



Special Section on CANS

Emotional response and visual attention to non-photorealistic images

David Mould^{a,*}, Regan L. Mandryk^b, Hua Li^a^a Carleton University, Canada^b University of Saskatchewan, Canada

ARTICLE INFO

Article history:

Received 2 November 2011

Received in revised form

30 March 2012

Accepted 31 March 2012

Available online 20 April 2012

Keywords:

Non-photorealistic rendering

Emotion

Affect

Arousal

Valence

Eye tracking

Gaze

ABSTRACT

Non-photorealistic rendering (NPR) algorithms are used to produce stylized images, and have generally been evaluated on the aesthetic qualities of the resulting images. NPR-produced images have been used for aesthetic and practical reasons in media intended to produce an emotional reaction in a consumer (e.g., computer games, films, advertisements, and web sites); however, it is not understood how non-photorealistic rendering affects the emotion portrayed in an image. We conducted a study of subjective emotional response and visual attention to five common NPR approaches, two blurring techniques, and the original image with 42 participants, and found that the NPR algorithms dampened participants' emotional responses in terms of arousal (activation) and valence (pleasure). Gaze data revealed that NPR rendering of images might reduce emotional response to an image by producing confusion, creating distracting visual artifacts, causing the loss of meaningful semantic information, or causing users to lose interest in the resulting image.

© 2012 Elsevier Ltd. All rights reserved.

1. Introduction

Non-photorealistic rendering (NPR) algorithms produce images in a wide range of expressive styles, including painting, drawing, and cartoons. Non-photorealistic images have some practical advantages over photographs, including a lack of distracting or irrelevant detail (e.g., in medical or archaeological illustrations), emphasis and clarification of crucial details (e.g., in caricature, maps, and technical illustrations), and, in some cases, ease of storage and reproduction. Beyond these practical advantages, NPR images have aesthetic benefits as the stylized images have an inherent beauty, vitality, and interest as compared to photographic images. Because of the practical and aesthetic advantages of stylized images, NPR algorithms have seen increasingly widespread use in computer games (e.g., the 2008 *Prince of Persia*, *Borderlands*, *Team Fortress 2*), films (e.g., *A Scanner Darkly*), and advertisements (e.g., Charles Schwab Investments), making the improvement of NPR algorithms an active area of research.

Researchers investing effort in improving NPR algorithms generally evaluate their results by considering the aesthetic quality of the resulting images. However, the use of the images for emotional media such as television, film, and advertisements means that the aesthetic quality of the image is not the only

consideration of success of an algorithm—we must also consider how the resulting image elicits an emotional response from the viewer. Media creators who intend to provoke an emotional response in the viewer of a stylized image need to know whether the use of a particular NPR algorithm alters the viewer's emotional response; however, researchers have little understanding of how the varying algorithms affect the perceived emotional content of the resulting image.

To determine how the perceived emotional content of the image is affected by different algorithms, we conducted a study investigating the emotional response of 42 participants to five well-known NPR algorithms (Haerberli's interactive painting [1]; the photo abstraction of Orzan et al. [2]; the abstract painting of Zhao et al. [3]; Secord's weighted Voronoi stippling [4]; and the line drawing technique of Kang et al. [5]), two blurring techniques that were used to reduce the image information in a systematic manner (uniform blur and salience-based blur), and the original image. Emotional response was measured using the three-dimensional model of emotion, which consists of arousal (level of activation), valence (positive versus negative), and dominance (level of control). Participants also rated the aesthetic quality of each image and ranked the techniques in a post-experiment questionnaire. Finally, participants' gaze was monitored when viewing the images to shed light on how what they were visually attending to may have affected their emotional response.

We found that the algorithms significantly affected user ratings of both arousal and valence. In general, the NPR techniques dampened the emotional response to the images, moving

* Corresponding author. Tel.: +1 613 520 2600x8975.

E-mail addresses: mould@scs.carleton.ca (D. Mould), regan@cs.usask.ca (R.L. Mandryk), hli1@connect.carleton.ca (H. Li).

participant responses toward neutral ratings. Image abstraction yielded emotional responses closest to the original image in both arousal and valence, whereas painterly styles showed the greatest flattening of emotional responses. Differences in arousal ratings between the algorithms grew larger as the images were more arousing, whereas differences in valence ratings between the algorithms were larger for low-valence and high-valence images than for neutrally valent images. Our results also show that the differential emotional responses cannot be solely attributed to information loss as a result of the filtering algorithms, and do not depend on whether participants had previously seen the original image. The gaze fixation results varied widely by image, but showed four main findings. First, some filters caused confusion in what the participants should have been attending to by making parts of the image indiscernible; second, some caused distraction by creating an area of visual interest where nothing salient had been present in the original image; third, some filters caused important semantic information in the original image to be ignored in the filtered image; and fourth, some filters made the resulting images uninteresting to the participants diminishing their engagement. Our results are of interest to NPR researchers, but are also of particular importance for artists, designers, and media creators who use the algorithms in media intended to produce an emotional response in a consumer, such as films, games, advertisements, and web sites.

This paper is an extended version of an earlier paper [6] which described the experiment and the emotional responses. The discussion of visual attention, constituting Section 7, is new to this version.

2. Related work

We first give an overview of non-photorealistic rendering and some of the work that has sought to evaluate synthetic images, followed by a primer on representations and measurement schemes for emotions, and finally a brief commentary on emotional responses to NPR.

2.1. Non-photorealistic rendering

Since the beginnings of NPR, myriad processes have been devised for synthesizing images in a wide range of different styles and traditional artistic media, including line art, mosaics, and painterly. While some methods are interactive and depend on user input to create images, others are automated and require only a scene description, either in the form of geometry or as an input image. Note that owing to our use of images as input, our study looked only at image processing stylization algorithms; we did not attempt to include interactive systems or rendering based on geometry. We lacked geometric descriptions of the scenes to be rendered, and eschewed interactive techniques to limit the potential for bias from the skill of the user. Even automatic techniques usually benefit from parameter tuning, which we also wanted to avoid; we generally used default parameters, but also manually created salience masks, dividing each image into a low-detail region (less salient) and a high-detail region (more salient). This weakly substitutes for per-image parameter tuning. Automatic determination of salience is a presently intractable problem; the issue was gracefully avoided by the gaze-directed abstraction scheme of DeCarlo and Santella [7], in which users would look at an image for a few seconds and then an automatic abstraction would take place: portions of the image that received less attention would be more aggressively abstracted.

Non-photorealistic images have been compared with manually created artistic images since the beginning of the field, yielding

the concept of the “NPR Turing Test” [8,9]. In this thought experiment, people are presented with images and asked to guess whether they were created by humans or by computers. One outcome of the observational study of Isenberg et al. [10] was the finding that, for the algorithms tested, people generally could distinguish between images created by computer and those created by hand. The NPR Turing Test is a benchmark we have not yet passed.

Non-photorealistic rendering has had different goals, including novelty, meeting the technical challenge, and comprehensibility of the resulting images. This last point is frequently used to motivate work in the area, with medical, archaeological, and technical illustration given as application domains; perceptual studies such as those of Winnemöller et al. [11,12] have shown improvements in subjects’ ability to recognize and comprehend objects depicted in abstracted styles compared with realistic depictions. NPR methods have generally not been cast by their creators as generating artistic or emotionally laden content; a possible exception is the work of Shugrina et al. [13], where the perceived emotional state of the viewer drives the appearance of the image. Nonetheless, researchers in NPR often seek to improve the aesthetic qualities of their synthetic images. Zhao and Zhu [3] use aesthetic arguments (about the appeal of abstract paintings) to motivate their algorithm, and perform a user study measuring the ease with which subjects could recognize objects in their synthetic semi-abstract paintings.

Duke et al. [14] investigated the emotional impact of a few types of images, though in an unsystematic way and without using images that were rated for affect: the coverage of emotional space was sparse and the range of styles covered was small. Our effort here explores the impact of a range of styles, on a set of images specifically chosen to provide full coverage of a 2D emotional space. We discuss possible parameterizations of emotional space next.

2.2. Affect and emotion

We are interested in understanding participants’ emotional responses to various non-photorealistic styles, so we must first consider how emotions are described and measured. The terms *affect*, *emotion*, and *mood* are often used interchangeably; we use *affect* to describe the low-level user responses to a stimulus (e.g., palms sweating, heart racing), *emotion* to describe the cognitive interpretation of the low-level responses (e.g., fear, surprise), and *mood* to describe the longer-term state of the user as they experience emotions. Affective responses are fleeting, emotions are short-lived, and moods change slowly over time. In this paper, we will use both affect and emotion to describe participant responses to the images.

2.2.1. Representing emotion

There have been two main approaches to describe emotions: categorical and dimensional. The categorical approach applies specific and discrete labels to various emotions through semantic labels (e.g., sadness, pride, fear) [15]. The dimensional approach [16,17] proposes that emotions can be represented by two primary orthogonal axes called arousal and valence. Valence describes the pleasure (positive) or displeasure (negative) of a feeling. Arousal is related to the energy or activation of the feeling and is typically described as low (e.g., sleepiness) to high (e.g., excitement) arousal. This arousal–valence space has been described as the circumplex model of emotion [16], and has been used to describe the categorical emotion labels. For example, “anger” would be a high-arousal, low-valence state, while “depression” would be a low-arousal, low-valence state. One criticism of the dimensional model is that the arousal and valence axes are not completely

independent [17]. For example, an emotion that is truly displeasurable is unlikely to also be very relaxing. A third axis, dominance, has been added to describe the feeling in terms of how controlled, influenced, or submissive it is, as compared to controlling, influential, or dominant [18]. Together, these three dimensions have been used in emotional assessment.

2.2.2. Measuring emotion

Although there are discrete methods of measuring emotional state (e.g., semantic differential scale [18]), we focus on the dimensional approaches using the three-axis model of emotional state.

Based on their circumplex model, Russell et al. used the arousal–valence space to create the affect grid [16]. The affect grid is a tool for quick assessment of affect in terms of arousal and valence. Participants place checkmarks in the squares of a grid in response to different stimuli. Avoiding semantic labels, the self-assessment manikin [19] is a 9-point pictorial scale used to self-report arousal, valence, and dominance using a series of five images, with blank images between. As shown in Fig. 1, the SAM provides a fast, easy, and non-linguistic way of assessing emotional state along three dimensions. Although subjective approaches are most commonly used to assess user emotional state, objective methods have been developed including using facial expression analysis, physiological signal analysis, and observational analysis [18]. The studies we describe in this paper use subjective self-assessment.

