Predicting Visual Importance Across Graphic Design Types

Camilo Fosco[1]**, Vincent Casser**[1]**, Amish Kumar Bedi**[2]**,**
Peter O'Donovan[2]**, Aaron Hertzmann**[3]**, Zoya Bylinskii**[3]**,**
[1]MIT [2]Adobe Inc. [3]Adobe Research
{camilolu, vcasser}@mit.edu {ambedi, podonova, hertzman, bylinski}@adobe.com

ABSTRACT

This paper introduces a Unified Model of Saliency and Importance (UMSI), which learns to predict visual importance in input graphic designs, and saliency in natural images, along with a new dataset and applications. Previous methods for predicting saliency or visual importance are trained individually on specialized datasets, making them limited in application and leading to poor generalization on novel image classes, while requiring a user to know which model to apply to which input. UMSI is a deep learning-based model simultaneously trained on images from different design classes, including posters, infographics, mobile UIs, as well as natural images, and includes an automatic classification module to classify the input. This allows the model to work more effectively without requiring a user to label the input. We also introduce Imp1k, a new dataset of designs annotated with importance information. We demonstrate two new design interfaces that use importance prediction, including a tool for adjusting the relative importance of design elements, and a tool for reflowing designs to new aspect ratios while preserving visual importance.[1]

ACM Classification Keywords

H.5.1 Information Interfaces and Presentation: Multimedia

Author Keywords

Graphic Designs; Importance; Saliency; Human Attention; Automated Design; User Interface for Design; Deep Learning

INTRODUCTION

Where a viewer looks on a poster or advertisement can determine whether the design is effective in communicating its message to the viewer, or if the viewer misses important concepts or details. Indeed, numerous companies offer eye tracking analyses of graphic designs, providing insight into the effectiveness of the design, but these typically depend on a separate eye tracking study for each design. Establishing visual hierarchy and importance is a common task for designers, ideally guiding the reader through the design from elements of

---

[1]The model, code, and importance dataset are available at: http://predimportance.mit.edu.
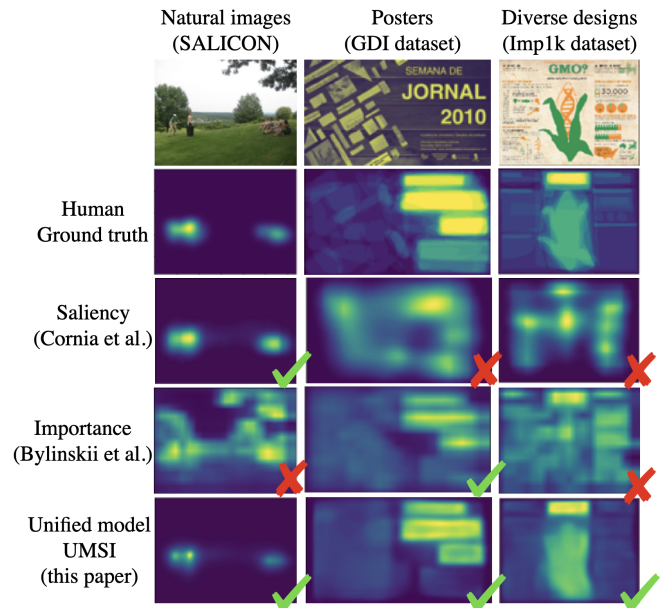
Figure 1. Our Unified Model of Saliency and Importance, UMSI, is able to predict human attention on both natural images and graphic designs. Its performance is comparable to state-of-the-art models specifically built for natural image saliency [13] and it outperforms the existing visual importance models [6, 28].

higher to lesser importance. Methods that predict visual importance accurately in real-time could thus be used in a number of UI applications, including providing real-time feedback to designers, as well as automatic reflow.

Prior approaches have tackled various aspects of attention modeling on graphic designs, including saliency on visualizations [26] and mobile UIs [15], visual flow on comics [9] and webpages [29], and importance prediction on visualizations [6] and posters [6, 28]. Unlike saliency or visual flow which model eye fixations and trajectories, respectively, importance identifies design elements of interest/relevance to the viewer. This makes it more practical as a building block for downstream design applications. Moreover, the narrow focus of prior work on particular design types, datasets, and tasks makes the models difficult to adapt to other problems and generalize to broader sets of design classes. Graphic designs in domains such as advertising, art and education show significant differences in terms of content, layout and appearance, which means that a designer would need to find a model trained specifically for their class of design.

This paper proposes a unified model for predicting importance in graphic designs from multiple design classes, as well as for predicting saliency in natural images. Our model, the Unified Model of Saliency and Importance (UMSI), is a deep neural network that automatically classifies the input design into one of five design classes before predicting the importance or saliency of the input image. Moreover, we introduce a new dataset, Imp1k, that contains 1000 annotated designs, covering webpages, movie posters, mobile UIs, infographics and advertisements. We show that our model either outperforms or remains competitive with state-of-the-art approaches on previous datasets of graphic designs and natural images (Figure 1). We also show that current state-of-the-art models struggle on Imp1k, highlighting a lack of generalization to diverse design classes. Our model successfully generalizes to the attention patterns that are unique to these different design classes.

This approach offers several benefits. First, our model can take many different kinds of designs and images as input, without having to train a separate model for each. A user does not need to explicitly label or classify the input design. Second, the model leverages training data more effectively by using a shared representation, learning what is both common and unique across datasets. Indeed, we show that our unified model gives superior performance on individual tasks as compared to models trained on individual tasks. Third, design represents a spectrum, and our model can generalize accordingly. Since our model trains on both designs and natural images, it works well on designs that include photographs, whether the photograph dominates the design or represents a small portion of it. While our method is already suitable for real applications, our training procedure can be generalized to more design categories in the future, simply by adding more classes.

Furthermore, we introduce two new interactive applications enabled by this model. In our first application, a designer can specify the target importance of an element in an existing design, and the algorithm updates the layout to achieve this goal. Second, we show a reflow application that converts a vector design into a new size and aspect ratio while maintaining the relative importance of design elements. Reflow is a critical problem for modern designers, who are typically tasked to retarget designs to a wide variety of display sizes and form factors. We show that using UMSI in these applications gives substantially better results than with existing baselines.

To summarize, our contributions include: UMSI, a unified model for predicting importance in different kinds of designs, and saliency in natural images; the Imp1k dataset, containing importance annotations for 1000 designs from 5 classes; a tool for revising a design to match target importance values; and a reflow application that uses UMSI to preserve importance.

RELATED WORK
Many previous studies analyze design perception, understanding, and memorability, e.g., [1, 2, 3, 17], including the use of eye tracking as an analysis tool. However, obtaining reliable eye movement data is typically far too costly and time-consuming to be used in most design scenarios.

Most prior work on automatic saliency prediction has focused on eye movements in natural images. Lately, these efforts have progressed significantly with the help of deep learning [4, 12, 13, 19, 22, 23]. When applied to graphic designs, models trained on natural images perform poorly [6, 16]. Specialized saliency models have been designed for visualizations [26], mobile UIs [9] and webpages [29, 31], but with limited generalization ability outside each design class.

