Non-Photorealistic Animation and Rendering
David Mould and Pierre Bénard (Editors)

EXPRESSIVE 2015

Drawing Characteristics for Reproducing Traditional Hand-Made Stippling

Domingo Martín[1] Vicente del Sol[1] Celia Romo[1] Tobias Isenberg[2]

[1] Department of Software Engeneering, University of Granada, Spain [2] INRIA, France

Abstract

We contribute an in-depth analysis of the characteristics of traditional stippling and relate these to common practices in NPAR stippling techniques as well as to the abilities and limitations of existing printing and display technology. In our work we focus specifically on the properties of stipple dots and consider the dimensions and attributes of pens and paper types used in artistic practice. With our analysis we work toward an understanding of the requirements for digital stippling, with the ultimate goal to provide tools to artists and illustrators that can replicate the stippling process faithfully in the digital domain. From the results of our study we provide a dataset for use in new example-based stippling techniques, derive a taxonomy of characteristics and conditions for the reproduction of stippling, and define future directions of work.

Categories and Subject Descriptors (according to ACM CCS): Computer Graphics [Computing methodologies]: Rendering—Non-photorealistic rendering

1. Introduction

Hand-made stippling is an artistic depiction style that represents images using dots. Usually, a pen is used to depose black ink on a white or clear paper. Although it seems to be an easy technique to master, it needs not only artistic talent but also a lot of training and much time for each individual image due to the need to place up to millions of dots. This traditional technique has many advantages: it is economic as it only relies on a single color (good for reproduction), it can represent not only tone but also shape and texture, and its dots do not impose an orientation which removes some visual artifacts. While stippling is no longer very commonly used, it can be found in some scientific domains as archeology, biology, entomology, etc., as well as in artistic drawing.

Similar to other traditional techniques of artistic expression or illustrative depiction, stippling has also been reproduced in the NPAR domain [SS02, DI13, KCW*13]. The support of computer processing has opened up new possibilities such as the creation frame-coherent animation using a 3D model as input [MFS03] to name just one example. One goal for most stippling approaches, however, remains the faithful replication of the traditional technique to be applied to new input images, for instance in situations when it is impossible to hire a professional stipple artist for this purpose.

It can be argued that, while many recent advances including example-based techniques have significantly pushed the quality level of computer-supported techniques, there are still numerous limitations of NPR stippling that remain to be addressed. To be able to make significant advances in the field in the future, we have thus embarked on a study of the traditional technique with the ultimate goal of being able to accurately reproduce traditional hand-made stippling. The goal in our work is to establish a clear set of conditions for an approach to reproduce traditional hand-made stippling. For this purpose we concentrate specifically on the generic problem of reproducing hand-placed stipple dots and the corresponding constraints that arise from the used traditional tools as well as the goals of digital presentation and reproduction. The results of our study of these low-level aspects of stipple dots can be used in future NPR techniques for faithful NPR stippling. In summary, our contributions are thus:

- A study about the appearance of hand-placed dots in traditional stippling and their physical characteristics. We examine dots based on the pens (type, nib sizes) and paper used by different artists to understand what dots result from different drawing materials. Based on this work we establish constraints for the shape, size, and color of dots when digitally reproduced.
- A discussion about the reproduction of stipple dots on

typical output devices. When transitioning from traditional hand-made stippling to digital reproductions we face a conversion from a continuous to a digital world. Based on this change we not only discuss the capabilities but also identify several important limitations to take into account for digital stippling.
- We use these capabilities and limitations to establish a set of characteristics and conditions for the reproduction of traditional stippling. This taxonomy allows us to revisit established and published stippling approaches to understand to what degree they are able to achieve the goal of a faithful reproduction.
- The taxonomy also allows us to clearly define future lines of research based on the state of the art. We thus propose and discuss essential goals in future research in simulating traditional hand-made stippling.

2. Previous Work in Digital/NPAR Stippling

Digital stippling is a technique that was first introduced to the NPAR literature in the late 1990s. Initially, researchers focused primarily on where to place stipple dots. Most early techniques were based on the concept of Centroidal Voronoi Diagrams (CVD, also called Lloyd's method [Llo82, MF92]): an initial dot distribution is created (e. g., depending on a condition such as tone), then the dot distribution's Voronoi Diagram is computed and the dots are moved to the centroids, a step which is repeated until the result is satisfying.

Deussen et al. [DHvOS99, DHvOS00] were the first to implement this process using an interactive system that used brushes to locally apply Lloyd's method. Secord [Sec02] generalized this process to use weighted centroidal Voronoi diagrams based on the local tone of the source image. Others such as Hiller et al. [HHD03] and Dalal et al. [DKLS06] then extended the general approach to be able to not only place circular dots but general shapes. Researchers have also investigated methods for obtaining a suitable initial point distribution, either to be used in CVD-based stippling or in its own right. Secord et al. [SHS02], for instance, probabilistically distributed primitives in image space—a method than can also be used for frame-coherent animations. In a related approach, Arroyo et al. [AML10] used a Monte Carlo technique sampling an adaptive probability density function. Related to these kind of stochastic methods is also the RenderBots system by Schlechtweg et al. [SGS05] that uses autonomous agents with random processing to place stipples.

Researchers generally investigated ways to avoid the visual problems in point placement that arise from the original version of Lloyd's method: chain artifacts that stipple illustrators aim to avoid [Hod03]. Balzer et al. [BSD09] presented a capacity-constrained way to create point distributions based on Lloyd's method that possess blue noise characteristics. Kopf et al. [KCODL06] presented a method based on non-repetitive Wang tiles and Poisson disk sampling that also produces point sets with blue noise characteristics and thus both avoids artifacts and facilitates an infinite yet smooth zoom into stipple images. Ascencio-Lopez et al. [ALMPHS10] similarly used Poisson disk sampling but with the goal to produce pleasing distributions at fast speeds. Finally, in an alternative approach, Deussen [Deu09] generalized Lloyd's original CVD method by using an energy-based optimization process, for example to also be able to produce point clusters instead of only evenly distributed point distributions.

While the techniques discussed so far—similar to the artistic example—concentrated on representing 2D still images using stippling, researchers also developed techniques for stipple rendering of 3D shapes. The transition to 3D models as the underlying input media not only opens up new possibilities not possible in traditional stippling (e. g., animation) but also raises the problem of frame-to-frame coherence. To solve it, Meruvia Pastor et al. [MS02, MFS03] used a particle system attached to the 3D object's surface to achieve smooth animations. Lu et al. [LTH*02, LMT*03], similarly, placed and tracked points on the surface and discussed hardware acceleration options. To facilitate the zooming into a model, Meruvia Pastor et al. [MPS04] demonstrated how to use distribution hierarchies with graph-based relaxation. In an approach that aims at simplifying the 3D computation, Yuan et al. [YNZC05] computed point distributions for 3D rendering in geometry-image space, achieving frame-to-frame coherence and benefiting from GPU acceleration. Vanderhaeghe et al. [VBTS07] went one step further and reverted back to computing the point locations in 2D space to optimize their 2D characteristics—yet ensuring that the moving point distribution behave correctly for 2D projections of animated (rigid or soft-body) 3D shapes. Approaches for other 3D surface models such as point-sampled geometry [XC04, ZS04] have also been discussed.

Two special forms of 3D object representation are implicit and volumetric models. For the creation of stippling for implicits, Foster et al. [FJW*05] randomly distributed a set of particles onto the model and then use attractor/repulser forces to move the points back to the implicit surface according to its changes as well as with respect to neighboring points. In contrast, Schmidt et al. [SIJ*07] used a real-time method to extract a low-quality base mesh from the implicit model and then placed surfels [PZvBG00] on the surface, which in turn carry a hierarchy of stipple dots. In a related approach for Hermite RBF implicits, Vital et al. [BMCS*10] densely sampled the implicit surface and then used these seed points to generate stipple dots for the rendering. For volumetric stippling, in contrast, Lu et al. [LME*02, LMT*03] randomly placed stipple points throughout the volume based on the data's features, out of which the stipple points to be shown at render time were selected based on the viewing conditions.

While the point distributions discussed so far were largely based on either random placement, noise qualities, or dedicated distribution processes, researchers recently started to derive the distributions from human input in form of example-based stippling. For example, Barla et al. [BBMT06] synthesized different styles of hatching and stippling based on extracting the drawing primitives and analyzing their neighborhood relationships. Inspired by earlier analyses of stipple aesthetics [MIA*07, MIA*08], Kim et al. [KMI*09] followed a similar goal and used a gray-level co-occurrence matrix (GLCM) to capture the statistics of stipple distributions for later synthesis. Martín et al. [MALI10, MALI11] concentrated on reproducing correct stipple distributions based on resolution considerations and scanned stipple dots that can faithfully reproduce the merging of stipple dots.

