arXiv:2406.06967v1 [cs.CV] 11 Jun 2024

Dual Thinking and Perceptual Analysis of Deep Learning Models using Human Adversarial Examples

Kailas Dayanandan
Indian Institute of Technology
Delhi, 110016
kailasd@gmail.com

Anand Sinha
Indian Institute of Technology
Delhi, 110016
anand.sinha85@gmail.com

Brejesh Lall
Indian Institute of Technology
Delhi, 110016
brejesh@ee.iitd.ac.in

Abstract

The dual thinking framework considers fast, intuitive processing and slower, logical processing. The perception of dual thinking in vision requires images where inferences from intuitive and logical processing differ. We introduce an adversarial dataset to provide evidence for the dual thinking framework in human vision, which also aids in studying the qualitative behavior of deep learning models. Our study also addresses a major criticism of using classification models as computational models of human vision by using instance segmentation models that localize objects. The evidence underscores the importance of shape in identifying instances in human vision and shows that deep learning models lack an understanding of sub-structures, as indicated by errors related to the position and number of sub-components. Additionally, the similarity in errors made by models and intuitive human processing indicates that models only address intuitive thinking in human vision.

1 Introduction

Dual thinking framework argues that humans have a fast intuitive system and a slow logical system, and has been widely mentioned in recent years [92, 90, 62, 19, 16]. A growing body of electrophysiological research suggests gist formation in initial $150 - 200ms$ using fast feed-forward processing, followed by slower iterative refinement [92, 35, 49, 86, 90, 85]. The dual thinking framework is not part of studies on visual perception due to the unavailability of methods to study them. Gestalt principles have been studied for over a century [94, 91] indicating its presence in human visual processing [71]. Gestalt theory proposes that humans perceive things as a whole rather than as a sum of parts [98, 72], which was a deviation from structuralism that considered analyzing parts as a key to understanding the whole object [94, 95]. The segmentation process in human vision requires explaining properties that help identify the validity of the possible groupings and the class of the instance, which the theories on visual perception need to address. The computational models using deep learning helped enhance the understanding of human vision [93, 49, 77, 48, 50], however, they face criticism for using classification models unlike human vision that localizes objects [12, 78, 8]. The comparative studies can help improve accuracy, robustness, and generalizability of models by incorporating strategies from human vision [65, 65, 31, 20, 57, 56, 27]. The analysis of computational aspects or the algorithmic component of Marr's three-level hypothesis [5, 59, 70] can help in better

Preprint. Under review.

integration of strategies, however, such analysis has been ignored in recent years as end-to-end deep learning models gained prominence [101, 68, 80]. The code for our work is shared at [1]

Our main contributions in this paper include

(a) We introduce the Human Confusion Dataset, a diverse dataset to study dual thinking and strategies in human vision. We leverage the dataset to study the behavioral properties of instance segmentation models by introducing a method to evaluate the models.

(b) We observe that logical processing in the dual thinking framework is essential to define human visual processing completely. Our study shows that logical processing is still nascent in deep learning models, and we provide insights into common logical errors in deep learning models. This can help the machine learning community critically examine these vulnerabilities.

(c) We show that instance segmentation models can serve as computational models of human vision and were able to observe common characteristics associated with human vision. The computational analysis of shape bias over texture in the formation of instances shows many advantages, and the wrong grouping of related sets of pixels indicates sub-component formation in intuitive processing.

2 Related Work

Deep learning models are not robust or generalizable despite attaining human-level performance on many tasks, leading to a comparison of strategies in deep learning models and human vision [7, 34, 40]. The comparisons are based on external behavior [32, 33, 30, 89] or by comparing internal representations [15, 44, 61, 74]. A significant finding in these studies is the texture bias in deep learning models [7, 34] and shape bias in human vision. The robust and accurate models had shape bias, which led to techniques to increase shape bias in deep learning models [34, 100, 54, 43, 83, 84]. However, models have shown good performance without shape information [13], and shape information can be extracted from later layers despite not being used for classification [40]. This indicates that shape may not be essential for discriminatory tasks. Subsequent work also shows that image stylization serves as a strong augmentation with shape bias emerging as a byproduct [63, 54]. This motivates the study of the importance of shape in instance segmentation tasks, which is similar to human vision that localizes objects. The computational analysis of segmentation task in cognitive science represents segmentation as interval graphs [2, 3, 29], which do not directly extend to image segmentation as instances can be composed of non-adjacent or disconnected segments (e.g., amodal closure in fig.6h), requiring separate analysis.

(a) (b) (c) (d) (e) (f)

Figure 1: (a) Global Shape [7] (b) Shape vs Texture (SIN) [34] (c) Edges [96] (d) Perceptual Abstraction [22] (e) Gazing patterns [46] (f) Gazing patterns [57]

Deep learning models made many recent advances, helping them approach human-level performance on many tasks [11, 10, 45, 33, 67, 69, 38] including on benchmark datasets like ImageNet [76]. This led to the introduction of many datasets that expose different vulnerabilities of deep learning models [1, 4, 41, 23]; including adversarial datasets like Image Net-A, consisting of natural images that were tough for the ResNet-50 [39]. However, they are easier for humans and lack instance-level annotations. With deep learning models making progress on many of these adversarial datasets [52, 53], there is a need for datasets that are challenging for human vision to evaluate the progress and understand the behaviors and strategies in human vision and deep learning models. Stylized Image Net (SIN) enforces conflict between texture and shape by synthetically transferring texture from one class to object of another class [34]; however, it lacks diversity as it can study only texture shape bias and does not have a correct ground truth class. In contrast, images in our dataset have a ground truth class (e.g., fig. 3b in our dataset is a dog and not a man with dog texture, whereas fig.1b

[1] https://github.com/kailasdayanandan/dual_thinking

from SIN is a cat shape with an elephant texture). The images in our datasets have a single stable state compared to bi-stable images. Similarly, another set of datasets focuses only on camouflaged or concealed objects (or figure-ground errors) [25, 51, 24, 60, 82]. These datasets cannot evaluate instance segmentation models trained on benchmark MS-COCO dataset [55] as the class of objects do not match. We do not constrain on any specific cue, which helps us address the gap in research about datasets that help us explore multiple aspects of human visual processing and the behaviors of deep learning models. Our dataset has images with inherent cue conflicts, with the cue for the correct perception contradicting the cue human vision uses for intuitive processing.

