

Capturing Local and Global Patterns in Procedural Content Generation via Machine Learning

Vanessa Volz¹, Niels Justesen¹, Sam Snodgrass¹, Sahar Asadi², Sami Purmonen²,
Christoffer Holmgård¹, Julian Togelius¹, Sebastian Risi¹
¹*modl.ai*, ²*King Digital Entertainment*

Abstract—Recent *procedural content generation via machine learning* (PCGML) methods allow learning from existing content to produce similar content automatically. While these approaches are able to generate content for different games (e.g. Super Mario Bros., DOOM, Zelda, and Kid Icarus), it is an open question how well these approaches can capture large-scale visual patterns such as symmetry. In this paper, we propose match-three games as a domain to test PCGML algorithms regarding their ability to generate suitable patterns. We demonstrate that popular algorithms such as Generative Adversarial Networks struggle in this domain and propose adaptations to improve their performance. In particular, we augment the neighbourhood of a Markov Random Fields approach to take into account not only local but also symmetric positional information. We conduct several empirical tests, including a user study that show the improvements achieved by the proposed modifications and obtain promising results.

Index Terms—procedural content generation, machine learning, global patterns, generative adversarial networks, markov random fields, Candy Crush Saga

I. INTRODUCTION

One approach to generating game content is training a machine learning (ML) model on available examples, and then sampling from the trained model for new game content that is similar to the original examples. This approach, generally referred to as Procedural Content Generation via Machine Learning (PCGML) [1], assumes that the ML algorithm will learn the appropriate invariants of the provided content, so that samples from the learned model retain the “style” of the original while introducing variety. For example, for platform game levels, you would want the model to learn the maximal lengths of gaps it can generate to still obtain playable levels, but to also introduce new challenges and scenarios (e.g. combinations of gaps) not found in the original levels.

Given that PCGML is a relatively recent research topic, it is perhaps not surprising that little research has been done on which invariants of the training contents are learned, and how to make the system learn the most relevant ones. In this paper, we investigate which types of patterns are captured by models trained on a large set of levels, and methods to capture more relevant patterns. We also provide a taxonomy of such methods in order to help structure research in this direction.

To test our methods we use levels from the mobile casual puzzle game Candy Crush Saga (CCS). Casual puzzle games as a genre are well-suited to PCG in general and PCGML in particular, given the huge demand for level content; e.g. more

than six thousand levels have been made for CCS. Just like in many other game genres, the quality of the levels is crucial to keep the player engaged. Interestingly, it is also a game genre for which, to our knowledge, few PCG studies and no PCGML studies exist in the literature.

Our results suggest that existing PCGML methods mostly focus on local patterns and are unable to reproduce global features such as symmetry. However, we find certain methodological advancements such as incorporating symmetric neighbours into Markov Random Fields, improve the methods. While these results are promising, consistently generating levels with the same complex local and global patterns as in the existing CCS levels is an important future research direction.

II. BACKGROUND

A. Patterns in Games

The term *pattern* has different connotations in different domains, but it generally describe regularities within a given object. In the context of frequent pattern mining, this mean sets of items, sub-sequences or sub-structures that occur multiple times. In the context of generation of abstract images, patterns have previously been classified as symmetry, repetition, and repetition with variation [2]. In games, the most prominent use of patterns is in *game design patterns* [3]. Game design patterns can take many shapes, including rather abstract patterns related to the overall game design, which we do not consider here. There are also more fine-grained patterns, including various spatial features observed in games such as first-person shooters [4] and platformers [5]. The type of patterns we address in this paper are fine-grained and visible spatial patterns, in particular relations between tiles.

While humans seem to have an intuitive understanding of what a visual pattern is, it is challenging to provide a general formal definition. In addition, games are ultimately made for humans – a formal definition that does not align with human perception thus would not be meaningful. Here, we choose to mainly rely on a user study to identify patterns accompanied by a simple context-specific description of the concept.

In particular, we focus on Candy Crush Saga (CCS), which is a free-to-play match-three puzzle game released by King¹ in 2012 and has since enjoyed major success as one of the top mobile casual games². In CCS, three or more candies (tiles)

¹<https://king.com/>

²data source: appannie.com

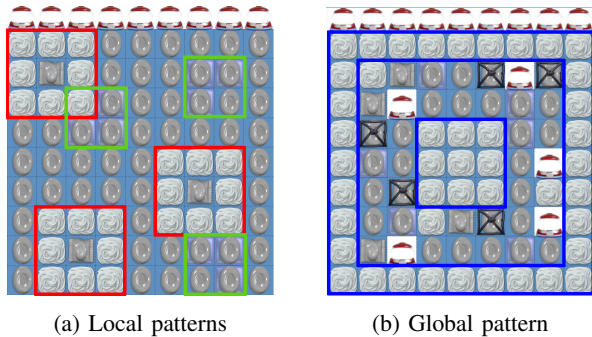


Fig. 1: Illustrative examples of patterns in synthetic CCS levels.

can be matched horizontally or vertically with neighbouring candies of the same colour [6]. When matched, candies disappear from the board. If there are no obstructing items, this causes existing candies to fall from the top and fill the resulting gaps, and new candies with random colours to be spawned (usually at the top of the board). The game introduces various constraints, obstacles and objectives that together define each level and thus create puzzles of varying difficulty. In the context of CCS, we use the terms:

- **Global pattern:** A pattern that can only be identified by looking at spatial structure of **all** elements on the board.
- **Local pattern:** A pattern that can be identified in a **small area** of the board.

