

Towards Friendly Mixed Initiative Procedural Content Generation: Three Pillars of Industry

Gorm Lai
Goldsmiths, University of London
New Cross, London, United Kingdom
glai001@gold.ac.uk

William Latham
Goldsmiths, University of London
New Cross, London, United Kingdom
w.latham@gold.ac.uk

Frederic Fol Leymarie
Goldsmiths, University of London
New Cross, London, United Kingdom
ffl@gold.ac.uk

ABSTRACT

While the games industry is moving towards procedural content generation (PCG) with tools available under popular platforms such as Unreal, Unity or Houdini, and video game titles like No Man's Sky and Horizon Zero Dawn taking advantage of PCG, the gap between academia and industry is as wide as it has ever been, in terms of communication and sharing methods. The authors have worked on both sides of this gap and in an effort to shorten it and increase the synergy between the two sectors have identified three design pillars for PCG using mixed-initiative interfaces. The three pillars are *respect designer control*, *respect the creative process* and *respect existing work processes*. *Respecting designer control* is about creating a tool that gives enough control to bring out the designer's vision. *Respecting the creative process* concerns itself with having a feedback loop that is short enough, that the creative process is not disturbed. *Respecting existing work processes* means that a PCG tool should plug in easily to existing asset pipelines. As academics and communicators, it is surprising that publications often do not describe ways for developers to use our work or lack considerations for how a piece of work might fit into existing content pipelines.

CCS CONCEPTS

• **Human-centered computing** → *HCI design and evaluation methods; Interactive systems and tools*; • **Software and its engineering** → *Interactive games*;

KEYWORDS

game design, procedural content, procedural content generation, games industry, mixed-initiative, co-creative agent, AI assisted game design

ACM Reference Format:

Gorm Lai, William Latham, and Frederic Fol Leymarie. 2020. Towards Friendly Mixed Initiative Procedural Content Generation: Three Pillars of Industry. In *International Conference on the Foundations of Digital Games (FDG '20)*, September 15–18, 2020, Bugibba, Malta. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3402942.3402946>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

FDG '20, September 15–18, 2020, Bugibba, Malta

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8807-8/20/09...\$15.00

<https://doi.org/10.1145/3402942.3402946>

1 INTRODUCTION

In recent years, both academia and the games industry have had an increased focus on PCG. However, except for isolated attempts such as Natural Motion [26], we need more effort in bringing academic research to use as inspirational or developmental material for the games industry. This is an opinion also expressed in Shaker et al. [29] in their interview with Andrew Doull, as he states: "There's a lot of interesting stuff happening on the academic side - getting this to percolate over to game development is going to be the real challenge." In game AI Togelius [35], and from the wider tech-industry Kaczmarczyk [20], have both lamented the lack of collaboration and understanding between academia and the industry. With over 30 years of combined work in the games industry, from indie-sized to AAA productions, the authors of this paper have the same experience.

The main contributions of this work are to distill such issues into three design pillars for creating tools for mixed-initiative procedural content generation (MI-PCG): *Respect designer control*, *respect the creative process* and *respect existing workflow*.

Respect designer control focuses on empowering the designer to be able to get their vision out. It asks the question, "does the algorithm provide enough control for the designer to express their vision?" *Respect the creative process* concerns itself with providing a short iterative loop that provides enough feedback to the user that it does not break their creative process. The last pillar, *respect existing work processes* focuses on the ease of embedding PCG tools into an organisation with an existing workflow that already combines a number of other tools. It is important to figure out exactly where a new tool fits into the workflow, who provides data for it, where the generated content goes next, and what and who is affected when the content is iterated upon. After introducing the pillars by referencing existing literature, we argue in case studies in sections 4.1 and 4.2 that these pillars are in fact useful to industry.

It is important to note we are not arguing other approaches to MI-PCG are not valid. The position as it is presented here, is that if you want to make a MI-PCG tool which considers an industry audience, then we put forward arguments that recommend you consider our three pillars. Research with a different focus exist outside this position and is not contrary to it.

