Discovery of Activities via Statistical Clustering of Fixation Patterns

Jeffrey B. Mulligan; NASA Ames Research Center; Moffett Field, CA / USA

Abstract

Human behavior often consists of a series of distinct activities, each characterized by a unique pattern of interaction with the visual environment. This is true even in a restricted domain, such as a pilot flying an airplane; in this case, activities with distinct visual signatures might be things like communicating, navigating, monitoring, etc. We propose a novel analysis method for gaze-tracking data, to perform blind discovery of these hypothetical activities. The method is in some respects analogous to recurrence analysis, which has previously been applied to eye movement data. In the present case, however, we compare not individual fixations, but groups of fixations aggregated over a fixed time interval (τ). We assume that the environment has been divided into a finite set of discrete areas-of-interest (AOIs). For a given time interval, we compute the proportion of time spent fixating each AOI, resulting in an N-dimensional vector, where N is the number of AOIs. These proportions can be converted to integer counts by multiplying by τ divided by the average fixation duration, a parameter that we fix at 283 milliseconds. We compare different intervals by computing the chi-squared statistic. The p-value associated with the statistic is the likelihood of observing the data under the hypothesis that the data in the two intervals were generated by a single process with a single set of probabilities governing the fixation of each AOI. We cluster the intervals, first by merging adjacent intervals that are sufficiently similar, optionally shifting the boundary between non-merged intervals to maximize the difference. Then we compare and cluster non-adjacent intervals. The method is evaluated using synthetic data generated by a hand-crafted set of activities. While the method generally finds more activities than put into the simulation, we have obtained agreement as high as 80% between the inferred activity labels and ground truth.

Introduction

It has long been known that patterns of eye fixations made when exploring a scene can reflect specific information-gathering goals. An early demonstration was Yarbus' studies of subjects viewing a painting, while trying to answer different questions about the contents of the scene[1]. Presumably similar task-dependent behaviors are manifested in actual work situations, such as piloting an aircraft. In this paper, we explore a method for automatically discovering the fixation behaviors linked to different work activities. The method is "blind," in that it does not require that the activities be identified ahead of time, rather it discovers the activities based on their unique pattern of eye fixations.

It is assumed that human behavior can be modeled as a sequence of distinct activities (see Figure 1). In this work, we assume that the number and nature of the activities is not known beforehand; the goal is to discover the activities, based on differences in the patterns of eye movements that are produced.

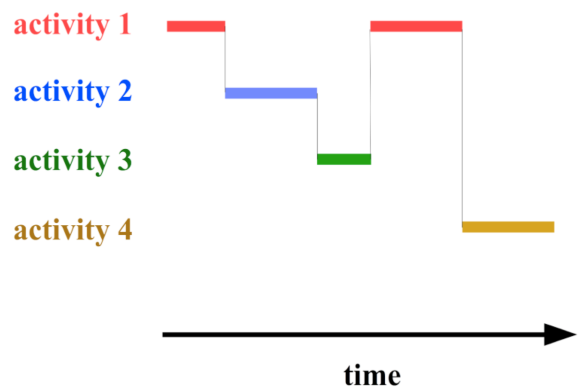

Figure 1: Cartoon illustrating the concept of sequential switching between activities. This model assumes that only a single activity is carried out at any time, so apparent "multi-tasking" can only occur via rapid switching.

The general strategy will be to choose a temporal window size, and then compare the eye movement patterns occurring in pairs of windows at different times. Repeating patterns are consistent with a recurring activity, while differing patterns indicate different activities. The critical parameter of the analysis is the duration of the temporal window: if it is longer than the length of time for which the subject persists in a given activity, then different activities will be blurred together, and we will not be able to resolve the individual activity transitions. Conversely, if the window is too small, then we may not be able to aggregate enough data to reliably differentiate between different activities. The success or failure of the method will therefore rely upon two parameters of human behavior: the duration of the activity segments, and the magnitudes of the differences in the eye movement patterns.

A key component of the approach is the comparison of two eye movement samples, or *scan paths*. Many methods have been proposed; a good summary is provided in a recent review paper [2]. Methods may be divided according to whether or not the temporal order of the fixations is considered; in the present work, we choose to disregard temporal order, based on the intuition that many activities will consist of gathering and integrating information prior to making a decision and/or executing an action. It does not really matter in what order the constituent elements of information are gathered, although it is likely that fixations to the action location will always follow fixations to the information sources (when these are distinct). In the present work, the intervals are treated simply as a "bag of fixations" - the data are completely characterized by the number of fixations to each area within the temporal window.

