From Genies Performing Magic to Sages Imparting Wisdom: A Value-Centred Survey of Music AI User Interfaces, Creative Affordances and Artist Objectives

Oliver Bown
Creative Technologies Research Lab
UNSW Sydney
Australia
o.bown@unsw.edu.au

Abstract

There are, at present, multiple fronts of debate about the ethics of Commercial Creative AI (ComCAI) product development. The issues of copyright and the right to know the AI provenance of work are dominant and well-defined. Meanwhile, a significant area of debate is trickier to define and analyse, relating to ComCAI's impact on the intrinsic values and virtues of human creative endeavour. In this paper I analyse this impact by surveying the emerging landscape of user interfaces (UIs) in the domain of commercial music AI products. I do this through a value-centred design lens, considering how it is possible to articulate intrinsic values in the arts. Specifically, I select two artistic objectives: the need for creative artists to create works that are differentiated from others; the need for deep cognitive creative engagement with the work being created. I ask how ComCAI UIs can enable these artistic objectives, first through a survey of UI types and their combination in products, second through user feature requests. This design analysis tentatively suggests that those ComCAI tools promising to support artists increasingly need to provide greater control, in order to satisfy their creative users, and this is focused towards either simply allowing post-generation editing, or improving generative control, which is in tension with the opaque nature of the generative systems. I propose that a value-centred design approach that highlights values such as differentiation and creative engagement helps to identify under-explored design spaces which have the potential to make better music AI experiences.

Keywords: Music AI, Commercial Creative AI, Value-Centred Design, User Interface Design.

Introduction

In the short space of time since the boom in Creative AI started, there are now numerous commercial products offering different content creation services; researchers in the field can shift from speculative explorations of Creative AI interfaces to field-studies of existing Creative AI tools in the hands of users (Holzapfel, 2023; Kaila et al., 2024). Simultaneously, public debate is grappling with concerns about the social impact of Creative AI (Bown, 2023).

This paper undertakes a form of value-centred design analysis (loosely inspired by (Cockton, 2005), and closely aligned to critical studies of Creative AI (Holzapfel, 2023; Kaila et al., 2023; Flick and Worrall, 2022; Rezwana and Maher, 2023; Piskopani et al., 2023)), reviewing the values being adopted in the Commercial Creative AI (ComCAI) sphere, and associated values being more widely articulated in public debates. In doing so I hope to expand critical discussion of commercial practices in AI applications in creative spheres (such as (Drott, 2018, 2021; Hodgson, 2021, 2020; Born, 2022)) to consider how ComCAI tools serve the needs of creators.

The contemporary commercial environment make this particularly important: Earlier technological revolutions such as photography, sound recording and synthesis may have been commercially driven but did not necessarily involve the billions now invested in start-ups looking to achieve platform status or serve existing platforms. The rise of social media, search and streaming services to create an attention economy, with negative social impacts such as social media addiction and filter bubbles, may seem far removed from art creation. But the current generation of ComCAI start-up companies appears closer aligned with modern platform business models than with tradition producers of creative tools, and with associated societal concerns (Crawford, 2021; Doctorow, 2023).

The most prominent ethical debates about ComCAI values have focused on, firstly, protecting copyright holders from exploitation by machine learning products (Zirpoli, 2023), and, secondly, the rights of people to know when AI has been used (for example as championed by the Human Artistry Campaign[1]). Yet public debates about ComCAI also commonly involve concerns that using Creative AI tools could undermine the forms of intrinsic value, both individual and social, that we place on artistic practices and experiences. This is an area that is much less well understood, thus harder to develop a shared conceptual framework for debate and evaluation, let alone formally quantify (though there is much work on the subject such as (Hesmondhalgh, 2013; Born, 2010)). Accordingly, companies' claims about enhancing creative expression and creative engagement, democratising access to creative fields, and advancing the value of creative arts remain inherently ungrounded and speculative.

[1] www.humanartistrycampaign.com

A common example of an ungrounded ComCAI value statement is the concept of "democratisation", which in this context means increasing the number of people who have access to a certain creative activity (as in the language of Adobe's CEO (Kumparak, 2020)). Others make appeals to "enhancing artistic creativity". Such claims can easily be problematised. Besides the fact that the usage of the term "democratisation" is etymologically inconsistent – democracy concerns the shared *control* of a system by people, whereas its use in ComCAI concerns increasing participation in restricted domains of activity – providing new automated means for art creation available to all doesn't have any self-evident bearing on outcomes such as artistic success, satisfaction, expression or engagement. ComCAI tools are unlikely to increase the number of economically successful artists, since that is set by other factors such as available streaming revenue (indeed it is more likely to diminish it by delivering more competing content). Speeding up creation seems of dubious value, except for established creators, and the refrain from critics has now become common that, in the viral words of author Joanna Maciejewska, "I want Al to do my laundry and dishes so that I can do art and writing, not for Al to do my art and writing so that I can do my laundry and dishes"[2]. And since ComCAI tools seem as likely to take away creative control from creators as to empower those creators, both differentiation and creative engagement seem as likely to be diminished as enhanced.

For the sake of imagining how these qualities are experienced in practice, consider abstract expressionist painting, as practiced by amateur artists. The ability to produce a completed artwork in an abstract expressionist style is arguably not where any barrier to entry into that culture of art practice lies. The values associated with this artworld include seeking one's own unique expression of abstraction in paint, undertaking to develop a personal relationship with the materials and processes, extending and contributing to an honoured tradition.

Arguably, then, ComCAI tools are subject to a form of "Red Queen effect" (Cliff and Miller, 1995): they neither democratise nor enhance creative practice, but shift the conditions under which people seek value from creative arts, finding new ways to achieve differentiation and creative engagement. At worst, the result is the opposite of what is claimed: instead of democratisation, we have "dullocratisation", art creation that is banal, disengaged and wasteful, perhaps even reproducing some of the negative effects of "doom scrolling" and fragmented engagement in cultural, artistic material (Dorsen, 2022). As a useful provocation, dullocratisation conjures images of artists generating art to post online, hardly even engaging with it themselves, or engaging with it only as surface material, at the expense of other frames of engagement, such as the underlying logic of form, process and relational context.

Putting this together (Figure 1) in an albeit highly simplified spectrum, debates seem to contrast positive "democratising" thinking with concerns of "dullocratisation", with forms of "Red Queen effect" thinking forming the middle ground.

Whichever way, claims of democratisation of creative expression by current ComCAI companies are easily dismissed as tokenistic and lacking any concerted interest in that goal, except indirectly through individual stories of creative empowerment. This points to the focus of the current paper. If the individual experiences of creators are the basis for developing the social value cases for ComCAI tools, then in this paper I seek to develop a richer analysis of what values could be expressed and identified as design goals, and then go on to consider how well these values are supported in current tools, specifically in interface design decisions.