Gestalt principles and Structuralism are prominent theories about human visual perception [8, 47]; however, they are insufficient to describe class identification and complex logical processing in human vision. The theories are also independent of findings about initial gist formation in $200ms$ observed in many electrophysiological studies [92, 49, 62, 35]. Our dataset facilitates the analysis of initial gist formation and logical processing in visual perception as multiple inferences are perceivable. It also helps evaluate the importance of shape and specific regions that human vision focuses on during visual processing [46, 57, 96] (fig. 1e, 1f). Our paper proposes a method to evaluate the model's qualitative behaviors and understand the inductive biases, which can help while designing systems for critical applications [21] in resource-constrained settings [6, 99].

3 Dual Thinking Framework

In this section, we discuss the computational advantages of using a dual-thinking framework for visual processing and the significance of existing theories on human vision within this framework.

Theorem *: If we have multiple possible groupings involving a common region, with different prior probabilities, the lowest average number of iterations to reach correct grouping or segmentation is achieved by completely assigning the component to highest probabilities sequentially.*

Let $p_1, p_2, p_3..p_n$ be the probability in the order of most probable grouping, such that $p_1 > p_2$ and $p_2 > p_3$ and so on and there are n possible groupings. The most probable grouping will be correct in the first iteration for p_1 fraction of the time, and the next highest probable grouping will take two iterations and be correct p_2 fraction of the time, and so on. The average length is given by

$$L_{avg} = 1p_1 + 2p_2 + 3p_3... + np_n = \sum_{i=1}^{n} ip_i \tag{1}$$

Refer Supplementary data for proof of the theorem by contradiction method. We can also note that the $\sum_{i=1}^{n} p_i = 1$ and when $p_1 \to 1$ other p_i tends to 0 for $i \neq 1$.

$$\lim_{p_1 \to 1} L_{avg} \to 1 \tag{2}$$

The intuitive processing in human vision is quite accurate; hence, we usually do not perceive the first inference. However, multiple inferences are perceivable in images in our dataset as the first inference is wrong and differs from the final inference. For example, in figure 6e referenced later, the circular region can have two possible groupings: a car tire or a mirror in the bike, with the car tire being the human intuitive perception and the mirror in the bike being the correct perception.

Sub-component Formation Let us consider a group of pixels **I** = $\{p_{i,j}\}$ where $p_{i,j}$ represent pixel at (i, j) position. Instance segmentation creates disjoint and nonempty subsets of pixels (s_i) that correspond to object instances or the background and satisfy the conditions that each pixel can belong to only one instance, and the segments must collectively cover the entire image (eq. 3).

$$S = \{s_0, s_1, ..s_T\} \quad where \quad s_i \cap s_j = \emptyset \quad and \quad \bigcup_{i=0}^{T} s_i = I \tag{3}$$

$$c_i = \{p_{i,j}\} \quad where \quad \bigcup_{i} c_i \quad \cup s_0 = I \quad and \quad c_i \cap c_j = \emptyset \tag{4}$$

In segmentation, the complexity or hardness is informally considered as the size of the search space [3, 2, 28], and the number of possible instances or the subsets possible for the set $\mathbf{I}$ is $O(2^K)$, where K is the number of pixels in the image. A common approach is to reduce the number of elements [14, 29]. We can observe that figure-ground extraction can reduce a significant number of pixels by ignoring the background (s_0). The formation of sub-components (C) using Gestalt principles (e.g., proximity and similarity) or edges can reduce possible combinations from $O(2^K)$ to $O(2^{card(C)})$. This sub-component formation is also amenable for iterative processing as rectification of incorrect grouping in pixel-based processing can result in redundant intermediate steps observed in segmentation models that use iterative processing [97, 87] (similar to fig.6e in the results section).

Instance and Class Identification The shape can help evaluate the possibility of the groupings in this search space to form an instance (e.g. in fig.6d, with more probable shapes that correspond to common postures can be prioritized for faster computation in intuitive processing (similar to Theorem 1), and can be evaluated till $\bigcup_i c_i \cup s_0 = I$ is satisfied (e.g. fig.6d). The space of all possible groupings of components can be searched optimally to identify instances using shape as it prioritizes sub-groups that fit the different object shapes or postures.

4 Dataset and Methods

In the first part, we introduce Human confusion dataset consisting of images adversarial for human vision to study dual thinking framework and deep learning models. We then propose a method to automatically evaluate model correctness and similarity to intuitive processing in human vision.

4.1 Human confusion dataset

Grouping elements with similar properties, including brightness, color, contrast, texture, etc, is called similarity (e.g., fig.2c). Continuity is important in human vision and is concerned with grouping elements that form smooth contours (e.g fig.2d), and proximity refers to closer elements likely to be grouped [94] (e.g fig.2b). Human vision focuses only on salient regions of the image (or foreground) for analysis, and an example of a figure-ground error is shown in fig.2a. Amodal closure denotes the completion of occluded elements [66, 64], and errors mainly occur when we perceive parts of different instances as a single instance (e.g., fig.2f). The errors can also be due to a combination of many properties, for example, similarity, proximity, and continuity in fig.2c,2d,2e. A comprehensive review of Gestalt principles is provided in [94, 95].