This last method uses halftoning, which some other stippling methods also employ. For example, Hausner [Hau05] extended error diffusion for generating point distributions for pointillist halftoning. Mould [Mou07] used a progressive distance calculation based on a graph representation of the source image, a method that specifically emphasizes wanted linear features using stipple chains. Li et al. [LM11] extended this structure-aware stippling approach using an error distribution scheme based on the importance of the different features to be reproduced. A special form of stippling related to such structure-preserving approaches is the reproduction of hedcut images. Here, the stipples are arranged along dedicated lines, related to hatching techniques. Examples of hedcut stippling were presented by Kim et al. [KSL*08], Kim et al. [KWME10], and Son et al. [SLKL11].

Virtually all these techniques concentrate on the placement of stipple dots, either for traditional 2D input or for three-dimensional shapes. The question of how to treat the reproduction of the dots themselves if often not raised, in many cases black circles or pixels are used to represent them. Questions of overlapping/merging dots as well as of whether stippling should be treated as a black-and-white technique are also not frequently discussed, with only few exceptions [KMI*09, MALI10, MALI11]). Even the suitable size of the dots is only rarely discussed. For instance, Deussen et al. [DHvOS99, DHvOS00] vary the dot size based on the image tone, while Secord [Sec02] mentioned size control as future work but implemented tone-based control in his demo tool—albeit only for circular dots. Martín et al. [MALI10, MALI11], finally, studied examples of scanned stipple dots and base their dot placement strategy on the physical size of real stipple dots, computing the respective resolutions accordingly. Based on a study of hand-made stippling, we thus extend this general approach and work toward establishing a new framework for digital stippling that focuses on the characteristics of the stipple dots, their physical sizes, the employed paper, and the used output media.

3. Traditional Hand-Made Stippling

Traditional hand-drawn stippling is produced by manually placing dots on paper with a pen. Typically, black ink and white paper are used. Stipple dots are typically placed intentionally one by one (in no particular order), while trying to avoid visual artifacts [Hod03] unless such artifacts are intended to represent specific features. Stipple dots can overlap each other—very dark zones in an image can be produced using many overlapping stipples. Figure 1 shows examples of hand-made stippling from three different artists, along with one detail view from each of these images.

Like other artistic techniques, stippling is a composition of three tasks: an artistic task, a procedural task, and an instrumental task. Only by mastering all three can an artist produce aesthetically pleasing depictions that convey the intended message (such as in an illustration). While the lowest-level instrumental task of placing dots is rather mechanical and can be acquired relatively easily, the other two tasks require an increasing amount of training and artistic skill.

In the *artistic task* the stipple artist selects what to represent and what to leave out, where detail needs to be provided and where a more abstract, simplified representation is sufficient, where to stick to the original source and where artistic freedom can be employed, etc. Martín et al. [MALI10, MALI11] mentioned the artistic task when they talked about "high-level processes" in hand-drawn stippling. Ultimately, these activities have a huge influence on the aesthetic and communicative qualities of the results and the resulting mental model of the viewer and require imagination, creativity, emotions, domain knowledge, etc. to master. They are thus beyond the scope of NPAR support at this time, and remain a challenge for future work to at least partially support with computer tools.

The *procedural task* refers to the way the ideas and concepts derived in the artistic task are converted into visual artifacts, i. e., the arrangement of dots on the paper. This task, consequently, is the one that most traditional digital stippling techniques have supported by deriving stipple dot distributions based on some input data. This task relates much to the unique style of a stipple artist: Kim et al. [KMI*09] pointed out in their conference talk that, based on their example-based method for capturing stipple distributions, stipple artists were able to recognize their own stippling style as well as that from colleagues. Beyond a personal style, the procedural task also refers to stylistic choices such as the regularity of the dot placement, the density of the dot placement, and the use of effects such as overlapping. These different strategies can be used, e. g., to replicate different materials and textures of the depicted objects. The existing approaches in the NPAR domain for supporting the procedural task of digital stippling have created varying rates of success: While Isenberg et al. [INC*06] found that it was quite easy to distinguish certain computer-created stippling images from hand-made ones in 2006, recent example-based stippling techniques would likely perform better if analyzed in a similar comparison today.

The *instrumental task*, finally, refers to the media and tools that are used in the stippling process and the low-level actions for placing the dots with the chosen tool onto the chosen medium. In this context it is important to mention that the human visual system works at different levels of details si-

(a) Hand-made example. © Randy Glass. (b) Hand-made example. © Miguel Endara. (c) Hand-made example. © Elena Piñar.

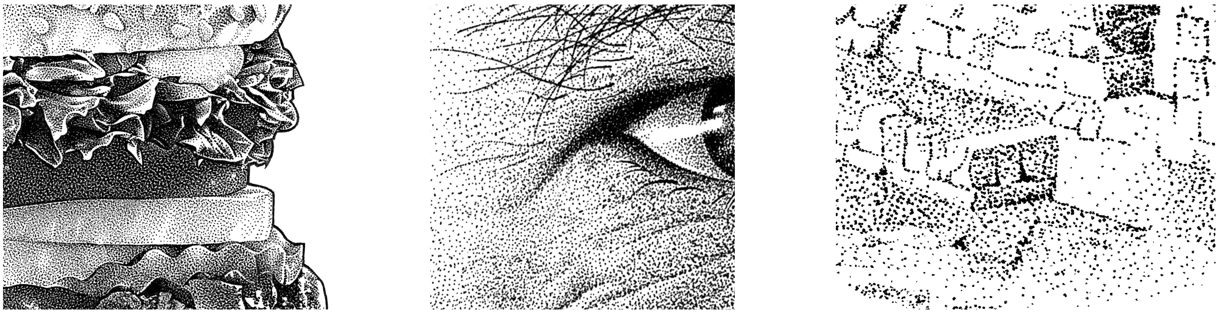

(d) Detail. © Randy Glass. (e) Detail. © Miguel Endara. (f) Detail. © Elena Piñar.

Figure 1: *Examples of hand-made stippling (some of the images have been down-sampled). All images used with permission.*

multaneously, ranging from the small part of the focus of attention to aspects of an image that are perceived about the rest [Mar82]. The choices and actions in the instrumental task thus not only affect the local appearance of each individual dot but also influence the overall appearance of the artwork. This means that, while certain details are important for specific image regions, the more global quality of dot placement are essential for the overall perception of the whole image.

To give a practical example, let's assume a stipple artist wants to reproduce a stippled representation of a landscape with mountains, trees, old houses, a river, etc. She starts by creating a mental model of the story she wants to tell and then decides what to depict based on what she sees in reality or on a picture, what not to show, and what to depict with less detail (i. e., the artistic task). She then starts placing stipple dots on paper, using different forms to place stipples guided by her experience and training that is reflected in her personal style (i. e., the procedural task). She also has made decisions on which pens to use (nib size), which paper (color, hot press vs. cold press paper) to draw on, as well as uses the pen with a given pressure and precision (i. e., the instrumental task).

As the artistic task is beyond the current abilities of NPAR techniques and because the procedural task is relatively well supported in our field already, we decided to focus on the instrumental task. Ultimately, our goal is to support stipple artists in the two lower-level tasks for computer-supported stippling, allowing them to concentrate on the artistic choices (like others have done for other artistic media such as watercolor [CAS*97] or pencils [AWI*09, SB99, SB00]). We thus focus in the physical characteristics of physically placed dots, their shape, their size, their color, as well as the possible relations between these characteristics. Similar to the lower-level characteristics of dot distributions, these attributes can have a significant influence on the final result, which motivates us in studying them to provide a solid foundation for a wide variety of existing and future digital stipple techniques.

4. From the Continuous to the Digital Domain

To establish this foundation it is necessary to understand to what degree a digital stipple dot can resemble a hand-made stipple dot, based on the type of digital reproduction. Therefore, to be able to transition/transfer from the continuous physical domain to the discrete digital domain, we need to understand both the characteristics of traditionally hand-made dots as well as the constraints of the digital reproduction such as spatial resolution, color resolution, reproduction medium (printing on paper vs. display on the screen), etc. as all these aspects influence the final perception of the result.

To obtain a solid basis for this study of realistic hand-placed stipple dots, we started by contacting two professional illustrators, Miguel Endara and Randy Glass,[†] as well as a

---

[†] http://miguelendara.com/ and http://www.randyglassstudio.com/ .