Differentiation and Creative Engagement

Given the above, I consider two perspectives related to intrinsic creative value that, despite being poorly understood, can be conceptually connected to the affordances and design considerations of ComCAI products. The hope is they enable more specific concepts that serve to analyse how creative practices using ComCAI tools are emerging and relating to value claims. The first comes from the now widely recognised theoretical groundwork laid most famously by Pierre Bourdieu, for whom a person's appreciation of creative arts is not merely a matter of individual pleasure but acts relationally: "Social subjects, classified by their classifications, distinguish themselves by the distinctions they make" (Bourdieu, 1984). For Bourdieu, the classical music connoisseur's deep engagement in the artform is lived and real, but also affords that person the ability to differentiate themselves socially. In the studies of French society that he undertook, this was about culture serving to reify class distinctions, but such thinking has since been extended widely. Individual subgenres of music can be seen to align with social groups (North and Hargreaves, 1999), and musical taste can be seen to inform judgement of and affiliation with others (Laplante, 2012). Meanwhile, there is a jostle *within* those groups for individual differentiation: individual creative actors compete within creative fields, choosing different strategies that depend on their positionality. There are extensive critical debates about Bourdieu's theories of taste (such as (Born, 2010)). For the purpose of this paper, it is not necessary to go into this detail: merely to accept, broadly, idea that a contemporary creative music practice necessitates the creation of difference or distinction, if not originality.

The second perspective is captured by McCarthy et al. (2001) as "cognitive growth" and has been treated in a number of ways by different scholars. Here the subject is more vague and dispersed, and I present a personal interpretation that draws on a set of related ideas, under the less spiritually loaded title "creative engagement". McCarthy et al. (2001) differentiate "cognitive growth" from the more immediate intrinsic value of "captivation". An art experience can be immediately captivating: a musical groove or melody triggers a pleasure response. But it can also lend itself to more enriching experiences, such as when a complex piece of music affords repeated listening to understand its structure. Theories of aesthetic perception in music have long attempted

[2]https://x.com/AuthorJMac/status/1773679197631701238

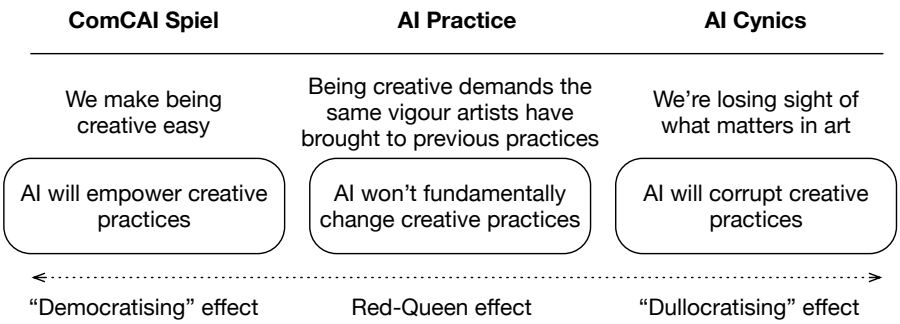

Figure 1: A simplified representation of a spectrum of positive, neutral and negative attitudes to the benefit of AI tools for creators. The neutral position presents overall creative benefits as invariant to the specifics of a technology.

to explain complexity in music in the idea that musical experience involves a complex interaction between cognitive domains. Huron (2006), for example, uses the two-system model of the brain – one fast, low-level and automatic, the other slow, conscious and reflective – to build a model of music experience in which these systems interact in a complex interplay of expectation and revelation. Meanwhile, the act of creative production, as described by numerous accounts of creative producers themselves (a popular introduction can be found in (Boden, 1990)), is an enriching experience in itself beyond the achievement of some creative product; an investigation or study of form or concept that provides immediate pleasure and arguably enhances one's understanding of the world.

In this paper, there is no attempt to think of differentiation and creative engagement in quantitative terms as suggested by a peer reviewer. But it is important to note that these aspects of a creative practitioner's outlook could conceivably be measured, and hence connected to attempts to quantitively measure stylistic difference as well as properties relating to the experience of creative aesthetic activities, as in the study of flow. These could subsequently be related to the goals and positionality of users. However, the present work is primarily aimed at mapping respective design and value spaces and to draw qualitative associations between these spaces.

Arguably, debates about the value of the arts are easily dominated by professional arts revenues and quality, and neglect the huge social and personal importance of these dimensions of creative arts practice. But both topics are invoked in public debates about AI's impact on the creative arts. In (Bown, 2023) I have argued that AI is seen by different communities both positively and negatively in the case of creative engagement, becoming a point of fissure, potentially diminishing creative engagement through alienating people from the work of creation, or enhancing it through new types of creative work or output. It may do both in different contexts. A range of critical positions against AI claim that it invites banality, an impoverished version of art creation, with people feeling that they can be creators without the hard work and nurturing of talent until recently required for creation (for one of many discussions in mainstream discourse, see (Shaffi, 2023)). Defenders of AI claim that it can't and won't undermine these values, because real artists will use it to push new boundaries (for one of many examples see the Dadabots keynote at the AI Music Creativity conference). Both differentiation and creative engagement are implicated in such discussions, usually indirectly.

Companies such as Boomy[3] and Splash[4], whose business is engaging non-musicians in new forms of musical creation, engage with concepts of music's value that can become a focal point for this public debate. Both emphasise connection with other music makers and audiences, and the idea of supporting people to become engaged and proactive members of a musical culture. Yet arguably both may actually diminish the underlying factors responsible for achieving this social engagement, for example in taking away individual creative expression and the understanding a creator has of the material they have created: what good is making something if you don't have anything to say about it? These business models pose many questions about creative engagement. How do factors such as effort, the nurturing of skill, or the search for unique forms of self-expression play into musical cultural connection? How much does it matter that creators understand the principles underlying the things they create? How is the intrinsic reward of creation impacted by the automation process, for example through one's ability to control specific details?

These are questions that seek to get to the heart of what music creators and communities value. They are slightly different in this sense from more familiar questions of use qualities (Löwgren, 2006) or even questions of design principles for creativity support tools (Shneiderman et al., 2006). This is because of the radical discontinuity that ComCAI tools present. Designers designing creativity tools build on an existing culture of practice. If improving a Digital Audio Workstation or a synthesiser, for example, they may think in terms of improving efficacy or discovery. But ComCAI tools' radical reconceptualisation of what it means to cre-

[3] boomy.com

[4] www.splashmusic.com

ate presents more fundamental questions of motivation. In this sense a value-based approach is different from a user-centered approach: it seeks to understand what purposes such tools should even have, let alone how well they achieve these purposes.

With rising communities of practice using Creative AI tools, we can now begin to look at how these dimensions of intrinsic value play out in the interaction between individual users' creative practices and the specific algorithmic affordances and user interface (UI) design concepts of ComCAI products. A plausible prediction, given the above theoretical themes, is that creative users will seek the means of differentiation and creative engagement from these tools. This conforms to the "defenders of AI" argument above, and the "Red Queen effect" take, according to which, since these are such foundational aspects of artistic behaviour, Creative AI tools won't dampen them, but neither will it help enhance them. Instead, creative communities will bend the use of Creative AI tools such that these qualities continue. Alternatively, such tools will reveal a range of different attitudes and creative strategies with relation to differentiation and creative engagement, depending on the user's situation.

Surveying Creative Value Affordances in ComCAI User Interface Designs

The evolving user interfaces of ComCAI tools, and more generally their inherent or designed affordances, show an emerging landscape of interaction possibilities that provides a useful starting point for analysing their changing creative use. Through a survey of ComCAI UI designs I focus on the rise of two interface types – path and prompt-based interfaces – what their real and perceived affordances are, and how they are being developed to interact with other interface types.