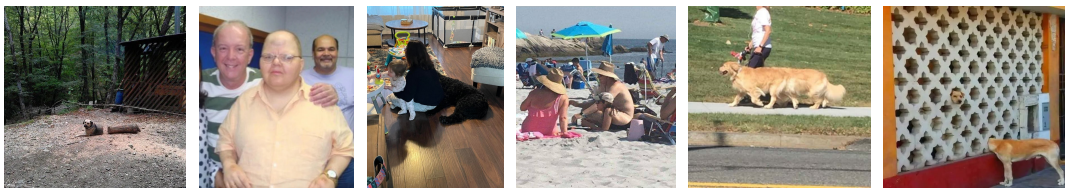

(a) (b) (c) (d) (e) (f)

Figure 2: (a) Figure ground and Similarity error with a part of background grouped with the dog due to similarity in color (b) Proximity error where the hand is grouped wrongly to the closer to the person in the middle. (c,d,e) Similarity, Continuity and Proximity related error occurring together involving similarity of color, the continuous regions forming smooth curves and sub-components being closer to each other. (f) Amodal Closure and Silarity related error where components from two instances that are apart grouped together as a single instance. (Refer Supple. for more details.)

Data Collection We crowd-source the initial data using Amazon Turk, where participants collect 150 confusing images by searching for "confusing photos", "images that need to be looked twice to understand," and similar terms, and were paid $30 each for the same. The initial three participants collected generic images, and the following four participants collected confusing images related to common objects. To expand the dataset, we engaged three additional participants to collect confusing images that were not part of the crowd-sourced set, as there was significant overlap and many synthetic images among the collected data using Amazon Turk, and were paid $300 each. In

total, ten participants contributed to the image collection, including seven from Amazon Turk and three for dataset expansion with a total payment of $1110.

Data Preparation We remove duplicate and synthetic images from the collected dataset and filter images where the instance segmentation task cannot identify the confusion (e.g., fig.3d the confusion concerns object interaction with each other or the environment and not related to instance segmentation). The final dataset consists of 1000 images without any NSFW content in correct perception with confusing instances belonging to classes in MS-COCO dataset [55]. We annotate each image on **(a)** Gestalt principles (figure-ground, proximity, continuity, similarity, and amodal closure) except modal closure to analyze qualitative behavior, **(b)** logical errors (size and count difference) to evaluate vulnerabilities, **(c)** instances with global shape of one instance but containing instances of different class to evaluate shape bias, and **(d)** images in wild that can help understand out of distribution behavior. The annotation also includes the operations and masks to check the correctness and similarity to the intuitive perception described in the next section.

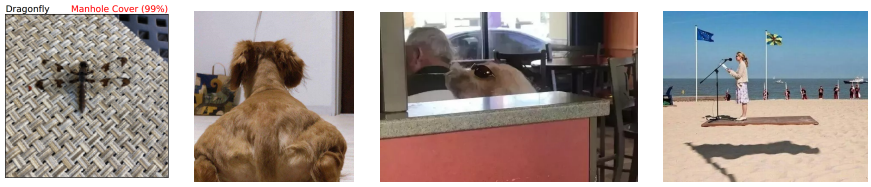

(a) (b) (c) (d)

Figure 3: (a) Images in ImageNet-A confuse deep learning models but not human vision [39] (b,c) Part of our dataset (d) Images that cannot be analyzed using segmentation models

4.2 Automatic Evaluation

In our analysis of images in the dataset, we observe that the errors in human perception mainly fall into the following categories **(a)** an object is missed out completely, **(b)** a part of an instance or background is wrongly grouped with another instance **(c)** two instances group to form a single instance **(d)** a part of the background identifies as an object or **(e)** an instance is misclassified to be of a different class. We evaluate the confusing region separately to prevent it from being overwhelmed by the overall error [18]. The basic algorithm that checks whether a part of the background is wrongly added to a confusing object instance is shown in alg.1 and an example in fig.4. We identify the confusing instance using the Intersection over Union (IoU) score with the instance mask and the model outputs. We use separate masks to identify human perception and correct perception to make evaluation independent of the size of the confusing region. The basic algorithm is modified to address the different types of errors described earlier. In case of errors, where the intuitive perception completely misses the confusing object instance (e.g., camouflaged animals), we skip the second part of the algorithm that evaluates the confusing region. In the case of errors, where two instances appear as a single instance, we run the basic algorithm using two sets of instance masks for each object to evaluate for correct perception and a mask with these two instances together to verify intuitive perception.

(a) (b) (c) (d) (e) (f)

Figure 4: An example showing automatic evaluation of correctness and similarity to human vision. (a) Input image (b) Mask of confusing Instance (c) Identified correct instance using IoU (d) Extracted Instance (e) Confusion region mask (f) Confusion Region Overlap = 0.013 < threshold.

Experimental Setting We use models from the MMDetection framework [17] trained on the MS-COCO dataset [55] and consider model outputs above the confidence threshold of 0.7. We select

Algorithm 1 Basic automatic evaluation algorithm

Input: modelOutput, instanceMask, confusionMask, includeConfusion, cThreshold
Output: correctStatus

matches ← []
for each $instance \in$ modelOutput **do**
  score ← IoU($instance$, instanceMask)
  if (score >match_threshold) **then**
    matches.insert(($instance$,score))
  end if
end for
matches.sort(key=score,reverse=True)
result = $False$
for each $match, score \in$ matches **do**
  overlap ← IoU(($match$*confusionMask, confusionMask)
  if ($overlap$ > cThreshold) **and** includeConfusion **then**
    result = $True$
  end if
  if ($overlap$ < cThreshold) **and not** includeConfusion **then**
    result = $True$
  end if
end for
return result

the thresholds for the automatic evaluation by comparing human annotations on a randomly selected subset of 250 images from our dataset on seven models. The threshold for recognizing instance is 73%, 25% overlap of confusion region not part of an instance, and 70% for being part of an instance. The automatic detection algorithm matched the human annotation decisions in 94.94% of cases. We evaluate models on 984 images except for 16 images with confusion related to the labels (e.g., fig3b). We conducted experiments on a Dell machine with an Intel Xeon-E5-v4 CPU and NVIDIA Titan X (12 GB GPU) running Ubuntu 20.04.6 LTS.

5 Results

In this section, we provide evidence for the dual thinking framework in human vision and discuss different strategies used in these stages. We then analyze deep learning models using the method described earlier.

