ZOØNIVERSE

Bayesian Active Learning for Probabilistic Galaxy Morphology

Mike Walmsley
with Chris Lintott, Lewis Smith, Yarin Gal, et al

University of Oxford

zooniverse.org @mike_w_ai

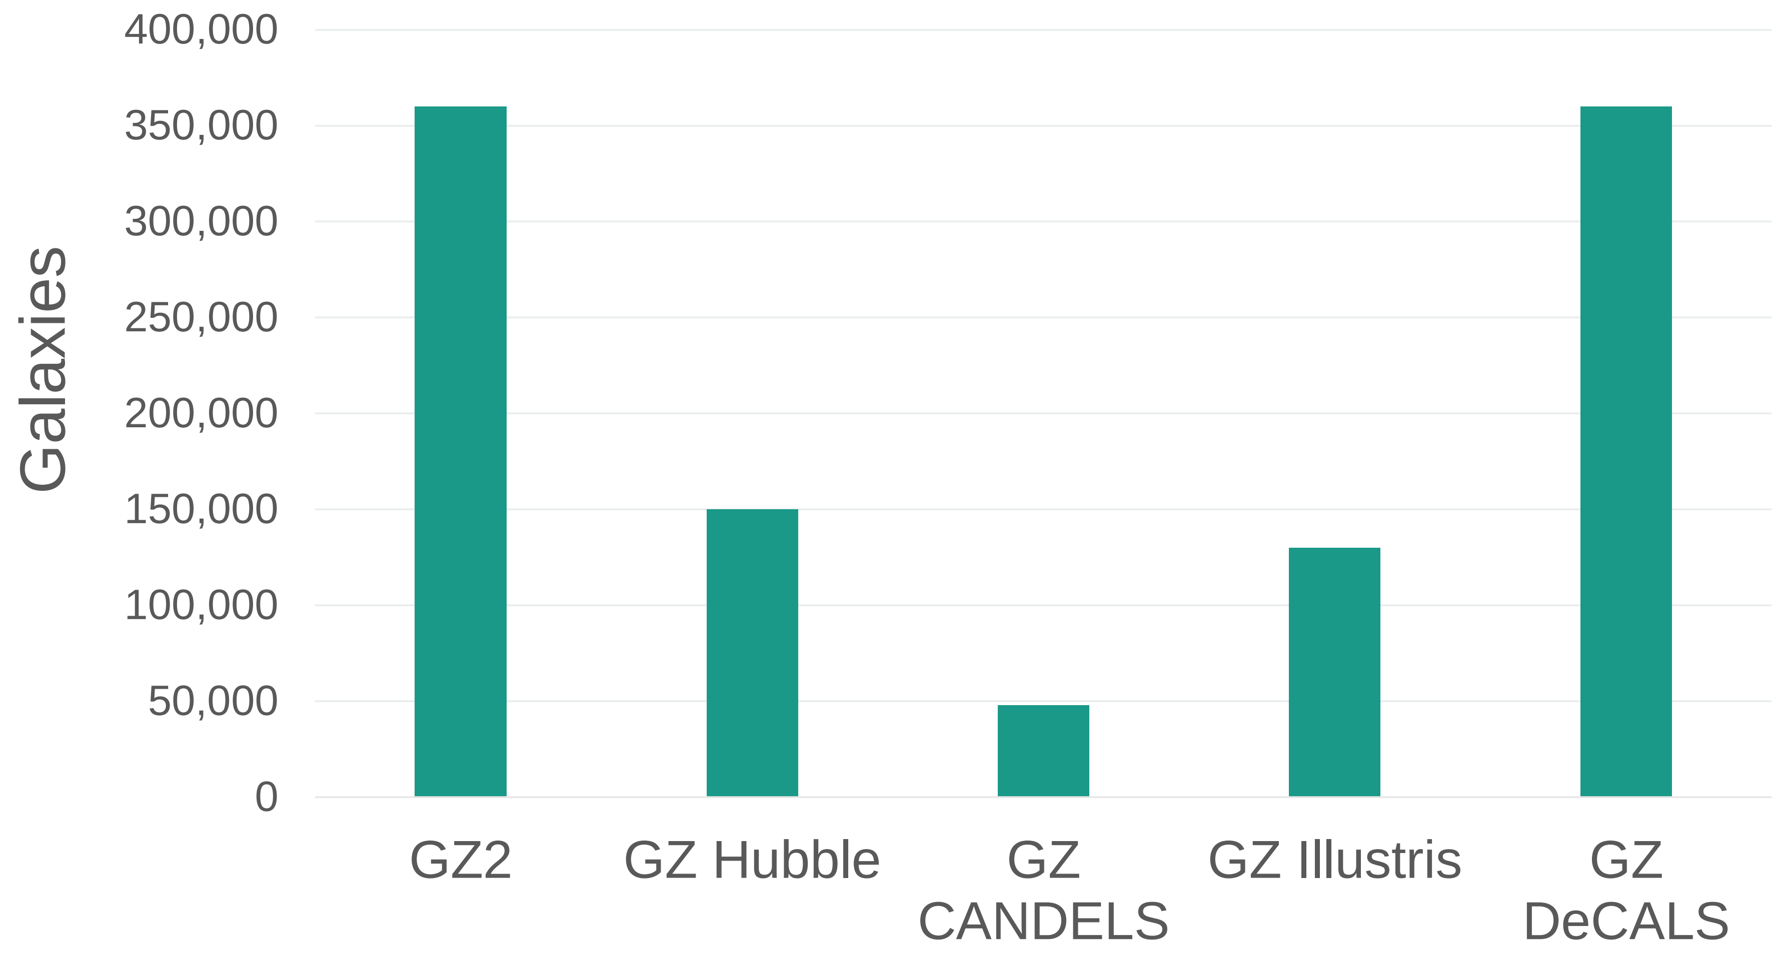
Galaxies
400,000
350,000
300,000
250,000
200,000
150,000
100,000
50,000
0
GZ2
GZ Hubble
GZ CANDELS
GZ Illustris
GZ DeCALS