2.3. Emotional response to graphics

There has been little work investigating a person's emotional response to computer-generated graphics. Hertzmann [20] motivates this issue well, pointing out that perceptual, aesthetic, and emotional content of artistic images are at least somewhat independent: “a work that is interesting is not necessarily beautiful”. For “interesting” we could substitute various other descriptive terms. There are numerous studies of the aesthetic properties of images, and perceptual studies are also common, but studies of emotional responses to NPR are rare. We previously mentioned the efforts of Duke et al. [14] in this regard. Shugrina et al. [13] used a 2D arousal–valence emotional space, but for image synthesis, not for evaluation of responses to other images. In fact, with their interactive system, detecting the user's emotional response to the image being created forges a probably undesirable feedback loop. Colton et al. [21] also generated NPR images based on emotion, but use automated facial expression analysis to determine a viewer's emotional state, which in turn determines the NPR approach to use for image generation. Mar et al. [22] investigated fundamental physiological reactions (in terms of blood oxygenation in the brain) to live-action or computer-animated agents and found that the response associated with the perception of agency was greater for the real versus the animated agents. Although they did not control the rendering of the animated agents, they used identical film clips from a movie that had been adjusted by animators to look like cartoons.

There has been some work investigating a person's emotional response to aspects of abstract art, which may inform our understanding of emotional response to stylized images. For example, Valdez and Mehrabian [23] describe a relationship between perceived valence and hue (wavelength), where they showed experimentally that blue was the most pleasant colour and that yellow was the least pleasant. They also showed that less bright and more saturated colours were more arousing. Their results were supported by Simmons in a study on the associations between colours and emotion [24]. There have also been observed relationships between shape and perceived emotion. For example, Ibanez [25] found that perceived valence corresponded with degree of symmetry in images that had their colour held constant. Mono [26] showed that circles, spirals, and shapes with smooth curves were more pleasant than shapes with hard angles. These results conform to earlier investigations into the emotion conveyed by drawings, where Hevner [27] found that curves denoted serenity, while jagged strokes and harsh angles denoted fury or agitation. In a similar vein, Halper et al. [28] found a relationship between line style and perceptions of safety—objects rendered using jagged lines were perceived as more dangerous than objects rendered using smooth lines. Line style was also related to perceptions of character strength (strong lines indicated strong characters).

3. Generating stimuli

3.1. Choosing the stimuli

To determine affective response to various NPR algorithms, we needed to choose a set of images to use as stimuli. The International Affective Picture System (IAPS) is a set of images that span emotional space. IAPS images were created for use as experimental stimuli, have been used in numerous studies, and provide normative ratings of emotion in terms of valence, arousal, and dominance for 956 colour photographs [29]. The advantage of using the IAPS photo set is that it is a standard set of images: all images have been rated for their emotional content through multiple studies and have been validated through numerous publications. The set is available to researchers by request.

We selected 18 IAPS images that spanned the arousal–valence space for use in our study. Images were chosen to represent nine specific arousal–valence locations, as shown in Fig. 2, including all combinations of low, neutral, and high arousal and valence. We were interested in whether the affective response to images of objects or scenes would be differentially affected by the various algorithms, so we chose an image at each of the nine locations that represented an object (e.g., tiger, gun) and one that represented a scene (e.g., beach, cemetery) for 18 images in total.

Due to regulations of IAPS use, we are unable to publish the images used. Fig. 2 shows a description of the image and the mean arousal, valence, and dominance of the images as provided

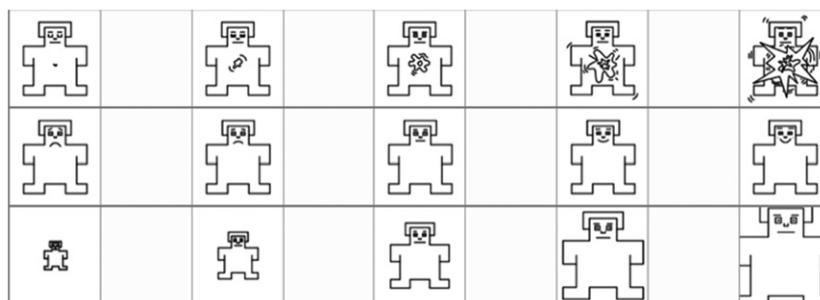


Fig. 1. The self assessment manikin 9-pt pictorial rating scale. Top: arousal; middle: valence; bottom: dominance.

by the IAPS documentation [29]. We also show schematic depictions of specific images in Section 7.

3.2. Rendering the stimuli

We rendered the stimuli in several styles using existing algorithms from the literature. We selected algorithms that were capable of operating automatically on images; our source data was in the form of images, so it would not have been possible to employ approaches that require geometry, and we wanted to avoid interactive algorithms in order to avoid the possibly confounding effects of the human user working through different interfaces.

We chose to employ three main styles: stippling, line art, and painterly rendering. We also used the generic photo abstraction method of Orzan et al., and we included two variants of blurring (uniform and salience-dependent) to provide a baseline abstraction. Stippling was one of the selected styles because of its long-standing interest to computer graphics practitioners. Line drawing and painterly rendering were chosen because of their long history and widespread use in non-photorealistic rendering. To obtain the stippled images, we used the classic stippling method of Secord [4]. The method of Kang et al. [5] was used to produce line art: among

automatic image-based methods, the results of this algorithm are unsurpassed. For the painterly algorithms, we faced a difficult decision because of the wide variety of methods and the different styles of output they generate; ultimately we decided to use two algorithms, the classic Haeberli “Paint by Numbers” [1] and the recent “Sisley” [3]. We consider Haeberli’s method to produce more representational images (closer to the original) while the output from Sisley is more abstract. Sisley also serves as a representative of a modern class of example-based techniques. Haeberli’s original system had a user guiding brushstroke placement, but we used a custom implementation of this semi-automatic system in which paint strokes were chosen from a database of six possible strokes and alpha blended onto an initially black canvas at random positions. For the blurred images, we used a Gaussian filter; the uniform blur used a radius-12 filter everywhere, while the salience-adaptive blur used radius-4 filter in the background followed by a second radius-4 filtering pass everywhere. This process ensured a clearer image in the foreground but avoided the appearance of a visible boundary between more and less strongly blurred regions.

For the Haeberli, Sisley, and blurring algorithms, we used binary salience mattes to distinguish between regions of high importance and regions of low importance. The mattes were created manually; our process was for some of the authors to paint candidate mattes for the images, then review and discuss the decisions for the mattes, and finally for one person to create a final matte for each image informed by the previous discussion. The same matte was then used for all three styles.

Example images illustrating the selected techniques are shown in Figs. 14 and 15. Fig. 3 shows the input to the algorithms: the original images and the hand-drawn salience masks. Note that these are illustrative only—the terms of use of the IAPS images do not allow them to be reproduced here. The results of the different rendering algorithms are shown in Fig. 4.

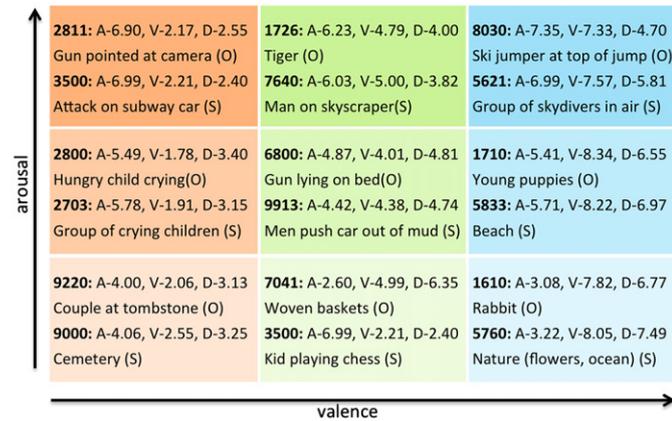


Fig. 2. The 18 IAPS image stimuli at nine different emotional locations in arousal-valence space. The vertical axis shows increasing arousal; the horizontal axis shows increasing valence.

4. Experiment

4.1. Task

The task consisted of participants rating their affective response to the images, which were rendered using the various algorithms. Participants began by completing an informed



Fig. 3. Input to the rendering algorithms. Above: original images. Below: hand-drawn salience masks.



Fig. 4. Output from the rendering algorithms. From the top: stippling (Secord); line drawing (Kang); painting (Haeberli); painting (Sisley); photo abstraction (Orzan); object blur; uniform blur.

consent, followed by the Ishihara Colour Plate Test to screen out participants who showed colour vision deficiencies. Participants then completed a training task of two neutral images (IAPS 1602,

2530) presented on a grey background. Images were presented for 10 s and participants were asked to describe the image verbally to the experimenter during this time. The description phase was followed by the rating phase; the image remained on the screen but was accompanied by the rating scales. After rating the image on all four scales, the user was presented with a grey mask and a submit button with instructions to press the button when they were ready to move on to the next trial.

Affective ratings were conducted using the self-assessment manikin 9-point pictorial scales [19]. Participants first rated the arousal of the image, followed by the valence, the dominance, and the aesthetic quality. Aesthetic quality was rated using a 9-point Likert scale. Only one rating scale appeared on the screen at any time and participants were required to provide a rating before moving on to the next rating scale (see Fig. 5).

After the training task, participants were presented with the experimental system, which was identical in appearance and function to the training system and differed only in the presented images. The 18 images described in the section on choosing the stimuli were presented in a block for each of the seven rendering techniques and the control condition (original image). The order of presentation of the seven techniques was counterbalanced using a Latin Square to avoid any effects of order of presentation. In addition, half of the participants saw the control image prior to any of the techniques, while the other half saw the control images last. We varied the position of the control images to see whether knowing the un-retouched content of the image affected user response to the various techniques.

The 18 images were presented in the same order for each technique so that the emotional content of the images did not vary too greatly from one trial to the next. We started each block with a neutral image, and ended with a relaxing image to ease participants into a new technique and leave the block with a relaxed and positive image. Beginning in the middle of arousal–valence space, participants progressed through the images in a counterclockwise spiral (see Fig. 2). Since images were repeated for each of the techniques, participants were required to describe the content of the image only for the technique that was presented first.