Dual thinking framework We could perceive multiple inferences, including the wrong first inference followed by a correct inference in our dataset, indicating the presence of a dual-thinking framework. The initial inference is corrected almost immediately (e.g., fig.5d,5c), which is in contrast to optical illusions used to demonstrate the shortcomings of human vision, where the decision does not usually change (fig.5a). In some images, the logical processing does not identify the errors; however, the users can recognize them retrospectively, which validates a common assumption in computational models of human vision about processing till a confidence threshold is attained [88]. For example, the camouflaged bird in fig.6g is not easily identifiable, but as these are not limitations of human vision, the users can identify the bird when mentioning the location.

Instance and Class Identification Shape as a method for grouping components to form instances is observable in grouping errors due to common posture (e.g. fig.6d), and the examples where two components from two instances with the shape of an instance (e.g., fig.5c). The priority for shape over color for grouping components in intuitive processing is observable in amodal closure errors, where sub-components of a single class with widely different colors are grouped as they have the shape of a single instance. The shape also had priority over length in grouping components that are widely separated (e.g., fig.6h). Our dataset provides evidence for the importance of shape in instance identification compared to existing studies that focus on shape for classification. We can observe the formation of sub-components, which is evident in errors that occur in groups of related pixels. Gestalt

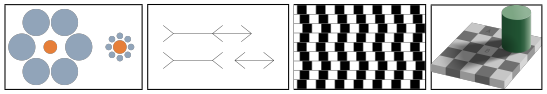

(a) Optical illusions are usually shortcomings of human vision (Ebbinghaus, Müller-Lyer, Café Wall, Checker Shadow Illusion.

(b) In dual thinking the false perception in intuitive processing is easily identified and immediately corrected.

(c) Shape can group components even without specific regions like eyes in intuitive perception.

(d) Shape is an important factor in grouping sub-components.

Figure 5: Multiple inferences indicates dual thinking framework.

principles and edge-based processing can help in sub-component formation; however, the adversarial examples only indicate sub-component formation but cannot distinguish between these approaches.

5.1 Logical processing in Models

The mistakes in intuitive processing in human vision are corrected during subsequent processing, whereas models face a significant limitation in their inability to identify and rectify logical errors. The models lack sub-component structure and depend on probabilistic grouping of components, leading to errors related to count mismatches, size differences, and failure to recognize impossible patterns.

Size Difference : Human vision can distinguish instances where objects of different sizes appear together with the shape of a single object. The initial intuitive processing using the overall shape is prone to these errors, which are corrected in the second stage of processing. Deep learning models had lower accuracy ($4-16\%$) and made errors similar to human intuitive processing ($61-80\%$), as they are unable to recognize the relative size of components (fig.6a, 7a).

Count Mismatch : The components grouped using the most probable shape tend to group nearby components without sub-component structure, and most models had poor accuracy ($0-11\%$). For example, in fig.6d, the common shape is used to form the first instance, and the remaining components form the next instance in human vision. In many cases, we also observe all the components to be grouped in intuitive processing (e.g., fig.6b,6c). The models seem to rely on a probabilistic grouping of the components without substructures and identify them as a single instance. Deep learning models made the same mistake as intuitive processing in human vision in $47-71\%$ of images (fig.7b).

Impossible Patterns : Human vision can identify objects in poses which is not feasible (e.g. fig.6d,6e). In the case of deep learning models, this capability is missing as it relies on probabilistic grouping. Deep learning models also split sub-components and partially assign them to two instances by splitting on proximity rather than assigning a sub-component entirely to one of the instances. The formation of sub-components in human vision can help avoid these errors.

Deep learning models lack sub-component structure and rely on grouping sub-components without logical underpinnings. The poor performance of models in size difference, count mismatch, and impossible patterns with errors similar to intuitive processing indicates the absence of the second stage of processing that helps detect and recover from these errors. The recent efforts to use language models to perform high-level reasoning can help realize the logical processing [36, 79], especially with the large language model's ability to capture information [42, 81] that logical processing requires.

5.2 Strategies and Performance

Our study observes that human vision uses a top-down approach similar to Gestalt theory, focusing on overall structure and shape, compared to deep learning models that focus more on details, similar to the structuralist theory of human vision. In many cases, the models generate outputs with both

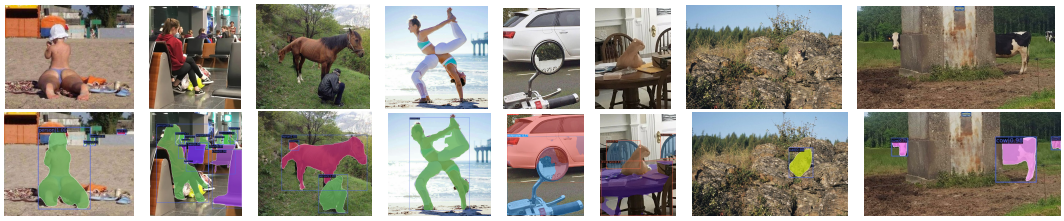

(a) (b) (c) (d) (e) (f) (g) (h)

Figure 6: Top row (a) Size Difference (b) Count Difference (c) Multiple shape match (d) Common Posture (e) Grouping Conflict (f) Texture over Shape (cover not cat) (g) Camouflage (h) Amodal closure - Bottom row Swin-T/YOLACT outputs

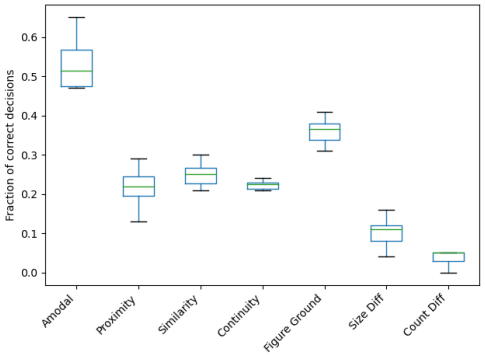

(a) Fraction of correct inference for different attributes (e.g. fig.6g, 6h are correct).

