

# Visipedia: collaborative harvesting and organization of visual knowledge

Pietro Perona  
California Institute of Technology

2014 Signal and Image Sciences Workshop  
Lawrence Livermore National Laboratory  
21 May, 2014







WIKIPEDIA  
The Free Encyclopedia

### navigation

- [Main page](#)
- [Contents](#)
- [Featured content](#)
- [Current events](#)
- [Random article](#)

### search

Go

Search

### interaction

- [About Wikipedia](#)
- [Community portal](#)
- [Recent changes](#)
- [Contact Wikipedia](#)
- [Donate to Wikipedia](#)
- [Help](#)

### toolbox

- [What links here](#)
- [Related changes](#)
- [Upload file](#)
- [Special pages](#)
- [Printable version](#)
- [Permanent link](#)
- [Cite this page](#)

# Beak

From Wikipedia, the free encyclopedia

*For other uses, see [Beak \(disambiguation\)](#).*

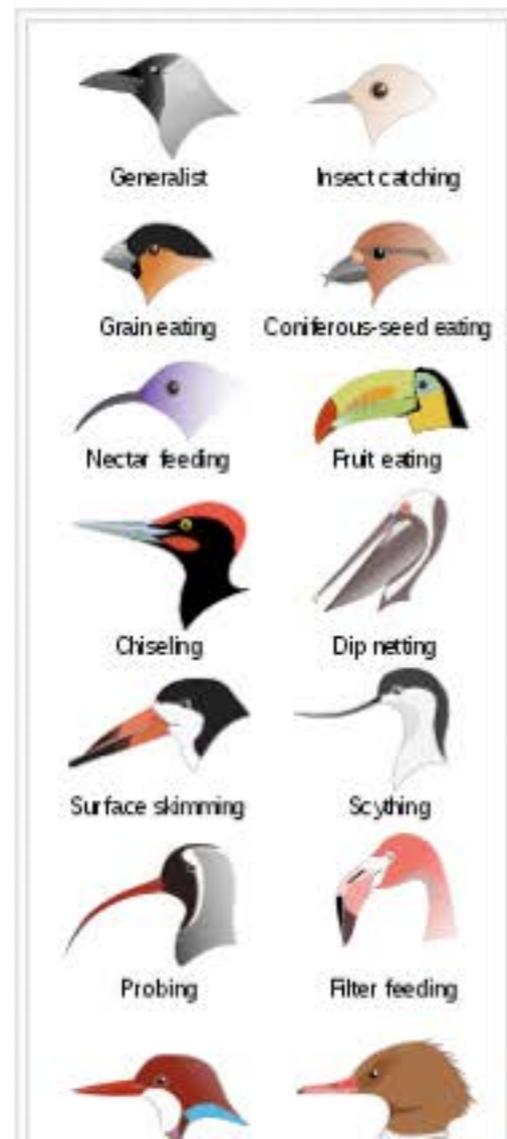
The **beak**, **bill** or **rostrum** is an external anatomical structure of [birds](#) which, in addition to eating, is used for [grooming](#), manipulating objects, killing prey, probing for food, [courtship](#) and feeding their young. The term also refers to a similar mouthpart in some [cephalopods](#), [cetaceans](#), [pufferfishes](#), [turtles](#), [Anuran tadpoles](#) and [sirens](#).



The bill of a scavenger—the [vulture](#).



The bill and knob of a [domesticated Chinese goose](#), highly exaggerated by farm selection.



## Contents [\[hide\]](#)

- [1 Anatomy](#)
- [2 Billing](#)
- [3 Terminology](#)
- [4 See also](#)
- [5 References](#)
- [6 External links](#)

## Anatomy [\[edit\]](#)

Beaks can vary significantly in size and shape from species to species. The beak is composed of an upper jaw called the [maxilla](#), and a lower jaw called the [mandible](#). The jaw is made





**WIKIPEDIA**  
*The Free Encyclopedia*

[article](#)

[discussion](#)

[edit this page](#)

[history](#)

# Cere

From Wikipedia, the free encyclopedia

The **cere** (from the Latin *cera*: wax)<sup>[1]</sup> or **operculum**<sup>[2]</sup> is a soft, fleshy swelling found on the beaks of certain birds. [Hawks](#), [parrots](#), [doves](#), [skuas](#) and [budgerigars](#) are among the birds that have ceres. The word 'cere' is often used synonymously with 'beak', although the two are not identical. The cere plays a role in indicating the reproductive stage of certain [dimorphic](#) birds, and also has a key function in respiration.

## Contents [\[hide\]](#)

- [1 Physical characteristics](#)
- [2 Role in respiration](#)
- [3 Role in indication of reproductive cycle](#)
- [4 References](#)
- [5 See also](#)

## Physical characteristics [\[edit\]](#)

The cere is located at the top of the beak,<sup>[2][3]</sup> and is often waxy in texture. The colour of the cere may vary from species to species, and also depends on the season.<sup>[4]</sup> The cere contains the [nares](#) (nostrils). The shape of the cere varies from species to species. In falcons, the opening of the nares is roughly circular in shape.<sup>[2]</sup>



The [Rock Pigeon](#), *Columba livia*. The cere is located at the top of the beak, and is a distinctive white-grey colour.

### navigation

- [Main page](#)
- [Contents](#)
- [Featured content](#)
- [Current events](#)
- [Random article](#)

### search




### interaction

- [About Wikipedia](#)
- [Community portal](#)
- [Recent changes](#)
- [Contact Wikipedia](#)
- [Donate to Wikipedia](#)
- [Help](#)

### toolbox

- [What links here](#)
- [Related changes](#)
- [Upload file](#)
- [Special pages](#)
- [Printable version](#)

from **John Stevenson** <solarjack@gmail.com>  
to Pietro Perona <perona@vision.caltech.edu>  
date Thu, Aug 6, 2009 at 7:50 AM  
subject Mushrooms galore at Caumsett

Can I eat this one? - Love John

Sent from my iPhone





WIKIPEDIA  
The Free Encyclopedia

[Learn more about citing Wikipedia.](#)

[Try Beta](#)

[Log in / create account](#)

[article](#)

[discussion](#)

[view source](#)

[history](#)

# Mushroom



From Wikipedia, the free encyclopedia

*For other uses, see [Mushroom \(disambiguation\)](#).*

A **mushroom** is the fleshy, [spore-bearing fruiting body](#) of a [fungus](#), typically produced above ground on soil or on its food source. The standard for the name "mushroom" is the cultivated white button mushroom, *Agaricus bisporus*, hence the word mushroom is most often applied to those fungi ([Basidiomycota](#), [Agaricomycetes](#)) that have a stem (*stipe*), a cap (*pileus*), and gills (lamellae, sing. [lamella](#)) on the underside of the cap, just as do store-bought white mushrooms.

The word "mushroom" can also be used for a wide variety of gilled fungi, with or without stems, and the term is used even more generally, to describe both the fleshy fruiting bodies of some [Ascomycota](#) and the woody or leathery fruiting bodies of some [Basidiomycota](#).



The mushroom *Amanita muscaria*, commonly

navigation

- [Main page](#)
- [Contents](#)
- [Featured content](#)
- [Current events](#)
- [Random article](#)

search

[Go](#)

[Search](#)

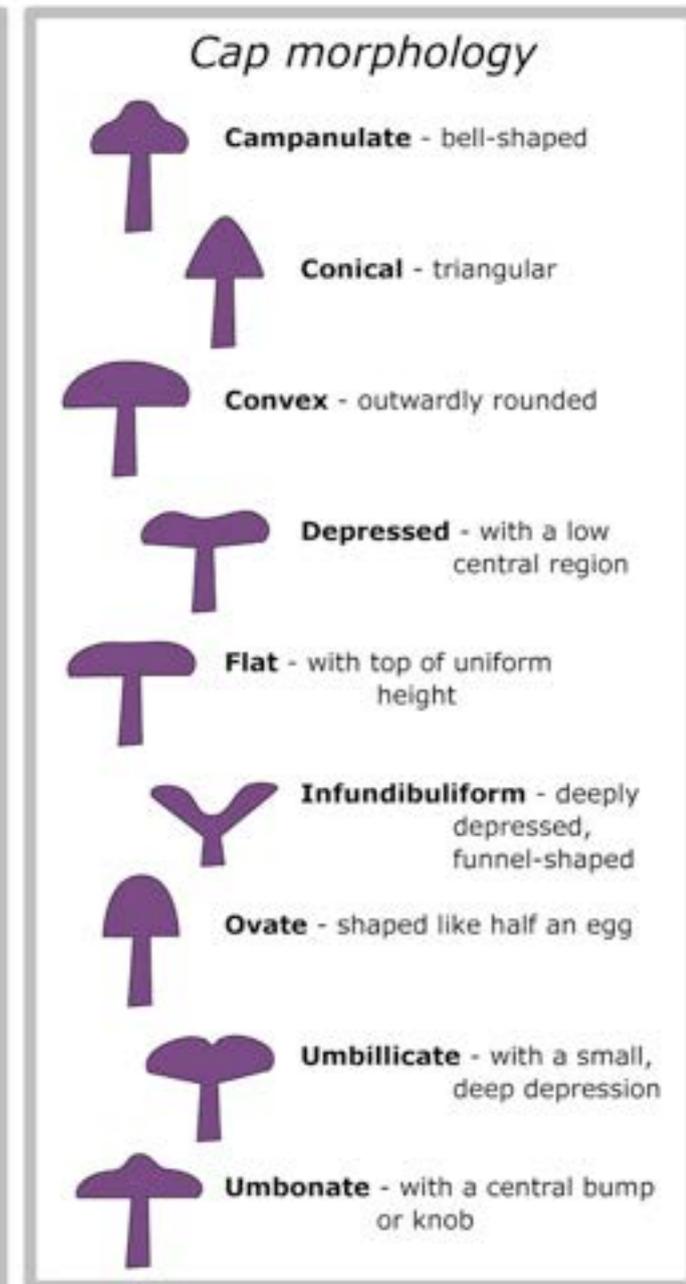
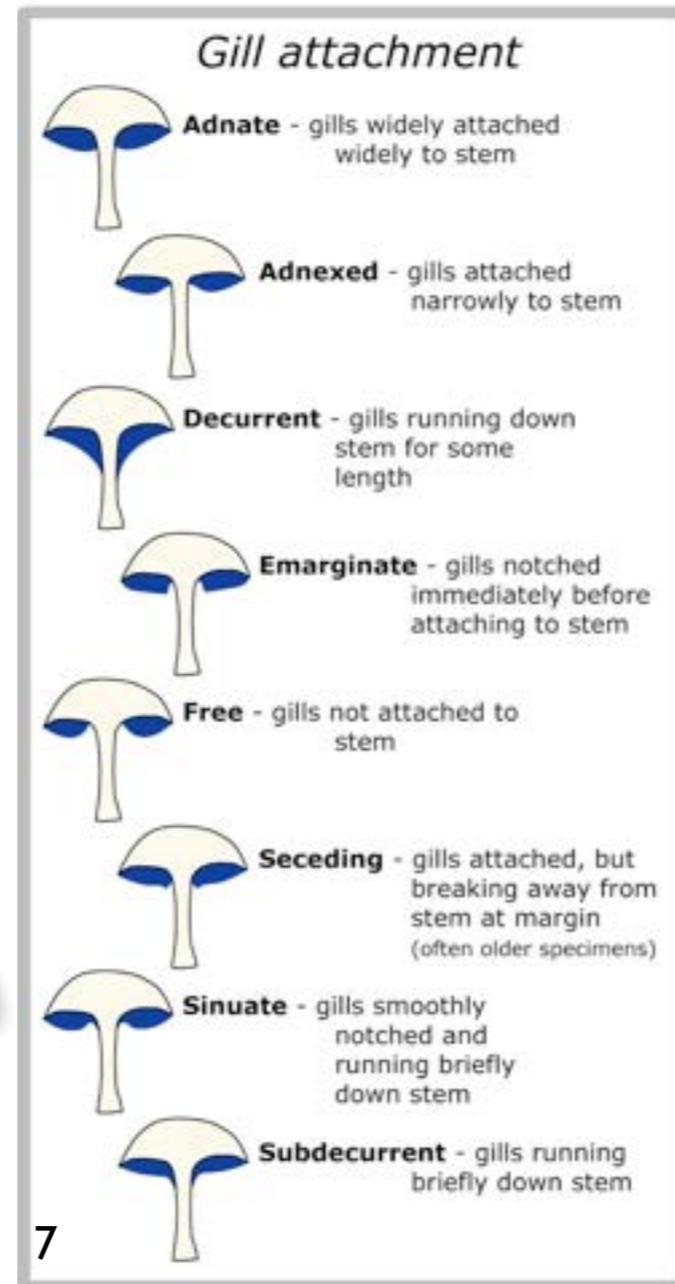
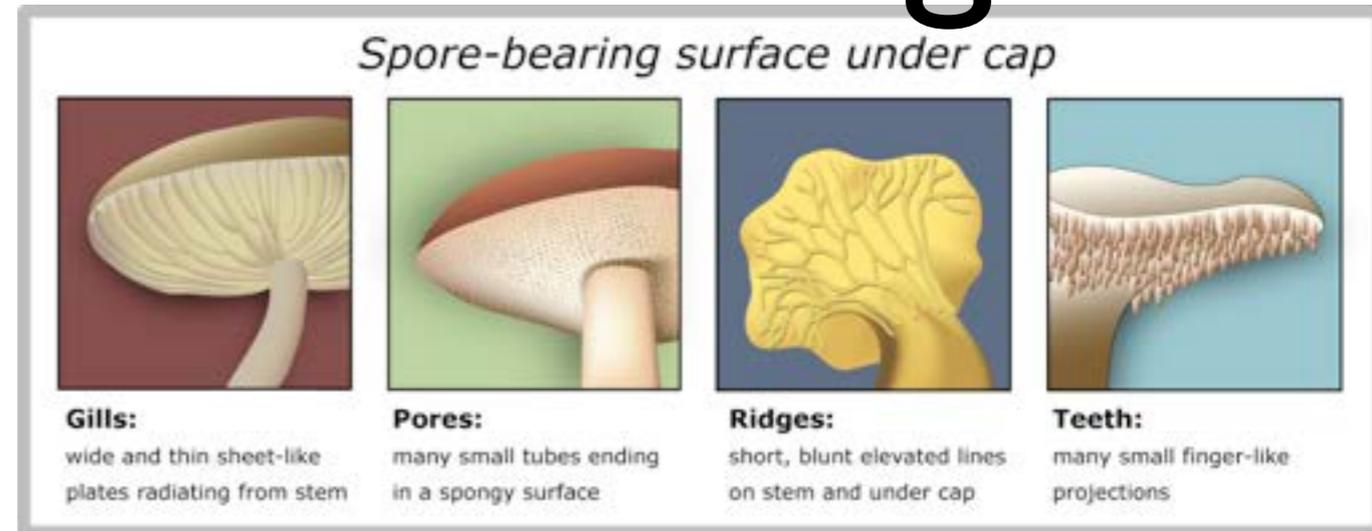
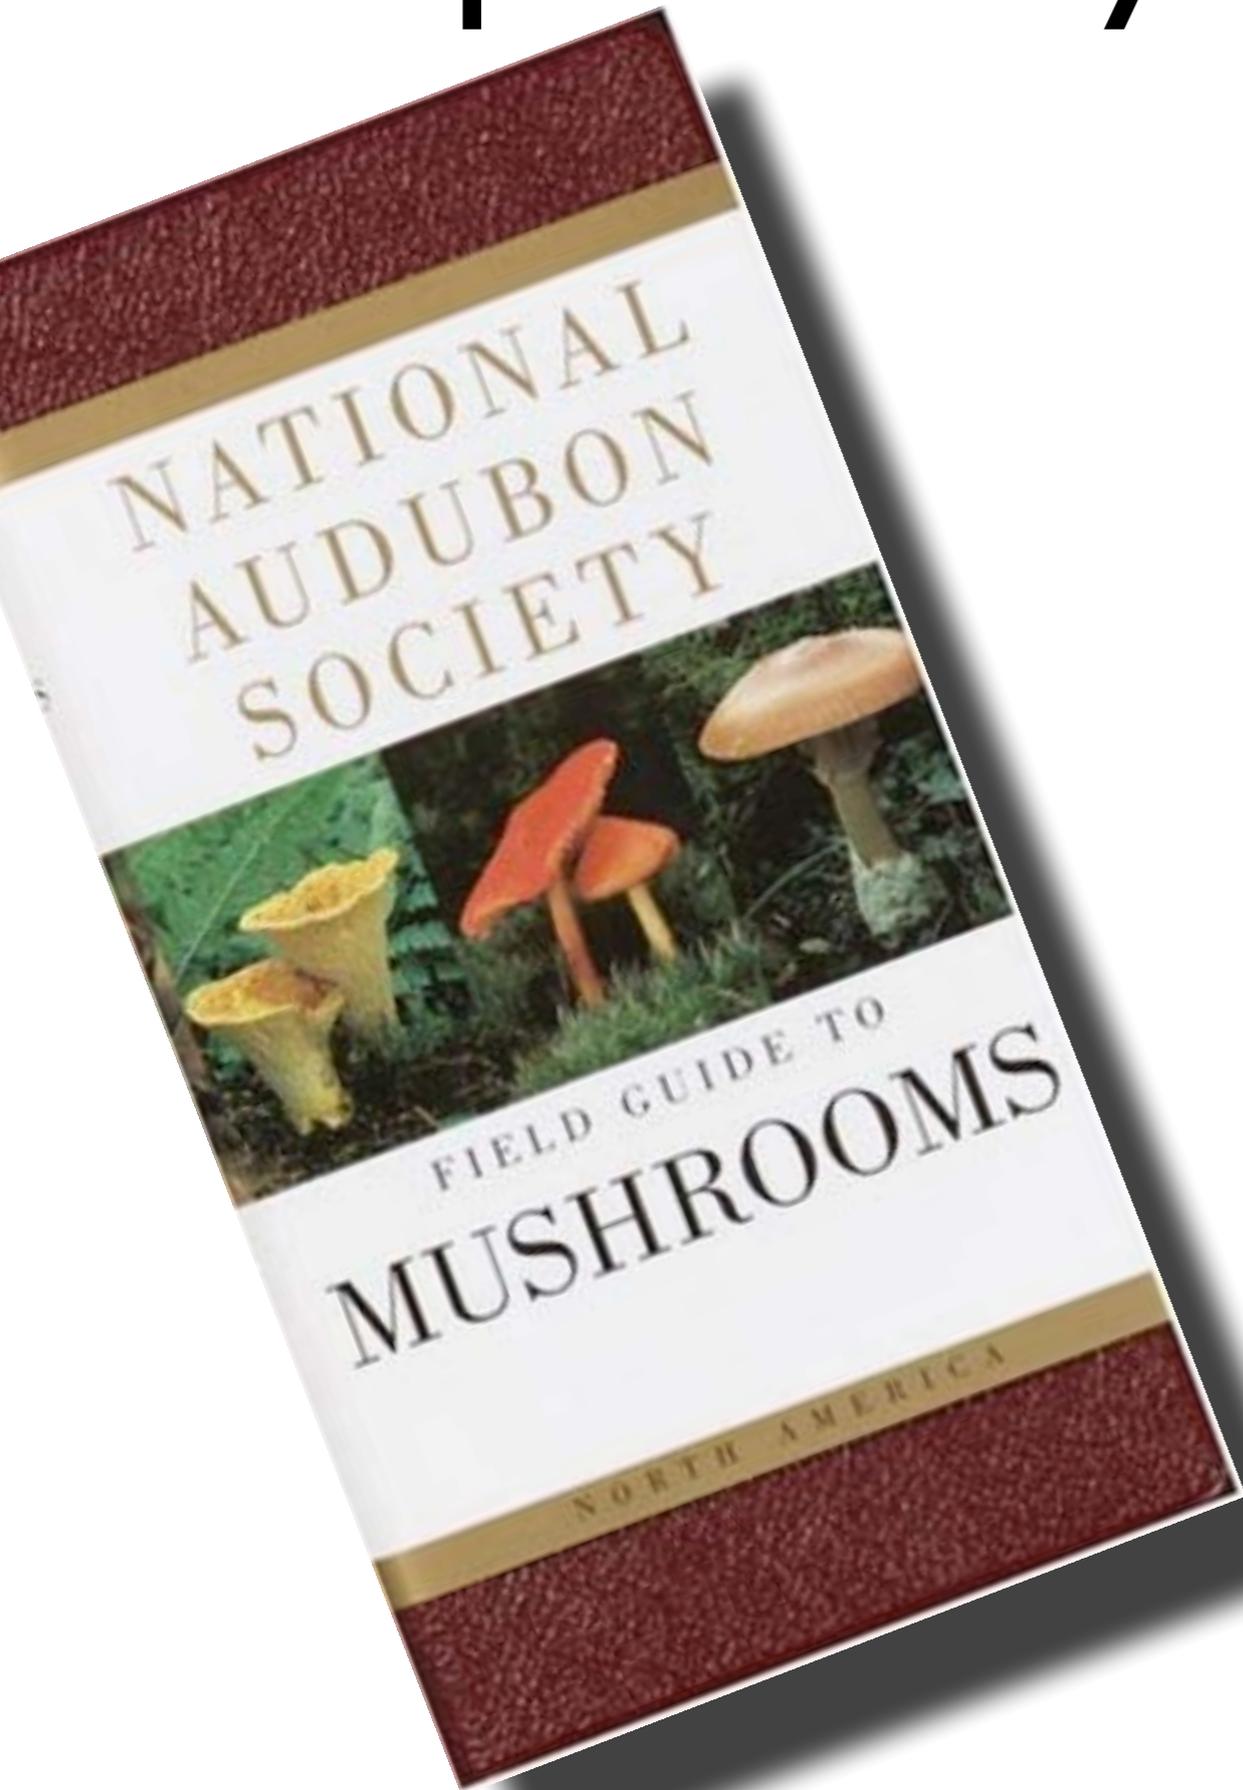
interaction

- [About Wikipedia](#)
- [Community portal](#)
- [Recent changes](#)
- [Contact Wikipedia](#)
- [Donate to Wikipedia](#)
- [Help](#)

toolbox

- [What links here](#)

# Wikipedia says: use a field guide



**Yet, the info is there...**

# *Amanita pantherina* var. *pantherina*

From Wikipedia, the free encyclopedia

(Redirected from [Amanita pantherina](#))

***Amanita pantherina*** var. ***pantherina***, also known as the "*European Panther*" and "*False Blusher*" due to its similarity to the true Blusher (*Amanita rubescens*), is a species of [Europe](#) and western Asia. Material described as *A. pantherina* in the [Americas](#) seems to belong to a number of distinct taxa only some of which have been described.

## Contents [\[hide\]](#)

- [1 Description](#)
- [2 Habitat and distribution](#)
- [3 Biochemistry](#)
- [4 See also](#)
- [5 References](#)
- [6 External links](#)

**European Panther**



**Visual expertise:  
not easily accessible to  
machines**

# Femur

From Wikipedia, the free encyclopedia

*For the **invertebrate femur**, see **Arthropod leg**.*

The **femur**, or **thigh bone**, is the most **proximal** (closest to the body) **bone** of the **leg** in **vertebrates** capable of **walking** or **jumping**, such as most land **mammals**, **birds**, many **reptiles** such as **lizards**, and **amphibians** such as **frogs**. In vertebrates with four legs such as **dogs** and **horses**, the femur is found only in the rear legs.

## Contents [hide]

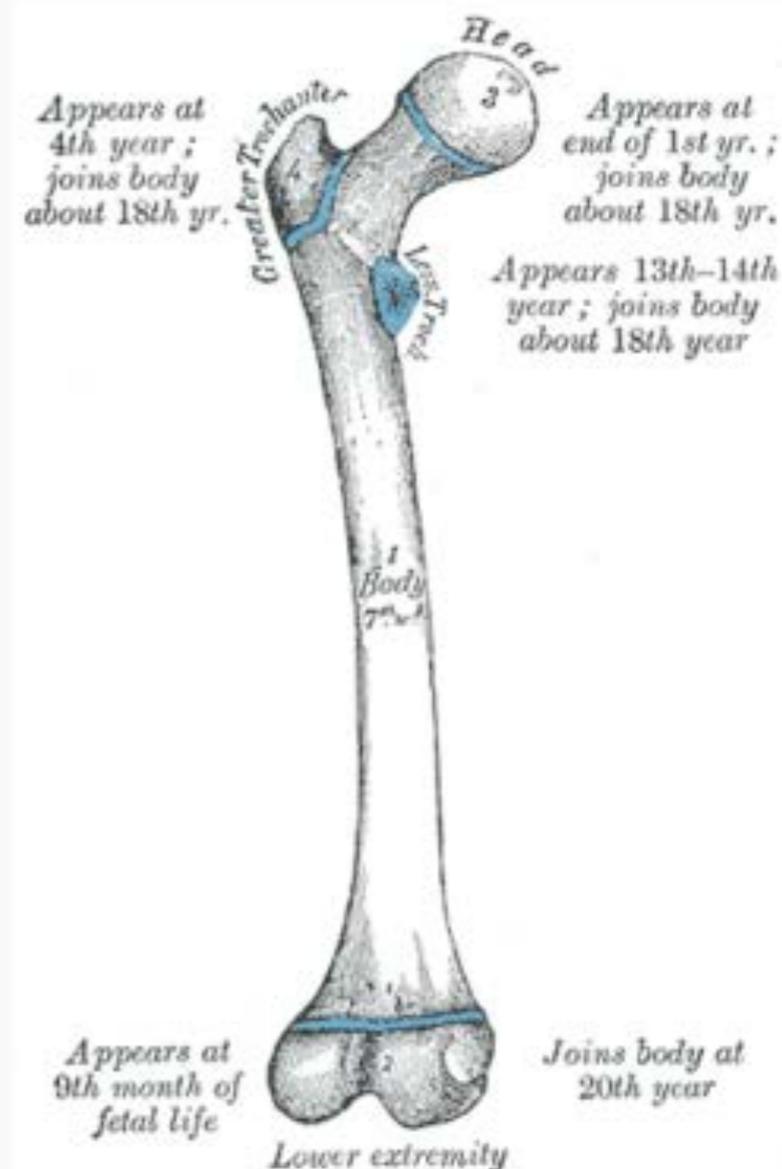
- [1 Human anatomy](#)
- [2 Evolutionary variation](#)
- [3 Etymology](#)
- [4 Additional images](#)
- [5 References](#)
- [6 External links](#)

## Human anatomy

[\[edit\]](#)

In **human anatomy**, the femur is the **longest** and largest bone. Along with the **temporal bone** of the **skull**, it is one of the two strongest bones in the body. The average adult male femur is 48 centimeters (18.9 in) in length and 2.34 cm (0.92 in) in diameter and can support up to 30 times the weight of an adult.<sup>[1]</sup> It forms part of the **hip** (at the **acetabulum**) and part of the **knee**, which is located above. There are four eminences, or protuberances, in the human femur: the **head**, the **greater trochanter**, the **lesser trochanter**, and the **lower extremity**. They appear at various times from just before birth to about age 14. Initially, they are joined to the main body of the femur with **cartilage**, which gradually

### Bone: Femur



Anterior view of the femur

**Latin** *os femoris*

**Gray's** *subject #59 242* 



**WIKIPEDIA**  
*The Free Encyclopedia*

### navigation

- [Main page](#)
- [Contents](#)
- [Featured content](#)
- [Current events](#)
- [Random article](#)

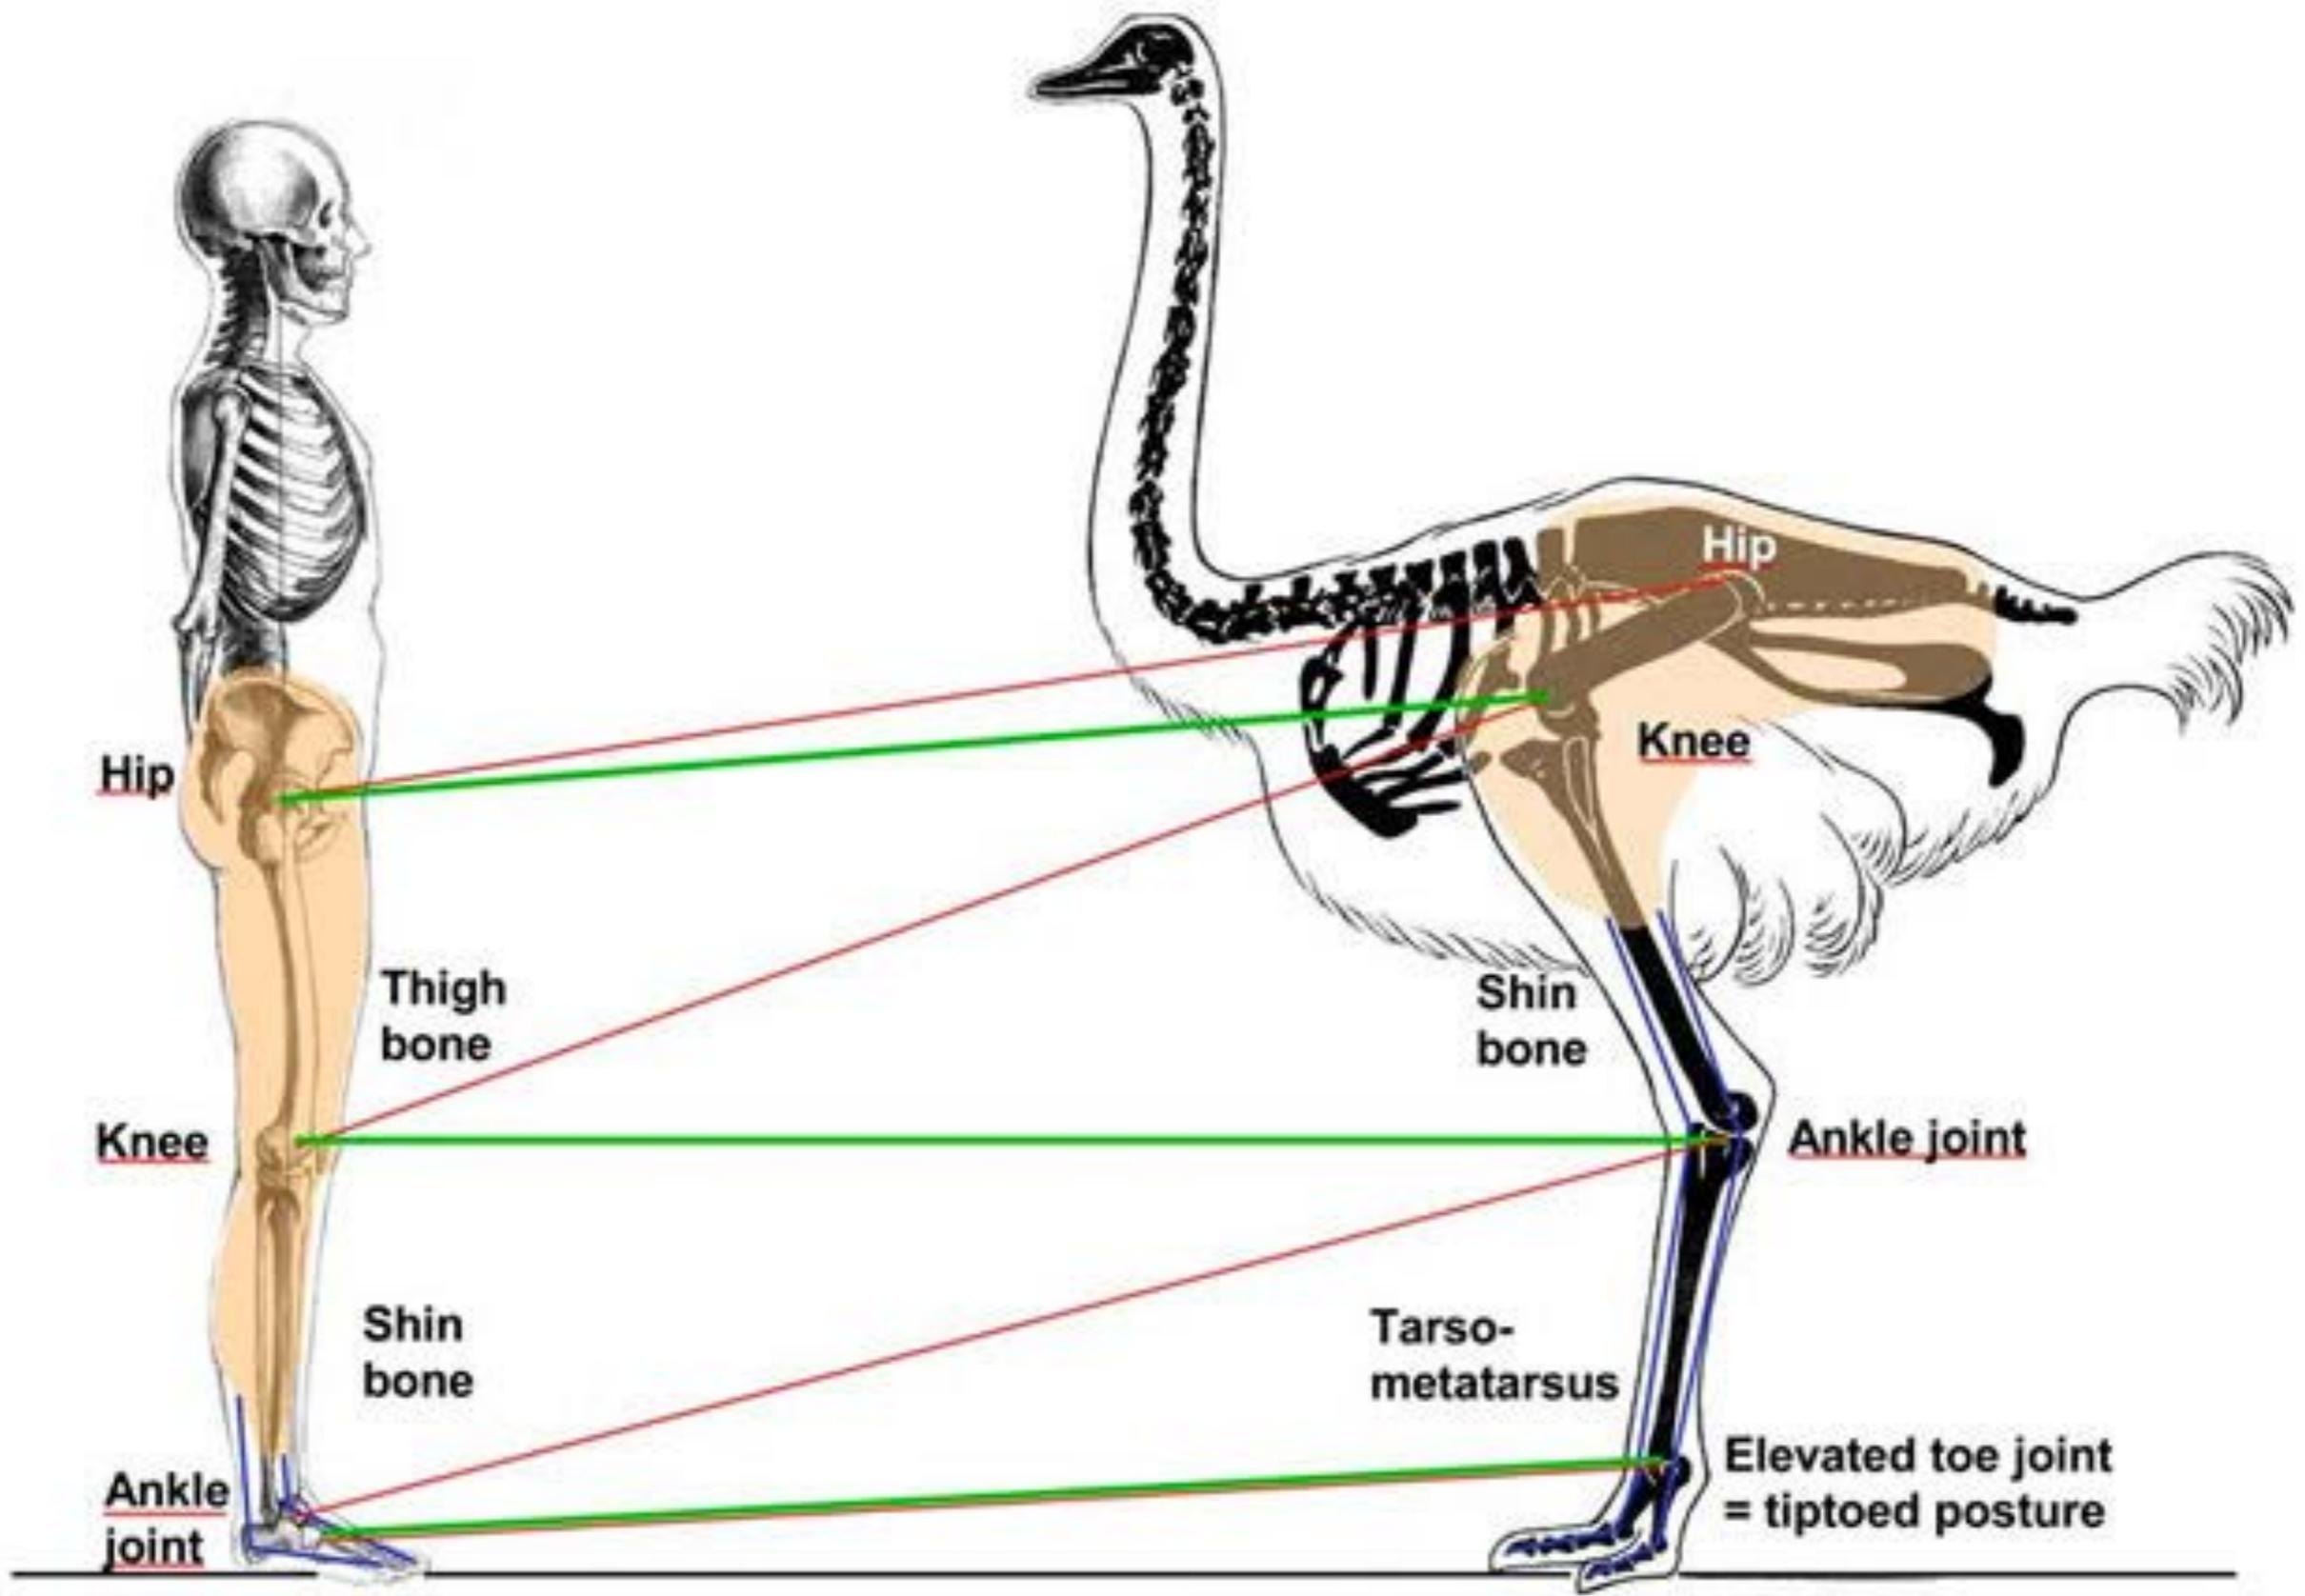
### search

### interaction

- [About Wikipedia](#)
- [Community portal](#)
- [Recent changes](#)
- [Contact Wikipedia](#)
- [Donate to Wikipedia](#)
- [Help](#)

### toolbox

- [What links here](#)
- [Related changes](#)
- [Upload file](#)
- [Special pages](#)
- [Printable version](#)
- [Permanent link](#)
- [Cite this page](#)

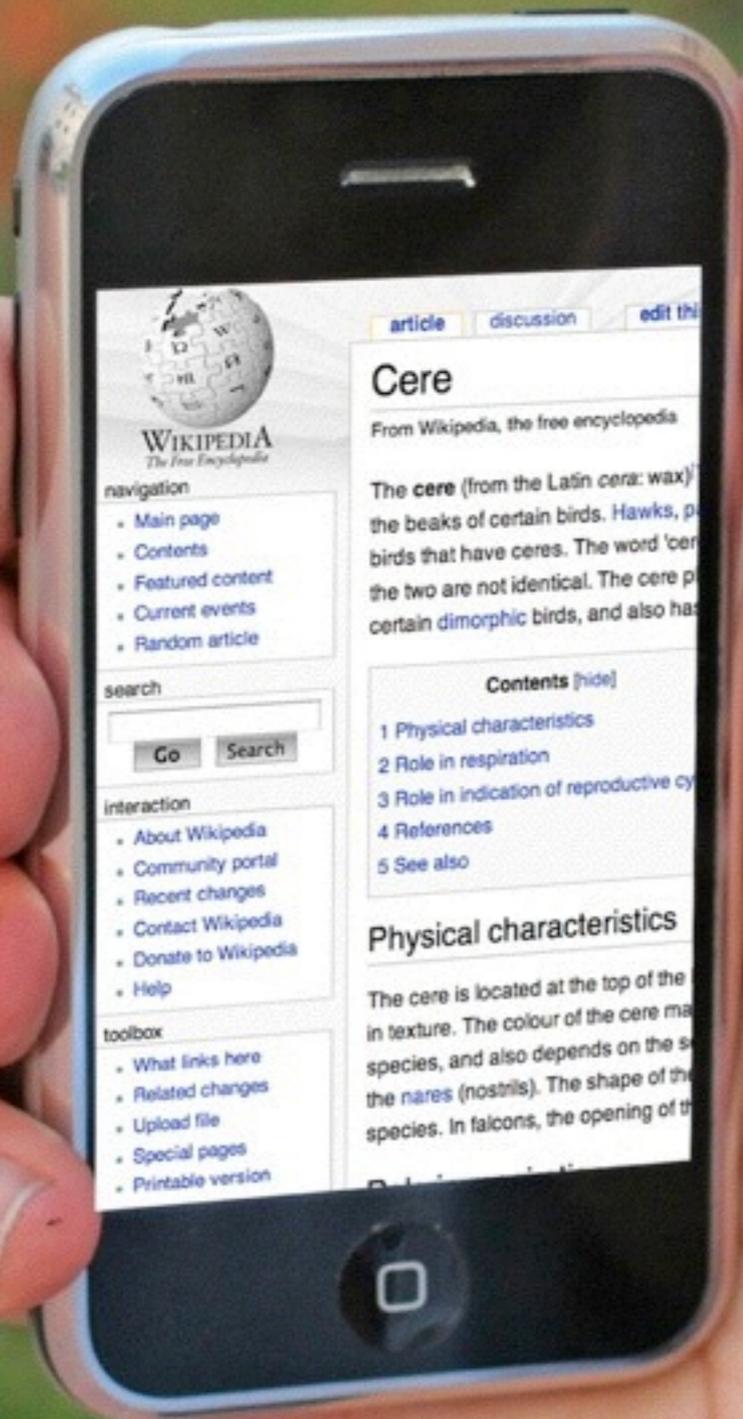


# Lessons:

- *Visual queries*
  - Easy for humans
  - Difficult for machines
- Pictures are *digital dark matter*
- Expert knowledge - how to collect?







# WIKIPEDIA



[Log in](#) / [create account](#)

[article](#) [discussion](#) [edit this page](#) [history](#)

## Cere

From Wikipedia, the free encyclopedia

The **cere** (from the Latin *cera*: wax)<sup>[1]</sup> or **operculum**<sup>[2]</sup> is a soft, fleshy swelling found on the beaks of certain birds. [Hawks](#), [parrots](#), [doves](#), [skuas](#) and [budgerigars](#) are among the birds that have ceres. The word 'cere' is often used synonymously with 'beak', although the two are not identical. The cere plays a role in indicating the reproductive stage of certain [dimorphic](#) birds, and also has a key function in respiration.

### Contents [\[hide\]](#)

- [1 Physical characteristics](#)
- [2 Role in respiration](#)
- [3 Role in indication of reproductive cycle](#)
- [4 References](#)
- [5 See also](#)

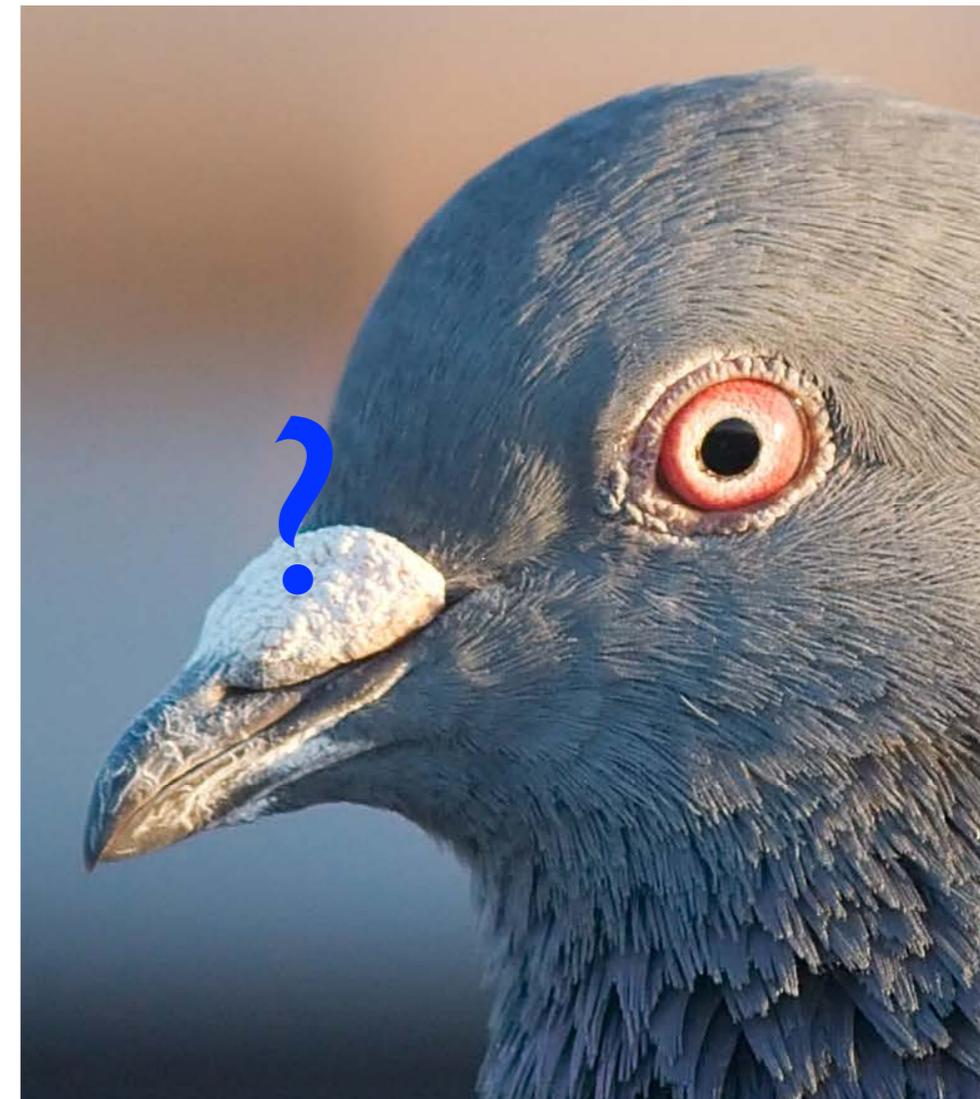
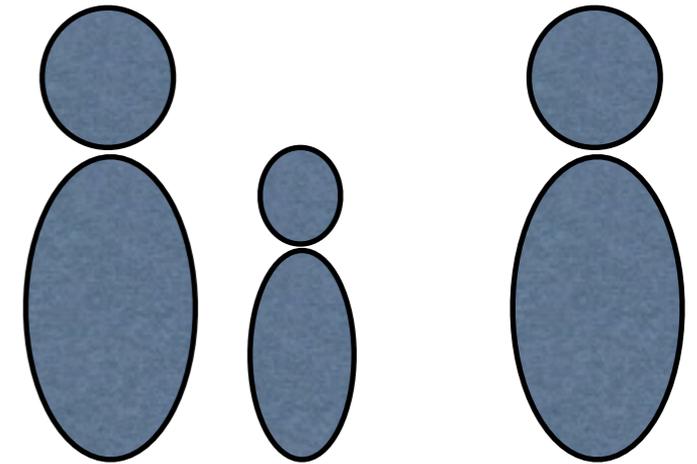
## Physical characteristics [\[edit\]](#)

The cere is located at the top of the beak,<sup>[2][3]</sup> and is often waxy in texture. The colour of the cere may vary from species to species, and also depends on the season.<sup>[4]</sup> The cere contains the [nares](#) (nostrils). The shape of the cere varies from species to species. In falcons, the opening of the nares is roughly circular in shape.<sup>[2]</sup>



The [Rock Pigeon](#), *Columba livia*. The cere is located at the top of the beak, and is a distinctive white-grey colour.

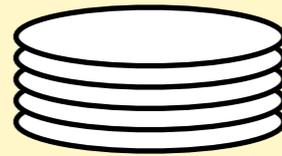
# Users



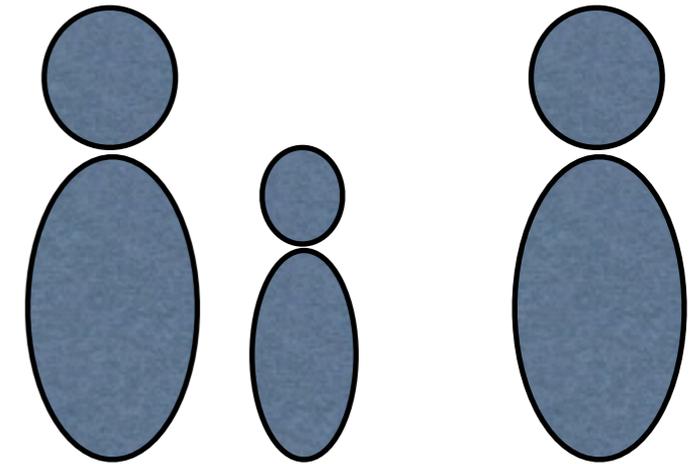
# WIKIPEDIA



# VISIPEDIA



# Users



[Log in / create account](#)

[article](#) [discussion](#) [edit this page](#) [history](#)

## Cere

From Wikipedia, the free encyclopedia

The **cere** (from the Latin *cera*: wax)<sup>[1]</sup> or **operculum**<sup>[2]</sup> is a soft, fleshy swelling found on the beaks of certain birds. [Hawks](#), [parrots](#), [doves](#), [skuas](#) and [budgerigars](#) are among the birds that have ceres. The word 'cere' is often used synonymously with 'beak', although the two are not identical. The cere plays a role in indicating the reproductive stage of certain [dimorphic](#) birds, and also has a key function in respiration.

**Contents** [\[hide\]](#)

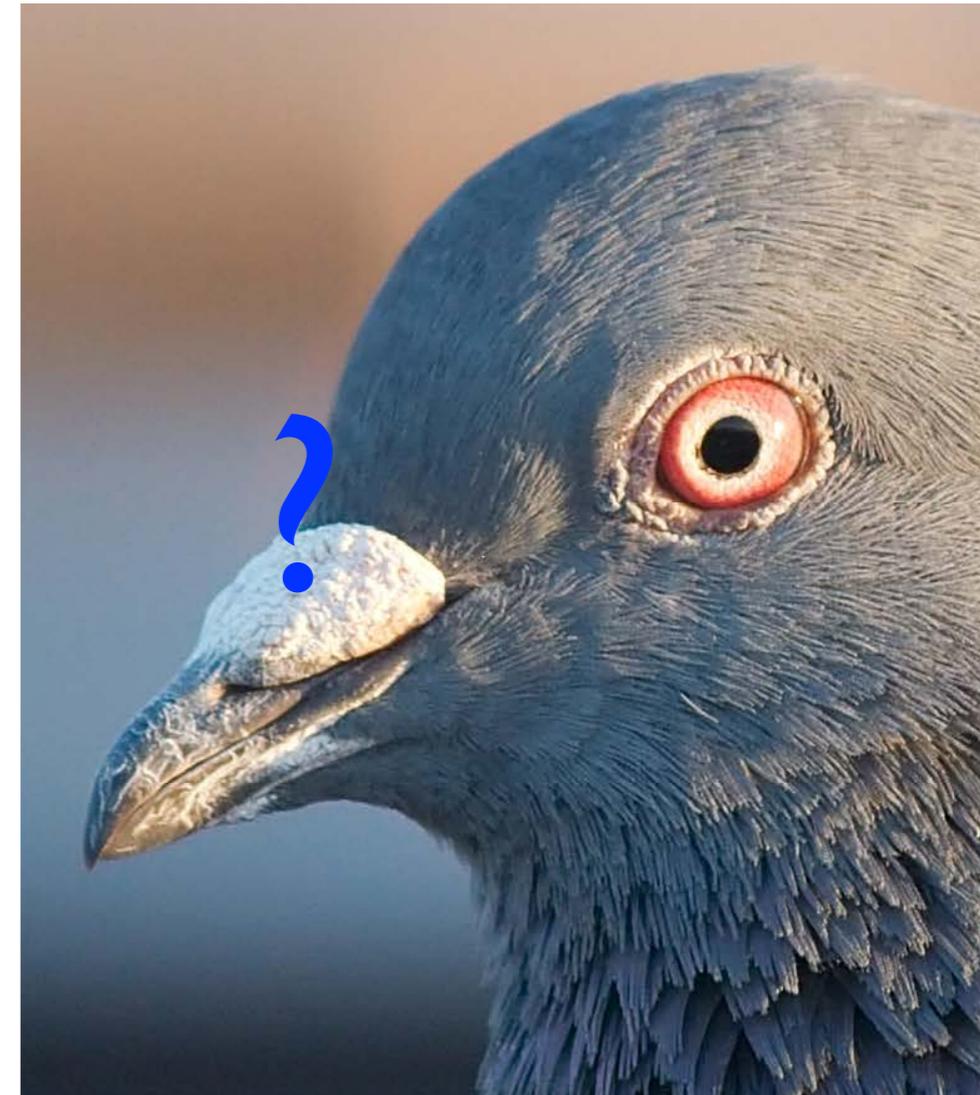
- 1 Physical characteristics
- 2 Role in respiration
- 3 Role in indication of reproductive cycle
- 4 References
- 5 See also



The [Rock Pigeon](#), *Columba livia*. The cere is located at the top of the beak, and is a distinctive white-grey colour.

### Physical characteristics [\[edit\]](#)

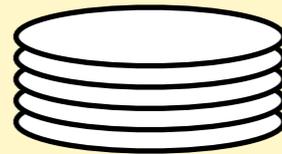
The cere is located at the top of the beak,<sup>[2][3]</sup> and is often waxy in texture. The colour of the cere may vary from species to species, and also depends on the season.<sup>[4]</sup> The cere contains the [nares](#) (nostrils). The shape of the cere varies from species to species. In falcons, the opening of the nares is roughly circular in shape.<sup>[2]</sup>



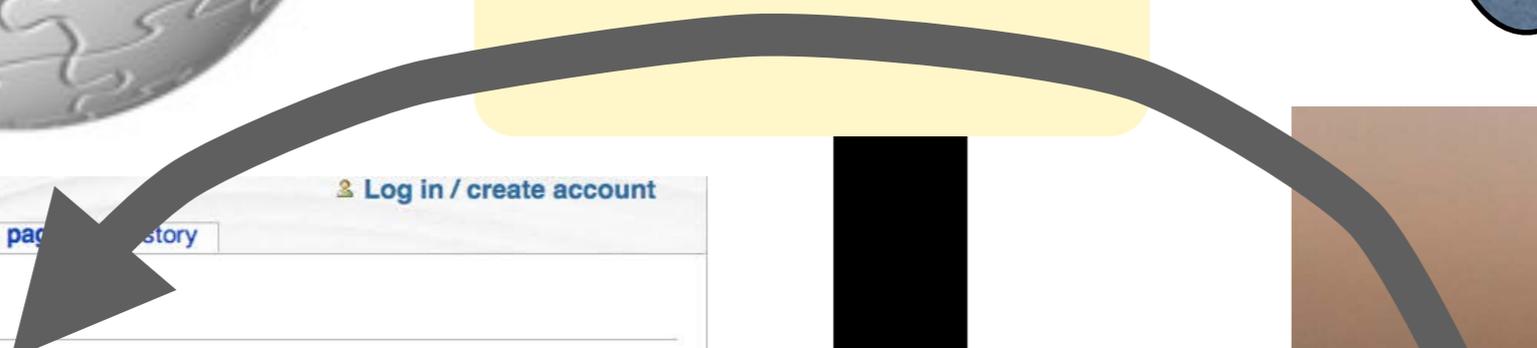
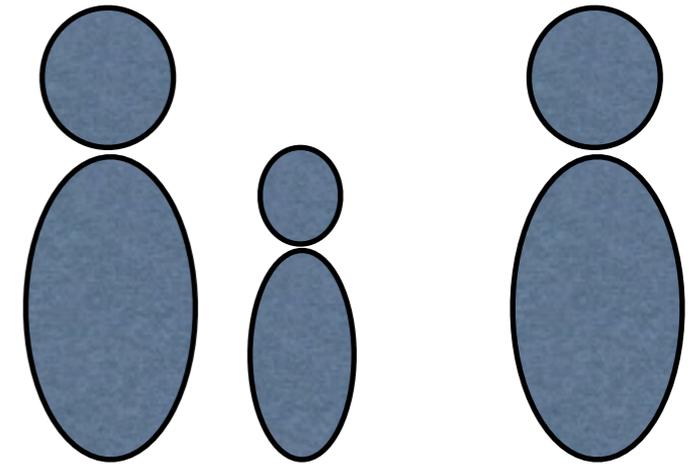
# WIKIPEDIA



# VISIPEDIA



# Users



[article](#) [discussion](#) [edit this page](#) [history](#) [Log in / create account](#)

## Cere

From Wikipedia, the free encyclopedia

The **cere** (from the Latin *cera*: wax)<sup>[1]</sup> or **operculum**<sup>[2]</sup> is a soft, fleshy swelling found on the beaks of certain birds. [Hawks](#), [parrots](#), [doves](#), [skuas](#) and [budgerigars](#) are among the birds that have ceres. The word 'cere' is often used synonymously with 'beak', although the two are not identical. The cere plays a role in indicating the reproductive stage of certain [dimorphic](#) birds, and also has a key function in respiration.

**Contents** [\[hide\]](#)

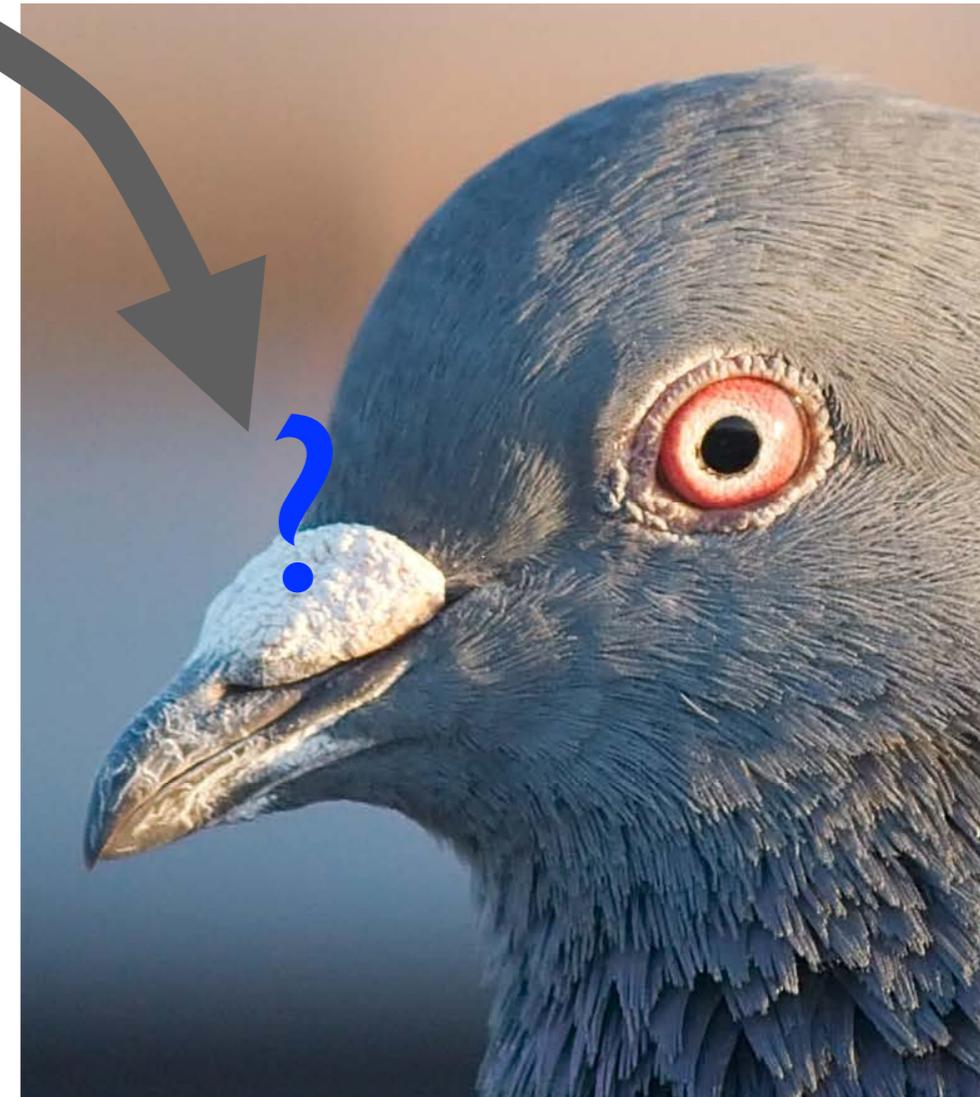
- 1 Physical characteristics
- 2 Role in respiration
- 3 Role in indication of reproductive cycle
- 4 References
- 5 See also



The [Rock Pigeon](#), *Columba livia*. The cere is located at the top of the beak, and is a distinctive white-grey colour.

### Physical characteristics

The cere is located at the top of the beak,<sup>[2][3]</sup> and is often waxy in texture. The colour of the cere may vary from species to species, and also depends on the season.<sup>[4]</sup> The cere contains the [nares](#) (nostrils). The shape of the cere varies from species to species. In falcons, the opening of the nares is roughly circular in shape.<sup>[2]</sup>



**How do we make this  
happen?**

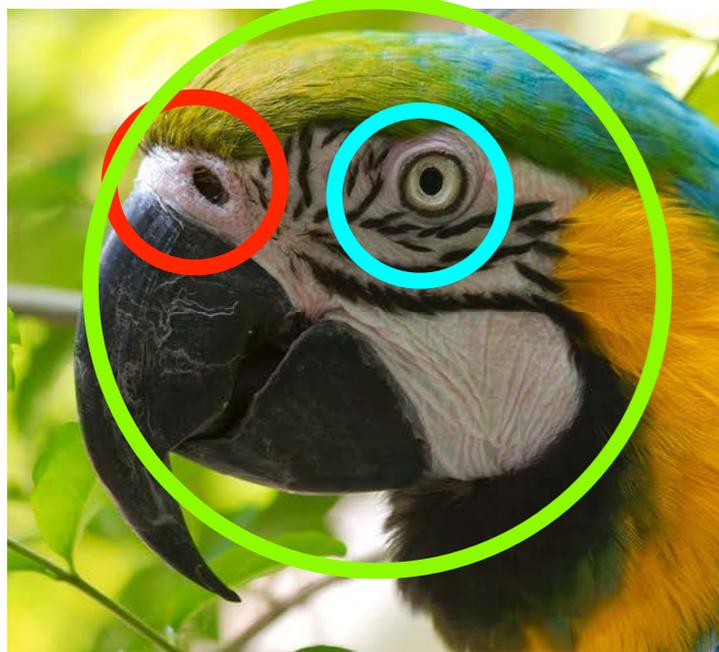
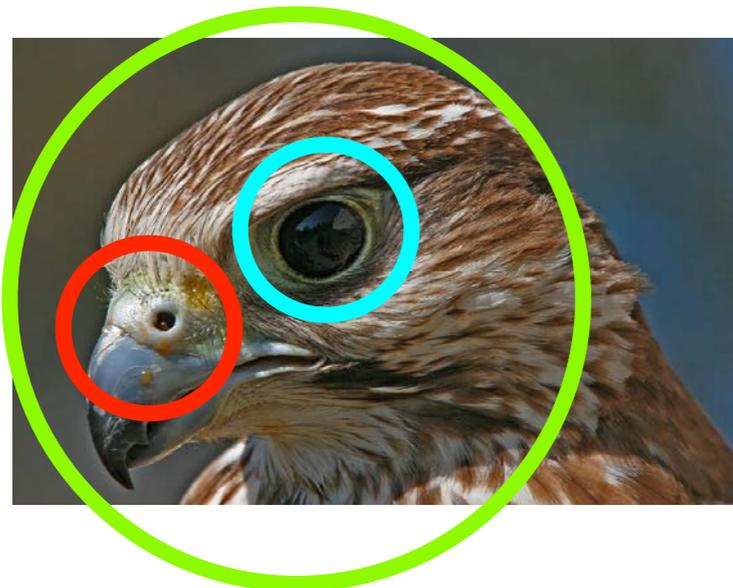
# 'Visipedia'



...

# 'Visipedia'

## Preprocessing

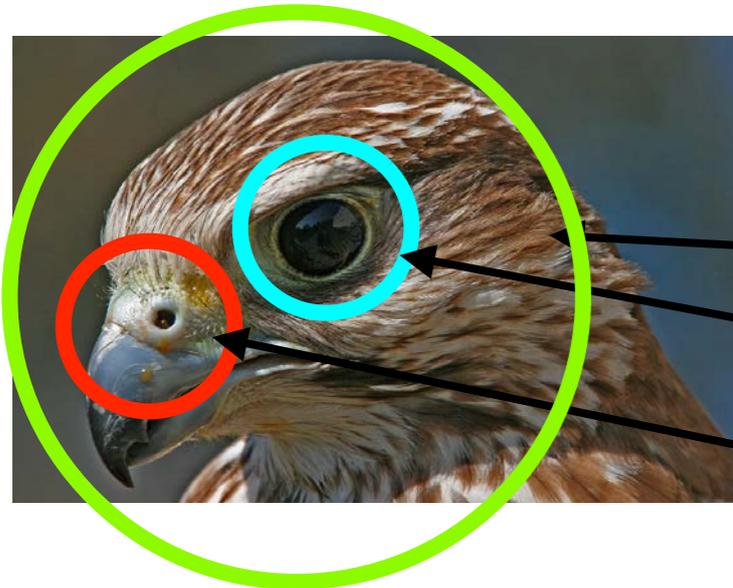


...

# 'Visipedia'

Preprocessing

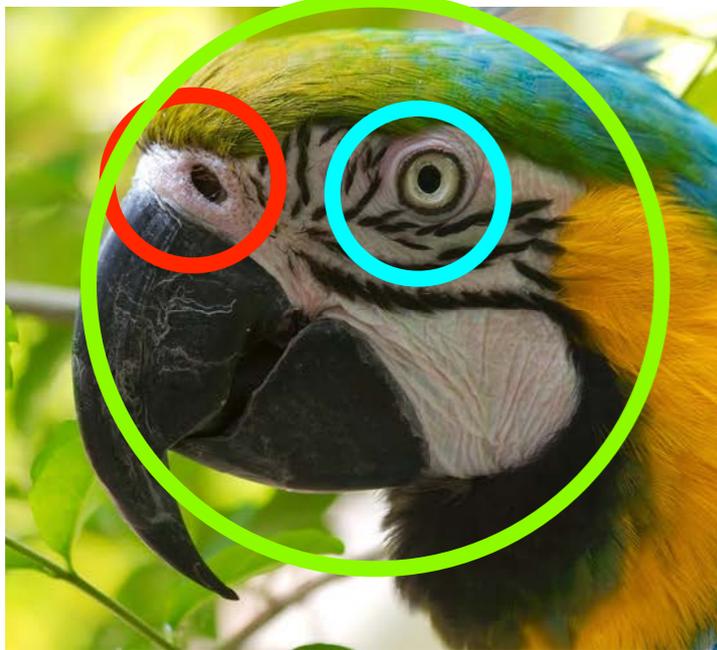
Expert



Head

Eye

Cere



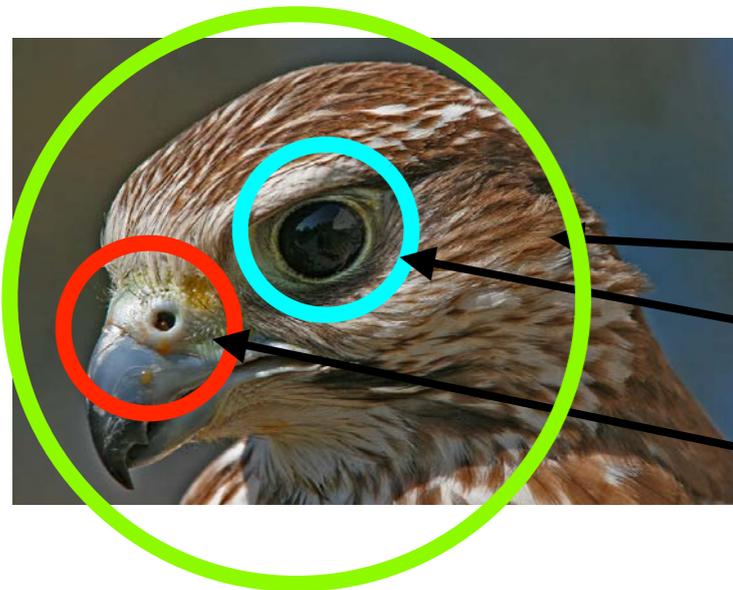
...

# 'Visipedia'

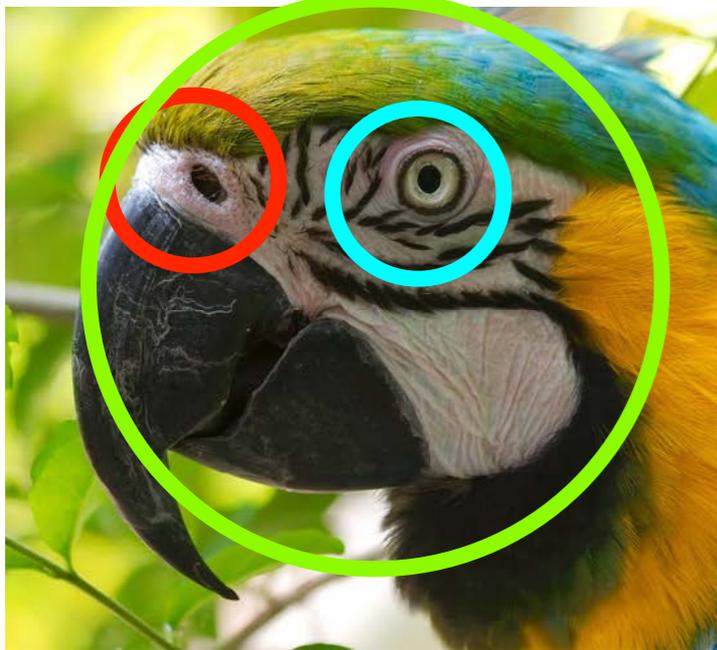
Preprocessing

Expert

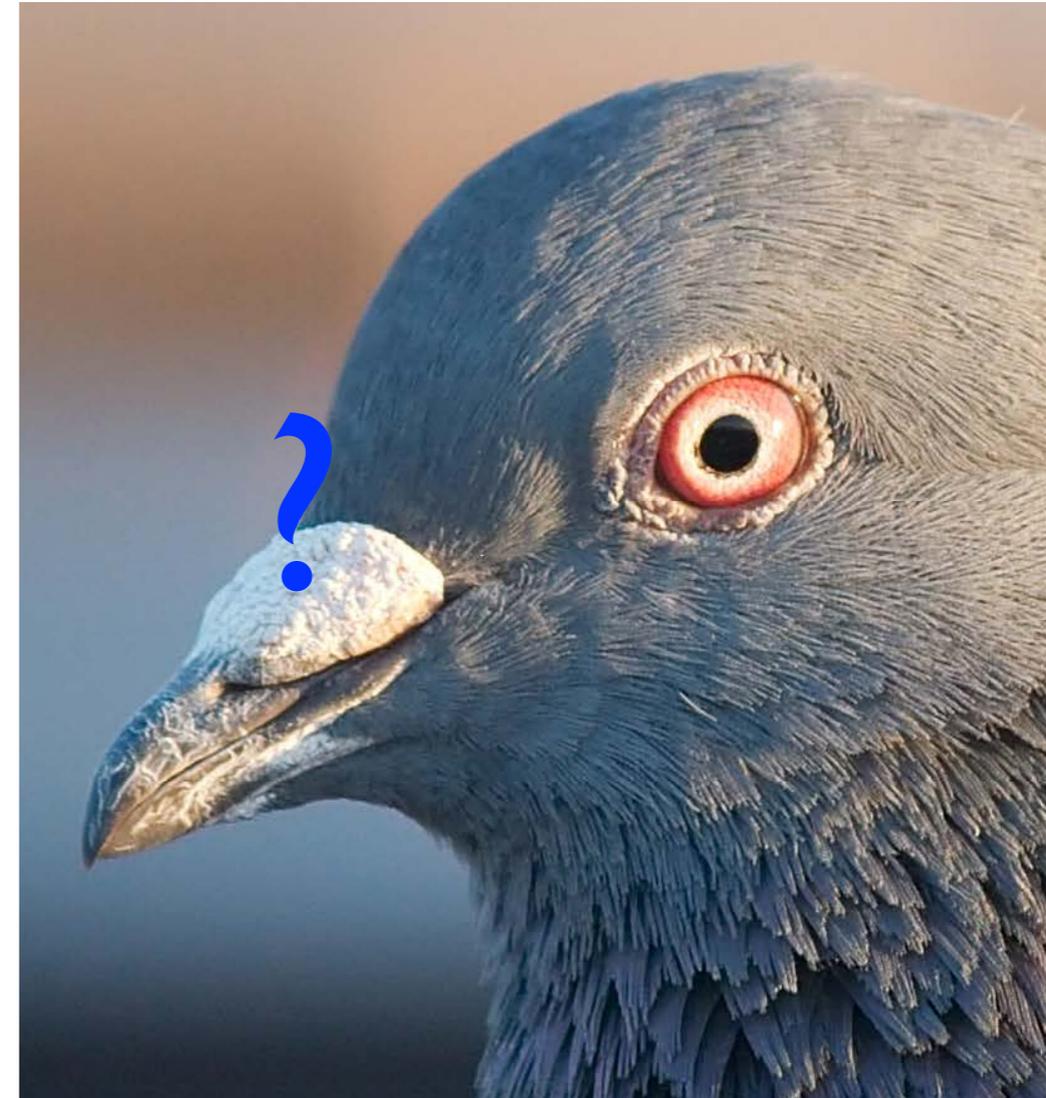
User



Head  
Eye  
Cere



...

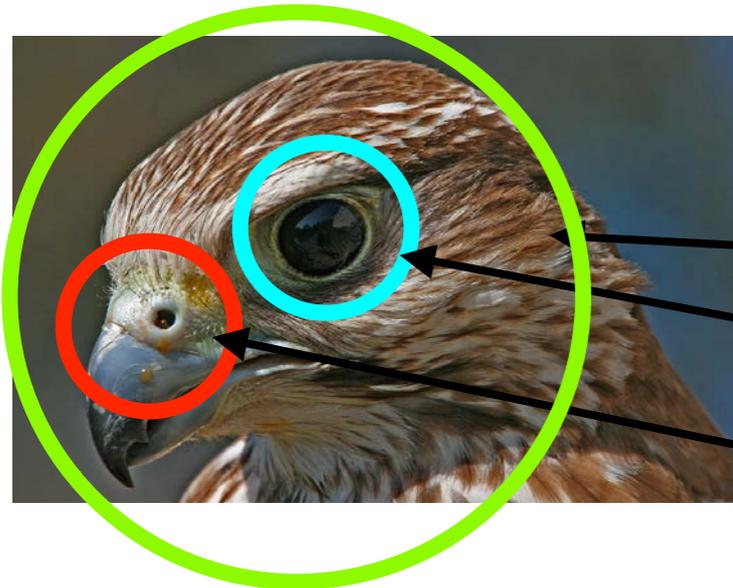


# 'Visipedia'

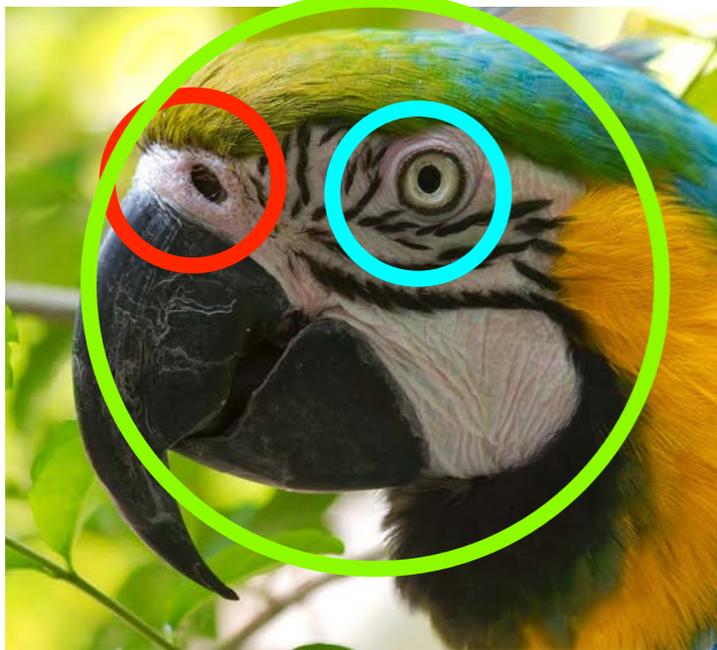
Preprocessing

Expert

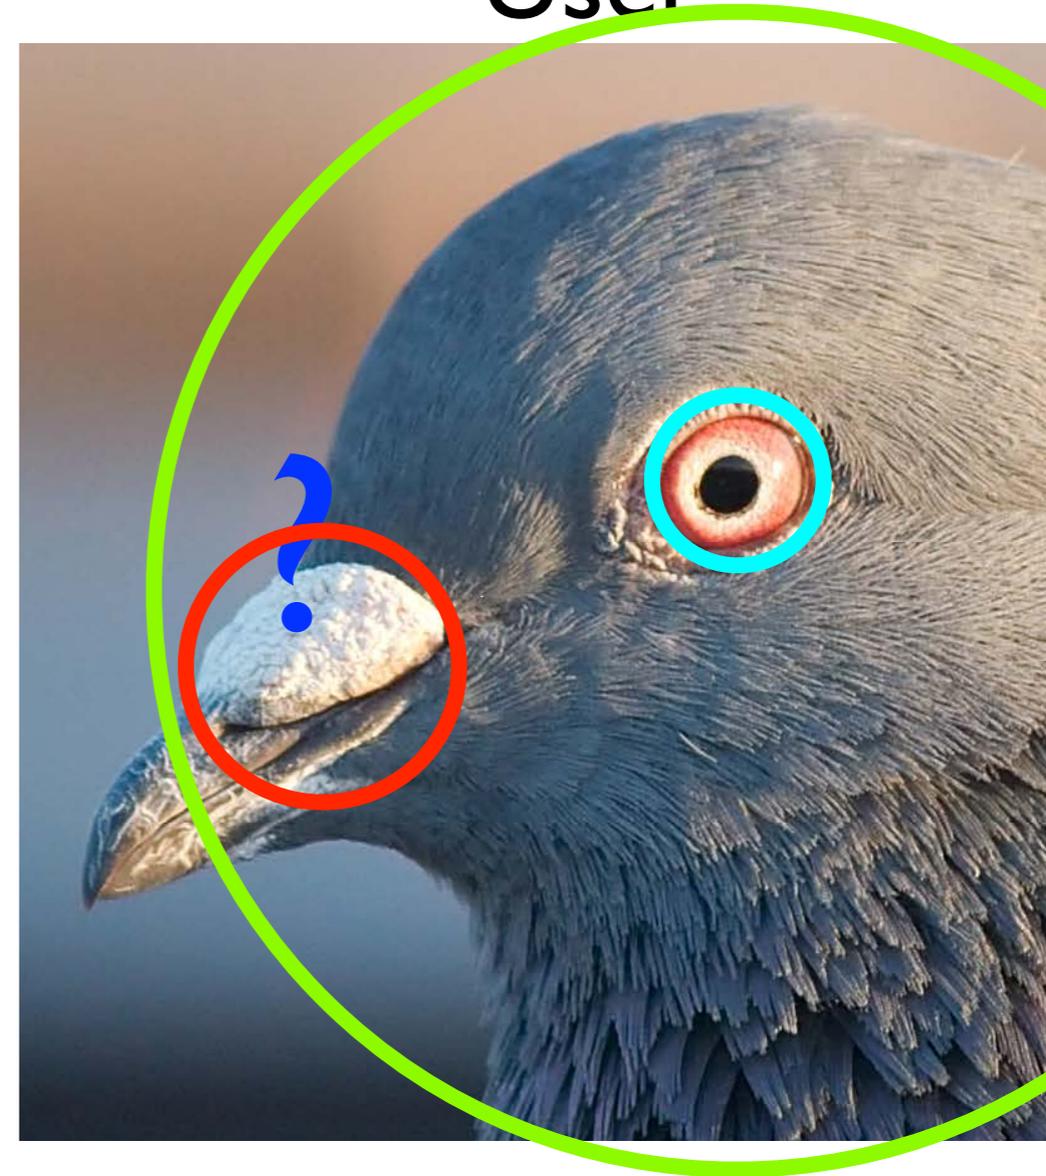
User



Head  
Eye  
Cere



...

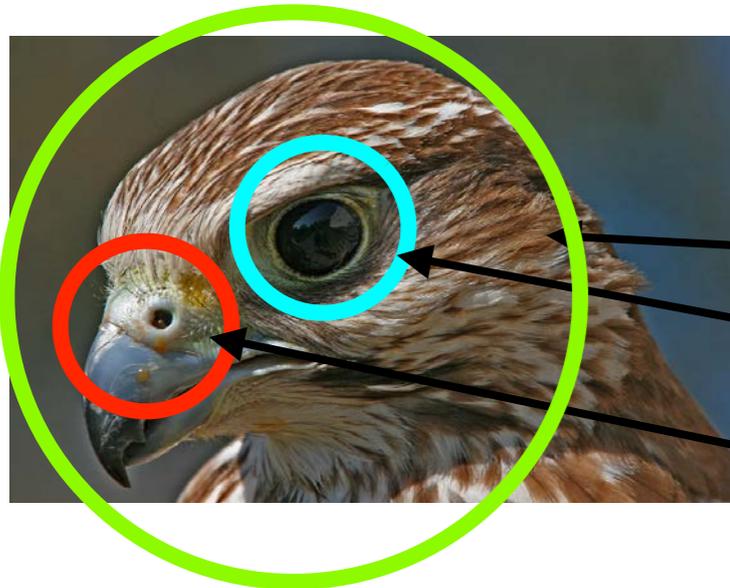


# 'Visipedia'

Preprocessing

Expert

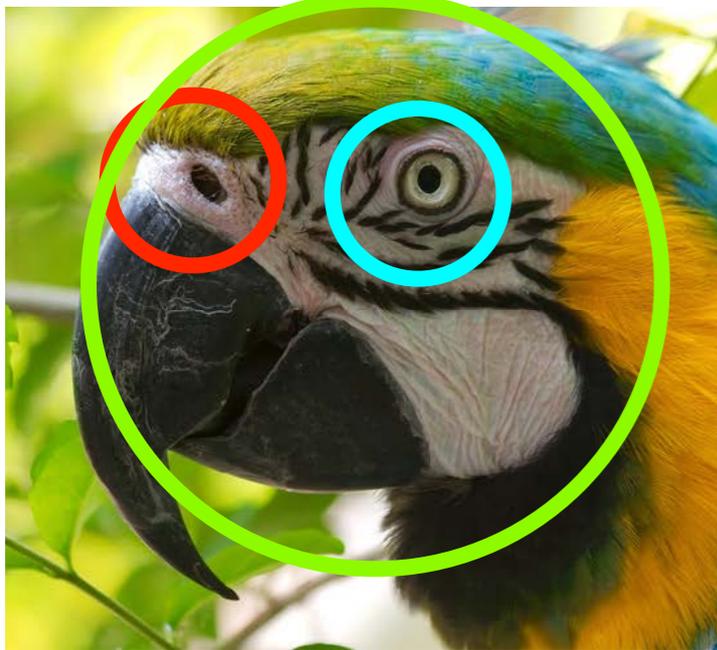
User



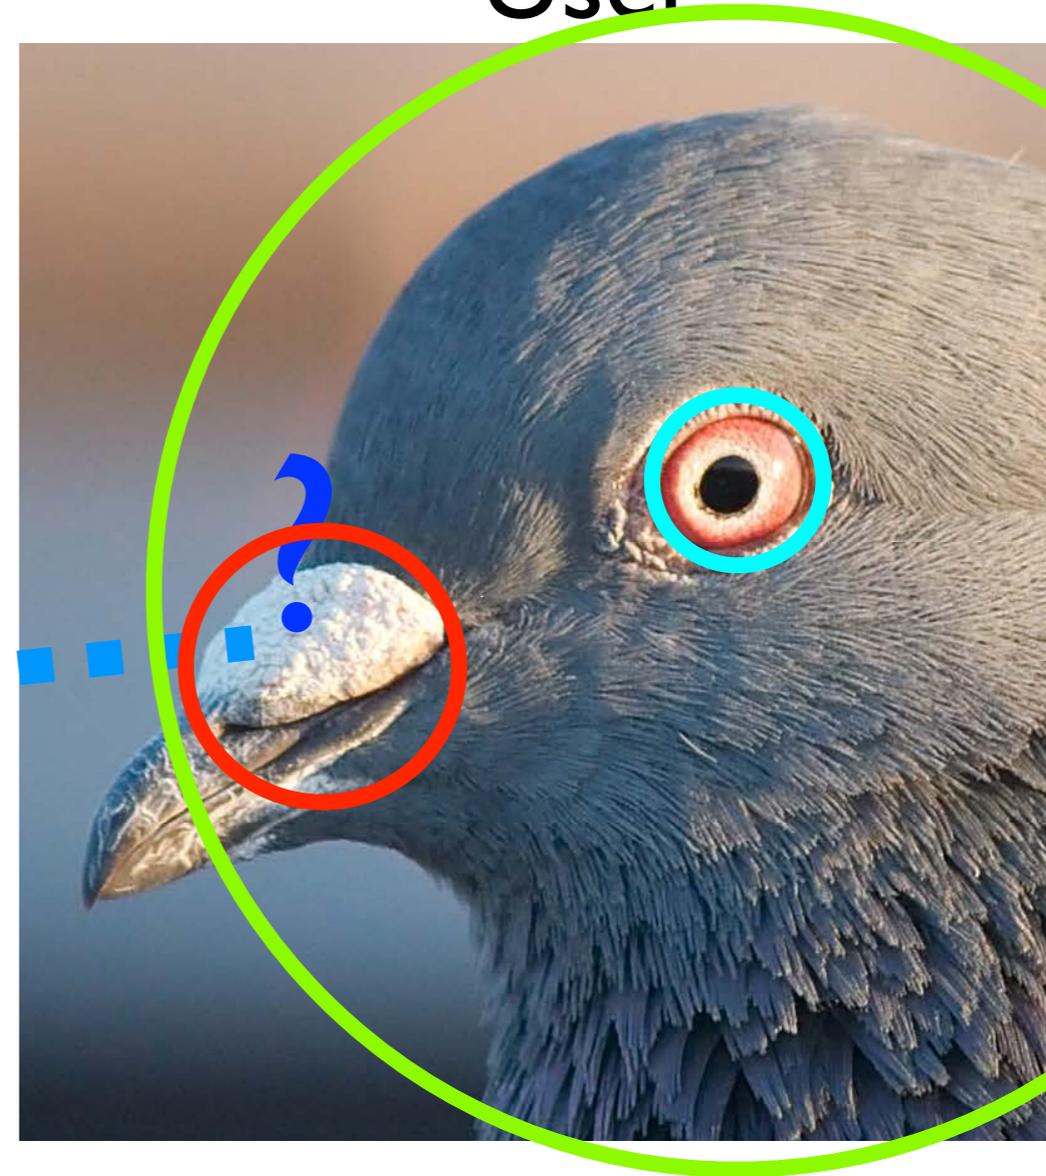
Head

Eye

Cere



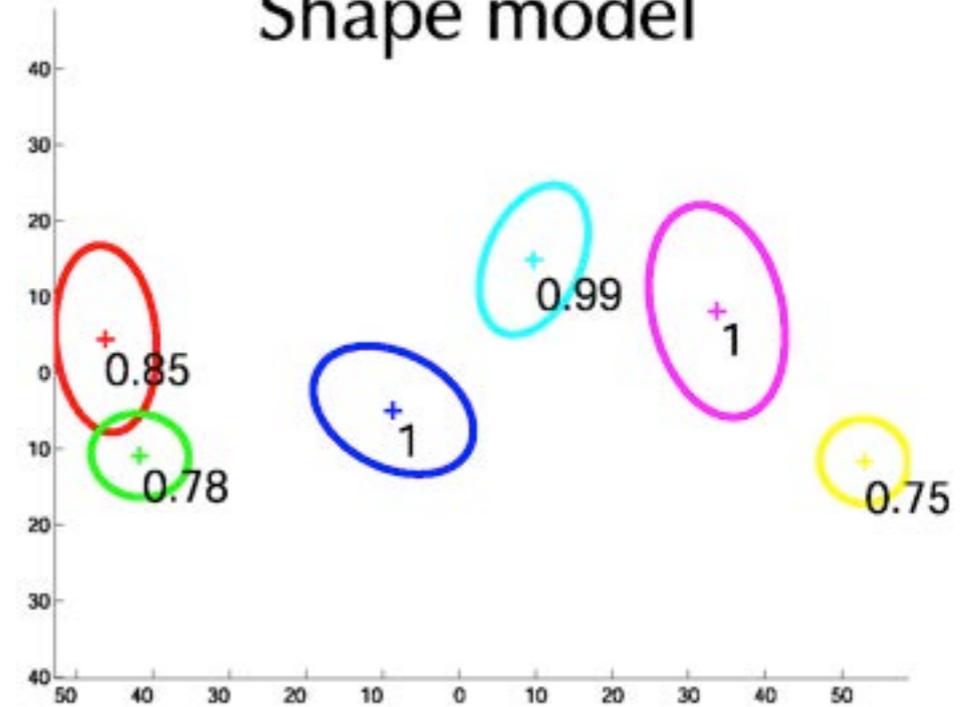
...



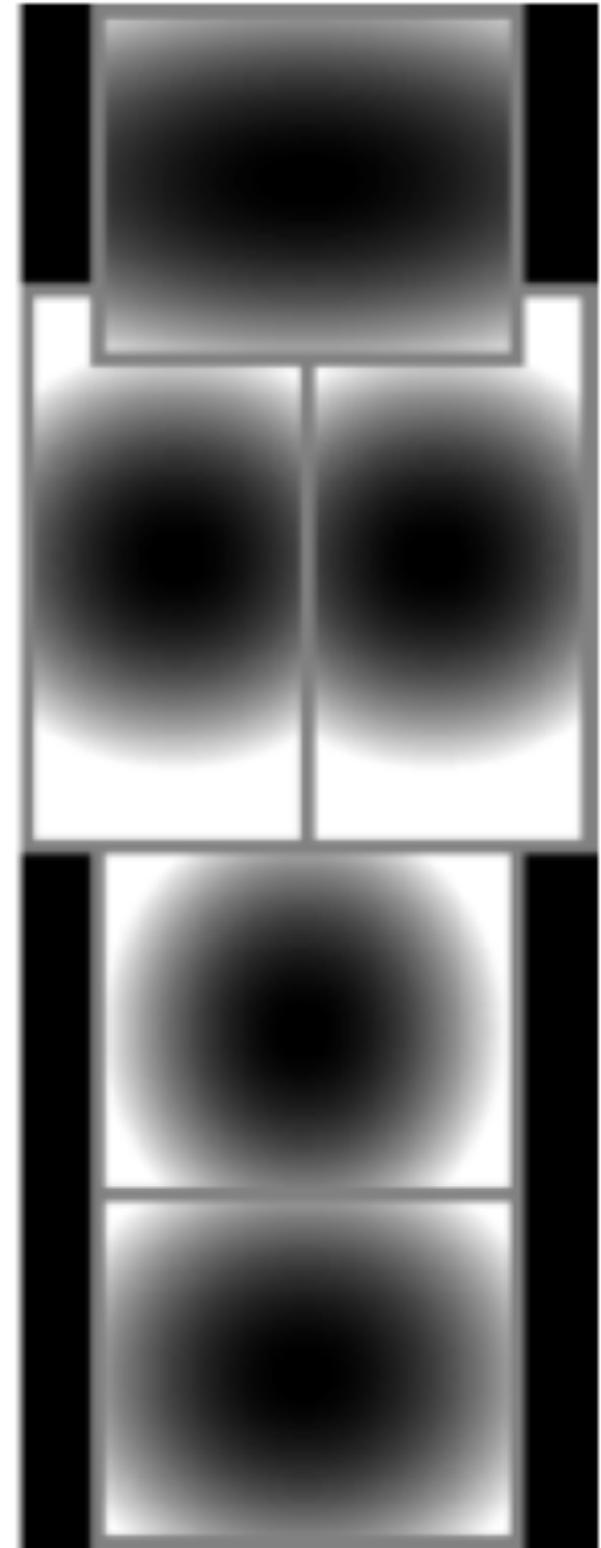
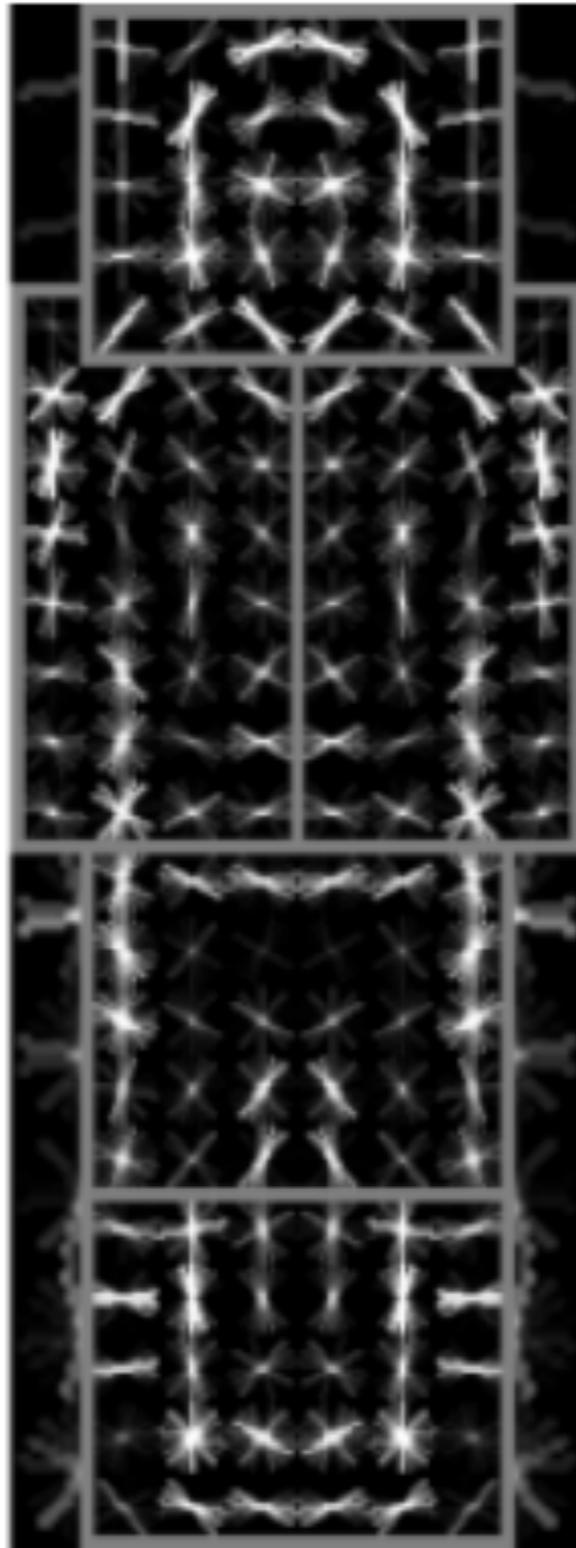
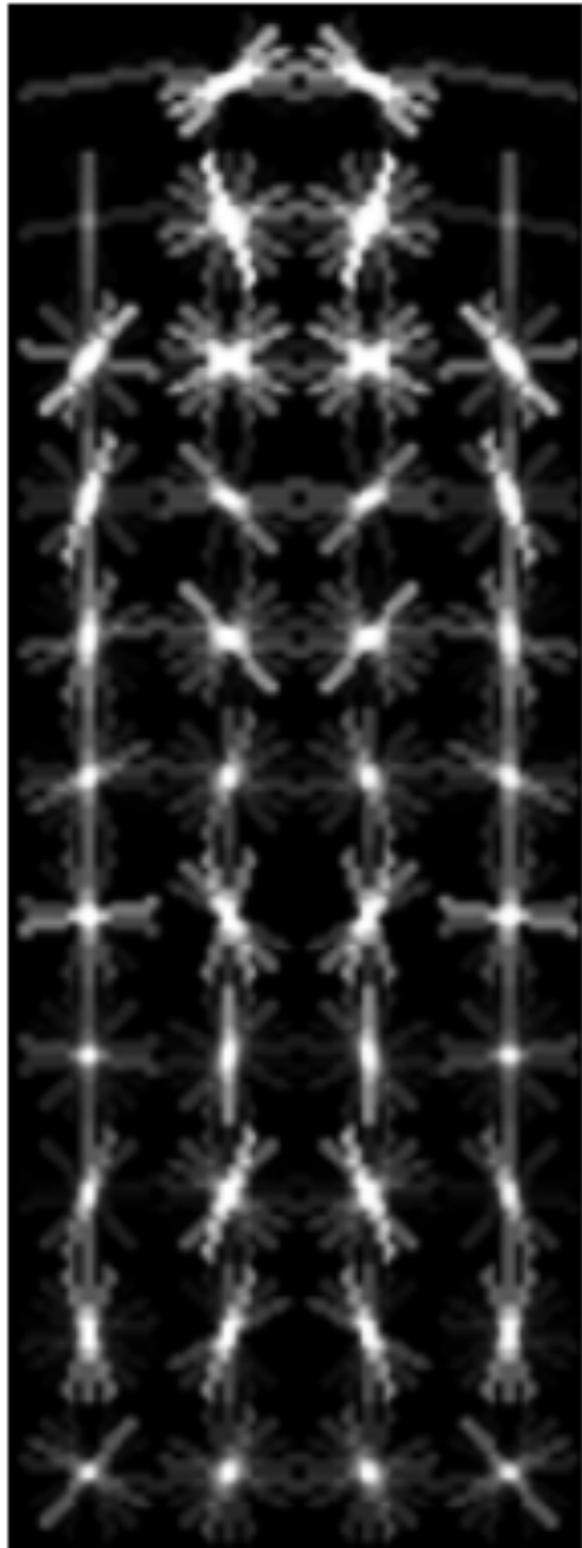
# Unsupervised learning



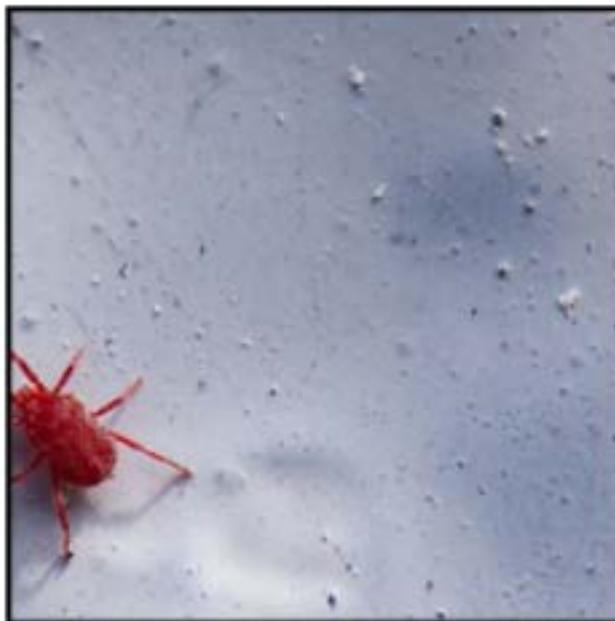
## Shape model



# Supervised learning



[Felzenszwalb et al. '10]



**mite**

**container ship**

**motor scooter**

**leopard**

	mite
	black widow
	cockroach
	tick
	starfish

	container ship
	lifeboat
	amphibian
	fireboat
	drilling platform

	motor scooter
	go-kart
	moped
	bumper car
	golfcart

	leopard
	jaguar
	cheetah
	snow leopard
	Egyptian cat



**grille**

**mushroom**

**cherry**

**Madagascar cat**

	convertible
	grille
	pickup
	beach wagon
	fire engine

	agaric
	mushroom
	jelly fungus
	gill fungus
	dead-man's-fingers

	dalmatian
	grape
	elderberry
	ffordshire bullterrier
	currant

	squirrel monkey
	spider monkey
	titi
	indri
	howler monkey

Need for human expertise



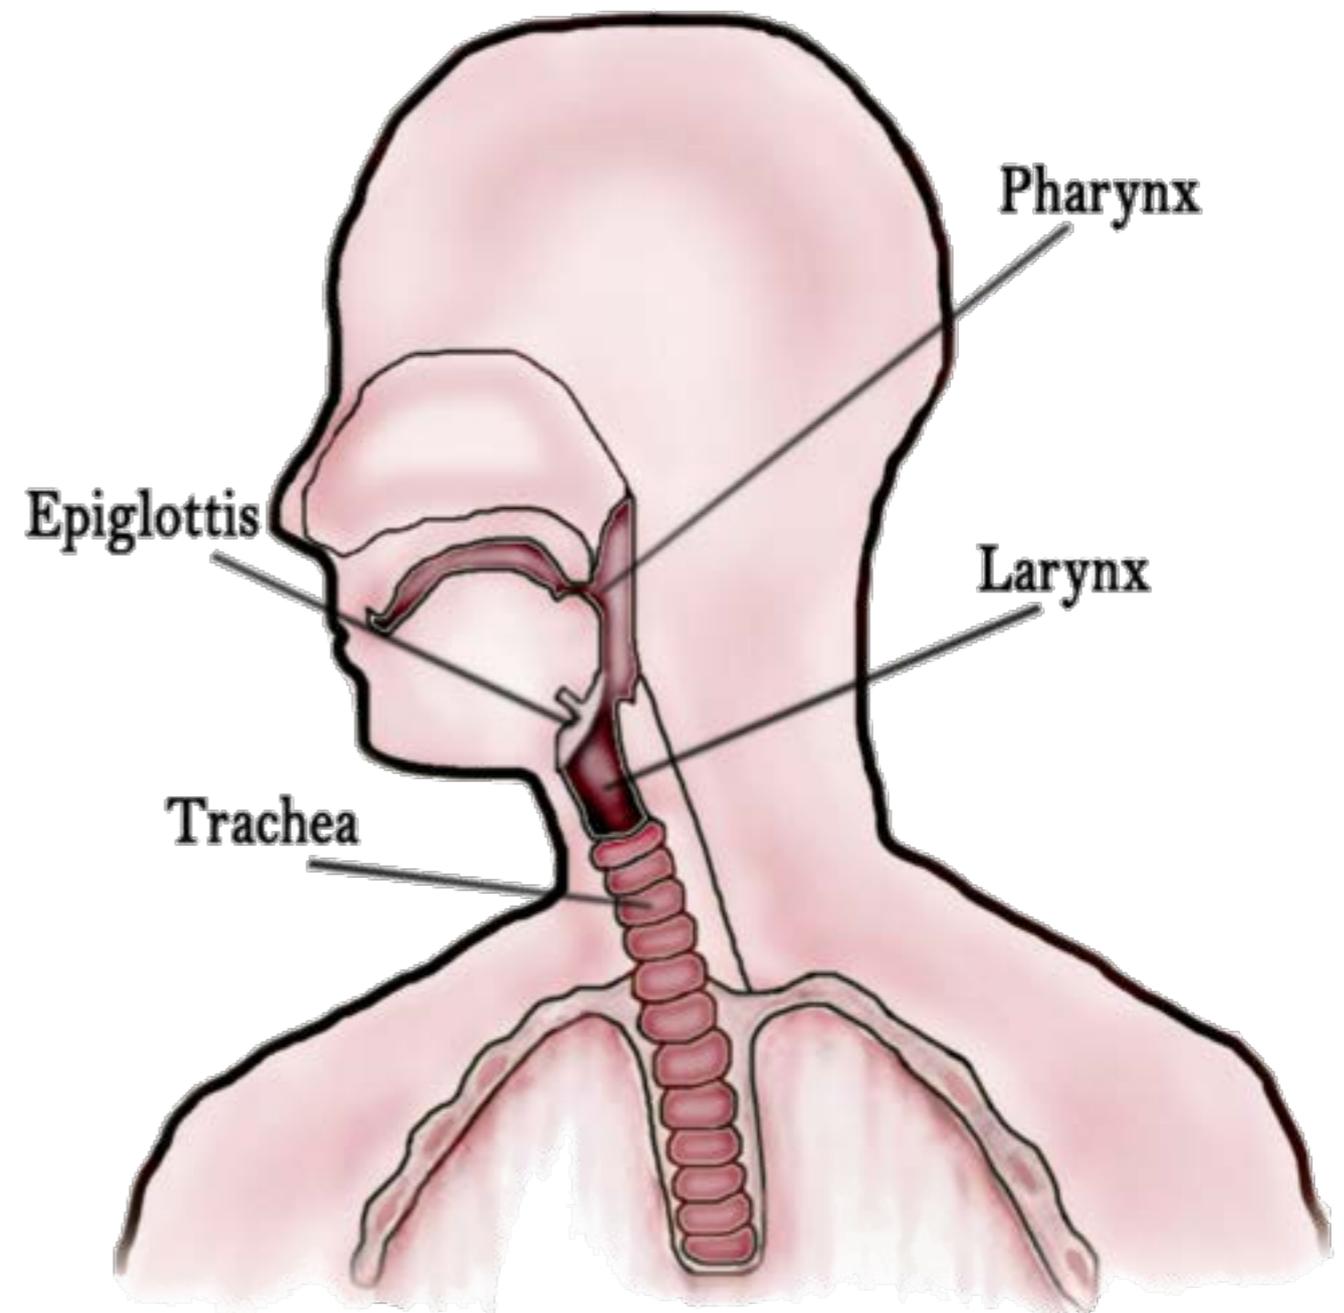
Need for human expertise

Throat



# Need for human expertise

Throat

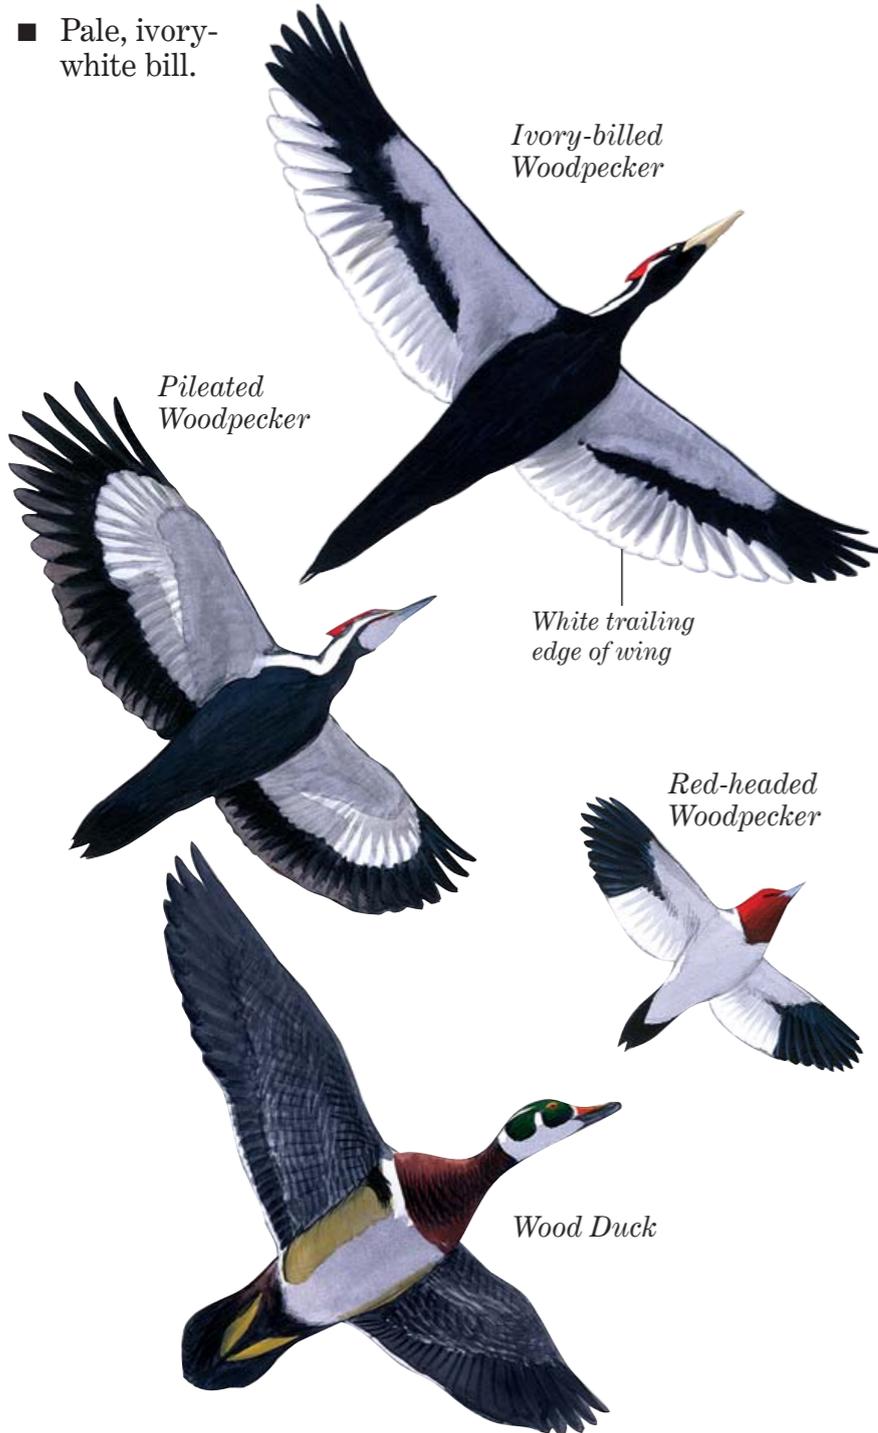


# Identifying Field Marks of an Ivory-billed Woodpecker and Similar Birds

## In flight - view from below

Distinct Ivory-billed Woodpecker characteristics:

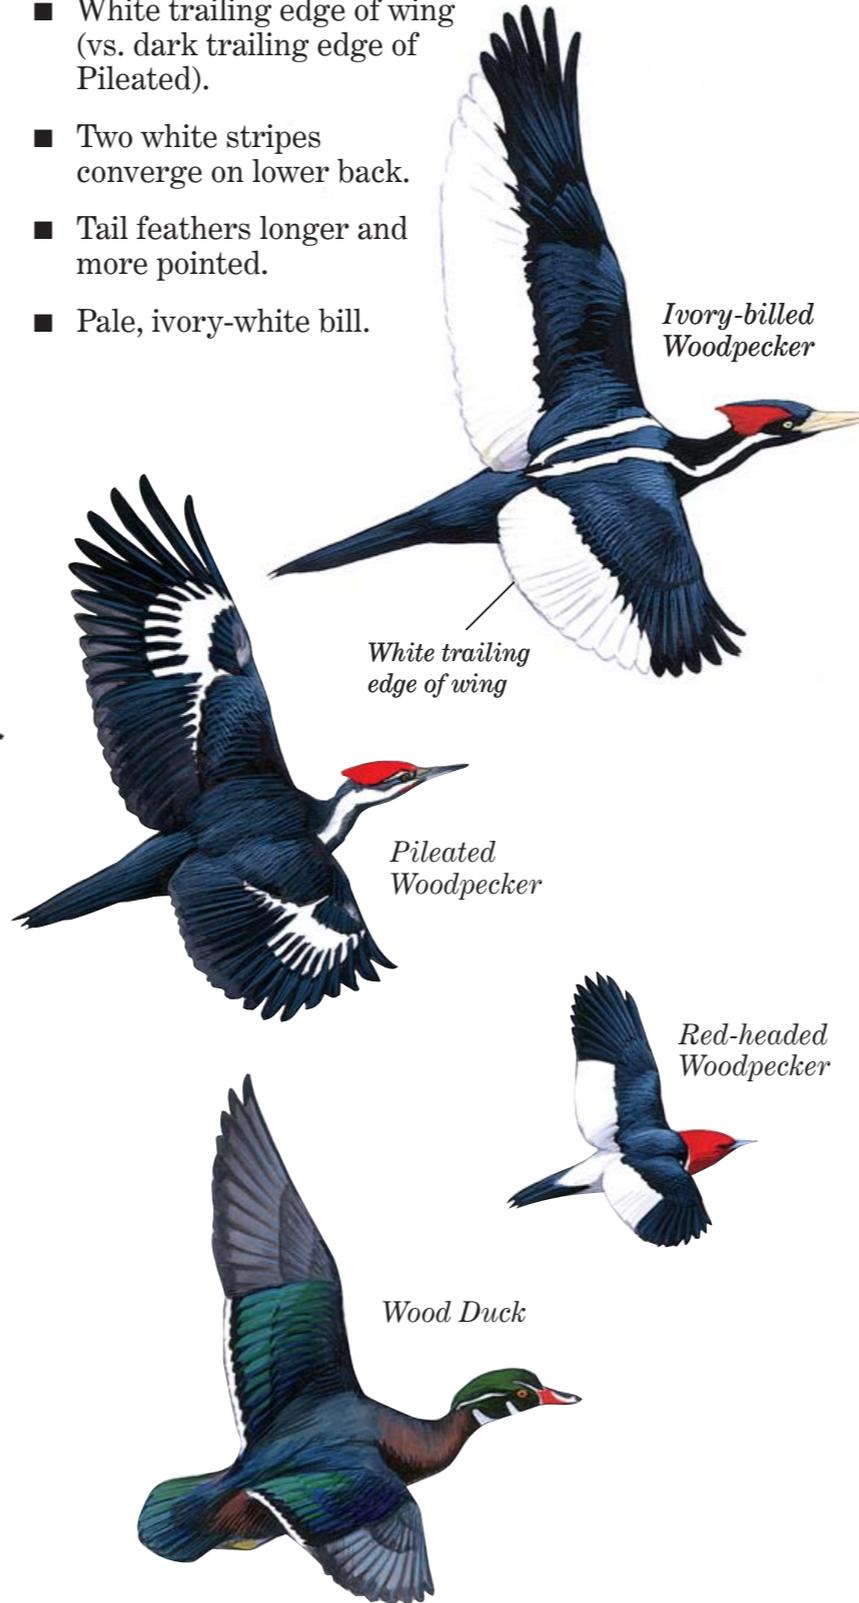
- White trailing edge of wing (vs. dark trailing edge of Pileated).
- Wing more slender than Pileated.
- Tail feathers longer and more pointed.
- Pale, ivory-white bill.



## In flight - view from above

Distinct Ivory-billed Woodpecker characteristics:

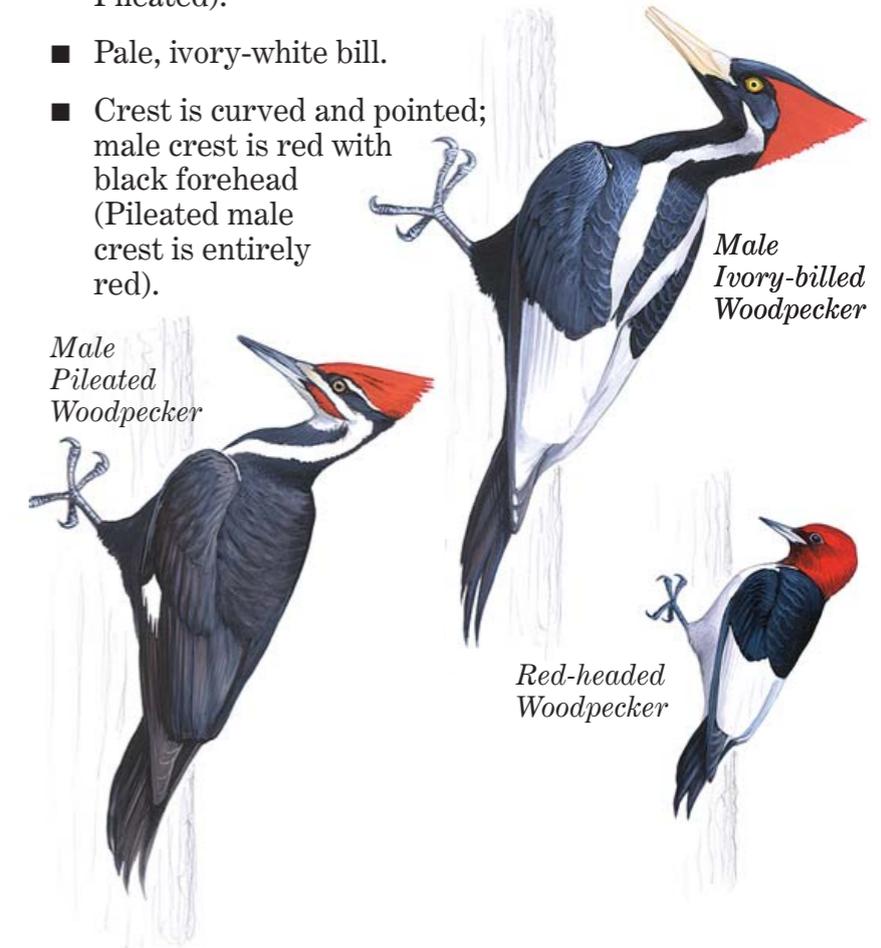
- White trailing edge of wing (vs. dark trailing edge of Pileated).
- Two white stripes converge on lower back.
- Tail feathers longer and more pointed.
- Pale, ivory-white bill.



## At rest

Distinct Ivory-billed Woodpecker characteristics:

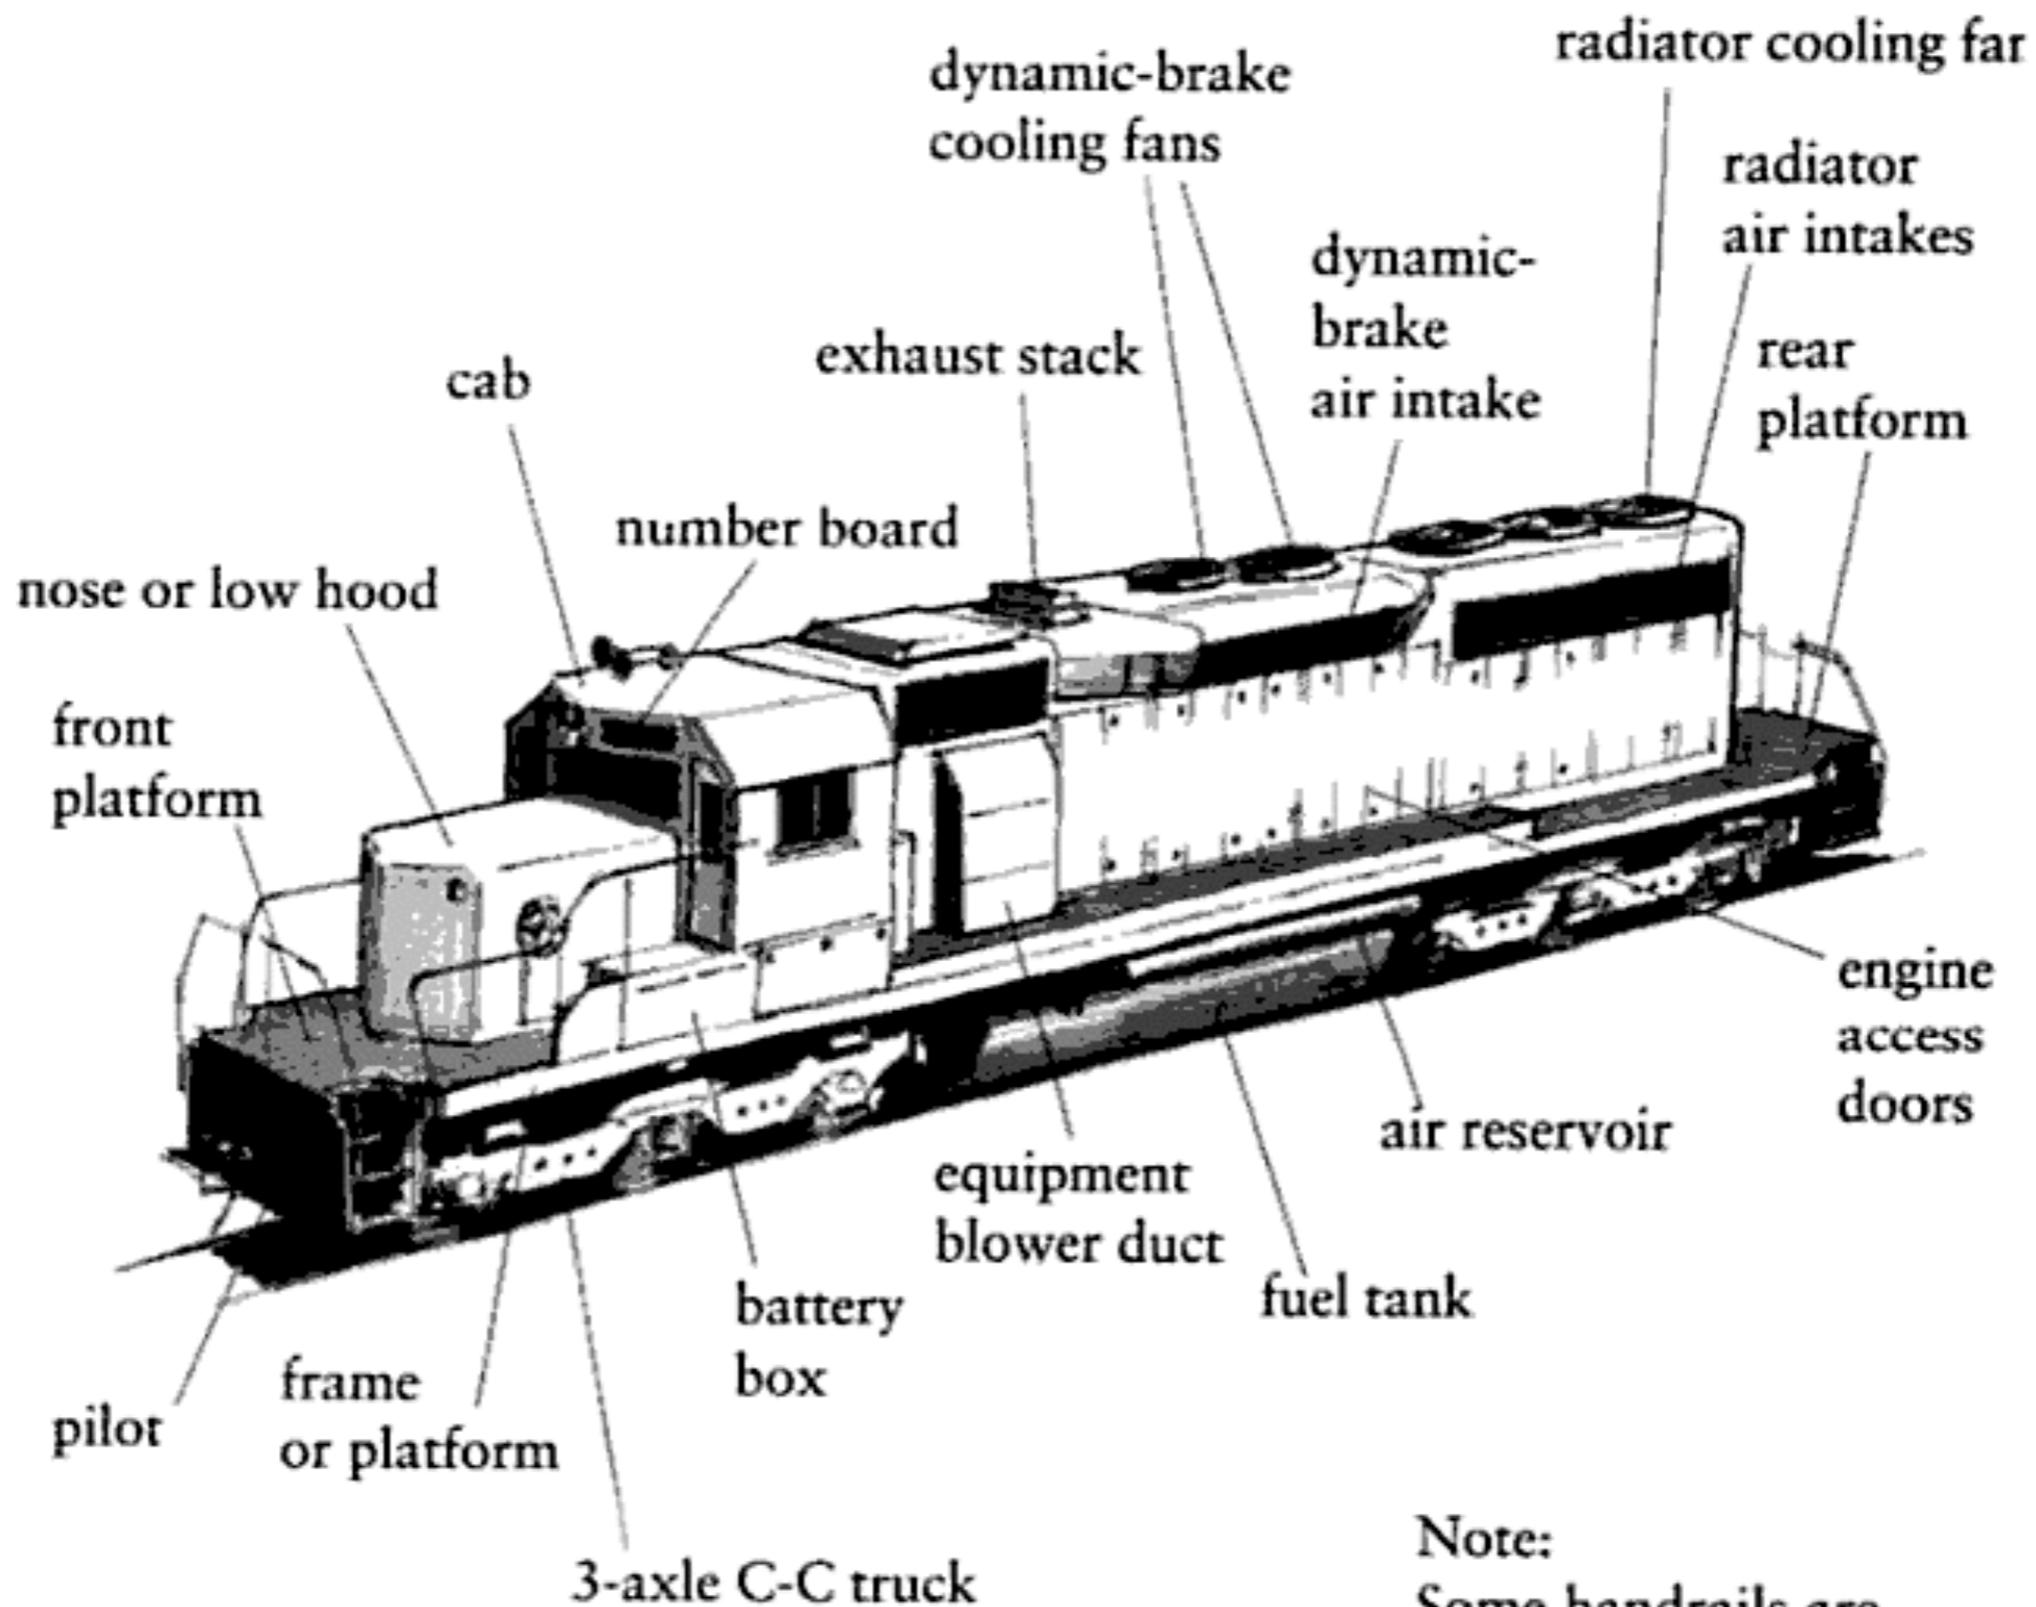
- Two white stripes converge on lower back.
- Entirely white secondary feathers give appearance of white "saddle" on back.
- Largely dark face and dark chin (vs. white chin of Pileated).
- Pale, ivory-white bill.
- Crest is curved and pointed; male crest is red with black forehead (Pileated male crest is entirely red).