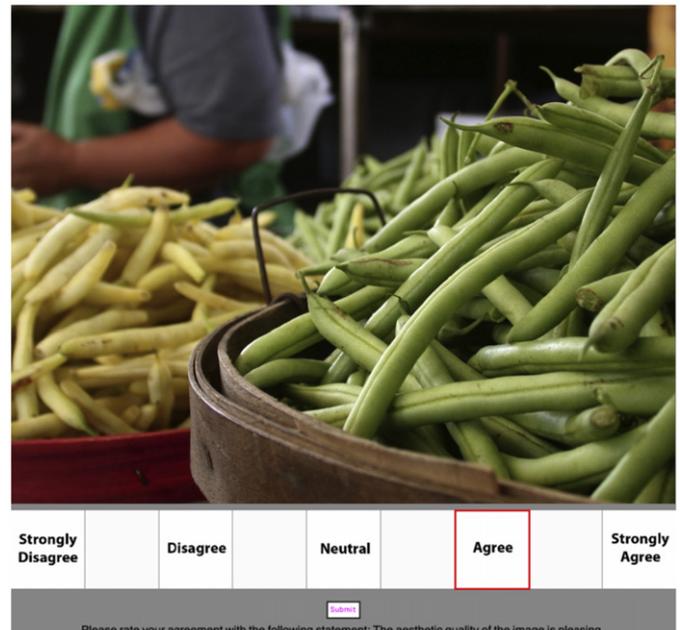


Fig. 5. Screenshot of experimental system with sample image. Participants rated the measures on a 9-pt scale by selecting the appropriate option and pressing “submit”.

After rating all images, participants completed a post-experiment questionnaire that gathered demographic information as well as preferences about the various techniques. The entire experiment took between 1 and 1.5 h to complete and participants were given \$15 to thank them for their participation. The experiment protocol was approved by the behavioural research ethics board at the University of Saskatchewan.

4.2. Apparatus

The experiment was conducted on a Windows 7 computer and a 24-in TFT display with a resolution of 1920×1200 pixels. The experimental software was written in Processing. All images were presented at a resolution of 1024×768 pixels. The system logged the information about participants, images, and ratings in a text file for subsequent analysis.

We collected gaze data for all of the participants to determine where on the images the participants were looking. Gaze data was collected using the Tobii 60XL integrated eyetracker. Gaze information was collected at 60 Hz for each eye and was analyzed using Tobii Studio Software. A gaze fixation was determined using the Tobii Fixation Filter classification algorithm that behaves like an I-VT filter that detects quick changes in the gaze point using a sliding window averaging method [30]. The velocity and duration thresholds for a single fixation were set at 35 pixels.

4.3. Participants

There were 42 participants, aged 18–33 (mean 24), of whom 21 were female. Participants all had normal or corrected-to-normal vision and did not have any colour vision deficiencies.

4.4. Data analyses

We conducted four separate Analysis of Variance tests (ANOVAs). After aggregating the ratings for all images, we conducted an overall repeated-measures MANOVA with one within-subjects factor (algorithm), one between-subjects factor (original image seen first or last) and four dependent measures (arousalRating, valenceRating, dominanceRating, aestheticRating). We also coded the images into three levels of arousal (low, neutral, and high) based on their IAPS ratings, and aggregated over these three levels rather than over all images. An RM-ANOVA with algorithm (8 levels) and imageArousal (3 levels) as within-subjects factors and original image position (2 levels) as a between-subjects factor on arousalRatings will be referred to as the ArousalANOVA. A similar process was undertaken for grouping the images according to valence and conducting a RM-ANOVA on valenceRatings (ValenceANOVA). Finally, we also aggregated separately over whether the images were objects or scenes and conducted an RM-MANOVA (ObjectMANOVA) with two within-subjects algorithms (algorithm, object versus scene), one between-subjects factor (original image position) and our four dependent measures. For all statistical tests, when the assumption of sphericity was violated, we used the Huynh–Feldt method of adjusting the degrees of freedom. Pairwise comparisons of significant results used the Bonferroni method of correcting for multiple tests with $\alpha = 0.05$.

Gaze fixation data for the first 14 participants were aggregated and visualized for each image using each rendering algorithm, resulting in 144 fixation visualizations. We used a standard heat map false-colour visualization; pixels that received greater attention were shaded red whereas pixels with less attention were shaded green, with varying colours in between. Pixels viewed less than a threshold appear black. Heat maps are an effective method of conveying the visual attention of participants viewing a stimulus [31]. Used in many applications and domains, heat maps

have helped researchers to interpret where users look on web-pages [32], while playing computer games [33], and in three-dimensional virtual environments [34]. Researchers have also used heat maps to study where users look on paintings from different periods, with varying visual complexity, and with different aesthetic appeals [35].

Our heat map was based on gaze duration, which shows the accumulated time participants spent looking at different areas of the images [31]. Because participants viewed each stimulus for a fixed duration, absolute gaze duration was used rather than relative gaze duration. The fixation heat maps were also processed to determine the percent of the pixels in the image that were viewed, detail-loss errors (pixels that were viewed in the original image, but not in the filtered versions), and distraction errors (pixels that were viewed in the filtered images that had not been viewed in the original image). We will discuss the outcomes in Section 7.

5. Results

In this section, we describe the results of our statistical tests. We summarize these results at the end of the section. Note that when we mention an interaction, it is significant ($p < 0.05$); we generally do not mention interactions that are not significant.

Do the images create the expected affective responses? To determine whether participants were responding to the processed images in a predictable manner, we looked at results from our ArousalANOVA and ValenceANOVAs. The ArousalANOVA showed a main effect of imageArousal ($F_{2,70.6} = 40.5$, $p \approx 0.000$). Bonferroni post hoc comparisons revealed that each arousal grouping was significantly different (all $p < 0.001$). The ValenceANOVA showed a main effect of imageValence ($F_{2,52.1} = 158.0$, $p \approx 0.000$). Bonferroni post hoc comparisons revealed that each valence grouping was significantly different (all $p < 0.001$). Thus, the images were producing consistent and predictable affective responses in the participants (see Fig. 6).

Are there overall differences in the affective responses to the various NPR algorithms? Using our overall RM-MANOVA, we found a significant effect of algorithm on arousalRating ($F_{4.8,194.7} = 12.3$, $p \approx 0.000$), valenceRating ($F_{7,280} = 6.2$, $p \approx 0.000$) and dominanceRating ($F_{4.1,166.6} = 2.6$, $p = 0.036$). For arousalRating, the original image was rated as more arousing than images produced with all other algorithms except Orzan. Also, the Orzan images were rated as more arousing than images produced using the Haeberli, Secord, and Sisley algorithms and images produced using the blur and object blur algorithms (see Fig. 7). For valenceRating, the images produced using the Kang and Secord algorithms were rated with lower valence (more negatively) than the original images and those produced using the Orzan or Sisley algorithms. In addition, images produced using blur were rated

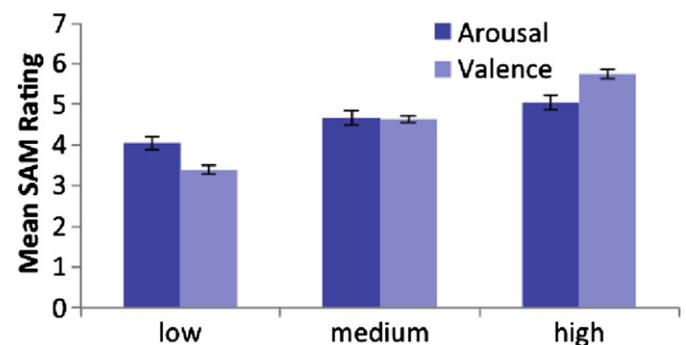


Fig. 6. Means \pm SE for arousal and valence ratings by the arousal or valence of the image stimuli.

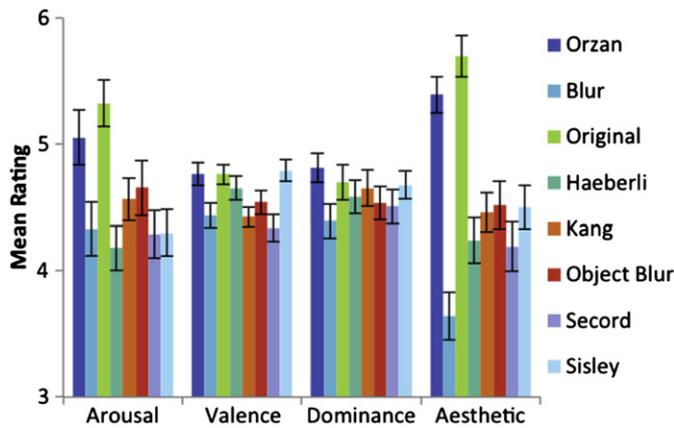


Fig. 7. Overall means \pm SE for arousal, valence, dominance, and aesthetic quality ratings.

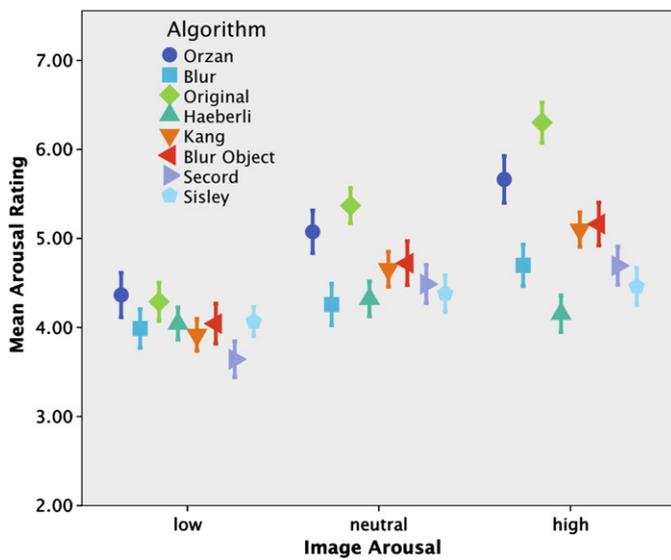


Fig. 8. Means \pm SE for arousal ratings by the arousal of the stimulus image separated by algorithm.

less valent than those produced using Sisley (see Fig. 7). For dominance, the images produced using blur were rated as less dominant than those produced using Orzan or Kang (see Fig. 7).

Were these differences affected by whether the participants saw the original image condition first or last? Our RM-MANOVA showed that there were no main effects of whether the original image was seen first or last or interactions of original image order and algorithm on any of the three affective measures (all $p > 0.05$).

