the Tractor: measuring astronomical objects using generative modeling

Dustin Lang
Princeton University

Berkeley Lab — 2012-02-07

Motivation

- Next generation of surveys:
  - Dark Energy Survey (DES),
  - Pan-STARRS,
  - Hyper-Suprime-Cam (HSC) survey,
  - Large Synoptic Survey Telescope (LSST)
- ...produce multi-epoch, multi-band, deep and highly blended images
- How do we [optimally] detect and measure astronomical objects in these image collections?

Motivation – Going deep

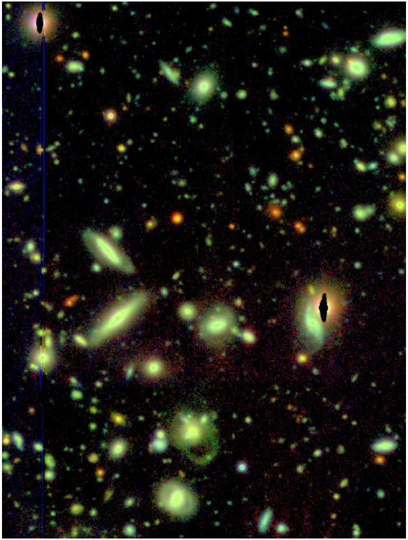

Motivation – for spectroscopic surveys

- For upcoming surveys:
  - eBOSS
  - Big-BOSS
- ...we need 14,000 square degrees of deep, multi-band imaging for target selection
- We might have to patch together existing and new surveys from different instruments
- (eg, PTF for r-band, CFHT for u and g, KPNO for i, z)
- How do we [optimally] detect and measure astronomical objects in these diverse image collections?

Tradition

- Traditional approach: some variant of:
  - make a canonical co-add
  - detect and deblend on canonical co-add
  - make a co-add for each band
  - do forced photometry on each band's co-add
- eg, CFHT-LS (Goranova *et al.* ; Gwyn, arXiv:1101.1084v2.), SDSS Stripe 82 (Annis, arXiv:1111.6619v2.; Huff, arXiv:1111.6958v1.)
- What's wrong with tradition?

There is no such thing as an optimal co-add

The Problem With Tradition

For detection of point sources, no weighted sum of images yields the total signal-to-noise available:

- Optimal point-source detection requires a matched filter
- Co-adding mixes the good- and bad-seeing images
- Can't match them both with a single filter; mismatched filter

But. . .

- It *is* possible to build an optimal detection map (Kaiser 2004; Lang *et al.* in prep) but it's not really an image, and it has bad resolution so isn't good for deblending or galaxy shape measurements.
- Take-home message: can do detection on co-adds, but making measurements (of galaxy shapes and brightnesses) on co-adds is fraught.

Idea – the Tractor

Simultaneous forward modeling

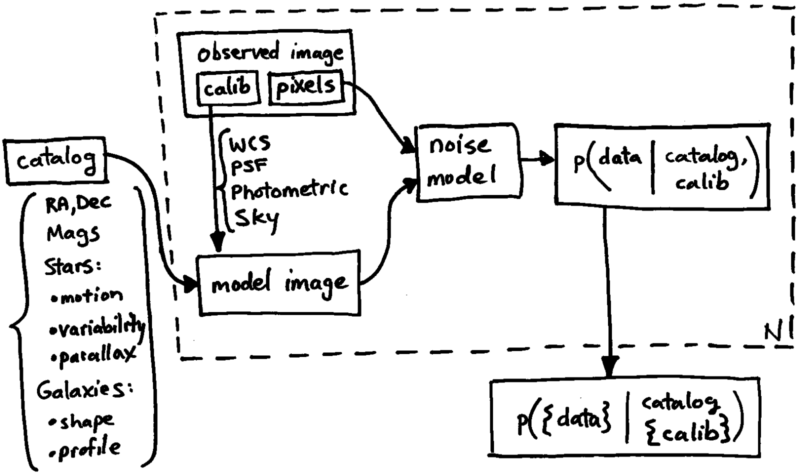

On a single galaxy

- Grab a single galaxy from the SDSS catalog:
  RA,Dec = (333.5550, 0.3644), r-band mag = 18.39;
  deVaucouleurs profile: $r_e = 0.7''$, $ab = 0.43$, $\phi = -136.1$ deg
  (in run 2728, camcol 4, field 236)
- What would that look like? At high res, with small PSF and no noise:

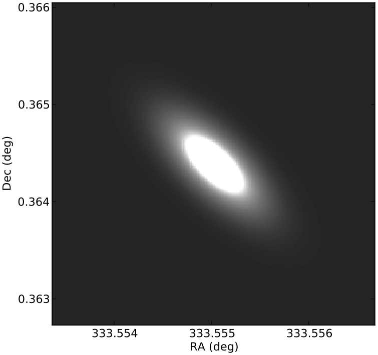

On a single galaxy

- Grab a single galaxy from the SDSS catalog:
  RA,Dec = (333.5550, 0.3644), r-band mag = 18.39;
  deVaucouleurs profile: $r_e = 0.7''$, $ab = 0.43$, $\phi = -136.1$ deg
  (in run 2728, camcol 4, field 236)
- What would that look like? At SDSS resolution, with small PSF and no noise (and SDSS WCS):

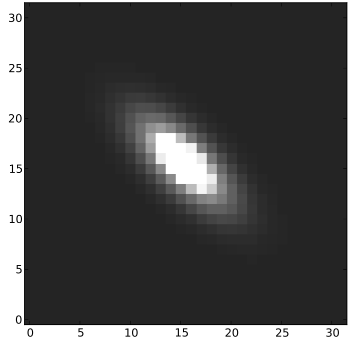

On a single galaxy

- Grab a single galaxy from the SDSS catalog:
  RA,Dec = (333.5550, 0.3644), r-band mag = 18.39;
  deVaucouleurs profile: $r_e = 0.7''$, $ab = 0.43$, $\phi = -136.1$ deg
  (in run 2728, camcol 4, field 236)
- What would that look like? At SDSS resolution, with the PSF from run/camcol/field 2728/4/236, and no noise:

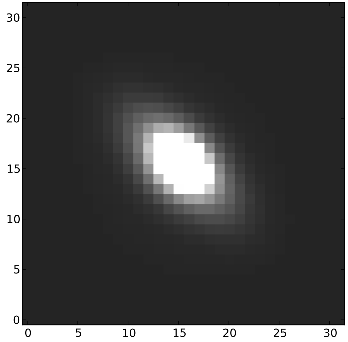

On a single galaxy

- Grab a single galaxy from the SDSS catalog:
  RA,Dec = (333.5550, 0.3644), r-band mag = 18.39;
  deVaucouleurs profile: $r_e = 0.7''$, $ab = 0.43$, $\phi = -136.1$ deg
  (in run 2728, camcol 4, field 236)
- What would that look like? In the actual SDSS image:

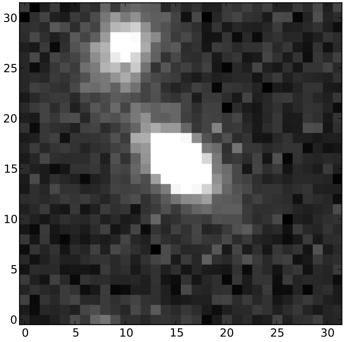

On a whole image

- The SDSS catalog has stars, exponential galaxies, deVaucouleurs galaxies, and composite (exp+deV) galaxies:

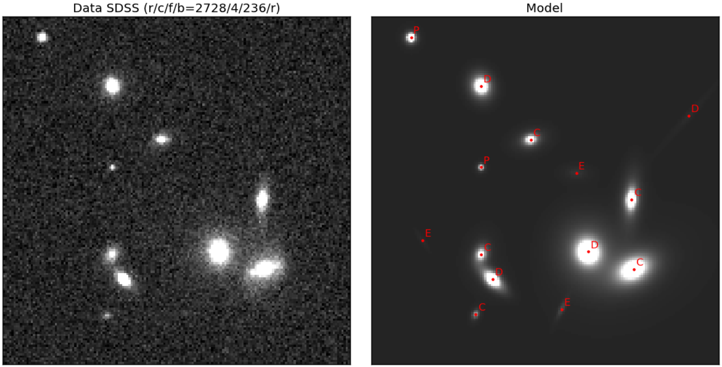

Noise model

- If our noise-free model were perfect, the only difference between the model and real image would be noise (photon noise from sky + source, readout noise)
- Pixelwise independent, ∼Gaussian noise
- Noise model: $p(\mathrm{data}|\mathrm{catalog}) \propto \exp(-\frac{1}{2}\chi^2)$

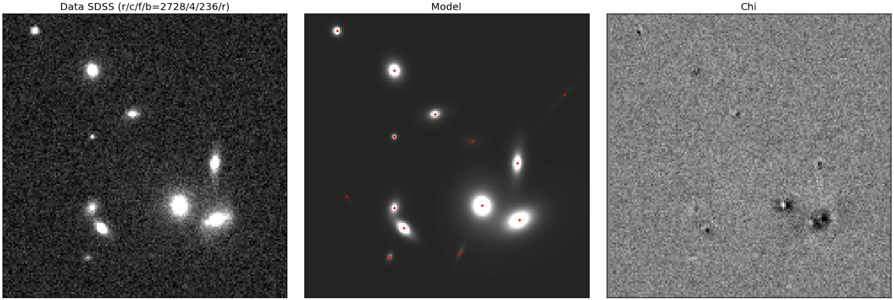

Multiple images

- Given a catalog and image calibration, we can predict a model image
- A noise model lets us assign a probability to the observed data given the model: usually chi-squared
- Multiple images? Produce model image for each image, apply the per-image noise model, and multiply the probabilities!

