

Salient Object Detection for Images Taken by People With Vision Impairments

Jarek Reynolds*, Chandra Kanth Nagesh*, and Danna Gurari
* denotes equal contribution
University of Colorado Boulder

Abstract

Salient object detection is the task of producing a binary mask for an image that deciphers which pixels belong to the foreground object versus background. We introduce a new salient object detection dataset using images taken by people who are visually impaired who were seeking to better understand their surroundings, which we call VizWiz-SalientObject. Compared to seven existing datasets, VizWiz-SalientObject is the largest (i.e., 32,000 human-annotated images) and contains unique characteristics including a higher prevalence of text in the salient objects (i.e., in 68% of images) and salient objects that occupy a larger ratio of the images (i.e., on average, ~50% coverage). We benchmarked ten modern models on our dataset. One method achieves nearly human performance while the rest struggle, mostly for images with salient objects that are large, have less complex boundaries, and lack text as well as for lower quality images. To facilitate future extensions, we share the dataset at <https://vizwiz.org/tasks-and-datasets/salient-object-detection>.

1. Introduction

Locating the most prominent foreground object in an image is a core computer vision problem, often referred to as salient object detection (as well as salient object segmentation and foreground object detection/segmentation) [6,10,30,37]. Our motivation is to have salient object detection models work well for images taken by people who are blind or who have low vision¹. Such a feature would offer several benefits to this community. Users of mobile phone applications which describe one’s visual surroundings (with help from remote humans or algorithms) [1,2,4,12,14,22,30,37] could use salient object detection to support privacy-preservation and impression management [3]. In particular, users could obfuscate all content except the foreground content of interest to avoid sharing information inadvertently

¹People with low vision have limited visual abilities that cannot be improved by wearing glasses or surgery.

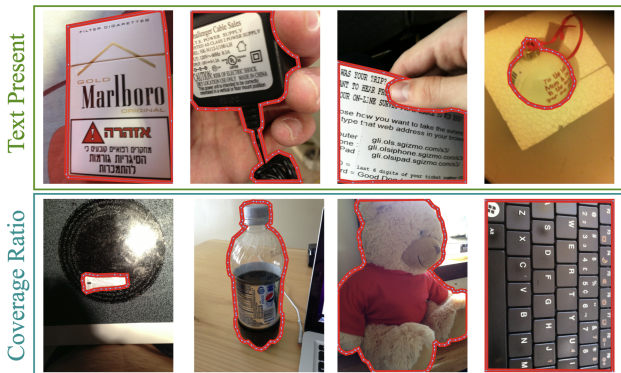


Figure 1. Images exemplifying unique aspects of our new VizWiz-SalientObject dataset when compared to other datasets. The salient objects commonly contain text and occupy a larger portion of the image (i.e., high coverage).

captured in the background of images [19].² Users of such applications also could use it to obfuscate background content so that visual descriptions focus only on the content of interest, the salient object [3]. Individuals also could edit their photos to crop the salient content, including for sharing on social media [38]. Finally, low vision individuals could use it to quickly magnify the content of interest, the foreground object, for deeper inspection [16,36].

While many salient object detection datasets have been created to enable progress in algorithm development [5,6,17,40], they are typically built using high-quality images collected from photo-sharing websites on the Internet. Yet, it has been shown that images taken by visually impaired photographers who are trying to learn about the content they photograph [19,20] often have distinct characteristics from images in mainstream computer vision datasets, including different types of content such as objects showing text [21] and many quality issues since photographers cannot verify the quality of images they take [11]. This begs a question of how well algorithms designed for and trained on existing datasets will generalize to the real-world challenges ob-

²Some organizations record submitted data for the potential that it could be useful evidence in legal contexts.

served in images taken by people with vision impairments.

To fill this gap, we introduce a new salient object detection dataset based on images captured in an authentic use case where visually impaired photographers shared their images to solicit assistance in learning about the visual content. In particular, we crowdsourced the collection of salient object annotations to create a final dataset consisting of 32,000 annotated images. Examples of annotated images are shown in Figure 1. We call our dataset VizWiz-SalientObject (or VizWiz-SO).

We also conduct a detailed analysis to reveal how this new dataset relates to existing datasets. When compared to seven existing salient object detection datasets, we observe VizWiz-SalientObject is the largest (i.e., 32,000 human-annotated images) and is unique in its higher prevalence of salient objects containing text (i.e., in 68% of images), occupying a larger ratio of the images (i.e., on average, $\sim 50\%$), and having less complex boundaries. When comparing our salient objects to the visual evidence needed to answer questions the photographers asked about their images (i.e., taken from the VizWiz-VQA-Grounding dataset [9]), we observe that over half the images have a salient object that matches the necessary visual evidence.

We then benchmark ten modern models on our dataset. Experiments reveal that they struggle for lower quality images as well as for images with salient objects that are large, have less complex boundaries, and lack text. To our knowledge, this work is the first to reveal these limitations of algorithms, and we attribute these findings in part to the distinct visual characteristics of our new dataset. Excitingly, one model nearly achieves human performance.

We expect this new dataset will support developing more generalized models that can also benefit related applications with similar real-world challenges (e.g., lower quality images, large objects, little object boundary complexity), including robotics, lifelogging, and other egocentric-based applications. We also anticipate that algorithmic improvements will benefit downstream applications from salient object detection, such as salient object tracking [15, 23, 28], visual question answering, and image captioning. To facilitate extensions, we publicly-share it at <https://vizwiz.org/tasks-and-datasets/salient-object-detection>.

