



Increasing generality in machine learning through procedural content generation

Sebastian Risi^{1,2} and Julian Togelius^{1,3}

Procedural content generation (PCG) refers to the practice of generating game content, such as levels, quests or characters, algorithmically. Motivated by the need to make games replayable, as well as to reduce authoring burden and enable particular aesthetics, many PCG methods have been devised. At the same time that researchers are adapting methods from machine learning (ML) to PCG problems, the ML community has become more interested in PCG-inspired methods. One reason for this development is that ML algorithms often only work for a particular version of a particular task with particular initial parameters. In response, researchers have begun exploring randomization of problem parameters to counteract such overfitting and to allow trained policies to more easily transfer from one environment to another, such as from a simulated robot to a robot in the real world. Here we review existing work on PCG, its overlap with current efforts in ML, and promising new research directions such as procedurally generated learning environments. Although originating in games, we believe PCG algorithms are critical to creating more general machine intelligence.

For several decades, procedural content generation (PCG) has been a feature of many video games (Table 1). PCG refers to the algorithmic creation of game content—not the game engine, but things such as levels, quests, maps, characters or even rules—either in runtime (as the game is being played) or at design time (as the game is made). There are several reasons why PCG is used in games: it can increase the replayability of a game as players are presented with a new experience every time they play, it can help to reduce production costs and disk storage space, and it enables new types of games built on the unique affordances of content generation.

Interestingly, developments in PCG and machine learning (ML) have started to influence each other in reciprocal ways. Procedural content generation via machine learning (PCGML)¹ refers to the use of ML to train models on existing game content, and then leverage these models to create novel content automatically. This can be done through simply sampling from the learned models, or through searching the artefact space implied by the model so as to optimize some objective. Interestingly, PCGML poses different and hard challenges compared to generating images, for example, because the produced content needs to function.

At the same time that PCG researchers are starting to incorporate these advances into their systems, interest in the ML community is increasing in PCG-inspired methods to improve the robustness of ML systems. One reason for this development is the growing evidence that, while ML methods perform well for tasks or in the environments they are trained on, they do not generalize well when that environment is changed or different from what is seen during training. Training neural networks with many free parameters and over long training times has led to state-of-the-art performance in many domains, but these solutions are typically overfitted to the particular training examples, achieving high accuracy on a training set but performing poorly on data not used for training. Particularly in deep reinforcement learning (RL), in which an agent has to learn in interaction with its environment, overfitting is rarely addressed but a significant problem. Take the very popular arcade learning environment as an example², a classic benchmark in RL

based on an emulation of the Atari 200 games console. Hundreds of games were made for that console; however, they all have fixed sets of levels and very little in the way of randomization. Training an agent to play a game in the arcade learning environment makes it liable to overfit not only to that particular game, but also to its levels and sequence of events.

The basic idea of employing PCG to address the generality problem in ML systems is to artificially create more training data or training situations. This way, ML-based systems can be biased towards learning general task properties instead of learning spurious elements found in the training examples. Methods include simpler approaches, such as data augmentation, that artificially increase the data used for training, or methods that train agents in a large number of training environments that include randomized elements. PCG methods have been extended to create complete maps for *Capture the Flag*³ or maps for two-dimensional video games⁴, and in a recent impressive demonstration of the advantage of training in procedurally generated environments, have allowed a robot hand trained in a simulation to manipulate a Rubik's cube in the real world⁵.

In this Review, we examine the history of PCG and the recent trends in hybridizing PCG methods with ML techniques. The goal is to share with the larger ML research community work from this exciting area that is just beginning to capture the interest of researchers outside of games but could ultimately encourage the emergence of more general artificial intelligence (AI). Additionally, we aim to supply ML researchers with a new toolbox to aid their generalization work, and games researchers and developers with new perspectives from ML. This Review points out convergent research interests and complementary methods in the two communities, details promising future research directions and under-explored research avenues enabled by more advanced PCG techniques.

Classic PCG

While possibly the first video game to include PCG dates from 1978 (*Beneath Apple Manor* by Don Worth for the Apple II), *Rogue* (1980) by Toy and Wichmann created an important design paradigm. In *Rogue*, the player explores a multi-level dungeon complex,

¹modl.ai, Copenhagen, Denmark. ²IT University of Copenhagen, Copenhagen, Denmark. ³New York University, New York, NY, USA. [✉]e-mail: sebr@itu.dk

Table 1 | A comparison of several methods for PCG and domain randomization described in this article

	Representation ^a		Generation method ^b					
	Learned	Hand-designed	Evolution	Learned	Gradient-based	Random search	Sampling	Rules
Classic PCG								
Standard constructive ^c (for example, <i>Rogue</i> , <i>Pitfall!</i> , <i>Civilization</i> , <i>Elite</i> , <i>Minecraft</i>)	○	●	○	○	○	○	○	●
Dawn Fortress ^d	○	●	○	○	○	●	○	○
Data augmentation and domain randomization ^e								
Simple data augmentation ^{12,13}	●	○	○	○	○	○	○	●
Uniform domain randomization ¹⁸	○	●	○	○	○	○	●	○
Guided domain randomization ²⁴	○	●	○	●	○	○	○	○
Automatic domain randomization ⁵	○	●	○	○	○	○	●	○
Search-based PCG ^f								
Standard search-based PCG ^{29,32,37,92}	○	●	●	○	○	○	○	○
PCGML ^g								
Standard PCGML ^{1,49,93}	●	○	○	○	○	○	●	○
PCGML with constrained sampling ⁵⁰	●	○	○	○	○	●	○	○
Latent variable evolution ^{h51,52,94,95}	●	○	●	○	○	○	○	○
PCGRL ⁱ⁵⁴	○	●	○	●	○	○	○	○
Generative playing networks ⁱ⁵³	●	○	○	○	●	○	○	○
Procedurally generated learning environments								
POET ⁵⁶ , MCC ⁵⁵	○	●	●	○	○	○	○	○
Progressive PCG ^{k28}	○	●	○	○	○	○	○	●

^aThe first two columns indicate the representation of the content, which is either designed by hand or learned through ML. ^bThe last columns indicate how the content is generated given a representation. ^cConstructive methods follow rules and do not do any resampling. The black circles show which attributes each method has. ^dDawn Fortress is an example of a generate-and-test (random search) method where the world is regenerated if it fails certain tests. ^eThe various forms of domain randomization use hand-coded representations and differ in whether they simply sample this space or perform some sort of search with a learned policy. ^fIn the search-based paradigm, a hand-coded representation is searched using an evolutionary algorithm. ^gMost PCGML approaches randomly sample a learned representation, whereas PCGML with constraints resample when constraints are not satisfied. ^hLatent variable evolution combines search-based PCG with a learned representation, for example, in the form of a GAN. ⁱPCGRL uses a policy learned by RL to search a hand-coded representation, whereas generative playing networks instead uses RL to test the artefacts and gradient descent to generate them. ^jPOET and MCC are fundamentally search-based methods, which include a learning agent inside the evaluation loop. ^kProgressive PCG uses a parameterizable constructive generator, coupled to a RL-based game-playing agent.

