

Sound of Contagion – An Artistic Research Project Exploring A.I. as a Creative Tool for Transmedial Storytelling

Wenzel Mehnert, Robert Laidlow, Chelsea Haith & Sara Laubscher

Introduction

Sound of Contagion (SoC) is a transmedial art project addressing the cultural narratives surrounding global diseases and pandemics through different media. It is borne out of the collaborative partnership between the University of the Arts in Berlin and the University of Oxford and explores the use of Artificial Intelligence (A.I.) in creative processes. We used texts written by the machine learning algorithm GPT-2, which was trained on narratives about pandemics from the last 2500 years. From the resulting text fragments, we built a storyworld that served as a basis for narratives, illustrations, musical compositions and a lecture performance.

Besides the aesthetic and creative output, the project is a research endeavour that explores the usage of A.I. as a creative tool to facilitate interdisciplinary working. Through the practical collaboration between artists, researchers and technology, the imaginary that A.I. can independently and consciously create art is shifting towards the notion that A.I. becomes a tool that inspires and encourages collective-creative processes. Thus, the work follows a collaborative paradigm and acknowledges A.I. as a creative tool with its own unique aesthetic.

This article is an insight into the project. It is a case study of how A.I. is used as a creative tool within an interdisciplinary collective working in the fields of literature studies, cultural studies, composition and visual arts.

About the project

The project Sound of Contagion (SoC) understands itself as a practice-based research project that uses artistic practices as an approach to understand A.I. as a creative tool. It is borne out of the Oxford / Berlin Creative Collaborations activities (supported and funded by the University of the Arts in Berlin (Germany) and the

University of Oxford (UK)), which fosters projects and research across the arts and humanities, and supported by the Minderoo Foundation through the Minderoo-Oxford Challenge Fund.

From the beginning, we focused on an interdisciplinary collaboration, as each member had a different artistic and academic background: Chelsea Haith, with a background in publishing, journalism and gender studies, and Wenzel Mehnert, with a background in cultural studies and narratology, represented the humanities. Sara Laubscher, working as a professional illustrator, and Robert Laidlow, composer, researcher and AI-expert, represent the artistic side. Through the interdisciplinary nature of the project, it started with several different research questions that were guiding the thought process without limiting the creativity and necessary openness of an artistic research project:

How can A.I. be used as a tool within an interdisciplinary, creative project? How do we acknowledge the unique agency of the tool? How do we integrate it within a human-technical collaboration?

The project started in the beginning of 2020, shortly after the Covid19 virus was first reported in Europe. We took this global transition into the unknown future of lockdowns and social distancing as the starting point for our endeavour. Through the international distribution of our project partners and the lockdown situations, we were forced to set up the project solely through digital communication technologies like Zoom and Miro. Some of us had not met in person, but only digitally.

Dataset, Data & Processes

Dataset

Cognitively occupied with the pandemic, the project became an opportunity to process and deal with the shifting situation through the experiences. We decided to work with narratives of previous generations and how they dealt with pandemics, epidemics or similar global phenomena. Principally led by Chelsea Haith at this stage, we selected texts, all dealing with pandemic and plagues, as the dataset to train the A.I. and curated the selection from the last 2500 years of pandemic literature. We permitted novels and plays only. One reason was the focus on narratives and reappearing tropes that could hint at cultural patterns, the other was, that we were interested to see, what a machine learning algorithm might do with these. The selection of dataset already had interesting consequences for the outputs generated.

We divided the texts into two data sets – Sophocles to the year 2000 was one data set, and texts published between 2000 to 2020 the other. This splits the total dataset roughly in half. This split provoked its own insights into cultural preoccupa-

tions in the last two decades as well as the increased popularity and thus production of dystopian and apocalyptic narratives in popular literature in the same period. As a total data set, these texts serve as cultural artifacts and represent the collective memory from which we draw and confront the changing environment. The earliest of our selection was David Mulroy's translation of Sophocles' *Oedipus Rex* (Sophocles, 2014 [429 BC]), while the latest was Lauren Beukes' *Afterland* (2021). These texts were used to train the A.I. model with the aim to generate new text fragments based on the dataset.

Data

To create an A.I. model that could generate new texts, based on existing texts referring to pandemics, we first had to make two decisions:

What would comprise the dataset upon which the A.I. was trained?

Which A.I. algorithm would we use?