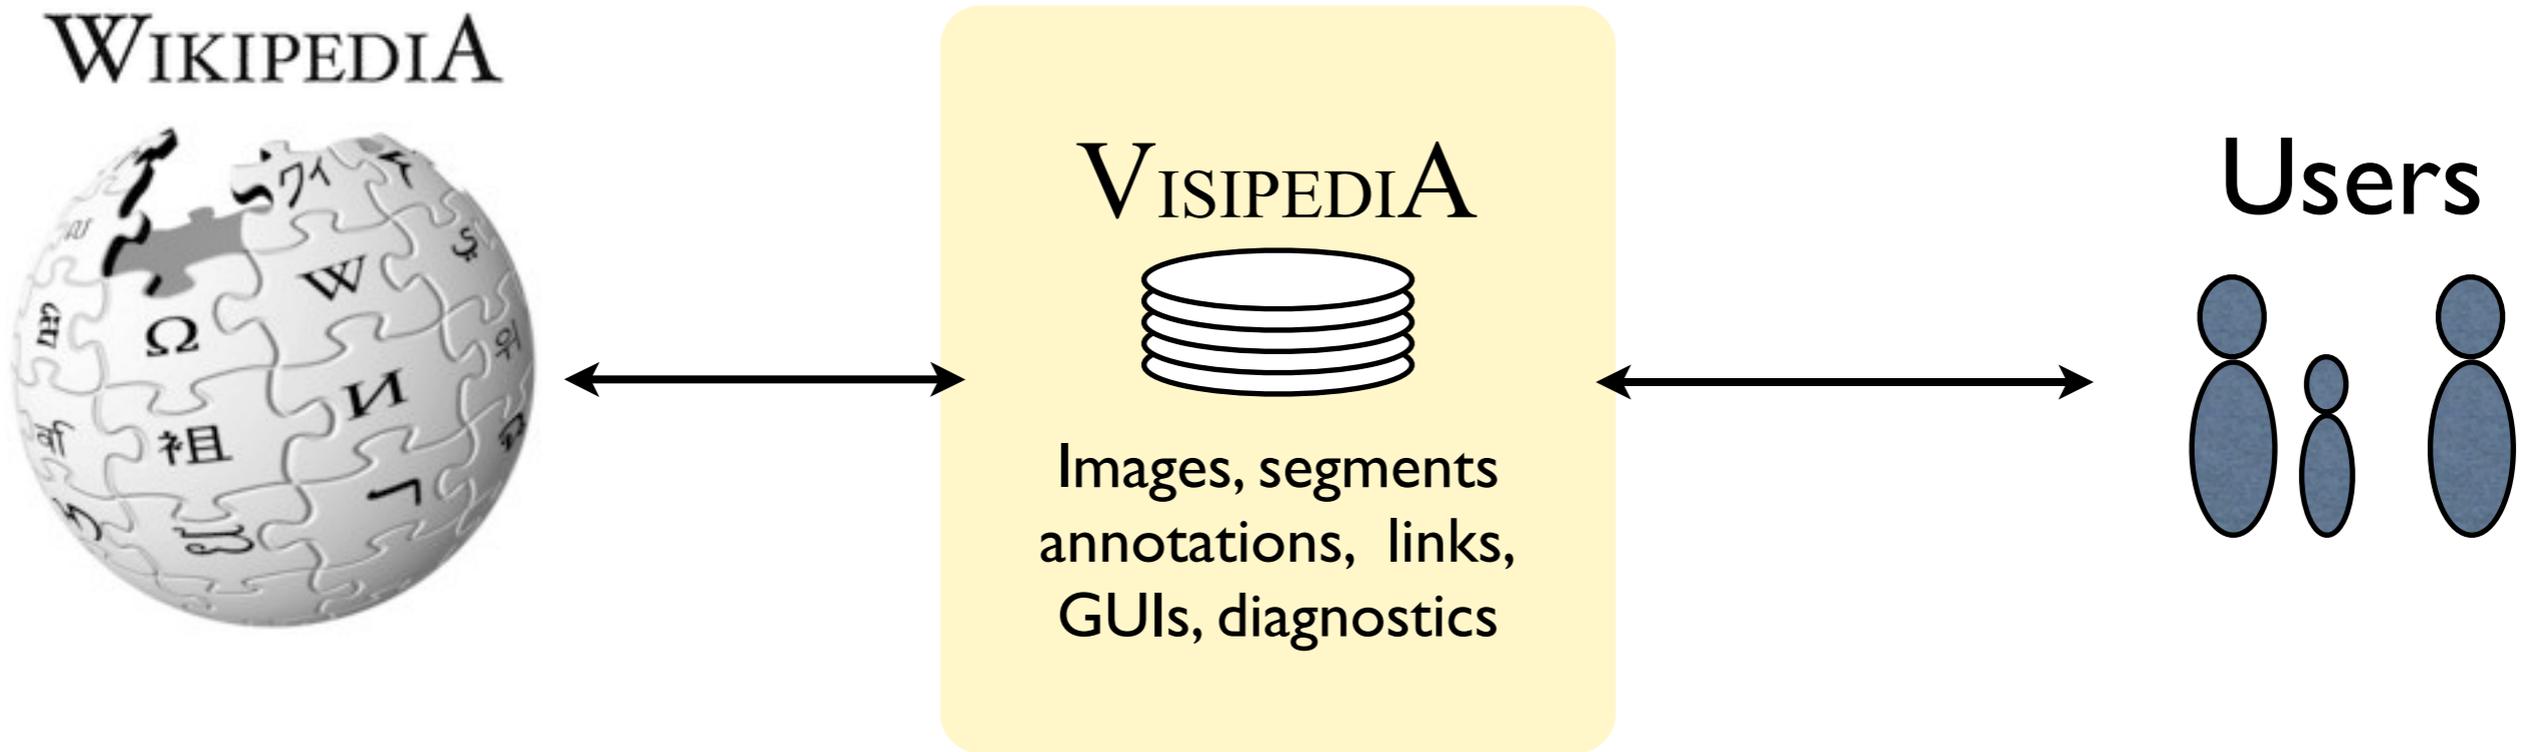
Female Head

- Female Ivory-bill crest is entirely black (female Pileated crest resembles male ivory-billed red crest with black forehead – use chin color as distinguishing feature)





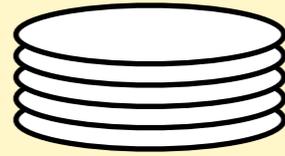
## THE PARTS OF A LOCOMOTIVE



WIKIPEDIA



VISIPEDIA



Images, segments  
annotations, links,  
GUIs, diagnostics

Users

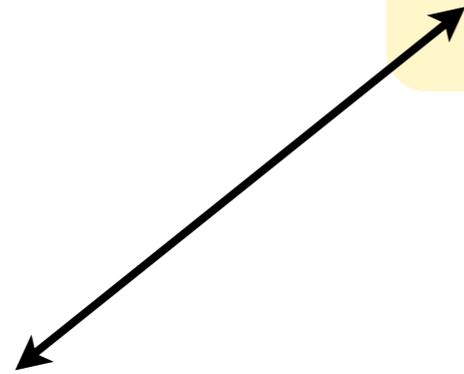
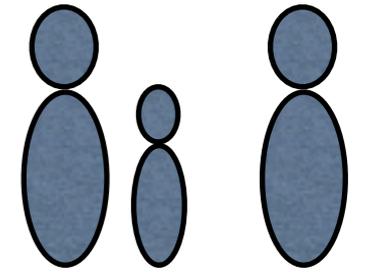
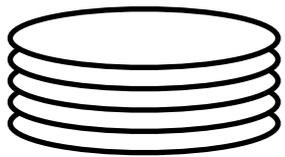


Image databases



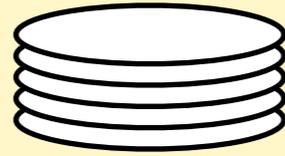
flickr®

IMAGENET

WIKIPEDIA



VISIPEDIA



Images, segments  
annotations, links,  
GUIs, diagnostics

Users

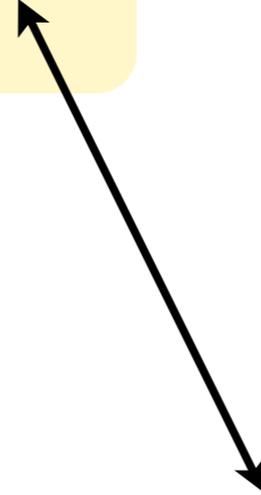
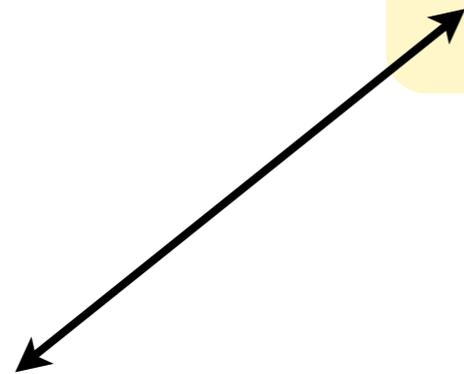
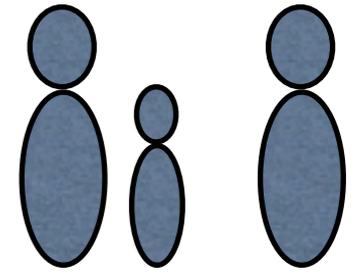
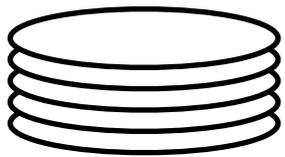


Image databases

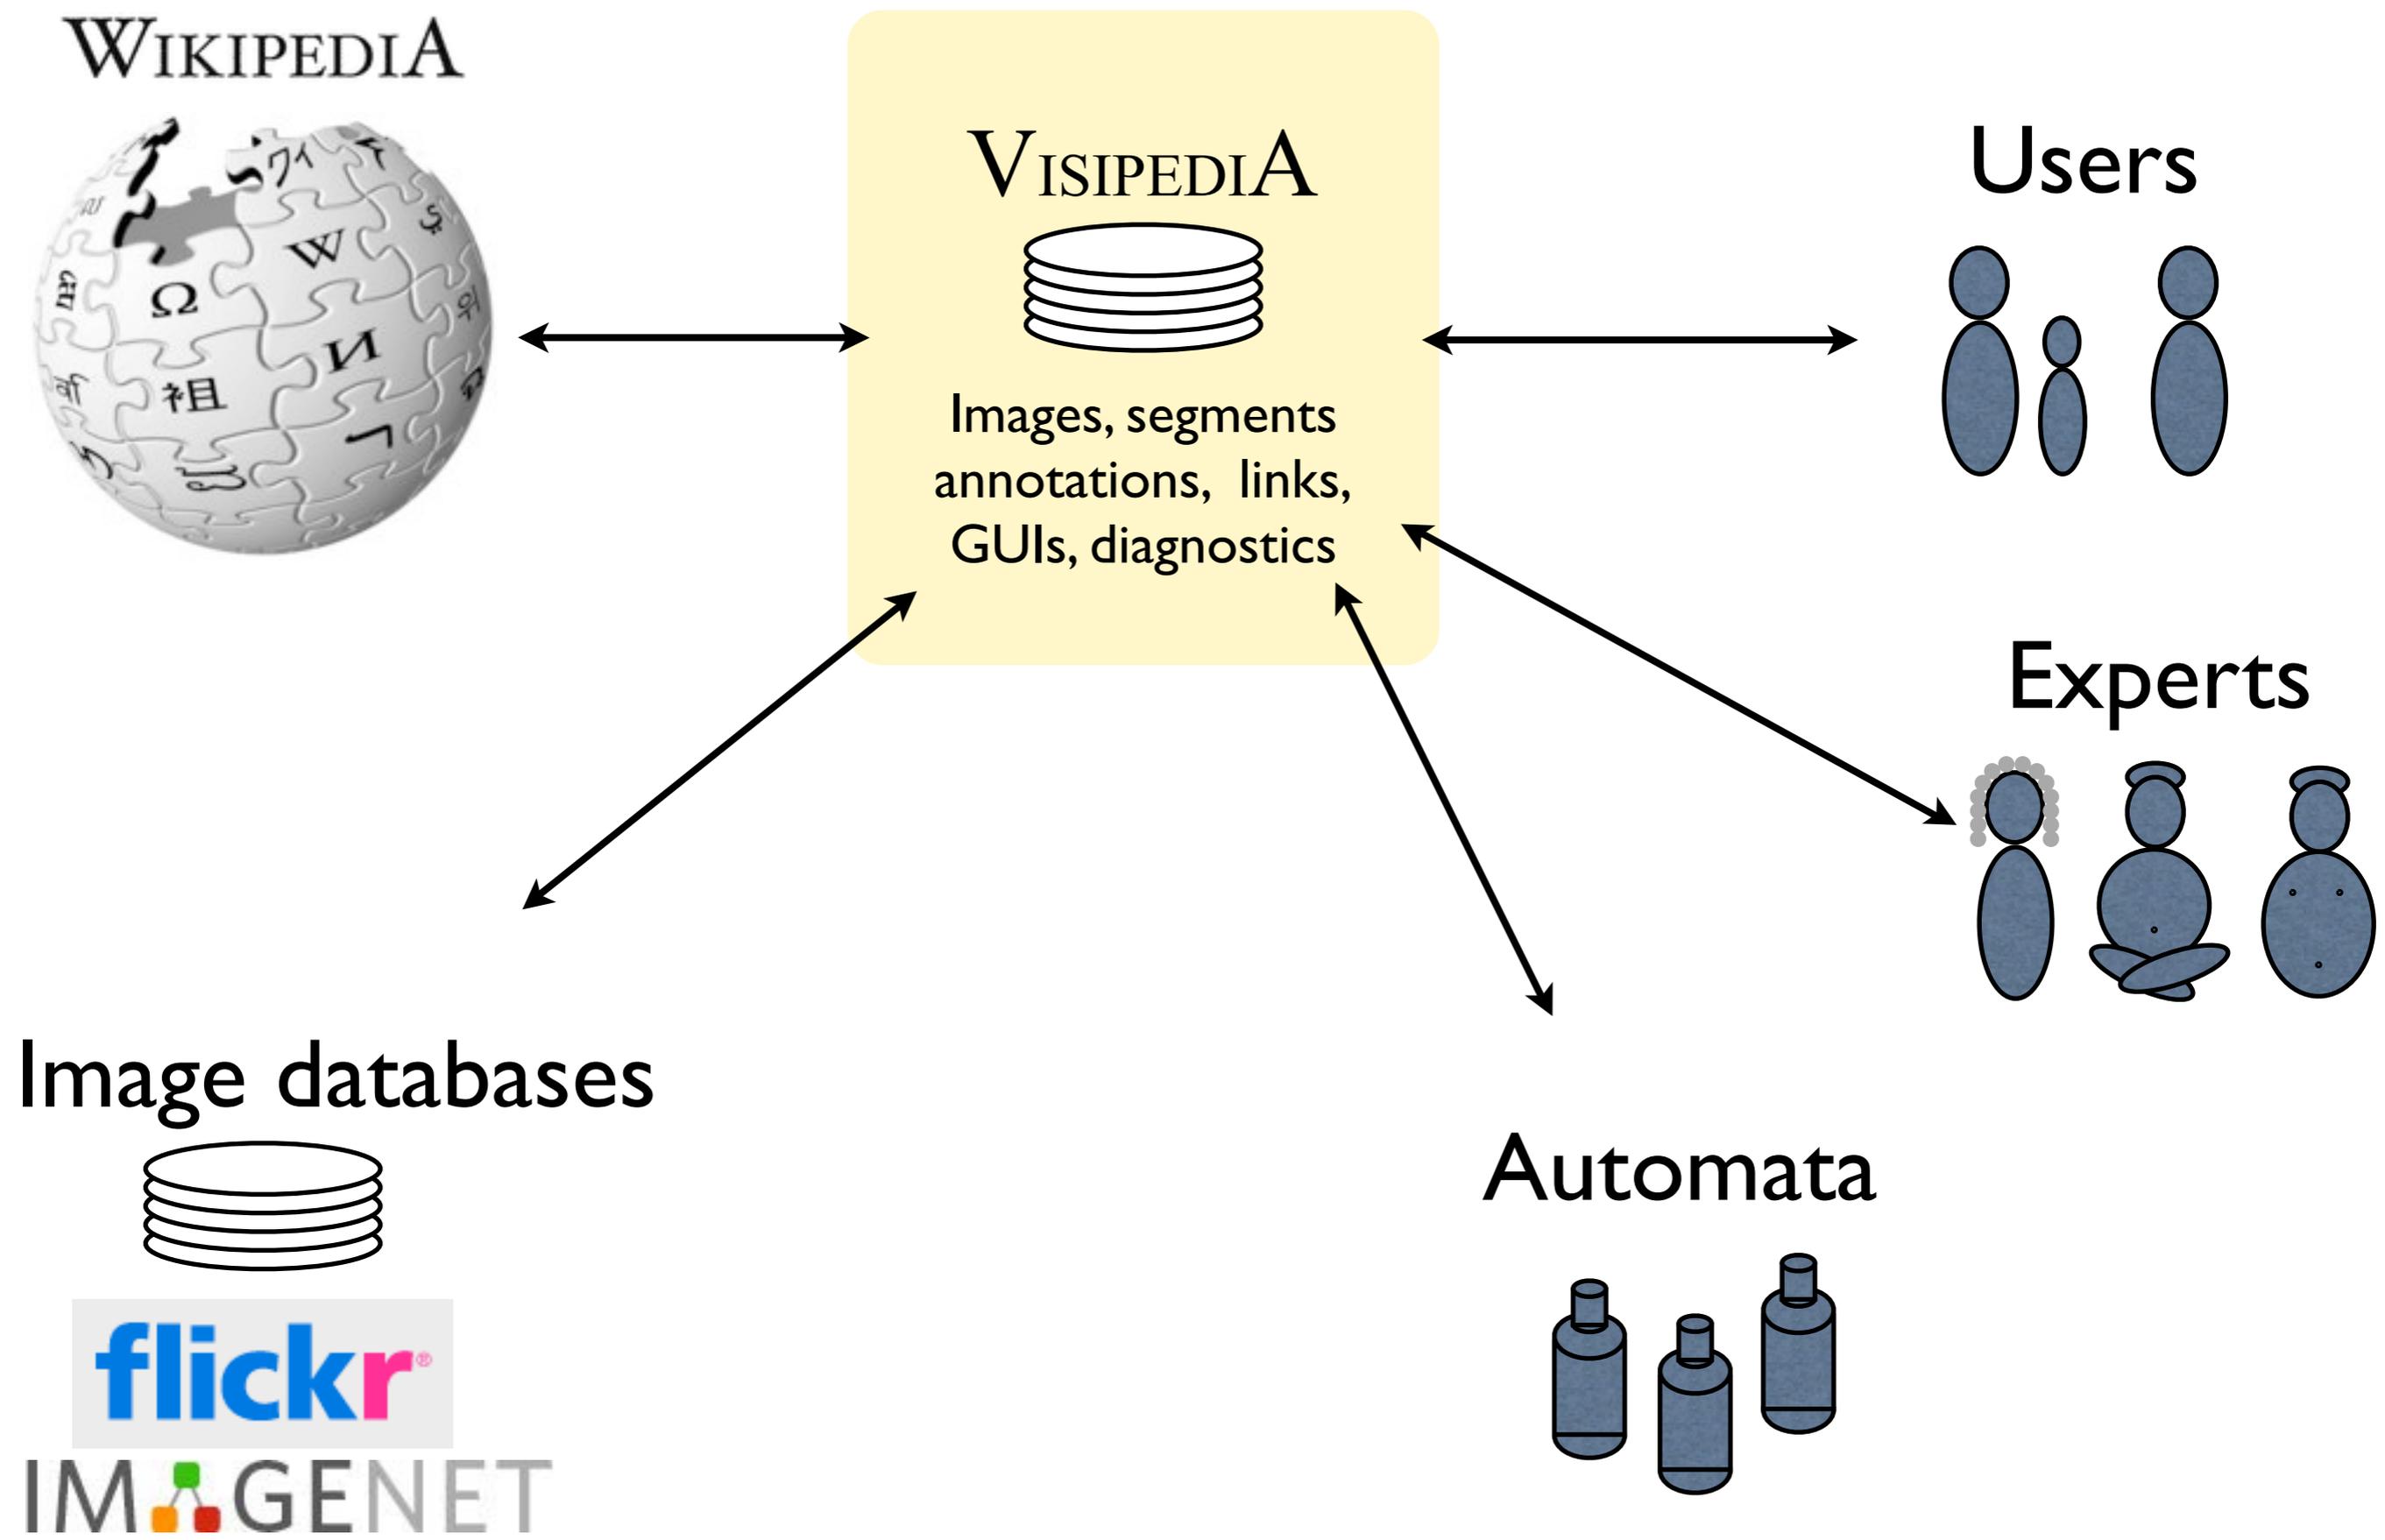


flickr®

IMAGENET

Automata

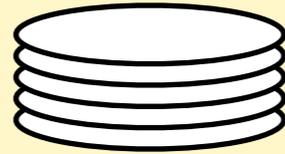




WIKIPEDIA

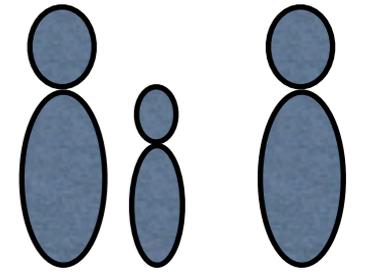


VISIPEDIA



Images, segments  
annotations, links,  
GUIs, diagnostics

Users



Experts

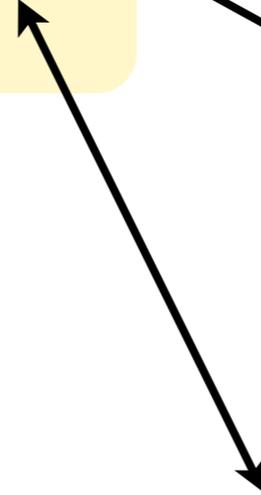
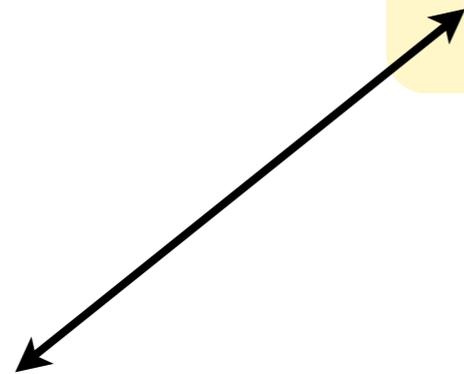
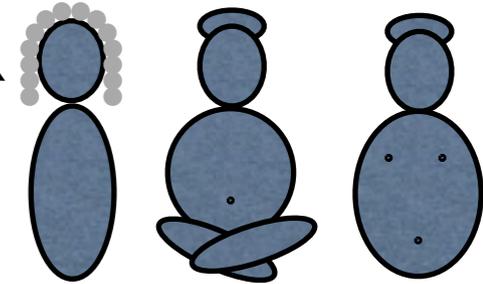
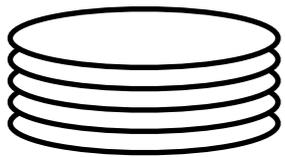
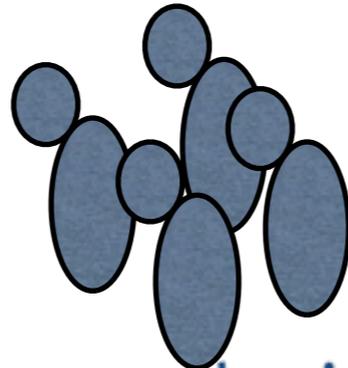


Image databases



Annotators

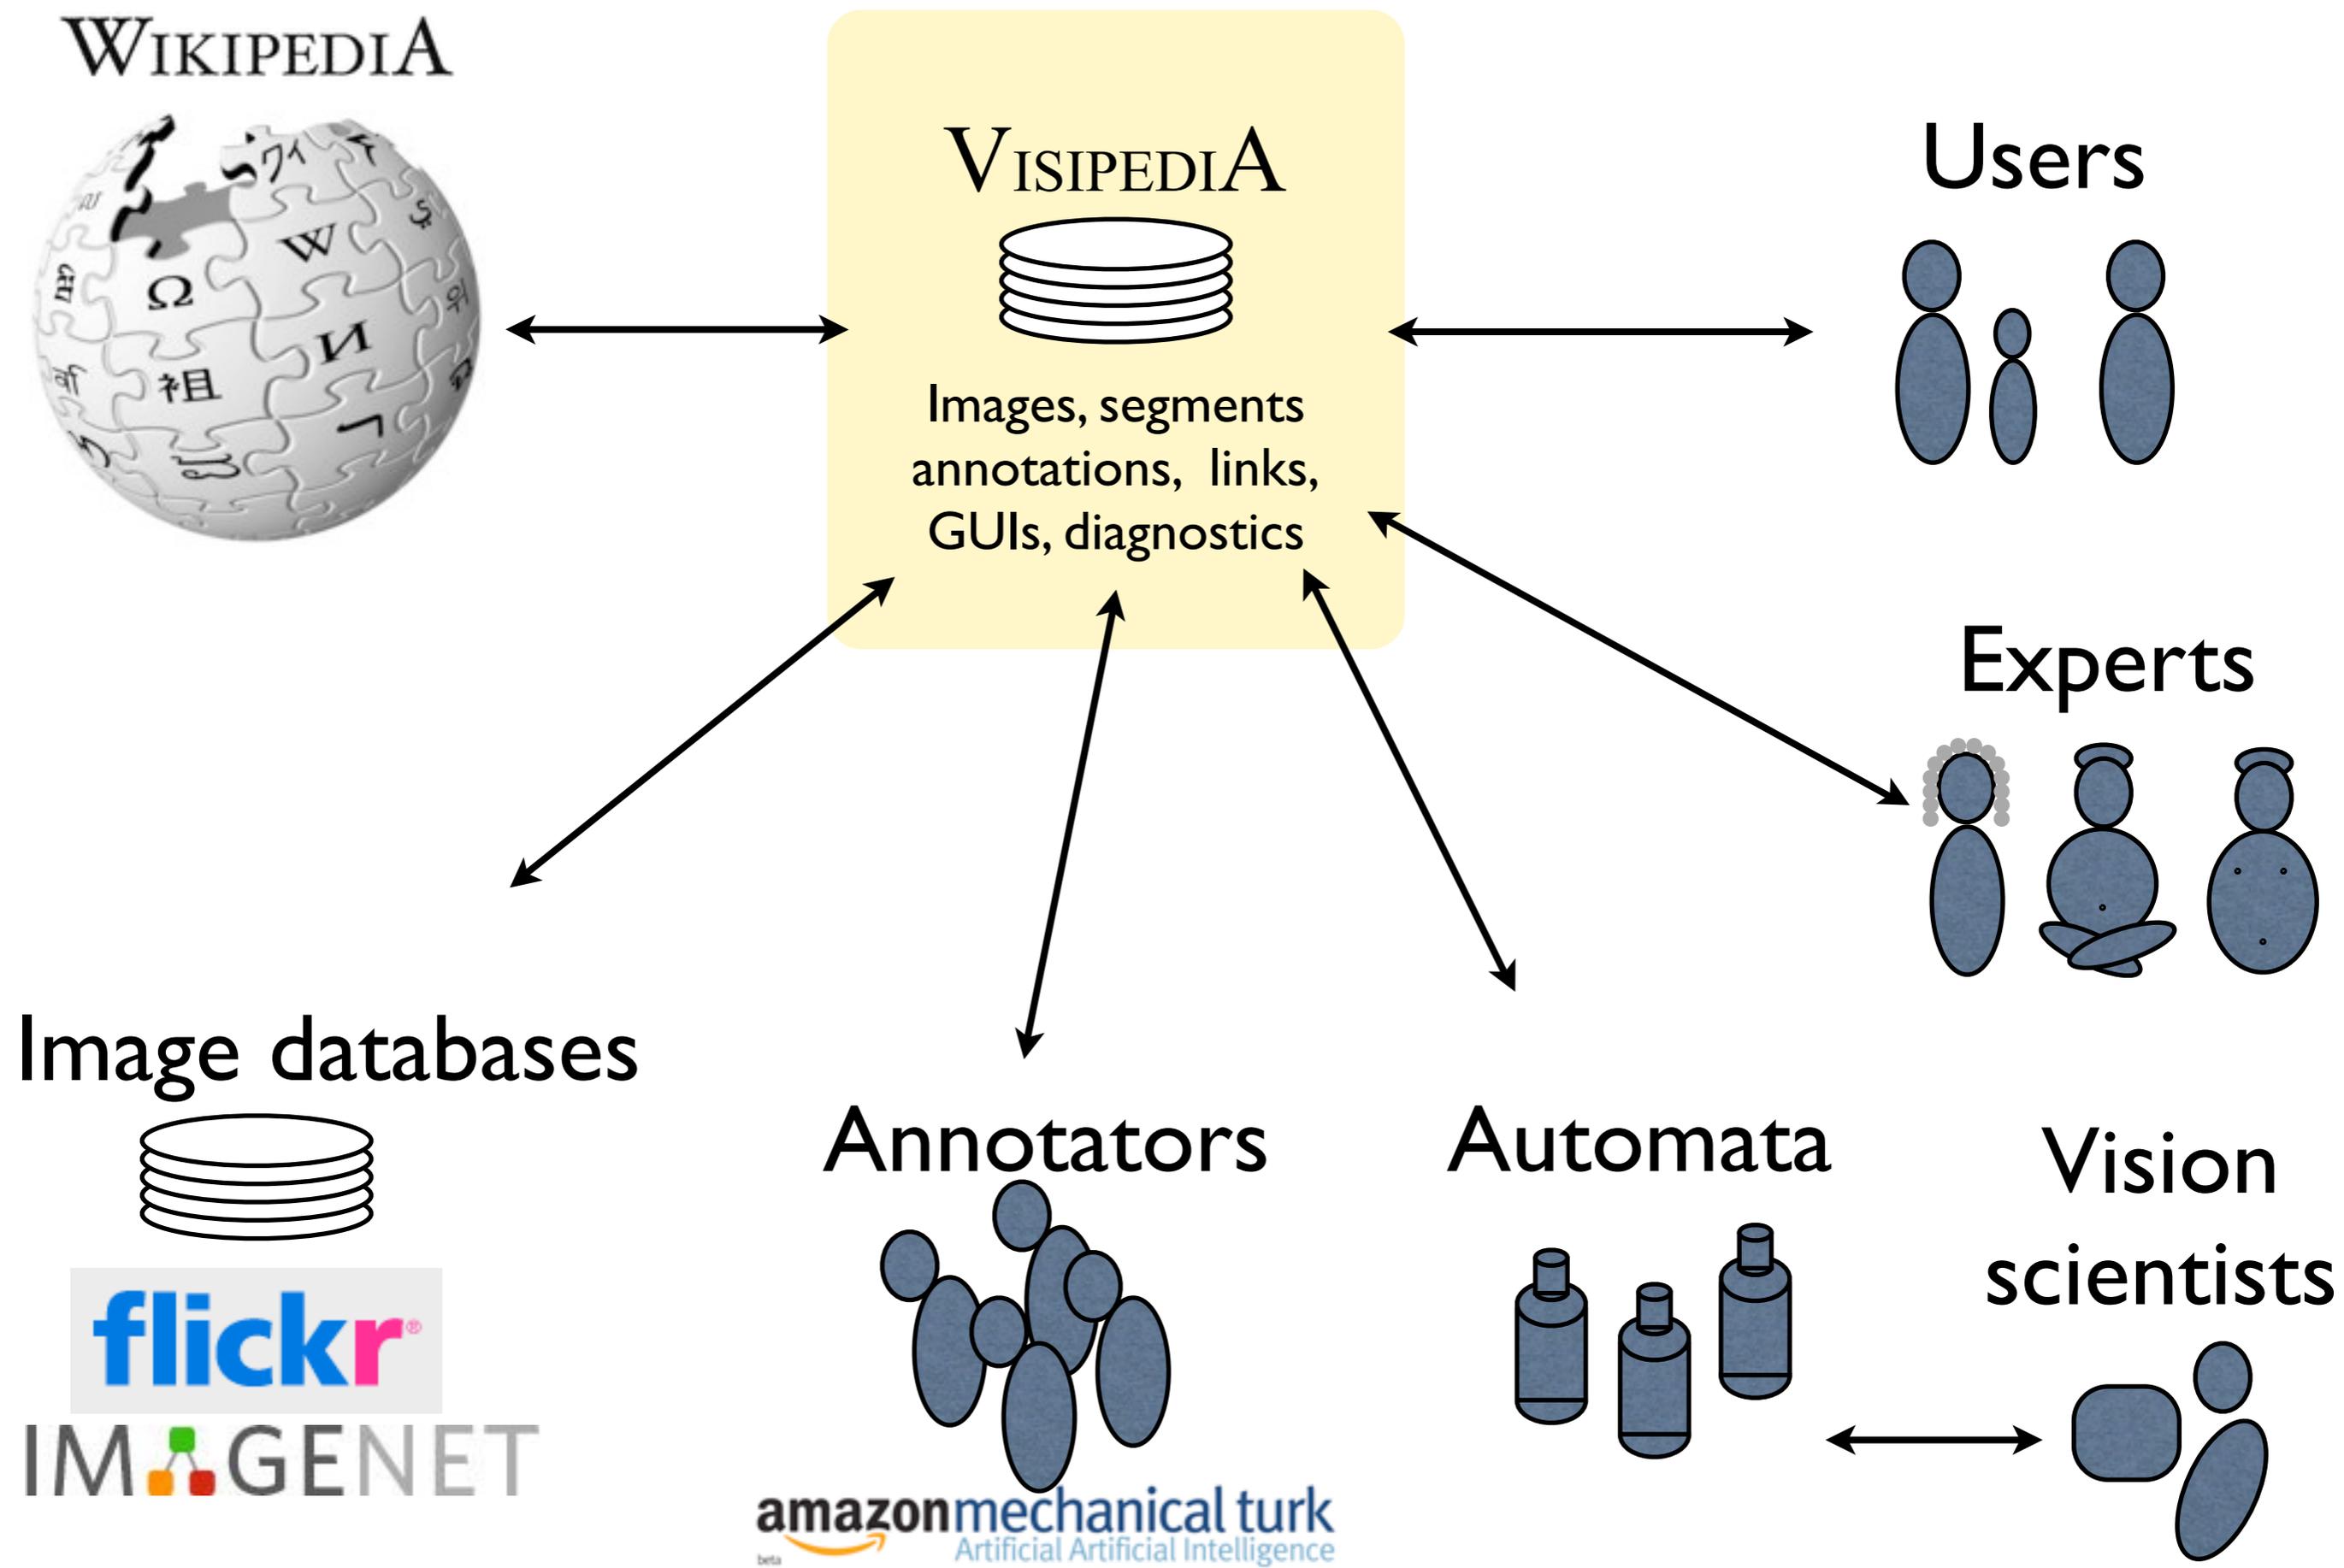


Automata



IMAGENET





# Crowdsourcing image annotation

[Welinder et al., NIPS2010]

# Indigo Bunting



Search

Photos | Groups | People

Everyone's Uploads

indigo bunting

SEARCH

Full Text | Tags Only  
Advanced Search

Sort: Relevant | Recent | Interesting

View: Small | Medium | Detail | Slideshow



From I Bird 2



From The Nature...



From Momba...



From davidcreebir...



From naturelover2...



From K\_Alinka



From davidcreebir...



From William ...



From violetfm



From jbobbe



From RitaK.



From mayalu



From ff151



From prairiedog



From redow



From ajnaturephot...



From Ken...



From R Hanson



From [Christine]



From davidcreebir...



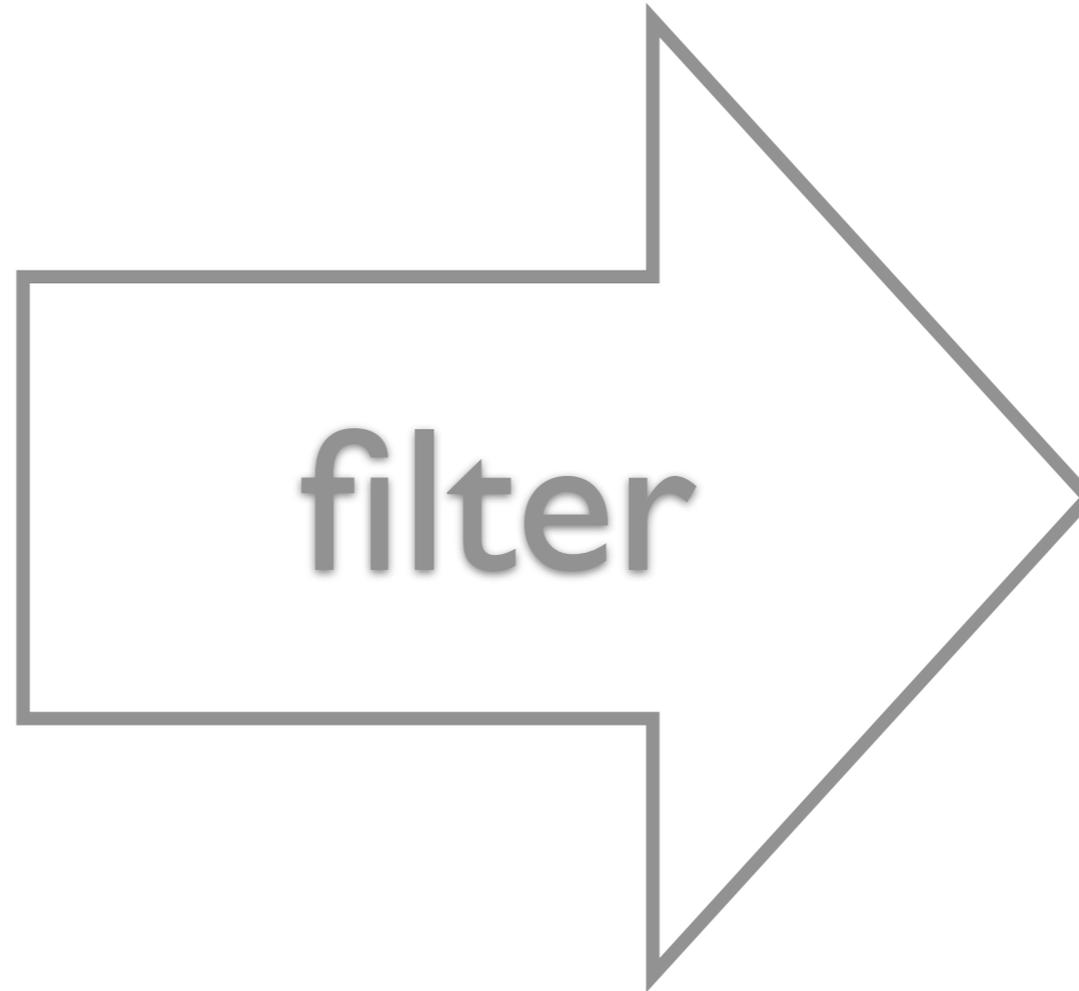
From reemac640

5,926 images

6000 images  
from flickr.com



# Building datasets



100s of  
training images

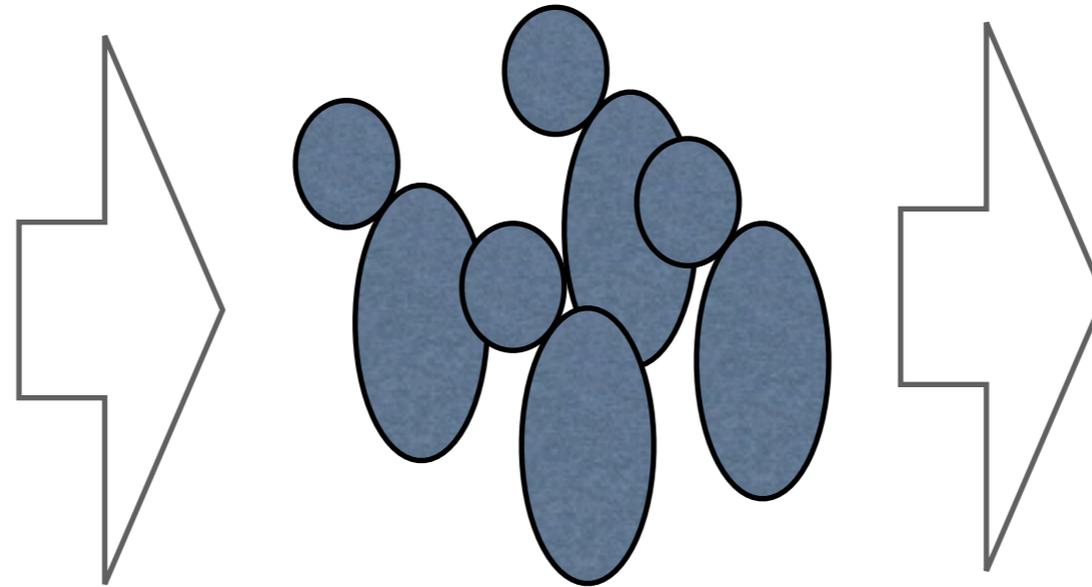


6000 images  
from flickr.com

# Building datasets

100s of  
training images

## Annotators



amazonmechanical turk  
beta Artificial Intelligence

Is there an Indigo bunting in the image?

# Find the Indigo Bunting



# Find the Indigo Bunting



A



A



A



A



A



A

# Find the Indigo Bunting



A B C D



A B C D



A B C D



A B C D



A B C D



A B C D

# Characterizing annotators: types of errors

Indigo Bunting?

Yes

No



Hit

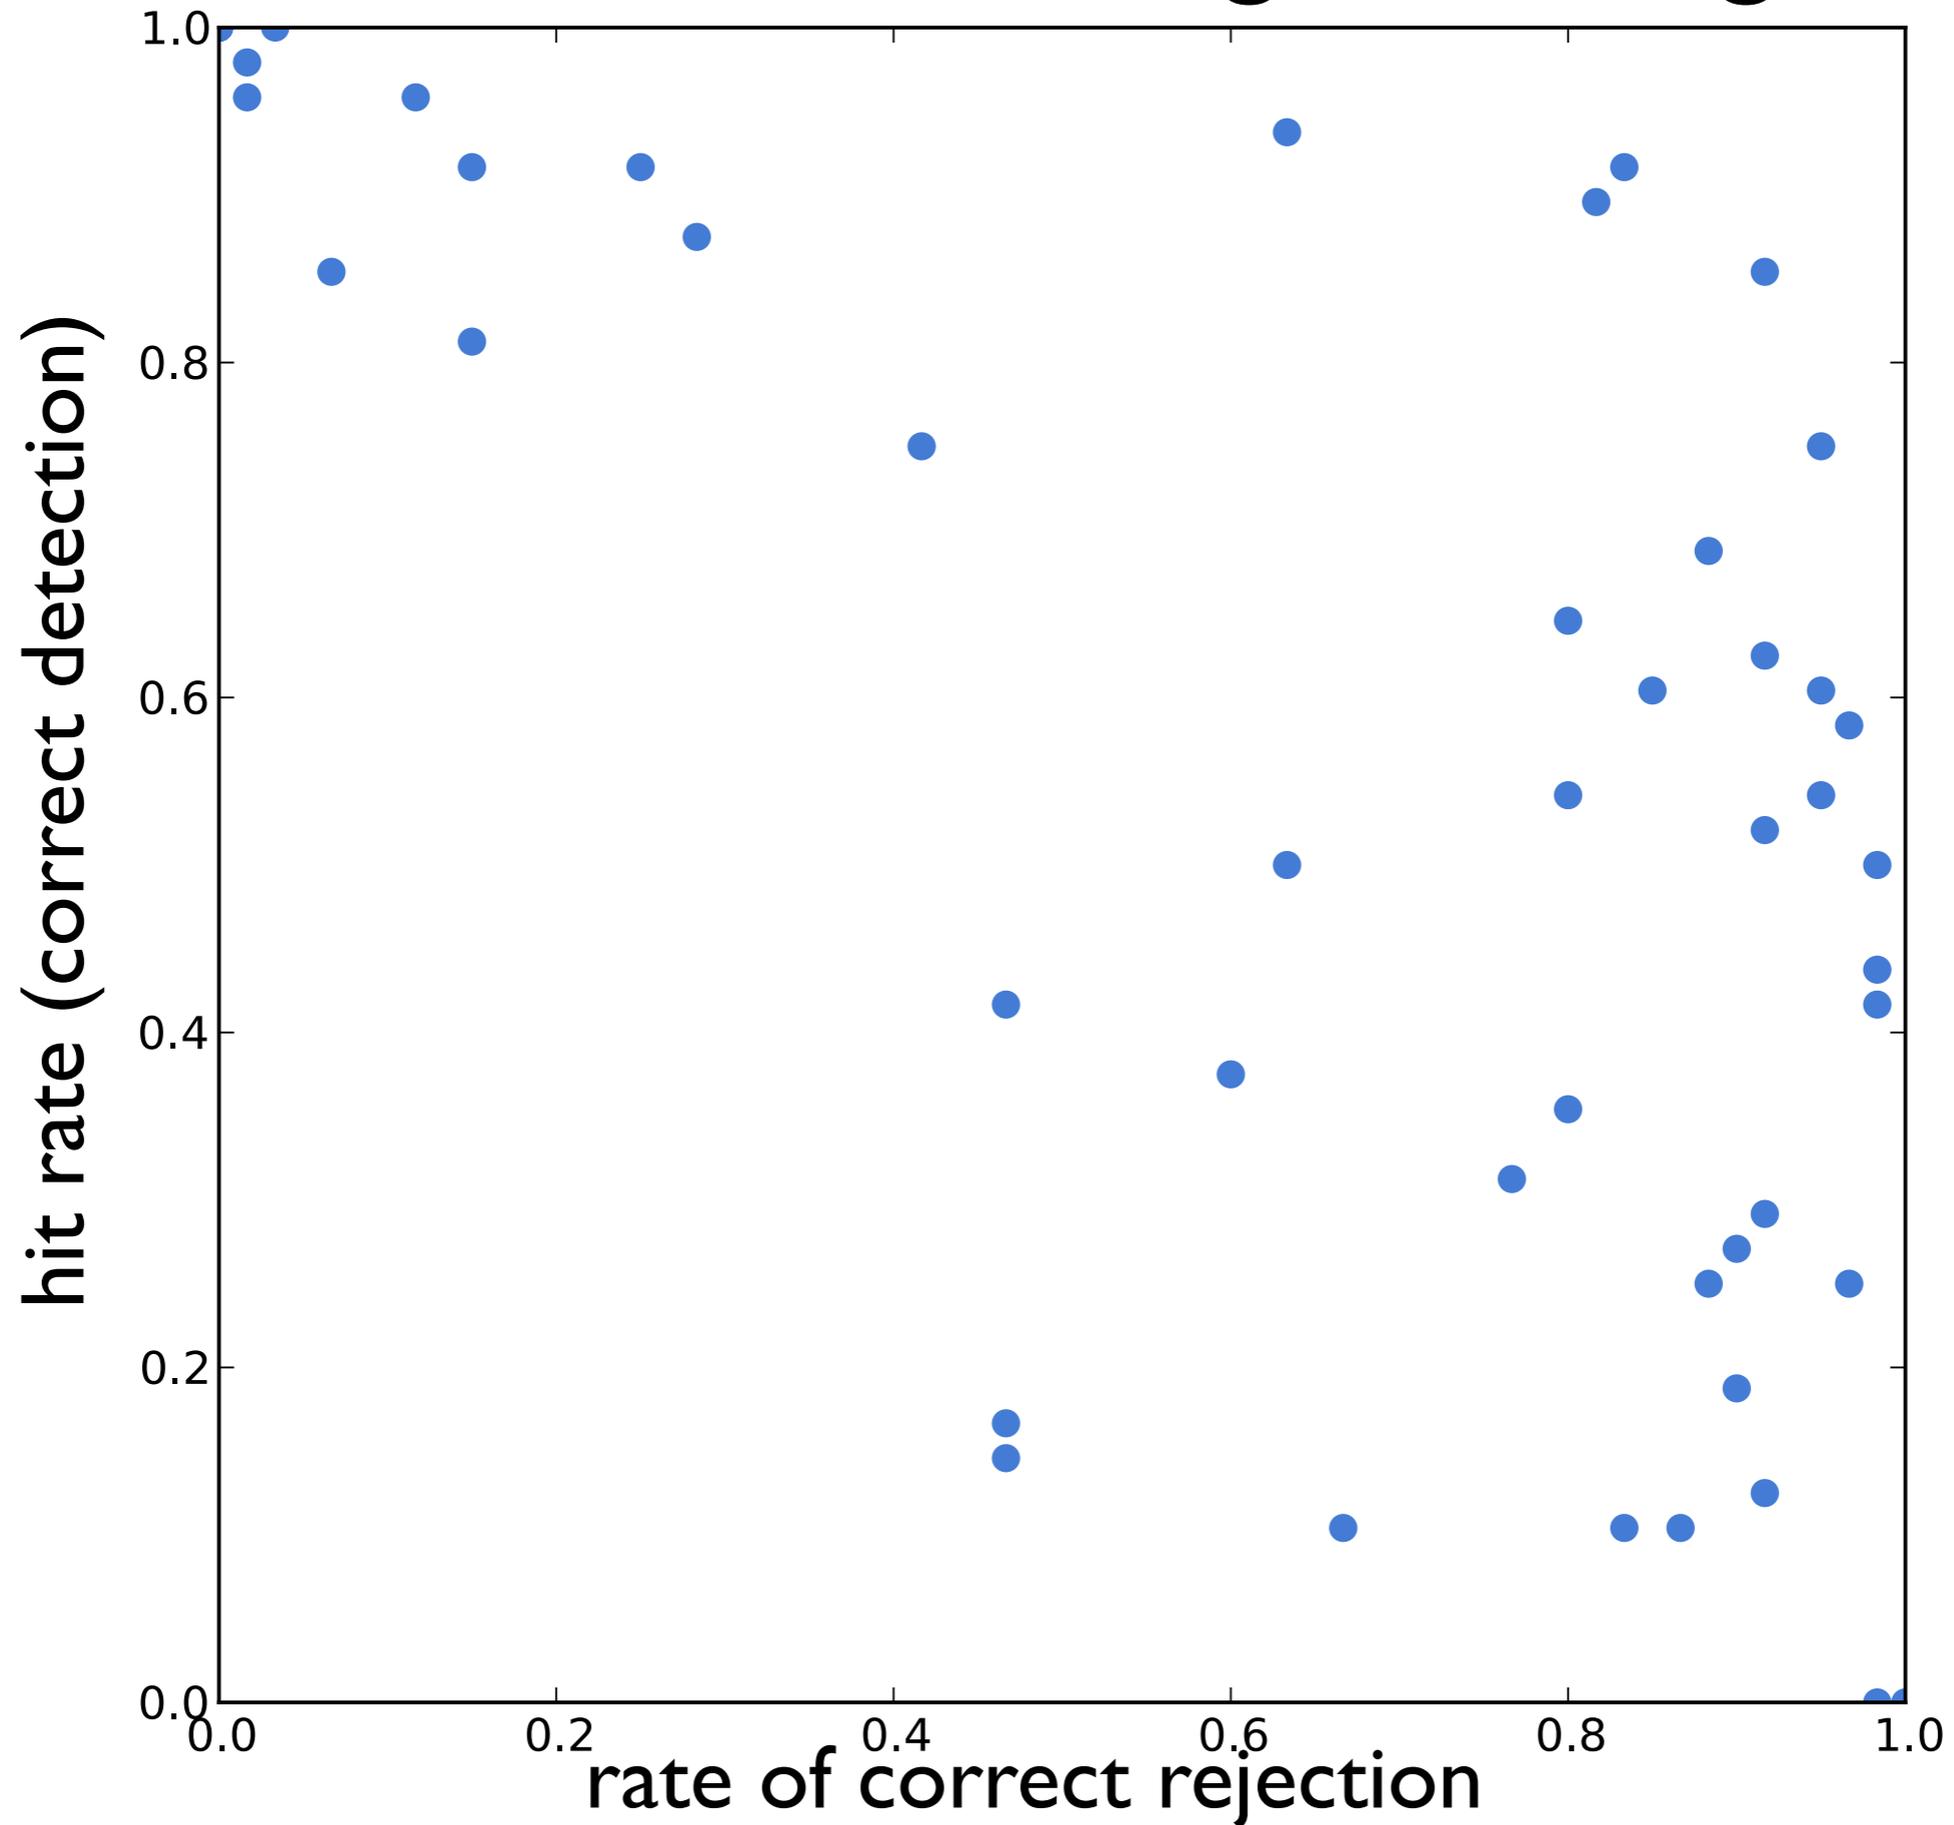
Miss



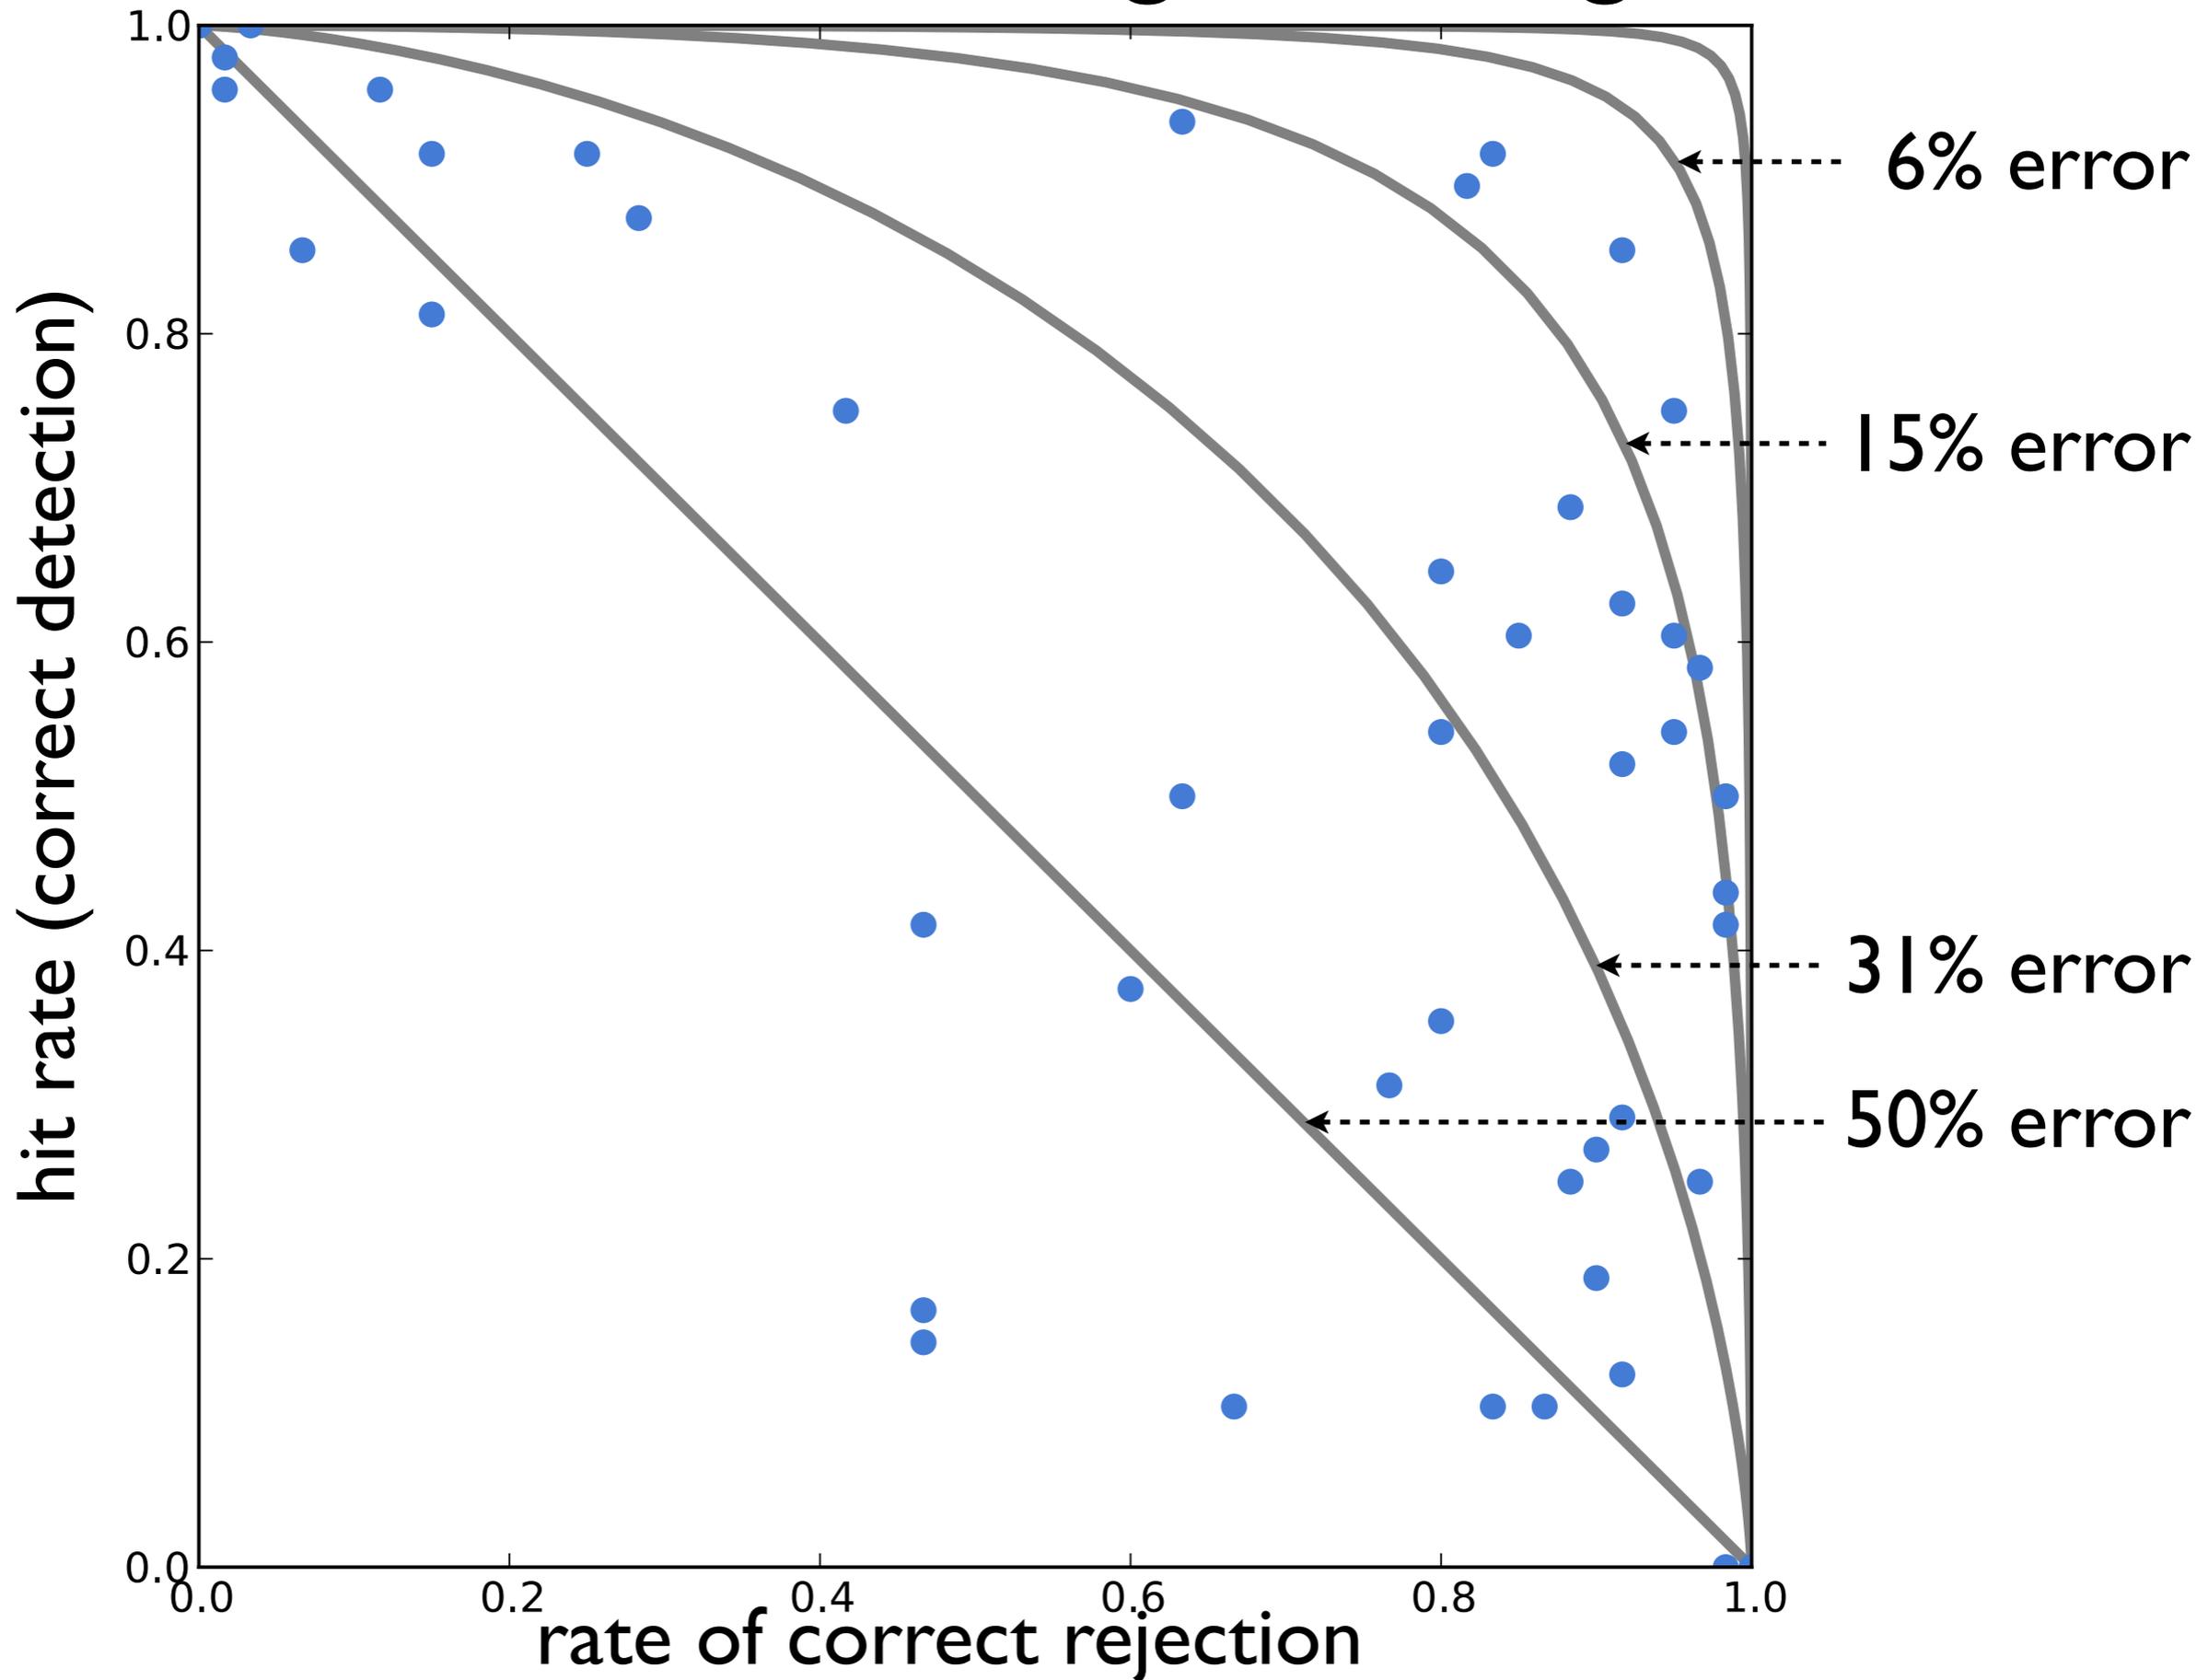
False Alarm

Correct  
Rejection

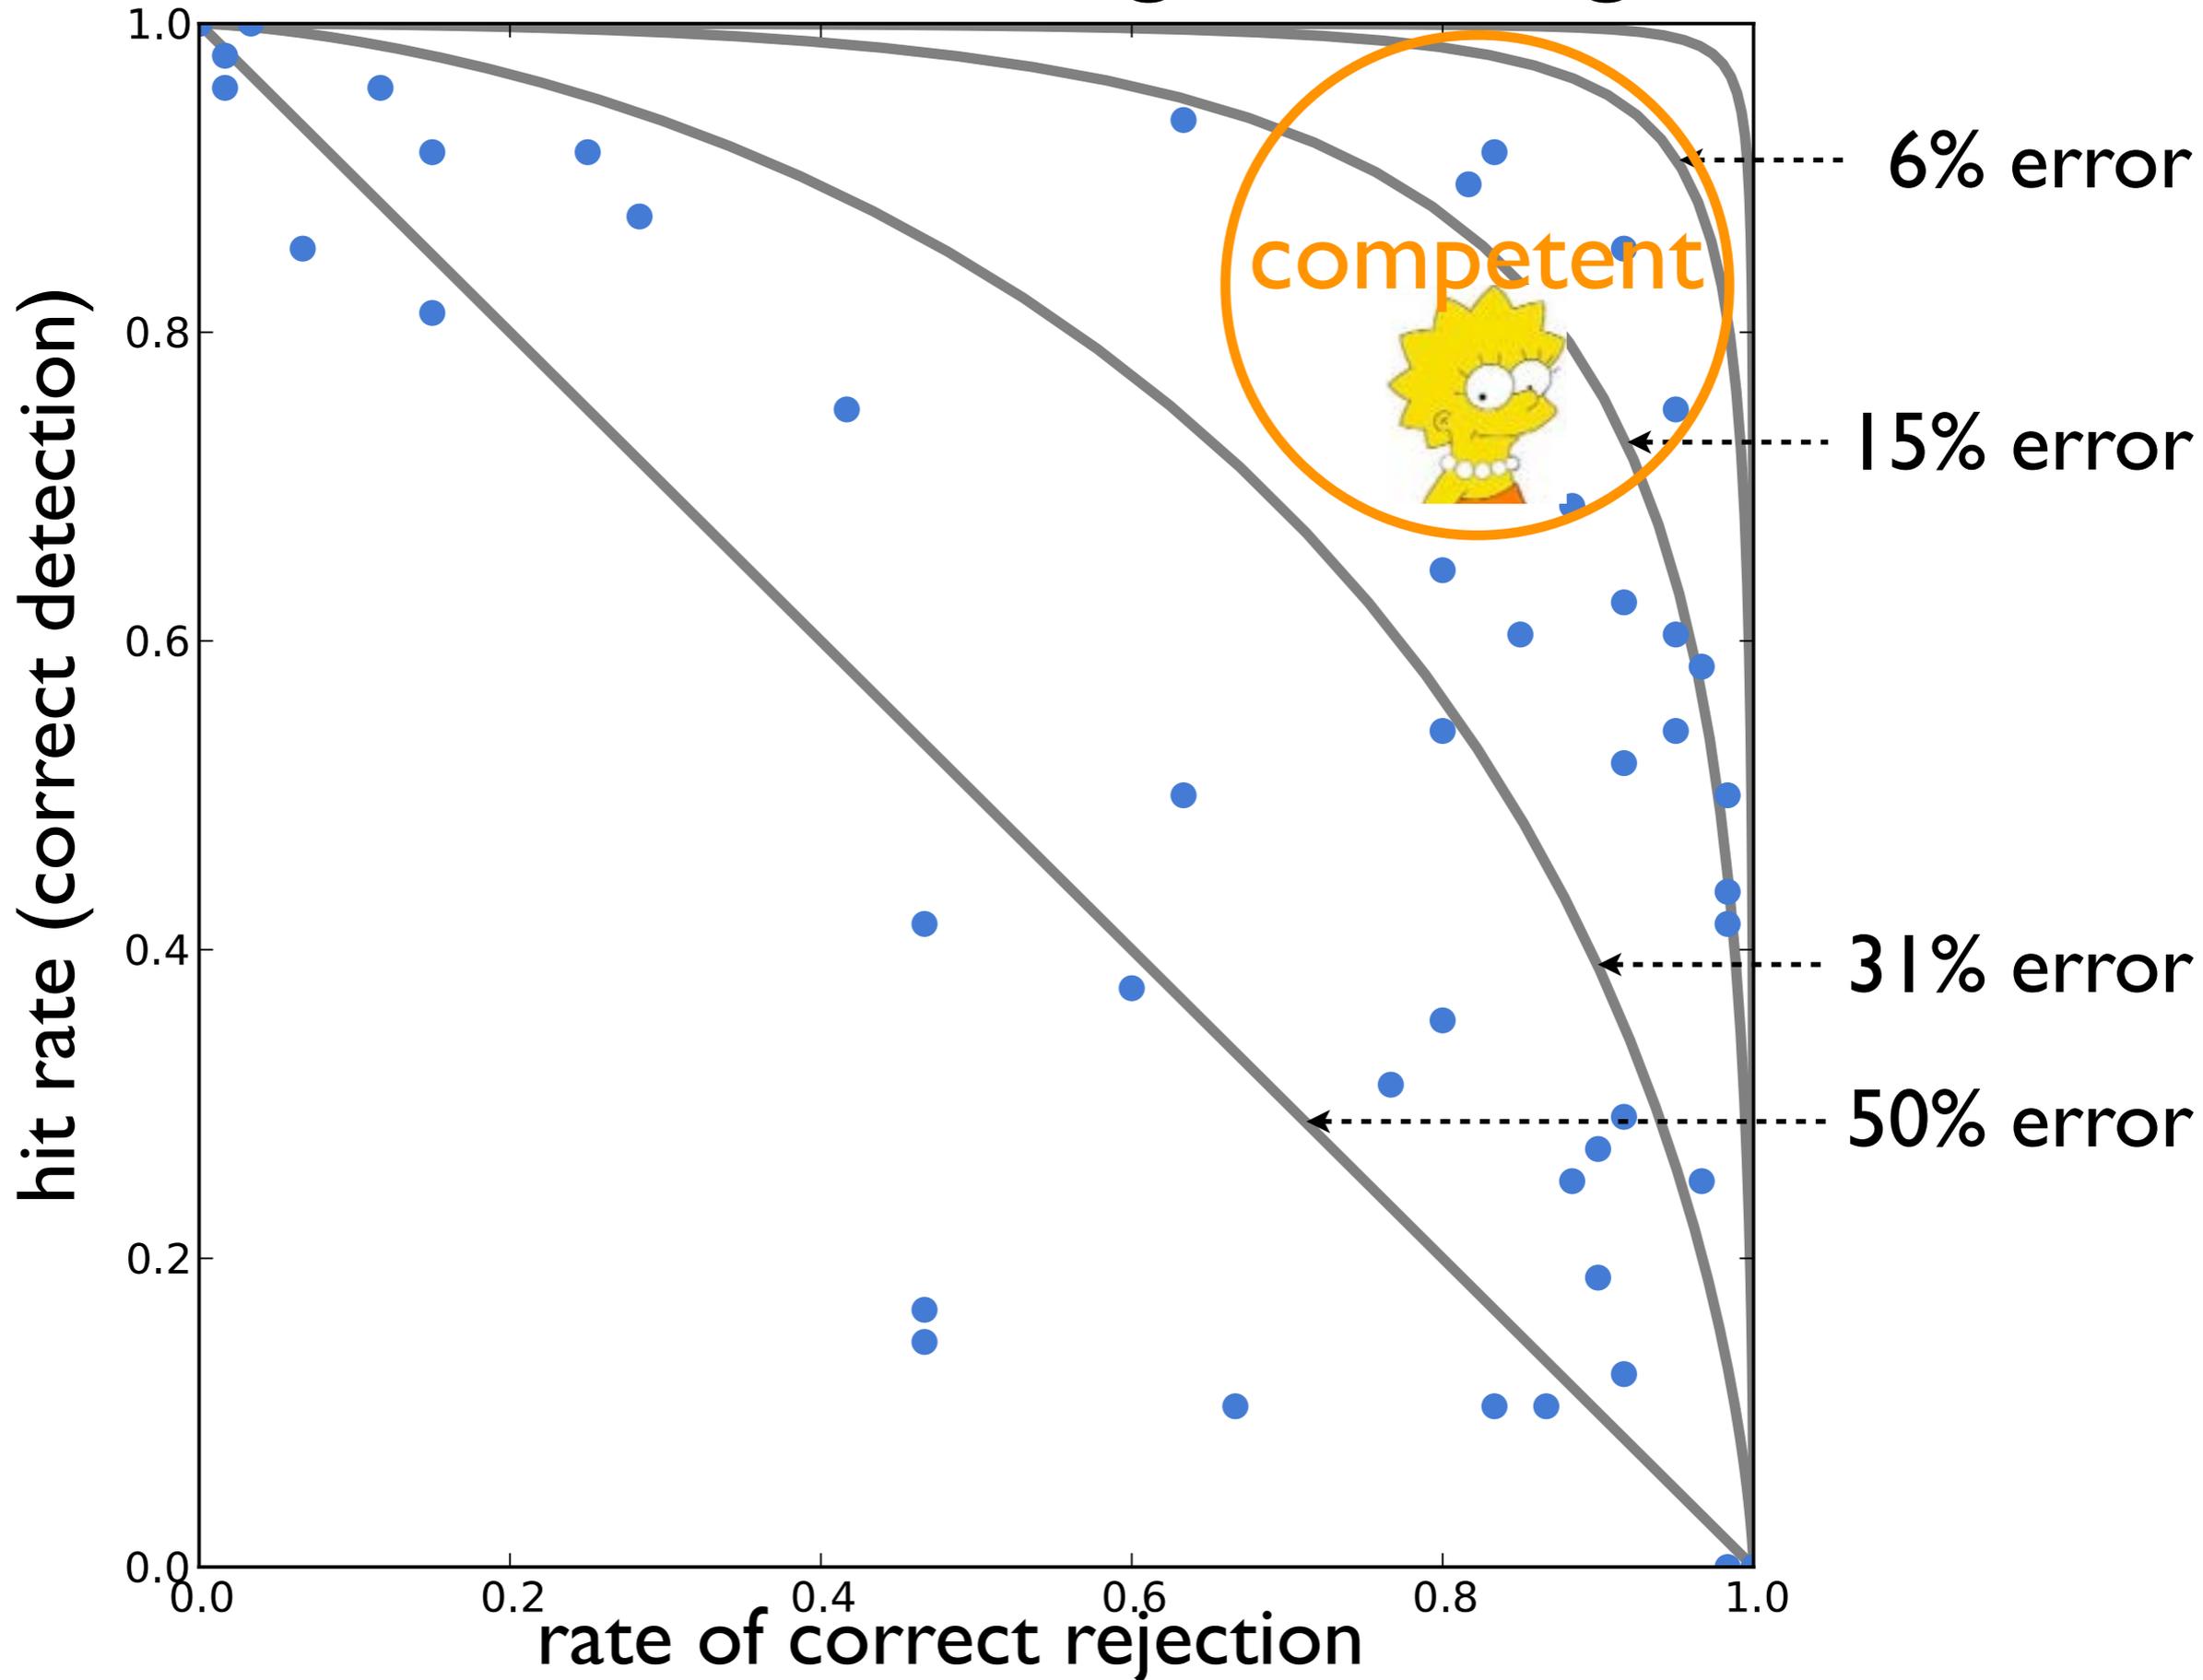
# Task: Find the Indigo Bunting



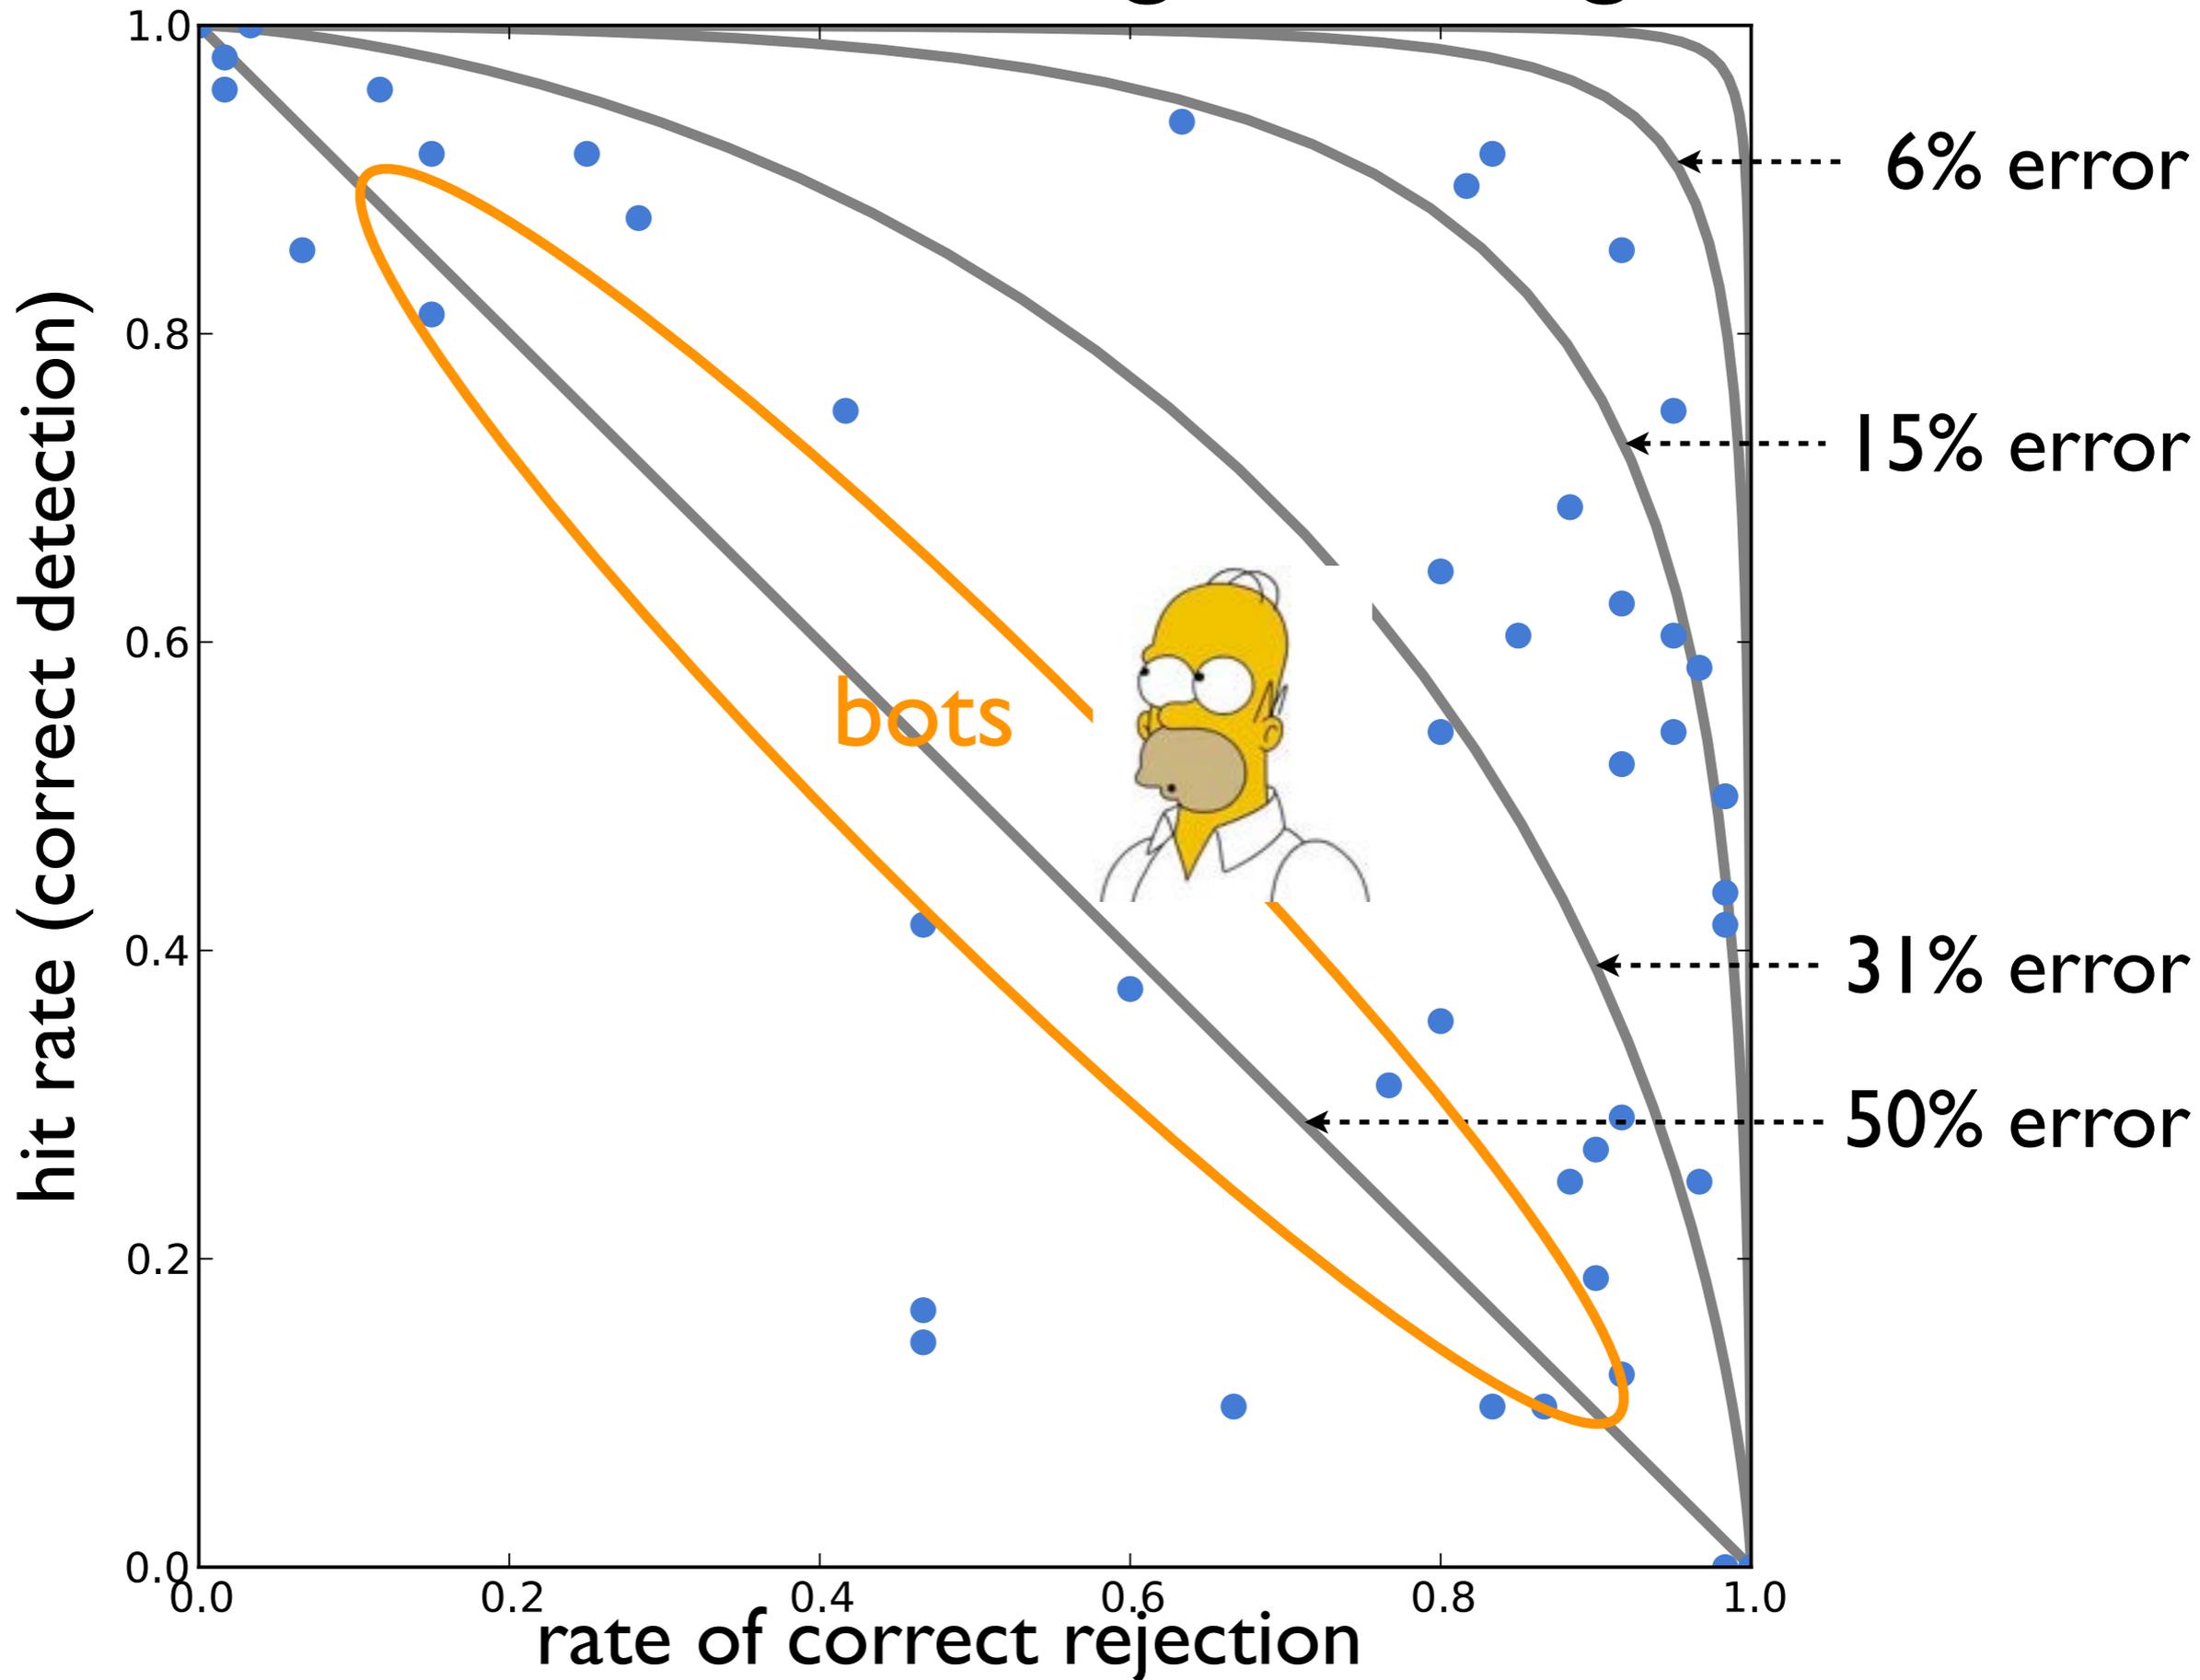
# Task: Find the Indigo Bunting



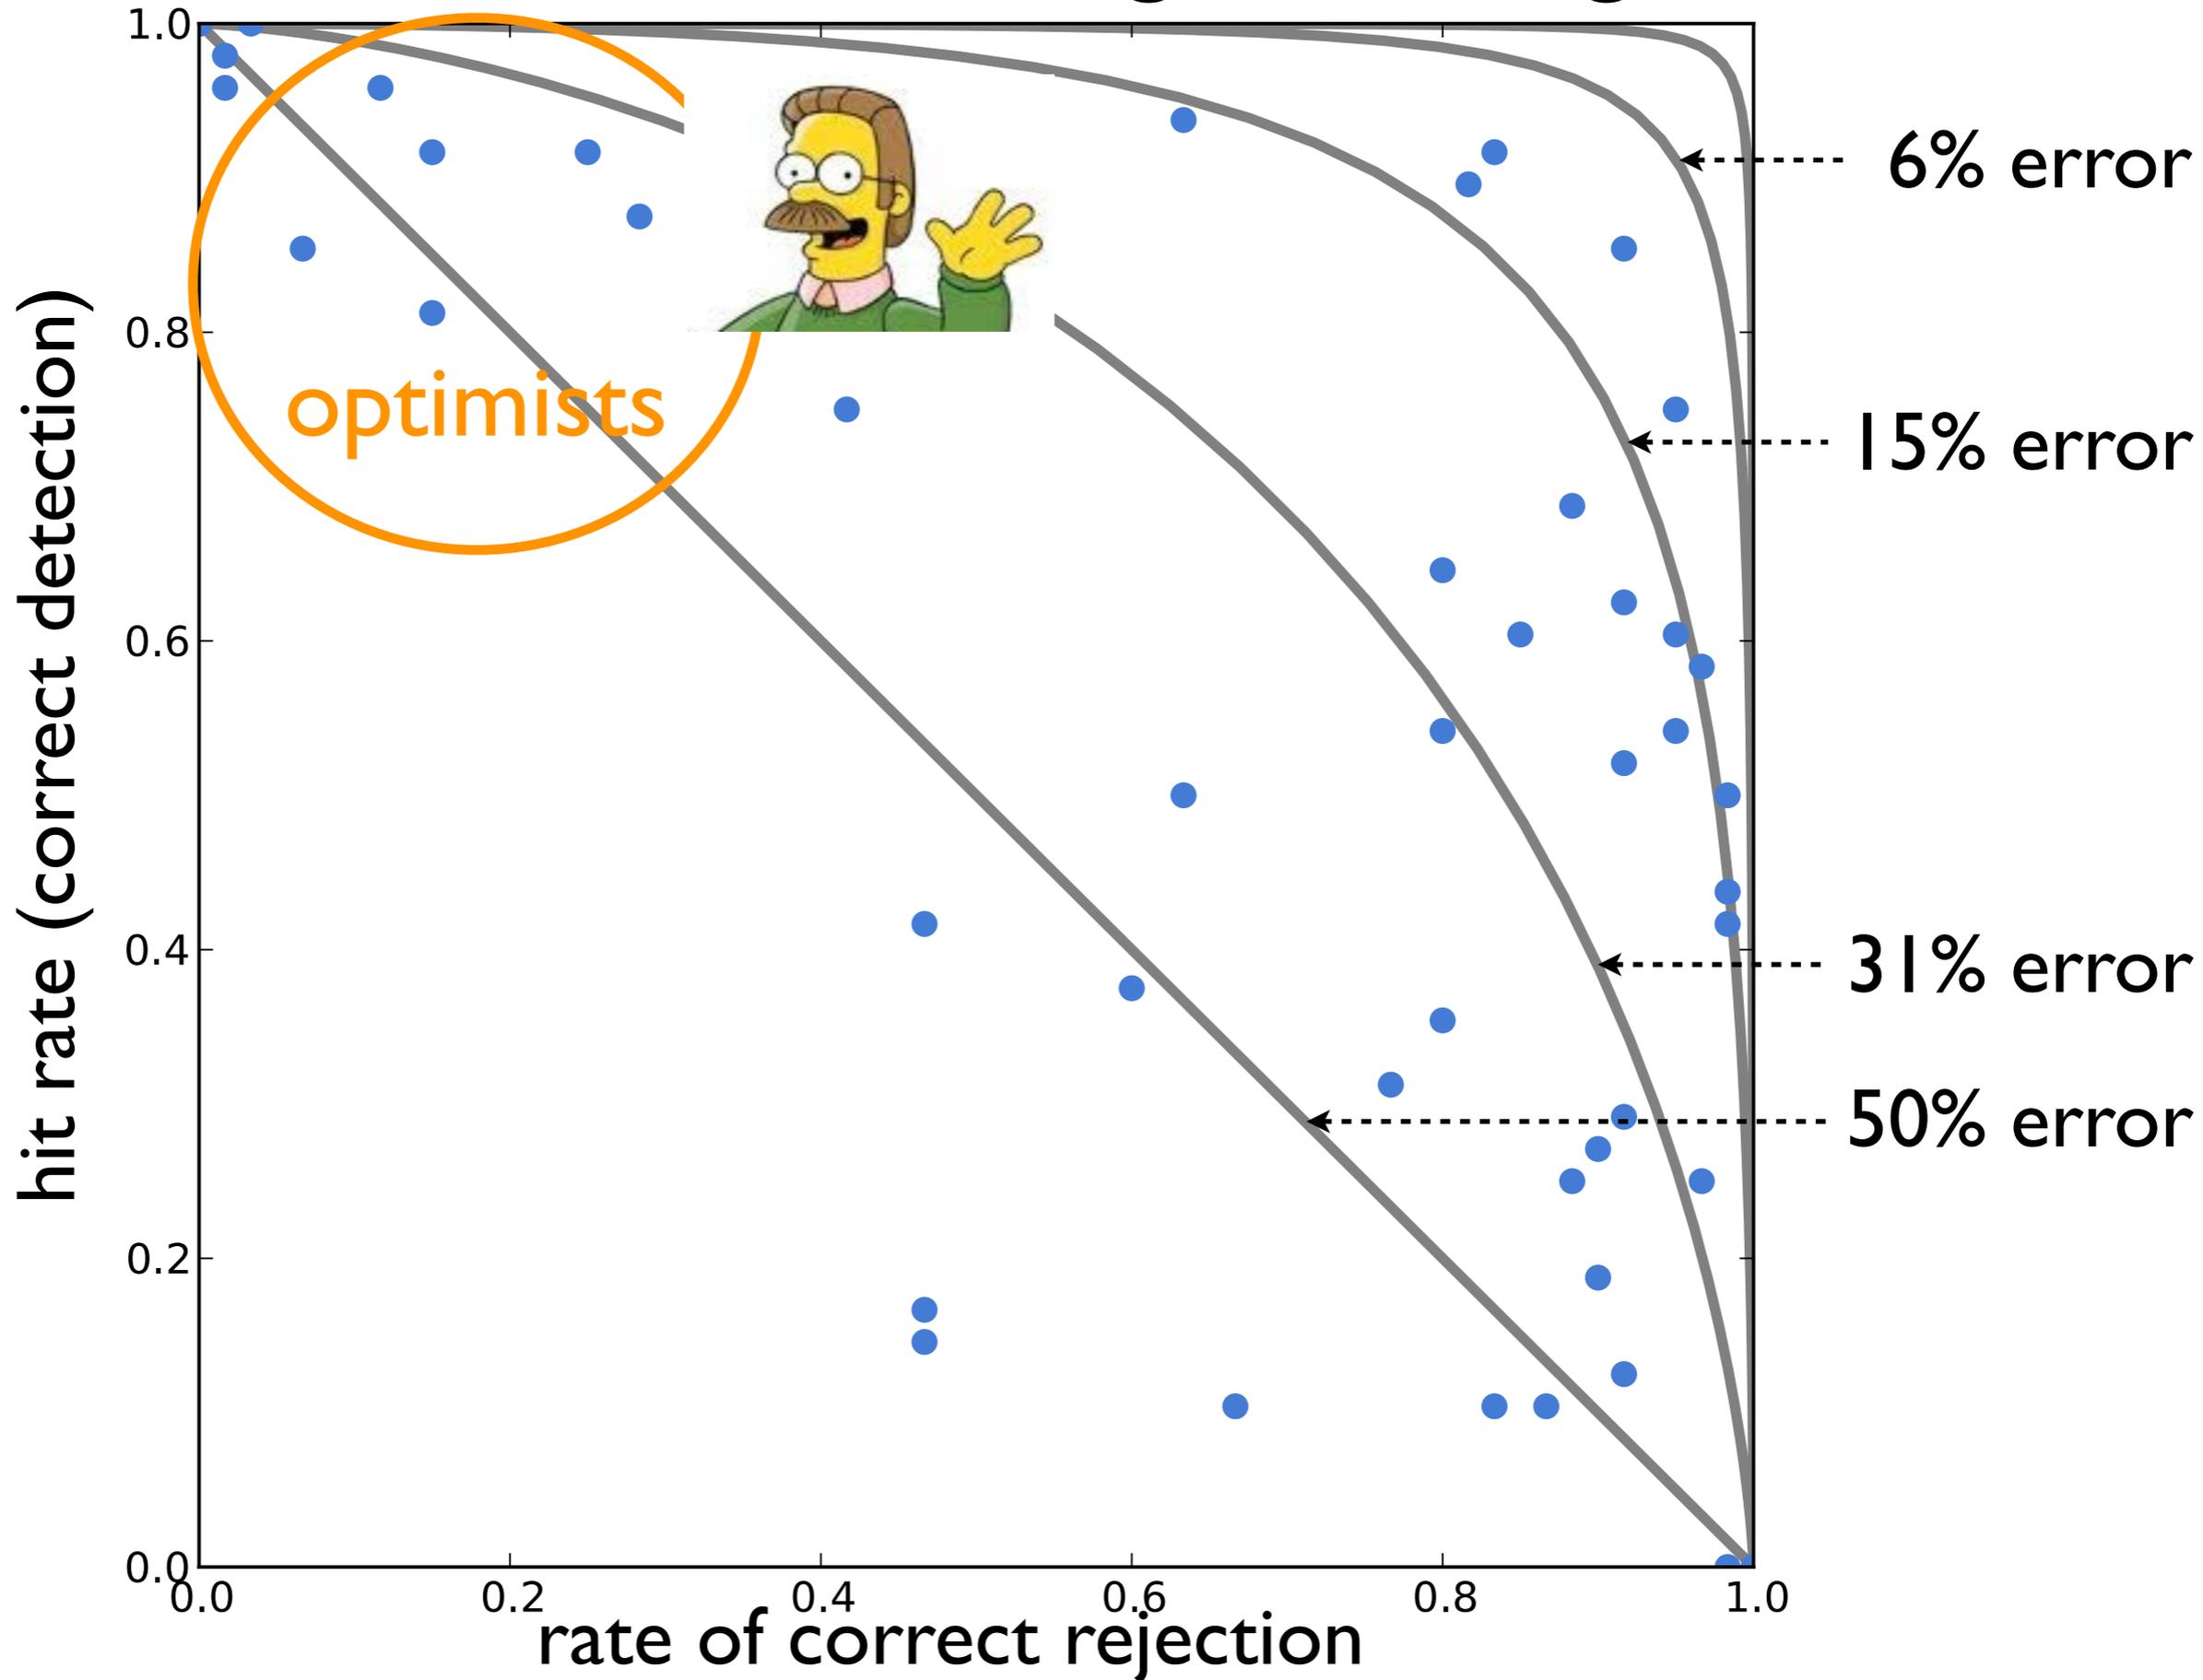
# Task: Find the Indigo Bunting



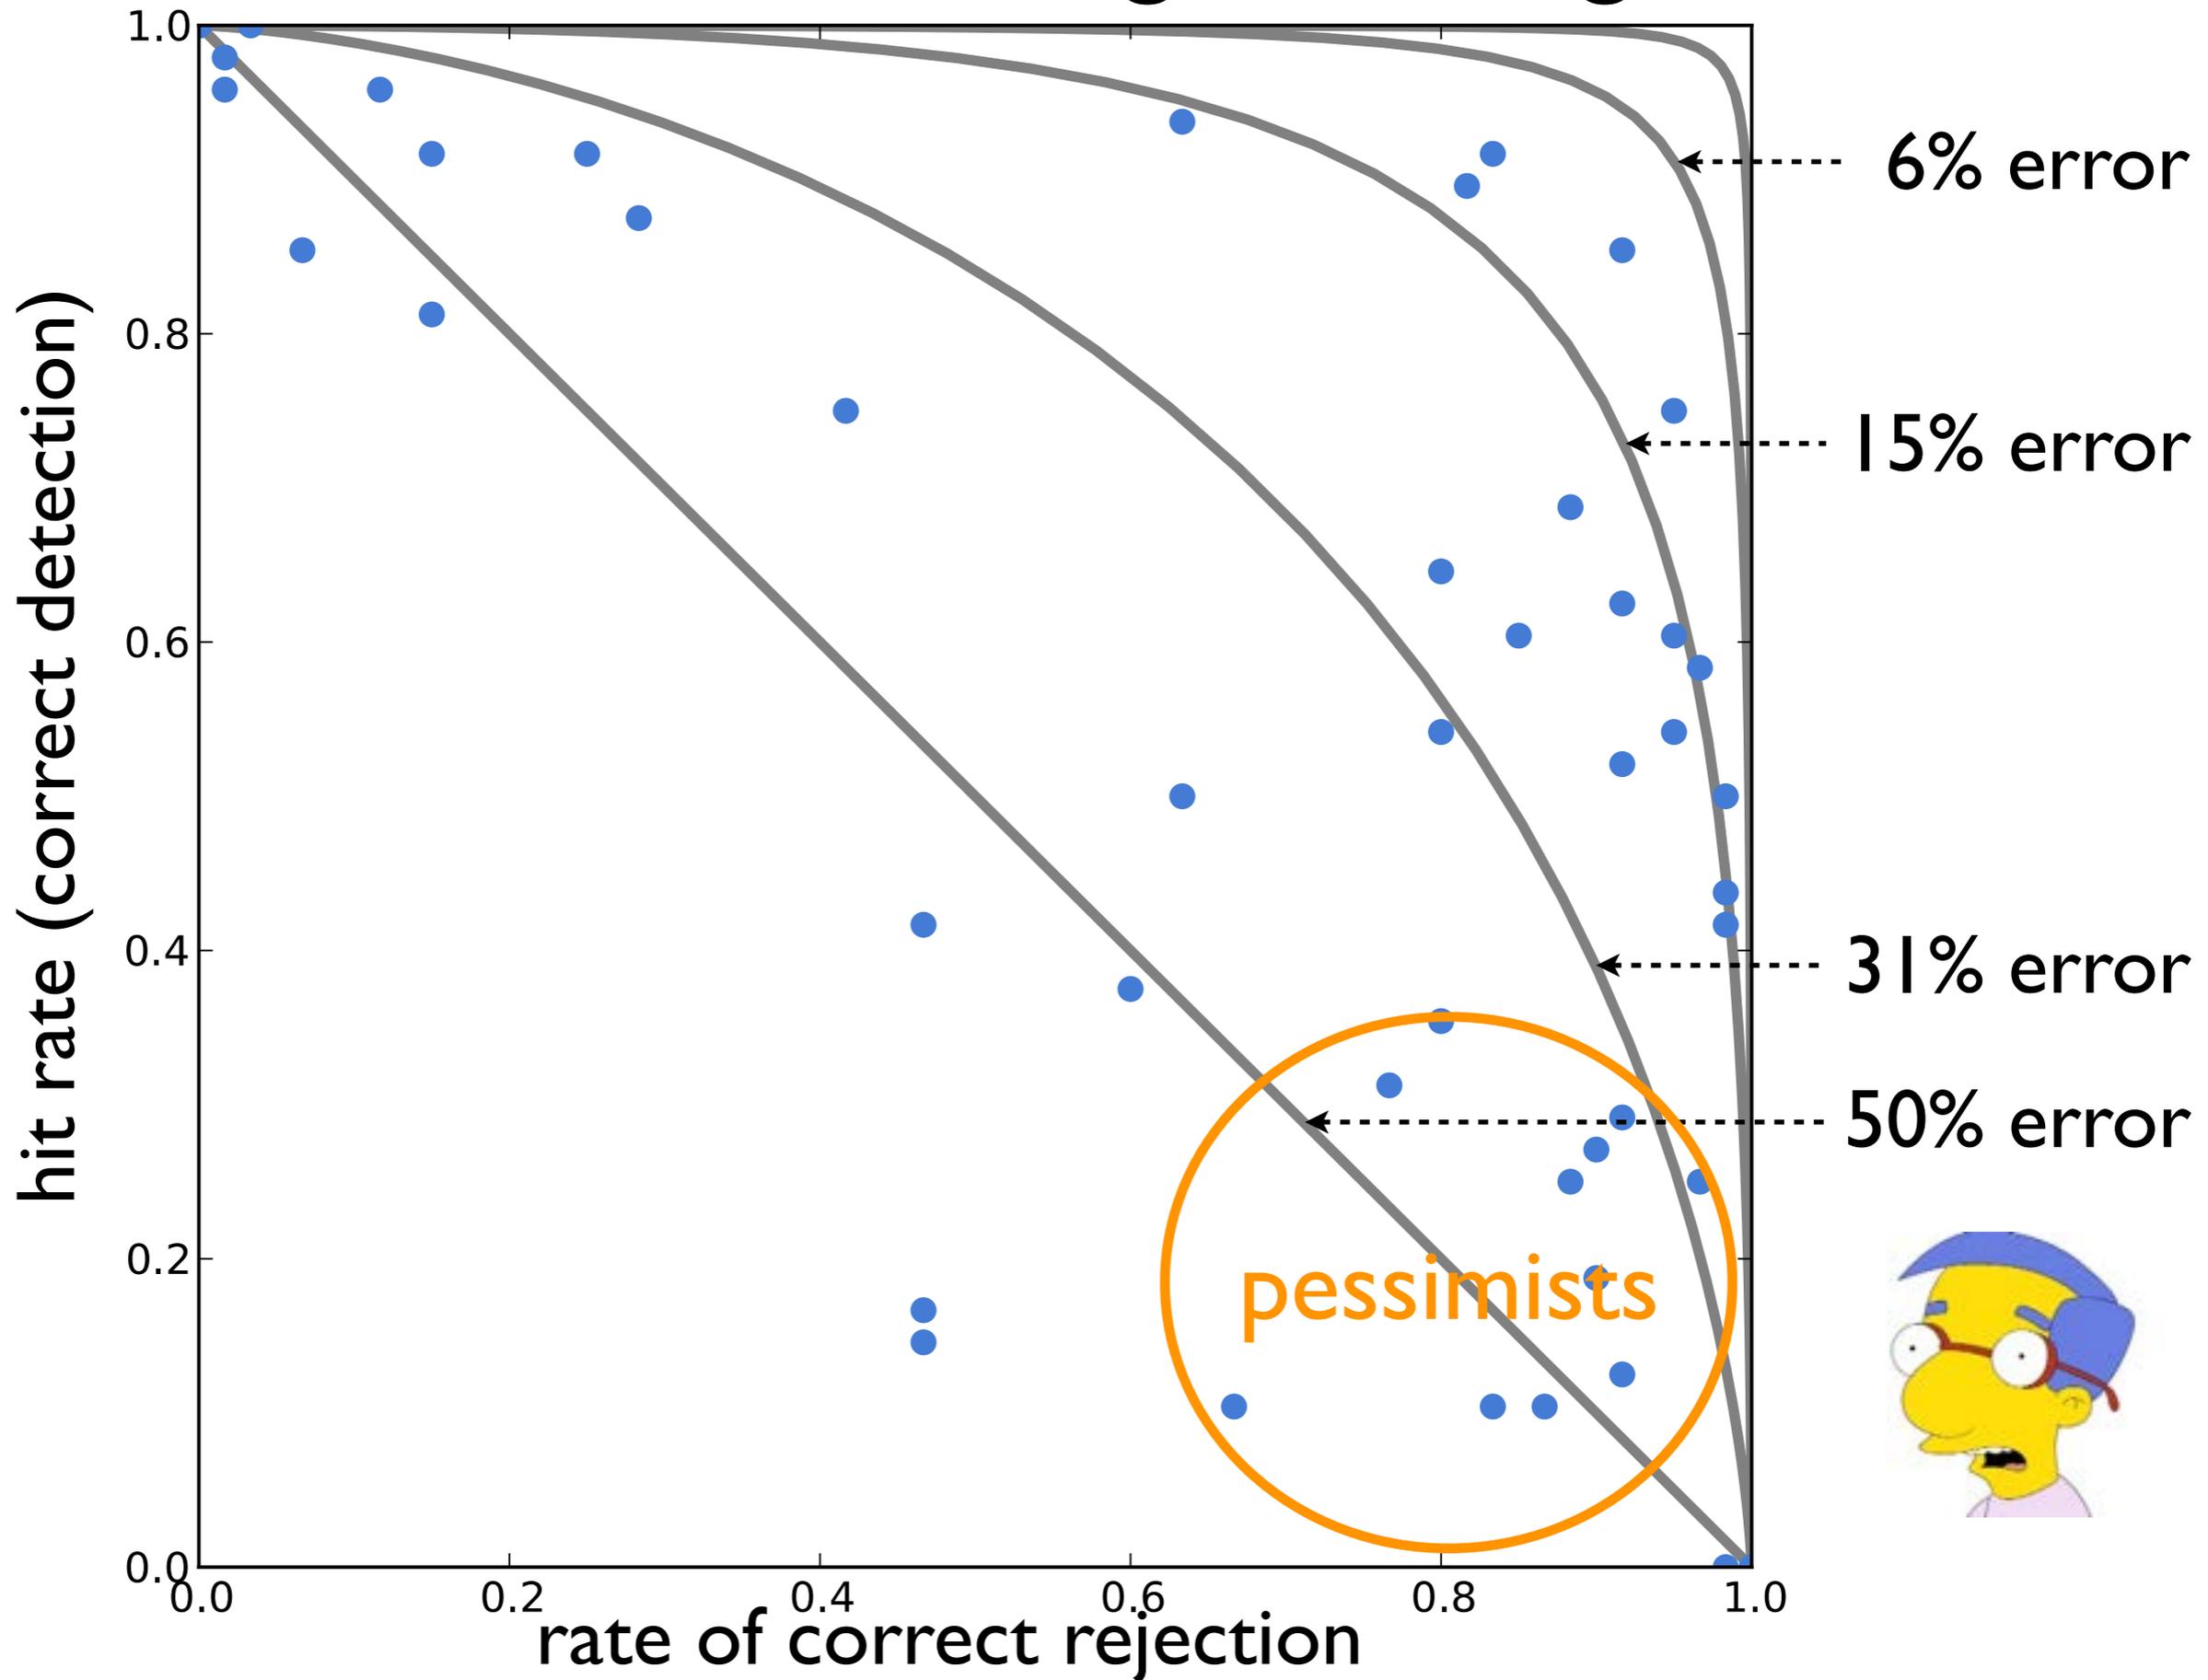
# Task: Find the Indigo Bunting



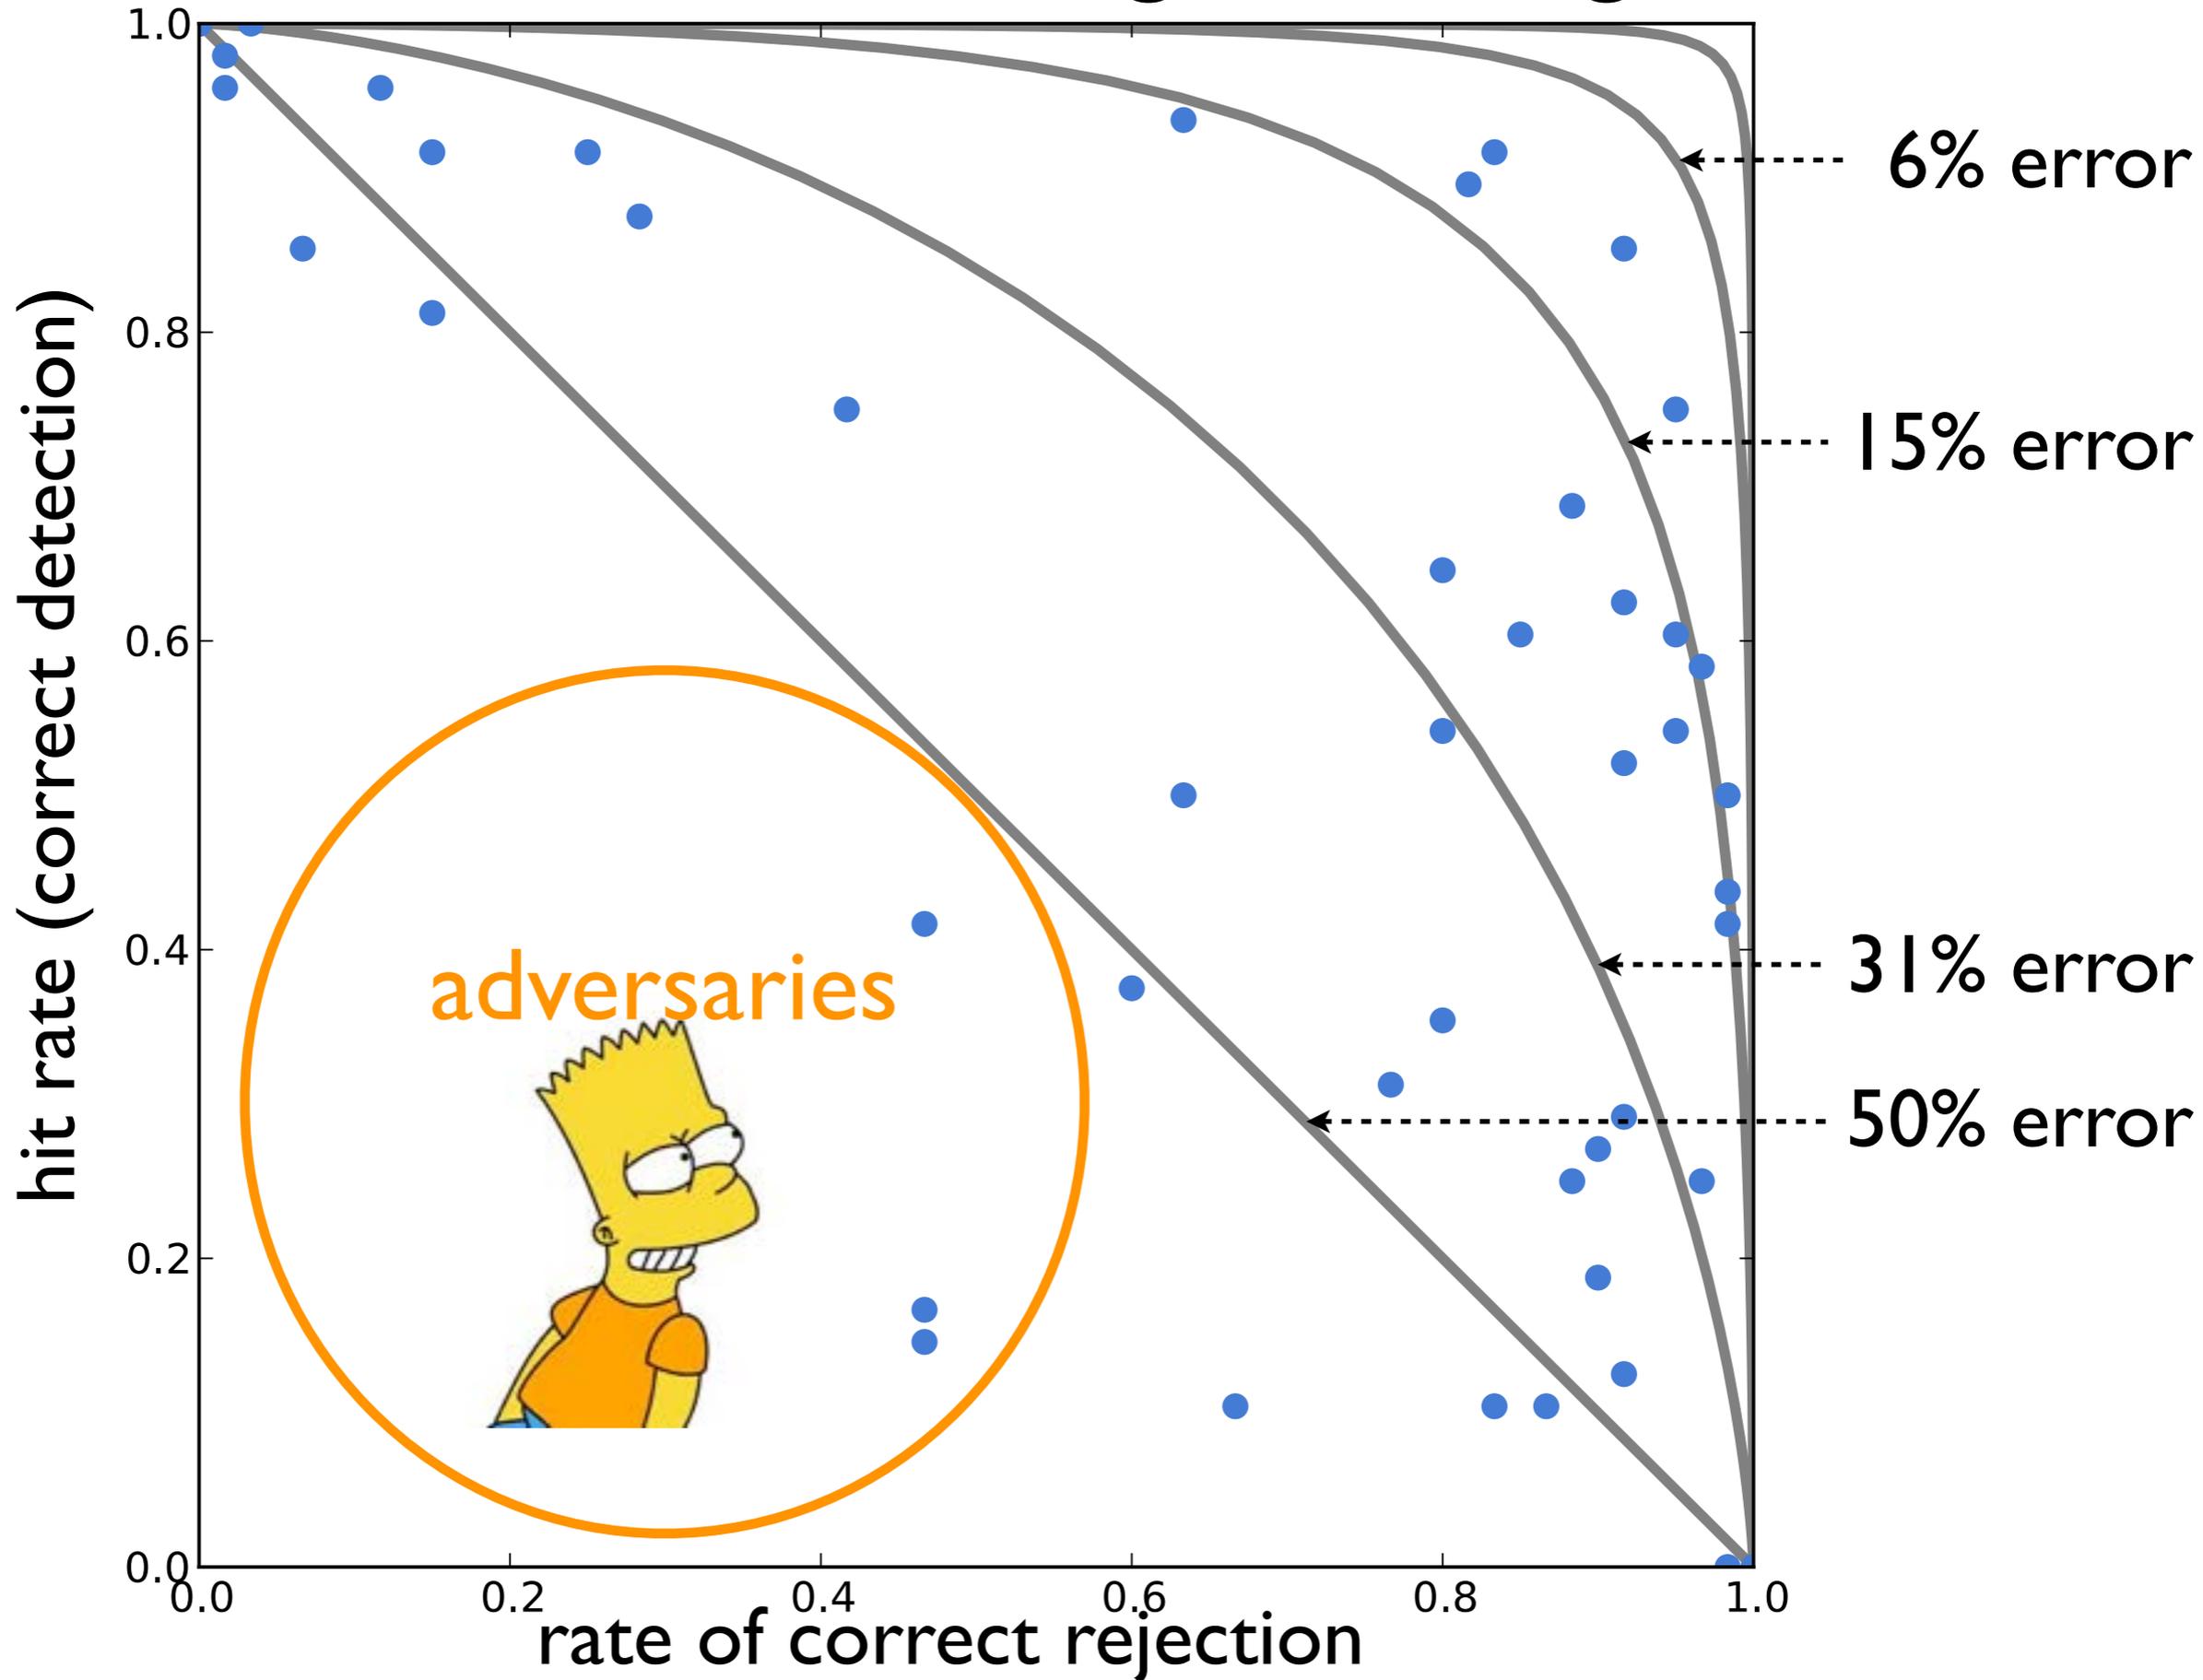
# Task: Find the Indigo Bunting



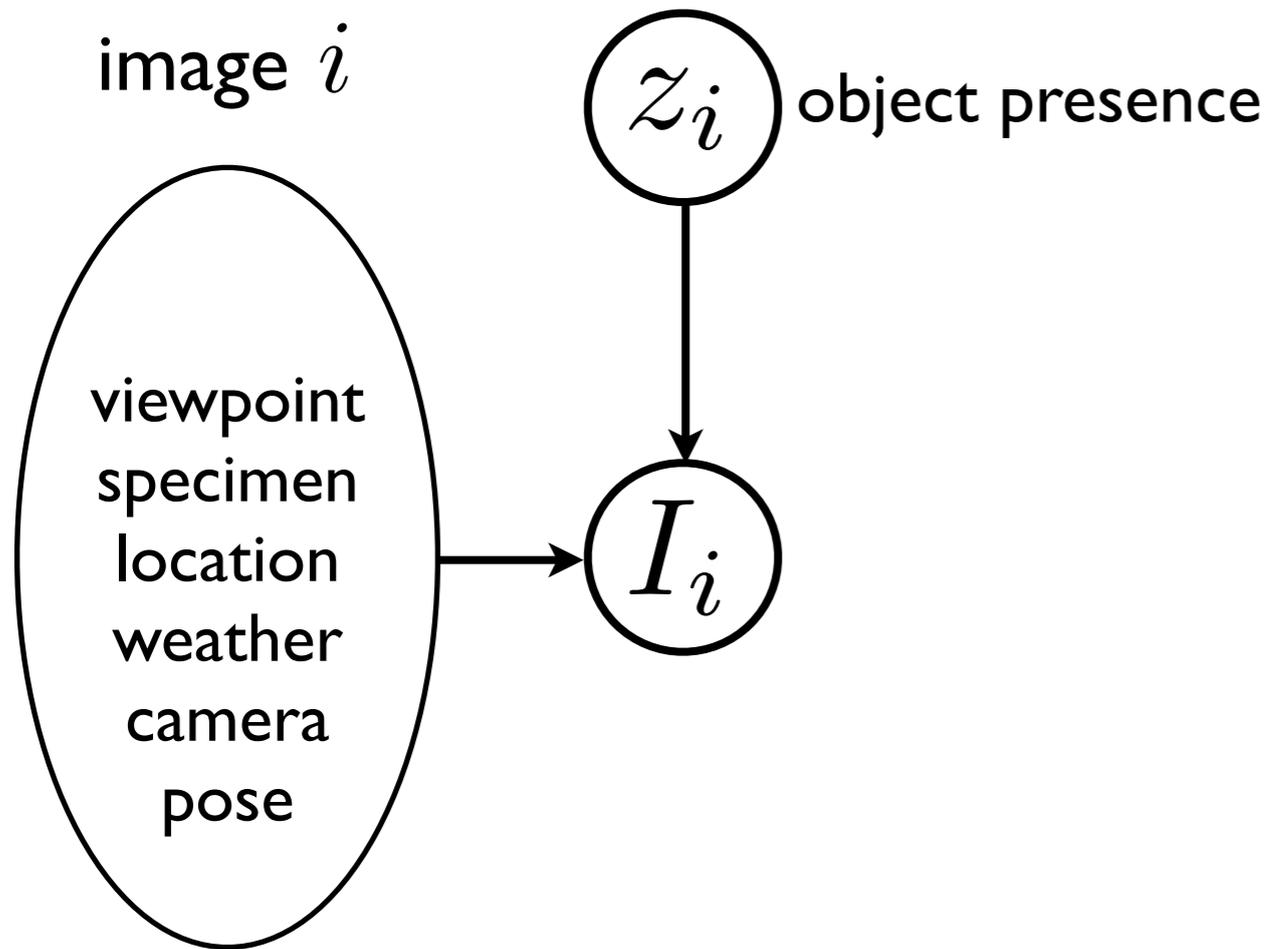
# Task: Find the Indigo Bunting



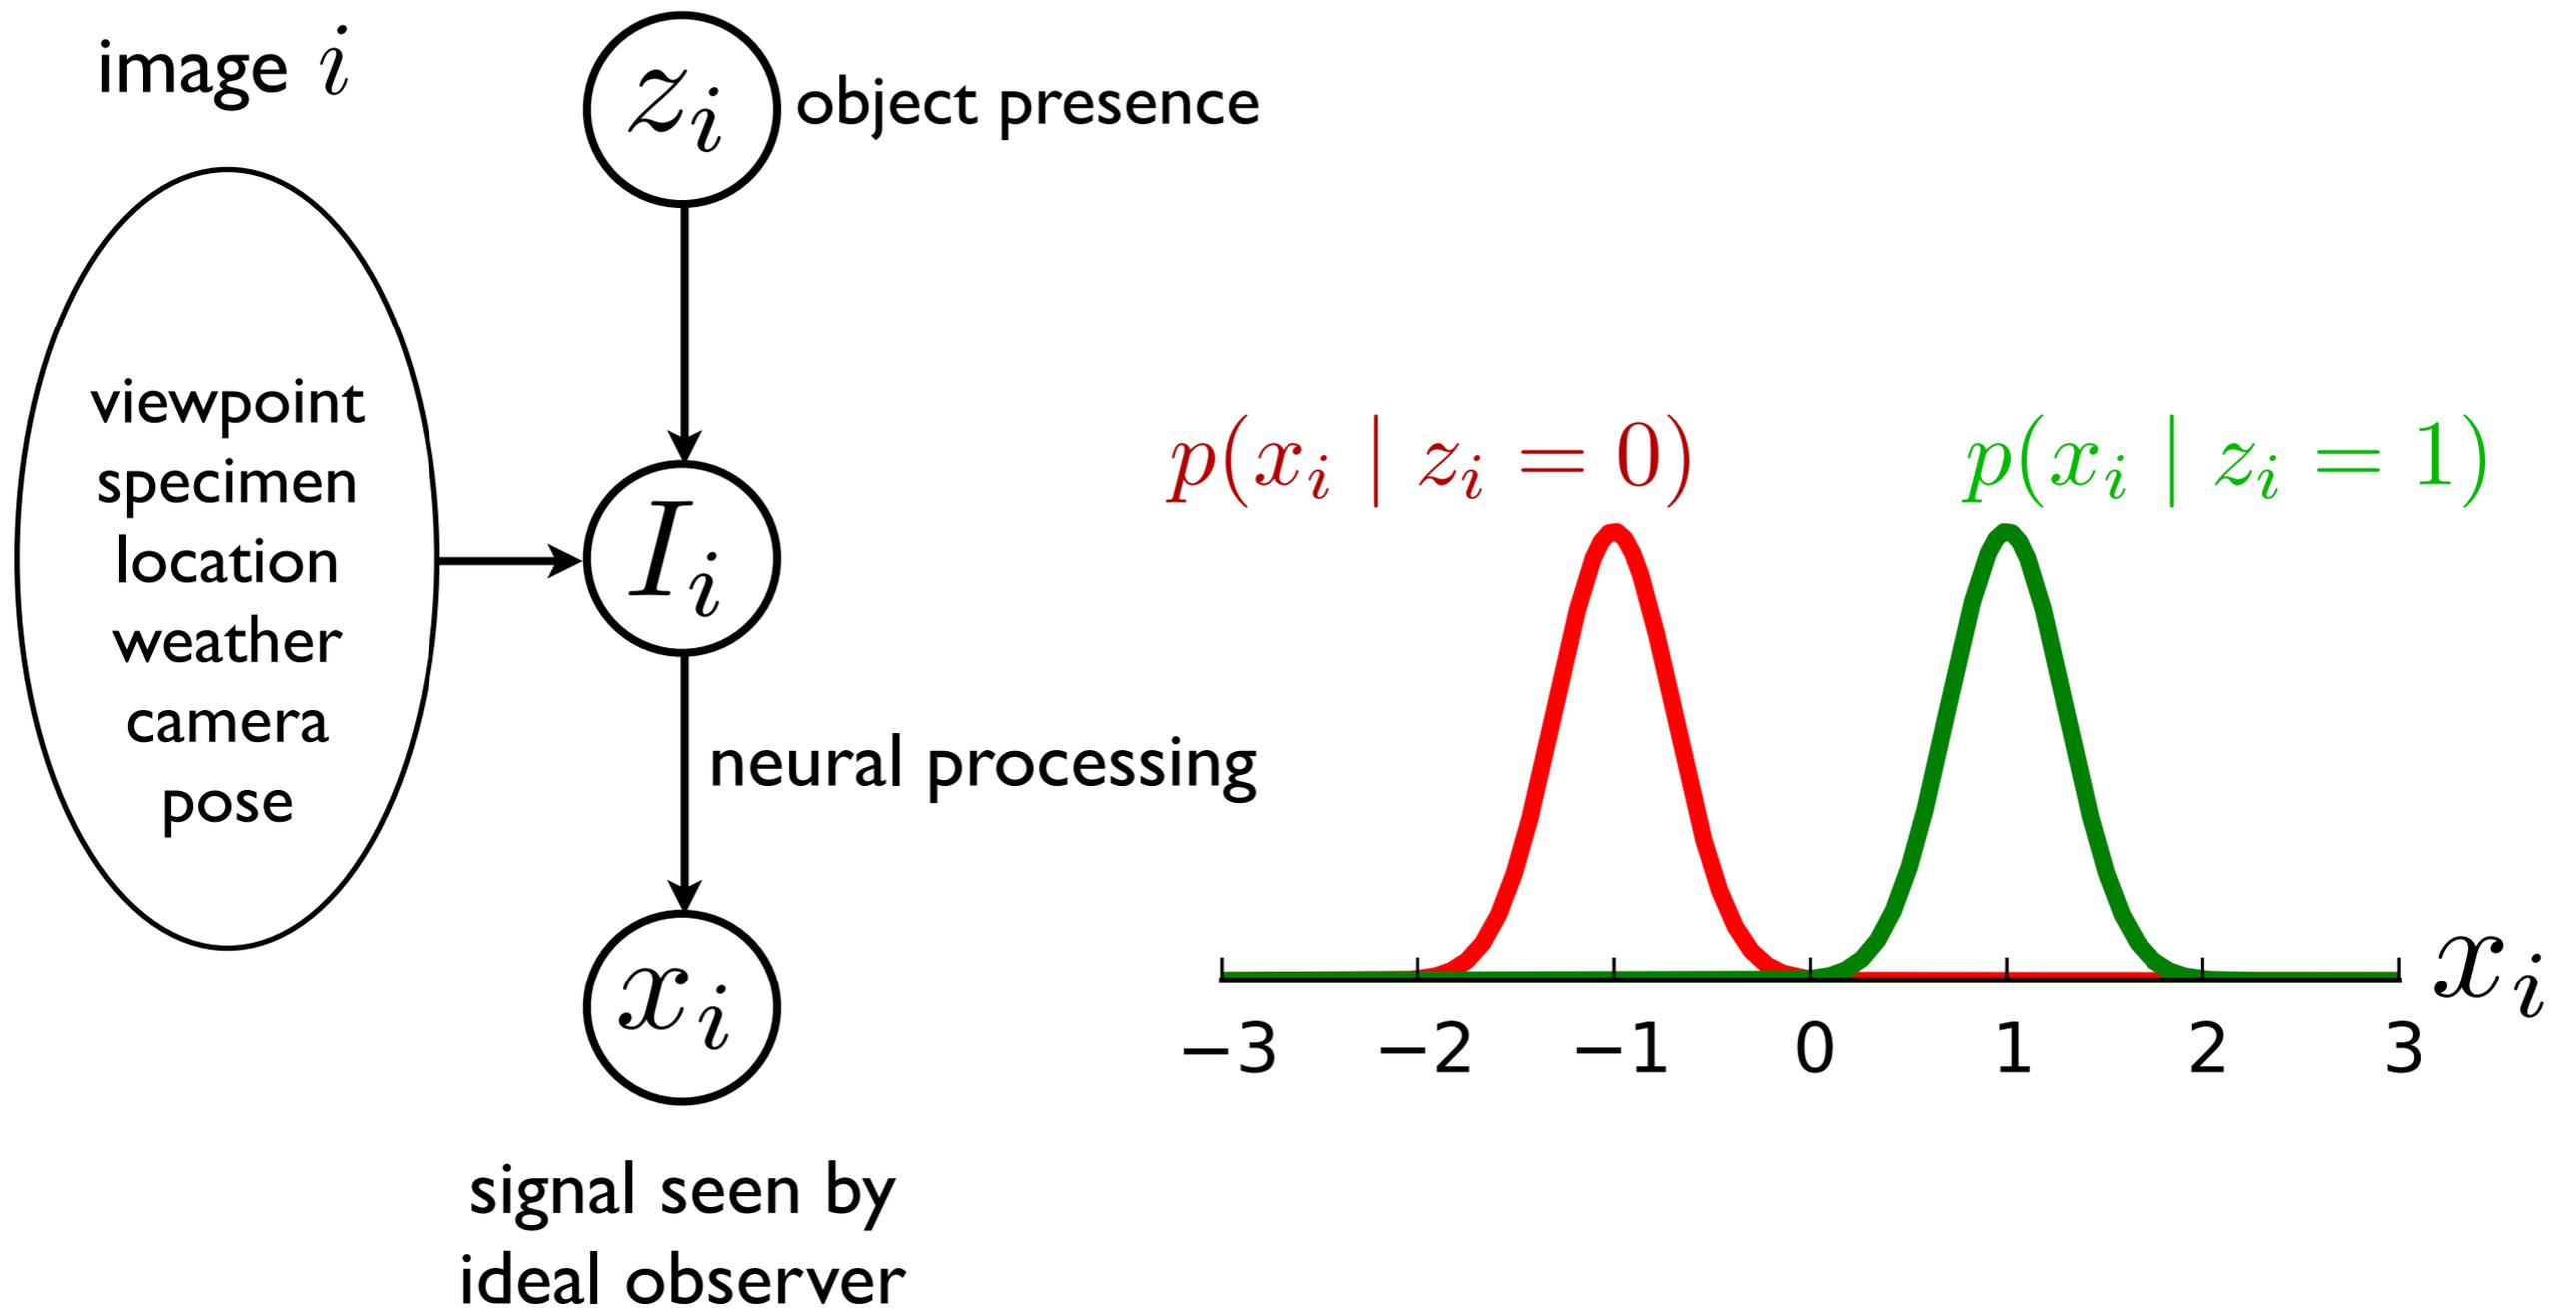
# Task: Find the Indigo Bunting



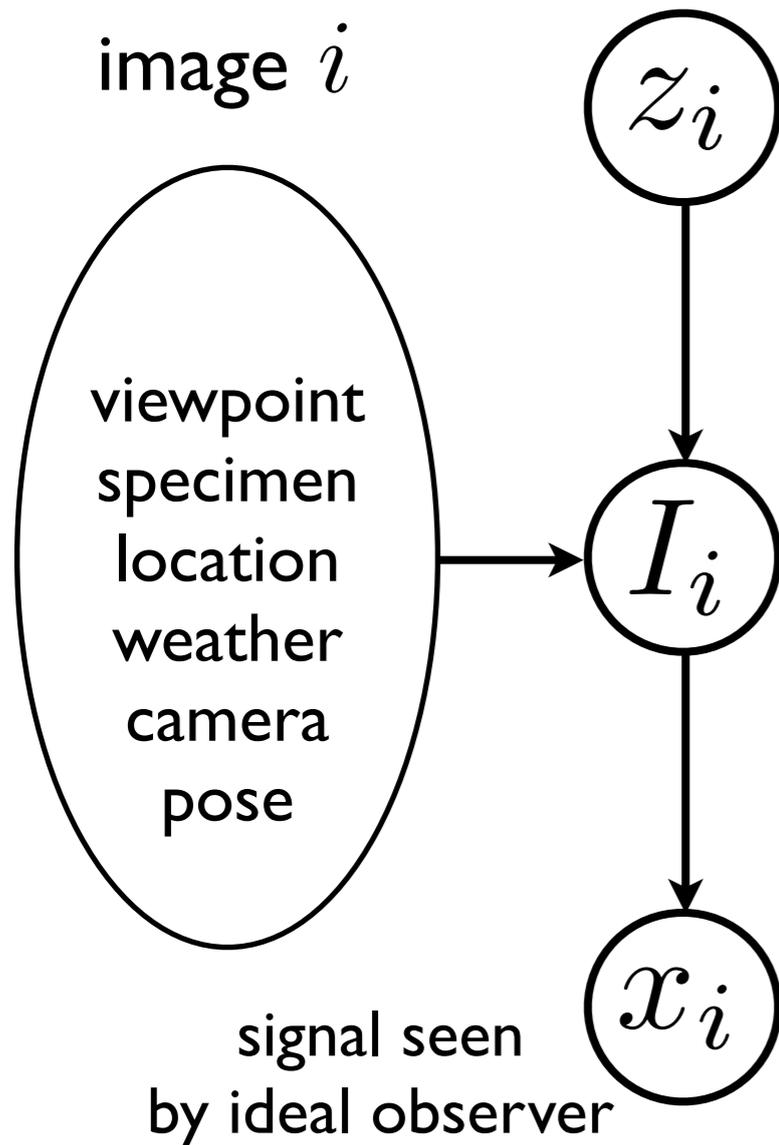
# Image difficulty and annotator competence



