Data Management Weak Lensing Updates

Jim Bosch
Princeton University

December 4, 2013

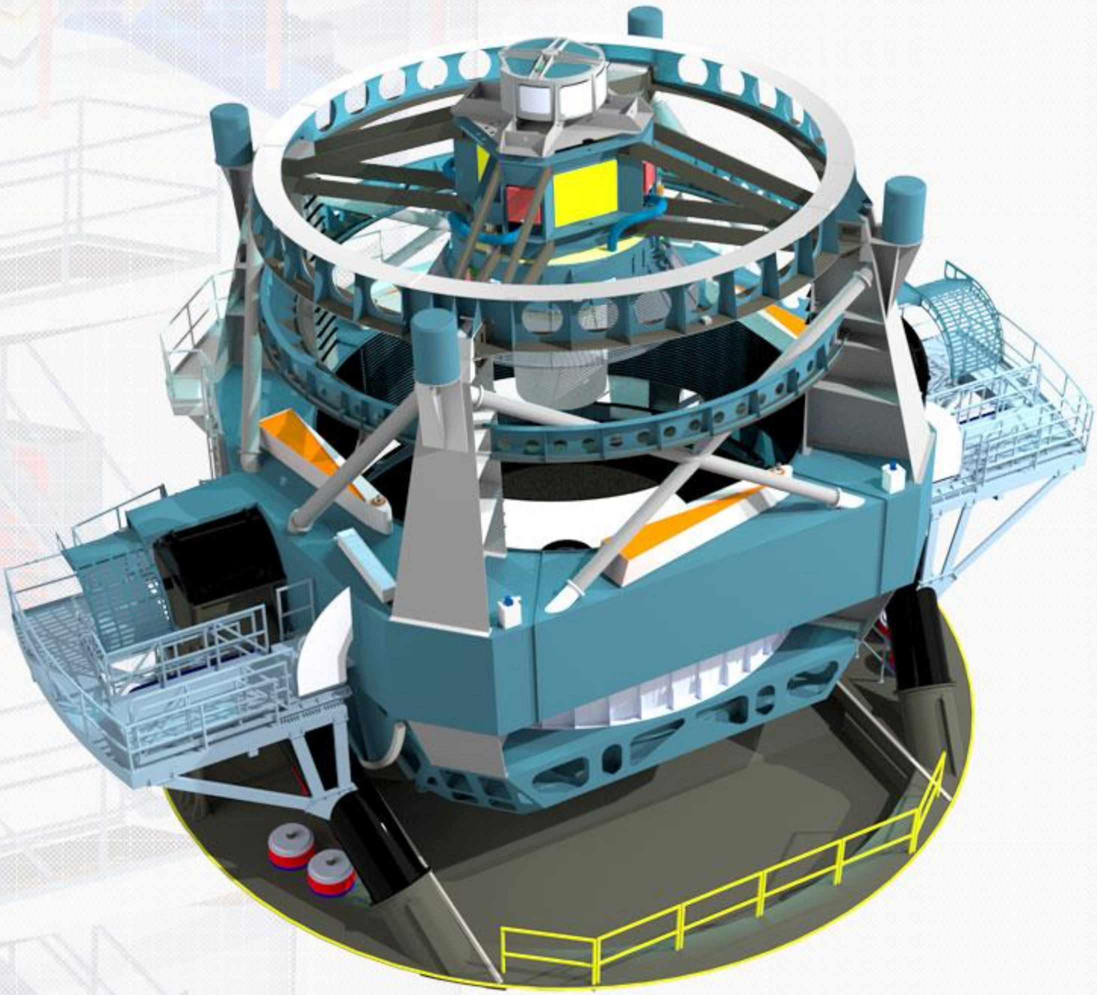

The Year in Reviews

LSST has undergone *26 reviews* this year (including the FDR happening now):

https://www.lsstcorp.org/reviews/hub/pastreviews

You probably only care about a few of them*:

- Data Products Definition (June 6-7)
- Photometric Calibration (July 9-10)
- ImSim (August 15-16)
- Core Data Processing Plan (September 19-20)

* More if you're interested in the details of how the hardware works

Documents! With Numbers!

One by-product of being reviewed is that we now have baseline plans for everything, and we have documents that describe those baseline plans. And those documents have numbers!