battling enemies and collecting treasures. As the creators did not want to author the dungeons themselves (they wanted to play the game and be surprised), they needed to create a dungeon generation algorithm; every time you play a game of *Rogue*, a new set of dungeons are generated. *Rogue* came to inspire a genre of games called roguelikes, which are characterized mainly by the use of runtime generation of content that is essential to gameplay. The highly successful *Diablo* series of games (Blizzard, 1997–2013), as well as platformers such as *Spelunky* (Mossmouth, 2008), are roguelikes.

While the PCG in *Rogue* was motivated by a need for replayability and unpredictability, another key reason for using PCG is to create game worlds that are larger than can fit in memory or on storage media. A paradigm-setting game here was *Elite* (Brabensoft, 1984), a spacefaring adventure game featuring thousands of planets that seemingly miraculously fit in memory on a Commodore 64, with 64 kilobytes of memory. Every time a star system was visited, the game would recreate the whole starsystem with planets, space stations and spacecraft, from a given random seed. This approach has later been used for games such as *No Man's Sky* (Hello Games, 2015), which famously contains more planets than you can visit in a lifetime, all with their own ecologies.

The strategy games in the very popular *Civilization* series also rely heavily on PCG, as a new world is created for the players to explore and contest every time a new game is created. Similarly, the open-world sandbox game *Minecraft* (Mojang, 2010) creates a completely new world at the start of each game session. Other games

use PCG in more peripheral roles, such as the sidequest generation (for example, creating an infinite supply of fetch quests through a guild system) in *The Elder Scrolls V: Skyrim* (Bethesda, 2011) (along with some earlier games in the series) and the pervasive generation of terrain features and vegetation in a large number of open-world three-dimensional games. PCG techniques are now so commonplace and reliable that it is more common than not to utilize them in many game genres.

Interestingly, PCG in video games is actually prefigured by certain pen-and-paper generators intended to be executed by humans with the help of dice or cards, including a dungeon generator for the classic *Dungeons and Dragons* (TSR, 1976) role-playing game⁶. Some recent board games that include aspects of PCG are *504* (2F Spiele, 2015) or *Betrayer at House on the Hill* (Avalon Hill, 2004).

The types of PCG that can be found in most existing games are called constructive PCG methods (Table 1). This means that the content generation algorithm runs in a fixed time, without iteration, and does not perform any search. For generating textures, heightmaps, and similar content, a commonly used family of algorithms are fractal noise algorithms such as Perlin noise⁷. Vegetation, cave systems, and similar branching structures can be efficiently generated with graphically interpreted grammars such as L-systems⁸. Other constructive methods that are borrowed from different fields of computer science, and were adapted to the needs of PCG in games, include cellular automata⁹ and other approaches based on local computation. Other constructive methods are based on rather

less principled and more game-specific methods. For example, *Spelunky* combines a number of pre-authored level chunks according to patterns that are designed so as to ensure unbroken paths from entrance to exit.

In this Review we focus on examples where there's a notion of 'environment' (typically through learning being centred on an agent in simulated physical world), but will throughout the text mention other learning settings where relevant for comparison. We will cover work coming out of both commercial game development, AI/ML research targeted at games, and AI/ML research targeted at other applications. Our (imperfect) distinction between PCG and other methods is that pure randomization/shuffling is not PCG; however, many PCG algorithms include randomness.

Data augmentation and domain randomization

While not necessarily called PCG in the ML community, the idea of data augmentation is essentially a simple form of constructive PCG. These methods aim to address the problem of overfitting in ML, that is, achieving a high accuracy on a training set but performing poorly on data not used for training^{10,11}.

Data augmentation methods increase the diversity in the dataset, not by collecting more data but by adding modified versions of the already existing data^{12,13}. Data augmentation is very common in supervised learning tasks, for example, through cropping, padding or adding noise to images in a dataset. It is common practice in ML and has resulted in substantially less overfitting and state-of-art results in a variety of domains^{12,14,15}.

A different form of data augmentation was introduced by Geirhos et al.¹⁶, in which the authors showed that training the same network architecture but with a stylized version of ImageNet images (for example, a cat with the texture of an elephant) can significantly increase the model's accuracy and robustness. In fact, the authors showed that a deep convolutional network trained on the standard ImageNet dataset mainly focuses on textures in images instead of their shape; training on the stylized version of ImageNet increases their shape bias and with that, their accuracy and robustness.

In the field of RL, domain randomization^{17–19} is a simple form of PCG and one way to counter overfitting in ML. The main idea of domain randomization is to train an agent in many simulated training environments, where certain properties are different in each environment. The goal is to learn a single policy that can work well across all of them. In addition to trying to encourage ML systems to be more robust and general, another use case of domain randomization is to facilitate the transfer of policies trained in a simulator to the real world^{18,20–22}. Training in a simulation instead of the real world has several advantages such as the training being faster, cheaper and more scalable, and having access to the ground truth.

In a promising demonstration of this approach, Tobin et al.¹⁸ trained an object detector on thousands of examples of objects with randomized colours, textures, camera positions, lighting conditions and so on in a simulator and then showed it can detect objects in the real world without any additional training. Another example is the work by Sadeghi et al.²⁰, who trained a vision-based navigation policy for a quadrotor entirely in a simulated environment with highly randomized rendering settings and then transferred this policy to the real world without further training.

Following Weng¹⁷, we can further divide domain randomization into three subgroups: uniform domain randomization, guided domain randomization, and automatic domain randomization. In uniform domain randomization, each parameter is uniformly sampled within a certain range. For example, in the work by Tobin et al.¹⁸, the size of objects, their mass or the amount of noise added to the camera image were drawn from a uniform distribution.

In the more sophisticated guided domain randomization, the type of randomization is influenced by its effect on the training process^{17,23–25}. The goal of this guided randomization is to save

computational resources by focusing the training on aspects of the task that actually increase the generality of the model. For example, instead of randomly applying pre-defined and hard-coded data augmentation methods, the approach AutoAugment²⁴ can learn new data augmentation techniques. These augmentation techniques are optimized for based on their validation accuracy on the target dataset. Such methods can be seen as a form of adaptive content generation; in the PCG literature there are approaches to PCG that adapt to an agent driven by, for example, Schmidhuber's theory of curiosity^{26,27}.