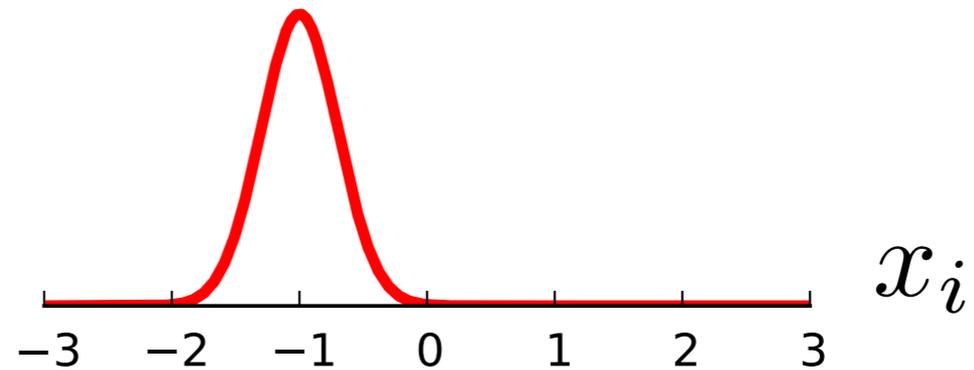
# Image difficulty and annotator competence



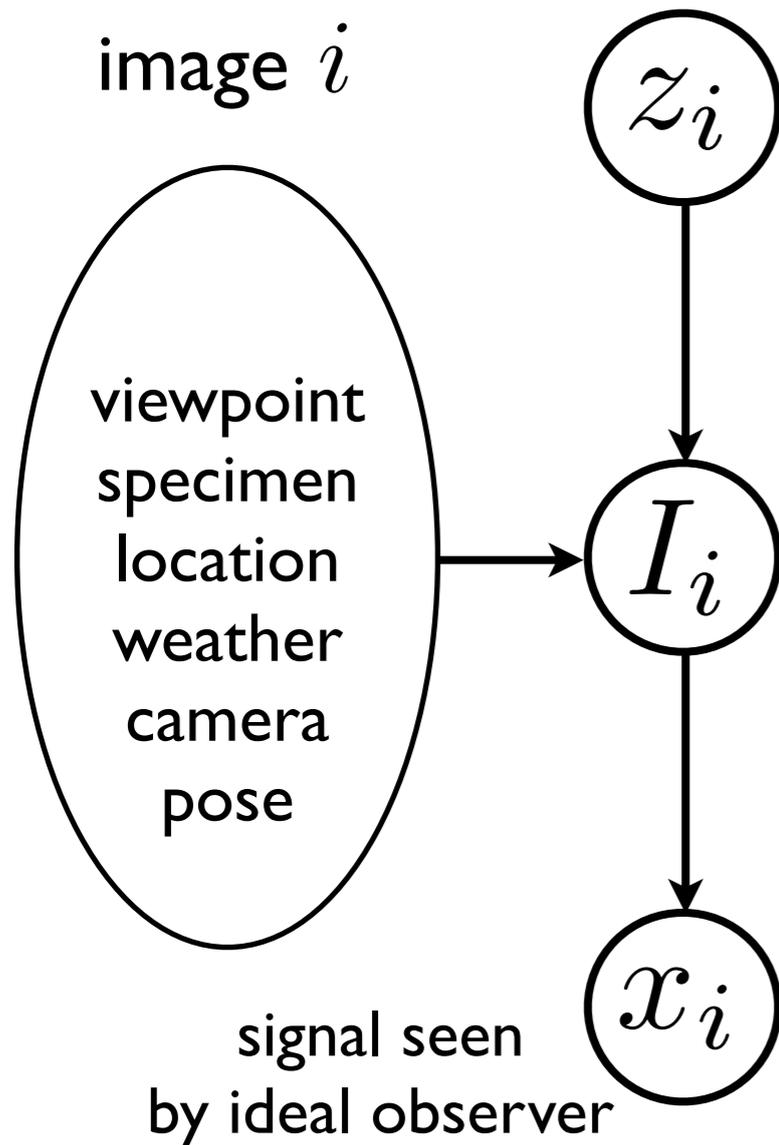
# Image difficulty and annotator competence



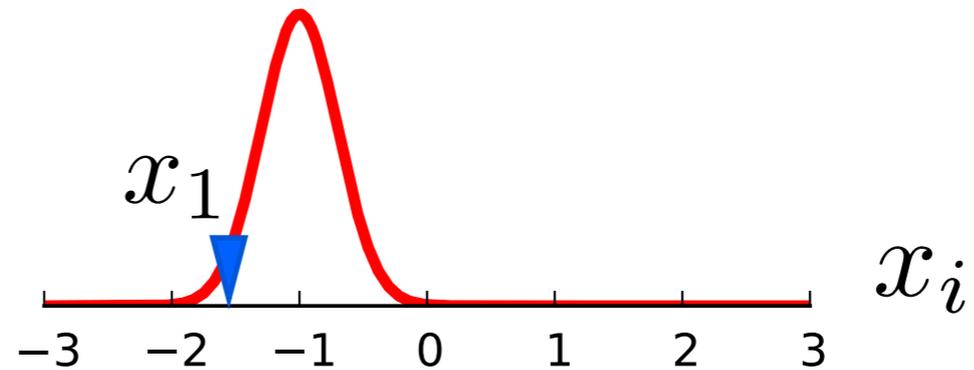
$$p(x_i | z_i = 0)$$



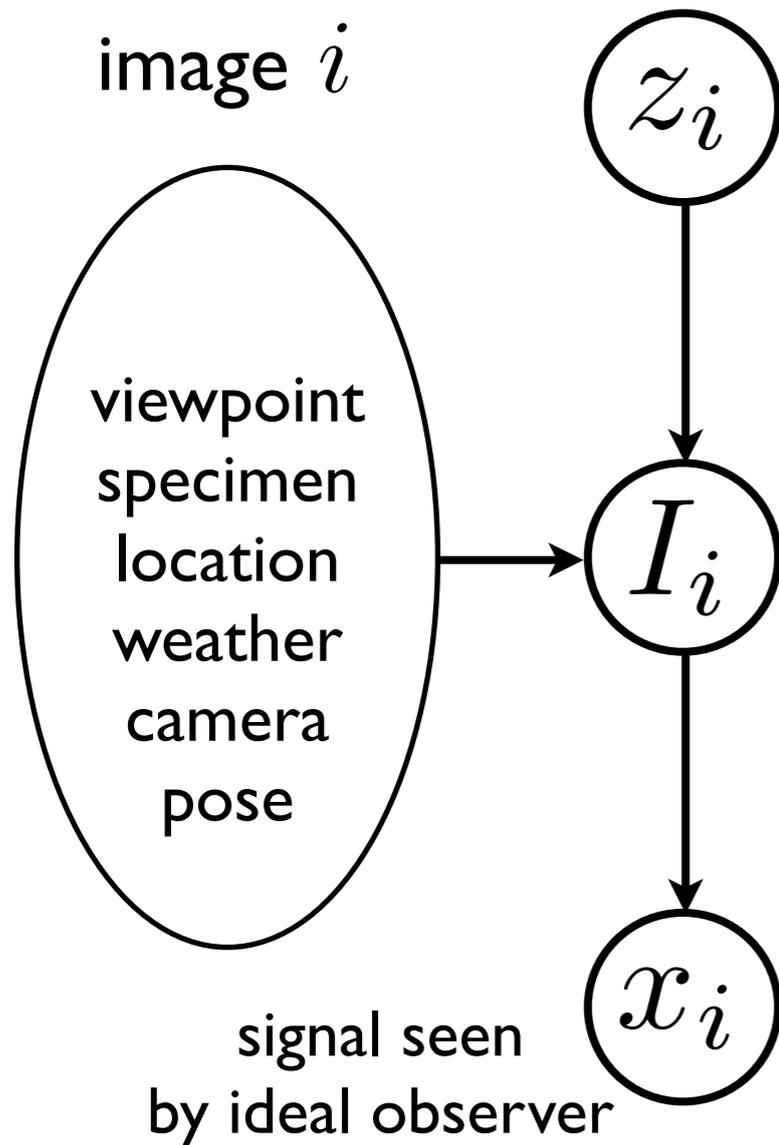
# Image difficulty and annotator competence



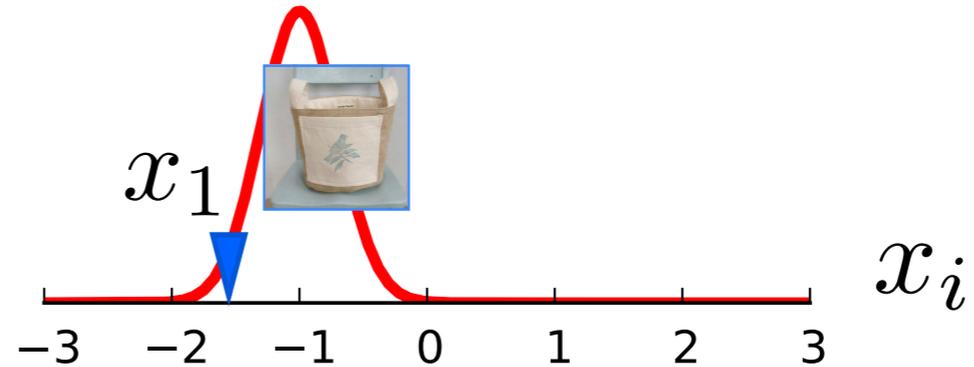
$$p(x_i | z_i = 0)$$



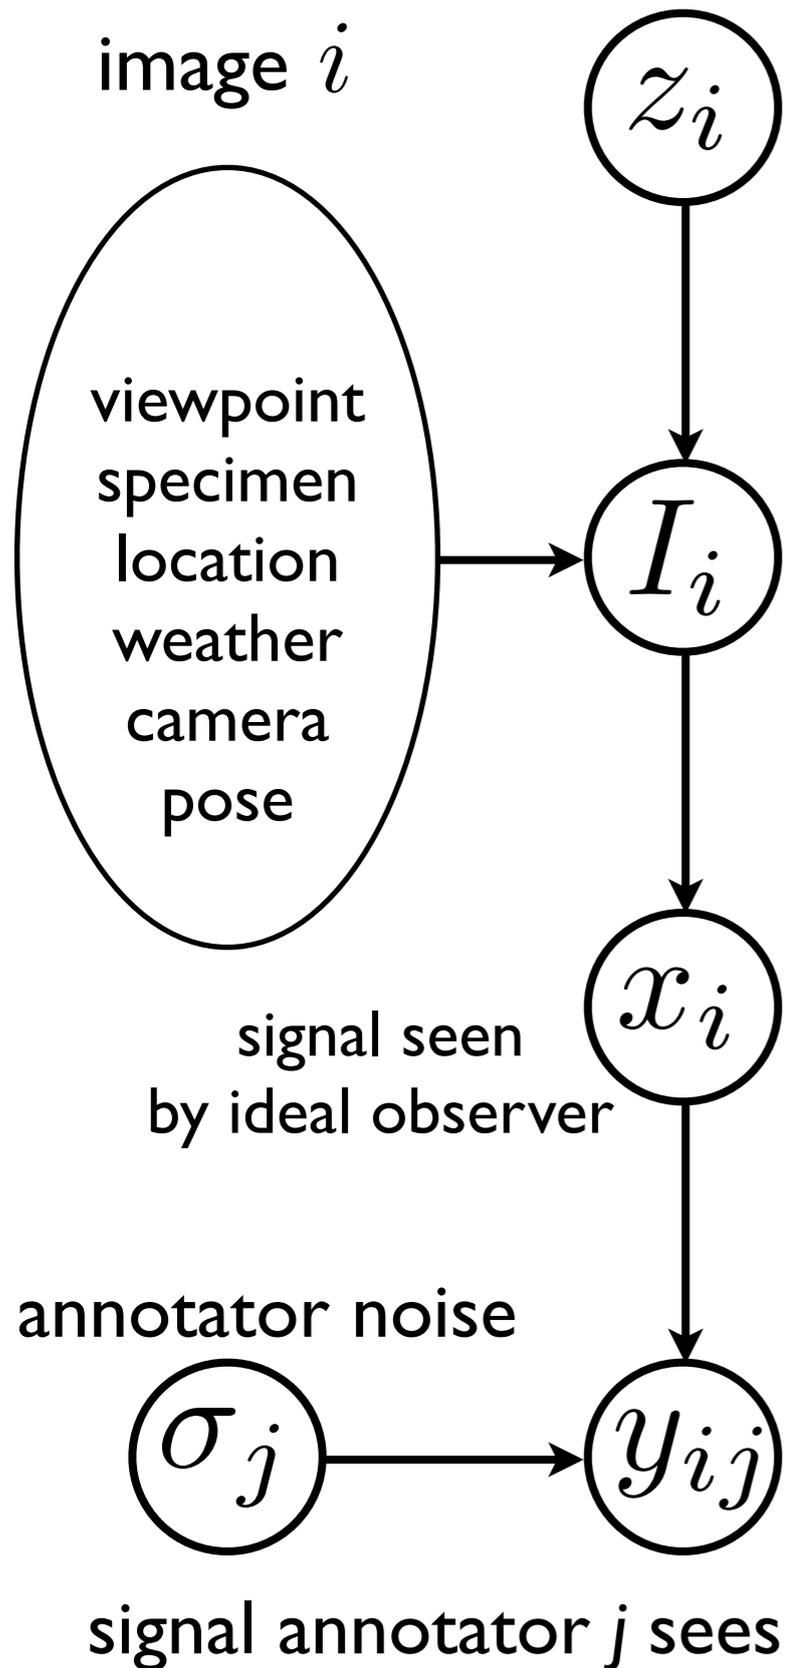
# Image difficulty and annotator competence



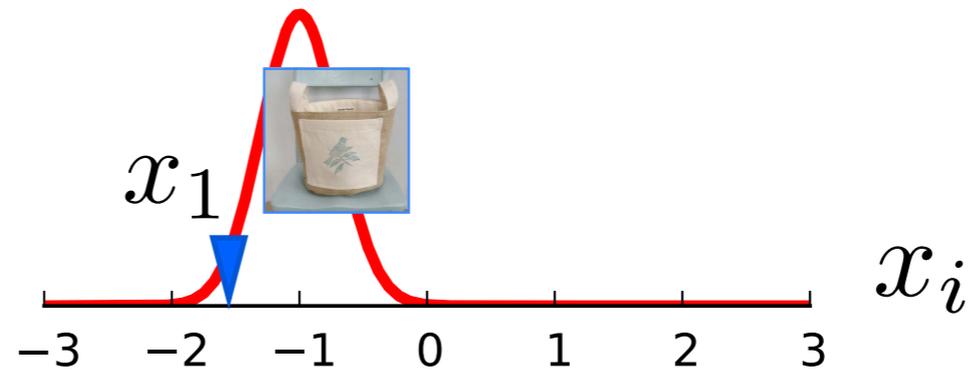
$$p(x_i | z_i = 0)$$



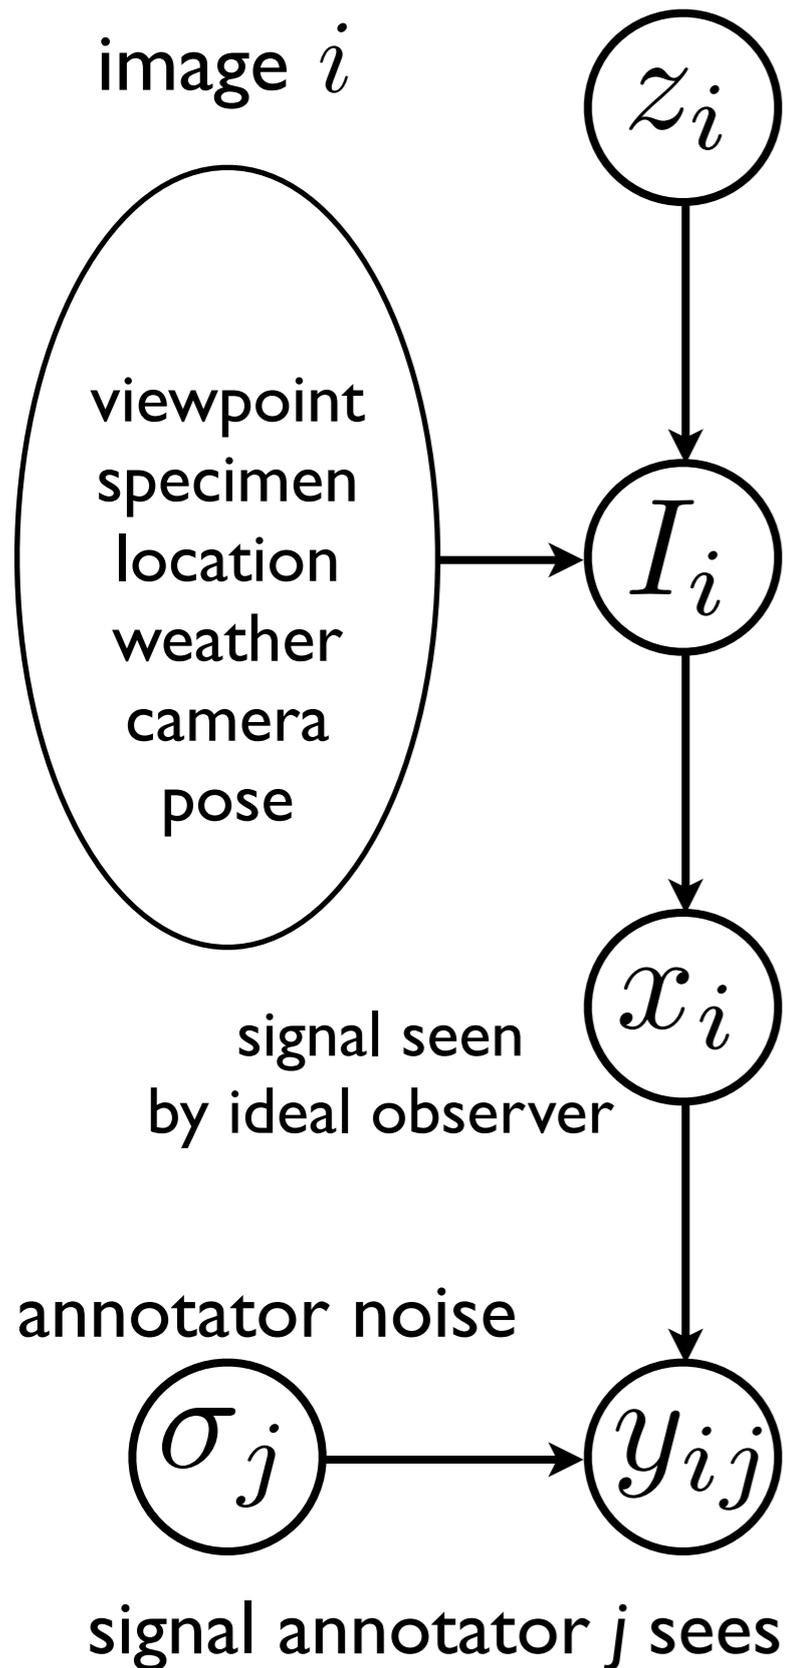
# Image difficulty and annotator competence



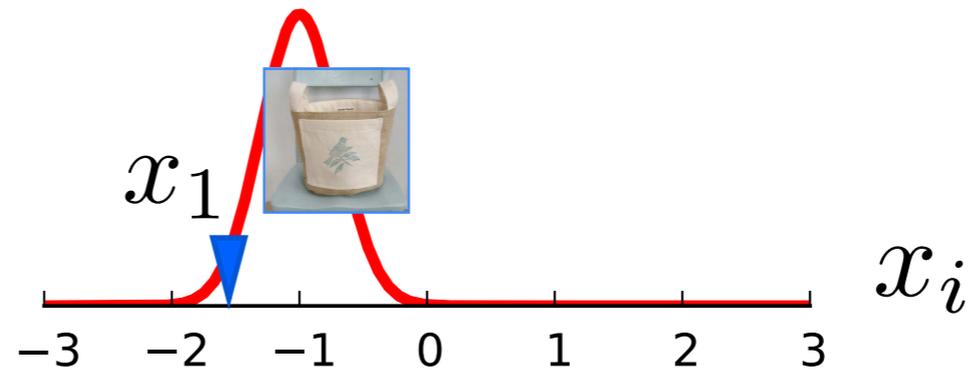
$$p(x_i | z_i = 0)$$



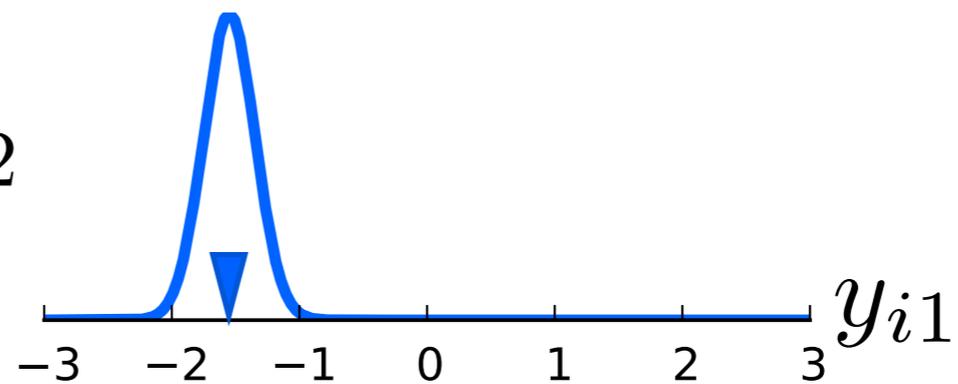
# Image difficulty and annotator competence



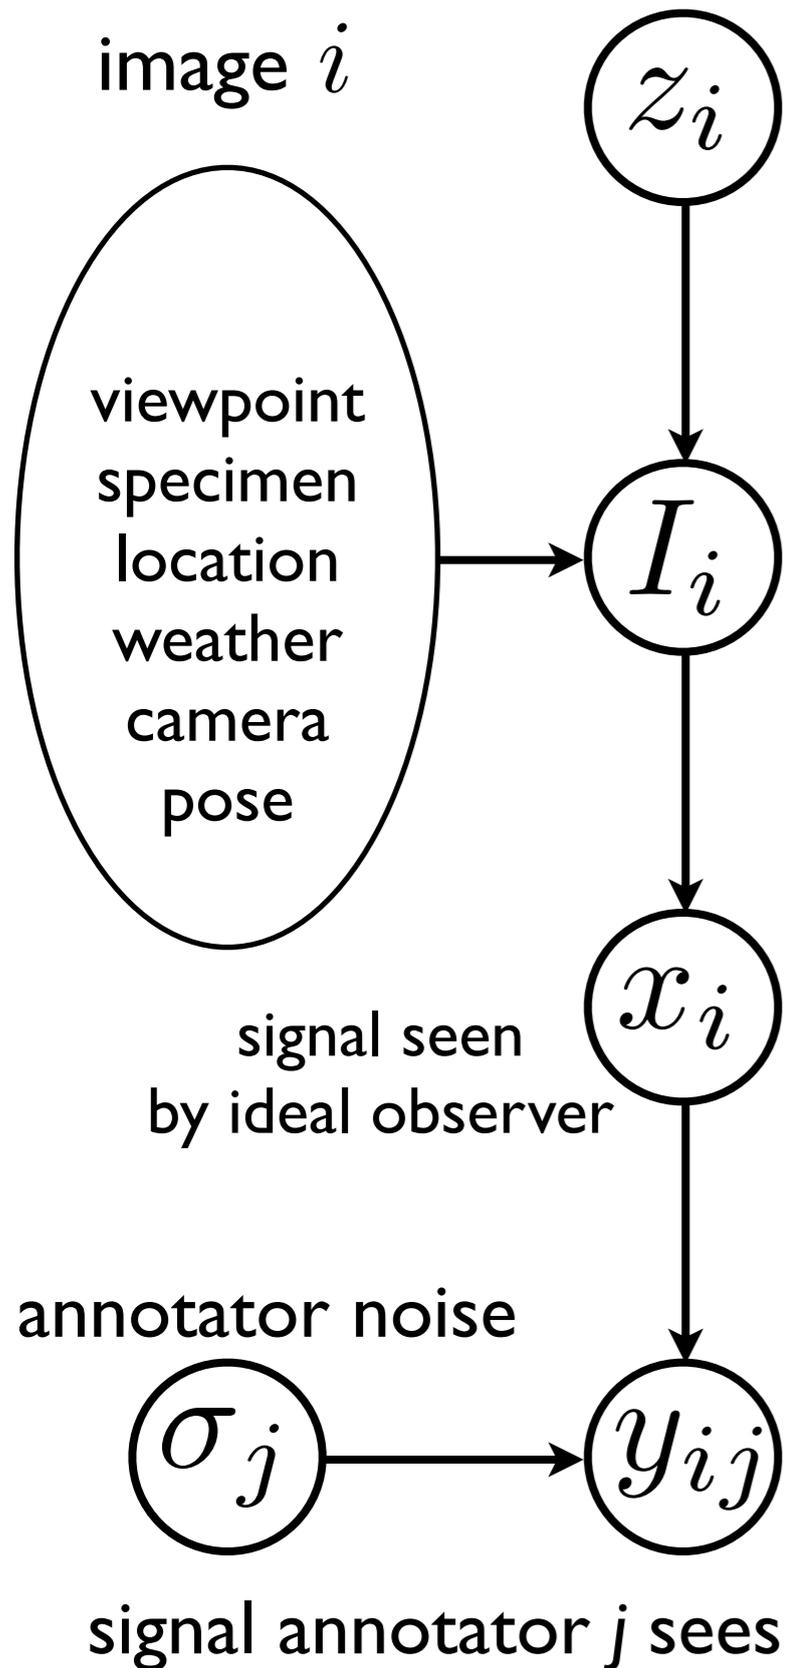
$$p(x_i | z_i = 0)$$



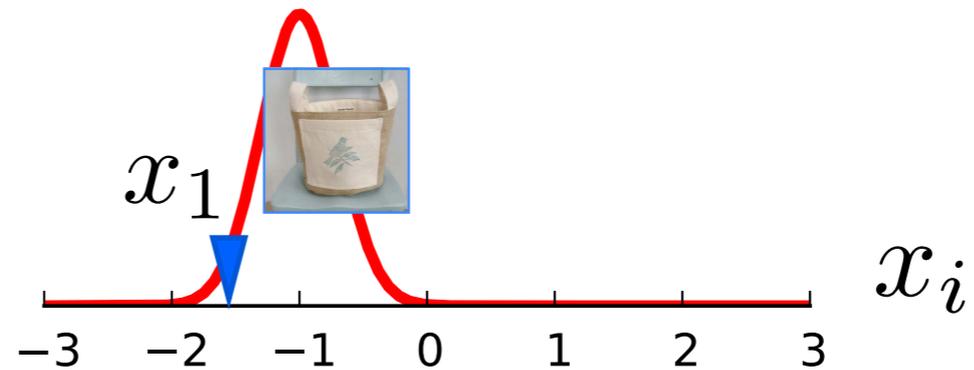
$$\sigma_1 = 0.2$$



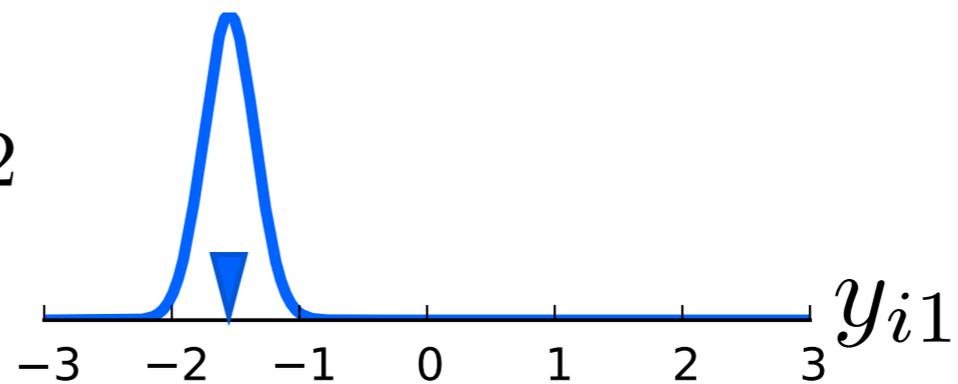
# Image difficulty and annotator competence



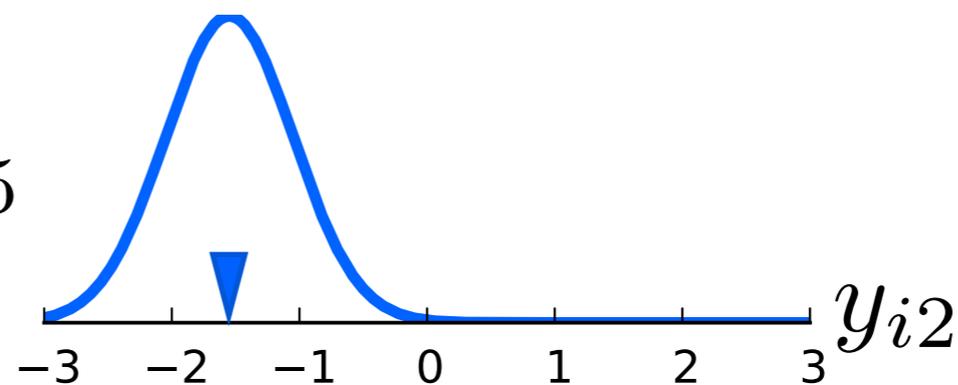
$$p(x_i | z_i = 0)$$



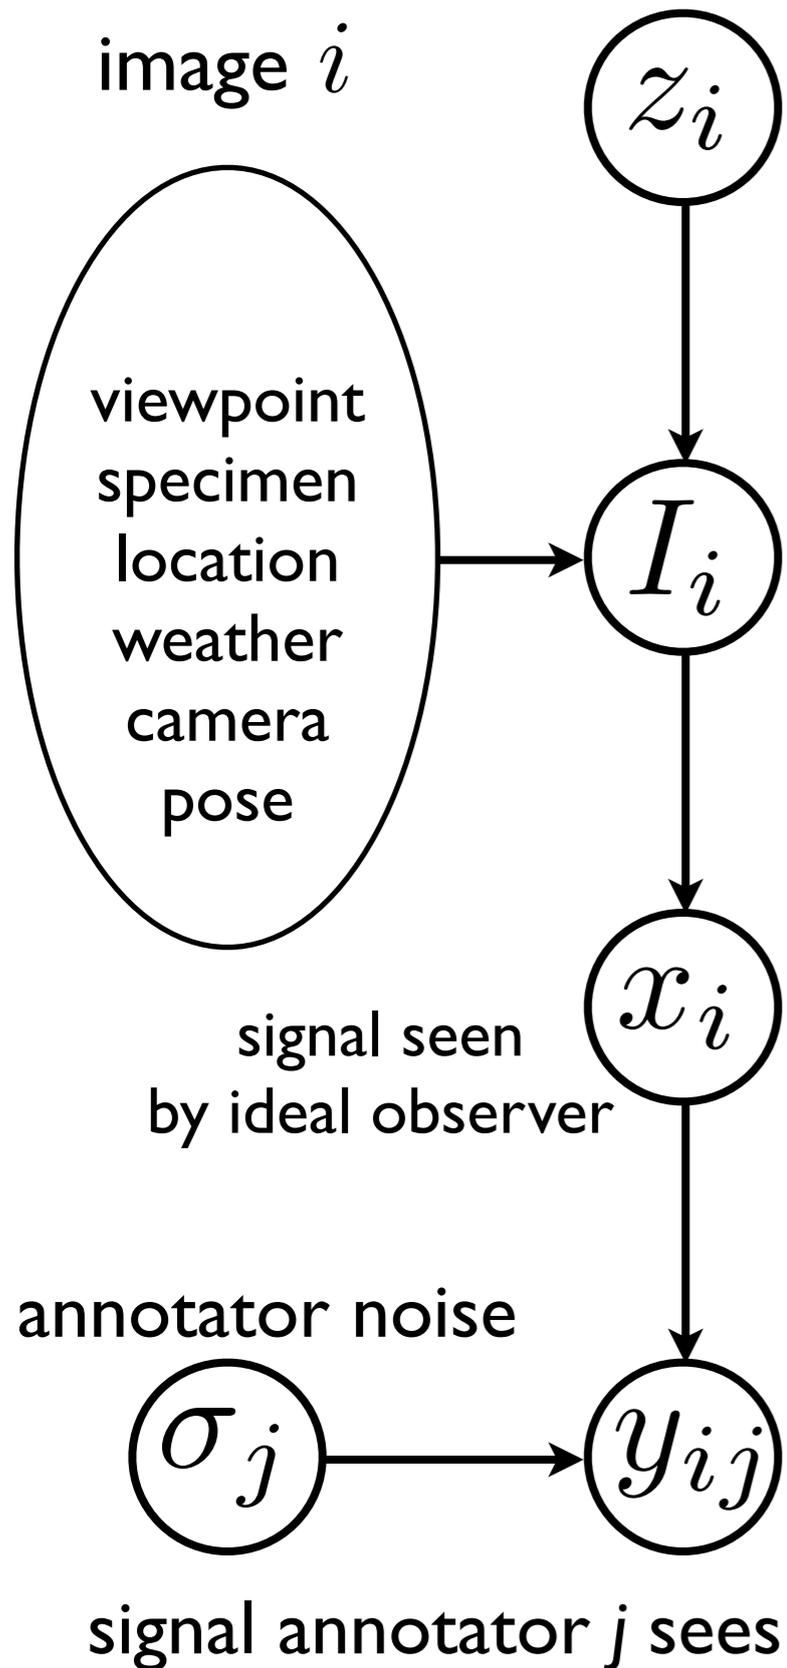
$$\sigma_1 = 0.2$$



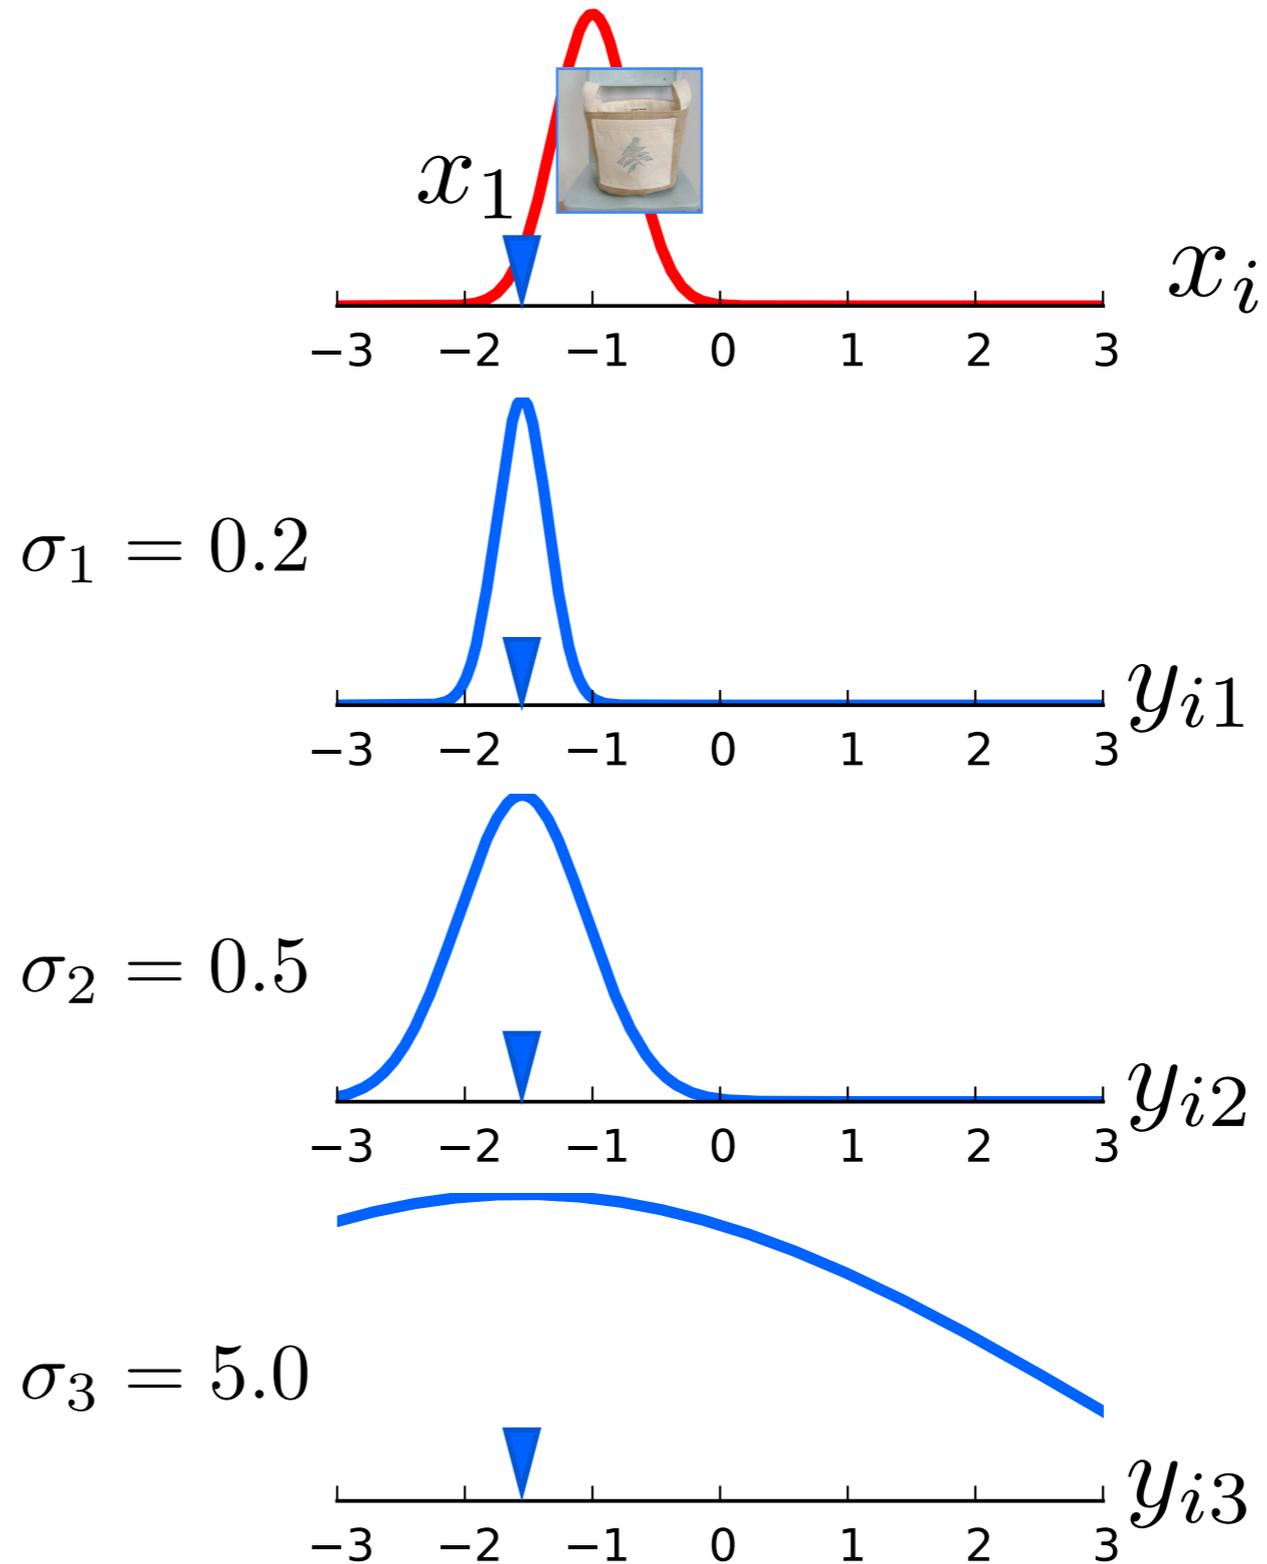
$$\sigma_2 = 0.5$$



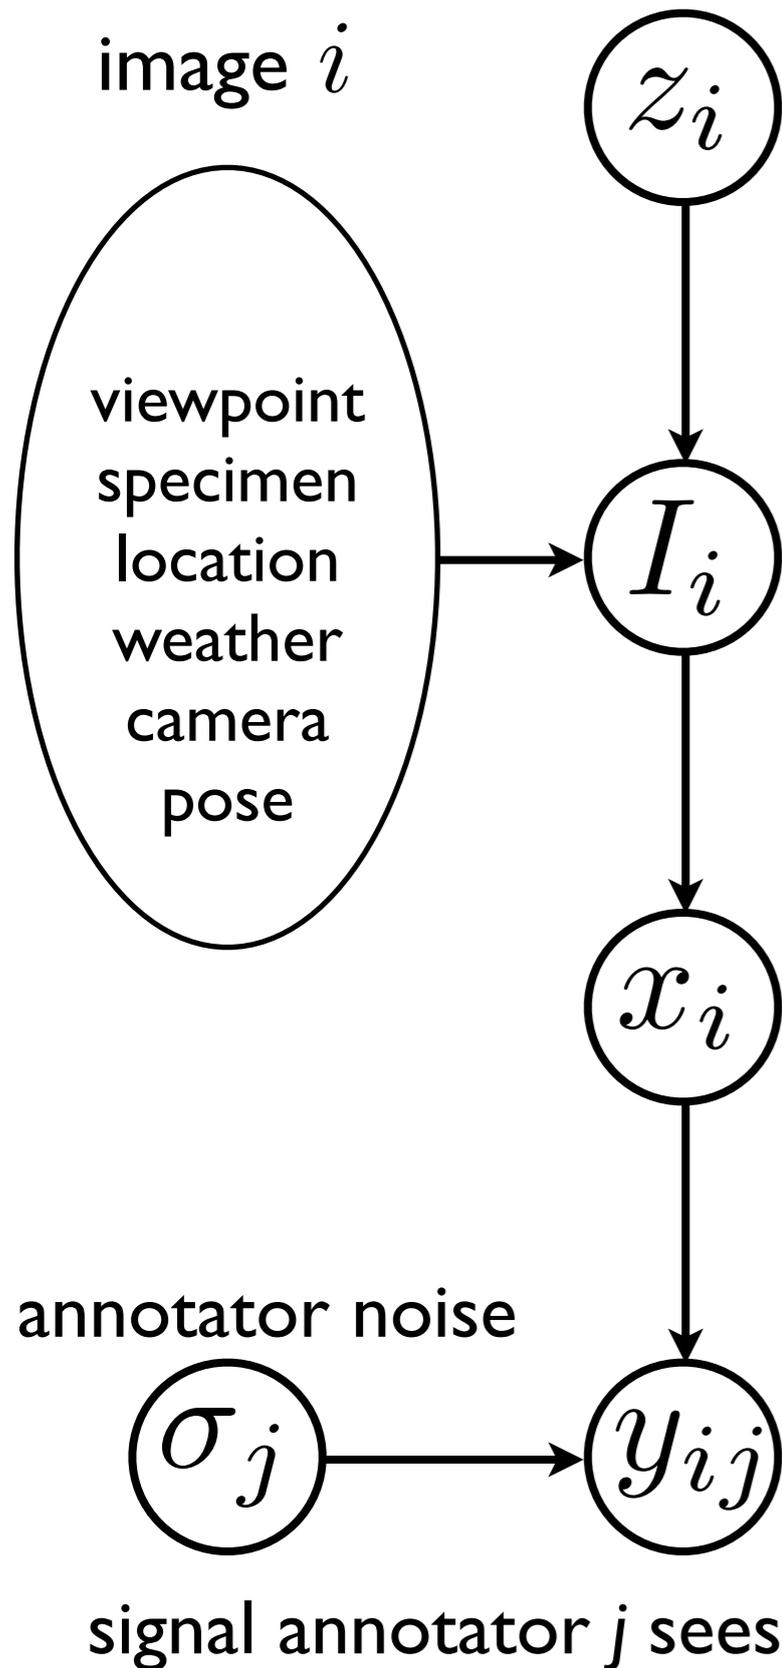
# Image difficulty and annotator competence



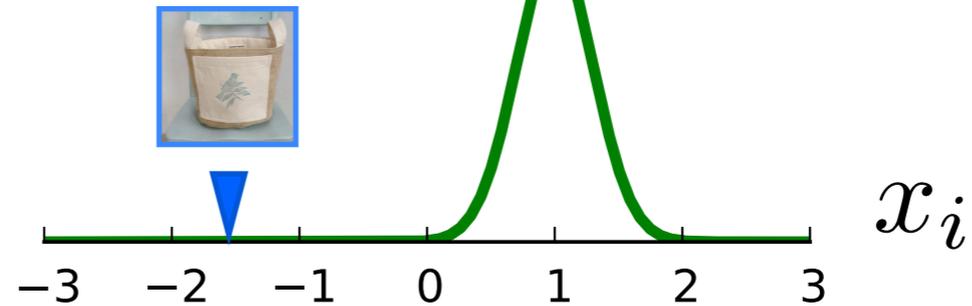
$$p(x_i | z_i = 0)$$



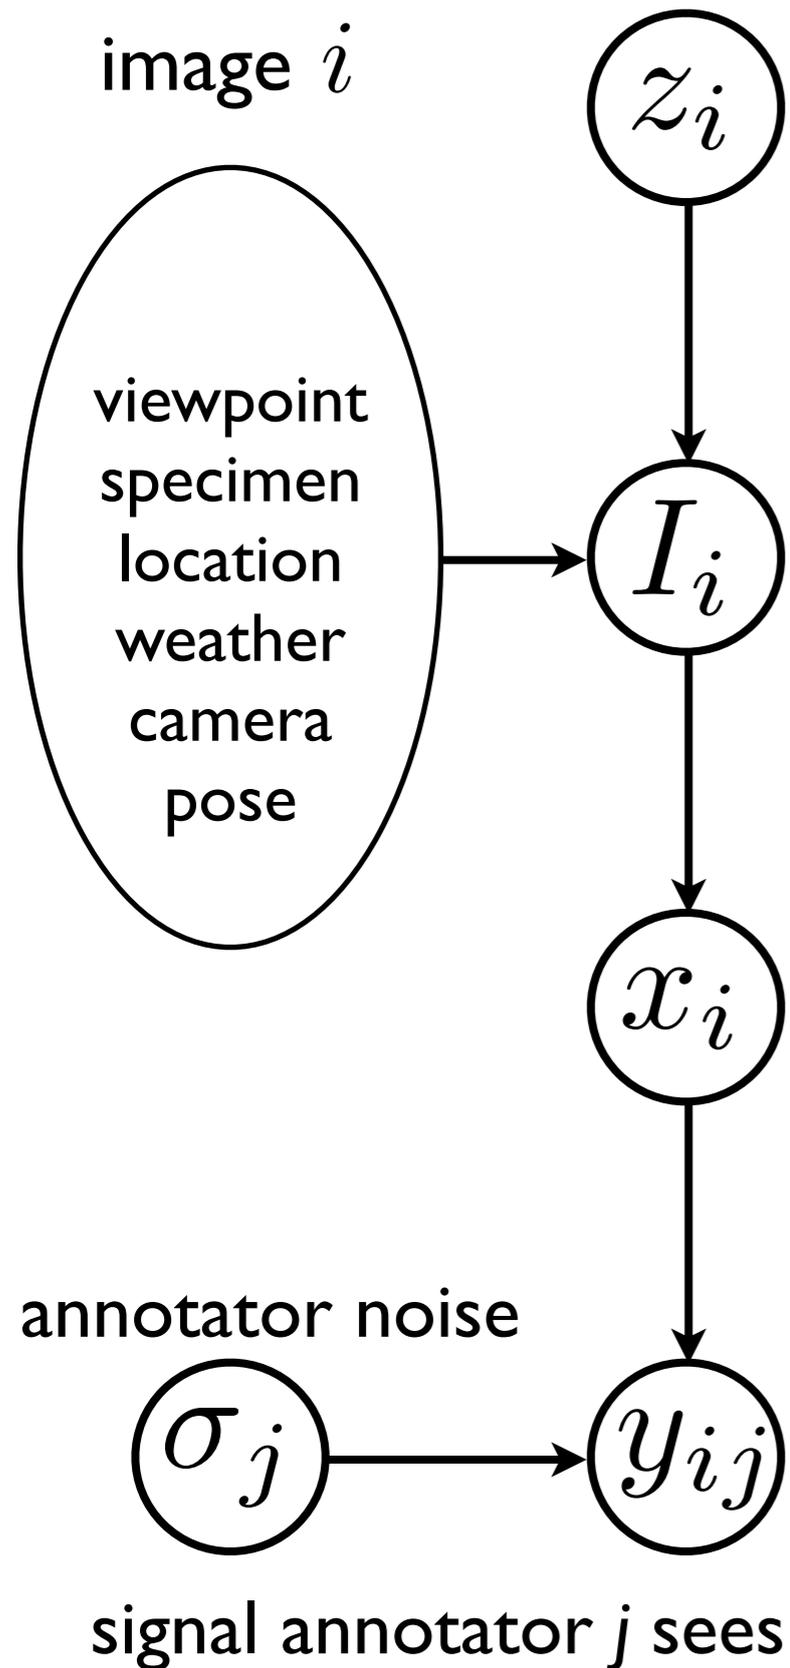
# Image difficulty and annotator competence



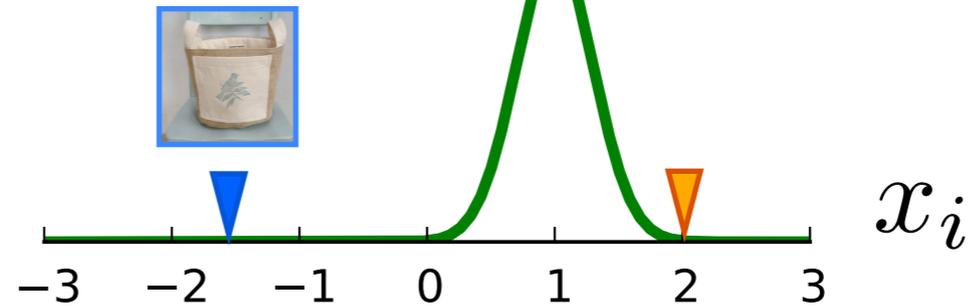
$$p(x_i | z_i = 1)$$



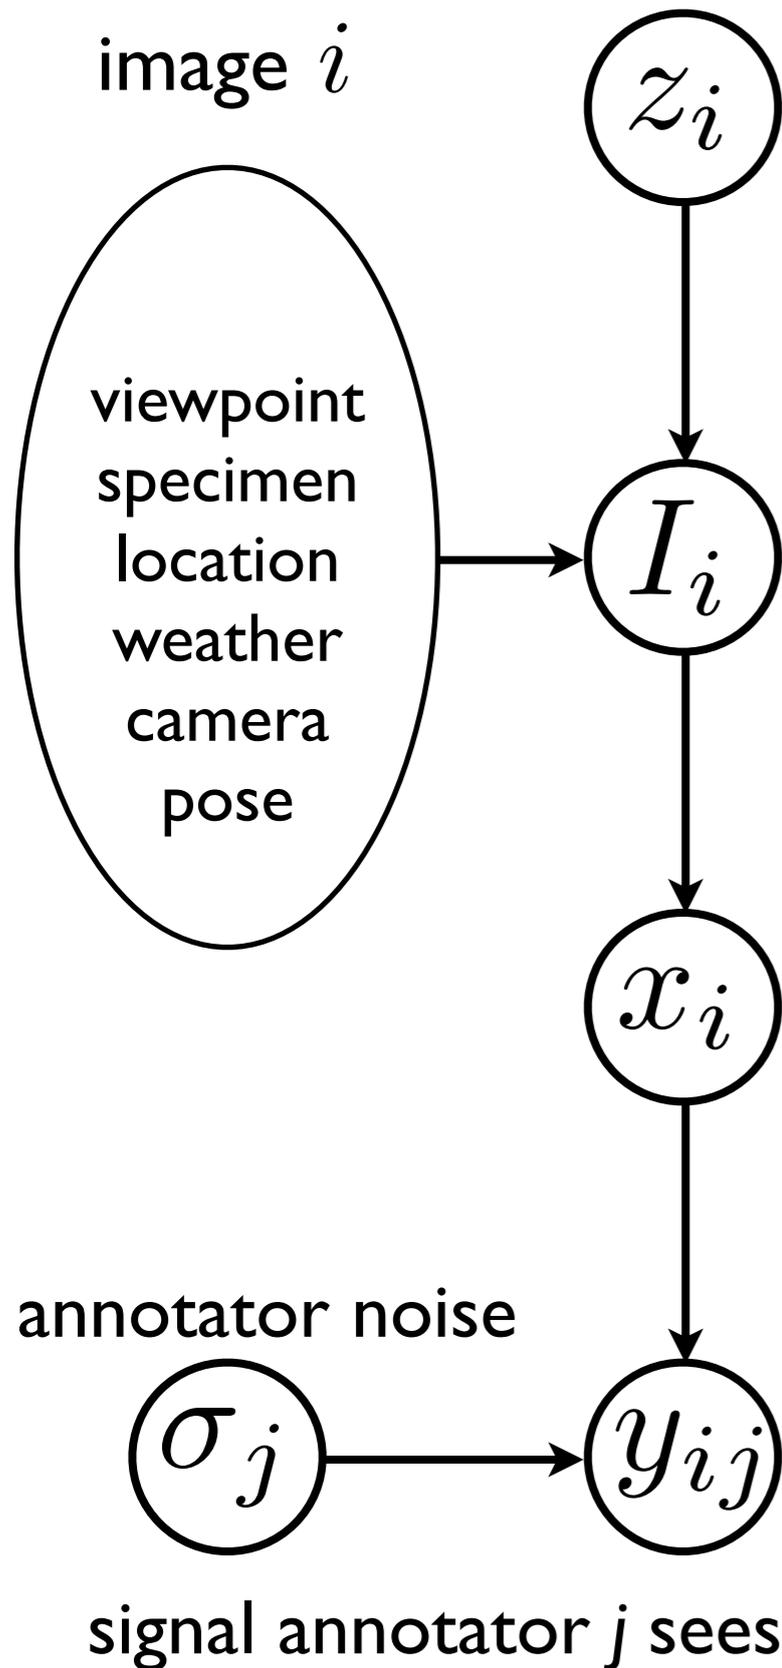
# Image difficulty and annotator competence



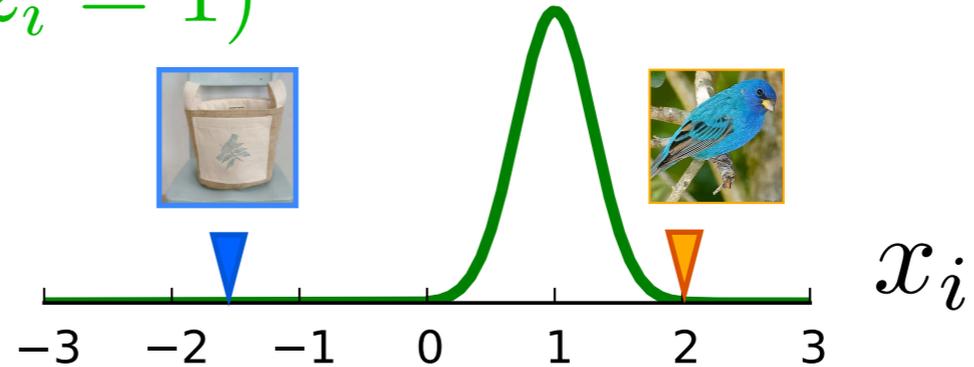
$$p(x_i | z_i = 1)$$



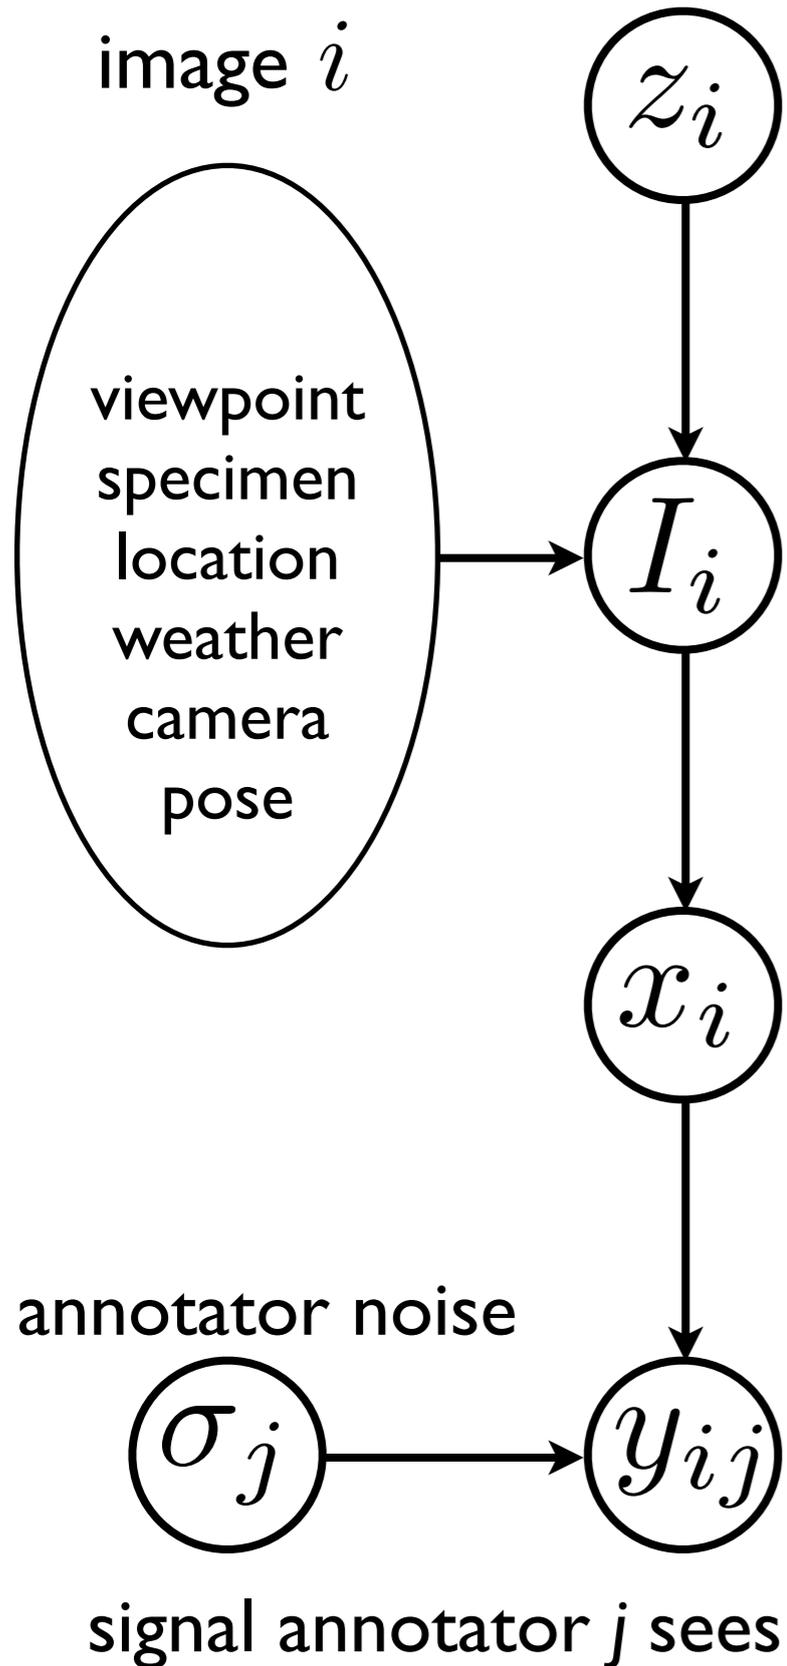
# Image difficulty and annotator competence



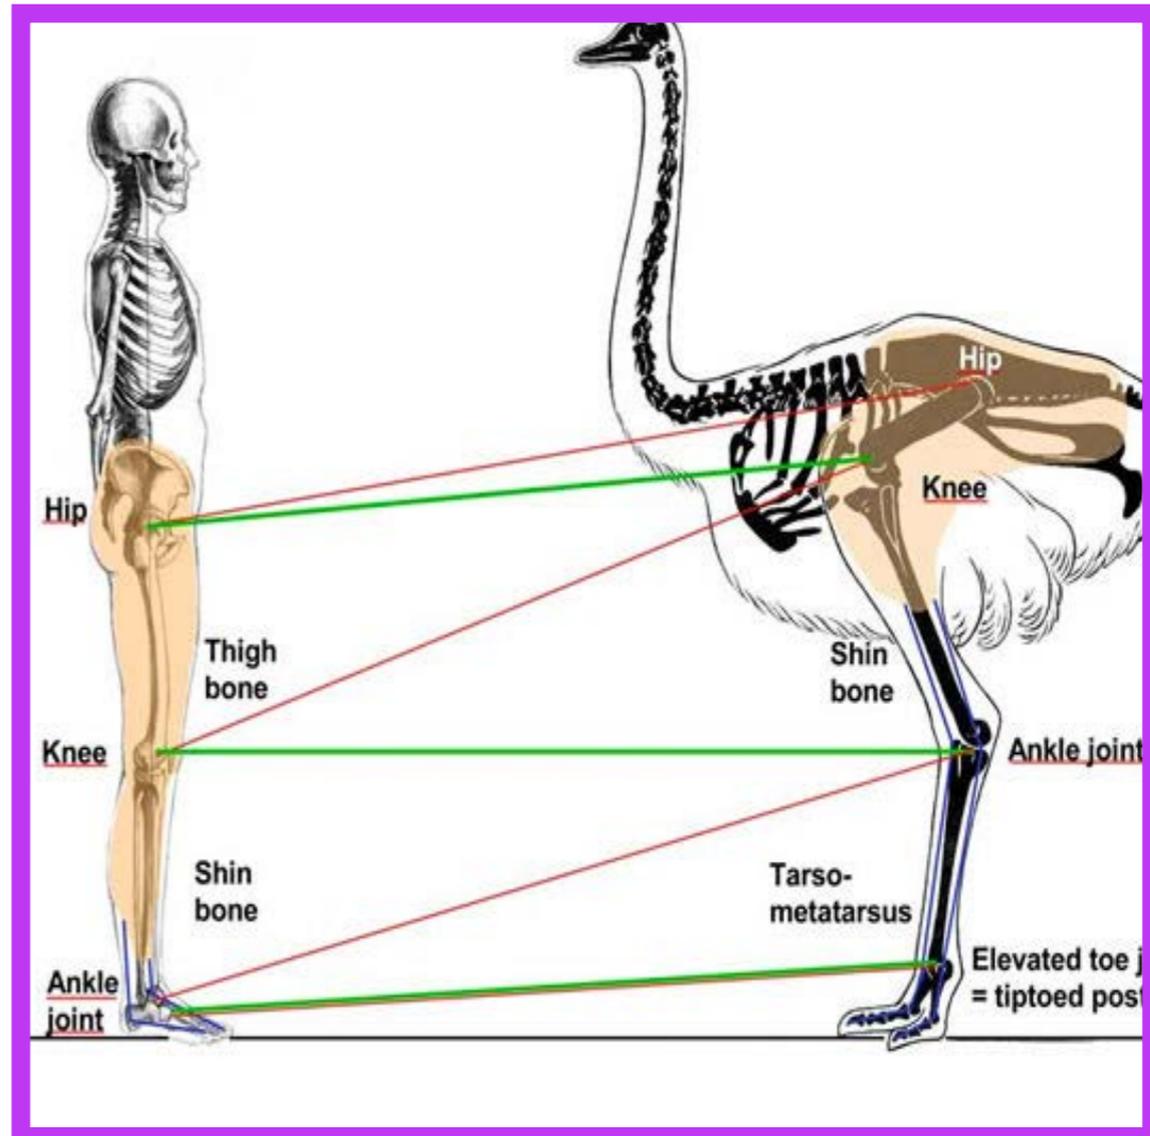
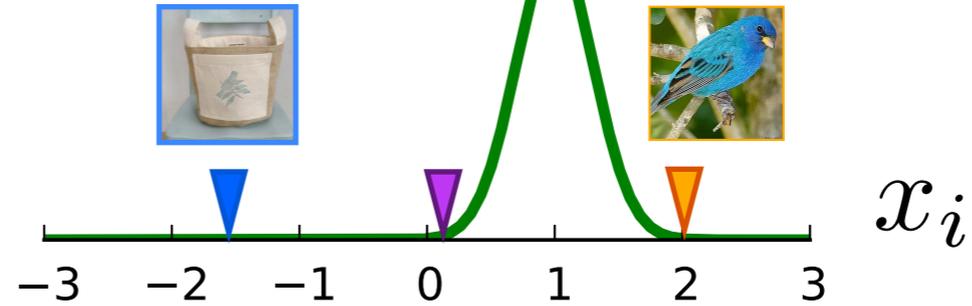
$$p(x_i | z_i = 1)$$



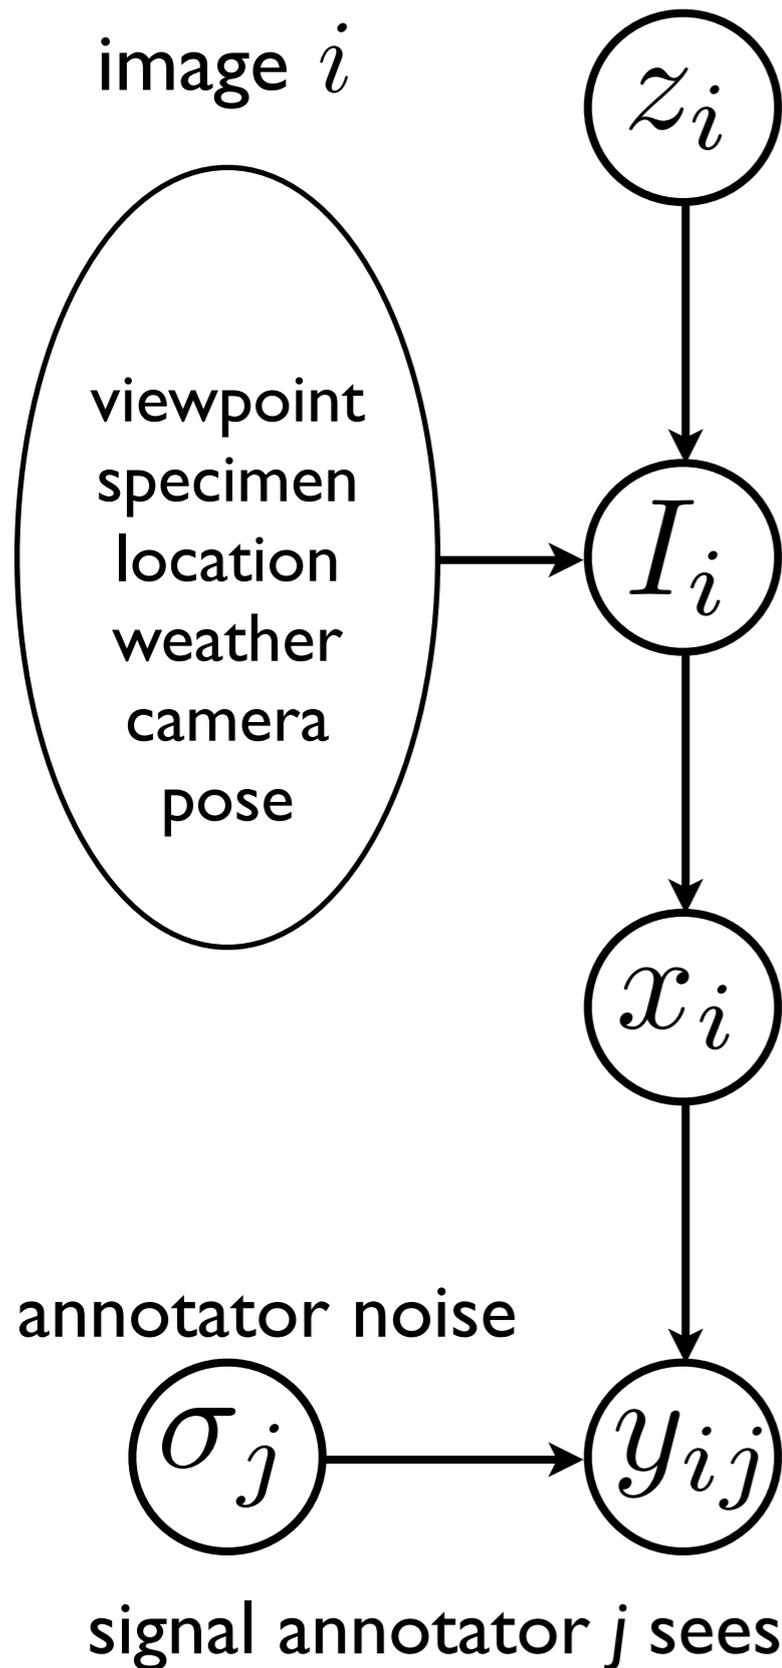
# Image difficulty and annotator competence



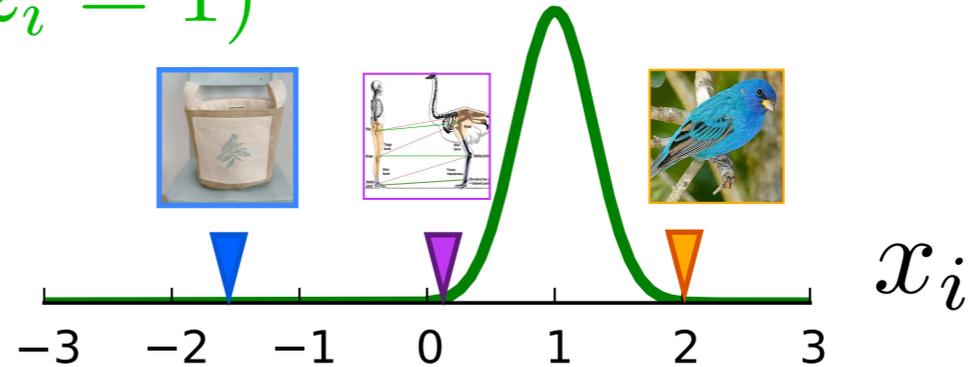
$$p(x_i | z_i = 1)$$



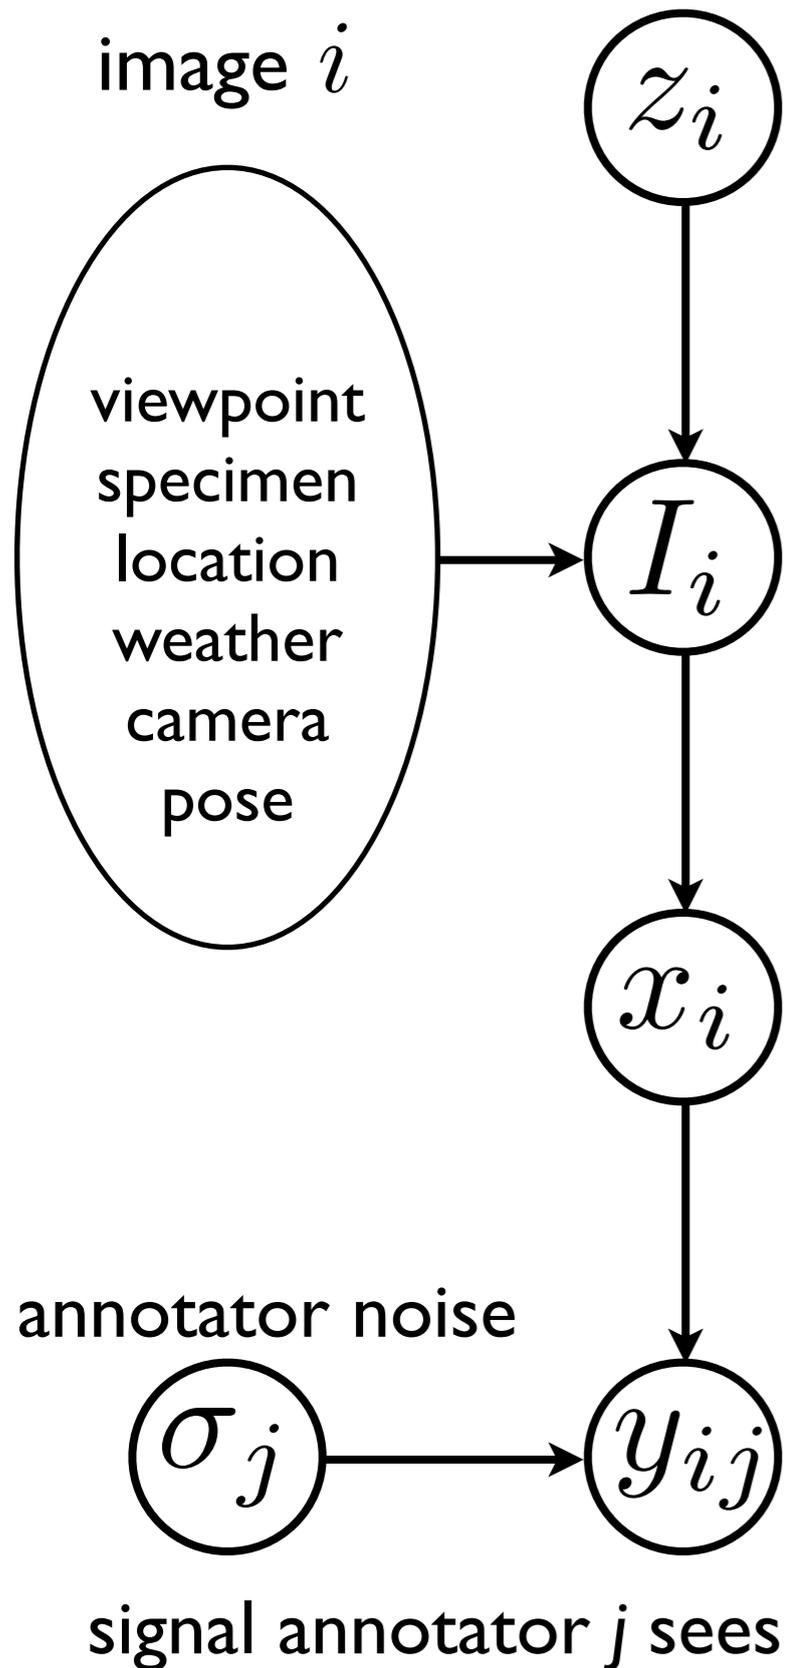
# Image difficulty and annotator competence



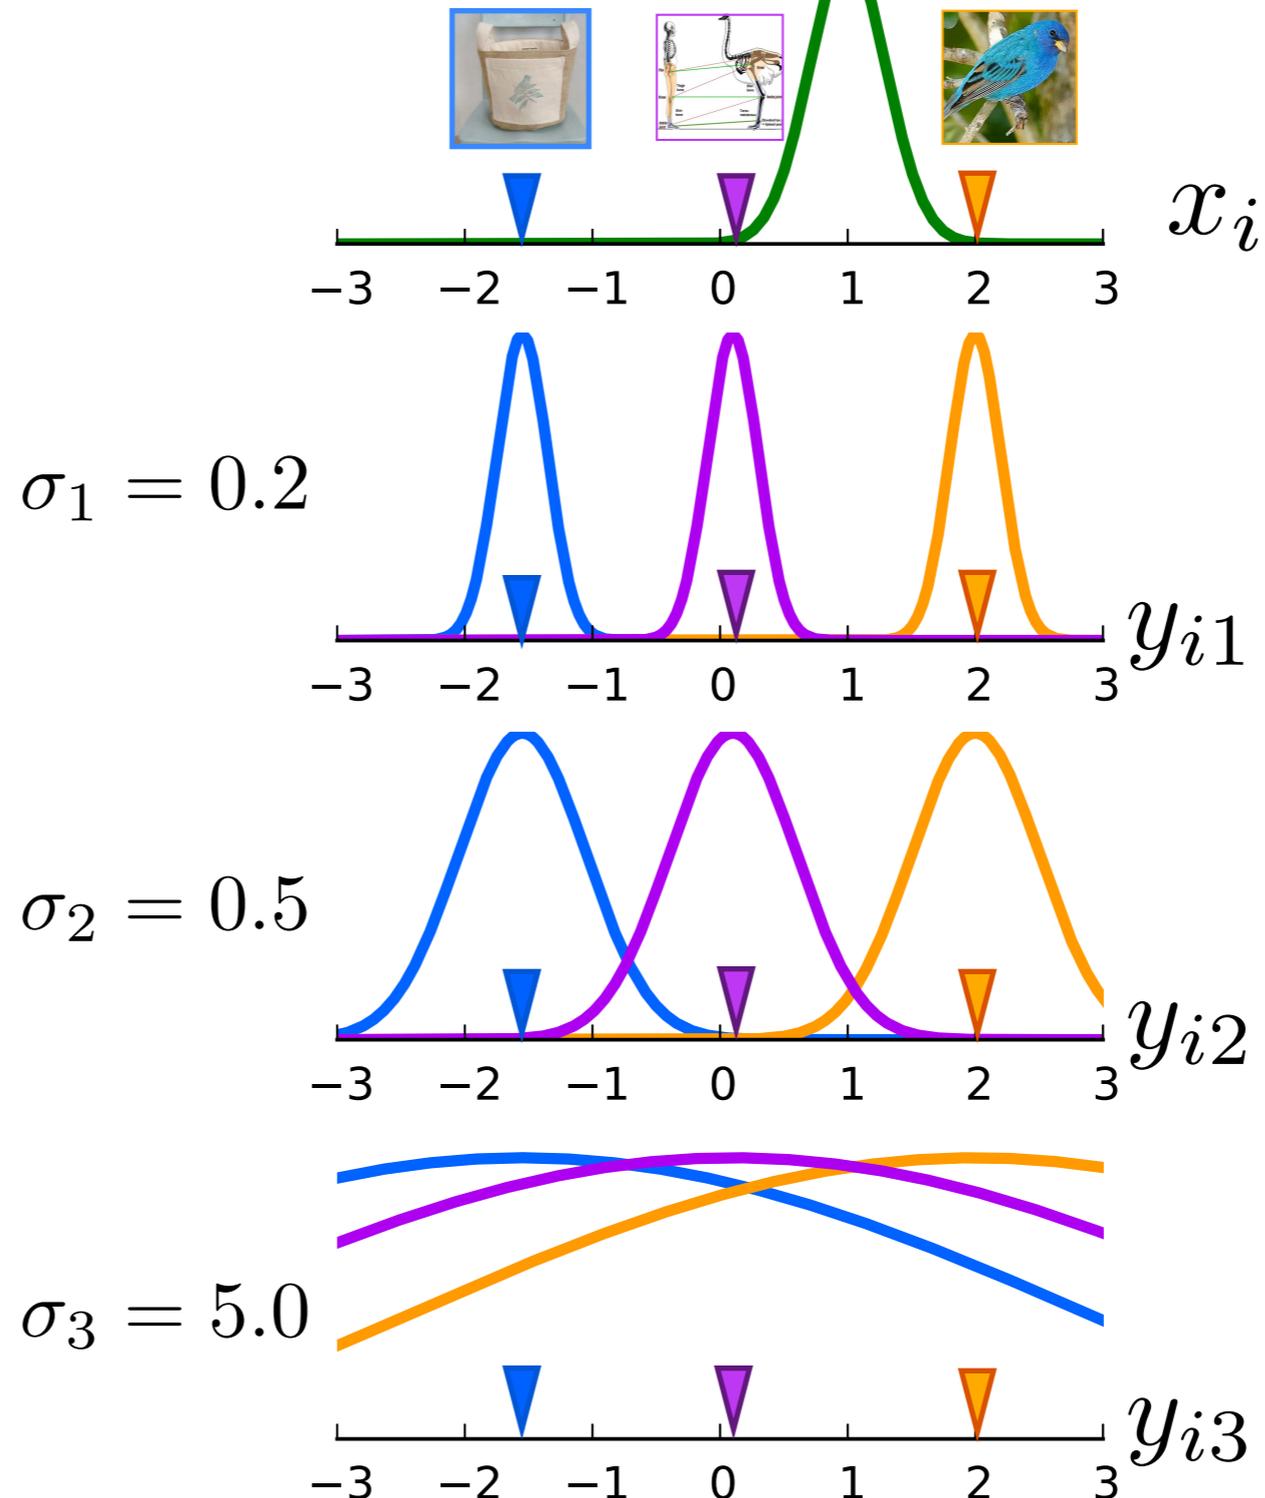
$$p(x_i | z_i = 1)$$



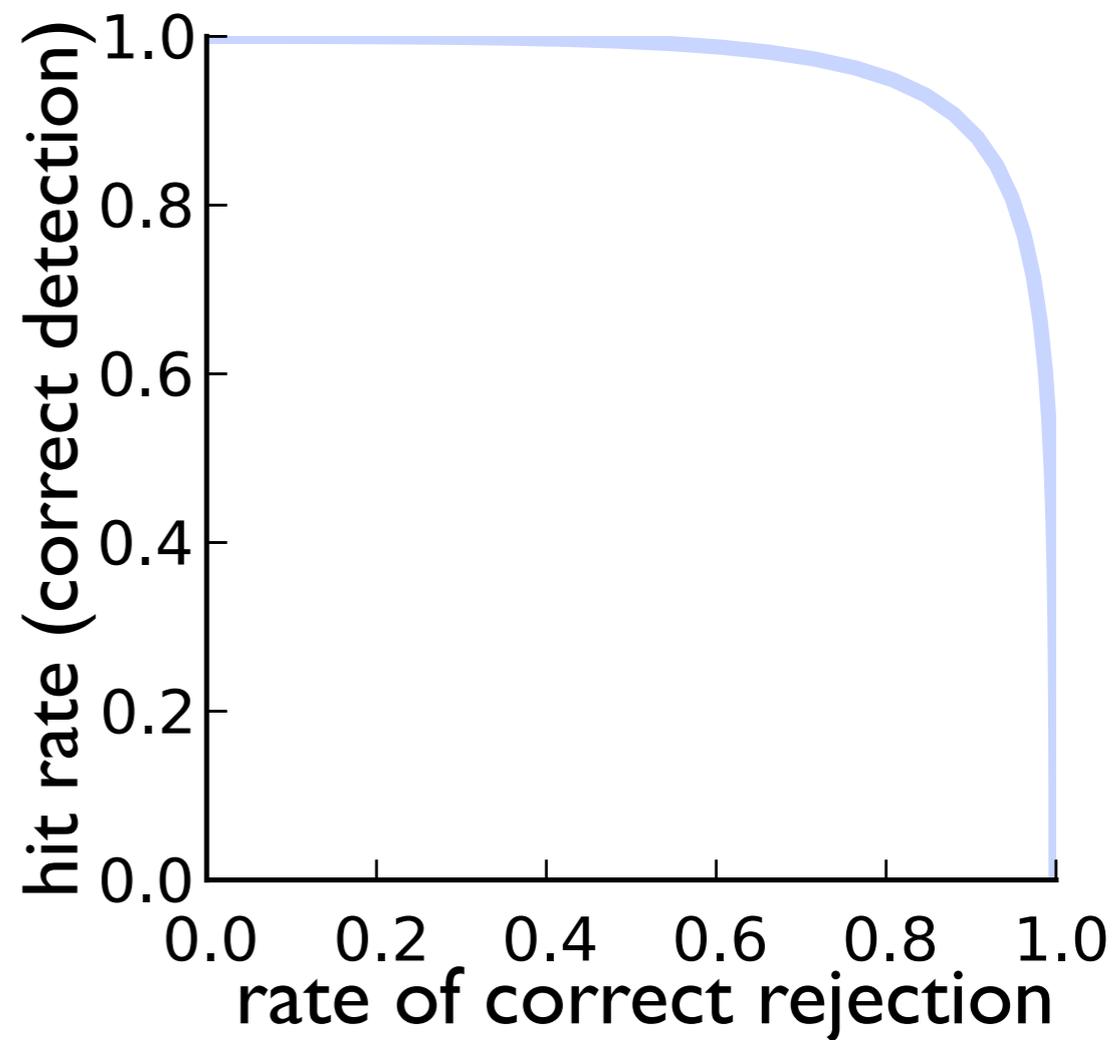
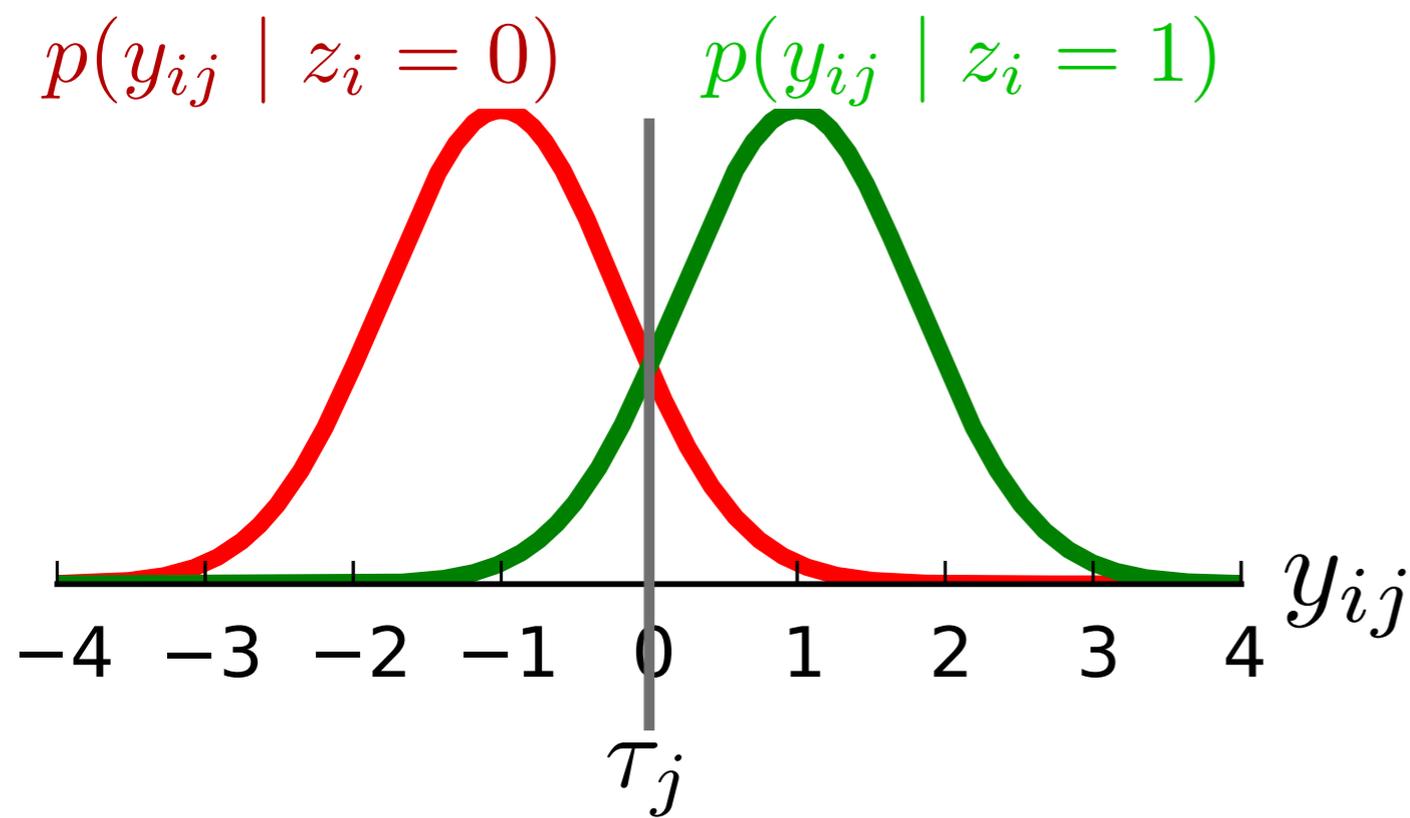
# Image difficulty and annotator competence



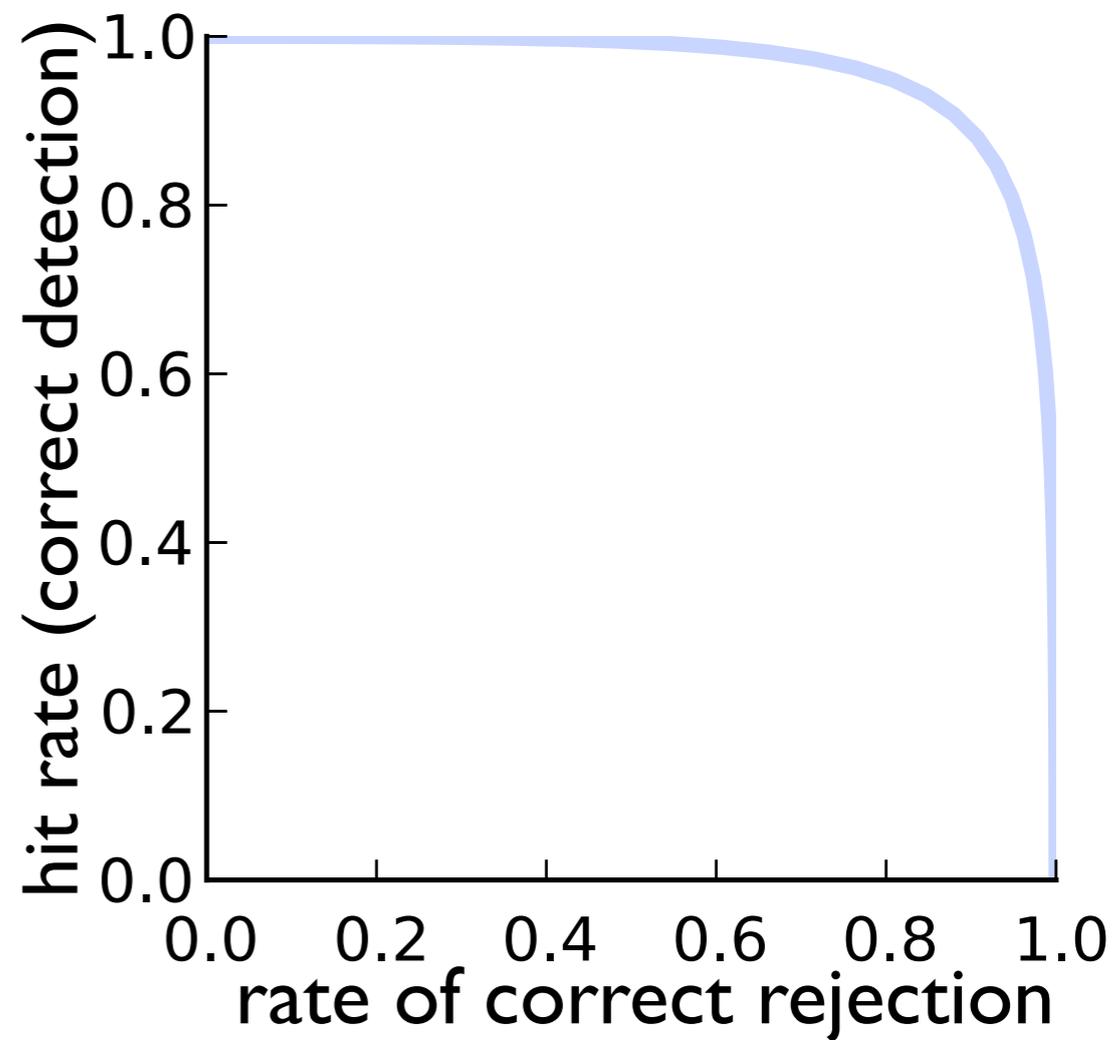
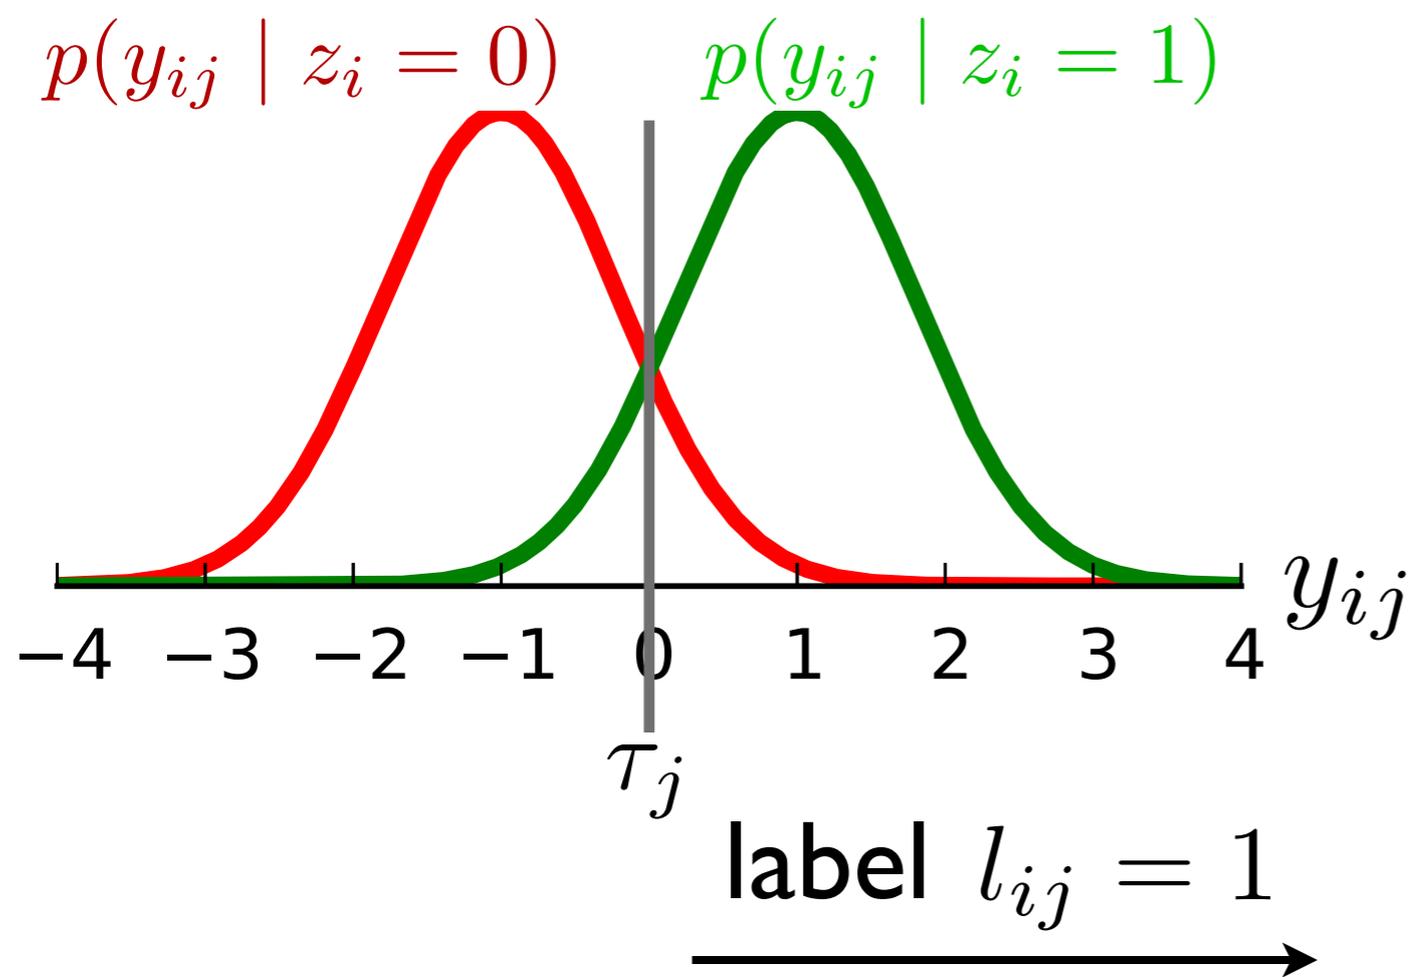
$$p(x_i | z_i = 1)$$



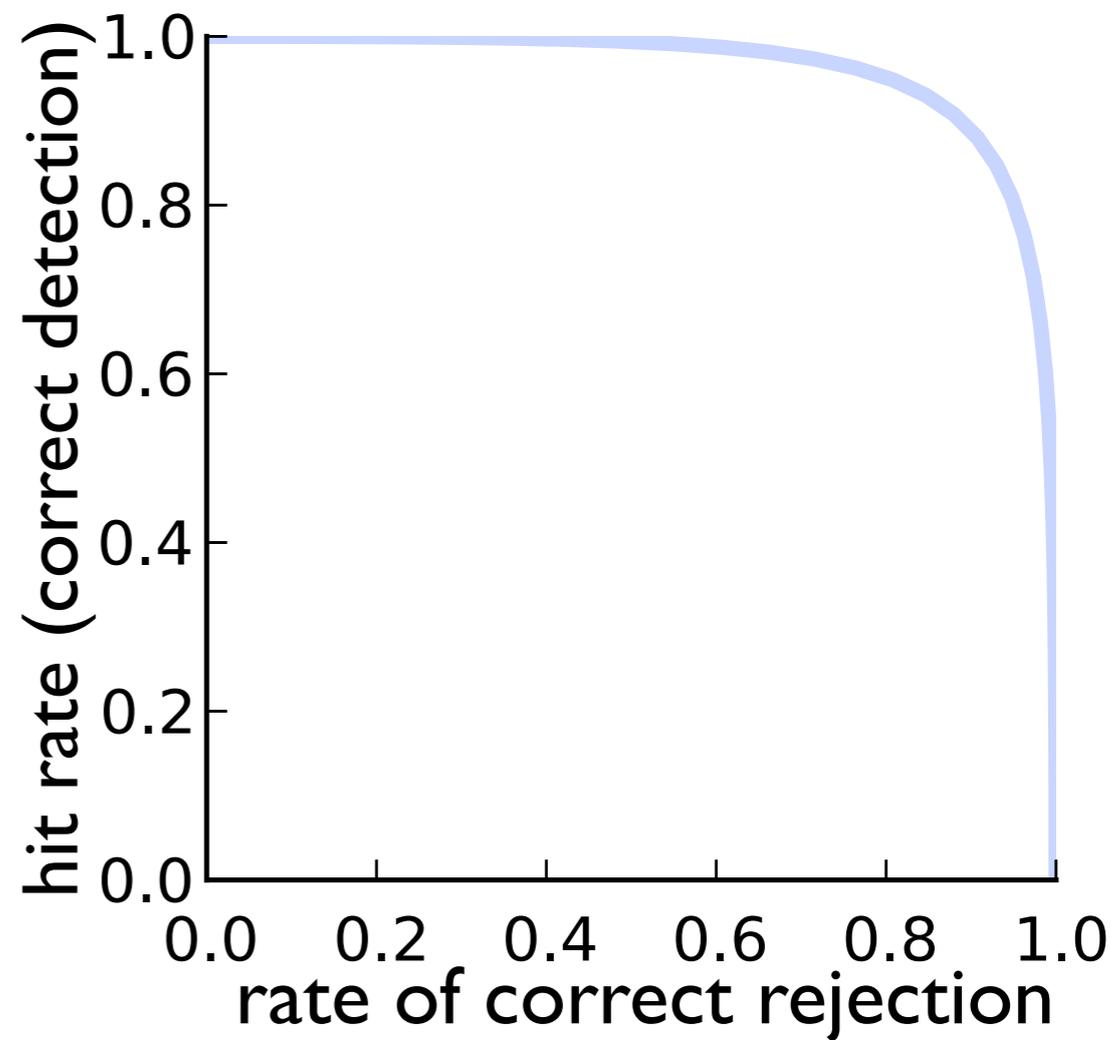
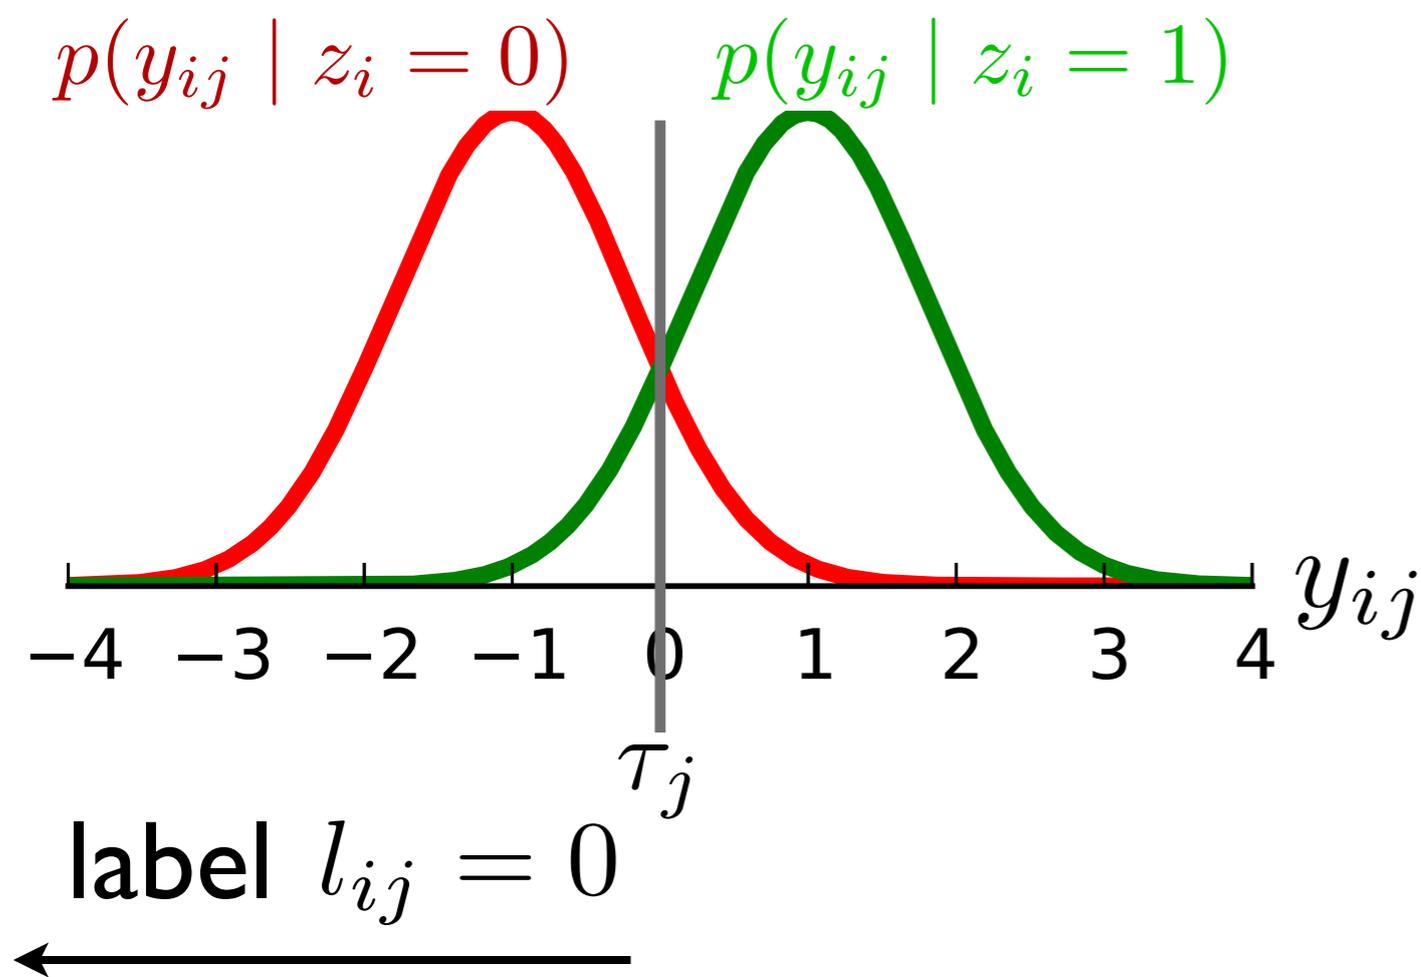
# Optimists and pessimists: $\tau_j$



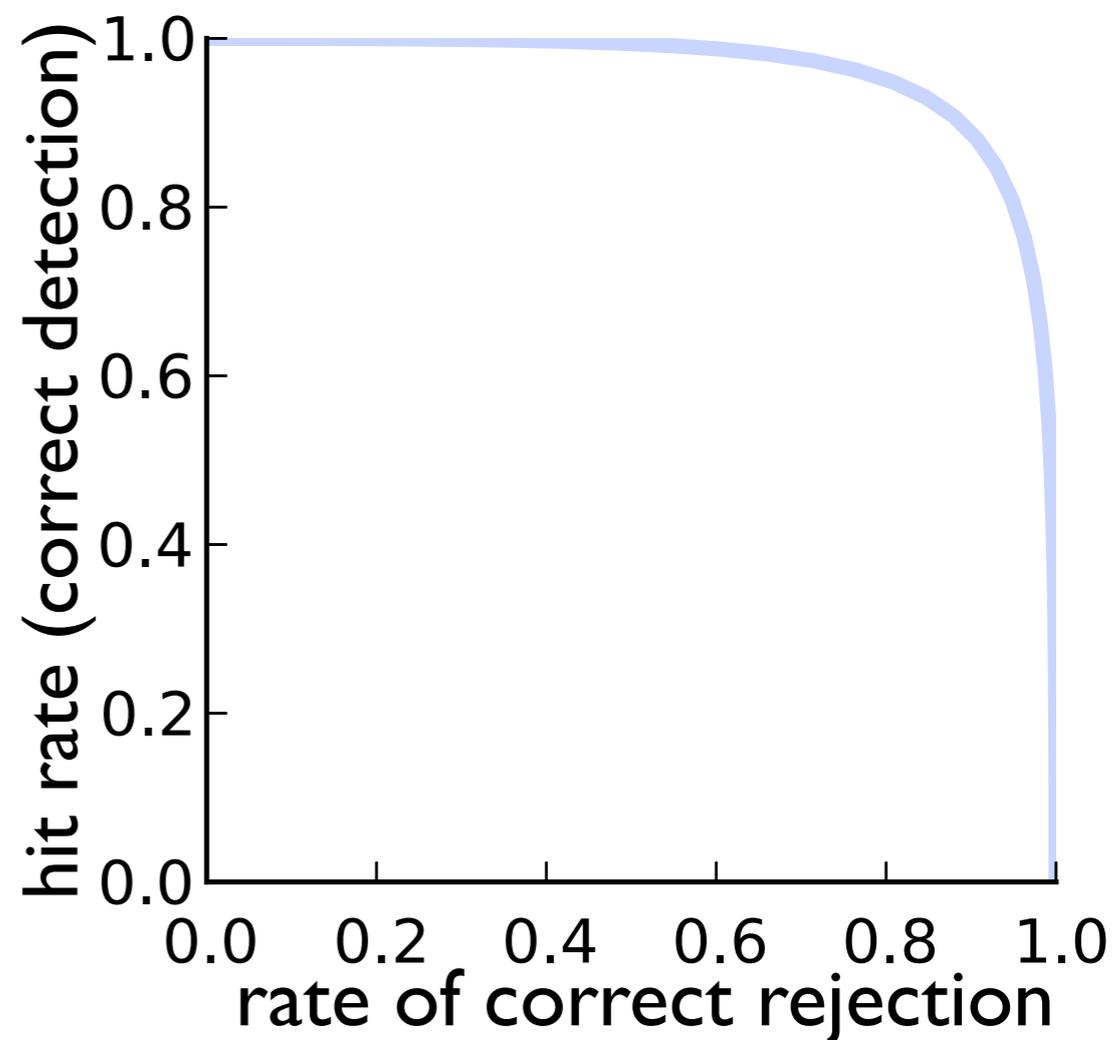
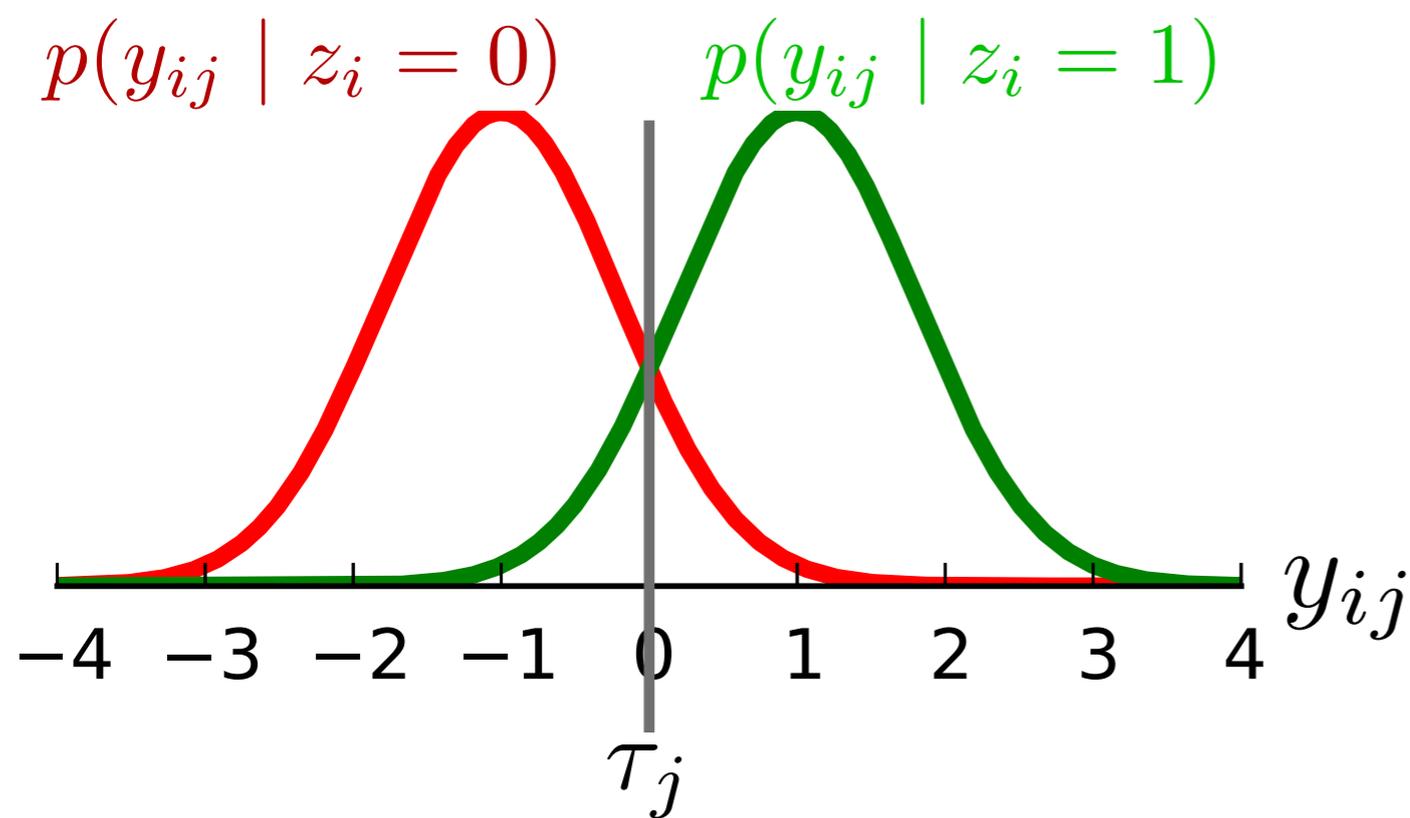
# Optimists and pessimists: $\tau_j$



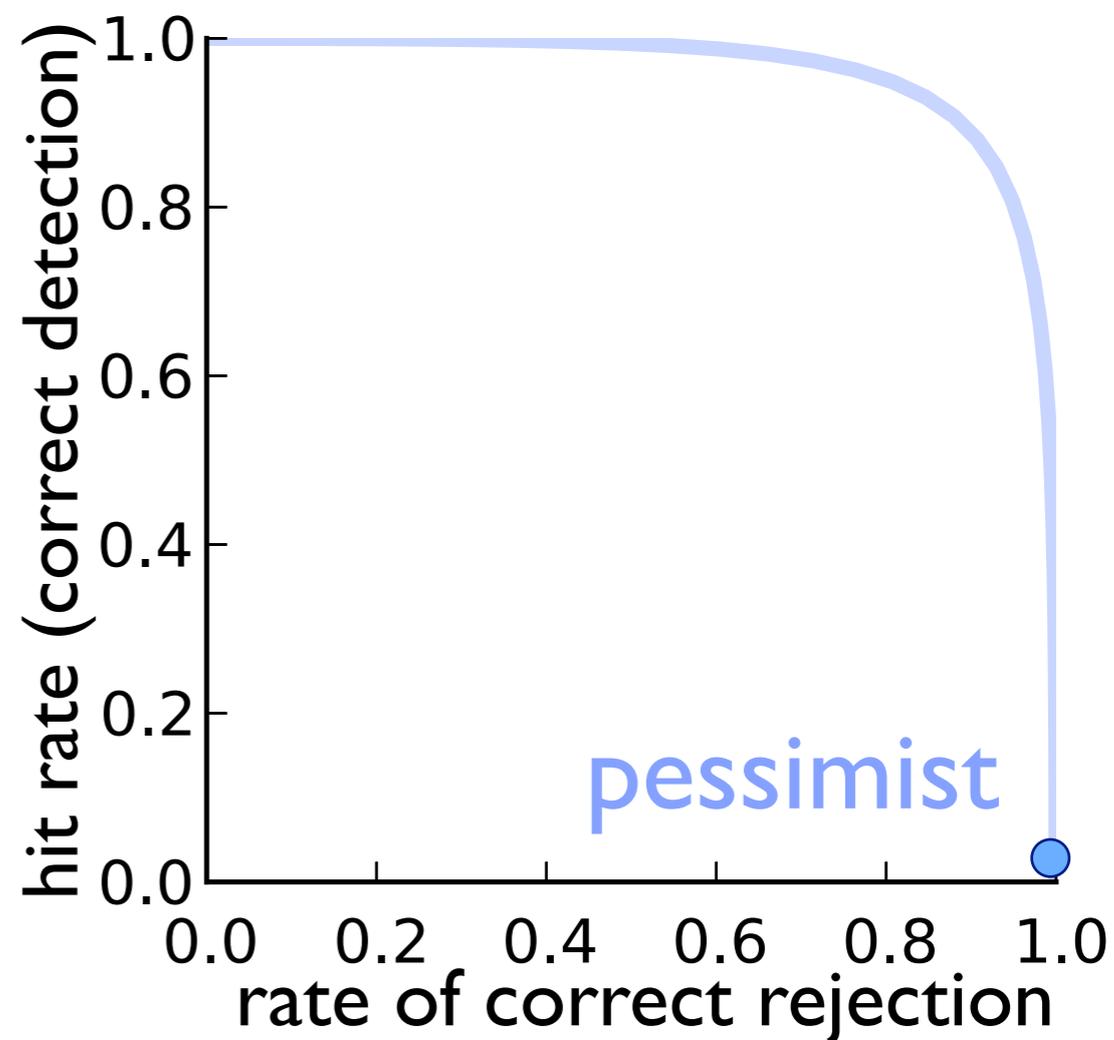
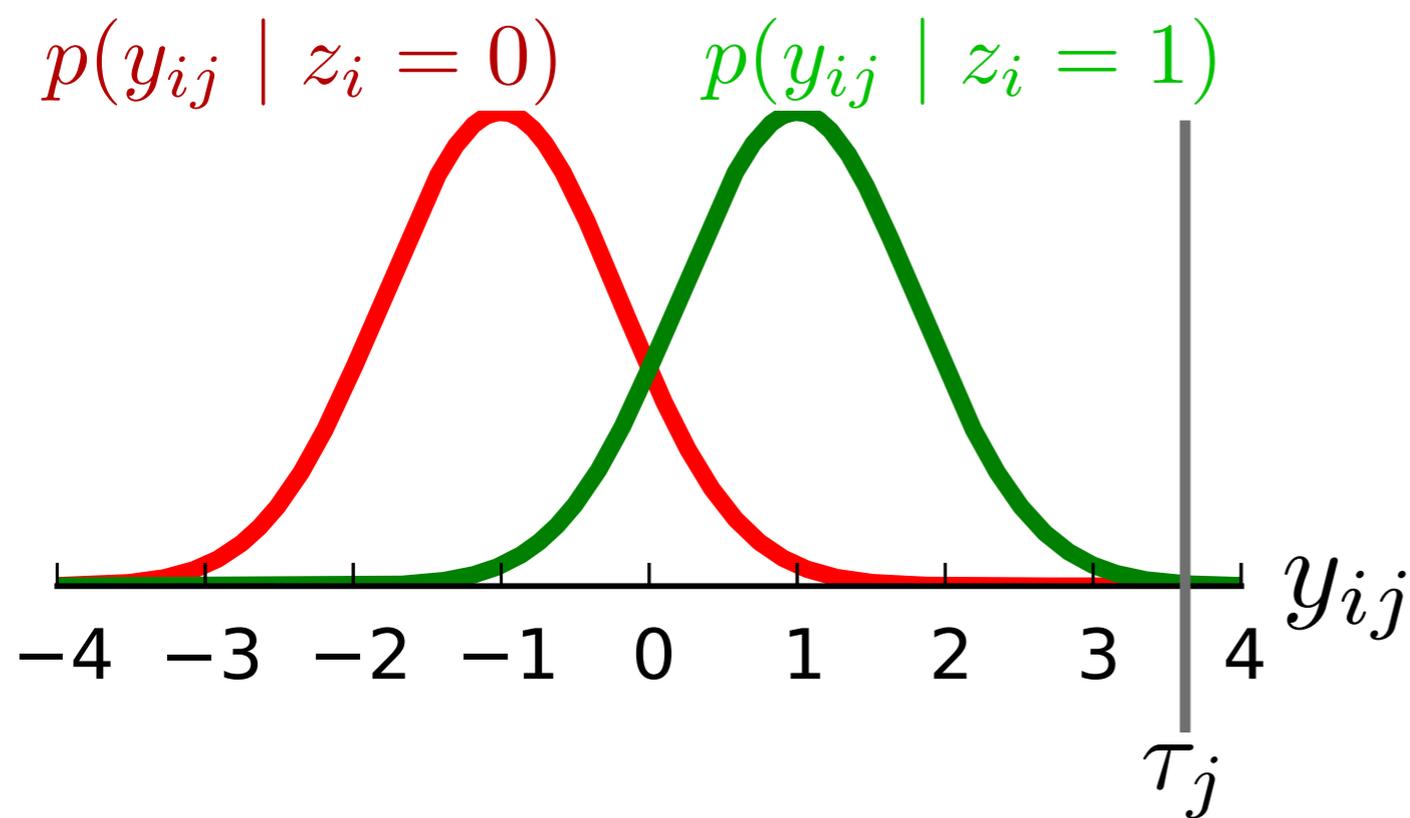
# Optimists and pessimists: $\tau_j$



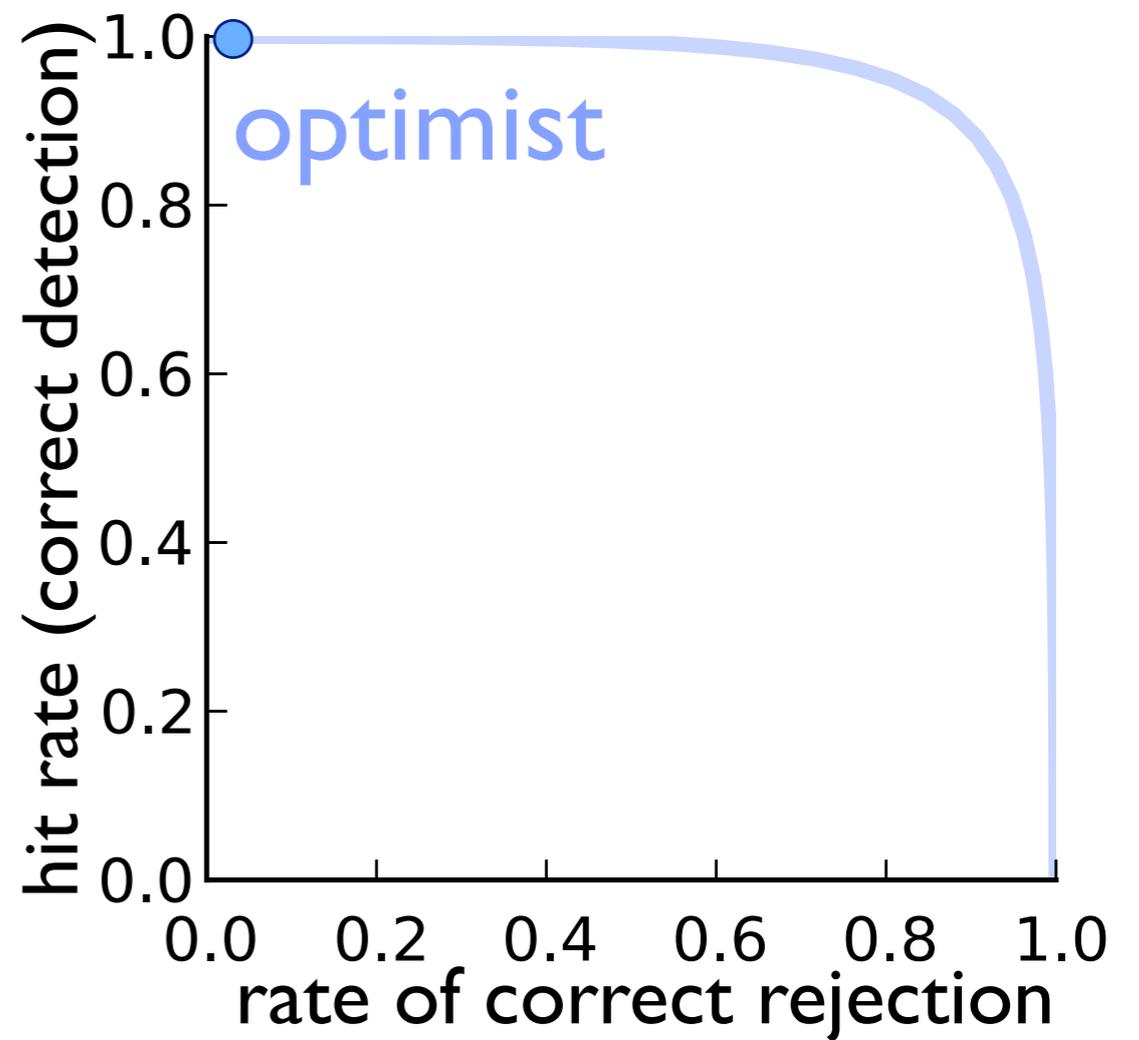
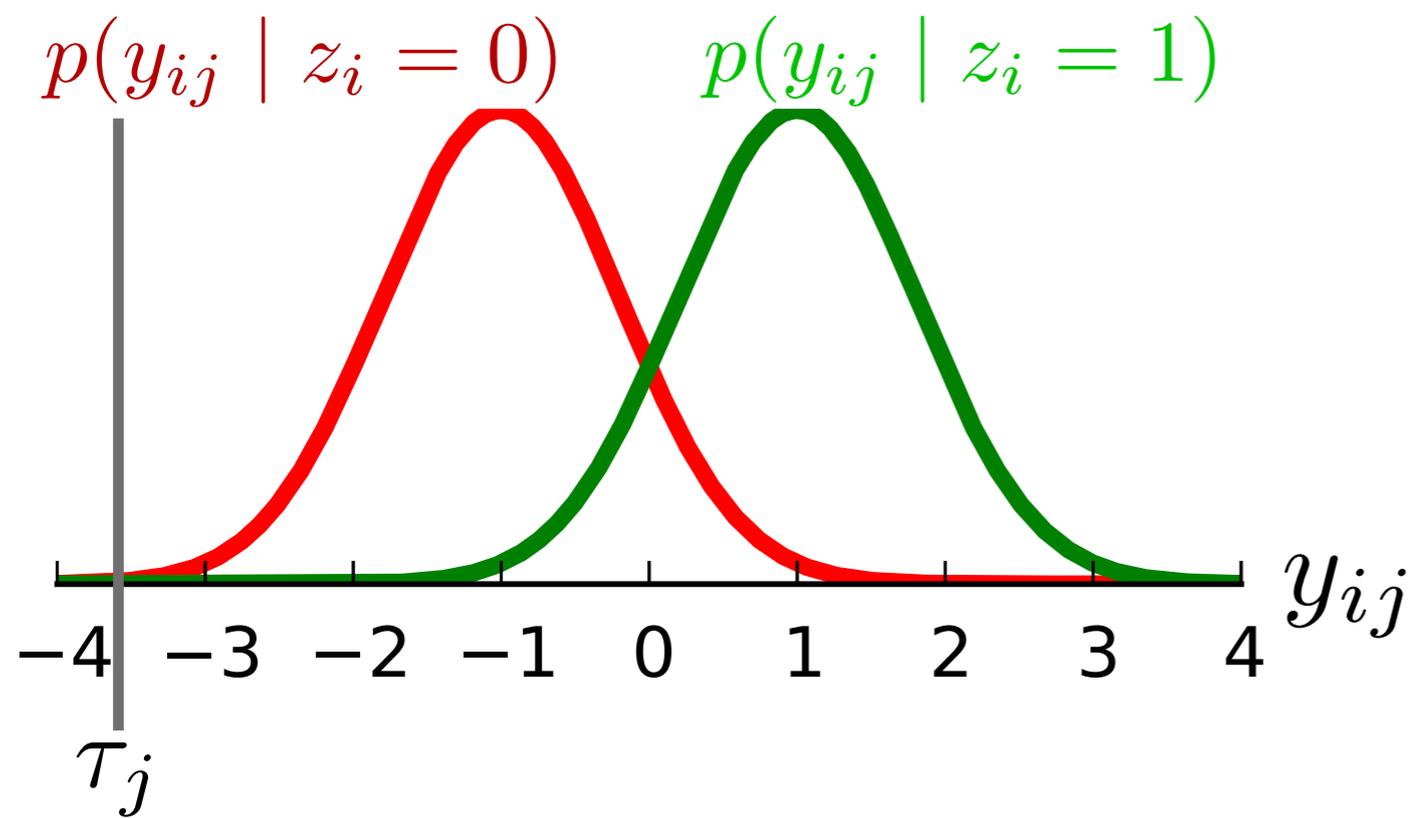
# Optimists and pessimists: $\tau_j$



