Enhancing Infographics Based on Symmetry Saliency

Kouhei Yasuda
Graduate School
of Frontier Sciences,
University of Tokyo
Kashiwa, Japan
kohemoko@gmail.com

Shigeo Takahashi
Department of Computer
Science and Engineering
University of Aizu
Aizu-Wakamatsu, Japan
takahashis@acm.org

Hsiang-Yun Wu
Department of Information and
Computer Science
Keio University
Yokohama, Japan
hsiang.yun.wu@acm.org

ABSTRACT

Image saliency is a biologically inspired concept for characterizing visual conspicuity of individual features in natural images, and provides us with a useful insight into the mechanism for directing instant visual attention from viewers. Nevertheless, this perceptual quality often remains to be further sophisticated especially for enhancing saliency in infographic images since they usually consist of relatively simple visual pattens that result in sharp image edges rather than smooth gradations in natural images. This paper presents a new approach to intentionally drawing visual attention for infographic images, in such a way that the corresponding important features naturally pop up in the image. The idea behind our approach is to introduce the concept of symmetry saliency for enhancing local symmetry inherent in such infographic images. This is accomplished by evaluating how much each image edge contributes to the symmetry saliency, and augmenting the corresponding image gradient in proportion to the amount of its contribution. The intensity field of the given image is then modulated with such enhanced image edges by solving the Poisson equation. Several examples together with statistics obtained through a user study demonstrate that our proposed approach successfully improves the readability of infographic images and effectively attracts visual attention to intended regions of interest.

CCS Concepts

•Human-centered computing → Information visualization; •Computing methodologies → Perception;

Keywords

Symmetry saliency; image enhancement; visual attention; infographics

1. INTRODUCTION

Image saliency [7] has been a primary measure for evaluating how much individual image features are conspicuous from a perceptual point of view. The associated computational model for the image analysis facilitates us to choose important regions that naturally pop up in the images. Furthermore, enhancing such image saliency leads us to the idea of intentionally drawing visual attention to specific regions of interest from viewers. This functionality is quite helpful when we have to improve the readability of visual images in such a way that every viewer can fully understand the associated contents in a perceptually plausible manner.

Nonetheless, the conventional saliency model basically assumes as input natural images that usually consist of relatively smooth gradation in the intensity value. This means that we have to incorporate additional mechanism when trying to analyze specific types of image that contain a large amount of sharp image edges. Infographic images are such examples in the sense that they usually consist of relatively simple patterns as visual metaphors and thus contain image edges as the boundary of these visual components. In this case, we have to think of such sharp image edges not only as individual features but also as a group so that we can extract specific patterns from the layout of such visual components.

The Gestalt principles [22] are a well-known set of psychological grouping laws that help us infer the underlying grouping structures of visual objects in the images. In practice, seven grouping principles are usually listed as the Gestalt laws: proximity, similarity, continuity, closure, area (figure and ground), common fate, and symmetry. Among these seven principles, the symmetry principle plays an important role because it allows us to schematize visual images by approximating them as pairs of symmetric patterns extracted from object silhouettes. This means that the symmetry is the most related to the aesthetics in the visual design and serves as the key for understanding schematized visual representations such as infographic images. Another examples on which the symmetry exerts a significant influence include facial images [18], architectural/artistic paintings [14], etc. Thus, incorporating image saliency arising from symmetric structures can lead us to a promising approach for improving the readability of visual images with sharp edges.

This paper presents a novel approach for improving the readability of such infographic images by enhancing local symmetry inherent in the layout of visual patterns. Our idea here is to extract symmetric patterns of image edges, as one of the grouping laws of the Gestalt principles. This is reasonable because infographic representations often employ aesthetic design rules that take into account symmetries in the layout of visual patterns as well as their shapes them-

VINCI '16, September 24-26, 2016, Dallas, TX, USA

DOI: http://dx.doi.org/10.1145/2968220.2968224

selves. Our computation proceeds by evaluating how much each image edge contributes to the saliency arising from local symmetry first, and then exaggerating its gradient of the intensity value according to its contribution. This is followed by an image transformation step, in which the field of intensity value is modulated using the technique similar to tone mapping for high dynamic range imaging. We also present several examples together with results of our user study, to demonstrate how our approach can successfully enhance the image saliency inherent in the underlying symmetric structures.

This paper is organized as follows: Section 2 provides a brief survey on existing work relevant to ours. Section 3 describes details of our proposed approach for enhancing features extracted from the local symmetry embedded in the input images. After presenting several examples together with results of our user study in Section 4, we conclude this paper and refer to possible future extensions in Section 5.

2. RELATED WORK

The concept of the saliency map originated in the work by Koch and Ullman [10]. In general, visual attention has been categorized into two types: bottom-up attention where visual interest arises from the visual stimuli obtained through retinal images [1, 16], and top-down attention where interest has been influenced by knowledge of the viewers about past experiences [20]. The computational model of the saliency map was initiated by Itti et al. [7], and after that a large number of its variants have been developed so far. These includes real-time computation based integral images [4], top-down attention systems [3], and modeling of task influence [15] to name a few.

Directing visual attention is one of the relevant applications based on the saliency model. This research issue is quite important since we often need to intentionally draw visual attention to specific image features in any type of visualization setup. A pioneering work was done by Kim and Varshney [9], where they solved an inverse problem of computing saliency maps from the given images in the context of volume visualization. This has been followed by, for example, semantic depth-of-field model based on saliency [19], saliency enhancement with spatiotemporal coherence [13], task-based saliency for directing top-down visual attention [6], etc. Jänicke and Chen [8] presented a saliency-based metric for evaluating the quality of visualization images, which allows us to choose the best visualization from a set of alternatives. Distorting spatial arrangement of depth cues also permits us to control the distribution of visual attention on perspective images [12, 25].

