
UniAR: A Unified model for predicting human Attention and Responses on visual content

Peizhao Li* ^{†1,2}, Junfeng He* ^{‡1}, Gang Li* ^{‡1}, Rachit Bhargava¹, Shaolei Shen¹, Nachiappan Valliappan¹, Youwei Liang^{†1,3}, Hongxiang Gu¹, Venky Ramachandran¹, Golnaz Farhadi¹, Yang Li¹, Kai J Kohlhoff¹, and Vidhya Navalpakkam¹

¹Google Research

²Brandeis University

³University of California San Diego

Abstract

Progress in human behavior modeling involves understanding both implicit, early-stage perceptual behavior, such as human attention, and explicit, later-stage behavior, such as subjective preferences or likes. Yet most prior research has focused on modeling implicit and explicit human behavior in isolation; and often limited to a specific type of visual content. We propose UniAR – a unified model of human attention and preference behavior across diverse visual content. UniAR leverages a multimodal transformer to predict subjective feedback, such as satisfaction or aesthetic quality, along with the underlying human attention or interaction heatmaps and viewing order. We train UniAR on diverse public datasets spanning natural images, webpages, and graphic designs, and achieve SOTA performance on multiple benchmarks across various image domains and behavior modeling tasks. Potential applications include providing instant feedback on the effectiveness of UIs/visual content, and enabling designers and content-creation models to optimize their creation for human-centric improvements.

1 Introduction

Implicit, early-stage perceptual behavior such as human attention is intricately linked with explicit, later-stage behavior such as subjective ratings/preferences. Yet prior research has often studied these in isolation. For example, there is a large body of work on predictive models of human attention that are known to be useful for various applications, ranging from basic attention/eye-movement research [32, 35], to optimizing interaction designs [4, 65, 9], enhancing webpage layouts [67, 85, 11], improving user experience in immersive environments [5] and improving natural image and photo quality by reducing visual distraction [1]. Prior research has also explored predicting other kinds of implicit human behavior such as the sequence/order in which items are viewed (attention scanpath) in natural images or webpages [22, 19], assessing visual importance in graphic designs [47, 63, 27], and understanding visual clutter [52, 79, 71].

Separately from implicit, early-perceptual behavior, there has also been research in modeling explicit, later-stage decision-making behavior such as subjective preferences [20] and aesthetic quality [37, 21, 55, 30]. Prior research has been further fragmented due to dedicated models focusing on specific combinations of behavior tasks, input domain (e.g., natural images, designs, and webpages), and task scenarios (e.g., free viewing, object searching, and question answering).

*Co-first authors, equal technical contribution

[†]Work done during an internship at Google Research

[‡]Corresponding authors, equal leading contribution: {junfenghe, leebird}@google.com

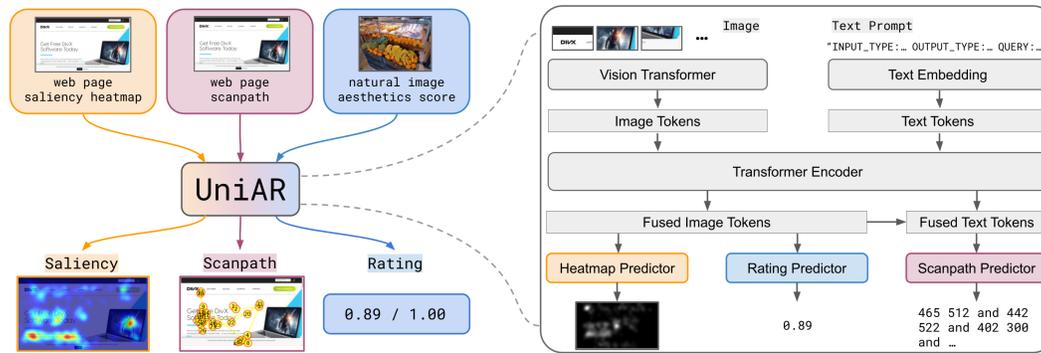


Figure 1: Overview of our UniAR model. UniAR is a multimodal model that takes an image (could be a natural image, screenshot of a webpage, graphic design, or UI) along with a text prompt as input, and outputs heatmaps of human attention/interaction, scanpath or sequence of viewing/interaction, and subjective preference/likes. Example inputs and corresponding outputs for saliency, scanpath, and rating are shown on the left side, and the detailed model architecture is shown on the right side.

To the best of our knowledge, a unified approach is still missing to modeling human visual behavior, ranging from implicit, early-perceptual behavior of what draws human attention, to explicit, later-stage decision-making on subjective preferences or likes.

In this paper, we ask the following question: *Can we build a unified model of human attention and preference behavior that reliably works across diverse types of visual content? If so, how does it compare with state-of-the-art (SOTA) models dedicated to specific domains and tasks?* Such a unified model could enable a wide variety of applications. For instance, it could augment human decision making and accelerate evaluation of effective UIs by not only predicting preferences as rewards, but also providing additional insights in the form of predicted human attention behavior.

UniAR. In this paper, we consider 11 public datasets consisting of different input domains/visual content (e.g., natural images, cartoons, art, graphic designs and webpages), behavior tasks (e.g., attention heatmaps, scanpath, likes/preference), and task-scenarios (e.g., free-viewing, object search, question answering). We introduce a new model, **UniAR** – A Unified model for predicting human Attention and Responses on visual content. UniAR is a multimodal transformer model that takes images and text prompts as input. The text prompt combines information about the input domain (e.g., natural image, graphics design or webpage), the desired behavior prediction task (e.g., attention heatmap or aesthetic score), and specifics of the task scenario when relevant (e.g., object name in an object search task). Our model generates predictions conditionally on these inputs. Experiments show that UniAR achieves SOTA performance across diverse datasets, spanning different input domains, behavior prediction tasks, and task scenarios.

Main contributions of this work are summarized below:

1. We proposed UniAR, a multimodal transformer model to predict different types of human behavior from attention to likes, across diverse types of visual content.
2. We trained UniAR on 11 benchmark datasets with different input domains (natural images, webpages, and graphic designs) and output behavior types (attention/importance heatmaps, viewing sequence or scanpath, and aesthetics/quality scores), and showed that UniAR, which is a single unified model, can outperform or perform comparably to SOTA models trained on specific tasks and datasets. We further showed that UniAR generalizes well to tasks with unseen input and output combinations, under a zero-shot setting.

We present various visualization results from UniAR on saliency/importance heatmap, scanpath, and ratings in Figure 2, compared to ground truth.

2 Related Work

Saliency prediction. Saliency or attention heatmap prediction is a common implicit behavioral task aimed at predicting which areas within an image are more likely to draw human attention. Saliency models can be helpful for understanding human visual attention, and have been used for applications such as evaluating the quality of UIs, optimizing content placement in graphic designs, and improving

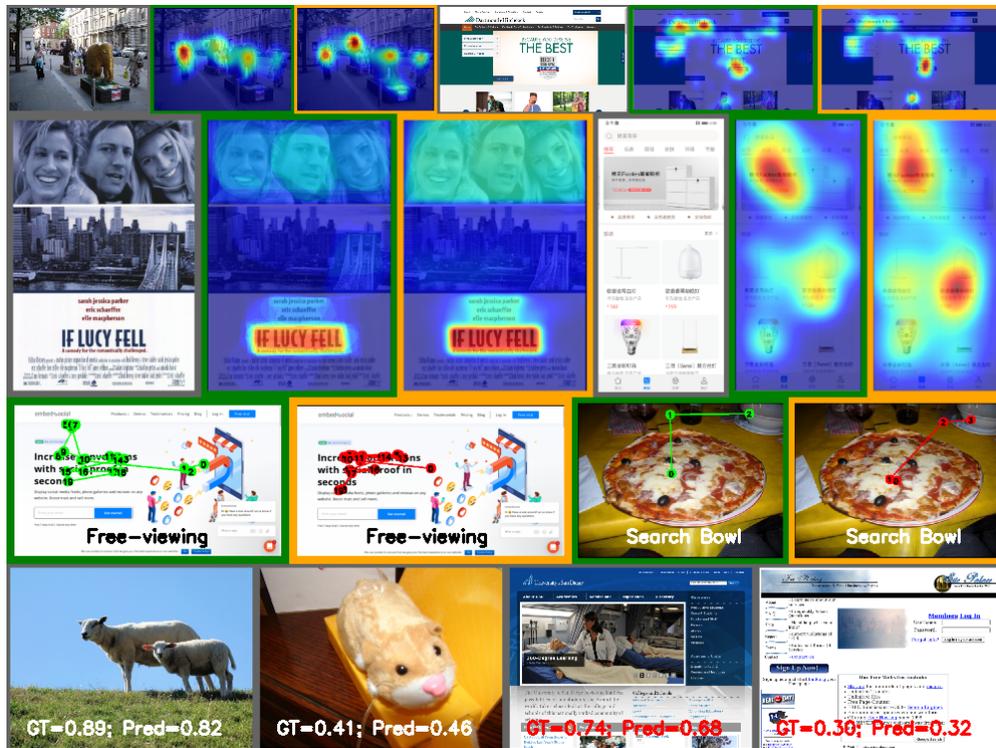


Figure 2: Examples of UniAR’s predictions across different tasks/domains. Images in green border are ground-truth, while images in orange border are UniAR’s predictions. **First row:** attention/saliency heatmap prediction on natural images (*Salicon*) and webpages (*WS-Saliency*). **Second row:** importance heatmap on graphic designs (*Imp1k*), and saliency heatmap on *Mobile UI*. **Third row:** scanpath-sequence during free-viewing of webpages (*WS-Scanpath*) and object-searching within images (*COCO-Search18*). **Fourth row:** preference/rating prediction for natural images (*Koniq-10k*) and webpages (*Web Aesthetics*).

perceptual quality of compressed images/videos (i.e., allocating more resources to visually important regions, and compressing the rest can preserve information while reducing bandwidth).