2. Related Work

Salient Object Detection Datasets. Over the past couple of decades, many datasets were introduced to facilitate improving the design of algorithms that address salient object detection problems. Several survey papers provide comprehensive characterizations of the tens of datasets designed for this task [5, 6, 17, 40]. A common observation is that datasets were artificially constructed around high-quality images which often feature salient objects in the center of the images with a high contrast against the background.

This is a mismatch from many real-world settings, especially for visual media taken by visually impaired photographers who often photograph distinct types of content with the aim to learn about that content. We introduce the first salient object detection dataset based on images taken by visually impaired people in an authentic use case where they were trying to learn about their visual surroundings. We show that, compared to seven modern datasets, our dataset is larger and has a high prevalence of salient objects containing textual information, occupying larger portions of the images, and having little boundary complexity.

Salient Object Detection Models. Novel models for performing salient object detection have been introduced for over 20 years, with the status quo since 2015 being that state-of-the-art methods employ neural networks trained on large-scale annotated datasets. Several survey papers provide comprehensive characterizations of the many models for this task [5, 6, 17, 40]. While convolutional neural network (CNN) based models became the mainstream method [31, 35, 41] in 2015, transformer based models [29, 42] have become the mainstream approach over the past few years. To assess how well modern methods perform on our new dataset, we benchmark ten modern methods. While most methods fall below human performance, struggling for lower quality images as well as for images with salient objects that are large, have little boundary complexity, and lack text, one method nearly achieves human performance.

Visual Assistance Applications. Automated salient object detection methods are desired by people with vision impairments to independently achieve many tasks. Examples include for privacy preservation and impression management [3], by obfuscating background content in images (or video frames) to avoid sharing personal information inadvertently [19]. This could benefit the 100s of thousands of users who submit millions of requests to numerous technologies [1, 2, 4, 12, 14, 22, 30, 37] to receive visual assistance for daily tasks such as deciding what to eat, wear, and buy [8, 19]. Such models would also support more efficient learning about the salient content in images by restricting image descriptions to focus only on that content [3] or rapidly magnifying that content for low vision users' inspection [16, 36]. Finally, such models would enable users to independently edit content in their images [38] and decide what to share to social media [47].

3. VizWiz-SalientObject Dataset

We now introduce our new salient object detection dataset, we call VizWiz-SalientObject (VizWiz-SO).

3.1. Dataset Creation

Image Source. We focus on images taken by visually impaired people who shared them in an authentic use case

where they were soliciting visual assistance. Specifically, we leverage the 39,181 labeled images from the VizWiz-Captions dataset, each of which is paired with five crowd-sourced captions [21]. Observing that images from these photographers can have severe quality issues resulting in no detectable salient object (e.g., extreme blur or inadequate illumination), we did not use the images which were captioned as follows by at least four of the five crowdworkers: “Quality issues are too severe to recognize visual content.” We also did not use the small images (i.e., both the height and width were less than 300 pixels) because of the challenges of collecting precise annotations for such images. This left us with 37,120 images for our annotation task.

Task Design. Our task interface for segmenting salient objects begins with a comprehensive instruction set at the top detailing both how to navigate the interface and how to complete challenging annotation scenarios. Next, it shows an image alongside two preliminary questions for verifying there is a single, unambiguous foreground object. The first question asks “Is the image showing a screenshot?” If the answer is “yes”, we conclude the image lacks a salient object. Next, we ask the more general, direct question of “Is there a single unambiguous foreground object?” An annotator is only prompted to segment the foreground object for images deemed by these preliminary questions to show a single, unambiguous foreground object.

To demarcate the boundary of the salient object, the interface collects a series of points that are connected into polygon(s). When segmenting the salient object, the annotator is required to remove any *holes* (e.g., donut) as well as capture all object parts when *occlusions* break a salient object into more than one polygon (e.g., hand obfuscates a pencil into two parts). The annotator also has an option to select a button indicating that the salient object occupies the full image.

Annotation Collection. We leveraged the benefits of an around-the-clock distributed workforce by crowdsourcing annotations via Amazon’s crowdsourcing marketplace, Amazon Mechanical Turk (AMT).

Although AMT can support our large-scale annotation needs, it brings concerns about annotation quality due to the anonymous nature of the workforce. Consequently, we implemented several measures to ensure the collection of high-quality annotations. First, we only accepted workers who had at least a 98% acceptance rate and had completed at least 500 Human Intelligence Tasks (HITs) on AMT. To encourage understanding of our task instructions, we also only accepted crowdworkers from the United States since that gave us confidence that they have English-language proficiency. We also required crowdworkers to pass a qualification test covering five challenging annotation scenarios, including annotating foreground objects with complex

boundaries, holes, and occlusions.

We employed 40 AMT crowdworkers who completed our qualification test to annotate all images. For each of the 37,120 images, we collected two annotations from the crowdworkers.³ During annotation collection, we monitored ongoing quality by tracking each worker’s performance with respect to their frequency of indicating the presence of full-screen annotations or no prominent foreground object as well as the level of detail they provided in their segmentations (e.g., high prevalence of triangles). Cumulatively, the crowdworkers took 1,290 annotation hours over 11 days to complete annotating the 37,120 images.

