

Considerations and Concerns of Professional Game Composers Regarding Artificially Intelligent Music Technology

Kyle Worrall and Tom Collins

Abstract—Artificially intelligent music technology (AIMT) is a promising field with great potential for creating innovation in music. However, the considerations and concerns surrounding AI-generated music from the perspective of professional video game composers have yet to be fully explored. In this study, 11 professional video game composers were interviewed to determine how they feel about AIMT and how this informs future research and tool design within the games industry. The interviews were analyzed using a reflexive thematic analysis (RTA) to identify key themes. The study found that while composers recognize the benefits of music AI, they have complex concerns beyond the obvious concerns for AI infringing on their agency and creativity. There is an inherent clash between the creative ego and music AI, which can make it difficult for composers to embrace this technology. Furthermore, a lack of standard technical knowledge, support, understanding, and trust of music AI is impeding tool use within the industry. These findings have implications for music AI researchers and industry practitioners. By better understanding the concerns and considerations of professional creatives, researchers can design and communicate their tools more effectively to music professionals. Moreover, this study lays the foundation for empirical research into the relationship between professional creatives and emerging AI technology – a topic that is under-emphasised in current research.

Index Terms—Games, Music, Artificial Intelligence, Thematic Analysis, Procedural Music Generation.

I. INTRODUCTION

VIDEO games have long been considered a medium that is well suited to music generation, due to their interactive, and unpredictable non-linear nature [12]. However, in comparison to visual and mechanical procedural content generation (PCG), and while it has already been proposed that procedural music generation (PMG) could empower composers through assistive means [44], the generation of music and sound has seen a much more limited scope of application in the video game industry. This is surprising, as the use of PCG in games has already been used to overcome repetition-based fatigue associated with visual assets in games [23]. Furthermore machine learning approaches have been used to speed up creative workflows for animators [6]. Meanwhile, in music for games, it is common for ~4 hours of music to be deployed and hence, to some extent, repeated across ~100 hours of gameplay [44], [70].

AI-driven music technologies (“music AI”) outside of games are advancing in their capabilities in performing cre-

ative tasks, such as composition [13], [27], musical in-filling [25], expressive rendering [30], [71], mastering [54], and mixing [40]. Furthermore, music AI has been shown to provide assistance to novices in musical co-creation [37]. The breadth of music AI applications here demonstrates some of the many ways that music AI could support professional composers. We note that with the exception of a few, largely rule-based examples such as the Dynamic Percussion System [9], [35], the use of music AI technology has not been adopted widely in the video games industry.

Industry practitioners and researchers highlight some limitations holding back procedural music/sound in the video games, such as: the need for robust timing systems and more audio programmers to support tools [66]; resource intensiveness of modern video games [44]; inconsistency in the quality of generative output [65]; a lack of human nuance or expression in computer-generated music [71]. Additionally, creatives in online forums and social media present a potential narrative of aversion to creative AI (AI that can be used in art, music or other creative tasks), citing ethical and legal concerns around its use (and music AI) to bear [13], [22], [73].

There is a lack of empirical research being conducted to investigate this seeming resistance to music AI and procedural music in games, from the viewpoint of an important stakeholder in this arena – the professional composer. In this paper, we lay a foundation of this empirical research by interviewing 11 professional video game composers of varying experience levels about their thoughts on AIMT in games (used hereafter to describe both PMG and music AI), in order to better understand the relationship between emerging AIMT and professional creatives. We present the results of a reflexive thematic analysis (RTA, [7], [8]) of the data in order to identify a variety of recurring themes and sub-themes, addressing the following two research questions:

- RQ1 How do composers feel about AIMT, and how does this inform future research?
- RQ2 What can we learn about professional workflow and technical knowledge to inform future intelligent music tool research/design?

The remaining sections of this paper are structured as follows: first, we provide a review of the literature surrounding technology acceptance, PCG and PMG. Second, we discuss the aims of our research and how the data were collected, and outline the qualitative approach used in our analysis. Third, we present the results of our RFA ([7], [8]), and discuss the 1195

This paper was produced by the IEEE Publication Technology Group. They are in Piscataway, NJ.

Manuscript received April 19, 2021; revised August 16, 2021.

codes, 29 sub-themes, and five themes that were identified (see Figure 1 for overview of themes across all the codes). Lastly, we conclude by discussing our findings in relation to the video game industry, their implications for researchers and industry practitioners, and the limitations of our work.

II. RELATED WORK

A. Technology Acceptance

Research into technology acceptance dates back to the 1980's with the Technology Acceptance Model (TAM), which states that two main factors in the acceptance of technology are perceived ease-of-use and perceived usefulness — though external variables such as social influence can play a role [16]. This was later expanded upon in two different papers: one found that specific determinants could affect the perceived ease-of-use as they develop over time, namely control, intrinsic motivation and emotion [59]; another proposed an updated version of the Technology Acceptance Model (TAM2) and tested it in four studies, where social influence processes (subjective norm, voluntariness of use and image), and cognitive instrumental processes (job relevance, output quality, result demonstrability, and perceived ease-of-use) were shown to significantly influence acceptance [60].

A unified model of technology acceptance (UMTA) has also been proposed and empirically validated, bringing together eight competing models of technology acceptance [61]. This research reveals four constructs that play a significant role in directly determining user acceptance and usage behaviour: performance expectancy, effort expectancy, social influence and facilitating conditions, while factors such as gender, age, experience and voluntariness of use could impact each of the four constructs to some degree [61]. As such, it is understood that technology is more likely to be accepted if it is easy to use, performs well, and public opinion of the technology is positive.

B. Procedural Content Generation

PCG is a well documented area of research, where a variety of computational approaches are applied to the generation of content in games. Examples of these approaches include: constructive PCG (which use a rules-based systems); PCG via machine learning [53], and search-based PCG (those using stochastic search-based algorithms) [56]. Within the industry, PCG has been leveraged to generate worlds [20], [39], geometry [48], world history [24] and determine resource/enemy placement throughout generative levels [10]. In particular, PCG is a very popular technique for independent games studios, where PCG is a “technical strategy for generating content despite limited production resources” [17, p.197]. PCG has also seen a wide range of applications at triple-A (AAA) companies, such as the generation of dungeons in *Bloodborne* [4], weapons in *Borderlands 3* [5] and the world in *Far Cry 5* [21].

In these examples, PCG has been used at AAA to develop games faster and with reduced costs [2], [3], [20] while preserving quality [2], [50], [64]. PCG's application has also led to unending game experiences for players and to supporting

live-service games (LSG) [68]. LSGs are a type of video game where content is delivered according to a continuing revenue model over the game's life span, often requiring more music with new content. Examples of LSGs include Massively multiplayer online games such as *Destiny 2* [18] and Multiplayer online battle arenas such as *League of Legends* [36].

PCG is often critiqued for the lack of diversity and polish in generated output, however, where generative content is compared to 10,000 bowls of oatmeal [14]. While PCG has advantages in a creative aspect, this can lead to criticism from players, but also likely affects player retention in games that lack diversity.

