UNFINISHED BUSINESS: SOME REFLECTIONS ON ADOPTING A PRACTICE-BASED APPROACH TO TECHNOLOGICAL RESEARCH AS AN ARTIST

Rodney Berry, Ernest Edmonds, Andrew Johnston
University of Technology Sydney
Creativity and Cognition Studios

ABSTRACT

This paper reflects upon aspects of my experience as an artist moving into research and my attempts to reconcile the two areas of activity and interest. It describes several tabletop augmented reality music systems in the context of my experience as an artist working in technology research environments. It is my intention to show how the relationship between creative practice and technological research has changed for me over time and how I have come to embrace a practice-based approach to research where creative practice takes a central and crucial role.

1. INTRODUCTION AND BACKGROUND

As first author, I write in the first person from my own experience and practice. However, the co-authors acknowledged above have provided discussion and mentorship that affect the form and content of this writing.

In 1998, I had an exhibition in Paris as part of the Virtual Worlds conference. The Director of the Media Integration and Communications Lab at Advanced Telecommunications Research (ATR) in Kyoto visited the show and talked with me at length about the work. He later invited me to come to ATR as a visiting researcher. It was there that I became involved with augmented reality interfaces that had both tangible and visual elements. The first such project was *The Augmented Groove* (Poupyrev et al. 2001), a demo developed in collaboration with Ivan Poupyrev and others. We presented it at Siggraph 2000 Emerging Technologies exhibition.

1.1. The Augmented Groove

This work consisted of a stack of old vinyl phonograph records, each printed with markers that could be identified and tracked by a computer. Each record, when visible to an overhead camera, caused a particular techno music track to be played. The position and rotations of the record controlled various live effect parameters such as filter sweeps, distortion, echo, reverb etc. When more than one record was visible to the system, the looping tracks could be mixed and effected live by several people at once.

The Augmented Groove explored representation and control where the player held something tangible and each axis of translation or rotation clearly controlled a separate effect parameter. For example, rotating the record clockwise and anti-clockwise changed the sound level of its particular track whereas lifting it high above the table would add distortion or other comparatively extreme effects. The former parameter change is persistent because a record can be left on the tabletop in its rotated state. However, the latter vertical movement requires physical effort to sustain by lifting of the record, so its corresponding parameter must eventually return to its state when left on the tabletop. It made sense to assign a more extreme effect parameter, such as distortion, that was interesting in moderation but not sustainable over time. In this way, the logic of the body and the physical space, including the tabletop, became part of the program by providing distinct affordances and constraints that mirrored their role in the actual music being made.

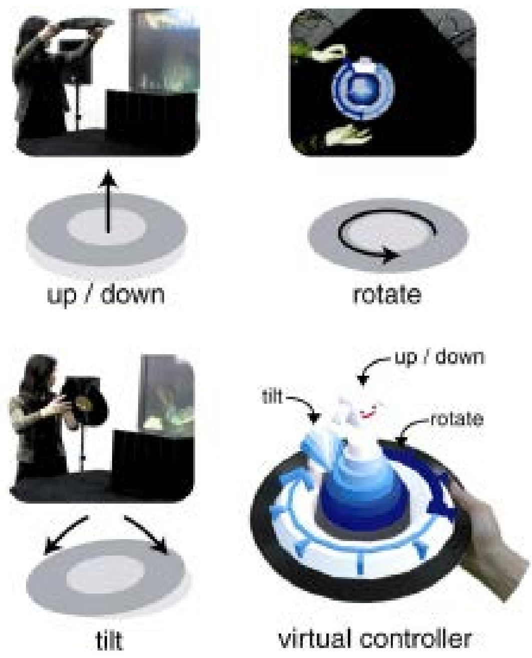

Figure 1 The Augmented Groove (2000)

As Paul Dourish points out (Dourish 2001), a tangible interface is rarely purely tangible. The very fact that tangible elements are physically present in the real world makes them also available to our other senses, especially the visual. In addition, fiducial tracking, via the Augmented Reality Toolkit (Kato and Billinghurst 1999), added extra visual elements to the interface. The view from the overhead camera was displayed on a screen at the rear of the table-top interacting space and animated 3d CG objects and characters appeared to be attached to the records via their visual markers. This a standard way of working with camera-based augmented reality, in this case the Augmented Reality Toolkit, to create a mixed field of view containing both real and virtual objects.

