THINKING BEYOND THE HUMAN: DESIGN APPROACHES FOR ROBOTS AND AI IN OPERA

EVELYN FICARRA
Sussex University, Brighton, UK
E.J.Ficarra@sussex.ac.uk

TIM HOPKINS
University of Sussex, Brighton, UK
T.Hopkins@sussex.ac.uk

ELIZABETH JOCHUM
Aalborg University, Aalborg, Denmark
jochum@ikp.aau.dk

CHRIS KIEFER
University of Sussex, Brighton, UK
c.kiefer@sussex.ac.uk

The publication of this article was made possible by the **Tele-Encounters: Beyond the Human project**, co-funded by the European Union's Creative Europe programme. The authors' research was presented during the "Beyond the Human: Theatre, Robots and Social Realities" International Symposium. Project coordinator: "George Ciprian" Theatre (Romania), with financial support from Buzău County Council. Project partners: UCAM Universidad Católica San Antonio de Murcia (Spain) and Industria Scenica (Italy). Dissemination partner: UNATC "I.L. Caragiale" (Romania).

Abstract: Opera has historically been a site for developing new and emerging technologies. More recently, opera has proven a rich environment for investigating future scenarios involving human-robot interaction and exploring the potential of robot performers. The explosion of enthusiasm for AI and the ubiquity of machine learning tools have created numerous possibilities for generating creative content in music, poetry, images, and choreography. This article outlines efforts to explore the creative possibilities of robots and AI through a dedicated workshop with researchers from diverse fields. Our goal is to begin thinking beyond the human to investigate what robot performers might look and behave like when designed from a relational rather than a representational approach. We provide a brief history of robots on stage, including the authors' prior work developing original performances featuring robots and human performers. We describe our iterative design approach and low-fi prototypes, summarizing key insights from the workshop. Finally, we analyze the workshop outcomes and consider the trade-offs of working with existing robots and the possibilities for working with open, configurable systems.

Keywords: robots, opera, dance, AI, puppetry, musical robots.

How to cite: Ficarra, E., Hopkins, T., Jochum, E., and Kiefer, C. (2023) "Thinking Beyond the Human: Design Approaches for Robots and AI in Opera," *Concept* 1(26), pp. 129-151.

Introduction

Live performance has long been a site for experimenting with new and emerging technologies. This is especially true for opera, where technology has historically functioned as both a catalyst for spectacular scenery and also as a tool for augmenting and extending human performance (Hensen 2016; Topey 2009). Corneille's 1650 opera *Andromede* used machines to represent gods, and Robert Lepage's 2012 production of Wagner's *Ring* cycle at the Metropolitan Opera House featured a giant robotic set: a platform made of twenty-four independently controlled, movable planks that functioned as a dynamic set piece. In the past two decades, robots have taken center stage as performers (Park 2012). Live performance has emerged as a site for developing and showcasing robot technologies, staging visions of human-robot futures, and exploring new configurations of human-robot interaction (Jochum 2016; Reilly 2019; Sone 2017). Advancements in robotics and AI are quickly expanding the possibilities for human performance and, in some

cases, creating a new class of performers. What, we ask, can contemporary robotics and AI technologies, together with new digital methods of singing synthesis, offer to modern-day opera making? How might these technologies inspire a new generation of opera composers, directors, choreographers, librettists, and scenographers and open up new creative dimensions for opera, both dramaturgically and musically? During this period of rapid development in robotics, how might opera serve as a testbed to explore novel modes of human-robot interaction (HRI) and collaboration, and help drive discoveries that could potentially feed back into robotics research and development?

This article describes our investigation into new ways of designing and programming robots for the stage that move beyond representation. We organized a three-day workshop with experts from different disciplines to explore *what happens when we look beyond the human?* Instead of using the human as the model for building social robots, what happens if we design robot performers along the axes of difference? We briefly summarize our previous work developing original dance and opera performances featuring robots and human performers. Using the workshop as a case study, we describe our iterative design approach for developing low-fi prototypes for robotic performers, summarizing key insights and presenting sketches for novel robot performers. Finally, we analyze the outcomes of the workshop and consider the trade-offs involved with working with existing robots and the possibilities for working with open, configurable systems.

Robots on Stage

On stage, robots typically function as metaphorical humans, frequently appearing as actors in scripted plays that reprise conventional narratives and tropes from science fiction about machines that rise up against humanity to take over the world (Poulten 2014; Reilly 2019). As Louis Chude Sokei observes, the stories we tell about robots and their material instantiations have shaped our understanding and expectations around technology throughout history: "The very notion of the machine or even industry would evolve from merely material descriptions, tools, and objects into signs and metaphors of abstract social and historical processes" (Sokei 2016, 36). Robots on stage always enact two, interconnected processes: the material instantiation of a machine (a system comprised of metal, wires, silicon chip, and sensors controlled by a computer) and its metaphorical resonances (the robot as symbol of what the machine might represent at a given time) (Jochum 2013). Robots on stage are typically pre-programmed or teleoperated to create the illusion of autonomy, but from a technical standpoint they function more like puppets or automatons (Jochum and Murphey 2014). Like their automaton predecessors that toured Europe during the Enlightenment, androids and other

humanoid technologies that populate theatre stages today both capture the public imagination and function as scientific arguments (Schaeffer 1999; Kang 2011).

With few exceptions, robots on theatre and opera stages largely imitate human behaviors, including speech, dialogue, and gestures. Even with recent advancements in AI and sensing technologies, our thinking about robots—and humanoids in particular—has not changed all that much in the last three hundred years. Contrary to Donna Haraway's speculation in the *Cyborg Manifesto*, computers and the recent wave of AI technologies have not ushered in revolutionary new ways of thinking about what robots might look like or how they might inhabit our world. As popular performances of android theatre and other humanoid robots on stage reveal (Jochum 2013; Poulten 2014; Sone 2017; Jochum and Derks 2019; Reilly 2019; Kaegi 2019), even when robots are not particularly realistic or human-like, robot characters still function as metaphorical humans. Although a robot's physical appearance might look very different from a human being, when it comes to spoken dialogue or singing, the default choice is usually to use either a human voice or an "electric" sounding voice to animate the robot. Dialogue is often pre-recorded or tele-operated, and the robot's voice is usually given very little attention.