Illustrator	Pen type	Nib size	Paper type
Randy Glass	Rotring Rapidograph	0.13mm	FLAX (Medium press)
Miguel Endara	Sakura Pigma Micron Pen 005	0.20mm	Strathmore 500 Illustration Board (Hot press)
Elena Piñar	Artline Drawing System	0.50mm	Canson Watercolor (Cold press)

Table 1: *Tool/material preferences of the three collaborating professional artists.*

part-time illustrator, Elena Piñar. Examples of their work are shown in Figure 1. Specifically, we asked these professionals about their use of paper and pen types for their stipple work (a summary of these preferences is shown in Table 1).

Based on this information we set out to create a dataset to allow us to compare the different pen and paper types. We decided focus on these aspects and leave out pressure and pen angle—pen angle does not seem to be frequently used in stippling and we kept pressure constant as we expected it to affect the outcome proportionally for different paper and pen types. We asked a skilled fine arts student[‡] (fourth year, for whom stippling is part of her education) to draw 5 different tones out of a gray ramp, using the three different pens from Table 1 and three different types of paper. The latter choice of paper is equally important as the pen type as, for example, cold press paper has texture and absorbs the ink quickly, while hot press paper is smooth and absorbs the ink more slowly. We thus used a cold press paper (Canson Watercolor; 370g/m^2), a medium press paper (Canson Graphics Art; 224g/m^2), and a hot press paper (Canson Technical Drawing; 160g/m^2).

We then digitized the resulting samples using an Epson Perfection V700 Photo scanner with an optical resolution of 4800 ppi.[§] Figure 2 shows a number of example sections from these scans,[¶] each a 1200 × 1200 pixel section from the scan. They are displayed at 8× magnification to allow us to examine the individual dot shapes in detail. These examples clearly show that stipple dots vary significantly in size and shape. As was to be expected, the size depends primarily on the size of the pen's nib, yet also varies for dots created with the same pen on the same paper. Moreover, we can see that both the size and the dot boundaries depend on the paper type, with cold press paper leading to more ink absorption artifacts than medium and cold press paper. This difference in diffusion/absorption can be seen, in particular, in Figures 2(g)–2(i)—the effect being stronger for pens with larger nib sizes.

To look at the individual dots even more closely we also show 240 × 240 pixel detail sections at a 40× magnification in Figure 3. For comparison, we included circles with diameters of 25 pixels, 38 pixels, and 94 pixels, respectively, corresponding to the sizes of the nibs of 0.13 mm, 0.2 mm, and 0.5 mm at the resolution of 4800 ppi. In the figure we can observe that the overall sizes of the actual stipples deviate from that of the pen's nib; in our case they are larger for the small and medium nib sizes and smaller for large nib size (it depends on the chosen pens as well as used pen orientation and pressure). Moreover, we can clearly see, in particular in Figures 3(g)–3(i), that the dots are never completely black. Instead, they exhibit a pattern of gray shades that depends on the underlying paper, like previously pointed out [MALI10, MALI11]. This is, in fact, to be expected as the paper-ink interaction is similar to that of watercolor, a field that has received much attention in NPAR in the past and for which elaborate schemes for the simulation of such diffusion/absorption patterns have been created (e. g., [CAS*97, DKMI13]).

Beyond this discussion of the dot's characteristics based on magnified scans we also wanted to understand how they are perceived by humans. For this purpose we recruited 11 unpaid volunteers (7 females; ages in the range of 22–25; fine arts graduate students from the local university population). All had experience in drawing with pen and ink, including stippling. We asked them to examine the previously created paper samples and to evaluate them based on their color (black vs. gray), shape (round vs. irregular), and size (constant vs. varying). For this purpose the participants were seated on a desk lit with a constant artificial fluorescent light source, and participants went through the stack of images in randomized order, filling out a questionnaire as they progressed.

Table 2 shows the results of this small perceptual survey. The data suggests that, in many cases, the stipple dots are indeed perceived as being black, but for some combinations of nib size and paper a considerable number of people perceives them as being rather gray—smaller dots more often seem to be perceived as being gray. It is interesting to note that, despite the discussed differences in diffusion/absorption for cold press paper vs. medium and hot press paper, our participants perceived the 0.5 mm pen on cold press paper as black, while 27% of our participants saw the dots of the same pen on hot press paper as gray. With respect to the dot size and shape, a large number of participants did perceive the irregularities in shape and size as we expected.

The results from both our detailed analysis of scanned samples of stipple dots as well as of our perceptual survey on what people perceive when they look at hand-made stippling suggests that—if we are interested in a faithful reproduction of stippling or in supporting artists with a tool—it is essential that we do not represent stipple dots exclusively as completely black circles or rounded shapes. Instead, we need to capture and reproduce the characteristics of hand-made stipples on paper. In particular the process of reproduction of the captured dots, however, highly depends on how and where the final result will be used so that we need to discuss the constraints of this reproduction process next.

[‡] To study the low-level aspects of instrumental task it is not necessary to rely on the artistic skills of a professional stipple artist.

[§] Please note that we make the important distinction between ppi as a unit when we talk about pixels in scanning and image processing and dpi when we talk about the dots in the printing process.

[¶] The full dataset is available as additional material.

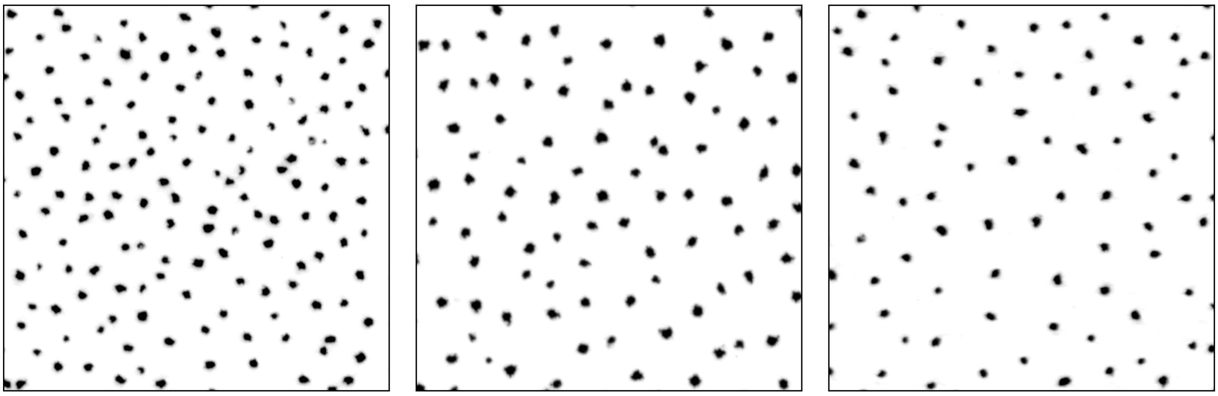

(a) Rotring rapidograph 0.13, cold press paper. (b) Rotring rapidograph 0.13, medium press paper. (c) Rotring rapidograph 0.13, hot press paper.

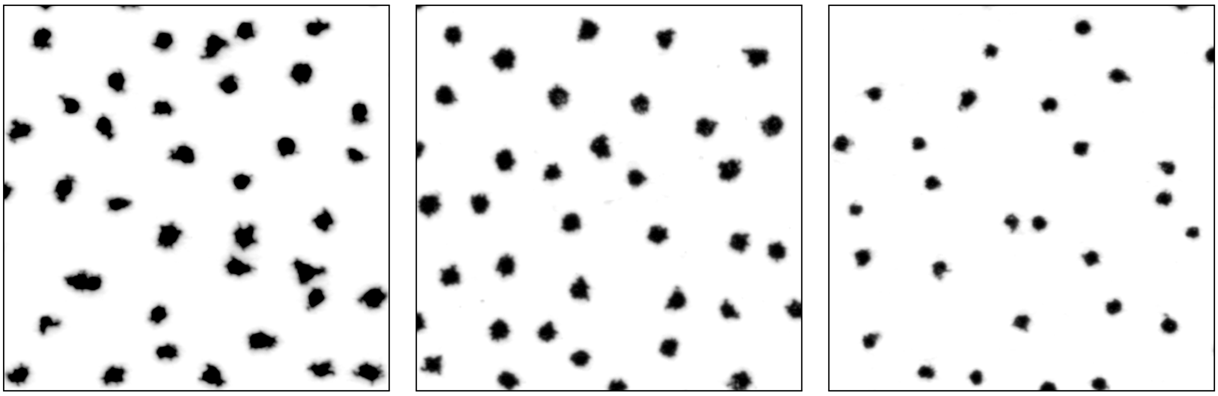

(d) Sakura Pigma 005, cold press paper. (e) Sakura Pigma 005, medium press paper. (f) Sakura Pigma 005, hot press paper.