Paths and Maps

A "path", as popularised by Bill Verplank in his Interaction Design Sketchbook (Verplank, 2009), is a design pattern that can be used when a user needs to be guided through a specific set of steps. The pattern can be found in many different kinds of software application, with common uses including selecting dates and locations for a flight booking, entering delivery and payment details for a purchase, or setting up a user account. Another term sometimes used for this kind of pattern is a "wizard", more commonly associated with tools that helps a user set-up or configure a series of settings (I generally stick with "path" in this paper). Design consultants Nielsen Norman describe the value of paths/wizards as follows:

> "Usually, people feel better when they are in control, but having too much freedom might get the user into trouble and will often cause them to waste time. Wizards allow the computer to control the flow of the dialogue, but this limitation of users' freedom can be liberating in cases where people don't care about their choices or don't know enough to make a decision" (Budiu, 2017).

Paths are found everywhere, but until the emergence of Creative AI tools, paths had little or no role to play in creative production software except in menial tasks such as exporting or importing data. The above quote indicates why this might be germane to generative AI. By contrast, software to support creative work is typified by "maps", Verplank's antithesis to the path, where a user has oversight and simultaneous access to a vast range of interface elements. Maps give you the full picture, enabling freedom of action. They allow multiple ways of doing the same thing, and enable many different things. They enable different artists to use creative tools in radically different ways. These are well-established values in the design of creativity support tools (Shneiderman et al., 2006). Maps bring the user closer to the forms of "direct manipulation" most commonly associated with creative work, the immediate manipulation of the creative object. In the digital world creative work largely takes place under the WIMP (windows, icons, menus, pointer) paradigm through map interfaces that offer a range of direct manipulation tools. In other words, maps are important in analysing current creative AI tools due to their apparent disappearance and the subsequent question: are they no-longer essential to creative software tools.

The rise of path-based interfaces in creative work is therefore surprising, in one sense, since in limiting user control and freedom, it is a radical departure from common design principles for creative tools. But the path is a natural, perhaps unavoidable, fit for automated creation tasks. Interfaces are highly constrained by the affordances of a technology, but they also offer many design choices that influence how that technology is used and experienced. It would be mistaken to think that creative artists and musicians always want freedom. On the contrary, numerous accounts, designs and theories support the idea that effective creative work proceeds through the decisive choice of constraints (Stokes and Fisher, 2005; Stokes, 2009; Biskjaer and Halskov, 2014). More generally, issues such as cognitive load can also mean that something that apparently enhances freedom, is actually limiting. As part of their famous design minimalism, Apple's mouse had one single button, not two or three, because they recognised the cognitive burden outweighed the interaction possibilities. Similarly, paths actively remove options.

Nevertheless, before the very recent rise of Commercial Creative AI, it is hard to think of examples in creative work where a path-based activity plays a significant role in the creation of content (a possible exception is in preset selection, such as choosing synthesiser sounds, where one chooses a bank of sounds and then a specific sound). In certain ComCAI applications, path-based interfaces have emerged as a popular design, and in many cases, the only conceivable design (due to the opacity of the generative process). For example, very similar music-generation paths can be found in at least 4 products (Beatoven[5], Bandlab[6], AIVA[7], Boomy[8]),

[5] www.beatoven.ai
[6] www.bandlab.com
[7] www.aiva.ai
[8] boomy.com

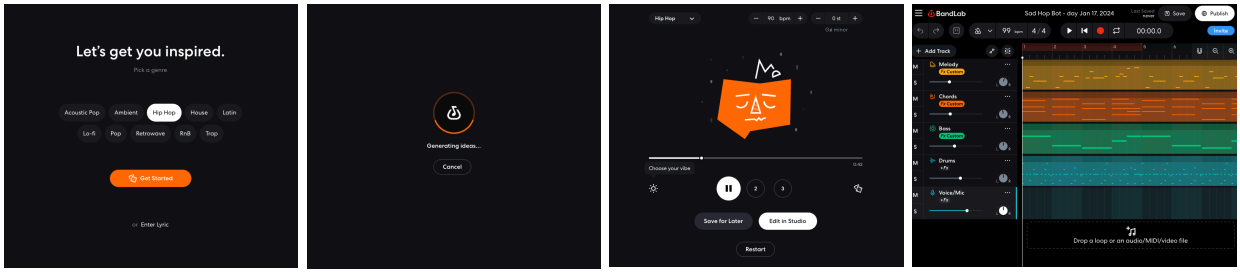

Figure 2: The Bandlab SongStarter generative music tool allows a user to generate original music by selecting only a genre, tempo and key. This can then be loaded into the Bandlab editor as a multitrack piano-roll project.

with important interface design variations. Exemplifying this design pattern, the AI "SongStarter" tool, found in the Bandlab online digital audio workstation (DAW), provides an interface for users to seed a composition in Bandlab using AI generated content (Figure 2). Users begin by selecting a genre. They can then audition three generated options. They can save any of these for later and then either regenerate, varying the genre selection, tempo and key if they wish, or they can go on to select a track to be opened in the Bandlab editor.

Tables 1 and 2 compare the options and interface design decisions available in path-based interfaces in four different ComCAI music products: Boomy, Bandlab, AIVA and Beatoven. In each product, a path-based generation process lies at the centre of the product's workflow, but each product offers very different variations on how that path is designed and integrated into other interface elements. The tables show the diversity of design considerations despite the common core, with different path control variables available. In Table 2 we see a range of approaches to the issue of how to enable users to edit via direct manipulation the outputs of the generation process.

Browsers

A browser, as used here, is an interface for searching or exploring a large corpus of existing material. Browsers are relevant in Creative AI because there is a continuum between browsing the different candidate outputs of a generation process (in many cases a generative process will produce a handful of options) and browsing other sources, such as human created work. Copyright free music AI applications like Mubert[9], Soundful[10] and Soundraw[11] have all integrated their generative processes into a browser-type interface, situating the act of generation within the broader context of searching or browsing for readymade options. Figure 3 shows Soundraw's browser interface, which almost conceals the generative process behind the browser. The user selects blends of genres and moods and the browser list regenerates on the fly. It is interesting to compare these tools, which are targeted at non-musicians in search of music (e.g., for a video or company branding) to the previously discussed tools Boomy, Bandlab and AIVA, which are more musician-focused. Beatoven is a bit different because it is also a copyright free music product (targeted at film-makers) but has chosen not to make a browser interface central.

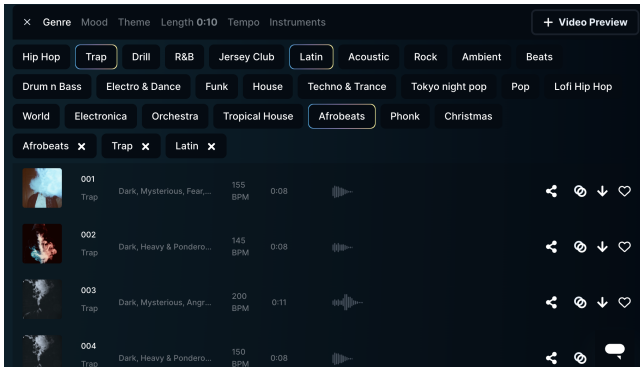

Figure 3: Soundraw's generation process is directly integrated with a browser. The user selects multiple tags representing genre, mood and theme, from which a selection of generated tracks is produced. This diminishes the perceived distinction between searching for music and generating music. The application domain is rapid and cheap access to copyright free music.