Through our workshop investigations, we aimed to uncover new approaches for designing and programming robots that avoid imitating humans. Instead of imitating outward behaviors—what we call the *representational* approach—the *relational* approach models processes such hearing, listening, and responding.[1] We were also interested in exploring the notion of virtuosity for robot performers, and were inspired to imagine what interactions between people and robots might look like beyond the imitation paradigm. We tasked ourselves to rethink what the relations between human performers and robot performers might be, not along the axis of familiarity, but along the axes of difference. Reimagining robot design and interaction from the relational approach might leverage a robot's unique traits and affordances, freeing the robot to perform on its own terms. What does vocal virtuosity mean for a robot? What other possibilities for robot voice might there be? These initial questions were the starting point for our workshop.

Robot Opera at University of Sussex

The Robot Opera[2] research project has been running at the Centre for Research in Opera and Music Theatre (CROMT) since 2017, with support from the Sussex Humanities Lab. Initiated by Evelyn Ficarra, the team of collaborators has included

1 We do not refer to Bourriaurd's relational aesthetics, but more by Jack Burnham's writing on System Esthetics. See Burnham 1968.

2 https://www.sussex.ac.uk/research/centres/centre-for-research-in-opera-and-music-theatre/research/projects/robot-opera

Ron Chrisley, Tim Hopkins, Ed Hughes, Chris Kiefer, Carol Watts, Kat Sinclair and others. The research project centres on creative exploration of human robot interaction, including how new technologies might contribute to novel forms of creative expression and challenge our understanding of voice, embodiment, and performance. Our experimentation was attended by close, scholarly reflection on topics of robotic otherness, artificial intelligence (AI), and human/robot hybridity, including the ethical dimensions of relationships between people and robots.

At the time of writing, the project has had a total of three research events, each of which featured live performances with robots followed by scholarly discussions. *Robot Opera* (2017) featured two Nao robots with performances of *O, One* (music and libretto by Evelyn Ficarra) and *The Opposite of Familiarity* (music by Ed Hughes, libretto by Eleanor Knight), and scholarly presentations by Ron Chrisley, Chris Kiefer and Thanos Polymeneas-Liontiris. *Robot Opera – What's Next?* (2019) was hosted by the Sussex Humanities Lab, featured a Pepper robot, a robot cello, a robot cat and two human performers. *Robo_Po /// Robo_Op* (2021), included robot poetry reading and an operatic performance featuring Cleo, a Mesmer realistic humanoid robot designed and built by Engineered Arts (Figure 1).

Figure 1. *Robo_Po /// Robo_Op* (2021) included robot poetry reading and an operatic performance featuring Cleo, a Mesmer humanoid robot designed and built by Engineered Arts. The Robot Opera research project has been running at the Centre for Research in Opera and Music Theatre (CROMT) since 2017, with support from the Sussex Humanities Lab. Initiated by Evelyn Ficarra, with Ron Chrisley, Tim Hopkins, Ed Hughes, Chris Kiefer, Carol Watts, Kat Sinclair and others.

The 2021 events were streamed live to an international audience and concluded with a discussion panel of artists, academics, and industry professionals.

The main questions guiding the Robot Opera project are: how would a robot sing if it sang like a robot, rather than emulating a human? How might the robot's own materiality produce an authentic voice? On a more philosophical level, when a robot sings, who is singing? What values, aesthetic choices, and biases are manifest when a robot sings? We know from prior experience with tele-operated androids and other humanoid social robots that human operators are at the center of much decision making. Important factors concerning how a robot looks or sounds, specific design choices (e.g. the options of male/female voices, regional accents, the interfaces for programming, recording and playback) involve implicit assumptions and decisions that have tremendous impact on the resulting behavior and interaction. Sometimes, these design choices are hard-wired into the system, and can become constraints when trying to work in new or unexpected ways with the robot. Teleoperation and pre-programming techniques ensure that robots perform safely and reliably on stage, but the limitations of working with ready-made robots (e.g. proprietary robots designed with specific purposes and applications in mind) often only come into full view after the creative team has spent significant time working with the robots during rehearsal. Theatre artists may nevertheless work around these constraints in creative and unexpected ways.

For the first three robot opera experiments, using the Nao, Pepper, and Mesmer robots, we were interested in "going under the hood" to work directly with the robot's own voice. From an operatic point of view, the initial obstacle for all three cases was that none of the robots were designed to sing. The robots' text-to-speech (TTS) software had been optimized for intelligible human speech, but not formatted for singing or other types of vocal expression. We had the option to import pre-recorded vocal material, but we wanted the robot sing with its own voice—the same voice it used for speaking. To address this challenge, we first attempted to engage with the constraints of the software. In contemporary music practice, there is a concept called "extended techniques," where a classical or traditional instrument is explored and played in unusual ways to create new sounds. We applied this technique to the robots. For the 2017 Nao project, collaborator Ron Chrisley adjusted the software through Python to control the pitch and rhythm of the robot. For that and subsequent projects, Evelyn Ficarra (not a programmer, but experienced in manipulating sound) spent hours experimenting with the speed and shape controls for these robot voices, and with typing extremely challenging gibberish words for the robot to reproduce, with startling and mostly humorous results. Translating the "extended technique" approach to robots using TTS also informed the 2019 project with the Pepper robot and Choregraphe software. In

2021, Ficarra extended this approach further to include recorded materials from the robot itself, both its speaking voice and the mechanical sounds created by the robot during movement. In addition, Ficarra visited the Engineered Arts robot factory in Cornwall (UK) where the robots are built, to record the sound of the robots and machinery as they were being manufactured. The thought was to incorporate the robot's materiality and physicality into the repertory of vocal sounds and utterances that Cleo (the Mesmer robot) could make. Additionally, Ficarra recorded sounds from a human opera singer (Loré Lixenberg) who was to perform alongside Cleo, to create a mash-up hybrid territory of computer / mechanical / human voice. Cleo is marketed as a 'realistic humanoid' robot. The robot's human hair and silicon skin make the most human looking of the robots used in the project thus far. The hybrid voice we developed for Cleo was tele-operated (Wizard-of-Oz style) during the performance. The guiding aesthetics were intended to reveal the robot's otherness rather than emphasise its human appearance. We found that the expanded technique enabled us to vastly extend the robot's expressive potential, and shifted how we thought about working in the next phase of the project.