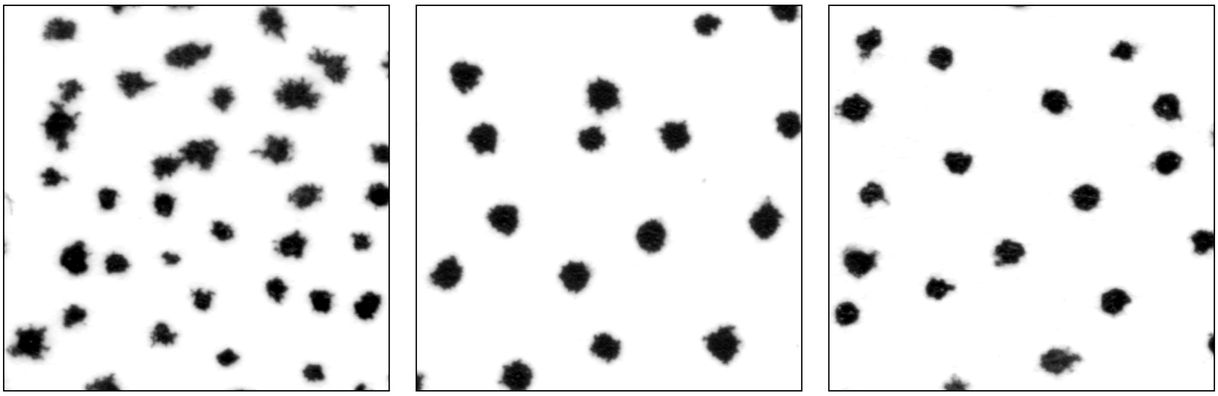

(g) Artline 0.5, cold press paper. (h) Artline 0.5, medium press paper. (i) Artline 0.5, hot press paper.

Figure 2: *Samples from the stipple dot experiments shown at 8× magnification.*

5. Dot Reproduction: Potential & Limitations

The results of digital stippling can be displayed on a variety of output media, the main types being presentation on a screen and print reproduction. These forms of reproduction, however, have quite different capabilities ad limitations. The same result such as in the form of a PDF document can even be intended to both being shown on the screen and being printed. We thus discuss the different capabilities and limitation next.

5.1. Printing

Let us first focus on printing because, in that case, the result is reproduced on paper, similar to the hand-made original.

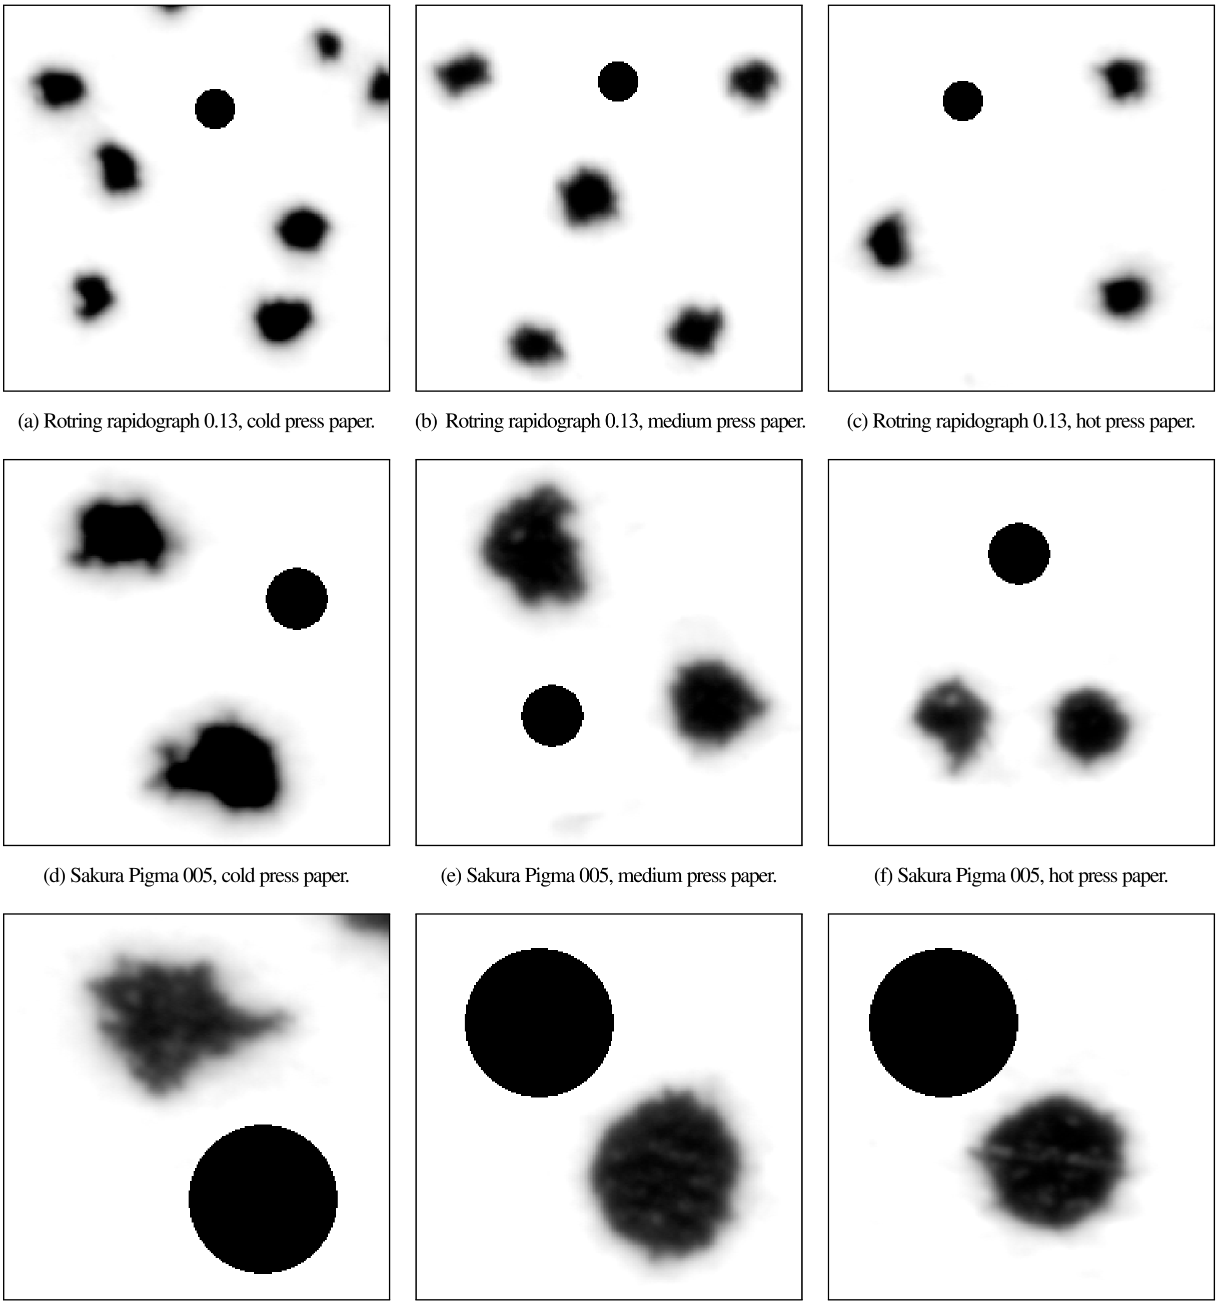

(a) Rotring rapidograph 0.13, cold press paper.

(b) Rotring rapidograph 0.13, medium press paper.

(c) Rotring rapidograph 0.13, hot press paper.

(d) Sakura Pigma 005, cold press paper.

(e) Sakura Pigma 005, medium press paper.

(f) Sakura Pigma 005, hot press paper.

(g) Artline 0.5, cold press paper.

(h) Artline 0.5, medium press paper.

(i) Artline 0.5, hot press paper.

Figure 3: *Details from Figure 2 at 40× magnification, with black circles added that represent the nib sizes of the used pens.*

While several traditional printing techniques exist, in today's digital world we typically use either laser printers or inkjet printers. Both place toner or ink dots on paper, either just in black or using CMYK primaries. All these technologies, however, share the same limitation that they can only use pure colors/toner/ink and are not able to directly produce shades of gray or of a color. This means that the result of any stippling technique that uses scanned or simulated grayscale stipple dots needs to undergo a conversion to a binary dot pattern that resembles the intended gray values. This conversion, of course, is subject to the printer's output resolution as we showcase next.