The choice of training data has been shown to make an enormous difference to the output of generative A.I. algorithms, even when other parameters are left the same (Leeming 2022; Mehri et al. 2017; Caillon & Esling 2021). Similarly, two different A.I. algorithms trained on the same kinds of dataset can provide entirely distinct generations; compare, for example, 'WaveNet' (Van den Oord et al. 2016) trained on Beethoven sonatas with e.g., PRiSM-SampleRNN re-implementation trained on the same (Melen 2020).

We first decided upon the choice of A.I. algorithm, which was to be GPT-2 Simple (Woolf 2019). GPT-2 Simple is a reimplementation of the popular and well supported GPT-2 A.I. algorithm (Radford et al. 2019), developed and released by OpenAI. We chose this algorithm because of its capacity to be fine-tuned to produce specific types of text, while still utilizing rules it has learned through general training. In this context, fine-tuning means that we began with a GPT-2 model that had already learned the 'rules' of English and then we directed this model specifically to learn to apply this to the context of our dataset. We felt this made it superior for the task we had in mind to other contemporary algorithms, which relied only on the fine-tuned dataset (e.g., WordRNN). We also had prior experience with this algorithm (de Roure et al. 2019) which helped in speeding up the early process. Finally, GPT-2 Simple is easily accessible online through Google Collab, and does not require a GPU on a personal computer to train. This allowed each team member to access it with their own machine during COVID-19 lockdowns, facilitating easier discussion and collaboration.

For the choice of dataset, we decided upon a set of 21 pandemic-related fictional texts, taken from across the last two millennia (as described above). This decision was reached as a compromise between wanting a substantial amount of data to train GPT-2 Simple on and our interest in being able to – or at least, thinking we were able

to – draw links between specific GPT-2 Simple generations and specific works in the dataset. We felt that identifying the A.I. algorithm's reference points would assist in the world-building and narrative probes element of the project.

Training the Models

We initially trained GPT-2 Simple several times using this dataset, creating a variety of models. Most important parameters included the base type of model that we then fine-tuned (124M or 355M), how the dataset was divided, and for how many steps the model was left to train before we sampled generations from it. From this initial experimentation and iteration, we chose a set of parameters we felt were most suitable for our project. This decision was made collaboratively and subjectively, according to our qualitative judgment of the models' text generations.

To create the text, we made the following decisions:

We used the 124M model as a base. Generations from this baseline model tended to reflect more strongly the ideas of pandemics that our dataset focussed upon.

We split the dataset (in full, 12MB in size), into two halves. These were texts written prior to the year 2000 (8.1MB in size) and post-2000 (5.7MB in size). Each dataset was therefore more homogenous than the full dataset, which allowed for finer control of text generation.

We trained the models for 4000 steps on Google Collab.

Once we had created the two fine-tuned models (pre-2000 and post-2000), we sampled text generations from them. These generations were both 'seeded' and 'not seeded'.

In this context, seeded means that we provided the beginning of a sentence and the A.I. algorithm continued from that point. Not seeded meant that the model was freely generating text. Seeding texts allowed us to 'point' the model towards talking about specific elements of pandemics, whereas not seeding allowed us to see how the model would 'write' without a specific context, after having been fine-tuned on a pandemic dataset. Generations also have a 'temperature' between 0 and 1 which is a rough indication of how much randomness is introduced into the system. The 'temperature' of the generations we created was generally 0.8 as this gave us interesting texts that made, in our opinion, at least some kind of sense. There was some variation here, with some generations used being as high as 0.95 or as low as 0.5.

On balance, the fine-tuning method seemed to us to be a success in that the models would consistently write about plagues, fevers, illness, chaos, death, destruction, displacement and other related themes, even when they were not seeded. While the models did not appear to overfit (copy exactly from the dataset), it was sometimes clear that some parts of the dataset had a greater prominence than others in its generations. A stand-out example would be Stephen King's *The Stand* (1990 [1978]) in the pre-2000 dataset, which was particularly long at 1152 pages.

Through this iterative process of training, sampling, dataset adjustment, and further training, we began to get a feel for how the two fine-tuned models wrote, and how we might get different kinds of responses out of them. We instructed the two models in exactly the same way, providing us with a contrasting pair of generations for each seed. The post-2000 model, for example, was prone to creating successive lines of dialogue. We used several texts from this model that were seeded with a single quotation mark and with single-word pronouns. The pre-2000 model more often described landscape in more poetic language. From the pre-2000 model we utilized more generations seeded with phrases such as 'the pandemic', 'the plague', and 'when it began'.

With the training of the models complete and several hundred generations sampled, the project moved to the next stage involving world-building, narrative probes, and transmedial collaboration.