Considering PCG use in relation to the unified model of technology acceptance [61], we can see that the use of PCG tools is bringing value to studios (high performance expectancy) and that they are developing their own tools to do this (which is high effort expectancy, but these techniques are very popular and hundreds of online resources are available to learn how to code them, which reduces the effort expectancy somewhat, and reduces the effort expectancy of level designers — the end users). Additionally, PCG has been widely adopted (regardless of concerns for the blandness of generative bowls of oatmeal, reducing social pressures and increasing the likelihood that this technology will continue to be accepted [61].

C. Procedural Music Generation

The impact and benefits of PMG are similar to those of visual/mechanical PCG, allowing game developers to computationally generate, perform or transform music in real-time in video games [44], [69]. With a notably small data allowance for audio in games [38], [55], [67], and the increasing need for variation [44], [70] and interactivity [43], [55], there is justification for the adoption of PMG in the video game industry.

Some successful instances of the use of PMG in games utilise approaches that do not require composers to change their workflow, or participate in technical work beyond what is expected of them already [9], [35], [65], which reduces effort expectancy [61]. As an example of successful symbolic music generation in industry, the Dynamic Percussion System is used in *Rise of the Tomb Raider* [45] and utilises machine learning to generate drums in real-time during gameplay, where the composer only needs to give the developers their percussion in symbolic form (MIDI files). This lowers the amount of effort required by the composer and does not add additional steps into their workflow [9], [35]. Weir on the other hand creates an audio domain music generation tool (PULSE) for *No Man's Sky* [41], where they specifically told the band providing the music (65daysofstatic) to compose an album as they typically would, but to deliver them in stems so that the musical elements could be manipulated to generate music in-game, without impacting the band's creative process [65].

In another example where the games composer was not involved, developers leveraged an older open-source tool for jazz solo generation known as Impro-visor [32] to overcome the lack of access to machine learning specialists for music

in indie game studios [1]. This decreases effort expectancy in another dimension by using functional, open-source code that already has a GUI rather than being left at command line [61]. While Plut & Pasquier outline a more comprehensive list of PMG systems in games, it is clear that when compared to PCG, PMG has seen much less application in the video game industry, giving the impression of less acceptance [44].

Outside of the industry, automatic music generation for games as a research topic has seen more interest in recent years, often focusing on rules-based approaches or less computationally taxing approaches such as Markov chains to decrease latency and computational costs in real-time [19], [51], which in turn increases performance expectancy [61]. However, there is a disconnect between industry and academia, so while research has been tackling game-specific use cases such as transition generation [15], generating music that adapts to gameplay [29], using gaussian mixtures to control melodic shape in generative music [72] or generating music to match nonplayer character relationships [63], procedural music systems that generate new material go largely ignored in industry. By comparison, the more popular use of procedural systems in video games is that of the 'transformational' approach [44], [69], that uses simple rules-based approaches or Finite State Machines (FSM) to transition between human authored tracks [28], [47], [52].

D. PMG Acceptance in Games

There are a variety of reasons that researchers and practitioners have identified as potential reasons for a lack of procedural music tools in the industry. The first is a lack of experienced developers for building or supporting PMG or AI-driven music tools in the video game industry [66], especially when video games have real-time computation and resources to consider when building PMGs [44]. This means potential users (game developers) have lower performance expectancy, and higher effort expectancy, as there is no support for using these tools, and using them without careful consideration can damage the game-play experience [61]. Compare this to non-games-related music AI tools, for instance, such as the Google Magenta Suite [46] and in-painting tools for Ableton [25], [26]. The use of AIMT is more adopted by Digital Audio Workstation developers than by games companies. This makes sense as these companies do not have to consider game rendering and frame drops as real-time concerns, just the typical real-time considerations of an audio engine. This lowers performance expectancy for the software developers using such tools.

Another reason that adoption of AIMT may be low is the inconsistency in the quality of generative output [14], [65] or a lack of human nuance in computer-generated music [71]. This is supported by anecdotal evidence such as composer discussions on social media, which outline the low quality and inconsistency of generative output as a factor against using these tools. This resonates with the unified model of technology acceptance, as low quality, un-nuanced output designates low performance expectancy, which will lower acceptance levels. A final reason that could potentially be lessening the

adoption of this technology is the ethical and legal concerns around creative AI (including music AI), where generative art tools (e.g., Midjourney) and LLMs (e.g., ChatGPT) have caused a stir in the creative community, for a variety of reasons based around copyright infringement concerns [13], [73], and concerns from creatives regarding misrepresentation and ownership [22]. These tie into the social influence aspect of the unified model of technology acceptance, as discourse surrounding generative art/music is volatile, and with such discourse comes social pressures to abstain from their use.

The above could explain the lack of adoption of AIMT in video games, but there is a clear lack of empirical evidence outlining how composers (and creatives) feel about these developments. Without looking to understand the end users of these tools, it is unlikely that they will be adopted. Highlighting the thoughts and feelings of such users, as we go on to do, could shape future tool design for increased technology acceptance, and fill this gap in the literature.

E. AIMT Acceptance

Research in the field of AI co-creation (the study of how humans and AI interact together) has shown that AI tools can support novices in composing tasks [37], however, research into the acceptance of AIMT for music professionals has yet to be fully explored.

Tsiros & Palladini investigate AI-assisted music production and propose a framework for how to design AI tools to be human-centric, noting the importance of lessening the following risks for users: AI making sub-optimal decisions, AI impacting engineer authority and control, AI forcing extreme change onto existing workflow [57]. Meanwhile, Vanka et al. study the opinions and thoughts of mixing engineers about how they use AI mixing tools and find that users can be separated into amateurs, semi-professionals and professionals, who all use intelligent mixing tools for different purposes [58]. They find that professionals are in favour of AI mixing tools, and that they use these tools as a way to speed up their workflow and experiment creatively. However, Vanka et al. state this is ultimately down to individual differences and that it is important that these tools integrate seamlessly with existing workflows, have a fine balance for control and automation, and become context-aware (as generic output tends not to be of interest to users) [58]. The above begins to provide a baseline for understanding AIMT acceptance, but does not tackle the realm of video game music, nor of professional composers, which this paper seeks to address.

F. Semi-Structured Interviews in Related Research

Semi-structured interviews are used in games research to explore the thoughts and opinions of game developers regarding particular topics, such as experiences using creative AI, namely Text-To-Image Generation (TTIG) models [62], and to investigate how terms are defined by those in the industry actually using them (such as in the case of quests, which can have a widely varied meaning across studios/games) [34]. Beyond their application in games, semi-structured interviews are used to interview visual artists and AI art communities to

further probe experiences using TTIG models [11], [33]. These applications demonstrate the suitability of semi-structured interviews, which offer reliably comparable data, but also the flexibility to introduce follow-up questions.

