Proceedings, The Thirteenth AAAI Conference on
Artificial Intelligence and Interactive Digital Entertainment (AIIDE-17)

Generative Design for Textiles: Opportunities and Challenges for Entertainment AI

Gillian Smith
Worcester Polytechnic Institute
Worcester, MA, USA
gmsmith@wpi.edu

Abstract

This paper reports on two generative systems that work in the domain of textiles: the Hoopla system that generates patterns for embroidery samplers, and the Foundry system that creates foundation paper piecing patterns for quilts. Generated patterns are enacted and interpreted by the human who stitches the final product, following a long and laborious, yet entertaining and leisurely, process of stitching and sewing. The blending of digital and physical spaces, the tension between machine and human authorship, and the juxtaposition of stereotypically masculine computing with highly feminine textile crafts, leads to the opportunity for new kinds of tools, experiences, and artworks. This paper argues for the values of textiles as a domain for generative methods research, and discusses generalizable research problems that are highlighted through operating in this new domain.

Introduction

Computational creativity and procedural generation typically involve producing digital artifacts, such as digital music (Cope 1996), digital art (Colton 2012), or game levels (Togelius et al. 2011). Systems that produce physical artifacts, such as vases (Horn et al. 2015) and paintings (Malina, McCorduck, and Cohen 1991) still often have an automated production phase, where the creation of the designed artifact is offloaded to another machine (a painting robot, or a 3D printer). Textile pattern generation sits in an interesting middle space: the design of the pattern itself is fully automated, while the actual creation of the artifact is performed by a human.

This is a meta-level design problem that involves procedurally creating a pattern that will allow a human to experience and make decisions on the creation of the final form. There are analogs in this new area of work to existing research in computational creativity. Textile pattern generation brings up questions of shared authorship, similar to mixed-initiative systems (Smith, Whitehead, and Mateas 2011; Liapis, Yannakakis, and Togelius 2013), but the human is responsible for final decisions and interpretation of the pattern into the final product. Producing textiles is a long and labor-intensive process, but also an enjoyable one performed by hobbyists and professional artists alike; thus, there are also parallels to experience-driven procedural generation (Yannakakis and Togelius 2011), where the focus is on creating interesting and valuable play processes. Finally, textile pattern generation must also be heavily product-driven: the aesthetics of the final piece are important and must be formally represented and evaluated, whether by machine or human curator. Textile pattern generation is thus an interesting, new domain for computational creativity that sits at the intersection of three major concerns of the field: authorship, process, and product.

As part of an exploration into textiles as a domain for generative design research, this paper reports on two experimental generative systems each for a different craft. Hoopla generates cross-stitch embroidery samplers using quotes and color palettes from web-based sources, and both the software and the finished pieces are intended as art pieces. Foundry generates quilt patterns (specifically, foundation paper piecing templates) using a parameterized, constraint-based generator, and is designed as a tool for exploring new quilt patterns. Resulting pieces from both of these systems have been shown in international venues. In this paper, I present the design of each system and lessons learned from their creation. The paper closes with an argument for textile arts as an interesting domain for entertainment AI, and for how both the open problems it reveals and values inherent in the craft can be illuminating for existing work in entertainment-focused AI.

Related Work

Hoopla and Foundry are by no means the first systems to generate textiles or textile patterns. Fiber artist Libs Elliott creates quilts that are designed by a program written in Processing using Joshua Davis's HYPE framework (Elliott 2017). However, this code only creates the form of the quilt without a formal representation for the pattern; she is left to reverse engineer the patterns herself before creating the quilt.

There is also commercial appeal to textiles that incorporate procedural design. Artist Fabienne Serriere generates algorithmic scarves, knit with patterns that are algorithmi-

Figure 1: A collection of stitched pieces created from patterns generated by Hoopla.

cally generated [1]. Similarly, Cody Lewis's string&&loop project generates knit socks with unique, procedural patterns [2]. Commercial company WOVNS produces woven textiles on demand, and includes several designs that are procedurally generated [3].