Prompts and Dialogues

Before the arrival of successful multimodal generative AI tools around 2021, prompt-based interfaces didn't exist in commercial tools (unless you consider search bars in the same category). The rise of deep-learning text-to-image/audio/video/etc. algorithms had led to the very rapid emergence of prompting as a creative practice and a standard interface. The prompt is an interesting new interface modality in that it implies the potential to do anything, like a genie (Plotkin, 2009) (complete user control), but offers no visibility into the potential things that can be done. Thus whilst maps attempt to reveal to the user the control they afford, and paths clearly reveal how they limit user con-

[9] mubert.com
[10] soundful.com
[11] soundraw.io

Table 1: Four products that use path interfaces to manage a generation task. The table indicates whether the produce includes the given user interface options. For genre/style, mood, instrumentation (inst), tempo, key and duration (dur), this refers to whether the path-interface includes these options. ”Iterable” refers to whether you can iterate (i.e., re-generate) on a given generated output. ”Browser” refers to whether the generation process results in a list of options. ”Target/input/prompt” refers to whether there is any way to steer the output to a reference track or descriptor or other form of target. These four products were selected because of their similar objective of supporting the creation or creative ideation process, and similar path-based generation process.

Product	**Genre/Style**	**Mood**	**Inst**	**Tempo**	**Key**	**Dur**	**Iterable**	**Browser**	**Target/Input/Prompt**
Boomy	+	-	+	+	-	-	+	-	-
Bandlab	+	+	-	+	+	-	-	+	-
AIVA	+	+	+	+	+	+	+	+	+
Beatoven	+	+	+	+	-	+	+	+	+

Table 2: For the same four products, how does the tool provide the user with an option to edit the resulting output? Any additional notes on the design of the interface are included in the final column.

Product	**Direct Manipulation Editor**	**Other Interface Design Notes**
Boomy	Simple rearrangement editor	Set densities, drum sounds, mix style
Bandlab	Full DAW editor	Mood represented as day/night, browser gives 3 options
AIVA	Built in simplified DAW	Allows different starting points such as ”from a style”, ”from an influence”, large array of meta parameters for generation (such as Harmonic Repetition).
Beatoven	Only edit intensity and mix points	Allows text prompt as an alternative to going through path.

trol, prompts risk the misleading user experience of offering complete control but not actually affording it. For example, I may prompt a text-to-audio system to "play Stairway to Heaven on a kazoo". The system may generate kazoo sounds playing a mournful melody but fail to incorporate the exact melody from the song (the two high-profile tools I tried didn't get close to producing a kazoo sound). This offers an example where, despite the elegant simplicity of a prompt-based interface, a path design may be preferable; to more honestly portray the capabilities of the text-to-audio interface, one could replace the freeform prompt with lists of options for instrumentation, genre, and so on. Indeed, conversational interfaces such as customer assistant chatbots often incorporate path elements into conversations as a way to improve user awareness of possibilities ("do you want to speak to sales, technical support or accounts? Please select one").

Prompt-based interfaces naturally imply new forms of *dialogic* interaction (Bown et al., 2020), whereby a creative user and an AI system might iteratively exchange creative artefacts and text, moving towards a creative goal. At the time of writing, dialogic interaction is poorly understood. Confounding the problem of prompts not accurately interpreting the intent of the prompter, dialogues promise to allow refinement through iteration. Generative AI tool makers have promoted the dialogic capabilities of their systems. For example, with their image generation capability built into Facebook and WhatsApp, Meta suggest to users that they can request refinements of generated images. Such refinement promises a solution to the randomness of prompt-based generation, with the creative user specifying changes to a generated image. Dialogue can also support exploration, such as when the creative user requests an array of diverse suggestions, or the system proposes new ways to reframe the creative task. But whilst the dialogic interaction paradigm now exists in embryonic form in tools such as Meta's and OpenAI's GPT-DALL-E integration, AI systems seem particularly bad at refinement, underlying which are severe limitations in systems' abilities to understand creative directives and objectives (Bown, Grace, Ocampo and Ibarolla, forthcoming).

Prompts may afford or at least appear to afford significantly greater user freedom than paths. They do not satisfy Nielsen's definition of wizards quoted above, with its focus on the reduction of user creative freedom. Yet both prompts and paths alike nevertheless provide little insight into the creative process leading to an output. They might be described instead as "overpromising genies", performing magic at your bidding, but in a way that is near impossible to control. Current attempts at dialogic interaction appear to compound these user experience problems.

Integrating Interface Types

Naturally creative tools integrate each of these interface types and others. For example, Soundraw's (see Figure 3) homepage and other user pages mix in generous arrays of readymade tracks appended below the primary interface. Users can click on a track and directly manipulate the energy level of different segments via a simplified editor interface. Users can also generate variations of tracks or start the generation process afresh. However, there is no easy way to refine the generation process. In Bandlab, the built-in path-based song generation tool "Song Starter", makes clear in its name that this tool can only be used to seed a new track, not at an iterative stage. Once the track is created it can be edited in a multichannel DAW project. The ability to manipulate the output is not available with all music generation tools. Like most path or prompt-based generation tools, the user is also invited to audition three generated variations, a brief foray into a browser interface.

Table 3 captures this broader landscape of integrated and mixed interface design elements in a wider set of commercially available tools. The table documents the target users for the tools along with the present user interface elements from the types listed above.

Arguably, the design of smart user interface integrations that best connect users with AI systems is the frontline of innovation in ComCAI, and the challenges are substantial. A picture is coming into view of how these different interface design elements might interrelate into natural combinations. Consider Figure 4, the generative apple, as one possible formulation of these interface integrations, placing the path or prompt as the central activity in a creative workflow. The user begins by triggering the generation of an output given a selection of parameters or a free text input. Upon generation, they select from a range of options (browser), then iterate in a range of ways which may include dialogue, direct editing, providing additional conditioning input (including in-painting, out-painting type interactions), or manipulating generation parameters. The flows between elements are likely to have a more complex structure, and in the extreme the interface better resembles a map, with path-based interactions occurring as less central subprocesses (as in an image effect dialog in Photoshop).