CCS levels often exhibit global patterns commonly considered aesthetically pleasing or interesting to the human eye. Examples include placing items on the board in recognisable shapes or in symmetric arrangements (cf. Figure 1b). There are also levels which repeat smaller, local patterns. A common example is that candies with additional beneficial effects are enclosed by obstructing items, making them harder to use (Figure 1a). Levels can display both local and global patterns.

The global aspect of CCS is different from standard applications of PCGML, such as Mario, where local tile neighbourhoods define the primary game elements (gaps, pipes, enemy groups). For this reason, we propose it as a test domain for PCGML algorithms to assess their ability to replicate different types of patterns, including local and global ones.

B. Patterns in PCGML

Algorithms performing Procedural Content Generation via Machine Learning (PCGML) [1] directly employ machine learning models as content generators for games. At its core, PCGML is applied to learn invariant patterns from a set of examples. The type and scale of patterns captured are largely determined by the underlying machine learning approach, the training data, and its representation. The majority of existing PCGML techniques represent the data and patterns as local relationships between small scale elements either sequentially [7]–[10] or as local neighbourhoods [8], [11]–[14]. These approaches have been shown to capture local patterns and small structures from the provided training data (e.g. pipes in *Super Mario Bros.* [7], [12], or pits with water in a generic platform game [11]). However, the ability of these approaches

to reproduce and model larger scale and global patterns is typically not explored. Instead, computational metrics (e.g. linearity and leniency [7], [8]) are used as proxies for global patterns. In this paper, we are interested in exploring the performance of these local models in domains where global patterns are a major aspect of the game.

A variety of extensions to local models have been developed to capture larger scale patterns better. Some have explored the idea of defining patterns explicitly either through clustering of, e.g., 3×3 tile structures from levels [15] or through training models on expert labelled patterns [16]. Others modified the generation step of the models to adhere to defined objectives and constraints [17], [18]. Lastly, some have tried capturing larger patterns by modelling levels at multiple resolutions or employing hierarchical models [15], [19]–[21]. While these extensions capture particular larger patterns or features, they are limited in what they each can capture; each extension was developed and tested with regard to specific patterns or features (e.g., 4×4 tile structures [15], item and enemy distributions [17], playability [18]).

A number of PCGML techniques have been proposed that try to learn from the entirety of a level (or large sections) at once. Lee et al. [22] train Deep Convolutional Neural Networks on full *StarCraft 2* maps to predict resource placement. Their model is trained on full maps but is meant to learn a specific feature and not a model of a full level. Autoencoders have been used for modelling and blending large sections from platforming levels [23]. However, they also rely on level features such as linearity and density as proxies for the global patterns. Shaker et al. [24] present a non-negative matrix factorisation approach that splits the levels into different layers corresponding to level structure and object placements. However, this approach is evaluated with level features as proxies to evaluate whether patterns are reproduced.

III. TOWARDS PATTERN-AWARE PCGML

As detailed above, most existing PCGML algorithms focus on capturing local patterns in games content. However, as we demonstrate in Section V, this does not necessarily ensure generating global patterns we see in typical CCS levels (Figure 3), such as symmetry. In the following, we present an overview of adaptations to improve existing PCGML methods in this regard. Three main approaches are identified, each of which relies on varying degrees of domain knowledge.

A. Enriching the Data

One way to encourage the generation of specific types of patterns is to enrich the data used for training. An **explicit** way of doing so is by labelling each example with the type of global pattern displayed. Conditional GANs [25], for example, are well suited to handle this kind of data. Models can also be trained on each class separately. Both options aim to strengthen the signals around global patterns that are present in the data. We investigate the latter option in this paper (Section IV-D).

A **data-driven** approach to enrich data could be to identify class labels automatically using unsupervised learning. If

limited labelled data is available, a corresponding approach with a classifier that encodes a **learned bias** is also possible.

B. Augmenting the Algorithm

Many of the algorithms discussed in Section II-B only take into account neighbouring patterns. Apart from enriching the training data, one strategy to improve a local model is to augment the algorithm to ensure it focuses on the desired patterns. If domain knowledge is available, this can be done **explicitly** by modifying the structure of the model to detect the desired patterns. For example, to generate outputs with symmetry, we can add positions that should be mirrored at a given position as input to the algorithm. We test this approach using MRFs as a baseline algorithm.

Another approach is to feed measures describing desired features (e.g. a symmetry score) to the algorithm, so that recognising the fitness of an individual based on desired features is facilitated. This can be done e.g. by giving additional inputs to a GAN’s discriminator. While these approaches are likely good at producing content exhibiting the desired features, they are heavily reliant on domain knowledge and the ability to characterise the features numerically.