The same augmented reality technique, along with overhead camera, table-top interactions and large screen placed behind the table, was used in The Music Table project (Berry et al. 2002).

Figure 2. The Music Table (2002)

1.2. The Music Table

The user or player of The Music Table creates melodic patterns by arranging blocks on the surface of a table. The overhead camera identifies and tracks the marked blocks and uses their positions to control a MIDI synthesiser. The position of blocks on the vertical axis (toward or away from the player) determines the pitch of the note associated with that block. Its horizontal position (to the left or right of the player) determines its position on a repeating timeline. In effect, the table-top and blocks act as a musical score in the physical world.

Visual augmentations, visible on the screen with each block, add information about the length and loudness of each note. Note lengths are changed by momentarily tilting a block to the left or right in order to increment or decrement the duration of the note. Each note can be made louder or softer by rotating clockwise or counter-clockwise as one might a traditional audio volume control. Because the timeline is relatively short, the system is suited to styles of music that involve repetition, such as minimalist instrumental music and dance music.

The project was shown at SIGGRAPH 2001 Emerging Technologies Exhibition as well as a number of other similar exhibitions attached to technology research and development conferences.

Both Augmented Groove and The Music Table were designed in such a way that they could be learned quickly and enjoyed by a naïve user. This reflects the context of both as primarily technology demonstrations for the SIGGRAPH Emerging Technologies exhibition, where the expected period of engagement is between thirty seconds and three minutes before the visitor moves on to the next thing in the show.

This also reflects the roles often assigned to artists and designers in a technology research environment.

1.3. Roles of Artists

A fortunate artist might be engaged as an 'artist in residence' and allotted budget and an army of underlings to do their bidding. This happened to some extent at ATR and other places where Japan's bubble economy and the technology boom in Silicon Valley contributed to a brief 'golden age' for media art in Japan and elsewhere. Research publications from these projects were often written by the artist's technologist collaborators or perhaps co-written with some input from the artist. More often than not, however, artists become decorators of other people's research. After the lab changed direction and the Art and Technology Project was no more, I and the two Japanese artists (who were employed as graphic designers), were largely tasked with creating compelling demonstrations for events like SIGGRAPH. As I became better at writing research papers and publishing, I (and other researchers) coached the other artists to do likewise. One of the artists is still working at a similar lab, but is now officially a researcher and publishes her own papers.

It is worth underlining the fact that I became a researcher through a chain of unplanned events and found myself in a community of practitioners who had already completed PhDs (or were in the process of doing so) and were socialised into the culture of research. It took some time for me to learn what was meant by such words as 'evaluation', or even 'research'. It took even longer to understand that the meanings were local to that community of practice and that I needed to learn the methods and metrics recognised by this particular community enough to be able to run my own experiments, analyse the results and publish the work.

1.4. Paradigms

We were conducting quantitative research according to a particular paradigm that felt at odds with the way I naturally work creatively with technological resources. Also, I felt that researchers around me were telling stories about their research that fit a particular narrative curve. The curve starts with a clearly defined problem; a review of literature; the formation of hypotheses from the literature and the testing of these hypotheses followed by evaluation. In reality, I observed them working in a much less rational way but they later altered the story to fit the classical scientific narrative. They were not being dishonest, just inclined to forget the parts of the process that do not fit in the narrative. I suspect that a large chunk of science goes undocumented because of this.

Thomas Kuhn (Kuhn 1970, pp. 43 - 51) stresses the importance of underlying paradigms and how they serve to unify particular research communities in the sense that such communities identify with a particular paradigm. Different research communities have different levels of agreement about what is interesting and how it should be investigated. In some, there is almost universal agreement and there is little discussion of research paradigms since the paradigm is largely implicit through one's membership of that community. For this reason, they have little need to examine the underlying paradigm or its philosophical basis. Vaishnavi and Kuechler (Vaishnavi and Kuechler 2008, p.8) describe such communities as tightly paradigmatic communities. Research into tangible interfaces and their use of representations draws on literature, theory and practice from a variety of communities that, as a research community, would best be termed "pre-paradigmatic" or "multi-paradigmatic" since, although its members might identify as a community, they are not necessarily bound together by a particular phenomenon of interest and accepted methods. This meant that, although the community in which I was located held to a particular paradigm, my field of study as a whole represents a number of communities loosely combined enough to have porous boundaries and to be open to new approaches.

