Negative Results

Computer Vision
Fall 2018
Columbia University

How are projects going?

Image Formation

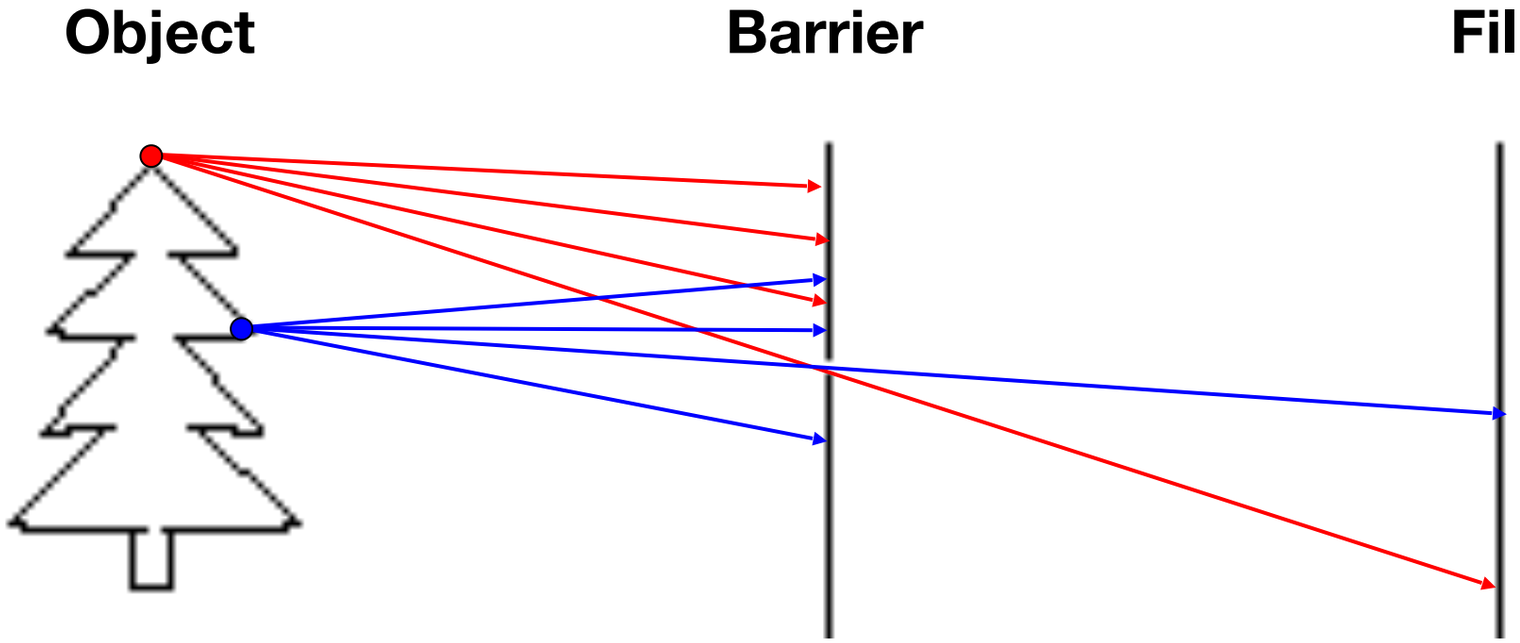

Slide credit: Steve Seitz

Emission Theory

Alternative theory that vision is accomplish by beams emitted from the eyeball

Proponents:
1. Plato
2. Leonardo da Vinci
3. Pythagoras
4. Galen
5. Over half of college educated adults in 2000

Fundamentally Misunderstanding Visual Perception. Winer et al

Slide credit: Alyosha Efros

Emission Theory

The “evidence:”

1. In near darkness, cat eyes are still visible, deer in headlights, also red eye

2. Taping the eye causes short flashes (don’t try it)

3. Evil eye, feel when somebody is looking at you

4. Elegance: similar to touch

Michelson - Morley Experiment

THE

AMERICAN JOURNAL OF SCIENCE.

[THIRD SERIES.]

ART. XXXVI.—*On the Relative Motion of the Earth and the Luminiferous Ether;* by ALBERT A. MICHELSON and EDWARD W. MORLEY.*

THE discovery of the aberration of light was soon followed by an explanation according to the emission theory. The effect was attributed to a simple composition of the velocity of light with the velocity of the earth in its orbit. The difficulties in this apparently sufficient explanation were overlooked until after an explanation on the undulatory theory of light was proposed. This new explanation was at first almost as simple as the former. But it failed to account for the fact proved by experiment that the aberration was unchanged when observations were made with a telescope filled with water. For if the tangent of the angle of aberration is the ratio of the velocity of the earth to the velocity of light, then, since the latter velocity in water is three-fourths its velocity in a vacuum, the aberration observed with a water telescope should be four-thirds of its true value.†

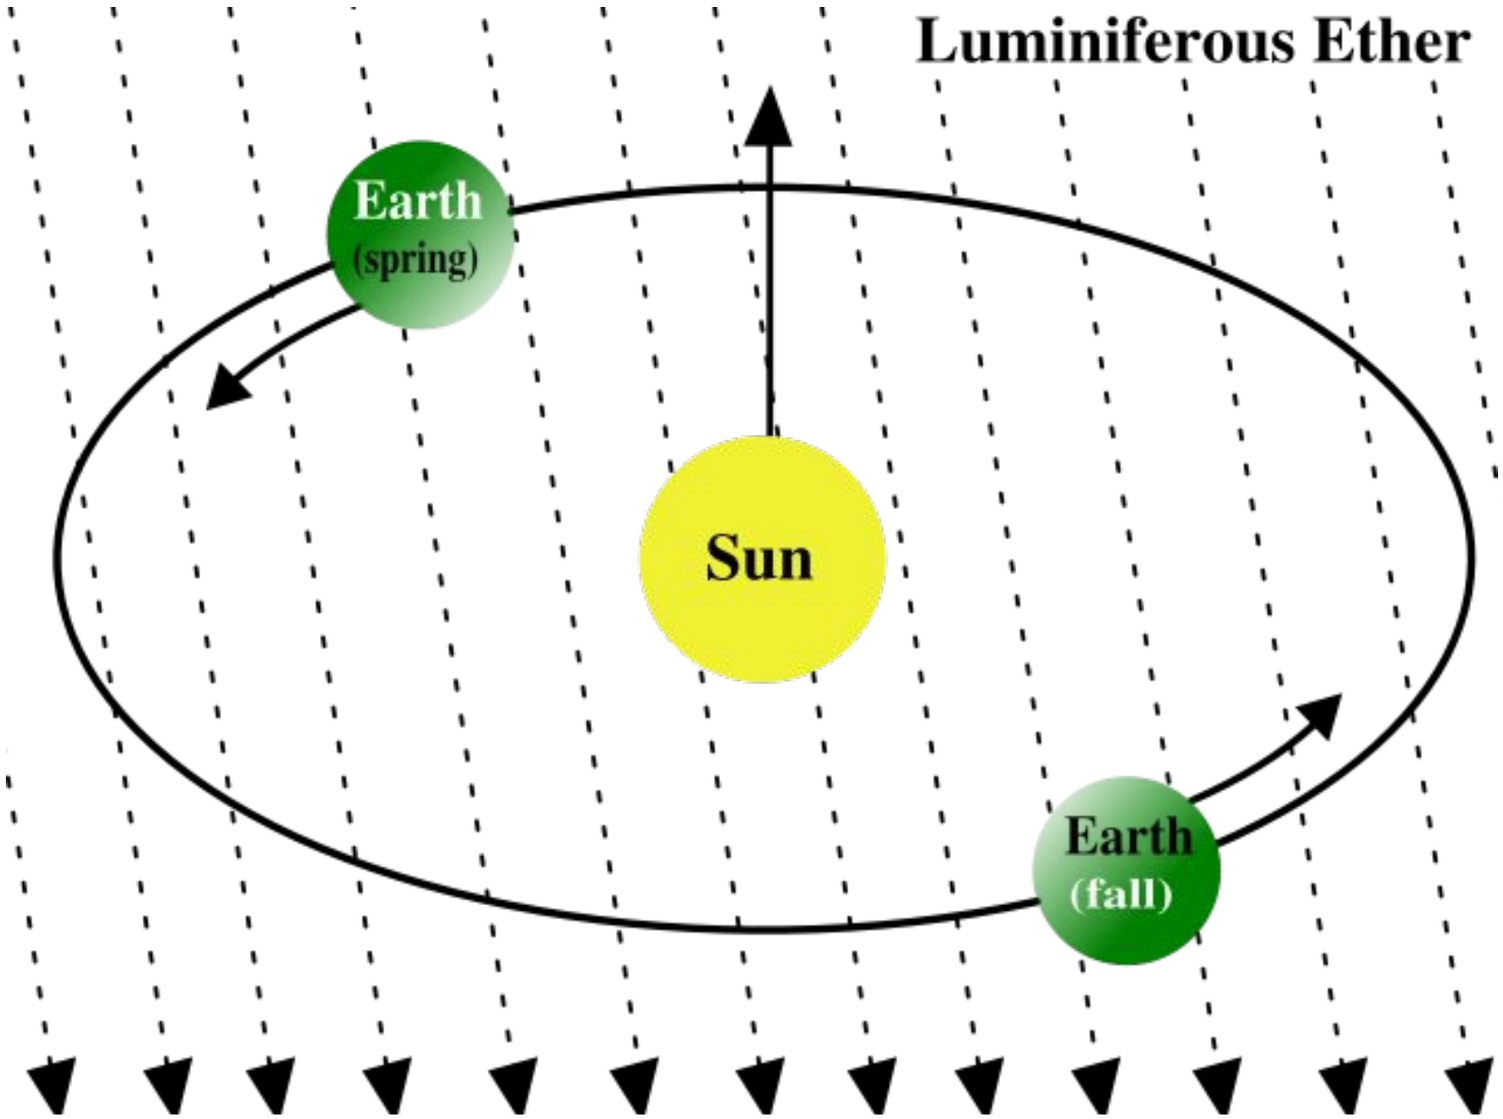

Michelson - Morley Experiment

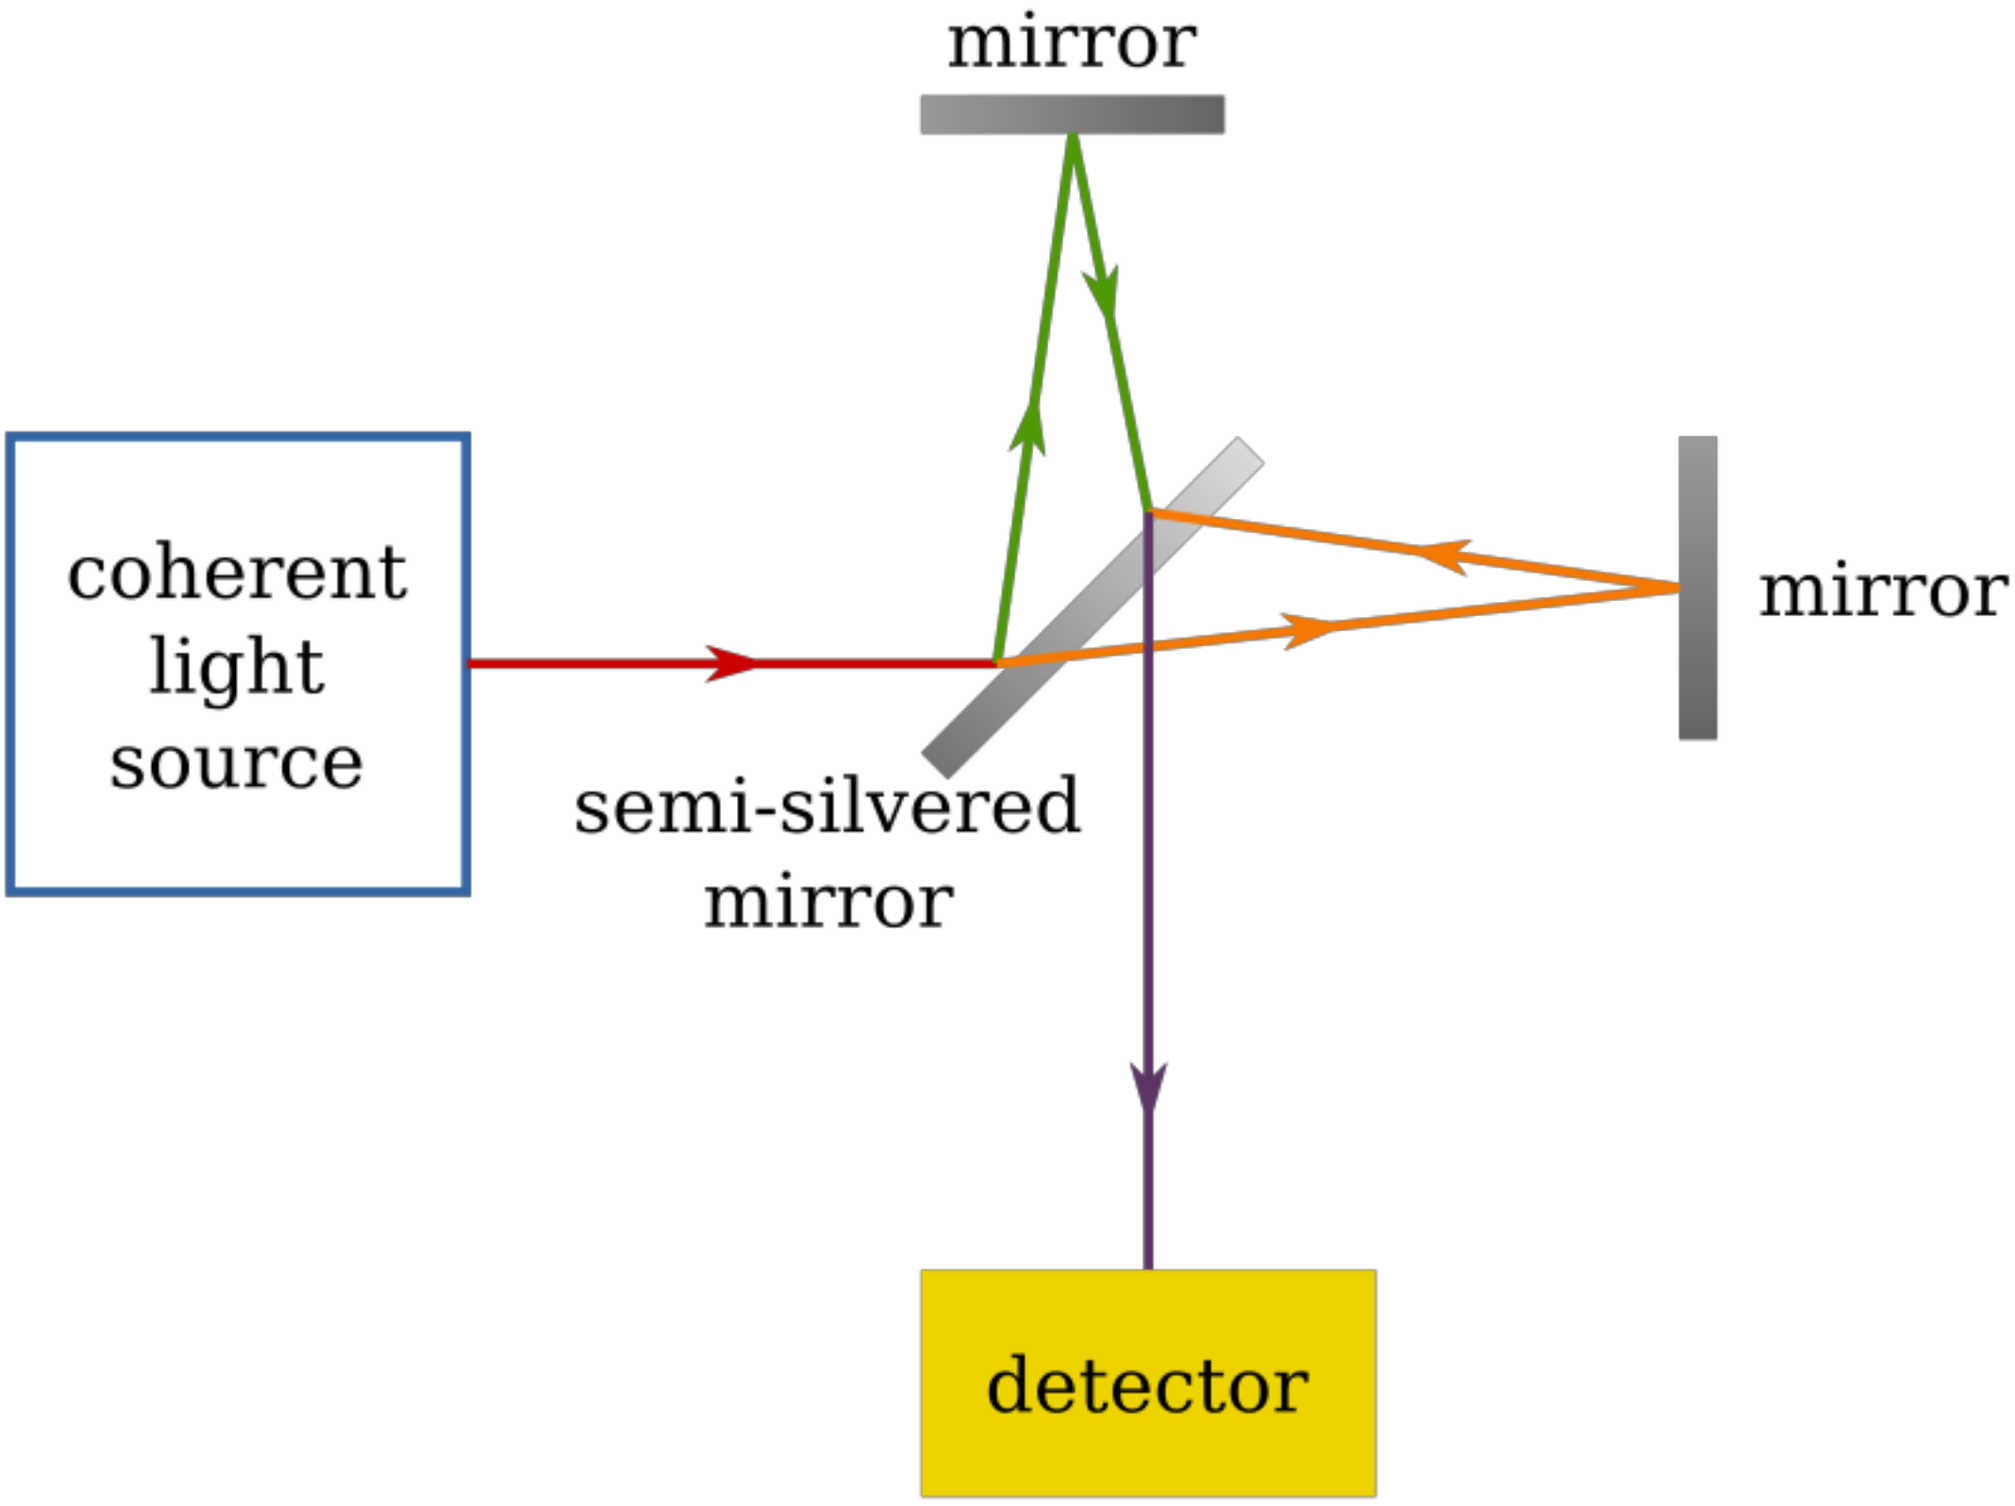

"If you torture the data long enough, it will confess to anything”