Did the affective ratings for the different algorithms change depending on the arousal level of the original image? In addition to showing that the imageArousal was yielding consistent arousalRatings, our ArousalANOVA also showed that there was a significant interaction between imageArousal and algorithm on the arousal ratings ($F_{4,3,489.4} = 11.8, p \approx 0.000$). As can be seen in Fig. 8, for low-arousal images, there was not a large difference in the arousal ratings for images from different algorithms. For neutral-arousal images, the differences became larger, and for high-arousal images, the arousalRating differences were largest. Some of the relationships visible in the figure are statistically significant differences, but many are not. For low-arousal images, there were no significant differences between the algorithms. For neutral-arousal images, the original images were more arousing than images produced using all other algorithms except Orzan;

Orzan images were more arousing than Haeberli, Sisley and blur images. For high-arousal images, the original images were more arousing than those produced with all other algorithms except Orzan, the Orzan images were more arousing than those from all other algorithms (but not the originals), the Haeberli algorithm's results were less arousing than all except blur, and the object Blurred images were more arousing than the Sisley and blur images. Other interactions were not statistically significant.

Did the affective ratings for the different algorithms change depending on the valence of the original image? Our ValenceANOVA showed that there was a significant interaction between image-Valence and algorithm on the valence ratings ($F_{11,0, 440.1} = 31.1, p \approx 0.000$). As seen in Fig. 9, for neutral-valence images, the differences between algorithms were small, but for low- and high-valence images, the differences between the algorithms were bigger. Details of statistically significant results are as follows: for low-valence images, the Haeberli and Sisley images were rated as more valent than all others except the blurred images, and the original images were rated as less valent than those of all other algorithms except Orzan. In addition, blur was more valent than object blur. For medium-valence images, the only differences were that blur was rated as less valent than Orzan or Secord. For high-valence images, the original image was rated as more valent than the images from all algorithms, Orzan images were more valent than all but the original image, object blurred images were more valent than those of all remaining algorithms except Sisley, and the Secord images were rated less valent than the images from any algorithms except blur and Haeberli. Other interactions did not reach statistical significance.

How did the algorithms fare in terms of aesthetic ratings? The RM-MANOVA described previously showed a main effect of algorithm on aestheticRating ($F_{5,3,211.5} = 23.8, p \approx 0.000$). The original images and the Orzan images were rated as having a higher aesthetic quality than images produced using all other algorithms. In addition, the blurred images were rated with lower aesthetic quality than all other images except those from the Secord algorithm (Haeberli images were marginal at $p = 0.058$).

Unlike the affective measures, the aesthetic ratings did change depending on whether participants saw the original images first or last. There was a main effect of original image order on aesthetic rating ($F_{1,40} = 7.2, p = 0.010$). Participants who saw the original images first tended to rate the aesthetic quality of all images lower

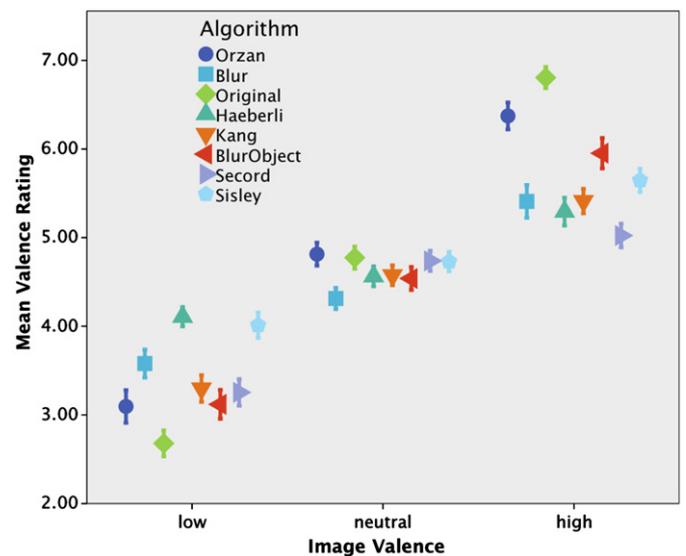


Fig. 9. Means \pm SE for valence ratings by the valence of the stimulus image separated by algorithm.

on average (mean=4.3, SE=0.17) than participants who saw the original image last (mean=4.9, SE=0.17). This main effect needs to be interpreted in light of a significant interaction of whether the original image was seen first or last and algorithm on aestheticRating ($F_{5,3,211.5}=3.2, p=0.011$). This interaction shows that although the participants who saw the control image first rated the images as having lower aesthetic quality for all algorithms, this difference was only significant for the Orzan, object Blur, and original images.

In addition, we also asked participants in a post-experiment questionnaire to choose their favourite and least-favourite algorithm. Participants overwhelmingly preferred the Orzan images (24/42 responses) with Sisley coming in a distant second choice (7/42). For least favourite, participants chose the blurred images most often (26/42 responses) with Secord coming in second (5/42).

Is there an overall difference in the affective ratings for images of objects and images of scenes? Our ObjectMANOVA revealed that there was a main effect of object or scene on valence ratings ($F_{1,40}=16.4, p \approx 0.000$). In general, participants rated the valence of objects lower (mean=4.5, SE=0.03) than the valence of scenes (mean=4.7, SE=0.03). This is not surprising as the IAPS-provided valence of the object images (mean=4.8, SE=0.04) was slightly lower than that of the scene images (mean=5.0, SE=0.04). There were no main effects of object or scene on any of the other measures (all $p > 0.05$). There was also an interaction effect of object or scene and algorithm on valence ratings ($F_{6,1,242.6}=16.4, p \approx 0.000$), meaning that the rating differences between the object and scene images were not consistent across all algorithms. Post hoc comparisons showed that there was a significant increase in the valence ratings of scenes over objects for all algorithms (all $p < 0.005$) except for Kang and Secord, where the valence ratings for scenes were lower (see Fig. 10). There were no interaction effects of object or scene and algorithm on arousal ratings or dominance ratings (all $p > 0.05$).

Summary of the results. Our results can be summarized into the following nine takeaway messages:

(1) The image choices were producing significant and predictable differences in the affective ratings, showing that our experimental stimuli were effective;

(2) Applying any of the algorithms except Orzan created less arousing images than the original, and the Orzan algorithm created more arousing images than all of the algorithms except Kang;

(3) The differences in arousal ratings between the algorithms became more apparent as the image itself was more arousing. There were no differences in the algorithms for low-arousal images, small differences for medium-arousal images, and large differences for high-arousal images;

(4) Images produced using Kang and Secord were less valent than images produced using Orzan, Sisley, or the original image;

(5) The differences in valence ratings between the algorithms were larger for low-valence and high-valence images than neutrally valent images;

(6) For Kang and Secord, the valence ratings were higher for objects over scenes, whereas the opposite was true for all other algorithms;

(7) The difference in arousal and valence ratings did not depend on whether participants saw the control images first or last;

(8) Participants preferred the Orzan and original images over images created by all other algorithms and least liked the blurred images; and

(9) The differences observed in all results cannot be solely attributed to information loss, as shown by our blurred image and blurred object algorithms.

6. Discussion

Our most general finding was that the rendered images produced flattened affect, as compared with the original images: arousal was reduced and valence was brought closer to neutral. This result was consistent across all the algorithms we tested, although some algorithms more strongly exhibited this flattening pattern. It is tempting to attribute this outcome to a failure on the part of the participants to recognize the content of the stylized images due to information loss; however, there are two reasons why information loss does not fully explain our results. First, were this explanation correct, we would have observed an effect of order of presentation of the original image on the affective ratings—participants who saw the original image first would have exhibited an improved ability to interpret stylized images. The absence of an order effect for any of our affective measures is evidence against this explanation. Second, we included the blurred images and blurred background images to specifically test the possibility that our results could be attributed to information loss. There was no consistent pattern where responses to the stylized images followed responses to the blurred images. In fact, affective responses to some algorithms more closely mirrored responses to the original image (e.g., Orzan). Some of the observed affective dampening can be attributed to information loss in the stylized images, as we will see in the next section; however, this is not the sole explanation and there are other factors that must be considered.

It might not be too surprising that synthetic non-photorealistic images do not have much emotional content. NPR has long been viewed as a scientific endeavour and technical challenge, and researchers have not often explicitly sought to induce emotional responses with their images. Nonetheless, work in the field has occasionally been motivated by the idea of creating more emotionally charged images (e.g., the recent work of Lopez-Moreno et al. [36]). The experimental data indicate that for a broad range of existing techniques, synthetic images have less emotional impact than the photographs from which they were derived. This points to an open problem for practitioners to address.

Among the algorithms tested, the photo abstraction method of Orzan et al. produced affective responses most similar to the original. Unlike the other methods which hid or removed most detail, this method preserved details, including colour gradients, region boundaries, and some high-frequency features; we speculate that the inclusion of a few details in addition to the large-scale content was responsible for this algorithm's success at evoking affective responses from the viewers. Although there is no advice in the literature on how the level of detail in images influences affective response, it stands to reason that preserving details will aid in preserving the emotional impact of the images.

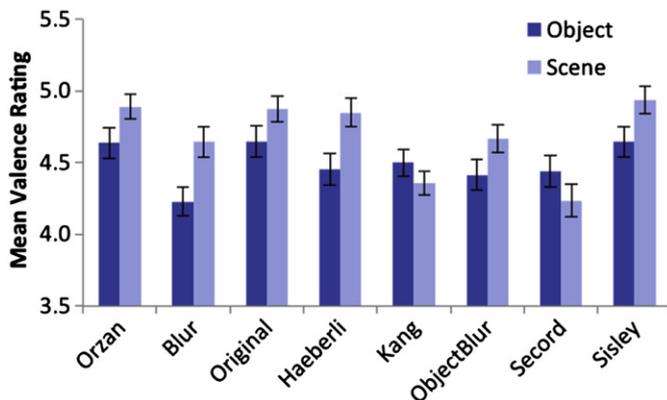


Fig. 10. Means \pm SE for valence ratings by algorithm and whether the stimulus image was an object or a scene.

There is a difference between the responses to “object” images (those concentrating on a distinct object or person) compared with the responses to “scene” images (those where large parts of the image are needed to establish context). For almost all rendering styles, and for the original images, scenes produced higher valence ratings than objects. This is not unexpected as the original images of scenes were rated with a slightly higher average valence in the IAPS database than the images of objects, and this difference may have simply carried over to the stylized images. It could also be that the loss of detail has greater impact on affective response to objects than scenes. For example, an image of a bunny might be more affected by the loss of detail in the algorithms than an image of a beach scene, where the general idea and tone of the image can be conveyed with much less detail. The differential affective response to stylized images of objects and scenes warrants future research, including questioning why this valence difference between objects and scenes was reversed in the case of the line drawing and stippling styles.