Another related approach is DeceptionNet²⁵, which is trained to find modifications to an image through distortion, changing the background and so on that make it harder for an image recognition network to output the correct classification. Both a recognition and deception network are alternatively trained such that the deception module becomes better in confusing the recognition module, and the recognition module becomes better in dealing with the modified images created by the deception module.

Very recently, OpenAI showed that a neural network that controls a five-fingered humanoid robot hand to manipulate a Rubik's cube can sometimes solve this task in the real world even though it was only trained in a simulated environment⁵. Key to this achievement in robotic manipulation was training the robot in simulation on a large variety of different environmental variations, similar to the domain randomization approaches already mentioned. Following related work in PCG for games²⁸, the ingredient to make this system work was to increase the amount of domain randomization, as the robot gets better and better at the task. For example, while the network was initially only tasked to control a Rubik's cube of 5.7 cm, later in training it had to deal with cubes that could range from 5.47–6.13 cm in simulation. Because the robot had to deal with many different environments, dynamics of meta-learning did emerge in the trained neural network; this allowed the robot to adapt to different situations during test time, such as the transfer to the real world. Automatic domain randomization is similar to guided domain randomization but focuses more on increasing the diversity of the training environments based on task performance, instead of sampling efficiently from a distribution of environments.

While current domain randomization methods are showing promising results, the PCG community has invented many sophisticated algorithms that—we believe—could greatly improve the generality of ML methods even further. As we discuss in the following sections, more recent work in PCG has focused on search-based approaches and on learning the underlying PCG representations through ML techniques.

AI-driven PCG

Given the successes of PCG in existing video games, as well as the perceived limitations of current PCG methods, the past decade has seen a new research field form around game content generation. The motivations include being able to generate types of game content that cannot currently be reliably generated, making game developments easier and less resource-intensive, enabling player-adaptive games that create content in response to player actions or preferences, and generating complete games from scratch. Typically, the motivations centre on games and players; however, as we shall see, many of the same methods can be used for creating and varying environments for developing and testing AI.

While there has been recent work on constructive methods, more work has focused on approaches based on search and/or ML.

Search-based PCG. In search-based PCG (Table 1), stochastic search/optimization algorithms are used to search for good content according to some evaluation function²⁹. Often, but not always, some type of evolutionary algorithms is used, due to the versatility of these algorithms. Designing a successful search-based content generation solution hinges on designing a good representation,



Fig. 1 | PCG-based games. **a, b**, In academia, PCG approaches have been used to produce complete and playable 3D games³³ (**a**) and rules for two-dimensional games²⁷ (**b**). **c**, PCG-enabled games include *Petalz*, in which players can collaboratively breed an unlimited variety of different procedurally generated flowers⁴². **d**, PCG also allows the creation of maps and character classes for first-person shooters³⁶. **e**, *Yavalath* is one of the few examples of commercially available games where the game rules are procedurally generated³².

which enables game content to be searched for. The representation affects, among other things, which algorithms can be used in the search process; if the content can be represented as a vector of real numbers, this allows for very strong algorithms such as CMA-ES³⁰ and differential evolution³¹ to be used. If the representation is, for example, a graph or a permutation, this poses more constraints on the search.

An early success for search-based PCG is Browne and Maire's work on generating board games, using a game description language capable of describing rules and boards for classical board games³². The initial population was seeded with a dozens of such games, including *Checkers*, *Connect Four*, *Gomoku* and *Hex*. The evaluation function was simulation-based; candidate games were evaluated through being played with a Minimax algorithm combined with a state evaluation function automatically derived for each game. The actual game evaluation is a combination of many metrics, including how often the game leads to a draw, how early in the game it is possible to predict the winner and the number of lead changes. This process, though computationally very expensive, came up with at least one game (*Yavalath*), which was of sufficient quality to be sold commercially (Fig. 1e).

While attempts to create complete video games including game rules through search-based methods have met mixed success^{27,33–35}, search-based PCG has been more effective in generating specific types of game content such as levels. We have seen applications to generating maps for the real-time strategy game *StarCraft*³⁶, and levels for the platform game *Super Mario Bros*³⁷, the first-person shooter *Doom*³⁸, and the physics puzzle game *Angry Birds*³⁹, among many similar applications. Search-based PCG has also been used for other types of game artefacts. In this paper, the term artefact refers

to objects made by an algorithm, such as particle effects for weapons⁴⁰, role-playing game classes⁴¹ or flowers⁴² (Fig. 1).

The most important component in a search-based content-generation pipeline is the evaluation function, which assigns a number (or vector) to how desirable the artefact is. In many cases, this is accomplished by playing through the content in some way and assigning a value based on characteristics of the gameplay, as in the *Yavalath* example in Fig. 1³²; other evaluation functions can be based on directly observing the artefact, or on some machine-learned estimate of, for example, player experience.

An emerging trend is to go beyond optimizing for a single objective and instead trying to generate a diverse set of artefacts that perform well. The goal here is to generate, for example, not just a single level but a set of levels that vary along various dimensions, such as the number of enemies or difficulty to solve for an A* algorithm⁴³. The Map-Elites algorithm, originally introduced to create more robust robot gaits⁴⁴, has been adapted to create sets of game levels that vary in what skills they require from the agent or what mechanics they feature⁴⁵.

An alternative to stochastic optimization algorithms is to use constraint satisfaction methods⁴⁶ such as answer set programming⁴⁷. Casting artefacts as answer sets can allow very efficient search for content that obeys specific constraints, but is hard to integrate with simulation-based evaluation methods. This paradigm is sometimes called solver-based PCG.

PCG via ML. ML methods such as generative adversarial networks (GANs)⁴⁸ have revolutionized the way we generate pictorial content, such as images of faces, letters and various objects. However, when

generating game content with some form of playability constraints (such as levels, maps or quests), things become more complicated because these types of content are in some ways more like code than images. An image of a face where the contours smudge just looks slightly off, whereas a level for *Super Mario Bros* with an impossibly long jump is not just a small defect, it's unplayable and therefore worthless. Similar functionality requirements can be found in level-like artefacts such as robot path planning problems, logic puzzles and quests¹. Therefore we call such content, in which some algorithmic way of verifying their functionality (for example, playability) exists, functional content.

Simply training a GAN on a large set of functional artefacts does not guarantee that the generator network learns to produce levels that fulfill these functionality requirements, nor that the discriminator learns to identify and check for those constraints. The result is often artefacts that look right but don't function well⁴. Another potential reason for the failure of ML-based methods to generate functional content is that methods such as GANs mostly learn local dependencies, whereas functionality in many types of content can depend on features that are far from each other, and/or counting the number of instances of a feature.

