

Design Process in Visual Programming: Methods for Visual and Temporal Analysis

Jack Armitage, Thor Magnusson

Intelligent Instruments Lab
Iceland University of the Arts
Reykjavik, Iceland
jack@lhi.is,
thor.magnusson@lhi.is

Andrew McPherson

Augmented Instruments Lab
Dyson School of Engineering
Imperial College London
London, UK
andrew.mcpherson@imperial.ac.uk

ABSTRACT

Visual programming languages, such as Pure Data (Pd) and Max/MSP, have been prevalent in computer music for nearly three decades. However, few shared and consistent research methods have emerged for studying the reproducible use of digital musical instrument (DMI) designers employing these languages. In this paper, we introduce straightforward methods for extracting design process data from Pd usage through automated version control and protocol-based annotation. This data enables visual and temporal analysis, which can reveal patterns of DMI design cognition and collaboration processes. Although our focus is on design, we believe that this approach could also benefit creativity studies and musicological analysis of the compositional process. We present the outcomes of a study involving four groups of DMI designers in a one-hour closed activity and demonstrate how these analysis methods can be used to gain additional insight by comparing them against participant survey data. In discussing how these methods could be enhanced and further developed, we address validity, scalability, replicability, and generalisability. Lastly, we examine motivations and challenges for DMI design cognition research.

1. INTRODUCTION

A significant proportion of DMI design researchers are music technologists who examine musical behaviour, rather than design technologists who investigate design behaviour, and the methods commonly employed in the field reflect this distinction. There are no frameworks that concentrate on or adequately acknowledge DMI design, nor are there any domain-specific ontologies or protocols for analysing DMI design processes. For instance, while numerous studies involve Pure Data and Max/MSP [1–3], and a considerable number of DMI design researchers have been utilising and teaching these platforms for decades, there are no common qualitative or quantitative methods in the field for analysing their usage. Consequently, each study involving the interpretation of these tools’ use must develop


its own research methods and tools to support them, yet other researchers do not benefit from any such tools. Furthermore, although there are numerous examples of technology probes in DMI design research, these are primarily applied to musical scenarios rather than design ones (for example, see [4, 5]). The reporting in such research papers focuses more on the probe design and construction process, and less on the activity or context design, which is crucial for investigating design [6].

In this paper, we present an extended analysis of the collaborative design process involved in creating a Pure Data (Pd) patch, as previously reported in [7]. The aim of this analysis was to reveal the typically concealed design processes occurring during Digital Musical Instrument (DMI) design. To achieve this, we developed visual and temporal methods, which, when combined with post-activity survey data, provided us with a more comprehensive insight into the events that transpired. We believe other researchers will find our methods easy to understand and valuable to apply in their own work.


2. ANALYSING DMI DESIGN PROCESSES

When it comes to design processes, there is significantly more DMI design literature prescribing them than analysing them [8]. In cases where design process analysis is conducted, it typically focuses on evaluating a specific design method [9] or tool [10], rather than examining the process involving embodied design, music cognition, craft practice, and decision-making under constraints. Although some examples of the latter are beginning to emerge [11], the methods required for advanced research in this area are not yet well-established within the DMI design research community. Furthermore, though the material at hand in this study is software, it is problematic to assume that software engineering definitions, theories and frameworks are applicable to the arts, where there are no problems to solve, only decisions to make [12].

An example of such research can be found outside of DMI design in the work of Delle Monache and Rocchesso [13]. Their work concentrates on investigating sound design cognition and employs methods from design studies, such as linkographic analysis [14] and ontology-based protocol analysis [15]. Probes and briefs are utilised to constrain design sessions and make them discretisable, the lat-



(a) Top view of the *Unfinished Instrument* with four mic capsules connected to the Karplus-Strong algorithms in Pure Data.



(b) Instrument and software tools in situ, with a second computer monitor showing Bela print debug and oscilloscope output.

Figure 1: Instrument (left) and workshop environment (right).

ter being a prerequisite for analysis. The aim of the analysis is to isolate *design moves* — discretised steps of design decision-making — and cross-reference them to form a linkograph, revealing the design state space and its non-linear progressions [14]. Additional ontologies can then be applied to address concerns such as design cognition and the impact of probes and briefs [16].

Although the methods of design studies have influenced our approach (for example, our study probes did discretise design moves), we did not pursue formal protocol or linkographic analysis, partly due to their logistical and labour overheads. More fundamentally, however, we lacked a model of DMI design on which to base any such protocol analysis. Instead, we opted to begin with observations, from which models could potentially be hypothesised later.