Early work focused on the importance of low-level image features in saliency prediction [32, 35, 36, 28, 42]. Recent approaches for saliency modeling use Convolutional Neural Networks (CNNs), Transformers, or a mixture of models [40, 41, 33, 51, 23, 25] as the backbone architecture to extract deep representations and predict the probability distribution over human gaze or fixations (computed as a Gaussian-blurred 2D heatmap, which aggregates all fixations from multiple human observers) [33, 61, 12, 42, 18]. Customized modules, such as 1×1 read-out convolutions [33] and Attentive ConvLSTM [18], have been introduced atop these CNNs to boost performance. Instead of the heatmap, regressing the Gaussian probability distribution has also been demonstrated as an alternative way for fixation predictions [61]. Chen et al. [12] propose to incorporate user profile information to personalize saliency predictions for each individual user.

Scanpath prediction. Unlike saliency, which predicts a heatmap/probability distribution over attention/importance, the goal here is to predict the sequence of eye movements as humans engage with visual content, offering insights into how individuals observe and comprehend visual information. With graphic designs as an example, predicting scanpaths can help optimize content placement, ensuring priority content captures attention first.

Prior work on human scanpath prediction has explored different task-scenarios, such as free-viewing, object searching, and visual question answering. Yang et al. [77] introduce a method utilizing inverse reinforcement learning for scanpath prediction during visual object searching. Continuing with this framework, Yang et al. [78] propose the concept of Foveated Feature Maps, enabling scanpath prediction when users search for an object that is not present in the initial image. To facilitate instruction following when performing a visual search over the image, Chen et al. [15] propose the use of a visual-question-answering model combined with a ConvLSTM module to predict the

distribution of fixation positions and duration in response to a question regarding the associated natural image. In recent work, Mondal et al. [54] propose employing the Transformer model to regress coordinates for each fixation within a scanpath dedicated to object searching. This regression is conditioned on the embedding of the object’s name from a pretrained language model.

Subjective rating prediction. Predicting explicit human responses such as subjective preference/likes can help to better assess image quality and improve graphic designs. These responses can be continuous or discrete ratings, and may reflect both technical quality and aesthetic quality of an image. Explicit human feedback has been used in many applications. Statistic-based methods [53, 83] and Convolutional Neural Networks-based methods [84, 68] are proposed, and recently Vision Transformer has also been adopted for this task [37].

Limitations of prior work. While there has been significant progress in modeling behavioral tasks such as saliency, scanpaths and subjective preferences, a key limitation is that prior approaches often focused on a dedicated model for each specific task x input domain. As a result, there are saliency models for natural images, scanpath prediction models for graphic designs, or subjective ratings/likes on webpages, but there isn’t a single unified model that generalizes across different tasks and domains. Instead of several dedicated per-task or per-domain models, our work seeks to build a single, unified model for these human-centered prediction tasks across diverse visual content.

Multi-tasking unified model for language & vision. There have been significant recent advances in large language models for natural language processing and vision-language learning [14, 57, 17, 3, 69, 62]. The underlying modeling recipe involves fine-tuning large transformer models on datasets containing a variety of recognition and reasoning tasks such as text summarization, sentiment analysis, machine translation for language models, and image captioning, question-answering, detection, and segmentation for vision-language models. These fine-tuned large models show strong generalization capacity across various tasks and data domains. Inspired by these generalizable models for language and vision, we propose UniAR – a unified model for predicting different types of human visual behavior (from attention to likes) on a variety of visual content.

3 Unifying Human Attention and Responses

Our model architecture along with example inputs and corresponding outputs is shown in Figure 1.

3.1 Model Architecture

Inspired by the recent progress in large vision-language models [14, 50, 45], We adopt a multimodal encoder-decoder transformer model to unify the various human behavior modeling tasks. The model takes two types of inputs: an image and a text prompt. Its architecture comprises of the following components: a Vision Transformer model [24] for image encoding, a word embedding layer to embed text tokens, and a T5 [60] Transformer encoder to fuse image and text representations. Additionally, it has three separate predictors: a heatmap predictor for attention/saliency heatmaps or visual importance heatmaps, a scanpath predictor for the sequence/order of viewing, and a rating predictor for quality/aesthetic scores of images or webpages. These predictors are described in Sections 3.2 to 3.4. Besides the architecture, the text prompt is designed to encode relevant information about the input domain (e.g., natural image, graphic design, webpages), the expected prediction type of the model (e.g., interaction heatmaps, sequence-of-viewing, or aesthetic score), and other task-related information such as viewing scenarios (e.g., free-viewing or object-searching), target object names, or questions to be answered, as described in Section 3.5. More details about model architecture including # of layers and layer size can be found in Appendix A.

To pretrain the model, we use both natural images from the WebLI dataset [14] and Web/mobile UI images [50], to ensure that the model can generalize to multiple domains. Image captioning and captioning for a screen region are used as the pretraining tasks, as in the original papers. To support sequence tasks involving prediction of gaze/interaction coordinates, such as scanpath prediction, we also add a pretraining task to predict the coordinates of the bounding box of relevant items given a text snippet and the screenshot (for webpage and mobile interface data).

3.2 Heatmap Predictor

Our model incorporates a heatmap head which is commonly used in attention/saliency research (i.e., predicting probability distribution of gaze over the input image). The heatmap prediction head takes

Table 1: List of all public datasets used to train our model. ‘# Image’ denotes the number of unique images in the entire dataset. Note that for annotation ‘scanpath,’ there are multiple scanpaths recorded from a group of users associated with one image, so ‘# Training Sample’ is much larger than ‘# Image.’ During training, we randomly sample from all training datasets with an equal sampling rate.

Dataset	Domain	Annotation	Viewing style	# Image	Resolution	# Training
<i>Salicon</i> [34]	Natural images	Attention heatmap	Free-viewing	15,000	640 × 480	10,000
<i>OSIE</i> [75]	Natural images	Attention heatmap	Free-viewing	700	800 × 600	500
<i>CAT2000</i> [6]	Natural images/cartoons/art...	Attention heatmap	Free-viewing	4000	1920 × 1080	2000
<i>WS-Saliency</i> [111]	webpage	Attention heatmap	Free-viewing	450	1,280 × 720	392
<i>Mobile UI</i> [47]	Mobile user interface	Attention heatmap	Free-viewing	193	Varied	154
<i>Imp1k</i> [27]	Graphic design	Importance heatmap	N/A	998	Varied	798
<i>WS-Scanpath</i> [111]	webpage	Scanpath (sequence)	Free-viewing	450	1,280 × 720	5,448
<i>FiWI</i> [67]	webpage	Attention heatmap	Free-viewing	159	1,360 × 768	121
<i>COCO-Search18</i> [16]	Natural images	Scanpath (sequence)	Object-searching	3,101	1,680 × 1,050	21,622
<i>KonIQ-10k</i> [31]	Natural images	Subjective rating	N/A	10,073	1,280 × 720	7,000
<i>Web Aesthetics</i> [21]	webpage	Subjective rating	N/A	398	1,280 × 768	398

the fused image tokens after the Transformer encoder, and processes the features via several read-out convolution layers, together with up-sampling so that the output will match the resolution of the input image. A sigmoid function is used at the end to ensure the generated values fall within the range $[0, 1]$ for each pixel.

In the experiments, we consider two different types of heatmaps, namely saliency and importance heatmap. A saliency heatmap is generated by aggregating human eye fixations from multiple participants viewing an image. On the other hand, importance heatmaps are obtained by participants highlighting or drawing bounding boxes to indicate the most critical design elements in a graphic design [27]. Each of these heatmaps reflects distinct aspects of human attention and preference.

The text prompt specifies which heatmap-type to generate for a given input sample, thereby allowing our model to predict a variety of heatmap prediction tasks (e.g., attention, interaction, importance etc.) using a single heatmap prediction head. We adopt a pixel-wise ℓ_2 loss function for the heatmap predictor during training.

3.3 Scanpath (Sequence) Predictor

The scanpath predictor takes both the fused image and text tokens after the Transformer encoder as input, and applies a Transformer decoder to generate the predicted scanpath.