The same effect has been found with other representations, such as long short-term memory networks⁴⁹ and Markov chains⁵⁰. One way of counteracting this effect is bootstrapping, where newly generated artefacts that are found to satisfy the functionality requirements are added back to the training set for continued training, thus biasing training specifically to functional artefacts⁴.

ML models can also be combined with search to improve their efficiency. One way to do this is to use the learned model as a representation for search-based PCG. The idea here is to use ML to find the general space of content that is roughly defined by the examples the model is trained on, and then search within that space. Using GANs, this could be done by searching the latent space; when training a GAN, a latent vector is used as input to the generator network. The latent space is defined by that input. Latent variable evolution refers to using evolutionary algorithms to search the latent space for artefacts that optimize some kind of objective function⁵¹. For example, latent variable evolution was used to generate new levels for *Super Mario Bros*, by first training a GAN on one-screen segments of most levels in the original game. The latent space was then searched for vectors that would maximize objectives such as that the segment should contain many jumps, or should not be winnable without jumping, or should be unwinnable⁵².

Functionality evaluation can be integrated into adversarial learning processes in other ways. Generative playing networks consist of a generator network that generates levels, and a RL agent that learns to play them⁵³. While the objective for the playing agent is simply to perform as well as possible on the level, the objective for the level-generating agent is to provide an appropriate level of challenge for the agent.

Another way of using ML for PCG is to use RL. The conceptual shift here is to see PCG as a sequential process, where each action modifies a content artefact in some way. The goal of the training process then becomes to find a policy that for any content state selects the next action so that it leads to maximum expected final content quality. For this training process to be useful, we will need the trained policy to be a content generator capable of producing diverse content, rather than simply producing the same artefact every time it is run. A recent paper articulates a framework for PCG via RL and proposes methods for ensuring that the policy has sufficiently diverse results in the context of generating two-dimensional levels⁵⁴. Two important lessons learned is to always start from a randomized initial state (which need not be a functional level) and to use short episodes, to prevent the policy from always converging on the same final level. (It is interesting to note that the issues with learning general policies in RL recur in trying to learn policies that create content that can help generalize RL policies.)

Compared to PCG based on supervised or self-supervised learning, PCG based on RL has the clear advantage of not requiring prior content to train on, but the drawback of requiring a reward function judging the quality of content. This is very similar in nature to the evaluation function in search-based PCG. Compared to search-based PCG, PCG via RL moves the time and computation expense from inference to training stage; whereas search-based PCG uses extensive computation in generating content, PCG via RL uses extensive computation to train a model that can then be used cheaply to produce additional content.

Procedurally generated learning environments

An exciting opportunity for PCG algorithms is to create the actual learning environments that scaffold the learning of artificial agents (Fig. 2). Similarly to how current ML methods are moving towards automating more and more facets of training (for example, meta-learning the learning algorithms themselves, learning network architectures instead of hand-designing them), the automated generation of these progressive curricula that can guide learning offers unique benefits.

One of the first examples of this idea is minimal criterion coevolution (MCC)⁵⁵. In MCC both the agent and the environment co-evolve to solve increasingly more difficult mazes. Recent work building on these ideas is POET⁵⁶, which deals with the more challenging OpenAI gym bipedal walker domain. POET is a good example of an approach in which solutions to a particular obstacle course can function as stepping stones for solving another one. In fact, for the most difficult environments (shown on the right in Fig. 2a) it was not possible to directly train a solution; the stepping stones found in other environments were necessary to solve the most ambitious course.

Importantly, procedurally generated training environments can also increase the generality of RL agents that typically overfit to their particular environment^{28,57,58}. Zhang et al.⁵⁷ showed that training on thousands of levels in a simple video game can allow agents to generalize to levels not seen before. Some domains in the OpenAI gym training environments include procedurally generated content, requiring the agents to learn more general strategies. For example, in the CarRacing-v0 environment⁵⁹, agents are presented with a new procedurally generated car racing track every episode and the final reward is the average reward over multiple rollouts. These procedurally generated environments required more sophisticated neural architectures to be solvable⁶⁰, highlighting their usefulness in testing the ability of the RL agents. A similar approach for encouraging the discovery of general policies worked well for evolving stable policies for a two-dimensional bipedal walker domain⁶¹. In addition to helping in supervised learning settings (see the 'Data augmentation and domain randomization' section), forms of data augmentation can also help RL agents to become more robust. Randomized environments are also present in the Arena multi-agent testbed⁶² and the MazeExplorer testbed⁶³, both built on first-person shooters. In the work by Cobbe et al.^{64,65} agents are trained in environments in which random rectangular regions of the environment are cut out and replaced by rectangles filled with random colours, which helps these agents to generalize better.

Jaderberg et al.³ relied on a PCG-based approach to allow RL agents to master Quake III Arena's *Capture the Flag* (Fig. 2b). In their work, agents were trained on a mixture of procedurally generated indoor and outdoor maps (with varying walls and flag locations), which allowed the agents to learn policies that are robust to variations in the maps or the number of players. This work also demonstrated another advantage of procedurally generated maps: because each map is different, agents learned to learn how to keep track of particular map locations (for example, the entrance to the two bases) through their external memory system.




Fig. 2 | Examples of learning environments created by PCG-based approaches. **a**, The POET algorithm learns to create increasingly complex environments for a two-dimensional bipedal walker together with their neural network controllers⁵⁶. **b**, Procedurally generated maps were one of the key ingredients to allow agents to master the Quake III *Capture the Flag* domain³. **c,d**, Increasing task complexity depending on the performance of the agent has shown to lead to more general solutions for controlling a robot hand for dexterous in-hand manipulating in simulation and in the real world^{5,66} (**c**), and video game playing²⁸ (**d**).

While the aforementioned work^{3,57} showed that training on a larger variety of environments can lead to more general agents, it did require a large number of training levels. In work from the PCG research community, Justesen et al.²⁸ introduced a progressive PCG approach (PPCG), which showed that performance of training agents can be increased while using less data if the difficulty of the level is changed in response to the performance of the agents (Fig. 2d). A similar approach was later adopted by OpenAI to train their humanoid robot hand (Fig. 4c) in increasingly more challenging environments^{5,66}.

In fact, there is some initial evidence that very varied training environments can also foster the emergence of meta-learning in recurrent neural networks, which allows adaption to situations not seen during training⁵. While approaches such as OpenAI's Rubik's cube solving robot hand hint at the potential of this approach, creating an encoding that can produce an even larger variety of different and effective training environments could have a substantial impact on the generality of the agents and robots we are able to train.