Feature	Cold press paper			Medium press paper			Hot press paper		
	0.13	**0.2**	**0.5**	**0.13**	**0.2**	**0.5**	**0.13**	**0.2**	**0.5**
Color (black / gray)	73% / 27%	82% / 18%	100% / 0%	100% / 0%	91% / 9%	100% / 0%	55% / 45%	55% / 45%	82% / 18%
Shape (regular / irregular)	55% / 45%	36% / 64%	27% / 73%	36% / 64%	45% / 55%	45% / 55%	64% / 36%	9% / 91%	0% / 100%
Size (constant / varying)	91% / 9%	36% / 64%	27% / 73%	55% / 45%	45% / 55%	55% / 45%	64% / 36%	27% / 73%	9% / 91%

Table 2: *Results of the perceptual survey for color, shape and size; depending on the paper type and the pen's nib sizes (in mm). The numbers indicate the percentage of people who answered in the respective fashion.*

For our discussion we assume that we are interested in reproducing the stipple image with a realistic spatial size (like also done by Martín et al. [MALI10, MALI11])—similar in dimension to hand-made originals such that the stipple dots have the correct size. Based on this assumption and the measurements discussed in the previous section we can try to analyze the effects of the printing process. Table 3 lists the number of horizontal or vertical dots that are available to represent a single stipple, depending on the printer's resolution and the used nib sizes. We can see that, for a 0.13 mm pen and a 300 dpi printer, less than 2 × 2 printer dots are available to completely represent the stipple. This is not only not enough to show a rounded dot, it is also certainly not sufficient to represent the complex shapes and texture of the stipples.

For today's printers which typically have a resolution of 1200 dpi the situation appears to be better: we have 6 × 6, 9 × 9, and 24 × 24 dots available for stipple sizes of 0.13 mm, 0.2 mm, and 0.5 mm, respectively. This may seem to be sufficient to capture the stipple shape, at least for the larger pen sizes. We have to take into account, however, that to be able to also represent the stipples' textures the printer has to use some form of halftoning [Uli87]. Because halftoning trades spatial resolution for color (grayscale) resolution, this process reduces the effective resolution that is available to represent stipples by at least 3 in each direction, thus to a maximum of 2 × 2, 3 × 3, and 8 × 8, respectively—in most cases not enough for an adequate reproduction of the stipple shapes. For a proper representation of 256 gray levels it would even be necessary to reduce the spatial resolution by 16 in both directions using AM halftoning [BR93, Gre99, Uli87] or by a factor of 8–12 using error diffusion [Agf94]. Moreover, the halftoning process introduces dot patterns that are supposed to emulate gray values. These patterns can still be perceived by humans at 1200dpi when looked at from a typical reading distance (25 cm–60 cm) [Gre99] and thus interfere with the perception of the stipple dot pattern.

The reason for this effect is the *visual acuity* of the human visual system [ICO84, Gre99]. Visual acuity is based on the limit of feature recognition for humans of 1 arc minute. For example, as Table 4 shows, at a reading distance of 50 cm people can distinguish features of approx. 0.14 mmm, corresponding to 183 dpi (or ppi). While this result would suggest that a 1200 dpi printer is sufficient for reproducing the gray values using halftoning, we have to take into account the size of the halftoning patterns that create the illusion of color shades and multiply the needed resolution by it as demonstrated in Table 4. This means that, for typical reading distances, at least a 2400 dpi printer is necessary such that the dot patterns from halftoning are no longer perceived [Gre99].

Nib	300 dpi	1200 dpi	2400 dpi	4800 dpi
0.13mm	1.54	6.14	12.28	24.57
0.20mm	2.36	9.45	18.90	37.38
0.50mm	5.91	23.62	47.24	94.49

Table 3: *Number of printer dots available to represent a stipple dot, depending on print resolution and pen's nib size.*

	25 cm	35 cm	50 cm	100 cm
mm per dot / pixel	0.07	0.10	0.14	0.28
dpi / ppi	366	261	183	91
needed dpi w/ halftoning, 8×	2926	2090	1463	732
needed dpi w/ halftoning, 16×	5852	4180	2926	1463

Table 4: *Maximum perceivable dot/pixel resolution depending on the viewing distance for 100% visual acuity. Resulting need for resolution in dpi when halftoning is used.*

5.2. Display on a Screen

A different situation arises when we try to display stipple images on screens. Regardless whether they are based on CRTs, LCDs, or OLEDs as their underlying technology, screens have the important benefit of being able to show color gradients—typically with at least 8 bits per color, newer devices with up to 12 bits per color. The same is true for e-ink displays, although the current technologies are limited to 4 bits or 16 shades of gray. They should thus all be able to depict the grayscale aspects of stipple dots well.

The problem arises due to the limited spatial resolution of today's screens as illustrated in Table 5. Current devices use spatial resolutions of around 100 ppi, even with newer UHD displays remaining at less than 200 ppi. Also the most recent e-ink displays have a maximum resolution of 212 ppi. Only some mobile phones have reached resolutions of 300 ppi or more, but their physical screen space is too small to realistically be considered for the display of stipple images.

Using 112 ppi as an example for a typical screen resolution,

Size	1360 × 768	1920 × 1080	2560 × 1600	3840 × 2160
15"	**104**	147	201	294
24"	65	**92**	126	184
27"	58	82	**112**	163
30"	52	73	**101**	147
32"	49	69	**94**	138

Table 5: *Resolutions in ppi for common monitor sizes and pixel counts. The most commonly used types in bold.*

Nib	76 ppi	92 ppi	112 ppi	138 ppi
0.13 mm	0.4	0.5	0.6	0.7
0.20 mm	0.6	0.7	0.9	1.1
0.50 mm	1.5	1.8	2.2	2.7

Table 6: *Number of (horizontal/vertical) pixels available to represent a single stipple dot, based on the pen's nib size.*

this means that each pixel has the size of 0.23 mm × 0.23 mm. It is thus not possible to use more than a single pixel to represent a stipple dot when showing stipple images at their correct size—less than one pixel is available to show stipples for pens with nib sizes of 0.13 mm and 0.2 mm. Table 6 shows a summary of the number of pixels available for a single stipple dot, depending on the pen's nib sizes and the display's resolution. In most cases there are less than 2 × 2 pixels available to display a single stipple, and even for the most high-resolution displays we have less than 3 × 3 pixels at our disposal. This means that, while screens would be able to show the grayscale properties well, they do not have a sufficient resolution to represent the irregular shapes of realistic stipple dots.

6. Taxonomy of Stippling Reproduction Goals

Based on these considerations on the capabilities and limitations of the different output media we can now propose a taxonomy of stipple reproduction goals, and give recommendations on how to achieve them. We start with a classification based on the aspects of stipple distribution/placement, as we reviewed them in Section 2:

1. **Stipple distribution quality:** Depending on the approach to place stipples, different quality classes can be distinguished:

   - example-based,
   - ensuring noise attributes, or
   - others.

It can be argued that example-based distributions (e. g., [BBMT06, KMI*09]) provide the best possible approximation of hand-drawn stipple distributions as they derive their dot patterns from exactly such hand-drawn examples. In contrast, stipple distributions that approximate certain noise characteristics (e. g., [BSD09,KCODL06,VBTS07,ALMPHS10]) could be seen as ones that aim for a generally good quality of stipple dot locations that avoids visual artifacts. Finally, other techniques such as those based on Lloyd's method (e. g., [DHvOS99, DHvOS00, Sec02]) often exhibit unwanted artifacts, unlike the human-made examples. Of course, there are other aspects of stipple dot distribution that we do not discuss here because they are beyond the scope of this paper. These aspects include, for instance, whether the distributions support the merging of points, whether they support animation, or whether they facilitate different levels of zoom.

We base the next classification on the discussions in this paper, namely the characteristics of the stipple dots:

2. **Stipple dot quality:** Depending on the desired level of reproduction for the stipple dots we distinguish whether:

   - stipple shape and texture are reproduced,
   - only the stipple shape is reproduced, or
   - neither stipple shape nor texture are reproduced.

As we showed, due to the existing technological limitations it is impossible to faithfully show stipple images on today's screen hardware using the intended (traditional) physical sizes. For print reproduction, however, we saw that—while today's printers with a typical resolution of 1200 dpi are still too limiting—it would be possible to faithfully show stippling at resolutions of 2400 dpi and above. Such resolutions exist in imagesetter hardware as it is used professionally. Previous approaches for digital stippling that would support such output are, in particular, Martín et al.'s [MALI10, MALI11] resolution-dependent grayscale stippling.

