Chris Hales

Artificial Intelligence: The Latent Revolution in Filmmaking

ADAMarts
Volume 2, 2020 • 21
Audiovisual Media Arts

Received: 01.03.2021
Accepted: 01.05.2021

Abstract

New practices in experimental film are emerging as a result of artificial intelligence (AI) developments in machine learning and generative adversarial networks (GANs). The author describes this as a Latent Revolution in filmmaking, complementing Lev Manovich's so-called Velvet Revolution in motion graphics that occurred in the mid-1990s. A review of current artistic and filmmaking practice using AI leads to the identification of the role of AI and GANs as creative partners, contributing their own ideas and imagery to complement those of the filmmaker. Experimental films produced recently using AI-based techniques are analysed; some of these could be said to be an extension of traditionally recognised practices, whereas others demonstrate new processes and result in novel genres of film that have little or no historical precedent. The primacy of latent space is highlighted as being a major characteristic of GANs that can lead to new kinds of creative imagery and in particular the 'latent spacewalk film'.

Keywords

experimental film, experimental animation, abstraction, neural networks, GANs, artificial intelligence, latent spacewalk

...

Introduction

Reflecting upon a decade of moving image developments, Lev Manovich (2007) considered motion graphics to be a totally novel type of filmmaking brought about in the mid-1990s by what he termed the 'Velvet Revolution' (paraphrasing the peaceful democratic changes in Czechoslovakia) of Adobe After Effects, which was the first popular software designed to create and combine animation, compositing, and special effects on a personal computer. Both the process and the types of video that emerged represented, according to Manovich, "a new hybrid visual language of moving images in general ... which is also common to a large proportion of short 'experimental' (i.e. non-commercial) films being produced for media festivals, the web, mobile media devices, and other distribution platforms" (Manovich, 2007, pp. 69–70).

Moving to the present day, this paper seeks to draw attention to what seems to be a new revolution that has emerged rapidly in the creative community yet for the time being has remained relatively unnoticed: the use of artificial intelligence (AI) by artists, creatives and filmmakers. The emphasis here is on experimental filmmaking and the article will be supported with first-hand experience gained by the author in teaching AI-filmmaking workshops at RISEBA and other universities. Since the creative adoption of AI is happening so recently and so rapidly, all that is possible here is to take a snapshot based on the moment in time at which this article is written—there is nonetheless a sense that the text will seem dated soon after it is published.

As a counterpoint to Manovich, the term proposed here is the 'Latent Revolution' brought about by AI. The term 'latent' reveals that the revolution is still in progress (at least at the time of writing—Manovich was writing retrospectively) and is occurring almost unnoticed, whilst at the same

Figure 1. Photographs of the author processed with DeepDream (left) and style transfer (right) in the style of Van Gogh's *Starry Night* (Riga: Chris Hales, 2021)

time it references the so-called 'latent space' which is a unique characteristic of neural networks. Latency in the context of digital and analogue media has also been addressed by Sean Cubitt (2014) albeit without reference to AI. Although the ensuing discussion is focused quite specifically on experimental filmmaking, it must be borne in mind that using AI for artistic purposes represents only a fraction of its application across all aspects of society.

There are specific reasons why AI has rapidly developed and expanded into creative fields such as experimental filmmaking, and why at the same time it seems to have caught us unawares. In 1989 film visionary Gene Youngblood acknowledged the growing importance of computer code to the filmmaker (Youngblood, 1989) and "intelligent behavior" was suggested in the 2003 exhibition catalogue of the monumental *Future Cinema* exhibition (Weibel, 2003, p. 597), yet at the same time Hamlyn and Smith's recent book (2018) on new perspectives and practices in experimental animation makes no reference whatsoever to artificial intelligence. As recently as 2018 the use of AI for creative purposes was in the hands of a privileged few who had access to substantial computing power and possessed high-level knowledge of complex computer programming techniques. Although the creative act was at that time out of the hands of the majority, the phenomenon entered mainstream culture due to *DeepDream* (acidic-coloured hallucinogenic imagery often including nightmarish animal faces, as in Figure 1), style transfer (applying the style of one image to a different image, also shown in Figure 1) and 'deep fakes' (the transfer of facial motion onto a different face, as illustrated in Figure 2). Within a short period of time these techniques (and others) were being applied to generate attention-grabbing imagery on a variety of easy-to-use websites and mobile apps, usually for entertainment purposes. The *MyHeritage* website, for example, which had already introduced an AI-based colourisation tool called *DeOldify* in 2020, offered a *Deep Nostalgia* service in February 2021 to animate the faces of photographs of ancestors using the same process as 'deep fake'—it is, basically, a tool that creates short animated films based on faces. *Deep Nostalgia* became the latest viral sensation (as had happened earlier with the mobile phone app *FaceApp*) and its popularity further demonstrates how AI-based image manipulation is becoming increasingly mainstream and its current capabilities are being widely understood at the end-user level. The real