A.I. as a creative tool

One of the guiding principles of the project was to treat the results of the A.I. as a unique form of creativity and acknowledging the creative agency of the algorithm with its own flaws and failures as a new aesthetic form. One of these "failures", in the context of fictional writing, is that the text itself does not hold any sense – as the author did not mean to create any sense in the first place. One figure of thought that helped us to understand our relation to this text comes from Roland Barthes famous essay *The death of the author* (1978 [1968]). Here Barthes proclaims that the origin of a text does not lie with the author, but with its destination: the reader (142 ff.). Barthes starts with a description of an author which he argues against:

The Author, when believed in, is always conceived of as the past of his own book: book and author stand automatically on a single line, divided into a before and an after. The Author is thought to nourish the book, which is to say that he exists before it, thinks, suffers, lives for it, is in the same relation of antecedence to his work as a father to his child. (145)

In our case, there is no author – It is not even dead as it has never been alive – and therefore no message or moral to convey is purposefully inscribed in the text. Furthermore, the A.I., particularly the GPT2 algorithm we used, still lacked a semiotic understanding. It did not make connections between the characters, or the places mentioned in the text fragments. Rather, the A.I. produced text fragments that are similar to the narratives used as training data, without understanding the context or the connections between the elements of the narratives. It recombined the passages

we fed it with, and the results are an arbitrary combination of different writings, like Barthes describes in the following:

A text is made of multiple writings, drawn from many cultures and entering into mutual relations of dialogue, parody, contestation, but there is one place where this multiplicity is focused and that place is the reader, not, as was hitherto said, the author. The reader is the space on which all the quotations that make up a writing are inscribed without any of them being lost; a text's unity lies not in its origin but in its destination" (148).

In the context of fictional writing, this means that the meaning of a text is given by the reader, who constantly recreates the sense and searches for coherence – this is particular true when reading text fragments created by a GPT2 algorithm, as they not only lack meaning but also a form of coherence. The non-coherent text created by the A.I. forces the reader to constantly seek for a meaning and brings the reader into the active role of interpretation and drawing connections. This is the unique aesthetic of text fragments produced by a GPT2 algorithm. Thus, as the A.I. was not capable of doing it, the practice of sense-making and creating coherence had to come from us.

Worldbuilding

For this reason and following the guiding principle described above, we considered all the text fragments as *Narrative Probes* (Fischer & Mehnert 2021). Narrative Probes are flash fiction like narratives that all probe into the same possible world, following a specific logic that we had to uncover. The narratologist Marie-Laury Ryan uses the metaphor of the text as windows to the extralinguistic realm of characters, objects, facts, and states of affairs serving as referents to the linguistic expressions. (Ryan 2001: 91) Reading and getting immersed into literature transports your mind into this extralinguistic realm, a world that constantly formed in your head as you read on. Ekman & Taylor call this process *readerly worldbuilding* and constitute that the reader carries out a construction of a world by adding pieces to a structure already in place evaluating the pieces in light of what is known about the world so far, and, conversely, evaluating the whole world in light of new pieces (Ekman & Taylor 2016: 11). With this perspective in mind, we treated the text fragments as a multitude of windows that showed us snippets of a storyworld that we as readers had to piece together. Therefore, through a structured deconstruction of the text fragments, we engaged in a worldbuilding process based on the textual cues that we found in the fragments, loosely building on the work by Wolf (2012) and Mehnert & Fischer (2021):

In the first step we detected the elements of the world. For this process, five cues were of specific importance to us: Places, events and references to time, characters mentioned, values and norms as well as specific objects.

In the second step, we enriched the elements with additional information that either came from the text or was already given by the logic (e.g. when the text refers to a leader, there is also a group of people following).

The third step was to create connections between the enhanced elements (e.g. if two characters were mentioned in two different text fragments, and we assume they both exist in one world, what is the connection between the two?).

Example

To give an example, we will explain the process with the following two text fragments:

Fragment I:

They came to a real leader, they said to each other. A real man, not a fake leader, not an actor who pretended to be a leader. He was real. He lived, and he led his people to a real leader. This real man was the one they called their real leader. He was named Mannum, and Mannum knew this about himself; he had practiced seduction and charm the long time. He was a good liar, a real one. The pandemic of war had begun.

Fragment II:

"Then, friend", said Piranesi, "make death eternal in the liquor laws, so that no sick person may return to drink after consuming what they drank".

