

# Defending Yarbus: Eye movements reveal observers’ task

Ali Borji<sup>a,\*</sup>, Laurent Itti<sup>a,b,c</sup>

<sup>a</sup>*Department of Computer Science, University of Southern California, 3641 Watt Way, Los Angeles, CA 90089*

<sup>b</sup>*Department of Neuroscience, University of Southern California, 3641 Watt Way, Los Angeles, CA 90089*

<sup>c</sup>*Department of Psychology, University of Southern California, 3641 Watt Way, Los Angeles, CA 90089*

---

## Abstract

In a very influential yet anecdotal illustration, Yarbus suggested that human eye movement patterns are modulated top-down by different task demands. While the hypothesis that it is possible to decode the observer’s task from eye movements has received some support (e.g., Iqbal & Bailey (2004); Henderson et al. (2013)), Greene et al. (2012) argued against it by reporting a failure. In this study, we perform a more systematic investigation of this problem, probing a larger number of experimental factors than previously. Our main goal is to determine the informativeness of eye movements for task and mental state decoding. We perform two experiments. In the first experiment, we re-analyze the data from a previous study by Greene et al. (2012) and contrary to their conclusion, we report that it is possible to decode the observer’s task from aggregate eye movement features slightly but significantly above chance, using a Boosting classifier (34.12% correct vs. 25% chance-level; binomial test,  $p = 1.0722e - 04$ ). In the second experiment, we repeat and extend Yarbus’ original experiment by collecting eye movements of 21 observers viewing 15 natural scenes (including Yarbus’ scene) under Yarbus’ seven questions. We show that task decoding is possible, also moderately but significantly above chance (24.21% vs. 14.29% chance-level; binomial test,  $p = 2.4535e - 06$ ). We thus conclude that Yarbus’ idea is supported by our data and continues to be an inspiration for future computational and experimental eye movement research. From a broader perspective, we discuss techniques, features, limitations, societal and technological impacts, and future directions in task decoding from eye movements.

*Keywords:* visual attention, eye movements, bottom-up saliency, top-down attention, free viewing, visual search, mind reading, task decoding, Yarbus

---

## 1. Introduction

Eyes are windows to perception and cognition. They convey a wealth of information regarding our mental processes. Indeed this has been elegantly demonstrated by seminal works of Guy T. Buswell (1935) and I. A. Yarbus (1967), who were the first to investigate the relationship between eye movement patterns and high-level cognitive factors. Yarbus recorded an observer’s eye movements (with his home-made gaze tracking suction cap device) while s/he viewed the I. E. Repin’s painting, “The Unexpected Visitor” (1884)<sup>1</sup>. He illustrated fixations of the observer as s/he viewed the painting under seven different instructions: 1) Free examination, 2) Estimate the material circumstances of the family, 3) Give the ages of the people, 4) Surmise (guess) what family had been doing before the arrival of the unexpected visitor, 5) Remember the cloths worn by the people, 6) Remember positions of people and objects in the room, 7) Estimate how long the visitor had been away from the family.

Yarbus’ results show striking differences in eye movement patterns across instructions, over the same visual stimulus. Early in the viewing period, fixations were particularly directed to the faces of the individuals

---

\*Corresponding author

*Email addresses:* borji@usc.edu (Ali Borji), itti@pollux.usc.edu (Laurent Itti)

<sup>1</sup>See Figure 4, image number 5.

in the painting and also observers showed a strong preference to look at the eyes more than any other features of the face. Yarbus concluded that the eyes fixate on those scene elements which carry useful information, thus showing where we look depends critically on our cognitive task. Further, Yarbus' experiments point towards the active nature of the human visual system, as opposed to passively or randomly sampling the visual environment. This active aspect of vision and attention has been extensively investigated by Dana Ballard, Mary Hayhoe, Michael Land and others who studied eye movements in the context of natural behavior. Please see Ballard et al. (1995); Hayhoe & Ballard (2005); Tatler et al. (2011); Tatler & Vincent (2009); Schütz et al. (2011); Navalpakkam & Itti (2005); Land & Hayhoe (2001); Land (2006); Itti & Koch (2001); Borji & Itti (2012) for recent reviews.

Two prominent yet contrasting hypotheses attempt to explain eye movements and attention in natural behavior. First, according to the cognitive relevance hypothesis, eyes are driven by top-down factors that intentionally direct fixations toward informative task-driven locations (e.g., in driving). Second, in the absence of such task demands (e.g., in scene free-viewing), eyes are directed to low-level image discontinuities such as bright regions, edges, colors, etc., so called salient regions. This is often referred to as the saliency hypothesis (Treisman & Gelade, 1980; Koch & Ullman, 1985; Parkhurst et al., 2002; Itti et al., 1998). Both hypotheses are likely to be correct, yet the relative contribution of top-down and bottom-up attentional components varies across daily behaviors. Conversely, by looking at eye movements, one could possibly infer the underlying factors affecting fixations (i.e., task at hand or mental state), or gain insights into what an observer is currently thinking. Active research is undergoing to discover the interplay between top-down task-driven factors and bottom-up stimulus-driven factors in driving visual attention and to assess the amount of information eye movements convey regarding mental thoughts.

Yarbus showed a proof of concept with a single observer but did not conduct a comprehensive quantitative analysis. Perhaps DeAngelus & Pelz (2009) were the first to confirm Yarbus' findings, with multiple observers viewing Repin's painting. Viewing times in their study were self-paced (9 to 50 seconds), and were significantly less than the enforced 3-minute viewing time of Yarbus' observer. DeAngelus and Pelz showed that observers' eye movement patterns were similar to those reported by Yarbus, with faces invariably fixated, and the overall viewing pattern varying with task instruction. A few of their observers, especially those with shorter viewing times, did not show dramatic shifts with instruction. The task "Give the ages of the people" resulted in the smallest inter-observer distance of all tasks, indicating that for this task the eye movement patterns were most similar among the observers. The "Estimate how long the visitor had been away from the family" task showed the most variability among observers, suggesting that observers used different viewing strategies to complete this task.

The general trend for fixations when viewing scenes to fall preferentially on persons within the scene had been shown previously by Buswell (1935). The tendency of observers to fixate on faces has recently been quantitatively confirmed by Cerf et al. (2009) and further supported by large-scale eye tracking studies (e.g., Judd et al. (2009); Subramanian et al. (2010)). Yarbus' results (along with DeAngelus & Pelz (2009)) indicate that, for extended viewing times, observers show a clear tendency to make repeated cycles of fixations between the key features of a face or a scene (a.k.a cyclic behavior). Both attention and face perception communities have been largely inspired by Yarbus' early insights (See Kingstone (2009)).

Castelhano et al. (2009) investigated how task instruction influences specific parameters of eye movement control. They asked 20 participants to view color photographs of natural scenes under two instruction sets: searching a scene for a particular item or remembering characteristics of that same scene. They found that viewing task biases aggregate eye-movement measures such as average fixation duration and average saccade amplitude. Mills et al. (2011) examined the influence of task set on the spatial and temporal characteristics of eye movements during scene perception. They found that task affects both spatial (e.g., saccade amplitude) and temporal characteristics of fixations (e.g., fixation duration).

Tatler et al. (2010) explored Yarbus' biography, his scientific legacy including his eye tracking apparatus, and his key contributions. They recorded eye movements of observers when viewing Yarbus' own portrait under the task conditions resembling Yarbus' questions with mild modifications. For example questions 4 and 7 were phrased as "Estimate what the person had been doing just before this picture was taken" and "Try to estimate how long this person had been away from home when this picture was taken and why he

had been away”, respectively. They showed that: 1) Yarbus’ findings generalize to a simpler visual stimulus, and 2) instructions influence where and which features an observer inspects in face viewing.

Betz et al. (2010) addressed how or whether at all high-level task information interacts with the bottom-up processing of stimulus-related information. They recorded viewing behavior of 48 observers on web pages for three different tasks: free viewing, content awareness, and information search. They showed that task-dependent differences in their setting were not mediated by a re-weighting of features in the bottom-up hierarchy, ruling out the weak top-down hypothesis. Consequently, they concluded that the strong top-down hypothesis, which proposes that top-down information acts independently of the bottom-up process, is the most viable explanation for their data. These results support Yarbus’ findings in that top-down factors influence where we look when viewing a scene.

Henderson et al. (2013) recorded eye movements of 12 participants while they were engaged in four tasks over 196 scenes and 140 texts: scene search, scene memorization, reading, and pseudo reading. They showed that the viewing tasks were highly distinguishable based on eye-movement features in a four-way classification. They reported a high task decoding accuracy above 80% using multivariate pattern analysis (MVPA) methods widely used in the neuroimaging literature. Their four tasks, however, are much coarser than Yarbus’ original questions, thus making the decoding problem effectively easier. Further, natural scenes and text used by Henderson et al. have dramatically different low-level feature distributions, which causes major differences in eye movement patterns (O’Connell & Walther, 2012; Harel et al., 2008), hence some of the decoding accuracy may be due to stimulus rather than task.

The list of studies addressing task decoding from eye movements and effects of tasks/instructions on fixations is not limited to the above works. Indeed, a large variety of studies have confirmed that eye movements contain rich signatures of the observer’s mental task, including: predicting search target (Rajashekar et al., 2006; Zelinsky & Samaras, 2008; Haji-Abolhassani & Clark, 2013; Zelinsky et al., 2013), predicting stimulus category (Harel et al., 2008; O’Connell & Walther, 2012; Borji et al., 2013b), predicting what number a person may randomly pick (Loetscher et al., 2010), predicting mental abstract tasks (Brandt & Stark, 1997; Mast & Kosslyn, 2002; Meijering et al., 2012; Ferguson & Breheny, 2011), predicting events (Peters & Itti, 2007; Bulling et al., 2011), classifying patients from controls (Tseng et al., 2012; Jones & Klin, 2013), and predicting driver’s intent (Lethaus et al., 2013). Several studies have investigated the role of eye movements in natural vision including: reading (Rayner, 1979; Reichle et al., 2003; Kaakinen & Hyönä, 2010; Clark & O’regan, 1998), visual search (Zelinsky, 2008; Torralba et al., 2006), driving (Land & Lee, 1994; Land & Tatler, 2001), tea making (Land & Rusted, 1999), sandwich making (Hayhoe et al., 2003), arithmetic and geometric problem solving (Epelboim & Suppes, 2001; Cagli et al., 2009), mental imagery (Kosslyn, 1994; Mast & Kosslyn, 2002), cricket (Land & McLeod, 2000), fencing (Hagemann et al., 2010), billiard (Crespi et al., 2012), drawing (Cagli et al., 2009), magic (Kuhn et al., 2008), shape recognition (Renninger et al., 2005), and walking and obstacle avoidance (Mennie et al., 2007).

Departing from the above studies arguing that it is possible to decode observers’ task from fixations (e.g., Iqbal & Bailey (2004); Henderson et al. (2013)) and Greene et al. (2012) have recently cast a shadow on task or mind state decoding by bringing counter examples. They conducted an experiment in which they recorded eye movements of observers when viewing scenes under four highly overlapped questions. Using three pattern classification techniques they were not able to decode the task significantly above chance using aggregate eye movement features (See Figure 4 in Greene et al.’s paper). They were, however, able to decode image and observer’s identity from eye movements above chance level. Task classification failure along with their finding that human judges could not tell the category of scanpath, led Greene et al. to conclude “The famous Yarbus figure may be compelling but, sadly, its message appears to be misleading. Neither humans nor machines can use scan paths to identify the task of the viewer”.