Hoopla

Hoopla is named for the traditional method in embroidery of finishing a piece in the hoop it is stitched in. Hooped embroidery is purely decorative, though the same embroidery techniques can be used with different finishing methods to embellish functional pieces such as clothing and housewares. Hoopla generates embroidery samplers, which are often used by novice stitchers to learn and practice new kinds of stitches and hone their skills. They often include features such as an alphabet (for practicing stitched fonts), decorative motifs, quotes, and borders.

The specific form of embroidery used in Hoopla is cross-stitch. Cross-stitch is a particular type of stitch in hand-embroidery characterized by two overlapping stitches that form an X shape. Cross-stitch samplers are often stitched using fine strands of embroidery floss on an even-weave fabric called Aida cloth, that has a specified number of pre-made holes per inch (the "count"). This helps guarantee that all stitches will be the same size and well-aligned. Children and novices may begin on plastic canvas, which is stitched using a thicker yarn and has large holes for blunt needles to pass through.

Hoopla generates sampler patterns that are then human-stitched (Figure 1). Completed samplers were generated and stitched for Science Gallery Dublin for their juried *Humans Need Not Apply* show in early 2017. These samplers were made on 22-count Aida cloth; exhibition visitors were able to generate patterns for 11-count plastic canvas (Figure 2).

[1] http://knityak.com/

[2] https://www.kickstarter.com/projects/860869947/stringandandloop-knitting-craft-with-code

[3] http://www.wovns.com/

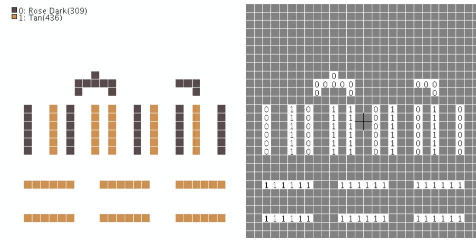

Figure 2: A pattern generated by a visitor to Hoopla at Science Gallery Dublin.

Spatial Representation

Hoopla is an art piece; thus, the goal is not to perfectly replicate existing cross-stitch samplers. Rather, the goal in Hoopla's design is to produce artifacts that are recognizably a part of the genre, while carrying their own aesthetic. The spatial representation in Hoopla is, therefore, loosely based on an informal review of existing, human-created samplers. Such samplers would often contain quotes surrounded by geometric motifs, alphabet sequences, or patterns to fill a space (a centuries-old example of such a sampler is in Figure 3).

Hoopla represents cross stitch samplers as a set of adjacent regions in a subdivided space. The space is recursively partitioned. Each partition can be either horizontally or vertically partitioned, and each parent partition is allowed at most three child partitions (for aesthetic reasons). The probability of a horizontal partitioning is 75%, weighted higher to produce more regions that might neatly fit text, and in keeping with the aesthetic goals of the project.

After subdivision, regions are tagged with a type:

- **Quote.** There is at most one quote region per sampler, and to be selected as a quote region the space must be at least large enough to contain 100 characters of stitched text. This length was selected to ensure that there are quotes that can potentially fill the space.
- **Text.** There are at most three text regions per sampler, containing randomly generated sequences of text: either segments of the alphabet, a numerical sequence, or randomly selected punctuation.
- **Motifs.** All remaining segments are tagged as motifs, which are then generated using the process described in the following section.

Generation Strategy

Hoopla uses a crowdsourced, data-driven approach for palettes and quotes, and an ad hoc constructive method for motifs and motif repetitions. The data-driven approach is intentionally simplistic, linking text and palettes via randomly selected keywords, as a commentary on how much meaning is missing from datasets.

Figure 3: A 1793 embroidery sampler by Maria Lalor, held in the collections of the Metropolitan Museum of Art.

Figure 4: Two sample color palettes that come from a palette search for the keyword "rose".

Quotes The first region on the sampler to be filled is the quote region. A random quote with a length that will fit into the region is selected from an online collection of old proverbs [4]. This quote is typeset to be centered in the region, using an arbitrarily selected font from the computer that is down-sampled to create a pixelated aesthetic for the font face. The contents of the quote then drives the remainder of the sampler generation process.

