

Fauxvea: Crowdsourcing Gaze Location Estimates for Visualization Analysis Tasks

Steven R. Gomez, Radu Jianu, Ryan Cabeen, Hua Guo, and David H. Laidlaw, *Fellow, IEEE*

Abstract—We present the design and evaluation of a method for estimating gaze locations during the analysis of static visualizations using crowdsourcing. Understanding gaze patterns is helpful for evaluating visualizations and user behaviors, but traditional eye-tracking studies require specialized hardware and local users. To avoid these constraints, we developed a method called Fauxvea, which crowdsources visualization tasks on the Web and estimates gaze fixations through cursor interactions without eye-tracking hardware. We ran experiments to evaluate how gaze estimates from our method compare with eye-tracking data. First, we evaluated crowdsourced estimates for three common types of information visualizations and basic visualization tasks using Amazon Mechanical Turk (MTurk). In another, we reproduced findings from a previous eye-tracking study on tree layouts using our method on MTurk. Results from these experiments show that fixation estimates using Fauxvea are qualitatively and quantitatively similar to eye tracking on the same stimulus-task pairs. These findings suggest that crowdsourcing visual analysis tasks with static information visualizations could be a viable alternative to traditional eye-tracking studies for visualization research and design.

Index Terms—Eye tracking, crowdsourcing, focus window, information visualization, visual analysis, user studies

1 INTRODUCTION

THE goal of this work is to make it easier to understand where people look in visualizations during analysis tasks. This gaze information could be helpful for improving visualization designs. An example application is verifying that increasing the size of a chart's guide marks, or repositioning them, increases how frequently a person gazes at these marks. Increased visual attention to these marks might affect how accurately a person interprets data in the chart. In addition, finding where people look could help researchers better understand visual analysis behavior and possibly improve their ability to identify low-level analysis tasks with visualizations, like finding extrema in a chart.

We present an evaluation of a crowd-based method for estimating gaze fixations for visualization research. The method builds on an earlier technique called the Restricted Focus Viewer (RFV) [1], an image viewer that simulates movement of the fovea by blurring the image and requiring viewers to deblur regions using the cursor. Essentially, the RFV requires a person to make manual interactions that are easily recorded and correspond to the areas of the image she wants to visually decode. We constructed and evaluated a Web-based version of the RFV called Fauxvea, with incremental improvements in the design of the focus window and data capture. The method enables large-scale gaze estimation experiments and can be used to crowdsource the

production of heatmaps showing spatial distributions of gaze fixations. Recently, Jiang et al. used a similar approach to crowdsource saliency maps for free-form viewing on natural images [2]. Here we evaluate the method in different experiments to test whether it is feasible for identifying where people look in static information visualization (infovis) charts during typical analysis tasks.

We demonstrate the method using workers on Amazon Mechanical Turk (MTurk) and evaluate how well gaze location estimates from these workers compare to eye-tracking data in corresponding laboratory studies. First, we compared Fauxvea fixation estimates to eye tracking from 18 participants for three common types of infovis charts (scatter plots, bar charts, and node-link diagrams) and photographs. Second, we compared the Fauxvea estimates with ones predicted by participants with expertise in vision and eye tracking, and found that even experienced individuals cannot predict fixations as consistently as data collected from a study using Fauxvea. Third, we reproduced findings from an existing study on tree layouts from Burch et al. [3] that involves a more complex visual analysis task than in the first experiment. In these experiments, we find that gaze locations on the visualizations by online participants are qualitatively and quantitatively similar to gazes from the eye-tracking studies.

The contributions of this work are threefold:

- S. R. Gomez, R. Cabeen, H. Guo, and D.H. Laidlaw are with the Department of Computer Science, Brown University, Providence, RI 02912. E-mail: {steveg, cabeen, huag, dhl}@cs.brown.edu.
- R. Jianu is with the School of Computing and Information Sciences, Florida International University, Miami, FL 33199. E-mail: rdjianu@cis.fiu.edu.

Manuscript received 1 Dec. 2014; revised 21 Oct. 2015; accepted 23 Oct. 2015. Date of publication 19 Feb. 2016; date of current version 4 Jan. 2017.

Recommended for acceptance by J. Heer.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TVCG.2016.2532331

- qualitative and quantitative evaluations of a method for crowdsourcing gaze location estimates during visualization analysis tasks; results from these experiments show fixation estimates are comparable to eye-tracking data on basic infovis analysis tasks (Section 4);
- an evaluation of how well experts can self-assess where others will gaze during visualization analysis tasks; we compare self-assessment to data collected using Fauxvea (Section 5);

- reproduced findings about visual exploration on tree layouts using the method instead of eye tracking for a more complex graph analysis task (Section 6).

Finally, we discuss limitations of the method and present opportunities for developing models of gaze that factor in both visualization stimuli and analysis tasks.

2 RELATED WORK

In this section, we describe how our proposed method relates to earlier process-tracking techniques, as well as other approaches for estimating gaze without an eye tracker.

2.1 Focus-Window Methods

The idea of restricting visual information to the location of a pointer and tracking its location has existed for decades. An early example is the MOUSELAB system, which was aimed at tracking a study participant's cognitive process during decision tasks involving information on a computer display [4]. In this system, boxes containing information appeared blank until the participant moved the mouse into one, which would reveal the information in that box.

Our method is more closely related to the Restricted Focus Viewer [1], an image viewer that requires the user to move the cursor in order to focus regions of the image. Unlike MOUSELAB, the RFV works with images that do not have predefined boundaries of information, so the mouse can be moved to focus any part of the image, and the image outside of the focus window is blurred. Cursor movements can be recorded and replayed as a proxy for actual gaze fixations. Fauxvea adapts this technique for the Web browser, with design changes that make it easier to use. Most significantly, the experiments we performed demonstrate that gaze estimates collected from online crowd workers – even in uncontrolled computing environments – are close to eye tracking for the visualization tasks we studied.

The aim of this paper is validating the RFV for visualization tasks using crowd workers. A similar cursor-based crowdsourcing method was recently used by Jiang et al. to collect a database of saliency annotations called SALICON [2]. The researchers focused on collecting saliency information on a large scale for natural images during unstructured *free viewing*, and found that crowdsourced data collected using the approach outperforms popular saliency models. Instead, we evaluate the approach for structured infovis stimuli and analysis tasks, for which *visualization-task-dependent* models of saliency do not yet exist, to the best of our knowledge. Experiments 1-3 evaluate the crowd-powered RFV for visualization from different angles, demonstrating its usefulness and comparing and contrasting gaze location estimates with controlled eye-tracking studies.

Previous evaluations of the RFV have validated the technique in controlled laboratory experiments and identified some general limitations. Blackwell et al. found that when people evaluated causal motions in diagrams of pulley systems, gaze patterns estimated by the RFV were similar to patterns collected from an eye-tracking study on the same stimuli [1]. Bednarik and Tukiaine [5] studied how participants in a controlled eye-tracking study used a Java software debugging environment with the RFV. They found that blurring affected how some users switched gaze between areas

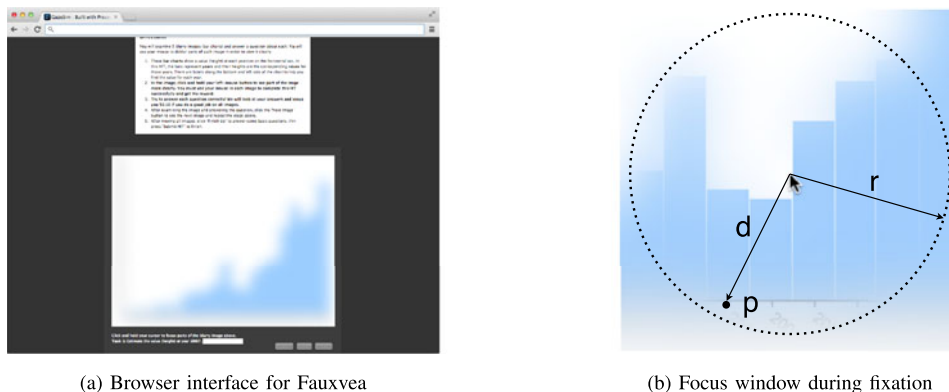
on the screen differently compared to eye tracking, but this behavior did not affect task performance; participants were able to extract the same information using the RFV as with a normal image viewer. Stimuli like coding environments or pulley diagrams differ from typical infovis charts in how directly they encode information, so we are motivated to study the focus-window method in this context. We find supporting evidence that the approach works even in realistic visualization scenarios involving moderately complex visual representations and tasks. For instance, as described in Section 6.3, we found similarities between eye tracking and our crowd-powered RFV (Fauxvea) in how people switched between areas of interest (AOIs) in tree diagrams during a graph analysis task.

Crowdsourcing might also help users of the RFV to select appropriate parameters for their experiments. Jones and Mewhort [6] found that badly chosen blur levels outside the focus window can affect scan paths. Earlier works have proposed guidelines for setting blur levels [1], [7], but it remains a challenge to apply these. Because picking blur parameters depends on the stimulus-task and not necessarily on individual differences, blur levels for stimuli and tasks could be tested rapidly, inexpensively, and at scale in pilot studies on MTurk. In our experiment, we chose a reasonable blur level after rapid testing on MTurk.

