The Divergent Interface: Supporting Creative Exploration of Parameter Spaces

Robert Tubb
Centre for Digital Music
Queen Mary University of London
London, E1 4NS
r.h.tubb@qmul.ac.uk

Simon Dixon
Centre for Digital Music
Queen Mary University of London
London, E1 4NS
simon.dixon@eecs.qmul.ac.uk

ABSTRACT

This paper outlines a theoretical framework for creative technology based on two contrasting processes: divergent exploration and convergent optimisation. We claim that these two cases require different gesture-to-parameter mapping properties. We present results from a user experiment that motivates this theory. The experiment was conducted using a publicly available iPad app: "Sonic Zoom". Participants were encouraged to conduct an open ended exploration of synthesis timbre using a combination of two different interfaces. The first was a standard interface with ten sliders, hypothesised to be suited to the "convergent" stage of creation. The second was a mapping of the entire 10-D combinatorial space to a 2-D surface using a space filling curve. This novel interface was intended to support the "divergent" aspect of creativity. The paths of around 250 users through both 2-D and 10-D space were logged and analysed. Both the interaction data and questionnaire results show that the different interfaces tended to be used for different aspects of sound creation, and a combination of these two navigation styles was deemed to be more useful than either individually. The study indicates that the predictable, separate parameters found in most music technology are more appropriate for convergent tasks.

Keywords

Creativity, Sound Synthesis, Mapping, Interfaces

1. INTRODUCTION

Much research into gestural control of synthesiser parameters has focussed on live expressive performance. A persuasive case has been made that one-to-one mappings found in the standard synthesiser interface are not optimal for this scenario [19, 13]. This research aims to apply a similar investigation to a different aspect of electronic music creation: that of creative sound design in the studio. A typical synthesiser or digital audio workstation (DAW) has tens, hundreds or even thousands of parameters. The sheer size of the combinatorial space of synthetic sound raises interesting questions: how do people navigate this vast space? What interface designs help or hinder "creative" cognition? Could the paths that people take through the space (i.e. the interaction data) provide a trace of the creative process in action?
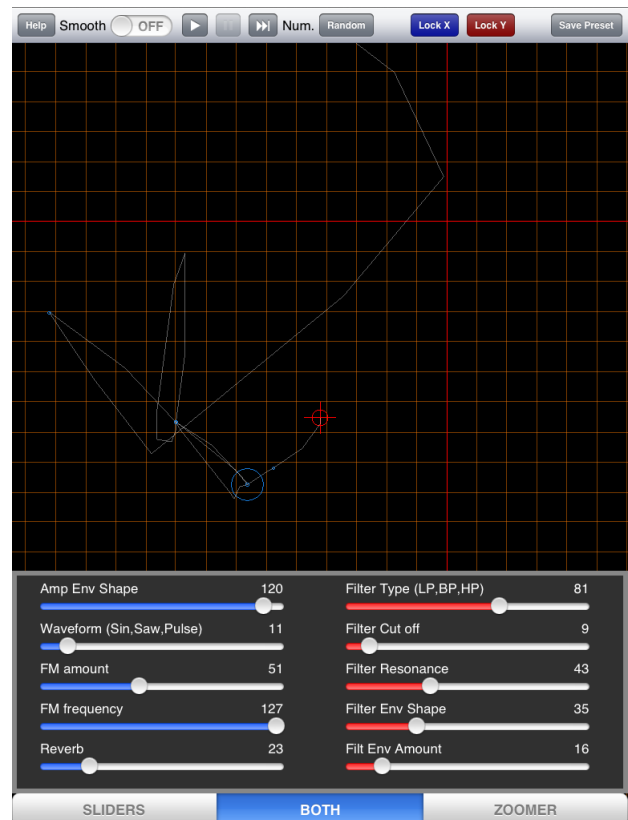
Figure 1: Sonic Zoom screenshot. The red crosshairs show the point corresponding to the slider settings, so scrolling the surface alters the synth timbre. The users path can be seen as a white line, with blue circles indicating points previously listened to.

In section 1.1 a model of creative processes is outlined. Section 1.2 then applies this model to the case of a finite continuous parameter space. Consideration is given to parameter mappings and dimension reduction techniques in section 1.3. An experiment investigating these issues is described in section 2.

