How do we look at images through polygonal apertures: The PolyCAT Dataset

Brahan Aklilu
brahan@post.bgu.ac.il
Ben-Gurion University of the Negev
Beer-Sheva, Israel

Ohad Ben-Shahar
ben-shahar@cs.bgu.ac.il
Ben-Gurion University of the Negev
Beer-Sheva, Israel

Abstract

We introduce PolyCAT, a public eye-tracking dataset for studying visual attention on natural images viewed through irregular polygonal apertures. Thirty participants viewed 600 images from the CAT2000 dataset through 27 polygon shapes of varying geometric properties while their eye movements were recorded with an EyeLink 1000+ on a 27″ 4K display at 70 cm viewing distance. PolyCAT provides sample-level gaze streams, fixation events, empirical saliency maps, scanpath sequences, and quality metrics. We report baseline saliency benchmarks and demonstrate that polygonal aperture geometry systematically affects fixation distributions, offering a controlled testbed for studying how viewing constraints modulate visual exploration.

Keywords

eye tracking, saliency, visual attention, polygonal aperture, gaze prediction, dataset

ACM Reference Format:
Brahan Aklilu and Ohad Ben-Shahar. 2026. How do we look at images through polygonal apertures: The PolyCAT Dataset. In *2026 Symposium on Eye Tracking Research and Applications (ETRA '26), June 01–04, 2026, Marrakesh, Morocco.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3797246.3806795

1 Introduction

Predicting where people look in natural scenes is a central problem in visual attention research. Over the past two decades, saliency models have progressed from biologically inspired feature maps [5] to deep networks that approach the inter-observer agreement ceiling on standard benchmarks [8]. This progress has been driven largely by publicly available eye-tracking datasets that provide ground-truth fixation data for model training and evaluation [2, 7].

Almost all existing saliency datasets present images in their entirety, typically on a uniform background and without explicit viewing constraints. In everyday vision, however, observers rarely see whole scenes at once. Objects and regions are partially revealed through partially opened windows, doorways, instrument panels, head-mounted displays, and other irregular openings that constrain what can be seen. How such aperture constraints interact with scene content to shape gaze allocation remains poorly understood, in part because there is no large-scale eye-tracking dataset that systematically manipulates the shape of the visible region.

We address this gap with **PolyCAT** (**Poly**gonal-aperture **C**AT2000 **A**ttention **T**racking), a dataset in which participants view images from the CAT2000 collection [2] through polygon-shaped apertures whose geometric properties—such as symmetry, convexity, and boundary irregularity—are controlled experimentally. The polygon set spans a diverse range of shapes, from simple symmetric references to complex irregular forms, enabling systematic investigation of how aperture geometry and image content jointly influence fixation distributions. By design, PolyCAT remains compatible with existing saliency infrastructure: it reuses CAT2000 images, follows common data formats, and provides standard evaluation splits.

This paper focuses on the *dataset* itself: its design, acquisition protocol, processing pipeline, and released artifacts. We also include a small set of baseline analyses to illustrate how PolyCAT can be used, but a comprehensive benchmark study is beyond the scope of this work and is left for future research.

The main contributions are:

(1) A publicly available eye-tracking dataset of polygonal-aperture viewing, comprising tens of thousands of image-viewing trials across hundreds of unique polygon–image pairings, recorded with an EyeLink 1000+ on a 27″ 4K display at 70 cm viewing distance.
(2) A principled polygon stimulus set spanning reference shapes, convexity-varied shapes, and computationally generated irregular forms, together with machine-readable geometry descriptors that enable controlled studies of aperture-geometry effects on gaze.
(3) Processed data products—including sample-level gaze streams, fixation tables, empirical saliency maps, scanpath sequences, and per-trial quality metrics—accompanied by detailed documentation, Croissant metadata, and example code for preprocessing and evaluation.
(4) Initial baseline saliency analyses, including a center-prior model and a state-of-the-art deep model, demonstrating that polygonal-aperture viewing poses a distinct challenge and highlighting opportunities for future aperture-aware attention models.