In summary, the effect of task on eye movement patterns has been confirmed by several studies. Despite the volume of attempts at studying task influences on eye movements and attention, fewer attempts have been made to decode observer’s task, especially on complex natural scenes using pattern classification techniques (i.e., the reverse process of task-based fixation prediction). However, there is of course a large body of work examining top-down attentional control and eye movements using simple stimuli and tasks such as visual search arrays and cueing tasks (e.g., Folk & Remington (1998), Folk et al. (1992), Egeth & Yantis

(1997), Yantis (2000), Duncan & Humphreys (1989), Bundesen et al. (2005), Sperling (1960), Sperling & Doshier (1986)). We attempt to thoroughly investigate the task decoding problem by analyzing previous data and findings of Greene *et al.* as well as our own collected data. We focus on Greene et al.’s study because we believe that their experimental design was best suited for task decoding and well in line with Yarbus’ original idea, yet they reported that decoding failed. Further, we study limitations and important factors in task decoding including features and methods used for this purpose. Finally, we discuss potential technological and societal impacts of task and mental state decoding.

## 2. Experiment One

Due to important implications of Greene et al.’s results, here we first reanalyze their data and then summarize learned lessons. They shared data of their third experiment with us which includes fixations of 17 observers viewing 20 grayscale images, each for 60 seconds<sup>2</sup>. They asked observers to view images under four questions: 1) Memorize the picture (memory), 2) Determine the decade in which the picture was taken (decade), 3) Determine how well the people in the picture know each other (people), and 4) Determine the wealth of the people in the picture (wealth). Table 1 shows the arrangement of observers over these tasks. Each observer did all four tasks but over different images. This results in  $17 \times 20 = 340$  scanpaths where each scanpath contains fixations of an observer over one image. The design was intentional to prohibit one observer seeing the same scene twice. Figure 1 demonstrates the stimuli used in this experiment.

images 1 ~ 5	6 ~ 10	11 ~ 15	16 ~ 20
4 O × T 1	4 O × T 2	4 O × T 3	4 O × T 4
4 O × T 2	4 O × T 3	4 O × T 4	...
5 O × T 3	5 O × T 4	...	
4 O × T 4	4 O × T 1		

Table 1: Arrangement of observers over tasks in Greene et al. (2012). O and T stand for observer and task, respectively.

Three factors may have caused task prediction failure in Greene et al.’s study: First and foremost, spatial image information is lost in the type of features they exploited (i.e., using histograms). This is particularly important since the first observation that strikes the mind from Yarbus’ illustration is spatial patterns of fixations<sup>3</sup>. Second, the importance of the classification technique may have been underestimated in Greene et al.’s study. In fact, they only tried linear classifiers (linear discriminant analysis, linear SVM, and correlational methods) and concluded that their failure in task decoding is independent of the classification technique. They made similar arguments for images and features. Third, in Greene et al.’s study, observers were partitioned across images. Thus image and observer idiosyncrasies might have effects on task decoding (Poynter et al., 2013; Risko et al., 2012; Chua et al., 2005). For example, one observer might not have the necessary knowledge regarding a task or an image may not convey sufficient information for answering questions. In what follows, we scrutinize these factors one by one.

Regarding the first factor, we use a simple feature which is the smoothed fixation map, downsampled to  $100 \times 100$  and linearized to a  $1 \times 10000$  D vector (feature type 1). Figure 2.A shows fixation maps for an example image. The fixation map reflects pure eye movement patterns. Additionally, we use histograms of Normalized Scanpath Saliency (NSS) proposed by Peters et al. (2005), using 9 saliency models<sup>4</sup>. This feature

<sup>2</sup>Please see the original paper for more details on the experimental setup. Greene et al., reported 16 observers on their paper (exp 3) but shared 17 with us and on their website <http://stanford.edu/~mrgreene/Publications.html>. Our results and conclusions are valid over selection of 16 subjects distributed equally across tasks.

<sup>3</sup>Greene et al. were able to decode the stimulus from the aggregate features. We suspect that using spatial patterns will lead to much higher accuracies as scanpaths on images are often quite different (e.g., O’Connell & Walther (2012); Harel et al. (2008)).

<sup>4</sup>Selected saliency models include: **AIM** (Bruce & Tsotsos, 2009), **AWS** (Garcia-Diaz et al., 2012), **GBVS** (Harel et al., 2006), **HouCVPR** (Hou & Zhang, 2007), **HouNIPS** (Hou & Zhang, 2008), **ITTI98** (Itti et al., 1998), **PQFT** (Guo & Zhang, 2010), **SEO** (Seo & Milanfar, 2009), and **SUN** (Zhang et al., 2008). For more details on these models, the interested reader is referred to Borji & Itti (2012) and Borji et al. (2012b). Note that *saliency* is not a unique measurement and may change from

reflects the *stimulus + behavior* effect and basically indicates which visual attributes may be important when an observer is viewing an image under a task. NSS values are activations at fixated locations from a saliency map that is normalized to have zero mean and unit standard deviation. For each image, NSS values are calculated and then the histogram of these values (using 70 bins) is considered as features. Thus for 9 models, this leads to a  $9 \times 70 = 630$  dimensional vector (feature type 2). Although our features are aggregates and histograms like Greene et al.’s, one critical difference is that the values which are aggregated reflect a spatial correlation between eye movements and spatial saliency features in each image. Thus, our features capture whether a task may lead an observer to allocate gaze differently over different types of salient image regions.

We also consider the first four features used in Greene et al. (2012) including the number of fixations, the mean fixation duration, the mean saccade amplitude, and the percent of the image area covered by fixations assuming a  $1^\circ$  fovea (feature type 3; dimensionality of 4). In addition, because it has been argued that the first few fixations over a scene may convey more information (Parkhurst et al., 2002), we form a 4th feature type which includes  $\langle x, y \rangle$  locations of the first 5 fixations (i.e., a 10D vector). Note that, in addition to these features, one could think of more complex features (e.g., scanpath sequence, NSS histograms on learned top-down task relevance maps, dwell times on faces, text, and human bodies, and temporal characteristics of fixations (Mills et al., 2011)) to obtain better accuracies. But as we show here, these simple features suffice to decode the task in this particular problem.

Regarding the second factor, we investigate other classification methods such as k-nearest-neighbor; kNN (Fix & Hodges, 1951) and Boosting (Freund & Schapire, 1997; Schapire, 1990) techniques which have been proven to be successful on different problems in machine learning, computer vision, and cognitive sciences<sup>5</sup>. The intuition is that for different problems, different classification methods may perform better. kNN is a classic non-parametric method for classification and regression problems. Given a distance metric (e.g., Euclidean distance), the kNN classifier predicts class label of a test sample as the majority vote of its k closest training examples in the feature space (i.e., the most common output among the neighbors). If  $k = 1$ , then the class label of the test sample is the same as its nearest neighbor. We also tried Boosting algorithms which are popular and powerful machine learning tools nowadays. The basic idea underlying Boosting algorithms is learning several weak classifiers (i.e., a classifier that works slightly better than chance) and combing their outputs to form a strong classifier (i.e., a meta-algorithm). The learning is done in an iterative manner. After adding a weak learner, the data is re-weighted to emphasize mistakes. Misclassified exemplars gain higher weight while correctly-classified exemplars lose weight. Here, we employ the RUSBoost (Random Under-Sampling Boost) algorithm (Seiffert et al., 2010) which uses a hybrid sampling/boosting strategy to handle class imbalance problem in data with discrete class labels. To better model the minority class, this algorithm randomly removes examples from the majority class until all classes have balanced number of examples (i.e., under-sampling). Due to the random sampling, different runs of this algorithm may yield different results. While the class imbalance (only one task has 5 subjects) is not a big issue in our data, we believe it is the ensemble of weak classifiers (here decision trees) that makes good prediction possible<sup>6</sup>.

With respect to the third factor, we conduct the following two analyses: 1) pooling data from all observers over all images and tasks (i.e.,  $17 \times 20$  scanpaths), and 2) treating each image separately. These analyses help disentangle the effects of image and observer parameters on task decoding.

### 2.1. Task decoding over all data

We trained multi-class classifiers to recover task (one out of four possible) from eye movement patterns. We follow a leave-one-out cross validation procedure similar to Greene et al. (2012). Each time we set one datapoint aside and train a classifier over the rest of data. The trained classifier is then applied to the set-aside datapoint. We repeat the same procedure over all 340 datapoints and report the average accuracy (i.e., 340 binary values). Decoding results are shown in Fig. 2.B.

---

one model to another. That is why here we employ several models instead of one.

<sup>5</sup>Boosting classifiers have been used for fixation prediction in free viewing tasks (e.g., Borji (2012) and Zhao & Koch (2012)).

<sup>6</sup>Please see Matlab documentation for `fitensemble` function.



Figure 1: Stimuli used in experiment 1. Easy and difficult scenes for task decoding are marked with blue and red boxes, respectively. Please see Table 3 (Appendix I) for performances of individual runs of the RUSBoost classifier. Average decoding accuracies (numbers after dash lines) are using feature type 3 over 50 RUSBoost runs. Numbers in brackets are classification accuracy using feature type 1 (over 50 RUSBoost runs). Original images are  $800 \times 600$  pixels.

Using kNN and feature type 1 (i.e., fixation map), we achieved average accuracy of 0.2412 ( $k=2$ ; binomial test,  $p = 0.89$ ). Feature type 2 leads to accuracy of 0.2353 ( $k=2$ ;  $p = 0.31$ ). Using feature type 3, we achieved accuracy of 0.3118 which is above Greene et al.’s results and is significantly better than chance ( $k=8$ ;  $p = 0.014$ ). Classification with feature type 4 leads to accuracy of 0.2441 ( $k=1$ ). Combination of features did not improve the results significantly. Fig. 2.B (bottom panel) shows kNN performance as a function of number of neighbors ( $k=2^n$ ,  $n = 0 \dots 6$ ). kNN classification performance levels here are for the best-performing value of  $k$ .

Using RUSBoost classifier with 50 boosting iterations and feature type 1, we achieved accuracy of 0.25 (non-significant vs. chance; binomial test,  $p = 0.6193$ ). We achieved accuracy of 0.2294 using feature type 2. Feature type 3 leads to accuracy of 0.3412 ( $p = 1.0722e - 04$ ). Finally, feature type 4 results in accuracy of 0.2176. Combination of all features did not increase the results significantly (accuracy of 0.3412 using all features).

**Bonferroni correction for multiple comparisons (Shaffer, 1995):** Here, we used two classifiers, 5 feature types (including combination of features), and 7 values of parameter  $k$  in kNN resulting in  $5 \times 7 + 5 = 40$  tries. We need to correct  $p$ -values for these comparisons. Thus, significance  $p$  is equal to  $0.05/40 = 0.0013$ . Using kNN, the best  $p$  value is 0.014 which is above the corrected significance level, therefore kNN does not yield statistically significant decoding accuracy that is strong enough given that we made 40 different attempts at decoding. Hence, we discard using kNN for the rest of the paper. Using the RUSBoost classifier (with feature type 3), however, results remain statistically significant after correction as  $p$  values are smaller than 0.0013 which indicates that task is decodable on this data significantly above the 25% chance level.

Results of this analysis indicate that spatial fixation patterns are not informative regarding the observer’s task when pooling all data (on Greene et al.’s data). Further, our results show that classification method is a key factor. For example, using the same 4 features employed by Greene et al. (feature type 3), we achieved better accuracies with kNN and Boosting classifiers. Note that, here we did not conduct an exhaustive search to find the best features or feature combinations. It might be possible to reach even higher accuracies with more elaborate feature selection strategies.

