

What Do Different Evaluation Metrics Tell Us About Saliency Models?

Zoya Bylinskii¹, Tilke Judd, Aude Oliva, Antonio Torralba, and Frédo Durand

Abstract—How best to evaluate a saliency model’s ability to predict where humans look in images is an open research question. The choice of evaluation metric depends on how saliency is defined and how the ground truth is represented. Metrics differ in how they rank saliency models, and this results from how false positives and false negatives are treated, whether viewing biases are accounted for, whether spatial deviations are factored in, and how the saliency maps are pre-processed. In this paper, we provide an analysis of 8 different evaluation metrics and their properties. With the help of systematic experiments and visualizations of metric computations, we add interpretability to saliency scores and more transparency to the evaluation of saliency models. Building off the differences in metric properties and behaviors, we make recommendations for metric selections under specific assumptions and for specific applications.

Index Terms—Saliency models, evaluation metrics, benchmarks, fixation maps, saliency applications

1 INTRODUCTION

AUTOMATICALLY predicting regions of high saliency in an image is useful for applications including content-aware image re-targeting, image compression and progressive transmission, object and motion detection, image retrieval and matching. Where human observers look in images is often used as a ground truth estimate of image saliency, and computational models producing a saliency value at each pixel of an image are referred to as saliency models.¹

Dozens of computational saliency models are available to choose from [7], [8], [11], [12], [37], but objectively determining which model offers the “best” approximation to human eye fixations remains a challenge. For example, for the input image in Fig. 1a, we include the output of 8 different saliency models (Fig. 1b). When compared to human ground truth the saliency models receive different scores according to different evaluation metrics (Fig. 1c). The inconsistency in how different metrics rank saliency models can often leave performance up to interpretability.

In this paper, we quantify metric behaviors. Through a series of systematic experiments and novel visualizations (Fig. 2), we aim to understand how changes in the input saliency maps impact metric scores, and as a result why models are scored differently. Some metrics take a probabilistic

approach to distribution comparison, yet others treat distributions as histograms or random variables (Section 4). Some metrics are especially sensitive to false negatives in the input prediction, others to false positives, center bias, or spatial deviations (Section 5). Differences in how saliency and ground truth are represented and which attributes of saliency models should be rewarded/penalized leads to different choices of metrics for reporting performance [8], [12], [42], [43], [54], [63], [82]. We consider metric behaviors in isolation from any post-processing or regularization on the part of the models.

Building on the results of our analyses, we offer guidelines for designing saliency benchmarks (Section 6). For instance, for evaluating probabilistic saliency models we suggest the KL-divergence and Information Gain (IG) metrics. For benchmarks like the MIT Saliency Benchmark which do not expect saliency models to be probabilistic, but do expect models to capture viewing behavior including systematic biases, we recommend either Normalized Scanpath Saliency (NSS) or Pearson’s Correlation Coefficient (CC).

Our contributions include:

- An analysis of 8 metrics commonly used in saliency evaluation. We discuss how these metrics are affected by different properties of the input, and the consequences for saliency evaluation.
- Visualizations for all the metrics to add interpretability to metric scores and transparency to the evaluation of saliency models.
- An accompanying manuscript to the MIT Saliency Benchmark to help interpret results.
- Guidelines for designing new saliency benchmarks, including defining expected inputs and modeling assumptions, specifying a target task, and choosing how to handle dataset bias.
- Advice for choosing saliency evaluation metrics based on design choices and target applications.

1. Although the term saliency was traditionally used to refer to bottom-up conspicuity, many modern saliency models include scene layout, object locations, and other contextual information.

• Z. Bylinskii, A. Oliva, A. Torralba, and F. Durand are with the Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139.
E-mail: {zoya, oliva, torralba, fredo}@csail.mit.edu.
• T. Judd is with Google, Zurich, Switzerland. E-mail: tilke.judd@gmail.com.

Manuscript received 14 Apr. 2016; revised 22 Dec. 2017; accepted 10 Mar. 2018. Date of publication 12 Mar. 2018; date of current version 13 Feb. 2019. (Corresponding author: Zoya Bylinskii.)

Recommended for acceptance by L. Zelnik-Manor.

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

Digital Object Identifier no. 10.1109/TPAMI.2018.2815601

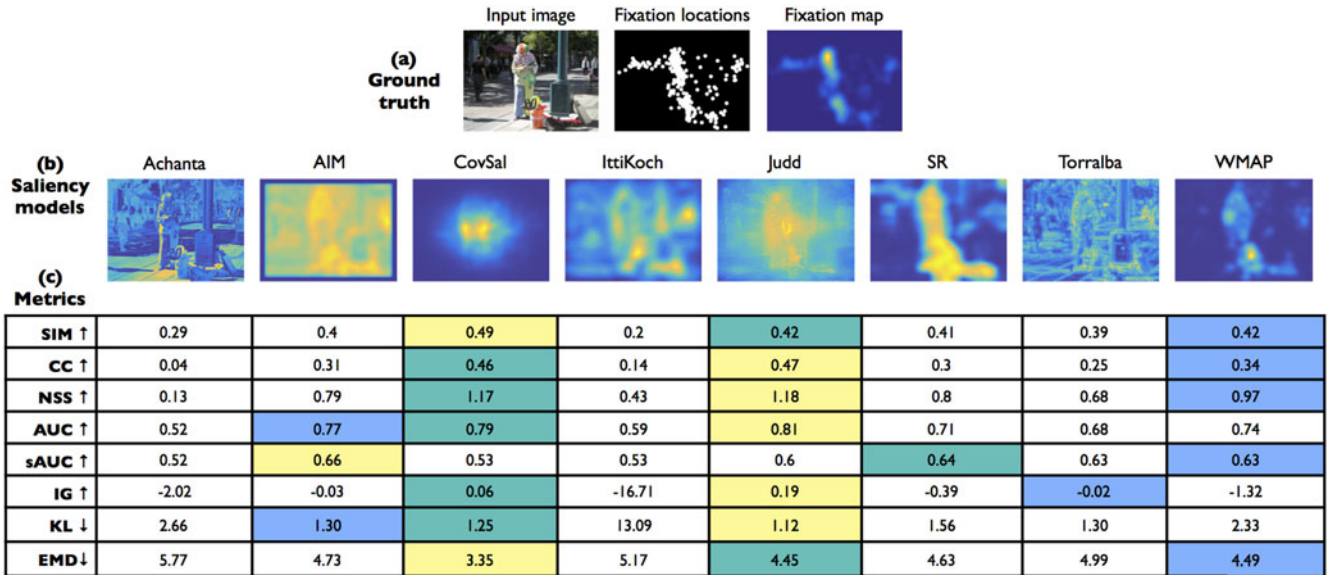


Fig. 1. Different evaluation metrics score saliency models differently. Saliency maps are evaluated on how well they approximate human ground truth eye movements (a), represented either as discrete fixation locations or a continuous fixation map. For a given image, saliency maps corresponding to 8 saliency models (b) are scored under 8 different evaluation metrics (c), 6 similarity and 2 dissimilarity metrics. We highlighted the top 3 best scoring maps (in yellow, green, and blue, respectively) under each metric (per row) for the particular image in (a).

2 RELATED WORK

2.1 Evaluation Metrics for Computer Vision

Similarity metrics operating on image features have been a subject of investigation and application to different computer vision domains [48], [70], [77], [88]. Images are often represented as histograms or distributions of features, including low-level features like edges (texture), shape and color, and higher-level features like objects, object parts, and bags of low-level features. Similarity metrics applied to these feature representations have been used for classification, image retrieval, and image matching tasks [65], [69], [70]. Properties of these metrics across different computer vision tasks also apply to the task of saliency modeling, and we provide a

discussion of some applications in Section 6.3. The discussion and analysis of the metrics in this paper can correspondingly be generalized to other computer vision applications.

2.2 Evaluation Metrics for Saliency

A number of papers in recent years have compared models across different metrics and datasets. Wilming et al. [82] discussed the choice of metrics for saliency model evaluation, deriving a set of qualitative and high-level desirable properties for metrics: “few parameters”, “intuitive scale”, “low data demand”, and “robustness”. Metrics were discussed from a theoretical standpoint without empirical experiments or quantification of metric behavior.

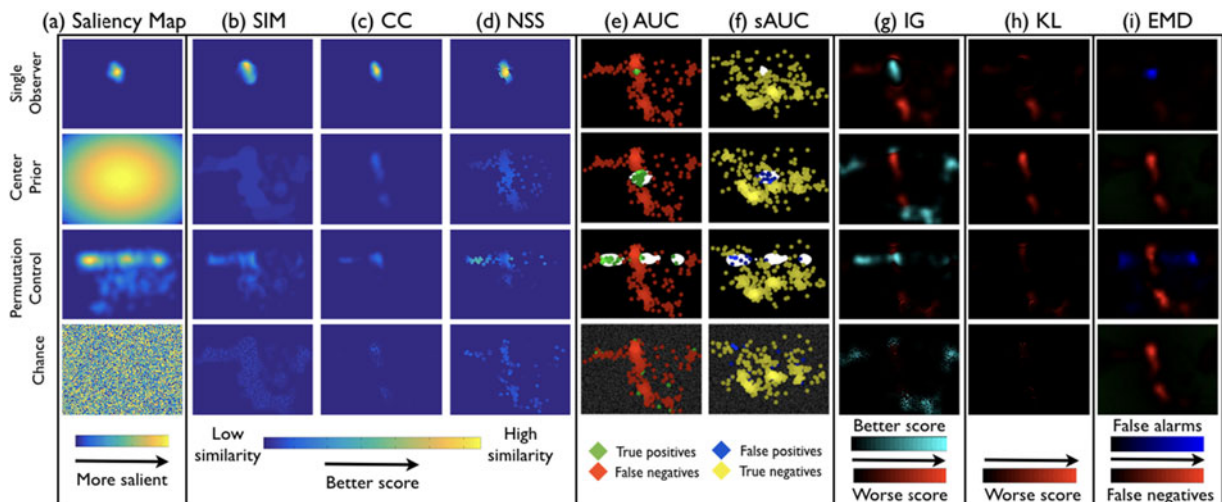


Fig. 2. A series of experiments and corresponding visualizations can help us understand what behaviors of saliency models different evaluation metrics capture. Given a natural image and ground truth human fixations on the image as in Fig. 1a, we evaluate saliency models, including the 4 baselines in column (a), at their ability to approximate ground truth. Visualizations of 8 common metrics (b-i) help elucidate the computations performed when scoring saliency models. In each visualization, higher density regions correspond either to matches or mismatches between the saliency map (from the first column) and the human ground truth (from Fig. 1a). The visualizations for all the metrics are explained in greater detail in Figs. 6-11 throughout the rest of the paper.

TABLE 1
The Most Common Metrics for Saliency Model Evaluation Are Analyzed in This Paper

Metric	Denoted	Used in evaluations
Area under ROC Curve	AUC	[8], [22], [23], [46], [54], [63], [82], [86]
Shuffled AUC	sAUC	[7], [8], [46], [63]
Normalized Scanpath Saliency	NSS	[7], [8], [22], [46], [54], [63], [82], [86]
Pearson's Correlation Coefficient	CC	[7], [8], [22], [23], [46], [63], [82]
Earth Mover's Distance	EMD	[46], [63], [86]
Similarity or histogram intersection	SIM	[46], [63]
Kullback-Leibler divergence	KL	[22], [46], [63], [82]
Information Gain	IG	[42], [43]

We include a list of the surveys that have used these metrics.

Le Meur and Baccino [54] reviewed many methods of comparing scanpaths and saliency maps. For evaluation, however, only 2 metrics were used to compare 4 saliency models. Sharma and Alsam [76] reported the performance of 11 models with 3 versions of the AUC metric on MIT1003 [38]. Zhao and Koch [86] analyzed saliency on 4 datasets using 3 metrics. Riche et al. [63] evaluated 12 saliency models with 12 similarity metrics on Jian Li's dataset [45]. They compared how metrics rank saliency models and reported which metrics cluster together, but did not provide explanations.

Borji, Sihite et al. [7] compared 35 models on a number of image and video datasets using 3 metrics. Borji, Tavakoli et al. [8] compared 32 saliency models with 3 metrics for fixation prediction and additional metrics for scanpath prediction on 4 datasets. The effects of center bias and map smoothing on model evaluation were discussed. A synthetic experiment was run with a single set of random fixations while blur sigma, center bias, and border size were varied to determine how the 3 different metrics are affected by these transformations. Our analysis extends to 8 metrics tested on different variants of synthetic data to explore the space of metric behaviors.

Li et al. [46] used crowdsourced perceptual experiments to discover which metrics most closely correspond to visual comparison of spatial distributions. Participants were asked to select out of pairs of saliency maps the map perceived to be closest to the ground truth map. Human annotations were used to order saliency models, and this ranking was compared to rankings by 9 different metrics. However, human perception can naturally favor some saliency map properties over others (Section 2.3). Visual comparisons are affected by the range and scale of saliency values, and are driven by the most salient locations, while small values are not as perceptible and don't enter into the visual calculations. This is in contrast to metrics that are particularly sensitive to zero values and regularization, which might nevertheless be more appropriate for certain applications, for instance when evaluating probabilistic saliency models (Section 6.4).