2 Related Work

Eye-tracking datasets for saliency. The MIT saliency benchmark [7] established a widely used evaluation framework with 300 natural images viewed by 39 observers, providing fixation maps and standardized metrics. CAT2000 [2] substantially expanded stimulus coverage with 4,000 images across 20 categories, viewed by 24 observers with 3 s free viewing per image. Subsequent datasets



This work is licensed under a Creative Commons Attribution 4.0 International License.
ETRA '26, Marrakesh, Morocco
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2519-7/2026/06
https://doi.org/10.1145/3797246.3806795

targeted specific aspects of gaze behaviour: SALICON [6] uses mouse-contingent viewing as a scalable proxy for eye movements, OSIE [14] augments images with object- and attribute-level annotations, and FIGRIM [3] incorporates memory judgments to relate gaze to subsequent recognition performance. Across these and other datasets, stimuli are presented as full, unoccluded images.

Several recent resources further broaden the scope of saliency evaluation—for example by introducing videos, egocentric recordings, or text-and-image combinations—but they likewise assume rectangular, globally visible stimuli. PolyCAT is complementary to this line of work: it retains the familiar natural-image setting while manipulating the shape of the visible region through polygonal apertures, creating conditions where only parts of the scene are accessible at any given fixation.

Viewing constraints and attention. The gaze-contingent display literature has long used artificial scotomas and moving windows to study the functional roles of central and peripheral vision [13]. Tunnel-vision paradigms restrict viewing to a window around the current fixation, whereas scene-completion paradigms reveal regions over time, probing how observers sample information under spatial constraints. Other work has examined attention in the presence of occluders, apertures, or limited fields of view in applied settings such as driving, interface design, and head-mounted displays. However, these studies typically employ circular or rectangular windows and are not organized as large-scale, publicly released datasets. Recent work has further shown that the stimulus center bias persists even when stimuli are spatially decoupled from the display center [11], raising the question of how this bias manifests when the visible region itself is irregularly shaped.

To our knowledge, no existing eye-tracking dataset systematically varies the geometric properties of *irregular* polygonal apertures while observers view natural images. PolyCAT fills this gap by combining controlled aperture geometry with standard saliency stimuli, enabling quantitative analyses of how viewing-region shape interacts with scene content to guide gaze.

Saliency evaluation metrics. Saliency models are commonly evaluated using a family of metrics that capture different aspects of agreement between predicted maps and human fixations [4]. AUC-Judd and shuffled AUC (sAUC) treat saliency as a binary classifier over fixation versus non-fixation locations; Normalized Scanpath Saliency (NSS) measures the standardized saliency at fixation points; Pearson's Correlation Coefficient (CC), KL divergence, and Similarity (SIM) compare predicted and empirical saliency distributions; and Information Gain (IG) quantifies improvement over a baseline model, often a center prior. We adopt this standard suite for our baseline analyses, adapting implementations to evaluate predictions strictly within the polygon-defined visible region.

3 Dataset Description

3.1 Participants

Thirty participants were recruited (ages 24–29, $M = 26.3$; 14 female, 16 male). All reported normal or corrected-to-normal vision and are included in the released dataset. The study was approved by the institutional ethics board and all participants provided informed consent. The data that was collected was de-identified and constructed according to applicable ethical guidelines and regulatory requirements.

3.2 Apparatus

Eye movements were recorded with an EyeLink 1000 Plus (SR Research Ltd.) in head-mounted configuration [12]. Stimuli were presented on a 27″ 4K monitor (3840 × 2160 pixels) at 144 Hz refresh rate. Viewing distance was 70 cm, yielding a spatial resolution of approximately 78.5 pixels/degree of visual angle. The first three participants (P01–P03) were recorded monocularly at 1000 Hz (right eye) during Part A (See below). Following a protocol update, their Part B sessions and all subsequent participants were recorded binocularly at 500 Hz per eye. Thirteen-point calibration and validation were performed at the start of each session, with drift correction before each trial.

3.3 Stimuli

We selected 600 images from the CAT2000 dataset [2] spanning six categories: Fractal, Object, Outdoor Natural, Random, Satellite, and Sketch. Images were displayed at native 3840×2160 resolution. Each image was paired with one of 27 polygonal apertures: the image was visible only within the polygonal region, with the surrounding area filled with a uniform neutral (gray) background. This produced 648 unique polygon–image combinations across the experiment. Figure 1 shows example stimuli with their polygonal apertures and fixation density maps.