Instead of relying on domain knowledge, another approach is to ensure that the input at least allows the algorithm to make connections between items at the scale of the desired global pattern. Examples of this **data-driven** approach could be adding a fully-connected layer as the first layer of the discriminator in a GAN. Another promising approach is adding the position of each input (e.g. as coordinates) to the input of a neural network. This has been explored before in the context of CPPNs [26], GANs [27], and their combination [28].

Between these two extremes lie approaches with a **learned bias**. Such a bias can be learned through labelled samples or adversarial training and then given to the model as an additional input. An example would be attention layers for a neural network as proposed in [29].

C. Filtering the Solutions

A third approach is filtering solutions, which is most straightforward if done **explicitly**, but it is conceivable to learn desired patterns and ways to identify them. Filtering can be executed at different times during the training process. Before training would mean creating a representation that only encodes solutions with the desired global patterns. In the case of symmetry, only half of the level could be generated and automatically mirrored to construct the final level. Repairing solutions to adhere to patterns (e.g. through mirroring) is also possible during or after training. A further option for filtering after training is applying a search algorithm to the space of generated content, e.g. latent vector evolution for GANs [12].

IV. EXPERIMENTAL SETUP

In this paper, we evaluate state-of-the-art PCGML methods with regard to how well they are able to capture global and local patterns in a game. We extend this evaluation to several adaptations discussed in Section III. We apply two

popular PCGML techniques, namely Markov Random Fields (MRF) [8] and Generative Adversarial Networks (GAN) [12], [19], [30] to CCS level generation. As argued in Section II-A, CCS levels exhibit several patterns at different scales. This review of the state-of-the-art is not exhaustive and only intended as an illustration of the potential shortcomings of popular methods on games like CCS.

In the following, we describe our experimental setup including the representation we use for CCS levels and the algorithms employed.

A. CCS Level Representation

CCS contains approximately 80 game elements with different characteristics. Some levels rely on unique mechanics or game elements, which makes them difficult to replicate. For this reason, with the help of experts, we selected a subset of published CCS levels that are more homogeneous. We selected only levels from a specific game mode (*Jelly*³) and discarded levels containing complex dynamic elements such as *frogs* and *conveyor belts*, resulting in the 504 levels we used in our experiments. There are still 51 unique items present in the reduced set of levels. Some items can be stacked on the same cell in the board, and we found 789 unique item stacks.

Most state-of-the-art PCGML approaches for tile-based content generation (Section II-B) are not equipped to scale up to this number of different tiles without incurring high computational costs. Thus, we introduce an abstract representation for CCS levels to reduce the representation complexity based on domain knowledge:

- **SHAPE**: Indicates which game board cells are non-void.
- **REGULAR**: The six types of regular candy that can be matched with other candies of the same colour.
- **SPECIAL**: Match-able items with additional effects
- **BLOCK**: Items that obstructs matches by occupying a cell
- **JELLY**: Items indicating cells where matches need to occur to win the level in *Jelly* game mode.
- **LOCK**: Items that obstructs matches by restricting the movement of items in the same cell

With the six categories mentioned above, each level is represented by a matrix with dimensions $9 \times 9 \times 6$. We use binary encoding to represent the occurrence of an item category in a given cell. However, in order to be able to ensure the validity of the generated CCS levels, we introduce the following post-processing method for the GAN-based approaches.

The first four layers cannot coexist in the same cell. The choice for each cell is determined by selecting the layer with the highest value.⁴ However, if none of the values is higher than a threshold (0.5 in this case), the cell is indicated empty. Locks are allowed to be placed only on cells that are not void or empty, and jelly is only placed on cells that are not void.

Additionally, in order to keep the complexity of the level representation low, we add the following mandatory post-processing steps to all levels in this experiment, including

³https://candycrush.fandom.com/wiki/Jelly_levels

⁴In case of ties, the order determines which layer is selected.

original ones: (1) Candies are spawned through candy cannons. Candy cannons are not included in this representation, so they are automatically placed above non-void cells in order to ensure that new candies are dropped. (2) Additional, complex dynamic elements such as portals and special candy cannons are removed from the game to avoid introducing unnecessary complexity. (3) As this paper focuses on visual patterns, we do not include further game meta-data such as the number of available moves in the representation.

B. Algorithms

This section introduces two baseline algorithms and our proposed extensions to them. We chose Markov Random Fields because they learn strictly local patterns, and GANs because they should be able to capture global patterns, as demonstrated in their application to face generation.