The line drawing and stippling styles are quite different: line drawing shows edges and largely preserves high-frequency details, while the stippling method indicates tone and better preserves low-frequency content. Nonetheless, both methods produced similar responses overall. We might attribute this to lack of colour, that being the main commonality between the two styles. The colour hypothesis provides a speculative explanation for the differing valences of objects and scenes for stippled and line images compared with the rest: more so than objects, scenes can be characterized and some content conveyed by large-scale use of recognizable colours, e.g., blue skies and oceans, grey clouds, or green vegetation.

Although we may have expected similar responses to the images that were blurred overall and those where the background was blurred more than the primary object, there is considerable difference between the uniformly blurred images and those informed by the mask. In terms of the aesthetic judgement, while the uniformly blurred style was by far the least liked, the object-blur style was competitive with the other sophisticated rendering algorithms, apart from Orzan et al.’s photo abstraction. Object blur also provided images with marginally higher arousal than most styles, as compared with uniform blur which yielded the least arousal. Altogether, this provides some support for the commonly held belief that image abstraction should pay attention to the content: less important content can be more abstracted than the more important content. Speculatively, the specific content may matter less than the simple fact of choosing some coherent subject for the image and portraying it more prominently than the background.

As a minor observation, we point out that the Sisley images did not have higher valence than the images produced by other algorithms. We had expected that the colour saturation shift employed in this algorithm would have an effect: the brighter colours would have seemed more cheerful to the participants, perhaps manifesting as an increase in reported valence for the neutral images. However, no such effect was detected. Fig. 7 shows that the mean valence ratings for Sisley images was comparable to Orzan and the original images; however, Fig. 9 further shows how the valence benefits of Sisley images were seen mainly for the low-valence images. Like Haerberli images, Sisley images showed considerable valence dampening towards neutral ratings (i.e., higher valence ratings for low-valence images and lower valence ratings for high-valence images).

7. Analysis of Gaze

As mentioned in Section 4.4, we aggregated the participants’ gaze data to produce a single heat map per stimulus. Because the heat maps show where the participants were looking as they

formed their impressions of the images, we hoped that an investigation of the heat maps would be able to explain some of the findings from the subjective questionnaires. Note that the heat maps display the gaze fixation of participants by estimating the foveated area. Through their peripheral vision, participants would still have seen image regions that were not fixated on.

Some quantitative data appears in Fig. 11. The top portion shows the percentage of pixels viewed, averaged across all images; notice that fewer pixels in the original images were viewed than in the filtered images. This can be a sign of confusion, discussed in greater detail below.

The lower portion of Fig. 11 summarizes mismatches between the heat maps from the original and filtered images, averaged over all image types. We identify a “detail loss” discrepancy as a pixel that was viewed in the original image but not viewed in the filtered image, and a “distraction” discrepancy as a pixel that was not viewed in the original but was viewed in the filtered image. We call an image element a distraction if it drew the viewers’ attention to an irrelevant part of the image; in fact, a more common cause of the discrepancies shown in Fig. 11 was “confusion”, characterized by a general tendency of the viewers to look at large portions of the image. We have used the term *confusion* because we believe the activity represents the viewers’ efforts to understand and interpret an unclear image.

The error bars of Fig. 11 mark one standard deviation above and below the mean. Their size conveys the enormous scale of the variation among images; this variation makes it very difficult to draw conclusions from the aggregate data. However, the individual images remain a rich source of information that we can discuss qualitatively. In the remainder of this section, we describe

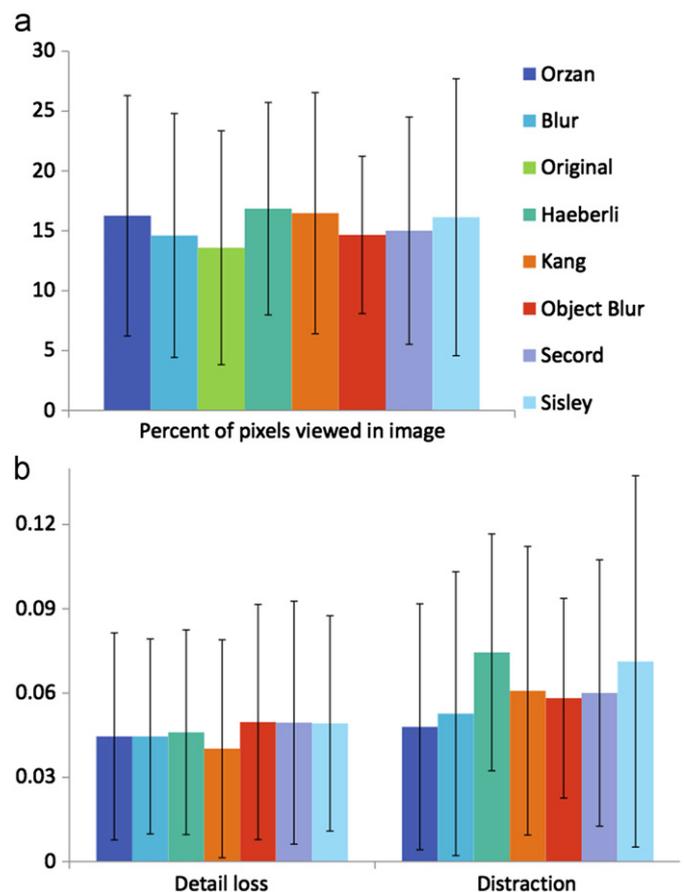


Fig. 11. Summary of gaze data per algorithm. (a) Average proportion of pixels viewed. (b) Average discrepancy between original and filtered images.

some of the mechanisms by which the emotional impact of the images was reduced.

The terms of use of the IAPS images do not permit their reproduction. However, our discussion of specific images will be easier to follow if the reader has some sense of the image composition. Accordingly, we created schematic versions of the images in question: major elements were indicated and labeled, and the resulting map of the image was overlaid on a massively blurred version of the original image.

7.1. Confusing images

Many images became *confusing*: the images or significant portions thereof became indiscernible, requiring the subjects to visually investigate in an effort to unravel the puzzle. We infer confusion from the heat maps where the subjects' gazes spread over a substantially larger portion of the image than was the case in subjects looking at the unfiltered image.

A mild example of confusion is illustrated in Fig. 12. The original image depicted a pair of mourners in a cemetery. The photo abstraction version of the image provoked specific examples of confusion: a fixation appeared on the background tombstone, easily dismissed when looking at the regular image; the flowers are indistinct in the abstraction and now demand some attention to identify them; and the facial expressions of the mourners are difficult to parse, leading to a much more intense study of this area than in the original.

Fig. 13 shows more widespread confusion. When looking at the original image, subjects concentrated mainly on the human figures: the two men pushing at the truck and the onlooker behind it. The perception of the ETF-filtered image is much less clear: while the two men still command a fair amount of attention, there is considerable effort put towards scanning the truck and even its shadow on the ground below. The wheel, of

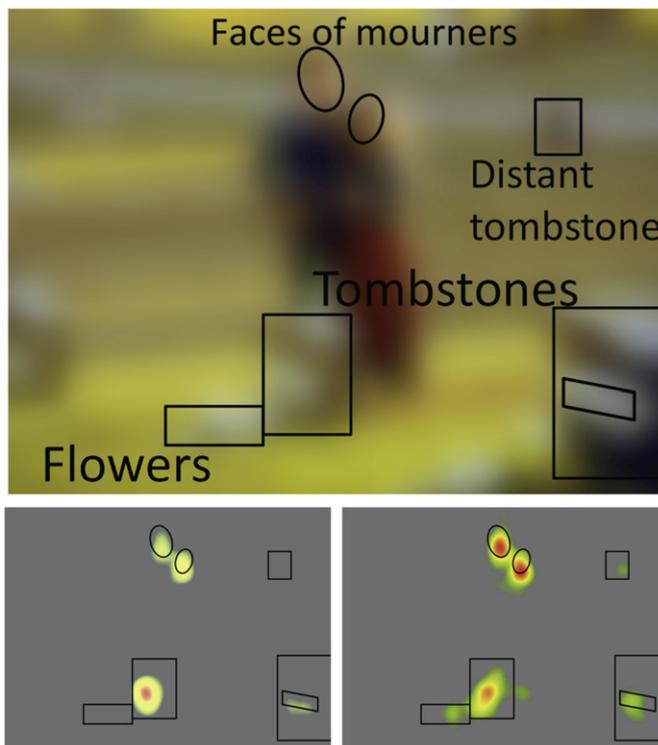


Fig. 12. Results for photo abstracted cemetery. Above: schematic; below left: heat map from viewing original image; below right: heat map from viewing filtered image.

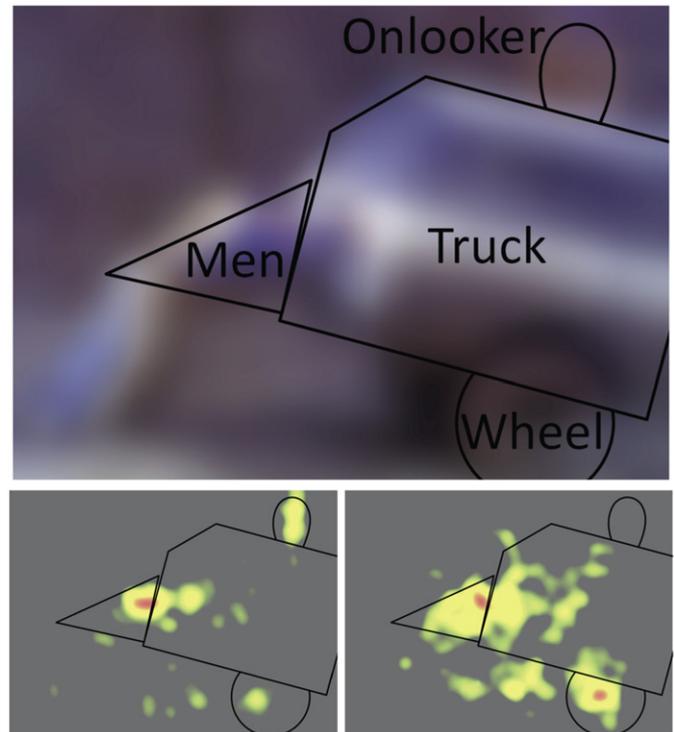


Fig. 13. Truck filtered with line drawing. Above: schematic; below left: heat map from viewing original image; below right: heat map from viewing filtered image.

little interest in the original, is a focus of attention. Finally, there is no fixation at all on the onlooker, suggesting that the subjects were not even able to identify this region of the image as containing a human figure.