3.4 Procedure

The experiment was divided into two parts (A and B), each containing 9 mini-blocks of 39 trials, for a total of 702 trials per participant. Of these, 648 were image-viewing trials and 54 were empty polygon trials (where the stimuli were polygons colored uniformly in gray with black edges). Each trial proceeded as follows:

(1) A small circular cue appeared at one of pre-determined 9 points on a grid (25%, 50%, 75% of the screen's height and width) locations.
(2) Upon stable fixation (verified by drift check), the stimulus was presented for 4 s.
(3) After a 0.5 s inter-trial interval, the next trial began.

To ensure attentive viewing, an old/new memory probe was presented after each mini-block: participants judged whether a briefly shown image had appeared in the preceding block. Mean probe accuracy was 60% (above the 50% chance level), with a correct rejection rate of 83% for novel images, confirming that participants attended to the stimuli. The released metadata includes all 702 trials per participant; fixation and saliency analyses are based on the 648 image-viewing trials. Parts A and B were conducted in separate sessions on different days.

3.5 Polygonal Apertures

The 27 polygonal apertures were designed to span a wide range of geometric properties—including symmetry, convexity, elongation, and boundary irregularity—so that their influence on fixation behavior can be studied under controlled conditions. The shapes

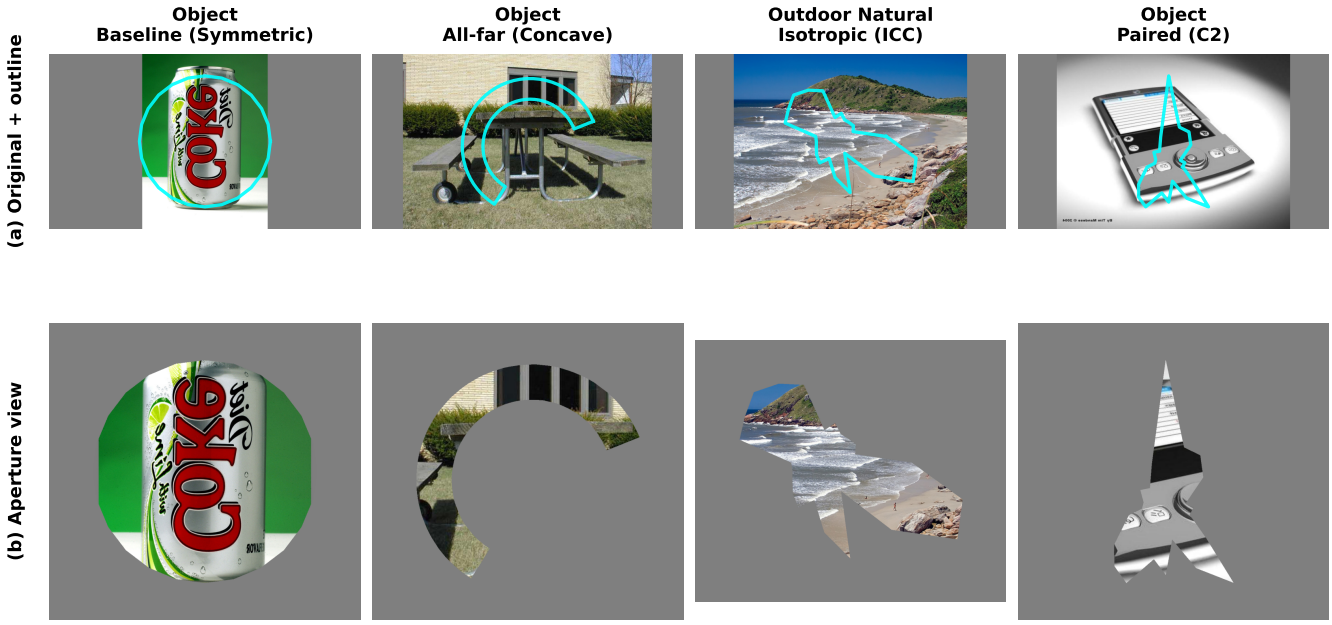


Figure 1: Example stimuli from PolyCAT. Each column shows a different polygon–image pair spanning the three polygonal condition groups: reference symmetric, convexity-varied (concave), and two irregular asymmetric shapes. Row (a): original image with polygon outline (cyan). Row (b): aperture view as presented to participants (cropped to polygonal region).

were produced using a computational generation and optimization procedure that targeted diverse geometric configurations while maintaining comparable area.