Annotation Post-Processing. We next analyzed the redundant annotations per image to determine how to use each annotated image in the final dataset. First, we removed 3,662 images for which workers agreed there was no single, unambiguous salient object, which occurred when both annotators either answered “Yes” to “Is the image a screenshot?” or “No” to “Is there a single most prominent foreground object?” Next, we manually inspected 7,443 images for which workers disagreed on the answers to either of the two preliminary questions and determined whether there is indeed a single, unambiguous object. Finally, with all images deemed to have a single, unambiguous salient object, we determined which annotation to assign as ground truth. To assist in this process, we computed the intersection over union (IoU) score between the two segmentations for all images with two or more segmentations. With $\text{IoUs} \geq 0.90$, we deemed both annotations high quality and randomly selected one as ground truth. For the remaining 2,951 images with $\text{IoUs} < 0.90$, we manually reviewed the annotations to decide whether one was correct or whether the image should be discarded due to foreground object ambiguity.

3.2. Dataset Analysis

We now characterize the VizWiz-SalientObject dataset (VizWiz-SO) and how it relates to existing datasets.

3.2.1 VizWiz-SalientObject vs Existing SOD Datasets

We first characterize our new dataset and how it is both distinct and similar to the following seven SOD datasets:

- *DUTS* [39]: commonly used to train state-of-the-art algorithms (e.g., [29, 31, 34, 35, 41, 42]) due to its large size paired with diverse saliency challenges.
- *DUT-OMRON* [44]: introduced in 2013, this older dataset was popular for a while.
- *ECSSD* [43]: designed to show complex scenes with textures and structures expected to be common in real-world salient object detection scenarios.

³For a subset of images, we collected four annotations to support further analysis of human performance (see the Supplementary Materials).

	DAVIS-S [46]	PASCAL-S [27]	HR [46]	ECSSD [43]	DUT-O [44]	UH [42]	DUTS [39]	Ours
Images	92	850	2,010	1,000	5,168	5,920	15,572	32,000
Text	13%	24%	15%	15%	11%	19%	13%	68%
MR	22%	31%	25%	9%	17%	35%	19%	1%
Holes	82%	50%	62%	29%	28%	75%	41%	4%
MW	1080	375	2704	300	300	3612	300	1296
MH	1920	500	3264	400	400	5000	400	968

Table 1. Characterization of our VizWiz-SO dataset and seven existing salient object detection datasets with respect to how many images are included (“Images”), percentage of images containing text in the salient objects (“Text”), percentage of images that have salient objects consisting of more than one region (“MR”), percentage of images that have holes in the salient objects (“Holes”), median image width (“MW”) and median image height (“MH”). As shown, our dataset is distinct in that it contains more images, more salient objects with text present, more salient objects consisting of one region, and fewer salient objects containing holes. (HR=HRSOD; UH=UHRSD)

- *PASCAL-S* [27]: derived from PASCAL VOC’s [13] validation set to facilitate salient object segmentation on a popular set of images in the computer vision community.
- *HRSOD* [46]: explicitly designed for salient object detection on high-resolution images. This is relevant for our use case since images taken by people with vision impairments often can be high resolution.
- *UHRSD* [42]: currently the largest ultra-high resolution salient object detection dataset, further supporting our use case since images can be ultra-high resolution.
- *DAVIS-S* [46]: derived from DAVIS [32], a densely annotated video segmentation dataset. This is relevant for our use case to analyze implications for video frames since visually impaired photographers often stream live video when using visual assistance technologies [2, 14].

Of note, images in six of these datasets originate from the Internet on photo-sharing websites such as Flickr [27, 39, 42–44, 46], and so likely are high quality since they were deemed of sufficient quality to upload to the Internet.⁴

For each salient object in every dataset, we characterize it in eight ways. Three measures detect the presence of particular properties for the salient object. These are whether the salient object *contains text*⁵, *consists of multiple regions*⁶, or *contains any hole(s)*. The next three measures characterize the salient region itself. First, we identify the position of an object within an image by measuring its *center of mass* relative to the image coordinates, resulting in x and y coordinate values in the range between 0 to 1. Next, we characterize the object’s *boundary complexity* by computing its isoperimetric inequality, which is the ratio of the object’s area to the length of its perimeter. Values range

⁴The origins of the images for [46] is not reported.

⁵We obfuscate all image content but the salient object and then check whether Microsoft Azure’s OCR API returns text.

⁶Multiple regions means there are multiple separate polygons. This can occur either because multiple salient objects were annotated or because of occlusions that lead to more than one region for a single salient object.

from 0 to 1, with larger values indicating simpler boundaries that are less jagged/dented (e.g., a circle). Finally, to gauge the relative size of a salient object in the image, we compute its *coverage ratio*, meaning the fraction of all image pixels that are occupied by the salient object’s pixels. The final two measures offer finer-grained details about typical image resolutions. We report the *median image width* and *median image height*.

We show summative statistics of our findings per dataset in Table 1 and Figure 2. In Table 1, we report how many images are in each dataset paired with the percentage having salient objects with text, multiple regions, and holes as well as summative statistics about image resolutions. *To our knowledge, this is the first paper that systematically characterizes SOD datasets based on whether text is present, holes are present, and multiple regions are present.* In Figure 2, we visualize statistics summarizing the values for each dataset’s salient objects with respect to center of mass, boundary complexity, and coverage ratio using boxplots.

A unique aspect of our VizWiz-SO dataset is that it features more salient objects with textual data. Specifically, 68% of salient objects in VizWiz-SO contain text while the dataset with the next highest prevalence of text, PASCAL-S [27], only has it for 24% of the images (Table 1). A gap of this magnitude (i.e., 44 percentage points) suggests that our new dataset offers a considerable domain shift in the salient object detection problem space. We suspect part of this shift stems from the types of salient objects included, with many more daily objects such as products (e.g., food packages) included in our VizWiz-SO dataset.

