Seeing Faces in Things: A Model and Dataset for Pareidolia

ECCV 2024







Mark Hamilton[1,2] · Simon Stent[3] · Vasha DuTell[1] · Anne Harrington[1] · Jennifer Corbett[1] · Ruth Rosentholtz[4] · William T. Freeman[1]

1 MIT · 2 Microsoft · 3 TOYOTA RESEARCH INSTITUTE · 4 NVIDIA

Face Pareidolia

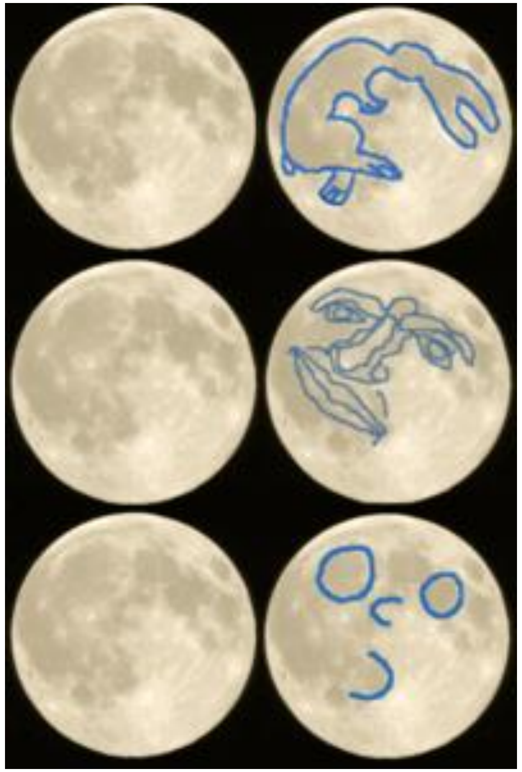

Our Contributions

The Faces in Things Dataset

Animal Faces → Pareidolia

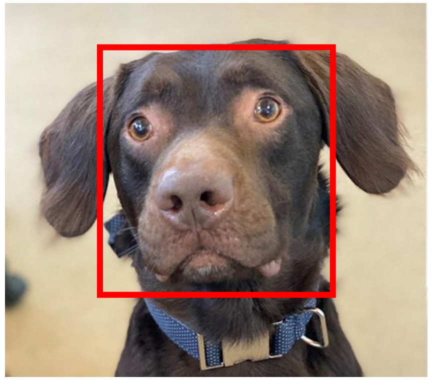

Mathematics of Pareidolia

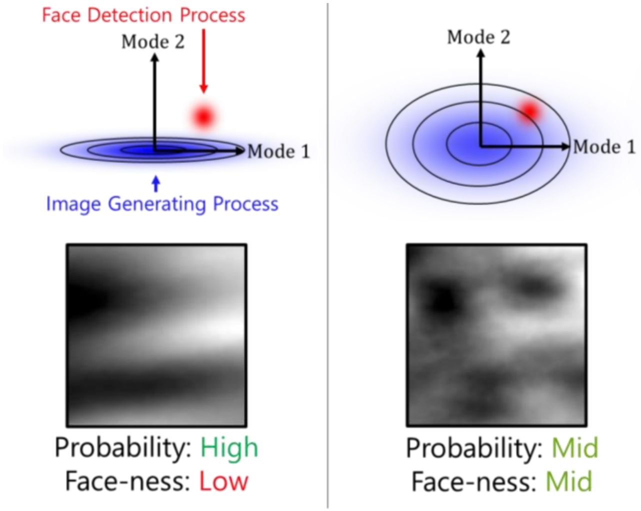

Human and Machine Verification of Theory

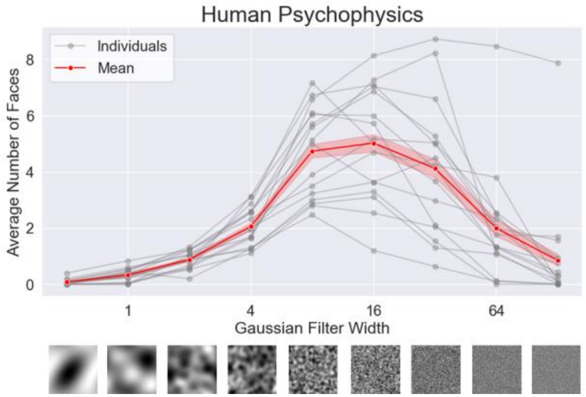

Our Contributions

The Faces in Things Dataset

