



HAL
open science

VETA: Visual eye-tracking analytics for the exploration of gaze patterns and behaviours

Sarah Goodwin, Arnaud Prouzeau, Ryan Whitelock-Jones, Christophe Hurter,
Lee Lawrence, Umair Afzal, Tim Dwyer

► **To cite this version:**

Sarah Goodwin, Arnaud Prouzeau, Ryan Whitelock-Jones, Christophe Hurter, Lee Lawrence, et al..
VETA: Visual eye-tracking analytics for the exploration of gaze patterns and behaviours. *Visual Informatics*, 2022, 6 (2), pp.1-13. 10.1016/j.visinf.2022.02.004 . hal-03673359

HAL Id: hal-03673359

<https://enac.hal.science/hal-03673359>

Submitted on 22 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

VETA: Visual Eye-Tracking Analytics for the Exploration of Gaze Patterns and Behaviours

Sarah Goodwin^a, Arnaud Prouzeau^{a,b,*}, Ryan Whitelock-Jones^a, Christophe Hurter^c, Lee Lawrence^a, Umair Afzal^{a,d}, Tim Dwyer^a

^a*Monash University, Clayton, Australia*

^b*Inria & LaBRI (University of Bordeaux, CNRS, Bordeaux-INP), Bordeaux, France*

^c*ENAC, Université de Toulouse, Toulouse, France*

^d*Swinburne University of Technology, Melbourne, Australia*

Abstract

Eye tracking is growing in popularity for multiple application areas, yet analysing and exploring the large volume of complex data remains difficult for most users. We present a comprehensive eye tracking visual analytics system to enable the exploration and presentation of eye-tracking data across time and space in an efficient manner. The application allows the user to gain an overview of general patterns and perform deep visual analysis of local gaze exploration. The ability to link directly to the video of the underlying scene allows the visualisation insights to be verified on the fly. The system was motivated by the need to analyse eye-tracking data collected from an ‘in the wild’ study with energy network operators and has been further evaluated via interviews with 14 eye-tracking experts in multiple domains. Results suggest that, thanks to state-of-the-art visualisation techniques and by providing context with videos, our system could enable an improved analysis of eye-tracking data through interactive exploration, facilitating comparison between different participants or conditions, thus enhancing the presentation of complex data analysis to non-experts. This research paper provides four contributions: 1) analysis of a motivational use case demonstrating the need for rich visual-analytics workflow tools for eye-tracking data; 2) a highly dynamic system to visually explore and present complex eye-tracking

*Corresponding author

Email address: arnaud.prouzeau@inria.fr (Arnaud Prouzeau)

data; 3) insights from our our applied use case evaluation and interviews with experienced users demonstrating the potential for the system and visual analytics for the wider eye-tracking community; and 4) a set of informed design recommendations for visual eye-tracking analysis.

Keywords: Eye Tracking, Visual Analytics, Spatial-temporal Visualisation

1. Introduction

Eye tracking is widely used in various domains. For instance, in marketing it can be used to understand which advertisement is more attractive [1] or in human factors research, it can be used to understand how pilots react when a critical event occurs [2]. Eye-tracking data consists of a record over time of eye position relative to the scene being viewed. Important features relate to *fixations* (steady gaze on a particular location for some period of time) and *saccades* (quick eye movements between locations). Most eye tracking occurs in controlled studies in laboratories that require external sensors to track gaze. However, with the growing development of wearable eye trackers, it is now possible to perform ‘in the wild’ studies (i.e. in the participant’s natural environment) by having participants wear glasses [3].

These ‘in the wild’ studies tend to be longer, producing more data, but also more noise as they are less controlled than studies in laboratories. In this type of study, we particularly focus on situations in which users have a complex workspace (i.e. more than one screen) but in which they are static most of the time (although rotation can be considered), like a control room or an aircraft cockpit. We, however, do not consider context in which users are moving, like someone playing tennis or a shopper moving in a supermarket (In this latter example, a participant looking at a specific shelf, and thus static in front of it could be considered). Current eye-tracking analysis tools provide limited solutions for the exploratory analysis and visualisation of such ‘wild’ data (e.g. Section 6.3.1). While standard software such as Tobii PRO Lab allows visualisation of gaze trajectories over time such software lacks the kind of holistic

25 visual workflow supporting all aspects of the exploratory sensemaking loop, from foraging through data, searching for evidence or schematising of knowledge to high-level sensemaking.

Motivated by an observational study in which we collected control-room operators' gaze data while they perform their actual tasks, we present Visual
30 Eye-Tracking Analytics (VETA), a system to visually explore and analyse eye-tracking data including, but not limited to, wearable eye tracking devices. This combines a number of distinct visual analytical features through a novel perspective of using eye-tracking data through spatial and temporal linked views, together with an immediate link between the data and the eye-tracking video,
35 to allow for video-driven analytics (see Section 4). It provides clutter reduction to highlight general patterns in the data, by using techniques such as interactive lenses [4] and edge bundling [5]. It extends typical eye-tracking analysis by allowing users to define and manipulate dynamic Areas Of Interest (AOI) and Time Windows of Interest (TWI). We also propose the predefined classification of saccades: 'glance' and 'short distance saccades' which were found to
40 be highly distinguishable and beneficial to filter and explore separately in our analysis. Finally, small multiple comparisons, annotation and the possibility to save and share the state of the analysis, allow for exploration, collaboration and sensemaking.

45 We assess the functionality of VETA, first, by demonstrating how we used the system to find new insights from our motivational use case with energy network operators (see Section 5). Then we discuss the results of interviews with 14 experienced eye-tracking users from different domains providing a richer understanding of the potential for the system to the wider community (see Section 6).
50 Results suggest that the interactive functionality, such as dynamic AOIs, linked spatial and temporal views, and the classification of saccades, would improve their exploration of eye tracking data. Other functionality, such as side-by-side comparison and the save mechanism, are also important to understand what are the common gaze patterns and the difference in behaviour between the participants or conditions in a study. Finally, the use of annotation and simple but
55

informative visualisations of spatial, directional and time-based gaze patterns allow analytical insights from complex eye-tracking data to be understood by non-expert users. From our evaluation of the software we propose a set of design guidelines for future work in visual exploratory analysis of eye tracking data (see
60 Section 7).

This paper enriches the growing body of knowledge for eye-tracking visualisation and analytics, with four specific contributions:

- (1) analysis of a use case of ‘wild’ data collected from a real-world control room environment motivating the need for a tool supporting rich visual-analytics
65 workflow for eye-tracking data;
- (2) the design and implementation of a **comprehensive and highly dynamic system** to visually explore and present complex eye-tracking data;
- (3) **rich and detailed insights** from our applied use case and from interviews with experienced eye-tracking analysts demonstrating the benefits and future
70 potential for the system;

~~(4) a set of **informed design guidelines** for exploratory visual analysis of complex eye-tracking data;~~

2. Motivational Use Case: Network Control Room

The design of VETA was motivated by the need for an efficient and inter-
75 active tool to analyse eye tracking data from a study¹ measuring the activities of operators in the Australian national energy network control room. Because of the complexity and duration of such activities, the data coming from such studies often contains noise and are hard to classify, and thus, difficult to handle using existing commercial analysis software (e.g. Tobii Pro Lab²).

~~80 2.1. Network Control Room Use Case~~

¹The study itself is not the focus of this paper, but it provides a valuable use case of a study ‘in the wild’ with a large amount of eye tracking data.

²<https://www.tobii.com/product-listing/tobii-pro-lab/>

In the control room there are 2-3 operators at any time, each of them has a workstation which is composed of 7 screens on two rows (see Fig. 4). The screens present numerous software applications simultaneously and the operators are able to organise these applications on any of their screens in a way that they believe maximises their understanding and interaction between those applications. As part of the study we recorded the gaze of the operator using a wearable eye tracker (Tobii Pro Glasses 2³). These wearable eye trackers have outward facing cameras that record video of what the operator was looking at, i.e. on their screens or any other elements of the environment. As per Tobii instructions, these were calibrated to 60 centimetres, which is similar to the distance that the operators typically sit away from the screens. They tend to work facing the screens most of the time.

The study was designed to provide quantitative data to allow us to derive insights and recommendations for improvements to the control systems and the operators' workspace. More precisely, we expected to use the eye tracking data to achieve four outcomes.

Outcome 14) To better comprehend the actions of operators during both routine activities and in response to exceptional events. In these contexts eye-tracking data allows us to understand operators' varying attention to different parts of the display over an extended time, such as operators having to transition repetitively between two applications that can be physically far from each other on the (large tiled-display) workspace. Notably, given the size of the screens, this can involve movement of the head in tandem with the eyes. For this work we are interested primarily in attention and therefore do not explore head and eye movement separately, however the techniques we describe could be extended to provide this separation, for example, to investigate visual field bias, i.e. issues of central versus peripheral vision. Whilst visual field biases can occur in these contexts where it is harder to attend to certain sections of their peripheral vision, given the operators have undergone prior training (albeit at different

³<https://www.tobiipro.com/product-listing/tobii-pro-glasses-2/>

110 levels) and they organised the applications along their preferences, therefore we assume that the levels of visual field biases are minimal.

~~Outcome 22)~~ To identify areas of the workspace that are rarely utilised versus locations for common activities e.g. *during this period operators never looked at this application, which is positioned in one of the main screens.*

115 ~~Outcome 33)~~ To compare the operators' reaction with the documented procedure, e.g. *after this alarm, the operator did check this information.*

~~Outcome 44) To understand gaze behaviour of experienced versus novice operators, e.g. novice operators spend more time scanning the network map to find a specific location, or spend~~

One specificity of the control room setup is the operators' workspace (i.e. 7
120 screens that can be considered as a single $3\text{m} \times 1\text{m}$ surface). Its large size makes visual search harder, and some visual behaviours can have a bigger impact on operators' performance compared to regular desktop screens. Glances (As defined by the ISO Norm 15007-1:2014, a glance is a quick look at an area of interest that may be comprised of more than one fixation and saccades to and
125 from it [6]. It duration is between 500 ms and 3 s)(~~i.e. performing a quick visual back and forth to another place of the workspace~~) between two applications, for instance, can take more time, cause divided attention and lead to errors. Short saccades are also important as they can indicate a period in which the operator is focused on a specific application, and can therefore lead to inattention to
130 activity happening in another part of the workspace. Being able to identify, isolate and analyse such specific gaze behaviours is important.

Whilst all operators involved in this control room study use the same applications and do similar tasks, each operator sets out their own workspace and their daily tasks are focused on different regions of the entire network, some
135 being larger than others. Operator tasks also depend heavily on the state of the network. The network is usually divided in sub-regions on which one operator focuses. However, in the case of an unexpected event, operators regroup to collaboratively elaborate an appropriate response. Directly comparing between operator's gaze behaviour is not possible without understanding the wider context.
140 AOIs for one participant are different from another, and can totally depend on

the applications used and their location on the workspace. Whilst applications tend to stay in one place when they are open, they can be resized and new application windows may be opened to do a particular task. Through our analysis we found that although types of operator task may be determined by patterns of gaze and behaviour, operators' tasks need to be analysed individually. Identifying whether specific data patterns related to a particular task needed to be verified using the eye tracking video to determine which application was open on that operator's screen at the time of recording.