In Fragment I we learn two important things about the world: (1) In this world exists a character called *Mannum*. We can say that *Mannum* is not an honest leader and a good liar, thus he is not to be trusted and rather an evil character. (2) We also have an event, the *pandemic of war* that stands in connection to *Mannum* as it has been mentioned right after. So in the world that we want to explore, there is some kind of a biological war going on but we do not know why and who else is involved.

In Fragment II, two characters are mentioned. The first character, named *Piranesi*, is talking about *making death eternal*. When enhancing this character, we assumed that *Piranesi* is a chemist who was responsible for creating a deadly disease. The second character in this Fragment is *friend*. We decided that *friend* is a reference to *Mannum*, the leader from Fragment I, by which we connected both fragments. Based on the cues given in these two text fragments a narrative about a conspiracy plotted by *Piranesi* and *Mannum* to create a biochemical weapon to start a war is formed.

Outcomes

Nested Narratives

Working with A.I. in such a way can be seen as an intriguing source of inspiration. It gives you material to play around with, to dig in, to take out, to transform, to iterate and to make sense of. The result of our sense-making process was a nested narrative, meaning a story within a story, within a story. Each part plays in a different time and features different characters, however, they are still connected with each other.

- Layer 1: 17th Century // Piranesi & Mannum plotting a war
- Layer 2: 1981 // A girl searching for her mother
- Layer 3: 2021 // The writer, trying to stay sane

The first part is a short story from the 17th century about the two characters introduced above, Piranesi & Mannum. Both are plotting a war using a biochemical weapon that fails and extinguishes Mannums kingdom and the people living within. In the second part, we follow the story of a little girl called Emily. At the beginning of the story, she reads a comic about the story of Piranesi & Mannum, this is how the first narrative is nested within the second. While she reads the story, a TV news anchor is presenting the latest development of an actual pandemic that is raging all over the world. We situated this situation in the year 1981 and the pandemic, which the news anchor is referring to, is the AIDS pandemic. Emily's narrative is again nested within a third narrative. The last one plays in 2021. Here we have text fragments about a person that we called The Writer. He is similar to our real lived experience in 2021. Due to the Covid-19 restriction and the social lockdown, he stays alone in a small apartment in New York. To stay sane, he starts to write a story about a young girl called Emily, which is how the second narrative is nested within the third. Furthermore, the writer mirrored our own experience, as we were in lockdown during the production phase of this project as well.

Illustrations

As mentioned in the introduction to this project, Sound of Contagion follows a transmedial approach to storytelling, thus we transformed the narratives into different other mediums beyond text. The first transformation was based on the illustration by Sara Laubscher. We picked key scenes of each nested narrative as prompts for Sara, who illustrated them in her own style. She got inspired and created the second output of our project.

The first illustration shows Mannum and Piranesi, sitting in an old library plotting how to use the biochemic weapon. In the second illustration we see Emily who is

sitting in front of a television set. The caption of this image is the report of the news anchor. We read it as the voiceover of the newscast. This section indicates that it was drawing from the part of the data set coming from the 20th century rather than from the Renaissance. The discourse around virus transmission, the effect on different geographical regions, and the conceptualisation of the pandemic as a weapon all point to the effect of social collapse during pandemic conditions. At the same time, we can identify reappearing tropes that often come up when we are dealing with topics of contagion, plagues or pandemics. The Writer, who became a point of self reflection but also the mastermind of our narrative, is depicted in the next illustration. Here we see the lonely writer sitting in his small apartment, haunted by his memories and trying to stay sane by coming up with stories about pandemics from different centuries.

Fig. 1: Illustrations from the Sound of Contagion project. Copyright: Sara Laubscher



As mentioned above, the illustrations were the first transformation of the story-world. The next transformation came in the form of music.

Musical Adaptation

The Sound of Contagion text was further adapted into several pieces of music by Robert Laidlow, and one by composer Marco Galvani. These include *The Writer*, premiered online during COVID-19 lockdowns and the piece *Disc Fragments*, both performed by Bandwidth Ensemble. The following is an insight into the creative process of Robert and how he reflects the process of working with the text fragments.