Figure 2. Stills from a short video sequence in which the impassive face of the Mona Lisa takes on facial expressions of the boy featured in Herz Frank's 1978 film *Ten Minutes Older* (Riga: Chris Hales, 2021)

revolution that is occurring, however, lies at the top end of the creative chain. Artists and creators are no longer disenfranchised and are now able to experiment with the *process* that creates the end product. There has been a sudden proliferation of easy-to-use software, combined with access to the necessary computing power, which enable the various capabilities of neural networks to be experimented with by non-programmers and creatives for their personal projects. How the field of experimental film is affected by this is the focus of the discussion that follows, subsequent to an introductory overview of AI-based terminology, techniques, processes and tools in the context of creative practice.

AI for Creatives and Filmmakers

Within the broad field of AI it would be more accurate to use the term 'machine learning' (ML) to describe the process by which a set of related data (images of faces for example) is trained into a neural network 'model' which is employed by software in a variety of ways. As a result of the algorithm and set of data on which it was trained, a model can be fed new data and will recognise patterns in it and can therefore make certain predictions and create its own patterns. Concerning machine learning itself, developments in generative adversarial networks (GANs) which pit a generator against a discriminator have proved to be easily applicable to artistic practice. Scientific papers demonstrating new varieties and modifications of GAN neural networks are being published almost daily, and it can be argued that the GAN is to the Latent Revolution what Adobe After Effects was to Manovich's Velvet Revolution.

Since the necessary datasets and models can be substantial, considerable computing power, usually beyond the capability of a desktop PC or laptop, becomes a necessity. Although plenty of open-source software code has been available on programming platforms such as *GitHub*, installing all the requisite and disparate software elements to successfully run platforms such as *Tensorflow* on a personal computer has always been a difficult and non-trivial task which would invariably require knowledge of a software language such as *Python*. This has proved an insurmountable obstacle for the majority of visual artists, and even if it was overcome the necessary computing power and available RAM memory would be limited to whatever the local CPU could muster: AI usually requires the computational power of graphics processing units (GPUs) which are many times more powerful than a typical central processing unit, the computer's CPU.

What has facilitated the recent uptake of machine learning techniques by creatives is the availability of software programmes that solve both these difficulties: firstly by offering cloud-based GPU time and secondly by easy-to-use software with

Figure 3. New styles of mittens created out of the latent space of a generative model trained by the author on 500 images of Latvian and Lithuanian mitten designs (Riga: Chris Hales, 2021)

understandable graphical user interfaces. The possibilities evolve almost daily but at the moment of writing *Runway ML* is such a software, originated by a graduate of NYU Tisch School of the Arts and specifically developed for ease of use by creatives and artists. Google's *Colaboratory* (*Colab*) *notebooks* offer an alternative, browser-based solution, which functions as a virtual machine: neither installation nor coding is required, yet the code is still revealed (and therefore demystified) in a series of annotated 'cells' that a user can activate in sequence. Many newly developed GANs and techniques published in scientific papers have their code made available in the form of a *Colab notebook*. Amongst other alternatives, Apple has introduced a software solution for its own platforms entitled *Create ML*.

Since it is more or less *de rigueur* for contemporary filmmakers to understand digital video editing software, the use of an additional software is by no means an insurmountable task. To make things easier, a great many datasets and pre-trained machine learning models have been made publicly available, meaning that creative outputs do not necessarily presuppose that filmmakers must harvest a dataset and carry out their own model training processes. This could be thought of as the equivalent of using the off-the-shelf filters/effects that come bundled with a video editing or post-production software: the starting point of any project in *Runway ML* is to select a model from a large variety offered by the software, some of which have been created by members of *Runway's* community. Undoubtedly, personalised models trained on a filmmaker's own dataset offer the best route towards uniqueness and originality but this requires a much greater commitment and the process is extremely time intensive and can require days of continual access to the computational power of a GPU: it is possible to do this with *Runway ML* but to do so would incur non-trivial costs. Similarly *Runway ML* permits the chaining of the output of one model to the input of another, which opens up fascinating possibilities for the creation of unexpected and unusual imagery—although to do this comes, again, at additional

usage costs.