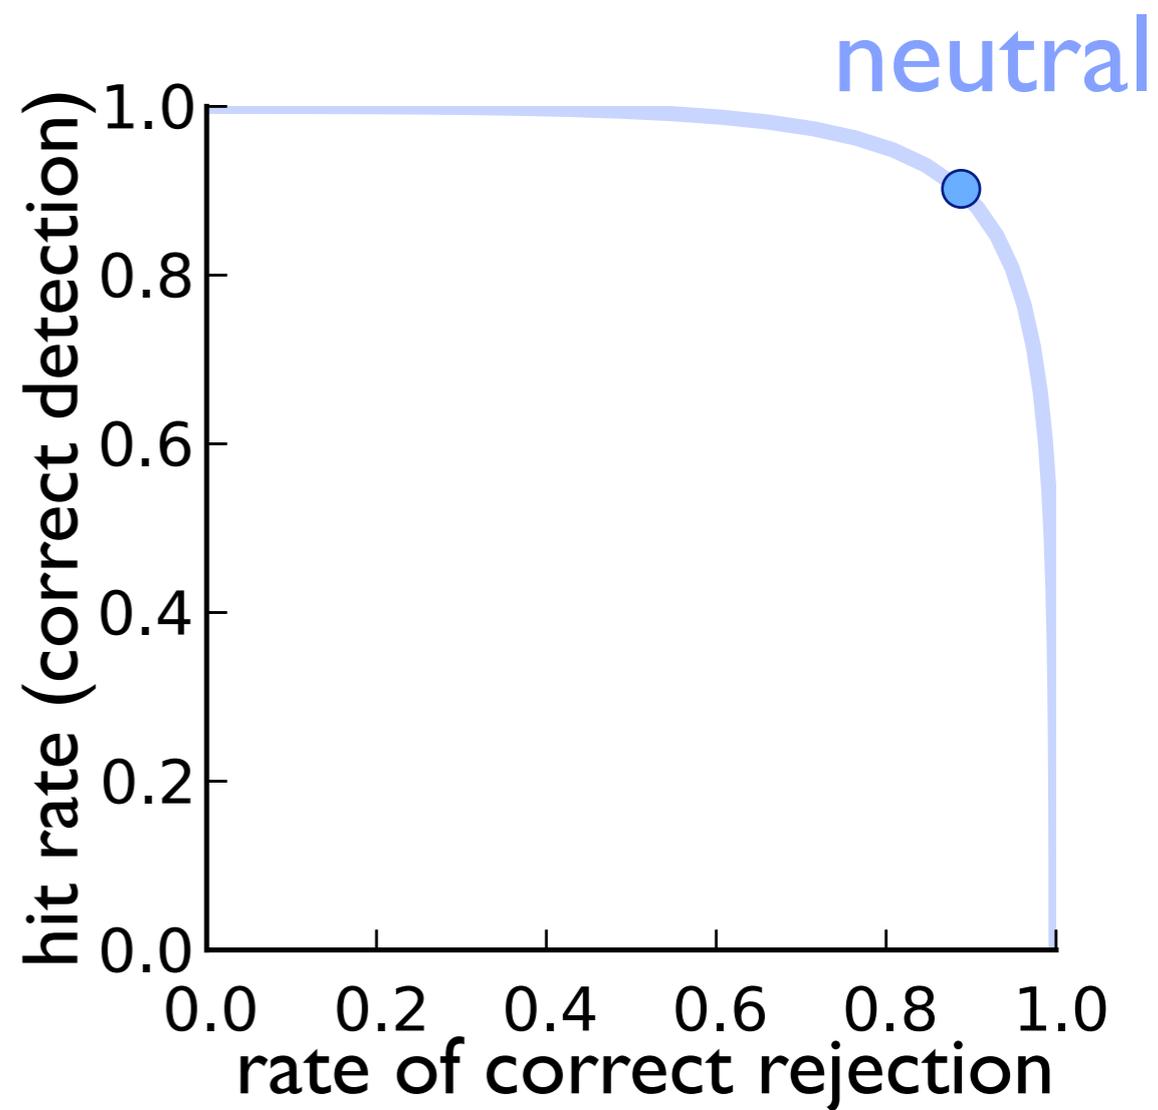
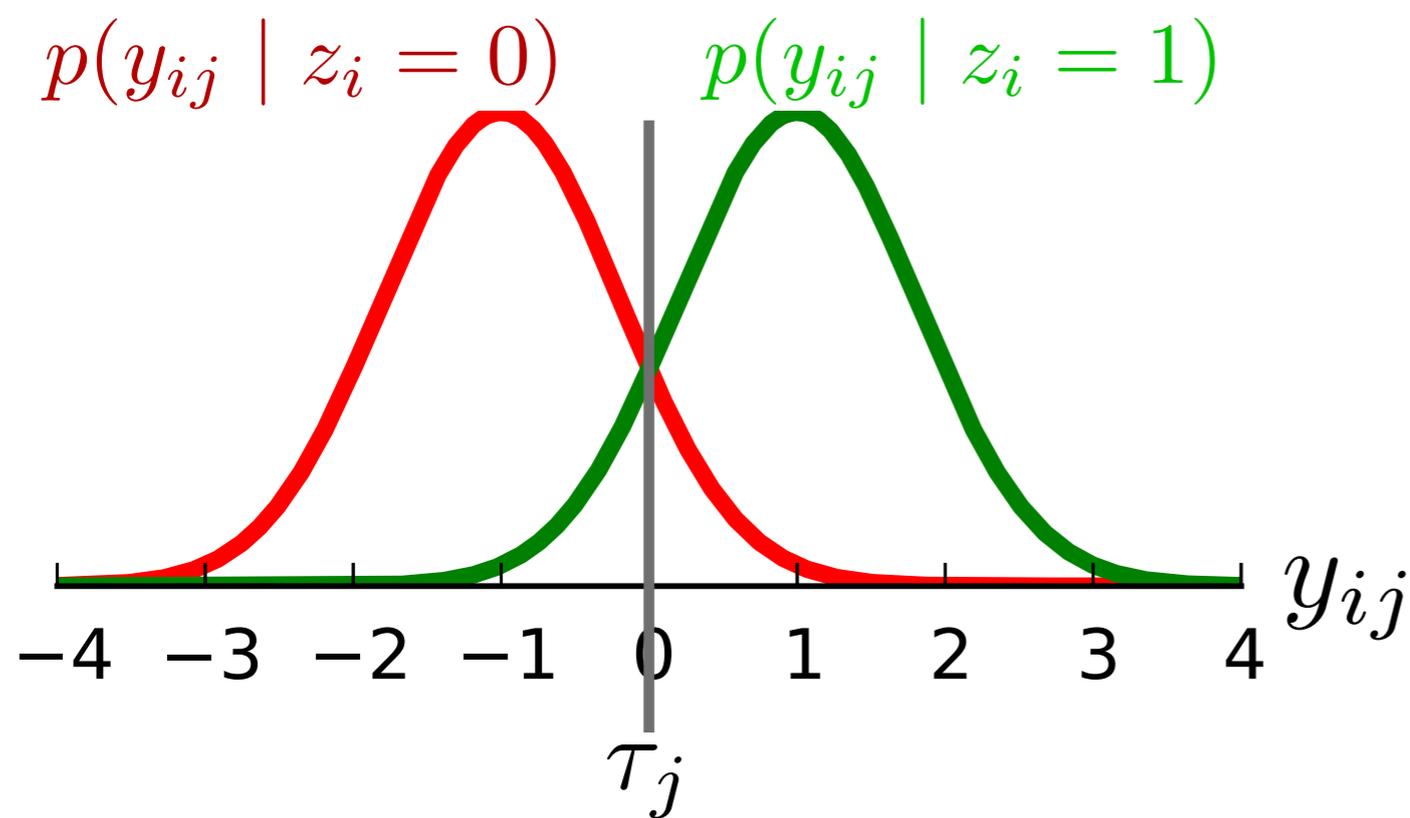
# Optimists and pessimists: $\tau_j$



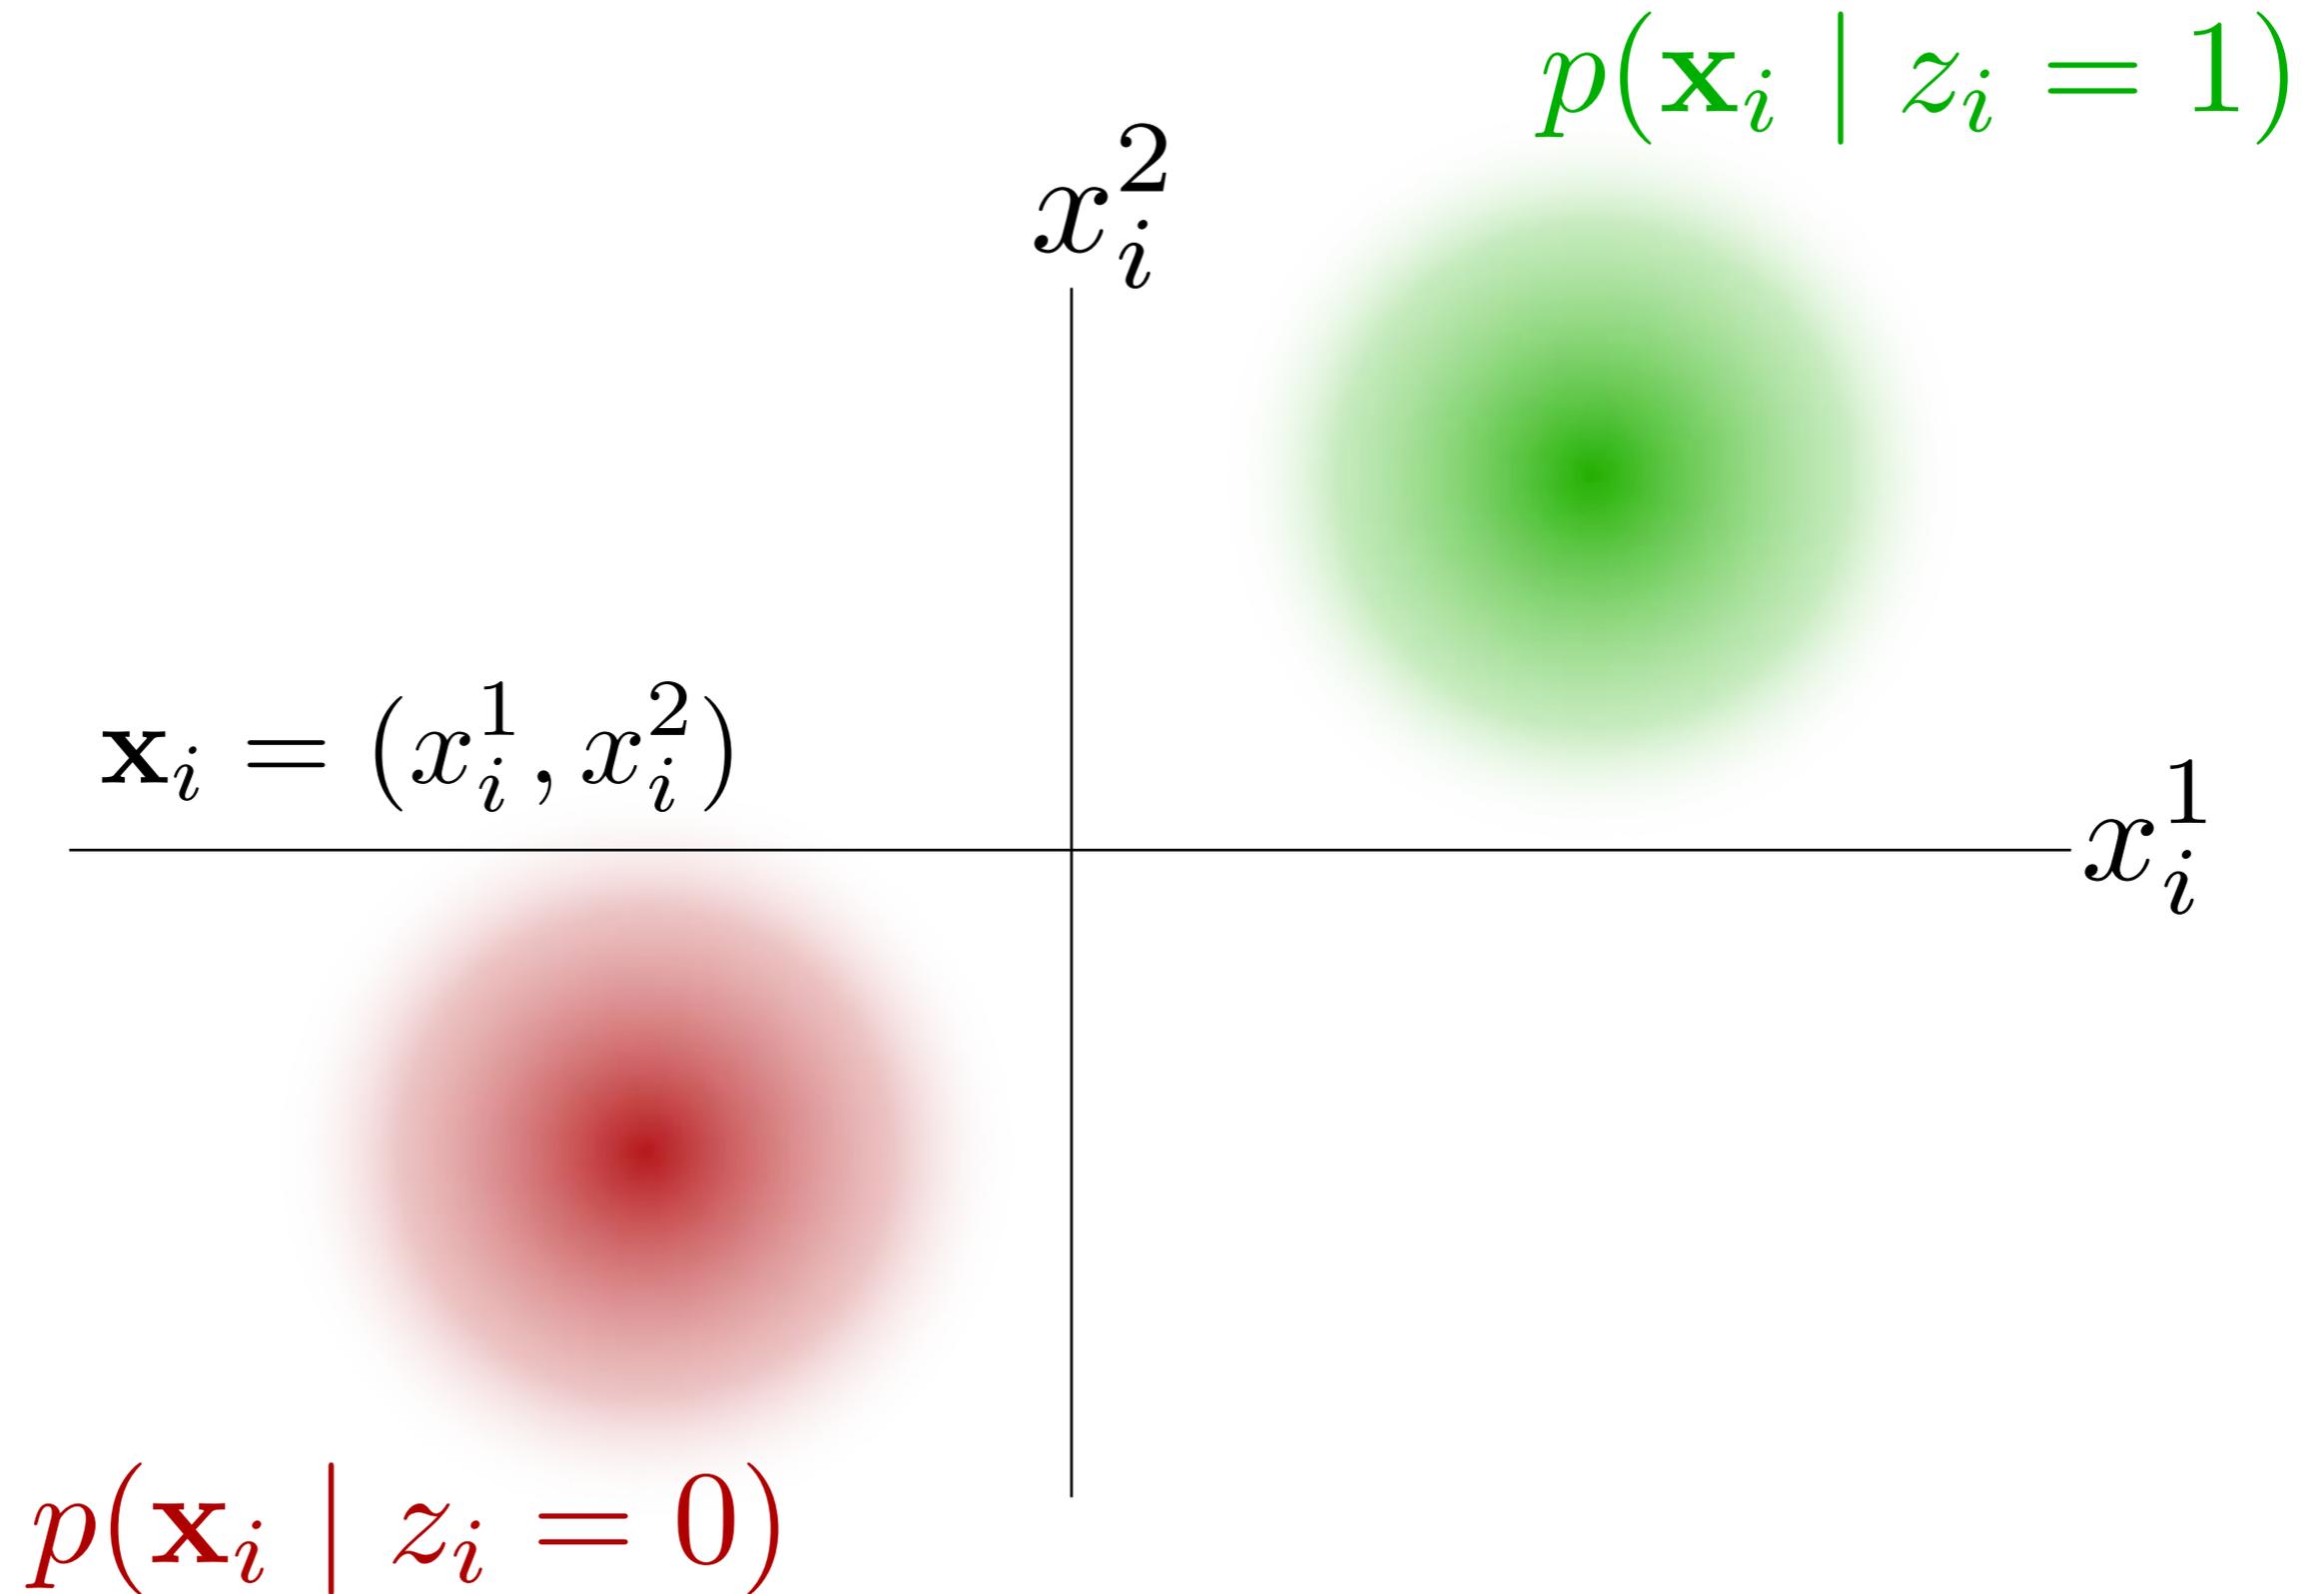
# Optimists and pessimists: $\tau_j$



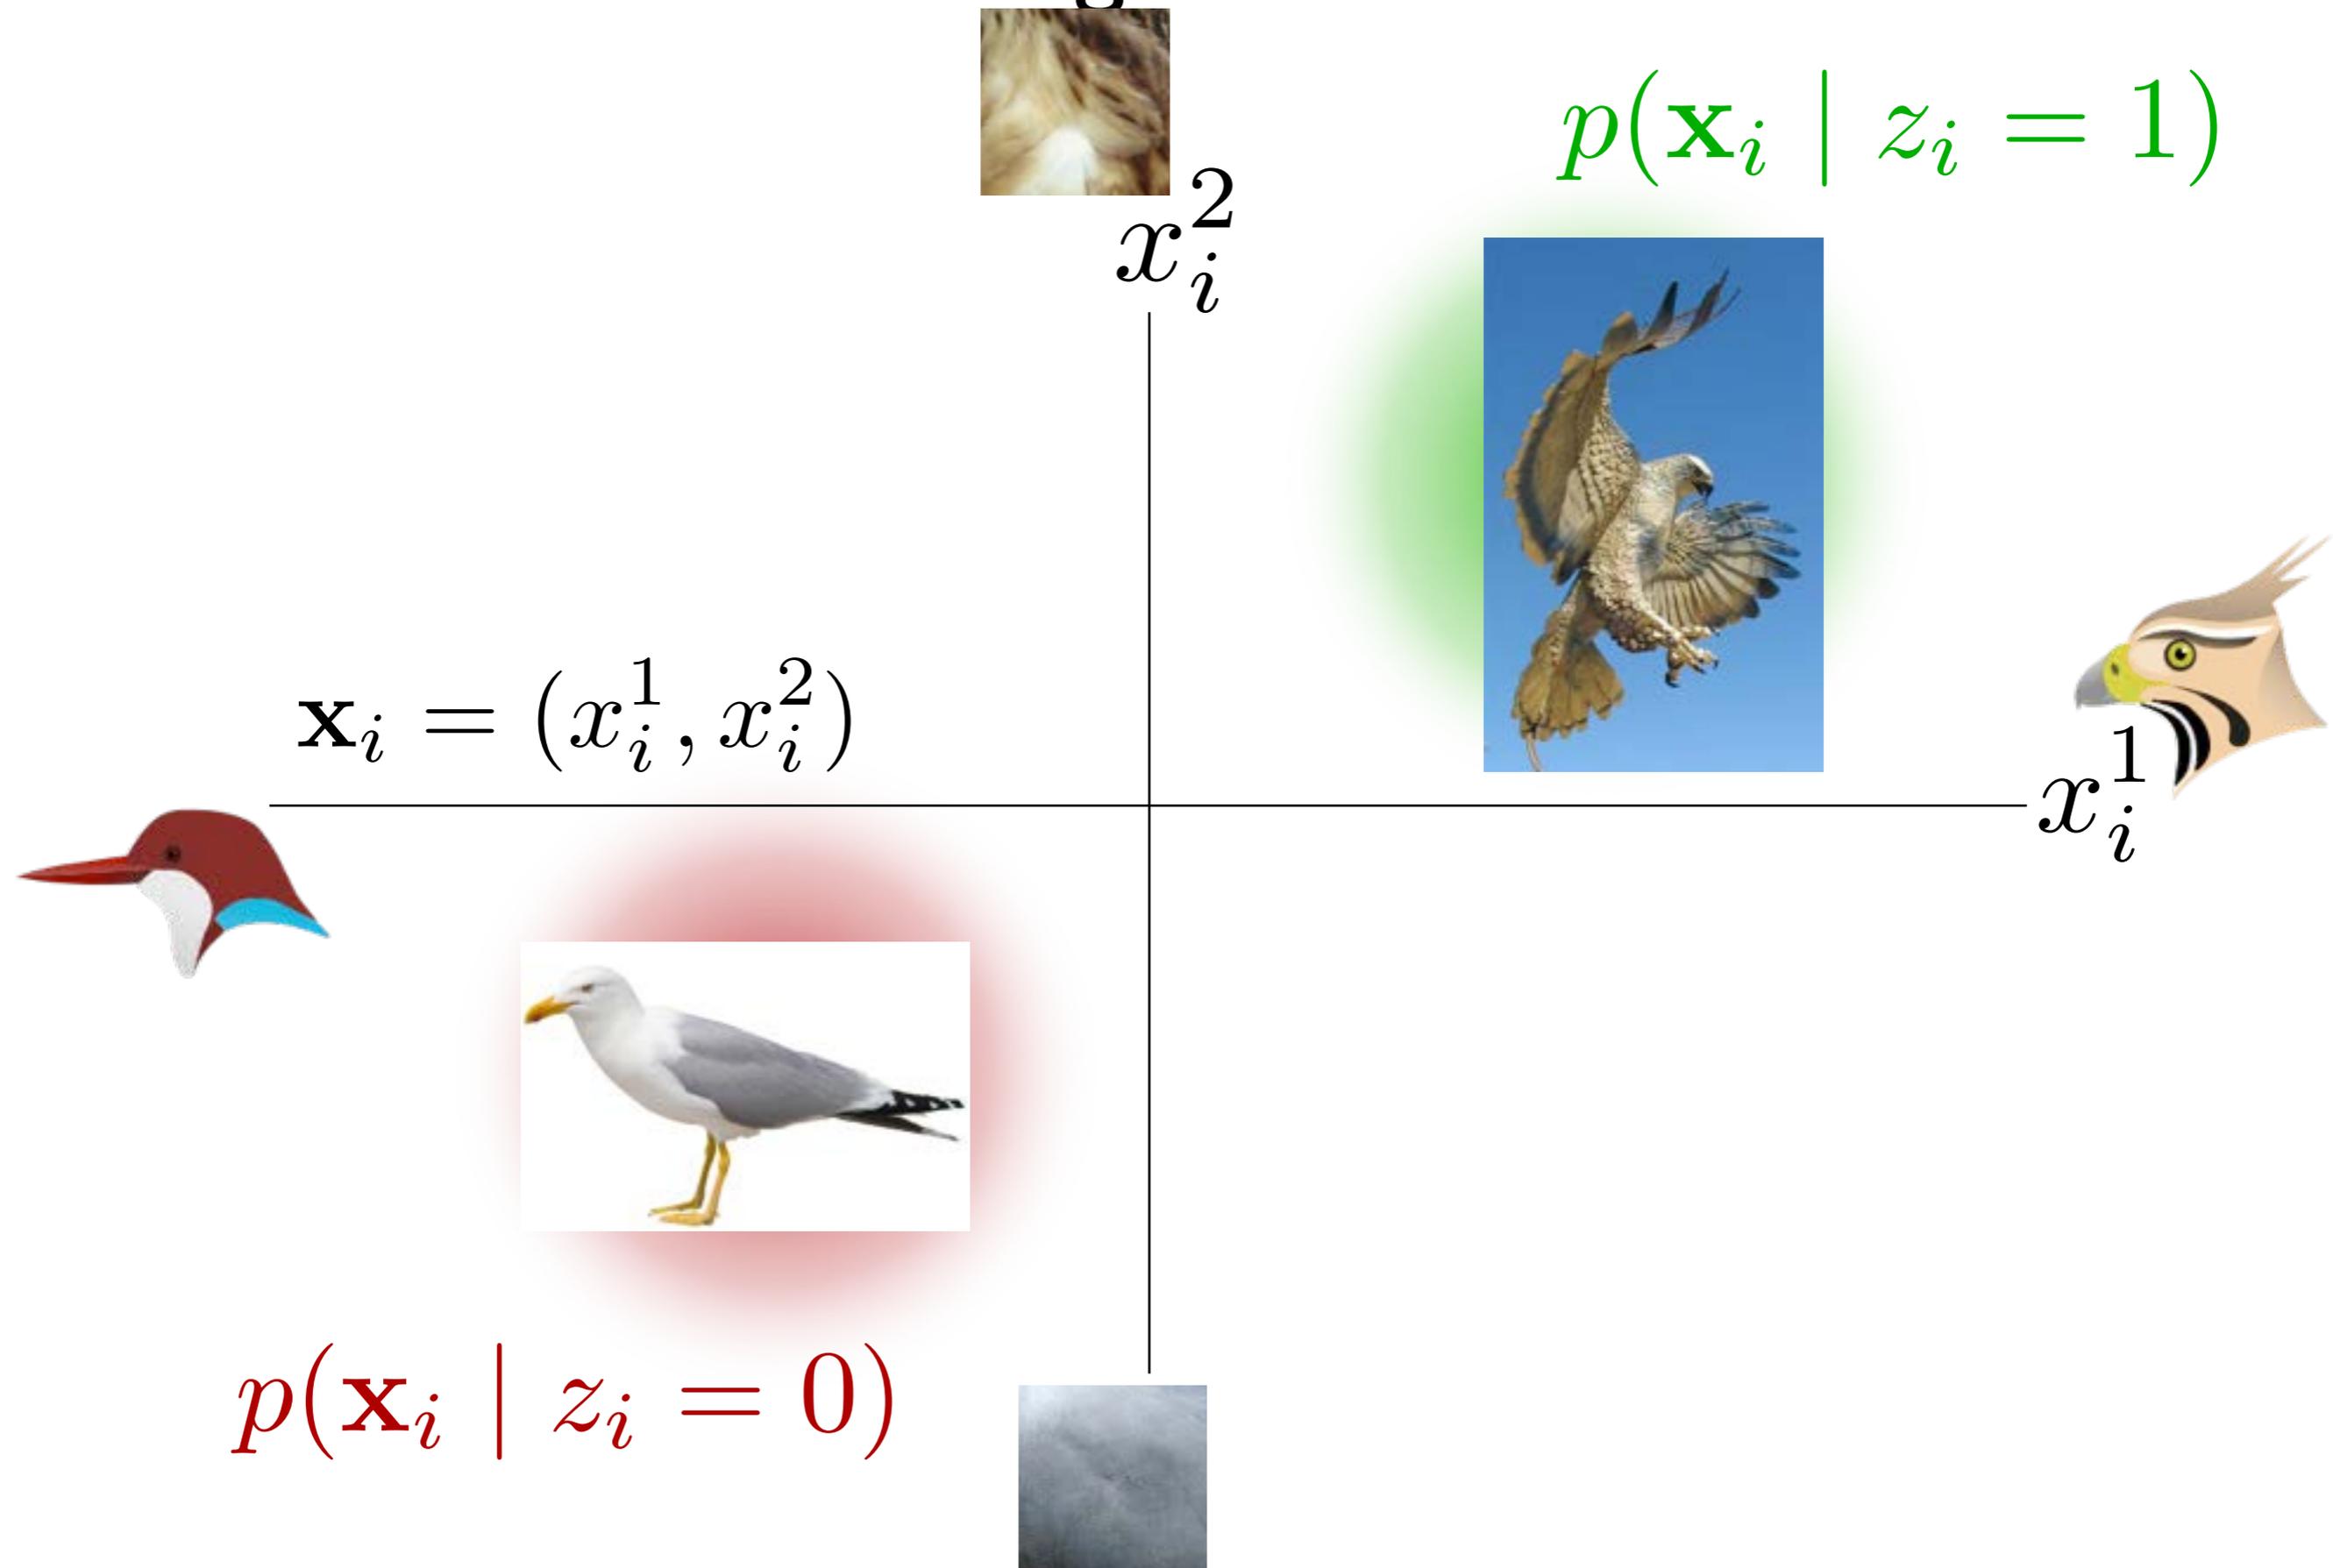
# Optimists and pessimists: $\tau_j$



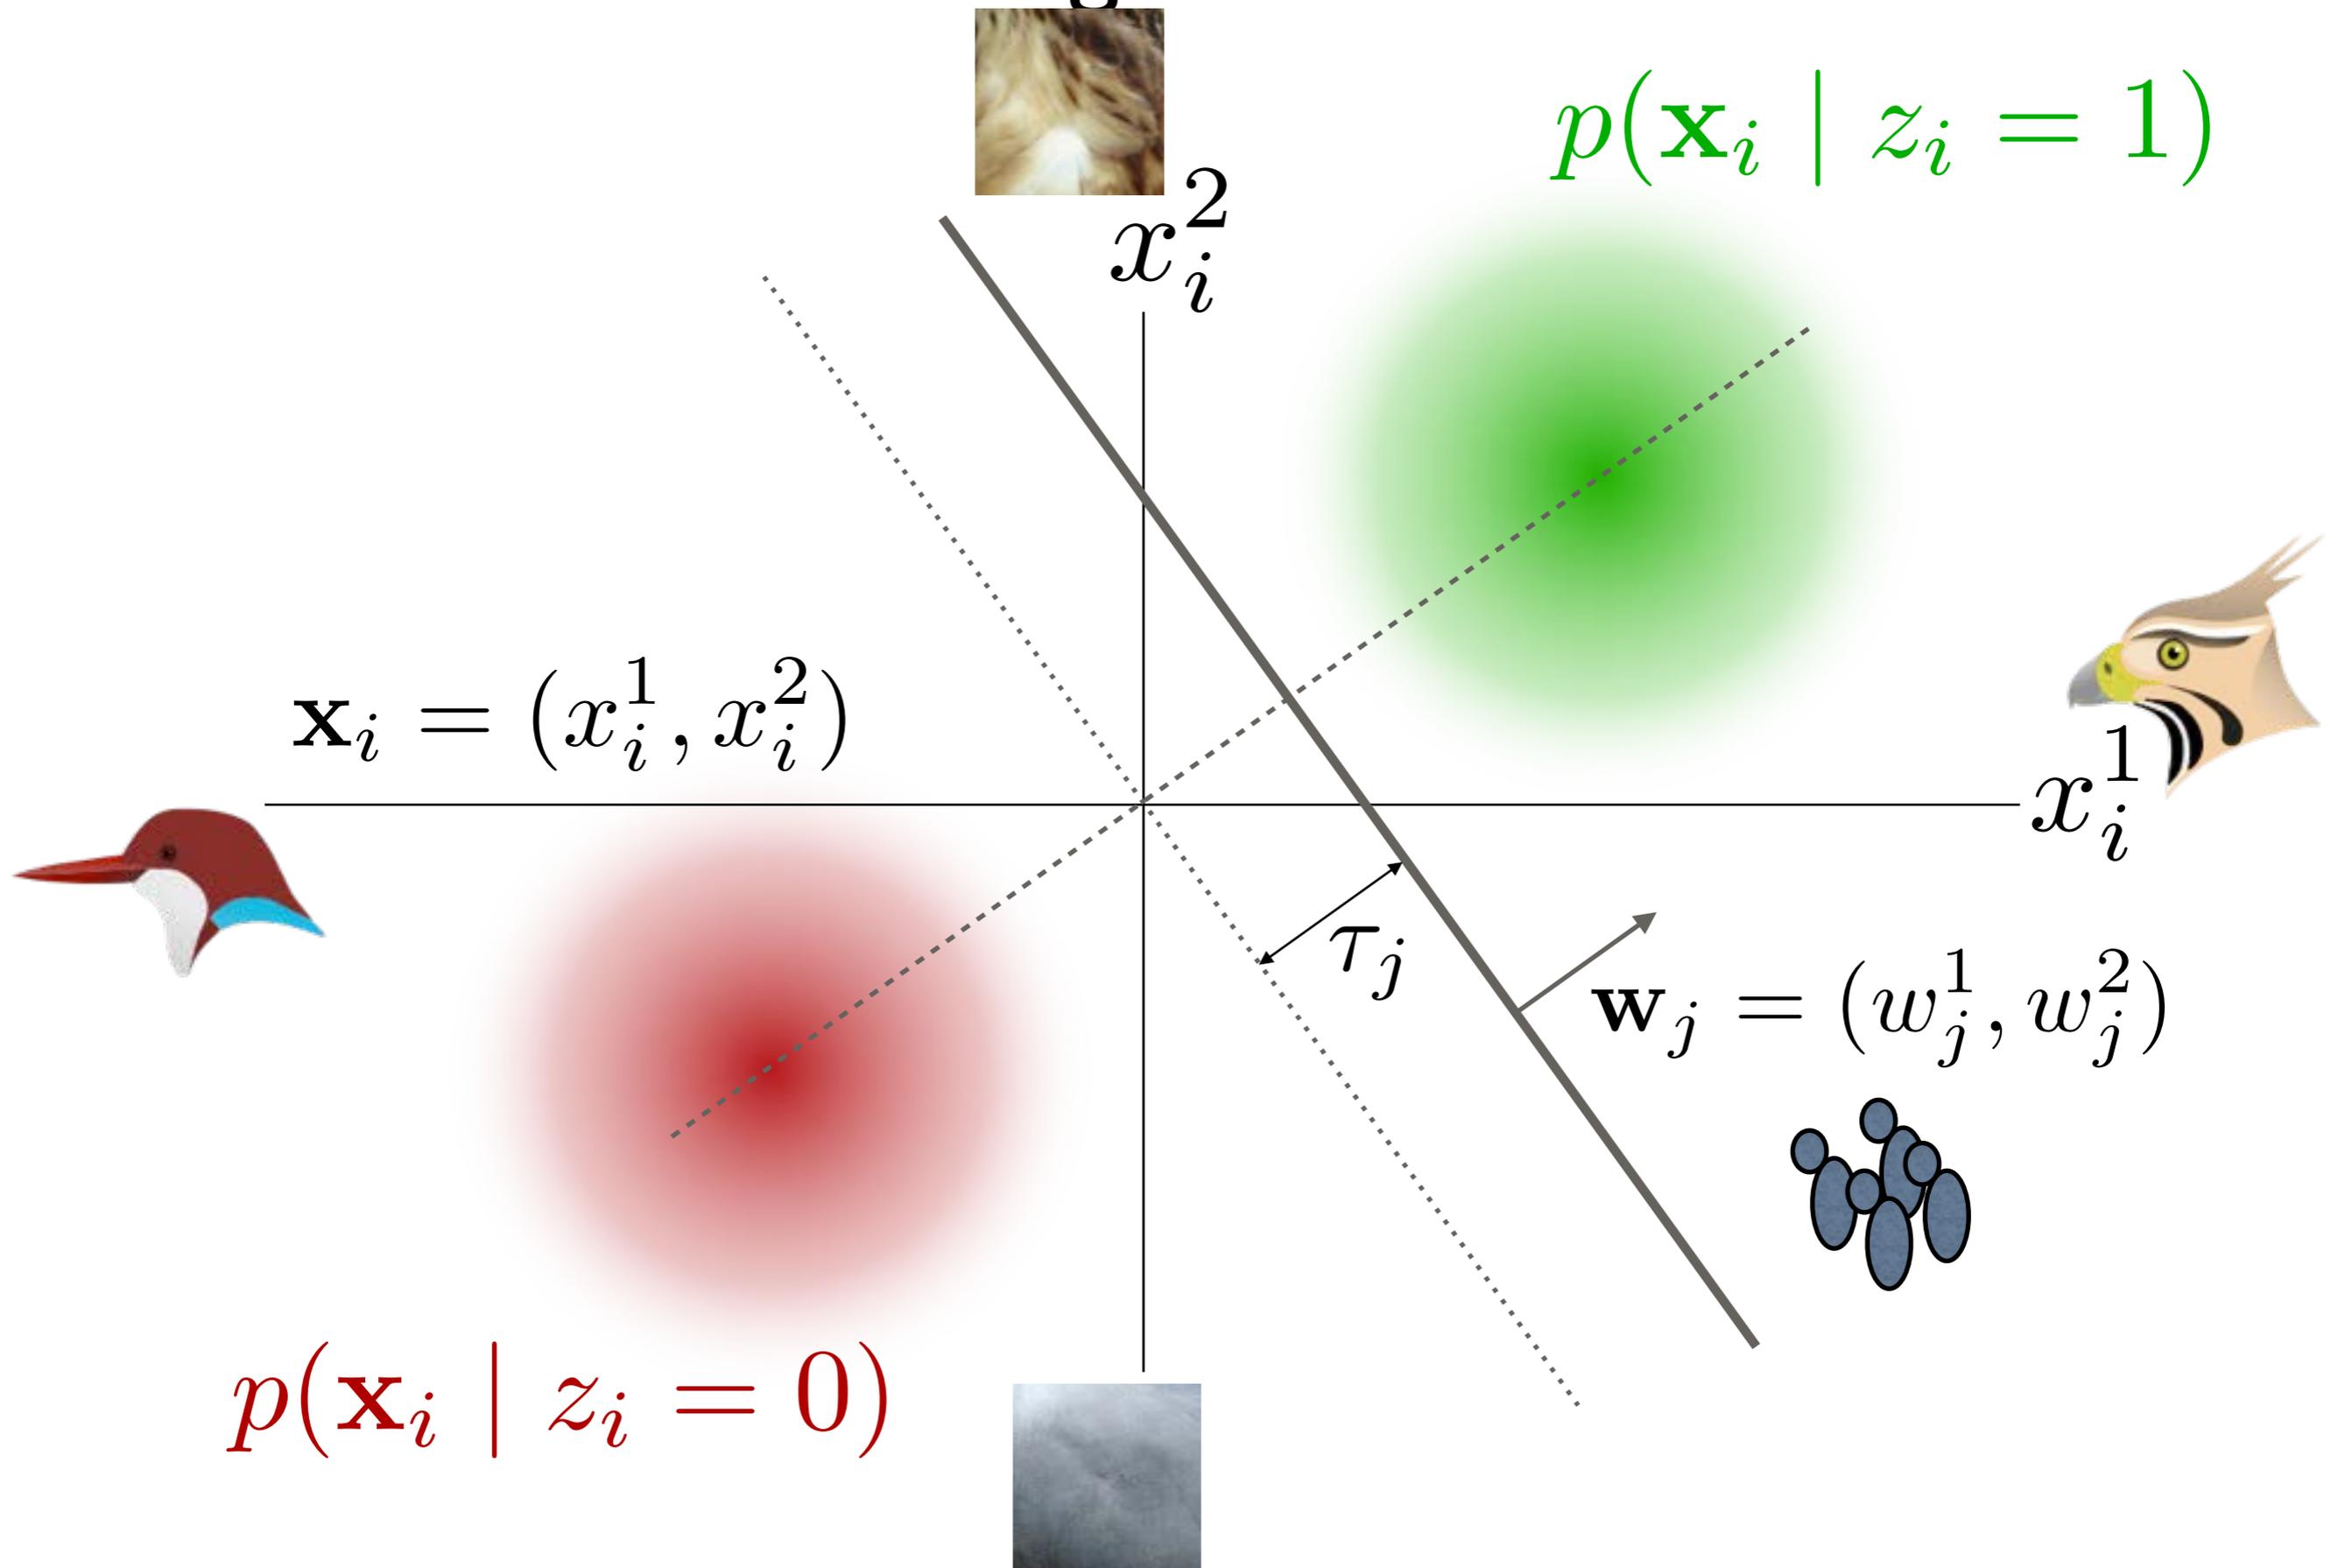
# Multidimensional signals and annotators



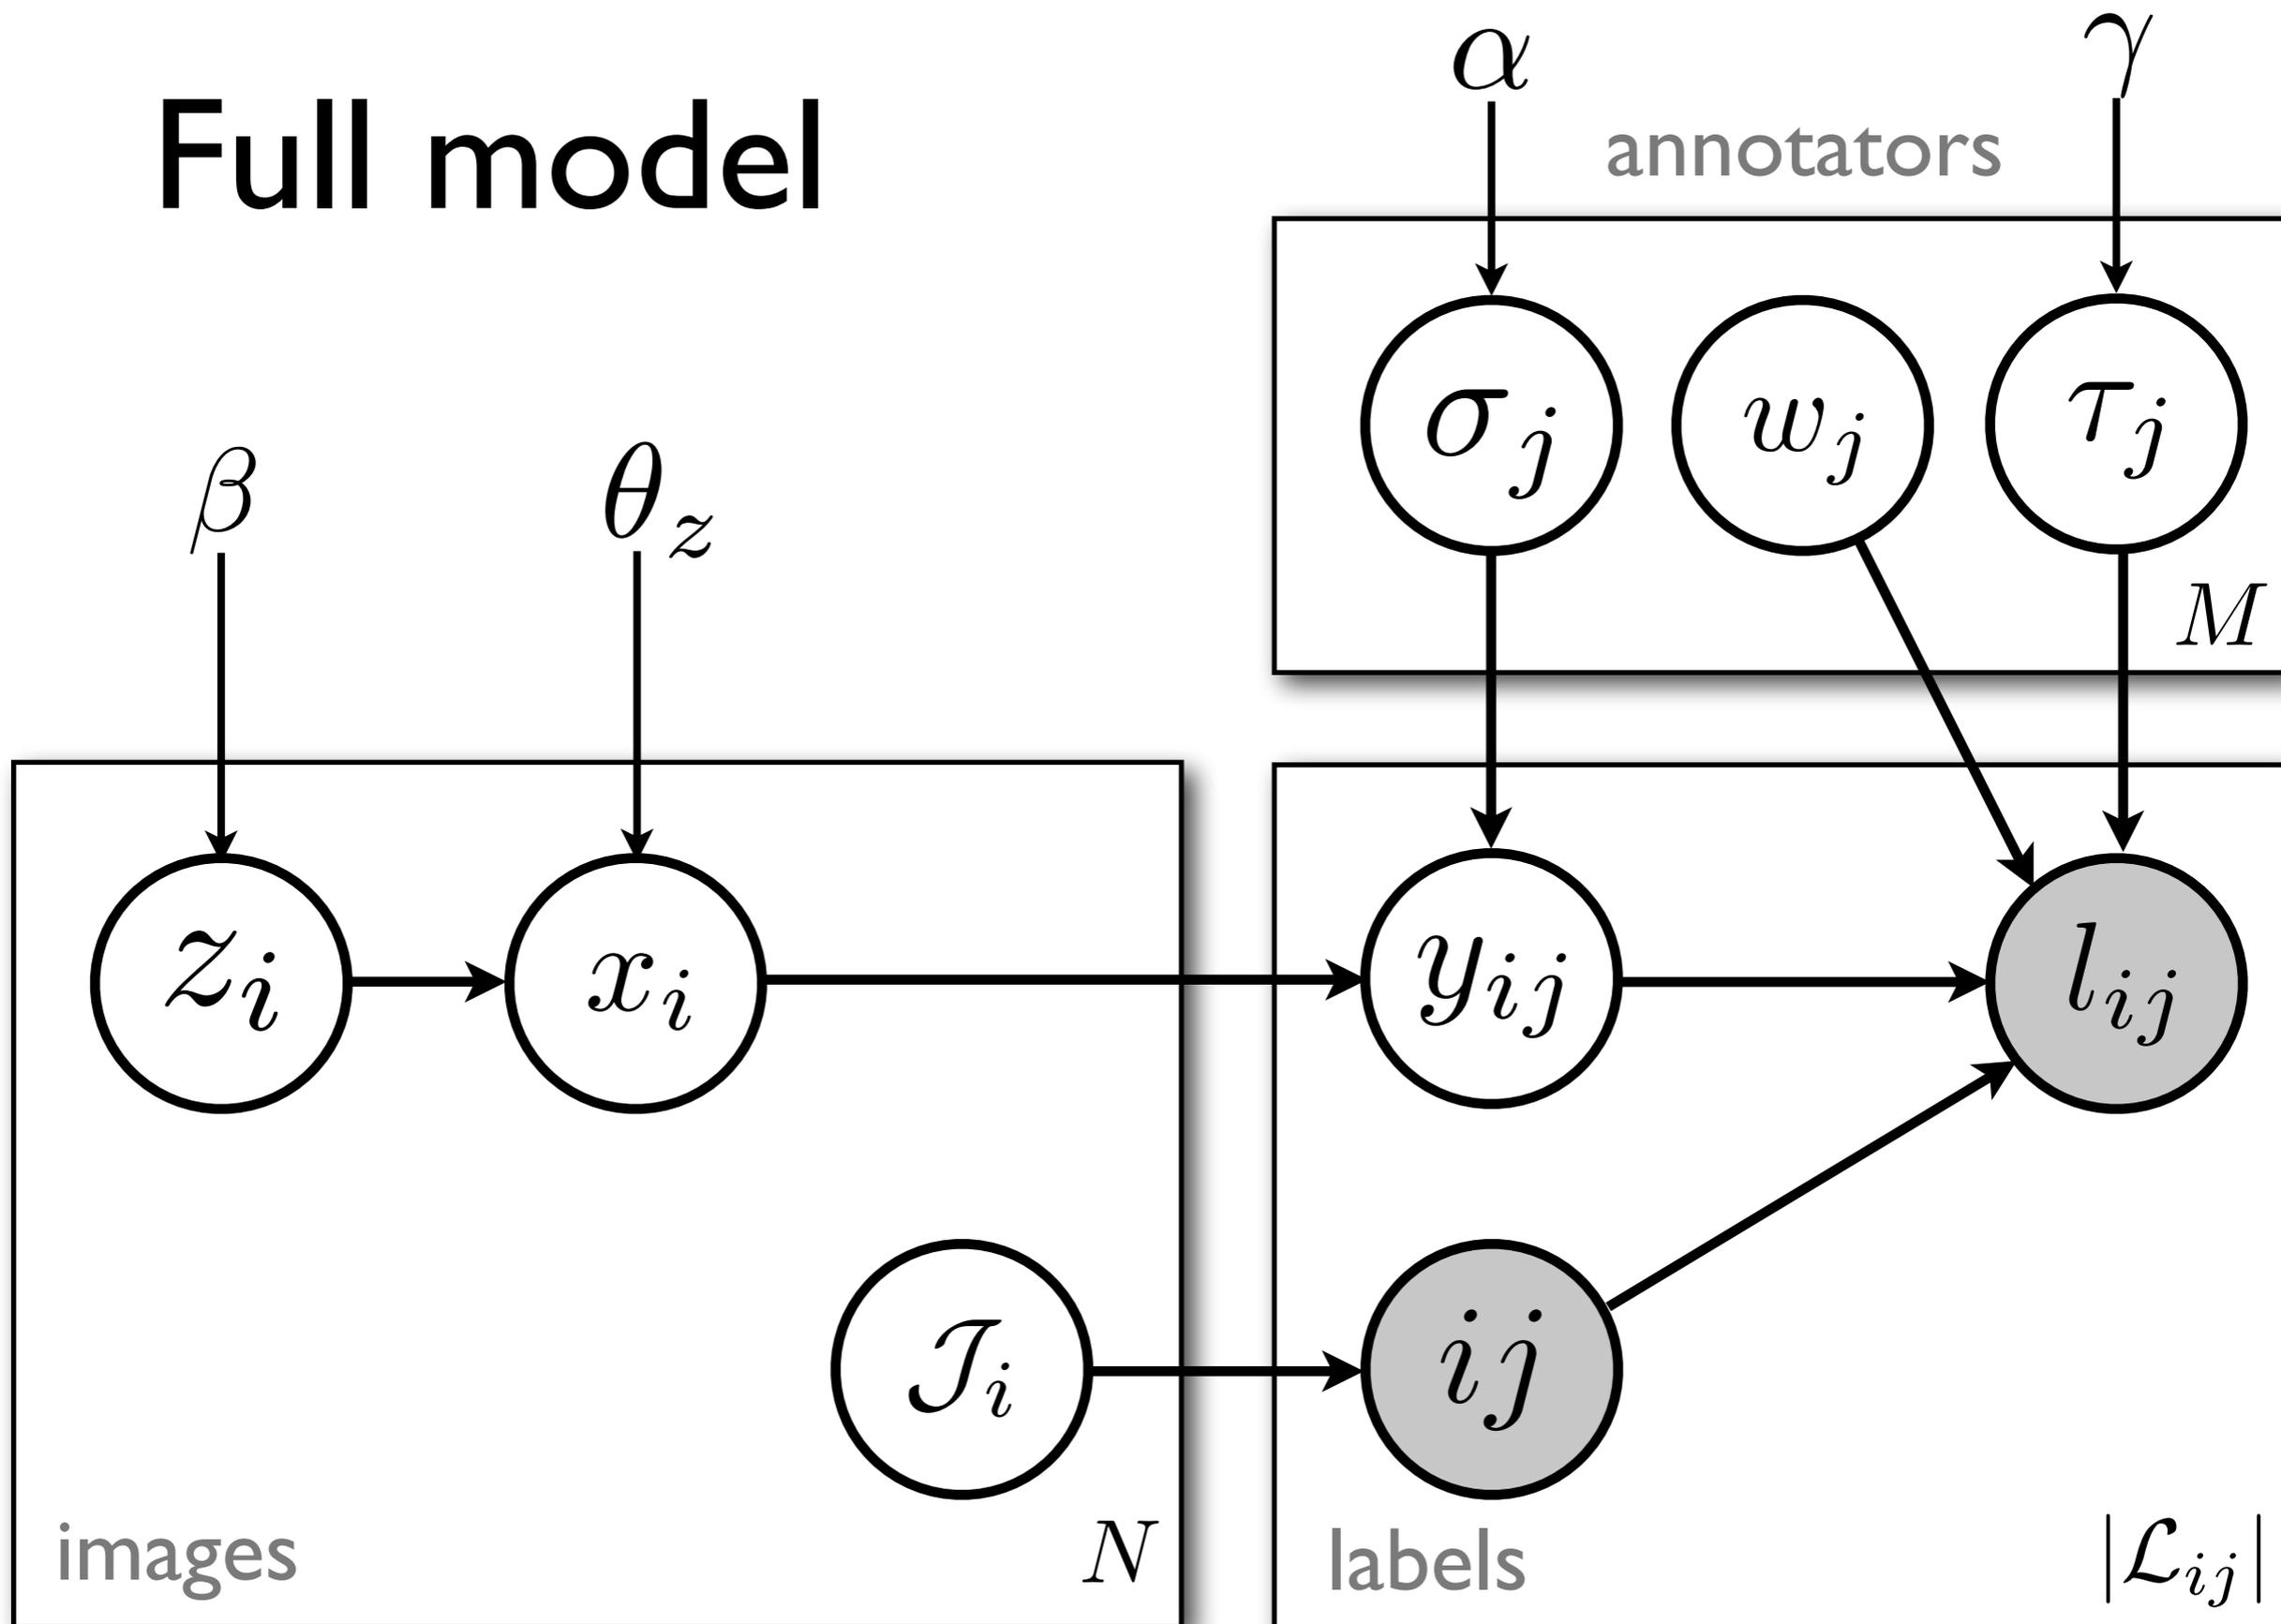
# Multidimensional signals and annotators



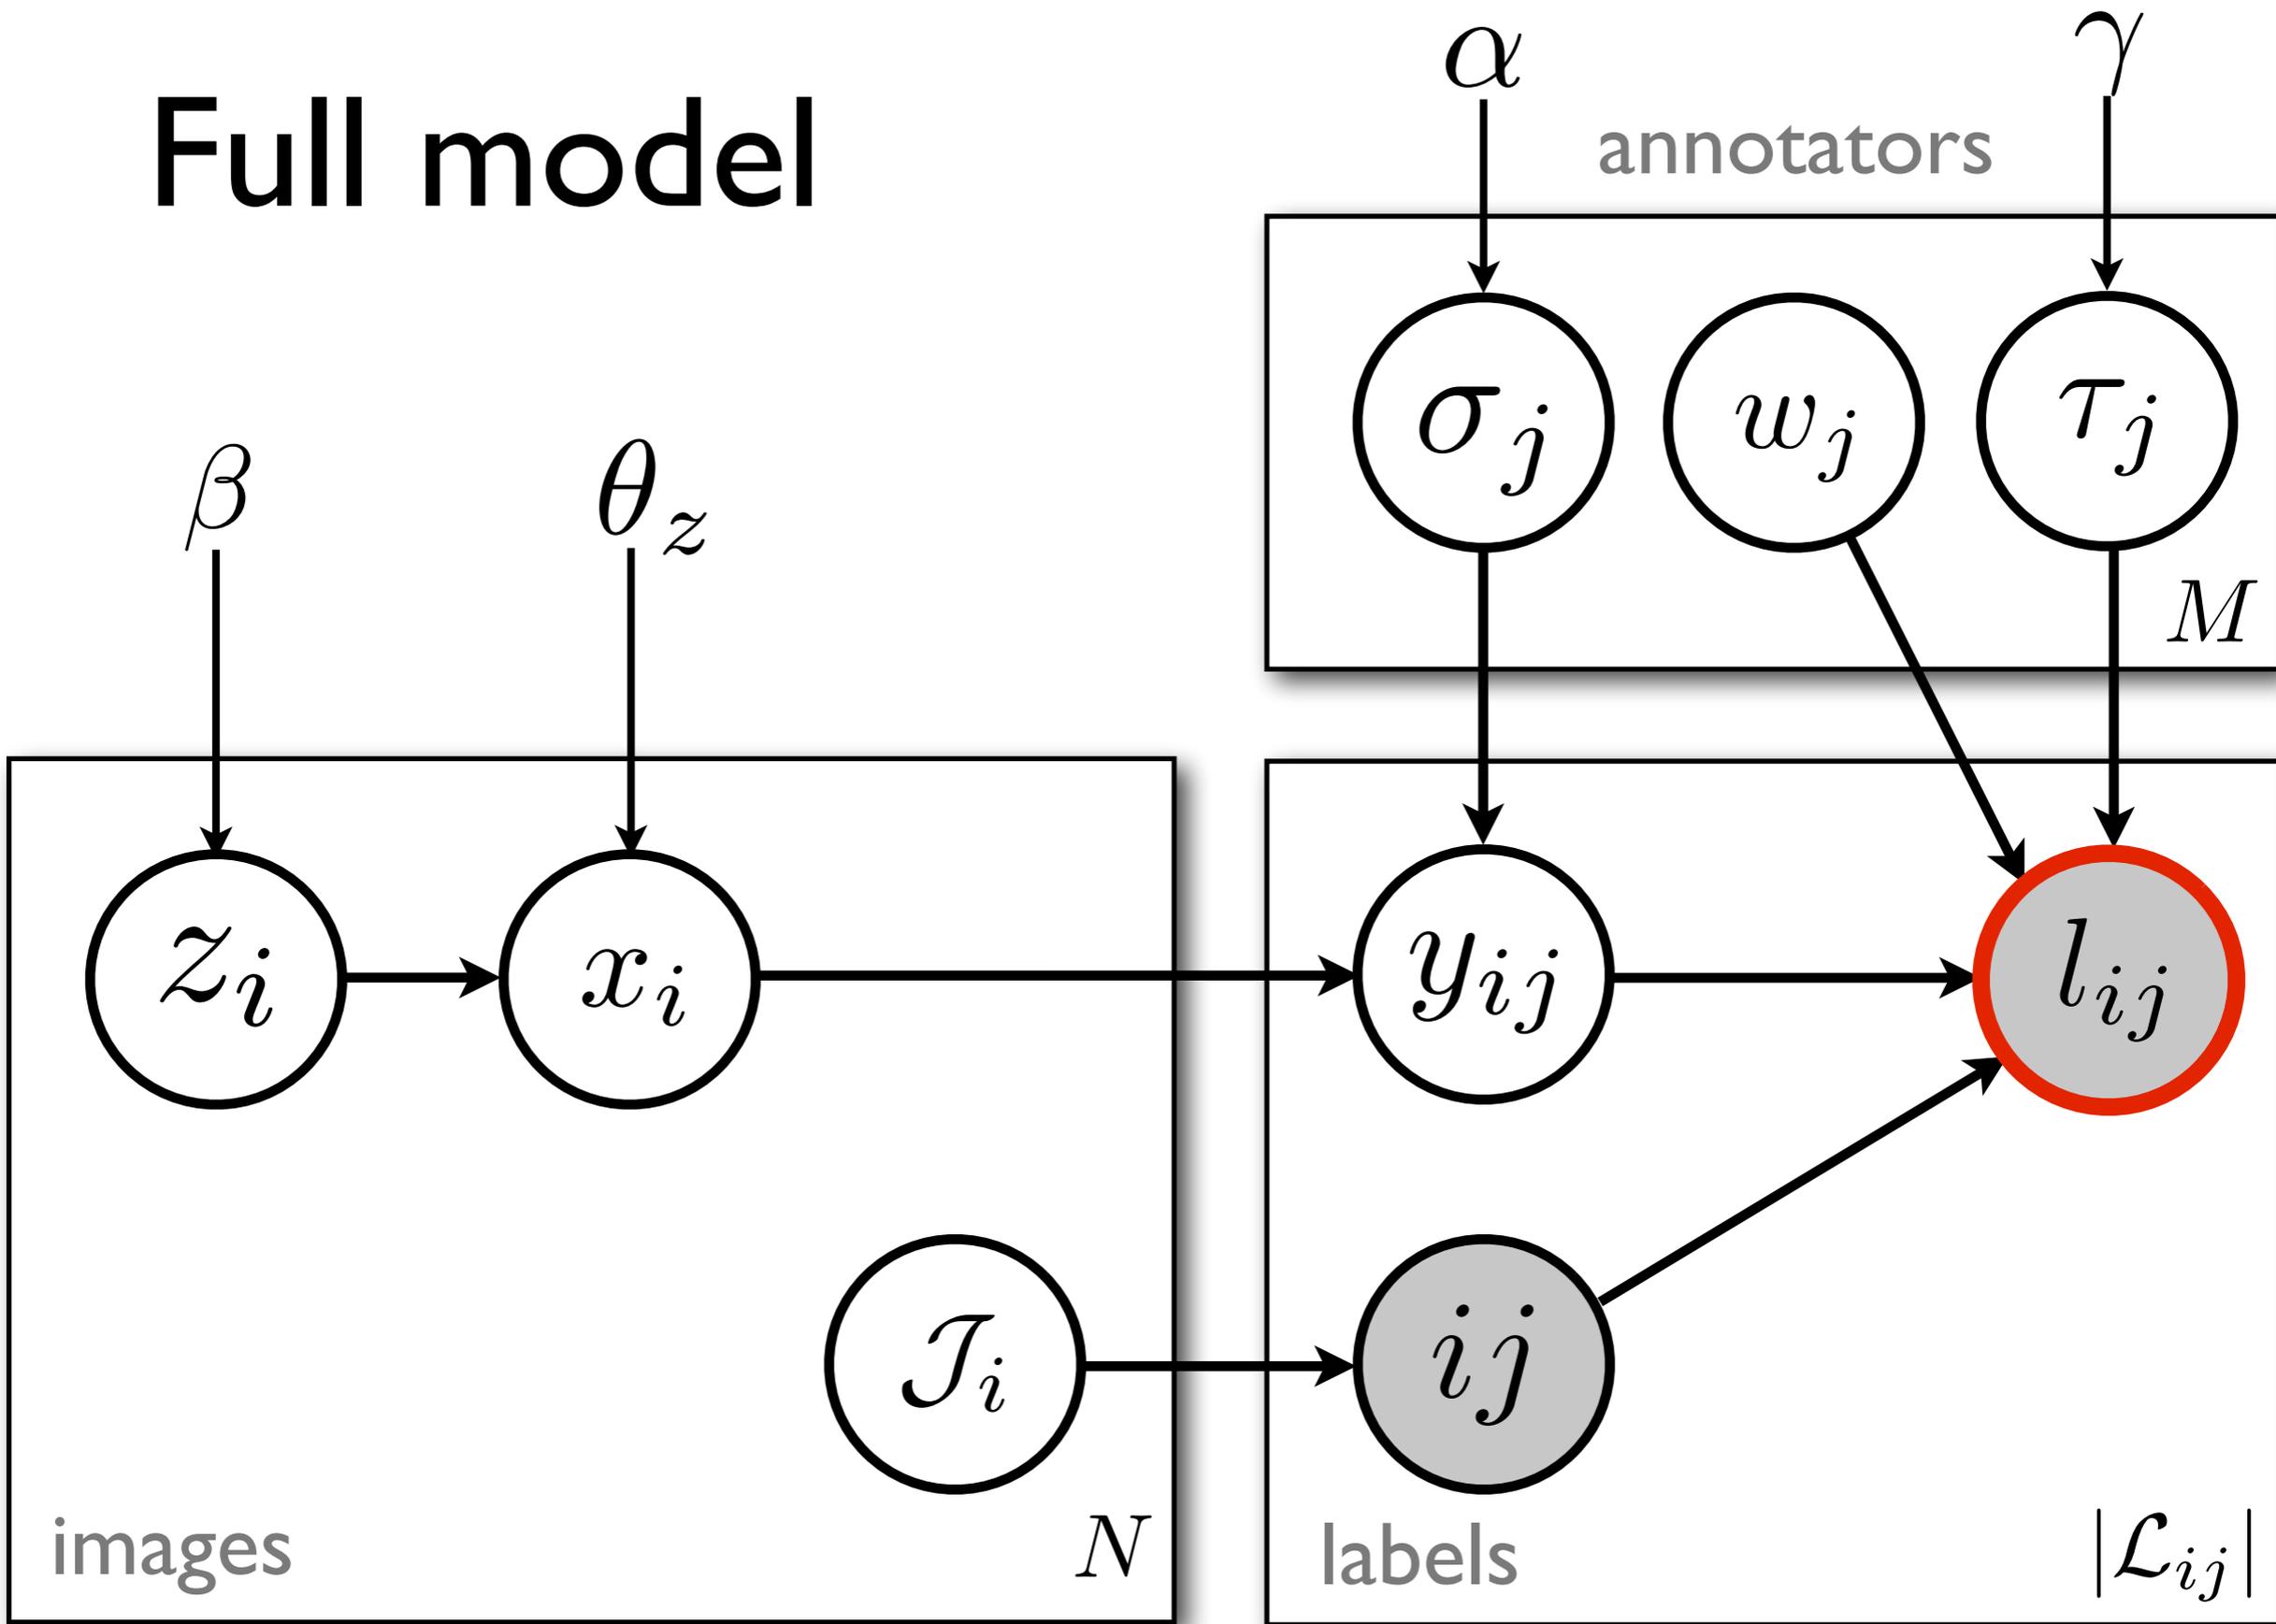
# Multidimensional signals and annotators