The polygon set comprises three main condition groups:

- **Reference baseline apertures (3):** A rectangle and two blob-like polygons (one symmetric, one mildly asymmetric), providing simple, familiar geometries as baseline conditions.
- **Convexity-varied apertures (3):** Three shapes of similar area spanning the convexity spectrum:fully convex, intermediate, and concave, enabling investigation of how boundary curvature and concavity influence gaze within the aperture.
- **Irregular asymmetric polygonal apertures (21):** A large set of diverse irregular polygons generated by a computational procedure that targeted seven distinct geometric configurations (×3 exemplars each), varying properties such as vertex distribution, elongation, and mass asymmetry to broadly sample the space of irregular polygonal geometries. All shapes in this group are convex.

4 Data Processing

4.1 Gaze Extraction

Raw EDF files were converted to sample-level gaze tables using custom Python scripts built on the `pyedfread` library. For binocular recordings, each time point yields separate left- and right-eye samples. A known issue with EyeLink binocular recordings is that fixation event coordinates in the EDF use HREF (head-referenced) coordinates rather than screen pixels. We addressed this by averaging raw sample positions over each fixation interval to obtain accurate screen coordinates. Gaze data were stored as tab-separated files with columns for timestamp, eye identity, pixel coordinates, pupil size, and validity flag. With the number of participants and trials, this totaled 83.8 million samples across 65 session files.

4.2 Fixation Detection

Fixations were detected using EyeLink's built-in saccade-based algorithm with default thresholds: 30 ° s^{-1} velocity and 8000 ° s^{-2} acceleration. The resulting fixation events were aligned to trial timing using EyeLink messages embedded during recording. Each fixation was mapped to its corresponding `trial_uid` and annotated with whether it fell within the polygonal aperture region. The dataset contains 410,166 fixations across 20,903 trials, with a mean of 19.6 fixations per trial (SD = 6.6) and a mean fixation duration of 322 ms (SD = 290 ms). Next to the fixation-saccade scan paths, we also include the raw sample-level gaze streams (Section 6) to enable users to apply alternative fixation-detection algorithms as needed.

4.3 Quality Metrics

We computed per-trial tracking ratio as the proportion of the 4 s stimulus interval covered by valid fixation events. A trial was considered valid if its tracking ratio exceeded 0.80. Across all 30 participants, the mean tracking ratio was 0.81 (SD = 0.09). Mean data loss (proportion of trials with no gaze data) was 0.6%. One participant (P13) showed the highest data loss at 6.4%, while 15 participants had zero data loss. Table 1 summarizes per-participant quality statistics.