Another unique aspect of VizWiz-SO is that far fewer images feature salient objects that consist of multiple regions; i.e., only 1% of images (Table 1). We suspect this distinction stems from our unique approach of adopting a rigorous annotation preprocessing step, where we require crowdworkers to verify images have one unambiguous salient object before allowing them to annotate images for use in our final dataset. Any remaining objects in our




Figure 2. Summary statistics for ours and seven other datasets with respect to four measures. Each box reveals statistics about all salient objects in a particular dataset with the central mark capturing the median value, box edges the 25th and 75th percentiles values, whiskers the most extreme data points not considered outliers, and individually plotted points the outliers. Our dataset is unique in that salient objects tend to have less complex boundaries, occupy larger portions of an image, and exhibit a greater diversity of sizes relative to the image.

dataset with multiple regions are therefore highly likely a result of occlusions breaking a single salient object into multiple polygons, which evidently is incredibly rare.

VizWiz-SO is also unique due to the rarity in which salient objects contain holes; i.e., only observed for 4% of images (Table 1). From visual inspection, we suspect this finding reflects a domain shift in the types of content found in the datasets. For example, objects in other datasets with holes include people riding bikes, people dancing, and animals in intricate poses. In VizWiz-SO, objects with holes include retail packaging made to hang from hooks, pairs of scissors, and coffee mugs.

A further distinction of our VizWiz-SO dataset is that the salient objects tend to have less complex boundaries (Figure 2). We suspect this is again because of a domain shift in the types of objects in our dataset, with many more human-made items, such as food packaging boxes and cans, that by design have simpler shapes.

A final distinction of salient objects in our VizWiz-SO is how much of the image they occupy (Figure 2). First, they tend to occupy a much larger amount of the image than observed in other datasets. Specifically, they on average occupy roughly half of all image pixels, with a mean coverage ratio of 0.5 and a median of 0.46. In contrast, the dataset with the next highest coverage ratio statistics is PASCAL-S [27], and over 75% of its images contain salient objects that occupy less than half of the image pixels. We attribute this distinction to the authentic use case of our dataset, where visually impaired photographers attempting to learn about the salient objects they are photographing seem to be taking zoomed-in or close-to-camera images of the content of interest. Another unique aspect of our salient objects, is that they exhibit a larger range of sizes, as shown by the gaps between the 25 and 75 percentile values of each box. For example, PASCAL-S features the next largest interquartile range with a 23% gap (i.e., 19% to 42%). In contrast,

the gap for VizWiz-SO is more than twice as large at 56% (i.e., 22% to 78%). Consequently, a unique challenge of our dataset for algorithms is that they no longer can assume a strong bias regarding a salient object’s relative size.

While our findings highlight that our VizWiz-SO dataset has many distinct characteristics, one commonality it has with most existing salient object detection datasets that we found surprising is that the salient objects typically occupy centered positions within an image. We observe this trend for all datasets except HRSOD in Figure 2. We found this surprising since visually impaired photographers cannot visually inspect their images to verify image quality. This finding suggests that visually impaired photographers have skills in conforming to the common photographer’s bias of centering contents of interest they are trying to photograph.

3.2.2 Salient Objects vs Answer Groundings for Visual Question Answering (VQA)

We next explore how the target content the photographers were asking about relates to an image’s salient object. To do so, we compare the annotations of the visual evidence needed to answer questions about the images, i.e., *answer groundings* provided in the VizWiz-VQA-Grounding dataset [9], to annotations of the salient objects. We first identified all images that were in common across the two datasets, yielding 6,540 images. For each image, we then measured the similarity between the answer grounding and salient object segmentations using the IoU metric. We visualize our results using a histogram where we categorize each image into one of ten interval bins starting with $\text{IoU}=[0.0, 0.1)$, incrementing in intervals of 0.1, and ending with $\text{IoU}=[0.9, 1.0)$. Results are shown in Figure 3.

Roughly half of the images have a high similarity between the salient object and answer grounding; e.g., 46% have an $\text{IoU} \geq 0.9$. This reveals that visually impaired photographers often tried to learn about the salient object

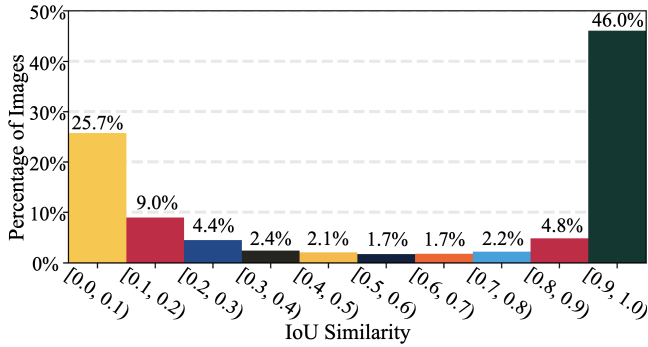


Figure 3. Frequency of observing different levels of similarity between the salient object and the visual evidence needed to answer the photographer’s question. As shown, the photographers often wanted to learn about the salient objects in their images.

when trying to get answers to their visual questions. This reinforces prior works’ findings that users of visual assistance technologies most commonly trying to learn about a (salient) object [7, 8, 18, 24, 45].

We also observe that roughly one quarter of the images have a very low similarity between the salient object and answer grounding; i.e., 25.7% of images had an IoU < 0.1. From manual review of 1,680 images with IoUs less than 0.1, we discovered that 95% (i.e., 1,599) have a salient object featuring a full-screen or large region while the answer grounding captures a small aspect of the salient object. Examples include expiration dates on food packages or the current page number of an open book. The remaining 5% (i.e., 81) of these images featured a answer grounding unrelated to the salient object.