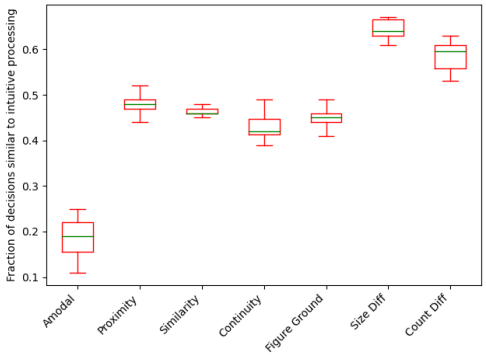

(b) Fraction of inferences that similar to mistakes in human intuitive processing (e.g. fig.6a, 6b)).

Figure 7: The model makes errors similar to human intuitive processing in count and has low accuracy, while they perform better on images with amodal closure errors due to low receptive field.

correct perception and perception similar to intuitive processing, as they do not learn the constraints in eq.3, a trivial logic in human visual processing. In general, the texture bias in deep learning models helped them perform slightly better on figure-ground errors ($25-41\%$) (e.g., fig.6f, 6g and 7a), and the smaller receptive field helps deep learning models avoid amodal closure errors in our adversarial dataset ($24-65\%$) (fig.7a). Despite better performance in these categories, none of the models could predict correctly on 47.87% of images, and swin-transformer was most accurate at 34% (fig.8). YOLACT model [9] designed for real-time operation with architecture that uses grouping by shape had a larger shape bias (as defined in [34]) of 9.2 on images with the global shape of one of class but comprised of two instances of different classes (e.g fig.5d) compared to next highest model GroIE that had a shape bias of only 5.5. In these images, intuitive processing prioritizes shape, whereas correct perception requires focusing on texture. YOLACT also had a higher ratio of human intuitive reference to correct inference overall and on Gestalt properties except for amodal closure errors. This along with shape bias shows the segmentation model's ability to model human vision and its characteristics. In general, models should be correct, robust, and generalizable and not limited by human performance or similarity in strategies to human vision. One of the challenges in realizing the second stage of processing is for the models to learn this as an emergent property using improved model architectures, appropriate loss function, and comprehensively represented training data.

Limitations : Ours is an adversarial dataset containing examples that require biasing on cues that are not commonly occurring; hence, better performance on this dataset may not directly translate to better metrics on benchmark datasets.

6 Conclusion

Our adversarial dataset provides insights into the dual-thinking framework in human visual processing, and computational analysis shows the advantages of human visual processing. The dataset highlighted

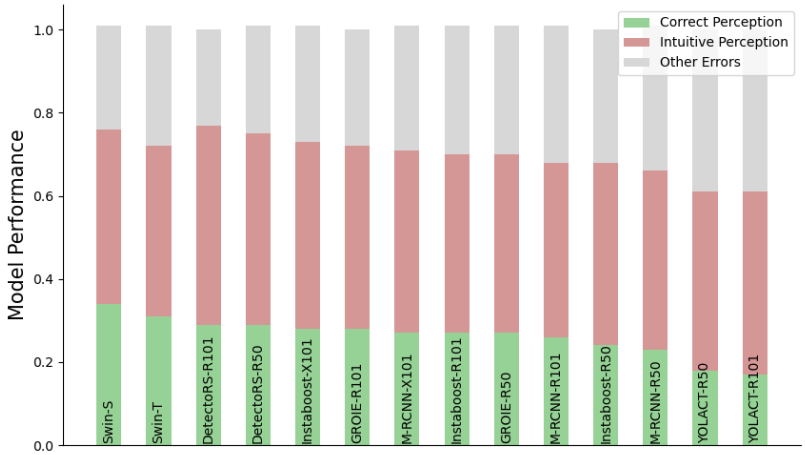

Figure 8: Overall model performance [75, 37, 9, 58, 73, 26]. The quantitative data for the plots are present in Supplementary data.

some of the vulnerabilities of deep learning models and research gaps in logical processing in cognitive science. The analysis of model behavior shows the limitations of models in performing logical processing present in the second stage of human vision. Our study shows that segmentation models can replicate characteristics of human vision, indicating their potential to serve as computational models for human vision. The dataset and the proposed approach can help evaluate changes in the qualitative behavior of models with changes in model architecture, training procedure, or dataset and augmentation approaches. Additionally, our dataset can help cognitive scientists evaluate new theories on human vision and validate assumptions about human visual processing in computational models of human vision.

Acknowledgments and Disclosure of Funding

The authors thank Dr.Sumeet Agarwal for initial guidance and discussions, Shubham Jain and Abhilash Kankokaran for insightful discussions and contributions. The authors also acknowledge the contributions of Tanushka, Nikhil VM, Shafi, Shahanas Babu, Abdul Vahab and others for their support in data collection and annotation.

References

[1] Amro Abbas and Stéphane Deny. Progress and limitations of deep networks to recognize objects in unusual poses. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023.

[2] Federico Adolfi, Todd Wareham, and Iris van Rooij. Computational complexity of segmentation. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, 2022.

[3] Federico Adolfi, Todd Wareham, and Iris van Rooij. A computational complexity perspective on segmentation as a cognitive subcomputation. *Topics in Cognitive Science*, 2022.

[4] Michael A Alcorn, Qi Li, Zhitao Gong, Chengfei Wang, Long Mai, Wei-Shinn Ku, and Anh Nguyen. Strike (with) a pose: Neural networks are easily fooled by strange poses of familiar objects. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4845–4854, 2019.

[5] Barton L Anderson. Can computational goals inform theories of vision? *Topics in Cognitive Science*, 7(2):274–286, 2015.

[6] Gregor Bachmann, Sotiris Anagnostidis, and Thomas Hofmann. Scaling mlps: A tale of inductive bias. *Advances in Neural Information Processing Systems*, 36, 2024.

[7] Nicholas Baker, Hongjing Lu, Gennady Erlikhman, and Philip J Kellman. Deep convolutional networks do not classify based on global object shape. *PLoS computational biology*, 14(12):e1006613, 2018.

[8] Valerio Biscione and Jeffrey S Bowers. Mixed evidence for gestalt grouping in deep neural networks. *Computational Brain & Behavior*, 6(3):438–456, 2023.