— *How to Lie With Statistics* by Darrell Huff

We prefer algorithms to data

Data

Features

Algorithm

Slide credit: Alyosha Efros


Data is messy

Recognition circa 2010

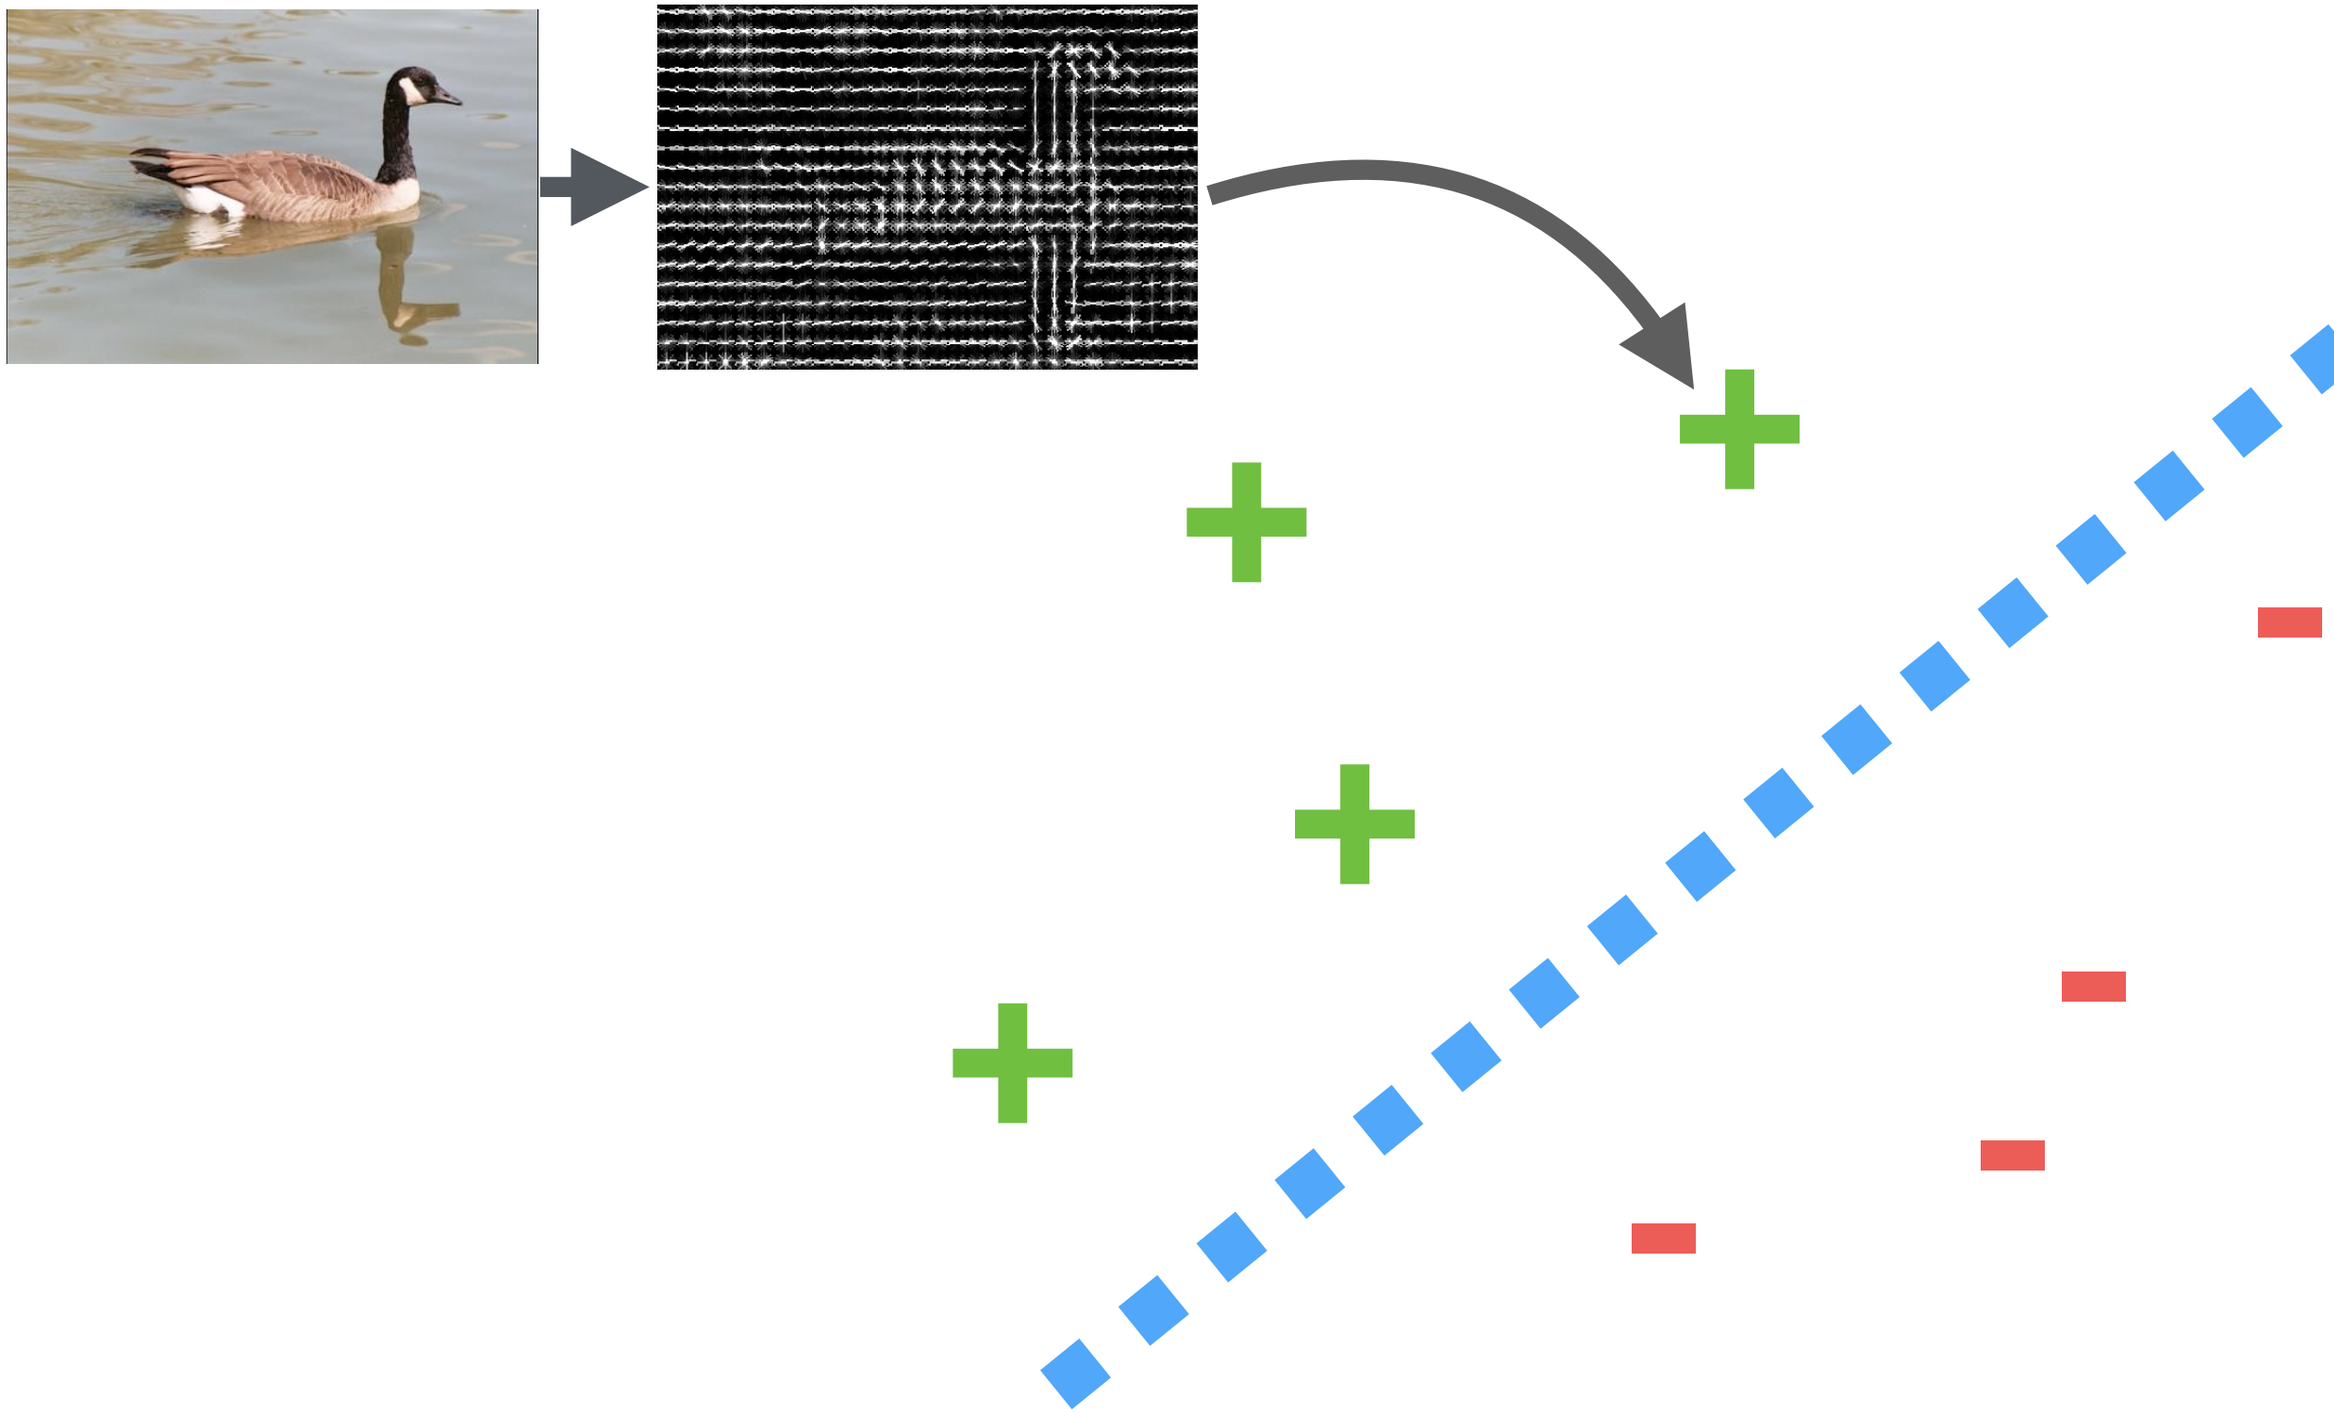

In 2013…

Chair

In 2013...

Car

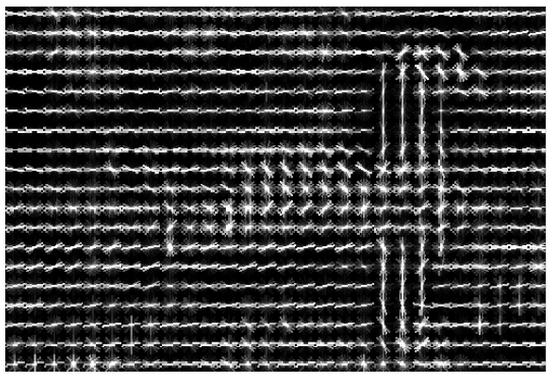

Recover
Image

Collect
Label

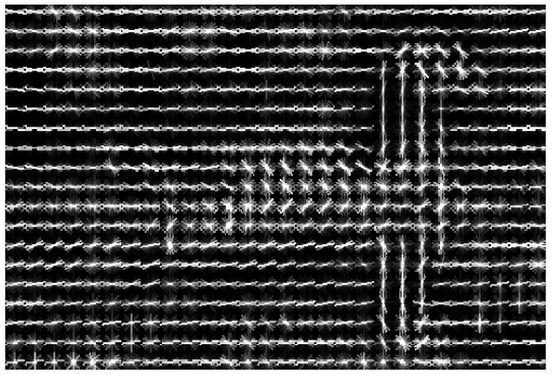

Recover
Image

Collect
Label

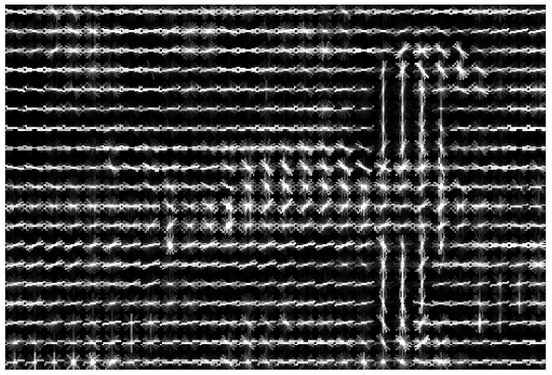

Recover
Image

Collect
Label

What do we need?

1. Algorithm to select examples for learning

2. Recover images from feature space

3. A very patient human annotator

What do we need?

1. Algorithm to select examples for learning **(???)**

2. Recover images from feature space **(???)**

3. A very patient human annotator **(me)**

Inverting Features

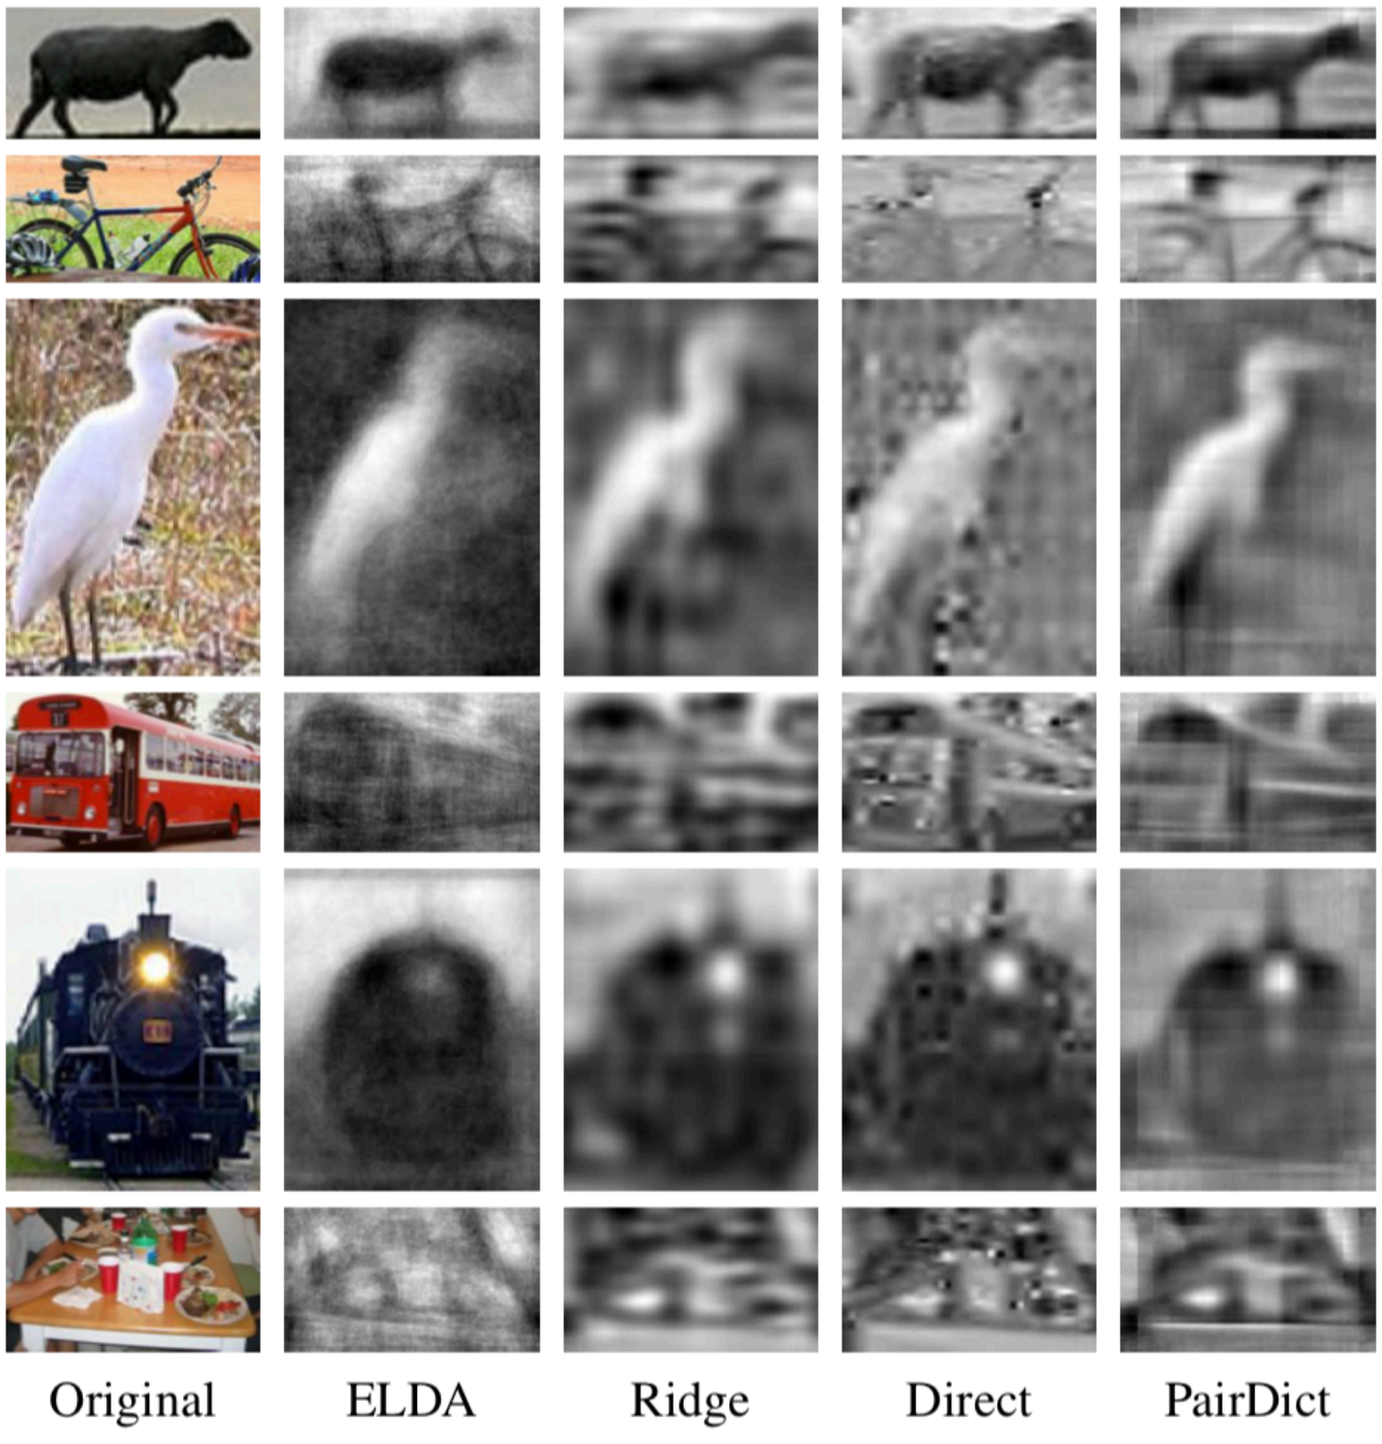

Category	ELDA	Ridge	Direct	PairDict	Glyph	Expert
aeroplane	0.433	0.391	0.568	**0.645**	0.297	0.333
bicycle	0.327	0.127	0.362	0.307	**0.405**	0.438
bird	0.364	0.263	**0.378**	0.372	0.193	0.059
boat	0.292	0.182	0.255	**0.329**	0.119	0.352
bottle	0.269	0.282	0.283	**0.446**	0.312	0.222
bus	0.473	0.395	**0.541**	**0.549**	0.122	0.118
car	0.397	0.457	**0.617**	0.585	0.359	0.389
cat	0.219	0.178	**0.381**	0.199	0.139	0.286
chair	0.099	0.239	0.223	**0.386**	0.119	0.167
cow	0.133	0.103	**0.230**	0.197	0.072	0.214
table	0.152	0.064	0.162	**0.237**	0.071	0.125
dog	0.222	0.316	**0.351**	0.343	0.107	0.150
horse	0.260	0.290	0.354	**0.446**	0.144	0.150
motorbike	0.221	0.232	**0.396**	0.224	0.298	0.350
person	0.458	0.546	0.502	**0.676**	0.301	0.375
pottedplant	0.112	0.109	**0.203**	0.091	0.080	0.136
sheep	0.227	0.194	**0.368**	0.253	0.041	0.000
sofa	0.138	0.100	0.162	**0.293**	0.104	0.000
train	0.311	0.244	0.316	**0.404**	0.173	0.133
tvmonitor	0.537	0.439	0.449	**0.682**	0.354	0.666
Mean	0.282	0.258	0.355	**0.383**	0.191	0.233