Markov Random Fields (MRFs) [31] are undirected graphical models that have been used extensively in texture synthesis and repair [32], [33]. MRFs model relationships between neighbouring positions (e.g., connected nodes in graphs or nearby pixels in textures). We build on previous work on level generation that used MRFs [8], in which the authors represent a level using a grid of tile types corresponding to different level elements (e.g., solid objects, empty space, items, etc.), and treat the four surrounding positions in the grid as the neighbourhood of the MRF. This setup allowed the MRF to learn a conditional probability distribution (CPD) describing the probability of tile types at a position given the tile types of the neighbouring positions (Figure 2a). We replicate this setup for our **local MRF**. To use this model with CCS, the six-layer representation is collapsed into a single-layer representation; the value at a position is the concatenation of values in each layer. In this way, each possible combination of values for the six layers is treated as a distinct tile type. For the **global MRF model**, the neighbourhood of the local MRF is modified by adding the vertically and horizontally symmetric positions relative to the current position to the neighbourhood. Figure 2b shows an example of this modified neighbourhood. By including non-local positions, we hypothesise the MRF will better capture global patterns such as symmetry. Other MRF extensions have tested longer-range [34], [35] and adaptive neighbourhoods [36]; but to our knowledge, this is the first time MRFs have been extended with symmetric neighbours.

Generative Adversarial Networks (GANs) are a class of unsupervised machine learning methods where two neural networks, a generator and a discriminator are trained in an adversarial setting [37]. The generator attempts to generate realistic data examples and the discriminator has to judge whether data examples come from the generator or a dataset of real examples. The Deep Convolutional GAN (DCGAN) [38] applies convolution and deconvolution to the discriminator and generator to enhance the results on image data. We build a DCGAN using **MarioGAN** [12] trained on levels from Super Mario Bros. The input data for MarioGAN was structured as 32×32 patches with one channel for each tile type (air, ground, etc.). In contrast to MarioGAN, we use 3×3 filters instead of




Fig. 2: The neighbours used in our local Markov Random Field approach (left), and in our global MRF approach that includes vertically and horizontally symmetric positions (right). White indicates the current position, blue indicates positions the current is conditioned on, and red indicates independent positions.

4×4 and produce 9×9 patches, which in the case of CCS is the entire level. All other hyper-parameters are kept the same. Each model is trained for 5,000 epochs. We hypothesise that modifying the discriminator to use 9×9 filters would enable it to learn kernels that capture global patterns. Ideally, two streams of convolution could be employed, one with 9×9 filters for global patterns and one with 3×3 filters for local patterns. For simplicity, we employ just one stream of 9×9 filters in the discriminator in the variant we call **GlobalGAN**.

C. Evaluation

Our evaluation focuses on identifying whether visual patterns present in the training data are found in generated levels. In particular, symmetry is evaluated as an example of a recognisable global pattern. We focus on symmetry as it is a common global pattern and is identifiable both visually and computationally. Still, as human perception can diverge from computational measures, we also conduct a user study.

1) *Visual analysis:* As we are targeting aesthetic patterns, a visual inspection of samples of generated levels is suitable in this case. Here, we use it to manually pre-screen promising generators to be evaluated further in the user study.

2) *Computational Measures:* Horizontal, vertical, and diagonal line symmetry (also known as reflection or mirror symmetry) are measured. These measures are chosen to capture level design preferences. Symmetry scores correspond to the ratio of positions on the board that are *identical* to one of their symmetric counterparts. For the diagonal symmetry score, both diagonal symmetry lines are considered at the same time, such that a position has to be identical to just one of the symmetric counterparts to count towards the score. Positions are considered to be identical if all layers are equal (see Section IV-A). Our data set of CCS levels has a median of 100% vertical symmetry and 55.6% horizontal symmetry.

3) *User study:* We conducted a user study among experts to identify and evaluate patterns that are not measurable computationally. In the following, we describe the design of

the questionnaire distributed among King employees.⁵

The participants were presented with a short explanation and several examples of what global and local patterns mean in the context of this study (cf. Section II-A). All levels were visualised using the original CCS sprites (see Figure 3). Based on this definition, participants were asked the questions below:

- 1) Indicate all levels that contain a global pattern (checkboxes next to level images).
- 2) Indicate all levels that contain one or more local patterns (checkboxes next to level images).
- 3) Which level looks less like a CCS level? (pairwise comparison of two level images)
- 4) Questions about experience of the participant in level design and game mechanics (multiple choice.)

D. Research Questions

We specifically want to investigate two types of adaptation suggested in Section III, namely *explicit algorithm augmentation* and *explicit data enrichment*.

RQ1: Does extending the neighbourhood of an MRF improve its ability to produce global patterns? To answer this question, we compare our extended MRF (GlobalMRF) with the standard MRF (LocalMRF).

RQ2: Does training a GAN on only a single type of global pattern improve its ability to produce this pattern? We compare generators trained with the GlobalGAN architecture - one on the full set of levels (GlobalGAN), one on only vertically symmetric ones (GlobalGAN-vert).

We add original levels from CCS as baselines (Original1 and Original2). The original levels are converted to the same level representation as described in Section IV-A in order to ensure comparability. Each generator produced 1,000 levels for visual and computational analysis. For the survey, we select three levels from each dataset, which are the levels with minimal, median and maximal vertical symmetry score (defined in Section IV-C2), respectively. This is intended to ensure that some, but not all levels, adhere to the most common and very obvious type of global pattern, i.e. vertical symmetry.

V. RESULTS

A. Visual Analysis

For each dataset we visually analyse the level generators. Here, we specifically focus on the generators' ability to create vertically symmetric levels. We choose to analyse the best-case scenario here, thus cherry-picking the level with the maximum vertical symmetry score for each set of levels (Figure 3).