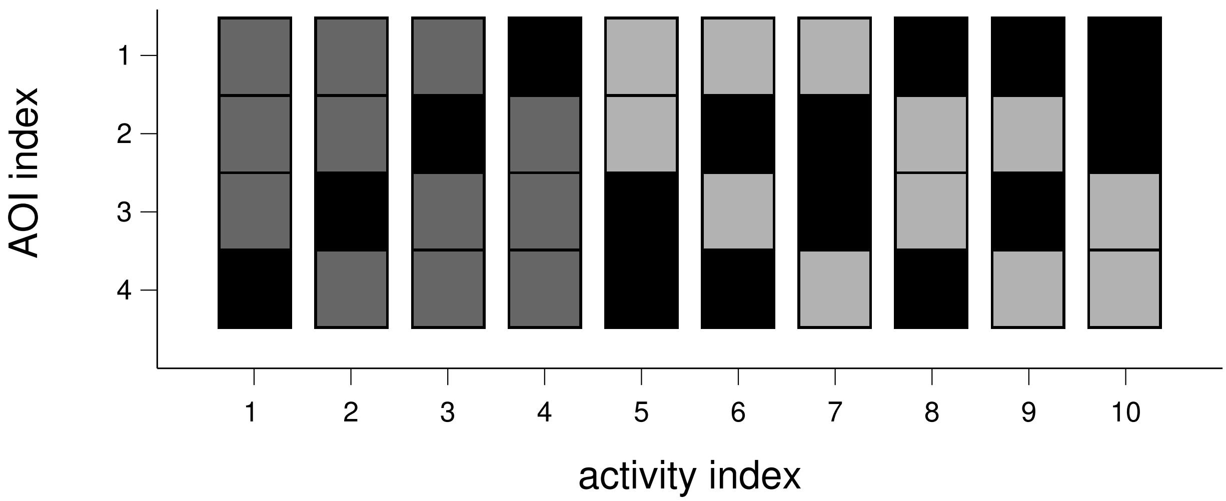

Figure 2: Graphical depiction of the probability vectors for the 10 synthetic activities. Activities 1-4 are the "leave-one-out" activities, in which three of the four AOIs are each visited with equal probability of $\frac{1}{3}$ (shown as dark gray); activities 5-10 are the "leave-two-out" activities, in which two of the four AOIs are visited with equal probability of $\frac{1}{2}$ (shown as light gray). Black cells represent a probability of 0.

Methods

We assume that the input data consist of a time series of region labels, representing the objects on which the subject's gaze falls at each time step. This assigment of labels to regions in the environment is often done by the eye-tracking system, based on a model specified by the experimenter prior to collecting the data. The regions are typically referred to as "Areas Of Interest" (AOIs). The raw data provided by the eye tracking system is assumed to include a time series of AOI labels at the sampling rate of the system. In the case of a system that detects fixations and provides summary data such as start time and duration, those data may be easily converted to a time series.

Generation of synthetic data

Synthetic data were generated using a first-order Markov process. This means that the probability that a particular AOI is fixated next depends only on the current AOI. If i is the index of the present AOI, then the probability that the next AOI has index j will be written as m_{ij}. The m_{ij} are the elements of the Markov matrix. The elements of the matrix are constrained in that each column must sum to 1, reflecting the fact that regardless of the index (column) of the current AOI, the next fixation must go *somewhere*. The asymptotic probabilities that a given area will be visited may be obtained by exponentiating the matrix; for small matrices, this process stabilizes after just a few iterations.

In pilot work, activities based on randomly generated Markov matrices were compared, but these proved difficult to discriminate using reasonable activity durations. To give the method a fighting chance, we hand-crafted a set of ten activities over a toy environment having 4 AOIs. Four activities consisted of visitation to three of the four AOIs, with equal probability, which we refer to as the "leave-one-out" activities. Additionally, six "leave-two-out" activities were created, having equal probabilities of visiting the two remaining AOIs. The Markov matrices were constructed such that transition probabilities were independent of the current state, making this effectively a 0-order Markov process. Figure 2 shows a graphical representation of these 10 activities, where the gray level represents the probability of visiting each of the 4 AOIs.

Simulated data were generated by running the model, simulating an eye tracker data rate of 60 samples per second. The fixation durations were uniformly distributed from 243 to 323 milliseconds (15-19 frames), with a mean duration of 283 milliseconds (17 frames).

Comparison of fixation sequences

Following a suggestion by Ahumada [3], the chi-squared statistic [4, 5] is used to compare short sequences of data. The null hypothesis is that a single process (defined as a set of probabilities for fixating the various AOIs) produced both sequences. The observations from both intervals are combined to form an estimate of those probabilities, and the deviations of the observations are tallied, producing a statistic that is small when the two intervals are similar, and large when they are different.

