Roadmap on AI technologies & applications for the Media Industry

Section: "Evolutionary Learning"

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 951911

info@ai4media.eu www.ai4media.eu

Authors	Georgios N. Yannakakis (University of Malta) Antonios Liapis (University of Malta)

This report is part of the deliverable D2.3 - "*AI technologies and applications in media: State of Play, Foresight, and Research Directions*" of the AI4Media project.

You can site this report as follows:

F. Tsalakanidou et al., Deliverable 2.3 - AI technologies and applications in media: State of play, foresight, and research directions, AI4Media Project (Grant Agreement No 951911), 4 March 2022

This report was supported by European Union´s Horizon 2020 research and innovation programme under grant number 951911 - AI4Media (A European Excellence Centre for Media, Society and Democracy).

Evolutionary Learning

Current status

As its name suggests, ***artificial evolution*** emulates the paradigm of evolution of the species, put forward in the 18th century by biologists such as Darwin and Lamarck, as a way to find good solutions when given specific goals and constraints. At the highest conceptual level, artificial evolution creates a population of initial solutions, evaluates how good they are in terms of an objective, then selects the best among them and stochastically adjusts their parameters (often recombining two or more solutions together) to create a new population of offspring. This evaluation, selection, genetic change and reinsertion (see Figure 1) is carried out over multiple generations until the population converges towards better solutions.

Evolutionary computation has over 50 years of history[1] and is one of the pillars of computational intelligence (along with machine learning and reinforcement learning). In terms of the problems that artificial evolution is often called to solve, extensive focus has been placed on evolving computer programs that are able to carry out computations with a fairly freeform underlying structure, as well as numerical optimisation tasks where the potential of artificial evolution to reach global optimal solutions is most advantageous. Unlike machine learning, artificial evolution explores the search space stochastically, often maintaining a population of solutions.

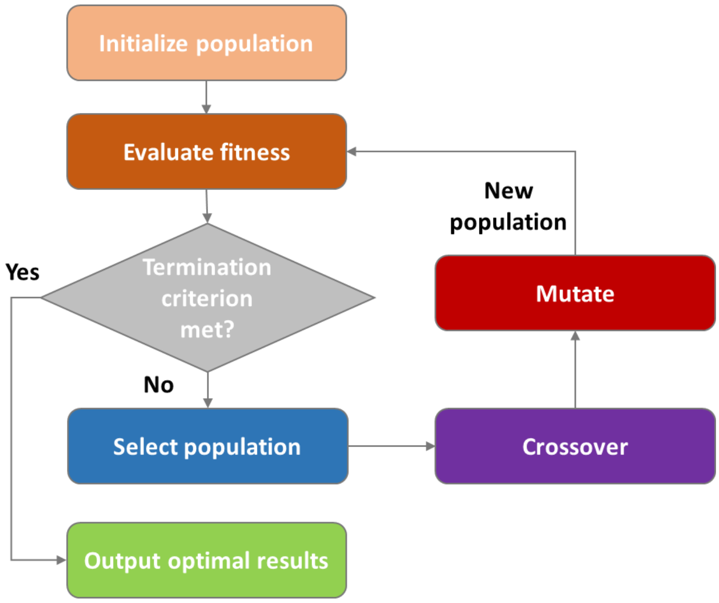

Figure 1: Evolutionary learning general concept[2].

In the media sector, artificial evolution has often been used to automatically or semi-automatically generate ***computational art***. Some of the earliest instances of evolutionary art

[1] K. A. De Jong. Evolutionary Computation: A Unified Approach. MIT Press, Cambridge, MA, USA. 2006

[2] Figure adapted from A. Doku, FitJSP - Fancy Interactive tool for Job-Scheduling problems (2020): https://blog.arinti.be/fitjsp-fancy-interactive-tool-for-job-scheduling-problems-791a9f6453ff

were the line drawings of Dawkin's *Biomorph*[3] (Figure 2) and Karl Sims' evolved computer programs that could generate 2D or 3D plants or 2D images[4]. The stochastic nature of artificial evolution, which allows for more wide exploration of the solution space compared to gradient descent, has made them particularly popular for the research of computational creativity[5].