2 BACKGROUND

In their book [27] Norman et al. introduce the two concepts, *the gulf of evaluation* and *the gulf of execution*. The gulf of evaluation is how easy it is for a user to perceive representations and feedback, and the gulf of execution is how easy it is to accomplish a given task. The gulf of execution is an important concept for respecting designer control. Shortening the gulf of evaluation and the gulf of

execution are important parts of respecting the creative process as both are part of the iterative creative process.

Moving on from general user experience to mixed-initiative (MI) interfaces, in one of the earliest publications on the subject Horvitz [13] lists twelve principles for MI interfaces with direct manipulation. The first of these deals with the question, “is the tool actually adding value?” A MI interface need to add tangible advantages over a well crafted user interface with no agent.

Building on that, Liapis et al. [22] lists six requirements and questions for designing MI-PCG tools. The one question that is most relevant to our work is “How can the method for control over content be balanced?”. It has a clear answer in the context of this paper, as we will see in the case study in section 4.1 and in the analysis in section 3.1. The follow up question, “How to resolve conflicts that arise due to the human stating conflicting desires?”, leads us invariably to the question of how can we give maximum control to the user without resorting to manual labour, or similarly, how can we create the most empathetic agent possible?

Regarding *respect existing work processes*, the authors have been able to find little literature on the analysis of asset pipelines or on MI interfaces that describe how such tools integrate with the workflow of the average game developer. In lieu of not being able to find many references we have made a case study (section 4.2), that will argue how important it is for a content creation (CC) tool to fit into existing work processes and asset pipelines.

3 THE THREE PILLARS

3.1 Respect Designer Control

When we invent a new tool for PCG, we must make sure that the tool gives enough control to the creator that it can bring out the designer’s vision. As Lambe [21] states “Manual creation’s strengths are PCG’s weaknesses”. Manual methods give the control the designer ultimately wants (as argued in section 4.1), but does not provide the automation and potential speed up in work process that some PCG tools can provide. Thus the impossible sounding goal is to create a tool with the speed of automation and control of manual labour. Gingold’s “magic crayons” [12] embody this approach as they “enable authors to obtain satisfactory results with a small amount of effort”, “are artistically expressive” and “are magic because they are imbued with the power of computation.” The MI magic crayon, is the empathetic tool that directly senses what a user is trying to do in a given context. In [23, 24] Liapis et al. get closer to this MI magic crayon by trying to learn the user’s preferences.

Liapis et al. [22] describe the “user fatigue” that sets in, when a tool does not allow the designer to converge fast enough on the wanted result, for example, when the designer is forced to go through too many iterations, if feedback is slow, when there are too many options or the interface requires a very specific input. These boil down to providing adequate controls for the designer to do what they need to in order to achieve a given goal. We see that the width of the gulfs of execution and evaluation are important factors in causing user fatigue.

Giving control by allowing the designer to overrule the MI-PCG agent is not equal to *respecting designer control*, as this takes away the co-creative human-agent relationship entirely. Instead some MI interfaces such as Tanagra [31, 32] allow local editing, while the

surrounding generative level geometry adapts to those changes to satisfy pre-defined constraints.

3.2 Respect the Creative Process

When making changes to procedurally generated content, the iteration must be so fast, that it can compete with the instant feedback of manual labour [5, 30]. The creative process is a feedback loop of trying something out, seeing the result, making changes, seeing the new result, making further changes, and so forth until the designer is satisfied. To stay focused on the task, it is important that this feedback loop is as short as possible. As shown in figure 1 Compton [6] distills the feedback loop into four steps: *build a hypothesis, modify the model, evaluate the result* and *update the model*. Compton defines these four steps as the “grokloop”, and considers it as being a way for the user to interact with a generative tool and examining its possibility space. She says “I found myself wanting a way to say *the speed of learning depends on how short the loop is*”. While Compton’s work is directed towards casual creators, her arguments are just as valid for game developers. *Respecting the creative process* echoes Compton’s call to make a grokloop that is as short as possible.




