Beyond the classical monoscopic 3D in graph analytics: an experimental study of the impact of stereoscopy

Nicolas Greffard*
DUKe Team - LINA
Polytech', University of Nantes

Fabien Picarougne†
DUKe Team - LINA
Polytech', University of Nantes

Pascale Kuntz‡
DUKe Team - LINA
Polytech', University of Nantes

ABSTRACT

Popular in the nineties, 3D visualization has since garnered much criticisms. While historically the vast majority of 3D visualization evaluations have been based on the classical Brunelleschi's perspective rendering, today new questions are raised with the recent access to high quality stereoscopic representations. Do these new interfaces offer improvements over the traditional 3D counterparts? This paper describes a comparative evaluation of monoscopic 3D and stereoscopic 3D for an important problem in graph visualization: the detection of communities. Our results show that stereoscopic 3D outperforms both 2D and monoscopic 3D in the task resolution. The 2D condition always yields the lower response times. Monoscopic and stereoscopic 3D obtain similar response times. A complementary study based on an original methodology reveals that these apparent similarities actually correspond to different interaction strategies -rotations around the graph-. Under the monoscopic condition, participants tend to explore the display space more thoroughly while under the stereoscopic condition, they seem to favor the adoption of a smaller set of viewpoints on the graph.

Index Terms: H.5.m [Information Interfaces and Presentation (e.g. HCI)]: Miscellaneous—

1 INTRODUCTION

3D has been considered as "a prejudicial epiphenomenon" by one of the major figures of the graph drawing community [8]. Even if pioneering works date back from the sixties [18], focus on 3D graph drawing was strong in the nineties when new 3D display hardware became available. Classical aesthetics were extended to 3D for the most employed models: 3D orthogonal drawing [2, 10], convex [4], straight line drawing [14]. Beside theoretical results for complexity and area or volume bounds [26, 14], various prototypes were developed. But it is clear that their expansion was short and, despite some recent attempts for applications (e.g. [29]), most of the current research still concerns 2D layouts. Even if 3D was shown to be preferred by users in some experiments in the context of information visualization [15], the main criticism for graph drawings is the lack of intelligibility of its renderings especially for dense and large graphs.

This criticism seems at variance with the intuitive idea that we can represent more data on a 3D layout by taking advantage of the depth perception. The mechanisms implied in depth perception have been widely studied in the literature [7, 19]. These so called depth cues come from different sources of information: ocular (accommodation and convergence), pictorial (shading, size and occlusion), dynamic (motion cues) and stereopsis (binocular disparity). While the relative importance of the different depth cues is still an open question, the current consensus seems to state that they are complementary, and that the quality of a 3D representation is strongly correlated with the effectiveness of the depth perception it conveys [3]. It is important to note that the vast majority of the 3D drawing approaches, which rely on tri-dimensional views based on the classical Brunelleschi's perspective rendering, ignores the binocular disparity.

However, recent technologies allow for stereoscopic representations of high quality which offer another experience of 3D. Compared to monoscopic 3D renderings, different experiments have highlighted several worthwhile properties: (1) it allows for a better estimation of relative depths [7], (2) it eases the visualization of specific objects through an important visual clutter [23], (3) it allows for the perception of camouflaged objects [31], (4) it eases the perception of curved surfaces and textures [16] and (5) it dampens the impact of several image degradations (noise, low contrast, low definition...) [23]. Even if some of them are still debated, perception in stereoscopic 3D is undoubtedly different from monoscopic 3D. Consequently, we claim that the "3D status" should be reconsidered in visual analytics especially for graph exploration.

1.1 Stereoscopy for graph exploration

A pioneering study about the impact of stereoscopy on graph analysis was carried out by Ware and Franck in 1996 [32]. The considered task was the identification of paths between two highlighted vertices, and the results in stereoscopy outperformed those in 2D for larger graphs (which were here bounded by 300 vertices). The experiment was reproduced ten years later [33] in an improved environment with a significantly higher resolution and screen refresh rate. The results showed that the parallax and stereoscopic depth cues seemed to be complementary and that they allowed users to carry out the task on graphs with up to 1000 vertices with less than 10% error. This is an order of magnitude greater than in 2D. More recently, Alper et al. [1] showed the benefits of the combination of 2D and stereoscopy (referred as 2.5 D) for neighbor and mutual neighbor identification and counting. By default, all the vertices are laid out in 2D, but when a group of vertices is selected, they are brought on a nearer plan in front of the user. Comparison of the 2.5D with a full stereoscopy showed no differences for the performance errors -except for a counting task of highlighted vertices- and showed lower response times for the 2.5 D.