## 2.2. Task decoding over single images

Task decoding accuracy highly depends on the stimulus set. For example, if an image does not have the necessary content that is called for by different tasks (in an extreme case, a blank image and tasks about age or wealth of people), it may not yield task-dependent eye movement patterns as strong as an image that has such content. That is, we expect that interaction between semantic image contents and task may give rise to the strongest eye movement signatures. Failure to decode task might thus be more likely if the stimuli do not support executing the task. This is particularly important since both Yarbus and Greene et. al., did not probe observers’ responses to see whether or not they were actually able to perform the task.

We train a RUSBoost classifier (with 50 boosting iterations) on 16 observers over each individual image and apply the trained classifier to the remaining observer over the same image (i.e., leave one observer out). We repeat this process for all 20 images. Using feature type 1, we achieve average accuracy of 0.3267 (over 50 runs and images). Feature type 3 resulted in accuracy of 0.3414 (See Appendix I (Table 3) for results of 50 runs). The maximum performance using this feature over runs was 0.3719 and the minimum was 0.3156. Using combination of all features (a feature vector of size  $10000 + 9 \times 70 + 4 + 10 = 10644D$ ) results in average accuracy of 0.3294. Examination of confusion matrices using RUSBoost and feature type 3 (Figure 2.C) shows above chance performance on diagonal elements with higher accuracies for memory and decade tasks. There is high confusion between wealth and other classes.

Average task decoding performance per image using feature type 3 is illustrated in Figure 2.C as well as in Figure 1. Using this feature, decoding accuracy is significantly above chance level for majority of the images, is non-significant vs. chance for one image, and is significantly below chance for three images (using t-test over 50 runs; see Appendix I). The easiest and most difficult stimuli using feature type 1 along with their scanpaths and confusion matrices (using a sample run of RUSBoost) are shown in Fig. 3.

Results of the second analysis support our argument that image content is an important factor in task decoding. Task decoding becomes very difficult if an image lacks diagnostic information relevant to the task

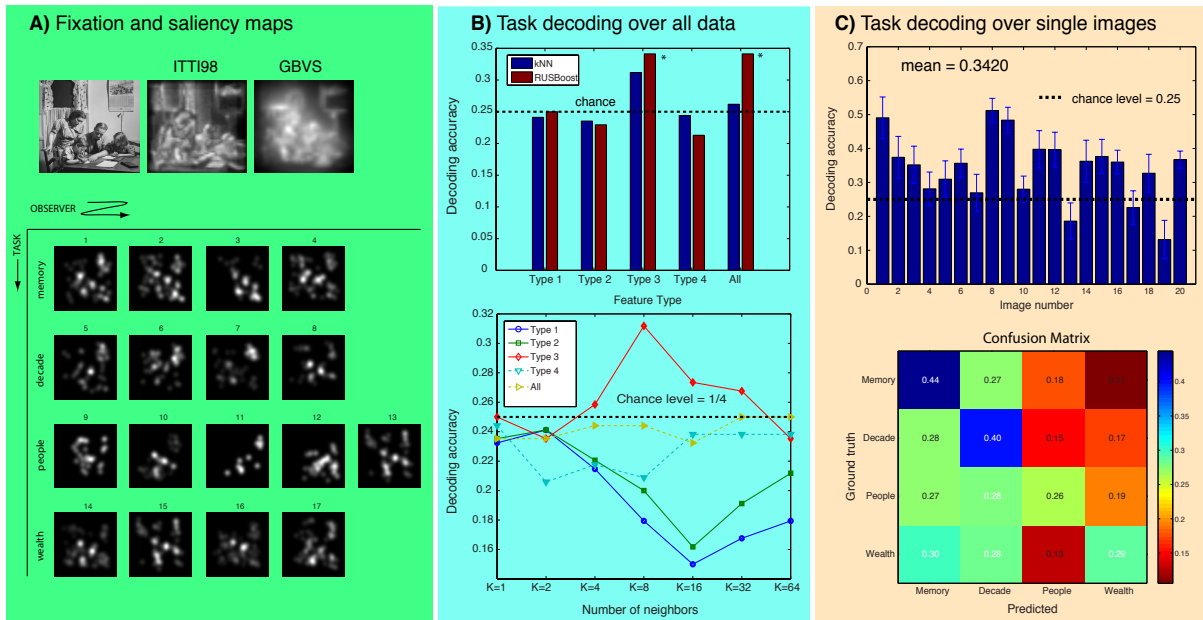


Figure 2: Results of experiment 1: A) Top: A sample image along with saliency maps using ITTI98 and GBVS models and its corresponding smoothed fixation maps (using Gaussian sigma 33 subtending about  $0.85^\circ \times 0.85^\circ$  of visual angle). Matlab code for generating the smoothed fixation map: `imresize(conv2(map, fspecial('gaussian', 200, 33)), [100 100], 'nearest')`. Numbers on top of fixation maps in the bottom panel show the observer's number (See Table 1). B) Top: Task decoding accuracy using individual features and their combination over all data. Stars indicate statistical significance versus chance using binomial test. Bottom: Effect of number of kNN neighbors on task decoding accuracy. C) Top: Average decoding accuracies over 50 runs of the RUSBoost classifier over individual images using feature type 3 (See Appendix I). Error bars indicate standard deviations over 50 runs. Bottom: Average confusion matrix (over 50 RUSBoost runs) averaged over all images.



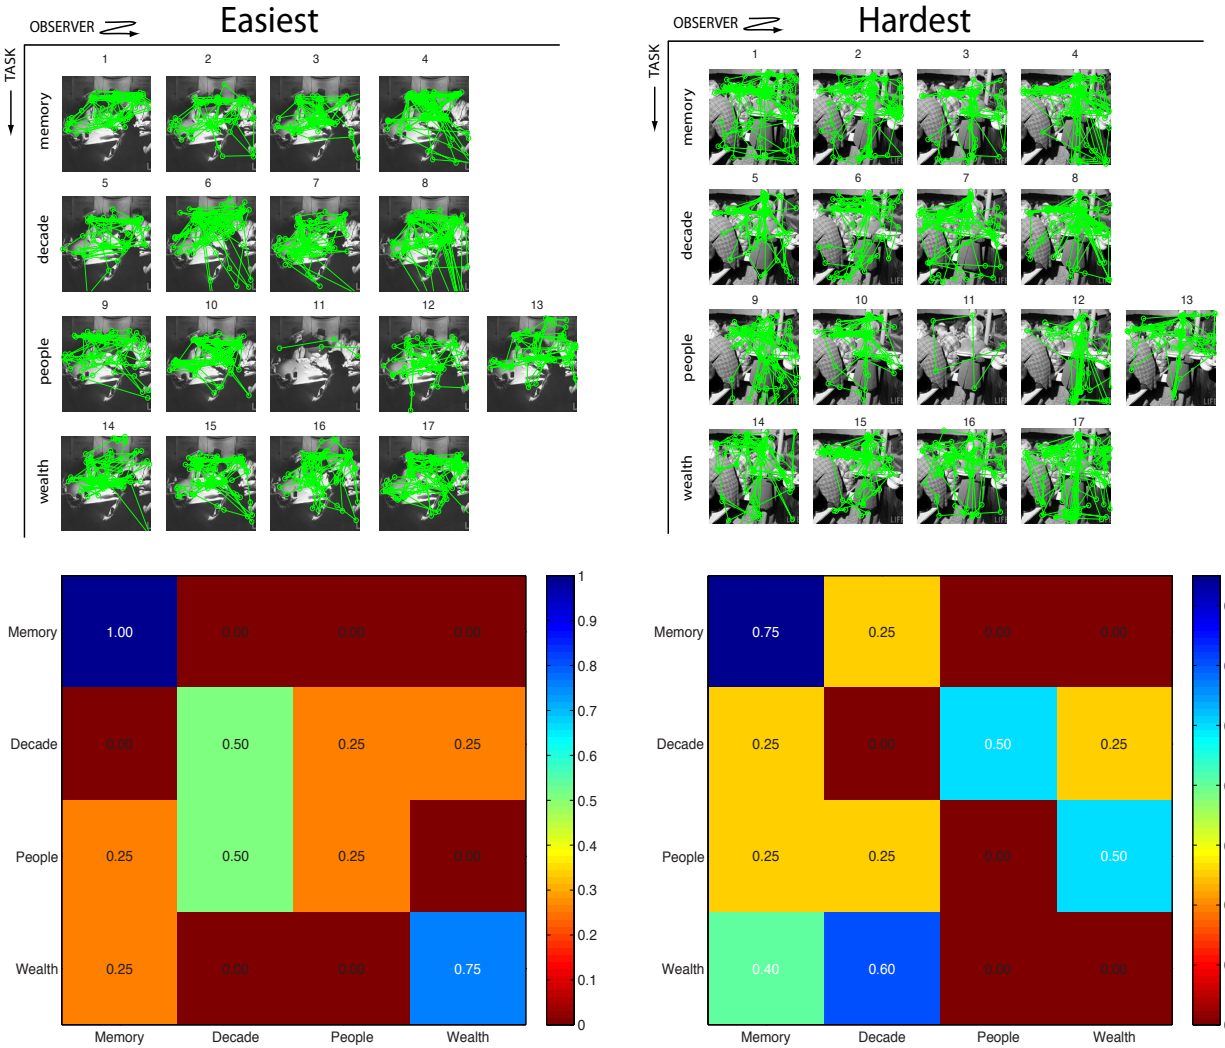


Figure 3: Easiest and hardest stimuli for task decoding in experiment 1 using feature type 1 over 50 RUSBoost runs. Confusion matrices are for a sample run of RUSBoost on each image using leave-one-out procedure.

(See Figure 3 for such an example). Further, when treating each image separately, scanpaths turn to become informative regarding the task for some images. Overall, results from this experiment suggest that it is possible to decode the task above chance from the same type of features used by Greene et al. (2012).

### 3. Experiment Two

The questions in the task set of Greene et al. are very similar to each other. For decade, memory and wealth tasks, observers have to look over the entire image to gain useful information. This causes very similar fixation patterns such that these patterns are not really differentiable to the naked eye (see Figure 3) or to classifiers, while, to some extent, clearly different patterns may have contributed to making Yarbus’ illustration so compelling. Here we aim to decode observer’s task from eye movements with particular emphasis on spatial fixation patterns (i.e., feature type 1) rather than aggregate features (type 3). While mean values of eye movement measures (i.e., feature type 3) can change as a function of task, the distributions of these values highly overlap across tasks (Henderson et al., 2013).

In our view an important limitation of Greene et al.’s study is that they did not use Yarbus’ original 7 tasks, as Yarbus might have reached different conclusions had he used different tasks. In this experiment, we thus seek to test the accuracy of Yarbus’ exact idea by replicating his tasks.

#### 3.1. Methods

#### 3.2. Participants

A total of 21 students (10 male, 11 female) from the University of Southern California (USC) participated. Students’ majors were Computer sciences, Neuroscience, Psychology, Mathematics, Cognitive sciences, Communication, Health, Biology, Sociology, Business, and Public relations. The experimental methods were approved by the USC’s Institutional Review Board (IRB). Observers had normal or corrected-to-normal vision and were compensated by course credits. Observers were in the age range between 19 to 24 (mean = 22.2, std = 2.6). They were naïve with respect to the purpose of the experiment.