In contrast to saliency, O'Donovan et al. [28] introduced the first approach using importance: they collected the GDI dataset, a crowdsourced dataset where people were asked to manually label what they thought was important in a design. With this data, they trained a model for importance prediction on graphic designs. Their method, however, requires annotation of design element position and alignment to generalize to new designs. Bylinskii et al. [6] followed up with the first end-to-end deep learning approach for importance prediction. Their model predicts importance from the original image using a Fully Convolutional Network [25], with pre-trained weights from semantic segmentation, fine-tuned on the GDI dataset. As in prior work [6, 28], we are using a viewer-identified notion of importance rather than a designer-specified notion. We are not explicitly taking into account aesthetics or design heuristics, but are focused on modeling the behavior of the viewer of the content.

Importance and saliency have previously been used for tasks such as retargeting [6, 28] and thumbnailing [20, 30]. For the design process, however, recent work focused on visualizing predicted importance as passive feedback for the designer [6]. In this work, we introduce applications of our model within interactive interfaces where the user can more actively engage with the importance model during the design process. Furthermore, our retargeting application works on vector designs, and is thus more practical for design applications, compared to image-based retargeting [6], and without requiring manual annotations as in prior work [28].

IMP1K DATASET
The previously available graphic design dataset with importance annotations is limited to posters and advertisements [28]. However, the structure of other types of designs like webpages and mobile UIs is quite different, and a model trained to predict importance on posters alone will not generalize. At the same time, there is great interest in predicting attention patterns on webpages and mobile UIs due to their widespread use and commercial impacts. Towards expanding the generalizability of an importance model, we collected a new graphic design dataset covering five diverse design classes and use cases: infographics representing design for knowledge transfer, webpages and mobile UIs representing design for utility, and advertisements and movie posters representing design for promotion. While prior work addresses importance for data visualizations [6], we omit them because their highly structured layouts (axes, data marks, etc.) make them less amenable to the design applications that are the focus of this work. The details of the data collection and annotation are provided below, along with an analysis of the importance patterns common to, and differing between, the design classes.

Dataset collection

We collected designs from existing research datasets so that the stimuli and annotations may be shared broadly with the research community. Infographics were sourced from Visually29k [33], webpages from ClueWeb09 [7], movie posters from Chu et al.'s movie poster dataset [11], mobile UIs from RICO [14], and advertisements from the Pitt Image Ads [18]. To compose our multi-class importance dataset, Imp1k (Fig. 2), we randomly sampled 200 designs per design class, after filtering out designs that had too few elements, skewed aspect ratios, or that were outliers among their design class.

For annotating the importance in each of the 1000 designs from our Imp1k dataset, we used the ImportAnnots UI [27]. Based on the original implementation of O'Donovan et al. [28], this UI provides participants with three options for annotation: (i) "regular stroke" involves painting with paintbrush-like strokes with the mouse button held down, (ii) "polygon stroke" allows drawing polygons through connecting vertices, and (iii) "fill stroke" fills an area delimited by the initial mouse click position and the most recent mouse position, allowing smooth tracing over continuous contours of a shape. Using any of these tools, a participant can produce a binary mask to annotate the most important design elements.

We deployed the ImportAnnots UI on Amazon's Mechanical Turk (MTurk) to collect importance data on all 1000 designs, splitting them into 10 designs of the same class per task (HIT), and requesting 30 participants per HIT. Participants could complete several HITs. Data from 249 individual participants was collected. To ensure good data quality, we additionally included 3 sentinels per HIT. These are simple, artificial designs where the importance annotation is expected to cover the only visible element in the design. We designed 40 such sentinels, which we manually annotated with ground truth importance. For the task to proceed, we automatically verified that participant annotations matched the ground truth annotations (with a computed intersection over union value over 0.6) for at least 2/3 of the sentinels. Participants were paid $1.00 for completing the task. For each of the 1000 designs collected, we computed an associated importance map by averaging the 25-30 individual binary annotations. While each participant's importance annotation is subjective, prior work has reported that the average over participants produces heatmaps that approximate attention maps [6, 21, 28].

Analysis

Each class in our dataset exhibits slightly modified patterns of importance (Fig. 3). We qualitatively observe that importance on movie posters is dominated by the title and human faces, while webpages tend to draw attention to the site name or company running the site, usually on the top left. Infographics have importance distributed across the full design, as do mobile UIs. Ads, on the other hand, concentrate most of their importance on a few elements in the center of the design. Across classes, human annotations tend to highlight a set of up to 7 different elements; more detailed annotations were rare in our data. We performed a basic analysis of text and face importance on different design types. Specifically, we used open source face [24] and text [32] detectors and

Figure 2. Examples of the 5 classes of the Imp1k dataset.

calculated the mean importance map value over the face and text bounding boxes per design. We observe that the relative importance of faces is highest on ads and movie posters, where they are nearly tied in importance. The importance of text is highest in ads, where it is perceived to be on average 32% more important than the text in infographics (the class with the second highest average importance of text). These differences in importance patterns across classes motivate the development of a model that can generalize to these designs, and tailor predictions for each design class.

UMSI MODEL

Motivated by these observations, we introduce a unified importance prediction model that is trained jointly on different design classes as well as natural images. The model contains a classification branch that infers the input design category, allowing it to produce results appropriate for the input design.

Model architecture

Given an input image, our objective is to predict a heatmap with an importance value assigned to every pixel. This is related to other image-to-image prediction tasks like saliency and segmentation, which motivate our architectural choices. Encoder-decoder architectures are common for such tasks, whereby features are first downsampled (encoded) before being upsampled (decoded) to produce the final prediction at the original image size. The previous state-of-the-art importance model [6] was based on a semantic segmentation architecture [25]. Similarly, we use a segmentation-inspired pyramid pooling module to leverage features at different scales and improve the fidelity of the heatmaps. The addition of a classification branch helps our importance model make use of class-specific information. The full architecture is visualized in Fig. 4 and the technical details of each component are detailed below.

Encoder. This is a low-level feature network that extracts basic image features. We used the Xception encoder, a high-performing state-of-the-art network common for segmenta-
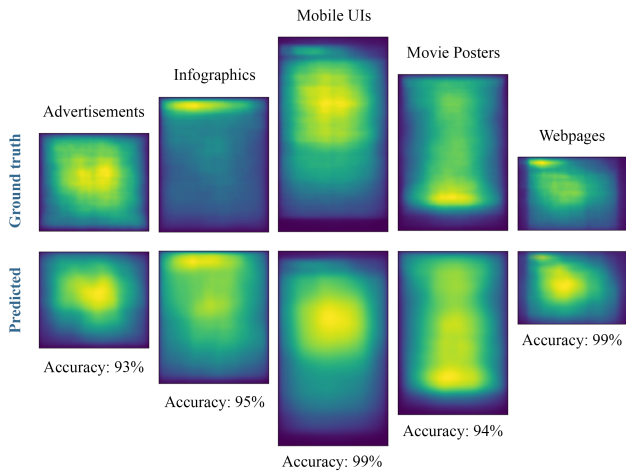
Figure 3. General importance trends on Imp1k across ground truth maps and our predicted maps. Our model captures the trends correctly, evidencing a pronounced class specialization. This is further supported by UMSI's high classification accuracy: our model learns to recognize the input's class and tailor its output to the particularities of that class.

tion [10] and saliency [8] tasks. It is composed of depthwise separable convolutions, which come with great computational savings at a minimal cost on accuracy. We modified this network to output feature maps by removing the final global average pooling and fully-connected layers. To obtain higher resolution maps, we reduced the strides of both the 1x1 residual connection and the last two max pooling layers to one.

