MATHEMATICS **MATTERS**

Listening with a digital ear:
How maths is changing the way we understand music

There is a major revolution happening in how people access music. In 2017, global digital revenues for recorded music were up by 19.1%, while physical revenues were down by 5.4%. Even more startling is the fact that streaming revenues rose by 41.4%, with Spotify alone having 83 million paying subscribers and over 159 million active users, and relative newcomer Apple Music already having 50 million paying subscribers. This digitalisation of music has only been made possible by recent technological advances, powered by a combination of increased computing power and applications of advanced mathematics.

One of the world's leading music research teams is based at the Centre for Digital Music at Queen Mary University of London. One member of the research team is Professor Elaine Chew who uses her knowledge and skill as a virtuoso pianist to inform her mathematical research into music.

> "One of the world's leading music research teams is based at the Centre for Digital Music at Queen Mary University in London."

Music Generation
There is widespread interest in the problem of generating music using artificial intelligence (AI). Applications of such technologies include game music that can adapt in real-time to the player's actions, or the mood of different scenarios, whether it be swashbuckling with Inigo Montoya or smelling the flowers with Ferdinand the Bull. A technical difficulty lies in making graceful musical transitions between contrasting scenarios. Another challenging application is to create artificial musical partners that can improvise reactively and coherently in partnership with human musicians. The technologies tend to draw heavily on human input, learning from human-made music patterns that sound good, and replicating these patterns in the music generation process.

In the last few years, more start-up companies specialising in music generation like Jukedeck and Melodrive have popped up, and technology leaders like Google have launched their own projects like Magenta, which uses deep learning techniques to develop computer-generated music and other forms of art. Deep learning approaches tend to use large databases of music to develop new computer generated music—this can have some promising results, however statistical learning from a large body of music is also prone to producing music which can sound like a bland average. Researchers, including members of Professor Chew's team, have worked on projects that overcome this problem using a variety of techniques.

Sparse Networks
The first project involved training the computer algorithms by using only a small amount of data to preserve a musician's individual style—taking an approach contrary to the current trend for using big data. The computers were only given a small number of pieces of music to analyse, meaning that they were working with what is known as a sparse network. The team then filled in the gaps using probabilistic rules as well as the rules of music theory. Using smaller amounts of data results in pieces of music that have far more distinctive character and depth than some of the pieces which have been generated by analysing huge libraries of pre existing music.

Inspired by Bach (or Haydn or Kabalevsky or Stravinsky …)
In a more recent music generation project the team took a completely different approach. One of the current challenges of computer generated music is that AI can produce single melodies, which are comprised of a stream of single notes, or single sequences of chords, but finds it extremely difficult to generate a longer piece of complex polyphonic music. For instance, it is difficult to generate convincing music that has simultaneously interacting

Institute of mathematics & its applications

streams or parts moving independently, or which has higher-level forms of organisation such as repetition and narrative structure throughout the whole piece.

In the MorpheuS project (the Marie Skłodowska-Curie Action fellowship project of Dr. Dorien Herremans at Professor Chew's lab), they overcame this problem by generating a new piece of music, based on a score by an existing composer. They took a piece by Bach (or Haydn or Kabalevsky or Stravinsky...), and used a computer to copy the entire rhythm, assigning random notes in the place of the original notes, while maintaining any repeated patterns within the piece. Assigning random notes produced unpleasant sounding music, so the research team used what is called a 'tension profile' to constrain the note assignments. Tension in a piece of music is an expression of how tense the piece sounds - this might sound like a vague and unquantifiable concept, however the research team have successfully developed a way to capture this mathematically, so that it can be described with accuracy and in a rich and multi-faceted way. Figure 1 shows the optimisation process as it iterates from a random set of notes towards ones that closely match the target tension profiles.

MorpheuS then took the piece of music, which had random notes and changed it, bit by bit, until it came close to matching the tension profile of Bach's original piece by using a mathematical optimisation algorithm called a variable neighbourhood search. The result was an extremely convincing piece of original music that was immediately playable, which had been developed using AI, but had used a pre-existing piece of music as its inspiration.