Emami and Hoberock [22] compared 9 evaluation metrics (3 novel, 6 previously-published) in terms of human consistency. They defined the best evaluation metric as the one which best discriminates between a human saliency

map and a random saliency map, as compared to the ground truth map. Human fixations were split into 2 sets, to generate human saliency maps and ground truth maps for each image. This procedure was the only criterion by which metrics were evaluated, and the chosen evaluation metric was used to compare 10 saliency models.

In this paper, we analyze metrics commonly used in other evaluation efforts (Table 1) and reported on the MIT Saliency Benchmark [12]. We include Information Gain (IG), recently introduced by Kümmerer et al. [42], [43]. To visualize metric computations and highlight differences in metric behaviors, we used standard saliency models for which code is available online. These models, depicted in Fig. 1b, include Achanta [2], AIM [9], CovSal [24], IttiKoch [41], [78], Judd [38], SR [68], Torralba [75], and WMAP [50]. Models were used for visualization purposes only, as the primary focus of this paper is comparing the metrics, not the models.

Rather than providing tables of performance values and literature reviews of metrics, this paper offers intuition about how metrics perform under various conditions and where they differ, using experiments with synthetic and natural data, and visualizations of metric computations. We examine the effects of false positives and negatives, blur, dataset biases, and spatial deviations on performance. This paper offers a more complete understanding of evaluation metrics and what they measure.

2.3 Qualitative Evaluation of Saliency

Most saliency papers include side-by-side comparisons of different saliency maps computed for the same images (as in Fig. 1b). Visualizations of saliency maps are often used to highlight improvements over previous models. A few anecdotal images might be used to showcase model strengths and weaknesses. Bruce et al. [10] discussed the problems with visualizing saliency maps, in particular the strong effect that contrast has on the perception of saliency models. We propose supplementing saliency map examples with visualizations of metric computations (as in Figs. 6-11) to provide an additional means of comparison that is more tightly linked to the underlying model performance than the saliency maps themselves.

3 EVALUATION SETUP

The choice of metrics should be considered in the context of the whole evaluation setup, which requires the following decisions to be made: (1) on which input images saliency models will be evaluated, (2) how the ground truth eye movements will be collected (e.g., at which distance and for how long human observers view each image), and (3) how the eye movements will be represented (e.g., as discrete points, sequences, or distributions). In this section we explain the design choices behind our data collection and evaluation.

3.1 Data Collection

We used the MIT Saliency Benchmark dataset (*MIT300*) with 300 natural images [12], [37]. Eye movements were collected by having participants free-view each image for 2 seconds (details in the online supplemental material, which can be found on the Computer Society Digital Library at <http://doi.ieeecomputersociety.org/10.1109/TPAMI.2018.2815601>),

averaging 4-6 fixations per participant. This is sufficient to highlight a few points of interest per image, and provides a testing ground for saliency models. Other tasks (e.g., visual search) differently direct eye movements and may require alternative modeling assumptions [11]. The free viewing task is most commonly used for saliency modeling as it requires fewest additional assumptions.

The eye tracking set-up, including participant distance to the eye tracker, calibration error, and image size affects the assumptions that can be made about the collected data. In the eye tracking set-up of the MIT300 dataset, one degree of visual angle is approximately 35 pixels. One degree of visual angle is typically used both (1) as an estimate of the size of the human fovea: e.g., how much of the image a participant has in focus during a fixation, and (2) to account for measurement error in the eye tracking set-up. The robustness of the data also depends on the number of eye fixations collected. In the MIT300 dataset, the eye fixations of 39 observers are available per image, more than in other datasets of similar size.

3.2 Ground Truth Representation

Once collected, the ground truth eye fixations can be processed and formatted in a number of ways for saliency evaluation. There is a fundamental ambiguity in the correct representation for the fixation data, and different representational choices rely on different assumptions. One format is to use the original fixation locations. Alternatively, the discrete fixations can be converted into a continuous distribution, a *fixation map*, by smoothing (Fig. 1a). We follow common practice² and blur each fixation location using a Gaussian with sigma equal to one degree of visual angle [54]. In the following section, we denote the map of fixation locations as Q^B and the continuous fixation map (distribution) as Q^D .

Smoothing the fixation locations into a continuous map acts as regularization. It allows for uncertainty in the ground truth measurements to be incorporated: error in the eye-tracking as well as uncertainty of what an observer sees when looking at a particular location on the screen. Any splitting of observer fixations in two sets will never lead to perfect overlap (due to the discrete nature of the data), and smoothing provides additional robustness for evaluation. In the case of few observers, smoothing the fixation locations helps to extrapolate the existing data.

On the other hand, conversion of the fixation locations into a distribution requires parameter selection and post-processing the collected data. The smoothing parameter can significantly affect metric scores during evaluation (Table 12a), unless the model itself is properly regularized.

The fixation locations can be viewed as a discrete sample from some ground truth distribution that the fixation map attempts to approximate. Similarly, the fixation map can be viewed as an extrapolation of discrete fixation data to the case of infinite observers.

Metrics for the evaluation of sequences of fixations are also available [54]. However, most saliency models and evaluations are tuned for location prediction, as sequences

2. Some researchers choose to cross-validate the smoothing parameter instead of fixing it as a function of viewing angle [42], [43].

TABLE 2
Different Metrics Use Different Representations of Ground Truth for Evaluating Saliency Models

Metrics	Location-based	Distribution-based
Similarity	AUC, sAUC, NSS, IG	SIM, CC
Dissimilarity		EMD, KL

Location-based metrics consider saliency map values at discrete fixation locations, while distribution-based metrics treat both ground truth fixation maps and saliency maps as continuous distributions. Good saliency models should have high values for similarity metrics and low values for dissimilarity metrics.

tend to be noisier and harder to evaluate. We only consider spatial, not temporal, fixation data.

4 METRIC COMPUTATION

In this paper, we study saliency metrics, that is, functions that take two inputs representing eye fixations (ground truth and predicted) and then output a number assessing the similarity or dissimilarity between them. Given a set of ground truth eye fixations, such metrics can be used to define scoring functions, which take a saliency map prediction as input and return a number assessing the accuracy of the prediction. The definition of a score can further involve post-processing (or regularizing) the prediction to conform it to known characteristics of the ground truth and ignore potentially distracting idiosyncratic errors. In this paper, we focus on the metric and not on the regularization of ground truth data.

We consider 8 popular saliency evaluation metrics in their most common variants. Some metrics have been designed specifically for saliency evaluation (shuffled AUC, Information Gain, and Normalized Scanpath Saliency), while others have been adapted from signal detection (variants of AUC), image matching and retrieval (Similarity, Earth Mover’s Distance), information theory (KL-divergence), and statistics (Pearson’s Correlation Coefficient). Because of their original intended applications, these metrics expect different input formats: KL-divergence and Information Gain expect valid probability distributions as input, Similarity and Earth Mover’s Distance can operate on unnormalized densities and histograms, while Pearson’s Correlation Coefficient (CC) treats its inputs as random variables.

One of the intentions of this paper is to serve as a guide to complement the MIT Saliency Benchmark, and to provide interpretation for metric scores. The MIT Saliency Benchmark accepts saliency maps as intensity maps, without restricting input to be in any particular form (probabilistic or otherwise). If a metric expects valid probability distributions, we simply normalize the input saliency maps, without additional modifications or optimizations.

In this paper we analyze these 8 metrics in isolation from the input format and with minimal underlying assumptions. The only distinction we make in terms of the input that these metrics operate on is whether the ground-truth is represented as discrete fixation locations or a continuous fixation map. Accordingly, we categorize metrics as *location-based* or *distribution-based* (following Riche et al. [63]). This organization is summarized in Table 2. In this section, we discuss the particular advantages and disadvantages of each metric, and present visualizations of the metric computations. Additional variants

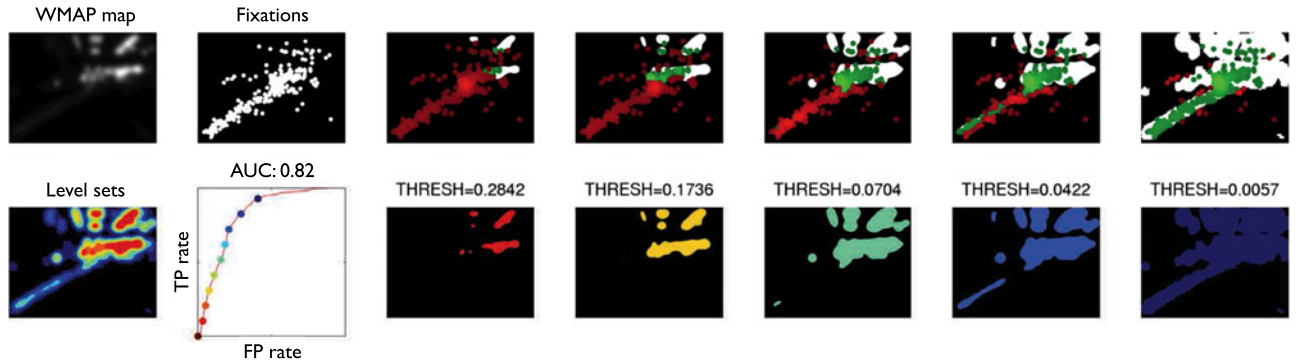


Fig. 3. The AUC metric evaluates a saliency map by how many ground truth fixations it captures in successive level sets. To compute AUC, a saliency map (left column) is treated as a binary classifier of fixations at various threshold values (THRESH) and an ROC curve is swept out (bottom left). Thresholding the saliency map produces level sets (columns 3-7). For each level set, the true positive (TP) rate is the proportion of fixations landing in the level set (green points). The false positive (FP) rate is the proportion of image pixels in the level set not covered in fixations (uncovered white regions). Fixations landing outside of a level set are false negatives (red points). We include 5 level sets colored to correspond to their TP and FP rates on the ROC curve. The AUC score for the saliency map is the area under the ROC curve.

and implementation details are provided in the online supplemental material.

4.1 Location-Based Metrics

4.1.1 Area Under ROC Curve (AUC): Evaluating Saliency as a Classifier of Fixations

Given the goal of predicting the fixation locations on an image, a saliency map can be interpreted as a classifier of which pixels are fixated or not. This suggests a detection metric for measuring saliency map performance. In signal detection theory, the Receiver Operating Characteristic (ROC) measures the tradeoff between true and false positives at various discrimination thresholds [25], [30]. The Area under the ROC curve, referred to as AUC, is the most widely used metric for evaluating saliency maps. The saliency map is treated as a binary classifier of fixations at various threshold values (level sets), and an ROC curve is swept out by measuring the true and false positive rates under each binary classifier (level set). Different AUC implementations differ in how true and false positives are calculated. Another way to think of AUC is as a measure of how well a model performs on a 2AFC task, where given 2 possible locations on the image, the model has to pick the location that corresponds to a fixation [43].

Computing True and False Positives. An AUC variant from Judd et al. [38], called *AUC-Judd* [12], is depicted in Fig. 3. For a given threshold, the true positive rate (*TP rate*) is the ratio of true positives to the total number of fixations, where true positives are saliency map values above threshold at *fixated pixels*. This is equivalent to the ratio of fixations falling within the level set to the total fixations.

The false positive rate (*FP rate*) is the ratio of false positives to the total number of saliency map pixels at a given threshold, where false positives are saliency map values above threshold at *unfixated pixels*. This is equivalent to the number of pixels in each level set, minus the pixels already accounted for by fixations.

Another variant of AUC by Borji et al. [7], called *AUC-Borji* [12], uses a uniform random sample of image pixels as negatives and defines the saliency map values above threshold at these pixels as false positives. These AUC implementations are compared in Fig. 4. The first row depicts the TP rate calculation, equivalent across implementations. The

second and third rows depict the FP rate calculations in *AUC-Judd* and *AUC-Borji*, respectively. The false positive calculation in *AUC-Borji* is a discrete approximation of the calculation in *AUC-Judd*. Because of a few approximations in the *AUC-Borji* implementation that can lead to suboptimal behavior, we report AUC scores using *AUC-Judd* in the rest of the paper. Additional discussion, implementation details, and other variants of AUC are discussed in the online supplemental material.

Penalizing Models for Center Bias. The natural distribution of fixations on an image tends to include a higher density near the center of an image [73]. As a result, a model that incorporates a center bias into its predictions will be able to account for at least part of the fixations on an image, independent of image content. In a center-biased dataset, a center prior baseline will achieve a high AUC score.

The shuffled AUC metric, *sAUC* [8], [20], [73], [74], [85] samples negatives from fixation locations from other images, instead of uniformly at random. This has the effect of sampling negatives predominantly from the image center because averaging fixations over many images results in the natural emergence of a central Gaussian distribution [73], [82]. In Fig. 4 the shuffled sampling strategy of *sAUC* is compared to the random sampling strategy of *AUC-Borji*.

A model that only predicts the center achieves an *sAUC* score of 0.5 because at all thresholds this model captures as many fixations on the target image as on other images ($TP\ rate = FP\ rate$). A model that incorporates a center bias into its predictions is putting density in the center at the expense of other image regions. Such a model will score worse according to *sAUC* compared to a model that makes off-center predictions, because *sAUC* will effectively discount the central predictions (Fig. 6). In other words, *sAUC* is not invariant to whether the center bias is modeled: it specifically penalizes models that include the center bias.