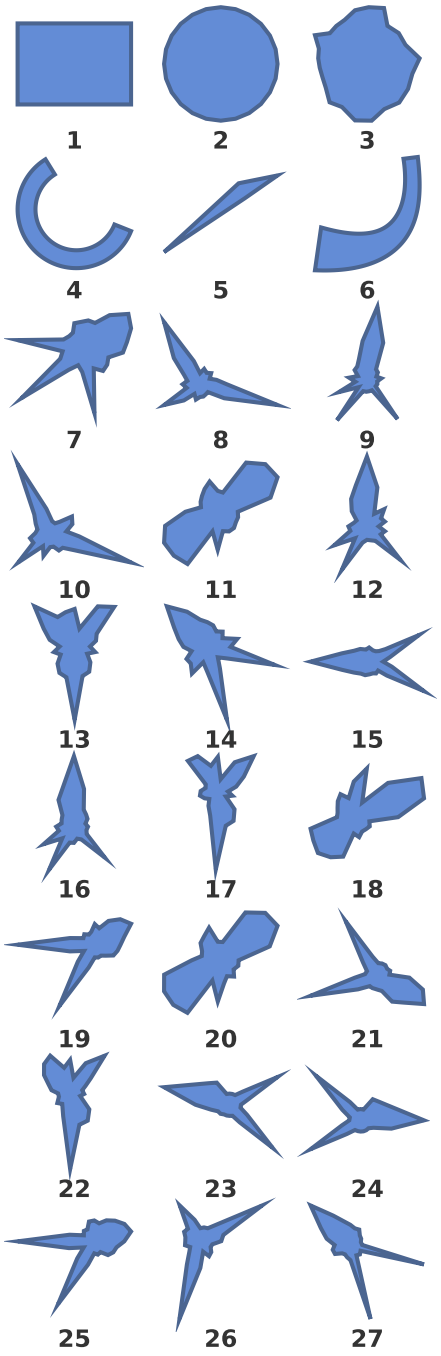


Figure 2: All 27 polygonal apertures used in the experiment, numbered 1–27. Shapes 1–3: reference baselines (rectangle, symmetric, asymmetric). Shapes 4–6: convexity-varied (concave, convex, intermediate). Shapes 7–27: irregular asymmetric polygons spanning seven geometric configurations with three exemplars each.

Table 1: Summary of data quality across 30 included participants.

Metric	Value
Participants included	30
Total trials	21,026
Valid trials (tracking ratio ≥ 0.80)	20,438 (97.2%)
Mean tracking ratio	0.81
Mean fixations per trial	19.6
Mean fixation duration	322.00 ms
Total fixations	410,166
Total gaze samples	83.8 M

4.4 Saliency Maps

Empirical saliency maps were generated for each of the 648 polygon–image combinations by aggregating fixation locations across all observers who viewed that combination. Fixation positions were convolved with a Gaussian kernel ($\sigma = 1°$, corresponding to approximately 78.5 pixels at 70 cm viewing distance). Note that the Gaussian kernel was applied identically regardless of polygon size. Maps were computed at the full stimulus resolution of 3840×2160 pixels. Values outside the polygonal aperture were set to zero, and each map was normalized to sum to one within the visible region. Maps were generated only for combinations with at least three observers. Fixations near polygon boundaries contribute density that extends beyond the visible region; this out-of-polygon density is zeroed before normalization.

5 Baseline Evaluation

We evaluated baseline and deep saliency models on all 648 polygon–image pairs, using the standard MIT saliency benchmark metrics [4]. All evaluations were performed at quarter resolution (960×540) for computational efficiency. Pairs are split into 492 training and 156 test pairs (by image identity) to support future model development; we report results on all pairs as well as the test set.

The **Uniform** baseline assigns equal probability to all pixels within the polygonal mask. The **Center Bias** baseline uses a Gaussian centered on the screen ($\sigma = 5°$), masked to the polygonal region, reflecting the well-documented tendency to fixate near the center of the display [13]. The **Gold (LOO)** baseline uses leave-one-observer-out cross-validation, providing an empirical ceiling on model performance. We additionally evaluate **DeepGaze IIE** [9], a state-of-the-art deep saliency model, under two conditions: *full*, where the model receives the original unoccluded image, and *masked*, where the model receives the polygon-masked stimulus (image visible inside the aperture, gray surround).

The Center Bias model achieves an AUC of 0.91, indicating that central fixation tendency is strong even within polygonal apertures. The Gold standard (LOO) yields an AUC of 0.97, sAUC of 0.81, and NSS of 4.56, establishing the inter-observer agreement ceiling. These values are comparable to those reported for the original CAT2000 dataset, suggesting that polygonal apertures do not fundamentally disrupt fixation consistency across observers. Performance is consistent across the train and test splits (Table 2), confirming the absence of data leakage.

Table 2: Saliency evaluation on 648 polygon–image pairs. ↑/↓: higher/lower is better. Bold and underline indicate best and second-best *models* (excluding the Gold ceiling).

Model	Input	AUC↑	sAUC↑	NSS↑	CC↑	KL↓	SIM↑	IG↑
Uniform	—	0.5000	0.5000	0.0000	0.0000	2.1907	0.1759	—
Center Bias	—	**0.9089**	0.5119	**2.1766**	**0.5568**	**0.9551**	**0.4482**	—
Gold (LOO)*	—	0.9710	0.8132	4.5604	0.9994	0.0020	0.9846	+1.6019
DeepGaze IIE [9]	full	0.8818	0.5375	1.4524	0.3663	1.3231	0.3373	−0.5777
DeepGaze IIE	masked	<u>0.8948</u>	**0.5694**	<u>1.6794</u>	<u>0.4101</u>	<u>1.2398</u>	<u>0.3600</u>	**-0.4445**

* *Gold (LOO) represents the inter-observer ceiling and is excluded from model ranking.*