#### 3.3. Apparatus

Participants sat 130 cm away from a 42 inch monitor screen so that scenes subtended approximately  $43^\circ \times 25^\circ$  of visual angle. A chin/head rest was used to minimize head movements. Stimuli were presented at 60Hz at resolution of  $1920 \times 1080$  pixels. Eye movements were recorded via an SR Research Eyelink eye tracker (spatial resolution  $0.5^\circ$ ) sampling at 1000 Hz.

#### 3.4. Materials

Stimuli consisted of 15 paintings (13 are oil on canvas, some are by I. E. Repin). Figure 4 shows stimuli including Repin’s painting used by Yarbus. We chose images such that a person<sup>7</sup> who could be construed as an unexpected visitor exists in all of them. Thus Yarbus’ questions are applicable to these images (e.g., more so on images 2, 3, and 11 and less so on images 4, 6, and 15).

#### 3.5. Procedure

We followed a partitioned experimental procedure similar to Greene et al., where observers answered questions on three sets of images (Table 2). Each set consists of five images corresponding to one row of Figure 4. In other words, no participant saw the same stimulus twice. Each image was shown for 30 seconds followed by a 5 seconds delay (gray screen). At the beginning of each session (5 images), the eye tracker was re-calibrated. Each observer viewed each set of 5 images only under one question. We used the seven questions of Yarbus’ study mentioned in the Introduction. Figure 5 illustrates eye movements of observers on 7 images.

---

<sup>7</sup>Or the dog in image number 8 in Figure 4.



Figure 4: Stimuli used in experiment 2. Images resemble Repin’s painting (image number 5) in that in all of the images exists a somewhat unexpected visitor. (Source: Courtesy of [www.ilyarepin.org](http://www.ilyarepin.org)). Three easiest and three most difficult stimuli are marked with blue and red boxes, respectively. Average decoding accuracies (numbers after dash lines) are using combination of feature types 1 and 2 over all RUSBoost runs. See Table 4 (Appendix II) for decoding results on individual RUSBoost runs. Numbers in brackets are classification accuracy using feature type 1.

### 3.6. Decoding Results

We employ the RUSBoost classifier with 50 boosting iterations as in the first experiment. Features consist of saliency maps of 9 models used in experiment 1 (footnote 4) plus additional 14 feature channels from the ITTI model including: ITTI-C, ITTI-CIO, ITTI-CIOLTXE, ITTI-E, ITTI-Entropy, ITTI-I, ITTI, ITTI-L, ITTI-O, ITTI-OLTXE, ITTI-Scorr, ITTI-T, ITTI-Variance, and ITTI-X. These feature channels extract different types of features that range from Intensity (I), Color (C), Orientation (O), Entropy (E), Variance, T-junctions (T), X-junctions (X), L-junctions (L), and Spatial Correlation (Scorr). Please see Itti et al. (1998) and Tseng et al. (2012) (and its supplement) for more details on these features and implementation details. ITTI and ITTI98 are different versions of the Itti *et al.* model, corresponding to different normalization schemes. In ITTI98, each feature map’s contribution to the saliency map is weighted by the squared difference between the globally most active location and the average activity of all other local maxima in the feature

images 1 ~ 5	6 ~ 10	11 ~ 15
3 O × T 1	3 O × T 2	3 O × T 3
3 O × T 2	3 O × T 3	3 O × T 4
3 O × T 3	3 O × T 4	3 O × T 5
3 O × T 4	3 O × T 5	...
3 O × T 5	...	...
3 O × T 6	...	...
3 O × T 7	3 O × T 1	3 O × T 2

Table 2: Arrangement of observers over tasks in experiment II. O and T stand for observer and task, respectively.



Figure 5: Eye movements of observers over stimuli in experiment 2 for 7 images. Note that each image was shown to an observer only under one question. Tasks are: 1) Free examination, 2) Give material circumstances (Wealth), 3) Estimate ages of the people, 4) Estimate the activity before the arrival of the visitor, 5) Remember cloths, 6) Remember positions of people and objects, and 7) Estimate how long the visitor had been away.

map (Itti et al., 1998). This gives rise to smooth saliency maps, which tend to correlate better with noisy human eye movement data. In the ITTI model (Itti & Koch, 2000), the spatial competition for saliency is much stronger which gives rise to much sparser saliency maps. Figure 6.A shows 23 saliency maps for a sample image.

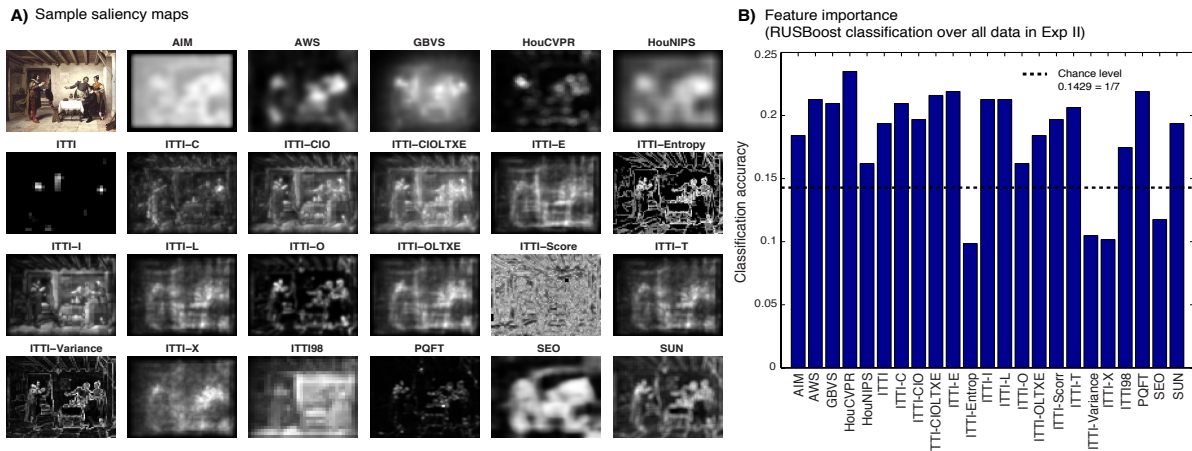


Figure 6: A) Saliency maps for a sample image used in the second experiment. Acronyms are: Intensity (I), Color (C), Orientation (O), Entropy (E), Variance, T-junctions (T), X-junctions (X), L-junctions (L), and Spatial Correlation (Scorr). B) Importance of saliency maps (feature type 2 using 70D NSS histograms) for task decoding. Here, a RUSBoost classifier (50 runs) was used over all data according to the analysis in section 3.6.1).

### 3.6.1. Task decoding over all data

Following the experiment one, we first pool all data and perform task decoding over all images and observers. We report results using a leave-one-out procedure. We have 21 observers, each viewing 15 images (each 5 images under a different question; 3 questions per observer) thus resulting in 315 scanpaths. Using feature type 1, we achieved average accuracy of 0.2421 which is significantly above chance<sup>8</sup> (binomial test,  $p = 2.4535e - 06$ ). Using feature type 2 (i.e., NSS histogram of 9 saliency models as in Exp 1) results in accuracy of 0.2254 ( $p = 5.6044e - 05$ ). Increasing the number of saliency models to 23 results in the same performance as when using 9 models. Combination of all features did not improve the results in this analysis. To evaluate the importance of saliency maps as features, we performed task decoding over all data using individual saliency features (i.e., NSS values; feature type 2) and RUSBoost classification (See section 3.6.1). Results are shown in Figure 6.B. Majority of the saliency models lead to above chance accuracy indicating informativeness of NSS histograms and low-level image features for task decoding.

**Bonferroni correction for multiple comparisons:** With the RUSBoost classifier, correcting for three features and their combination,  $p$  values have to be smaller than  $0.05/4 = 0.0125$  which is the case here using all feature types. Thus, we can safely conclude that task is decodeable from eye movements on our data using spatial fixation patterns and NSS histograms (as opposed to experiment 1).

### 3.6.2. Task decoding over single images

Three observers viewed each image under one question thus resulting in 21 datapoints per image (i.e., 3 observers  $\times$  7 questions). Note that each set of three observers were assigned the same question (Table 2). RUSBoost classifier and feature type 1 results in average accuracy of 0.2724 over 50 runs and 15 images. Using first two feature types (a  $10000 + 23 \times 70 = 11610$ D vector) results in average performance of 0.2743. Over all runs (i.e., table rows), the minimum accuracy (average over all 15 images) is 0.2540 and maximum

<sup>8</sup>We obtained accuracy of  $0.2399 \pm 0.0016$  (mean + std) over 60 runs of the RUSBoost classifier.

accuracy is 0.3079. Note that our accuracies are almost two times higher than the 14.29% chance level (i.e., 1/7). Easy and difficult stimuli for task decoding are shown in Figure 4. See Table. 4 (Appendix II) for results of individual runs of the RUSBoost classifier over individual images. The easiest and hardest stimuli using feature type 1 along with their scanpaths and confusion matrices are shown in Fig. 8.

To measure the degree to which tasks differ from each other, we show in Figure 7.A the distribution of fixations over all images with the same task. Each element shows the amount of overlap in two questions. To generate this plot, we first normalize each map to  $[0\ 1]$  and then subtract maps from each other. Hence brighter blue and red regions mean higher difference between two tasks. It shows profound differences among tasks 3 (estimating ages), 4 (estimating activity) and 7 (estimating away time) to other tasks. Task 1 (free examination) is more similar to other tasks. The reason might be because people look everywhere in images including faces and people which are also informative objects for other tasks. Task 2 (estimating wealth) and task 6 (remembering positions) show smaller difference to other tasks probably because observers inspect the entire image in two tasks. Figure 7.B shows the confusion matrix averaged over 15 images and 50 RUSBoost runs using feature type 1. We observe high accuracies for task 3 (estimating age), task 5 (remembering cloths), and task 7 (estimating how long the visitor has been away) but low accuracy for the free-viewing task. There is a high confusion between task 2 and tasks 6 and 1 and also between task 1 and task 7.

Results of the two analyses in second experiment, in alignment with DeAngelus & Pelz (2009), confirm that eye movements are modulated top-down by task demands in a way that task can be predicted from eye movement patterns. We found that spatial fixation patterns, which were not much informative over Greene et al.’s data, suffice to decode the task on our data. We expect to gain even higher task decoding accuracies by using other eye movement statistics, such as fixation durations or amplitudes which have been shown to be different across Yarbus’ questions (DeAngelus & Pelz, 2009).

#### 4. Discussion and Conclusion

What do we learn from the two experiments in this study? Successful task decoding results provide further evidence that fixations convey diagnostic information regarding the observer’s mental state and task<sup>9</sup>, consistent with the cognitive relevance theory of attention (See Hayhoe & Ballard (2005)). This means that top-down factors in complex tasks systematically influence the viewer’s cognitive state and his thought processes. Our results support previous decoding findings mentioned in the introduction section (e.g., some over more controlled stimuli such as predicting search target (Rajashekar et al., 2006)).

We demonstrated that it is possible to reliably infer the observer’s task from Greene et al.’s data using stronger classifiers. Classification was better when we treated images individually. Although we were able to decode the task from Greene et al.’s data, making strong arguments regarding feasibility of task decoding on this data is difficult mainly due to the small size of this dataset. We think to gain better insights, larger datasets for task decoding are necessary. Such datasets allow break down of data into (larger) separate train and test sets. Parameters of a classifier can be optimized using the train set and the resultant classifier can be evaluated on the test set. Performing the analysis in this manner eliminates the need for correction for multiple comparisons, hence allowing one to try possibly thousands of possible classifiers and parameters.