Invariance to Monotonic Transformations. AUC metrics measure only the relative (i.e., ordered) saliency map values at ground truth fixation locations. In other words, the AUC metrics are ambivalent to monotonic transformations. AUC is computed by varying the threshold of the saliency map and comparing true and false positives. Lower thresholds correspond to measuring the coverage similarity between distributions, while higher thresholds correspond



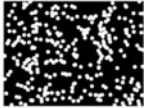
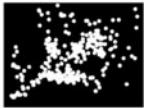
Ground truth fixation locations	Saliency map level sets				
Fixations 	TP rate = 0.10	TP rate = 0.30	TP rate = 0.52	TP rate = 0.69	TP rate = 0.92
Non fixated locations (b) FP for AUC-Judd 	FP rate = 0.06	FP rate = 0.21	FP rate = 0.33	FP rate = 0.42	FP rate = 0.54
Random samples (c) FP for AUC-Borji 	FP rate = 0.07	FP rate = 0.21	FP rate = 0.36	FP rate = 0.47	FP rate = 0.59
Shuffled samples (d) FP for sAUC 	FP rate = 0.20	FP rate = 0.63	FP rate = 0.75	FP rate = 0.80	FP rate = 0.89
Legend:	◆ True positives	◆ False negatives	◆ False positives	◆ True negatives	

Fig. 4. How true and false positives are calculated under different AUC metrics (AUC-Judd, AUC-Borji, sAUC): (a) In all cases, the true positive rate is calculated as the proportion of fixations falling into a thresholded saliency map (green over green plus red). These fixations are superimposed on 5 level sets of the CovSal [24] saliency map. (b) In AUC-Judd, the false positive rate is the proportion of non-fixated pixels in the thresholded saliency map (blue over blue plus yellow). (c) In AUC-Borji, this calculation is approximated by sampling negatives uniformly at random and computing the proportion of negatives in the thresholded region (blue over blue plus yellow). (d) In sAUC, negatives are sampled according to the distribution of fixations in other images instead of uniformly at random. Saliency models are scored similarly under the AUC-Judd and AUC-Borji metrics, but differently under sAUC due to the sampling of false positives.

to measuring the similarity between the peaks of the two maps [23]. Due to how the ROC curve is computed, the AUC score for a saliency map is mostly driven by higher thresholds: i.e., the number of ground truth fixations captured by the peaks of the saliency map, or the first few level sets (Fig. 5). Models that place high-valued predictions at fixated locations receive high scores, while low-valued predictions at non-fixated locations are mostly ignored (Section 5.2).

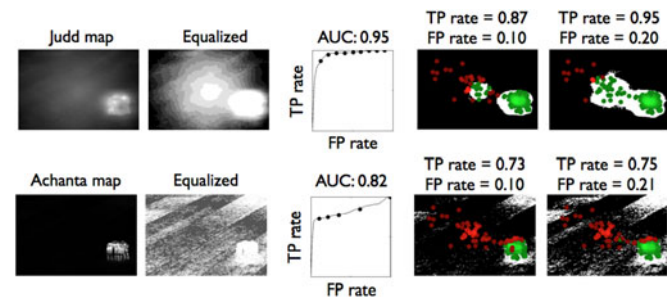


Fig. 5. The effect of true positive (TP) and false positive (FP) rates on AUC score. We compare the Judd [38] and Achanta [2] models by visualizing (left to right) the original saliency maps, equalized saliency maps (to highlight the level sets), ROC curves, and the first two level sets of both maps. Overlaid on the level sets are true positives (fixations classified salient) in green and false negatives (fixations classified non-salient) in red. Points on the ROC curve correspond to TP and FP rates of different level sets (both axes span 0 to 1). Judd accounts for more fixations in the first few level sets than Achanta, achieving a higher AUC score overall. The AUC score is driven most by the first few level sets, while the total number of levels sets and FP in later level sets have a significantly smaller impact. Achanta has a smaller range of saliency values, and thus fewer points on the ROC curve.

4.1.2 Normalized Scanpath Saliency (NSS): Measuring the Normalized Saliency at Fixations

The Normalized Scanpath Saliency, *NSS*, was introduced to the saliency community as a simple correspondence measure between saliency maps and ground truth, computed as the average normalized saliency at fixated locations [60]. Unlike in AUC, the absolute saliency values are part of the normalization calculation. *NSS* is sensitive to false positives, relative differences in saliency across the image, and general monotonic transformations. However, because the mean saliency value is subtracted during computation, *NSS* is invariant to linear transformations like contrast offsets. Given a saliency map P and a binary map of fixation locations Q^B :

$$NSS(P, Q^B) = \frac{1}{N} \sum_i \bar{P}_i \times Q_i^B \quad (1)$$

where $N = \sum_i Q_i^B$ and $\bar{P} = \frac{P - \mu(P)}{\sigma(P)}$,

where i indexes the i th pixel, and N is the total number of fixated pixels. Chance is at 0, positive *NSS* indicates correspondence between maps above chance, and negative *NSS* indicates anti-correspondence. For instance, a unity score corresponds to fixations falling on portions of the saliency map with a saliency value one standard deviation above average.

Recall that a saliency model with high-valued predictions at fixated locations would receive a high AUC score even in the presence of many low-valued false positives

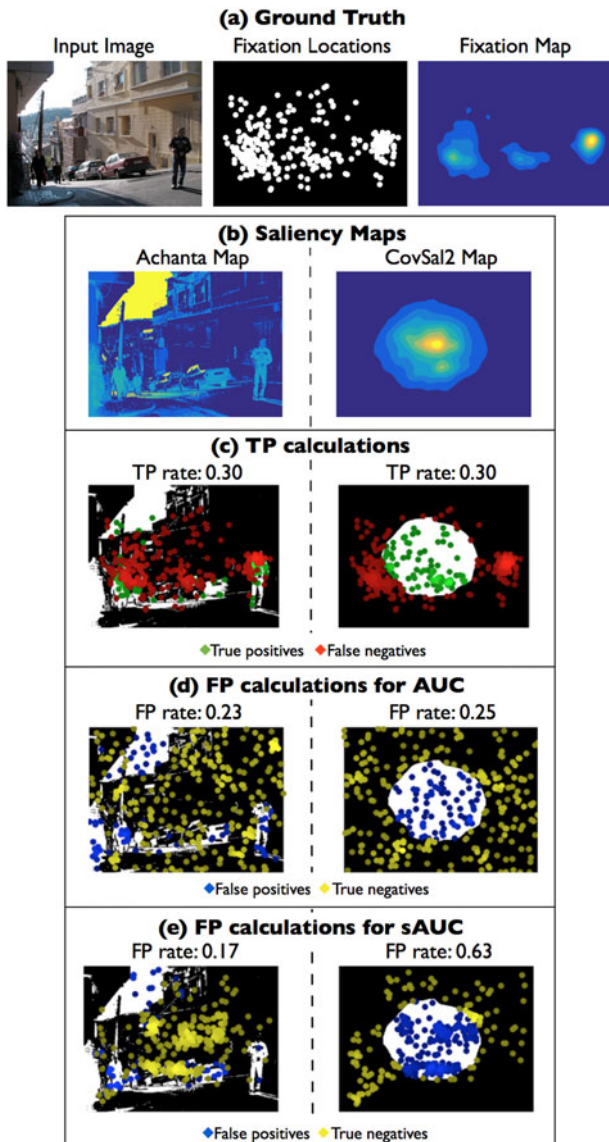


Fig. 6. Both AUC and sAUC measure the ability of a saliency map to classify fixated from non-fixated locations (Section 4.1.1). The main difference is that AUC prefers maps that account for center bias, while sAUC penalizes them. The saliency maps in (b) are compared on their ability to predict the ground truth fixations in (a). For a particular level set, the true positive rate is the same for both maps (c). The sAUC metric normalizes this value by fixations sampled from other images, more of which land in the center of the image, thus penalizing the rightmost model for its center bias (d). The AUC metric, however, samples fixations uniformly at random and prefers the center-biased model which better explains the overall viewing behavior (e).

(Fig. 7d). However, false positives lower the normalized saliency value at each fixation location, thus reducing the overall NSS score (Fig. 7c). The visualization for NSS consists of the normalized saliency value (\bar{P}_i) at each fixation location ($Q_i^B = 1$).

4.1.3 Information Gain (IG): Evaluating Information Gain Over a Baseline

Information Gain, IG , was recently introduced by Kümmerer et al. [42], [43] as an information theoretic metric that measures saliency model performance beyond systematic bias (e.g., a center prior baseline). Given a binary map of fixations Q^B , a saliency map P , and a baseline map B :

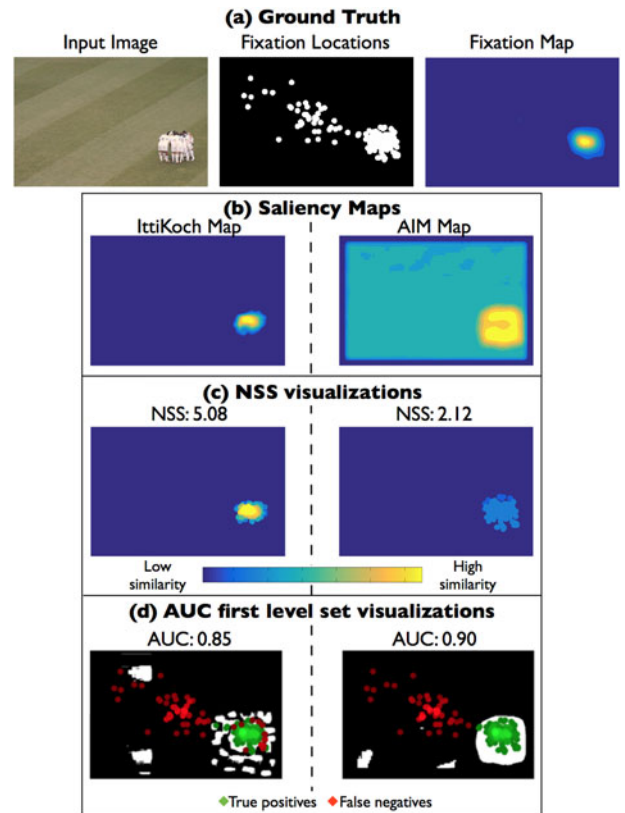


Fig. 7. Both AUC and NSS evaluate the ability of a saliency map (b) to predict fixation locations (a). AUC is invariant to monotonic transformations (Section 4.1.1), while NSS is not. NSS normalizes a saliency map by the standard deviation of the saliency values (Section 4.1.2). AUC ignores low-valued false positives but NSS penalizes them. As a result, the rightmost map has a lower NSS score because more false positives means the normalized saliency value at fixation locations drops (c). The AUC score of the left and right maps is very similar since a similar number of fixations fall in equally-sized level sets of the two saliency maps (d).

$$IG(P, Q^B) = \frac{1}{N} \sum_i Q_i^B [\log_2(\epsilon + P_i) - \log_2(\epsilon + B_i)], \quad (2)$$

where i indexes the i th pixel, N is the total number of fixated pixels, ϵ is for regularization. This metric measures the average information gain of the saliency map over the center prior baseline at fixated locations (at $Q^B = 1$), in bits per fixation.

IG assumes that the input saliency maps are probabilistic, properly regularized and optimized to include a center prior [42], [43]. A score above zero indicates the saliency map predicts the fixated locations better than the center prior baseline. This score measures how much image-specific saliency is predicted beyond image-independent dataset biases, which in turn requires careful modeling of these biases.

We can also compute the information gain of one model over another to measure how much image-specific saliency is captured by one model beyond what is already captured by the other. The example in Fig. 8 contains a visualization of the information gain of the Judd model over the center prior baseline and IttiKoch models. In red are image regions where the Judd model underestimates saliency relative to each model, and in blue are image regions where the Judd model achieves a gain in performance. The human under the parachute is salient under the center prior model, while the Judd model underestimates this saliency (red); the Judd

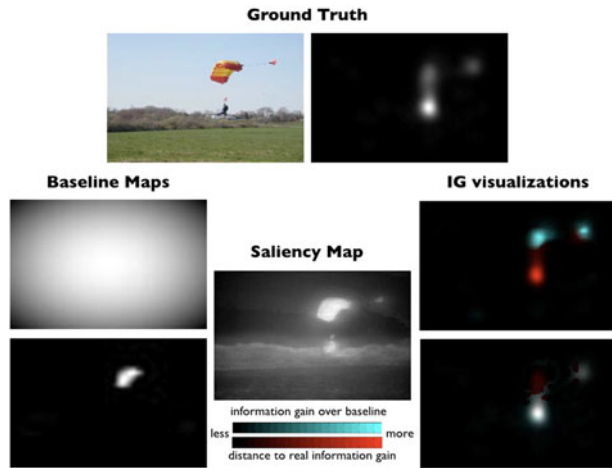


Fig. 8. We compute the information of one model over another at predicting ground truth fixations. We visualize the information gain of the Judd model over the center prior baseline (top) and the bottom-up IttiKoch model (bottom). In blue are the image pixels where the Judd model makes better predictions than each model. In red is the remaining distance to the real information gain: i.e., image pixels at which the Judd model underestimates saliency.

model has positive information gain over the center prior on the parachute (blue). On the other hand, the bottom-up Itti-Koch model captures the parachute but misses the person in the center of the image, so in this case the Judd model achieves gains on the central image pixels but not on the parachute. We refer the reader to [43] for detailed discussions and visualizations of the IG metric.

