ALTER

Robert Laidlow (Composer) leading the PRiSM (Practice & Research in Science & Music) team

Alter for mezzo-soprano and ensemble utilises several generative machine learning algorithms as collaborative and interactive tools in the compositional process. It combines models in the symbolic-generative (MuseNet), audio-generative (WaveNet), and text-generative (WordRNN and GPT-2) domains to create a musical structure defined by the machine learning process, where the results of neural networks at different stages of training exist on a spectrum between being showcased without alteration to being radically transformed by the composer.

Important to note is that *Alter* was written as one of several commissions focussing on the life and writings of Ada Lovelace. The piece was premiered in 2019 at the Barbican Centre as part of an event titled 'Imagining the Analytical Engine', and as such also has a connection with Babbage's theorised Analytical Engine, which Lovelace postulated might one day be capable of writing music of its own. These connections will be utilised as the catalyst for much of the work's narrative and machine learning processes.

Alter's structure is a direct embodiment of the machine learning process. The piece is divided into three parts, during which each of the algorithms undergo some audible transformation or 'upgrade'. The musical surface also reflects these alterations, becoming less inhibited and more exploratory as the work progresses – for example, the first section features micro-repeats and stuttering, channelling the sound of algorithms at the beginning of their training process, which are liable to become stuck in a groove. These sections begin roughly at Bars 2, 64, and 150 and are usually joined by or preceded by an interlude featuring the Lovelace Engine instrument (see later). In terms of the work's narrative, it follows the development of a fictional artificial mind which becomes increasingly self-aware.

The piece makes use of OpenAI's MuseNet algorithm in various ways. In 2019 Christine Payne kindly finetuned MuseNet on a selection of my previous scores for a project with the BBC Philharmonic Orchestra. I found this algorithm which had attempted to learn and compose in 'my style' to have enormous compositional potential – far beyond the remit of that particular project – so I have employed it once again throughout *Alter*. MuseNet's generations provide what I call a 'hidden layer' throughout the work: a compositional backbone that informs every decision made in the piece, though the musical surface itself is not usually the verbatim output of a machine learning algorithm. More specifically, it provided several harmonic seeds which I orchestrated, magnified, repeated or transformed in other ways throughout. These seeds were created by priming MuseNet on a piece for solo harp by composer John Thomas (1826-1913) which was dedicated to Ada Lovelace, who sponsored his study at the Royal Academy of Music. While the prompt was Thomas, MuseNet was instructed to continue in my style (which does not sound like Thomas), resulting in a fascinating array of outputs that nonetheless seemed to return to certain ideas again and again. It is these ideas that I isolated and turned into *Alter*'s hidden layer of machine learning-generated music. The most prominent of these is the alternation between the Ab(11) and A(#11) which is present in some form throughout, but is particularly audible from Bar 86 onwards and around Bar 205. More subtly, it can be found underpinning structural harmonic motion i.e. the move from tonal centre A to Ab from the beginning of the piece to Bar 36.

In harmony with the piece's narrative structure, use of MuseNet transforms and becomes more prominent as the work progresses. By Bar 181, its generations are not only providing a hidden layer but also an audible layer to the work's surface. This duet between soloist and her pre-recorded self is composed by prompting MuseNet with a melody, choosing subjectively from its generations and then layering that generation and the initial prompt atop one another to create an interwoven structure. Later, from Bar 215, the musical content of the work is taken over by MuseNet, as each instrument as voice performs its own line generated by the algorithm independently, creating a complex machine learning polyphony.

In the audio domain, DeepMind's WaveNet algorithm is used in various guises throughout the piece's electronics track. This largely corresponds with *Alter*'s narrative structure of gradual self-realisation. In the piece's introduction, I sampled WaveNet's example outputs of the algorithm speaking in invented languages, as heard on its release blog. These were cut up into small phrases and into bodily sounds (throat, lips, breathing) generated by the algorithm, and reconstituted into the track. The acoustic music of the first section that follows continues this trend through becoming repeatedly stuck on individual notes, syllables or ideas. Later in the piece, WaveNet is utilised to render to speech text that was generated by machine learning (see later), and still later I applied a vocoder to the rendered speech. In this way, I hope to show a development from nonsensical language to real speech to singing. As mentioned, the final sung duet between singer and electronics is pre-recorded and is not rendered using artificial intelligence: in the future, this is an element of the piece I would like to change. At the time, I did not have access to an algorithm such as OpenAI Jukebox that could convincingly render a given text as sung melody. Either way, this journey is heard throughout the piece.

Alter's text was entirely generated by machine learning utilising different datasets and approaches to priming. The generations of each section were primed by the section(s) that preceded it, which we felt gave the work an overall tonal consistency. The piece's first and second sections are mostly generated using the Word-RNN algorithm, trained at first only on Lovelace's correspondence, and then additionally with a larger dataset of nineteenth-century writing. This increased size in dataset is noticeable between parts one and two, though there was a considerable amount of cherry-picking by to whittle down the enormous generations to something that is singable in just a few minutes. The text in the work's third section is generated by OpenAI's GPT-2. I was particularly taken with these generations, which are not cherry-picked, because of the way they appeared to force GPT-2 to act. Usually, GPT-2 produces sensible enough output, veering towards the safe 'article' or 'conversation' style of text. As can be seen from the work's text, priming it with the output from another (less sophisticated) algorithm appeared to push it in a very strange direction which was artistically exciting. Despite their abstract and elusive nature, the generations clearly pick up upon and expand the themes of self-awareness and self-doubt that were present in the cherry-picked Word-RNN outputs forming sections one and two. With the recent release of GPT-3, it would be fascinating to input these GPT-2 generations as prompt and see what happens.

A discussion of *Alter* cannot be complete without mention of the 3D-printed percussion instrument designed specifically for the piece: the Lovelace Engine. During the composition of *Alter*, we brought together experts from many fields – music, computer science, history, engineering – in the same way that Lovelace herself brought together her expertise across disciplines to inform her work. An early idea we pursued was that of creating a new instrument with modern technology, inspired by Lovelace, Babbage and the Analytical Engine. Its relationship to machine learning in the piece is one partly of theatricality: the Lovelace Engine plays whenever the algorithms are 'upgrading' themselves between sections. In this way, I wanted to link the very physical, steam-powered, notion of computing that Lovelace and Babbage thought would be the future, with the actual technology that can generate music and text today. In this way, the Engine serves to show how forward-thinking Lovelace was, but also as a warning that today's hype may have no effect on tomorrow's technology.