For a ‘Sound of Contagion’ event at Oxford University in November 2021, I (Robert Laidlow) composed a piece for tenor and synthesizer called *Disc Fragments*, which used texts generated as part of this project. This piece is in seven movements, the first three and last three being somewhat symmetrical – that is, Movement 1 is similar to 7, 2 to 6, and 3 to 5. This arch form was a parallel to the nested narrative form that Chelsea Haith assembled from ‘GPT-2’ texts we had generated earlier. The texts for six movements (all except Movement 4) were generated by fine-tuned ‘GPT-2’. These texts dictated the form and material of the music; here, the first movement, *At Delphi*, is discussed as an example. There was a certain mysticism that intrigued me in this text:

Fragment III:

At Delphi

A Prayer

Let not the Impossible Him

The Impossible Him take what he has done

Take it,

Steal it,

put the word in your mind,

take it,

the other choice.

The use of ‘Delphi’, ‘the Impossible Him’, and ‘put the words in your mind’ implied to me that ‘GPT-2’ could be read as imitating an oracle. I wanted the music to create a sense of ritual and of unseen pattern, similar to (for example) Messiaen’s *Quartet for the End of Time*. The voice part is constructed from a series of interlocking patterns. Ten pitches repeat, superimposed on a rhythmic pattern repeating every nine notes. The text is simply applied to this pattern. The synthesizer is set to an organ voice and uses chorale-style material generated by AI, using an algorithm developed

by Omar Peracha (<https://omarperacha.github.io/make-js-fake/>). This repeats itself every twenty-three beats.

In this A.I. generated text, among others, I was struck by its sudden ending. When 'GPT-2' is generating, the user instructs how many characters to generate. Due to its transformer architecture, 'GPT-2' does not plan forward. It only looks back to what has already been generated whenever it generates a new token. When it reaches its arbitrary character limit set by the user, it simply stops. It cannot plan a 100-character length 'story', for example, because it is not able to plan forward. In this movement the music simply stops when the text comes to an end, which does not coincide with the end of any pattern described above. The inspiration for such a technique comes also from composers such as Birtwistle (i.e., *Carmen Arcadiae Mechanicae Perpetuum*) and Edmund Finnis (i.e., *The Air, Turning*).

Lecture Performance

As already hinted at in the paragraph above, the last transformation of the work came in form of a lecture performance at the University of Oxford in November 2021. The performance was a mixture of different media and perspectives. The first part introduced the process, the creation of the fragments and explained the worldbuilding process. In the second part we read the narratives, presented the illustrations, and performed the musical pieces on stage. Throughout the Performance we presented different perspectives of the project through different means of interaction: while Robert and the musicians were on stage, both, the reading of the fragments done by Wenzel as well as a contextualisation of the fragments by Chelsea, came via zoom. After the presentation, there was a Q&A session in which the audience got invited to discuss with the performers, pose questions to the project and engage in a discussion.

A.I. for transmedia collaborations

Through this project, we found that A.I. acted as a useful way to facilitate collaborations across different media and disciplines. We were each able to use algorithms generating text and images, despite each having very different levels of previous experience with algorithmic processes or computer programming in general. This allowed us to contribute meaningfully to artistic areas that were not naturally our own. Going forward, it seems as though A.I. might be useful as a means to 'try' another collaborator's discipline. This might allow collaborators to much more quickly engage in higher-level discussions and decision-making.

In addition, using A.I. as part of a collaborative artistic research project inherently mandated many creative decisions to which we all contributed. During the

training portion these included decisions such as what comprised the dataset or datasets, the type of model we were going to fine-tune, what our preferred parameters were when training the models, and how we were going to sample generations from these models. Later, they included which elements of the generations we would keep or discard, and how much (if at all) we would adapt or edit them. Since there was no dedicated human writer on the team, these decisions were made collectively.

Utilising A.I. as part of a creative process also differed from a normal collaborator in the sense that A.I. does not fundamentally understand the context of what it is writing. While it can create plausible-sounding sentences and stories, it is not aware that it is doing this, and there is no intentionality behind any of the ‘creative’ decisions it makes. This is very different to most human writers. We therefore felt empowered to make radical decisions when treating the text, including transforming it into entirely different media, but also felt liberated in leaving it in its found-state, even if that found-state was nonsensical and narratively underdeveloped. None of our readings of the material it generated could be ‘right’ or ‘wrong’. This is especially true of the GPT-2 model we were using when compared to models released closer to the time of this paper, which have shown significant improvement in understanding underlying context of text and ability to develop logical statements (Griffiths 2022).

Conclusion

In this project, we explored the use of A.I. as a creative tool for transmedial storytelling. The project points to the cultural-semantic flaws, which are peculiar for the creative results of A.I. works and fills this gap with the sensemaking work done by human actors. SoC challenges the myth that A.I. could replace authors, poets or other creators of art, entertainment or fiction in general. Instead, by using them in a creative process, A.I.s make an important contribution to new art forms, foster inspiration for artistic practices and thus promote independent modes of expression.