1.1 Creative Cognition

J.P.Guilford [9] was one of the first prominent psychologists to draw attention to creativity as something that could be studied scientifically. He characterised the creative process as a combination of "convergent" and "divergent" thinking. Divergence is the generation of many provisional candidate solutions to a problem, whereas convergence is the narrowing of the options to find the most appropriate solution.

Whilst this is a simple model of an immensely complex phenomenon, most modern theories have similar processes present in some form, sometimes referred to by different names such as "Generative" and "Evaluative". Campbell [3] and Simonton [20] have considered creativity as a Darwinian process, and similarly propose a process of idea mutation (or recombination) and idea selection.

The Geneplore model [27] also features two complementary processes. The "Gene-" stage is the generation of "pre-inventive structures", fluid collections of provisional, experimental concepts. The "-plore" stage is the exploration and evaluation of those concepts in more detail. Flexible alternation between these two processes is important.

Whilst creativity is still a controversial topic, the most basic definition is relatively well agreed on: that creative products are both *novel* and *valuable*. Divergence could be seen as novelty generation, and convergence the identification and/or optimisation of value. This value could be some objectively measurable quantity (in science and engineering applications) or a subjective aesthetic quality (artistic applications).

1.2 Conceptual/Parameter Space Search

Creativity is also being studied in the context of artificial intelligence: a field known as Computational Creativity. By attempting to build artificial systems that exhibit creative behaviour, we may form models of how creativity might function in our own minds. Wiggins' Creative Systems Framework (CSF) [28] is a more formal descendent of Boden's theories of artificial creativity [2]. In this framework creativity is seen as a way of extending conceptual space: using a traversal mechanism that produces a concept falling outside of the existing space (an "aberration"), but is nevertheless seen as valuable and appropriate according to the evaluation function of the domain.

This research aims to find out what traversal mechanisms are most effective for generating aberrations, and hence innovations in sound design and music. There are good reasons why music technology should be a useful tool for studying the creative process. "Liveness", i.e. the ability to hear results in near real-time is important for creative flow [17], therefore many electronic musicians carry out their creative work whilst actually manipulating some kind of interface. In addition, most abstractions of musical data are handled within the machine, having a large effect on the mental representation of the material [4]. Therefore the parameter space of the sequencers, synthesisers and effects is closely linked to to the conceptual space of the electronic music, hence could be considered as a microcosm of a conceptual space in general.

Considering this framework, we can further pin down what divergent and convergent strategies might mean. Convergence, as the word's use in optimisation literature would suggest, is analogous to gradient descent. A solution is continuously improved until it appears to sit at the optimum of some measure of fitness. Possible evidence for this kind of behaviour in musical parameter space is presented later. Divergence, however, is a more stochastic exploration of the space (uninformed by the fitness measure) such that local optima can be escaped, and new regions can be explored and subsequently optimised.

An interesting critique of convergence-only algorithms is found in neuroevolution literature [15]. This research demonstrated that, for even well-defined problem spaces such as maze navigation and simulated biped walking tasks, novelty search can significantly outperform objective-based search. In domains where originality is often seen as intrinsically desirable, such as art and music, novelty driven exploration

	Div.	Conv.	Perf.
Predictability & Control	×	✓	✓
Separability	×	✓	×
Distance Preservation	✓	✓	✓
Location Preservation	✓	✓	✓
Dimension Reduction	✓	×	✓
Speed	✓	×	✓
Space Reduction	×	×	✓

Table 1: Desirable properties of controller mappings for different creative stages: divergent exploration (Div.), convergent honing (Conv.) and expressive performance (Perf.).

should be even more appropriate. Evaluating musical interfaces by exclusively testing goal directed behaviour [26] may therefore miss some of the picture.

A final consideration for navigating complex search spaces is that of cognitive load. There is considerable evidence that two decision-making systems exist within the brain. One fast, parallel, inflexible and unconscious (implicit or system 1), and the other slower, serial, flexible and deliberate (explicit or system 2) [5] [14] [21]. Given the rapid nature of system 1, it clearly behoves the interface designer to enable implicit processing wherever possible, in order to free up explicit resources that may be required to engage in higher-level artistic decisions.