— *2.2. Exploratory Data Analysis*

150

As our study was performed 'in the wild' (i.e. in the control room while operators were performing real-time monitoring of the network), we could not control the content of the study. Operators' activities depended on the real-time situation and their habits. Thus, collected data is rather unstructured and noisy, and cannot be analysed with predefined assumptions about what task the operator is performing at the time. Initial visual exploration of the dataset is required to identify AOIs and TWIs containing interesting activity, before comparative or summative analysis can be performed.

Pirolli and Card divided the process of exploratory data analysis in two main loops [7]: (1) the foraging loop, in which the user actively seeks information, by searching and filtering the data, and (2) the sensemaking loop, in which the user extract pieces of knowledge from the data and understand the connections between them.

In each of these loops, visualisation plays an important part. In the foraging loop, it can allow for visual filtering and highlighting, but also for visual aggregation of similar data in order to highlight outliers, this helps the user to extract the relevant data points [8]. In Jigsaw [9], Stasko et al. proposed to use connected views to interactively highlight the connection between the different data point and their dimensions.

In the sensemaking loop, visualisation allows the user to offload information from their working memory, by, for instance, allowing annotations and manual

170

selection of information. To facilitate this offloading, but also support collaboration, Dunne et al. proposed to save intermediate state of the analysis and expose the exploration process using a graph [10].

175 The overall workflow envisaged by Pirolli and Card links the *foraging* and *sensemaking* loops with a cycle of *schematizing and search for evidence*, as summarised in Fig. ???. Schematizing elevates observations from the foraging stage to knowledge ready for sensemaking, while search for evidence occurs when the analyst returns to foraging to assist further knowledge discovery. We broadly classify analysis
180 activity for ‘in the wild’ eye-tracking data into these four main sets of activity within the Pirolli and Card sensemaking loop, which for brevity in our later analysis we label FOR, SCH, EVI, and SEN, as follows:

FORaging: Analysts need to be able to filter and aggregate data interactively. Each session in our study created a large amount of eye tracking data. In order to
185 understand patterns and structure in the gaze trajectories—but also to highlight outliers—it is important to be able to visually filter and aggregate them [8].

SCHematizing and Search for EVIDence: analysts need to be able to flexibly encode activity to highlight and organise patterns (*SCHematizing*), and interpret them in the context of the original activity (*search for EVIDence*), for
190 example, through viewing the patterns against the original video footage of the analyst’s point of view [7].

SENsemaking: When dealing with a large quantity of data, analysts’ working memory relating visual patterns (or *evidence*) to insights (*validated hypotheses*), can quickly become overloaded [7]. Annotation can be used to offload this
195 working memory directly to the visualisation. As analysis is detailed, takes time and needs to be reproducible it is important to be able to save the state of analysis including the encoding and analyst’s notes. It is also important to be able to easily share between collaborators in order for other members of the team to observe and comprehend the steps of analysis that have been performed
200 and interpret the patterns with their own knowledge [10].

3. Related Work

In this section we report on previous work that has inspired us to build our system. Although no single prior system provides the full set of features that we require nor the novel features we have developed, VETA takes advantages
205 of existing gaze data visualisation and analytic tools and extends them with novel interactive features (e.g. interactive lenses, animations, bundling, small multiples).

In traditional eye-tracking analysis one would segment the session into pre-defined TWI, and define AOI in the workspace. Typical metrics include the time
210 spent on each AOI during each TWI, or the number of transitions between AOIs. Current eye-tracking analysis software provides tools to define these TWIs and AOIs, as well as overview visualisations such as heatmaps and gaze plots. Our system is not based on the exploration of static AOIs, but on the combination of recorded video and dynamic Focus plus Context (F+C) techniques such as
215 interactive lenses to leverage the user’s ability to gain a better understanding of recorded gaze datasets. With complex and noisy data the context is fundamental in understanding the detail of locally investigated eye-tracking data. F+C techniques has previously been shown to reduce the noise and facilitate visual clustering of the data (e.g. MoleView [11])

Gaze data visualisation and analytic tools: Eye-tracking technology
220 has a long history of capture and analysis of user gaze behaviours [12, 13]. Commercial solutions are available, covering the whole data processing pipeline, from data capture to cleaning, processing and presentation, including Tobii Lab ⁴, SMI be Gaze; Blickshift ⁵ and Eyevido ⁶. Open-source solutions exist
225 such as Ogama [14] and Web-based solutions are also available [15]. These gaze data analytics tools collect a large body of records consisting of fixation points linked by paths. These tools can visualise and extract relevant user patterns

⁴<https://www.tobii.com/>

⁵<https://www.blickshift.com/>

⁶<https://eyevido.de/>

Table 1: Summary table of systems providing both point-based and AOI-based visualisations of gaze data extracted from Blascheck et al. [26], with the addition of VETA. Note: The purpose of this table is not to show the complete classification, but only the dimension for which there differences between the systems.

	Gaze Data		Visualisation	
	Temporal	Spatial	Animated	Static
Stellmach et al. [27]	●	●	○	●
Weibel et al. [28]	●	○	●	●
Pfeiffer [29]	○	●	●	○
VETA	●	●	●	●

via heatmaps [16], scanpaths [17, 13], AOIs [18, 19], Sankey diagrams [20] or can automatically annotate interesting video sections [21]. Andrienko et al. used
230 geographical visual analytics approaches to explore gaze data based on AOI [22]. Blascheck et al. investigated multiple source aggregation among eye tracking data and think-aloud to evaluate an interactive system [23]. Other systems use AOI and space-time cube data representations to investigate eye-tracking datasets [24, 25].

235 Blascheck et al. proposed a classification of eye tracking visualisation systems in a recent survey [26]. They first classified the systems into *Point-based* and *AOI-based* visualisations, and only 3 systems propose both type of visualisations, which is also the case of VETA. Table 1 proposed a summary of Blascheck et al. classification for this three systems and we compare them with
240 VETA. Stellmach et al. proposed a system to visualise both spatial and temporal gaze data for 3D Virtual Environments [27]. The system provides an object-based visualisations of the gaze as well as a surface one (i.e. a heatmap) and a timeline indicating which object is in focus, all of this is static. In their system [28], Weibel et al. focused mostly on temporal data and visualise them
245 using a timeline and scanpath that can be animated. Finally, Pfeiffer proposed animated techniques to visualise spatial and temporal gaze data in 3D environments [29]. The author used both 3D heatmaps and 3D scanpaths to

represent the gaze data. Compared to these systems, VETA can manage both temporal and spatial data and uses both animated and static visualisations. The reason to propose a various panel of techniques is to allow users to choose the most appropriate to their data exploration. Compared to the large majority of these systems, our work does not focus on the use of static AOIs but on the combination of dynamic and interactive lenses together with the ability to switch to the recorded video to support the understanding and interpretations of the data.

Focus plus Context trajectoryflow analytic techniques: Heatmaps, scan-paths, transition matrices and charts of AOIs provide useful visualisation techniques to support insight discovery from gaze data. However, they always provide a trade-off between the visualisation of the raw data (often dense and complex) and a given level of abstraction which gives a specific point of view of the data. One possible solution to such limitations is to use Focus plus Context (F+C) techniques which show dual data visualisations at different levels of detail. F+C techniques emphasise specific areas [30] while showing contextual information. A survey across application domains is available which also investigates lens deformations [4].

Lense have been widely used as a focus+context technique, especially with graphs mostly to locally filter edges that can quickly clutter the visualisation. Lenses can act on two different elements, nodes or edges. Tominski et al. presented a lens of each type: one that brought neighboring nodes closer to the node inside the lens and one that remove all the non-local edges that went through it (a local edge is an edge linked to a node inside the lens) [31]. Most latter work tended to focus on edges as nodes position can encode data and thus represent information (which is the case for gaze trajectories). With Edgelens [32], Wong et al. proposed to curve the edges that went through the lens instead of removing them. The same concept is adapted to only non-local edges by Schmidt et al. [33]. In EdgeAnalyzer [34], Panagioditis et al. explored the use of multiple lenses in order to perform set operations (e.g. to filter the edges that went through two different lenses). The techniques also allow to fil-

ter inside the lens using different parameters including their direction and other
280 non-geometrical data dimensions. Finally it allowed to bundle locally edges in-
side the lenses. Finally, with MultiLens, Kister et al. combine most of these
different lens techniques in a customizable one [35].

Trajectories can be considered as a specific type of graph in which nodes rep-
resent waypoints and edges is the path interpolation between them. Early works
285 on lenses for trajectories mostly focuses on 3D flows visualisations. Fuhrmann
and Groller proposed a 3D lens which act as a clipping plane and show trajec-
tories with a higher resolution when they are behind the lens [36]. Mattausch
et al. proposed a similar solution with a 3D box as a lens [37]. In Flowlens
[38], Gasteiger et al. do not filter trajectories with the lens, but change its
290 encoding and add another data dimension locally that they consider as contex-
tual information. In TrajectoryLenses [39], Krüger et al. took inspiration from
EdgeAnalyzer and explore the use of several lenses to perform set operations on
2D trajectories. They allow the user to filter trajectories by origin, destination
or waypoints, all of which defined using lenses. Their system also provide a
295 timeline to visualise trajectories distribution by hours, days or months. Finally,
Ferreira et al. presented RoseTrajVis [40], a techniques that provides a sum-
mary of the information about the trajectories that go through a lens instead
of filtering. A colored rose diagram shows the distribution of direction of the
trajectories and the average speed for each.

300 We took inspiration from most of these work on lenses in the design of our in-
teractive AOIs and extend them by adapting them for gaze data. We explore, in
this paper, the use of these interactive AOIs to perform complex filtering opera-
tion and show contextual information regarding users' gaze behaviour (including
their use of glances), and its spatial and temporal distribution. In addition to
305 interactive lenses, we use other techniques to help the exploration and the un-
derstanding of gaze dataset using, for instance, animated transition [41] and
trajectories bundling [42].

VETA uses F+C to allow the user to investigate AOIs while keeping contextual
information and showing gaze flow directions to indicate where the gaze comes

310 from and goes to. While trajectories/flows can be visualised statically thanks
to their characteristics: origin, destination, intensity and clusters [43], other
techniques use animation—such as particle systems [44]—to show flow parameters.
VETA extends known features with the combination of animation and interactive
F+C techniques.

315 **Animation and view linking:**—Animated visualisations are common in
data exploration [41] and they are attractive for data presentation and communication
purposes [45]. For instance, animated transitions can aid the perception of
changes of the underlying data [46, 47, 48, 11, 49] and thus they have been
widely used for animated visualisation. Heer and Robertson [50] introduced a
320 taxonomy of transition types and found that animated transitions can significantly
improve graphical perception. As a global level, effective animations have
to follow two principles: *Congruence* and *Apprehension* [51]. The former
suggests that the animation must relate to the information to be conveyed,
while the latter suggests that the animation should be easy to comprehend.
325 VETA uses animation to ease transitioning between different visual mapping,
which helps the user keep track of the data of interest. Such smooth transition
can also operate to smoothly remove clutter in dense views thanks to visual
simplification.