Figure 1: Visualising the grokloop – adapted with permission of the author [6, Fig. 8.2]

The time spent on the grokloop is not only important when interacting with an agent, but also when training this agent. This is the point where machine learning (ML) algorithms with long training times can fall short [5]. In their breakthrough paper on interactive machine learning (IML) Fails et al. [10] introduce the principle “fast and focused”. They focus on empowering non-technical users to quickly create classifiers for ML based image classification algorithms, and they describe the principle as upheld if the user is able to quickly create a classifier, while being focused on the classification problem itself, and not the underlying algorithms. Their variant of the grokloop is ML specific and specialises to *manually classify, generate classifier* and *evaluate classifier*. Through the principle of “fast and focused”, a successful application will be able to perform this loop at interactive frame rates and break out of it entirely within minutes. The principles of IML are in opposition to

classic ML algorithms where a high level of technical skill is often required, and training data often needs to be fed to the algorithm for a long time, without an interactive interface to test when the user is happy with the result of the learning.

3.3 Respect Existing Work Processes

To facilitate maximum adoption of PCG tools in the industry, it is important that innovations plug directly into existing workflows. To make a tool easy and cost-effective to adopt, it should only affect the workflow of the people on the team using it. For example, if two artists A and B work on the same asset at different points in the pipeline, and B is dependent on A's output, a PCG tool to optimise part of A's workflow should output assets in a format that B can work with. It is worth remembering that while most academic research is done on an individual level or in small teams, many video games are made by teams orders of magnitude larger. For example over 3000 people worked on Red Dead Redemption 2 [11, 34]. Even with smaller teams of tens or hundreds of people using a plethora of tools, changing workflow becomes a potentially costly and risky proposition. If a research team wants to maximise the chance their research will be adopted into a game production pipeline, they need to adopt the path of least resistance. Examples of commercial PCG tools with a focus on easy integration with existing asset pipelines can be seen in table 1.

4 CASE STUDIES

We now discuss two case studies to further argue the validity of our three pillars. The first study focuses on the improvements on the Evolutionary Dungeon Designer [7] made throughout a number of works [1–4]. These improvements mainly relate to the first pillar, *respect designer control*, though the second pillar, *respect the creative process*, is also expressed. The second case study focuses on the third pillar, *respect existing work processes*, by examining twelve randomly selected PCG tools and which features, if any, they emphasise to ensure they fit into existing work processes.

4.1 Evolutionary Dungeon Designer

Dalhskog et al. [7], Alvarez et al. [1, 2], and Baldwin et al. [3, 4], created a series of works based on the procedural generation of dungeons via a MI interface. The initial work [7] analyzed 91 games published between 1975-1993 and extracted patterns that were transformed into mechanical game design patterns to be used for procedural content. Through user surveys they documented the improvements made [1, 2, 4]. In the first of those works [4], 3 out of 5 users requested some sort of local control, so that edited parts were not modified by the evolutionary search algorithm. This supports the pillar *respecting designer control*. It can be argued that the tool already upholds the first pillar, as the user is able to manually edit parts, or the user could keep generating new searches until they found something they were happy with. However, these arguments are easily dismissed. In the former case, the user would have to design the whole dungeon by hand, in which case the co-creative agent has been entirely omitted. In the latter case, in theory the user could keep using the tool to search for new layouts until they found something pleasing enough to be a starting point for manual editing. However, there is no guarantee that the algorithm would