1.3 Mapping, Dimensionality Reduction and Divergent Strategies

Given an interface, and a variety of parameter to gesture mappings, an interesting question to ask is how different mapping geometries may enable a user to traverse and evaluate the parameter space, such that local minima can be escaped and optimal locations found quickly.

A detailed analysis of mapping techniques is beyond the scope of this paper; good treatments of the geometry of performance mappings are given in [7] and [23]. Since it is difficult to provide any mapping that satisfies all criteria for all situations, it is useful to indicate how they become more or less important for different stages of music creation. Table 1 enumerates three creative stages, and which mapping properties may suit them best. For example, predictability of a mapping (a result of geometrical properties such as linearity and smoothness) is clearly important for both converging on solutions and live performance: when unpleasant surprises are unwelcome. However the separability of parameters can have a negative effect on performance [12], but certainly becomes useful when fine tuning details.

Dimension reducing mappings are useful for both divergence and performance. In particular, there are numerous advantages to a 2-D representation of a sound space. There is a compelling metaphor of exploration of physical terrain. The gestural control can be made completely consistent with maps applications, a widespread, familiar and efficient interaction style. Activity within the space is able to be visualised, such as the path that has been explored. This builds a memorable "geography" of the sound-space, and may take advantage of specialised areas of the brain dedicated to visuospatial processing. Favourite presets can appear as points in space, and can be recalled in a more integrated way than using a drop down text list. Previous efforts in this area take a timbre space approach, generated by multi-dimensional scaling (e.g. SoundExplorer [29] and ISEE [25]). Alternatively a 2-D subspace can be generated by interpolating between existing preset points (e.g. Bencina's metasurface [1], the preset explorer in [24] and

AD 1	I am familiar with music software and sound synthesis.
AD 2	The ability to retrace my steps using the history path was useful.
AD 3	The correspondence between the sliders and the grid was understandable.
AD 4	Scrolling a greater distance on the grid seemed to correspond to larger difference in the sound.
AD 5	The ability to see other presets laid on the grid was useful.
AD 6	The range of sounds was too limited/poor quality to be able to judge the eventual usefulness of the interface.
AD 7	The Zoomer was an improvement on just using a randomiser.
AD 8	The combination of Zoomer and Sliders was better than either individually.
AD 9	I enjoy "happy accidents" in the creative process

**Table 2: Questions requiring a 5 point agree/disagree answer.**

the "nodes" object in Max/MSP). One criticism of many of these techniques from a creative systems point of view is that they build the low-dimensional space from pre-existing favourites, and render the rest of the space inaccessible. Therefore they may lower the probability of discovering aberrant points that are essential for novelty generation.

Another possible advantage of dimensionality reduction is that it can confound the separability of individual dimensions, and encourage users to shift from an analytic to a holistic processing mode [12]. This may take advantage of the fast, intuitive parallel-processing brain system, and free up working memory for other tasks, such as critical listening [16]. Given a simpler method of auditioning sounds, producers may be more able to evaluate the overall timbre and the way it fits into the music as a whole, rather than concentrating on manipulating just one aspect of it.

So, let us define a "divergent interface". It should enable speed, dimension reduction and repeatability, whilst preserving access to all possibilities. It may intentionally sacrifice predictability and separability. If the musician is looking to the instrument for inspiration, it makes no sense to enforce predictability. Put simply, if you do not know where you want to go it scarcely matters if you don't know how to get there. Preserving locality may be useful, such that users can explore the neighbourhood of a sound to create elaborations. Another important consideration is that once an interesting sound has been discovered, it should not incur too much effort to swap to a convergent interaction mode. The next section describes an attempt to create and evaluate such an interface.