Visual simplification tools:—Gaze datasets are often dense with many
330 fixations and saccades, which can create visually cluttered views. Specific F+C
techniques showing animated visual simplification [11] can help to partially
address this, to make patterns in the data more visible. More generally with
trajectory data, researchers explored how to use the third dimension to visualise
individual trajectories at the same time as aggregated ones [52], and they also
335 explored the use of interactive lenses to show details at specific locations [39].
This has been adapted for eye-tracking data in SaccadeLenses [53].

A class of techniques known as edge or trailset bundling can simplify the
dense and complex data produced by eye-tracking systems, Lhuillier et al. [42]
and Zhou et al. [54] provide a survey. A large range of algorithms exist to create
340 such bundling of trailset data, beginning with an assumption of a hierarchical

data structure [55], to more recent, faster and more generally applicable scalable methods [56, 57, 42]. Recent advanced techniques capture the statistical properties of the visually aggregated bundles [58] or a patchwork of multiple technique aggregation [59]. Such techniques have also been applied to static visualisation of bundled set of user gaze when investigating junk charts [60]. VETA extends such technique with its usage for real-time edge bundling technique [61] to aid the user interpretation.

4. Design and Development

VETA provides a combination of exploratory visualisation and interaction capabilities to allow the user to gain a quick visual overview of the data as well as to explore user gaze behaviour in detail. This has been implemented in the following manner. A video showing the system functions and interactions accompanies can be found in our supplementary web link⁷.

4.1. Data Processing and Implementation

Three kinds of data are utilised: (1) eye tracking data of a sequence of fixations, each containing time of recording, an X and Y coordinate in space, and a $\Delta time$ representing how long the subject fixated at that point. This data is stored as a CSV file, (2) a background image as spatial reference to give the user context and (3) a video of the subjects' actions can be attached from the original eye-tracker recording. The eye tracking data can be explored without the mapped image or the video, however these provide further context for understanding the behaviour and verifying patterns and insights.

Mapping gaze data onto a reference image of a real-world environment is an increasingly common technique included in eye tracking software. Tobii Pro and other eye-tracking providers allow for tracking data to be mapped onto a still image of the environment, and gaze data of computer screens will usually be calibrated to be recorded in terms of screen coordinates, while gaze data

⁷<https://youtu.be/NsMP6eDtvGM>

on a scrollable webpage might be given in page coordinates, accounting for scrolling. This can be done manually, or using image processing algorithms to
370 find matching regions on the reference image [62].

The mapped gaze coordinates produced by the eye tracking software can then be formatted for importing into VETA. We used Tobii Pro Studio for the initial mapping of the data which is then exported to a CSV files which can be imported in the VETA application.

375 For the VETA application, we implemented the application in Processing 3.5.3⁸ for speed of data processing as well as flexibility of the graphical interface. The Processing Video library is used for accessing video files, the Java Swing library for file access and the Java OpenCL⁹ library for accessing GPU power for implementing real-time dynamic edge-bundling. It was developed and runs
380 well on a standard computer. The code is open and available on GitHub¹⁰.

4.2. Interface Design

Data is presented spatially on a *canvas* (i.e. the area of the workspace in which the gaze points are displayed) and temporally on the *timeline*.

The prototype includes controls for filtering and manipulating parameters
385 of the data as well as a legend to distinguish colours and filter by colour. These features can be turned off and on, and also have keyboard shortcuts. Eye-tracking events are shown spatially on the canvas on top of the image to which the data is mapped, as explained in Sec. 4.1. Eye-tracking events are also shown chronologically on the timeline beneath the canvas (e.g. see Fig.3-c). Each time-
390 slice is presented as a vertical rectangle, where the width is proportional to the length of the recording divided by the screen width.

Interaction in the canvas—i.e. filtering data or changing the colour mode of presentation—are reflected in both the canvas and the timeline through linked views. Visible data can be filtered spatially by the creation of user-defined *lenses*

⁸<https://processing.org/>

⁹<http://www.jocl.org/>

¹⁰<https://github.com/rtwhi2/VETA>

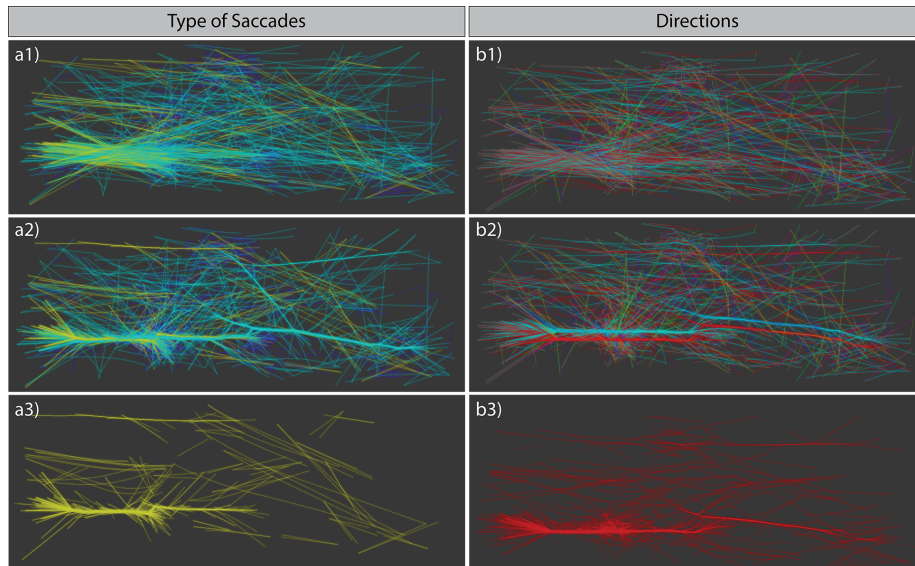


Figure 1: Filter and aggregation of saccades using type and directional colour modes a) by saccade type: short (blue), *glance* (yellow) and other (turquoise) and b) coloured by direction with 8 subclasses (refer to colour wheel in section 4.2.1. 1) all saccades, 2) bundling with colour mode, 3) bundling filtered by colour, a. all *glances*, b. all right saccades

395 defining AOIs on the canvas. Visible data can also be filtered via the timeline to a specific TWI using interactive range sliders. The user can zoom into this selected TWI. A horizontal line (e.g. see Fig.4-a2) divides the timeline. Above the line provides visual information about fixations (in each AOI). Below the line provides information about activate saccades. Filtered data is shown in
400 grey. User defined AOIs are also presented on the timeline in horizontal rows to allow for exploration of AOI information over time (e.g. see Fig.4-a2). Finally, the user has the ability to switch between the data-exploration mode and the eye-tracking video at any time to verify their insights.

4.2.1. Visual Encoding and Clutter Reduction

405 Immediately on top of the mapped image all fixations in the full dataset are displayed for *context*. This provides an overview of the activity to enable

natural clusters to be identified instantaneously. The user has the ability to turn this context layer on and off during their exploration process. These fixations are shown on the canvas as transparent circles (with 20% transparency). The appearance of the circles can be altered to present them either equally sized or sized proportionally to the length of the fixation. The maximum circle size can also be adjusted manually.

On top of the context layer, saccades between the fixations are presented as lines linking fixation locations on the canvas. We distinguish two specific subtypes of saccades that were identified during our design phase: *Short-distance saccades* and *glances* (Fig. 1-a). [Glance is a common type of gaze pattern used to analyse users' behavior in different domain: piloting aircrafts \[63\], construction \[64\] and command and control contexts \[65\]. It can be used to differentiate between novice or expert users, to understand better their attention \(and when it is misplaced or blurred\), or to analyse and their performance for a specific task.](#)

Short-distance saccades often represent the subject's eye moving within a small area (e.g. an AOI), rather than movement between components or to new locations [13]. The distribution of short saccades can therefore help highlight where natural AOIs supported by the data exist, yet they provide little information about the interaction between those AOIs. By default short saccades are defined as 100 pixels in length as this was determined as appropriate for our use case but the user has the ability to alter this value. If the minimal distance is reduced to zero no sub-selection takes place and all saccades are classed as *other saccades*.

A common pattern observed was a *glance*, where the subject fixates in one area, then moves to fixate in another and then immediately moves back to near where they were fixating immediately before. If the subject repeatedly moves back and forth between two components, without looking anywhere else, these are considered as *glance* saccades. We found that this pattern is fairly common in our data and provides distinct information about the relationship between components. For example, subjects might *glance* to a component that holds

information necessary for a task in another component, or that might contain an alarm message that needs to be checked regularly. Since these two types give
440 different information from other saccades, it is possible to filter them separately.

Saccades can be filtered by *type*. Type is also distinguished through the use of colour: short (blue), *glance* (yellow) and other (turquoise). Apart from these default colours for saccades, further colour modes allow for more in-depth exploration. To allow for directional gaze patterns to be explored, a directional
445 colour schema is provided. A colour wheel is divided into 8 classes, to allow for visual clustering of saccades by general direction of eye movement. Using this schema, horizontal and vertical movement in different directions can be visually distinguished by colour, as shown in Fig. 1-b, where red indicates the right saccades and turquoise the left ones, etc.

450 Direction of saccades can also be presented using animation, where fireflies animate along the saccade rather than shown as a solid line. Finally, a custom colour schema provides the ability for the user to define their own categories (using an additional column in the input file). This custom mode could be used for common tasks or TWIs for instance.

455 To allow for overarching patterns to be observed in the saccade trajectories even when significant amounts of data are visible, dynamic force-directed edge *bundling* can be performed on the fly [61] (Fig. 1-a2/b2). This is applied on the saccades that are currently visible for instance on all saccades, those in a certain TWI or those within a particular AOI. Bundling is also effected by the
460 colour schema: if colour is by type, the *glance* and normal saccades will bundle separately, and if by direction, lane-forming directional edge-bundling is used so that the weight of transitions in each direction can be compared. The degree of rigidity in the saccades being bundled can be controlled by the K parameter through the interface. Bundling can be reset at anytime. Fig. 1-a2 shows how
465 bundling linked with the colour mode can help to remove saccade clutter. More specifically, Fig.1-b2 demonstrates that directional edge-bundling can be used to help identify user behaviour patterns in the data visually. By clicking on one colour in the legend, it is possible to show only the saccades of that colour

(Fig 1-a/b3).

470 4.2.2. Exploring Areas of Interest

Areas of interest (AOI) are defined by the creation of interactive lenses that can be manipulated manually using the mouse. The user has the ability to create, reshape, select or delete them. They are shaped as superellipses, ranging from diamonds to rectangles, where the user has the flexibility to control the size, shape and aspect ratio. In each AOI the percentage of the fixation duration is displayed as percentage value. As this is moved around the screen, this value recalculates on the fly.

Gaze datasets are often dense with many fixations and saccades, which can create visually cluttered views. The use of interactive lenses in for eye tracking have shown to be promising [53]. AOI can be used to reduce data clutter and filter the data for exploration in several ways. Using AOI saccades can be filtered using multiple filter modes: within an AOI, around an AOI, from/to one AOI or between multiple AOIs. AOIs can also be used to remove fixations shown on the canvas to remove outliers or unnecessary data. Fixations can be removed from view using two options: within or outside an AOI. This action can be reset at anytime to return to the full dataset.