All these experiments are concerned with local tasks (e.g. vertex or short path identification). However it is well-known in graph visual analytics [21] that "higher level" tasks (e.g. community detection) are involved in the exploration process. In a previous paper, we compared user performances for counting the communities on graph layouts obtained with a classical energy model and displayed in 2D, monoscopic 3D and stereoscopy [15]. The results showed that 2D outperformed the other renderings for "easy" graphs with a well-defined structure in communities, whereas for "complex" graphs with community overlapping, the best results

---

*e-mail: nicolas.greffard@univ-nantes.fr
†e-mail: fabien.picarougne@univ-nantes.fr
‡e-mail: pascale.kuntz@univ-nantes.fr

were obtained with stereoscopy. Whatever the result quality, there was no significant difference between the response times for the two 3D conditions and the shortest times were always reached in 2D. This first experiment has underlined the potential of stereoscopy for graph exploration. However, as a pioneering work, the experimental protocol was totally unsupervised and the classes of complexity ("easy", "medium", "complex"), which were discovered *a posteriori* during the performance analysis, had a heterogeneous statistical distribution of the graphs which makes the validation tricky.

1.2 Content of the paper

We here focus on the community detection task in a node-link layout as it is one of the major tasks in the exploration process for highlighting the overall organization of the graph topology. The main contributions of the paper are about two complementary questions.

The first question reexamines the debate "3D or not 3D" set in the early 2000s [5]: stereoscopic 3D, monoscopic 3D or not 3D? We have repeated the experience quoted above with a new protocol and new participants. We have only considered a greater number of graphs with a certain degree of community overlapping: 270 graphs were shown for the three conditions to 18 participants. And, for a pairwise comparison each layout was examined twice for the three conditions. The results confirm the previous ones: performances with stereoscopic 3D outperform those with both monoscopic 3D and 2D, and the time responses are similar for the two 3D conditions and smaller for the 2D one.

The second question is concerned with the user interactivity during the exploration process: are the interactions with the layouts similar in stereoscopic 3D and monoscopic 3D ? If the short response time can be easily explained for the 2D condition by an absence of interactions, we explore whether the time similarity for the two 3D conditions corresponds to similar interaction behaviors or not. Interaction is indeed a crucial component in a data mining process [17], but its understanding in data analytics is still in its infancy, and as far as we know, it has never been analyzed for graph exploration in stereoscopy.

In our experiment, interactions on the layouts were intentionally restricted to rotations around the graph barycenters. We filmed the participants, and the video recordings have suggested interaction differences between the two restitutions. In stereoscopic 3D, participants seem to focus their attention on restricted areas of the layouts, whereas in monoscopic 3D their exploration seems more uniform. To quantify these differences, we first compare the total distance traveled by the mouse pointer and the proportion of its direction changes during the layout explorations. Then, we compute heatmaps of the time spent by each user on the different graph viewpoints. Each heatmap is deduced from a triangulation of a bounded sphere of each graph layout. Our analysis confirms the observed behavioral differences between the two 3D conditions and the results suggest different strategies for the community detection task in monoscopic 3D and in stereoscopy.

The rest of the paper is organized as follows. Section 2 describes the experimental design. Section 3 compares the community detection efficiency of the three experimental conditions and Section 4 details the analysis of the interactions in 3D.

2 EXPERIMENTAL DESIGN

Due to the lack of a standard terminology, let us first precise the three experimental conditions considered in this paper. A 2D display is restricted to a classical Euclidean bi-dimensional space. It does not require any depth cue. A monoscopic 3D display (referred to as "mono 3D" in the following) results from a classical Brunelleschi linear perspective projection. A stereoscopic 3D display ("stereo 3D") is deduced from a mono 3D display to which a binocular disparity is added. The different depth cues involved in the evaluation are summarized in Table 1.

2.1 Graph layout and interaction

It is well-known that the interpretation of a graph closely depends on the layout algorithm [27]. The importance of the community detection in graph mining has led to the development of various algorithms (eg: [24]). However, many of them are intrinsically dependent on the bi-dimensionality of the embedding space, and their comparison with mono and stereo 3D restitutions would imply deeply different layouts. There is no doubt about the interest of such a comparison but it sets open questions which go beyond the objectives of this paper. Consequently, to limit bias as much as possible, we have selected a force-directed approach which is directly applicable to the three experimental conditions. The Fruchterman-Reingold algorithm is one of the most popular algorithms for creating straight-line drawings of undirected graphs which highlight communities [25, 9]. We first compute a layout of each graph in respectively R^2 and R^3 with the same algorithm with 10 000 iterations to limit any convergence bias. Then, we project the vertex coordinates on the screen -associated to a Cartesian coordinate system centered on the screen center- with a classical perspective projection. For the stereo 3D, two viewpoints are used with a slight horizontal shift to mimic the actual separation between the eyes according to the recommendation of [13]. We use a $120Hz$ display and the synchronization between each eye and the corresponding projection is achieved with shutter glasses allowing a smooth frame rate of $60Hz$ per eye.