[9] Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. Yolact: Real-time instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9157–9166, 2019.

[10] Victor Boutin, Lakshya Singhal, Xavier Thomas, and Thomas Serre. Diversity vs. recognizability: Human-like generalization in one-shot generative models. *Advances in Neural Information Processing Systems*, 35:20933–20946, 2022.

[11] Victor Boutin and FEL Thomas. Diffusion models as artists: Are we closing the gap between humans and machines? In *International Conference on Machine Learning*, 2023.

[12] Jeffrey S Bowers, Gaurav Malhotra, Marin Dujmović, Milton Llera Montero, Christian Tsvetkov, Valerio Biscione, Guillermo Puebla, Federico Adolfi, John E Hummel, Rachel F Heaton, et al. Deep problems with neural network models of human vision. *Behavioral and Brain Sciences*, pages 1–74, 2022.

[13] Wieland Brendel and Matthias Bethge. Approximating cnns with bag-of-local-features models works surprisingly well on imagenet. In *International Conference on Learning Representations*, 2018.

[14] Michael R Brent. Speech segmentation and word discovery: A computational perspective. *Trends in Cognitive Sciences*, 3(8):294–301, 1999.

[15] Charles F Cadieu, Ha Hong, Daniel LK Yamins, Nicolas Pinto, Diego Ardila, Ethan A Solomon, Najib J Majaj, and James J DiCarlo. Deep neural networks rival the representation of primate it cortex for core visual object recognition. *PLoS Comput Biol*, 10(12):e1003963, 2014.

[16] Chen Chen, Kerstin Hammernik, Cheng Ouyang, Chen Qin, Wenjia Bai, and Daniel Rueckert. Cooperative training and latent space data augmentation for robust medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 149–159. Springer, 2021.

[17] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, et al. Mmdetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019.

[18] Bowen Cheng, Ross Girshick, Piotr Dollár, Alexander C Berg, and Alexander Kirillov. Boundary iou: Improving object-centric image segmentation evaluation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15334–15342, 2021.

[19] Kahneman Daniel. Thinking, fast and slow, 2017.

[20] Joel Dapello, Tiago Marques, Martin Schrimpf, Franziska Geiger, David Cox, and James J DiCarlo. Simulating a primary visual cortex at the front of cnns improves robustness to image perturbations. *Advances in Neural Information Processing Systems*, 33:13073–13087, 2020.

[21] Kailas Dayanandan and Brejesh Lall. Enabling multi-modal conversational interface for clinical imaging. In *Extended abstracts of the 2020 CHI conference on human factors in computing systems*, 2024.

[22] Doug DeCarlo and Anthony Santella. Stylization and abstraction of photographs. *ACM transactions on graphics (TOG)*, 21(3):769–776, 2002.

[23] Yinpeng Dong, Qi-An Fu, Xiao Yang, Tianyu Pang, Hang Su, Zihao Xiao, and Jun Zhu. Benchmarking adversarial robustness on image classification. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 321–331, 2020.

[24] Deng-Ping Fan, Ge-Peng Ji, Ming-Ming Cheng, and Ling Shao. Concealed object detection. *IEEE transactions on pattern analysis and machine intelligence*, 44(10):6024–6042, 2021.

[25] Deng-Ping Fan, Ge-Peng Ji, Guolei Sun, Ming-Ming Cheng, Jianbing Shen, and Ling Shao. Camouflaged object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2777–2787, 2020.

[26] Hao-Shu Fang, Jianhua Sun, Runzhong Wang, Minghao Gou, Yong-Lu Li, and Cewu Lu. Instaboost: Boosting instance segmentation via probability map guided copy-pasting. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 682–691, 2019.

[27] Thomas Fel, Ivan F Rodriguez Rodriguez, Drew Linsley, and Thomas Serre. Harmonizing the object recognition strategies of deep neural networks with humans. *Advances in neural information processing systems*, 35:9432–9446, 2022.

[28] Nicholas T Franklin, Kenneth A Norman, Charan Ranganath, Jeffrey M Zacks, and Samuel J Gershman. Structured event memory: A neuro-symbolic model of event cognition. *Psychological Review*, 127(3):327, 2020.

[29] Karl J Friston, Noor Sajid, David Ricardo Quiroga-Martinez, Thomas Parr, Cathy J Price, and Emma Holmes. Active listening. *Hearing research*, 399:107998, 2021.

[30] R Geirhos, CR Medina Temme, J Rauber, HH Schütt, M Bethge, and FA Wichmann. Generalisation in humans and deep neural networks. In *Thirty-second Annual Conference on Neural Information Processing Systems 2018 (NeurIPS 2018)*, pages 7549–7561. Curran, 2019.

[31] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, 2020.

[32] Robert Geirhos, Kristof Meding, and Felix A Wichmann. Beyond accuracy: quantifying trial-by-trial behaviour of cnns and humans by measuring error consistency. *Advances in Neural Information Processing Systems*, 33:13890–13902, 2020.

[33] Robert Geirhos, Kantharaju Narayanappa, Benjamin Mitzkus, Tizian Thieringer, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. Partial success in closing the gap between human and machine vision. *Advances in Neural Information Processing Systems*, 34:23885–23899, 2021.

[34] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations*, 2018.

[35] Tijl Grootswagers, Amanda K Robinson, and Thomas A Carlson. The representational dynamics of visual objects in rapid serial visual processing streams. *NeuroImage*, 188:668–679, 2019.

[36] Tanmay Gupta and Aniruddha Kembhavi. Visual programming: Compositional visual reasoning without training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14953–14962, 2023.

[37] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.

[38] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.

[39] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15262–15271, 2021.

[40] Katherine L Hermann, Ting Chen, and Simon Kornblith. The origins and prevalence of texture bias in convolutional neural networks. *Advances in neural information processing systems*, 2020.

[41] Hossein Hosseini and Radha Poovendran. Semantic adversarial examples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1614–1619, 2018.