2.2 Estimating Gaze on the Web

User interfaces have been developed to collect gaze estimates using a Web browser, but these have focused on domains outside of visualization. Much of this work is based on findings about the relationship between gaze and cursor movements (e.g., [8]), which are easy to track in web applications. Other studies using web search tasks in lab settings have identified specific types of eye-mouse coordination patterns [9], [10] and demonstrated the predictive power of cursor actions for estimating gaze [11]. Huang et al. performed an eye-tracking study relating cursor activity to gaze in search engine results pages (SERPs), then followed up with a large-scale study of cursor tracking that linked cursor movements and results examination behaviors in SERPs [12]. Our method also uses cursor actions to predict gaze, but we make use of deliberate cursor presses and releases rather than hover locations in order to measure start and end times for gaze fixations.

A web-based system related to a moving focus window is ViewSer, which helped researchers study how remote users examine SERPs without eye-tracking [13]. The interface blurs DOM elements in the page corresponding to search results, and deblurs results when users hover over them with the cursor. One limitation of this method for evaluating visualization analysis is that it can only deblur entire DOM elements. Even if visualization components do correspond with DOM elements, e.g., using D3 [14], the size of the the deblurred component might be large enough that the hovered location does not reflect where the user is gazing at a useful level of precision. With Fauxvea, the deblurring area is based on a simple model of the human fovea. Because the focus region becomes more blurred away from its center, the user must press the cursor near the pixels she wants to see clearly. Therefore, the precision of Fauxvea for



(a) Browser interface for Fauxvea

(b) Focus window during fixation

Fig. 1. Left: Fauxvea interface showing analysis task instructions, the blurred image viewer, and an input field for the task answer. This example shows a bar chart task from Experiment 1. Right: The focus window during a Fauxvea fixation. All pixels outside radius r are fully blurred, and pixels inside are blended between the blurred image and the focused one. The blend ratio for each pixel p is proportional to its distance d from the cursor location.

estimating gaze is linked to a parameter in the model and is not dependent on the way DOM elements are rendered.

Gaze locations in video frames were crowdsourced using a novel video interface. Rudoy et al. asked workers on Amazon Mechanical Turk (“Turkers”) to watch videos then report text codes that randomly appeared on the display in different parts of the image [15]. This allowed researchers to look up an approximate region each Turker was gazing at on a given frame based on the specific code he reported. One limitation of this technique is achieving high spatial resolution of gaze estimates. Codes cannot be so close to one another that a person cannot identify them quickly. Fauxvea has a similar limitation: users might gaze at locations that are within the focus region without bothering to refocus precisely where they are attending. In practice, we find that users like to refocus directly on interesting parts of the focus region.

2.3 Crowdsourcing Visual Analysis Tasks

Recently, crowdsourcing platforms have been used to evaluate visualizations with scalable, non-expert populations. Some notable examples include Heer and Bostock’s reproduction of classic graphical perception results [16], Kong et al.’s study on TreeMap design [17], evaluations by Kosara and Ziemkiewicz of visual metaphors and percentage value reading [18], and Ziemkiewicz et al.’s study of the effects of individual differences on visualization performance [19]. This line of work has provided valuable examples and guidelines for crowdsourcing visualization analysis tasks, but they largely focus on evaluating the speed and accuracy of Turkers’ task performance as outputs. Instead, we estimate gaze locations with Fauxvea using additional data from the task execution, e.g., cursor presses that facilitate performing a task.

3 DESIGN AND METHODS

We adapted the RFV into a web-based application called Fauxvea that estimates gaze fixations during visualization analysis tasks without using eye-tracking hardware. By design, tasks on the RFV interface can be performed in parallel by remote users, or crowdsourced as human intelligence tasks (HITs) on MTurk as Jiang et al. demonstrated [2]. The two main objectives for the Fauxvea prototype were: (1)

enabling the collection of data that is comparable to eye tracking during analysis of a static visualization; (2) enabling scalable experiments with remote users, like crowdsourced participants or remote domain experts who are unavailable for local eye-tracking studies.

Comparable to eye tracking. The goal of this work is to make gaze data and metrics more accessible to visualization designers and evaluators. We are mainly interested in the location and duration of fixations, where the eye is focused in the field of view and has the highest visual acuity, during visualization analysis. The “eye-mind” hypothesis [20] suggests this data identifies areas of a visualization that a person cognitively processes during the analysis task. However, a person can attend to areas outside her fixation or be blind to visual information near her fixation (“inattention blindness”), as Simons and Chabris demonstrated [21]. One must be cautious to interpret RFV results with the understanding that covert attention cannot be detected in this process-tracking scheme.

The Fauxvea viewer is fully blurred until the user presses the cursor in the viewport. The time and location of each cursor press are recorded as the start time, end time, and location point of a fixation estimate. This is more precise than determining a fixation based on the speed of a hovering cursor, as in the original RFV. While the cursor is pressed, image details directly under the cursor are revealed within a focus radius, as shown in Fig. 1. The blur approximates how details in one’s peripheral vision appear when the fovea is fixated elsewhere in the field of view. We use a radius instead of the original RFV’s rectangular window with steps of blur. The idea of a focus spotlight is similar to other approaches in foveated imaging and Focus+Context techniques in information visualization, like semantic depth of field [22]. We note that the information loss that occurs in peripheral vision and how it affects visual search are not fully understood. Roseholtz and others have argued that blur is too simple of a model and that other summary statistics may be computed over a pooling region in one’s vision [23]. Incorporating different models of lossy visual information into an RFV-like interface is an open challenge beyond the scope of this work.

For the experiments described later, the focus radius is equal to 133 pixels (1/6 the width of each stimulus). This parameter is chosen to approximate to the extent of the fovea, which is between 1-2 degrees of the field of view [24]. The Fauxvea focus region does not move if the user drags,

forcing the user to release before pressing in a new location. This lets Fauxvea record fixation start and end times. The interface does not support zoom or pan operations, though scrolling in the browser window will not impact the interface. Within the focus radius, each pixel has a color that is a blend of corresponding colors in the original image and blurred image. The blend ratio for each pixel is proportional to its distance from the cursor press location; pixels outside the focus radius are fully blurred.

For the purpose of tracking fixations during visual analysis tasks, the visualization should be blurred enough that the task is impossible to answer correctly without fixating using the cursor. We expect users to fixate in the image using either: (1) previous knowledge of the image type (e.g., where guide marks might exist in a chart), (2) interesting low-resolution details in the blurred image or in the blended focus radius of a previous fixation location. In the Fauxvea prototype, images are blurred as a preprocessing step. For the experiments described later in this paper, all stimuli are blurred with a Gaussian filter that we selected following a pilot study.

Scalable. Fauxvea is designed to support scalable, online experiments related to visualization analysis. In addition to the image browser, the webpage includes task instructions, controls to navigate between tasks, and an input field for task answers. Cursor interactions and answers to task questions are stored on the client during the task, then sent as a transaction to our database when the task is completed. Full histories of task executions are collected for each user.

Evaluation methods. We ran three experiments to evaluate the validity of the method from different angles as a viable alternative to eye tracking for visualization. First, we collected fixation data using eye tracker with participants performing analysis tasks on basic infovis charts; we compared these fixations to estimates collected online with our method on the same stimuli and tasks using workers on MTurk. Second, we evaluated how well self-assessment works as an alternative to eye tracking or Fauxvea for predicting gaze. Third, we used our method to reproduce findings about visual exploration on tree layouts from an eye-tracking study by Burch et al. [3] to evaluate the method in a realistic scenario with a more complex analysis task.

4 EXPERIMENT 1

In Experiment 1, we performed in parallel an eye-tracking study and an online study using Fauxvea with workers recruited on MTurk. Both studies asked participants to perform a set of visual analysis tasks for image stimuli. Participants were asked one question per image that required them to inspect the image. In the eye-tracking study, participants viewed the stimuli with a normal image viewer while the eye tracker collected data. In the MTurk study, Turkers used the Fauxvea interface and pressed the cursor on the interface while inspecting each image to focus the viewer.

We hypothesize that fixation data collected from both studies will be comparable both qualitatively (**H1a**, **H1b**) and quantitatively (**H2**).

H1a For each stimulus, the two distributions of fixation locations from eye tracking and Fauxvea studies are qualitatively similar.

H1b For each stimulus, the two distributions show patterns that are related to the corresponding analysis task.

H2 Quantitatively, the similarity between the two distributions for each stimulus is significantly higher than the similarity between the eye-tracking distribution and random fixations drawn from a null distribution.

We evaluate **H1a** and **H1b** in Section 4.6 by generating and interpreting heatmaps of fixation locations using data from each study. We evaluate **H2** by applying a distance function (described later in Section 4.4) that compares two fixation distributions.

4.1 Stimuli and Tasks

Three of the most common types of information visualizations were chosen for this experiment: bar charts, scatter plots, and node-link diagrams. A fourth stimulus type, photographs, were also selected from a dataset by Judd et al. [25] and serve in contrast to structured charts in our experiment. We used five images of each type in this experiment, resulting in 20 unique stimuli. All images were scaled to a width of 800 pixels, and the heights ranged from 600-623 pixels.

Each of the visualizations was created programmatically using D3 and Vega. Each bar chart and scatter plot shows 20 samples of a quadratic polynomial with noise added to each value. No axis titles are rendered in the charts. Each node-link diagram showed a graph of 20 nodes with average degree of 3. Networks of this size have been studied in previous eye-tracking experiments [26]. Blurred versions of all stimuli were created using ImageMagick. Additionally, we chose a visual analysis task for each type.