On the other hand, we need a different model for computing the saliency if we handle a specific type of image that contains features associated with grouping laws of the Gestalt principles, such as infographic images. Several computational models for Gestalt related saliency have been proposed that employ the laws of continuity [21], symmetry [11], similarity [24], and mixture of proximity, similarity, and closure [26]. Our approach specifically employs the model by Kootstra et al. [11] for computing symmetric saliency embedded in the given images.

3. OUR APPROACH

As mentioned earlier, our approach has its basis on the computational model for symmetry saliency [11], which employs an operator for detecting local symmetry [17]. In this model, we need to measure the amount of local symmetry at each pixel, which is influenced by image edges within its local neighborhood. This is accomplished by computing how much an individual image edge contributes to the local symmetry within its vicinity. We then enhance an individual edge by accentuating its image gradient in proportion to the degree of its contribution to the symmetry saliency. Note that we conduct this process at several resolutions of the images individually through the Gaussian pyramid, in the same way as done when computing conventional image saliency [7], and thus we can incorporate an appropriate range of frequency components in this computation. We then try to modulate the field of image intensity by respecting the enhanced gradients associated with the image edges. This has been achieved by solving the Poisson equation, with reference to conventional tone mapping techniques used in high-dynamic range imaging.

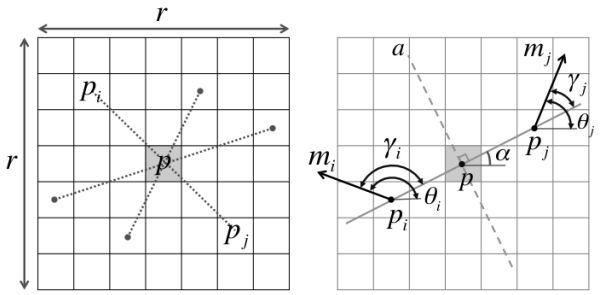

(a) (b)

Figure 1: A symmetry kernel consisting of an $r \times r$ grid of pixels. (a) We consider a pair of pixels p_i and p_j that are symmetric with respect to the center pixel p, and (b) then compute the associated angles θ_i and θ_j and edge sharpness m_i and m_j to compute the symmetry saliency $s(i, j)$ in terms of p_i and p_j.

3.1 Computing Symmetry Saliency Maps

First of all, we limit ourselves to grayscale images as input. Let us consider an $r \times r$ grid of pixels on the input image as shown in Figure 1. Here, we call this rectangular grid a *symmetry kernel* [17], where r is set to 25 by default.

Our first step is to apply Sobel filters at every pixel to compute the sharpness of its vertical and horizontal edges, where the corresponding filters are represented by the following 3×3 masks:

$$I_x = \begin{bmatrix} 1 & 0 & -1 \\ 2 & 0 & -2 \\ 1 & 0 & -1 \end{bmatrix}, \quad I_y = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix} \tag{1}$$

This means that we compute a weighted sum of the 3×3 grid of pixels around the target pixel while referring to the matrices of weight values shown above. We then compute the sharpness and orientation of the edge at pixel p_i as:

$$m_i = \sqrt{I_x(p_i)^2 + I_y(p_i)^2}, \quad \text{and} \tag{2}$$

$$\theta_i = \tan^{-1}\Big(\frac{I_y(p_i)}{I_x(p_i)}\Big), \tag{3}$$

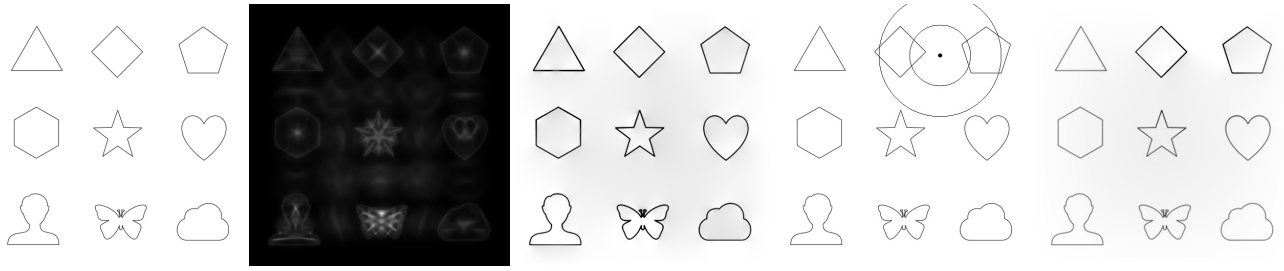

(a) (b) (c) (d) (e)

Figure 2: A simple example. (a) An input image. (b) Symmetry saliency map. (c) Entirely enhanced image. (d) Contour lines representing the region for local enhancement. (e) Locally enhanced image.

where $I_x(p_i)$ and $I_y(p_i)$ indicate the sharpness of horizontal and vertical image edges at pixel p_i.

The symmetry saliency at pixel p, when sampling a symmetric pair of pixels p_i and p_j, can be given by:

$$s(p_i, p_j) = d(p_i, p_j, \sigma) \cdot c(p_i, p_j) \cdot \log(1+m_i) \cdot \log(1+m_j). \quad (4)$$

Here, the weight function $d(p_i, p_j, \sigma)$ is defined in such a way that it increases as the distance between p_i and p_j (i.e., $|p_j - p_i|$) becomes smaller. In practice, $d(p_i, p_j, \sigma)$ is set to be

$$d(p_i, p_j, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp(-\frac{|p_j - p_i|^2}{8\sigma^2}), \quad (5)$$

where $\sigma = 32$ by default. Moreover, $c(p_i, p_j)$ is defined as:

$$c(p_i, p_j) = (1 - \cos(\gamma_i + \gamma_j)) \cdot (1 - \cos(\gamma_i - \gamma_j)). \quad (6)$$

Here, we assume that α represents an angle spanned by the x-axis and line $\overline{p_i p_j}$, as shown in Figure 1(b), which means that γ_i and γ_j correspond to $\theta_i - \alpha$ and $\theta_j - \alpha$, respectively.