User interactions with a layout are restricted to rotations around the center of gravity of the vertex positions (graph barycenter). More precisely, users can perform a motion on the graph by using an arcball technique [28] that allows them to spin their viewpoints (1° for 10 pixels) with a classical wireless mouse whose driver parameters (sensitivity, material acceleration off) are constant during the experimentation. This technique ensures that users keep the graph in sight contrary to free view techniques, and some experiments suggest that rotational object motion conveys better depth perception than translational object motion [6]. Other interactions are implemented in our experimental framework. But due to the absence of previous works, we have preferred to first focus on the most important interaction for the task.

2.2 Apparatus

The layouts are displayed in shades of white on a black background using an anti aliasing algorithm to improve the quality of the display. The visual restitution is on a white painted wall by an ACER H5360 3D projector ($2.30 \times 1.30m^2$ screen) with a resolution of 1280×720 pixels (view angle of 0.05° for a pixel in the center of the screen). Like most recent experiments with stereo 3D, we use active stereoscopy with Nvidia 3D Vision Shutter glasses which are worn throughout the experiment to maintain the same luminosity level for each layout. Computations ran on an Intel Core 2 Duo ($3.00Ghz$) E8400 processor, with 4 GB of RAM and an NVidia Quadro FX 3800 GPU.

Participant answers (number of detected communities) are entered with a touch screen tablet PC: different numbers are proposed (between 1 and 15, plus "don't know") and participants touch the corresponding number.

2.3 Graph databases

Roughly speaking, the objective was to generate graphs with more or less complex structures in communities. Different models were proposed in the literature [11]. Here we used a classical simple one which generates different topologies depending on the proportion of edges within (resp. between) the communities [12]. The generic model $G(k; nv; p_{int}; p_{ext})$ depends on four parameters : the number

Depth Cue	3D Stereo	3D Mono	Implementation
Binocular Disparity	Yes	No	Two horizontally separated parallel viewpoints along with active stereo glasses
Motion Cues	Yes	Yes	Graph rotation in the display space with a mouse
Relative Size*	Yes	Yes	Linear perspective projection (OpenGL)
Occlusion	Yes	Yes	3D rendering and Fruchterman layout in 3D allowing occlusion
Shading*	Yes	Yes	Gouraud shading with a light source located above the viewpoint

Table 1: List of the main depth cues involved in the evaluation. (*) A 16× anti aliasing algorithm enhances the quality of these cues.

k of *a priori* communities, the number nv of nodes per community, the probability p_{int} (resp. p_{ext}) of an edge between two nodes belonging to the same community (resp. different communities).

In a previous work, we tested various parameter combinations [15]. We here restrict ourselves to combinations which generate a class of graphs with a high probability of community structure which is not obvious to detect. We generated 135 graphs from all the combinations of the following parameters: $k \in [4 \dots 12]$, $nv = 30$, $p_{int} \in [0.6; 0.7; 0.8]$ and $p_{ext} \in [0.04; 0.05; \dots; 0.08]$. The parameter nv is constant to ensure that the task is only affected by the complexity ratio $pext/pint$ (overlapping factor) and the number of communities k. Class size variations should be explored in the next future but our initial trials have shown that the impact of community overlapping is more important.

2.4 Experimental procedure

The 135 graphs were randomly distributed in 9 non overlapping subsets of 15 graphs. Each subset was viewed under the 3 conditions (2D, mono 3D, stereo 3D) by two participants. Consequently, there were $9 \times 2 = 18$ participants who each analyzed $15 \times 3 = 45$ layouts.

The 2D condition was always viewed between the two 3D conditions, and to limit bias the order was inverted for the participants having the same subset (stereo 3D-2D-mono 3D and mono 3D-2D-stereo 3D).

Participants were asked to estimate the number of communities displayed as fast as possible. At the end of each condition, a message appeared to indicate the end and the beginning of the following one. The average task completion time was 15 minutes (28 minutes at most).

Before the experiment, a few questions were asked to participants to gather their experience on visualization. A description sheet was handed out to briefly explain the experimental process, and a quick demo presented the different viewing methods with an easily readable layout. Then, to get familiar with the system, the participants had to go through a training session of 3 layouts of increasing complexity for each condition. We took advantage of this training session to ensure that each participant had an acceptable stereoscopic acuity by asking, for several pairs of nodes, which were the closer and farther away.

