Link to data: https://atreus.informatik.uni-tuebingen.de/seafile/d/8e2ab8c3fdd444e1a135/?p= %2FRandomFernScanpathClassification&mode=list

Ferns for area of interest free scanpath classification

Wolfgang Fuhl University Tübingen, Perception Engineering Germany wolfgang.fuhl@uni-tuebingen.de

Alexander Lotz Daimler AG, Truck Product Engineering Stuttgart rene_alexander.lotz@daimler.com

Nora Castner University Tübingen, Perception Engineering Germany castnern@informatik.uni-tuebingen. de

Wolfgang Rosenstiel University Tübingen, Technical **Computer Science** Germany Wolfgang.Rosenstiel@uni-tuebingen.

Thomas Kübler University Tübingen, Perception Engineering Germany thomas.kuebler@uni-tuebingen.de

Enkelejda Kasneci University Tübingen, Perception Engineering Germany enkelejda.kasneci@uni-tuebingen.de

de

ABSTRACT

Scanpath classification can offer insight into the visual strategies of groups such as experts and novices. We propose to use random ferns in combination with saccade angle successions to compare scanpaths. One advantage of our method is that it does not require areas of interest to be computed or annotated. The conditional distribution in random ferns additionally allows for learning angle successions, which do not have to be entirely present in a scanpath. We evaluated our approach on two publicly available datasets and improved the classification accuracy by ≈ 10 and ≈ 20 percent.

CCS CONCEPTS

 Computer systems organization → Embedded systems; Re*dundancy*; Robotics; • **Networks** \rightarrow Network reliability;

KEYWORDS

Eye tracking, scanpath analysis, random ferns

ACM Reference Format:

Wolfgang Fuhl, Nora Castner, Thomas Kübler, Alexander Lotz, Wolfgang Rosenstiel, and Enkelejda Kasneci. 2019. Ferns for area of interest free scanpath classification. In 2019 Symposium on Eye Tracking Research and Applications (ETRA '19), June 25-28, 2019, Denver, CO, USA. ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3314111.3319826

1 **INTRODUCTION**

The way we direct our eyes at can tell us much more than what we are looking at. Naturally, gaze behavior reflects the interplay of cognitive as well as sensory processes. Consequently, patterns of fixations and saccades, known as the scanpath, can offer insight into the order and nature of information processing. Over multiple

ETRA '19, June 25-28, 2019, Denver, CO, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6709-7/19/06...\$15.00

https://doi.org/10.1145/3314111.3319826

domains, eye tracking studies have identified patterns in gaze behavior. The literature found that these patterns could characterize specific subject groups (e.g., novices and experts), or experimental settings, such as the task assigned to a subject.

For instance in art, eye scanning differences have been found between professional and naive art viewers for both realistic and abstract art [Zangemeister et al. 1995]. Also, top-down beliefs as well as bottom-up feature attention can affect the gaze behavior on artworks [Locher et al. 2015; Massaro et al. 2012].

Additionally in the medical domain, scanpath differences can reflect the professional as well as the treatment aspects. Scanpath differences between novices and experts have been found in microneurosurgeons [Eivazi et al. 2012; Kübler et al. 2015a] as well as radiologists [Manning et al. 2006; Van der Gijp et al. 2017]. It was also found that dental students who received a specialized radiography training course could be accurately identified from their scanpaths [Castner et al. 2018]. On the treatment aspect, gaze behavior differences from healthy controls have been measured in both patients suffering from schizophrenia [Loughland et al. 2002] and autism spectrum disorder [Horley et al. 2004; Pelphrey et al. 2002]. Hence, scanpaths can likely be employed towards more refined training, diagnostic, and treatment protocols.

In driving, scanpaths have been used to robustly determine safe or unsafe driving in people with visual field defects [Kasneci et al. 2014; Kübler et al. 2015b]. In addition, they can be used in driver assistance systems to indicate the take-over readiness [Braunagel et al. 2017] or cognitive load [Palinko et al. 2010], and fatigue [Ji et al. 2004].

Interestingly, most of the studies mentioned above concentrate on finding statistically significant differences in individual scanpath metrics. Therefore, there is a large and ever growing body of scanpath comparison and classification methodology: Ranging from simple statistics to state-of-the-art machine learning (see [Anderson et al. 2015] for a review).

RELATED WORKS 2

In 1935, Buswell hypothesized that "the mental set obtained by the directions given [...] obviously influences the characteristics of the perceptual process" [1935]. A finding followed up in late

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.