Inverting Features

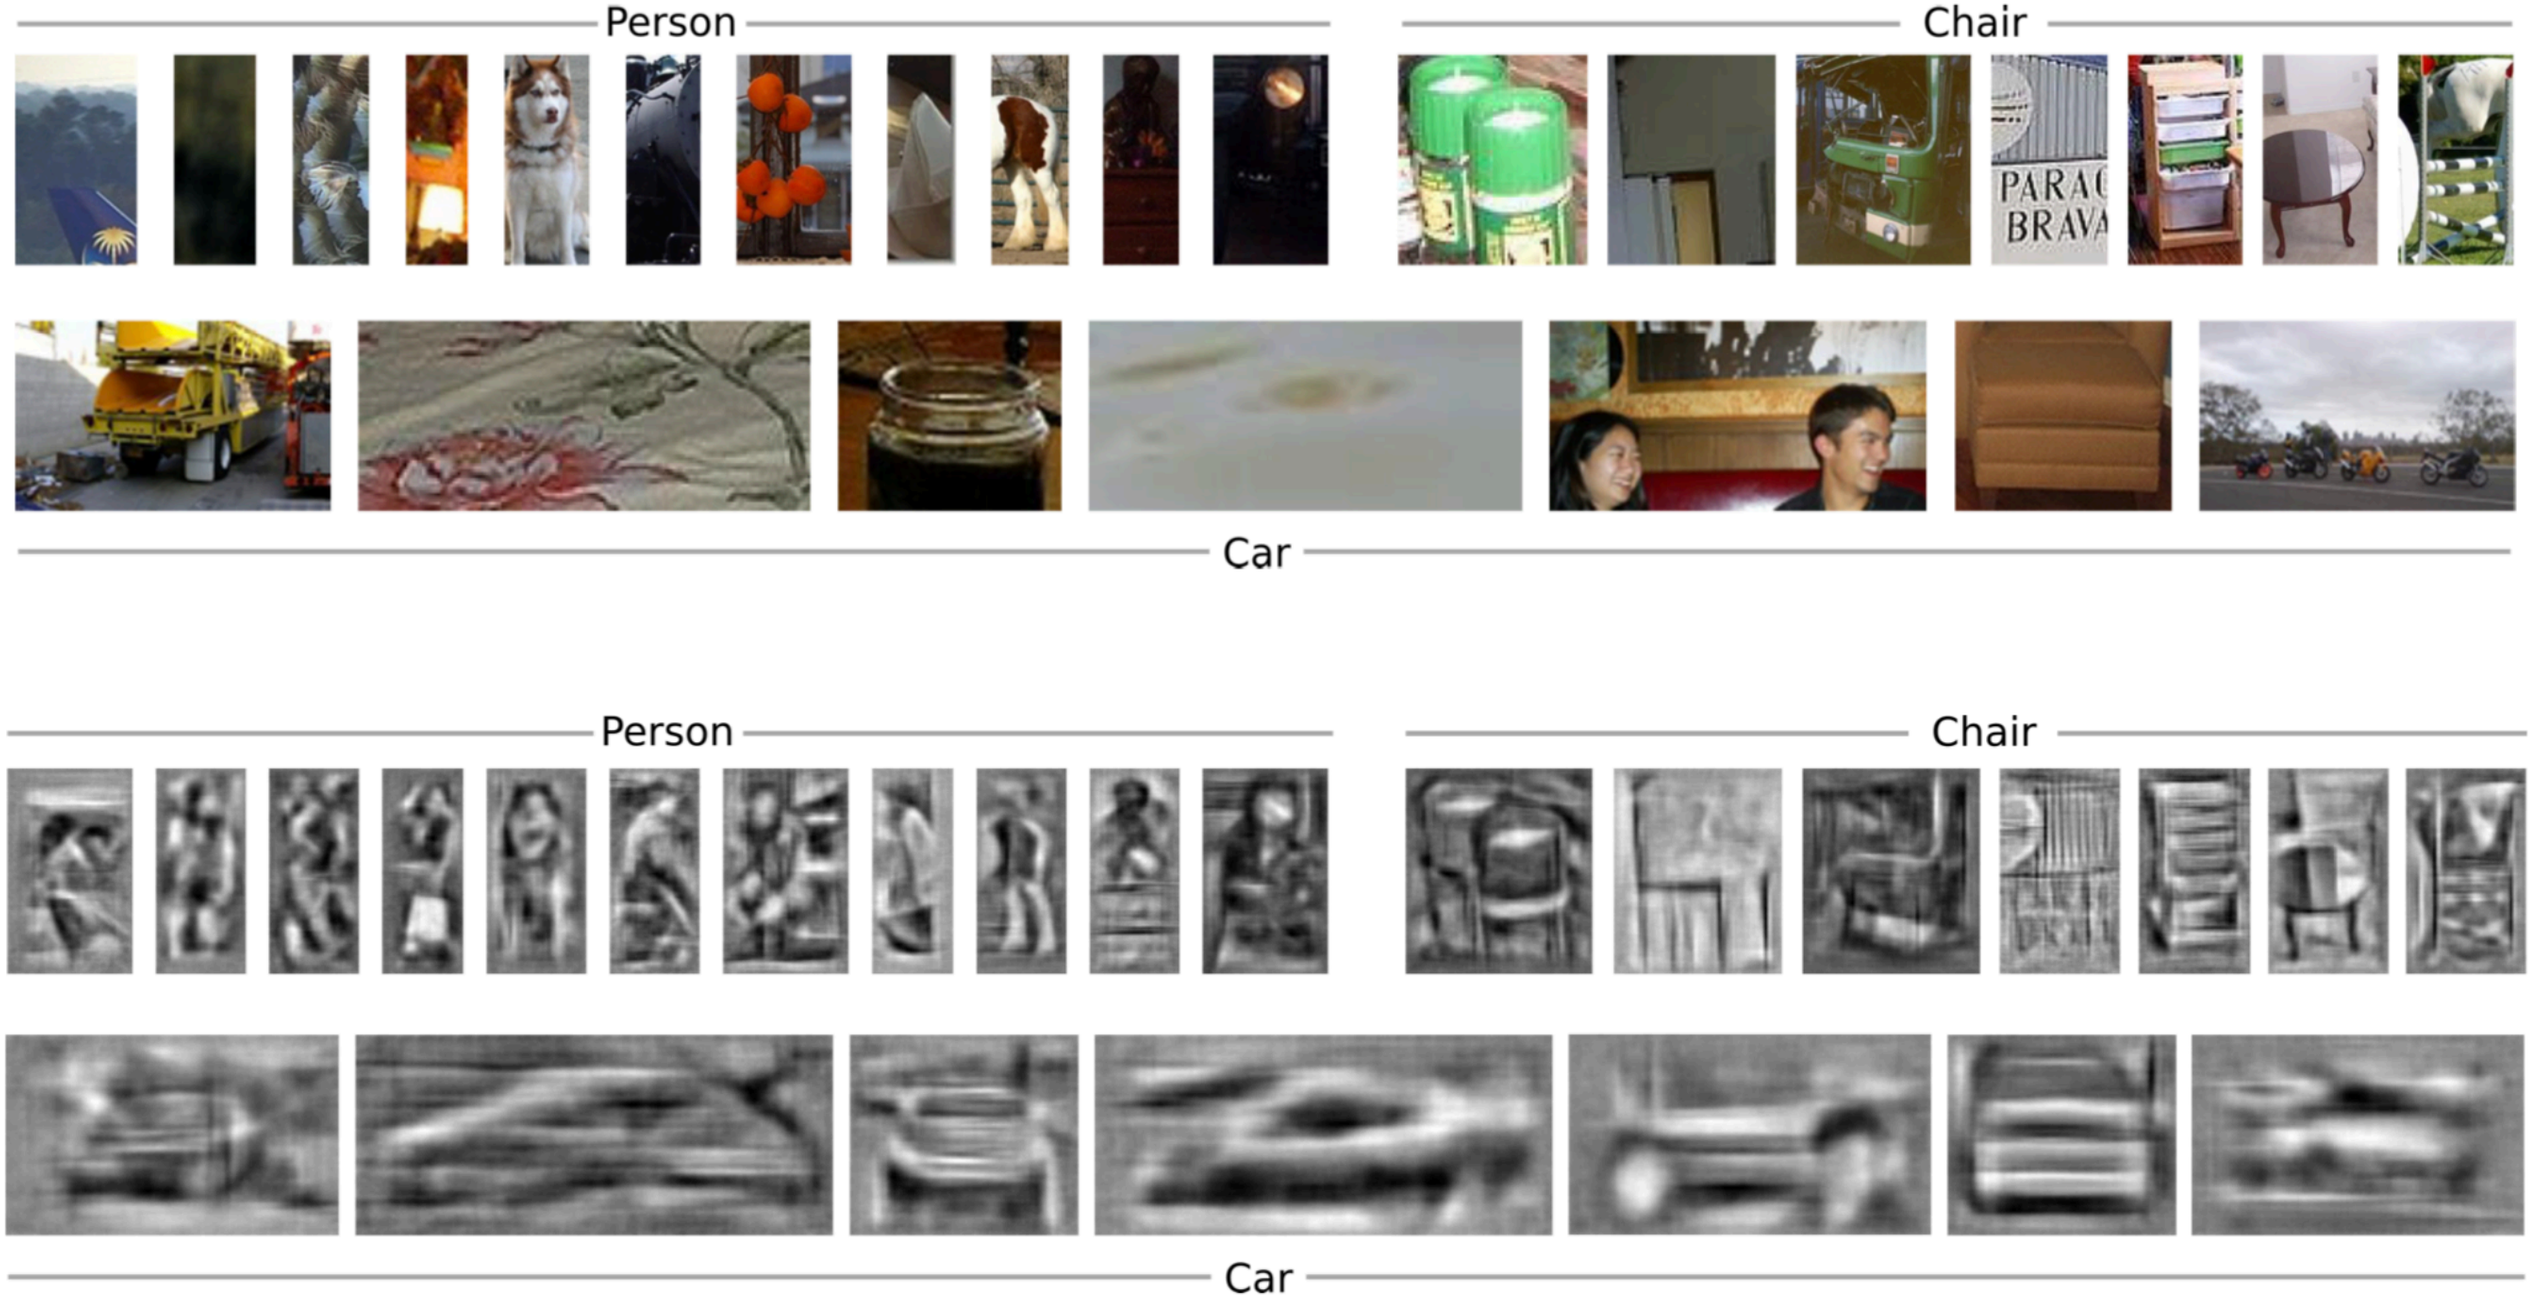

What do we need?

1. Algorithm to select examples for learning **(???)**

2. Recover images from feature space **(my inversion)**

3. A very patient human annotator **(me)**

SVMs (linear)

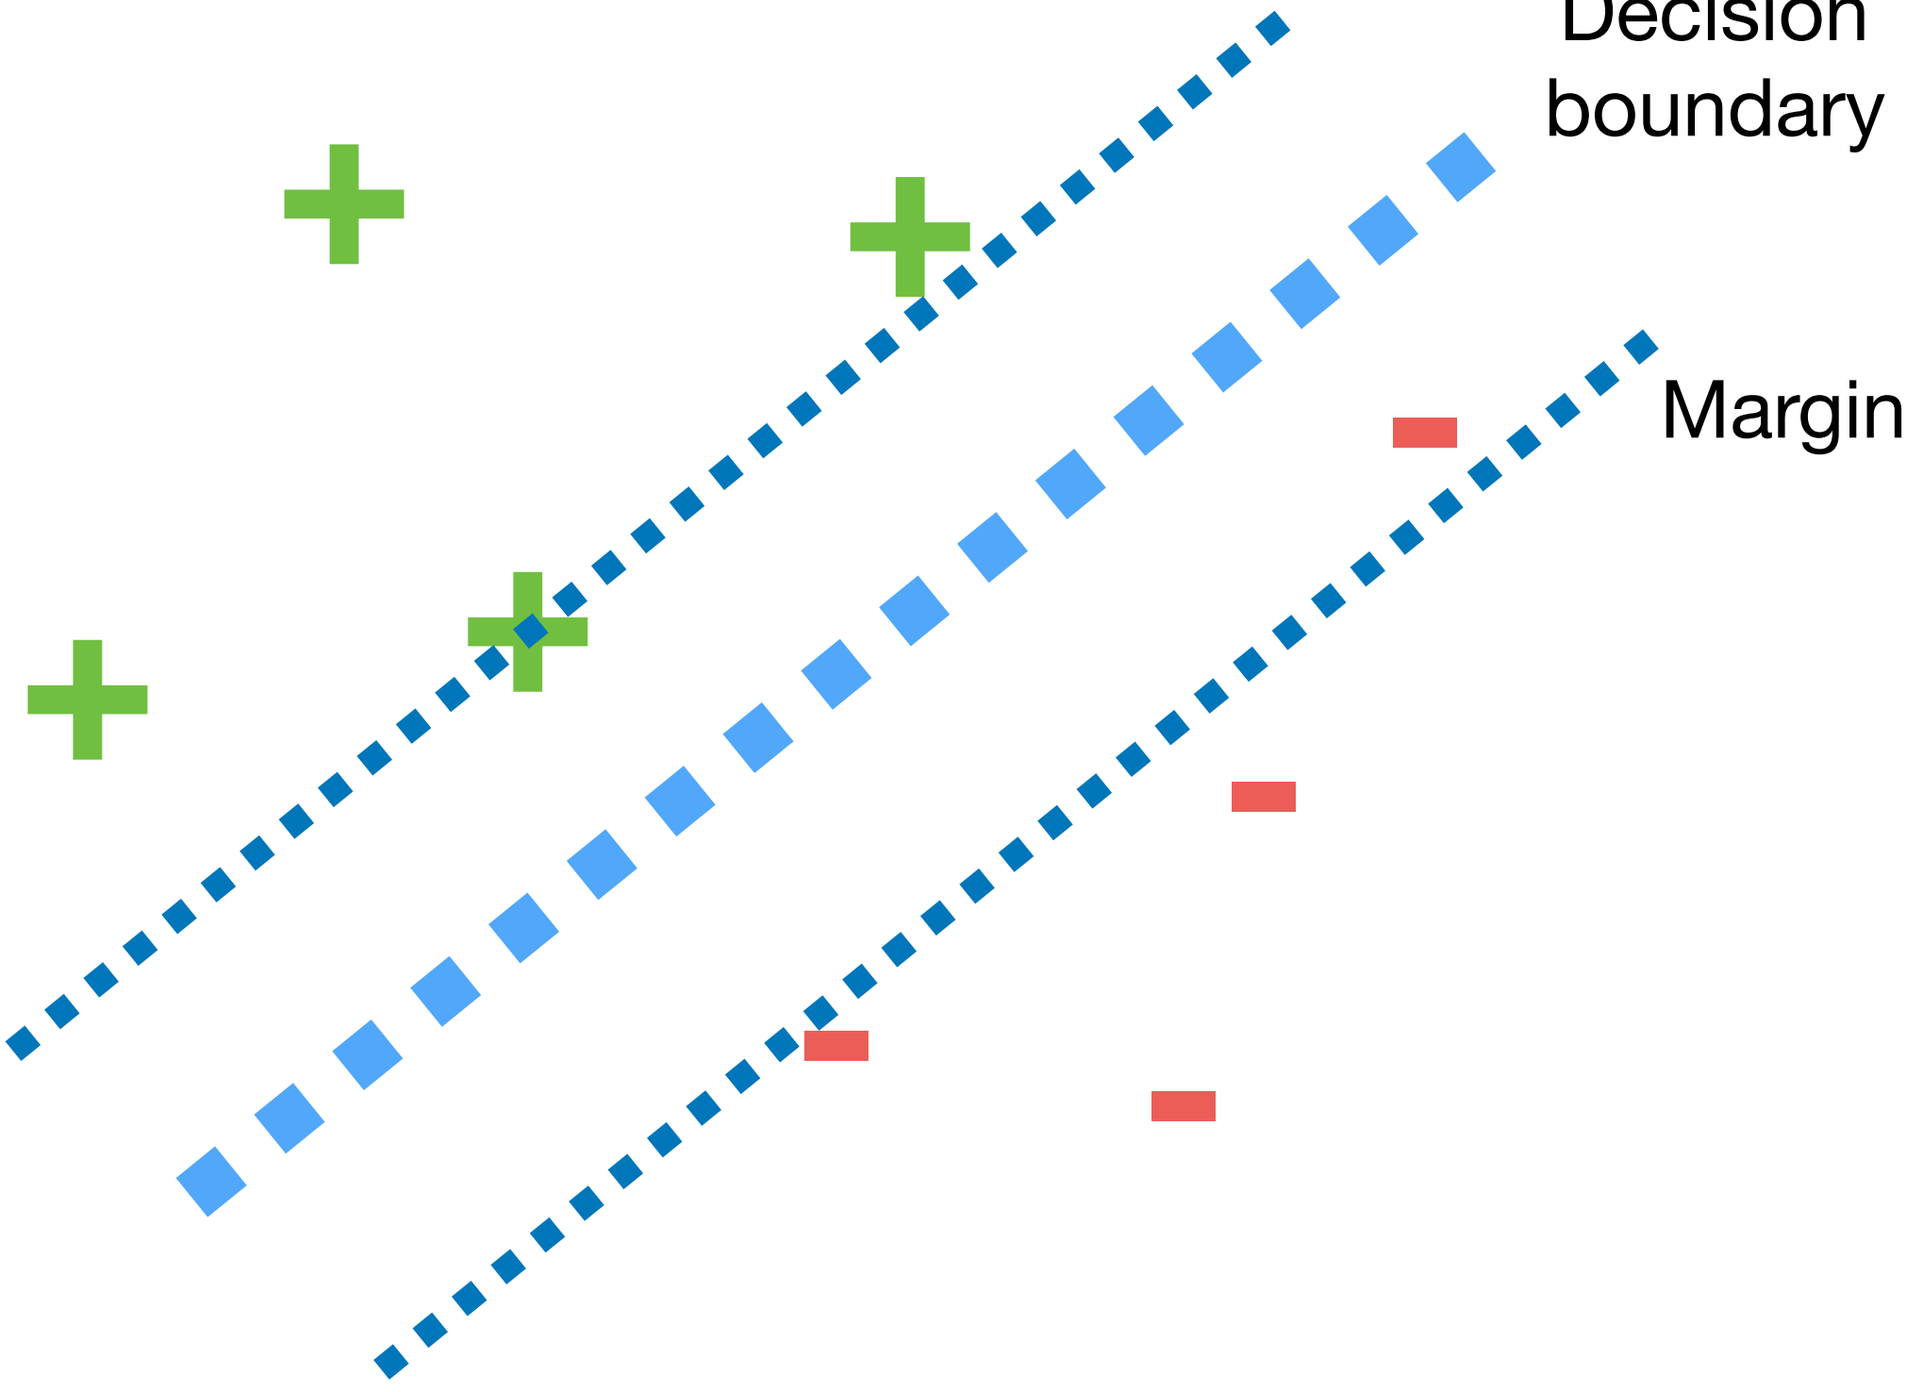

SVMs (linear)

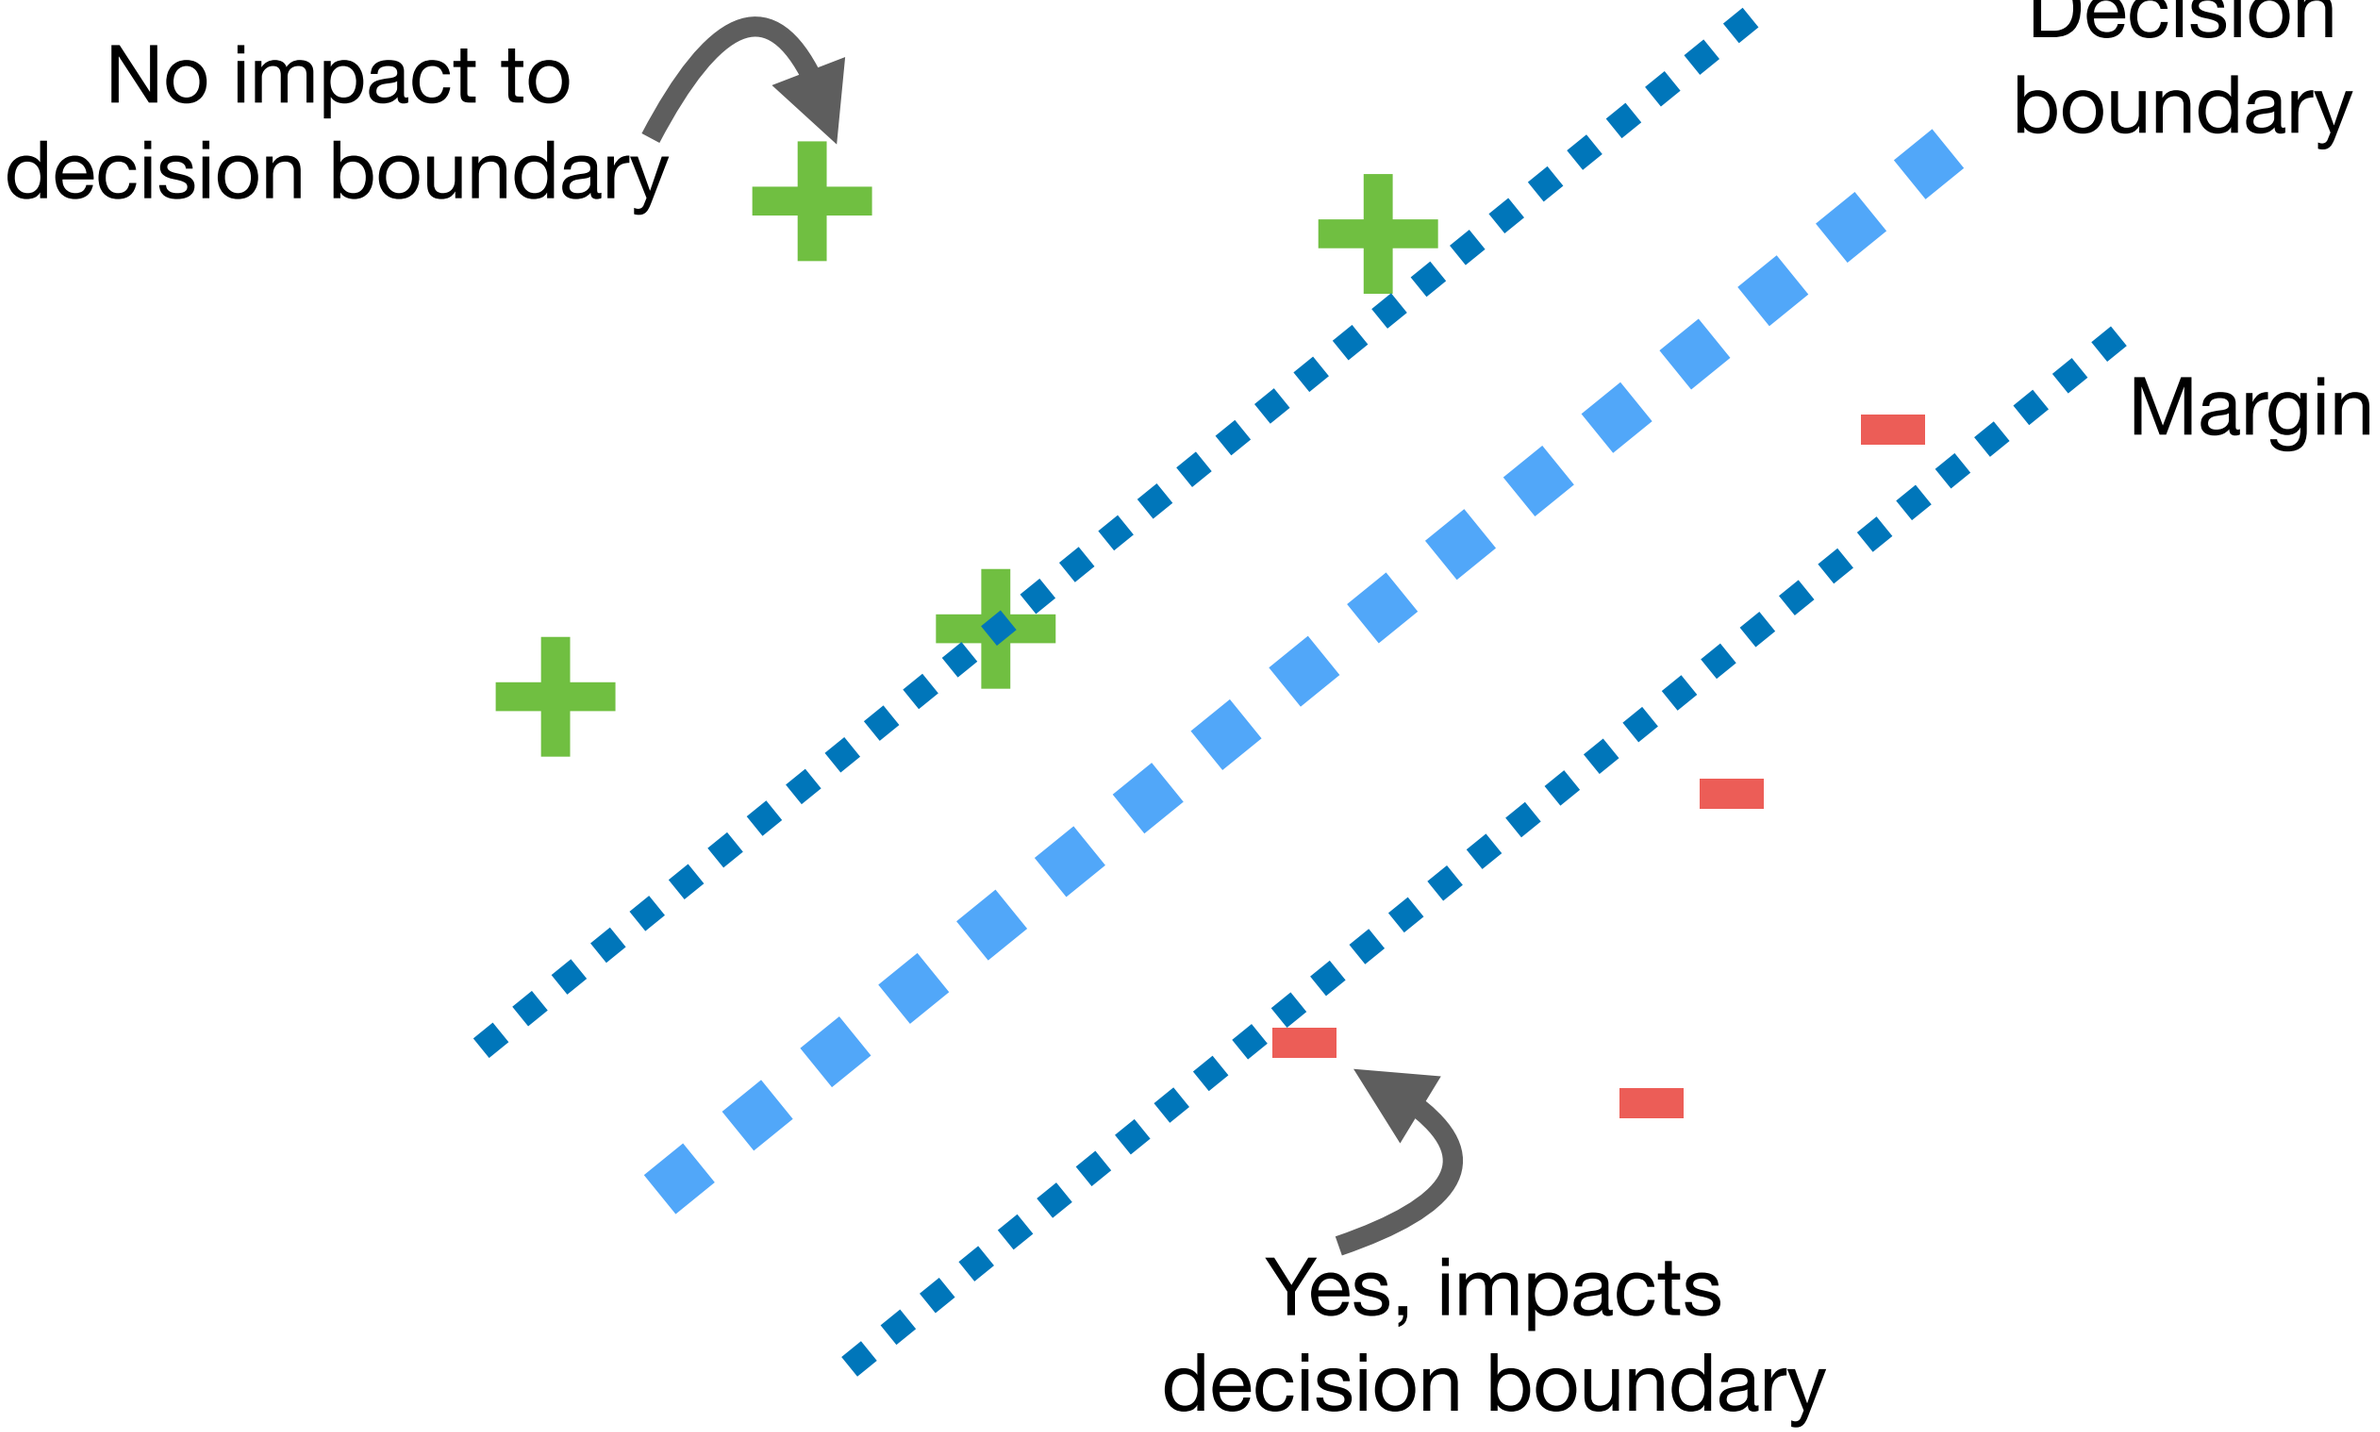
No impact to
decision boundary
Decision
boundary
Margin
Yes, impacts
decision boundary