The total amount of local symmetry $S_l(p)$ can be computed by summing up $s(p_i, p_j)$ for every pair of p_i and p_j within the symmetry kernel of p as:

$$S_l(p) = \sum_{(p_i, p_j) \in \Gamma} s(i, j), \quad (7)$$

where Γ represents the symmetry kernel of p. Notice that l represents the layer ID of the image in the Gaussian pyramid. In this approach, the local symmetry $S_l(p)$ is calculated independently at each layer, and then the local symmetry values are summed up at all the levels through the entire Gaussian pyramid by equalizing the resolution (i.e., sampling rate) of the image at the respective layers. Figure 2(b) presents the symmetry saliency map of the input image shown in Figure 2(a), where white pixels indicate high saliency regions.

3.2 Modulating Gradients of Image Edges

For improving the readability of infographic representations, we modulate the gradients of image edges according to their contribution to the symmetry saliency of the images. First, we compute the directional gradient values in terms of the pixel coordinates. Let us denote the intensity value at the pixel coordinates (x, y) by $h(x, y)$. The gradient g and its direction v at (x, y) are defined as

$$g(x, y) = \| \nabla h(x, y) \| \quad \text{and} \quad (8)$$

$$v(x, y) = \nabla h(x, y) / \| \nabla h(x, y) \|, \quad (9)$$

respectively, where the partial derivatives of h can be calculated as finite differences. We are now ready to consider how we can modulate the gradient at each pixel. Suppose that we consider the pixel coordinates $c = (i, j)$ and its $r \times r$ symmetry kernel. This means that pixel $q = (x, y)$ within the symmetry kernel will contribute the symmetry saliency at the center pixel c. Let us denote this contribution by $t(q, c)$, which can be derived from Eq. (4). The total contribution of q over the entire image is given by

$$T(q) = \sum_{c \in image} t(q, c). \quad (10)$$

After having accumulated the contribution of pixels c over the entire image, we normalize the contribution values at each pixel by referring to the maximum and minimum contribution values $T_{\max}$ and $T_{\min}$, respectively. Let us denote the normalized contribution of pixel q as $\overline{T}(q) \in [0, 1]$. The gradient $g(x, y) = g(q)$ will be transformed to:

$$g'(q) = ((b - a) \cdot \overline{T}(q) + a) \cdot g(q). \quad (11)$$

Note that a and b indicate the minimum and maximum enhancement ratio of the modulated gradient, and we set $a = 1.0$ and $b = 3.0$ unless stated otherwise, while we employ $a = 1.0$ and $b = 4.0$ for producing results in Figures 2(c)-(e).

3.3 Reconstructing Fields of Image Intensities

For finally enhancing the image edges based on the symmetry saliency, we employ conventional gradient-based techniques for compressing the high dynamic range of image intensity [2, 23]. In the original dynamic range compression, they tried to compress contrasts inherent in image edges while maximally retaining textures with small gradients. On the other hand, we employ this technique to modulate the field of image intensity values while respecting enhanced image edges as described above.

The formulation of our technique can be summarized as follows. Recall that the intensity value and enhanced gradient at pixel $q = (x, y)$ are $h(x, y)$ and $g'(x, y)$ (cf. Eq. (11)), respectively. What we have to do here is to find the optimized integration $h'(x, y)$ as the following integral:

$$h'(x, y) = \underset{h}{\operatorname{argmin}} \iint \|\nabla h(x, y) - g'(x, y)\|^2 dx dy \quad (12)$$

This eventually amounts to solving the Poisson equation:

$$\nabla^2 h' = \operatorname{div} g'. \quad (13)$$

In order to solve this equation, we employed the successive overrelaxation (SOR) method, which is an extension of the

Gauss-Seidel method. This allows us to improve the readability of infographic images by enhancing the image edges that contribute to the local symmetry. Figure 2(c) shows such an example where image edges over the entire domain have been enhanced according to the amount of their contribution to the symmetry saliency.

3.4 Enhancing Local Image Regions

It is also possible to limit our edge enhancement process to a specific region when we want to intentionally direct visual attention from viewers. This is achieved by assigning more weight to the contribution of that region by convolving a Gaussian kernel around the center pixel $c = (i, j)$ in that region, as follows:

$$T(q) = \sum_{c \in image} w(c) \cdot t(q, c), \tag{14}$$

where $w(c)$ indicates the Gaussian kernel centered at $c = (i, j)$. Of course, we can control the range of this Gaussian kernel by adjusting its standard deviation σ. By employing this updated amount of contribution to local symmetry at q, we can modulate the field of image intensity values again by solving Poisson equation in Eq. (13). Figure 2(e) shows such an example where the edges of the input image have been specifically augmented in the neighborhood of the specific position outlined by contour lines in Figure 2(d).

3.5 Extending Formulation to Color Images

Finally, we describe how we can extend our formulation to color images. Basically, even for color images, we can still follow the same computational process as we did for grayscale images, except for the definition of symmetry saliency. Indeed, we replace the grayscale symmetry operator in Eq. (4) with the following color symmetry operator [5]:

$$s'(p_i, p_j) = \sum_{(k_i, k_j) \in K} d(p_i, p_j, \sigma) \cdot c(p_i, p_j, k_i, k_j) \cdot \log(1 + m_i(k_i)) \cdot \log(1 + m_j(k_j)), \tag{15}$$

where K represents the set of all combinations of colors from red (R), green (G), and blue (B), i.e., $K = \{(R, R), (R, G), (R, B), (G, R), (G, G), (G, B), (B, R), (B, G), (B, G)\}$. This means that $c(p_i, p_j, k_i, k_j)$ can be computed by first referring to the edge orientations of color k_i at pixel p_i and color k_j at pixel p_j, and then calculating Eq. (6). After having retrieved $m_i(k_i)$ and $m_j(k_j)$ as the sharpness of the image edges of color k_i at pixel p_i and color k_j at pixel p_j, respectively, we can finally accumulate symmetry saliency values for all the colors by computing Eq. (7).