[42] Linmei Hu, Zeyi Liu, Ziwang Zhao, Lei Hou, Liqiang Nie, and Juanzi Li. A survey of knowledge enhanced pre-trained language models. *IEEE Transactions on Knowledge and Data Engineering*, 2023.

[43] Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Adversarial examples are not bugs, they are features. *Advances in neural information processing systems*, 32, 2019.

[44] Georgin Jacob, RT Pramod, Harish Katti, and SP Arun. Qualitative similarities and differences in visual object representations between brains and deep networks. *Nature communications*, 12(1):1–14, 2021.

[45] Priyank Jaini, Kevin Clark, and Robert Geirhos. Intriguing properties of generative classifiers. In *International Conference on Learning Representations*, 2024.

[46] Tilke Judd, Krista Ehinger, Frédo Durand, and Antonio Torralba. Learning to predict where humans look. In *2009 IEEE 12th international conference on computer vision*, pages 2106–2113. IEEE, 2009.

[47] Been Kim, Emily Reif, Martin Wattenberg, Samy Bengio, and Michael C Mozer. Neural networks trained on natural scenes exhibit gestalt closure. *Computational Brain & Behavior*, 4(3):251–263, 2021.

[48] Junkyung Kim, Drew Linsley, Kalpit Thakkar, and Thomas Serre. Disentangling neural mechanisms for perceptual grouping. In *International Conference on Learning Representations*, 2019.

[49] Gabriel Kreiman and Thomas Serre. Beyond the feedforward sweep: feedback computations in the visual cortex. *Annals of the New York Academy of Sciences*, 1464(1):222–241, 2020.

[50] Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *Behavioral and brain sciences*, 40:e253, 2017.

[51] Trung-Nghia Le, Yubo Cao, Tan-Cong Nguyen, Minh-Quan Le, Khanh-Duy Nguyen, Thanh-Toan Do, Minh-Triet Tran, and Tam V Nguyen. Camouflaged instance segmentation in-the-wild: Dataset, method, and benchmark suite. *IEEE Transactions on Image Processing*, 31:287–300, 2021.

[52] Minjong Lee and Dongwoo Kim. Robust evaluation of diffusion-based adversarial purification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 134–144, 2023.

[53] Yanjie Li, Bin Xie, Songtao Guo, Yuanyuan Yang, and Bin Xiao. A survey of robustness and safety of 2d and 3d deep learning models against adversarial attacks. *ACM Computing Surveys*, 56(6):1–37, 2024.

[54] Yingwei Li, Qihang Yu, Mingxing Tan, Jieru Mei, Peng Tang, Wei Shen, Alan Yuille, and Cihang Xie. Shape-texture debiased neural network training. In *International Conference on Learning Representations*, 2021.

[55] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.

[56] Drew Linsley, Alekh K Ashok, Lakshmi N Govindarajan, Rex Liu, and Thomas Serre. Stable and expressive recurrent vision models. *Advances in neural information processing systems*, 2020.

[57] Drew Linsley, Dan Shiebler, Sven Eberhardt, and Thomas Serre. Learning what and where to attend. In *International Conference on Learning Representations*, 2019.

[58] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021.

[59] Bradley C Love. The algorithmic level is the bridge between computation and brain. *Topics in cognitive science*, 7(2):230–242, 2015.

[60] Yunqiu Lv, Jing Zhang, Yuchao Dai, Aixuan Li, Bowen Liu, Nick Barnes, and Deng-Ping Fan. Simultaneously localize, segment and rank the camouflaged objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11591–11601, 2021.

[61] Johannes Mehrer, Courtney J Spoerer, Nikolaus Kriegeskorte, and Tim C Kietzmann. Individual differences among deep neural network models. *Nature communications*, 11(1):5725, 2020.

[62] Yalda Mohsenzadeh, Sheng Qin, Radoslaw M Cichy, and Dimitrios Pantazis. Ultra-rapid serial visual presentation reveals dynamics of feedforward and feedback processes in the ventral visual pathway. *Elife*, 7:e36329, 2018.

[63] Chaithanya Kumar Mummadi, Ranjitha Subramaniam, Robin Hutmacher, Julien Vitay, Volker Fischer, and Jan Hendrik Metzen. Does enhanced shape bias improve neural network robustness to common corruptions? In *International Conference on Learning Representations*, 2021.

[64] Micah M Murray, Deirdre M Foxe, Daniel C Javitt, and John J Foxe. Setting boundaries: brain dynamics of modal and amodal illusory shape completion in humans. *Journal of Neuroscience*, 24(31):6898–6903, 2004.

[65] Lukas Muttenthaler, Lorenz Linhardt, Jonas Dippel, Robert A Vandermeulen, Katherine Hermann, Andrew Lampinen, and Simon Kornblith. Improving neural network representations using human similarity judgments. *Advances in Neural Information Processing Systems*, 36, 2024.

[66] Bence Nanay. The importance of amodal completion in everyday perception. *i-Perception*, 9(4), 2018.

[67] Alice J O'Toole and Carlos D Castillo. Face recognition by humans and machines: three fundamental advances from deep learning. *Annual Review of Vision Science*, 7:543–570, 2021.

[68] Giuseppe Papari and Nicolai Petkov. Adaptive pseudo dilation for gestalt edge grouping and contour detection. *IEEE transactions on image processing*, 17(10):1950–1962, 2008.

[69] Connor J Parde, Virginia E Strehle, Vivekjyoti Banerjee, Ying Hu, Jacqueline G Cavazos, Carlos D Castillo, and Alice J O'Toole. Twin identification over viewpoint change: A deep convolutional neural network surpasses humans. *ACM Transactions on Applied Perception*, 20(3):1–15, 2023.

[70] David Peebles and Richard P Cooper. Thirty years after marr's vision: Levels of analysis in cognitive science, 2015.

[71] James R Pomerantz and Mary C Portillo. Grouping and emergent features in vision: toward a theory of basic gestalts. *Journal of Experimental Psychology: Human Perception and Performance*, 37(5):1331, 2011.

[72] James R Pomerantz, Lawrence C Sager, and Robert J Stoever. Perception of wholes and of their component parts: some configural superiority effects. *Journal of Experimental Psychology: Human Perception and Performance*, 3(3):422, 1977.