Within the realm of ready-made models, a variety of different functions are commonly available which are trained with and acting upon data that might be in the form of text, visual imagery, and audio, and many of these have applications for filmmakers. Three common functions are generators, identifiers and stylisers. The former are GAN-based neural networks that can create new and original instances from the trained model content, with or without a prompt from the user, and their outputs can be readily seen in current websites such as those entitled This Person (or horse, or cat, etc.) Does Not Exist. GANs are equally effective using text-based models to produce sentences, stories and scripts. A recent (2020) GAN-based architecture of note is NVIDIA's *StyleGAN2*, which was developed on high-quality images and is able to generate impressively photorealistic face images down to the detail in hair and skin. Identifiers (or classifiers) could almost be thought of as having the inverse function and are most commonly employed to recognise objects or describe the content of a still image or video frame, or to find and identify faces and facial emotions. Stylisation, in particular transformation between images of different styles, is arguably the most well-known creative application of machine learning and, as alluded to earlier, is readily available through a variety of websites and apps. In the basic technique specific models are pre-trained to specialise in a certain artistic style, for example Claude Monet or Vincent Van Gogh, and the user need only supply an image indicative of content (image-specific features) in order to create a result in which the chosen style is applied to the content. A multi-style model enables the output to be produced as a weighted mixture of styles, whereas 'arbitrary' style models permit the user to supply not only the content image but also the requisite style. Stylisers are frequently applied to faces—the abovementioned *FaceApp*, for example, includes use of the style-based *CycleGAN*—and audio style transfer is also possible, although for the time being it is much less common.

Some ML techniques simply perform the function of video post-production tools and although of interest to filmmakers, they tend to replicate already existing non-AI functionality but do it differently and potentially more successfully. An example would be *ESRGAN*, which can upscale image/video resolution (so-called 'super-resolution'). AI-based processes can also be used to improve image quality, apply colourisation, and perform auto-cropping and green-screening. Additional AI tools exist to predict depth in an image, to animate the depth within a 2D image, and to perform motion capture. All of these tools might have an application in any kind of film, including those intended to be experimental artworks.

The key characteristic of generative neural networks is the concept of latent space—hence the term *Latent Revolution* suggested here. This will be discussed in more detail later, but a latent image can be considered a representation that is not mathematically modelled (as might be the case in 3D computer graphics) but is a vector-defined method by which the image can be formed from the manipulation of aspects of the dataset on which the GAN was trained. A neural network typically has many layers, each of which can make its own particular contribution to a generated image: novel images are therefore created from these generative models that are related to, but different from, those in the original dataset. Figure 3 shows outputs from a model trained by the author. Given sufficient computing power, it is possible to 'find' the equivalent of a user-supplied input image within the latent space manifold in a process called projection, which opens up possibilities to subsequently modify it in various ways—for example, aging a face or changing its expression is possible once the latent face representation has been found.

There is an additional characteristic of note, which is that many of the GAN-based visual outputs are generated in a square format, primarily because the model was itself trained on square format images. Whereas square negatives have always been popular in photography, the dimensions are more problematic with video and not commonly recognised as a valid format. This is a practical consideration that filmmakers, at least for the time being, need to take into account. Another practical issue is that not all models are specifically made to act upon video files—the majority of visual models function on a single image at a time, meaning that sequences of images need to be processed one by one and the outputs combined into a single video file when processing is complete.

AI as Creative Partner

Manovich not only examines his Velvet Revolution from the point of view of its audiovisual outputs, but also by analysing the new working practices it engenders:

"In the case of After Effects, the working method that it puts forward is neither animation, nor graphic design, nor cinematography, even though it draws from all these fields. It is a new way to make moving image media. Similarly, the visual language of media produced with this and similar software is also different from the languages of moving images that existed previously."

Lev Manovich (2007, p. 74)

Artificial intelligence also has the potential to fundamentally change the manner in which experimental filmmakers formulate and create their works—by acting as a creative partner. Although mainstream films are essentially collaborations between large teams of specialists, the experimental filmmaker has traditionally been a loner with a personal vision to express, concerned with all aspects of the filmmaking process from cinematography to acting, only occasionally bringing in a specialist for aspects such as the soundtrack (Malcolm le Grice, for example, recruited Brian Eno for the soundtrack of his 1970 film *Berlin Horse*). Formal partnerships or teamwork in experimental film are rare (an example being the Surrealist filmmakers the Themersons) and even informal collaborations are unusual—*Meshes of the Afternoon* (1943) is recognised as Maya Deren's creation even though it had cinematography and an acting role by her then-husband Sasha Hammid. Access to funding, equipment and facilities, and even screening opportunities have historically been difficult for the experimental filmmaker, although these particular aspects have become more democratised in the digital era.