The computation is performed using the fixation counts, i.e. the number of fixations made to each AOI. Let a_i and b_i be the the sets of fixations counts for the two intervals, where the subscript i indexes the AOI, ranging from 1 to N. We begin by computing the total fixations for each interval,

$$T_a = \sum_{i=1}^{N} a_i \qquad T_b = \sum_{i=1}^{N} b_i. \tag{1}$$

The expected fraction of fixations made to each AOI under the null hypothesis, f_i, is computed by averaging the two intervals,

$$f_i = \frac{a_i + b_i}{T_a + T_b}. \tag{2}$$

The expected number of fixations to each AOI in each interval is then computed by multiplying the expected fraction times the total number of fixations in each interval:

$$e_{a,i} = f_i * T_a, \qquad e_{b,i} = f_i * T_b. \tag{3}$$

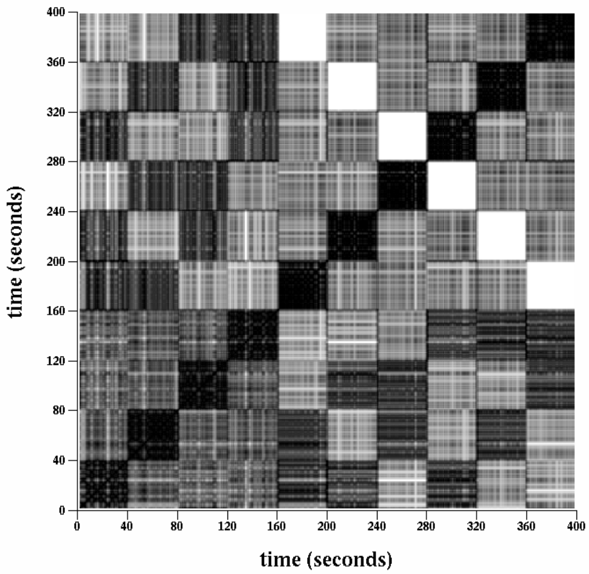

Figure 3: Matrix of chi-squared values.

Finally, the test statistic is formed by summing the squared deviations between the observed and expected values, normalized by the expected value:

$$s = \sum_{i=1}^{N} \left[\frac{(a_i - e_{a,i})^2}{e_{a,i}} + \frac{(b_i - e_{b,i})^2}{e_{b,i}} \right]. \quad (4)$$

This statistic is approximately distributed as the χ^2 distribution with $N-1$ degrees of freedom [5, 6].

To perform the analysis, an interval duration τ is first selected. In earlier versions of this work, a set of partially-overlapping intervals was considered, with each interval advanced by $\frac{\tau}{4}$. This was done to improve temporal resolution for activity changes. Within an interval, the numbers of 60 Hz samples for each AOI were totalled, and a fixation count was obtained by dividing this number by 17 (the number of samples in the mean fixation duration of 280 milliseconds). Thus, long fixations, produced when the Markov model repeated a particular AOI one or more times, were counted as multiple fixations. The chi-squared statistic was then computed for all interval pairs.

Figure 3 shows a matrix of chi-squared statistics computed in this way. In the figure, the numerical values of the statistics have been converted to gray levels, scaled so that a value of 0 is mapped to black, and the highest value is mapped to white. The 10 by 10 block structure of the image reflects the 10 synthetic activities (diagrammed in Figure 2), which were run in sequence, each for a duration of 40 seconds. The main diagonal consists of 10 mostly-black blocks, showing that low chi-squared values are obtained when a sample from a given activity is compared to another sample of the same activity. The matrix can be further subdivided into a 4 block by 4 block sub-matrix in the lower left-hand (corresponding to comparisons between the 4 leave-one-out activities), and a 6 block by 6 block sub-matrix in the upper right (corresponding to the 6 leave-two-out activities). This 6 by 6 block shows a white "anti-diagonal" - these blocks correspond to comparisons of samples from activity pairs that have no shared AOIs. The other off-diagonal blocks in this upper sub-matrix are relatively bright, indicating good discriminability between pairs of leave-two-out activities with one shared AOI. This can be contrasted with the relatively dark off-diagonal blocks in the lower left sub-matrix, showing relatively poor discrimination between pairs of leave-one-out activities, as each pair necessarily has two shared AOIs. The 4 block by 6 block sub-matrices in the upper left and lower right correspond to comparisons between a leave-one-out activity (visiting 3 AOIs) and a leave-two-out activity (visiting 2 AOIs). Discriminability is relatively poor when 2 AOIs are shared, and relatively good when only 1 AOI is shared.

Although clustering can be performed on the full matrix, as shown in Figure 3, this method was abandoned, for a couple of reasons. First, it is computationally expensive to compute the full matrix. Second, the fact that overlapping subsequences are not statistically independent invalidates that standard chi-squared calculation described above, and while it may be possible to correct for the partial overlap, it seemed an unnecessary complication. Instead, a procedure was developed to refine the location of the boundary between adjacent subsequences.