[73] Siyuan Qiao, Liang-Chieh Chen, and Alan Yuille. Detectors: Detecting objects with recursive feature pyramid and switchable atrous convolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10213–10224, 2021.

[74] Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? *Advances in Neural Information Processing Systems*, 34, 2021.

[75] Leonardo Rossi, Akbar Karimi, and Andrea Prati. A novel region of interest extraction layer for instance segmentation. In *2020 25th international conference on pattern recognition (ICPR)*, pages 2203–2209. IEEE, 2021.

[76] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.

[77] Andrew Saxe, Stephanie Nelli, and Christopher Summerfield. If deep learning is the answer, what is the question? *Nature Reviews Neuroscience*, 22(1):55–67, 2021.

[78] Thomas Serre. Deep learning: the good, the bad, and the ugly. *Annual review of vision science*, 5(1):399–426, 2019.

[79] Zhenwei Shao, Zhou Yu, Meng Wang, and Jun Yu. Prompting large language models with answer heuristics for knowledge-based visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14974–14983, 2023.

[80] I-Chao Shen and Wen-Huang Cheng. Gestalt rule feature points. *IEEE Transactions on Multimedia*, 17(4):526–537, 2015.

[81] Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, et al. Large language models encode clinical knowledge. *Nature*, 620(7972):172–180, 2023.

[82] Przemysław Skurowski, Hassan Abdulameer, J Błaszczyk, Tomasz Depta, Adam Kornacki, and P Kozieł. Animal camouflage analysis: Chameleon database. *Unpublished manuscript*, 2(6):7, 2018.

[83] Ajay Subramanian, Elena Sizikova, Najib Majaj, and Denis Pelli. Spatial-frequency channels, shape bias, and adversarial robustness. *Advances in Neural Information Processing Systems*, 36, 2024.

[84] Mingjie Sun, Zichao Li, Chaowei Xiao, Haonan Qiu, Bhavya Kailkhura, Mingyan Liu, and Bo Li. Can shape structure features improve model robustness under diverse adversarial settings? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7526–7535, 2021.

[85] Hanlin Tang, Martin Schrimpf, William Lotter, Charlotte Moerman, Ana Paredes, Josue Ortega Caro, Walter Hardesty, David Cox, and Gabriel Kreiman. Recurrent computations for visual pattern completion. *Proceedings of the National Academy of Sciences*, 115(35):8835–8840, 2018.

[86] Simon Thorpe, Denis Fize, and Catherine Marlot. Speed of processing in the human visual system. *nature*, 381(6582):520–522, 1996.

[87] Nikhil Kumar Tomar, Debesh Jha, Michael A Riegler, Håvard D Johansen, Dag Johansen, Jens Rittscher, Pål Halvorsen, and Sharib Ali. Fanet: A feedback attention network for improved biomedical image segmentation. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.

[88] Alexander Tscshantz, Beren Millidge, Anil K Seth, and Christopher L Buckley. Hybrid predictive coding: Inferring, fast and slow. *PLOS Computational Biology*, 19(8):e1011280, 2023.

[89] Shikhar Tuli, Ishita Dasgupta, Erin Grant, and Tom Griffiths. Are convolutional neural networks or transformers more like human vision? In *Proceedings of the Annual Meeting of the Cognitive Science Society*, 2021.

[90] Ruben S van Bergen and Nikolaus Kriegeskorte. Going in circles is the way forward: the role of recurrence in visual inference. *Current Opinion in Neurobiology*, 65:176–193, 2020.

[91] Eline Van Geert and Johan Wagemans. Prägnanz in visual perception. *Psychonomic Bulletin & Review*, pages 1–27, 2023.

[92] Rufin VanRullen. The power of the feed-forward sweep. *Advances in Cognitive Psychology*, 3(1-2):167, 2007.

[93] Lukas Vogelsang, Sharon Gilad-Gutnick, Evan Ehrenberg, Albert Yonas, Sidney Diamond, Richard Held, and Pawan Sinha. Potential downside of high initial visual acuity. *Proceedings of the National Academy of Sciences*, 115(44):11333–11338, 2018.

[94] Johan Wagemans, James H Elder, Michael Kubovy, Stephen E Palmer, Mary A Peterson, Manish Singh, and Rüdiger von der Heydt. A century of gestalt psychology in visual perception: I. perceptual grouping and figure–ground organization. *Psychological bulletin*, 138(6):1172, 2012.

[95] Johan Wagemans, Jacob Feldman, Sergei Gepshtein, Ruth Kimchi, James R Pomerantz, Peter A Van der Helm, and Cees Van Leeuwen. A century of gestalt psychology in visual perception: Ii. conceptual and theoretical foundations. *Psychological bulletin*, 138(6):1218, 2012.

[96] Dirk B Walther, Barry Chai, Eamon Caddigan, Diane M Beck, and Li Fei-Fei. Simple line drawings suffice for functional mri decoding of natural scene categories. *Proceedings of the National Academy of Sciences*, 108(23):9661–9666, 2011.

[97] Wei Wang, Kaicheng Yu, Joachim Hugonot, Pascal Fua, and Mathieu Salzmann. Recurrent u-net for resource-constrained segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2142–2151, 2019.

[98] Max Wertheimer. Experimentelle studien uber das sehen von bewegung. *Zeitschrift fur psychologie*, 61, 1912.

[99] Yufei Xu, Qiming ZHANG, Jing Zhang, and Dacheng Tao. Vitae: Vision transformer advanced by exploring intrinsic inductive bias. In *Advances in Neural Information Processing Systems*, 2021.

[100] Zhenlin Xu, Deyi Liu, Junlin Yang, Colin Raffel, and Marc Niethammer. Robust and generalizable visual representation learning via random convolutions. In *International Conference on Learning Representations*, 2021.

[101] Charles T Zahn. Graph-theoretical methods for detecting and describing gestalt clusters. *IEEE Transactions on computers*, 100(1):68–86, 1971.