In the second experiment, we showed that it is possible to decode the task using Yarbus’ original tasks, almost twice above chance, much better than using Greene et al.’s tasks. These results are in line with findings of DeAngelus & Pelz (2009). While our results are significantly above chance, it might be still possible to obtain better accuracies by exploiting even more informative features and other types of classification techniques. Our investigation on task decoding using 5-second time slots (i.e., 1st 5-seconds, 2nd 5-seconds, ...) suggest that accuracies might be higher for early fixations but this needs further investigation. We also found that decoding accuracy critically depends on three factors: 1) task set (how separable they are),

---

<sup>9</sup>Note that here we used task and cognitive state interchangeably. There are however subtle differences. Cognitive state refers to the state of a person’s psychological condition (e.g., confusion, preoccupation, wonder, etc.). By task here referred to a well-defined question that observers should try to answer (e.g., estimating age, search for an object, reading, etc.).

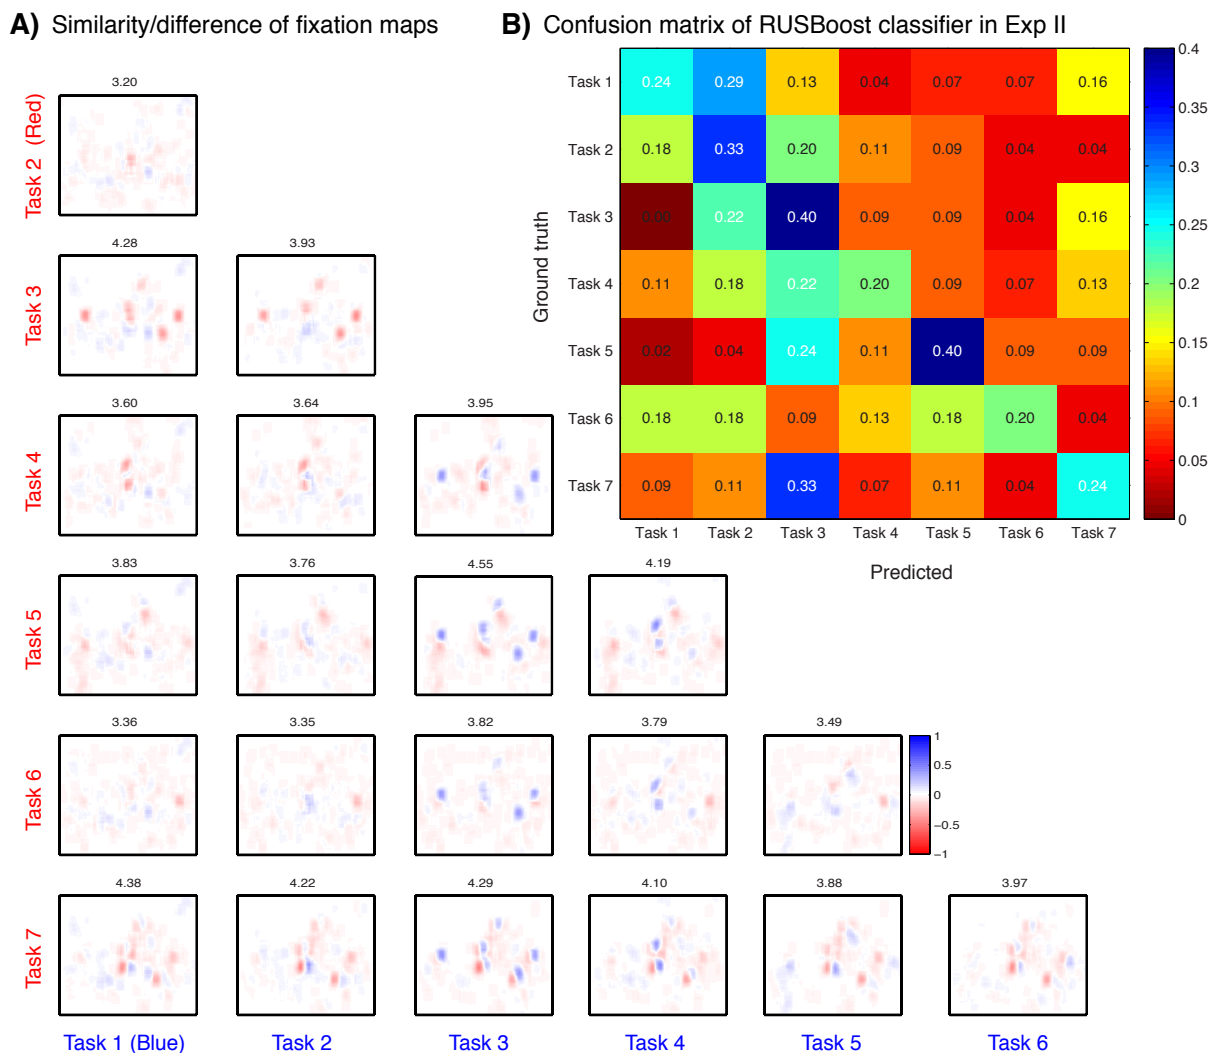


Figure 7: A) Similarity/difference of tasks from human fixation maps in experiment 2. Brighter red or blue regions mean higher difference. Values close to zero mean less difference. The numbers on top of each image show the sum of the absolute differences between two fixation maps. B) Confusion matrix of the RUSBoost classifier averaged over 50 RUSBoost runs each on a single image using feature type 1.

2) stimulus set (whether a scene has sufficient information or not), and 3) observer’s knowledge (whether observers understand questions).

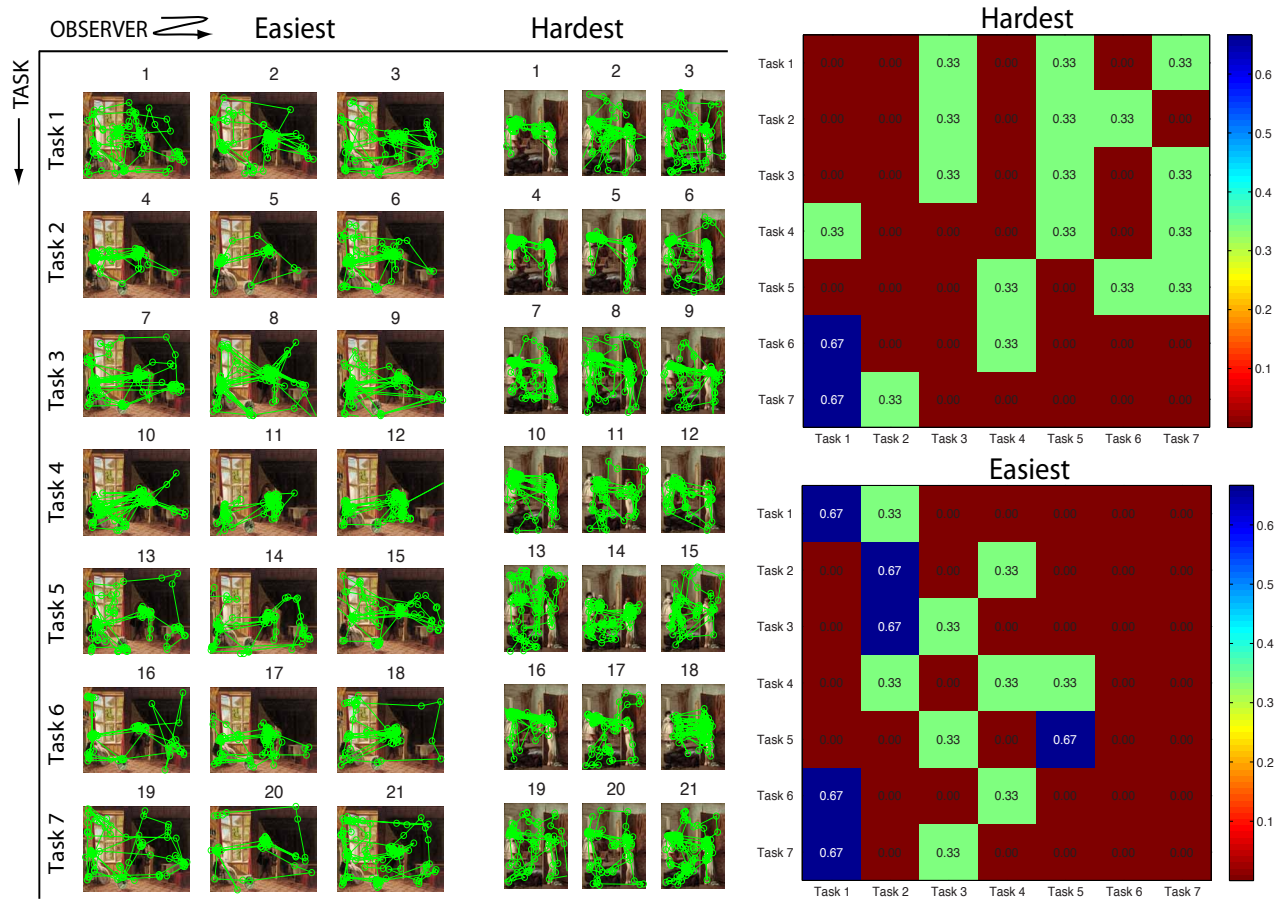


Figure 8: Easiest and hardest stimuli for task decoding in experiment 2 using feature type 1 over 50 RUSBoost runs. Confusion matrices are for a sample run of RUSBoost on each image using leave-one-out procedure.

Is it always possible to decode task from eye movements? We argue that there is no general answer to this type of pattern recognition questions. Answers depend on the used stimuli, observers, and questions. One could choose tasks such that decoding becomes very hard even with sophisticated features and classifiers; we found that this is the case on Greene et al.’s data. In particular, on the type of tasks and scenes used here majority of fixations are attracted to faces and people which causes a huge overlap across tasks. In some easier scenarios, where tasks are more different, very simple features might suffice to decode the task accurately (e.g., Henderson et al. (2013)). In the extreme simplest case, one can imagine a task like this: a person on the left side of the screen and a dog on the right side, with observers’ tasks being: 1) How old is the person? and 2) What breed is the dog? Obviously answering these tasks demands looking at the person for the first question and looking at the dog for the second question which results in 100% task decoding accuracy (for a rational observer) just from eye movement locations. One can also choose images from which task decoding is very difficult because they contain little information that is directly relevant to the task. This was also found in our results, as some images yield more accurate task decoding than others. One could also recruit observers who don’t understand the question. So far none of the works mentioned in the present study have analyzed the observers’ answers on tasks. So, the failure in task decoding might be simply due to the observer’s disability to extract useful information from the scene.

Since the parameter space is large, making strong arguments regarding impossibility of task decoding



(see, e.g., Greene et al.’s claim “static scan paths alone do not appear to be adequate to infer complex mental states of an observer” in their abstract) seems to be very difficult and needs a systematic probing of the whole parameter space (or a theoretical proof). On the other hand, to prove that task decoding on a particular setting is feasible, one only needs to find a working set of parameters (and it suffices; after accounting for multiple comparisons and following a cross-validation procedure (Salzberg, 1997)). The latter is the common practice in pattern recognition community. Please note that our results also do not imply that it is always possible to decode the task. The counter example proposed by Greene et al. was found to not hold in our analysis.

As a control analysis, Greene et al., asked some human participants to look at eye movements of their observers and guess which task the observers have been doing. They showed that similar to classifiers, participants also failed in task decoding. Failure of humans to decode the task by looking at eye movement patterns (experiment 4 in Greene et al.) does not necessarily mean that fixations lack task-relevant information. Indeed, there are some cases in vision sciences where machine learning techniques outperform humans, in particular over large datasets (e.g., frontal face recognition, defect detection, cell type differentiation, DNA microarray analysis, etc.).