All fixations within an AOI are highlighted on the canvas using white opaque circles. These are 30% of the size of the context circles shown in the background. This allows the user to identify visual clusters within the AOI easily, alleviating issues of occlusion due to overlapping fixations. Background context fixations can be turned off or on to aid comprehension and exploration at this time.

When an AOI is activated the interactions between the AOI and surrounding space are highlighted in the foreground. Selected fixations outside the AOI are shown with 80% transparency, where three types of fixation are distinguished through the use of colour; Fixations which occurred shortly *before* a visit (maroon), fixations shortly *after* (green), and fixations representing a *glance* both in and out of the AOI (where the subjects attention moves from the AOI, to the location and back to a location within the AOI) (yellow), as shown in Fig. 2.




Figure 2: Left: the generic view provides the fixation before and after each fixation inside the AOI. Right: the alternate view provides all the fixations in a chosen amount of time before and after each fixation inside the AOI (in this example, the chosen amount of time is 5 seconds).

Two different interpretations of ‘shortly before’ are applied. In the *generic*
 500 *view*, the immediate sequence of fixations is used, with only a fixation immedi-
 ately before or after a fixation within the AOI being marked (Fig. 2-left). In
 the *alternative view*, a user-controlled window of time is provided, where all
 fixations within a given number of seconds before or after a fixation within the
 AOI are highlighted (Fig. 2-right). To quickly distinguish between these two
 505 modes a circle is used to present fixations in the generic view, whilst a cross is
 used in the alternative view.

Each AOI is identifiable in the timeline, allowing the user to see AOI selec-
 tions over time. The fixation section of the timeline, above the horizontal line,
 is split by the number of AOIs on the screen. These appear in order of cre-
 510 ation. The selected AOI is highlighted, and the selected filter and colour modes
 are reflected both in the canvas and timeline. To allow for quick exploration
 of selected data over time, hovering the mouse over a fixation activates a line
 connecting the fixation to the timeline. To provide a visual overview of the
 distribution of transitions the selected fixations can be presented spatially, by
 515 direction or by time. The “by direction” mode maps the surrounding fixations
 around the AOI perimeter (Fig. 3-b and Fig. 4), whilst “by time” maps the
 surrounding fixations chronologically along the timeline (Fig. 3-c). These his-
 togram views are implemented using animated transitions to allow the fixation

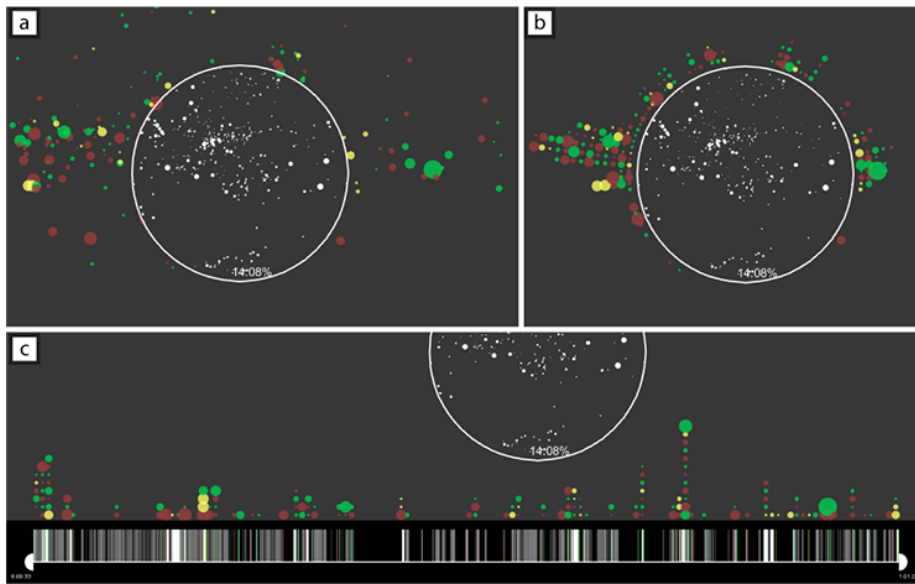


Figure 3: Circular AOI presenting fixations within the AOI in white and selected fixations outside the AOI in three colours: Shortly before (maroon), shortly after (green) and *glances*, both out then in again (yellow). These surrounding fixations sized by duration of fixation time can be presented (a) spatially, (b) by direction or (c) by time (where the white bars indicate a fixation at that time).

locations to be visually followed, as shown in the accompanying video. Such
520 animated transitions have been shown to aid the perception of changes of the
underlying data [46, 47, 48, 11, 49].

To allow scan paths of the subjects' attention to be explored. A *travel-line*
can be drawn between the 'before' fixation and the matching 'after' fixation.
In the alternative view, where multiple before and after fixations can exist for
525 each visit to the AOI, the entire saccade *trajectory*, within the window of time
around a visit to the AOI defined by the user, can be drawn. This allows scan
paths to be investigated, however, large AOIs or TWIs may lead to the view
being overly cluttered.

4.2.3. Provenance and Sense-making

530 In order to ensure the system is useful in practice we have incorporated a
number of features to aid dataset comparison, user collaboration and report
making. VETA provides the ability for users to explore multiple datasets in
one view in juxtaposition. The screen can be divided into four small-multiple
sections, and each small multiple can present subsets of the same dataset (i.e.
535 different AOIs or TWIs) or different datasets (i.e. different participants). Se-
lection and filtering within each small multiple can be manipulated separately
or they can be explored all together. The user has the ability to zoom into one
view for exploration and move back to the multi-view easily. At present VETA
is limited to four juxtaposed small-multiple views but this could be extended if
540 necessary and if system resources allow.

The user has the ability to load different datasets and to save workspace
settings, this includes all selections and filtering, including AOI, TWI, bundling
and colour mode settings. Saved workspaces can be loaded into VETA at any-
time and when saved or loaded they appear as a sequence of small thumbnails on
545 the right side of the screen. This allows for a smooth exploration process being
able to move between workspaces quickly. In addition to saving settings record
making is essential in the data exploration process. VETA allows user-defined
text notes to be created at any time, anywhere on the canvas or linked to the

550 timeline. These notes appear as the top layer and can be edited or deleted at
anytime. This not only allows the provenance of the exploration process to be
recorded but allows for collaboration with others. Finally, the workspace can
be saved by the user at anytime.

5. Use Case: Analysis of a Network Control Room

In the context of the study presented in Section 2 we recorded gaze (using
555 a wearable eye tracker) of a network operator during the first hour of their
shift, termed as the *handover* session between the night and day shift, on two
consecutive days. During the handover, the night shift operator gives a summary
to the day shift operator (our participant) of what happened during the night
and of what is expected to happen during the day. The day shift operator then
560 sets up their workspace and starts working. It is important to note that Day 1 of
the study was the first day of the day shift cycle of our participant, which means
they had just had 5 days off and therefore had limited situational awareness of
the network. We used VETA to compare the gaze between these two sessions
to assess the impact of this lack of situational awareness on operators activities.
565 For this analysis we refer to each screen as Screens 1-7 as shown in (Fig. 4-
a1). For each action, we give in parenthesis the functionality we used from the
system ~~and which step from the sensemaking loop it is~~. In the last subsection,
the comparison, we link our findings with the *outcomes* from Section 2. Each
session lasted around 1 hour and we have around 2500 fixations in each, which
570 represent almost one fixation per second. On the first day, there were 326
transitions between two screens, which is 5.4 transitions per minute. On the
second day, there were 256 transitions between two screens, which represents
4.2 transitions per minute.

5.1. Analysis of Day 1

575 Using VETA, we loaded the mapped data with the appropriate background
image. We started by displaying only the short saccades (Saccades filtering—




Figure 4: Comparison of the handover session on Day 1 (column a) and 2 (column b) for the same operator. (a-b1) Directional bundling of all saccades for both sessions. (a-b2) Fixations (green, yellow, maroon) and short saccades (blue) with 4 AOIs for both sessions. (a-b3) Fixations and short saccades with 4 AOIs for the Phase 1. (a-b4) Fixations and short saccades with 4 AOIs for both sessions. (a-b3) Fixations and short saccades with 4 AOIs for the Phase 2. Annotations have been added for descriptive purpose and enlarged for readability.

FOR), to see which area of the workspace the participant focused on the most (Fig. 4-a2). We could see that the participant mostly focused on 4 screens during the session (see Fig. 4-a1 for screen number). We then created an interactive
580 AOI for each screen to get a more accurate idea of relative time spent across screens: Screen 2 (42%), Screen 7 (18%), Screen 5 (12%), and Screen 1 (11%) (Interactive AOIs—**FOR**). By looking at the timeline, we saw a 5 minute period during which no data was recorded. By looking at the video, we can easily see that the participant took a coffee break during this pause (Video exploration
585 —**EVI**). Considering the evident difference in screen focus between these times we decided to divide the session into two phases for analysis: *Phase 1* before and *Phase 2* after the break.

We focused first on Phase 1, and thus adjust the slider on the timeline to only display gaze data during this period (Fig. 4-a3) (Time filtering—**SEN**).
590 The AOIs gave us the percentage of screen usage. In this phase, the participant was 45% focused on Screen 7, 27% on Screen 1, and around 10% on Screens 2 and 5. We then displayed the long saccades and filtered the ones which were not between the AOIs (Saccades filtering—**FOR**). We observed very few transitions between the different screens. This is confirmed by visualising the fixations
595 sequentially (i.e. using the colour codes differentiating before/after/glance, as shown in Fig. 3) for each AOI (Before/After glances—**SCH**). This allowed us to also notice that the participant glanced a few times at the Screen 6 while working on Screen 7. A look at the timeline confirms that during this period, the participant focused on each of the 4 screens for long periods of time, which there-
600 fore reduce the need to transition between them. A look at the video showed us that the participant was actually setting up the different applications during this time (Video exploration—**EVI**). They had to open and prepare around 15 applications, which explains the long periods of time spent on each screen. This process was mostly managed from Screen 7 on which the applications were
605 initially opened.

Keeping the same AOIs, we then adjusted the timeline to Phase 2 (Fig. 4-a4) (Time filtering—**SEN**). The screen usage radically changed: Screen 2 is

used around 50% of the total time, while screens 1, 5, and 7 are each used for around 10% of the time. When looking at the transitions (Saccades filtering **—FOR**), they mostly happened between Screen 2 and screens 1 and 7. By activating the directional histogram around the AOI on Screen 2, we can see that more than half of the transitions with Screen 1 are glances (Histogram view **—EVI**). Checking the video allowed us to see that the participant had the real-time network diagram and the alarm summary applications on Screen 2 (Video exploration **—EVI**). We assume they are using this screen to gain an awareness of what happened before their shift (e.g. previous alarms, part of the network under maintenance), which is one of the main goals of the handover session. This required the participant to glance at Screen 1 to get additional information about the different events that happen in the previous days. Transition to Screen 7 to look at a geographical map with natural risk information (e.g. lightning, fire etc) on the monitored sector took more time than a simple glance.