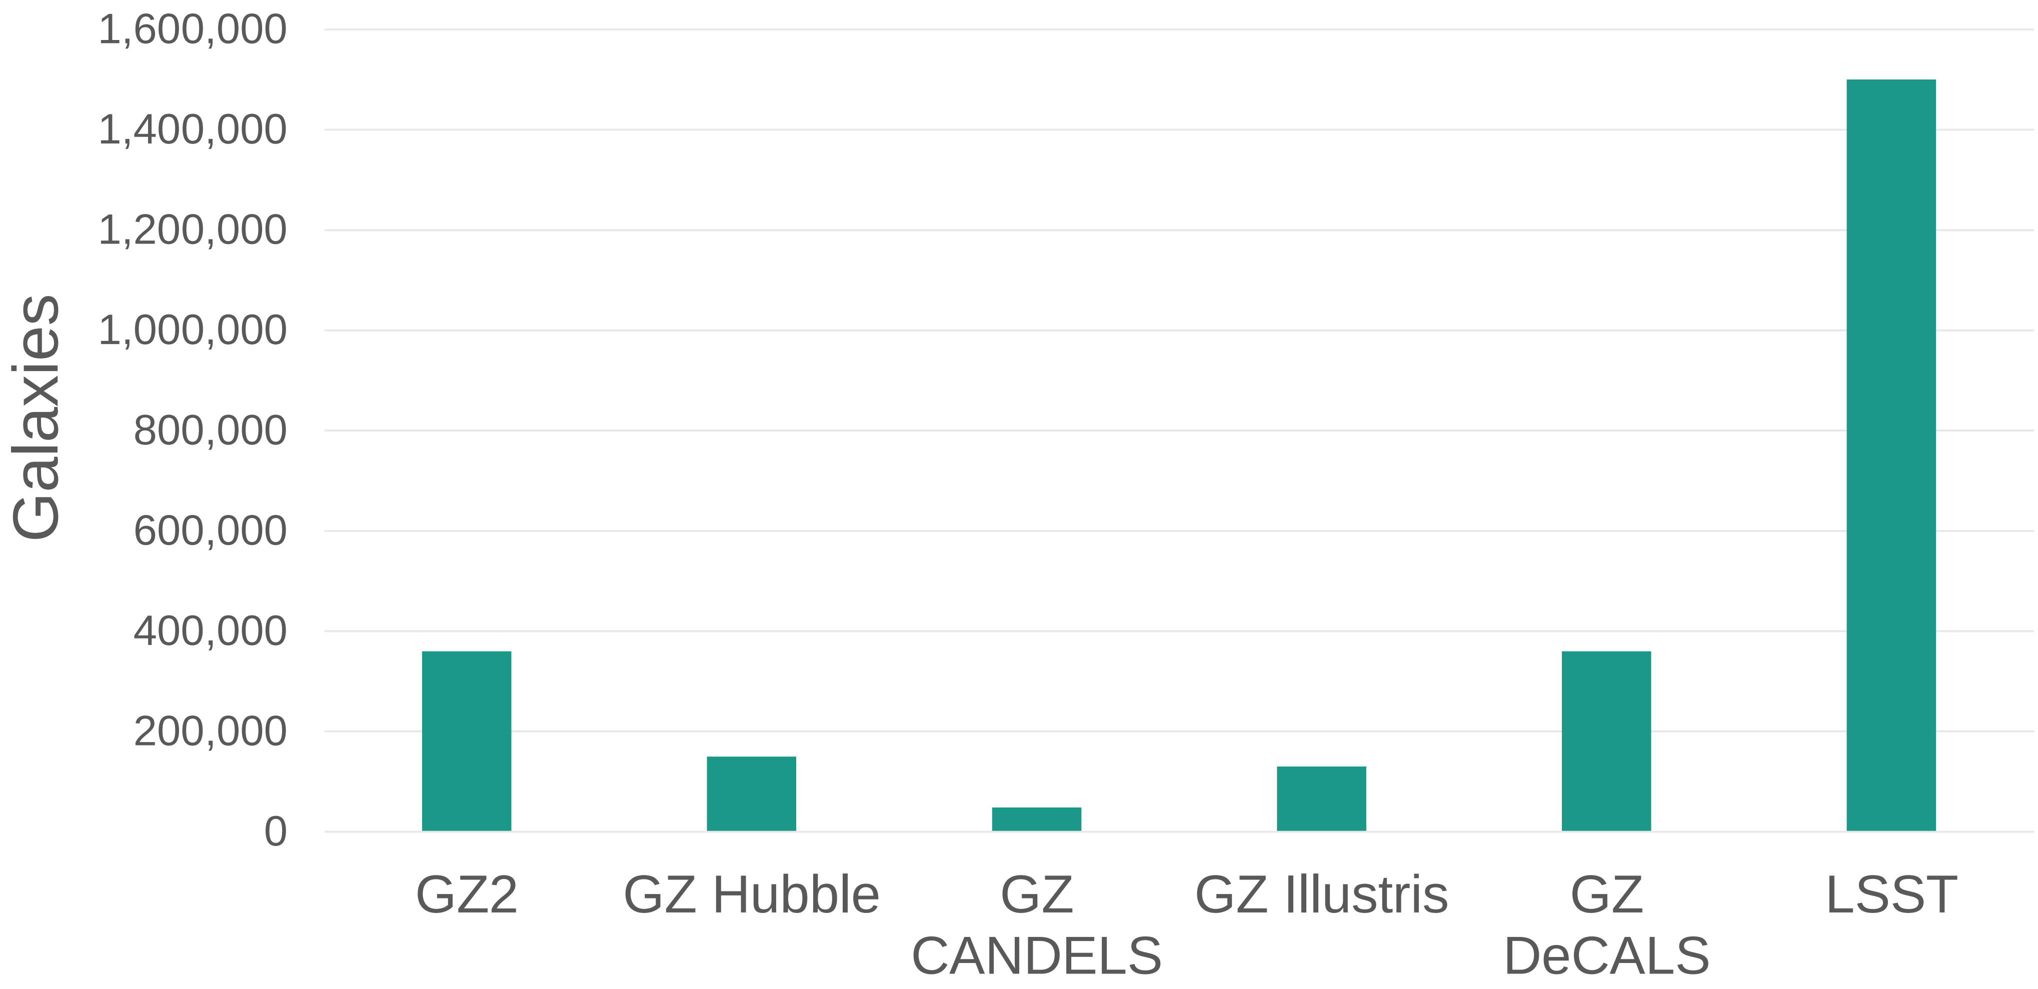
Galaxies
1,600,000
1,400,000
1,200,000
1,000,000
800,000
600,000
400,000
200,000
0
GZ2
GZ Hubble
GZ CANDELS
GZ Illustris
GZ DeCALS
LSST

