

Drawing: Research, Theory, Practice
Volume 5 Number 2

© 2020 Intellect Ltd Article. English language. https://doi.org/10.1386/drtp_00033_1

ANDREW STARKEY
University of Aberdeen

KATE STEENHAUER
Kate Steenhauer Limited

JACK CAVEN
Star Refrigeration

Painting Music: Using artificial intelligence to create music from live painted drawings

Keywords

artificial intelligence
live drawing
music improvisation
drawing performance
multidisciplinary
collaboration
figurative drawing

Abstract

This article describes the development of artificial intelligence (AI) techniques that monitor a painting or drawing evolving in real time and produce musical notes that relate to the individual elements of art as the artwork develops on the canvas. The article describes the practical approach required to capture the artwork unfolding in real time and then describes the framework used to develop the correlations between

visual art and music. The AI technique exploits these areas of similarity within the two distinct artforms in order to respond to the live-painted elements and produce musical notes that reflect the development of the evolving artwork. A prototype of this system was implemented in a live stage performance at Aberdeen May Festival 2019 whose narrative centred on the question Is AI good or bad? Other outputs of this project are a 20-minute film and a body of (tangible) visual artwork for digital platforms and gallery environments informed and inspired by AI. The integration of these disciplines through AI transforms a static artform into one that is dynamic, interactive, transformational, transient and temporal.

Introduction

The combination of visual art and music has an unequivocally compelling nature, with remarkable areas of similarity between the two artforms. Many artists (Kandinsky 1910; Vi et al. 2017) have explored multisensory ways of perceiving and, to a certain extent, have examined the analogy of mark-making to sound. This project, however, uses cutting-edge artificial intelligence (AI) techniques to examine the correlation and harmonious balance between calligraphic-style ink drawings and music, leading to the translation (conversion) of tangible marks into audible sounds (and in theory vice versa; however that is not described in this paper).

Calligraphic-style mark-making (using the fluid medium of Japanese ink) lends itself greatly for creating a musical dialogue on paper consisting of a dynamic sequence of movements unfolding in time and pulsating in rhythm while capturing the essence of a pose, gesture or any other subject matter.

AI has previously been applied to, for example, the problem of individual letter recognition (Graves and Schmidhuber 2009; Graves et al. 2009), and this project extends this in an intelligent manner to some of the fundamental parameters underlying visual art. However, the project also looks to explore and reflect on the impact of AI on society using the visual and aural outputs of the process. Previous studies (Fast and Horvitz 2016; Kelley et al. 2020) have shown how the perception of AI by the public has changed over time, and that although the public has been more optimistic than pessimistic when discussing AI, the fear of loss of control of AI has been increasing in recent years.

In practical terms, this means that the goal of the project is not to produce a process that always results in harmonious music and melodies, but that the AI can make mistakes. This question is interesting in a musical context since purely harmonious notes may be considered uninteresting by some people but beautiful by others. By the same token, 'mistakes' in the musical process could result in an interesting and novel musical form for some listeners but be jarring for others to listen to. The

exploration of this question is at the heart of this project and intends to tap into the public's unconscious fear – or excitement – of AI and the change that this will bring to all aspects of society.

Similarly, in terms of the visual component, whilst the purity of the drawing discipline is not compromised, conflict transpires through the subject of each drawing, but also through the application of dramaturgical drawing gestures that 'destroy' or transform each drawing through rolling, brushing paint, pouring water or scrunching up the drawing itself. The choice depends on what type of energy, drawing effect and ultimately what meaning in the scene is called for. The drawing can be partially or completely obliterated by this process, or more often is transformed into abstract marks that force the audience to consider the meaning of the drawing and also the emotional reaction produced by the development of the drawing.

Transforming a relatively conventional artwork that an audience has 'invested' in has a pronounced effect, often evoking a strong (and sometimes immediate) emotional reaction, for example, expressed as a gasp or shock running through the show. This also serves to reinforce the underlying philosophy of the drawing process combined with the musical output, that is, that the drawing is not a static artform but is dynamic, interactive, transient, transformational and temporal.