What is revolutionary about the use of AI is that it is not merely a new type of tool but is a genuine creative partner that contributes ideas and content to the production of an artistic artefact, whether that be text, image or film. Artist and programmer Casey Reas explains that GANs "assist with creating unexpected images, unlike any that have been created before. They can be unlike photographs and paintings – they are truly something new" (Reas, 2019). Arthur Miller's recent book *The Artist in the Machine* (2019) analyses the nature of creativity and consciousness and looks at each step in the recent development of creative AI and the evolution of GANs. Miller investigates the authorship and/or co-authorship of creative works produced using AI, which is timely given that in 2018 the auction house Christie's in New York sold *Portrait of Edmond de Belamy*, created from a GAN, for $432,500. This somewhat controversial sale led to discussions about the nature of copyright ownership and whether the work was produced by a human artist, a machine, or the programmers of the GAN that generated the work. Miller concludes his book by suggesting that in the future machines will be fully creative and that computers will not only be capable of

Figure 4. Ten outputs from a model trained by the author on 400 images of the old-fashioned Supilinn district of wooden houses in Tartu, Estonia. The original training images were mostly sourced online (Riga: Chris Hales, 2021)

entertaining each other but will entertain humans to such an extent that some people might even come to prefer computer-generated creativity.

Whereas Miller tends to differentiate between visuals (artwork), text and music, filmmaking is to a certain extent a synthesis of all three. From my own anecdotal experience with teaching experimental film workshops based on AI techniques, a fruitful starting point for ideas is to enlist a text-based AI such as the website *Talk to Transformer* to generate meaningful sentences, prose and poetic fragments that might inspire or form part of the experimental film. This might be just a catchy title, some poetic phrases, multiple lines of dialogue, or even an entire screenplay: the experimental science fiction short film *Sunspring* (2016) was produced in a 48-hour challenge from a script generated by a GAN trained on science fiction novels and has been described as "hilarious and intense" (Newitz, 2016). The soundtrack of an experimental film is an additional component that can be wholly or partly created by AI, for example using audio style transfer, which dispenses with the skills of an audio specialist whilst producing unconventional results. Using a style transfer process exemplifies a human-machine partnership because it would be the filmmaker who chooses and supplies the audio samples and specifies how much the proportion of style and content is biased one way or the other. Regarding the predominantly visual aspects of filmmaking, numerous techniques and processes will be discussed below, but what is abundantly clear is that 'AI as creative partner' undoubtedly represents a novel method of making moving image media—in the same way that Adobe After Effects brought about new ways of working in Manovich's Velvet Revolution of the mid-1990s.

AI-based Experimental Films in Relation to Earlier Practices

Although the AI experimental film certainly entails a new approach with novel types of outcome, its multiple techniques need to be examined in detail and in the context of previous classifications of experimental film to determine whether these AI-made films represent something new or are actually an extension of traditional practices. In classifying experimental films Bordwell

and Thompson (2013), known mostly for their interest in traditional cinema, propose just the two categories of abstract form and associational form (which they exemplify by 'poetic film'). Peter Weibel (2002), a media artist/curator more in tune with contemporary media art, considered experimental approaches based on materiality (emphasising the celluloid or videotape itself, or the technology of its representation), multiple narratives, multiple screens, time and space, sound, expanded cinema, and found footage (reframing or reinterpreting the original footage). To these, the structural film should be added: the term is wide-encompassing (Gidal, 1979) but would include the analytical calculation and logic of a film's montage according to a set of rules, often mapped out in structural diagrams and 'scores'. Experimental animation could also be ventured as a separate category and indeed it is a form of practice that seems relevant to some of the AI-based films to be examined later.

The abstract form, certainly, can be considered as one of the direct links between new AI practices and previous movements in experimental film. GANs in general seem to follow an abstraction tendency when interpolation is made through their latent spaces even when using quite representational model imagery, and this is particularly apparent when a GAN has been trained on too wide a variety of data. The webpage *Generative Engine*, which responds to typed characters and words by generating images (using the text-to-image generator *AttnGAN*), attempts to visually represent the typed words but the results are rarely if ever representational and resemble a fusion of abstraction and Post-Impressionism—it could even be said that the *Generative Engine* has its own distinctive non-representational artistic style. AI imagery produced as abstraction seems successful because it is somewhat forgiving: exact representation is less important than the style. Additionally, in terms of the image generation process it is of note that much abstract film was historically made, like AI-generated imagery, without the use of a camera.

Considering digital ones and zeros as the essential raw material of films produced by AI, materiality—as a traditional experimental filmmaking categorisation—can also be directly mapped to machine learning outputs, particularly in the case of the well-known *DeepDream*-generated imagery. The originator of the process, Alex Mordvintsev, refers to its aesthetic qualities as being "similar to hallucinogenic experiences" (Mordvintsev, 2021). Infinite zoom films can be produced with *DeepDream* by moving further and further into the noise of the network without losing visual output quality. This method intentionally reveals the inner workings of the GAN's calculations, allowing users to examine specific layers and stages of the neural network process. *DeepDream* might be considered a continuation of 'glitch' in the sense that the imperfections in the machine are elevated to artistic status—certainly it is a process that focuses on the inner workings of the digital 'brain' and could be classified as an approach that foregrounds the digital materiality within a neural network.