While working with Cleo, we frequently encountered the phrase "Oh no one has ever asked us about that before!" from the engineers who assisted on the project. Working with proprietary robots, which are built and developed by private companies, meant that a lot of aesthetic decisions are made in advance. We were somewhat constrained in terms of the level of experimentation and how far we could push these limits. Although it was possible to import new sounds, we could never really "get our hands" on the algorithm controlling the robot's voice—a challenge made more difficult by the fact that the voice function is subcontracted to another company, creating one further remove. This is a general difficulty that artists encounter when working with proprietary tools (and especially software) in creative practice: artists become subject to the underlying logic of the system, and those limitations only become clear when trying to do something unexpected that the original designers did not intend. To address this limitation, and to explore how far we could push beyond the human, we decided to re-think our approach and devise new ways of exploring vocal virtuosity for robot performers. We were also keen to explore the newly-available machine learning tools for generating creative content in music, poetry, images, and choreography. What follows is an account of our efforts to systematically explore what robot performers might look, behave, and sound like when designed from a relational, rather than a representational approach.

Workshop

We organized a three-day expert workshop on robots, AI, opera, and music at the Research Laboratory for Art and Technology at Aalborg University (AAU) together with researchers from CROMT. Experts were selected for their diversity of profiles and expertise and willingness to explore. The primary goals of the workshop were: 1) knowledge sharing, 2) idea generation, and 3) practical development towards joint collaboration and research on robots, AI, music, opera, and embodied interaction. The workshop was funded by a grant from the AI Research Cluster and the AI for the People Bridging Projects at Aalborg University.

Participants

The 11 participants were all professors and artists from three universities in Europe and UK, spanning various disciplines including computer music, composition, opera directing, AI and machine learning, embodied interaction, human computer interaction, robotics, puppetry, media technology, dance, theatre, and interaction design.[3]

Workshop Structure

During the three-day workshop, we shared our research through hands-on/body-on activities and group work to investigate prototypes for artificial performers that could perform alongside human performers. The workshop was held in a converted power plant, which functions both as an arts center and teaching venue, and is located in close proximity to the fab-labs. Workshop participants were encouraged to stay the entire three days, with some attending for only a few sessions. A core group of nine researchers participated for all three days. While some of the researchers had previously worked with one another, many did not know each other before arrival.

The workshop was divided into eight sessions, each of which was themed and loosely facilitated, and there was a progression so the activities built on top of the previous session (Figure 2). The idea was to form small groups of 3 or 4 researchers that would develop small material investigations through physical materials and software that would result in a lo-fi prototype. Session #1 was instrumental in setting the tone for the subsequent sessions and workflow. The evaluation during Session #7 was important for reflecting and summarizing the important take-aways to inform our future work.

3 Participants from AAU: Elizabeth Jochum, Daniel Overholt, Brian Bemman, Chen Li, Dimitris Chrysostomou, Cumhur Erkut; from Sussex University: Evelyn Ficarra, Chris Kiefer, Tim Hopkins; from University of Gothenburg: Kivanc Tatar, and independent artist Valeria Rizzo.

Day 1	Day 2	Day 3
Session #1 Knowledge sharing	**Session #4** Lo-fi prototyping	**Session #7** Evaluation
Session #2 Body-on workshop	**Session #5** Stop-motion prototyping	**Session #8** Next steps
Session #3 Research talks	**Sesssion #6** Mini-showings & reflections	

Figure 2. *Workshop schedule*

Session 1: Knowledge Sharing. This session asked each of the participants to present a 10-15 minutes long exercise, hands-on tutorial, or demonstration of their artistic or research practice. This could take the form of leading participants in exercise utilizing a piece of software, tool, instrument, or group activity. Slide presentations were welcome, but the focus was predominantly on teaching/sharing through hands-on demonstrations. Workshop participants provided some basic background and introduction to the activities, but the focus was primarily on knowledge sharing through doing. The goal of the introductory session was to generate a shared "tool box" / "toy box" that we would use in our material investigations towards the development of an artificial performer. Presentations included Dall-E, poetry generated using Chat GPT, experiments with drawing robots, algorithmic music presentation, physical and vocal improvisation, live coding demonstration, puppetry, and devised theatre/staging. We used the "sand box" metaphor from design research where people could "play without consequences" and gain experience in building, taking apart, and rebuilding their designs. Below are short summaries of some the exercises from Session #1.

Physical and Vocal Improvisation (Evelyn): I led two warm-up games, the first as a name game/icebreaker, and the second as a mechanism for getting people to listen/observe closely and have fun. Participants are asked to use their bodies as music notation, which is a kind of code, so effectively we are embodying and performing code. The exercise also creates irregular rhythms and keeps people on their toes. The second task is a version of the children's Telephone Game, where someone whispers something around a circle to see how it transforms by

the time it gets back to the first speaker. Instead of whispers, the game is played out loud, with any kind of sound, using your voice. Generally, the sound evolves as it moves around, because humans can't imitate each other exactly, try as they might. This exercise is an elementary lesson in composition, demonstrating how sonic materials develop and transform.

Operatic Voice Instrument (Chris): The idea behind this task was to bring in some themes from recent computer music research and to make an improvisatory group exercise. I presented an interactive voice synthesiser that uses "concatenative synthesis," a technique where a collection of sounds are segmented into short chunks and organised according to various musical qualities (pitch, texture, similarity, etc). Once organised, the player can navigate through these chunks, recombining them into new sounds. Concatenative synthesis isn't a new technique, but it has been revitalised by its inclusion in FluCoMA (Fluid Corpus Manipulation Project) framework (Tremblay 2021), with the benefit of new machine listening techniques.[4] The FluCoMA framework was used to analyse around 40 minutes of opera singing, and organise the sounds on a 2D plane, so that a player could trigger the sounds by navigating a path through them using a gamepad joystick. The session began with a demo and explanation of the system, before moving into group discussion and exploration.

