

Co-creating Platformer Levels with Constrained Adversarial Networks

Paolo Morettin^{a,b}, Andrea Passerini^b and Stefano Teso^b

^aKU Leuven, Celestijnenlaan 200 A box 2402, 3001 Leuven, Belgium

^bUniversity of Trento, Via Sommarive 9, 38123 Povo, Trento, Italy

Abstract

Given the success of deep generative models in many creative tasks, it is natural to ask how to best leverage them to support human designers. We study this problem in the context of mixed-initiative design of platformer levels, a paradigmatic co-creative task. This setting is especially challenging, because – like all functional content – platformer levels must satisfy complex validity constraints, like coherency and playability. We explore mixed-initiative interaction with constrained adversarial networks (CANs), a class of deep generative models that synthesize structures satisfying one or more validity constraints. As such, CANs can be used to complete user-supplied partial levels while retaining full control of the constraints to be applied. We go one step beyond, and consider the issue of *customizing* a pre-trained CAN to some target design task at hand and to the designer’s preferences. We discuss how to achieve this by combining CANs with coactive learning, a very natural mixed-initiate interaction protocol that acquires the necessary supervision from the designer in a transparent manner. Finally, we illustrate how to extend coactive learning to acquire informative supervision in the form of interpretable constraints.

Keywords

mixed-initiative interaction, generative adversarial networks, functional content, constraints

1. Introduction

Deep generative models have achieved remarkable results in image synthesis [1, 2], music generation [3], style transfer [4], and many other creative tasks. It is only natural to ask how these models, which are designed to work in full autonomy, can be adapted for assisting and complementing human designers.

To ground the discussion, we focus on a concrete application, namely mixed-initiative design of 2-D platformer game levels [5, 6, 7].

This and other forms of semi-autonomated content generation are very relevant for the game industry, as they promise to lower production costs and facilitate scalability [8]. A key feature of this problem is that – in contrast to, e.g., natural images – game levels are *functional* [7], that is, they obey complex validity constraints like coherence (for instance props must be complete, non-flying enemies must touch the ground) and playability (the goal tile must be reachable), see Figure 1. Standard deep generative models, however, do not handle validity constraints and hence struggle to generate valid objects.

This motivates us to explore mixed-initiative level design with *constrained adversarial networks* (CANs) [9]. CANs are generative adversarial networks (GANs) [1] designed specifically for functional objects. Like other generative models, once trained on existing content, CANs can be invoked to create novel

Joint Proceedings of the ACM IUI 2021 Workshops, April 13-17, 2021, College Station, USA

✉ paolo.morettin@cs.kuleuven.be (P. Morettin);
andrea.passerini@unitn.it (A. Passerini);
stefano.teso@unitn.it (S. Teso)

🆔 0000-0003-4321-5215 (P. Morettin);
0000-0002-2765-5395 (A. Passerini);
0000-0002-2340-9461 (S. Teso)

© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings
(CEUR-WS.org)






Figure 1: Example Mario levels generated by a constrained adversarial network. Left: Pipe props are constrained to be complete and symmetric (coherency). Right: The rightmost column must be reachable from the starting position, regardless of gaps (playability). Images taken with permission from [9].

content from scratch [1] or to finish incomplete sketches supplied by a designer [10]. As a bonus, during interaction the designer can turn on and off different constraints, allowing her to synthesize different types of structures and boost diversity [9].

In this setup the model is trained to generate content that mimics the training data, which might be quite different from the content needed by the designer for the task at hand. Compatibly with recent work [11, 12], we argue that effective content synthesis requires *customizing* the generative model for the target domain. To this end, we propose to leverage *coactive learning* (CL) [13], an interaction protocol whereby the machine acquires pairwise preferences (that is, “object x is better than object x' ”) by tracking the modifications made by the designer to its own suggestions. The generative model is then periodically adjusted to comply with the collected preferences. Coactive learning integrates seamlessly into mixed-initiative workflows and is completely transparent to the designer [14, 15].

Given the tension between the amount of feedback that can reasonably be acquired during a design session and the number of preferences needed to adjust deep generative models like CANs [12], we also discuss how to extend CL to collect both preferences and con-

straints. In our view, constraints are a powerful form of feedback as they provide information about entire *sets* of candidate levels: levels that *do* satisfy a constraint are preferred to levels that *do not*, all else being equal. Importantly, constraints can be imposed on top of both outputs (e.g., tile and enemy patterns) and disentangled latent representations of the generative models.

