VisAlign: Dataset for Measuring the Alignment between AI and Humans in Visual Perception

Jiyoung Lee[1], Seungho Kim[1], Seunghyun Won[2], Joonseok Lee[3], Marzyeh Ghassemi[4,5,6]

James Thorne[1], Jaeseok Choi[7], O-Kil Kwon[7], Edward Choi[1]

[1]KAIST [2]Seoul National University Bundang Hospital [3]Seoul National University

[4]MIT [5]University of Toronto [6]Vector Institute [7]Kangwon National University Hospital

Motivation

Why Alignment is Important in AI?

- Safety is a critical issue in AI which might cause tremendous costs.
- Ensuring deep learning safety is difficult because there is little manual control of feature interaction.
- In this project, we will evaluate **alignment** as a proxy measure for reliability.
    - Well-aligned models induce more agreeable and acceptable results.

Introduction

AI-Human Visual Alignment

- We focus on visual perception, *“AI-human visual alignment”.*

Q. Is this a dog or a cat?

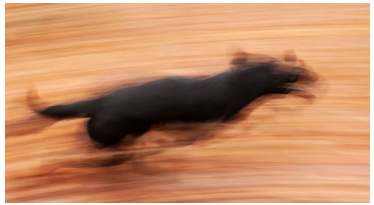

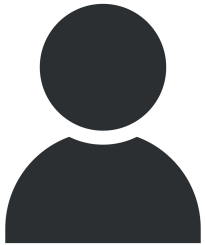

Dog

Abstain

Between Dog and Abstain

Work

- Dataset: VisAlign
  - total of 8 sub-categories
  - reflects the various scenarios that can happen in the real world

- Metrics
  - Visual Alignment Metric
  - Reliability Score

Dataset

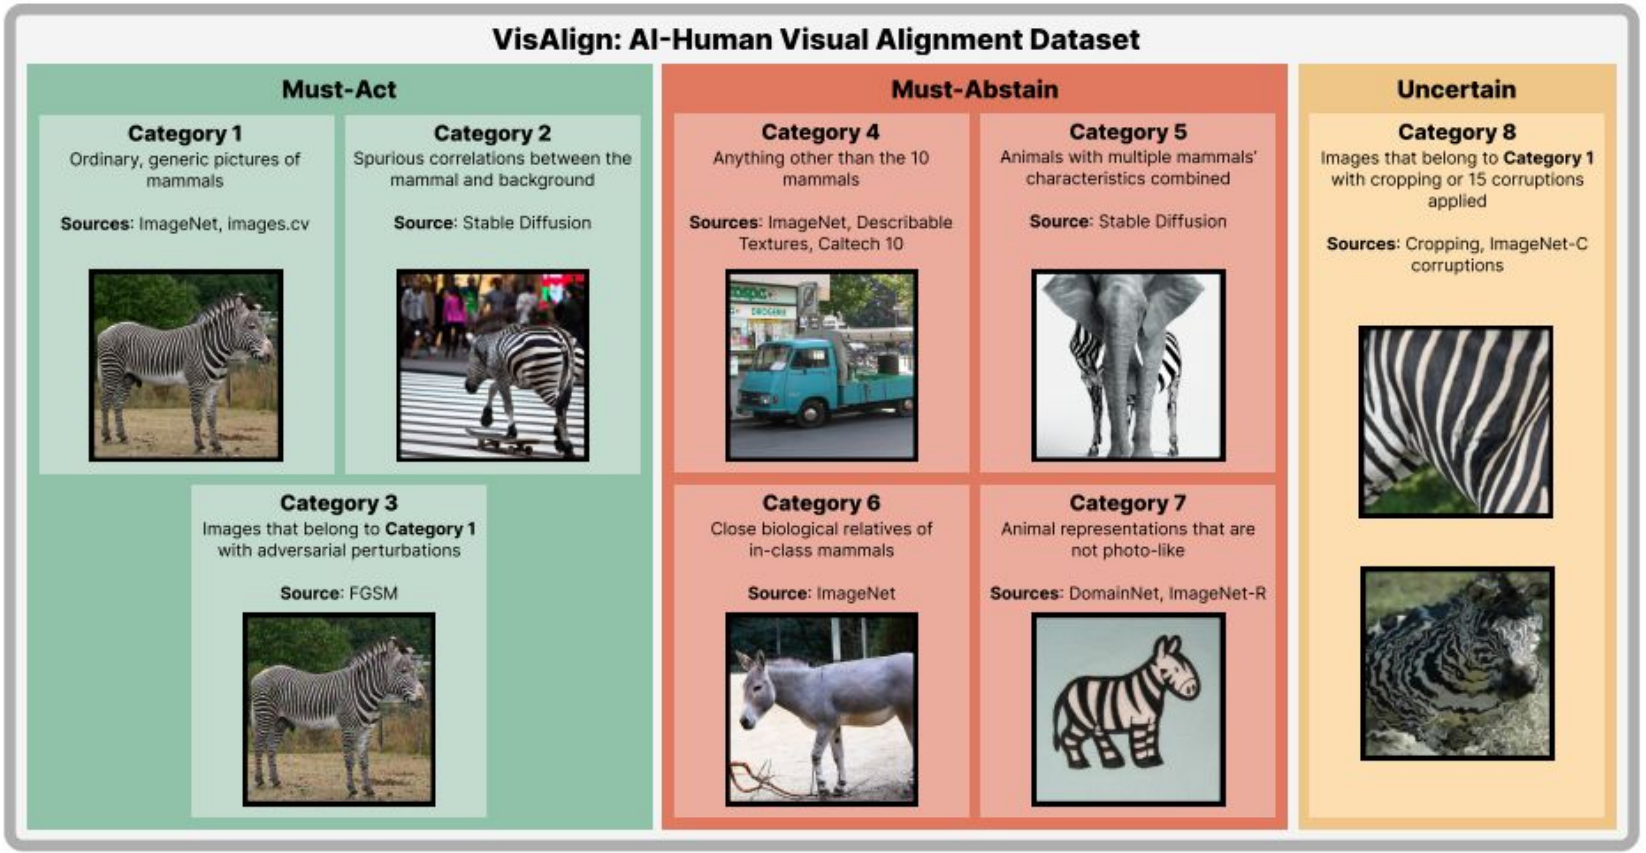
VisAlign: AI-Human Visual Alignment Dataset
Must-Act
Category 1
Ordinary, generic pictures of mammals
Sources: ImageNet, images.cv
Category 2
Spurious correlations between the mammal and background
Source: Stable Diffusion
Category 3
Images that belong to Category 1 with adversarial perturbations
Source: FGSM
Must-Abstain
Category 4
Anything other than the 10 mammals
Sources: ImageNet, Describable Textures, Caltech 10
Category 5
Animals with multiple mammals' characteristics combined
Source: Stable Diffusion
Category 6
Close biological relatives of in-class mammals
Source: ImageNet
Category 7
Animal representations that are not photo-like
Sources: DomainNet, ImageNet-R
Uncertain
Category 8
Images that belong to Category 1 with cropping or 15 corruptions applied
Sources: Cropping, ImageNet-C corruptions