2.5 Participants

Aged from 20 to 44 with a right handed use of the mouse, all the participants (13 males, 5 females) were computer science students or researchers. Four subjects had never visualized any stereoscopic material, and two were not familiar with 3D software such as video games. Fortuitously, no left-handed people took part to the experiment but it allowed us to discard some potential bias due to differences in visual perception between left-handed and right-handed people [20] which might impact depth perception.

3 COMMUNITY DETECTION EFFICIENCY AND RESPONSE TIMES

For each layout, the error in the community detection is defined by the difference $E_k = |k_p - k_{ans}|$ between the *a priori* number of communities k_p from the graph generation model and the participant answer k_{ans}.

Over all the trials, the average error of the 3 conditions are: (2D) $E_k2D = 2.7$, (mono 3D) $E_kM = 2.5$, (stereo 3D) $E_kS = 1.8$. And there is a statistically significant difference between the three conditions (with a Kruskal-Wallis test: $X^2(2) = 18.96$, $p_{value} < 0.0001$). A pairwise post-hoc analysis with a Bonferroni correction shows that the 2D and the mono 3D conditions are associated with a significantly higher E_k error than the stereo 3D condition (resp. $W(269) = 79770$, $z_{value} = 4$, $p_{value} < 0.0001$, and $W(269) = 24560$, $z_{value} = 4.16$, $p_{value} < 0.0001$). The comparison between the 2D and mono 3D conditions yields no statistical significance.

The response times are the following: (2D) 10sec; (mono 3D) 17sec, (stereo 3D) 17.3sec. These response times are significantly different (with a Kruskal-Wallis test: $X^2(2) = 139.5$, $p_{value} < 0.0001$) and they are significantly shorter for the 2D condition whereas there is no statistical difference between the two 3D conditions (with a Bonferroni correction: $W(269) = 55062$, $z_{value} = -9.91$, $p_{value} < 0.0001$ for 2D vs stereo 3D ; $W(269) = 53949$, $z_{value} = -10.52$, $p_{value} < 0.0001$ for 2D vs mono 3D ; $W(269) = -3098$, with $z_{value} = -1.21$, and $p_{value} = 0.22$ for mono 3D vs stereo 3D).

These results lead to the following question: do the similar response times in the two different 3D conditions correspond to similar interaction behaviors or not? In the following, we compare how the participants explore the different regions of the graph layouts in 3D. Let us remark that due to our probabilistic model of graph generation, a restrictive subset of layouts (63) *a posteriori* appeared easy for the task with a response time smaller than 3s for one of the two 3D conditions. We do not consider them for the interaction study. Note that their cancellation from the database does not change the conclusion of the response time comparison (TM = 18.4sec, TS = 18.6sec).

4 INTERACTIONS IN 3D

As recalled by Keim et al [17], “visual analytics is more than only visualization”. It combines visualization with data analysis and human factors, and interaction plays a key role in the data exploration process. There is a huge literature in HCI dedicated to the analysis of interactions with a mouse. But, in numerous studies, the mouse trajectories are restricted to point-and-select tasks and the optimal trajectories are known. However, in data exploration, the optimal sequence of interactions to solve a task is most of the time *a priori* unknown. As shown by Keim and others, visual analytics consists of a feedback loop where interactions alter the visual restitution to enhance the user understanding of the data (see visual analytics pipeline [30]). Therefore, the study of interactions is problem and data specific. But, to the best of our knowledge, there is no consensual framework for studying interactions in the context of visual analytics. Thus, in this section we present an original approach for the community detection problem in 3D graph layouts. Two interaction aspects are analyzed: the mouse pointer movements and the viewpoints on the layouts taken by the users.

4.1 Pointer movements

We compare the total distance traveled by the pointer and the proportion of direction changes. For both 3D condition and each layout we recorded the coordinates of the mouse pointer on the screen every millisecond. The pointer is not visible to the users. Let us denote by $O = (0,0)$ the origin of the screen center. For the measure computation we have retained the coordinates of the position changes. The starting position ($x_0 = 0, y_0 = 0$) of the pointer is centered in O and (x_i, y_i), $i = 1 \dots n$, denotes the i-th position of the pointer s.t. $(x_i, y_i) \neq (x_0, y_0)$ with n the number of changes between two consecutive positions of the mouse. Once a pointer movement (x_i, y_i) is recorded, the pointer position is reset to (x_0, y_0) to allow infinite movements along any direction (which would otherwise be bounded by the screen resolution).