In the following, we overview GANs and CANs. Next, we discuss mixed-initiative level design with CANs and distinguish between strategies based on conditional generation and adaptive strategies based on coactive learning and constraint acquisition. In Section 4 we discuss related work and then conclude with some final remarks.

2. Background

Generative Adversarial Networks (GANs) [1] are a popular class of generative models where two neural networks, the *generator* g and the *discriminator* d , are trained jointly in an adversarial fashion. Specifically, the discriminator d is trained to distinguish “real” objects in the training set \mathcal{X} from “fake” objects synthesized by the generator g . At the same time, g is trained to output objects that fool d . The generator is typically implemented as

a feed-forward or a deconvolutional neural network that maps random vectors \mathbf{z} (sampled from some simple distribution, like a multivariate normal) into objects \mathbf{x} . The discriminator is a feed-forward or convolutional neural network that takes objects \mathbf{x} and outputs their probability of being real or fake.

Training is typically performed using (variants of) stochastic gradient descent by interleaving updates to g and d , and it proceeds until the synthesized objects look indistinguishable from those in the training set \mathcal{X} . Under idealized assumptions [1], the learned generator eventually recovers the data distribution. Once trained, the generator g can be used to synthesize *new* objects \mathbf{x} by simply sampling random vectors \mathbf{z} and computing their image under the generator, that is, $\mathbf{x} = g(\mathbf{z})$. This operation is extremely efficient. This form of training has the advantage of not requiring to compute an objective based on likelihood or other probabilistic metrics, which are in general intractable and thus approximated. For this reason, GANs have found successful applications in domains where formalizing the target distribution is hard.

GANs and other standard deep generative models are not designed to generate valid structures [9]. Indeed, generating objects consistent with a constraint involves first discovering that such a constraint exists (at least implicitly) by looking at the data, and generative models are not designed to carry out this non-trivial task. Furthermore, if the training data contains infeasible examples (because of, e.g., mistakes during the data collection), then the validity constraint cannot be acquired perfectly even in principle.¹

¹It can be shown that in this case GANs will acquire a wrong validity constraint [9].

2.1. Constrained Adversarial Networks

Constrained Adversarial Networks (CANs) address this issue by taking both training examples and validity constraints into consideration. CANs train the generator g so that it jointly maximizes the probability of fooling the discriminator d and the probability of synthesizing valid objects. This is achieved by augmenting the adversarial loss used in standard GANs (denoted by ℓ_{adv}) with the *semantic loss* (SL), a technique proposed in [16] to encourage neural networks to output predictions consistent with constraints. Letting ψ be a user-supplied validity constraint (encoded as a propositional logic formula over the elements of \mathbf{x}), CANs train g and d to optimize (resp. minimize and maximize) the following expression:

$$\ell_{adv}(g, d) + \lambda \cdot SL_{\psi}(g) \quad (1)$$

where $\lambda > 0$ is a hyper-parameter controlling the importance of the constraint. The SL is designed to be large whenever the probability that the generator outputs a valid object is small. To this end, the SL is defined as the negative logarithm of this probability:

$$SL_{\psi}(g) = -\log \left[\sum_{\mathbf{x} \text{ satisfies } \psi} P_g(\mathbf{x}) \right] \quad (2)$$

The sum measures the total probability allocated by the generator to valid objects. Evaluating the sum involves enumerating *all* possible objects (e.g., tile arrangements) \mathbf{x} , checking which ones are feasible, and computing their probability with respect to the generator (written $P_g(\mathbf{x})$). Since the number of possible objects is typically exponential, naïve enumeration is infeasible. Knowledge compilation (KC) [17] is thus used to compile the sum into a *compact* polynomial (or more precisely, an arithmetic circuit) that can be evaluated efficiently during training. KC works

by leveraging distributivity to rewrite $SL_\psi(g)$ as compactly as possible, enormously speeding up evaluation. This is achieved by identifying shared sub-components and compactly representing the factorized expression using a DAG.

If ψ is very complex, the polynomial output by KC may be large. This is not a huge issue during training, which is performed once (or infrequently) on powerful machines, but it could be problematic for inference. A major advantage of CANs is that – as in regular GANs – synthesizing new objects boils down to a simple forward pass over the generator, independently of the KC polynomial, which can thus be thrown away after training.