4. RESULTS

This section presents several examples of enhanced images obtained using the proposed approach, together with results of our user study and discussions.

4.1 Image Enhancement Examples

Our prototype system was implemented on a laptop PC with Intel Core i7 CPU (1.7GHz, 4MB cache) and 8GB RAM; the source code is written in C++ using OpenCV library for handling images. Overall, it took approximately one minute or a little more to finish up the modulation for one image in this computational environment.

We took as inputs images of statistical charts, a node-link diagram, a tag cloud, and a city map, as shown in the 1st row of Figure 3. Note that the first two are grayscale images while the last two images are in color. We first computed the symmetry saliency map of each input image through its Gaussian pyramid representation using the conventional algorithm [11] as shown in the 2nd row of Figure 3, and then enhanced the image edges in proportion to its contribution to the symmetry saliency as shown in the 3rd row of Figure 3.

We also tried to locally enhance the image edges derived from the symmetry saliency so as to intentionally direct visual attention to a specific region from viewers. The top row of Figure 3 also indicates contour lines, which corresponds to level sets of Gaussian functions we employed for locally enhancing the symmetry saliency in the respective images. The associated transformations were reflected in locally accentuated images where image edges are naturally modulated according to the position of the weighting Gaussian functions, as shown in the 4th row of Figure 3. We again computed the symmetry saliency maps of these locally enhanced infographic images as shown in the bottom row of Figure 3, which helps assess the effects of such local edge enhancements. In practice, the comparison between the symmetry saliency maps in the 2nd and 5th rows confidently predicts that our approach for locally enhancing symmetry saliency will successfully attract viewers' visual attention to the specified region, from a computational point of view. It is also noteworthy that our image enhancer smoothly augmented the underlying symmetry saliency over the image, and thus produces naturally looking visual representations regardless of the types of infographic images.

4.2 User Study

We also conducted a user study in order to evaluate practical effects of the proposed image enhancement. Our user study was conducted in two parts. The first part is to seek feedback from participants through the online questionnaires, to identify how much the image readability is improved while unnatural artifacts are incorporated according to the degree of image enhancement. The second part is to perform eye-tracking studies for investigating how much we can draw visual attention on specified regions of interest.

The objective of the first part is to identify the proper range of the enhancement ratio b in Eq. (11), where we fix a to be 1.0. In this part, we recruited 21 participants (7 females and 14 males) where half of them major in computer graphics, image processing, and other relevant fields, and their ages ranged from 20 to 59. As the first example, we generated enhanced grayscale images of a node-link diagram with several enhancement ratios $b = 1.25$, 1.5, 2.0, 2.5, 3.0, 3.5, and 4.0 as shown in Figure 4. Here, the participants are requested to see the original image (Figure 4(a)) and the enhanced images (Figures 4(b)-(h)) all at once, and answer whether the readability of each enhance image is improved or not and whether it contains unnatural artificial effects. Figure 4 shows the statistics of the responses from the participants. Note that the percentage of "None" represents the ratio of the participants who could not recognize any readability improvements and artificial effects on all the enhanced images. We also ask the participants to carry out the same evaluation for the tag cloud color images as shown in Figure 5, as the second example.