1.5. Practice-based Research

As I became involved in the work of the Creativity and Cognition Studios, I began to learn about approaches to research that place creative practice firmly at the centre of the research activity. I was motivated by my previous experience to revisit aspects of my earlier work and examine it through a fresh set of cultural lenses. Eventually I decided to return to study in CCS's doctoral programme. In effect, I feel there is unfinished business in the earlier work that I believe can be resolved by adopting a different research approach.

Edmonds and Candy (Edmonds and Candy 2010) present a model (see Figure 3) for practice-based research consisting of three domains of activity: Practice, Theory and Evaluation.

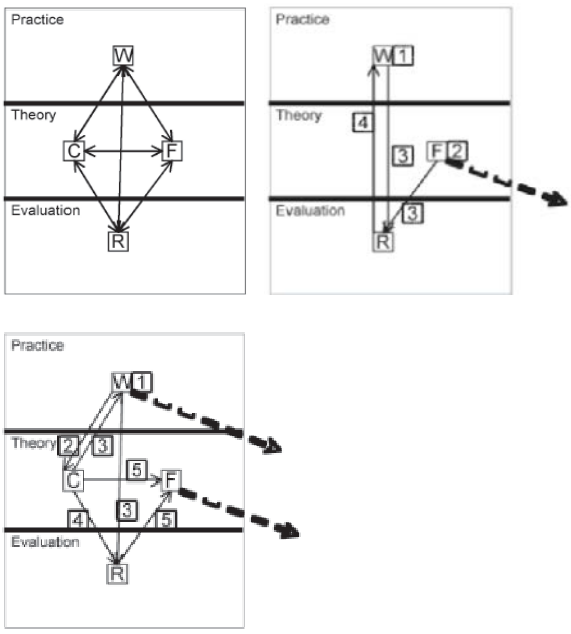

Figure 3. Three different trajectories for PBR

Practice results in Works (artefacts, installations, exhibitions, performances, etc.); Theory gives us Criteria (**C**) (design rules, guidelines etc.) and Frameworks (**F**); and Evaluation (**E**) produces Results (**R**). All are influenced by one another and the researcher might start at any point and proceed to any others, depending on their focus. In practical terms, Edmonds and Candy identify several typical trajectories based on their observations of Doctoral students at CCS. For example, one might start with theory via practice and evaluation to test and modify the theory. Or, one might start with practice and move through evaluation and theory-making in a number of iterated loops.

My chosen approach is closest to the second trajectory and my goal is to produce theory and design criteria out of an iterative process of prototyping and testing of tangible music systems with visual augmentations.

My starting point is a tangible system, very much like The Music Table, that uses the arrangement of blocks on a table-top to create *bytebeat*, a relatively recent computer music practice that traditionally relies on the computer keyboard. The focus is on capturing and reflecting upon very early stages of an iterative prototyping process whose ultimate goal is not a 'usable' commercial system. The system is developed purely to uncover design criteria and middle range theory (Merton 1968) for similar systems in the future. By middle range theory, I mean substantive theory that applies specifically to the situation at hand and is less concerned with broad applicability. Put another way, I am prepared to sacrifice a certain amount of traditional scientific validity in order to gain specific relevance to the subject of my study.

At this stage of my process, I will iteratively prototype, reflect on what is done and found, and analyse these reflections in much the same way as I will later in the user studies with more mature versions of the prototypes.

The product of the research will be presented as an art work in a variety of contexts. One context will be that of an installed work in an art gallery. Another form of the work is as a live performance instrument where the user forms and modifies bytebeat strings on-the-fly.

2. BYTEBEAT

Bytebeat is a genre of electronic music, made by forming mathematical expressions where the symbol, ***t*** represents an endless looping count of integers from **0-255**. The solution for each new value of ***t*** is fed to the computer's audio output via whatever housekeeping code is required by the programming language and its host system. For example: ***t*** on its own will produce a rising sawtooth wave, whereas ***t * 2*** will produce the same wave an octave higher (because the frequency is doubled). Traditionally, it is integer-based and uses a C-like infix syntax for the expressions. However, many also use a postfix or Reverse Polish Notation syntax and there are variations that use floating point arithmetic instead (*floatbeat*).

