Design and Evaluation of Multifield Visualisation Techniques for 2D Vector Fields

Chris C. van Egmond
Graphics Group, Department
of Computer Science
University of Auckland
Auckland, New Zealand
cvan680@aucklanduni.ac.nz

Charl P. Botha
Graphics Group, Department
of Mediamatics
TU Delft
Delft, The Netherlands
c.p.botha@tudelft.nl

Burkhard C. Wünsche
Graphics Group, Department
of Computer Science
University of Auckland
Auckland, New Zealand
burkhard@cs.auckland.ac.nz

ABSTRACT

The visualisation of vector fields is essential for many applications in science, engineering and biomedicine. A large number of vector icons has been developed, but little research has been done on their effectiveness, especially when visualising multiple vector fields simultaneously. We apply research in visualisation and cognitive science to identify four classes of post-processing techniques for visualising two 2D vector fields simultaneously: *blending*, *overlay*, *bump mapping*, and *masking*. We apply these four post-processing methods to Line Integral Convolution (LIC) textures and thus develop several novel multi-field visualisation techniques. We evaluate their effectiveness with a user study requiring participants to locate and classify singularities, and to rate each method on its effectiveness and aesthetic value. The results of the study suggest that *blending* is the most effective technique to combine multiple vector field visualisation textures, while *masking* performs worst. There is some evidence that visualisations with smooth colour changes are perceived as visually more attractive, and that aesthetics increases the perceived effectiveness of a visualisation technique.

Categories and Subject Descriptors

I.3.3 [**Picture/Image Generation**]: Display algorithms; I.3.7 [**Three-Dimensional Graphics and Realism**]: Color, shading, shadowing, and texture; I.6.8 [**Simulation and Modeling**]: Types of Simulation—*Visual*

Keywords

User study, vector field visualisation, flow visualisation, multifield visualisation

1. INTRODUCTION

Vector field visualisation is an important task in scientific visualisation. A large variety of techniques has been pro-

posed ranging from vector glyphs [12] and textural representations [26, 2], to the visualisation of structural information such as the vector field topology [19, 7].

One central issue of current research is multi-field visualisation [18]. Novel image acquisition and simulation techniques have made is possible to record a large number of co-located data fields. Visualising them simultaneously facilitates the investigation of correlations and relationships between the fields [24]. An example is the visualisation of multi-field cardiac imaging data (MRI, PET, CT, DTI, tagged MRI, molecular imaging) in order to analyse the relationship between cardiac function, structure, anatomical changes, metabolic activity, blood perfusion, and cellular remodelling.

In this paper we investigate texture-based visualisation techniques for displaying two 2D vector fields simultaneously. Section 2 reviews previous work on vector field visualisation and multi-field visualisation. Section 3 explains the design of our multi-field visualisation techniques. Section 4 presents an evaluation of the techniques with results being discussed in section 5. We conclude this paper and give an outlook on future work in section 6.

2. LITERATURE REVIEW

Texture-based representations are popular for vector field visualisation due to their high spatial resolution (vector information is displayed for many points of the domain with comparatively little screen space) and the ease with which they can be applied to different 2D and 3D geometries [15].

One of the earliest classes of techniques in this domain is "Spot noise", which distributes a large number of spots over a surface and transforms them according to the underlying vector field [4]. A contrasting approach is to place streamlines close together until the visualisation domain is covered and then separating them [10].

The arguably most popular texture-based vector field visualisation technique is Line Integral Convolution (LIC) [2]. The method places short streamlines over each pixel of an input texture (usually white noise), and then applies a convolution kernel to the texels covered by each streamline. LIC has several advantages such as indicating vector field direction in every point, indicating local information such as convergence, producing 2D or 3D textures which can be visualised using standard rendering techniques, and producing images with a fine spatial resolution. Moreover, unlike streamlines, LIC does not depend on the selection of good seed points.