User Discussion of Interfaces and Control

Users' own discussion of these tools on the Discord discussion forums for various products, particularly on specific feature request threads, give some insight into specific experiences of music creation with these tools. The general and feature request channels for the above projects, over a 1 year period (06/2023 - 06/2024), in the region of 1,000 posts, were studied both via a manual review and using keyword searches in order to identify discussion related to users' interest in differentiation and creative engagement. The purpose of this review is not to make qualitative statistical inferences about what a broad base of users experience or want – since forum posts are highly subject to multiple biases such as towards specific self-selecting individuals – but to capture specific user experience stories. The themes I have applied to forum comments are derived from the above analysis of design considerations and comments were sorted accordingly. These and the analysis presented here will be used to inform a more detailed forthcoming study of attitudes to generative music AI tools which does seek qualitative relevance. The five themes are: **(1) comments discussing overall generation quality**; **(2) comments discussing better *control* of the generation process**; **(3) comments dis-**

Table 3: Integration of different design elements in commercial music AI products. These tools were selected to expand the original set with a wider range of interface styles. DSE refers to "Domain Specific Editor", meaning an editor that allows limited edits of specific aspects of the generated music. For example, in Soundraw, the editor allows the user to change the "energy" of pre-generated segments. "Stems" refers to the option to download individual stems (stereo-mixes of the sub-tracks) or other component materials such as MIDI. Browsers are common, and the copyright-free music tools, targeted at non-musicians (Mubert, Soundful, Soundraw) make heavier use of Browsers and less heavy use of Maps and full-blown Editors. All companies have produced original blends of interface types. Note that these are subjective evaluations: there are many ambiguities and wide variations in how these interface types are manifest. More generally we can also note the sense of exploratory design variation, represented by the diverse range of combinations here. Lastly note that paths and prompts have become relatively interchangeable. For example, Udio and Suno have text-to-audio at their heart, but in their UI design have offered guides for what to include in prompts via their GUI design.
* At the time of writing products are under rapid development. Udio and Suno (which appeared after the start of this research), include in-painting and variation methods that allow a richer iterative approach, approaching both Map and Dialogue design elements, but since both tools generate full audio outputs, they have limited capability to offer editors or stems. However, Dialogic interaction remains elusive: forms of user iteration afford feeding different inputs to seed or define regions for generation, but still struggle with the idea of prompt- or path-based refinement.

Product	**Primary Users**	**Path**	**Prompt**	**Browser**	**Map**	**Editor**	**DSE**	**Stems**	**Dialogue**
AIVA	Commercial creators	+	+	+	+	+	-	+	-
Bandlab	Amateur/pro creators	+	+	+	+	+	-	+	-
Beatoven	Media end users	+	+	+	+	-	+	-	-
Boomy	Amateur creators	+	-	-	-	+	+	-	-
Udio	Creators and consumers	+	+	+	*	-	+	-	*
Suno	Creators and consumers	+	+	+	*	-	+	-	*
Mubert	Media end users	+	+	+	-	-	-	-	-
Soundful	Media end users	+	-	+	-	-	-	-	-
Soundraw	Media end users	+	+	+	-	-	+	-	-

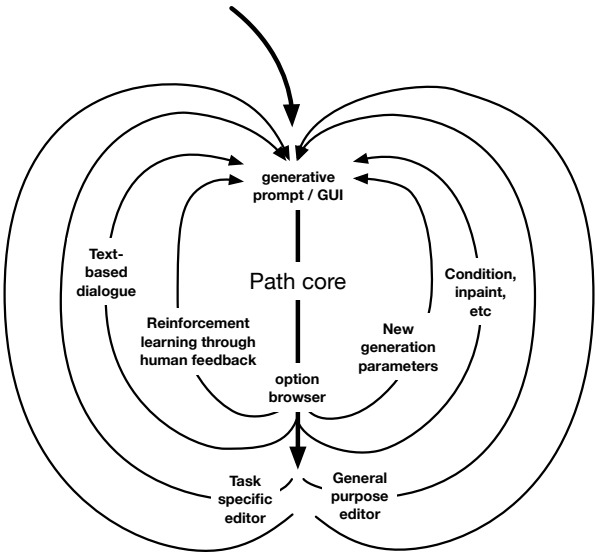

Figure 4: The generative apple. Given the centrality of path and prompt-based forms of interaction to automated generation tasks, it is possible to formulate the integration of different interaction types in terms of an iterative cycle of generation around them.

cussing design relating to the user's own *understanding* of the generation process; **(4) other ways to seed the generation process**; and **(5) the ability to edit the output of the generation process**. These themes in turn relate in varying degrees to issues of differentiation and creative engagement, discussed in the following section. These comments give some indirect insight into creators' attitudes towards the value they gain from creating with music AI tools. In most cases I have paraphrased rather than provided direct quotes to protect users' anonymity.

In relation to theme (1), a number of comments referred to the need for generative tools to generate better outputs overall, or at least statistically speaking. By contrast, other users said that they were happy with the hit-or-miss nature of generation. A user of AIVA states that it produces workable pieces about 10% of the time, but stresses that this is not a problem. On the contrary they see it as a powerful idea generator. This portrays a workflow with a low success rate but a quick turnover enabling discovery. Such attitudes are interesting with relation to differentiation and creative engagement in terms of how users' approach the immediate enjoyment of the generation process, and the time commitment that they make to nurturing works of distinction and asserting their own creative agency. The time a user such as this spends searching, even if in an opaque slot-machine style interaction, can still be considered effort invested in seeking differentiation.

Relating to themes (1), (2) and (3), one user of Boomy

discusses how they are struggling to find good outputs and outlines a number of variations they've tried to make. They point out that not knowing how the generation process works makes it hard to understand what needs to be done, and put it down to the 'luck of the draw'. This aligns with the 'slot machine' nature of Creative AI tools involving two elements: first that they lack direct control of specific elements, and second, relatedly, that in general that outcomes are not what they expect or disorienting in terms of how the outputs related to their inputs. The comment also relates to the fact that with software-as-a-service products a user can never be certain that the generation hasn't changed from one day to the next, amplifying any user experience of a lack of control (in generative AI, where companies are training and switching models frequently, this has become a significant user-experience and user-trust frontier: users of OpenAI's ChatGPT have at several points reported sudden drops in generation quality).

Contrasting with other comments we can explore a possible tension between views here. Another AIVA user comments that they hope to see better generation (90% success rate), with tracks ready to go straight to Spotify, specifying that they *shouldn't* need to go ahead and edit the results, the sounds should already be great. The tension is that clearly such users still value the generative output enough to use it, yet they feel the need to work on the output to get it up to standard.

Requests for better control of the generation process were complicated by a lack of knowledge about how the generation process worked (theme 3), and therefore what kind of control was possible. Another AIVA user's comments suggest that the need for creative control emerges with the need for differentiation within their own repertoire: they identify that they have made two successful tracks but want to get two more that *don't* sound the same, but they are stuck with how to do so.

Several other comments seek improved control, specifically seeking to be able to edit the output. Another commenter combines these and several other factors, praising the introduction of new control elements and good sounding synths, but stating that the results are still generally poor. They request more instrumental options, greater control of the mix, and options to control time signature and chord structure (e.g., 12-bar blues), with the generative system filling in the gaps.

Another AIVA user gives an example of how the limits of text-to-music prompting are not explicated through the interface, also relating to theme 3. Whilst text-to-music tools are impressive at expressing genre or instrumentation decisions, there is little evidence showing they can accurately interpret directives about note, key or chord. They state surprise that prompts directly specifying a chord progression do not result in the correct output.

The fourth theme is **(4) other ways to seed the generation process**. This is again already interspersed in the comments mentioned above and in addition several specific requests were seen in specific product discussion forums for in-painting or the ability to upload files as seeds.

Lastly, **(5) the ability to edit the results** was present in several requests, as already indicated above. An AIVA user requests that the generated output better resembles a DAW, making it easier to select between different instruments.