The found footage film tradition is also highly relevant where AI is concerned, most evidently in the sense that GANs can be trained from datasets fed with images which may be user-supplied, taken from public-domain image datasets, or scraped from online sources such as *Google Images* or *Flickr*. Outputs generated by the model would inevitably reflect and reinterpret the originals, as shown in Figure 4, although in terms of copyright it is not possible to digitally deconstruct a trained model to reveal the original images on which it was trained. Found footage may play a different role when a video file is specified as an input to an existing model which responds by generating new or modified output. A simple example would be using the *Deoldify* model to artificially colour vintage black-and-white films. Anna Ridler's film *The House of Usher* (Ridler, 2017) illustrates a deeper and

more complex approach. Ridler takes as input the 1928 silent film *Fall of the House of Usher* and by an iterative process, samples of the original frames are manually redrawn in pen and ink, output is generated using the *Pix2Pix* model, and the process is repeated to distress and transform the original. The underlying process is reminiscent of Malcolm Le Grice's 16mm colour film *Berlin Horse* (1970) referred to earlier, which reinterprets original sequences (which include early newsreel footage) through refilming, multiple superimpositions and colour transformations.

Structuralism also has its relevance to the AI film. More often than not, however, it seems to be manifested as a stylistic secondary characteristic that is not necessarily inherent to the GAN-based process. For example, Derrick Schultz describes his original film *You Are Here* (2020) as structuralist seemingly because there is a continual transformation (interpolation) of forms obeying a set rhythm and timing. This structured steady pace is due to authorial choice rather than a pattern suggested by the AI which has functioned to provide the visual content. In a similar fashion, when imagery is generated in real time by *AttnGAN* in response to a user typing words as input, the responses (i.e. the resultant film sequences) follow the pace of the user's typing of letters and words rather than creating any montage patterns of their own. Undoubtedly trained models will become generally available that can produce their own temporal structures and a pointer in this direction might lie with a film by Damien Henry based on footage filmed looking out of moving train windows. Henry explains the function of his model thus:

"First, it learns how to predict the next frame of the videos, by analyzing examples. Then it produces a frame from a first picture, then another frame from the one just generated, etc. The output becomes the input of the next calculation step. So, excepting the first one that I chose, all the other frames were generated by the algorithm. … The results are low resolution, blurry, and not realistic most of the time. But it resonates with the feeling I have when I travel in a train. It means that the algorithm learned the patterns needed to create this feeling. Unlike classical computer generated content, these patterns are not chosen or written by a software engineer."

Damien Henry (2017)

Even though the process simply generates one frame after another without the complexity of montage patterns, there is a structural methodology exemplified by the procedure of feeding output frames back into the generator. A work such as this could convincingly be categorised as falling within the field of structural filmmaking: it also typifies the human and AI working together in a mutually creative partnership.

New Practices and Processes Using AI

Damien Henry's process gives a hint of a new type of film and the primacy of *the model* again becomes apparent, in particular the originality that can emerge from training one's own neural network on a personal dataset (such as scans of a sketchbook, a painting portfolio, selfies, a stamp collection, etc.) For the time being, however, this is still a step too far for many creatives who prefer off-the-shelf functionality. As suggested earlier, one of the most popular visual applications of AI in recent times has been the use of a style transfer technique to generate new imagery in the style of others. In this case the AI acts as a 'pastiche machine'. Although a new phenomenon, minimal creative effort is required to produce pastiched film sequences using this method and the results have limited artistic value, but there is potential for greater sophistication. Derrick Schultz has used *BigBiGAN* (which generates similar images based on an input image) to create films in the style of Stan Brakhage, Maya Deren, Mary Ellen Bute, Len Lye and Norman McLaren. The work of

Stan Brakhage, one of the most well-known of all experimental filmmakers, has also been the subject of two related pieces by Casey Reas entitled *Earthly Delight 1.1* and *Earthly Delight 2.1* (2019) in which new filmic imagery was generated from Reas' own ML models trained on scans of collaged vegetation. Here, Reas transcends style transfer by producing a contemporary technological interpretation of the core idea of Brakhage's film *The Garden of Earthly Delights* (1981). Reas investigates the filmmaking process rather than the resultant content and his generated images show similarity to the frames of the original but have a different visual quality which goes well beyond the trivial operation of a 'pastiche machine'.