Animal Faces → Pareidolia

Mathematics of Pareidolia

Human and Machine Verification of Theory

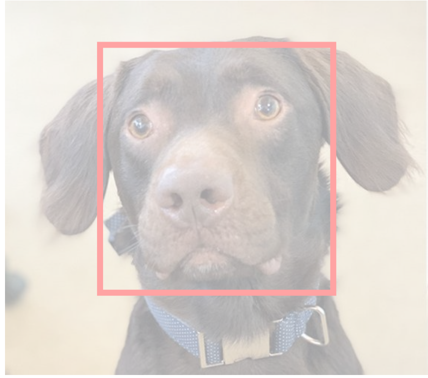

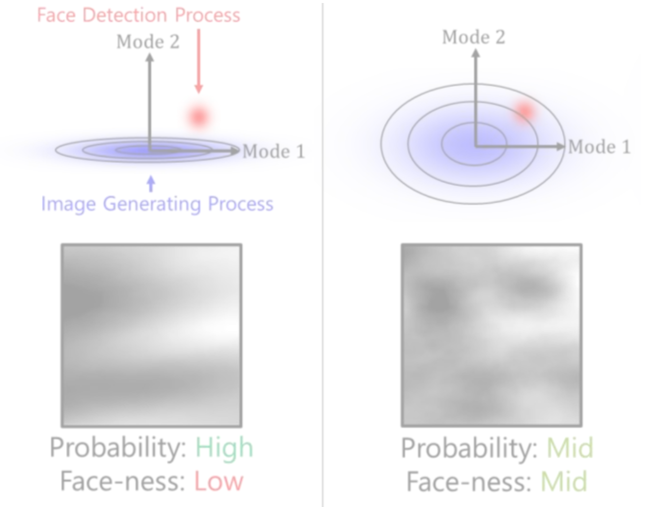

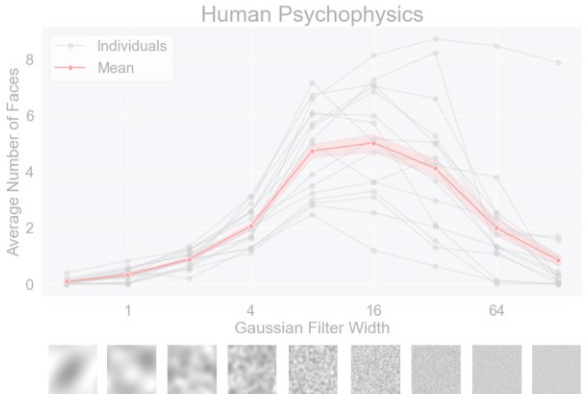

Faces in Things Dataset

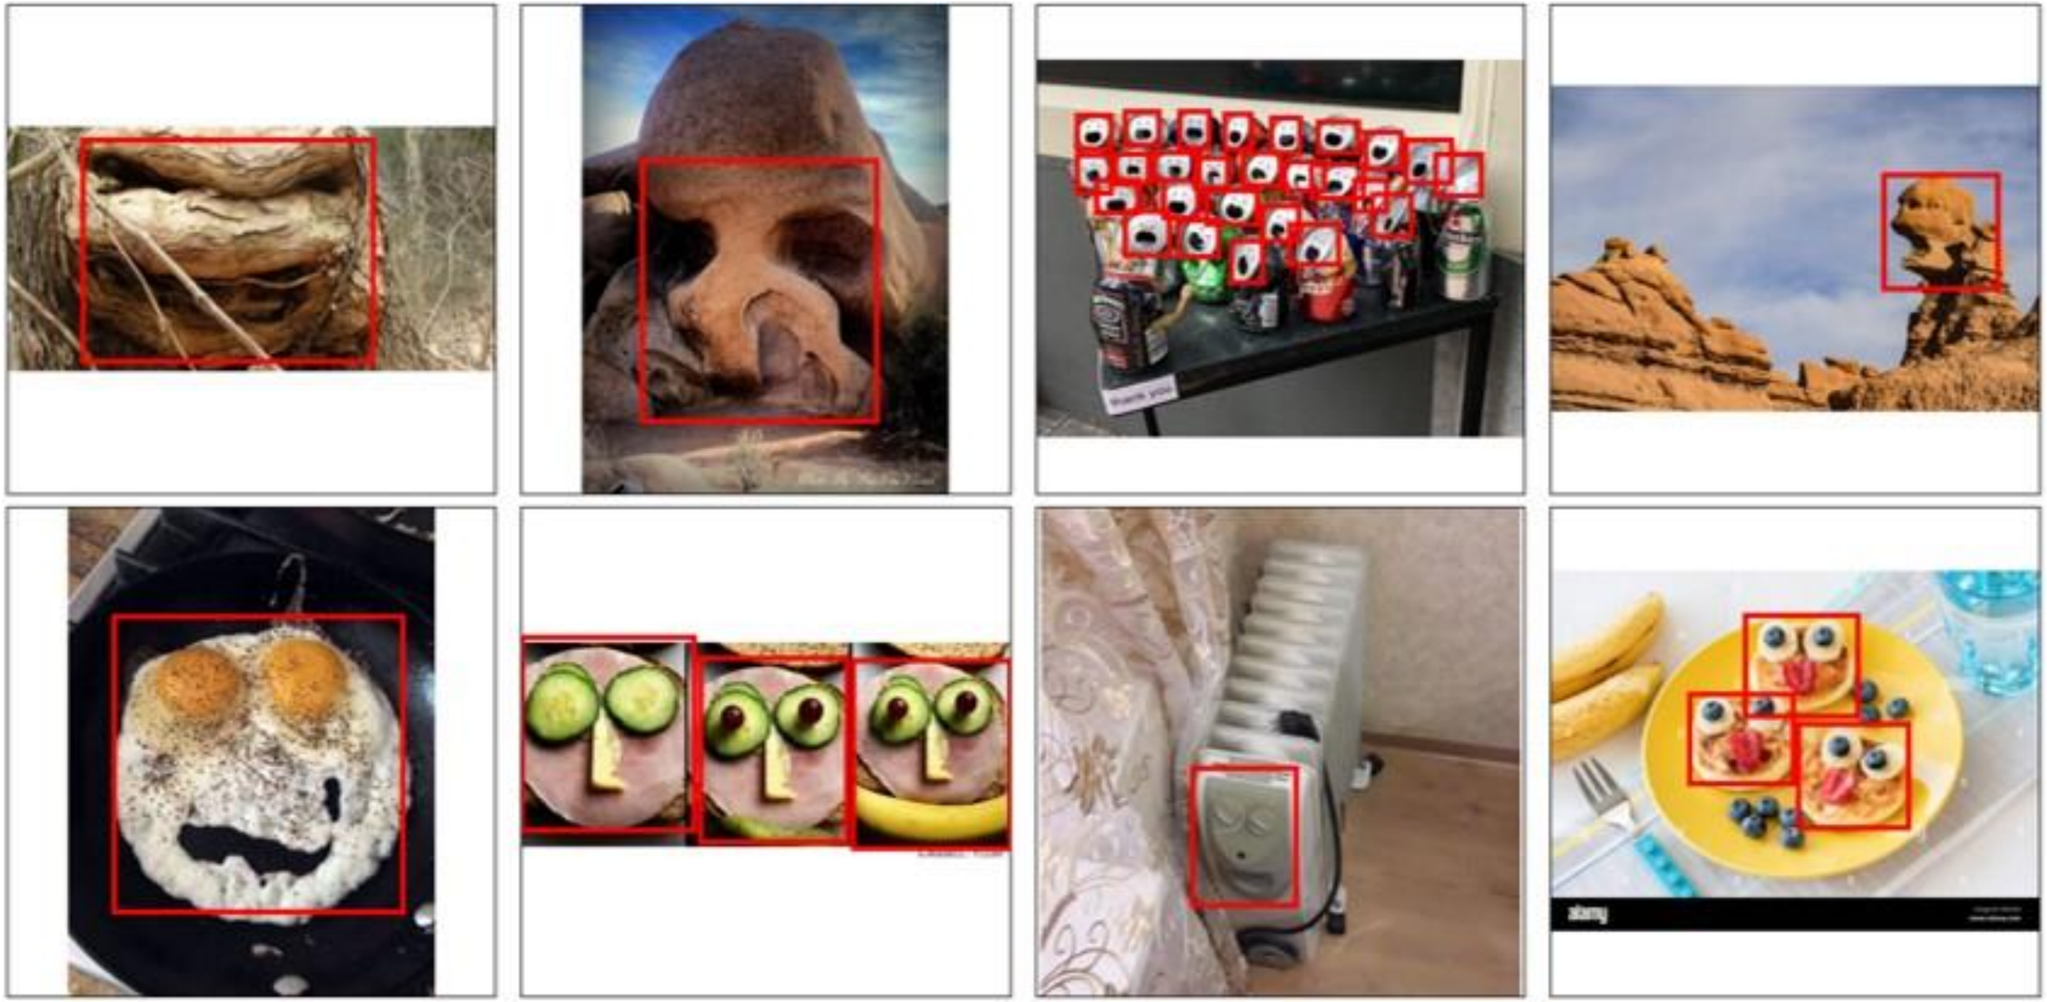

Faces in Things Dataset

- 5,000 Images
- Human Labeled Bounding Boxes

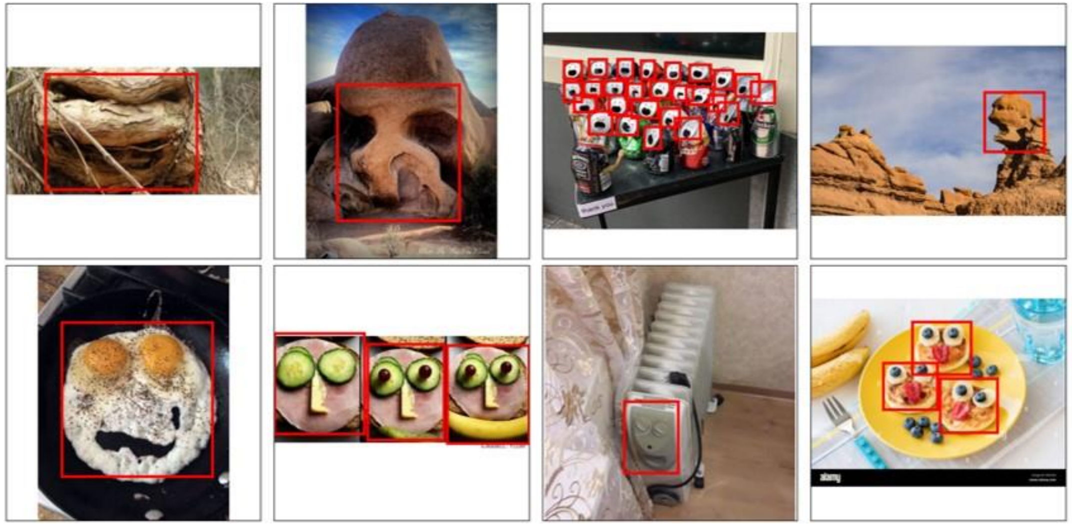

Average Faces

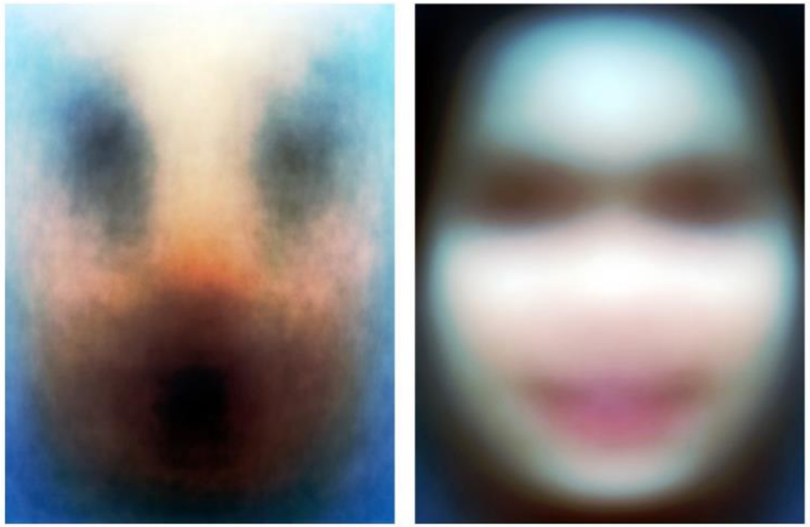

Pareidolic Human

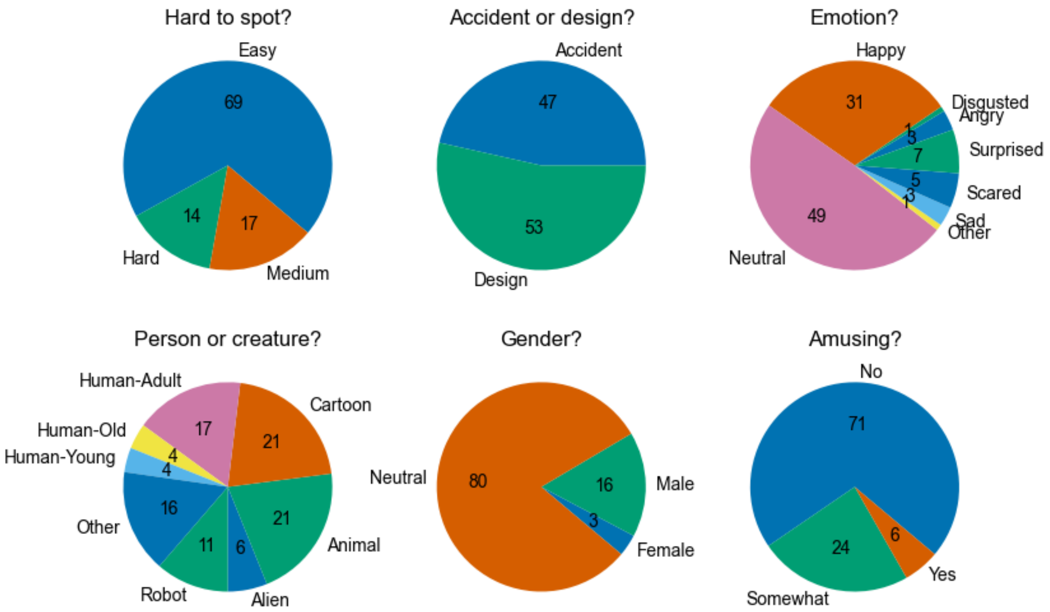

Do SOTA face detectors experience pareidolia?

Retina Face

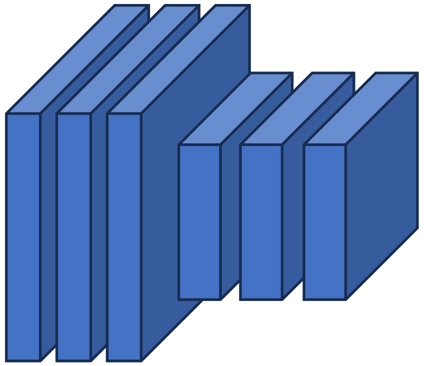

Do SOTA face detectors experience pareidolia?