Palettes The words contained in the selected quote are saved and one of them is randomly selected to query a color API that features crowdsourced colors and palettes that are semantically tagged [5]. The goal in this stage is for the colors chosen for the sampler to be explainable for the main message it communicates. For example, the proverb "The fairest rose at last is withered." may produce palettes with rose colors, if the keyword 'rose' is selected at random (see Figure 4). Once selected, each color in the palette is converted to the LAB color space and the closest matching DMC thread color is selected [6]

Motifs Motifs are generated as collections of arbitrary-length lines (horizontal, vertical, and diagonal) that are then be tiled across the space. Lines were selected as the motif primitive due to the common technique in embroidery to half-stitch lines in one direction, and then complete the cross-stitch by half-stitching in the other direction. When the motifs are tiled, they can be mirrored or rotated, and have an arbitrary size border between the motif repetitions. Many of the motifs are reminiscent of the output of the famous Commodore 64 "10PRINT" program (Montfort et al. 2012)[7].

Creating Collections

Hoopla creates not just individual embroidery samplers, but collections that are intended to be displayed together. These collections aim to have internal consistency via repeated motifs, words, and colors, yet sufficient variety. Each sampler is linked to another via a common keyword. For each quote and palette added to the system, keywords from them are added to a list that is then selected from when creating a new sampler pattern. Similarly, motifs will occasionally be copied from one sampler to another. For example, in the example shown in Figure 1, the word "self" is repeated between the largest and smaller of the pieces.

User Participation

As part of Hoopla's exhibition at Science Gallery Dublin, visitors to the gallery were given the ability to generate their own samplers and stitch them. There was no guarantee that users would have any experience with embroidery, so samplers were designed for small, plastic canvas squares (which is easier to hold and stitch). For logistical and budget reasons, a limited palette of DMC flosses was made available to visitors. Because of the size of the canvas, quotes would not fit onto these small squares; however, short keywords from palettes would occasionally fit. Visitors were given a simple interface to Hoopla that supported generating new samplers and selecting one to complete themselves. As of this writing, samplers generated by Science Gallery visitors are available online. [8]

Foundry

Foundry is a parameterized tool for generating quilt patterns based on foundation paper piecing (Figure 5). Quilts are typically three layers of fabric that are stitched together: the top layer is usually sewn together out of many other fabrics in decorative patterns, the bottom layer can also be pieced but is commonly just one large piece of fabric, and the middle layer (called "batting") is an insulating piece of fabric that is selected to provide the appropriate weight, softness, and warmth. Foundry's purpose is designing the top, pieced layer of the quilt.

Foundation paper piecing is a traditional technique for piecing fabric together into blocks of fabric that, when put together, form an entire quilt top. Using this technique, the quilter sews fabric onto a paper "foundation" in a specific order, folding and trimming fabric as they go. Though the technique has limitations and can be time consuming, it offers a great deal of assistance with geometric precision, especially for complex designs or shapes that don't have right angles (which are easier to cut by hand).

[4] https://proverbs-api.herokuapp.com/api/proverb

[5] http://www.colr.org

[6] DMC is a major manufacturer of embroidery floss, and has a palette of 471 colors. An online RGB-to-DMC color table is available online: https://www.csh.rit.edu/ vance/pages/color.html

[7] This was not intentional during design, but deliberately kept when the similarity to "10PRINT" mazes was noticed.