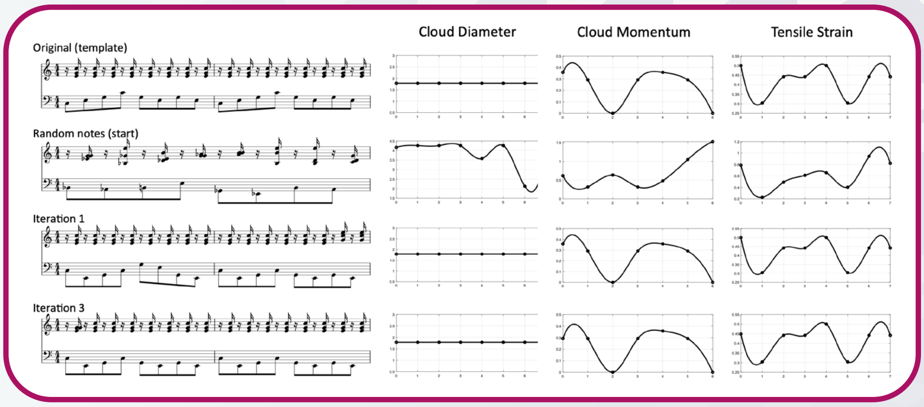

▲ Figure 1. The optimisation process employed by MorpheuS takes as input a template piece, then randomly re-assigns all the notes, and iterates towards the template tension profiles (cloud diameter, cloud momentum, and tensile strain) using a variable neighbourhood search algorithm. The MorpheuS project is described in youtu.be/l_qBcUjnbzo.

Capturing expression within music

One of the very human aspects of music making is that a player can impart a great deal of expression into a piece during a performance. Two different pianists might be playing the same score, while at the same time producing two very different interpretations of the same piece. If a computer was to automatically play a piece of piano music from a score, it would undoubtedly sound lifeless and flat compared to a real live performance. Until now, this type of expression within music has been very difficult to quantify, let alone analyse.

Professor Chew and her research team have developed ways of representing these subtleties of expression by using a variety of mathematical models which capture the different changing features of a performance such as timings, tension and loudness. Once they have been captured numerically, it enables detailed analysis of a performance in a way which hasn't been done before. It has been possible to analyse extreme timing changes within a piece, comparing them with changes in tension, enabling researchers to quantify when a musician is able to delay the playing of a particular note, perhaps for dramatic effect. Figure 2 shows where time is taken as it corresponds to the level of dissonance in the music. This type of knowledge will help musicians to learn from other performers in a scientific and quantifiable way, improving the theory informing music education. It will also help general listeners appreciate better the kinds of expressive nuances musicians introduce in performance.

Expressivity in music mimics rhythms and gestures in the real world such as physical movements, speech or birdsong, and physiological timings. Thus, mathematical models of musical expressivity can also help us make strides towards understanding and modelling movement and control, the timing of expressive speech, as well as biological rhythms.