SVMs (linear)

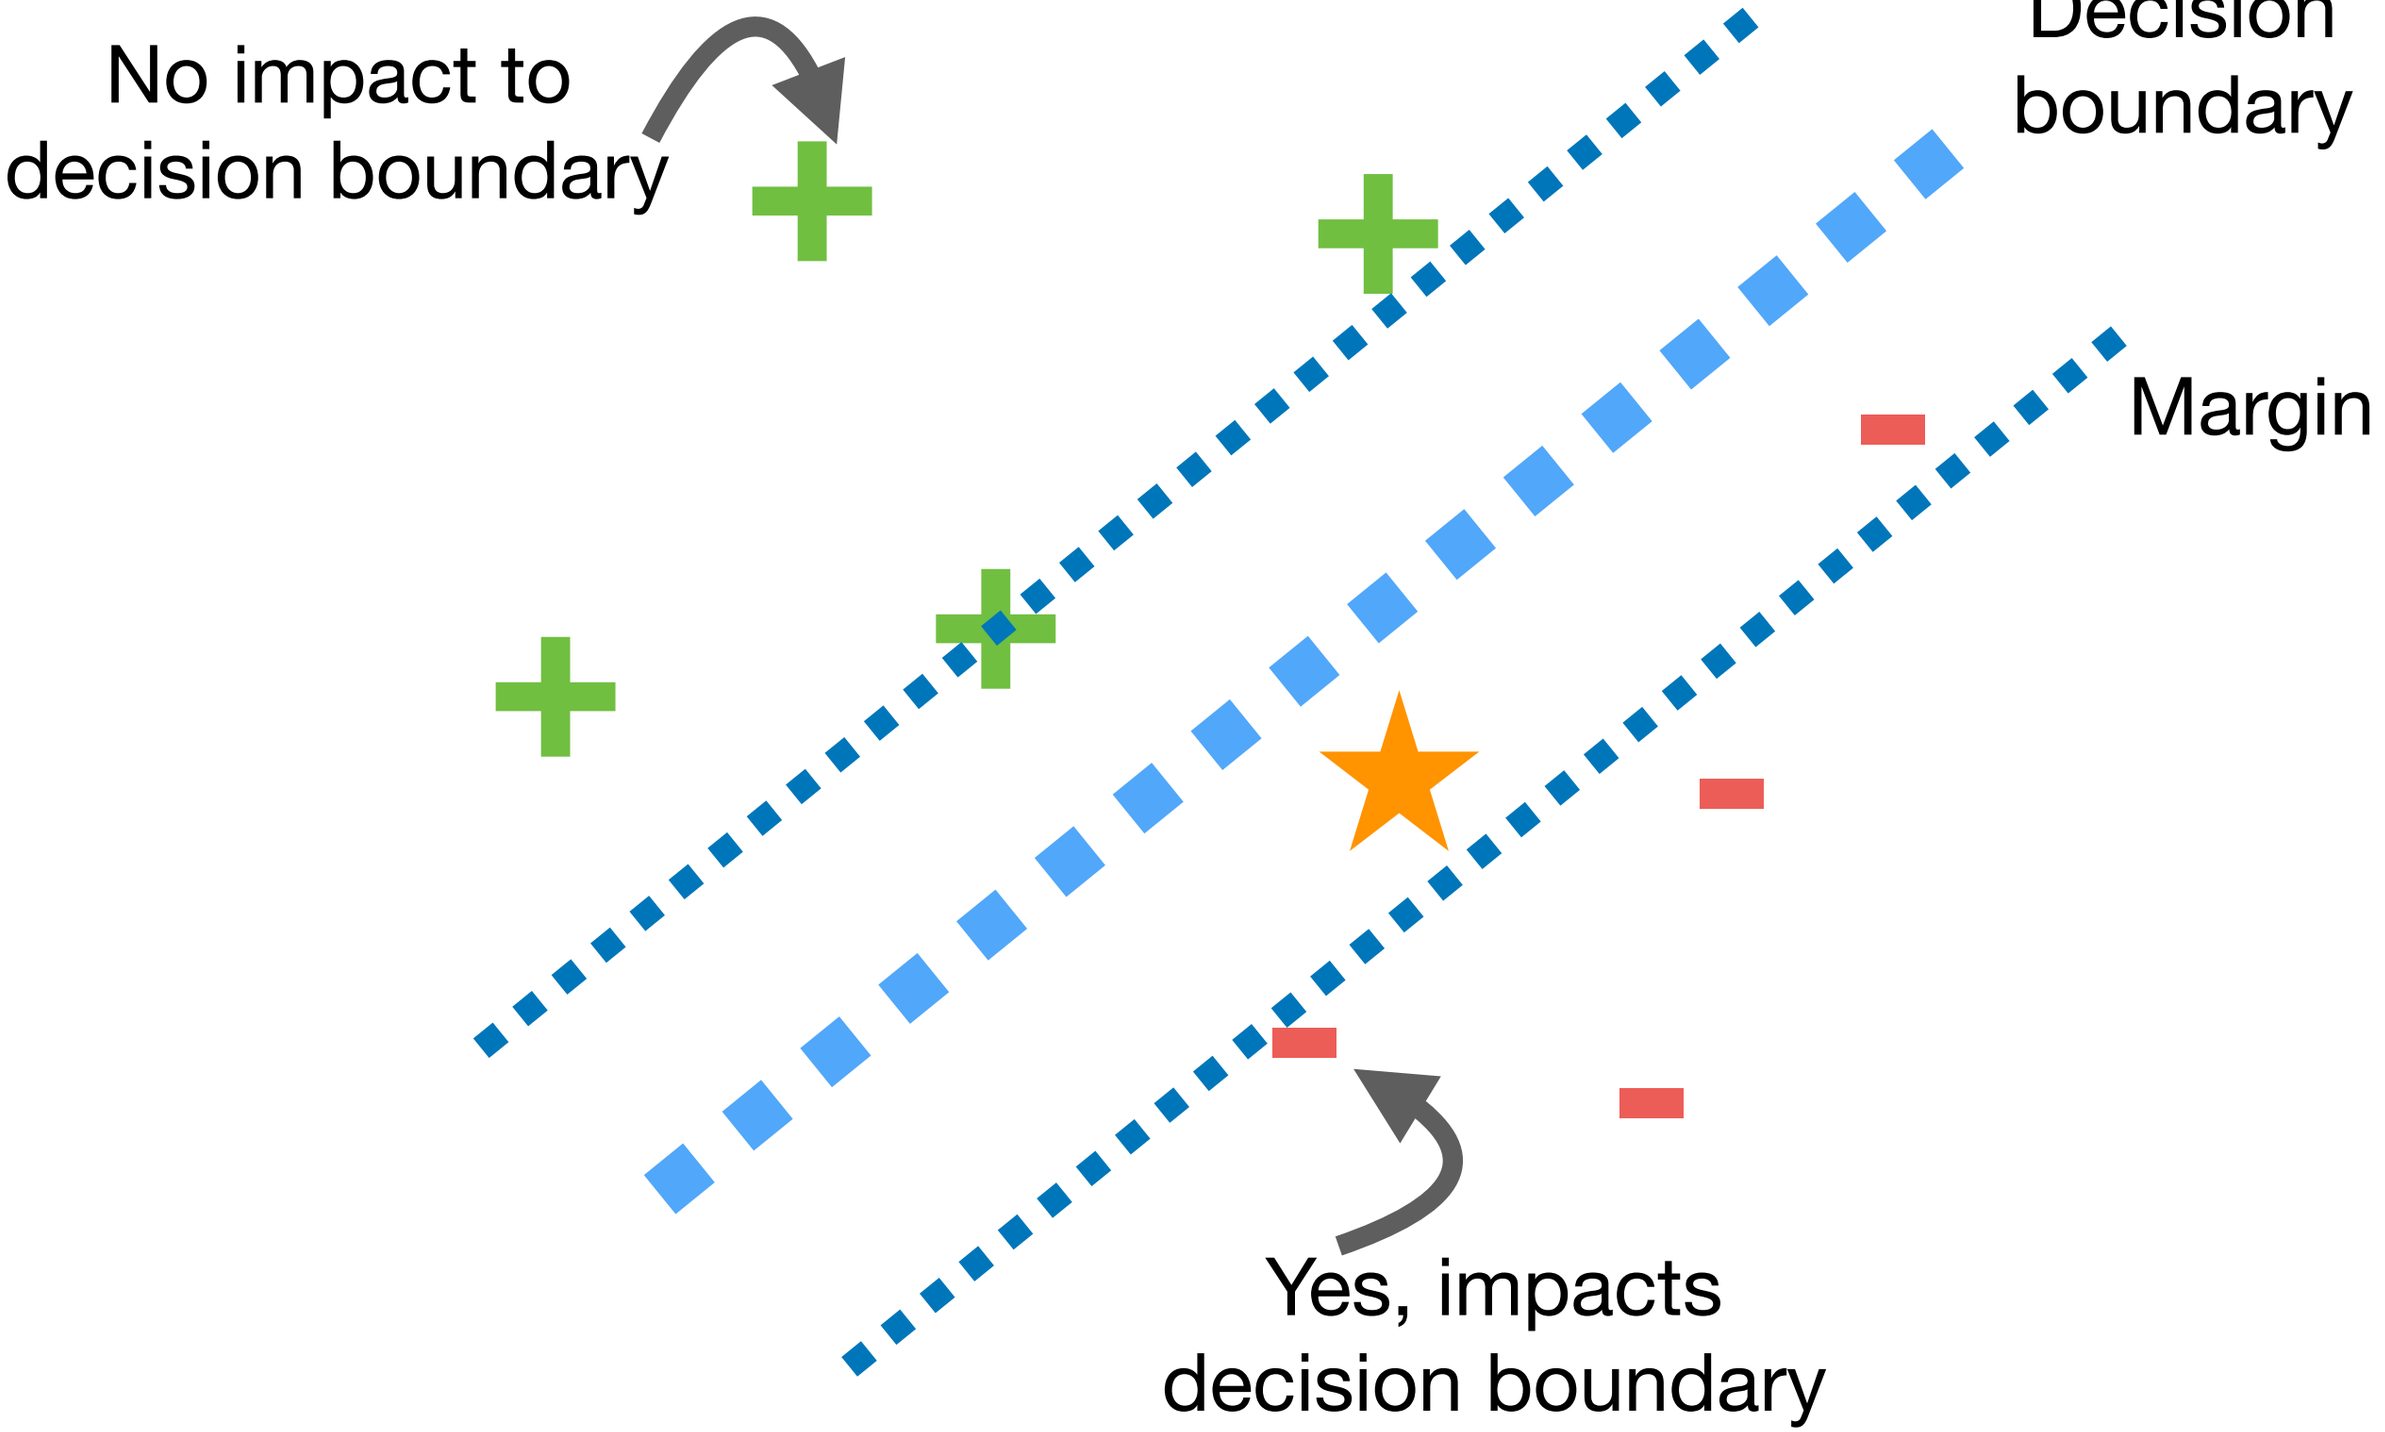
No impact to
decision boundary
Decision
boundary
Margin
Yes, impacts
decision boundary