2. THE "SONIC ZOOM" EXPERIMENT

Sonic Zoom is an iPad app made publicly available on the Apple App Store. In the application two interfaces are provided. The first is a reasonably standard set of ten sliders (sending 7-bit MIDI continuous control (CC) values), used to control the timbre of a subtractive synthesiser. The second interface is a scrollable, zoomable surface: a map of every possible slider combination (referred to from now on as the "Zoomer"). The mapping from the high dimensional parameter space to the two dimensional surface uses two space-filling curves. Hilbert curves [10] constructed from Gray codes [8] with long bit runs [6] can be used to preserve locality as much as is possible, whilst still maintaining ac-

Figure 2: Questionnaire responses to agree/disagree Likert items. Neutral response is centred.

cess to the entire space. A more detailed description of this Hilbert curve mapping can be found in [22].

Since the Hilbert curve maps the entire space, scrolling in 2-D immediately corresponds to movements of the ten sliders and vice-versa. Of course, the space is huge: in the case of ten 7-bit midi parameters, each axis contains $2^{5 \times 7} \simeq 10^{10}$ points. Zooming functionality is therefore essential: a pinch-out gesture zooms into a smaller area. Due to the Hilbert curve's locality properties, this area will correspond to a smaller 10-D hypercube which can then be explored in further detail. This means convergent navigation is also provided by the zoomer. The disadvantage of the Hilbert curve is its lack of linearity: it is impossible to predict what scrolling in a certain direction will do to the sound.

A user study was performed to compare the Zoomer with the Sliders. The app was downloaded by over 1000 users. Over 40 hours of interaction time was logged, and 238 completed questionnaires were obtained.

The users were simply instructed to search for sounds they liked, or thought were useful or interesting. They were told to make sure to save favourites as presets. A number of simple sequences could be skipped through so that listeners would not get too tired of one repeating pattern, but the sequences themselves were not editable, so as to restrict interaction to timbre adjustments. The different interfaces were presented individually and in combination for 5 minutes each, in a random order. After the timed session a questionnaire was presented, and on completion further features were unlocked: such as the ability to show and hide the two interfaces, and MIDI connectivity. Users agreed to a consent statement before their interactions were logged.

3. RESULTS
3.1 Questionnaire

Tables 2 and 3 show the questions asked at the end of the timed sessions. Figures 2 and 3 show the results as diverging stacked bar charts [18]. Results where the user had answered every question the same were discarded.

Most respondents were clearly very familiar with electronic music (AD 1). The participants self select, so some bias in favour of novel interfaces can be expected.

Positive responses to this application include the ability to see the presets as points in space, and to see your "undo" path (AD 2 and 5). The question of whether the mapping was understandable seemed inconclusive (AD 3), but most users did get a sense of the locality property (AD 4) (further

SZ 10	The best interface for discovering interesting sounds quickly was...
SZ 11	The best interface for fine tuning a sound was...
SZ 12	Interface that I felt more in control using...
SZ 13	The interface that felt more creative was...
SZ 14	Interface better for generating new ideas...
SZ 15	Interface better for performing live would be...
SZ 16	Overall, the interface I preferred using was...

Table 3: Questions requiring a 5 point sliders vs. Zoomer answer.

work could investigate this further by obtaining similarity judgements). The Zoomer was deemed more useful than a mere randomiser (AD 7). The strongest result of all (albeit to a heavily loaded question!) was that people highly value happy accidents in the creative process (AD 9).

Particularly of interest was the hypothesis that sliders would be preferred for convergent tasks and the Zoomer preferred for divergent. Responses to SZ 10 and 14 (divergent aspects) contrast sharply with SZ 11 and 12 (convergent tasks). There was a large significant difference between the means of these two properties (difference $= -2.6$ $p < 0.01$), confirming this hypothesis. Most participants felt that the Zoomer was the more creative (SZ 13) which may reflect the popular identification of creativity with novelty and divergent thinking, or simply the fact that new experiences with novel technologies can be inspiring in themselves.

3.2 Interaction Logs

3.2.1 Sound Discovery Rates

Our hypothesis was that if more presets were saved in a particular mode, it might indicate that this interface was best for locating good sounds quickly. The total numbers of presets saved in each different session are shown in Table 4. The upper row of values show the totals when the users spent 5 minutes on each interface, the lower shows the number of saves during the subsequent free-use period. For timed sessions, most presets were saved in the zoomer only mode, indicating that this may have been the fastest interface for sound discovery. However these results are not statistically significant (p > 0.05), as the number of saves per user is rather low and highly variable. Greater incentive to find as many sounds as possible may have improved the experiment in this regard. The large number of saves in the combination interface after the experiment reveal that people much preferred the combination, given the choice.