ZOONIVERSE

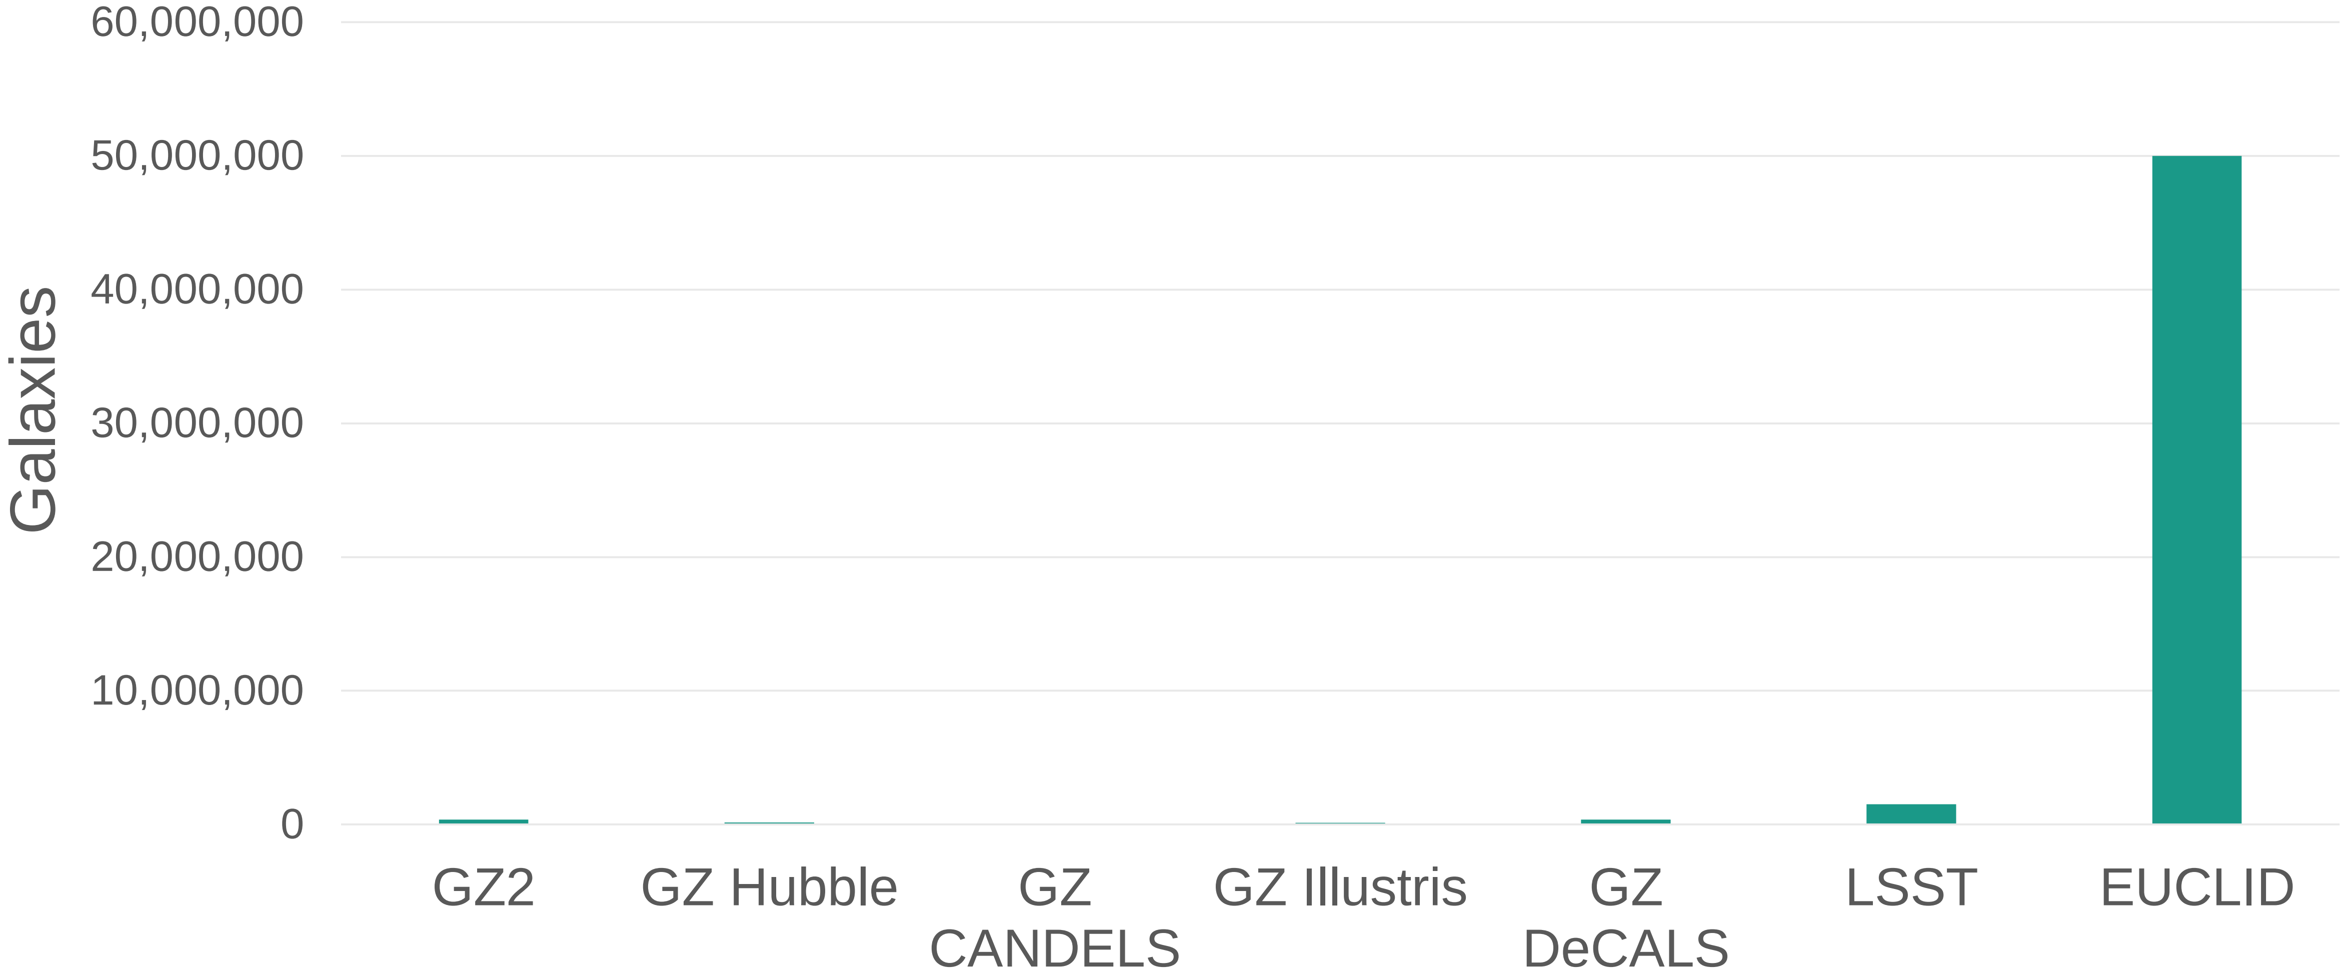

Galaxies
60,000,000
50,000,000
40,000,000
30,000,000
20,000,000
10,000,000
0
GZ2
GZ Hubble
GZ CANDELS
GZ Illustris
GZ DeCALS
LSST
EUCLID

BCNN with Two Twists:

1. Embrace Label Uncertainty

2. Apply Active Learning

Posteriors for Votes

- Our CNN can learn from uncertain labels and make probabilistic predictions $p(k|w)$

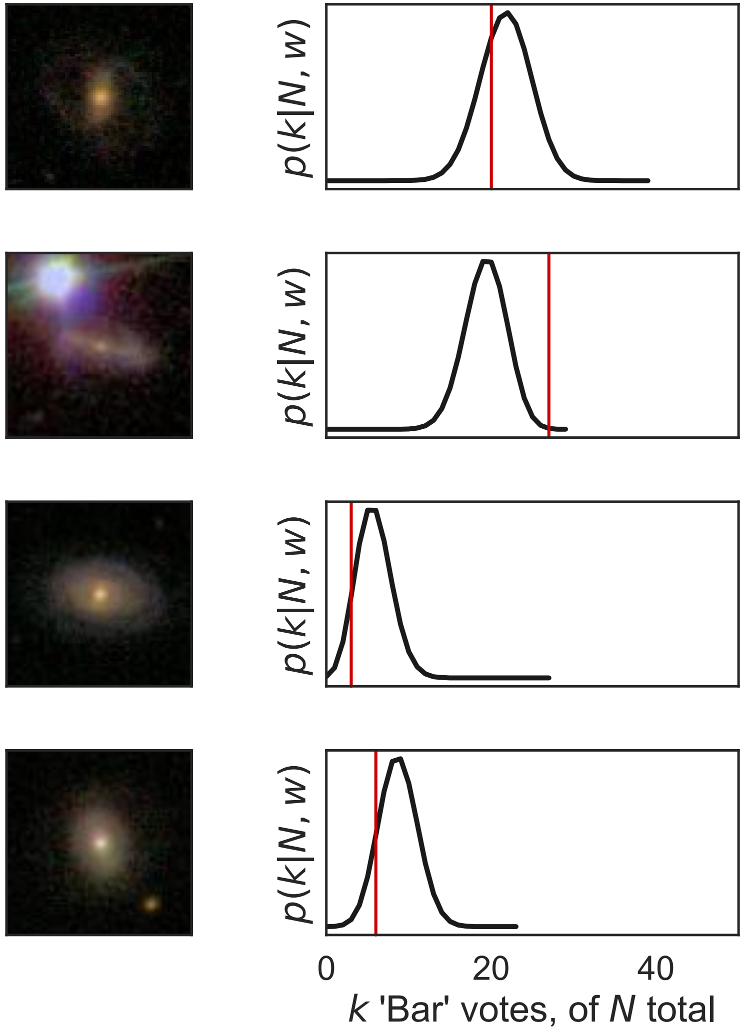

1 Model

Mike \

Probabilistic CNN

Volunteers N
Responses k
Typical vote prob. ρ
Galaxy x
CNN output $f^w(x)$

N volunteers and k responses ≈ N trials and k successes

Volunteer model: $p(k \mid x, N) = Bin(k \mid \rho, N)$

How likely is each ρ given observed k, N?

Log Likelihood $\mathcal{L} = \log[\, Bin(k \mid \rho, N)\,] = k \log(\rho) + (N - k) \log(1 - \rho) + C$

Predict $f^w(x) = \hat{\rho}$ and maximise likelihood of $\hat{\rho}$

Posteriors for Votes

- Our CNN can learn from uncertain labels and make probabilistic predictions $p(k|w)$

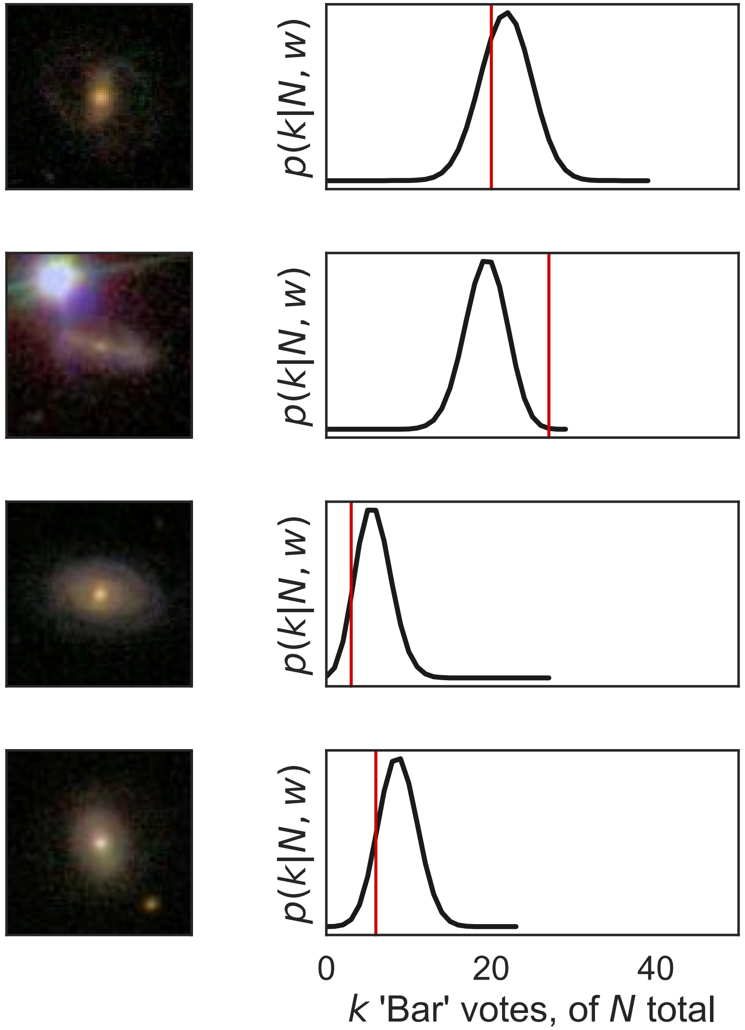

1 Model

Mike \

For details on BCNN, see Y. Gal (2016)

Posteriors for Votes

- Our CNN can learn from uncertain labels and make probabilistic predictions $p(k|w)$

