How AI is Enabling a Creativity Renaissance

Ben Falchuk
Systems, Cybersecurity and Machine Learning Dept.
Peraton Labs
Basking Ridge, NJ USA
e-mail: bfalchuk@peratonlabs.com

***Abstract*—Artificial Intelligence (AI) has infiltrated many aspects of our lives, in both recognizable and invisible ways. Deep learning and sophisticated new information technologies allow the deployment of AI at massive scales, and media giants like Facebook and Google have whole-heartedly adopted AI to exploit their massive data sets in the pursuit of their economic goals. Enterprise and consumer tools and apps are embracing these techniques too. But has this proliferation given users a breadth of creative aptitudes akin to – say - those of Leonardo da Vinci? In this article, we present a survey of impressive AI tools for creativity that provide users with profound new creative powers. We illustrate the sweeping breadth of potential for a human-AI creativity renaissance.**

Keywords-Artificial Intelligence (AI); User Interface (UI); creativity; Human-Machine Interface (HMI); web application.

I. INTRODUCTION

Artificial General Intelligence (AGI) refers to machine intelligence that achieves human-like or better cognitive abilities, and features planning, learning and reasoning. While AGI is considered a sort of moonshot, we already find ourselves amidst extremely impressive Artificial Narrow Intelligences (ANI) whose capabilities are superior to human capabilities in pre-defined realms, such as chess and mathematics. Surprisingly, many AI scientists believe that there is a greater than 50% chance that AGI will be achieved in the next 45 years or less [1]. The term artificial superintelligence refers to AI that is autonomous, self-improving, and vastly superior to humans in most every endeavor. Whether or not AGI will comprise a threat to humanity or exhibit consciousness are hotly debated topics among philosophers and computer scientists [2].

Deep learning is a computational AI technique in which so-called Artificial Neural Networks (ANN) – many layers deep – form a basis for predicting results from inputs [3][4]. Deep ANN's are essentially processing engines based loosely upon the mammalian cerebral cortex (synapses, potential functions, etc.) in which an input layer feeds into intermediate layers and finally into an output layer. Each layer has many "neurons" at which mathematical adjustment continually takes place during training. A principle goal is to train the ANN sufficiently so that it can make useful predictions about data it has never seen before. In so-called supervised learning, an ANN learns to produce ideal outputs from labeled data (e.g., artworks, voices, etc.). In unsupervised learning, the data is unlabeled and the goal is to uncover latent patterns and clusters. Classification and prediction are critically important in realms such as meteorology, navigation, and medicine, and it is clear that even narrow AI technologies are changing our lives profoundly. The impact of narrow AI is seen in the data sciences, social sciences, and engineering, and is an underpinning of social network feeds, self-driving vehicle navigation systems, and Natural Language Processing (NLP) to name just a few. The application of AI in traditional creative fields, such as painting and writing, is relatively nascent, but growing rapidly in both theory and practice [5]. A technique called Generative Adversarial Networks (GAN) is often used to generate things that "look authentic". It does so by pitting two ANN's against each other in a sort of zero-sum game of generation and recognition [3].

In this article, we do not attempt to answer the question, "what is art?" nor distinguish between amateur and professional human artists. Instead, we focus on how the creative class (as defined in [6]) can now collaborate with AI's in new ways. The works included herein are organized by category and are either recent advances, or representative of the field; we have favored accessible systems that readers can try themselves. Older technologies are mentioned when they serve to illustrate the field's vector of progress. The remainder of this article is structured as follows. Section II provides our survey. Acknowledgements and conclusions close the article.

II. AI CREATIVITY TECHNOLOGIES

We will never know what da Vinci – the quintessential *Renaissance man* - might have created had he been armed with 21st century tools, such as Photoshop, Internet, and AutoCAD. This section surveys how applications of AI and human-AI collaborations can help to turn merely creative people into latent da Vinci's. The quest to understand approaches for instilling creativity in computer AI is not new (e.g., [63][64]); this article assembles recent, noteworthy and accessible systems across several domains of creativity.