Several concerns exist that need to be carefully thought about before conducting a task decoding experiment using eye movements. Here we followed the procedure by Greene et al., in which: 1) no observer viewed the same image twice, and 2) the same scene was shown under multiple questions. The first rule aims to eliminate memory biases. The second rule ensures that the final result is not due to differences in stimuli. Yarbus (1967) and DeAngelus & Pelz (2009) violated the first rule where the same observers viewed the images under the same questions. Henderson et al. (2013) violated the second rule in which different questions were asked over different images (which might be the reason why they obtained such high accuracies above 80%). Another possibly important factor affecting task decoding results is eye tracking accuracy. This is particularly important when tasks are very similar to each other. One other concern regards selection of the stimulus set. If the stimulus set includes many images containing people, faces, and text, which capture a large portion of fixations in a task independent manner, then there is basically not much information left helping task decoding. The last concern is about the suitability of features. In some scenarios, especially in dynamic environments (e.g., watching a video, driving a car, etc.) the type of features employed here may not be suitable for task decoding. In particular, spatial information is reduced to one fixation per frame. This requires temporal processing of features to see which places (or in what order) observers have visited the locations.

Here, we showed that task is decode-able on static images by a more systematic and exhaustive exploration of the parameter space including features, classifiers, and new data. Pushing deeper into real-time scenarios, using joint online analysis of video and eye movements, we have recently been able to predict – more than one second in advance – when a player is about to pull the trigger in a flight combat game, or to shift gears in a car racing game (Peters & Itti, 2007). We have been also able to predict next fixation of a video game player for such games as running a hot-dog stand (Borji et al., 2013a) and Super Mario Cart (Borji et al., 2012a). In a similar approach where our computational models provide a normative gold standard against one particular individual’s gaze behavior, we have demonstrated a system which can predict, by recording an observer’s gaze for 15 minutes while one watches TV, whether one has ADHD (Tseng et al., 2012). These preliminary results clearly demonstrate how computational attention models can be used jointly with behavioral recordings to infer some internal state of a person, from a short-term intention (e.g., pull the trigger) to long-term characteristics (e.g., like ADHD).

Beyond scientific value, decoding task from eye movements has practical applications. Potential technological applications include: wearable visual technologies (smart glasses like Google Glass), smart displays, adaptive web search, marketing, activity recognition (Fathi et al., 2011; Pirsiavash & Ramanan, 2012; Albert et al., 2012), human-computer interaction, and biometrics. Portable electronic devices such as smart-phones, tablets, smart glasses with cameras are becoming increasingly popular (See Windau & Itti (2013) for an example study). Enabling eye tracking on these devices could be used to predict the user’s intent one step ahead and provide him necessary information in a more efficient and adaptive manner. This could be augmented with approaches that use non-visual information on cell phones such as accelerometer data or global

positioning systems (e.g., Albert et al. (2012)). Another area of applicability is assistant systems especially for elderly and disabled users (e.g., in driving or other daily life activities (Bulling et al., 2011; Doshi & Trivedi, 2012, 2009)). Here, we focused on predicting observer’s task. Some studies have utilized eye movements to tap into mental states such as confusion and concentration (Victor et al., 2005; Griffiths et al., 1984), arousal (Woods et al., 1978; Subramanian et al., 2010), or deception (Kuhn & Tatler, 2005). Eye movements can also be utilized as a measure of learning capacity in category learning and feature learning (e.g., Rehder & Hoffman (2005); Chen et al. (2013)) and expertise (e.g., Vogt & Magnussen (2007); Bertram et al. (2013); Jarodzka et al. (2010)).

From a societal point of view, reliable fixation-based task decoding methods could be very rewarding. One area of application is patient diagnosis. Several high-prevalence neurological disorders involve dysfunctions of oculomotor control and attention, including Autism Spectrum Disorder (ASD), Attention Deficit Hyperactivity Disorder (ADHD), Fetal Alcohol Spectrum Disorder (FASD), Parkinson’s disease (PD), and Alzheimer<sup>10</sup>. Diagnosis and treatment of these disorders are becoming a pressing issue in today’s society (See (Klin et al., 2009; Jones & Klin, 2013)). For example, about 1 in 6 children in the U.S. had a developmental disability, such as intellectual disabilities, cerebral palsy, and autism, in 2006-2008. Reliable and early diagnosis of these disorders boils down to accessing observers’ internal thought processes and their cognitive states. This is where our task decoding framework becomes relevant and could potentially replace or complement existing clinical neurological evaluation, structured behavioral tasks, and neuroimaging technique which are currently expensive and time-consuming. We believe that the type of methods discussed here along with low-cost non-invasive eye tracking facilities offer considerable promise for patient screening. However, to make it happen in the future, high-throughput and robust task-decoding methods need to be devised. One direction could be augmenting eye movements with purely biological cues such as pupils, sweating, heart rate, and breathing for this purpose.

## Acknowledgments

This work was supported by the National Science Foundation (grant number CMMI-1235539), the Army Research Office (W911NF-11-1-0046 and W911NF-12-1-0433), and U.S. Army (W81XWH-10-2-0076). The authors affirm that the views expressed herein are solely their own, and do not represent the views of the United States government or any agency thereof. Authors would like to thank Michelle R. Greene and Jeremy Wolf for sharing their data with us. We also thank Dicky N. Sihite for his help on parsing the eye movement data. Our code and data is publicly available at <http://ilab.usc.edu/borji/Resources.html>.

Commercial relationships: none.

Corresponding author: Ali Borji.

Email: [borji@usc.edu](mailto:borji@usc.edu).

Address: University of Southern California, Hedco Neuroscience Building, 3641 Watt Way, Los Angeles, CA 90089-2520, USA

## References

- Albert, M., Toledo, S., Shapiro, M., & Kording, K. (2012). Using mobile phones for activity recognition in parkinson’s patients. *Frontiers in Neurology.*, 3.
- Ballard, D., Hayhoe, M., & Pelz, J. (1995). Memory representations in natural tasks. *Journal of Cognitive Neuroscience.*, 7, 66–80.
- Bertram, R., Helle, L., Kaakinen, J. K., & Svedström, E. (2013). The effect of expertise on eye movement behaviour in medical image perception. *Plos One.*, 8.

---

<sup>10</sup>Visit here for the prevalence statistics: <http://www.cdc.gov/ncbddd/>

- Betz, T., Kietzmann, T., Wilming, N., & König, P. (2010). Investigating task-dependent top-down effects on overt visual attention. *Journal of Vision.*, *10*.
- Borji, A. (2012). Boosting bottom-up and top-down visual features for saliency estimation. *CVPR 2012*, .
- Borji, A., & Itti, L. (2012). State-of-the-art in modeling visual attention. *IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI)*, *35*, 185–207.
- Borji, A., Sihite, D. N., & Itti, L. (2012a). Probabilistic learning of task-specific visual attention. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, .
- Borji, A., Sihite, D. N., & Itti, L. (2012b). Quantitative analysis of human-model agreement in visual saliency modeling: A comparative study. *IEEE Trans. Image Processing.*, *22*, 55–69.
- Borji, A., Sihite, D. N., & Itti, L. (2013a). What/where to look next? modeling top-down visual attention in complex interactive environments. *IEEE Transactions on Systems, Man, and Cybernetics, PART A-SYSTEMS AND HUMANS.*, .
- Borji, A., Tavakoli, H. R., Sihite, D. N., & Itti, L. (2013b). Analysis of scores, datasets, and models in visual saliency prediction. *ICCV*, .
- Brandt, S. A., & Stark, L. W. (1997). Spontaneous eye movements during visual imagery reflect the content of the visual scene. *Journal of Cognitive Neuroscience*, *9*, 27–38.
- Bruce, N., & Tsotsos, J. (2009). Saliency, attention, and visual search: An information theoretic approach. *Journal of Vision.*, *9*.
- Bulling, A., Ward, J. A., Gellersen, H., & Tröster, G. (2011). Eye Movement Analysis for Activity Recognition Using Electrooculography. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *33*, 741–753.
- Bundesen, C., Habekost, T., & Kyllingsbæk, S. (2005). A neural theory of visual attention: bridging cognition and neurophysiology. *Psychological review*, *112*, 291.
- Buswell, G. (1935). *How people look at pictures..* Chicago.: University of Chicago Press.
- Cagli, R. C., Coraggio, P., Napoletano, P., Schwartz, O., Ferraro, M., & Boccignone, G. (2009). Visuomotor characterization of eye movements in a drawing task. *Vision Research*, *49*.
- Castelhano, M., Mack, M., & Henderson, J. (2009). Viewing task influences eye movement control during active scene perception. *Journal of Vision*, *9*, 1–15.
- Cerf, M., Frady, E. P., & Koch, C. (2009). Faces and text attract gaze independent of the task: Experimental data and computer model. *Journal of Vision*, *9*.
- Chen, L., K.M., M., Blair, M., Watson, M., & Wood, M. (2013). Temporal characteristics of overt attentional behaviour during category learning. *Attention Perception & Psychophysics*, *19*, 244–256.
- Chua, H. F., Boland, J. E., & Nisbett, R. E. (2005). Cultural variation in eye movements during scene perception. *Proceedings of the National Academy of Sciences of the United States of America*, *102*, 12629–12633.
- Clark, J. J., & O’regan, J. K. (1998). Word ambiguity and the optimal viewing position in reading. *Vision Research*, *39*, 843–857.
- Crespi, S., Robino, C., O., S., & deSperati C. (2012). Spotting expertise in the eyes: billiards knowledge as revealed by gaze shifts in a dynamic visual prediction task. *Journal of Vision.*, *12*.
- DeAngelus, M., & Pelz, J. B. (2009). Top-down control of eye movements: Yarbus revisited. *Visual Cognition*, *17*, 790–811.

- Doshi, A., & Trivedi, M. M. (2009). On the roles of eye gaze and head dynamics in predicting driver’s intent to change lanes. *Intelligent Transportation Systems, IEEE Transactions on*, *10*, 453–462.
- Doshi, A., & Trivedi, M. M. (2012). Head and eye gaze dynamics during visual attention shifts in complex environments. *Journal of Vision*, *12*.
- Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological review*, *96*, 433.
- Egeth, H. E., & Yantis, S. (1997). Visual attention: Control, representation, and time course. *Annual review of psychology*, *48*, 269–297.
- Epelboim, J., & Suppes, P. (2001). A model of eye movements and visual working memory during problem solving in geometry. *Vision Research*, *41*, 1561–1574.
- Fathi, A., Farhadi, A., & Rehg, J. M. (2011). Understanding egocentric activities. *ICCV*, .
- Ferguson, H. J., & Breheny, R. (2011). Eye movements reveal the time-course of anticipating behaviour based on complex, conflicting desires. *Cognition*, *119*.
- Fix, E., & Hodges, J. L. (1951). Discriminatory analysis, nonparametric discrimination: Consistency properties. *Technical Report 4, USAF School of Aviation Medicine, Randolph Field, Texas.*, .
- Folk, C. L., & Remington, R. (1998). Selectivity in distraction by irrelevant featural singletons: evidence for two forms of attentional capture. *Journal of Experimental Psychology: Human perception and performance*, *24*, 847.
- Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human perception and performance*, *18*, 1030.
- Freund, Y., & Schapire, R. E. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences.*, *1*, 119–139.
- Garcia-Diaz, A., Fdez-Vidal, X. R., Pardo, X. M., & Dosil, R. (2012). Saliency from hierarchical adaptation through decorrelation and variance normalization. *Image and Vision Computing.*, *30*, 51–64.
- Greene, M., Liu, T., & Wolfe, J. (2012). Reconsidering yarbus: A failure to predict observers’ task from eye movement patterns. *Vision Research*, *62*, 1–8.
- Griffiths, A. N., Marshall, R. W., & Richens, A. (1984). Saccadic eye movement analysis as a measure of drug effects on human psychomotor performance. *Br J Clin Pharmacol.*, *18*, 73S–82S.
- Guo, C., & Zhang, L. (2010). A novel multiresolution spatiotemporal saliency detection model and its applications in image and video compression. *IEEE Trans. on Image Processing.*, *19*.
- Hagemann, N., Schorer, J., Cañal-Bruland, R., Lotz, S., & Strauss, B. (2010). Visual perception in fencing: do the eye movements of fencers represent their information pickup? *Attention, perception & psychophysics*, *72*, 2204–2214.
- Haji-Abolhassani, A., & Clark, J. J. (2013). A computational model for task inference in visual search. *Journal of Vision.*, *13*, 1–24.
- Harel, J., Koch, C., & Perona, P. (2006). Graph-based visual saliency. *Advances in Neural Information Processing Systems (NIPS).*, *19*, 545–552.
- Harel, J., Moran, C., Huth, A., Einhaeuser, W., & Koch, C. (2008). Decoding what people see from where they look: predicting visual stimuli from scanpaths. *Lecture Notes in Artificial Intelligence (WAPCV).*, .