Live Coding (Kivanc): This presentation demonstrated an autoencoder tool that was used in a live performance by Kivanç Tatar titled *Coding the Latent No.1*. The performance brings together live coding and AI audio generation, where the live performance becomes the site to experimentally "hack" AI models by exploring the latent audio space in real time. According to Tatar, "the approach builds on musical improvisation practice, while emphasizing proactivity, over reactivity and immediate musical reactions" (Tatar 2022). Thus, the performance becomes a site for the performer to "hack" the AI model in a live concert setting while also proactively planning compositional layers. The performance is a mixture of planned, unexpected, and disruptive tendencies. This first edition of the *Coding the Latent* series invited the audience to witness a glitchy and noisy world of AI synthesized sounds, spatialized on 47 speakers.

4 The Fluid Corpus Manipulation project (FluCoMa) instigates new musical ways of exploiting ever-growing banks of sound and gestures within the digital composition process, by bringing breakthroughs of signal decomposition DSP and machine learning to the toolset of techno-fluent computer composers, creative coders and digital artists (https://learn.flucoma.org).

Computer Music (Brian): This presentation introduced participants to the work of composer Milton Babbitt, a composer known for composing twelve-tone and serial music that is both complex and highly constrained (Bemman and Meredith 2018). Brian provided a detailed introduction to Babbitt's compositional process which derives from mathematical constructs. He described the techniques Babbitt developed, focusing in particular on the all-partition array, time-point system, and equal-note-value strings used in later works. Then, Brian described a proposed procedure for automating Babbitt's compositional process using these techniques, and demonstrating the procedure to automatically generate an entirely new musical work in the style of the composer. This presentation introduced another way of working with algorithms in musical compositionthat contrasted with the autoencoder tools in the Live Coding presentation.

Chat-GPT and Dall-E (Chen & Dimitris): This session consisted of three parts. Part one was a presentation on industrial Virtual Assistant (VA) for HRI. The presenters introduced recent work carried out in developing an industrial virtual assistant (Li et al. 2023), designed to assist in human-robot interaction (HRI) within the industrial domain. The presentation covered topics such as human attention tracking and industrial task assistance, demonstrating how AI can enhance efficiency and productivity in industrial settings. Part two was a DALL-E Image Generation Demonstration. DALL-E is an AI model developed by OpenAI that utilizes text inputs to generate corresponding images. Participants provided text descriptions, and DALL-E generated images based on those descriptions (Figure 3). This demonstration highlighted the ability of AI to understand and translate textual prompts into visual representations. Part Three used the Shakespearean Poetry Generation using BiLSTM Neural Network. Each participant was asked to write a sentence of their choice. These sentences were collected and used as inputs to feed a Bidirectional Long Short-Term Memory (BiLSTM) neural network. The BiLSTM model had been trained on William Shakespeare's Sonnets, allowing it to learn the style and language of Shakespearean poetry. Using the given input sentences, the neural network generated new poems in the style of Shakespeare. This activity showcased the ability of AI to simulate and generate artistic content in a specific style based on textual inputs.

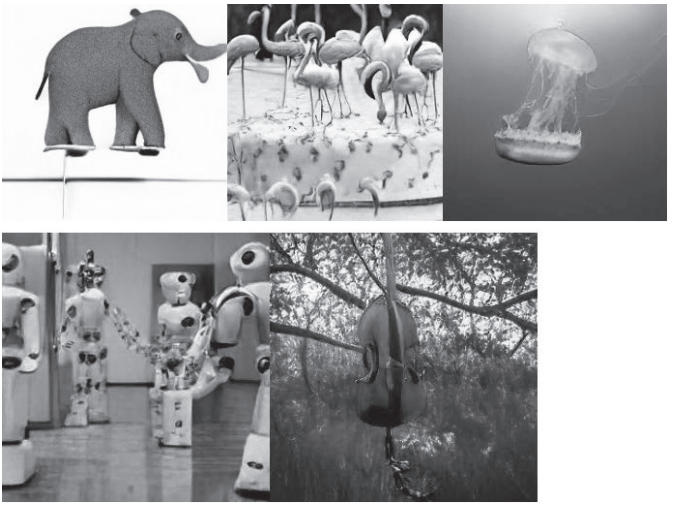

Figure 3. Workshop experiments with DALL-E, an AI model developed by OpenAI. DALL-E utilizes text inputs to generate corresponding images. Workshop participants provided text descriptions, and DALL-E generated images based on those descriptions.

Session #2: Body-On Workshop

The first session consisted in demonstrations of software tools and approach, but the focus was not on physical embodiment (except for the introductory session on voice and sound with Evelyn). The second session maintained a focus on embodied interaction with materials, space, and one another. This consisted of two activities: one on puppetry, and one on devised theater/theatrical composition.

Puppetry (Elizabeth): My session was centered on puppetry. The field of puppetry has a long and rich history of generating expressive movement using relatively simple control mechanisms. We worked with small, soft, deformable stick puppets made of foam. Participants were asked to spend 2 minutes exploring all the different ways that the object could move. They were then asked to work in pairs with their puppets through movement only (no sound or spoken words), to communicate, collaborate, and explore through the affordances of the simple objects. In a very short period of time, stories, characters, idiosyncratic behaviors began to emerge through interaction.

Devised Lyric Theatre (Tim)*:* I invited the group to participate in a quickly-assembled opera staging in the workshop room, set to 2 x 1 min sections of an audio recording of the Overture to Handel's opera *Radamisto* (1720). Before hearing the music, participants were guided/instructed in sequences of movement and intention in two parts: Part One was towards a group activity; Part Two was towards more individual trains of action/thought. Part One involved two groups facing each other across the space and then crossing to the other side, passing by each other repeatedly, gradually developing moments of accidental proximity into moments of contact – involving light touch, eye contact, etc. This kind of activity is sometimes used to explore choreography through a discipline called *contact improvisation* (Koteen and Stark Smith). In Part Two, I introduced individual actions to different parts of the space, in series. Individuals were given specific tasks, moods, or partner relationships, based on very pedestrian movements (e.g. facing a blackboard; crossing to a chair and sitting on it; facing a partner, then turning one's back on them) layering spatial and behavioural relationships over time. We ran parts 1 and 2 as a sequence – first with no music, then again with Handel's music as accompaniment. The exercise was particularly valuable for workshop participants with less experience with performance and gave them a sense of togetherness an understanding of creating structures with rule-based systems enacted through pedestrian movements, and a sense of how music can transform both experience and structure.