DeepGaze IIE applied to the full (unoccluded) image achieves an AUC of 0.88 but falls below the Center Bias baseline on all other metrics (NSS 1.45 vs. 2.18, CC 0.37 vs. 0.56), with negative Information Gain (−0.58). This is expected: the model predicts saliency for the entire scene, while participants can only fixate within the polygonal aperture. When provided with the polygon-masked stimulus, DeepGaze IIE improves across all metrics (AUC 0.89, NSS 1.68, CC 0.41, IG −0.44), confirming that the aperture information helps. However, even the masked condition remains substantially below Center Bias and far below the Gold ceiling, indicating significant room for improvement in aperture-aware saliency prediction. Figure 3 provides a qualitative comparison of model predictions.

6 Dataset Structure

PolyCAT is organized into the following directories:

- **`data/metadata/`** Core metadata tables: `participants.csv` (30 rows with demographics and inclusion status), `sessions.csv` (session selection and timing), `trials.csv` (21,026 rows with full trial-level detail including a `split` column assigning each image to "train" or "test"), and `quality_metrics.csv` (per-participant quality statistics). The train/test split (492/156 polygon–image pairs) is by image identity, ensuring no image appears in both sets.
- **`data/gaze/`** Sample-level gaze streams as TSV files (one per session, 65 files total). Columns: participant, part, session, trial identifier, timestamp, eye, pixel coordinates, pupil size, validity.
- **`data/fixations/`** Fixation events as CSV (`fixations_all.csv` and per-participant files). Each fixation includes trial identifiers, timing relative to stimulus onset, pixel and degree coordinates, and an `in_polygon` flag.
- **`data/saliency_maps/`** Empirical saliency maps in NumPy (`.npy`) and PNG formats, organized by polygon ID and image ID (648 maps).
- **`data/scanpaths/`** Per-trial ordered fixation sequences as CSV files.
- **`data/stimuli/`** Polygon definitions (27 JSON files with vertex coordinates) and geometry manifest.
- **`benchmarks/`** Baseline evaluation results and evaluation scripts.
- **`code/`** Preprocessing, evaluation, and visualization scripts.

All coordinate data use a top-left origin pixel system on the 3840×2160 display. Degree-of-visual-angle coordinates are provided relative to screen center. A MATLAB-compatible export (`fixations.mat`) mirrors the CAT2000 format for direct compatibility with existing saliency evaluation toolboxes. The full dataset is documented with a data dictionary, file format specifications, and a Croissant metadata descriptor [1]. Data are released under a CC BY 4.0 license.

7 Discussion

Polygonal apertures and fixation behavior. The strong Center Bias performance (AUC = 0.91) indicates that the tendency of participants to fixate near the center of the display persists even when the image content is only partially visible. However, the polygon geometry modulates this tendency: the diverse set of irregular shapes, which vary systematically in symmetry, convexity, and boundary complexity, provides the experimental leverage needed to study how aperture geometry shapes gaze allocation beyond simple central tendency. The poor performance of DeepGaze IIE on the full image (IG < 0) highlights that standard saliency models, trained to predict fixations on standard unconstrained scenes, do not transfer well to aperture-constrained viewing. Adapting these models to reason about visible regions represents an open challenge for which PolyCAT may be constructive. Detailed analysis of shape-specific effects on fixation distributions is beyond the scope of this dataset paper and is the subject of ongoing work. Recent modeling of space-dependent saccadic biases in polar coordinates [10] suggests that the interplay between launch-site position and aperture boundaries may further shape scanpath dynamics, a direction enabled by PolyCAT's full gaze stream and scanpath data.

Data quality. The mean tracking ratio of 0.81 and 97.2% valid trial rate indicate high overall data quality. The fixation statistics (mean 19.6 fixations per 4 s trial, mean duration 322 ms) are consistent with typical free-viewing behavior on natural scenes. The small number of early participants recorded monocularly does not affect the fixation-level analyses, as coordinates are derived from the recorded eye in all cases.

