

The Meaning in "the Mix": Using Ethnography to Inform the Design of Intelligent Tools in the Context of Music Production

Glenn Mcgarry

Mixed Reality Laboratory, The University of Nottingham,
Nottingham, UK, glenn
mcgarry@nottingham.ac.uk

Andy Crabtree

Mixed Reality Laboratory, The University of Nottingham,
Nottingham, UK, andy
crabtree@nottingham.ac.uk

Alan Chamberlain

Mixed Reality Laboratory, The University of Nottingham,
Nottingham, UK, alan
chamberlain@nottingham.ac.uk

Christopher Greenhalgh

Mixed Reality Laboratory, The University of Nottingham,
Nottingham, UK, christopher
greenhalgh@nottingham.ac.uk

ABSTRACT

In this paper we report on two ethnographic studies of professional music producers at work in their respective studio settings, to underpin the design of intelligent tools and platforms in this domain. The studies are part of a body of work that explores this complex and technically challenging domain and explicates the ways in which the actors involved address the tension between artistic and engineering practices. This report focusses on the flow of work in the creation of a song in a digital audio workstation (DAW), which often eschews the technical requirement to document the process to maintain a creative "vibe", and the troubles that occur in subsequent stages of the production process in which complex and often messy compositions of audio data are handed over. We conclude with implications for metadata used in the process, namely the labelling and organisation of audio, to drive tools that allow more control over the creative process by automating process tracking and documenting song data provenance.

CCS CONCEPTS

• **Human computer interaction (HCI)**; • **Ethnographic studies**; • **Sound and music computing**;

KEYWORDS

Music production, Workflows, Metadata, Design Ethnography

ACM Reference Format:

Glenn Mcgarry, Alan Chamberlain, Andy Crabtree, and Christopher Greenhalgh. 2021. The Meaning in "the Mix": Using Ethnography to Inform the Design of Intelligent Tools in the Context of Music Production. In *Audio Mostly 2021 (AM '21)*, September 01–03, 2021, virtual/Trento, Italy. ACM, New York, NY, USA, ?? pages. <https://doi.org/10.1145/3478384.3478406>

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

AM '21, September 01–03, 2021, virtual/Trento, Italy

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8569-5/21/09...\$15.00

<https://doi.org/10.1145/3478384.3478406>

1 INTRODUCTION

As we design and develop intelligent tools and platforms, we need to be able to fully understand and appreciate the complexities and issues associated with the domain in which we are working. Each domain can come with its own set of individual characteristics which can inform the design and interactional aspects of the system that we would hope to develop. One such domain with an arguably individual set of characteristics is Music Production - an increasingly digital and complex enterprise that involves using multiple different software 'plugins', workflows/techniques, hardware/software; and which exists in a context with multiple 'actors' with differing expertise across a range of time frames and physical spaces.

The complexity of this setting is explicated in our previous work which demonstrates that labelling and metadata of production assets can be both physical and digital in the production space [26] and be spread across a range of media, channels and products when music is consumed [7]; and that while one might think of music as somewhat ephemeral [31], there are multiple views and framings of a given piece which may and do involve process tracking, views of the digital audio workstation (DAW), audio-based engagement, stems, tracks, mixes, temporal framings, visualisations and text.

How do we deal with the huge amount of data that one could capture within these settings? How do we start to understand and unpack what is useful and why? As Barkati writes, "without a knowledge model it would be impossible to give traces a meaning, nor to predetermine what kind of traces is worth capturing. To achieve this model, professional knowledge must be identified, listed, and characterised with domain experts, defining a digital music production Knowledge Level" [2]. We argue that studies, such as the ones we present here, can and do offer opportunities that the digital audio systems design community should take note of. This is by no means a provocative statement, a criticism, or a call to action, but rather a chance to engage with ways of doing design that begins with empirical understandings of real-world work practices. Our studies garner these understandings and engage with designers through ethnomethodologically (EM) informed ethnographies, or design ethnographies for short [9], which have a long tradition in the fields of CSCW (Computer Supported Cooperative Work) and (HCI Human Computer Interaction). We hope that in presenting our work we will be able to both support and strengthen digital audio systems design.

2 RELATED WORK

In recent years, developments in knowledge representation of the music production domain [13, 16, 17, 21, 42, 43] has underpinned demonstrations of intelligent data driven tools to support music-making [38], including modelling and automating the mixing process [19, 22, 23, 34, 41]; automating audio editing [14, 15, 18, 20]; and for tracking production activities and audio data provenance [1–3]. However, in order to comprehensively support music production activities it is essential that intelligent production tools are developed with an empirical understanding of real-world music production practices [34] and to what extent these may be usefully automated [28]. A requirement exists, then, for a foundational understanding of music production work, which we argue is best brought about through the ethnographic tradition of study that has been established in the field of HCI.