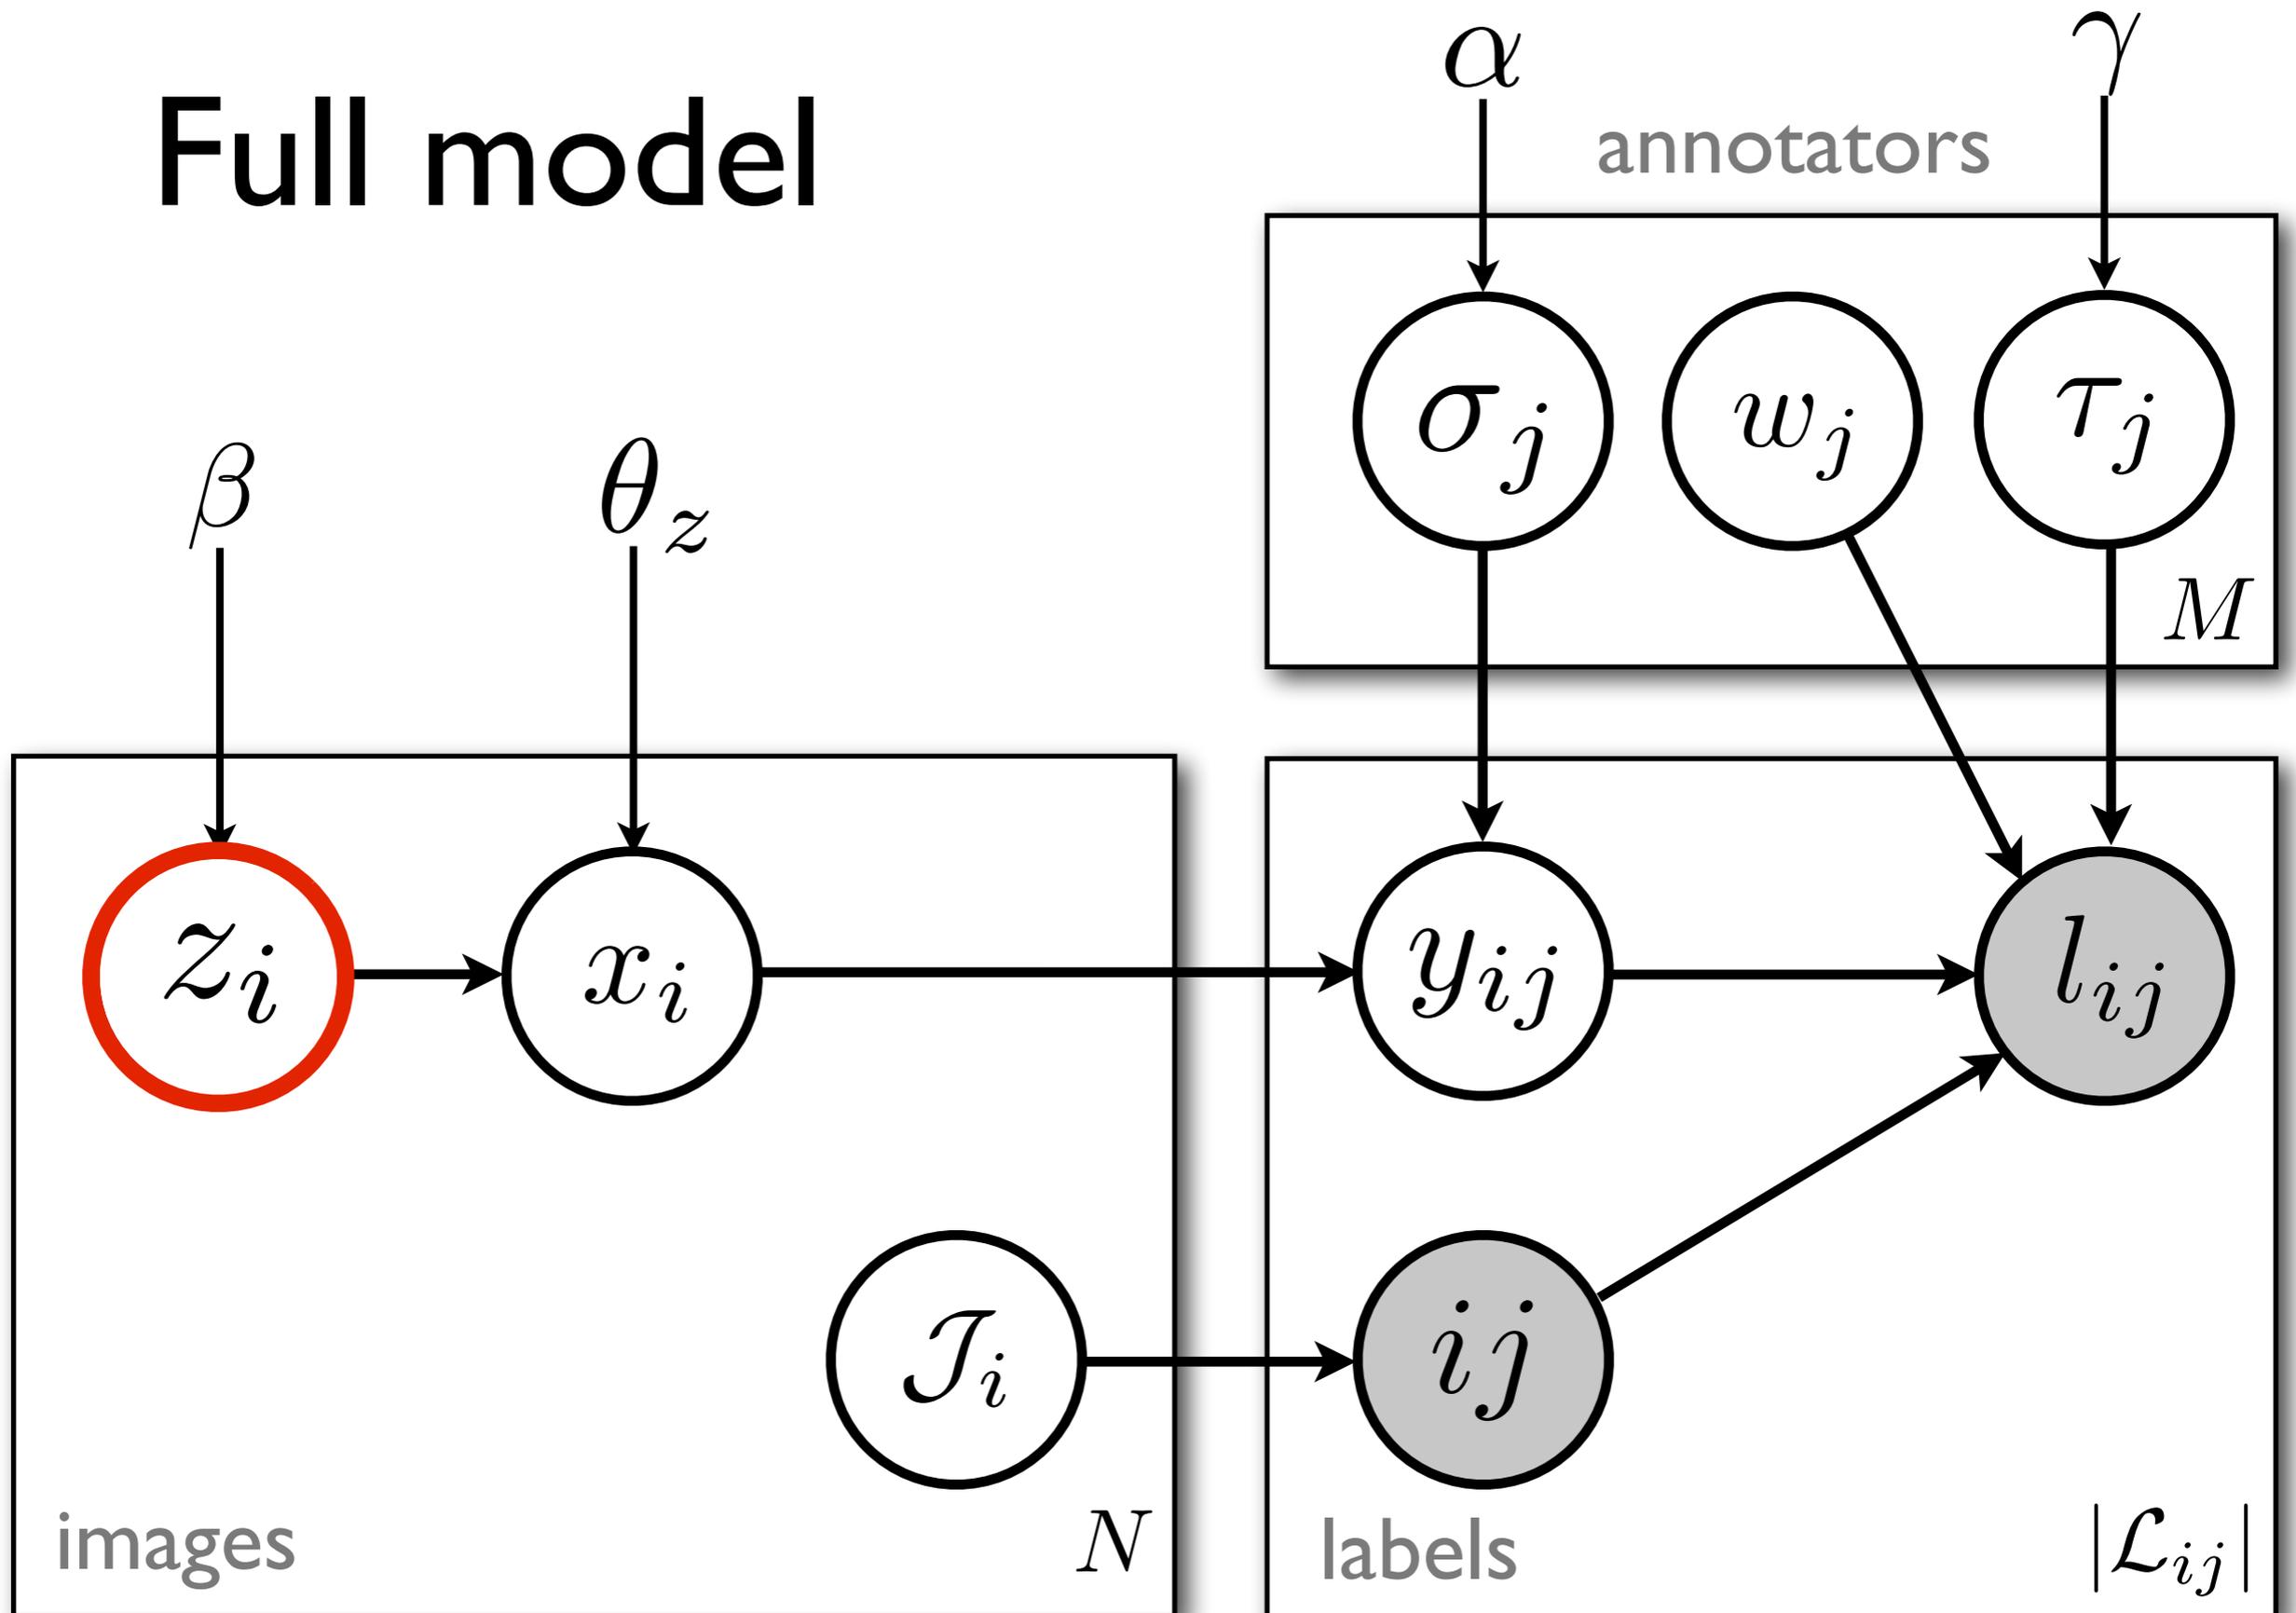
# Full model



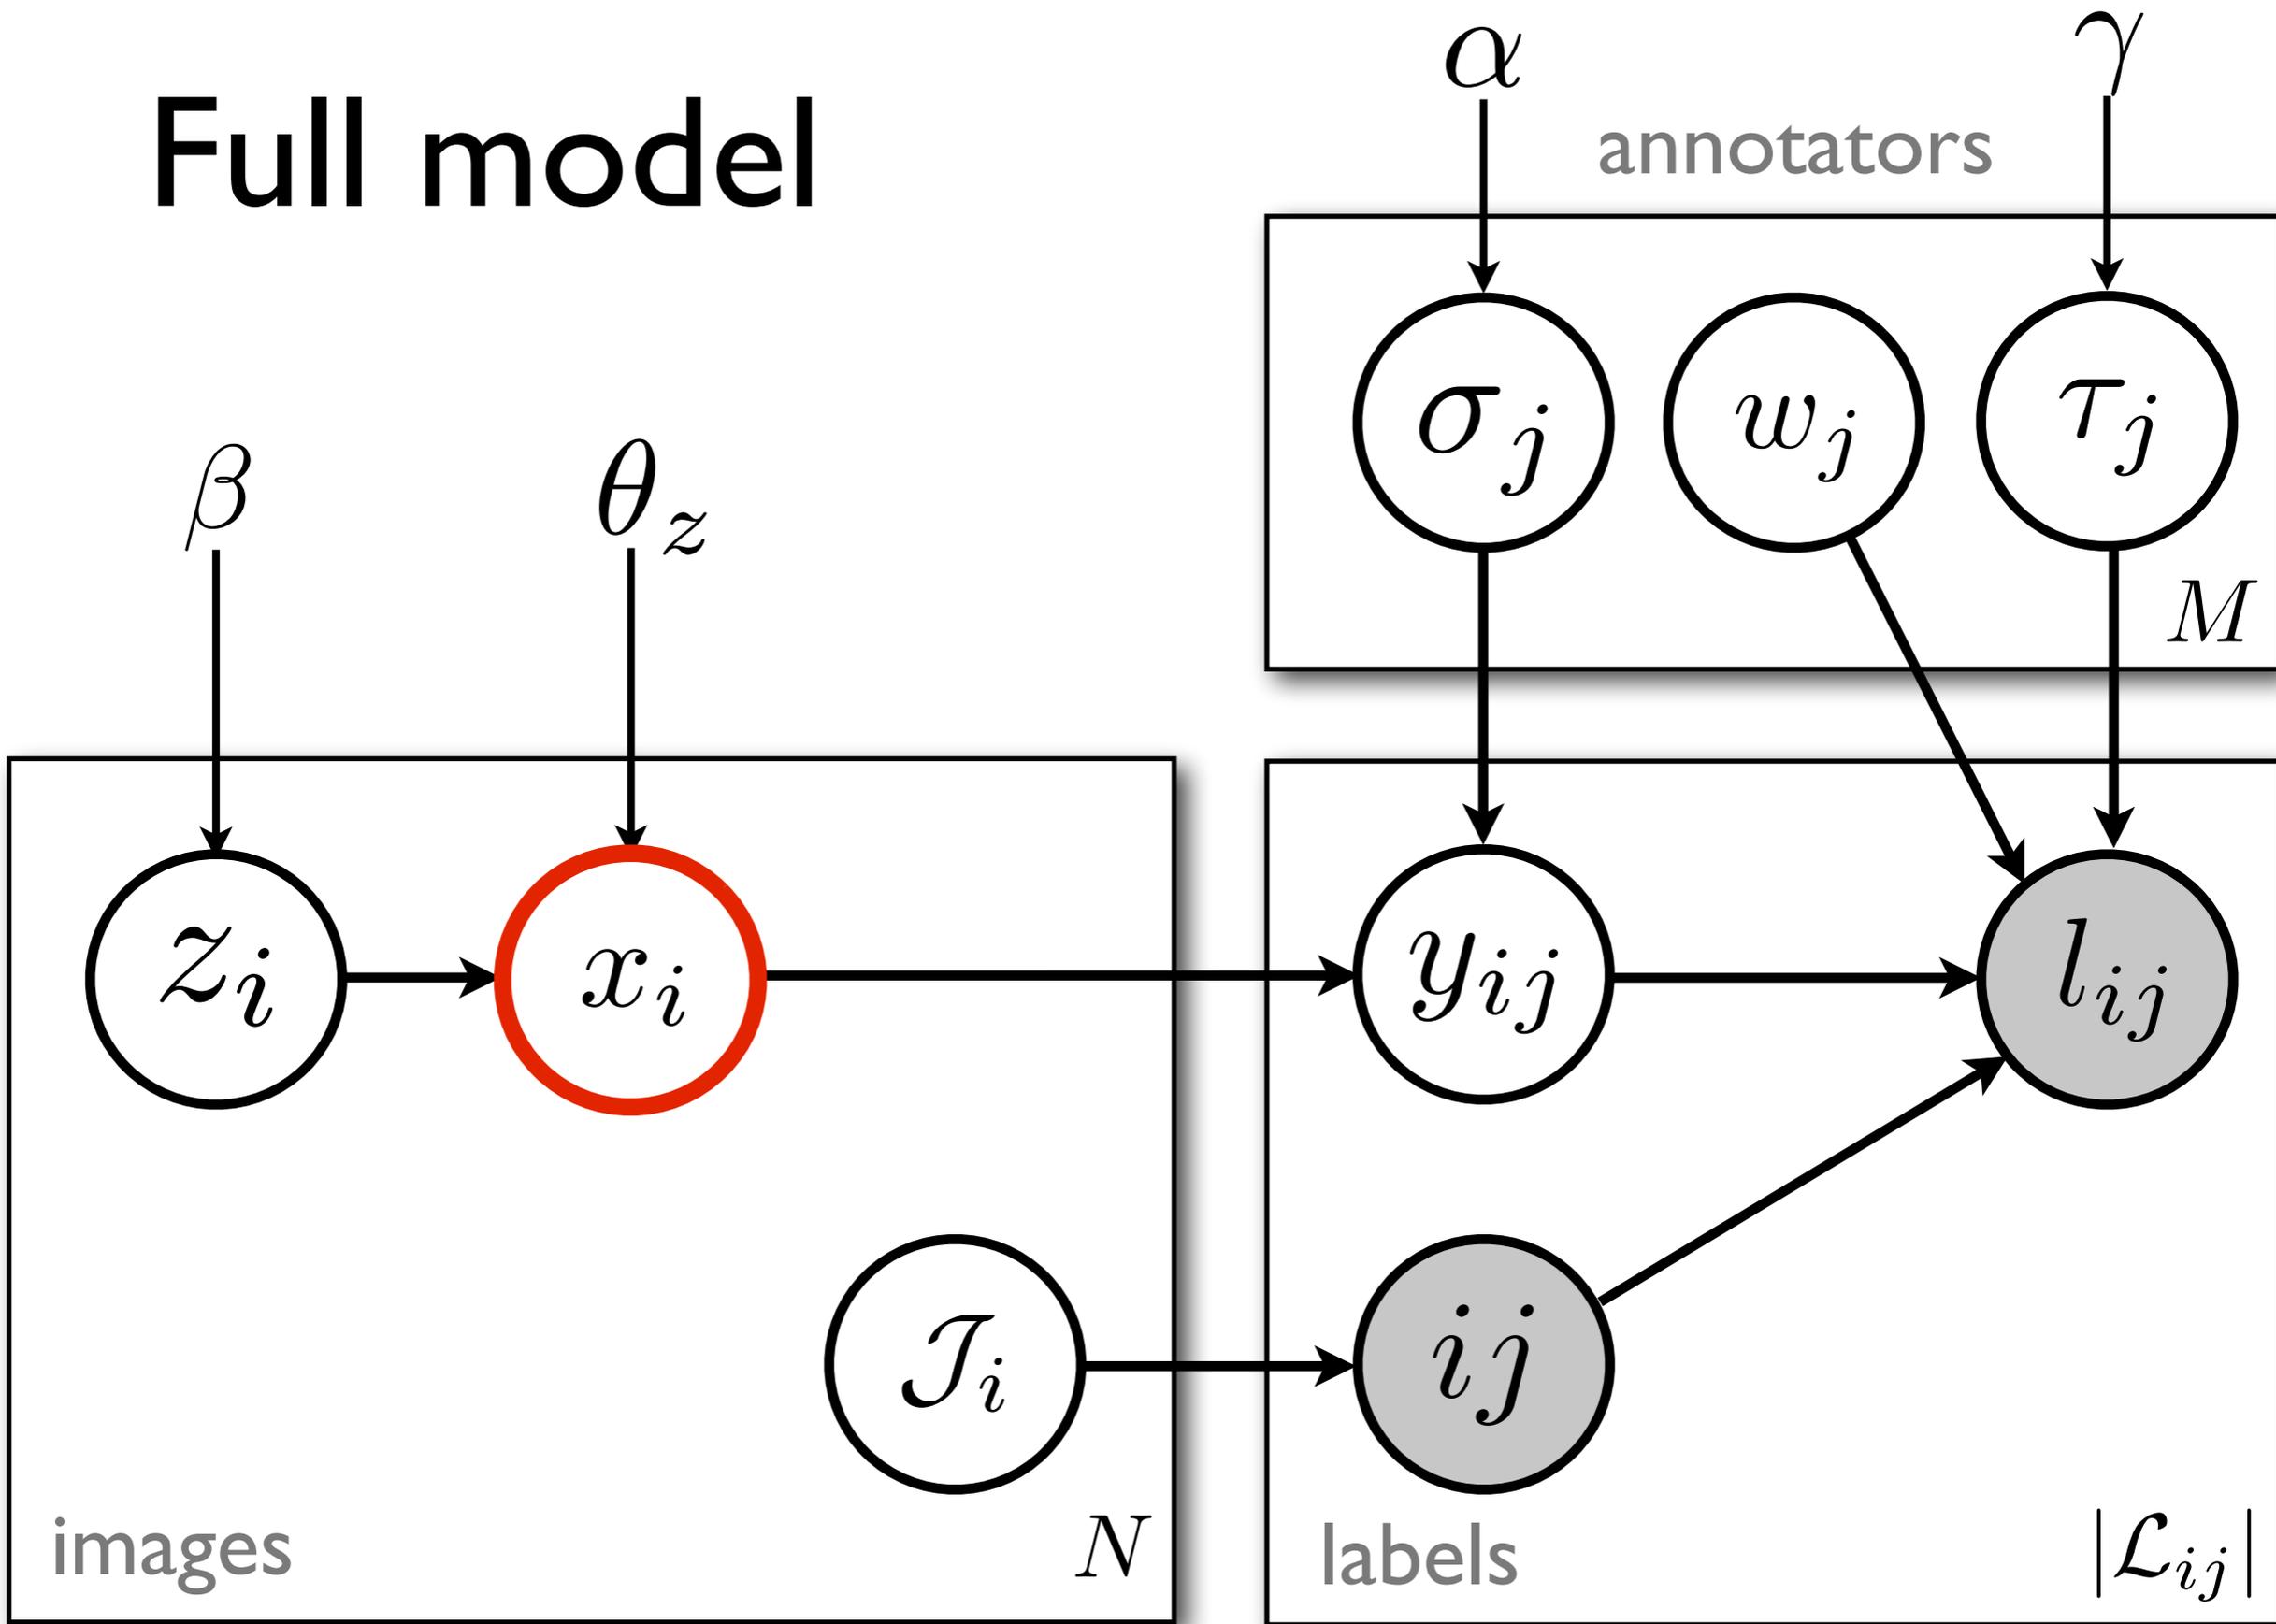
# Full model



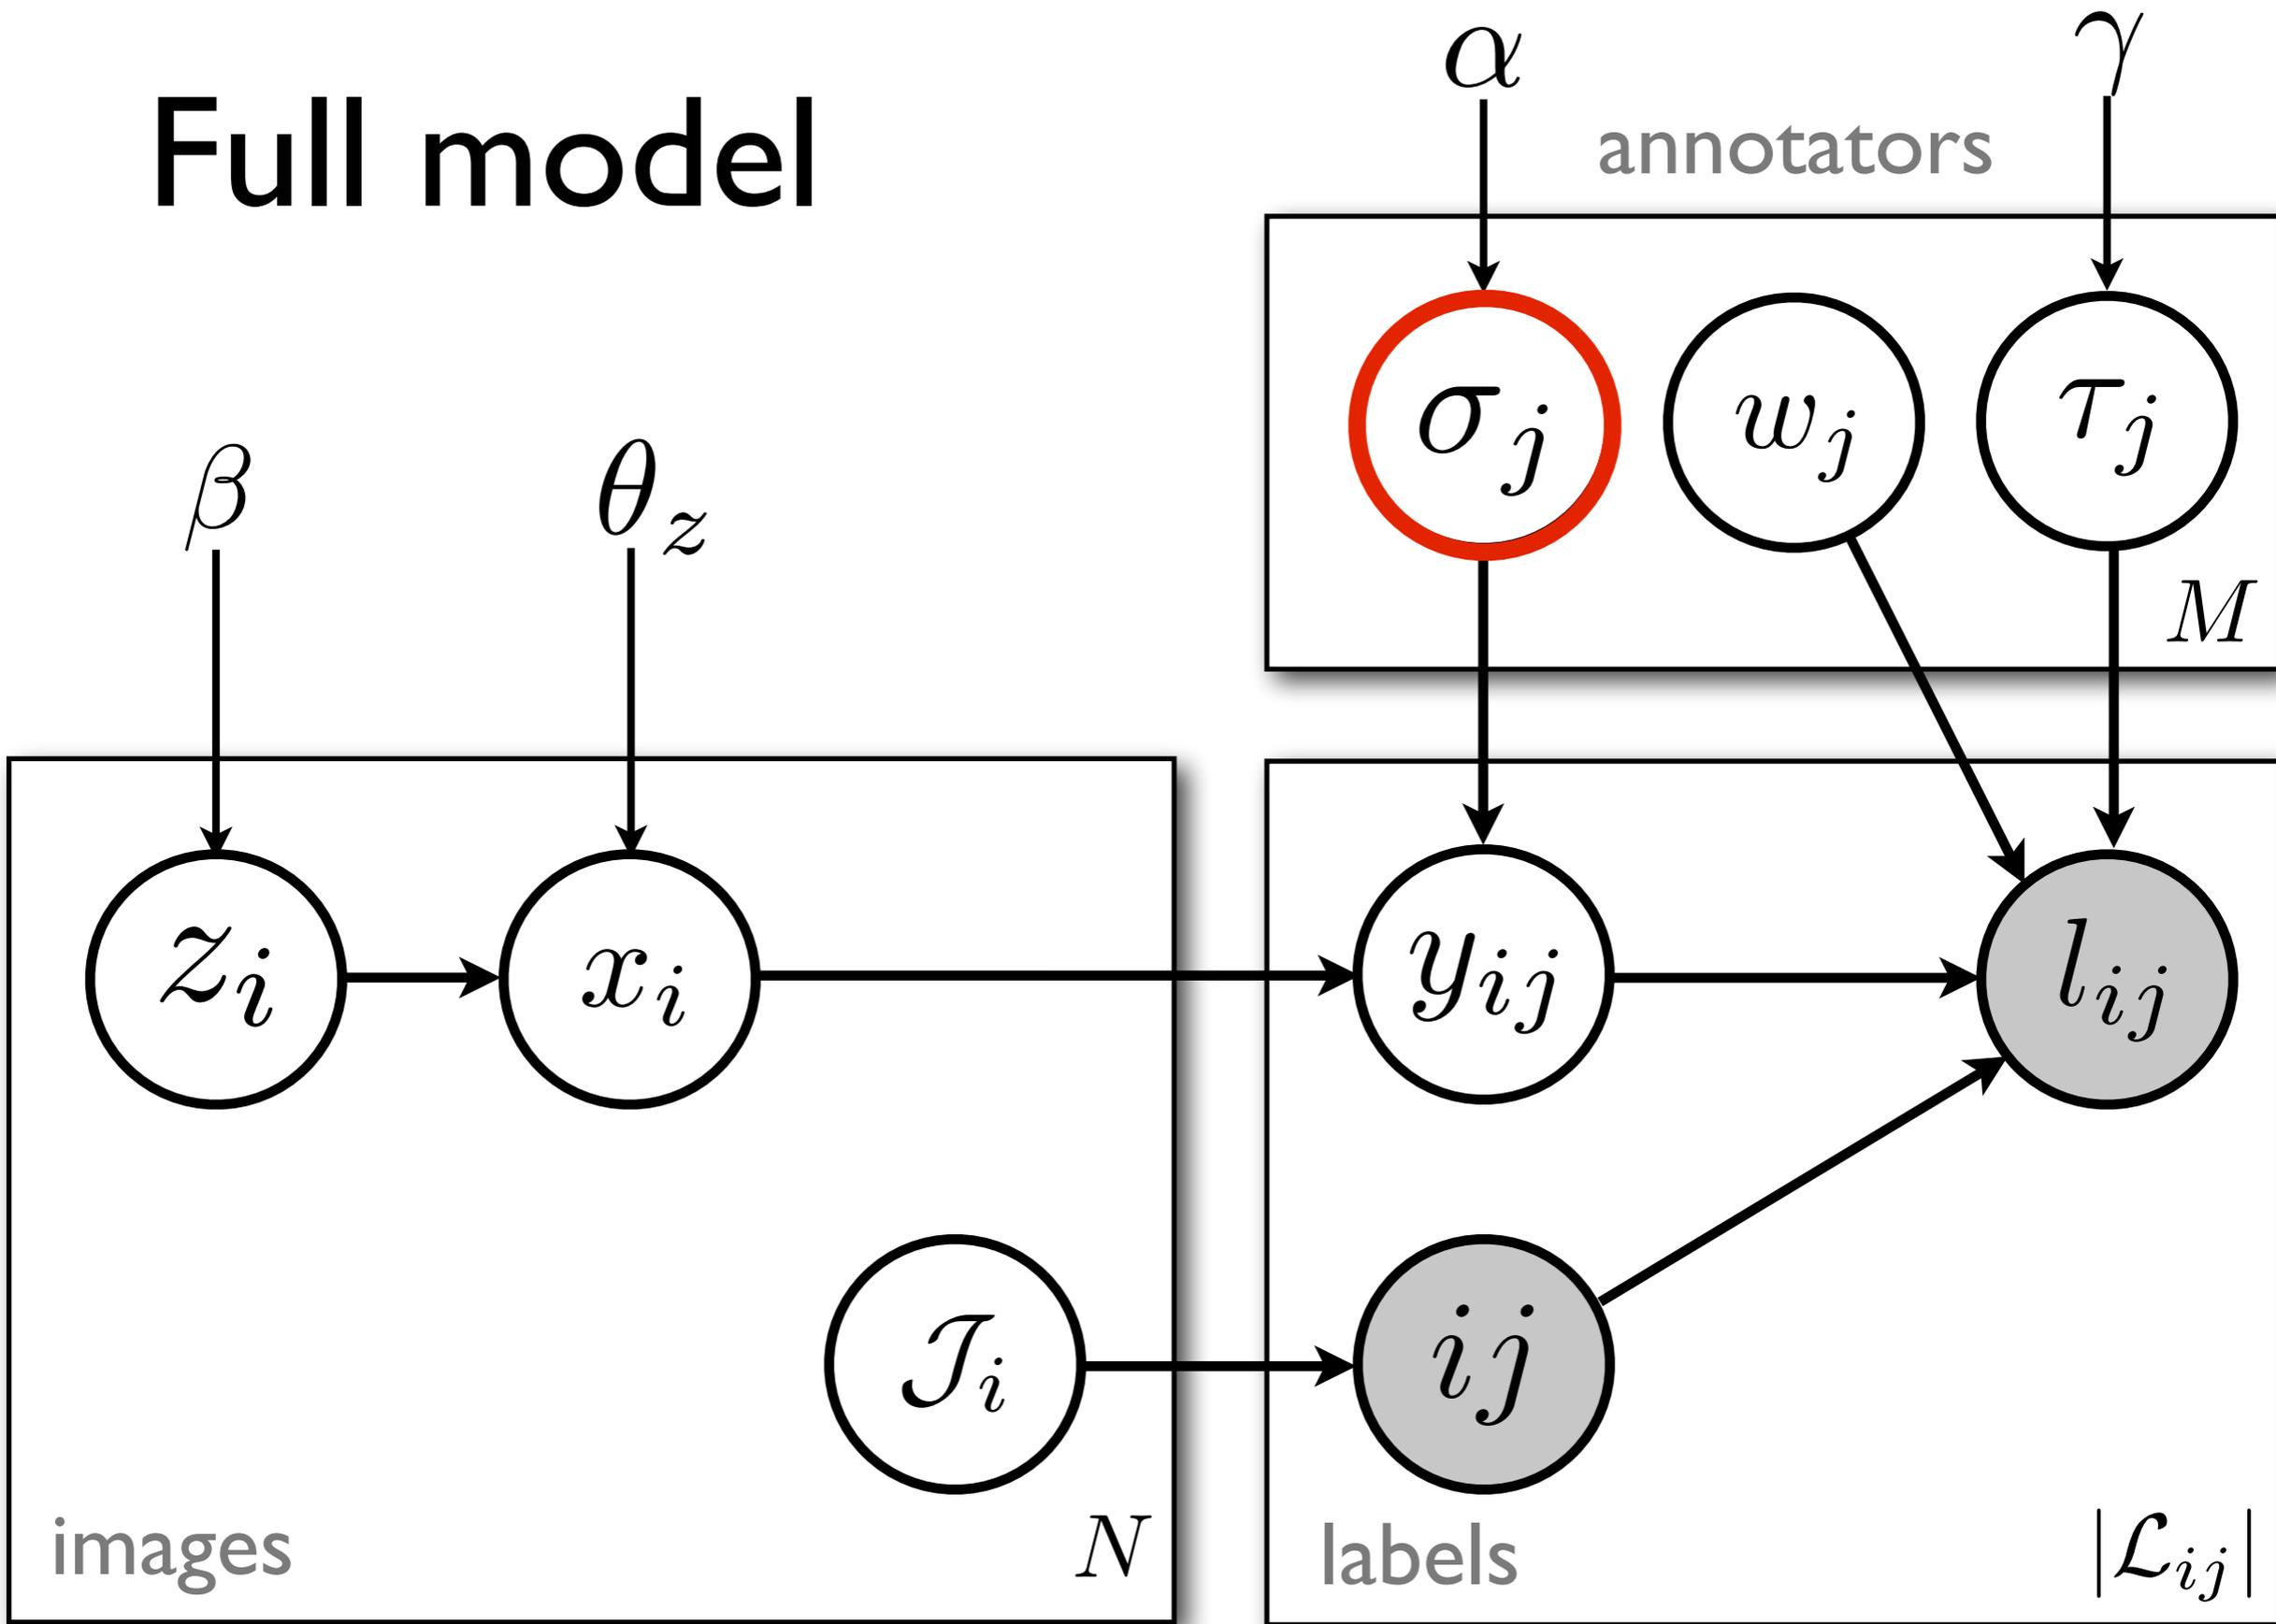
# Full model



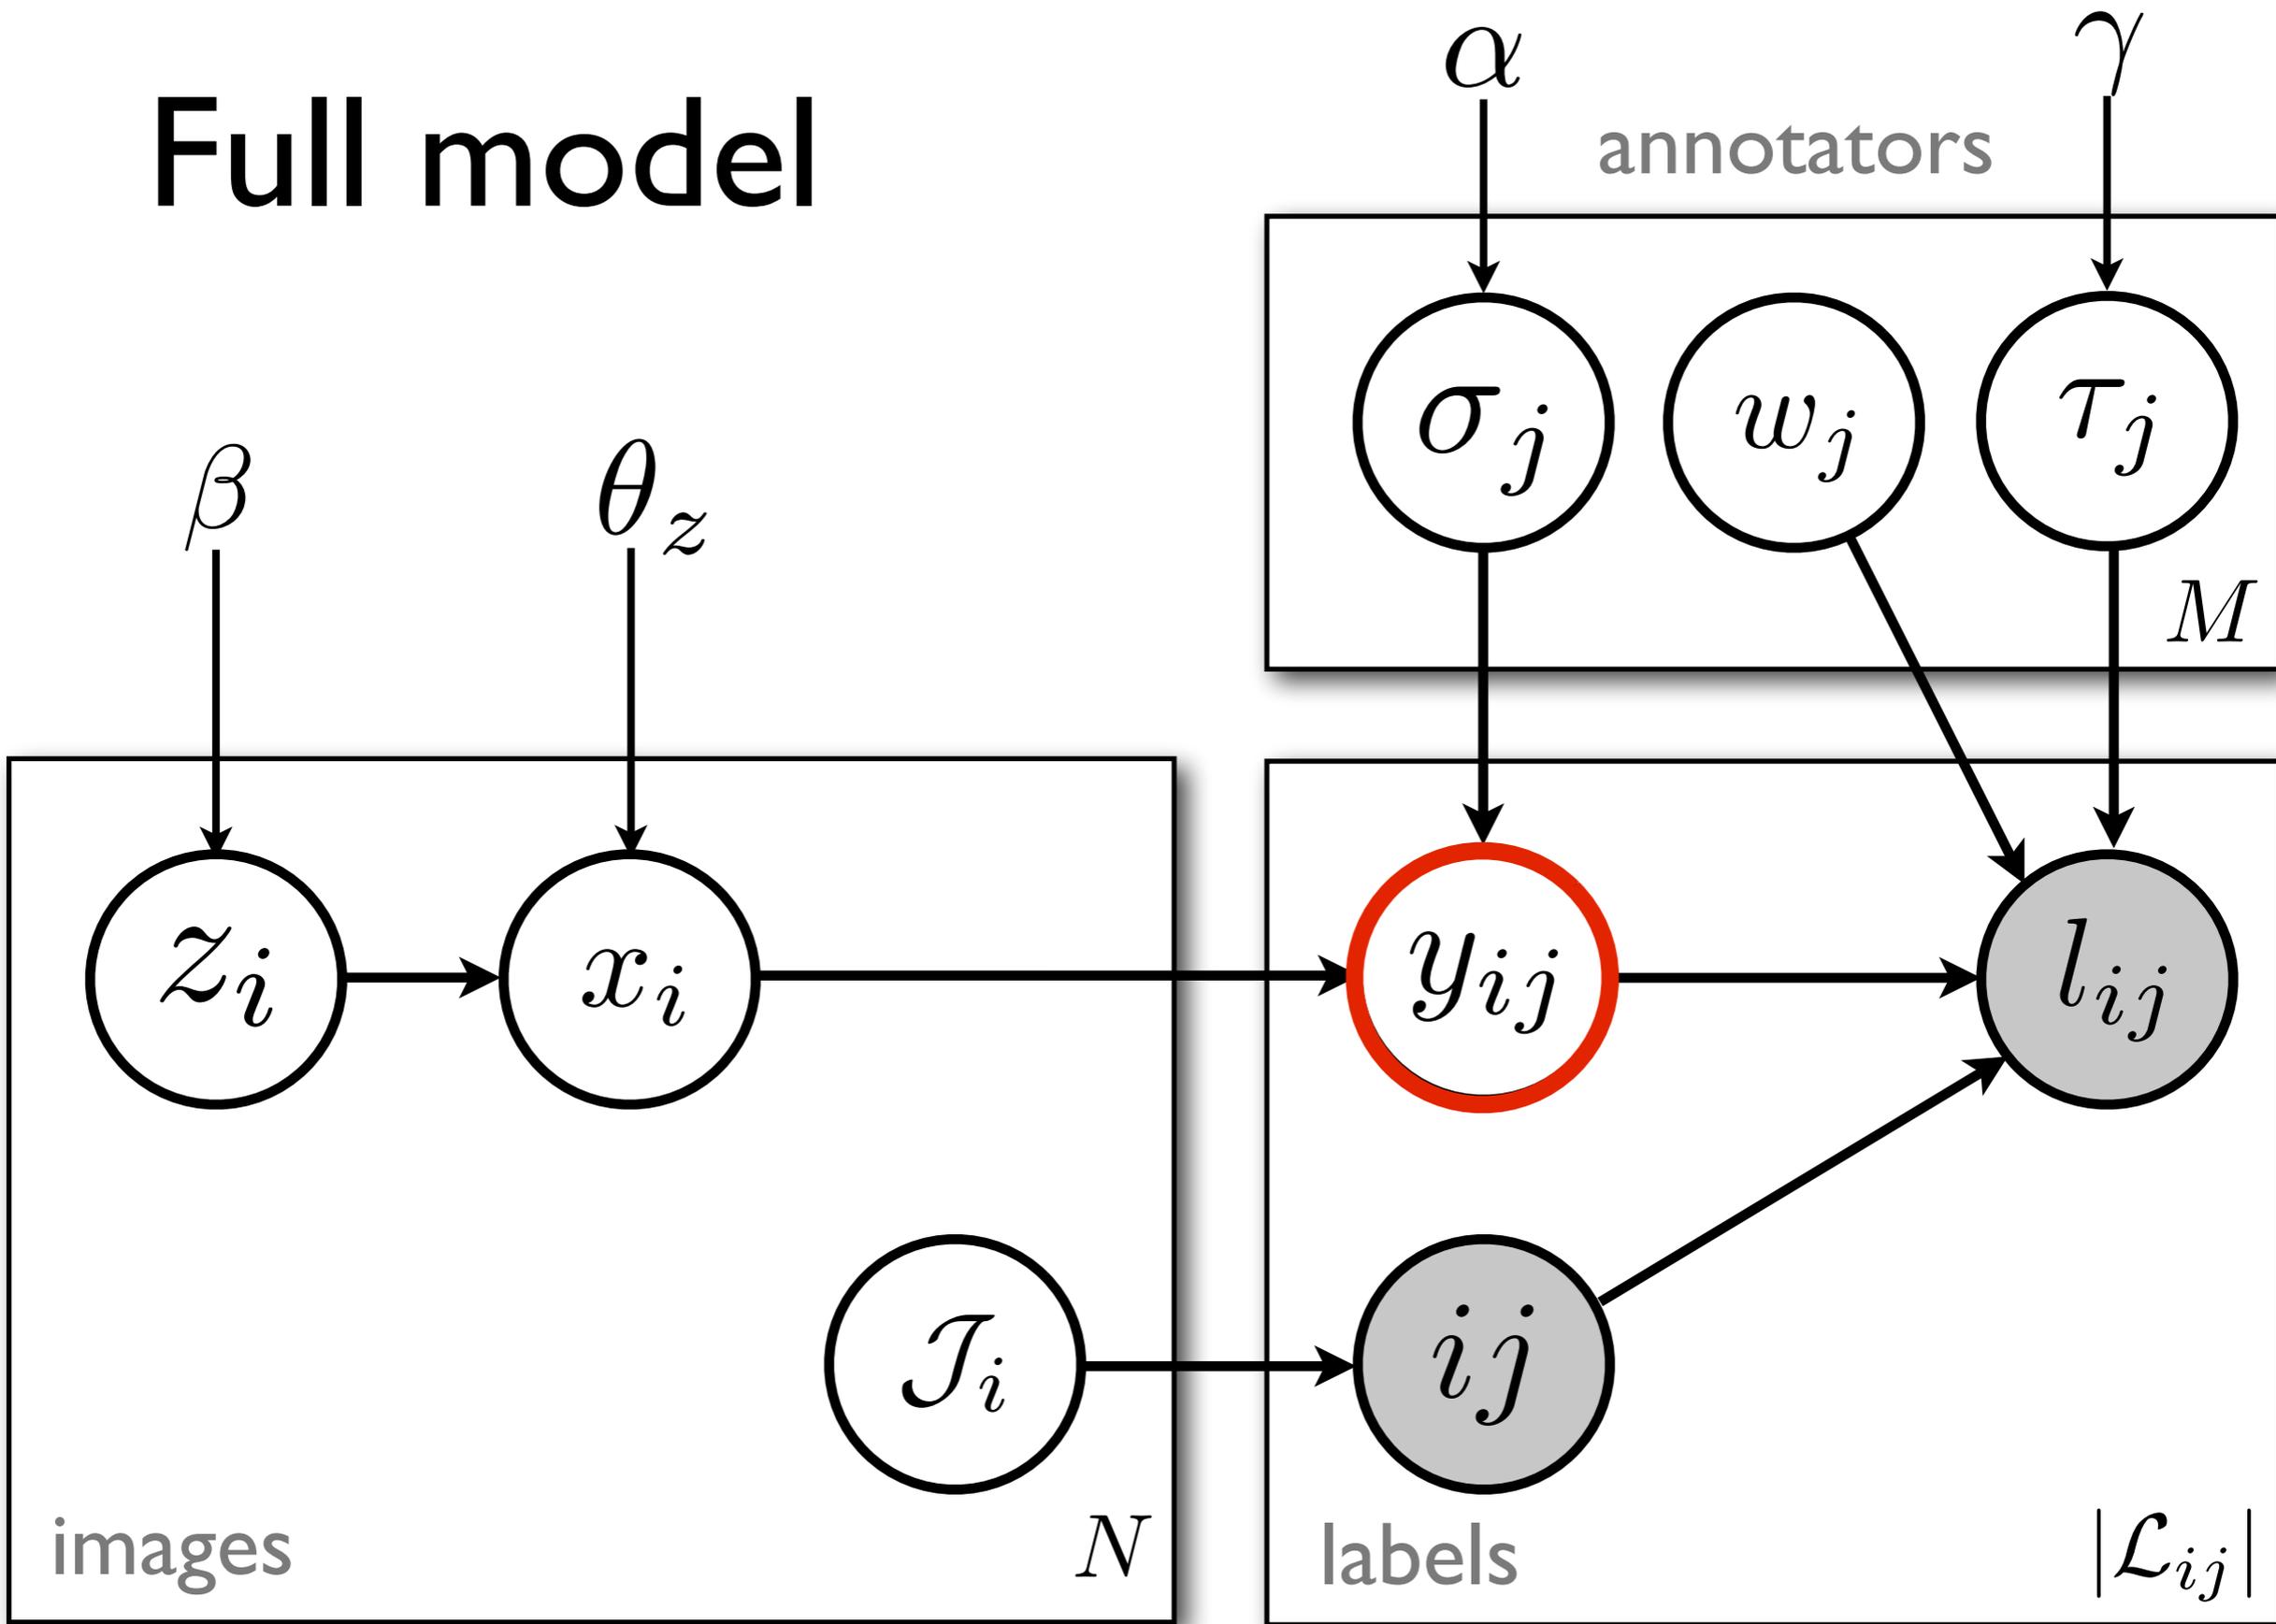
# Full model



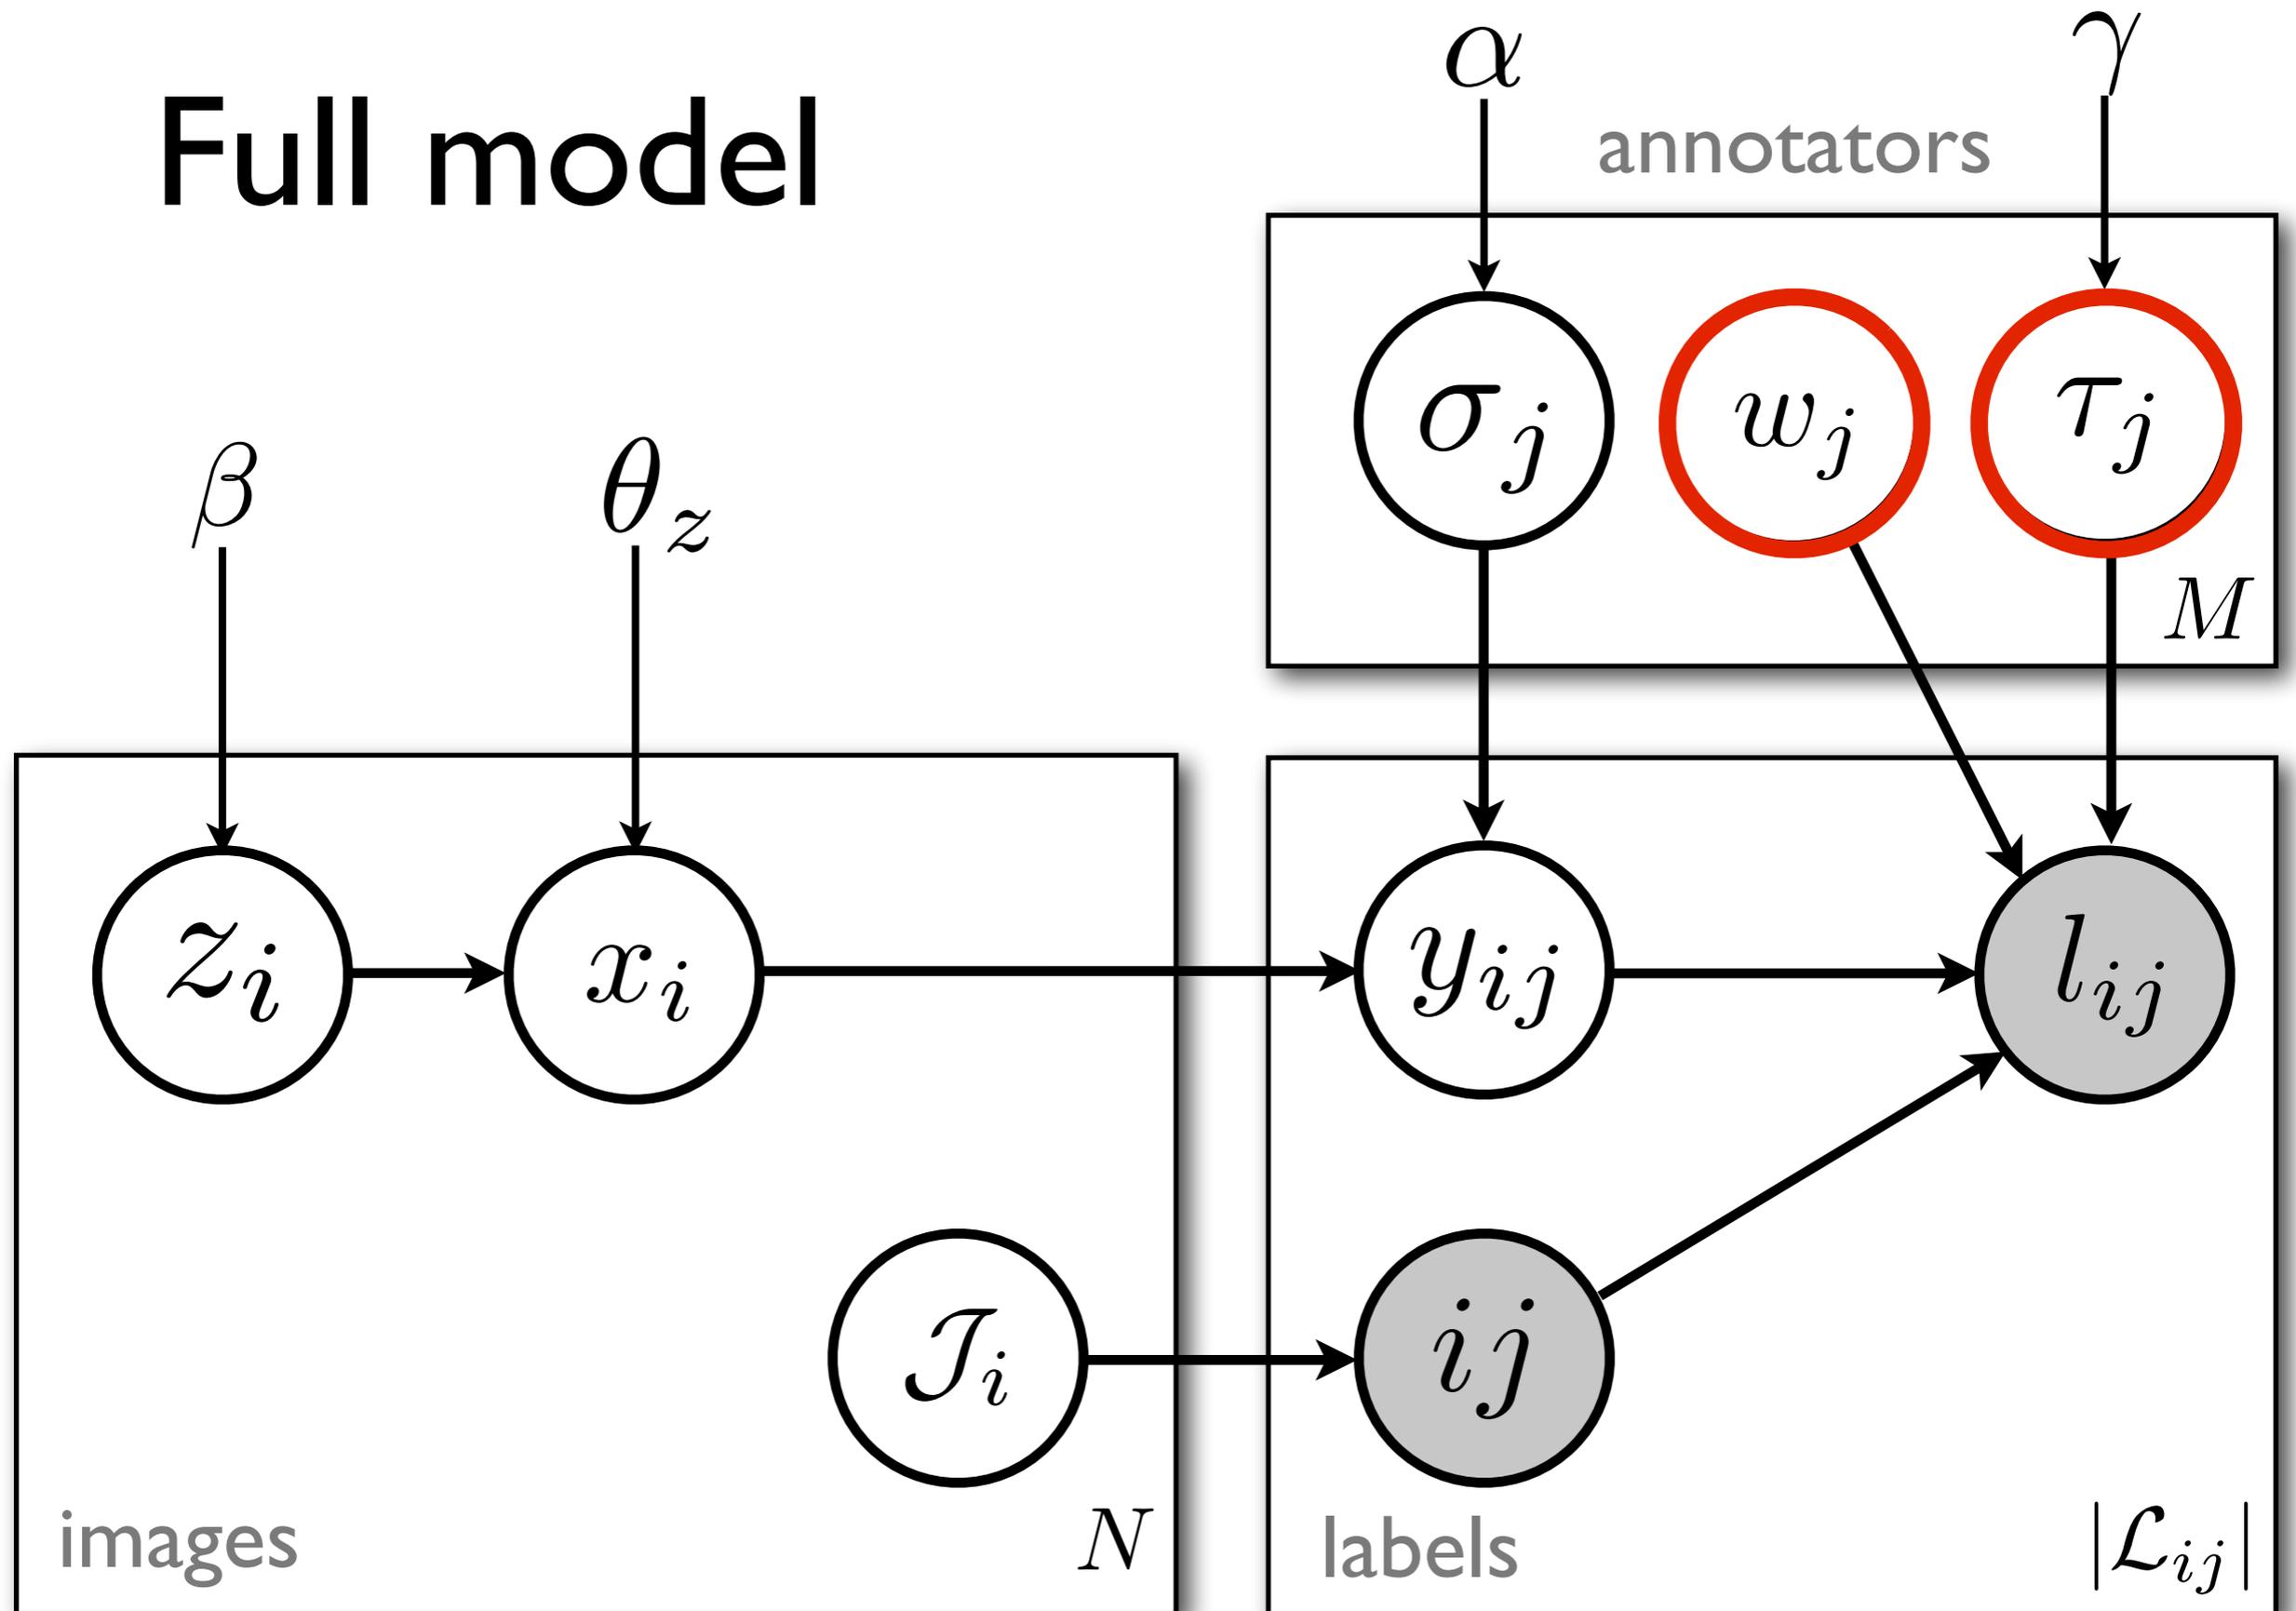
# Full model



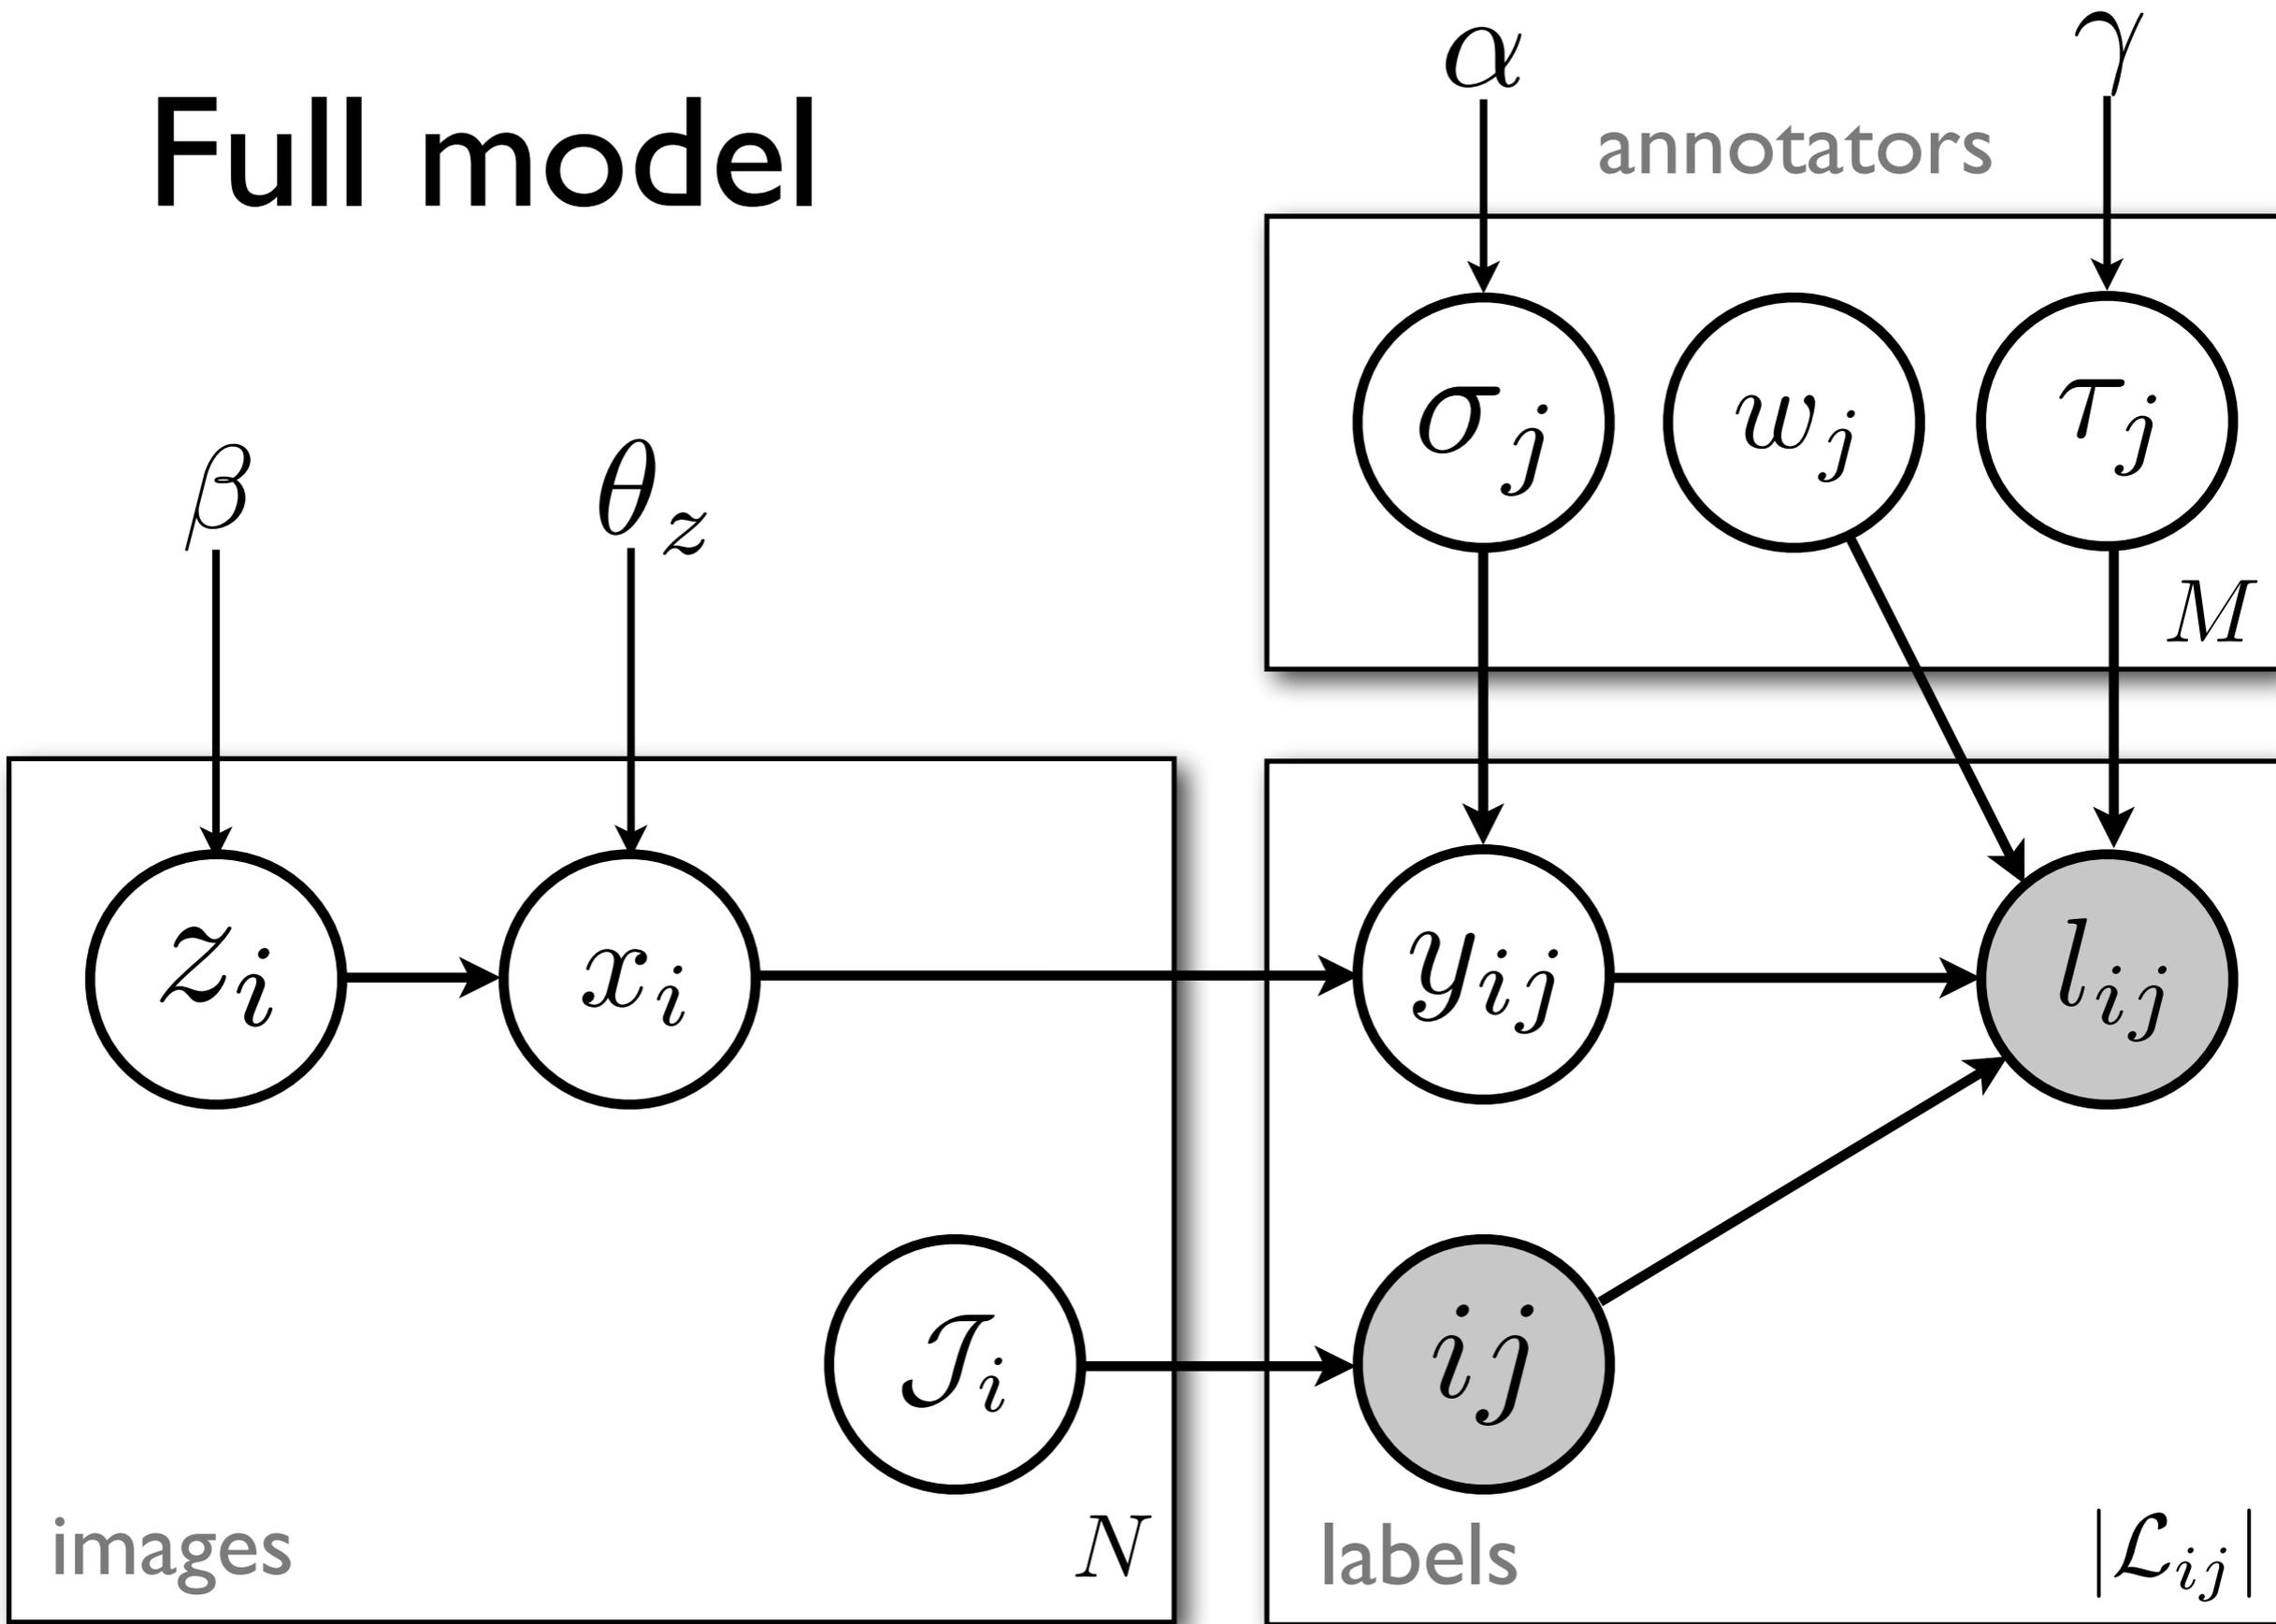
# Full model



# Full model



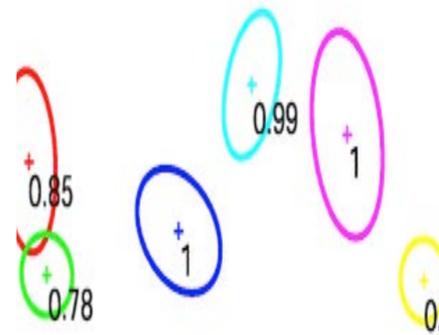
# Full model



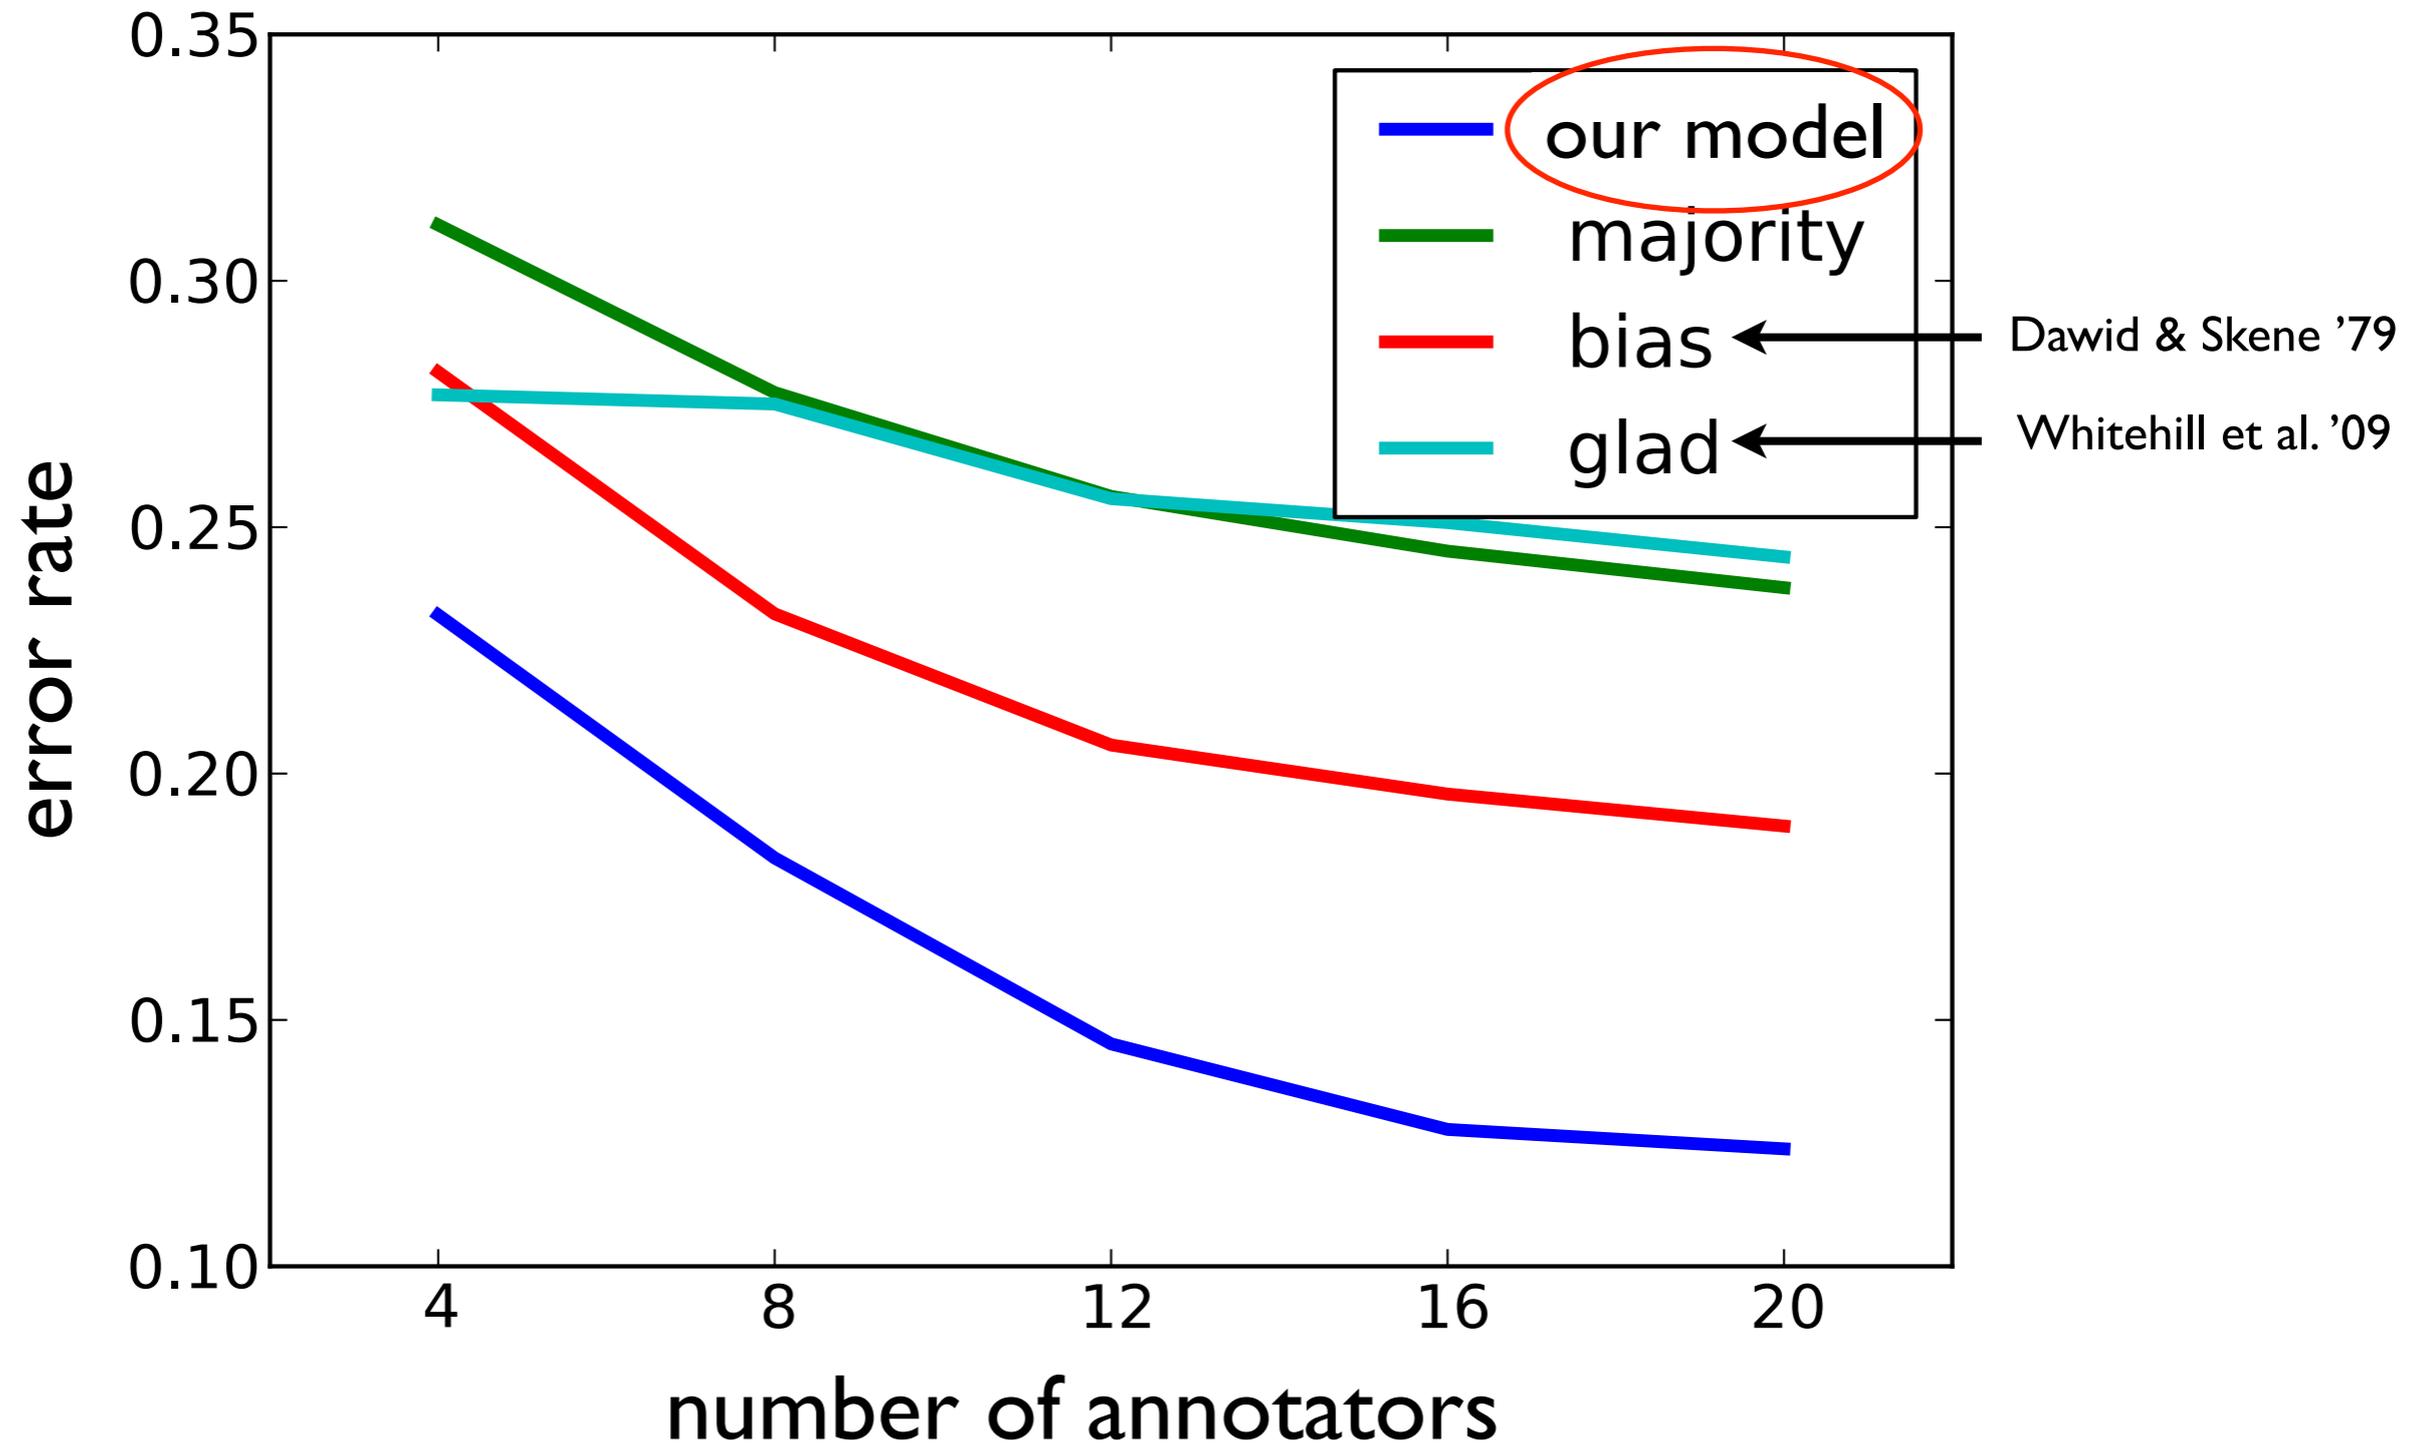
# Experiments: estimating $z_i$

Indigo Bunting

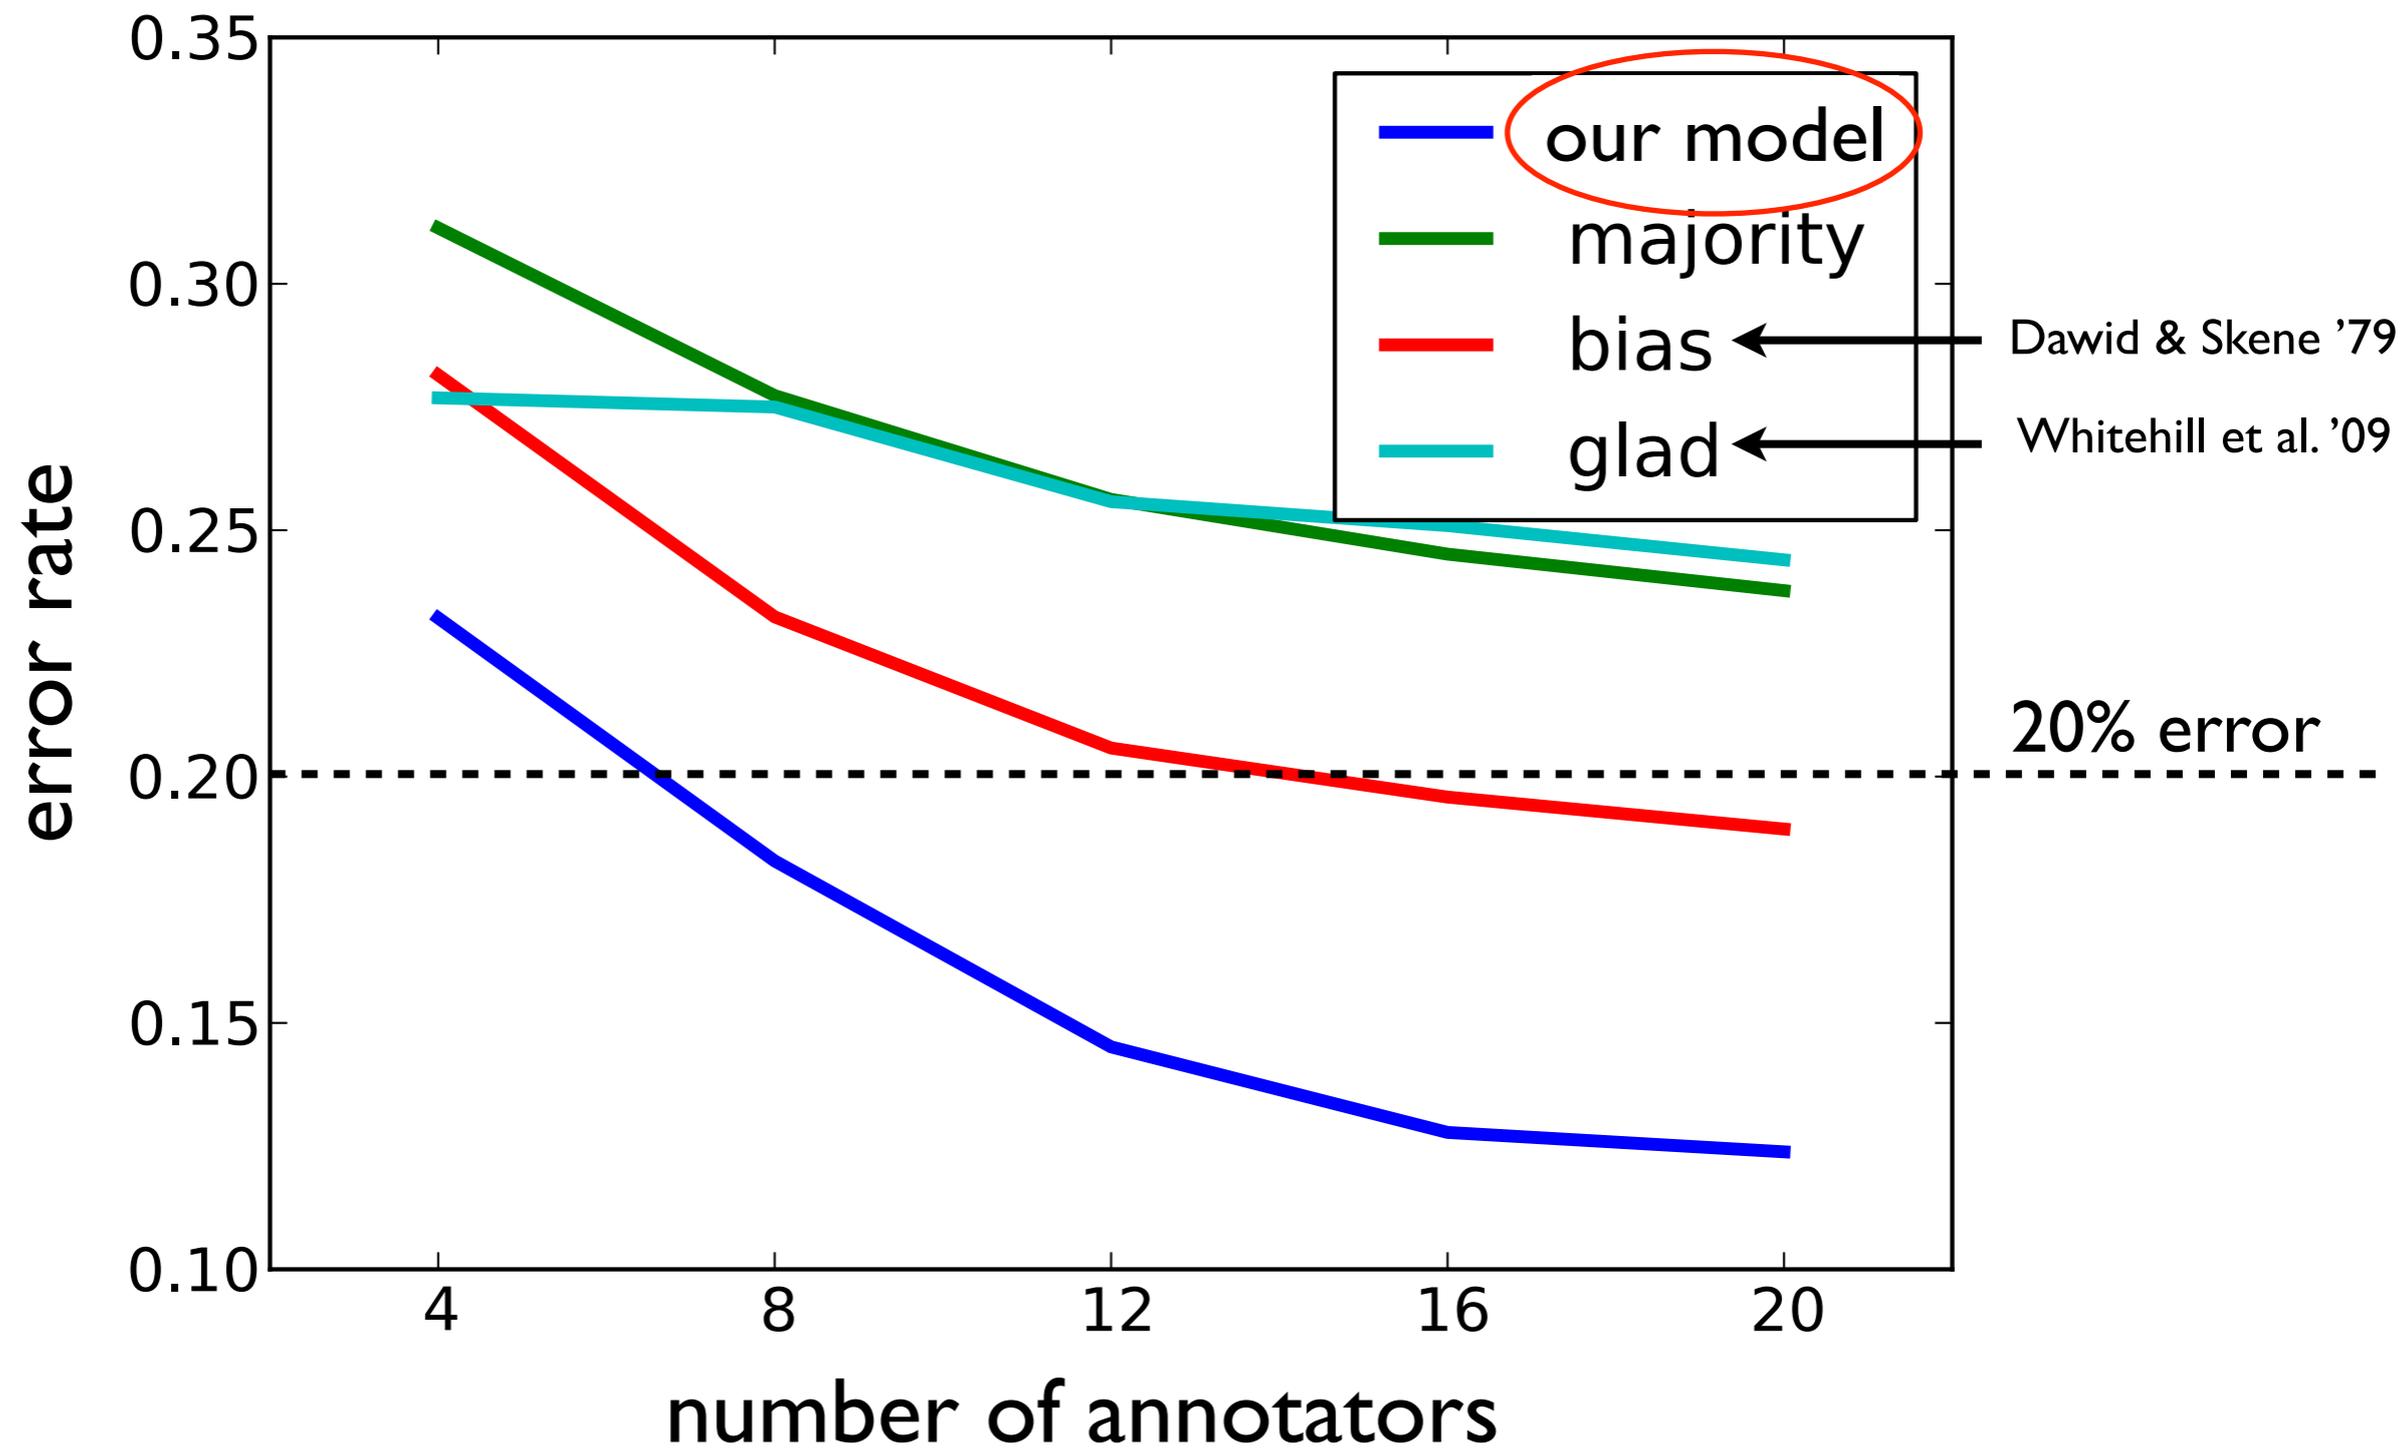
Blue Grosbeak



# Experiments: estimating $z_i$



# Experiments: estimating $z_i$

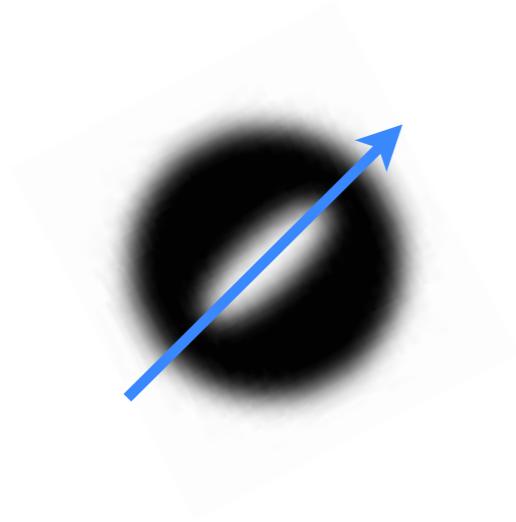


# Sanity Check 1: Rotated Ellipses Experiment



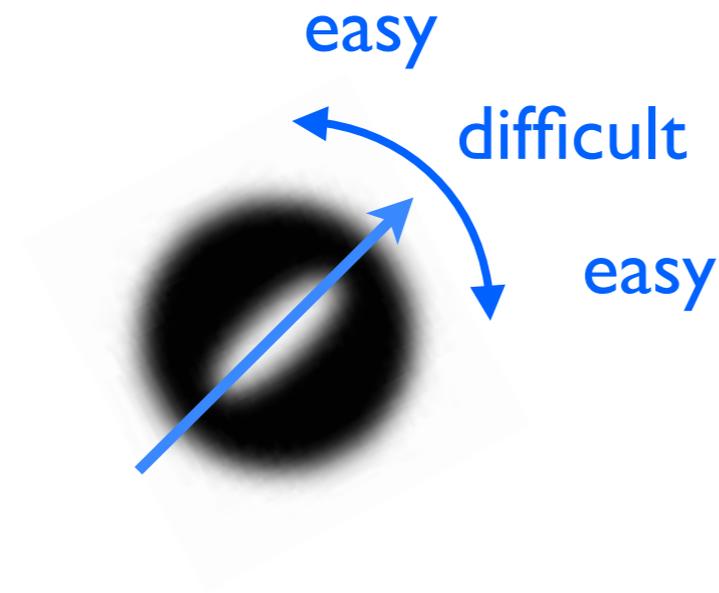
Is the ellipse oriented vertically?

# Sanity Check 1: Rotated Ellipses Experiment



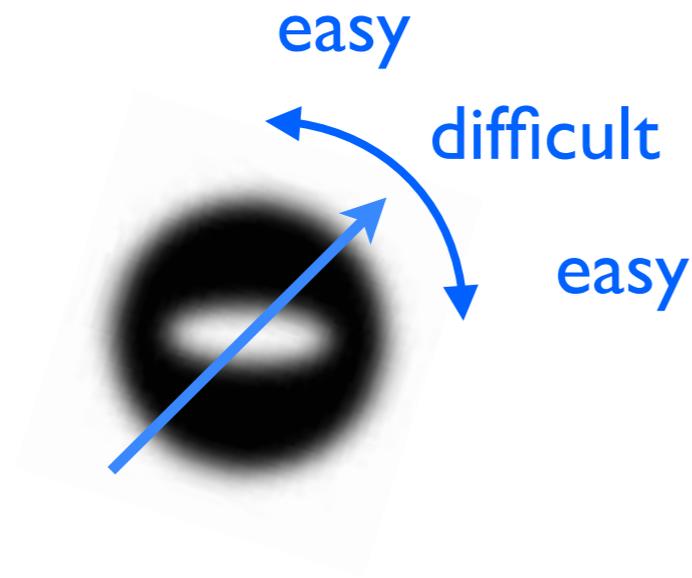
Is the ellipse oriented vertically?

# Sanity Check 1: Rotated Ellipses Experiment



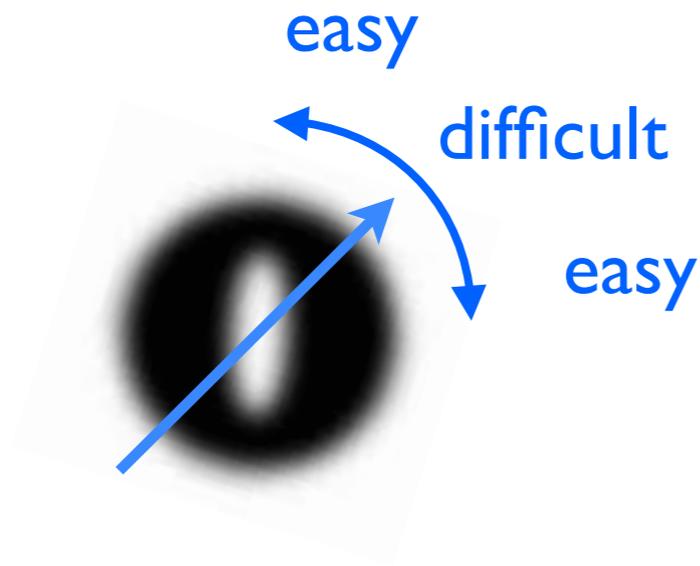
Is the ellipse oriented vertically?

# Sanity Check 1: Rotated Ellipses Experiment



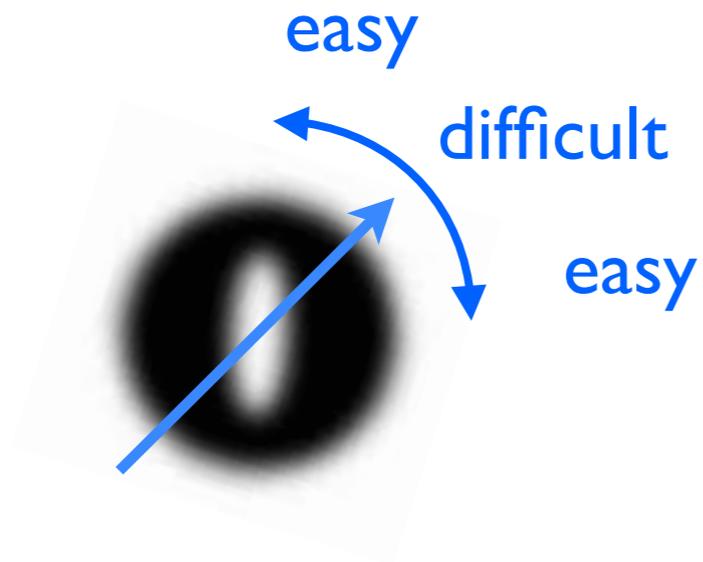
Is the ellipse oriented vertically?

# Sanity Check 1: Rotated Ellipses Experiment



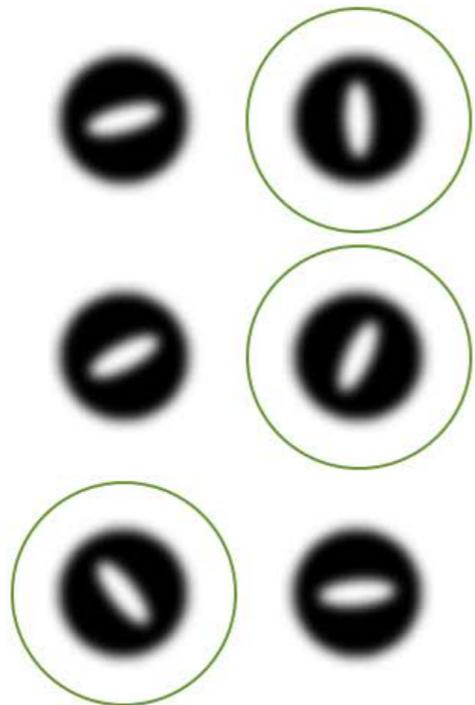
Is the ellipse oriented vertically?

# Sanity Check 1: Rotated Ellipses Experiment

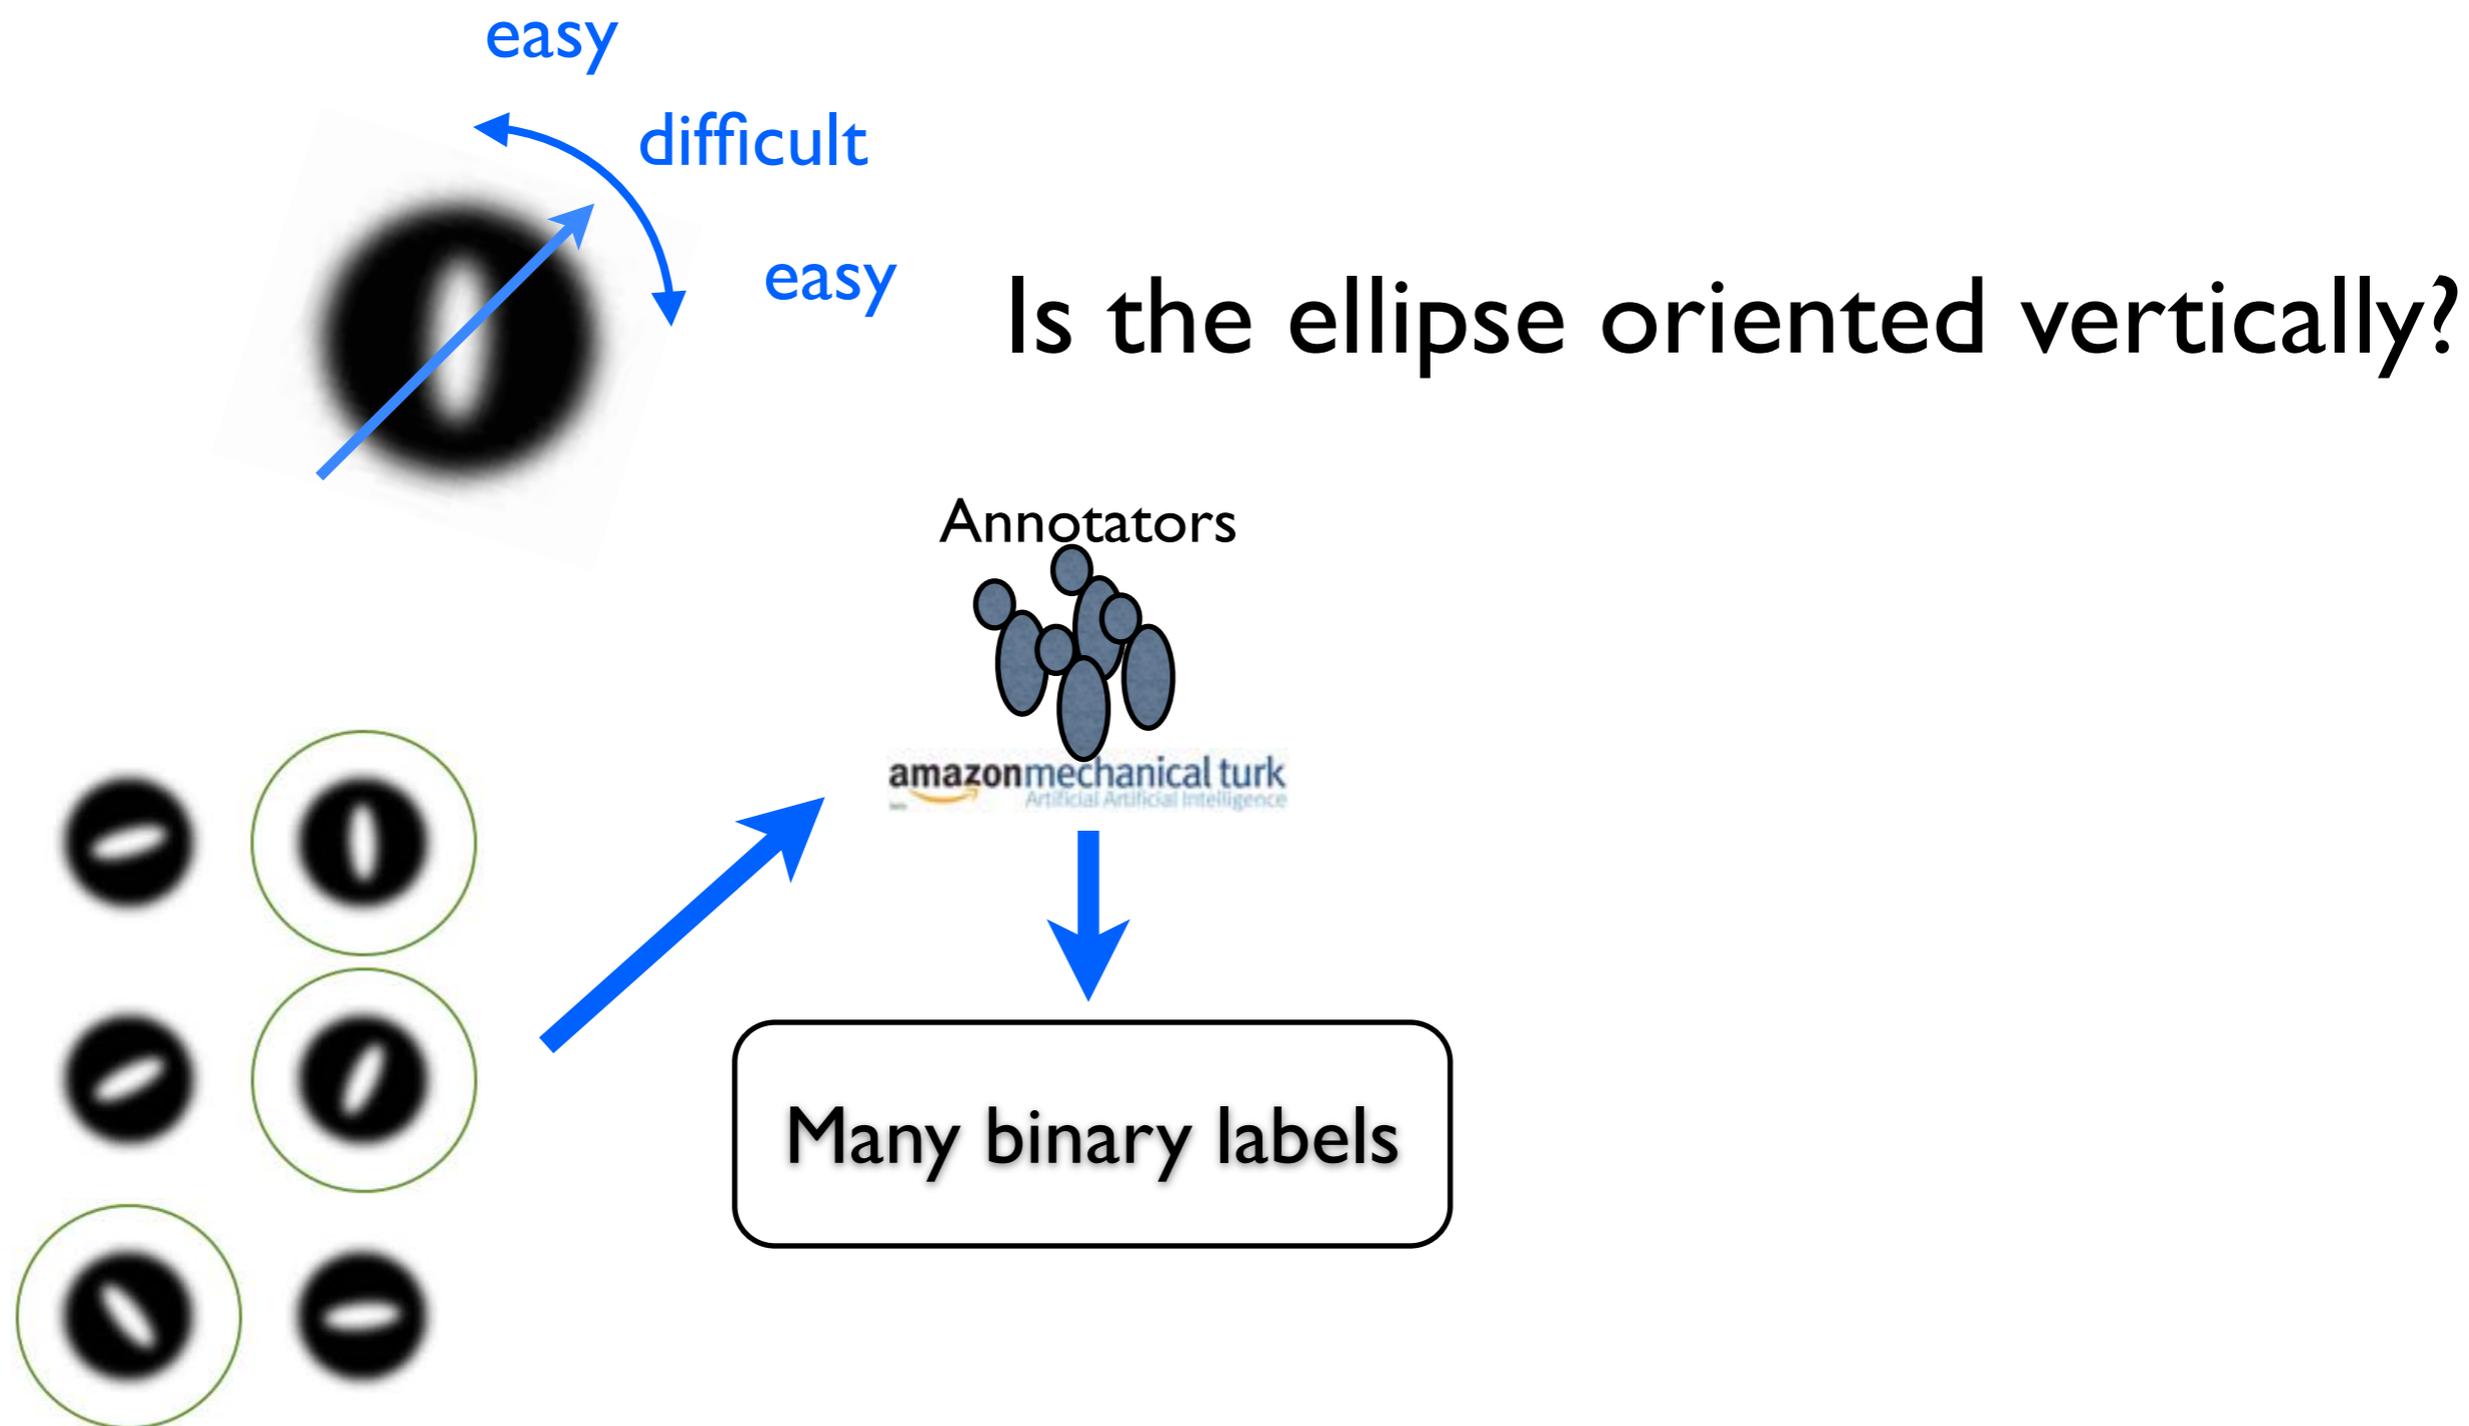


Is the ellipse oriented vertically?

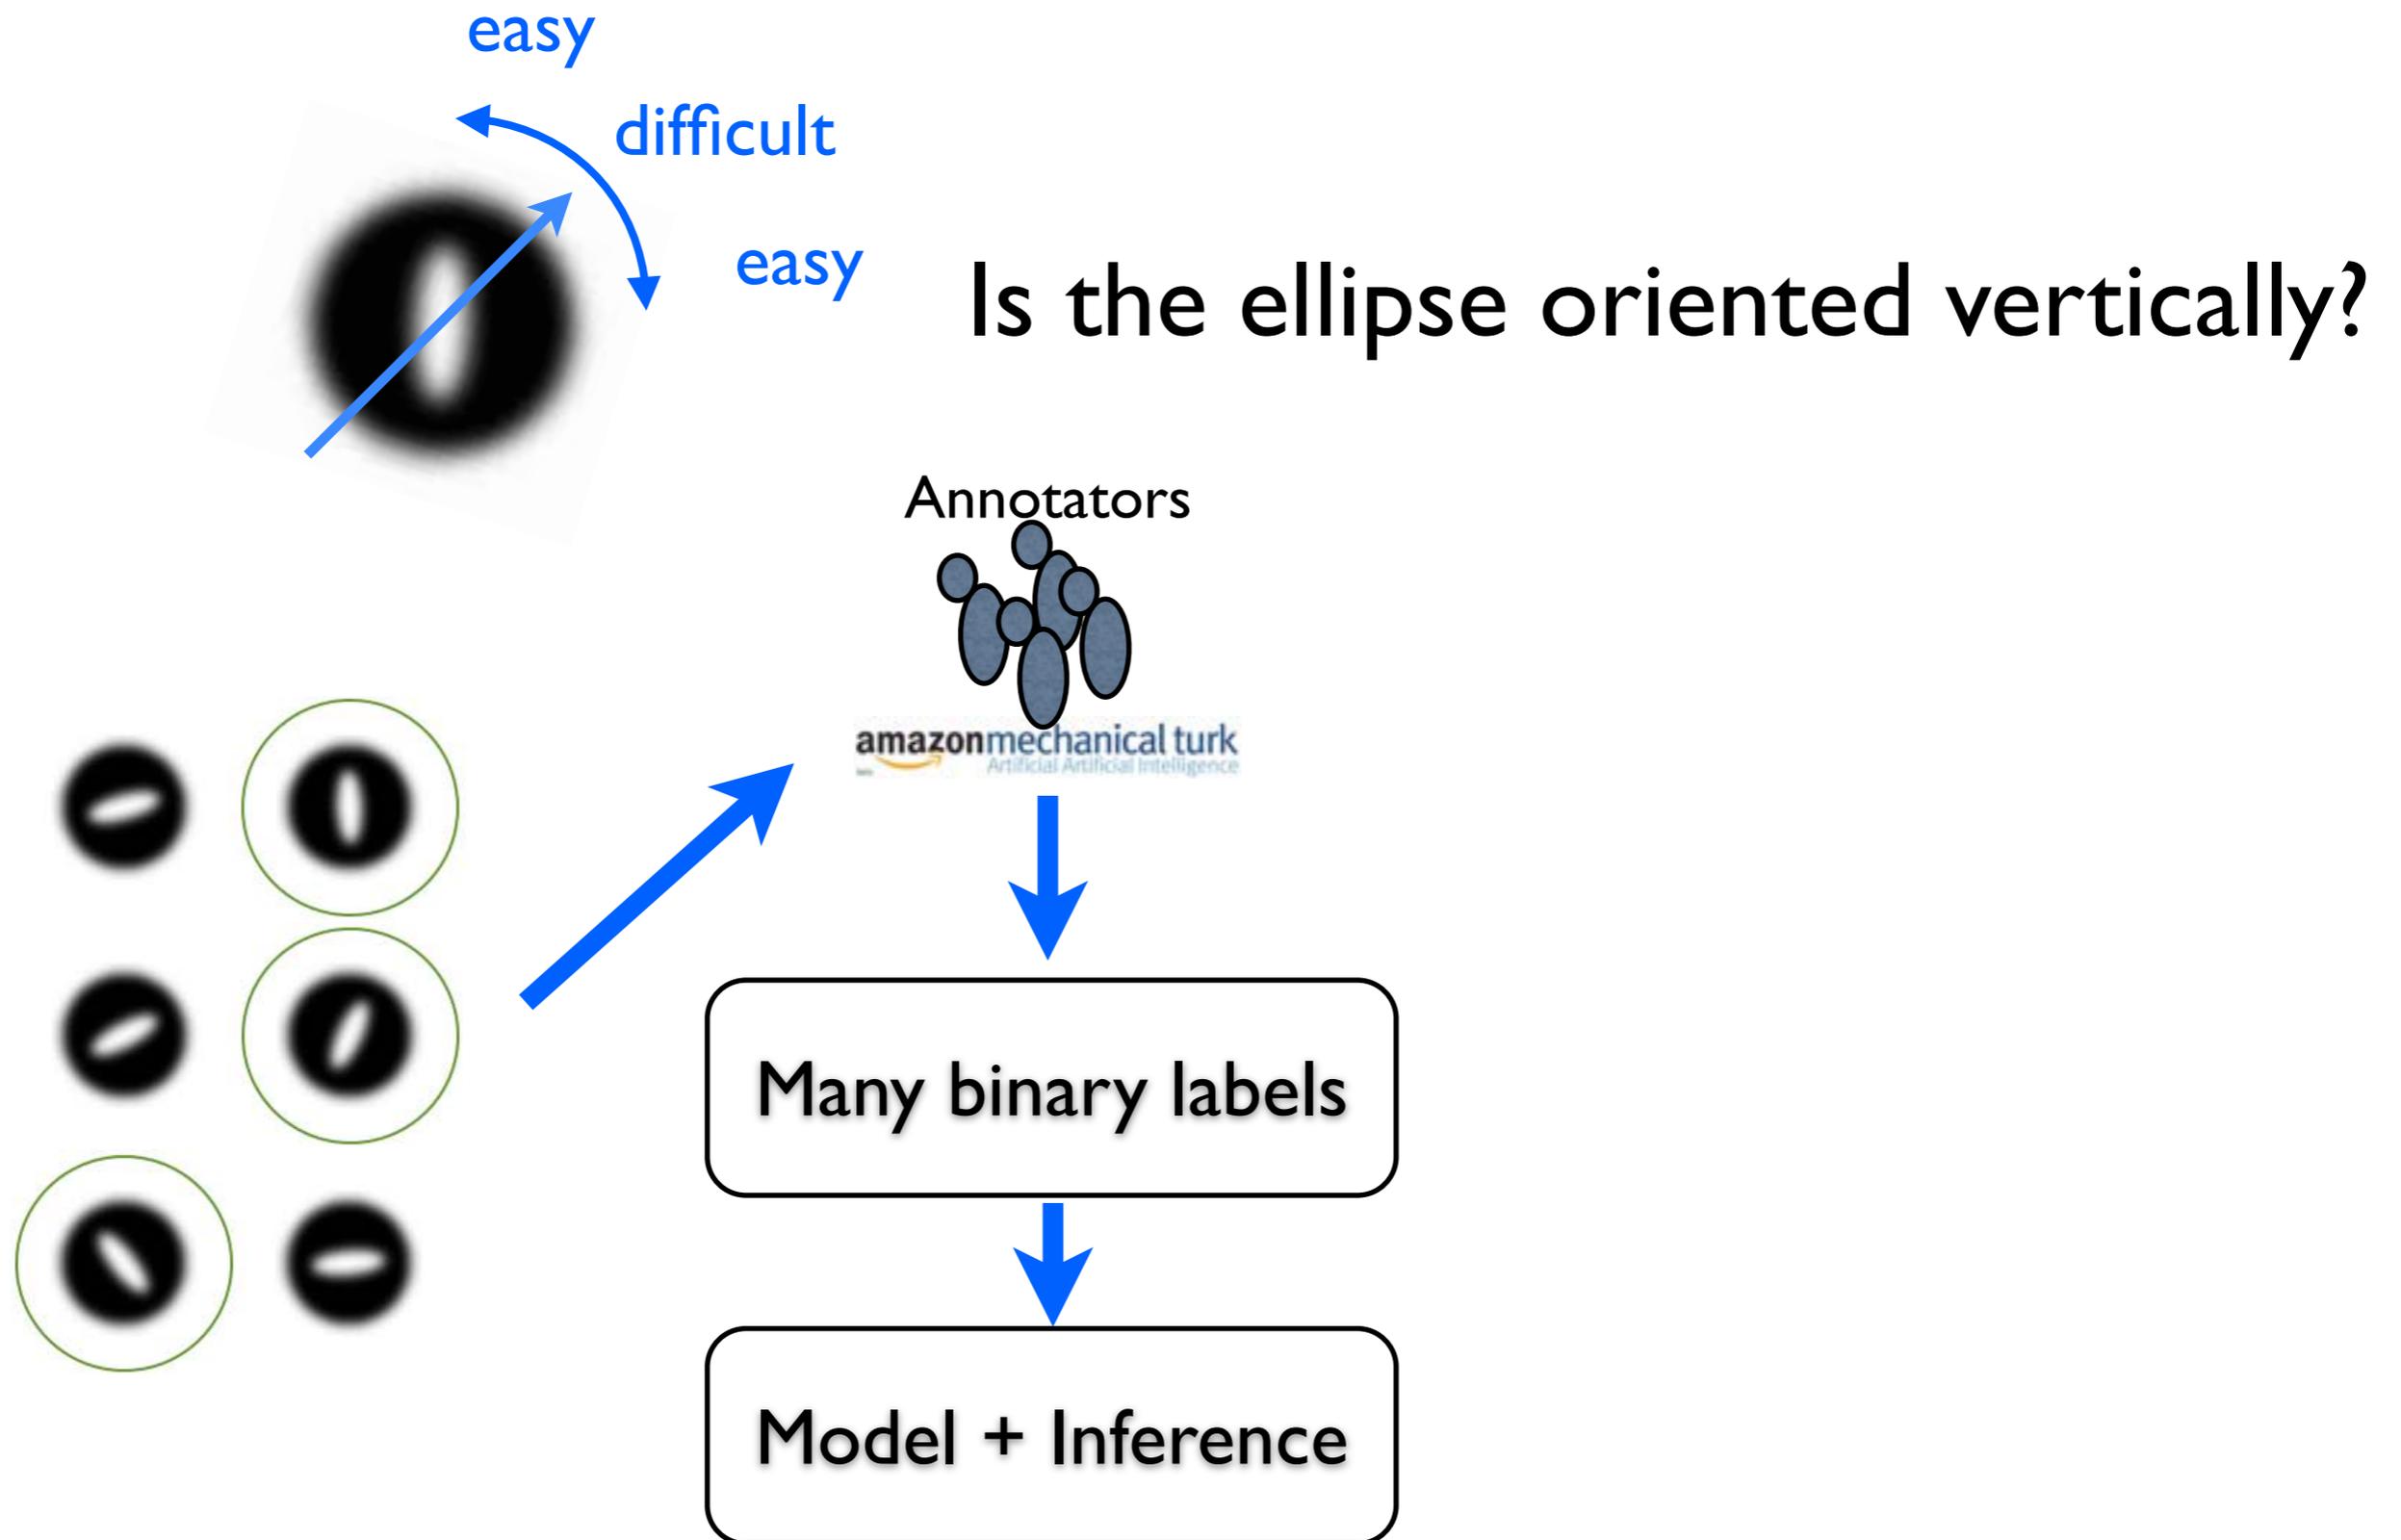
Annotators



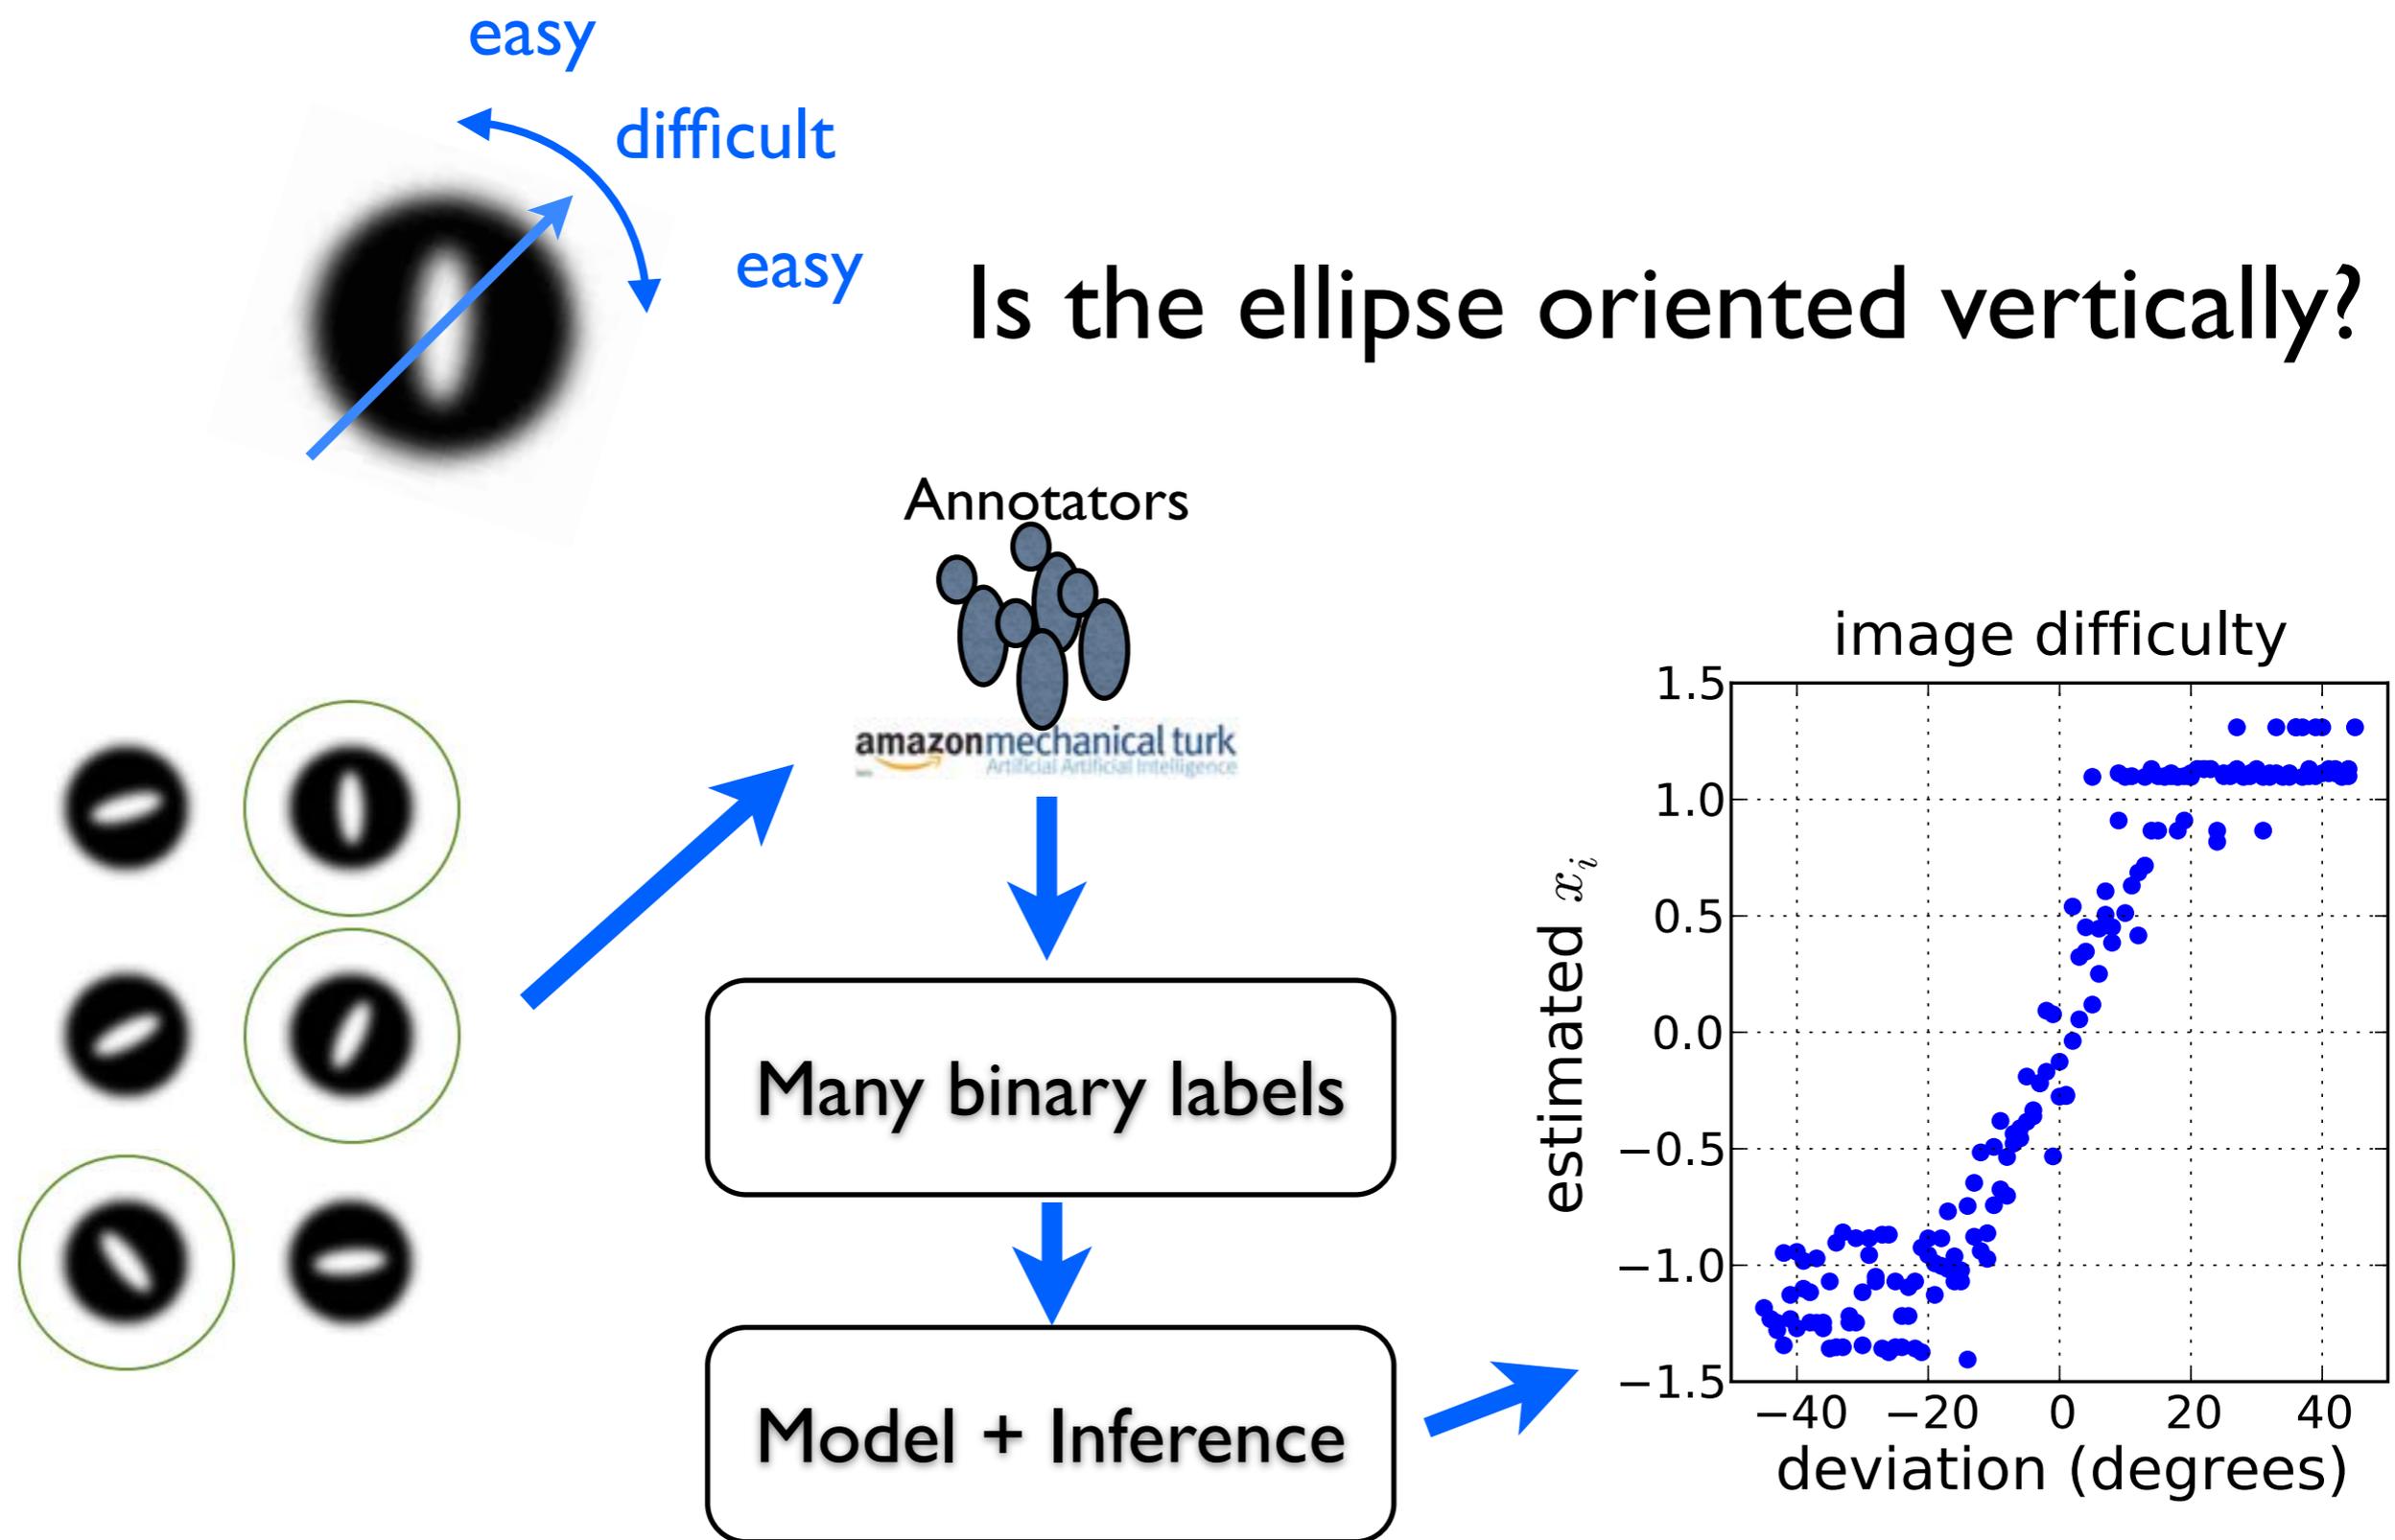
# Sanity Check 1: Rotated Ellipses Experiment



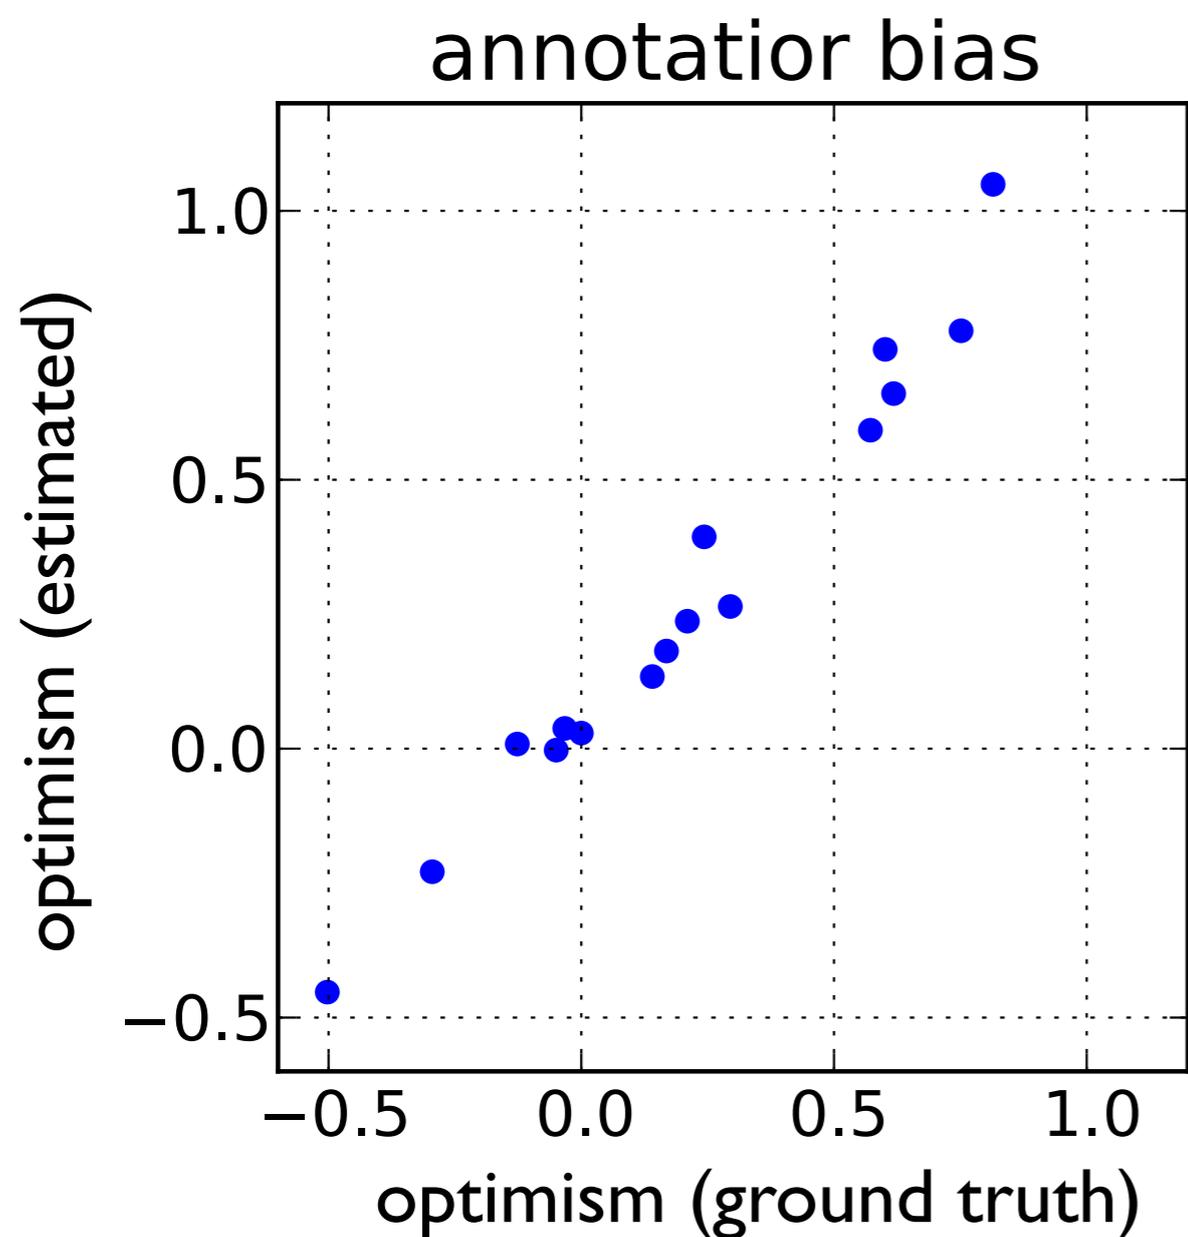
# Sanity Check 1: Rotated Ellipses Experiment



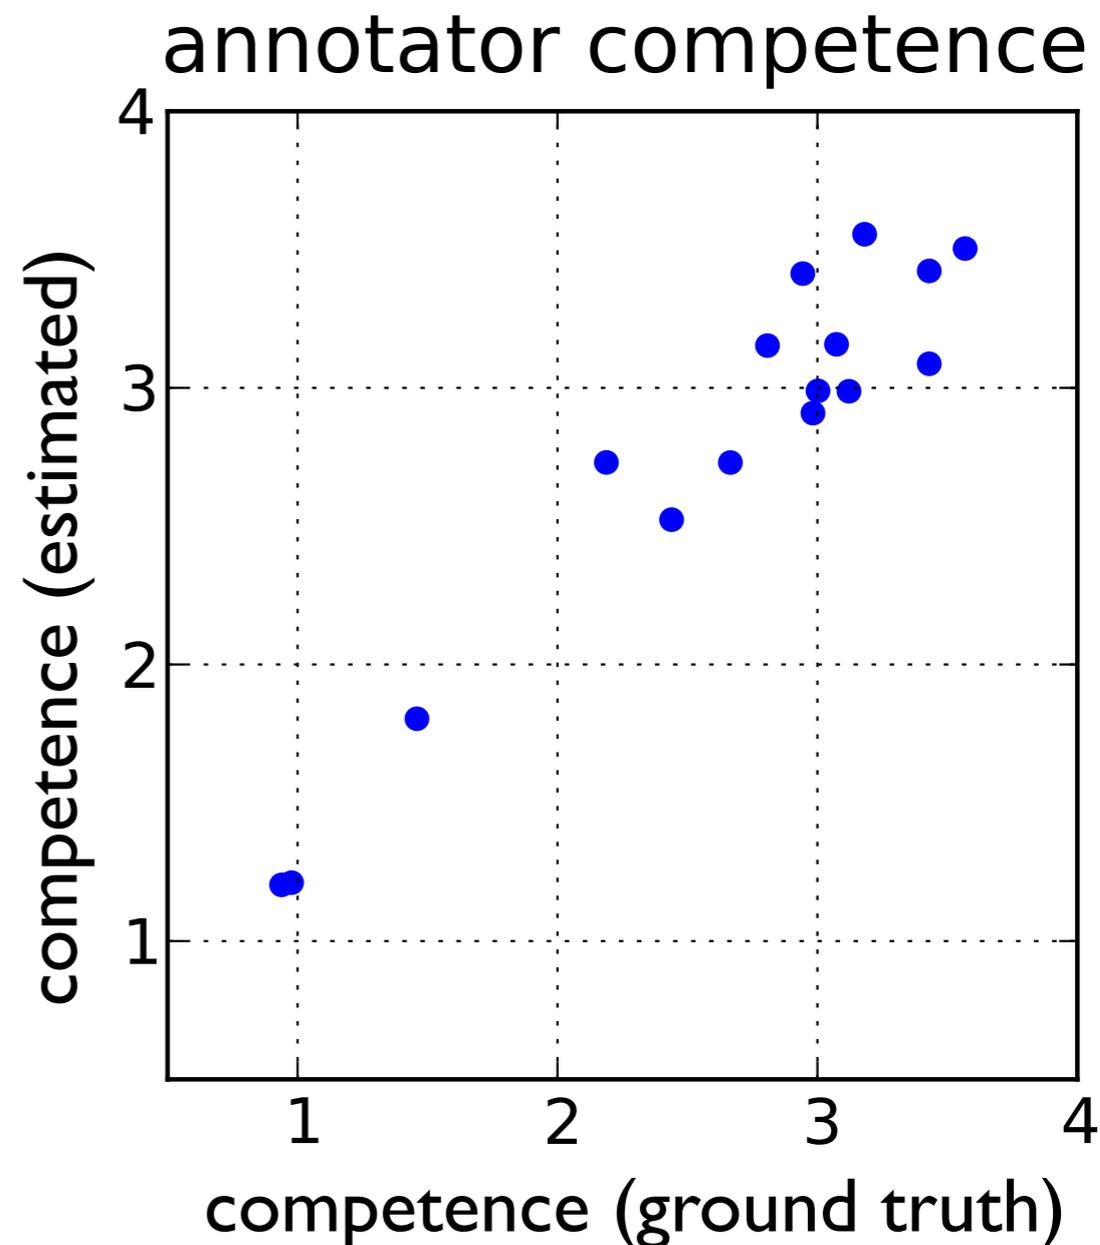
# Sanity Check 1: Rotated Ellipses Experiment



# Sanity Check 2: Rotated Ellipses Experiment



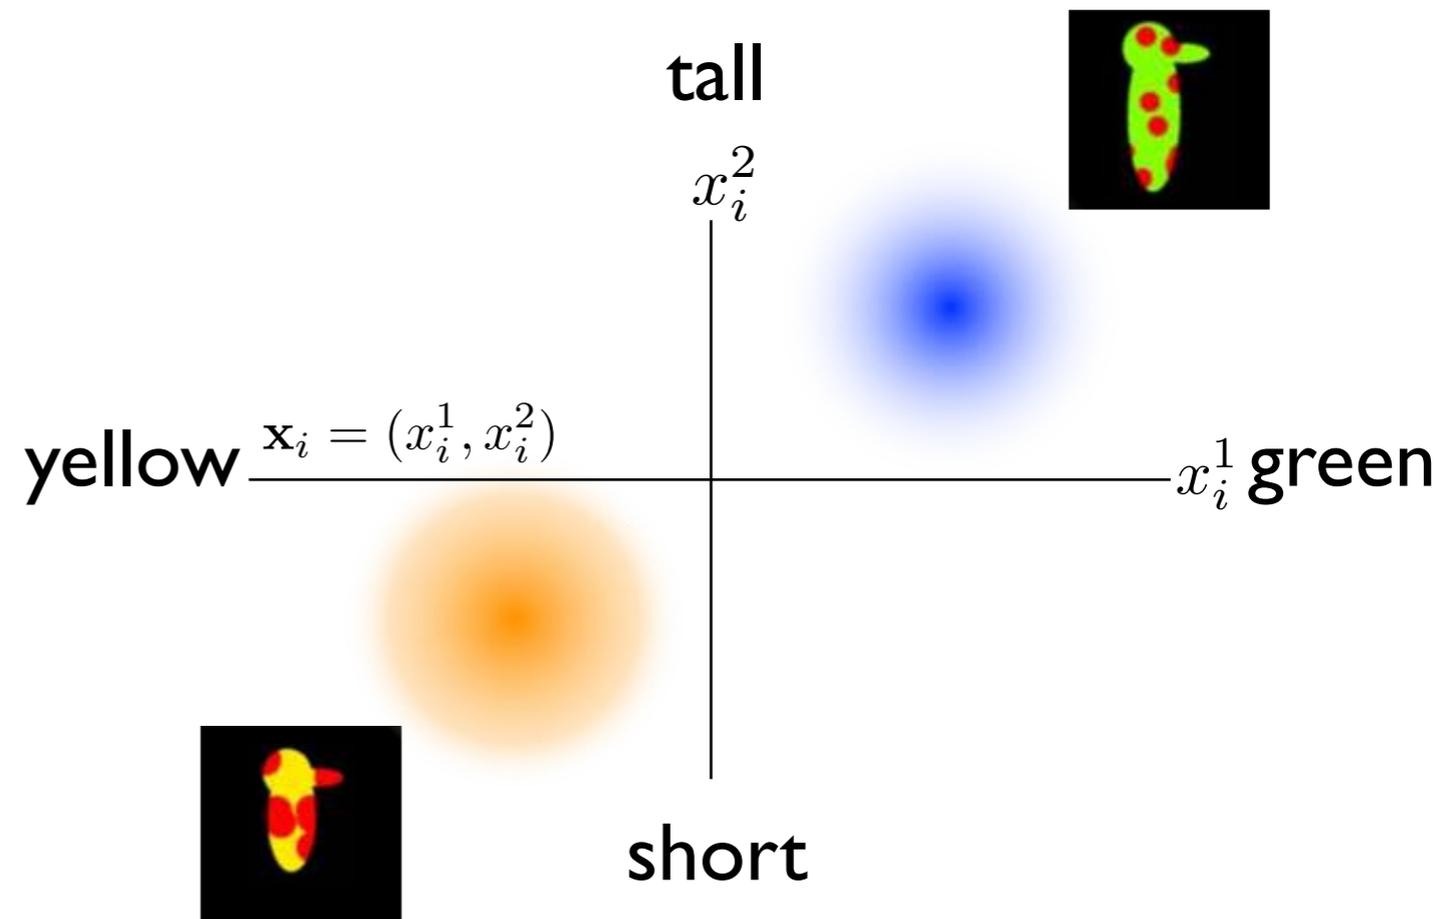
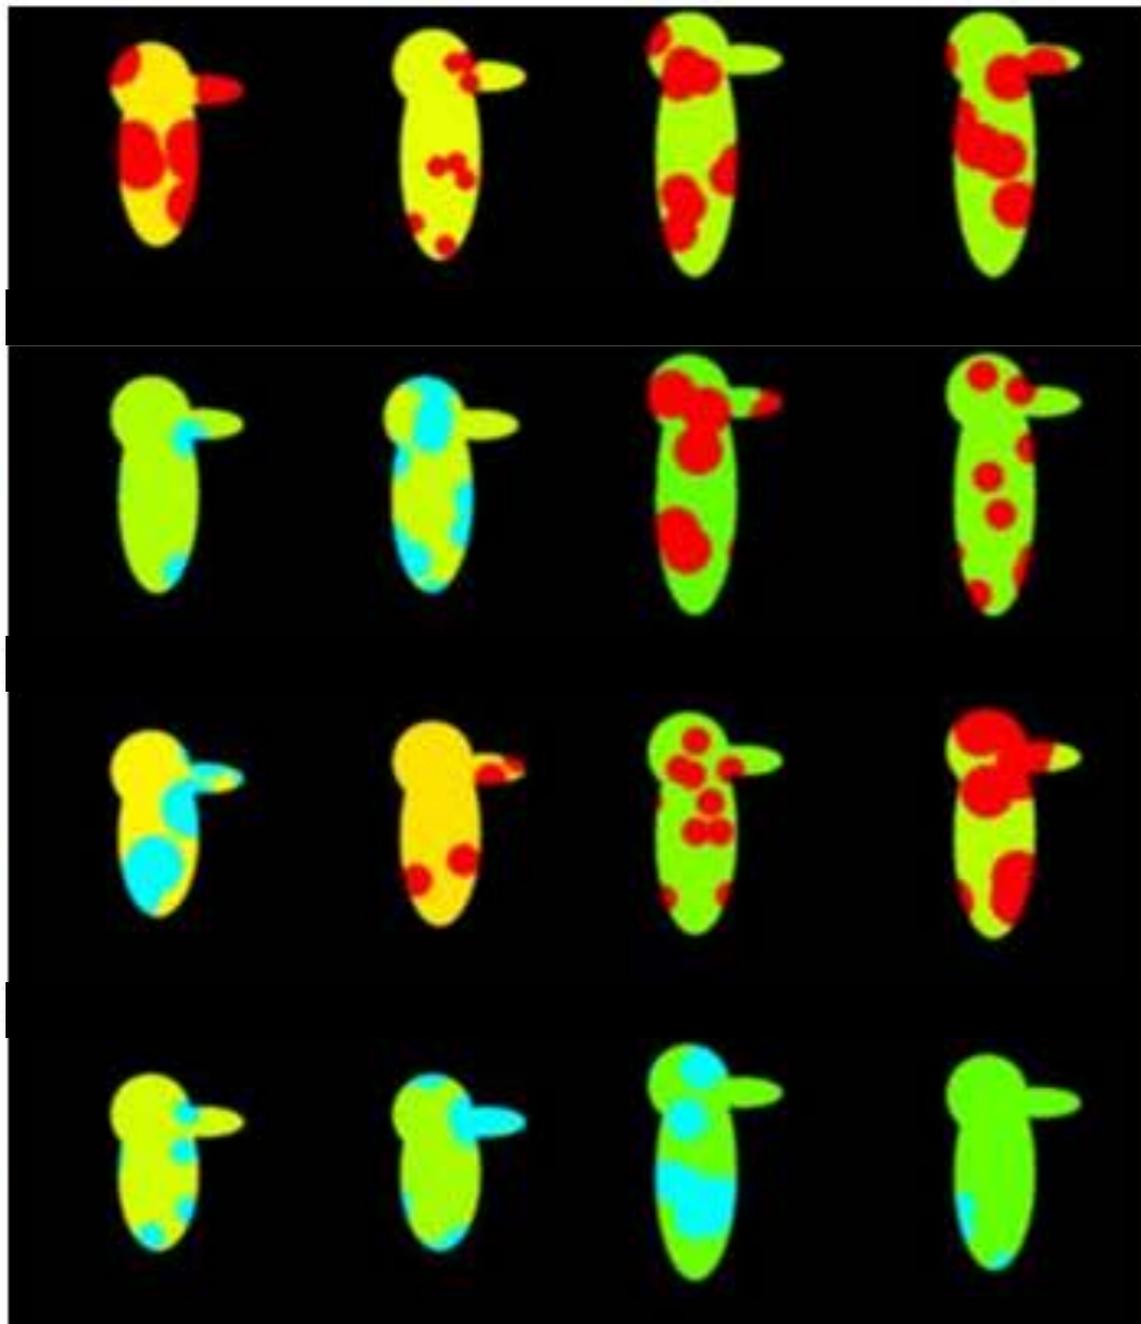
$\tau_j$



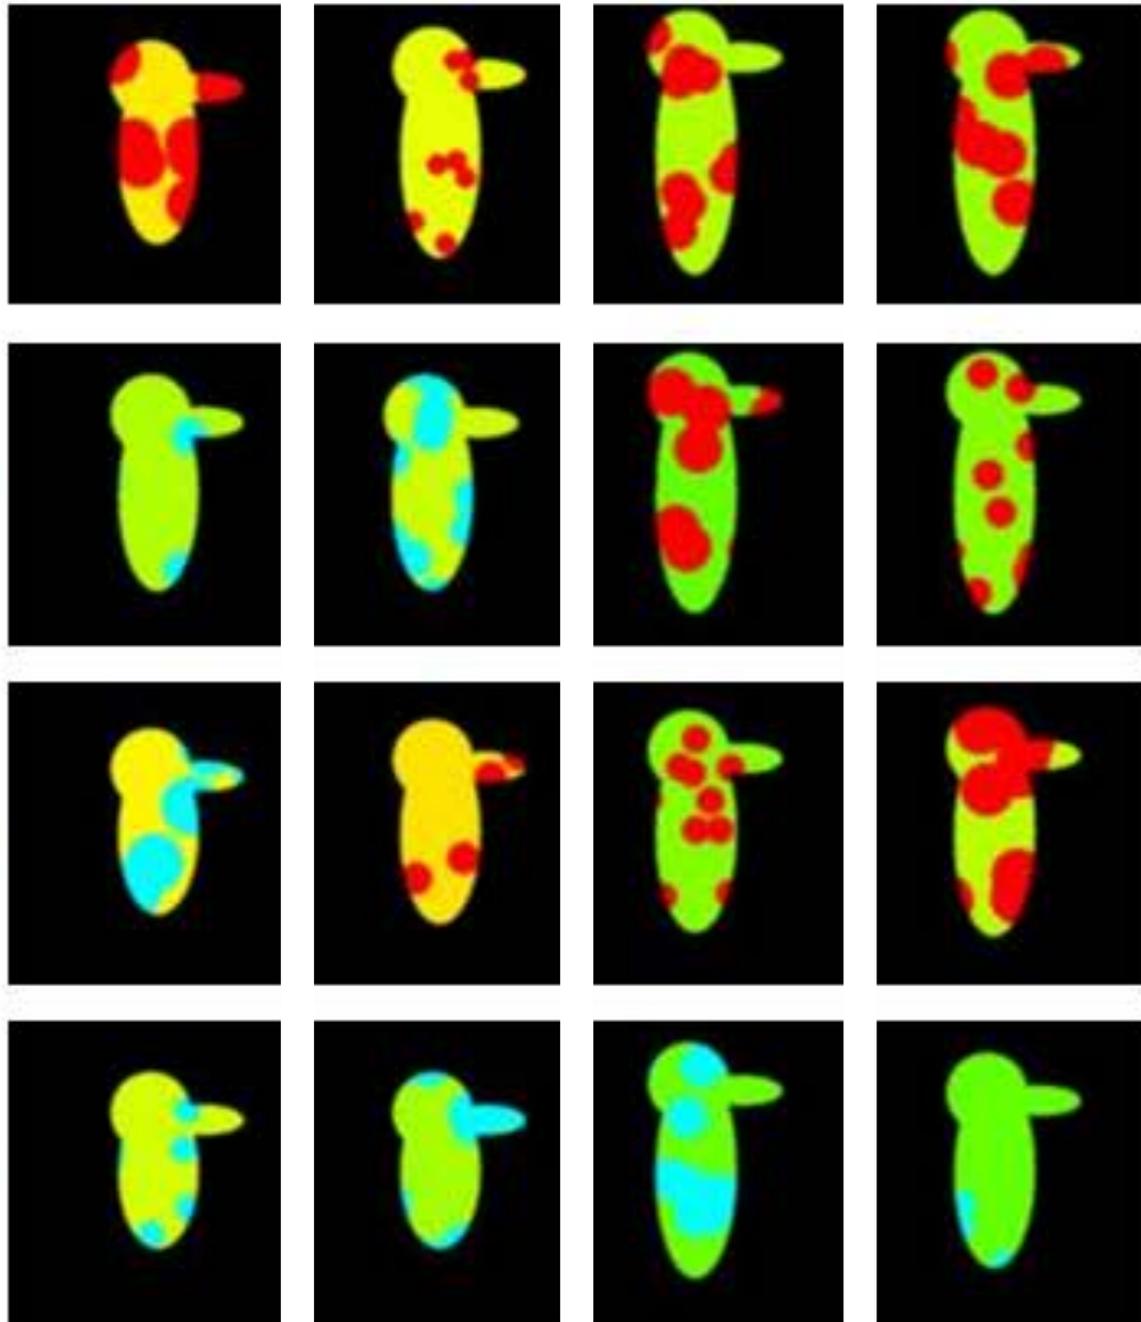
$1/\sigma_j$

# Greebles: estimating multiple attributes

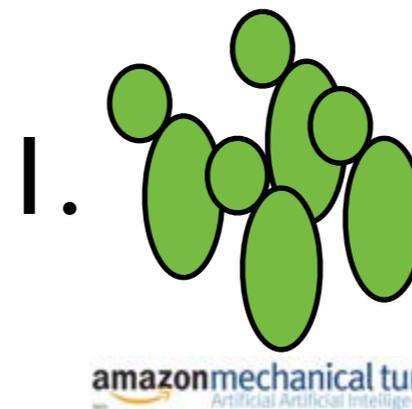
Synthetic figures distinguished by height and color



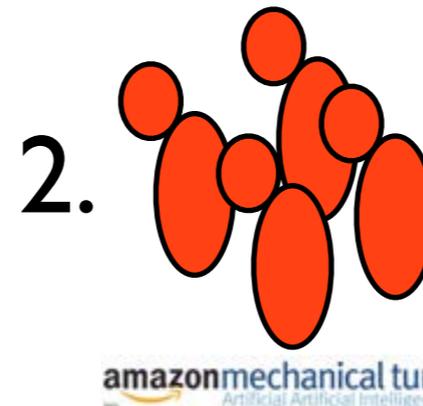
# Greebles: estimating multiple attributes



Send to two groups of annotators

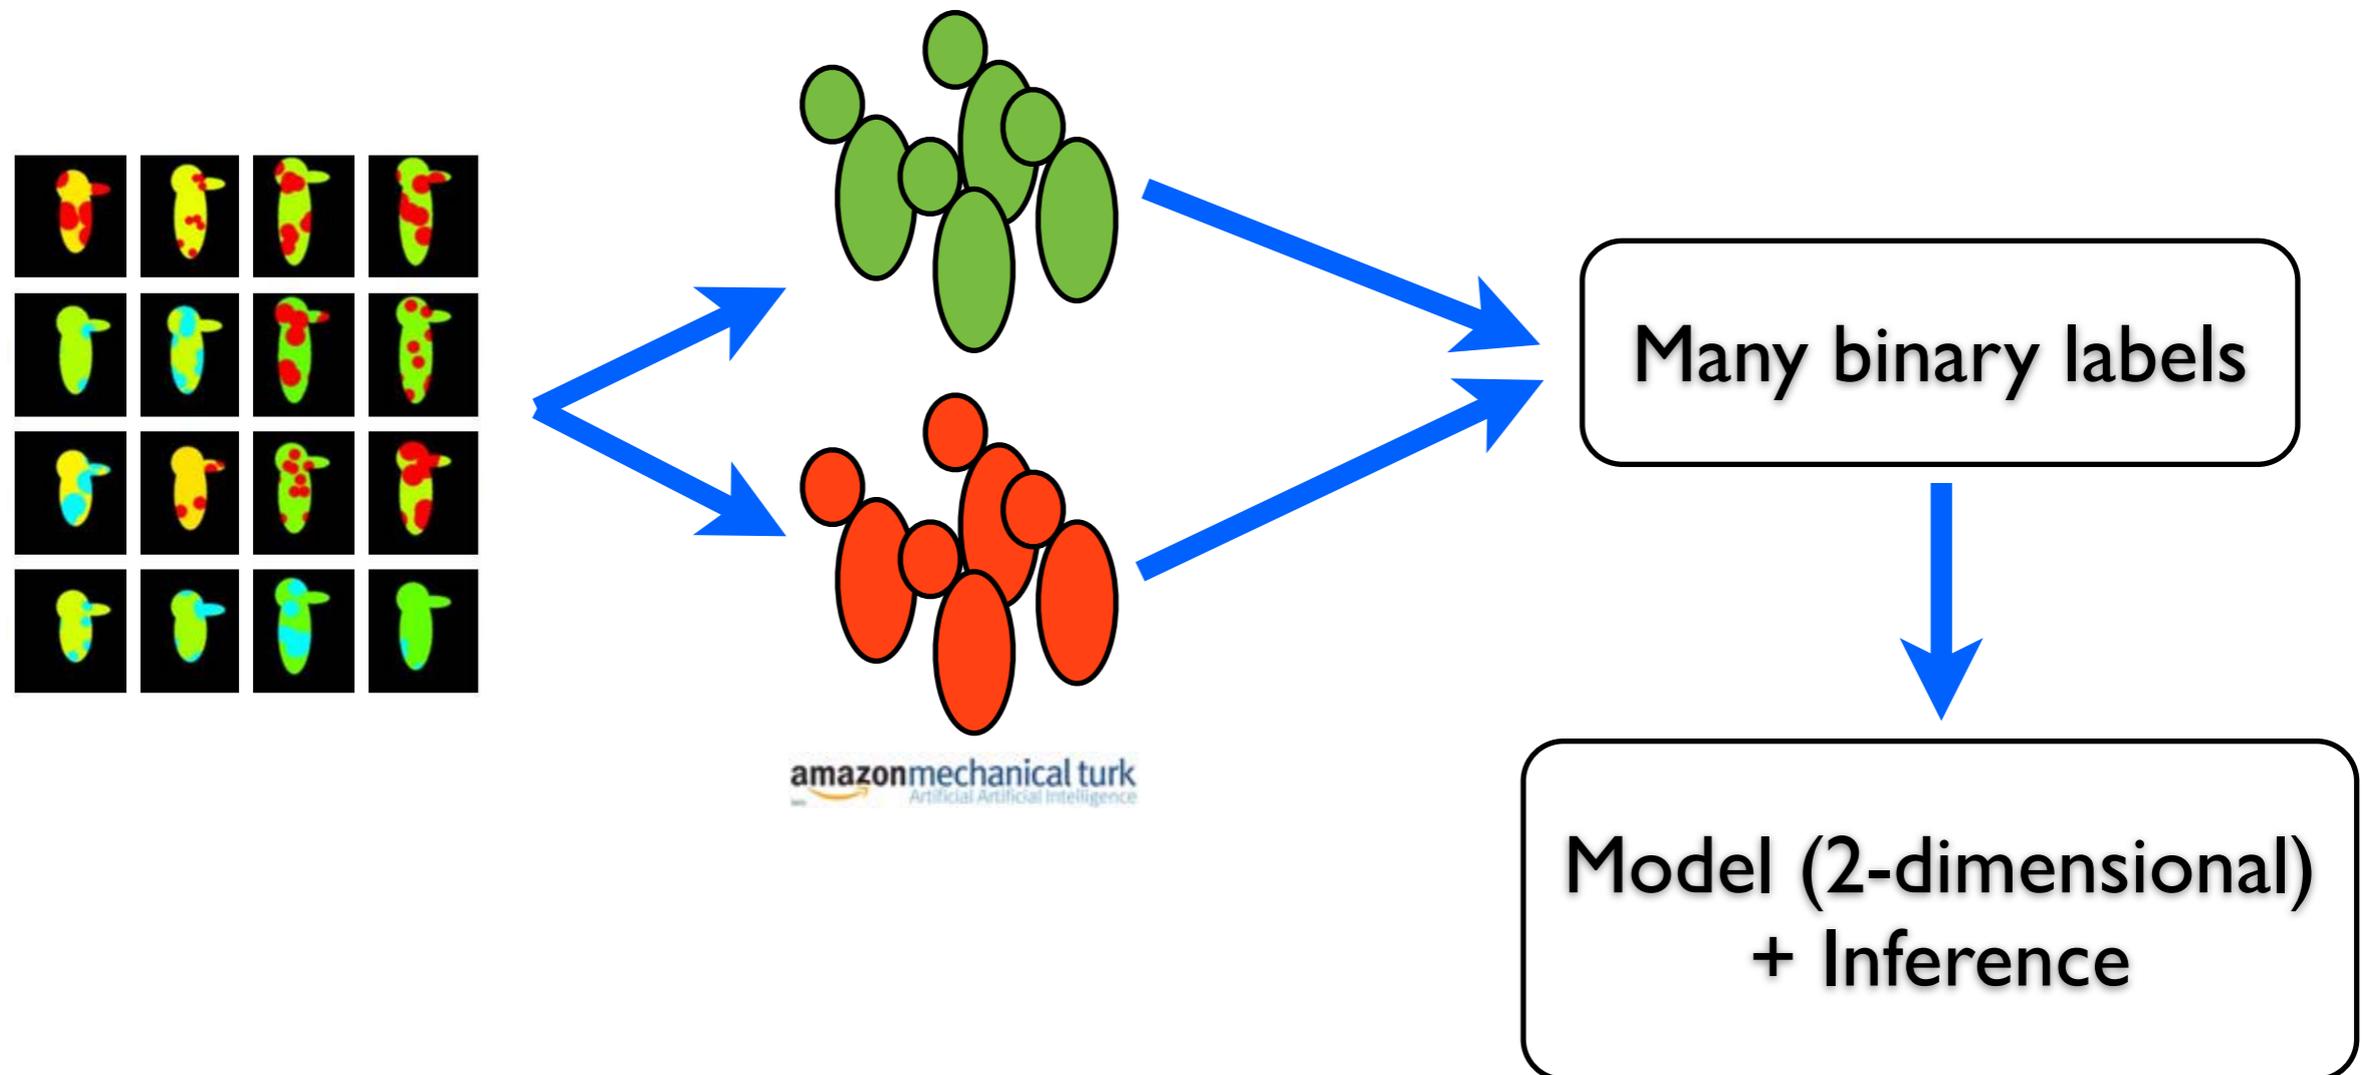


Select green greebles

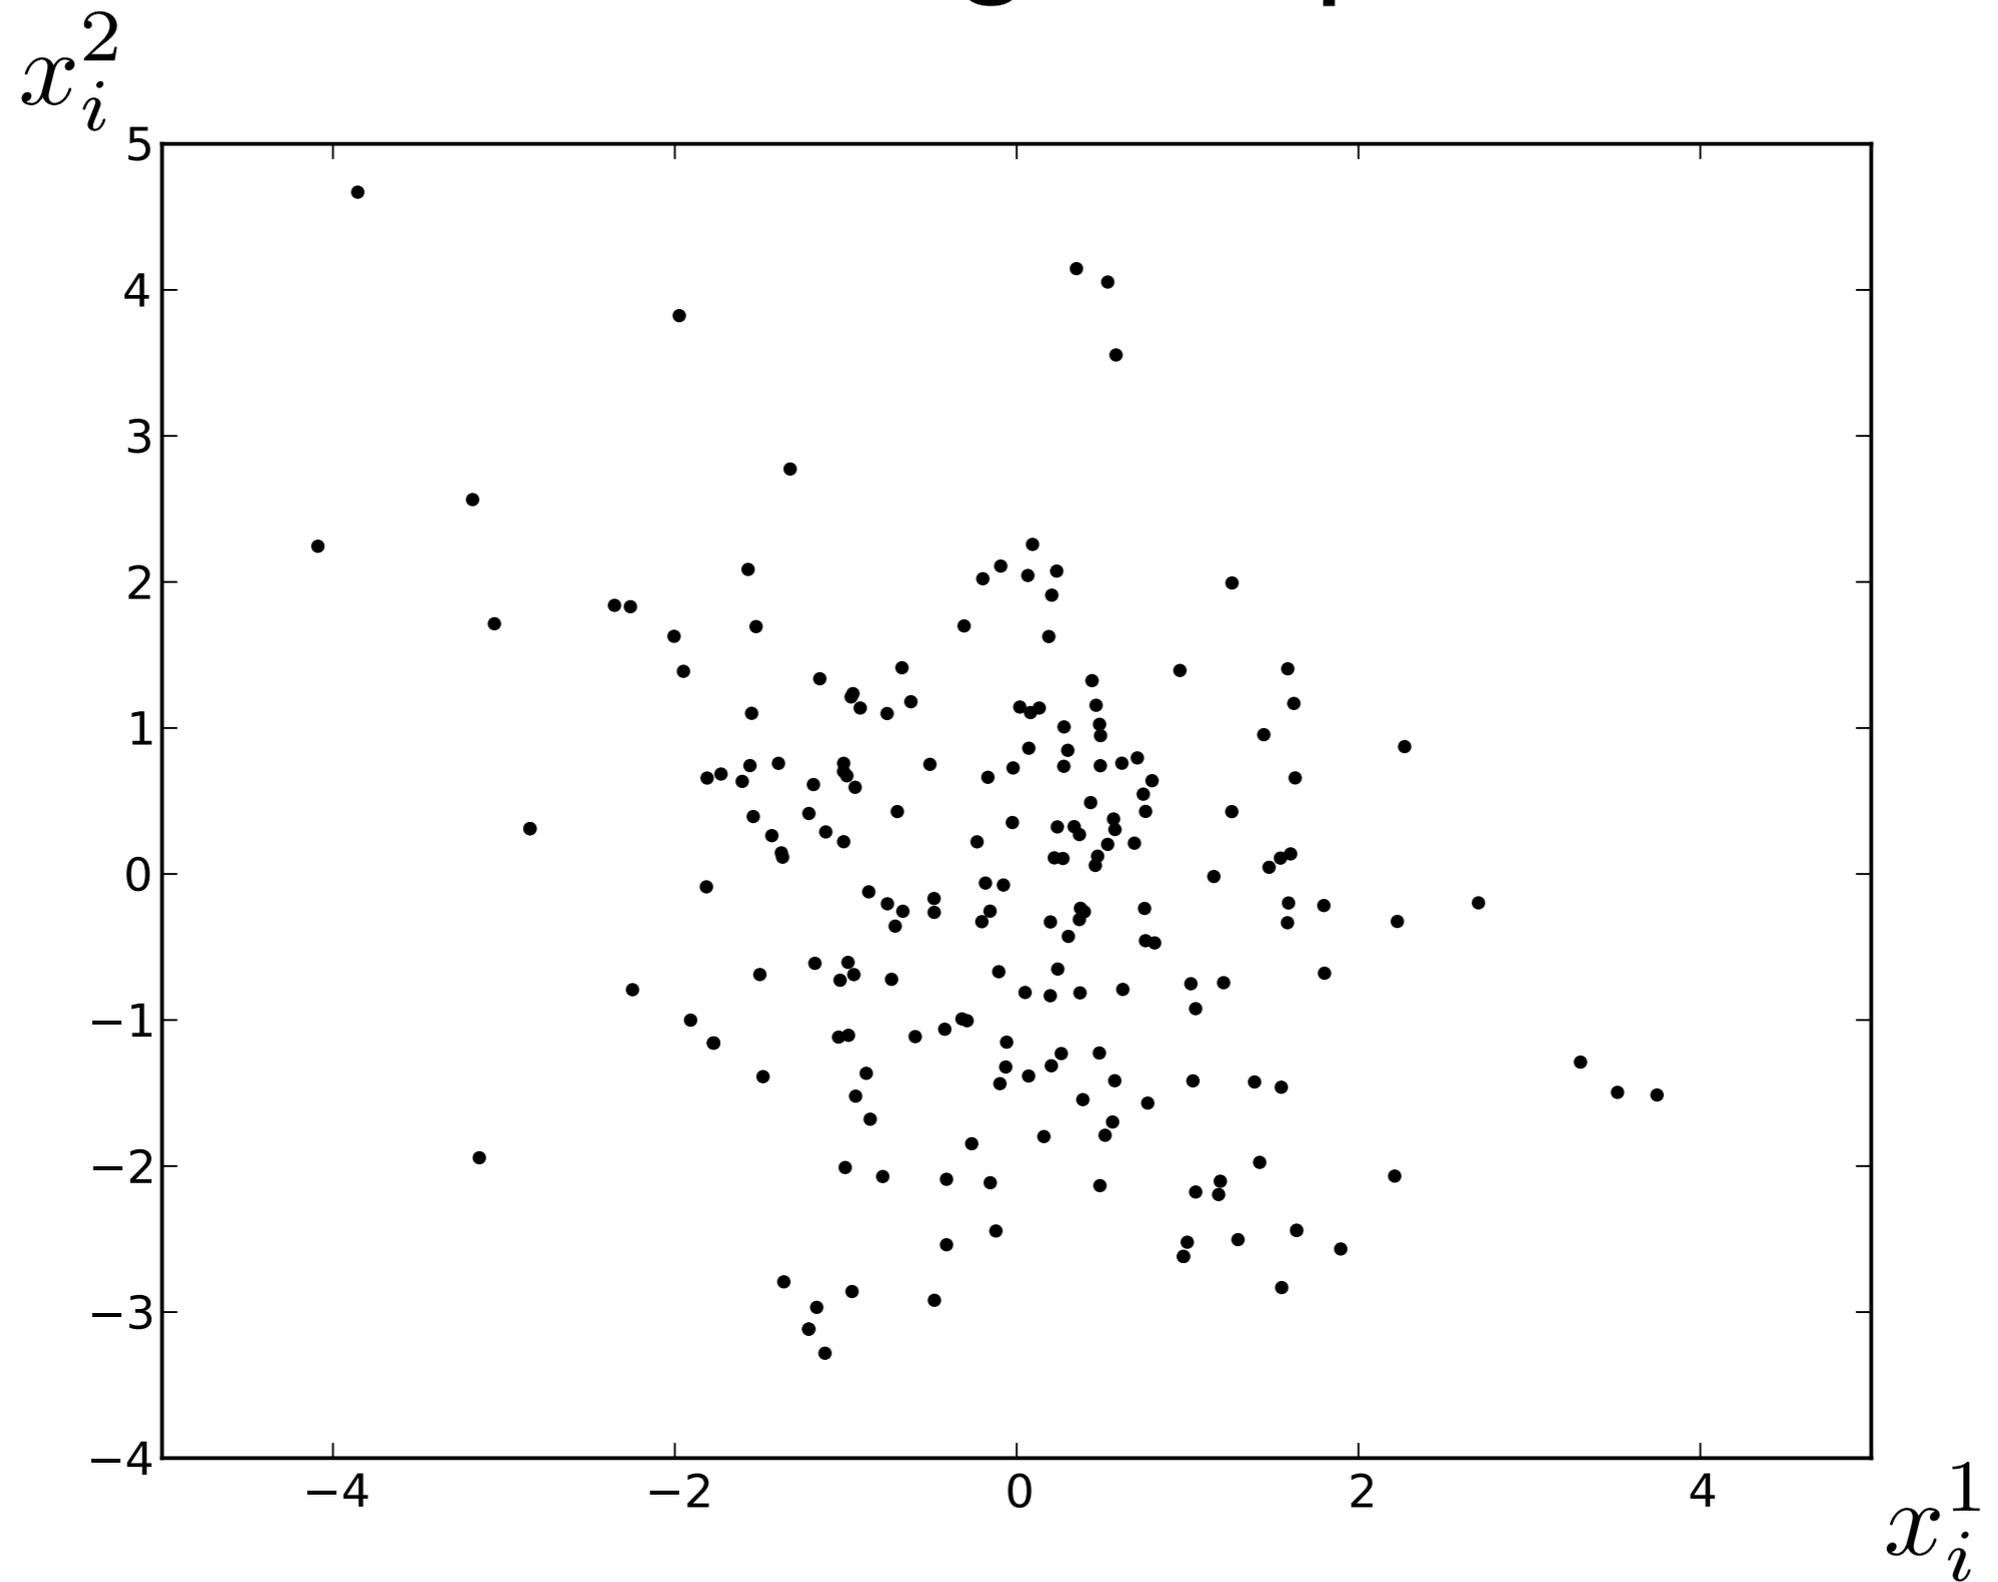


Select tall greebles

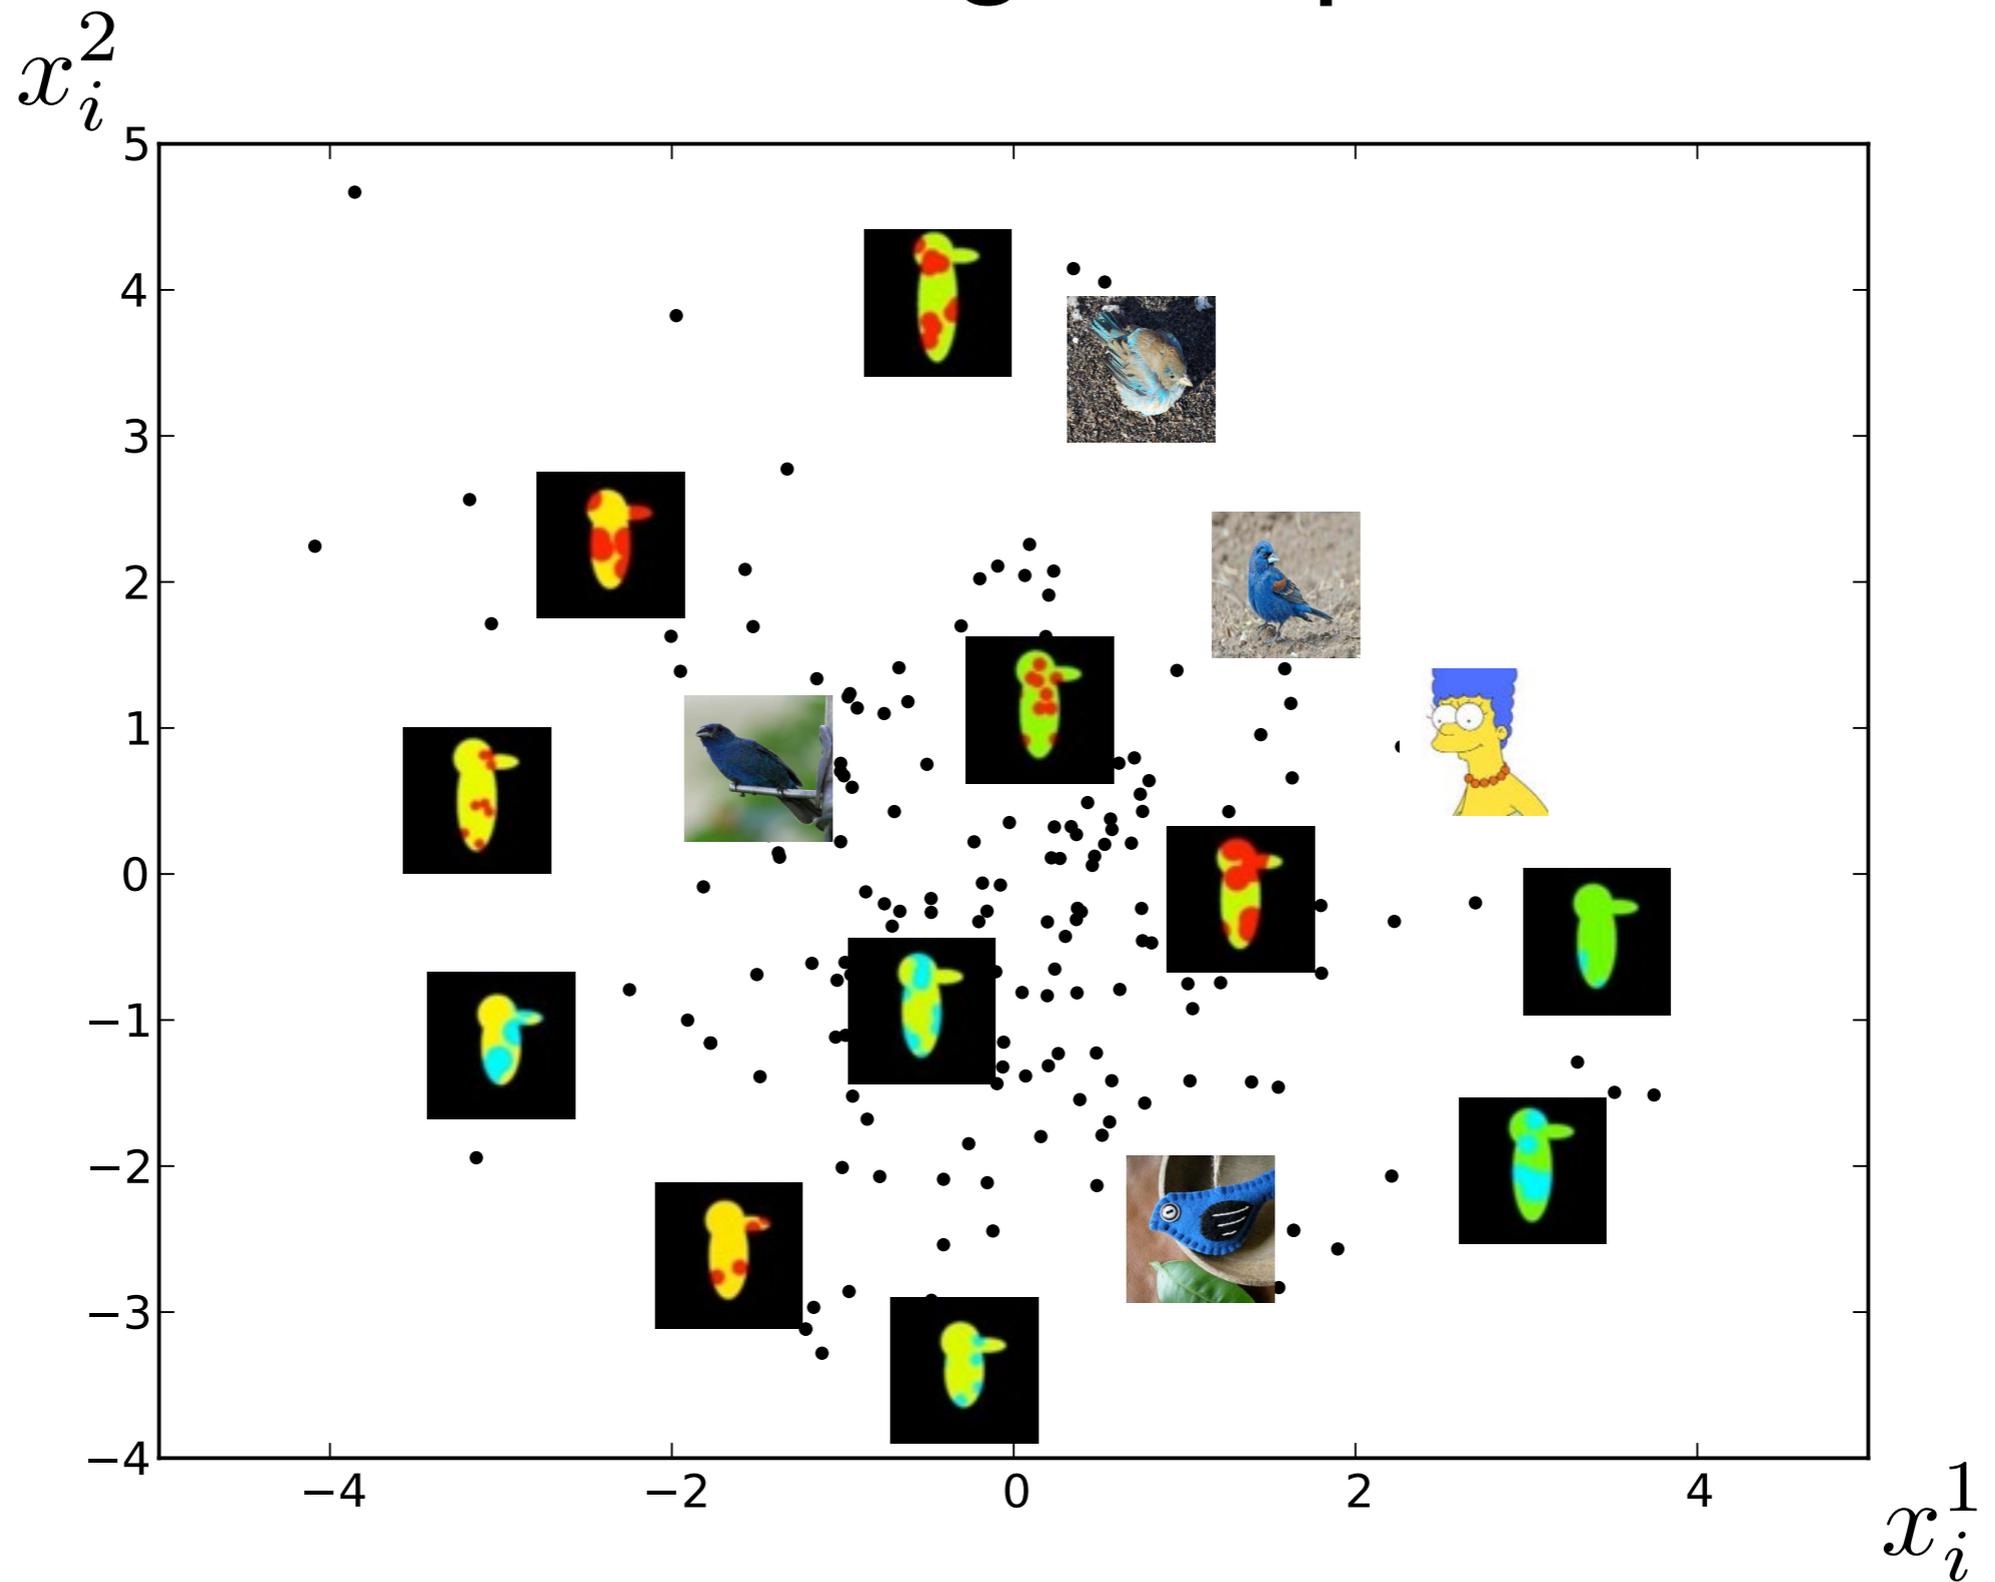
# Greebles: estimating multiple attributes



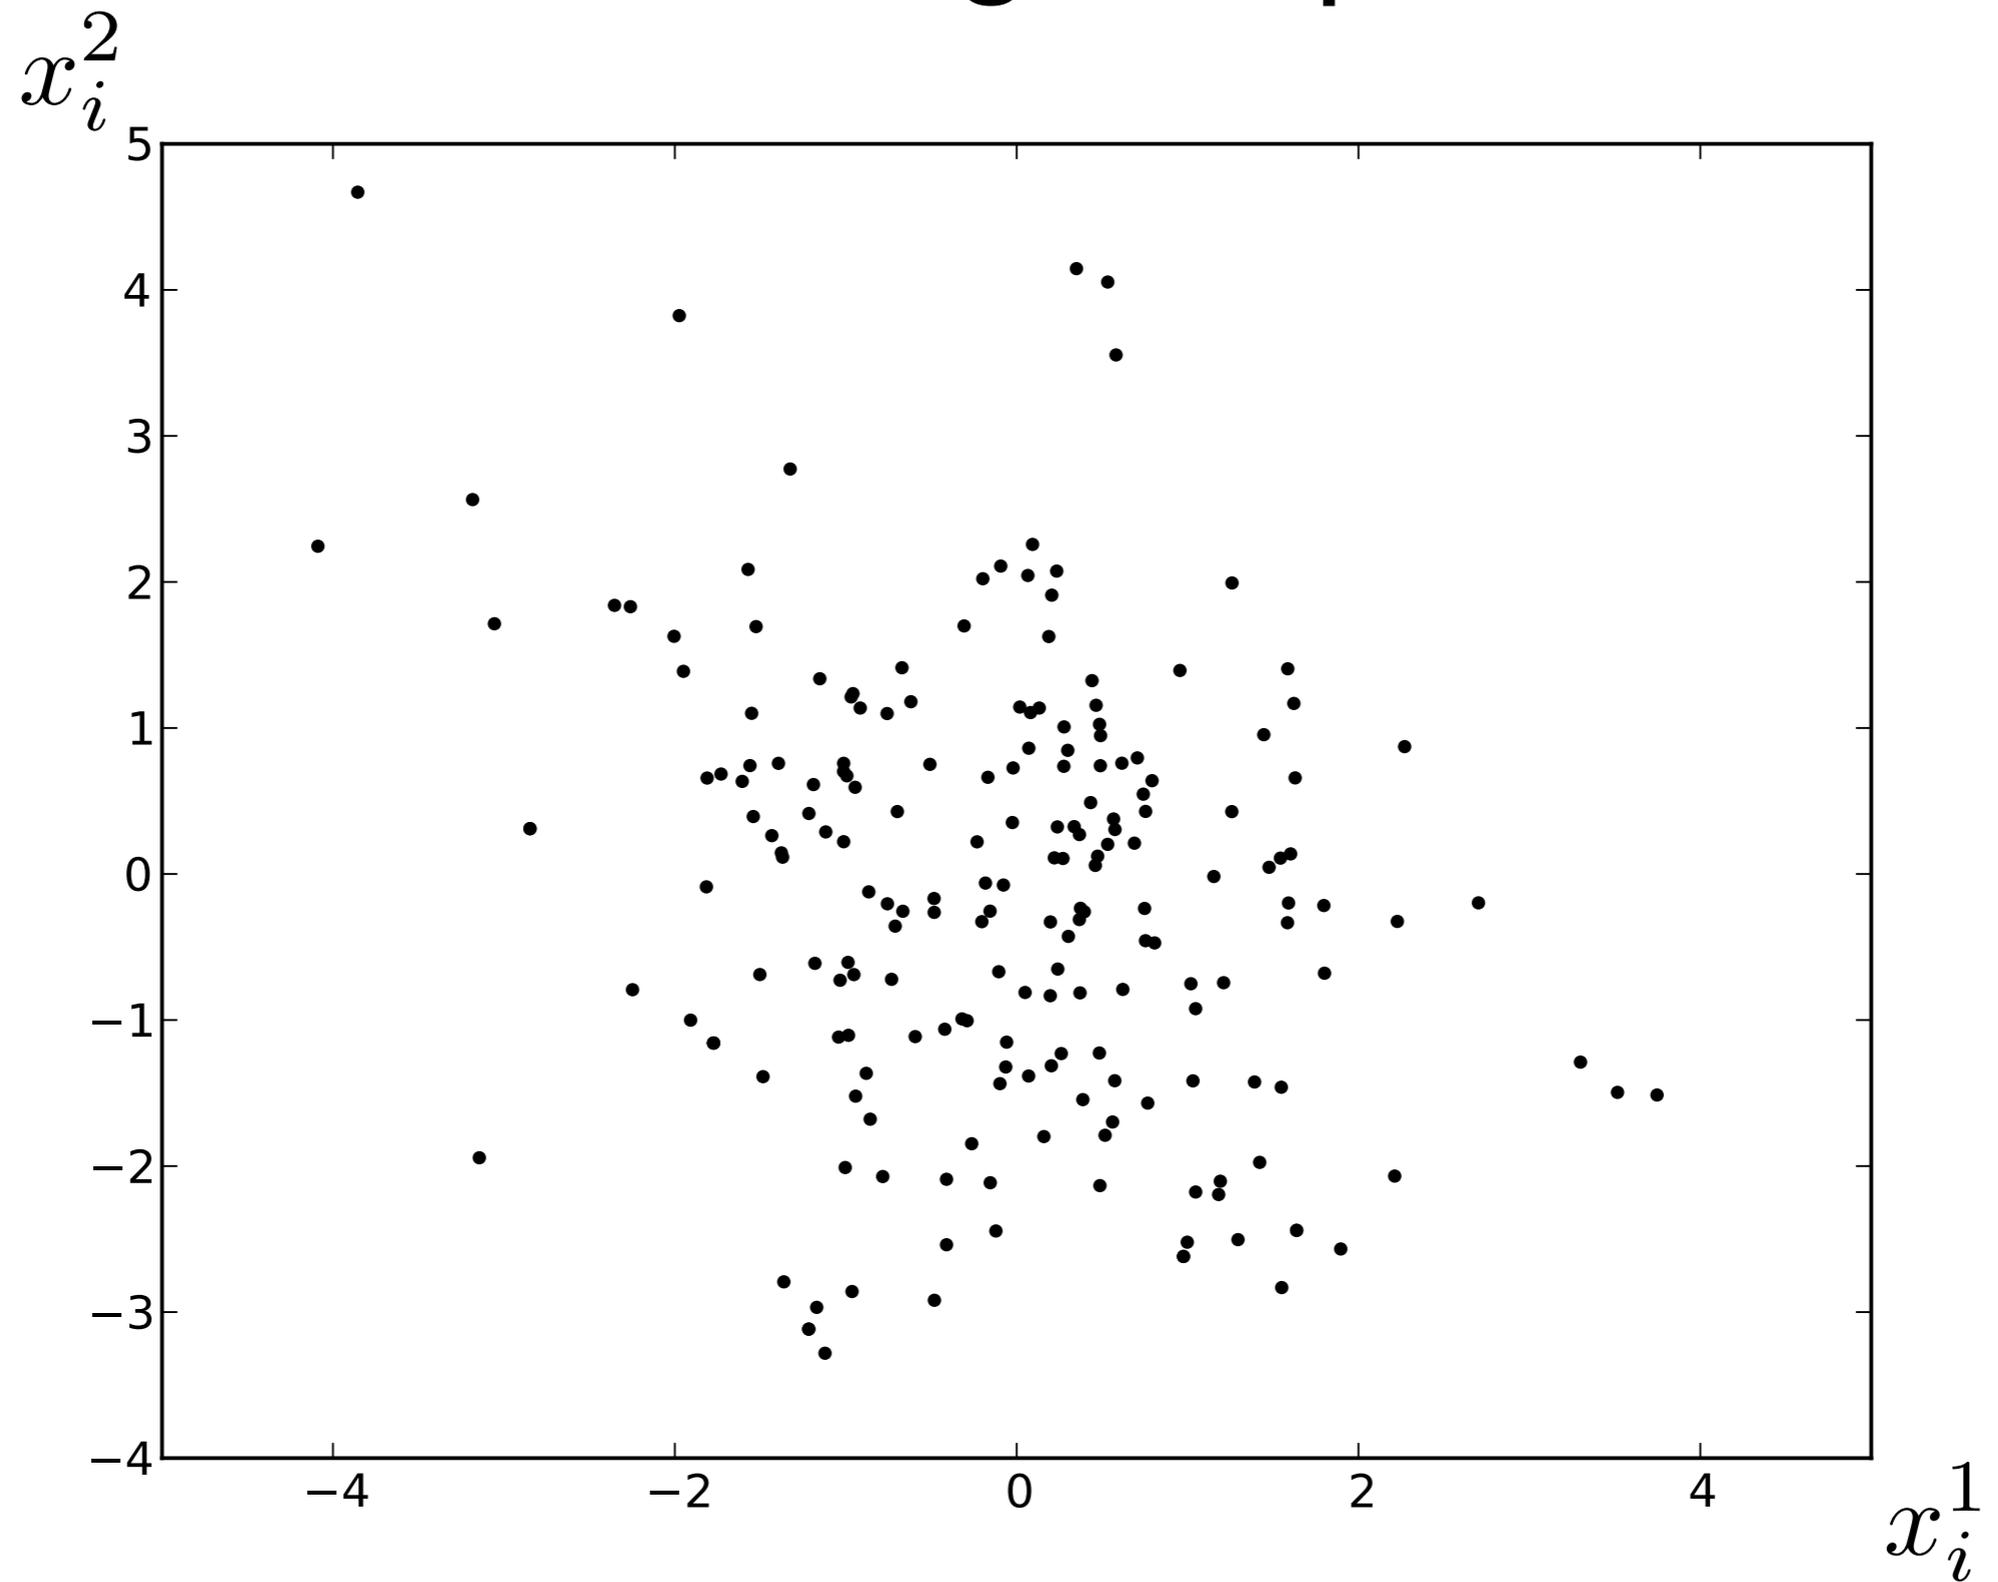
# Greebles: estimating multiple attributes