5.2. Analysis of Day 2

We then loaded the data of the second day of recording and mapped this to the same background image, as the participant arranged their workspace in very similar ways on both days (Fig. 4-b2). When looking at the timeline, we can see there is a break similar to day 1 after the first 15 minutes. A quick look at the video showed us that again the participant took a coffee break (Video exploration **—EVI**). The overview of the short saccades seemed to be slightly different than in day 1 (Saccades filtering **—SCH**). We placed the same AOIs and observed different focus time proportion: the participant spent around 32% of the total time on Screen 5, 17% on both Screen 2 and 7 and only 6% on Screen 1 (Interactive AOIs **—EVI**).

We focused the timeline on Phase 1 (before the break), and observed a pattern similar to the previous day (Fig. 4-b3) (Time filtering **—SEN**). The participant spent 30% of the time on Screen 7, with several glances to Screen 6, 20% to Screen 5, 7% to Screen 1 and 4% to Screen 2. By briefly looking at the video, we could see that the activity was very similar to the previous day

(Video exploration—**EVI**). The participant set up their workspace during this phase.

640 Most of the difference with the previous day happened in Phase 2. Instead of mostly focused on Screen 2, the participant spent 40% of the time on Screen 5, and only 25% on Screen 2. They focused on Screen 7 10% of the time and Screen 1 4% (Fig. 4-b4). Interestingly, when looking at the timeline, we could see that most of the time they spent on Screen 2 happened in the beginning of
645 the Phase 2, with glances to Screen 1. The video helped us confirm that their activity was similar to the day before (Video exploration—**EVI**), they were gaining an awareness of events that had happened prior to their shift. The rest of the phase is spent mostly on Screen 5. On this screen, the participant had placed mostly typical ‘office’ applications, such as their e-mail client and web
650 browser. Checking the video showed us that during the rest of the phase, it was a quiet period for real-time operation tasks and therefore the participant was mostly occupied (with some exceptions) with e-mails, research for longer-term projects and other long-term planning tasks. Tasks expected during quiet periods of normal operation.

655 5.3. Comparison

The participant is expected to do the same tasks in both handover sessions. We could see a similar overall pattern in time management: after a verbal handover with the night operator, the participant started by setting up their workspace and then took a quick break prior to starting their daily tasks. How-
660 ever, the gaze trajectory tells a different story (**Following Outcome 3**). This is highlighted by a side-by-side comparison of the trajectories of both sessions aggregated using directional bundling (Fig. 4-a1 & b1) (**Interactive bundling—SEN**). We can observe that on Day 1, there is a clearly visible trajectory from Screen 1, to 2 and on to 5 and 7. On Day 2 the main branch of the trajectory
665 goes from Screen 6 to Screen 7(**Following Outcome 1**). There is a small one from Screen 1 to 2, but less distinct than on Day 1 (**Following Outcome 3**). Analysis of the monitoring of the control room sessions in the afternoons of both

days, with the same participant, showed us a very similar trajectory pattern to the latter stage of Phase 2 of the handover session on Day 2.

670 The reasons for the difference between the two sessions lies in the fact that the participant is coming back from the required 5 days break in their shift cycle. On the first day, they need to gain a far greater awareness of what happened in the last few days and what was anticipated to occur in the next few days. By contrast, on Day 2, they only need to get an awareness of what happened
675 over night. As the night had been relatively quiet and nothing unexpected had occurred, this takes only a short amount of time at the beginning of Phase 2. During the rest of this session they undertook the regular work that is expected during ‘uneventful’ network monitoring times (similar to what we observed in the afternoon sessions, as part of the wider study). Interestingly, the handover
680 session on Day 1 seemed to leave a more positive imprint on the participant (as shown in the survey data), likely this is because they had had a break and the handover was more out of the ordinary than the subsequent day. Using the annotation functionalities, we added these insights directly overlaying the data, similar to what has been done in Fig. 4 (Annotation—[SEN](#)).

685 **6. Experts’ Feedback**

To evaluate our system and gather feedback from other eye-tracking data analysts, we ran 14 individual feedback sessions. Participants were all eye tracking data users (i.e. they had used eye tracking data for their own work) with different levels of expertise. All participants did one session. This study was
690 approved by Monash University Human Research Ethics Committee (Project 23671). The purpose of the session was to show the functionalities of the tool with our dataset, understand the activities of the participant and get their feedback about the tool itself.

6.1. Procedure

695 All the sessions were performed remotely (due to COVID-19 restrictions) using video conference software. The first stage of the session started with a

brief introduction of the purpose of the tool and the study. Then, participants introduced themselves, their main activities, and how they use eye tracking (i.e. for what goal, which software tools they used, etc.). In the second stage of the interview we showed a video of the functionality of VETA, pausing when necessary. Finally, the third stage consisted of a discussion around the different ways such a tool could help them analyse their own data. One experimenter shared their screen with the tool running to answer specific questions regarding any functionalities or demonstrate specific processes. In this session, we used data from the use case presented in Sections 2 and 5. A session lasted roughly one hour and was audio and video recorded. After the session an online questionnaire was sent to the participants. This consisted of questions about the specific tools they use in their analysis and to gather post-interview feedback, allowing the participant some time to reflect.

6.2. Participants

We recruited 14 participants (Table 2 for details about their occupation and type of studies done with eye trackers), 11 males and 3 females, 2 aged between 18 and 24, 6 between 25 and 34, 4 between 35 and 44, 1 between 45 and 54, and 1 over 65. 4 participants declared using eye tracking data at least once a week, 4 at least once a month and 3 just once a year. 6 out of 14 participants mostly used eye tracking data in controlled experiments as a measure to perform comparisons, for instance to compare graph visualisations (P2), or to test different interfaces (P4). 9 out of 14 mostly used it in more observational studies of, for instance, students following a class (P5) or air force pilots training (P7). The analysis took between days for some participants (5/14), to several weeks (4/14) or months (2/14) for others.

6.3. Reflections

During our interviews, participants explained how they currently use eye tracking data, gave feedback on the tools and how it could adapt to their own use case. Overall, all participants found it could be useful for their analysis to

Table 2: Expertise of the 14 participants of the expert feedback sessions and the types of studies in which they use eye trackers (O: Observational, C: Controlled)

ID	Domain	Studies
P1	Energy	O
P2	Information Visualisation	C
P3	Human-Computer Interaction	C
P4	Accounting	C
P5	Learning Analytics	O
P6	Neuroergonomics for Aerospace	C/O
P7	Human Factor and Human-Computer Interaction	C/O
P8	Machine Learning and Sensor Computing	O
P9	Information Visualisation	C
P10	Eyetracking Data Quality and AR UX	C
P11	AR/VR Innovations	O
P12	Simulation-based education	O
P13	Software Developer	O
P14	AR/VR Developer	O

“previsualise the data without any processing and explore the raw data” (P6) (**FOR**), “gain some intuition” (P4) (**EVI**) and “identify interesting events in the data faster” (P9) (**SCH-SEN**).

6.3.1. Current analysis tools

730 A bit less than half of the participants (6/14) reported using applications provided by the eye tracker manufacturers for their analysis. However, most participants also rely on *self-developed tools* (11/14). The reasons for this being that commercial software mostly handles only basic measures (2/14), and basic visualisations (3/14) and lacks important functionality for data exploration like
735 filtering data outside AOIs (P2) (**FOR**), setting up accurate TWIs (P7) (**FOR**), and good support for dynamic content like video (P9) (**EVI**). Participants implement advanced metrics on their self-developed tools, like entropy (P7), and

novel visualisations such as time plot [66] (P7), edge bundling of saccades [67] (P6), and space time cube [68] (P9). P14 also mentioned developing tools to
740 handle 360 degree images.

6.3.2. Data Exploration

The potential for eye tracking data exploration was identified by all the participants. They argued that the main feature that supports it is the interactive AOIs (~~FOR~~) (10/14). It is very easy to create them, and then compare the
745 quantity of gaze movement within them (P3). P5 and P7 also appreciated that it is possible to visualise only the transition between them (~~FOR~~). P6, who is particularly interested in visual screening of instruments by pilots, suggested that it would be interesting to be able to filter saccades that go through AOIs in a specific order (~~FOR/SCH~~). In general the transitions between inside and
750 outside of the AOIs were also considered as “*helpful to know the glances coming in and out of the areas of interests*” (P1) (~~EVI~~), more specifically the colour coding of the position of the fixation before going into the AOI and after (10/14). The histogram of such fixations around the AOI made it easy to understand (P4, P6, P10) (~~SEN~~). P10 stated that it was “*clear for people who do not want to
755 dive deep in the data*” and P4 imagined it would be convenient to ‘brush’ AOIs across the entire workspace with this histogram to identify interesting areas early on (~~FOR~~).

Filtering and bundling are functionality aimed at reducing visible clutter, however, participants thought more could be done to reduce it (~~FOR~~). P2
760 proposed to filter saccades inside AOIs, and P6 suggested to cluster fixation more. The latter is justified by the uncertainty of the position of each fixation themselves (due to error in the measurement, mapping, etc.). While our system proposes various shapes for areas of interest, P9 suggested that being able to construct arbitrary polygonal shapes would to be more precise and better suited
765 to their needs.

Another functionality considered interesting for data exploration is the timeline of fixations (10/14) (~~FOR/EVI~~). The linking of the timeline with the spa-

tial visualisation (i.e. colour replicated in the timeline) gives a good indication of the distribution of the different types of saccades, it is *“very useful as it helps to visualise changes in gaze patterns in a window”* (P8) (EVI). More specifically, P2 suggested that being able to see the temporal distribution of before and after fixations (in the context of an AOI) provides good insights regarding users’ strategies (SCH). P6 added that this was helped by the animated histogram which transitioned between the spatial view and the timeline (SCH). The definition of time windows around specific events can be important to understand specific aspects of the data: *“Sometimes, you want to explore 10 seconds or 1 minute before and after an event”* (P7) (EVI). The same participants found our tool allows for an accurate definition of such time windows compared to current manufacturer tools. As some of the events this participant focused on can repeat several times in one session, P7 would like to be able to define several of these windows in the timeline (TWIs).

Finally, 10/14 participants valued the ability to link the real video of the trial to explore their eye tracking datasets (EVI). Indeed, as P5 explained, lots of use cases imply a dynamic background, which is necessary to provide context information but hard to visualise with the data. In VETA, it is thus useful to have access to the raw video in order to *“help answer questions about specific behaviour”* (P7) (SCH) and to *“link the data with the precise stimuli”* (P8). P7, who analysed individual pilots behaviour, also envisioned using it to first provide quick analysis of the user behaviour (FOR), before visualising the data (EVI). Similarly, P8, who uses eye tracking to classify learning behaviour, considered the video as more important, and would use the data as a secondary source of information, a way of giving a *“bigger picture”* to the video (SEN). Finally, P10 suggested that it would be a good way of identifying when participants have to move their head to look from one area to another, which can be physically tiring (EVI).