8 Conclusion and Future Work

We presented PolyCAT, a public eye-tracking dataset that introduces polygonal apertures as a controlled manipulation for studying eye movements and visual attention on natural scenes. With 30 participants, 648 polygon–image combinations, 410,166 fixations, and comprehensive processed data products, PolyCAT provides a rich testbed for saliency modeling under viewing constraints. Baseline evaluation with various models demonstrates that existing saliency

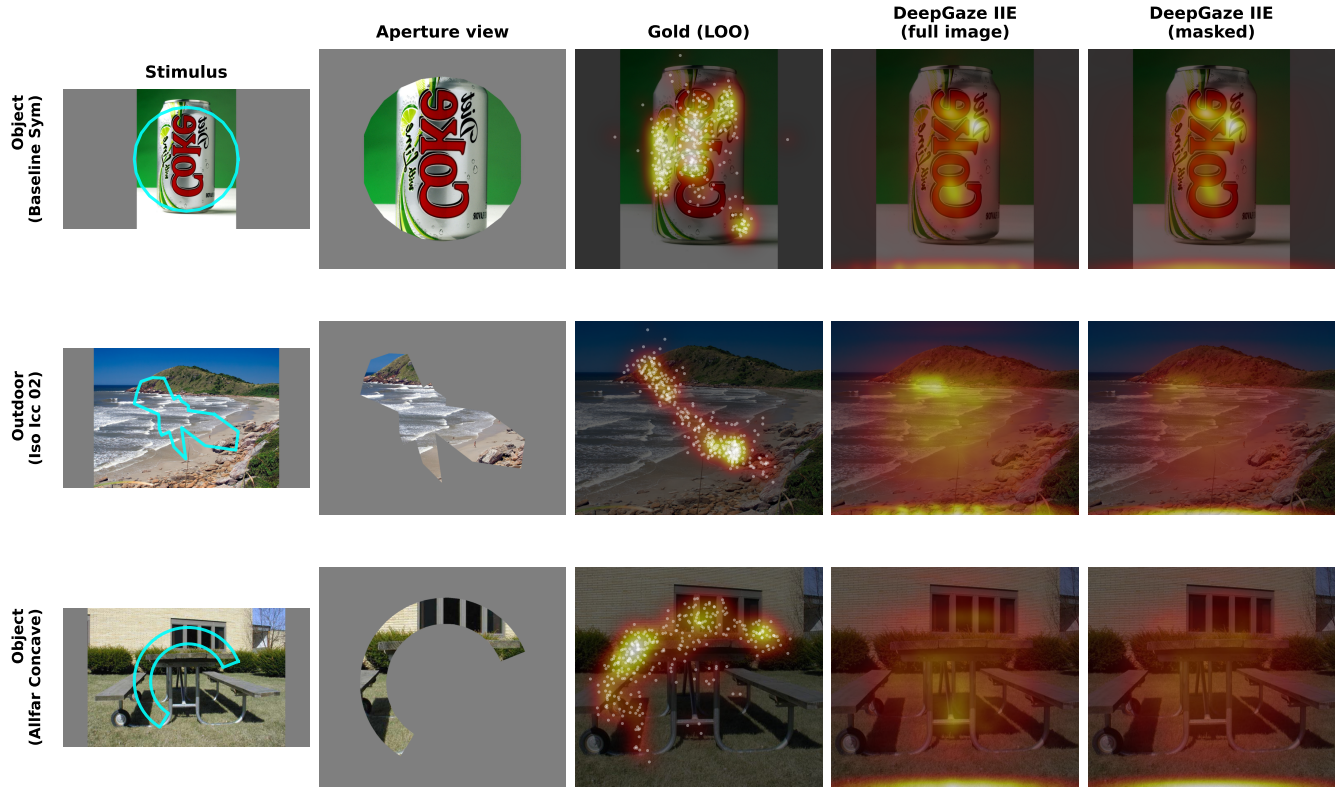


Figure 3: Qualitative comparison of saliency predictions across three polygon–image examples. Columns from left to right: original image with polygon outline, aperture view as presented to participants, Gold (leave-one-out) empirical map with individual fixation points (white dots), DeepGaze IIE prediction from the full (unoccluded) image, and DeepGaze IIE prediction from the polygon-masked stimulus. All saliency overlays are cropped to the polygonal region.

theories struggle with aperture-constrained viewing, motivating future work on aperture-aware attention models.