3. DMI ACTIVITY & PROBE DESIGN

As previously mentioned, we have discussed this study in another publication, focusing on the themes of bricolage and liveness in digital and hybrid lutherie [7]. A comprehensive account of the workshop activity and environment, instrument design and editing workflow, as well as the participants, groupings, and data collection methods can be found in that paper. In this section, we provide a summary of the aforementioned aspects and emphasise the elements specific to the analysis methods presented.

A one-hour workshop activity involved four groups of three DMI designers working with a simple modular probe requiring creative and technical intervention (Figure 1). Groups were chosen primarily to enable comparison with another previous study [17]. The goal was to explore the instrument and develop its character, gaining insight into how the material environment affects design idea generation, exploration, decision making, and development. The brief for the one-hour activity was simply to continue “finishing” the DMI in any manner the group deemed appropriate. The point was not to see if participants could complete this brief within one hour, which is impossible, but to see how they would approach the task and what they would do with the time available. For more on our motivations towards

one-hour activities, see [18]. Participants later summarized their work, completed a survey, and were debriefed.

The *Unfinished Instrument*, constructed from laser-cut modular parts, houses four microphone capsules connected to a Bela device running a Pd patch (Figure 2). Interaction with the mics excites four Karplus-Strong vibrating string models. The Pd patch preprocesses each mic signal and mixes the four string sounds, providing debugging facilities.

Pd usually allows objects and connections to be added, edited, or removed while running. However, the patch was running on a remote Bela device, so the workflow for editing differed. The update procedure involved:

1. Editing the patch using the desktop Pd application.
2. Saving the patch to update the instrument.
3. Automatically copying the patch to the embedded device and recompiling.
4. Continuing to interact with the instrument after recompilation.

3.1 Data Capture and Annotation

To capture the design process data, we utilised Sparkle-share¹ to automatically version control the Pd patches whenever the participants saved their work, a necessary step for updating the Bela devices they were working with. To analyse this data, we initially extracted the timestamp and size of change from each git commit. Additionally, we checked out individual git commits at the conclusion of each session and exported them as vector-based PDF files, which we could then manually annotate using another software. For this study, we opted for a simple annotation schema based on the positions and edit status of the Pd objects.

- Light grey annotations indicate the original position of the Pd objects, as all groups began with the same patch.
- Dark grey annotations signify objects that have been moved.
- Green annotations represent newly added objects.

¹ <https://www.sparkleshare.org/>

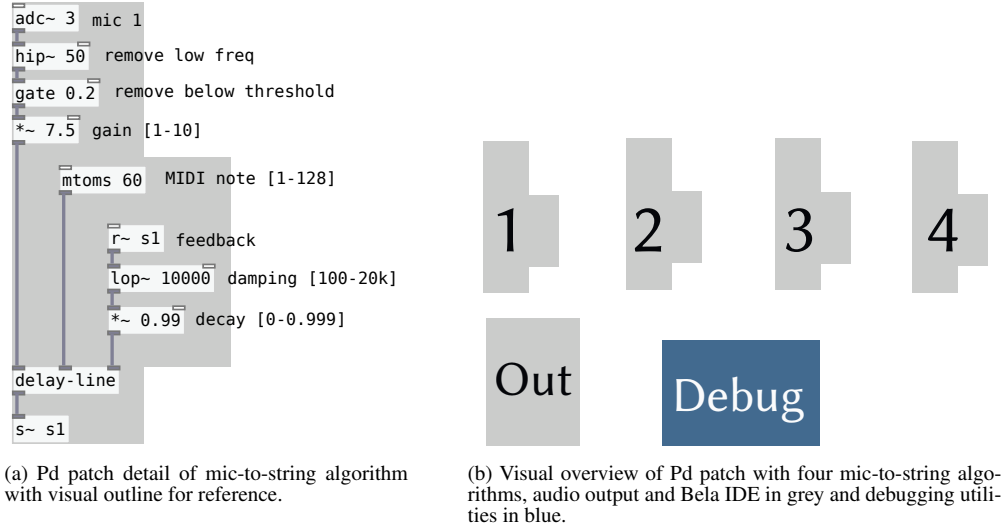


Figure 2: Pure Data patch mic-to-string algorithm detail (left) and patch overview (right).

- Red annotations denote existing objects that have been edited.
- Blue annotations highlight the debugging tools present in the original patch.