Given the immense space of possibilities, the resulting sounds are often surprisingly musical, if a little nasty. I speculate that this is because integers are more likely to yield frequency relationships in simple whole-number ratios, corresponding to the natural harmonic series. Simple ratios will also make rhythmic events more danceable – although the number of legs the dancer needs may vary.

Bytebeat first appeared in 2011 with a YouTube video, subsequent blog posts and short paper by Ville-Matias Heikkilä (Heikkilä 2011), otherwise known as Viznut.

The idea quickly spread and there was a flurry of activity to discover and share new expressions. Programs appeared for just about every platform, including browser-based ones that made bytebeat even more accessible.

Aesthetically, bytebeat practice tends to go in two directions. The initial push at the time of its discovery was, and continues to be, the challenge of making the most interesting musical result from the shortest string of symbols. This aesthetic is shared by the, by then well-established, demo scene.

The second aesthetic direction for bytebeat quickly emerged as programmers began to make bytebeat programs evaluate the strings in real time without waiting for compilation of the code. Changes made on the command line had immediate musical (or at least sonic) consequences. If the first direction is concerned with composition, the second direction is more about performance and improvisation. This led to bytebeat's absorption of many aspects of the live coding movement. Proficiency (even virtuosity) at the computer keyboard is part of the live coding movement's aesthetic and certainly applies to live bytebeat performance. Copy, Cut and Paste, as well as repeated Undo operations with the keyboard, allow for quick changes and recall of earlier states.

The live coders' issue of what to show on the big screen drew practitioners' attention to how the interface looked to the audience. It was also quickly discovered that the process could produce visual patterns as well as sonic and has even been used to drive a Tesla coil (Montfort, Stayton, and Fedorova 2014). Because bytebeat uses 8bit sound, it made sense for the visuals to be 8 bit as well, or at least evocative of early video games and an overall glitch aesthetic prevalent at the time.

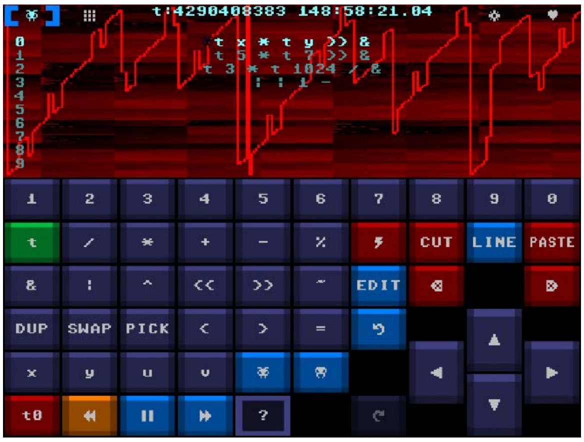

Figure 4. Screenshot of iOS bytebeat application: *The Glitch Machine* by Madgarden

Bytebeat programs, at the time of this writing, can be found for most platforms. They typically offer a keyboard interface or, for phones and tablets, a set of touch buttons for entering strings. The look and feel is often that of early video games with coarse, pixelated graphics and a limited palette of intensely saturated colours. In addition, most have some kind of representation of the waveform and/or spectrum of the output audio. The overall effect is one of sensory overload and a riot of noise and colour.

An important part of the bytebeat experience is that, although simple multiplications and divisions of ***t*** produce fairly predictable results, it is hard, even for an experienced practitioner to predict the results of an unfamiliar string. Having only existed as a genre since 2011, there may still be room for newcomers to contribute and perhaps discover something as yet overlooked. The practice is more akin to circuit-bending where one does not necessarily need to know what one is doing to discover interesting sonic possibilities.

3. TANGIBLE BYTEBEATS

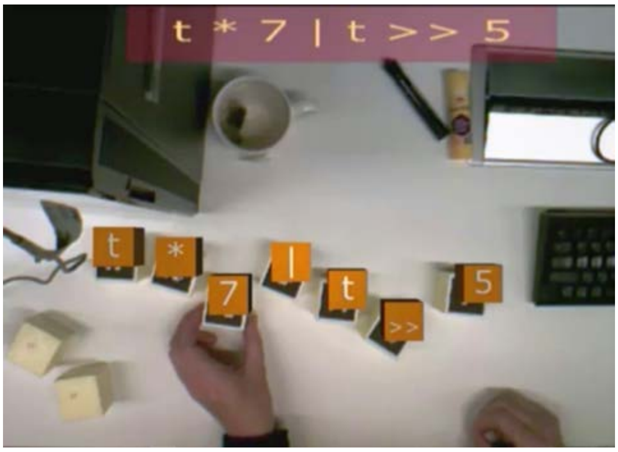

Figure 5. Tangible Bytebeats screenshot

Tangible Bytebeats consists of a tabletop and a collection of 45mm cubic foam blocks, each marked with fiducial markers. An overhead camera is set up so it can see the working area and is connected to a computer. A monitor sits on the tabletop so that the markers and augmentations are visible together on the screen, along with the user's hands when manipulating the blocks. The physical set-up is similar to the author's Music Table project. The software is developed using C++ and TouchDesigner.

