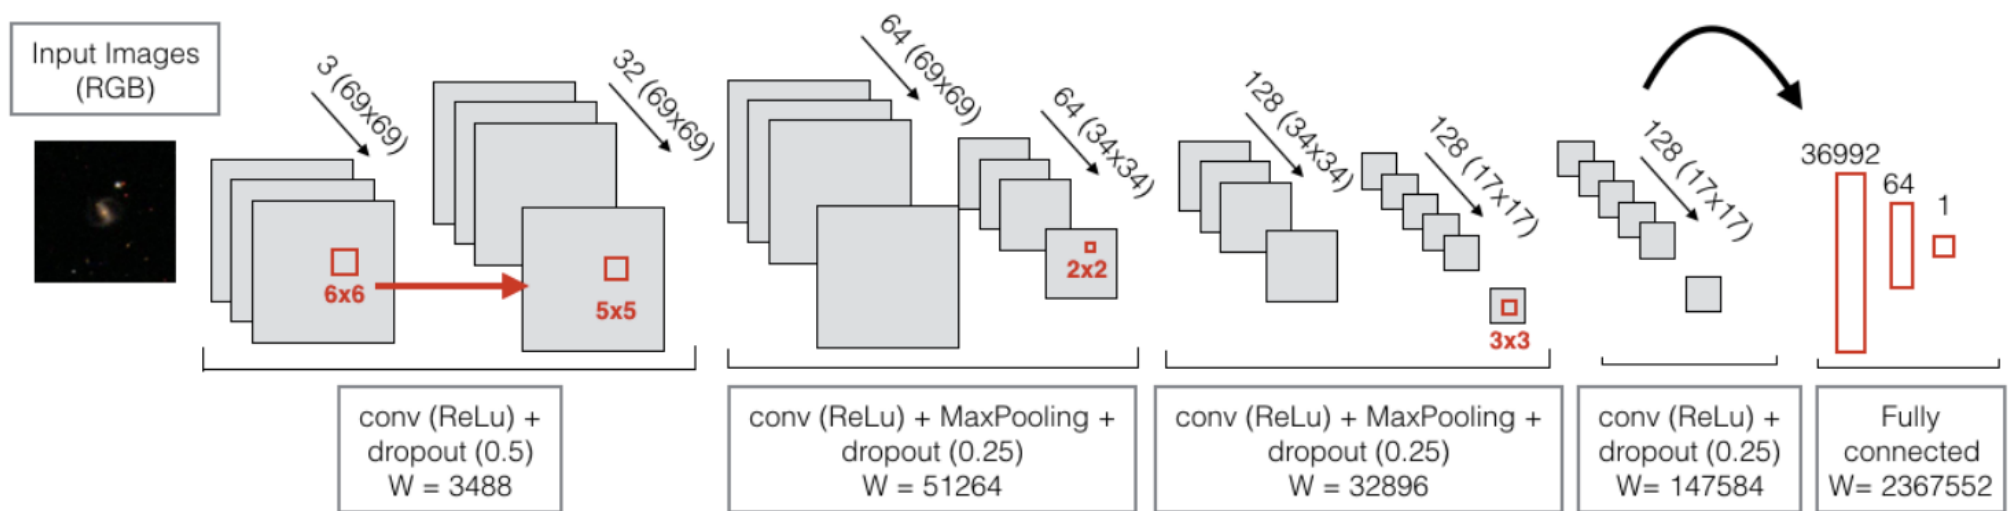


TUTORIALS

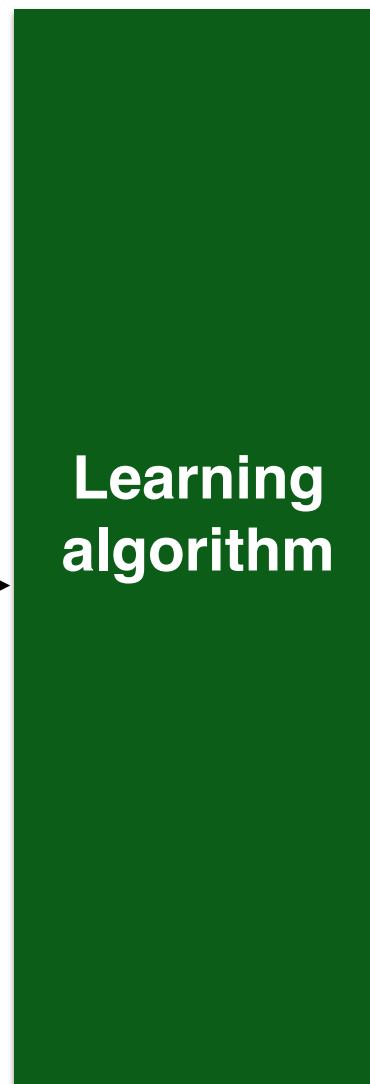
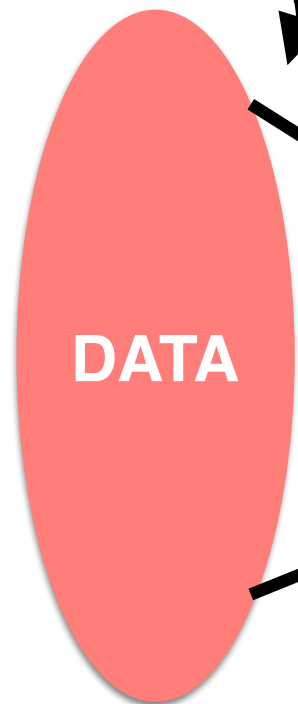
- GO HERE AND INSTALL PYTHON CONDA ENVIRONMENT. SHOULD BE VERY STRAIGHTFORWARD.
- DO NOT DOWNLOAD ANY DATA SINCE I AM CHANGING THE TUTORIALS TO SIMPLER ONES
- MORPHOLOGY CLASSIFICATION WITH / WITHOUT DEEP LEARNING
- A BIT OF VAEs...MAYBE

PART IV: IMAGE 2 IMAGE NETWORKS +
INTRODUCTION TO GENERATIVE
MODELS

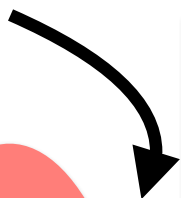
EXAMPLE OF VERY SIMPLE CNN



Raw data

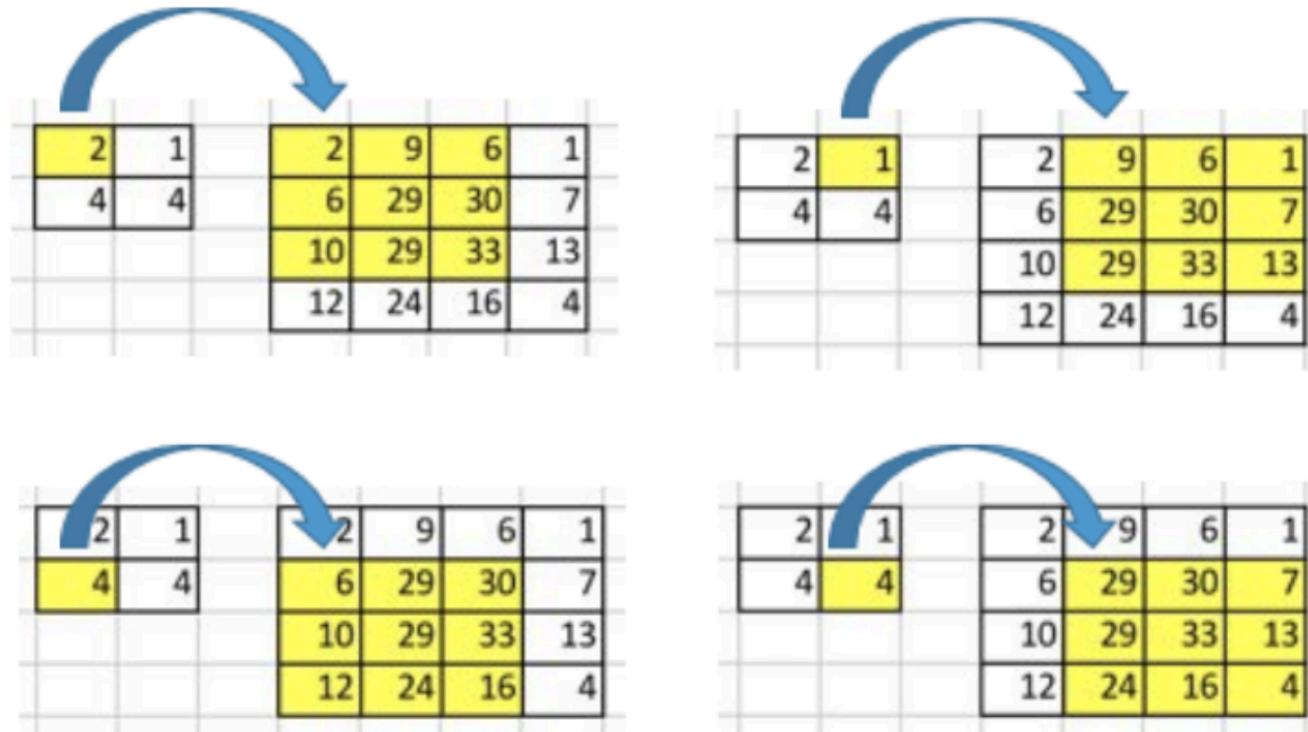


prediction



TRANSPOSED CONVOLUTION

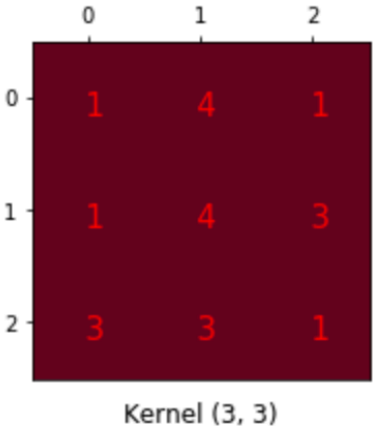
ALLOWS TO INCREASE THE SIZE



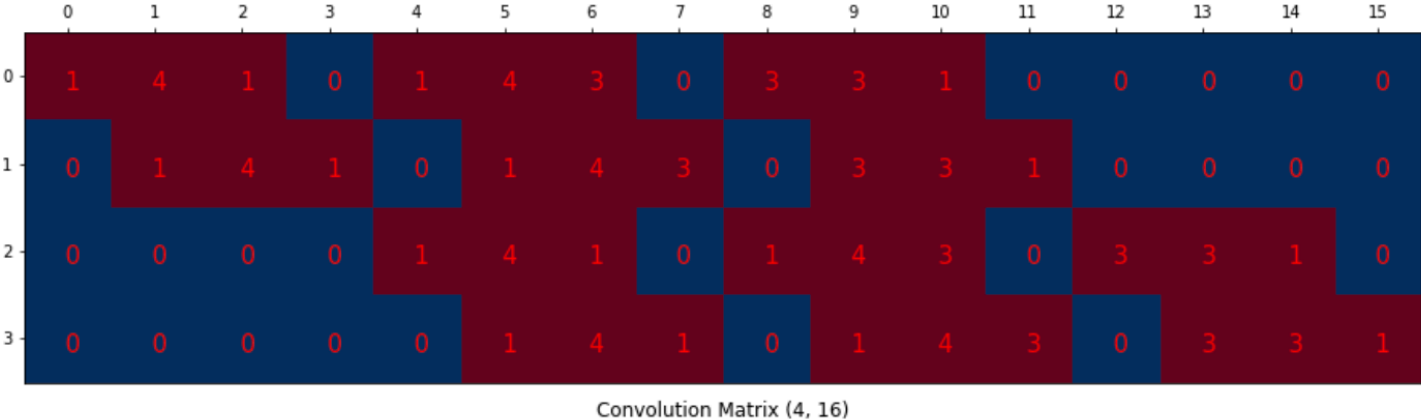
Going Backward of Convolution

EXAMPLE TAKEN FROM HERE

CONVOLUTION MATRIX

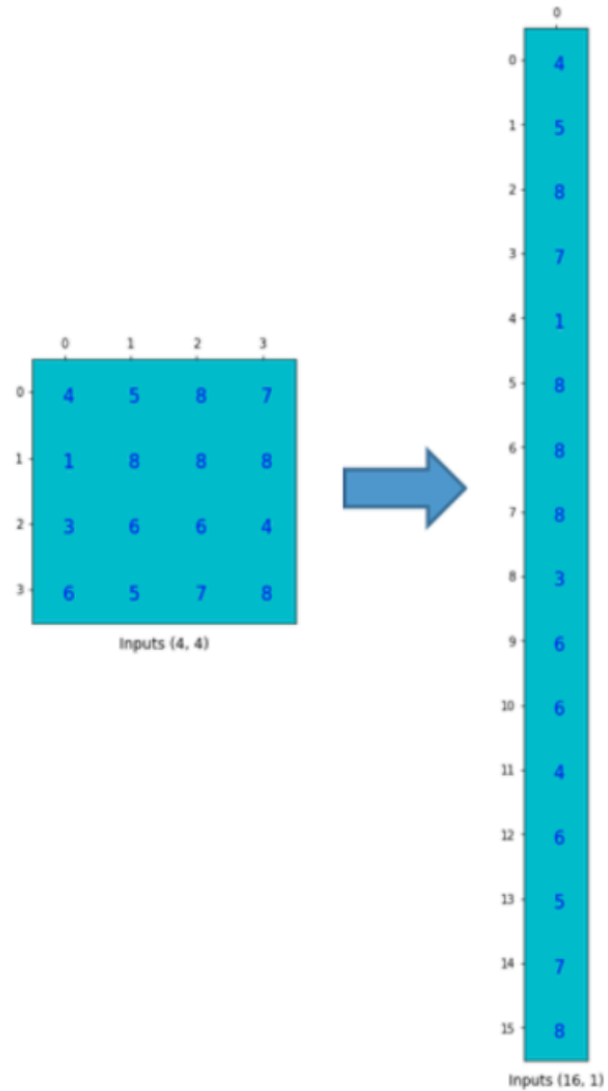


THE KERNEL CAN BE ARRANGED IN FORM OF A MATRIX:



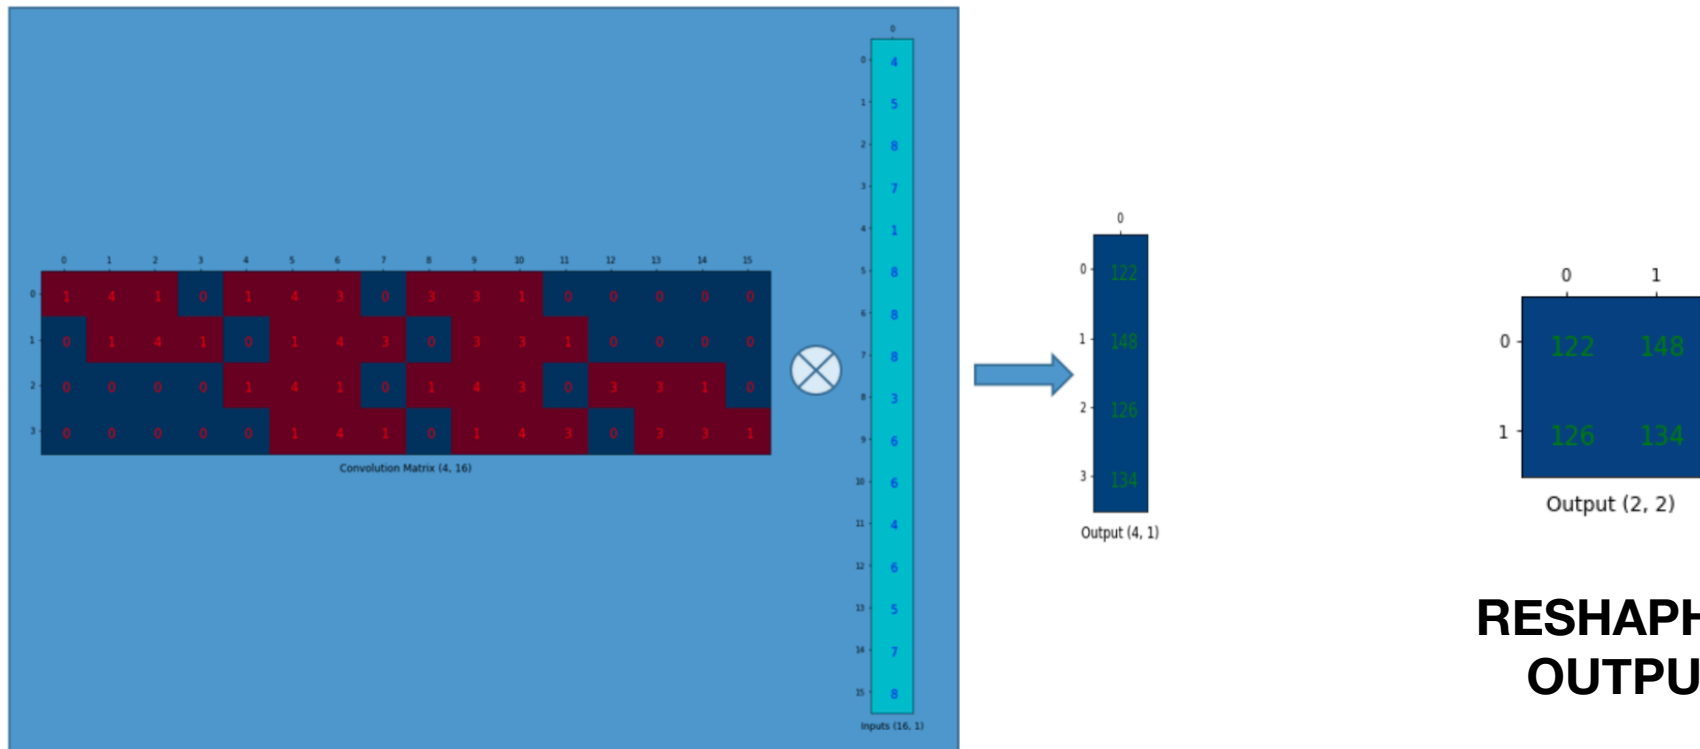
EXAMPLE TAKEN FROM HERE

THE INPUT IS FLATTENED INTO A COLUMN VECTOR



EXAMPLE TAKEN FROM HERE

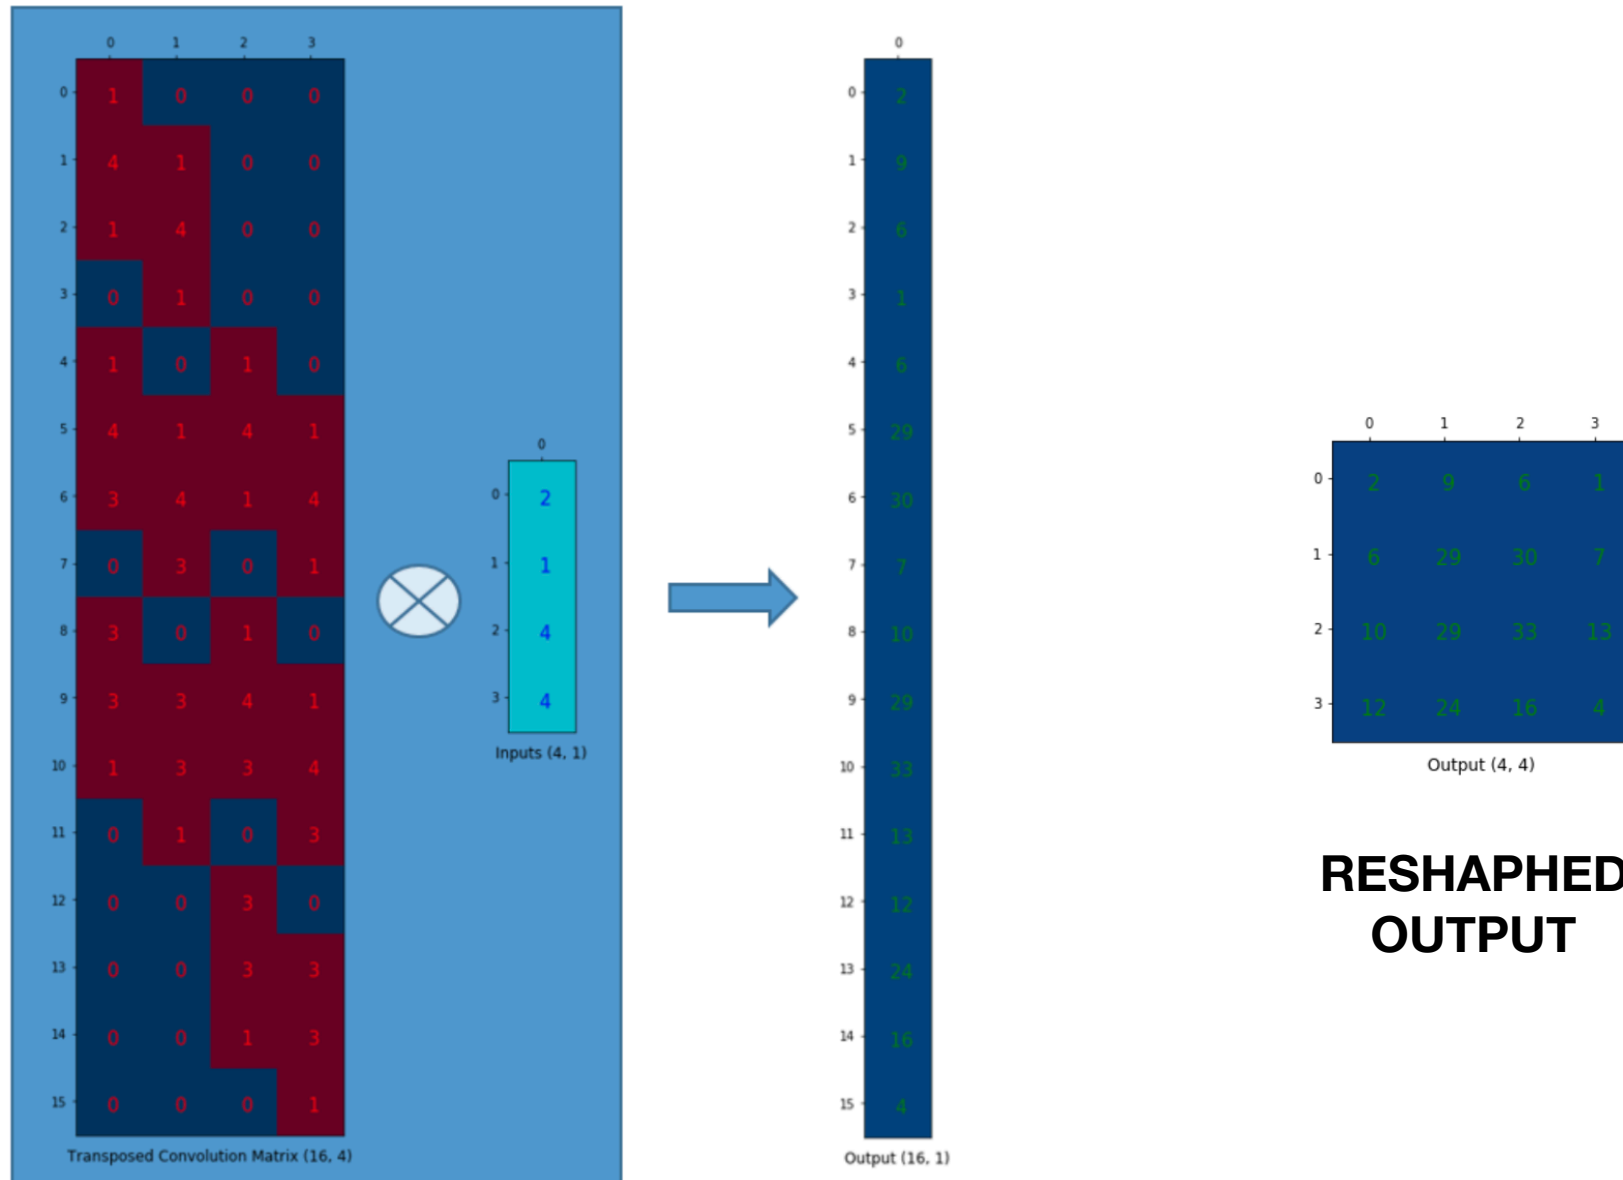
THE CONVOLUTION IS TRANSFORMED INTO A PRODUCT OF MATRICES



**RESHAPED
OUTPUT**

EXAMPLE TAKEN FROM HERE

THE TRANSPOSED CONVOLUTION IS THE INVERSE OPERATION



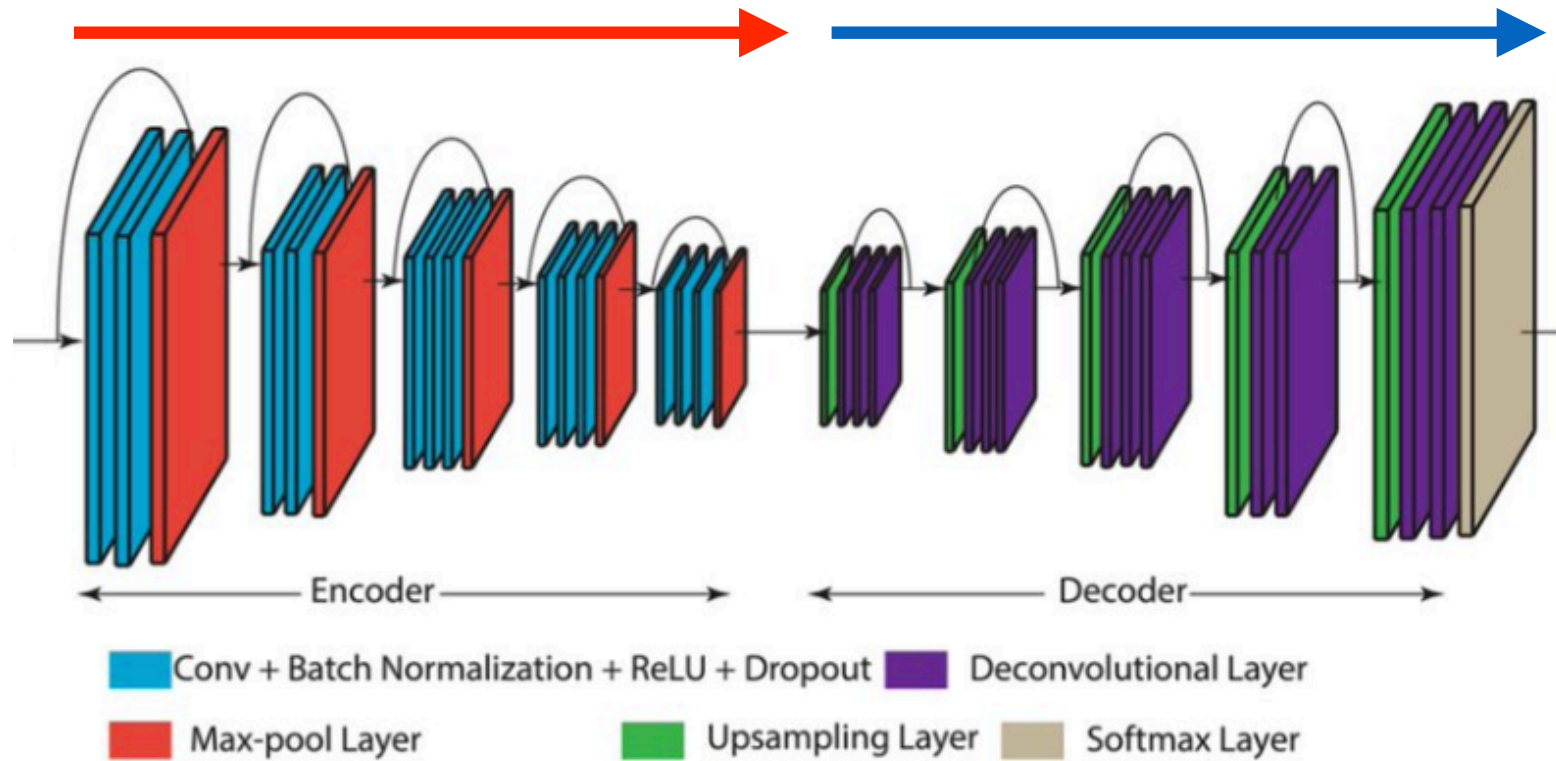
**RESHAPED
OUTPUT**

EXAMPLE TAKEN FROM HERE

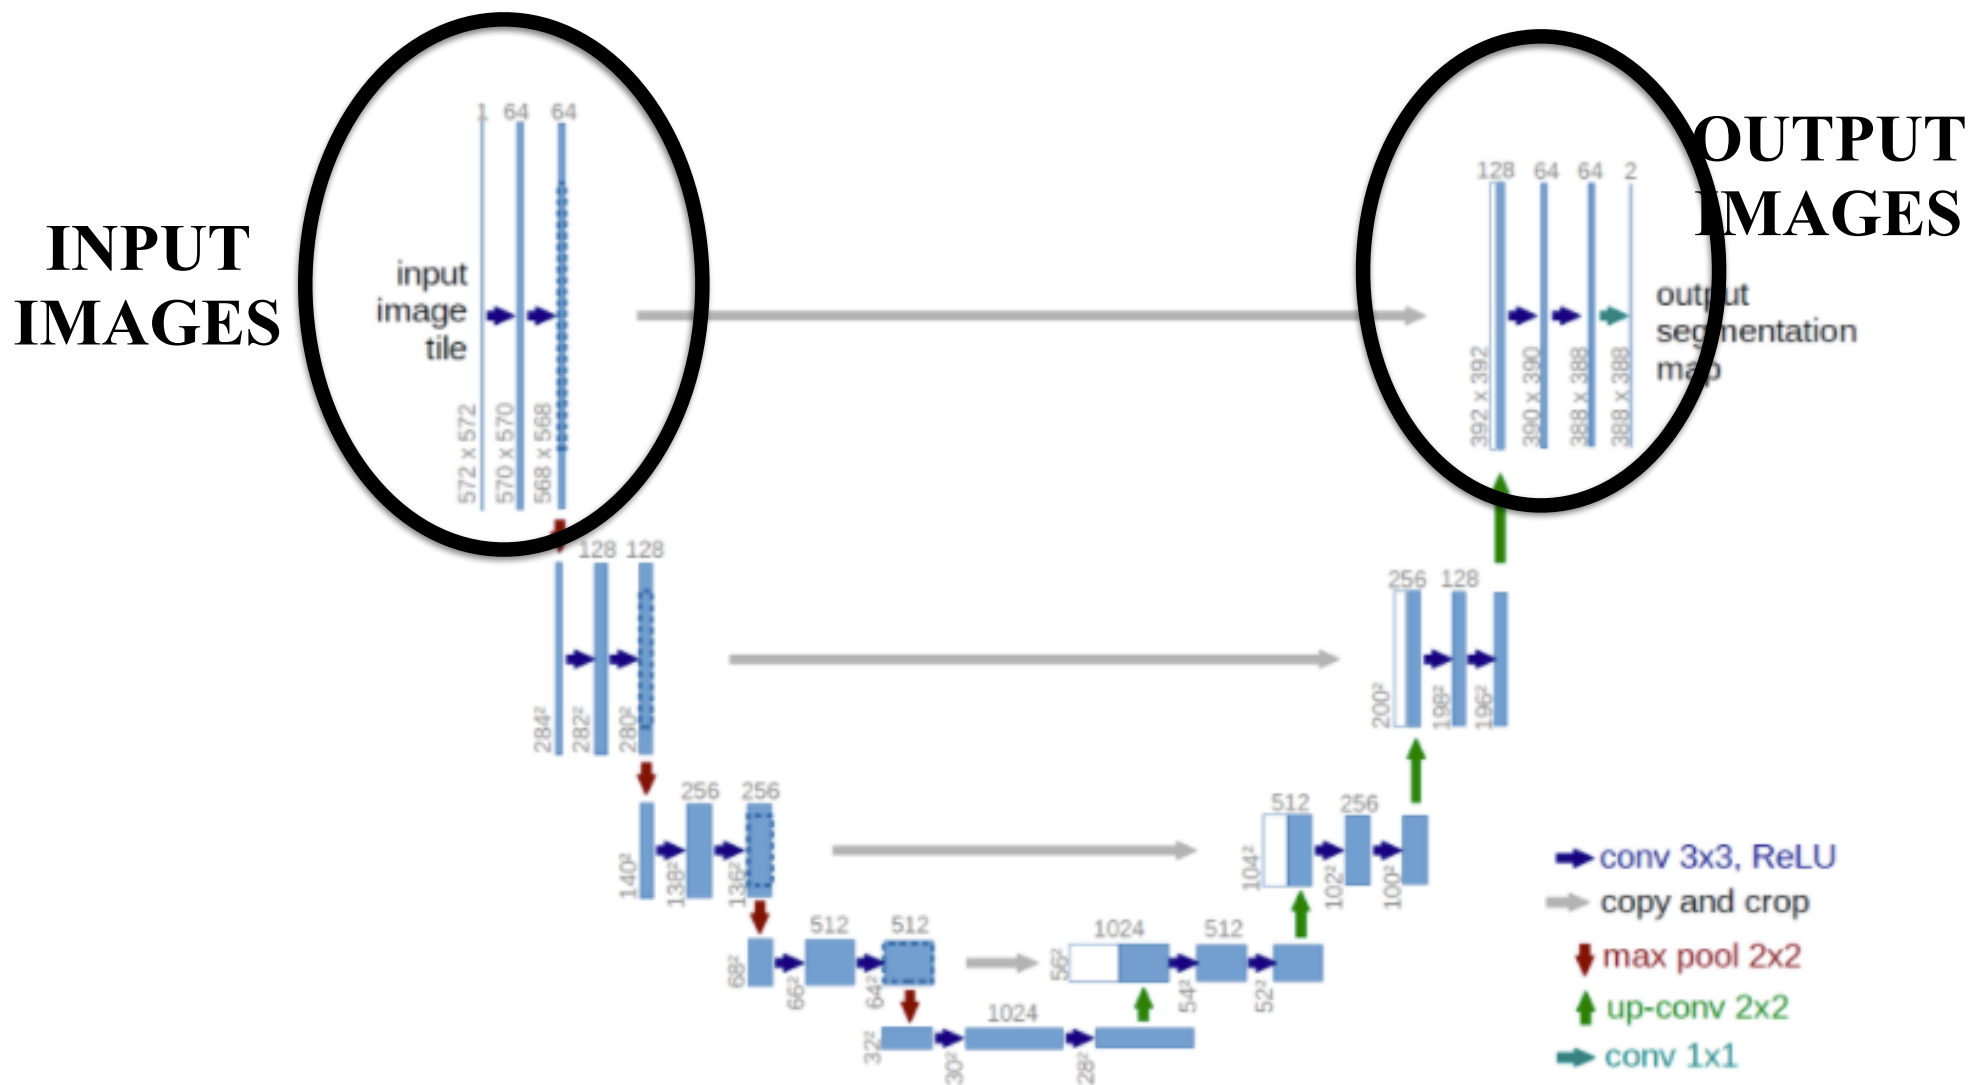
DECODER-ENCODER NETWORKS

ENCODING

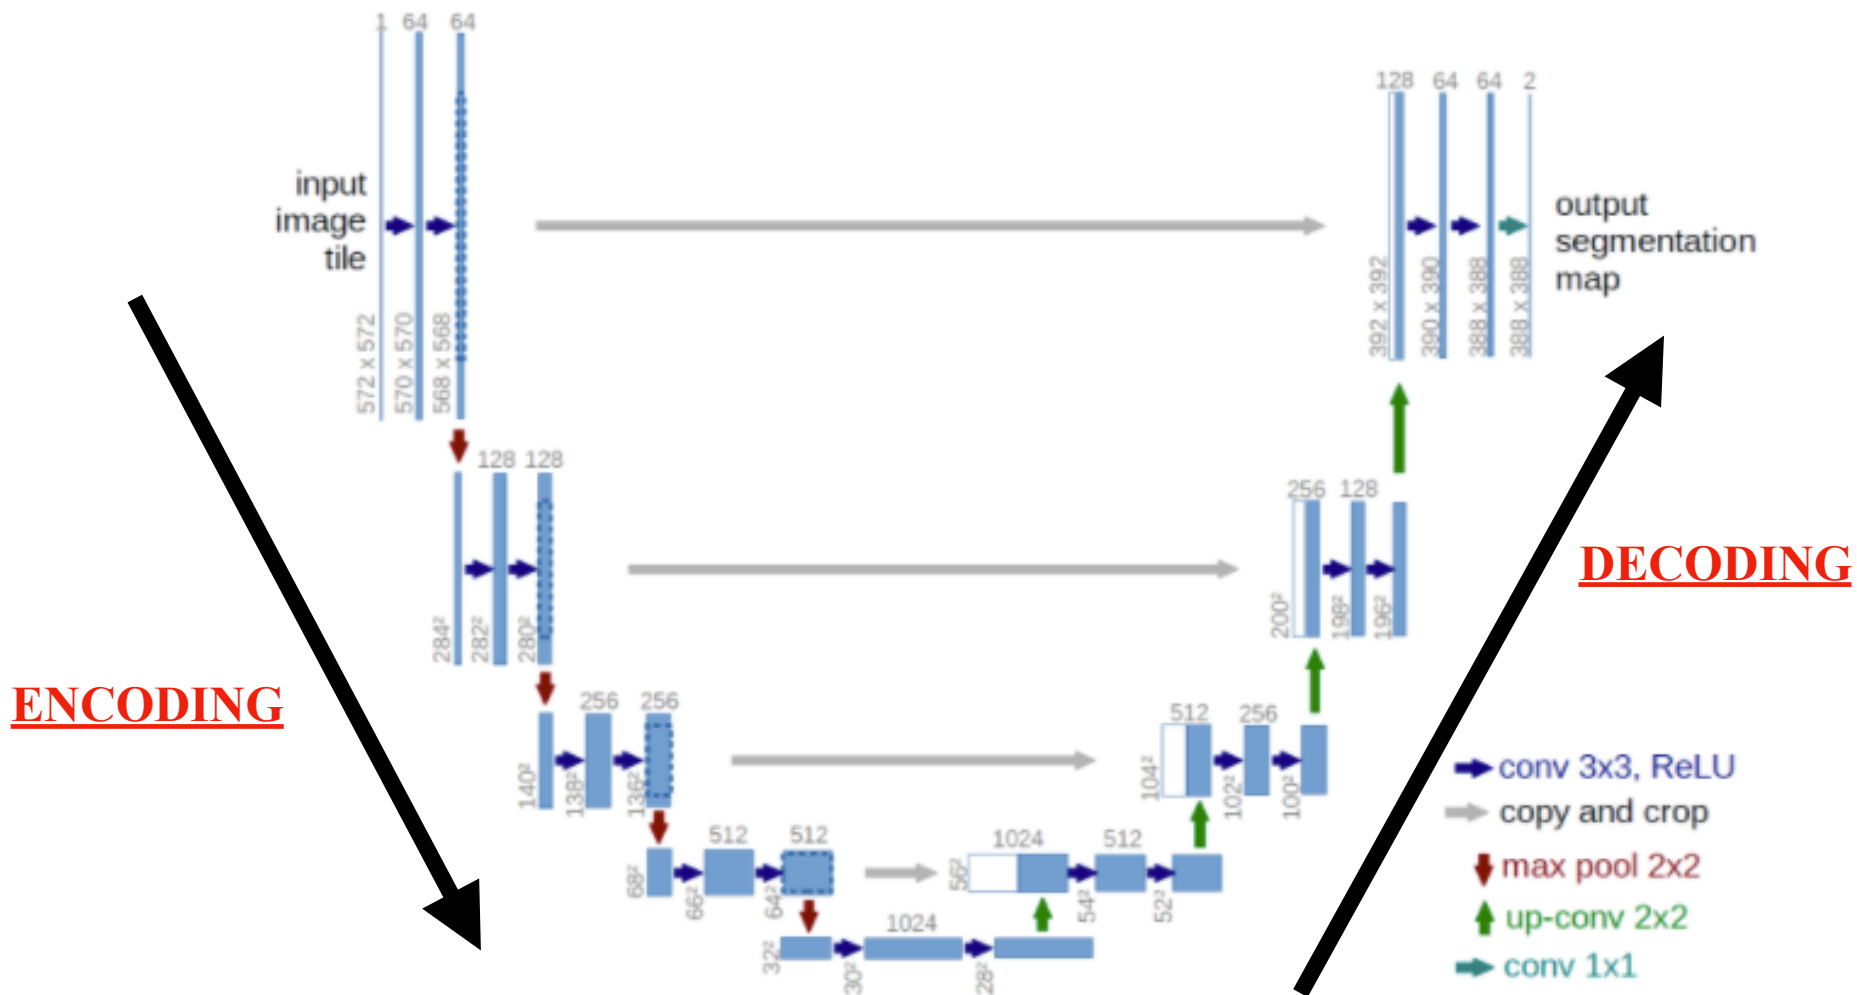
DECODING



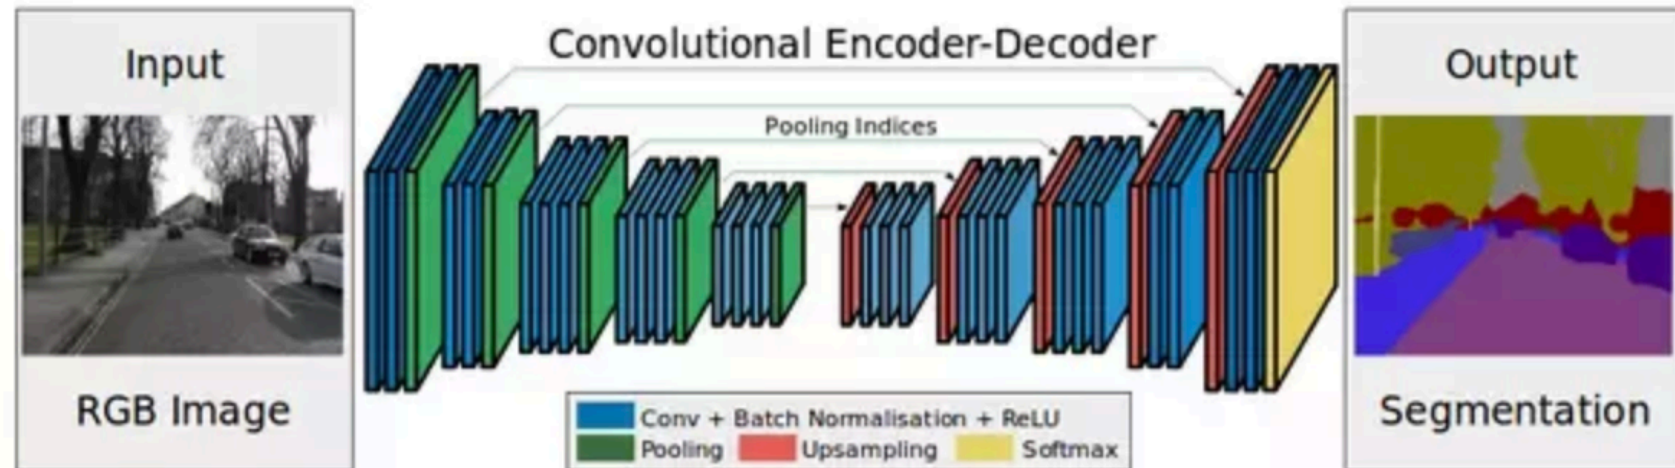
ENCODING-DECODING TO EXTRACT IMAGE FEATURES: U-NET



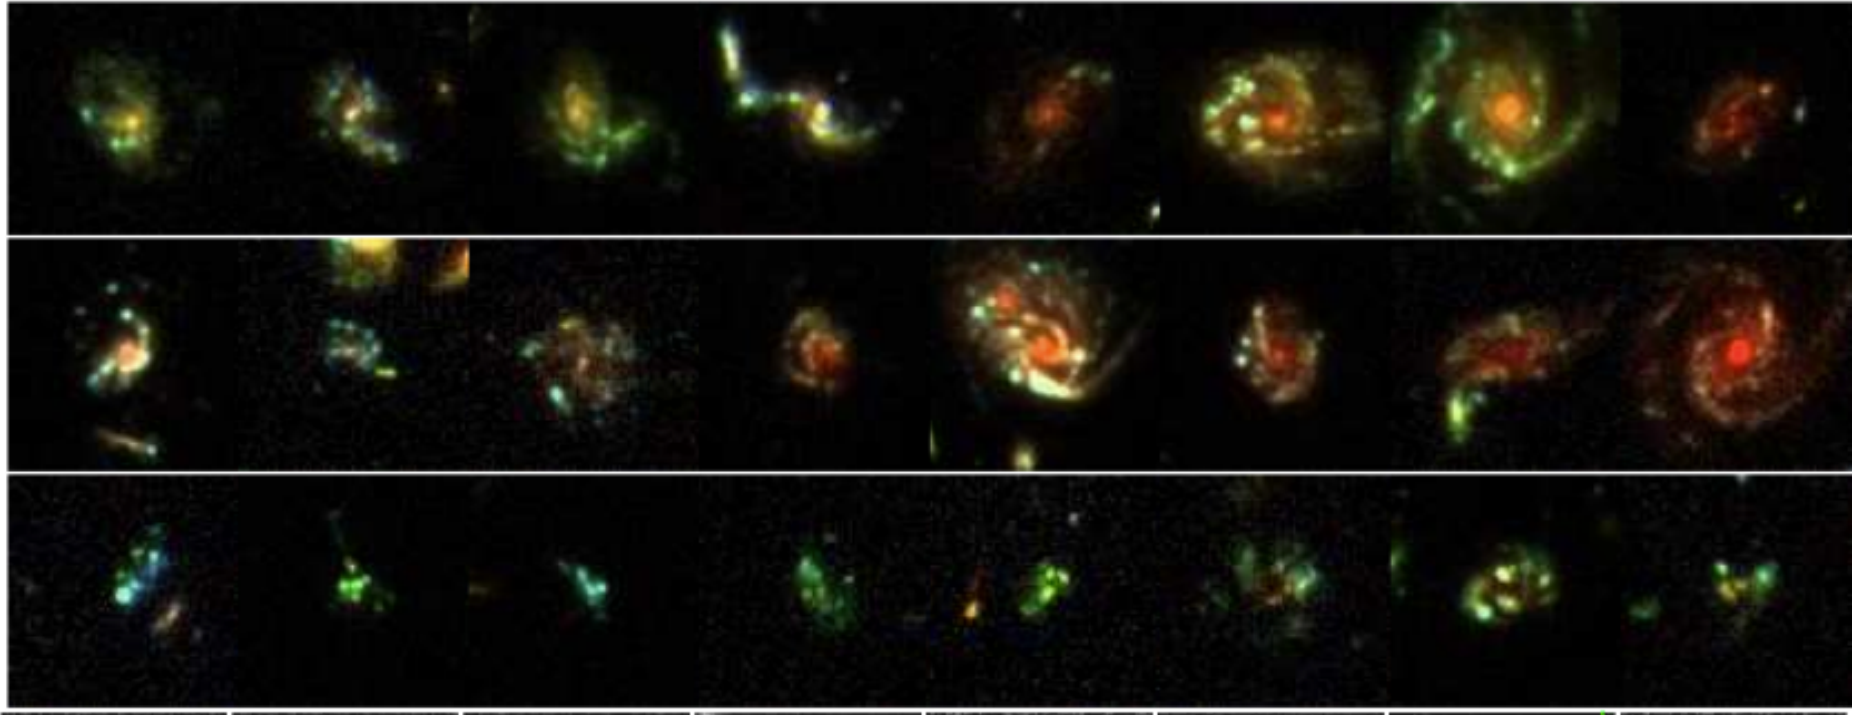
ENCODING-DECODING TO EXTRACT IMAGE FEATURES: U-NET



MAIN APPLICATION IS IMAGE SEGMENTATION



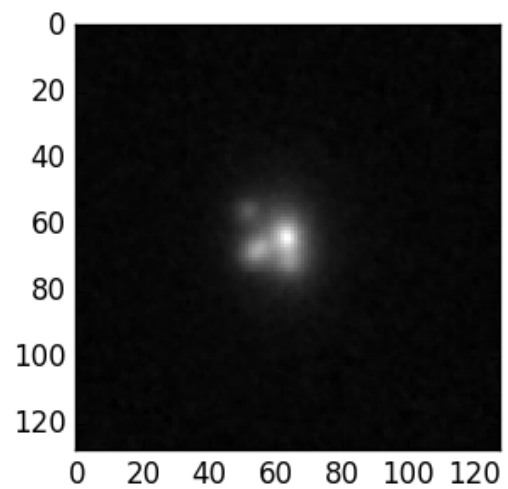
CLUMP DETECTION



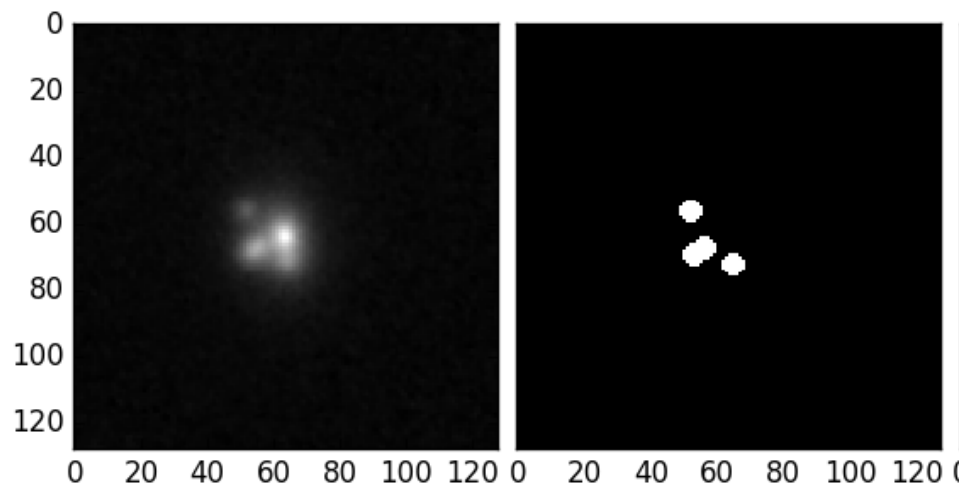
Guo+15,18

HIGH REDSHIFT GALAXIES PRESENT CLUMPS - THEIR
ROLE IN BULGE FORMATION IS DEBATED

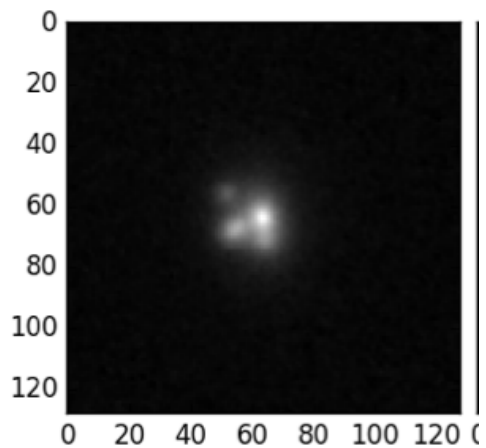
**VERY SIMPLE SERSIC
ANALYTIC
SIMULATIONS
+ UNRESOLVED
CLUMPS**



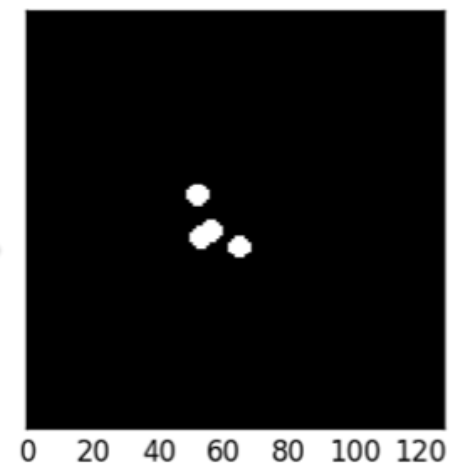
**VERY SIMPLE SERSIC
ANALYTIC
SIMULATIONS
+ UNRESOLVED
CLUMPS**



**CLUMP
POSITION IS
KNOWN**



U-NET



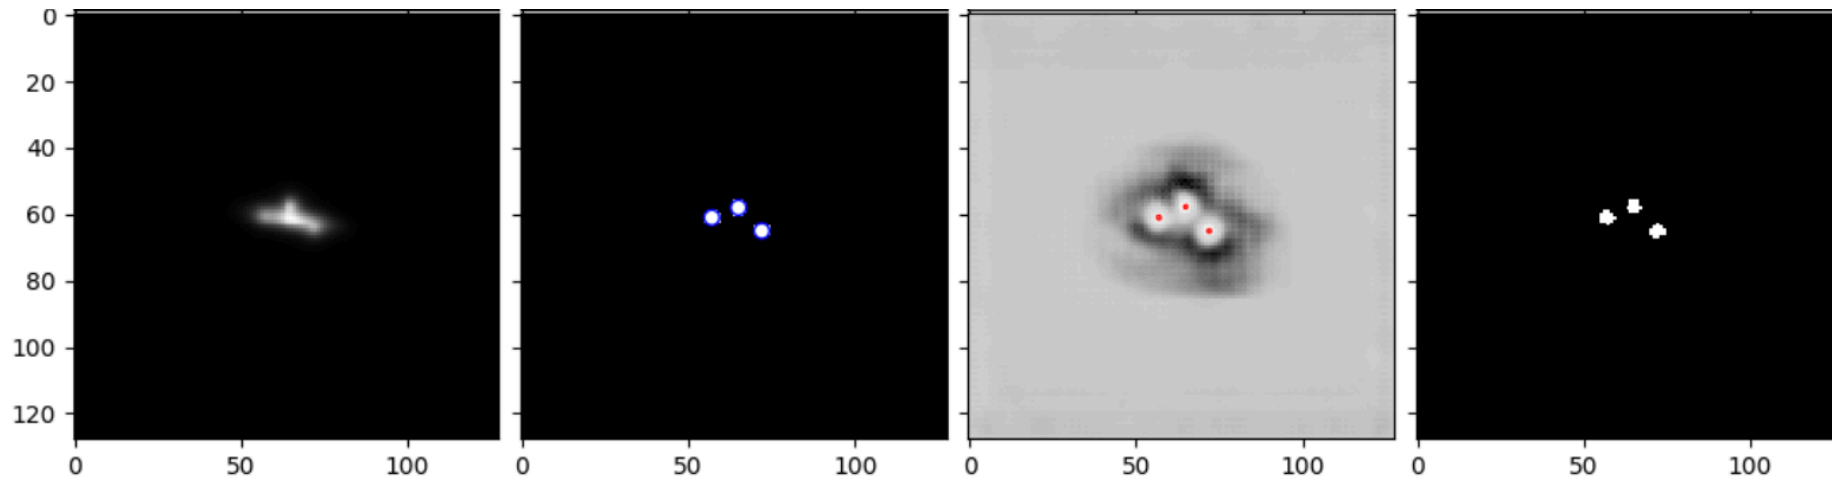
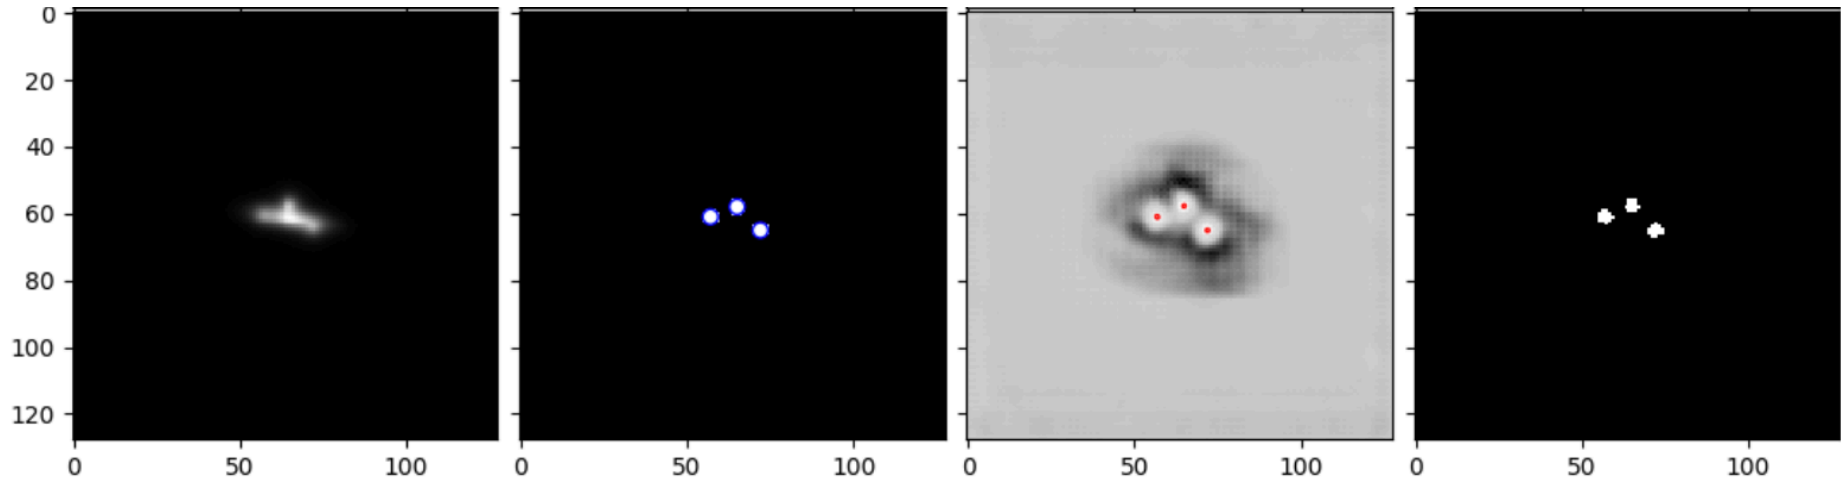
LEE, MHC, PRIMACK, GUO+

**SIMULATED
GALAXY**

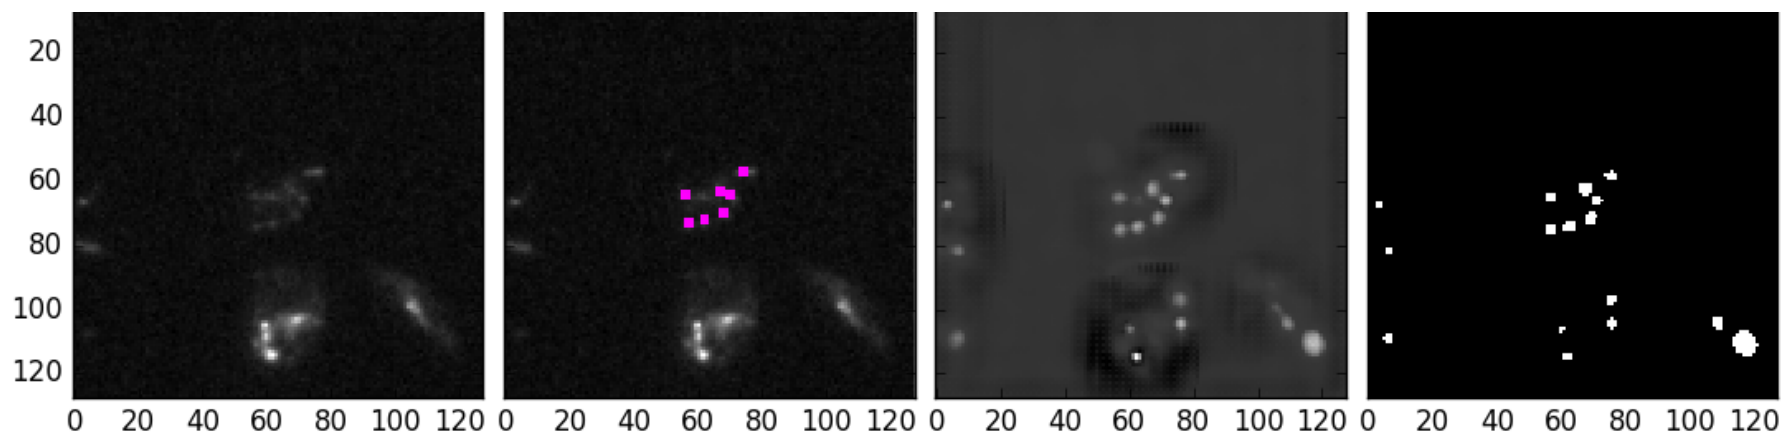
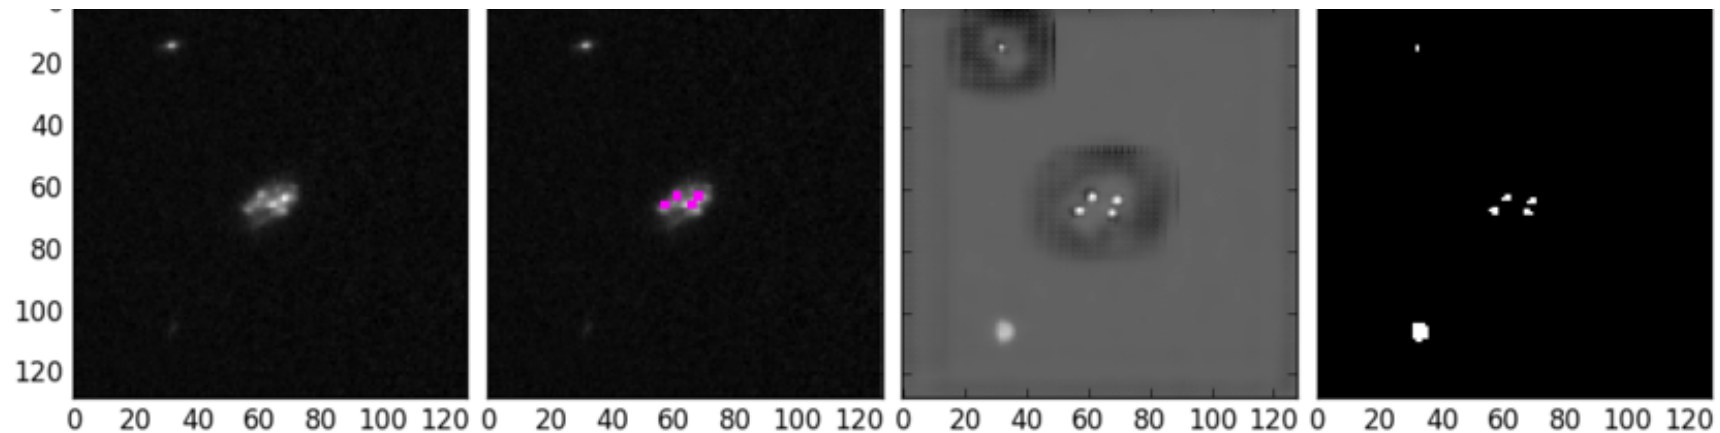
**CLUMP
MASK**

**NETWORK
PROBA
MAP**

**NETWORK
OUTPUT
THRESHOLDED
(SEXTRACTOR)**



SEEMS TO WORK REASONABLY WELL ON REAL OBSERVATIONS

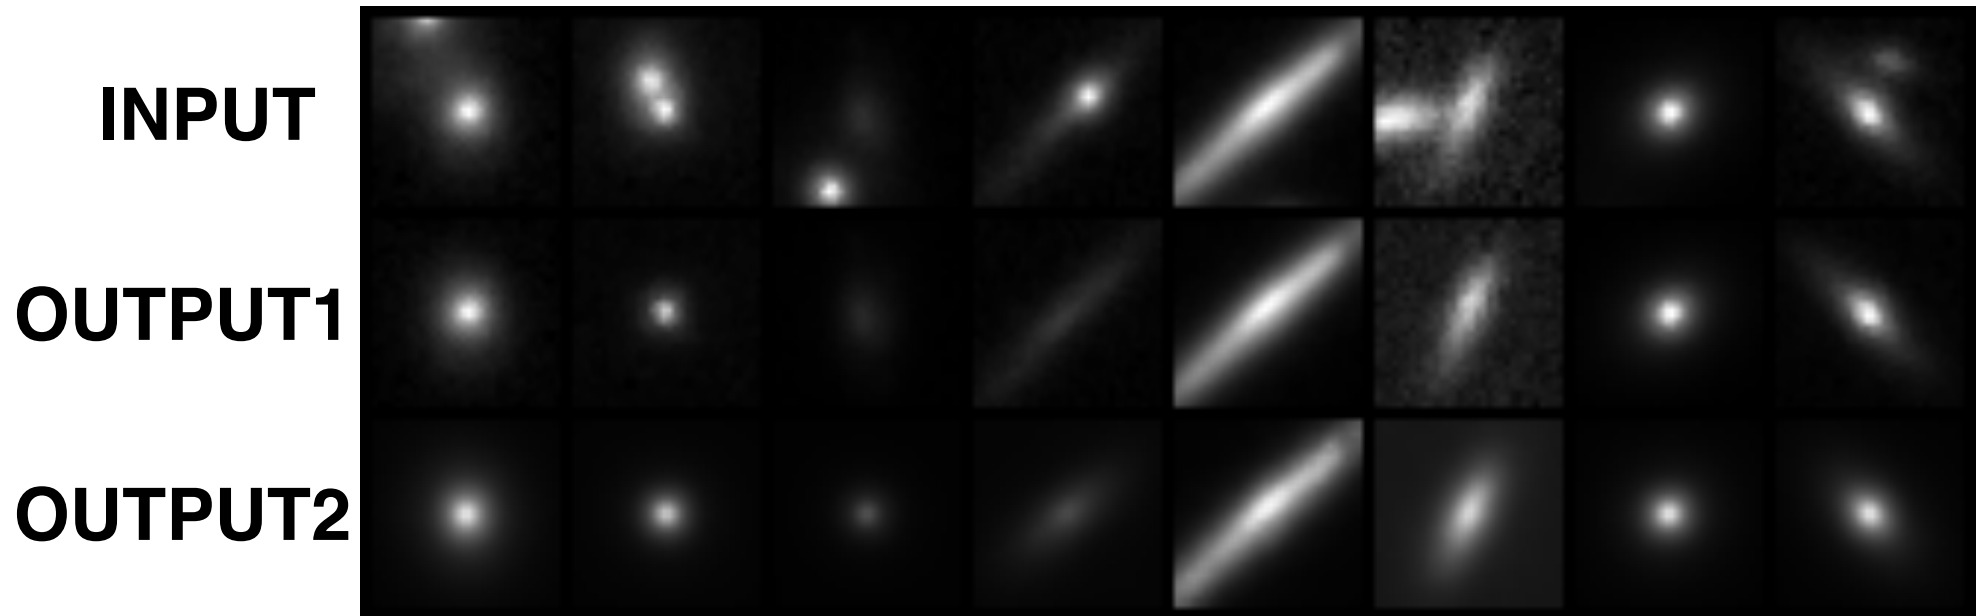


**OBSERVED
GALAXY**

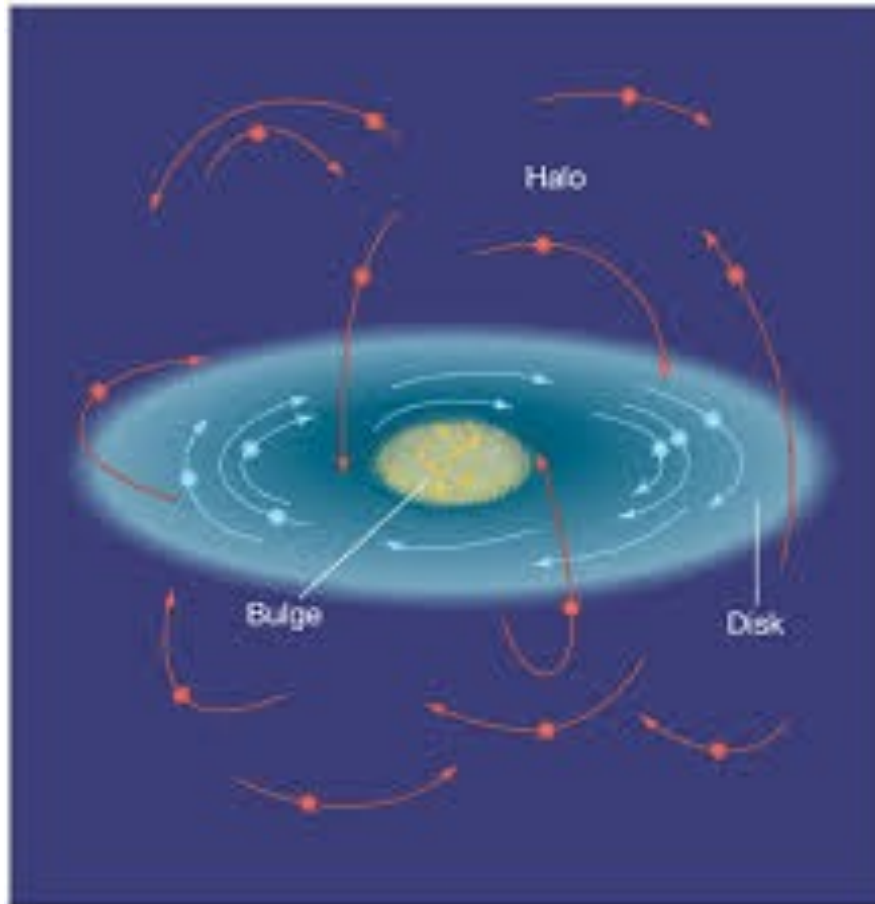
**CLUMPS
DETECTED BY
GUO+**

U-NET OUTPUTS

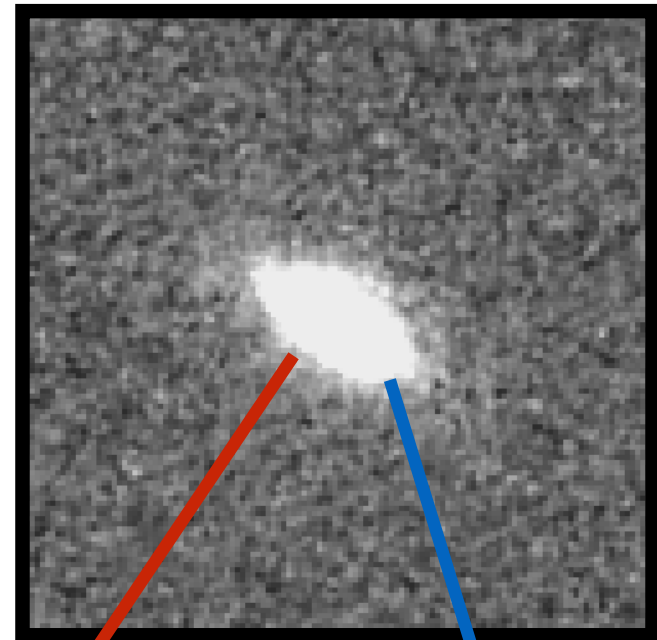
U-NET FOR GALAXY DEBLENDING



GALAXIES HAVE TWO COMPONENTS WITH DIFFERENT ASSEMBLY HISTORIES



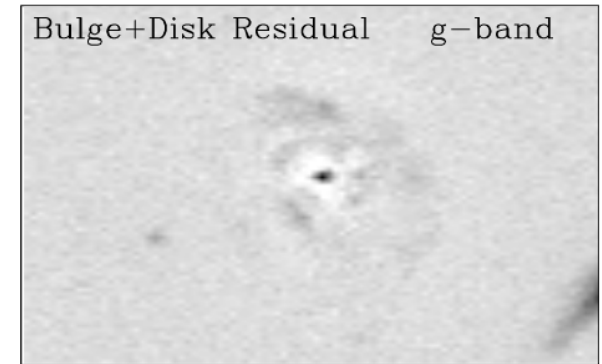
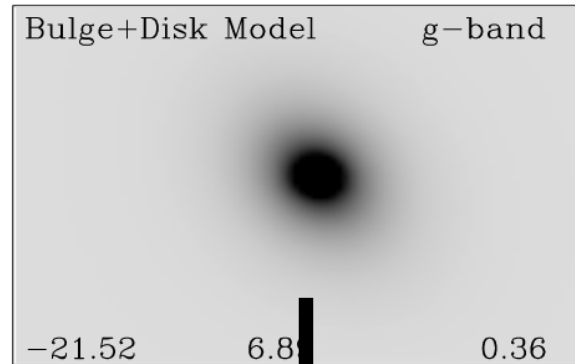
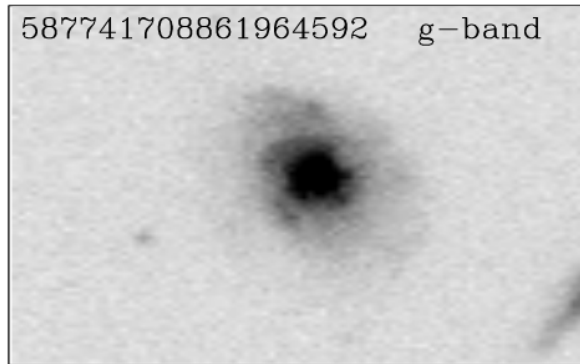
OBSERVED GALAXY



BULGE?

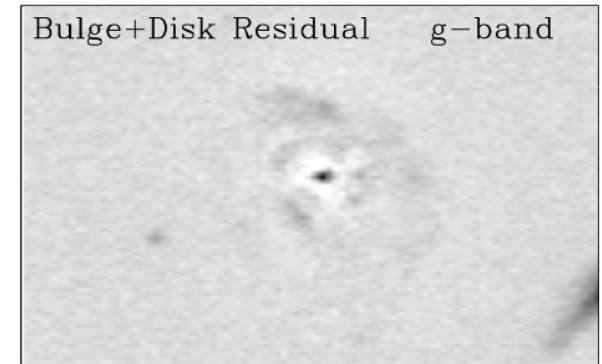
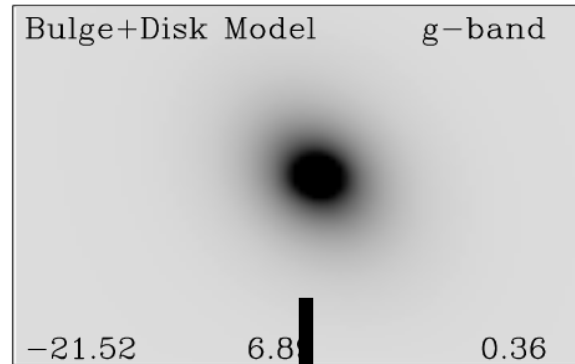
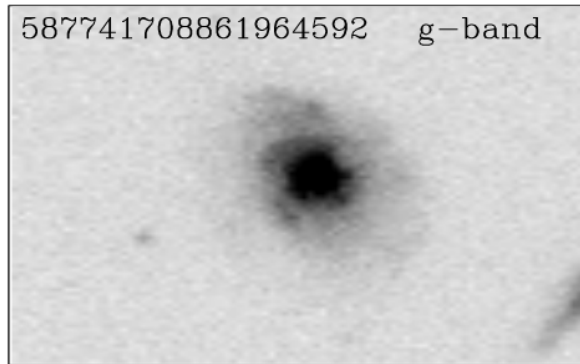
DISK?