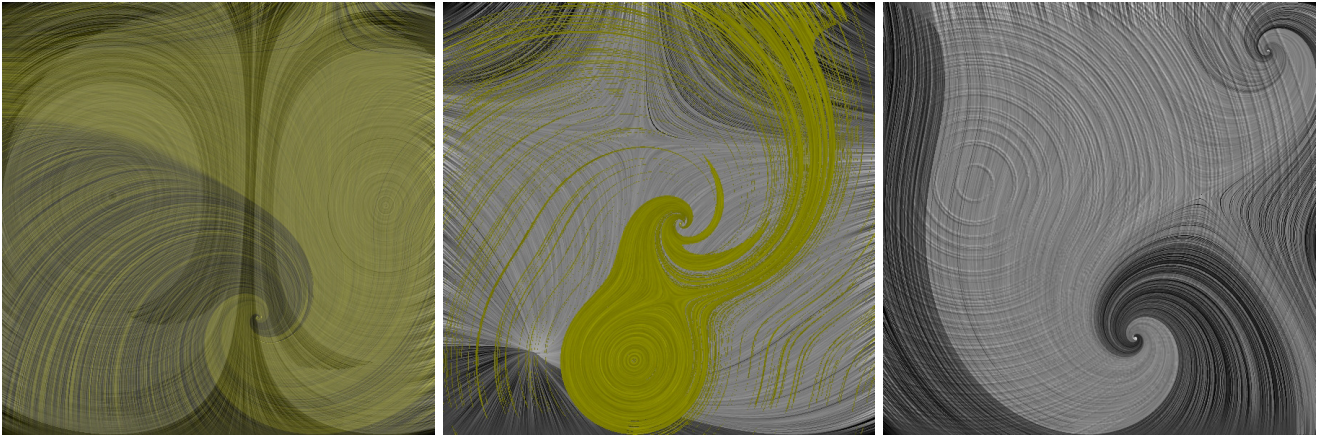Multi-field datasets contain several fields defined over the

Figure 1: **Different pairs of 2D vector fields visualised simultaneously using the blend method (left), the overlay method (middle), and the smooth bump map method (right).**

same domain. Blaas et al. introduce interaction techniques for selecting and comparing components of interest [1]. Kniss et al. present novel transfer functions for direct volume rendering high-dimensional volume data sets [13]. Weigle uses rectangular texture elements to encode multiple variables [27]. Fuchs and Hauser give an overview of state-of-the-art techniques for visualizing multi-field data [6].

Much can be learned from observing the occurence of complex textures in other application fields. Interrante suggests to harness the intricate variety and subtle richness of detail of natural texture patterns [9]. Using artistic concepts and associating vector information with brush strokes in an oil-painting makes it possible to combine multiple layers of visual elements to a surface texture encoding multiple fields, while still allowing perception of each layer [11].

Little work has been done on comparing the effectiveness of vector field visualisation techniques. Laidlaw et al. investigated the effectiveness of six visualisation methods for two dimensional vector fields by means of a user study [14]. Urness et al. present strategies for visualising multiple 2D vector fields by overlaying texture and glyph-based techniques, but do not evaluate their algorithms. Our research identifies key concepts for combining texture-based visualisations, presents several new techniques, and evaluates them with a user study.

3. DESIGN

We created LIC textures of vector fields using Voreen [16] and used concepts from graphics and cognitive psychology to combine them. Except for the bump map, all vector fields were rendered using an implementation of the fast LIC algorithm [22] in Voreen (sampling each pixel and using a large convolution kernel $k = 100$). Vector field magnitude is usually encoded by colour, but in order to compare the different techniques we only display vector field direction and use colour to perceptually differentiate multiple fields.

3.1 Blend

Alpha blending is frequently used in computer graphics to simulate transparencies or combine different effects such as shading and texture mapping. In order to achieve effective blending we exploit chromatic adaptation, i.e., the human visual system's ability to adjust to changes in illumination

and preserve the appearance of object colours [5]. We hence combine two textures I_1 and I_2 by encoding them with different hues (gray and yellow) and, as in the original LIC paper, vary the intensity for the white noise texture. The textures are then blended using an alpha value $\alpha \in [0.4, 0.6]$.

$$I_{\text{result}}(x,y) = \alpha I_1(x,y) + (1 - \alpha) I_2(x,y) \qquad (1)$$

The left hand image of figure 1 shows an example.