Faces in Things Dataset

Retina Face

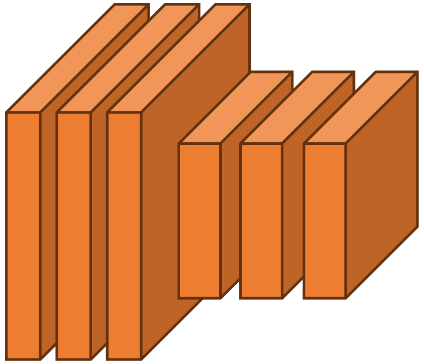

9% AP

36% AP

When Trained on Pareidolic Images

Measuring the Effect of Training Data Intervention

Control Group Dataset

Retina Face

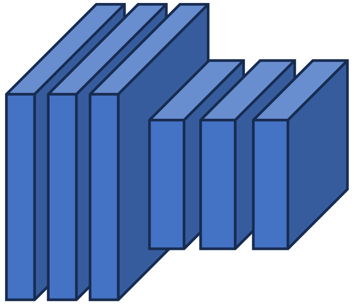

Treatment: Sobel Filtering

Retina Face

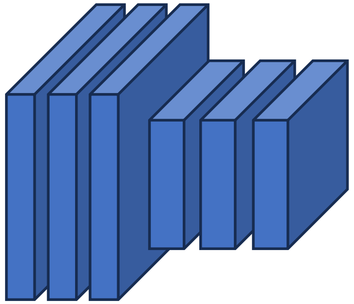

Treatment: Animal Faces

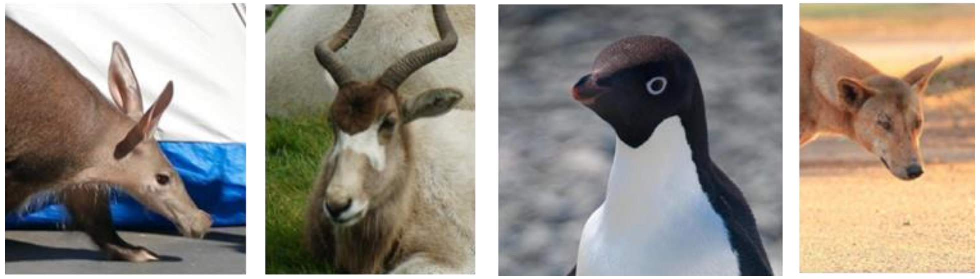

Retina Face

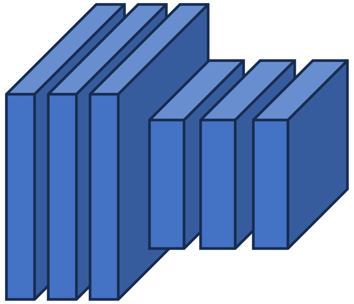

Measuring the Effect of Training Data Intervention

Faces in Things Dataset

Control Retina Face

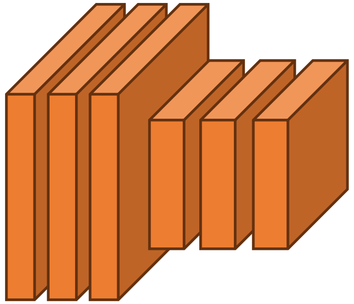

↑ 0%