- Marginalising over weights (BCNN) lets us predict votes *over all CNN we might have trained*

$$p(k|D) = \int p(v|w)\, p(w|D) dw$$

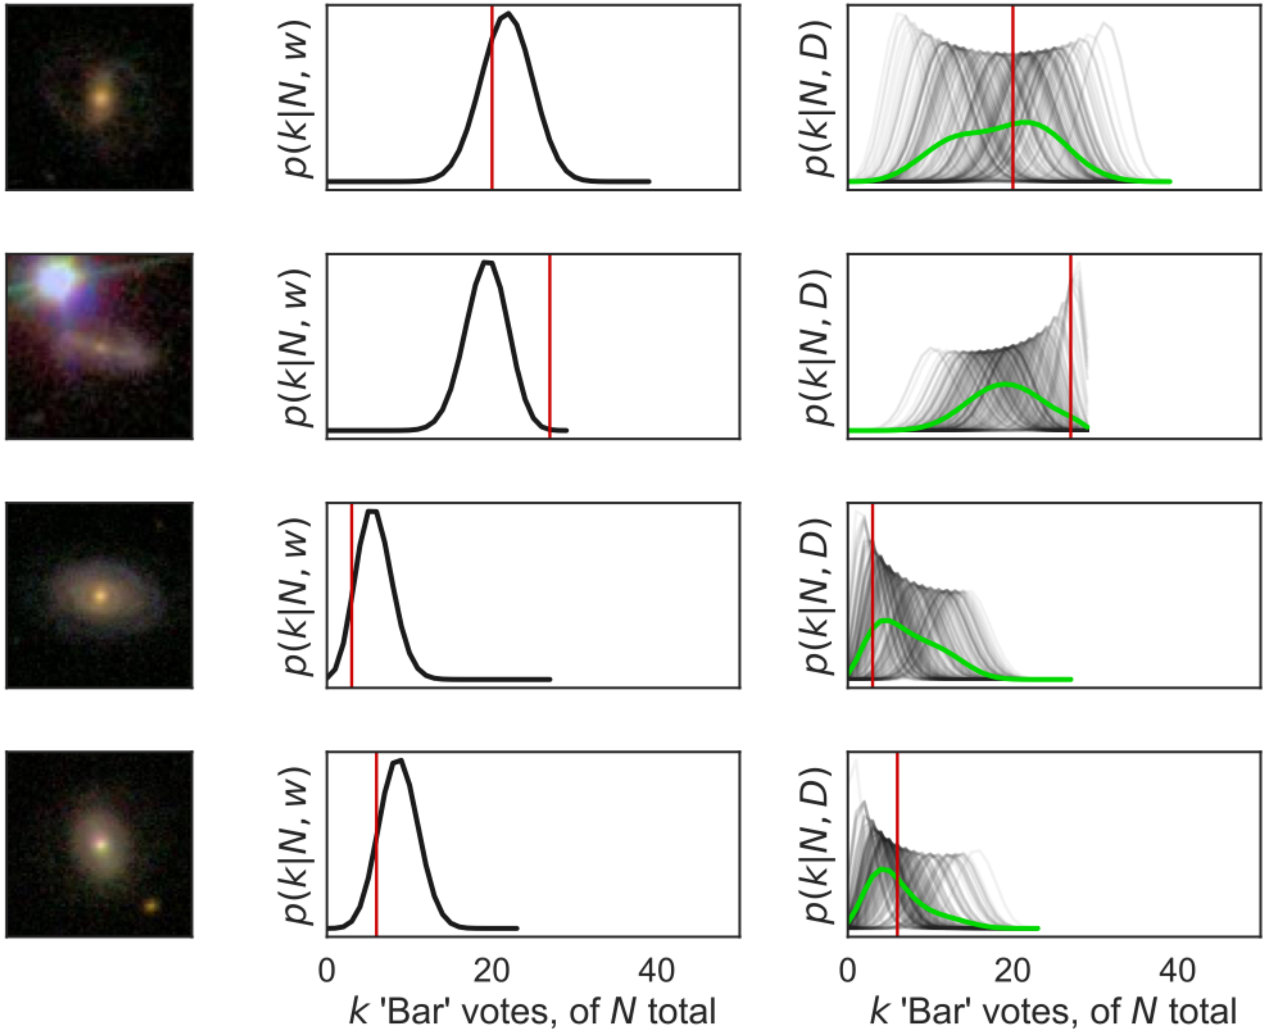

1 Model

30 “Models” (BCNN)

ZOØNIVERSE

For details on BCNN, see Y. Gal (2016)

Posteriors for Votes

- Marginalising over weights (BCNN) lets us predict votes *over all CNN we might have trained*