A scanpath is defined as a sequence of 2D locations $(x_1, y_1), (x_2, y_2), \dots, (x_N, y_N)$ with a total of N fixations, capturing the temporal aspect of attention and visual exploration. The subsequent fixations are conditional on all the previous fixations, thus fitting an autoregressive model with conditional generation. Inspired by previous literature, we use Transformer decoder for object detection and other localization tasks [13, 50], and therefore generate the location end-to-end with Transformer and text generation. In general, we let the Transformer predict the sequence of coordinates as characters in a string one after one and readout the locations from the generated text subsequently.

We spatially decompose the entire image into $1,000 \times 1,000$ bins with equal interval, and map each coordinate x_n or y_n to its nearest bin $\tilde{x}_n, \tilde{y}_n \in \mathbb{Z}$ in the range $[0, 999]$.

To formulate the target sequence for teacher-forcing training, we put a special token ‘<extra_id_01>’ at the start of each target sequence, and attach another special token ‘<extra_id_02>’ at the end, to indicate the entire scanpath sequence. We concatenate location coordinates with a separation word ‘and’. Let y indicate the target sequence with length $3N + 1$ (corresponding to N fixations), we have the target sequence for teacher-forcing training as follows (\leftrightarrow indicates the line changing due to the paper format):

$$y = \langle \text{extra_id_01} \rangle \tilde{x}_1 \tilde{y}_1 \text{ and } \tilde{x}_2 \tilde{y}_2 \text{ and } \dots \text{ and } \tilde{x}_N \tilde{y}_N \langle \text{extra_id_02} \rangle.$$

The training objective is to maximize the log-likelihood of tokens conditioned on the input image and all preceding tokens in ground-truth scanpath string, i.e.,

$$\max \sum_{j=1}^{3N+1} w_j \log P(\tilde{y}_j | x, y_{1:j-1}), \quad (1)$$

where x is the input image and text prompt, and y is the target sequence associated with x . w_j is the weight for the j -th token. We use a unified weight for each token in experiments.

Decoding during inference. During inference, it is not strictly guaranteed that the generated string will exactly follow the format of the target sequence, especially when the number of fixations is relatively large (e.g., $N \geq 30$), resulting in an invalid sequence to readout. To deal with the possible invalid cases, in practice, we identify two special tokens '<extra_id_01>' and '<extra_id_02>' in the predicted string if available, and extract the context between these two special tokens. Then we split the extracted string with the separator word 'and'. For a pair of two tokens from the beginning, we check if they are both numerical. If so, we add one fixation with this coordinate after mapping them to the original resolution, then iteratively move to the following, and if not, we terminate the decoding process and keep the existing sequence. If there is no fixation available in the predicted string, we mark the scanpath as invalid. During training, we observe that the valid rate ($\# \text{valid scanpaths} / \# \text{scanpaths generated}$) of scanpath decoding quickly converges to 1, meaning every predicted scanpath will contain valid fixation(s).

Compared to the scanpath predictors in GazeFormer [54] which predicts each point as 2D continuous numbers instead of two text tokens, our model seems less intuitive. However, as shown in Section 4.3, the proposed scanpath predictor works quite well. Moreover, one advantage of the current design of the scanpath predictor is that it can be easily extended to predict other types of human behavior sequences, e.g., text sequence, or 1-D number sequence, with minor modifications on the sequence output format. This flexibility is important for a unified model like ours.

3.4 Rating Predictor

This prediction head takes image tokens after the Transformer encoder module, and processes the features via a few convolution and connected layers. An ℓ_2 loss is used for training the rating predictor with rating data.

3.5 Text Prompt

To enhance the model's ability to generalize across a variety of visual content and task scenarios, we integrate specific task instructions into the model via text prompts. The prompts used in UniAR are structured as follows:

INPUT_TYPE:<input_type> OUTPUT_TYPE:<output_type> QUERY:<query>.

We fill <input_type> with string taken from {natural image | webpage | graphic design | mobile user interface} and <output_type> taken from {saliency heatmap | importance heatmap | aesthetics score | scanpath}. We append a query in string Query:<query> to the prompt if a task-specific query is available, for example, the object name to search, or the question to answer, depending on the use case. For example, an example full prompt is "INPUT_TYPE: natural image OUTPUT_TYPE: scanpath QUERY:searching a bowl", which guides the model to predict scanpath output on a natural image under the task of "searching a bowl". The prompt we use is modularized and can easily adapt to different types of datasets and scenarios.

4 Experiment

4.1 Protocol

Datasets. Please refer to Table 1 for all public datasets we consider in training and benchmarking. For more dataset processing details, please refer to Appendix B.

Benchmarks. We reuse benchmarks from recent literature for model comparison purposes. We adopt the benchmarks for *WS-Saliency* and *WS-Scanpath* from Tables 3 and 7 in Chakraborty et al. [11] respectively, *Mobile UI* from Table 2 in Leiva et al. [47], *Imp1k* from Table 2 in Fosco et al. [27], *OSIE* from Table 4 in Chen et al. [12], *Salicon* from Table 1⁴ in Reddy et al. [61], *COCO-Search18* from Table 1 in Mondal et al. [54], *KonIQ-10k* from Table 2 in Ke et al. [37], and *Web Aesthetics* from Table 4 in Delitzas et al. [21]. We also provide model results from some other papers for comparison [25, 33]. The baseline results for *CAT2000* are from *CAT2000* Leaderboard⁵.

⁴There are two versions of Salicon data: Salicon 2015 and Salicon 2017. The results are on Salicon 2017.

⁵CAT2000 Leaderboard: https://saliency.tuebingen.ai/results_CAT2000.html

Evaluation metrics. Inheriting from the above benchmarks, we consider the following evaluation metrics. **CC** [44]: Pearson’s Correlation Coefficient is used here to measure the linear relationship in all pixel values between the predicted and ground-truth saliency heatmaps; **KLD** [39]: the metric to use KL-Divergence between the predicted heatmap and ground-truth heatmap to measure the distribution discrepancy, with the prediction used as the target distribution. **AUC-Judd** [35]: Area under ROC curve (AUC) in the variant from Judd et al. [35] treating the heatmap prediction as binary classification with various thresholds. The specific calculations of true positive and false positive rates can be referred to [10]. **sAUC** [7]: the shuffled AUC metric samples negatives from other images for AUC calculation. **NSS** [59]: Normalized Scanpath Saliency is the average saliency strength (pixel values in the predicted heatmap) at all ground-truth fixation locations. **SIM** [64]: Similarity is computed as the sum of the minimum values among the normalized prediction and ground-truth heatmaps. **RMSE**: the root mean square error between the predicted and ground-truth heatmaps. **R-Squared** (R^2): the coefficient of determination applied to all values in the heatmap. **SemSS** [78]: Semantic Sequence Score converts each fixation to an ID decided by a semantic segmentation over the image, and compares two strings with a string-matching algorithm [56]. **SemFED** [78]: similar to SemSS, Semantic Fixation Edit Distance uses Levenshtein distance for string matching [48]. **SequenceScore**: similar to SemSS, but instead of using a semantic segmentation map, the Sequence Score uses the clustering results from a MeanShift clustering to segment the image and map the fixation to their ID. **MultiMatch** [22]: MultiMatch is the average of four metrics of scanpath, namely: **Shape**, **Direction**, **Length**, and **Position**, characterizing the similarity between the predicted scanpath and its ground-truth. **SRCC** and **PLCC**: refer to Spearman’s rank correlation coefficient and Pearson’s linear correlation coefficient, respectively, used to quantify the quality of predicted ratings.

Experimental benchmarks for one task among different datasets may not have uniform evaluation metrics but most of their metrics are shared. For saliency and importance heatmap predictions, we resize the predicted heatmap back to its original image resolution for evaluation.

4.2 Model Training

We pretrain the model on a series of pre-training tasks, including Web/Mobile UI understanding [50] and natural image captioning [14]. Subsequently, we fine-tune the entire model using the Adafactor optimizer with a learning rate of 0.1, batch size of 128, and image resolution of 512×512 . All images maintain their aspect ratio and are padded to fit the training resolution. The model uses ViT B16 as the vision encoder and T5 base as the Transformer encoder of the image and text tokens, resulting in a total of 848 million parameters. The model is implemented in JAX. We use 64 Google Cloud TPU v3 to train UniAR for 20k steps in 12 hours.

Datasets mixture. As we are combining a series of public datasets, in every iteration, for all training datasets in Table 1, we employ a random sampling strategy that ensures an equal sampling rate across these datasets. This approach guarantees that each dataset has an equal probability of contributing a sample, irrespective of its sample volume.

Table 2: Subjective rating prediction results on Natural image dataset *KonIQ-10k* and webpage dataset *Web Aesthetics*.