III. METHOD

A. Data Collection

We use social media platforms (Twitter/Facebook) to recruit 11 professional video game composers for semi-structured interviews. As an indication of participants' level of expertise, all participants have worked on at least one published game. The supplementary materials provided with this manuscript show that of our 11 participants, seven identify as male, three identify as female and one identifies as non-binary. This is a fairer representation than in the industry itself, where 84% of the industry identifies as male [49].

In terms of experiences that might inform their responses, all participants have multiple years of experience. For composing experience, participants self-report that they work on projects of the following levels: 27.3% at the AAA level, 27.3% between the indie and Midcore levels (where Midcore describes small medium enterprises that produce professional games on smaller budgets than AAA) and 45.4% entirely on indie titles. For adjacent experience, 45.5% of participants have worked, or currently work in audio-adjacent roles in the video game industry, where 50% of those 6 participants work for a AAA company.

Participants are asked ten questions ranging from topics such as their level of experience, specific considerations relating to their role, thoughts on their workflow, programming knowledge and then opinions on a range of existing music AI (see Fig. 1). These questions are designed to introduce participants to a range of different music AI tools likely to be of differing levels of controversy, so that we can gauge differences in their response based on what kind of role the tool plays (i.e. composition, mixing, mastering, humanising). This research is given ethical approval by the University of York Computer Science Ethics Committee and participants are each compensated £25 for their time.

B. Data Analysis

These interviews are conducted via Zoom, and the ~13 hours of audio recordings are transcribed and anonymised. During initial coding of data in our RTA, we use an inductive approach, whereby we allow the data to determine the themes, rather than approaching the data with preconceived notions and theories [8]. When interpreting the codes, we use a latent approach, whereby we read into the subtext of the data to find underlying meaning, rather than relying only on stated opinions/thoughts [8].

These decisions are made as we are interested in all of the information that the participants provide, not just the codes that related to preconceived theories (such as the TAM and UMTA).

1) *Potential Bias*: Data analysis for this project is carried out by a single researcher, meaning that there is a chance that researcher bias may affect the resulting codes and themes. The typical approach to minimise this type of bias is to have two or more researchers code the data and then measure for inter-rater reliability in the codes using Cohen's Kappa. This is not possible in this case, therefore a reflexive approach to thematic analysis is used, whereby once the codes and candidate themes are identified, the researcher then takes the candidate themes back to the raw data and then checks that the themes are representative of the data. This minimises the chance that researcher's preconceived notions will lead to themes that do not fit the data. Furthermore, during data coding, an inductive approach to thematic analysis is used, whereby the data is allowed to shape the themes, instead of a deductive approach, where the researcher begins with notions derived from existing theory. This choice decreases the likelihood that researcher bias could influence the results of this study.

IV. RESULTS

As a result of the RTA, we identify 1199 codes, which can be grouped into 250 broader codes. These broader codes can then be grouped into 30 sub-themes, which fit under five thematic umbrellas (with some overlap — see Table I).

Five themes were identified in this RTA. Three of the themes create a narrative that addresses RQ1, and two address RQ2. These themes are outlined below, and the coverage of the themes among the 1199 codes can be seen in Table II.

A. Benefits Of Music AI

The first overarching theme identified in the data explains that composers can mainly see and are excited by the benefits to workflow and creativity that could come with adopting music AI. Participants construct the benefits of music AI in a few ways. Firstly, by focusing on what they could perceive as workflow benefits, where music AI is part of a toolbox that saves composers time and supports them in completing tasks efficiently.

A really useful way for the AI to help out [...] You can get it to finish or at least get started on the tasks that I'm having trouble with, which in this case would be the fine tuning. [P4: Indie-Midcore Male]

[F]rankly, except to make the composer job easier, right? [...] the more things I can automate, the more creative I can be, because then I'm just focused on creativity.[P7: AAA Male]

Participants constructed a second sub-narrative of music AI as a creative benefit, where music AI performs the role of prompting composers who are stuck with writer's block rather than finishing the music for them, or where it is used for musical exploration that leads to new, novel and unexpected musical ideas in their work.

I think that could be fun to use. I think sometimes if you're a bit creatively stuck it could be a really good prompt, the same way that writers use prompts to

Sub-themes	<i>n</i> = 1199	Thematic Group
AI can benefit workflow	70	Benefits
AI can benefit creativity	60	Benefits
AI can improve accessibility	11	Benefits
AI could benefit the industry	16	Benefits
AI hard to visualise/understand	28	Understanding/Trust
Tools need to communicates musically	15	Understanding/Trust
Familiarity = Acceptance	1	Understanding/Trust
Curious about AI capability	4	Understanding/Trust
Ethical Concerns about music AI	42	Concerns
Bad AI can damage player experience	3	Concerns
AI could change composition process	78	Concerns
AI does not generate original, human quality music	191	Concerns
Concerns about job security	11	Concerns
Concerns about musical representation/ownership	83	Concerns
Technical concerns for music AI in games	58	Concerns
Lack of standard support/knowledge	100	Support
Assistants vs AI	6	Support
All composers needs are different	17	Support
music AI needs to easily fit into existing workflows	27	Support
Game devs and composers communicate differently	26	Support
Generating less melodic music is less problematic	31	Ego
Deep rooted psychological need to control	63	Ego
Desire for non-compositional AI	33	Ego
music AI clashes with the creative ego	133	Ego
The quality of music AI is improving, but not good enough	10	Ego
music AI is not well suited to music	3	Ego
If it benefits the game/players	29	Ego
Humanity in Music	16	Ego
AI is not worth the time/effort/cost	34	Ego

TABLE I

THE SUB-THEMES IDENTIFIED IN THE DATA, THEIR COUNT ACROSS THE DATA SET AND THE THEMATIC GROUPINGS THAT THEY REPRESENT.

Themes	Data Coverage
Ego vs AI	29.36%
Understanding/Trust	4%
Standardisation in Support	14.68%
Concerns about AIMT	38.87%
Benefits of AIMT	13.09%

TABLE II

PERCENTAGE OF CODES WITHIN EACH THEMATIC GROUPING IN THE DATA.

write, you know, . . .and that will kind of be a similar thing. [P6: AAA Female]

Furthermore, participants constructed this theme by adding in additional thoughts about how music AI can be a benefit to accessibility for novices (supporting existing research into musical co-creation as a tool for novices [37], and by mentioning how music AI could be a potential solution to legal issues around music used during the streaming of games, an issue being explored by startup Infinite Album.

Yeah, that would be cool. I think that’s another example of it making it more accessible. Cause like you say, depending on music theory knowledge, you may not even be aware that inversion exists or what inversions are or how to create them and stuff like that. [P6: AAA Female]

B. Complex Concerns About Music AI

The second overarching theme identified in the data explains that composers have multifaceted, and deeply complex concerns about music AI that go beyond the obvious concerns

for job security. Participants construct a narrative theme of the complex concerns about music AI, the first subset of which relate to concerns about the computational implications of music AI use in games, and inconsistent musical quality and lack of human nuance in the output, all of which can potentially have a severely negative impact on the gameplay experience, and subsequently on players as a whole.