Although the fields of HCI and CSCW have been significant in providing detailed understandings of the intersections between humans and computers in different work domains [30], music production remains a particularly challenging domain to study. *Firstly*, the problem space is broad and complex with methods that include ‘traditional’ processes and legacy technologies, novel methods and modern digital tools, and any combination between [6, 40]; and *secondly*, professional places of music production, such as the recording studio, are notoriously difficult to access as study site [39]. Given this barrier, surveys and interviews are commonly deployed in order to garner insights into professional music production practices [32, 33, 35, 36], however, there are limitations to what can be explicated about the work from the interviewee’s account, unless it is conducted ethnographically from “within the flow of work” [9]. Nonetheless, a few ethnographic studies of music production exist that highlight issues in the hand-over of audio data from the recording studio to the mix engineer [26]; the use of socially networked technologies to support distributed and collaborative production activities [27]; and the ad hoc contingencies required to leverage functionality of incompatible production tools [5].

Our studies - which are part of a set of studies that are reported more fully elsewhere [24, 26] - contribute to the body of ethnographic work in this domain and begin to map out the problem space within the scope of *professional* music production in different contexts and settings.

3 THE STUDIES

In this section we present a summary of two of our studies that extend the scope and context of music production practice, beyond the setting of the traditional recording studio and its processes (again, these are reported elsewhere [24, 26]), to ‘in-the-box’ music production; methods in which songs are created from end-to-end using the virtual studio environment of the digital audio workstation (DAW).

The participants were recruited through the researcher’s professional network and the study settings selected opportunistically based on some basic criteria intended to define the scope of the studies:

Firstly, the participants’ primary professional occupation involved music production. This addresses a gap in existing studies that centre on amateur practice [5, 27], however as we shall see, in reality this facet of the modern music production landscape is often mixed.



Figure 1: The producer’s home studio

Secondly, that the study scenario involved a real-world musical project. Per design ethnography’s requirements, this criterion ensures that the interactions, troubles, and contingent actions are captured that are otherwise missed by, for example, representations of work rendered in process models and technical documents, or the documentary accounts of interviews.

Thirdly, that the settings and scenarios were sufficiently different, so that a cross-sectional view across the breadth and depth of the problem space can begin to be mapped out. Recruitment resulted in the completion of studies with professional producers in three locations: a ‘traditional’ recording studio; a home studio; and a shared collaborative project studio. Here we report on the latter two studies

The researcher was immersed in the two studio settings to observe the work practices of the two professional music producers and to understand and report their working methods as the participants themselves do, including the accountability of the organisational character of artefacts and interactions in the setting.

The research data gathered includes 20 hours of video recordings of production activities captured in real-time, field notes, and photographs of the settings. From the analysis of these research data, the studies make reportable a range of music making activities using DAW software, which in the following sections are illustrated through workflows, descriptions and some extracts of the video transcription that demonstrate how audio is reused and assembled in different stages of the production and how different orientations to order in these stages propagate through and cause tension in later stages of the process.

4 THE PROFESSIONAL HOME STUDIO

The first of the studies observes a professional electronic dance music producer, ‘Steve’, working alone to create ‘tunes’ using only a DAW in his home studio (Figure 1).

Steve has over two decades of experience as professional DJ and music producer in the Drum and Bass genre and is a partner in an independent record label through which his music is marketed and released for publishing.

Figure 2 represents Steve’s process which involves a range of activities that centre on assembling, layering, and shaping sound into a song.

For the purposes of this paper, we focus on two of the core activities: Preparation and Setup; and Creating the ‘Tune’.

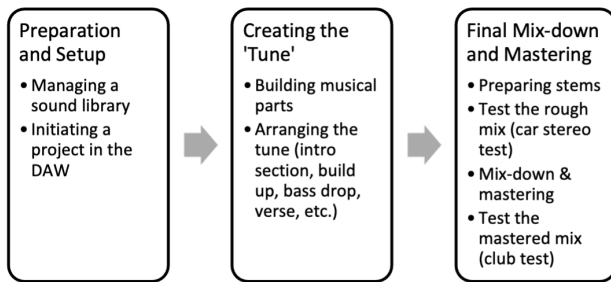


Figure 2: An overview of the producer's workflow

4.1 Preparation and Setup

4.1.1 Managing a Sound Library. Audio clips are used extensively in creating a tune and so part of Steve's studio preparations involve sourcing these from multi-media, for example film dialogue; commercial sample CD's; or experimentation with sound recorded acoustically or generated using the DAW's synthesizer plug-ins, known as Virtual Instruments (VI's). Steve organises audio on the DAW PC into named folders (FX, drums, *etc.*) and uses automated analysis software to append tempo and key signature information to audio file names, which provides for the retrieval and auditioning of audio, using the PC's native search tools, when creating the tune; a provision for which we will see later in the process.