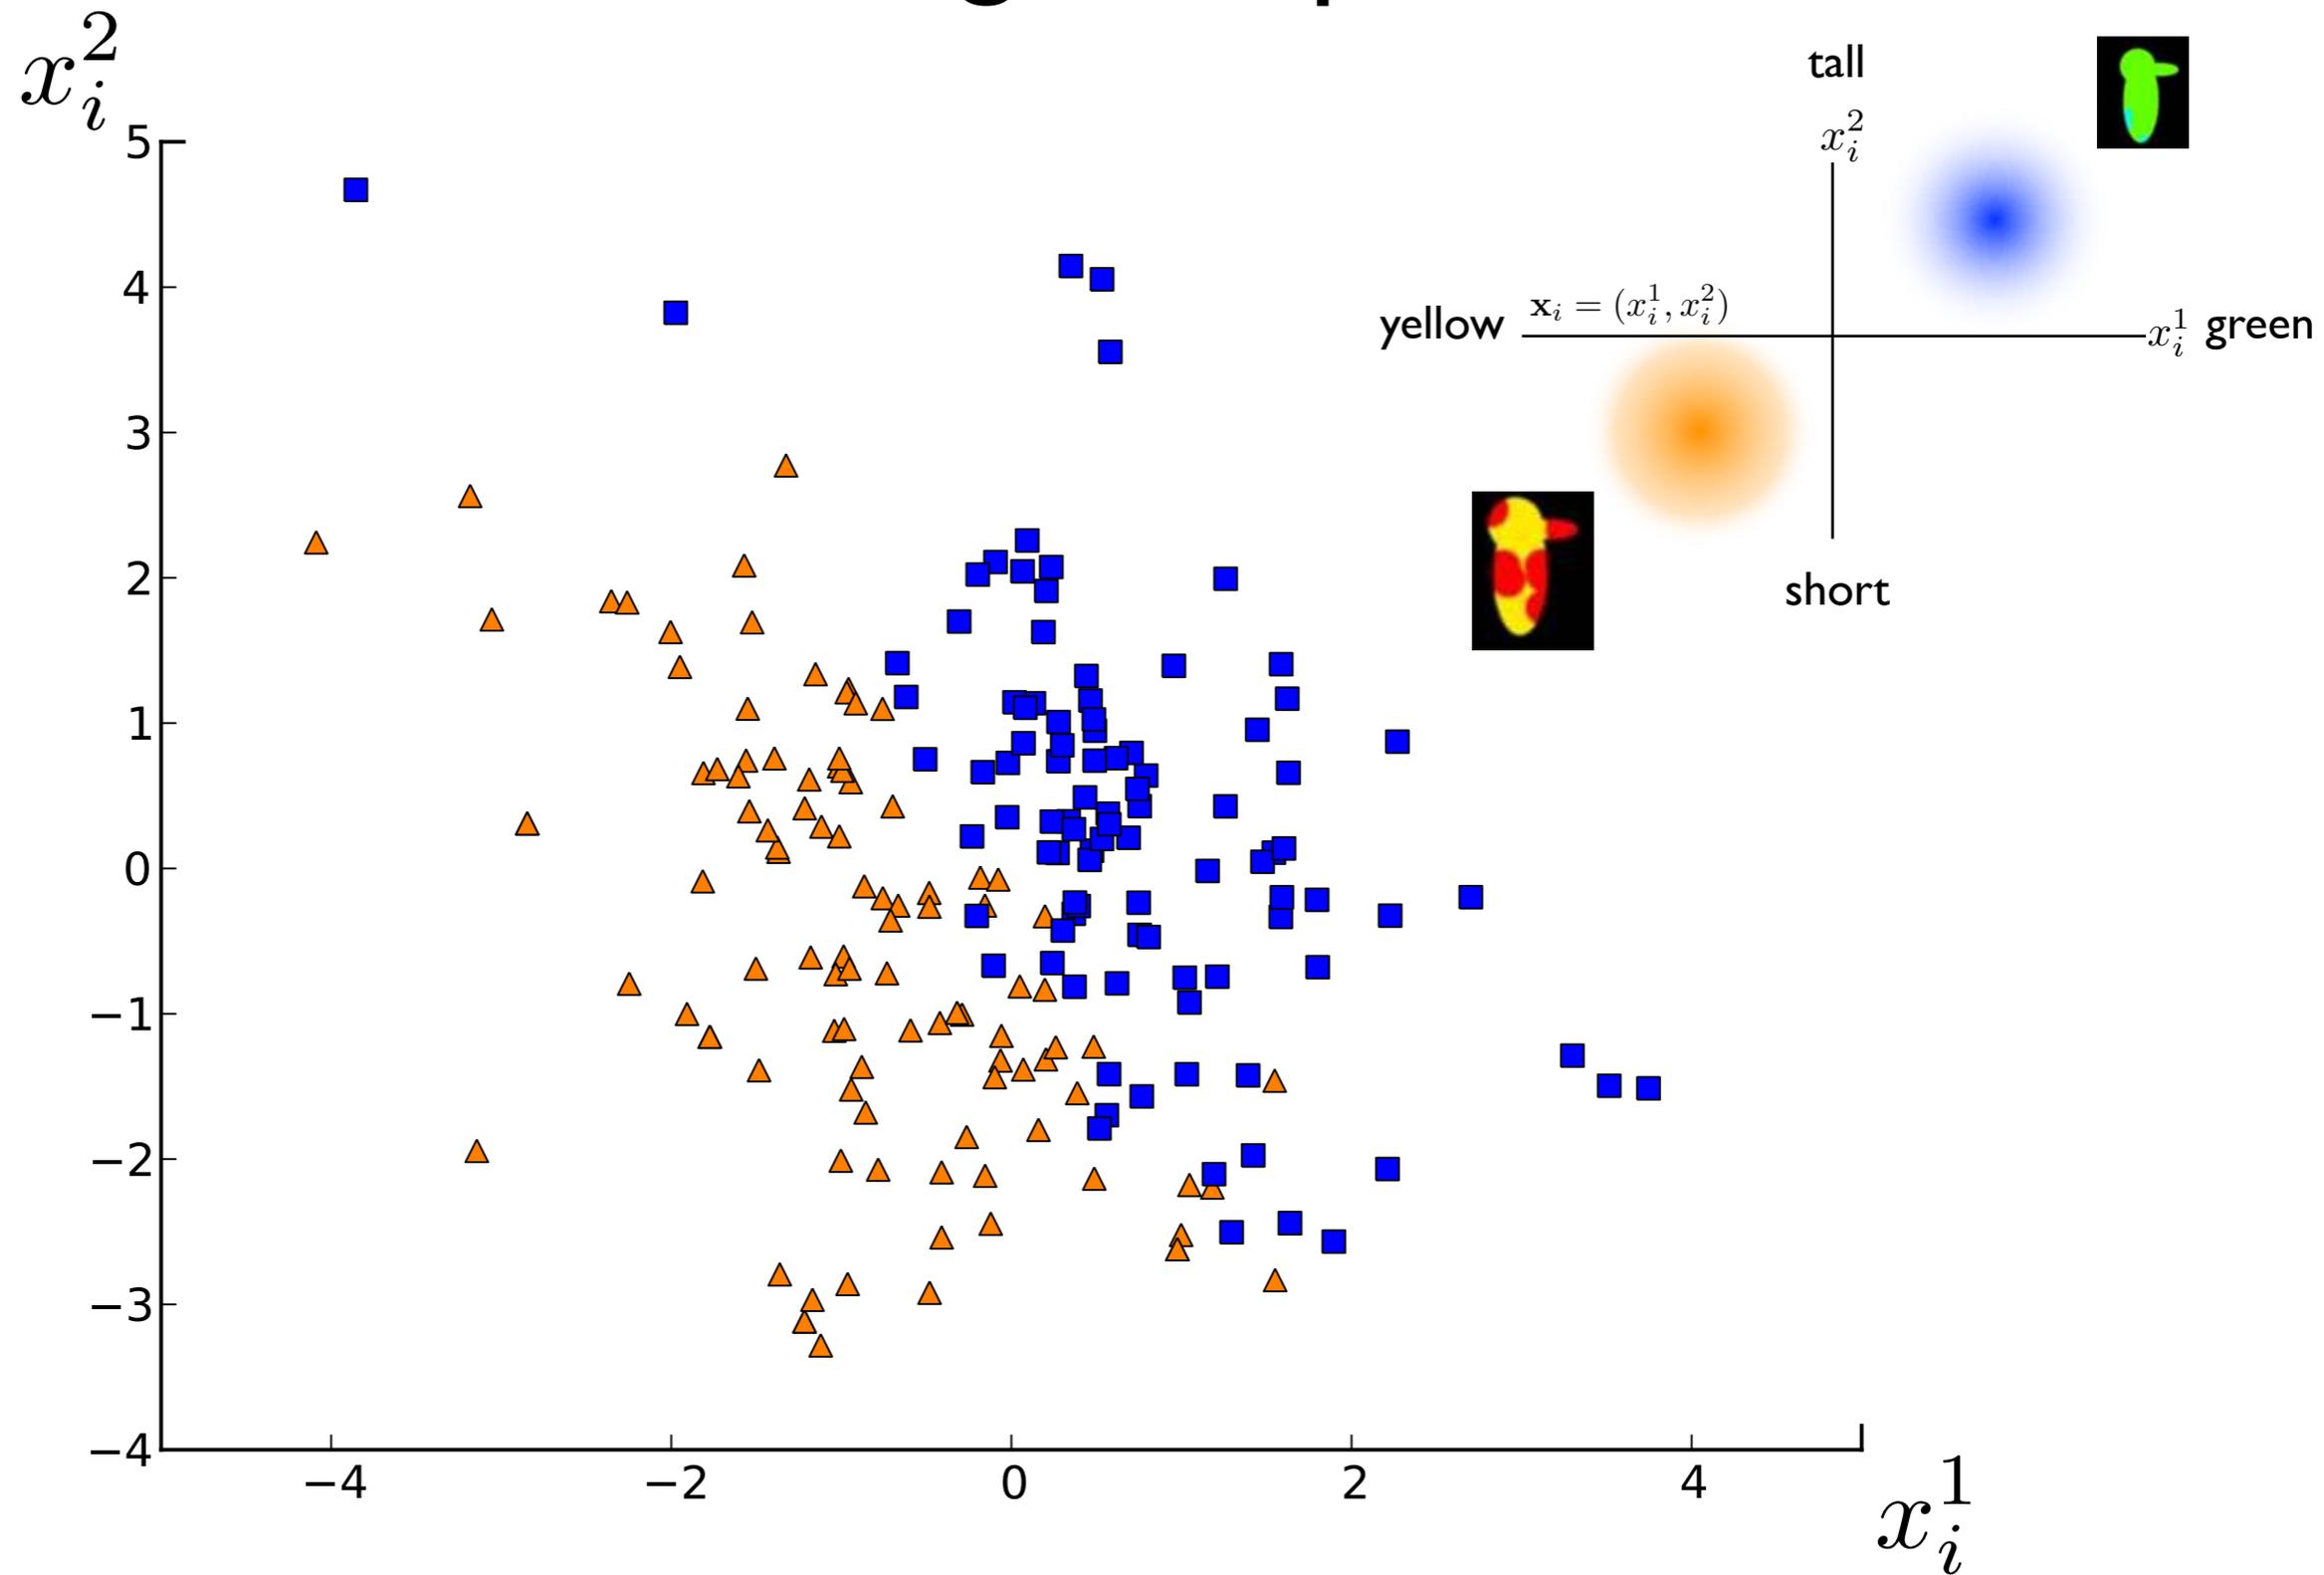
# Greebles: estimating multiple attributes



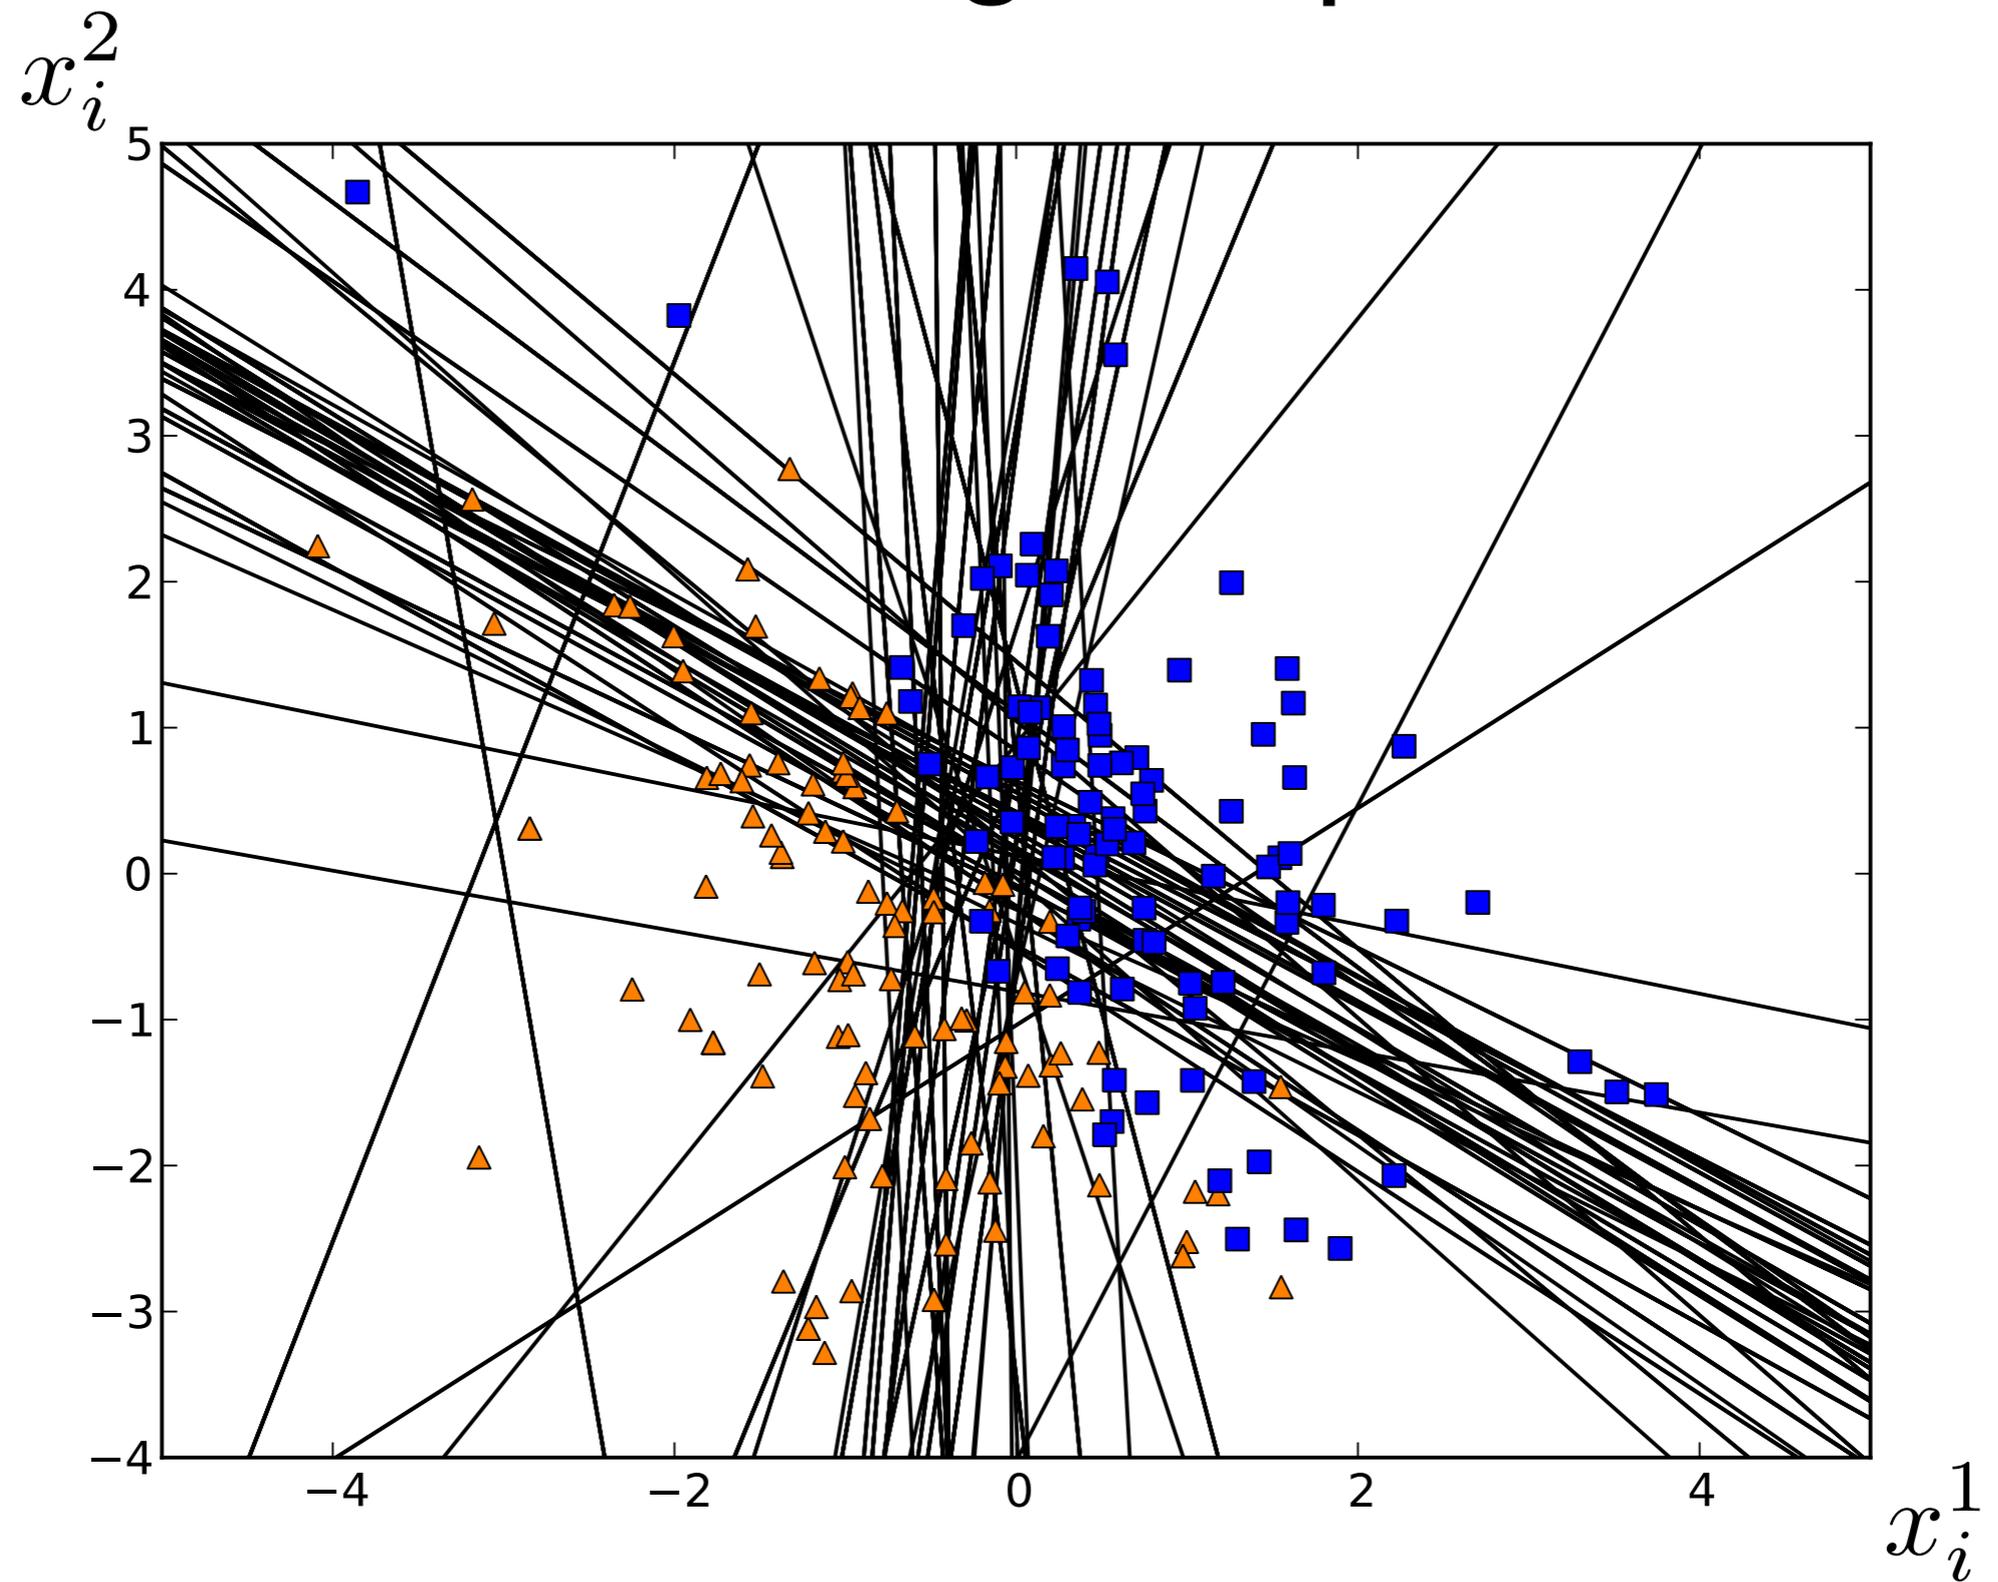
# Greebles: estimating multiple attributes



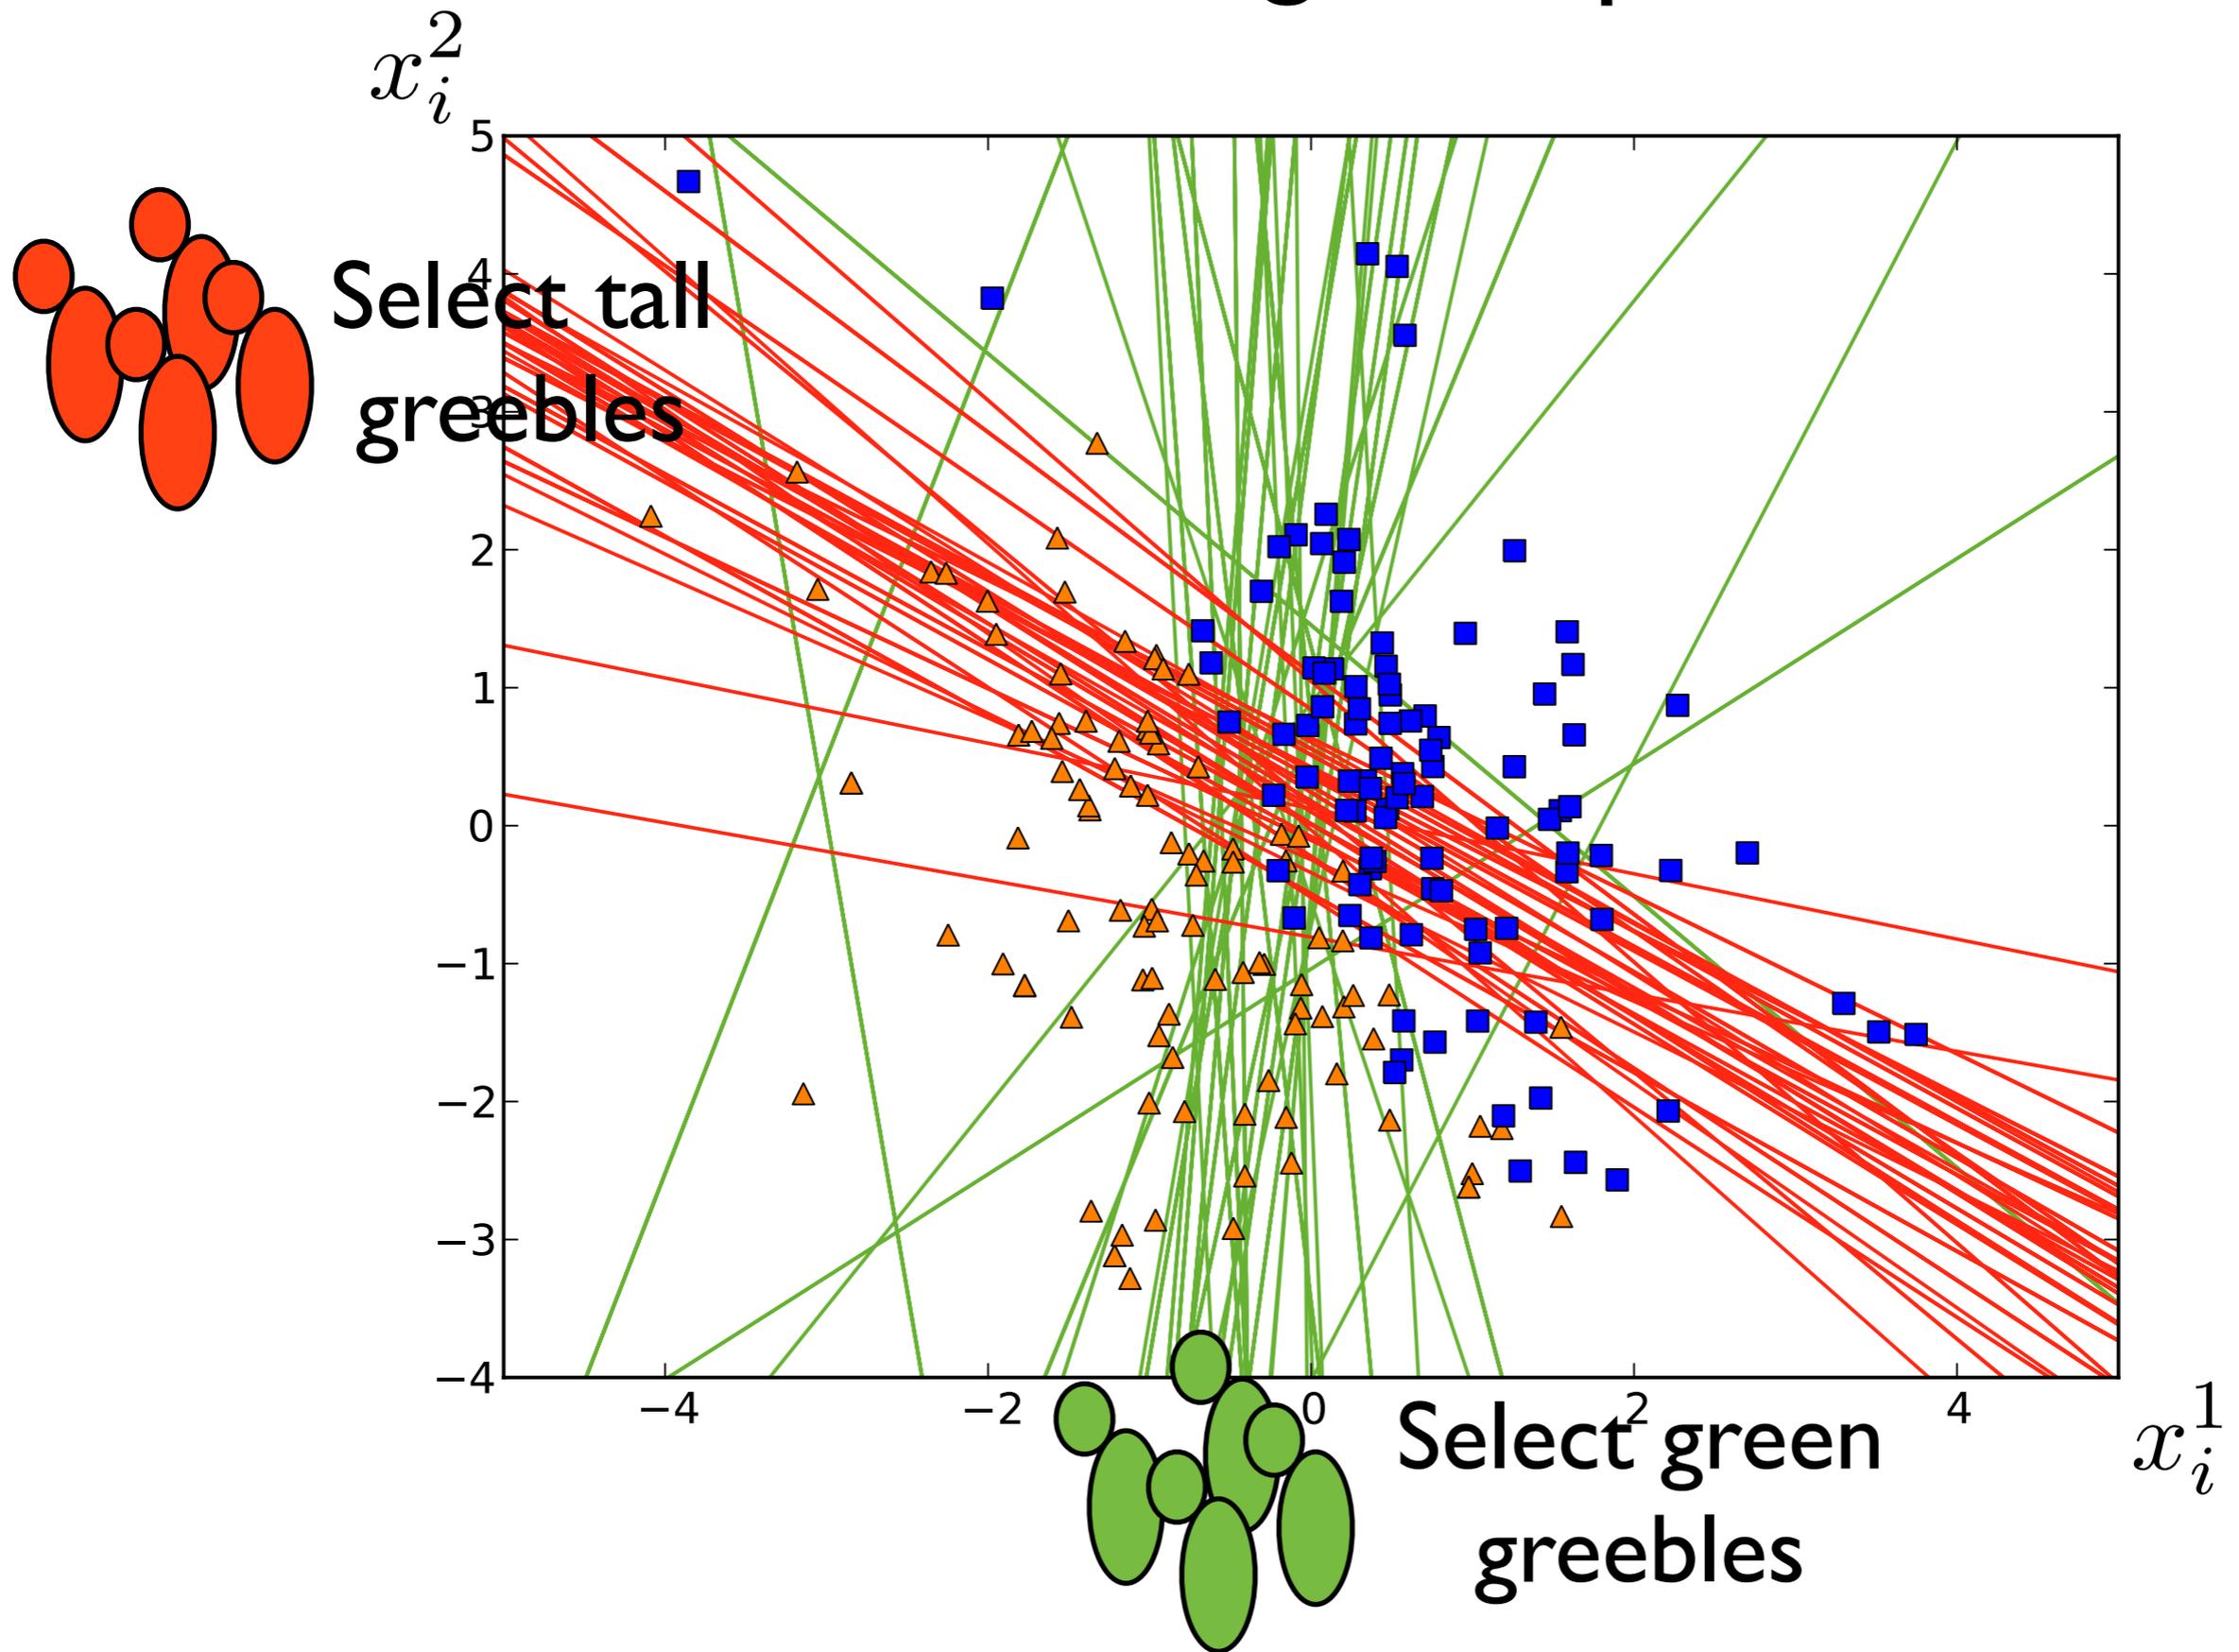
# Greebles: estimating multiple attributes



# Greebles: estimating multiple attributes



# Greebles: estimating multiple attributes



# Last Experiment: Waterbirds

## Mallard



## American Black Duck



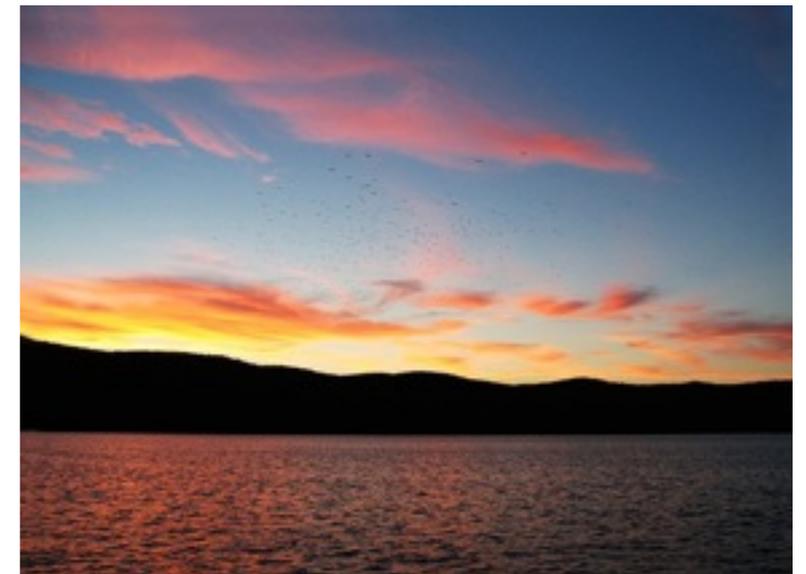
## Canada Goose



## Red Necked Grebe

A screenshot of a Wikipedia article titled "Mushroom". The article includes a navigation sidebar on the left with links like "Main page", "Contents", and "Search". The main text defines a mushroom as the fleshy, spore-bearing fruiting body of a fungus. An image of a red mushroom with white spots is included, with a caption: "The mushroom *Amanita muscaria*, commonly known as 'fly agaric'".

## Non-bird



# DUCKS

Mallard



American Black Duck



Canada Goose



Red Necked Grebe

Learn more about citing Wikipedia. Try Beta Log in / create account

article discussion view source history

## Mushroom

From Wikipedia, the free encyclopedia

*For other uses, see Mushroom (disambiguation).*

A **mushroom** is the fleshy, *spore-bearing fruiting body* of a *fungus*, typically produced above ground on soil or on its food source. The standard for the name "mushroom" is the cultivated white button mushroom, *Agaricus bisporus*, hence the word mushroom is most often applied to those fungi (Basidiomycota, Agaricomycetes) that have a stem (*stipe*), a cap (*pileus*), and gills (lamellae, sing. lamella) on the underside of the cap, just as do store-bought white mushrooms.

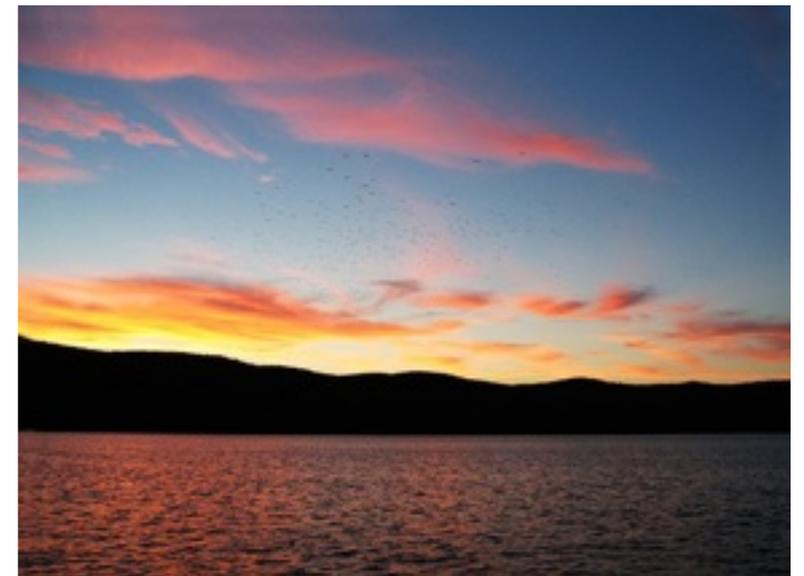
The word "mushroom" can also be used for a wide variety of gilled fungi, with or without stems, and the term is used even more generally, to describe both the fleshy fruiting bodies of some *Ascomycota* and the woody or leathery fruiting bodies of some *Basidiomycota*, depending upon the context of the word.

Forms deviating from the standard morphology usually have more specific names, such as "puffball", "stinkhorn", and "morel", and gilled mushrooms themselves are often called "agarics" in reference to their similarity to *Agaricus* or their placement in the order *Agaricales*. By extension, the term "mushroom" can also designate the entire fungus when in culture or the thallus (called a *mycelium*) of species forming the fruiting bodies called mushrooms, or the species itself.

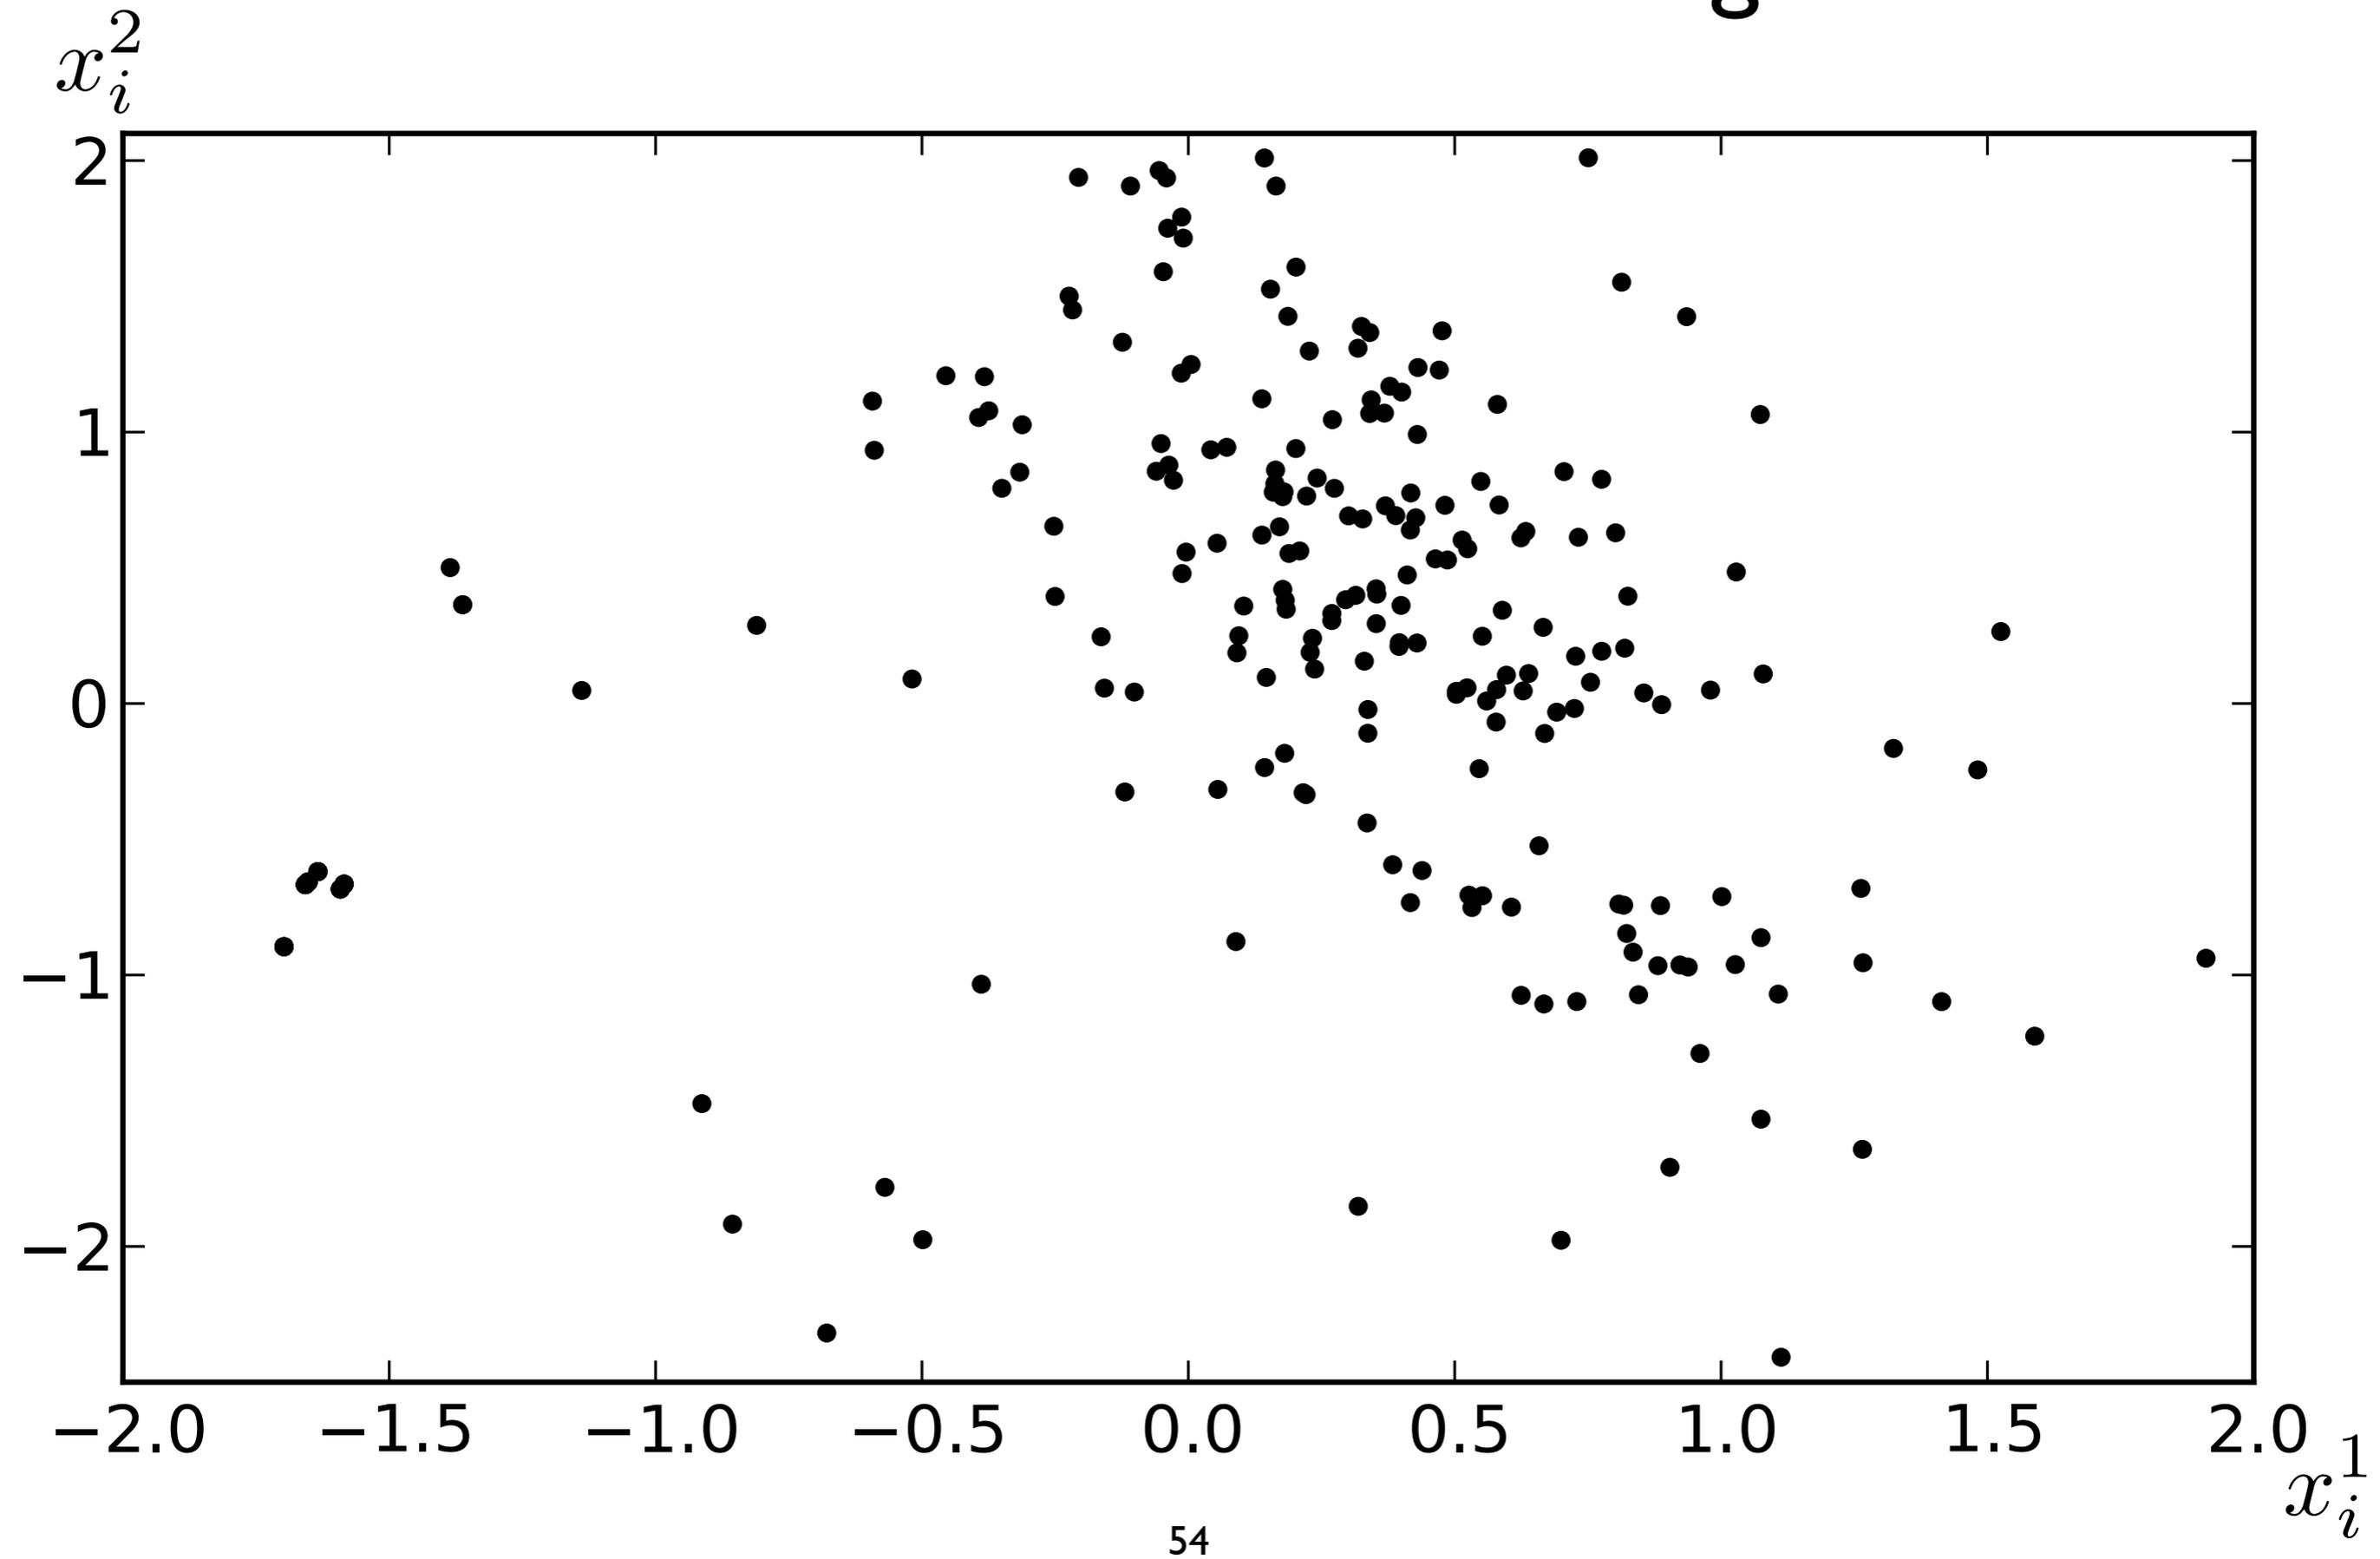


The mushroom *Amanita muscaria*, commonly known as "fly agaric"

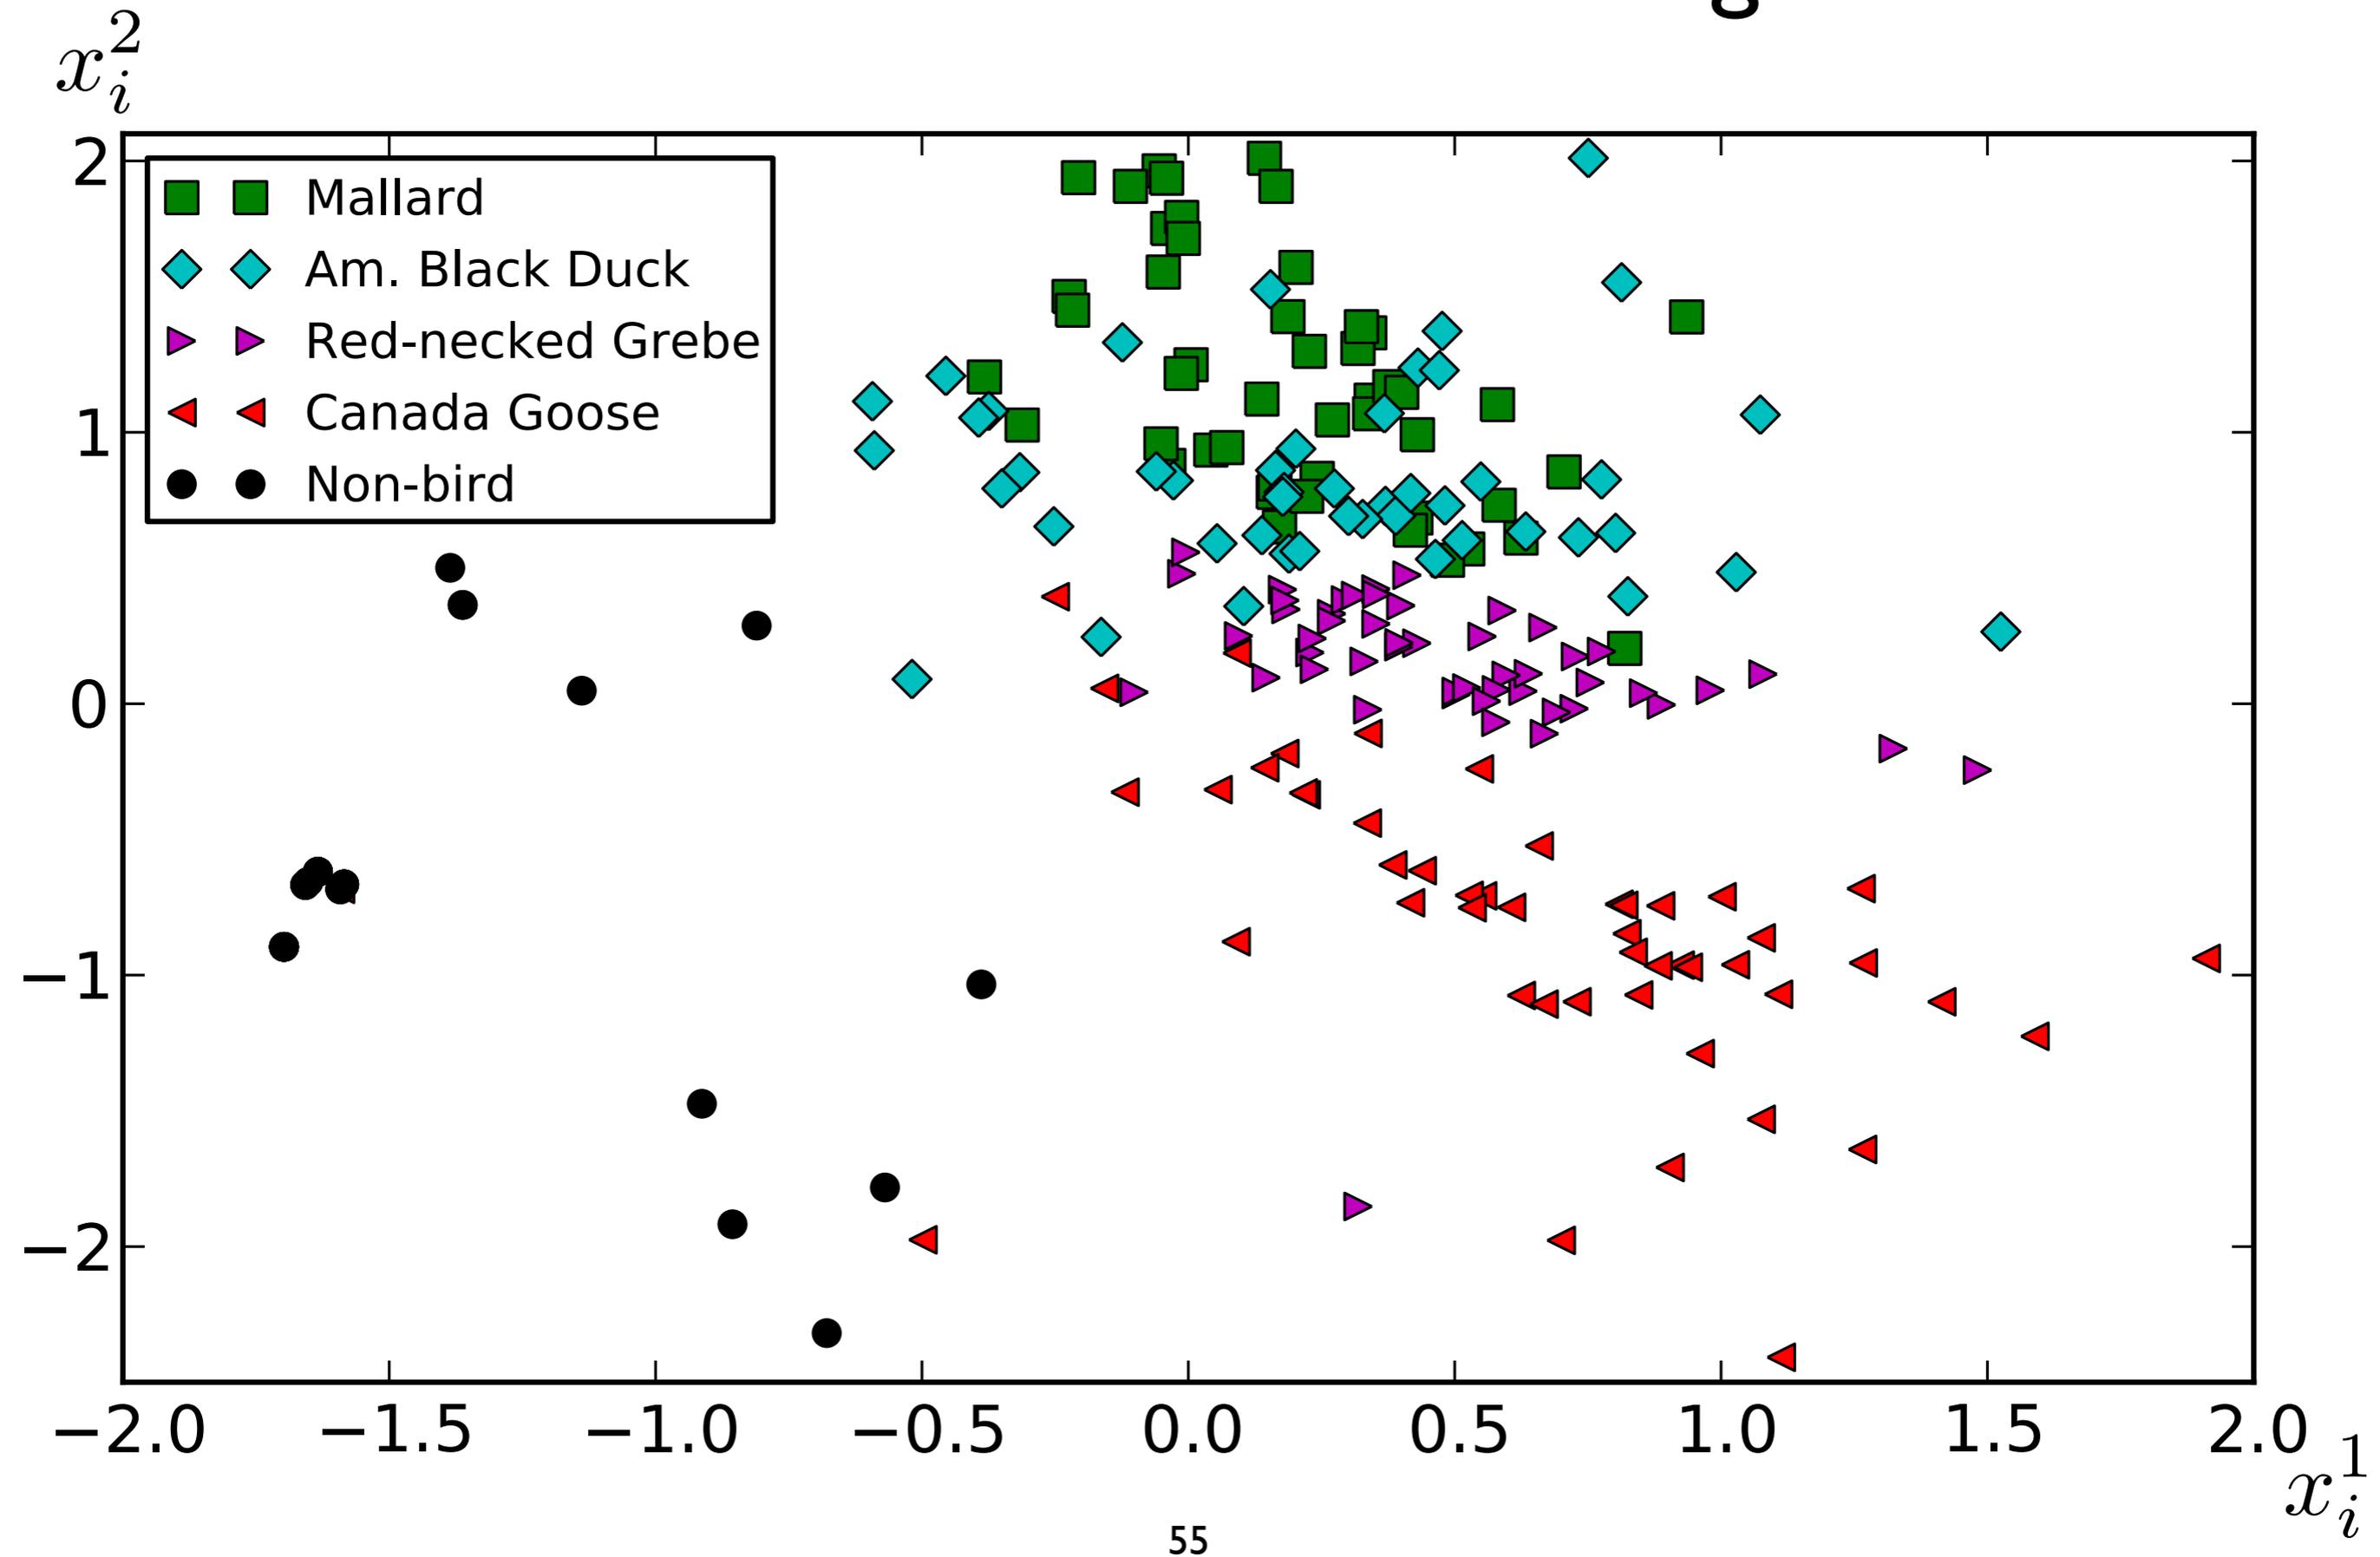
Non-bird



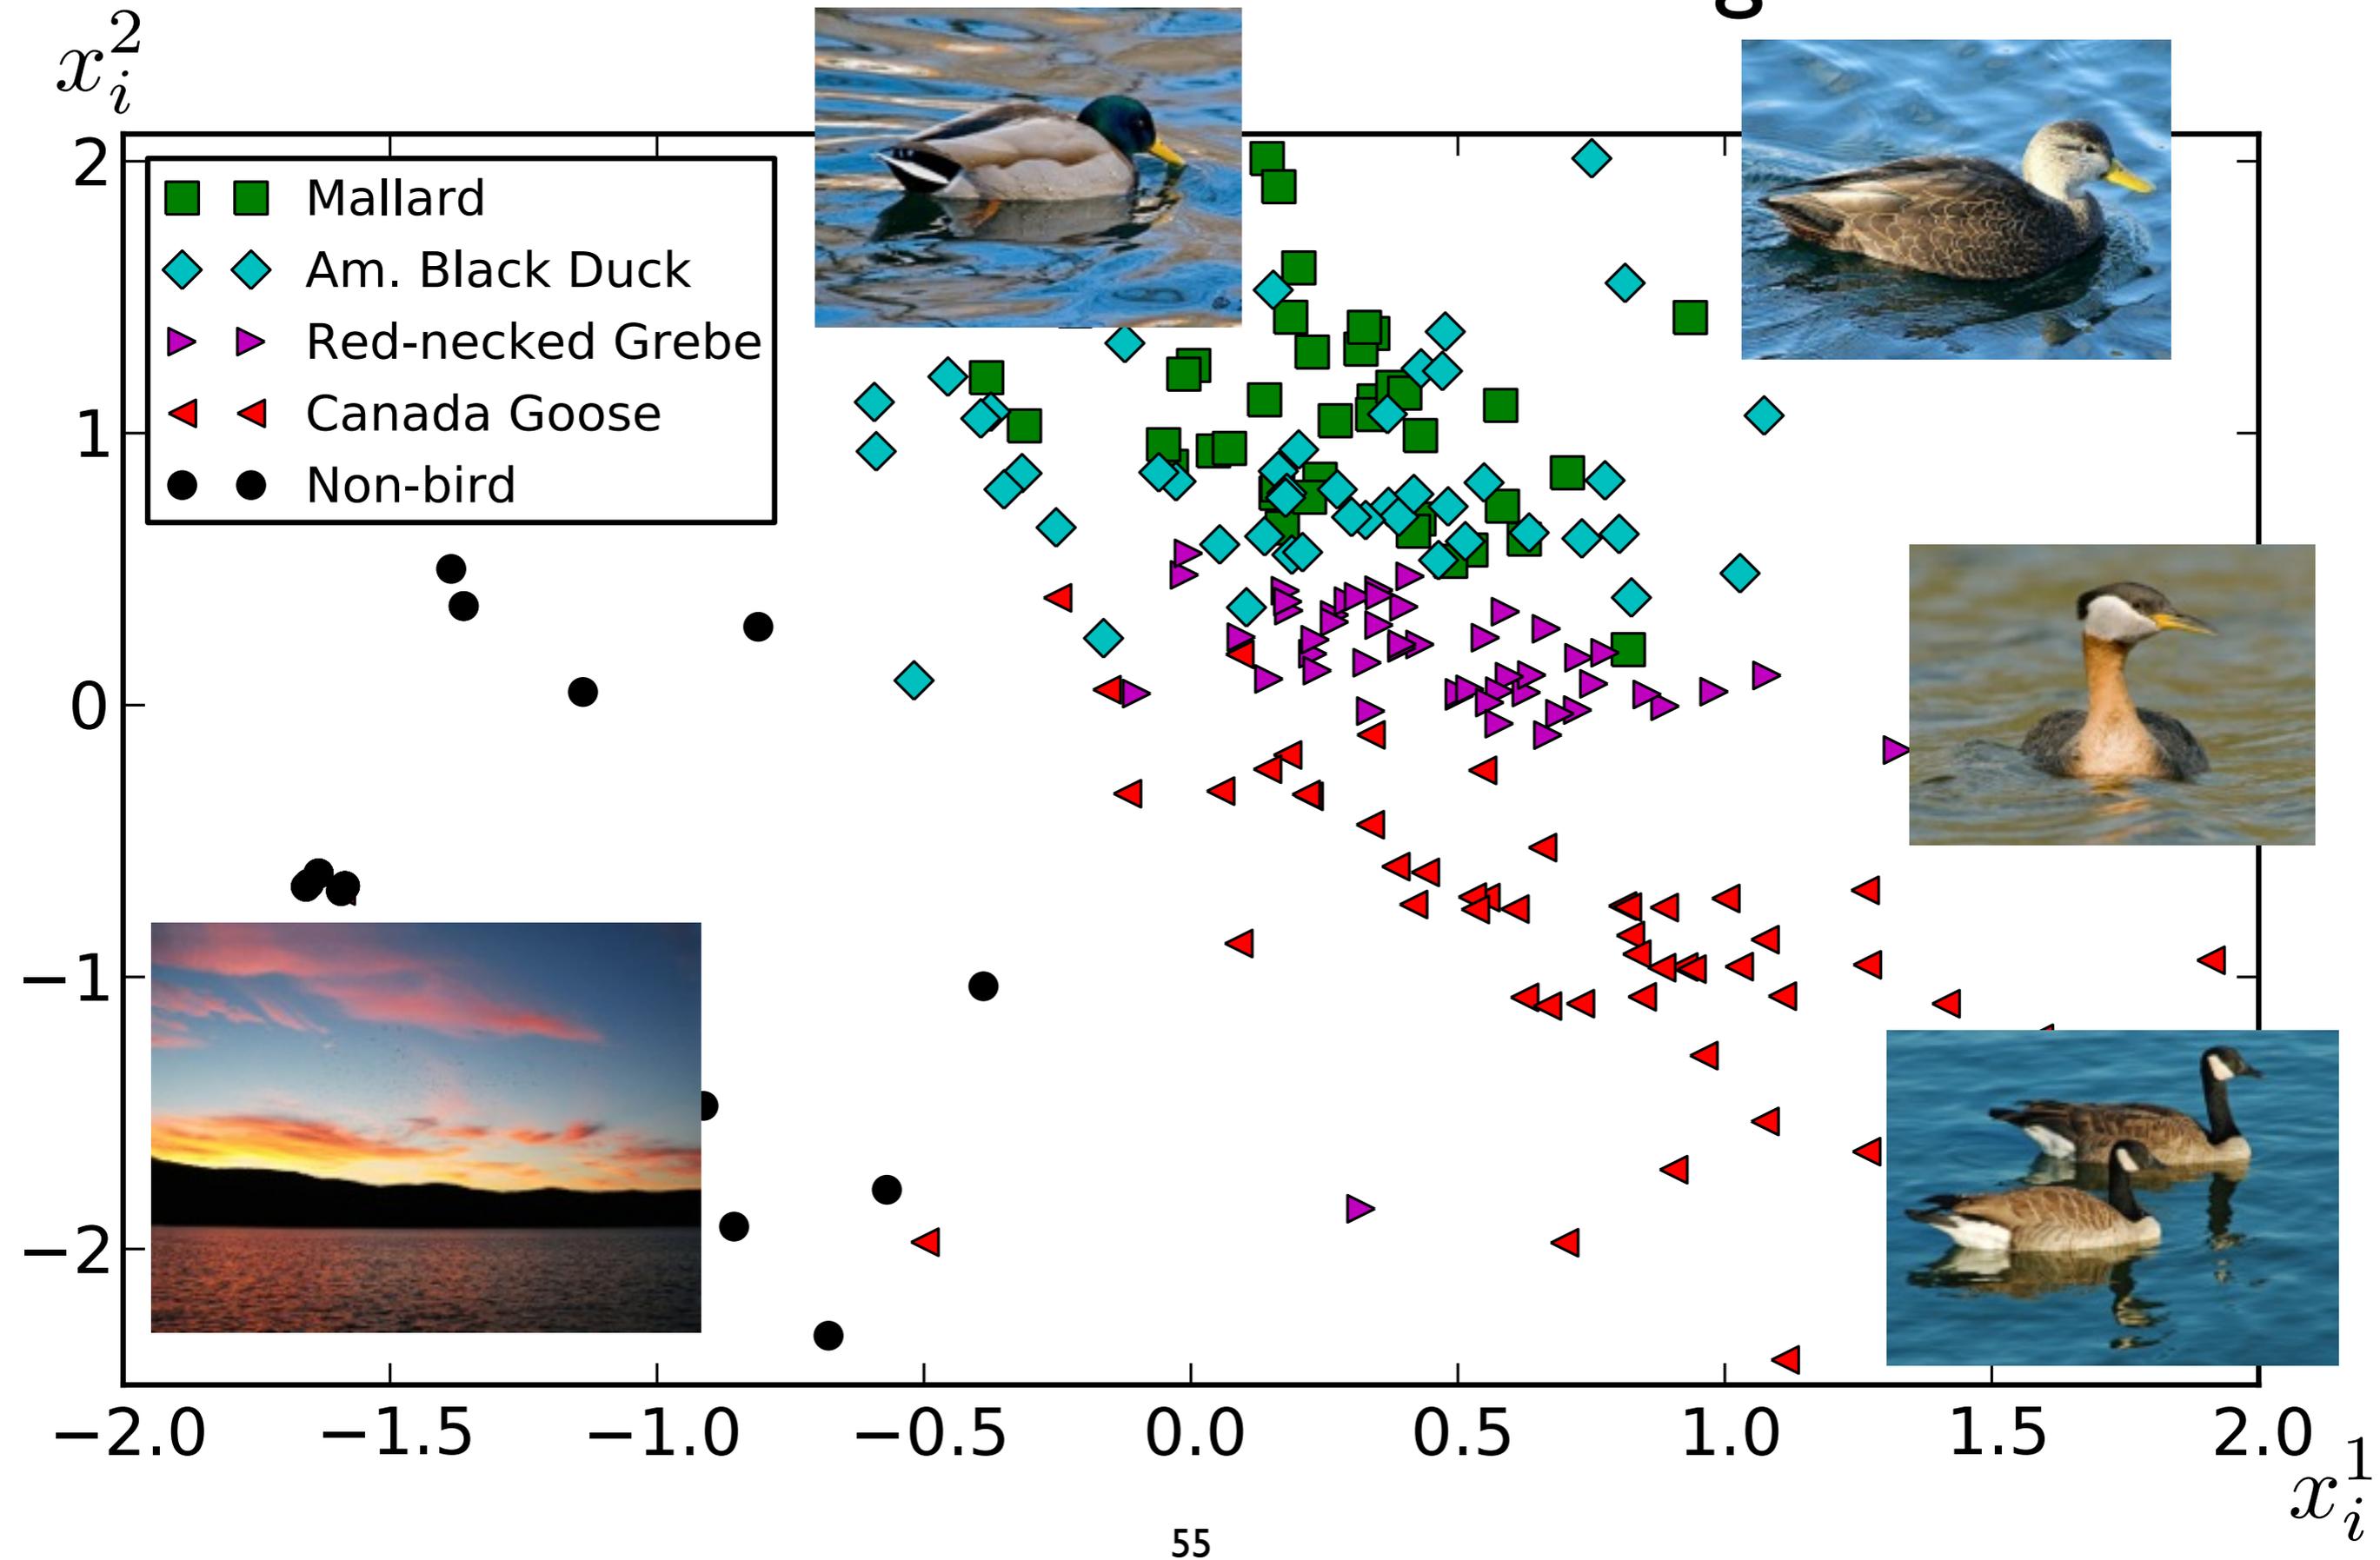
Is there a duck in the image?



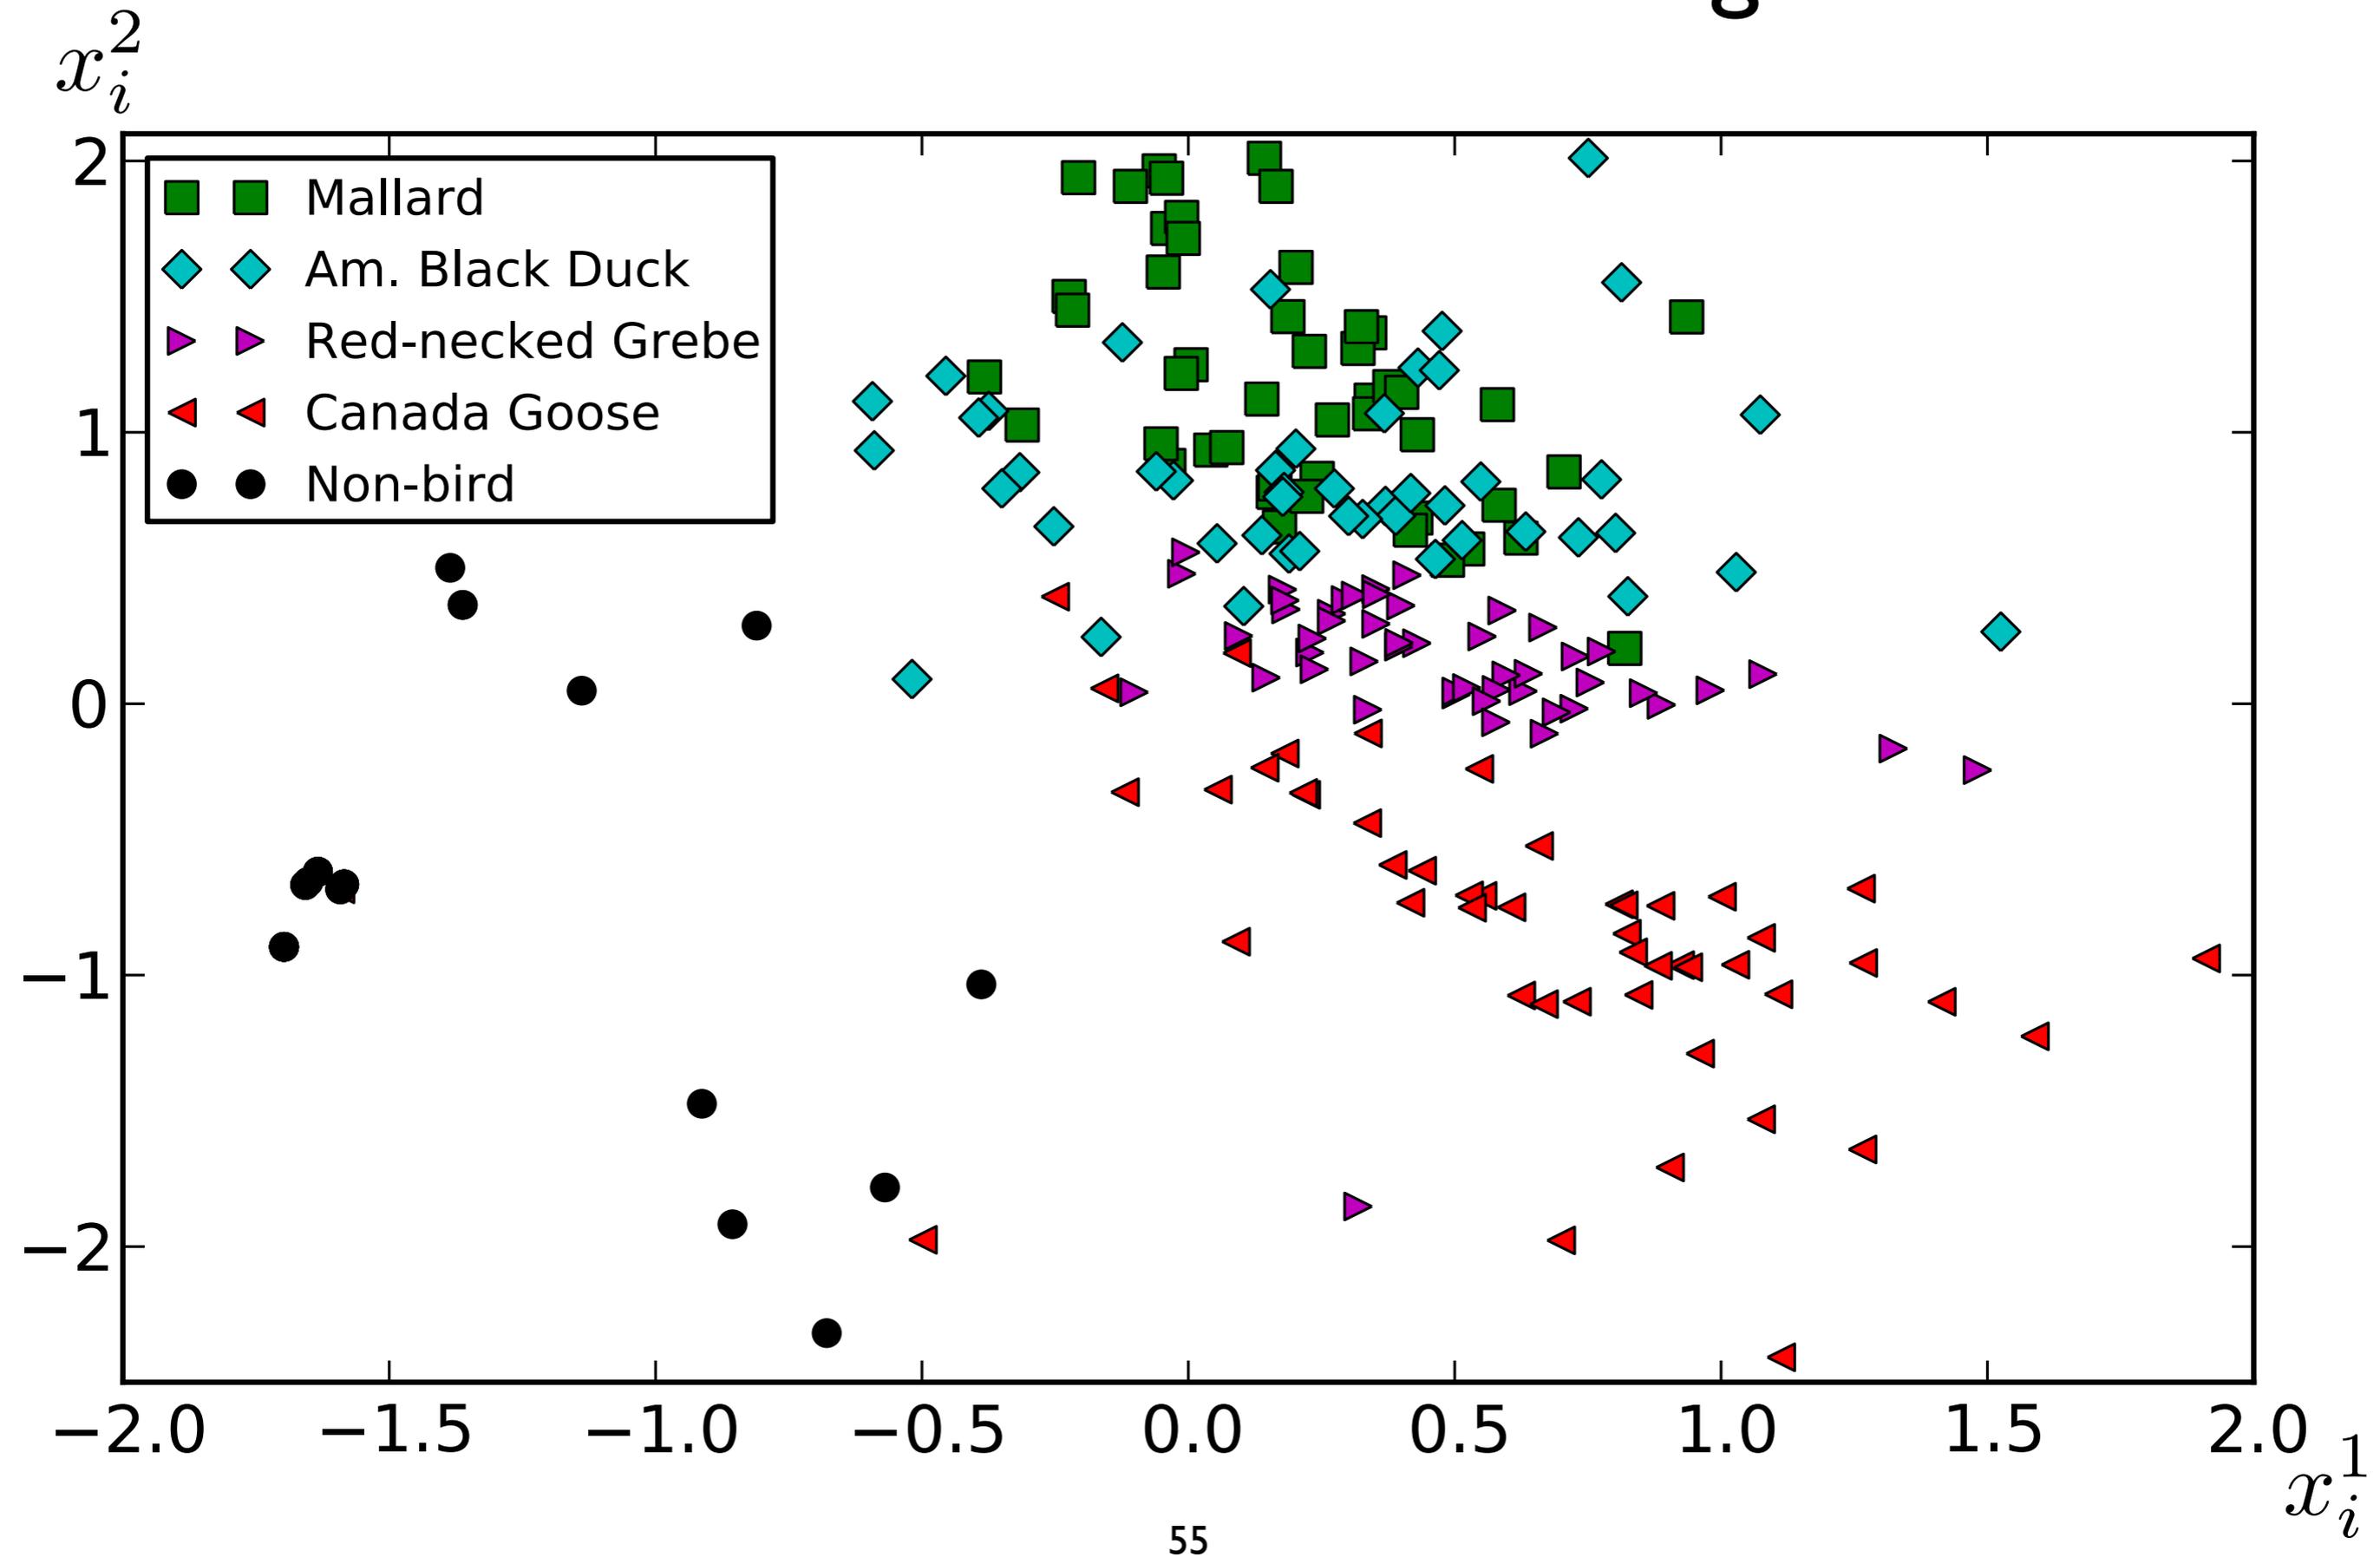
# Is there a duck in the image?



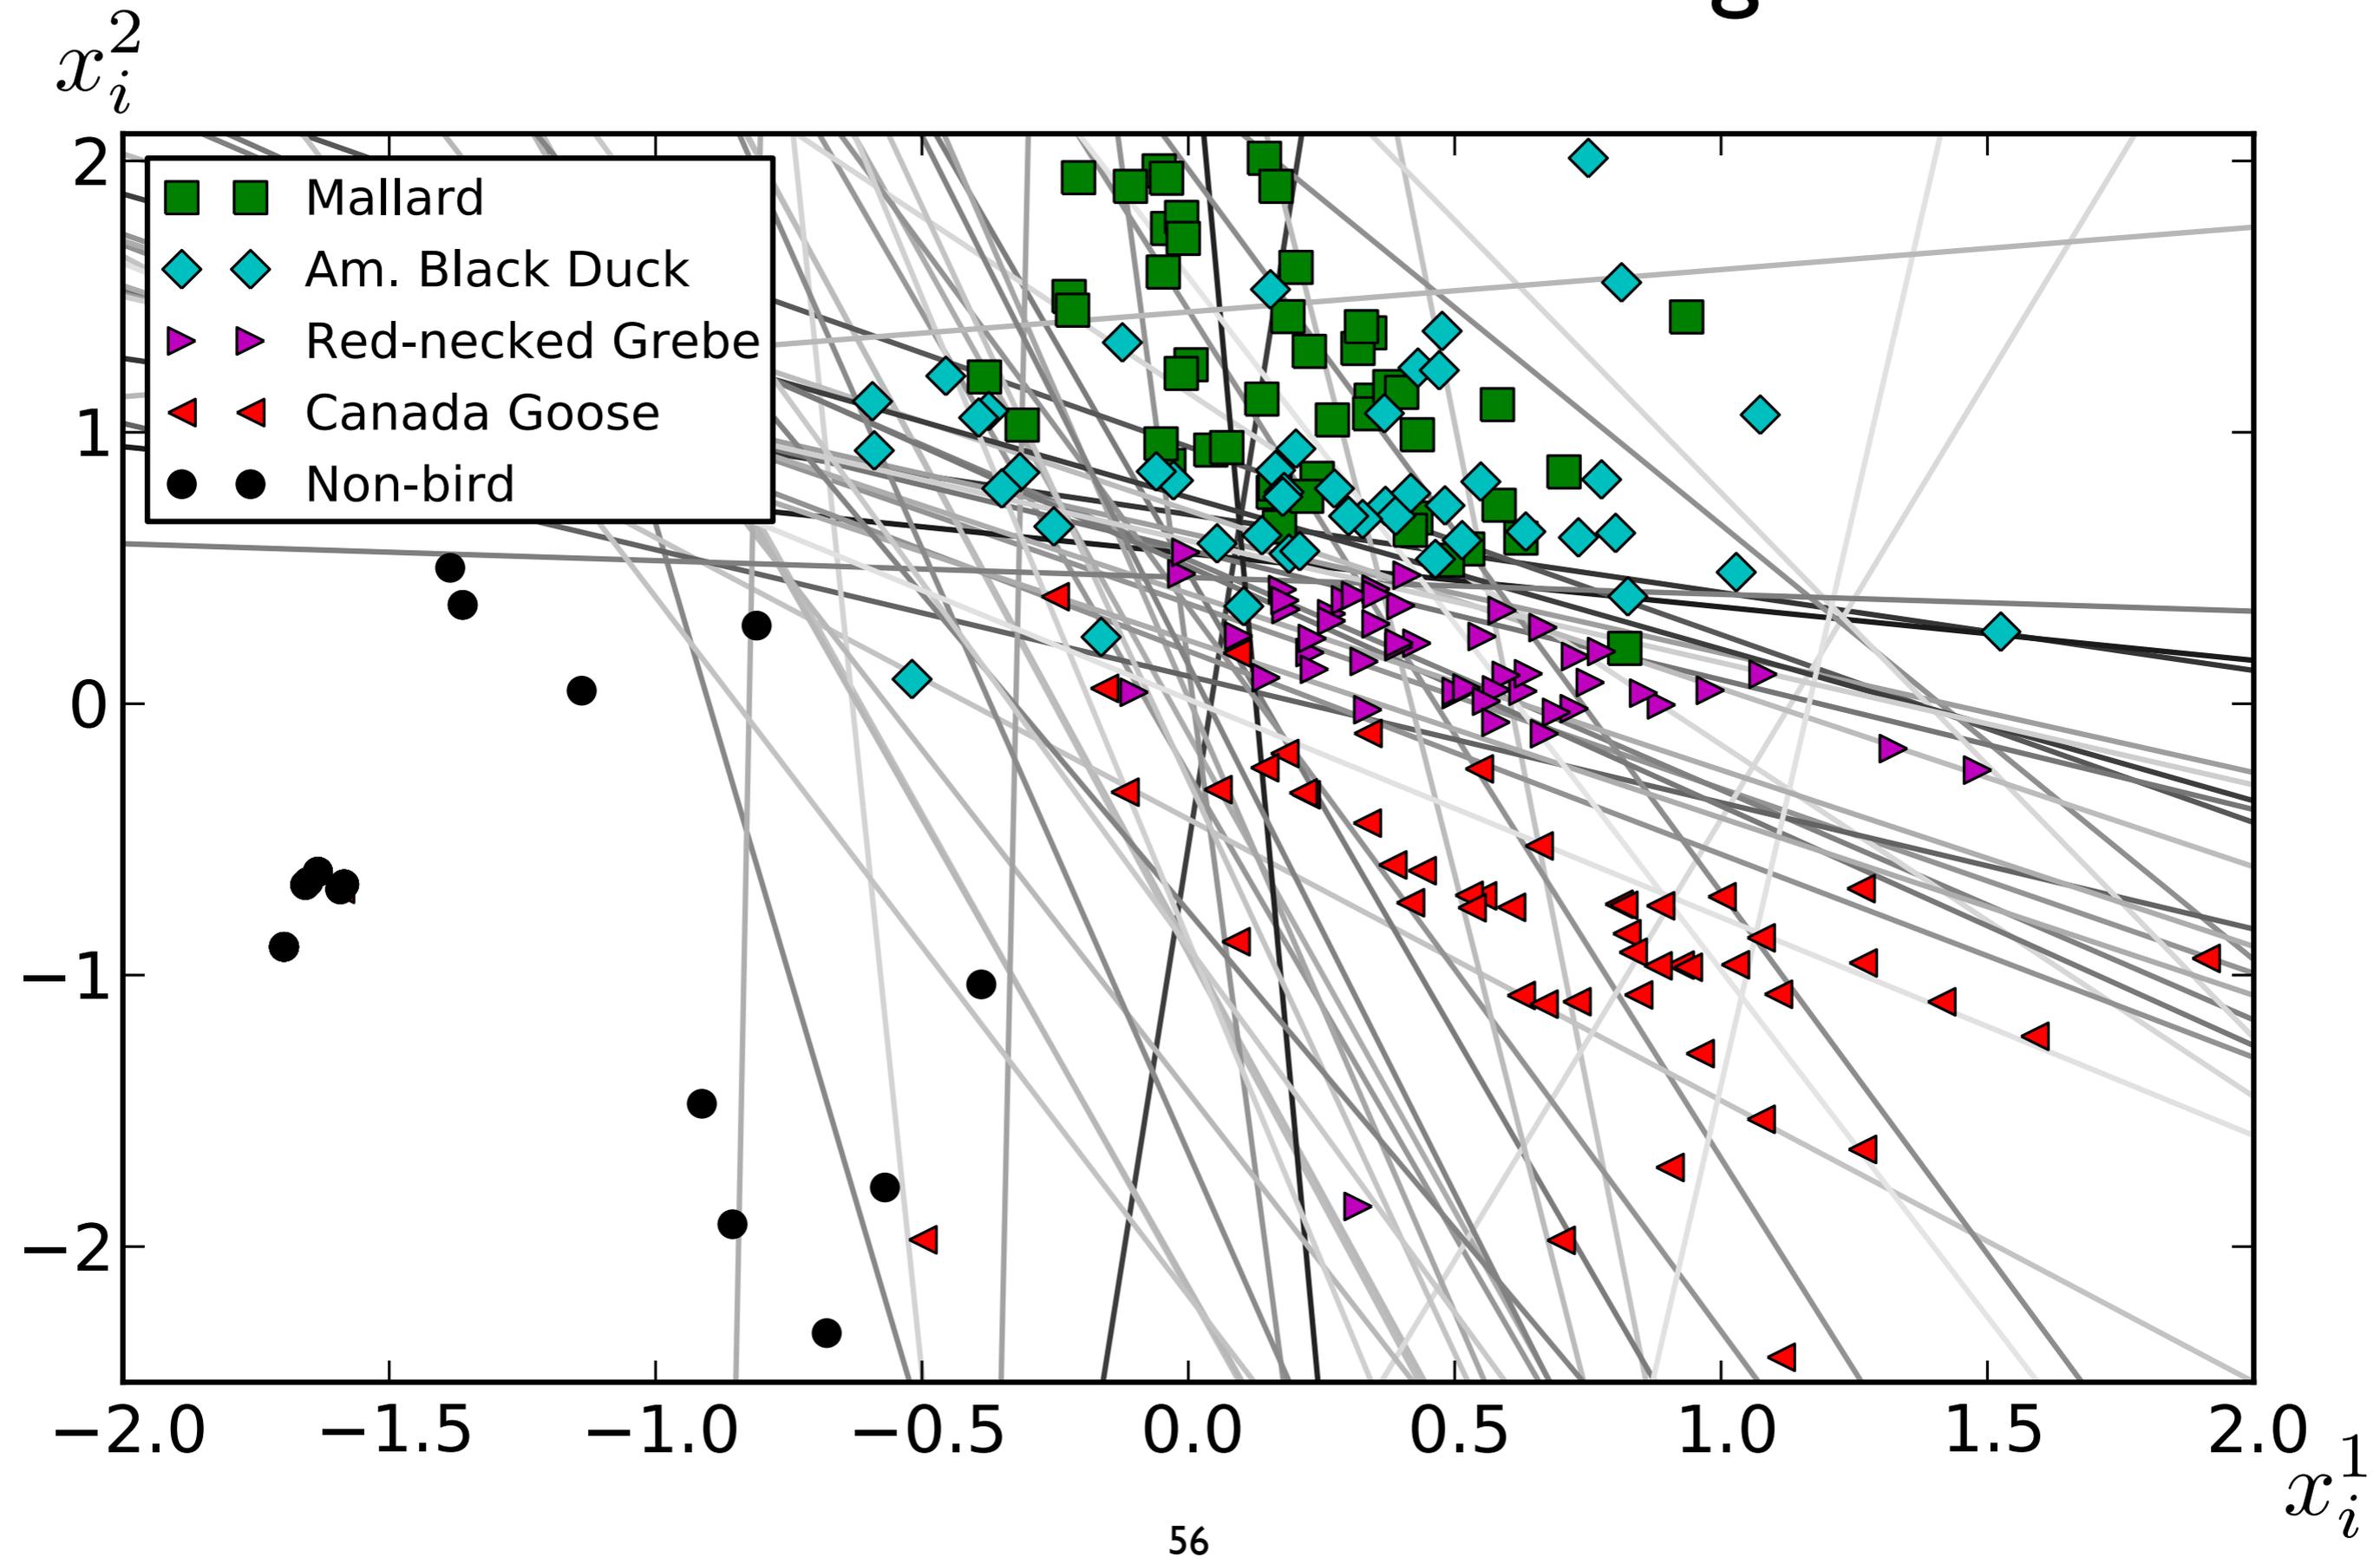
# Is there a duck in the image?



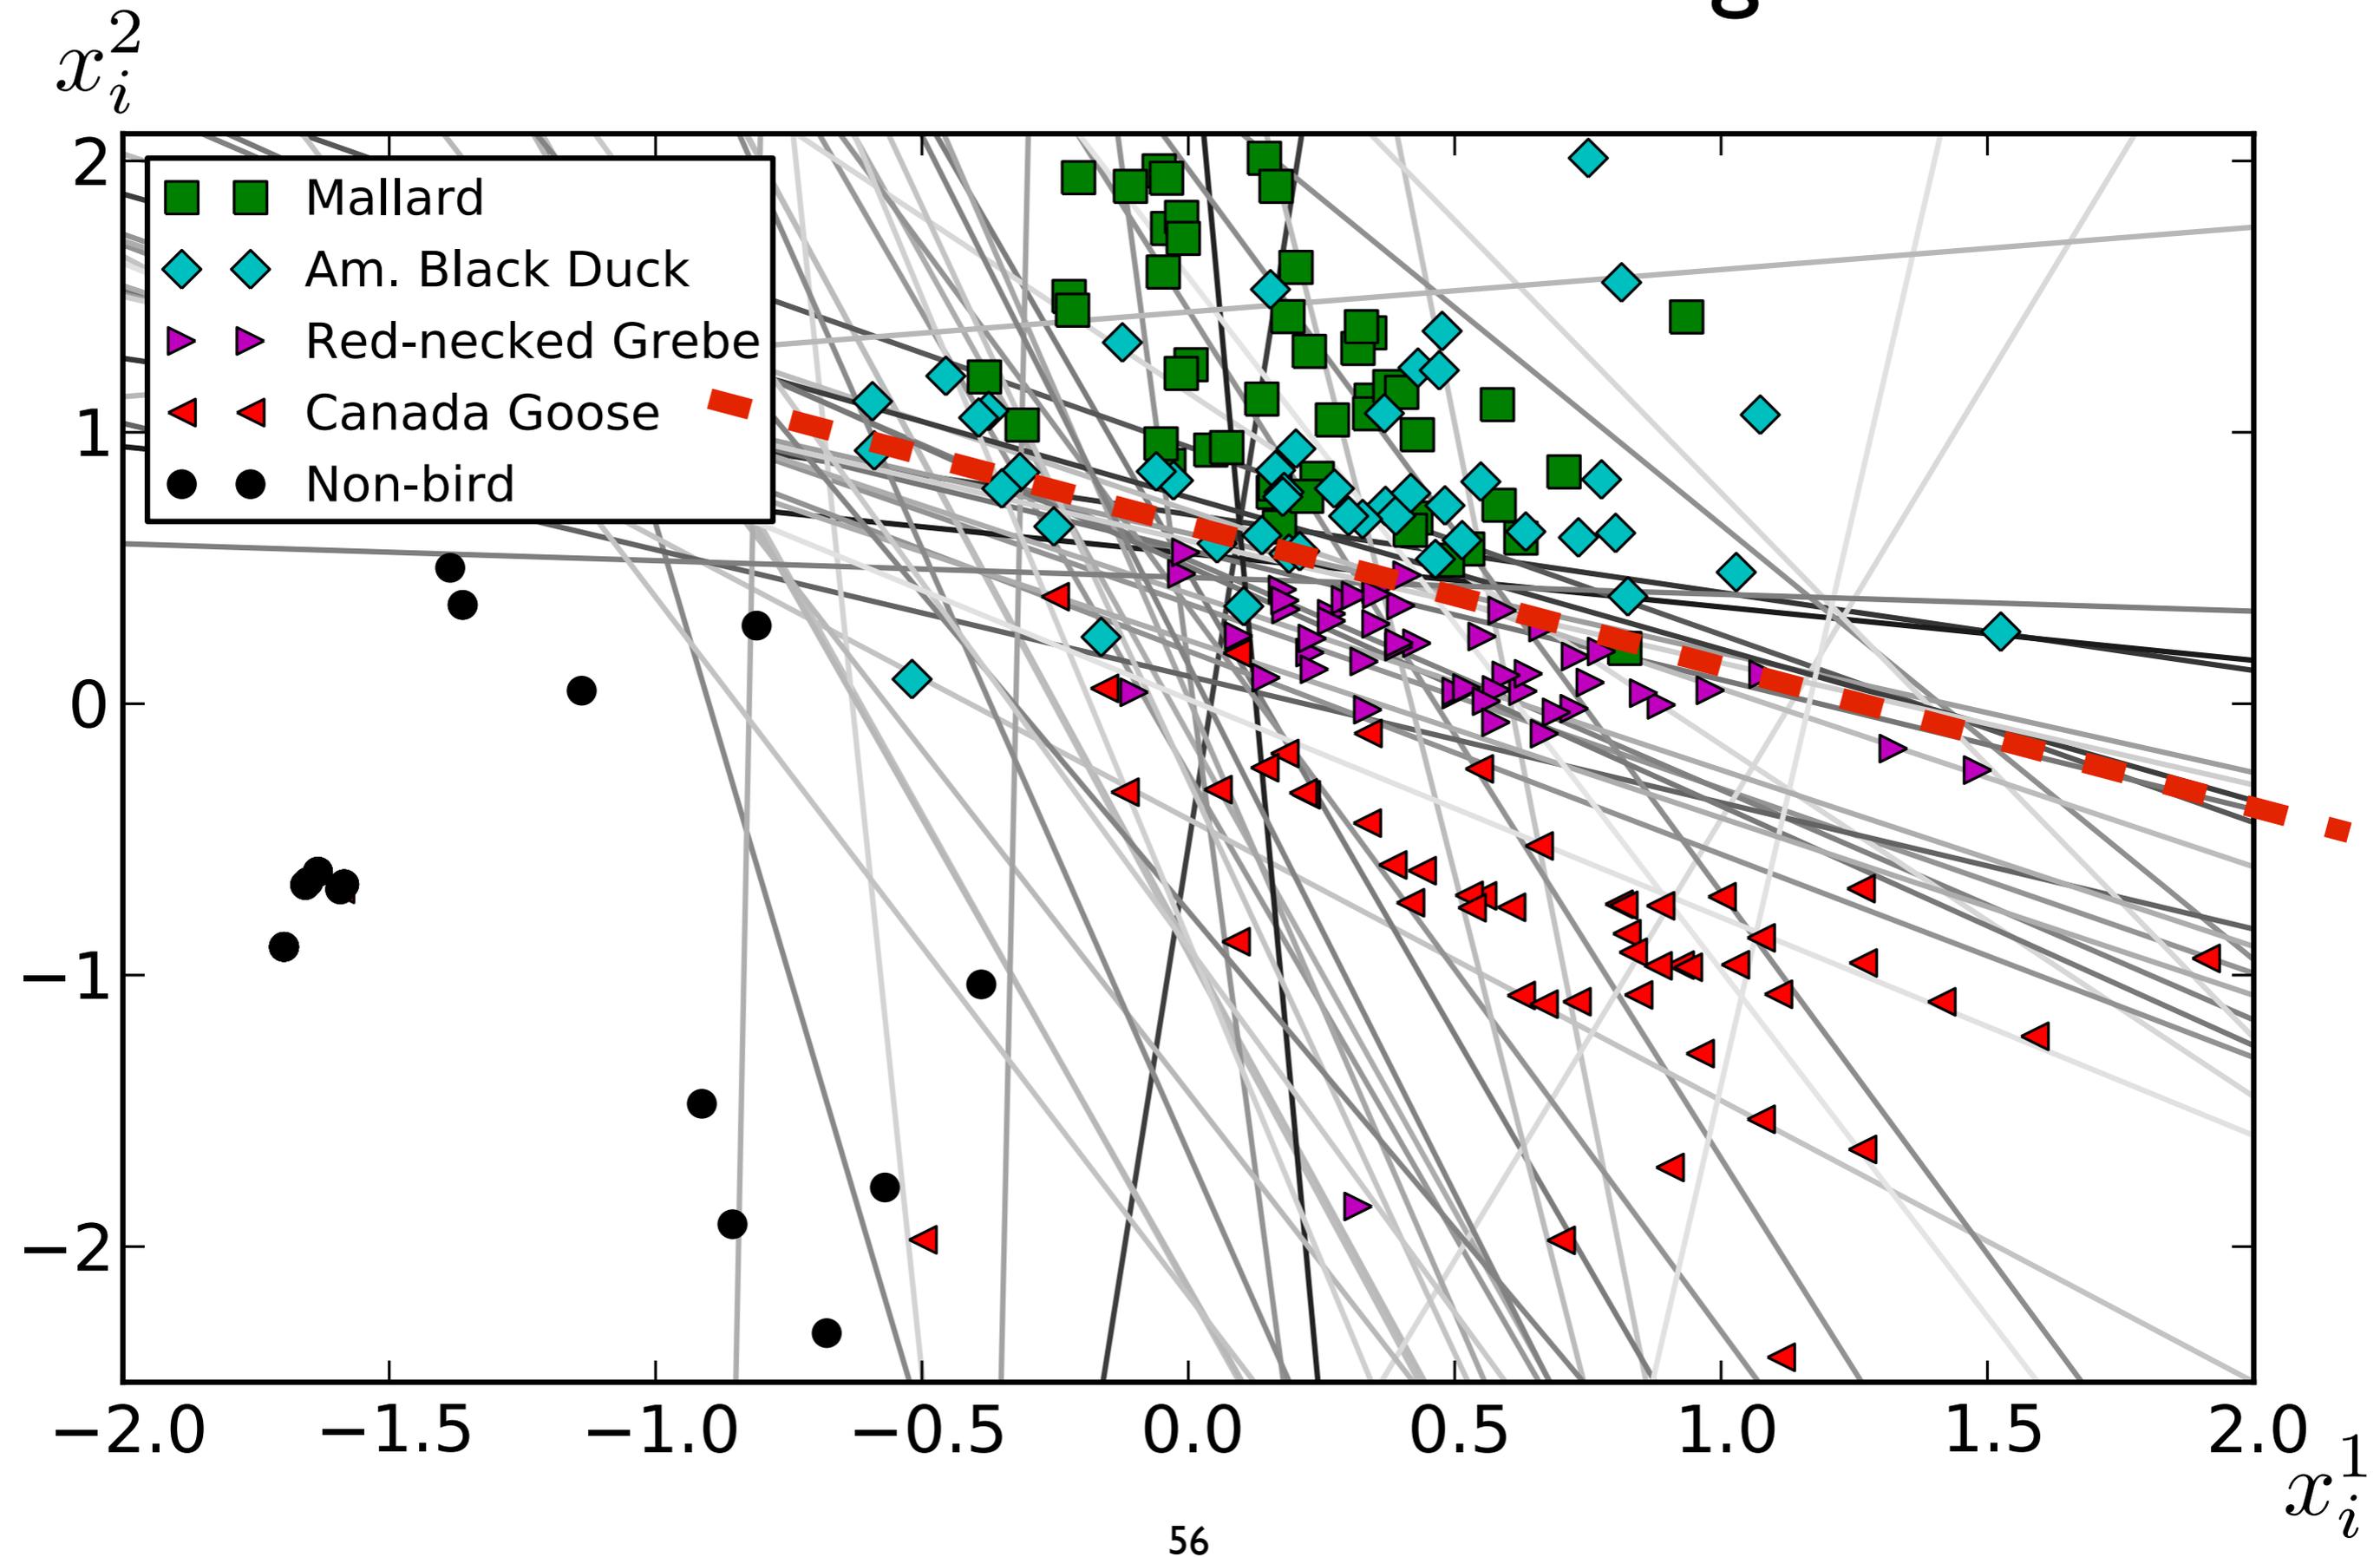
# Is there a duck in the image?



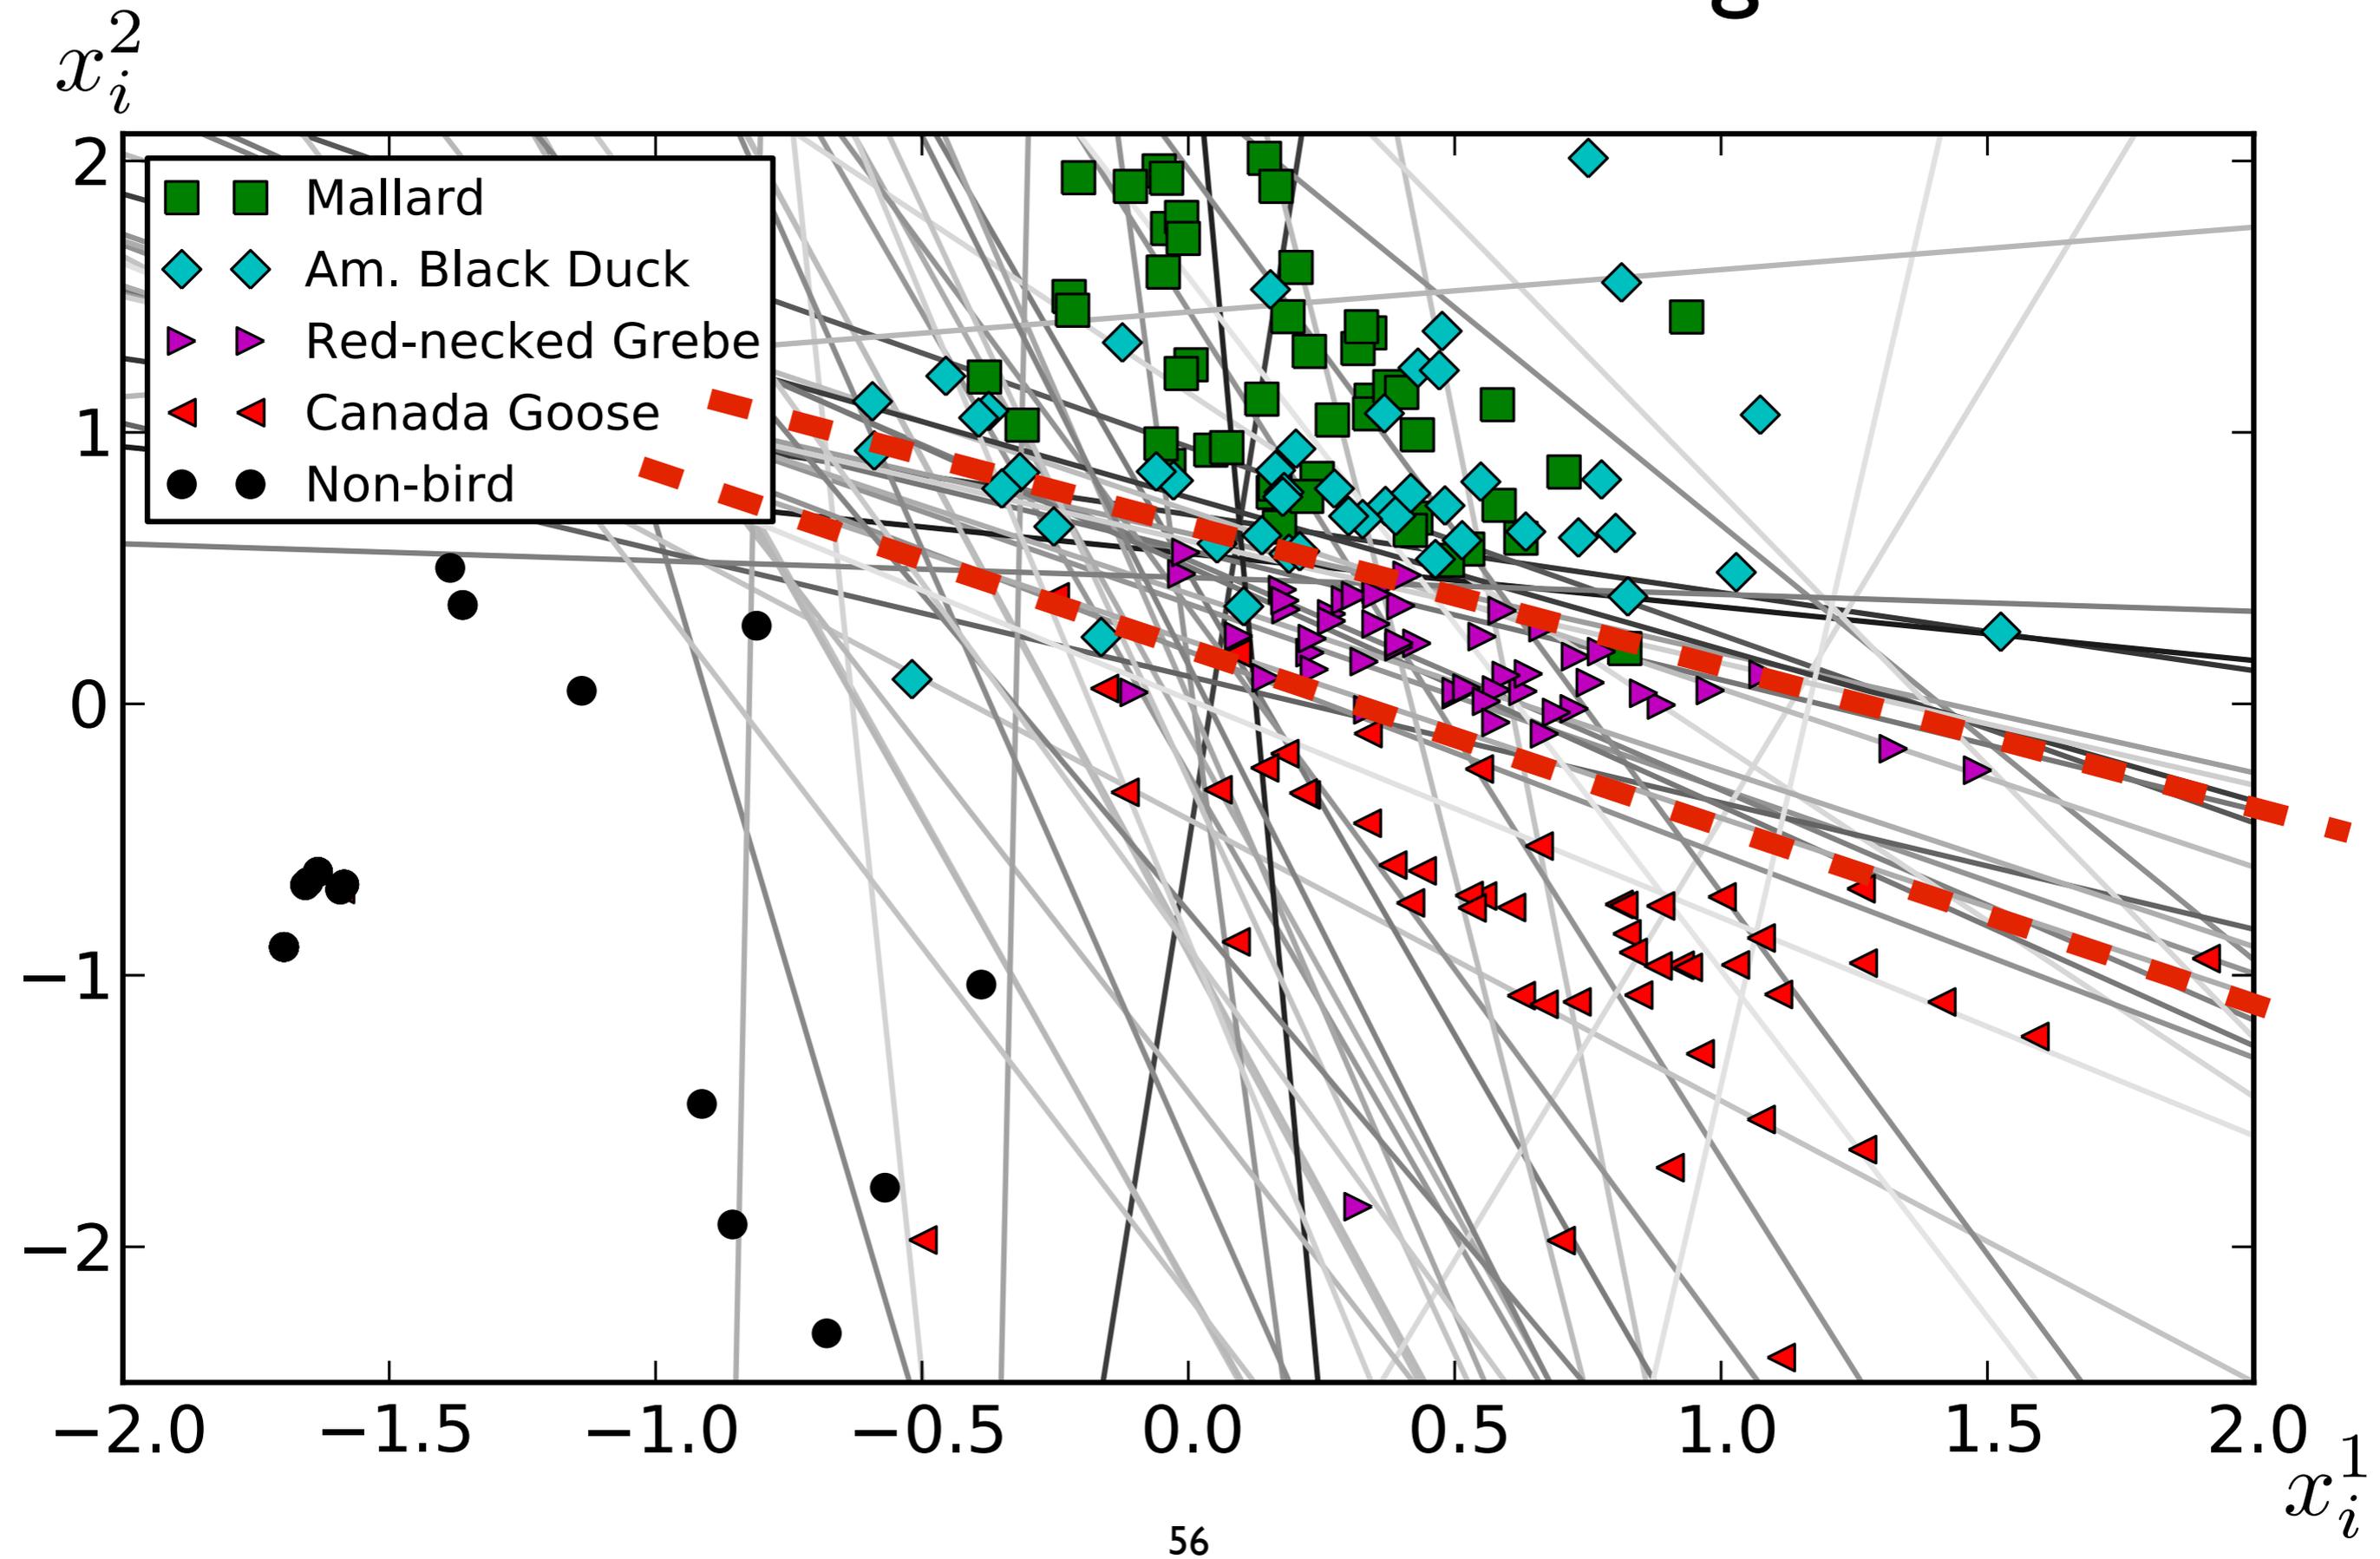
# Is there a duck in the image?



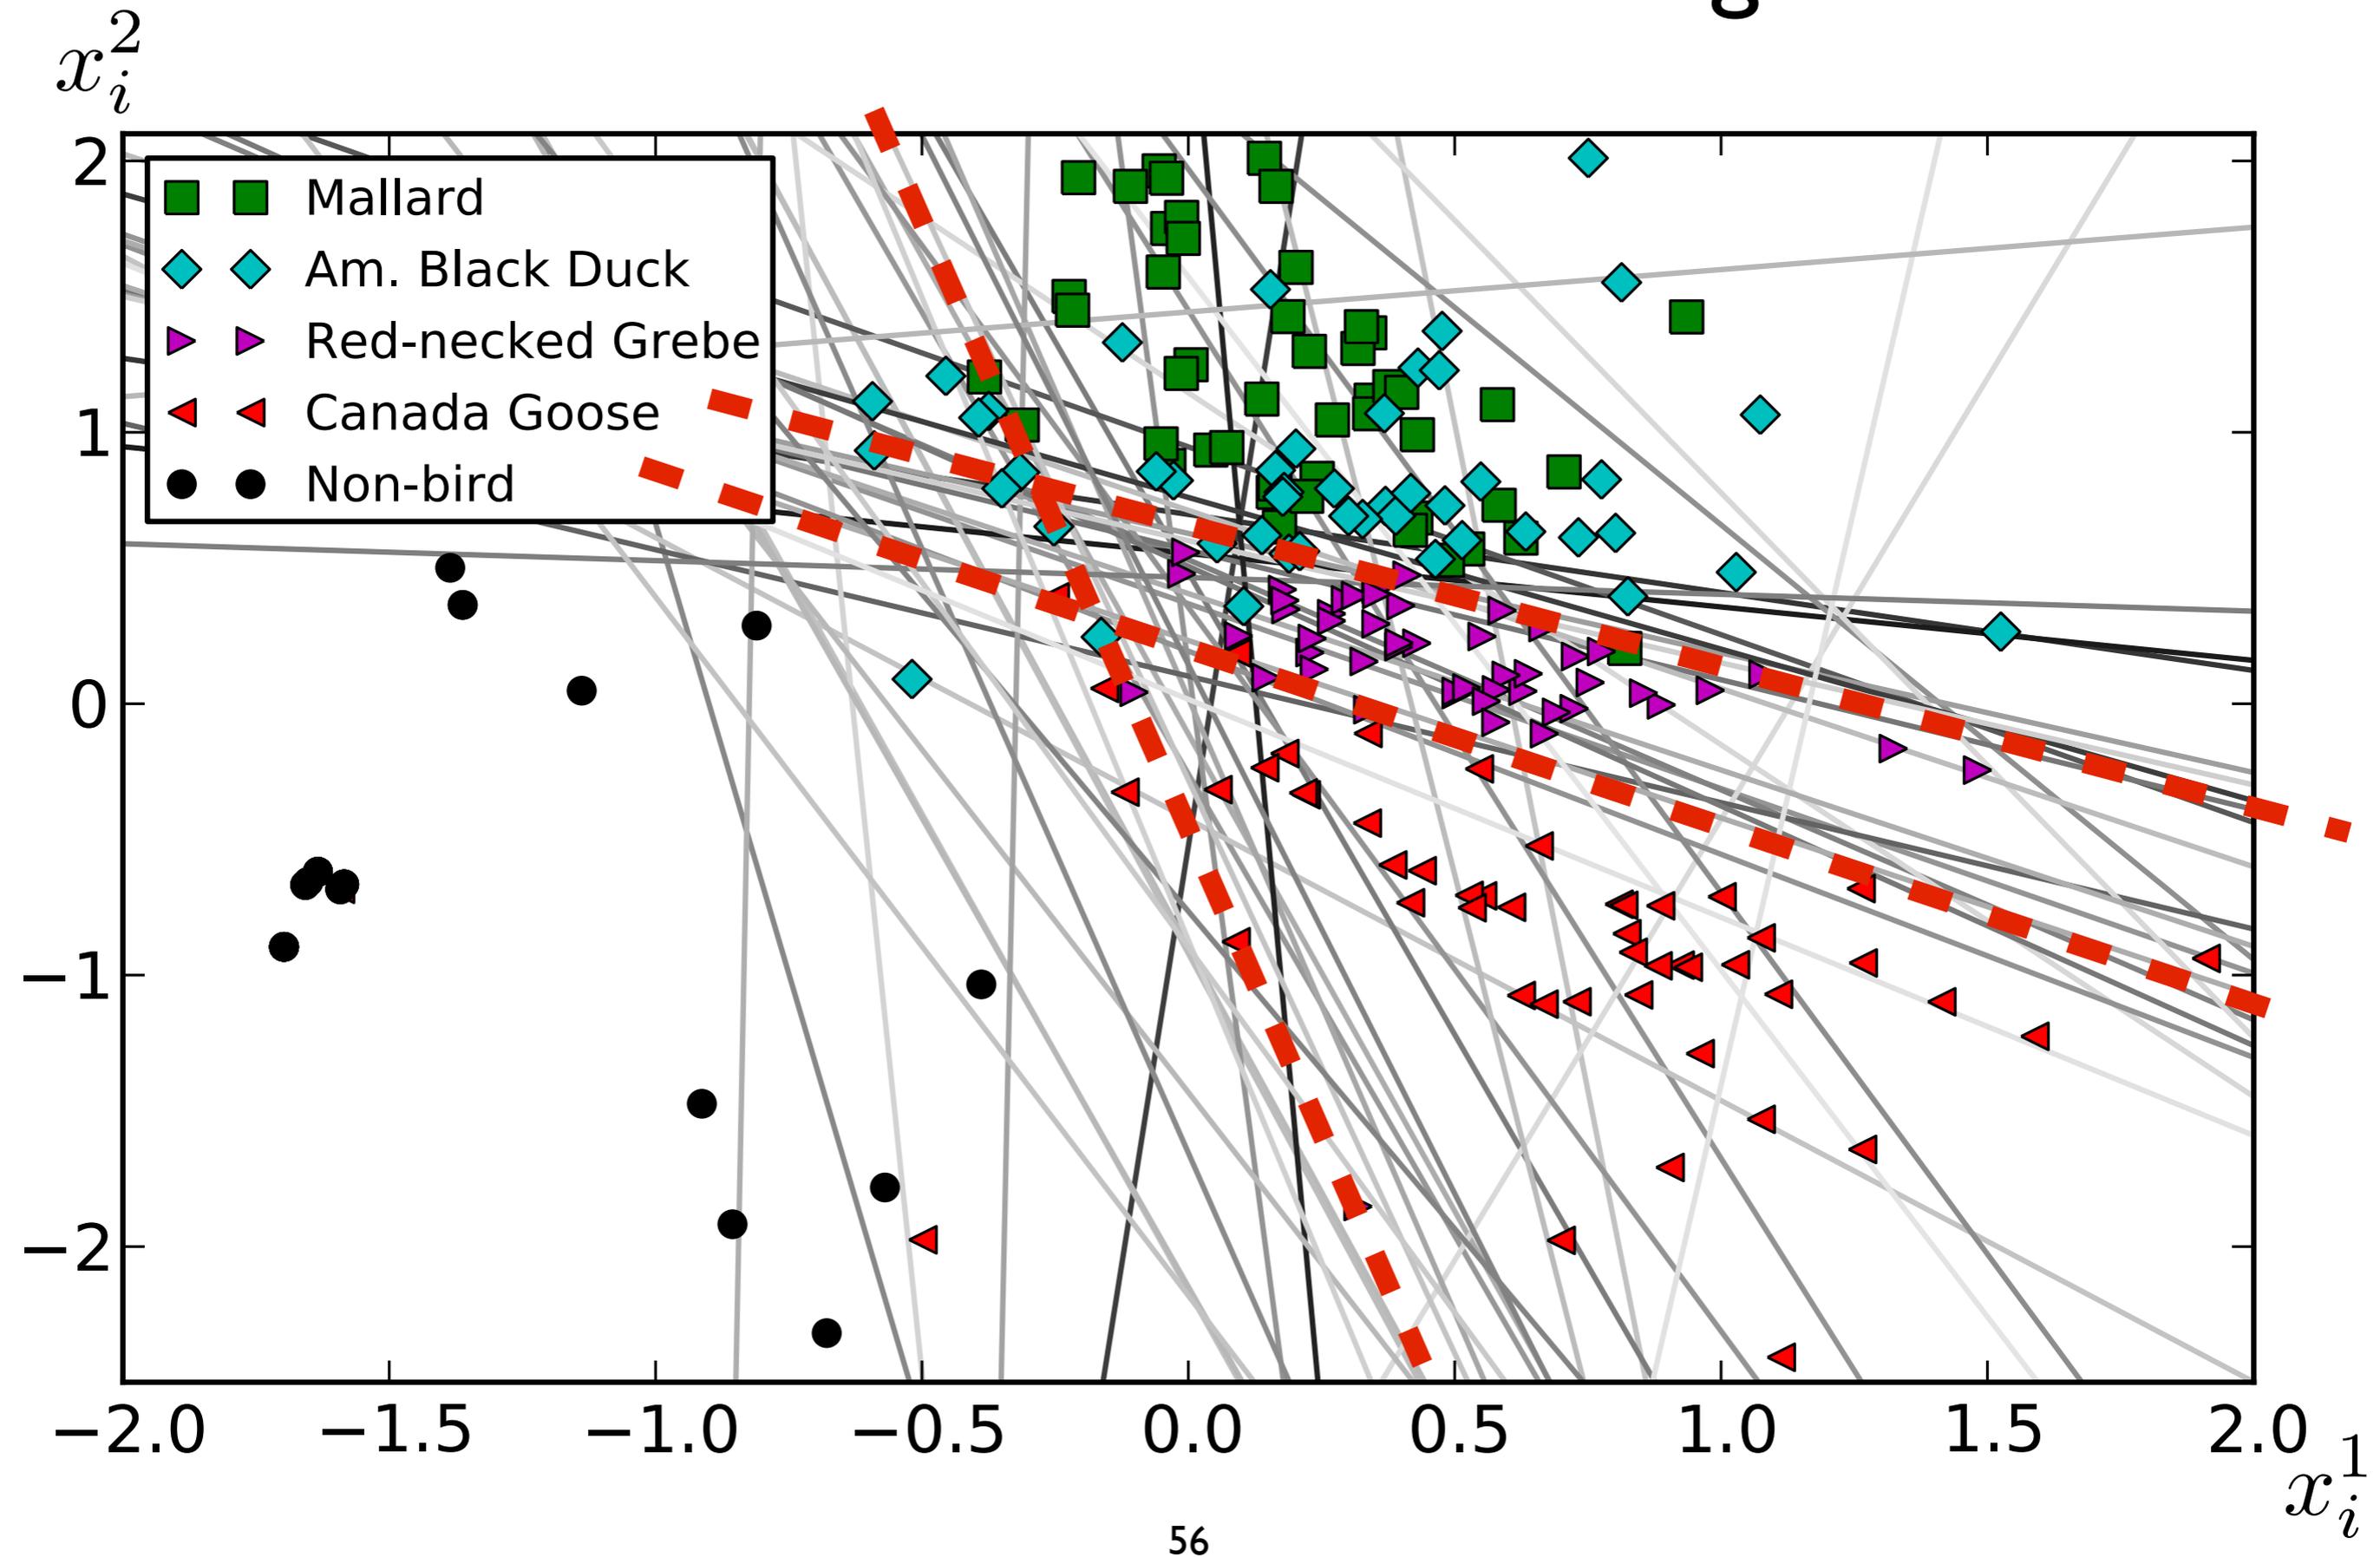
# Is there a duck in the image?