Session #3 Research Talks. This session was open to students and members of the public. Invited experts gave presentations on previous works with opera robots, AI and coding in live performance, and robots in music and performance. The talks were followed by a Q&A session which ranged across topics such as the nature of creativity and issues of gender and robotics in fictional/theatrical narratives.

Session #4 Lo-Fi Prototyping & Session #5 Stop Motion Prototyping

We self-organized into groups of three, according to which aspects of the technology, interaction or tools introduced in Sessions #1 and #2 we were most interested in exploring further. Groups documented their processes in sketches, photos and short video clips, and wandered around the room to get inspiration from other groups. The outcomes of these individual working groups are reported in the next section.

Session #6 Mini-Showing and Reflections

We invited a professional performer (Valeria Rizzo) trained in dance, puppetry, and circus to experiment with the lo-fi prototypes. From our prior

experiences with robot performances, we knew the value of working with performers early in the process: it can help open up important conversations about how performers can work relationally with technology, rather than being forced to adapt to the constraints of the system later in the process. We opted to have someone who was not part of the development but who could give meaningful feedback on what we had done so far and suggest avenues we might explore. We have observed that designers/developers can become very attached to their ideas or the vision of the thing they are trying to build, without considering how the device might be viewed differently, or activated differently in performance by a human operator/performer who was not part of the design process. The feedback from the professional performer was not meant to be formal feedback on the design, but more exploratory and open-ended.

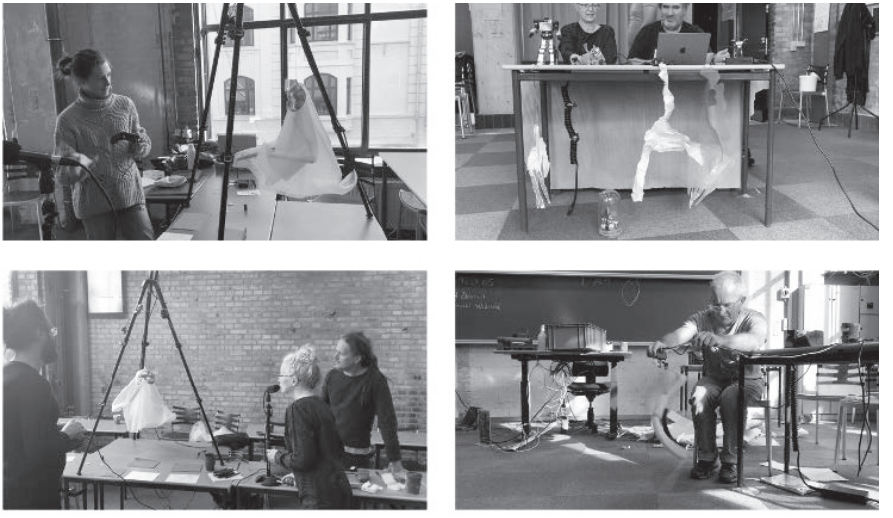

Figure 4. The three working groups produced three different prototypes of robotic systems for performance based on interaction, music/sound and material. Clockwise from top left: Valeria Risso interacts with Isolde; Evelyn and Cuhmur animate the Puppet Panorama; Tim Hopkins with Twisty; Kivanç Evelyn and Chris interact with Isolde.

Prototypes

Three working groups produced three different prototypes of robotic systems for performance based on interaction, music/sound, and material. We titled them: Isolde, the Puppet Panorama, and Twisty. All of the prototypes were based on Arduinos and servo motors to activate the movement (Figure 4).

Isolde (Chris, Dan, Elizabeth). This system was directly inspired by the vocal and physical improv activities in Session #1. We thought about the very first workshop activities and asked – *how do we build a robot that would be able to participate in the improvised warm up activities led by Evelyn* (e.g. the Telephone Game)? Like puppeteers, we thought first about what the robot should *do*, rather than what it should look like or what story it should tell. Participating in the Telephone Game requires that an individual be able to: 1) make a sound based on a sound made by another person (hearing, repeating, vocalizing); and 2) make a gesture on its own, that could be performed at the same time as the sound (gesturing). We opted to work with an adaptation of the Operatic Voice Instrument that Chris introduced in Session #1. The "body" was made of a piece of lightweight netting (tulle) that could be activated very simply with a lightweight motor from underneath or above. We co-opted a tripod and attached it, marionette-like to different points on the tripod, with small sticks attached to fishing wire to move the material. It was important that the material be activated directly by sound (listening) and also capable of generating its own sound. We used the vocal database based on concatenative synthesis based on previously recorded snippets of vocal improvisation from the *Robot Opera*. One performer could navigate the database physically by moving a joystick of a game controller. The sound from the voice synthesiser was analysed using MFCC features, and coefficients from these features were mapped onto the motion of four servo motors. A musical phrase would trigger the robot to move (gesturing) and the movement of the servos would make sound. The motor sounds were picked-up by contact microphones attached directly to the servo motors, amplifying the robot's natural, embodied "voice" (vocalization) and linking the robot's/puppet's physical gesture and vocalization. The motor sounds were analysed using Shannon Entropy[5] of low and high frequencies, and these signals were added to the gamepad motion. This created a feedback loop, creating emergent behaviour patterns in the puppet. None of the movement sequences were very long, the result was short, jerky movements of the motors (made to appear more fluid by the tulle). We later added a microphone, whose signal was mixed in with the motor sounds. This created two modalities for interaction and playing the Telephone Game with the robot—either through navigating the vocal instrument database with the joystick, or directly with their voice (Figure 5). This last feature was important because it moved from a dyadic interaction to a triadic interaction, creating the possibility for more performers to interact with the robot in real time, as in the improv game. At some point we

5 Shannon entropy is part of communication theory that focuses on the minimum amount of information required for transmitting a signal.

attached a carnival mask to the tulle, giving the character a "face" which helped the human performers focus their attention when interacting with the puppet.