This project has been developed with a focus on live stage performance with the outputs from the artist unfolding on a projector screen in real time and the musical outputs from the AI process amplified through loudspeakers. This multidisciplinary collaboration celebrates the art of live drawing and highlights the dynamic capacity and interplay between drawing and music.

The application of AI in this novel manner reveals new insights into connections between seemingly unrelated compositional elements of these two art forms, exploring harmonic or dissonant comprehensions of the constituent forms and attributes of the pieces.

Drawing in a performative setting

Within the project, the artwork consists of painted drawings using calligraphic-style mark-making in Japanese ink, as can be seen in Figures 1–4, which show scenes from the performance at the 2019 May Festival at the University of Aberdeen. The system is initially trained and tested through digitized static photos of previously developed artworks, and accordingly, a system is developed that is able to interact with live-evolving (and thus wet) painted drawings (where the artist is the instigator with AI outputs to follow).

Whilst promoting drawing as a significant discipline in its own right, the oriental drawing materials and techniques are crucial to the goals and performative nature of this project, possessing the following characteristics:



Figure 1: Kate Steenhauer (May 2019). Top Dog (Japanese ink, 700 × 560 mm). © Kate Steenhauer.




Figure 2: Kate Steenhauer (May 2019). *The Physical Brain* (Japanese ink, 600 × 550 mm). © Kate Steenhauer.



Figure 3: Kate Steenhauer (May 2019). *Child's Play* (Japanese ink, 560 × 700 mm). © Kate Steenhauer.



Figure 4: Kate Steenhauer (May 2019). *Community* (Japanese ink, 1100 × 550 mm). Circles in top panels illustrate art elements in image analysis with (left) dominance of low frequency and (right) dominance of high frequency. © Kate Steenhauer.

- Minimalistic, 'clean', one-layer mark-making allows a good basis to break down and analyse a piece of art according to the fundamental elements and principles of art, and useful as a framework to correlate them to its musical counterparts as described in more detail below.
- These inks have a relatively high saturation/purity or strength. Each drawing is monochromatic in nature with a range of tonal values obtained by diluting the ink with water.
- The ink's fluidity and calligraphic tool enable dynamic, varied and descriptive mark-making, which therefore lends itself to a varied musical output.

Real-time data capture

Within the live performative setting, data is gathered in the following way:

- *Image capture*: The system requires a laptop to be connected to a camera system that is situated immediately above the canvas. The camera used in the project was a GoPro Hero 4 Silver. This is connected to the laptop via Bluetooth or WiFi and allows the laptop to take high-quality (static) photographs of the canvas (12 MP). This process is controlled by the operator of the laptop, who monitors the artist and takes photos (following a predetermined signal by the artist) after each element in the drawing is painted. This is to ensure that the hands of the artist (or hair or head) are not captured in the image.
- *Image analysis*: The laptop is running concurrently a computer programme that has been developed in Matlab (2019), a powerful software package commonly used within academia and the engineering field. This programme is continually running and monitors for new images to be saved into a specified folder used by the GoPro camera and will then subsequently begin the process of analysing the images. The first step of the entire process is to manually identify where the boundaries of the canvas lie, so that anything outside of this boundary is ignored in subsequent image analysis.

The programme then analyses the image following the removal of anything outside the boundary of the canvas and also compares the image against the previous image (if there is one) in order to clarify the details of the new element added to the canvas. The image is then reduced to the size of the new element rather than analysing the image in the context of the whole canvas; this allows individual elements to be more easily compared against each other. Image analysis of this new element is undertaken and includes the following image processing outputs:

- *Analysis of the amount of whitespace (i.e. no paint) in the element.* This is calculated by literally counting the number of pixels in the image that can be considered to be white and can be converted into a ratio so that this is independent of the size of the painted element.
- *Analysis of the amount of painted canvas in the element.* This is calculated by subtracting the number of white pixels in the image from the total number of pixels in the image and can be converted into a ratio so that this is independent of the size of the painted element.
- *Analysis of the variance of density of colour in the element.* This is accomplished by analysing the image once it is converted into a greyscale and binning the output so that the first bin represents the colour that is closest to white, with other bins representing increasing strength in the painted colour. As before, this can be normalized so that the bins represent a ratio of the whole image so that this is independent of the size of the painted element.
- *Two-dimensional (2D) frequency analysis of the element.* This is accomplished by undertaking 2D Fourier transform of the element (Frigo and Johnson 1998). The subsequent output is then binned into ten separate bins that reflect the frequency density of the painted element in ten equally spaced bins that range from low frequency to high frequency.