3. Mixed-initiative Level Design with CANs

We are concerned with mixed-initiative design of 2-D platformer levels. In this setting, a human designer and a generative model take turns in proposing modifications to a shared level map [5, 6, 7]. The designer can reject or modify any suggestions made by the model. The process can be considered a success if the machine proposes useful suggestions (that is, modifications that improve the overall quality of the level) thus saving time and resources, or if it inspires the human designer to create an overall better level [18, 19].

Co-designing with a pre-trained model

Perhaps the most straightforward co-design strategy is to use a pre-trained generative model to synthesize levels *conditioned* on the user’s input. The simplest setup is *inpainting*: in this case, the user supplies an incomplete level and the network fills in the missing parts based on the context [10]. CANs, in particular, can be easily adapted to inpainting under constraints [9]. In a second, orthogonal approach,

the designer specifies some desired features of the output and lets the network generate a level compatible with those features. For instance, one could condition the generator to output objects from a specific class (e.g., underground vs. open air levels) [20], to include a given quantity of certain elements (enemies, amount of water) [21], or to satisfy expected length/playtime or leniency (a proxy of expected difficulty) requirements [22].

CANs support an additional form of conditional generation that makes it possible during synthesis to enable or disable different groups of constraints. A technical description of this mechanism can be found in [9]. Hence CANs not only inherit the conditional capabilities of traditional GANs, but also support the generation of objects conditioned on properties expressed in logical terms. Using this technique, the designer can generate diverse levels by turning on and off constraints like “*a room contains a boss fight if and only if it contains a treasure chest*” or “*if the exit is locked then the room must contain a key somewhere*”. The requirement is that these optional constraints are all baked into the CAN generator during training.

In complex co-creative scenarios like level generation, the output of the model is unlikely to satisfy all of the designer’s desiderata from the get go. To solve this issue, recent work has investigated strategies that allow the designer to iteratively refine an initial suggestion by interacting with the generative model [23, 24, 25]. In this sequential setting, a crucial aspect is how to ensure consistency of the generated object with respect to previous iterations. This is particularly challenging when the designer’s feedback is expressed in natural language and the output is relatively unstructured [26]. Since game levels are inherently structured it is easier to unambiguously describe the desired changes, to maintain consistency with respect to previous iterations, and to identify possibly con-

tradictory feedback from the user. For instance, it is much easier to effectively account for feedback like “*turn all the water tiles into lava*” or “*the first half of the level should not contain enemies*” than “*the subject in the photo should smile*”.

Co-designing and customization So far, we considered interaction based on conditioning a pre-trained generative model during synthesis. Since the model generates content that mimics the training data, unless the training data is designed appropriately, the synthesized content will not fit the target application. This introduces a fundamental tension between the effort required to create the training data and the effort saved by using the generative model [12].

A sensible option is then to collect feedback on the model’s suggestions while interacting with the designer, and then to incorporate the latter into the generative model itself [15, 14, 11]. We propose to do so by leveraging *coactive learning* [13, 14], an interactive learning protocol in which the machine iteratively suggests content to the user, the user improves – even by changing just a few tiles – the suggested content, and the process repeats. The meaning of “improvement” is entirely defined by the designer’s preferences and requirements. At the end of each iteration, the machine has access to a suggested object x and an improved configuration x' , and extracts a pairwise preference of the form “ x' is better than x ”. Such preferences are collected in a data set and then used to adjust the generative model. A simple procedure to do this is to re-train the current generative model to better comply to the collected preferences using a ranking loss, as proposed in [27]. As interaction proceeds, more preferences are collected and the model progressively aligns to the designer’s needs. Notice that coactive learning is completely transparent to the user and integrates natu-

rally in mixed-initiative interaction via manipulative interfaces [14], which are commonplace in level co-design [5, 6, 7].

The nature of the feedback highly impacts the performance of the model. Preferences are very effective for capturing the designer’s needs [28], however they do not explain *why* one object is preferred to the other. As a concrete example, if the user is presented with a Super Mario level x and improves it by removing some coins from it, leading to x' , the system does not know whether the coins were removed because they were unreachable, because there were too abundant already, because they did not look good, etc. This kind of explanatory supervision conveys a lot of information [29] and can dramatically improve the speed of adaptation, especially for data hungry models like GANs. We propose to do so by combining constraint acquisition [30] with CANs. The human designer not only can adjust the generative model by modifying the presented output, but is also empowered with tools for specifying constraints on the desired output. Of course, encoding preferences in formal/logical terms can be hard for the end user. Nonetheless, there exists a body of work on translating natural language into various formal representations [31, 32, 33]. Most crucially, the system must be able to learn the building blocks of this interaction language, that is, the logical predicates that the designer uses to specify her needs. While many predicates of practical interest can be specified a priori, like the presence or number of specific elements in a portion of the level, we can go further and provide an interface for learning new predicates from the user. As shown in [9], it is indeed possible to use learned (neural) predicates in the CAN framework, opening up the possibility of extending the base language with new high-level concepts. For instance, the designer may want to specify that a portion of the level referenced by its coordinates represents a castle,

isCastle(x_1, y_1, x_2, y_2). Even if the concept is unknown to the system, it can be learned from user examples and possibly refined with interactions where the user is presented with some positive examples and labels whether they represent a castle or not.