3.2.2 Interface Preference for Divergent and Convergent Traversal

Was divergent or convergent behaviour detectable from the interaction data? One indication of this was the average zoom level at which people scrolled around compared to the average zoom level at which they saved a preset. The hypothesis would be that people zoomed in to hone the sound before saving. The total amount of time users spent scrolling at 7 different zoom scales is shown in figure 4. The zoom levels are the logarithm of the scale factor rounded to the nearest integer. Data from before the zoom functionality was first used were omitted from the summation. Users showed a clear preference for larger scales, despite the unpredictable timbre changes: they spent 200 minutes scrolling at the largest scale (where sliders change by 64 MIDI CC units per grid division), and only 50 minutes at the lowest scale (1 MIDI CC unit per division).

With Zoomer and sliders present, the zoom functionality was used less: this indicated that when the sliders were

Figure 3: Interface preference responses. Darker bars indicate preference for sliders, lighter for the Zoomer. "No preference" is centred. 11 and 12 reveal slider preference for convergent properties, 10 and 14 show Zoomer preference for divergent.

Interface	Sliders	Combination	Zoomer
No. saves (timed)	452	406	488
No. saves (free.)	9	366	51

Table 4: Total number of presets saved for the three interface views, during timed stages and after the completed experiment.

present the zoom functionality was eschewed in favour of convergence using separate parameters. However, the difference between means was not significant (p = 0.1204).

A more significant trend is seen when both interfaces were on screen, by investigating which interface was being used immediately before and after saving a preset. The hypothesis was that users would exhibit a repeating diverge - converge - save approach, therefore the interface used immediately after saving would be the one preferred for diverging, and the one immediately prior would be the one preferred for converging. Table 5 shows the results, indicating that people were about six times more likely to follow a Sliders - Save - Zoomer pattern than the reverse, supporting the hypothesis. It is hard to confirm this by analysing the timbral paths, as they are hugely different for the different interfaces, the average distance-per-event was 60 for the Zoomer and 10 for the sliders, so it could be said that the Zoomer was intrinsically more rapid, random and divergent. Further work will attempt to extract better path features for this analysis. Path properties are easier to analyse for a single interface: Figure 5 shows that, for the Zoomer only sessions, the average speed of scrolling tends to reduce by about a factor of two as the user converges.

Another thing to note is that presets were around five times more likely to be saved during continuous Zoomer use than during Slider use. It could be argued that this merely shows an overall preference for using the Zoomer, but the total interaction time was only 2:1 in favour of the Zoomer, so overall it still seems more prolific.

Before Save	Zoomer	Sliders	Zoomer	Sliders
After save	Zoomer	Sliders	Sliders	Zoomer
Total	708	127	58	333

Table 5: Which interface was used immediately before and after saving favourites. This includes free interaction after the timed sessions.

Figure 4: Histogram of the time users spent at each of seven zoom levels. Vertical lines show the means of zoom levels when presets were saved and when scrolling the grid for both the Zoomer and Combination stages.
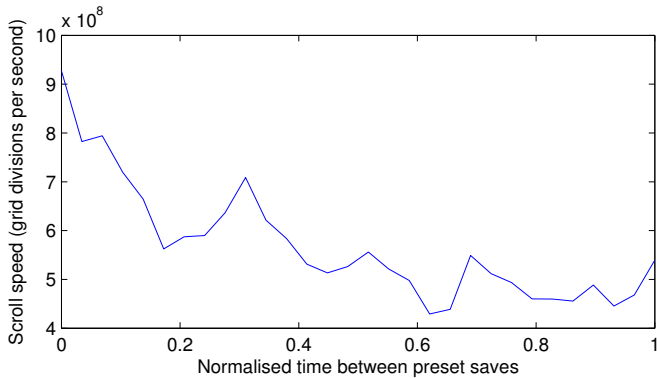

Figure 5: Scrolling speed between preset saves, averaged across all pairs of consecutive saves. Time is normalised such that the time of the previous save = 0 and time of next save = 1.