One noticeable theme emerging from the current use of AI is what could be termed the 'face film'. This might be allied to the popular culture of taking and sharing 'selfie' photographs and might even be classed as a post-internet practice, but there is little or no precedent of a significant corpus of experimental films that explore the human face. The 'face film' has been popularised by the fun and familiarity of manipulating one's own facial representation or those of friends or celebrities in an uncomplicated manner using a variety of recent popular phone apps and webpages. In addition to the entertainment value, faces have gained their prominence due to the fact that they were amongst the first huge image datasets made available for models to be trained from, in addition to the fact that there are clear real-world applications. Undoubtedly this accounts for the square image format on which so many models have been trained, as a face fits conveniently into a square frame. Facial models and their neural networks are highly developed and progress has been rapid in face detection and alignment using a 68-facial-landmark model, leading to facial emotion recognition, face recognition, changing characteristics of a face such as expression or age, and animating a source face according to the motion of a driving video. NVIDIA's improved *StyleGAN2* face models, released in 2020 at 1024-pixel resolution, seem to have further encouraged the 'face film'. A key ingredient is manipulation of the latent space of the neural network to mould latent faces. Mario Klingemann has explored the potential of *StyleGAN2* latent faces in numerous works, some of which run as real-time generative installations such as *Memories of Passers-by* (2018) and *StyleGAN2 - mapping music to facial expressions in real time* (2020). His *Neural Face* and *Alternative Face* (both 2017) are based around a model trained on seven videos of the face of French singer Françoise Hardy. The novelty will probably soon wear off with regard to these face films, except perhaps in regard of the use of a so-called 'first order

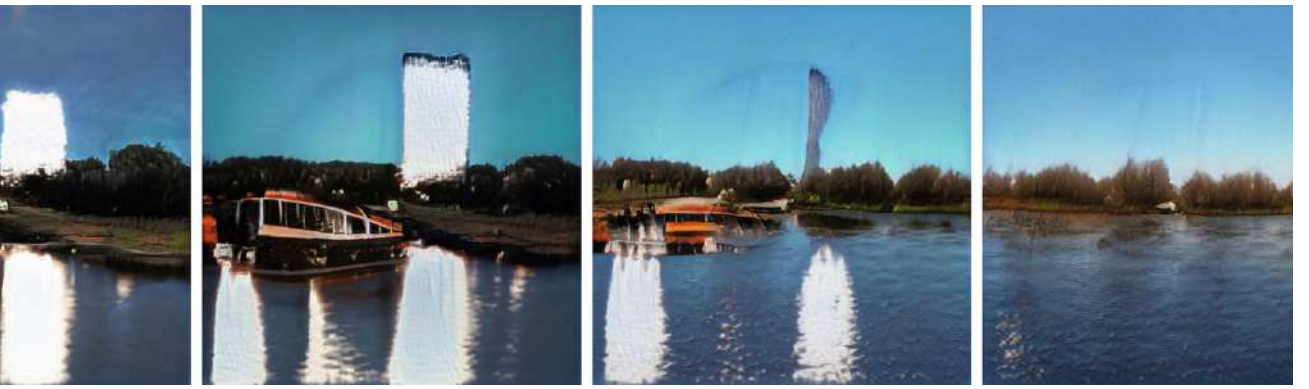

Figure 5. Frames extracted from a latent spacewalk sequence created by the author from his GAN model trained on images of the River Emajogi in Estonia (Riga: Chris Hale, 2021)

motion model' to combine a source face image with motion patterns from the video of a different face. These types of video, as explained earlier, entered mainstream culture around 2018 under the term 'deep fakes', and the moral implications they raise—in contrast to the trivial effort required to produce them—means they will undoubtedly remain prominent in the near future.

A second theme newly emerging from AI arises with the foregrounding of text. In regard of historical experimental films, those in which text plays a significant role represent a very minor aspect of the field of practice, although the identification of a genre of 'text film' has been suggested by MacDonald (1995) and Knowles (2015). Significant examples in which text takes prominence include Peter Rose's *Secondary Currents* (1982) and some of the early films of Peter Greenaway. Basing a film around AI-generated text has the potential to become a significant new methodology: the text could represent the first stage of the process from which the audiovisual aspects would provide an artistic response or, vice versa, the text could emerge as an accompanying response to visual input. In my own teaching experience, AI-generated text has been used surprisingly frequently by students as an approach to creating an experimental film, not necessarily as a means to create a lengthy script such as *Sunspring* but by using the AI as a creative partner that suggests resonant poetic phrases (or even just a title) which become an inspiration for imaginative visual responses from the filmmaker. These films might be considered hybrids that combine both new and old filmmaking processes and serve as a reminder that the AI need not necessarily dominate the final whole in order for it still to have played an integral part in its creation.