The total distance Dt covered by the pointer is defined by $Dt = \sum_{i=0}^{n} \sqrt{x_i^2 + y_i^2}$. Results for the two 3D conditions are given in Table 2. The difference is not statistically significant ($W(206) = 44893$, $z_{value} = 1.5937$, $p_{value} = 0.1$). This observation is consistent with the similarity of the response times.

Condition	*min*	*max*	*avg*	*std*
Stereo 3D	682	45200	7398	6371
Mono 3D	564	38143	7651	5719

Table 2: Statistics of the total distance (Dt) measure.

The so-called “alternation ratio” (AR) evaluates the proportion of direction changes. It is inspired by the “movement direction changes” and the “orthogonal direction changes” [22] commonly used in HCI to evaluate distortions relative to an optimal trajectory. Here the presence/absence of a direction change is defined by a function ang_i of the angle between the line passing through the origin and the pointer position and the x-axis :

$$ang_{i>1} = \begin{cases} 1 & \text{if } sign(cos(x_i, y_i)) \neq sign(cos(x_{i-1}, y_{i-1})) \\ 1 & \text{if } sign(sin(x_i, y_i)) \neq sign(sin(x_{i-1}, y_{i-1})) \\ 0 & otherwise \end{cases} \quad (1)$$

The alternation ratio is the ratio between the number of direction changes and the number of position changes: $AR = \frac{\sum_{i=1}^{n} ang_i}{n-1}$. Results are given in Table 3. The difference is significant ($W(206) = 35157$, $z_{value} = -6.4035$, $p_{value} < 0.001$): position changes are more often associated with direction changes in 3D stereo than in 3D mono. Under the 3D mono condition, participants tend to move the pointer in a “linear fashion” while under the 3D stereo condition, the pointer positions are prone to backtracking. This observation suggests that, as observed on the videos, the exploration of the graphs differ between the conditions.

4.2 Viewpoint analysis

In order to compare the behaviors at a macroscopic level, we analyze the time spent by users on the different parts of the 3D layouts. The spatial time distribution is represented by a heatmap computed from a triangulation of a bounding sphere of the graph layout.

More precisely, for a given layout, we first compute a triangulation of its minimal radius bounding sphere with a fixed number of equally sized triangles (set here to 8000 which is a good trade-off between the required precision and the computation time) (figure 2a). Then, for each viewpoint, we classically deduce the “visible triangles” which are the triangles on the sphere belonging to the cone defined by the viewpoint position and the tangent to its direction on the sphere (figure 2b). The times spent on the viewpoint positions are added for each visible triangle. The distribution of the time values on the triangle set is represented by a heatmap (figure 2): the red (resp. blue) zones correspond to areas of the layout which are often (resp. rarely) in the field of view.

Condition	*min*	*max*	*avg*	*std*
Stereo 3D	0.14	0.79	0.38	0.11
Mono 3D	0.07	0.62	0.31	0.10

Table 3: Statistics of the alternation ratio (AR) measure.

The spatial Shannon’s entropy of the time values allows the highlighting of the presence of macroscopic patterns: $H = -\sum_{j=1}^{m} p(t_j) logp(t_j)$ with m the number of distinct observed time values t_j and $p(t_j)$ their observed frequencies on the triangulated sphere. Highest values of the entropy correspond to a uniform distribution. The average entropies on the 207 retained layouts are respectively $HM = 5.12$ for mono 3D and $HS = 4.67$ for stereo 3D and this difference is statistically significant (with a Wilcoxon ranksum test: $W(206) = 14939$, $z_{value} = 8.66$, $p_{value} < 0.0001$). This result seems to confirm our initial hypothesis deduced from videos: user observations are more uniformly distributed in mono 3D than in stereo 3D.

A further analysis of the individual heatmaps helps to precise the different interaction strategies: in mono 3D, participants tend to explore the display space more thoroughly by adopting numerous different viewpoints, whereas in stereo 3D, they tend to focus on a few distinct areas. Typical behaviors are illustrated on figure 2. The time value histogram for mono 3D shows a trend towards a uniform distribution which corresponds to a homogeneous coloration in the heatmap except for a blind area (dark blue). The histogram for stereo 3D highlights a large set of smaller time values and a peak of high values, respectively associated on the heatmap with a large blue area and a quite small red one. The AR measure mentioned in section 4.1 corroborates this analysis: a large AR value (stereo 3D) is associated with more directional changes which consequently lead successive viewpoints to gravitate around the same area. A small RA value (mono 3D) is associated with less directional changes, meaning that successive viewpoints tend to get farther away in a monotonous fashion.