Multiple images

SDSS 2728/4/236 SDSS 4868/4/31 CFHT-LS 850994p

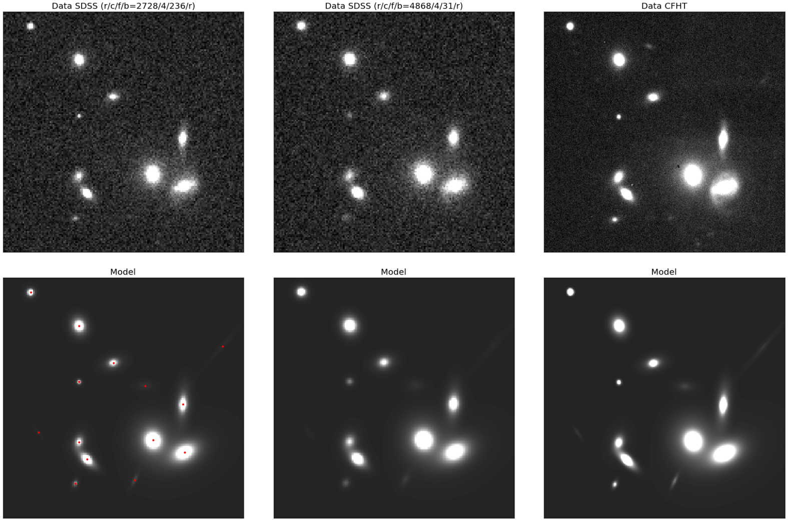

Multiple images

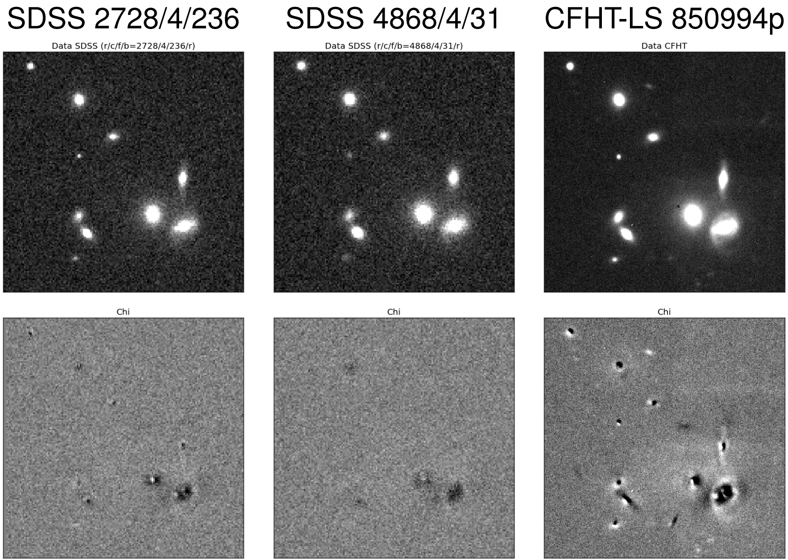

Multiple bands

Some options:

- Share the galaxy shape parameters between the bands, but fit separate magnitudes per band
  - one image much deeper than the rest: it drives the fit and you get roughly forced photometry
- Closer to forced photometry: fit on a canonical band, "pin" the galaxy shapes, and then fit mags per band
- Fit separate galaxy shapes and mags per band

Idea

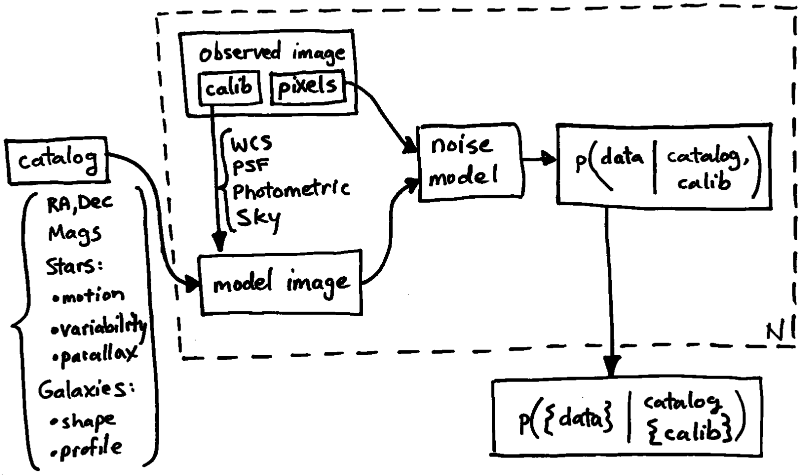
observed image
calib
pixels
WCS
PSF
Photometric
Sky
catalog
RA,Dec
Mags
Stars:
•motion
•variability
•parallax
Galaxies:
•shape
•profile
model image
noise model
p(data | catalog, calib)
N
p({data} | catalog {calib})

Idea

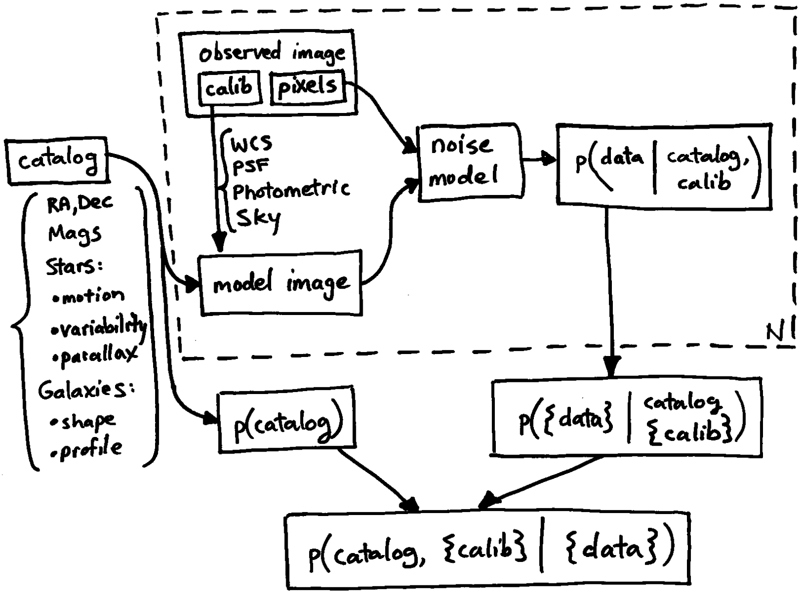
observed image
calib
pixels
WCS
PSF
Photometric
Sky
model image
noise model
p(data | catalog, calib)
N
catalog
RA,Dec
Mags
Stars:
•motion
•variability
•parallax
Galaxies:
•shape
•profile
p(catalog)
p({data} | catalog {calib})
p(catalog, {calib} | {data})