Clustering

The procedure used to generate the results shown in the rest of this paper began with a set of non-overlapping adjacent intervals of duration τ. A first pass attempted to merge adjacent intervals by computing the chi-squared statistic, and merging when the associated p value was above a threshold. The intervals were processed in temporal order, and the process repeated until no more merges could be performed.

The significance level chosen to reject a merge is a critical parameter. A high value of the chi-squared statistic, producing a low p-value, says that it is unlikely that the two samples were produced by the same process. When we merge two intervals, we are effectively accepting the null hypothesis; therefore, to be conservative, we would choose a large critical p-value rather than a small one. For example, if we were to use a small p-value, such as 0.01, we might accept a merge of two intervals with a chi-squared p-value of 0.02, meaning that we would be assuming that the two samples came from a single process in spite of the fact the the odds of observing the data in that case would only be 1 in 50.

These merge passes can be optionally interleaved with a boundary adjustment phase. In this phase, each pair of adjacent intervals is examined to see whether a better separation could be obtained by shifting the boundary between them. A set of possible boundary locations centered on the current boundary is exhaustively searched, by computing the chi-squared statistic for each division, and choosing the one with the largest value. The search space extended by $\frac{\tau}{2}$ in each direction, unless the interval length is less than τ, in which case only half of the interval is searched. The interval pairs were considered in order, and intervals whose boundaries were adjusted were added to a list for reconsideration in a subsequent pass. Typically, the first pass will adjust almost all of the interval boundaries, reflecting the fact that the true activity boundaries rarely align with a multiple of τ. Most, but usually not all of these will be stable on the second pass, so the third pass will reconsider a relatively small number of interval pairs, corresponding to those containing an interval that is adjusted on the second pass. The process terminates when a pass makes no adjustments. This entire process is relatively time-consuming, which is why we also computed the results obtained when it is omitted.

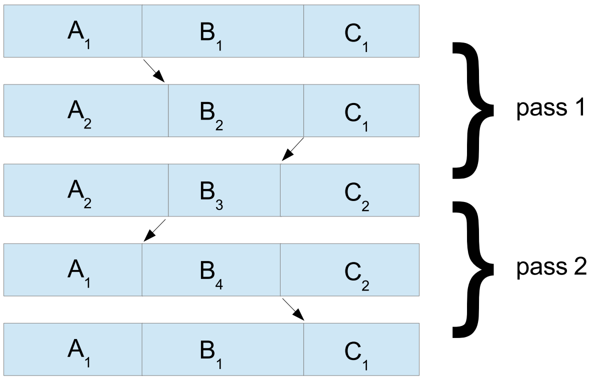

Figure 4: Cartoon illustrating rare but occasionally-occurring instability in the interval boundary adjustment procedure.

The basic method just described occasionally fails to terminate. We present here a description of one pattern that has been observed, which is illustrated in Figure 4. We assume there are three adjacent intervals that we call A_1, B_1, and C_1. The first pass adjusts the boundary between A_1 and B_1, resulting in new intervals A_2 and B_2, and then adjusts the boundary between B_2 and C_1, resulting in new intervals B_3 and C_2. The second pass adjusts the boundary between A_2 and B_3, moving it *back to its original location*, resulting in the original interval A_1, along with new interval B_4. Finally, the boundary between B_4 and C_2 is adjusted, and is also returned to its original location, resulting in intervals B_1 and C_1 once again. Thus, after two passes the partitions have been changed on each pass but have returned to the original configuration; as the standard algorithm iterates until a pass is performed with no changes, it will never terminate with this input. This problem was solved in a somewhat arbitrary way, by keeping a count of the number of passes, and reducing the boundary point search space by this count. Eventually this will fail to consider one of the shifts, and the process will then terminate.

When all possible merges of adjacent intervals have been performed, we search for the underlying activities by examining pairs of non-adjacent intervals for possible clustering. We first create one activity for each interval, initializing the probabilities from the observed fixation distributions. Then we construct the matrix of chi-squared statitistics and associated probabilities for all activity pairs. The entries on the diagonal (corresponding to comparing an activity with itself) have the statistic set to 0, but the probability is set to -1, indicating an invalid entry. We then search the matrix of p-values for the highest value. If that value exceeds a pre-determined threshold, then the two activities are assumed to have been generated by the same underlying process, and the fixation counts are pooled. The row and column of the matrix corresponding to the second activity are deleted, and those corresponding to the first activity are recomputed using the pooled counts. After the merge, there is one fewer activity, and the matrix is reduced in size by one. This process is repeated until there are no entries that exceed the threshold.