Some of these documents are actually worth reading:

- Data Products Definition Document (LSE-163): what DM will provide to science users
- DM Software Development Roadmap (LDM-240): when DM will implement certain features

Get these from DocuShare button at http://www.lsstcorp.org (or ask me or someone else in the project to get them for you)

What "Baseline Plan" Does and Doesn't Mean

- There is no guarantee that this is what the LSST Project will actually do.
  - If you hate our plan, you can still work to change it!
  - We might change it ourselves if we decide something else will work better, even if you love it; baseline plans are not requirements.
- This is what we've used to project costs for the Project, and those are now fixed.
  - If you want to do something else, it can't be more expensive.
- These are plans for the LSST Project, not DESC or other science collaborations.
  - DESC can do additional processing later, and the Project has long considered it a requirement to support that sort of work from a technical (but not necessarily financial) standpoint.

Baseline Plans for Shape Measurement

- Limitations of coadd-based measurement
- Data products produced via MultiFit methods
- Shape measurement for weak lensing
  - The algorithmic landscape, and LSST's approach to it
  - Computational challenges and optimization
- Implementation: current prototypes
- Roles for DESC and the LSST Project in shape measurement

Coadd Limitations: Optimal Weights

How do we optimally combine these two images?

	Exposure Time	Seeing
A	30s	0.6”
B	100s	1.0”

Kaiser: formally possible; weight each spatial frequency separately

- requires stationary noise, spatially constant PSF, no bad pixels, no distortion between images...
- involves convolving each image with its PSF – complicates measurement
- worth investigating, but never demonstrated in practice

Just a signal-to-noise issue – not a source of systematic error

Coadd Limitations: Correlated Noise

Warping and resampling images correlates pixel noise

- correlations are spatially coherent across the image: dangerous for weak lensing
- measurement in the presence of correlated noise is more challenging (most often it's just ignored, which at the very least will produce incorrect uncertainty estimates)
- propagating full covariance information through the coaddition process is impractical
  - for lanczos2 resampling, we'd have approximately 12 unique nonzero covariance elements per coadd pixel

Potentially a source of systematic error

Coadd Limitations: PSF

- The effective point-spread function on a coadd is discontinuous at the edges of the input frames
- When pixels are rejected from the coadd due to clipping or masks, there is no well-defined effective PSF on the coadd
  - worst case: stars that are saturated only in good seeing

Partial solution from Jee and Tyson (2011): *compute PSF model on the coadd by warping/coadding PSF models from individual frames (a.k.a. StackFit, CoaddPsf)*

- only valid when object footprints do not contain missing pixels or frame boundaries
- unknown whether this can fully account for DCR effects
- still a substantial improvement over attempting to do PSF determination from coadd star images

Coadd Limitations: Detector Effect

In most modern thick CCDs, the presence of charge in a pixel can significantly distort the electric field in the chip, redistributing charge in a nonlinear way (this is *not* a convolution)

When operating on a single exposure, we need to be able to model this effect and apply it in the forward direction, which is difficult enough.

In order to build a coadd that accounts for these effects, we'd have to apply the inverse operations, which is inherently less robust and introduces additional noise correlations.