Product	Content Pipeline Support
Arbaro[8]	Exports to Povray, DXF and Wavefront Obj 3D model formats
Dreamscape[28]	Integrates with 3D Studio Max
Houdini Engine[17]	Integrates with common CC tools such as Maya, Cinema4D and Unity
Mixamo[16]	Exports to formats used in Unity, Unreal and Blender
Speed Tree[19]	Integrates with Unity, Unreal and Lumberyard. Provides custom SDK
Substance Designer[14]	Integrates with Maya, Houdini, Unity, Unreal and more
Substance Painter[15]	Exports to FBX, Wavefront Obj, X and more 3d model formats
Tree It[33]	Integrates with Maya, Cinema4D, and exports to common 3d model formats
Vue[9]	Extensive import and export support for 3D geometry and 2D Maps
World Creator[36]	Exports to formats compatible with common CC tools
World Machine[25]	Integrates with Maya, Cinema4D
Xfrog[18]	

Table 1: Advertised asset pipeline support by PCG tools

converge within reasonable time, so most likely user fatigue as described in section 3.1 would set in. Also, 4 out of 5 users argued they would like to have a way of highlighting the design patterns. Seen through the lens of *respecting the creative process*, this is a way of shortening the grokloop.

More of the suggested improvements were implemented and described by Alvarez et al. [1]. Along with further improvements Alvarez et al. [2] includes a user survey, where the authors conclude: "To a certain extent, controllability is preferred than expressivity, as the users continuously try to impose their vision, which is a non-trivial task for automated systems to capture, thus, the users are more likely to sacrifice to a certain degree expressivity and exploration of the tool by gaining control over the generated content." This expresses the pillar, *respect designer control*.

4.2 PCG Middleware

Table 1 shows custom content pipeline support as described by the products' own advertising on the web. We have trusted this advertising and it has not been verified, as in this context, the important part is the intent that the product's description declares.

Most of the larger commercial software packages list one or more sales pitches about how well their software integrates with and can be customized to a games studio's pipeline. For example on the Houdini Engine [17] web page, it says "Houdini Engine supports a deep integration of Houdini and its procedural workflow within the larger framework of a Studio pipeline." This re-assures decision makers that the tool will optimise and not disrupt work processes.

All twelve pieces of software provide some sort of integration with a studio's workflow, even if in the case of freeware like Arbaro [8], it is just a simple export functionality. We conclude that providing support for integration with a studio's content pipeline is a central feature to be advertised by PCG tool-makers.

5 CONCLUSION

We have argued through referencing literature and two separate case studies, that at least three central design pillars should be taken into consideration when researchers work on MI-PCG tools which consider a games industry audience. These three pillars are *respect designer control*, *respect the creative process* and *respect existing workflow*. While the focus of most MI works is on attributes that fall within the pillar of *respect the creative process*, such as a short grokloop and visual feedback, most works disregard *respect designer control* by resorting to letting the designer lock down and manually edit parts of a piece of content. This inspires us to think much can be gained by focusing on increasingly empathetic magic crayon [12] agents.

Regarding *respecting existing work processes*, we should begin to think about how PCG tools integrate with existing content pipelines. This can be done as high-level thinking about where in the content pipeline a tool fits in, who it will be used by and how the exported content will be used further down the line. These questions inform us about the possible use and important features of a tool. Knowing that artists, programmers and level designers often work on the same data at different stages in the pipeline, can affect whether we put the emphasis on art or game logic, who it affects when new edits are applied to existing data, what data the tool imports/exports and consequently what sort of data is processed and how abstractions will be modelled. On a more practical level, some tools in table 1 support data making roundtrips, so the same data can be exported as well as imported. For example, a heightmap generated as part of a terrain asset, might be exported and fine-tuned in an image editing program, then re-imported into the terrain generation tool, where the modified heightmap will affect existing 3D geometry. Similarly, if a dungeon is updated and re-exported, we should consider what workers will need, further down the content pipeline to update game logic and re-apply art. By portraying our tools in a real content pipeline, we can focus the scope of our work.

It is our hope that this paper can contribute to a discussion about how we can improve MI tools for PCG and reach across the divide from academia to people with a similar interest in the games industry.