We also summarize the similarities and differences between POET, MCC and PPCG in Table 1. While all three approaches use hand-designed representations, PPCG does not evolve the levels but instead uses a ruled-based generator.

Opportunities and challenges

We believe the idea of automatically and procedurally generating learning environments with the right complexity that scaffold the learning of autonomous agents is an exciting research direction that can help overcome some of the constraints that impede generalization and open-ended learning in current AI. This research direction is similar to what has been proposed in PCG research before, and also to the idea of AI-generating algorithms⁶⁷. We identify six main open challenges that we believe are essential in pushing the field of PCG forward, allowing it to realize its promise to create more adaptive and lifelong learning ML agents.

Learning from limited data. When generating images with ML, it is common practice to train the model on thousands, maybe even millions, of images⁶⁸. However, such amounts of high-quality

data are rarely available when developing a game, or even in a finished game. For example, the original *Super Mario Bros* game has 32 levels, resulting in a few hundred screens' worth of content. For some games, a large amount of user-generated content is available online, but this content can be of very variable quality. And when creating, for example, a robot-learning benchmark from scratch, creating scenarios to train a content model on can be a substantial time investment. Bootstrapping in PCG⁴ (see subsection 'PCG via ML') can help overcome this problem of content shortage, and various data augmentation could also help but learning to generate new content from limited data is still a significant challenge. More research is needed on how to learn from little data, and on how to learn generative models based on many different types of data. For example, by training a model on lots of available benchmark rules to learn generic patterns, it should be possible to generate environments for a new benchmark.

Generating complete games. While PCG techniques have shown impressive results for particular types of content in particular game genres, there has been much less progress on the harder problem of generating complete games. Browne and Maire's work from 2010 (discussed above³²), which resulted in a well-reviewed board game that is sold in stores, remains the gold standard. Generating complete video games^{33–36,47,56,69} (Fig. 1) or card games⁷⁰ seems to be a much harder challenge, with the results often being unplayable or uninteresting. Methods that have been tried include constraint satisfaction through answer set programming as well as evolutionary search. This is partly because of these games are very complex, and partly because it is very hard to find good evaluation metrics for complete games. Yet, generating complete challenges, including rules, topology, visuals and so on, seems to be a crucial part of a process where we gradually scale up challenges for agents that are capable of completing not just one challenge, but multiple ones.

PCG via ML could be a potentially promising approach to tackle this challenge. For example, Fan et al.⁷¹ very recently showed that a neural network can learn from crowd-sourced elements such as descriptions of locations and characters to create multiplayer text adventure games. This idea of leveraging and integrating real-world

data to create games (also known as data games), was first proposed by Gustafsson et al.⁷² and later extended to procedurally generate simple adventures games using open data from Wikipedia⁷³. Another example of how to leverage advances in ML for PCG is the recent *AI Dungeon 2* text adventure game⁷⁴. In this game, players can type in any command and the system can respond to it reasonably well, creating the first never-ending text adventure. The system is built on OpenAI's GPT-2 language model⁷⁵, which was further fine-tuned on a number of text adventure stories. This work also highlights that ML techniques combined with PCG might lead to completely new types of games that would not have been possible without advanced AI methods.

Lifelong generation for lifelong learning. The problem of lifelong learning is that of continuously adapting and improving skills over a long lifetime of an agent, comprising many individual episodes, though not necessarily divided into episodes as currently thought of^{76–78}. This would require an agent to build on previously learned skills as it faces increasingly harder or more complex, or just more varied, challenges. Lifelong learning is a problem, or maybe rather a setting, whose popularity has seemingly waxed and waned (under different names) as subsequent generations of researchers have discovered this challenge and then understood how hard it is. Within the artificial life community, the challenge of simulating open-ended evolution is closely related to that of lifelong learning. The idea behind open-ended evolution is to try to computationally replicate the process that allows nature to endlessly produce a diverse set of interesting and complex artefacts. Environments such as *Tierra*⁷⁹ and *Avida*⁸⁰ were early attempts at realizing that possibility.

The procedural generation of environments and challenges is a great opportunity for lifelong learning, and might even be a precondition for lifelong learning to be practically possible. It is possible that earlier attempts to realize lifelong learning have had limited success partly because the environments lacked sufficient challenges of the right complexity. The POET system shows one way of co-creating environments with agents⁵⁶. However, there is a great outstanding research challenge in devising mechanisms for gradually growing or complexifying environments (see the following subsection) so as to generate the right problems at the right time for agents to continually learn.

New PCG-based RL benchmarks. A variety of benchmarks have been proposed to test the generalization abilities of RL algorithms. Justesen et al.²⁸ used procedurally generated levels in the general video game AI (GVG-AI) framework⁸¹ to study overfitting of RL algorithms to different level distributions. In a similar vein to the work by Justesen et al.²⁸, levels in the *CoinRun* platform game are procedurally generated to quantify the ability of RL algorithms to generalize to never-before-seen levels^{64,65}. Another procedurally generated environment is the Unity game engine-based Obstacle Tower environment⁸², which requires increasingly complex skills such as locomotion, planning and puzzle-solving. Others have recently combined the Unity environment with GVG-AI, creating UnityVGDL⁸³, which allows ML agents in Unity to be tested on a large selection of games.

Other setups that do not use PCG include the work by Nichol et al.⁸⁴, in which *Sonic the Hedgehog* levels were separated into a training and test set to investigate how well RL algorithms generalize. In the Psychlab environment⁸⁵, agents are tested on known tasks from cognitive psychology, such as visual search or object tracking, making the results from simulated agents directly comparable to human results.

We propose the creation of PCG-based benchmarks in which the agent's environment and reward are non-stationary and become more and more complex over time. A starting point could be PCG approaches that are able to evolve the actual rules of a game

(see subsection 'Generating complete games'). New rules could be introduced based on agents' performance and estimates of their learning capacity. Adaptation within trials is as important as adaptation between trials: a generator could generate increasingly difficult games, which are different enough in each trial that a policy that would not adapt within a trial would fail. The Animal-AI Environment⁸⁶, in which agents have to adapt to unforeseen challenges based on classical tests from animal cognition studies, shares similar ideas with the benchmarks we are proposing here but does not focus on procedurally generated environments and tasks.

From simulation to the real world. Procedurally generated environments have shown their potential in training robot policies that can cross the reality gap. Promising work includes approaches that try to learn the optimal parameters of a simulator, so that policies trained in that simulator work well with real data^{87,88}. However, current approaches are still limited to lab settings, and we are far from being able to train robots that can deal with the messiness and diversity of tasks and environments encountered in the real world.