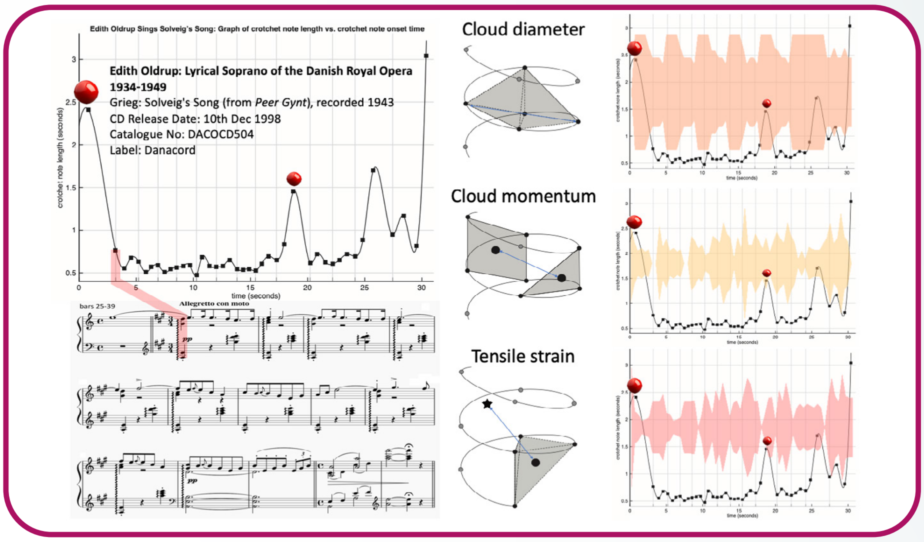

▲ Figure 2. Beat durations in Edith Oldrup's rendition of a segment of Edvard Grieg's "Solveig's Song" -- musical tipping points are marked by red cue balls -- and the tension profiles (cloud diameter, cloud momentum, and tensile strain) of this section of the score to help explain the singer's expressive choices. For more examples of tipping points, see https://bit.ly/TippingPoints-playlist

Using the research to understand abnormal heart rhythms

In recent work, Professor Chew and her team have shown that abnormal heart rhythms have musical properties and can be represented in ways similar to music, opening up new ways to analyse electrocardiographic data. Figure 3 shows how the abnormal heart rhythms of atrial fibrillation can be encoded using music notation and embedded naturally into a new piece collaged from an existing piece of music.

Professor Chew is a new awardee of a European Research Council Advanced Grant to carry forward this work in the project COSMOS: Computational Shaping and Modelling of Musical Structures, where she and her research team will be designing new analytical techniques to understand music structures as they are conveyed and experienced in performance, and also as they are found in unusual sources such as arrhythmic heartbeats.

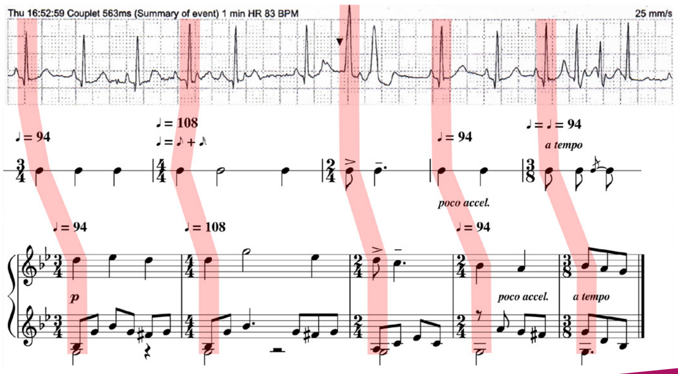

Figure 3. The musical rhythms of a short ECG recording of atrial fibrillation transcribed and embedded in a piece comprised of fragments from Bach's "Siciliano" collaged by Chew. See animation at vimeo.com/221351463; a description of one transcription process can be found in vimeo.com/226516952. Also for more examples of arrhythmia music, see bit.ly/ArrhythmiaMusic.

TECHNICAL SUPPLEMENT

One way in which performers add expression to their music is to delay playing a note for dramatic effect. But what is it that enables a musician to make extravagant gestures like taking a lot of time at a certain point in their performance and not at another? In order to tackle this problem, Canishk Naik for his summer internship at Professor Chew's lab asked people to listen to a piece of music and mark when they thought there was a significant transition, a tipping point, in the music. The researchers then analysed the music mathematically in order to quantify how the parameters of tension, loudness and tempo changed throughout the piece. They then used change-point analysis to look at the significant transition points and see how they were controlled by the underlying parameters of tension, loudness and timing. The aim was to find the parameter changes that lead to tipping points in the piece of music. The results showed that almost all popular tipping points corresponded to large timing deviations or to significant changes in a tension parameter.

This is an example of how mathematics is being used to analyse something that would previously have been seen as purely artistic, whereas musical expression can now be quantified and analysed.

References

Chew, E. (2018). Notating disfluencies and temporal deviations in music and arrhythmia. Music and Science, Vol. 1, published online 24 Sep 2018, pp. 1-22. url: http://journals.sagepub.com/doi/10.1177/2059204318795159

Chew, E. (2016). Playing with the edge: Tipping points and the role of tonality. In Stephen McAdams, David Temperley, Alexander Rozin (eds.): Milestones in Music Cognition Special Issue, Music Perception, 33(3):344-366. doi: 10.1525/MP.2016.33.03.344

Chew, E., A. R. J. François, J. Liu and A. Yang (2005). ESP: A Driving Interface for Musical Expression Synthesis. In Proceedings of the International Conference on New Interfaces for Musical Expression (NIME), pp. 224-227, Vancouver, B.C., Canada, May 26-28, 2005.