Dataset	Method	SRCC \uparrow	PLCC \uparrow
<i>KonIQ-10k</i> [31] (Natural image)	BRISQUE [53]’12	0.665	0.681
	ILNIQE [83]’15	0.507	0.523
	HOSA [74]’16	0.671	0.694
	BIECON [38]’16	0.618	0.651
	WaDIQaM [8]’17	0.797	0.805
	PQR [81]’17	0.880	0.884
	SFA [49]’18	0.856	0.872
	DBCNN [84]’18	0.875	0.884
	MetalQA [86]’20	0.850	0.887
	BIQA (25 crops) [68]’20	0.906	0.917
	MUSIQ-single [37]’21	0.905	0.919
	UniAR	0.905 -0.11%	0.925 +0.65%
<i>Web Aesthetics</i> [21] (webpage)	Rating-based Calista [21]’23	-	0.770
	Comparison-based Calista [21]’23	-	0.820
	UniAR	0.811	0.839 +2.31%

4.3 Experiment Results

We present the results of UniAR for predicting heatmaps, scanpath-sequences as well as ratings across domains and datasets in Tables 2 to 4, in comparison with the baselines that are trained on a specific domain, task, or dataset.

Heatmap prediction. Table 3 shows the performance of UniAR across 7 public benchmarks. A complete version of results including all baselines and metrics is presented in Table 6 in Appendix C. Among these datasets, which vary in domains (natural images, webpages, and graphic designs) and tasks (heatmaps of attention, and importance), UniAR achieves SOTA performance compared to strong baselines, and outperforms previous SOTAs in many cases. On *Mobile UI* and *Imp1k* datasets, UniAR outperforms previous SOTA across every metric. Out of the 27 metrics listed in Table 3,

Table 3: Heatmap prediction results on 7 public datasets across natural images, art, cartoons, mobile UIs, and webpages (Please refer to Table 6 in Appendix C for complete baselines & metrics). For *Imp1k* we predict the importance heatmap, while for the remaining datasets, we predict the attention/saliency heatmap. For each dataset and metric, the best result is in **bold**, second best is in **blue**, and our method is highlighted in green. For our model, the relative performance change compared to the best result is noted. Note that the metric values for baseline models are obtained from existing references as described in the "Benchmarks" paragraph. "-" means the metrics are not reported in references. Also note that there are two versions of Salicon data, Salicon 2015 and Salicon 2017. The results in this table are on Salicon 2017.

Dataset	Method	CC \uparrow	KLD \downarrow	AUC-Judd \uparrow	NSS \uparrow
<i>WS-Saliency</i> [11] (Webpage)	SAM-ResNet [18]'18	0.596	1.506	0.795	1.284
	EML-NET [33]'20	0.565	2.110	0.790	1.277
	UMSI [27]'20	0.444	1.335	0.757	1.042
	DI Net + WS [76]'19	0.798	0.690	0.852	1.777
	AGD-F (W/o-L) [11]'22	0.815	0.637	0.858	1.802
	UniAR	0.827 +1.47%	0.299 -53.06%	0.860 +0.23%	1.783 -1.05%
<i>FiWI</i> [67] (Webpage)	AGD-F [11]'22	0.735	-	0.767	1.606
	EML-NET [33]'20	0.661	0.603	0.847	1.653
	EML-NET + <i>Salicon</i> [33]'20	0.689	0.567	0.848	1.722
	Chen et al. [12]'23	0.699	0.564	0.851	1.752
	UniAR	0.734 -0.14%	0.571 +0.29%	0.859 +0.94%	1.838 +4.91%
<i>Mobile UI</i> [47] (Mobile interface)	ResNet-Sal [47]'20	0.657	-	0.692	0.704
	SAM-S2015 [18]'18	0.477	-	0.650	0.537
	SAM-S2017 [18]'18	0.834	-	0.723	0.839
	SAM-mobile [47]'20	0.621	-	0.666	0.655
	UniAR	0.879 +5.40%	0.115	0.756 +4.56%	1.008 +20.14%
<i>CAT2000</i> [43] (Natural images, cartoons, art...)	DeepGaze II [42]'17	0.795	0.382	0.864	1.962
	UNISAL [25]'20	0.740	0.470	0.860	1.936
	DeepGaze III [51]'21	0.819	0.345	0.869	2.112
	SalFBNet [23]'22	0.703	1.198	0.855	1.879
	UniAR	0.870 +6.23%	0.613 +77.68%	0.877 +0.92%	2.338 +10.70%
<i>Salicon</i> [34] (Natural images)	SimpleNet w. ResNet-50 [61]'20	0.895	0.211	0.868	1.881
	SimpleNet w. PNASNet-5 [61]'20	0.907	0.193	0.871	1.926
	MDNSal [61]'20	0.899	0.217	0.868	1.893
	UNISAL [25]'20	0.880	0.226	0.867	1.923
	EML-NET [33]'20	0.890	0.204	0.802	2.024
	UniAR	0.900 -0.77%	0.214 +10.88%	0.870 -0.11%	1.946 -3.85%
<i>OSIE</i> [70] (Natural images)	SAM-ResNet [18]'18	0.758	0.480	0.860	1.811
	UMSI [27]'20	0.746	0.513	0.856	1.788
	EML-NET [33]'20	0.717	0.537	0.854	1.737
	Chen et al. [12]'23	0.761	0.506	0.860	1.840
	UniAR	0.742 -2.50%	0.583 +21.46%	0.862 +0.23%	1.789 -2.77%
Dataset	Method	CC \uparrow	KLD \downarrow	RMSE \downarrow	R^2 \uparrow
<i>Imp1k</i> [27] (Graphic design)	SAM [18]'18	0.866	0.166	0.168	0.108
	UMSI-nc [27]'20	0.802	0.177	0.152	0.095
	UMSI-2stream [27]'20	0.852	0.168	0.141	0.105
	UMSI [27]'20	0.875	0.164	0.134	0.115
	UniAR	0.904 +3.31%	0.124 -25.00%	0.079 -41.04%	0.823 +615.65%

UniAR achieves the best result in 17 of them and the second best in 6 cases. In summary, UniAR experimentally demonstrates promising performance in saliency modeling in various fields.

Scanpath prediction. In Table 4, scanpath-sequence prediction results are shown for two datasets: COCO-Search18 [16] (scanpath in natural images for object searching) and *WS-Scanpath* [11] (scanpath on webpages under free viewing). On both the datasets, UniAR performs comparably to baselines, and further outperforms the baselines on all the metrics on *WS-Scanpath*. Among the 5 reported metrics in Table 4, our model achieved the best result in 4 of them. A complete version of results including all baselines and metrics is presented in Table 7 in Appendix C.

Score prediction. In Table 2, we present rating prediction results on two datasets: *KonIQ-10k* [31] on natural images and *Web Aesthetics* [21] on webpages. UniAR achieves the best results for PLCC metrics on both datasets and the second best for SRCC on *KonIQ-10k*. Note that in Ke et al. [37], a multi-scale version of MUSIQ performs slightly better than UniAR on SRCC (0.916 vs 0.905). However, since UniAR does not use multi-scale inputs, we did not include those results.

Table 4: Scanpath (sequence) prediction results on natural image and digital design datasets. Please refer to Table 7 in Appendix C for complete baselines & metrics.

Dataset	Method	SemSS \uparrow	SemFED \downarrow	Sequence Score \uparrow	MultiMatch \uparrow
<i>COCO-Search18</i> [16] (Natural images, object searching)	IRL [77]'20	0.481	2.259	-	0.833
	Chen et al. [15]'21	0.470	1.898	-	0.820
	FFM [78]'22	0.407	2.425	-	0.808
	Gazeformer [54]'23	0.496	1.861	-	0.849
	UniAR	0.521 +5.04%	2.004 +7.68%	-	0.874 +2.94%
<i>WS Scanpath</i> [11] (webpage, free-viewing)	SceneWalker [66]'20	-	-	0.194	0.716
	AGD-F (w. layout) [11]'22	-	-	0.203	0.719
	AGD-S (w/o layout) [11]'22	-	-	0.221	0.745
	AGD-S (w. layout) [11]'22	-	-	0.224	0.755
	UniAR	-	-	0.267 +19.20%	0.887 +17.48%

Transferring knowledge between tasks. We test UniAR’s ability to generalize and transfer to unseen tasks/domain combinations. Our model is trained on certain combinations of task and image domains, and tested on new, unseen combinations of behavior tasks and image domains.

Our experiment uses *WS* (webpage) and *COCO-Freeview* (free-viewing on natural image) datasets, where *WS-saliency*, *WS-scanpath*, *CC saliency* and *CC scanpath*, are saliency and scanpath data for *WS* and *COCO-Freeview* data respectively. We test the model performance on scanpath prediction on the webpage scanpath data (*WS-Scanpath* dataset). We consider three different training scenarios: (1) Using scanpath data from natural image (*COCO-Freeview*); (2) Combining scanpath from natural image (*COCO-Freeview*) with saliency heatmaps from webpage (*WS*); (3) Employing both scanpath and saliency heatmap from natural image (*COCO-Freeview*), augmented with saliency heatmap from webpage (*WS*). Each scenario maintains some relevance to our test set by either sharing the same task or image domain, but never both.

In Table 5, we show experimental results for the above scenarios, and also attach the baseline results on this test set and the results from UniAR (full training data) as reference. As shown in Table 5, our third training scenario: leveraging scanpath and saliency heatmaps from *COCO-Freeview* and saliency heatmap from *WS*, shows good results against previous SoTA [11], despite the model not having seen webpage scanpath data during training. Prediction performance declines in scenarios 1 & 2, which take more limited datasets, but remain competitive.