Exact local patterns are difficult to detect in all the levels. However, in all generated levels (Figures 3a-3d), locks and blockers appear in groups instead of on their own. This is a local pattern reflected in the original levels as well (Figures 3e-3f). This was expected as all methods have previously been shown to be able to reproduce local patterns (cf. Section II-B).

⁵The full questionnaire can be accessed at <https://forms.gle/nwAyPqckAZiEvBbr7>.

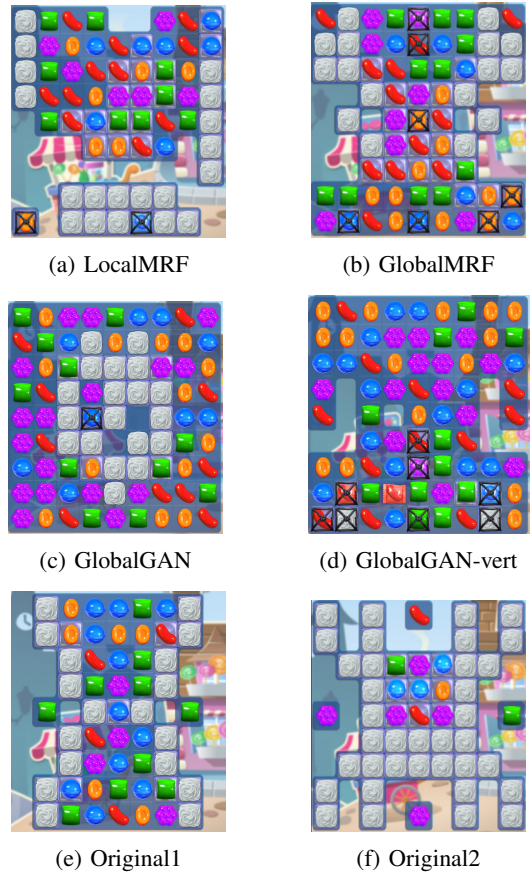


Fig. 3: Visualisation of most vertically symmetric CCS level from each set of levels.

The generated levels (Figures 3a-3d) do show recognisable vertical symmetry. While the addition of information (*RQ1*: GlobalMRF vs. LocalMRF) seems to improve the generator's ability to generate symmetric levels recognisably, the difference seems to be less stark when enriching the data (*RQ2*: GlobalGAN vs. GlobalGAN-vert).

The results suggest that the generators are able to capture some global patterns (namely vertical symmetry), but not consistently. No other global patterns were observed, but it is unclear whether this is due to the small sample size.

B. Computational Measures

As the above analysis is anecdotal, we seek to validate the results using computational measures, i.e. symmetry scores as defined in Section IV-C2. To this end, we generated 1000 levels with each generator and computed their symmetry scores. The resulting scores are shown in Figure 4, together with the symmetry measures computed for the original levels.

Most of the original levels are vertically symmetric, with some apparent variety (Figure 4a). This variety is even bigger across the other symmetry dimensions (Figures 4b and 4c). The generators are able to learn from this data and create individual levels with reasonable symmetry scores. However, their symmetry score distributions are vastly different from the training level - showing much less variety and lower scores

across the board. Adapting PCGML methods so they can also capture the variety of original levels along several dimensions would be an interesting direction for future research, for example, by guiding the sampling of original levels appropriately.

Notice that LocalMRF has very low scores for horizontal and vertical symmetry, as is expected given the strictly local neighbourhood of the model. It is, however, slightly better in terms of diagonal symmetry. This could be a result of a cascading effect caused by sampling with the local neighbourhood, but further investigation is needed.

RQ1: Extending the neighbourhood with symmetric neighbours (GlobalMRF) noticeably increases the symmetry scores along the horizontal and vertical axes. Diagonal information was not added, resulting in a decreased diagonal symmetry score. This highlights a potential problem with this type of adaptation - only global patterns that are explicitly captured will be encouraged with this method, while others (diagonal symmetry in this case) could be negatively affected. We leave further exploration of this for future work.

RQ2: The GlobalGAN shows average scores across the board, and a definite improvement in vertical symmetry is achieved by training only on vertically symmetric levels (GlobalGAN-vert). The vertical scores also vary less, but still quite considerably. However, the score improvement is not major, which may be because most original levels are already vertically symmetric. The performance in other symmetry dimensions is not affected. This suggests that the vertically symmetric levels show sufficient variety to be able to capture other types of symmetry. If the resulting sample was too small or too homogeneous, we would likely see adverse effects. Further work is needed to confirm this conclusion in general.

C. Survey Results

We obtained 60 responses from King employees to our user study. We received answers from people working on various games, of which 63% work on match-three games and almost 40% work on CCS specifically. The participants have different roles in the company, with 14 working in level design.

The original levels have more recognisable global patterns according to the study. This aligns with our results from both visual and computational analysis. The generators are mostly not able to produce global patterns (Figure 5a) except for GlobalMRF which creates levels with a comparatively high vertical symmetry (Figure 4a), reflected also in the study results.