If we relax the goal of capturing and reproducing the stipple texture, we can still aim for trying to reproduce the stipple shapes. This has no effect on the applicability for on-screen display as the needed spatial resolution does not change. In the printing process, however, we can use 1 bit black-and-white representations, derived from thresholded stipple scans or simulations. These are simply re-scaled by the printer without halftoning, leading to a faithful reproduction of the stipple shapes. In the past, this approach was used by Kim et al. [KMI*09] and Martín et al. [MALI10, MALI11] as well as in many traditional reproductions of hand-made stippling.

Finally, we can relax the goals of reproducing the stipple shape and texture completely—meaning that we only use a chosen stipple distribution to place simple dots, circles, or pixels. This approach is not only possible for print reproduction but also for on-screen display as we have seen in Table 4. In fact, this approach is interesting as it implies that we use the stipple distribution as a form of halftoning. In the past, this approach has been used not only by early approaches that used circular stipple dots (e. g., [DHvOS99, DHvOS00, Sec02]) but also, for example, by Kopf et al. [KCODL06], Schmidt et al. [SIJ*07], Kim et al. [KMI*09], and Ascencio-Lopez et al. [ALMPHS10] as they placed stipple dots in the form of pixels or very small points.

In fact, based on our discussion one could argue that—for small pens with approx. 0.1 mm nib sizes, hot press paper to avoid fuzzy boundaries, and printers or displays with 300 dpi resp. 300 ppi resolutions—each stipple would be equivalent to a dot or pixel in size. This could be interesting, in particular, for inkjet printers as they do deposit real ink on real paper. The caveat in this case, however, is that such printers (and equivalent displays) use a regular grid on which the dots (or pixels) are arranged. This means that this regular arrangement would only not be apparent for sparse stipple densities.

Of course, all our discussion of stipple dot properties and the quality constraints w.r.t. their reproduction rests on the assumption that we aim to reproduce them at the correct spatial size—equivalent to what a traditional artist would have done on paper. Because this may not always be the goal for digital stippling, we can add the following classification:

3. **Stippling reproduction size:** Depending on the intended spatial dimensions for the final product, we distinguish whether:
   - the stippling image and the stipple dots should be shown approximately at the same dimensions at which equivalent traditional artworks/illustrations would have been created or
   - other sizes are acceptable.

If the former is the case, then all our consideration provided in this paper and summarized in the previous element of the taxonomy apply. If one does not feel bound by these constraints, however, we can use stippling in a more flexible way, such as as a means to distribute certain shapes in the plane (e. g., [HHD03, DKLS06]) or as the basis to derive patterns for halftoning. Similarly, approaches that involve animations, stippling textures on 3D shapes (e. g., [LTH*02, LMT*03, MS02, MFS03, MPS04, YNZC05, XC04, ZS04]), or stippling inside volumes (e. g., [LME*02, LMT*03]) are less likely to be bound by a 1:1 mapping of spatial sizes. However, such freedom implies that the produced stippling is more removed from the artistic example, which then only serves as an inspiration for such techniques. Nevertheless, we can still discuss the distribution quality (classification # 1 in our taxonomy) and, to some degree, the dot quality (classification # 2 in our taxonomy).

7. Discussion, Implications, and Future Work

Ultimately, we thus raise the question of the intentions and goals in non-photorealistic rendering. As we demonstrated in this paper, the goal of a faithful reproduction of stippling according to the artistic example can be one goal, for example when we aim to create tools to be used by artists and illustrators. In that case, however, we have to carefully study the traditional technique and consider the implications of the entire toolchain, from the creation of digital elements to the final reproduction of the results. While there are certainly many other valid goals and motivations for NPR work, if we assume the faithful reproduction (such as for tools intended to be used by artists or illustrators) as our goal we need to discuss a number of implications.

Because current display hardware does not have the sufficient resolution to display faithfully created stipple images, it also does not make sense to use displays as the primary output medium. However, displays certainly need to be used as a tool for the production of digital stippling by artists and illustrators. The respective tools thus not only have to work at a much higher resolution than can be displayed on the display, but the tool also has to show both a correctly scaled representation of the produced work as well as zoomed-in versions to see the detail. Only with such a focus+context view will it be possible to the artists to work at both the overview and detail levels in the way they are used to in their traditional practice. Morever, the tools have to provide means to output the result in several different versions: a high-reolution grayscale (color) version, a high-resolution 1 bit version that only captures the stipple shapes, and potentially a version that uses the 1 stipple = 1 pixel convention. A vector graphic output (as it is sometimes advocated for NPR work [ICCS05]) would be useful if it captures the shape of the stipples at a high resolution, which would avoid the re-sampling of 1 bit raster images at the printer. In the future, maybe there will be vector graphic languages that also are capable to capture the texture faithfully—some initial attempts in this direction have already been presented [SLWS07, OBW*08, BEDT10, JCW11, BB13, BDF14].

The implications for print reproduction arise from certain assumed standards in the publishing domain. In particular, publishers frequently ask for 300 ppi images for the inclusion in material intended to be printed (in academic publishing and elsewhere). As we have seen, this resolution is neither sufficient to reproduce stippling in full fidelity, nor does it suffice for a shape-only black-and-white reproduction. In these cases one has to try to convince the publishers of the specific needs of stippling as a medium, and at least aim for 1200 ppi in 1 bit black-and-white mode to be able to capture the stipple shapes—a representation that uses similar or less bandwidth than a 300 ppi grayscale or full color image [ICCS05]. If publishers do not impose a limit such as 300 ppi for images, of course, one can and should embed results with higher resolutions since imagesetter hardware supports resolutions of 2400 dpi and above.

These implications for print reproduction and display on a screen also affect those media forms that were created for both forms of output, such as PDF or Postscript documents. Here we face the diverging capabilities/constraints that cannot be met at the same time. A good compromise seems to be to use 1 bit black-and-white images as done in Figures 1(c) and (f) in this paper—it both prints well and can also be used to display good versions on a screen through interpolation, at the intended or at a zoomed-in scale, albeit at the expense of loosing the stipple textures. Future media may provide different versions of the embedded images based on the present output device and the currently employed zoom.

It may also be worth considering the implications the decisions that are made at the instrumental task level. For example, the use of pens with small nib sizes and hot press papers not only increase the possible stipple dot numbers/frequencies but also lead to a darker tone, more similar sizes, and more regular shapes. Such characteristics better suit the existing forms of reproduction and produce stipple dots similar to what has been used in NPR in the past, albeit these are still not perfect circles in pure black.

Finally, in the future it may also be interesting to consider dedicated output devices that physically reproduce some of the low-level aspects of the instrumental task. For example, plotters could be equipped with a real pen and used for realistic output of stipple images, similar to what has been explored with robots for NPR painting [Hag92, DLPT12, LMPD15], portrait drawing [TL13], or Chinese painting [YS05].

8. Conclusion

The evolution of techniques and computing hardware allows us to (re-)produce stippling at an ever-increasingly quality. While researchers concentrated mostly on stipple point distributions in the past, we showed that it is equally important to consider the reproduction of the dots themselves. We discussed, in particular, the constraints that we need to consider when producing stippling that is faithful to its hand-made counterpart—depending on the used materials/pens on the one hand and the reproduction type on the other hand.

While we did not present a dedicated technique to reproduce realistic dots, we provide the dataset of high-resolution scans of stipple gray-ramps that we discuss in this paper under a creative-commons license for future study. We hope that these samples can be used in future example-based techniques. For example, the dots could be used directly as a set of discrete examples or as a basis to come up with an example-based stipple dot synthesis technique.

Our work not only provides a better insight on the size, shape, and texture of real stipple dots but, with our taxonomy, we also discuss the question of goals and intentions of NPR work. Our discussion was mostly driven by the ultimate goal of producing tools to be used by artists and illustrators—a goal that has been one of the driving forces of much of the work in NPR. We are convinced that we have made quite some progress toward this goal with our discussion and the insights that we present in this paper.

As we also discussed in Section 3, both our work and much of the past work in NPR stippling only addresses the instrumental task of stippling and leaves the artistic and procedural tasks largely untouched. The (at least partial) support of these activities in NPR tools, we believe however, are necessary to create tools that will have a practical impact for artists and illustrators and should thus become a much more important aspect of future NPR work.