A user of Udio asks for "the ability to download stems which are essential for mixing, mastering, sync licensing, use in video, TV and advertisement production. The ability to download the track and stems in WAV format. WAV is essential for professional level commercial use. The creation and ability to download MIDI of each individual track in the song individually." In Udio's case, since the model directly generates the audio content in a single generation process, this would require complex post-processing[12]. An experienced user would know how to do this in a separate piece of software.

The comments discussed above also imply a wide range of attitudes and approaches to creative practice. With the exception of comments relating to item (1) – the quality of generated outputs – these comments can be generally associated with a strong focus on improving control, either during the generation process, or iterations of it thereof, or in the ability to edit the results. Even when generation quality was discussed, it seemed generally to be in this context of users wanting to go on and edit the resulting work.

What of users who simply wanted the system to do the creation work for them? This remains hard to understand, but one user's comments on the Boomy Discord server are interesting. They refer to the underlying, implied promise of Boomy, that it acts as a platform enabling people to easily release music and achieve musical success. They state that the quality of the music they've generated is on them, not Boomy, but imply that they would expect Boomy's direct feed to streaming platforms to result in more plays. This can be considered in terms of a new tension: the user clearly shows a sense of ownership of the music, and a responsibility for its quality, even if Boomy offers fairly limited control over its outputs (and asserts attribution to them on streaming platforms). But the user looks to Boomy to provide a path to success. This is understandable given Boomy's narrative marketing focus on musical success rather than music creation for its own sake. What is unclear is whether the user thinks Boomy ought to be doing more promotion, or whether they believe that Boomy-produced music is inherently more hit-worthy.

To augment these comments further, I compare them to an experienced experimental AI music creator and critic, Eryk Salvaggio, in an interview about his practice. Salvaggio emphasises the experimental nature of his work, grounded in artistic movements such as plunderphonics, considering ways in which he can stimulate Creative AI tools to generate specific forms. But he also notes the limits to creation, both in relation to the novelty of the experience, and the control afforded:

> "The productivity does become repetitive after a set amount of time, as with any tool. Because what the tool is offering you in terms of newness, in terms of getting in there and shaking things up is extremely lim-

[12] At the time of final revision of this paper, Udio had added this feature.

ited. Dall-E lets you type text in and that's all you can do... fundamentally you're dealing with a highly limited environment and you're negotiating with that environment to wrestle stuff and after time you get tired of wrestling when you can open up a DAW and do your own stuff."

Meanwhile, although (non-musician) media end users' (film-makers, advertisers) comments are not considered in this work, the design of tools for those users seemingly don't need to cater as well for such needs, except in functional ways such as enabling the mapping of music to a film's arc. The fact that these tools are centred on a browser experience corresponds with the expectation that many such users don't strongly desire such personalised involvement, although they may still value the idea that the music they are using is not repeated elsewhere.

Connecting Values, Use-Cases, Interfaces and Technologies

How do these attitudes, user experience issues and specific user interface design decisions relate to each other, and potentially relate back to the value principles of democratisation, differentiation and creative engagement, and the prognosis for the impact of creative AI tools? Given the complexity of the topic and the ambiguous nature of concepts involved, the following response is necessarily speculative and incomplete, focusing on mapping concepts rather than an operationalised set of definitions and relations. This work is necessary to seed future analysis of generative AI practices.

We have seen in a range of leading creative AI companies a dramatic shift in creative user interfaces from maps and direct manipulation to path and prompt-based designs, potentially incorporating browsers, editors, variations on maps and iterative tools such as prototypical dialogic interfaces. At one level these designs make creating music trivial, but they remove users' creative autonomy and precise control, and obfuscate the creative process. Thus the leading interface design issue for creative AI tools has been defined as the problem of the slot machine (Mass et al., 2023; Dorsen, 2022). We can also see this struggle through user feature requests that sought better generation control, better understanding of the generation process and the ability to edit outputs, while other users simply asked for better generation *quality*. It was also notable that some users took the slot machine effect in their stride, still asserting their creative agency in the process of selecting and steering.

From a simplified sociocultural perspective, all of these users are simply trying to get the respective systems to make the best or most suitable outputs as efficiently as possible. Accordingly, generative AI needs to either just create great outputs on request, or enable better user control to finish the job. But we know that individual creative goals are socioculturally rich and complex: people may wish to feel a sense of authorship, and not be labelled an imitator; they may want to nurture individual creative styles that position themselves (both through association and differentiation) in a cultural field; they may attach meaning to the processes of creation, display expertise, talent, virtuosity, and effort.

Through the co-evolution of interface designs and user behaviour there is evidence of the ComCAI market struggling with the issue of control. Udio, at the time of reviewing this paper, introduced new features allowing users to seed generations with their own audio, and allowing stem-separated downloads of the resulting output. In the image generation market, which is further ahead, companies like Adobe and Leonardo are racing to integrate user control features like in-painting and image-seeding tools. In the music AI space, AIVA is notable because it generates MIDI, which affords greater integrated edit-ability. However, this presents a dilemma for the company: should it build it's own editing tools from the ground up, a complex task involving a great cost and risk of a poor user experience? By comparison, Bandlab does not face this dilemma, since it is already the market leading online DAW. Meanwhile, others like Soundraw and Beatoven have attempted to design novel simplified interface that are fit for dedicated purposes – the enaction of simple musical edits by non-musicians – also a risky strategy in a poorly mapped design space.

In Figure 5 I speculatively plot how different interface types might be understood to offer user control and the potential for expression in different ways. I focus on "expression" here with the assumption that it is key to differentiation and, to a lesser extent creative engagement. At the heart of this mapping is the path-map distinction. Paths appear particularly constrained: there is little the user can control, and correspondingly little potential for expression. Maps, and more generally direct manipulation interfaces, stand in contrast. They allow fine control of creative outputs and consequently facilitate high expression. In this mapping, prompts are represented as presenting a dilemma. They potentially offer control, but because of the slot machine effect and relative lack of autonomy of the prompter in defining the features of the output, they offer limited expression. If combined with other interface elements and processes, such as the ability to edit, or improved future dialogic interfaces, then expression becomes increasingly possible. Meanwhile the plot provokes the question, can interfaces that don't offer control still offer expression? Figure 5 can be cross-checked against Tables 1, 2 and 3 and future work could potentially seek to measure the relationship between different interface elements and dimensions of user experience.

Bringing together these points, seemingly at least some segment of the ComCAI sector is developing according to the Red Queen expectation: the promise of instant generation is countered by a range of means of control, meaning that, although artists' processes may have sped up, and they may find they can do more things, they still have work to do and expertise to develop with respect to their positioning in the field in which they work and the creative goals they have.