4.2 Distribution-Based Metrics

The (location-based) metrics described so far measure the accuracy of saliency models at predicting discrete fixation locations. If the ground truth fixation locations are interpreted as a sample from some underlying probability distribution, then another approach is to predict the distribution directly instead of the fixation locations. Although we can not directly observe this ground truth distribution, it is often approximated by Gaussian blurring the fixation locations into a fixation map (Section 3.2). Next we discuss a set of metrics that measure the accuracy of saliency models at approximating the continuous fixation map.

4.2.1 Similarity (SIM): Measuring the Intersection Between Distributions

The similarity metric, *SIM* (also referred to as *histogram intersection*), measures the similarity between two distributions, viewed as histograms. First introduced as a metric for color-based and content-based image matching [66], [72], it has gained popularity in the saliency community as a simple comparison between pairs of saliency maps. *SIM* is computed as the sum of the minimum values at each pixel, after normalizing the input maps. Given a saliency map P and a continuous fixation map Q^D :

$$\begin{aligned} SIM(P, Q^D) &= \sum_i \min(P_i, Q_i^D) \\ \text{where } \sum_i P_i &= \sum_i Q_i^D = 1, \end{aligned} \quad (3)$$

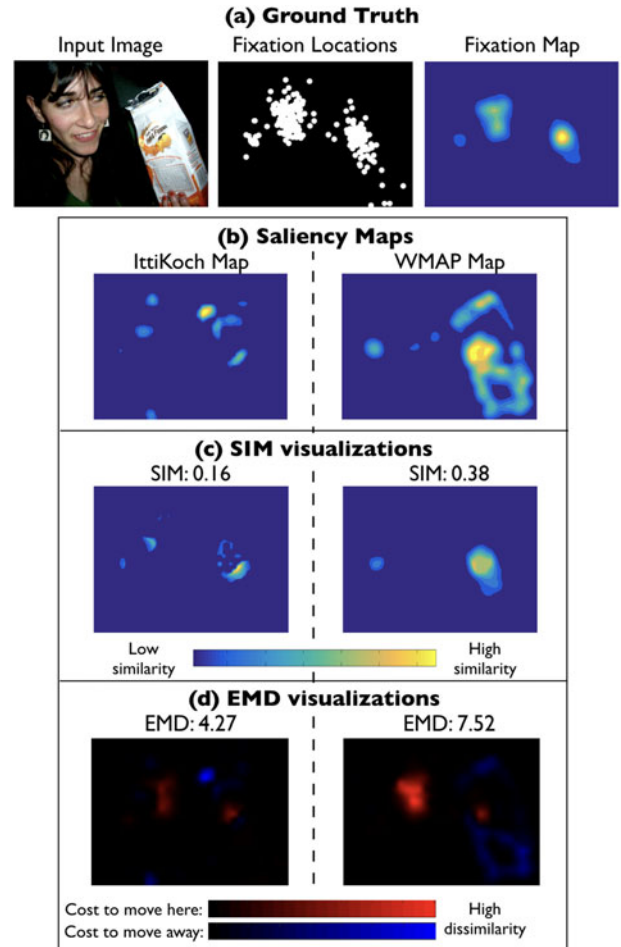


Fig. 9. The EMD and SIM metrics measure the similarity between the saliency map (b) and ground truth fixation map (a). EMD measures how much density needs to move before the two maps match (Section 4.2.4), while SIM measures the direct intersection between the maps (Section 4.2.1). EMD prefers sparser predictions, even if they do not perfectly align with fixated regions, while SIM penalizes misalignment. The saliency map on the left makes sparser predictions, resulting in a smaller intersection with the ground truth, and lower SIM score (c). However, the predicted density in this map is spatially closer to the ground truth density, achieving a better EMD score than the map on the right (d).

iterating over discrete pixel locations i . A SIM of one indicates the distributions are the same, while a SIM of zero indicates no overlap. Fig. 9c contains a visualization of this operation. At each pixel i of the visualization, we plot $\min(P_i, Q_i^D)$. Note that the model with the sparser saliency map has a lower histogram intersection with the ground truth map. *SIM* is very sensitive to missing values, and penalizes predictions that fail to account for all of the ground truth density (Section 5.2).

Effect of Blur on Model Performance. The downside of a distribution metric like *SIM* is that the choice of the Gaussian sigma (or blur) in constructing the fixation and saliency maps affects model evaluation. For instance, as demonstrated in the synthetic experiment in Fig. 12a, even if the correct location is predicted, *SIM* will only reach its maximal value when the saliency map's sigma exactly matches the ground truth sigma. The *SIM* score drops off drastically under different sigma values, more than the other metrics. Fine-tuning this blur value on a training set with similar parameters as the test set

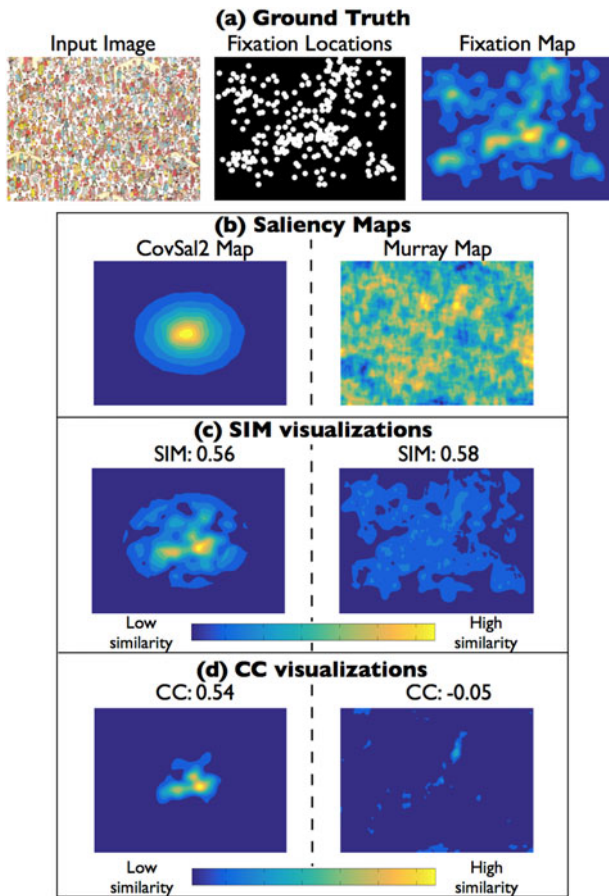


Fig. 10. The SIM and CC metrics measure the similarity between the saliency map (b) and ground truth fixation map (a). SIM measures the histogram intersection between two maps (Section 4.2.1), while CC measures their cross correlation (Section 4.2.2). CC treats false positives and negatives symmetrically, but SIM places less emphasis on false positives. As a result, both saliency maps have similar SIM scores (c), but the saliency map on the right has a lower CC score because false positives lower the overall correlation (d).

(eyetracking set-up, viewing angle) can help boost model performances [12], [37].

The SIM metric is good for evaluating partial matches, where a subset of the saliency map accounts for the ground truth fixation map. As a side-effect, false positives tend to be penalized less than false negatives. For other applications, a metric that treats false positives and false negatives symmetrically, such as CC or NSS, may be preferred.

4.2.2 Pearson's Correlation Coefficient (CC): Evaluating the Linear Relationship Between Distributions

The Pearson's Correlation Coefficient, CC, also called *linear correlation coefficient* is a statistical method used in the sciences to measure how correlated or dependent two variables are. Interpreting saliency and fixation maps, P and Q^D , as random variables, CC measures the linear relationship between them [55]:

$$CC(P, Q^D) = \frac{\sigma(P, Q^D)}{\sigma(P) \times \sigma(Q^D)}, \quad (4)$$

where $\sigma(P, Q^D)$ is the covariance of P and Q^D . CC is symmetric and penalizes false positives and negatives equally.

It is invariant to linear (but not arbitrary monotonic) transformations. High positive CC values occur at locations where both the saliency map and ground truth fixation map have values of similar magnitudes. Fig. 10 is an illustrative example comparing the behaviors of SIM and CC: where SIM penalizes false negatives significantly more than false positives, but CC treats both symmetrically. For visualizing CC in Fig. 10d, each pixel i takes on the value:

$$V_i = \frac{P_i \times Q_i^D}{\sqrt{\sum_j (P_j^2 + (Q_j^D)^2)}}. \quad (5)$$

Due to its symmetric computation, CC can not distinguish whether differences between maps are due to false positives or false negatives. Other metrics may be preferable if this kind of analysis is of interest.

4.2.3 Kullback-Leibler Divergence (KL): Evaluating Saliency with a Probabilistic Interpretation

Kullback-Leibler (KL) is a broadly-used information theoretic measure of the difference between two probability distributions. In the saliency literature, depending on how the saliency predictions and ground truth fixations are interpreted as distributions, different KL computations are possible. We discuss a few alternative varieties in the online supplemental material. To avoid future confusion about the KL implementation used, we can refer to this variant as *KL-Judd* similarly to how the AUC variant traditionally used on the MIT Benchmark is denoted AUC-Judd. Analogous to our other distribution-based metrics, KL-Judd takes as input a saliency map P and a ground truth fixation map Q^D , and evaluates the loss of information when P is used to approximate Q^D :

$$KL(P, Q^D) = \sum_i Q_i^D \log \left(\epsilon + \frac{Q_i^D}{\epsilon + P_i} \right), \quad (6)$$

where ϵ is a regularization constant.³ KL-Judd is an asymmetric dissimilarity metric, with a lower score indicating a better approximation of the ground truth by the saliency map. We compute a per-pixel score to visualize the KL computation (Fig. 11d). For each pixel i in the visualization, we plot $Q_i^D \log \left(\epsilon + \frac{Q_i^D}{\epsilon + P_i} \right)$. Wherever the ground truth value Q_i^D is non-zero but P_i is close to or equal to zero, a large quantity is added to the KL score. Such regions are the brightest in the visualization. There are more bright regions in the rightmost map of Fig. 11d, corresponding to ground truth areas that were not predicted salient. Both models compared in Fig. 11 are image-agnostic: one is a chance model that assigns a uniform value to each pixel in the image, and the other is a *permutation control* model which uses a fixation map from another randomly-selected image. The permutation control model is more likely to capture viewing biases common across images. It scores above chance for many metrics in Table 3. However, KL is so sensitive to zero-values

3. The relative magnitude of ϵ will affect the regularization of the saliency maps and how much zero-valued predictions are penalized. The MIT Saliency Benchmark uses MATLAB's built-in $\text{eps} = 2.2204e-16$.

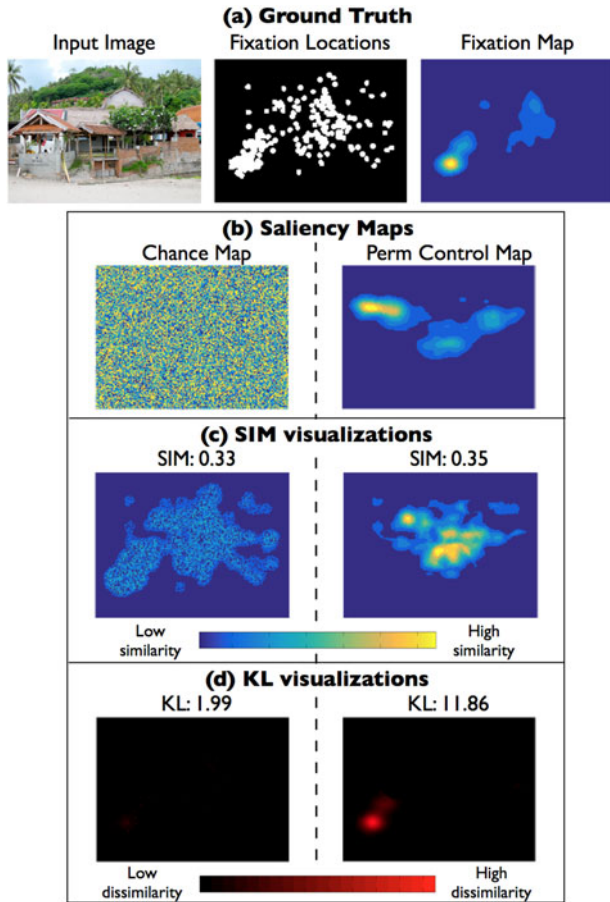


Fig. 11. The SIM and KL metrics measure the similarity between the saliency map (b) and ground truth fixation map (a), treating the former as the predicted distribution and the latter as the target distribution. SIM measures the histogram intersection between the distributions (Section 4.2.1), while KL measures an information-theoretic divergence between them (Section 4.2.3). KL is much more sensitive to false negatives than SIM. Both saliency maps in (b) are image-agnostic baselines, and receive similar scores under the SIM metric (c). However, because the map on the left places uniformly-sampled saliency values at all image pixels, it contains fewer zero values and is favored by KL (d). The rightmost map samples saliency from another image, resulting in zero values at multiple fixated locations and a poor KL score (d).

that a sparse set of predictions is penalized very harshly, significantly worse than chance.

4.2.4 Earth Mover’s Distance (EMD): Incorporating Spatial Distance into Evaluation

All the metrics discussed so far have no notion of how spatially far away the prediction is from the ground truth. Accordingly, any map that has no pixel overlap with the ground truth will receive the same score of zero,⁴ regardless of how predictions are distributed. Incorporating a measure of spatial distance can broaden comparisons, and allow for graceful degradation when the ground truth measurements have position error.