3.2 Overlay

Figure-ground perception describes the observation that an object can be instantly perceptually separated from its background [21]. This is due to physically different attributes of a figure and its background but is also influenced by size, angle, and association with meaningful shapes. We employ this concept to *overlay* two LIC textures by rendering them using different hues (like above) and replacing each pixel of texture I_1 with the corresponding pixel of texture I_2, where it exceeds a threshold t:

$$I_{\text{result}}(x,y) = \begin{cases} I_1(x,y) & \text{if } I_2(x,y) \leq t \\ I_2(x,y) & \text{if } I_2(x,y) > t \end{cases} \qquad (2)$$

The image in the middle of figure 1 shows an example.

3.3 Bump Map

Shape perception is dominated by the curvature of the silhouette contour (figure-ground boundary) and 3D surface shading [8]. We exploit this property by combining two textures by representing one of them using shading differences which simulate surface deformations. This concept is called *bump mapping* in computer graphics. In our implementation we create a bump map by rendering the second vector field using a much lower sampling rate of the LIC algorithm, and applying a Roberts' Cross diagonal gradient filter [20] R to calculate edges. This gradient image is multiplied with the first image, resulting in the second vector field appearing as "bumps". In our subsequent experiments we will use a *crisp* version explained above (equation 3), and a *smooth* version where the bumpmap is first convolved with a 3 by 3 Gaussian kernel G (equation 4 and figure 1 (right)).

$$I_{\text{result}}(x,y) = I_1(x,y)(I_2 * R)(x,y) \qquad (3)$$

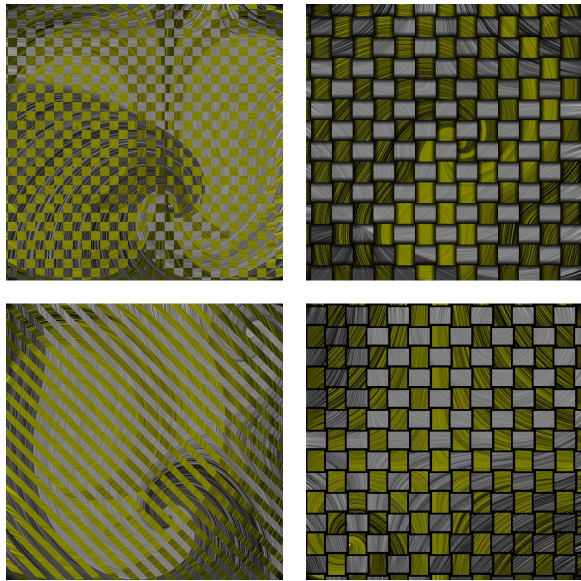$$I_{\text{result}}(x,y) = I_1(x,y)((I_2 * R) * G)(x,y) \qquad (4)$$

Figure 2: Four visualisations showing two vector fields simultaneously using different pattern masks. From the top left image, in clockwise order: a fine checkerboard pattern, a gradual shadow weaving pattern, a solid shadow weaving pattern, and a fine diagonal pattern.

3.4 Pattern Mask

The concept of Gestalt originates from the fine arts and expresses the notion that the "whole contains more information than the parts". An example is the perception of a circular arrangement of symbols as a circle. Perception of Gestalt is influenced by proximity, similarity, continuation, closure, symmetry, and the law of Prägnanz, which states the the eyes tend to see the simplest and most stable figure [21]. We exploit Gestalt perception by using different masks, which subdivide the domain of two vector fields, and show for each section only one vector field. By using proximity and continuation the human visual system should be able to mentally interpolate a vector field over a region where it is not displayed. Figure 2 shows four examples of masking patterns.

Checkerboard pattern: These patterns have been successfully used in medical imaging for image fusion [23]. We investigate a coarse and a fine pattern in order to determine the influence of size on the human visual system's ability to mentally interpolate textures.