A. Painting, Drawing and Sculpture

Artistic painting and drawing techniques are thousands of years old but today we find deep learning and AI technologies are breaking new ground. In 2018, an artwork called La Famille de Belamy sold for over $400,000 in auction [7]. It was created by a GAN trained on 15,000 painted portraits spanning several periods (see Figure 1, left). The dramatic sale marked an early and significant moment, and illustrated the commercial potential of AI art. More

 ISBN: 978-1-61208-870-9

recently, an AI artwork by a robot called Sofia sold for $688,888 [62]. Meanwhile, Google DeepDream is a means for visualizing the internal state of a convolutional neural network (CNN) as it learns to classify images [8]. DeepDream is used by artists as a means to create surreal animated journeys "through the layers of a neural network" by recursively processing images, zooming in, and creating new variants (see Figure 1, right). Each variant attempts to re-create particular types of elements, such as trees or dogs (as directed by the artist), giving the resulting animations a meandering creative feel.

Figure 1. La Famille de Belamy (left); DeepDream image (right) [7,8]

DeepArt.io is an algorithmic artistic style transfer algorithm that outputs new works of art in the style of other artworks (on which it has been trained), and gained notoriety for its ability to rapidly transfer the style of the painter Van Gogh to arbitrary images [9]. DeepArt.io and other style transfer techniques are leveraged by artists in unique and beautiful ways (e.g., dinosaurs made of flowers [10]) and are embedded in scores of popular iOS and Android mobile apps, such as the popular PicsArt editor.

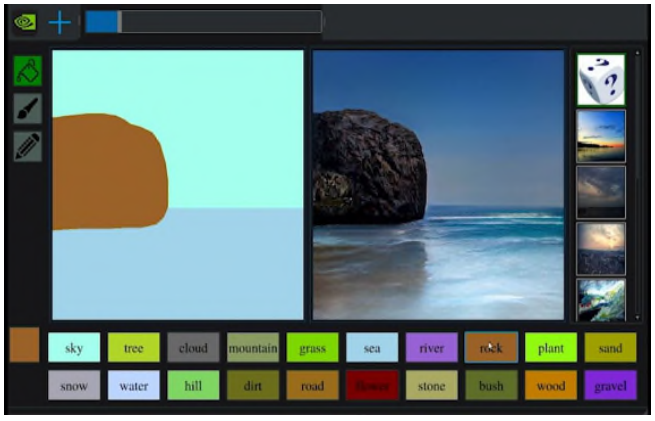

Figure 2. GauGAN: human user's broad strokes (left), corresponding AI-generated scene (right) [9]

Research in GAN's and techniques for converting segmentation maps into lifelike imagery has resulted in a tool named GauGAN [11]. Human artists working with this tool need merely to sketch a broadly-stroked image and ask GauGAN to fill in the details, texture, reflection, and colors, which it does by referencing its vast training set of images. In such a human-AI collaboration, then, the human steers the broad strokes while the AI creates details (see Figure 2).

In a similar vein, the Art42 webapp leverages GAN's trained on Cubist period art to provide a continuous stream of convincing new AI-generated artworks on each access [12]. Each output is a sort of Cubist daydream, vaguely familiar yet swirly and free in a way that belies its AI origins. Elsewhere, AI's are infiltrating the sculptural arts and seem poised to jump into the 3rd dimension. For example, *RobotSculptor* explores how humans can teach robotic arms and hands to sculpt clay in an optimal repeatable fashion by finding the robotic motions to optimally satisfy human objectives for the surface [13]. AI is also dramatically improving aspects of 3d printing through fault diagnosis and property predictions [14].

B. Writing, Poetry and Illustrated stories

The expression of human thought as written stories is transcendent, and yet even here, AI is encroaching. The OpenAI Generative Pre-trained Transformer (GPT) technologies are language prediction models capable of producing human-like text, with massive billion-parameter neural networks [15]. Write with Transformer is an online demo illustrating the breadth of GPT-2 by autocompleting lines of human-written text with phrasings that mostly make sense [16]. This is possible, in part, because it has sampled and learned from 8 million Internet web pages. Figure 3 shows exemplary results in which human inputted text is plain and GPT-2 contributions are highlighted in color.

The flowers in the garden are very special. Their scent is so different from anything I have ever smelled before. But when I approached the
- plant, I could sense a familiar fragrance.
- garden a couple of weeks ago and placed a single rose..
- flowers with the glass lid in I was surprised to find out..

Figure 3. Write with Transformer sample [15].