Faces in Things Dataset

Retina Face + Sobel

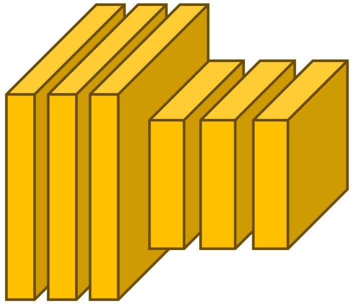

↑ 21%

Faces in Things Dataset

Retina Face + Animal

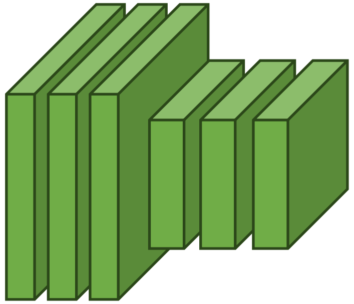

↑ 92%

Measuring the Effect of Training Data Intervention

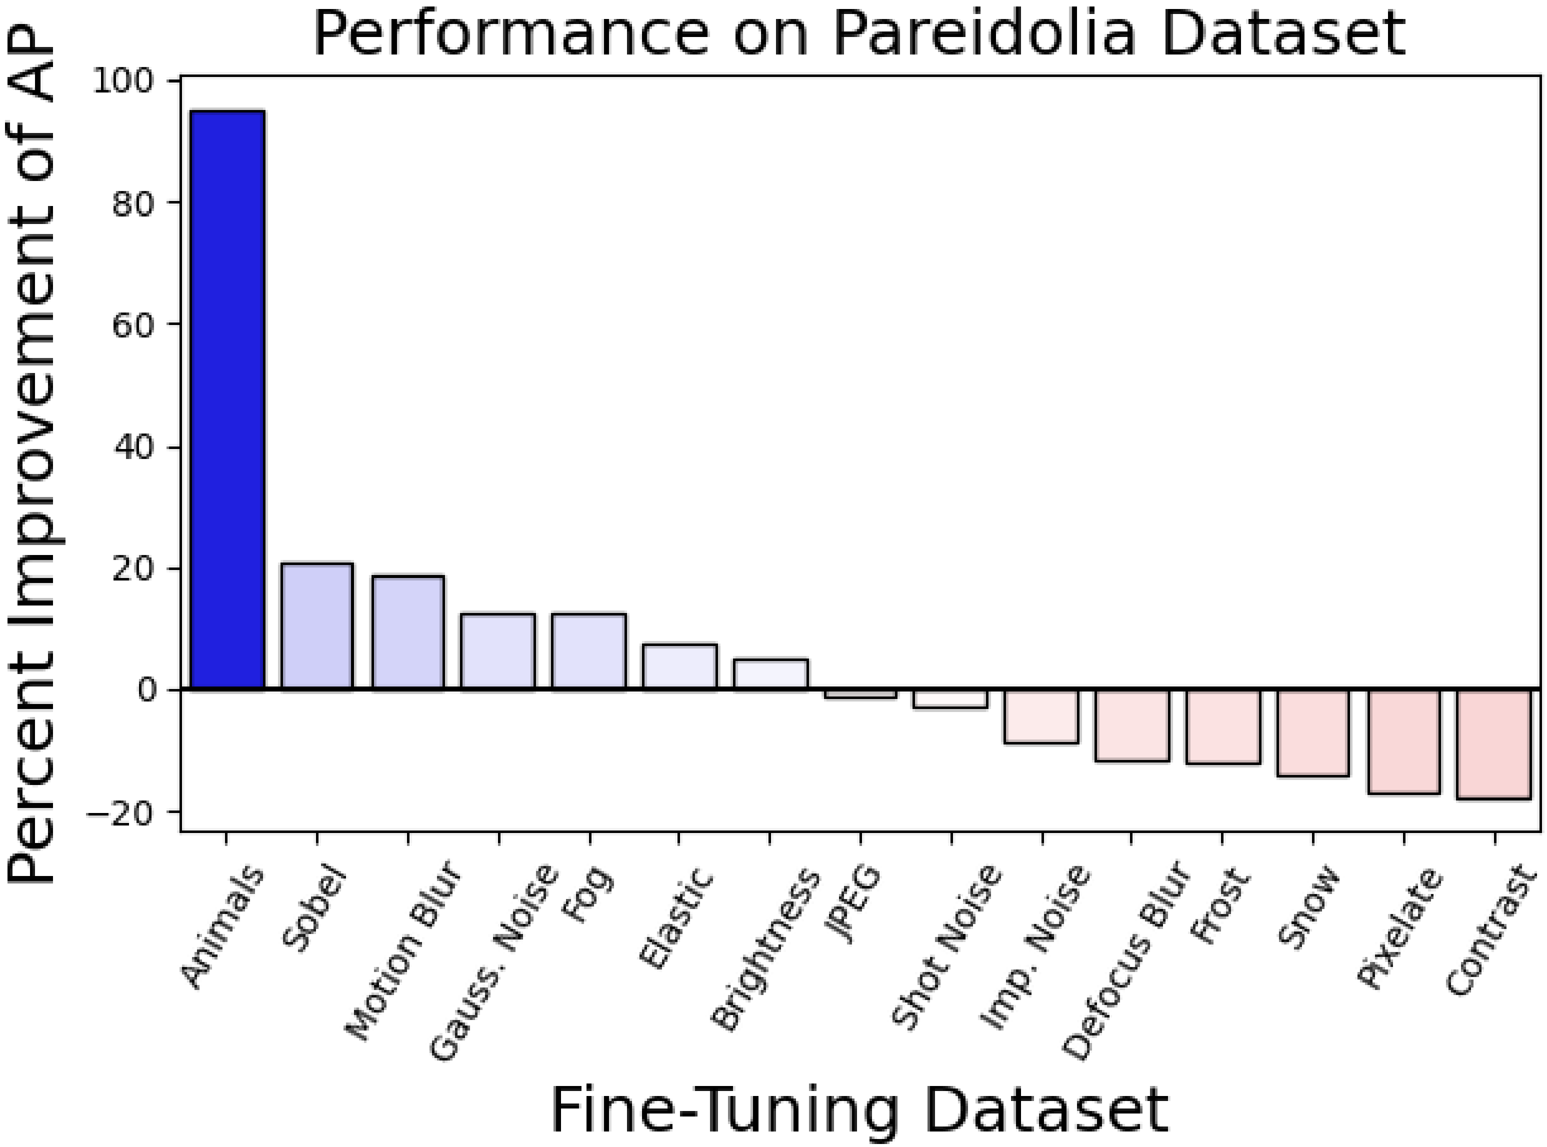

Trained Networks Group Animal and Pareidolic Faces Together

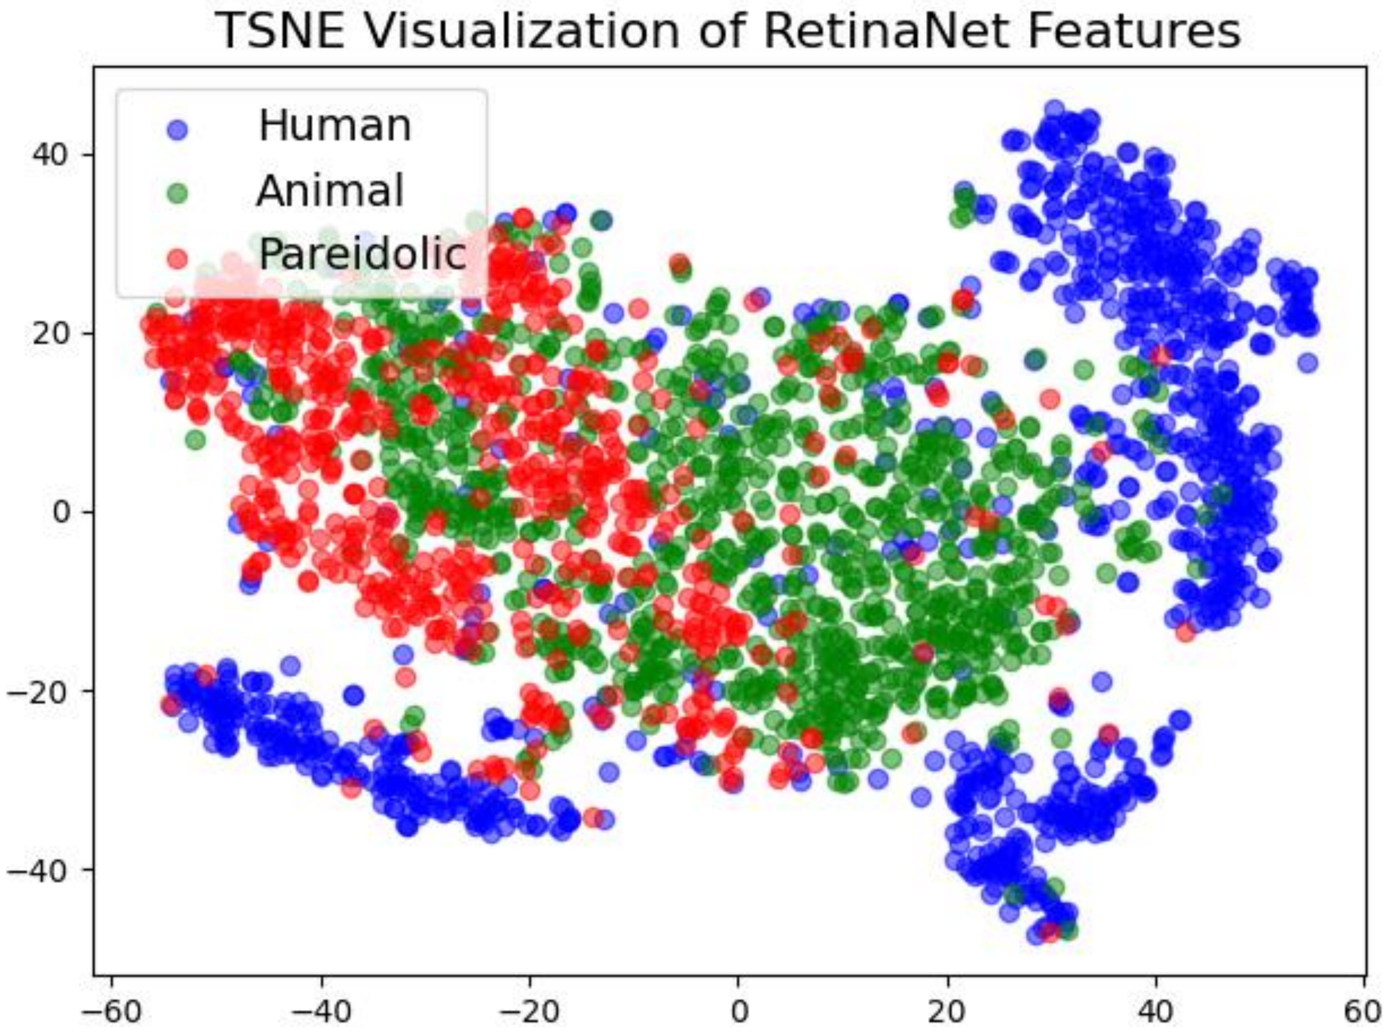

Finding Pareidolic Doppelgangers

Animal Face

Closest Pareidolic Face

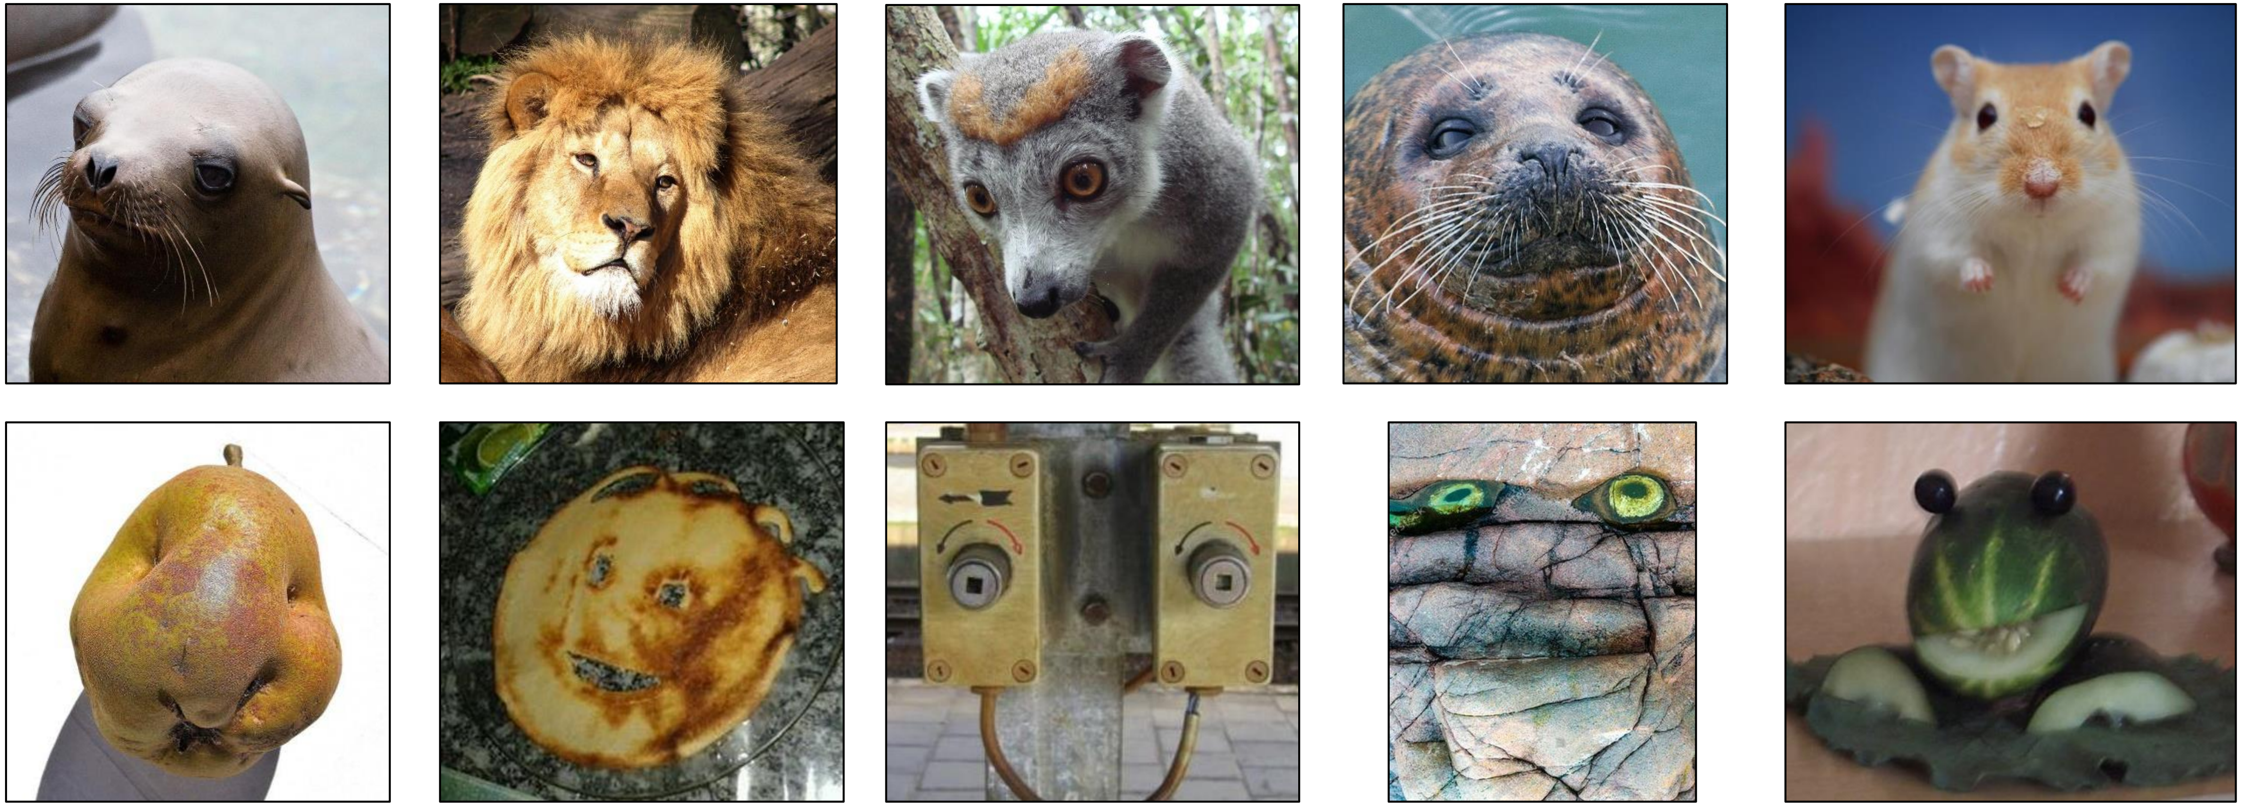

A Simple Mathematical Model of Pareidolia

"What should the frequency of my image generator be to maximize pareidolia?"