[8] https://hoopladublin.tumblr.com/

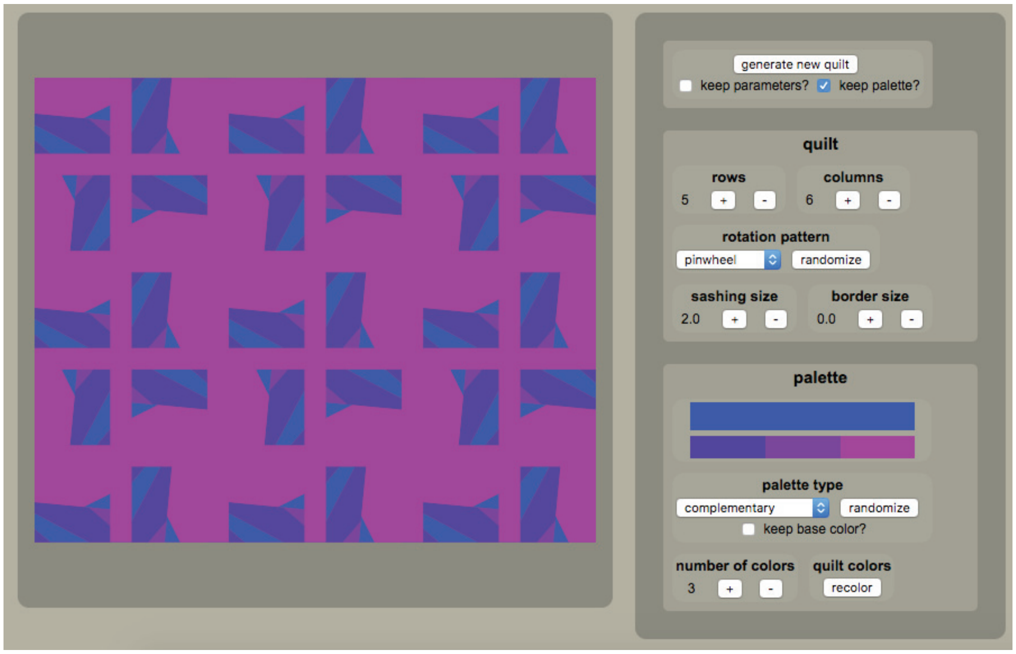

Figure 5: Interface to the Foundry quilt generation tool.

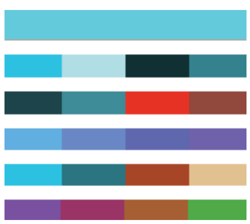

Figure 6: Generated color palettes in Foundry, each with 5 colors and sharing a common base color. From top to bottom: the base color, monochromatic, complementary, analogous, split complementary, and rainbow.

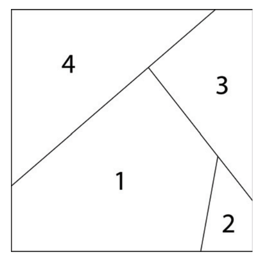

Figure 7: A foundation paper piece template. The first seam sewn is between regions 1 and 2, then region 3 is added, and finally region 4.

Representation

Foundry represents quilts as a series of equally sized blocks that are each paper pieced and that share a common color palette. These blocks can be tiled across the quilt with varied spacing (referred to as "sashing" in quilting) in between them. Blocks can be rotated in different combinations to produce emergent patterns in the overall quilt.

Generation Strategy

Foundry generates color palettes using a constraint-based system. The user can select the type and size of the color palette, and the system will generate a base color and sequence of colors based on the base color, by constraining the hue, saturation, and value of each color in the palette with respect to the base color (Figure 6).

The palette size can range from 2 to 8 colors. Using a constraint-based system for generating color palettes makes it easy to update palettes when varying their size, as well as generate legal palettes that exclude certain ranges of colors that come across as muddy or dusty (such as colors with low saturation and medium value).

Blocks are generated by performing a binary space partition, with seam start and end points biased towards being placed in the middle 70% of the parent seam line in order to avoid bulky corners with many overlapping seams (Figure 7). Each region is assigned a color from the generated palette at random, and can be randomly re-colored at any time during design.

Parameterized Design

As opposed to Hoopla, which is intended primarily as an art piece, Foundry was created as a tool for designers to explore new quilt ideas. Thus, Foundry is heavily parameterized:

- *Palette size.* A palette can be between 2 colors and 8 colors in size.
- *Palette type.* A palette can be analogous, monochromatic, complementary, split complementary, or rainbow.

Figure 8: “Collaboration”, a quilt designed with Foundry.