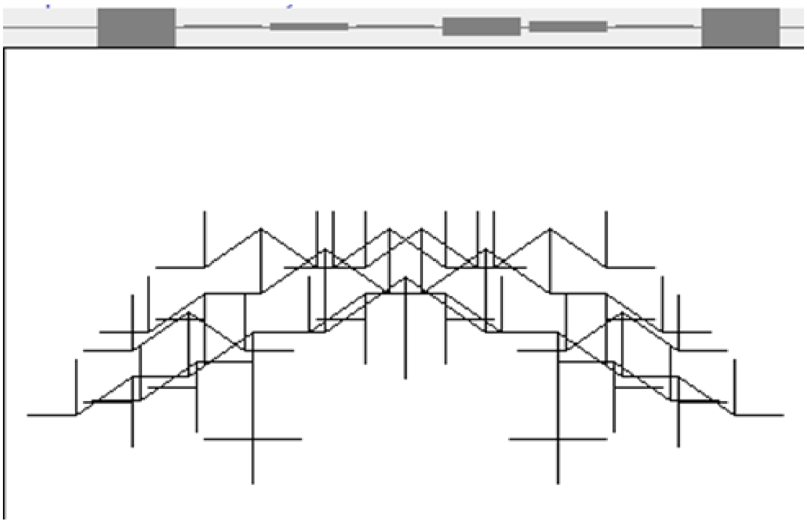

Figure 2: Dawkin's Biomorph[6], arguably the first instance of evolutionary art.

Computational creativity is "*the philosophy, science and engineering of computational systems which, by taking on particular responsibilities, exhibit behaviours that unbiased observers would deem to be creative*"[5]. Research in computational creativity often follows the paradigms established around human creativity, such as the concept of "p-creativity" (***psychological creativity***), when a creator considers their creations novel and valuable regardless of whether others would agree, and "h-creativity" (***historical creativity***), which introduces previously unimagined ideas or inventions into the world. Applications of computational creativity focus on artistic expression, including visual art (see for example, Figure 2, Figure 3 and Figure 4), music, narrative, humor and poetry[7]. Unlike numerical optimisation, however, formulating what constitutes "quality", "novelty" or a "valid solution" in computational creativity and evolutionary art raises significant challenges and debates[8].

---

[3] R. Dawkins. The blind watchmaker: Why the evidence of evolution reveals a universe without design. W.W. Norton & Co, New York, NY, USA. 1987.
[4] K. Sims. Artificial evolution for computer graphics. In Proceedings of the 18th SIGGRAPH. Association for Computing Machinery, New York, NY, USA, 1991.
[5] S. Colton, R. Lopez de Mantaras & O. Stock. Computational Creativity: Coming of Age. AI Magazine, 30(3), 11. 2009.
[6] Image source: Wikipedia - https://en.wikipedia.org/wiki/File:BiomorphBounce.png
[7] S. Colton & G. Wiggins. Computational creativity: The final frontier?. Frontiers in Artificial Intelligence and Applications. 2012.
[8] G. Ritchie. Some empirical criteria for attributing creativity to a computer program. Minds and Machines, 17:76–99, 2007

Figure 3: PicBreeder[9] uses neuroevolution to produce images that multiple users can interact with, evolve further, and evaluate through a public website.

When applying evolutionary computation to the media sector and towards computational creativity more broadly, a major challenge is assessing the quality of generated content in such aesthetic-oriented, subjective domains, in order e.g. to guide the generator towards better content. It is not surprising that some of the early work in evolutionary art rely on a human curator to guide evolution; similar practices in leveraging humans to perform interactive evolution are still popular today (an example can be seen in Figure 3).

Research challenges

A core challenge of applying evolutionary computation in the media sector remains the ***evaluation of quality***. To guide artificial evolution towards better content, it is common to use existing corpora such as human ratings[10], classifiers between man-made and generated images[11], object recognition[12] (see also Figure 4) or models that match images with language[13].