You lose the composer’s tricks like all the flurries. . .but from my own standpoint, I don’t think it would be something that I would personally use, because everybody’s writing style is very different. [P1: Indie Male]

The second subset of the concerns of composers about music AI form a narrative about how such technology could have ethical and legal implications surrounding job security, musical ownership, misrepresentation (of a composer’s quality or competence if an AI makes a mistake) and how AI could potentially disrupt the way composers work, and damage their enjoyment of the music creation process.

At the end of the day, I think it would be something I wouldn’t particularly be comfortable using because you could definitely lose a lot of the composer’s identity. . .because they might write bass very differently from an AI or they might do things very differently. So again, I think the composition might lose a lot of its identity. [P1: Indie Male]

[Regarding AI voice] there’s severe legal issues with the idea of trying to market something based on the voices of existing people. Because fundamentally

you're going to... run into potential fraudulent cases.
[P8: AAA Male]

C. *The Clash Between Ego And AI*

The third overarching theme identified in the data describes an inherent clash between creative ego and music AI. Participants construct the theme of the clash between ego and AI in a variety of sub-themes that outline a compelling narrative, whereby composers disliked how generative AI steps on their toes in composition tasks, but had a contrasting desire for AI to handle non-composing tasks. When combined with the sub-theme suggesting that participants found AI use for less complex and melodic music to be more suitable than when used in melodic or interactive music, this suggests that the clash between the creative ego and AI not only exists, but could be mediated by how closely the music AI infringes on the tasks with which the individual identifies. In other words, a mix engineer may dislike a mixing AI more than a composing or sketching AI, because although each task is arguably as nuanced and complex as the others, their self-identity and in some way their self-worth is wrapped in their vocation. In response to being asked what they wish an AI could be used for:

[H]aving some feedback on mixes and they're [the AI] being like, ah, yes, clashing frequencies, like we suggest cutting here, boosting that and like, I never learned mixing that's the one thing that I wish I had some professional training. [P9: Indie Non-Binary]

Furthermore, this lack of openness to specific vocationally aligned AIMT seems to scale with the importance of the musical element being generated for a composer. For example, the majority of composers were not open to the generation of melody, going so far as to describe melody as important, special, a communication between client and composer, and as having soul. In comparison, participants were more open to chord generation (although some mentioned that generating musical inversions can change the whole feeling of a piece) and even more were open to bass or pad generation to support composition. As such this implies that the less ego-centric the element being generated, the less likely the composer is to find issue with the generation of the element.

I love to establish that melody before we rip it apart. Yeah. I would say just so it's in there somewhere because maybe that's just a sense of pride like, you wrote this melody and you really want it to be there and they want it to stand out... Cause sometimes coming up with just that melody between you and the client can be so special and important that you don't want ever bury it under iterations... Yeah, chords are chords but a melody has that soul, you know? [P11: Indie Female]

No, not interested... you know, melody is as simple as just any string of notes. It doesn't matter what the string of notes is. That can be a melody is it a good melody? How do I discern that? Like how did, how would a machine discern good from

bad? I don't know that anyone knows what makes a good melody... I would never want to rely on a machine to give me options because either I've heard it before a million times, because all it's going to do is regenerate the same kind of thing. I'm looking for something no one else has found... so I don't see that as being useful at all. Waste of time. [P7:AAA Male]

The clash between ego and AI is further supported by the way that participants described a desire for bespoke tools with lots of finer controls as part of the interface of music AI tools, which allow the users improved and more nuanced output, and in their concerns about musical ownership and misrepresentation (through lack of quality, or musical mistakes).

[The key to] feeling confident in those choices is being able to quickly scrub through them and audition them and make sure that I'm okay with that combination. And then in some cases almost being able to disallow certain combinations which then gets a lot more complicated to figure out. [P10: Indie Male]

A second part of this theme's narrative is identified in the sub-themes of AI not being suited to music at all and that there is humanity in art (whereby the participants were expressing that art and music are human expressions and pointless when carried out by an AI). This further supports the clash between creative ego and AI, as computers have been shown to perform almost as well as humans for some tasks (e.g., [71]), yet creatives fully believe that computers are incapable of participating in something in which they find value.

[There's a] camp that... art is a human-only expression and therefore a computer can never do it. [P10: Indie Male]

When do we cease to bring the humanity to whatever art we're making by allowing machines to make the art for us? I mean, what's the point of art in that case? [A]rt is arguably a communication from a human being to another human being... it's not, let the computer tell us what's good. [P7: AAA Male]

D. *No Standardisation Of Support And Understanding*

The fourth overarching theme identified in the data indicates that there is a lack of standardisation in the technical support provided within the industry to support composers, especially given how practices vary from studio to studio. This is especially true for AIMT as there is already a clear need for more audio programmers in the industry to support audio teams, but AIMT also requires machine learning or AI audio specialists as part of this support structure.

This theme also highlights a lack of standardisation of technical knowledge and language among composers, which has previously been identified around the term "procedural" [44]. This lack of standardisation inhibits composers' ability to understand AIMT clearly, which in turn lowers their trust and understanding in it.

Beyond the support needed to allow composers to easily experiment with AIMT, there are further operational support requirements that are created when using AIMT.

For example, if an algorithm generates symbolic music such as MIDI, then the game needs to either include samples to render this music in real-time, or the composer needs to book studios/musicians to make the AI music match the quality of the pre-recorded assets. If the tool is used early enough, then this could form part of the normal recording process, but if not, this then provides additional support needs for the user. This observation resonates with industry practitioners and researchers, who discuss the lack of audio programmers/audio-focused AI tool developers to support development and use of procedural audio tools in industry [66].

[Audio programming] is not really canonised as a field anyways. Like talking to the audio programmers that I've met, it's like, it wasn't even a dedicated profession. [P10: Indie Male]

If you need to do things like live recordings. . . you're probably going to need to send some music to live players [musicians]. . . if that's a requirement, then you're going to need to hire some players [musicians]. [P3: Indie Male]

Additionally, a sub-theme of “human assistance” emerges that supports the disparity in support structure between the levels in industry (in the composers’ own studios and not in game studios). In this sub-theme, composers at the AAA level are less likely to be open to using AIMT, as many of their tasks are already automated through human assistants (which they see as ideal, and not requiring a learning curve like new technology). Furthermore, the sub-themes of “technical concerns”, “no standardisation in understanding/language”, and “the disparity between the language of game developers and composers” form a narrative that demonstrate the hurdles keeping composers (and creatives in general) from learning about AIMT, as they are not only facing a lack of support, but a lack of cohesive language to use to ask and learn about new technology.