Additionally, we gathered information about the participants' level of experience with Pd / Max/MSP and conducted a post-activity survey for each participant, inquiring about their process and outcomes. To summarise the relative level of experience for each group:

G1 had the least experience with Pd.

G2 each possessed at least one year of experience.

G3 all had some Pd experience.

G4 included two out of three participants with experience.

	Total updates	Updates / min	Avg. size	Comp. time (%)
G4	46	0.76	20.5	8.9
G2	44	0.73	11.5	8.6
G3	30	0.50	32.5	5.8
G1	21	0.35	52.0	4.1

Table 1: Summarial statistics for each group, ordered by number of total updates (first column).

4. OUTCOMES

Figure 3 illustrates the final states of the Pd patches for each group. The overall architecture of the patches remained unaltered. G1, G2, and G4 utilised the debugging facilities (blue) of the Bela IDE. G1, G2, and G4 modified (red) the threshold of the sensors to achieve enhanced sensitivity. G1 and G2 endeavoured to convert sensor signals into control structures (green), while G4 incorporated additional delay lines (green). G3 also attempted to introduce extra delay lines (green), but removed their work before the conclusion.

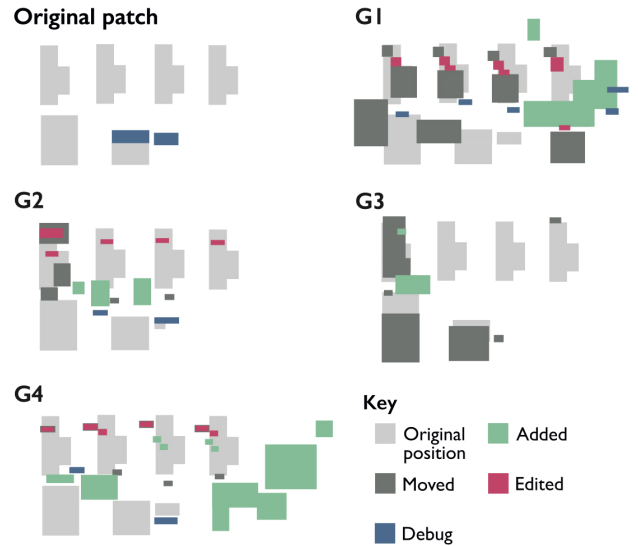


Figure 3: Visual summary of the final states of the Pd patches for each group.




Figure 4: Pd patch updates over the 60 minute activity period for each Group. Bar height represents git commit size.

Figure 4 and Table 1 illustrate the frequency and size of patch updates throughout the 60-minute activity. During the initial ten minutes, only a few updates occurred, with the rate and size of updates generally increasing towards the end of the activity. Most groups exhibited consistency in the frequency and size of their updates, except for G1.

G1 made infrequent and large updates, and ceased editing for nearly 20 minutes in the middle of their session. Both G2 and G4 had the highest number of updates; however, G4's updates were almost twice the size of G3's on average. G3 had the second-largest average update size. Since updates took approximately seven seconds from saving the patch to hearing the output, G4 spent almost 9% of their session waiting for updates to occur.

4.1 Types of Design Interventions

Comparing Figure 3, Figure 4, Table 1, and Table 2 provides a comprehensive view of the various design interventions that took place, as well as how these were influenced by the probe and activity design choices, in addition to the group collaboration constraints. Overall, there was a noticeable absence of high-level design exploration throughout the sessions. High-level interventions might have encompassed re-architecting the patch, modifying existing abstractions or developing new ones, or altering the physical model's structure and parameters. Aside from the time constraints, this indicates that the scenario was ineffective in facilitating such operations. The lack of interaction with the damping and decay parameters was unexpected, given that they were relatively accessible and significantly affected the timbre and envelope of the sound.

The most prevalent types of design intervention involved mappings. As detailed in [7], examples of these encompassed generating control signals from audio signals and utilising them to modify mappings, cross-mapping sensors and delay lines in conjunction, and incorporating additional delay lines and mapping those atop the existing patch. While some of these alterations did significantly impact timbre, the resulting DMIs remained sufficiently similar across the group to facilitate comparison of their mappings. Connecting wires in a semi-random manner necessitates minimal forethought, is straightforward to execute, and can yield novelty, which may partially account for the inclination towards this type of activity.