More generally, we observe that the IoU scores follow a U-shaped distribution with only a small portion of images having middling scores; e.g., 7.9% (i.e., 511) have an IoU ≥ 0.3 and < 0.7 . Among these, the salient object always contained the answer grounding region. Two primary trends led to these less common IoU scores. The first is that larger answer grounding regions occur with smaller salient objects. Examples include brands of cereal, types of soda, and denominations of currency. The second trend was for salient objects featuring holes. The VizWiz-VQA-Grounding dataset did not account for holes in their annotation task, and this often led to lower scores when the salient object did contain holes.

Altogether, these findings highlight that a valuable step for tackling many of this population’s VQA goals is to initially locate the salient object. That is because the answer will likely only be grounded in the salient object *or* the background rather than their intersection.

4. Algorithm Benchmarking

We finally analyze modern salient object detection algorithms to show how they perform on our new dataset. We conducted all experiments on a Nvidia A100 GPU.

4.1. Experimental Design

Dataset Splits. We adhere to the standard splits adopted for all VizWiz-based datasets (e.g., [21]), which translates to approximately a 60/20/20 training, validation and test split for our VizWiz-SO dataset; i.e., 19,116, 6,105, and 6,779 images respectively.

Evaluation Metrics. We evaluate each model with respect to five popular metrics for salient object detection models: Mean Absolute Error (MAE), Structure Measure (S_m), Mean F-Measure (F_m), Enhanced Alignment Measure (E_m), and Intersection over Union (IoU).

Algorithms. We benchmark the following ten methods from the past four years to assess the difficulty of our new dataset for modern salient object detection models:

- *Boundary Aware Segmentation Network (BASNet)* [34], 2019: can run at 70fps, making it suitable for real-world applications like those motivated by our dataset.
- *Fusion, Feedback and Focus Network (F3Net)* [41], 2020: was state-of-the-art on five datasets.
- *U2 Network (U2Net)* [35], 2020: is compact (4.7MB), making it suitable for resource-constrained devices such as smartphones (i.e., the image source for our use case).
- *Visual Saliency Transformer (VST)* [29], 2021: was state-of-the-art, with a purely transformer architecture.
- *Pyramidal Feature Shrinking Network (PFSNet)* [31], 2021: was state-of-the-art on five datasets.
- *Pyramid Grafting Network (PGNet)* [42], 2022: state-of-the-art on five datasets using a one-stage framework based on a transformer and CNN backbone.
- *Dichotomous Image Segmentation (DIS)* [33], 2022: state-of-the-art for related task of segmenting foreground objects (even if not salient) in high resolution images.
- *Salient Object Detection via Integrity Learning (ICON)* [48], 2022: state-of-the-art at time of publication.
- *Extreme Attention Guided Salient Object Tracing Network (TRACER)* [26], 2022: the current state-of-the-art model on DUTS challenge.
- *Revisiting Image Pyramid Structure for High Resolution Salient Object Detection (InSPyReNet)* [25], 2022: the current 2nd ranked model on DUTS challenge.

We further characterize each model in Table 2 by identifying datasets used for model training, image size used for

		HP	VST [29]	PGNet [42]	DIS [33]	ICON [48]	TRACER [26]	IPR [25]	IPR-FT	IPR-S	IPR-DS
Attr.	Training set	-	D	DH	DIS5K	D	D	DHU	VW	VW	DV
	Input size	-	224 ²	224 ² , 1024 ²	1024 ²	352 ²	640 ²	384 ²	384 ²	384 ²	384 ²
	Size (MB)	-	171	280	169	251	630	351	351	351	351
VizWiz-SO	$MAE \downarrow$	0.02	0.17	0.21	0.36	0.20	0.21	0.03	0.03	0.04	0.04
	$S_m \uparrow$	0.92	0.65	0.62	0.46	0.69	0.67	0.90	0.91	0.89	0.90
	$F_m \uparrow$	0.96	0.83	0.79	0.61	0.86	0.83	0.95	0.95	0.94	0.94
	$E_m \uparrow$	0.97	0.76	0.74	0.55	0.73	0.65	0.94	0.94	0.93	0.94
	$IoU \uparrow$	0.94	0.73	0.67	0.49	0.70	0.68	0.93	0.93	0.90	0.91

Table 2. Analysis of existing models that we benchmark on our VizWiz-SO dataset, including both off-the-shelf models as well as the top-performing algorithm fine-tuned (-FT), trained from scratch using VizWiz-SO (-S), and trained from scratch using both DUTS and VizWiz-SO (-DS). We first report differentiating attributes between the algorithms and then present the model performance with respect to five widely-used metrics. (HP=Human Performance; IPR=InSPyReNet; DH=D+HR; DHU=D+HR+UH; VW=VizWiz-SO; DV=DUTS+VizWiz-SO; D=DUTS-TR [39]; VW=VizWiz-SO; HR=HRSOD; UH=UHRSD [46])

model training, and model size. All models predict segmentation maps showing pixel brightness, which we then convert into binary masks.

Humans. We evaluate human performance to establish an upper bound for automated methods by comparing the two annotations in cases where their IoU is greater than 0.90.⁷

4.2. Performance for Off-The-Shelf Models

We first evaluate all models in their original design. Key results are shown in Table 2, and the rest are in the Supplementary Materials due to space constraints.

We observe that the top-performing model, InSPyReNet [25], nearly reaches human performance. For example, the gap in MAE performance is 0.01. Towards understanding why other models may perform far worse than human performance, we later explore whether it may be due to the training data or the architectures themselves.