The distinction is far less clear when it comes to local patterns (Figure 5c). Local patterns were recognised much less often than global ones across all levels. The original levels mostly still have an advantage, but GlobalMRF is a close third.

Note, however, that human perception is influenced by visual aspects such as colours. So in some cases, structural patterns may exist, but are not easily recognisable by a human. We tested this by adding levels to the questionnaire that were deliberately constructed to contain a specific pattern. For example, even though a global pattern is clearly identifiable in Figure 6a, it is not very noticeable when visualised as a CCS

level (Figure 6b). As a consequence, only eight participants indicated that they saw a global pattern in the level.

We conclude that humans are not always able to explicitly identify patterns. For the evaluation of our research questions, we thus relied more on pairwise comparisons that only require intuition instead of an explicit detection of patterns.

RQ1: In the pairwise comparison as depicted in Figure 5b, we can clearly see that most respondents considered the LocalMRF levels less likely to be CCS levels when compared with GlobalMRF. In both cases most respondents identify the generated level as dissimilar from CCS levels when compared with an original level. Still, based on this data, some of the respondents could see generated levels be part of the corpus of CCS levels, even when compared with originals. Levels from the augmented algorithm GlobalMRF were preferred more often. We conclude that adapting the MRF algorithm by adding symmetrical inputs improved its ability to create global patterns that make them recognisable as CCS levels.

RQ2: The pairwise comparison (Figure 5d) is also favourable for the adaptation here, GlobalGAN-vert when compared to the baseline algorithm GlobalGAN. The adaptation is detected as an atypical CCS level less often. We thus also deem this adaptation to be beneficial for creating recognisable global patterns.

VI. DISCUSSION

Based on our visual, computational, and survey results, we conclude that the implemented adaptations are indeed beneficial in improving the underlying algorithms' ability to produce recognisable global patterns. However, our main intention in this paper is not to promote a specific adaptation, but to demonstrate the untapped potential for improving existing PCGML methods specifically with regards to their ability to produce global patterns. This is a claim we can make, even if our exact results are not necessarily generalizable across applications and / or PCGML algorithms. We assume that adaptations suitable for GANs should also be suitable for other neural network-based approaches (Variational Autoencoders, Convolutional Neural Networks) and would expect similar results. Sequential models (Markov Chains, Long Short-Term Memory networks) probably require more specific adaptations (like GlobalMRF), because generation and modelling take place in sequence. However, as shown in this paper, encoding domain knowledge can greatly improve the obtained results.

Our observations regarding the fact that GANs seem not to be able to produce symmetric levels consistently are interesting because one of the algorithm's most popular applications is generating faces which usually display imperfect vertical symmetry. We investigated whether a lack of consistent symmetric training signal in CCS could explain this by training on vertically symmetric levels exclusively. While we observed small improvements in this regard (Figure 4a), this explanation evidently does not fully account for the lack of symmetry in the output. There are a number of potential explanations, but one avenue that should be investigated more is how the discretisation of the level representation (Section IV-A) affects




Fig. 4: Boxplots of symmetry scores of each set of levels. All generators are able to create more or less symmetric levels, with the GlobalMRF being able to create vertical and horizontal symmetric levels better than any other approach. The vertical symmetry scores of the GAN approach can be improved slightly by exclusively training on vertical levels.