In terms of more nuanced aspects, most groups adjusted the input signal processing parameters in an effort to enhance the responsiveness of the inputs. The groups did not delve beyond modifying the parameters for gate threshold and gain. It is likely that they considered the improvements attained to be 'good enough' to proceed with, given the limited time available for exploration. Although there were instances of groups finalising certain changes and attempting to refine them further, they did not make significant progress in this regard.

4.2 Outcomes by Group

Table 2 presents quotes from each group, which can be further compared with their individual design activities. G1 initially expressed frustration with the instrument's quality

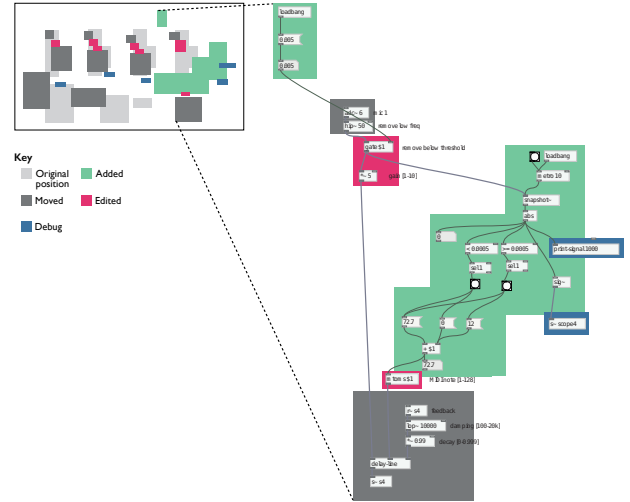


Figure 5: G1's Pd patch detail.

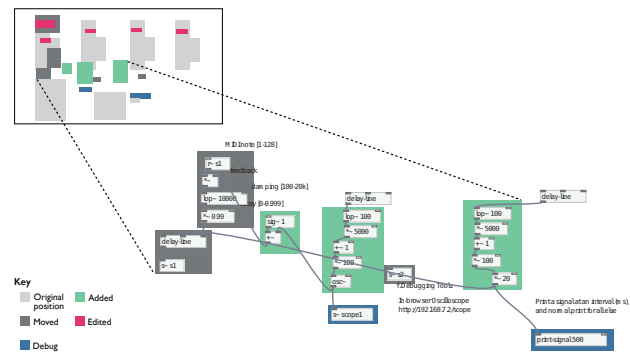


Figure 6: G2's Pd patch detail.

and the workflow's speed. They experimented with optimising sensitivity and altered the tonality of pitches from major to minor. Halfway through the session, they resorted to pen and paper to sketch out their idea and its implementation (see the gap in the middle of Figure 4). This period lasted 20 minutes, after which they encountered difficulties in debugging their implementation. They attempted to convert a sensor signal into a control signal by thresholding it, intending to control a pitch-shifting effect (Figure 5). Ultimately, they were somewhat frustrated with their result, stating that the process was too slow or that they needed more time.

G2 initially expressed dissatisfaction with the sensor response, yet remained enthusiastic about exploring the instrument. They employed cross-mapping of sensors to enhance the sound, but encountered difficulties in deriving control values from the sensors. To achieve variety, they reduced the delay time, modified the string pitches, and experimented with acoustic feedback (Figure 6). By the end, they believed they had made some progress, but felt that the group process and the update procedure hindered their pace.

G3 analysed the instrument architecture and discussed ideas for modifying it. Verbal discussion was prominent in their group process, which occasionally appeared to result in hesitation and withdrawal. At times, they were silent

	Goals	Group Process	Design Process	Outcomes
G1	<i>Find a fun way to play the physical thing, adapt to that.</i> <i>More melodic, harmonic control.</i>	<i>Difficult [...] in the flat hierarchy there was a tendency for agenda pushing.</i> <i>Efficient, once we realised the limitations [...] we found a common goal.</i>	<i>There was a tendency to focus on screen, rather than playing.</i> <i>Slower than desired. Pd became too high level for some operations.</i>	<i>It ended up a multi-interactive instrument rather than a hit instrument.</i> <i>[...] we didn't have time to develop our ideas and experiment.</i>
G2	<i>Just to explore and see what's possible.</i> <i>Find out what could be interesting.</i>	<i>We all took turns to explore the instrument and added ideas on each topic.</i> <i>I usually work alone. The group element affected the ideas being made.</i>	<i>It developed through small changes [...] Little steps at a time.</i> <i>Compile stage limited flow. A desktop patch w/ OSC streams would be good.</i>	<i>I like the instrument more now... But we didn't get to the final creation.</i> <i>Brainstorming amongst experienced people is helpful for this kind of project.</i>
G3	<i>Our goals focused on the interaction (how to play the instrument).</i> <i>Improve existing "problems" (unresponsive to touch) and create new sounds.</i>	<i>Perhaps better suited to working independently.</i> <i>The need for approval of team meant that ideas were held back.</i>	<i>Partly based on a deterministic approach. Partly a "trial/error" method.</i> <i>Tried to test fast and decide based on the results.</i>	<i>From a keyboard to an instrument with continuous and discrete interaction.</i> <i>Only a minor improvement compared to the starting point.</i>
G4	<i>To make interesting sounds with a definable interactive capabilities.</i> <i>Explore sounds and interactions made inspired by the physical setup.</i>	<i>I ended up being the player, the others were focused on the design.</i> <i>Much more intentional and planned with the group than if I was alone.</i>	<i>It developed however lack of debug lead to some issues.</i> <i>We found ways to actuate the sensors tapping w/ pencil, singing, etc.</i>	<i>Functionality was built but not really fine tuned.</i> <i>Would have loved more time. Didn't even start on pitch manipulation.</i>