We found that all models except InSPyReNet [25] perform worse on VizWiz-SO than on datasets which they were originally evaluated. For example, the MAE and S_m performance observed by PGNet [42] on DUTS-TE is 0.028 and 0.912 respectively versus 0.2123 and 0.6233 respectively for our dataset. These findings underscore that our dataset can offer distinct challenges from existing SOD datasets. Analysis in subsequent sections reveal what makes our new dataset challenging for many modern models.

We also observe that the performance of methods designed for real-world, practical challenges, specifically to run fast (i.e., BASNet [34]) and be compact (i.e., U2Net [35]), are among the poorer-performing methods. A tangentially important direction for future research will therefore be to ensure that models not only produce accurate saliency maps but also meet these practical needs.

⁷We chose this threshold to match what was used for establishing ground truth during dataset creation.

4.3. Performance When Training on VizWiz-SO

We now explore whether training on our new dataset boosts the performance for the top-performing algorithm, InSPyReNet [25]. We analyze three variants: (1) pretrained InSPyReNet model fine-tuned on VizWiz-SO (IPR-FT), (2) InSPyReNet algorithm trained from scratch on VizWiz-SO (IPR-S), and (3) InSPyReNet algorithm trained from scratch on DUTS [39] and VizWiz-SO (IPR-DS). Overall, we observe slightly worse or comparable performance from these variants, suggesting that the in-domain training data from VizWiz-SO is not necessary. Only leveraging the greater diversity of training data VizWiz-SO provides, by fine-tuning on Viz-Wiz-SO, yields a very slight advantage.

4.4. Fine-grained Analysis

We next analyse what makes our dataset challenging for modern algorithms. To do so, we divide the test set according to the following *five* factors, with the first four based on metadata from Section 3.2 to characterize our dataset:

- *Text Presence*: two groups based on whether text is present in the salient object.
- *Coverage Ratio (Coverage)*: three groups based on the 33rd and 66th quartile values in our dataset. All images with coverage ratio less than 0.32 has *small* coverage, between 0.32 and 0.62 has *medium* coverage, and greater than 0.62 has *large* coverage.
- *Boundary Complexity (Boundary)*: two groups by splitting them around the mean score for boundary complexity (i.e., 0.66) with *high* complexity when the score is less than the mean and *low* complexity otherwise.
- *Image Resolution (Resolution)*: two groups by splitting images around whether they are high resolution, as defined by whether the image width and height are both at least 1080 and 1920 respectively.

		VST	PGNet	DIS	ICON	TRACER	IPR	IPR-FT	IPR-S	IPR-DS
		[29]	[42]	[33]	[48]	[26]	[25]			
Text Present	True	0.13	0.16	0.32	0.15	0.18	0.02	0.02	0.03	0.03
	False	0.24	0.29	0.40	0.24	0.25	0.09	0.09	0.10	0.11
Coverage	Small	0.11	0.12	0.10	0.12	0.13	0.07	0.07	0.08	0.09
	Medium	0.09	0.15	0.25	0.10	0.13	0.05	0.05	0.06	0.07
	Large	0.30	0.35	0.70	0.35	0.39	0.08	0.09	0.10	0.10
Boundary	High	0.12	0.16	0.21	0.17	0.14	0.05	0.05	0.06	0.07
	Low	0.21	0.25	0.48	0.26	0.24	0.08	0.07	0.08	0.08
Resolution	High	0.16	0.22	0.37	0.21	0.19	0.09	0.09	0.11	0.11
	Low	0.17	0.21	0.35	0.20	0.21	0.08	0.08	0.10	0.10
Quality	Good	0.14	0.17	0.30	0.16	0.19	0.04	0.03	0.05	0.06
	Poor	0.27	0.34	0.50	0.33	0.32	0.10	0.10	0.11	0.11

Table 3. Fine-grained analysis of existing models with respect to presence of text on the salient object (“Text Present”), relative size of the salient object in the image (“Coverage”), relative complexity of the salient object’s boundary (“Boundary”), and image quality (“Quality”) using the $MAE \downarrow$ metric. As shown, the models perform worse when salient objects lack text, occupy a large portion of the image, and have less complex boundaries as well as when the image quality is poor. (IPR=InSPyReNet)

- *Quality*: leveraging metadata from prior work [21], which indicates how many of the five crowdworkers indicated an image as insufficient quality to recognize the content, we split the images into groups with *good* quality being when none of the crowdworkers indicate insufficient quality and *poor* otherwise.

To our knowledge, this paper is the first to systematically characterize SOD algorithm performance with respect to these criteria. We again split results with key ones shown in Table 3 and the rest shown in the Supplementary Materials due to space constraints.

In terms of text presence, we see that the models perform better when there is text present as opposed to when there is none. For example, the performance drops by 0.11 for the second best model, VST, and 0.07 for the top-performing model InSPyReNet. We suspect visual patterns that arise with text may serve as a valuable cue to models in locating salient objects.

Next, we see that as the coverage ratio of the salient objects increase, the models tend to perform worse with all models performing the worst on the large objects. For instance, the second best model, VST, has a performance dropoff of 0.19 when predicting images with small coverage ratios as opposed to large coverage ratios. In contrast, the top-performing model InSPyReNet only has a drop of 0.01, which we suspect is a large reason behind this model’s superior performance. We suspect this performance gap arises in part from existing datasets largely lacking such large salient objects, which both could have affected what algorithms were designed to handle as well what they could learn from the data they observed.