$$p(k|D) = \int p(v|w)\, p(w|D)dw$$

- BCNN posteriors are much better calibrated

Result: Morphology catalog with trustworthy uncertainties

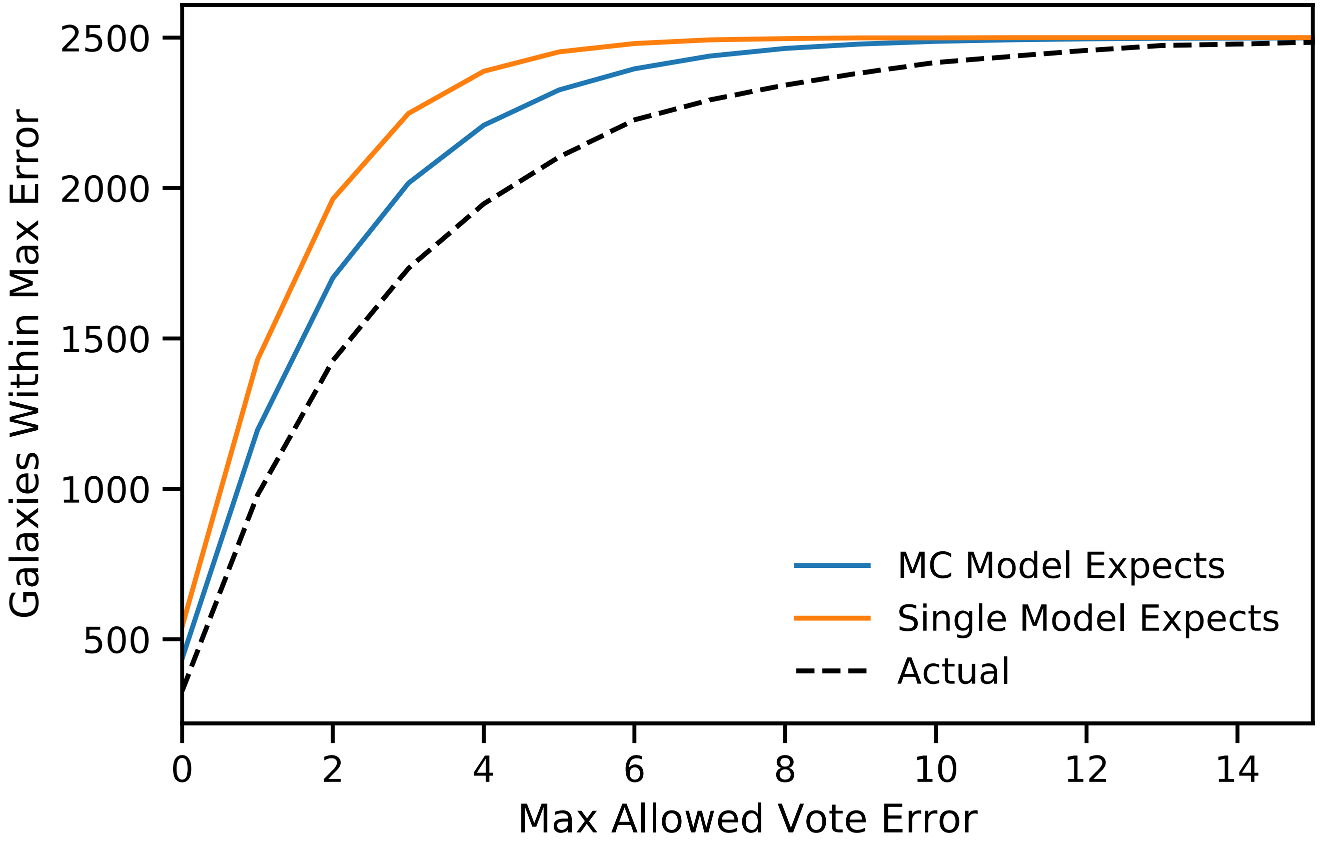

Uncertainty Matters – Predicting Expert Opinion

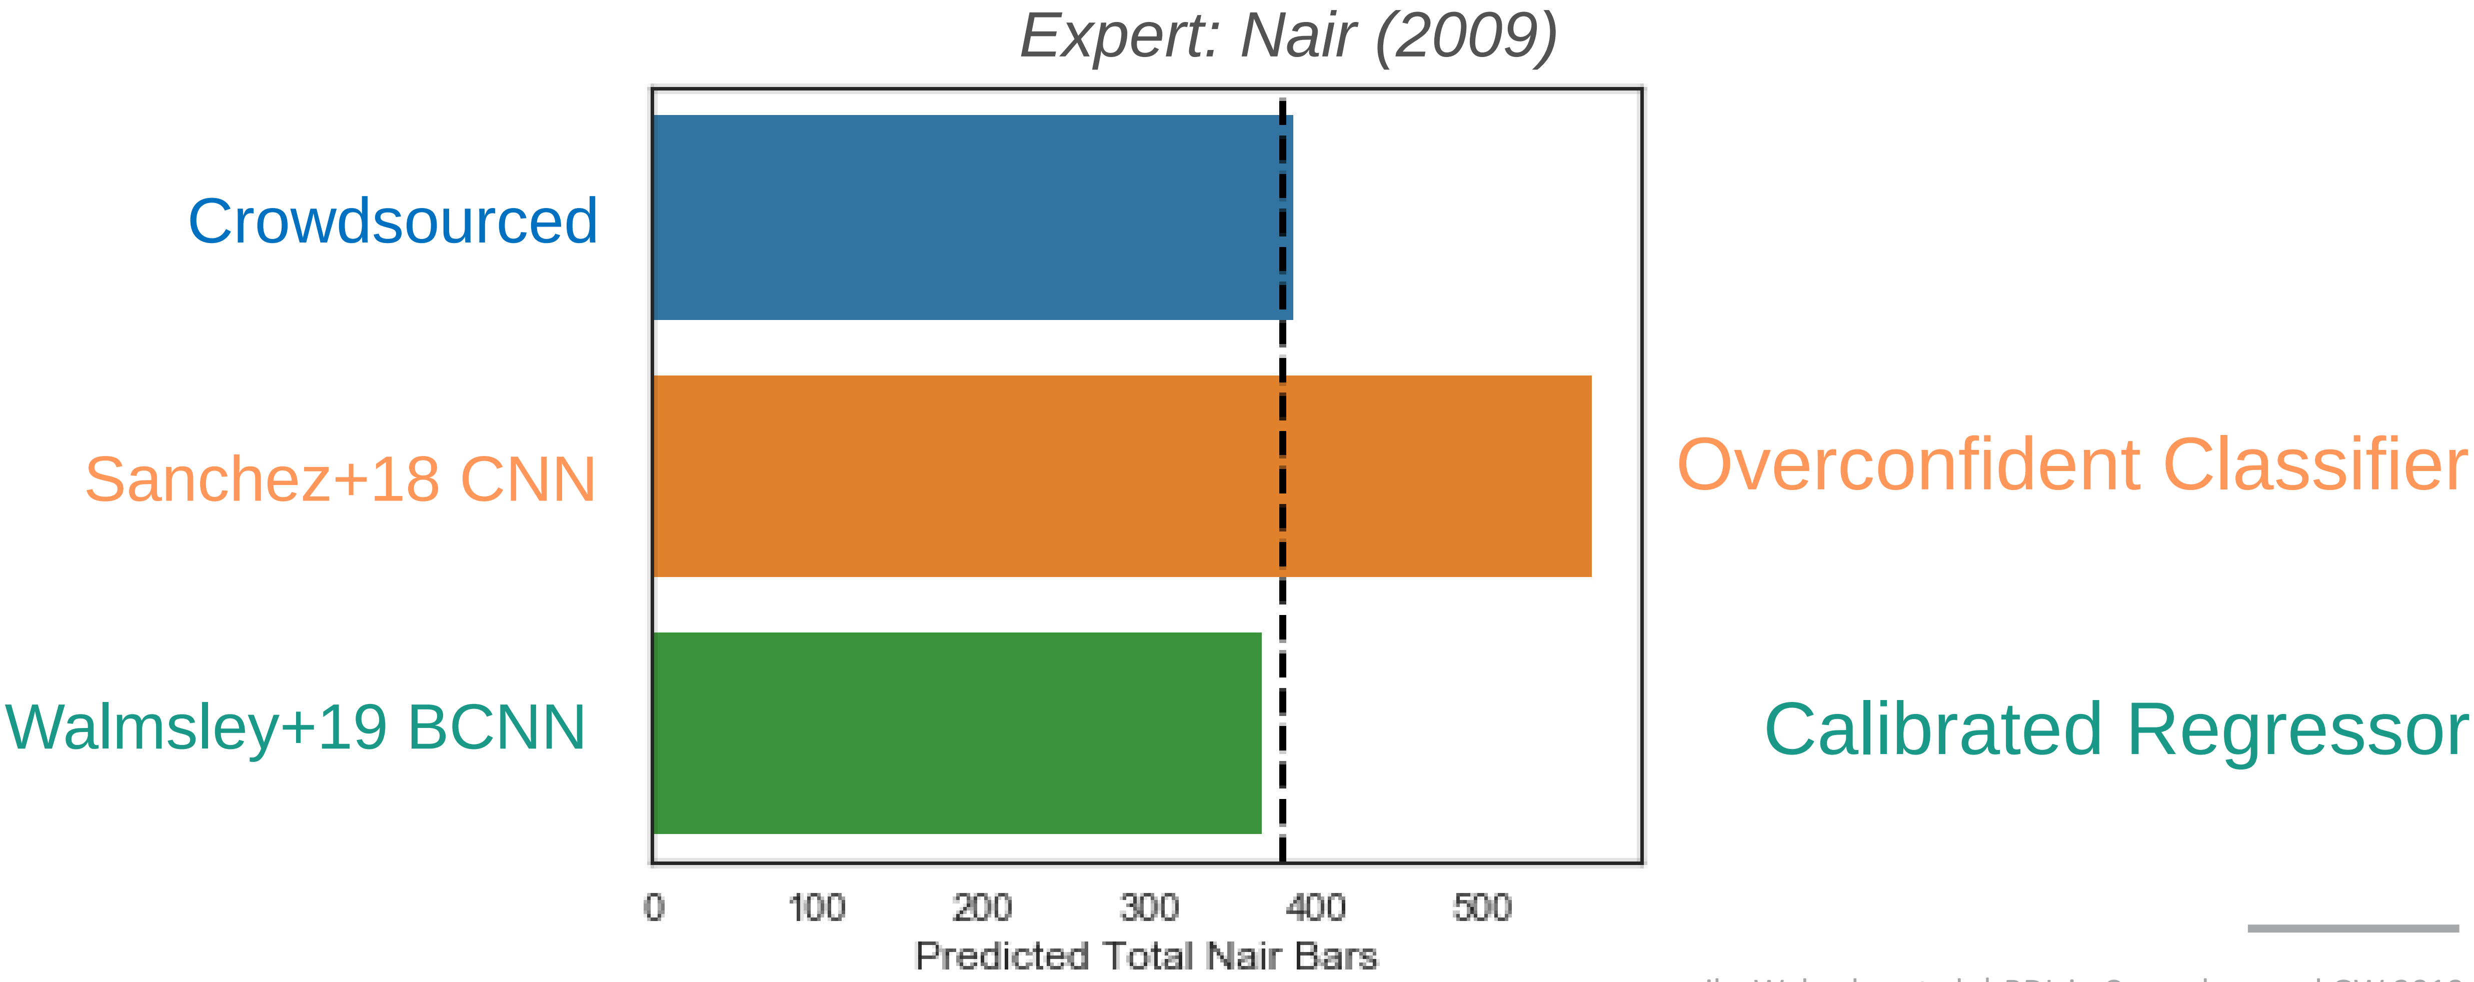

ZOØNIVERSE

BCNN with Two Twists:

Predictions must be uncertain:

Drawing scientific conclusions relies on quantifying errors

1. Embrace Label Uncertainty

2. Apply Active Learning

Principles for Choosing Data

Mutual Information $I[X,Y] = H[p(X)] - E_{p(Y)}H[P(X|Y)]$

Information Gain $U(x) = H[p(\theta|D)] - E_{p(y|x,D)}H[p(\theta|D,x,y)]$

For details on Mutual Information, see N. Houlsby (2014)

Principles for Choosing Data

Information Gain

$$U(x) = H[p(\theta|D)] - E_{p(y|x,D)} H[p(\theta|D,x,y)]$$

$$= I[\theta, y \mid D, x]$$

…rearrange by symmetry of I…

$$= I[y, \theta| D, x]$$

$$= H[p(y|x,D)] - E_{p(\theta|D)} H[p(y|x,\theta)]$$

For details on Mutual Information, see N. Houlsby (2014)

Mutual Information

Mutual Information I
Entropy H
Votes k
Weights θ
Training data D

Pick galaxies where the models **confidently** disagree.

$$I = -\int H[p(k|\theta)]\; p(\theta|D)\, d\theta + H\left[\int p(k|\theta)\, p(\theta|D)\, d\theta\right]$$

Each model is confident…

…but they give different answers

Only possible because we:

- Think about labels probabilistically, $p(k|\theta)$
- Approximate training many models with BCNN, $p(\theta|D)$

$$H = -\sum_{i}^{N} p(k=i) \log p(k=i)$$

For details on Mutual Information, see N. Houlsby (2014)