Atrous Separable Pyramid Pooling (ASPP). To aggregate image information at multiple scales, we used an ASPP module [10] after our encoding module. This module applies several dilated convolutions with different dilation rates in parallel to a given set of feature maps. The output maps are then concatenated and passed through a projecting convolutional layer to reduce the channel dimensionality of the feature block. We used an ASPP module with 4 parallel convolutional layers with 256 filters, with dilation rates of 1, 6, 12 and 18.

Decoder. Our decoder is a set of convolutions followed by upsampling layers and dropout. This allows the feature maps to be scaled up to the original image size, without a significant increase in parameters. We use three different blocks: the first two composed of two convolutional layers followed by a 2x2 upsampling layer, and the last composed of one convolutional and one upsampling layer. A 1x1 convolutional layer then transforms the feature maps into our final importance map. The convolutional layers have kernels of size 3 and the dropout applied is 0.3. We use ReLU activations throughout.

Classification. The output of our encoder is also passed in parallel to a classification submodule. It is composed of two 3x3 convolutional layers, followed by a convolutional layer with stride 3 to reduce dimensionality, then another 3x3 convolutional layer. The resulting feature maps are globally pooled and a dense layer with dropout transforms the maps into a 1x1x256 feature vector C_f. This vector is used both (i) for directly outputting the predicted image class, and (ii) for intro-

ducing class-specific information for importance prediction. In the first case, C_f is fed to a dense layer with 6 outputs and softmax activation, to produce a 1x6 vector of class probabilities. In the second case, C_f is resized to the required dimensions and concatenated with the output from the ASPP module along the channel dimension as part of the Concatenation layer.

Model training
Our challenge is to produce a model with good performance both on natural images and graphic designs, which tend to have quite different layouts and attributes (e.g., graphic elements, text regions, etc.). Moreover, the available ground truth attention data on natural images and graphic designs has been collected using different means, producing differences in formats and data quantity. The largest available dataset commonly used for training saliency models on natural images is SALICON, containing 20K images with mouse movements aggregated into saliency maps [19]. The only available importance dataset for graphic designs until this paper was GDI [28], containing 1000 designs with binary annotations aggregated into importance maps. Our Imp1k dataset also contains 1000 designs, but split across 5 distinct classes, while GDI is composed mainly of ads and posters from Flickr. Training a model to learn the attention patterns on both natural images and graphic designs, while using differently sized and formatted datasets, involved a specialized procedure, detailed below.

Training procedure. We first pre-train our model on SALICON for 10-15 epochs. This allows our model to learn about image features broadly relevant to saliency. We then fine-tune on a graphic design importance dataset - which is either GDI or Imp1k, depending on the particular evaluation (next section). To prevent the network from "forgetting" saliency while learning design importance, we mix in 160 new SALICON images during each fine-tuning epoch. This number maintains class balance when training on Imp1k, since the training set consists of 80% of 200 designs per class.

Training details. We use the 10K training, 5K validation, and 5K test splits from SALICON [19]. For GDI, we use the same training/testing split as in [6]. For Imp1k, we define a test set by randomly selecting 20% of images from each class. We train with KL and CC losses, with coefficients of 10 and -3, respectively. A binary cross-entropy loss with a weight of 5 is used for the classification submodule. The learning rate is $1e^{-4}$, and is decreased by a factor of 10 every 3 epochs. To limit overfitting, we use a dropout of 0.3 on all layers. We train with a batch size of 8 using the Adam optimizer.

MODEL EVALUATION
We next compare with existing methods, on our dataset as well as previous importance datasets. We compare different variants, including retraining previous models on our data, and training subsets of our model (e.g., without classification), in order to evaluate each element of our method separately.

Importance prediction
We evaluated our model on two importance datasets, GDI and Imp1k, the results of which can be found in Tables 1 and 2, respectively. On both datasets we compare against, and
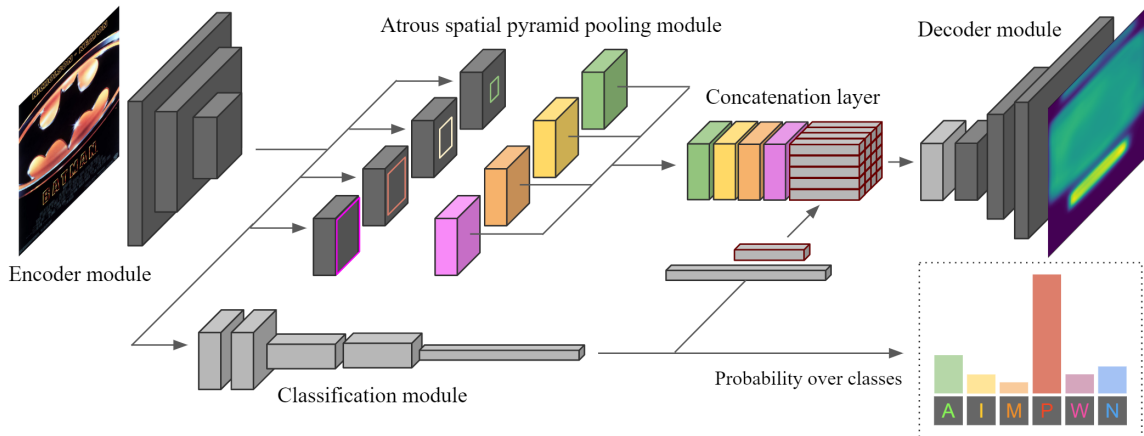
Figure 4. The architecture of our Unified Model of Saliency and Importance (UMSI). Our model leverages Xception-based encoder and decoder modules to generate accurate maps with few parameters. An atrous spatial pyramid pooling (ASPP) module captures features at different scales, and a separate classification module classifies the input into one of six categories, represented by colored letters in the figure: Ads (A), Infographics (I), Mobile-UIs (M), Movie posters (P), Webpages (W) and Natural Images (N). Central to our model is the concatenation layer, which combines the outputs from the ASPP with class-specific features obtained from the classification module and allows for tailored heatmaps for each input class.

outperform, the prior state-of-the-art importance model [6] across all four evaluation metrics. On the GDI dataset, UMSI notably achieves an improvement of 54% over the prior model according to the Pearson's Correlation Coefficient (CC metric), the recommended metric for saliency evaluation [5].

	$R^2 \uparrow$	$RMSE \downarrow$	$CC \uparrow$	$KL \downarrow$
O. et al., Auto [28]	0.539	0.212	-	-
O. et al., Full [28]	0.754	0.155	-	-
B. et al. [6]	0.576	0.203	0.596	0.409
B. et al. + O. [6]	0.769	0.150	-	-
SAM [13]	0.693	0.146	0.884	0.102
UMSI	**0.781**	**0.120**	**0.915**	**0.086**

Table 1. Evaluation of models on importance prediction on the GDI dataset. We compare our model, UMSI, to the prior state-of-the-art importance models [6] and [28], and to a strong saliency model finetuned on importance [13]. Our model outperforms all other alternatives.