CANs readily provide an interface for learning from such constraints, hence once collected these constraints can simply be used to re-train the CAN until it complies with them. It is to be expected that very complex constraints require an expensive compilation step and many re-training epochs. We admit for adaptation to occur in the background or during periodic sleep cycles, rather than *during* the design session as in CL. Delayed learning of this kind is completely sound from a machine learning perspective [34] and widely adopted in settings characterized by computationally demanding *concept drifts* like anomaly detection [35] and some reinforcement learning applications [36, 37].

4. Related Work

Procedural (game) content generation (PCG) has been traditionally addressed with search-based or rule-based methods. In recent years, the focus has shifted toward techniques based on machine learning [38, 8] and deep learning [7]. Game content can be divided into *functional* content, like game levels, game mechanics or behavioural rules for non-player characters, and *cosmetic* content, such as textures, music and sound effects. Functional content is arguably more challenging for automatic generators. While tasks in computer vision or natural language processing are supported by very large and accessible datasets, the training data for the generation of functional content for games is typically very scarce. Game level generation is by far the most studied problem in this area, with most works focusing on 2D tile-based video games. Although

annotated datasets of levels exist [39], they are relatively small in size and cover a handful of popular games. The available data is often insufficient for effectively training a fully autonomous deep generative model and, most crucially, different approaches must be adopted for training a generator for a novel game with no data available. While some works mitigate this problem with bootstrapping techniques [40] or by allowing for diverse sources of supervision (like gameplay videos [41] or transfer learning across different games [42, 43]), involving human interaction in the training of these systems is likely a necessity for all but the most trivial settings.

Automated creation tools are unlikely to output content that fits the target application perfectly. Successful applications of computer-aided design (CAD) tools have spurred research in developing mixed-initiative co-creation techniques for level generation. This interaction paradigm is not only deemed effective for reaching satisfiable end results, but it is also useful in fostering the creativity of the human designer [18].

Coactive learning was first proposed in the context of interface optimization [44] and information retrieval [13]. A useful feature of CL is that it integrates seamlessly with mixed-initiative interaction: preferences are extracted whenever the human supervisor modifies, explicitly or implicitly, any suggestion made by the machine. For this reason, CL was adopted in constructive preference elicitation with manipulative interaction [14]. Our proposed approach is directly inspired by this line of work.

Recently, Guzdial et al. [11] introduced an interaction protocol reminiscent of CL. The difference is that, whereas in CL the model used to synthesize structures and the model used to learn from the designer's feedback are the very same, Guzdial et al. allow the two models to be different [11]. From a learning perspective this is sub-optimal, as the two models might make different mistakes and have

different biases, and therefore feedback useful for one model might be less (than) useful for the other. Finally, our idea of integrating interpretable constraint acquisition into coactive learning extends prior work in constructive preference elicitation [15] by combining it with constraint learning [30] and explanatory interactive learning [29, 45].

5. Conclusion

We discussed mixed-initiative level design with constrained adversarial networks. Our contribution is conceptual: on the one hand, we show that these models – which are designed for autonomously generating functional content – can be used for helping human designers, and on the other that they can be adapted to the task at hand by combining them with coactive learning and constraint acquisition. Of course, these insights must be validated through extensive experiments and user studies. This is left to future work.

Acknowledgments

This work has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. [694980] SYNTH: Synthesising Inductive Data Models). The research of ST and AP was partially supported by TAILOR, a project funded by EU Horizon 2020 research and innovation programme under GA No 952215.