4.1.2 Initiating a Project in the DAW. Due to the flexibility that is designed-in to most DAW platforms to support a range of production activities, part of the producer's process is to configure the DAW for the activity at hand. In the following extract, Steve talks through this activity with the researcher (Glenn) as he sets up the DAW ready to begin creating a new tune. This includes configuring signal routing to include 'busses' which group sounds to the same channel (or sub-mix) for processing and volume control; and side chaining, in which signals can be routed for cross channel control of audio processing (Extract 1).

Extract 1 - Setting up the DAW to "get the vibe going"

Steve: So generally, what I always do is kind of, set up all my busses and stuff first. . . I put a couple of busses in here [creates a bus and labels it 'Ds'(drums)] a erm. . . call it like a track spacer bus [creates a bus and labels it 'Ts'(trackspacer)] [creates a bus and labels it 'bass'] Bass bus. . .

Glenn: So, busses are like, sub-groups or sub-mixes?

Steve: Yeah so what I'll do basically [points to first track in DAW project window] so I'll have like, maybe kind of pads or intro noises and then. . . .send them. . . .to the trackspacer bus, which (is) basically (this): [drags plug-in preset from menu to 'Ts' bus; opens interface] Waves Factory Trackspacer (plug-in). . . .that keeps everything under the drums, and if I start like that then it. . . to be honest the end mix it makes it much easier because I'm never fighting for anything, you know what I mean. There's nothing worse than killing er. . . the vibe of everything by constantly trying to tweak everything as you go

along. We're trying. . . what you kind of do is get the vibe going first and then really worry about that (the mix) at the end . . . I work much faster and . . . here you go [clicks on a mixer context menu to select routing option] I just sidechain the drum bus with the trackspacer so now any of the drums I put through will make all the rest of the music (levels) duck (in response to the drum levels) . . .

Glenn: Gotcha.

Steve: Yeah? Simple and that's pretty much a default thing that I always do.

The producer's methods for setting up the DAW are oriented toward scaffolding the following tune creation activity with a configuration that prospectively provides order and supports the flow of work. Creating the tune involves some experimentation with sound, as we describe in the next section, and Steve's configuration provides a way of semi-automatically balancing the mix of sounds as they are created (and sometimes discarded) so that the flow of creative work is less frequently interrupted by the need to balance, or "tweak" the mix.

4.2 Creating the Tune

This activity is made up of several constituent actions that essentially involve the creation and layering of 'musical parts' (drums, bass, synth *etc.*) and sections of the tune. (Figure 3).

4.2.1 Building the Tune in Layers. To create each part, Steve goes through a process of searching, browsing, and auditioning sounds from the sound library or the array of Virtual Instrument (VI) plug-ins in the DAW to find sounds that might fit with the central musical idea of the tune. Selected sounds are then imported to the DAW and musical parts are then created using an iterative process of shaping and treating sounds through experimentation with layering, editing, and digital signal processing (DSP); arranging the part musically in terms of rhythm and/or melody, usually in short playback loops of 4, 8, or 16 bars; and employing referencing and analysis techniques that compare the evolving tune with existing music recordings, in order to balance the mix of parts and/or reshape musical parts if necessary. Once enough of the core musical parts are in place - for example, the drums, the bass, and other instrumentation that establishes a musical theme or motif - then, Steve begins to build the tune in sections.

4.2.2 Building the Tune in Sections. This essentially involves copying and pasting sound clips in bulk and arranging them along the timeline of the DAW to begin building the song up in sections (Figure 4).

This approach switches from creating musical parts in a playback loop to arranging sounds as musical events in relation to the full runtime of the tune. For example, the musical parts might cue in one at a time to build the introduction section; create climactic points in the song (for example the bass 'drop'). This process also occasionally requires the addition of other musical parts to emphasise climactic points in the song and which uses the same methods as in building the initial layers of the tune, as illustrated in Extract 2