(a) Stimulus

(b) Example Saccades of one participant

(c) Similar patterns

Figure 1: (a) the stimulus image, Tintoretto's "Last Supper", and (b) saccades of one subject in blue. (c) similar saccade patterns (in green) found in another participant are shown overlaying the subject in (b)'s scanpath.

sixties/early seventies by Yarbus [1967] and Noton & Stark [1971] when they measured gaze pattern differences as an effect of contextual information. Here, the term *scanpath* was initially used, and the method for comparing multiple scanpaths was executed manually: Using semantic understanding of the image. Even today, some aspects of manual scanpath comparison are still used; such as AOIs hand-labeled by the experimenter, which can be tedious, task dependent, and subjective [Jarodzka et al. 2010].

Only in the nineties were automated metrics initially proposed [Brandt and Stark 1997]. Since then, a wide array of methodology for automated scanpath comparison has evolved [Anderson et al. 2015]. More recently, machine-learning based approaches with impressive results have emerged [Crabb et al. 2014; French et al. 2017; Hoppe et al. 2018; Kübler et al. 2017; Zhang and Le Meur 2018]. In general, they can handle the task of distinguishing relevant eve movement patterns from an overall high level of noise. However, how the eye movement trajectories should be encoded to enable efficient machine learning is still an area of open debate. While some algorithms rely heavily on a massive agglomeration of time-aggregated features or complete-sequence alignment [Burch et al. 2018; Cristino et al. 2010; Dewhurst et al. 2018; Hoppe et al. 2018], the use of gaze transitions (i.e., the shift of gaze between two targets) has become one of the most popular features. This strategy is notable, as a chain of cognitive associations between gaze targets can be modeled this way. Also, while Hidden Markov Models (HMMs) are still the most common approach [Coutrot et al. 2018; Ellis and Stark 1986; Hacisalihzade et al. 1992], other methods have emerged that extend these patterns in length to span multiple subsequent fixations and saccades. Finding patterns in such lengthier sequences is of special importance, as the specificity of these patterns for a specific task or subject group is likely increased and, therefore, they are highly useful for the classification task [Kübler et al. 2017].

3 METHOD

We analyzed the patterns that represent the angles between successive saccades. An example is shown in Figure 1, where similar angle patterns are a subsequence, as apparent over multiple participants. Thus, our algorithm searches for repetitive patterns over subjects. These repetitions are then used to classify the scanpath into a category. Since such patterns are subject to large variations between individuals and even between repeated trials-not to mention eye tracker inaccuracies-comparing angular patterns requires a slack range in which two angles are considered similar. Therefore, each angle in a pattern has an assigned and variable angular tolerance range. By employing ferns [Bosch et al. 2007; Ozuysal et al. 2010], we can allow assigning probability values to each permutation of such a pattern, which makes it possible to correctly classify incomplete patterns. In the following section, each step of our approach is described in detail.

3.1 Saccade sequence angles



Figure 2: Different options for calculating angles between subsequent saccades. On the left, the absolute, counterclockwise angle (red part of a circle) to the positive x-axis (green dashed line) is shown. The right side shows the relative, counter-clockwise angle between subsequent saccades.

Figure 2 shows two possible approaches to compute the angles between saccades from a sequence of saccades. The major difference between these approaches is whether the angles are calculated relative to an absolute reference frame or relative to the preceding saccade. Both can likely be valid approaches. The latter adds invariance to rotations of the pattern, which could be attractive for data from head-mounted eye-trackers. Whereas the former is able to distinguish between rotated representations of the pattern.

In our evaluation, both approaches performed almost equally for both data sets evaluated. Since the angle between saccades is invariant to rotation of the eye tracker–a challenge especially with head-mounted eye trackers–we believe it is the more robust approach. However, our implementation provides both computation models. In the following, these angles will be denoted as $\alpha_{i,fi}$, with the indices *i* and *f i* denoting the index and the assigned fern index (i.e., the position index in the fern where it is used), respectively.

Figure 3 shows the angular tolerance ranges in which saccades would be considered similar to the current pattern. In our implementation, those tolerances can be in the range of 0° to 359°. Thus, a

Ferns for area of interest free scanpath classification

ETRA '19, June 25-28, 2019, Denver, CO, USA



Figure 3: Angular ranges (teal part of a circle) for saccade successions.

feature that characterizes a pattern would consist of an angle $\alpha_{i,fi}$ and a similarity range $\rho_{i,fi}$ denoted as $(\alpha, \rho)_{i,fi}$. Then, a feature vector F_{fi} is the concatenation of one or more of these tuples.