References

- [1] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial nets, in: *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [2] A. Odena, C. Olah, J. Shlens, Conditional image synthesis with auxiliary classifier gans, in: *International conference on machine learning*, PMLR, 2017, pp. 2642–2651.
- [3] K. Kumar, R. Kumar, T. de Boissiere, L. Gestein, W. Z. Teoh, J. Sotelo, A. de Brébisson, Y. Bengio, A. C. Courville, Melgan: Generative adversarial networks for conditional waveform synthesis, *Advances in Neural Information Processing Systems* 32 (2019) 14910–14921.
- [4] J.-Y. Zhu, T. Park, P. Isola, A. A. Efros, Unpaired image-to-image translation using cycle-consistent adversarial networks, in: *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2223–2232.
- [5] G. Smith, J. Whitehead, M. Mateas, Tanagra: A mixed-initiative level design tool, in: *Proceedings of the Fifth International Conference on the Foundations of Digital Games*, 2010, pp. 209–216.
- [6] A. Liapis, G. N. Yannakakis, J. Togelius, Sentient sketchbook: computer-assisted game level authoring (2013).
- [7] J. Liu, S. Snodgrass, A. Khalifa, S. Risi, G. N. Yannakakis, J. Togelius, Deep learning for procedural content generation, *Neural Computing and Applications* (2020) 1–19.
- [8] A. Summerville, S. Snodgrass, M. Guzdial, C. Holmgård, A. K. Hoover, A. Isaksen, A. Nealen, J. Togelius, Procedural Content Generation via Machine Learning (PCGML), *IEEE Transactions on Games* 10 (2018) 257–270.
- [9] L. Di Liello, P. Ardino, J. Gobbi, P. Morettin, S. Teso, A. Passerini, Efficient generation of structured objects with constrained adversarial networks,

- Advances in Neural Information Processing Systems 33 (2020).
- [10] J. Yu, Z. Lin, J. Yang, X. Shen, X. Lu, T. S. Huang, Generative image inpainting with contextual attention, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 5505–5514.
- [11] M. Guzdial, N. Liao, M. Riedl, Co-creative level design via machine learning, arXiv preprint arXiv:1809.09420 (2018).
- [12] I. Karth, A. M. Smith, Addressing the fundamental tension of pcgml with discriminative learning, in: Proceedings of the 14th International Conference on the Foundations of Digital Games, 2019, pp. 1–9.
- [13] P. Shivaswamy, T. Joachims, Coactive learning, Journal of Artificial Intelligence Research 53 (2015) 1–40.
- [14] P. Dragone, S. Teso, A. Passerini, Constructive preference elicitation, Frontiers in Robotics and AI 4 (2018) 71.
- [15] S. Teso, P. Dragone, A. Passerini, Coactive critiquing: Elicitation of preferences and features, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 31, 2017.
- [16] J. Xu, Z. Zhang, T. Friedman, Y. Liang, G. Broeck, A semantic loss function for deep learning with symbolic knowledge, in: International Conference on Machine Learning, 2018, pp. 5498–5507.
- [17] A. Darwiche, P. Marquis, A knowledge compilation map, Journal of Artificial Intelligence Research 17 (2002) 229–264.
- [18] G. N. Yannakakis, A. Liapis, C. Alexopoulos, Mixed-initiative co-creativity (2014).
- [19] P. Karimi, J. Rezwana, S. Siddiqui, M. L. Maher, N. Dehbozorgi, Creative sketching partner: an analysis of human-ai co-creativity, in: Proceedings of the 25th International Conference on Intelligent User Interfaces, 2020, pp. 221–230.
- [20] M. Mirza, S. Osindero, Conditional generative adversarial nets, arXiv preprint arXiv:1411.1784 (2014).
- [21] A. Hald, J. S. Hansen, J. Kristensen, P. Burelli, Procedural content generation of puzzle games using conditional generative adversarial networks, in: International Conference on the Foundations of Digital Games, 2020, pp. 1–9.
- [22] J. Schrum, V. Volz, S. Risi, Cppn2gan: Combining compositional pattern producing networks and gans for large-scale pattern generation, arXiv preprint arXiv:2004.01703 (2020).
- [23] A. El-Nouby, S. Sharma, H. Schulz, D. Hjelm, L. El Asri, S. Ebrahimi Kahou, Y. Bengio, G. W. Taylor, Keep drawing it: Iterative language-based image generation and editing, in: Neural Information Processing Systems: Visually Grounded Interaction and Language Workshop, 2018.
- [24] A. El-Nouby, S. Sharma, H. Schulz, D. Hjelm, L. E. Asri, S. E. Kahou, Y. Bengio, G. W. Taylor, Tell, draw, and repeat: Generating and modifying images based on continual linguistic instruction, in: Proceedings of the IEEE International Conference on Computer Vision, 2019, pp. 10304–10312.
- [25] Y. Cheng, Z. Gan, Y. Li, J. Liu, J. Gao, Sequential attention gan for interactive image editing, in: Proceedings of the 28th ACM International Conference on Multimedia, 2020, pp. 4383–4391.
- [26] Y. Liu, M. De Nadai, D. Cai, H. Li, X. Alameda-Pineda, N. Sebe, B. Lepri, Describe what to change: A text-guided unsupervised image-to-image translation approach, in: Proceedings of the 28th ACM International Conference on Multimedia, 2020, pp. 1357–1365.
- [27] E. Heim, Constrained generative ad-