Human writers can use GPT-like AI to help them produce new kinds of creative output – never before possible - involving textual phrases. In one light-hearted experiment an ANN generated new Candy Heart messages ("Love bun", "Call me", etc.) by leveraging its training on existing candies [17]. Remixing language on custom source materials and then applying language models can yield new results in the style of the source. An Amazon-funded company trained Alexa's NLP on the full Harry Potter corpus and then used the resulting ANN to generate an entirely new chapter. The text sounded quite a bit like J. K. Rowling's voicing but lacked a singular train-of-thought [18]. Fascinating advances in ANN encodings allow AI to generate coherent novel sentences that interpolate between two given sentences [19]. Such a capability could be imagined as part of a collaborative human-AI poem-writing process in which the human user bounds the prose and the AI fills in the middle parts. When a different ANN was trained on the corpus of Shakespeare sonnets and asked to generate new ones it

 ISBN: 978-1-61208-870-9

adequately captured the aspects of both rhyme and meter [20]. On the other hand, expert human evaluators were able to distinguish the sonnets as fakes largely due to degraded readability and emotional pull. A system called Hafez combines finite state machinery with deep learning, yielding an ability to create a rhyming poem about an arbitrary word, such as *tree* [21]. A short story called *The Day a Computer Writes a Novel* – the result of a human-AI collaboration - made it through the first round of a Japanese literary contest in the mid 2010's. In other work, when researchers got both GPT-3 and humans to write college essays and had the results graded by experts, GPT-3 got passing marks, and even more praise than human essays in some respects. On the other hand, the AI failed to create strong narratives "incorporating the 5 senses", and was at times vague or awkward [22].

The AI that underpins the Verse by Verse application was trained on full-text poetry of more than 12 classical era poets and is very effective at generating a line of poetry to follow any line written by a human (given a preference for a particular poet). In this way, human and AI may forge a new poem, line by line, that has the voice of both the human and a historical poet, such as Dickenson or Whitman [23]. A Google Arts and Culture experiment called Poem Portraits fuses visual art with poetry creation [24]. It not only creates a long-running poem in 19th century style around user-inputted words, but it creates a visual portrait of each contributor's face. Its ANN is trained on over 25 million words of poetry using a long short-term memory recurrent neural network. A related project named Please Feed the Lions employed AI to create a novel poem based on viewer inputs, and projected the poem onto a huge lion sculpture in London's Trafalgar Square. These are good ways to introduce people to the potential of AI poetry.

Figure 4. Images created by DALL-E in response to the prompt, "an armchair in the shape of an avocado" [26]

OpenAI's DALL-E tool can generate images from almost any text descriptions by leveraging GPT-3 and Image GPT [25]. The resulting images are impressive. In particular, many results seem to exhibit a styling that is somehow clever, belying the computational nature of their inception. For example, the text, "an armchair in the shape of an avocado" produces numerous images of exactly that but the images seem more like the result of human brainstorming than those of a computer algorithm (see Figure 4). Indeed, almost anything that one types to DALL-E is similarly visualized. It is easy to imagine human-AI collaborative loops in which a human author's poem is illustrated in near real-time (as it is typed) by an AI partner. The images, in turn, spark new textual ideas for the human.

C. Culinary Arts

This section highlights advances in AI related to the culinary arts. Unscientific experiments with GPT-3 have shown that it can output textual recipes, but its lack of knowledge of procedure or chemistry make the results more silly than useful. For example, one might get a recipe for "watermelon cookies" that makes very little practical sense [27]. On the other hand, food scientists have trained an ANN on vineyard weather and irrigation parameters to predict wine aroma profiles, a practical tool that could give growers and winemakers powerful insights for business operations and new product design [28]. Meanwhile, researchers found over 300 unique ingredients by training an ANN on millions of recipes from Recipe1M. They then recovered a scored set of over 300,000 food pairings from these ingredients [29]. In principle, this kind of AI could be effective as an assistant to a human chef in a food preparation use case.

Figure 5. The AI-generated recipe "Caymanian Plantain Dessert", as cooked by Engadget's T.O'Brien [60]

Elsewhere, a different platform for creativity in support of culinary recipes employs machine learning, Bayesian probability, chemoinformatics, traditions, and crowdsourced seedling recipes, and leverages IBM Watson. The resulting recipes show apparent novelty and are rated as "very creative" by domain experts (e.g., "Caymanian Plantain Dessert", created with ~17 ingredients, is depicted in Figure 5) [30]. IBM published a full book of these AI recipes in 2015 and it has mostly 4 and 5 star reviews on Amazon. Google and Sony have also pointed their AI prowess at the culinary arts and, in 2021, helped human bakers innovate a new dessert type [61].