TouchDesigner is a procedural flow-chart-style development environment for real-time interactive 3D graphics. It has a collection of operators, represented by rectangular tiles on the screen and connected by graphical wires.

There are operators for building procedural 3D models and surfaces; lighting, texturing, rendering and compositing; manipulating live 'channels' of real-time data; creating and modifying data tables and text (including Python scripts and GLSL shaders) as well as handling data from other programs and devices. A separate program (Berry et al. 2008) uses ARToolkit to identify and track markers. The tracking data, along with the camera image and lens distortion data, are sent via shared memory to TouchDesigner where the actual modelling, rendering and compositing take place. Programming of all interaction also takes place in TouchDesigner.

The graphical display elements are currently simple cubes (Figure 5) that each correspond to one of the markers.

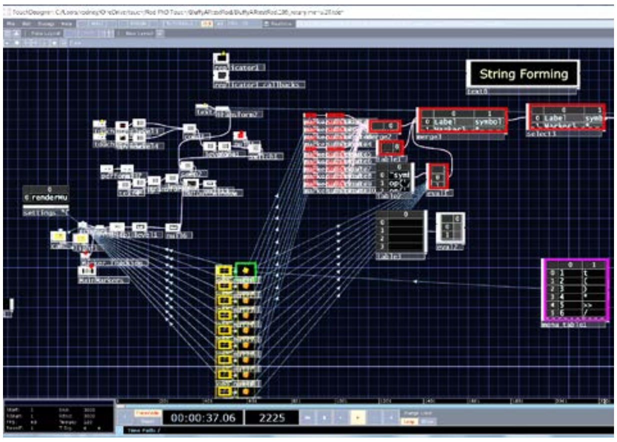

Figure 6. TouchDesigner Screenshot

Although the cubes follow the xyz translations of the markers, they do not rotate with the markers but remain square-on in relation to the camera, directly above the tabletop. Currently, each marker block has two modes of display corresponding to the two modes of function for each block. In Operator Mode, the block shows one of the following bytebeat operators:

t + - * / >> << ( ) % ^ & |

They are selected by rotating the block clockwise or anticlockwise on the tabletop.

In Number Mode, the block displays an integer between 1 and 256, once again, numbers are incremented and decremented by rotating the block in relation to the tabletop. Switching from one mode to the other is achieved by raising the block up toward the camera and back down again.

This allows the user to arrange blocks on the tabletop to form bytebeat expressions such as **t * t >> 20.** The strings are displayed at the top of the screen and also sent to a bytebeat generator, also built within TouchDesigner. A valid expression will generate sound from the computer's audio output. At present, the system does not support the full bytebeat syntax and will require a redesign of the sound generator components. It does at least work well enough to demonstrate the concept.

It is important to point out that the interest in bytebeat began as a divergence from the design for something more like the original Music Table. The idea of using the blocks to create and edit bytebeat expressions came in a perverse intuitive flash. Aesthetically, there is a tension between a wholly keyboard-based practice that is inherently nerdy and arcane, and the addition of a very kindergarten user interface to replace the keyboard. What happens to this tension as the player explores and gets used to the system is one of the stories enacted in the process of investigation.

This is also informed and inspired by Mitch Resnick's *Lifelong Kindergarten* work at MIT Media lab (Resnick 1998) and a belief that kindergarten-style *thinking through your fingers* has a place in the adult world. I think that this is something musicians do when using physical musical instruments and that instruments often used by composers, not just as controllers but as tangible representations.

4. NEXT STEPS

Now the prototype has been through a few iterations, I am documenting my in-process reflections and charting out the emerging design criteria from this process.