These results led us to hypothesize that participants are more efficient in the stereo 3D condition by adopting a smaller set of viewpoints. They are probably able to gather the information required to identify the communities by looking through the visual clutter. This hypothesis is in accordance with the work of Merritt [23] who showed that stereoscopy can help distinguish objects through a visual clutter.

In addition, the equality of the total distances suggests that there is a non negligible amount of movements which have no substantial effect on the viewpoint positions. These could be attributed to an attempt by the participants to maintain a constant motion to stimulate the motion parallax to further enhance their depth perception.

5 CONCLUSION

The results presented in this paper confirm that stereoscopic 3D outperforms monoscopic 3D in a community detection task where the graphs are subjects to large amounts of community overlapping. Moreover, we show that the performance gain of stereoscopic 3D

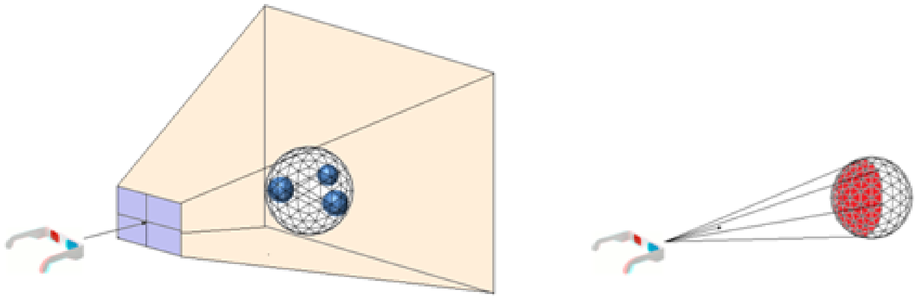

(a) The triangulation of the bounding sphere of the graph (represented from the participant's viewpoint)

(b) For each viewpoint taken by the participant, we increment a value associated to each visible triangle by the duration spent on this viewpoint.

Figure 1: Heatmap generation

against monoscopic 3D does not come with a trade-off between efficiency and response times. An analysis of the interactions involved in the task reveals that these apparently similar response times are actually associated with different interaction behaviors. Under the monoscopic condition, participants tend to explore the display space more thoroughly while under the stereoscopic condition, they seem to favor the adoption of a smaller set of viewpoints on the graph.

This observation raises new questions about graph drawings in stereoscopic 3D. The fact that a smaller set of viewpoints is required to grasp the structure of the graph in our task suggests that it might be possible to find an "optimal" viewpoint. 3D layout algorithms usually allow the drawing to expand equally onto the 3 dimensions. As a result, the vertices distribution across each dimension are roughly equivalent. With a well-defined viewpoint, the optical axis (depth) could be used to highlight particular structures (in a similar fashion to the work of Alper et al. [1]) by allowing the vertices to span a greater space across that dimension.

Generally speaking, we believe that this study provides strong evidences that stereoscopic 3D should be added to the 2D contenders in the everlasting "3D or not 3D?" [5] debate. In addition, we trust that extensive studies on the interactions are primordial to better understand the underlying cognitive processes of visual analytics which could, in turn, allows us to refine our visualization techniques. But such breakthrough would first require the formalization of a global framework to allow for the comparison between the visual representations and interaction techniques for high-level tasks encountered in data exploration.

REFERENCES

[1] B. Alper, T. Hollerer, J. Kuchera-Morin, and A. Forbes. Stereoscopic highlighting: 2d graph visualization on stereo displays. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2325–2333, 2011.

[2] G. D. Battista, M. Patrignani, and F. Vargiu. A split&push approach to 3d orthogonal drawing. In *Proceedings of the 6th International Symposium on Graph Drawing*, GD '98, pages 87–101, London, UK, 1998. Springer-Verlag.

[3] I. Cho, W. Dou, Z. Wartell, W. Ribarsky, and X. Wang. Evaluating depth perception of volumetric data in semi-immersive vr. In *Proceedings of the International Working Conference on Advanced Visual Interfaces*, AVI '12, pages 266–269, New York, NY, USA, 2012. ACM.

[4] M. Chrobak, M. T. Goodrich, and R. Tamassia. Convex drawings of graphs in two and three dimensions (preliminary version). In *Proceedings of the twelfth annual symposium on Computational geometry*, SCG '96, pages 319–328, New York, NY, USA, 1996. ACM.

[5] A. Cockburn and B. McKenzie. 3d or not 3d?: evaluating the effect of the third dimension in a document management system. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '01, pages 434–441, New York, NY, USA, 2001. ACM.

[6] V. Corinilleau-Prs and J. Droulez. The visual perception of three-dimensional shape from self-motion and object-motion. In *Vision Research*, volume 34, pages 2331–2336, 1994.