An intriguing opportunity is to train policies in much more diverse simulated environments than have been explored so far, with the hope that they will be able to cope better with a wider range of tasks when transferred to real physical environments. Both the Unity Simulation environment and Facebook's AI Habitat are taking a step in this direction. With Unity Simulation, Unity is aiming for simulation environments to work at scale, allowing developers to build digital twins of factories, warehouses or driving environments. Facebook's AI Habitat is designed to train embodied agents and robots in photorealistic three-dimensional environments to ultimately allow them to work in the real world.

In addition to developing more sophisticated ML models, one important research challenge in crossing the reality gap is the content gap⁸⁸. Because the synthetic content that the agents are trained on typically only represents a limited set of scenarios that might be encountered in the real world, the agents are likely to fail if they encounter situations that are too different from what they have seen before.

How to create PCG approaches that can limit this content gap and create large and diverse training environments, which prepare agents well for the real world tasks to come, is an important open research direction.

PCG beyond game-like environments. While PCG, as surveyed in this Review, is a set of techniques originally developed for video games with wide applicability in game-like environments, the ideas behind PCG can be generalized to many other domains. The idea of creating diverse artificial problems so as to enhance learning has wide applicability. For example, neural architecture search is a domain with some similarities to games, where the 'content' (network architectures) is evaluated on its functionality. Interestingly, many of the same techniques work well for neural architecture search and game content generation^{89–91}. Looking further afar, PCG techniques might play a role in many ML domains that at a first glance have very little in common with games.

Received: 30 November 2019; Accepted: 29 June 2020;

Published online: 3 August 2020

References

- Summerville, A. et al. Procedural content generation via machine learning (PCGML). *IEEE Trans. Games* **10**, 257–270 (2018).
- Bellemare, M. G., Naddaf, Y., Veness, J. & Bowling, M. The arcade learning environment: an evaluation platform for general agents. *J. Artif. Intell. Res.* **47**, 253–279 (2013).
- Max Jaderberg, M. et al. Human-level performance in 3D multiplayer games with population-based reinforcement learning. *Science* **364**, 859–865 (2019).
- Torrado, R. et al. Bootstrapping conditional gans for video game level generation. Preprint at <https://arxiv.org/abs/1910.01603> (2019).