MultiFit vs. Coadd Measurement

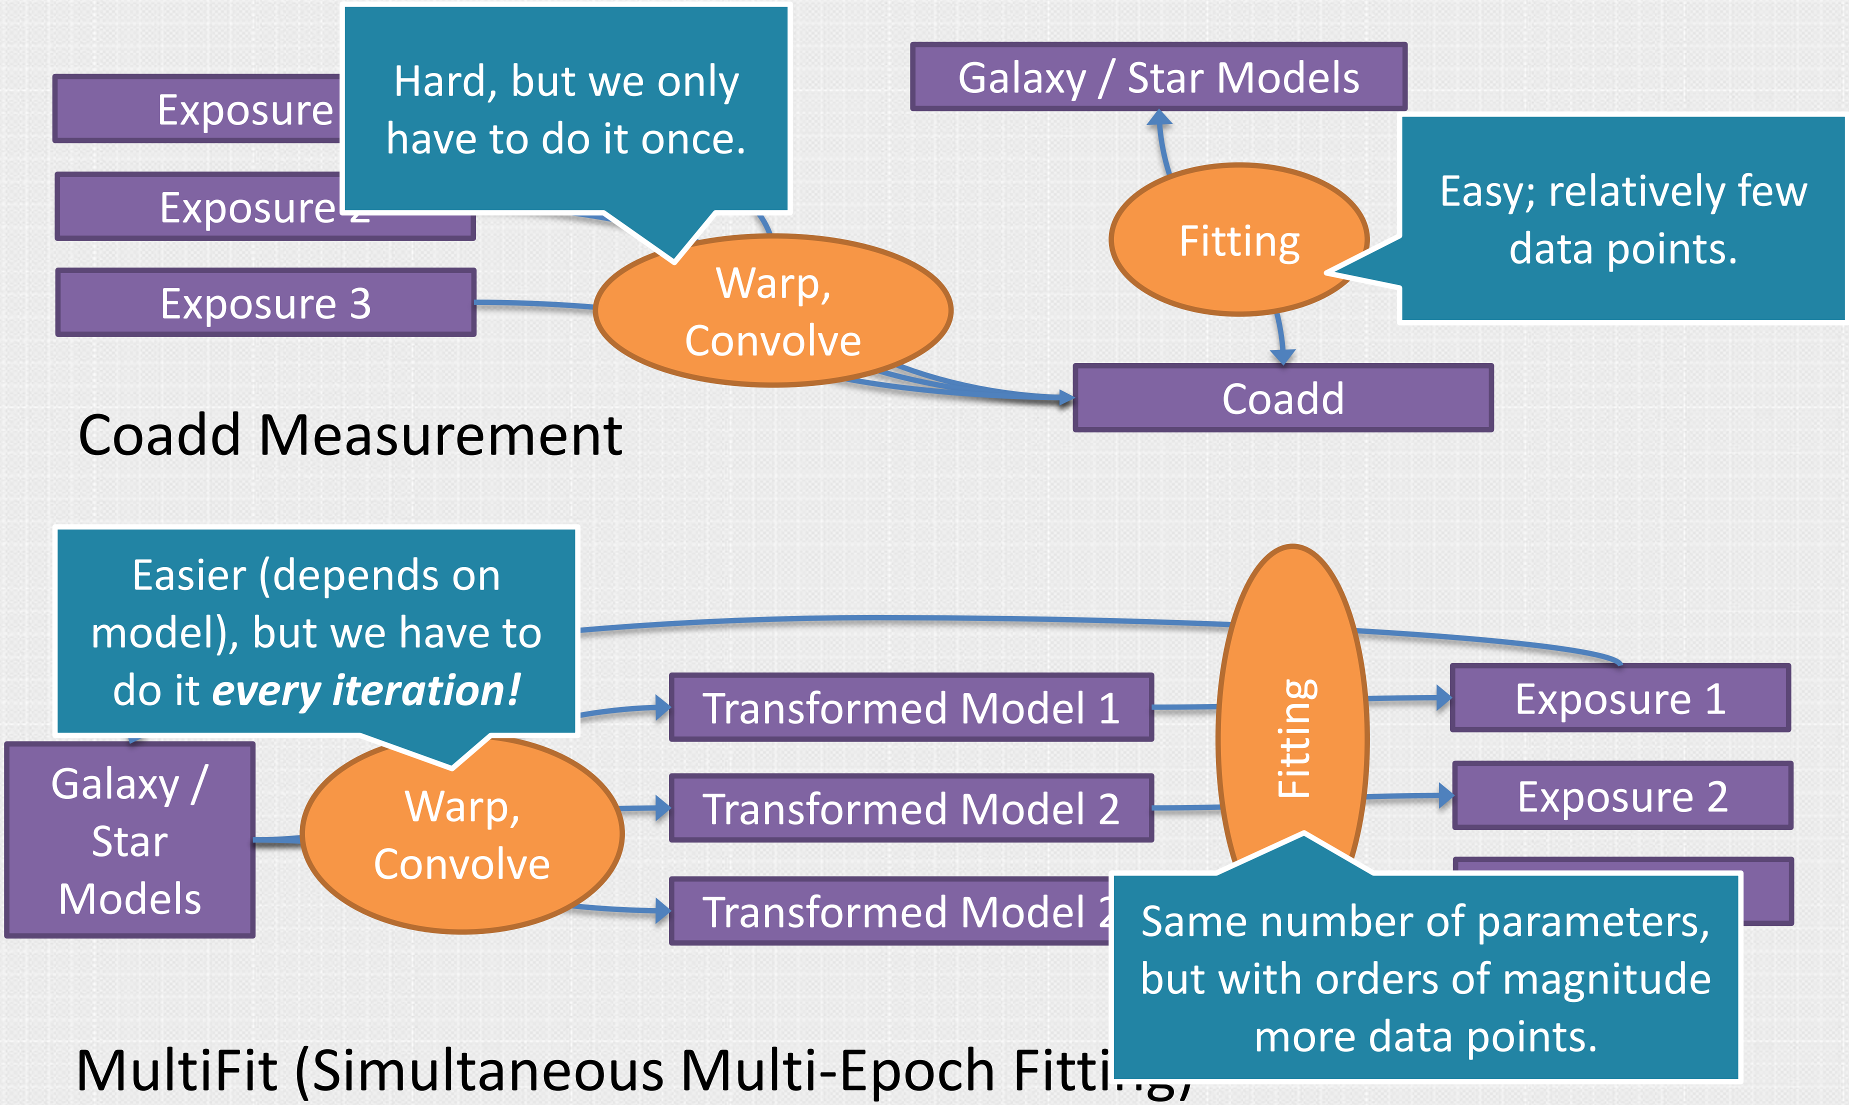

Data Products

From Data Products Definition Document (Section 5.2.1):