- *Dimensions.* A quilt can be between 1 and 8 blocks wide, and between 1 and 8 blocks tall.
- *Sashing and Borders.* Borders refer to long pieces of solid fabric around the outside edge of the quilt. Sashing refers to the spacing in between each block. The user can vary the size of borders and sashing independently from each other.
- *Rotation Pattern.* Blocks can be aligned identically, or set in one of the following patterns: sequential, where each subsequent block in a row is rotated by 90 degrees from the previous one; x-symmetry or y-symmetry, where the quilt mirrors blocks on its central x and y axis; or pinwheel, where 2x2 sets of blocks are organized in a rotational pinwheel pattern.

In addition to these parameters, users have the ability to manipulate a palette independent from the block design, generate new palettes from the same base color, or randomize the layout of blocks within the same rotation style.

Exhibitions and Shows

Works produced from patterns generated by Hoopla and Foundry have been evaluated through inclusion in international, juried shows. In the case of Hoopla, the embroidered work and generator was accepted into an exhibition of AI-made art, where the jury and exhibition staff were well aware of the procedural nature of the work, and stated in personal communication that it was being accepted on the strength of both its visual aesthetics and the blending of machine and human authorship in a tangible piece. In the case of Foundry, the “Collaboration” quilt was accepted to a juried show organized by the Modern Quilt Guild at QuiltCon 2015 (Figure 8); while the artist statement made clear that the quilt was designed by a generator, it is not clear to what extent this influenced the decision to include the quilt in the show, as no written or verbal feedback is provided to artists with regard to the jurying system. “Collaboration” has also been shown at the Maynard ArtSpace in Maynard, MA, USA, as part of a curated show featuring work from local artists. The curators had no knowledge that the quilt was machine-generated.

Technical Challenges and Opportunities

Hoopla and Foundry are initial explorations into generative design for textiles. Though the systems provide limited technical innovation, reflecting on these systems reveals interesting open questions in generative design that are applicable in a broad set of domains.

Color Palette Generation

Color palette generation is key for any visual media. Foundry and Hoopla use two different approaches, detailed above. Hoopla sacrifices human control over the color palette in favor of gaining the ability to create a relationship (tenuous though it can be) between text elements in the piece and colors. Foundry provides greater human control over the color palette, but sacrifices the ability for the system to build a relationship between color choices and other design elements.

Other computationally creative systems also create color palettes in the course of their generation process. Fully autonomous systems such as The Painting Fool (Colton 2012), though, do not need to be accountable for the decisions made by the system to a human author who will continue create the final product. Other systems that have shared authorship, such as mixed-initiative tools for game generation, could pay more attention to color and visual aspects of design in their process. Currently, such tools typically avoid the question of color altogether (and, indeed, most visual design considerations) by depending upon a pre-made set of art assets and focusing their AI systems on the content’s structure.

There is also a need to create generative models for color that take into account affect, culture, and visual communication. Existing work in color psychology (Ou et al. 2004) and human-computer interaction (Bartram, Patra, and Stone 2017) provide a foundation for such work.

Generating Collections

Most work in procedural content generation is for creating individual levels or games (Togelius et al. 2011; Hendrikx et al. 2013), or for creating level progressions that are appropriate to increasing player skill (Butler et al. 2013). In computational creativity, algorithms may produce multiple artifacts that are recognizably similar, but there is little work in creating small, tightly curated collections where each piece is intended to go with another.

Yet, in all creative fields, there are examples of such collections. Triptychs–collections of three related pieces that often have one focal piece and two supplemental panels–are common in visual arts such as painting and photography. Musicians create albums, where songs are intentionally crafted to fit together and be experienced in a particular order. Game designers create small collections of levels that are thematically related, such as Sonic the Hedgehog’s “zones” (Team 1991). Each individual piece in a collection can stand on its own, but is altered or even strengthened when experienced in proximity to its partner pieces.

Foundry does not currently create collections, though it is intended in future work to create collections of blocks that can be coexist in a single quilt. Hoopla creates collections through attempting to build loose semantic relationships between the pieces via color and text. In general, generating collections of pieces regardless of the domain requires a means for explicitly modeling thematic relationships between those pieces and for the system to reason about the role the piece plays in the larger collection.