Puppet Panorama (Cumhur, Evelyn, Chen): We wanted to explore how simple materials in fairly random shapes could come alive through movement. Chen brought with him a lot of small robot toys, so the piece partly became about the juxtaposition of these "ready-made" toy robots with roughly made abstract puppets. We made three robot/puppets: one out of lightweight packing paper, one from plastic mesh and plastic, one from three strips of coloured paper. We arranged the readymade plastic robots in a line across the table, whilst the homemade paper and mesh robots were suspended from the edge of the table. Cumhur took charge of downloading / installing / testing the software (Dynamixel Wizard 2 and a self-made Python interface available on a university network) to activate the servo motors using the computer. Three puppeteers controlled the puppets: Cumhur was at the computer, Chen was attaching different puppets to the servo motor, and Evelyn was pushing the plastic robots with a stick. Each suspended robot was activated in turn by attaching the body directly to the servo motor, so that it could spin and seem to come to life. (If we had had three motors, we could have done some choreography between them perhaps). Meanwhile, on the top of the table, the plastic robots were moving towards the edge of the table and were eventually pushed off, apart from the dog, who stayed on top, observing the scene. There was a juxtaposition between the toy robots and the more fragile homemade quality of the paper/mesh robots. In the performance, all three dangling robots broke (this was not intentional!) and fell to the floor. We responded by pushing the plastic robots off the table as in a "mass suicide," and the piece came to be about destruction. Ironically, the more "advanced" looking robots (the plastic ones) were animated with very crude technology (pushing the puppets with a stick), while the fragile paper/mesh abstract shapes had the high tech computer manipulation. Valeria commented that when the second robot was spinning they happened to touch the hand of the first robot, and this looked like an intentional moment. This comment made us realise that even chance behaviour of the materials can lead to dramatic action, which is why it is so useful to play with the materials and see how they behave before fixing them into specific shapes. Evelyn was interested in trying to make a "voice" for the robots out of the sounds of their materials. During the software download / operation activity, which can feel like a "lull" as the programmers do their work, Evelyn took the materials (mostly different kinds of paper) to a different space and recorded sounds from them. Ultimately these sounds were not used in the "performance," since we shared the work with Valeria before the sounds were ready to be deployed:

Twisty (Kivanc, Brian, Tim): We began from two directions: making a foam/stick puppet, (from Elizabeth's puppetry exercise), and sourcing a database derived from sound worlds of opera recordings. From this potentially vast field, we drew in particular on music from scenes depicting loss of control or self-destructive autonomy (often featured in opera narratives as a kind of punishment for unsanctioned gratification or intensity—nineteenth century opera has lots of this). Behind this choice lay speculation about how emotional registers supposedly hosted by opera material might manifest in a puppet animated by motors controlled by musical data, and how as animateurs our decisions become dramaturgically consequent. We "robotized" the puppet with a motor. After initial problems getting the motor and control communication to work, the physical characteristics of the robot puppet started to feed back. Rotations acted on the foam unevenly, twisting the lower structure into itself. Valeria noted limitations of the puppet, and we made adjustments, to find similar issues in play, but acting within slightly more interesting physical forms, with more potential. Meanwhile, Kivanç continually modified the opera soundworld in various ways, allowing it to listen to itself and transform further sonic responses by adjustin values and range. These modifications produced a number of expressive effects, for example, solo singing voices might be prominent for a while, then recede into a different, spectral space. Necessarily a bit preoccupied with functionality, we were able even at this stage to sense the potential for these interactions to accumulate significance through different choices about design and the scope of interaction. The puppet had a kind of comic physicality—its contortions speaking to other genres—clowning for example, or Punch and Judy-esque violence. This concept was realised further through Valeria's spontaneous ventriloquising of the puppet, improvising the voice of the puppet as it vengefully struck its operator.

These lo-fi prototypes were not meant as characters in the traditional sense, but rather as material investigations into the nature of sound and movement and how artificial performers and human performers might interact. These performers came to be not because we were trying to imitate human behaviors, but because we were attempting to model different relational processes—*hearing, listening, responding, echoing*. Achieving this meant we had to creatively and jointly *think through* the different materiality of the puppets/robots in new ways, asking questions, sketching ideas, roughly prototyping both puppet bodies and lines of code on the fly. We note how this "sketching" process, largely rooted in improvisation, is fundamentally different from how humanoid robots or scripted robot characters are usually programmed.

Discussion

Our initial workshop marked a first step in a longer development process for a research project on robot voice and opera. In our effort to move beyond the human, we found that approaching the design of robot characters from a relational approach was key to unlocking our creativity and charting new pathways of working in an inspired way across disciplines. We conducted a structured evaluation using the PQRS method (Praise, Query, Reflection, Synthesis). Four themes emerged: 1) Participatory Design; 2) Environment; 3) Materials Matter; and 4) Challenges. We briefly summarise these themes, with the hope that others might carry this work forward and build upon the lessons learned.

Participatory Design. We did not design the workshop structure explicitly with a Participatory Design framework. However, reflecting on our experiences and the outcomes, this turned out to be an important aspect of the workshop. The focus was never on technology *per se*, but on the knowledge sharing of the diverse experts and fields in the room. As one workshop participant noted, the workshop was very much "in the spirit of participatory design—working with people instead of supporting them. This is a small but important difference." Equal space and value were given to technical and non-technical components and exercises, and there was sufficient time to explore each of these topics on their own. There was no rush to start immediately programming a robot, or working with a pre-trained AI model. The composition of the workshop participants was also key to creating this participatory and collaborative atmosphere. Some of the more technically-oriented folks had little experience with opera or theatre, while the puppeteers and theatre makers had limited programming experience. The hands-on activities put everyone on a similar level from the start, so participants could start exploring together, rather than dividing up into "tech" and "non-tech" camps. This collaborative spirit came as a pleasant surprise to many of the participants, many of whom had come with different expectations:

- *I expected the practical work would be different. Never expected to be programming an arduino.*

- *I do not normally work with electronics much even though I am expert in it, but there is always so much to do on the software side. Now I have had hands-on experience of how fun it could be. What's nice is that I can imagine getting my hands on things and looking into them. Trying out. And I have an intention to get into robotics and art, even though I didn't expect it.*

The prompts and activities introduced in Session #1 generated many ideas and tools that were picked up and developed generatively throughout by the different groups. The focus on having hands-on activities (as opposed to only presentations), and tools that were "at hand" to work with on the spot, was extremely helpful in launching the groups and creating an atmosphere that encouraged experimentation. Also important was the range of presentations and expertise in the room: we didn't just start with a set of AI tools or robot platform (in fact, there were no arduinos/servo motors in the room when we started—we figured out along the way that we needed them).