Extract 2 Auditioning and importing sounds

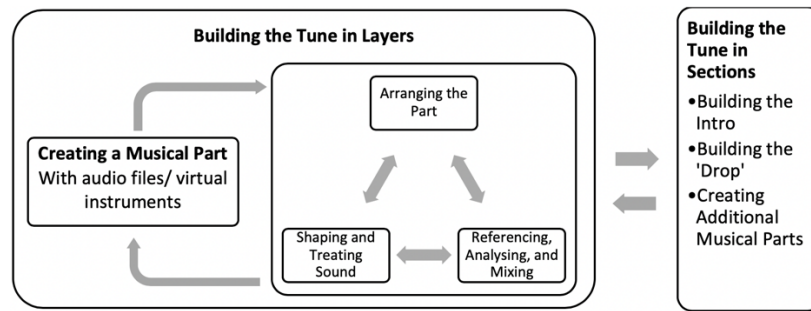


Figure 3: The activities involved in creating the tune



Figure 4: Duplicating the arrangement of musical parts and clips along the DAW timeline



Figure 5: Auditioning sound files

Steve has mostly completed layering parts and arranging the tune's intro and is adding some sound effects to emphasise a climactic build up to the end of the intro. To do this he browses and retrieves sounds from a file folder named 'FX I ALWAYS USE'.

Steve: [listens to the introduction; stops playback near the end] right, let's work on the build-up next [minimises the DAW window, closes an open explorer window] (sighs) told you it's boring [eyes scan desktop for a folder; clicks on 'FX I ALWAYS USE' folder to open an explorer window to view contained files]

Steve: [begins auditioning sounds: clicking on a file in the folder to open and audition the sound in media player (Figure 5)]

Steve: (inaudible) Want one of them in there (that will go) in the effects section, quickly. [drags file to a new DAW track; creates a new bus and routes all new tracks to it, names bus 'fx'; starts playback; duplicates the sound effect clip and drags to position later in the timeline; stops playback]

An observation of interest about this process, is the orientation to the organisation of sounds, which in turn is continuously evolving. On importing the additional sounds to the DAW, for example, as Steve imports audio files the DAW automatically creates new tracks in the arrangement before he manually routes the audio outputs of those tracks to a buss (labelled 'fx') as part of the mix. The DAW also automatically labels the new tracks with the name of

the imported file which the producer can, of course, however his is deferred to a later stage of the production process (preparing 'stems' for the final mix down) and is set aside at the immediate stage to not interrupt the flow of creative ideas in the process. In this way labelling propagates through the process, but not always predictably, as during experimentation sounds may be discarded or replaced, leaving their imprint in the form of a track label. For Steve this is perhaps not problematic as the artistic meaning in the mix of audio artefacts is accounted for and available to him in other ways – the arrangement and order of the tracks, the groupings of audio facilitated through the configuration of busses. In the longer term, however, this can be problematic as automatically generated labels might not fully communicate artistic intent to other parties who might be involved in the process nor be useful to coordinating later stages of the production processes, as we observed in the second study.

5 THE PROJECT STUDIO

This study observes a professional producer, Mark, doing a 'client mix' of a demo song in a shared project studio (Figure 6), as part of an ongoing musical production with a collaborating artist.

Mark is also a professional with over two decades of experience producing music in the Hip-Hop genre and is also a university lecturer in music technology. Little information could be disclosed about the collaborating artist, other than he created his own song demos using DAW software and is presumably a novice music producer. Mark's involvement in the project is to facilitate what he



Figure 6: The Project Studio Setting

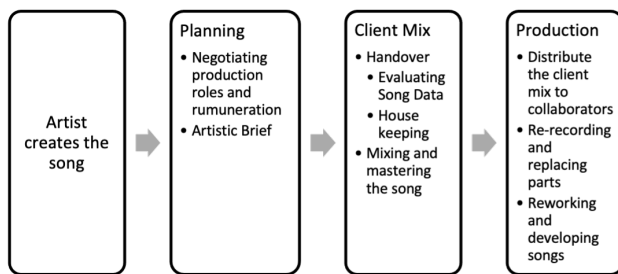


Figure 7: The process for distributed collaboration in music production

described as the ‘iterative’ development of the demos to a professional level of production quality, which at the time of the study prospectively included activities such as re-recording musical parts, altering to the musical arrangement, and doing professional mixes of the songs (Figure 7).

For the purposes of this paper, we focus on a transitional phase of this process, the ‘client mix’, and drill down further to the handover activity.

5.1 Client Mix (Handover)

The practical reasoning for doing a client mix of each song is firstly, to allow Mark to become familiar with the artist’s songs, to consider how they might be developed and what production activities these might involve; and secondly, to produce and distribute an enhanced mix of the demos, to ‘create a vibe’ amongst potential collaborating performers.