E. Lack Of Understanding And Trust In AIMT

The fifth overarching theme suggests that the technical nature of AIMT and lack of consistency in its attendant terminology make it difficult for creative individuals to gain confidence and proficiency in this domain. The distinct way that participants construct this theme is rooted in their ethical concerns (often stemming from a lack of understanding of how AI models are trained), combined with their difficulty visualising what AIMT can do. This is further impacted by the lack of musical language used in existing tools. All this ferments a distrust in AI, especially when considering the lack of consistency in output quality, and the current AI-negative narrative present on social media, meaning creatives can have a bias against using AI tools.

That one's a trickier one because it's hard to kind of envisage what that would actually, how that would function or sort of what it would sound like. [P6: AAA Female]

If we relate this distrust and lack of understanding back to the UMTA [61], then we can see how not being able to understand a tool leads to distrust, and then this distrust leads to negative bias impacting social influence, and lowering the likelihood of AIMT being accepted.

V. DISCUSSION

How the perception and understanding of music AI impacts the acceptance and use of AIMT within the video game industry and among professional composers is of central interest to researchers of music AI and music informatics, as well as in multiple application domains, such as AI musical co-creation and game audio. To our knowledge, prior to this paper, there is no empirical research looking into professional composers’ opinions on AIMT, and what we can learn about their technical knowledge, to better inform future design of and research in AIMT. Some music AI research involves asking professional composers for their opinions of generative output [42], however more often than not, music AI are evaluated by music students (due to their accessibility and musical knowledge) [30], [31], [71].

In this paper, we outline the five identified themes and now we discuss their relationship to our initial research questions:

- RQ1 How do composers feel about AIMT, and how does this inform future research?
- RQ2 What can we learn about professional workflow, technical knowledge and tool use to inform future intelligent music tool research/design?

With regards to RQ1, we find that while composers can see some benefits to music AI (mainly assisting to prompt creativity and speed up their workflow), there are two main themes holding back AIMT acceptance among professionals. The first is that they have complex and multifaceted concerns regarding AIMT, limited not only to the quality of generative output, but including: ethical concerns about training data¹; societal concerns about misuse by game developers; concerns about misrepresentation by and ownership of musical material generated based on their own work; concerns for job security; concerns for technical constraints at run time in games; and worries about how their workflow will be affected.

The second factor holding back the acceptance and use of AIMT is an inherent clash between the creative ego and AI, which scales depending on how closely the tool infringes upon tasks with which the creative individual identifies (in this case, compositional AI tools were rejected more than mixing and humanising tools). The ego in this case desires very fine-tuned control over AIMT that are non-composition focused, and do not infringe upon the music composition process. Furthermore, the adversity of reactions to generative music is linked to three other aspects: firstly, the complexity of genre of music being generated (where more ambient soundscape generation was more acceptable than melodic classical music generation); secondly, the medium for which the music is intended (where game composers see games as harder to generate for due to

¹While we are aware of AI approaches that do not involve training data, our non-technical participants do not and ethical sourcing of training data was a major concern for them.

their adaptive, interactive nature when compared to “meditation music”, which is noted to be simplistic, and “trailer music”, which is seen as formulaic by our participants; thirdly, the extent to which the material being generated is melodic – for example, participants communicate melodies as important, and a connection between the developer and composer, and as the most important part of a track, meaning composers are less accepting of melody generation than chord/bass generation).

Our findings resonate with many of the opinions held by practitioners and theorists, such as the inconsistency in quality of output [65]), originality of generative output [13], [73], resource-intensiveness of modern video games [44], the need for robust timing systems [66], and ethical and legal considerations around misuse of creative AI [22].

The participants’ concerns about the originality and quality of output, resource-intensiveness of AIMT, and misrepresentation (in musical style but also of themselves as a composer) in our findings demonstrate that composers have low performance expectancy when it comes to AIMT. Additionally, the concerns we find about ethical and societal misuse and job security, build a picture of negative social influence creating bias against AIMT. Lastly, concerns for their workflow being disrupted and their enjoyment of the process being reduced shows that our participants have high effort expectancy. These factors likely lead to lower acceptance of AIMT.

Our findings extend beyond these ideas, demonstrating that in addition to the concerns described about AIMT, existing tools are likely not being designed in a way that considers the inherent clash between AIMT and creative ego. By designing around this consideration, we can ameliorate the negative impact that social factors play in reducing acceptance, making them less abrasive to creatives, and increasing technology acceptance for AIMT.

With regards to RQ2, we find that the use of AIMT is being stymied by three separate factors. The first is a lack of standardisation in support for music professionals within game studios. While all composers’ needs differ, support varies not only based on the level of the game studio (i.e. AAA/midcore/indie), but across studios at the same level. This lack of standardisation in support mirrors the practitioners’ concerns about lack of audio programmers to build and support tool use in the games industry [66]. Furthermore, this lack of standardisation increases effort expectancy for composers, as the support structure is not in place across the industry to help the less technical creatives to utilise new technology, especially when machine learning or other intelligent technologies such as AI or reinforcement learning are involved.

The level at which a composer is working also seems to affect AIMT acceptance. At the AAA level, game composers are more likely to have assistants or teams to handle tasks, whereas indie game composers do not, and as such the latter were more likely to be accepting of AI tools that can support their needs that an assistant would meet if they could afford one (as this reduces their personal effort expectancy in their role and frees up time for creative work), whereas at AAA companies, human assistants could be handling these tasks already, meaning AAA game composers are potentially less open to these assistive technologies.

A second factor to consider is language. Composers want tools to fit easily into their workflow while providing high quality, original and stylistically appropriate output, however these tools need to be designed around the clash between ego and AI that is described above. Furthermore, participants articulated a desire for tools to communicate “in their language”, as there is a disparity between how game developers and composers communicate, and tools often use quantitative scales as a method of communicating musical features, which differs from the qualitative language that composers may use. This difference in the way that tools and users communicate is disruptive to workflow and requires experimentation from users, increasing effort expectancy, but also reducing perceived performance expectancy if the user misunderstands the language the tool is using. This likely leads to lower acceptance of AIMT [61].

Additionally, there is a lack of standardisation in understanding of and trust in AIMT among composers. This differs from the lack of standardisation in support, as this theme relates more to the idea that the way the tools we design communicate ideas to the end user. Composers are often non-technical; as such, they find AIMT use hard to visualise, especially when it comes to the training of models on data and how models can work with their own provided music. This in turn leads to a lack of trust in AIMT, as it is hard for users to trust what they do not understand. This lack of general understanding of AIMT likely lowers technology acceptance, as it increases effort expectancy and lowers performance expectancy, as composers do not fully understand how AIMT is designed to work [61]. Furthermore, this problem is made more complicated, as there is a lack of consistent terminology within the industry (i.e. procedural as noted by [44]). This lack of consistent language makes it hard for composers/participants to communicate their concerns or problems in a clear way. A final thought on tool design is that often composers relate AIMT to non-AI driven tools that already exist, and when doing so they are more positively disposed toward AIMT. This demonstrates the potential that better understanding can play in increasing trust in AIMT and also reducing the negative impact of social factors [61].