Above all else, the most paradigm-shifting aspect behind the architecture of a GAN and its potential use for creative purposes lies with the latent data manifold or latent space. Its novelty and significance have led to the term *Latent Revolution* being used in this article to describe the effect of AI on experimental filmmaking. As tentatively explained earlier, latent space is a somewhat conceptual and abstract data representation into which a GAN generator can map points using complex vector arithmetic. This offers much variety in the style and characteristics of the generated image and permits transitions (interpolations) by creating a linear path through the latent dimension between

the points that generated two images and then creating an animated sequence from all the images produced by the points along the path. This was demonstrated in the influential scientific paper by Radford et al. (2016) which discusses the concept of walking in the landscape of the latent space: the term 'latent spacewalk films' can be used to describe films produced by interpolation within the latent space of a neural network. The importance of the vector (which is typically 100-dimensional) to the generative process seems to connect directly to Sean Cubitt's assertion, made before the recent growth in creative AI, that the "vector aesthetic is our orientation to the future that already reads the present as sheer motion, sheer change, the fearful allure of the unknowable that cannot be eradicated by any machinery of governance" (Cubitt, 2014, p. 263). Figure 5 illustrates a traversal of latent space in a landscape model trained by the author: each individual frame of the interpolation is newly created by the GAN, which can autonomously create its own random sequences (noise loops) or can have sequences of keyframes specified by human selection.

The latent spacewalk film is essentially a completely new type of experimental animation and although traditional animators would recognise interpolation as being a synonym for the process of inbetweening between key frames, each spacewalk frame is a new creation in its own right and not a superficial compromise between the starting and ending key frames. If an image-based GAN has undergone significant training (this is currently measured in days or even weeks) with a homogeneous dataset, each generated frame can have a photorealistic quality, yet at the same time a model may be deliberately undertrained or mistrained to produce imagery with an unpredictable painterly quality. Artists have been quick to explore the creative potential this offers, for example Jake Elwes' 45-minute video *Latent Space* (2017) is a dreamlike exploration of the latent space of a model trained on the 14.2 million images of the *imageNET* dataset. Many 'face films' use the interpolation process and these could be considered a distinct subset in their own right—*StyleGAN2's* photorealistic latent face interpolations are a compelling example. An encouragement to the creative exploration of latent space lies with the fact that the spacewalking process itself has become simple to accomplish by non-programmers: *Runway ML* software, for example, can be used to configure and export these sequences, reducing the creation of a latent spacewalk film to a simple exercise that code-allergic art students have no difficulties in accomplishing, and it comes as no surprise that such films prove popular amongst students of my experimental AI-filmmaking workshops. Enhancements to the creative process of latent spacewalking continue apace through research such as *GANSpace* (2020), which aims to identify which layers and vectors in a GAN create specific change in the synthesised image, offering greater authorial control over characteristics such as viewpoint, aging, lighting, and time of day. Network blending, network bending, and GAN rewriting are amongst other current areas of interest that attempt to push the boundaries of GAN research.

Conclusion

Lev Manovich's *Velvet Revolution* was founded on newly developed moving image techniques of layering, transparency and compositing, and "software such as After Effects became a Petri dish where computer animation, live cinematography, graphic design, 2D animation, and typography started to interact, creating new hybrids" (Manovich, 2007, p. 70). These techniques became firmly established and emerging now is a Latent Revolution in which neural networks are teaming up with filmmakers and artists as their creative partners and GANs have become a new Petri dish where latent spaces are being explored to form new kinds of creative outcomes. Artificial intelligence is undoubtedly making a contribution to new

modes of experimental filmmaking, aided by recent software that is easy to use and offers access to powerful GPUs. Techniques such as the creation of 'deep fake' video sequences are already available to all and are used for malign purposes just as much as for creative ones, and ethical questions are being raised across the entire spectrum of AI which will require continued attention as new uses and misuses of the technology arise.

Although structuralist and materialist influences are still apparent in certain AI-based films, as are connections to abstraction and found footage, new filmic characteristics have emerged and others will undoubtedly continue to develop—and like any mode of art practice, artists will constantly find ways to use techniques in unexpected and unusual ways. For the time being, creating an animation purely from the latent space still offers fascination; however, it will soon be considered one tool in a wider armoury. Self-trained models will become a commonplace medium of unique personal expression, and the technique of chaining models will offer greater originality and unexpected outcomes as machines pass ideas to each other. Applying multiple approaches and a variety of models within a single film also has much potential: Kira Bursky's experimental film *Lessons From My Nightmares* (2020) is an interesting example that uses numerous GAN models—several trained by Bursky on personal imagery such as selfies—to create latent spacewalks, to assist animated walk cycles, to create backgrounds and to aid 3D depth estimation. Manipulation in After Effects software makes a significant contribution to the final form of Bursky's film, a reminder perhaps that there should not be too much emphasis on the final film being entirely a product of AI. Hybrids composed of both traditional and AI processes will become the norm and in these cases the fruits of the Velvet and Latent Revolutions would meld together in a harmonious whole.