5.1.1 Evaluating Song Data. For each song that is handed over, Mark undertakes a technical evaluation of the song data that the artist has created, which in this case are compositions of audio clips assembled using DAW software (similar to the results of Steve’s creative methods in the home studio study). This enables Mark to assess the potential workload implications surrounding, for example, compatibility between DAW platforms and, crucially, the make-up of the song’s composite musical parts. In this scenario, the song data that the artist has handed over was created in a compatible DAW, done so without the kinds of labelling and organisation




Figure 8: The handed over song with the DAW’s automatic labelling intact.

of sounds that is useful in communicating information about the musical composition through to the mixing process (Extract 3).

Extract 3 Evaluating song data and planning the mix

Mark: I think what the next step is, and I just need to be clear about this - . . . because there’s a potentially big workload implication in how I go about this next step. . . - I’ve got to label all the crap. I’m going to move things around, so I know where they’re at.

Glenn: Let me take a snapshot of that (DAW screen), the fact that it says, ‘audio 15, 17, 25’, randomly, that’s of interest (Figure 8).

Mark: Yeah, and it’s not good housekeeping, I’d murder my students for it. But whatever, I’m not going to criticise.

Glenn: No, but you’ve still got to contend with it?

Mark: Yeah, I don’t know what’s what, what’s vocals. . . so I’m going to move them and colour code them into rhythm section, vocals, instruments, and shit. . . (or) stuff, you know what I mean? . . . so, when I get going (with the mixing) I want to move fast, I don’t want to be going where’s that?... I think that’s the process now.

5.1.2 Housekeeping. After the evaluation of the song data, Mark goes through a process of housekeeping, to bring about some organisation prior to doing the client mix, and to unpick and make sense of the composition. This involves a process of isolating playback, or soloing, of the tens of tracks in turn, in order to discern their content and place in the musical composition (e.g. lead vocal, vocal harmonies ‘hi’ and ‘low’, bass line etc.) and them applying appropriate labelling, reordering and grouping tracks, and colour coding tracks to reorganise them (Extract 4).

Extract 4 Color coding as part of housekeeping

Glenn: Do you have any particular preference for colour coding?

Mark: I tend to go for blues for vocals or bluey type of colours and brown, earthy things for rhythm sections. . . .

Mark: . . . So, I’m going to distinguish female vocals and backing vocals from (vocal) effects really [selects all clips on tracks 18-28 (female vocals); applies indigo colour option]

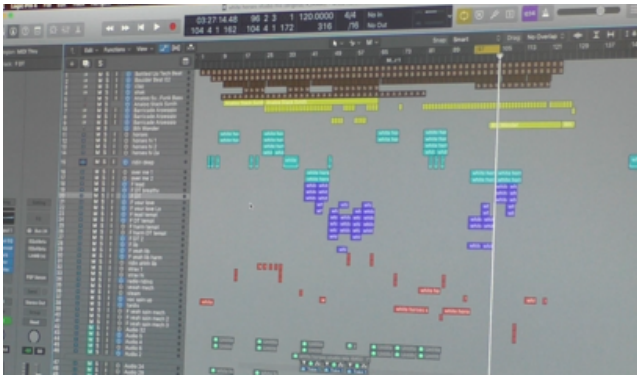


Figure 9: "We've got some organisation". The reordered, labelled, and color-coded composition tracks.

Mark: [selects all clips on tracks 11-17 (male vocals); applies light blue colour option]

Mark: And then. . . .these are all (vocal effects) [selects all clips on tracks 29-40 (vocal effects); applies red colour option]

Mark: Right so. . . we've got some organisation (Figure 9).

This task took approximately 2 hours of Mark's time in the studio to unpick and reorganise before being able to embark on the client mix task, which, it should be noted, was only the first of a set of several similarly produced songs that artist had handed over.

The key observation from this example is that processual information can be of relative importance to the situation and orientation of work. For the collaborating artist and for Steve in the home studio study, the task of labelling and organising musical elements is secondary to the process of using the DAW to be musically creative. In contrast, for Mark these kinds of information are essential for communicating the artistic intent of a song through to a new and more technical phase of the production process, also facilitated in the DAW. Mark demonstrates that this information can be re-constituted and the meaning in the mix can be accounted for and interpreted from the remaining song data, however this is costly overhead of work. In the next section we discuss the implications of these findings further

6 DISCUSSION

The examples that we present in this paper are in no way intended to address the entire problem space, which is now vast in its scope due to the ubiquity of the DAW and the novel production possibilities that digital audio technology has brought about. Further still, the full scope of music production methods that were observed and explicated in these studies are not included in the space available here, but are available elsewhere [24]. Even so, further mapping out of this domain is most certainly required and in doing so, the generalisable features of music production will inevitably surface - and with few examples required where studies attend to the recognisably natural order of a setting [12].