- *Bar charts:* “Estimate the value (height) at year 2008.” The domain in each chart represents years from 1993 to 2013, and the year in the task description changed between images.
- *Scatter plots:* “Estimate the (x, y) position of the biggest outlier in this data trend. For example, $(3.5, 14.8)$.”
- *Node-link diagrams:* “What is the fewest number of edges to travel between the red marks A and B?” Each image shows a different graph layout and has two randomly selected nodes colored red and labelled A and B.
- *Photos:* “Estimate the average age (years) of all people in the photo.” Each photo contains one or more people.

Tasks at this level of complexity have been used in eye-tracking studies involving visualization analysis (e.g., [26]).

4.2 Eye-Tracking Study (ET)

We recruited 18 participants (14 male, 4 female) for the eye-tracking portion of the experiment. Participants were undergraduate and graduate students, except two who were not students. The eye-tracker used in our study was a contact-free RED 125 Hz from SMI. Stimuli were displayed on a 1,600 x 900 pixel monitor with participants seated approximately 30 inches from the monitor. In order to faithfully replicate the Fauxvea browser setup, the visual interface was made to look like the Fauxvea webpage (see Fig. 1) but with

unblurred stimuli. The unblurred stimuli were shown at the same pixel resolution as used in the browser setup.

After a minimal introduction and eye-tracking calibration, participants were shown all 20 stimuli in succession and were asked to provide verbal responses to the task questions.

4.3 MTurk Study (MT)

We created four different HITs on MTurk and recruited 100 Turkers to complete each. Each HIT corresponded to one of four stimulus-task types: bar charts, scatter plots, node-link diagrams, and photographs. In each HIT, participants looked at five images and performed the corresponding visual analysis tasks described earlier. All participants were located within the United States.

Participants were then asked to inspect five visualizations of the same type one at a time before answering the associated question and moving on. The instructions for each HIT briefly described the image type and task. Participants were instructed to press and hold the cursor over the blurred image to reveal details. Based on results from a pilot study with 41 Turkers, we determined that training materials beyond the instructions were not necessary for these tasks. We were cautious not to suggest analysis strategies for completing these tasks. Participants could advance to the next image in the sequence after any amount of time by providing an answer to the question and clicking a button on the webpage. Participants were paid \$0.15 for completing the HIT, plus \$0.10 bonus if all answers were accurate according to an expert reviewer. The bonus is an incentive for Turkers to use the cursor interface thoroughly and helps control data quality.

4.4 Comparing Eye Tracking to Fauxvea Estimates

Distance scores were computed between the eye-tracking and crowdsourced gaze data. Low distance scores indicate high similarity between the gaze locations in both data sets. For each image, we considered two sets of points: the union of all cursor press locations by Turkers using Fauxvea, and the union of all fixation locations by the eye-tracked participants. For each image, the analysis followed these steps:

- 1) For both sets, estimate probability density functions for the pixel locations using kernel density estimation (KDE) with a Gaussian kernel. This gives spatially smooth representations of the fixation data.
- 2) Discretize each smooth representation of the gazes on the original pixel grid. This creates two histograms, H_{ET} and H_{MT} .
- 3) Compute the distance between H_{ET} and H_{MT} .

This approach is similar to a previous study comparing gaze maps [15].

In this experiment, we tested several distance functions to compare H_{ET} and H_{MT} . In the remainder of this paper, we report results from the χ^2 goodness-of-fit test and a symmetric version of Kullback-Leibler (KL) divergence. Both are off-the-shelf techniques that have been previously used to quantify differences between gaze sets [15] and between saliency maps and human fixation maps [27], [28]. Other metrics including Earth Mover's Distance (EMD) and Area Under the Curve (AUC) variations have also been used and combined to evaluate saliency models [29]

and are applicable to our study; we limited the metrics to χ^2 and KL for simplicity.

4.5 Comparing Eye Tracking to Random Gazes

For hypothesis **H1**, we try to reject the null hypothesis that Fauxvea gaze estimates are spatially uncorrelated with actual eye-tracking fixations. In this section, we describe "null" distributions, or baselines, for gazes that we expect to be less similar to ET than MT is. In Section 7.2, we discuss how building models of gaze during visualization tasks could help us test more realistic null hypotheses.

We expect the distance between a real gaze map and a random gaze map to be significantly larger than the distance between corresponding ET and MT gazes for a visualization task. We considered several baseline gaze distributions that we believe are unlikely to be correlated spatially with eye-tracking fixations during visualization tasks:

- *Grid*, where fixations are evenly distributed in the stimulus.
- *Uniform*, where fixations are equally likely in any part of the image. Rudoy et al. compared χ^2 distances between eye tracking and crowdsourced fixations with their method to distances between ET and uniform random fixations [15]. This is a baseline model for saliency ("Chance") in the MIT Saliency Benchmark [29].
- *Centered Gaussian*, where fixations are normally distributed in the center of the image. Judd et al. showed that the center of a photograph is a good a priori estimate of gaze location [25]. This is a baseline model for saliency ("Center") in the MIT Saliency Benchmark [29].
- *Uniform + Centered Gaussian*, which is a combination of the uniform and centered Gaussian distributions.

In addition to the above baselines, we compute outputs from a visual saliency model (*Saliency*) that is task-agnostic and compare these heatmaps to our ET gazes. The motivation for this step is to see how an off-the-shelf saliency detector compares to Fauxvea for predicting gaze during predefined analysis tasks. There are many saliency detectors available that take images as inputs and output smoothed saliency heatmaps; in this experiment, we demonstrate using Judd et al.'s model [25] that is trained using a benchmark set of eye-tracking data, and is therefore transparent for others to use. We report distances from ET to each of these gaze distributions in Section 4.6.

We computed distances from ET to each baseline.

- For Grid: a set of points were generated forming a $n \times n$ grid on the stimulus, where n is the square root of the number of fixations in the gaze data.
- For Uniform, Centered, or Uniform+Centered: a set of points was sampled from the distribution, using as many fixations as in the gaze data.

The distance between these baseline point data and the gaze data was computed using the algorithm described in Section 4.4. For the baselines involving a sampling procedure, distances were computed for 100 sampling iterations for each stimulus, then averaged.

For the Saliency baseline, we computed the average distance between ET and the model-generated saliency map

TABLE 1
Summary of Turkers From Experiment 1

Task	Participants			Fixations		Familiarity with the Image Type
	Total	Age	Mouse / Trackpad / Other	Total	Per task	Never / Sometimes / Often
Bars	98	$\mu = 29.1, \sigma = 7.7$	77.6% / 18.3% / 4.1%	4,216	$\mu = 9.9, \sigma = 7.3$	52.0% / 40.8% / 7.2%
Scatter	98	$\mu = 27.5, \sigma = 6.9$	68.4% / 28.6% / 3.0%	7,484	$\mu = 24.9, \sigma = 20.5$	57.1% / 34.7% / 8.2%
Node-link	96	$\mu = 28.2, \sigma = 7.3$	74.0% / 24.0% / 2.0%	4,520	$\mu = 10.2, \sigma = 10.0$	63.5% / 27.1% / 9.4%
Photos	100	$\mu = 29.3, \sigma = 9.7$	73.0% / 24.0% / 3.0%	6,314	$\mu = 14.3, \sigma = 11.7$	6.0% / 43.0% / 51.0%

“Fixations” refers to the number of cursor presses that are used to focus on the stimulus. “Total” is the number of fixations for all participants on all five stimuli in each category. “Per task” shows the mean and standard deviation for the number of fixations per user, per stimulus.

for all stimuli. To compute each single distance score for a stimulus, we used the algorithm in Section 4.4 to get a normalized histogram of the ET gazes, then we normalized the model-generated saliency map as a histogram before applying the distance function.

4.6 Results

In this section, we report findings from our comparison of eye tracking and Fauxvea estimates for basic infovis charts and tasks.

Summary statistics for our data collection experiments are shown in Table 1. Turkers performed 392 HITs from four different stimuli-task types. Eight Turkers submitted HITs without performing any Fauxvea cursor presses; therefore, their data are not included in our analysis.