For the grayscale image of a node-like diagram, the most

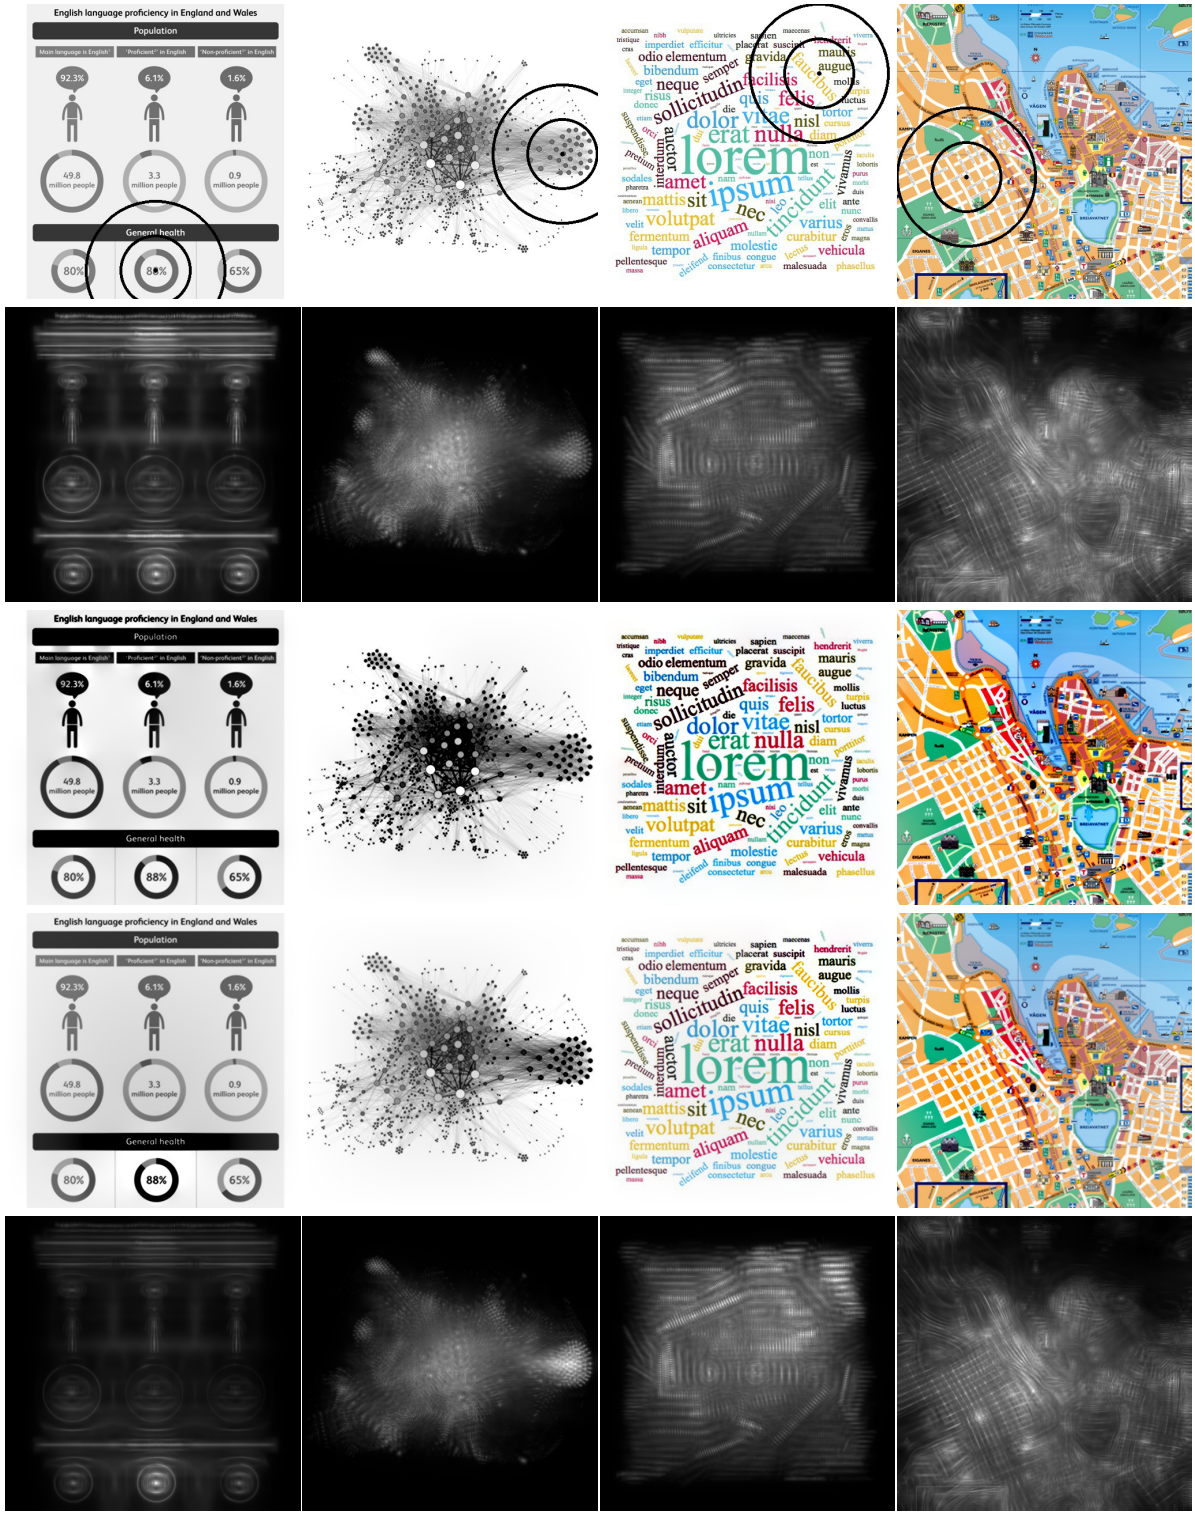

(a) (b) (c) (d)

Figure 3: Enhancing symmetry saliency in infographic images: (a) Statistical charts (grayscale). (b) Node-link diagram (grayscale). (b) Tag cloud (color). (c) City map (color). 1st row: Input images (with contour lines for local symmetry enhancement). 2nd row: Symmetry saliency maps. 3rd row: Entirely enhanced images. 4th row: Locally enhanced images. 5th row: Symmetry saliency maps of the locally enhanced images.