An important step will be to bring in other composers and sound artists to play with the system. This research will use qualitative data collection and analysis. Rather than a large sample of naïve subjects, I will work with a smaller number of highly expert and articulate participants. The chief data collection method will be the recorded video and audio protocols of participants' reflections while composing with the prototypes and from semi-structured interviews with the participants. This approach is largely informed by Ericsson and Simon (Ericsson and Simon 1992) as well as Schön (Schön 1983).

Analysis of data follows a Grounded Theory approach based on the Constructivist interpretation of the method described by Charmaz (Charmaz 2006) and others. Grounded Theory Methods are characterized by the following:

- The goal is to build middle-range theory grounded in the data itself rather than from the literature.
- Analysis begins as soon as the first data is collected. Subsequent collection is iteratively informed by the emerging analysis.
- The literature review proper happens after data gathering and analysis has already begun. The idea is to give the data a chance to speak and for the researcher to keep a fresh perspective on what is observed without forcing one's analysis with predetermined theoretical frameworks. In a sense, the literature is treated more as data and less as a sacred source of wisdom.

I am using this particular approach because the study is exploratory in nature and does not readily lend itself to a classical quantitative approach. There is also a good fit between the nature of my own art practice and the prototyping process as a way of discovering/constructing and developing ideas. There is also fit between art practice and the grounded theory approach of inductive, constant comparison and movement toward higher levels of abstraction in the analysis of data. The on-going dialogue with one's data is, for me at least, analogous to the artist's ongoing conversation with his or her materials.

5. CONCLUSIONS

Over the years, working in creative practice and technology research has given me a kind of split personality that I am only now beginning to re-unify through a practice-based approach to research. The purpose of this writing has been to describe several similar creative projects and how my approach has changed toward the research process from which they emerged. For me, the important learning is that it is possible to do research without putting creative practice aside and that creative practice has a place at the centre of the research process.

6. REFERENCES

Berry, R., I. Poupyrev, M. Tadenuma, and N. Tetsutani. 2002. "Inside the score: music and augmented reality." *International Workshop on Entertainment Computing Technologies and Applications*.

Berry, Rodney, Janaka Prasad, Jörg Unterberg, Wei Liu, and Adrian David Cheok. 2008. "Augmented Reality for Non-Programmers." ACM SIGGRAPH, Los Angeles, August 15.

Charmaz, Kathy. 2006. *Constructing grounded theory : a practical guide through qualitative analysis*. London: SAGE.

Dourish, Paul. 2001. *Where the action is : the foundations of embodied interaction*. Cambridge, Mass.: MIT Press.

Edmonds, Ernest, and Linda Candy. 2010. "Relating theory, practice and evaluation in practitioner research." *Leonardo* 43 (5):470-476.

Ericsson, K. Anders, and Herbert Alexander Simon. 1992. *Protocol analysis : verbal reports as data (Revised edition)*. Rev. ed. Cambridge, Mass: MIT Press.

Heikkilä, Ville-Matias. 2011. "Discovering novel computer music techniques by exploring the space of short computer programs." *arXiv preprint arXiv:1112.1368*.

Kato, H., and M.; Billinghurst. 1999. "Marker tracking and HMD calibration for a video-based augmented reality conferencing system." 2nd IEEE and ACM International Workshop on Augmented Reality, 1999. (IWAR '99), San Francisco, CA.

Kuhn, Thomas S. 1970. *The structure of scientific revolutions*. 2nd ed. Chicago: : University of Chicago Press.

Merton, Robert King. 1968. *Social theory and social structure*. 1968 enl. ed. New York,: Free Press.

Montfort, Nick, Erik Stayton, and Natalia Fedorova. 2014. "The Trope Tank: A Laboratory with Material Resources for Creative Computing." *Texto Digital* 10 (2):53-74.

Poupyrev, I., R. Berry, M. Billinghurst, H. Kato, K. Nakao, L. Baldwin, and J. Kurumisawa. 2001. "Augmented Reality Interface for Electronic Music Performance." proceedings of the 9th International Conference on Human-Computer Interaction (HCI International 2001).

Resnick, M. 1998. "Technologies for Lifelong Kindergarten." *Educational Technology Research and Development* 46 (4).

Schön, Donald A. 1983. *The reflective practitioner : how professionals think in action*. New York: Basic Books.

Vaishnavi, Vijay, and William Kuechler. 2008. *Design science research methods and patterns : innovating information and communication technology*. Boca Raton: Auerbach Publications.