Low Frequency

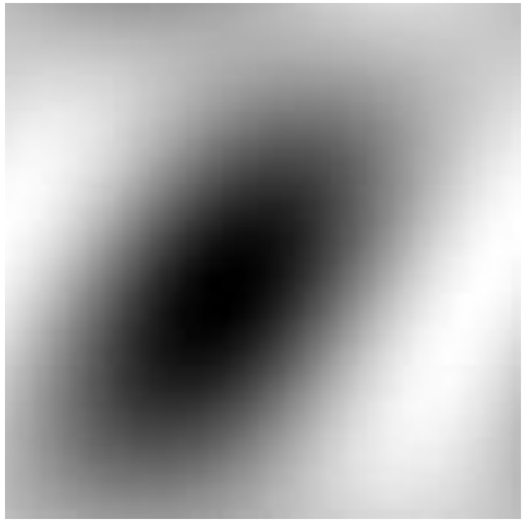

Mid Frequency

High Frequency

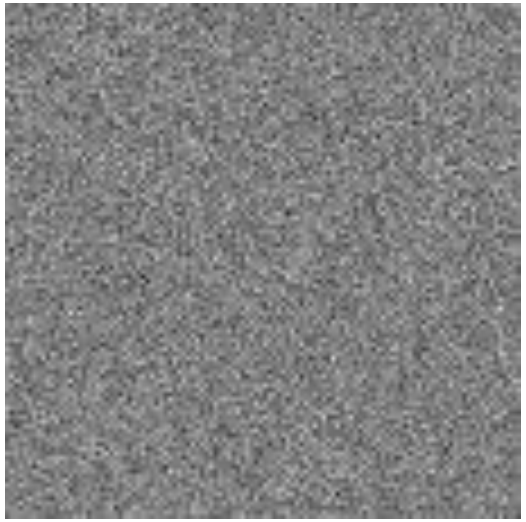

A Simple Mathematical Model of Pareidolia

"What should the frequency of my image generator be to maximize pareidolia?"

Low Frequency

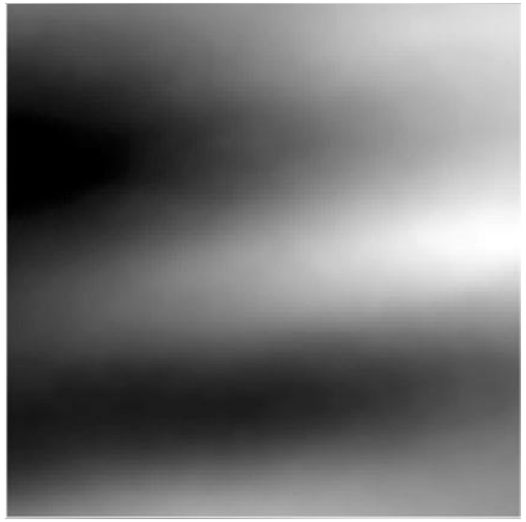

Not rich enough to generate faces!

Mid Frequency

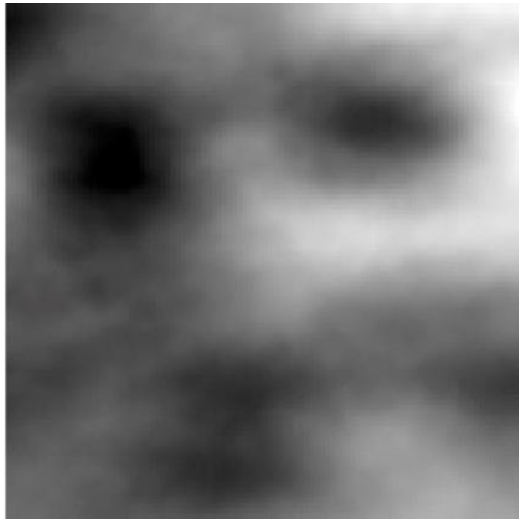

Rich enough to form faces, likely enough to happen.

High Frequency

Very rich distribution, but faces are very unlikely!