Results

$$x \sim \mathcal{N}(0_d, I_d)$$

$$\phi^{-1}(x)$$

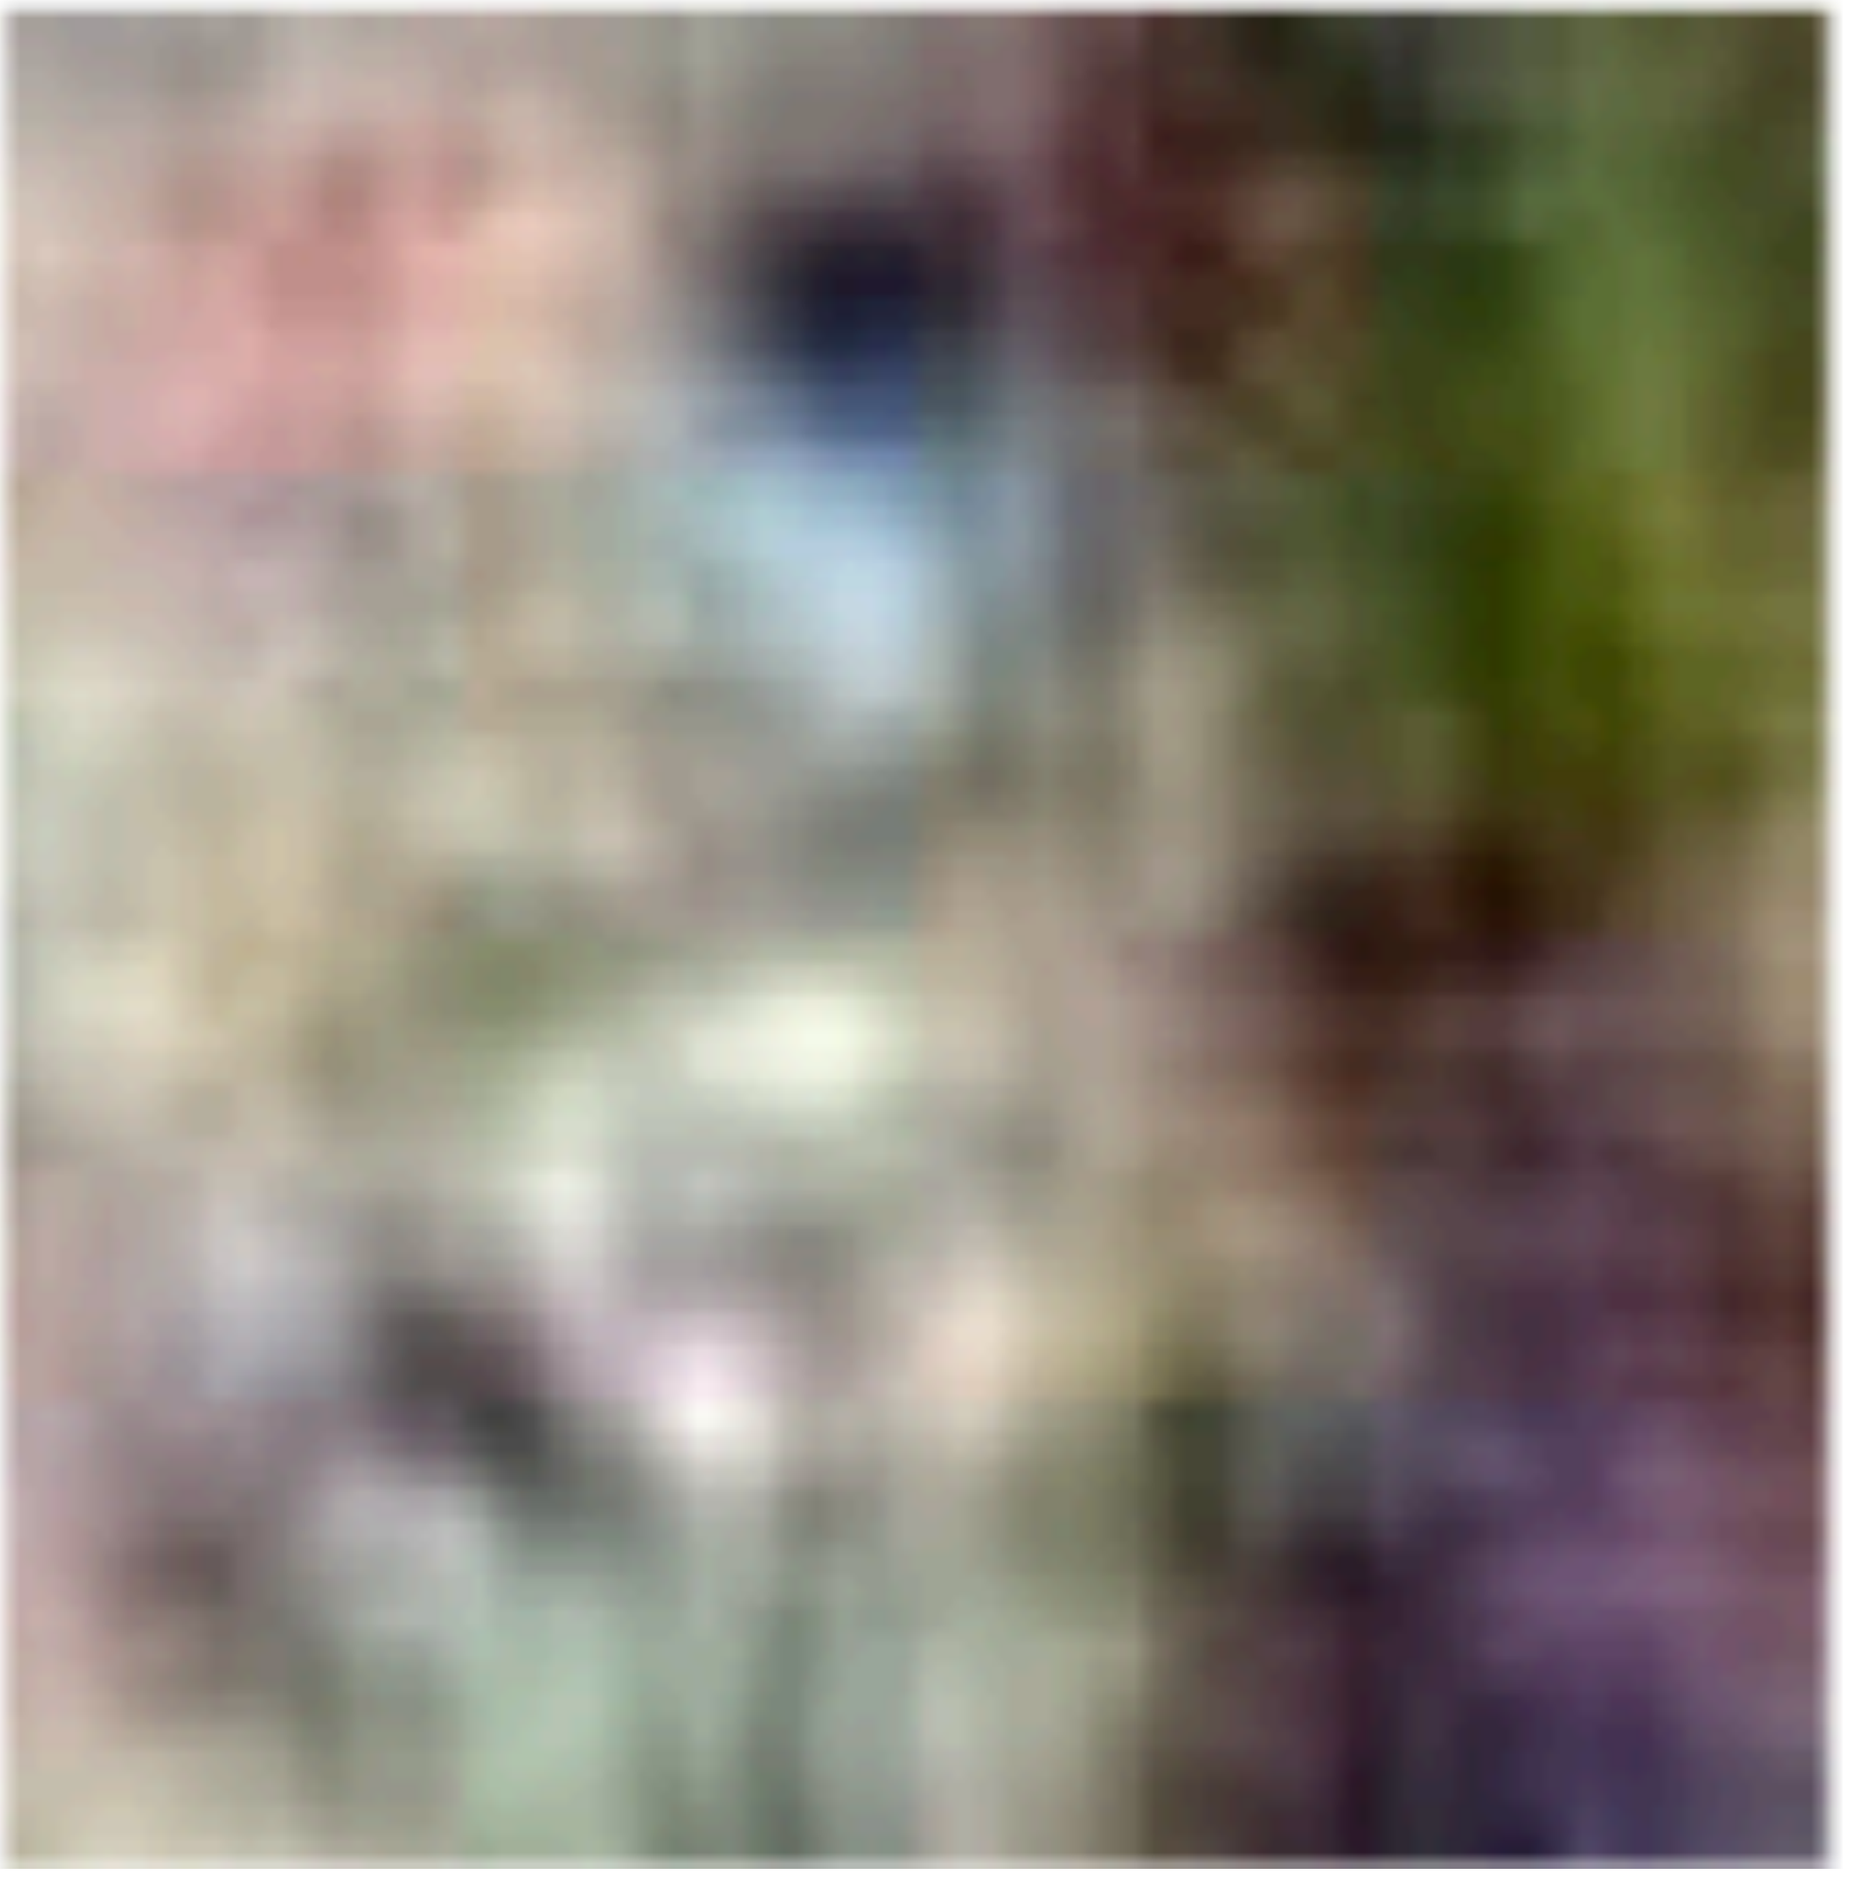

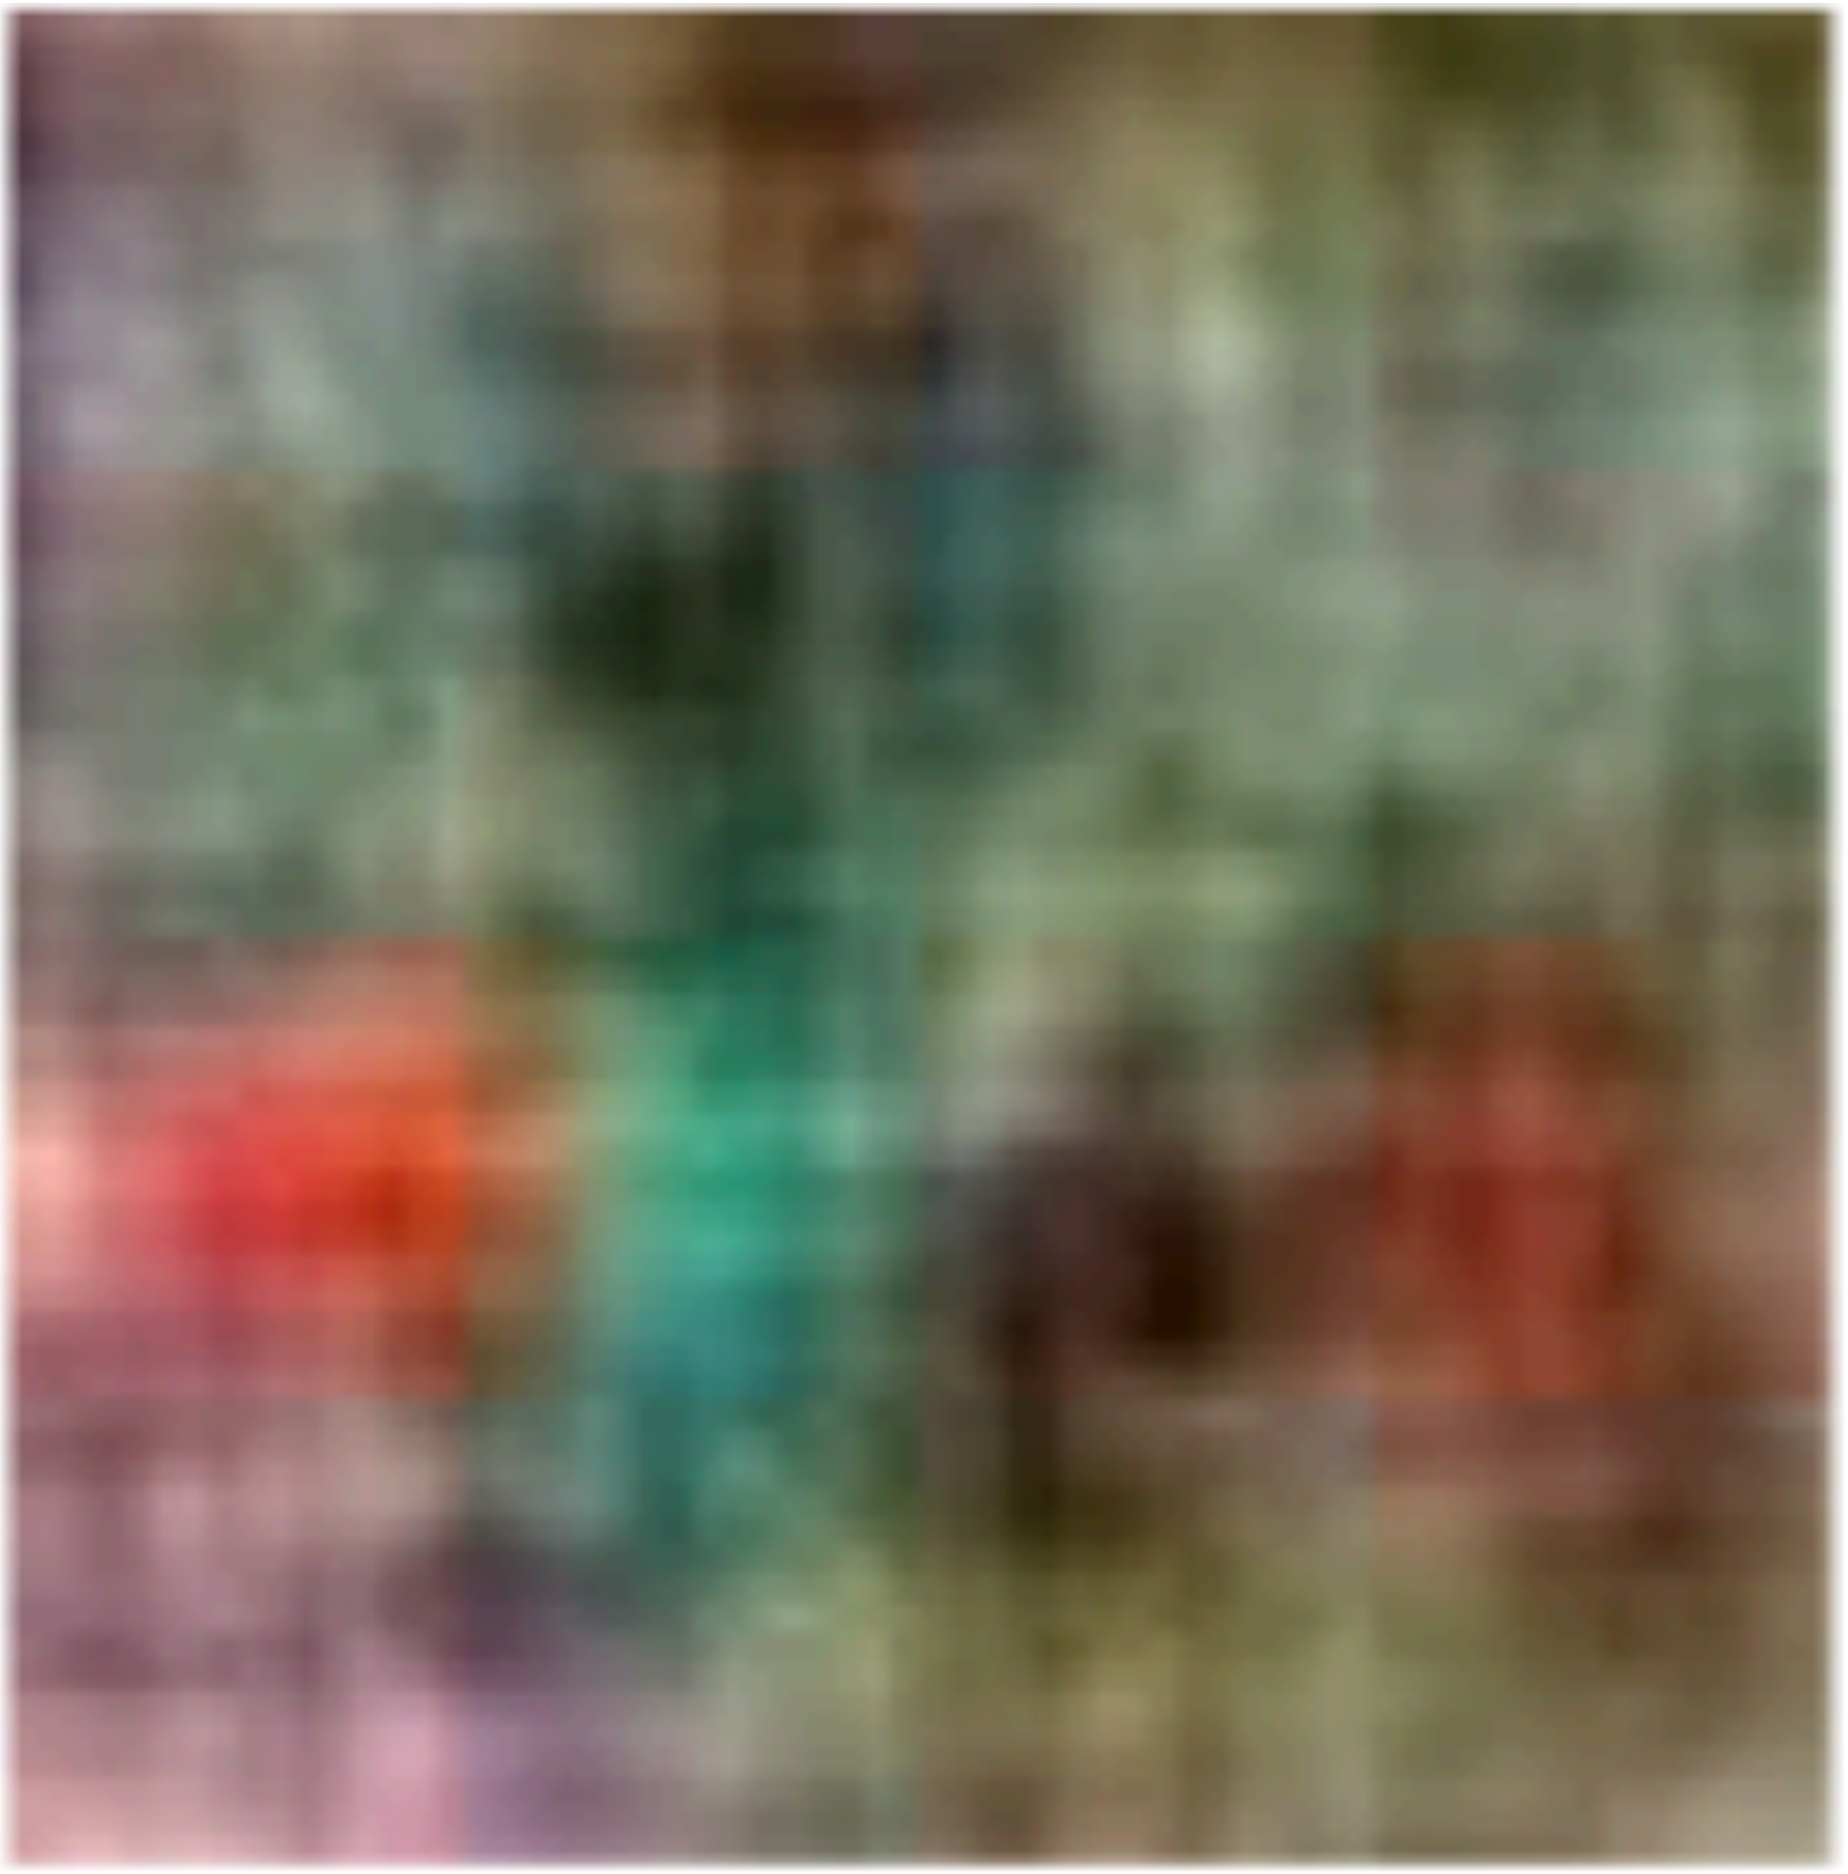

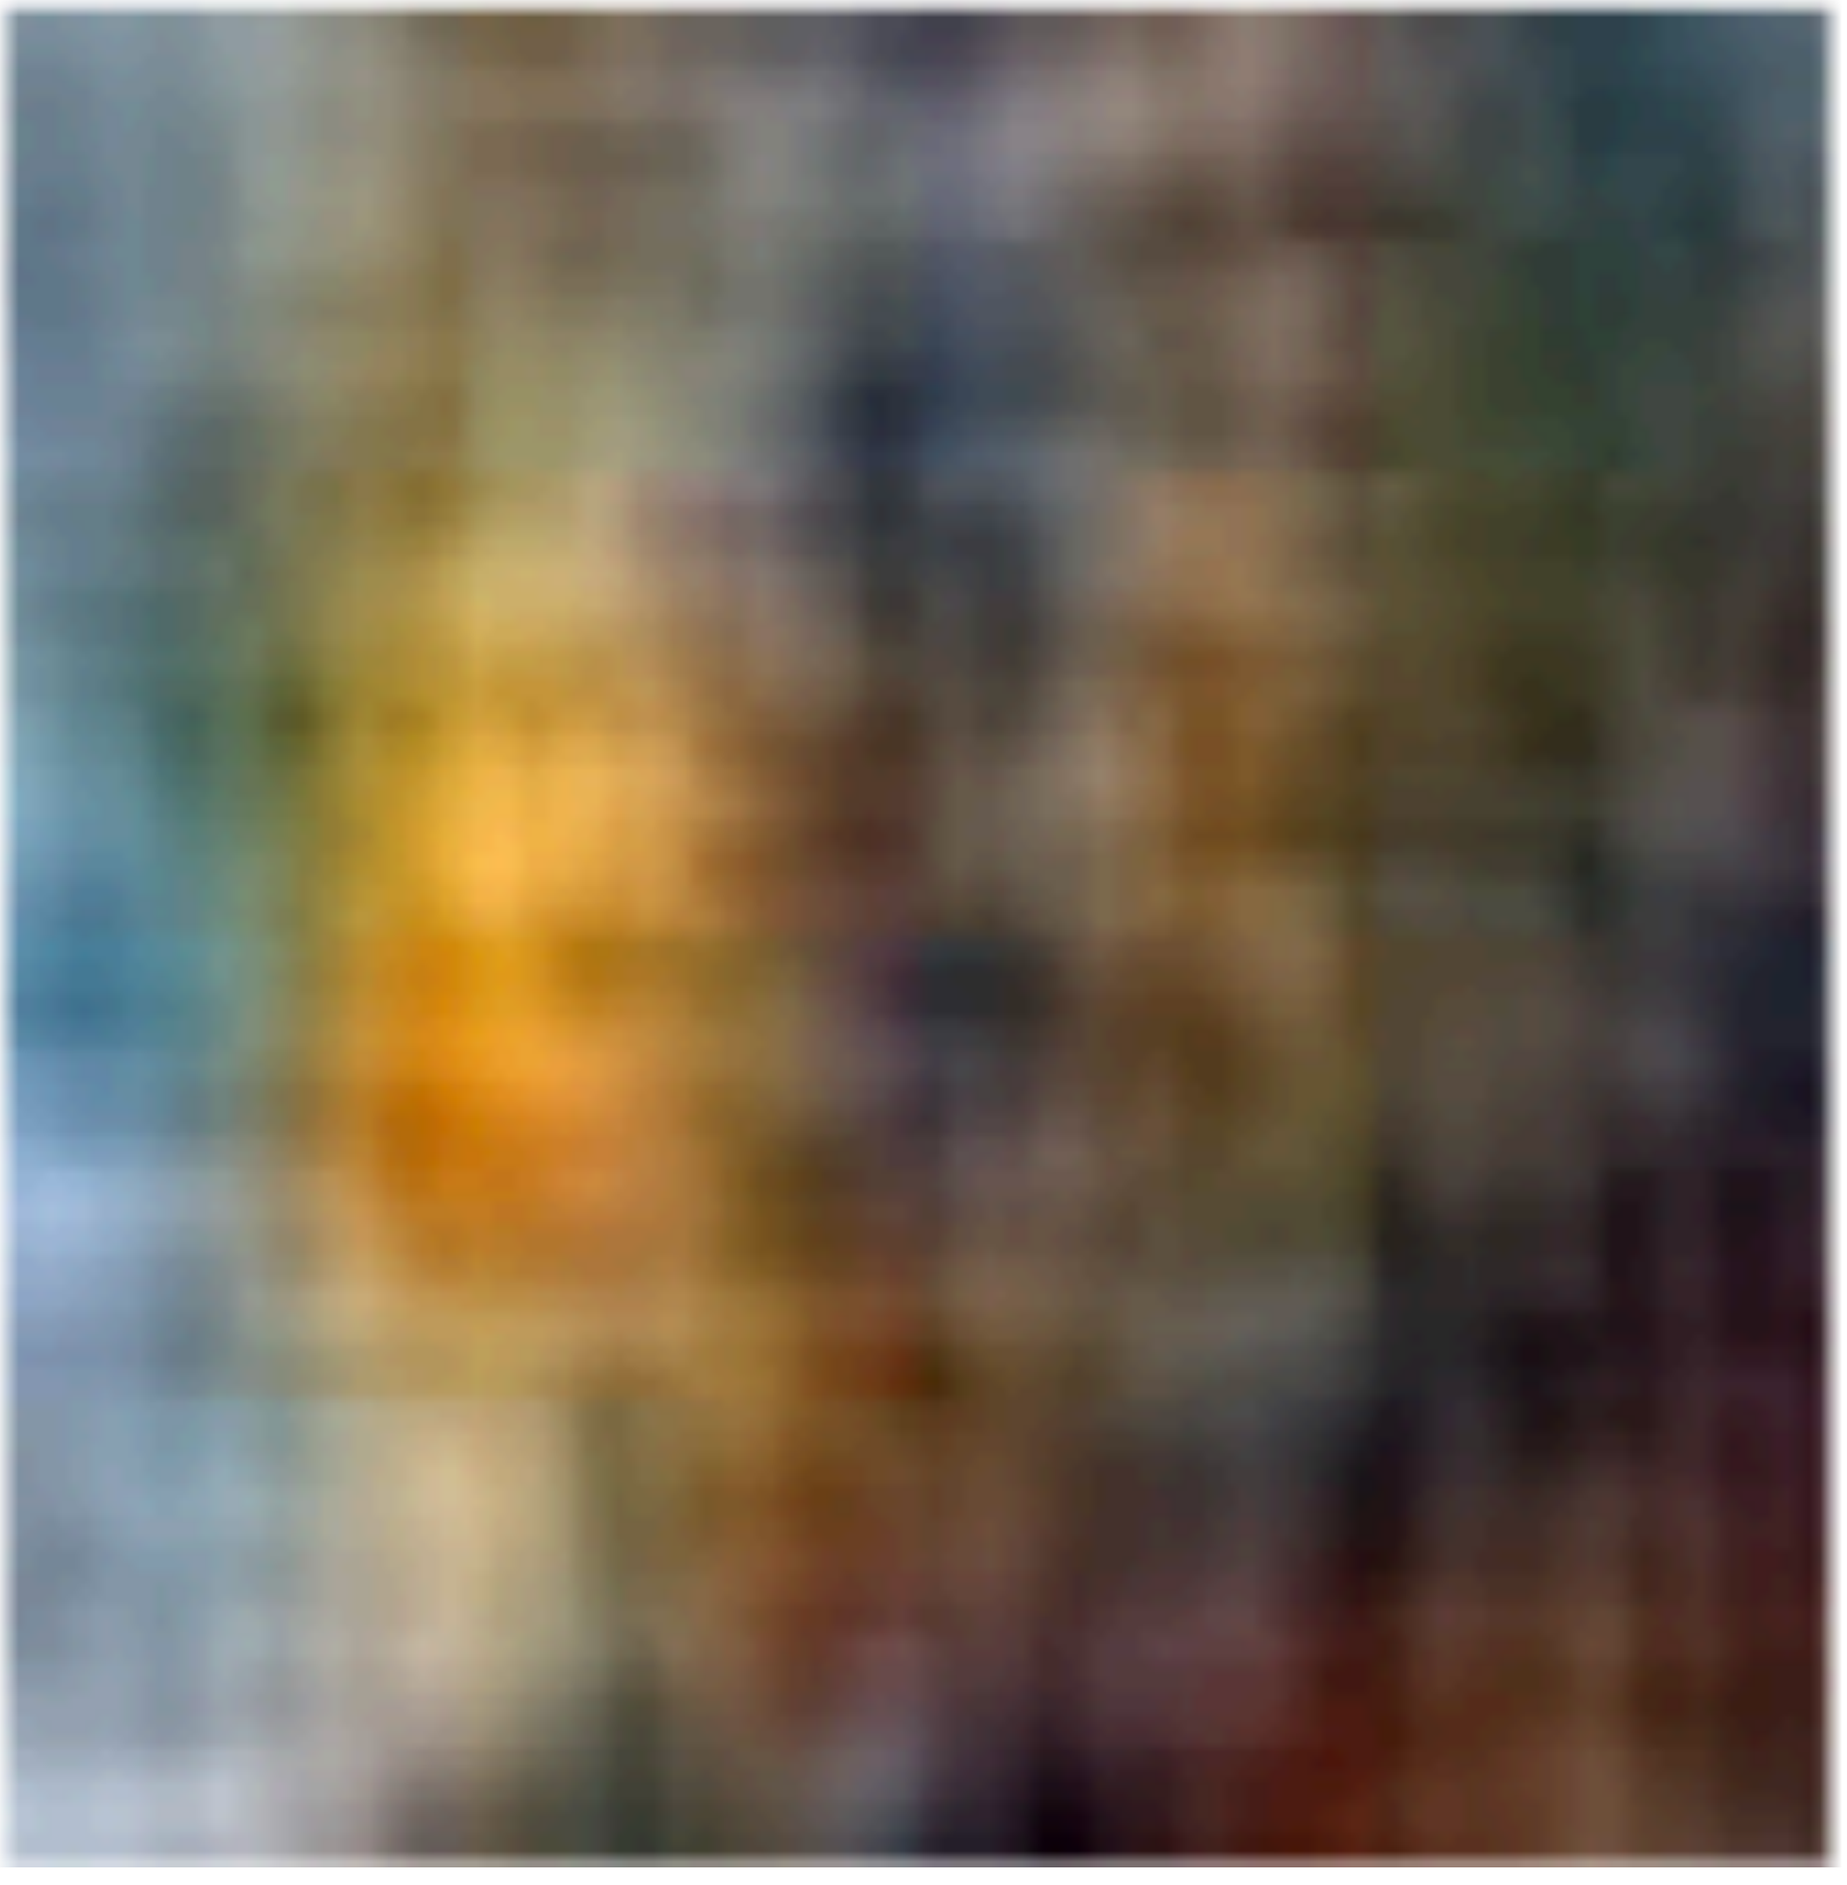

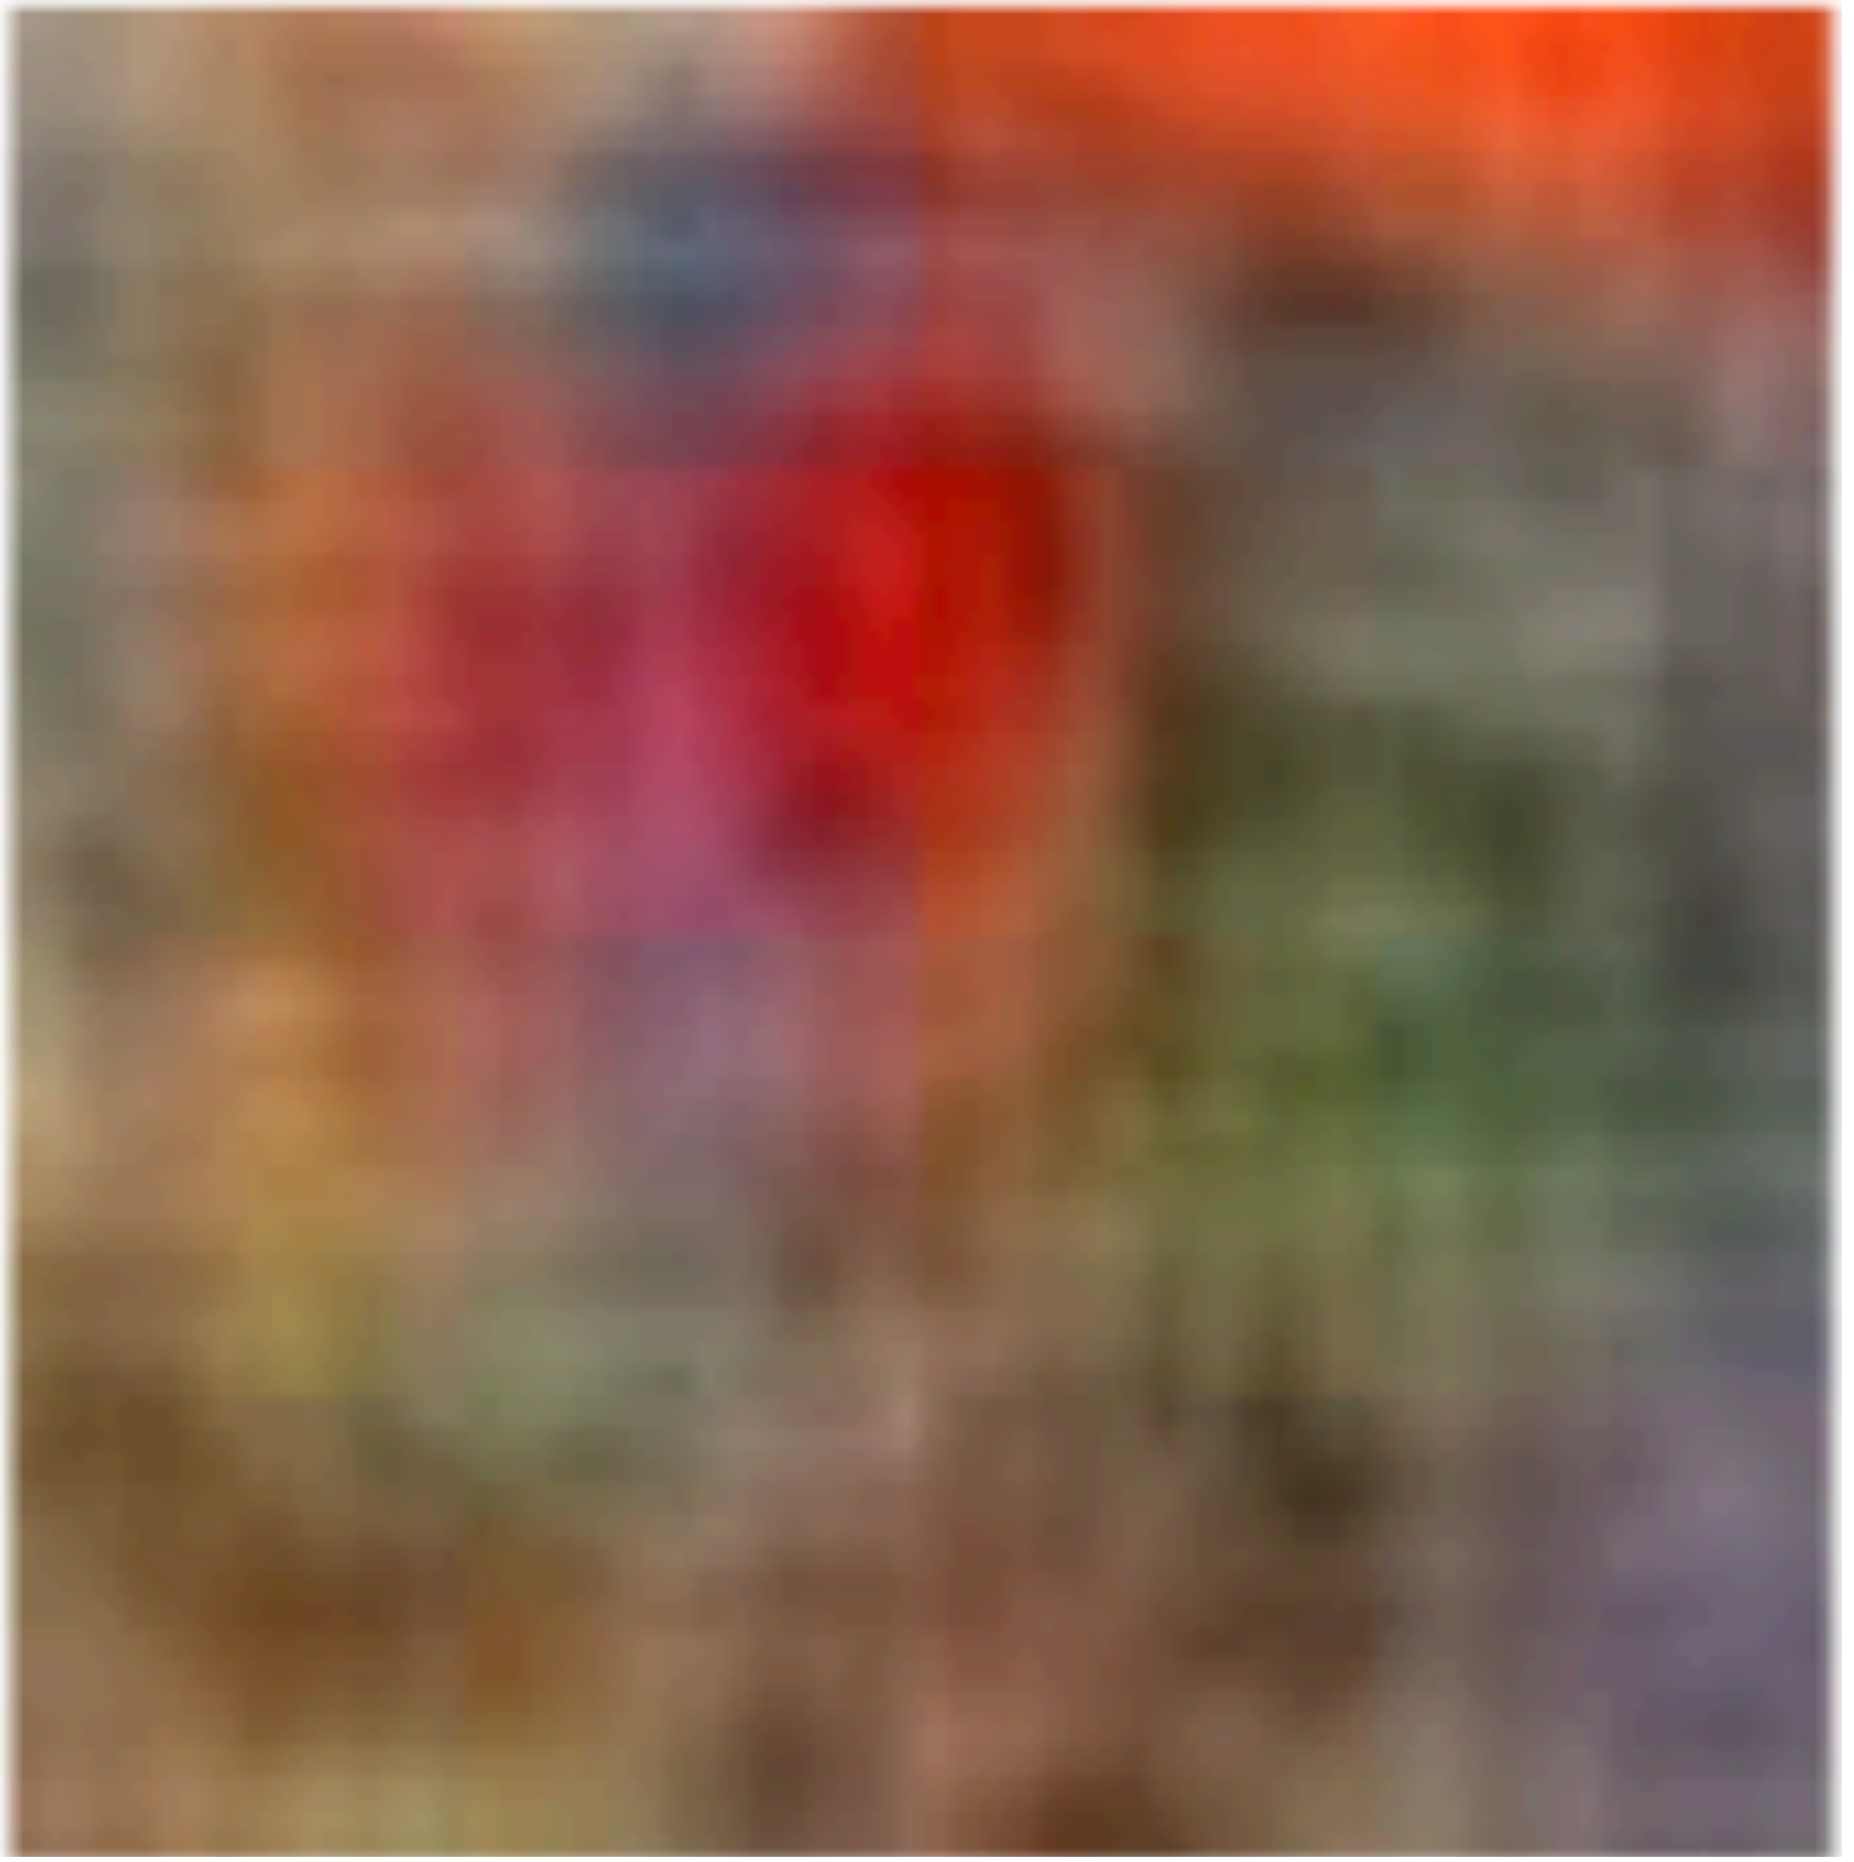

Classification Images

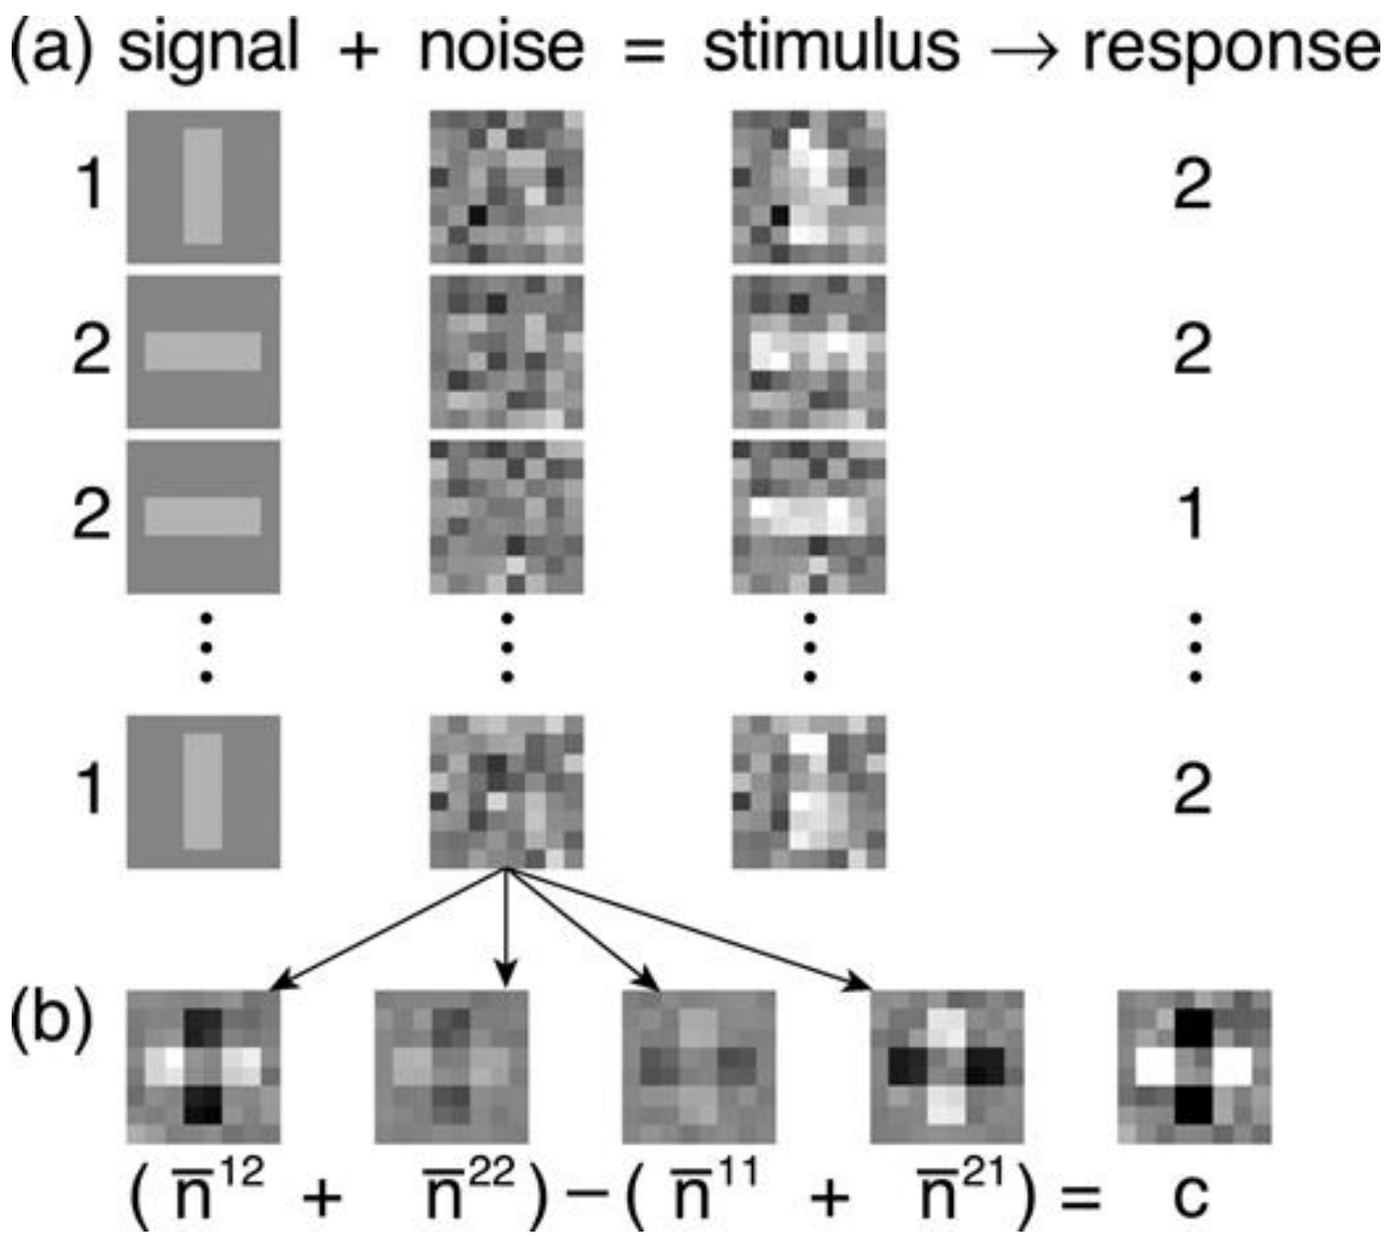