Acknowledgments

We thank Miguel Endara, Randy Glass, and Elena Piñar for the permission to use their artwork in our paper. We also thank Ana Belén Cruz for the samples we used in our discussion and Germán Arroyo for feedback on the manuscript. This work has partially been funded by the Spanish Ministry of Economy and Competitiveness through the project TIN 2014-60956-R with FEDER funds.

References

[Agf94] AGFA: *An Introduction to Digital Scanning*. Tech. rep., Agfa, 1994.

[ALMPHS10] ASCENCIO-LOPEZ I., MERUVIA-PASTOR O., HIDALGO-SILVA H.: Adaptive incremental stippling using the Poisson-disk distribution. *Journal of Graphics, GPU, and Game Tools 15*, 1 (2010), 29–47. doi> 10.1080/2151237X.2010.10390650

[AML10] ARROYO G., MARTÍN D., LUZÓN M. V.: Stochastic generation of dots for computer aided stippling. *Computer-Aided Design and Applications 7*, 4 (2010), 447–463. doi> 10.3722/cadaps.2010.447-463

[AWI*09] ALMERAJ Z., WYVILL B., ISENBERG T., GOOCH A. A., GUY R.: Automatically mimicking unique hand-drawn pencil lines. *Computers & Graphics 33*, 4 (Aug. 2009), 496–508. doi> 10.1016/j.cag.2009.04.004

[BB13] BARLA P., BOUSSEAU A.: Gradient art: Creation and vectorization. In *Image and Video based Artistic Stylisation*, Rosin P., Collomosse J., (Eds.), vol. 42 of *Computational Imaging and Vision*. Springer, London, Heidelberg, 2013, ch. 8, pp. 149–166. doi> 10.1007/978-1-4471-4519-6_8

[BBMT06] BARLA P., BRESLAV S., MARKOSIAN L., THOLLOT J.: *Interactive Hatching and Stippling by Example*. Tech. rep., INRIA, 2006.

[BDF14] BENJAMIN M. D., DIVERDI S., FINKELSTEIN A.: Painting with triangles. In *Proc. NPAR* (2014), ACM, New York, pp. 13–20. doi> 10.1145/2630397.2630399

[BEDT10] BEZERRA H., EISEMANN E., DECARLO D., THOLLOT J.: Diffusion constraints for vector graphics. In *Proc. NPAR* (2010), ACM, New York, pp. 35–42. doi> 10.1145/1809939.1809944

[BMCS*10] BRAZIL E. V., MACÊDO I., COSTA SOUSA M., VELHO L., DE FIGUEIREDO L. H.: A few good samples: Shape & tone depiction for hermite RBF implicits. In *Proc. NPAR* (2010), ACM, New York, pp. 7–15. doi> 10.1145/1809939.1809941

[BR93] BLATNER D., ROTH S.: *Real World Scanning and Halftones*. Peachpit Press, 1993.

[BSD09] BALZER M., SCHLÖMER T., DEUSSEN O.: Capacity-constrained point dostributions: A variant of Lloyd's method. *ACM Transactions on Graphics 28*, 3 (Aug. 2009), 86:1–86:8. doi> 10.1145/1531326.1531392

[CAS*97] CURTIS C. J., ANDERSON S. E., SEIMS J. E., FLEISCHER K. W., SALESIN D. H.: Computer-generated watercolor. In *Proc. SIGGRAPH* (1997), ACM, New York, pp. 421–430. doi> 10.1145/258734.258896

[Deu09] DEUSSEN O.: Aesthetic placement of points using generalized Lloyd relaxation. In *Proc. CAe* (2009), Eurographics Association, Goslar, Germany, pp. 123–128. doi> 10.2312/COMPAESTH/COMPAESTH09/123-128

[DHvOS99] Deussen O., Hiller S., van Overveld C. W. A. M., Strothotte T.: Computer-generated stipple drawings. In *Proc. Workshop on Vision, Modelling and Visualization* (1999), pp. 329–338.

[DHvOS00] Deussen O., Hiller S., van Overveld C., Strothotte T.: Floating points: A method for computing stipple drawings. *Computer Graphics Forum 19*, 3 (Sept. 2000), 41–50. doi> 10.1111/1467-8659.00396

[DI13] Deussen O., Isenberg T.: Halftoning and stippling. In *Image and Video based Artistic Stylisation*, Rosin P., Collomosse J., (Eds.), vol. 42 of *Computational Imaging and Vision*. Springer, London/Heidelberg, 2013, ch. 3, pp. 45–61. doi> 10.1007/978-1-4471-4519-6_3

[DKLS06] Dalal K., Klein A. W., Liu Y., Smith K.: A spectral approach to NPR packing. In *Proc. NPAR* (2006), ACM, New York, pp. 71–78. doi> 10.1145/1124728.1124741

[DKMI13] DiVerdi S., Krishnaswamy A., MAch R., Ito D.: Painting with polygons: A procedural watercolor engine. *IEEE Transactions on Visualization and Computer Graphics 19*, 5 (May 2013), 723–735. doi> 10.1109/TVCG.2012.295

[DLPT12] Deussen O., Lindemeier T., Pirk S., Tautzenberger M.: Feedback-guided stroke placement for a painting machine. In *Proc. CAe* (2012), Eurographics Association, Goslar, Germany, pp. 63–70. doi> 10.2312/COMPAESTH/COMPAESTH12/025-033

[FJW*05] Foster K., Jepp P., Wyvill B., Costa Sousa M., Galbraith C., Jorge J. A.: Pen-and-ink for BlobTree implicit models. *Computer Graphics Forum 24*, 3 (Sept. 2005), 267–276. doi> 10.1111/j.1467-8659.2005.00851.x

[Gre99] Green P.: *Understanding Digital Color*, 2nd ed. GATF Press, 1999.

[Hag92] Haggerty M.: Robotic painting: Art creating art. *IEEE Compter Graphics and Applications 12*, 4 (July/Aug. 1992), 8–10. doi> 10.1109/MCG.1992.10027

[Hau05] Hausner A.: Pointillist halftoning. In *Proc. CGIM* (2005), ACTA Press, Calgary, AB, Canada, pp. 134–139.

[HHD03] Hiller S., Hellwig H., Deussen O.: Beyond stippling – Methods for distributing objects on the plane. *Computer Graphics Forum 22*, 3 (Sept. 2003), 515–522. doi> 10.1111/1467-8659.00699

[Hod03] Hodges E. R. S. (Ed.): *The Guild Handbook of Scientific Illustration*, 2nd ed. John Wiley & Sons, Hoboken, NJ, USA, 2003.

[ICCS05] Isenberg T., Carpendale M. S. T., Costa Sousa M.: Breaking the pixel barrier. In *Proc. CAe* (2005), Eurographics Association, Goslar, Germany, pp. 41–48. doi> 10.2312/COMPAESTH/COMPAESTH05/041-048

[ICO84] ICO Visual Functions Committee: *Visual Acuity Measurement Standard*. Standard, International Council of Ophthalmology, 1984.

[INC*06] Isenberg T., Neumann P., Carpendale S., Costa Sousa M., Jorge J. A.: Non-photorealistic rendering in context: An observational study. In *Proc. NPAR* (2006), ACM, New York, pp. 115–126. doi> 10.1145/1124728.1124747

[JCW11] Jeschke S., Cline D., Wonka P.: Estimating color and texture parameters for vector graphics. *Computer Graphics Forum 30*, 2 (Apr. 2011), 523–532. doi> 10.1111/j.1467-8659.2011.01877.x

[KCODL06] Kopf J., Cohen-Or D., Deussen O., Lischinski D.: Recursive Wang tiles for real-time blue noise. *ACM Transactions on Graphics 25*, 3 (July 2006), 509–518. doi> 10.1145/1141911.1141916

[KCW*13] Kyprianidis J. E., Collomosse J., Wang T., , Isenberg T.: State of the “art”: A taxonomy of artistic stylization techniques for images and video. *IEEE Transactions on Visualization and Computer Graphics 19*, 5 (May 2013), 866–885. doi> 10.1109/TVCG.2012.160

[KMI*09] Kim S., Maciejewski R., Isenberg T., Andrews W. M., Chen W., Costa Sousa M., Ebert D. S.: Stippling by example. In *Proc. NPAR* (2009), ACM, New York, pp. 41–50. doi> 10.1145/1572614.1572622

[KSL*08] Kim D., Son M., Lee Y., Kang H., Lee S.: Feature-guided image stippling. *Computer Graphics Forum 27*, 4 (June 2008), 1209–1216. doi> 10.1111/j.1467-8659.2008.01259.x

[KWME10] Kim S., Woo I., Maciejewski R., Ebert D.: Automated hedcut illustration using isophotes. In *Proc. Smart Graphics* (2010), Springer, Berlin/Heidelberg, pp. 172–183. doi> 10.1007/978-3-642-13544-6_17

[Llo82] Lloyd S. P.: Least squares quantization in PCM. *IEEE Transactions on Information Theory 28*, 2 (Mar. 1982), 129–137. doi> 10.1109/TIT.1982.1056489

[LM11] Li H., Mould D.: Structure-preserving stippling by priority-based error diffusion. In *Proc. Graphics Interface* (2011), Canadian Information Processing Society, Waterloo, ON, Canada, pp. 127–134.