In general, fixations were distributed at similar locations between the eye-tracking (ET) and Fauxvea (MT) studies for all infovis stimuli. Heatmaps of fixations collected in Experiment 1 are shown in Fig. 2, along with saliency maps generated from the Judd model. Each row shows a sample

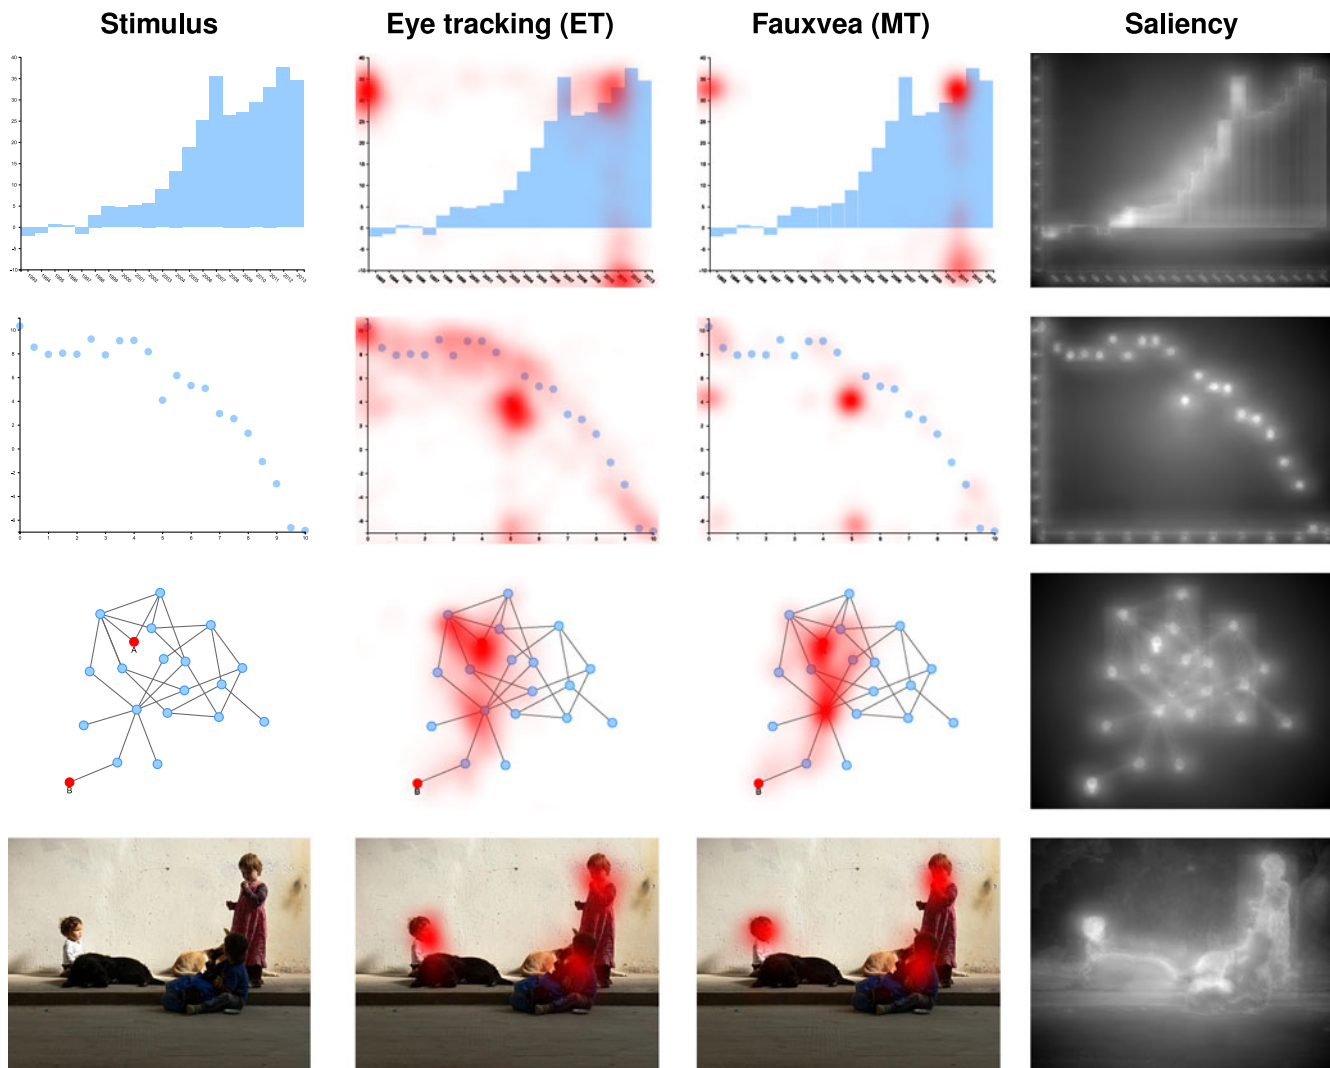


Fig. 2. Comparison of eye-tracking gazes, Fauxvea gaze estimates from Turkers, and visual saliency maps. Red overlays show maps of fixation locations by 18 eye-tracking participants (middle-left) and between 96-100 Turkers per stimulus type (middle-right). Saliency maps (right) were computed from a visual saliency model [25], but models like these do not account for predefined analysis tasks.




Fig. 3. Fauxvea estimates are significantly more similar to eye tracking (Experiment 1) than each other baseline is ($p < .001$ for each). Smaller scores indicate more similarity. Error bars show ± 1 standard error.

visualization from our experiment, along with overlays of gaze data collected in ET and MT studies. Red overlays show normalized maps of fixation locations both by eye-tracking participants and Turkers. All stimuli and heatmaps from Experiment 1 are available online at: <http://bit.ly/fauxvea-sup>.

Similarities in these heatmaps between conditions support **H1a**. In most cases, white spaces in a visualization are not fixated on in either eye tracking or Fauxvea, and the most relevant marks for the analysis task are fixated on most heavily. Evidence supporting **H1b** is clearest in the heatmaps of bar chart and scatter plot, where specific axis labels corresponding to correct task responses (i.e., column heights or (x, y) coordinate values) are fixated on heavily while the others are largely ignored. Heatmaps can also illustrate what visual-analysis strategies are used to complete tasks with less structure that do not use guide marks. For example, it is clear that people primarily fixate on faces in both eye-tracking and Fauxvea results to answer the photograph task “Estimate the average age of all people in the photo” and not other context clues (see Fig. 2 for an example).

We found quantitative evidence that fixation estimates made with Fauxvea are more similar to eye tracking than the baseline estimates we tested. While it is not surprising that Fauxvea performs better than random and task-agnostic gaze estimates, the results confirm a basic requirement for the method and also demonstrate how to quantitatively compare two sets of fixation locations. As we discuss in Section 7, opportunities exist to use this evaluation approach to compare new models of gaze against each another.




Fig. 4. Pair-wise χ^2 distances between eye tracking (ET) and Fauxvea gaze estimates on Mechanical Turk (MT) for all 20 stimuli. As a sanity check, we compared each ET dataset to each MT dataset from Experiment 1 and visualized the distance scores in a matrix. We expect that when using a reasonable distance metric, the smallest distances (darkest cells) will appear on the diagonal, where ET and MT are compared for the same stimulus.

The distance between ET and MT (0.39 using the symmetric KL function, 0.23 using the chi-squared function) is significantly less than the distance between ET and each of the baselines ($p < .001$ for all paired, two-tail t-tests), which supports **H2**. Fig. 3 shows the average distance for both symmetric KL divergence and χ^2 distance between all ET and MT data, and the difference between ET data and each of the baselines we considered.

Table 2 shows the average distance scores for both metrics between the ET data for the four image-task types and each of four baseline null gaze distributions. These values suggest which null distribution is most fair to sample for random comparisons against Fauxvea gaze estimates, for each of the four stimuli-task types we evaluated. In general, fixations on charts with guide marks near the edge of the image (e.g., bar charts and scatter plots) are most similar to samples from a grid-based or uniform distribution rather than a distribution with higher likelihood near the image center. Marks like axes are critical for decoding information,

TABLE 2
Distance from Eye-Tracking Data on Different Visualization Types to Random Gazes from Four Baseline Distributions

Image type	Symmetric KL Divergence				χ^2 Distance			
	Uniform+Gaussian	Gaussian	Uniform	Grid	Uniform+Gaussian	Gaussian	Uniform	Grid
Eye tracking								
Scatter	1.09	8.10	0.75	0.73	0.83	1.35	0.65	0.64
Bars	1.60	9.84	1.08	1.07	1.11	1.63	0.84	0.84
Node-link	1.20	2.25	1.52	1.51	0.90	0.96	1.13	1.13
Photos	0.98	3.90	0.96	1.07	0.75	0.98	0.87	0.86

For each distance function, bold values show the distribution that most closely fits the image type (smallest distance score). These values suggest which null distribution is the fairest to sample for baseline comparisons against Fauxvea gaze estimates, for each of the four stimuli-task types we evaluated.

but where a person attends is usually task-dependent. By contrast, photographers tend to frame the most interesting parts of the image near the center.

5 EXPERIMENT 2

We performed a small-scale follow-up experiment to test whether people with experience and interest in eye tracking are able to reliably predict fixation locations for the tasks and stimuli in Experiment 1. We wanted to know whether it is viable for a person to self-assess where gazes happen during a visualization task instead of running a crowdsourced study with Fauxvea or performing an eye-tracking study.

5.1 Methods

We recruited six participants (four male, two female) who were researchers in human-computer interaction or computer vision at a major research university in the United States. Each was right-handed, had normal or corrected-to-normal vision, and identified himself as having experience or interest in learning where people look in images. The ages of participants ranged from 25-35 years ($M = 28.7$, $SD = 4.13$). All participants passed an Ishihara test for normal color vision.

In the first task, each participant was seated about 18 inches from a 24" 1,920 x 1,200 pixel monitor and viewed each of the 20 task-stimulus pairs from Experiment 1. For each task, participants were asked to select five or more locations in the stimulus they believe people performing the task will fixate on. Participants indicated their fixation locations by clicking a cursor at locations inside the image. The tasks were presented in the same order that participants in the eye-tracking condition (ET) in Experiment 1 performed them. We recorded each predicted location.

In the second task, participants were seated at the same display and shown eye tracking (ET) and Fauxvea (MT) gaze heatmaps for each task-stimulus pair, side by side. The ET and MT heatmaps were generated from data collected in Experiment 1. For each of the 20 pairs, each participant was asked to click on the heatmap she believed was generated from real eye-tracking data; the position of the ET heatmap – left or right – was randomly assigned between stimuli. Participants could also select a "Too close to call" button if they could not identify the ET heatmap. We scored how accurate each person was in selecting the ET heatmaps in the set of stimuli.