 ISBN: 978-1-61208-870-9

Alternatively, ANN's trained on recipes, food imagery, and topic networks can perform so-called image-to-recipe in which AI predicts the ingredients of photographed foods [31]. Such a technology could be the basis for an AI sous-chef that watches food preparation, infers ingredients, finds food pairings, and interacts with the chef using natural language.

D. Photography, Portraiture,and Beyond

This section describes the promise of AI to not only enhance photography and portraits but also enable new forms of visual portraiture, not even remotely possible in da Vinci's time. Super Resolution (SR) is a pragmatic technique (offered by Adobe and others) in which ANN's trained in narrow image categories (e.g., cars, animals) can upscale small images to much larger resolution ones by "hallucinating" - and filling in - convincing details [32][33]. To some extent, SR frees the human artist from practical scale constraints.

GANPaint employs ANN GAN's to improve the manner and consistency in which synthetic elements can be added to photographs [34][35]. Artists can now augment imagery with convincing new semantic elements, such as adding a convincing new window into a photo of a window-less kitchen. Figure 6 illustrates a GANPaint use case in which the human artist broadly strokes a tree into the image in order to replace the tower; the AI then performs a seamless replacement. In principle, this technique also works with facial features and portraits.

Figure 6. GANPaint – original urban photo with CN Tower (left); image after "painting a tree" atop the tower (right) [35]

The StyleGAN (and StyleGAN2) generator is an efficient way to generate unique high quality facial images while controlling aspects of style, such as hair or facial features [36]. The best of these results have proven to be stunningly detailed and extremely convincing. The webapp This Person Does Not Exist displays a stream of StyleGAN-generated faces, each evoking an eerie humanity [37]. Exemplary images from this new kind of portraiture are illustrated in Figure 7.

Artbreeder - a GAN that "breeds" new images from chosen ones – enables a new form of human-AI portraiture process [38]. For example, a human artist who has created a portrait image can breed it towards the features of a mountain lion. To do this, she chooses mountain lion images as parent images for the portrait and asks artbreeder to evolve her original portrait [39]. Artbreeder (see also *picbreeder*) encourages a human-AI collaboration during which the artist and an AI hone new images in a series of interactive steps. The results are often visually stunning AI-generated portraits [40].

Figure 7. Two non-existent people from This Person does not Exist [37]

So-called style-transfer GAN's can leverage large homogeneous training sets to amazing effect. In 2019 a webapp called Toonify used GAN's and trained on characters from animated films [41]. It could then use this network to apply Disney-esque facial features to real photographs, essentially "finding and assembling" features in cartoon images to correspond to features of the photograph. The resulting images (Figure 8) were largely perceived as cheerful and surreal and were so popular and interesting that the developer's demo website was quickly out of bandwidth.

The notion of portraiture is changing to include dynamic visual effects. Artists can now animate the mouths and faces of static portraits to match a target audio (e.g., using temporal annotated GAN) to very dramatic – and potentially politically dangerous - effect [42][43]. These animations are sometimes called deepfakes (i.e., they use "deep" learning).

Figure 8. Exemplary Toonify output (right) from photo input (left) [36].

In a related vein, a brilliant AI-driven augmented reality prototype finds a body in any photo or artwork and then animates it to walk right off the page before your eyes [44] (Figure 9, shown from third person). This new form of art allows creators to think in multiple modalities beyond 2d (and would almost certainly be considered sorcery by da Vinci).

 ISBN: 978-1-61208-870-9

Figure 9. AI turns a 2d image into an animated 3d scene on an augmented reality platform [44].

E. *Music*

Music is a vital part of all cultures. Da Vinci himself invented various musical technologies and was a good musician. Today, deep learning is encroaching on - and contributing to - musical culture. OpenAI has demonstrated a Jukebox capability comprised of an ANN trained on 1.2 million songs, lyrics, and genres. The ANN can predict compressed audio tokens [45] and JukeBox has created convincing new songs with good coherence, although lacking a certain degree of musical structure. Meanwhile, an app named DeepSinger – trained on 92 hours of singing data – can synthesize singing voices in multiple languages using a feed-forward transformer [46].

The classic rock singer Freddie Mercury is immortalized in the AI webapp called FreddieMeter in which an AI judges precisely how much a viewer sounds like Freddie when she is prompted to sing his songs. The convolutional encoder underpinnings of FreddieMeter represent a novel approach to voice processing [47].