The Earth Mover’s Distance, *EMD*, measures the spatial distance between two probability distributions over a region. It was introduced as a spatially robust metric for image matching [59], [66]. Computationally, it is the minimum cost

4. Unless the model is properly regularized to compensate for uncertainty.

of morphing one distribution into the other. This is visualized in Fig. 9d where in blue are all the saliency map locations from which density needs to be moved, and in red are all the fixation map locations where density needs to be moved to. The total cost is the amount of density moved times the distance moved, and corresponds to brightness of the pixels in the visualization. It can be formulated as a transportation problem [17]. We used the following linear time variant of EMD [59]:

$$\widehat{EMD}(P, Q^D) = \min_{\{f_{ij}\}} \sum_{i,j} f_{ij} d_{ij} + \left| \sum_i P_i - \sum_j Q_j^D \right| \max_{i,j} d_{ij}$$

under the constraints:

$$(1) f_{ij} \geq 0 \quad (2) \sum_j f_{ij} \leq P_i \quad (3) \sum_i f_{ij} \leq Q_j^D, \quad (7)$$

$$(4) \sum_{i,j} f_{ij} = \min \left(\sum_i P_i, \sum_j Q_j^D \right),$$

where each f_{ij} represents the amount of density transported (or the *flow*) from the i th supply to the j th demand and d_{ij} is the *ground distance* between bin i and bin j in the distribution. Eq. (7) is therefore attempting to minimize the amount of density movement such that the total density is preserved after the move. Constraint (1) allows transporting density from P to Q^D and not vice versa. Constraint (2) prevents more density to be moved from a location P_i than is there. Constraint (3) prevents more density to be deposited to a location Q_j^D than is there. Constraint (4) is for feasibility: such that the amount of density moved does not exceed the total density found in either P or Q^D . Solving this problem requires global optimization on the whole map, making this metric quite computationally expensive.

A larger EMD indicates a larger difference between two distributions while an EMD of zero indicates that two distributions are the same. Generally, saliency maps that spread density over a larger area have larger EMD values (i.e., worse scores) as all the extra density has to be moved to match the ground truth map (Fig. 9). EMD penalizes false positives proportionally to the spatial distance they are from the ground truth (Section 5.2).

5 ANALYSIS OF METRIC BEHAVIOR

This section contains a set of experiments to study the behavior of 8 different evaluation metrics, where we systematically varied properties of the predicted maps to quantify the differential effects on metric scores. We focus on the metrics themselves, without assuming any optimization or regularization of the input maps. This most closely reflects how evaluation is carried out on the MIT Saliency Benchmark, which does not place any restrictions on the format of the submitted saliency maps. Without additional assumptions, our conclusions about the metrics should be informative for other applications, beyond saliency evaluation.

5.1 Scoring Baseline Models

Comparing metrics on a set of baselines can be illustrative of metric behavior and be used to uncover the properties of saliency maps that drive this behavior. In Table 3 we include the scores of 4 baseline models and an upper bound for each

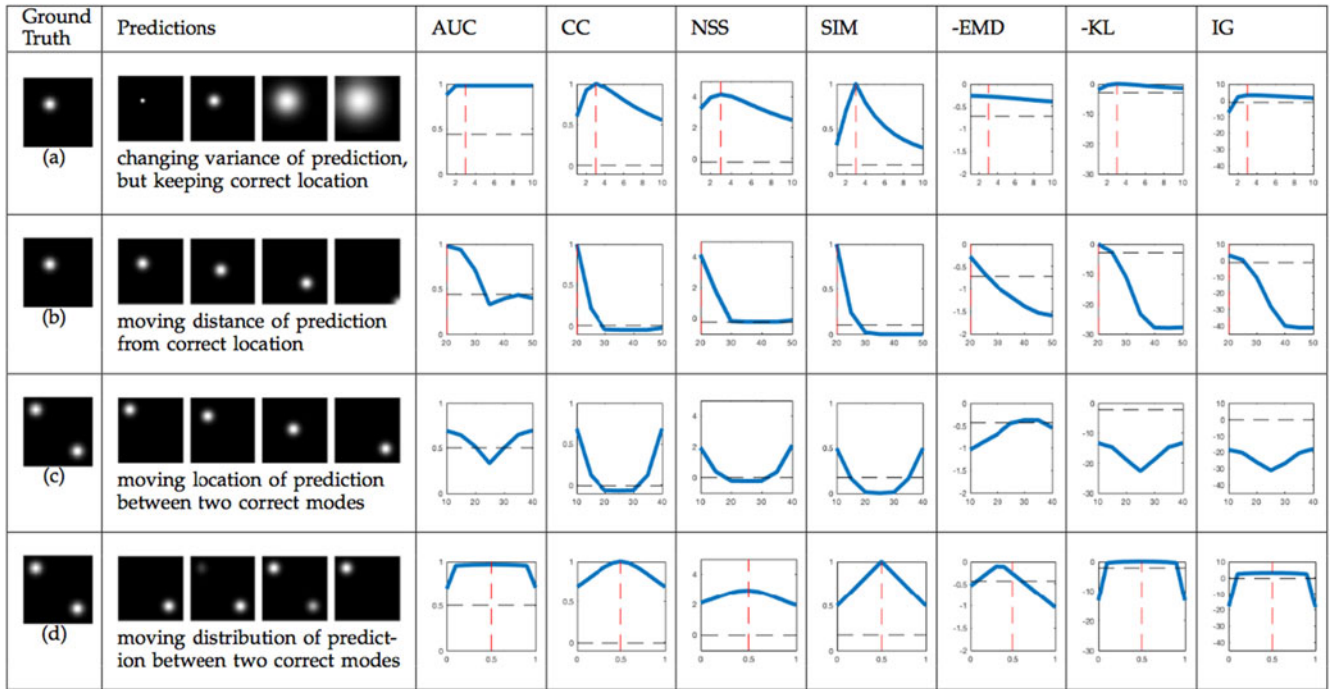


Fig. 12. We systematically varied parameters of a saliency map in order to quantify effects on metric scores. Each row corresponds to varying a single parameter value of the prediction: (a) Variance, (b-c) location, and (d) relative weight. The x -axis of each subplot spans the parameter range, with the dotted red line corresponding to the ground truth parameter setting (if applicable). The y -axis is different across metrics but constant for a given metric. The dotted black line is chance performance. EMD and KL y -axes have been flipped so a higher y -value indicates better performance across all subplots.

metric. The *center prior* model is a symmetric Gaussian stretched to the aspect ratio of the image, so each pixel's saliency value is a function of its distance from the center (higher saliency closer to center). Our *chance* model assigns a uniform value to each pixel in the image. An alternative chance model that also factors in the properties of a particular dataset is called a *permutation control*: it is computed by randomly selecting a fixation map from another image. It has the same image-independent properties as the ground truth fixation map for the image since it has been computed with the same blur and scale. The *single observer* model uses the fixation map from one observer to predict the fixations of the remaining observers (1 predicting $n - 1$). We repeated this leave-one-out procedure and averaged the results across all observers.

To compute an upper bound for each metric we measured how well the fixations of n observers predict the fixations of another group of n observers, varying n from 1 to 19 (half of the total 39 observers). Then we fit these prediction scores to a power function to obtain the limiting score of *infinite observers* (details available in the online supplemental material). This is useful to obtain dataset-specific bounds for metrics that are not otherwise bounded (i.e., NSS, EMD, KL, IG), and

to provide realistic bounds that factor in dataset-specific human consistency for metrics where the theoretical bound may not be reachable (i.e., AUC, sAUC).

There is a divergent behavior in the way the metrics score a center prior model relative to a single observer model. The center prior captures dataset-specific, image-independent properties; while the single observer model captures image-specific properties but might be missing properties that emulate average viewing behavior. In particular, the single observer model is quite sparse and so achieves worse scores according to the KL, IG, and SIM metrics.

Similarly, we compare the chance and permutation control models. Both are image-independent. However, the chance model is dataset-independent, while the permutation control model captures dataset-specific parameters. CC, NSS, AUC, and EMD scores are significantly higher for the permutation control, pointing to the importance under these metrics, of fitting parameters of a particular dataset (including center bias, blur, and scale). On the other hand, KL and IG are sensitive to insufficient regularization. As a result, the permutation control model, which has more zero values, performs worse than the chance model.

TABLE 3
Performance of Saliency Baselines (as Pictured in Fig. 2) with Scores Averaged Over MIT300 Benchmark Images

Saliency model	Similarity metrics					Dissimilarity metrics		
	SIM \uparrow	CC \uparrow	NSS \uparrow	AUC \uparrow	sAUC \uparrow	IG \uparrow	KL \downarrow	EMD \downarrow
Infinite Observers	1.00	1.00	3.29	0.92	0.81	2.50	0	0
Single Observer	0.38	0.53	1.65	0.80	0.64	-8.49	6.19	3.48
Center Prior	0.45	0.38	0.92	0.78	0.51	0	1.24	3.72
Permutation Control	0.34	0.20	0.49	0.68	0.50	-6.90	6.12	4.59
Chance	0.33	0.00	0.00	0.50	0.50	-1.24	2.09	6.35

TABLE 4
Properties of the 8 Evaluation Metrics in This Paper (Under Our Implementations)

	AUC	sAUC	SIM	CC	KL	IG	NSS	EMD
Implementation								
Bounded	✓	✓	✓	✓				
Location-based, parameter-free	✓	✓				✓	✓	
Local computations, differentiable			✓	✓	✓	✓	✓	
Symmetric			✓	✓				✓
Behavior								
Invariant to monotonic transformations	✓	✓						
Invariant to linear transformations (contrast)	✓	✓		✓			✓	
Requires special treatment of center bias		✓				✓		
Most affected by false negatives			✓		✓	✓		
Scales with spatial distance								✓

One possible meta-measure for selecting metrics for evaluation is how much better one baseline is over another (e.g., [22], [53], [61]). However, the optimal ranking of baselines is likely to be different across applications: in some cases, it may be useful to accurately capture systematic viewing behaviors if nothing else is known, while in another setting, specific points of interest are more relevant than viewing behaviors.

5.2 Treatment of False Positives and Negatives

Different metrics place different weights on the presence of false positives and negatives in the predicted saliency relative to the ground truth. To directly compare the extent to which metrics penalize false negatives, we performed a series of systematic tests. Starting with the ground truth fixation map, we progressively removed different amounts of salient pixels: pixels with a saliency value above the mean map value were selected uniformly at random and set to 0. We then evaluated the similarity of the resulting map to the original ground truth map and measured the drop in score with 25, 50, and 75 percent false negatives. To make comparison across metrics possible, we normalized this change in score by the score difference between the infinite observer limit and chance. We call this the *chance-normalized score*. For instance, for the AUC-Judd metric the upper limit is 0.92, chance is at 0.50, and the score with 75 percent false negatives is 0.67. The chance-normalized score is: $100\% \times (0.92 - 0.67)/(0.92 - 0.50) = 60\%$. Values for the other metrics are available in Table 5.

KL, IG, and SIM are Most Sensitive to False Negatives. If the prediction is close to zero where the ground truth has a non-zero value, the penalties can grow arbitrarily large under KL, IG, and SIM. These metrics penalize models with false negatives significantly more than false positives. In Table 5, KL and IG scores drop below chance levels with only 25 percent false negatives. Another way to look at this is that these metrics' sensitivity to regularization drives their evaluations of models. KL and IG scores will be low for sparse and poorly regularized models.

AUC Ignores Low-Valued False Positives. AUC scores depend on which level sets false positives fall in: false positives in the first level sets are penalized most, but those in the last level set do not have a large impact on performance. Models with many low-valued false positives (e.g., Fig. 7) do not incur large penalties. Saliency maps that place different amounts of density but at the correct (fixated) locations will receive similar AUC scores (Fig. 12d).

NSS and CC are Equally Affected by False Positives and Negatives. During the normalization step of NSS, a few false positives will be washed out by the other saliency values and will not significantly affect the saliency values at fixated locations. However, as the number of false positives increases, they begin to have a larger influence on the normalization calculation, driving the overall NSS score down.

By construction, CC treats false positives and negatives symmetrically. NSS is highly related to CC, as a discrete approximation to CC (see online supplemental material). NSS behavior will be very similar to CC, including the treatment of false positives and negatives.

EMD's Penalty Depends on Spatial Distance. EMD is least sensitive to uniformly-occurring false negatives (e.g., Table 5) because the EMD calculation can redistribute saliency values from nearby pixels to compensate. However, false negatives that are spatially far away from any predicted density are highly penalized. Similarly, EMD's penalty for false positives depends on their spatial distance to ground truth, where false positives close to ground truth locations can be redistributed to those locations at low cost, but distant false positives are highly penalized (Fig. 9).

5.3 Systematic Viewing Biases

Common to many images is a higher density of fixations in the center of the image compared to the periphery, a function of both photographer bias (i.e., centering the main

TABLE 5
Metrics Have Different Sensitivities to False Negatives

Map	EMD	CC	NSS	AUC	SIM	IG	KL
	↓	↑	↑	↑	↑	↑	↓
Orig	0.00 (0%)	1.00 (0%)	3.29 (0%)	0.92 (0%)	1.00 (0%)	2.50 (0%)	0.00 (0%)
-25%	0.13 (2%)	0.85 (15%)	2.66 (19%)	0.85 (17%)	0.78 (33%)	-1.78 (114%)	2.55 (122%)
-50%	0.16 (3%)	0.70 (30%)	2.18 (34%)	0.77 (36%)	0.59 (61%)	-6.35 (237%)	5.64 (270%)
-75%	1.09 (17%)	0.50 (50%)	1.57 (52%)	0.67 (60%)	0.45 (82%)	-10.65 (352%)	8.18 (391%)

We sorted these metrics in order of increasing sensitivity to 25, 50, and 75 percent false negatives, where EMD is least, and KL is most, sensitive. Scores are averaged over all MIT300 fixation maps. Below each score is the percentage drop in performance from the metric's limit, normalized by the percentage drop to chance level.