Diagonal, horizontal and vertical pattern: We experimented with different masking pattern orientations. In our user study we use a coarse and a fine version for a diagonal pattern, and a coarse version of a horizontal and vertical pattern.

Weaving pattern: Weaving patterns are distinctive and most users are familiar with them. We use a gradual shadow and a solid shadow checkerboard weaving pattern, and a longer weaving pattern using a gradual shadow. The use of shadows improves contrast between regions at the expense of obscuring continuation of texture patterns.
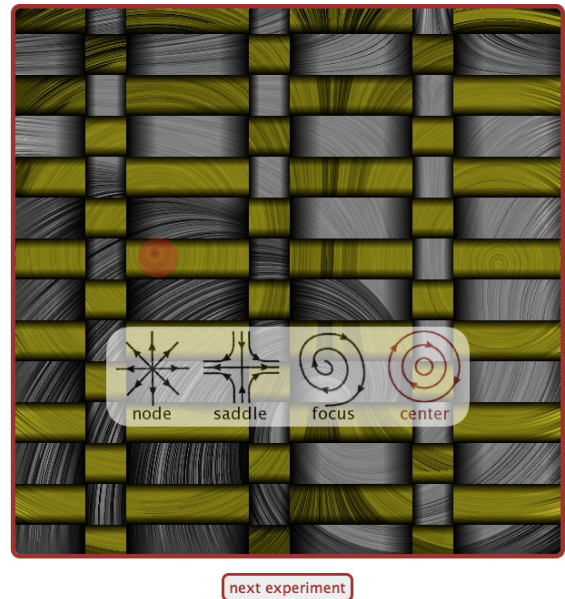
4. EVALUATION

next experiment

Figure 3: Web interface used in the user study. The screen shot shows a critical point in the yellow vector field (indicated as transparent red circle) located by clicking on the canvas. The click has triggered a context menu enabling the user to select the type of critical point. After locating and classifying all identified critical points clicking the button at the bottom displays the next experiment.

In order to evaluate the above presented visualisation techniques we created seven unique pairs of different vector fields. Each vector field was represented by a 600×600 pixel texture, and contained 1 to 3 critical points (nodes, center, focus or saddle). Consequently each simultaneous visualisation of two vector fields contained 2 to 6 critical points. The precise location and type of each critical points was stored in a database, and compared with the user input. A total of thirteen different visualisation techniques, using the four concepts introduced above, were evaluated.

4.1 Methodology

We created 91 visualisations by using seven pairs of vector fields and thirteen visualisation techniques employing the four concepts introduced in section 3. An online survey was created showing for each of the 13 visualisation techniques one image displaying one randomly selected pair of vector fields (each vector field occurred equally often). The first page of the survey contained a quick explanation of the tasks (see subsection 4.2).

Removing Bias: The participants were not given any form of training other than an introductory text and example images. In order to counter "learning bias", the order in which the visualisations were tested was randomised for each participant. Also, to avoid bias based on some vector field potentially being more difficult than others, an image was rendered for each combination of visualisation methods and vector field pairs.

Participant Demographics: The user study had 55 participants: 27 participants with experience in the field of

data visualisation ("experts") and 28 "non-experts". Most of the "experts" were researchers or students from the computer graphics research groups at the University of Auckland and the University of Technology, Delft. The survey took roughly 15 minutes. Participation was voluntary, anonymous, and no compensation was given.

4.2 Required Participant Tasks

To test the effectiveness of each of the visualisation methods, participants were required to perform the following tasks for each visualisation:

4.2.1 Locate and Classify Critical Points

For each visualisation participants were required to locate and classify all critical points by first clicking on the precise location and then selecting the type of critical point (see figure 3). The visualisations were presented on a 600 by 600 pixels HTML5 canvas element. A click near a critical point triggered a context menu showing the four singularity types. We chose as admissible target area a circle with a radius of 22 pixels. In a preceding pilot study this provided sufficient distance between features which could be mistaken as critical points, while still taking into account errors due to rushed motions and lack of image information (e.g., when using masking patterns). The process was repeated until a participant clicked the "next experiment" button to indicate that he/she could not find any more singularities. The location of a click, as well as the selected singularity type, was recorded using JavaScript.