DrumBot demonstrates the nascent promise of AI-based musical accompaniment; in this webapp prototype, an AI derives a drum track to accompany any simple melodic sequence made by a human artist [48]. Figure 10 illustrates the simple user interface of DrumBot.

The webcam webapp called Air Guitar creates guitar music driven by the position of the musician's hands in the webcam [49]. In fact, there is a plentiful variety of similar such demonstrators spanning many body parts, gestures, and musical instruments.

Deep Music Visualizer employs a GAN to visualize music with imagery from trained datasets; the ever-changing AI-generated imagery is responsive to parameters of the music, such as tempo and timbre [50]. In this way, the AI is contributing a lucid visual accompaniment that is sensitive to the context of the music. Lucid Sonic Dreams is a similar music visualizer [51].

By leveraging the Google Magenta technology (in which a deep learning system has made advances in the length of musical compositions that it can generate [65]) a team has trained an AI to compose new songs in the style of existing rock artists such as the Beatles. A recent initiative resulted in a convincing new Nirvana song whose words and music were AI-generated [66].

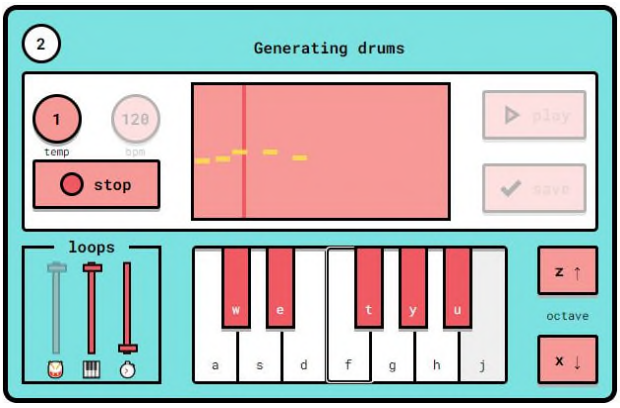

Figure 10. DrumBot web application creates a drum track to the melody track inputted by the user (shown in yellow dashes) [48].

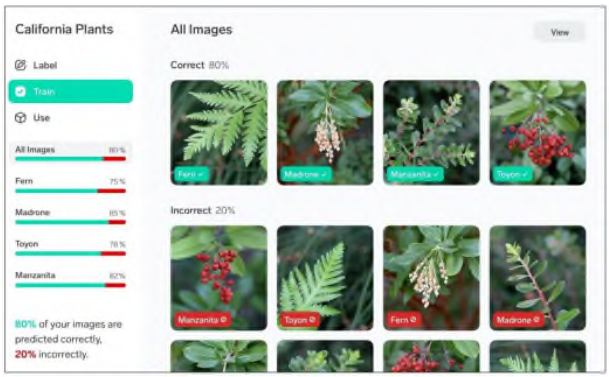

Figure 11. A user interface of Lobe, image classification use case, uses familiar widgets like gauges to convey feature detection [56].

F. *Frameworks*

Where Da Vinci would have employed easels and art supply bags, today's artists are now beginning to employ AI-based frameworks. These software tools serve as scaffolding for creative computation and abstract away unnecessary technical concepts. Their emergence is a sign that human-AI creativity is gaining importance, and that a purely coding-centric view of AI is yielding to a higher-level view more amenable to creative non-coders. Both SageMaker and AutoML are frameworks allowing users to train ANN's, use existing models (e.g., vision, NLP), and deploy them, using a simple UI [52][53].

MediaPipe promises cross-platform tools for creative visual apps, such as those leveraging pose, face detection, and object tracking [54], and deployments that work on web and mobile platforms alike. Figure 12 illustrates MediaPipe body and face tracking. Runway is a framework tailored to creative pros and their workflows, such as animation and green screen video-editing [55]. Lobe is a code-free AI model maker by Microsoft. It currently supports only image classification but the Lobe UI (see Figure 11) is notable for how well it abstracts away complexity and makes ANN training into something simple, elegant and fun [56].