Table 5: Experiments on transferring knowledge from other domain/task combinations to *WS-Scanpath* dataset for scanpath predictions. *CC* = *COCO-FreeView* dataset.

Training Set	Sequence Score \uparrow	MultiMatch \uparrow
<i>WS-Scanpath</i> (previous SoTA [11])	0.224	0.755
<i>WS-Scanpath</i> (ours)	0.261	0.894
UniAR full model	0.267	0.887
<i>CC scanpath</i> (ours)	0.196	0.836
<i>CC scanpath</i> + <i>WS-Saliency</i> (ours)	0.190	0.858
<i>CC saliency/scanpath</i> + <i>WS-Saliency</i> (ours)	0.231	0.857

5 Limitations and Future Work

When modeling human preferences and behavior, it is important to carefully consider ethics and AI principles, and acknowledge dataset limitations.

Ethics and AI principles. Modeling any aspect of human behavior should adhere to ethical guidelines on data collection and applications, and be conducted in a transparent way, including clarifying the limitations of the model when replicating human preferences. It should keep humans in the loop, as the model prediction is intended as a reference, not as a means to replace evolving human preferences with a synthetic guide. We take these ethical considerations and AI Principles into account, ensuring that the model usage remains socially beneficial and responsible.

Aligning with human preference. While UniAR is trained to predict human preferences and behavior, we recognize that using it as a reward model may lead to reward hacking, which would make it less representative of genuine human preference. We suggest considering techniques [46, 26] to mitigate this.

Representing diverse human preferences. Humans have diverse preferences for subjective notions like image attractiveness. Without personalization, the model converges towards a more uniform notion of preference - a common concern for ML models. To promote visual diversity when using UniAR, we propose two strategies: (1) using the model in a hybrid manner, providing insights to

help humans make decisions in applications like web or visual content optimization, and (2) develop personalized models based on our initial unified model, which will help generate more diverse predictions based on user attributes.

Adjusting to evolving human preferences. Human preferences evolve over time, and to remain accurate, the model has to adjust accordingly. Updates can be achieved by fine-tuning with more recent data. Future work will include updating the training data, and exploring concepts like continual learning techniques [72], to keep the model up to date.

Dataset limitations. Using diverse, representative datasets is important to minimize potential biases in the model. In this paper, we focus on a proof-of-concept for unified modeling of human attention/preference behavior, based on existing, publicly available datasets. Below is a listing of annotator demographics, as described in the original papers.

1. *WS-saliency* [11]: "A total of 41 participants (19 females, 22 males; age range 17-23; with normal or corrected-to-normal vision) participated in our data collection."
2. *Mobile UI* [47]: "Thirty participants (12 male, 18 female). [...] The average age was 25.9 (SD=3.95). The participants had normal vision (8) or corrected-to-normal-vision (22). Twenty of the 22 wore glasses and the remaining two wore contact lenses."
3. *Implk* [27]: "The data of 43 participants (29 male, most in their 20s and 30s) were used in the resulting analyses."
4. *FiWI* [67]: "11 students (4 males and 7 females) in the age range of 21 to 25 participated in data collection. All participants had normal vision or corrective visual apparatus."

Despite some balance in male vs. female participants, the age distribution is skewed towards participants in their 20s. This is likely because most data were collected at universities. Crowdsourced datasets like Salicon and Koniq-10K are expected to cover a wider range of age and other attributes. Future work will focus on gathering a more diverse and representative dataset.

Improve accessibility. UniAR predicts human attention, trained on datasets collected from humans not experiencing visual impairments beyond corrective lenses. UniAR cannot model behavior of, for example, blind and low-vision users directly, but it can still benefit them by acting as an accessibility tool for highlighting important areas of a webpage for a screen reader. One way to enhance the accessibility of the model is using multi-modal preference modeling, which incorporates not just visual cues, but also how users interact with content through screen readers, voice commands, and other assistive technologies. Collaborating with accessibility experts and organizations could help improve future iterations of our work.

6 Conclusion

We developed a multimodal, unified model UniAR to predict different types of implicit and explicit human behavior on visual content, from attention to subjective preferences/likes, using image and text prompts. This model, trained on diverse public datasets across natural images, graphic designs, webpages and UIs, effectively predicted human attention heatmaps, scanpath sequences, and aesthetic or quality scores. Our model achieved SOTA performance across multiple benchmarks and tasks. We plan to explore more behavior tasks and domains in future work.

References

- [1] Kfir Aberman, Junfeng He, Yossi Gandelsman, Inbar Mosseri, David E Jacobs, Kai Kohlhoff, Yael Pritch, and Michael Rubinstein. Deep saliency prior for reducing visual distraction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [2] Hossein Adeli, Françoise Vitu, and Gregory J Zelinsky. A model of the superior colliculus predicts fixation locations during scene viewing and visual search. *Journal of Neuroscience*, 2017.
- [3] Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*, 2023.
- [4] Saskia Bakker and Karin Niemantsverdriet. The interaction-attention continuum: Considering various levels of human attention in interaction design. *International Journal of Design*, 2016.
- [5] Leonardo Bonanni, Chia-Hsun Lee, and Ted Selker. Attention-based design of augmented reality interfaces. In *CHI'05 extended abstracts on Human factors in computing systems*, 2005.
- [6] Ali Borji and Laurent Itti. Cat2000: A large scale fixation dataset for boosting saliency research. *arXiv preprint arXiv:1505.03581*, 2015.
- [7] Ali Borji, Hamed R Tavakoli, Dicky N Sihite, and Laurent Itti. Analysis of scores, datasets, and models in visual saliency prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, 2013.
- [8] Sebastian Bosse, Dominique Maniry, Klaus-Robert Müller, Thomas Wiegand, and Wojciech Samek. Deep neural networks for no-reference and full-reference image quality assessment. *IEEE Transactions on image processing*, 2017.
- [9] Zoya Bylinskii, Nam Wook Kim, Peter O'Donovan, Sami Alsheikh, Spandan Madan, Hanspeter Pfister, Fredo Durand, Bryan Russell, and Aaron Hertzmann. Learning visual importance for graphic designs and data visualizations. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*, 2017.
- [10] Zoya Bylinskii, Tilke Judd, Aude Oliva, Antonio Torralba, and Frédo Durand. What do different evaluation metrics tell us about saliency models? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- [11] Souradeep Chakraborty, Zijun Wei, Conor Kelton, Seoyoung Ahn, Aruna Balasubramanian, Gregory J Zelinsky, and Dimitris Samaras. Predicting visual attention in graphic design documents. *IEEE Transactions on Multimedia*, 2022.
- [12] Shi Chen, Nachiappan Valliappan, Shaolei Shen, Xinyu Ye, Kai Kohlhoff, and Junfeng He. Learning from unique perspectives: User-aware saliency modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [13] Ting Chen, Saurabh Saxena, Lala Li, Tsung-Yi Lin, David J Fleet, and Geoffrey E Hinton. A unified sequence interface for vision tasks. *Advances in Neural Information Processing Systems*, 2022.
- [14] Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. Pali: A jointly-scaled multilingual language-image model. *arXiv preprint arXiv:2209.06794*, 2022.
- [15] Xianyu Chen, Ming Jiang, and Qi Zhao. Predicting human scanpaths in visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [16] Yupei Chen, Zhibo Yang, Seoyoung Ahn, Dimitris Samaras, Minh Hoai, and Gregory Zelinsky. Cocossearch18 fixation dataset for predicting goal-directed attention control. *Scientific Reports*, 2021.
- [17] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*, 2022.
- [18] Marcella Cornia, Lorenzo Baraldi, Giuseppe Serra, and Rita Cucchiara. Predicting Human Eye Fixations via an LSTM-based Saliency Attentive Model. *IEEE Transactions on Image Processing*, 2018.
- [19] Filipe Cristino, Sebastiaan Mathôt, Jan Theeuwes, and Iain D Gilchrist. Scanmatch: A novel method for comparing fixation sequences. *Behavior Research Methods*, 2010.