We started with a GPT-2 algorithm, trained on a vast amount of fictional literature about pandemics and plagues, that generated several incoherent text fragments. Through the process of worldbuilding, linking the fragments and creating an imaginary world, we acknowledged the unique aesthetic of the A.I., its flaws and particular form of nonsense-writing, and made the tool part of the creative collaboration between different disciplines and practitioners. By documenting the process, we showed that A.I. is not an autonomous actor, that creates art on its own, but instead, that it has the potential to become a tool in an emerging toolbox and that it needs new artistic approaches to make use of this tool and shape it to its own desires.

Last but not least, this insight into the process also serves as an inspiration for continuous, creative engagements with A.I. software. As past events have shown – and surely future development will proof – A.I. technology becomes a fixed part in

our everyday life and in the workflow of designers and artists. Therefore, It is necessary to explore approaches of adaptation, appropriation and domestication within new routines and creative processes.

Bibliography

- Barthes, Roland (1978): *Image-Music-Text*. Hill and Wang.
- Beukes, Lauren (2021): *Afterland*. Penguin.
- Caillon, Antoin / Philippe Esling (2021): "RAVE: A variational autoencoder for fast and high-quality neural audio synthesis". *ArXiv*. arXiv:2111.05011 (cs.LG).
- De Roure, David / Emily Howard / Robert Laidlow / Pip Willcox (2019): "Alter: An Ensemble Work Composed With and About AI" *Proceedings of DMRN+14: Digital Music Research Network One-Day Workshop 2019*. Queen Mary University of London.
- Ekman, Stefan / Audrey Isabel Taylor (2016): "Notes Toward a Critical Approach to Worlds and World-Building". *Nordic Journal of Science Fiction and Fantasy Research* (3)12.
- Fischer, Nele / Wenzel Mehnert (2021): "Building possible worlds. A speculation based research framework To reflect on images of the future". *Journal of Futures Studies* 25(3): 25–38.
- Griffiths, Max (2022): "Is LaMDA sentient?" *A.I. & Society* <https://doi.org/10.1007/s00146-022-01559-z>.
- King, Stephen (1990): *The Stand*. Doubleday.
- Leeming, Zakiya (2022): *Conception to Construction: Compositional decision-making informed by theories, patterns and processes in the physical, life, and computer sciences*. Doctoral thesis, the Royal Northern College of Music in collaboration with Manchester Metropolitan University.
- Mehri Soroush / Kundan Kumar / Ishaan Gulrajani / Rithesh Kumar / Shubham Jain / Jose Sotelo / Aaron Courville / Yoshua Bengio (2017): SampleRNN: An Unconditional End-To-End Neural Audio Generation Model. *ArXiv*. arXiv:1612.07837v2 (cs.SD).
- Melen, Christopher (2020): *PRiSM SampleRNN*. RNCM PRiSM. <https://www.rncm.ac.uk/research/research-centres-rncm/prism/prism-collaborations/prism-samplelenn/> (21 January, 2024).
- Radford, Alec / Jeffrey Wu / Rewon Child / David Luan / Dario Amodei / Ilya Sutskever (2019): "Language Models are Unsupervised Multitask Learners". <https://d4mucfpsywv.cloudfront.net/better-language-models/language-models.pdf> (21 January, 2024).
- Ryan, Marie-Laure (2001): *Narrative as Virtual Reality: Immersion and Interactivity in Literature and Electronic Media*. The Johns Hopkins University Press.
- Sophocles (2014): *Oedipus Rex*. CreateSpace Independent Publishing Platform.

Van den Oord Aaron / Sander Dieleman / Heiga Zen / Karen Simonyan / Oriol Vinyals / Alex Graves / Nal Kalchbrenner / Andrew Senior / Koray Kavukcuoglu (2016): *WaveNet: A Generative Model for Raw Audio*. ArXiv. arXiv:1609.03499v2 (cs.SD).

Wolf, Mark J. P. (2012): *Building Imaginary Worlds*. Routledge.

Woolf, Max (2019): *GPT-2-Simple* <https://github.com/minimaxir/gpt-2-simple> (21 January, 2024).

Discography

Birtwistle, Harrison (1978): *Carmen Arcadiae Mechanicae Perpetuum*.

Finnis, Edmund (2016): *The Air, Turning*.

Messiaen, Olivier (1941): *Quartet for the end of time*.