Environment. The physical environment of the workshop was important to establishing collaboration and also making people comfortable exploring how to work in new ways with new materials. The room was a flexible space with tables and chairs that could be rearranged on the fly. It was also important that we were all in the room together, and the space was large enough so each group had ample working area, and there were also corners and places for participants to have some down time, either to work on code, work through an idea, or just relax in the presence of the flurry of workshop activity and participate passively at some times. The time scales of working with theatre, choreography, puppet making, and coding all vary, and it is important that participants can feel connected to the work. It was key that participants felt they had agency to either "lean-in" or "lean-back" during different throughout the workshop days. The room was within walking distance of a lab where we could check out materials and tools (soldering irons, arduinos, speakers, a tripod). In hindsight, this was key to our successful prototypes, because we could realise some ideas that we would not have been otherwise able to.

Materials Matter. All of the workshop leaders had experience with existing (and in some cases very sophisticated) robot platforms used for commercial development (such as Nao, Pepper, Sawyer, Baxter, Mesmer, UR3 and UR10 industrial robot arms). The decision to work with bottom-up approaches, animating very rough and crude materials (mesh, paper, foam, tulle, and string) to investigate their movement and sonic/vocal potentials, opened up more space in the studio/workshop to think about strategies of interacting.

- *Tech-wise, it was nice that there were a lot of pre-made, ready-made stuff in the room, and the bits to collect them together. That enabled us to work, and to be able to get creative, and start to fine tune it.*

In another way, we had too many ideas to choose from and some ideas and tools that were introduced were left by the wayside. For example, Chen and Dimitris introduced Chat GPT exercise with specific poetry, keywords, and prompts. Evelyn later experimented with composing some text scores based on the results of the experiments in Session #1, but there was not enough time to explore them fully.

Challenges. Bringing disparate research fields together and bridging methods and practice is not without challenges. Most of the challenges related to how to best structure the time that allows for different kinds of creative work, and how to best document the workshop outcomes so that we can continue working and iteratively develop over a series of workshops. How do we balance the time scales needed for development of new processes and prototypes, with the activity of combining/recombining existing technologies? In all creative work, everyone needs down time. Can we allocate a specific table or space for rest/reflection, with materials that help participants dream something on to the page, to help keep the creative energy and ideas flowing? We found that it takes substantial time to generate a shared understanding and shared vocabulary. For example, are we really in agreement about our understanding of what a voice is? Is voice something that we make with our bodies, like a cricket does? Or is it a mouth? Does voice always imply words and communication, or can it just be sound? The same is true for metaphorical understandings; we might describe that something electronic sounds like "*this*…", but these definitions might differ from person to person. Taking time to listen, watch, share, absorb, test and discuss materials and processes together can help to produce shared reference points and define creative directions. Workshops that hope to break new ground require time for pondering and exploring; we need to continually ask ourselves, "Is this what we want? What else could there be?"

Conclusions

Based on our prior work with "ready made" robot platforms and experimental work with choreography and opera, we set out to develop a workshop format that would allow us think beyond the human, and to explore new tools with the idea of designing our own prototypes. Reviewing the outcomes of the three-day workshop on robots, AI, opera, and music, we identified the trade-offs involved with working with existing robots and the possibilities for working with open, configurable systems. Robotics companies invest huge resources and person hours in designing and building sophisticated robots with capabilities that are impressive

on many levels. However, the tradeoff is a certain rigidity and opacity in design and capability. The aesthetics of these robots are fixed, and for a creative team devising new operatic work, this can have a huge impact on the dramaturgies and operatic language (both musical and theatrical). Creating our own performing robots in an exploratory way can be enormously freeing and exciting. At the same time, it is also resource and time intensive, and it can feel like there is a mountain to climb before we might get to something that could work on a professional stage. Reflecting on our experience with both approaches, we conclude that building robotic performers in a relational way, with input from a multidisciplinary team, from the ground up, asking questions about our materials and starting from basic principles of music, movement, and expression is hugely rewarding. This open process builds on the creativity of the participants in a generative way, which we believe is more likely to yield original and unexpected results and to create genuinely new forms, dramaturgies, and music for opera. The workshop method allowed us to explore this task in a relational way, by modelling processes and exploring through the material, rather than imitating explicit human behaviors or starting with a narrative, concept, or script.

Robots onstage present new dramaturgical considerations and provide a site for rethinking and reconfiguring the evolving relationship between humans and machines. In order to move beyond the human, we have to get under the hood and model examine the underlying paradigms that shape human-machine interactions in creative ways. This is the goal of our experimental workshops: to explore existing technologies to challenge traditional hierarchies and working practices in both opera and AI through exploratory and participatory design methods.

References:

1. Bell, J. (2001) *Puppets, Masks and Performing Objects*. Cambridge: MIT Press.
2. Bemman, B., and Meredith, D. (2018) "Generating new musical works in the style of Milton Babbitt," in *Computer Music Journal*, 42(1), pp. 60-79.
3. Burnham, J. (1968) "Systems Esthetics," in *Artforum*, September 1968, pp. 30-35.
4. Francis, P. (2011) *Puppetry: A Reader in Theatre Practice*. Bloomsbury.
5. Henson, K. (editor) (2016) *Technology and the Diva: Sopranos, Opera, and Media from Romanticism to the Digital Age*. Cambridge: Cambridge University Press (Cambridge Studies in Opera).
6. Hoffman, G. and Ju, W. (2014) "Designing Robots with Movement in Mind," in *Journal of Human-Robot Interaction*, 3(1), pp. 89-122.
7. Jochum, E. (2013) "Deus Ex Machina: Towards an Aesthetics of Autonomous and Semi-Autonomous Machines," Doctoral Dissertation, University of Colorado. https://scholar.colorado.edu/concern/graduate_thesis_or_dissertations/qf85nb374
8. Jochum, E., Millar, P., and Nuñez, D. (2017) "Sequence and Chance: Design and Control methods for entertainment robots," in *Robotics and Autonomous Systems*, 87, pp. 372-380.