[7] J. Cutting. How the eye measures reality and virtual reality. *Behavior Research Methods, Instrumentation, and Computers*, 29:29–36, 1997.

[8] P. Eades. On the future of graph drawing. In *Invited talk, 18th International Symposium on Graph Drawing*, 2010.

[9] P. Eades and X. Lin. Spring algorithms and symmetry. *Theoretical Computer Science*, 240:379–405, 1999.

[10] P. Eades, A. Symvonis, and S. Whitesides. Three-dimensional orthogonal graph drawing algorithms. *Discrete Applied Mathematics*, 103(1-3):55–87, 2000.

[11] S. Fortunato. Community detection in graphs. *Physics Reports*, 486(3):75–174, 2010.

[12] J. Garbers, H. J. Prmel, and A. Steger. Finding clusters in vlsi circuits. In *Proceedings of ICCAD'90*, pages 520–523, 1990.

[13] S. Gateau and R. Neuman. Stereoscopy from xy to z. In *Courses of SIGGRAPH Asia*, 2010.

[14] E. D. Giacomo, G. Liotta, and H. Meijer. Computing straight-line 3d grid drawings of graphs in linear volume. *Computational Geometry*, 32(1):26 – 58, 2005.

[15] N. Greffard, F. Picarougne, and P. Kuntz. Visual Community Detection: An Evaluation of 2D, 3D Perspective and 3D Stereoscopic Displays. In M. van Kreveld and B. Speckmann, editors, *Proccedings of the 19th International Symposium on Graph Drawing*, LNCS 7034, pages pp. 215–225, Pays-Bas, 2011. Springer.

[16] N. Holliman. 3d display systems. Technical report, Department of Computer Science, Univ. Durham, 2005.

[17] D. A. Keim, F. Mansmann, J. Schneidewind, J. Thomas, and H. Ziegler. Visual data mining. chapter Visual Analytics: Scope and Challenges, pages 76–90. Springer-Verlag, Berlin, Heidelberg, 2008.

[18] A. Kolmogorov and Y. Barzdin. About realization of sets in 3-dimensional space, problems in cybernetics., 1967.

[19] M. S. Landy, L. T. Maloney, and M. J. Young. Psychophysical estimation of the human depth combination rule. volume 1383, pages 247–254. SPIE, 1991.

[20] N. Le Bigot and M. Grosjean. Effects of handedness on visual sensitivity in perihand space. *PloS one*, 7(8):e43150, 2012.

[21] B. Lee, C. Plaisant, C. S. Parr, J.-D. Fekete, and N. Henry. Task taxonomy for graph visualization. In *Proceedings of the 2006 AVI workshop on BEyond time and errors: novel evaluation methods for information visualization*, BELIV '06, pages 1–5, New York, NY, USA, 2006. ACM.

[22] I. S. MacKenzie, T. Kauppinen, and M. Silfverberg. Accuracy measures for evaluating computer pointing devices. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '01, pages 9–16, New York, NY, USA, 2001. ACM.