For a saccade sequence to match a specific pattern/feature vector, all consecutive angles between the saccades must meet the angular range criteria defined by the feature tuples.

3.2 Random ferns

The definition of a fern [Bosch et al. 2007; Ozuysal et al. 2010] is a group of conditionally dependent features. Each feature evaluates either to 1 or zero. The combination of multiple ferns is done under the assumption that each group of features is conditionally independent. Figure 4 shows three features (F_{fi}) wherein each feature



Figure 4: A set of features in which the feature length is two and the second index specifies the position in the ferm (fi). The first index (i) specifies the sequence of conditions $((\alpha, \rho)_{i, fi})$.

consists of two angle - range tuples. The first feature is assigned to the index one (the left side of Figure 4), the central feature to index two, and the last feature (the right side of Figure 4) to index three. In a fern [Bosch et al. 2007; Ozuysal et al. 2010], this set of features is used to compute an index to a conditional distribution p(T|F), with F denoting the evaluated features and T is the target label for classification. Therefore, each class receives a probability based on the occurrence of features within the class. Figure 5 shows the evaluation process of a fern on three different input sequences. These features are evaluated by the fern and matched to the input scanpaths. In the case that a feature is successfully matched, it is encoded by a 1, otherwise by a 0 at the respective fern index. This step builds a binary-encoded number, which is used as an index to a conditional distribution (indicated by the black bars in figure 5). This probability relates directly to the probability that the seen scanpath can be assigned to a specific output label. To distinguish between several class labels, multiple probability distributions need to be learned. An example of multi-class classification is shown in Figure 6, with three classes distinguished based on their assigned relative probabilities. The upper scanpath (in green) has the highest probability for class 2, the central scanpath (in red) is assigned to class 1, and the bottom scanpath (in blue) receives the highest probability for class 3. Learning these probability distributions is a similar process to evaluating them. Where the features on the



Figure 5: Three different scanpaths are evaluated based on three features F_{fi} . The result of this evaluation yields the index in the probability distribution.



Figure 6: Class label selection based on multiple conditional probability distributions.

training data set are evaluated, and the corresponding entry in the histogram (later transformed to the probability distribution) is increased. After learning, the histogram bins are normalized to form a probability distribution. A uniform distribution must be assumed for initialization, since several of the ferns in combination serve as classifiers and the probabilities of the individual ferns are multiplied. If this step is not performed, zero probabilities can occur, which can influence the classification negatively.

3.3 Training

The training of ferns consists of three parts, namely feature selection, distribution training, and fern combination selection. Due to the massive amount of possible features, brute-force for all possible patterns is only applicable for short feature length. Therefore, we only compute features up to a length of three and only inspect angle successions that occur in the training data scanpaths.

First, we extract all possible angles of length one to three. Since a feature consists of tuples $((\alpha, \rho)_{i,fi})$, we have to assign each angle $(\alpha_{i,fi})$ a range $(\rho_{i,fi})$. This assignment is done by inspecting all possible ranges per angle and evaluating those on the training dataset with a step width of one degree: Running in polynomial time (360^{length}) . We define this evaluation as the feature score.

$$FS_j = |\sum (\alpha, \rho)(T_j) == 0 - \sum (\alpha, \rho)(T_j) == 1|$$
(1)

For each feature and each class (j), this score can be computed using Equation 1. Features for one specific class that result in a score that deviates by 10% from the average score of all classes were added to the feature set of that class. Afterward, the features are sorted based on this score to select the best features directly.

The second part of our training is the selection of a set of ferns. The conditional distribution size per fern is equal to the number of features, and the amount of distributions is equal to the class labels in the training set. Therefore, if seven classes have to be distinguished, each distribution has a size of 2^F , with F denoting the amount of used features, and each fern uses seven of these distributions. The 2 is due to each feature evaluating as either one or zero. Afterward, we select one feature per class from the feature pool and add it to the feature set of the fern. The conditional distribution is filled with a uniform distribution and trained on the training dataset based on the occurrence of the different combinations statistically. Since this step does not necessarily result in an optimal feature combination, we evaluate all possible and select the best. This process can also be done in polynomial time (featurepoolsize)^{Classlabels}, since the class labels are fixed. In our implementation, we compute this step using threads to reduce the training time.