9. Jochum, E., and Derks, J. (2019) "Tonight We Improvise! Real-time tracking for human-robot improvisational dance," in *MOCO'19: Proceedings of the 6th International Conference on Movement and Computing*, pp. 1-11.
10. Jochum, E., and Murphey, E. (2014) "Programming Play: Puppets, Robots, and Engineering," in Posner, D., Orenstein, C. & Bell, J. (editors), *Routledge Companion to Puppetry and Material Performance*. Routledge, pp. 308-321.
11. Kang, M. (2011) *Sublime Dreams of Living Machines*. Cambridge: Harvard University Press.
12. Li, C., Hansen, A.K., Chrysostomou, D., Bøgh, S., and Madsen, O. (2022) "Bringing a Natural Language-enabled Virtual Assistant to Industrial Mobile Robots for Learning, Training and Assistance of Manufacturing Tasks," in *2022 IEEE/SICE International Symposium on System Integration* (SII), pp. 238-243.
13. Nikolaidis, R., and Weinberg, G. (2010) "Playing with the masters: A model for improvisatory musical interaction between robots and humans," in *19th International Symposium in Robot and Human Interactive Communication*, Viareggio, Italy, 2010, pp. 712-717.
14. Poulten, C. (2014) "From Puppet to Robot: Technology and the Human in Japanse Theatre," in Posner, D., Orenstein, C., and Bell, J. (editors), *Routledge Companion to Puppetry and Material Performance*. Routledge, pp. 308-321.
15. Reilly, K. (2019) "Robots and Anthropomorphism in Science Fiction Theatre: from Rebellion to Domesticity and Back Again," in Wyants, N. (editor) *Media Archaeology and Intermedial Performance Deep Time of the Theatre*. London: Palgrave Macmillan, pp. 193-210
16. Schaeffer, S. (1999) "Enlightened Automata," in Clark, W., Golinski, J., Schaffer, S. (editors), *The Sciences in Enlightened Europe*. Chicago: Chicago University Press. pp. 126-165.
17. Sokei, L.C. (2015) *The Sound of Culture: Diaspora and Black Technopoetics*. Wesleyan University Press.
18. Sone, Y. (2017) *Japanese Robot Culture: Performance, Imagination, and Modernity*. New York: Palgrave Macmillan.
19. Tatar, K., Cotton, K., and Bisig, D. (2023) "Sound Design Strategies for Latent Audio Space Explorations Using Deep Learning Architectures," in Proceedings of Sound and Music Computing. arXiv:2305.15571
20. Topey, Peter (2009). "Disembodied Performance," Master's Thesis, MIT.
21. Tremblay, P.A., Roma, G., and Green, O. (2021). "Enabling Programmatic Data Mining as Musicking: The Fluid Corpus Manipulation Toolkit," in *Computer Music Journal,* 45(2), pp. 9-23.

Online references:

22. Koteen, D., and Stark Smith, N., (2023) *Contact Quarterly*, [Online]. Available at: https://contactquarterly.com/contact-improvisation/about/ (Accessed: 3 August 2023)
23. Park, E. (2011) *The Ring Transformed* [Online]. Available at: https://www.metopera.org/user-information/nightly-met-opera-streams/articles/the-ring-transformed/ (Accessed: 3 August 2023).
24. Tatar, K. (2022) *Coding the Latent*. [Online]. Available at: https://kivanctatar.com/Coding-the-Latent (Accessed: 11 August 2023).

Performance references:

25. *The Uncanny Valley* (2019) Directed by Stefan Kaegi [Gessnerallee, Zurich].
26. *O, One* (2017) Music and libretto by Evelyn Ficarra. Directed by Evelyn Ficarra [Centre for Research in Opera and Music Theatre, Sussex].
27. *The Opposite of Familiarity* (2017) Music by Ed Hughes, libretto by Eleanor Knight. Directed by Evelyn Ficarra [Centre for Research in Opera and Music Theatre, Sussex].

28. *Robot Opera – What's Next?* (2019) Music and libretto by Loré Lixenberg, Anton Lukoszevieze and Sam Bilbow. Directed by Evelyn Ficarra [Centre for Research in Opera and Music Theatre, Sussex].
29. *Robo_Po /// Robo_Op* (2021) Music and libretto by Evelyn Ficarra. Directed by Tim Hopkins. [Centre for Research in Opera and Music Theatre, Sussex].

Evelyn Ficarra is Professor of Interdisciplinary Composition (Music) and Co-Director of the Centre for Research in Opera and Music (CROMT) at the University of Sussex. Composer, sound artist and educator; Evelyn's work finds expression across a range of forms including music theatre, multimedia, installation, dance, film and the concert hall. Primary research interests include: investigations into sound objects in different contexts (theatre, installation); electroacoustic and multi-disciplinary improvisation, and robot opera. Evelyn started the robot opera research project at CROMT in 2017, exploring issues of embodiment, performance and vocality through the lens of radical technical innovation in opera.

Tim Hopkins is Senior Research Fellow, Music, University of Sussex, and has a background in opera, music theatre and theatre direction, scenic design, moving image, performing, in the UK, Europe, and US. He joined Sussex in 2007 as AHRC Fellow in the Creative and Performing Arts. He is interested in opera and music theatre as a technological adventure and societal forum, researching new forms in performance, experimentation and collaboration, exploring how the art-form responds to technologies and their meanings. His own work often links historic and current practice, new spaces technology has created and the evolving artistic toolbox, to explore impacts of technological transformation itself—cultural, emotional, political—and imagine possible futures for opera.

Elizabeth Jochum is an Associate Professor and leads the RELATE Research Laboratory for Art and Technology at Aalborg University (Denmark). Her research uses the visual and performing arts as catalysts for re-thinking how we design and implement robots and other assistive technologies. From industrial robots to exoskeletons, Dr. Jochum's work involves transdisciplinary collaboration in human-robot interaction, health, and engineering to develop creative, value-sensitive, and human-centered approaches to ensure the technologies we build address the real needs of people. She coordinates the European ABRA project on Artificial Biology, Robotics and Art, and the ImprovAIze project on machine learning, dance improvisation, and wearable technologies. She was named one of the top 50 Women in Robotics by Robohub in 2021.

Chris Kiefer is a Senior Lecturer in Music Technology, Experimental Music Technologies Lab (EMUTE) at University of Sussex. Chris is a computer-musician and musical instrument designer, specialising in musician-computer interaction, physical computing, and machine learning. He performs with custom-made instruments including malleable foam interfaces, touch screen software, interactive sculptures, and a modified self-resonating cello. His work concentrates on machine learning and complex systems in new instrument design. He co-runs the AHRC Feedback Musicianship Network.