6.3.3. Data Comparison

To compare data, and more generally gaze behaviour, is an important task in eye tracking data analysis [69]. Participants mentioned comparing between different conditions, or comparing directly between different participants (~~SEN~~).
800 Edge bundling to aggregate saccades was considered by several participants as a good way of visually comparing different gaze behaviour, to highlight general patterns, and to detect outliers (6/14) (~~EVI~~), it also “helps to get rapid insights on the connections between different zones” (P6) (~~SEN~~). Participants also suggest that more formal comparison can be done with VETA using the small
805 multiple functionality (9/14) (~~SCH-SEN~~). P2 stated that “Doing a side-by-side comparison for multiple participants at a time saves more time to identify user patterns”. While it is possible to load several users into our system, a majority of participants thought it would be useful to have an explicit participant management system (6/14). It would allow the ability to visualise specific
810 conditions, or even trials (in the context of a controlled experiment). P7 suggested it would be interesting to be able to select one pattern in one user and show all the other users who have the same pattern. P9 went a bit further and proposed that the tool could classify the participants in function of their gaze behaviour. On a similar idea, P11 proposed to be able to compare participants’
815 gaze to a “standard” trajectory, to, for instance, check if they follow the appropriate procedure responding to one event, or to compare to expert behaviour (~~EVI-SEN~~).

Most of the functionalities of VETA allow for visual comparisons, yet visualisation is not the only thing the participants need to perform their analysis.
820 The majority of participants expect to be able to get *quantitative* metrics about participants’ behaviour (7/14) (~~EVI~~). Eye tracking manufacturer software provide basic ones, but when needing more advanced metrics, participants had to develop their own tools. We asked which metrics participants would need, and we received a wide range of answers, including gaze distance inside AOIs (P2),
825 number of transitions between AOIs (P4) and entropy (P7). Yet, they acknowl-

edged it would be difficult to include all metrics in to one tool as people would always ask for more, and that other tools already provide these. To solve this problem, P7 suggested to implement a link from the visualisation tool to open source applications that provide eye tracking data metrics, such as Ogama¹¹.

830 6.3.4. Data Presentation

To present the results of the data analysis process to non-experts, visual representation of the insights are often used. Such visualisations are often hard to produce especially to illustrate complex gaze behaviours. For instance, P6 has to present findings regarding gaze behaviour of pilot trainees to their in-
835 structors, and to this end, provide them with “*Annotated improved heatmaps*”, which highlight important components of the gaze behaviour (~~SCH~~). In our system, complexity of the visualisation can be decreased using edge bundling and filtering, while annotation can be used to explicitly emphasise interesting components of the visualisation (7/9 participants found it useful). P7 went
840 further by explaining that VETA is simple enough to be used by non-experts allowing them to interact with the visualisations and thus grasp a better understanding of gaze data (~~FOR SEN~~). P9, however, warned that the amount of functionality provided could be overwhelming at first for non (eye tracking) experts.

845 7. Discussion

Throughout the design, development and evaluation of VETA, we have gained a wide understanding of the benefits of visual analytics for eye-tracking data.

Whilst we acknowledge there is further user-interface design and usability
850 work to develop the digital prototype further, we have identified clear potential for using such a comprehensive visual analytics system for complex eye-tracking data, for a wide variety of domains. Here, we reflect on our findings ~~and propose~~

¹¹<http://www.ogama.net/>

several design guidelines (shown in italics) for about visual eye-tracking analytics. Each of the guidelines cover all four of the stages (**FOR**, **EVI**, **SCH**, **SEN**)
855 from the visual analytics sense-making loop (see Sec. 2).

7.1. Fluid and Dynamic Exploration

~~Analysts need to be able to link the spatial and temporal dimensions of the data: We found that the linking of spatial and temporal dimensions through the visualisations provided useful information about the transitions between time and space and allowed the dynamic ability to link between AOIs and TWIs succinctly. All experts found this feature useful and would explore the data both over time and space.~~

~~Analysts need to be able to interactively create and manipulate areas/time windows of interest:~~
An important part of the data exploration process is to actively seek information, by searching and filtering the data [7]. In the context of eye tracking, the most common method of filtering the data is to create AOIs (and TWIs), for example, these may be defined prior to study to test hypotheses. However, to support more exploratory analysis after the study is performed, it is important to allow the analyst to interactively create and manipulate them. In our proto-
865 type the dynamic AOIs allows for a highly interactive exploration of data. The expert evaluation indicated that these dynamic AOIs enhanced the experience of AOI analysis over using traditional eye tracking software. Some also felt that hand drawn AOIs would also be useful.

We found that the linking of spatial and temporal dimensions through the
875 visualisations provided useful information about the transitions between time and space and allowed the dynamic ability to link between AOIs and TWIs succinctly. All experts found this feature useful and would explore the data both over time and space.

~~Analysts need to be able to filter and aggregate information quickly and interactively:~~
880 In order for analysts to understand patterns and structure in the gaze trajectories—but also to highlight outliers—it is important to be able to visually filter and aggregate them interactively [8] in general, with edge bundling [70] and with

gaze data [69]. We found the classification of saccades into ‘glances’ and ‘short distance’ were beneficial as a form of aggregation and provide additional context in our analysis. Bundling saccades on-the-fly provides the ability to see the bundling happening and thus helps the user understand how saccades are aggregated. In addition, combining type and directional bundling with their colour modes enables visual clarity to patterns (Fig. 3a).

Moreover the real simplicity of the directional histogram and time histogram, with animated transitions, were not only one of our favourite features of the interface, but a real hit with the experts. This demonstrates the power of simple visualisation of complex data. Many of the participants expressed that these simple, but dynamic views could be highly useful in exploring and understanding gaze behaviour in their own analysis.

~~*Analysts need to be able to access the context of the session easily.*~~ In traditional eye tracking analysis software it can be hard to link the real-time video with the gaze trajectory data during the analysis. It is therefore difficult to infer the behaviour of participants at a specific time thus understanding their actions, and motivations in the context of the entire session. In VETA, in addition to the typical mapped background image to provide context, we proposed the ability to display the video of the eye tracking device (or screen capture) at anytime during analysis. This allows for the analyst to better understand what the participant is doing during a specific period of time and verify the accuracy of the exploration insights.

The quick transition from the data to the video was appreciated by the majority of the experts who saw a way of validating insights and more generally, to complete the bigger picture provided by the data with details regarding the behaviour of the participants. In our design, switching between video and data visualisation is designed to fill the screen and provide short snippets of video context to validate insights. We used a fade-in transition between the background image and the video, but, to avoid too much disagreement due to the context switching, it would also be possible to use image recognition to detect the position of the frame of the video that user wants to watch and then inter-

polate a transition between the background image and this frame [71]. Another
915 solution would be to have a small insert screen showing the video that the user
can maximise and minimise when needed.

7.2. The Need for Comparison

~~Analysts need to be able to compare gaze behaviours:~~ Although the video and
fixation information were found to be useful context, the interviews also iden-
920 tified the importance of customised metrics to understand the context numeri-
cally. In the motivational use case, we were interested in dynamically exploring
the patterns of gaze behaviour and chose to add percentages of fixation time per
AOI as this was interesting for our analysis of times spent at each location of
their workspace. But many of our experts identified useful and different metrics
925 to add to the interface. Providing further metrics which calculate dynamically
could certainly be highly beneficial to the visual analytics process, yet further
analysis of current systems and a wider group of eye tracking users would be
needed to understand which measurements would be most useful in this context.
One suggested option is to link VETA to open software such as Ogama, which
930 would allow users to define their own measures. It could also be interesting
to implement more complex measures and provide sequence analysis to extract
viewing patterns [72]. Such measure has been shown to be beneficial in the
context of air traffic control [73, 74].

Whilst in our use case comparing between participants was not anticipated,
935 we did choose to implement small multiple views for comparing gaze patterns
of one participant across multiple sessions and TWIs (e.g. Fig. 4). Through our
evaluation we found that the small multiple comparisons were highly beneficial,
and the ability to switch quickly between small multiple and specific context is
also needed.

940 Most of the experts noted that they have to compare multiple participants
in their analysis. This could involve embedding a small participant manage-
ment system. This would allow users to aggregate and visualise participants, or
conditions at the same time or in small multiples. In addition, it would also be

possible to explore machine learning algorithms for the automatic classification
945 of participants.

7.3. A Comprehensive and Collaborative Tool

~~Analysts need to be able to save and share analytical state and record insights:~~

One important aspect of exploratory data analysis is the ability to save/load
and share the state of the analysis during and after the analysis process. This
950 allows collaborators to easily grasp the main steps of analysis that have been
performed [10]. When dealing with a large quantity of data, analysts' working
memory can quickly become overloaded [7]. Saving insights through making
annotation records is useful to offload this working memory directly to the vi-
sualisation.

955 In our system, we provide a flexible solution for this context and it worked
well for our needs. Further work, noted from our interviews could include at-
taching annotations to specific data points or time stamps.

The ability to enable the data to be tagged or coded with information would
also be useful and could be utilised in the custom colour mode for distinguishing
960 different tasks for example.

In addition to sharing the digital state of the system, the system has the abil-
ity to save a static image. For our motivational use case we needed to produce
a report for non-expert users, so simple and good quality static imagery was
needed. The ability for the system to be able to produce simple visualisations
965 of complex eye-tracking data was verified through our interviews.

Finally, we acknowledge that there is a learning curve for the features of
the system, and that some aspects are more intuitive than others, yet we found
that the learning time is balanced by the time gained in the analysis of the
data. In our use case (see Section 5) we found that the video link for verifying
970 operator tasks meant this could be performed in a fraction of the time spent on
similar activities in traditional software, and that many more insights could be
identified through visual encoding and reducing visual clutter.

8. Conclusions

In this paper we introduce the VETA system, which provides a visual analytics workflow for exploratory analysis of in-the-wild eye tracking data. The VETA system comprehensively covers the key aspects of the visual analytics sense-making loop and is highly dynamic, to enable the visual exploration and presentation of complex eye-tracking data. We document insights from our applied use case and interviews with experienced eye-tracking analysts demonstrating the benefits and future potential for the system. Through our evaluation we identified a set of design guidelines for visual eye-tracking analysis and many further use cases and potential collaborations. We look forward to exploring the use of VETA in wider contexts.

Acknowledgements

The observational study that motivated the initial development of VETA was funded by the Australian Energy Market Operator (AEMO). We would like to thank all AEMO participants and all eye tracking experts who participated in the evaluation. We also acknowledge the use of Monash Business Behavioural Laboratory equipment in this project.

Author Contribution

Sarah Goodwin: Conceptualization, Data curation, Formal analysis, Funding acquisition, Investigation, Methodology, Project administration, Supervision, Validation, Visualization, Writing - original draft/review editing. **Arnaud Prouzeau:** Conceptualization, Data curation, Formal analysis, Funding acquisition, Investigation, Methodology, Project administration, Supervision, Validation, Visualization, Writing - original draft/review editing. **Ryan Whitelock-Jones:** Investigation, Software, Validation, Writing - original draft. **Christophe Hurter:** Conceptualization, Investigation, Methodology, Visualization, Writing - original draft/review editing. **Lee Lawrence:** Funding acquisition, Investigation, Writing - original draft/review editing. **Umair Afzal:**

Data curation, Investigation, Validation. **Tim Dwyer:** Conceptualization, Funding acquisition, Investigation, Methodology, Writing - original draft/review editing.