Musical notes are then subsequently produced based on this image analysis and also based on previous elements painted on the canvas (discussed further below).

This process continues until the operator indicates that the last element in the drawing has been done and the drawing is complete. This gives a final section of musical notes to be added and played by the computer, and then the music will stop.

Correlation between visual art and music

Basic (classical) definitions of the elements and principles of art are used to establish an initial framework that can subsequently be correlated to music theory. Examples of the relationships that are proposed by the authors between these worlds are shown in Table 1.

There are clearly many more attributes that can and should be used, and this will form the focus of further work. However, the attributes described in Table 1 allow some of the fundamental properties of the painted elements to be found and then matched against elements that are considered by the authors to be equivalent in the musical world.

The image analysis gives an output that is a ratio of the amount of white space (or non-painted area) for the new element. This is considered to be equivalent to the amount of silence in any new notes to be created. Similarly, for the ratio of painted pixels for the new element, this is considered to be equivalent to the velocity of musical notes to be created. The tonal value in the image, measured by the strength of colour, can be considered to be inversely related to the pitch range of the musical

Measured artwork attribute	Music counterpart
White (non-painted) space	Silence
Painted pixels	Velocity
Tonal value	Pitch range
Frequency	Number of notes and duration

Table 1: System attributes used in correlating art with music.

notes, so that strong tonal strength represents low pitches and light tonal strength represents high pitches. Finally, the frequency analysis of the element can be used to give a guidance to the system for the number of notes to produce when creating a musical motif and the duration of the notes. This means that higher frequencies in the painted element (i.e. a more intricate image) will result in more notes in a musical motif that will be played at a faster pace, and vice versa.

Clearly, some of these correlations are subjective, with different interpretations possible. However, even if different approaches were to be applied, the resultant musical output would be similar as long as the individual elements identified above were matched against each other in some manner.

The programme was then able to determine for any given painted element the attributes in music using the above mapping and could then generate a sequence of notes to be played in real time.

Development of AI solution

The requirements for the real-time solution are for a system that can interpret the visual (physical) components of an element within a work of art and to match those elements against their musical equivalents. This process is made up of several steps, namely image analysis of a captured photograph of canvas, understanding and categorization of image analysis, reinterpretation of this categorization into musical notes, production of musical notes in combination with other painted elements.

This leads to the development and implementation of a Self-Organizing Map (SOM; Kohonen 2001), a type of neural network and AI technique that is loosely based on the way that our own brain works and that is able to automatically categorize input data that can then subsequently be used to make decisions based on the categorization made.

Since the correlations between the visual art and music worlds has been defined as in Table 1, this means that the inputs into the SOM allow it to categorize the painted element that has just been

created and to then use that categorization to define musical outputs using the correlations defined above. For example, an element such as that encircled in the top left panel of Figure 4 is relatively simple with a relatively large amount of whitespace and a block-type structure. This implies that the musical element should also be simple, with few notes, with the possibility of silence and, due to the size and tone of the painted element, should be of low frequency. However, an element as shown in Figure 4 has a more complex structure, indicating more musical notes to make up a harmony. As the tonal value of this second piece is higher, this means that the notes produced will begin at a higher pitch.

The analysis of the SOM using the quantization error (Kohonen 2009) of the node that the element maps to (i.e. the decision made by the AI of which previously seen images the current image most closely resembles) allows a range to be defined for each of the musical outputs that are equivalent to the image analysis undertaken, since the image features that act as inputs into the AI process have direct musical equivalents as described and discussed above. Within these ranges, the AI can then choose values for each of the musical equivalences defined above so that the process can begin the construction of the musical notes.