STANDARD MODEL FITTING



[10 parameter model]

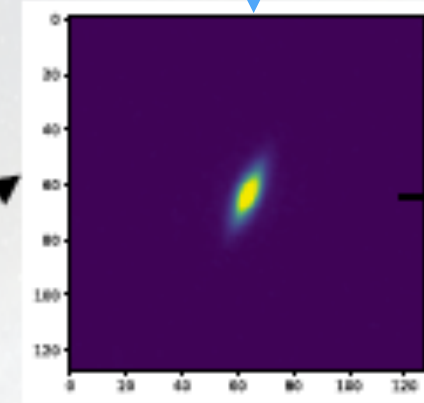
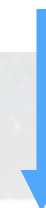
STANDARD MODEL FITTING



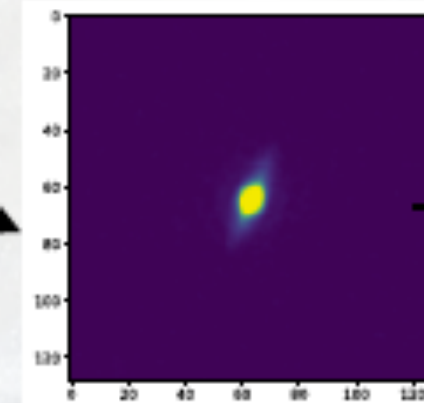
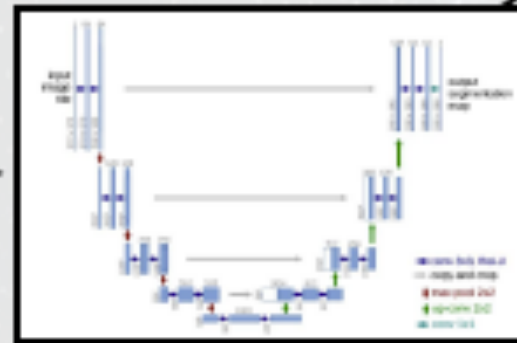
[10 parameter model]

LARGE AMOUNT OF DEGENERACIES

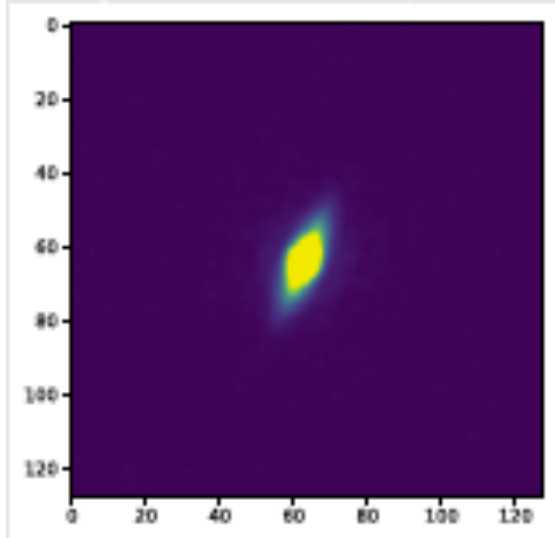
DISK



DeepLegato_2c

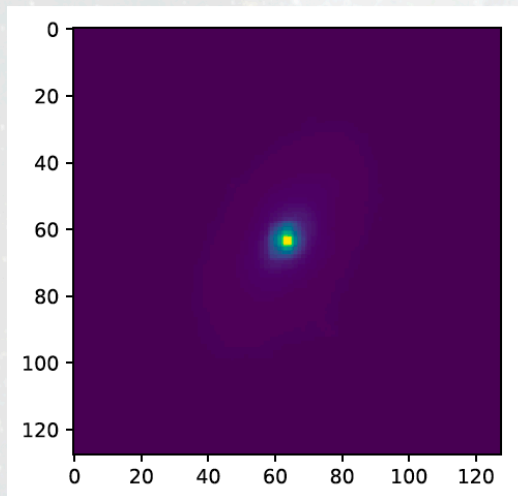
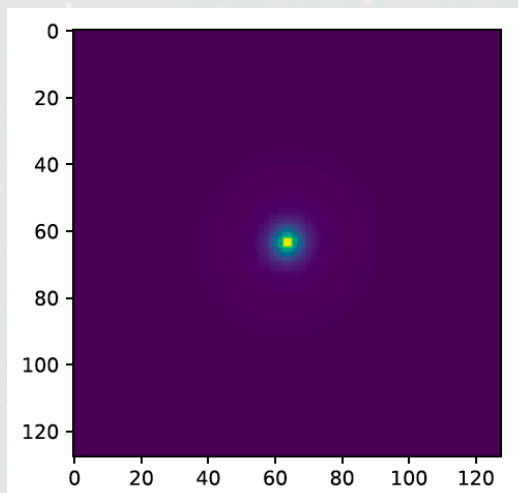


BULGE

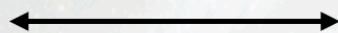


BULGE+DISK

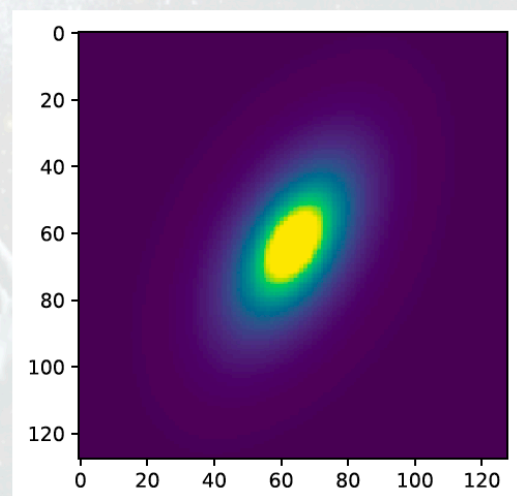
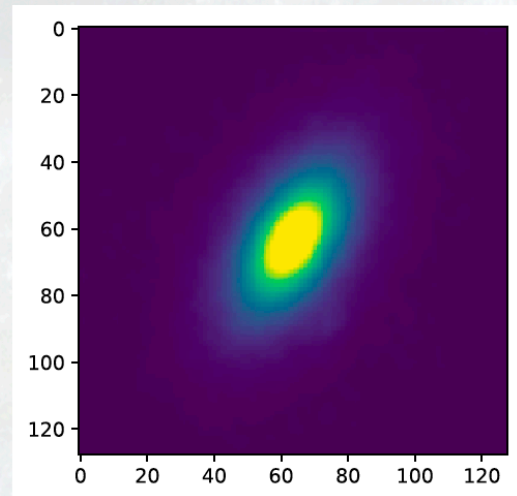
best Bulge 9



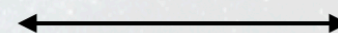
Predicted



best Disk 9

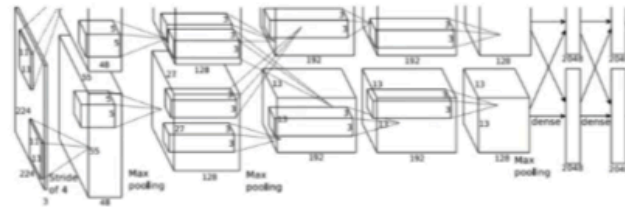


Validation

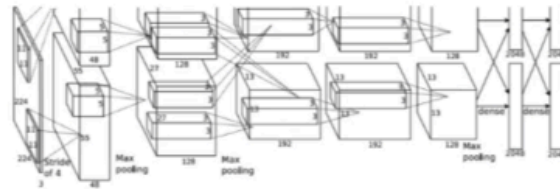


A WORKD ON R-CNNs..

JUST FOR YOUR RECORDS...



CAT: (x, y, w, h)



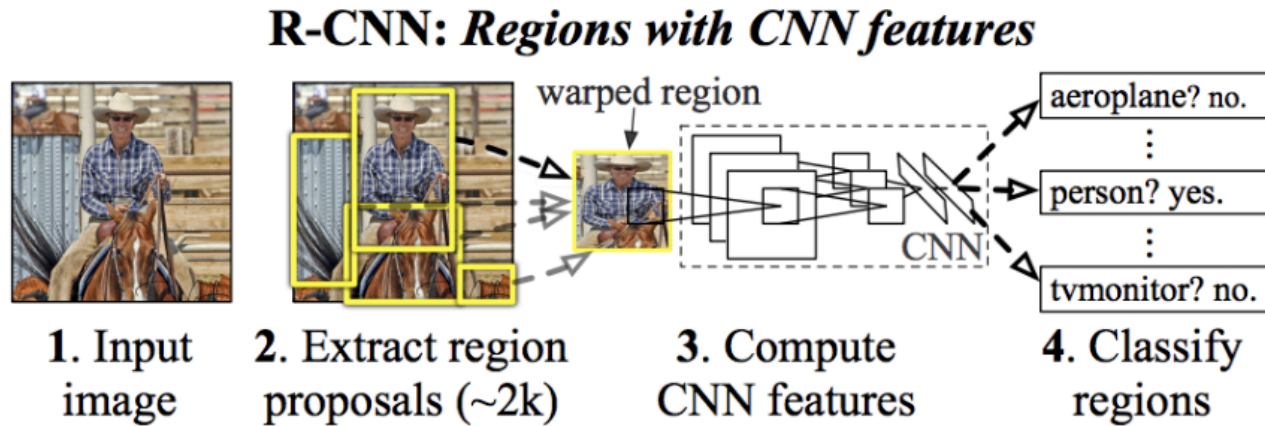
DUCK: (x, y, w, h)

DUCK: (x, y, w, h)

....

A WORKD ON R-CNNs..

JUST FOR YOUR RECORDS...



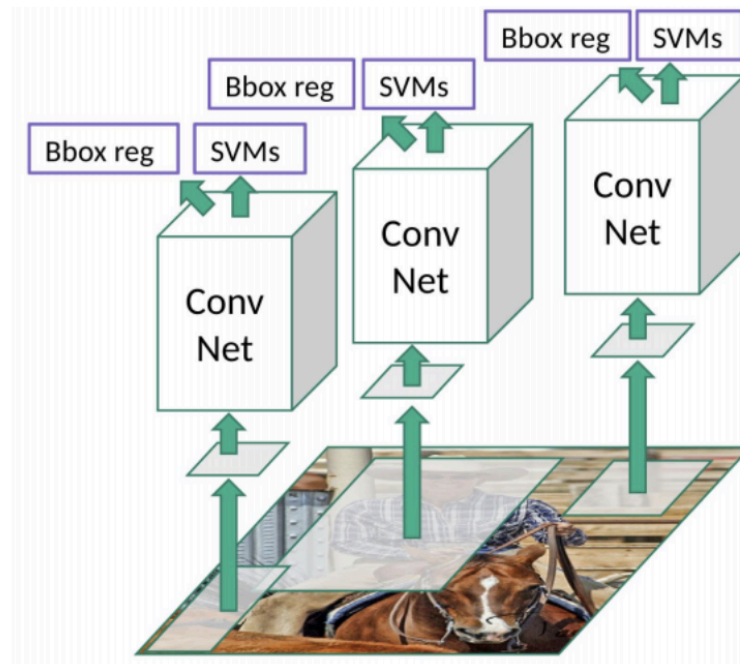
R-CNN

A WORKD ON R-CNNs..

JUST FOR YOUR RECORDS...



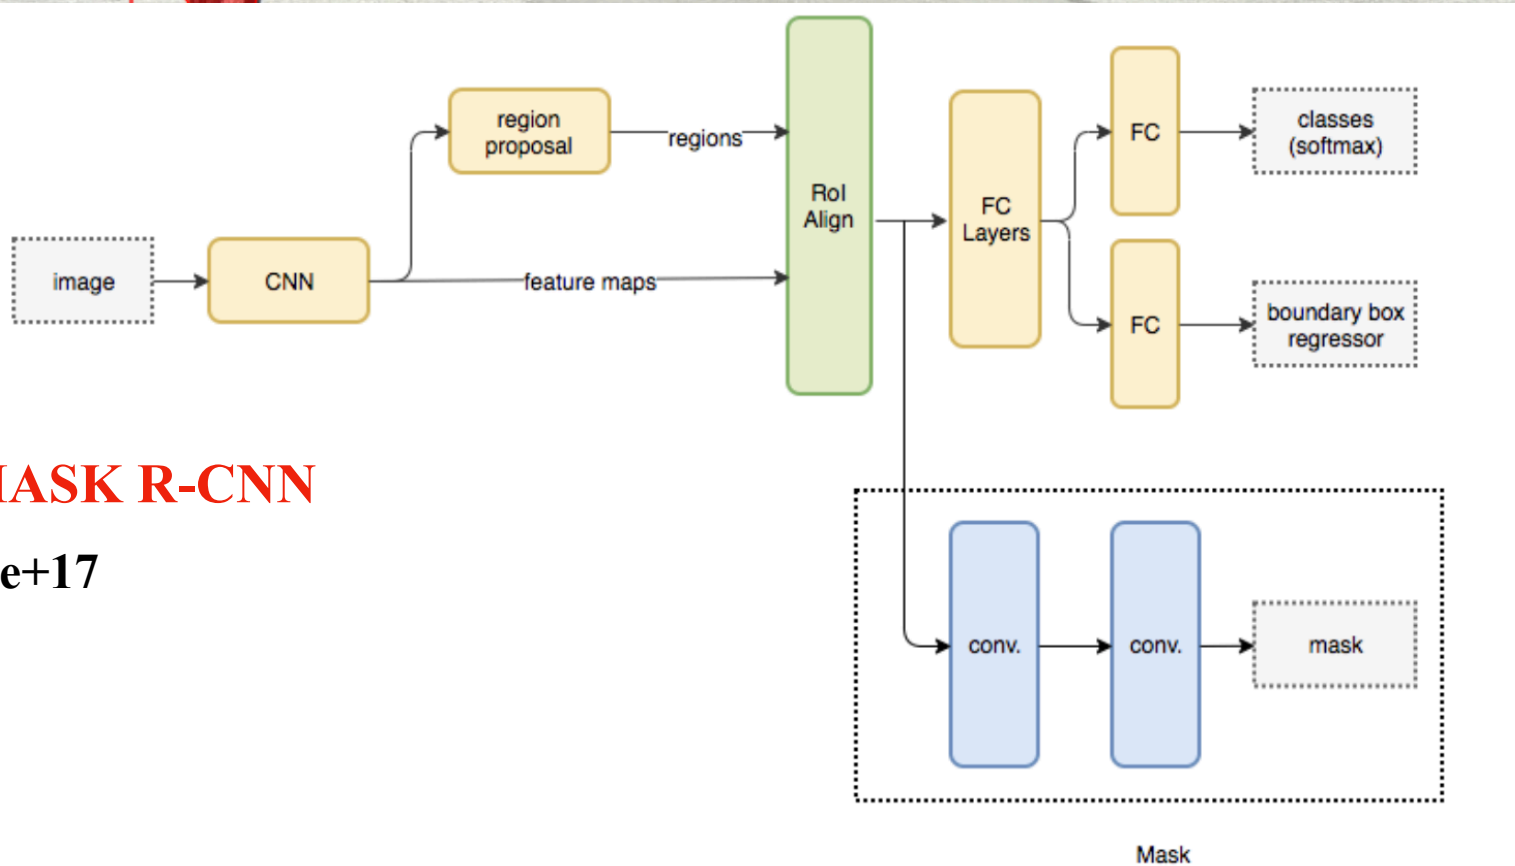
1. Input image



R-CNN

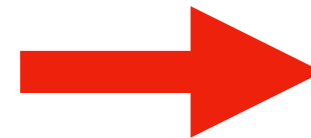
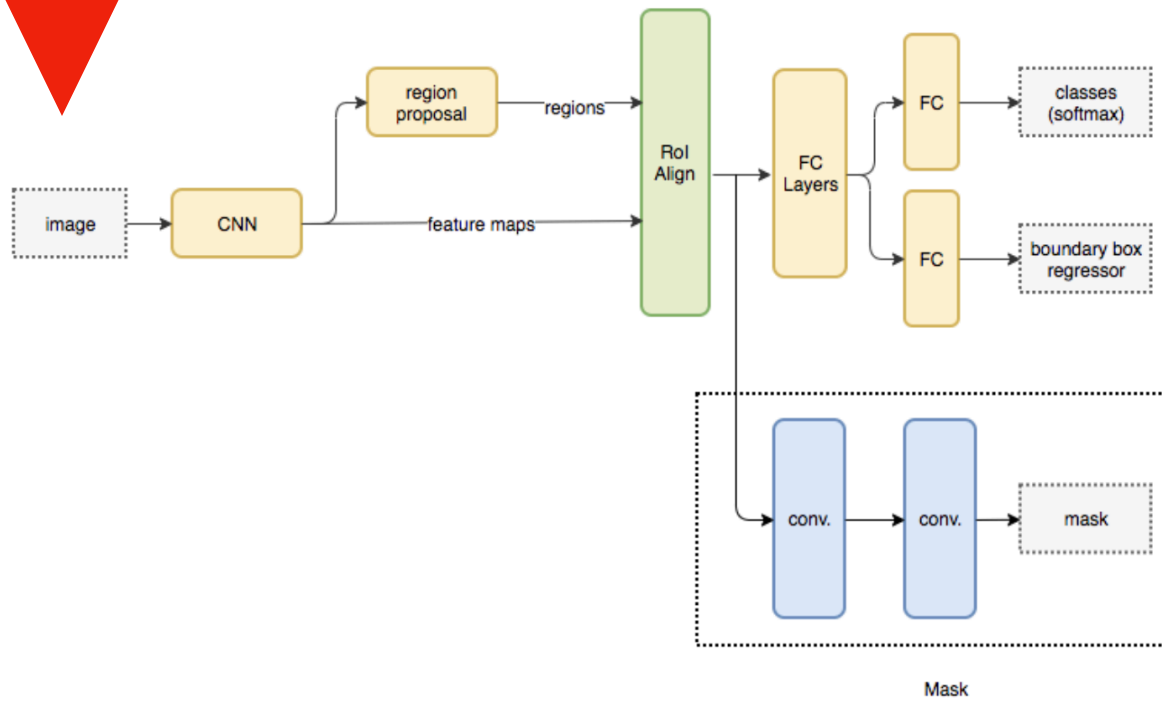
eroplane? no.
:
erson? yes.
:
vmonitor? no.
. Classify
regions

SIMULTANEOUS DETECTION + SEGMENTATION + CLASSIFICATION



MASK R-CNN

He+17

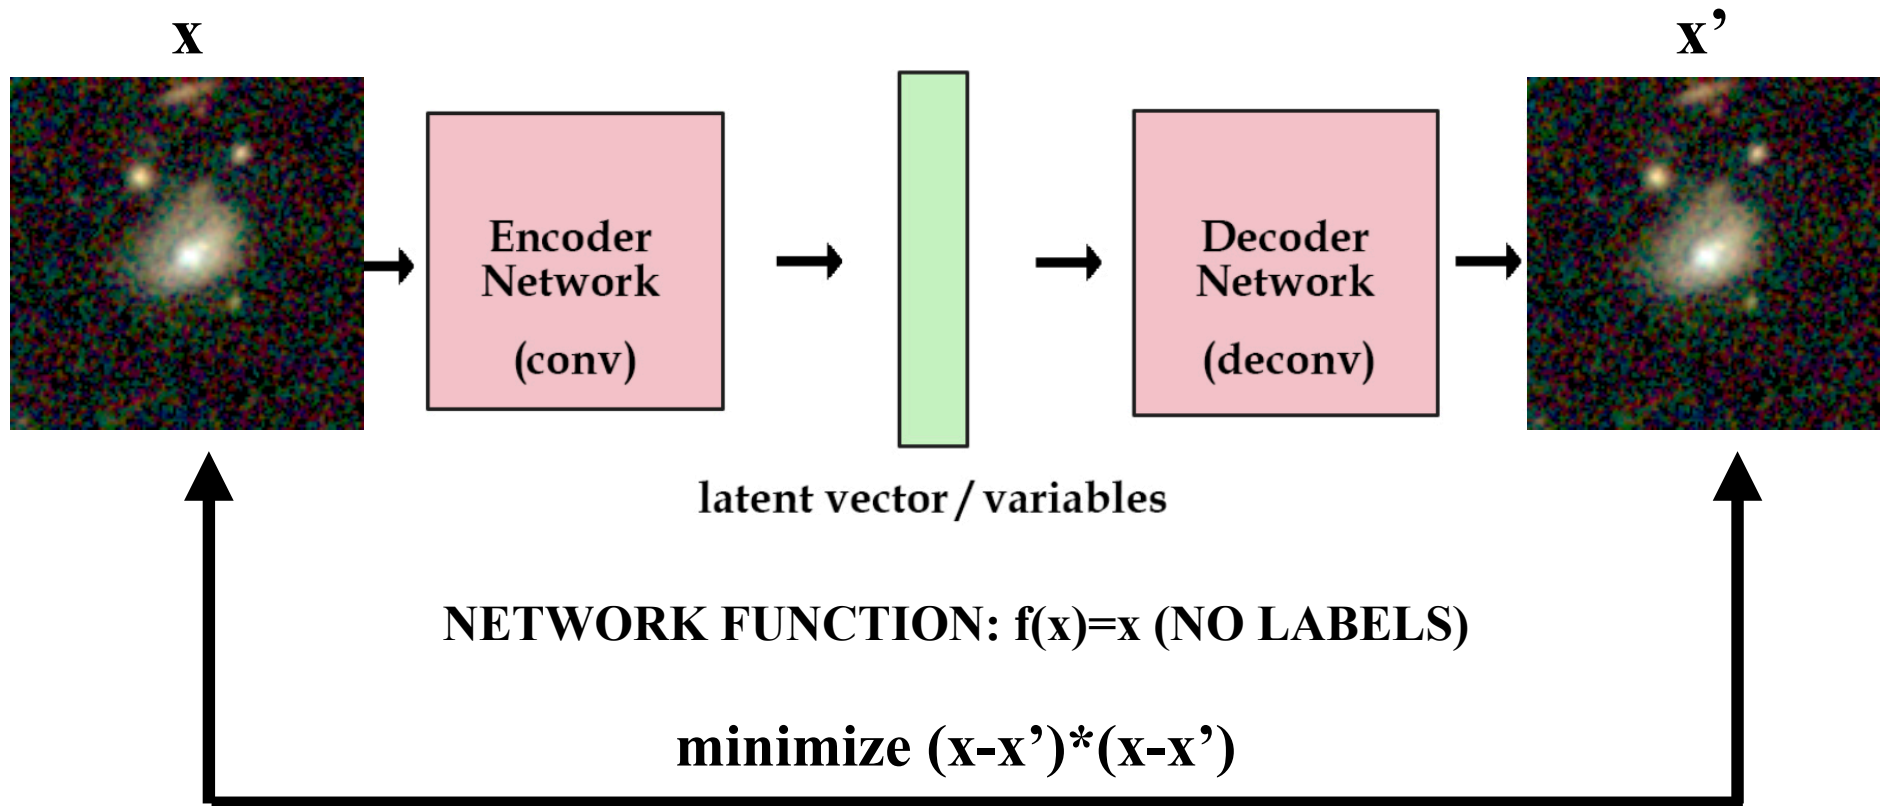


- SEG-MAP**
- PHOTOMETRY**
- MORPHOLOGY**
- PHOTO-Z**
- ...

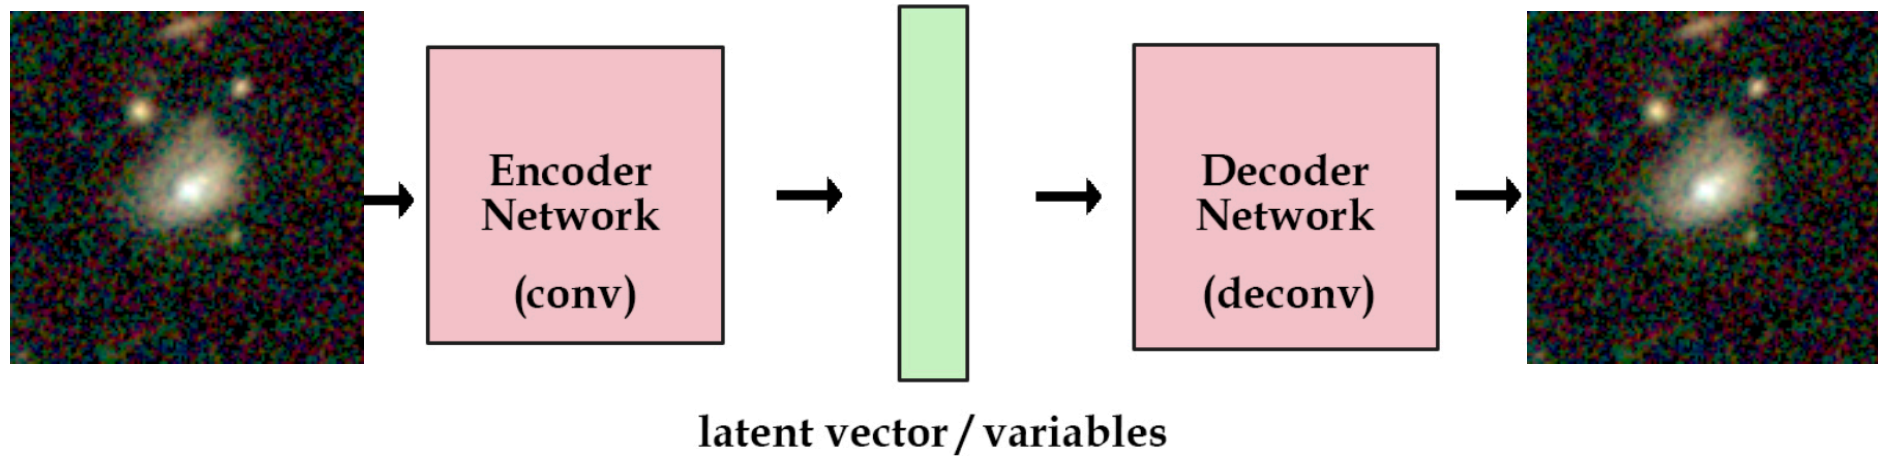
INTRODUCTION TO UNSUPERVISED DEEP LEARNING: GENERATIVE MODELS



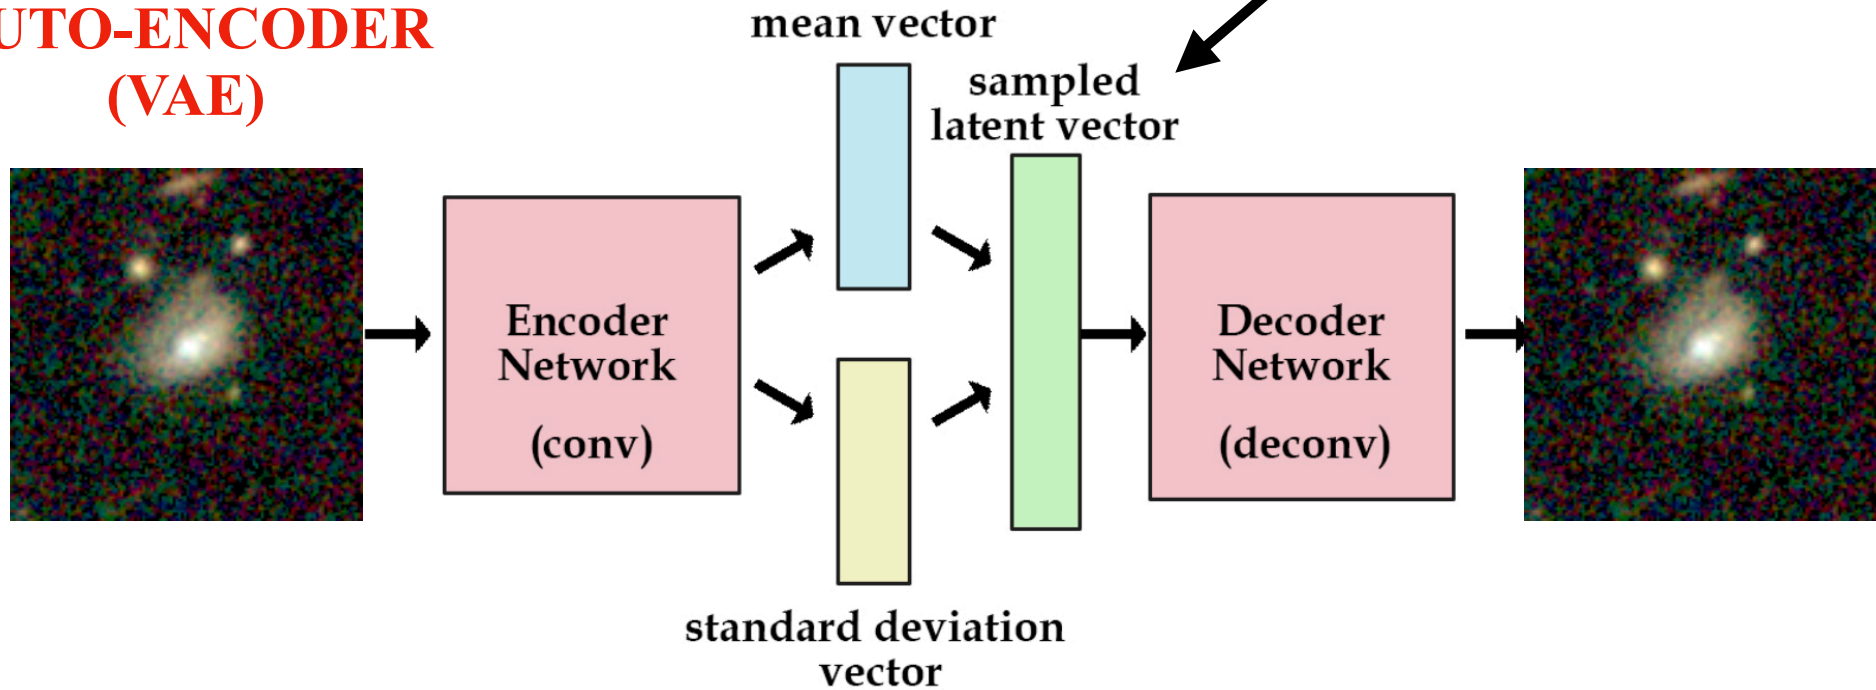
AUTO-ENCODER



AUTO-ENCODER



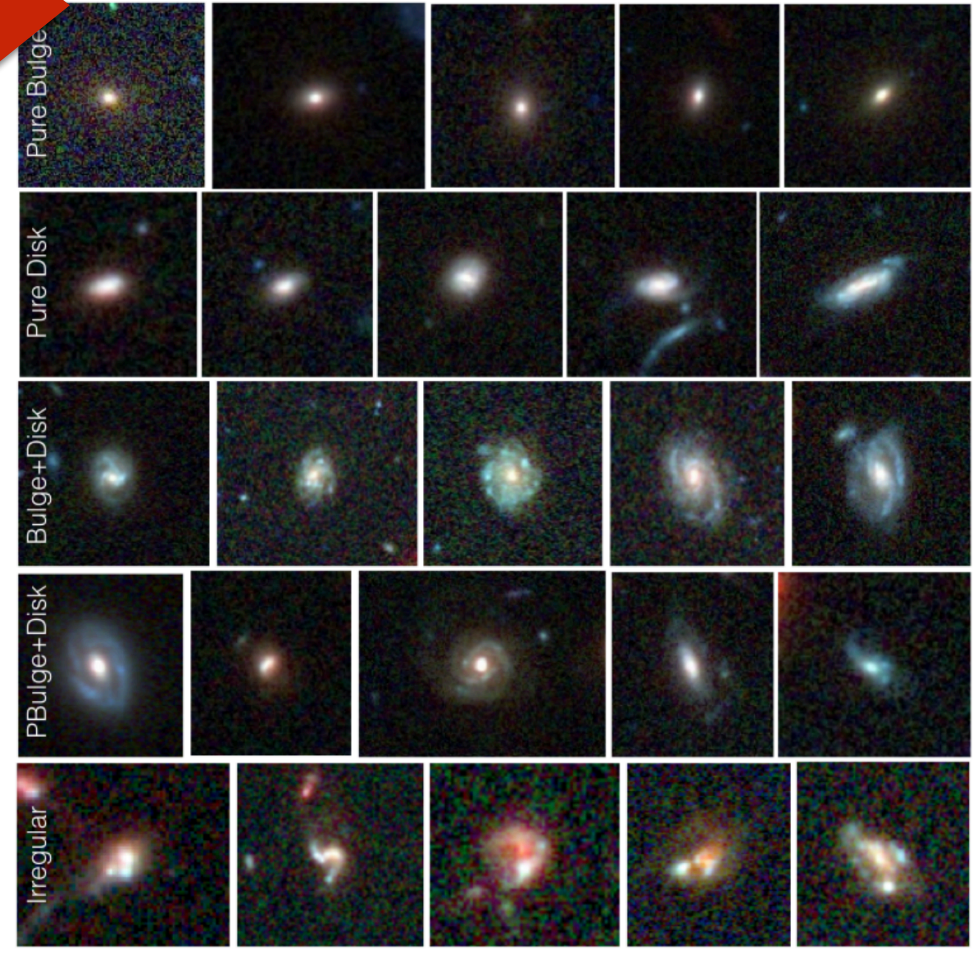
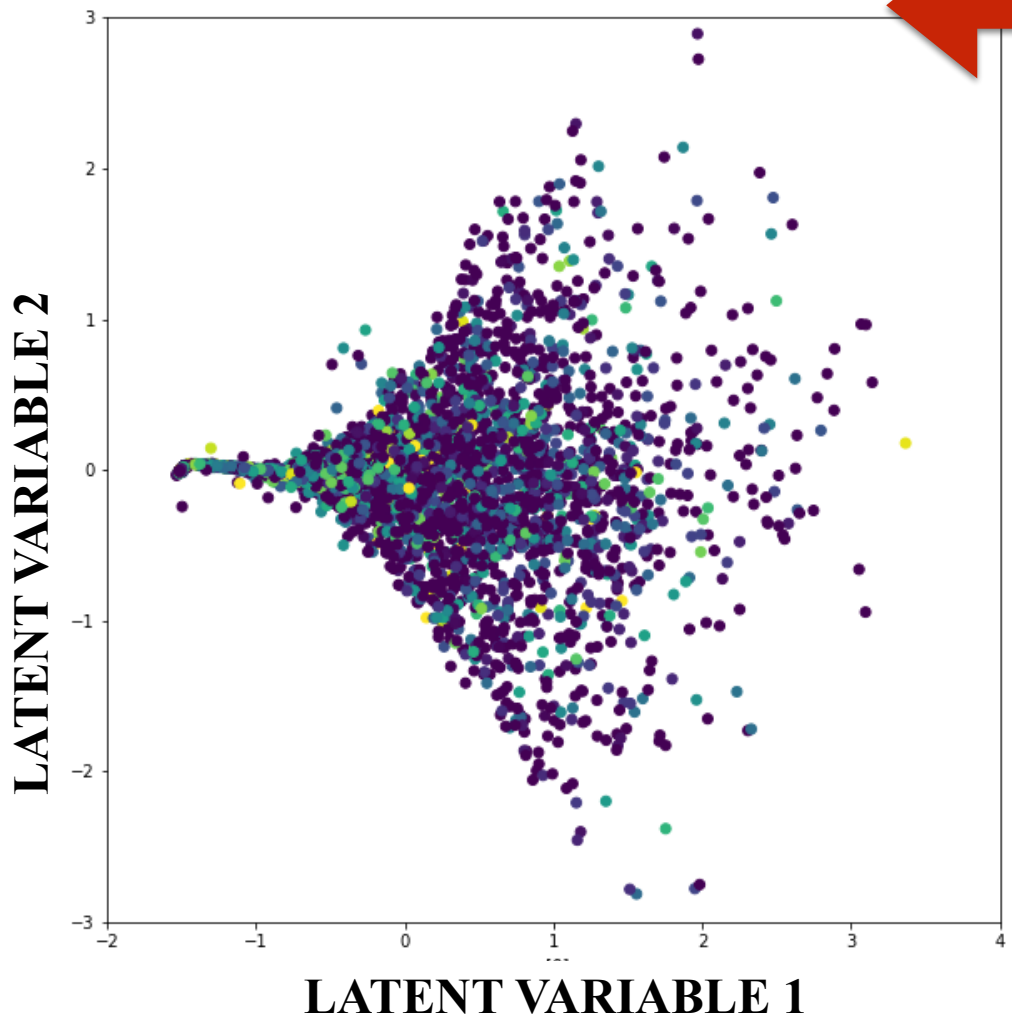
VARIATIONAL AUTO-ENCODER (VAE)



VAE DERIVED LATENT SPACE FOR CANDELS GALAXIES [H BAND]

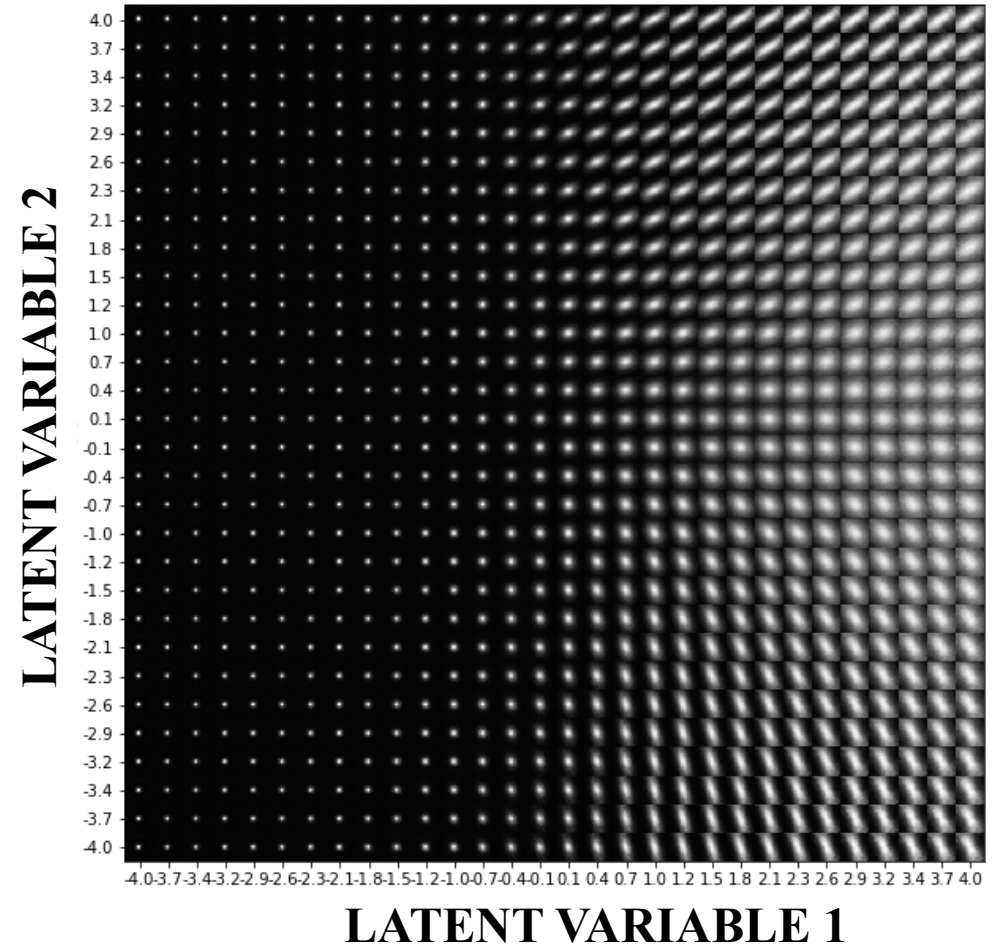
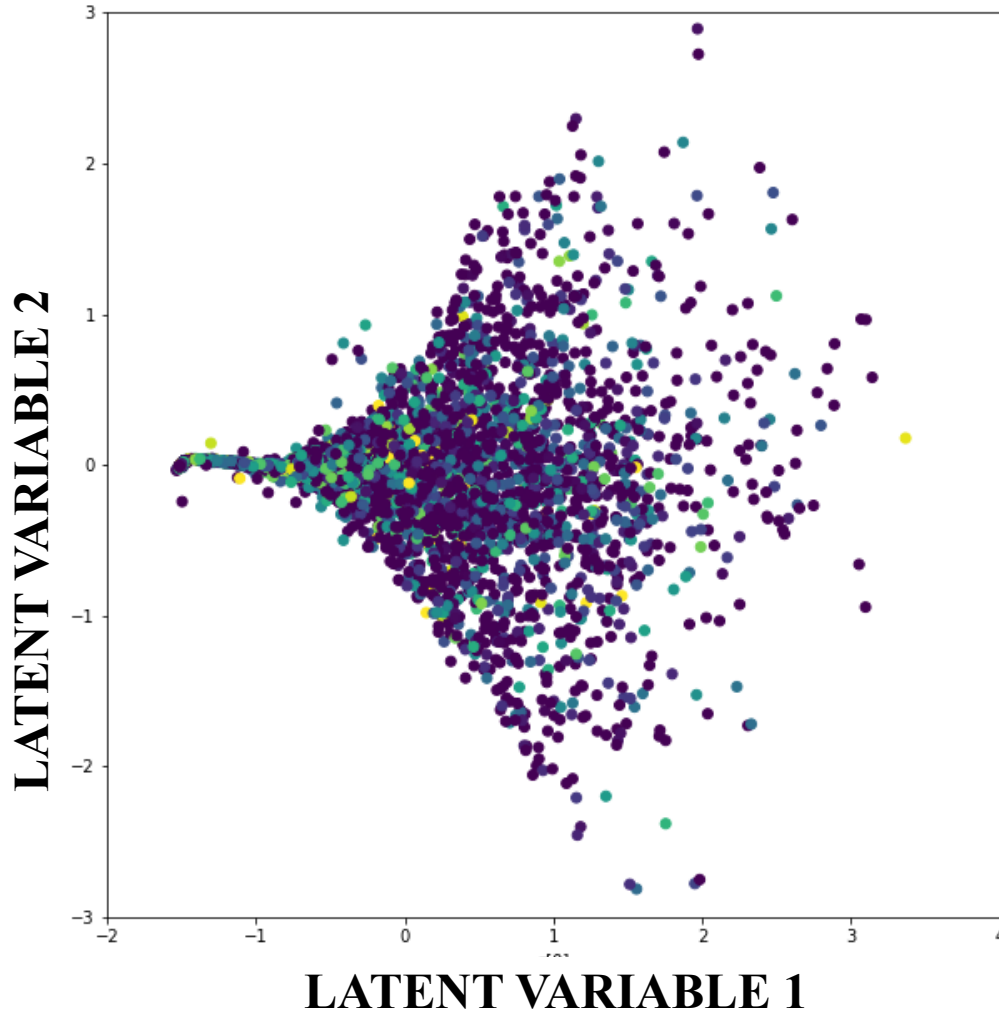
PRELIMINARY!