5.2 Results

The results from this study are primarily qualitative and anecdotal for two reasons: (1) recruiting a large sample size of people with experience in vision or eye tracking is difficult; (2) when asked to select fixation locations, most participants selected only the minimum number of locations we requested. Therefore, the data are sparse.

We examined the results from the first task by visualizing all fixations predicted by the expert participants and looking for patterns in how participants chose points across stimuli. We created a filterable visualization of these fixation predictions that is available online at: <http://bit.ly/fauxvea-sup>. Most participants identified similar key areas in each stimulus, but predictions varied in how people attend to areas of the stimuli that are visual salient but

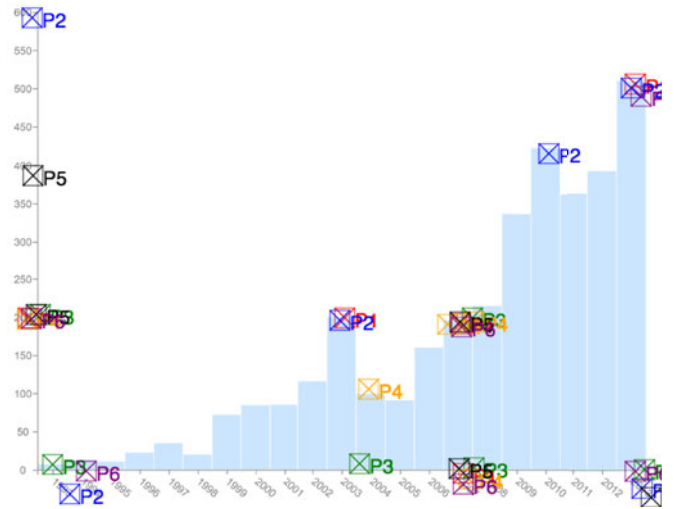


Fig. 5. Predicted fixation locations for the task "Estimate the value (height) at year 2007" by participants (P1-P6, marked with unique colors) in Experiment 2.

irrelevant to the particular tasks. Fig. 5 shows an example where five of six participants predicted that others would fixate on the top of the column, x -axis guide marks, and y -axis guide marks corresponding to the task; however, there was little consensus on which other bars or guide marks people would attend. One participant (P2), a graduate student who studies eye tracking, did not predict any fixations on the task-specific areas of the bar chart. In a follow-up interview, P2 indicated that she focused on marking only visually salient regions.

Our main insight is that even with expertise in thinking about where others will gaze in an image, as an individual it is difficult to predict how a population of people will gaze during a visual analysis task. Several participants commented that they made predictions by first solving the task on their own, then reporting where they gazed during that trial; however, this strategy limits the evaluator to only one perspective and is not viable for estimating gazes from a population that might analyze a visualization in different ways.

In the second task, the experts correctly identified the real eye-tracking heatmap with 68.3 percent accuracy on average. They incorrectly identified the Fauxvea heatmaps as eye tracking 23.3 percent of the time, and 8.3 percent of the time they selected "Too close to call". In a follow-up interview, most participants indicated that when heatmaps were noisier it was an indicator of real eye tracking. P5, who had run eye-tracking studies prior to our experiment, commented that adding random noise to the Fauxvea heatmaps could make them look closer to the eye-tracking heatmaps.

6 EXPERIMENT 3

We ran a third experiment to see if the Fauxvea method is able to reproduce findings about visual exploration behaviors from an existing eye-tracking study using tree visualizations. This follow-up was aimed at evaluating the external validity of Fauxvea beyond the basic visualization interpretation tasks in Experiment 1. We reproduced the task and three stimuli from an eye-tracking study of traditional, orthogonal, and radial tree layouts [3]. Burch et al. note that

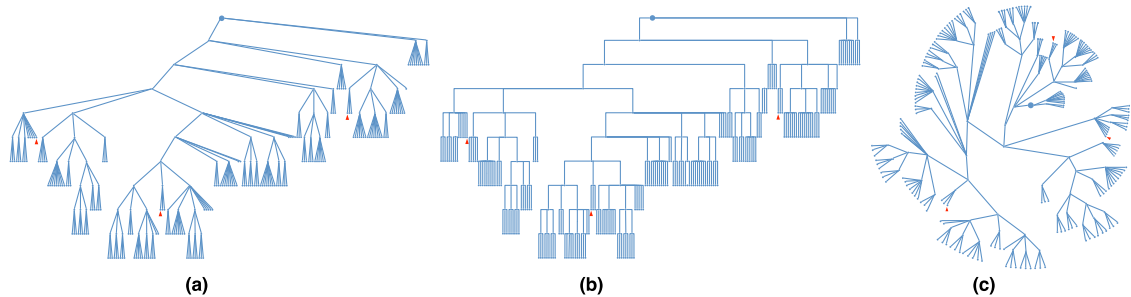


Fig. 6. Three stimuli for Experiment 3: (a) traditional, (b) orthogonal, and (c) radial tree layouts. The root node is indicated by a larger circle mark, and the target nodes for the common-ancestor task are indicated by red arrows.

these layouts are “frequently used in many application domains, they are easy to implement, and they follow aesthetic criteria for tree drawing”.

In general, we hypothesize that fixation locations and transition frequencies reported in the original study will be reproduced by running a similar experiment with Fauxvea (**H3**). In addition, we hypothesize that the three findings in the section titled “Analysis of Exploration Behavior” in [3] will be corroborated with data collected from a reproduction of the experiment using Fauxvea instead of an eye tracker (**H4a**, **H4b**, **H4c**).

- H3 Transition frequencies between predefined areas of interest in the original study will be similar to transition frequencies reproduced using Fauxvea.
- H4a Participants will jump more frequently between leaf nodes that are near each other in the traditional layout compared to the orthogonal and radial layouts.
- H4b The pixel distance between the marked leaf nodes will affect the transition frequency.
- H4c Participants viewing the radial layout will transition back from the root node to AOI 2 more frequently compared to the traditional and orthogonal layouts.

We evaluate **H3** in Section 6.3 by analyzing the most frequent destination AOI from each source AOI. We evaluate **H4a**, **H4b**, and **H4c** in Section 6.3 by analyzing specific patterns in the corresponding transition tables.

Testing these hypotheses will help us evaluate whether Fauxvea can reproduce findings about visual exploration of visualizations without using an eye tracker. We note that we do not compute distance scores, as we did in Experiment 1, because the fixation data from the eye-tracking study were not available. Furthermore, rather than measuring how closely the Fauxvea estimates match the eye-tracking heatmaps, the main goal of this experiment is to corroborate or reject the findings from [3] using a similar analysis.

6.1 Stimuli and Tasks

Participants were shown tree diagrams with marks that indicated the root node and three target nodes. Stimuli were composed of three tree layouts: traditional, orthogonal, and radial. The layouts differ in how nodes and edges are aligned. These layouts are shown in Fig. 6. The participants were asked to find the *least common ancestor* (LCA) of the target nodes in each tree. The instructions included a definition of the LCA written in plain English.

We asked participants to report the coordinates of the LCA for each tree they analyzed, so we added an interaction

to the Fauxvea interface that lets users find the coordinates over the cursor location. While interacting with the interface, the user can type the ‘Return/Enter’ key to place a mark under the cursor and its (x, y) coordinates are displayed on the screen. In this way, participants can quickly find the coordinates of locations in the image without interfering with cursor presses.

6.2 MTurk Study

We created three different HITs on MTurk corresponding to the three tree layouts and recruited 85 Turkers to complete each. We restricted each participant to one HIT only because the same underlying graph data is visualized in each HIT. All participants were located within the United States.

In each HIT, participants performed three training tasks using Fauxvea and were shown example trees with the LCA labelled to help them understand the task. For the fourth task, participants completed the task for the test stimulus that was replicated from the Burch et al. study. The test stimulus did not have the LCA labelled. During each HIT, participants could advance to the next image in the sequence after any amount of time by providing an answer to the question and clicking a button on the webpage. They were not allowed to revisit past images after moving on. We expected that each HIT would take 4-5 minutes to complete and paid each Turker who completed a HIT \$0.45. A \$0.15 bonus was offered to each Turker who identified the LCA correctly according to an expert reviewer.

After the visual analysis tasks, we collected demographic information about participants’ age, sex, and cursor device (mouse, trackpad, or other), as well as how often they look at images like these (“never”, “sometimes”, “often”) and general feedback about the strategy each Turker used in the LCA task.

6.3 Results

Each task response was considered accurate if the reported coordinates for the LCA were within 10 pixels of the known answer. Turkers who were accurate were given a \$0.15 bonus as an incentive to be thorough. The average task accuracy for the HITs differed: Turkers were most accurate with the orthogonal layout (50.6% = 43/85), slightly less accurate with the traditional layout (41.2% = 35/85), and least accurate with the radial layout (20% = 17/85).