In this manner, the SOM was initially trained by feeding it over more than 200 digitized static art images (captured by camera) previously created by Steenhauer. It should be noted that these artworks were different from those painted in the live performance, meaning that the SOM had not been trained on artworks or elements that were used in a real-time environment. Each artwork itself is broken up into many individual components as well as the whole composition generating an image database of over 1000 pieces.

The categorisations (or attributes) described in Table 1 were applied within these pieces of art using image analysis techniques, obtaining numerical data for each painted element that was used to train a SOM. The SOM was made up of 25 neurons in total in a 5×5 arrangement, trained and tested, simulating the task that will be carried out in a live production.

Construction of musical notes

Once the process has defined the musical parameters with which to construct the music, the process needs to use a different technique to produce the actual notes. This uses a Markov Chain (MC) (Gagnius 2017) model, which is a statistical approach for giving a probability of the next state space. This can be used in conjunction with historical pieces of music in order to learn the probabilities of different notes following each other, or for which notes to use when constructing chords and harmonies, and has, for example, previously been applied in this manner for the analysis of Beethoven's *Ode to Joy* (Linskens 2014).

This approach requires a number of pieces of music to exist (e.g. several songs by The Beatles), and these are then analysed in terms of the notes that make up the musical piece. The MC model can then analyse all the notes within the songs presented to give a probability of choosing a second note given a first note. The first note is determined from the musical outputs of the SOM process as defined above, and then a random number generator is used to select the second note according to the MC model. This process can be continued for as many notes as are required for this image and the required number of musical notes to be created.

The MC model approach can also be used to define which notes are used for harmonics or for chords. The MC model will learn the probability of notes being used as a chord or harmonic from the existing songs analysed, which can then be used in a similar way to give a choice of note to be used in real time by the process.

It should be noted that, due to the random nature of this process, the same notes are not chosen each time. In addition, the probability of notes being chosen can also be altered by choosing different songs or compositional styles in the training of the MC model. For example, in the 2019 Aberdeen May Festival live performance, separate MC models were trained on collections of songs by The Beatles, by Phillip Glass, and by Beethoven for different scenes within the show. Each scene within the performance, therefore, had music created in real time by an MC model using songs from different artists.

This step was undertaken to short-cut the musical knowledge that would ordinarily be required to create musical notes in real time that would have a harmonious sound. However, this is a limitation of the current approach, and further work is required to develop an AI alternative that will have a rudimentary understanding of musical rules to decide what notes to produce. This would also enhance the musical outputs of the project and increase the merit of the work to the audience.

The process then will choose whether to create a new melody or to create harmonies or chords, depending on whether painted elements have already been analysed and, if so, then on a random choice so that the process does not choose the same each time. Again, in the future, this will be improved to include an analysis of the evolving composition of the artwork in order to make this decision.

The system thus developed is therefore capable of exploiting the areas of similarity defined in Table 1 in order to produce novel musical notes that reflect the painted elements in real time.

A Musical Instrument Digital Interface (MIDI) synthesizing toolbox is then used to produce the final aural output.

However, as discussed earlier, the goal of the project is not always to produce perfectly harmonious music; and the random nature of the choice of notes in the creation of a musical motif, the length of notes within this motif and the notes chosen when playing chords or harmonies will sometimes result in music that is not pleasant to listen to. This is entirely by design and is to force the

1.
2.
3.
4.
5.
6.
7.
8.
9.
10.
11.
12.
13.
14.
15.
16.
17.
18.
19.
20.
21.
22.
23.
24.
25.
26.
27.
28.
29.
30.
31.
32.
33.
34.
35.
36.

listener to consider the question of whether the AI is doing a 'good' or 'bad' job in creating the music. One listener may react differently to another since this is an entirely subjective question and will also depend on the listener's view on AI itself – that is, it will bring to the surface the individual's excitement or fear over the increasing usage of AI in all parts of our daily lives.