The implications of these findings are that researchers and developers should be designing AIMT to perform paracompositional tasks that can support composers creatively (e.g., expressive rendering or mixing), or add to existing music while guaranteeing consistent high-quality and original output in the style of the user, all while fitting easily in to a composers workflow, if researchers wish for their tool to be used by professionals.

Additionally, in order to increase trust and decrease negative social influence, developers should be very clear about what data they require from the composer in order to generate in their style, and that AIMT should communicate in more qualitative language and be very accessible for users with non-technical backgrounds. Finally, developers would benefit from ongoing communications with creative users during tool development, to increase familiarity with the tool, which will likely lead to increased trust.

Thus, we shed light on how professional game composers

view music AI, while also informing future research and tool design. We would like to underline the importance of designing tools around the end user (creatives), where an effective AIMT necessitates an extra level of communication or negotiation between the end user and the developers regarding features, or language use to increase acceptance, even at the cost of difficulty for the developer. We finish by outlining some limitations in our approach, and ideas for future work in this domain.

A. Limitations

1) *Researcher Bias*: A weakness of this research is that only one researcher conducts the RTA. Without a second researcher to conduct the data analysis, we are unable to use Cohen's Kappa to measure for agreement in the codes/themes. This means that there is likely some researcher bias affecting the results of this study. However, due to the nature of RTA, the role of researcher bias on the results is expected to some degree in their identification of codes and themes (with *identification of codes and themes* being used to show that the researcher plays an active part in the process of creating themes and codes, when compared to the commonly used *emerging themes* of other qualitative approaches – which disregard the active role that researcher bias plays in the creation of themes). Furthermore, RTA as an approach does help to mitigate this issue somewhat, as candidate themes are compared against the uncoded original data, allowing the opportunity to review the themes, ensure that they explain the data and see any potential bias.

2) *Participant Bias*: As noted in the introduction, music (more generally, creative) AI is a controversial topic among music-making communities, especially in the advent of tools such as Midjourney and ChatGPT. The participants we did manage to recruit could still be biased against the use of AI in music making, but they are perhaps less biased and more open-minded to the possibility than composers who refuse to participate, one of which told us they would not consider discussing it with us, such was their impression of the level of controversy surrounding the topic.

3) *Representation Bias*: While the spread of indie-AAA composers in this study is somewhat representative of the spread of work self-reported in industry surveys (GSC), and every effort has been taken to recruit as many AAA composers as possible, there are only two participants that compose music for AAA titles. This means that what has been seen here may not be truly representative of the views of AAA composers, and that some findings may be over generalised, such as AAA composers being less open to the use of AI by having assistants, which has been the case in these interviews.

B. Future Work

Future research could look to better-understand how AIMT affects creativity/productivity among non-novice composers, which would offer valuable insights into the value of AIMT. By using an experimental design that allows participants the opportunity to work with AIMT over multiple sessions, research could evaluate how participants' familiarity with

AIMT affects productivity, while also gaining insight as to how acceptance changes as composers get more familiar with tools.

Additionally, further research could explore the clash between creative ego and AIMT, by allowing participants to work with AIMT that generate different musical elements (i.e. percussion, harmony, melody, bass) and having them grade the output quality of the content, with aims of improving our understanding of how AIMT acceptance scales based on the type of musical content being generated.

Finally, further research could perform analyses of existing AIMT such as Infinite Album, AIVA, the Magenta Suite, in order to establish what the common practices are within the industry, such as the interface designs, the language the tools use, and whether the tool is para-compositional, or replaces the composer's work. This could lead to a better understanding of how our findings could be applied in industry.

VI. CONCLUSION

In this paper, we outline five themes that constitute a novel understanding of the answer to the question "Why has AIMT not been adopted more widely in video games or by professional composers?". This is the beginning of empirical research into an under-researched topic, where we place composers in the focus and ask them about their thoughts and feelings on music AI. We find that composers have multi-faceted concerns, and that the creative ego is not always factored in by the developers of these tools. By providing a standardised support structure within the industry, as well as working towards making AIMT more understandable and consistent in its attendant terminology, we may be able to mitigate these issues in future.

ACKNOWLEDGMENT

This work was supported and funded by the EPSRC Centre for Doctoral Training in Intelligent Games & Games Intelligence (IGGI) [EP/S022325/1].

REFERENCES

- [1] Abel, J (2021). Genesis Noir: Bringing Jazz Improvisation To Gaming. [Online]. Available: <https://www.youtube.com/watch?v=enccMn3LObU>
- [2] Amato, A. (2017). Procedural Content Generation in the Game Industry. In Korn, O. and Lee, N., editors, *Game Dynamics: Best Practices in Procedural and Dynamic Game Content Generation*, pages 15–25. Springer International Publishing, Cham.
- [3] Barriga, N. A. (2019). A short introduction to procedural content generation algorithms for videogames. *International Journal on Artificial Intelligence Tools*, 28(02): 1930001. Publisher: World Scientific.
- [4] "Bloodborne," FromSoftware, Tokyo, Japan, 2015.
- [5] "Borderlands 3," Gearbox Quebec, Quebec, Canada, 2019.
- [6] Bosque, G. (2022). Cinematic Animation in RPGs: Motion Capture and Emotion Recognition. Intelligent Games and Game Intelligence Seminar.
- [7] Braun, V. and Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2):77–101. Publisher: Routledge.
- [8] Braun, V. and Clarke, V. (2013). *Successful Qualitative Research*. Sage Publications Ltd, London: UK.
- [9] Brown, D (2019). The Dynamic Percussion System: A procedural music engine for video games. ImproTech.
- [10] Bucklew, B. and Grinblat, J. (2019). Math for Game Developers: End-to-End Procedural Generation in Caves of Qud. [Online]. Available: <https://www.youtube.com/watch?v=jV-DZqdKlnE>.