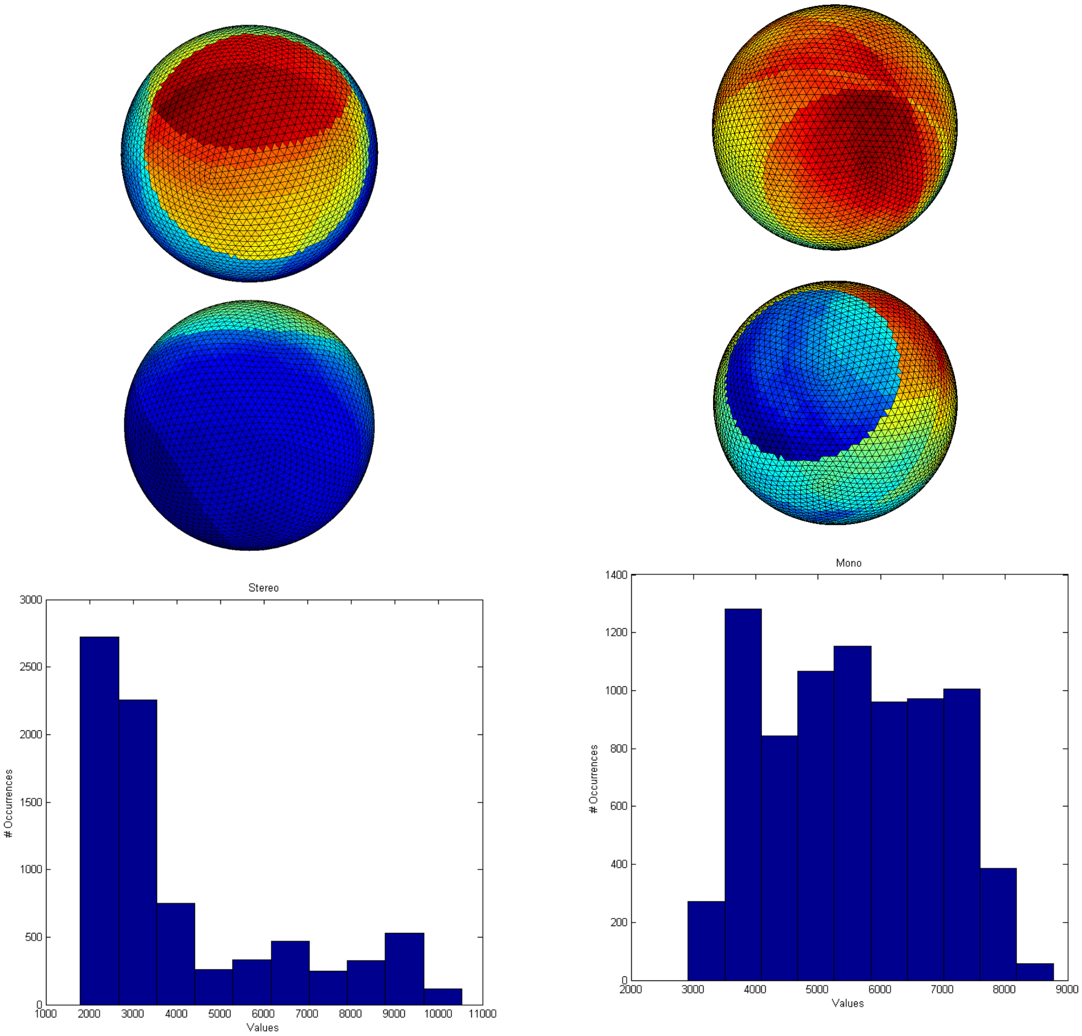

(a) Front (top) and back (bottom) views of a 3D heatmap for the stereo 3D condition. Two main points of views sharing a common area of interest can be seen while a large area remains less explored. This behavior is also apparent on the histogram of the time values associated with the heatmap with a peak of small values.

(b) Front (top) and back (bottom) views of a 3D heatmap for the mono 3D condition. The whole space has been thoroughly explored on the front side. Solely a small area on the back side has been explored significantly less. This behavior is also apparent on the histogram of the time values associated with the heatmap as it tends towards a uniform distribution of the values.

Figure 2: Examples of the resulting 3D heatmap and their corresponding histograms for the two 3D conditions.

[23] J. O. Merritt. Evaluation of stereoscopic display benefits. in: Introduction to stereoscopic displays and applications. In *SPIE The International Society for Optical Engineering*, 1991.

[24] A. Noack. An energy model for visual graph clustering. In *Proceedings of the 11th International Symposium on Graph Drawing (GD 2003)*, pages 425–436. Springer, 2004.

[25] A. Noack. Modularity clustering is force-directed layout. *Phys. Rev. E*, 79:026102, Feb 2009.

[26] M. Patrignani. Complexity results for three-dimensional orthogonal graph drawing. *Journal of Discrete Algorithms*, 6(1):140–161, 2008.

[27] H. C. Purchase. *Experimental human-computer interaction: a practical guide with visual examples*. Cambridge University Press, 2012.

[28] K. Shoemake. Arcball: a user interface for specifying three-dimensional orientation using a mouse. In *Proceedings of the conference on Graphics interface '92*, pages 151–156, San Francisco, CA, USA, 1992. Morgan Kaufmann Publishers Inc.

[29] A. Teyseyre and M. Campo. An overview of 3d software visualization. *IEEE Trans. on Visualization and Computer Graphics*, 15(1):114–135, 2009.

[30] J. J. Van Wijk. The value of visualization. In *Visualization, 2005. VIS 05. IEEE*, pages 79–86. IEEE, 2005.

[31] S. G. Wardle, J. Cass, K. R. Brooks, and D. Alais. Breaking camouflage: Binocular disparity reduces contrast masking in natural images. *Journal of Vision*, 10(14), 2010.

[32] C. Ware and G. Franck. Evaluating stereo and motion cues for visualizing information nets in three dimensions. In *ACM Transactions on Graphics*, volume 15, pages 121–139. 1996.

[33] C. Ware and P. Mitchell. Visualizing graphs in three dimensions. In *ACM Transactions on Applied Perception*, volume 5, pages 2–15. 2008.