2 DIMENSIONS

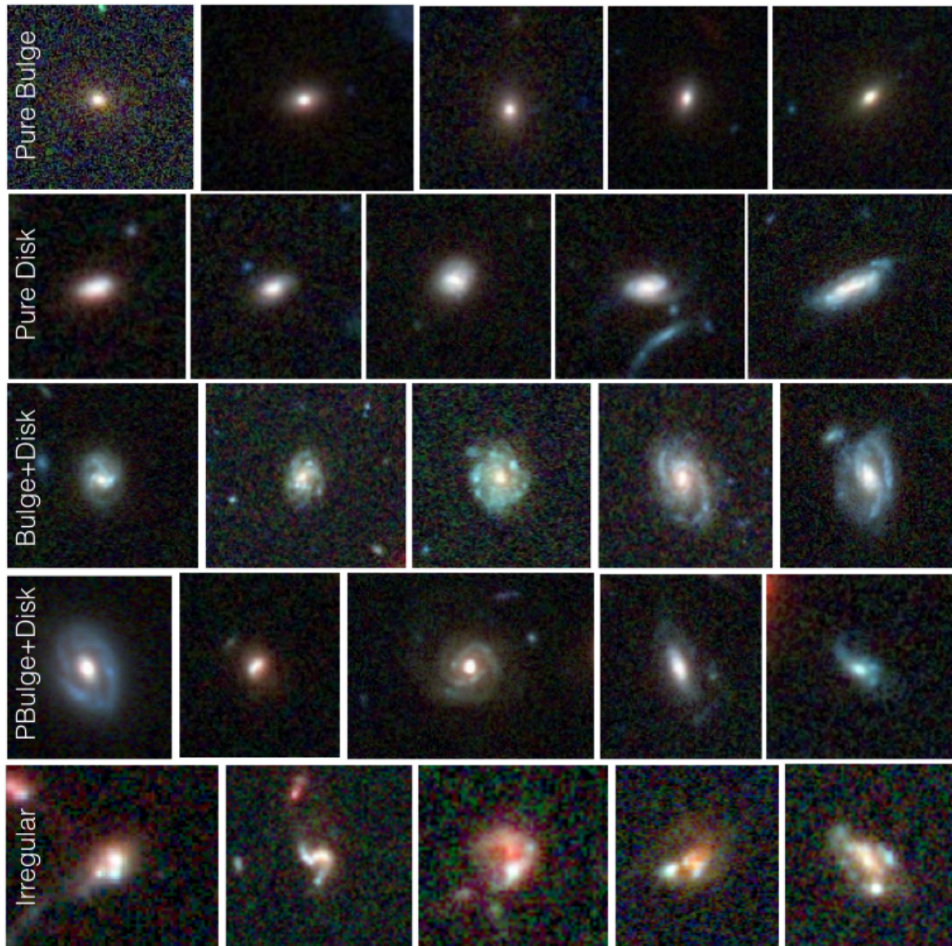


VAE DERIVED LATENT SPACE FOR CANDELS GALAXIES [H BAND]

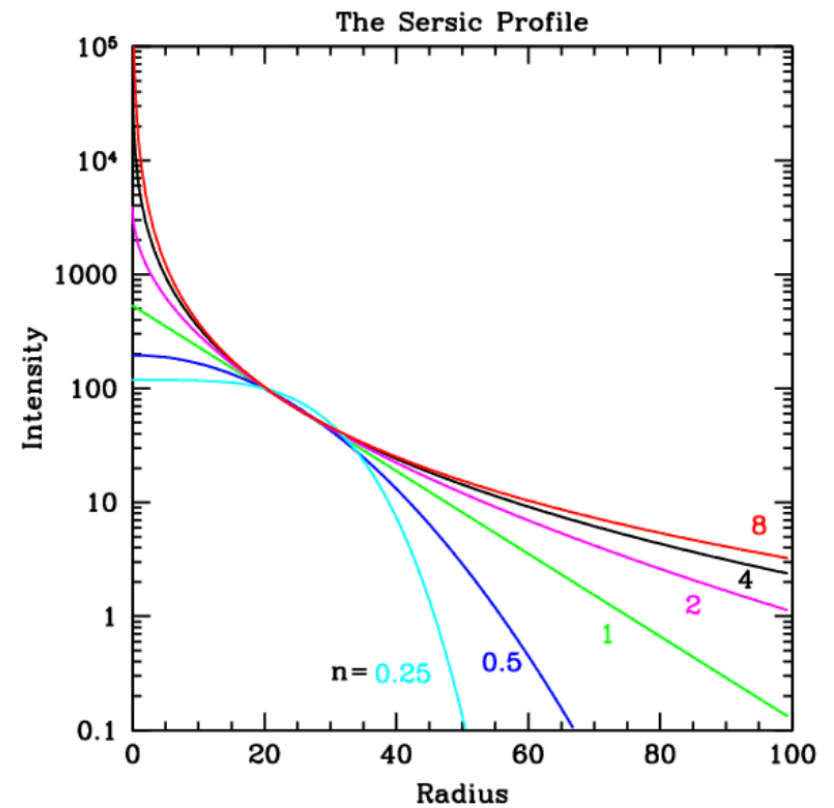
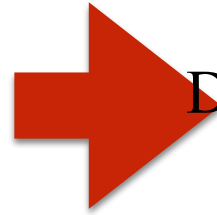
PRELIMINARY!



INTERPOLATION IN THE LATENT SPACE GENERATES GALAXIES WITH DIFFERENT PROPERTIES

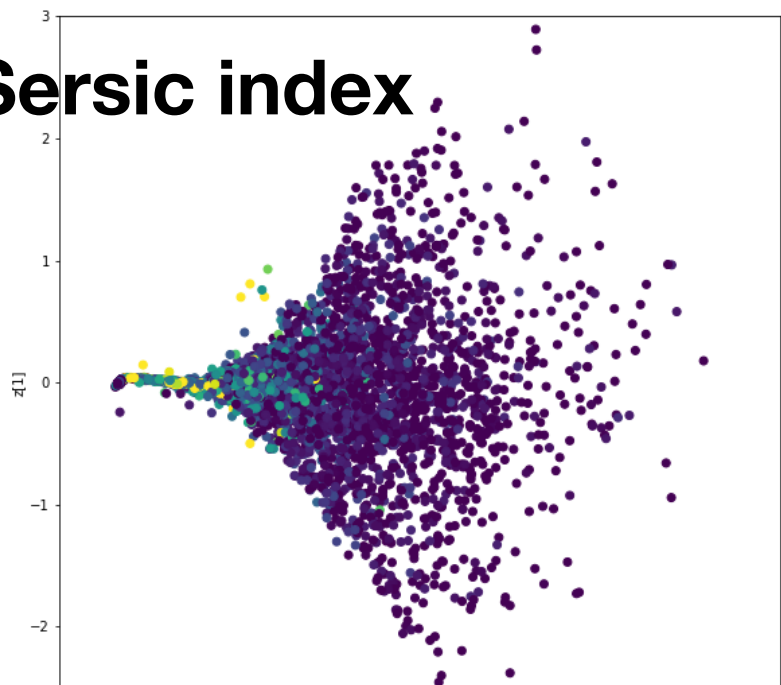


TRADITIONALLY GALAXY
STRUCTURE IS
DESCRIBED WITH AT LEAST 4
PARAMETERS

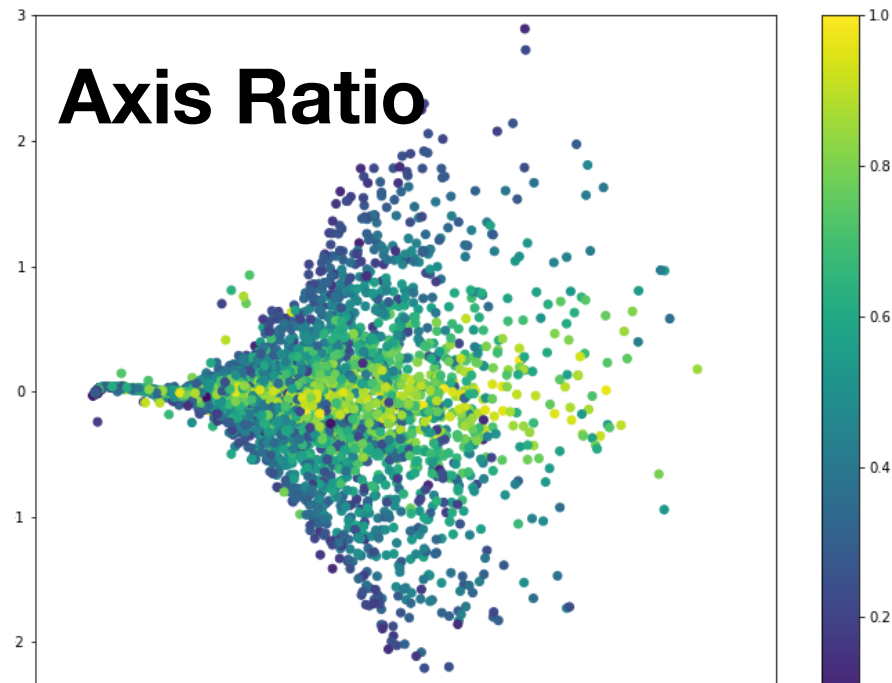


Effective Radius, Sersic Index, Axis-Ratio, Position Angle

Sersic index

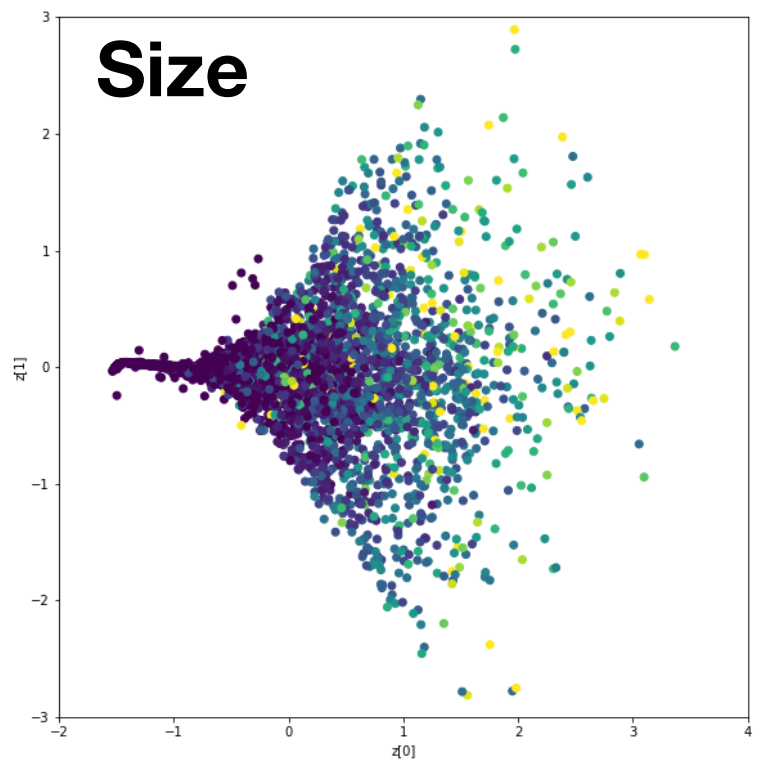


Axis Ratio

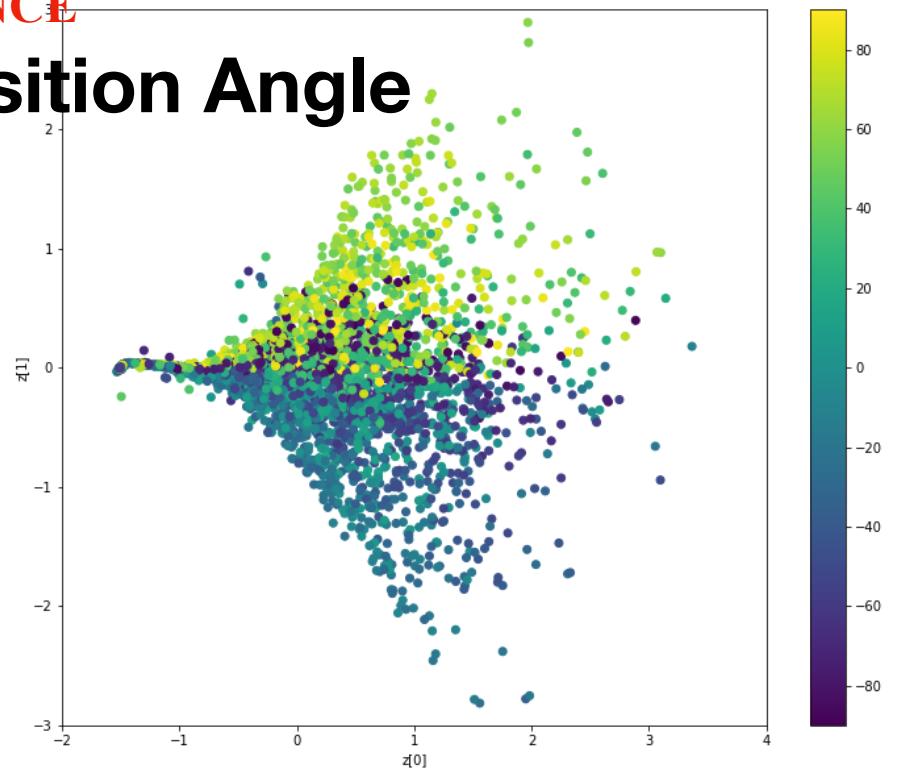


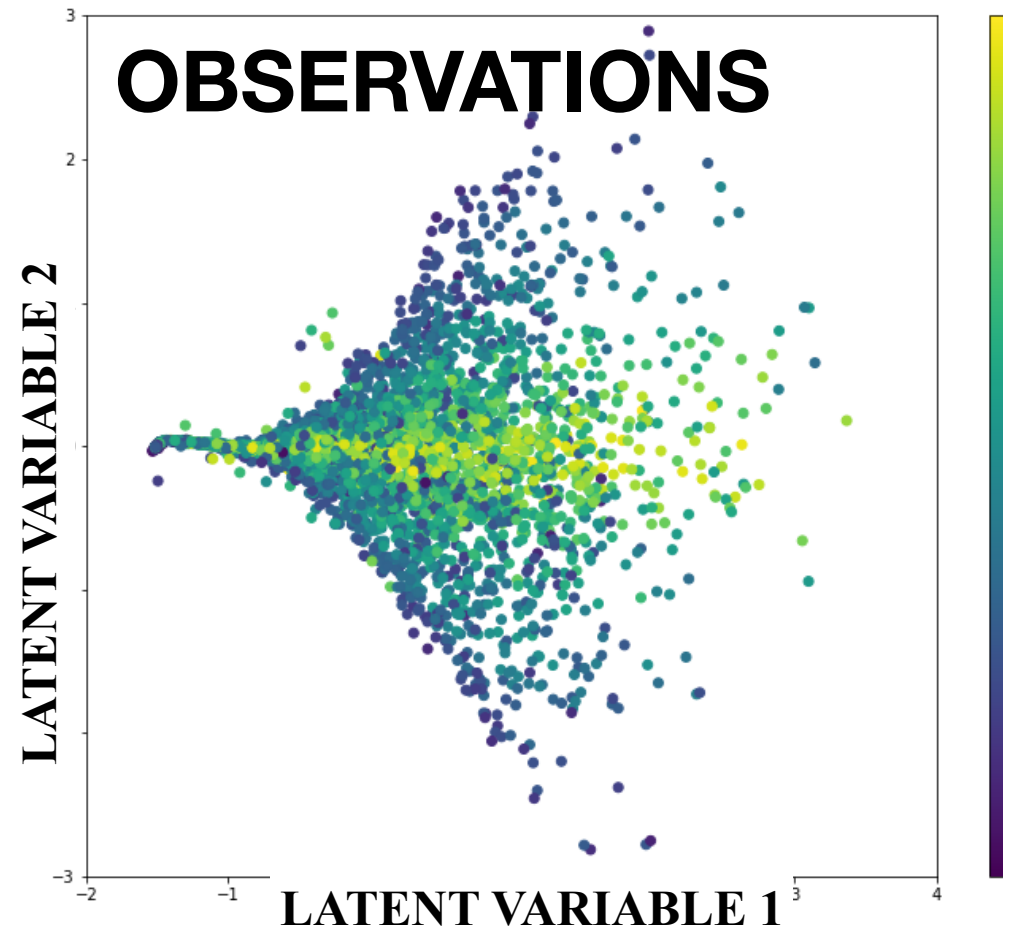
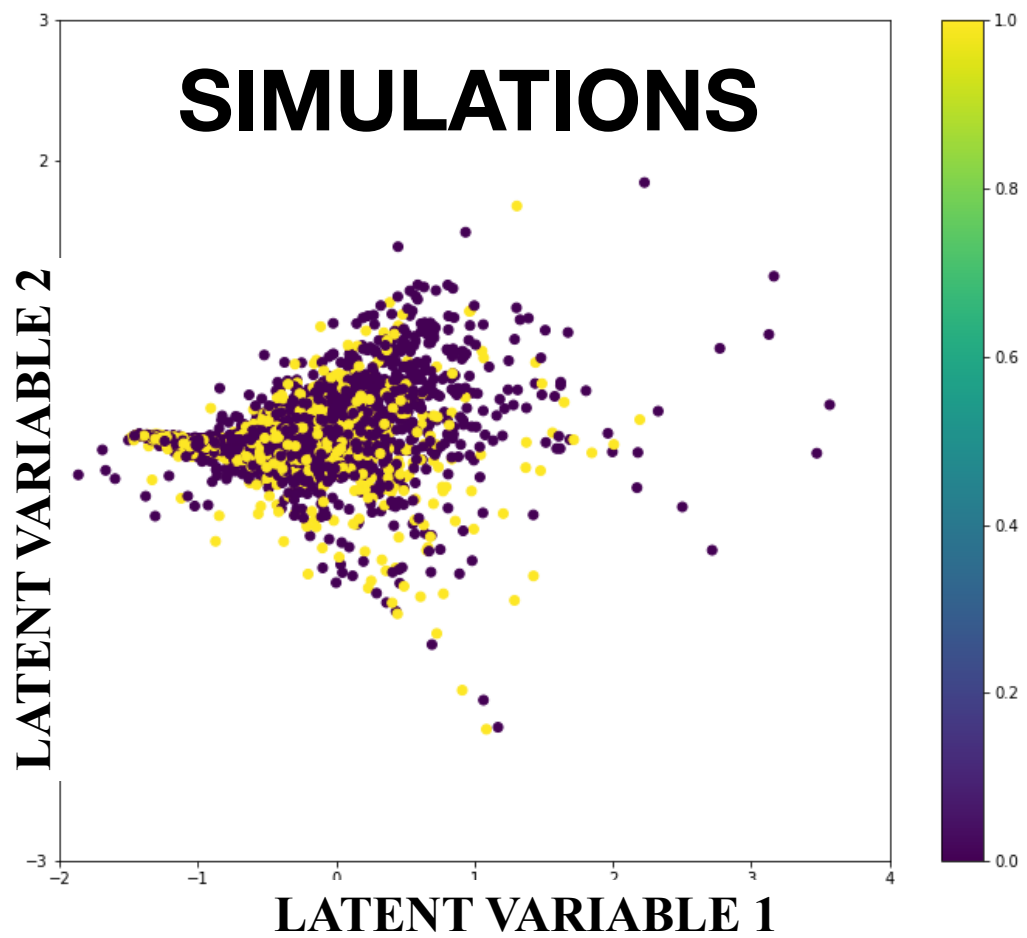
**EXPLORE THE MORPHOLOGICAL DISTRIBUTION OF GALAXIES
AT A GLANCE**

Size

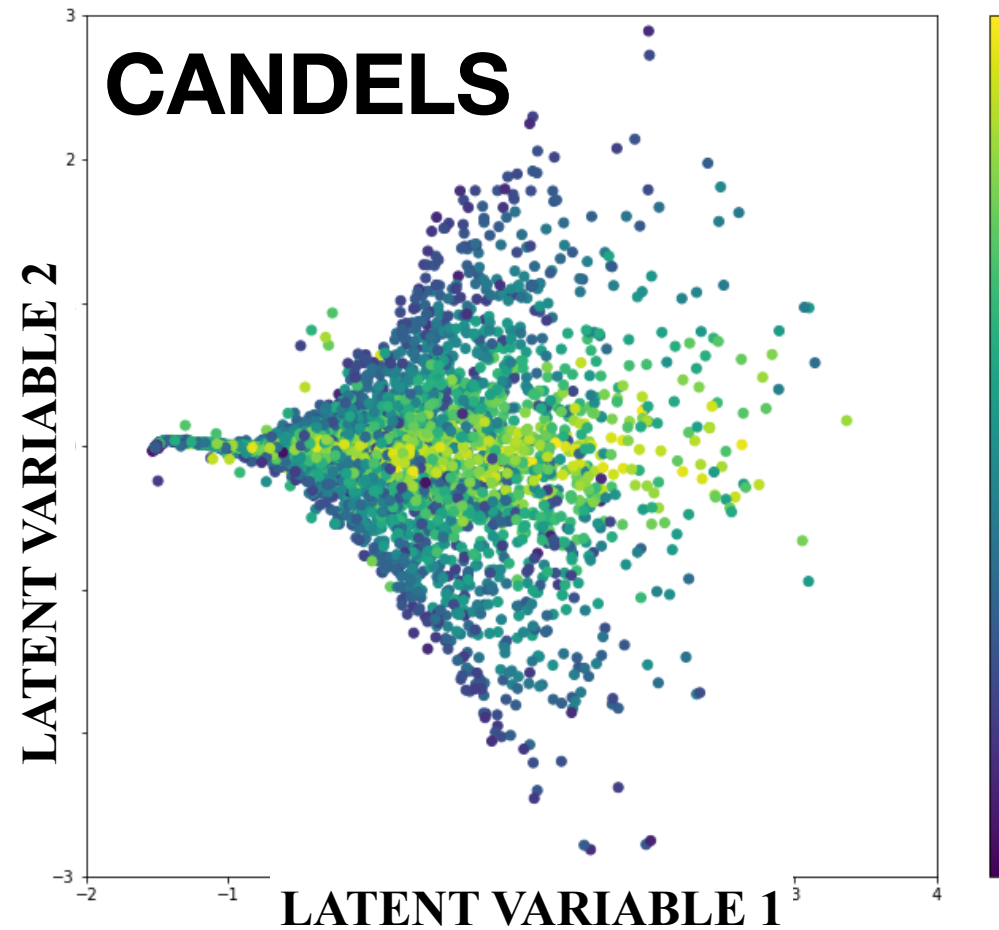
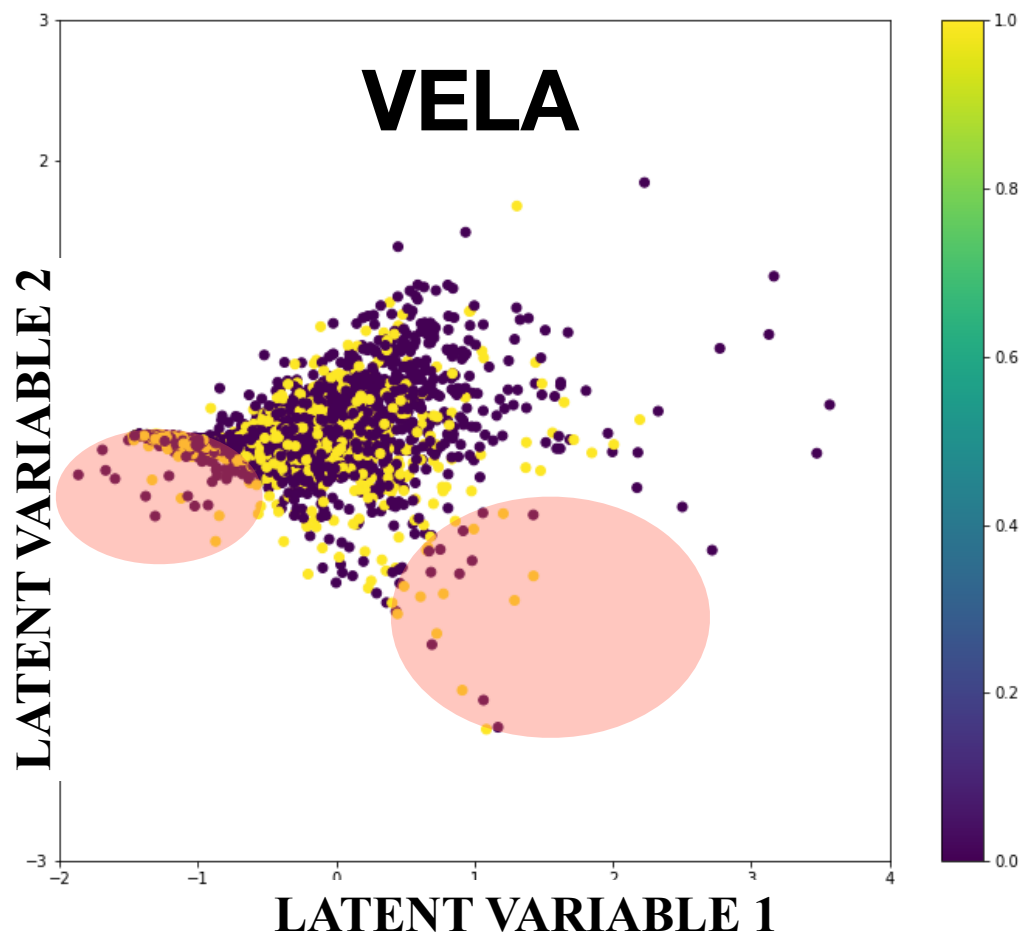


Position Angle



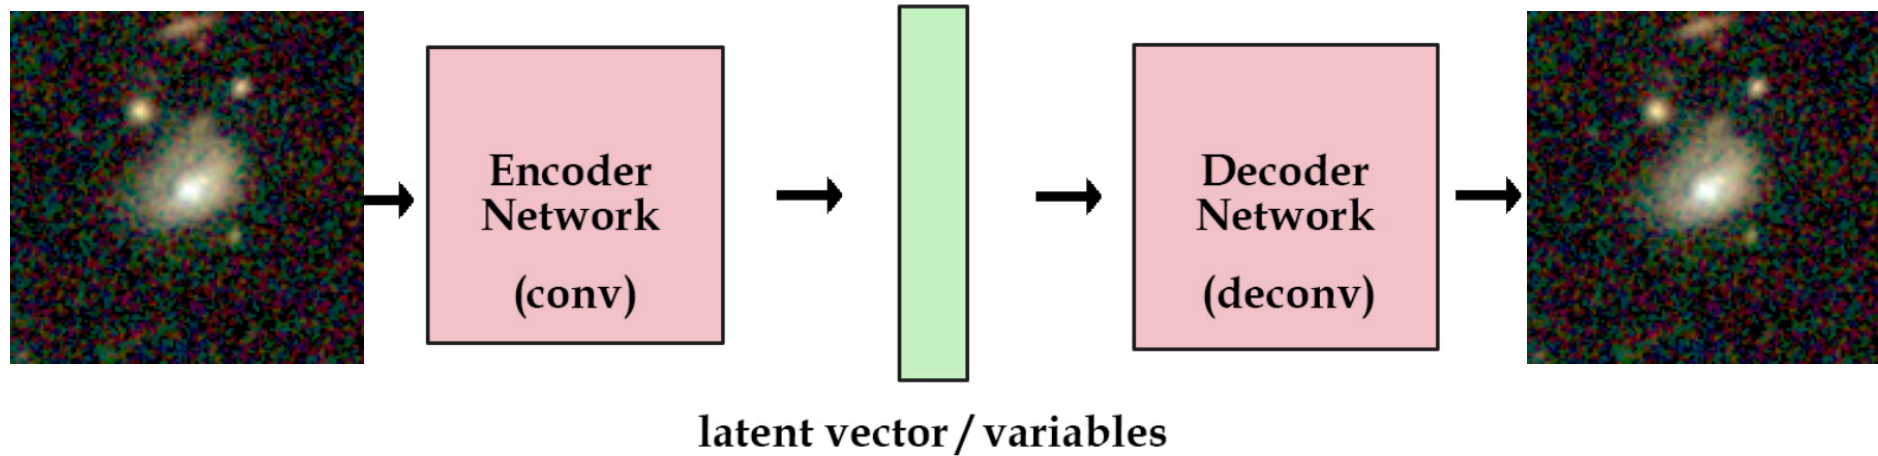


**HOW SIMULATIONS POPULATE
THE LATENT SPACE?**

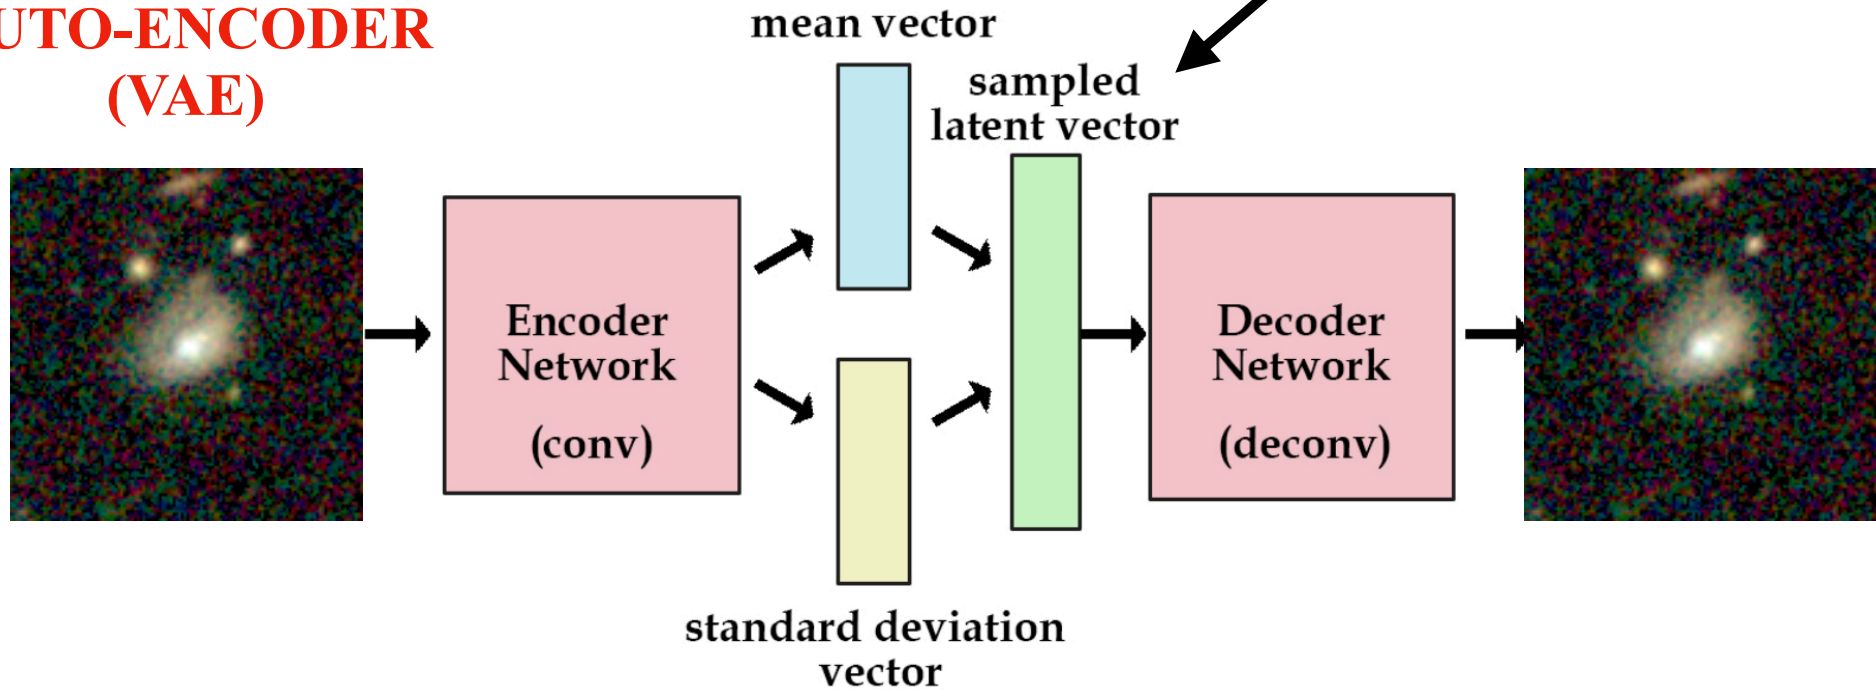


**HOW SIMULATIONS POPULATE
THE LATENT SPACE?**

AUTO-ENCODER



VARIATIONAL AUTO-ENCODER (VAE)



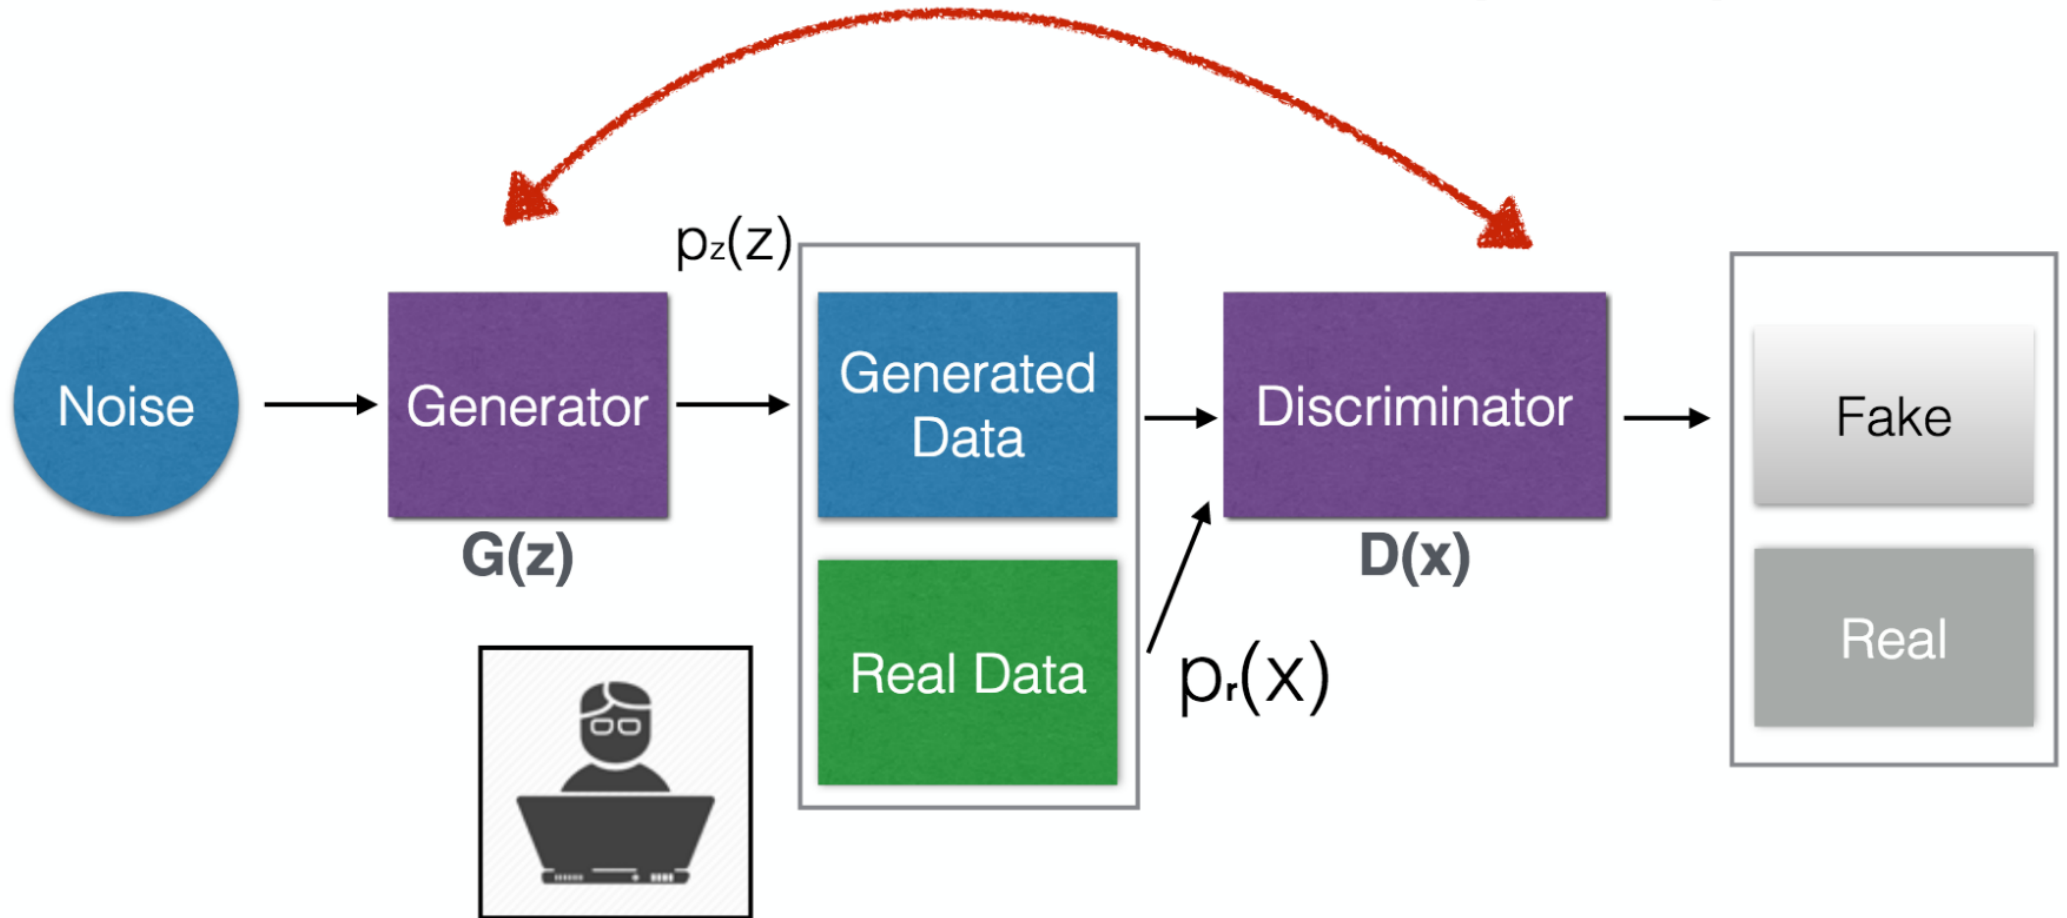
GENERATIVE ADVERSARIAL NETWORKS

(Goodfellow+14)

5570 citation in 4 years!