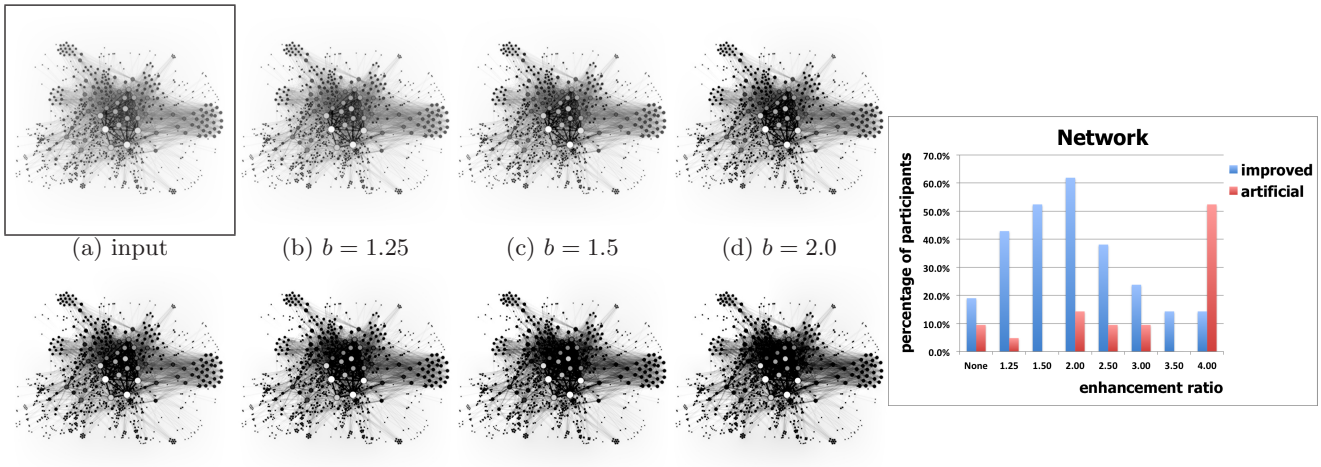

(a) input (b) $b = 1.25$ (c) $b = 1.5$ (d) $b = 2.0$

(e) $b = 2.5$ (f) $b = 3.0$ (g) $b = 3.5$ (h) $b = 4.0$ (i) statistics

Figure 4: The input node-link diagram image and their enhanced versions with enhancement ratios $b = 1.25$, 1.5, 2.0, 2.5, 3.0, 3.5, and 4.0 in Eq. (11).

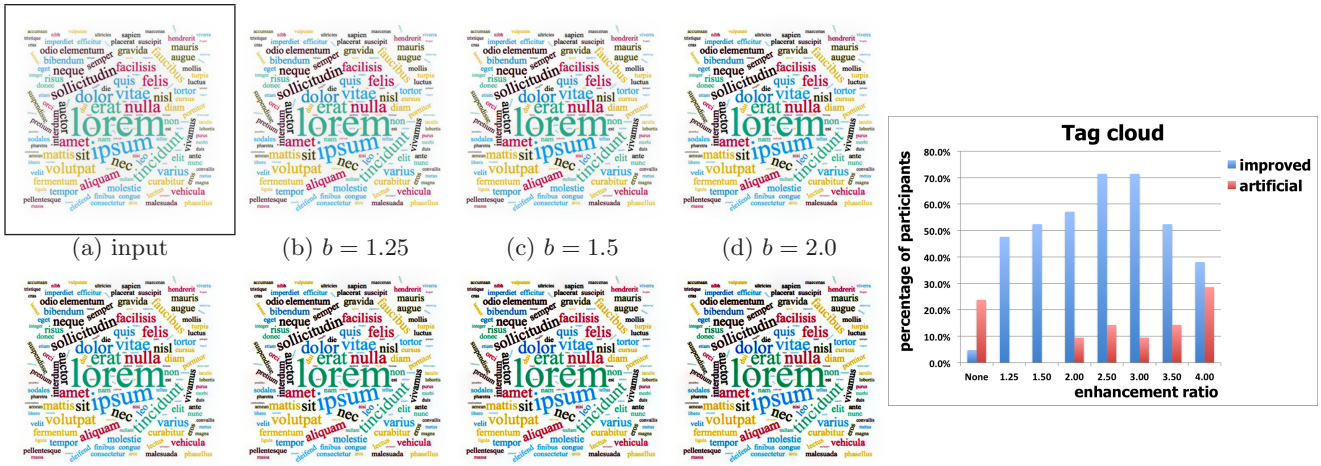

(a) input (b) $b = 1.25$ (c) $b = 1.5$ (d) $b = 2.0$

(e) $b = 2.5$ (f) $b = 3.0$ (g) $b = 3.5$ (h) $b = 4.0$ (i) statistics

Figure 5: The input tag cloud images and their enhanced versions with enhancement ratios $b = 1.25$, 1.5, 2.0, 2.5, 3.0, 3.5, and 4.0 in Eq. (11).

participants liked the image with the enhancement ratio $b = 2.0$ while the majority of them feel excessive intensity emphasis when $b = 4.0$. On the other hand, for the color image of a tag cloud, the participants still preferred enhanced images with $b = 2.5$ and $b = 3.0$ and found all the enhanced images less artificial. In practice, we presented a few more infographic images to the participants in this first part, and recognized that these trends are common in the statistics both for grayscale and color images. This suggested that we can employ $b = 2.0$ for grayscale infographic images and $b = 2.5$ or 3.0 for color images, as the compromise between the effects and artifacts of the proposed image enhancement.

In the second part, we evaluated how much we could intentionally direct visual attention to specific regions on the input infographic images through eye-tracking experiments. In this part, we collected 12 participants (2 females and 10 males) with normal or corrected to normal vision, and their ages ranges from 22 to 35. Participants are asked to freely look at a set of original images, a set of entirely enhanced images, and a set of locally enhanced images, while the orders of the three images sets and the images in each set were both randomly shuffled. In this experiment, each image was presented to the participants for 5 seconds immediately after an initial gaze fixation on a cross mark for 3 seconds. Note that we spared a short break after an individual set of images were presented to sufficiently refresh the participants.

We employed the Tobii X120 eye-tracker to record the transition of the eye-gaze fixation over the respective images. Here, we conducted comparisons between original, entirely enhanced, and locally enhanced images as exhibited in Figure 6, where, for each image, we produced two heat maps that represent the relative durations of gaze fixation for the first half (from 0.5 to 2.5 seconds) and second half (from 2.5 seconds to 4.5 seconds), respectively. Note that we skipped the first 0.5 seconds to exclude influence from the initial gaze fixation at the image center, while the last 0.5