The third and last step of the training is the selection of a combination of multiple ferns. The final classification is performed by multiplication of all fern probabilities and selection of the highest ranking probability as the final label. Since we have already evaluated all possible combinations of features per fern, which results in a fern pool to directly select from. Here, we inspect the probability per class, which our existing detector provides in comparison to the possible new fern in the detector. For the calculation, the existing detector consisting of several ferns is evaluated on the training data as well as the possible new fern. Afterward, the average difference between both results is computed and the fern with the highest difference is selected as the new member of the detector. This process starts with the selection of the best fern and ends after ten ferns are selected. The amount of ten ferns was selected empirically and can be changed in our implementation.

4 EVALUATION

For the evaluation and comparison to the state-of-the-art, we used two publicly available data sets. The first data set is the task classification in the second experiment from the Defending Yarbus (DY) paper [Borji and Itti 2014]. It contains seven tasks and 45 recordings for each of them: Data is provided as fixation coordinates. The second data set is the task classification from Reconsidering Yarbus (RY) paper [Greene et al. 2012], in which four tasks were distinguished and \approx 80 recordings are provided per task: Data is provided as gaze points. Since our algorithm needs fixations as input, we computed the mean velocity for the gaze points per trial and used this value as a threshold to separate fixations and saccades. This method is not an accurate event detection, but was performed to make the results easily reproducible without relying on algorithm and parameter choices for event detection. As our focus is identifying each saccade and not about e.g. accurate fixation duration, such a coarse-but simple-event detection is sufficient for this application. For each data set, we performed ten fold cross-validation;

Table 1: The classification accuracy in % of our approach in comparison to the state-of-the-art for task classification. ANG-X and ANG-V specify the angle computation models to the x-axis or between two saccades respectively.

| | Datasets | |
|-----------------------|--------------|--------------|
| Method | DY (7 Tasks) | RY (4 Tasks) |
| Chance level | 14.3 | 25.0 |
| Fern _{ANG-X} | 43.57 | 54.06 |
| Fern _{ANG-V} | 41.42 | 54.37 |
| [Kübler et al. 2017] | 24.2 | 34.4 |
| [Borji and Itti 2014] | 28.71 | 34.1 |
| [Greene et al. 2012] | - | 27.1 |
| [Kanan et al. 2014] | - | 33.0 |

each fold contained an equal amount of recordings for each task. Therefore, each fold in DY dataset contained four recordings of each task and seven recordings for RY dataset. The non-integer divisible share was added to the training data. Table 1 shows our results in comparison to the state-of-the-art. As evident, our approach outperforms the competitor algorithms. Moreover, our algorithm does not require areas of interest (AOIs). This limitation to scanpath algorithms was already treated in [Kübler et al. 2017], in which the AOIs are computed based on the data. Since both angle calculation functions deliver approximately the same results, it is not possible to make a statement regarding the robustness based on this evaluation (Table 1); further research has to be conducted to evaluate this aspect. One disadvantage of our algorithm is the long training time due to the polynomial complexity of selecting suited patterns. However, it can be solved by selecting only a subset of the training data. Another issue is that long recordings will require a windowing function.

5 CONCLUSION

We proposed a novel algorithm for scanpath classification. It uses angle and angle-range tuples as features and assigns probabilities based on conditional distributions. Our algorithm outperformed the state-of-the-art on two publicly available datasets. However, the disadvantages of our algorithm are the polynomial complexity of the training and therefore, the high computational demand.

Since our approach is independent of AOIs and the features are based on angles that are calculated between saccades, the next planned step is to create a data set using head-mounted eye trackers. This data would further elaborate and verify that angle computation method between saccades and whether rotation invariance is of importance for data including head movements. In addition, we intend to evaluate optimal time segment sizes based on eye tracker frame rates as well as for different tasks,e.g., expert and novices, driver state detection, etc.

ACKNOWLEDGMENTS

Work of the authors is supported by the Institutional Strategy of the University of Tübingen (Deutsche Forschungsgemeinschaft, ZUK 63) Ferns for area of interest free scanpath classification