Further observed trends are that performance drops for salient objects with lower boundary complexity and poorer

quality images. We suspect this is due to domain shifts between our dataset and prior datasets that affect what algorithms were designed for and what they could learn from the training data.

5. Conclusions

This work’s contributions are: (1) a new SOD dataset originating from a practical use case (Section 3.1), (2) comparison of this dataset to seven popular SOD datasets to reveal how it fills important gaps of existing SOD datasets (Section 3.2.1), (3) comparison of this dataset to the VQA grounding dataset to reveal how it relates (Section 3.2.2), and (4) benchmarking of seven modern algorithms on this dataset to reveal limitations of modern SOD algorithms (Section 4). To our knowledge, this work is the first to (1) systematically analyze SOD datasets based on whether text is present, holes are present, multiple regions are present, and their relationship to VQA tasks (Section 3.2) and (2) systematically analyze SOD algorithms with respect to whether text is present, image quality, as well as levels of an object’s boundary complexity and image coverage ratio (Section 4). Results reveal domain shifts between our SOD dataset and existing SOD datasets as well as limitations of many SOD algorithms: they struggle for salient objects that are large, with simpler boundaries, and lack text as well as for lower quality images. Results also show that one algorithm achieves nearly human performance.

Acknowledgments. This project was supported in part by a National Science Foundation SaTC award (#2148080) and Amazon Mechanical Turk. We thank Leah Findlater and Yang Wang for contributing to this research idea.

References

- [1] Aipoly. Aipoly homepage. <https://www.aipoly.com/>, 2020. (Accessed on 01/08/2020). 1, 2
- [2] Aira. Aira homepage. <https://aira.io/>, 2020. (Accessed on 01/08/2020). 1, 2, 4
- [3] Rahaf Alharbi, Robin N Brewer, and Sarita Schoenebeck. Understanding emerging obfuscation technologies in visual description services for blind and low vision people. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2):1–33, 2022. 1, 2
- [4] BeSpecular. Bespecular. <https://www.bespecular.com/>, 2020. (Accessed on 01/08/2020). 1, 2
- [5] Ali Borji, Ming-Ming Cheng, Qibin Hou, Huaizu Jiang, and Jia Li. Salient object detection: A survey. *Computational visual media*, pages 1–34, 2019. 1, 2
- [6] Ali Borji, Ming-Ming Cheng, Huaizu Jiang, and Jia Li. Salient object detection: A benchmark. *IEEE transactions on image processing*, 24(12):5706–5722, 2015. 1, 2
- [7] Erin Brady, Jeffrey P Bigham, et al. Crowdsourcing accessibility: Human-powered access technologies. *Foundations and Trends® in Human-Computer Interaction*, 8(4):273–372, 2015. 6
- [8] Erin L Brady, Yu Zhong, Meredith Ringel Morris, and Jeffrey P Bigham. Investigating the appropriateness of social network question asking as a resource for blind users. In *Proceedings of the 2013 conference on Computer supported cooperative work*, pages 1225–1236. ACM, 2013. 2, 6
- [9] Chongyan Chen, Samreen Anjum, and Danna Gurari. Grounding answers for visual questions asked by visually impaired people. *arXiv preprint arXiv:2202.01993*, 2022. 2, 5
- [10] Ming-Ming Cheng, Niloy J Mitra, Xiaolei Huang, Philip HS Torr, and Shi-Min Hu. Global contrast based salient region detection. *IEEE transactions on pattern analysis and machine intelligence*, 37(3):569–582, 2014. 1
- [11] Tai-Yin Chiu, Yinan Zhao, and Danna Gurari. Assessing image quality issues for real-world problems. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3646–3656, 2020. 1
- [12] Ned Desmond. Microsoft’s Seeing AI founder Saqib Shaikh is speaking at Sight Tech Global. 1, 2
- [13] Mark Everingham, Luc Gool, Christopher K. Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *Int. J. Comput. Vision*, 88(2):303–338, jun 2010. 4
- [14] Be My Eyes. Bringing sight to blind and low-vision people. <https://www.bemyeyes.com/>, 2020. (Accessed on 01/08/2020). 1, 2, 4
- [15] Deng-Ping Fan, Wenguan Wang, Ming-Ming Cheng, and Jianbing Shen. Shifting more attention to video salient object detection. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8546–8556, 2019. 2
- [16] American Federation for the Blind. Low vision optical devices. 1, 2
- [17] Ashish Kumar Gupta, Ayan Seal, Mukesh Prasad, and Priytee Khanna. Salient object detection techniques in computer vision—a survey. *Entropy*, 22(10), 2020. 1, 2
- [18] Danna Gurari, Kun He, Bo Xiong, Jianming Zhang, Mehrnoosh Sameki, Suyog Dutt Jain, Stan Sclaroff, Margrit Betke, and Kristen Grauman. Predicting foreground object ambiguity and efficiently crowdsourcing the segmentation (s). *International Journal of Computer Vision*, 126(7):714–730, 2018. 6
- [19] Danna Gurari, Qing Li, Chi Lin, Yinan Zhao, Anhong Guo, Abigale Stangl, and Jeffrey P Bigham. Vizwiz-priv: A dataset for recognizing the presence and purpose of private visual information in images taken by blind people. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 939–948, 2019. 1, 2
- [20] Danna Gurari, Qing Li, Abigale J Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P Bigham. Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3608–3617, 2018. 1
- [21] Danna Gurari, Yinan Zhao, Meng Zhang, and Nilavra Bhattacharya. Captioning images taken by people who are blind. In *European Conference on Computer Vision*, pages 417–434. Springer, 2020. 1, 3, 6, 8
- [22] iDentifi. identifi. <http://getidentifi.com/>, 2020. (Accessed on 01/08/2020). 1, 2
- [23] Ge-Peng Ji, Keren Fu, Zhe Wu, Deng-Ping Fan, Jianbing Shen, and Ling Shao. Full-duplex strategy for video object segmentation. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4902–4913, 2021. 2
- [24] Hernisa Kacorri, Kris M Kitani, Jeffrey P Bigham, and Chieko Asakawa. People with visual impairment training personal object recognizers: Feasibility and challenges. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 5839–5849, 2017. 6
- [25] Taehun Kim, Kunhee Kim, Joonyeong Lee, Dongmin Cha, Jiho Lee, and Daijin Kim. Revisiting image pyramid structure for high resolution salient object detection. In *Proceedings of the Asian Conference on Computer Vision*, pages 108–124, 2022. 6, 7, 8
- [26] Min Seok Lee, Wooseok Shin, and Sung Won Han. Tracer: Extreme attention guided salient object tracing network, 2022. 6, 7, 8
- [27] Yin Li, Xiaodi Hou, Christof Koch, James M. Rehg, and Alan L. Yuille. The secrets of salient object segmentation. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 280–287, 2014. 4, 5
- [28] Nian Liu, Kepan Nan, Wangbo Zhao, Xiwen Yao, and Junwei Han. Learning complementary spatial-temporal transformer for video salient object detection. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–11, 2023. 2
- [29] Nian Liu, Ni Zhang, Kaiyuan Wan, Ling Shao, and Junwei Han. Visual saliency transformer, 2021. 2, 3, 6, 7, 8
- [30] LookTel. Looktel recognizer. 1, 2
- [31] Mingcan Ma, Changqun Xia, and Jia Li. Pyramidal feature shrinking for salient object detection. *Proceedings of*