- [11] Buraga, A. (2022) The Emergence of the Type-Generated AI Art Community: A Netnographic and Content Analysis Approach. Master's Thesis. Malmö University, Sweden.
- [12] Collins, K. (2009). An Introduction to Procedural Music in Video Games. *Contemporary Music Review*, 28:5–15.
- [13] Collins, T. and Laney, R. (2017). Computer-Generated Stylistic Compositions with Long-Term Repetitive and Phrasal Structure. *Journal of Creative Music Systems*, 1(2). Number: 2.
- [14] Compton, K. (2016). So you want to build a generator... [Online]. Available: <https://galaxykate0.tumblr.com/post/139774965871/so-you-want-to-build-a-generator>.
- [15] Cutajar, S. (2020). *Automatic Generation of Dynamic Musical Transitions in Computer Games*. PhD Thesis, The Open University.
- [16] Davis, F. D. (1989). Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology. *MIS Quarterly*, 13(3):319–340. Publisher: Management Information Systems Research Center, University of Minnesota.
- [17] D'Errico, M. (2015). Worlds of Sound: Indie Games, Proceduralism, and the Aesthetics of Emergence. *Music, Sound, and the Moving Image*, 9(2):191–206.
- [18] "Destiny 2," Bungie Inc, Washington, USA, 2017.
- [19] Engels, S., Tong, T., and Chan, F. (2015). Automatic real-time music generation for games. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 11. Issue: 1.
- [20] Etienne Carrier (2018). Procedural World Generation of Far Cry 5. | [Online]. Available: <https://www.youtube.com/watch?v=JBp8zLVsgg>
- [21] "Far Cry 5," Ubisoft, Saint-Mandé, France, 2018.
- [22] Flick, C. and Worrall, K. (2022). The Ethics of Creative AI. In Vear, C. and Poltronieri, F., editors, *The Language of Creative AI: Practices, Aesthetics and Structures*, Springer Series on Cultural Computing, pages 73–91. Springer International Publishing, Cham.
- [23] Greuter, S. and Nash, A. (2014). Game Asset Repetition. In *Proceedings of the 2014 Conference on Interactive Entertainment*, IE2014, pages 1–5, New York, NY, USA. Association for Computing Machinery.
- [24] Grimblat, J. (2018). Procedurally Generating History in Caves of Qud. [Online]. Available: <https://www.youtube.com/watch?v=H0sLa1y3BW4>.
- [25] Guo, R., Simpson, I., Kiefer, C., Magnusson, T., and Herremans, D. (2022). MusIAC: An Extensible Generative Framework for Music Infilling Applications with Multi-level Control. In Martins, T., Rodríguez-Fernández, N., and Rebelo, S. M., editors, *Artificial Intelligence in Music, Sound, Art and Design*, Lecture Notes in Computer Science, pages 341–356, Cham. Springer International Publishing.
- [26] Hadjeres, G. (2021). The Piano Inpainting Application | Sony CSL: Etincelle #14.
- [27] Huang, C.-Z. A., Vaswani, A., Uszkoreit, J., Shazeer, N., Simon, I., Hawthorne, C., Dai, A. M., Hoffman, M. D., Dinculescu, M., and Eck, D. (2018). Music Transformer. *arXiv:1809.04281 [cs, eess, stat]*. arXiv: 1809.04281.
- [28] Hughes, S and Korb, D. (2021) HADES Special with Audio Director Darren Korb. The Sound Architect Podcast S04, E01.
- [29] Hutchings, P. E. and McCormack, J. (2020). Adaptive Music Composition for Games. *IEEE Transactions on Games*, 12(3):270–280. Conference Name: IEEE Transactions on Games.
- [30] Jeong, D., Kwon, T., Kim, Y., Lee, K., and Nam, J. (2019a). VirtuosoNet: A hierarchical RNN-based system for modelling. In *International Society for Music Information Retrieval (ISMIR)*, pages 908–915.
- [31] Jeong, D., Kwon, T., Kim, Y., and Nam, J. (2019b). Graph neural network for music score data and modeling expressive piano performance. In *International Conference on Machine Learning*, pages 3060–3070. PMLR.
- [32] Keller, R. M., Morrison, D. R., and College, H. M. (2007). A Grammatical Approach to Automatic Improvisation. In *Proceedings of the 4th Sound and Music Computing Conference*, pages 330–337, Greece.
- [33] Ko, H., Park, G., Jeon, H., Jo, J., Kim, J. & Seo, J. Large-scale text-to-image generation models for visual artists' creative works. *Proceedings Of The 28th International Conference On Intelligent User Interfaces*. pp. 919-933 (2023)
- [34] Kristen, K., Guzdial, M. & Sturtevant, N. The definition-context-purpose paradigm and other insights from industry professionals about the definition of a quest. *Proceedings Of The AAAI Conference On Artificial Intelligence And Interactive Digital Entertainment*. 17 pp. 107-114 (2021), Issue: 1
- [35] Lamperski, P. and Tahouri, B. (2015). Real-time Procedural Percussion Scoring in 'Tomb Raider's' Stealth Combat. [Online]. Available: <https://www.gdcvault.com/play/1023215/Real-time-Procedural-Percussion-Scoring>.
- [36] "League of Legends," Riot Games, LA, USA, 2009.
- [37] Louie, R., Coenen, A., Huang, C. Z., Terry, M., and Cai, C. J. (2020). Novice-AI Music Co-Creation via AI-Steering Tools for Deep Generative Models. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–13, Honolulu HI USA. ACM.
- [38] Marks, A. (2008). *The Complete Guide to Game Audio: For Composers, Musicians, Sound Designers, Game Developers*. Focal Press.
- [39] McKendrick, I. (2017). Continuous World Generation in No Man's Sky. [Online]. Available: <https://www.youtube.com/watch?v=sCRzxEEcO2Y>.
- [40] Moffat, D. AI Music Mixing Systems. *Handbook Of Artificial Intelligence For Music: Foundations, Advanced Approaches, And Developments For Creativity*. pp. 345-375 (2021), Publisher: Springer
- [41] "No Man's Sky," Hello Games, Guildford, UK, 2016.
- [42] Oore, S., Simon, I., Dieleman, S., Eck, D., and Simonyan, K. (2020). This time with feeling: Learning expressive musical performance. *Neural Computing and Applications*, 32(4):955–967. Publisher: Springer.
- [43] Phillips, W. (2014). *A Composer's Guide to Game Music*. Publisher: MIT Press.
- [44] Plut, C. and Pasquier, P. (2020). Generative music in video games: State of the art, challenges, and prospects. *Entertainment Computing*, 33:100337.
- [45] "Rise of the Tomb Raider," Crystal Dynamics, CA, USA, 2015.
- [46] Roberts, A., Engel, J., Mann, Y., Gillick, J., Kayacik, C., Nørly, S., Dinculescu, M., Radebaugh, C., Hawthorne, C., and Eck, D. (2019). Magenta studio: Augmenting creativity with deep learning in ableton live.
- [47] Rogier, Joroen (2022) Better Adaptive Music for Starship Troopers Audio Developer Conference. 14-16 November. London, UK.
- [48] Sander van Rossen (2020). Tools Summit: Geometry in Milliseconds: Real-Time Constructive Solid Geometry. [Online]. Available: <https://www.youtube.com/watch?v=Iqmg4gblreo>.
- [49] Schmidt, B. (2021). Game Audio Industry Survey 2021. Game Sound Con. [Online]. Available: <https://www.gamesoundcon.com/post/game-audio-industry-survey-2021>.
- [50] Shaker, N., Togelius, J., and Nelson, M. J. (2016). *Procedural content generation in games*. Publisher: Springer.
- [51] Shapiro, I. and Huber, M. (2021). Markov Chains for Computer Music Generation. *Journal of Humanistic Mathematics*, 11(2):167–195.
- [52] Somberg, G. (2018) Game Audio Programming Audio Developer Conference. 19-21 November. CodeNode, London, UK.
- [53] Summerville, A., Snodgrass, S., Guzdial, M., Holmgård, C., Hoover, A., Isaksen, A., Nealen, A. & Togelius, J. Procedural Content Generation via Machine Learning (PCGML). *IEEE Transactions On Games*, 10, 257-270 (2018,9), Conference Name: IEEE Transactions on Games
- [54] Sterne, J. and Razlogova, E. (2019). Machine learning in context, or learning from landr: artificial intelligence and the platformization of music mastering. *Social Media+ Society*, 5(2):2056305119847525. Publisher: SAGE Publications, Sage UK: London, England.
- [55] Stevens, R. and Raybould, D. (2015). *Game Audio Implementation: A Practical Guide Using the Unreal Engine*. Publisher: Focal Press.
- [56] Togelius, J., Yannakakis, G., Stanley, K. & Browne, C. Search-Based Procedural Content Generation. *Applications Of Evolutionary Computation*. 6024 pp. 141-150 (2010),
- [57] Tsiros, A. & Palladini, A. Towards a Human-Centric Design Framework for AI Assisted Music Production. *Proceedings Of The International Conference On New Interfaces For Music Expression*. (2020)
- [58] Vanka, S., Safi, M., Rolland, J. & Fazekas, G. (2023) Adoption of AI Technology in the Music Mixing Workflow: An Investigation. [Online]. Available: arXiv:2304.03407.
- [59] Venkatesh, V. (2000). Determinants of Perceived Ease of Use: Integrating Control, Intrinsic Motivation, and Emotion into the Technology Acceptance Model. *Information Systems Research*. Publisher: INFORMS.
- [60] Venkatesh, V. and Davis, F. D. (2000). A Theoretical Extension of the Technology Acceptance Model: Four Longitudinal Field Studies. *Management Science*, 46(2):186–204. Publisher: INFORMS.
- [61] Venkatesh, V., Morris, M. G., Davis, G. B., and Davis, F. D. (2003). User Acceptance of Information Technology: Toward a Unified View. *MIS Quarterly*, 27(3):425–478. Publisher: Management Information Systems Research Center, University of Minnesota.
- [62] Vimpari, V., Kultima, A., Hämäläinen, P. & Guckelsberger, C. "An Adapt-or-Die Type of Situation": Perception, Adoption, and Use of Text-To-Image-Generation AI by Game Industry Professionals. *ArXiv Preprint ArXiv:2302.12601*. (2023)