GENERATIVE ADVERSARIAL NETWORKS

(Goodfellow+14)

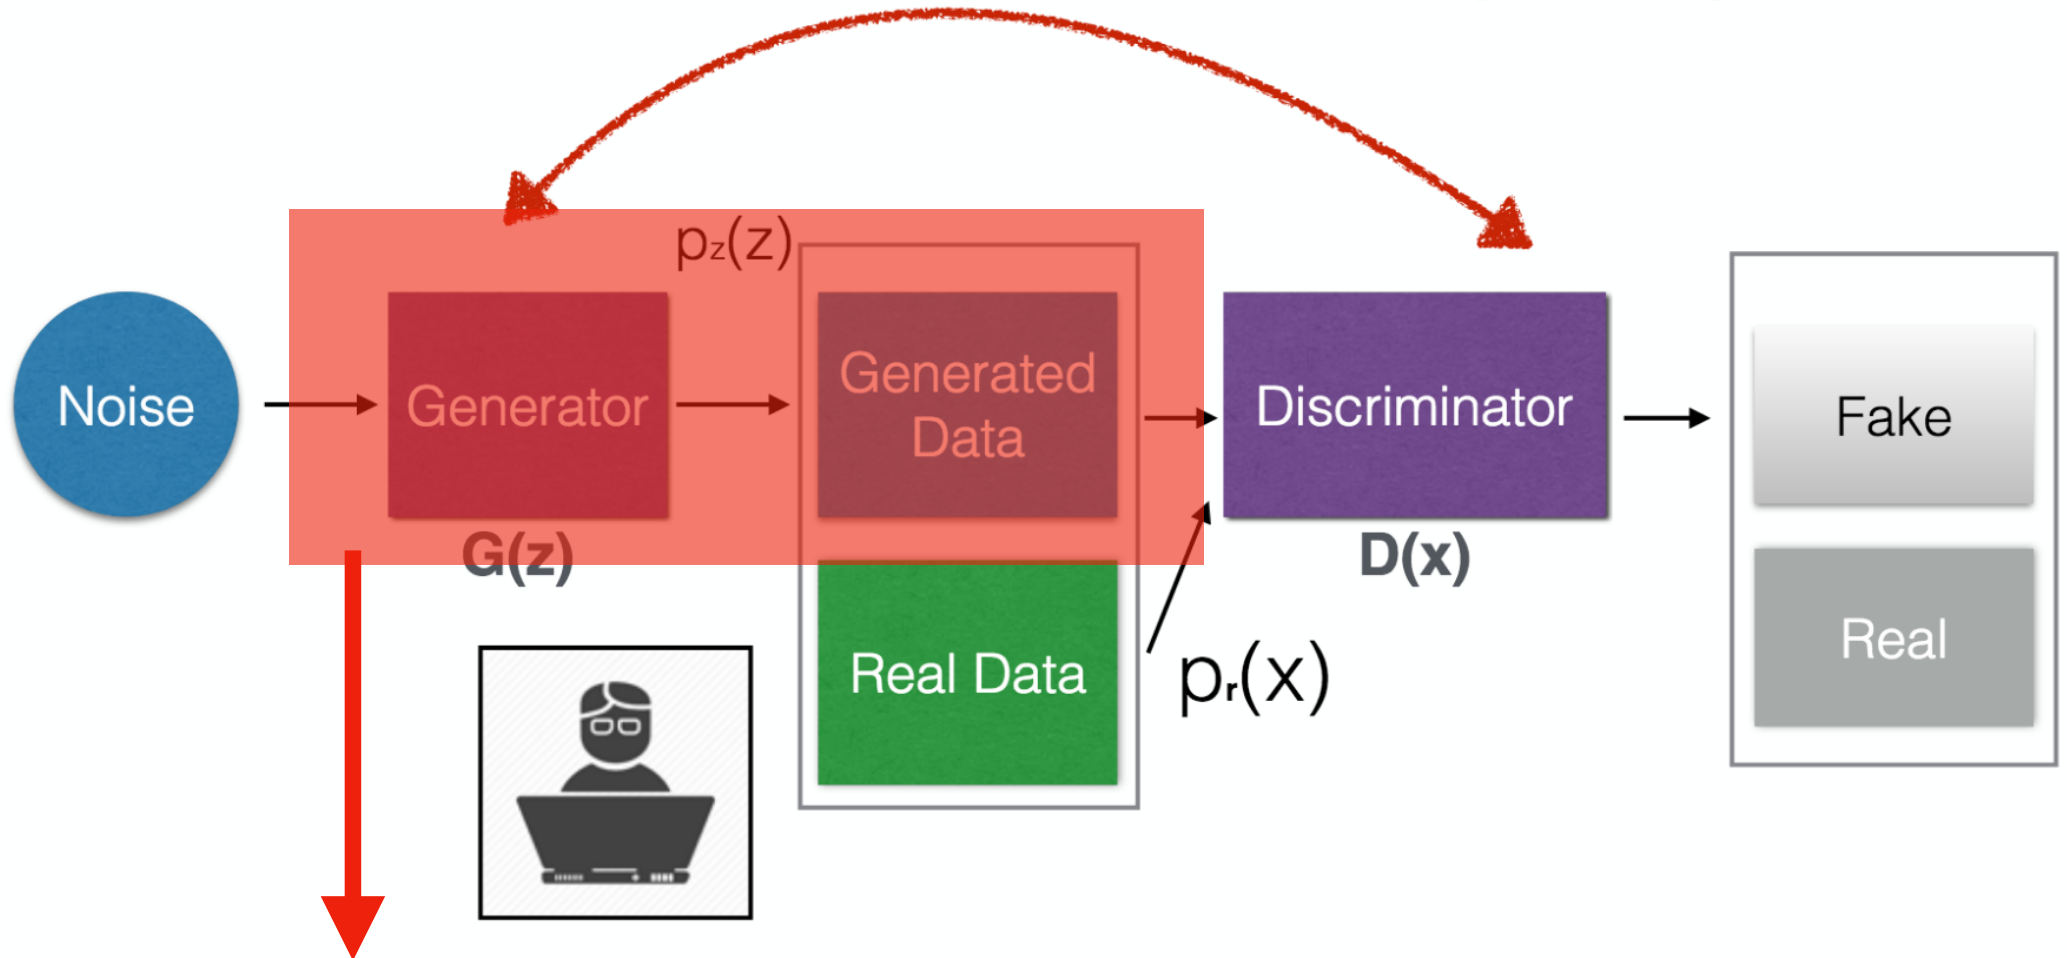


TWO COMPETING NETWORKS

GENERATIVE ADVERSARIAL NETWORKS

(Goodfellow+)

TWO COMPETING NETWORKS

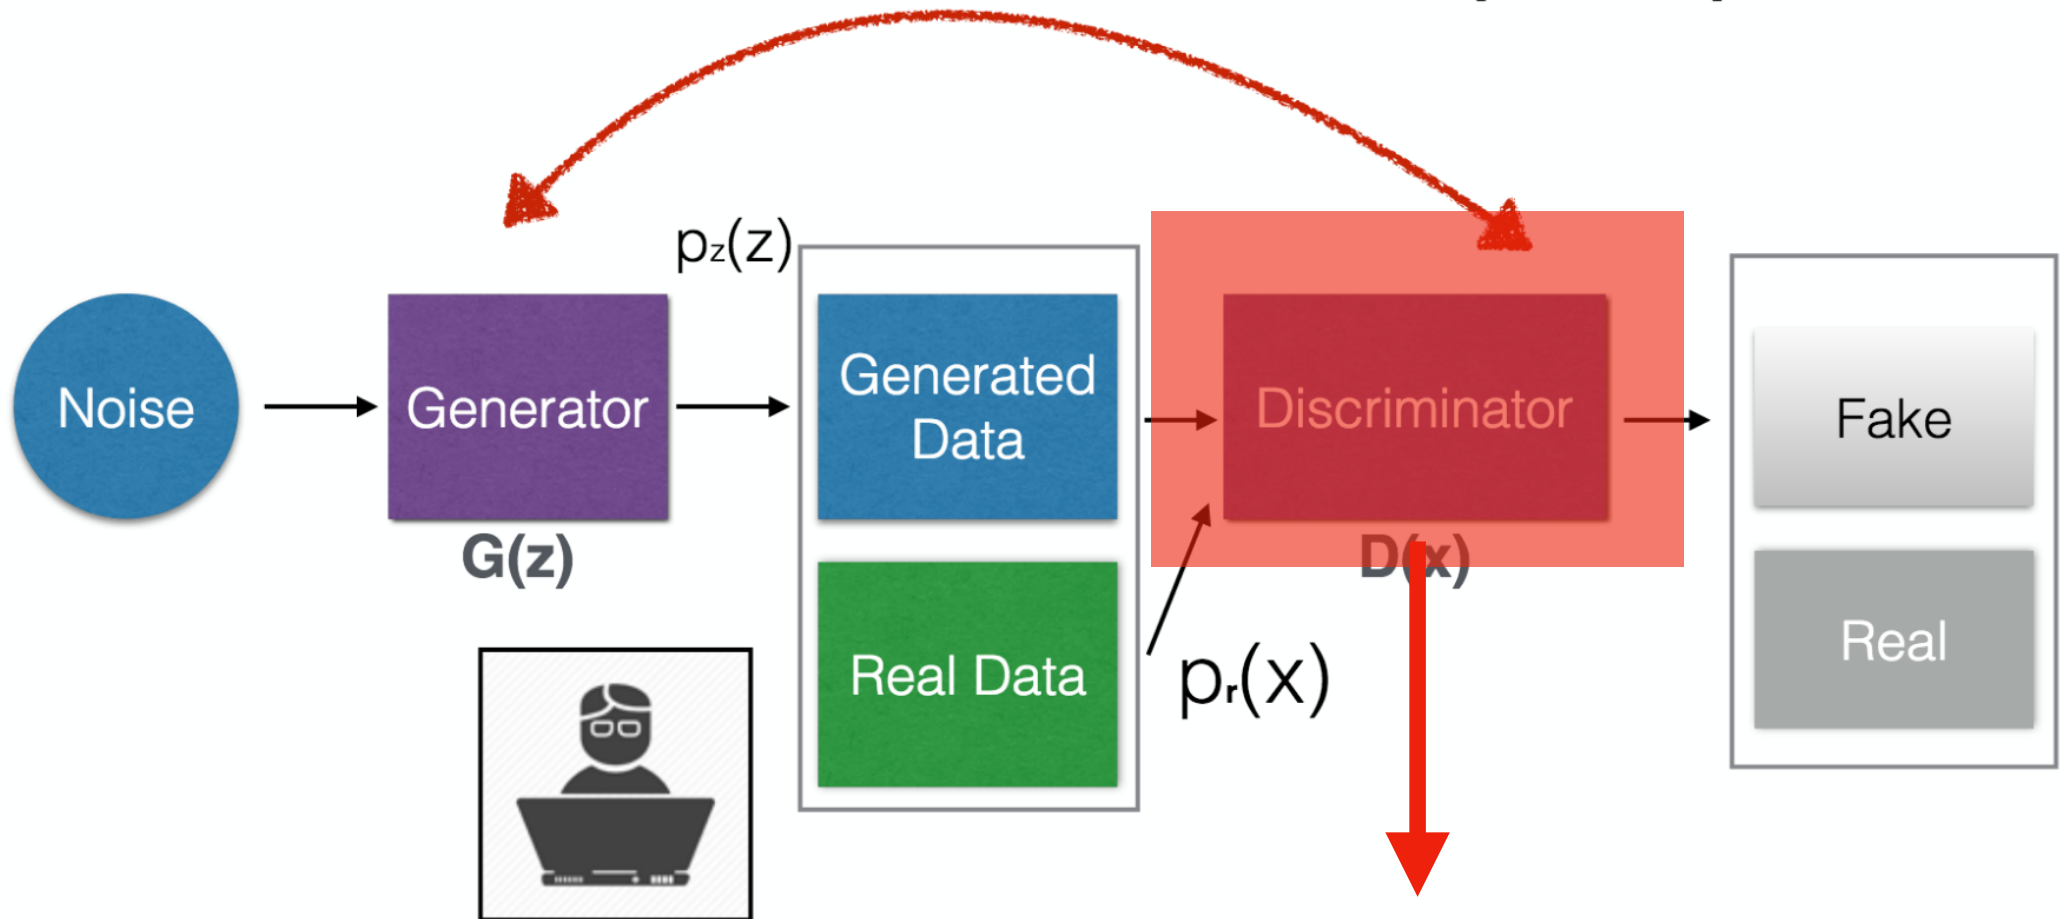


Every 2 iterations the generator is trained to force the discriminator to classify as real

GENERATIVE ADVERSARIAL NETWORKS

(Goodfellow+)

TWO COMPETING NETWORKS



Every 2 iterations the discriminator is trained to force to distinguish between real and fake

PROGRESSIVE GROWING OF GANs FOR IMPROVED QUALITY, STABILITY, AND VARIATION

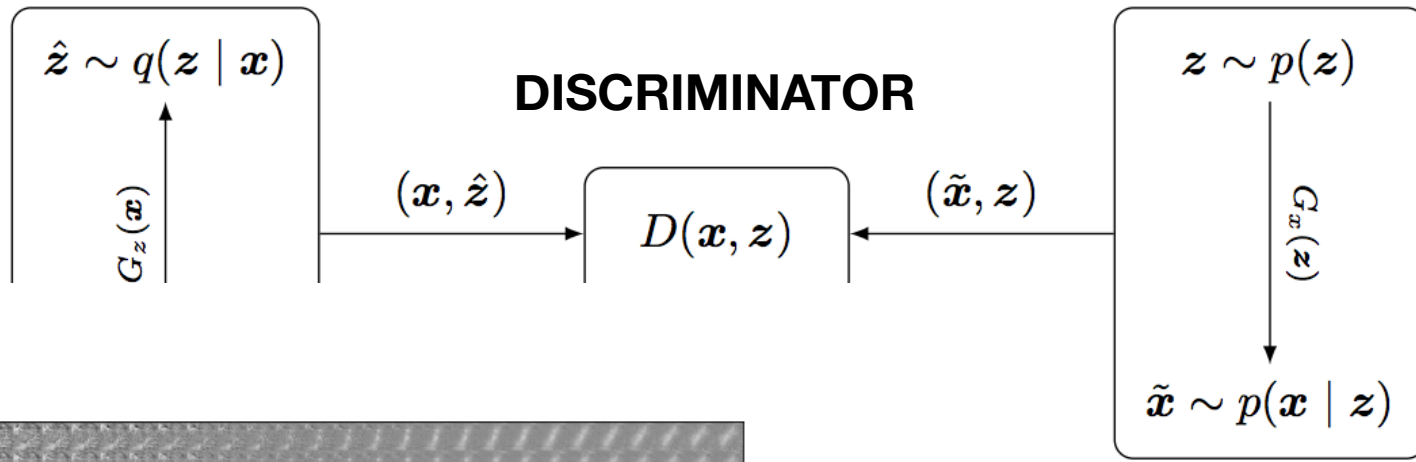
Tero Karras
NVIDIA

Timo Aila
NVIDIA

Samuli Laine
NVIDIA

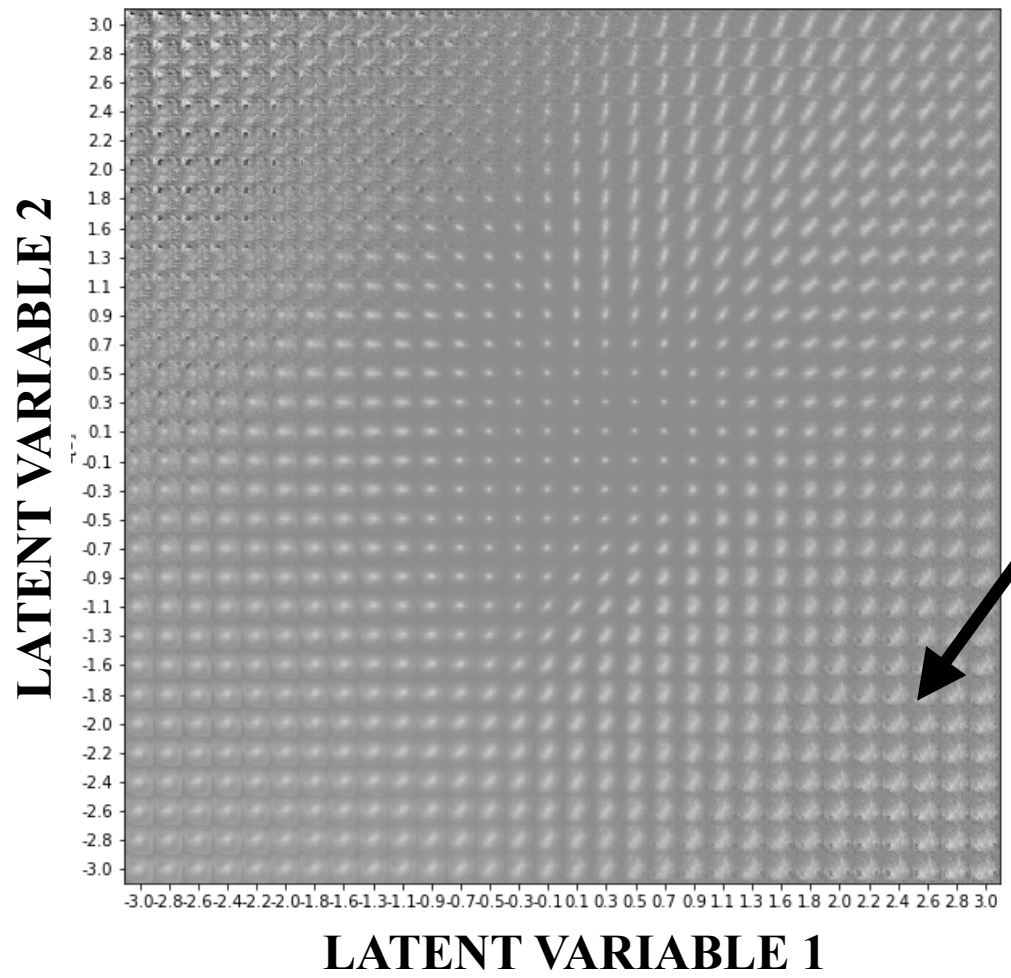
Jaakko Lehtinen
NVIDIA
Aalto University

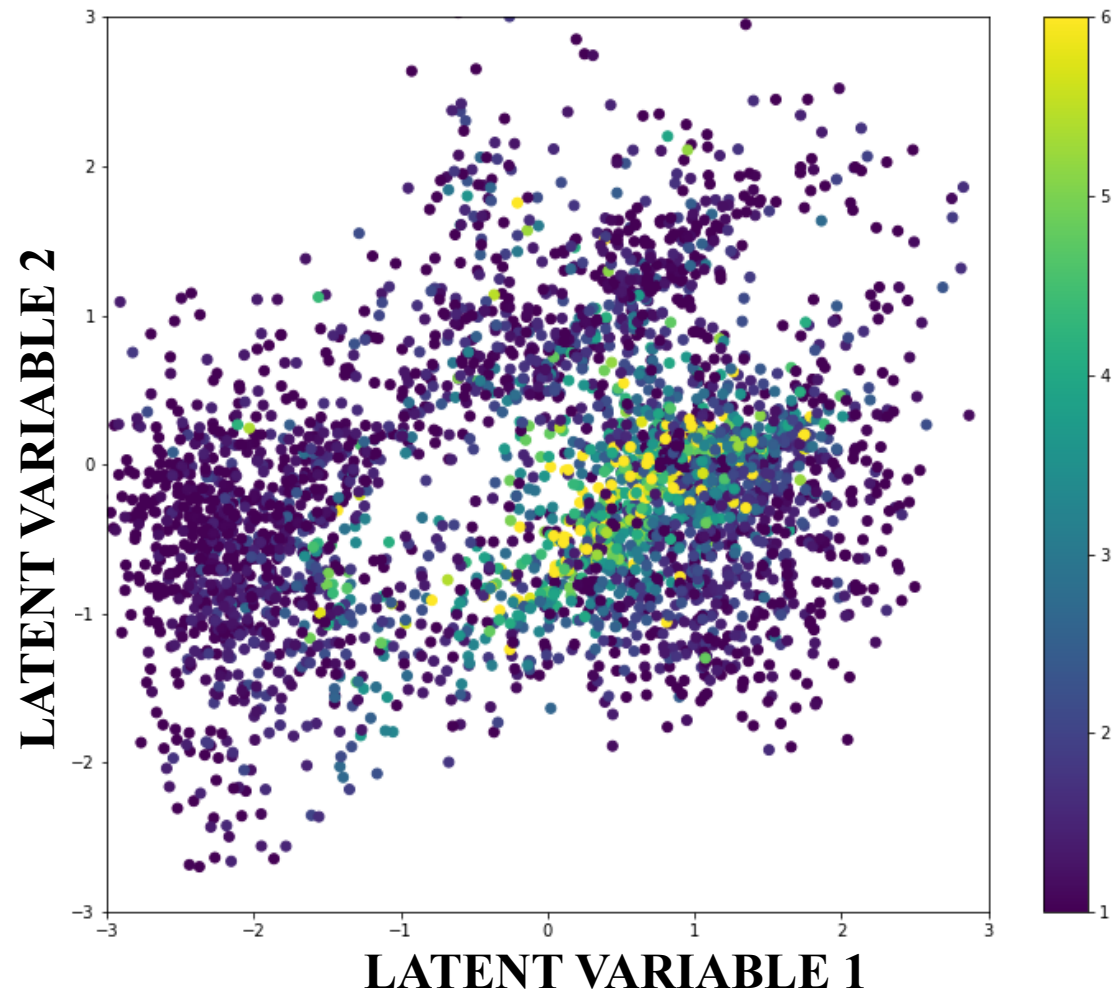
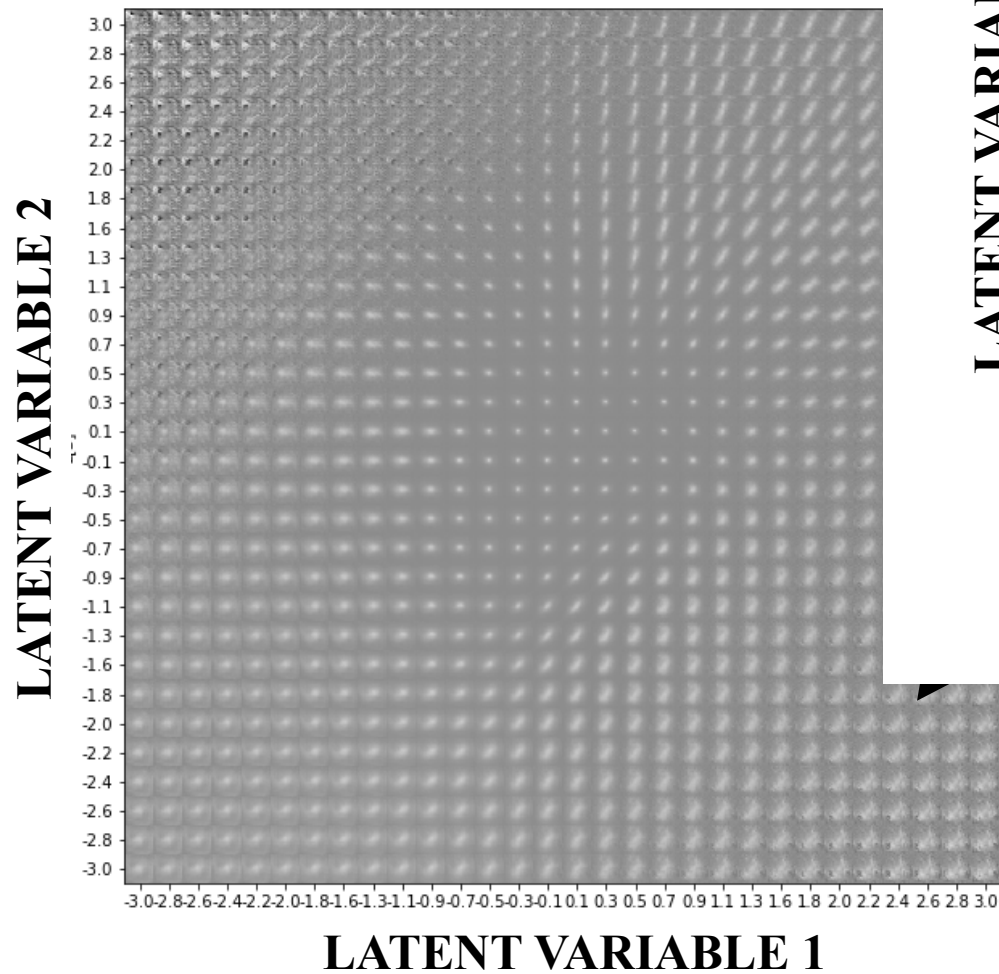
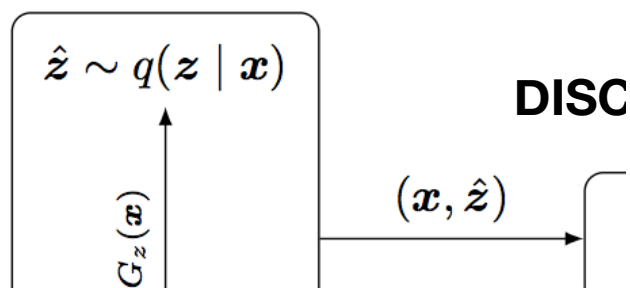




Adversarial Latent Inference (ALI) game. **Dumoulin+17**

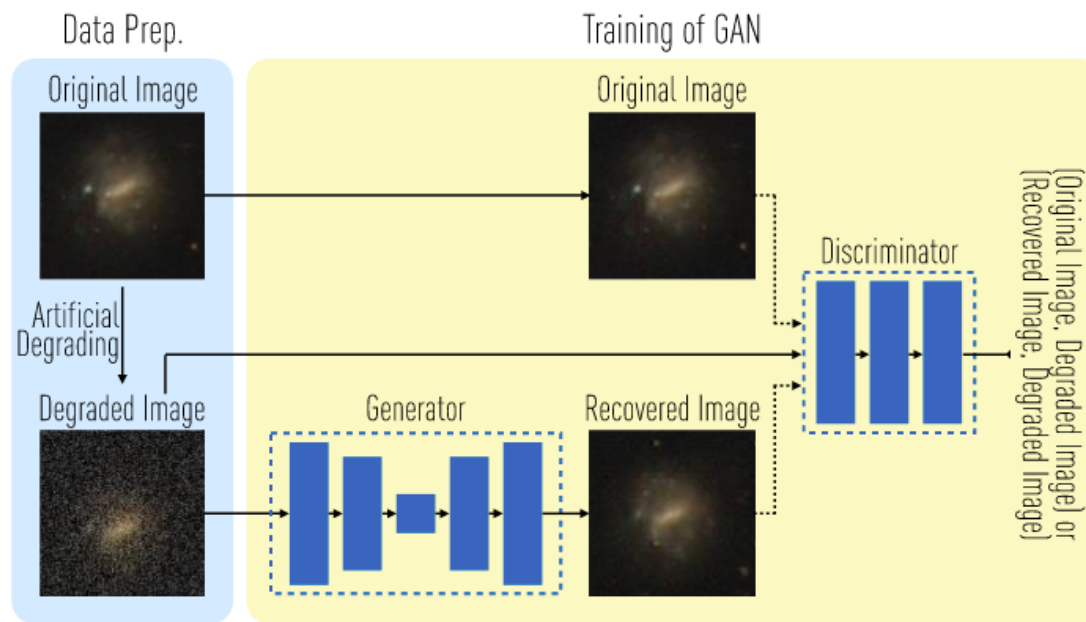
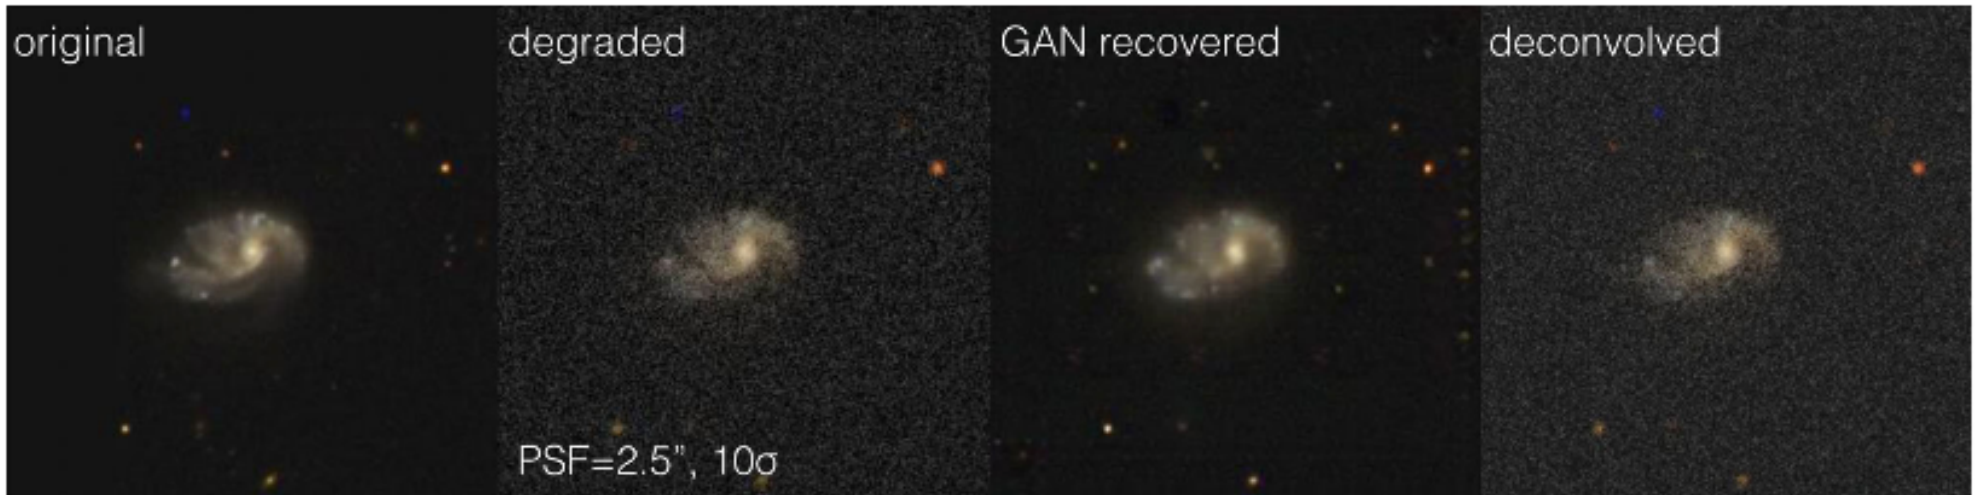
GENERATES MORE COMPLEX MODELS





