), 133–177, (2003).

[39] Oscan Thörn, Peter Fögel, Peter Knudsen, Luis de Miranda, and Alessandro Saffiotti, 'Anticipation in collaborative music performance using fuzzy systems: a case study', in *Proc. of 31st Swedish AI Society Workshop*, (2019).

[40] Oscan Thörn, Peter Knudsen, and Alessandro Saffiotti, 'Human-robot artistic co-creation: a study in improvised robot dance', in *Int Symposium in Robot and Human Interactive Communication*, (2020).

[41] William F Walker, 'A computer participant in musical improvisation', in *Proc of the ACM SIGCHI Conf on Human factors in computing systems*, pp. 123–130, (1997).

[42] Nan-ning Zheng, Zi-yi Liu, Peng-ju Ren, Yong-qiang Ma, Shi-tao Chen, Si-yu Yu, Jian-ru Xue, Ba-dong Chen, and Fei-yue Wang, 'Hybrid-augmented intelligence: collaboration and cognition', *Frontiers of Information Technology & Electronic Engineering*, **18**(2), 153–179, (2017).