---

[9] J. Secretan, N. Beato, D.B. D'Ambrosio, A. Rodriguez, A. Campbell, J. T. Folsom-Kovarik, and K. O. Stanley. 2011. Picbreeder: A case study in collaborative evolutionary exploration of design space. Evol. Comput. 19, 3 (Fall 2011), 373–403. Image source: http://picbreeder.org/ (accessed 15 Dec. 2021)

[10] S. Baluja, D. Pomerleau, and T. Jochem. Towards automated artificial evolution for computer-generated images. Musical networks, pages 341–370, 1999

[11] P. Machado, J. Romero, A. Santos, A. Cardoso, and A. Pazos. On the development of evolutionary artificial artists. Computers & Graphics, 31(6):818–826, 2007.

[12] J. Correia, P. Machado, J. Romero, and A. Carballal. Evolving figurative images using expression-based evolutionary art. In Proceedings of the International Conference on Computational Creativity, pages 24–31, 2013.

[13] D. Norton, H. Darrell, and D. Ventura. Establishing appreciation in a creative system. In Proceedings of the International Conference Computational Creativity, pages 26–35, 2010.

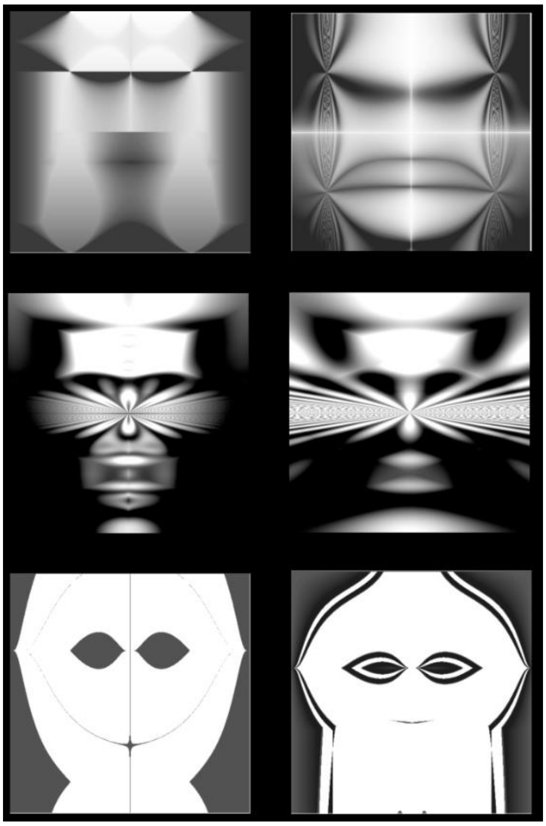

Figure 4: Evolved grayscale images guided by an evaluation based on the certainty of a trained classifier of human faces[14].

Treating evolutionary search as an *optimiser*, when it comes to creative media content generation, can be limiting and short-sighted. It is arguably impossible to adequately address the problem of searching for a “best” solution in domains and problem spaces where “good” is a subjective notion, as well as one that is deeply related to the context of use, intent, or current trends in a community of human (or AI) artists. Searching for solutions that are different from each other[15] can somehow mitigate this issue, by attempting to explore the space as thoroughly as possible rather than towards short-term exploitation. In experiments with robotics, divergent search has shown to be efficient in handling deceptive problems, where the final goal can only be reached by passing through “bad” parts of the space according to a pre-constructed, quantifiable notion of goodness[15]. However, what constitutes novelty in divergent search algorithms, as well as how quality can be formulated and maintained in quality-diversity algorithms remain open research challenges that can be as difficult to tackle as formulating an objective function for such subjective domains.