Overall, we found that the distributions of fixations from Experiment 3 on the three stimuli were similar to those published in Burch et al. [3]. In Fig. 7, we show heatmaps of fixations from Experiment 3 alongside those from the original

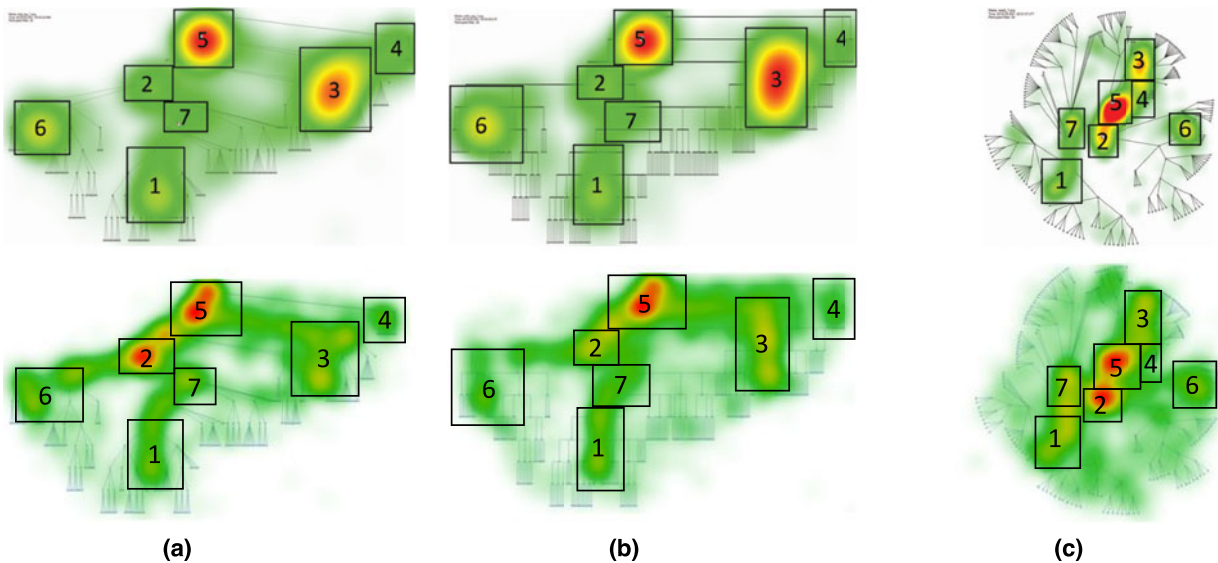


Fig. 7. Comparison of results from Experiment 3 with the results reported by Burch et al. Data in columns (a), (b), and (c) correspond to traditional, orthogonal, and radial layout conditions. The top row shows eye-tracking heatmaps from [3], © IEEE 2011. The bottom row shows the corresponding heatmaps we computed in Experiment 3.

study. The top row shows fixation heatmaps and the AOIs specified in the original study. The second row shows heatmaps we generated from the data collected in Experiment 3. We implemented a heatmap renderer using a color map that approximates the visualization technique in the earlier work; therefore, some visual differences between these charts may be due to implementation differences.

We observed several similarities between the heatmaps from the original study and Experiment 3. In all heatmaps, AOI 5, which contains the root node of the tree, is fixated on heavily. This makes sense because locating the root node is critical for the task of finding the LCA of the target nodes. Some subtrees and leaf nodes were essentially ignored in both the original study and in Experiment 3. These include areas that are dense with nodes and edges but are not parts of the visualization that one must attend to find the LCA (e.g., between AOI 1 and AOI 3 in the traditional and orthogonal layouts). This suggests Turkers are able to focus on the task at hand and do not spend effort fixating on areas of the visualization that are irrelevant to the task, even if those areas are comparably visually salient in the blurred viewer. The fact that the heatmaps are similar between studies, and that they show evidence participants fixate on task-specific areas, supports both **H1a** and **H1b**. We did not evaluate **H2** for Experiment 3 because the raw eye-tracking data needed to compute quantitative distance scores were not available.

We also observed some differences between the fixation maps from these experiments. In general, the eye-tracking fixations appear more focused and less spread out than the Fauxvea fixations. This contrasts earlier findings in Experiment 1. An exception to this is the set of Fauxvea fixations that occur along edges in the traditional and orthogonal layouts (e.g., between AOI 2 and AOI 6, and between AOI 1 and AOI 7). Another noticeable area where heatmaps differ in the traditional layout is AOI 2, which is fixated on relatively more than AOI 3 in our experiment and less than AOI 3 in the original experiment. In this layout, Turkers were much more likely to transition from AOI 7 to AOI 2 than to any other AOI; in the original eye-tracking study, AOI 2

was also the most frequent transition from AOI 7, but AOI 1 and AOI 3 are other common destinations (each with > 10 percent frequency).

All observed transition frequencies, which can be thought of as probabilities, between specific AOIs are shown in Fig. 8. The top row shows probabilities from the original eye-tracking study, and the second row shows probabilities from the data in our experiment. The similarity in transition tables between experiments suggest that participants explored the AOIs using similar patterns. For 14 out of the 21 source AOIs in the three layouts (67 percent), the most frequent destination AOI (highlighted in green) was the same in both the original results and our results. This is much better than the 3 or 4 matches ($16.7\% = 1/6$) we expect if the most frequent destination from each AOI were randomly chosen from the remaining six. Therefore, we find support for **H3**. The cells highlighted in yellow show where the most likely destination AOI is different between the experiments. In all of these cases, the most likely transition in our experiment was the second most likely in the original experiment data.

Examining the transition probabilities, we did not find support for **H4a**. Burch et al. found that the probability from AOI 1 to AOI 6 (and vice versa) was 19 percent (17 percent) for the traditional layout, in contrast to 7 percent (6 percent) in our study, which is comparable to the orthogonal layout results from both studies. Transition probabilities for these AOIs in the radial layout are comparable between studies.

We found partial support for **H4b**. We re-examined the transitions that supported this hypothesis in the original study. The Fauxvea transition probability from AOI 1 to AOI 3 (and vice versa) is less than the probability from AOI 1 to AOI 6 (and vice versa) in the traditional (3 percent (2 percent) compared to 7 percent (6 percent)) and orthogonal (4 percent (3 percent) compared to 7 percent (3 percent)) layouts. The distance between AOIs 1 and 3 is greater than the distance between AOIs 1 and 6. However, for the radial layout, we found that the probability from AOI 1 to AOI 3 (and vice versa) was not necessarily less than the probability from AOI 1 to AOI 6 (and vice versa), as in the eye-tracking

traditional		source							
		outside	AOI 1	AOI 2	AOI 3	AOI 4	AOI 5	AOI 6	AOI 7
destination	outside	-	59%	35%	59%	22%	35%	36%	47%
	AOI 1	17%	-	11%	3%	4%	8%	17%	13%
	AOI 2	6%	3%	-	0%	0%	4%	31%	18%
	AOI 3	26%	5%	11%	-	40%	24%	0%	11%
	AOI 4	6%	1%	0%	10%	-	17%	3%	0%
	AOI 5	25%	3%	15%	26%	33%	-	12%	5%
	AOI 6	12%	19%	9%	0%	0%	11%	-	5%
	AOI 7	7%	10%	20%	1%	0%	1%	2%	-

orthogonal		source							
		outside	AOI 1	AOI 2	AOI 3	AOI 4	AOI 5	AOI 6	AOI 7
destination	outside	-	55%	38%	41%	45%	42%	66%	43%
	AOI 1	17%	-	10%	1%	3%	7%	2%	16%
	AOI 2	12%	5%	-	2%	0%	5%	10%	25%
	AOI 3	28%	3%	16%	-	36%	32%	5%	6%
	AOI 4	4%	0%	2%	11%	-	7%	0%	0%
	AOI 5	16%	3%	21%	43%	15%	-	15%	8%
	AOI 6	17%	7%	5%	1%	0%	4%	-	2%
	AOI 7	5%	27%	10%	1%	0%	3%	2%	-

radial		source							
		outside	AOI 1	AOI 2	AOI 3	AOI 4	AOI 5	AOI 6	AOI 7
destination	outside	-	56%	17%	12%	8%	16%	33%	40%
	AOI 1	21%	-	10%	1%	1%	4%	4%	15%
	AOI 2	16%	14%	-	4%	3%	22%	36%	28%
	AOI 3	15%	1%	5%	-	23%	31%	4%	0%
	AOI 4	1%	0%	2%	43%	-	16%	7%	2%
	AOI 5	17%	8%	42%	35%	62%	-	15%	15%
	AOI 6	14%	1%	13%	4%	4%	9%	-	0%
	AOI 7	16%	20%	10%	0%	0%	2%	0%	-

traditional		source							
		outside	AOI 1	AOI 2	AOI 3	AOI 4	AOI 5	AOI 6	AOI 7
destination	outside	-	68%	40%	59%	27%	50%	49%	39%
	AOI 1	11%	-	5%	2%	3%	2%	6%	4%
	AOI 2	20%	7%	-	2%	1%	11%	23%	41%
	AOI 3	18%	3%	7%	-	29%	18%	6%	5%
	AOI 4	4%	0%	0%	3%	-	9%	0%	0%
	AOI 5	25%	3%	20%	27%	31%	-	11%	4%
	AOI 6	10%	7%	13%	2%	5%	6%	-	4%
	AOI 7	8%	9%	12%	2%	0%	0%	3%	-

orthogonal		source							
		outside	AOI 1	AOI 2	AOI 3	AOI 4	AOI 5	AOI 6	AOI 7
destination	outside	-	49%	57%	65%	56%	58%	68%	47%
	AOI 1	10%	-	2%	3%	0%	3%	3%	26%
	AOI 2	12%	2%	-	0%	0%	11%	12%	18%
	AOI 3	23%	4%	3%	-	23%	14%	4%	2%
	AOI 4	6%	0%	0%	10%	-	3%	0%	1%
	AOI 5	29%	2%	27%	19%	20%	-	8%	3%
	AOI 6	10%	7%	2%	1%	0%	5%	-	1%
	AOI 7	6%	32%	7%	0%	0%	1%	2%	-

radial		source							
		outside	AOI 1	AOI 2	AOI 3	AOI 4	AOI 5	AOI 6	AOI 7
destination	outside	-	61%	42%	28%	11%	26%	64%	42%
	AOI 1	24%	-	8%	2%	0%	2%	3%	18%
	AOI 2	20%	9%	-	5%	2%	29%	17%	26%
	AOI 3	12%	4%	3%	-	23%	23%	1%	2%
	AOI 4	1%	0%	1%	9%	-	7%	5%	0%
	AOI 5	16%	2%	20%	46%	51%	-	7%	8%
	AOI 6	14%	2%	5%	5%	11%	4%	-	0%
	AOI 7	11%	20%	17%	1%	0%	6%	1%	-