Indeed, PolyCAT enables several research directions: (1) adapting existing saliency models to predict fixations under partial occlusion, (2) studying how polygon geometry—symmetry, convexity, and irregularity—modulates gaze allocation, (3) developing aperture-aware attention models that jointly reason about image content and viewing geometry, and (4) investigating scanpath dynamics under spatial constraints. Future work may also extend the paradigm to dynamic apertures that change during viewing, increasing ecological validity for settings such as natural occlusion and head-mounted displays. Expanding stimulus coverage to the full CAT2000 collection and varying the engagement task (e.g., visual search versus pure free viewing) would further clarify how task demands interact with aperture geometry to shape gaze behavior.

The dataset, code, and documentation are publicly available under a CC BY 4.0 license at **HuggingFace** and also available on the project website at The **I**nterdisciplinary **C**omputational **V**ision **L**aboratory website (**ICVL Laboratory Website**), and we hope that making it available can advance these and many other research directions.

References

[1] Mubashara Akhtar, Omar Benjelloun, Costanza Conforti, Luca Foschini, Joan Giner-Miguelez, Pieter Gijsbers, Sujata Goswami, Nitisha Jain, Michalis Karamousadakis, Michael Kuchnik, Satyapriya Krishna, Sylvain Lesage, Quentin Lhoest, Pierre Marcenac, Manil Maskey, Peter Mattson, Luis Oala, Hamidah Oderinwale, Pierre Ruyssen, Tim Santos, Rajat Shinde, Elena Simperl, Arjun Suresh, Goeffry Thomas, Slava Tykhonov, Joaquin Vanschoren, Susheel Varma, Jos van der Velde, Steffen Vogler, Carole-Jean Wu, and Luyao Zhang. 2024. Croissant: A Metadata Format for ML-Ready Datasets. arXiv:2403.19546 [cs.LG] https://arxiv.org/abs/2403.19546

[2] Ali Borji and Laurent Itti. 2015. CAT2000: A Large Scale Fixation Dataset for Boosting Saliency Research. In *CVPR 2015 Workshop on Future of Datasets*. http://saliency.mit.edu/cat2000

[3] Zoya Bylinskii, Phillip Isola, Constance Bainbridge, Antonio Torralba, and Aude Oliva. 2015. Intrinsic and Extrinsic Effects on Image Memorability. *Vision Research* 116 (2015), 165–178. doi:10.1016/j.visres.2015.03.005

[4] 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.

[5] Laurent Itti, Christof Koch, and Ernst Niebur. 1998. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20, 11 (1998), 1254–1259.

[6] 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 (CVPR)*. 1072–1080.

[7] Tilke Judd, Krista Ehinger, Frédo Durand, and Antonio Torralba. 2009. Learning to predict where humans look. In *IEEE International Conference on Computer Vision (ICCV)*. 2106–2113.

[8] Matthias Kümmerer, Thomas S.A. Wallis, and Matthias Bethge. 2015. Information-theoretic model comparison unifies saliency metrics. In *Proceedings of the National Academy of Sciences*, Vol. 112. 16054–16059.

[9] Akis Linardos, Matthias Kümmerer, Ori Press, and Matthias Bethge. 2021. DeepGaze IIE: Calibrated prediction in and out-of-domain for state-of-the-art saliency modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 1604–1613. https://arxiv.org/abs/2105.12441

[10] Rotem Mairon and Ohad Ben-Shahar. 2025. The Polar Saccadic Flow model: Re-modeling the center bias from fixations to saccades. *Vision Research* 228 (2025), 108546. doi:10.1016/j.visres.2025.108546

[11] Rotem Mairon and Ohad Ben-Shahar. 2025. Stimulus Center Bias Persists Irrespective of Its Position on the Display. *Journal of Eye Movement Research* 18, 6 (2025). doi:10.3390/jemr18060077

[12] SR Research Ltd. 2023. EyeLink 1000 Plus Technical Specifications. https://www.sr-research.com/eyelink-1000-plus/.

[13] Benjamin W. Tatler. 2007. The central fixation bias in scene viewing: Selecting an optimal viewing position independently of motor biases and image feature distributions. *Journal of Vision* 7, 14 (2007), 4.

[14] Jianxiong Xu, Ming Jiang, Shuo Wang, Mohan S. Kankanhalli, and Qi Zhao. 2014. Predicting Human Gaze Beyond Pixels. *Journal of Vision* 14, 1 (2014), 28. doi:10.1167/14.1.28