The examples from the studies that we present here are intended to sensitise designers to the different orientations of music production work in which the DAW is used and how we might characterise the organisation of each stage of the process. What we hope we have made explicit here is that the issues we have highlighted are not necessarily a result of bad production practices, or a need for education, but that organisation of different kinds are at play and at different points in the process (the locally organised sound libraries, groupings, and visual representations of audio and their arrangements, labelling, and colour coding). To the trained eye the unlabelled, unordered layers of sound making up a composition might represent a chaotic and amateurish effort. However, there is a situational "order at all points" [37] that is missed in the formal representations of professional production practice - for example process models and professional accounts. Further to this, it is these formal representations that underpin much of the DAW design in its emulation of traditional recording studio equipment and workflows [4].

The troubles that our studies have explicated then, highlight that the human-readable metadata - in the form of labelling of the musical elements of the composition displayed in the DAW - are essential to the coordination of production process, as the meanings and artistic intent that they convey to others involved in the process are embodied within the organisation of production resources. Furthermore, this embodied organisation holds together and enables novel production methods, such as in-the-box music production, to exist outside of the traditional production practices that are centred on the recording studio, which are themselves bound by a standard of process tracking and documentation [29] that would be otherwise unworkable in a small-scale production. Even so, for such novel production practices it should not be assumed that organisational information will be generated by the human in the loop, nor that all processes are of a piece in terms of the practical requirements for processual information.

This has significant implications for the prospective design of intelligent tools that might seek to address the kinds of troubles seen in the handover of song data, which is that there is a need for DAW-based tools that can capture the creative process and communicate the underlying artistic intent through a combination of application design that is sympathetic to creative workflows and uses a minimally obtrusive level of automation. We observed, for example, that in the tune creation process that file names, meaningless or otherwise, propagate through the process via an automated features of the DAW when audio data is imported into the process. However, this propagation is not exploited in different directions to the flow of data so that, for example, newly created tracks (which may remain in the song or be deleted in experimentation) would inherit the names applied to the buss that it is assigned to, or from its classification in a sound library.

Through ethnographic methods of inquiry, then, we have not only demonstrated a need for the prospective design of intelligent tools for process tracking, but also the levels of automation required and at what points in the process, the kinds of information that are useful to capture, and at which stages this information should be captured and subsequently reused. We hope that our studies will promote future work in this domain, including "in the wild" studies

[8, 10, 11] that provoke new ways of thinking about and designing intelligent tools for music production.

7 CONCLUSION

In this paper, we have demonstrated that ethnographic studies have been able to provide more detailed understandings of the value of processual information in music production and how metadata as an organising resource propagates through the process. Through these insights, we have offered implications for the prospective design of intelligent tools to support the production process through potentially high levels of automated tracking of targeted aspects of the creative process. This also has broader implications for the design of intelligent tools for automating tasks in other domains that are becoming increasingly digital through, for example interacting with IoT-based technologies. Understanding the practices of people in their day-to-day settings, "in the wild" [8, 11] across a range of contexts is key to the development and design of applications using AI in the creative industries and other domains. These understandings give us a greater insight into the role that HCI and Human-Centred approaches can play in the development of intelligent systems in the area of music production, creation and performance [25, 26]. We hope, then, that our work will motivate and provoke conversations, discussion and future work in the music production/creation domain and will be of interest to researchers in other disciplines, research communities and industry.

ACKNOWLEDGMENTS

This research was funded by the UK EPSRC project Fusing Semantic and Audio Technologies for Intelligent Music Production and Consumption (FAST) (EP/L019981/1). Dr Alan Chamberlain's part in this work was supported by the Engineering and Physical Sciences Research Council [grant number EP/T51729X/1] projects RCUK Catapult Researchers in Residence award Digital - Disruptive Beats - Music - AI - Creativity - Composition and Performance, [grant number EP/V00784X/1] UKRI Trustworthy Autonomous Systems Hub and [grant number EP/S035362/1] PETRAS 2.