Table 2: Selected quotes from different members of each group about the activity, based on a post-activity survey.

and did not interact with the instrument. Their primary concept involved combining discrete and continuous interactions to produce a rich and sustained sound. However, the overall outcomes were deemed to be insignificant, and the technical aspects of the activity seemed to be in conflict with the collaborative aspects.

G4 thoroughly investigated the interaction possibilities of the instrument. In contrast to G1-3, they displayed a particular interest in the instrument's physicality. Participant C, who had no prior experience with Pd, opted to assume the role of a performer, while participants A and B modified the patch. The group was quick to generate ideas and was the only group to implement a significant patch update within the first ten minutes. Their goal was to create sustained sounds, prompting them to explore the addition of delay lines and acoustic feedback. Participants A and B believed they made progress, but lacked the time to refine their work. Participant C provided a performer-centric perspective on the activity.

5. DISCUSSION

5.1 Findings about DMI Design Process

In this paper, we introduce various design data metrics, which can be compared with qualitative data to gain insights into the DMI design process. Despite having only one hour to work with, we observed numerous intriguing patterns in the design process data. From the annotated visual summaries, we gleaned that DMI designers quickly identify points of high leverage, as they demonstrated by adjusting the microphone sensitivity. From the summarial statistics, we observed that groups making smaller, more frequent changes achieved greater flow in their design process. The Pd patch update frequency versus size plots allowed us to identify periods of ideation and discussion, as

well as instances where groups engaged in more hands-on design activity. Overall, our interpretation of this study [7], along with other studies [1], highlights how the available tools significantly influence the design process.

In particular, comparing the design data with qualitative data confirms findings from our previous work [7], regarding the impact of liveness on DMI design and the influence of group dynamics on the design process. Both of these works emphasise that practitioners face considerable challenges with current DMI design tools in terms of rapid iteration and real-time collaboration. We provide a detailed critique of these issues in [19], and based on this critique, we propose designing hybrid handcrafting tools and comparing how digital luthiers work with them [20].

5.2 Method Limitations and Improvements

In this study, we conducted a high-level comparison across DMI design groups, without delving into more granular observations at the level of design episodes or design moves. This was primarily due to time constraints, and further investigation is necessary to validate these methods at all levels of design.

Manually saving the Pd patch was essential in this case as saving triggered updates to the Bela device, and additionally, a save typically marked a design move. However, automating this process in the background could potentially yield even finer data. This would enable breaking down discretised design moves [14] into continuous design gestures, for example to capture hesitancy, backtracking, and other design behaviours. Such an approach would also allow for finer temporal alignment between discussion and design activity.

We utilised git commit size as a proxy for the scale and frequency of design moves, but this approach requires further validation and could be improved in various ways. The

annotation process was performed manually, but it could potentially be automated, perhaps by using one of the GUI forks of Pd to embed annotations directly within the editor. Custom visualisations for real-time display of design process data have great potential not just as a design feedback tool, but also for pedagogical supervision [21].

The temporal plots of Pd design activity revealed gaps when groups were away from the keyboard and mouse, engaged in ideation and discussion. It would be interesting to plot metrics such as words per minute, or even conduct protocol or linkographic analyses of the discussions, alongside the Pd design process data.

Lastly, we believe that the methods presented in this study could also be effective for textual programming languages (for example, SuperCollider has a history class), natural language programming with large language models (LLMs) [?] and other types of digital lutherie tools including hybrid handcraft interfaces [6,20]. This would be possible as long as there is a reliable way to snapshot the design space and the designer’s intentions throughout the session.