- [63] Washburn, M. and Khosmood, F. (2020). Dynamic Procedural Music Generation from NPC Attributes. In *International Conference on the Foundations of Digital Games*, pages 1–4, Bugibba Malta. ACM.
- [64] Watkins, R. (2016). *Procedural Content Generation for Unity Game Development*. Publisher: Packt Publishing Ltd.
- [65] Weir, P. (2017). The Sound of 'No Man's Sky'. | GDC Vault. [Online]. Available: https://www.youtube.com/watch?v=zKJ_XuQijiw
- [66] Weir, P., Timea Farkas, Kyle Worrall, and Alessandro Coronas (2021). Procedural Audio in Narrative Game Design Panel. [Online]. Available: <https://www.youtube.com/watch?v=AllYuKKxks8>
- [67] Wilde, M (2014). *Audio Programming for Interactive Games*. Publisher: Routledge.
- [68] Wilson, C (2019). Designing 'Path of Exile' to Be Played Forever. [Online]. Available: <https://www.youtube.com/watch?v=tmuy9fyNUjY>
- [69] Wooller, R., Brown, A., Miranda, E., and Diederich, J. (2005). A framework for comparison of process in algorithmic music systems. *Generative Arts Practice*, pages 109–124. Publisher: Creativity & Cognition Studio Press.
- [70] Worrall, K. (2021). Remaking Music for Modern Sensibilities: A case study in the evolution of music design for a more critical audience. | North American Conference on Video Game Music.
- [71] Worrall, K., Yin, Z., and Collins, T. (under review). Cue Free Express: A cue-free model for expressive music performance.
- [72] Xiang, Z. and Guo, Y. (2021). Controlling Melody Structures in Automatic Game Soundtrack Compositions With Adversarial Learning Guided Gaussian Mixture Models. *IEEE Transactions on Games*, 13(2):193–204.
- [73] Yin, Z., Reuben, F., Stepney, S., and Collins, T. (2021). "A Good Algorithm Does Not Steal – It Imitates": The Originality Report as a Means of Measuring When a Music Generation Algorithm Copies Too Much. In Romero, J., Martins, T., and Rodríguez-Fernández, N., editors, *Artificial Intelligence in Music, Sound, Art and Design*, volume 12693, pages 360–375. Springer International Publishing, Cham.

BIOGRAPHIES

Kyle Worrall received his MSc degree in Sound and Music for Interactive Games from Leeds Beckett University in 2016. He is a PhD researcher with the Centre for Doctoral Training in Intelligent Games and Game Intelligence, based in the Music Computing and Psychology Lab at the University of York, specialising in music generation and expressive rendering for video games.

Tom Collins is an Associate Professor in the Frost School of Music at the University of Miami, and is also heavily involved in the music cooperative MAIA, Inc. Tom studied Music at Cambridge, Mathematics and Statistics at Oxford, and did his PhD on automatic pattern discovery and music generation at the Open University. He has held multiple positions in the US and Europe.

VII. APPENDIX

This appendix includes the questions presented to the composers throughout the semi-structured interview (Fig. 1) and the demographic information of the participants (Fig. 2).

Semi-Structured Interview Questions

1. How long have you been working in video game music? If you have experience in music for films or linear media, then how long have you done that?
2. What specific considerations do you have in mind when working on game music?
3. Do you have any concerns when it comes to how long a player may be listening to your music during extended gameplay segments and repetition?
4. Please describe your workflow when it comes to composing music for games.

Prompt:

- *What macro-tasks would you say makes up the process of creating a piece of music? I.e. drafting, composing, mixing, mastering etc.*
- *What smaller micro-tasks would you say that these larger tasks can be broken up into when describing your workflow?*
- *Do you spend time writing various versions of musical sections due to the nonlinear nature of video games?*

5. How much if any experience in programming (coding) do you have? And if you tend to have a hand in how your music is implemented, please give details on the process.
6. Do you have any opinions or concerns relating to procedural music systems in games? Please give details. And your thoughts on Music AI?
7. If you look at the list of things that AI can be used for in Music (on the attached document), what do you think about each of these ideas/demos?
 - a. Melody Generation
 - b. Texture/Pad Generation
 - c. Chord Generation
 - d. Musical In-Painting
 - e. Automatic Expression for Musical Passages
 - f. Generating Small Ensembles to support melody
 - g. Generating your music in different styles or genres?
8. Is there a specific aspect of music composition for which you would be most open to AI assistance?
9. What do you think about the usability of the interface designs for the tools shown on the attached document?
10. Would you be willing to participate in another study in the future where you get to try out a music AI tool so we can assess how they impact workflow?

Fig. 1. The questions that were used in the semi-structured interviews with the professional game composers.




Fig. 2. Demographic Information of Participants.