Balancing activity sequences

One of the challenges of this approach is correctly locating the activity transition boundaries. This is especially true when we begin with a large temporal window duration τ. When an interval straddles a transition between two activities, it might appear to have been generated by a third activity that looks like the average of the first two. For example, for the activities shown in Figure 2, an interval containing data from both activities 9 and 10 (visiting AOIs 2 and 4, and 3 and 4, respectively) could be quite similar to activity 4, which visits AOIs 2, 3, and 4. Therefore, in studying the performance of the system it is important that the synthetic data that we use for evaluation contain all possible transitions. The data shown in Figure 3 only sample 9 of the 45 possible transitions (90 if order is considered).

Sequences that are balanced in this way are known as *De Bruijn sequences* [7, 8]. For our test case, we construct a sequence by first creating a fully-connected graph with 10 nodes, representing the 10 activities. We then wish to find an *Eulerian circuit* of this graph, i.e. a traversal that uses each edge exactly once. Such a circuit does not exist in this case, because all of the nodes have an odd number of edges. (For an Eulerian circuit to exist, there can be no more than two.) However, it is possible to construct a circuit that traverses each edge twice, once in each direction. There is in fact a very large number of such circuits, and we constructed one arbitrarily. In order to sample each of the 90 ordered transitions, each synthetic data sequence consisted of a series of 91 40-second intervals in which a single activity was used to generate fixations. The duration of individual fixations was jittered as described above, sampling from a uniform distribution ranging from 15 to 19 frames. In the course of the sequence, 9 of the activities were repeated 9 times, while the first activity was repeated at the end, giving it an extra interval. As we simulated a 60 Hz eye-tracking system, each 40-second interval was made up of 2400 samples, and a complete record of the 91 intervals contained 218,400 samples. Ten different sequences were generated for use in the analysis.

Two parameters controlled the behavior of the analysis: τ, the initial duration of the analysis interval, and p_θ, the probability threshold for the chi-squared analysis. The data were analyzed using various values for these two parameters, with four values of τ corresponding to 5, 10, 20 and 40 fixations, and values for p_θ of 0.01, 0.02, 0.05, 0.1 and 0.2. A higher value of p_θ corresponds to a more conservative criterion for merging activities.

Results

The number of discovered activities is plotted as a function of the threshold probability p_θ in Figure 5. It can be seen that the number of activities grows rapidly as p_θ is increased, and that even for a seemingly liberal merging criterion of $p_\theta = 0.01$ more than 20 activities are usually discovered. (Recall that the true value is 10.)

The number of activities, however, does not really give us a fair assessment of how well the method is performing, because many of the activities are represented by short intervals, and so make up a small part of the data. To obtain a measure that takes this into account, we first determine which of our discovered activities correspond most closely to the ground truth activity probabilities shown in Figure 2. To do this, we take the probability vector of each of the ground truth activities in turn, and search the set of discovered activities for the one whose vector of probabilities is closest it. (We use the Euclidean distance between the probability vectors to define ”closest.”) This discovered activ-

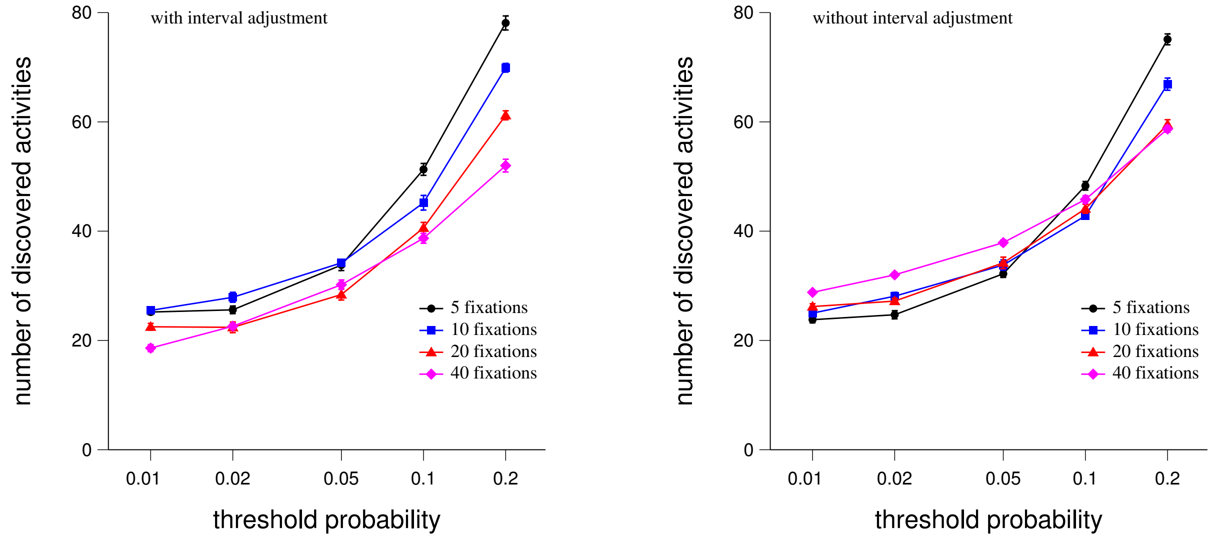

Figure 5: Number of discovered activities is plotted as a function of the threshold probability p_θ used in the chi-squared analysis. The different curves show the results for different values of the initial interval length τ. The left panel shows results obtains using the interval boundary adjustment step, while the right panel shows the results with that step omitted.