3.2.3 Analysis of Slider Adjustments

Did slider use show any indication of optimisation-style behaviour? Figure 6 shows the absolute size of consecutive adjustments in the case where 3 direction changes were made before the interaction with this slider was ended. The different distributions seem to indicate different interaction stages. From the second movement onwards, the average sizes of adjustments get progressively smaller. A tentative model to explain this would consist of three basic stages:

1. Initial "effect query": the humped distribution, and fact that there are hardly any small adjustments indicates that this might be just an exploratory enquiry: "what happens when I move this?".

2. "Slow scan": the second movement possesses a fairly uniform distribution, indicating that the desired sound could be anywhere along the length of the slider. This

Figure 6: Histogram of the size of consecutive slider adjustments, in the case of 3 changes of direction (first top, last bottom). Thin vertical line shows the mean.

might indicate a more careful search. This stage usually takes more time.

3. "Honing": there can be several of these stages becoming smaller and more focussed on the eventual preference. The last movement is usually very small (M < 30) indicating that the desired setting has just been found, but has been overshot somewhat.

Progression from stage 2 to stage 3 can be seen for all interactions, up to around 7 direction changes. So, we speculate that slider interaction may reveal a smaller, one dimensional microcosm of divergent and convergent behaviour. Most of the time however, little exploration is necessary with a slider: the most common number of direction changes is one or none.

4. CONCLUSION: NO NEED TO LEAVE SERENDIPITY TO CHANCE

The strongest result of this experiment was that, when both a predictable one-to-one mapping interface is combined with an unpredictable, exploratory interface, clear asymmetry in interface preference is seen before and after the locating of favourites. This asymmetry seems well explained with reference to divergent and convergent search strategies.

Combining this finding with the questionnaire feedback we can claim that:

- Divergent exploration and convergent honing behaviour can be detected in interaction logs.

- Different parameter navigation strategies are suited to different stages of the creative process.

- Users will naturally use the most suitable interface for these strategies, given the choice, therefore the ability to switch between navigation styles is important.

So, even in an uncontrolled experiment such as this, some clues as to musicians' creative processes can be obtained. This raises the intriguing question of whether comparing brain imaging and interaction logs could reveal correlations,

or if artificial creative systems could learn something from the statistical properties of human search strategies.

What is missing from this experiment is some attempt to evaluate the value of the discovered sounds, and if adding a divergent component had a positive effect on the *quality*, as well as the quantity of the discoveries. A social media aspect could be introduced to enable users to rate each others presets (in the manner of Amabile's consensual assessment technique [11]).

Musicians often admit to a role for unpredictability and serendipity in their work. Due to music technology's roots in the recording studio, and the engineer's tendency to think in terms of goal oriented tasks, there has perhaps been a lack of acknowledgement of the more serendipitous aspects of creation when designing interfaces and controller mappings. The happy accidents that do emerge are often seen as uncontrollable by-products, and not something possible to design for. The results of this experiment indicate that, whilst completely removing convergent control of individual parameters would certainly be a bad idea, deliberate design according to the considerations in section 1.3 may unlock divergent traversal strategies and potentially more innovation.

4.1 Links

A video demonstration of Sonic Zoom is provided at `http://y2u.be/485FnfJOuhI`

The app can be obtained free of charge from `http://appstore.com/soniczoom`

5. REFERENCES

[1] R. Bencina. The Metasurface - applying natural neighbour interpolation to two-to-many mapping. In *Proceedings of the 2005 International Conference on New Interfaces for Musical Expression (NIME05), Vancouver, BC, Canada*, 2005.

[2] M. A. Boden. *The creative mind: Myths and mechanisms*. Abacus, 1992.

[3] D. T. Campbell. Blind variation and selective retentions in creative thought as in other knowledge processes. *Psychological review*, 67(6):380, 1960.

[4] M. Duignan, J. Noble, and R. Biddle. Abstraction and activity in computer-mediated music production. *Computer Music Journal*, 34(4):22–33, 2010.

[5] J. S. B. Evans. In two minds: dual-process accounts of reasoning. *Trends in cognitive sciences*, 7(10):454–459, 2003.