Intuition behind the 🐪

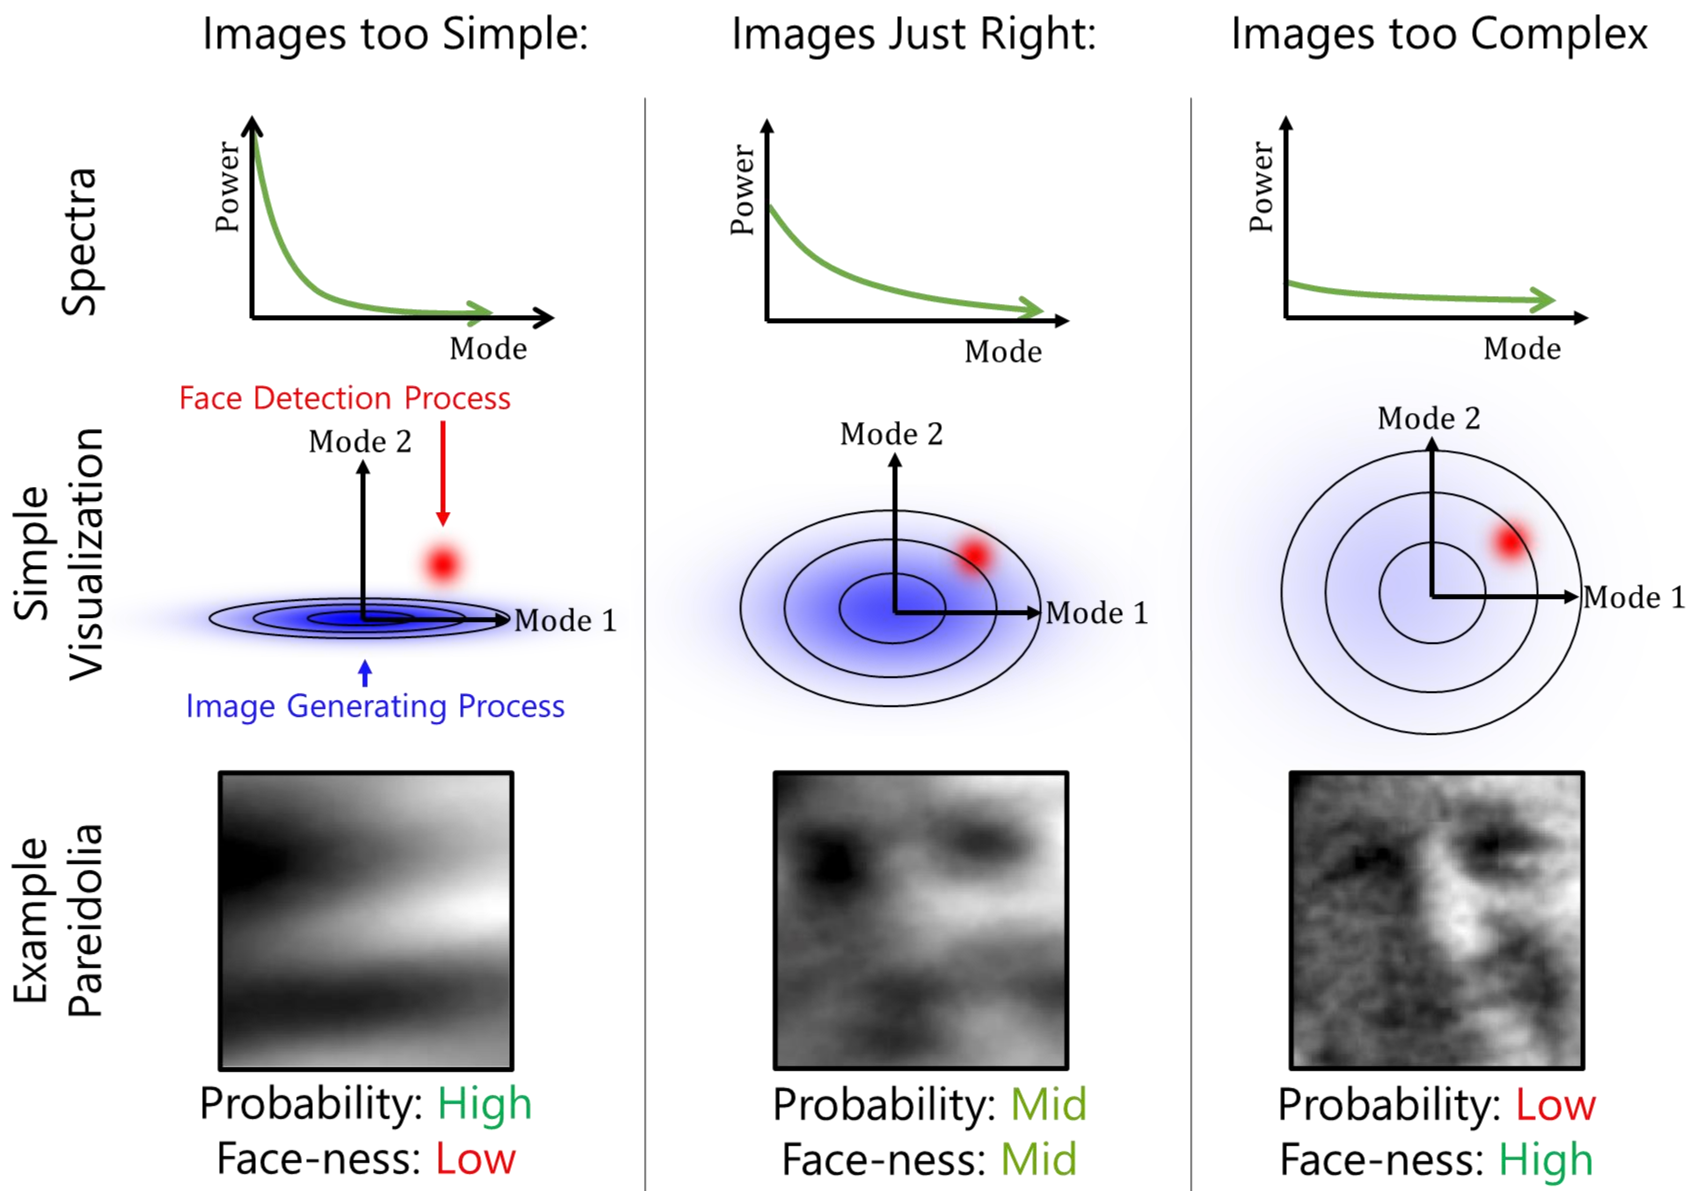

A Simple Mathematical Model of Pareidolia