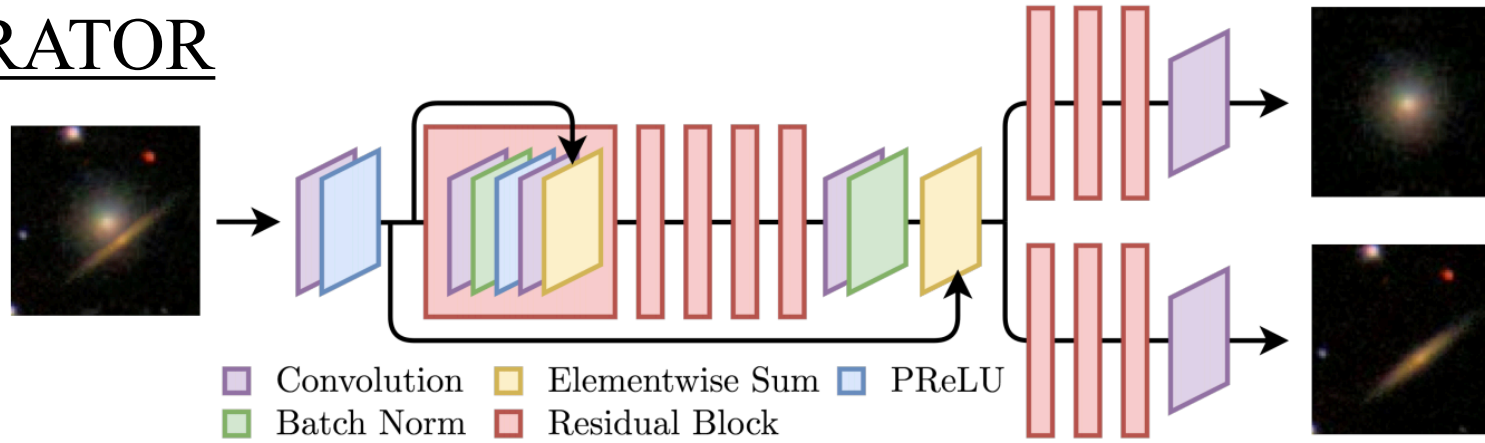
MORE DIFFICULT TO INTERPRET

GANs: OTHER APPLICATIONS IN ASTRONOMY

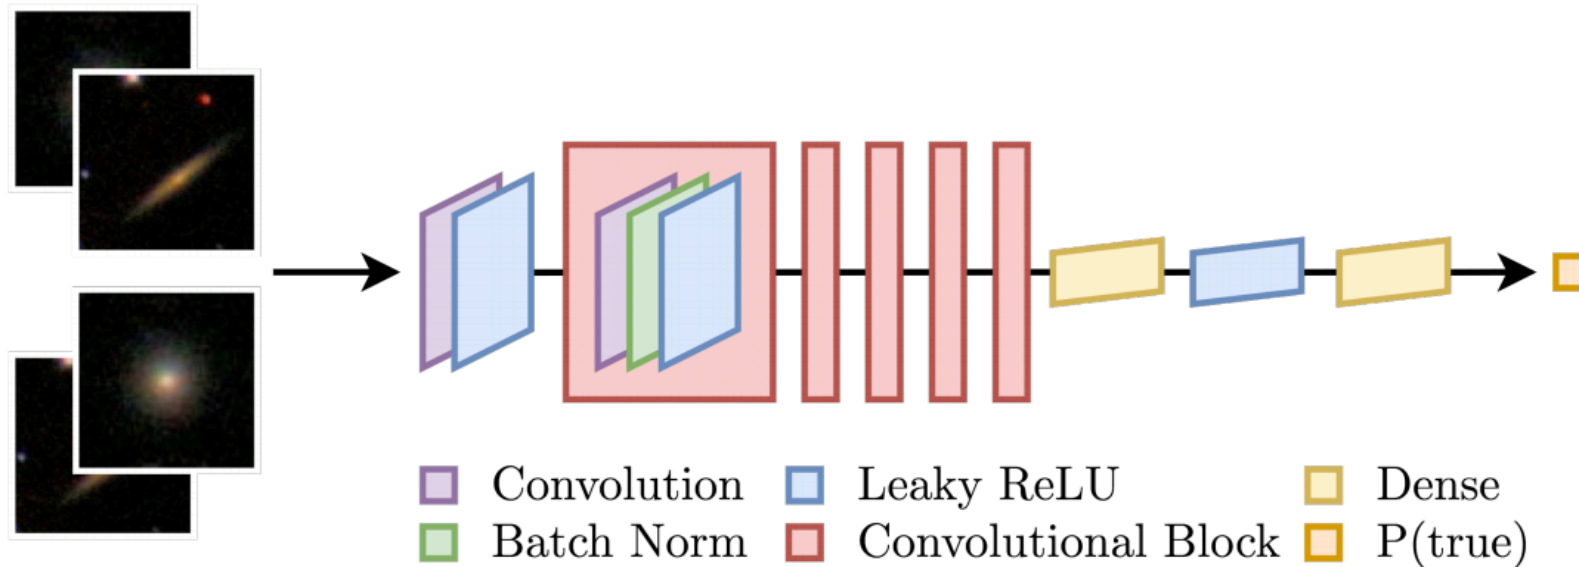


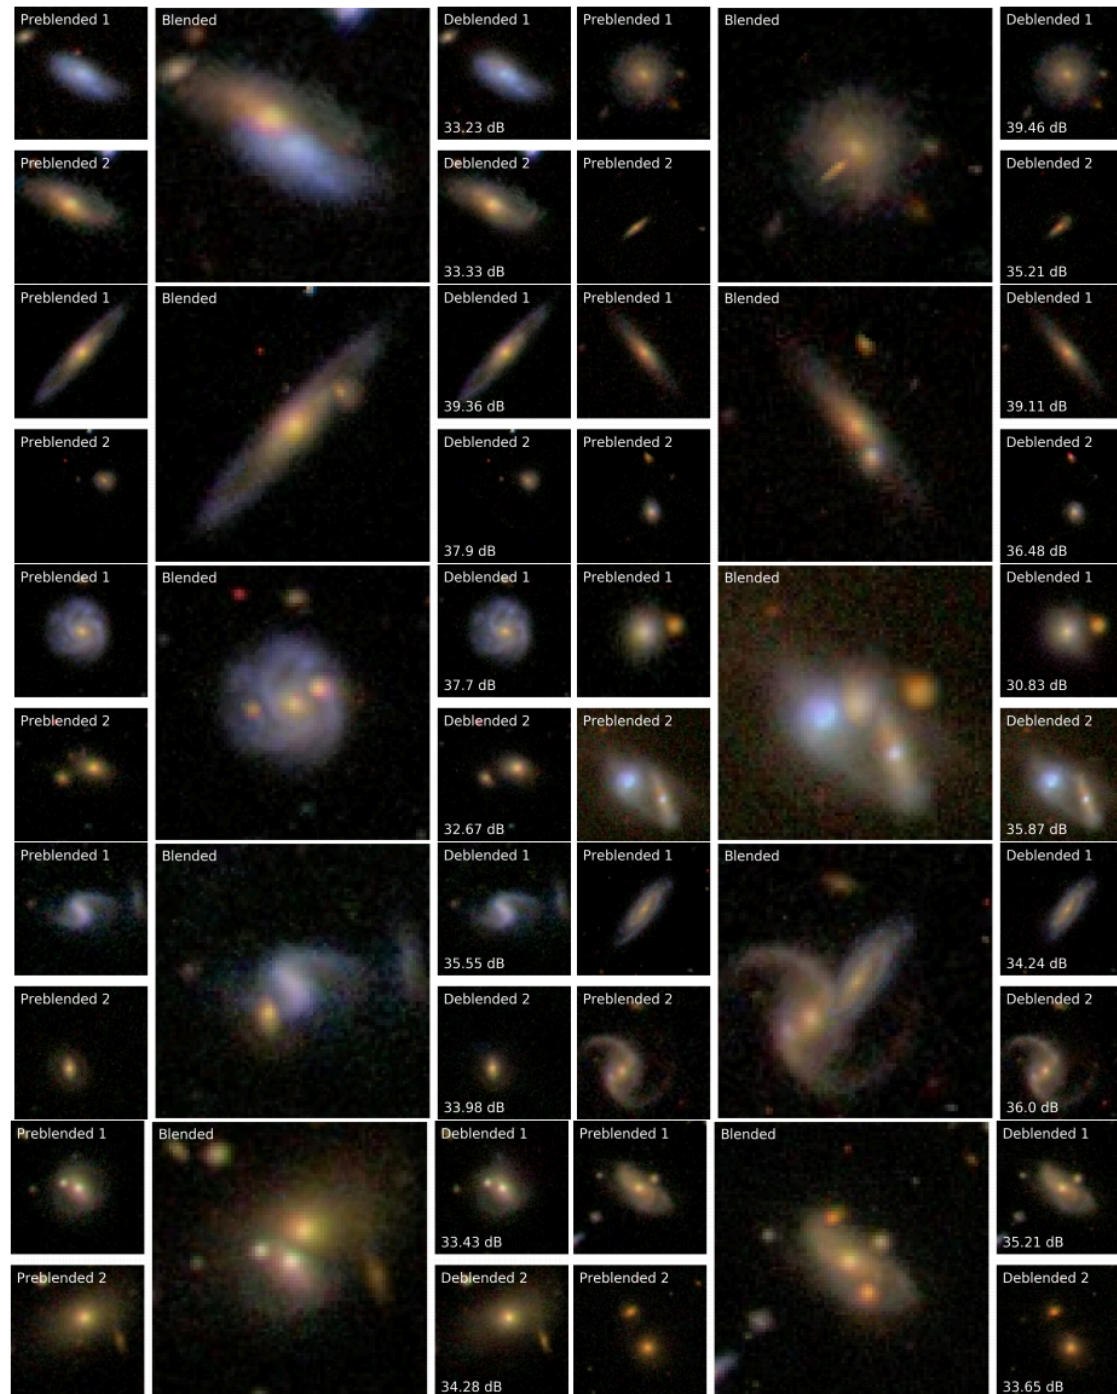
Schawinsky+17

GENERATOR



DISCRIMINATOR





ANOMALY DETECTION WITH GANs

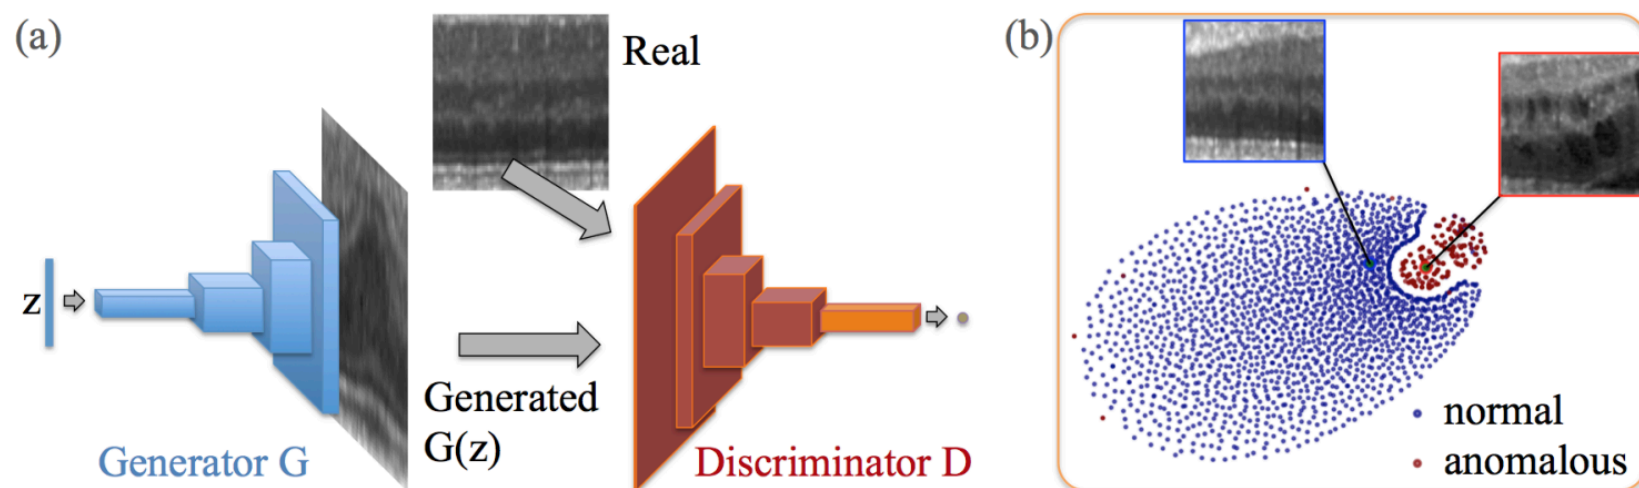


Fig. 2. (a) Deep convolutional generative adversarial network. (b) t-SNE embedding of normal (blue) and anomalous (red) images on the feature representation of the last convolution layer (orange in (a)) of the discriminator.

ANOMALY DETECTION WITH GANs

ANOMALY SCORE

(a)

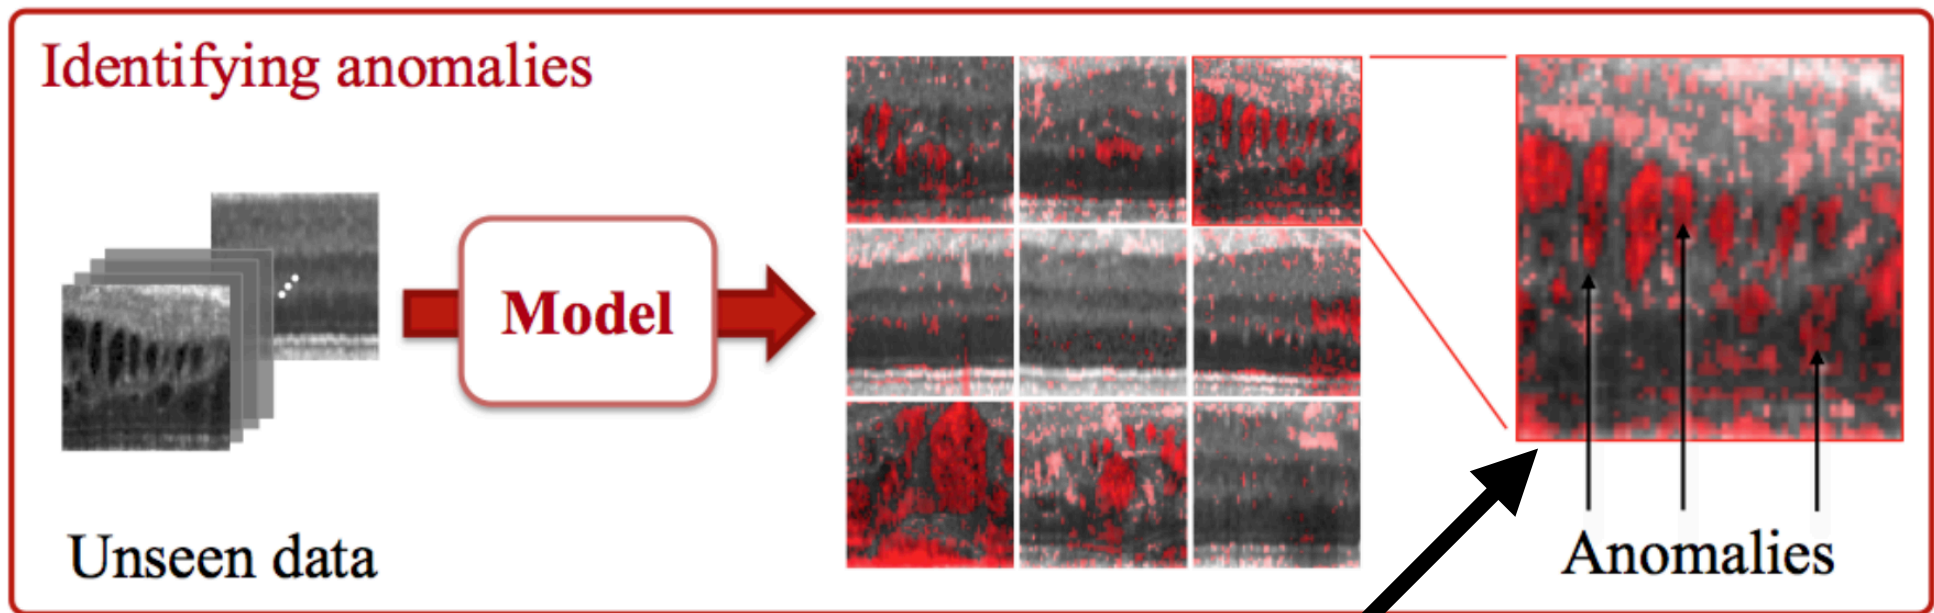
$$A(\mathbf{x}) = (1 - \lambda) \cdot R(\mathbf{x}) + \lambda \cdot D(\mathbf{x}),$$



Fig. 2. (a) Deep convolutional generative adversarial network. (b) t-SNE embedding of normal (blue) and anomalous (red) images on the feature representation of the last convolution layer (orange in (a)) of the discriminator.

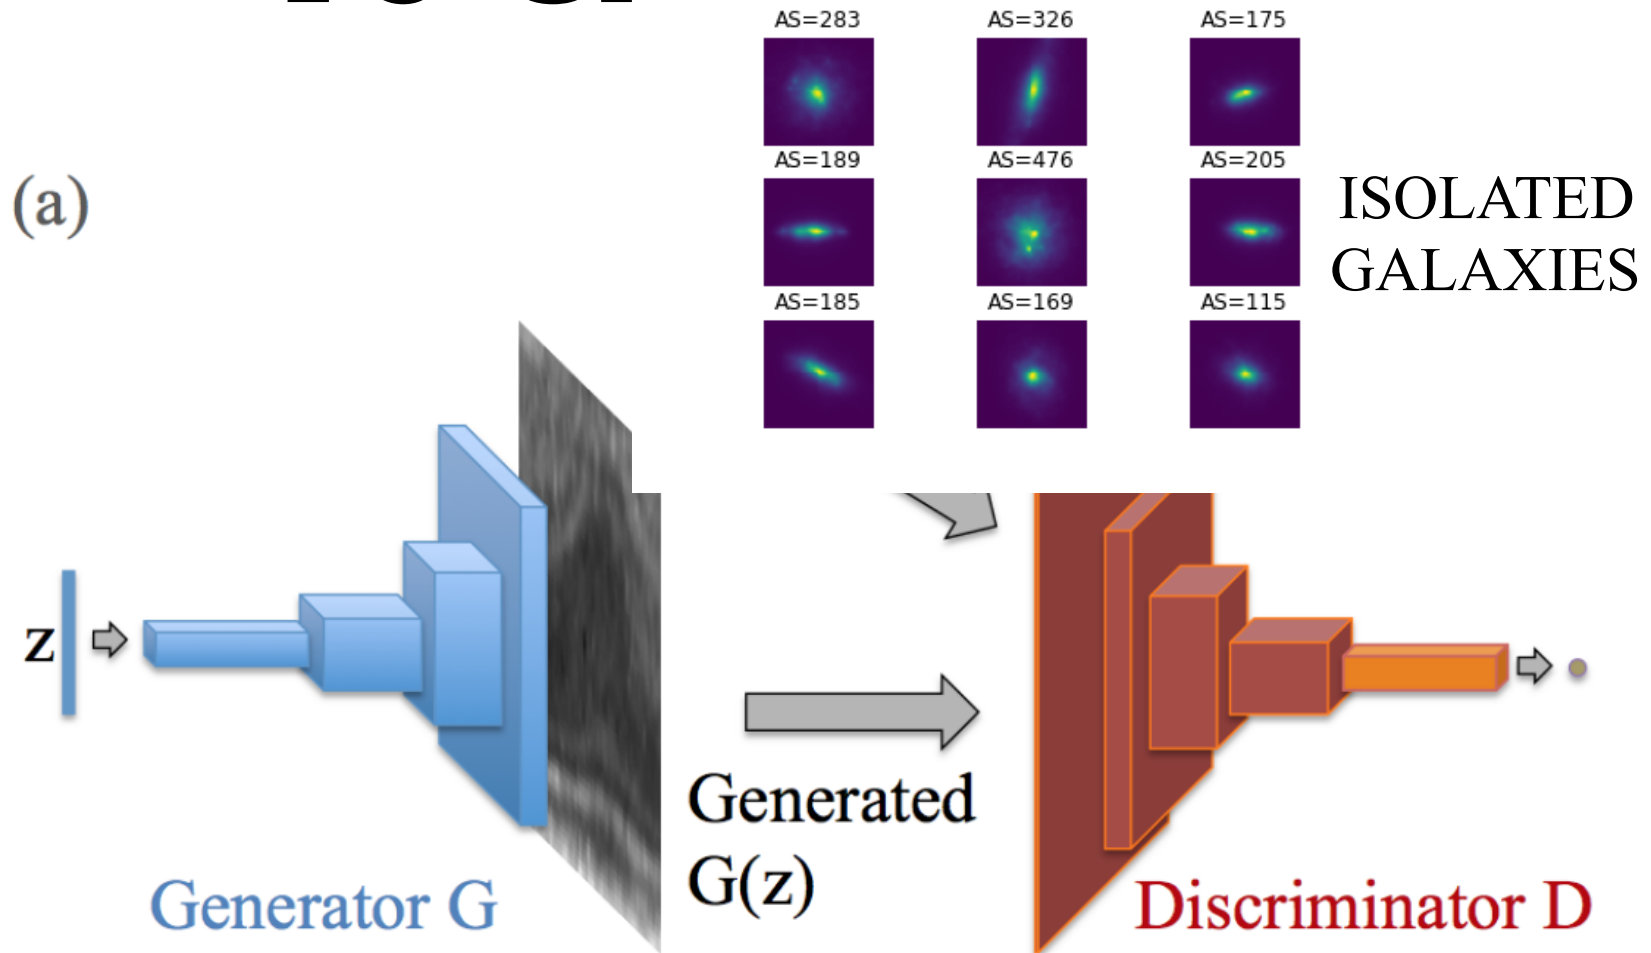
$$\mathbf{x}_R = |\mathbf{x} - G(\mathbf{z}_\Gamma)|$$

DIFFERENCE BETWEEN GENERATED FROM LATENT SPACE
AND INPUT

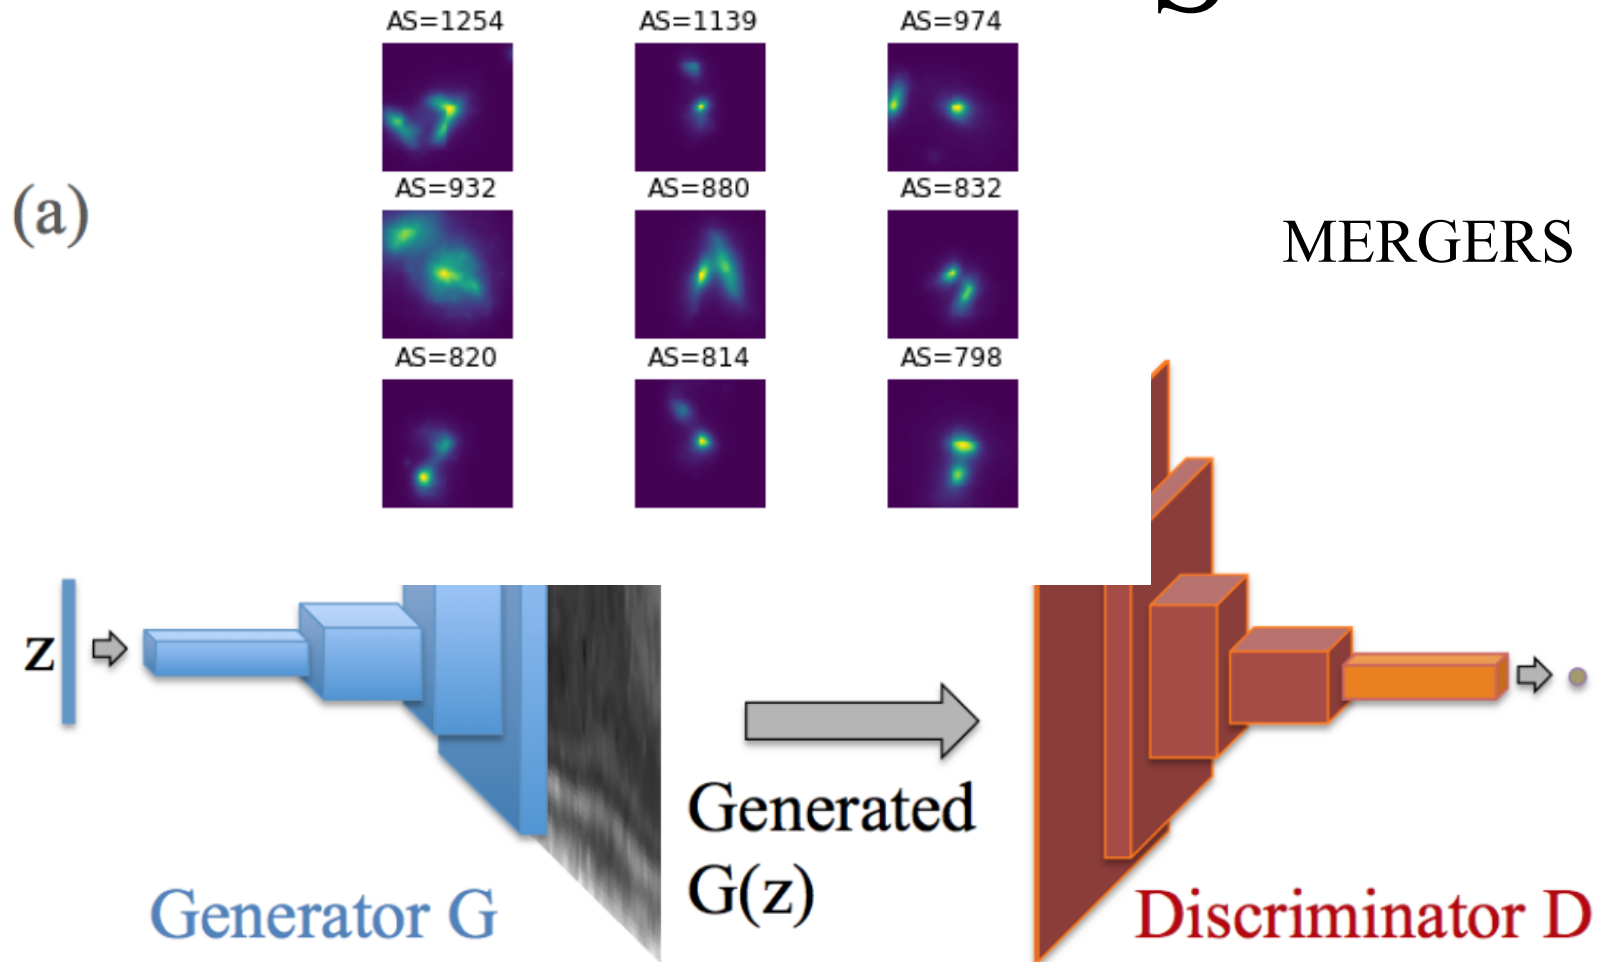


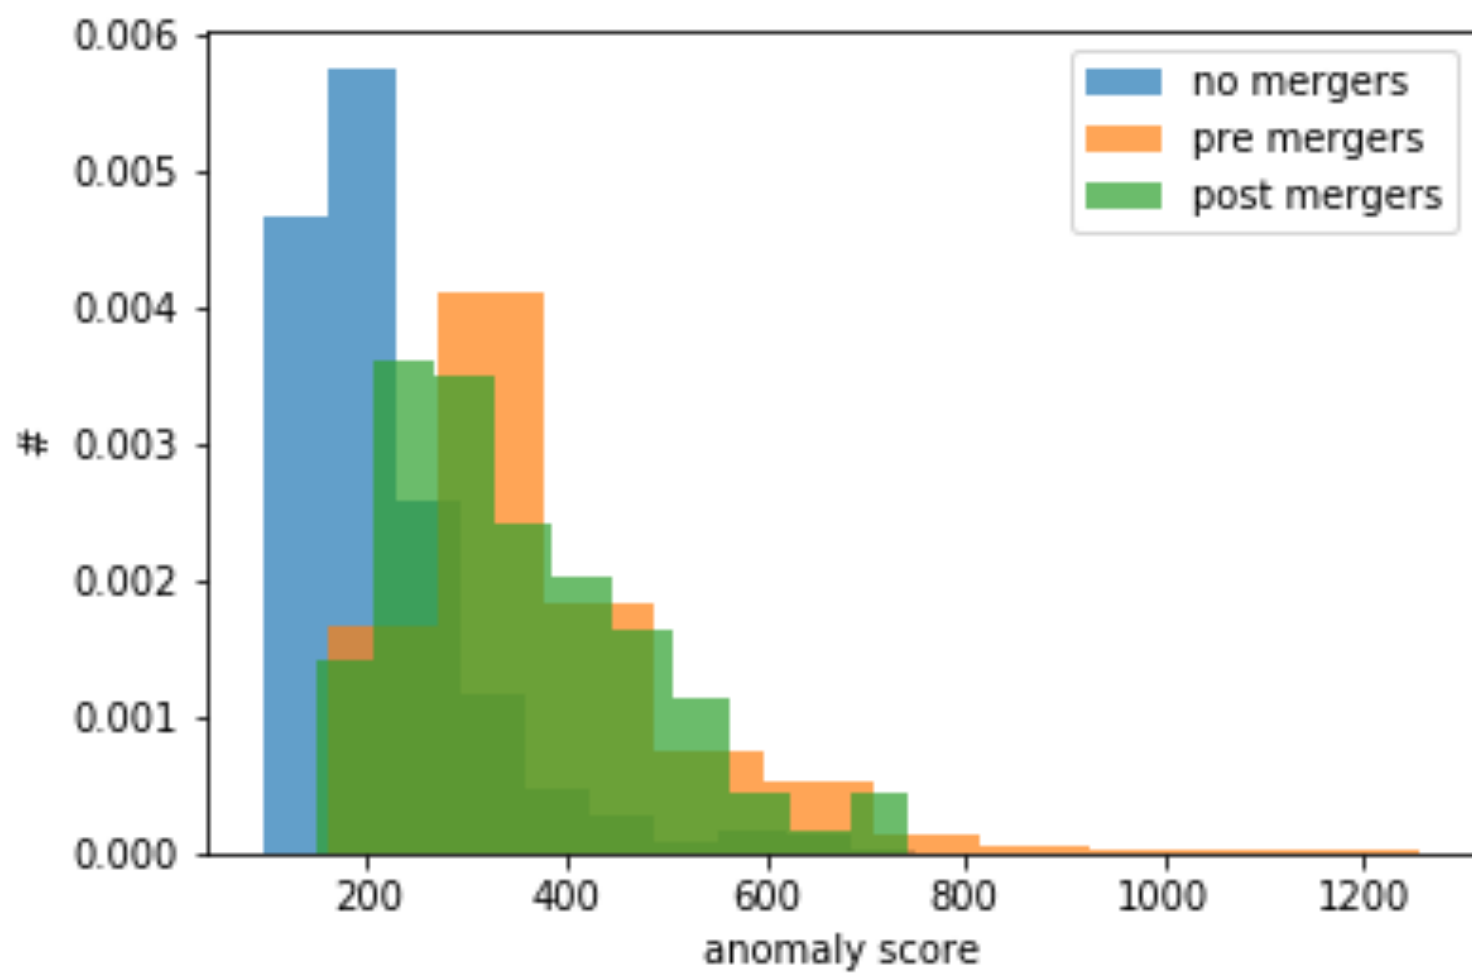
IDENTIFY THE ANOMALOUS REGIONS IN THE IMAGE

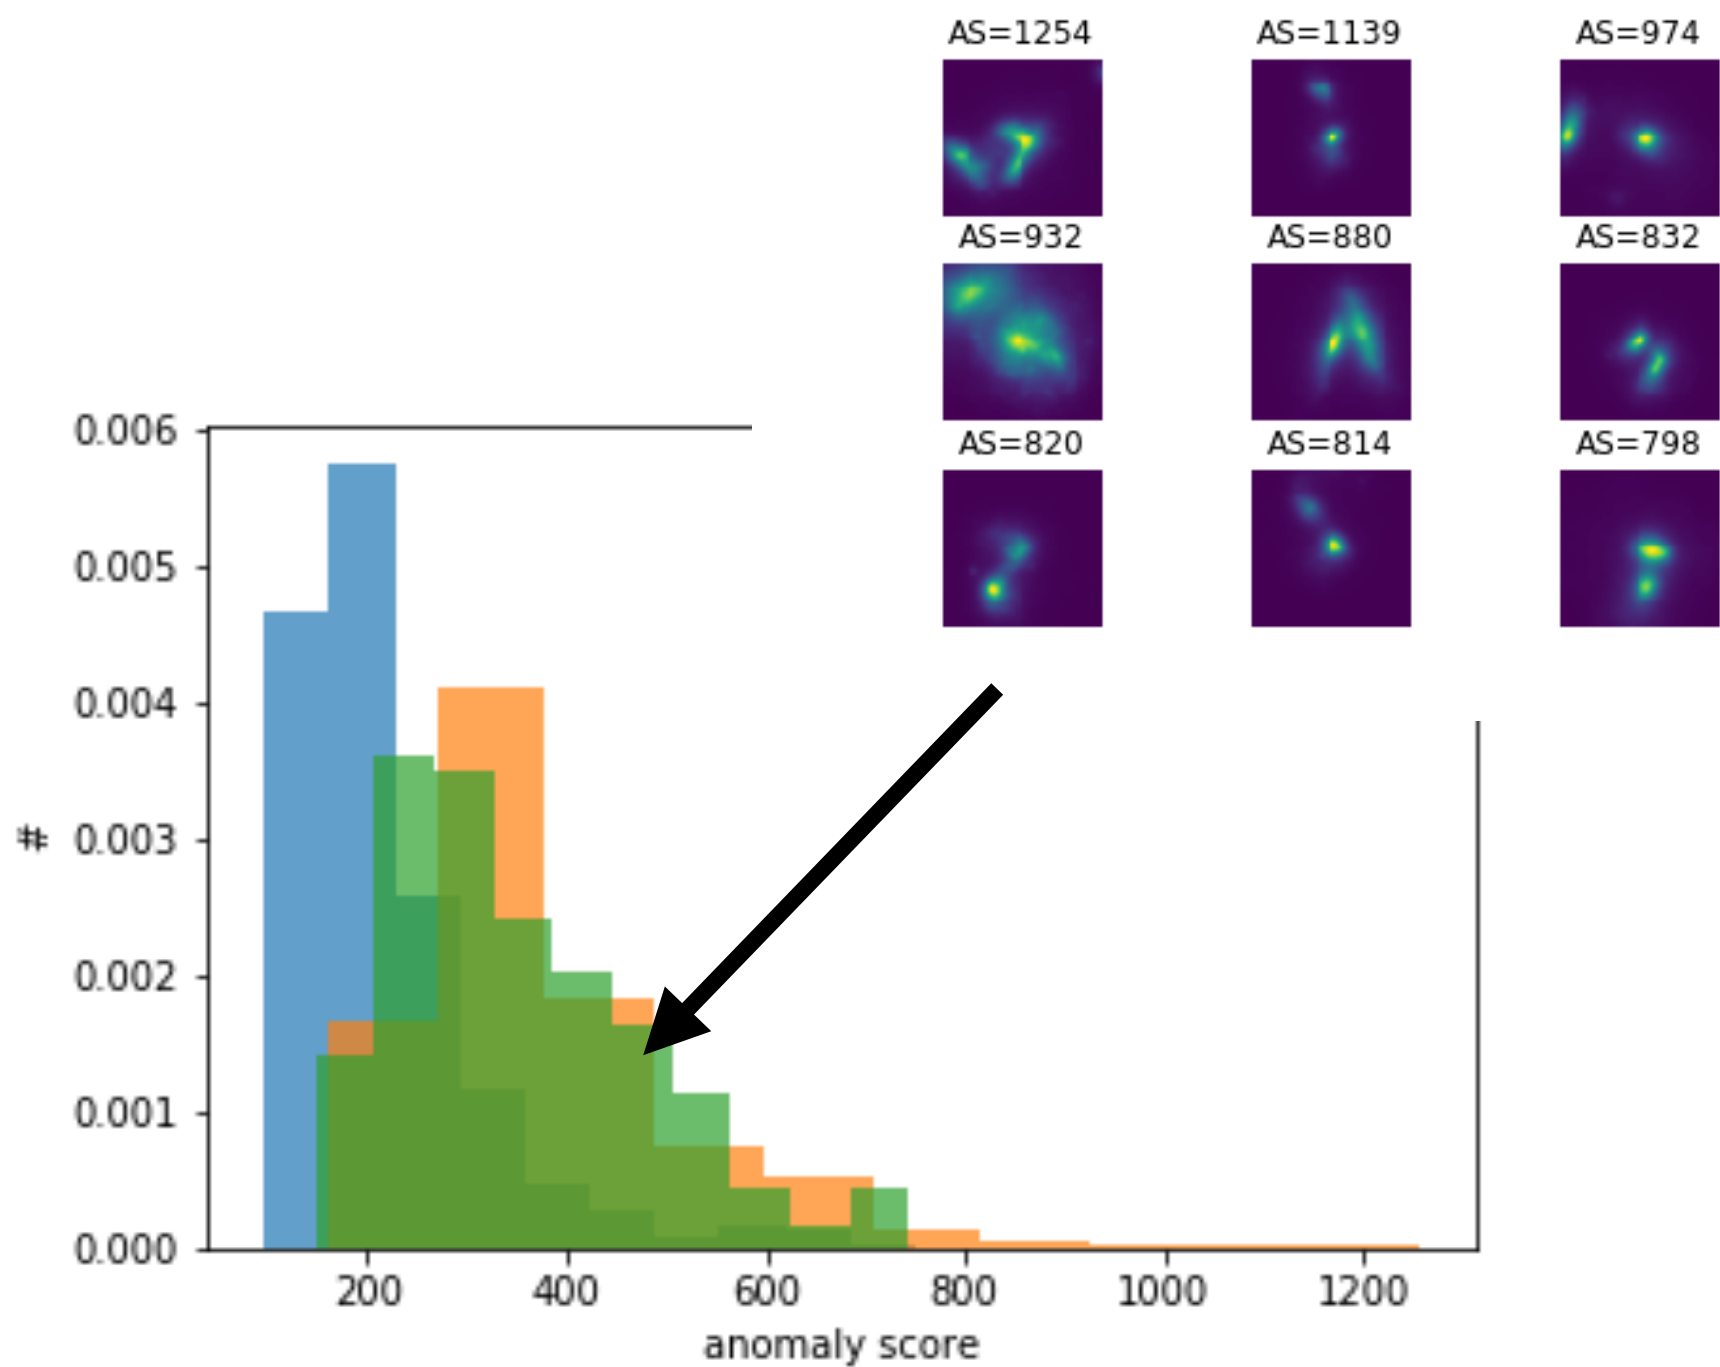
AN EXAMPLE APPLIED TO GALAXIES



AN EXAMPLE APPLIED TO GALAXIES







PART V: SOME PRACTICAL CONSIDERATIONS

HOW DO I KNOW THAT I HAVE
REACHED CONVERGENCE?

FOR HOW MANY EPOCHS DO
I HAVE TO TRAIN?

THERE IS NO “MAGIC RULE” TO MY KNOWLEDGE.

NORMALLY A DECISION IS TAKEN BASED
ON THE MONITORING OF THE VALIDATION LOSS

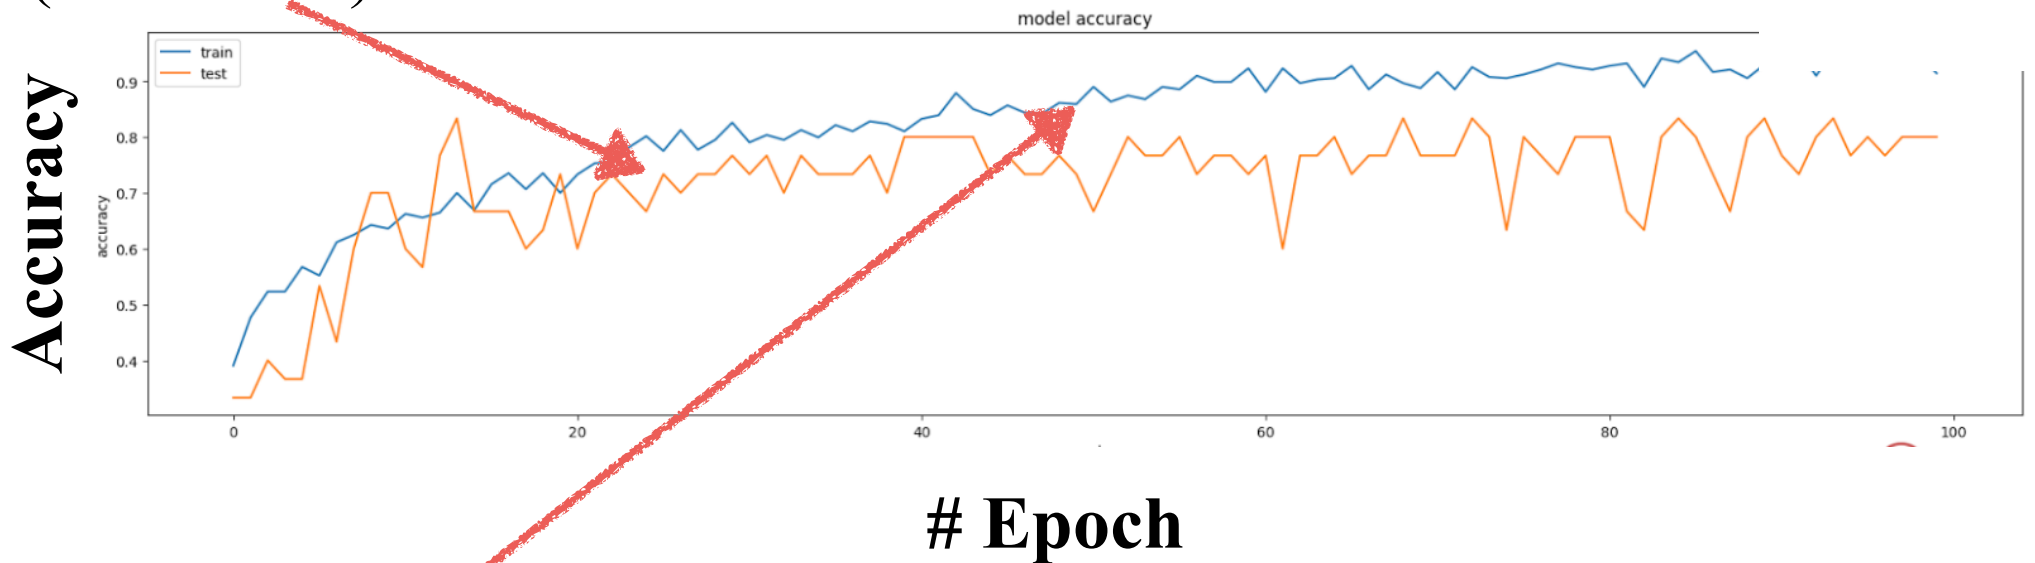
... “IF THE VALIDATION LOSS DOES NOT CHANGE MORE
THAN EPSILON OVER THE LAST N EPOCHS”...