	$R^2 \uparrow$	$RMSE \downarrow$	$CC \uparrow$	$KL \downarrow$	$Acc \uparrow$
B. et al. [6]	0.072	0.181	0.758	0.301	-
B. et al. x5 [6]	0.061	0.205	0.732	0.388	-
SAM [13]	0.108	0.168	0.866	0.166	-
UMSI-nc	0.095	0.152	0.802	0.177	-
UMSI-2stream	0.105	0.141	0.852	0.168	0.91
UMSI	**0.115**	**0.134**	**0.875**	**0.164**	**0.98**

Table 2. Evaluation of models on importance prediction on the Imp1k dataset. We compare our model, UMSI, to the prior state-of-the-art model trained on the entire dataset (B. et al. [6]) and on each class separately (B. et al. x5 [6]). We include SAM [13], a strong saliency model also fine-tuned on the Imp1k dataset. UMSI-nc corresponds to UMSI without a classification module, while UMSI-2stream is an alternative architecture with classification. UMSI outperforms all other alternatives.

We additionally surpass the performance of the "Full" O'Donovan model [28], a non-automatic approach which requires human annotations of the graphical element locations at test-time. The poorer-performing, fully-automatic O'Donovan

model is also included for comparison. The best-performing model out of the alternatives was a combination of the two prior models [5] and [28], which was previously reported in [5]. UMSI outperformed this combined model too.

As another comparison point, we trained the state-of-the-art saliency model SAM [13] to predict importance, by first pre-training on SALICON, and then fine-tuning on the GDI and Imp1k datasets, to report performances on each dataset. On the Imp1k dataset, we also re-trained five instances of the prior state-of-art importance model [6] on each of the 5 design classes, separately ("B. et al. x5"). From Tables 1 and 2, we see that UMSI outperforms these alternative re-trained models, demonstrating the contributions of our architecture design, beyond the training data alone.

Example predictions can be found in Fig. 7. Training on both natural images and graphic designs allows our model to correctly distribute importance in images that contain different amounts of text and visuals. In these cases, a saliency model might over-predict visual regions, whereas an importance model trained only on posters and ads (such as [6]) struggles to generalize to more complex designs.

Saliency prediction
We also evaluated our model on the SALICON test set [19]. On saliency, our model performs comparably to state-of-the-art. UMSI obtains a CC of 0.782 and KL of 0.341, compared to a CC of 0.811 and KL of 0.324 for SAM [13]. Qualitatively, we observe very similar patterns as SAM on the output heatmaps: our model correctly detects people and faces, and identifies elements of high contrast (see Figs. 1 and 5a).

The ability of our model to generalize to natural images and graphic designs alike make it usable within a typical design workflow. As a proof-of-concept, we generated some designs with a SALICON image as background, and collected importance annotations on designs corresponding to different stages of the design process. On natural images, our model predicts
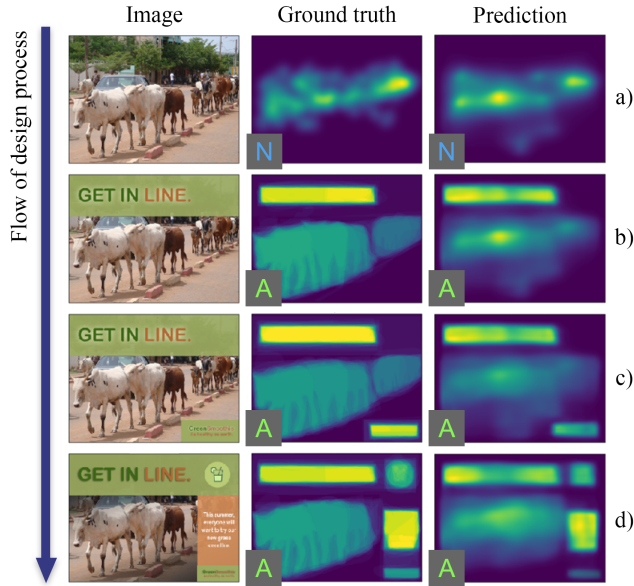
Figure 5. Our model updates importance predictions throughout the design process. Letters on the lower left represent ground truth and predicted design classes. (a) On a background natural image, UMSI predicts traditional saliency. This can inform the designer about which parts to emphasize or de-emphasize. (b) The designer adds a title element to the image. The model predicts the design is an advertisement and that the title is now most important. (c-d) As the designer adds more elements, the output map correctly captures their importance and diminishes the importance of the image background.
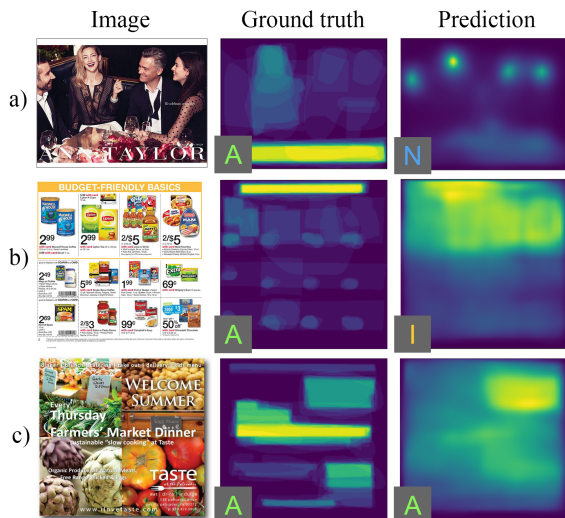

Figure 6. Failure modes of UMSI. Letters on the lower left represent ground truth and predicted design classes. (a) Our model classifies the input as a natural image, missing the subtle text overlay, and over-emphasizing the faces. (b) Our model missclassifies this ad as an infographic due to the multitude of details. Nevertheless, importance is still highest at the top of the design. (c) A complex and busy design poses a challenge for our model. Although the classification is correct, the model selects a relevant but secondary text overlay as most important.

accurate saliency maps. As the designs become more complex, our model switches to distributing importance across the image and design elements (Fig. 5). We also observe a diminished text bias, as our model accurately recognizes the importance of secondary design elements over text. We note that excessive text bias was one of the limitations of [6].

Classification

Our classification module achieves an average accuracy of 95% when determining which of 5 classes the designs from Imp1k belong to (per-class accuracy scores are reported in Fig. 3). From Fig. 3, we see that our model has learned to capture general trends of importance that differentiate the 5 design classes. We include some interesting failure cases of our model, when it misclassifies designs, in Fig. 6.

We evaluated the contribution of our classification module to importance prediction (Table 2). To do so, we trained an additional two versions of UMSI, one without the classification module at all (UMSI-nc), and one where the classification module does not feed directly back into the importance prediction (as in Fig. 4), but branches off after the ASPP module, and still affects the features learned by both the encoder and ASPP module (UMSI-2stream). Both alternatives to our architecture performed worse, where not having a classification stream at all (UMSI-nc) affected the scores most.

APPLICATIONS

In this section we present two new applications of our importance prediction model. These are proofs-of-concept intended to demonstrate how such a model can help support iterative design and automate some common design workflows.

Model-assisted interactive design

Previous work used importance prediction within a design tool as passive feedback in the form of real-time importance visualization [6]. Here we consider a use case where the user can more actively engage with the importance model during the design process to receive design suggestions. We developed a bare-bones prototype with minimal design support to evaluate the ability of our model, when coupled with an optimization procedure, to actively adjust a design according to user-imposed constraints.

The workflow allows a user to manipulate design elements on a canvas, and to receive immediate feedback about the predicted importance of each element, similar to the visualizations in [6]. However, different from [6] is a new ability to directly interact with the importance scores of each design element, allowing the user to specify constraints to increase or decrease the importance of design elements.