In addition, the development of this system gives the ability to the artist to create a multifaceted and multidisciplinary artistic expression through both visual and aural dimensions, over which the artist has full control, as well as the ability to explore the relationship between the drawings, the musical output of the system and the audience themselves. This means that mark-making has a direct implication not only for the composition of the drawing but also on the development of the musical piece, and in turn the audience. For example, composing the same drawing but starting and ending with different elements will result in a very different musical output, depending on the characteristics of these elements. Therefore, although drawing is often considered a static art form, the musical output that is based on the mark-making transforms this into a temporal and transient output that is dynamic and depends on the order and choice of mark-making by the artist.

The artist can control the musical output by the choice of compositional elements – for example, the scene shown in Figure 3 comprises a number of figures that force the AI to create a number of musical motifs and chords that are similar in their tonal range, whereas the scene shown in Figure 2 has very distinct compositional elements that will result in a much different musical output.

The nature of the drawing process means that the mark-making will always be different even if the same drawing is undertaken. This is also the case for the musical output due to the built-in random process for the development of the notes and chords. This lends the entire system to a performance with no two performances being alike and results in additional dimensions to the generated artistic piece.

Conclusion

Project *Painting Music* uses AI to create music from live-painted drawings. The system employs a SOM to analyse individual elements of an artwork as they are being drawn and, based on this, produces music in real time and unique for each production. The system shows that it is possible to create a truly multidisciplinary artform where tangible visual marks in one domain are used to drive the novel creation of audible sounds/music in another.

Painting Music is a live stage performance, but other outputs of the project include a 20-minute film and a body of visual artwork. This work is a prototype that will be built upon by incorporating a greater number of art elements, as well as the basic principles of art, and subsequently enhancing the correlation to the fundamental parameters of music theory. In addition, it will look to enhance

the process of musical improvisation using theory and practice from this field and building AI approaches that will remove the need for heuristics in the music generation process.

The system developed thus far allows the artist to directly control the musical output of the system by the choice of compositional elements and the sequence in which they are drawn and thus in turn their relationship with the audience. This gives an additional dimension to the artistic capabilities of the piece and moves visual art as a traditionally static experience into a dynamic interactive artform that is temporal and transient in nature. This is emphasized in the production by the transformation of the artwork once it is complete in a traditional sense, as illustrated in the figures in this article, giving an additional dimension to the artistic capabilities of the overall production.

Further work is required to enhance the integration of both art and music theory into the process and will look to involve experts from both fields to develop the AI engine that drives the improvisation of the music that is created in real time. This is also important in terms of the question of novelty and innovation in the artistic piece. Currently the focus on novelty is with the artist's process and product; however, the project would greatly benefit from the integration of musical theory into the AI process. The integration would also create the possibility of undertaking the reverse of the system described in this article – to create novel painted drawings in real time from the sounds/music as they are being played.

Painting Music offers a unique experience and taps into the current demands for interactive and evolving 'stories' that contain elements of risk and reveal the artistic process to the audience. The production process brings attention to the art of live drawing and highlights drawing's dynamic capacity with other artforms, technology and their relationship with its audience. Intrigue within each production is created, hooking the audience through our own questions centred around the narrative *Is AI good or bad?* and exploring thoughts and fears over the application and impact of AI and its prevalence both today and in the future.