4.2.2 User Rating of Perceived Effectiveness and Aesthetic Value

Upon completing the above tasks participants were shown all 13 previous visualisations in identical order on a single page. Users had to rate each visualisation both on its *perceived effectiveness* (ease to locate and classify critical points) and *visually attractiveness*. The questions were rephrased as statements (e.g., "I found this visualisation effective") and ratings were performed using seven-point Likert scales ranging from -3 (strongly disagree) to 3 (strongly agree).

The user rating was performed to detect any discrepancies between participants' perception of *effectiveness* and the actual test results. The aesthetic value of a visualisation techniques is of interest since it can effect its usage in commercial (business) applications. In addition studies have shown that the aesthetic value of a visualisation is related to users' ability and willingness to successfully fulfill visualisation tasks [3, 17].

4.2.3 Measured Data

For each visualisation the selected location and type of critical point and the user rating was stored. In addition the time period for each visualisation task was measured. From this data we derived statistical measures, such as the percentage of found critical points, the percentage of correct classifications, the number of wrongly located or classified critical points.

5. RESULTS AND DISCUSSION

5.1 Locating and Classifying Critical Points

For each participant we computed the percentage of correctly located and classified critical points, and the percent-

age of incorrect choices (i.e., either wrong position or wrong classification). Table 1 shows in the first two columns for each of these measures the mean value and standard deviation.

In order to quantify the overall effectiveness and accuracy of each method, we calculated F_1-scores [25] using equation 5. Precision is the number of correctly classified critical points divided by the total number of clicks, and recall is the number of correctly classified critical points divided by the total number of critical points that should have been found.

$$F_1 = 2 \times \frac{precision \times recall}{precision + recall} \qquad (5)$$

Figure 4 shows a plot of the mean and 95% confidence interval of the F_1-score for each visualisation technique. The *blend* method has the highest score, closely followed by the *overlay* and *crisp bump map* methods. The masking methods perform far worse than the other three methods; even the best performing masking method (the coarse diagonal pattern) has a significantly lower $\bar{F_1}$ than any of the non-masking techniques (the smooth bump map performs only slightly better).

The results indicate that the human visual system's ability of Gestalt continuation are not sufficient to reliably detect features in a vector field visualisation. However, with roughly 80% correctness results are sufficiently good to justify further investigation in selected application cases where the other visualisation techniques are not practical (e.g., if additional scalar fields are encoded by colour).

With regard to different masking patterns the coarse diagonal, vertical, and horizontal patterns worked best and also had a high perceived effectiveness. The weaving pattern with gradient and the fine checkerboard pattern rated worst. The results indicate that visual interpolation of missing data is most effective when using relatively large connected regions without boundaries.

We also measured the precision of correctly located critical points. Results were surprisingly similar. The fine diagonal pattern mask performed best with an average location error of 4.61 pixels, followed by *overlay* (4.65 pixels), *blend* (4.78 pixels) and *smooth bumpmap* (4.87 pixels). The worst location error was observed for the gradual weaving pattern (6.20 pixels).

5.2 Response Time

We evaluated the efficiency of the presented visualisation techniques by measuring the time to complete the visualisation task (fourth column of table 1). The results of one participant were removed since the measured time was more than 2 hours, which suggests that the user interrupted the study at least once. Figure 5 demonstrates similar results as for the effectiveness. The response time for the *blend*, *overlay* and *bump map* methods are similar and faster than the best performing masking method (in this case the coarse diagonal pattern). The worst performing method is the weaving pattern with a gradual shadow. The results suggest that vector field perception by mental interpolation of patterns is considerable less efficient than perception using other visual attributes such as shading, chromatic adaptation, and figure-ground perception.