- versarial networks for interactive image generation, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10753–10761.
- [28] G. Pigozzi, A. Tsoukias, P. Viappiani, Preferences in artificial intelligence, *Annals of Mathematics and Artificial Intelligence* 77 (2016) 361–401.
- [29] P. Schramowski, W. Stammer, S. Teso, A. Brugger, F. Herbert, X. Shao, H.-G. Luigs, A.-K. Mahlein, K. Kersting, Making deep neural networks right for the right scientific reasons by interacting with their explanations, *Nature Machine Intelligence* 2 (2020) 476–486.
- [30] L. De Raedt, A. Passerini, S. Teso, Learning constraints from examples, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [31] J. Dzifcak, M. Scheutz, C. Baral, P. Schermerhorn, What to do and how to do it: Translating natural language directives into temporal and dynamic logic representation for goal management and action execution, in: *2009 IEEE International Conference on Robotics and Automation*, IEEE, 2009, pp. 4163–4168.
- [32] S. Tellex, T. Kollar, S. Dickerson, M. Walter, A. Banerjee, S. Teller, N. Roy, Understanding natural language commands for robotic navigation and mobile manipulation, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 25, 2011.
- [33] A. Brunello, A. Montanari, M. Reynolds, Synthesis of ltl formulas from natural language texts: State of the art and research directions, in: *26th International Symposium on Temporal Representation and Reasoning (TIME 2019)*, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- [34] M. Zinkevich, J. Langford, A. Smola, Slow learners are fast, *Advances in neural information processing systems* 22 (2009) 2331–2339.
- [35] M. Odiathevar, W. K. Seah, M. Freat, A hybrid online offline system for network anomaly detection, in: *2019 28th International Conference on Computer Communication and Networks (ICCCN)*, IEEE, 2019, pp. 1–9.
- [36] M. Moravčík, M. Schmid, N. Burch, V. Lisý, D. Morrill, N. Bard, T. Davis, K. Waugh, M. Johanson, M. Bowling, Deepstack: Expert-level artificial intelligence in heads-up no-limit poker, *Science* 356 (2017) 508–513.
- [37] N. Brown, T. Sandholm, Superhuman ai for heads-up no-limit poker: Libratus beats top professionals, *Science* 359 (2018) 418–424.
- [38] M. Hendrikx, S. Meijer, J. Van Der Velden, A. Iosup, Procedural content generation for games: A survey, *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 9 (2013) 1–22.
- [39] A. J. Summerville, S. Snodgrass, M. Mateas, S. O. n'on Villar, The VGLC: The Video Game Level Corpus, *Proceedings of the 7th Workshop on Procedural Content Generation* (2016).
- [40] R. R. Torrado, A. Khalifa, M. C. Green, N. Justesen, S. Risi, J. Togelius, Bootstrapping conditional gans for video game level generation, in: *2020 IEEE Conference on Games (CoG)*, IEEE, 2020, pp. 41–48.
- [41] M. Guzdial, M. O. Riedl, Game level generation from gameplay videos., in: *AI-IDE*, 2016, pp. 44–50.
- [42] S. Snodgrass, S. Ontanon, An approach to domain transfer in procedural content generation of two-dimensional videogame levels, in: *Twelfth Artificial Intelligence and Interactive Digital Entertainment Conference*, 2016.

- [43] A. Sarkar, S. Cooper, Blending levels from different games using lstms., in: *AIIDE Workshops*, 2018.
- [44] K. Gajos, D. S. Weld, Preference elicitation for interface optimization, in: *Proceedings of the 18th annual ACM symposium on User interface software and technology*, 2005, pp. 173–182.
- [45] W. Stammer, P. Schramowski, K. Kersting, Right for the right concept: Revising neuro-symbolic concepts by interacting with their explanations, *arXiv preprint arXiv:2011.12854* (2020).