REFERENCES

- Nicola C Anderson, Fraser Anderson, Alan Kingstone, and Walter F Bischof. 2015. A comparison of scanpath comparison methods. *Behavior research methods* 47, 4 (2015), 1377–1392.
- Ali Borji and Laurent Itti. 2014. Defending Yarbus: Eye movements reveal observers' task. Journal of vision 14, 3 (2014), 29–29.
- Anna Bosch, Andrew Zisserman, and Xavier Munoz. 2007. Image classification using random forests and ferns. In 2007 IEEE 11th international conference on computer vision. Ieee, 1–8.
- Stephan A Brandt and Lawrence W Stark. 1997. Spontaneous eye movements during visual imagery reflect the content of the visual scene. Journal of cognitive neuroscience 9, 1 (1997), 27–38.
- Christian Braunagel, David Geisler, Wolfgang Rosenstiel, and Enkelejda Kasneci. 2017. Online Recognition of Driver-Activity Based on Visual Scanpath Classification. *IEEE Intelligent Transportation Systems Magazine* 9, 4 (2017), 23–36.
- Michael Burch, Kuno Kurzhals, Niklas Kleinhans, and Daniel Weiskopf. 2018. EyeMSA: exploring eye movement data with pairwise and multiple sequence alignment. In Proceedings of the 2018 ACM Symposium on Eye Tracking Research & Applications. ACM, 52.
- Guy Thomas Buswell. 1935. How people look at pictures: a study of the psychology and perception in art. (1935).
- Nora Castner, Enkelejda Kasneci, Thomas Kübler, Katharina Scheiter, Juliane Richter, Thérése Eder, Fabian Hüttig, and Constanze Keutel. 2018. Scanpath comparison in medical image reading skills of dental students: distinguishing stages of expertise development. In Proceedings of the 2018 ACM Symposium on Eye Tracking Research & Applications. ACM, 39.
- Antoine Coutrot, Janet H Hsiao, and Antoni B Chan. 2018. Scanpath modeling and classification with hidden Markov models. *Behavior research methods* 50, 1 (2018), 362–379.
- David P Crabb, Nicholas D Smith, and Haogang Zhu. 2014. What's on TV? Detecting age-related neurodegenerative eye disease using eye movement scanpaths. *Frontiers in Aging Neuroscience* 6 (2014), 312.
- Filipe Cristino, Sebastiaan Mathôt, Jan Theeuwes, and Iain D Gilchrist. 2010. ScanMatch: A novel method for comparing fixation sequences. *Behavior research methods* 42, 3 (2010), 692–700.
- Richard Dewhurst, Tom Foulsham, Halszka Jarodzka, Roger Johansson, Kenneth Holmqvist, and Marcus Nyström. 2018. How task demands influence scanpath similarity in a sequential number-search task. Vision research 149 (2018), 9–23.
- Shahram Eivazi, Roman Bednarik, Markku Tukiainen, Mikael von und zu Fraunberg, Ville Leinonen, and Juha E Jääskeläinen. 2012. Gaze behaviour of expert and novice microneurosurgeons differs during observations of tumor removal recordings. In Proceedings of the Symposium on Eye Tracking Research and Applications. ACM, 377–380.
- Stephen R Ellis and Lawrence Stark. 1986. Statistical dependency in visual scanning. Human factors 28, 4 (1986), 421–438.
- Robert M French, Yannick Glady, and Jean-Pierre Thibaut. 2017. An evaluation of scanpath-comparison and machine-learning classification algorithms used to study the dynamics of analogy making. *Behavior research methods* 49, 4 (2017), 1291–1302.
- Michelle R Greene, Tommy Liu, and Jeremy M Wolfe. 2012. Reconsidering Yarbus: A failure to predict observers' task from eye movement patterns. Vision research 62 (2012), 1–8.
- Selim S Hacisalihzade, Lawrence W Stark, and John S Allen. 1992. Visual perception and sequences of eye movement fixations: A stochastic modeling approach. IEEE Transactions on systems, man, and cybernetics 22, 3 (1992), 474–481.
- Sabrina Hoppe, Tobias Loetscher, Stephanie A Morey, and Andreas Bulling. 2018. Eye movements during everyday behavior predict personality traits. *Frontiers in human neuroscience* 12 (2018), 105.
- Kaye Horley, Leanne M Williams, Craig Gonsalvez, and Evian Gordon. 2004. Face to face: visual scanpath evidence for abnormal processing of facial expressions in social phobia. *Psychiatry research* 127, 1-2 (2004), 43–53.
- Halszka Jarodzka, Kenneth Holmqvist, and Marcus Nyström. 2010. A vector-based, multidimensional scanpath similarity measure. In Proceedings of the 2010 symposium on eye-tracking research & applications. ACM, 211–218.
- Qiang Ji, Zhiwei Zhu, and Peilin Lan. 2004. Real-time nonintrusive monitoring and prediction of driver fatigue. IEEE transactions on vehicular technology 53, 4 (2004), 1052–1068.
- Christopher Kanan, Nicholas A Ray, Dina NF Bseiso, Janet H Hsiao, and Garrison W Cottrell. 2014. Predicting an observer's task using multi-fixation pattern analysis. In Proceedings of the symposium on eye tracking research and applications. ACM, 287–290.
- Enkelejda Kasneci, Katrin Sippel, Kathrin Aehling, Martin Heister, Wolfgang Rosenstiel, Ulrich Schiefer, and Elena Papageorgiou. 2014. Driving with binocular visual field loss? A study on a supervised on-road parcours with simultaneous eye and head tracking. *PloS one* 9, 2 (2014), e87470.
- Thomas Kübler, Shahram Eivazi, and Enkelejda Kasneci. 2015a. Automated visual scanpath analysis reveals the expertise level of micro-neurosurgeons. In MICCAI workshop on interventional microscopy. 1–8.