Fig. 8. Comparison of results from Experiment 3 with the eye-tracking results reported in [3]. Data in columns marked (a), (b), and (c) correspond to traditional, orthogonal, and radial layout conditions. Rows 1 and 2 respectively show transition probabilities between AOIs in both eye tracking and from Experiment 3. Transition probabilities to or from areas outside any AOI are grayed out. In general, transitions in both experiments were similar and fell into two cases: (1) green cells indicate where the most likely destination AOI from a source in Experiment 3 is also the *most likely* in the eye-tracking results; (2) yellow cells indicate where the most likely destination AOI from a source in Experiment 3 was the *second most likely* in the eye-tracking results.

study: 4 percent (2 percent) compared to 2 percent (3 percent). In fact, the distances from AOI 1 to AOI 3 and AOI 6 are not as different in the radial layout and in the traditional and orthogonal ones (see Fig. 6). We discuss possible explanations for these differences in Section 7.1.

Finally, we found strong support for **H4c**. In our experiment, the probability from AOI 5 to AOI 2 is 29 percent in the radial layout but only 11 percent in both traditional and orthogonal layouts. This is comparable to Burch et al.'s probabilities for this transition: 22 percent (radial), 4 percent (traditional), and 5 percent (orthogonal).

7 DISCUSSION

7.1 Visual Exploration Behaviors

We noticed a few differences in how eye-tracking and Fauxvea fixations are distributed spatially for visualization tasks. In Experiment 1, we found that eye-tracking gazes generally occur over wider regions of the image and appear more spread out than Fauxvea gaze estimates (see Fig. 2). There are several possible explanations for this behavior: (1) People do not look at a singular point of interest for long; instead, their gaze hovers around that point. (2) Holding a cursor at one pixel location over time requires less effort than gazing at one location for the same amount of time. (3) The time and effort needed to move and press the cursor is greater than a saccade of the eye. (4) For crowd workers, there is an opportunity cost to being slow or getting distracted by irrelevant image details. (5) Eye trackers can have errors due to both calibration and moments when the eye tracker cannot find the eye. Therefore, recorded coordinates might be inexact.

Our findings in Experiment 3, which involves a more complex task, show a different pattern: in some cases, Fauxvea gaze estimates are more spread out than eye-tracking gazes (see the radial layout in Fig. 7). It is possible that Turkers using the Fauxvea interface had less experience with this task compared to the participants in the eye-tracking study and therefore spend more time exploring the diagram. Turkers

also used the interface to fixate on edges in the tree diagrams more than eye-tracking participants, which supports the idea that they focus on tracing paths to complete the task. Eye-tracking participants with the normal viewer, on the other hand, can make saccades between nodes and may rely on peripheral vision to view edges. In both populations, similar hot spots related to the task appear in the heatmaps.

7.2 Quantitative Comparisons

In Experiment 1, we compared Fauxvea fixations against eye-tracking fixations using a quantitative approach similar to an earlier evaluation by Rudoy et al. [15]. We evaluated an additional distance function and tested additional baseline distributions of random fixations, plus a grid baseline and a saliency model. We discuss our findings below.

Distance metrics. We explored two measures of similarity between smoothed representations of gaze locations: a symmetric version of KL divergence and χ^2 distance. Fig. 4 shows one of these metrics (χ^2) computed between the ET and MT data we collected, for each pair of stimuli-tasks. As a sanity check, we were interested in seeing how values on the diagonal – which are distances between corresponding visualization tasks in the two conditions – compare to distances between unrelated visualization tasks. We note that this type of matrix should not necessarily be symmetric across the diagonal, because the columns (Fauxvea) represent a different modality for which fixations were collected compared to the rows (eye tracking). The matrix shows that the diagonal is in fact darker than any single row. This is also visual evidence of hypothesis **H1b**: the fixations from both ET and MT are linked to underlying visualization tasks.

Generating baseline gazes. We evaluated Fauxvea quantitatively by computing how much closer to eye tracking Fauxvea's fixation locations are compared to fixations drawn from null distributions that represent where people might look without regard to the visualization task. Creating realistic computational models of gaze during visualization tasks is an open problem. Task-aware models could replace the

need for gaze-estimation methods with humans in the loop, or provide stronger baselines for evaluating new estimation methods.

We used an off-the-shelf, state-of-the-art model of visual saliency [25] and found that the maps it generates from the visualization stimuli in Experiment 1 are not much closer quantitatively to eye tracking than the other null distributions, like Grid and Centered Gaussian. This is not surprising because people do not necessarily attend to salient regions that are irrelevant to the analysis task they are given. It is possible that people with experience in vision and eye tracking could identify where people will look during tasks, but as we found in Experiment 2, self-assessment of these areas is not consistent even among experts.

7.3 Similarity to Eye Tracking

This paper focused on validating a method that lets others collect gaze-estimate data and gaze-related metrics for visualization use. Poole and Ball summarized three types of eye-movement metrics [30] that we explored in Experiments 1-3. *Fixation-derived*. In Experiments 1 and 3, we found that estimating the spatial distribution of fixations is feasible. The time spent per Fauxvea fixation was measured, but we have not yet validated that fixation times in AOIs are correlated between Fauxvea and eye-tracking modalities. *Saccade-derived*. Saccade speed and location are difficult to estimate with existing process-tracking techniques. Visual information is not decoded during saccades and, therefore, cannot influence more easily tracked human behaviors (e.g., moving a cursor). *Scan-path-derived*. Experiment 3 demonstrated that transition frequencies between AOIs can be reasonably estimated with Fauxvea, but more work is needed to validate that full scan paths using Fauxvea are correlated with eye-tracking scan paths.

One must be careful when interpreting metrics involving individual Fauxvea fixations, since Fauxvea fixations might be more coarse-scale than eye-tracking fixations. We found that the duration of a Fauxvea fixation is on average longer than an eye-tracking fixation. One explanation for this is that when a person examines an area of a visualization in a normal image viewer, she might make several fixations near the same area; in Fauxvea, this might be replaced by a single, longer fixation due to the added cost of refocusing, as described in Section 7.1.

Using any form of eye tracking or gaze estimation to understand a person's cognitive activities with a visualization depends on the "eye-mind" hypothesis – that what a person gazes indicates her foremost cognitive process [20], [30]. The dissociation between where a person fixates and her covert attention could lead researchers to misinterpret where she attends. In particular, visualization analysis sometimes requires keeping several pieces of information in mind while solving a task, so it is possible the hypothesis does not hold for some tasks, as Kim et al. found [31]. We believe follow-up interviews or questionnaires could help verify cognitive processes that are apparent in gaze traces.

7.4 Limitations and Open Challenges

We evaluated Fauxvea using straightforward tasks, each lasting five minutes or less, on static visualizations. Many

visualization use cases fall outside of these constraints and have not been tested with Fauxvea, such as exploratory analysis scenarios that last tens of minutes or longer and use interactive tools. Interactive visualizations are particularly difficult to study with a RFV-based method because the viewer requires cursor interactions that might conflict with underlying interactions. Mixing both focus interactions and interactions in the underlying visualization could disrupt a user's analysis workflow. Another limitation of using a RFV interface during visualization evaluation is that it could discourage users from participating in the evaluation. For instance, Bednarik and Tukiaine [5] reported that some users disliked the RFV interface.

We identified several opportunities to build on the Fauxvea method and related approaches. First, we did not analyze fixation sequences, or *scan paths*, in this work. In a preliminary analysis, we clustered Fauxvea and eye-tracking sequences from Experiment 1 using dynamic time warping (DTW) and found some evidence that scan paths from both conditions cluster together on the same task. However, these results were not conclusive because of the small sample size of eye-tracking scan paths. Second, we believe the data collected from Fauxvea experiments could be used to improve the baselines mentioned in Section 7.2 and create computational models of gaze for visualization. In turn, these models could bootstrap the evaluation of gaze-estimation user interfaces that have humans in the loop, like Fauxvea. In addition, we have not evaluated Fauxvea using remote visualization-domain experts. It might be possible to collect fixation estimates from expert populations that were previously inaccessible for traditional eye-tracking lab studies, enabling studies with larger sample sizes of domain experts than are currently possible.