[6] L. Goddyn and P. Gvozdjak. Binary gray codes with long bit runs. *The Electronic Journal of Combinatorics*, 10(27):1, 2003.

[7] C. Goudeseune. Interpolated mappings for musical instruments. *Organised Sound*, 7(2):85–96, 2002.

[8] F. Gray. Pulse code communication, Mar. 17 1953. US Patent 2,632,058.

[9] J. P. Guilford. *The nature of human intelligence*. McGraw-Hill (New York), 1967.

[10] C. H. Hamilton and A. Rau-Chaplin. Compact Hilbert indices: Space-filling curves for domains with unequal side lengths. *Information Processing Letters*, 105(5):155–163, 2008.

[11] B. A. Hennessey and T. M. Amabile. Consensual assessment. *Encyclopedia of creativity*, 1:347–359, 1999.

[12] A. Hunt and R. Kirk. Mapping strategies for musical performance. *Trends in Gestural Control of Music*, 21, 2000.

[13] A. Hunt, M. Wanderley, and R. Kirk. Towards a model for instrumental mapping in expert musical interaction. In *Proceedings of the 2000 International Computer Music Conference*, pages 209–212, 2000.

[14] D. Kahneman. *Thinking, fast and slow*. Farrar, Straus and Giroux, 2011.

[15] J. Lehman and K. O. Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223, 2010.

[16] J. Mycroft, J. D. Reiss, and T. Stockman. The influence of graphical user interface design on critical listening skills. In *Sound and Music Computing (SMC), Stockholm*, 2013.

[17] C. Nash and A. Blackwell. Liveness and flow in notation use. In G. Essl, B. Gillespie, M. Gurevich, and S. O'Modhrain, editors, *Proceedings of the International Conference on New Interfaces for Musical Expression (NIME)*, Ann Arbor, Michigan, May 21-23 2012. University of Michigan.

[18] N. B. Robbins and R. M. Heiberger. Plotting Likert and other rating scales. *JSM Proceedings, Section on Survey Research Methods, Alexandria, VA*, pages 1058–1066, 2011.

[19] J. B. Rovan, M. M. Wanderley, S. Dubnov, and P. Depalle. Instrumental gestural mapping strategies as expressivity determinants in computer music performance. In *Proceedings of the Kansei, The Technology of Emotion Workshop,(Genova-Italy)*, 1997.

[20] D. K. Simonton. *Origins of genius: Darwinian perspectives on creativity*. Oxford University Press, 1999.

[21] K. E. Stanovich and R. F. West. Individual differences in reasoning: Implications for the rationality debate? *Behavioral and brain sciences*, 23(5):645–665, 2000.

[22] R. Tubb and S. Dixon. Sonic Zoom: A zoomable mapping of a musical parameter space using Hilbert curves. *Computer music journal*, 38(3):Forthcoming, http://tinyurl.com/lyakhlz, 2014.

[23] D. Van Nort, M. M. Wanderley, and P. Depalle. On the choice of mappings based on geometric properties. In *Proceedings of the 2004 conference on New Interfaces for Musical Expression (NIME)*, pages 87–91. Shizuoka University of Art and Culture, 2004.

[24] J. J. Van Wijk and C. van Overveld. Preset based interaction with high dimensional parameter spaces. *Kluwer International Series in Engineering and Computer Science.*, pages 391–406, 2003.

[25] R. Vertegaal and E. Bonis. ISEE: An intuitive sound editing environment. *Computer Music Journal*, 18:21–29, 1994.

[26] M. M. Wanderley and N. Orio. Evaluation of input devices for musical expression: Borrowing tools from HCI. *Computer Music Journal*, 26(3):62–76, 2002.

[27] T. B. Ward, S. M. Smith, and R. A. Finke. Creative cognition. In R. J. Sternberg, editor, *Handbook of Creativity*. Cambridge University Press, Cambridge, 1999.

[28] G. A. Wiggins. A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems*, 19(7):449–458, 2006.

[29] M. J. Yee-King. *Automatic sound synthesizer programming: techniques and applications*. PhD thesis, University of Sussex, September 2011.