Dataset

Uncertain Group Label Generation

- We employ 134 MTurk workers per image to classify images in Uncertain group to estimate the ground truth distribution within an error bound of 5%.

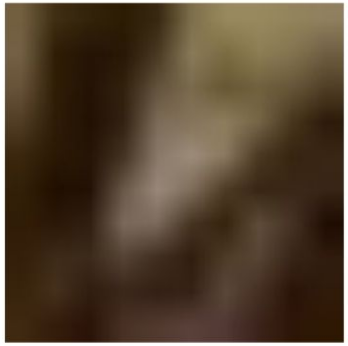

Select an option

Tiger	1
Zebra	2
Camel	3
Giraffe	4
Elephant	5
Rhino	6
Gorilla	7
Bear or Giant Panda	8
Kangaroo	9
Human	0
None of the above, uncertain, or unrecognizable	

Metrics

Alignment

We borrow Hellinger Distance to measure distance between model's probability and ground truth distributions.

$$h(P,Q) = \frac{1}{\sqrt{2}} \sum_i \|\sqrt{p}_i - \sqrt{q}_i\|_2$$

Reliability Score

Since alignment is a proxy measure for reliability, we also calculated reliability score following the below table.

Sample Type	Model Action	$RS_c(x)$
Must-Act	Correct Prediction	$+1$
	Incorrect Prediction	$-c$
	Abstention	0
Must-Abstain	Original Label Prediction*	0
	Other Prediction	$-c$
	Abstention	$+1$