[14] Image source: P. Machado. 2021. Evolutionary art and design: representation, fitness and interaction. Proceedings of the Genetic and Evolutionary Computation Conference Companion. Association for Computing Machinery, New York, NY, USA, 1002–1031.
[15] J. Lehman and K. O. Stanley, “Abandoning objectives: Evolution through the search for novelty alone,” Evolutionary computation, vol. 19, no. 2, 2011

Societal and media industry drivers

Vignette: Designing new levels in a video game

Kiko is a game developer who wishes to design a new level for their upcoming game "*The Fabulous Journey*". Kiko works at a small game studio with only three developers (one programmer, one artist, and one level designer). Due to their small team, if they wish to publish "*The Fabulous Journey*" within the next two years they must rely on procedural content generation. Therefore, "*The Fabulous Journey*" is a series of spatial challenges and Kiko is tasked with designing these individual levels taking advantage of the power of the AI tool at their disposal. This puts low demands on the artist with regards to creation of art assets, character designs and animations since content is re-used throughout the levels, while the programmer can focus on the gameplay mechanics that are persistent throughout all levels. Kiko sits down to come up with an idea for the next level, starting up the AI tool and loading Kiko's profile, which contains all the history of their interaction and past designed and/or generated levels. Kiko considers some of the planned mechanics that the programmer has suggested, and wants to highlight a "wall-jump" mechanic that has not been very often used so far in past levels. The programmer has already updated the AI tool to use the current version of the game which uses the new mechanic, and has added a logger to count uses of "wall-jump" during a play session. Kiko uses the graphic user interface of the AI tool to specify the constraints for the next task: "use [wall-jump] [at least] [5] times" (text in brackets are options in drop-down lists). Kiko also uses the graphic user interface to specify that they want to explore levels of different length, and with different numbers of enemies. Kiko could also change what would qualify as the "best" level, but they keep the same metric that the studio has been using throughout production, which is the level that has the highest score after a simulation with an AI agent.

With everything setup, Kiko presses the *Generate* button and after a short while, a number of level layouts start appearing on the screen. Kiko can wait until the system has produced all the best levels with few enemies, many enemies, short length, long length and any combination thereof. Kiko can choose the computer-generated levels, and get summary statistics from the simulation (such as score of the agent at the end of the level, number of times that each mechanic was used, number of deaths by enemies, number of enemies killed etc.). Moreover, Kiko can choose to watch the AI agent's playthrough by re-running the simulation. Finally, Kiko can choose to play through the level themselves. Kiko can also stop the AI generation earlier, and select some of the levels that they prefer. Then Kiko can continue running the automated process: the AI will generate levels more similar to Kiko's selections. Alternatively, Kiko can enter an editing tool and directly modify the levels generated by the AI. Once Kiko has finished editing, they can playtest the level themselves or have an AI agent playtest it and report some summary statistics of the simulation. Kiko can export the levels they created or some hand-picked generated levels, and add them to the current version of the game. The programmer, artist, and also Kiko can further edit these in future iterations (e.g. after more content or more code has been produced) with or without the use of the AI tool.

Future trends for the media sector

Evolutionary computation is a powerful tool for exploring a large variety of designs. Coupled with other computational intelligence algorithms, such as recommender systems and latent representations, evolutionary computation can produce high-quality and personalised content that is appropriate to show to a designer during their workflow. Often, such AI-provided ideation mechanisms are used in early conceptual phases, allowing the designer to see many options but leaving more room for human creativity and control during later stages of the design where getting the details right is critical. Moreover, allowing the designer to keep their own artistic vision of the final product is imperative: this can be facilitated by interfaces that customise the initial parameters for exploration, by choosing specific examples that the user would prefer the AI to move towards, or by manually editing interim products to help the AI start from a better seed and move towards a better direction. All of these modes of interaction are described in the vignette above.

Based on these properties and requirements for integrating evolutionary algorithms in the design process, future research trends in this vein for the media sector will have to address four main issues: (a) the ***type of representations*** that evolution can explore, (b) the way in which ***quality and diversity*** are calculated, (c) ways of ***modelling designers*** in order for the AI to produce more personalised artefacts, and (d) ***interaction paradigms*** for the users to be able to view, control, and make use of the generated artefacts. The first two goals are likely best tackled through deep learning methodologies which can produce a more compact representation that evolution can more easily explore (compared to e.g. pixel-level representations), while supervised and unsupervised learning can be used to train predictive models of an artifact's quality and a dataset's diversity respectively. These predictive measures of quality and diversity can be used instead of the current mathematically defined formulas for quality-diversity evolutionary algorithms. Early work has already started to explore this direction, in non-media domains[16].