- Thomas C Kübler, Enkelejda Kasneci, Wolfgang Rosenstiel, Martin Heister, Kathrin Aehling, Katja Nagel, Ulrich Schiefer, and Elena Papageorgiou. 2015b. Driving with glaucoma: task performance and gaze movements. *Optometry and Vision Science* 92, 11 (2015), 1037–1046.
- Thomas C Kübler, Colleen Rothe, Ulrich Schiefer, Wolfgang Rosenstiel, and Enkelejda Kasneci. 2017. SubsMatch 2.0: Scanpath comparison and classification based on subsequence frequencies. *Behavior research methods* 49, 3 (2017), 1048–1064.
- Paul Locher, Elizabeth Krupinski, and Alexandra Schaefer. 2015. Art and authenticity: Behavioral and eye-movement analyses. Psychology of Aesthetics, Creativity, and the Arts 9, 4 (2015), 356.
- Carmel M Loughland, Leanne M Williams, and Evian Gordon. 2002. Visual scanpaths to positive and negative facial emotions in an outpatient schizophrenia sample. *Schizophrenia research* 55, 1-2 (2002), 159–170.
- David Manning, Susan Ethell, Tim Donovan, and Trevor Crawford. 2006. How do radiologists do it? The influence of experience and training on searching for chest nodules. *Radiography* 12, 2 (2006), 134–142.
- Davide Massaro, Federica Savazzi, Cinzia Di Dio, David Freedberg, Vittorio Gallese, Gabriella Gilli, and Antonella Marchetti. 2012. When art moves the eyes: a behavioral and eye-tracking study. *PloS one* 7, 5 (2012), e37285.
- David Noton and Lawrence Stark. 1971. Scanpaths in eye movements during pattern perception. Science 171, 3968 (1971), 308–311.
- Mustafa Ozuysal, Michael Calonder, Vincent Lepetit, and Pascal Fua. 2010. Fast keypoint recognition using random ferns. *IEEE transactions on pattern analysis and machine intelligence* 32, 3 (2010), 448–461.
- Oskar Palinko, Andrew L Kun, Alexander Shyrokov, and Peter Heeman. 2010. Estimating cognitive load using remote eye tracking in a driving simulator. In Proceedings of the 2010 symposium on eye-tracking research & applications. ACM, 141–144.
- Kevin A Pelphrey, Noah J Sasson, J Steven Reznick, Gregory Paul, Barbara D Goldman, and Joseph Piven. 2002. Visual scanning of faces in autism. *Journal of autism and* developmental disorders 32, 4 (2002), 249–261.
- A Van der Gijp, CJ Ravesloot, H Jarodzka, MF Van der Schaaf, IC Van der Schaaf, Jan PJ van Schaik, and Th J Ten Cate. 2017. How visual search relates to visual diagnostic performance: a narrative systematic review of eye-tracking research in radiology. Advances in Health Sciences Education 22, 3 (2017), 765–787.
- Alfred L Yarbus. 1967. Eye movements during perception of complex objects. In Eye movements and vision. Springer, 171–211.
- WH Zangemeister, Keith Sherman, and Lawrence Stark. 1995. Evidence for a global scanpath strategy in viewing abstract compared with realistic images. *Neuropsychologia* 33, 8 (1995), 1009–1025.
- A Tianyi Zhang and B Olivier Le Meur. 2018. How Old Do You Look? Inferring Your Age from Your Gaze. In 2018 25th IEEE International Conference on Image Processing (ICIP). IEEE, 2660–2664.