5.3 Effectiveness and Aesthetics Ratings

The last two columns of table 1 show the user ratings for

	Singularities Correct (in%)		Singularities False (in%)		F_1-Score		Time (in s)		Effectiveness user rating (in range [-3,3])		Visual Attractiveness user rating (in range [-3,3])	
	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ
Blend	94.6	14.7	1.4	5.7	0.952	0.118	24.89	13.60	1.56	1.44	1.11	1.65
Overlay	89.5	20.3	0.6	3.1	0.919	0.179	24.62	14.45	1.78	1.36	1.06	1.45
Crisp bumpmap	88.7	19.7	1.4	5.7	0.917	0.173	25.33	15.58	1.43	1.50	0.87	1.61
Smooth bumpmap	84.7	21.9	4.0	11.2	0.870	0.192	24.46	12.28	1.43	1.53	0.94	1.67
Mask: f. diagonal	80.1	24.5	6.3	16.0	0.832	0.228	32.07	17.06	-0.02	1.67	-0.41	1.55
Mask: c. diagonal	81.0	22.7	2.6	8.2	0.864	0.183	27.02	13.86	-0.11	1.58	-0.48	1.52
Mask: vertical	81.5	24.1	3.0	11.1	0.858	0.201	30.49	21.21	-0.04	1.53	-0.41	1.64
Mask: horizontal	80.3	23.8	2.9	9.6	0.855	0.188	31.58	41.47	0.11	1.56	-0.39	1.45
Mask: f. checker	79.3	24.8	1.5	6.4	0.842	0.199	32.66	22.05	-0.56	1.57	-0.83	1.54
Mask: c. checker	77.6	26.7	5.6	15.8	0.822	0.240	33.47	22.14	-0.20	1.55	-0.57	1.45
Mask: long weave	78.2	25.8	4.4	16.1	0.835	0.230	31.84	14.37	-1.22	1.46	-1.06	1.55
Mask: solid weave	77.6	25.2	4.2	15.0	0.830	0.212	33.69	18.18	-1.09	1.43	-1.13	1.48
Mask: grad. weave	77.4	27.4	2.7	10.4	0.825	0.234	36.26	24.26	-1.37	1.41	-1.00	1.67

Table 1: Results of the user study. The rows represent the visualisation methods (c. and f. stand for coarse and fine patterns). The columns show the mean and standard deviation of respectively 1) fraction of critical points that were correctly classified, 2) fraction of clicks that were not on a critical point, 3) the calculated F-score, 4) the time in seconds it took participants to complete a visualisation, 5) the rating given for effectiveness, and 6) the rating given for visual attractiveness.

the effectiveness and visual attractiveness of each visualisation technique. The ratings that participants gave for the effectiveness are largely in line with the results for the error metric in equation 5, except that users have rated the *overlay* method highest for effectiveness ($\mu = 1.78$) followed by *blend* ($\mu = 1.56$), while the test results suggest that *blend* is the most effective method followed by *overlay*. The smooth and crisp bump map follow with $\mu = 1.43$. All of the masking methods were rated significantly lower on the effectiveness scale, the horizontal pattern being rated the most effective with $\mu = 0.11$, and the gradual weaving pattern the least with $\mu = -1.37$. Similar results are obtained for the aesthetics ratings.

The smooth bump maps were rated as more aesthetic than the coarse bump maps and as similarly effective, even though their actual performance was inferior to the coarse bump maps. In general methods with smooth colour changes obtained a higher rating for visual attractiveness.

A contradiction between user perception and performance was also observed for coarse and fine masking patterns. Whereas the fine checkerboard masking pattern performed better in terms of correctly and falsely classified singularities, users perceived the coarse masking pattern as more effective and visually more attractive.

The results suggest that aesthetics increases the perceived effectiveness of a visualisation technique.