A perspective that takes into account differentiation recognises that, to some extent, the need for artists to control and exert creative agency in the work they produce is driven by the desire for differentiation, not simply by objective measures of the quality of work output. In the dismissal of AI art as derivative, the need for people make work that

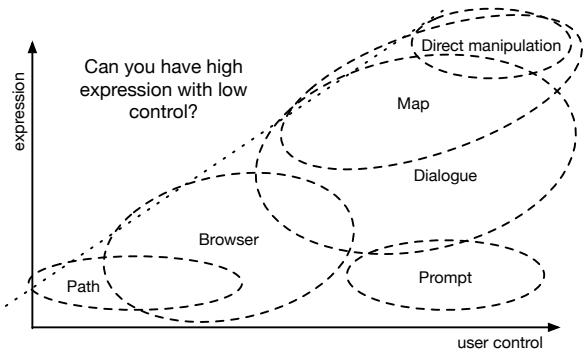

Figure 5: For creators, how can different interface types support a user's creative expression given the level of control they afford (without additional editing of the output)? Different interface types are speculatively plotted to consider what affordances they can potentially offer in these two dimensions (this not based on user data). Can users achieve expression without an interface that affords control?

doesn't sound like 'generic AI music', there is both an appeal to upholding extant aesthetic qualities but also a more complex process of redefining what those qualities should be as part of a process of differentiation. This is most pronounced in creative cultures such as the avant garde or alternative music that actively challenging mainstream practices with experimental practices. The desire to differentiate can also be expressed more indirectly in the displays of investment of time and effort in creative work, perhaps more evident in amateur and professional, rather than commercial and casual, creative spheres, where identity and the struggle for personal expression take the fore.

A perspective that takes into account creative engagement recognises that, to some extent, the stimulation that occurs during the creative process can be critical both as an end in itself or as a means to an end, that is, as part of a creative ideation process in which the creator discovers things as they go. We can see in both the selected user interface designs and user experiences that the slot machine nature of generative AI systems can be a block to creativity but can also be stimulating, presenting the user with new creative outcomes that inspire new ideas or directions, thanks to, rather than in spite of, the random nature of results (and bearing in mind the slot machine metaphor also invokes an image of generative AI practices as deeply engaging in the wrong way, addictive and mind-numbing – i.e., the AI-cynic view). In the domain of language model interaction this can go so far as tangible learning or formal ideation on the part of the user. But beyond utilitarian thinking, we have many accounts from artists of the value of the narrative behind the work and the process, the symbolic value of the elements involved (such as the provenance of sounds), the psychogeography of places where creation took place, painstaking processes of organisation, forms of meditation, derivé and so on. These and many other concepts provide dimensions of creative engagement that cannot be reduced to the production of output artefacts.

Thus creative engagement as a value could be achieved through supporting the user's learning: a chord sequence completion algorithm, say, could not only take prompts but also offer the user alternatives and explanations: this gives you a turn that was commonly used by Bowie, this chord can be used to ask a question, etc. It could also lead to designs that focus on the user experience of the process, for example allowing the user to collect artefacts. Tools that enable custom fine tuning point in this direction.

The Red Queen perspective portrays an expanded and more complicated view of creative practice in which concepts such as differentiation and creative engagement sit alongside efficiency and empowerment. But it can also accommodate an expansion of user demographics in which the concept of democratisation appears unevenly. Generative AI fits naturally into cultures of casual creators (Compton and Mateas, 2015), in which creation forms part of the everyday without necessarily involving overt attention to the creation process, expression, effort and skill, as in Instagram. A community of users of Udio and Suno, who use it to write spontaneous songs about life events, such as a job promotion, to share with friends exemplify this type of practice. Effort may be low, but there is still casual potential for differentiation and creative engagement.

These considerations help inform how creators of music AI tools can better embrace a diversity of creative values. The analysis helps us think of the problem space as follows (Figure 6). At one end is the hard problem of better generation. Although this can be couched in user-centred design issues, grounded in the psychology of arts and creativity (understanding what people consider "good"), past experience has shown that it is dominated by AI advances. At the other end is the editability of the resulting creation. Some systems output audio (Udio, Suno), others output entire DAW projects (Bandlab, AIVA). But all have the potential to be further integrated into an artists' work (generated audio can be treated as samples, or it can also be unmixed using further AI stem separation software).

In the middle lies a far more ambiguous design space of "control", tightly integrated with the affordances offered by any particular generative algorithm. In-painting and other forms of seeding have become central forms of control for generative AI tools, as has the ambiguous set of affordances associated with text-to-audio generation. This sets up a design challenge based around a trade-off between control and generative capability: how do we expand the design front to combine greater control and generative capability, or are there hard limits (Figure 6, centre plot)? But meanwhile, around this core design space is an open design space ill-defined potentials that may satisfy creative user experience in a range of contexts (Figure 6, corners).

Conclusion: Towards More Sagely Wizards

A value centred approach based on differentiation and creative engagement seeks to expand this design space further.

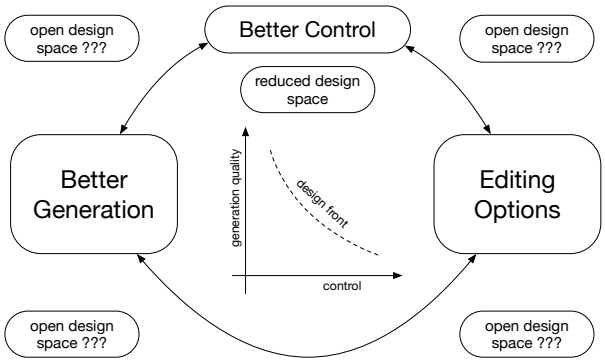

Figure 6: A reduced design space emerges due to the technology-led approach to ComCAI, caught between a triangle of better control of the generation process, better generation itself, and editing options for the resulting media. Putting aside editing options, which take us outside of the AI space, a challenging front is presented between generation and control. Other areas of the design space are underexplored.

In his foundational work on value-centred design, Cockton's (Cockton, 2005) first three design principles are pertinent:

- "Design is the intentional creation of value; other ill-focused creation succeeds through 'dumb luck'."
- "Intended value must be specified during a process of opportunity identification."
- "Value is created in the world, not in the system, and must be evaluated there, not in interaction."

(Cockton, 2005)

With this in mind, current AI tool development seems to be constrained to a limited design space, driven by technological potentials and a narrow understanding of creative values. This potentially shuts down alternative possibilities for how AI can expand its value benefits. The result appears to be that Creative AI branches in two directions: one in which AI is used in entirely functional ways, to produce music required for a purpose at negligible cost; the other in which artists use AI but in ways in which they maintain control of their artistic expression, supporting their differentiation, and creative engagement. In either case, and regardless of the potential harms of such development, AI music companies claims of supporting creative expression or meaningfully democratising music creation are tenuous.

Taking Cockton's call to specify values and to understand how they play out in the world, it is worth considering how creative engagement specifically (under which I have collapsed many different concepts) can become an actively supported value-goal for creative music AI tools. For example, in the act of path- or prompt-based music creation, how could an AI system inform the user of potentials, or other relevant information? How could the user of AIVA, who could not get their prompt with specific musical note directions to work, gain a better understanding of what prompts can achieve? How could the user of Boomy gain a better understanding of what qualities needed to be varied to achieve a differentiated result that they considered quality? How could a creator use AI not to directly generate outputs, but to support their understanding of what creation possibilities are available to them?

Perhaps in the creative space the wizard could do with a rethink. He could be more sagely, or more of a muse. With the idea that users of the most recent and advanced tools, Udio and Suno, might be casual creators engaging with AI creation for immediate listening pleasure, there is potentially a very large untapped design space of creative engagement. If those users don't come to create, in the sense we are familiar with – to act as expressive artists – but as prosumers, creating for the immediate pleasure of listening, or indeed simply requesting creation to occur, then the vision of such companies could be rethought, de-emphasising expression. These are potentially the most disruptive application areas, since they do fulfil the promise of musicians literally being replaced: while TV music was already a corporatised field, the daily listening of music audiences was not. If this points to a significant potential change, then we can consider the value-goal of creative engagement in a different use case scenario, but with potentially greater importance.