 ISBN: 978-1-61208-870-9

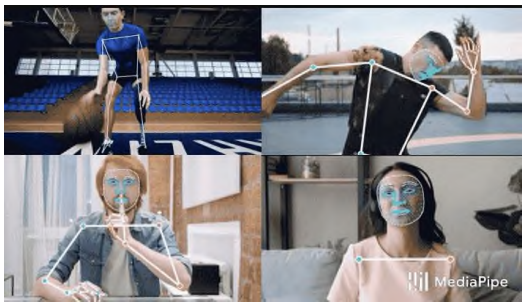

Figure 12. MediaPipe body and face tracking (body pose estimators shown as white skeletal lines, facial features as blue) [54]

TensorFlow is a machine learning underpinning for an ever-changing landscape of creative projects. Its capabilities include pose, hand, and body tracking, and deep learning, and all of these can run directly in most desktop web browsers [57][58][59].

III. CONCLUSION

The barriers to human-AI collaborations are rapidly lowering thanks in part to ever-advancing UI designs, hardware and GPU's, and Internet and web accessibility and protocols. This article has highlighted only a small subset of tools that illustrate the vast potential of human-AI creativity. Our future work may include further detailed analysis of particular creative domains and the AI techniques that support them, particularly on web platforms. Even in their present (sometimes) nascent states, these technologies have already enabled new forms of art - ranging from painting to poetry, cooking, and music - that both inspire and boggle the mind. We observe the following trends:

- Increasing potential for human-AI collaborative flows in which each actor contributes unique value
- Improved accessibility to human-AI tools thanks user interface and web technologies
- Single-domain tools currently dominate the scene, while tools with domain-crossing skills are more rare

On the other hand, daunting research challenges remain, such as training set bias, fairness, accessibility, and sustainability. Furthermore, it is unclear how the notions of consciousness, ethics, language, empathy, personality, and experience relate to a machine's ability to perform as a creative partner. The nascent field of human-AI creativity will continue to be driven by the combined advances from the sciences, philosophy, and the digital arts. We are not all da Vinci's just yet but we may be inching closer. Yet even more likely, AI will move us in altogether unpredictable creative directions.

ACKNOWLEDGMENT

We thank the reviewers for their valuable comments.

REFERENCES

[1] K. Grace, J. Salvatier, A. Dafoe, B. Zhang and O. Evans, "Viewpoint: When Will AI Exceed Human Performance? Evidence from AI Experts" , Journal of AI Research, Vol.62, 2018

[2] N. Bostrom, Superintelligence: Paths, Dangers, Strategies. Oxford University Press, Oxford, UK, 2014.

[3] A. Burkov, The Hundred Page Machine Learning Book, self-published, 2019.

[4] J. Schmidhuber, Deep Learning in Neural Networks: An Overview, TR IDSIA-03-14 / arXiv: 1404.7828 v4, 2014.

[5] A. Miller, The artist in the machine: the world of AI-powered creativity, MIT Press, Cambridge, 2019.

[6] R. Florida, Rise of the Creative Class, Basic Books, New York, 2012.

[7] BBC News, *Portrait by AI program sells for $432,000*, Oct. 2018. [Online]. Available from: https://www.bbc.com/news/technology-45980863 [retrieved: June 2021]

[8] A. Mordvintsev, C. Olah and M. Tyka, *DeepDream - a code example for visualizing Neural Networks*, 2015 [Online]. Available from: https://ai.googleblog.com/2015/07/ [retrieved: June 2021]

[9] L. Gatys, A. Ecker and M. Bethge, A neural algorithm of artistic style, arXiv:1508.06576v2, 2015.

[10] A. Liszewski, *A Neural Network Turned a Book of Flowers Into Shockingly Lovely Dinosaur Art*, June 2017. [Online]. Available from: https://gizmodo.com/a-neural-network-turned-a-book-of-flowers-into-shocking-1796221045 [retrieved: June 2021]

[11] NVIDIA Inc., *GauGAN Beta*. [Online]. Available from: http://nvidia-research-mingyuliu.com/gaugan [retrieved: June 2021]

[12] V. Vieriu, *Art42 - Infinite cubist art*. [Online]. Available from: https://art42.net/ [retrieved: June 2021]

[13] Z. Ma, S. Duenser, C. Schumacher, R. Rust, M. Bacher, F. Gramazio, M. Kohler and S. Coros, "RobotSculptor: Artist-directed robotic sculpting of clay", Symp. On Computational Fabrication, 2020.

[14] G. D.Goh, S. Sing and W. Yeong, "A reivew on machine learning in 3d printing", AI Review, Vol.54, 2021.

[15] M. Hutson, "Robo-writers: the rise and risks of language-generating AI", Nature, V.591. [Online]. Available from: https://nature.com/articla/d41586-021-00530-0 [retrieved: June 2021]

[16] Write with Transformer, https://transformer.huggingface.co .