6. CONCLUSIONS AND FUTURE WORK

Based on properties of the human visual system we have developed four classes of techniques for visualising two 2D vector fields simultaneously: *blending* (alpha blending of two LIC textures), *overlay* (replacement of selected pixels in one LIC texture with those form another LIC texture), *bump*
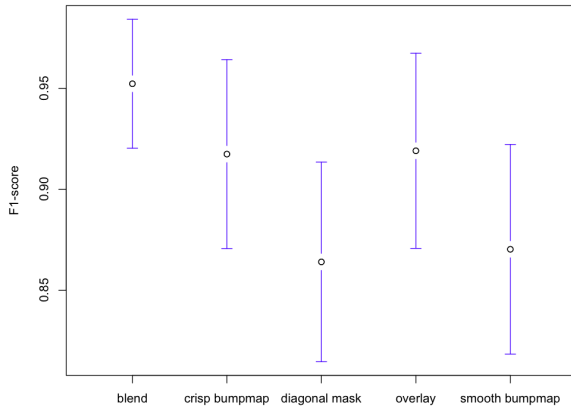

Figure 4: Mean and 95% confidence interval for the F-score. Both bump mapping techniques are shown as their performance differs significantly. For the masking technique, only the result for the best performing masking pattern (coarse diagonal) is shown.
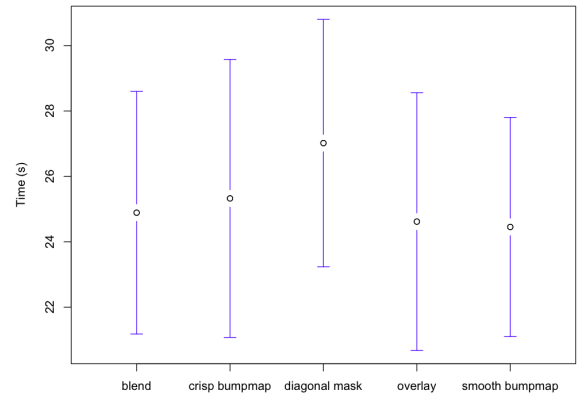

Figure 5: Mean and 95% confidence interval of the time taken per experiment in seconds. For the masking technique, only the results for the best performing pattern (coarse diagonal) are shown.

mapping (placement of an LIC texture on top of another LIC texture using shading differences simulating height variations) and *pattern masking* (division of an image into regions, each of which displays one field exclusively). The use of bump mapping and masking seems to be new in this application domain.

We evaluated the techniques with a user study and showed that all four concepts aid users with finding critical points in a vector field. Blending was most effective, followed by overlay and "crisp" bump maps. Similar results were obtained for efficiency. There were some contradictions between perceived and actual effectiveness. The results suggest that aesthetics increases the perceived effectiveness of a visualisation technique, and that visualisations with smooth colour changes are perceived as visually more attractive.

In future work we want to test the presented visualisations using a larger variety of visualisation tasks, and we want to investigate in more detail the relationship between aesthetics, perceived effectiveness and actual effectiveness.

7. REFERENCES

[1] J. Blaas, C. P. Botha, and F. H. Post. Interactive visualization of multi-field medical data using linked physical and feature-space views. In *Proceedings of EuroVis '07*, pages 123–130, 2007.

[2] B. Cabral and L. C. Leedom. Imaging vector fields using line integral convolution. In *Proceedings of SIGGRAPH '93*, pages 263–270. ACM, 1993.

[3] N. Cawthon and A. Moere. The effect of aesthetic on the usability of data visualization. In *Proceedings of InfoVis '07*, pages 637–648, July 2007.

[4] W. C. de Leeuw and J. J. van Wijk. Enhanced spot noise for vector field visualization. In *Proceedings of Visualization '95*, pages 233–239, 1995.

[5] M. Fairchild. *Color appearance models*, volume 3. Wiley, 2005.

[6] R. Fuchs and H. Hauser. Visualization of multi-variate scientific data. *Computer Graphics Forum*, 28(6):1670–1690, 2009.

[7] J. Helman and L. Hesselink. Representation and display of vector field topology in fluid flow data sets. *Computer*, 22(8):27–36, Aug. 1989.

[8] G. W. Humphreys, editor. *Understanding vision: an interdisciplinary perspective*. Blackwell, 1992.