Posterior probability

- Noise model gives us likelihood $p(\mathrm{data} \mid \mathrm{catalog}, \mathrm{calib})$
- We really want posterior $p(\mathrm{catalog} \mid \mathrm{data})$
- Bayes to the rescue:

$$\mathrm{posterior} = \frac{\mathrm{likelihood} \times \mathrm{prior}}{\mathrm{evidence}}$$

$$p(\mathrm{catalog} \mid \mathrm{data}) = \frac{p(\mathrm{data} \mid \mathrm{catalog})\, p(\mathrm{catalog})}{\mathrm{p(data)}}$$

  and we can ignore $p(\mathrm{data})$ here.
- Key: prior $p(\mathrm{catalog})$: penalize unlikely individual objects, and encourage simplicity in the catalog (number and complexity of sources)
- Occam's Razor for model selection

Tuning

- So far, we are just rendering the SDSS catalog
- But since we have a scalar objective function,

$$p(\mathrm{catalog} \,|\, \mathrm{data})$$

  we can optimize the catalog to better match the observed images
- Instead of starting with the SDSS catalog, we could have started from scratch
- We can also optimize the calibration parameters:

$$p(\mathrm{catalog}, \mathrm{calib} \,|\, \mathrm{data})$$

  or marginalize over them.

Tuning structure

- Need to optimize the model structure as well as the parameter values
- (number and type of sources; PSF, astrometry, and sky)
- Heuristics needed here!

Implementation – Tuning parameters

Current optimization approach (naive):

- For each catalog parameter, take a small step
- Render the new model image and compute the finite difference approximate derivative
- Build a big sparse matrix: $A_{ij} = \mathrm{dpixel}_i / \mathrm{dparam}_j$
- Solve (least squares) $(Aw)X = \chi$, with per-pixel standard deviations w
- X is a parameter update
- Try to step in direction X, or $X/2$, $X/4$, ...
- Repeat

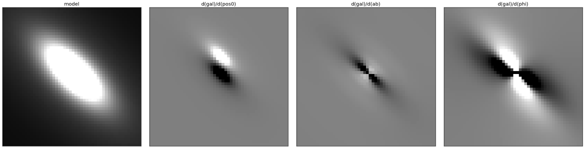

Implementation – Tuning structure

A naive version of model switching:

- try switching a star $\Leftrightarrow$ exponential or deVaucouleurs galaxy $\Leftrightarrow$ composite galaxy
- (or try creating a new source where there are positive residuals)
- re-optimize everything after the change
- accept if posterior probability is better

Implementation – Rendering model images

- Represent PSF as mixture of Gaussians (fit to pixelized model via EM)
- Represent galaxy profiles (exp, deV) as mixtures of Gaussians
- Convolution is analytic!
- Rendering models involves lots of Gaussian evaluations
- (GPU?)