- [20] Benjamin de Haas, Alexios L. Iakovidis, D. Samuel Schwarzkopf, and Karl R. Gegenfurtner. Individual differences in visual salience vary along semantic dimensions. *Proceedings of the National Academy of Sciences*, 2019.
- [21] Alexandros Delitzas, Kyriakos C Chatzidimitriou, and Andreas L Symeonidis. Calista: A deep learning-based system for understanding and evaluating website aesthetics. *International Journal of Human-Computer Studies*, 2023.
- [22] Richard Dewhurst, Marcus Nyström, Halszka Jarodzka, Tom Foulsham, Roger Johansson, and Kenneth Holmqvist. It depends on how you look at it: Scanpath comparison in multiple dimensions with multimatch, a vector-based approach. *Behavior Research Methods*, 2012.
- [23] Guanqun Ding, Nevrez İmamoğlu, Ali Caglayan, Masahiro Murakawa, and Ryosuke Nakamura. Salfbnet: Learning pseudo-saliency distribution via feedback convolutional networks. *Image and Vision Computing*, 2022.
- [24] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.
- [25] Richard Droste, Jianbo Jiao, and J Alison Noble. Unified image and video saliency modeling. In *The European Conference on Computer Vision*, 2020.
- [26] Luca Eyring, Shyamgopal Karthik, Karsten Roth, Alexey Dosovitskiy, and Zeynep Akata. Reno: Enhancing one-step text-to-image models through reward-based noise optimization. *arXiv preprint arXiv:2406.04312*, 2024.
- [27] Camilo Fosco, Vincent Casser, Amish Kumar Bedi, Peter O'Donovan, Aaron Hertzmann, and Zoya Bylinskii. Predicting visual importance across graphic design types. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, 2020.
- [28] Jonathan Harel, Christof Koch, and Pietro Perona. Graph-based visual saliency. *Advances in Neural Information Processing Systems*, 2006.
- [29] Sen He, Hamed R Tavakoli, Ali Borji, Yang Mi, and Nicolas Pugeault. Understanding and visualizing deep visual saliency models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [30] Weng Khuan Hoh, Fang-Lue Zhang, and Neil A Dodgson. Salient-centeredness and saliency size in computational aesthetics. *ACM Transactions on Applied Perception*, 2023.
- [31] Vlad Hosu, Hanhe Lin, Tamas Sziranyi, and Dietmar Saupe. Koniq-10k: An ecologically valid database for deep learning of blind image quality assessment. *IEEE Transactions on Image Processing*, 2020.
- [32] Laurent Itti, Christof Koch, and Ernst Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1998.
- [33] Sen Jia and Neil D.B. Bruce. Eml-net: An expandable multi-layer network for saliency prediction. *Image and Vision Computing*, 2020.
- [34] Ming Jiang, Shengsheng Huang, Juanyong Duan, and Qi Zhao. Salicon: Saliency in context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [35] Tilke Judd, Krista Ehinger, Frédo Durand, and Antonio Torralba. Learning to predict where humans look. In *IEEE International Conference on Computer Vision*, 2009.
- [36] Tilke Judd, Frédo Durand, and Antonio Torralba. A benchmark of computational models of saliency to predict human fixations. *MIT Technical Report MIT-CSAIL-TR-2012-001*, 2012.
- [37] Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. Musiq: Multi-scale image quality transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021.
- [38] Jongyoo Kim and Sanghoon Lee. Fully deep blind image quality predictor. *IEEE Journal of selected topics in signal processing*, 2016.
- [39] Solomon Kullback. *Information Theory and Statistics*. 'Dover Publications, Inc.', 1959.
- [40] Matthias Kümmerer, Lucas Theis, and Matthias Bethge. Deep gaze i: Boosting saliency prediction with feature maps trained on imagenet. *arXiv preprint arXiv:1411.1045*, 2014.

- [41] Matthias Kümmerer, Thomas SA Wallis, and Matthias Bethge. Deepgaze ii: Reading fixations from deep features trained on object recognition. *arXiv preprint arXiv:1610.01563*, 2016.
- [42] Matthias Kümmerer, Thomas SA Wallis, Leon A Gatys, and Matthias Bethge. Understanding low-and high-level contributions to fixation prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, 2017.
- [43] Matthias Kümmerer, Zoya Bylinskii, Tilke Judd, Ali Borji, Laurent Itti, Frédo Durand, Aude Oliva, Antonio Torralba, and Matthias Bethge. Mit/tübingen saliency benchmark. <https://saliency.tuebingen.ai/>, 2018.
- [44] Olivier Le Meur, Patrick Le Callet, and Dominique Barba. Predicting visual fixations on video based on low-level visual features. *Vision Research*, 2007.
- [45] Kenton Lee, Mandar Joshi, Iulia Raluca Turc, Hexiang Hu, Fangyu Liu, Julian Martin Eisenschlos, Urvashi Khandelwal, Peter Shaw, Ming-Wei Chang, and Kristina Toutanova. Pix2Struct: Screenshot parsing as pretraining for visual language understanding. In *Proceedings of the 40th International Conference on Machine Learning*, 2023.
- [46] Seung Hyun Lee, Yinxiao Li, Junjie Ke, Innfarn Yoo, Han Zhang, Jiahui Yu, Qifei Wang, Fei Deng, Glenn Entis, Junfeng He, et al. Parrot: Pareto-optimal multi-reward reinforcement learning framework for text-to-image generation. In *European Conference on Computer Vision*, 2025.
- [47] Luis A Leiva, Yunfei Xue, Avya Bansal, Hamed R Tavakoli, Tuğçe Körođlu, Jingzhou Du, Niraj R Dayama, and Antti Oulasvirta. Understanding visual saliency in mobile user interfaces. In *International Conference on Human-Computer Interaction with Mobile Devices and Services*, 2020.
- [48] Vladimir I Levenshtein et al. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet Physics Doklady*. Soviet Union, 1966.
- [49] Dingquan Li, Tingting Jiang, Weisi Lin, and Ming Jiang. Which has better visual quality: The clear blue sky or a blurry animal? *IEEE Transactions on Multimedia*, 2018.
- [50] Gang Li and Yang Li. Spotlight: Mobile UI understanding using vision-language models with a focus. In *The Eleventh International Conference on Learning Representations*, 2023.
- [51] Akis Linardos, Matthias Kümmerer, Ori Press, and Matthias Bethge. Deepgaze iie: Calibrated prediction in and out-of-domain for state-of-the-art saliency modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021.
- [52] Alaster J Meehan and Joanne B Culpepper. Clutter estimation and perception. *Optical Engineering*, 2016.
- [53] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spatial domain. *IEEE Transactions on image processing*, 2012.
- [54] Sounak Mondal, Zhibo Yang, Seoyoung Ahn, Dimitris Samaras, Gregory Zelinsky, and Minh Hoai. Gazeformer: Scalable, effective and fast prediction of goal-directed human attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [55] Bence Nanay. Aesthetic attention. *Journal of Consciousness Studies*, 2015.
- [56] Saul B Needleman and Christian D Wunsch. A general method applicable to the search for similarities in the amino acid sequence of two proteins. *Journal of Molecular Biology*, 1970.
- [57] OpenAI. Gpt-4 technical report, 2023.
- [58] Junting Pan, Cristian Canton Ferrer, Kevin McGuinness, Noel E O'Connor, Jordi Torres, Elisa Sayrol, and Xavier Giro-i Nieto. Salgan: Visual saliency prediction with generative adversarial networks. *arXiv preprint arXiv:1701.01081*, 2017.
- [59] Robert J Peters, Asha Iyer, Laurent Itti, and Christof Koch. Components of bottom-up gaze allocation in natural images. *Vision Research*, 2005.
- [60] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 2020.
- [61] Navyasri Reddy, Samyak Jain, Pradeep Yarlagadda, and Vineet Gandhi. Tidying deep saliency prediction architectures. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020.

- [62] Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024.
- [63] Claudia Roda. Human attention and its implications for human–computer interaction. *Human attention in digital environments*, 2011.
- [64] Yossi Rubner, Carlo Tomasi, and Leonidas J Guibas. The earth mover’s distance as a metric for image retrieval. *International Journal of Computer Vision*, 2000.
- [65] Eldon Schoop, Xin Zhou, Gang Li, Zhouong Chen, Bjoern Hartmann, and Yang Li. Predicting and explaining mobile ui tappability with vision modeling and saliency analysis. In *Conference on Human Factors in Computing Systems*, 2022.
- [66] Lisa Schwetlick, Lars Oliver Martin Rothkegel, Hans Arne Trukenbrod, and Ralf Engbert. Modeling the effects of perisaccadic attention on gaze statistics during scene viewing. *Communications Biology*, 2020.
- [67] Chengyao Shen and Qi Zhao. Webpage saliency. In *Proceedings of the European Conference on Computer Vision*, 2014.
- [68] Shaolin Su, Qingsen Yan, Yu Zhu, Cheng Zhang, Xin Ge, Jinqiu Sun, and Yanning Zhang. Blindly assess image quality in the wild guided by a self-adaptive hyper network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [69] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- [70] Nachiappan Valliappan, Na Dai, Ethan Steinberg, Junfeng He, Kantwon Rogers, Venky Ramachandran, Pingmei Xu, Mina Shojaeizadeh, Li Guo, Kai Kohlhoff, et al. Accelerating eye movement research via accurate and affordable smartphone eye tracking. *Nature Communications*, 2020.
- [71] Ronald Van den Berg, Frans W Cornelissen, and Jos BTM Roerdink. A crowding model of visual clutter. *Journal of Vision*, 2009.
- [72] Liyuan Wang, Xingxing Zhang, Hang Su, and Jun Zhu. A comprehensive survey of continual learning: Theory, method and application. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [73] Wenguan Wang and Jianbing Shen. Deep visual attention prediction. *IEEE Transactions on Image Processing*, 2017.
- [74] Jingtao Xu, Peng Ye, Qiaohong Li, Haiqing Du, Yong Liu, and David Doermann. Blind image quality assessment based on high order statistics aggregation. *IEEE Transactions on Image Processing*, 2016.
- [75] Juan Xu, Ming Jiang, Shuo Wang, Mohan S Kankanhalli, and Qi Zhao. Predicting human gaze beyond pixels. *Journal of Vision*, 2014.
- [76] Sheng Yang, Guosheng Lin, Qiuping Jiang, and Weisi Lin. A dilated inception network for visual saliency prediction. *IEEE Transactions on Multimedia*, 2019.
- [77] Zhibo Yang, Lihan Huang, Yupei Chen, Zijun Wei, Seoyoung Ahn, Gregory Zelinsky, Dimitris Samaras, and Minh Hoai. Predicting goal-directed human attention using inverse reinforcement learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [78] Zhibo Yang, Sounak Mondal, Seoyoung Ahn, Gregory Zelinsky, Minh Hoai, and Dimitris Samaras. Target-absent human attention. In *Proceedings of the European Conference on Computer Vision*, 2022.
- [79] Chen-Ping Yu, Dimitris Samaras, and Gregory J Zelinsky. Modeling visual clutter perception using proto-object segmentation. *Journal of vision*, 2014.
- [80] Dario Zanca, Stefano Melacci, and Marco Gori. Gravitational laws of focus of attention. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [81] Hui Zeng, Lei Zhang, and Alan C Bovik. A probabilistic quality representation approach to deep blind image quality prediction. *arXiv preprint arXiv:1708.08190*, 2017.
- [82] Jianming Zhang and Stan Sclaroff. Saliency detection: A boolean map approach. In *Proceedings of the IEEE International Conference on Computer Vision*, 2013.