[17] J. Shane, *Okay GPT-3: Candy hearts!*, Feb., 2021. [Online]. Available from: https://janellecshane.substack.com/ p/okay-gpt-3-candy-hearts [retrieved: June 2021]

[18] Botnik Inc., *Harry Potter*. [Online]. Available from: https://botnik.org/content/harry-potter.html [retrieved: June 2021]

[19] R. Bowman, L. Vilnis, O. Vinyals, A. Dai, R. Jozefowicz and S. Bengio, Generating Sentences from a Continuous Space, arXiv:1511.06349v4, 2016.

[20] J. Lau, T. Cohn, T. Baldwin, J. Brooke and A. Hammond, "Deep-speare: A joint neural model of poetic language, meter and rhyme", Proc. 56th Annual Meeting of the ACM Linguistics, Melbourne, 2018.

[21] M. Ghazvininejad, X. Shi, Y. Choi and K. Knight, "Generating Topical Poetry", Proc. Empirical Methods in NLP, Austin, 2017.

[22] EduRef, *What Grades can AI get in College*?, [Online]. Available from: https://www.eduref.net/features/what-grades-can-ai-get-in-college/ [retrieved: June 2021]

[23] Google Research, *Verse by Verse*. [Online]. Available from: https://sites.research.google/versebyverse/ [retrieved: June 2021]

[24] S. Devlin and R. Goodwin, *Poemportraits*. Available from: https://experiments.withgoogle.com/poemportraits [retrieved: June 2021]

[25] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen and I. Sutskever, Zero-shot text-to-image generation, arXiv:2021.12092, 2021.

[26] OpenAI, *DALL-E:creating images from text*, [Online]. Available from: https://openai.com/blog/dall-e/ [retrieved: June 2021]

[27] J. Schachter, Twitter feed, [Online]. Available from: https://twitter.com/joshu/status/1286009425220706305 [retrieved: June 2021]

[28] S. Fuentes, E. Tongson, D. Torrico and C. Gonzalez Viejo, "Modeling Pinot Noir Aroma Profiles Based on Weather and Water Management Information Using Machine Learning Algorithms", Foods 9(1), 2021.

[29] D. Park, K. Kim, Y. Park, J. Shin and J. Kang, "KitcheNette: Predicting and Ranking Food Ingredient Pairings using Siamese Neural Networks", Proc. Int'l. Joint Conf. on AI (IJCAI-19), Macao, 2019.

[30] L. Varshney, F. Pinel, K.R. Varsheny, D. Bhattacharjya, A. Schorgendorder and Y-M..Chee, "A big data approach to the computational creativity: The curious case of Chef Watson", IBM Journal of Research and Development, 63(1), 2019

[31] M. Serifovic, *Image-to-recipe translation with Deep Convolutional Neural Networks*, 2018. [Online]. Available from: https://towardsdatascience.com/this-ai-is-hungry-b2a8655528be [retrieved: June 2021]

[32] C. Ledig et al, Photo-realistic image super-resolution using a generative adversarial network, arXiv:1609.04802v5, 2017.

[33] DeepAI, *Super Resolution API*. [Online]. Available from: https://deepai.org/machine-learning-model/torch-srgan [retrieved: June 2021]

[34] D. Bau, H. Strobelt, W. Peebles, J. Wulff, B. Zhou, J-Y Zhu and A. Torralba, "Semantic Photo Manipulation with a Generative Image Prior", Proc. ACM SIGGRAPH, Los Angeles, 2019.

[35] MIT and IBM, *GANPaint Studio*, [Online]. Available from: https://ganpaint.io [retrieved: June 2021]

[36] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen and T. Aila, Analyzing and improving the image quality of StyleGAN, arXiv: 1912.04958v2, 2019.

[37] NVIDIA Research, *This person does not exist*. [Online]. Available from: https://thispersondoesnotexist.com [retrieved: June 2021]

[38] B. Uterwijk, *A.I. generated portraits*. [Online]. Available from: https://www.basuterwijk.com/portfolio [retrieved: June 2021]

[39] Artbreeder (formerly Ganbreeder). [Online]. Available from: https://www.artbreeder.com [retrieved: June 2021]

[40] picbreeder, [Online]. Available from: http://picbreeder.org/intro.php [retrieved: June 2021]

[41] J. Pinkney, *Making Toonify Yourself*. [Online]. Available from: https://www.justinpinkney.com/making-toonify [retrieved: June 2021]

[42] K. Vougioukas, S. Petridis and M. Pantic, Realistic speech-driven facial animation with GAN's, arXiv:1906.06337v1, 2019.