Modes of practice are also changed by the advent of AI. Future experimental filmmakers may need to exchange their traditional cinematographic skills for the curation and creation of datasets and the training of personalised models: ironically, the deployment of these models might in fact result in the creation of 'self-pastiche' films. Debates about the stage at which the creativity contributed by the machine enables it to be called an artist in its own right will continue for a while yet, but a stage of creative partnership has been reached already. With a few exceptions, experimental filmmakers have traditionally been solitary, working alone on all aspects of their productions. Now these filmmakers can escape their lonely existence and work in tandem with an ever-willing creative partner—in the form of their preferred GAN or neural network that will contribute its own ideas and imagery to the co-production of the whole.

...

References

Bordwell, D., Thompson, K., 2013. *Film Art: An Introduction. 10th Edition*. New York: McGraw-Hill.

Bursky, K., 2020. *Lessons From My Nightmares*. Available at: https://www.youtube.com/watch?v=xu4m9_3HcHY [Accessed 20 February 2021].

Cubitt, S., 2014. *The Practice of Light: A Genealogy of Visual Technologies from Prints to Pixels*. Cambridge, MA: MIT Press.

Elwes, J., 2017. *Latent Space* [video]. Available at: https://www.jakeelwes.com/project-latentSpace.html [Accessed 20 February 2021].

Generative Engine [website]. Available at: https://experiments.runwayml.com/generative_engine [Accessed 20 February 2021].

Gidal, P., 1976. *Structural Film Anthology*. London: BFI.

Hamlyn, N., Smith, V., 2018. *Experimental and Expanded Animation. New Perspectives and Practices*. Cham: Palgrave Macmillan.

Henry, D., 2017. *Video generated by an algorithm in one shot. No edit or post-processing*. Available at: https://www.youtube.com/watch?v=lr59AhOPgWQ [Accessed 20 February 2021].

Härkönen, E., Hertzmann, A., Lehtinen, J., and Paris, S., 2020. GANSpace: Discovering Interpretable GAN Controls. *ArXiv*, abs/2004.02546.

Klingemann, M., 2020. *StyleGAN2 - mapping music to facial expressions in real time*. Available at: https://www.youtube.com/watch?v=A6bo_mIOt00 [Accessed 20 February 2021].

Knowles, K., 2015. Performing Language, Animating Poetry: Kinetic Text in Experimental Cinema. *Journal of Film and Video*, Vol. 67, No. 1, pp. 46–59.

MacDonald, S., 1995. *Screen Writings: Scripts and Texts by Independent Filmmakers*. Berkeley: University of California Press.

Manovich, L., 2007. After Effects, or Velvet Revolution. *Artifact*, Vol. 1, No. 2, pp. 67–75.

Miller, A. I., 2019. *The Artist in the Machine. The World of AI-Powered Creativity*. Cambridge, MA: MIT Press.

Mordvintsev, A., 2021. Interviewed for *Click: The Virtual Race* (*Zwift*), S13 E02, 18 January. BBC World Television.

MyHeritage [website]. Available at: https://www.myheritage.com [Accessed 20 February 2021].

Newitz, A., 2016. *Movie written by algorithm turns out to be hilarious and intense*. Available at: https://arstechnica.com/gaming/2016/06/an-ai-wrote-this-movie-and-its-strangely-moving [Accessed 20 February 2021].

Radford, A., Metz, L., and Chintala, S., 2016. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *CoRR*, abs/1511.06434.

Reas, C., 2019. Interviewed by Caroline Menezes for *Studio International*, 21 May. Available at: https://www.studiointernational.com/index.php/casey-reas-interview-computer-art-coding [Accessed 20 February 2021].

Ridler, A., 2017. *The House of Usher*. Available at: https://vimeo.com/217670143 [Accessed 20 February 2021].

Schultz, D., *Artificial Images* [website]. Available at: https://artificial-images.com [Accessed 20 February 2021].

Talk To Transformer [website]. Available at: https://app.inferkit.com/demo [Accessed 20 February 2021].

This Person Does Not Exist [website]. Available at: https://thispersondoesnotexist.com/ [Accessed 20 February 2021].

Weibel, P., 2002. Narrative Theory: Multiple Projection and Multiple Narration (Past and Future). In: Rieser, M., Zapp, A. (Eds.) *New Screen Media: Cinema/Art/Narrative*. London: BFI, pp. 42–53.

Weibel, P., 2003. The Intelligent Image: Neurocinema or Quantum Cinema. In: Shaw, J., Weibel, P. (Eds.) *Future Cinema, the Cinematic Imaginary after Film*. Cambridge/London: MIT Press, pp. 594–601.

Youngblood, G., 1989. Cinema and the Code. *Leonardo*, Supplemental Issue 2, pp. 27–30.