5.3 Validity

In terms of the validity of our annotations, we employed an exceedingly straightforward empirical annotation schema in this instance, which showcased activity as raw editing of the Pd patch. Consequently, it was not difficult for an individual researcher to determine validity. If we were to implement a more advanced design analysis protocol on the design data, such as the function-behaviour-structure (FBS) ontology [22], the validity of the annotations might be compromised by individual researcher training with the protocol and potential biases. Such an analysis would necessitate training more than one researcher in protocol-based annotation and establishing an appropriate arbitration process to resolve disagreements. Although this approach is well-documented in design and other fields, it could prove more challenging to apply in DMI design studies, particularly when more subjective aspects of design are considered, as well as in cases involving musical gesture, interaction, and performance.

More generalised protocols might not adequately encapsulate these types of activities, and domain-specific protocols may first need to be proposed and validated, which is a more ambitious but rewarding objective. We have elsewhere proposed a scale-based ontology of DMI design that defines macro, meso and micro scales of design details [19], and we apply this ontology to compare this study with others [6, 17] in [18].

Another approach to validating annotations involves comparing them with participant verbalisations during the activity via think-aloud instructions, or interview or survey data after the activity, which we partially implemented in this study. However, think-aloud methods are ultimately of limited use during DMI design activity, which is inherently tacit and embodied, and deals with subtle nuances and emotional responses. Instead, activity participants could review, narrate, and annotate their session videos during the post-activity interview (perhaps even in slow motion), employing video-cue recall, which has already been used

in an augmented DMI evaluation setting [23]. Video-cue recall could even be integrated into the study probes and interfaces themselves, enabling real-time recall and review during the activity, which would, of course, affect the design process itself, perhaps in intriguing ways.

5.4 Scalability

Addressing the scalability of these methods, we commenced this paper by emphasising the absence of standard research approaches in this domain. A method is not scalable if only specific research groups can utilise it, due to limited access to appropriate resources or (translated) documentation. We perceive that the most scarce resource across all research groups in DMI design is the researcher’s time. We have endeavoured to use this resource effectively, primarily by demonstrating that the methods can be employed in brief, closed activities, which are the least burdensome for researchers, as they utilise the same setup and location, and do not require an extended duration to execute. The other significant time burden for researchers lies in the analysis phase. We have aimed to emphasise automation in the collection of design process data, and in the preceding subsection, we address how automation of annotation could be enhanced. To supplement this theme, we believe that if sufficient high-quality DMI design process data were collected and published as a dataset, machine learning models could be trained and distributed, enabling researchers worldwide to instantaneously annotate their data. However, we believe the issues discussed subsequently would need to be addressed first.

5.5 Replicability

We do not propose that researchers should attempt to replicate this specific activity directly; instead, we recommend that future work employ our methods as a foundation for designing replicable studies observing DMI design, addressing the limitations we have already emphasised. Another methodological issue affecting replication in our work was the presence of the researcher in the activity, which was necessary for methodological and technical reasons in this instance, but is nonetheless problematic in the way it influences participant behaviour. Eliminating the need for this would enhance the chances of replication in terms of comparing across groups within a study, as well as comparing across studies, which is the ultimate objective. Methodologically, researchers can still observe design sessions in real-time remotely via camera, although this may still disrupt the design activity. Technically, researchers can minimise the need to teach, assist, and support participants during sessions by incorporating necessary resources into the probes and activity environment, which in turn requires more pilot testing and refinement of study probes. Revisiting researcher training, research methods workshops would serve as an excellent means for researchers to experiment with the methods and discuss them in detail.

5.6 Generalisability

The primary challenge in generalising the proposed methods lies in their closed and concise nature, which we have

emphasised as a necessary requirement for scalability, but this fact cannot be ignored. The brevity of these methods results in design process data being skewed towards familiarisation and exploration. It is difficult to envisage filtering or abstracting these responses; however, if such a process were possible, it would render study outcomes more comparable to actual practice.

Numerous significant DMI design studies take place “in-the-wild” [24–26], involving open-ended activities and longitudinal encounters between various ecologies, including makers, musicians, composers, audiences, and more. Naturally, a comprehensive account of the DMI design process must also encompass longitudinal investigation, particularly to understand the evolution of these processes and the gradual development of expertise related to the perception and manipulation of subtle details.