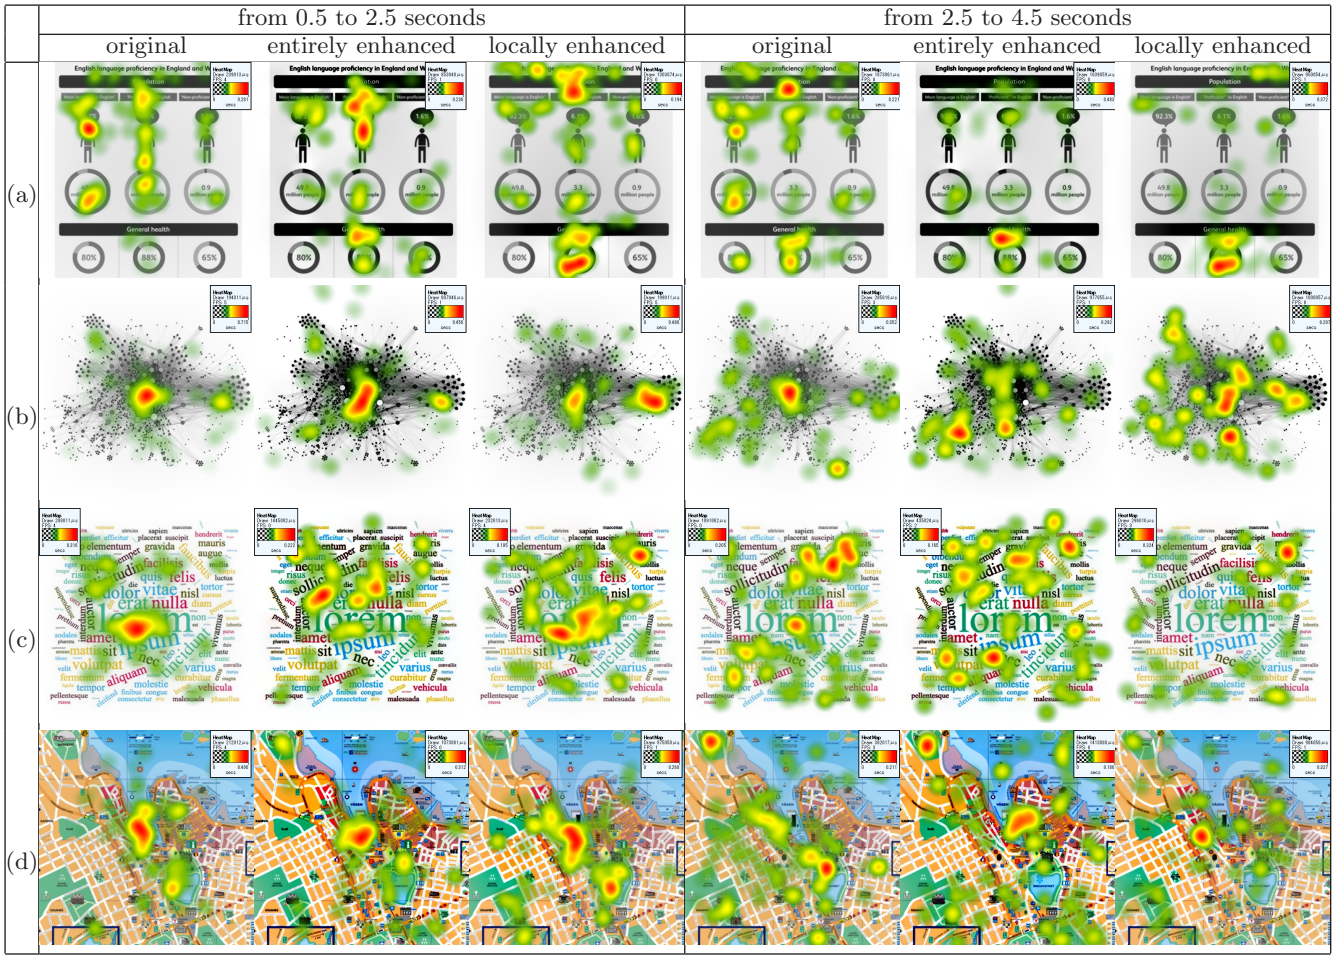

Figure 6: Heat maps representing the relative durations of gaze fixation from 0.5 to 2.5 seconds (left table) and 2.5 to 4.5 seconds (right table). In each table, three columns correspond to the original (left), entirely enhanced (center), and locally enhanced (right) images, respectively, while the four rows (a)-(d) correspond to the input images in Figure 3.

seconds to equalize the periods of image presentation in the first and second halves. We also set $b = 3.0$ in Eq. (11) for enhancing all the images both entirely and locally. As shown in the color legends, the color of each heat map changes from green to yellow to red as the fixation duration increases.

Overall, the results show that each locally enhanced image synthesized using our approach could draw more visual attention to the specified region of interest already in the first half. On the other hand, an entirely enhanced image usually directed attention to all the major edges, and thus the eye-gaze distribution was scattered rather uniformly when compared with that for the locally enhanced one. One interesting exception is the city map case, where the eye-gaze distributions were almost the same for the three images in the first half, while in the second half the locally enhanced image drew more attention to the specified region. Actually, our observation suggests, as the image contents become more complicated, it takes more time until the distributions of eye-gaze fixation differ among the three images. This can be justified from that we need more time to direct visual attention in the order of Figures 6 (a), (b), (c), and (d).

Last but not least, all the participants were asked whether they could identify differences among the three images after their tasks, and surprisingly none of them could not do so. This implies that our synthesized images can naturally draw attention from them without creating noticeable artifacts.

4.3 Discussions

One possible failure case is infographic images with annotation texts, in which we fail to sufficiently direct visual attention to specific patterns. In practice, annotation texts are more likely to hold attention from viewers, and some special schemes may be necessary for fully enhancing the visual patterns inherent in infographic images while avoiding conflicts with the texts. Another possible problem is that we cannot sufficiently attract visual attention if the input image only has dull edges or little local symmetry. Adaptively adjusting the degree of symmetry enhancement over the image still remains as our future problems, especially for handling images having both smooth graduation and sharp edges as we often encounter in the context of scientific visualization.