Chuan, C.-H. and E. Chew (2011). Generating and Evaluating Musical Harmonizations that Emulate Style. Computer Music Journal, 35(4): 65-82. doi: 10.1162/COMJ_a_00091

Déguernel, K., E. Vincent and G. Assayag (2018). Probabilistic Factor Oracles for Multidimensional Machine Improvisation. Computer Music Journal, 42(2).

Ghassemi, N. H., J. Hannink, C. F. Martindale, H. Gaßner, M. Müller, J. Klucken, and B. M. Eskofier (2018). Segmentation of Gait Sequences in Sensor-Based Movement Analysis: A Comparison of Methods in Parkinson's Disease. Sensors, 18(145), doi:10.3390/s18010145.

Herremans, D. and E. Chew (2017). MorpheuS: generating structured music with constrained patterns and tension. IEEE Transactions on Affective Computing. doi: 10.1109/TAFFC.2017.2737984

Herremans, D., and E. Chew (2016). Clouds and vectors in the spiral array as measures of tonal tension. In Proceedings of the 14th International Conference for Music Perception and Cognition (ICMPC), pp. 25, July 5-9, 2016, San Francisco, California, USA.

Killick, R., and I. A. Eckley (2014). changepoint: An R package for changepoint analysis. Journal of Statistical Software, 58(3): 1-19.

Kosta, K., O. Bandtlow, E. Chew (2015). A Change-point Approach Towards Representing Musical Dynamics. In Tom Collins, David Meredith, Anja Volk (eds.): Mathematics and Computation in Music—5th International Conference, MCM2015, London, UK, Jun 22-25, 2015, Proceedings. LNAI 9110, pp. 179-184, Switzerland: Springer. doi: 10.1007/978-3-319-20603-5; ebook isbn: 978-3-319-20603-5; softcover isbn: 978-3-319-20602-8

Naik, C., E. Chew (2017). Tipping points, pulse elasticity, and tonal tension: An empirical study on what generates tipping points. Late Breaking/Demo Session, International Conference on Music Information Retrieval (ISMIR), Suzhou China, Oct 23-28, 2017.

Pachet, F., A. Papadopoulos, P. Roy (2017). Sampling variations of sequences for structured music generation. In Proceedings of the International Conference on Music Information Retrieval, pp. 167-173.

Stowell, D., V. Morfi, L. F. Gill. (2016). Individual Identity in Songbirds: Signal Representations and Metric Learning for Locating the Information in Complex Corvid Calls. In Proceedings of Interspeech, pp. 2607-2611.

Sturm, B. L., O. Ben-Tal, Ú. Monaghan, N. Collins, D. Herremans, E. Chew, G. Hadjeres, E. Deruty & F. Pachet (2018). Machine learning research that matters for music creation: A case study. Journal of New Music Research, published online 3 Sep 2018, doi: 10.1080/09298215.2018.1515233

Sturm, B., J. P. Santos, O. Ben-Tal, I. Korshunova (2016). Music transcription modelling and composition using deep learning. ArXiv: 1604.08723.

Expert

Professor Elaine Chew

This article was written in 2018. Professor Chew is now a senior researcher at the French National Centre for Scientific Research (CNRS) - UMR9912 / Science and Technologies of Music and Sound (STMS) Laboratory located at the Institut de Recherche et Coordination Acoustique/Musique (IRCAM) in Paris, France, where she is Principal Investigator of the European Research Council (ERC) Advanced Grant project COSMOS: Computational Shaping and Modeling of Musical Structures (https://cosmos.cnrs.fr) and ERC Proof of Concept project HEART.FM: Maximizing the Therapeutic Potential of Music through Tailored Therapy with Physiological Feedback in Cardiovascular Disease (https://bit.ly/HeartFM-about), and Visiting Professor of Engineering in the Faculty of Natural and Mathematical Sciences at King's College London, UK.

The IMA would like to thank Professor Elaine Chew, for her help in the preparation of this document.