Classification images: A review. Richard F. Murray

White noise in different spaces

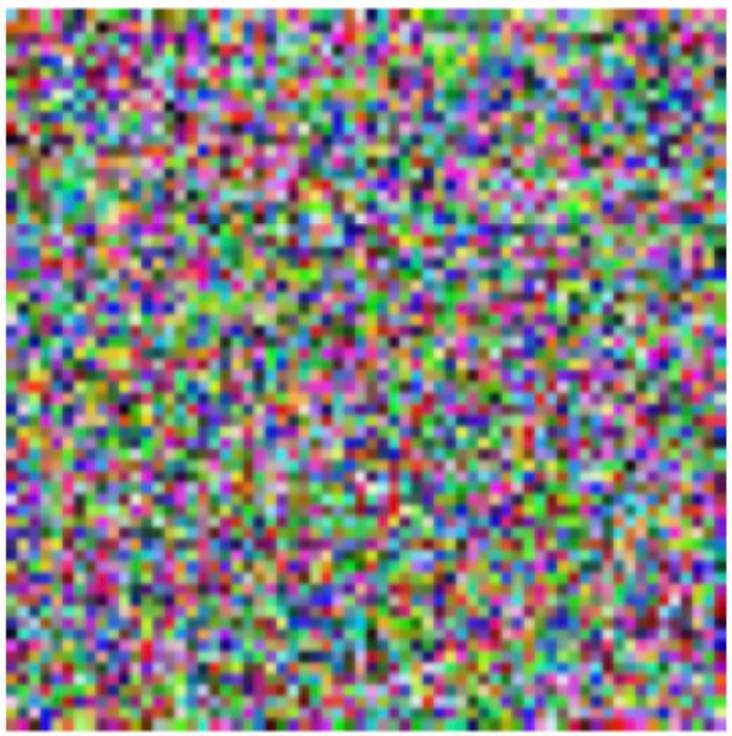

(a) RGB

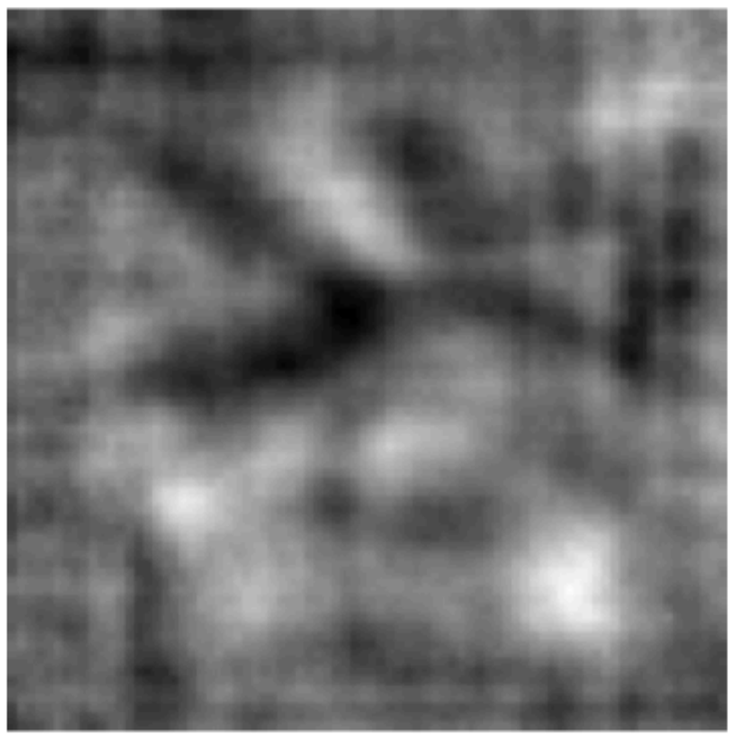

(b) HOG

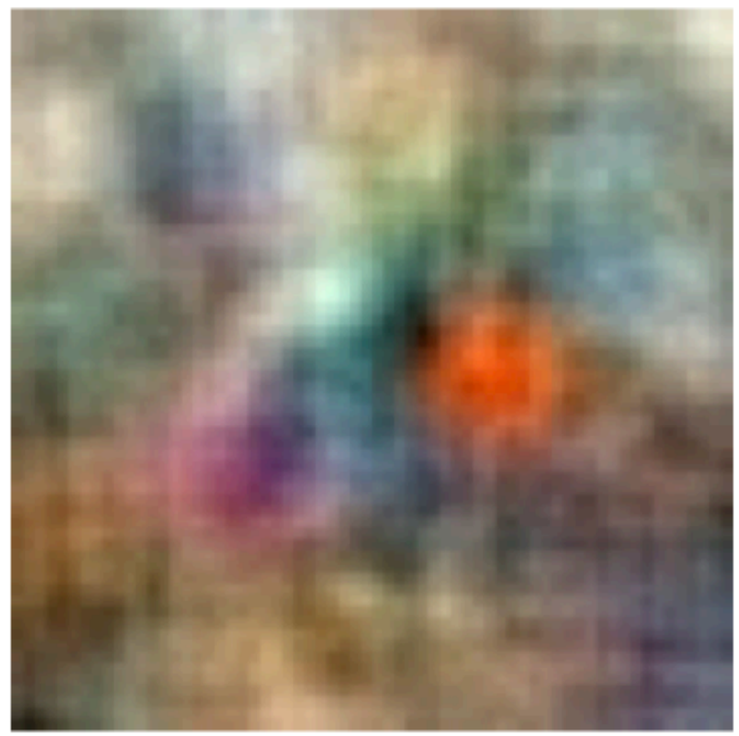

(c) CNN

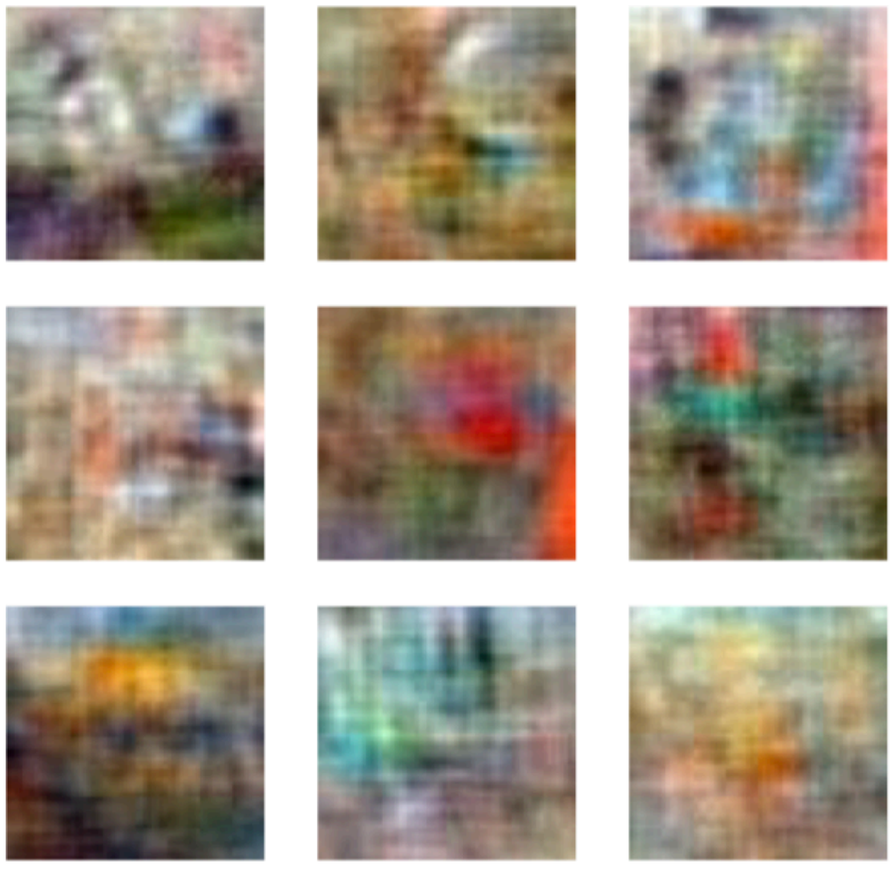

Do this 100,000 times…

“Is this a car?”