ACKNOWLEDGMENTS

This work was funded by the EPSRC Centre for Doctoral Training in Intelligent Games & Game Intelligence (IGGI) EP/L015846/1.

REFERENCES

- [1] Alberto Alvarez, Steve Dahlskog, Jose Font, Johan Holmberg, and Simon Johanson. 2018. Assessing aesthetic criteria in the evolutionary dungeon designer. (2018), 1–4. <https://doi.org/10.1145/3235765.3235810>
- [2] Alberto Alvarez, Steve Dahlskog, Jose Font, Johan Holmberg, Chelsi Nolasco, and Axel Österman. 2018. Fostering creativity in the mixed-initiative evolutionary dungeon designer. In *Proceedings of the 13th International Conference on the Foundations of Digital Games - FDG '18*. ACM Press, New York, New York, USA, 1–8. <https://doi.org/10.1145/3235765.3235815>
- [3] Alexander Baldwin, Steve Dahlskog, Jose M. Font, and Johan Holmberg. 2017. Mixed-initiative procedural generation of dungeons using game design patterns. *2017 IEEE Conference on Computational Intelligence and Games, CIG 2017* (2017), 25–32. <https://doi.org/10.1109/CIG.2017.8080411>
- [4] Alexander Baldwin, Steve Dahlskog, Jose M. Font, and Johan Holmberg. 2017. Towards pattern-based mixed-initiative dungeon generation. *20506* (2017), 1–10. <https://doi.org/10.1145/3102071.3110572>
- [5] Michael Buttner. 2019. I3D'19 Keynote: "Machine Learning for Motion Synthesis and Character Control. At 17:21. <http://www.youtube.com/watch?v=zuvmQxcCOM4&t=1760s>. Accessed: 2019-07-25.
- [6] Kate Compton. 2019. *Casual Creators*. Ph.D. Dissertation. University of California, Santa Cruz.
- [7] Steve Dahlskog and Julian Togelius. 2015. Patterns, Dungeons and Generators. *Foundations of Digital Games* June (2015).
- [8] Wolfram Diestel. 2003. *Arbaro*. <http://arbaro.sourceforge.net/>. Accessed: 2020-01-31.
- [9] e-on software. [n.d.]. *Vue*. <http://info.e-onsoftware.com/more-info-vue>. Accessed: 2020-01-31.
- [10] Jerry Alan Fails and Dan R. Olsen. 2003. Interactive machine learning. In *Proceedings of the 8th international conference on Intelligent user interfaces - IUI '03*. ACM Press, New York, New York, USA, 39. <https://doi.org/10.1145/604045.604056>
- [11] Rockstar Games. [n.d.]. *Red Dead Redemptions 2 Credits*. <http://www.rockstargames.com/reddeadredemption2/thankyou>. Accessed: 2020-01-31.
- [12] Chaim Gingold. 2003. *Miniature Gardens & Magic Crayons : Games , Spaces , & Worlds .*, 116 pages. <http://levitylab.com/cog/writing/thesis/>
- [13] Eric Horvitz. 1999. Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems the CHI is the limit - CHI '99*. ACM Press, New York, New York, USA, 159–166. <https://doi.org/10.1145/302979.303030>
- [14] Adobe Inc. [n.d.]. *Substance Designer*. <http://www.substance3d.com/products/substance-designer/>. Accessed: 2020-01-31.
- [15] Adobe Inc. [n.d.]. *Substance Painter*. <http://www.substance3d.com/products/substance-painter/>. Accessed: 2020-01-31.
- [16] Adobe Inc. 2008. *Mixamo*. <http://www.mixamo.com/#/>. Accessed: 2020-01-31.
- [17] Side Effects Software Inc. [n.d.]. *Houdini Engine*. <http://www.sidefx.com/products/houdini-engine/>. Accessed: 2020-01-31.
- [18] Xfrog inc. [n.d.]. *Xfrog*. <http://xfrog.com/>. Accessed: 2020-01-31.