Active Learning Results

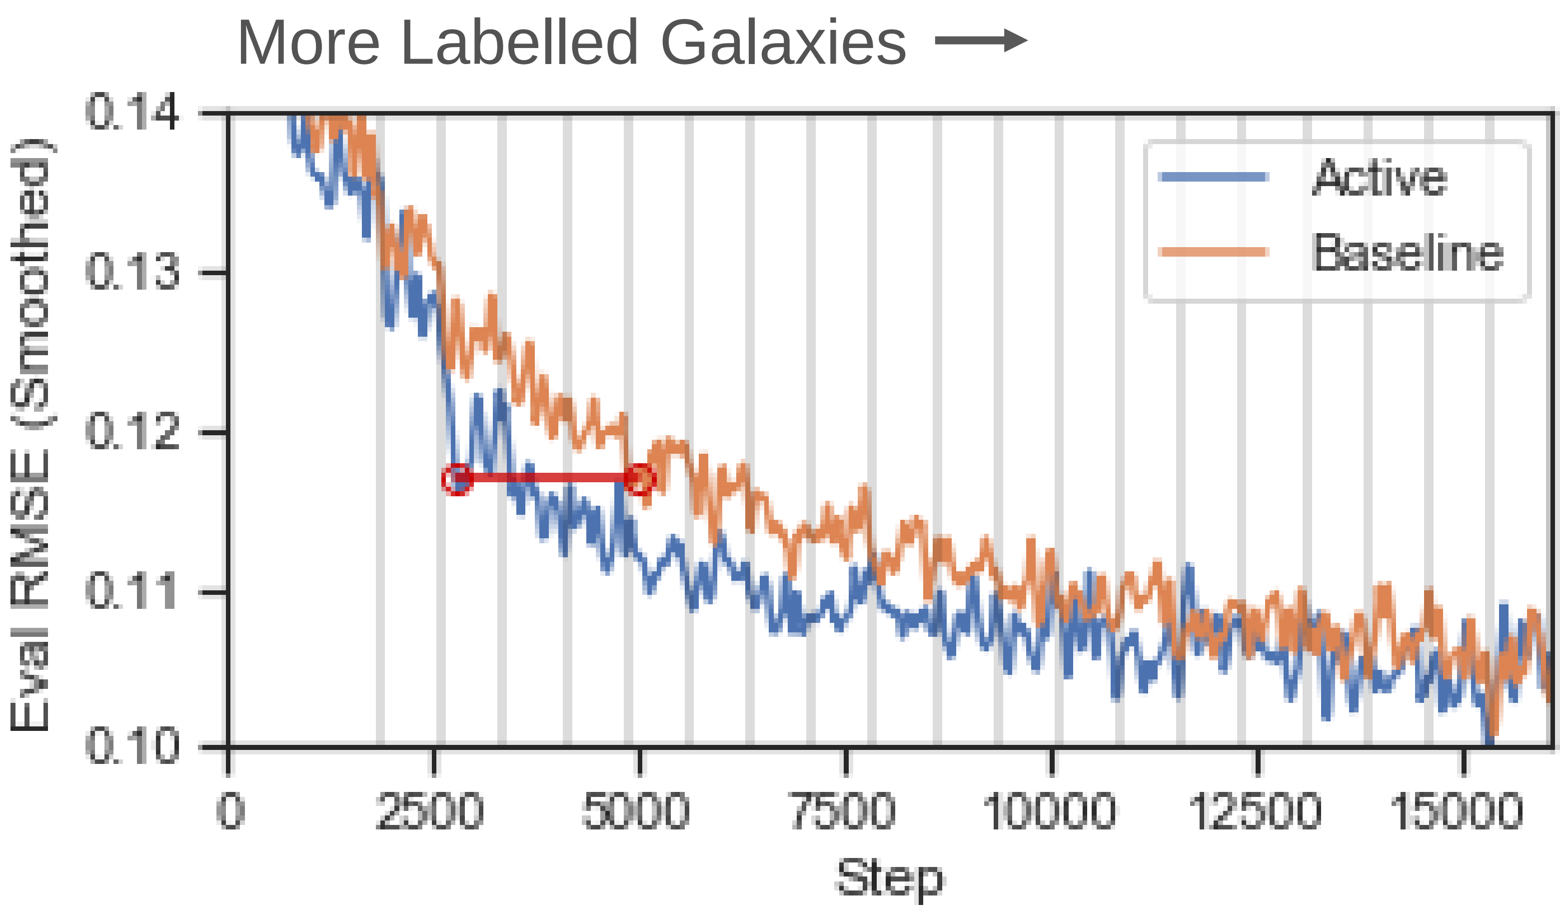

Selected Galaxies for “Smooth?”

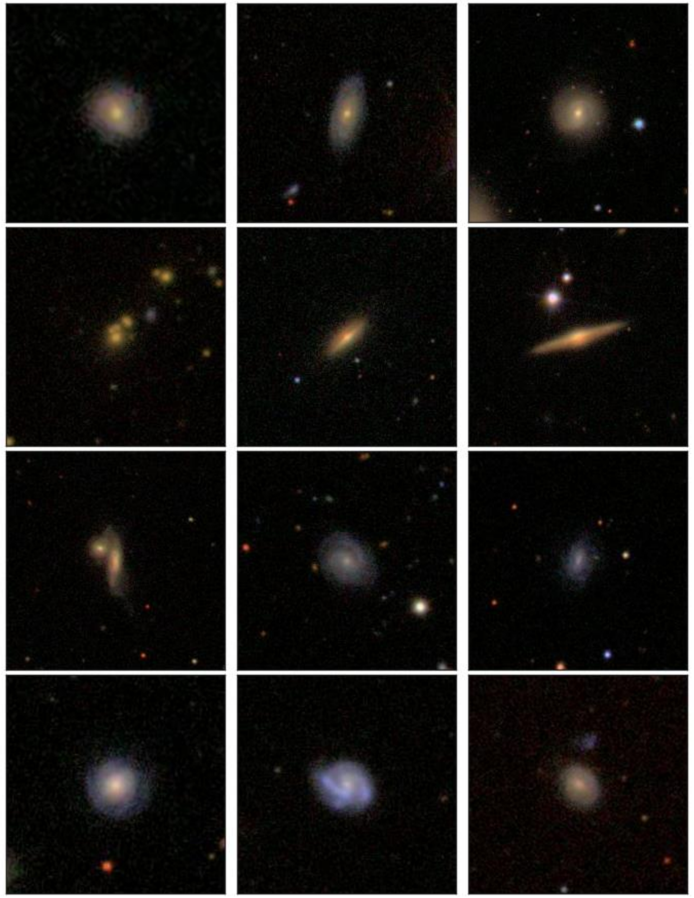

High mutual information

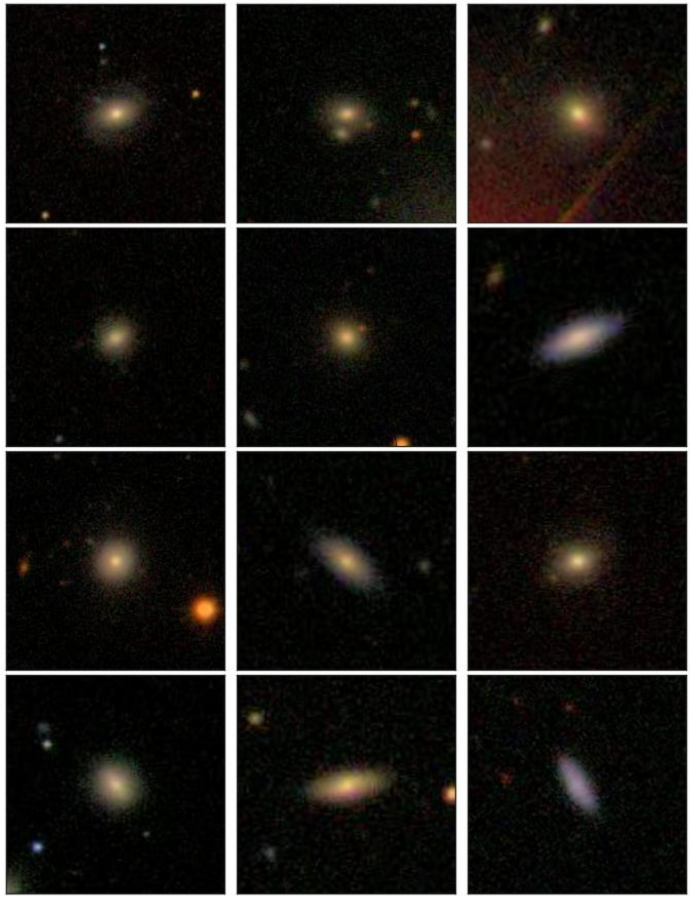

Low mutual information

Selected Galaxies for "Smooth?"