7.2. Distracting images

As a special case of confusion, we mention *distraction*, in which a formerly non-salient portion of the image becomes interesting because of artifacts introduced by the filtering process. That is, it is not that the filter enhanced structures already present to bring them above a salience threshold; rather, the filter introduced detail not formerly present, which then was perceived as interesting. Note that we do not suggest this is an altogether undesirable feature in a non-photorealistic image filter; rather, we claim that the abstract structures thus introduced are likely to diminish the emotional impact of the original semantically meaningful image content.

We illustrate distraction with the examples in Figs. 14 and 15. Fig. 14 contains a cemetery image, with elements of confusion as well as spurious fixations. Fig. 15 shows an image of a gun lying on a sheet; the wide distribution of fixations suggests confusion, but confusion is contraindicated by the simplicity of the image. Many minor fixations appear in regions which superficially appear entirely empty. We therefore advance the explanation that the subjects are examining the complex high-frequency texture introduced by the Sisley filter.

Details of the Haeberli-filtered cemetery image are shown in Fig. 16. There is considerable confusion over the mourners—even more pronounced than in the Orzan-filtered image, as if the subjects are having difficulty even identifying the faces. In addition, we see two minor fixations. While these locations in the original image have little to draw the eye, the filtering introduced prominent, moderately high-contrast individual strokes to these spots. This introduced detail is spurious, not based on any feature of the original image.

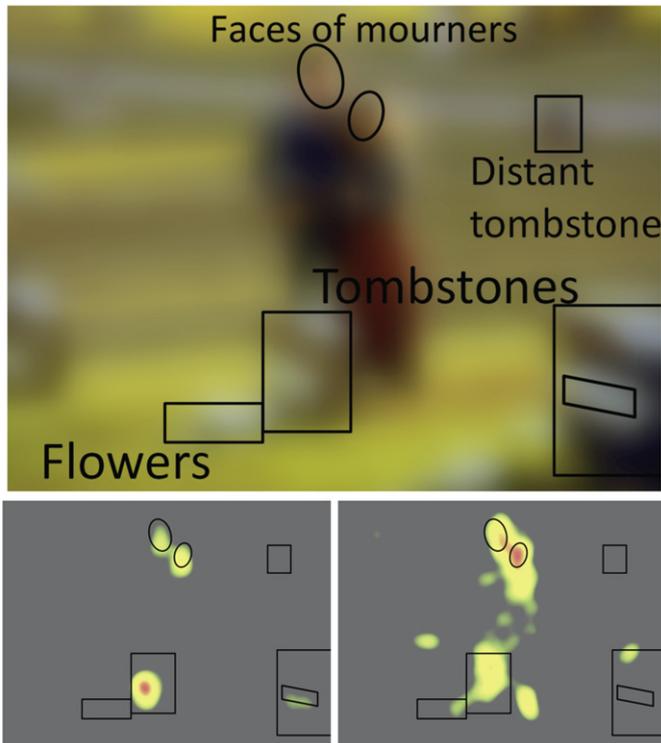


Fig. 14. Results for cemetery in Haerberli style. Above: schematic; below left: heat map from viewing original image; below right: heat map from viewing filtered image.

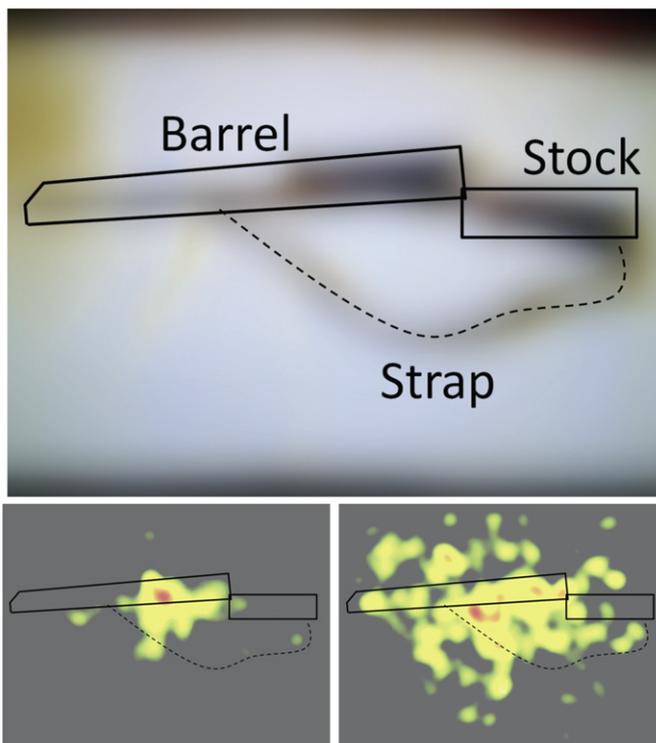


Fig. 15. Results for gun filtered with Sisley abstract painting. Above: schematic; below left: heat map from viewing original image; below right: heat map from viewing filtered image.

Details of a portion of the Sisley-filtered gun image are shown in Fig. 17. From the heat map we can see that this image largely lacks specific fixations, which we might identify as confusion. However, the original image is quite high-contrast and its subject

readily identifiable in all filtered images. In the heat map of the Sisley-filtered image, some fixations do correspond to features of the image, but others do not. We show a specific example of an area that was visually investigated by the participants: it was empty before filtering, and is largely still empty, but now contains irregular high-frequency texture and subtle shading effects. Note that the detail shown is not unique either in containing this kind of texture or in having been studied by the participants. Because of the lack of features in these areas, and the unlikelihood of confusion, we suggest that the details introduced by the Sisley filter caused distraction in the viewers.

7.3. Loss of semantic detail

We inferred confusion and distraction when subjects looked at non-salient portions of the image. The opposite kind of discrepancy can also appear, when subjects looking at the filtered images ignore formerly salient parts of the image. We saw an example of this in the ETF-filtered truck, where the onlooker was ignored; however, we do not believe the onlooker was an important contributor to the emotional content of the image. We next discuss two examples where crucial semantic details have been lost, illustrated in Figs. 18 and 19.

In Fig. 18, the original image depicted a skier about to make a steep descent towards a ski jump with the camera positioned behind the shoulder of the skier. The jump is vital to understanding the scene, but in the Haerberli-filtered version, it is largely ignored. Even the fixation on the skier is weakened; the focus of attention is now the orange tent, which, while visually striking by virtue of its bright colour, is not important to the imminent action.

For Fig. 19 the original image showed a group of crying children. There are five major points of fixation for the original image: four of the children's faces, and an outstretched pair of hands. Subjectively, our assessment is that the face of the right-most child is the most expressive; however, the Sisley version loses this, the more distant face, and the supplicating hands as well. Also, the facial expressions have been lost even though they are still recognizable as faces and hence loci of attention. Refer to the faces in the stylized images in Fig. 4 for a sense of how well facial expressions are preserved.

7.4. Uninteresting images

Finally, we discuss a general problem where subjects did not examine the filtered images as closely as the originals. We call such images "uninteresting". While there may not be a particular loss of semantic detail, and subjects' ability to interpret the images may be largely unimpaired, their engagement with the image has diminished. Uninterestingness is characterized in the heat maps by a substantial reduction in the portion of the image inspected.

Two examples where the filtered image has become uninteresting are shown in Figs. 20 and 21. Fig. 20 shows a beach scene; the ETF-filtered version produces the same general fixation pattern, but the degree of fixation is less, especially in the cliff region. Fig. 21 shows a nature scene; the version of the image treated with the photo abstraction filter elicited much less interest overall, with increased attention to the far shore of the lake.

One aspect of the beach image is the shift of focus from the lower right, in the original, to the lower middle-right, in the filtered image. Fig. 22 shows the details of these fixations: attention has shifted from the beachgoer to the prominent but more abstract beach umbrella. The abstraction of the man is less discernible, though possibly still identifiable in context.

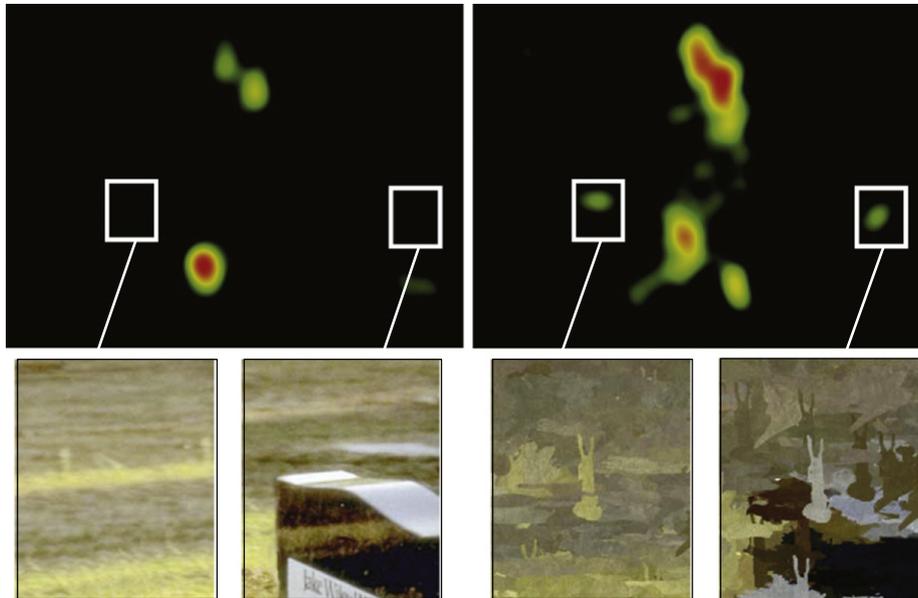


Fig. 16. Distractions in the Haeberli cemetery image. Left: original; right: filtered.

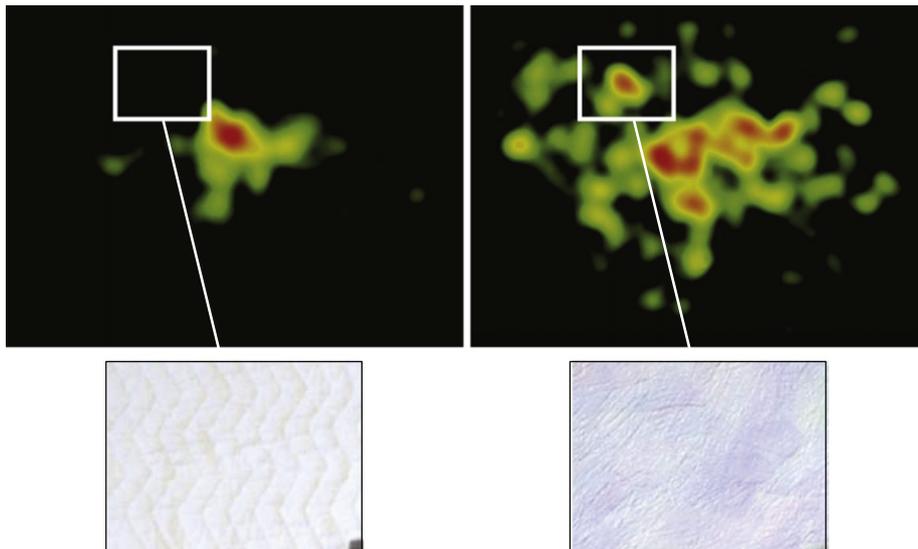


Fig. 17. Distractions in the Sisley gun image. Left: original; right: filtered.