This paper hopefully makes a contribution to the public debate about the value offerings AI music companies claim. It specifies two key values that are implied in current public debates, and shows how Creative AI tools ambiguously serve those values, and how the design space is currently structured, driven explicitly by technological capabilities, then far more implicitly by such values. This highlights problems with value claims made by music AI companies. It argues that we should better articulate these values in public debates with a longer-term goal of having them established in commercial music AI design thinking.

Acknowledgments

This research was funded by the ERC Advanced Grant "Music and Artificial Intelligence: Building Critical Interdisciplinary Studies" (MusAI), European Research Council grant agreement no. 101019164, PI Prof. Georgina Born, hosted by UCL, 2021–26. I thank the MusAI research group, especially Georgina Born, Rebecca Fiebrink, Owen Green, Christopher Haworth, Aaron Einbond, Eric Drott, Darci Sprengel, Bob Sturm, Jonathan Sterne, and Artemi Maria Gioti, and with the help of Eylul Bombaci, for the opportunity to formulate these ideas in a dedicated workshop on the interfaces of music AI. I also thank the editors and two anonymous reviewers for excellent feedback on earlier manuscripts.

References

Biskjaer, M. and Halskov, K. (2014). Decisive constraints as a creative resource in interaction design. *Digital Creativity*, 25(1):27–61.

Boden, M. (1990). *The Creative Mind*. George Weidenfeld and Nicholson Ltd.

Born, G. (2010). The social and the aesthetic: For a post-bourdieuian theory of cultural production. *Cultural Sociology*, 4(2):171–208.

Born, G. (2022). Recommender systems as cultural technologies: steps to an expanded definition of culture in/and ai. In *Workshop on Cultures in AI/AI in Culture at NeurIPS 2022*.

Bourdieu, P. (1984). *Distinction: A social critique of the judgement of taste*. Harvard University Press.

Bown, O. (2023). Music AI's Potential Impact: Scoping the terms of the debate about value. In *Proceedings of the Conference on AI and Musical Creativity 2023*. https://aimc2023.pubpub.org/pub/rwi3v7tb.

Bown, O., Grace, K., Bray, L., and Ventura, D. (2020). A speculative exploration of the role of dialogue in human-computer co-creation. In *Proceedings of the 11th International Conference on Computational Creativity*.

Budiu, R. (2017). Wizards: Definition and design recommendations.

Cliff, D. and Miller, G. F. (1995). Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations. In *Advances In Artificial Life*, pages 200–218. Springer.

Cockton, G. (2005). A development framework for value-centred design. In *CHI'05 extended abstracts on Human factors in computing systems*, pages 1292–1295.

Compton, K. and Mateas, M. (2015). Casual creators. In *ICCC*, pages 228–235.

Crawford, K. (2021). *The Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.

Doctorow, C. (2023). The 'enshittification' of tiktok. *Wired*.

Dorsen, A. (2022). Ai is plundering the imagination and replacing it with a slot machine. *Bulletin of the Atomic Scientists*.

Drott, E. (2018). Why the next song matters: Streaming, recommendation, scarcity. *Twentieth-Century Music*, 15(3):325–357.

Drott, E. (2021). Copyright, compensation, and commons in the music ai industry. *Creative Industries Journal*, 14(2):190–207.

Flick, C. and Worrall, K. (2022). The ethics of creative ai. In *The Language of Creative AI: Practices, Aesthetics and Structures*, pages 73–91. Springer.

Hesmondhalgh, D. (2013). *Why music matters*. John Wiley & Sons.

Hodgson, T. (2020). Quantifying music: imagined metrics in digital startup culture. *Culture, Theory and Critique*, 61(4):424–439.

Hodgson, T. (2021). Spotify and the democratisation of music. *Popular Music*, 40(1):1–17.

Holzapfel, A. (2023). Introducing political ecology of creative-ai. In *Handbook of Critical Studies of Artificial Intelligence*, pages 691–703. Edward Elgar Publishing.

Huron, D. (2006). *Sweet Anticipation*. MIT Press.

Kaila, A.-K., Holzapfel, A., and Jääskeläinen, P. (2024). Gardening frictions in creative ai: Emerging art practices and their design implications. In *15th International Conference on Computational Creativity*.

Kaila, A.-K., Jääskeläinen, P., and Holzapfel, A. (2023). Ethically aligned stakeholder elicitation (ease): Case study in music-ai. In *Proceedings of the International Conference on New Interfaces for Musical Expression (NIME)*.

Kumparak, G. (2020). Adobe cto says ai will 'democratize' creative tools. Tech Crunch.

Laplante, A. (2012). Who influence the music tastes of adolescents?: a study on interpersonal influence in social networks. In *Proceedings of the second international ACM workshop on Music information retrieval with user-centered and multimodal strategies*, pages 37–42. ACM.

Löwgren, J. (2006). Articulating the use qualities of digital designs. In Fishwick, P., editor, *Aesthetic Computing*, pages 383–403. MIT Press.

Mass, P., Carey, F., Wheeler, C., Saatchi, E., Billington, P., and Shamash, J. Y. (2023). To infinity and beyond: Show-1 and showrunner agents in multi-agent simulations. *Self-published*.

McCarthy, K. F., Ondaatje, E. H., Zakaras, L., and Brooks, A. (2001). *Gifts of the muse: Reframing the debate about the benefits of the arts*. Rand Corporation.

North, A. C. and Hargreaves, D. J. (1999). Music and adolescent identity. *Music education research*, 1(1):75–92.

Piskopani, A. M., Chamberlain, A., and Ten Holter, C. (2023). Responsible ai and the arts: The ethical and legal implications of ai in the arts and creative industries. In *Proceedings of the First International Symposium on Trustworthy Autonomous Systems*, pages 1–5.

Plotkin, R. (2009). *The genie in the machine: how computer-automated inventing is revolutionizing law and business*. Stanford University Press.

Rezwana, J. and Maher, M. L. (2023). User perspectives on ethical challenges in human-ai co-creativity: A design fiction study. In *Proceedings of the 15th Conference on Creativity and Cognition*, pages 62–74.

Shaffi, S. (2023). 'it's the opposite of art': why illustrators are furious about ai. *The Guardian*.

Shneiderman, B., Fischer, G., Czerwinski, M., Resnick, M., Myers, B., Candy, L., Edmonds, E., Eisenberg, M., Giaccardi, E., Hewett, T., et al. (2006). Creativity support tools: Report from a US national science foundation sponsored workshop. *International Journal of Human-Computer Interaction*, 20(2):61–77.

Stokes, P. D. (2009). Using constraints to create novelty: A case study. *Psychology of Aesthetics, Creativity, and the Arts*, 3(3):174.

Stokes, P. D. and Fisher, D. (2005). Selection, constraints, and creativity case studies: Max beckmann and philip guston. *Creativity Research Journal*, 17(2-3):283–291.

Verplank, B. (2009). Interaction design sketchbook. Online resource.

Zirpoli, C. T. (2023). Generative artificial intelligence and copyright law. CRS Legal Sidebar Prepared for Members and Committees of Congress.