The application then depends on an optimization procedure to generate new design variants, which are scored by the importance model. For this demo, we chose a genetic algorithm that makes adjustments to the visual arrangement (scale and location) of the design elements and selects adjustments that reduce the gap between the current predicted values and the target values specified by the user. As the optimization algorithm itself is not a contribution of this paper, we leave its implementation details to the Supplemental Material.
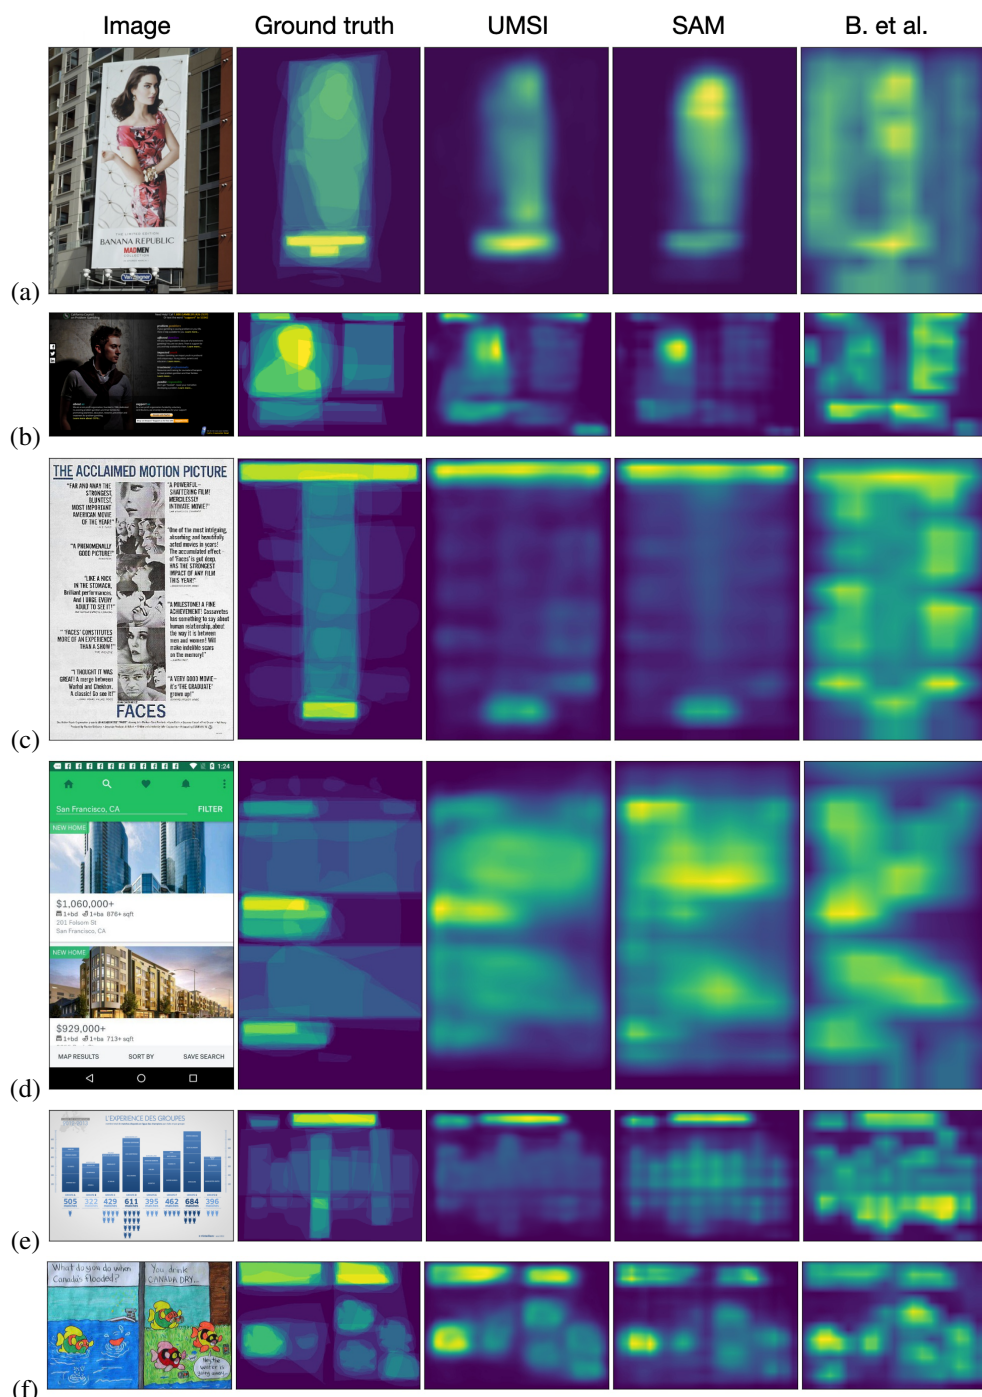
Figure 7. Example images and model predictions from the Imp1k dataset. (a) In this advertisement, the brand is rated as the most important part of the design by human annotators (ground truth). (b) For this webpage and the advertisement above, the natural image saliency model (SAM) over-focuses on the faces. (c) A movie poster, where the prior state-of-the-art importance model (B. et al. [6]) over-predicts all the text elements as having similar importance. (d) A mobile UI, composed of text and images in equal parts, where SAM focuses on the images and B. et al. distributes importance evenly. In all the above examples, UMSI correctly detects the most important design elements, regardless of their spatial location on the image. (e) A challenging infographic for all the models, where correctly predicting importance requires an understanding of the data presented. (f) A drawing that does not belong to any other design class, but nevertheless generates correct predictions of importance by the first two models.

UI design. We developed a simple UI for this task which allows drag-and-drop placement of graphical elements (Fig. 8A-B). A real-time (TCP network-based) interface between our front-end UI and our back-end importance model allows for remote computation on GPU-enabled computing infrastructure. As users manipulate the design, continuous execution of the importance model in the back-end allows real-time updates of the importance predictions, visualized as a heatmap (Fig. 8C). We also compute the individual importance of all the design elements by averaging the importance values inside the mask of each design element. We plot these element-wise importance values as an interactive bar plot (Fig. 8D). The user can then set individual target importance values for each design element, by adjusting the level of the respective bar in the plot. This interaction triggers an optimization procedure that tries to reduce the discrepancy/gap between the current predicted and target importance values (Fig. 8E).

User studies. We evaluated the ability of our importance model to accurately guide the optimization procedure towards design variants matching user-specified constraints. Using 6 initial designs, we selected 3 elements per design that we wanted to separately increase the importance of. Specifically, we used the interactive bar plot within our UI to constrain the importance of the chosen design element to have maximal importance in the final design. We then launched the optimization procedure for 5 consecutive runs, to produce 5 possible design variants that meet the specified constraints. We produced a total of 90 design variants = 6 (initial designs) x 3 (separate constraints) x 5 (runs). We then launched the ImportAnnots UI to collect importance maps (25 participants worth of annotations per design) for all 90 design variants, as well as for the 6 initial designs. We compared these ground truth importance maps to the ones predicted by our model and used by the optimization procedure to produce each design variant. The average CC score across all 90 variants was 0.928 (where the upper bound of CC is 1.0), indicating a high correspondence between the model and ground truth. On average, across the 6 designs x 3 constraints per design, the optimization algorithm succeeded in increasing the importance of the target element in 87% of the runs.