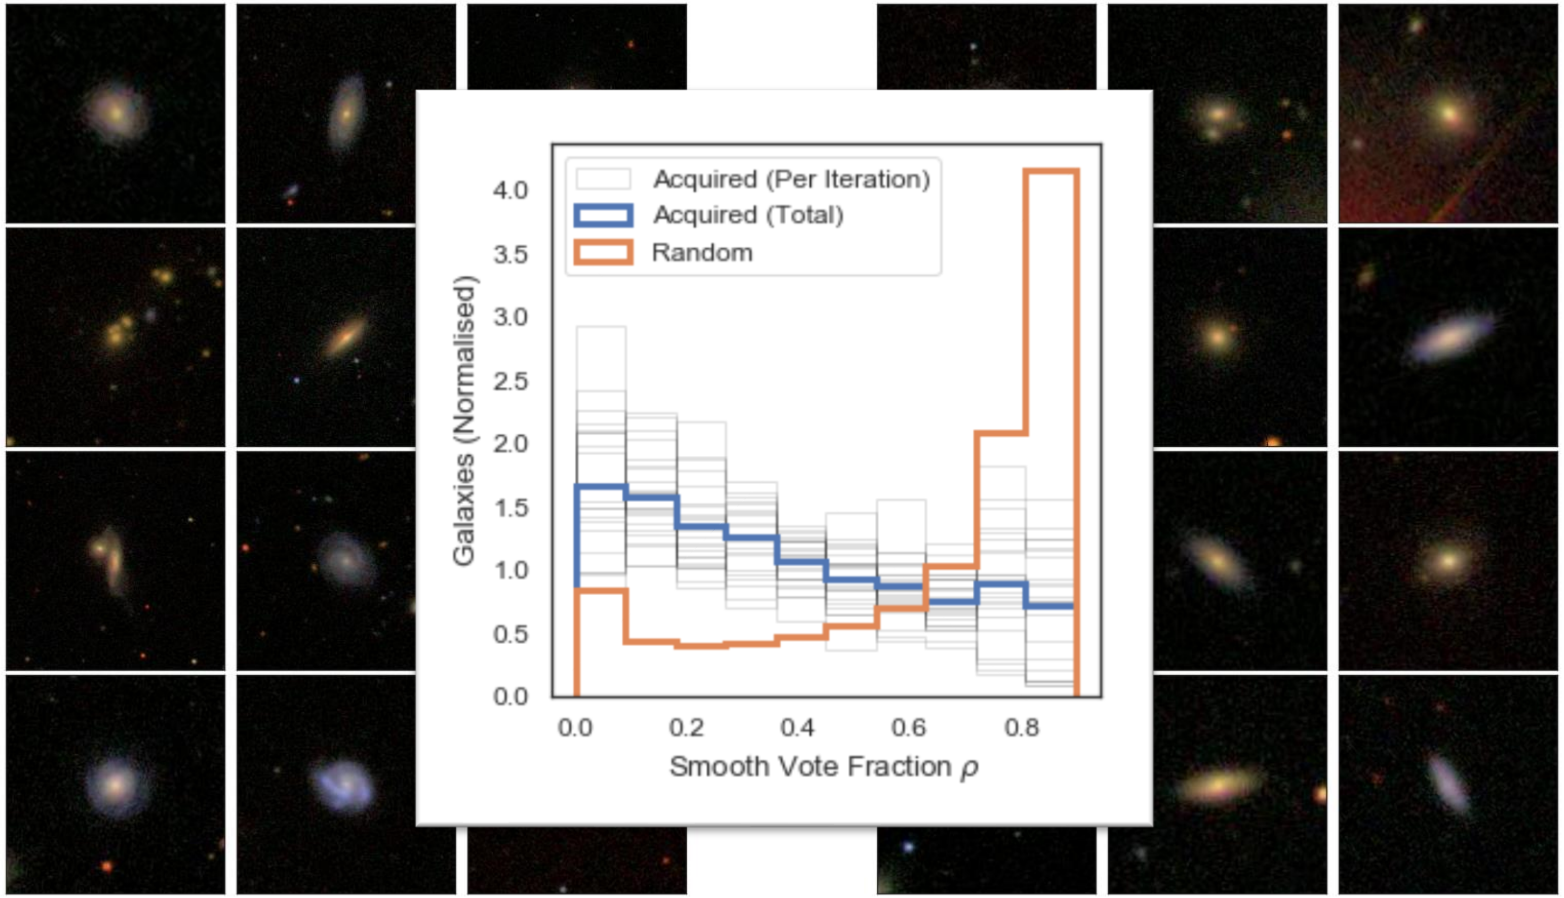

High mutual information

Low mutual information

Selected Galaxies for “Bar?”

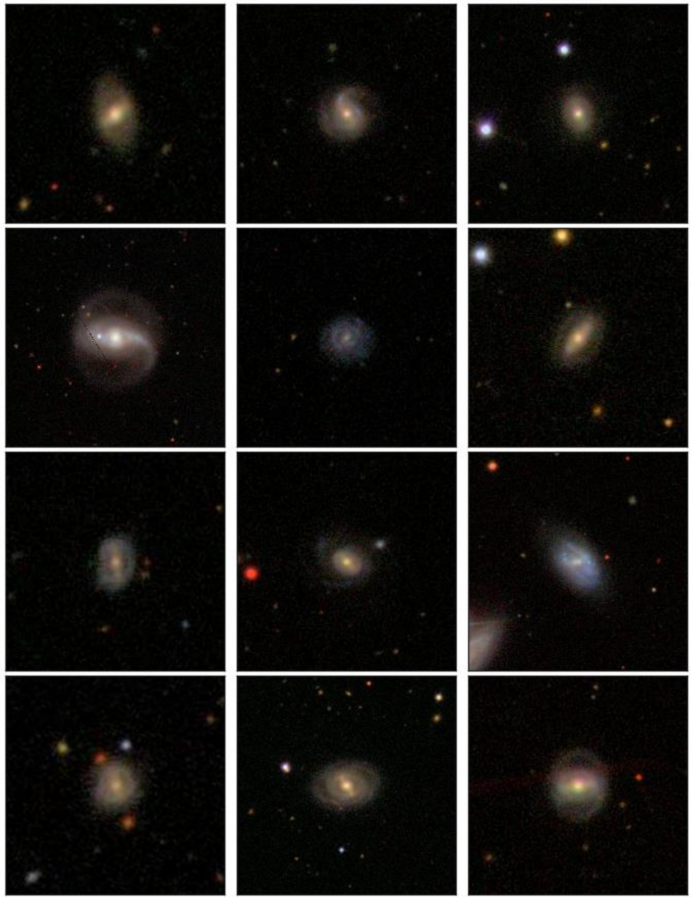

High mutual information

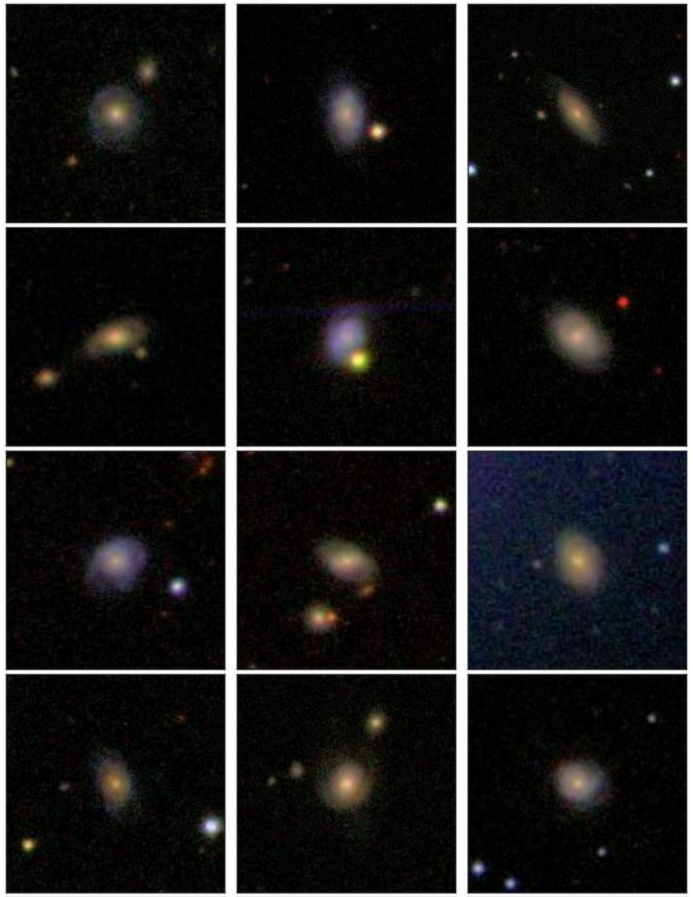

Low mutual information

Selected Galaxies for “Bar?”