[9] V. Interrante. Harnessing natural textures for multivariate visualization. *IEEE Comput. Graph. Appl.*, 20(6):6–11, Nov. 2000.

[10] B. Jobard, G. Erlebacher, and M. Y. Hussaini. Lagrangian-eulerian advection of noise and dye textures for unsteady flow visualization. *IEEE Transactions on Visualization and Computer Graphics*, 8(3):211–222, July 2002.

[11] R. M. Kirby, H. Marmanis, and D. H. Laidlaw. Visualizing multivalued data from 2d incompressible flows using concepts from painting. In *Proceedings of Visualization '99*, pages 333–340, Oct. 1999.

[12] R. V. Klassen and S. J. Harrington. Shadowed hedgehogs: a technique for visualizing 2D slices of 3D vector fields.

[13] J. Kniss, S. Premoze, M. Ikits, A. Lefohn, C. Hansen, and E. Praun. Gaussian transfer functions for multi-field volume visualization. In *Proceedings of IEEE Visualization '03*, pages 497–504, 2003.

[14] D. H. Laidlaw, R. M. Kirby, C. D. Jackson, J. S. Davidson, T. S. Miller, M. da Silva, W. H. Warren, and M. J. Tarr. Comparing 2d vector field visualization methods: A user study. *IEEE Transactions on Visualization and Computer Graphics*, 11(1):59–70, 2005.

[15] R. S. Laramee, H. Hauser, H. Doleisch, B. Vrolijk, F. H. Post, and D. Weiskopf. The state of the art in flow visualization: Dense and texture-based techniques. *Computer Graphics Forum*, 23(2):203–221, 2004.

[16] J. Meyer-Spradow, T. Ropinski, J. Mensmann, and K. Hinrichs. Voreen: A rapid-prototyping environment for ray-casting-based volume visualizations. *IEEE Comput. Graph. Appl.*, 29(6):6–13, Nov. 2009.

[17] A. V. Moere and H. Purchase. On the role of design in information visualization. *Information Visualization*, 10(4):356–371, Oct. 2011.

[18] R. Moorhead, C. Johnson, T. Munzner, H. Pfister, P. Rheingans, and T. Yoo. Visualization research challenges: a report summary. *Computing in Science Engineering*, 8(4):66–73, Aug. 2006.

[19] J. Reininghaus, C. Lowen, and I. Hotz. Fast combinatorial vector field topology. *IEEE Transactions on Visualization and Computer Graphics*, 17(10):1433–1443, Oct. 2011.

[20] L. G. Roberts. Machine perception of three-dimensional solids. Technical report, Massachusetts Institute of Technology, May 1963.

[21] H. R. Schiffman. *Sensation and Perception: An Integrated Approach*. John Wiley & Sons, 4^{th} edition, 1996.

[22] D. Stalling and H.-C. Hege. Fast and resolution independent line integral convolution. In *Proceedings of SIGGRAPH '95*, pages 249–256. ACM, 1995.

[23] R. Stokking, I. G. Zubal, and M. A. Viergever. Display of fused images: Methods, interpretation, and diagnostic improvements. *Seminars in Nuclear Medicine*, 33(3):219–227, 2003.

[24] T. Urness, V. Interrante, E. Longmire, I. Marusic, S. O'Neill, and T. W. Jones. Strategies for the visualization of multiple 2d vector fields. *IEEE Comput. Graph. Appl.*, 26(4):74–82, July 2006.

[25] C. J. van Rijsbergen. *Information Retrieval*. Butterworth, 2nd edition, 1979.

[26] J. J. van Wijk. Spot noise texture synthesis for data visualization. *Proceedings of SIGGRAPH '91*, 25(4):309–318, July 1991.

[27] C. Weigle, W. Emigh, G. Liu, R. Taylor, J. Enns, and C. Healey. Oriented sliver textures: A technique for local value estimation of multiple scalar fields. In *Proc. of Graphics Interface*, pages 163–170, 2000.

In *Proceedings Visualization '91*, pages 148–153. IEEE Computer Society Press, 1991.