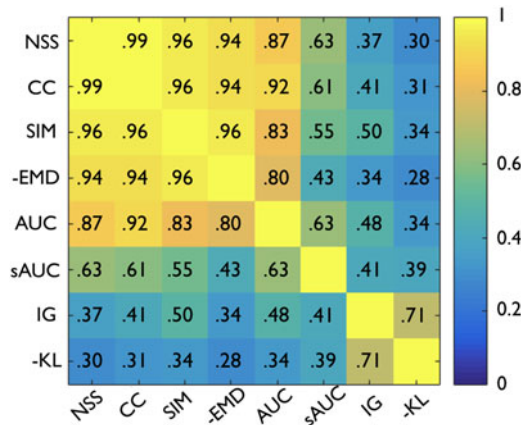


Fig. 13. We sort the saliency models on the MIT300 benchmark individually by each metric, and then compute the Spearman rank correlation between the model orderings of every pair of metrics (where 1 indicates perfect correlation). The first 5 metrics listed are highly correlated with each other. KL and IG are correlated with each other and uncorrelated with the other metrics, due to their high sensitivity to zero-valued predictions at fixated locations. The sAUC metric is also different from the others because it specifically penalizes models that have a center bias.

subject) and observer viewing biases. The effect of center bias on model evaluation has received much attention [11], [21], [44], [57], [62], [73], [74], [86]. In this section we discuss center bias in the context of the metrics in this paper.

sAUC Penalizes Models that Include Center Bias. The sAUC metric samples negatives from other images, which in the limit of many images corresponds to sampling negatives from a central Gaussian. For an image with a strong central viewing bias, both positives and negatives would be sampled from the same image region (Fig. 6), and a correct prediction would be near chance. The sAUC metric prefers models that do not explicitly incorporate center bias into their predictions. For a fair evaluation under sAUC, models need to operate under the same assumptions, or else their scores will be a function of whether or not they incorporate center bias.

IG Provides a Direct Comparison to Center Bias. Information gain over a center prior baseline provides a more intuitive way to interpret model performance relative to center bias. If a model can not explain fixation patterns on an image beyond systematic viewing biases, such a model will have no gain over a center prior.

EMD Spatially Hedges its Bets. The EMD metric prefers models that hedge their bets if all the ground truth locations can not be accurately predicted (Fig. 12c). For instance, if an image is fixated in multiple locations, EMD will favor a prediction that falls spatially between the fixated locations instead of one that captures a subset of the fixated locations (contrary to the behavior of the other metrics).

A center prior is a good approximation of average viewing behavior on images under laboratory conditions, where an image is projected for a few seconds on a computer screen in front of an observer [5]. A dataset-specific center prior emerges when averaging fixations over a large set of images. Knowing nothing else about image content, the center bias can act as a simple model prior. Overall if the goal is to predict natural viewing behavior on an image, center bias is part of the viewing behavior and discounting it entirely may be suboptimal. However, different metrics make different assumptions about the models: sAUC penalizes models

that include center bias, while IG expects center bias to already be optimized for. These differences in metric behaviors have led to differences in whether models include or exclude center bias. As a result, model rankings according to a particular metric can often be dominated by the differences in modeled center bias (Section 5.4).

5.4 Relationship Between Metrics

As saliency metrics are often used to rank saliency models, we can measure how correlated the rankings are across metrics. This analysis will indicate whether metrics favor or penalize similar behaviors in models. We sort model performances according to each metric and compute the Spearman rank correlation between the model orderings of every pair of metrics to obtain the correlation matrix in Fig. 13. The pairwise correlations between NSS, CC, AUC, EMD, and SIM range from 0.80 to 0.99. Because of these high correlations, we call this the *similarity cluster* of metrics. CC and NSS are most highly correlated due to their analogous computations, as are KL and IG (see online supplemental material).

Driven by extreme sensitivity to false negatives, KL and IG rank saliency models differently than the similarity cluster. Viewed another way, these metrics are worse behaved if saliency models are not properly regularized. For these metrics, a zero-valued prediction is interpreted as an impossibility of fixations at that location, while for the other metrics, a zero-valued prediction is treated as less salient. These metrics have a natural probabilistic interpretation and are appropriate in cases where missing any ground truth fixation locations should be highly penalized, such as for detection applications (Section 6.3). Changing the regularization constant ϵ in the metric computations (Eqs. (2) and (6)) or regularizing the saliency models prior to evaluation (as in [43]) can reduce score differences between KL, IG, and the similarity cluster.

Although EMD is the only metric that takes into account spatial distance, it nevertheless ranks saliency models similarly as the other similarity cluster metrics. This is likely the case for two reasons: (i) like the similarity cluster metrics, EMD is also center biased (Table 3, Section 5.3) and (ii) current model mistakes are often a cause of completely incorrect prediction rather than imprecise localization (note: as models continue to improve, this might change).

Shuffled AUC (sAUC) has low correlations with other metrics because it modifies how predictions at different spatial locations on the image are treated. A model with more central predictions will be ranked lower than a model with more peripheral predictions (Fig. 6). Shuffled AUC assumes center bias has not been modeled, and penalizes models where it has. For these reasons, sAUC has been disfavored by some evaluations [10], [46], [54]. An alternative is optimizing models to include a center bias [37], [38], [42], [43], [58], [86]. In this case, the metric can be ambivalent to any model or dataset biases.

Saliency metrics are much more correlated once models are optimized for center bias, blur, and scale [37], [42], [43]. As a result, the differences between the metrics in Fig. 13 are largely driven by how sensitive the metrics are to these model properties. It is therefore valuable to know if different models make similar modeling assumptions in order to interpret saliency rankings meaningfully across metrics.

5.5 Comparisons to Related Work

Riche et al. [63] correlated metric scores on another saliency dataset and found that KL and sAUC are most different from the other metrics, including AUC, CC, NSS, and SIM, which formed a single cluster. We can explain this finding, since KL and sAUC make stronger assumptions about saliency models: KL assumes saliency models have sufficient regularization (otherwise false negatives are severely penalized) and sAUC assumes the model does not have a built-in center bias. These results explain the divergent rankings across metrics, as commonly evaluated saliency models may not share the same modeling assumptions.

Emami and Hoberock [22] used human consistency to compare 9 metrics. In discriminating between human saliency maps and random saliency maps, they found that NSS and CC were the best, and KL the worst. This is similar to the analysis in Section 5.1.

Li et al. [46] used crowd-sourced experiments to measure which metric best corresponds to human perception. The authors noted that human perception was driven by the most salient locations, the compactness of salient locations (i.e., low false positives), and a similar number of salient regions as the ground truth. The perception-based ranking most closely matched that of NSS, CC, and SIM, and was furthest from KL and EMD. However, the properties that drive human perception could be different than the properties desired for other applications of saliency. For instance, for evaluating probabilistic saliency maps, proper regularization and the scale of the saliency values (including very small values) can significantly affect evaluation. For these cases, perception-based metrics might not be as meaningful.

We propose that the assumptions underlying different models and metrics be considered more carefully, and that the different metric behaviors and properties enter into the decision of which metrics to use for evaluation (Table 4).

6 RECOMMENDATIONS FOR DESIGNING A SALIENCY BENCHMARK

Saliency models have evolved significantly since the seminal IttiKoch model [41], [78] and the original notions of saliency. Evaluation procedures, saliency datasets, and benchmarks have adapted accordingly. Given how many different metrics and models have emerged, it is becoming increasingly necessary to systematize definitions and evaluation procedures to make sense of the vast amount of new data and results [11]. The MIT Saliency Benchmark is a product of this evolution of saliency modeling; an attempt to capture the latest developments in models and metrics. However, as saliency continues to evolve, larger more specialized datasets may become necessary. Based on our experience with the MIT Saliency Benchmark, we provide some recommendations for future saliency benchmarks.

6.1 Defining Expected Input

As observed in the previous section, some of the inconsistencies in how metrics rank models are due to differing assumptions that saliency models make. This problem has been emphasized by Kümmerer et al. [42], [43], who argued that if models were explicitly designed and submitted as *probabilistic models*, then some ambiguities in evaluation would

disappear. For instance, a probability of zero in a probabilistic saliency map assumes that a fixation in a region is impossible; under alternative definitions, a value of zero might only mean that a fixation in a particular region is less likely. Metrics like KL, IG, and SIM are particularly sensitive to zero values, so models evaluated under these metrics would benefit from being regularized and optimized for scale. Similarly, knowing whether evaluation will be performed with a metric like sAUC should affect whether center bias is modeled, because this design decision would be penalized under this metric. A saliency benchmark should specify what definition of saliency is assumed, what kind of saliency map input is expected, and how models will be evaluated. The online supplemental material includes additional considerations.

6.2 Handling Dataset Bias

In saliency datasets, dataset bias occurs when there are systematic properties in the ground-truth data that are dataset-specific but image-independent. Most eye-tracking datasets have been shown to be center biased, containing a larger number of fixations near the image center, across different image types, videos, and even observer tasks [7], [8], [14], [16], [33], [36]. Center bias is a function of multiple factors, including photographer bias and observer bias, due to the viewing of fixed images in a laboratory setting [73], [83]. As a result, some models have a built-in center bias (e.g., Judd [38]), some metrics penalize center bias (e.g., sAUC), and some benchmarks optimize models with center bias prior to evaluation (e.g., LSUN [31]). These different approaches result from a disagreement in where systematic biases should be handled: at the level of the dataset, model, or evaluation. For transparency, saliency benchmarks should specify whether the submitted models are expected to incorporate center bias, or if dataset-specific center bias will be accounted for and subtracted during evaluation. In the former case, the benchmark can provide a training dataset on which to optimize center bias and other image-independent properties of the ground truth dataset (e.g., blur, scale, regularization), or else share these parameters directly.

The MIT Saliency Benchmark provides the MIT1003 dataset [38] as a training set to optimize center bias and blur parameters, and for histogram matching (scale regularization).⁵ Both MIT300 and MIT1003 have been collected using the same eye tracker setup, so the ground truth fixation data should have similar distribution characteristics, and parameter choices should generalize across these datasets.

The first saliency models were not designed with these considerations in mind, so when compared to models that had incorporated center bias and other properties into saliency predictions, the original models were at a disadvantage. However, the availability of saliency datasets has increased, and many benchmarks provide training data from which systematic parameters can be learned [12], [31], [34]. Many modern saliency models are a result of this data-driven approach. Over the last few years, we have seen fewer differences across saliency models in terms of scale, blur, and center bias [12].

5. Associated code is provided at https://github.com/cvzoya/saliency/tree/master/code_forOptimization.

6.3 Defining a Task for Evaluation

Saliency models are often designed to predict general task-free saliency, assigning a value of saliency or importance to each image pixel, largely independent of the end application. Saliency is often motivated as a useful representation for image processing applications such as image re-targeting, compression, and transmission, object and motion detection, and image retrieval and matching [6], [35]. However, if the end goal is one of these applications, then it might be easier to directly train a saliency model for the relevant task, rather than for task-free fixation prediction. Task-based, or application-specific, saliency prediction is not yet very common. Relevant datasets and benchmarks are yet to be designed. Evaluating saliency models on specific applications requires choosing metrics that are appropriate to the underlying task assumptions and expected input.

Consider a detection application of saliency such as object and motion detection, surveillance, localization and mapping, and segmentation [1], [15], [26], [27], [40], [47], [56], [84]. For such an application, a saliency model may be expected to produce a probability density of possible object locations, and be highly penalized if a target is missed. For this kind of probabilistic target detection, AUC, KL, and IG would be appropriate. EMD might be useful if some location invariance is permitted.

Applications including adaptive image and video compression and progressive transmission [28], [32], [51], [81], thumbnailing [52], [71], content-aware image re-targeting and cropping [3], [4], [64], [67], [79], rendering and visualization [39], [49], collage [29], [80] and artistic rendering [18], [38] require ranking (by importance or saliency) different image regions. For these applications, when it is valuable to know how much more salient a given image region is than another, an evaluation metric like AUC (that is ambivalent to monotonic transformations of the input map) is not appropriate. Instead, NSS or SIM would provide a more useful evaluation.

6.4 Selecting Metrics for Evaluation

This paper showed how metrics behave under different conditions. This can help guide the selection of metrics for saliency benchmarks, depending on the assumptions made (e.g., whether the models are probabilistic, whether center bias is included, etc.). Future benchmarks should make explicit any assumptions along with the expected saliency map format.