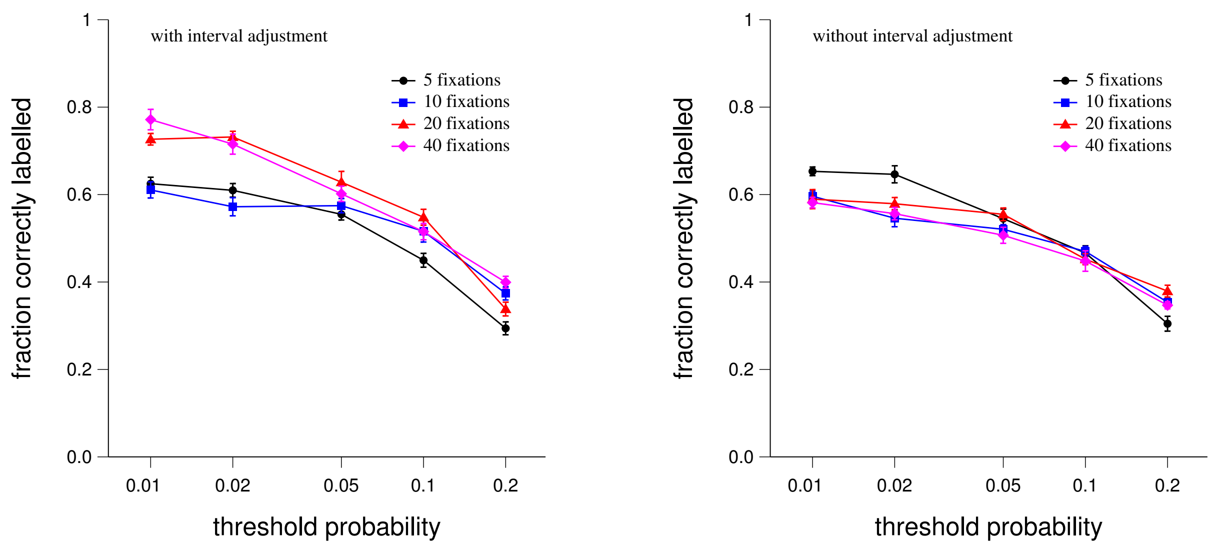

Figure 6: Proportion of samples correctly-labeled is plotted as a function of threshold probability p_θ, for different values of the initial interval length τ as in Figure 5. Note that the best performance is obtained when interval boundary adjustment is used (left panel), for the smallest value of p_θ (0.01) and the largest value of τ (40 fixations).

ity is then associated with the ground truth activity. Thus, from the set of discovered activities, we obtain a subset of 10 that we assume to be "correct." We compute the fraction of samples contained within these correctly labelled intervals; this is plotted as a function of threshold probability in Figure 6. While most of these correctly-labeled samples were indeed generated by the associated ground truth activity, there is the possibility that errors in activity boundaries may cause a small number of samples to be mis-labeled, so strictly speaking these numbers are really upper bounds.

A visualization of the labeling is provided in Figure 7. This figure represents the first 26 minutes of a single instance of a simulation of 60 minutes of behavior. The top row, labeled (a), is an idealized version of the ground truth. Each 40 second interval is rendered a graphical representation of the generating probabilities previously shown in Figure 2, with color tints added to aid identification. This sequence of activities was generated by an instance of a De Bruijn sequence beginning 1, 4, 7, 10, 3, 6, 9, 2, ... The second row, labeled (b), is similar to panel (a), but here the lightnesses represent not the generating probabilities, but rather the observed probabilities for one particular simulation; these differ from the generating probabilities due to stochastic variation. The