- [83] Lin Zhang, Lei Zhang, and Alan C Bovik. A feature-enriched completely blind image quality evaluator. *IEEE Transactions on Image Processing*, 2015.
- [84] Weixia Zhang, Kede Ma, Jia Yan, Dexiang Deng, and Zhou Wang. Blind image quality assessment using a deep bilinear convolutional neural network. *IEEE Transactions on Circuits and Systems for Video Technology*, 2018.
- [85] Quanlong Zheng, Jianbo Jiao, Ying Cao, and Rynson WH Lau. Task-driven webpage saliency. In *Proceedings of the European Conference on Computer Vision*, 2018.
- [86] Hancheng Zhu, Leida Li, Jinjian Wu, Weisheng Dong, and Guangming Shi. Metaiqa: Deep meta-learning for no-reference image quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.

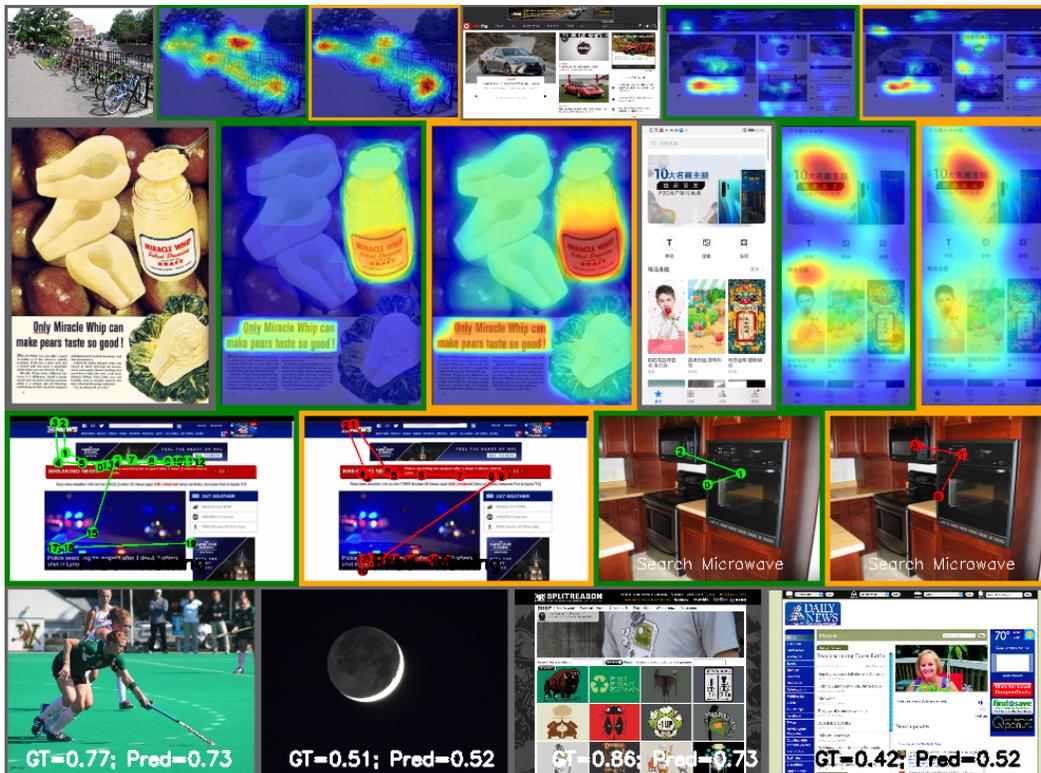


Figure 3: Another set of visualizations on UniAR’s predictions. Images in green border are ground-truth, while images in orange border are UniAR’s predictions. **First row:** saliency heatmap on *Salicon* and *WS-Saliency*. **Second row:** importance heatmap on *Implk*, and saliency heatmap on *Mobile UI*. **Third row:** free-viewing scanpath on *WS-Scanpath* and object-searching scanpath on *COCO-Search18*. **Fourth row:** rating prediction on *Koniq-10k* and *Web Aesthetics* datasets.

A Model Details

The main model components consist of a ViT B16 encoder for image encoding, a T5 base encoder for mixing image and text tokens, and three predictors for rating, heatmap, and scanpath prediction, respectively.

Vision Transformer and T5 Encoder. The ViT B16 encoder uses 16×16 patch size, 12 layers with 12 heads, MLP dimension 3,072, and hidden dimension 768. The T5 base encoder uses 12 layers with 12 heads and MLP dimension 2,048 and hidden dimension 768.

Score Predictor. The score predictor consists of four convolutional layers with Layer Normalization and ReLU activation. The filter size, kernel size, and strides are $[768, 384, 128, 64]$, $[2, 2, 2, 2]$, $[1, 1, 1, 1]$, respectively. Three dense layers of size 2,048, 1,024 and 1 are used to generate a scalar with ReLU activations for the first two layers, and sigmoid for the last.

Heatmap Predictor. The heatmap predictor consists of two convolution layers with filter size, kernel size, and stride as $[768, 384]$, $[3, 3]$, $[1, 1]$, respectively. It then uses four de-convolution layers to up-sample to the required output resolution, with the filter size, kernel size, and stride as $[768, 384, 384, 192]$, $[3, 3, 3, 3]$, $[2, 2, 2, 2]$, respectively. Each de-convolution layer is with two read-out convolution layers of kernel size 3 and stride 1. Layer Normalization and ReLU are used for each layer. In the end, two read-out convolution layers and a final sigmoid activation are used to generate the heatmap prediction.

Scanpath Predictor. The scanpath predictor is implemented using a T5 base decoder with 12 layers of 12 heads and MLP dimension 2,048 and hidden dim 768. Output token length is 64.

We combine the losses from the three predictors, i.e., sequence cross-entropy loss, heatmap L2 loss, and score L2 loss using weights $[1, 500, 50]$ empirically (to make them at the similar scale).

B Dataset Processing

In this section, we describe some of the key dataset processing details.

Table 6: The full table for heatmap prediction results on 7 public datasets spanning digital design and natural images, benchmarking with eight metrics in total. The details on datasets and evaluation metrics can be found in Section 4.1. For *Imp1k* dataset we predict the importance heatmap, while for the rest of the four datasets, we predict attention/saliency heatmap. Our method is highlighted with green background. For each dataset and each metric, the best result in the current column is in **bold**, and the second best result is in **blue**. For our model, the relative performance change compared to the second best result (or the best result if we are not the best) in % is noted.