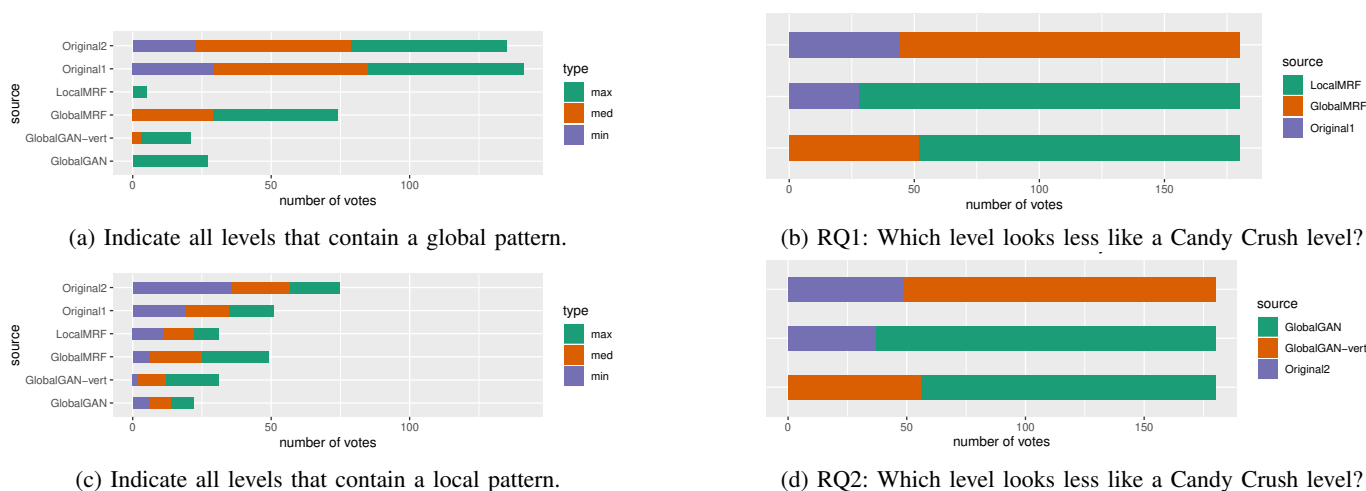


Fig. 5: Questionnaire results. (a, c) Participants were presented with three levels from each method with the minimal, median and maximal vertical symmetry score. GlobalMRF is better than any of the other methods but the original levels still contain more recognizable global and local patterns. The pairwise comparisons between the different methods (b, d) show that levels generated by GlobalMRF are more likely considered CCS levels than any of the other approaches but are still considered to look less like CCS levels than the original levels.

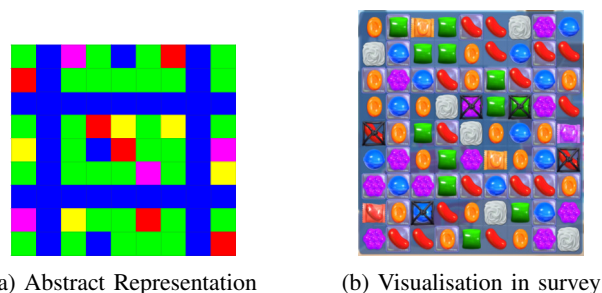


Fig. 6: Validation example from the questionnaire. While specific patterns are clearly visible when shown in colour (a), they become less obvious when shown as game tiles (b).

the output. GANs have been applied to level generation in various tile-based games, all requiring such discretisation.

A more general observation based on our list of potential pattern-aware adaptations to PCGML algorithms in Section III, is that they all require varying levels of domain knowledge. However, even domain experts have difficulty identifying patterns (cf. Section V-C). It is thus important to consider that such required domain knowledge might not always be available when applying a PCGML algorithm. This result also highlights the importance of testing PCG approaches on different domains in order to be able to assess them independently of potential unconscious biases in the training data and / or algorithms. We are looking to confirm our results on other games where (1) non-exact patterns are important (like imperfectly symmetric faces) and that (2) have reasonable complexity so the application of PCGML is warranted.

Further, several of the pattern-aware adaptations mentioned in Section III resemble some techniques from other machine

learning domains. It therefore seems worthwhile to investigate further which other adaptations are made in related fields that operate on data with patterns.

Additionally, it is important to note that evaluating generated content in PCG computationally is still an open challenge. This is especially true if the measure is intended to target something that is not easily formalised and / or relies on human perception. Even visual patterns, as targeted in this paper, proved difficult to evaluate despite the fact that humans do not need to play the game to identify them.

Lastly, this paper focused only on visual patterns, i.e. patterns observable from the representation of the generated content. However, ultimately, the goal of PCGML is to create content that not only looks, but also plays similarly to the given set of levels. It is thus important to study further how visual patterns relate to gameplay patterns in different game genres. We also plan to conduct a deeper investigation on the expressive range of different PCGML methods.

VII. CONCLUSIONS

The results in this paper suggest that match-three games such as CCS provide a challenge for current state-of-the-art PCGML methods since they often include larger scale (global) structural patterns. Several different ways to improve existing methods were outlined, and some investigated further. The implemented adaptations proved to be successful in our experiments, and related research questions should be pursued further. However, significant work remains to be done, especially for generalising, formalising, extending and rigorously testing pattern-aware adaptations of existing PCGML algorithms. In addition, we plan to explore further characterisation and similarity measures to pick up on global and local patterns while minimising the required domain knowledge. It would be interesting to investigate further how the discrete representation of the generated content affects the performance of the respective algorithm, especially for GANs. Lastly, we plan to study whether levels with similar visual patterns also result in similar gameplay.

REFERENCES

- [1] A. Summerville *et al.*, "Procedural Content Generation via Machine Learning (PCGML)," *IEEE Transactions on Games*, vol. 10, no. 3, pp. 257–270, sep 2018.
- [2] K. O. Stanley and R. Miikkulainen, "A taxonomy for artificial embryogeny," *Artificial Life*, vol. 9, no. 2, pp. 93–130, 2003.
- [3] S. Bjork and J. Holopainen, *Patterns in game design*. Charles River Media Hingham, 2005, vol. 11.
- [4] K. Hullett and J. Whitehead, "Design patterns in fps levels," in *Foundations of Digital Games*, 2010, pp. 78–85.
- [5] S. Dahlskog and J. Togelius, "Patterns and procedural content generation: revisiting mario in world 1 level 1," in *First Workshop on Design Patterns in Games*, 2012, pp. 1–8.
- [6] P. Eisen, "Simulating human game play for level difficulty estimation with convolutional neural networks," Master's thesis, KTH, School of Information and Communication Technology (ICT), 2017.
- [7] A. Summerville and M. Mateas, "Super Mario as a string: Platformer level generation via LSTMs," *DiGRA / FDG*, 2016.
- [8] S. Snodgrass and S. Ontanón, "Learning to generate video game maps using markov models," *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 9, no. 4, pp. 410–422, 2016.