References

- 1005 [1] M. Wedel, R. Pieters, Eye tracking for visual marketing, *Foundations and Trends® in Marketing* 1 (4) (2008) 231–320. doi:10.1561/17000000011.
URL <http://dx.doi.org/10.1561/17000000011>
- [2] F. Dehais, M. Causse, N. Régis, E. Menant, P. Labedan, F. Vachon, S. Tremblay, Missing critical auditory alarms in aeronautics: Evidence for
1010 inattentive deafness?, *Proc. of the Human Factors and Ergonomics Society Annual Meeting* 56 (1) (2012) 1639–1643. arXiv:<https://doi.org/10.1177/1071181312561328>, doi:10.1177/1071181312561328.
URL <https://doi.org/10.1177/1071181312561328>
- [3] S. D. Starke, C. Baber, N. J. Cooke, A. Howes, Workflows and individual differences during visually guided routine tasks in a road traffic
1015 management control room, *Applied Ergonomics* 61 (2017) 79 – 89.
doi:<https://doi.org/10.1016/j.apergo.2017.01.006>.
URL <http://www.sciencedirect.com/science/article/pii/S0003687017300121>
- 1020 [4] C. Tominski, S. Gladisch, U. Kister, R. Dachsel, H. Schumann, Interactive lenses for visualization: An extended survey, *Comput. Graph. Forum* 36 (6) (2017) 173–200. doi:10.1111/cgf.12871.
URL <https://doi.org/10.1111/cgf.12871>
- [5] H. Zhou, P. Xu, Y. Xiaoru, Q. Huamin, Edge bundling in information
1025 visualization, *Tsinghua Sci. Tech.* 18 (2) (2013) 148–156.
- [6] ISO 15007:2020: Road vehicles — Measurement and analysis of driver visual behaviour with respect to transport information and control systems,

Standard, International Organization for Standardization (ISO), Geneva, CH (Aug. 2020).

- 1030 [7] P. Pirolli, S. Card, The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis, in: Proc. of International Conference on Intelligence Analysis, 2005, pp. 2–4.
- [8] R. Amar, J. Eagan, J. Stasko, Low-level components of analytic activity in information visualization, in: IEEE Symposium on Information Visualization. INFOVIS 2005., 2005, pp. 111–117.
1035
- [9] J. Stasko, C. Görg, Z. Liu, Jigsaw: Supporting investigative analysis through interactive visualization, *Information Visualization* 7 (2) (2008) 118–132. arXiv:<https://doi.org/10.1057/palgrave.ivs.9500180>, doi:10.1057/palgrave.ivs.9500180.
1040 URL <https://doi.org/10.1057/palgrave.ivs.9500180>
- [10] C. Dunne, N. Henry Riche, B. Lee, R. Metoyer, G. Robertson, Graphtrail: Analyzing large multivariate, heterogeneous networks while supporting exploration history, in: Proc. SIGCHI Conference on Human Factors in Computing Systems, CHI '12, ACM, New York, NY, USA, 2012, p. 1663–1672.
1045 doi:10.1145/2207676.2208293.
URL <https://doi.org/10.1145/2207676.2208293>
- [11] C. Hurter, A. Telea, O. Ersoy, Moleview: An attribute and structure-based semantic lens for large element-based plots, *IEEE Trans. Vis. & Comp. Graphics* 17 (12) (2011) 2600–2609. doi:10.1109/TVCG.2011.223.
1050 URL <https://doi.org/10.1109/TVCG.2011.223>
- [12] A. T. Duchowski, A breadth-first survey of eye-tracking applications, *Behavior Research Methods, Instruments, & Computers* 34 (4) (2002) 455–470. doi:10.3758/BF03195475.
URL <https://doi.org/10.3758/BF03195475>

- 1055 [13] K. Holmqvist, R. Andersson, Eye-tracking: A comprehensive guide to methods, paradigms and measures, Oxford University Press, NY, 2017.
- [14] A. Voßkühler, V. Nordmeier, L. Kuchinke, A. M. Jacobs, Ogama (open gaze and mouse analyzer): Open-source software designed to analyze eye and mouse movements in slideshow study designs, Behavior Research Methods
1060 40 (4) (2008) 1150–1162. doi:10.3758/BRM.40.4.1150.
URL <https://doi.org/10.3758/BRM.40.4.1150>
- [15] M. Burch, A. Kumar, N. Timmermans, An interactive web-based visual analytics tool for detecting strategic eye movement patterns, in: Proc. 11th ACM Symposium on Eye Tracking Research & Applications, ETRA
1065 '19, ACM, New York, NY, USA, 2019. doi:10.1145/3317960.3321615.
URL <https://doi.org/10.1145/3317960.3321615>
- [16] A. A. Bojko, Informative or misleading? heatmaps deconstructed, in: J. A. Jacko (Ed.), Human-Computer Interaction. New Trends, Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp. 30–39.
- 1070 [17] J. H. Goldberg, X. P. Kotval, Computer interface evaluation using eye movements: methods and constructs, International Journal of Industrial Ergonomics 24 (6) (1999) 631 – 645. doi:[https://doi.org/10.1016/S0169-8141\(98\)00068-7](https://doi.org/10.1016/S0169-8141(98)00068-7).
URL <http://www.sciencedirect.com/science/article/pii/S0169814198000687>
1075 S0169814198000687
- [18] M. Burch, A. Veneri, B. Sun, Eyeclouds: A visualization and analysis tool for exploring eye movement data, in: Proc. 12th International Symposium on Visual Information Communication and Interaction, VINCI'2019, ACM, New York, NY, USA, 2019. doi:10.1145/3356422.3356423.
1080 URL <https://doi.org/10.1145/3356422.3356423>
- [19] K. Holmqvist, J. Holsanova, M. Barthelson, D. Lundqvist, Reading or scanning? A study of newspaper and net paper reading,

- in: J. Hyönä, R. Radach, H. Deubel (Eds.), *The Mind's Eye*, North-Holland, Amsterdam, 2003, pp. 657 – 670. doi:<https://doi.org/10.1016/B978-044451020-4/50035-9>.
1085 URL <http://www.sciencedirect.com/science/article/pii/B9780444510204500359>
- [20] A. Fabian, I. Schilstra, P. Lukas, S. Bryczkowski, M. Park, M. Burch, Interactive data visualization tool for eye movement data.
- 1090 [21] K. Kurzhals, N. Rodrigues, M. Koch, M. Stoll, A. Bruhn, A. Bulling, D. Weiskopf, Visual analytics and annotation of pervasive eye tracking video, in: *Proc. ACM International Symposium on Eye Tracking Research and Applications (ETRA)*, 2020.
- [22] G. Andrienko, N. Andrienko, M. Burch, D. Weiskopf, Visual analytics methodology for eye movement studies, *IEEE Trans. Vis. & Comp. Graphics* 18 (12) (2012) 2889–2898.
1095
- [23] T. Blascheck, M. John, K. Kurzhals, S. Koch, T. Ertl, Va2: A visual analytics approach for evaluating visual analytics applications, *IEEE Trans. Vis. & Comp. Graphics* 22 (1) (2016) 61–70.
- 1100 [24] K. Kurzhals, D. Weiskopf, Space-time visual analytics of eye-tracking data for dynamic stimuli, *IEEE Trans. Vis. & Comp. Graphics* 19 (12) (2013) 2129–2138.
- [25] P. K. Muthumanickam, K. Vrotsou, A. Nordman, J. Johansson, M. Cooper, Identification of temporally varying areas of interest in long-duration eye-tracking data sets, *IEEE Trans. Vis. & Comp. Graphics* 25 (1) (2019) 87–97.
1105
- [26] T. Blascheck, K. Kurzhals, M. Raschke, M. Burch, D. Weiskopf, T. Ertl, Visualization of eye tracking data: A taxonomy and survey, *Computer Graphics Forum* 36 (8) (2017) 260–284. arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13079>, doi:10.1111/cgf.13079.
1110 URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13079>

- [27] S. Stellmach, L. Nacke, R. Dachsel, Advanced gaze visualizations for three-dimensional virtual environments, in: Proceedings of the 2010 Symposium on Eye-Tracking Research and Applications, ETRA '10, Association for Computing Machinery, New York, NY, USA, 2010, p. 109–112. doi:10.1145/1743666.1743693.
1115 URL <https://doi.org/10.1145/1743666.1743693>
- [28] N. Weibel, A. Fouse, C. Emmenegger, S. Kimmich, E. Hutchins, Let's look at the cockpit: Exploring mobile eye-tracking for observational research on the flight deck, in: Proceedings of the Symposium on Eye Tracking Research and Applications, ETRA '12, Association for Computing Machinery,
1120 New York, NY, USA, 2012, p. 107–114. doi:10.1145/2168556.2168573.
URL <https://doi.org/10.1145/2168556.2168573>
- [29] T. Pfeiffer, Measuring and visualizing attention in space with 3d attention volumes, in: Proceedings of the Symposium on Eye Tracking Research and Applications, ETRA '12, Association for Computing Machinery, New York,
1125 NY, USA, 2012, p. 29–36. doi:10.1145/2168556.2168560.
URL <https://doi.org/10.1145/2168556.2168560>
- [30] A. Cockburn, A. Karlson, B. B. Bederson, A review of overview+detail, zooming, and focus+context interfaces, ACM Comput. Surv. 41 (1). doi:
1130 10.1145/1456650.1456652.
URL <https://doi.org/10.1145/1456650.1456652>
- [31] C. Tominski, J. Abello, F. van Ham, H. Schumann, Fisheye tree views and lenses for graph visualization, in: Tenth International Conference on Information Visualisation (IV'06), 2006, pp. 17–24. doi:10.1109/IV.2006.54.
- [32] N. Wong, S. Carpendale, S. Greenberg, Edgelens: an interactive method for
1135 managing edge congestion in graphs, in: IEEE Symposium on Information Visualization 2003 (IEEE Cat. No.03TH8714), 2003, pp. 51–58. doi:10.1109/INFVIS.2003.1249008.