# Is there a duck in the image?



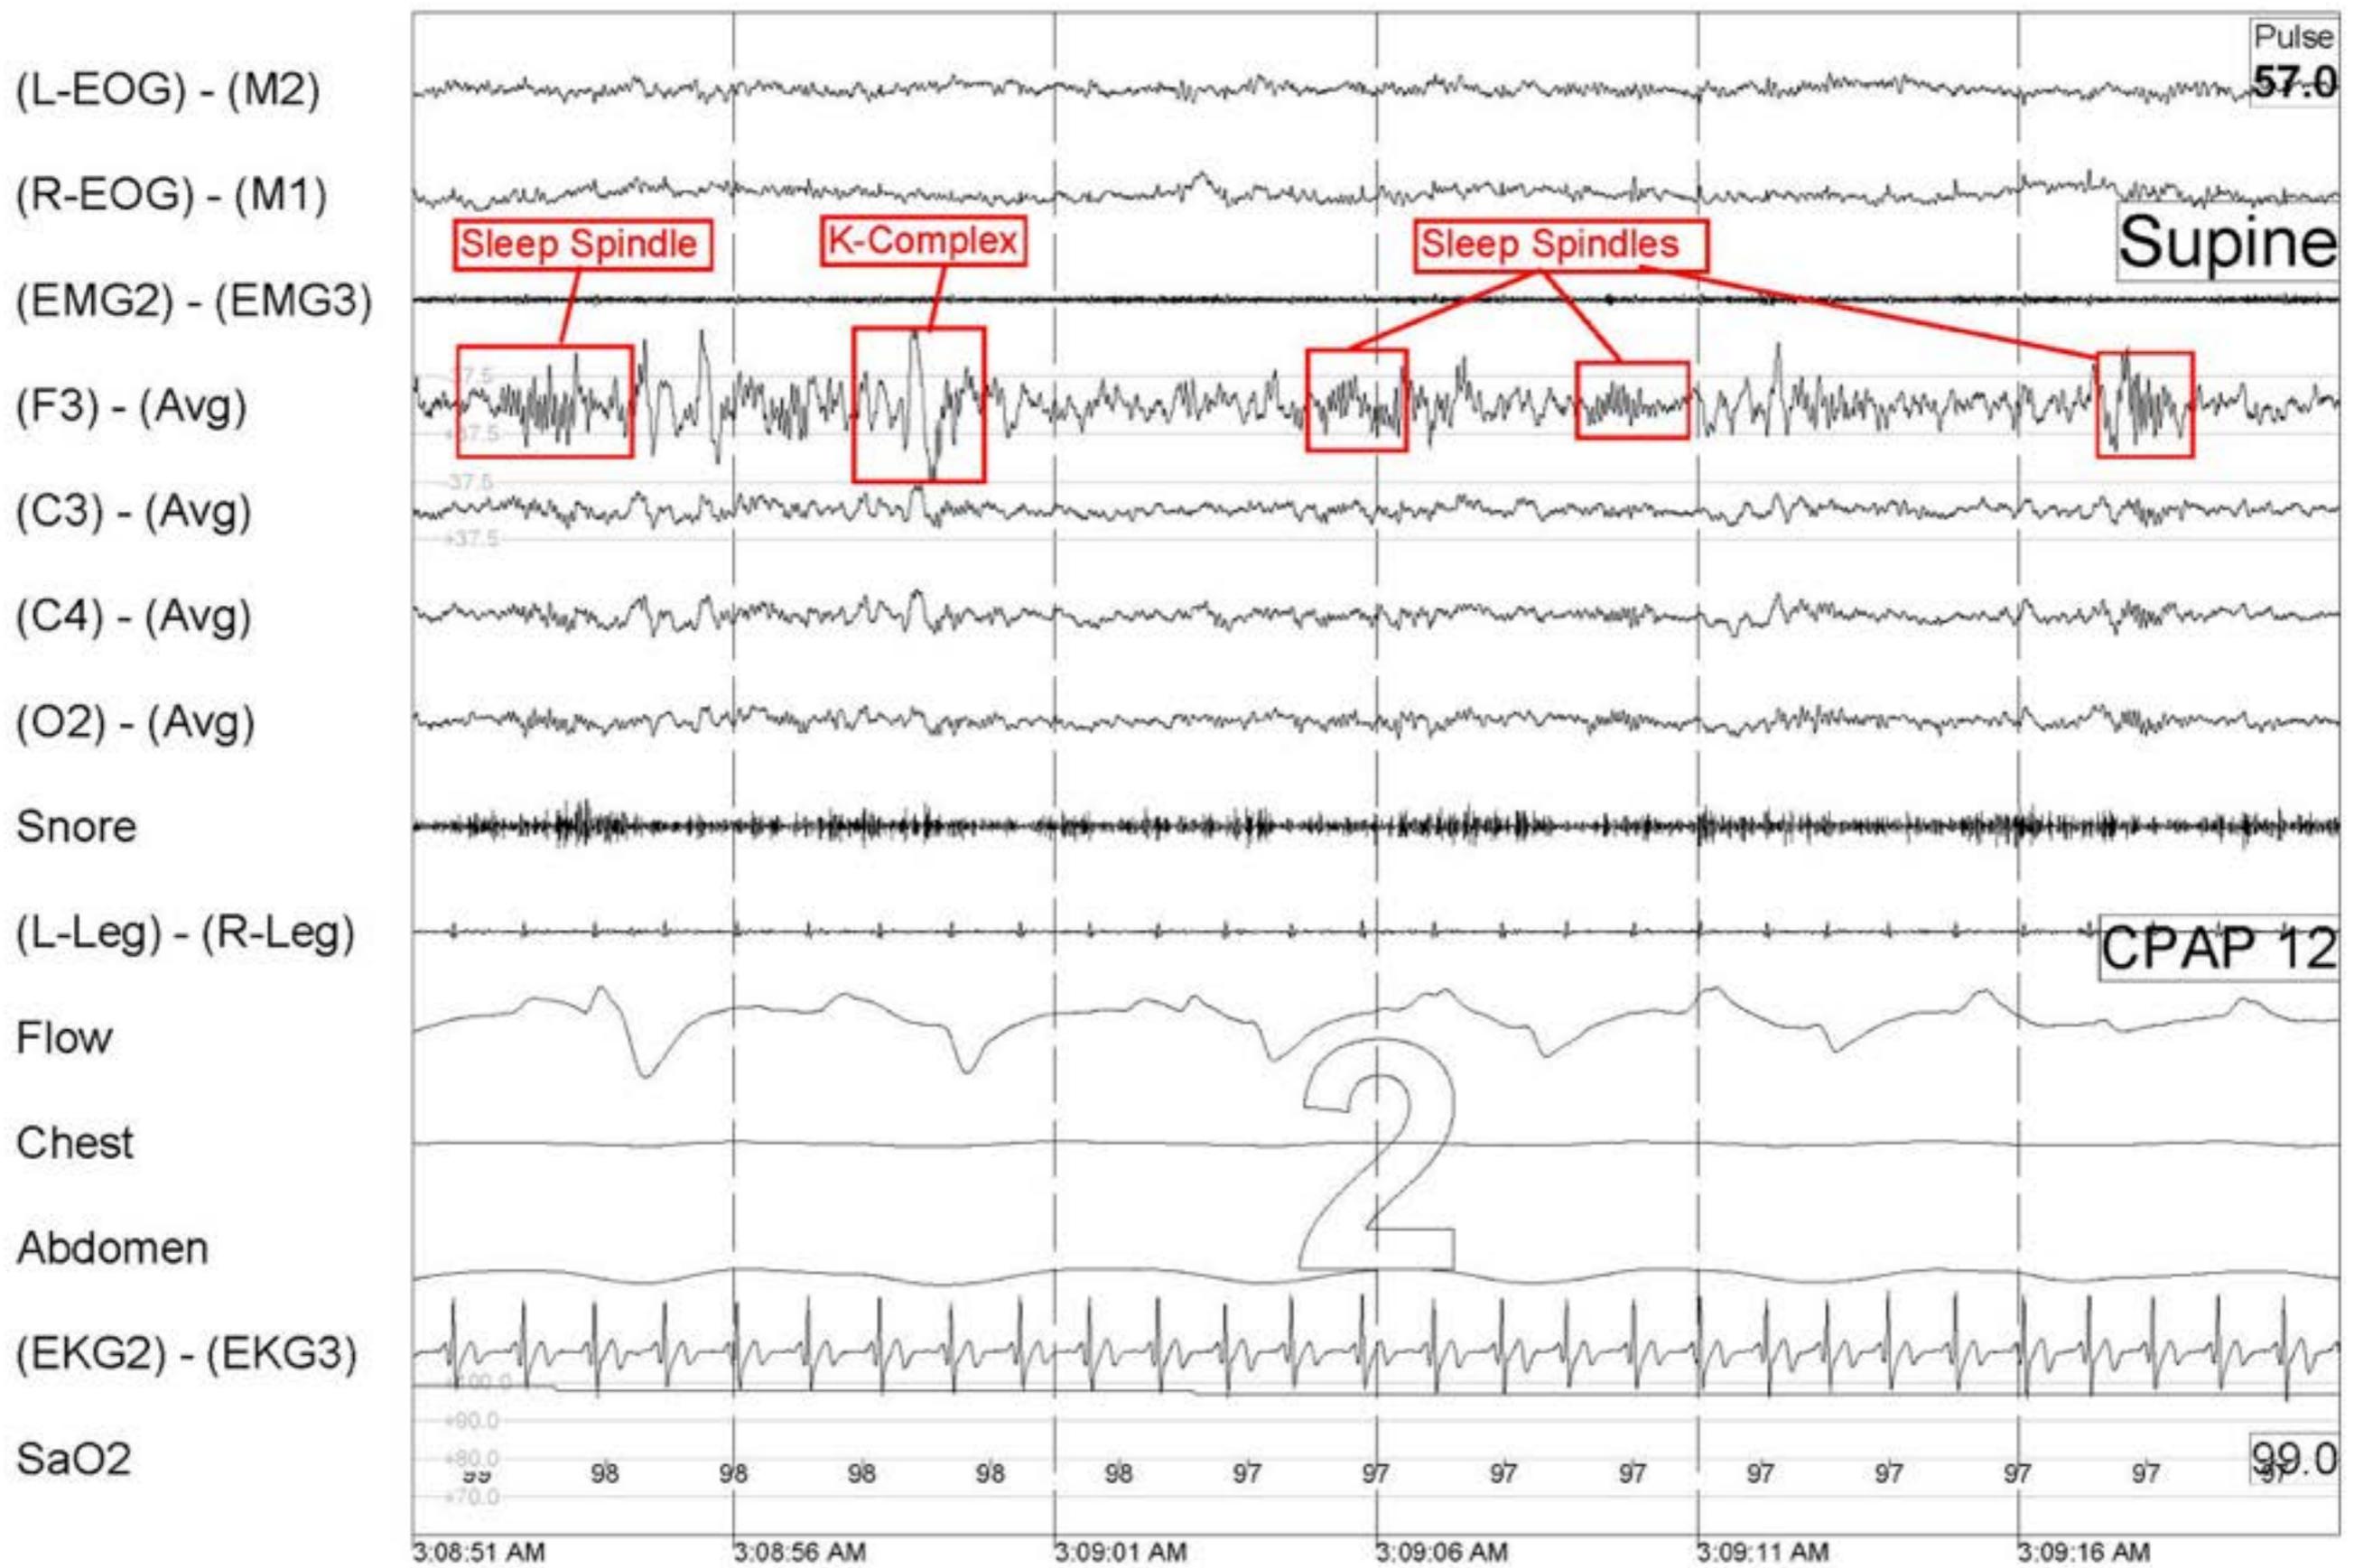
# Is there a duck in the image?



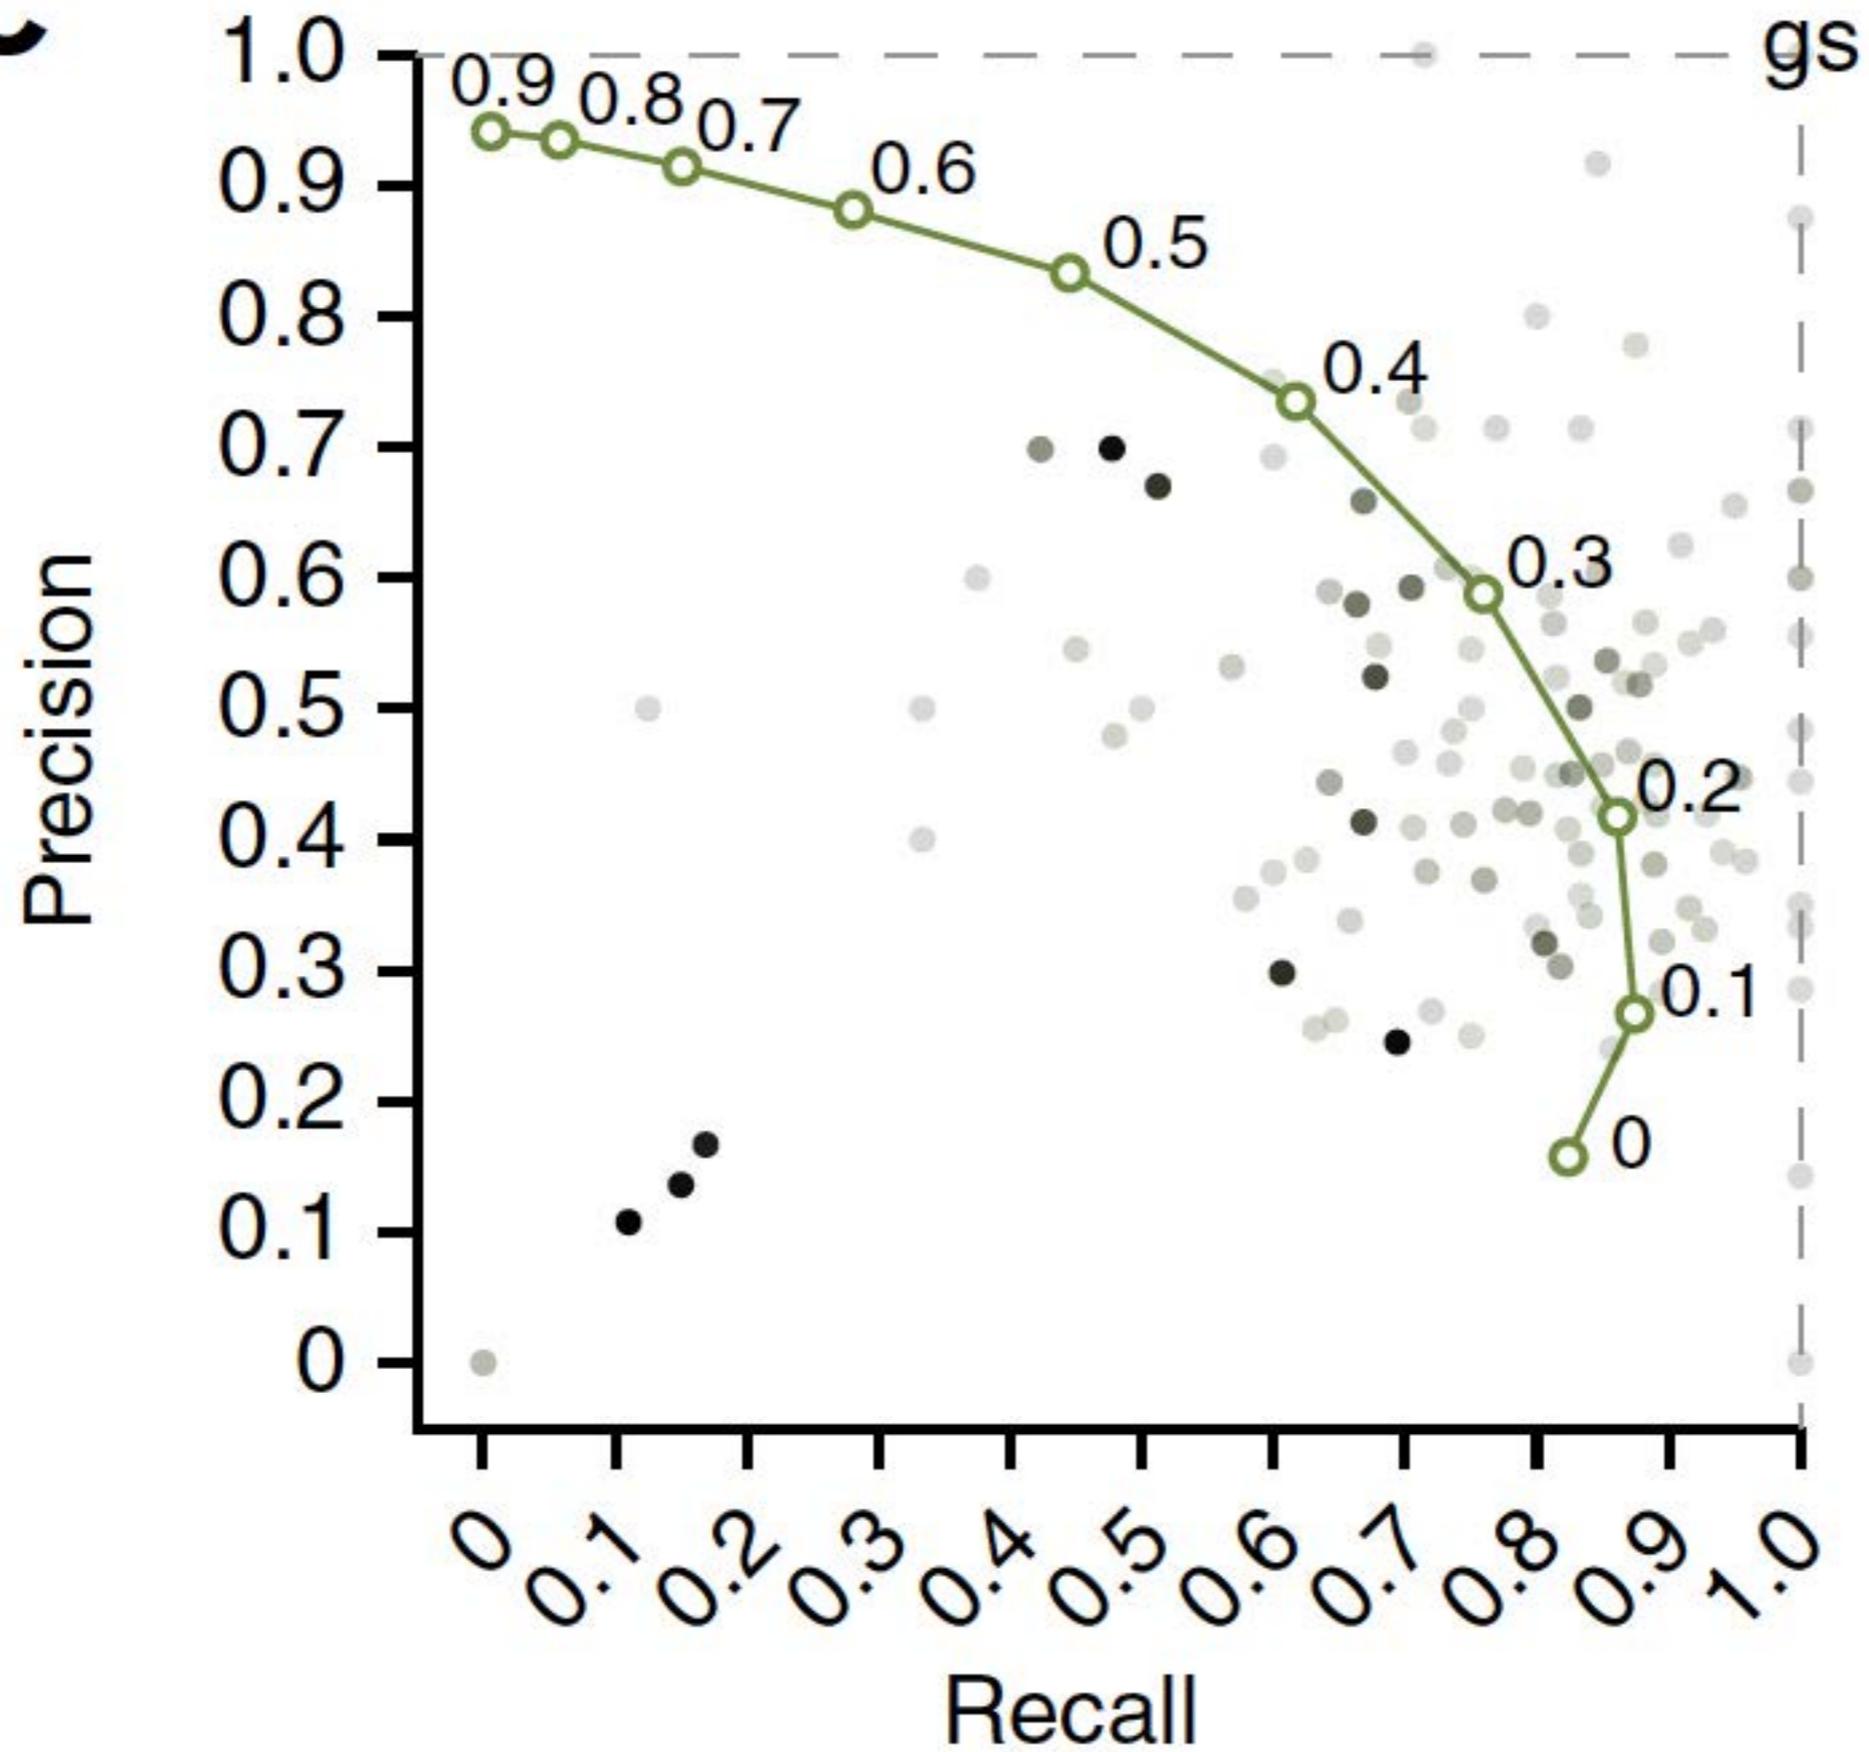
# The first application: sleep annotation

[Warby et al. 2014]





[Warby et al. 2014]

**C**

[Warby et al. 2014]



# Adaptive on-line annotation

 vs  what is the bias?

# Adaptive on-line annotation

- How many labels do we need until we are certain of an annotation?
- Coin flipping:  vs  what is the bias?

# Adaptive on-line annotation

- How many labels do we need until we are certain of an annotation?
- Coin flipping:  vs  what is the bias?



# Adaptive on-line annotation

- How many labels do we need until we are certain of an annotation?
- Coin flipping:  vs  what is the bias?



# Adaptive on-line annotation

- How many labels do we need until we are certain of an annotation?
- Coin flipping:  vs  what is the bias?



# Adaptive on-line annotation

- How many labels do we need until we are certain of an annotation?
- Coin flipping:  vs  what is the bias?



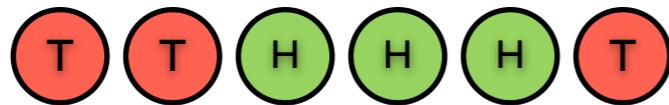
# Adaptive on-line annotation

- How many labels do we need until we are certain of an annotation?
- Coin flipping:  vs  what is the bias?



# Adaptive on-line annotation

- How many labels do we need until we are certain of an annotation?
- Coin flipping:  vs  what is the bias?



# Adaptive on-line annotation

- How many labels do we need until we are certain of an annotation?
- Coin flipping:  vs  what is the bias?



# Adaptive on-line annotation

- How many labels do we need until we are certain of an annotation?
- Coin flipping:  vs  what is the bias?



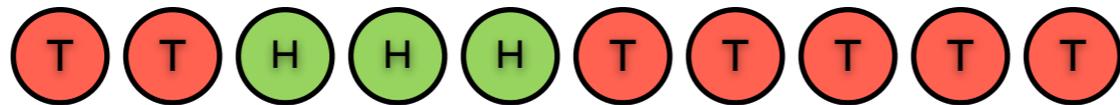
# Adaptive on-line annotation

- How many labels do we need until we are certain of an annotation?
- Coin flipping:  vs  what is the bias?



# Adaptive on-line annotation

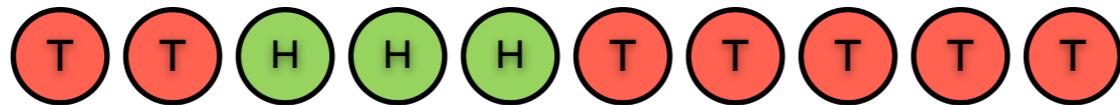
- How many labels do we need until we are certain of an annotation?
- Coin flipping:  vs  what is the bias?



# Adaptive on-line annotation

- How many labels do we need until we are certain of an annotation?

- Coin flipping:  vs  what is the bias?

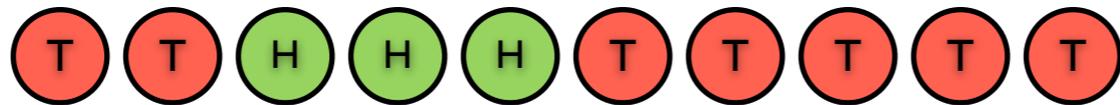


- Wald (1945): sequential probability ratio test (SPRT)

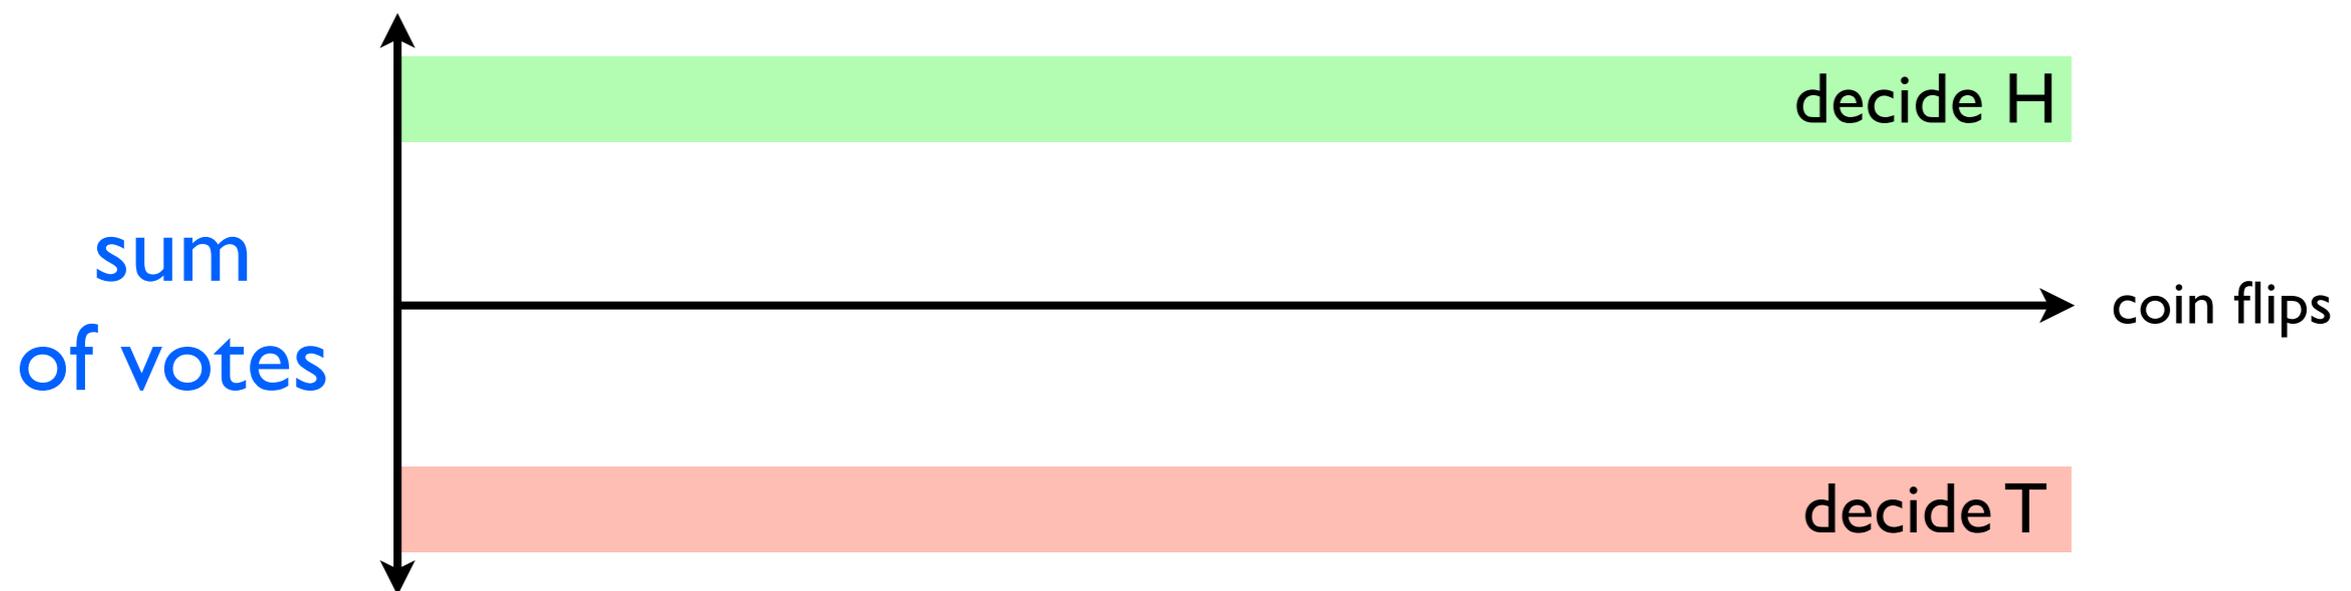
# Adaptive on-line annotation

- How many labels do we need until we are certain of an annotation?

- Coin flipping:  vs  what is the bias?



- Wald (1945): sequential probability ratio test (SPRT)



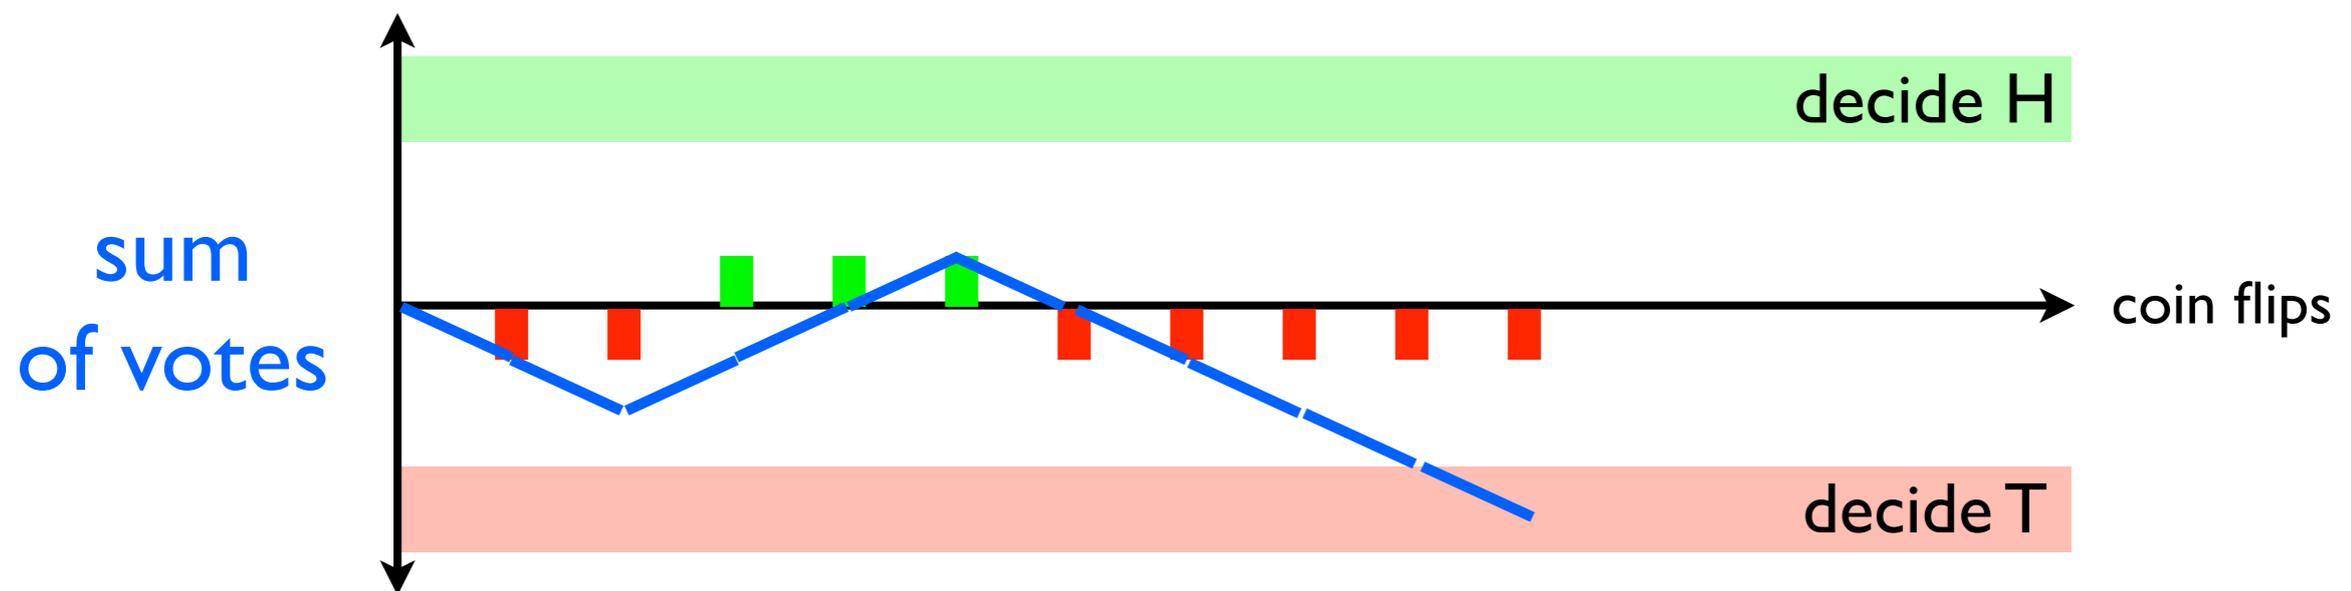
# Adaptive on-line annotation

- How many labels do we need until we are certain of an annotation?

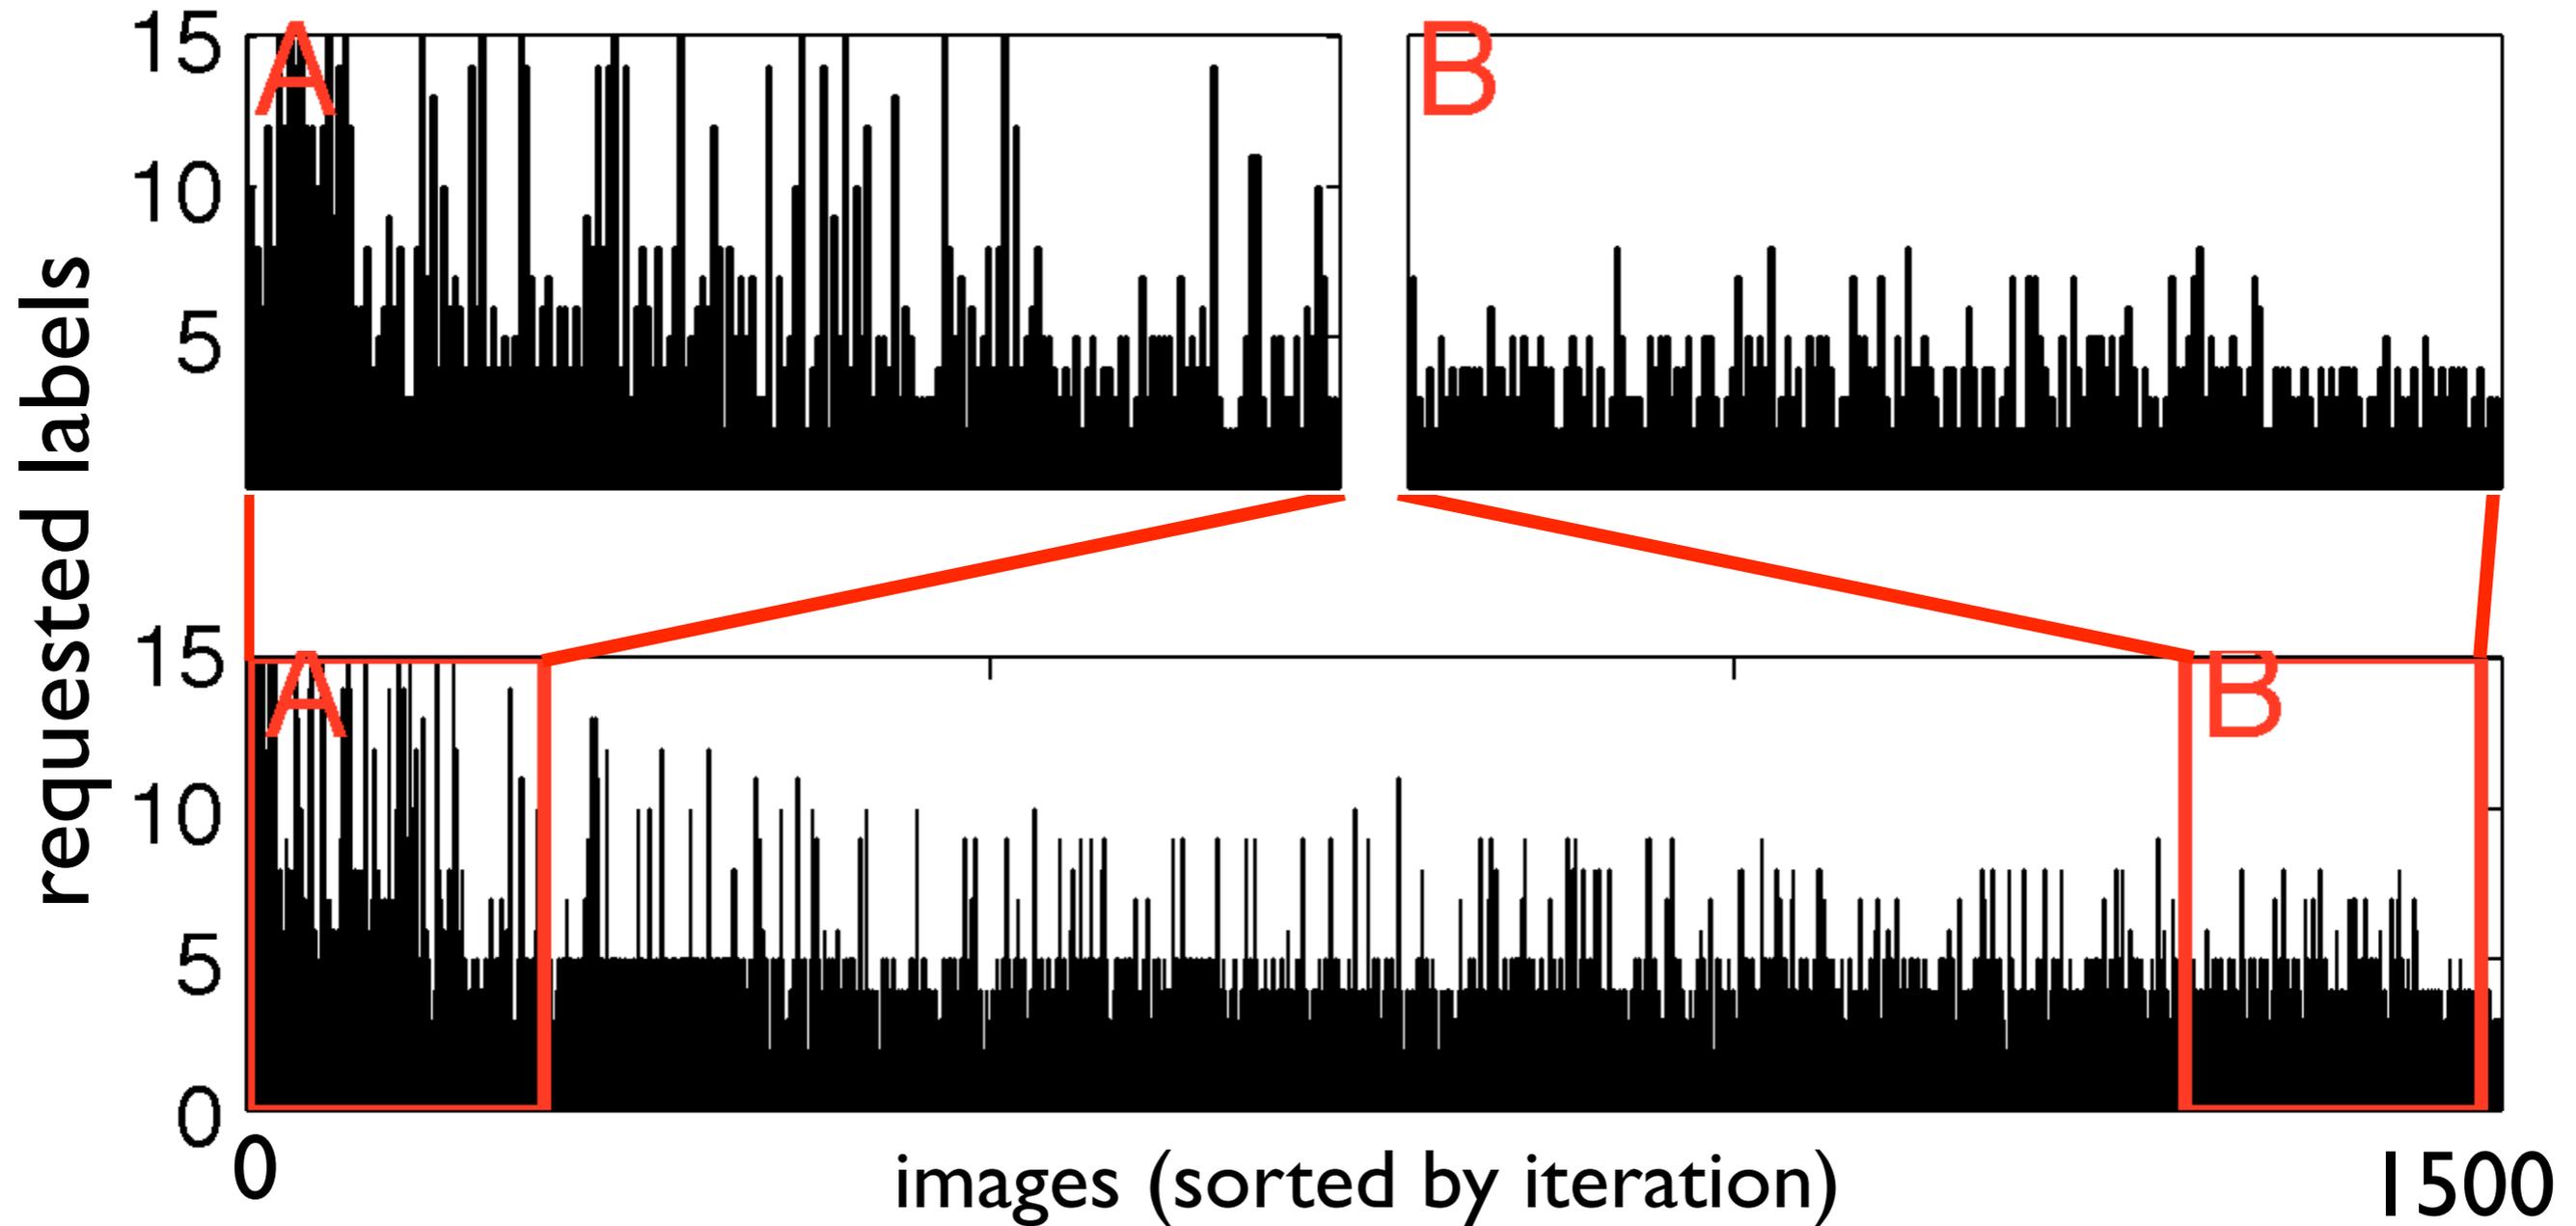
- Coin flipping:  vs  what is the bias?

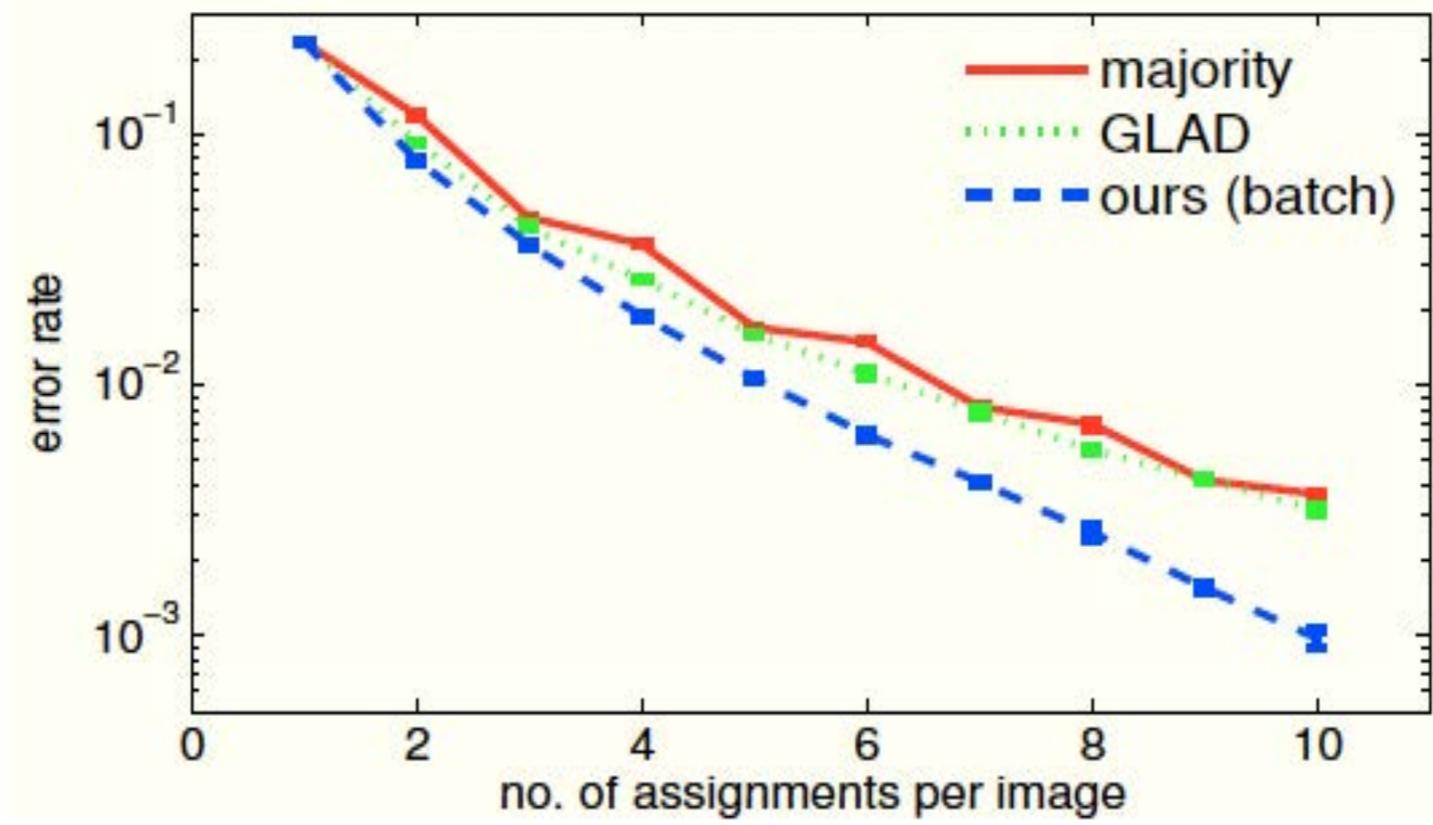


- Wald (1945): sequential probability ratio test (SPRT)



# Example





# Location annotations



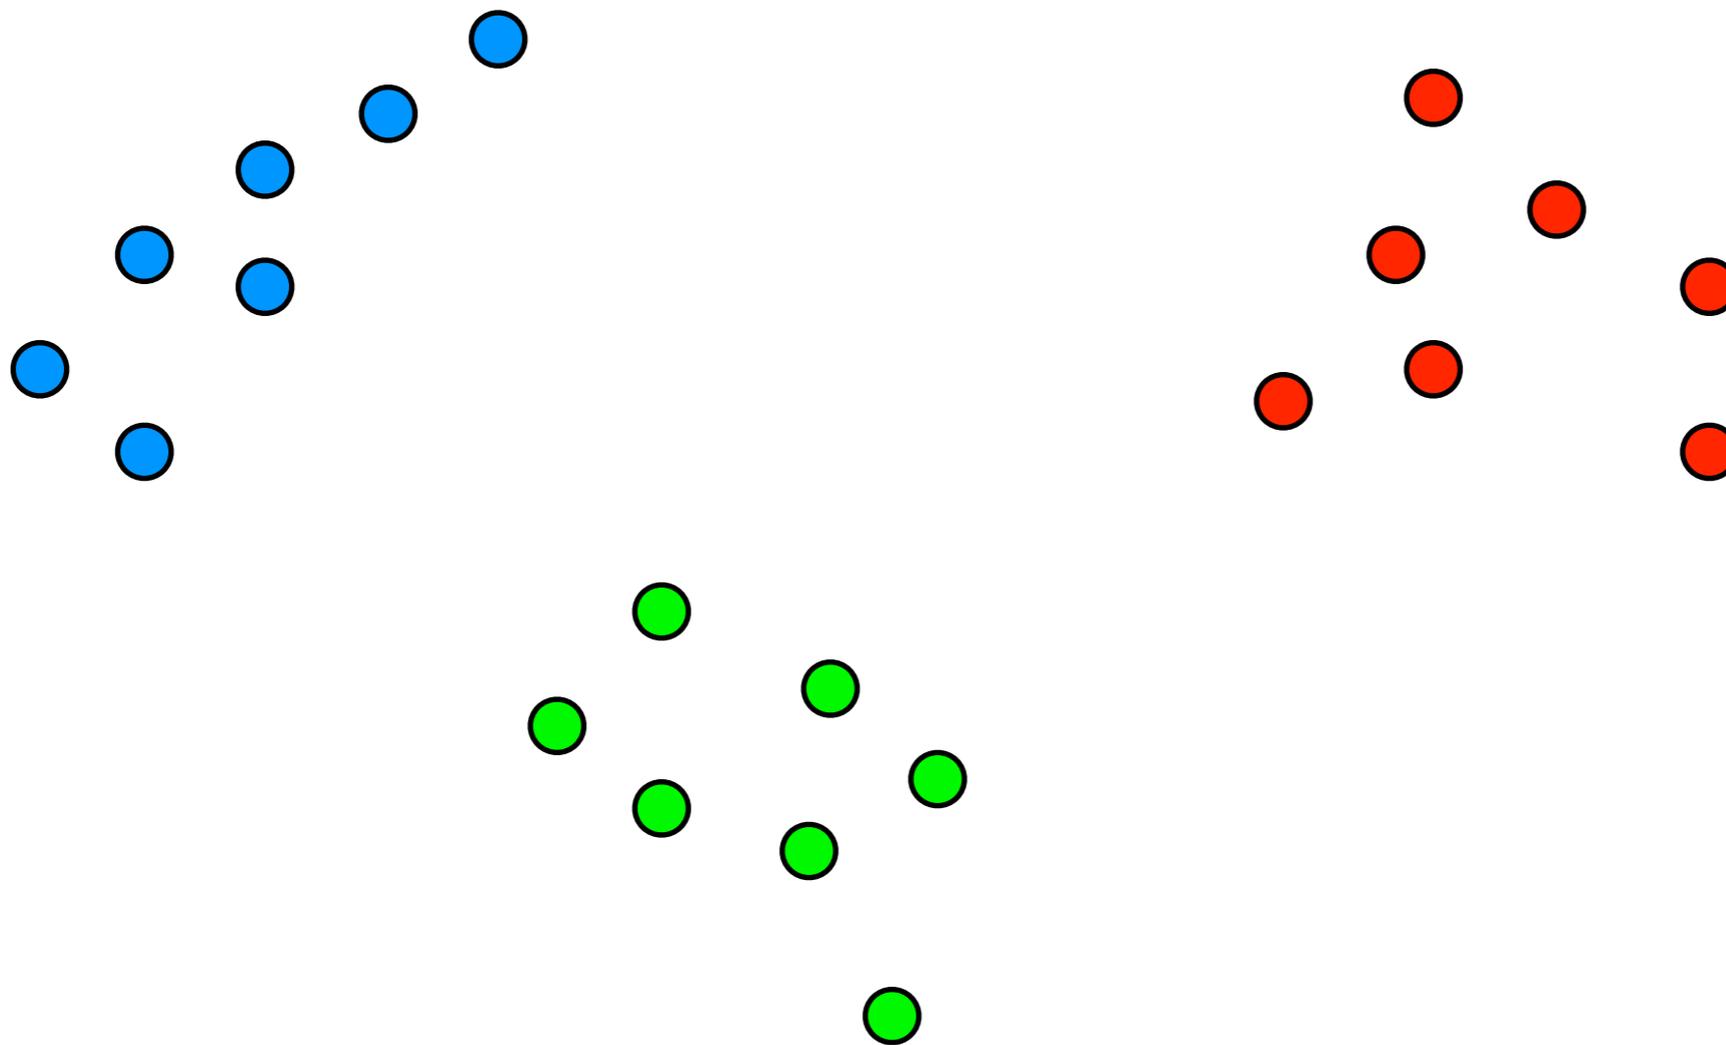
# Crowdclustering

[Gomes et al., NIPS 2011]

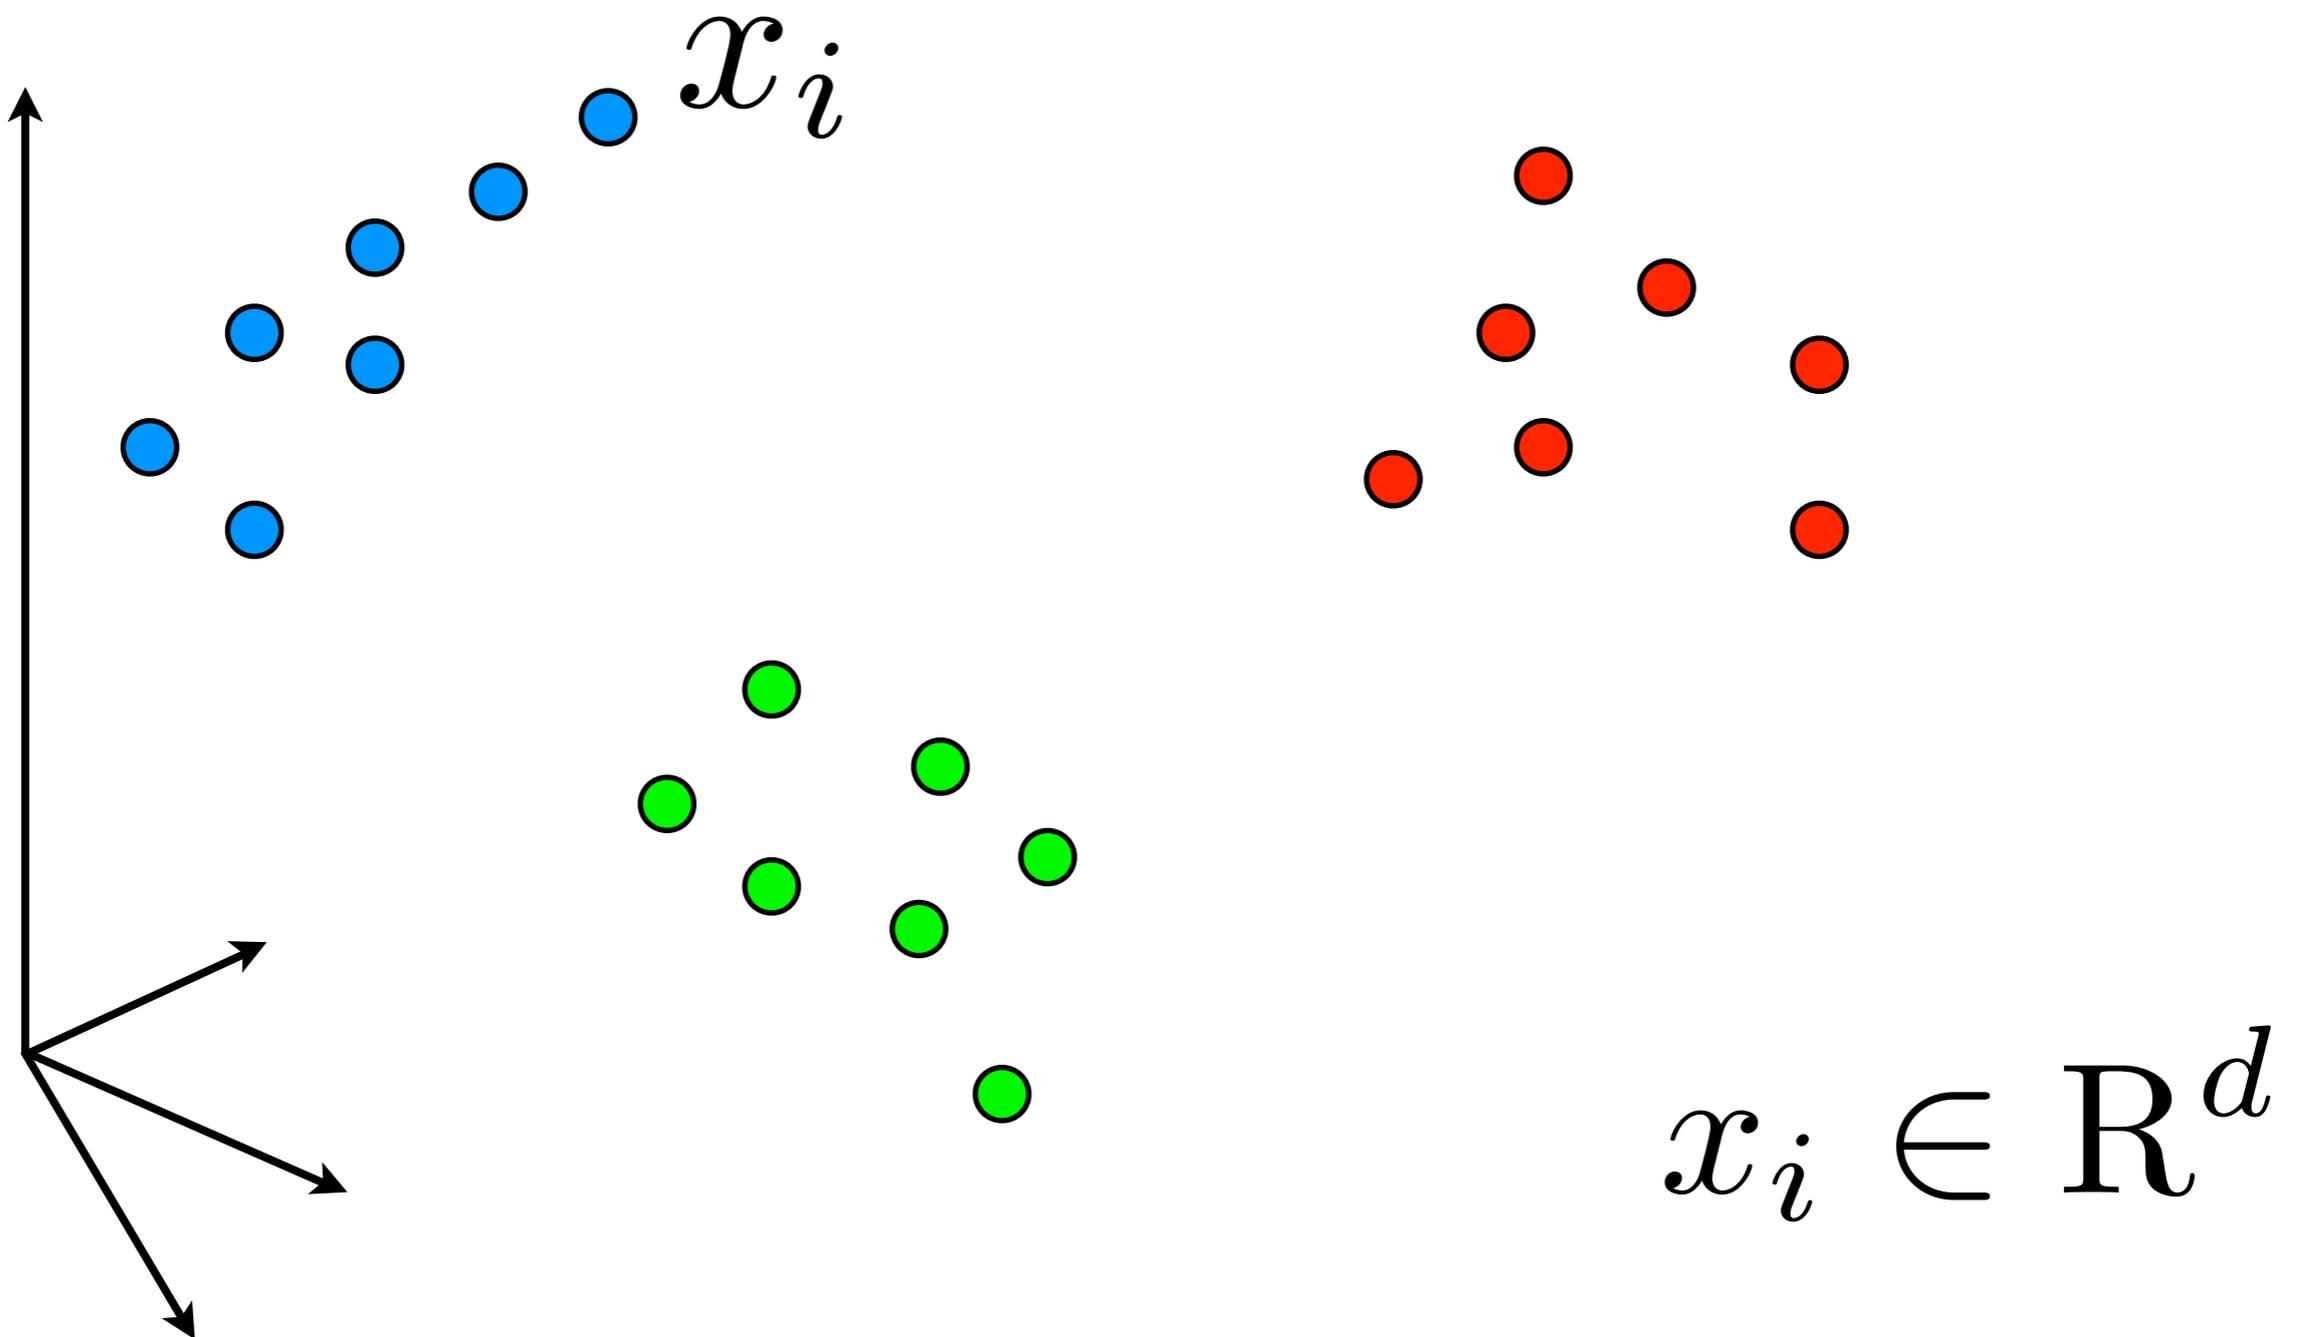
# Organizing visual knowledge



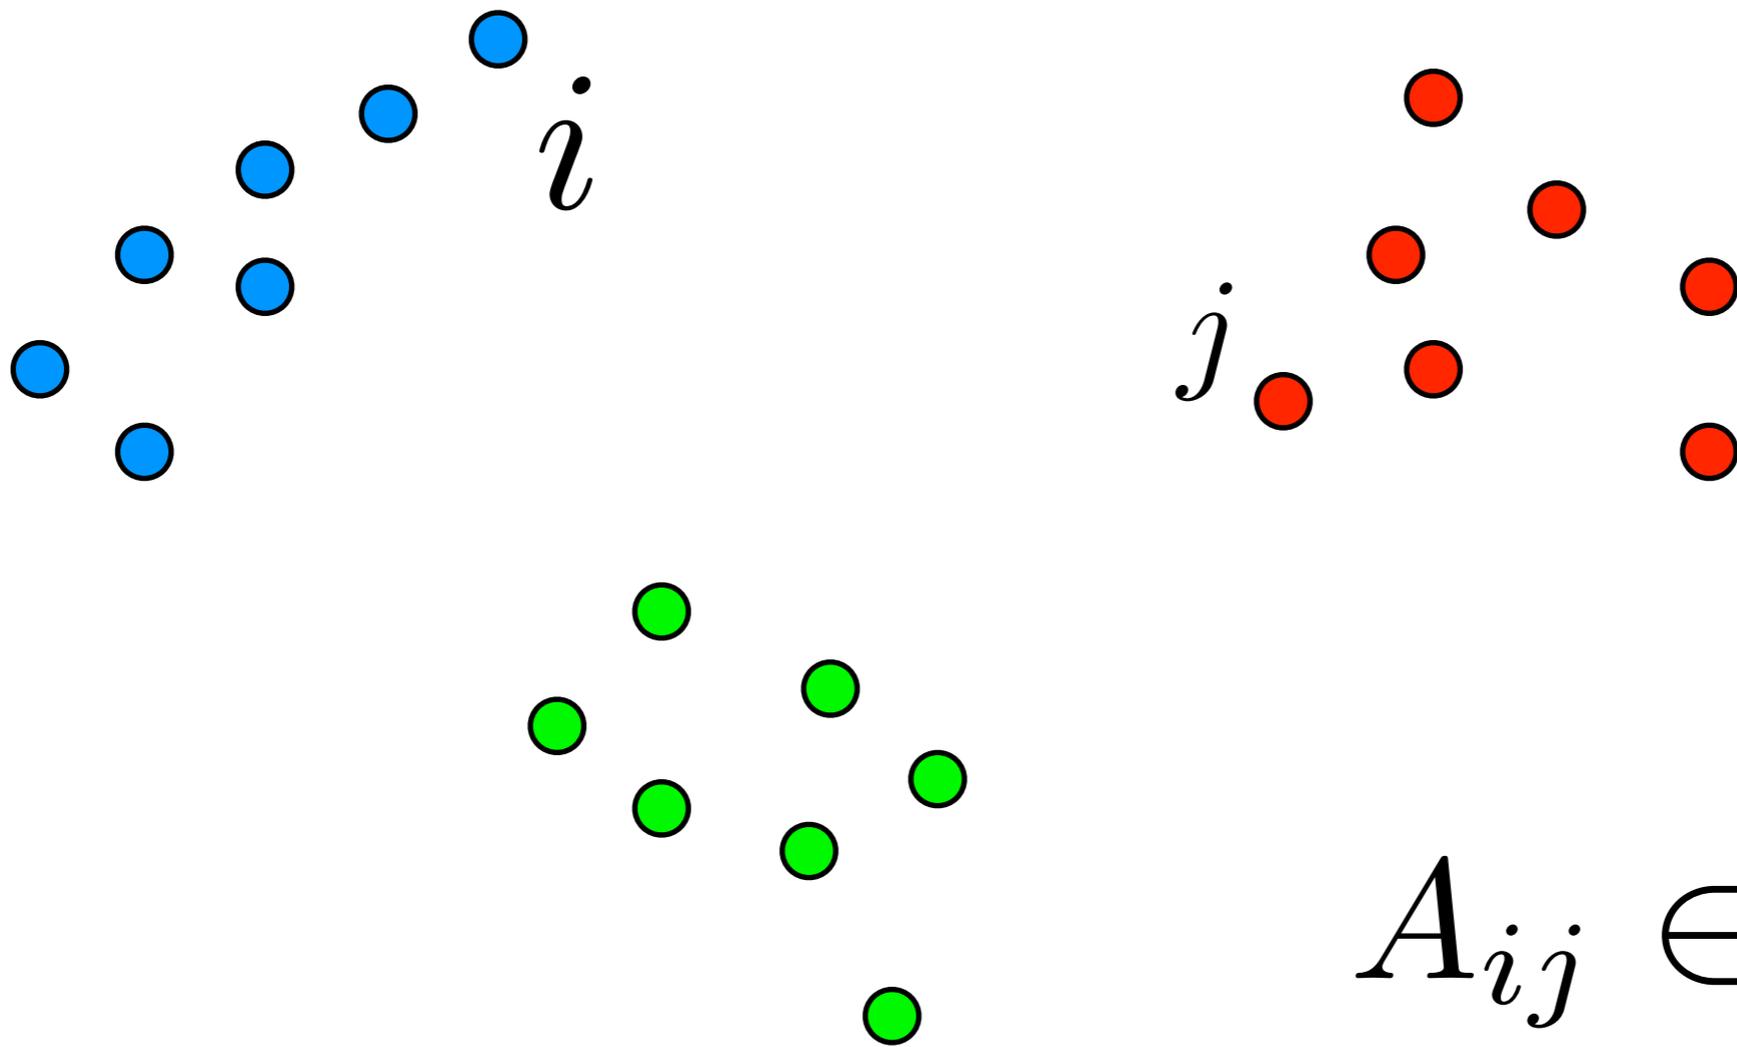
# Clustering



# Need



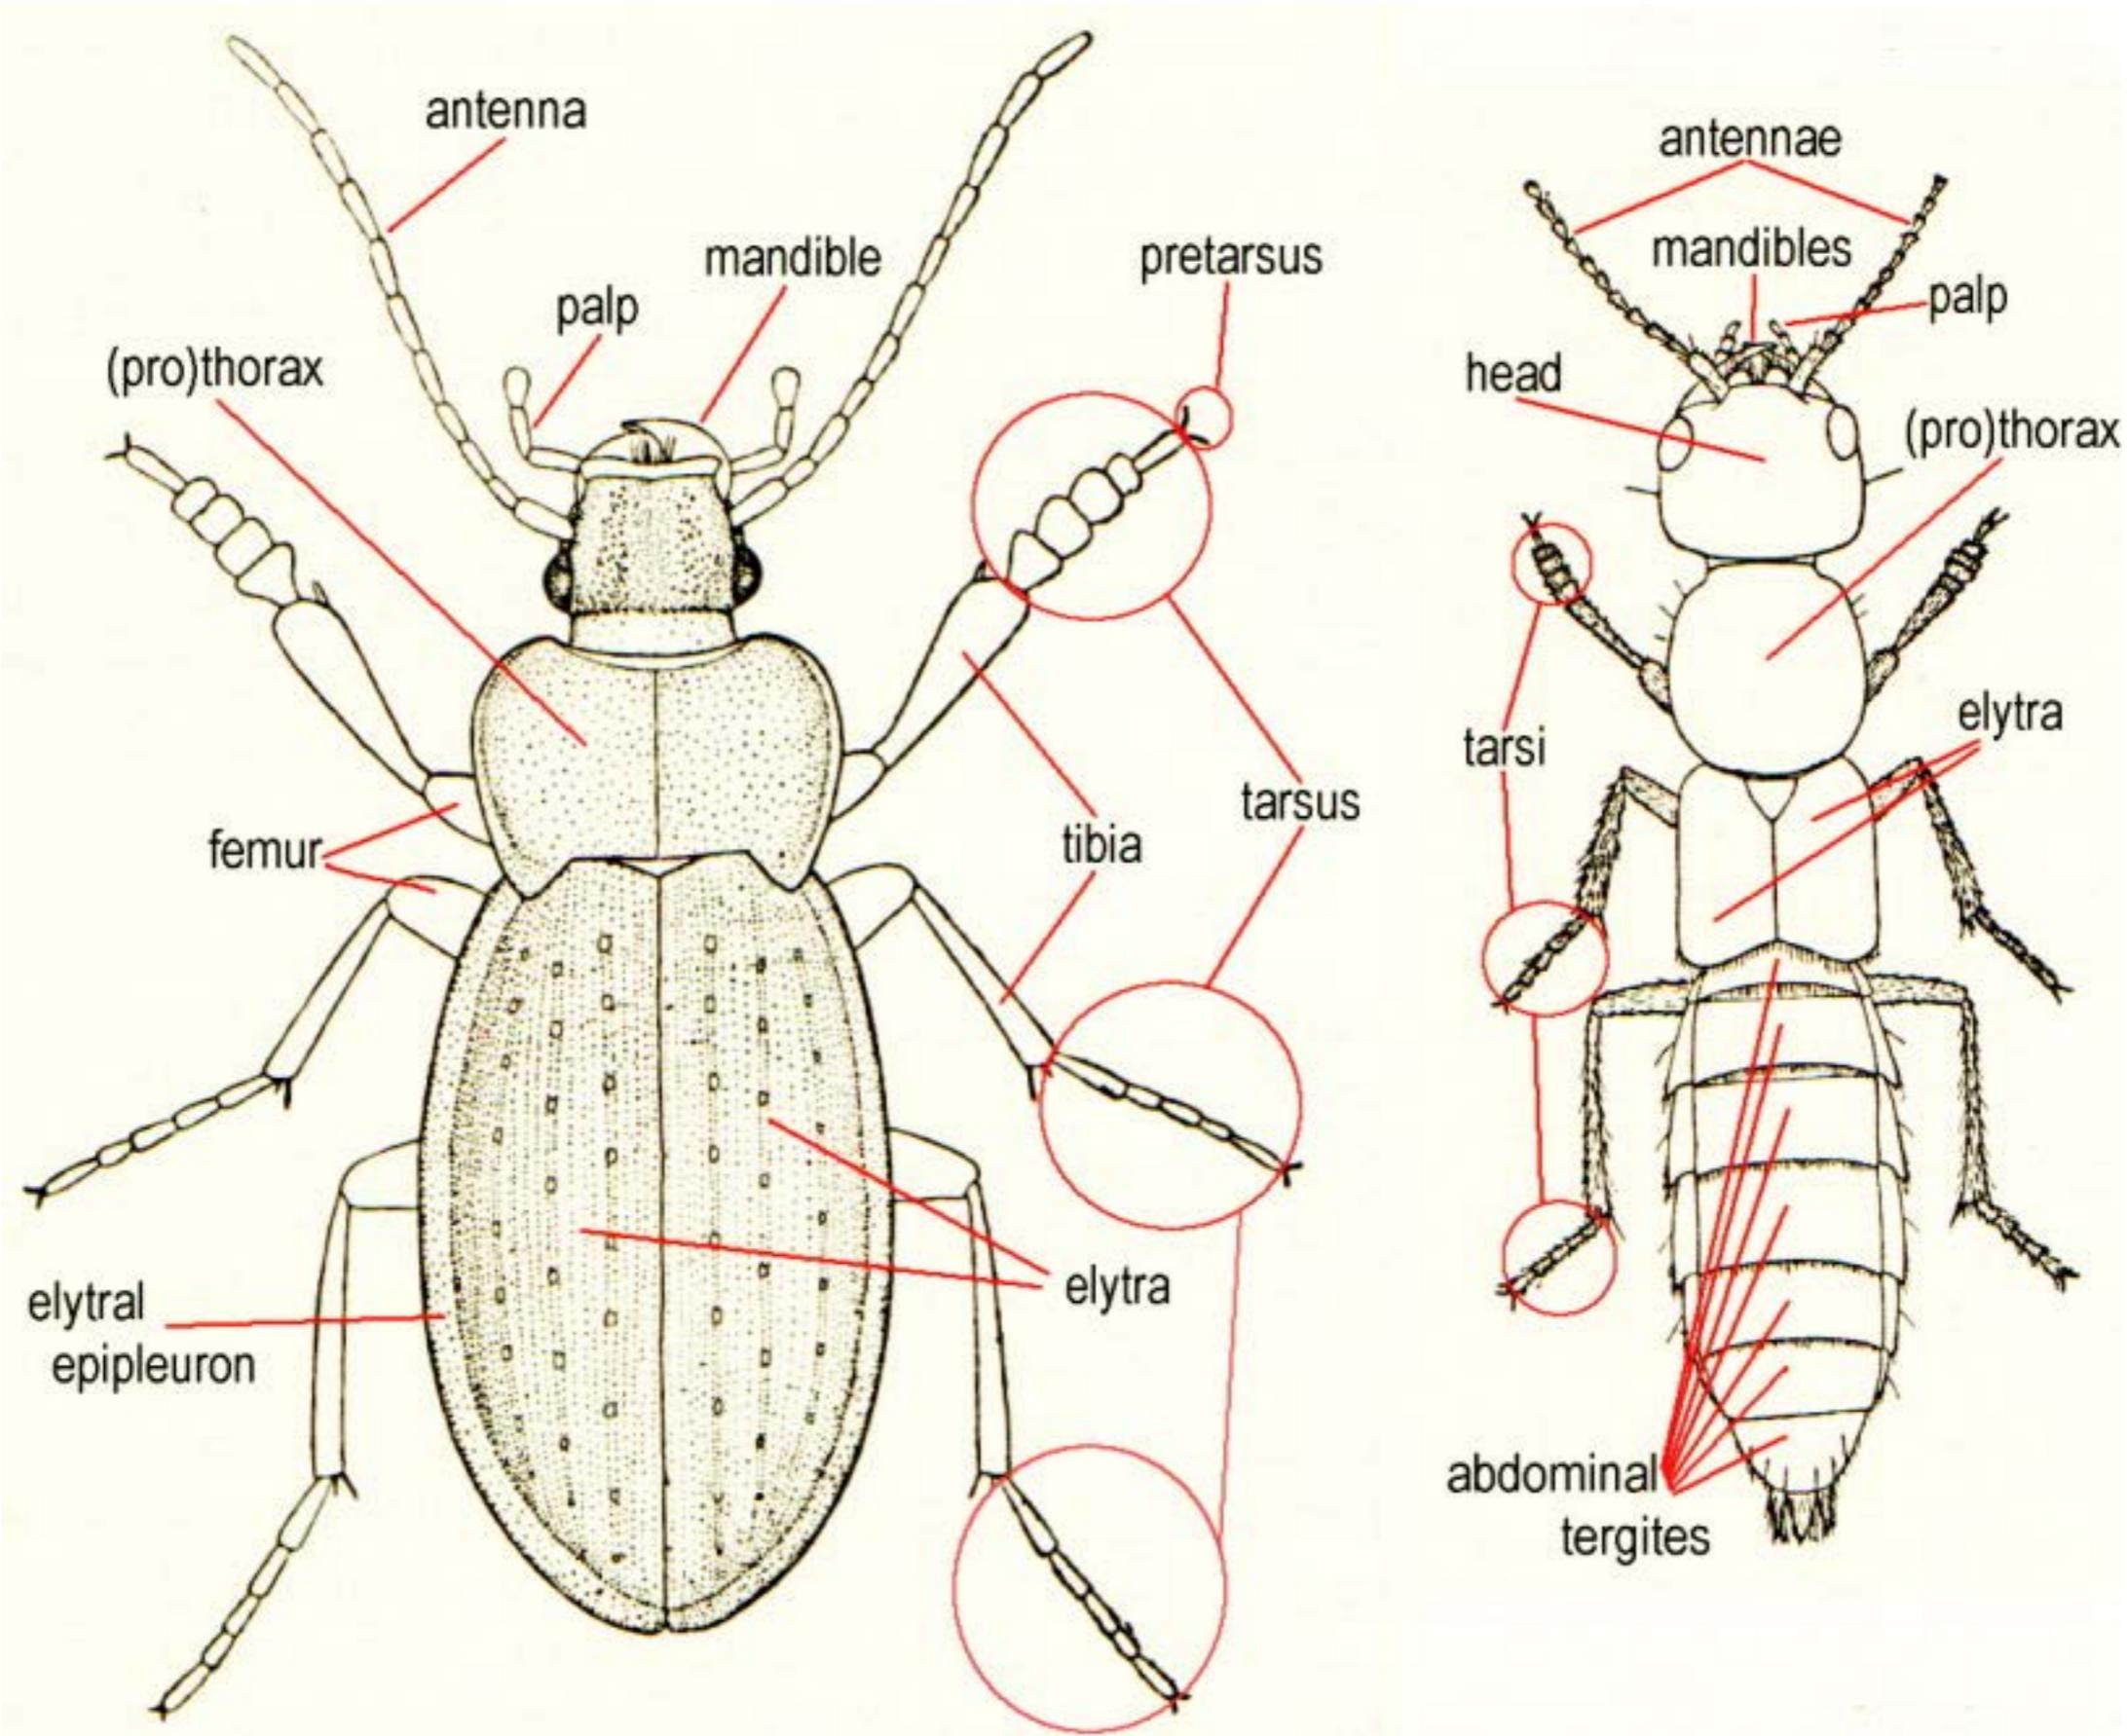
# Need

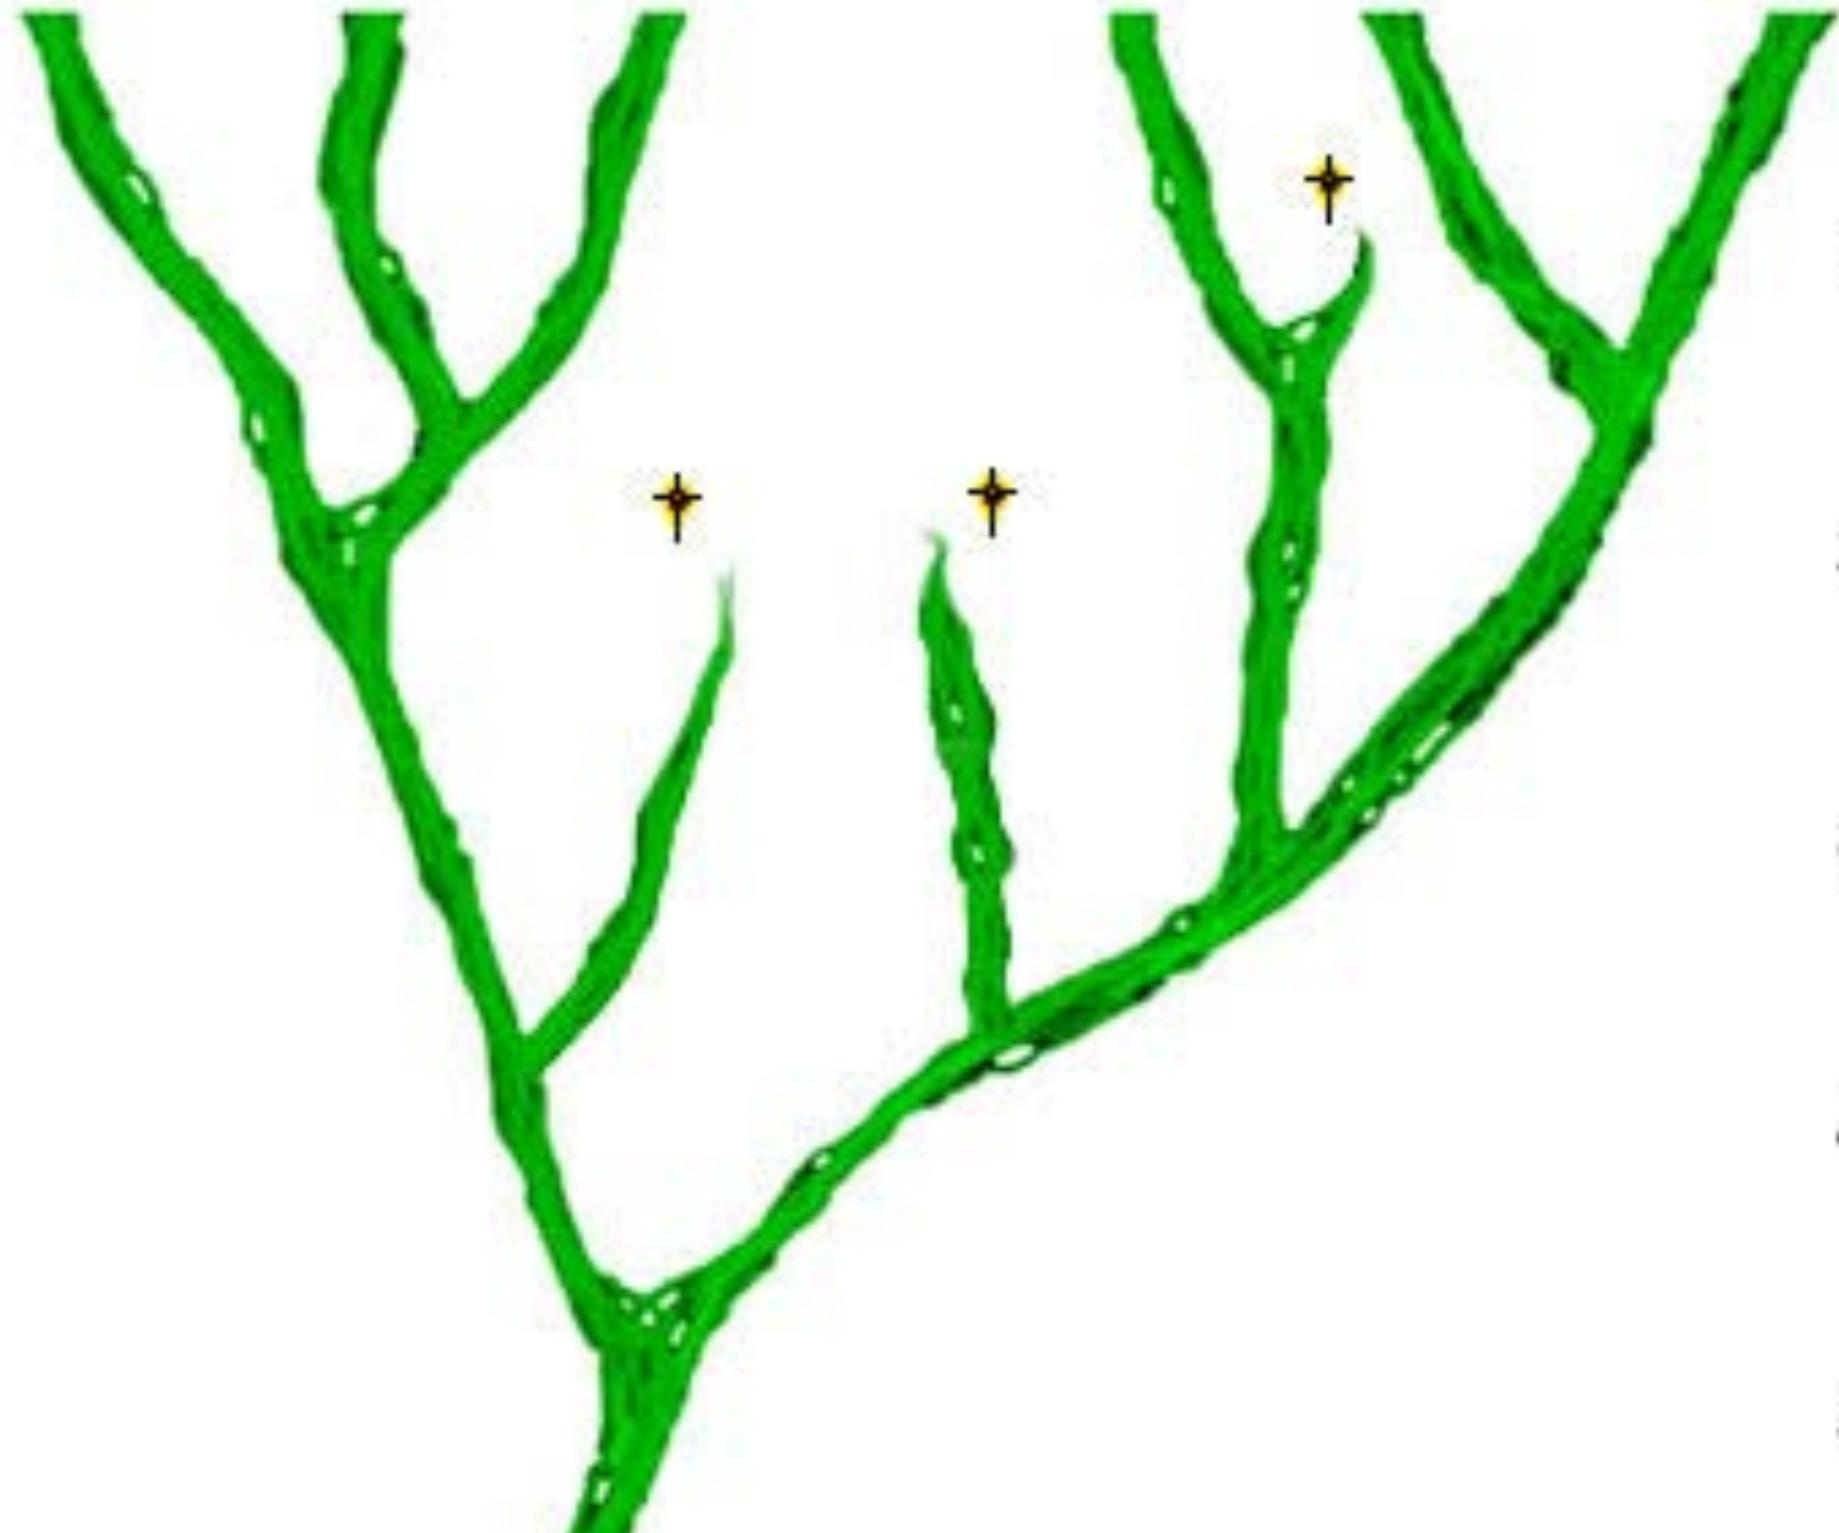


# Cluster this...



# Metric for clustering





**1 million years ago**

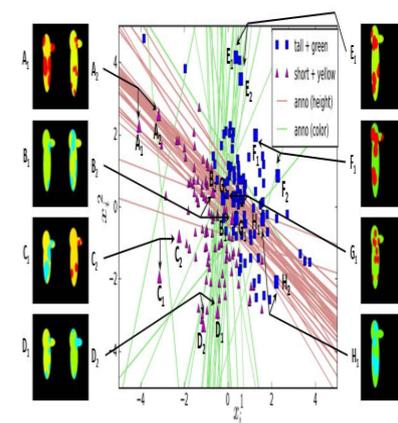
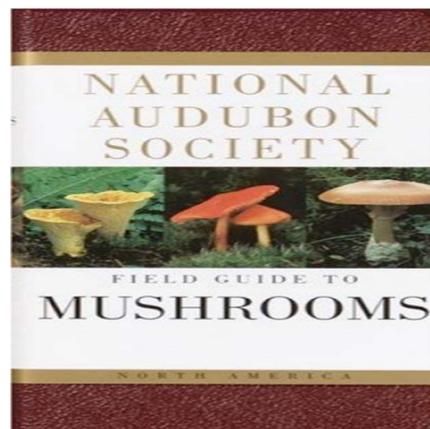
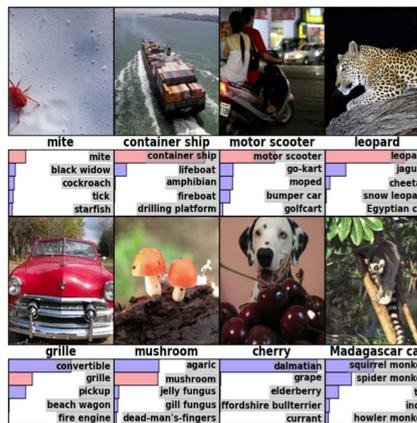
**2 million years ago**

**3 million years ago**

**4 million years ago**

**5 million years ago**

Showing 1 - 18 of 67,880 Results, sorted by 'Price: High-to-Low'  
[3 June 2011]



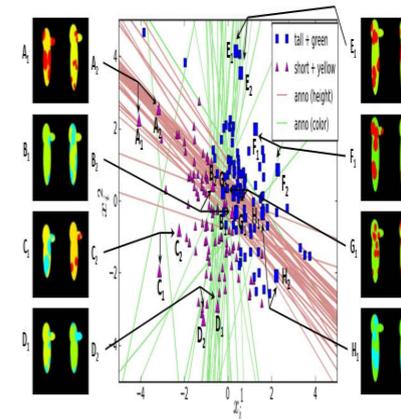
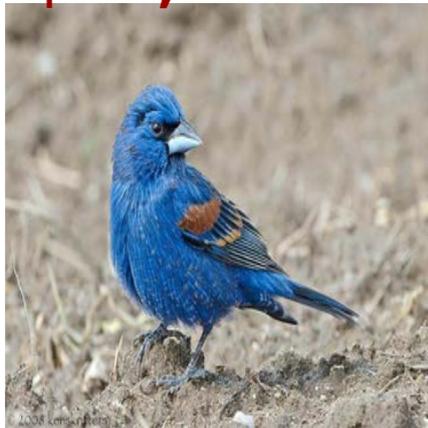
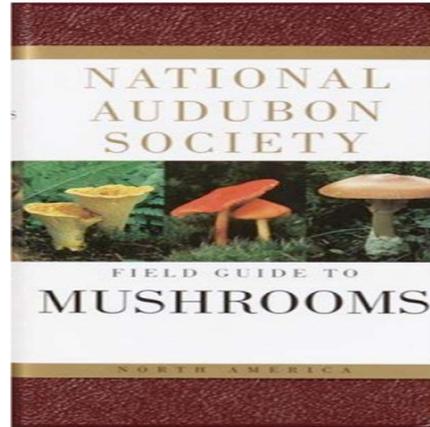
Showing 1 - 18 of 67,880 Results, sorted by 'Price: High-to-Low'  
[3 June 2011]



\$85,633.93



\$86,999.99



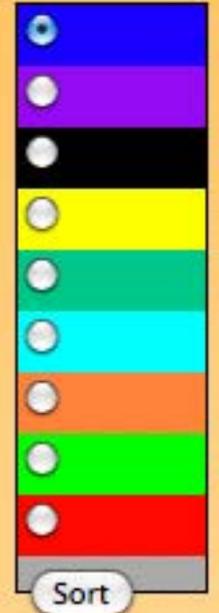
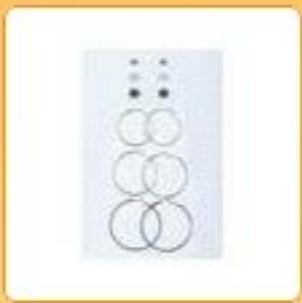


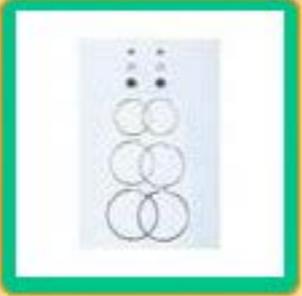




# Challenges

- Lots of images
- Annotators working on small subsets
- Categorization criteria may differ

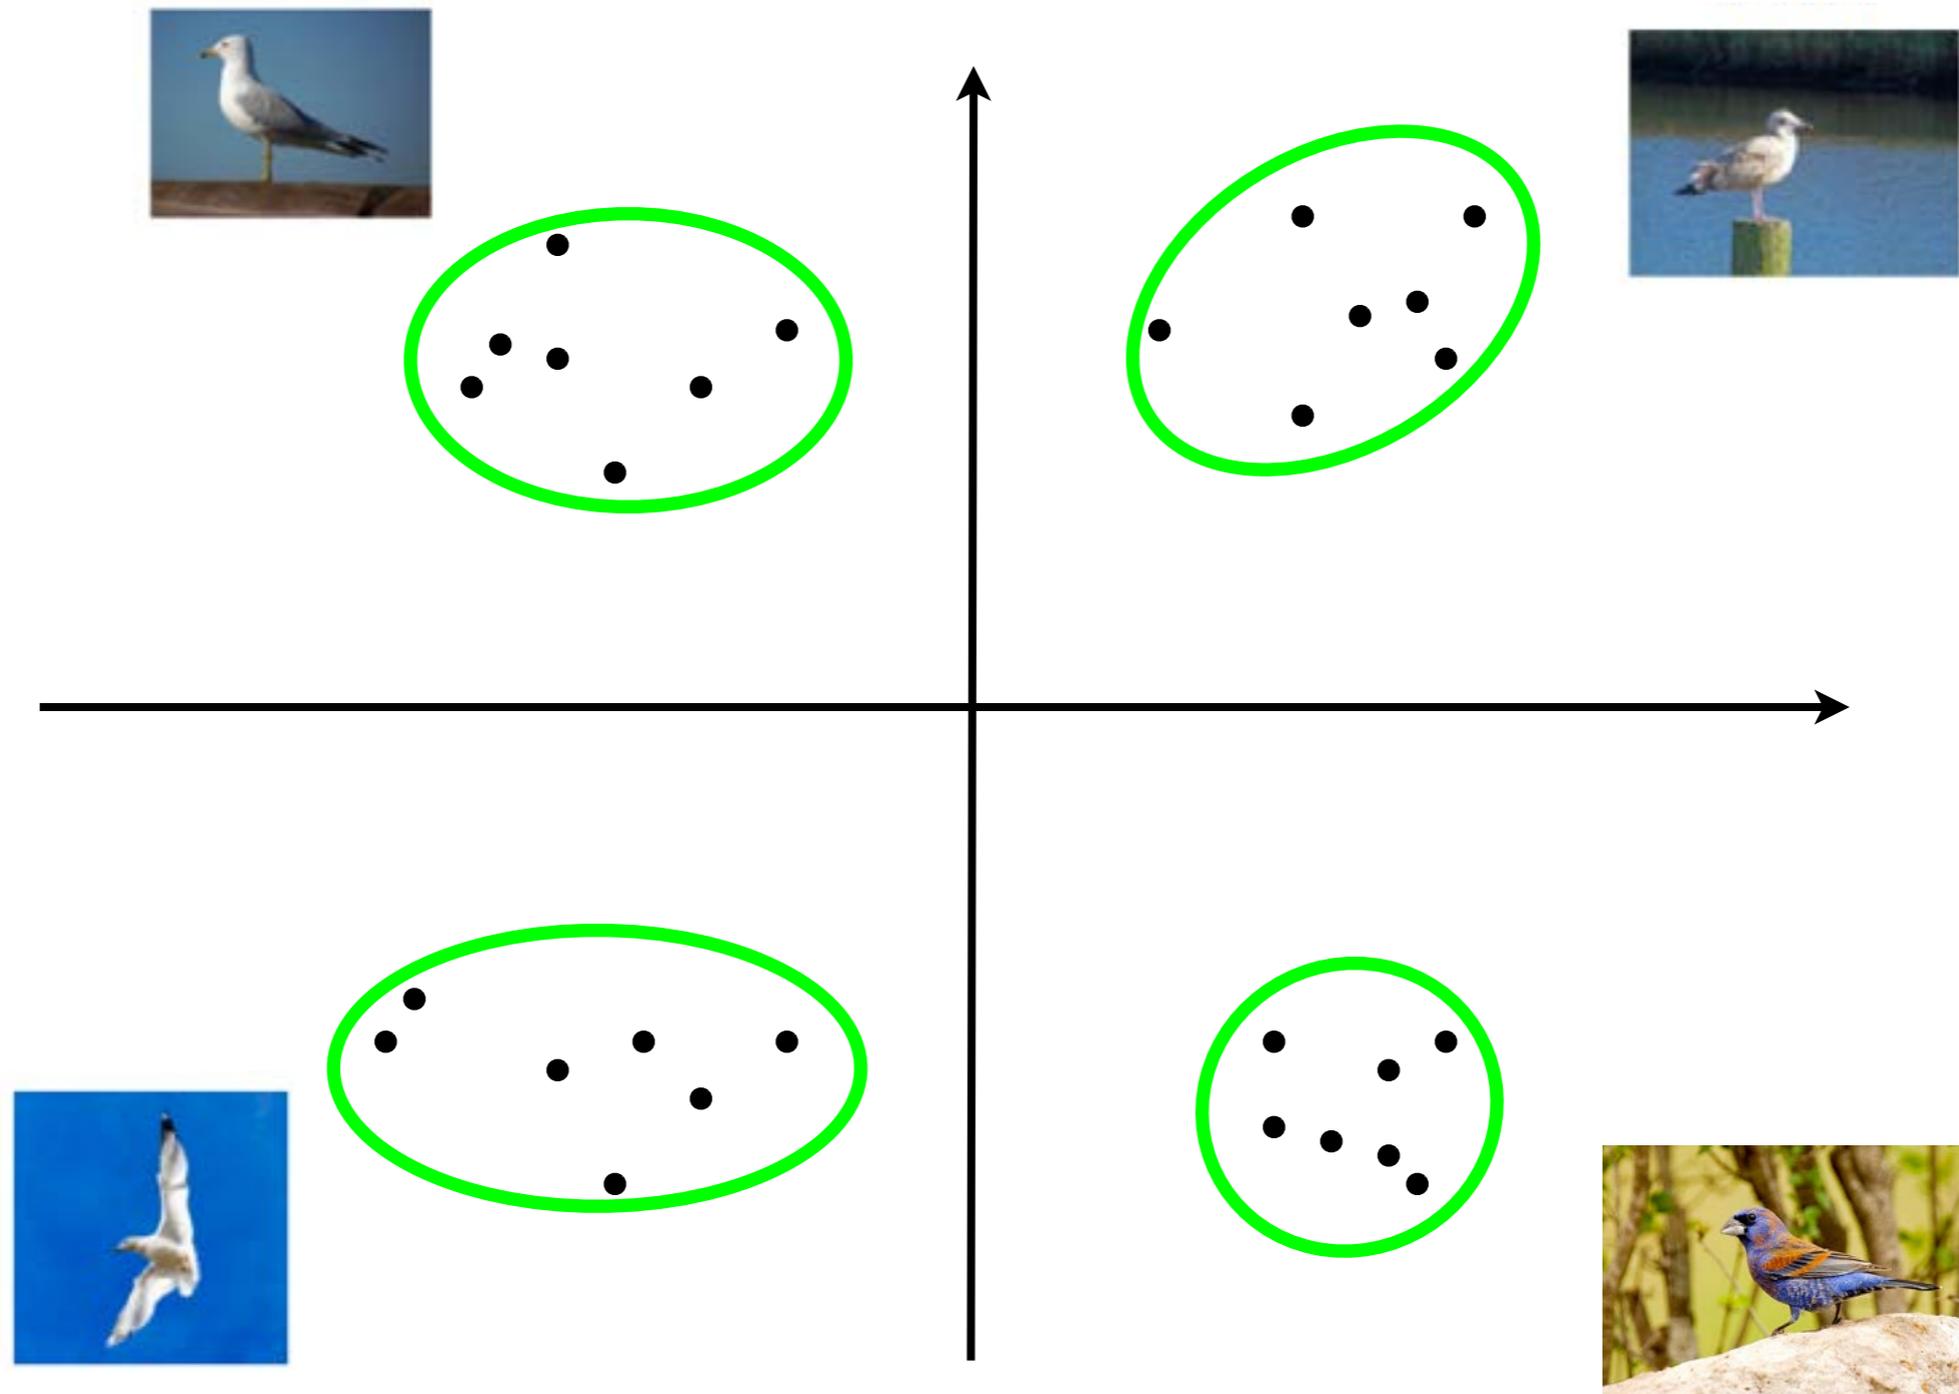




- Blue
- Purple
- Black
- Yellow
- Green
- Cyan
- Orange
- Light Green
- Red

Sort

# Generative Model



- Assume images drawn from clusters in an embedding space
- Each annotator corresponds to an inner product in that space

# How do we aggregate the results from the crowd?



cluster 1



cluster 2



- Annotator 1: sensitive to ground vs. air

# How do we aggregate the results from the crowd?



cluster 1

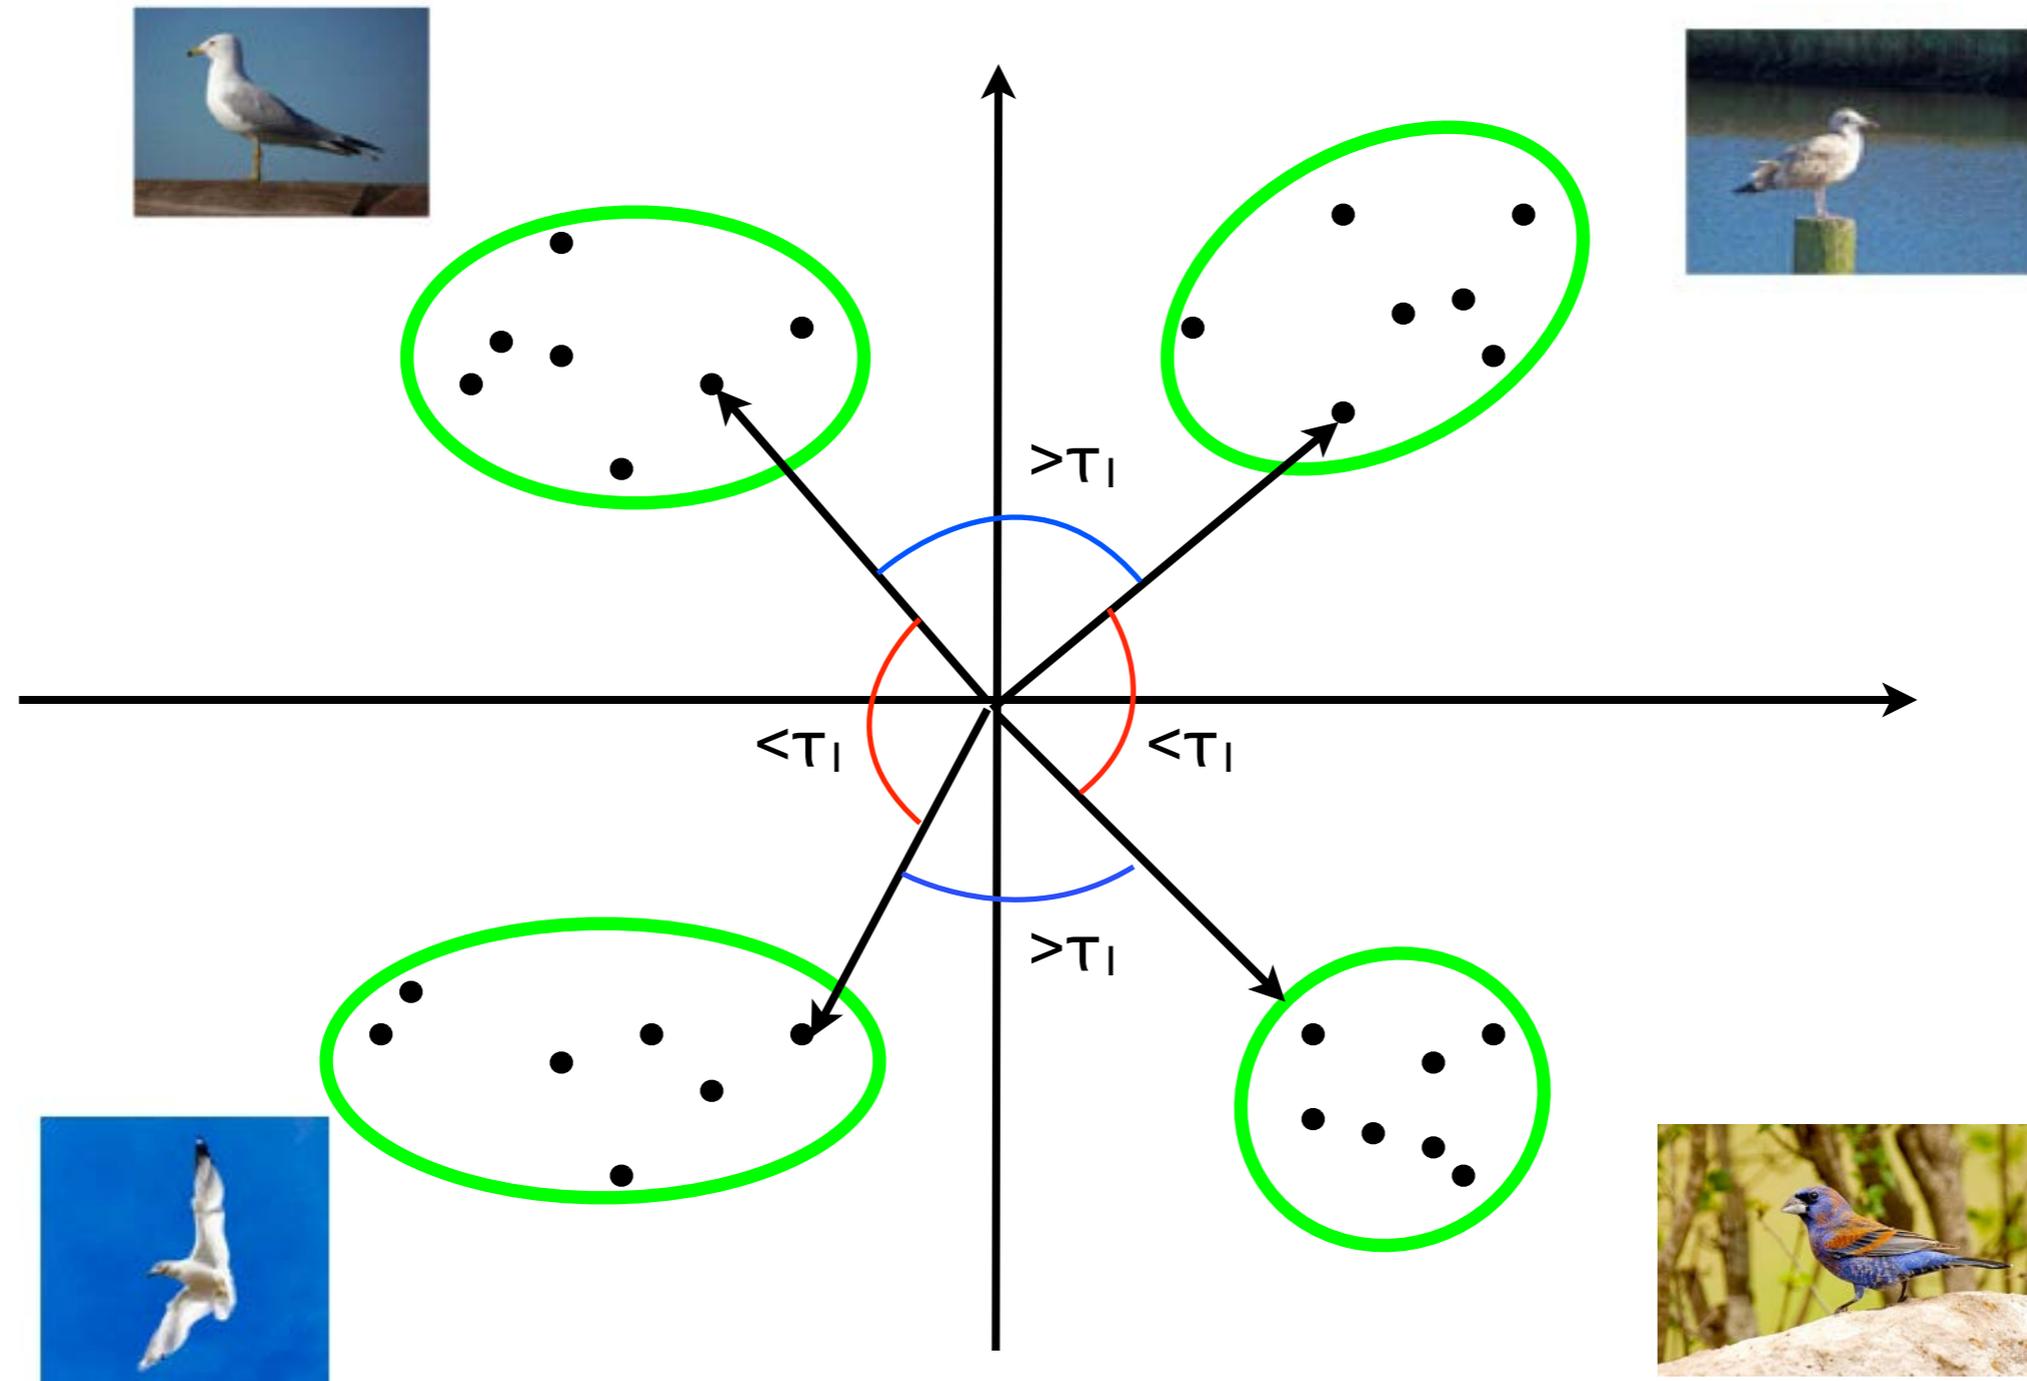


cluster 2

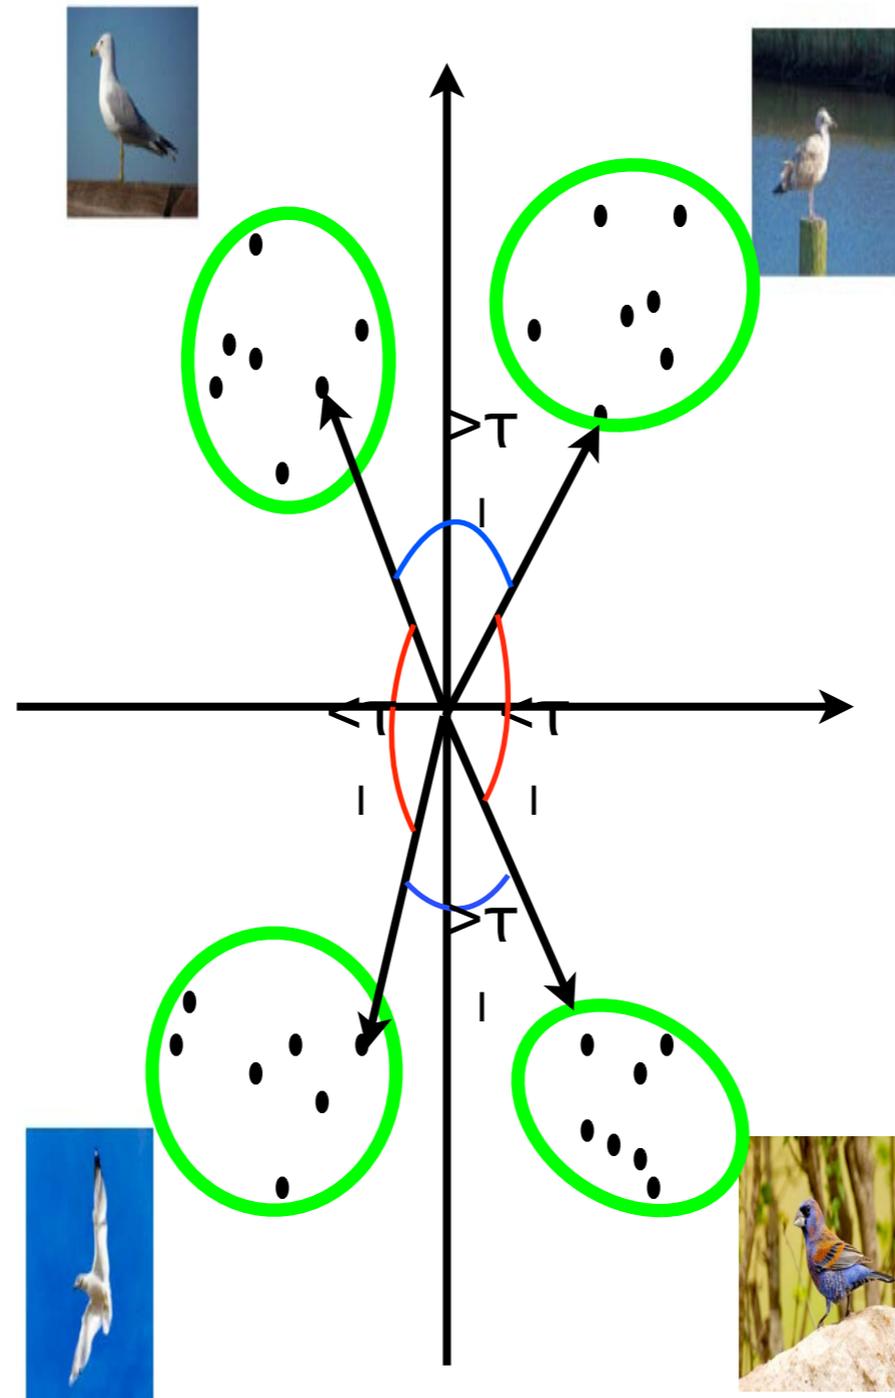


- Annotator 2: sensitive to left vs right

# Generative Model

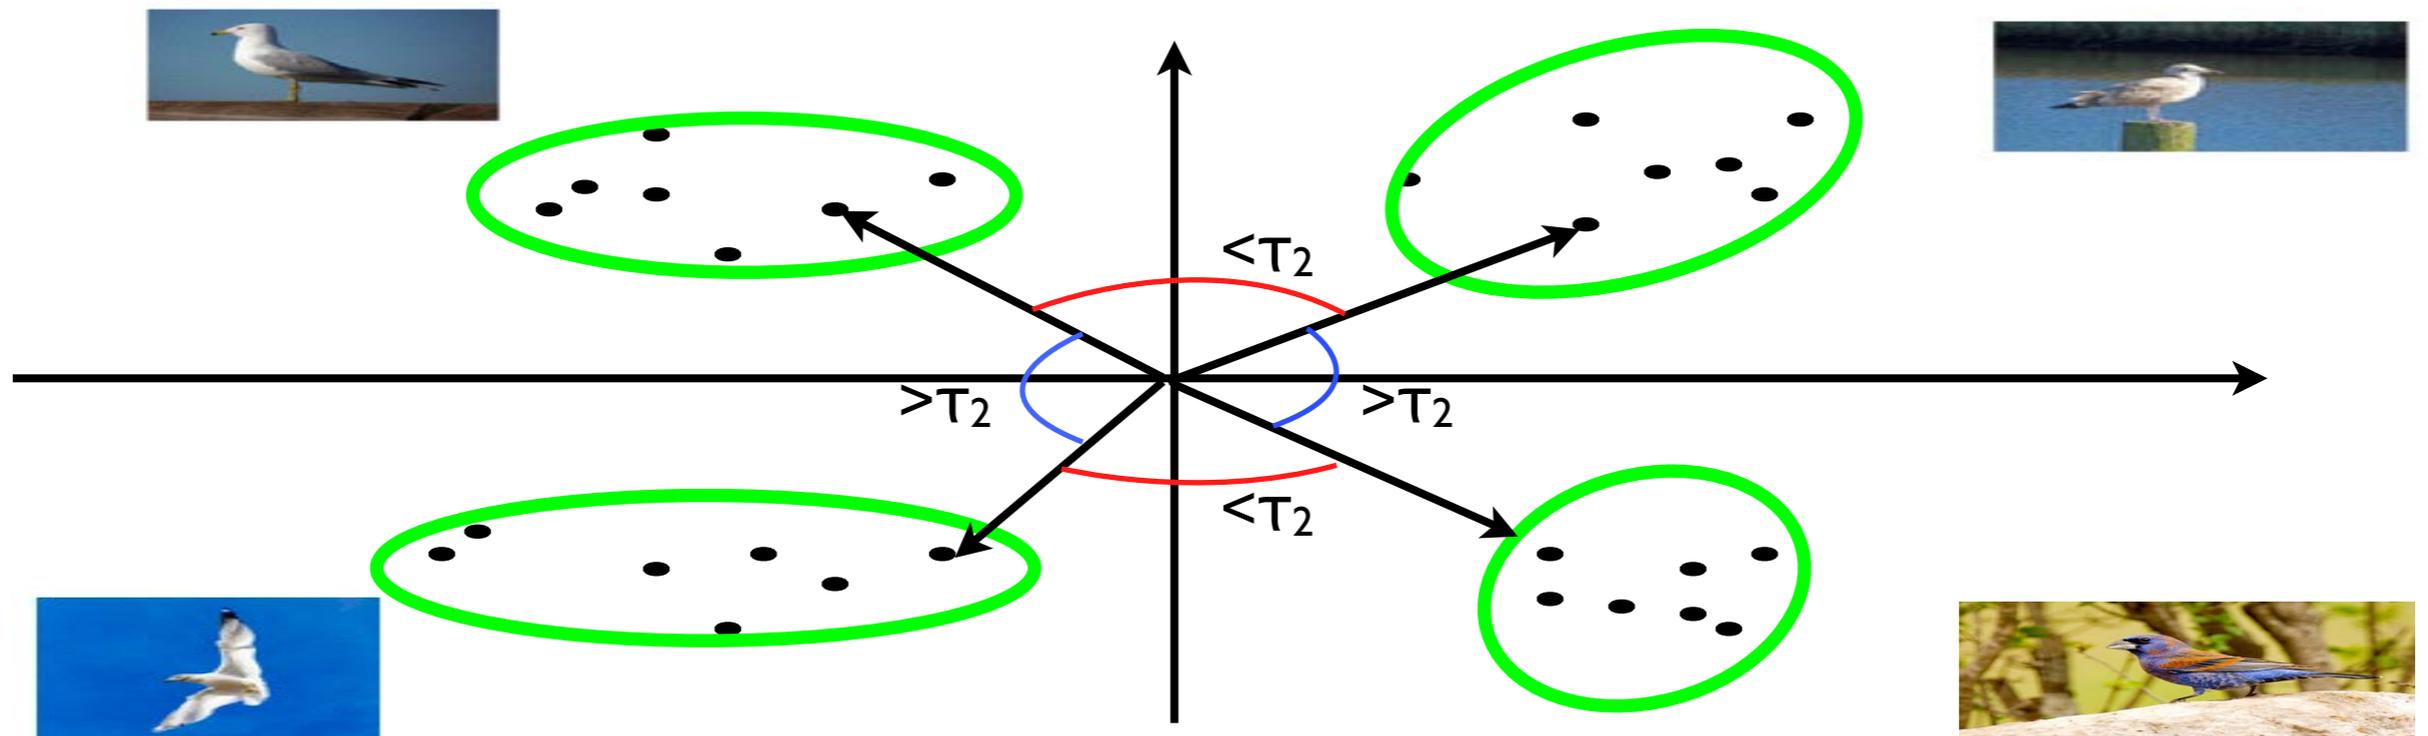


# Generative Model

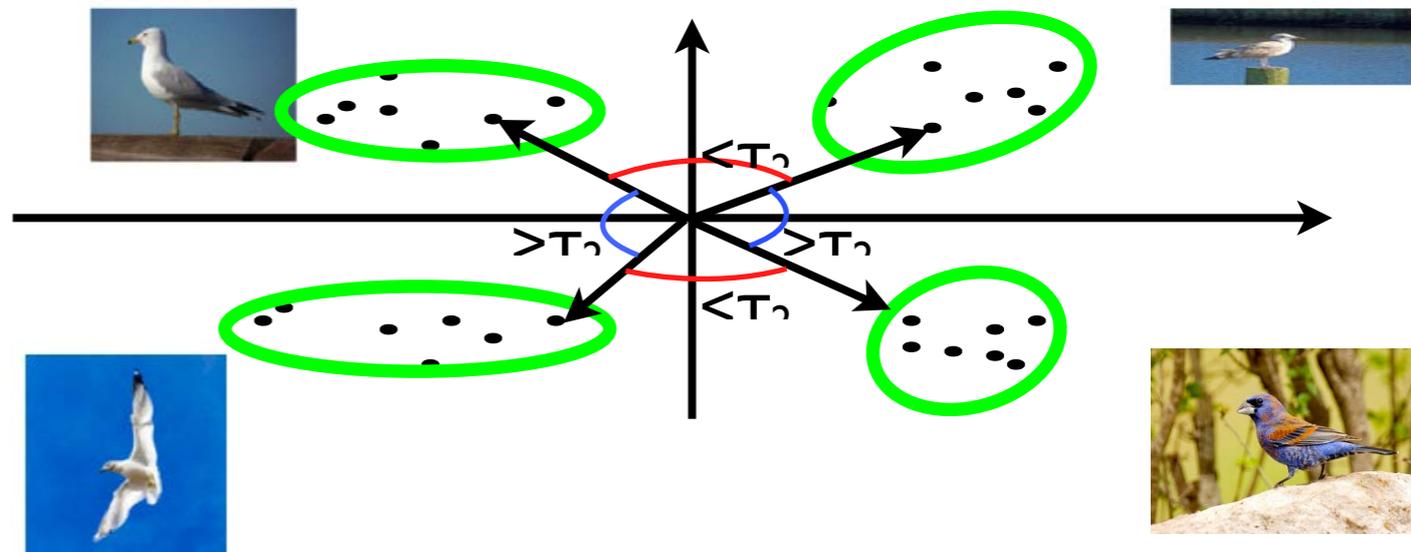
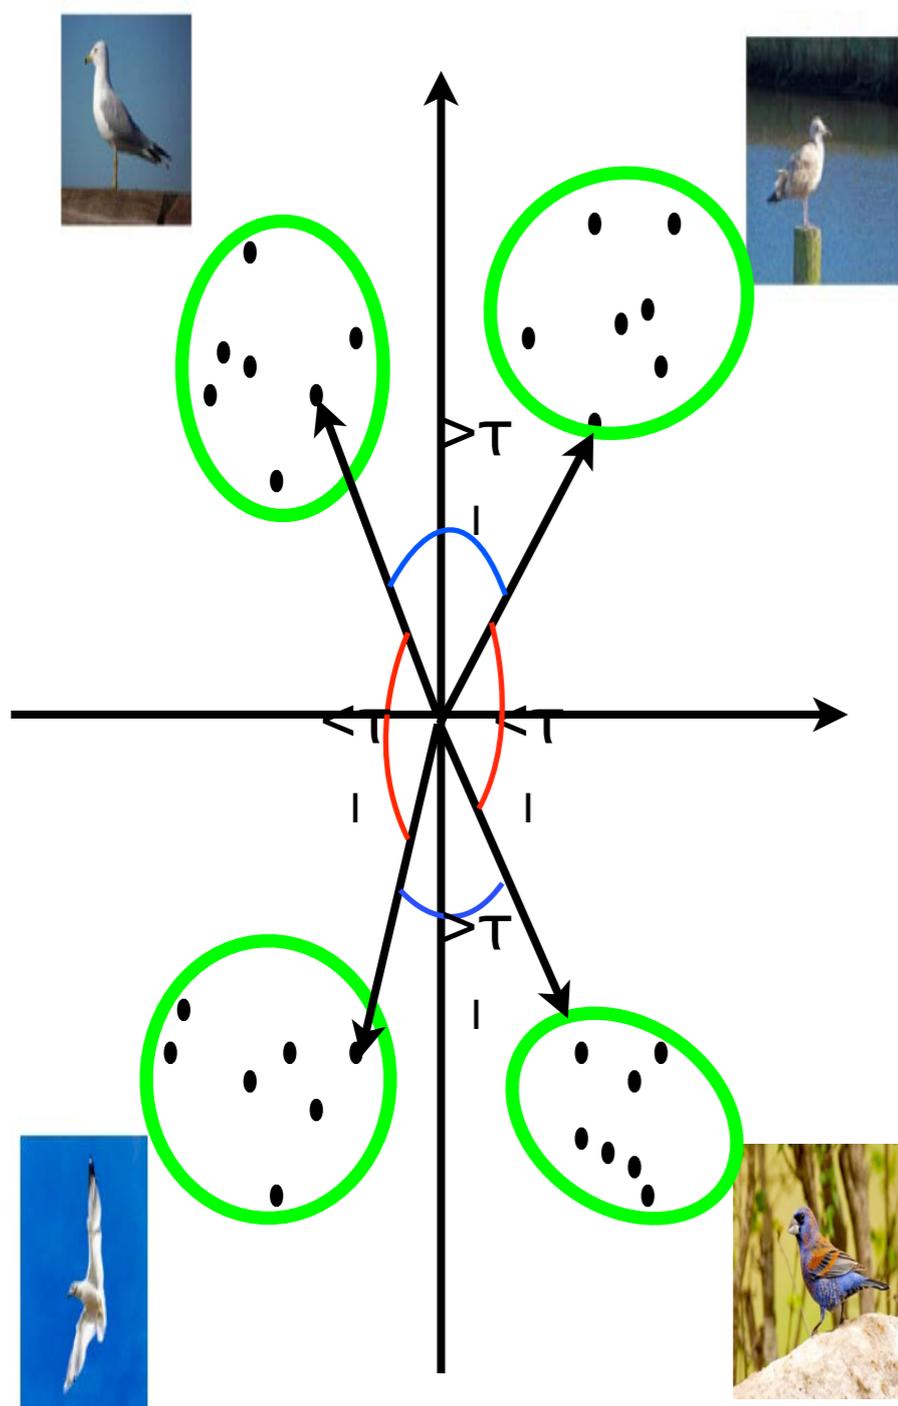