- Hayhoe, M., & Ballard, D. (2005). Eye movements in natural behavior. *Trends in Cognitive Sciences*, *9*, 188–194.
- Hayhoe, M. M., Shrivastava, A., Mruczek, R., & Pelz, J. B. (2003). Visual memory and motor planning in a natural task. *Journal of Vision*, *3*.
- Henderson, J., Shinkareva, S., Wang, J., Luke, S., & Olejarczyk, J. (2013). Predicting cognitive state from eye movements. *Plos One.*, *8*.
- Hou, X., & Zhang, L. (2007). Saliency detection: A spectral residual approach. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Hou, X., & Zhang, L. (2008). Dynamic visual attention: Searching for coding length increments. *Advances in Neural Information Processing Systems (NIPS)*., (pp. 681–688).
- Iqbal, S., & Bailey, B. (2004). Using eye gaze patterns to identify user tasks. In *The Grace Hopper Celebration of Women in Computing*.
- Itti, L., & Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision research.*, *40*, 1489–1506.
- Itti, L., & Koch, C. (2001). Computational modelling of visual attention. *Nature reviews. Neuroscience*, *2*, 194–203.
- Itti, L., Koch, C., & Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *20*, 1254–1259.
- Jarodzka, H., Scheiter, K., Gerjets, P., & Van Gog, T. (2010). In the eyes of the beholder: How experts and novices interpret dynamic stimuli. *Learning and Instruction*, *20*, 146–154.
- Jones, W., & Klin, A. (2013). Attention to eyes is present but in decline in 2-6-month-old infants later diagnosed with autism. *Nature, advance online publication*.
- Judd, T., Ehinger, K., Durand, F., & Torralba, A. (2009). Learning to predict where humans look. In *International Conference on Computer Vision (ICCV)*.
- Kaakinen, J. K., & Hyönä, J. (2010). Task effects on eye movements during reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *36*, 1561–1566.
- Kingstone, A. (2009). Taking a real look at social attention. *Journal of Vision*, *19*.
- Klin, A., Lin, D., Gorrindo, P., Ramsay, G., & Jones, W. (2009). Two-year-olds with autism orient to non-social contingencies rather than biological motion. *Nature*, *459*.
- Koch, C., & Ullman, S. (1985). Shifts in selective visual attention: Towards the underlying neural circuitry. *Human Neurobiology.*, *4*, 219–227.
- Kosslyn, S. M. (1994). *Image and Brain: The Resolution of the Imagery Debate*. MIT Press, Cambridge, MA.
- Kuhn, G., & Tatler, B. W. (2005). Magic and fixation: now you dont see it, now you do. *Perception.*, *34*, 1155–1161.
- Kuhn, G., Tatler, B. W., Findlay, J. M., & Cole, G. G. (2008). Misdirection in magic: Implications for the relationship between eye gaze and attention. *Visual Cognition*, *16*.
- Land, M. (2006). Eye movements and the control of actions in everyday life. *Progress in Retinal and Eye Research.*, *25*, 296–324.

- Land, M. F., & Hayhoe, M. (2001). In what ways do eye movements contribute to everyday activities? *Vision research*, *41*, 3559–3565.
- Land, M. F., & Lee, D. N. (1994). Where we look when we steer. *Nature*, *369*, 742–744.
- Land, M. F., & McLeod, P. (2000). From eye movements to actions: how batsmen hit the ball. *Nature neuroscience*, *3*, 1340–1345.
- Land, M. F., & Tatler, B. W. (2001). Steering with the head: The visual strategy of a racing driver. *Current Biology*, *11*, 1215–1220.
- Land, M. N., M., & Rusted, J. (1999). The roles of vision and eye movements in the control of activities of daily living. *Perception*, *28*, 1311–1328.
- Lethaus, F., Baumann, M. R. K., Köster, F., & Lemmer, K. (2013). A comparison of selected simple supervised learning algorithms to predict driver intent based on gaze data. *Neurocomputing*, *121*, 108–130.
- Loetscher, T., Bockisch, C., Nicholls, M., & Brugger, P. (2010). Eye position predicts what number you have in mind. *Current Biology*, *20*, 264–265.
- Mast, F., & Kosslyn, S. (2002). Eye movements during visual mental imagery. *TICS*, *6*.
- Meijering, B., van Rijn, H., Taatgen, N. A., & Verbrugge, R. (2012). What eye movements can tell about theory of mind in a strategic game. *Plos One*, *7*.
- Mennie, N., Hayhoe, M., & Sullivan, B. (2007). Look-ahead fixations: Anticipatory eye movements in natural tasks. *Experimental Brain Research*, *179*, 427–442.
- Mills, M., Hollingworth, A., Van der Stigchel, S., Hoffman, L., & Dodd, M. D. (2011). Examining the influence of task set on eye movements and fixations. *Journal of Vision*, *8*.
- Navalpakkam, V., & Itti, L. (2005). Modeling the influence of task on attention. *Vision Research*, *45*, 205–231.
- O’Connell, T., & Walther, D. (2012). Fixation patterns predict scene category. In *Proceedings of Vision Science Society (VSS)*.
- Parkhurst, D., Law, K., & Niebur, E. (2002). Modeling the role of salience in the allocation of overt visual attention. *Vision Research*, *42*, 107–123.
- Peters, R., & Itti, L. (2007). Congruence between model and human attention reveals unique signatures of critical visual events. In *Advances in Neural Information Processing Systems (NIPS)*.
- Peters, R. J., Iyer, A., Itti, L., & Koch, C. (2005). Components of bottom-up gaze allocation in natural images. *Vision Research*, *45*, 2397–2416.
- Pirsiavash, H., & Ramanan, D. (2012). Detecting activities of daily living in first-person camera views. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on* (pp. 2847–2854). IEEE.
- Poynter, W., Barber, M., Inman, J., & Wiggins, C. (2013). Individuals exhibit idiosyncratic eye-movement behavior profiles across tasks. *Vision Research*, *89*.
- Rajashekar, J., Bovik, L., & Cormack, A. (2006). Visual search in noise: Revealing the influence of structural cues by gaze-contingent classification image analysis. *Journal of Vision*, *17*, 379–386.
- Rayner, K. (1979). Eye guidance in reading: Fixation locations within words. *Perception*, *8*, 21–30.

- Rehder, B., & Hoffman, A. B. (2005). Eyetracking and selective attention in category learning. *Cognitive Psychology*, *Na*.
- Reichle, E. D., Rayner, K., & Pollatsek, A. (2003). The e-z reader model of eye movement control in reading: Comparisons to other models. *Behavioral and Brain Sciences*, *26*, 445–476.
- Renninger, L. W., Coughlan, J. M., Verghese, P., & Malik, J. (2005). An information maximization model of eye movements. In *Advances in Neural Information Processing Systems (NIPS)*.
- Risko, E. F., Anderson, N. C., Lanthier, S., & Kingstone, A. (2012). Curious eyes: Individual differences in personality predict eye movement behavior in scene-viewing. *Cognition*, *122*, 86–90.
- Salzberg, S. (1997). On Comparing Classifiers: Pitfalls to Avoid and a Recommended Approach. *Data Mining and Knowledge Discovery*, *1*, 317–328.
- Schapire, R. E. (1990). The strength of weak learnability. *Machine Learning*, *2*, 197–227.
- Schütz, A., Braun, D., & Gegenfurtner, K. (2011). Eye movements and perception: A selective review. *Journal of Vision*, *11*, 1–30.
- Seiffert, C., Khoshgoftaar, T. M., Van Hulse, J., & Napolitano, A. (2010). RUSBoost: A Hybrid Approach to Alleviating Class Imbalance. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, *40*, 185–197.
- Seo, H., & Milanfar, P. (2009). Static and space-time visual saliency detection by self-resemblance. *Journal of Vision*, *9*.
- Shaffer, J. (1995). Multiple hypothesis testing. *Annual Review of Psychol.*, *46*, 561–584.
- Sperling, G. (1960). The information available in brief visual presentations. *Psychological monographs: General and applied*, *74*, 1.
- Sperling, G., & Doshier, B. A. (1986). Strategy and optimization in human information processing, .
- Subramanian, R., Sebe, N., Kankanhalli, M., & Chua, T. (2010). An eye fixation database for saliency detection in images. In *European Conference on Computer Vision (ECCV)*.
- Tatler, B., Hayhoe, M., Land, M., & Ballard, D. (2011). Eye guidance in natural vision: Reinterpreting salience. *Journal of vision*, *11*, 1–23.
- Tatler, B., & Vincent, B. (2009). The prominence of behavioural biases in eye guidance. *Visual Cognition*, *17*, 1029–1054.
- Tatler, B., Wade, N., Kwan, H., Findlay, J., & Velichkovsky, B. (2010). Yarbus, eye movements, and vision. *i-Perception*, *1*, 7–27.
- Torralba, A., Oliva, A., Castelhana, M. S., & Henderson, J. M. (2006). Contextual guidance of eye movements and attention in real-world scenes: the role of global features in object search. *Psychological review*, *113*, 766–786.
- Treisman, A., & Gelade, G. (1980). A feature integration theory of attention. *Cognitive Psychology*, *12*, 97–136.
- Tseng, P., Cameron, I. G. M., Pari, G., Reynolds, J. N., Munoz, D. P., & Itti, L. (2012). High-throughput classification of clinical populations from natural viewing eye movements. *Journal of Neurology*, .
- Victor, T. W., Harbluk, J. L., & Engstrm, J. A. (2005). Sensitivity of eye-movement measures to in-vehicle task difficulty. *Transportation Research Part F*, *8*, 167–190.