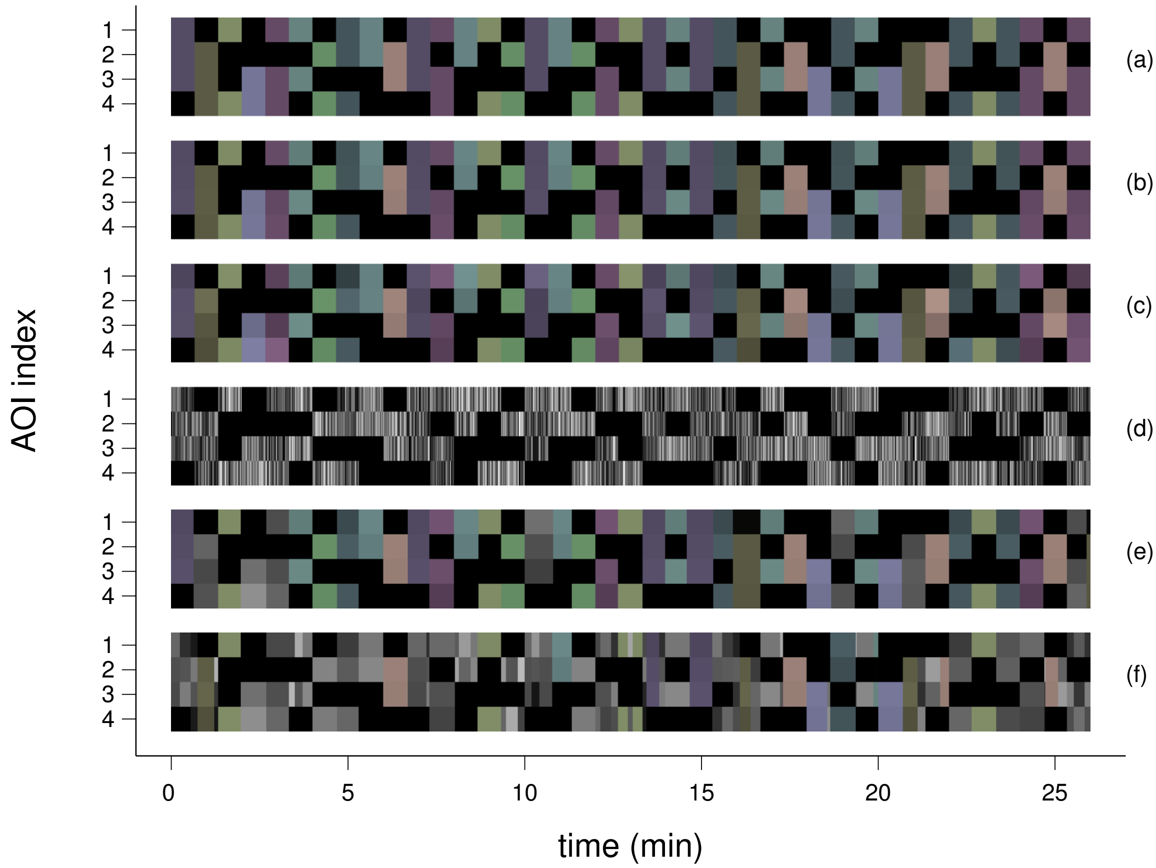

Figure 7: Graphical representation of activity sequences, i.e. the probabilities associated with each of the 4 AOIs as a function of time. The plot shows the first 26 minutes of one of the ten hour-long records of synthetic data used in the analysis. Panel (a): idealized representation of ground truth. Each 40 second interval is rendered a graphical representation of the generating probabilities previously shown in Figure 2. Color tints have been added to aid identification. Panel (b): similar to panel (a), but here the lightnesses represent not the generating probabilities, but rather the observed probabilities for this particular simulation, which differ from the generating probabilities due to stochastic variation. The probabilities have been pooled across different intervals containing the same activity. This represents the best result that we could hope to obtain from a single data set. Panel (c): similar to panel (b), but without pooling across non-adjacent intervals. Panel (d): an approximation of the raw data input to the method. (A small amount of subsampling has been performed.) Panel (e): results of the method for $p_\theta = 0.01$ and $\tau = 40$ fixations, with interval adjustment. Color tints have been applied to the correctly identified activities, as described in the text. This is an example of the best performance obtained to date (approx 80% correct). Panel (f): similar to panel (e), this panel shows results obtained with $p_\theta = 0.2$ and $\tau = 40$ fixations.