REFERENCES

- [1] Karim Barkati, Alain Bonardi, Antoine Vincent, and Francis Rousseaux. 2012. GAMELAN: A Knowledge Management Approach for Digital Audio Production Workflows. *Proceedings of Artificial Intelligence for Knowledge Management Workshop of ECAI 2012*: 3–8.
- [2] Karim Barkati and Francis Rousseaux. 2013. Considering music production and culture management as an emerging digital ecosystem. In *Proceedings of the Fifth International Conference on Management of Emergent Digital EcoSystems - MEDES '13*, 45–53. <https://doi.org/10.1145/2536146.2536151>
- [3] Karim Barkati and Francis Rousseaux. 2014. How to Understand Digital Studio Outputs: The Case of Digital Music Production. In *Artificial Intelligence for Knowledge Management*. 151–169. https://doi.org/10.1007/978-3-642-54897-0_9
- [4] Adam Patrick Bell. 2008. Can we afford these affordances? GarageBand and the double edged sword of Digital Audio Workstations. *Action, Criticism & Theory for Music Education* 7, 1: 105–138.
- [5] Phillip Brooker and Wes Sharrock. 2013. Remixing Music Together: The Use and Abuse of Virtual Studio Software as a Hobby. In *Ethnomethodology at Play. Directions in Ethnomethodology and Conversation Analysis*, Peter Tolmie and Mark Rouncefield (eds.). Ashgate, 135–155.
- [6] D Byrne. 2013. *How Music Works*. Canongate.
- [7] Alan Chamberlain and Andy Crabtree. 2016. Searching for music: understanding the discovery, acquisition, processing and organisation of music in a domestic setting for design. *Personal and Ubiquitous Computing* 20, 4: 559–571. <https://doi.org/10.1007/s00779-016-0911-2>
- [8] A Crabtree and A Chamberlain. 2013. Introduction to the special issue of "The Turn to The Wild". *ACM Transactions on Computer-Human Interaction* 20, 3. <https://doi.org/10.1145/2491500.2491501>
- [9] Andrew Crabtree, Mark Rouncefield, and Peter Tolmie. 2012. *Doing Design Ethnography*. Springer London, London. <https://doi.org/10.1007/978-1-4471-2726-0>
- [10] Andy Crabtree, Steve Benford, Chris Greenhalgh, Paul Tennent, Matthew Chalmers, and Barry Brown. 2006. Supporting Ethnographic Studies of Ubiquitous Computing in the Wild. *Proceedings of the 6th Conference on Designing Interactive Systems (DIS'06)*: 60–69. <https://doi.org/10.1145/1142405.1142417>
- [11] Andy Crabtree, Peter Tolmie, and Alan Chamberlain. 2020. "Research in the Wild": Approaches to Understanding the Unremarkable as a Resource for Design. In *Into the Wild: Beyond the Design Research Lab*, Alan Chamberlain and Andy Crabtree (eds.). Springer International Publishing, Cham, 31–53. https://doi.org/10.1007/978-3-030-18020-1_3
- [12] Andy Crabtree, Peter Tolmie, and Mark Rouncefield. 2013. "How Many Bloody Examples Do You Want?" Fieldwork and Generalisation. *ECSCW 2013: Proceedings of the . . .*, September: 21–25. https://doi.org/10.1007/978-1-4471-5346-7_1
- [13] György Fazekas, Yves Raimond, Kurt Jacobson, and Mark Sandler. 2010. An Overview of Semantic Web Activities in the OMRAS2 Project. *Journal of New Music Research* 39, 4: 295–311. <https://doi.org/10.1080/09298215.2010.536555>
- [14] György Fazekas and Mark Sandler. 2007. Structural decomposition of recorded vocal performances and it's application to intelligent audio editing. 1–20.
- [15] György Fazekas and Mark B. Sandler. 2007. Intelligent Editing of Studio Recordings with the help of Automatic Music Structure Extraction. *Proceedings of the 122nd AES Convention*.
- [16] György Fazekas and Mark B. Sandler. 2013. Describing audio production workflows on the semantic web. In *International Workshop on Image Analysis for Multimedia Interactive Services*, 1–4.
- [17] György Fazekas and Mark B Sandler. 2011. The Studio Ontology Framework. *Information Retrieval*, Ismir: 471–476.
- [18] Steven Hargreaves, Anssi Klapuri, and Mark Sandler. 2012. Structural Segmentation of Multitrack Audio. *IEEE Transactions on Audio, Speech, and Language Processing* 20, 10: 2637–2647. <https://doi.org/10.1109/TASL.2012.2209419>
- [19] Jordan V Hawkins. 2013. Automating Music Production with Music Information Retrieval. Ohio State University.
- [20] Tetsuro Kitahara, Masataka Goto, Kazunori Komatani, Tetsuya Ogata, and Hiroshi G. Okuno. 2005. Instrument Identification in Polyphonic Music: Feature Weighting with Mixed Sounds, Pitch-Dependent Timbre Modeling, and Use of Musical Context. *Proceedings of the 6th International Society for Music Information Retrieval Conference (ISMIR)*: 558–563.
- [21] Sefki Kolozali. 2014. A framework for automatic ontology generation based on semantic audio analysis. 1–10.
- [22] Brecht De Man, Brett Leonard, Richard King, and Joshua D Reiss. 2014. An Analysis and Evaluation of Audio Features for Multitrack Music Mixtures. *15th International Society for Music Information Retrieval Conference ISMIR 2014*, Ismir: 137–142.
- [23] Brecht De Man and Joshua D Reiss. 2013. A Knowledge-Engineered Autonomous Mixing System. In *Audio Engineering Society Convention 135*.
- [24] Glenn McGarry. 2020. Studies of the Social Character of Metadata in Modern Music Production. University of Nottingham.
- [25] Glenn McGarry, Alan Chamberlain, Andy Crabtree, and Christopher Greenhalgh. 2021. Placing AI in the Creative Industries: The Case for Intelligent Music Production. . 562–572. https://doi.org/10.1007/978-3-030-78635-9_72
- [26] Glenn McGarry, Peter Tolmie, Steve Benford, Chris Greenhalgh, and Alan Chamberlain. 2017. 'They're all going out to something weird': Workflow, Legacy and Metadata in the Music Production Process. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing - CSCW '17*, 995–1008. <https://doi.org/10.1145/2998181.2998325>
- [27] Sean McGrath, Alan Chamberlain, and Steve Benford. 2016. Making Music Together. In *Proceedings of the Audio Mostly 2016*, 186–193. <https://doi.org/10.1145/2986416.2986432>
- [28] David Moffat and Mark B. Sandler. 2019. Approaches in Intelligent Music Production. *Arts* 8, 4: 125. <https://doi.org/10.3390/arts8040125>
- [29] Bobby Owsinski. 2005. *The Recording Engineers Handbook*. ArtistPro.
- [30] Philippe Palanque. 2018. Engineering Automations: From a Human Factor Perspective to Design, Implementation and Validation Challenges. In *Proceedings of the ACM SIGCHI Symposium on Engineering Interactive Computing Systems*, 1–2. <https://doi.org/10.1145/3220134.3223044>
- [31] Konstantinos Papangelis, Alan Chamberlain, and Hai Ning Liang. 2016. New directions for preserving intangible cultural heritage through the use of mobile technologies. *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct, MobileHCI 2016*: 964–967. <https://doi.org/10.1145/2957265.2962643>
- [32] Pedro Pestana. 2013. Automatic Mixing Systems Using Adaptive Digital Audio Effects. UNIVERSIDADE CATOLICA PORTUGUESA.
- [33] Pedro Pestana and Joshua D Reiss. 2014. Intelligent Audio Production Strategies Informed by Best Practices. *AES 53rd International Conference on Semantic Audio*: 1–9.
- [34] Joshua D. Reiss. 2011. Intelligent systems for mixing multichannel audio. In *2011 17th International Conference on Digital Signal Processing (DSP)*, 1–6. <https://doi.org/10.1145/2491500.2491501>