In such scenarios, designing customised DMI design probes may not even be necessary. Instead, the DMI designer practitioner’s own working environment and DMI tools and prototypes could be considered as probes, which could be augmented to produce valid DMI design process data. A prime opportunity to study DMI design in practice would be to investigate the replication of the exact same DMI by the same practitioner or by a group of practitioners [27]. Demonstrating the effectiveness of these methods in-the-wild would then enable hybrid study designs that combine short, closed activities with potentially unfamiliar probes, alongside extended and situated practice.

Our work primarily focused on design processes. Nevertheless, we would like to emphasise that the methods we propose may be equally applicable to creativity studies or musicological analysis of the genesis of a musical composition. Conventional musicological analysis has acknowledged the significance of sketches, notes, and experiments undertaken by the composer. Consequently, the visual analysis of the history and experimentation within the creative process could serve as a valuable supplement to such research.

6. CONCLUSION

There is no diagram or model for the subtle and nuanced craft, which is best left as “undecidable” [28]. No two DMIs are identical either. Thus, how can we make any kind of generalisation regarding DMI design by comparing the design process across them? The aim of this research is not to “solve” DMI design or propose a single design model that explains everything, any more than the goal of music cognition and psychology is to solve music. We view all models as culturally situated and contextual, and the extent to which they generalise depends as much on the naivety of the researcher proposing or using them as on any inherent universality. Nonetheless, we believe that emphasising the comparison of DMI design *processes* enables researchers to circumvent, to a degree, the non-generalisability of the *content* of the DMIs themselves, which, being necessarily specialised, is esoteric to each instrument and practice.

In this research, our primary focus has been on the DMI design processes rather than the tools. However, it is ev-

ident from the literature that each tool embodies its own epistemology [12], idiomaticity [2], and mythology [3]. Can the methods presented in this study provide a novel approach to empirically compare the influences of each DMI design tool’s affordances? Do expert DMI designers develop strategies to overcome these influences, and is it possible to demonstrate this by comparing their design processes with those of novices, beginners, and intermediates? More broadly, do commonalities in DMI design processes exist within specific geographies or cultures, or even across them? Given that digital luthiers often design for themselves as musicians, how can we capture and compare embodied DMI design cognition with embodied music cognition? Addressing these questions will require a substantial amount of rigorously scrutinised and unambiguous evidence, but it is likely to reveal many intriguing insights along the journey.

Acknowledgments

Thanks to the anonymous reviewers for their helpful comments and suggestions.

This research is supported by the European Research Council (ERC) as part of the Intelligent Instruments project (INTENT), under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No. 101001848).

This research was also supported by EPSRC in the UK, under the grants EP/L01632X/1 (Centre for Doctoral Training in Media and Arts Technology) and EP/N005112/1 (Design for Virtuosity).

7. REFERENCES

- [1] A. McPherson and G. Lepri, “Beholden to our tools: Negotiating with technology while sketching digital instruments,” in *Proceedings of the International Conference on New Interfaces for Musical Expression*, R. Michon and F. Schroeder, Eds. Birmingham, UK: Birmingham City University, Jul. 2020, pp. 434–439.
- [2] A. McPherson and K. Tahiroğlu, “Idiomatic Patterns and Aesthetic Influence in Computer Music Languages,” *Organised Sound*, p. 20, 2019.
- [3] J. Snape and G. Born, “Max, Music Software, and the Mutual Mediation of Aesthetics and Digital Technologies,” in *Music and Digital Media: A Global Anthropology*, G. Born, Ed. In Preparation.
- [4] R. Jack, J. Harrison, and A. McPherson, “Digital musical instruments as research products,” in *Proceedings of the International Conference on New Interfaces for Musical Expression*, R. Michon and F. Schroeder, Eds. Birmingham, UK: Birmingham City University, Jul. 2020, pp. 446–451.
- [5] R. Jack, “Tangibility and Richness in Digital Musical Instrument Design,” Thesis, Apr. 2019.
- [6] J. Armitage, T. Magnusson, and A. McPherson, “Studying Subtle and Detailed Digital Lutherie: Motivational Contexts and Technical Needs,” in *Proc. New Interfaces for Musical Expression*.