Too Simple | Just Right | Too Complex

What the Model Says:

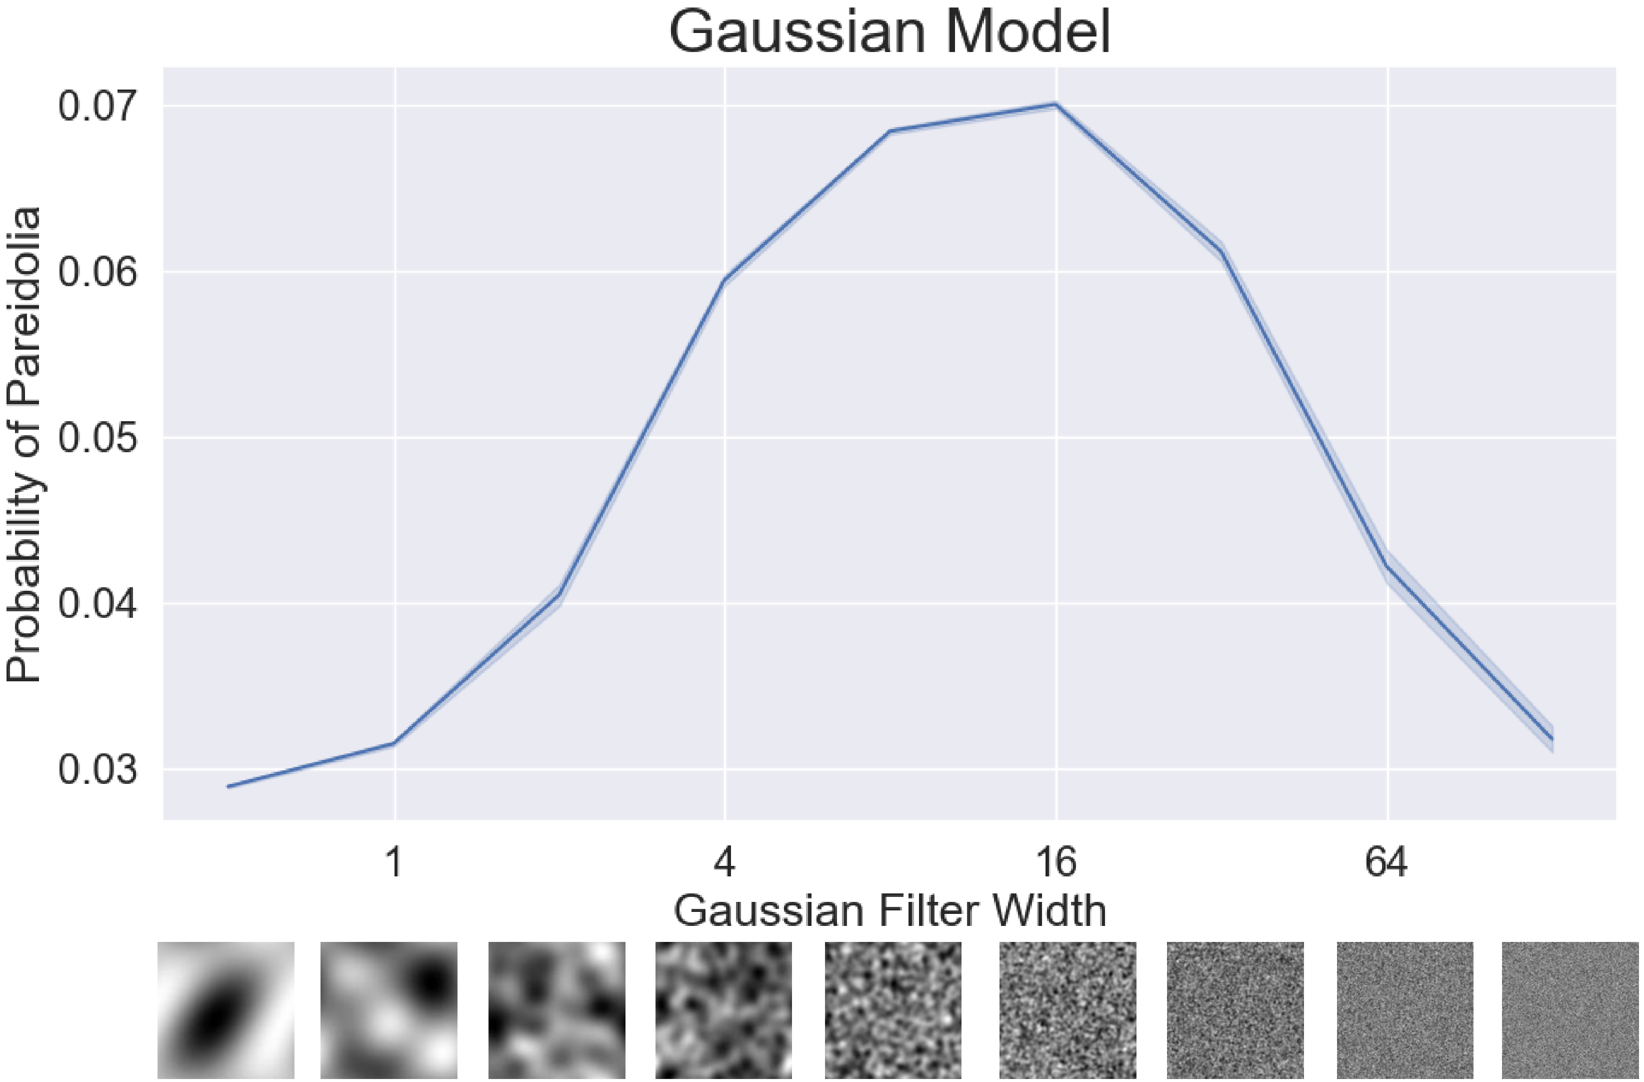

What the Data Says:

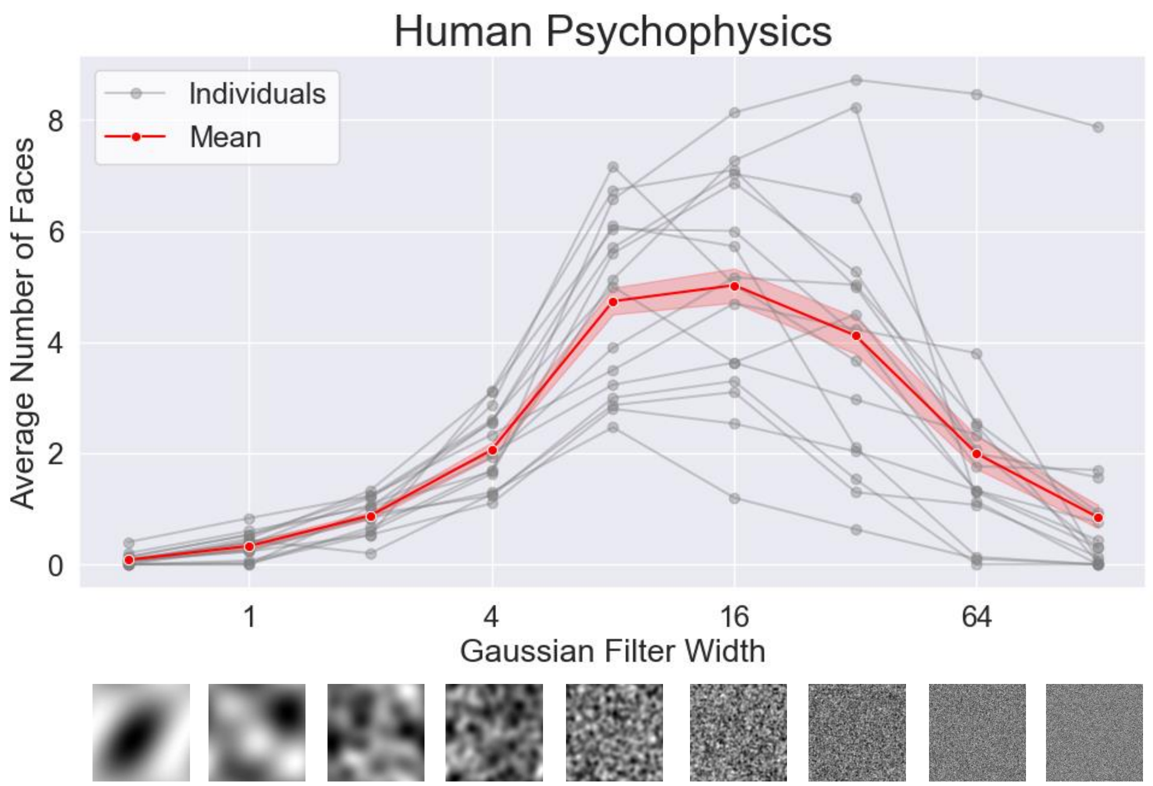

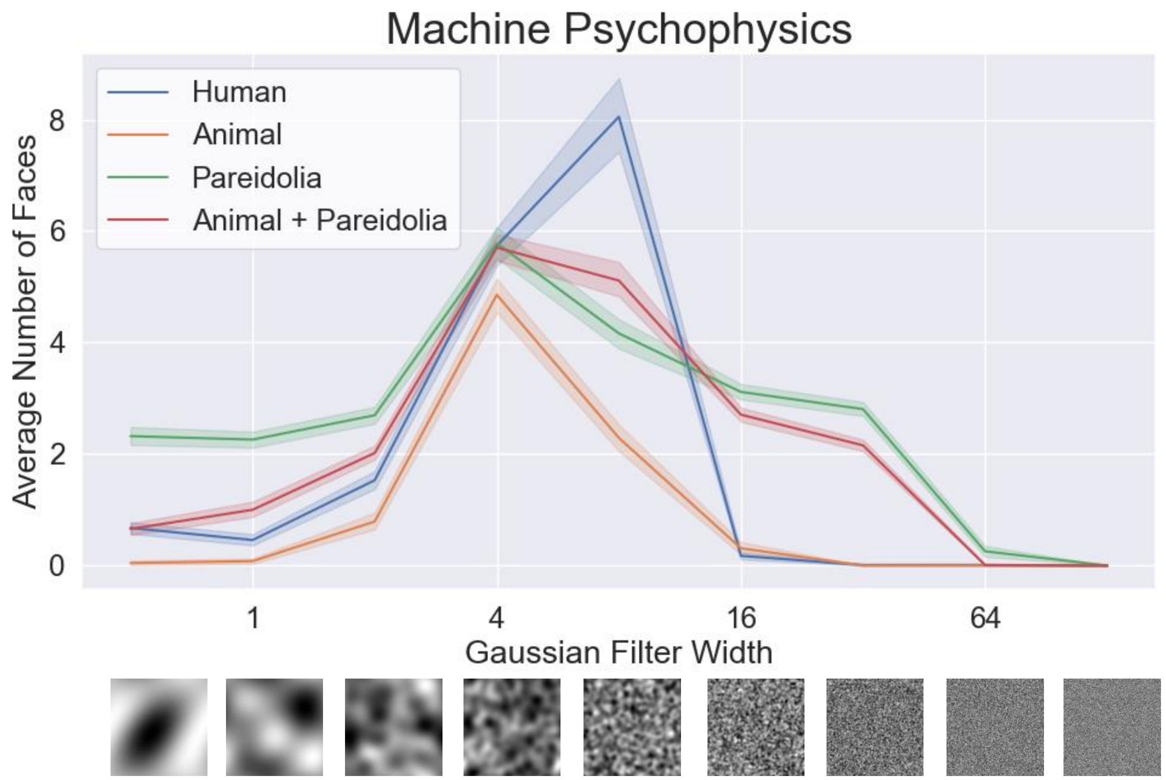

Our Contributions

The Faces in Things Dataset

Animal Faces → Pareidolia

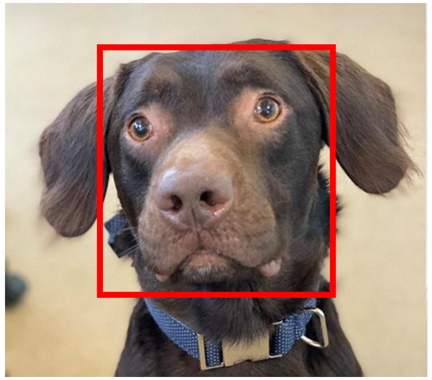

Mathematics of Pareidolia

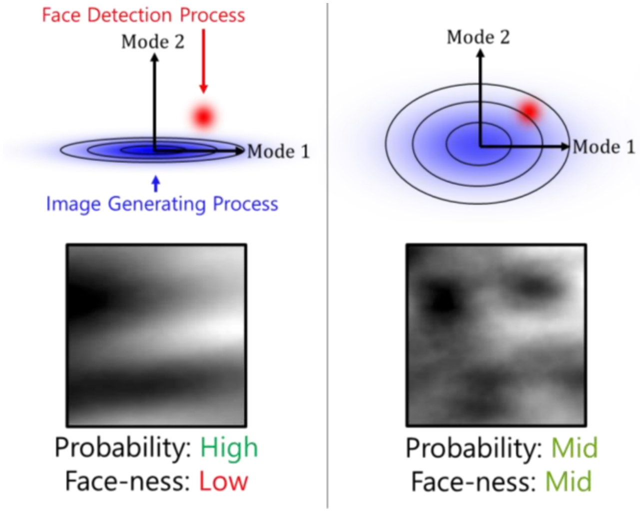

Human and Machine Verification of theory

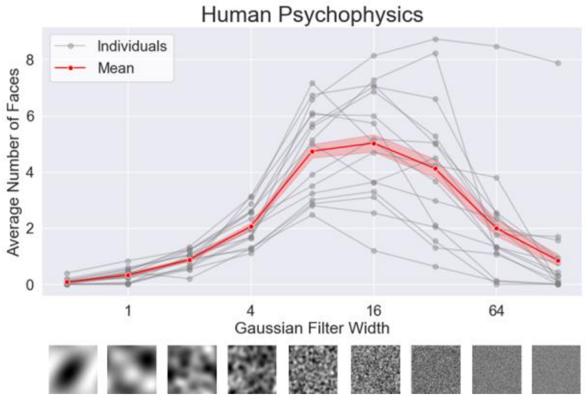

Seeing Faces in Things: A Model and Dataset for Pareidolia

ECCV 2024

aka.ms/faces-in-things