5. Akkaya, I. et al. Solving Rubik's cube with a robot hand. Preprint <https://arxiv.org/abs/1910.07113> (2019).
6. Smith, G. An analog history of procedural content generation. In *Proc. 10th Int. Conf. Foundations of Digital Games* (FDG, 2015).
7. Perlin, K. An image synthesizer. *ACM Siggraph Comp. Graph.* **19**, 287–296 (1985).
8. Prusinkiewicz, P. Graphical applications of L-systems. In *Proc. Graphics Interface* Vol. 86 247–253 (ACM, 1986).
9. von Neumann, J. et al. The general and logical theory of automata. In *Cerebral Mechanisms in Behavior* (ed. Jeffress, L. A.) 1–41 (Wiley, 1951).
10. Hardt, M., Price, E. & Srebro, N. Equality of opportunity in supervised learning. In *30th Conf. Neural Information Processing Systems* 3315–3323 (NeurIPS, 2016).
11. Lin, J., Camoriano, R. & Rosasco, L. Generalization properties and implicit regularization for multiple passes sgm. In *Int. Conf. Machine Learning* 2340–2348 (ICML, 2016).
12. Krizhevsky, A., Sutskever, I. & Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems* 1097–1105 (NeurIPS, 2012).
13. Simard, P. Y., Steinkraus, D. & Platt, J. C. Best practices for convolutional neural networks applied to visual document analysis. In *Proc. Seventh International Conference on Document Analysis and Recognition* (IEEE, 2003).
14. Perez, L. & Wang, J. The effectiveness of data augmentation in image classification using deep learning. Preprint at <https://arxiv.org/abs/1712.04621> (2017).
15. Cui, X., Goel, V. & Kingsbury, B. Data augmentation for deep neural network acoustic modeling. *IEEE/ACM Trans. Audio Speech Lang. Proces.* **23**, 1469–1477 (2015).
16. Geirhos, R. et al. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In *Proc. Seventh Int. Conf. Learning Representations* (ICLR, 2019).
17. Weng, L. Domain randomization for sim2real transfer. *GitHub* <https://lilianweng.github.io/lil-log/2019/05/05/domain-randomization.html> (2019).
18. Tobin, J. et al. Domain randomization for transferring deep neural networks from simulation to the real world. In *IEEE/RSJ Int. Conf. Intelligent Robots and Systems* 23–30 (IEEE, 2017).
19. Tobin, J. Beyond domain randomization. *GitHub* <https://sim2real.github.io> (2019).
20. Sadeghi, F. & Levine, S. CAD2RL: real single-image flight without a single real image. In *Proc. Robotics: Science and Systems Conf.* (RSS, 2017).
21. Tremblay, J. et al. Training deep networks with synthetic data: bridging the reality gap by domain randomization. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition Workshops* 969–977 (IEEE, 2018).
22. Prakash, A. et al. Structured domain randomization: bridging the reality gap by context-aware synthetic data. In *2019 Int. Conf. Robotics and Automation* 7249–7255 (IEEE, 2019).
23. Yu, W., Liu, C. K. & Turk, G. Policy transfer with strategy optimization. In *Proc. Seventh Int. Conf. Learning Representations* (ICLR, 2019).
24. Cubuk, E. D., Zoph, B., Mane, D., Vasudevan, V. & Le, Q. V. Autoaugment: learning augmentation policies from data. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition* (IEEE, 2019).
25. Zakharchov, S., Kehl, W. & Ilic, S. Deceptionnet: network-driven domain randomization. In *Proc. IEEE Int. Conf. Computer Vision* (IEEE, 2019).
26. Schmidhuber, J. Curious model-building control systems. In *Proc. IEEE Int. Joint Conf. Neural Networks* 1458–1463 (IEEE, 1991).
27. Togelius, J. & Schmidhuber, J. An experiment in automatic game design. In *IEEE Symp. Computational Intelligence and Games* 111–118 (IEEE, 2008).
28. Justesen, N. et al. Illuminating generalization in deep reinforcement learning through procedural level generation. In *NeurIPS 2018 Workshop on Deep Reinforcement Learning* (NeurIPS, 2018).
29. Togelius, J., Yannakakis, G. N., Stanley, K. O. & Browne, C. Search-based procedural content generation: a taxonomy and survey. *IEEE Trans. Comput. Intell. AI Games* **3**, 172–186 (2011).
30. Hansen, N. & Ostermeier, A. Completely derandomized self-adaptation in evolution strategies. *Evolut. Comput.* **9**, 159–195 (2001).
31. Storn, R. & Price, K. Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces. *J. Global Optim.* **11**, 341–359 (1997).
32. Browne, C. & Maire, F. Evolutionary game design. *IEEE Trans. Comput. Intell. AI Games* **2**, 1–16 (2010).
33. Cook, M., Colton, S. & Gow, J. The angelina videogame design system—part I. *IEEE Trans. Comput. Intell. AI Games* **9**, 192–203 (2016).
34. Cook, M. & Colton, S. Multi-faceted evolution of simple arcade games. In *2011 IEEE Conf. Computational Intelligence and Games* 289–296 (IEEE, 2011).
35. Nielsen, T. S., Barros, G. A. B., Togelius, J. & Nelson, M. J. Towards generating arcade game rules with VGDL. In *2015 IEEE Conf. Computational Intelligence and Games* 185–192 (IEEE, 2015).
36. Togelius, J. et al. Controllable procedural map generation via multiobjective evolution. *Genet. Program. Evol. Mach.* **14**, 245–277 (2013).
37. Dahlskog, S. & Togelius, J. A multi-level level generator. In *IEEE Conf. Computational Intelligence and Games* 1–8 (IEEE, 2014).
38. Cachia, W., Liapis, A. & Yannakakis, G. N. Multi-level evolution of shooter levels. In *Proc. Eleventh AAAI Conf. Artificial Intelligence and Interactive Digital Entertainment* (AAAI, 2015).
39. Calle, L., Merelo, J. J., Mora-García, A. & García-Valdez, J.-M. Free form evolution for Angry Birds level generation. In *Int. Conf. Applications of Evolutionary Computation* 125–140 (Springer, 2019).
40. Hastings, E. J., Guha, R. K. & Stanley, K. O. Automatic content generation in the galactic arms race video game. *IEEE Trans. Comput. Intell. AI Games* **1**, 245–263 (2009).
41. Pantaleev, A. In search of patterns: disrupting RPG classes through procedural content generation. In *Proc. Third Workshop on Procedural Content Generation in Games* 4 (ACM, 2012).
42. Risi, S., Lehman, J., D'Ambrosio, D. B., R. Hall, R. & Stanley, K. O. Petalz: search-based procedural content generation for the casual gamer. *IEEE Trans. Comput. Intell. AI Games* **8**, 244–255 (2015).
43. Gravina, D., Khalifa, A., Liapis, A., Togelius, J. & Yannakakis, G. N. Procedural content generation through quality diversity. In *2019 IEEE Conf. Games* 1–8 (IEEE, 2019).
44. Cully, A., Clune, J., Tarapore, D. & Mouret, J.-B. Robots that can adapt like animals. *Nature* **521**, 503–507 (2015).
45. Khalifa, A., Green, M. C., Barros, G. & Togelius, J. Intentional computational level design. In *Proc. Genetic and Evolutionary Computation Conf.* 796–803 (ACM, 2019).
46. Mateas, M., Smith, G. & Whitehead, J. Tanagra: reactive planning and constraint solving for mixed-initiative level design. *IEEE Trans. Comput. Intell. AI Games* **3**, 201–215 (2011).
47. Smith, A. M. & Mateas, M. Answer set programming for procedural content generation: a design space approach. *IEEE Trans. Comput. Intell. AI Games* **3**, 187–200 (2011).
48. Goodfellow, I. et al. Generative adversarial nets. In *Advances in Neural Information Processing Systems* 2672–2680 (NeurIPS, 2014).
49. Summerville, A. J. & Mateas, M. Mystical tutor: a magic: the gathering design assistant via denoising sequence-to-sequence learning. In *Twelfth AAAI Conf. Artificial Intelligence and Interactive Digital Entertainment* (AAAI, 2016).
50. Snodgrass, S. & Ontanón, S. Controllable procedural content generation via constrained multi-dimensional markov chain sampling. In *Proc. Twenty-Fifth Int. Joint Conf. Artificial Intelligence* 780–786 (ICJAI, 2016).
51. Bontrager, P., Roy, A., Togelius, J., Memon, N. & Ross, A. Deepmasterprints: generating masterprints for dictionary attacks via latent variable evolution. In *IEEE 9th Int. Conf. Biometrics Theory, Applications and Systems* 1–9 (IEEE, 2018).
52. Volz, V. et al. Evolving mario levels in the latent space of a deep convolutional generative adversarial network. In *Proc. Genetic and Evolutionary Computation Conf.* 221–228 (ACM, 2018).
53. Bontrager, B. & Togelius, J. Fully differentiable procedural content generation through generative playing networks. Preprint at <https://arxiv.org/abs/2002.05259> (2020).
54. Khalifa, A., Bontrager, P., Earle, S. & Togelius, J. PCGRL: procedural content generation via reinforcement learning. Preprint at <https://arxiv.org/abs/2001.09212> (2020).
55. Brant, J. C. & Stanley, K. O. Minimal criterion coevolution: a new approach to open-ended search. In *Proc. Genetic and Evolutionary Computation Conf.* 67–74 (ACM, 2017).
56. Wang, R., Lehman, J., Clune, J. & Stanley, K. O. Paired open-ended trailblazer (poet): endlessly generating increasingly complex and diverse learning environments and their solutions. In *Proc. Genetic and Evolutionary Computation Conf.* 142–151 (ACM, 2019).
57. Zhang, C., Vinyals, O., Munos, R. & Bengio, S. A study on overfitting in deep reinforcement learning. Preprint at <https://arxiv.org/abs/1804.06893> (2018).
58. Ruderman, A. et al. Uncovering surprising behaviors in reinforcement learning via worst-case analysis. In *Safe Machine Learning Workshop at ICLR* (ICLR, 2019).
59. Klimov, O. CarRacing-v0. *OpenAI* <https://gym.openai.com/envs/CarRacing-v0/> (2016).
60. Ha, D. & Schmidhuber, J. Recurrent world models facilitate policy evolution. In *Advances in Neural Information Processing Systems* 2450–2462 (NeurIPS, 2018).
61. Ha, D. Evolving stable strategies. *Otoro* <http://blog.otoro.net/2017/11/12/evolving-stable-strategies/> (2017).
62. Lukasiewicz, T., Song, Y., Wu, L. & Xu, Z. Arena: a general evaluation platform and building toolkit for multi-agent intelligence. In *Proc. 34th National Conf. Artificial Intelligence* (AAAI, 2020).
63. Harries, L., Lee, S., Rzepecki, J., Hofmann, K. & Devlin, S. Mazeexplorer: a customisable 3D benchmark for assessing generalisation in reinforcement learning. In *2019 IEEE Conf. Games* 1–4 (IEEE, 2019).
64. Cobbe, K., Klimov, O., Hesse, C., Kim, T. & Schulman, J. Quantifying generalization in reinforcement learning. In *Proc. Int. Conf. Machine Learning* 1282–1289 (ICML, 2019).