8 CONCLUSION

We evaluated a crowd-powered method called Fauxvea that estimates gaze locations during analysis tasks with static information visualizations without using an eye tracker. The method was adapted from an earlier one that has not been evaluated in the context of analysis tasks for information visualization. We ran three experiments to evaluate the method, including a reproduction of earlier eye-tracking findings about tree visualizations. In Experiment 1, we found quantitative and qualitative evidence that Fauxvea fixations from many Turkers are similar, and often less noisy, compared to fixations collected in a parallel eye-tracking study with a typical number of participants. In Experiment 2, we found that self-assessment of fixations by individuals can be inconsistent; therefore, using Fauxvea with crowdsourced workers is likely to be a more reliable approximation of eye tracking. In Experiment 3, we found that the way people transition their gaze between AOIs using Fauxvea is similar to eye tracking, but comparing full scan paths remains an open challenge.

Our method is a practical alternative to eye tracking for finding task-specific areas of interest in static visualizations. Creating a robust computational model of gaze for visualization tasks is an open problem, and data collected using our method might be helpful in constructing such a model.

ACKNOWLEDGMENTS

This work was supported in part by US National Science Foundation (NSF) award IIS-10-16623. All opinions, findings, conclusions, or recommendations expressed in this document are those of the authors and do not necessarily reflect the views of the sponsoring agencies.

REFERENCES

- [1] A. F. Blackwell, A. R. Jansen, and K. Marriott, "Restricted focus viewer: A tool for tracking visual attention," in *Proc. 1st Int. Conf. Theory Appl. Diagrams*, 2000, pp. 162–177.
- [2] M. Jiang, S. Huang, J. Duan, and Q. Zhao, "Salicon: Saliency in context," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, Jun. 2015, pp. 1072–1080.
- [3] M. Burch, N. Konevtsova, J. Heinrich, M. Hoeflerlin, and D. Weiskopf, "Evaluation of traditional, orthogonal, and radial tree diagrams by an eye tracking study," *IEEE Trans. Vis. Comput. Graph.*, vol. 17, no. 12, pp. 2440–2448, Dec. 2011.
- [4] E. J. Johnson, J. W. Payne, D. A. Schkade, and J. R. Bettman, "Monitoring information processing and decisions: The MouseLab system," *Office Naval Res.*, Tech. Rep. 89-4, 1989.
- [5] R. Bednarik and M. Tukiaine, "Validating the restricted focus viewer: A study using eye-movement tracking," *Behavior Res. Methods*, vol. 39, no. 2, pp. 274–282, 2007.
- [6] M. N. Jones and D. J. K. Mewhort, "Tracking attention with the focus-window technique: The information filter must be calibrated," *Behavior Res. Methods, Instrum. Comput.*, vol. 36, no. 2, pp. 270–276, 2004.
- [7] A. R. Jansen, A. R. Blackwell, and K. Marriott, "A tool for tracking visual attention: The restricted focus viewer," *Behavior Res. Methods, Instrum. Comput.*, vol. 35, no. 1, pp. 57–69, 2003.
- [8] M. C. Chen, J. R. Anderson, and M. H. Sohn, "What can a mouse cursor tell us more?: Correlation of eye/mouse movements on web browsing," in *Proc. Extended Abstracts Human Factors Comput. Syst.*, 2001, pp. 281–282.
- [9] K. Rodden and X. Fu, "Exploring how mouse movements relate to eye movements on web search results pages," in *Proc. Web Inf. Seeking Interaction Workshop*, 2007, pp. 29–32.
- [10] K. Rodden, X. Fu, A. Aula, and I. Spiro, "Eye-mouse coordination patterns on web search results pages," in *Proc. Extended Abstracts Human Factors Comput. Syst.*, 2008, pp. 2997–3002.
- [11] Q. Guo and E. Agichtein, "Towards predicting web searcher gaze position from mouse movements," in *Proc. Extended Abstracts Human Factors Comput. Syst.*, 2010, pp. 3601–3606.
- [12] J. Huang, R. W. White, and S. Dumais, "No clicks, no problem: Using cursor movements to understand and improve search," in *Proc. SIGCHI Conf. Human Factors Comput. Syst.*, 2011, pp. 1225–1234.
- [13] D. Lagun and E. Agichtein, "Viewer: Enabling large-scale remote user studies of web search examination and interaction," in *Proc. 34th Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2011, pp. 365–374.
- [14] M. Bostock, V. Ogievetsky, and J. Heer, "D3: Data-driven documents," *IEEE Trans. Vis. Comput. Graph.*, vol. 17, no. 12, pp. 2301–2309, Dec. 2011.
- [15] D. Rudoy, D. B. Goldman, E. Shechtman, and L. Zelnik-Manor, "Crowdsourcing gaze data collection," in *Proc. Conf. Collective Intell.*, 2012, pp. 1–8.
- [16] J. Heer and M. Bostock, "Crowdsourcing graphical perception: Using mechanical turk to assess visualization design," in *Proc. SIGCHI Conf. Human Factors Comput. Syst.*, 2010, pp. 203–212.
- [17] N. Kong, J. Heer, and M. Agrawala, "Perceptual guidelines for creating rectangular treemaps," *IEEE Trans. Vis. Comput. Graph.*, vol. 16, no. 6, pp. 990–998, Nov. 2010.
- [18] R. Kosara and C. Ziemkiewicz, "Do mechanical turks dream of square pie charts?" in *Proc. 3rd Workshop: BEyond Time Errors: Novel Eva. Methods Inf. Vis.*, 2010, pp. 63–70.
- [19] C. Ziemkiewicz, A. Ottley, R. J. Crouser, A. R. Yauilla, S. L. Su, W. Ribarsky, and R. Chang, "How visualization layout relates to locus of control and other personality factors," *IEEE Trans. Vis. Comput. Graph.*, vol. 19, no. 7, pp. 1109–1121, Jul. 2013.
- [20] M. A. Just and P. A. Carpenter, "Eye fixations and cognitive processes," *Cognitive Psychol.*, vol. 8, pp. 441–480, 1976.
- [21] D. J. Simons and C. F. Chabris, "Gorillas in our midst: Sustained inattention blindness for dynamic events," *Perception-London*, vol. 28, no. 9, pp. 1059–1074, 1999.
- [22] R. Kosara, S. Miksch, and H. Hauser, "Semantic depth of field," in *Proc. IEEE Symp. Inf. Vis.*, 2001, p. 97.
- [23] R. Rosenholtz, J. Huang, A. Raj, B. J. Balas, and L. Ilie, "A summary-statistic representation in peripheral vision explains visual search," *J. Vis.*, vol. 12, no. 4, pp. 1–17, 2012.
- [24] J. Johnson, et al., *Designing with the Mind in Mind: Simple Guide to Understanding User Interface Design Rules*. San Mateo, CA, USA: Morgan Kaufmann, 2010.
- [25] T. Judd, K. Ehinger, F. Durand, and A. Torralba, "Learning to predict where humans look," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2009, pp. 2106–2113.
- [26] M. Pohl, M. Schmitt, and S. Diehl, "Comparing the readability of graph layouts using eyetracking and task-oriented analysis," in *Proc. Eurographics Conf. Comput. Aesthetics*, 2009, pp. 49–56.
- [27] T. Judd, F. Durand, and A. Torralba, "A benchmark of computational models of saliency to predict human fixations," Massachusetts Institute of Technology, Cambridge, MA USA, Tech. Rep. MIT-CSAIL-TR-2012-001, 2012.
- [28] L. Zhang, M. H. Tong, T. K. Marks, H. Shan, and G. W. Cottrell, "Sun: A Bayesian framework for saliency using natural statistics," *J. Vis.*, vol. 8, no. 7, p. 32, 2008.
- [29] Z. Bylinskii, T. Judd, A. Borji, L. Itti, F. Durand, A. Oliva, and A. Torralba. (2012). Mit saliency benchmark [Online]. Available: <http://saliency.mit.edu>
- [30] A. Poole and L. J. Ball, "Eye tracking in human-computer interaction and usability research: Current status and future prospects," in *Encyclopedia of Human-Computer Interaction*, C. Ghaoui, Ed. Idea Group, Hershey, PA USA, 2005.
- [31] S.-H. Kim, Z. Dong, H. Xian, B. Upatising, and J. S. Yi, "Does an eye tracker tell the truth about visualizations?: Findings while investigating visualizations for decision making," *IEEE Trans. Vis. Comput. Graph.*, vol. 18, no. 12, pp. 2421–2430, Dec. 2012.



Steven R. Gomez received the ScM degree in 2011 and is currently working toward the PhD degree in the Department of Computer Science at Brown University. He is a member of the Visualization Research Lab at Brown University. His research interests include visualization design, evaluation, and human-computer interaction.



Radu Jianu received the PhD degree in computer science from Brown University. He is currently an assistant professor in the School of Computing and Information Sciences at Florida International University. His research interests are in visualization evaluation, applications of eye tracking to data visualization, and applications of visualization to disaster management and mitigation.



Ryan Cabeen received the BS degree from the California Institute of Technology in 2005, the ScM degree in 2012, and is currently working toward the PhD degree in the Department of Computer Science at Brown University. He was at the UCLA's Laboratory of Neuro Imaging. His research interests are in applications of computational modeling and visualization to scientific and medical imaging.



Hua Guo is currently working toward the PhD degree in the Department of Computer Science at Brown University. She is a member of the Visualization Research Lab at Brown University. Her research interest is to inform and automate visualization design and evaluation using human-centered approaches.



David H. Laidlaw received the PhD degree in computer science from the California Institute of Technology, where he also did post-doctoral work in the Division of Biology. He is a professor in the Computer Science Department at Brown University. His research centers on applications of visualization, modeling, computer graphics, and computer science to other scientific disciplines. He is a senior member of the IEEE and the IEEE Computer Society.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.