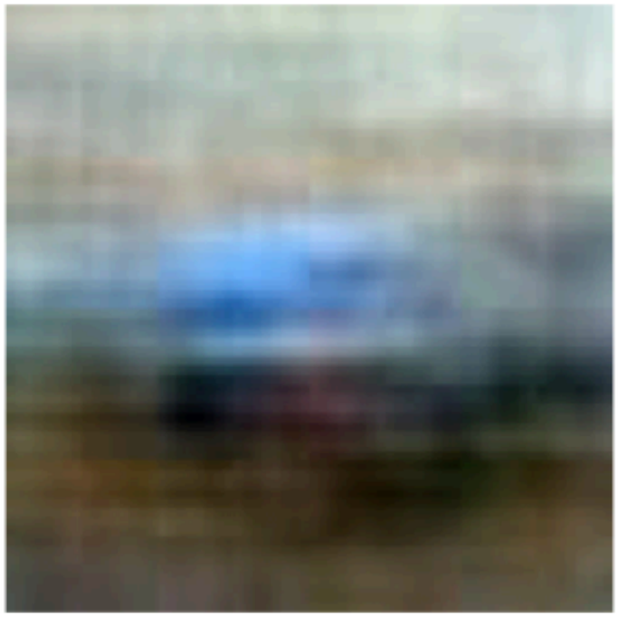

Car

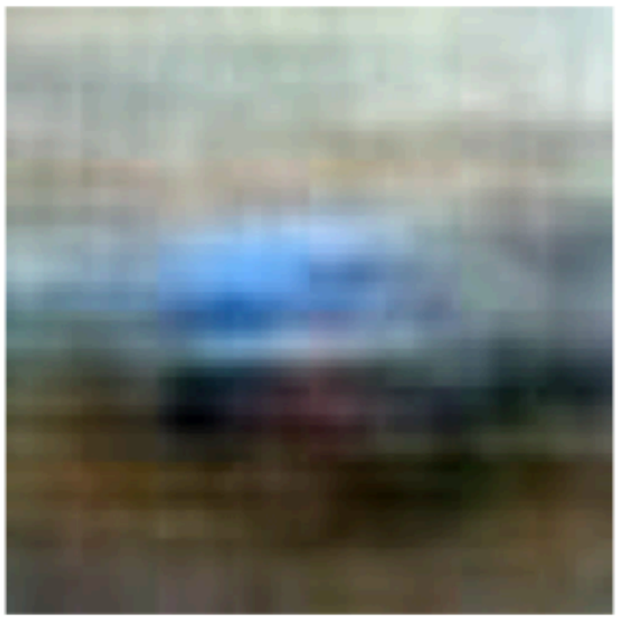

Car

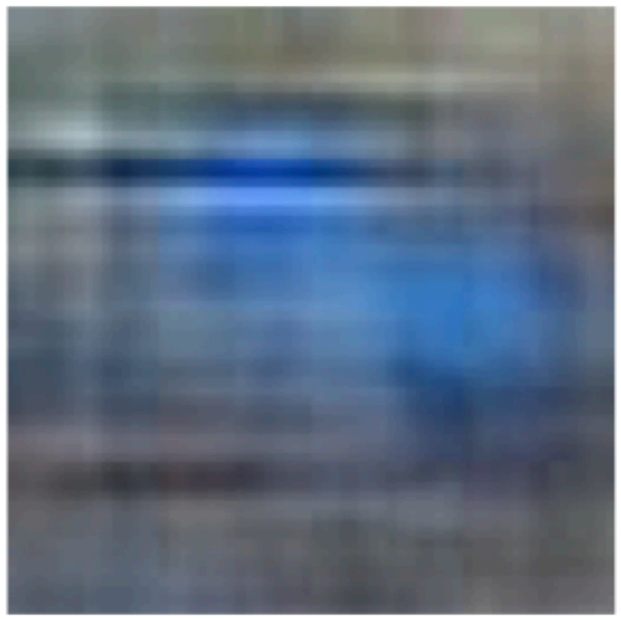

Television

Person

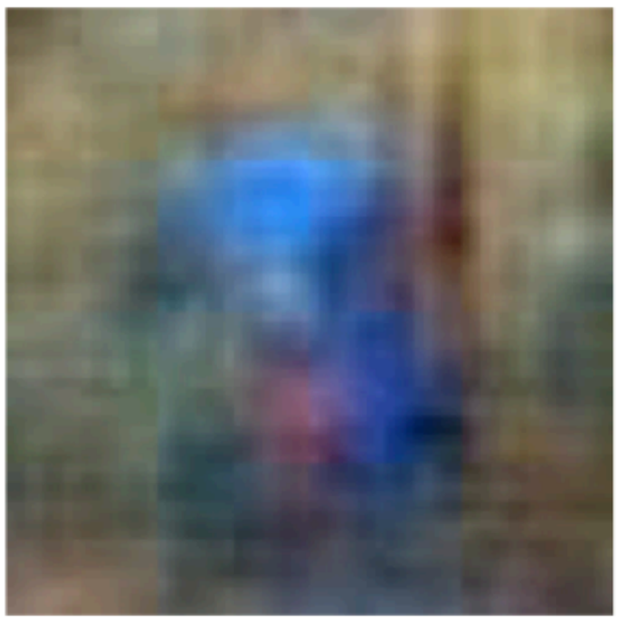

Bottle

Fire Hydrant

“Is this a sports ball?”

(a) India

(b) United States

“Is this a sports ball?”

(a) India

(b) United States

Top retrievals from classification image

Car

Person

Bottle

Fire Hydrant

Television

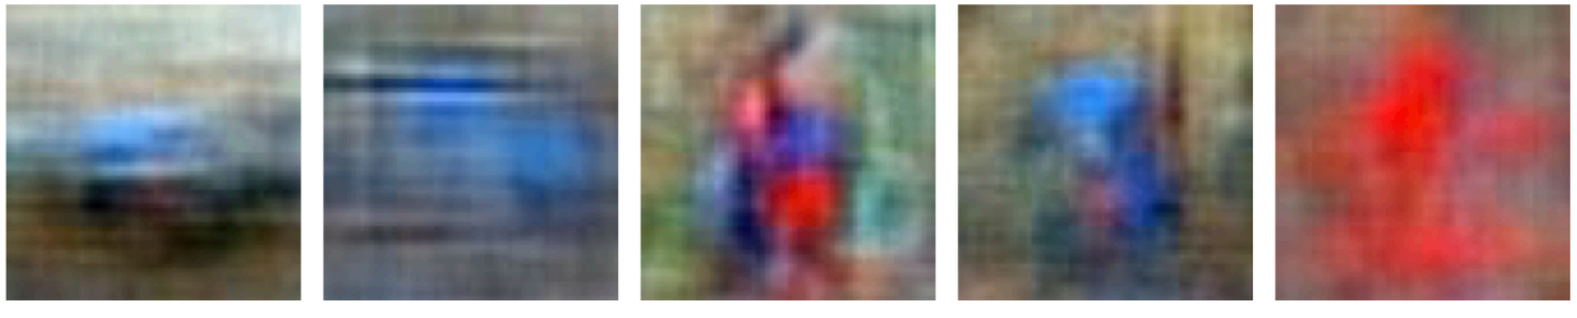

Car Television Person Bottle Fire Hydrant

Not going to beat state-of-the-art here…

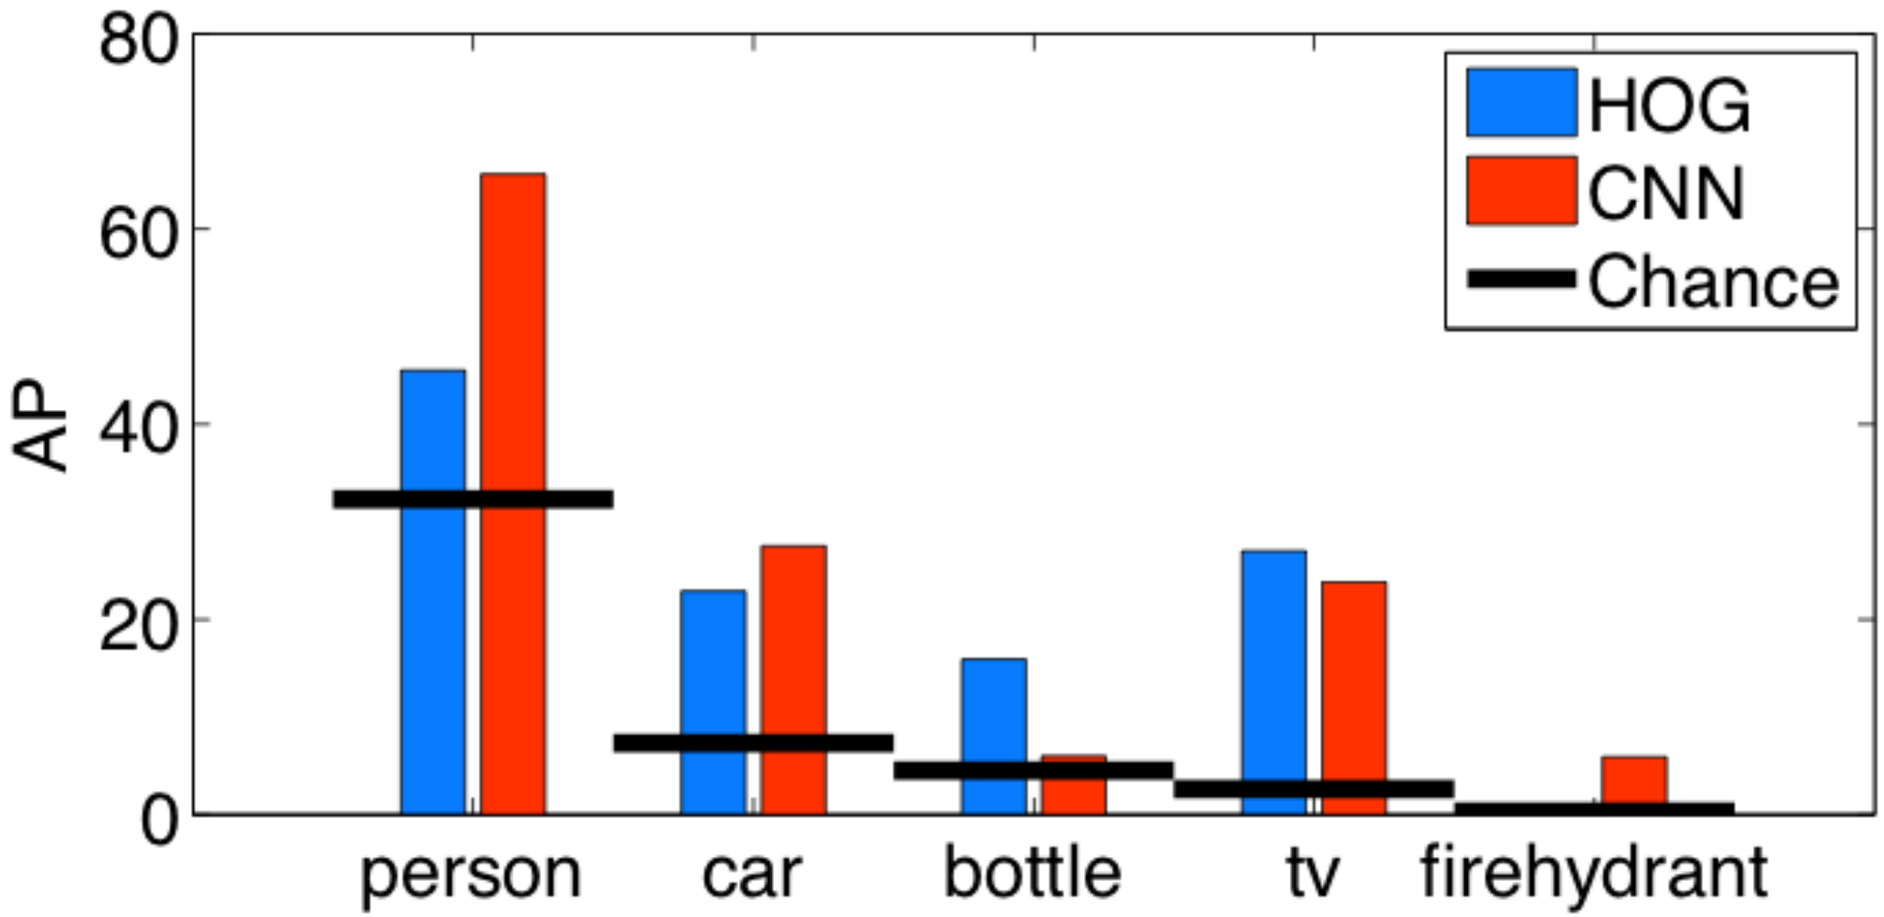

Inverting Features

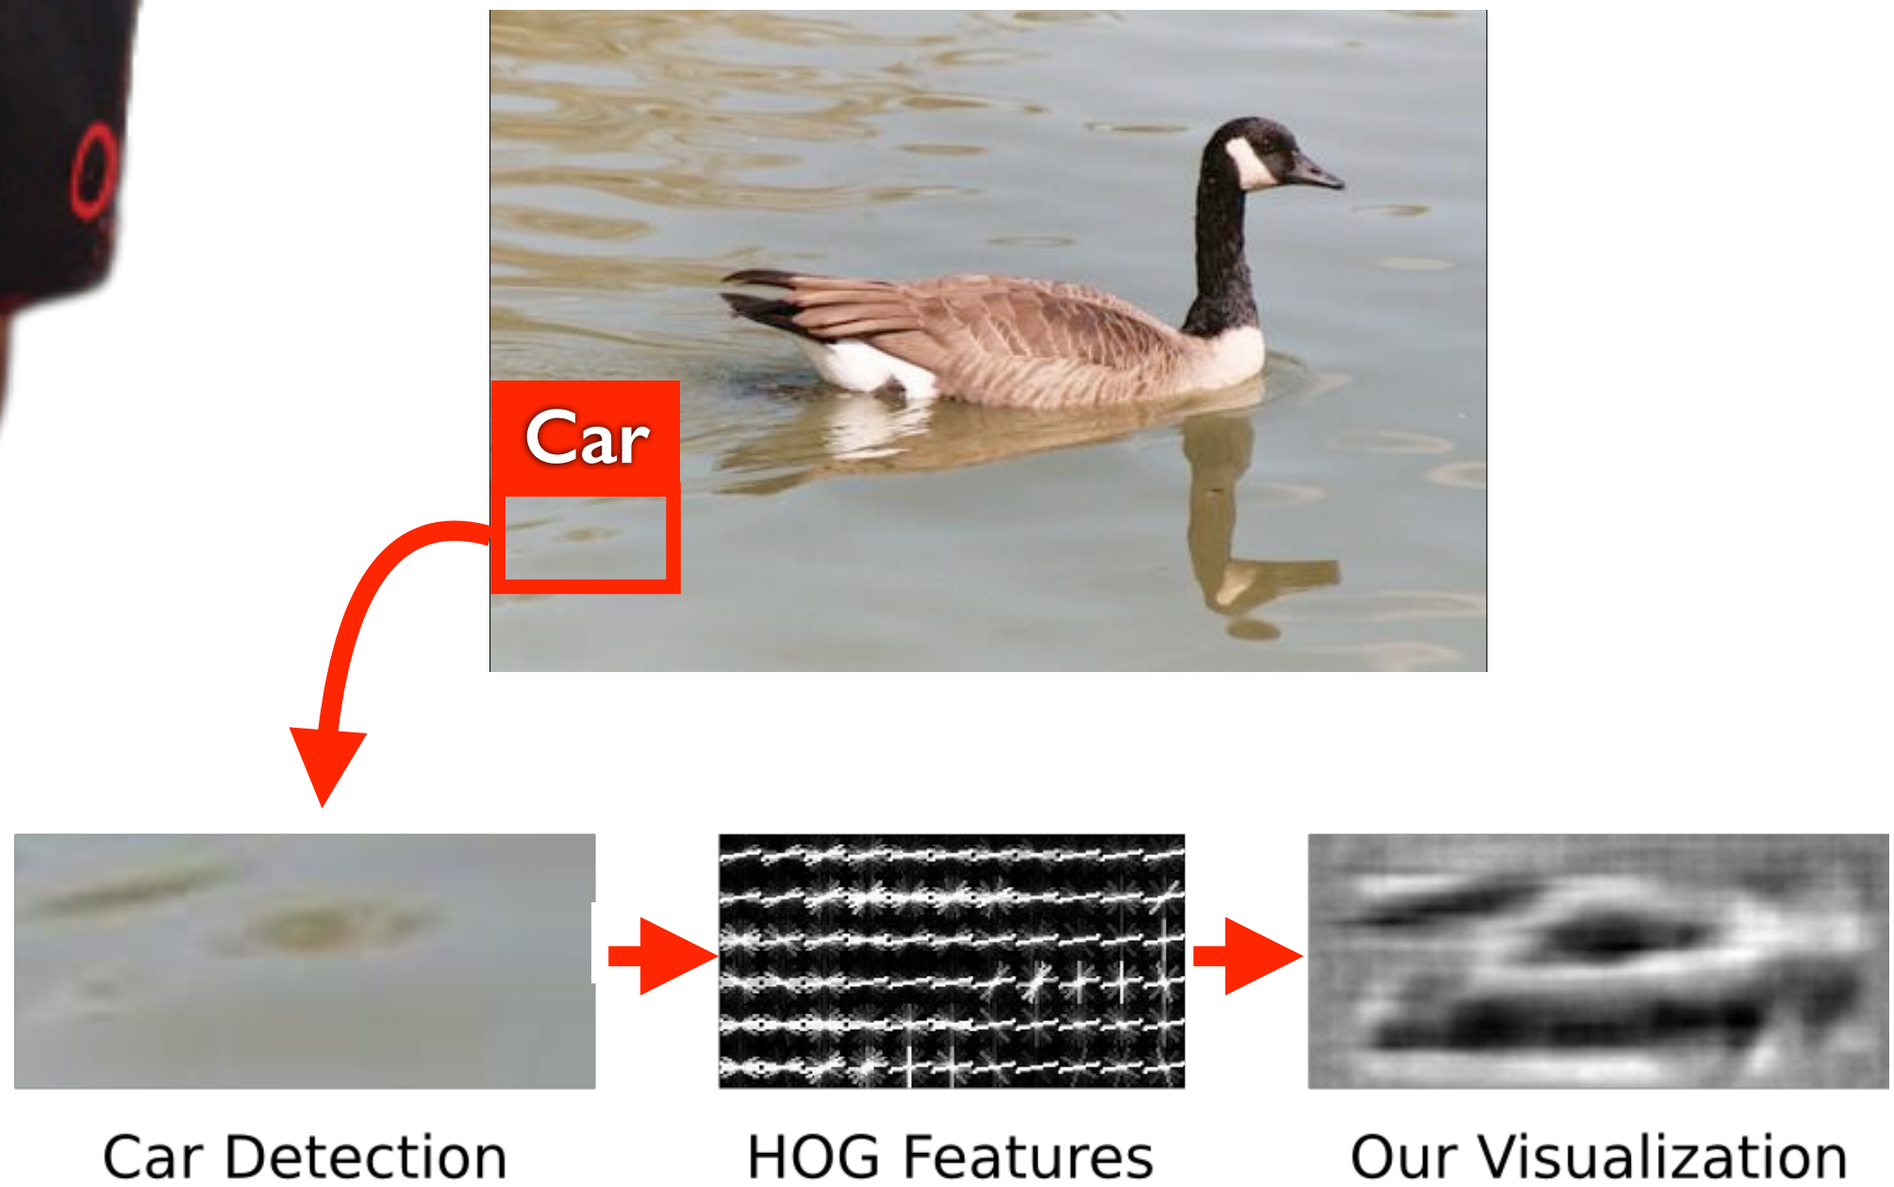

Vondrick, Khosla, Malisiewicz, Torralba. ICCV 2013

My mistake:

All these interesting detours kept cropping up, and I ignored them

The good scientist

The most exciting phrase to hear in science, the one that heralds new discoveries, is not “Eureka!” but “That’s funny…”

— Isaac Asimov

The good scientist

- Develops a hypothesis, but pivots with new data
  - Conviction to test hypothesis, but know when to refine theory

- Collects and explores tons of natural data
  - Real world data is messy, but that is key problem

- Remains curious about unusual experimental results
  - Need solid experiments so unusual is not just a bug

- Healthy dosage of self-doubt
  - And you resolve your doubt by collecting evidence

A good scientist is like a good machine learning model:

- They both fit the hypothesis to data
- They both favor the simple hypothesis (Occam's razor)

Example: ResNet

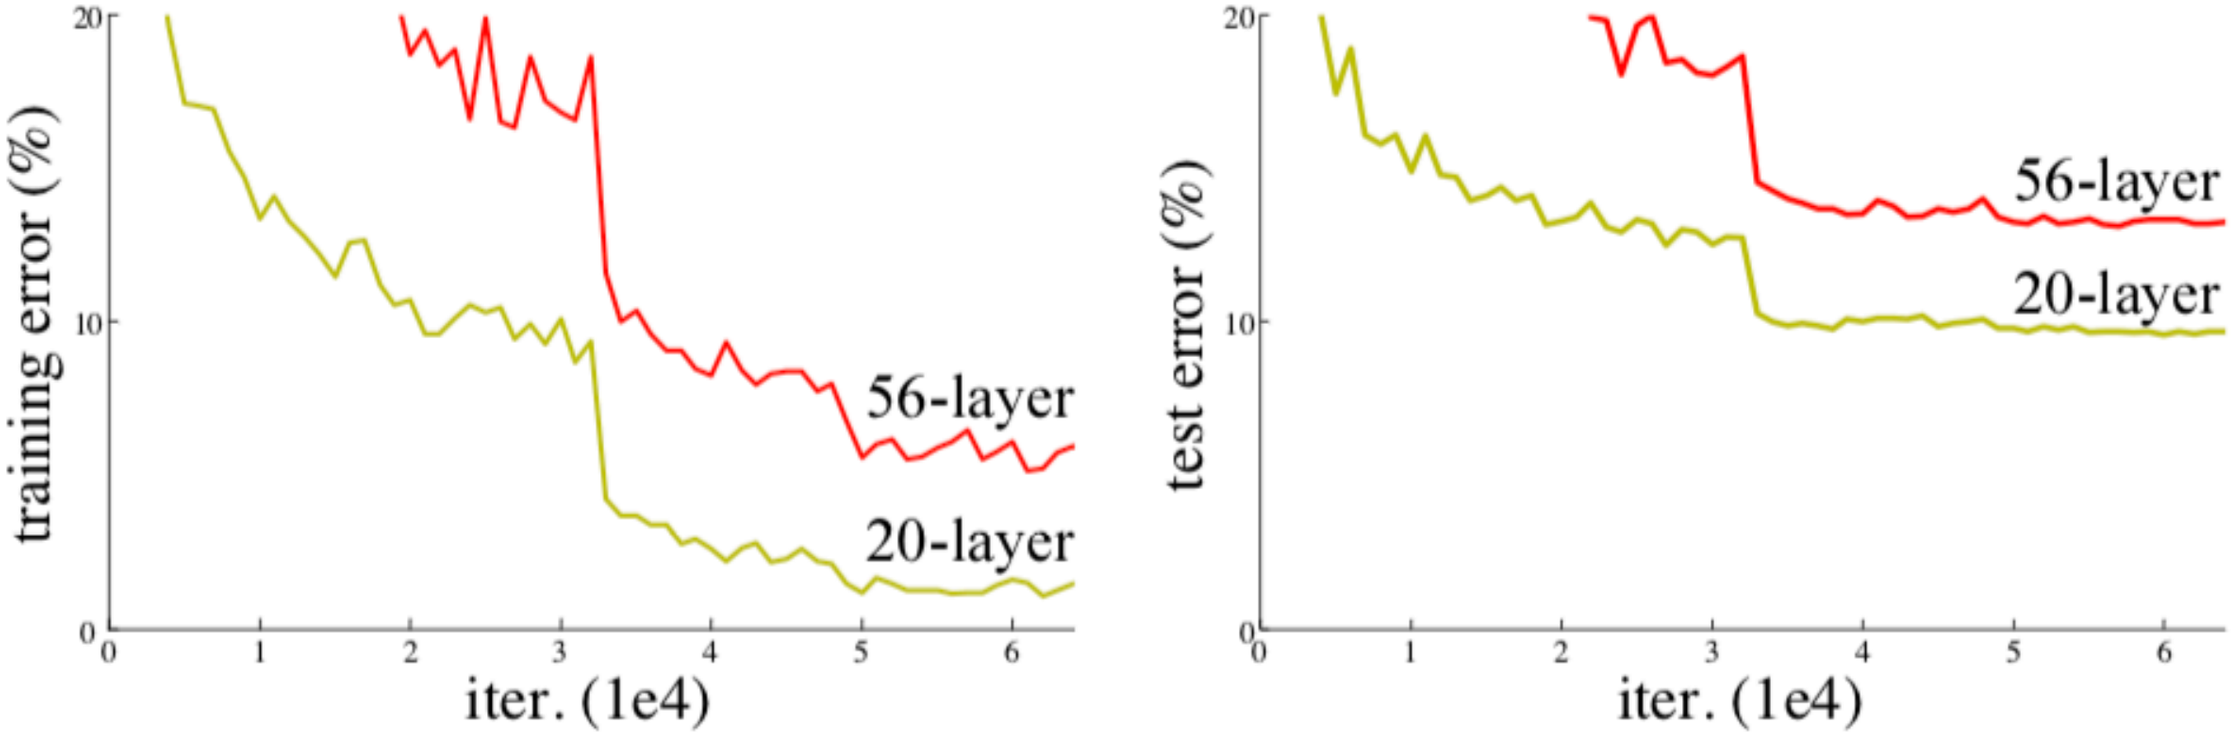

Figure 1. Training error (left) and test error (right) on CIFAR-10 with 20-layer and 56-layer "plain" networks. The deeper network has higher training error, and thus test error. Similar phenomena on ImageNet is presented in Fig. 4.