The MIT Saliency Benchmark assumes that all fixation behavior, including any systematic dataset biases, are accounted for by the saliency model. Capturing viewing biases is part of the modeling requirements. Metrics like shuffled AUC will downrank models that have a strong center bias. Saliency models submitted are not necessarily probabilistic, so they might be unfairly evaluated by the KL, IG, and SIM metrics that penalize zero values (false negatives), unless they are first regularized and pre-processed as in Kümmerer et al. [43]. AUC, which is ambivalent to monotonic transformations, has begun to saturate on the MIT Saliency Benchmark and is becoming less capable of discriminating between different saliency models [13]. The Earth Mover's Distance (EMD) is computationally expensive to compute and difficult to optimize for.

Take-Aways: For a benchmark operating under the same assumptions as the MIT Saliency Benchmark, we recommend reporting either CC or NSS. Both make limited assumptions about input format, and treat false positives and negatives symmetrically. For a benchmark intended to evaluate saliency maps as probability distributions, IG and KL would be good choices; IG specifically measures prediction performance beyond systematic dataset biases.

6.5 Looking Ahead

The top models on the MIT Saliency Benchmark are currently all neural networks, trained on large datasets of attention data [34], and often pre-trained on even larger datasets of natural scenes and objects [19], [87]. Because largely the same datasets are used for training, these models adopt similar biases. As a result, differences in modeling assumptions (e.g., whether or not to include center bias) are becoming less pronounced. Models are picking up on the same salient image regions, particularly objects and faces, and making the same mistakes, such as missing objects of gaze and action [13]. Metrics like AUC are becoming less capable of discriminating between these models, and are saturating on the MIT Saliency Benchmark. The difference between models is becoming more fine-grained: e.g., in how different models assign importance values to different regions of the image. For instance, although the top models can detect all the faces in an image, how can we differentiate between the models that correctly weight the saliency across the faces? These finer-grained differences are easier to measure using metrics like NSS. Moreover, these advances call for larger benchmarks, more specific tasks, and finer-grained evaluation procedures (as discussed in [13]).

7 CONCLUSION

We provided an analysis of the behavior of 8 evaluation metrics to make sense of the differences in saliency model rankings. Properties of the inputs affect metric scores differently: how the ground truth is represented, whether the prediction includes dataset bias, whether the inputs are probabilistic, whether spatial deviations exist between the prediction and ground truth. Knowing how these properties affect metrics, and which properties are most important for a given application can help with metric selection. Other considerations include whether the metric computations are expensive, local, and differentiable, which would influence whether a metric is appropriate for model optimization. Take-aways about the metrics are included in Table 6.

We considered saliency metrics from the perspective of the MIT Saliency Benchmark, which does not assume that saliency models are probabilistic, but does assume that all systematic dataset biases (including center bias, blur, scale) are taken care of by the model. Under these assumptions we found that the Normalized Scanpath Saliency (NSS) and Pearson's Correlation Coefficient (CC) metrics provide the fairest comparison. Closely related mathematically, their rankings of saliency models are highly correlated, and reporting performance using one of them is sufficient.

However, under alternative assumptions and definitions of saliency, another choice of metrics may be more appropriate. Specifically, if saliency models are evaluated as probabilistic models, then KL-divergence and Information Gain (IG) are recommended. Arguments for why it might be

TABLE 6
A Brief Overview of the Metric Analyses and Discussions Provided in This Paper, Highlighting Some of the Key Properties, Features, and Applications of Different Evaluation Metrics

Metric	Quick take-aways
Area under ROC Curve (AUC)	Historically the most commonly-used metric for saliency evaluation. Invariant to monotonic transformations. Driven by high-valued predictions and largely ambivalent of low-valued false positives. Currently saturating on standard saliency benchmarks [12], [13]. Good for detection applications.
Shuffled AUC (sAUC)	A version of AUC that compensates for dataset bias by scoring a center prior at chance. Most appropriate in evaluation settings where the saliency model is not expected to account for center bias. Otherwise, has similar properties to AUC.
Similarity (SIM)	An easy and fast similarity computation between histograms. Assumes the inputs are valid distributions. More sensitive to false negatives than false positives.
Pearson's Correlation Coefficient (CC)	A linear correlation between the prediction and ground truth distributions. Treats false positives and false negatives symmetrically.
Normalized Scanpath Saliency (NSS)	A discrete approximation of CC that is additionally parameter-free (operates on raw fixation locations). Recommended for saliency evaluation.
Earth Mover's Distance (EMD)	The only metric considered that scales with spatial distance. Can provide a finer-grained comparison between saliency maps. Most computationally intensive, non-local, hard to optimize.
Kullback-Leibler divergence (KL)	Has a natural interpretation where goal is to approximate a target distribution. Assumes input is a valid probability distribution with sufficient regularization. Misdetections are highly penalized.
Information Gain (IG)	A new metric introduced by [42], [43]. Assumes input is a valid probability distribution with sufficient regularization. Measures the ability of a model to make predictions above a baseline model of center bias. Otherwise, has similar properties to KL.

preferable to define and evaluate saliency models probabilistically can be found in [42], [43]. Specific tasks and applications may also call for a different choice of metrics. For instance, AUC, KL, and IG are appropriate for detection applications, as they penalize target detection failures. In applications where it is important to evaluate the relative importance of different image regions, such as for image-retargeting, compression, and progressive transmission, metrics like NSS or SIM are a better fit.

In this paper we discussed the influence of different assumptions on the choice of appropriate metrics. We provided recommendations for new saliency benchmarks, such that if designed with explicit assumptions from the start, evaluation can be more transparent and reduce confusion in saliency evaluation.

We also provide code for evaluating and visualizing the metric computations⁶ to add further transparency to model evaluation and to allow researchers a finer-grained look into metric computations, to debug saliency models and visualize the aspects of saliency models driving or hurting performance.

ACKNOWLEDGMENTS

The authors would like to thank Matthias Kümmerer and other attendees of the saliency tutorial at ECCV 2016 for helpful discussions about saliency evaluation.⁷ Thank you also to the anonymous reviewers for many detailed suggestions. ZB was supported by a postgraduate scholarship (PGS-D) from the Natural Sciences and Engineering Research Council of Canada. Support to AO, AT, and FD was provided by the Toyota Research Institute / MIT CSAIL Joint Research Center.

6. <https://github.com/cvzoya/saliency>

7. http://saliency.mit.edu/ECCVTutorial/ECCV_saliency.htm

REFERENCES

- [1] R. Achanta, F. Estrada, P. Wils, and S. Susstrunk, "Salient region detection and segmentation," in *Proc. Int. Conf. Comput. Vis. Syst.*, 2008, pp. 66–75.
- [2] R. Achanta, S. Hemami, F. Estrada, and S. Susstrunk, "Frequency-tuned salient region detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 1597–1604.
- [3] R. Achanta and S. Süssstrunk, "Saliency detection for content-aware image resizing," in *Proc. 16th IEEE Int. Conf. Image Process.*, 2009, pp. 1005–1008.
- [4] S. Avidan and A. Shamir, "Seam carving for content-aware image resizing," *ACM Trans. Graph.*, vol. 26, no. 3, 2007, Art. no. 10.
- [5] M. Bindermann, "Scene and screen center bias early eye movements in scene viewing," *Vis. Res.*, vol. 50, pp. 2577–2587, 2010.
- [6] A. Borji and L. Itti, "State-of-the-art in visual attention modeling," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 1, pp. 185–207, Jan. 2013.
- [7] A. Borji, D. N. Sihite, and L. Itti, "Quantitative analysis of human-model agreement in visual saliency modeling: A comparative study," *IEEE Trans. Image Process.*, vol. 22, no. 1, pp. 55–69, Jan. 2012.
- [8] A. Borji, H. R. Tavakoli, D. N. Sihite, and L. Itti, "Analysis of scores, datasets, and models in visual saliency prediction," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2013, pp. 921–928.
- [9] N. D. B. Bruce and J. K. Tsotsos, "Saliency, attention, and visual search: An information theoretic approach," *J. Vis.*, vol. 9, no. 3, 2009, pp. 1–24.
- [10] N. D. B. Bruce, C. Wloka, N. Frosst, S. Rahman, and J. K. Tsotsos, "On computational modeling of visual saliency: Examining what's right, and what's left," *Vis. Res.*, vol. 116, pp. 95–112, 2015.
- [11] Z. Bylinskii, E. M. DeGennaro, R. Rajalingham, H. Ruda, J. Zhang, and J. K. Tsotsos, "Towards the quantitative evaluation of visual attention models," *Vis. Res.*, vol. 116, pp. 258–268, 2015.
- [12] Z. Bylinskii, T. Judd, A. Borji, L. Itti, F. Durand, A. Oliva, and A. Torralba, MIT Saliency Benchmark. (2014). [Online]. Available: <http://saliency.mit.edu/>
- [13] Z. Bylinskii, A. Recasens, A. Borji, A. Oliva, A. Torralba, and F. Durand, "Where should saliency models look next?" in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 809–824.
- [14] R. L. Canosa, J. Pelz, N. R. Mennie, and J. Peak, "High-level aspects of oculomotor control during viewing of natural-task images," in *Proc. SPIE*, 2003, pp. 240–251.