KEEPING TRACK OF PERFORMANCE DURING TRAINING

THE LEARNING HISTORY

[MODEL ACCURACY AS A FUNCTION OF EPOCH]

(validation)



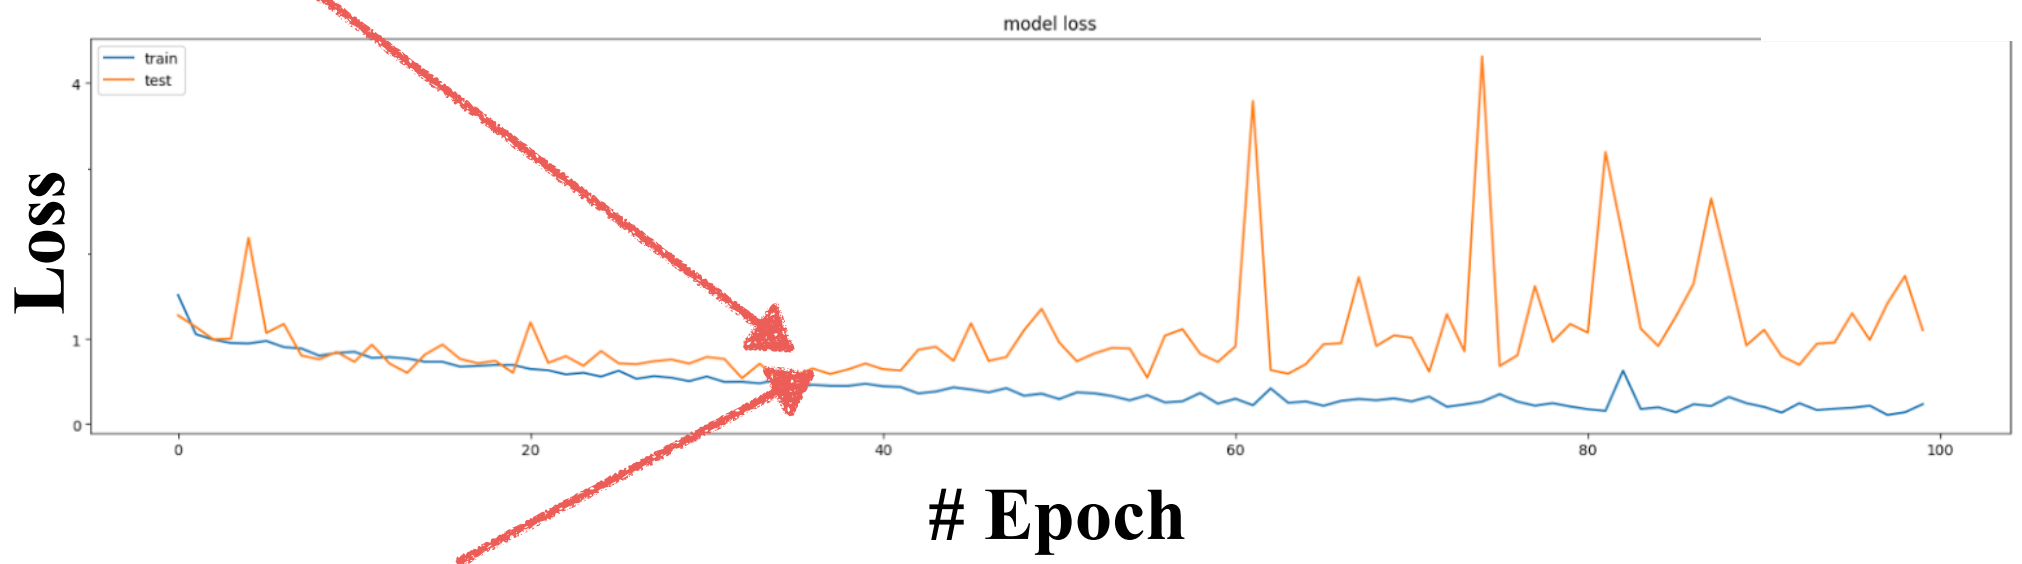
(training)

KEEPING TRACK OF PERFORMANCE DURING TRAINING

THE LEARNING HISTORY

[MODEL LOSS AS A FUNCTION OF EPOCH]

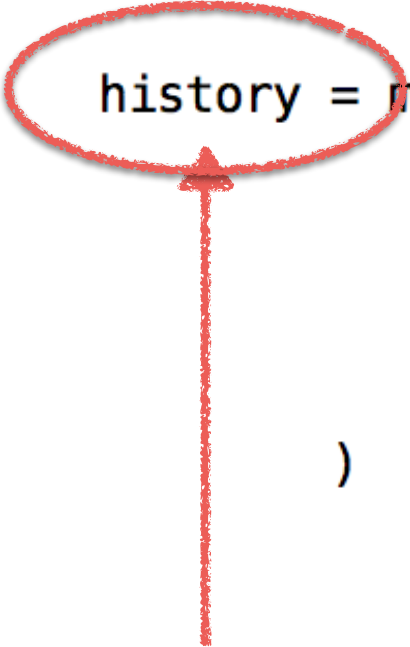
(validation)



(training)

IMPLEMENTATION IN KERAS

```
history = model.fit_generator(  
    datagen.flow(X_train, Y_train, batch_size=batch_size),  
    samples_per_epoch=X_train.shape[0],  
    nb_epoch=nb_epoch,  
    validation_data=(X_val, Y_val),  
    callbacks=[earlystopping, modelcheckpoint]  
)
```



THE TRAINING RETURNS A **HISTORY** OBJECT

IMPLEMENTATION IN KERAS

THE TRAINING RETURNS A **HISTORY** OBJECT

```
from keras.models import Sequential

print(history.history.keys())
['acc', 'loss', 'val_acc', 'val_loss']

plt.plot(history.history['acc'])
plt.plot(history.history['val_acc'])

plt.plot(history.history['loss'])
plt.plot(history.history['val_loss'])
```

“LIVE” LEARNING HISTORY

TensorBoard SCALARS IMAGES GRAPHS > INACTIVE

Show data download links
 Ignore outliers in chart scaling
Tooltip sorting method: default

Smoothing 0.6

Horizontal Axis **STEP** RELATIVE WALL

Runs
Write a regex to filter runs

train
 eval

TOGGLE ALL RUNS

Filter tags (regular expressions supported)

accuracy 1
cross entropy 1

cross entropy

run to download CSV JSON

Name	Smoothed Value	Value	Step	Time	Relative
eval	0.02591	0.02550	170.0	Mon Sep 12, 15:40:41	8s
train	0.02851	0.03362	166.0	Mon Sep 12, 15:40:40	7s

EARLY STOPPING IN KERAS

set patience (number of epochs to monitor)

```
from keras.callbacks import EarlyStopping
```

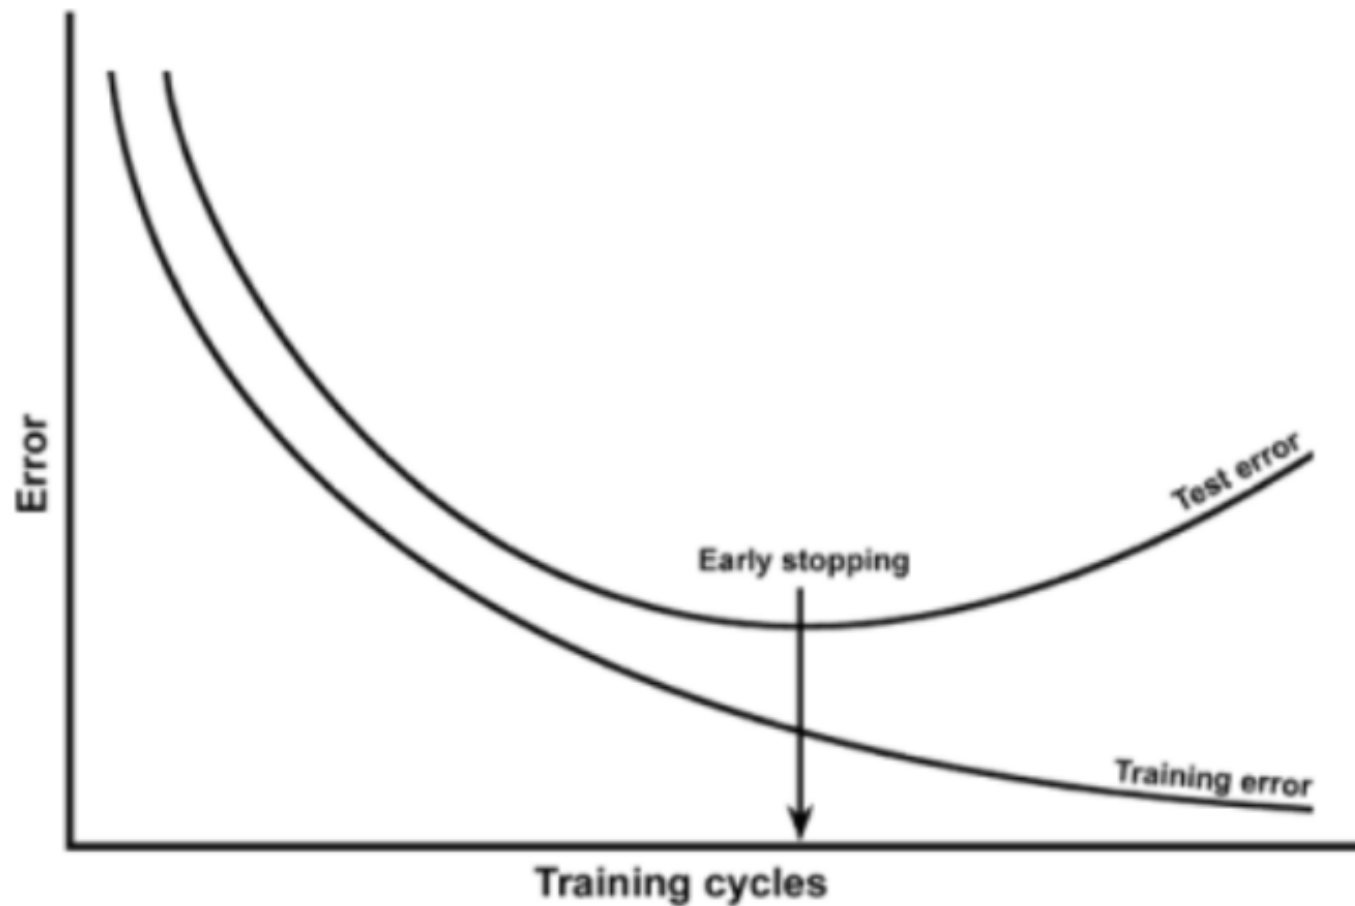
```
patience_par=10
```

```
earlystopping = EarlyStopping( monitor='val_loss',patience =  
patience_par,verbose=0,mode='auto' )
```

```
history = model.fit(X, Y, validation_split=0.33, epochs=150, batch_size=10,  
verbose=0, callbacks=[earlystopping])
```

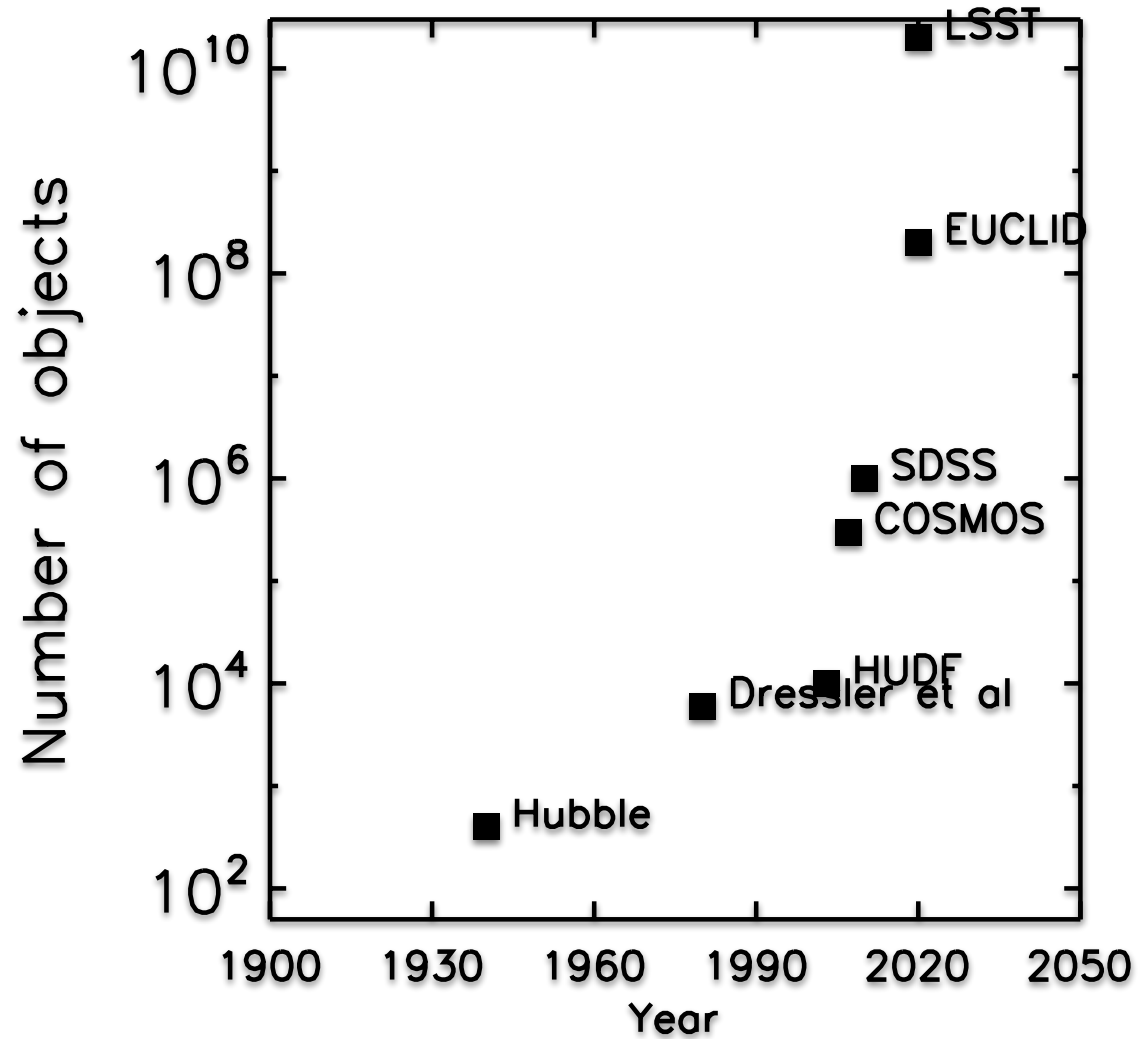
parameter to monitor [typically
validation loss]

EARLY STOPPING HELPS ALSO PREVENTING OVERFITTING...



HOW LARGE DOES MY DATASET NEED TO BE?

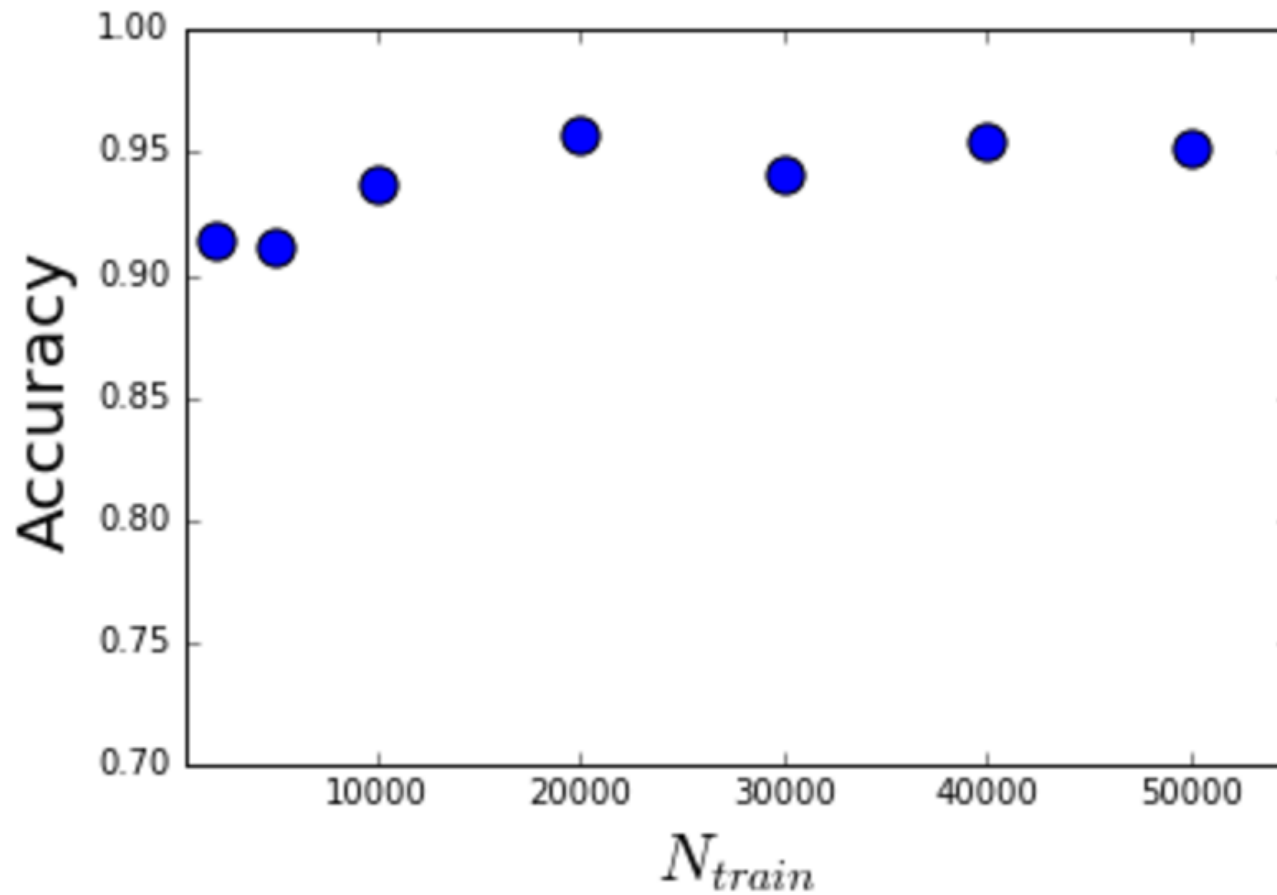
DOES THAT MEAN THAT IF I DO NOT HAVE MILLIONS
OF LABELLED EXAMPLES ALL THIS IS USELESS?



THE SIZE OF DATASETS
IS
INCREASING FAST BUT
WE DO
NOT ALWAYS HAVE
MILLIONS OF
LABELLED
EXAMPLES TO TRAIN

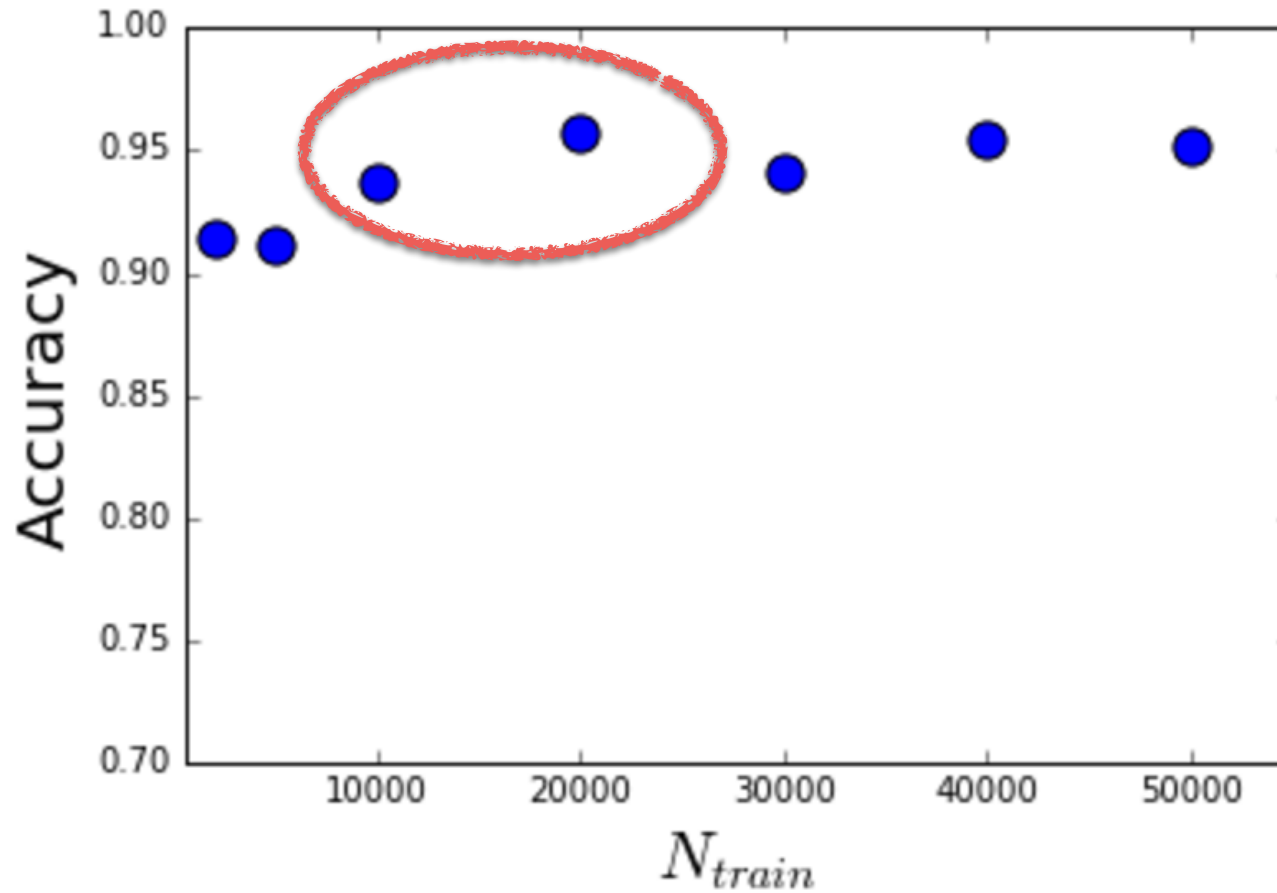
...

NOT REALLY...



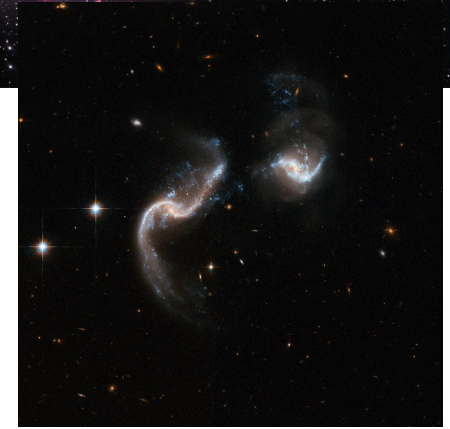
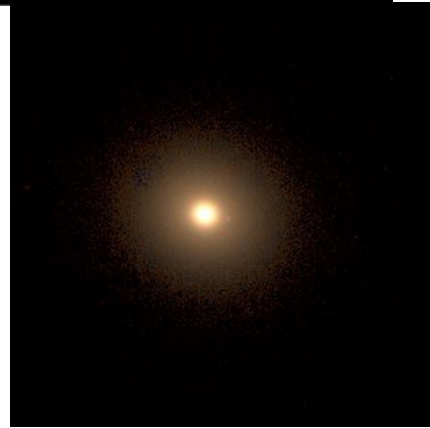
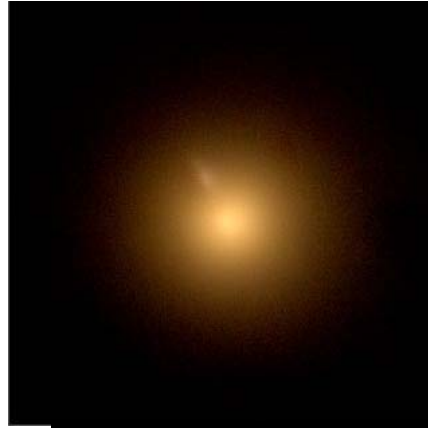
ACCURACY OF MORPHOLOGICAL CLASSIFICATIONS
OF GALAXIES AS A FUNCTION OF THE SIZE OF THE
TRAINING SET

NOT REALLY...



ACCURACY OF MORPHOLOGICAL CLASSIFICATIONS
OF GALAXIES AS A FUNCTION OF THE SIZE OF THE
TRAINING SET

WHY?



WHY?



ship



dog



deer



bird



ship



cat



dog



dog



horse



horse



ship



frog



bird



ship



bird



cat



automobile



ship



deer



truck



dog



deer



automobile



horse



WHY?

IMAGES USED IN
ASTROPHYSICS ARE
SIMPLE...



ship



dog



deer



bird



ship



cat



dog



dog



horse



horse



ship



frog



bird



ship



bird



cat



automobile



ship



deer



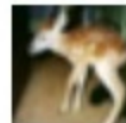
truck



dog



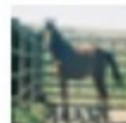
deer



automobile



horse



SOME TRICKS FOR “SMALL” DATASETS

DATA AUGMENTATION

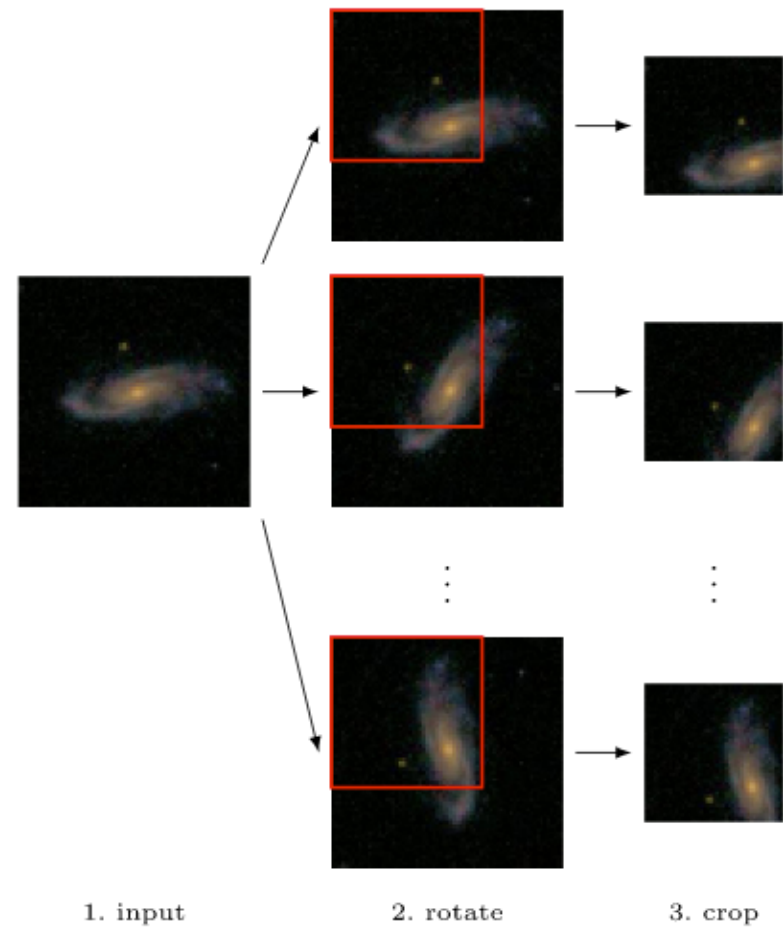
ANOTHER WAY TO REDUCE OVER-FITTING IS TO
“AUGMENT” THE SIZE OF THE DATASET AVAILABLE FOR
TRAINING

FOR MANY APPLICATIONS THE CLASSIFICATION SHOULD

BE INDEPENDENT TO:

- TRANSALTIONS
- ROTATIONS
- SCALINGS
- ETC...

DATA



CLASSIFICATION

Dieleman+15

FOR MANY APPLICATIONS THE CLASSIFICATION SHOULD BE INDEPENDENT TO:

- TRANSLATIONS
- ROTATIONS
- SCALINGS
- ETC...

DATA AUGMENTATION



FOR MANY APPLICATIONS THE CLASSIFICATION SHOULD
BE INDEPENDENT TO:

- TRANSLATIONS
- ROTATIONS
- SCALINGS
- ETC...

DATA AUGMENTATION

CAN BE DONE “ON THE FLY” DURING THE TRAINING PHASE

KERAS IMPLEMENTATION:

DEFINES
THE RANGE OF
PERTURBATIONS
APPLIED TO
IMAGES DURING
TRAINING

```
datagen = ImageDataGenerator(  
    featurewise_center=False,  
    samplewise_center=False,  
    featurewise_std_normalization=False,  
    samplewise_std_normalization=False,  
    zca_whitening=False,  
    rotation_range=45,  
    width_shift_range=0.05,  
    height_shift_range=0.05,  
    horizontal_flip=True,  
    vertical_flip=True,  
    zoom_range=[0.75,1.3])
```

DATA AUGMENTATION

SOMETIMES ADDING NOISE INCREASES THE CLASSIFICATION / REGRESSION ACCURACY!

STANDARD DEVIATION OF GAUSSIAN NOISE

```
model = Sequential()  
model.add(GaussianNoise(0.01, input_shape=(img_channels, img_rows, img_cols)))  
# ...  
model.compile(optimizer='adam', loss='categorical_crossentropy', metrics=['accuracy'])
```

A LAYER OF GAUSSIAN NOISE ADDED IN THE MODEL

DOMAIN ADAPTATION (or knowledge transfer)

THE CONVOLUTIONAL PART OF A CNN IS
A FEATURE EXTRACTOR

DOMAIN ADAPTATION (or knowledge transfer)

THE CONVOLUTIONAL PART OF A CNN IS
A FEATURE EXTRACTOR

IN THAT RESPECT, THEY ARE VERY FLEXIBLE ...

DOMAIN ADAPTATION (or knowledge transfer)

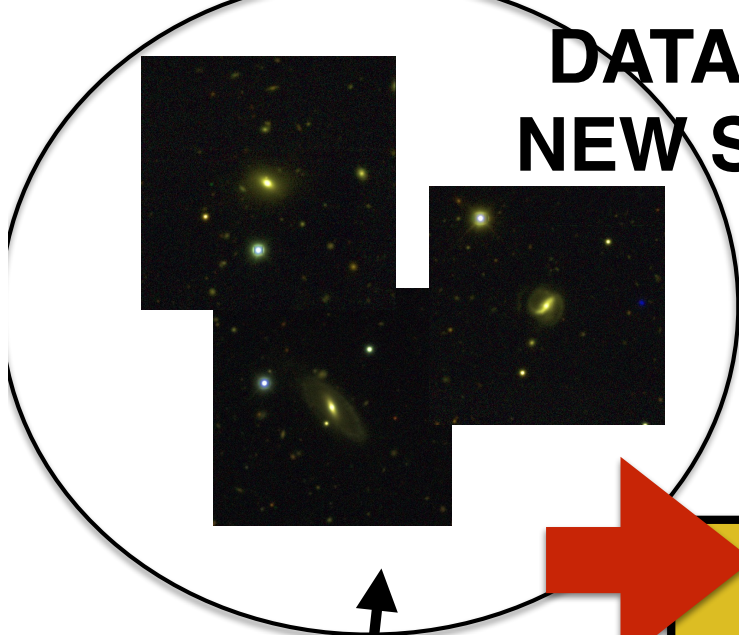
EVEN IF OUR TRAINING SET IS NOT SO LARGE ...

WE CAN USE A CNN PRE-TRAINED ON A LARGER SAMPLE

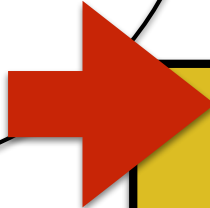
DEPENDING ON HOW SIMILAR BOTH DATASETS ARE, WE
CAN:

- RECYCLE THE SAME FEATURES
- FINE-TUNING THE WEIGHTS

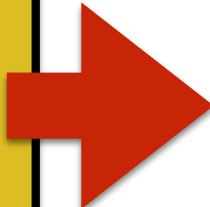
**DATA FROM
NEW SURVEY**



How robust to different datasets?
Do we always need a big training
set?



**DEEP-LEARNING
BASED
MACHINE**

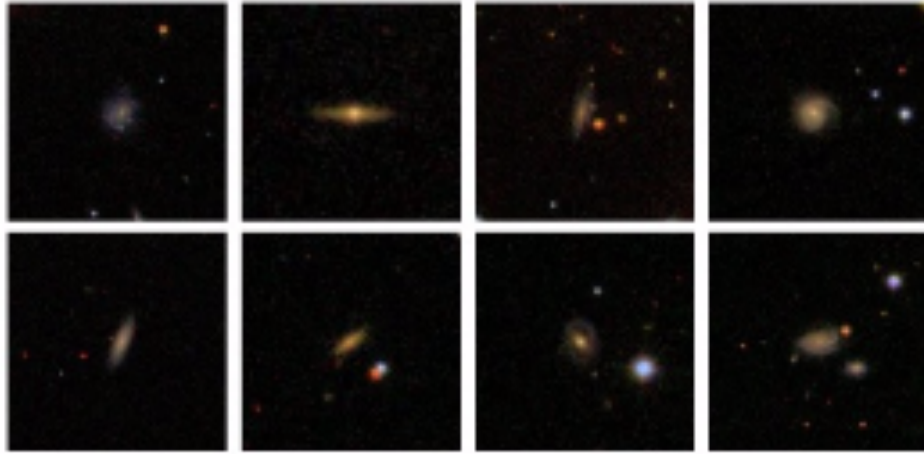


“Improved”
Galaxy ZOO like
classifications for
for the entire
sample

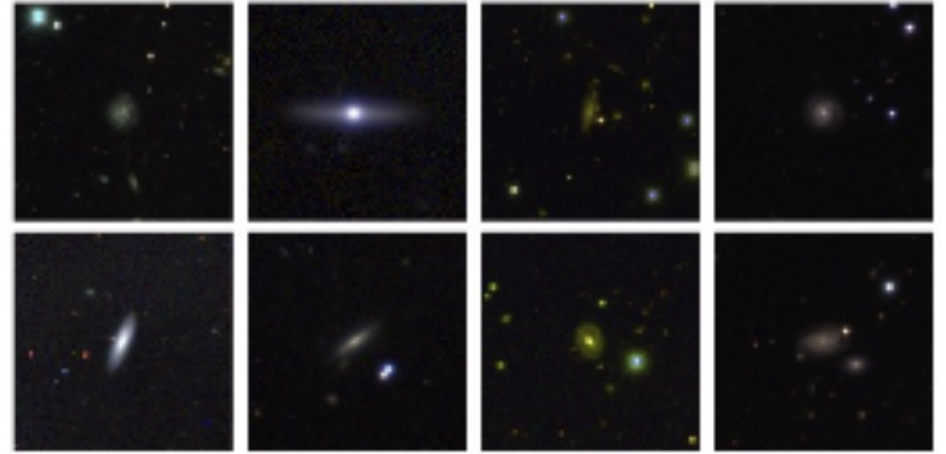
Transfer knowledge?