The last two goals are as reliant on traditional human-computer interaction and user modelling as they are on explainable AI research[17]. Research in this vein has so far focused on which of the large number of generated content to show to a user, given the cognitive overload of too many options. Moreover, early work has focused on personalising designer models in terms of different prescriptive quality and diversity dimensions, based on which AI generated solutions the user tends to select over unselected ones[18]. However, with the introduction of machine learning in all aspects of the design pipeline (including the future trends of combining evolution and latent vector representations discussed above), explaining to the designer why the level is considered of good quality – or more importantly, why the system considers the level

[16] A. Cully. Autonomous skill discovery with quality-diversity and unsupervised descriptors. In Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '19). Association for Computing Machinery, New York, NY, USA, 81–89. 2019

[17] J. Zhu, A. Liapis, S. Risi, R. Bidarra and G. M. Youngblood, "Explainable AI for Designers: A Human-Centered Perspective on Mixed-Initiative Co-Creation," Proceedings of the IEEE Conference on Computational Intelligence and Games, 2018.

[18] A. Liapis, G. N. Yannakakis and J. Togelius. Designer modeling for Sentient Sketchbook. Proceedings of the 2014 IEEE Conference on Computational Intelligence and Games, 2014.

appropriate for this designer – is vital for the designer to be able to trust the system and use it for their creative work. Especially in domains where artistic vision and creativity are paramount, the only opportunity for AI to be able to prompt co-creativity[19] is by providing ***an AI partner that can be both supportive but also "honest" and transparent.***

Goals for next 10 or 20 years

Public and academic attention in generative art has been largely driven by deep-learning-based architectures. One can only expect that short-term future accomplishments in creative domains such as images, music, and text generation will largely rely on these corpora-driven, trained models and AI methods such as transformers[20], generative adversarial networks[21] and style transfer[22]. However, ensuring that evaluation of the quality of generated media is scientifically robust and replicable remains an open challenge that will need to be addressed in the next 10 years. Theoretical constructs from computational creativity research[8] can shed important light on identifying the novelty and quality of such generated artworks. Standards for what constitutes original and "authentic" are necessary not only for the purposes of ascertaining intent and creativity but also for handling Intellectual Property and financial gain.

Beyond short-term accomplishments, integrating deep learning paradigms in interactive evolution is expected to lead to more promising long-term accomplishments. The main benefit of artificial evolution and specifically divergent or quality-diversity search would be that it explores the possible set of solutions better, and can thus lead to more varied outcomes than the gradient-based methods used currently. A major drawback of evolutionary computation is that this additional exploration comes at a computational cost. Long-term hardware and engineering developments that can parallelise such search processes can benefit the media sector both by training better deep learning models through neuroevolution[23] and by generating a more diverse set of creative artefacts based on pre-trained representations[24] and driven by pre-trained or custom-trained evaluations of quality, diversity, appropriateness, personal preference, and more.

[19] G. N. Yannakakis, A. Liapis and C. Alexopoulos. Mixed-Initiative Co-Creativity. Proceedings of the 9th Conference on the Foundations of Digital Games. 2014

[20] T. Brown, et. al. Language Models are Few-Shot Learners. Advances in Neural Information Processing Systems 33. 2020.

[21] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial networks. In Proceedings of the International Conference on Neural Information Processing Systems, 2014

[22] Z. Hu, J. Jia, B. Liu, Y. Bu, and J. Fu. Aesthetic-Aware Image Style Transfer. In Proceedings of the 28th ACM International Conference on Multimedia. Association for Computing Machinery, New York, NY, USA, 3320–3329. 2020.

[23] K. O. Stanley and R. Miikkulainen. Evolving neural networks through augmenting topologies. Evol. Comput. 10, 2 (Summer 2002), 99–127. 2002.

[24] M. C. Fontaine and S. Nikolaidis. Differentiable Quality Diversity. Proceedings of the 35th Conference on Neural Information Processing Systems. 2021.

AI4media

ARTIFICIAL INTELLIGENCE FOR THE MEDIA AND SOCIETY

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 951911

info@ai4media.eu www.ai4media.eu