My experience in getting computer vision to work

- Start with an idea — Bigger data! Deeper models!
- Try very, very hard to get it work.
- Discover something unusual or curious. If you don’t find anything unusual, you haven’t tried hard enough.
- Isolate the unusual thing. Use simple experiments and clear visualizations. Study it. Make sure not a bug.
- Capitalize on it. You might give up your original idea, and that’s ok.

How to find unusual things

- Get very familiar with your data
- Create lots of qualitative visualizations
- Collect lots of numbers
- New lenses to view data have historically lead to revolutions

What to do with a negative result?

- ~~Don’t tell anyone~~
- You need to answer:
  - Why doesn’t it work?
  - What are the implications of this not working?
- Tell people & me that

Paper and Report Writing

Many slides from Bill Freeman

A paper's impact on your career

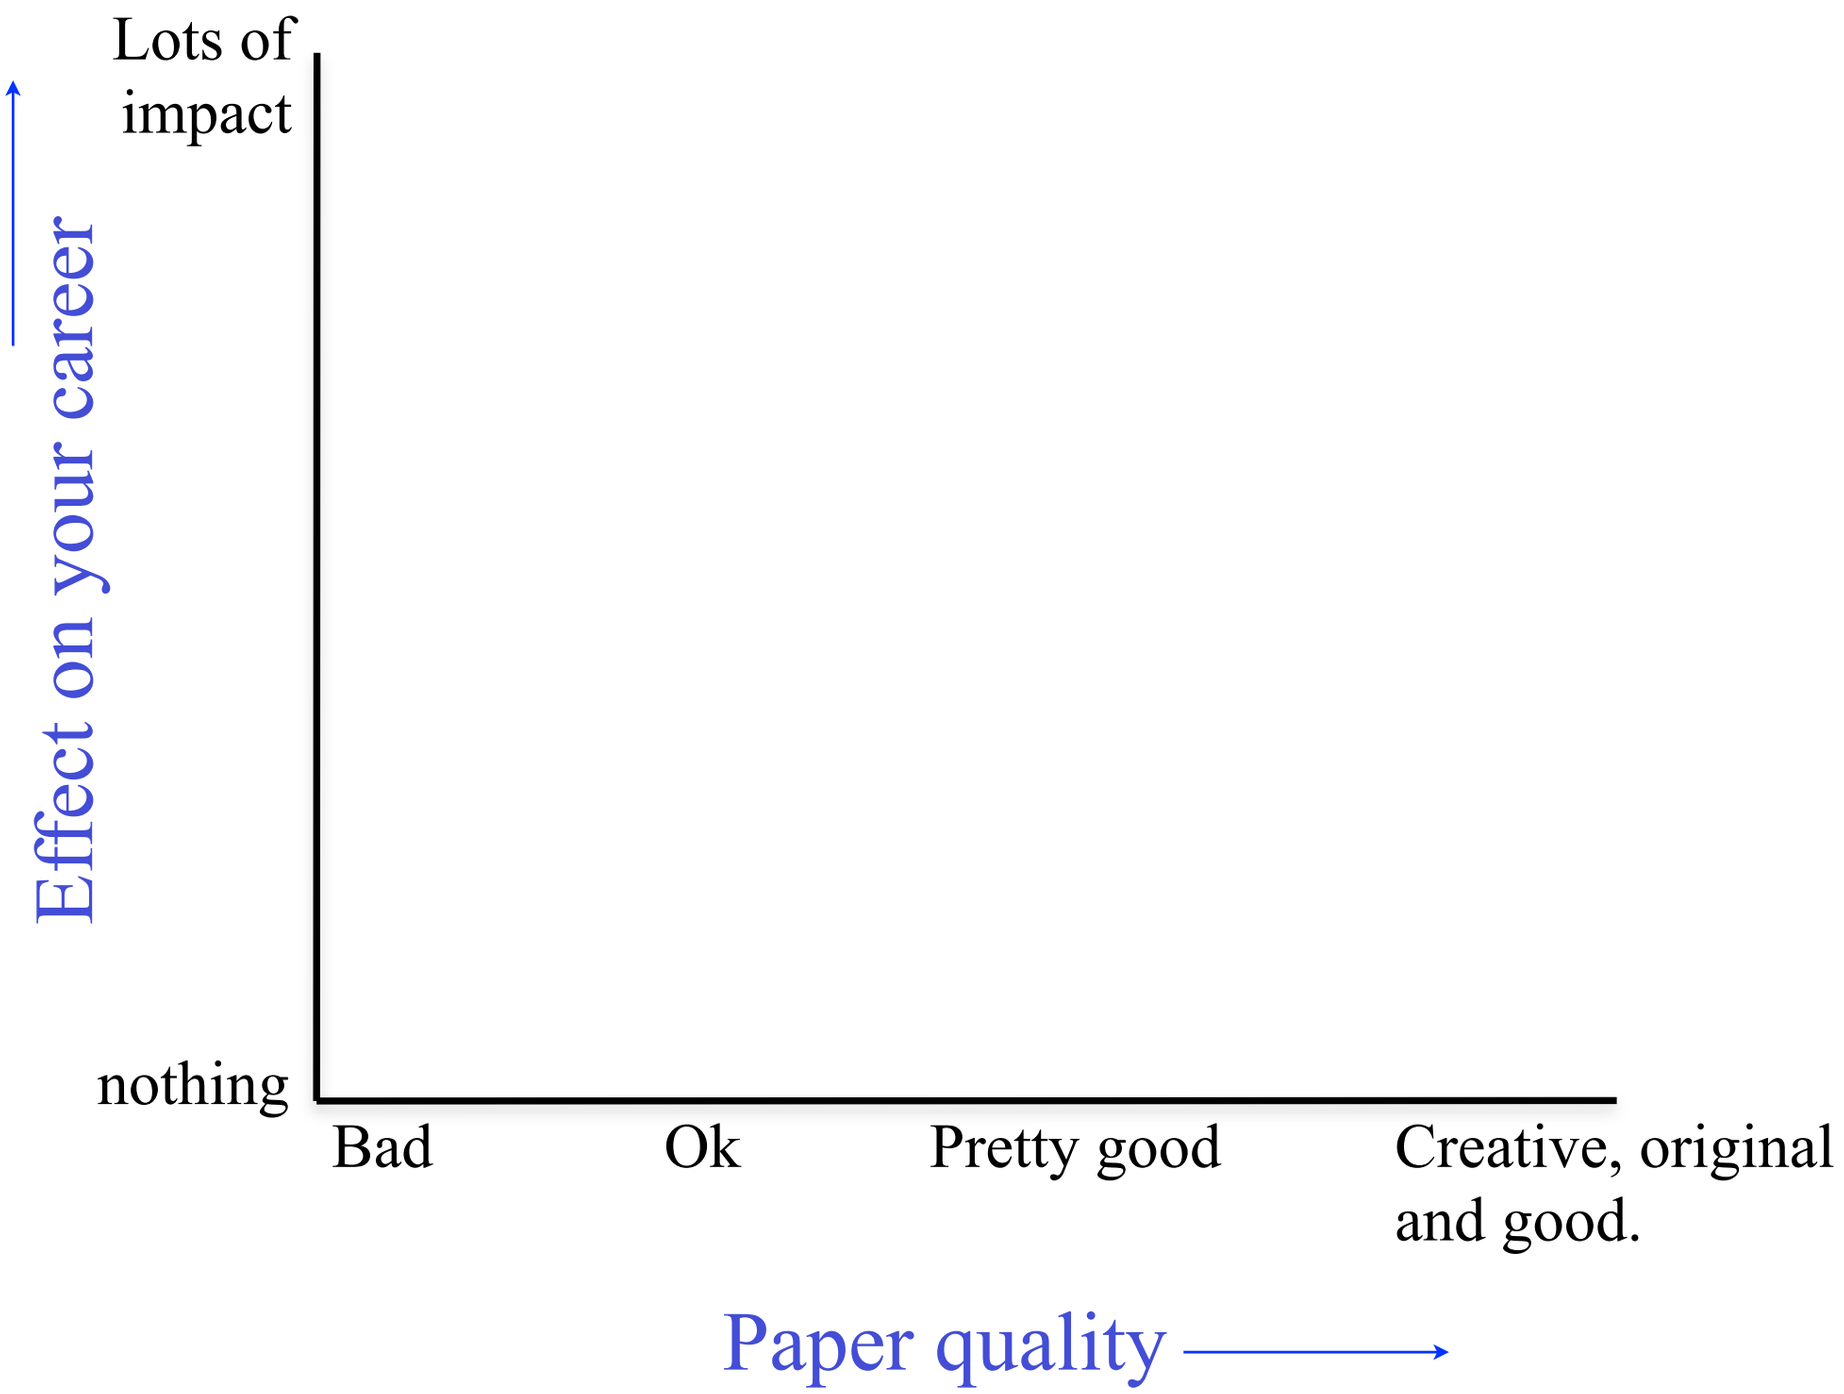

Slide credit: Bill Freeman

A paper's impact on your career

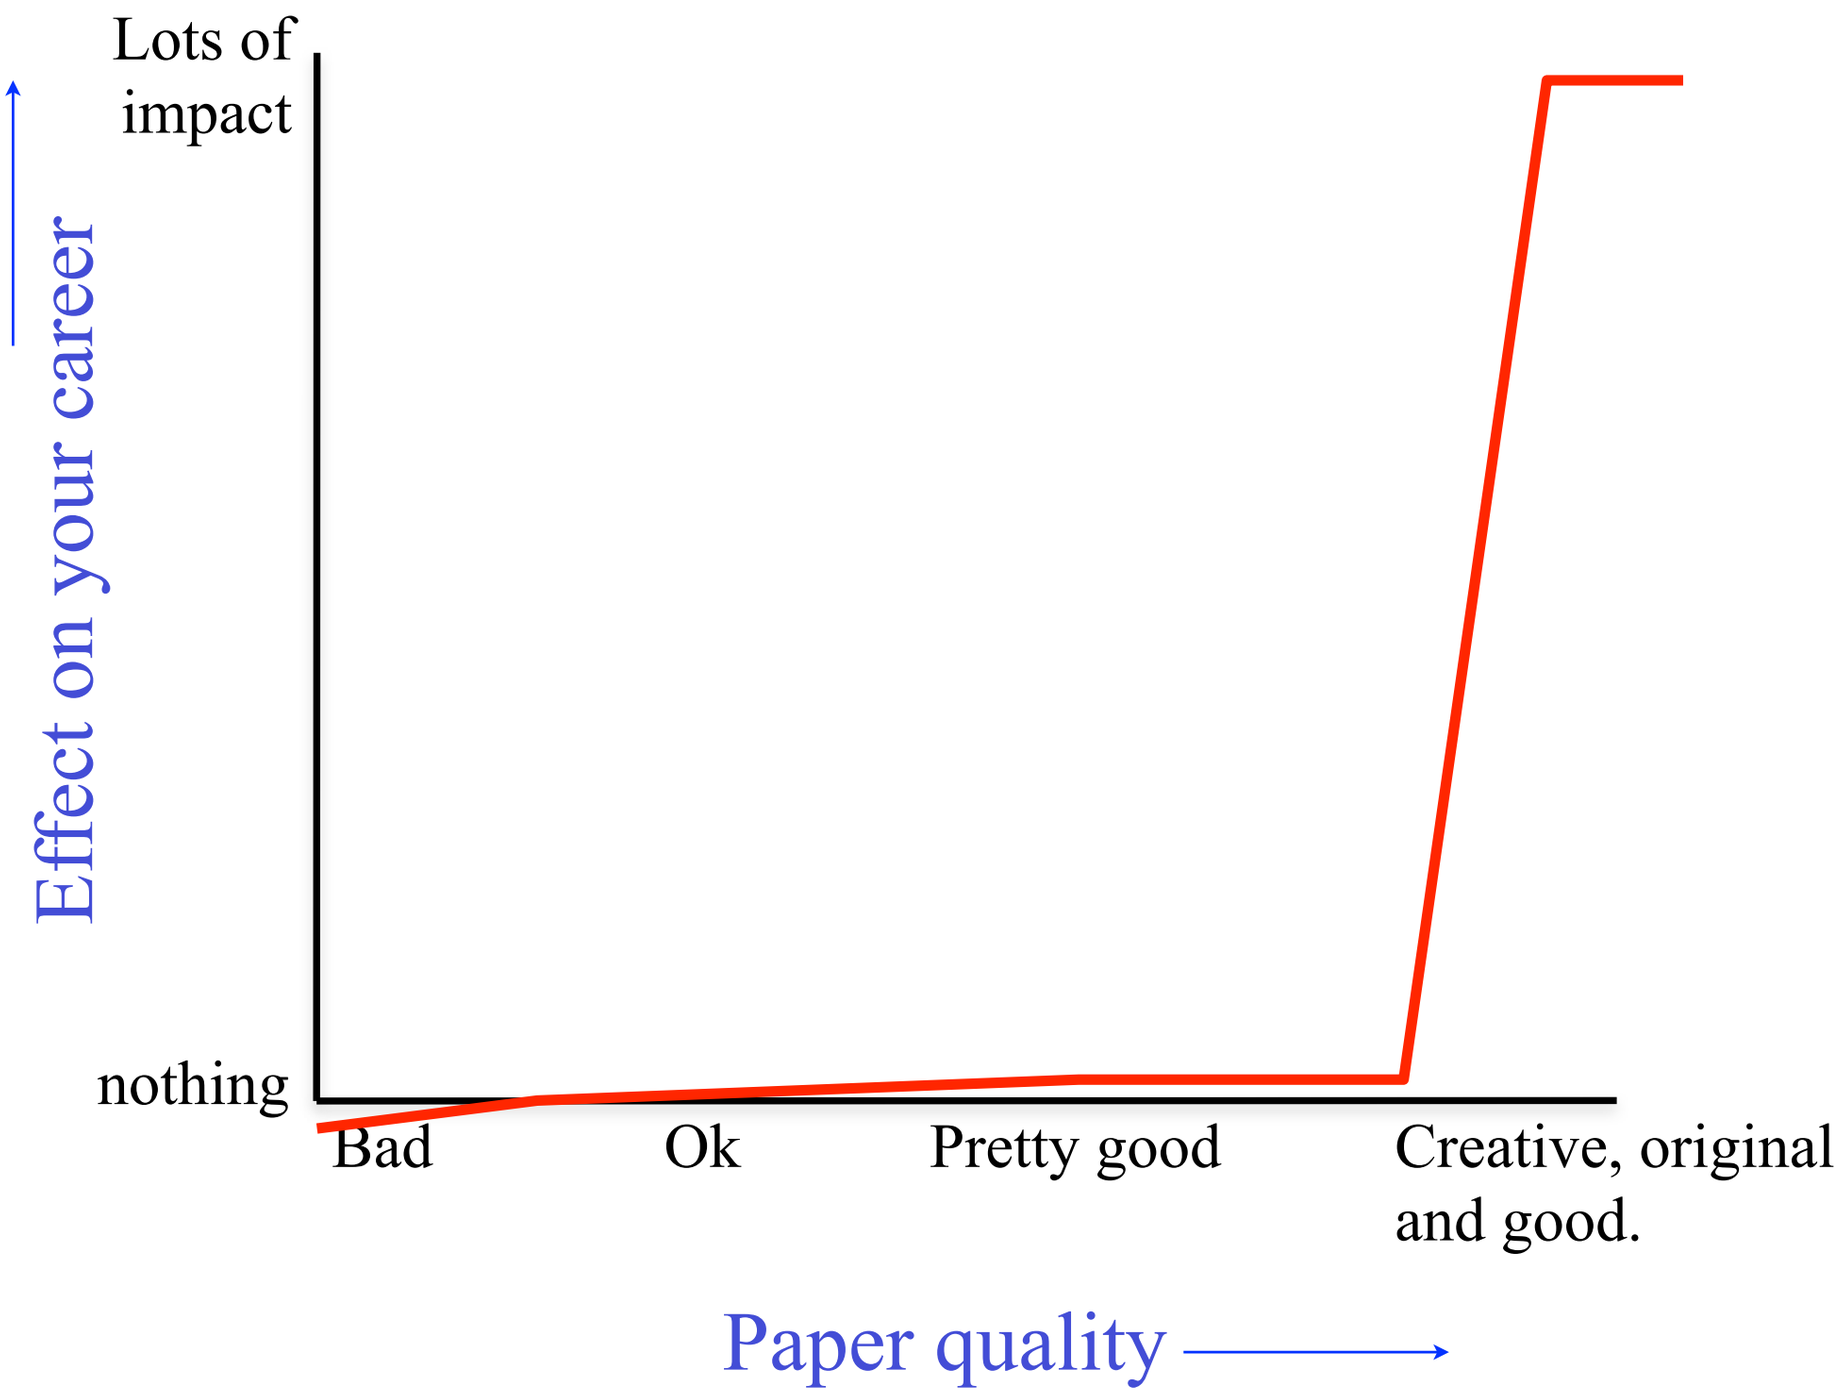

Slide credit: Bill Freeman

Our image of the research community

- Scholars, plenty of time on their hands, pouring over your manuscript.

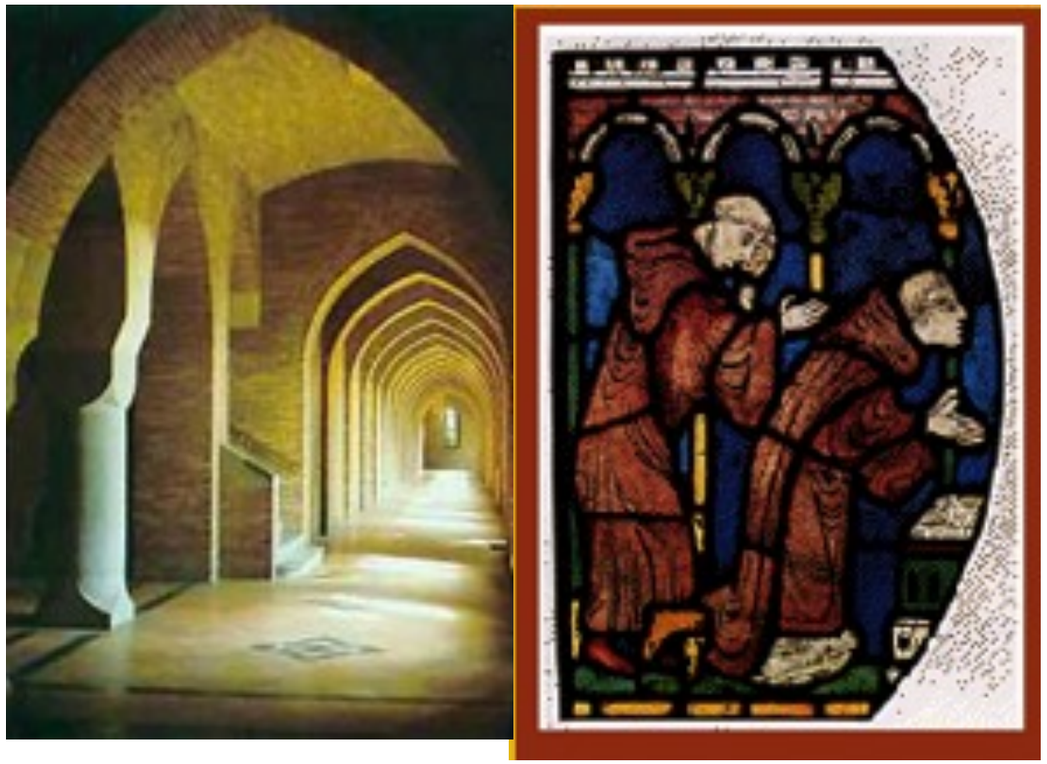

Slide credit: Bill Freeman

The reality:
more like a large, crowded marketplace

http://ducksflytogether.wordpress.com/2008/08/02/looking-back-khan-el-khalili/

Slide credit: Bill Freeman

Paper Organization

- Introduction
- Related Work
- Method
- Experiments
- Discussion

Paper Organization

- Introduction: motivation, what you will do
- Related Work: what has been tried before
- Method: clearly explain main idea
- Experiments: evidence for the idea
- Discussion: so what? larger implications

Ted Adelson on paper organization.

(1) Start by stating which problem you are addressing, keeping the audience in mind. They must care about it, which means that sometimes you must tell them why they should care about the problem.

(2) Then state briefly what the other solutions are to the problem, and why they aren't satisfactory. If they were satisfactory, you wouldn't need to do the work.

(3) Then explain your own solution, compare it with other solutions, and say why it's better.

(4) At the end, talk about related work where similar techniques and experiments have been used, but applied to a different problem.

Since I developed this formula, it seems that all the papers I've written have been accepted. (told informally, in conversation, 1990).

Slide credit: Bill Freeman

Treat the reader as you would a guest in your house

Anticipate their needs: would you like something to drink?
Something to eat? Perhaps now, after eating, you'd like to rest?

Slide credit: Bill Freeman