Loss of semantic information might be considered a special case of images having become “uninteresting”. However, we reserve the diagnosis of uninterestingness for cases in which there is no identifiable semantic element that has vanished.

7.5. Summary

The heat maps provide some insight into why the emotional impact of the original images was weakened by the filters. The filters often produced confusing images, which demanded attention to interpret. In some cases, semantic information was lost completely, making the images more difficult or impossible to appreciate. Distractors were sometimes introduced. Finally, images sometimes became less interesting, perhaps due to overall loss of nonspecific details.

This list of sources of discrepancy may be instructive for designers of future non-photorealistic image filters. However, we want to caution against interpreting the implied advice too broadly. This paper examined the emotional impact of images, but images

can have aesthetic appeal or be used to communicate independently from the emotional content of their subject matter. As a specific example, we can cite the Sisley image of the gun: while the abstraction process muted the emotional impact of the scene, we can infer that it increased the viewers’ interest in the image.

8. Conclusions

This paper investigated emotional responses to computer-generated non-photorealistic images. We conducted a 42-subject study measuring valence, arousal, dominance, and aesthetics over a set of 18 images shown in eight different styles: five existing image-based algorithms were used, plus two variants of blurring, plus the original photographic image. The 18 original images were from the IAPS dataset and had been rated for affective content; our participants’ responses were consistent with the initial rating.

We found that the use of NPR algorithms significantly affected participants’ reported experiences of valence and arousal. Across

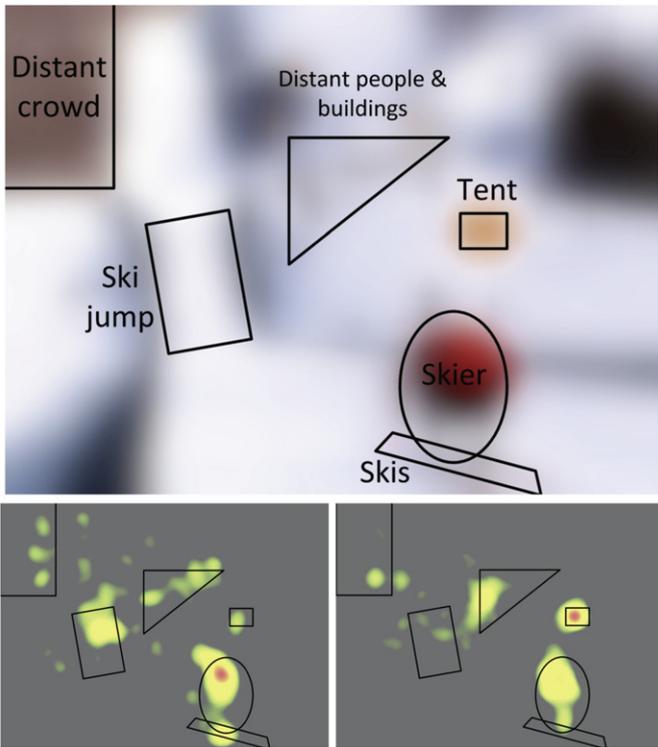


Fig. 18. Skier filtered in Haerberli style. Above: schematic; below left: heat map from viewing original image; below right: heat map from viewing filtered image.

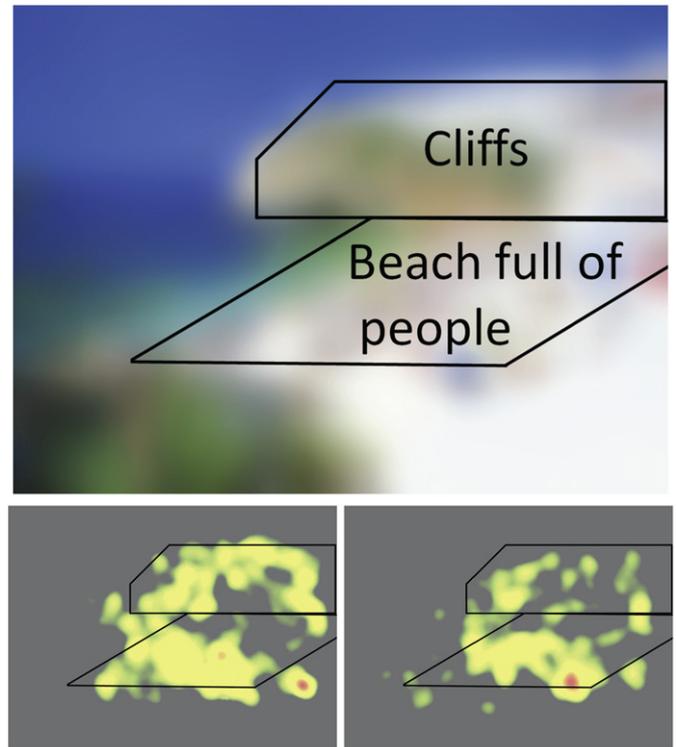


Fig. 20. Results for beach in line drawing style. Above: schematic; below left: heat map from viewing original image; below right: heat map from viewing filtered image.

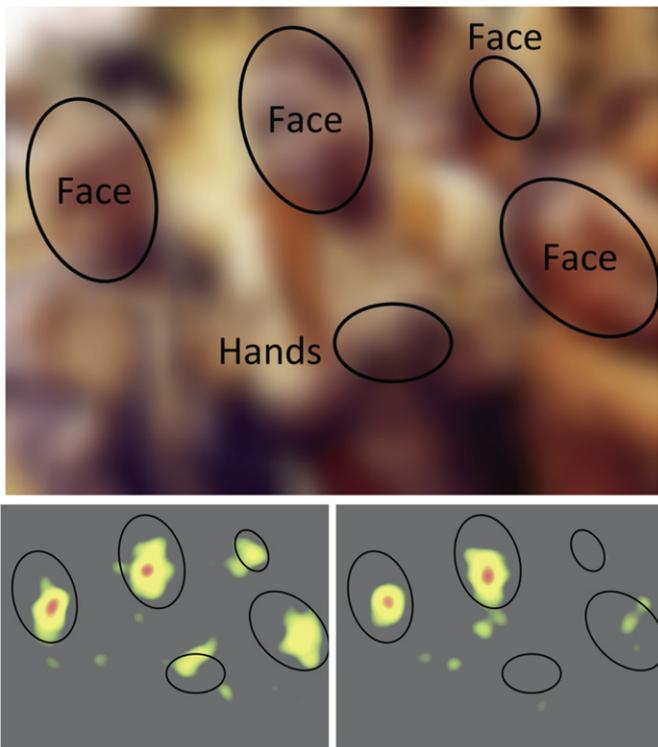


Fig. 19. Results for crying children in Sisley style. Above: schematic; below left: heat map from viewing original image; below right: heat map from viewing filtered image.

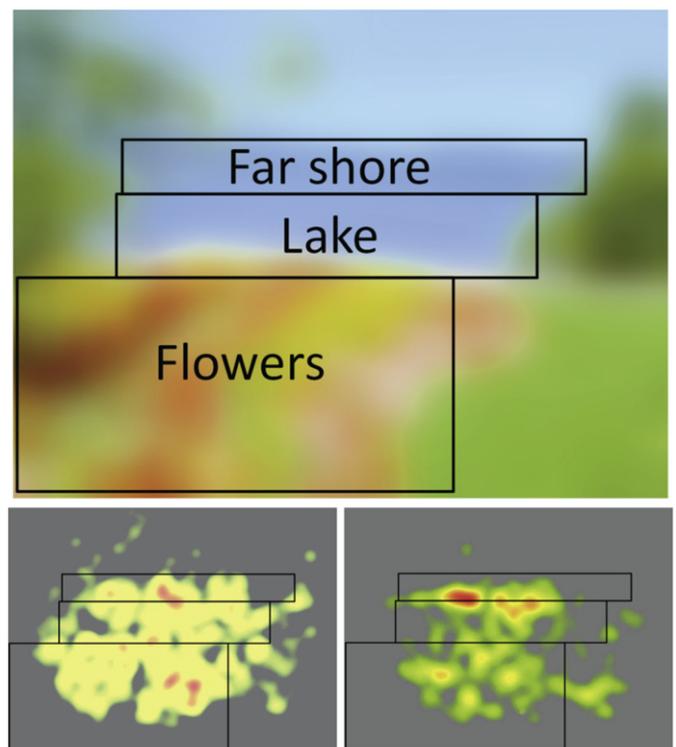


Fig. 21. Results for nature scene with Orzan abstraction. Above: schematic; below left: heat map from viewing original image; below right: heat map from viewing filtered image.

all algorithms, emotional responses were muted, being shifted from more strongly felt emotions towards a neutral state. Nonetheless, the emotional responses were never suppressed entirely, nor are the reduced intensities of emotions attributable to loss of

detail in the rendered images. Among the algorithms investigated, the photo abstraction of Orzan et al. best preserved emotional responses, while the painterly algorithms exhibited the most

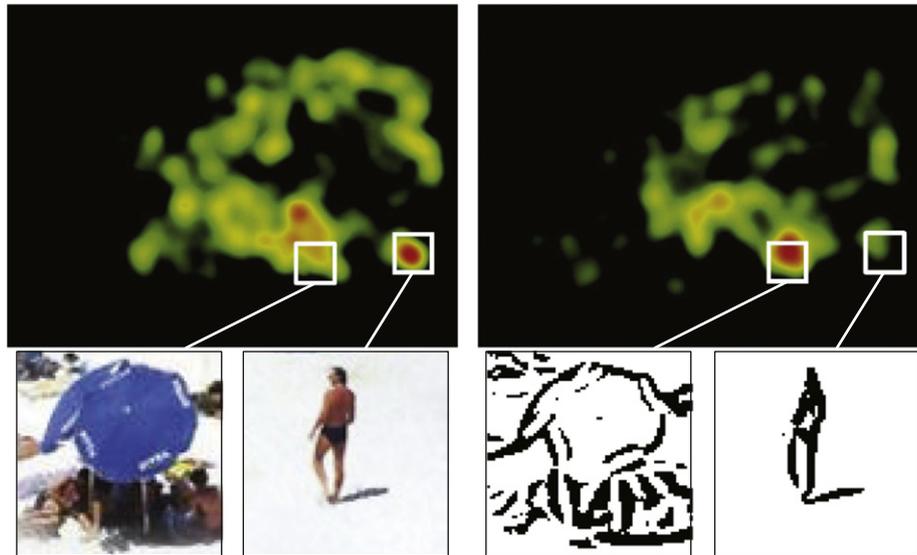


Fig. 22. Above: heat map for original beach scene; heat map for line drawing of beach scene. Below: details of fixation points; originals, left, and line drawings, right.

dampening. Heat maps of the participants' gaze locations revealed that the filtered images sometimes contained confusing or distracting elements, or lost detail making them difficult to properly interpret. We hope that our results will provoke further investigation of emotional responses to NPR images, and that they will inspire researchers in the NPR community to devise techniques that can retain or even amplify the emotional content of the input.