Discussion. Some example results are provided in Fig 9. The top example contains a simple advertisement where a user wants to emphasize the discount. Our application automatically rearranges the elements such that the design element selected is enlarged and takes a more central position. In the bottom example, a user wants to put more emphasis on the location of the advertised event, while also decreasing the prominence of the event title. In the automatic redesign, text about the event time and date is separated from the location, creating more room to enlarge the relevant information.

In this section we provided a means by which our importance model can be coupled with a separate optimization procedure to offer design suggestions in an interactive design tool. The design tool itself is bare-bones and does not optimize for balance, symmetry, or other aesthetic properties, and so the returned results are not guaranteed to be good quality designs. Despite this, we were able to guarantee, with 87% success rate,
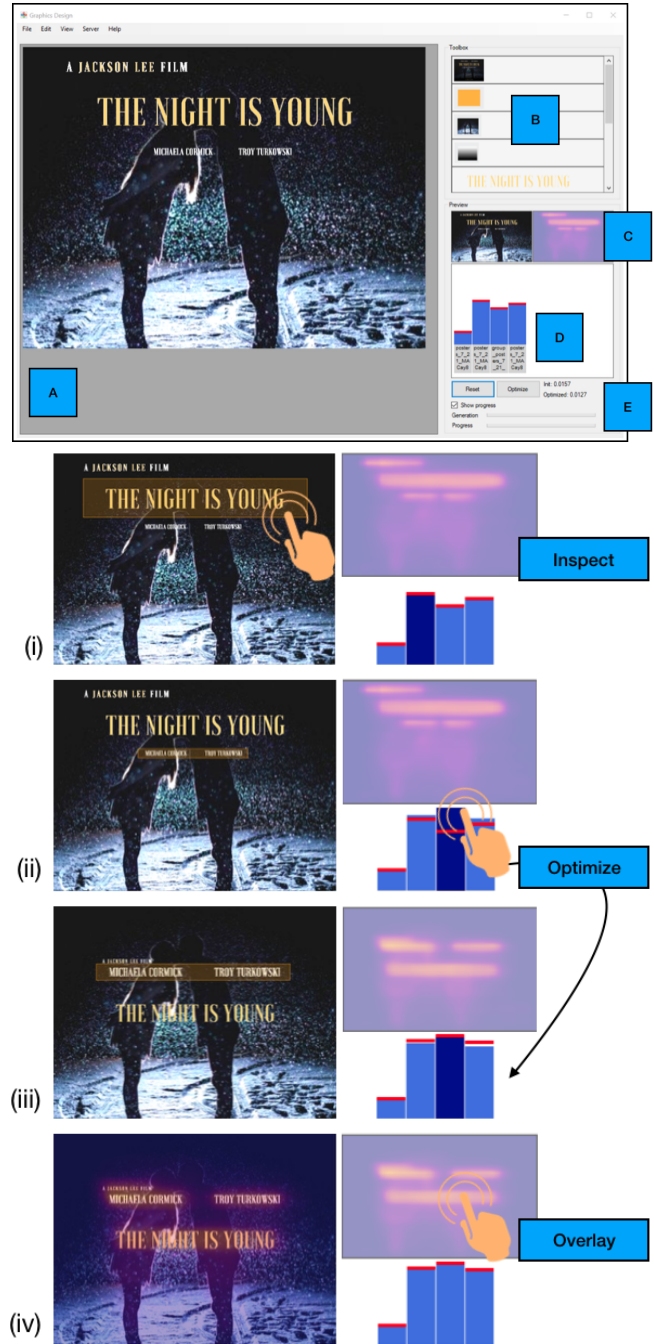

Figure 8. At the top we demonstrate our model-assisted interactive design UI: A) Canvas onto which a user can drag a vector graphics file. B) A set of layers that can be used to compose design elements together. C) A preview of the interactively-computed importance map for the current design. D) An interactive bar plot listing all the design elements and their predicted importance scores. E) A progress bar for the optimization procedure launched when users adjust the interactive bar plot values. Below the UI is a sample interaction sequence: (i) Clicking on a design element highlights the element's predicted importance score in the interactive bar graph. (ii) Making adjustments to the bar graph launches an optimization procedure that automatically adjusts the design (iii) until the target importance values are reached for all the elements. (iv) Selecting the importance heatmap overlays it on the final design.

Figure 9. **Model-assisted design examples. Left: initial designs. Right: designs after optimization. We maximize the importance of elements marked in green. The algorithm accomplishes this through spatial rearrangement (top) and/or changing proportions (bottom). Changing target importance of one element commonly also affects neighboring elements to preserve their importance levels as closely as possible while giving room for the other element to gain importance.**

that the importance objectives specified by the user were met, in accurately increasing the importance of specified design elements. These initial investigations show promise for the utility of an importance model in future design tools. How best to incorporate AI and other forms of automatic assistance and collaboration in interfaces is a substantial open problem for the field.

Model-assisted design reflow

We also investigated the benefits of using importance prediction as a back-end mechanism for creating different sized variations of an input design. This is required for adapting designs to alternative form factors and devices. We propose using the importance scores of design elements computed from the importance heatmaps to guide repositioning and rescaling of the elements when reflowing a graphic design (Fig. 10).

We worked with the developers of a commercially-available layout application to integrate our importance models into design reflow. The application lets users create or upload their existing designs and perform basic operations, like placing, scaling, and grouping graphical elements. Once the user is content with the initial design, the design is sent as an image to our back-end importance model. The predicted importance heatmap is then used to rank the graphic elements by importance, by computing the average importance value over each element's extent. Our proposition is that by preserving the relative importance values of all the design elements in the reflow result, we maintain the designer's intention in allocating attention across design elements.

Reflow algorithm. Given an input design and a new target design size (Fig. 10a), the reflow algorithm selects a new

position and scale for each design element (Fig. 10b). For designs of a few different sizes and numbers of elements, we manually composed templates indicating the importance rank of each element in that design (examples in Fig. 10c). We composed these templates based on professional designs of similar sizes and number of elements. The templates are stored in the reflow application. At test-time, we retrieve the template matching the input design in number of design elements, and with similar aspect ratio to the target design size. Next, we map each element from the initial design to a placeholder in the template with a matching importance rank. This step preserves the importance ranks of the elements from the original design.

User studies. To evaluate whether using our importance model at the back-end of a layout application indeed improves the final reflow results, we ran two user studies, with crowdworkers and with professional designers, respectively. For 17 different graphic designs, we provided study participants with three automatically-computed design variants in an alternative aspect ratio. Participants were asked to rank the variants from 1 (best) to 3 (worst), based on personal judgement without further guidance. These variants corresponded to: (i) a baseline implementation of reflow without using importance[2], and results from the importance-aided reflow algorithm described above, using: (ii) the previous state-of-the-art importance model [6] and (iii) our UMSI model. The ordering of the variants was randomized across designs and participants. We ran the study on Amazon's MTurk, and recruited 100 participants, each presented with 9 randomly-sampled designs and 3 repeat designs, and compensated \$1.50. To ensure high quality data, we only kept the participants that consistently ranked the variants on 2/3 of the repeat designs. The data of 43 participants (29 male, most in their 20s and 30s) were used in the resulting analyses. Based on an average of 23 participant rankings for each of the 17 designs, we found that the reflow results using the UMSI model performed significantly better ($p < 0.05$, Bonferonni-corrected) than the other two reflow variants (Fig. 11a). The difference between importance-based reflow using [6] and the baseline was not significant, however. We repeated the task with 6 professional designers recruited from personal networks (3 male, between 25-45 years old, with a variance of 3-20 years of design experience). Each professional was asked to rank the three design variants for all 17 designs, and was compensated \$5.00. Based on their average rankings of the design variants, we confirm that the reflow results using the UMSI model performed better than the other two reflow variants (Fig. 11b). Examples of reflow results from the three methods compared can be found in Fig. 10b,d,e, along with their average ranks from the MTurk study.