65. Cobbe, K., Hesse, C., Hilton, J. & Schulman, J. Leveraging procedural generation to benchmark reinforcement learning. Preprint at <https://arxiv.org/abs/1912.01588> (2019).
66. Andrychowicz, O. M. et al. Learning dexterous in-hand manipulation. *Int. J. Robot. Res.* **39**, 3–20 (2020).
67. Clune, J. AI-GAs: AI-generating algorithms, an alternate paradigm for producing general artificial intelligence. Preprint at <https://arxiv.org/abs/1905.10985> (2019).
68. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
69. Nelson, M. J. & Mateas, M. Towards automated game design. In *Congress of the Italian Association for Artificial Intelligence* 626–637 (Springer, 2007).
70. Font, J. M., Mahlmann, T., Manrique, D. & Togelius, J. A card game description language. In *European Conf. Applications of Evolutionary Computation* 254–263 (Springer, 2013).
71. Fan, A. et al. Generating interactive worlds with text. In *Proc. AAAI Conf. Artificial Intelligence* 1693–1700 (AAAI, 2020).
72. Friberger, M. G. et al. Data games. In *Proc. Fourth Workshop on Procedural Content Generation in Games* (ACM, 2013).
73. Barros, G. A. B., Liapis, A. & Togelius, J. Playing with data: procedural generation of adventures from open data. In *Proc. First International Joint Conf. DiGRA and FDG* (DiGRA, 2016).
74. Walton, N. *AI Dungeon 2* <https://aidungeon.io/> (2019).
75. Radford, A. et al. Language models are unsupervised multitask learners. *OpenAI Blog* **1**, 9 (2019).
76. Thrun, S. & Mitchell, T. M. Lifelong robot learning. *Robot. Autonom. Sys.* **15**, 25–46 (1995).
77. Parisi, G. I., Kemker, R., Part, J. L., Kanan, C. & Wermter, S. Continual lifelong learning with neural networks: a review. *Neural Netw.* **113**, 54–71 (2019).
78. Soltoggio, A., Stanley, K. O. & Risi, S. Born to learn: the inspiration, progress, and future of evolved plastic artificial neural networks. *Neural Netw.* **108**, 48–67 (2018).
79. Ra, T. S. An evolutionary approach to synthetic biology: Zen and the art of creating life. *Artif. Life* **1**, 179–209 (1993).
80. Adami, C., Brown, C. T. & Kellogg, W. in *Artificial Life IV* (Brooks, R. A. & Maes, P.) 377–381 (MIT Press, 1994).
81. Perez-Liebana, D. et al. General video game AI: a multi-track framework for evaluating agents, games and content generation algorithms. *IEEE Trans. Games* **11**, 195–214 (2019).
82. Juliani, A. et al. Obstacle tower: a generalization challenge in vision, control, and planning. <https://arxiv.org/abs/1902.01378> (2019).
83. Johansen, M., Pichlmair, M. & Risi, S. Video game description language environment for unity machine learning agents. In *2019 IEEE Conf. Games* 1–8 (IEEE, 2019).
84. Nichol, A., Pfau, V., Hesse, C., Klimov, O. & Schulman, J. Gotta learn fast: a new benchmark for generalization in RL. Preprint at <https://arxiv.org/abs/1804.03720> (2018).
85. Leibo, J. Z. et al. Psychlab: a psychology laboratory for deep reinforcement learning agents. Preprint at <https://arxiv.org/abs/1801.08116> (2018).
86. Beyret, B. et al. The animal-AI environment: training and testing animal-like artificial cognition Preprint at <https://arxiv.org/abs/1909.07483> (2019).
87. Ruiz, N., Schuler, S. & Chandraker, M. Learning to simulate. In *Proc. Int. Conf. Learning Representations* (ICLR, 2019).
88. Kar, A. et al. Meta-sim: learning to generate synthetic datasets. In *Proc. Int. Conf. Computer Vision* 4550–4559 (ICCV, 2019).
89. Bayer, J., Wierstra, D., Togelius, J. & Schmidhuber, J. Evolving memory cell structures for sequence learning. In *Int. Conf. Artificial Neural Networks* 755–764 (Springer, 2009).
90. Elsken, T., Metzen, J. H. & Hutter, F. Neural architecture search: a survey. *J. Mach. Learn. Res.* **20**, 55 (2019).
91. Such, F. P., Rawal, A., Lehman, J., Stanley, K. O. & Clune, J. Generative teaching networks: accelerating neural architecture search by learning to generate synthetic training data. In *Proc. Int. Conf. Machine Learning* (ICML, 2020).
92. Togelius, J., De Nardi, R. & Lucas, S. M. Towards automatic personalised content creation for racing games. In *IEEE Symp. Computational Intelligence and Games* 252–259 (IEEE, 2007).
93. Summerville, A. & Mateas, M. Super mario as a string: Platformer level generation via LSTMs. In *Proc. DiGRA/FDG Joint Conf.* (AAAI, 2016).
94. Bontrager, P., Lin, W., Togelius, J. & Risi, S. Deep interactive evolution. In *Int. Conf. Computational Intelligence in Music, Sound, Art and Design* 267–282 (Springer, 2018).
95. Zaltron, N., Zurlo, L. & Risi, S. CG-GAN: an interactive evolutionary GAN-based approach for facial composite generation. In *Thirty-Fourth AAAI Conf. Artificial Intelligence* (AAAI, 2020).
96. Karavolos, D., Liapis, A. & Yannakakis, G. N. A multi-faceted surrogate model for search-based procedural content generation. *IEEE Trans. Games* <https://doi.org/10.1109/TG.2019.2931044> (2019).

Acknowledgements

We would like to thank all the members of modl.ai, especially N. Justesen, for comments on earlier drafts of this manuscript. We would also like to thank A. Wojcicki, R. Canaan, S. Devlin, N. Ruiz, R. Saremi and J. Clune.

Author contributions

Both authors (S.R. and J.T.) contributed equally to the conceptualization and writing of the paper.

Competing interests

The authors declare a potential financial conflict of interest as co-founders of modl.ai ApS.

Additional information

Correspondence should be addressed to S.R.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© Springer Nature Limited 2020