- [19] Inc. (IDV) Interactive Data Visualization. 2002. *SpeedTree*. <http://store.speedtree.com/>. Accessed: 2020-01-31.
- [20] Lisa C. Kaczmarczyk. 2015. When worlds collide. *ACM Inroads* 6, 3 (aug 2015), 42–43. <https://doi.org/10.1145/2794296>
- [21] Ichiro Lambee. 2012. *Procedural Content Generation: Thinking With Modules*. http://www.gamasutra.com/view/feature/174311/procedural_content_generation_php Accessed: 2020-02-01.
- [22] Antonios Liapis, Gillian Smith, and Noor Shaker. 2016. *Chapter 11: Mixed-initiative content creation*. Springer International Publishing, Cham, 195–214. https://doi.org/10.1007/978-3-319-42716-4_11
- [23] Antonios Liapis, Georgios N. Yannakakis, and Julian Togelius. 2012. Adapting Models of Visual Aesthetics for Personalized Content Creation. *IEEE Transactions on Computational Intelligence and AI in Games* 4, 3 (sep 2012), 213–228. <https://doi.org/10.1109/TCAIG.2012.2192438>
- [24] Antonios Liapis, Georgios N. Yannakakis, and Julian Togelius. 2014. Designer modeling for Sentient Sketchbook. In *2014 IEEE Conference on Computational Intelligence and Games*. IEEE, 1–8. <https://doi.org/10.1109/CIG.2014.6932873>
- [25] World Machine Software LLC. [n.d.]. *World Machine*. <http://www.world-machine.com/>. Accessed: 2020-01-31.
- [26] Natural Motion. [n.d.]. *Natural Motion*. <http://www.naturalmotion.com>. Accessed: 2020-01-06.
- [27] Donald A. Norman and Stephen W. Draper (Eds.). 1986. *User Centered System Design: New Perspectives on Human-computer Interaction*. CRC Press.
- [28] Sitni Sati. [n.d.]. *DreamScape*. <http://www.afterworks.com/DreamScape.asp>. Accessed: 2020-01-31.
- [29] Noor Shaker, Julian Togelius, and Mark J Nelson. 2014. PCG Book - Appendix A Game-designer interviews To. In *Procedural content generation in games: A textbook and an overview of current research*. Springer, Chapter Interviews, 225–237.
- [30] Karl Sims. 1987. Locomotion of Jointed Figures over Complex Terrain.
- [31] Gillian Smith, Jim Whitehead, and Michael Mateas. 2010. Tanagra: A Mixed-Initiative Level Design Tool. *Proceedings of the 2010 International Conference on the Foundations of Digital Games FDG 2010* (2010), 209–216. <https://doi.org/10.1145/1822348.1822376>
- [32] Gillian Smith, Jim Whitehead, and Michael Mateas. 2011. Tanagra: Reactive planning and constraint solving for mixed-initiative level design. *IEEE Transactions on Computational Intelligence and AI in Games* 3, 3 (2011), 201–215. <https://doi.org/10.1109/TCAIG.2011.2159716>
- [33] Evolved Software. [n.d.]. *Tree It*. <http://www.evolved-software.com/treetit/treetit>. Accessed: 2020-01-31.
- [34] Rockstar Studios. 2018. *Red Dead Redemption 2*. Game [PlayStation 4, Xbox One, Microsoft Windows]. Rockstar Games.
- [35] Julian Togelius. 2014. Why academics and game industry don't collaborate on AI, and how we could improve the situation. <http://togelius.blogspot.com/2014/10/why-academics-and-game-industry-dont.html>. Accessed: 2020-04-19.
- [36] Bite The Bytes UG. [n.d.]. *WorldCreator*. <http://www.world-creator.com/>. Accessed: 2020-01-31.