Online links to outputs of *Painting Music*

1-minute highlights live stage performance

<https://www.youtube.com/watch?v=vNCp6wBmf7U>

1-minute highlights of film

<https://www.youtube.com/watch?v=rLptcJZauSA&t=>

20-minute film

<https://www.youtube.com/watch?v=VvmyCOPMvoQ&feature=youtu.be>

Project website and visual artwork

<https://katesteenhauer.com/painting-music>

References

- Fast, E. and Horvitz, E. (2017), 'Long-term trends in the public perception of artificial intelligence', AAAI, pp. 963–69.
- Frigo, M. and Johnson, S. G. (1998), 'FFTW: An adaptive software architecture for the FFT', in Proceedings of the International Conference on Acoustics, Speech, and Signal Processing, 15 May, Piscataway, NJ: IEEE Signal Processing Society, pp. 1381–84.
- Gagniuc, P. A. (2017), *Markov Chains: From Theory to Implementation and Experimentation*, Hoboken, NJ: John Wiley & Sons.
- Graves, A., Liwicki, M., Fernandez, S., Bertolami, R., Bunke, H. and Schmidhuber, J. (2009), 'A novel connectionist system for improved unconstrained handwriting recognition', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31:5, pp. 855–68.
- Graves, A. and Schmidhuber, J. (2009), 'Offline handwriting recognition with multidimensional recurrent neural networks', in Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams and A. Culotta (eds), *Advances in Neural Information Processing Systems 22*, Vancouver, BC: Neural Information Processing Systems Foundation, pp. 545–52.
- Kandinsky, W. (1910), *Concerning the Spiritual in Art*, New York: Dover Fine Art.
- Kelley, P. G., Yang, Y., Heldreth, C., Moessner, C., Sedley, A., Kramm, A., Newman, D. and Woodruff, A. (2020), '"Happy and assured that life will be easy 10 years from now": Perceptions of artificial intelligence in 8 countries', arXiv:2001.00081, <https://arxiv.org/pdf/2001.00081.pdf>. Accessed 16 October 2020.
- Kohonen, T., Nieminen, I. T. and Honkela, T. (2009), 'On the quantization error in SOM vs. VQ: A critical and systematic study', in Proceedings of the 7th International Workshop on Advances in Self-Organizing Maps, 8–10 June, Heidelberg: Springer-Verlag, pp. 133–44.
- Kohonen, T., Schroeder, M. R. and Huang, T. S. (2001), *Self-Organizing Maps*, 3rd ed., Heidelberg: Springer-Verlag.
- Linskens, E. (2014), *Music Improvisation Using Markov Chains*, Maastricht: Maastricht University.
- Matlab (2019), *Matlab Release 2019*, Natick, MA: The MathWorks, Inc.
- Vi, C. T., Ablart, D., Gatti, E., Velasco, C. and Obrist, M. (2017), 'Not just seeing, but also feeling art: Mid-air haptic experiences integrated in a multisensory art exhibition', *International Journal of Human-Computer Studies*, 108, December, pp. 1–14.

Suggested citation

Starkey, Andrew, Steenhauer, Kate and Caven, Jack (2020) 'Painting Music: Using artificial intelligence to create music from live painted drawings', *Drawing: Research, Theory, Practice*, 5:2, pp. 209–224, doi: https://doi.org/10.1386/drtip_00033_1

CONTRIBUTOR DETAILS

Andrew Starkey is a senior lecturer in the School of Engineering at the University of Aberdeen. He has over twenty years of experience in artificial intelligence (AI), and his research focus is on automated data analysis using AI techniques, explainable AI, virtual reality and autonomous learning for robots.

Contact: School of Engineering, University of Aberdeen, Aberdeen AB24 3UE, UK.
E-mail: a.starkey@abdn.ac.uk

 <https://orcid.org/0000-0003-1797-8737>

Kate Steenhauer is a visual artist whose practice explores the dynamic and interactive capacity of drawing in dialogue with other artforms, technology and their relationship with its audience. This moves visual art as a traditionally static experience into a dynamic interactive artform that is temporal, transient and transformational. Multidisciplinary collaborations include dance, opera, music, sound, verbatim and artificial intelligence. These productions can be viewed through the mediums of theatre, film, (live) audio-visual installation and artwork. She also has a large portfolio capturing contemporary life at some of Scotland's most iconic industries using printmaking and oil painting techniques.

Contact: Kate Steenhauer Limited, 12 Donbank Terrace, Aberdeen AB24 2SJ, UK.
E-mail: katesteenhauer@hotmail.com
Web address: <http://katesteenhauer.com>

Jack Caven is a software engineer and aspiring entrepreneur. He has recently graduated from the University of Aberdeen, where he completed a Masters in Mechanical Engineering. For his final-year honours thesis, Jack undertook the challenge of applying artificial intelligence to convert the process of a painting being constructed into music. This project has since turned into his first entrepreneurial endeavour: *Painting Music*.

Contact: Star Refrigeration.
E-mail: jack_caven@btinternet.com

Andrew Starkey, Kate Steenhauer and Jack Caven have asserted their right under the Copyright, Designs and Patents Act, 1988, to be identified as the authors of this work in the format that was submitted to Intellect Ltd.