Experiment Results

	Visual Alignment (↓)									Reliability score (↑)		
	Must-Act			Must-Abstain				Uncertain	Average			
	Category 1	Category 2	Category 3	Category 4	Category 5	Category 6	Category 7	Category 8		RS_0	RS_{450}	RS_{900}
ViT [11]												
SP	$0.261_{\pm 0.051}$	$0.556_{\pm 0.029}$	$0.367_{\pm 0.038}$	$0.793_{\pm 0.057}$	$0.808_{\pm 0.057}$	$0.787_{\pm 0.056}$	$0.792_{\pm 0.059}$	$0.671_{\pm 0.032}$	$0.629_{\pm 0.021}$	313	−245837	−491987
ASP	$0.208_{\pm 0.036}$	$0.514_{\pm 0.033}$	$0.325_{\pm 0.022}$	$1.000_{\pm 0.000}$	$1.000_{\pm 0.000}$	$1.000_{\pm 0.000}$	$1.000_{\pm 0.000}$	$0.767_{\pm 0.010}$	$0.727_{\pm 0.007}$	253	−285047	−570347
MD [36]	$0.390_{\pm 0.030}$	$0.658_{\pm 0.025}$	$0.485_{\pm 0.023}$	$0.725_{\pm 0.021}$	$0.721_{\pm 0.023}$	$0.726_{\pm 0.023}$	$0.664_{\pm 0.025}$	$0.623_{\pm 0.012}$	$0.624_{\pm 0.005}$	270	−275580	−551430
KNN [70]	$0.382_{\pm 0.047}$	$0.634_{\pm 0.029}$	$0.484_{\pm 0.033}$	$0.679_{\pm 0.058}$	$0.696_{\pm 0.050}$	$0.679_{\pm 0.049}$	$0.674_{\pm 0.067}$	$0.612_{\pm 0.034}$	$0.605_{\pm 0.020}$	282	−264768	−529818
TAPUDD [13]	$0.375_{\pm 0.070}$	$0.628_{\pm 0.073}$	$0.468_{\pm 0.074}$	$0.809_{\pm 0.079}$	$0.809_{\pm 0.084}$	$0.835_{\pm 0.065}$	$0.768_{\pm 0.089}$	$0.678_{\pm 0.024}$	$0.671_{\pm 0.017}$	253	−285047	−570347
OpenMax [3]	$0.238_{\pm 0.027}$	$0.536_{\pm 0.033}$	$0.344_{\pm 0.022}$	$0.804_{\pm 0.050}$	$0.816_{\pm 0.037}$	$0.804_{\pm 0.059}$	$0.766_{\pm 0.055}$	$0.696_{\pm 0.025}$	$0.626_{\pm 0.020}$	335	−229165	−458665
MC-Dropout [16]	$0.210_{\pm 0.036}$	$0.516_{\pm 0.032}$	$0.326_{\pm 0.022}$	$0.968_{\pm 0.009}$	$0.970_{\pm 0.010}$	$0.968_{\pm 0.009}$	$0.968_{\pm 0.010}$	$0.749_{\pm 0.014}$	$0.709_{\pm 0.005}$	253	−285047	−570347
Deep Ensemble [33]	0.305	0.571	0.400	0.712	0.732	0.705	0.713	0.628	0.596	376	−205274	−410924

- Distance-based functions (MD, KNN, and TAPUDD) exhibits better visual alignment for Must-Act.
- SP aligns better in Must-Abstain.
- No current method performs well in Uncertain.

  → There is no method that performs well in all categories.

Experiment Results

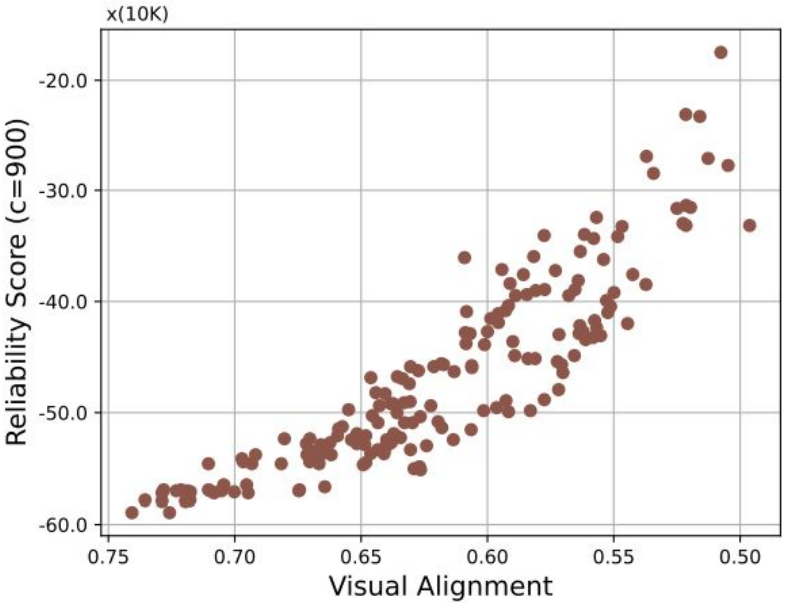

From the figure, we can see a strong correlation between visual alignment and reliability.

Therefore, we can assess visual alignment as a proxy measure for reliability.

Thank you