- the AAAI Conference on Artificial Intelligence*, 35(3):2311–2318, May 2021. 2, 3, 6
- [32] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 4
- [33] Xuebin Qin, Hang Dai, Xiaobin Hu, Deng-Ping Fan, Ling Shao, and Luc Van Gool. Highly accurate dichotomous image segmentation, 2022. 6, 7, 8
- [34] Xuebin Qin, Deng-Ping Fan, Chenyang Huang, Cyril Diagne, Zichen Zhang, Adrià Cabeza Sant’Anna, Albert Suàrez, Martin Jagersand, and Ling Shao. Boundary-aware segmentation network for mobile and web applications, 2021. 3, 6, 7
- [35] Xuebin Qin, Zichen Zhang, Chenyang Huang, Masood Dehghan, Osmar R. Zaiane, and Martin Jagersand. U2-net: Going deeper with nested u-structure for salient object detection. *Pattern Recognition*, 106:107404, 2020. 2, 3, 6, 7
- [36] Abigale J Stangl, Esha Kothari, Suyog D Jain, Tom Yeh, Kristen Grauman, and Danna Gurari. Browsewithme: An online clothes shopping assistant for people with visual impairments. In *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility*, pages 107–118, 2018. 1, 2
- [37] TapTapSee. Taptapsee. 1, 2
- [38] Violeta Voykanska, Shiri Azenkot, Shaomei Wu, and Gilly Leshed. How blind people interact with visual content on social networking services. In *Proceedings of the 19th acm conference on computer-supported cooperative work & social computing*, pages 1584–1595, 2016. 1, 2
- [39] Wang, Lijun, Lu, Huchuan, Wang, Yifan, Feng, Mengyang, Wang, Dong, Yin, Baocai, Ruan, and Xiang. Learning to detect salient objects with image-level supervision. In *CVPR*, 2017. 3, 4, 7
- [40] Wenguan Wang, Shuyang Zhao, Jianbing Shen, Steven C. H. Hoi, and Ali Borji. Salient object detection with pyramid attention and salient edges. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 1, 2
- [41] Jun Wei, Shuhui Wang, and Qingming Huang. F3net: Fusion, feedback and focus for salient object detection, 2019. 2, 3, 6
- [42] Chenxi Xie, Changqun Xia, Mingcan Ma, Zhirui Zhao, Xiaowu Chen, and Jia Li. Pyramid grafting network for one-stage high resolution saliency detection, 2022. 2, 3, 4, 6, 7, 8
- [43] Qiong Yan, Li Xu, Jianping Shi, and Jiaya Jia. Hierarchical saliency detection. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1155–1162, 2013. 3, 4
- [44] Yang, Chuan, Zhang, Lihe, Lu, Huchuan, Ruan, Xiang, Yang, and Ming-Hsuan. Saliency detection via graph-based manifold ranking. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 3166–3173. IEEE, 2013. 3, 4
- [45] Xiaoyu Zeng, Yanan Wang, Tai-Yin Chiu, Nilavra Bhattacharya, and Danna Gurari. Vision skills needed to answer visual questions. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2):1–31, 2020. 6
- [46] Yi Zeng, Pingping Zhang, Jianming Zhang, Zhe Lin, and Huchuan Lu. Towards high-resolution salient object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7234–7243, 2019. 4, 7
- [47] Yuhang Zhao, Shaomei Wu, Lindsay Reynolds, and Shiri Azenkot. The effect of computer-generated descriptions on photo-sharing experiences of people with visual impairments. *Proceedings of the ACM on Human-Computer Interaction*, 1(CSCW):1–22, 2017. 2
- [48] Mingchen Zhuge, Deng-Ping Fan, Nian Liu, Dingwen Zhang, Dong Xu, and Ling Shao. Salient object detection via integrity learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 6, 7, 8