[LME*02] Lu A., Morris C. J., Ebert D. S., Rheingans P., Hansen C.: Non-photorealistic volume rendering using stippling techniques. In *Proc. VIS* (2002), IEEE Computer Society, Los Alamitos, pp. 211–218. doi> 10.1109/VISUAL.2002.1183777

[LMPD15] Lindemeier T., Metzner J., Pollak L., Deussen O.: Hardware-based non-photorealistic rendering using a painting robot. *Computer Graphics Forum 34*, 2 (May 2015). doi> 10.1111/cgf.12562

[LMT*03] Lu A., Morris C. J., Taylor J., Ebert D. S., Hansen C., Rheingans P., Hartner M.: Illustrative interactive stipple rendering. *IEEE Transactions on Visualization and Computer Graphics 9*, 2 (Apr.–June 2003), 127–138. doi> 10.1109/TVCG.2003.1196001

[LTH*02] Lu A., Taylor J., Hartner M., Ebert D. S., Hansen C. D.: Hardware-accelerated interactive illustrative stipple drawing of polygonal objects. In *Proc. VMV* (2002), Aka GmbH, pp. 61–68.

[MALI10] Martín D., Arroyo G., Luzón M. V., Isenberg T.: Example-based stippling using a scale-dependent grayscale process. In *Proc. NPAR* (2010), ACM, New York, pp. 51–61. doi> 10.1145/1809939.1809946

[MALI11] Martín D., Arroyo G., Luzón M. V., Isenberg T.: Scale-dependent and example-based stippling. *Computers & Graphics 35*, 1 (2011), 160–174. doi> 10.1145/1809939.1809946

[Mar82] Marr D.: *Vision*. The MIT Press, 1982.

[MF92] McCool M., Fiume E.: Hierarchical Poisson disk sampling distributions. In *Proc. Graphics Interface* (1992), Morgan Kaufmann Publishers Inc., San Francisco, pp. 94–105.

[MFS03] Meruvia Pastor O. E., Freudenberg B., Strothotte T.: Real-time animated stippling. *IEEE Computer Graphics and Applications 23*, 4 (July/Aug. 2003), 62–68. doi> 10.1109/MCG.2003.1210866

[MIA*07] Maciejewski R., Isenberg T., Andrews W. M., Ebert D. S., Sousa M. C.: Aesthetics of hand-drawn vs. computer-generated stippling. In *Proc. CAe* (2007), Eurographics Association, Goslar, Germany, pp. 53–56. doi> 10.2312/COMPAESTH/COMPAESTH07/053-056

[MIA*08] Maciejewski R., Isenberg T., Andrews W. M., Ebert D. S., Costa Sousa M., Chen W.: Measuring stipple aesthetics in hand-drawn and computer-generated images. *IEEE Computer Graphics and Applications 28*, 2 (Mar./Apr. 2008), 62–74. doi> 10.1109/MCG.2008.35

[Mou07] Mould D.: Stipple placement using distance in a weighted graph. In *Proc. CAe* (2007), Eurographics Association, Goslar, Germany, pp. 45–52. doi> 10.2312/COMPAESTH/COMPAESTH07/045-052

[MPS04] Meruvia Pastor O. E., Strothotte T.: Graph-based point relaxation for 3D stippling. In *Proc. Mexican International Conference on Computer Science* (2004), IEEE Computer Society, Los Alamitos, pp. 145–152. doi> 10.1109/ENC.2004.1342599

[MS02] Meruvia Pastor O. E., Strothotte T.: Frame-coherent stippling. In *Proc. Eurographics 2002, Short Presentations* (2002), Eurographics Association, Goslar, Germany, pp. 145–152.

[OBW*08] Orzan A., Bousseau A., Winnemöller H., Barla P., Thollot J., Salesin D.: Diffusion Curves: A vector representation for smooth-shaded images. *ACM Transactions on Graphics 27*, 3 (Aug. 2008), 92(1)–92(8). doi> 10.1145/1360612.1360691

[PZvBG00] Pfister H., Zwicker M., van Baar J., Gross M.: Surfels: Surface elements as rendering primitives. In *Proc. SIGGRAPH* (2000), ACM, New York, pp. 335–342. doi> 10.1145/344779.344936

[SB99] Sousa M. C., Buchanan J. W.: Computer-generated graphite pencil rendering of 3D polygonal models. *Computer Graphics Forum 18*, 3 (Sept. 1999), 195–207. doi> 10.1111/1467-8659.00340

[SB00] Sousa M. C., Buchanan J. W.: Observational model of graphite pencil materials. *Computer Graphics Forum 19*, 1 (Mar. 2000), 27–49. doi> 10.1111/1467-8659.00386

[Sec02] Secord A.: Weighted Voronoi stippling. In *Proc. NPAR* (2002), ACM, New York, pp. 37–43. doi> 10.1145/508530.508537

[SGS05] Schlechtweg S., Germer T., Strothotte T.: Renderbots—Multi agent systems for direct image generation. *Computer Graphics Forum 24*, 2 (June 2005), 137–148. doi> 10.1111/j.1467-8659.2005.00838.x

[SHS02] Secord A., Heidrich W., Streit L.: Fast primitive distribution for illustration. In *Proc. EGWR* (2002), Eurographics Association, Goslar, Germany, pp. 215–226. doi> 10.2312/EGWR/EGWR02/215-226

[SIJ*07] Schmidt R., Isenberg T., Jepp P., Singh K., Wyvill B.: Sketching, scaffolding, and inking: A visual history for interactive 3D modeling. In *Proc. NPAR* (2007), ACM, New York, pp. 23–32. doi> 10.1145/1274871.1274875

[SLKL11] Son M., Lee Y., Kang H., Lee S.: Structure grid for directional stippling. *Graphical Models 73*, 3 (May 2011), 74–87. doi> 10.1016/j.gmod.2010.12.001

[SLWS07] Sun J., Liang L., Wen F., Shum H.-Y.: Image vectorization using optimized gradient meshes. *ACM Transactions on Graphics 26* (July 2007), 11:1–11:7. doi> 10.1145/1276377.1276391

[SS02] Strothotte T., Schlechtweg S.: *Non-Photorealistic Computer Graphics. Modeling, Animation, and Rendering*. Morgan Kaufmann, San Francisco, 2002. doi> 10.1016/B978-1-55860-787-3.50019-0

[TL13] Tresset P., Leymarie F. F.: Portrait drawing by Paul the robot. *Computers & Graphics 37*, 5 (Aug. 2013), 348–363. doi> 10.1016/j.cag.2013.01.012

[Uli87] Ulichney R.: *Digital Halftoning*. MIT Press, 1987.

[VBTS07] Vanderhaeghe D., Barla P., Thollot J., Sillion F. X.: Dynamic point distribution for stroke-based rendering. In *Rendering Techniques* (2007), Eurographics Association, Goslar, Germany, pp. 139–146. doi> 10.2312/EGWR/EGSR07/139-146

[XC04] Xu H., Chen B.: Stylized rendering of 3D scanned real world environments. In *Proc. NPAR* (2004), ACM, New York, pp. 25–34. doi> 10.1145/987657.987662

[YNZC05] Yuan X., Nguyen M. X., Zhang N., Chen B.: Stippling and silhouettes rendering in geometry-image space. In *Proc. EGSR* (2005), Eurographics Association, Goslar, Germany, pp. 193–200. doi> 10.2312/EGWR/EGSR05/193-200

[YS05] Yao F., Shao G.: Painting brush control techniques in Chinese painting robot. In *Proc. ROMAN* (2005), IEEE Computer Society, Los Alamitos, pp. 462–467. doi> 10.1109/ROMAN.2005.1513822

[ZS04] Zakaria N., Seidel H.-P.: Interactive stylized silhouette for point-sampled geometry. In *Proc. GRAPHITE* (2004), ACM, New York, pp. 242–249. doi> 10.1145/988834.988876