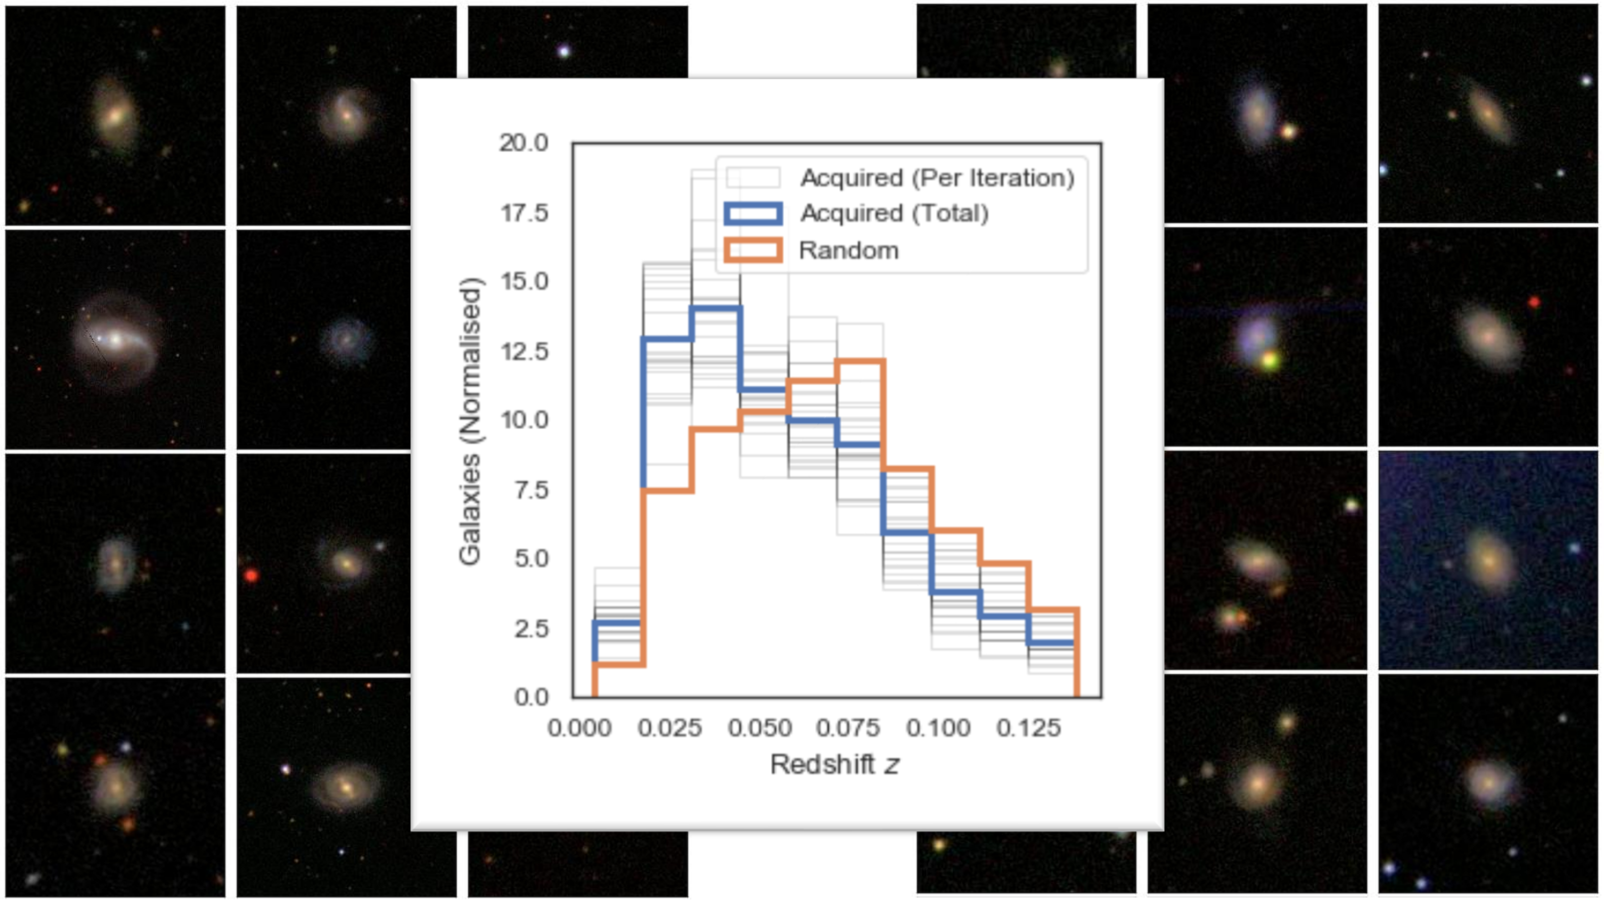

High mutual information

Low mutual information

Live on Galaxy Zoo

Model retrains and requests new classifications weekly

New surveys get classified in weeks, not years

Every galaxy seen by at least 3 volunteers

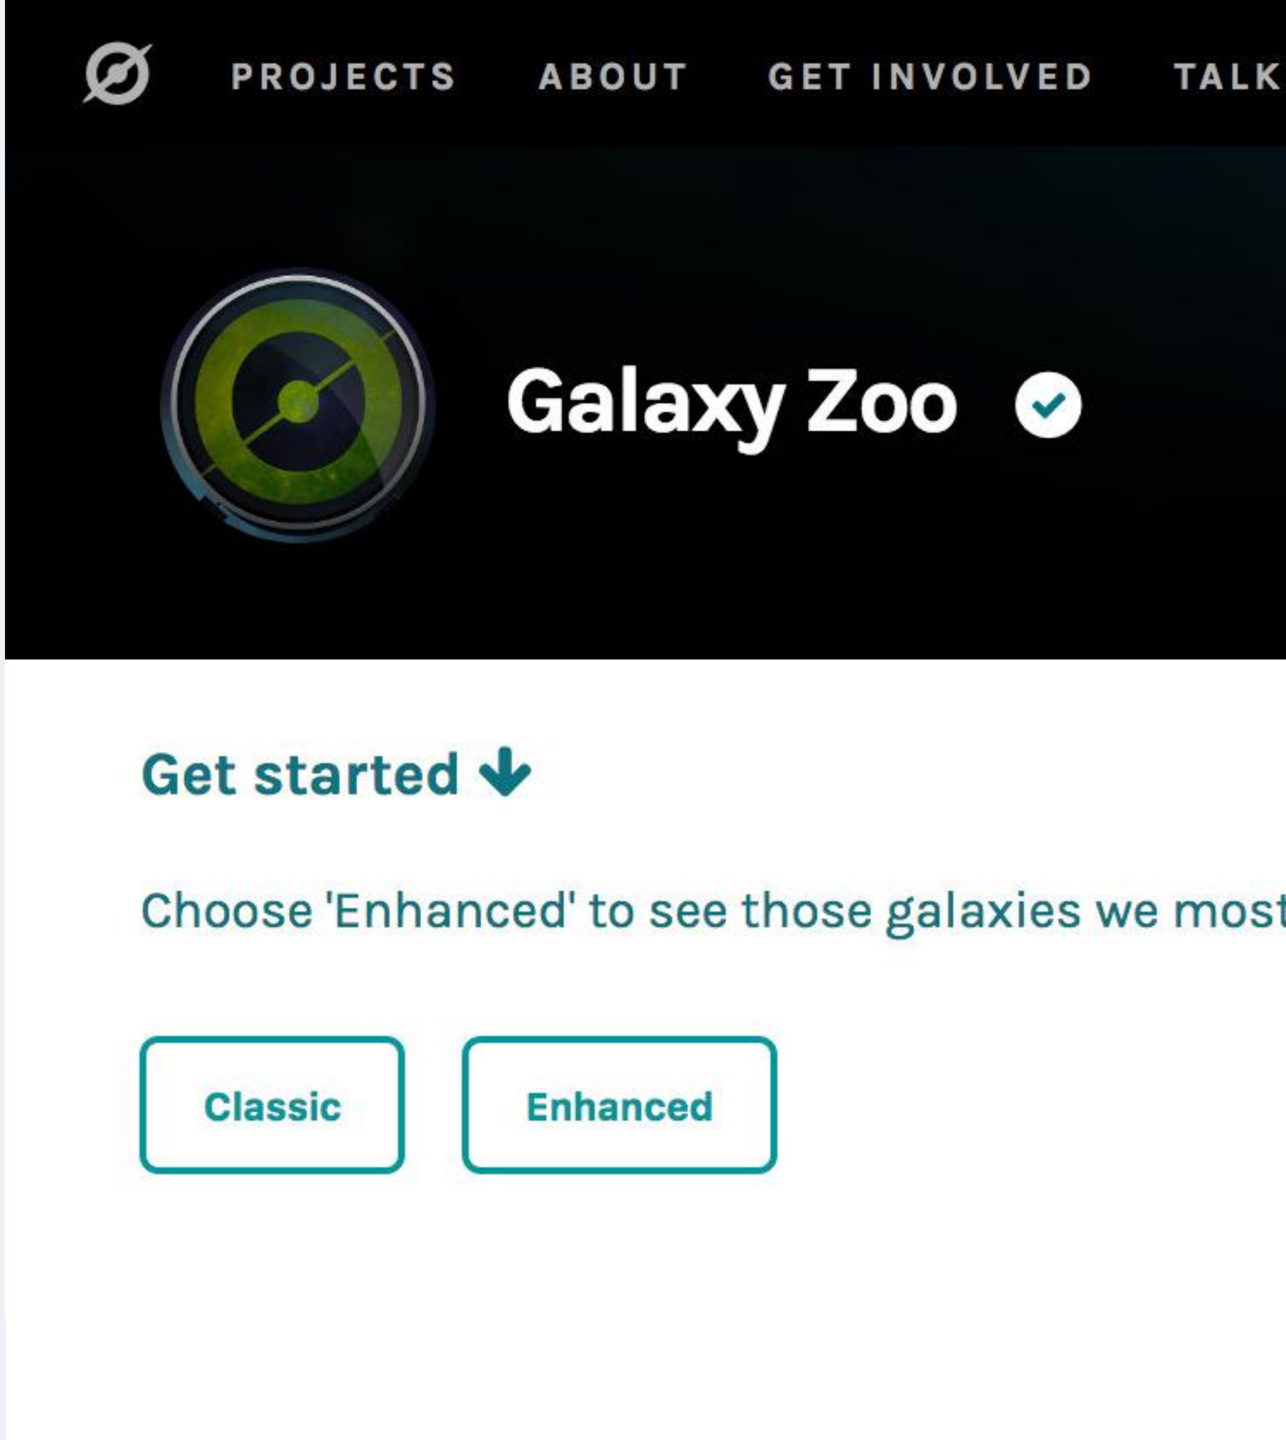

ZOØNIVERSE

arxiv: 1905.07424

@mike_w_ai
@chrislintott
@yaringal
lewis.smith@kellogg.ox.ac.uk