- [//doi.org/10.1109/ICDSP.2011.6004988](https://doi.org/10.1109/ICDSP.2011.6004988)
- [35] David Ronan, Hatice Gunes, and Joshua D. Reiss. 2017. Analysis of the Subgrouping practices of professional mix engineers. In *142nd Audio Engineering Society International Convention 2017, AES 2017*.
- [36] David M. Ronan, Brecht De Man, Hatice Gunes, and Joshua D. Reiss. 2015. The impact of subgrouping practices on the perception of multitrack mixes. *139th Audio Engineering Society International Convention, AES 2015*: 1–9.
- [37] Harvey Sacks. 1984. 2. Notes on methodology.
- [38] Mark Sandler, David De Roure, Steven Benford, and Kevin Page. 2019. Semantic Web Technology for New Experiences Throughout the Music Production-Consumption Chain. In *2019 International Workshop on Multilayer Music Representation and Processing (MMRP)*, 49–55. <https://doi.org/10.1109/MMRP.2019.00017>
- [39] Paul Thompson and Brett Lashua. 2014. Getting It on Record: Issues and Strategies for Ethnographic Practice in Recording Studios. *Journal of Contemporary Ethnography* 43, 6: 746–769. <https://doi.org/10.1177/0891241614530158>
- [40] A Watson. 2014. *Cultural production in and beyond the recording studio*. Routledge.
- [41] Thomas Wilmering, György Fazekas, and Mark B. Sandler. 2012. High-Level Semantic Metadata for the Control of Multitrack Adaptive Digital Audio Effects. In *Proceedings of the 133rd AES Convention*.
- [42] Thomas Wilmering, György Fazekas, and Mark B. Sandler. 2013. The Audio Effects Ontology. *Proceedings of the 14th International Society for Music Information Retrieval Conference (ISMIR-2013)*.
- [43] Thomas Wilmering and Mark B Sandler. 2013. Semantic Metadata for Music Production Projects. First International Workshop on Semantic Music and Media SMAM2013.