- [9] S. Dahlskog, J. Togelius, and M. J. Nelson, "Linear levels through n-grams," *International Academic MindTrek*, 2014.
- [10] A. K. Hoover, J. Togelius, and G. N. Yannakis, "Composing video game levels with music metaphors through functional scaffolding," in *First Computational Creativity and Games Workshop. ACC*, 2015.
- [11] I. Karth and A. M. Smith, "Wavefunctioncollapse is constraint solving in the wild," in *Foundations of Digital Games*, 2017, pp. 1–10.
- [12] V. Volz *et al.*, "Evolving mario levels in the latent space of a deep convolutional generative adversarial network," in *Genetic and Evolutionary Computation Conference*, 2018, pp. 221–228.
- [13] S. M. Lucas and V. Volz, "Tile pattern kl-divergence for analysing and evolving game levels," in *Genetic and Evolutionary Computation Conference*, 2019, pp. 170–178.
- [14] R. Jain *et al.*, "Autoencoders for level generation, repair, and recognition," in *ICCC Workshop: Computational Creativity and Games*, 2016.
- [15] S. Snodgrass and S. Ontanón, "A hierarchical MdMC approach to 2D video game map generation," in *AIIDE*, 2015.
- [16] M. Guzdial *et al.*, "Explainable PCGML via Game Design Patterns," in *AIIDE Workshops*, 2018.
- [17] A. J. Summerville, S. Philip, and M. Mateas, "Mcmcts pcg 4 smb: Monte carlo tree search to guide platformer level generation," in *AIIDE*, 2015.
- [18] S. Snodgrass and S. Ontanón, "Controllable procedural content generation via constrained multi-dimensional markov chain sampling," in *IJCAI*, 2016, pp. 780–786.
- [19] J. Gutierrez and J. Schrum, "Generative adversarial network rooms in generative graph grammar dungeons for the legend of zelda," *arXiv preprint arXiv:2001.05065*, 2020.
- [20] A. Summerville and M. Mateas, "Sampling hyrule: Sampling probabilistic machine learning for level generation," in *AIIDE*, 2015.
- [21] M. Guzdial and M. Riedl, "Game level generation from gameplay videos," in *AIIDE*, 2016.
- [22] S. Lee *et al.*, "Predicting resource locations in game maps using deep convolutional neural networks," in *AIIDE*, 2016.
- [23] A. Sarkar, Z. Yang, and S. Cooper, "Controllable level blending between games using variational autoencoders," in *EXAG Workshop at AIIDE*, 2019.
- [24] N. Shaker and M. Abou-Zleikha, "Alone we can do so little, together we can do so much: A combinatorial approach for generating game content," in *AIIDE*, 2014.
- [25] M. Mirza and S. Osindero, "Conditional generative adversarial nets," *arXiv preprint arXiv:1411.1784*, 2014.
- [26] K. O. Stanley, "Compositional pattern producing networks: A novel abstraction of development," *Genetic Programming and Evolvable Machines*, vol. 8, no. 2, pp. 131–162, 2007.
- [27] R. Liu *et al.*, "An intriguing failing of convolutional neural networks and the coordconv solution," in *Advances in Neural Information Processing Systems 31*, 2018, pp. 9605–9616.
- [28] J. Schrum, V. Volz, and S. Risi, "CPPN2GAN: Combining compositional pattern producing networks and gans for large-scale pattern generation," in *Genetic and Evolutionary Computation Conference*, 2020.
- [29] H. Zhang *et al.*, "Self-attention generative adversarial networks," in *Int. Conference on Machine Learning*, vol. 97, 2019, pp. 7354–7363.
- [30] R. R. Torrado *et al.*, "Bootstrapping conditional gans for video game level generation," *arXiv preprint arXiv:1910.01603*, 2019.
- [31] P. Clifford, "Markov random fields in statistics," *Disorder in physical systems: A volume in honour of John M. Hammersley*, vol. 19, 1990.
- [32] E. Levina, P. J. Bickel *et al.*, "Texture synthesis and nonparametric resampling of random fields," *The Annals of Statistics*, vol. 34, no. 4, pp. 1751–1773, 2006.
- [33] M. Hassner and J. Sklansky, "The use of markov random fields as models of texture," in *Image Modeling*. Elsevier, 1981, pp. 185–198.
- [34] Y. Li and D. P. Huttenlocher, "Sparse long-range random field and its application to image denoising," in *European Conference on Computer Vision*. Springer, 2008, pp. 344–357.
- [35] J. Sun and M. F. Tappen, "Learning non-local range markov random field for image restoration," in *CVPR 2011*. IEEE, 2011, pp. 2745–2752.
- [36] S. Suwanwimolkul *et al.*, "An adaptive markov random field for structured compressive sensing," *IEEE Transactions on Image Processing*, vol. 28, no. 3, pp. 1556–1570, 2018.
- [37] I. Goodfellow *et al.*, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [38] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," in *ICLR*, 2016.