**Human classifications
from existing survey**

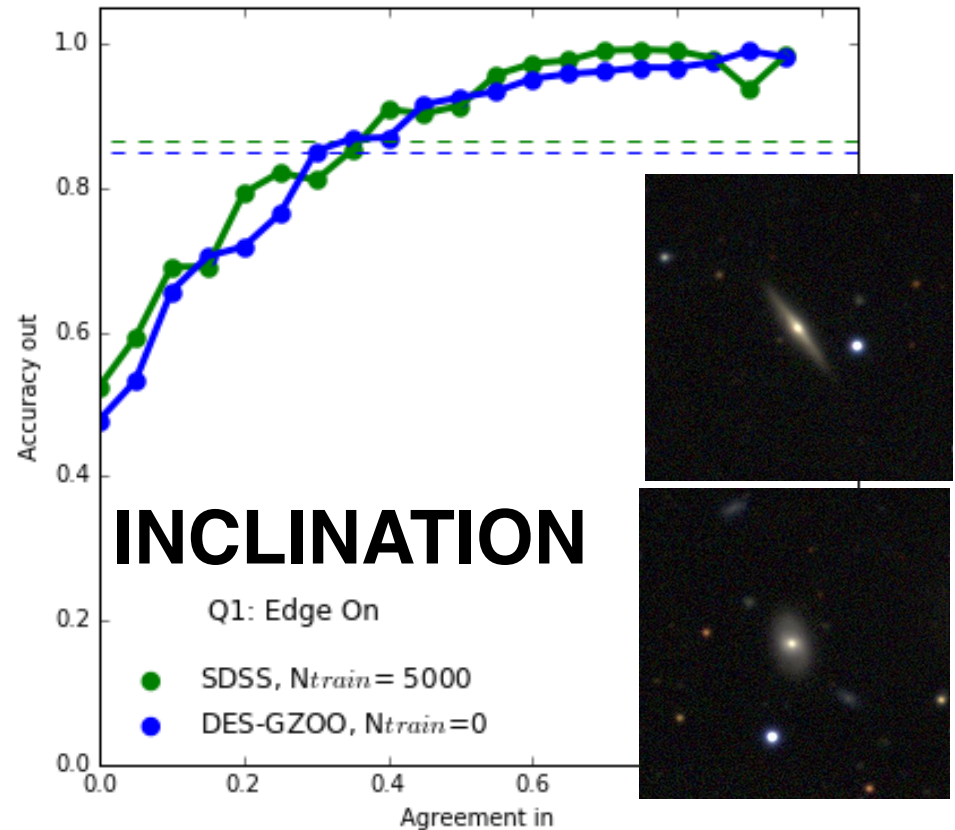
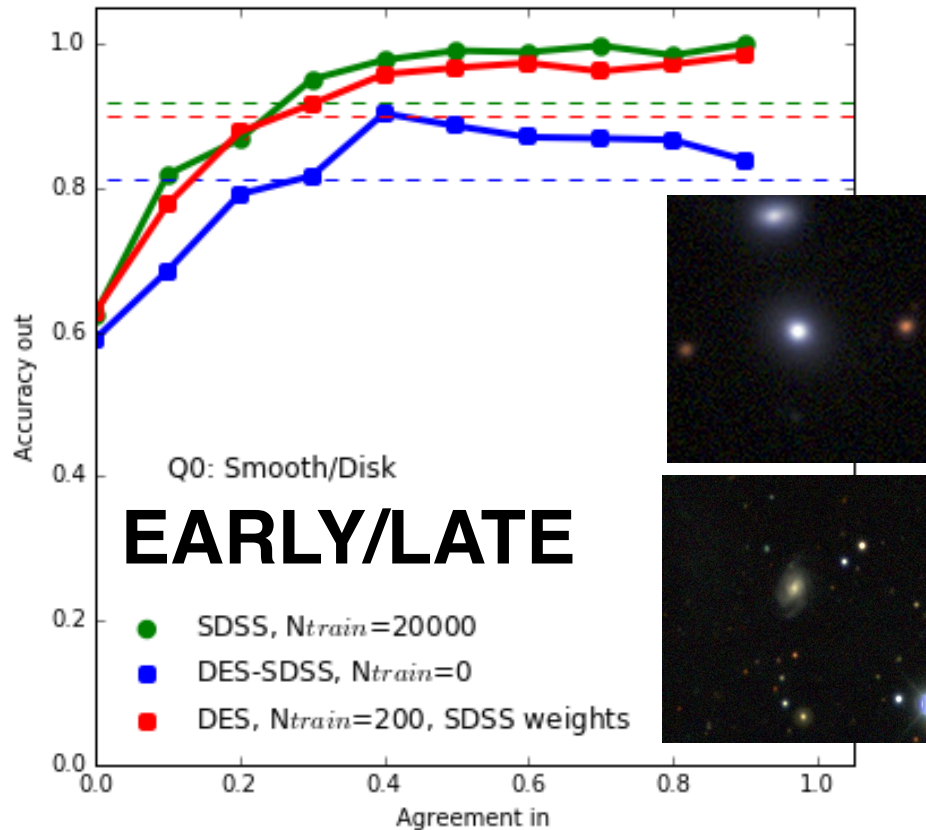


SDSS



DES

Knowledge Transfer from SDSS to DES



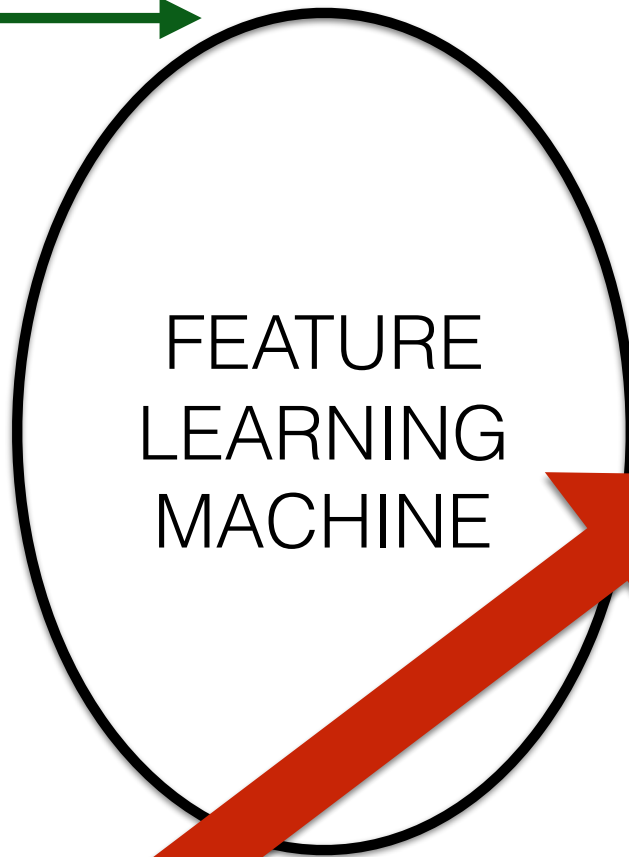
Only 200 (1%!) objects classified in DES are needed to reach an accuracy $>90\%$ if a machine trained on the SDSS is used

For some properties, i.e. EDGE-ON galaxies. No training at all is needed to go from SDSS to DES

TRAINING:
simulations of
analytic profiles
with PSF, noise
effects

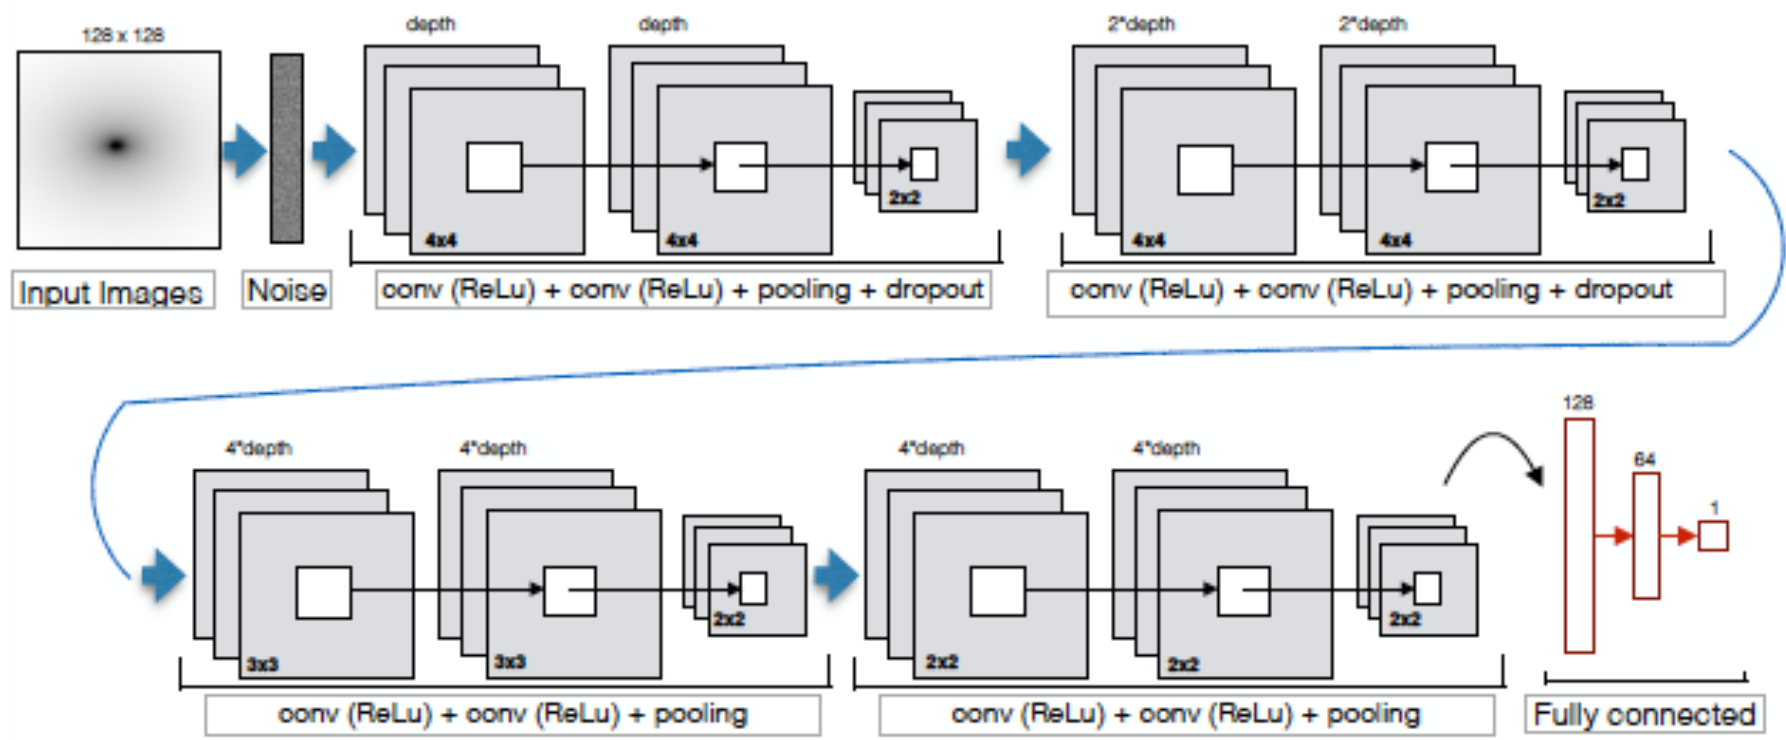
**(no limits
on size)**

**DOMAIN
ADAPTING
FROM
SIMULATIONS**



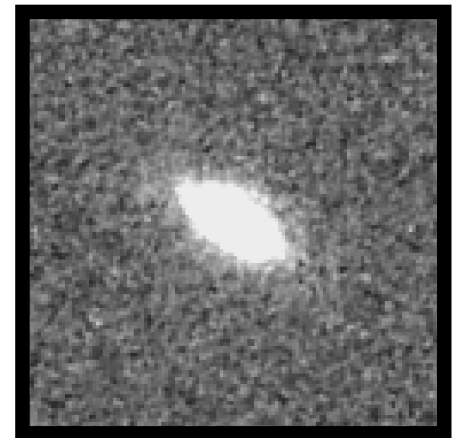
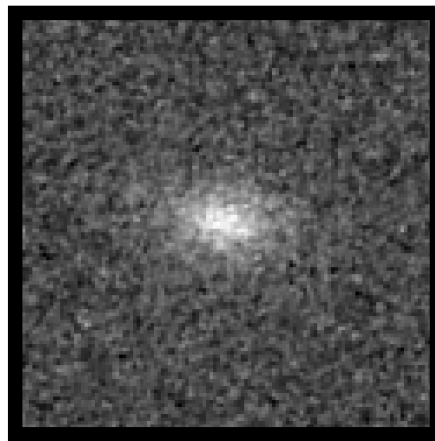
DATA:
HST deep field
observations
CANDELS

- Flux
- Sersic Index
- Radii
- b/a
-
-
-



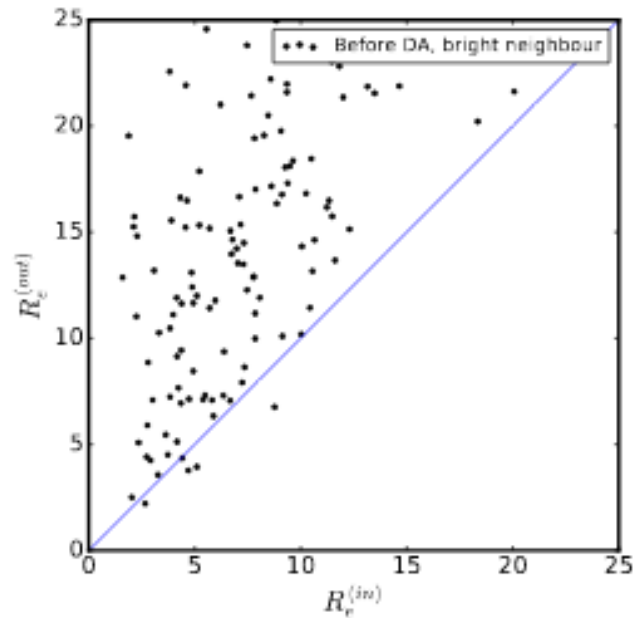
Standard analytic profiles

- 100.000-300.000 galaxies [GALSIM]
 - Real HST background added + PSF (F160)
 - Random distribution of parameters (uniform):
 - $18 < \text{Mag} < 24$, $0 < \text{BT} < 1$, $< \text{Nb} <$, $\text{Nd} = 1$, $0.2 < \log(\text{rb}) < 1.3$,
 $0.2 < \log(\text{rd}) < 1.5$, $0.05 < \text{eb} < 0.95$, $0.05 < \text{ed} < 0.95$, $0 < \text{PA} < 180$
 - 64*64 stamps
 - **FULLY IDEALISTIC -
NO COMPANIONS
NO IRREGULARS
NO CLUMPY!**



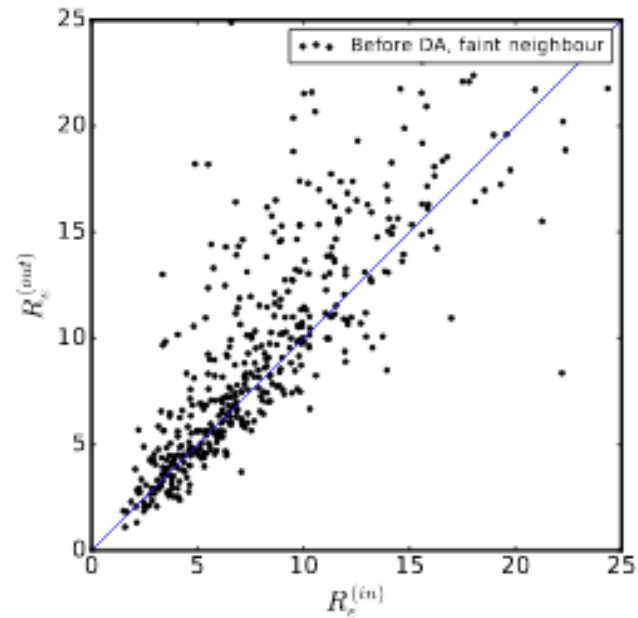
MORPHOMETRY OF REAL GALAXIES TRAINED ON ANALYTIC PROFILES

BRIGHT NEIGH.



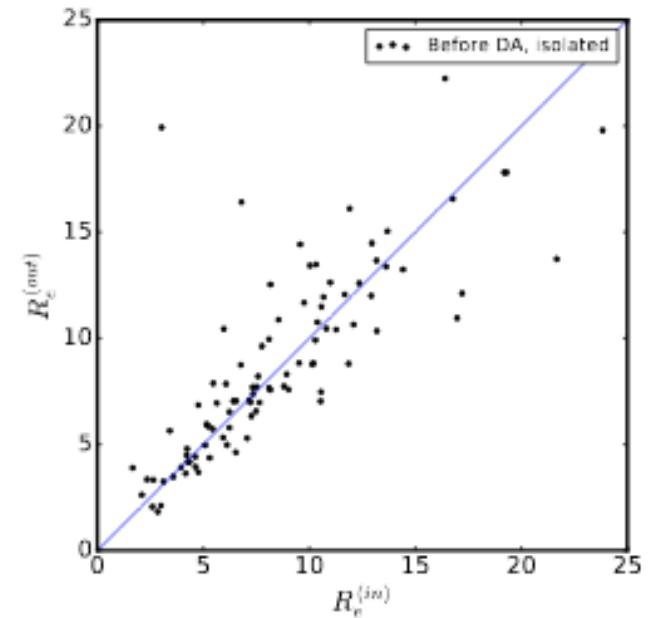
(c) BDA bright neighbours

FAINT NEIGH.



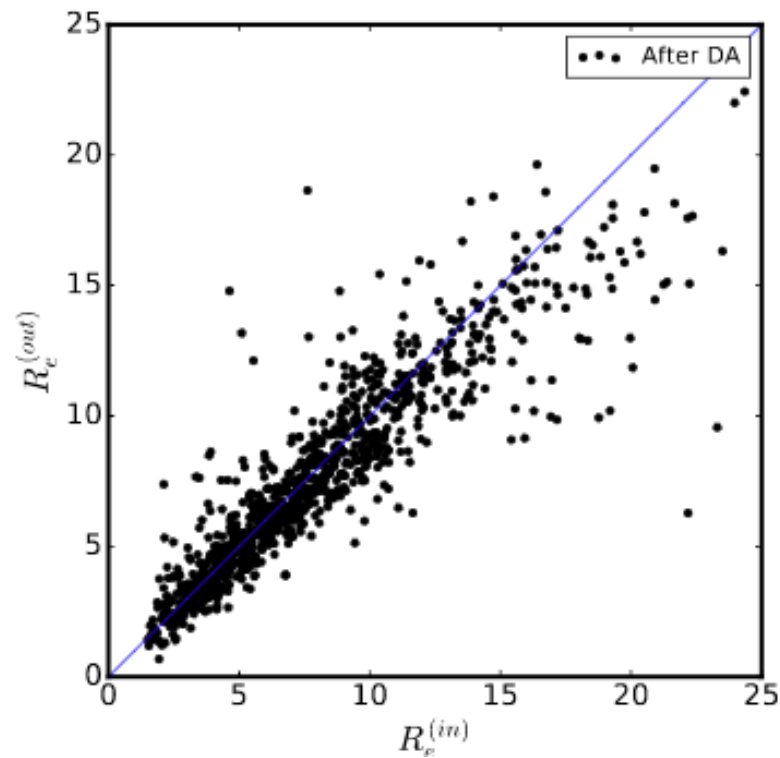
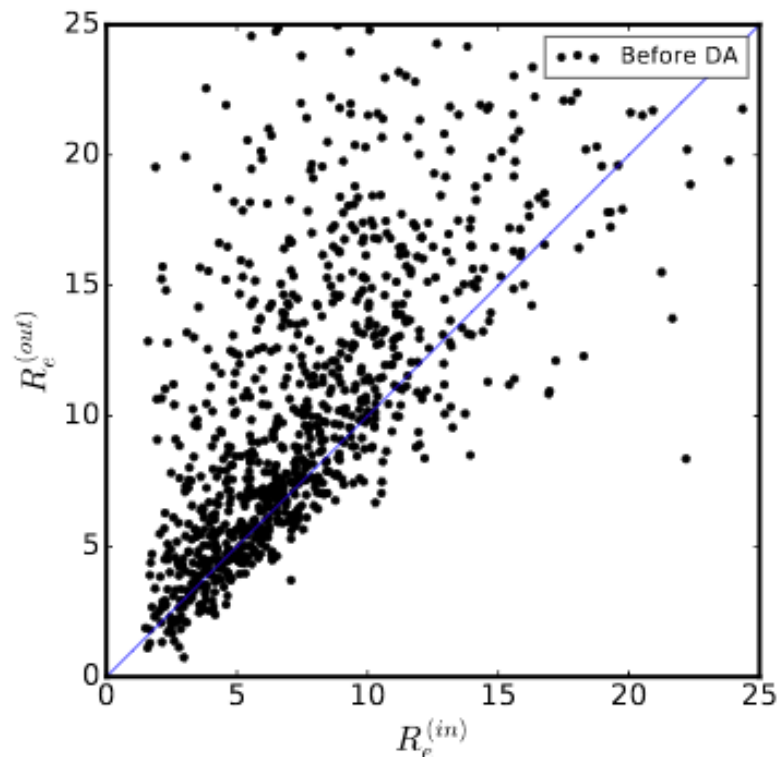
(d) BDA faint neighbours

ISOLATED

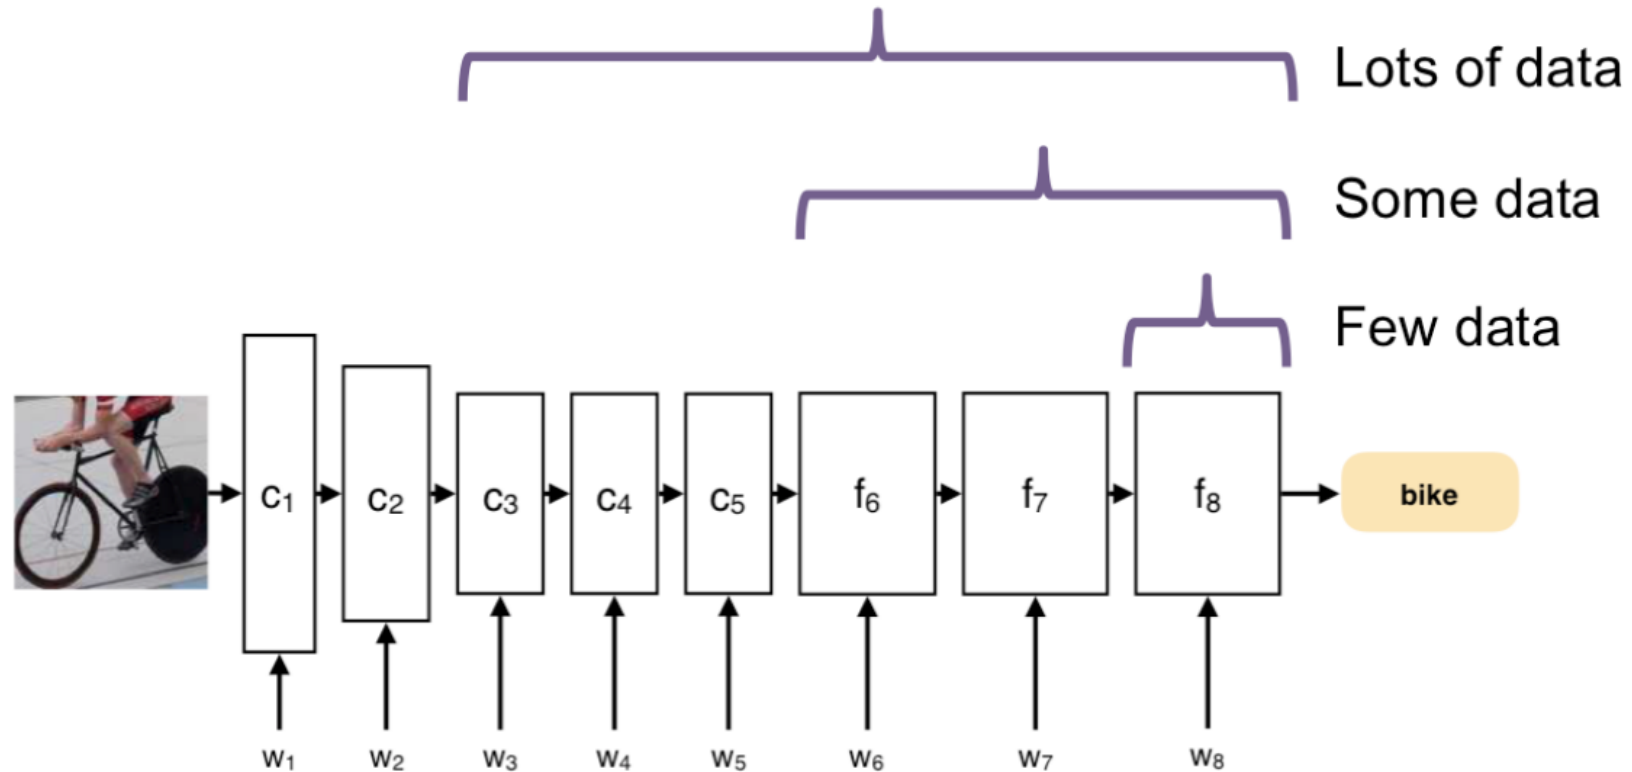


(e) BDA isolated galaxies

DOMAIN ADAPTATION: 0.1% OF “REAL” GALAXIES



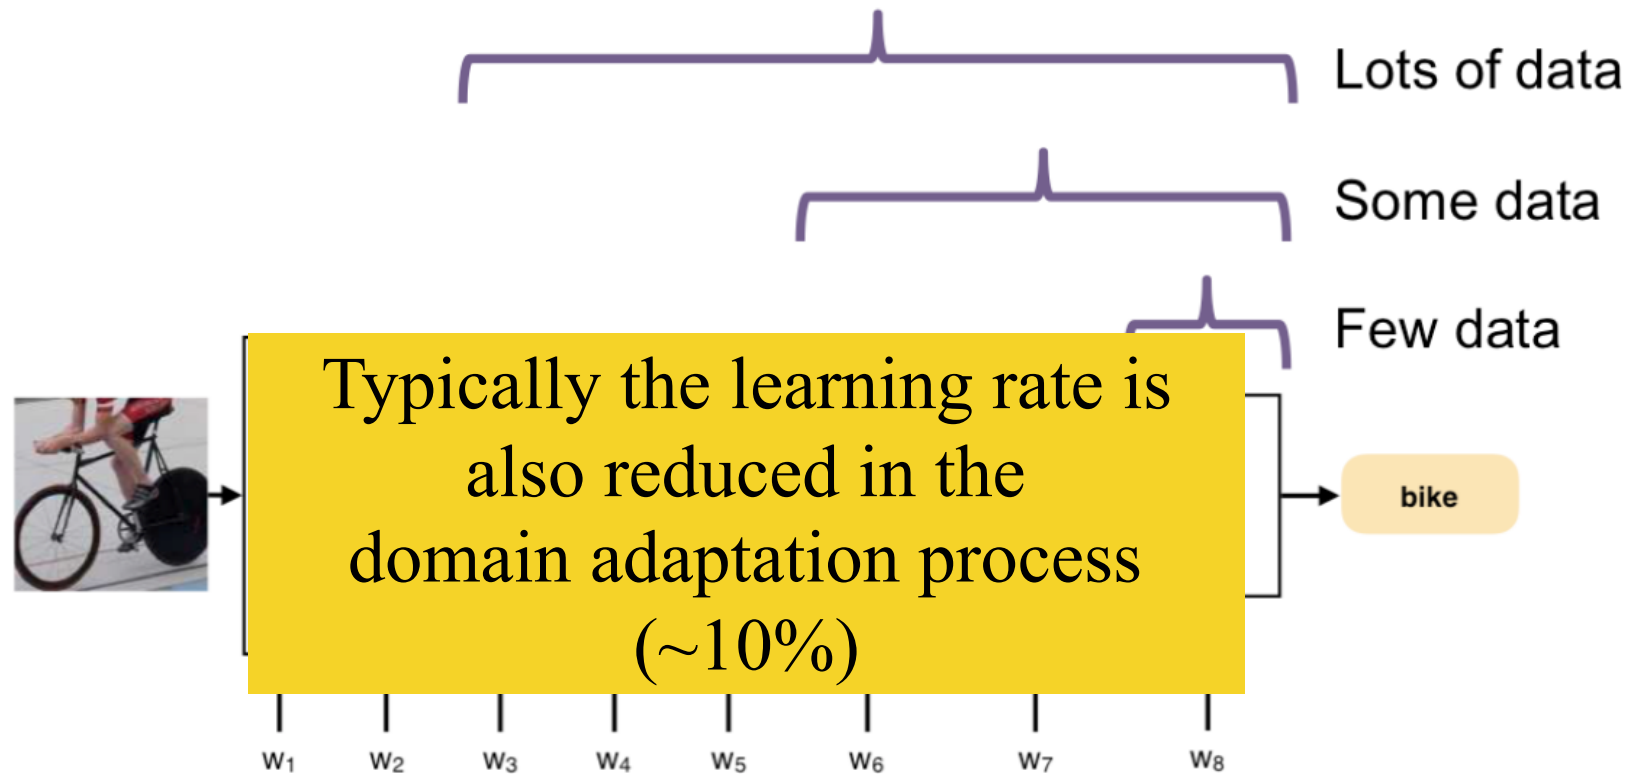
HOW MANY LAYERS “DOMAIN ADAPT” ?



DEPENDING ON HOW MUCH SIMILAR BOTH DATASETS ARE:

- FINE-TUNE ONLY FULLY CONNECTED LAYERS
- FINE-TUNE A FEW LAYERS
- FINE-TUNE ALL LAYERS

HOW MANY LAYERS “DOMAIN ADAPT” ?



DEPENDING ON HOW MUCH SIMILAR BOTH DATASETS ARE:

- FINE-TUNE ONLY FULLY CONNECTED LAYERS
- FINE-TUNE A FEW LAYERS
- FINE-TUNE ALL LAYERS

KERAS IMPLEMENTATION

“FULL” DOMAIN
ADAPTATION

learning
rate

```
model.save_weights(test_name+".hd5", overwrite=True)
```

```
keras.optimizers.Adam(lr=0.001, beta_1=0.9, beta_2=0.999, epsilon=None, decay=0.0, amsgrad=False)
```

```
model.load_weights(model_name+".hd5")
```

start training

load weights
of previous
training

```
history = model.fit_generator(  
    datagen.flow(X_train, Y_train, batch_size=batch_size),  
    samples_per_epoch=X_train.shape[0],  
    nb_epoch=nb_epoch,  
    validation_data=(X_val, Y_val),  
    callbacks=[earlystopping, modelcheckpoint]  
)
```

KERAS IMPLEMENTATION

“PARTIAL” DOMAIN ADAPTATION

```
model.save_weights(test_name+".hd5", overwrite=True)
```

```
keras.optimizers.Adam(lr=0.001, beta_1=0.9, beta_2=0.999, epsilon=None, decay=0.0, amsgrad=False)
```

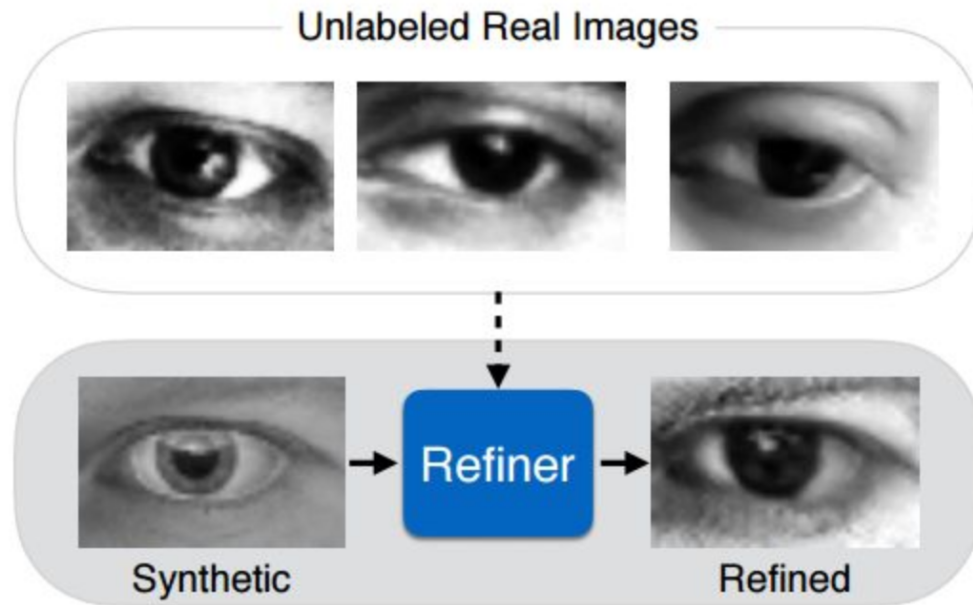
```
.. -----  
model.load_weights(model_name+".hd5")
```

```
all_layers = model.layers
```

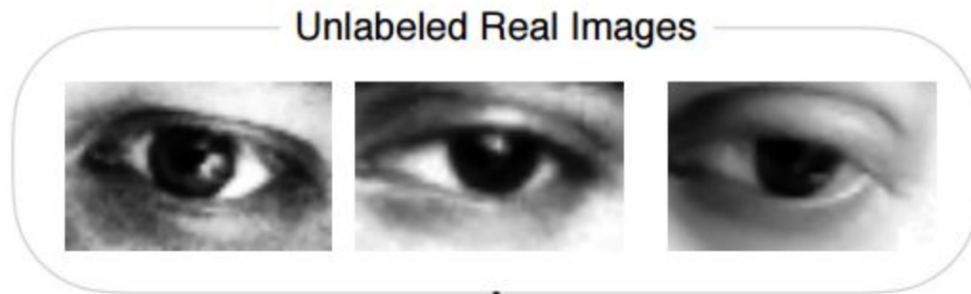
```
for i in range(model.layers.index(mid_start)):  
    all_layers[i].trainable = False
```

Freeze
some
layers

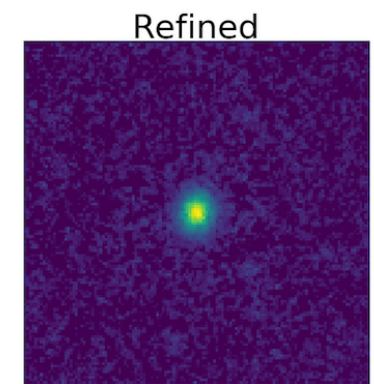
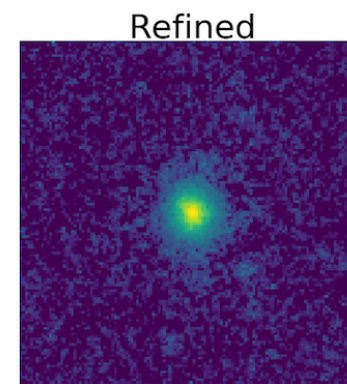
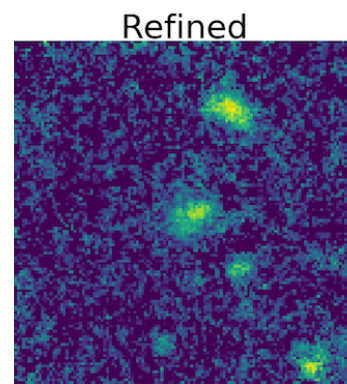
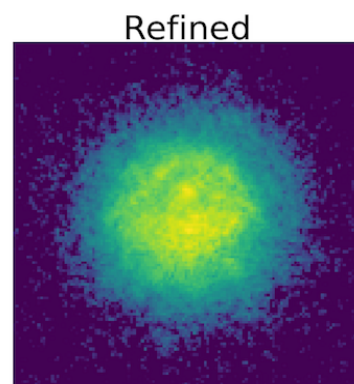
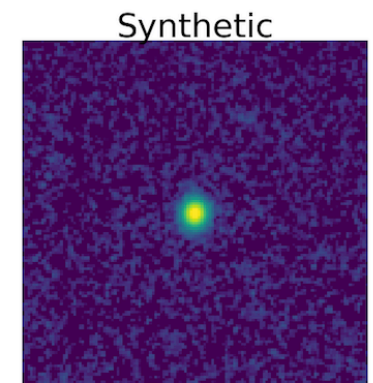
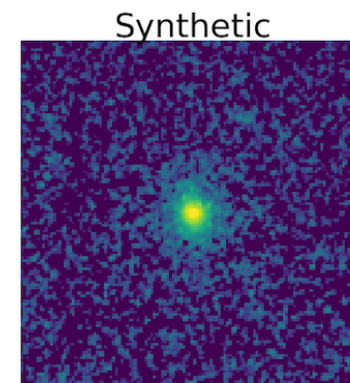
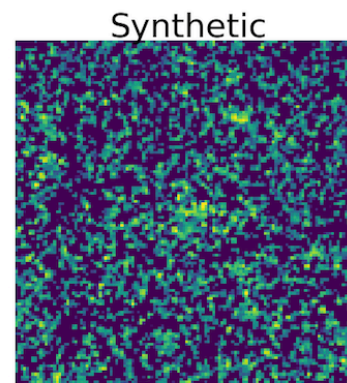
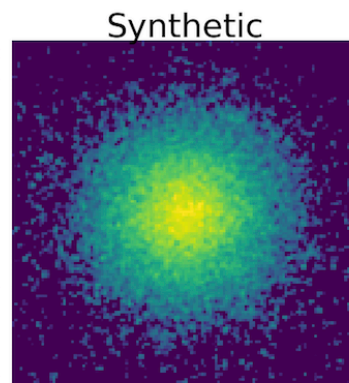
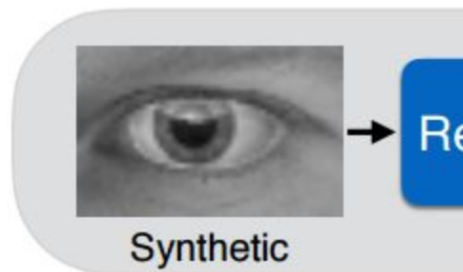
USE GANs?



USE GANs?



VERY PRELIMINAR WORK
IN CANDELS



HOW DO I SETUP MY CNN? HYPER-PARAMATER SEARCH

UNFORTUNATELY, THERE IS NO MAGIC RECIPE



HOW DO I SETUP MY CNN?

UNFORTUNATELY, THERE IS NO MAGIC RECIPE

SOME GENERAL ADVICES:

- START WITH SOMETHING THAT WORKS
- CAPACITY INCREASES WITH LAYERS, CHANNELS, RECEPTIVE FIELD
- USE PRIOR KNOWLEDGE OF THE EXPECTED SCALE CONTAINING MEANINGFUL INFORMATION
- GRID SEARCH...
- AVERAGING

CREDIT

HOW DO I SETUP MY CNN?

TYPICAL PARAMETERS

NUMBER OF HIDDEN LAYERS?

- START WITH FEW LAYERS AND INCREASE COMPLEXITY IF NEEDED
- ADD MORE LAYERS, CHECK PERFORMANCE
- ADD MORE NEURONS, CHECK PERFORMANCE

HOW DO I SETUP MY CNN?

TYPICAL PARAMETERS

ACTIVATION FUNCTION

- PRIORITY TO ReLU
- TRY EVENTUALLY LeakyReLU, PreLU

HOW DO I SETUP MY CNN?

TYPICAL PARAMETERS

OPTIMIZER

- NOT REALLY A RULE FOR THIS...
- SGD, ADAM THE ONES I MOSTLY USED. ADAM MORE ROBUST TO “NAN” LOSSES

HOW DO I SETUP MY CNN?

TYPICAL PARAMETERS

LEARNING RATE

- THIS IS ONE OF THE MOST TWEAKED PARAMETERS
- THE LEARNING RATE SHOULD BE LARGE AT THE BEGINNING AND SMALL TOWARDS THEN WHEN CLOSER TO THE MINIMUM

HOW DO I SETUP MY CNN?

TYPICAL PARAMETERS

LEARNING RATE

- THIS IS ONE OF THE MOST TWEAKED PARAMETERS
- THE LEARNING RATE SHOULD BE LARGE AT THE BEGINNING AND SMALL TOWARDS THEN WHEN CLOSER TO THE MINIMUM

HOW DO I SETUP MY CNN?

LEARNING RATE

- THE DECAY OF THE LEARNING RATE CAN BE SET :
 - STEP DECAY [DECAY BY FIX AMOUNT EVERY FEW EPOCHS]
 - EXPONENTIAL DECAY $\lambda = \lambda_0 e^{-kt}$
 - INVERSE DECAY $\lambda = \frac{\lambda_0}{1 + kt}$

HOW DO I SETUP MY CNN?

LEARNING RATE

- THE DECAY OF THE LEARNING RATE CAN BE SET :
 - STEP DECAY [DECAY BY FIX AMOUNT EVERY FEW EPOCHS]
 - EXPONENTIAL DECAY $\lambda = \lambda_0 e^{-kt}$
 - INVERSE DECAY $\lambda = \frac{\lambda_0}{1 + kt}$

- WHEN OPTIMIZING THE NETWORK
 - FIRST PERFORM A COARSE SEARCH, USUALLY RECYCLING SOMETHING THAT WORKED, RUN A FEW EPOCHS
 - THEN, LONGER TRAINING, FINER SEARCH

FINAL THOUGHTS

- START SIMPLE - KEEP ALWAYS YOUR SCIENCE IN MIND. NEW ALGORITHMS ARE COOL BUT COULD BE OVERKILL.
- HOW DO I KNOW WHICH ALGORITHM TO USE?
 - DON'T KNOW
 - DEEP LEARNING: IMAGE, SPECTRAL DATA WITH SPATIAL CORRELATIONS EASILY IDENTIFIED BY YOUR EYE BUT DIFFICULT TO DESCRIBE
- STAY CONNECTED! (OR MAKE FRIENDS WHICH ARE CONNECTED). MOST OF THE THINGS I HAVE SHOWN HERE WILL BE UPDATED IN A FEW MONTHS

FINAL THOUGHTS

- TRY TO CATCH THE ATTENTION OF COMPUTER SCIENTISTS WITH CHALLENGING PROBLEMS. WE DO HAVE MANY IN ASTRONOMY!
- SPECIFIC PROPERTIES OF ASTROPHYSICAL DATA:
 - OBJECTS WITH NO BOUNDARIES
 - HUGE DYNAMICAL RANGE
 - LOW S/N REGIME
- ALSO THEY ARE UNIQUE:
 - BIG-DATA VOLUMES
 - PUBLICLY AVAILABLE ALMOST INSTANTANEOUSLY
 - NO COMMERCIAL VALUE / ETHICAL PROBLEMS

THE FUTURE IS VERY
EXCITING!

MOST OF THE PROCESSING WE DO TODAY ON IMAGES CAN
BE ACTUALLY SOLVED WITH AI...