Modeling Process

Hoopla and Foundry both focus on the final product being created, but my experience in creating the pieces themselves (rather than the software) brought into sharp relief the importance of modeling *process* as well as product. In procedural content generation, the closest analog to this may be experience-driven PCG (Yannakakis and Togelius 2011), which advocates for creating play experiences that induce particular affective responses. However, experience-driven PCG focuses on personalization of experience in real-time, given the digital and rapid nature of the games used in such projects. Craft projects are slow, unfolding over several days or weeks. There is pleasure to be found in the repetitive, tangible actions required by the craft, and satisfaction in the slow reveal as individual stitches combine, slowly and sometimes unexpectedly, to produce a finished piece.

Long-term processes such as these have greater scope than current experience-driven work. Outside of craft, we see similar kinds of long-term processes in games that are played over long periods such as MMOs, or in distance-based games that were played by mail (old-fashioned as they may now seem). Adapting user models for entertainment AI to long-term behavior in physical spaces is an open research problem revealed by working in textiles.

A Vision for Textiles and Entertainment AI

In addition to the technical challenges raised by Hoopla and Foundry, these two projects also prompt reflection on the values associated with textiles and the potential they create for new work in entertainment AI.

Tradition

Fiber arts such as quilting, embroidery, and weaving have a long–in some cases millenia-old–history and set of traditions. Common patterns have been extensively cataloged and indexed (Brackman 1993), providing a rich set of data and ideas to draw on. These patterns are commonly used and/or intentionally subverted in contemporary work. In addition to design patterns related to the product being created, there are also well-documented, formal processes for the creation of textiles. Techniques for textiles have been honed and adapted to new technologies over time, from new dyeing processes to sewing machines to personal manufacturing for creating sewing templates and custom tools (Behuniak-Long 1994). While improvisational methods exist for all crafts, traditional methods for quilting, embroidery, and weaving are highly structured and often internally repetitive. Textile craft practices can lend themselves well to formalization.

Feminist Scholarship

Textiles have a long history as so-called “women’s work”. As opposed to fine art, where recognition and acceptance is subject to the institutional sexism found in funding bodies, museums, and educational practices (of Women in the Arts 2017), textile arts and crafts enjoy more participation from women and have often been used as a medium for feminist art (Bernick 1994).

Entertainment AI in this domain has exciting possibilities for new forms of expression, juxtaposing the masculinized, digital domain of computation with the feminine domain of hand crafts. There is also the potential to reach new audiences with AI; in addition to being seen as a women’s craft, there is also a long history of people of color working in textiles, such as the African American Quilters of Gee’s Bend (Beardsley 2002).

Community and Storytelling

Textiles have long been used to tell stories. From The Bayeaux Tapestry’s [9] epic and literal retelling of Anglo-Saxon history to a modern day quilt made to commemorate a wedding, textiles often provide a visual representation of story and community. Additionally, there is a strong tradition of craft as a community-oriented practice, such as quilting bees that involve many community members coming together to create a single quilt. Many techniques for these crafts are so formally defined that it can be difficult to tell where one artist’s work ends and another’s begins. This community-based storytelling inherent in fiber arts aligns nicely with the aims of interactive storytelling research, and offers yet another avenue for new experiences that can emerge from the integration of AI and textiles.

Conclusion

This paper has presented two generative systems that work in the domain of textiles: Hoopla, a cross-stitch sampler generator; and Foundry, a paper-pieced quilt pattern generator. It is not my goal in this paper to argue for significant technical innovation in either Foundry or Hoopla. Both use relatively simple, well-established techniques for generating patterns in their domain. Though, it is worth noting that even these simple techniques can produce striking results.

Rather, the main goal with this paper is to use these two initial explorations into generative design for textiles as illustrative examples in discussing the ways in which operating in this domain can highlight underexplored areas of work in the field. Via lessons learned from these explorations into generative textiles, I have identified open technical challenges and discussed how the values associated with textiles can inform new strands of research in entertainment-focused AI. Many of these lessons learned from Hoopla and Foundry can be applied back to more common domains for generative design. Further, this paper introduces the notion of non-digital, process-heavy leisurely activities as a valuable domain for study, that can provide new kinds of interactive entertainment experiences.