- Vogt, S., & Magnussen, S. (2007). Expertise in pictorial perception: eye-movement patterns and visual memory in artists and laymen. *Perception*, *36*, 91–100.
- Windau, J., & Itti, L. (2013). Situation awareness via sensor-equipped eyeglasses. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- Woods, D. J., Beecher, G. P., & Ris, M. D. (1978). The effects of stressful arousal on conjugate lateral eye movement. *Motivation and Emotion*, *2*, 345–353.
- Yantis, S. (2000). Goal-directed and stimulus-driven determinants of attentional control. *Attention and performance*, *18*, 73–103.
- Yarbus, A. (1967). *Eye movements and vision*. New York: Plenum.
- Zelinsky, G., Peng, Y., & Samaras, D. (2013). Eye can read your mind: Decoding gaze fixations to reveal categorical search targets. *Journal of Vision*, *13*.
- Zelinsky, G. J. (2008). A theory of eye movements during target acquisition. *Psychological review*, *115*, 787–835.
- Zelinsky, Z. W., G., & Samaras, D. (2008). Eye can read your mind: Decoding eye movements to reveal the targets of categorical search tasks [abstract]. *Journal of Vision*, *8*, 380.
- Zhang, L., Tong, M. H., Marks, T. K., Shan, H., & Cottrell, G. W. (2008). Sun: A bayesian framework for saliency using natural statistics. *Journal of Vision*, *8*.
- Zhao, Q., & Koch, C. (2012). Learning visual saliency by combining feature maps in a nonlinear manner using adaboost. *Journal of Vision*, *12*.

4.1. *Appendix I: Task decoding accuracies over single images in experiment I using RUSBoost.*

4.2. *Appendix II: Task decoding accuracies over single images in experiment II using RUSBoost.*





	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
R1	0.1905	0.3333	0.3333	0.2381	0.2381	0.3333	0.1429	0.1905	0.2857	0.4286	0.1905	0.3810	0.2381	0.2857	0.2857
R2	0.1905	0.3333	0.2381	0.3333	0.1905	0.2857	0.0476	0.2857	0.2857	0.3810	0.1429	0.3810	0.2857	0.3810	0.2381
R3	0.1905	0.3333	0.3810	0.2381	0.2381	0.2857	0.0952	0.2381	0.3333	0.3810	0.1429	0.4286	0.2381	0.3333	0.2381
R4	0.0952	0.2857	0.3333	0.1905	0.2381	0.1905	0.0952	0.2381	0.2381	0.4286	0.1429	0.3333	0.3810	0.3810	0.2381
R5	0.1905	0.3810	0.3810	0.2381	0.1905	0.2381	0.1429	0.2381	0.2857	0.4762	0.2381	0.3810	0.3333	0.2857	0.2381
R6	0.1429	0.4762	0.3333	0.2381	0.2381	0.1905	0.0952	0.1905	0.2857	0.3810	0.0952	0.3810	0.2857	0.3333	0.3333
R7	0.0952	0.3333	0.3333	0.2381	0.3333	0.2381	0.0952	0.2381	0.2857	0.3810	0.2381	0.4286	0.2857	0.3333	0.2381
R8	0.1905	0.3333	0.3333	0.2381	0.2381	0.3333	0.1905	0.3333	0.2381	0.4286	0.2857	0.3810	0.2857	0.3810	0.2857
R9	0.1429	0.3333	0.2857	0.2857	0.2381	0.2857	0.1429	0.2857	0.2857	0.4286	0.1429	0.3333	0.2857	0.3810	0.2381
R10	0.0952	0.4286	0.3333	0.1905	0.2857	0.2381	0.0952	0.2381	0.2857	0.3810	0.1429	0.4286	0.2381	0.4286	0.2381
R11	0.1905	0.3810	0.2857	0.2381	0.2381	0.2381	0.1429	0.2381	0.2381	0.4286	0.1905	0.4286	0.2381	0.3810	0.1905
R12	0.0952	0.3333	0.3333	0.1429	0.2857	0.3333	0	0.2857	0.2857	0.4286	0.1429	0.3810	0.2381	0.2857	0.2381
R13	0.1429	0.3810	0.3810	0.2381	0.2381	0.2857	0.1429	0.3333	0.2381	0.3810	0.1429	0.3810	0.2857	0.3810	0.2381
R14	0.2381	0.4762	0.3810	0.2381	0.2381	0.2381	0.0476	0.2857	0.2381	0.4286	0.0952	0.3333	0.2381	0.3810	0.2857
R15	0.2381	0.3810	0.2381	0.2857	0.2381	0.1905	0.1429	0.2857	0.2381	0.4286	0.1429	0.3333	0.3333	0.2381	0.2857
R16	0.1905	0.3810	0.2857	0.1905	0.3810	0.2381	0.0476	0.2381	0.2857	0.4762	0.1429	0.3333	0.1905	0.3333	0.2857
R17	0.1905	0.4286	0.2857	0.1905	0.2857	0.2381	0.0476	0.2857	0.2857	0.3810	0.1429	0.3810	0.2381	0.4286	0.2857
R18	0.1429	0.3333	0.3333	0.1905	0.2857	0.2857	0.0952	0.3333	0.2857	0.4762	0.1429	0.3810	0.2381	0.4286	0.2381
R19	0.1429	0.4286	0.2381	0.2381	0.2381	0.2857	0.1429	0.2857	0.2857	0.3810	0.1429	0.4286	0.2857	0.3810	0.1905
R20	0.0952	0.3810	0.3810	0.2381	0.2381	0.3333	0.1429	0.2857	0.2857	0.4762	0.1429	0.4286	0.2857	0.2857	0.2381
R21	0.1429	0.3333	0.1905	0.1905	0.2857	0.2857	0.1429	0.3333	0.2857	0.4286	0.1429	0.3810	0.1905	0.3810	0.2381
R22	0.1905	0.3333	0.2857	0.1905	0.2381	0.3333	0.0952	0.2857	0.2381	0.3810	0.0952	0.4286	0.2381	0.3810	0.2381
R23	0.1429	0.3333	0.3810	0.2381	0.1905	0.3333	0.0952	0.3333	0.4286	0.4286	0.0476	0.4286	0.2857	0.2857	0.2381
R24	0.1905	0.4286	0.2857	0.2857	0.2381	0.2381	0.1429	0.2857	0.3333	0.4286	0.1905	0.3810	0.2381	0.2857	0.2857
R25	0.1429	0.4286	0.2857	0.2857	0.2857	0.2857	0.0476	0.2857	0.1905	0.3810	0.0952	0.3810	0.2857	0.3333	0.2381
R26	0.1429	0.3333	0.3333	0.2381	0.3333	0.3333	0.0952	0.2381	0.2857	0.4286	0.0952	0.3810	0.2381	0.3810	0.2381
R27	0.2857	0.3333	0.3333	0.2857	0.2857	0.3333	0.1905	0.2857	0.3333	0.4286	0.2381	0.3810	0.2381	0.3333	0.2857
R28	0.1429	0.3810	0.2381	0.1905	0.2857	0.3333	0.1905	0.2857	0.3333	0.3810	0.0952	0.3810	0.2857	0.3810	0.1905
R29	0.1429	0.3333	0.4286	0.2857	0.2381	0.3810	0.0476	0.1905	0.2857	0.4286	0.0476	0.4286	0.2857	0.4286	0.2381
R30	0.1429	0.4286	0.3333	0.2381	0.2857	0.2857	0.1429	0.2857	0.3810	0.4762	0.1905	0.3810	0.2857	0.4286	0.3333
R31	0.1429	0.4286	0.3810	0.1905	0.2381	0.1905	0.1429	0.2857	0.2381	0.3810	0.0952	0.3810	0.2857	0.3333	0.1905
R32	0.1429	0.3810	0.2857	0.2857	0.2857	0.1905	0.0952	0.2381	0.2857	0.3810	0.1905	0.4286	0.2381	0.3810	0.2857
R33	0.1429	0.3333	0.3333	0.3333	0.2381	0.3333	0.1429	0.2381	0.2857	0.4286	0.2857	0.4286	0.1905	0.3810	0.2381
R34	0.1429	0.4286	0.2857	0.2381	0.2857	0.2857	0.0476	0.1905	0.2857	0.4286	0.1429	0.4286	0.2381	0.3810	0.2381
R35	0.1905	0.3333	0.3810	0.2381	0.2381	0.2857	0.1429	0.2381	0.2857	0.4762	0.1905	0.3333	0.2381	0.4286	0.1905
R36	0.1429	0.3810	0.3333	0.2381	0.2857	0.3333	0.0952	0.2381	0.2857	0.4762	0.1905	0.3810	0.1905	0.4286	0.2381
R37	0.1429	0.3333	0.3333	0.1905	0.2857	0.2381	0.0476	0.3333	0.3333	0.4286	0.2381	0.3810	0.2381	0.3810	0.3333
R38	0.1905	0.3810	0.3810	0.2381	0.3333	0.2857	0.0476	0.2857	0.2381	0.3333	0.0952	0.4286	0.2857	0.4286	0.1905
R39	0.1429	0.3810	0.3333	0.1429	0.2381	0.3333	0.0952	0.2381	0.2857	0.3810	0.2381	0.3810	0.2381	0.3333	0.2381
R40	0.1905	0.2857	0.2381	0.1429	0.2857	0.3333	0.0476	0.2857	0.3333	0.3810	0.0952	0.4286	0.2381	0.4286	0.2857
R41	0.1429	0.3333	0.3333	0.2857	0.2381	0.3333	0.0476	0.2381	0.3333	0.3810	0.1429	0.3810	0.1905	0.4286	0.1905
R42	0.1905	0.3333	0.2381	0.3333	0.1905	0.2857	0.0476	0.2857	0.2857	0.3810	0.1429	0.3810	0.2857	0.3810	0.2381
R43	0.1429	0.3810	0.3810	0.2381	0.2381	0.2857	0.1429	0.3333	0.2381	0.3810	0.1429	0.3810	0.2857	0.3810	0.2381
R44	0.2381	0.4762	0.3810	0.2381	0.2381	0.2381	0.0476	0.2857	0.2381	0.4286	0.0952	0.3333	0.2381	0.3810	0.2857
R45	0.2381	0.3810	0.2381	0.2857	0.2381	0.1905	0.1429	0.2857	0.2381	0.4286	0.1429	0.3333	0.3333	0.2381	0.2857
R46	0.1905	0.3810	0.2857	0.1905	0.3810	0.2381	0.0476	0.2381	0.2857	0.4762	0.1429	0.3333	0.1905	0.3333	0.2857
R47	0.1905	0.4286	0.2857	0.1905	0.2857	0.2381	0.0476	0.2857	0.2857	0.3810	0.1429	0.3810	0.2381	0.4286	0.2857
R48	0.1429	0.3333	0.3333	0.1905	0.2857	0.2857	0.0952	0.3333	0.2857	0.4762	0.1429	0.3810	0.2381	0.4286	0.2381
R49	0.1429	0.4286	0.2381	0.2381	0.2381	0.2857	0.1429	0.2857	0.2857	0.3810	0.1429	0.4286	0.2857	0.3810	0.1905
R50	0.1905	0.3810	0.2381	0.2857	0.2857	0.2381	0.1429	0.3333	0.2381	0.4286	0.1429	0.3810	0.3333	0.2857	0.2857
Avg.	0.1648	0.3733	0.3152	0.2352	0.2619	0.2771	<b>0.1019</b>	0.2724	0.2828	<b>0.4162</b>	0.1514	0.3867	0.2600	0.3648	0.2505
p-val.	0.0004	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.2602	0.0000	0.0000	0.0000	0.0000

Table 4: Performance of the RUSBoost classifier for task decoding in experiment 2 using one observer out procedure. Columns represent images 1 to 15 and each row corresponds to an individual run. Chance level is at 14.29%. Results are using the first two feature types (i.e., a 11610D vector). Shown by the last row p-values (across RUSBoost runs), decoding is significantly above chance for some images, is significantly below chance for image number 7, and is non-significant versus chance for image number 11 (using t-test).