- [33] S. Schmidt, M. A. Nacenta, R. Dachsel, S. Carpendale, A set of multi-
1140 touch graph interaction techniques, in: ACM International Conference
on Interactive Tabletops and Surfaces, ITS '10, Association for Computing
Machinery, New York, NY, USA, 2010, p. 113–116. doi:10.1145/
1936652.1936673.
URL <https://doi.org/10.1145/1936652.1936673>
- [34] A. Panagiotidis, H. Bosch, S. Koch, T. Ertl, Edgeanalyzer: Exploratory
1145 analysis through advanced edge interaction, in: 2011 44th Hawaii Inter-
national Conference on System Sciences, 2011, pp. 1–10. doi:10.1109/
HICSS.2011.178.
- [35] U. Kister, P. Reipschläger, R. Dachsel, Multilens: Fluent interaction with
1150 multi-functional multi-touch lenses for information visualization, in: Pro-
ceedings of the 2016 ACM International Conference on Interactive Surfaces
and Spaces, ISS '16, Association for Computing Machinery, New York, NY,
USA, 2016, p. 139–148. doi:10.1145/2992154.2992168.
URL <https://doi.org/10.1145/2992154.2992168>
- [36] A. Fuhrmann, E. Groller, Real-time techniques for 3d flow visualization,
1155 in: Proceedings Visualization '98 (Cat. No.98CB36276), 1998, pp. 305–312.
doi:10.1109/VISUAL.1998.745317.
- [37] O. Mattausch, T. Theußl, H. Hauser, E. Gröller, Strategies for interac-
tive exploration of 3d flow using evenly-spaced illuminated streamlines, in:
1160 Proceedings of the 19th Spring Conference on Computer Graphics, SCCG
'03, Association for Computing Machinery, New York, NY, USA, 2003, p.
213–222. doi:10.1145/984952.984987.
URL <https://doi.org/10.1145/984952.984987>
- [38] R. Gasteiger, M. Neugebauer, O. Beuing, B. Preim, The flowlens: A focus-
1165 and-context visualization approach for exploration of blood flow in cerebral
aneurysms, IEEE Transactions on Visualization and Computer Graphics
17 (12) (2011) 2183–2192. doi:10.1109/TVCG.2011.243.

- [39] R. Krüger, D. Thom, M. Wörner, H. Bosch, T. Ertl, Trajectorylenses – a set-based filtering and exploration technique for long-term trajectory data, *Comp. Graphics Forum* 32 (2013) 451–460. `arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12132`, `doi:https://doi.org/10.1111/cgf.12132`.
URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12132`
- [40] A. Ferreira, A. P. Afonso, L. Ferreira, R. Vaz, Visual analytics of trajectories with rosetrajvis, *Big Data Research* 27 (2022) 100294. `doi:https://doi.org/10.1016/j.bdr.2021.100294`.
URL `https://www.sciencedirect.com/science/article/pii/S2214579621001118`
- [41] F. Chevalier, N. H. Riche, C. Plaisant, A. Chalbi, C. Hurter, Animations 25 years later: New roles and opportunities, in: *Proc. International Working Conference on Advanced Visual Interfaces, AVI '16*, ACM, New York, NY, USA, 2016, p. 280–287. `doi:10.1145/2909132.2909255`.
URL `https://doi.org/10.1145/2909132.2909255`
- [42] A. Lhuillier, C. Hurter, A. Telea, FFTEB: Edge bundling of huge graphs by the Fast Fourier Transform, in: *Proc. IEEE PacificVis*, 2017.
- [43] N. Andrienko, G. Andrienko, P. Gatalsky, Exploratory spatio-temporal visualization: an analytical review, *Journal of Visual Languages & Computing* 14 (6) (2003) 503–541, *visual Data Mining*. `doi:https://doi.org/10.1016/S1045-926X(03)00046-6`.
URL `https://www.sciencedirect.com/science/article/pii/S1045926X03000466`
- [44] R. Scheepens, C. Hurter, H. Van De Wetering, J. J. Van Wijk, Visualization, selection, and analysis of traffic flows, *IEEE Trans. Vis. & Comp. Graphics* 22 (1) (2016) 379–388.
- [45] G. Robertson, R. Fernandez, D. Fisher, B. Lee, J. Stasko, Effectiveness

- of animation in trend visualization, *IEEE Trans. Vis. & Comp. Graphics* 14 (6) (2008) 1325–1332.
- [46] M. Shanmugasundaram, P. Irani, The effect of animated transitions in zooming interfaces, in: *Proc. working conference on Advanced visual interfaces*, 2008, pp. 396–399.
- 1200 [47] N. Elmqvist, P. Dragicevic, J.-D. Fekete, Rolling the dice: Multidimensional visual exploration using scatterplot matrix navigation, *IEEE Trans. Vis. & Comp. Graphics* 14 (6) (2008) 1539–1148.
- [48] C. Hurter, B. Tissoires, S. Conversy, Fromdady: Spreading aircraft trajectories across views to support iterative queries, *IEEE Trans. Vis. & Comp. Graphics* 15 (6) (2009) 1017–1024. doi:10.1109/TVCG.2009.145.
URL <https://doi.org/10.1109/TVCG.2009.145>
- 1205 [49] C. Hurter, R. Taylor, S. Carpendale, A. Telea, Color tunneling: Interactive exploration and selection in volumetric datasets, in: *Proc. 2014 IEEE Pacific Visualization Symposium, PACIFICVIS '14*, IEEE Computer Society, USA, 2014, p. 225–232. doi:10.1109/PacificVis.2014.61.
URL <https://doi.org/10.1109/PacificVis.2014.61>
- [50] J. Heer, G. Robertson, Animated transitions in statistical data graphics, *IEEE Trans. Vis. & Comp. Graphics* 13 (6) (2007) 1240–1247.
- 1215 [51] B. Tversky, J. B. Morrison, M. Betrancourt, Animation: can it facilitate?, *International journal of human-computer studies* 57 (4) (2002) 247–262.
- [52] C. Tominski, H. Schumann, G. Andrienko, N. Andrienko, Stacking-based visualization of trajectory attribute data, *IEEE Trans. Vis. & Comp. Graphics* 18 (12) (2012) 2565–2574. doi:10.1109/TVCG.2012.265.
- 1220 [53] R. Krueger, S. Koch, T. Ertl, Saccadelenses: interactive exploratory filtering of eye tracking trajectories, in: *IEEE Second Workshop on Eye Tracking and Visualization*, 2016, pp. 31–34.

- [54] H. Zhou, P. Xu, X. Yuan, H. Qu, Edge bundling in information visualization, *Tsinghua Science and Technology* 18 (2) (2013) 145–156. doi:10.1109/TST.2013.6509098.
- 1225
- [55] D. Holten, J. J. van Wijk, A user study on visualizing directed edges in graphs, in: *Proc. ACM CHI*, 2009, pp. 2299–2308.
- [56] C. Hurter, O. Ersoy, A. Telea, Graph bundling by kernel density estimation, *CGF* 31 (3) (2012) 435–443.
- 1230
- [57] M. van der Zwan, V. Codreanu, A. Telea, CUBu: Universal real-time bundling for large graphs, *IEEE Trans. Vis. & Comp. Graphics* 22 (12) (2016) 2550–2563.
- [58] C. Hurter, S. Puechmorel, F. Nicol, A. Telea, Functional decomposition for bundled simplification of trail sets, *IEEE Trans. Vis. & Comp. Graphics* 24 (1) (2018) 500–510.
- 1235
- [59] Y. Wang, M. Xue, Y. Wang, X. Yan, B. Chen, C. Fu, C. Hurter, Interactive structure-aware blending of diverse edge bundling visualizations, *IEEE Trans. Vis. & Comp. Graphics* 26 (1) (2020) 687–696.
- [60] T. Andry, C. Hurter, F. Lambotte, P. Fastrez, A. Telea, Interpreting the Effect of Embellishment on Chart Visualizations, *Association for Computing Machinery*, New York, NY, USA, 2021.
- 1240
- URL <https://doi.org/10.1145/3411764.3445739>
- [61] C. Hurter, O. Ersoy, A. Telea, MoleView: An attribute and structure-based semantic lens for large element-based plots, *IEEE Trans. Vis. & Comp. Graphics* 17 (12) (2011) 2600–2609.
- 1245
- [62] A. Bykowski, S. Kupiundefinedski, Automatic mapping of gaze position coordinates of eye-tracking glasses video on a common static reference image, in: *Proc. 2018 ACM Symposium on Eye Tracking Research and Applications*, ETRA '18, ACM, NY, USA, 2018. doi:10.1145/3204493.3208331.
- 1250
- URL <https://doi.org/10.1145/3204493.3208331>

- [63] K. Kilingaru, J. W. Tweedale, S. Thatcher, L. C. Jain, Monitoring pilot “situation awareness”, *Journal of Intelligent & Fuzzy Systems* 24 (3) (2013) 457–466.
- [64] M. Sears, O. Alruwaythi, P. Goodrum, Evidence of inconsistent results using current eye tracking glance and visit analysis standards, *Automation in Construction* 132 (2021) 103951. doi:<https://doi.org/10.1016/j.autcon.2021.103951>.
URL <https://www.sciencedirect.com/science/article/pii/S0926580521004027>
- [65] C. Sharma, P. Bhavsar, B. Srinivasan, R. Srinivasan, Eye gaze movement studies of control room operators: A novel approach to improve process safety, *Computers Chemical Engineering* 85 (2016) 43–57. doi:<https://doi.org/10.1016/j.compchemeng.2015.09.012>.
URL <https://www.sciencedirect.com/science/article/pii/S0098135415003075>
- [66] K.-J. Rähkä, A. Aula, P. Majaranta, H. Rantala, K. Koivunen, Static visualization of temporal eye-tracking data, in: M. F. Costabile, F. Paternò (Eds.), *Human-Computer Interaction - INTERACT 2005*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2005, pp. 946–949.
- [67] V. Peysakhovich, C. Hurter, Scan path visualization and comparison using visual aggregation techniques, *Journal of Eye Movement Research* 10 (5). doi:[10.16910/jemr.10.5.9](https://doi.org/10.16910/jemr.10.5.9).
URL <https://bop.unibe.ch/JEMR/article/view/3730>
- [68] B. Bach, P. Dragicevic, D. Archambault, C. Hurter, S. Carpendale, A descriptive framework for temporal data visualizations based on generalized space-time cubes, *Computer Graphics Forum* 36 (6) (2017) 36–61. arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12804>, doi:[10.1111/cgf.12804](https://doi.org/10.1111/cgf.12804).
URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12804>

- 1280 [69] T. Blascheck, K. Kurzhals, M. Raschke, M. Burch, D. Weiskopf, T. Ertl,
State-of-the-art of visualization for eye tracking data, Proceedings of EuroVis 2014.
- [70] C. Hurter, A. Telea, O. Ersoy, MoleView: An Attribute and Structure-Based Semantic Lens for Large Element-Based Plots, IEEE Trans. Vis. &
1285 Comp. Graphics 17 (12) (2011) 2600–2609. doi:10.1109/TVCG.2011.223.
- [71] A. W. Fitzgibbon, Stochastic rigidity: image registration for nowhere-static scenes, in: Proc. Eighth IEEE International Conference on Computer Vision. ICCV, Vol. 1, 2001, pp. 662–669.
- [72] W. C. Wilson, Activity pattern analysis by means of sequence-alignment
1290 methods, Environment and Planning A: Economy and Space 30 (6) (1998) 1017–1038. arXiv:<https://doi.org/10.1068/a301017>, doi:10.1068/a301017.
URL <https://doi.org/10.1068/a301017>
- [73] C. Westin, K. Vrotsou, A. Nordman, J. Lundberg, L. Meyer, Visual scan
1295 patterns in tower control: Foundations for an instructor support tool, in: SESAR Innovation Days, 2019.
- [74] K. Vrotsou, A. Nordman, A Window-based Approach for Mining Long Duration Event-sequences, in: C. Turkay, K. Vrotsou (Eds.), EuroVis Workshop on Visual Analytics (EuroVA), The Eurographics Association, 2020.
1300 doi:10.2312/eurova.20201087.