Discussion. While we evaluated the benefit of guiding reflow using an importance model within a particular design tool, we believe that the observed gains in the quality of reflow results are encouraging. Future reflow applications could incorporate importance as an additional constraint (i.e., by preserving the importance ranks of design elements in the original and redesign versions), when combined with balance, symmetry, and other common aesthetic metrics. Indeed, one

---

[2]An existing constraint-based engine that was the default in the commercially-available layout application we used.
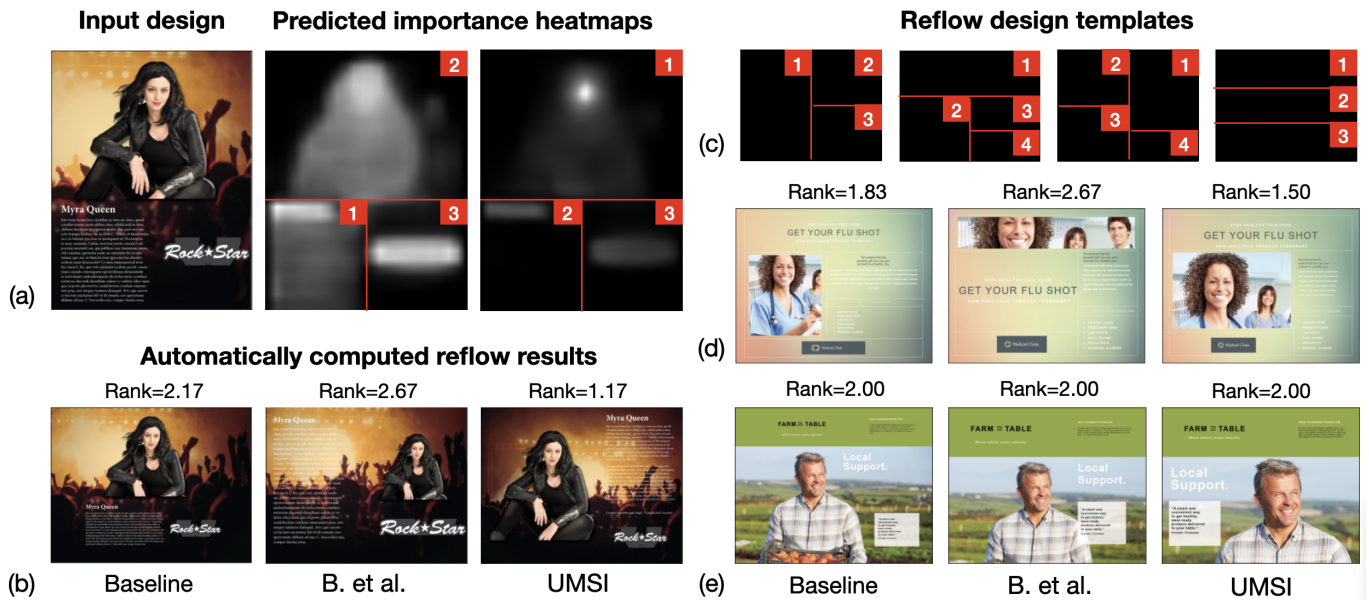
Figure 10. Model-assisted design reflow. (a) An input design and importance heatmaps predicted by B. et al. [6] and our UMSI model, respectively. These heatmaps are used to rank the individual design elements by importance (numbered insets). (b) During reflow, design elements are then moved and rescaled by being mapped to placeholders in retrieved templates. (c) Template examples. The first template was used for the results in (b), while the second template was used in (d) and (e). (d,e) More examples of reflow results, comparing a baseline (described in text), to importance-aided reflow with the model B. et al. [6] and with UMSI. Crowdsourced ranking scores are included above each variant (the lower, the better).
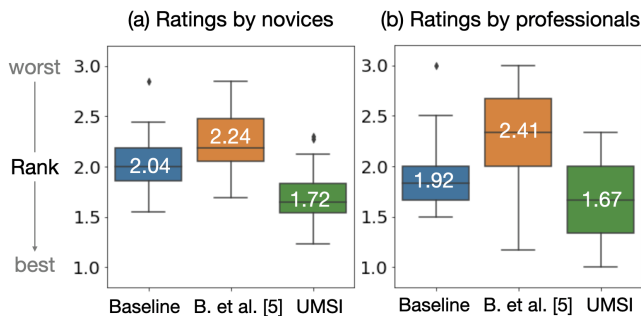

Figure 11. User study results for model-assisted design reflow. We compare a baseline (no importance model used) to reflow aided by a previous importance model [6] and UMSI, respectively. All 3 reflow variants were ranked from 1 (best) to 3 (worst) for 17 different designs by an average of 23 novice crowdworkers (a) and 6 professional designers (b). UMSI-aided design reflow results were rated significantly better than the alternatives by both participant pools.

of our professional designer participants, who was naive to the purpose of the study and the sources of the retargeting results, described that what makes up a good reflow result is *"the hierarchy of elements, making sure we read the most important thing first."* While our approach only modifies element size and position, future work could consider colors, fonts, and other features to affect visual importance during reflow.

Notably, as seen from our user study, we achieved improved reflow results with our proposed UMSI importance model, but not with the prior state-of-the-art model [6]. The UMSI model generalizes better to more diverse and complex designs and its predictions are less biased towards text (Fig. 10b,d) - properties that make our importance model more amenable to

the reflow application. In Fig. 10e, we see an example where the reflow results of all methods were scored comparably, with the key difference being that each successive model zoomed in more on the man's face. UMSI was trained on natural images in addition to graphic designs, and as a result has learned that faces are very salient to human observers, and should be prioritized. The baseline model produces a result that is most like a crop of the input design. The original designs and more reflow results can be found in the Supplemental Material.

CONCLUSION

In this paper, we presented the first Unified Model of Saliency and Importance, capable of approximating a notion of human attention on natural scenes and graphic designs alike. We showed that our model is competitive with state-of-the-art models that have been specialized for different tasks: predicting natural scene saliency, and graphic design importance, respectively. Not only can our model generalize to these broadly different input modalities, it can also provide accurate fine-grained predictions on different design classes, which makes two new interactive design applications possible. We presented a model-assisted design application that automatically adjusts the elements in a vector design to meet user-specified importance constraints; and a graphic design reflow application that automatically adjusts the locations and sizes of design elements to fit to new aspect ratios. The model architecture presented in this paper is not limited to the design classes in Imp1k. The training procedure can be adapted, given a new training set, to any number of additional design classes. This opens up the possibility of continually improving the current UMSI model to make it a one-stop shop for predicting attention on any images, natural and designed.