8.1. Future work

Our results are the first to show that people's emotional reactions to stylized images change with the use of different NPR algorithms. These results open a number of research opportunities in this space.

Subjective evaluation is a good approach for understanding participants' attitudes and opinions and provided significant and consistent results in our study. Prior research has shown that emotional responses to pictorial stimuli from the International Affective Picture System can also be measured via objective physiological response [37]. One main advantage of physiological measures of emotional response is that the affective response is accessed directly and not mediated by cognitive processes. We plan to add the objective measurement of emotional reaction via physiological measures to determine if there are low-level responses to the various NPR algorithms.

Our study examined reactions to still images; however, NPR algorithms have been used in emotionally rich animated media such as computer games, films, and advertisements. We plan to extend our work to examine emotional response to animated clips of stylized images, and to consider more ecologically valid stimuli such as dramatic or narrative film clips that include sound. Even within still images, more testing remains to be done. We have by no means exhausted the range of existing NPR algorithms, and emotional responses to other algorithms could be studied; in particular, we would include the Empathic Painting system [13] in any future study. We would also want to conduct a deeper investigation of example-based techniques to see whether more sophisticated methods, e.g., incorporating semantic information [38] can better preserve emotional content. In general, though, we fear that emphasizing artifacts of the artistic media, such as definite strokes, will still distract the viewer.

Finally, we compared emotional responses to non-photorealistic images against emotional responses to photographic images.

It would be worthwhile to perform an experiment comparing computer-generated artistic images with hand-drawn artistic images.

Acknowledgements

This work received financial support from NSERC and from GRAND. We would like to thank Brett Taylor, Brett Watson, Emma Cey and all of our participants. We would also like to thank the authors of the algorithms used in the study for making their implementations available.

References

- [1] Haeberli P. Paint by numbers: abstract image representations. In: Proceedings of the 17th annual conference on computer graphics and interactive techniques. SIGGRAPH '90. New York, NY, USA: ACM; 1990. p. 207–14.
- [2] Orzan A, Bousseau A, Barla P, Thollot J. Structure-preserving manipulation of photographs. In: Proceedings of the 5th international symposium on non-photorealistic animation and rendering. NPAR '07. New York, NY, USA: ACM; 2007. p. 103–10.
- [3] Zhao M, Zhu SC. Sisley the abstract painter. In: Proceedings of the 8th international symposium on non-photorealistic animation and rendering. NPAR '10. New York, NY, USA: ACM; 2010. p. 99–107.
- [4] Secord A. Weighted voronoi stippling. In: Proceedings of the 2nd international symposium on non-photorealistic animation and rendering. NPAR '02. New York, NY, USA: ACM; 2002. p. 37–43.
- [5] Kang H, Lee S, Chui CK. Coherent line drawing. In: Proceedings of the 5th international symposium on non-photorealistic animation and rendering. NPAR '07. New York, NY, USA: ACM; 2007. p. 43–50.
- [6] Mandryk RL, Mould D, Li H. Evaluation of emotional response to non-photorealistic images. In: Proceedings of the 9th international symposium on non-photorealistic animation and rendering. NPAR '11. ACM; 2011. p. 7–16.
- [7] DeCarlo D, Santella A. Stylization and abstraction of photographs. *ACM Trans Graph* 2002;21:769–776.
- [8] Strothotte T, Schlechtweg S. *Non-photorealistic computer graphics*. Morgan Kaufman Publishers; 2002.
- [9] Gooch AA, Long J, Ji L, Estey A, Gooch BS. Viewing progress in non-photorealistic rendering through heinlein's lens. In: Proceedings of the 8th international symposium on non-photorealistic animation and rendering. NPAR '10. New York, NY, USA: ACM; 2010. p. 165–71. ISBN 978-1-4503-0125-1.
- [10] Isenberg T, Neumann P, Carpendale S, Sousa MC, Jorge JA. Non-photorealistic rendering in context: an observational study. In: Proceedings of the 4th international symposium on non-photorealistic animation and rendering. NPAR '06. New York, NY, USA: ACM; 2006. p. 115–26.
- [11] Winnemöller H, Feng D, Gooch B, Suzuki S. Using NPR to evaluate perceptual shape cues in dynamic environments. In: Proceedings of the 5th international symposium on non-photorealistic animation and rendering. NPAR '07. New York, NY, USA: ACM; 2007. p. 85–92.

- [12] Winnemöller H, Olsen SC, Gooch B. Real-time video abstraction. *ACM Trans Graph* 2006;25:1221–1226.
- [13] Shugrina M, Betke M, Collomosse J. Empathic painting: interactive stylization through observed emotional state. In: *Proceedings of the 4th international symposium on non-photorealistic animation and rendering*. NPAR '06. New York, NY, USA: ACM; 2006. p. 87–96.
- [14] Duke D, Barnard P, Halper N, Mellin M. Rendering and affect. *Comput Graphics Forum* 2003;22(3):359–368.
- [15] Ekman P. Basic emotions. John Wiley & Sons, Ltd; 2005 p. 45–60. (Chapter 3).
- [16] Russell JA, Weiss A, Mendelsohn GA. Affect grid: a single-item scale of pleasure and arousal. *J Pers Soc Psychol* 1989;57(3):493–502.
- [17] Lang PJ. The emotion probe, studies of motivation and attention. *Am Psychol* 1995;50(5):372–385.
- [18] Coan JA, Allen JJB. *Handbook of emotion elicitation and assessment*. Oxford University Press; 2007.
- [19] Bradley MM, Lang PJ. Measuring emotion: the self-assessment manikin and the semantic differential. *J Behav Therapy Exp Psychiatry* 1994;25(1):49–59.
- [20] Hertzmann A. Non-photorealistic rendering and the science of art. In: *Proceedings of the 8th international symposium on non-photorealistic animation and rendering*. NPAR '10. New York, NY, USA: ACM; 2010. p. 147–57.
- [21] Colton S, Valstar MF, Pantic M. Emotionally aware automated portrait painting. In: *Proceedings of the 3rd international conference on digital interactive media in entertainment and arts*. DIMEA '08. New York, NY, USA: ACM; 2008. p. 304–11.
- [22] Mar RA, Kelley WM, Heatherton TF, Macrae CN. Detecting agency from the biological motion of veridical vs animated agents. *Soc Cognitive Affective Neurosci* 2007;2(3):199–205.
- [23] Valdez P, Mehrabian A. Effects of color on emotions. *J Exp Psychol General* 1994;123(4):394–409.
- [24] Simmons DR. The association of colours with emotions: a systematic approach. *J Vision* 2006;6(6):251.
- [25] Ibanez J. Showing emotions through movement and symmetry. *Comput Human Behav* 2011;27(1):561–567.
- [26] Mono R. *Design for product understanding: the aesthetics of design from a semiotic approach*. Liber; 1997.
- [27] Hevner K. Experimental studies of the affective value of colors and lines. *J Appl Psychol* 1935;19(4):385–398.
- [28] Halper N, Mellin M, Herrmann CS, Linneweber V, Strothotte T. Towards an understanding of the psychology of non-photorealistic rendering. *Comp Visualistics Media Inf Virtual Communities* 2003;11:67–78.
- [29] Lang PJ, Bradley MM, Cuthbert BN. *International affective picture system (IAPS): affective ratings of pictures and instruction manual*. University of Florida; 2008.
- [30] Tobii. *Tobii studio 2.2 user manual*. <www.tobii.com/Global/Analysis/Downloads/User_Manuals_and_Guides/Tobii_Studio2.2_UserManual.pdf>; 2010. Last visited January 25, 2012.
- [31] Bojko AA. Informative or misleading? Heatmaps deconstructed *Human Comp Interact New Trends* 2009;5610:30–39.
- [32] Nielsen J, Pernice K. *Eyetracking web usability*. New Riders Press; 2009.
- [33] Johansen SA, Noergaard M, Rau J. Can eye tracking boost usability evaluation of computer games? In: *CHI 2008 workshop on evaluating user experience in games*; 2008. p. 4.
- [34] Stellmach S, Nacke L, Dachsel R. Advanced gaze visualizations for three-dimensional virtual environments. In: *Proceedings of the 2010 symposium on eye-tracking research & applications*. ETRA '10. New York, NY, USA: ACM; 2010. p. 109–12.
- [35] Wallraven C, Cunningham DW, Rigau J, Feixas M, Sbert M. Aesthetic appraisal of art—from eye movements to computers. In: *Computational aesthetics*; 2009. p. 137–44.
- [36] Lopez-Moreno J, Jimenez J, Hadap S, Reinhard E, Anjyo K, Gutierrez D. Stylized depiction of images based on depth perception. In: *Proceedings of the 8th international symposium on non-photorealistic animation and rendering*. NPAR '10. New York, NY, USA: ACM; 2010. p. 109–18.
- [37] Lang PJ, Greenwald MK, Bradley MM, Hamm AO. Looking at pictures: affective, facial, visceral, and behavioral reactions. *Psychophysiology* 1993;30(3):261–273.
- [38] Zhao M, Zhu SC. Portrait painting using active templates. In: *Proceedings of the 9th international symposium on non-photorealistic animation and rendering*. NPAR '11. New York, NY, USA: ACM; 2011. p. 117–24. ISBN 978-1-4503-0907-3.