- [7] J. Armitage and A. McPherson, “Bricolage in a hybrid digital lutherie context: A workshop study,” in *Proceedings of the 14th International Audio Mostly Conference: A Journey in Sound*. Nottingham United Kingdom: ACM, Sep. 2019, pp. 82–89.
- [8] J. Armitage, F. Morreale, and A. McPherson, “The finer the musician, the smaller the details: NIMEcraft under the microscope,” in *Proc. New Interfaces for Musical Expression*. Copenhagen, Denmark: Aalborg University Copenhagen, 2017, pp. 393–398.
- [9] T. J. Barraclough, D. A. Carnegie, and A. Kapur, “Musical Instrument Design Process for Mobile Technology,” 2015.
- [10] F. Calegario, M. M. Wanderley, S. Huot, G. Cabral, and G. Ramalho, “A Method and Toolkit for Digital Musical Instruments: Generating Ideas and Prototypes,” *IEEE MultiMedia*, vol. 24, no. 1, pp. 63–71, Jan. 2017.
- [11] T. J. West, M. Wanderley, and B. Caramiaux, “Making mappings: Examining the design process,” in *Proceedings of the International Conference on New Interfaces for Musical Expression*, R. Michon and F. Schroeder, Eds. Birmingham, UK: Birmingham City University, Jul. 2020, pp. 291–296.
- [12] T. Magnusson, “Epistemic Tools: The Phenomenology of Digital Musical Instruments,” Ph.D. dissertation, University of Sussex, 2009.
- [13] S. Delle Monache and D. Rocchesso, “Cooperative Sound Design: A Protocol Analysis,” in *Proceedings of the Audio Mostly 2016*, ser. AM ’16. New York, NY, USA: ACM, 2016, pp. 154–161.
- [14] G. Goldschmidt, *Linkography: Unfolding the Design Process*. MIT Press, 2014.
- [15] J. W. Kan and J. S. Gero, “Using the FBS ontology to capture semantic design information in design protocol studies,” in *About: Designing. Analysing Design Meetings*. CRC Press, 2009, pp. 213–229.
- [16] —, “Acquiring information from linkography in protocol studies of designing,” *Design Studies*, vol. 29, no. 4, pp. 315–337, 2008.
- [17] J. Armitage and A. McPherson, “Crafting Digital Musical Instruments: An Exploratory Workshop Study,” in *Proc. New Interfaces for Musical Expression*, Blacksburg, Virginia, USA, 2018.
- [18] J. Armitage, “Subtlety and detail in digital musical instrument design,” Thesis, Queen Mary University of London, Apr. 2022.
- [19] J. Armitage, T. Magnusson, and A. McPherson, “A Scale-Based Ontology of Digital Musical Instrument Design,” in *Proc. New Interfaces for Musical Expression*.
- [20] —, “Sculpting Algorithmic Pattern: Informal and Visuospatial Interaction in Musical Instrument Design,” in *Proc. Sound and Music Computing Conference*.
- [21] A. G. Zhang, Y. Chen, and S. Oney, “VizProg: Identifying Misunderstandings By Visualizing Students’ Coding Progress,” in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, ser. CHI ’23. Association for Computing Machinery, pp. 1–16. [Online]. Available: <https://dl.acm.org/doi/10.1145/3544548.3581516>
- [22] K. Dorst and P. E. Vermaas, “John Gero’s Function-Behaviour-Structure model of designing: A critical analysis,” *Research in Engineering Design*, vol. 16, no. 1-2, pp. 17–26, Nov. 2005.
- [23] F. Morreale, A. Guidi, and A. P. McPherson, “Magpick: An augmented guitar pick for nuanced control,” in *Proceedings of the International Conference on New Interfaces for Musical Expression*, M. Queiroz and A. X. Sedó, Eds. Porto Alegre, Brazil: UFRGS, Jun. 2019, pp. 65–70.
- [24] Y. Rogers, “Interaction design gone wild: Striving for wild theory,” *Interactions*, vol. 18, no. 4, pp. 58–62, 2011.
- [25] L. S. Pardue, “Violin augmentation techniques for learning assistance,” PhD Thesis, Queen Mary University of London, 2017.
- [26] J. Harrison, “Instruments and access: The role of instruments in music and disability,” Ph.D. dissertation, Queen Mary University of London, London, 2020.
- [27] L. Zayas-Garin, J. Harrison, R. Jack, and A. McPherson, “DMI Apprenticeship: Sharing and Replicating Musical Artefacts,” in *International Conference on New Interfaces for Musical Expression*. PubPub, Apr. 2021.
- [28] S. Kettley, “The foundations of craft: A suggested protocol for introducing craft to other disciplines,” *Craft Research*, vol. 3, no. 1, pp. 33–51, 2012.