REFERENCES

1. Michelle A Borkin, Zoya Bylinskii, Nam Wook Kim, Constance May Bainbridge, Chelsea S Yeh, Daniel Borkin, Hanspeter Pfister, and Aude Oliva. 2016. Beyond memorability: Visualization recognition and recall. *IEEE transactions on visualization and computer graphics* 22, 1 (2016), 519–528.

2. Michelle A Borkin, Azalea A Vo, Zoya Bylinskii, Phillip Isola, Shashank Sunkavalli, Aude Oliva, and Hanspeter Pfister. 2013. What makes a visualization memorable? *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2306–2315.

3. Zoya Bylinskii, Michelle A Borkin, Nam Wook Kim, Hanspeter Pfister, and Aude Oliva. 2015a. Eye fixation metrics for large scale evaluation and comparison of information visualizations. In *Workshop on Eye Tracking and Visualization*. Springer, 235–255.

4. Zoya Bylinskii, Tilke Judd, Ali Borji, Laurent Itti, Frédo Durand, Aude Oliva, and Antonio Torralba. 2015b. Mit saliency benchmark. (2015).

5. Zoya Bylinskii, Tilke Judd, Aude Oliva, Antonio Torralba, and Frédo Durand. 2019. What do different evaluation metrics tell us about saliency models? *IEEE transactions on pattern analysis and machine intelligence* 41, 3 (2019), 740–757.

6. Zoya Bylinskii, Nam Wook Kim, Peter O'Donovan, Sami Alsheikh, Spandan Madan, Hanspeter Pfister, Fredo Durand, Bryan Russell, and Aaron Hertzmann. 2017. Learning visual importance for graphic designs and data visualizations. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. ACM, 57–69.

7. Jamie Callan, Mark Hoy, Changkuk Yoo, and Le Zhao. 2009. Clueweb09 data set. (2009).

8. Pat Sukhum Yun Bin Zhang Nanxuan Zhao Aude Oliva Zoya Bylinskii Camilo Fosco, Anelise Newman. 2020. How much time do you have? Modeling multi-duration saliency. *Proceedings of the IEEE conference on computer vision and pattern recognition* (2020).

9. Y. Cao, R. Lau, and A. B. Chan. 2014. Look Over Here: Attention-Directing Composition of Manga Elements. *ACM Transactions on Graphics (Proc. of SIGGRAPH 2014)* 33 (2014). Issue 4.

10. Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. 2017. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587* (2017).

11. Wei-Ta Chu and Hung-Jui Guo. 2017. Movie Genre Classification based on Poster Images with Deep Neural Networks. (2017).

12. Marcella Cornia, Lorenzo Baraldi, Giuseppe Serra, and Rita Cucchiara. 2016. A deep multi-level network for saliency prediction. In *2016 23rd International Conference on Pattern Recognition (ICPR)*. IEEE, 3488–3493.

13. Marcella Cornia, Lorenzo Baraldi, Giuseppe Serra, and Rita Cucchiara. 2018. Predicting human eye fixations via an lstm-based saliency attentive model. *IEEE Transactions on Image Processing* 27, 10 (2018), 5142–5154.

14. Biplab Deka, Zifeng Huang, Chad Franzen, Joshua Hibschman, Daniel Afergan, Yang Li, Jeffrey Nichols, and Ranjitha Kumar. 2017. Rico: A mobile app dataset for building data-driven design applications. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. ACM, 845–854.

15. Prakhar Gupta, Shubh Gupta, Ajaykrishnan Jayagopal, Sourav Pal, and Ritwik Sinha. 2018. Saliency Prediction for Mobile User Interfaces. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 1529–1538.

16. Michael J Haass, Andrew T Wilson, Laura E Matzen, and Kristin M Divis. 2016. Modeling human comprehension of data visualizations. In *International Conference on Virtual, Augmented and Mixed Reality*. Springer, 125–134.

17. Lane Harrison, Katharina Reinecke, and Remco Chang. 2015. Infographic aesthetics: Designing for the first impression. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 1187–1190.

18. Zaeem Hussain, Mingda Zhang, Xiaozhong Zhang, Keren Ye, Christopher Thomas, Zuha Agha, Nathan Ong, and Adriana Kovashka. 2017. Automatic understanding of image and video advertisements. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1705–1715.

19. Ming Jiang, Shengsheng Huang, Juanyong Duan, and Qi Zhao. 2015. Salicon: Saliency in context. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1072–1080.

20. Binxing Jiao, Linjun Yang, Jizheng Xu, and Feng Wu. 2010. Visual summarization of web pages. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*. ACM, 499–506.

21. Nam Wook Kim, Zoya Bylinskii, Michelle A Borkin, Krzysztof Z Gajos, Aude Oliva, Fredo Durand, and Hanspeter Pfister. 2017. BubbleView: an interface for crowdsourcing image importance maps and tracking visual attention. *ACM Transactions on Computer-Human Interaction (TOCHI)* 24, 5 (2017), 36.

22. Matthias Kümmerer, Lucas Theis, and Matthias Bethge. 2014. Deep gaze i: Boosting saliency prediction with feature maps trained on imagenet. *arXiv preprint arXiv:1411.1045* (2014).

23. Matthias Kümmerer, Thomas SA Wallis, and Matthias Bethge. 2016. DeepGaze II: Reading fixations from deep features trained on object recognition. *arXiv preprint arXiv:1610.01563* (2016).

24. Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. 2016. Ssd: Single shot multibox detector. In *European conference on computer vision*. Springer, 21–37.

25. Jonathan Long, Evan Shelhamer, and Trevor Darrell. 2015. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3431–3440.

26. Laura E Matzen, Michael J Haass, Kristin M Divis, Zhiyuan Wang, and Andrew T Wilson. 2017. Data visualization saliency model: A tool for evaluating abstract data visualizations. *IEEE transactions on visualization and computer graphics* 24, 1 (2017), 563–573.

27. Anelise Newman, Barry McNamara, Camilo Fosco, Yun Bin Zhang, Pat Sukhum, Matthew Tancik, Nam Wook Kim, and Zoya Bylinskii. 2020. TurkEyes: A Web-Based Toolbox for Crowdsourcing Attention Data. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (2020).

28. Peter O'Donovan, Aseem Agarwala, and Aaron Hertzmann. 2014. Learning layouts for single-page graphic designs. *IEEE transactions on visualization and computer graphics* 20, 8 (2014), 1200–1213.

29. Chengyao Shen and Qi Zhao. 2014. Webpage saliency. In *European conference on computer vision*. Springer, 33–46.

30. Jaime Teevan, Edward Cutrell, Danyel Fisher, Steven M Drucker, Gonzalo Ramos, Paul André, and Chang Hu. 2009. Visual snippets: summarizing web pages for search and revisitation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2023–2032.

31. Quanlong Zheng, Jianbo Jiao, Ying Cao, and Rynson WH Lau. 2018. Task-driven webpage saliency. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 287–302.

32. Xinyu Zhou, Cong Yao, He Wen, Yuzhi Wang, Shuchang Zhou, Weiran He, and Jiajun Liang. 2017. East: an efficient and accurate scene text detector. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*. 5551–5560.

33. Spandan Madan* Adria Recasens* Kimberli Zhong Hanspeter Pfister Fredo Durand Aude Oliva Zoya Bylinskii*, Sami Alsheikh*. 2017. Understanding infographics through textual and visual tag prediction. In *arXiv preprint arXiv:1709.09215*. `https://arxiv.org/pdf/1709.09215`