Dataset	Method	CC ↑	KLD ↓	AUC-Judd ↑	sAUC ↑	SIM ↑	NSS ↑	RMSE ↓	R^2 ↑
<i>Mobile UI</i> [47] (Mobile interface)	Itti et al. [32]'98	0.082	-	0.223	-	0.558	0.126	-	-
	BMS [82]'13	0.131	-	0.249	-	0.206	0.138	-	-
	GBVS [28]'06	0.580	-	0.666	-	0.709	0.591	-	-
	ResNet-Sal [47]'20	0.657	-	0.692	-	0.734	0.704	-	-
	SAM-S2015 [18]'18	0.477	-	0.650	-	0.562	0.537	-	-
	SAM-S2017 [18]'18	0.834	-	0.723	-	0.819	0.839	-	-
	SAM-mobile [47]'20	0.621	-	0.666	-	0.664	0.655	-	-
	UniAR	0.879 +5.40%	0.115	0.756 +4.56%	-	0.832 +1.59%	1.008 +20.14%	0.116	0.777
<i>WS-Saliency</i> [11] (webpage)	Itti et al. [32]'98	0.367	0.840	0.710	0.661	-	0.769	-	-
	Deep Gaze II [41]'16	0.574	3.449	0.815	0.644	-	1.380	-	-
	SalGAN + WS [58]'17	0.637	0.622	0.818	0.703	-	1.458	-	-
	DVA [73]'17	0.571	0.701	0.805	0.711	-	1.260	-	-
	UAVDVSM [29]'19	0.519	0.858	0.739	0.668	-	1.133	-	-
	SAM-ResNet [18]'18	0.596	1.506	0.795	0.717	-	1.284	-	-
	EML-NET [33]'20	0.565	2.110	0.790	0.702	-	1.277	-	-
	UMSI [27]'20	0.444	1.335	0.757	0.698	-	1.042	-	-
	TaskWebSal-FreeView [85]'18	0.525	0.784	0.769	0.714	-	1.107	-	-
	SAM-ResNet + WS [18]'18	0.718	0.994	0.828	0.725	-	1.532	-	-
	DI Net + WS [76]'19	0.798	0.690	0.852	0.739	-	1.777	-	-
	AGD-F (W/o-L) [11]'22	0.815	0.637	0.858	0.753	-	1.802	-	-
	UniAR	0.827 +1.47%	0.299 -53.06%	0.860 +0.23%	0.779 +3.45%	0.737	1.783 -1.05%	0.085	0.691
<i>FiWI</i> [67] (webpage)	DeepGaze II [41]'16	0.488	-	0.797	0.625	-	1.229	-	-
	SAM-ResNet [18]'18	0.595	-	0.791	0.673	-	1.246	-	-
	UMSI [27]'20	0.457	-	0.755	0.675	-	0.938	-	-
	AGD-F [11]'22	0.735	-	0.767	0.748	-	1.606	-	-
	EML-NET [33]'20	0.661	0.603	0.847	0.675	-	1.653	-	-
	EML-NET + Salicon [33]'20	0.689	0.567	0.848	0.697	-	1.722	-	-
	Chen et al. [12]'23	0.699	0.564	0.851	0.704	-	1.752	-	-
	UniAR	0.734 -0.14%	0.571 +0.29%	0.859 +0.94%	0.749 +0.13%	0.627	1.838 +4.91%	0.082	0.544
<i>Salicon</i> [34] (Natural images)	SimpleNet w. ResNet-50 [61]'20	0.895	0.211	0.868	-	0.786	1.881	-	-
	SimpleNet w. PNASNet-5 [61]'20	0.907	0.193	0.871	-	0.797	1.926	-	-
	MDNSal [61]'20	0.899	0.217	0.868	-	0.797	1.893	-	-
	UNISAL [25]'20	0.880	0.226	0.867	0.725	0.771	1.923	-	-
	EML-NET [25]'20	0.890	0.204	0.802	0.778	0.785	2.024	-	-
	UniAR	0.900 -0.77%	0.214 +10.88%	0.870 -0.11%	0.753 -3.32%	0.791 -0.75%	1.946 -3.85%	0.079	0.823
	SALICON [34]'15	0.685	0.575	0.846	-	0.600	1.641	-	-
<i>OSIE</i> [70] (Natural images)	SAM-ResNet [18]'18	0.758	0.480	0.860	-	0.648	1.811	-	-
	UMSI [27]'20	0.746	0.513	0.856	-	0.631	1.788	-	-
	EML-NET [33]'20	0.717	0.537	0.854	-	0.619	1.737	-	-
	Chen et al. [12]'23	0.761	0.506	0.860	-	0.652	1.840	-	-
	UniAR	0.742 -2.50%	0.583 +21.46%	0.862 +0.23%	0.745	0.640 -1.84%	1.789 -2.77%	0.103	0.558
<i>CAT2000</i> [43] (Natural image)	ICF [42]'17	0.780	0.445	0.856	0.619	0.670	1.959	-	-
	DeepGaze II [42]'17	0.795	0.382	0.864	0.650	0.687	1.962	-	-
	UNISAL [25]'20	0.740	0.470	0.860	0.668	0.663	1.936	-	-
	DeepGaze IIE [51]'21	0.819	0.345	0.869	0.668	0.706	2.112	-	-
	SalFBNet [23]'22	0.703	1.198	0.855	0.633	0.643	1.879	-	-
	UniAR	0.870 -0.92%	0.613 +13.96%	0.877 +0.81%	0.615 -7.93%	0.752 +6.52%	2.338 +10.70%	-	-
	Gold Standard	0.969	0.089	0.916	0.787	0.866	2.743	-	-
<i>Imp1k</i> [27] (Graphic design)	Bylinskii et al. [9]'17	0.758	0.301	-	-	-	0.181	0.072	-
	Bylinskii et al. [9]'17	0.732	0.388	-	-	-	0.205	0.061	-
	SAM [18]'18	0.866	0.166	-	-	-	0.168	0.108	-
	UMSI-nc [27]'20	0.802	0.177	-	-	-	0.152	0.095	-
	UMSI-2stream [27]'20	0.852	0.168	-	-	-	0.141	0.105	-
	UMSI [27]'20	0.875	0.164	-	-	-	0.134	0.115	-
	UniAR	0.904 +3.31%	0.124 -25.00%	-	-	0.836	-	0.079 -41.04%	0.823 +61.65%

Imp1k. In the *Imp1k* dataset, we observe some resolution mismatch between the image and its ground-truth importance map. To unify the image and the ground-truth into the same resolution, we find out the lower resolution (with a smaller area) between these two and downsample the larger one into the lower resolution.

WS-Scanpath. We asked Chakraborty et al. [11] for their code on evaluating the scanpath, and follow them to MeanShift clustering to generate spatial bins for the metric **SequenceScore**. We also follow their criteria to select the number of clusters in MeanShift clustering which will be used in the evaluation later on.

Mobile UI. We contacted the authors but the original dataset partition is missing. We randomly partitioned the dataset into a training and a testing set with the number of instances following Leiva et al. [47]. We calculate the saliency results using images without padding.

COCO-Search18. For the scanpath results, we follow the updated experimental results in the appendix of GazeFormer [54]. The original results in the main paper are impacted by an image padding issue, as reported in their GitHub repo.

C Full Results and More Visualizations

We attach another set of visualizations of UniAR's predictions in Figure 3 including heatmap, scanpath, and rating predictions.

Table 7: The full table for scanpath prediction results on natural image and digital design datasets, with object-searching and free-viewing tasks.

Dataset	Method	SemSS \uparrow	SemFED \downarrow	Sequence Score \uparrow	Shape \uparrow	Direction \uparrow	Length \uparrow	Position \uparrow	MultiMatch \uparrow
<i>COCO-Search18</i> [16] (Natural image, object searching)	IRL [77]'20	0.481	2.259	-	0.901	0.642	0.888	0.802	0.833
	Chen et al. [15]'21	0.470	1.898	-	0.903	0.591	0.891	0.865	0.820
	FFM [78]'22	0.407	2.425	-	0.896	0.615	0.893	0.850	0.808
	Gazeformer [54]'23	0.496	1.861	-	0.905	0.721	0.857	0.914	0.849
	UniAR	0.521 +5.04%	2.004 +7.68%	-	0.946 +4.53%	0.724 +0.42%	0.924 +3.47%	0.901 -1.42%	0.874 +2.94%
<i>WS Scanpath</i> [11] (webpage, free-viewing)	Itti et al. [32]'98	-	-	0.177	0.781	0.676	0.778	0.594	0.707
	MASC [2]'17	-	-	0.169	0.788	0.580	0.818	0.514	0.717
	SceneWalker [66]'20	-	-	0.194	0.843	0.616	0.842	0.562	0.716
	G-Eymol [80]'19	-	-	0.218	0.820	0.673	0.816	0.681	0.748
	AGD-F (w. layout) [11]'22	-	-	0.203	0.787	0.642	0.771	0.677	0.719
	AGD-S (w/o layout) [11]'22	-	-	0.221	0.814	0.663	0.805	0.698	0.745
	AGD-S (w. layout) [11]'22	-	-	0.224	0.820	0.677	0.813	0.708	0.755
UniAR	-	-	0.267 +19.20%	0.967 +14.71%	0.826 +22.01%	0.960 +14.01%	0.794 +12.15%	0.887 +17.48%	

We present full tables of UniAR's performance on heatmap and scanpath predictions in Tables 6 and 7, with more baselines and a complete set of evaluation metrics. UniAR offers consistently good predictions on three tasks across multiple datasets, compared to the ground-truths.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Contributions and scope are accurately reflected.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Limitations are discussed in Section 5.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: : No new theory in this paper.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.

- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: See experiments details.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: Not applicable for now.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.

- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: See Section 4.1

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[No\]](#)

Justification: We follow the existing references for experiment results on benchmark datasets, where only average metric number, but no error bar, are reported.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: See Section 4.1

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.

- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: We followed the code of ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: See Section 5

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [Yes]

Justification: See Section 5.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.

- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: They are all publicly available data with proper license.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: No new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No new human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Not applicable.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.