- [15] C.-K. Chang, C. Siagian, and L. Itti, "Mobile robot vision navigation & localization using gist and saliency," in *Proc. IEEE/RISJ Int. Conf. Intell. Robots Syst.*, Oct. 2010, pp. 4147–4154.
- [16] A. D. F. Clarke and B. W. Tatler, "Deriving an appropriate baseline for describing fixation behaviour," *Vis. Res.*, vol. 102, pp. 41–51, 2014.
- [17] G. B. Dantzig, "Application of the simplex method to a transportation problem," *Activity Anal. Prod. Allocation*, T.C. Koopmans. Publishers: John Wiley & Sons, Inc., pp. 359–373, 1951.
- [18] D. DeCarlo and A. Santella, "Stylization and abstraction of photographs," *ACM Trans. Graph.*, vol. 21, no. 3, pp. 769–776, 2002.
- [19] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 248–255.
- [20] W. Einhäuser and P. König, "Does luminance-contrast contribute to a saliency for overt visual attention?" *Eur. J. Neuroscience*, vol. 17, pp. 1089–1097, 2003.
- [21] W. Einhäuser, M. Spain, and P. Perona, "Objects predict fixations better than early saliency," *J. Vis.*, vol. 8, no. 14, pp. 1–26, 2008.
- [22] M. Emami and L. L. Hoberock, "Selection of a best metric and evaluation of bottom-up visual saliency models," *Image Vis. Comput.*, vol. 31, no. 10, pp. 796–808, 2013.
- [23] U. Engelke, H. Liu, J. Wang, P. Le Callet, I. Heynderickx, H.-J. Zepernick, and A. Maeder, "Comparative study of fixation density maps," *IEEE Trans. Image Process.*, vol. 22, no. 3, pp. 1121–1133, Mar. 2013.
- [24] E. Erdem and A. Erdem, "Visual saliency estimation by nonlinearly integrating features using region covariances," *J. Vis.*, vol. 13, no. 4, pp. 1–20, 2013.
- [25] T. Fawcett, "An introduction to ROC analysis," *Pattern Recognit. Lett.*, vol. 27, no. 8, pp. 861–874, 2006.
- [26] S. Frintrop, "General object tracking with a component-based target descriptor," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2010, pp. 4531–4536.
- [27] S. Frintrop, P. Jensfelt, and H. Christensen, "Simultaneous robot localization and mapping based on a visual attention system," in *Attention in Cognitive Systems. Theories and Systems from an Interdisciplinary Viewpoint*, pp. 417–430, Berlin, Germany: Springer, 2007.
- [28] W. S. Geisler and J. S. Perry, "A real-time foveated multiresolution system for low-bandwidth video communication," *Proc. SPIE: Human Vis. Electron. Imag.*, vol. 3299, pp. 294–305, 1998.
- [29] S. Goferman, A. Tal, and L. Zelnik-Manor, "Puzzle-like collage," *Comput. Graph. Forum*, vol. 29, no. 2, pp. 459–468, 2010.
- [30] D. M. Green and J. A. Swets, *Signal Detection Theory and Psychophysics*, Hoboken, NJ, USA: Wiley, 1966.
- [31] Princeton Vision Group, NUS VIP Lab, and Bethge Lab, Large-scale scene understanding challenge: Saliency prediction. *Technical report*, 2016. [Online]. Available: lsun.cs.princeton.edu/challenge/2016/saliency/saliency.pdf
- [32] L. Itti, "Automatic foveation for video compression using a neurobiological model of visual attention," *IEEE Trans. Image Process.*, vol. 13, no. 10, pp. 1304–1318, Oct. 2004.
- [33] L. Itti and P. F. Baldi, "Bayesian surprise attracts human attention," *Adv. Neural Inform. Process. Syst.*, vol. 18, pp. 547–554, 2006.
- [34] M. Jiang, S. Huang, J. Duan, and Q. Zhao, "Salicon: Saliency in context," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 1072–1080.
- [35] T. Judd, "Understanding and Predicting Where People Look in Images," PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 2011.
- [36] T. Judd, F. Durand, and A. Torralba, "Fixations on low-resolution images," *J. Vis.*, vol. 11, no. 4, 2011.
- [37] T. Judd, F. Durand, and A. Torralba, "A benchmark of computational models of saliency to predict human fixations," in MIT, Tech. Rep. MIT-CSAIL-TR-2012-001, 2012, <http://hdl.handle.net/1721.1/68590>
- [38] T. Judd, K. Ehinger, F. Durand, and A. Torralba, "Learning to predict where humans look," in *Proc. IEEE 12th Int. Conf. Comput. Vis.*, 2009, pp. 2106–2113.
- [39] Y. Kim and A. Varshney, "Saliency-guided enhancement for volume visualization," *IEEE Trans. Vis. Comput. Graph.*, vol. 12, no. 5, pp. 925–932, Sep./Oct. 2006.
- [40] D. Klein, S. Frintrop, et al., "Center-surround divergence of feature statistics for salient object detection," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2011, pp. 2214–2219.
- [41] C. Koch and S. Ullman, "Shifts in selective visual attention: towards the underlying neural circuitry," *Human Neurobiology*, vol. 4, pp. 219–227, 1985.
- [42] M. Kümmerer, T. Wallis, and M. Bethge, "How close are we to understanding image-based saliency?" arXiv preprint arXiv:1409.7686, 2014.
- [43] M. Kümmerer, T. Wallis, and M. Bethge, "Information-theoretic model comparison unifies saliency metrics," *Proc. Nat. Acad. Sci. United States America*, vol. 112, no. 52, 2015, pp. 16054–16059.
- [44] O. Le Meur, P. Le Callet, D. Barba, and D. Thoreau, "A coherent computational approach to model bottom-up visual attention," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 5, pp. 802–817, May 2006.
- [45] J. Li, M. D. Levine, X. An, X. Xu, and H. He, "Visual saliency based on scale-space analysis in the frequency domain," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 4, pp. 996–1010, Apr. 2013.
- [46] J. Li, C. Xia, Y. Song, S. Fang, and X. Chen, "A data-driven metric for comprehensive evaluation of saliency models," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 190–198.
- [47] T. Liu, J. Sun, N.-N. Zheng, X. Tang, and H.-Y. Shum, "Learning to detect a salient object," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2007, pp. 1–8.
- [48] Y. Liu, D. Zhang, G. Lu, and W.-Y. Ma, "A survey of content-based image retrieval with high-level semantics," *Pattern Recognit.*, vol. 40, no. 1, pp. 262–282, 2007.
- [49] P. Longhurst, K. Debattista, and A. Chalmers, "A GPU based saliency map for high-fidelity selective rendering," in *Proc. 4th Int. Conf. Comput. Graph. Virtual Reality Vis. Interaction Africa*, 2006, pp. 21–29.
- [50] F. Lopez-Garcia, X. Ramon Fdez-Vidal, X. Manuel Pardo, and R. Dosić, "Scene recognition through visual attention and image features: A comparison between sift and surf approaches," T. P. Cao, ed., in *Object Recognition*, Rijeka, Croatia: InTech, pp. 185–200, 2011.
- [51] Y.-F. Ma, X.-S. Hua, L. Lu, and H.-J. Zhan, "A generic framework of user attention model and its application in video summarization," *IEEE Trans. Multimedia*, vol. 7, no. 5, pp. 907–919, Oct. 2005.
- [52] L. Marchesotti, C. Cifarelli, and G. Csurka, "A framework for visual saliency detection with applications to image thumbnailing," in *Proc. IEEE 12th Int. Conf. Comput. Vis.*, 2009, pp. 2232–2239.
- [53] R. Margolin, L. Zelnik-Manor, and A. Tal, "How to evaluate foreground maps?" in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 248–255.
- [54] O. Le Meur and T. Baccino, "Methods for comparing scanpaths and saliency maps: Strengths and weaknesses," *Behavioral Res. Methods*, vol. 45, no. 1, pp. 251–266, 2013.
- [55] O. Le Meur, P. Le Callet, and D. Barba, "Predicting visual fixations on video based on low-level visual features," *Vis. Res.*, vol. 47, no. 19, pp. 2483–2498, 2007.
- [56] V. Navalpakkam and L. Itti, "An integrated model of top-down and bottom-up attention for optimizing detection speed," in *Proc. IEEE Comput. Society Conf. Comput. Vis. Pattern Recognit.*, 2006, vol. 2, pp. 2049–2056.
- [57] D. Parkhurst, K. Law, and E. Niebur, "Modeling the role of salience in the allocation of overt visual attention," *Vis. Res.*, vol. 42, no. 1, pp. 107–123, 2002.
- [58] D. J. Parkhurst and E. Niebur, "Scene content selected by active vision," *Spatial Vis.*, vol. 16, no. 30, pp. 125–154, 2003.
- [59] O. Pele and M. Werman, "A linear time histogram metric for improved sift matching," in *Proc. Eur. Conf. Comput. Vis.*, 2008, pp. 495–508.
- [60] R. J. Peters, A. Iyer, L. Itti, and C. Koch, "Components of bottom-up gaze allocation in natural images," *Vis. Res.*, vol. 45, no. 18, pp. 2397–2416, 2005.
- [61] J. Pont-Tuset and F. Marques, "Supervised evaluation of image segmentation and object proposal techniques," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 7, pp. 1465–1478, Jul. 2016.
- [62] L. W. Renninger, J. Coughlan, P. Verghese, and J. Malik, "An information maximization model of eye movements," in *Proc. Adv. Neural Inform. Process. Syst.*, pp. 1121–1128, 2004.
- [63] N. Riche, M. Duvinage, M. Mancas, B. Gosselin, and T. Dutoit, "Saliency and human fixations: State-of-the-art and study of comparison metrics," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2013, pp. 1153–1160.
- [64] M. Rubinstein, A. Shamir, and S. Avidan, "Improved seam carving for video retargeting," in *Proc. ACM SIGGRAPH Papers*, 2008, Art. no. 16.
- [65] Y. Rubner and C. Tomasi, *Perceptual Metrics for Image Database Navigation*, New York, NY, USA: Springer, 2001.
- [66] Y. Rubner, C. Tomasi, and L. J. Guibas, "The earth movers distance as a metric for image retrieval," *Int. J. Comput. Vis.*, vol. 40, pp. 99–121, 2000.
- [67] A. Santella, M. Agrawala, D. DeCarlo, D. Salesin, and M. Cohen, "Gaze-based interaction for semi-automatic photo cropping," in *Proc. SIGCHI Conf. Human Factors Comput. Syst.*, 2006, pp. 771–780.
- [68] H. J. Seo and P. Milanfar, "Static and space-time visual saliency detection by self-resemblance," *J. Vis.*, vol. 9, no. 12, pp. 1–27, 2009.

- [69] A. K. Sinha and K. K. Shukla, "A study of distance metrics in histogram based image retrieval," *Int. J. Comput. Technol.*, vol. 4, no. 3, pp. 821–830, 2013.
- [70] A. W. M. Smeulders, M. Worring, S. Santini, A. Gupta, and R. Jain, "Content-based image retrieval at the end of the early years," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 12, pp. 1349–1380, Dec. 2000.
- [71] B. Suh, H. Ling, B. B. Bederson, and D. W. Jacobs, "Automatic thumbnail cropping and its effectiveness," in *Proc. 16th Annu. ACM Symp. User Interface Softw. Technol.*, 2003, pp. 95–104.
- [72] M. J. Swain and D. H. Ballard, "Color indexing," *Int. J. Comput. Vis.*, vol. 7, no. 1, pp. 11–32, 1991.
- [73] B. W. Tatler, "The central fixation bias in scene viewing: Selecting an optimal viewing position independently of motor biases and image feature distributions," *J. Vis.*, vol. 7, no. 14, pp. 1–17, 2007.
- [74] B. W. Tatler, R. J. Baddeley, and I. D. Gilchrist, "Visual correlates of fixation selection: effects of scale and time," *Vis. Res.*, vol. 45, no. 5, pp. 643–659, 2005.
- [75] A. Torralba, A. Oliva, M. S. Castelhan, and J. M. Henderson, "Contextual guidance of eye movements and attention in real-world scenes: The role of global features in object search," *Psychological Rev.*, vol. 113, no. 4, pp. 766–786, Oct. 2006.
- [76] P.-H. Tseng, R. Carmi, I. G. M. Cameron, D. P. Munoz, and L. Itti, "Quantifying center bias of observers in free viewing of dynamic natural scenes," *J. Vis.*, vol. 9, no. 7, 2009, Art. no. 4.
- [77] N. Vasconcelos, "On the efficient evaluation of probabilistic similarity functions for image retrieval," *IEEE Trans. Inf. Theory*, vol. 50, no. 7, pp. 1482–1496, 2004.
- [78] D. Walther and C. Koch, "Modeling attention to salient proto-objects," *Neural Netw.*, vol. 19, no. 9, pp. 1395–1407, 2006.
- [79] D. Wang, G. Li, W. Jia, and X. Luo, "Saliency-driven scaling optimization for image retargeting," *Vis. Comput.*, vol. 27, no. 9, pp. 853–860, 2011.
- [80] J. Wang, L. Quan, J. Sun, X. Tang, and H-Y Shum, "Picture collage," in *Proc. IEEE Comput. Society Conf. Comput. Vis. Pattern Recognit.*, pp. 347–354, vol. 1, 2006.
- [81] Z. Wang, L. Lu, and A. C. Bovik, "Foveation scalable video coding with automatic fixation selection," *IEEE Trans. Image Process.*, vol. 12, no. 2, pp. 243–254, Feb. 2003.
- [82] N. Wilming, T. Betz, T. C. Kietzmann, and P. König, "Measures and limits of models of fixation selection," *PLoS One*, vol. 6, 2011, Art. no. e24038.
- [83] C. Wloka and J. K. Tsotsos, "Overt fixations reflect a natural central bias," *J. Vis.*, vol. 13, no. 9, pp. 239–239, 2013.
- [84] K. Yun, Y. Peng, D. Samaras, G. J. Zelinsky, and T. Berg, "Studying relationships between human gaze, description, and computer vision," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2013, pp. 739–746.
- [85] L. Zhang, M. H. Tong, T. K. Marks, H. Shan, and G. W. Cottrell, "Sun: A bayesian framework for saliency using natural statistics," *J. Vis.*, vol. 8, no. 7, pp. 1–20, 2008.
- [86] Q. Zhao and C. Koch, "Learning a saliency map using fixated locations in natural scenes," *J. Vis.*, vol. 11, no. 3, pp. 1–15, 2011.
- [87] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva, "Learning deep features for scene recognition using places database," in *Proc. 27th Int. Conf. Neural Inf. Process. Syst.*, 2014, pp. 487–495.
- [88] B. Zitova and J. Flusser, "Image registration methods: a survey," *Image Vis. Comput.*, vol. 21, no. 11, pp. 977–1000, 2003.



Zoya Bylinskii received the BS degree in computer science and statistics from the University of Toronto, and the MS degree in computer science from MIT under the supervision of Antonio Torralba and Aude Oliva. She is working toward the PhD degree in computer science in the Massachusetts Institute of Technology, advised by Frédo Durand and Aude Oliva. She works on topics with the interface of human and computer vision, on computational perception and cognition. She is an Adobe research fellow (2016), a recipient of the Natural

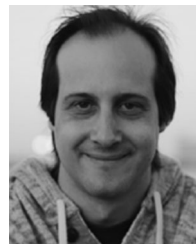
Sciences and Engineering Research Council of Canada's (NSERC) post-graduate doctoral award (2014-6), Julie Payette NSERC Scholarship (2013), and a finalist for Google's Anita Borg Memorial Scholarship (2011).



Tilke Judd received the BS degree in mathematics from the Massachusetts Institute of Technology (MIT), and the MS and PhD degrees in computer science from MIT, in 2007 and 2011, respectively, supervised by Frédo Durand and Antonio Torralba. She is a product manager with the Google in Zurich. During the summers of 2007 and 2009 she was an intern with Google and Industrial Light and Magic respectively. She was awarded a National Science Foundation Fellowship for 2005-2008 and a Xerox Graduate Fellowship in 2008.



Aude Oliva received the French baccalaureate degree in physics and mathematics, the BSc degree in psychology, and the two MSc degrees and PhD degree in cognitive science from the Institut National Polytechnique of Grenoble, France. She is a principal research scientist with the MIT Computer Science and Artificial Intelligence Laboratory (CSAIL). She joined the MIT faculty in the Department of Brain and Cognitive Sciences in 2004 and CSAIL in 2012. Her research on vision and memory is cross-disciplinary, spanning human perception and cognition, computer vision, and human neuroscience. She is the recipient of a National Science Foundation CAREER Award in Computational Neuroscience (2006), the Guggenheim fellowship in Computer Science (2014), and the Vannevar Bush faculty fellowship in Cognitive Neuroscience (2016).



Antonio Torralba received the degree in telecommunications engineering from Telecom BCN, Spain, in 1994 and the PhD degree in signal, image, and speech processing from the Institut National Polytechnique de Grenoble, France, in 2000. He is a professor of electrical engineering and computer science with the Massachusetts Institute of Technology. From 2000 to 2005, he spent postdoctoral training with the Brain and Cognitive Science Department and the Computer Science and Artificial Intelligence Laboratory, MIT. He received the 2008 National Science Foundation (NSF) Career award, the best student paper award at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) in 2009, and the 2010 J. K. Aggarwal Prize from the International Association for Pattern Recognition (IAPR).



Frédo Durand received the PhD degree from Grenoble University, France, in 1999, supervised by Claude Puech and George Drettakis. He is a professor in electrical engineering and computer science with the Massachusetts Institute of Technology, and a member of the Computer Science and Artificial Intelligence Laboratory (CSAIL). From 1999 till 2002, he was a post-doc in the MIT Computer Graphics Group with Julie Dorsey.

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