5. CONCLUSION

This paper has presented a novel approach to enhancing

infographic images by referring to their symmetry saliency. Following the Gestalt principle that the symmetry significantly attracts our visual attention, we enhance the image edges arising from local symmetry by taking advantage of conventional model for computing such symmetry saliency. This has been accomplished by enhancing image edges in proportion to the degree of their contribution to local symmetry, and modulating the field of intensity values so that we can accommodate enhanced image edges within the image quantization levels through the Poisson equation solver. Synthesized image examples together with the results of our user study justify the mechanism of our approach.

6. ACKNOWLEDGMENTS

This work has been partially supported by JSPS KAKENHI Grant Numbers 16H02825, 26730061, and 15K12032, and MEXT KAKENHI Grant Number 25120014.

7. REFERENCES

[1] J. Duncan and G. W. Humphreys. Visual search and stimulus similarity. *Psychological Review*, 96:433–458, 1989.

[2] R. Fattal, D. Lischinski, and M. Werman. Gradient domain high dynamic range compression. *ACM Transactions on Graphics*, 21(3):249–256, 2002.

[3] S. Frintrop, G. Backer, and E. Rome. Goal-directed search with a top-down modulated computational attention system. In *Proceedings of the Annual Meeting of the German Association for Pattern Recognition*, volume 3663 of *Springer Lecture Notes in Computer Science*, pages 117–124, 2005.

[4] S. Frintrop, M. Klodt, and E. Rome. A real-time visual attention system using integral images. In *International Conference on Computer Vision Systems*, 2007.

[5] G. Heidemann. Focus-of-attention from local color symmetries. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(7):817–830, 2004.

[6] S. Hillaire, A. Lecuyer, T. Regia-Corte, R. Cozot, J. Royan, and G. Breton. Design and application of real-time visual attention model for the exploration of 3D virtual environments. *IEEE Transactions on Visualization and Computer Graphics*, 18(3):356–368, 2012.

[7] L. Itti, C. Koch, and E. Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11):1254–1259, 1998.

[8] H. Jänicke and M. Chen. A salience-based quality metric for visualization. *Computer Graphics Forum*, 29(3):1183–1192, 2010.

[9] Y. Kim and A. Varshney. Saliency-guided enhancement for volume visualization. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):925–932, 2006.

[10] C. Koch and S. Ullman. Shifts in selective attention: Towards the underlying neural circuitry. *Human Neurobiology*, 4:219–227, 1985.

[11] G. Kootstra, B. de Boer, and L. Schomaker. Predicting eye fixations on complex visual stimuli using local symmetry. *Cognitive Computation*, 3(1):223–240, 2011.

[12] K. Mashio, K. Yoshida, S. Takahashi, and M. Okada. Automatic blending of multiple perspective views for aesthetic composition. In *Proceedings of the 10th International Symposium on Smart Graphics (Smart Graphics 2010)*, volume 6133 of *Springer Lecture Notes in Computer Science*, pages 220–231, 2010.

[13] E. Mendez, S. Feiner, and D. Schmalstieg. Focus and context in mixed reality by modulating first order salient features. In *Proceedings of the 10th International Symposium on Smart Graphics (Smart Graphics 2010)*, volume 6133 of *Sprinter Lecture Notes in Computer Science*, pages 232–243, 2010.

[14] L. Nan, A. Sharf, K. Xie, T.-T. Wong, O. Deussen, D. Cohen-Or, and B. Chen. Conjoining Gestalt rules for abstraction of architectural drawings. *ACM Transactions on Graphics*, 30(6):185, 2011.

[15] V. Navalpakkam and L. Itti. Modeling the influence of task on attention. *Vision Research*, 45:205–231, 2005.

[16] H. C. Nothdurft. The role of features in preattentive vision: comparison of orientation, motion and color cues. *Vision Research*, 33(14):1937–1958, 1993.

[17] D. Reisfeld and Y. Yeshurun. Preprocessing of face images: Detection of features and pose normalization. *Computer Vision and Image Understanding*, 71(3):413–430, 1998.

[18] G. Rhodes, F. Proffitt, J. M. Grady, and A. Sumich. Facial symmetry and the perception of beauty. *Psychonomic Bulletin & Review*, 5(4):659–669, 1998.

[19] Z. Su and S. Takahashi. Real-time enhancement of image and video saliency using semantic depth of field. In *Proceedings of International Conference on Computer Vision Theory and Applications (VISAPP 2010)*, pages 370–375, 2010.

[20] A. M. Treisman and G. Gelade. A feature-integration theory of attention. *Cognitive Psychology*, 12:97–136, 1980.

[21] Z. Wang and B. Li. A two-stage approach to saliency detection in images. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, 2008 (ICASSP 2008)*, pages 965–968, 2008.

[22] C. Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann Publishers Inc., 3rd edition, 2013.

[23] T. Weyrich, J. Deng, C. Barnes, S. Rusinkiewicz, and A. Finkelstein. Digital bas-relief from 3D scenes. *ACM Transactions on Graphics*, 26(3):32, 2007.

[24] J. Wu and L. Zhang. Gestalt saliency: Salient region detection based on Gestalt principles. In *Proceedings of the 20th IEEE International Conference on Image Processing (ICIP 2013)*, pages 181–185, 2013.

[25] K. Yoshida, S. Takahashi, H. Ono, I. Fujishiro, and M. Okada. Perceptually-guided design of nonperspectives through pictorial depth cues. In *Proceedings of 7th International Conference on Computer Graphics, Imaging and Visualization (CGiV2010)*, pages 173–178, 2010.

[26] J.-G. Yu, G.-S. Xia, C. Gao, and A. Samal. A computational model for object-based visual saliency: Spreading attention along Gestalt cues. *IEEE Transactions on Multimedia*, 18(2):273–286, 2016.