probabilities have been pooled across different intervals containing the same activity. This represents the best result that we could hope to obtain from a single data set, in which the segmentation of activities matches exactly the ground truth, repeated activities are correctly clustered, and the underlying probabilities are estimated by pooling over the entire cluster. The third row, labeled (c), is similar to panel (b), but without pooling across non-adjacent intervals. This shows the stochastic variation in the observed AOI proportions between different instances of an activity. For example, activity 1 (tinted purple) occurs five times, at minute 0, 6.7, 10, 13.3 and 14.7. Each instance has slightly different observed probabilities, and indicated by the intensities of the corresponding squares. For example, the segment beginning at the 10 minute mark has a higher than average visitation of AOI 1, while the segment beginning at 13.3 minutes visits AOI 1 less than average. The fourth row, labeled (d), shows the raw fixation data, subsampled slightly to a resolution appropriate for the figure. These data are the input to the method. The fifth and sixth rows, labeled (e) and (f), show the results of applying the method for two sets of parameter values. Panel (e) shows results obtained using $p_\theta = 0.01$ and $\tau = 40$ fixations, which have produced the best results to date (as seen above in Figure 6), with approximately 80% of the intervals labeled correctly. Of the 5 intervals generated from activity 1 discussed above, it can be seen that the first, second, fourth and fifth have been correctly clustered, as indicated by the purple tint, which the third (at the 10 second mark) was considered to be a distinct activity, with a higher probability associated with AOI 1. It can also be noted that the interval associated with this activity is longer than the ground truth value of 40 seconds. Finally, panel (f) shows results obtained with a more conservative criterion for merging intervals, with $p_\theta = 0.2$ and $\tau = 40$ fixations. While some of the intervals are still correctly labeled, it can be seen that there are many more short intervals that have been classified as distinct activities.

Discussion

The results demonstrate that the method is capable of blind discovery of activities, provided the true activities are sufficiently distinct. It remains an empirical question whether real human activities meet this criterion. It should also be noted that there is a trade-off between the distinctness of the activities and the durations over which single activities are perseverated: less distinct activities can be discriminated if a longer time window is used over which to amass statistics. But to be useful in a given domain of human behavior, this window should not be longer than the minimum duration of an activity segment; this duration is likely to vary across domains.

The analysis presented here treats each interval as a "bag of fixations," without regard to the order in which the various AOIs are fixated. In real world human activities, it is possible that this simplification will result in throwing away useful information present in the fixation sequence. The analysis could be extended to incorporate order in a fairly straightforward way, by tabulating transitions from one AOI to another. In the example considered here, instead of having 4 AOIs, we would need to consider 16 ordered transitions. For the simple, hand-crafted activities used in our simulations, there are 9 transitions with non-zero probabilities for each of the leave-one-out activities, and 4 transitions for each of the leave-two-out activities. This illustrates a fundamental problem: a more detailed activity model (incorporating transition probabilities) necessarily has more parameters to estimate, and will therefore require more data (i.e., longer intervals generated by single activities). A similar problem was studied by Kontsevich [**?**], who developed a statistical model of human behavior in a penny-matching game. The model had a variable parameter, the number of previous plays determining the next play. When this number was very small, the model was simple and easy to accurately estimate, but did not have good predictive power, because it neglected important components of the behavior. Conversely, when a large number of previous plays was incorporated into the model, performance was poor because it was impossible to accurately estimate the parameters from the available data. A "sweet spot" was found using a small number of previous responses, which balanced complexity of the model against availability of data to use to estimate the parameters, allowing the computer to model (and beat) the human.

Summary

A method has been presented for the blind discovery of behavioral activities on the basis of eye fixation patterns. As presented here, the method relies upon the proportion of fixations made to distinct areas-of-interest (AOIs); more complicated models of the behavior incorporating order of visitation and transition probabilities are possible, but require more data. The method is potentially applicable to a variety of operator monitoring scenarios.

References

[1] Yarbus, A. L., [*Eye movements and vision*], Plenum Press (1967). translated from Russian by B. Haigh.

[2] Anderson, N. C., Anderson, F., Kingstone, A., and Bischof, W. F., "A comparison of scanpath comparison methods," *Behav. Res.* **47**, 1377–1392 (2015).

[3] Ahumada, Jr., A. J. personal communication (2015).

[4] Pearson, K., "On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling," *Philosophical Magazine, Series 5* **50**(302), 157–175 (1900).

[5] Chi-squared test, "Chi-squared test — Wikipedia, the free encyclopedia," (2018). [Online; accessed 13-September-2018].

[6] Chi-squared distribution, "Chi-squared distribution — Wikipedia, the free encyclopedia," (2018). [Online; accessed 13-September-2018].

[7] De Bruijn sequence , "De bruijn sequence — Wikipedia, the free encyclopedia," (2019). [Online; accessed 22-February-2019].

[8] Aguirre, G. K., Mattar, M. G., and Magis-Weinberg, L., "de bruijn cycles for neural decoding," *Neuroimage* **56**(3), 1293–1300 (2011).

Author Biography

Jeffrey B. Mulligan received the A.B. degree in physics from the Harvard University (1980), and the M.A. and Ph.D. degrees in psychology from the University of California at San Diego (1982, 1986). Since then he has worked as a computer engineer at the NASA Ames Research Center, in Moffett Field, CA. His work has focused on visual perception and eye movements, and applications to the design of aerospace information displays. He is a member of the Vision Sciences Society, the Optical Society, and the Society for Information Display.