Acknowledgments

Hoopla was partially funded by Science Gallery Dublin for their show, *Humans Need Not Apply*.

[9] The Bayeaux Tapestry is actually an embroidered piece, though its name implies it as a woven tapestry.

References

Bartram, L.; Patra, A.; and Stone, M. 2017. Affective Color in Visualization. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, 1364–1374. New York, NY, USA: ACM.

Beardsley, J. 2002. *The Quilts of Gee's Bend*. Tinwood Books.

Behuniak-Long, S. 1994. Preserving the Social Fabric: Quilting in a Technological World. In Torsney, C. B., and Elsley, J., eds., *Quilt Culture: Tracing the Pattern*. University of Missouri Press.

Bernick, S. E. 1994. A quilt is an art object when it stands up like a man. *Quilt Culture: Tracing The Pattern* 134–150.

Brackman, B. 1993. *Encyclopedia of pieced quilt patterns*. American Quilter's Society.

Butler, E.; Smith, A. M.; Liu, Y.-E.; and Popovic, Z. 2013. A mixed-initiative tool for designing level progressions in games. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*, 377–386. ACM.

Colton, S. 2012. The Painting Fool: Stories from Building an Automated Painter. In *Computers and Creativity*. Springer Berlin Heidelberg. 3–38.

Cope, D. 1996. Experiments in musical intelligence.

Elliott, L. 2017. Libs Elliott: Technology + Tradition, Making Code Into Physical Quilts.

Hendrikx, M.; Meijer, S.; Van der Velden, J.; and Iosup, A. 2013. Procedural Content Generation for Games: A Survey. *ACM Transactions on Multimedia Computing, Communications and Applications* 9(1).

Horn, B.; Smith, G.; Masri, R.; and Stone, J. 2015. Visual Information Vases: Towards a Framework for Transmedia Creative Inspiration. In Toivonen, H.; Colton, S.; Cook, M.; and Ventura, D., eds., *Proceedings of the Sixth International Conference on Computational Creativity*, 182–188.

Liapis, A.; Yannakakis, G. N.; and Togelius, J. 2013. Sentient sketchbook: Computer-aided game level authoring. In *Proceedings of ACM Conference on Foundations of Digital Games*.

Malina, R. F.; McCorduck, P.; and Cohen, H. 1991. Aaron's Code: Meta-Art, Artificial Intelligence and the Work of Harold Cohen. *Leonardo* 24(5):628.

Montfort, N.; Baudoin, P.; Bell, J.; Bogost, I.; Douglass, J.; Marino, M. C.; Mateas, M.; Reas, C.; Sample, M.; and Vawter, N. 2012. *10 PRINT CHR $(205.5+ RND (1));: GOTO 10*. Cambridge, MA: MIT Press.

of Women in the Arts, N. M. 2017. Get The Facts.

Ou, L.-C.; Luo, M. R.; Woodcock, A.; and Wright, A. 2004. A study of colour emotion and colour preference. Part I: Colour emotions for single colours. *Color Res. Appl.* 29(3):232–240.

Smith, G.; Whitehead, J.; and Mateas, M. 2011. Tanagra: Reactive Planning and Constraint Solving for Mixed-Initiative Level Design. *IEEE Transactions on Computational Intelligence and AI in Games (TCIAIG), Special Issue on Procedural Content Generation* 3(3).

Team, S. 1991. Sonic the Hedgehog (Genesis).

Togelius, J.; Yannakakis, G.; Stanley, K.; and Browne, C. 2011. Search-Based Procedural Content Generation: A Taxonomy and Survey. *Computational Intelligence and AI in Games, IEEE Transactions on* 3(3):172 –186.

Yannakakis, G., and Togelius, J. 2011. Experience-Driven Procedural Content Generation. *IEEE Transactions on Affective Computing* 2(3):147–161.