Example

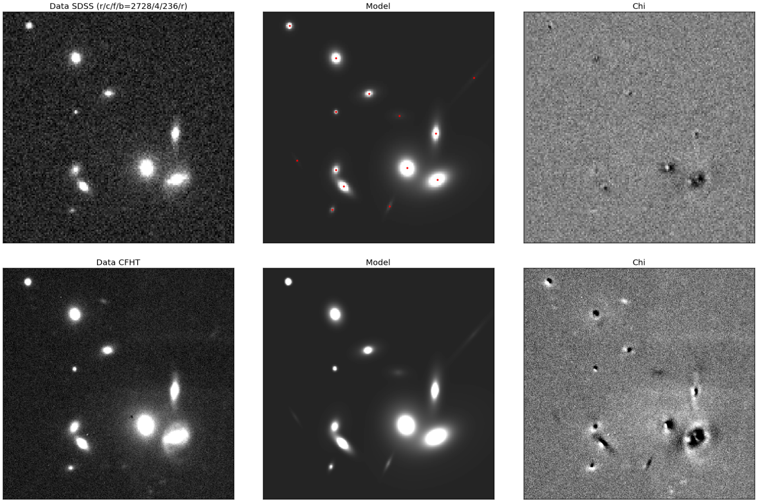

Example

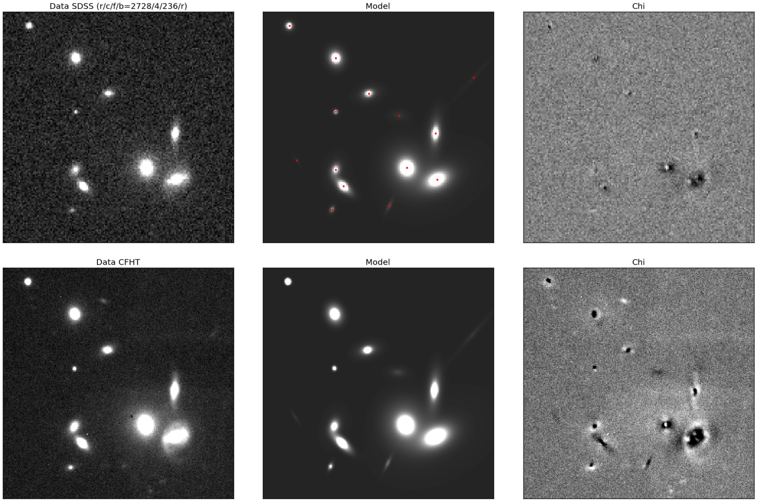

Example

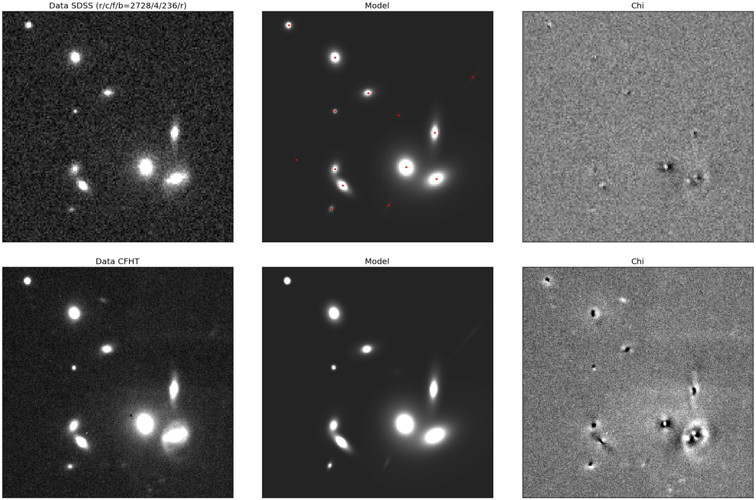

Example

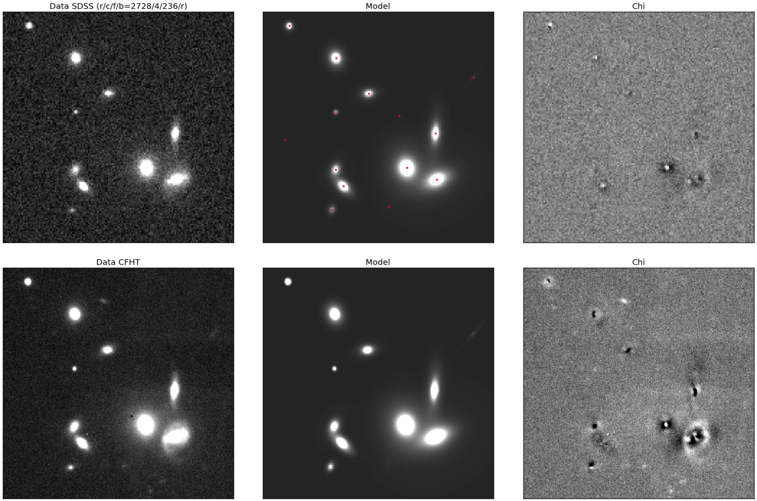

Example

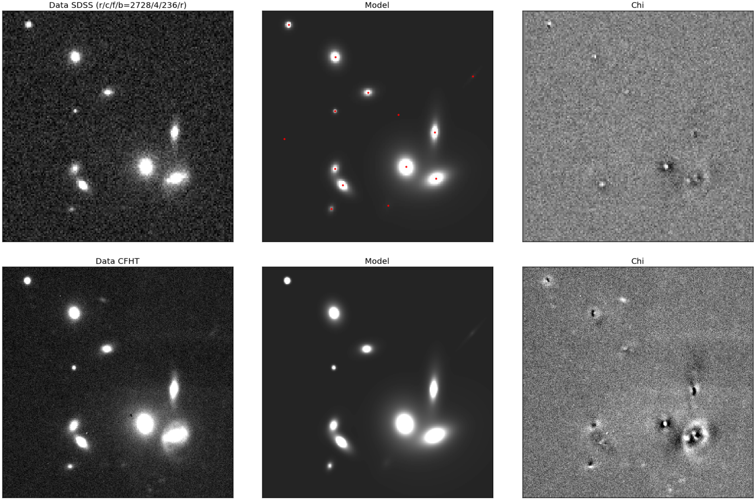

Example

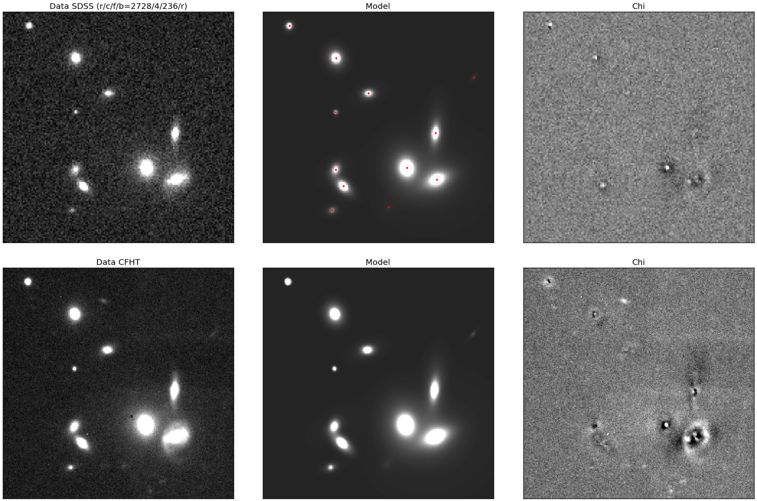

Next. . .

- CFHT CS82 + SDSS Stripe 82
- Schlegel's Amazing Technicolor Dream Survey (PTF + CFHT + KPNO)
- Hubble?

Future work

- Annotating galaxy catalogs (easy)
- Sampling (easy-ish)
- Better sky (medium)
- Regularized PSF model (medium)
- Proposing catalog structural changes (medium)
- Flexible, regularized galaxy models (Bayesian nonparametric) (medium?)
- Summarizing and propagating covariance (deep+hard)
- Layering cosmology on top (eg, weak lensing) (hard)

Thanks!

Time for questions and discussion!