[43] D. Coldewey, "Mona Lisa frown: machine learning brings old paintings and photos to life", Techcrunch, May 22, 2019.

[44] C-Y. Weng, B. Curless and I.K. Shlizerman, *Photo wake-up: 3d character animation from a single photo*. [Online]. Available via: https://grail.cs.washington.edu/projects/wakeup [retrieved: June 2021]

[45] OpenAI, *Juebox*. [Online]. Available from: https://openai.com/blog/jukebox/ [retrieved: June 2021]

[46] Y. Ren et al, *DeepSinger: Singing voice synthesis with data mined from the web*. [Online]. Available from: https://speechresearch.github.io/deepsinger/ [retrieved: June 2021]

[47] Google, *Freddie Meter A.I. Experiment*. [Online]. Available from: https://freddiemeter.withyoutube.com/ [retrieved: June 2021]

[48] Drumbot, [Online]. Available from: https://drumbot.glitch.me/ [retrieved: June 2021]

[49] K. Newcombe, *Webcam Air Guitar*. [Online]. Available from: https://codepen.io/kevinnewcombe/pen/QRZdYP [retrieved: June 2021]

[50] M. Siegelman, *Deep Music visualizer*. [Online]. Available from: https://github.com/msieg/deep-music-visualizer [retrieved: June 2021]

[51] M. Alafriz, *Introducing "Lucid Sonic Dreams": Sync GAN Art to Music with a Few Lines of Python Code*. [Online]. Available from: https://towardsdatascience.com/introducing-lucid-sonic-dreams-sync-gan-art-to-music-with-a-few-lines-of-python-code-b04f88722de1 [retrieved: June 2021]

[52] Google Cloud, *AutoML*. [Online]. Available from: https://cloud.google.com/automl [retrieved: June 2021]

[53] Amazon, *SageMaker*. [Online]. Available from: https://aws.amazon.com/sagemaker/ [retrieved: June 2021]

[54] Google, *MediaPipe on Github*. [Online]. Available from: https://github.com/google/mediapipe [retrieved: June 2021]

[55] Runway | Make the impossible, [Online]. Available from:: https://runwayml.com/ [retrieved: June 2021]

[56] Microsoft, *Lobe desktop application*. [Online]. Available from: https://lobe.ai/ [retrieved: June 2021]

[57] V. Dibia, *Handtrack.js*, TensorFlow Blog [Online]. Available from: https://blog.tensorflow.org/2019/11/handtrackjs-tracking-hand-interactions.html [retrieved: June 2021]

[58] TensorFlow, *Pose estimation with TensorFlow*. [Online]: https://www.tensorflow.org/lite/examples [retrieved: June 2021]

[59] deeplearn.js, cghawthorne GitHub. [Online]. Available from: https://github.com/cghawthorne/deeplearnjs [retrieved: June 2021]

[60] T. O'Brien, *Cooking with Watson: Caymanian Plantain Dessert*, Engadget, June 2015.

[61] S. Robinson and A. Behbi, *How Sweet it is: using Cloud AI to whip up new treats with Mars Maltesers*, [Online]. Available from: https://cloud.google.com/blog, 2021 [retrieved: June 2021]

[62] M. Ives, "The Latest Artist Selling NFTs? It's a Robot", New York Times, March 25, 2021.

[63] M. A. Boden, Creativity and artificial intelligence, Artificial Intelligence, 102(1-2), pp. 347-356, 1998.

[64] J. Rowe, D. Partridge, "Creativity: a survey of AI approaches", Artifical Intelligence Review, 7(1), 2009.

[65] C-Z. A.Huang, A. Vaswani, J. Uszkoreit, N. Shazeer, I. Simon, C. Hawthorne, A. Dai, M. Hoffman, M. Dinculescu and D. Eck, Music Transformer, arXiv: 1809.04281v3

[66] K. Grow, *In Computero: Hear How AI Software Wrote a New Nirvana Song*, [Online]. Available from: https://www.rollingstone.com/music/music-features/nirvana-kurt-cobain-ai-song-1146444/ [retrieved: June 2021]