- Annotator I: sensitive to ground vs. air

# Generative Model



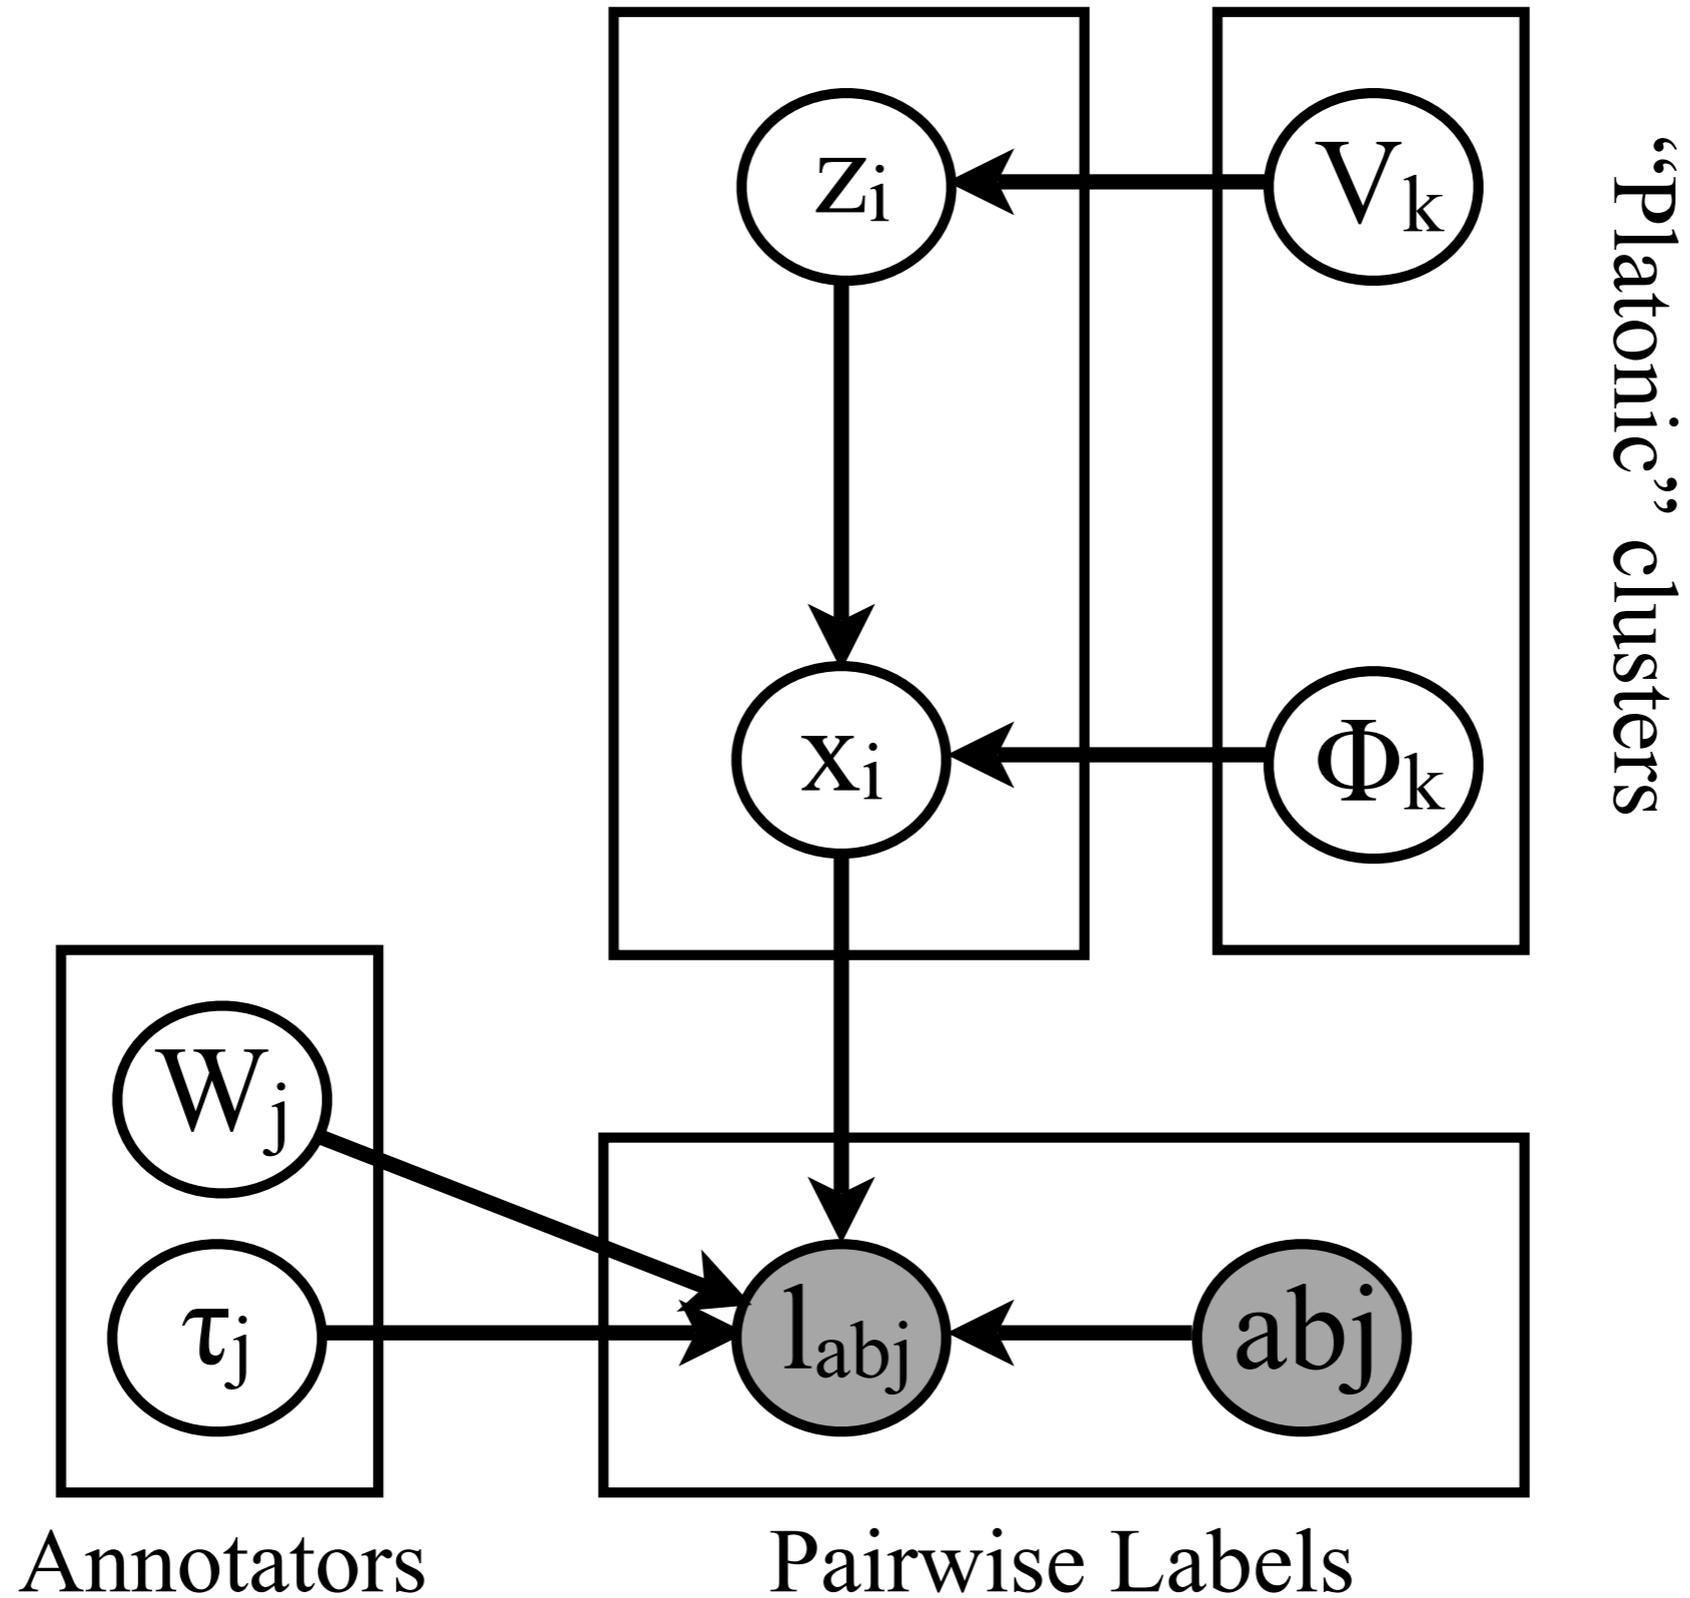
- Annotator 2: sensitive to left vs. right



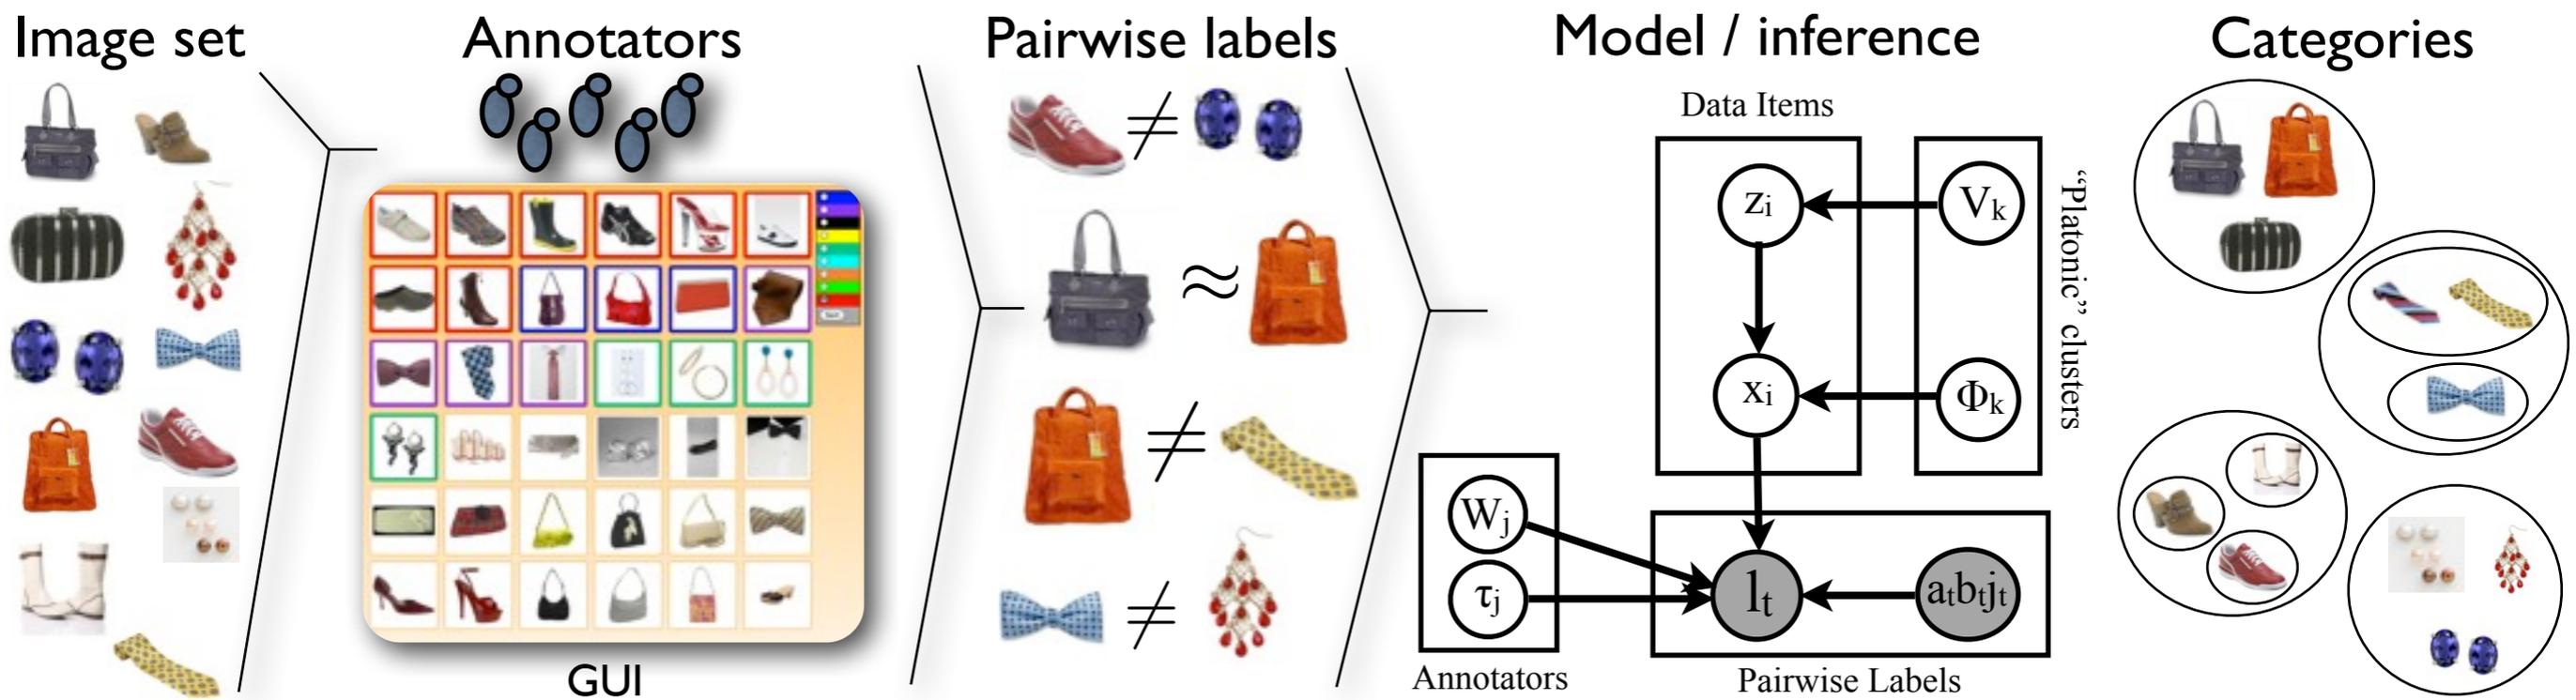
$$p(l_{ab2} = 1 | \mathbf{x}_a, \mathbf{x}_b, \mathbf{W}_2, \tau_2) = \frac{1}{1 + \exp\{- (\mathbf{x}_a^T \mathbf{W}_2 \mathbf{x}_b - \tau_2)\}}$$

# Model

Data Items



# Pipeline



# Experiments

1



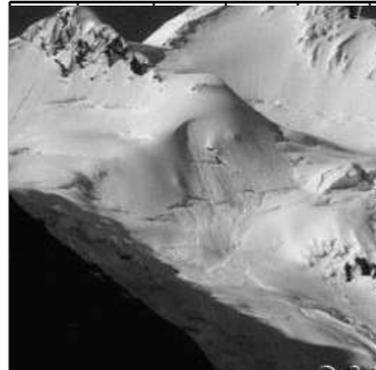
2



3



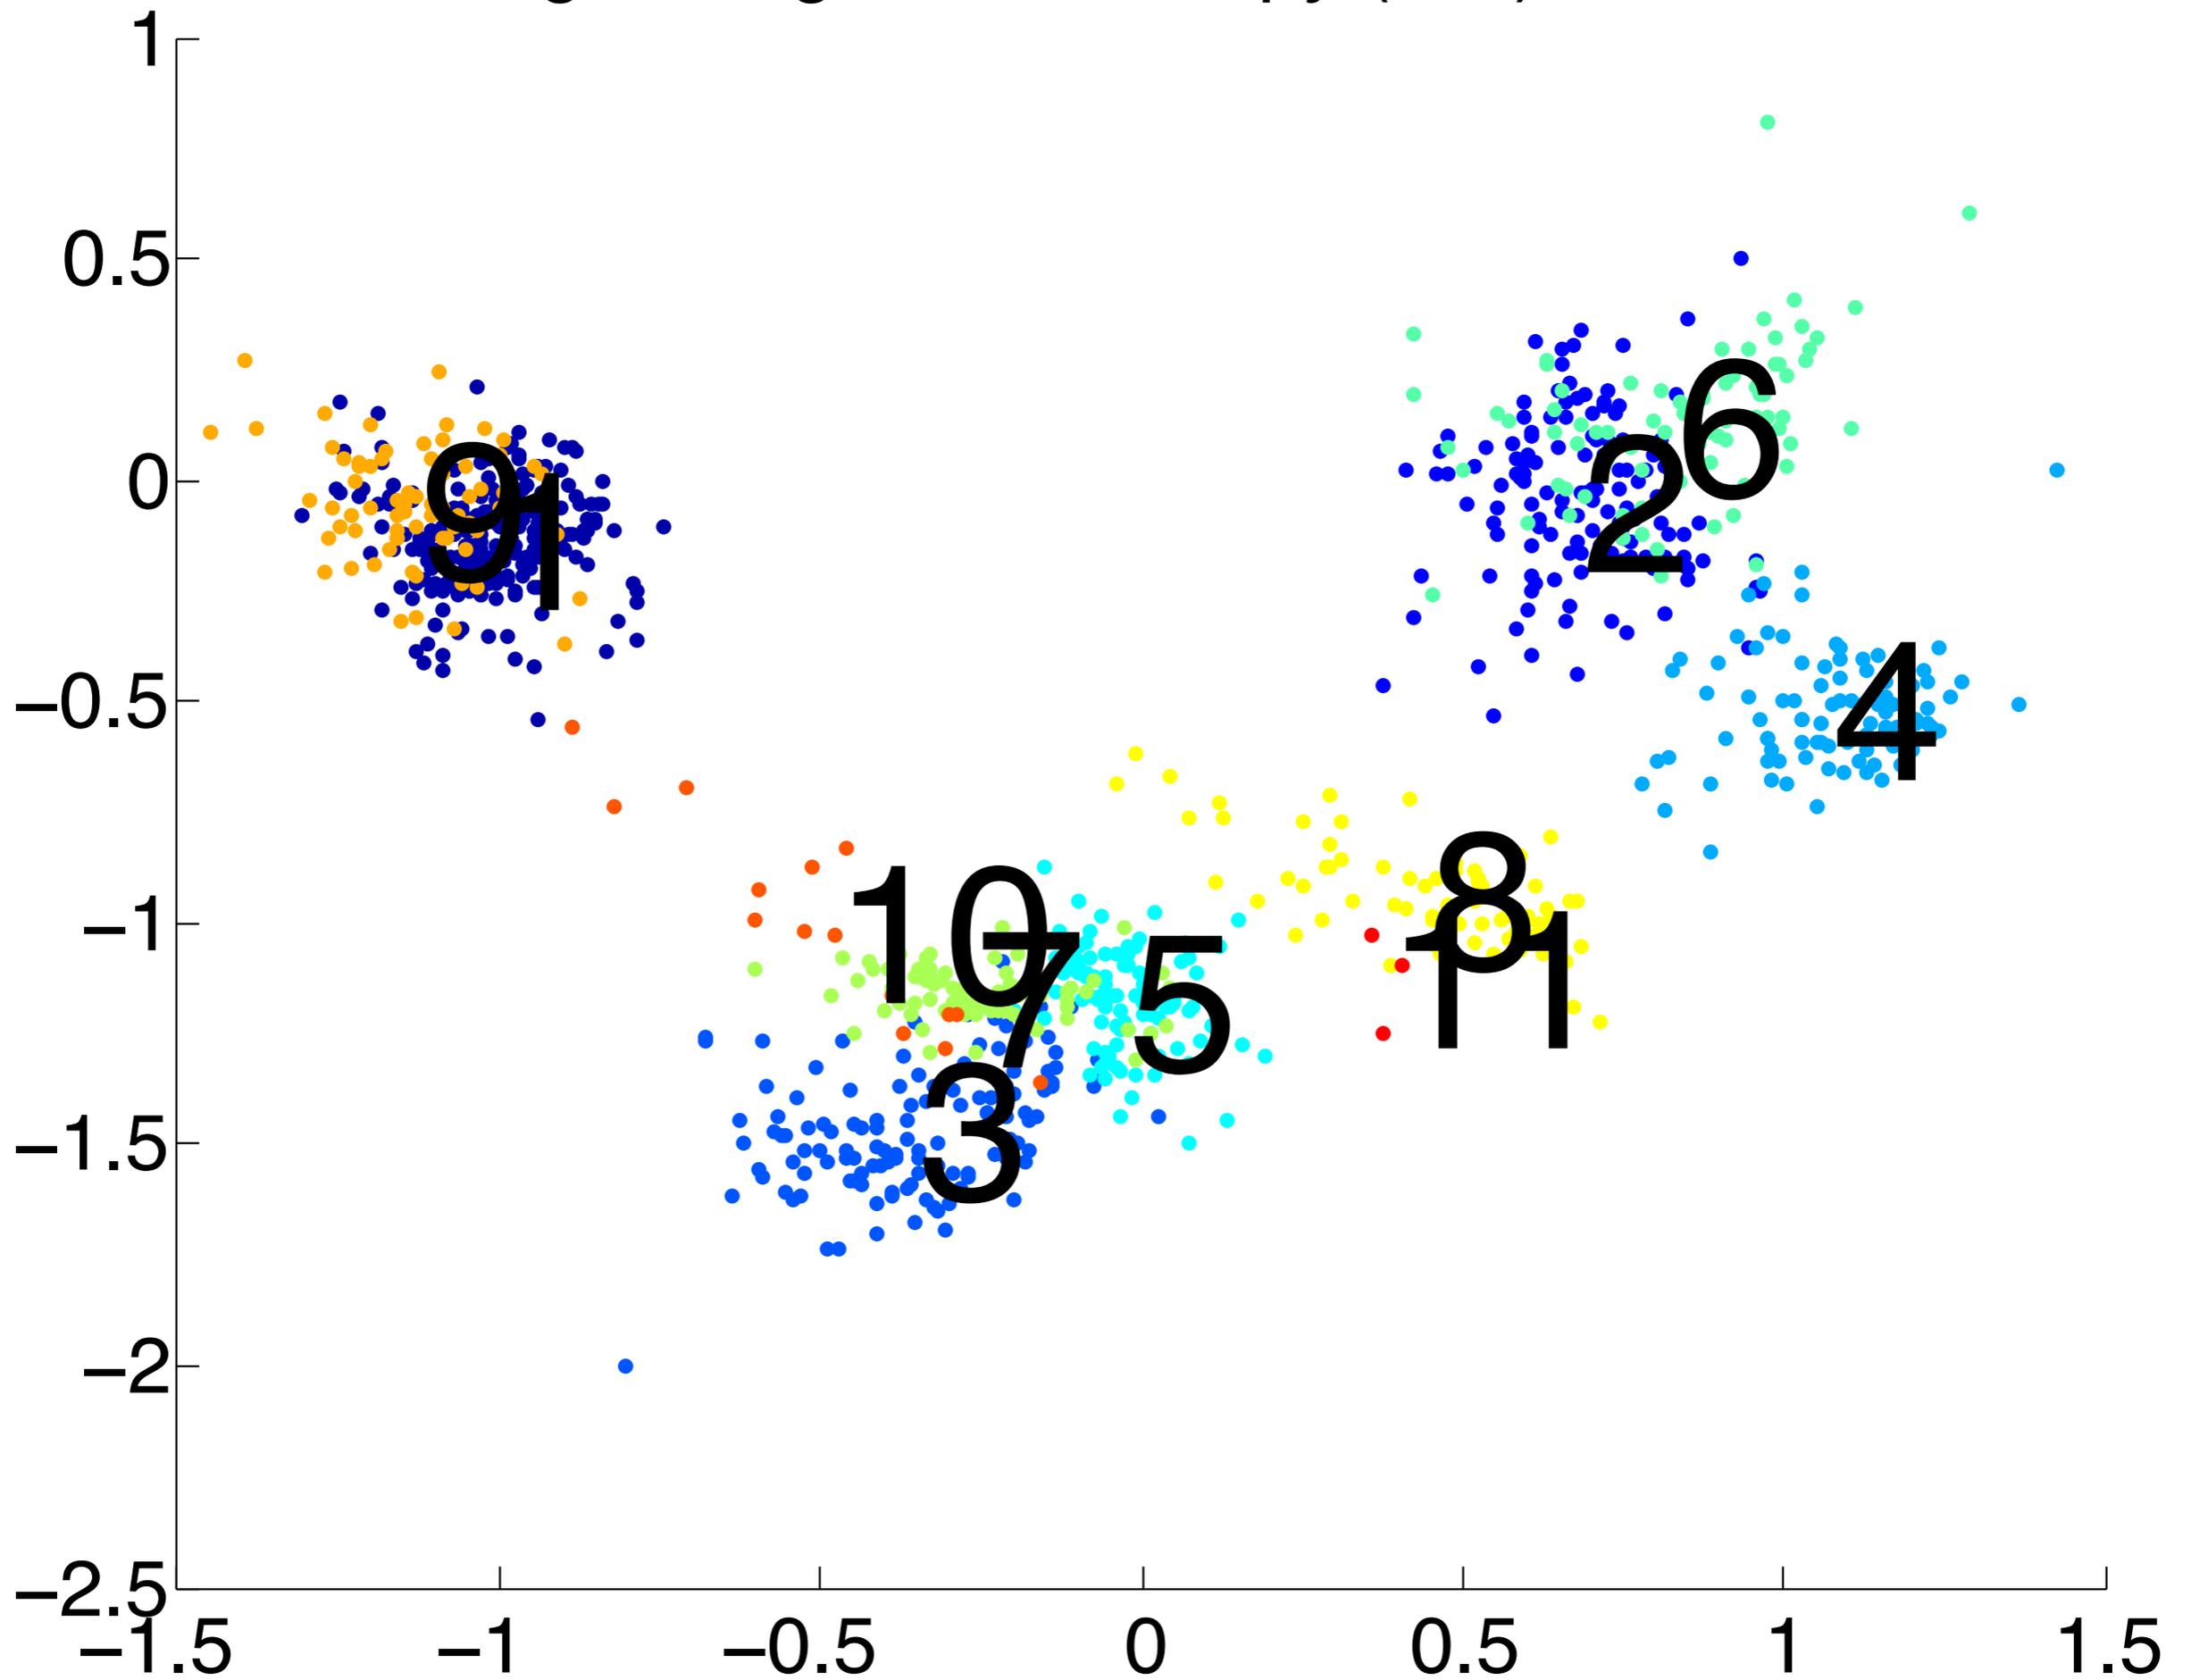
4



5



Average assignment entropy (bits): 0.0029653



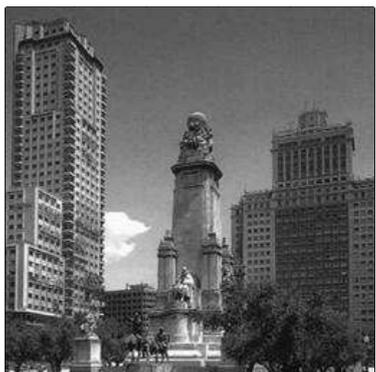
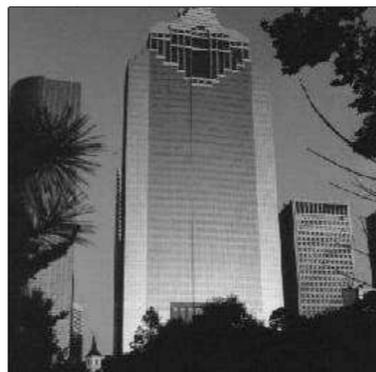
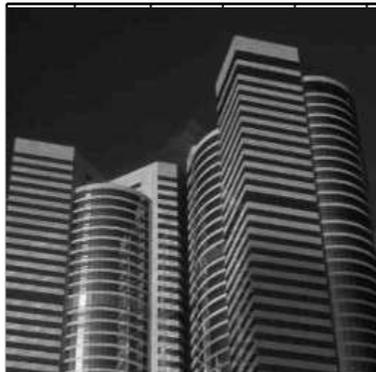
1



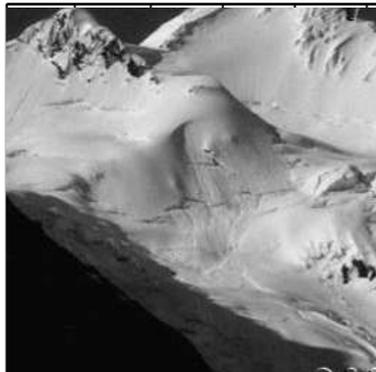
2



3

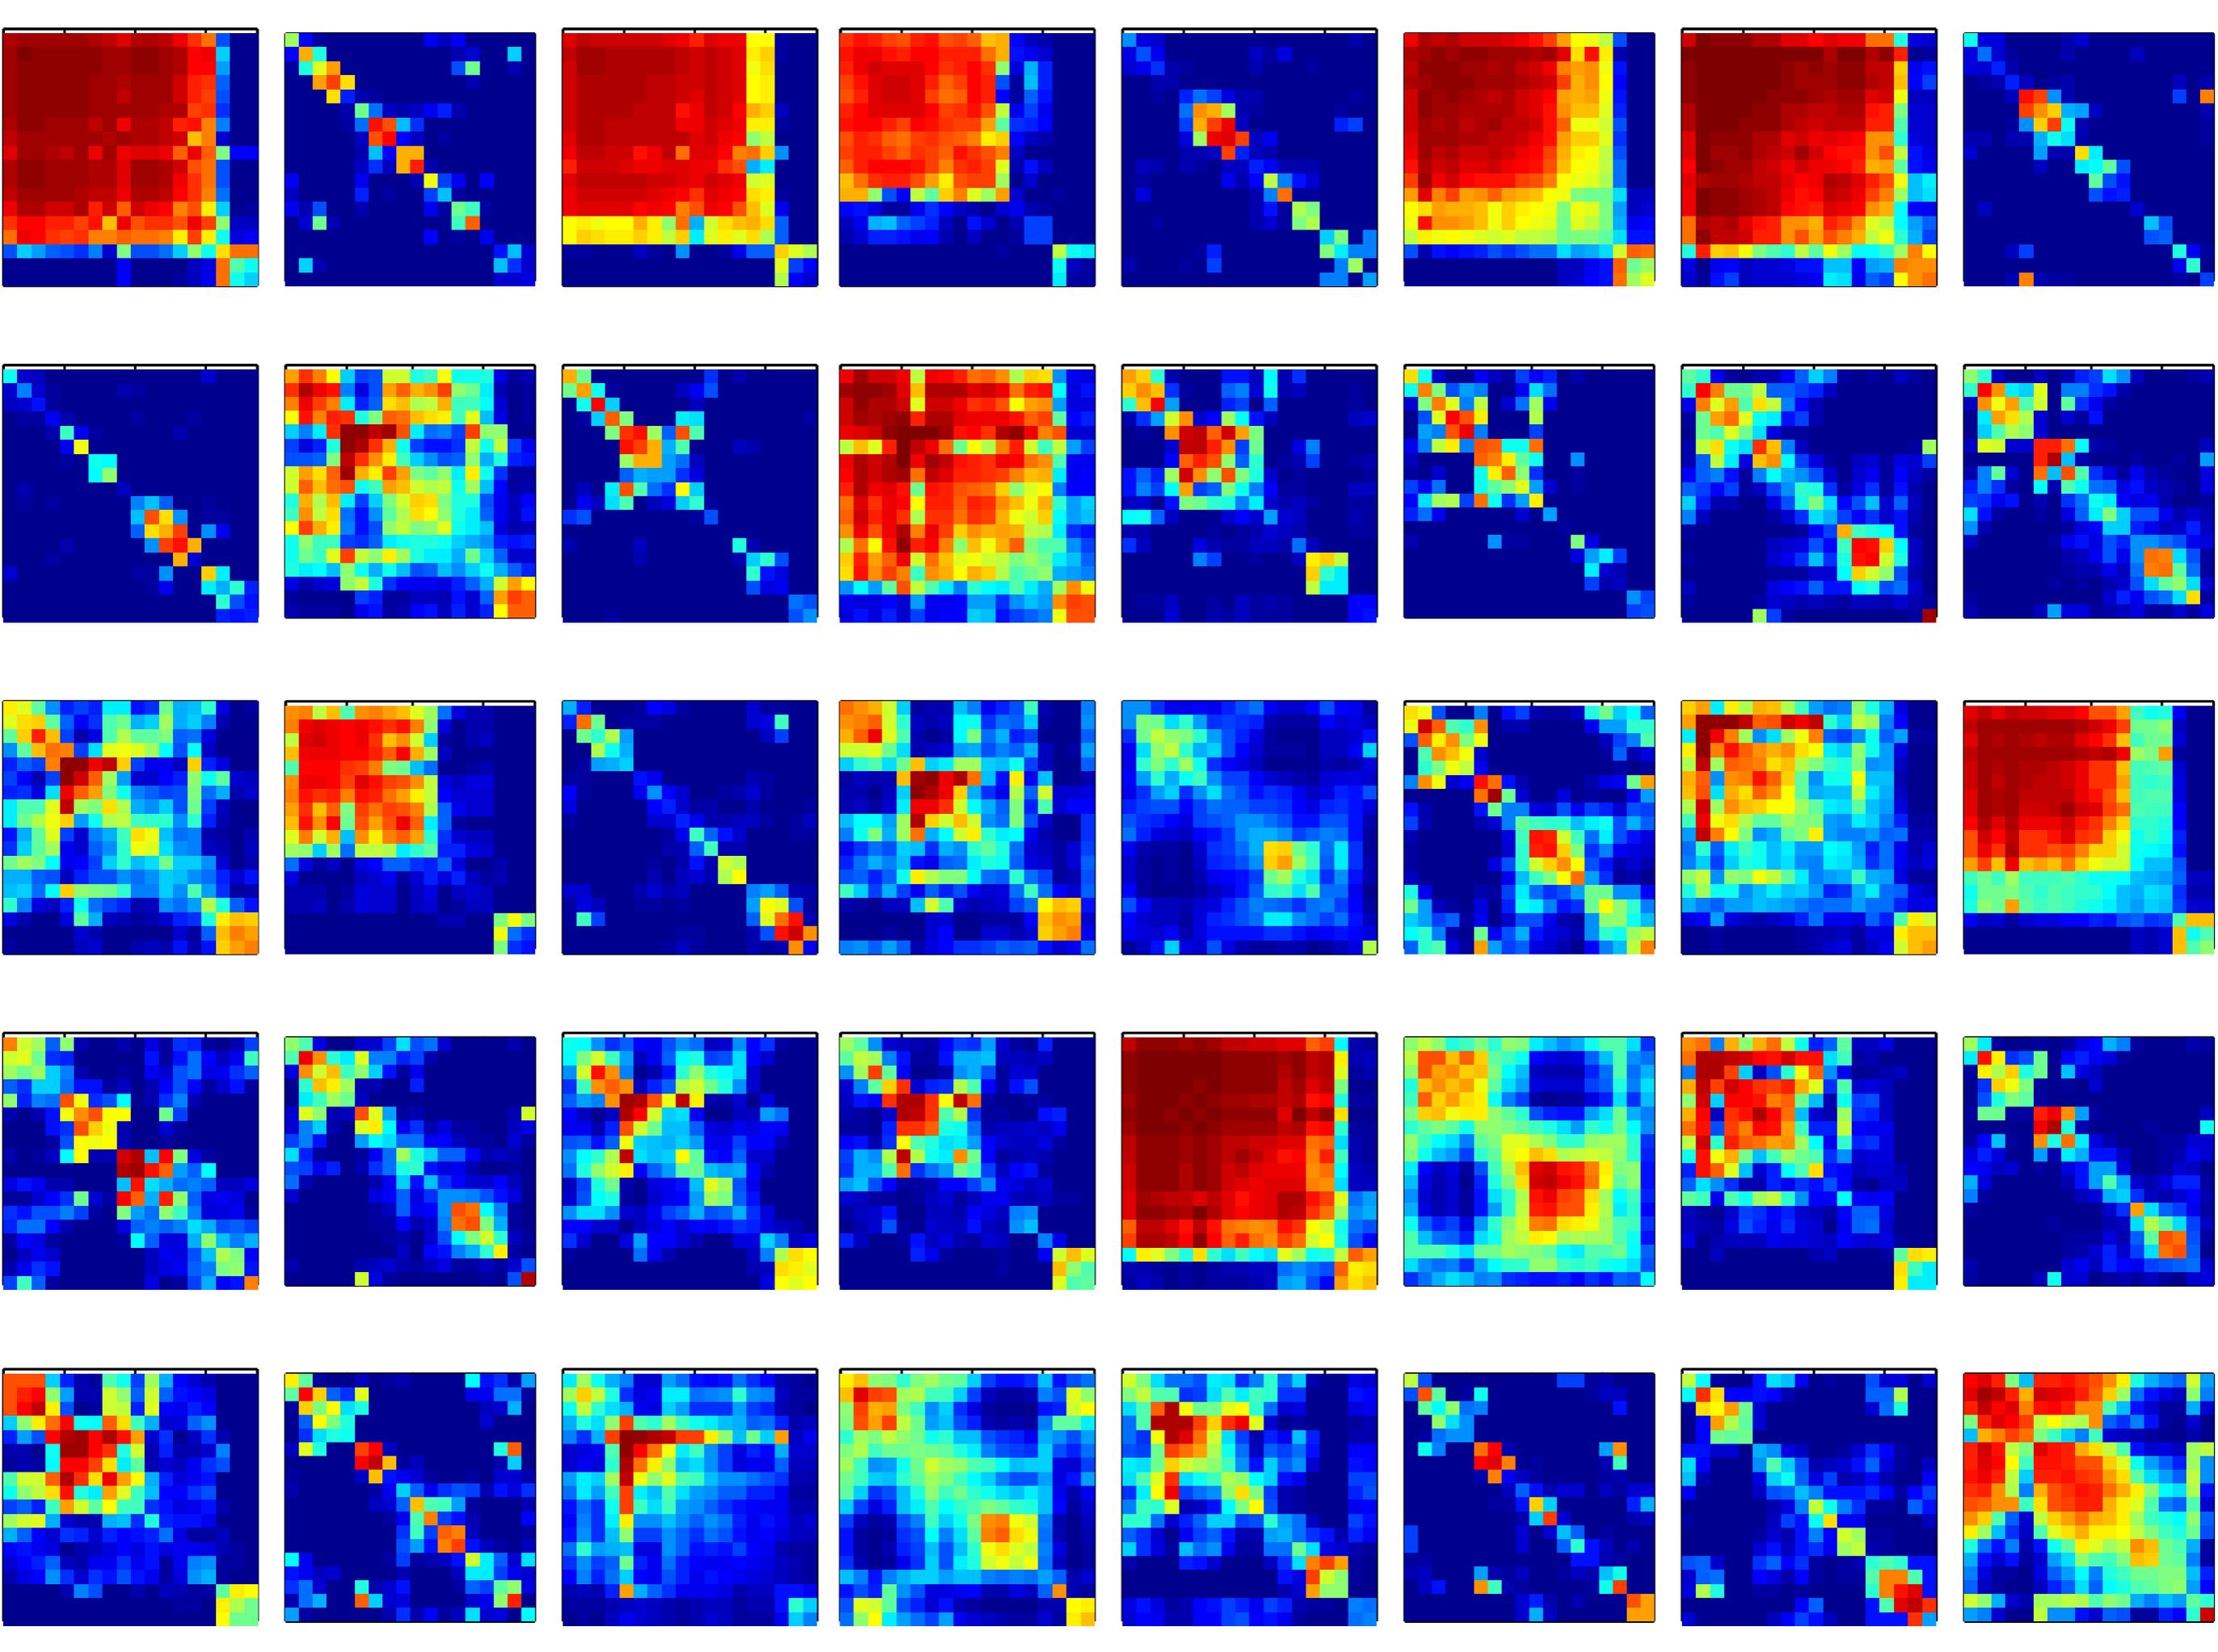


4



5













1.1



1.2



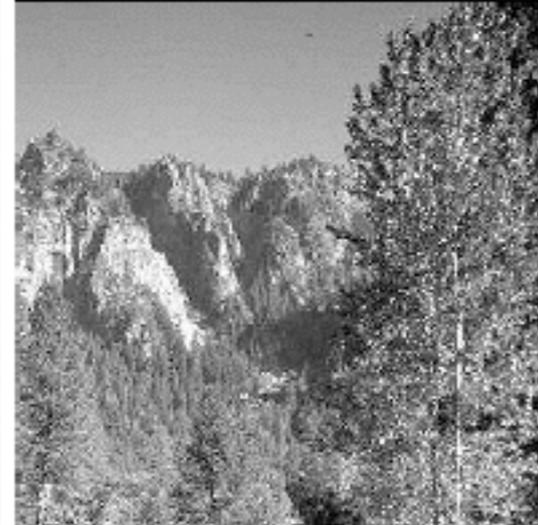
1.3



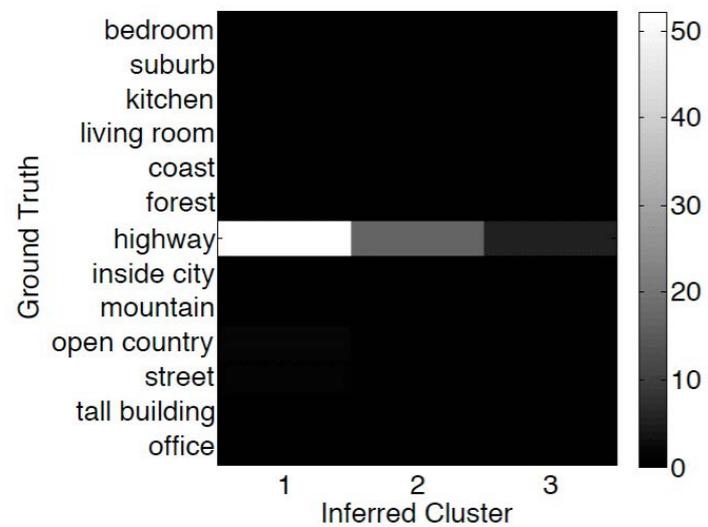




4.1







8.1

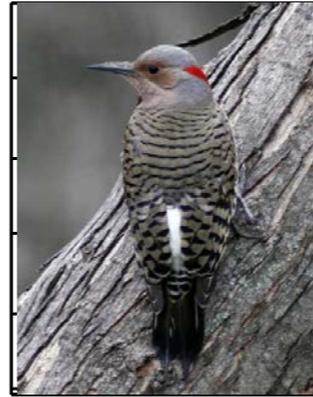
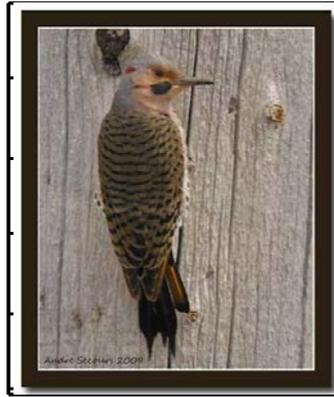


8.2

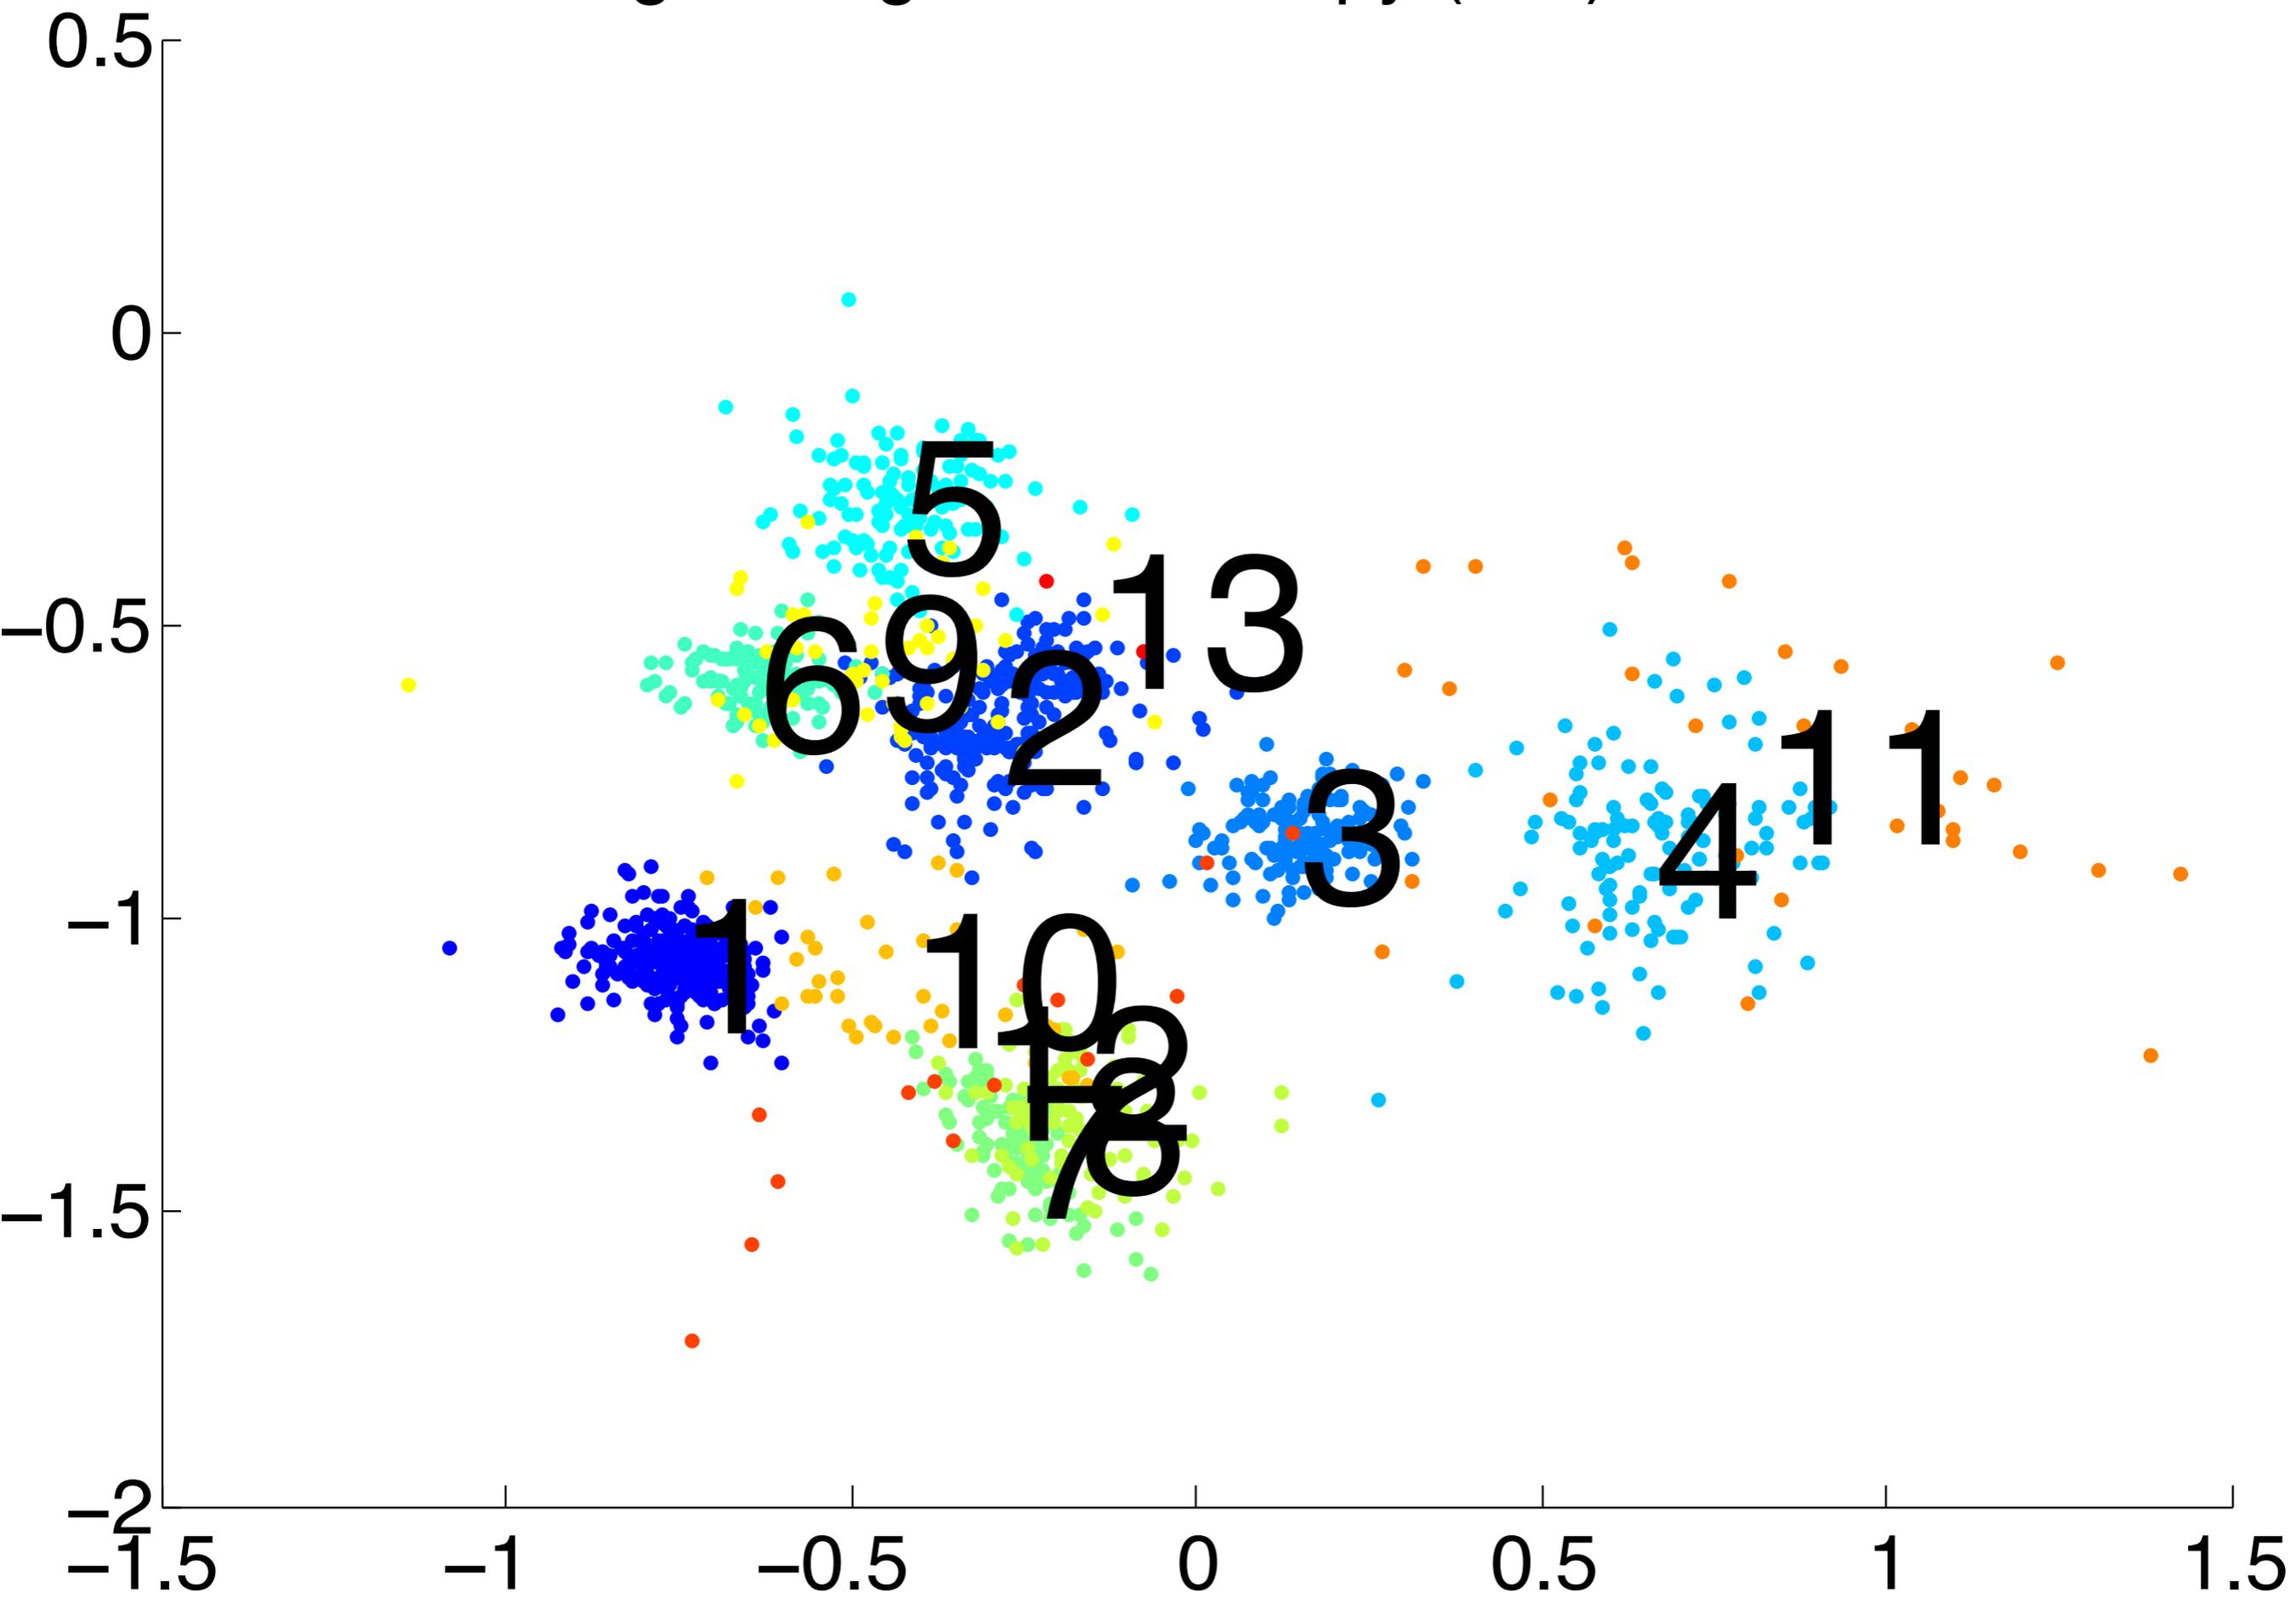


8.3

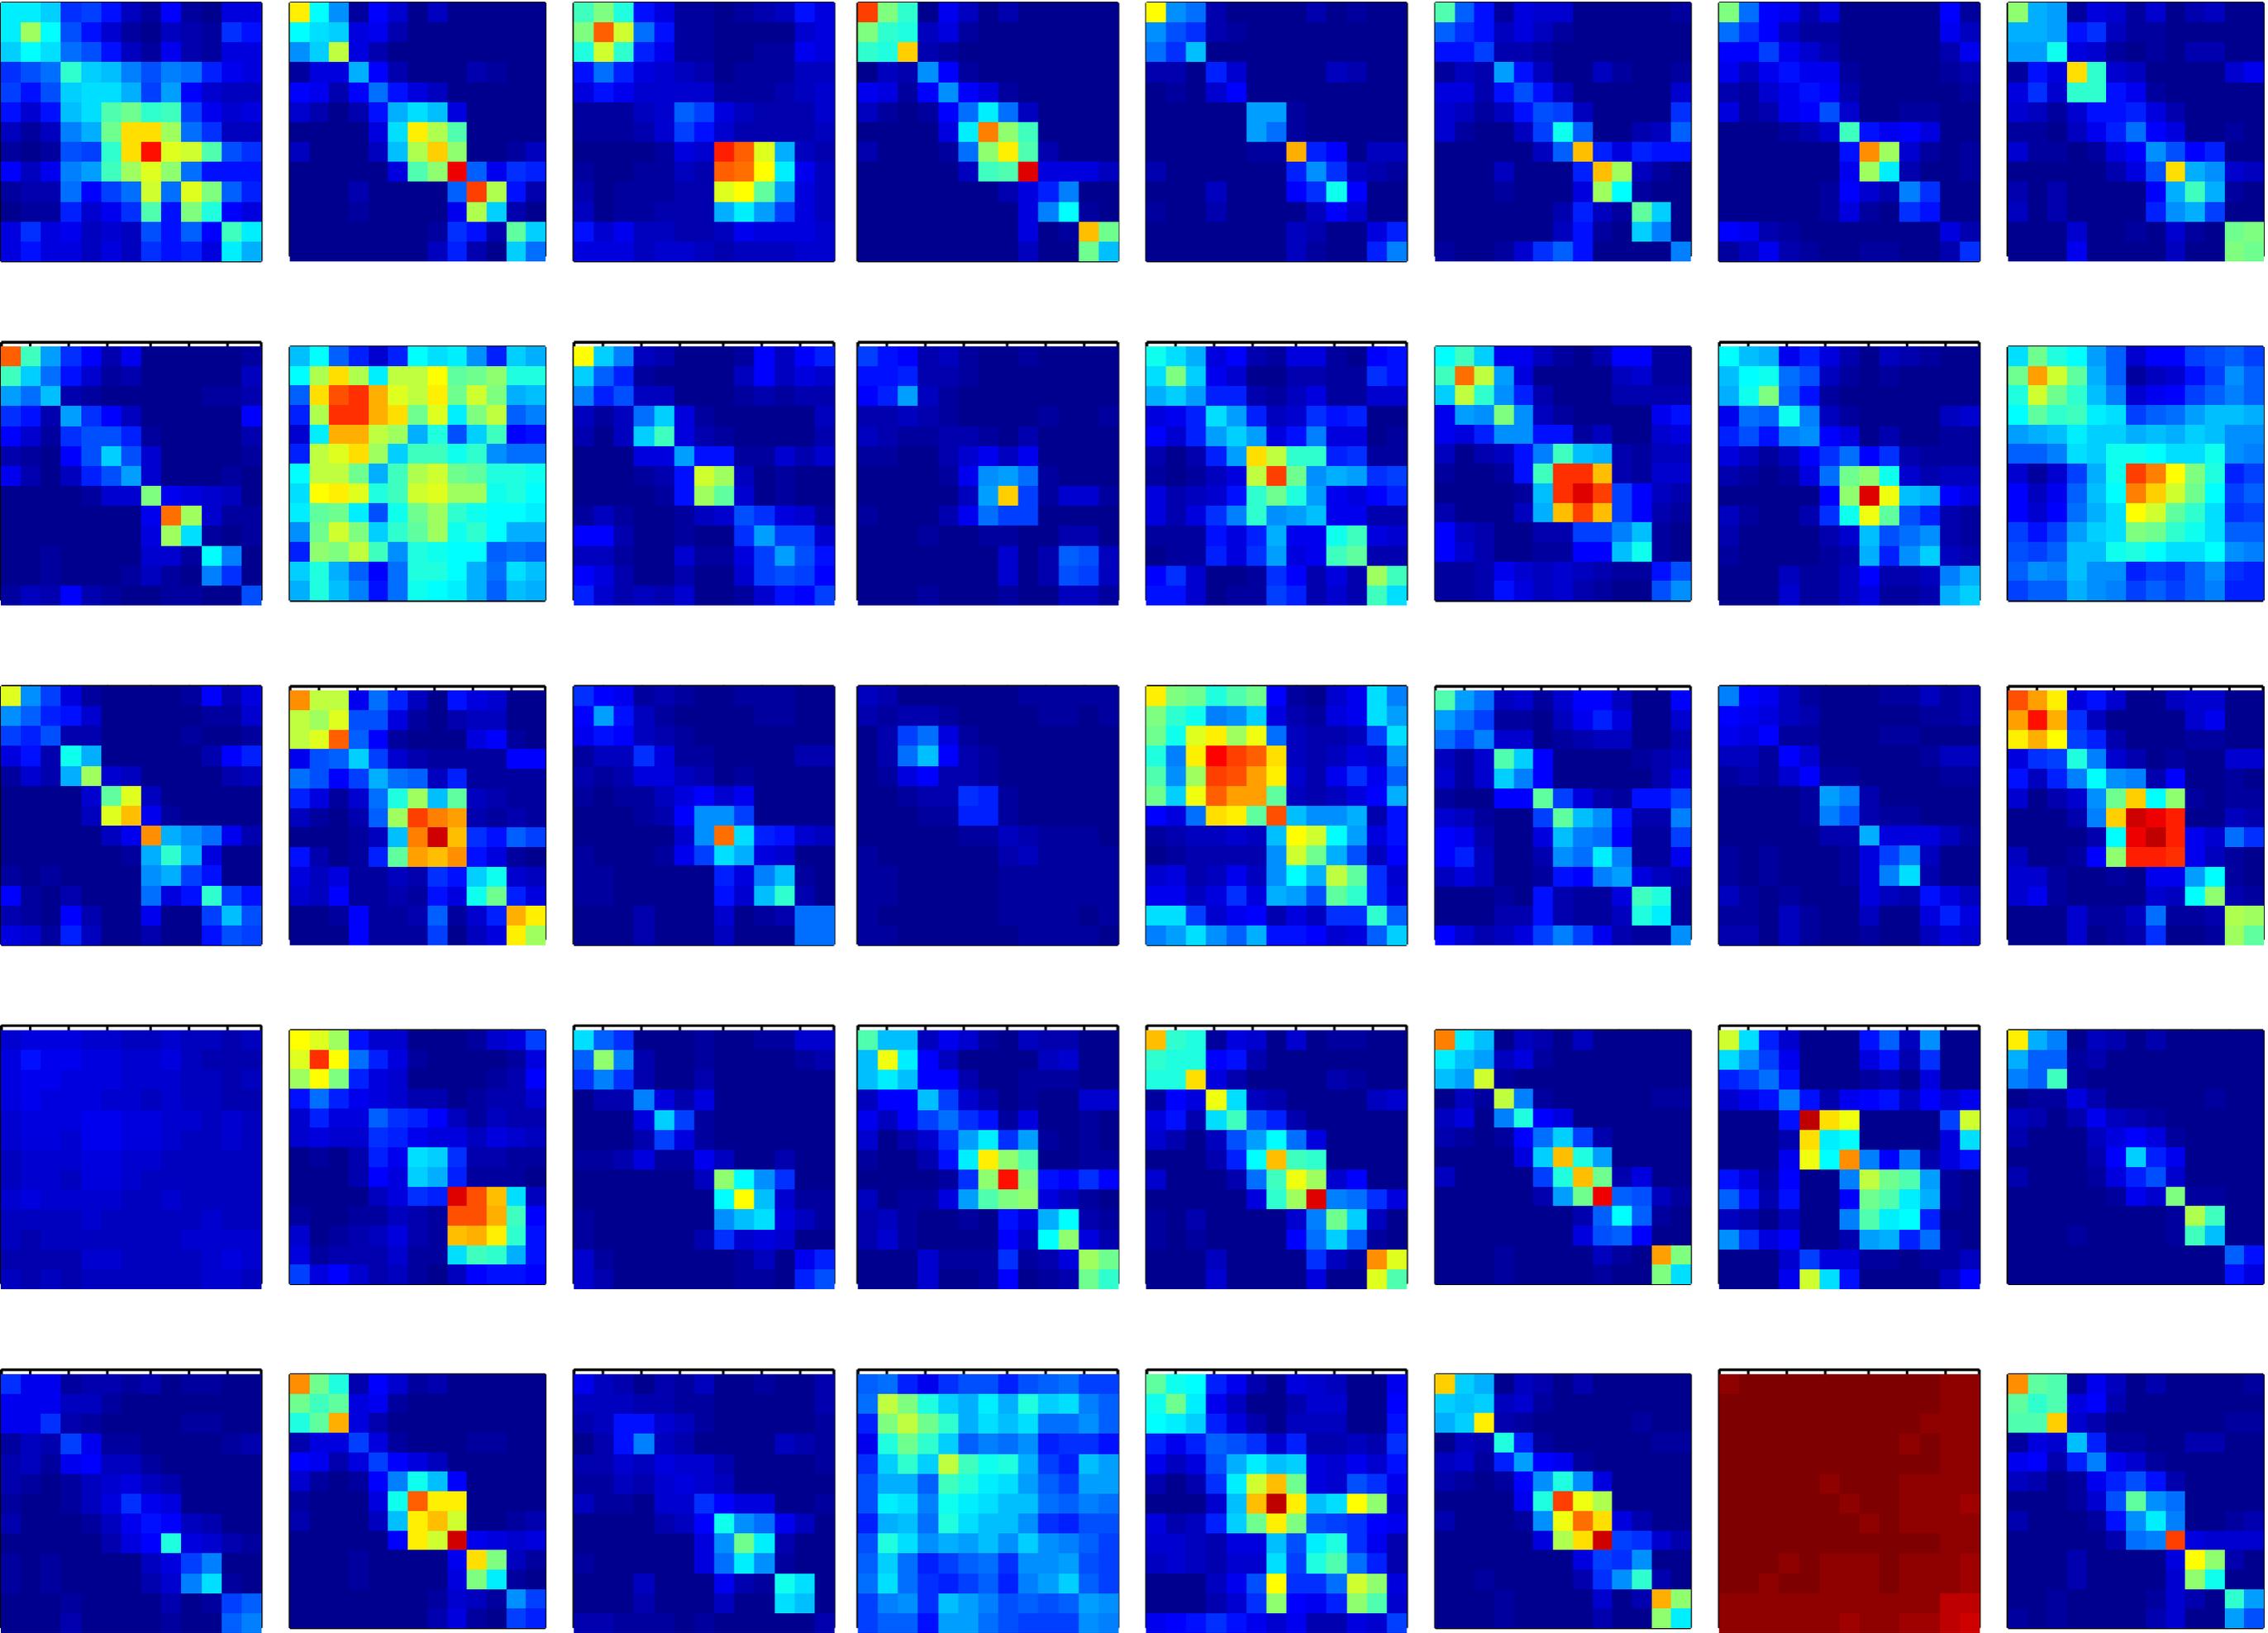




Average assignment entropy (bits): 0.004792

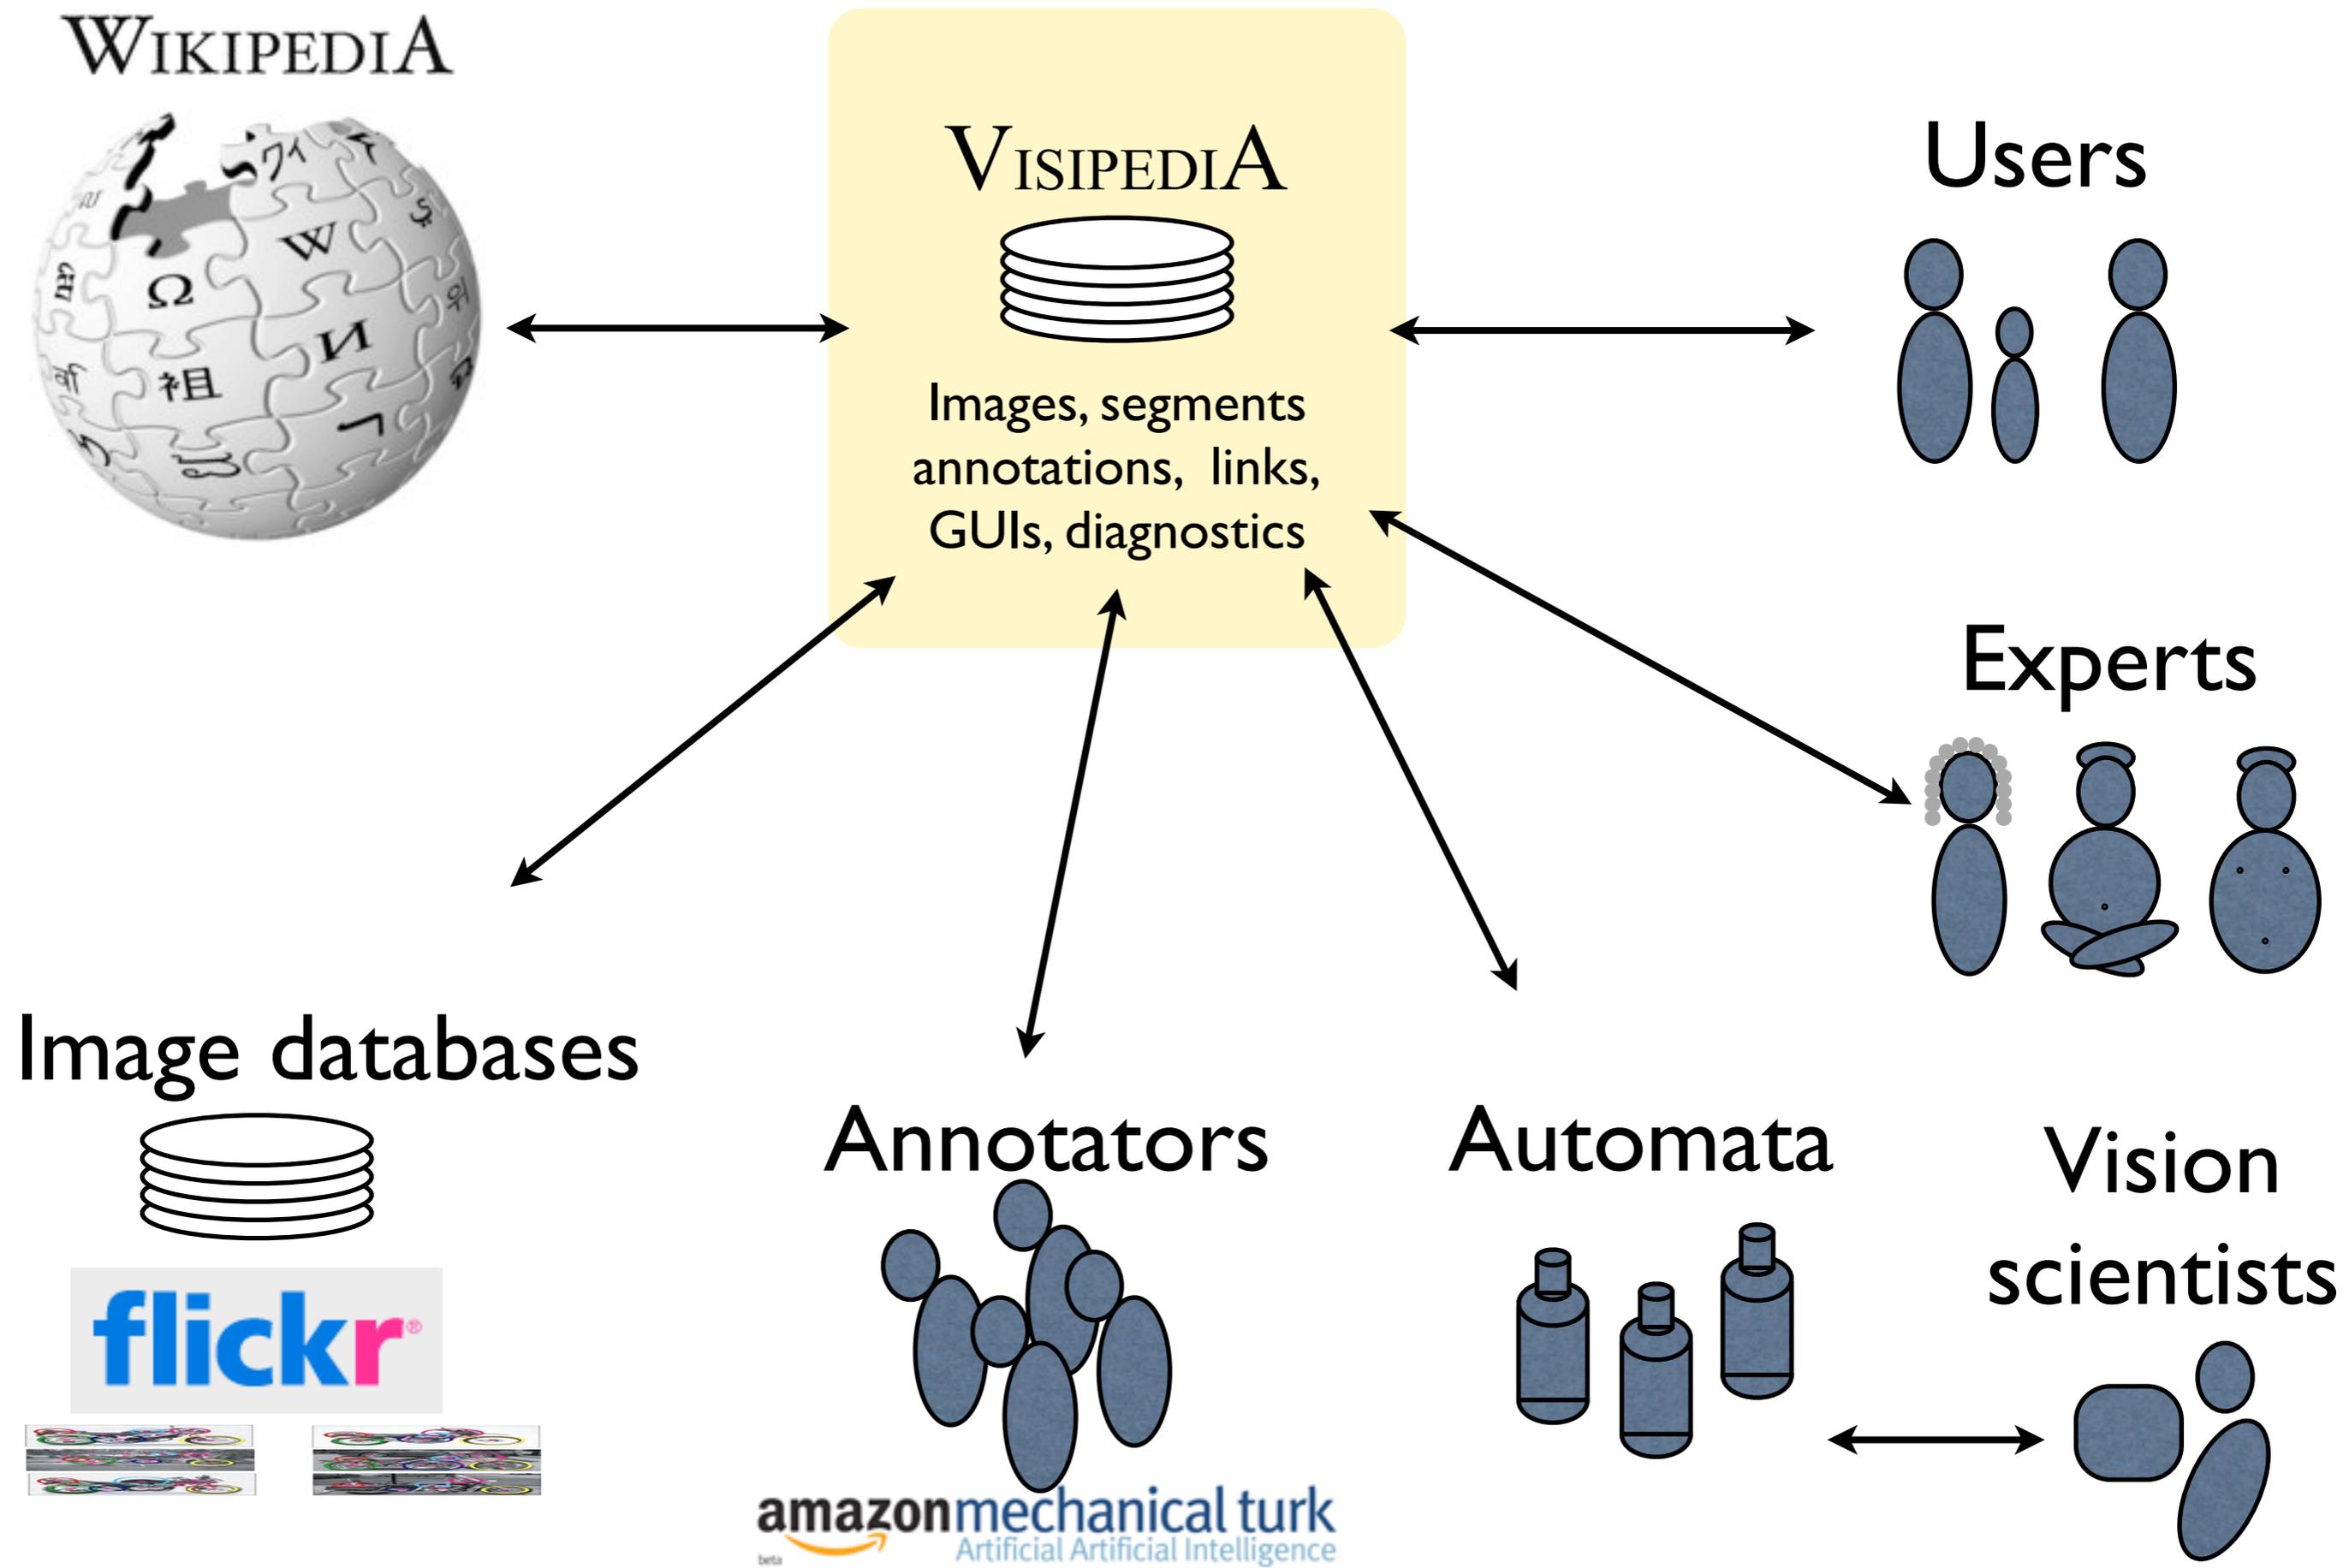


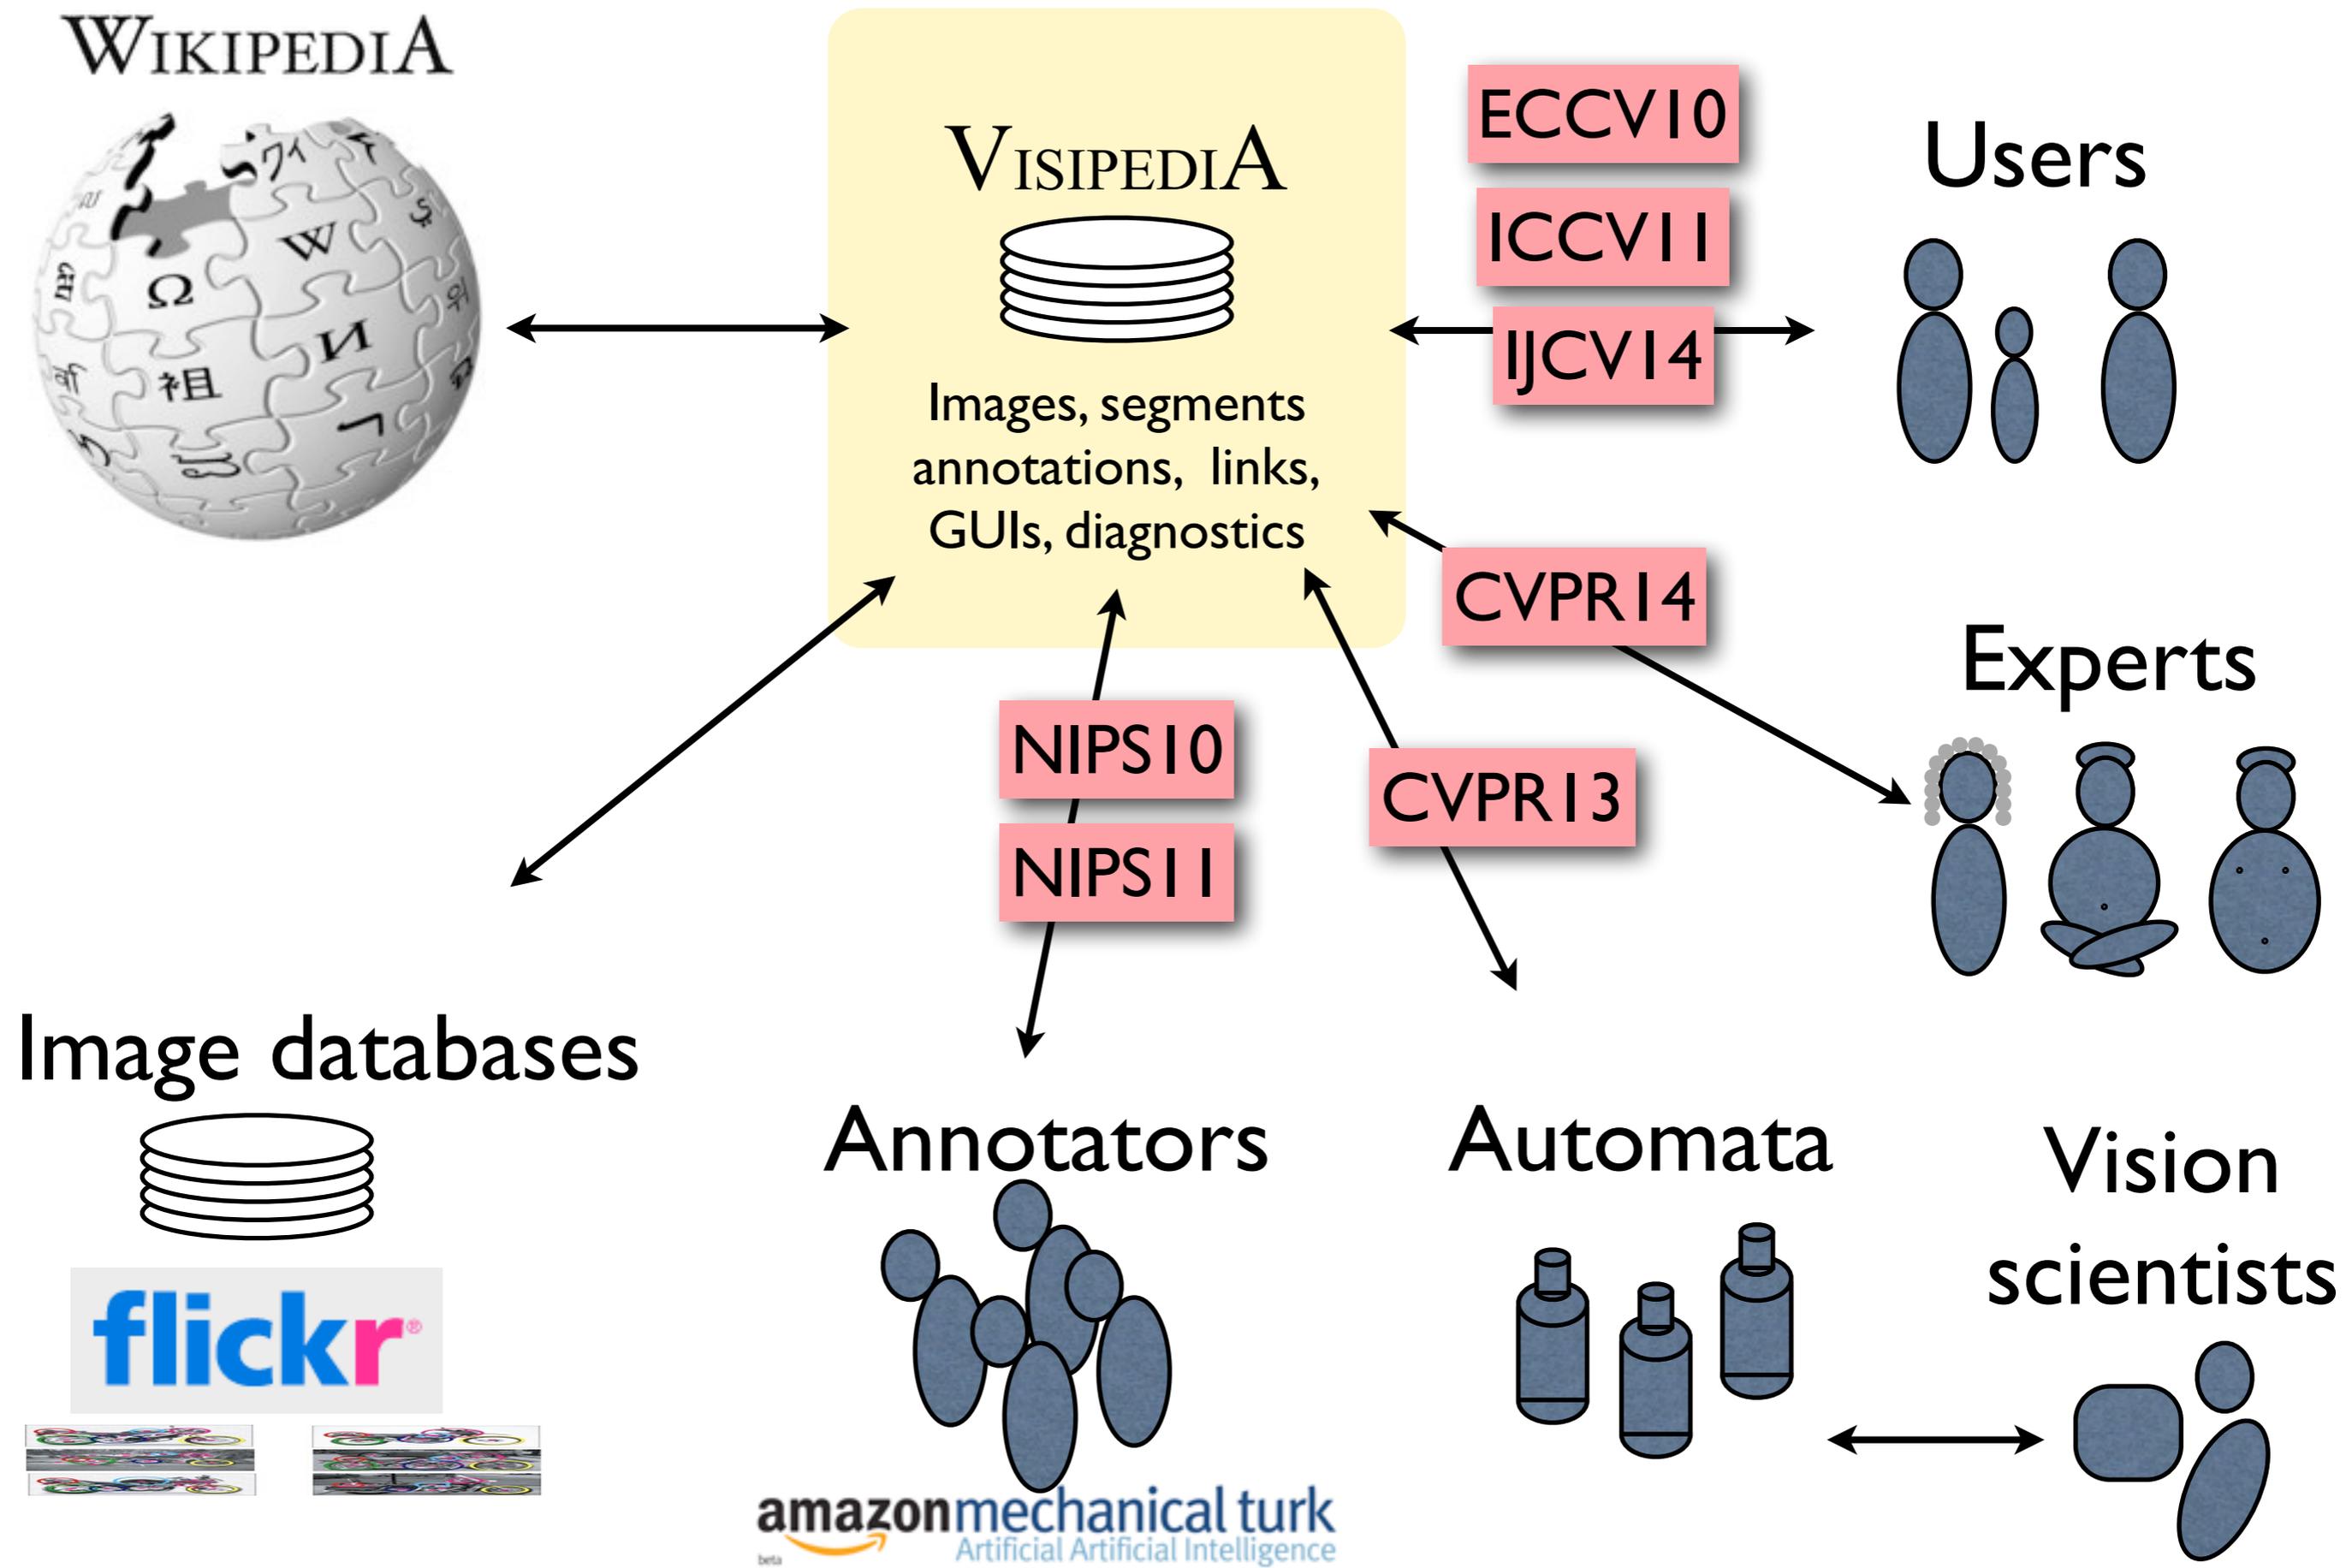


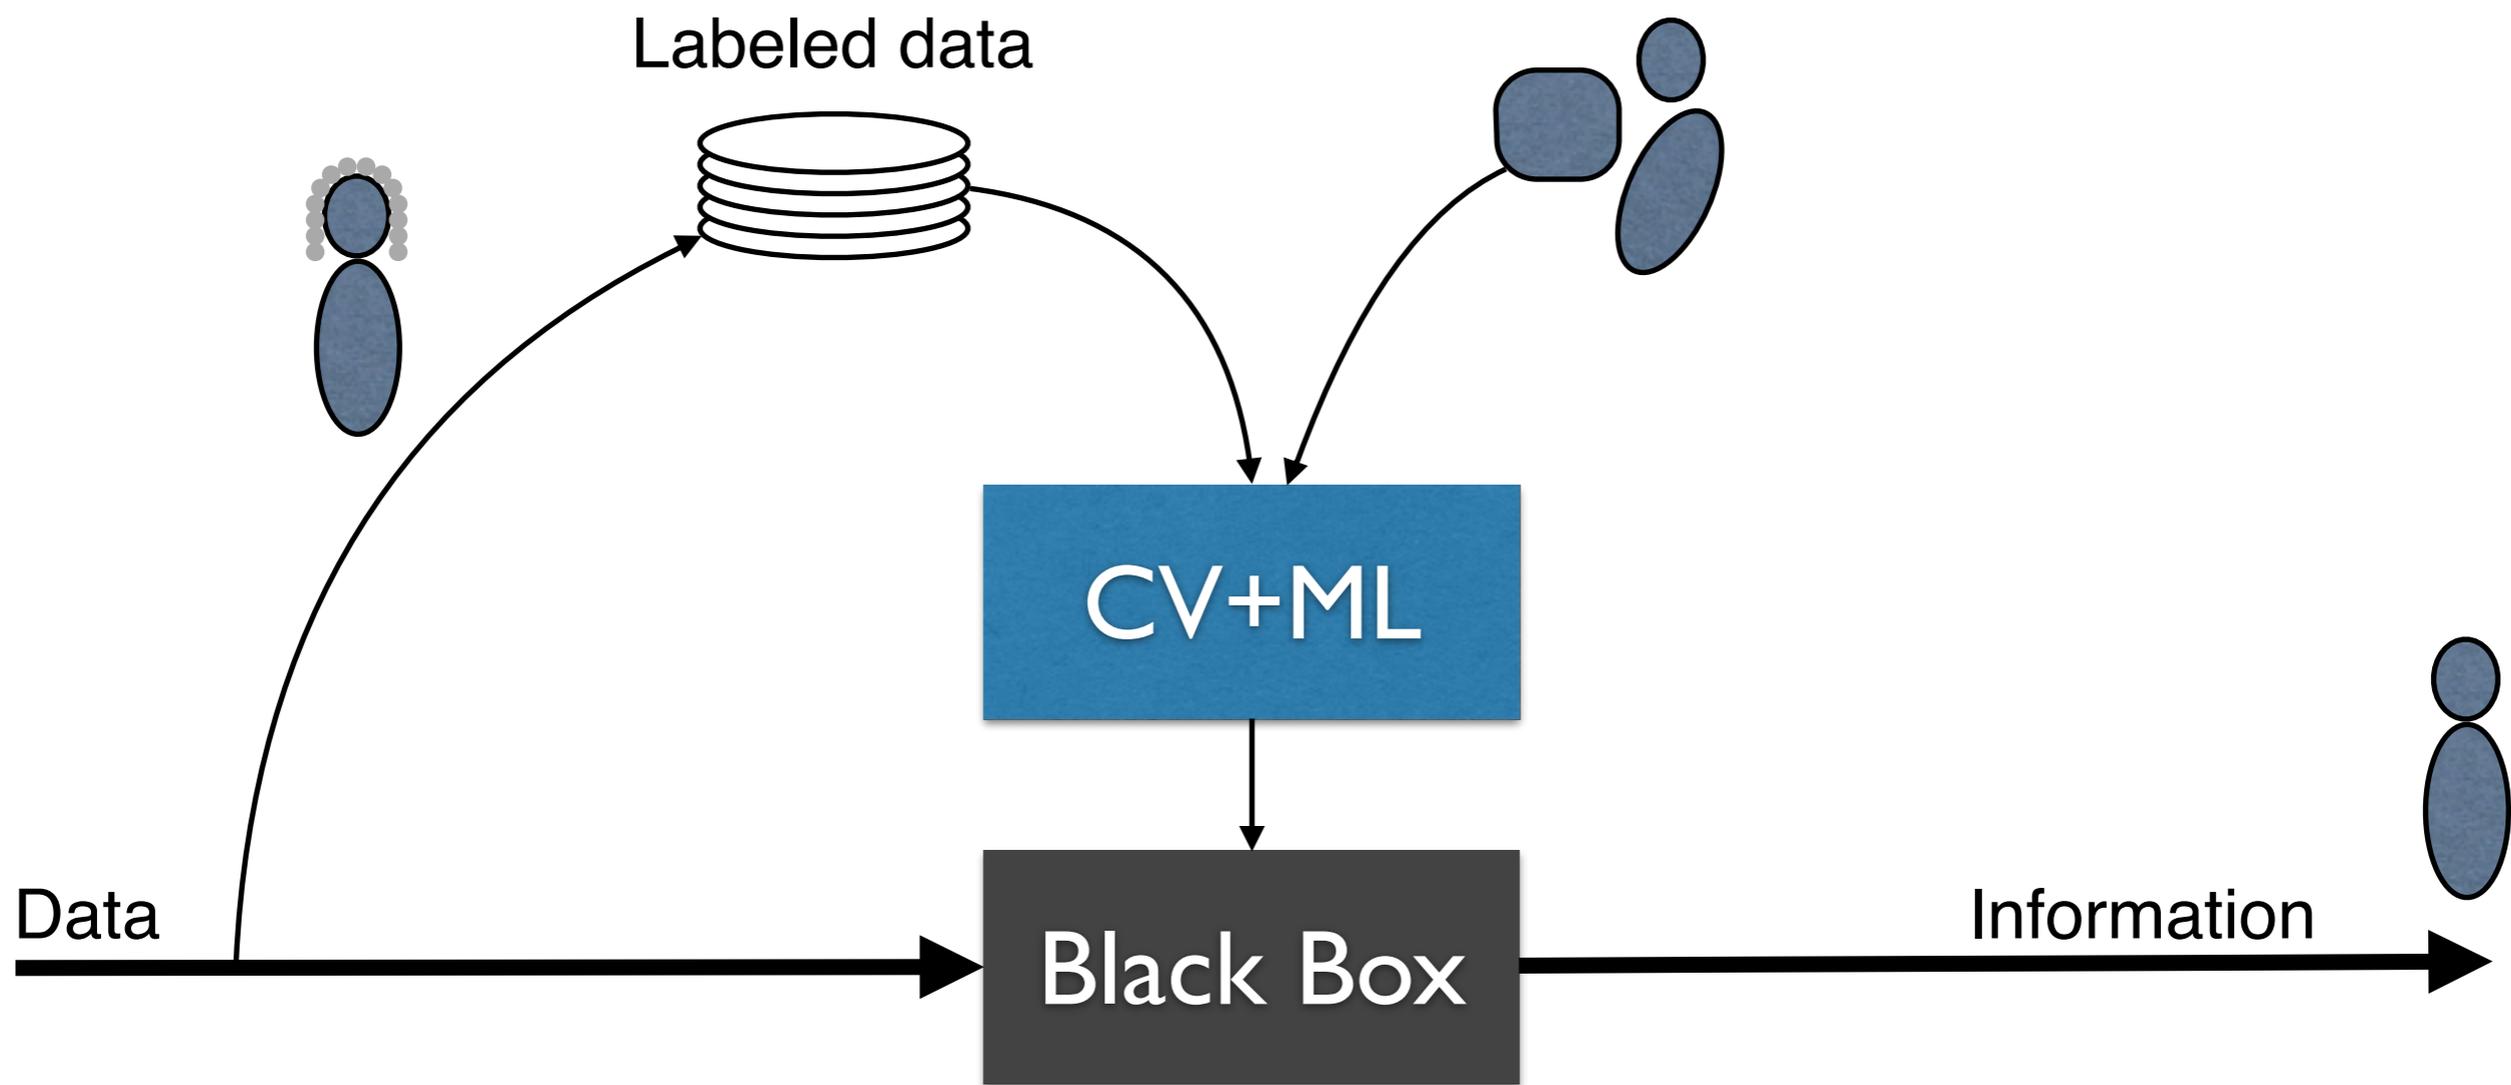


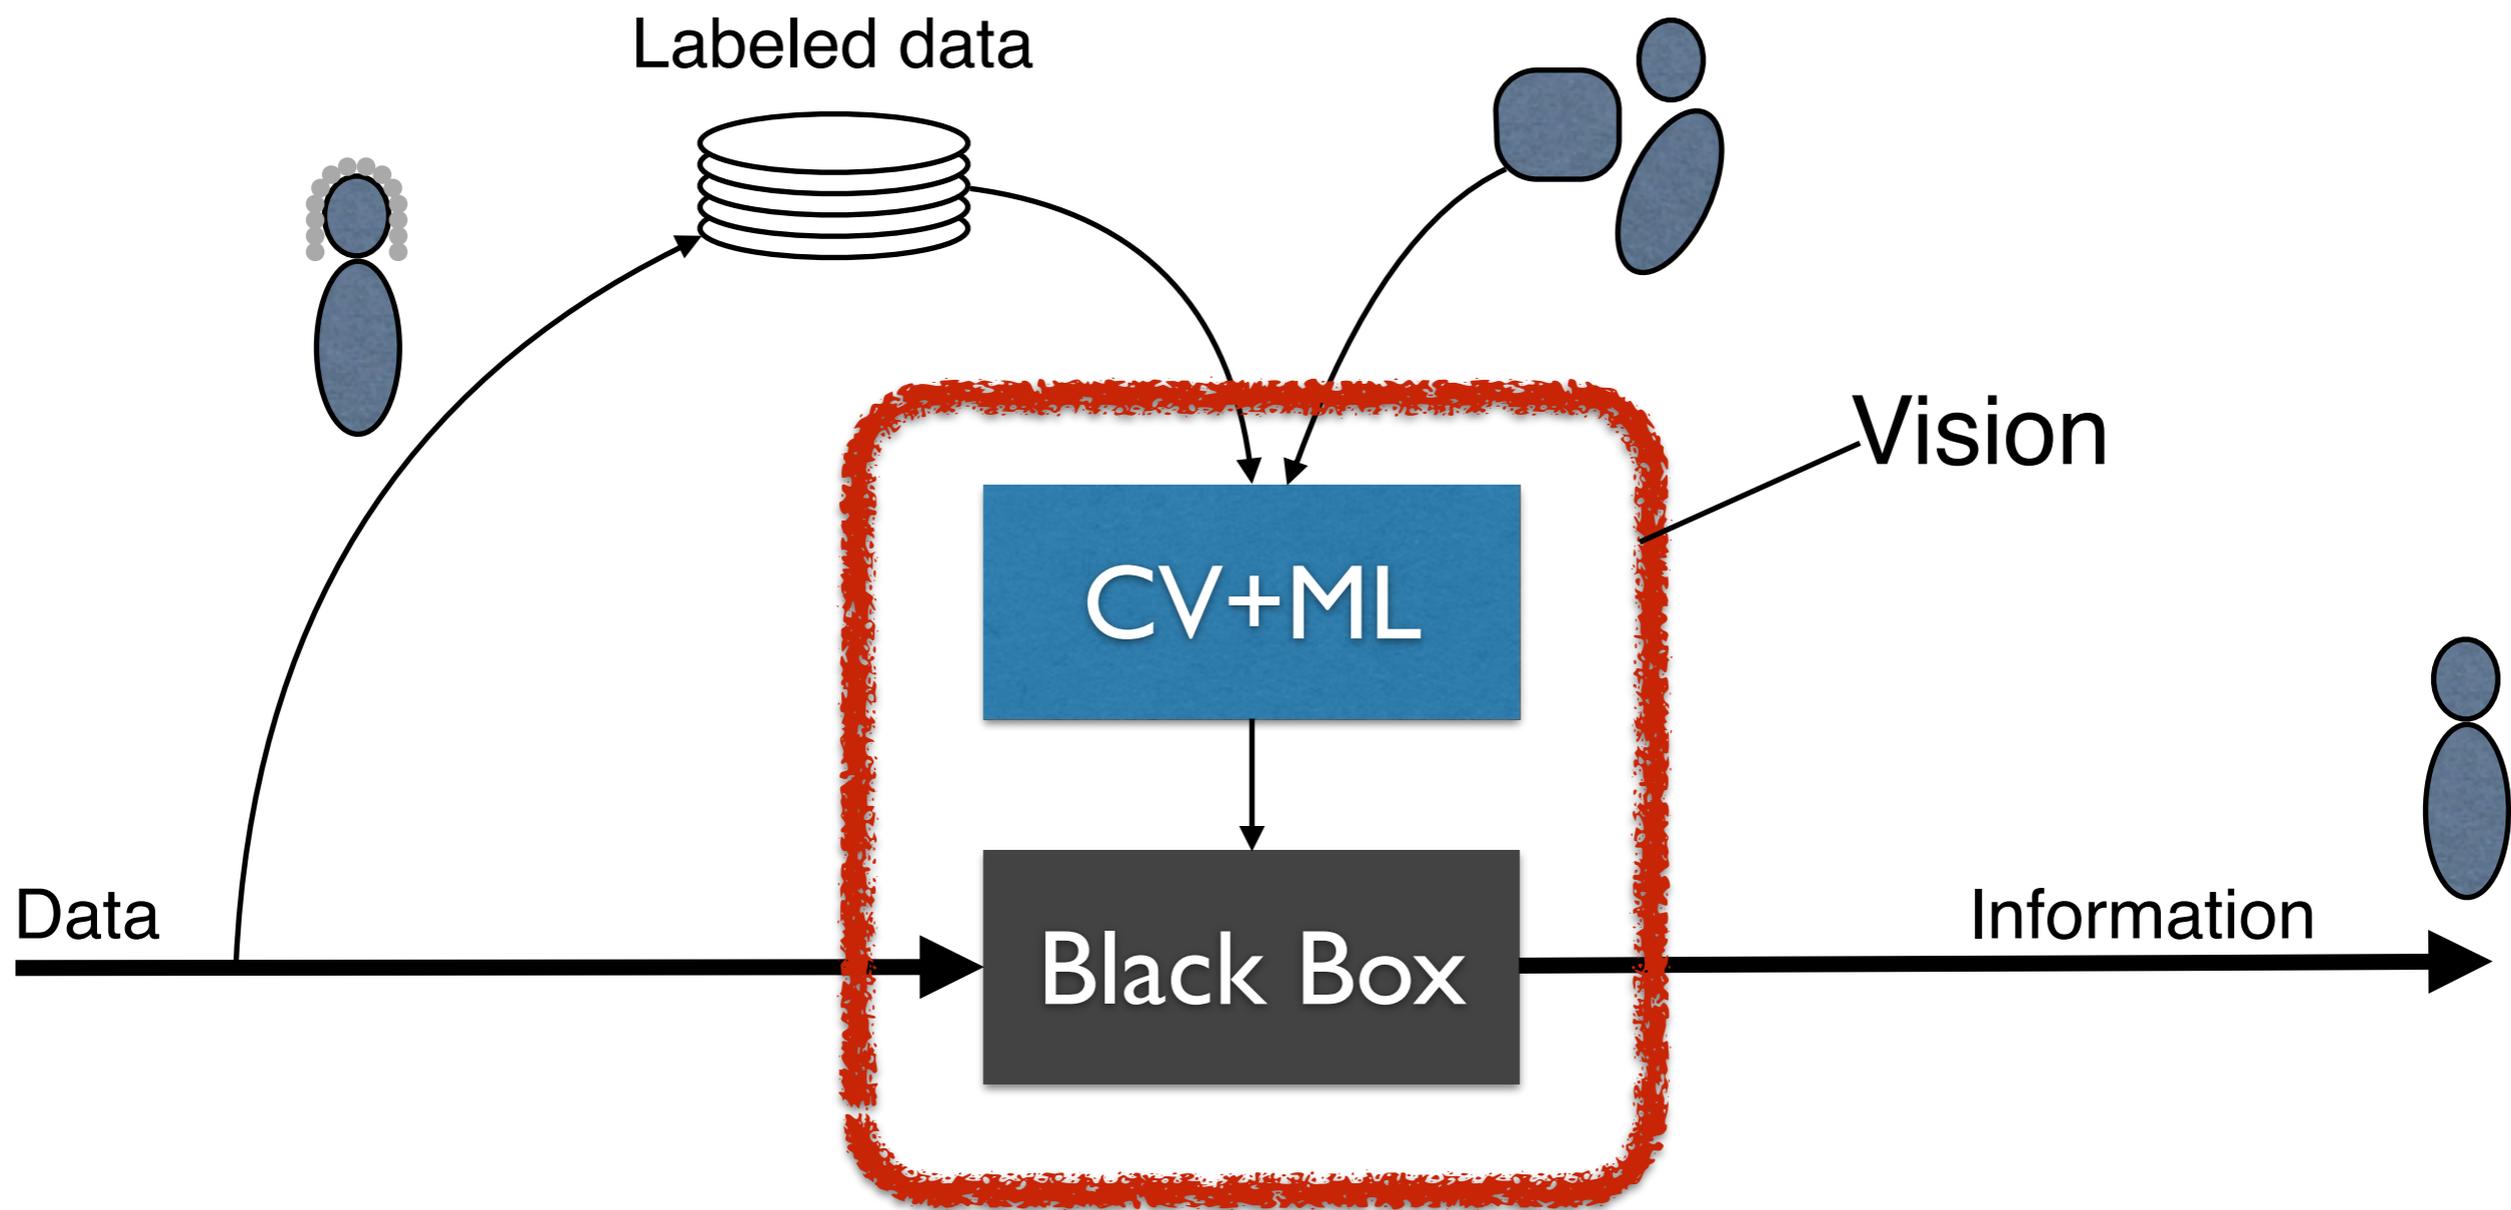
# Crowdclustering

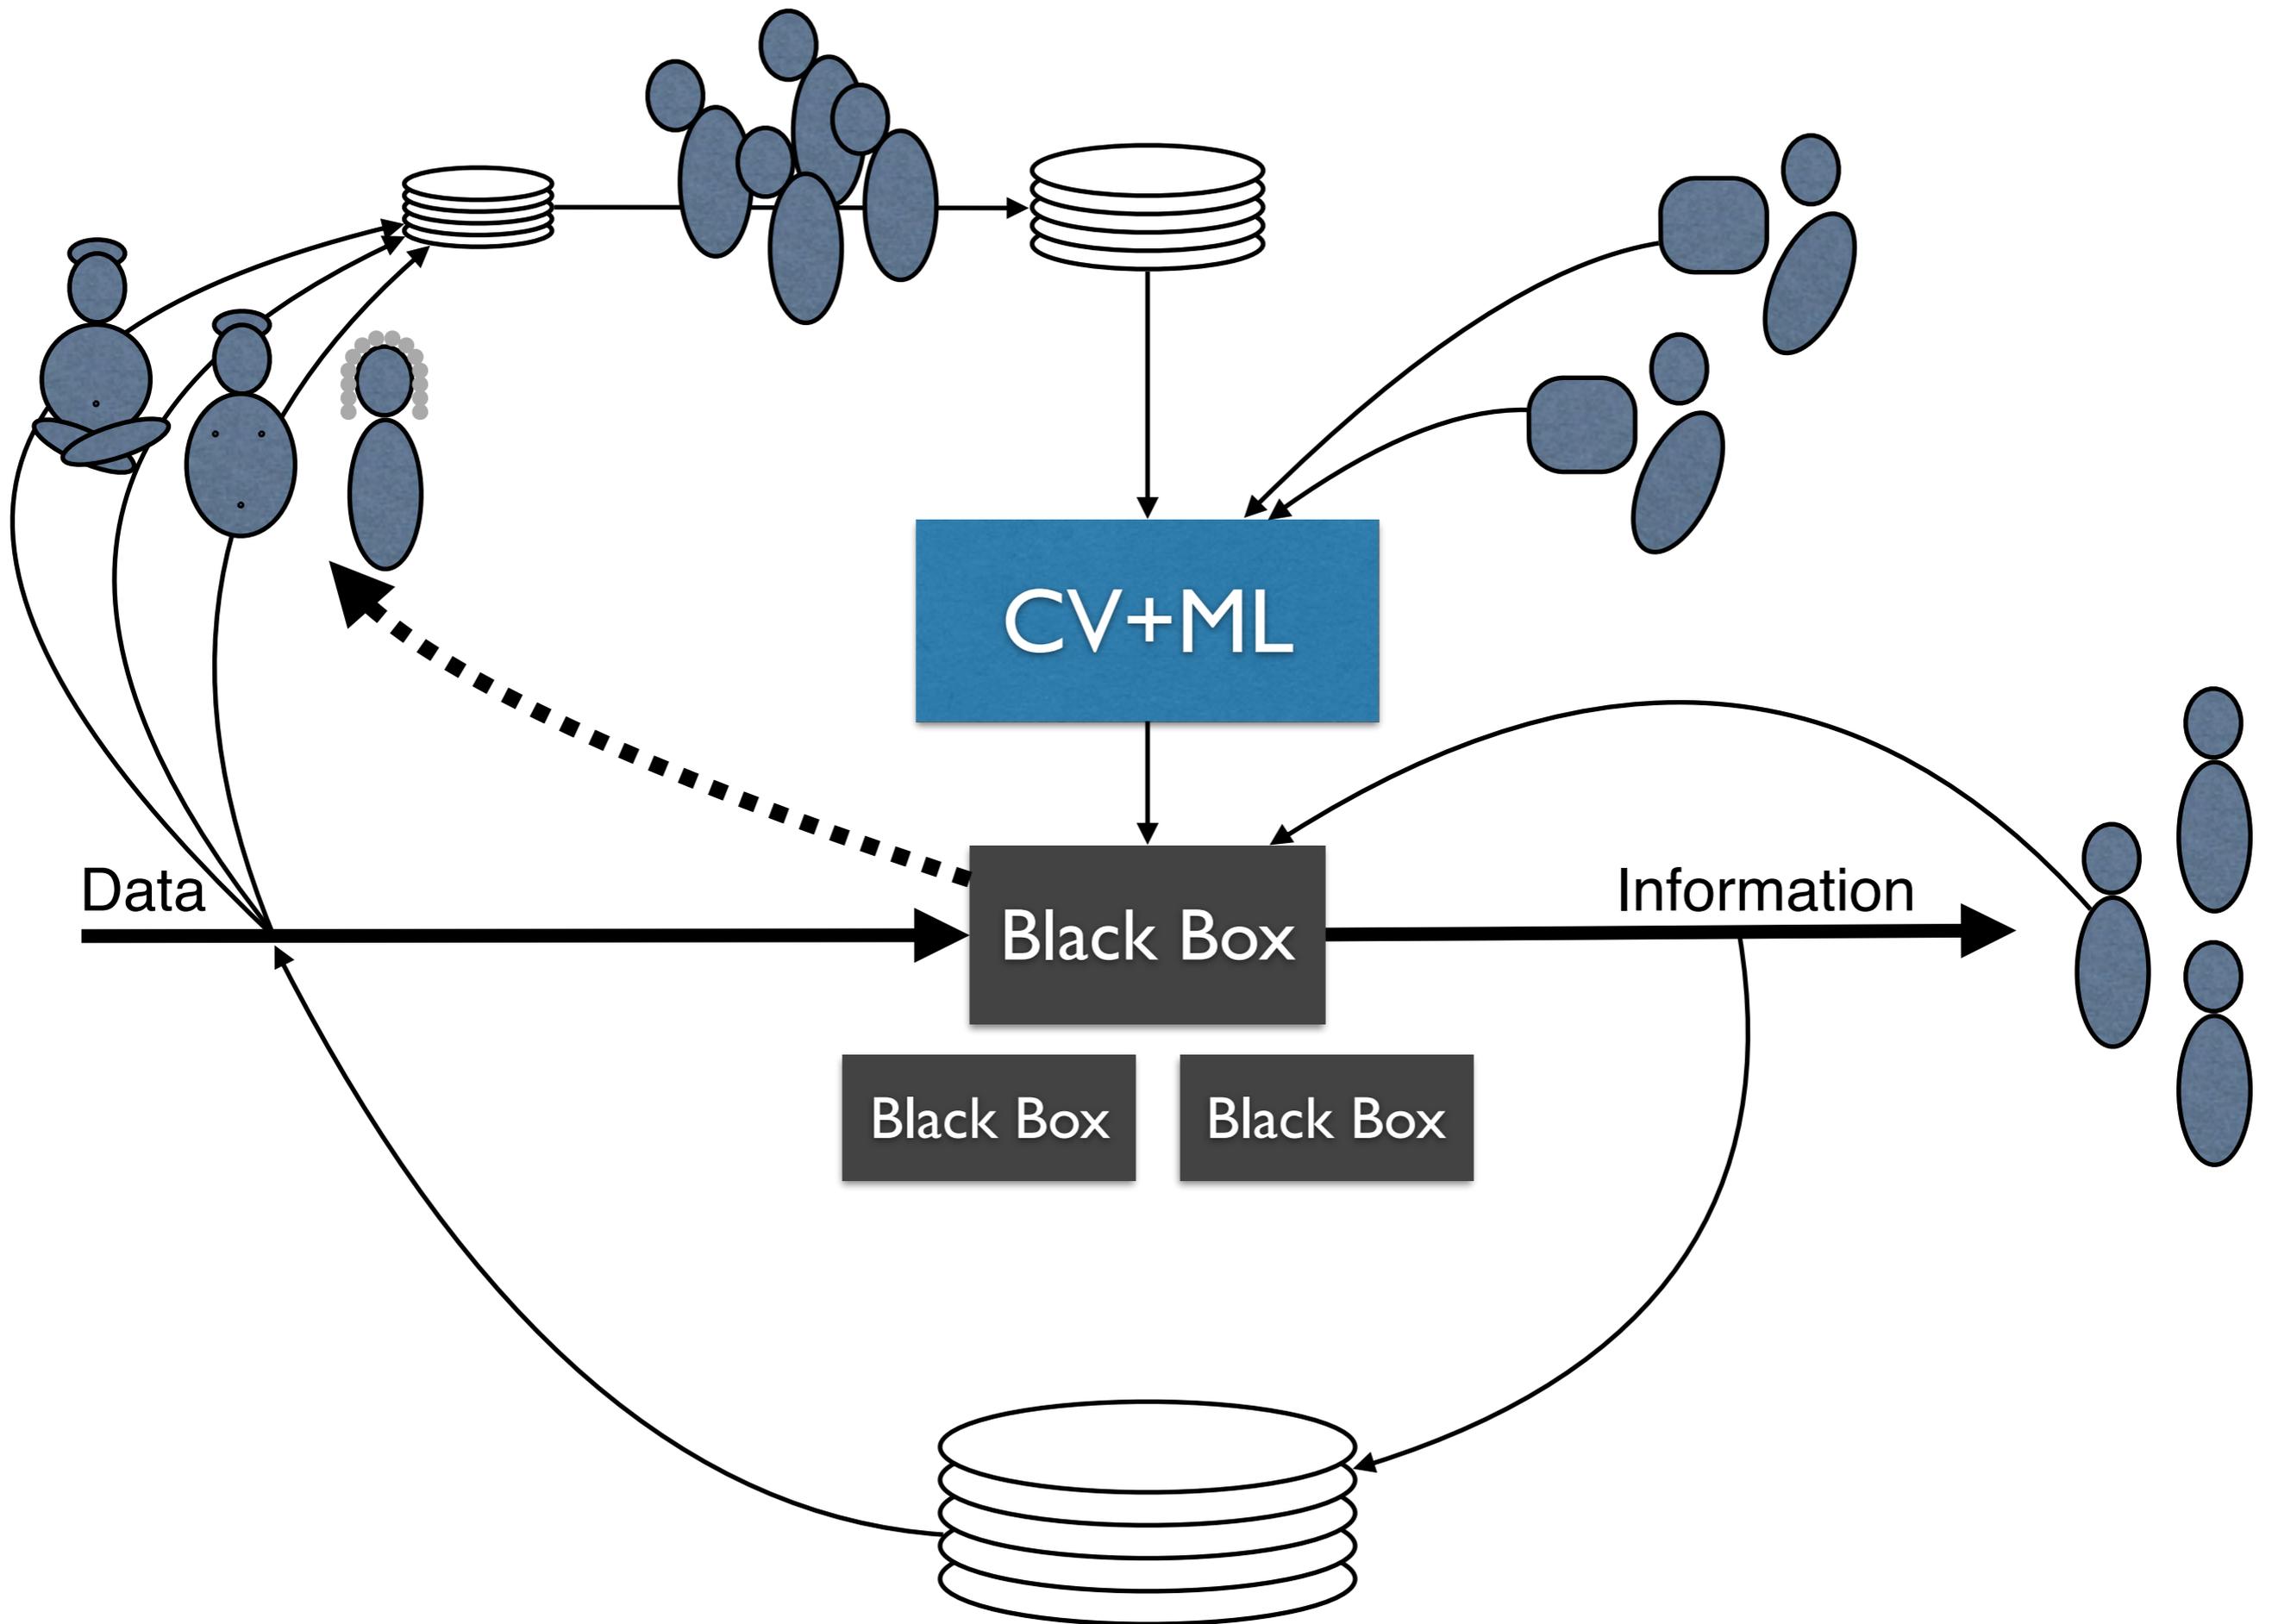
- Experiments  $\sim 10^4$  images,  $\sim 10^2$  workers
- Works well
- Better than one expert
- Any discoveries?



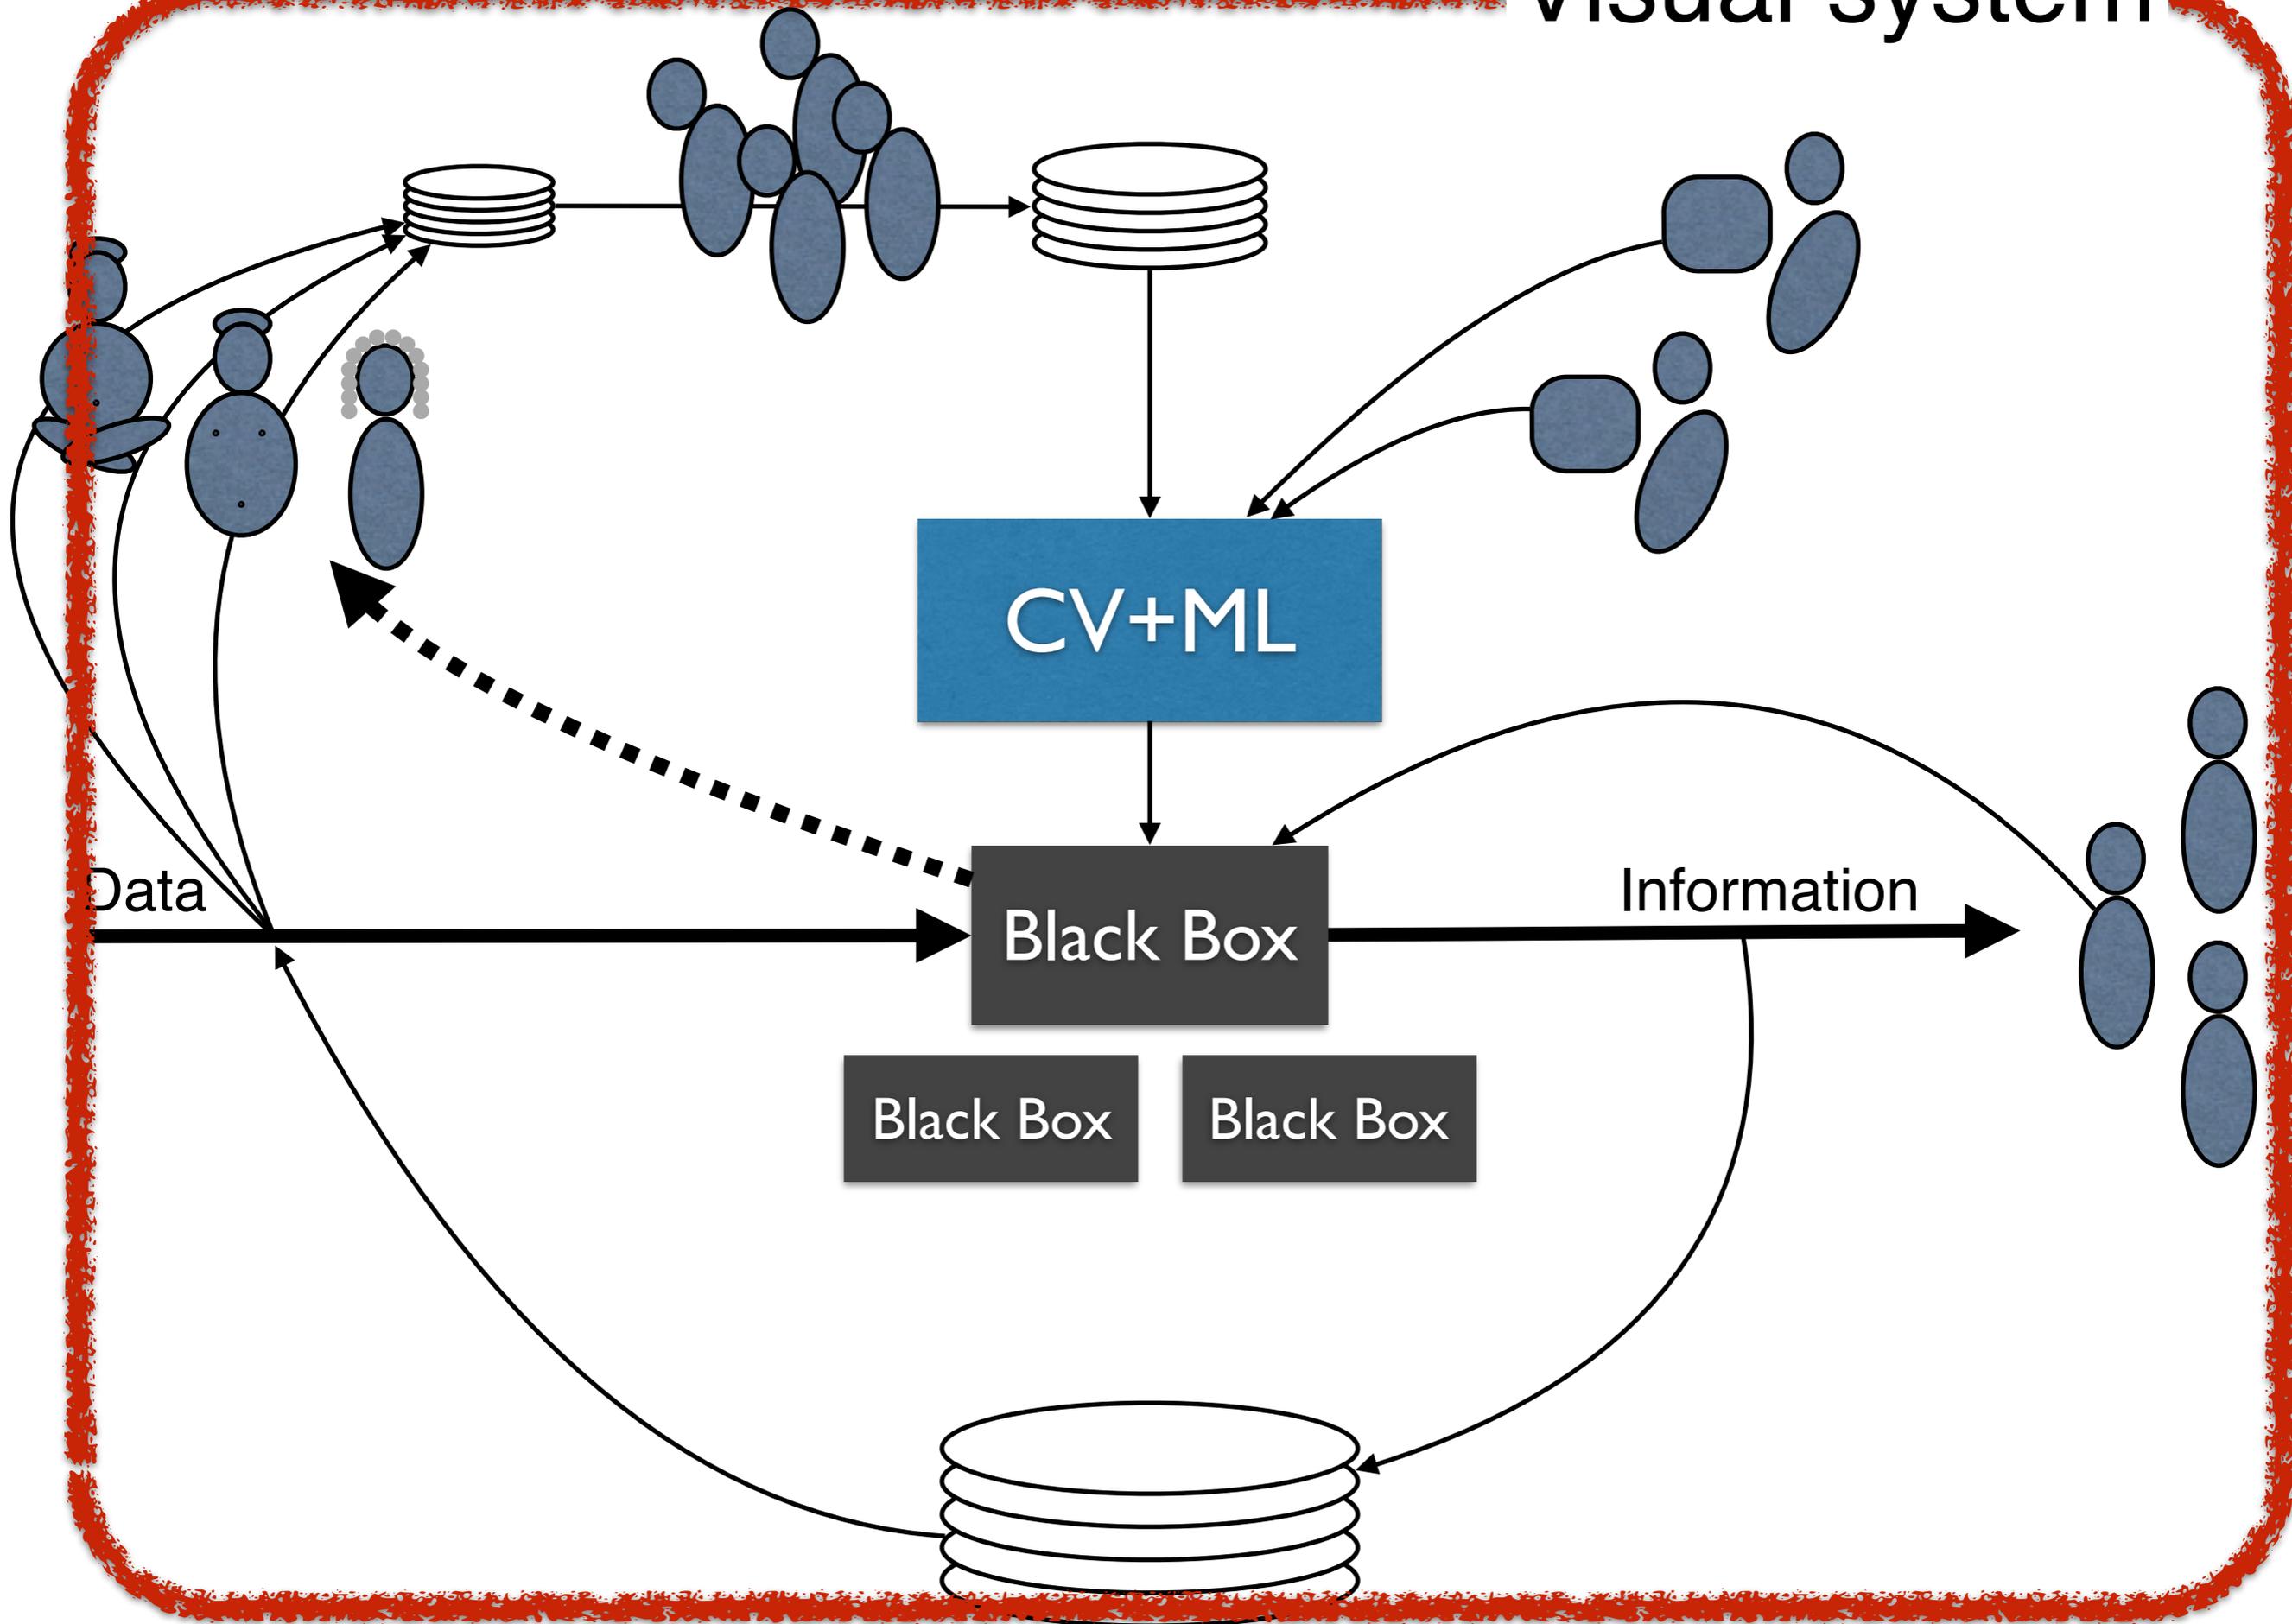








# Visual system



# Summary

- Vision and Knowledge
- Visipedia
- Crowdsourcing visual processing
- Crowdclustering

<http://www.vision.caltech.edu/visipedia/>

# Collaborators

- Serge Belongie
- Steve Branson
- Catherine Wah
- Peter Welinder
- Ryan Gomes

<http://www.vision.caltech.edu/visipedia/>