- ***Bulge-disk model fit*.** The object is modeled as a sum of a de Vaucouleurs (Sersic n = 4) and an exponential (Sersic n = 1) component. This model is intended to be a reasonable description of galaxies. The object is assumed not to move. The components share the same ellipticity and center. The model is independently fit to each band. There are a total of 8 free parameters, which will be simultaneously constrained using information from all available epochs for each band. Where there's insufficient data to constrain the likelihood (e.g., small, poorly resolved galaxies, or very few epochs), priors will be adopted to limit the range of its sampling. In addition to the maximum likelihood values of fitted parameters and their covariance matrix, a number (currently planned to be 200, on average) of independent samples from the likelihood function will be provided. These will enable use-cases sensitive to departures from the Gaussian approximation, with shear measurement being the primary use case. A permissible descope, in case of insufficient storage, will be not to sample the posterior for u and y bands.

Data Products: Galaxy Model Fitting

- Current baseline is similar to *lens*fit (Miller 2013):
  - restricted bulge+disk model
    - same ellipticity for both components
    - radius ratio between components is fixed
    - amplitude of components can vary independently
  - Bayesian: posterior = prior $\times$ likelihood
  - will sample from posterior, rather than rely on maximum likelihood fit with 2^{nd}-derivative covariance estimate
- Will explore modifications with more/fewer free parameters
- Will store posterior samples in the database, using a weak prior; users can apply a stronger prior later if desired

Shape Measurement: Algorithmic Landscape

- weighted moments + corrections
  - KSB 1995 (etc)
- forward fitting, Bayesian sampling
  - *lens*fit (Miller 2013)
- forward fitting, max-likeli
  - e.g. *im3shape* (Zuntz 2013)
- Fourier-domain methods
  - Bernstein & Armstrong 2013
- object-stacking methods
  - Lewis 2009 (etc)

Easy computationally; no longer state-of-the-art in systematics control

State-of-the-art in systematics control, battle-tested on real data, and likely the most computationally challenging

Unproven on real data; likely easier computationally

Shape Measurement: Work Yet To Do

- *lens*fit shear biases were acceptable for CFHTLenS; will need improvement for LSST due to smaller statistical uncertainties
  - better priors?
  - use posterior samples differently?
  - more flexible galaxy models?
  - marginalization over PSF/distortion uncertainties?
  - fitting multiple galaxies simultaneously?

Most of these increase the number of parameters in the galaxy model.

- *lens*fit is already too slow to be practical for LSST as-is
  - approximately 700ms per epoch (in many-epoch limit)
  - LSST computation budget for shape measurement is ~17ms per epoch – we need to speed things up by a factor of 40.

Shape Measurement: *lens*fit in detail

- Galaxy models are evaluated and convolved directly in Fourier-space, using pre-computed profiles

Fast, and by no means naive – but can we do better?

- Centroid and amplitude parame marginalized, reducing dimensionality of parameter space

- Radius and complex ellipticity are sampled on an adaptive grid
  - requires approximately 500 samples per galaxy – much less than MCMC sampling, which can require up to 8000
  - if we need to add more parameters to the model, a grid approach doesn't scale

Monte Carlo sampling *would* scale better with additional parameters

Multi-Shapelet Model Evaluation

Approximate Sérsic profiles using sums of Gaussians:

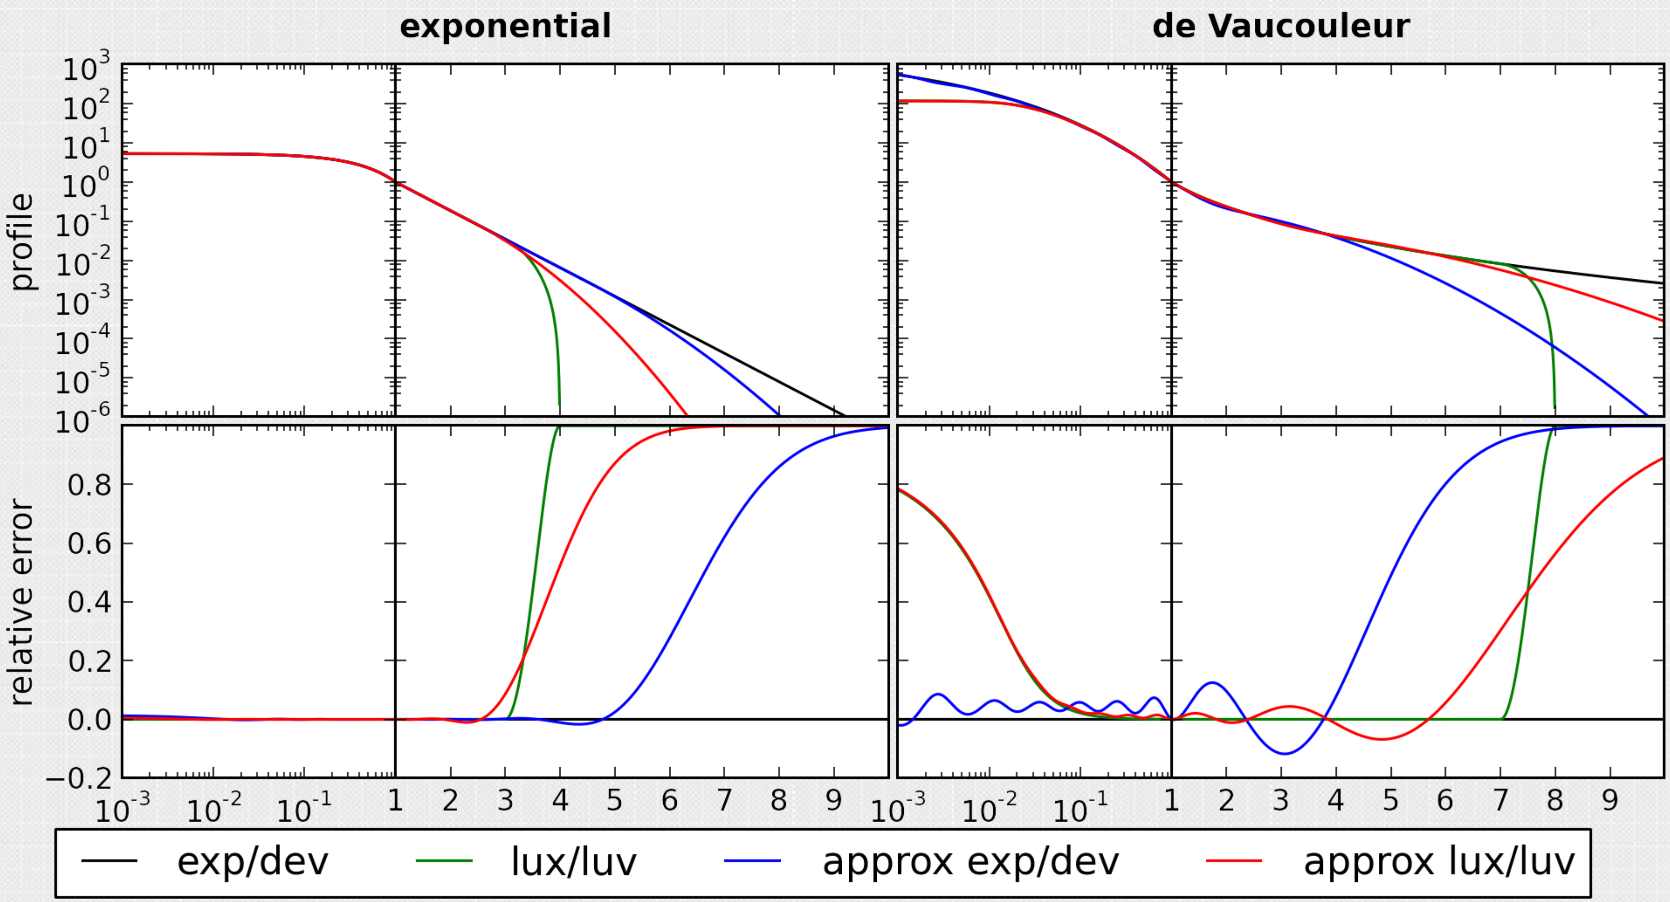

See Hogg and Lang (2013), Bosch (2010) for more information

Multi-Shapelet Model Evaluation

...and approximate the PSF using sums of shapelets:

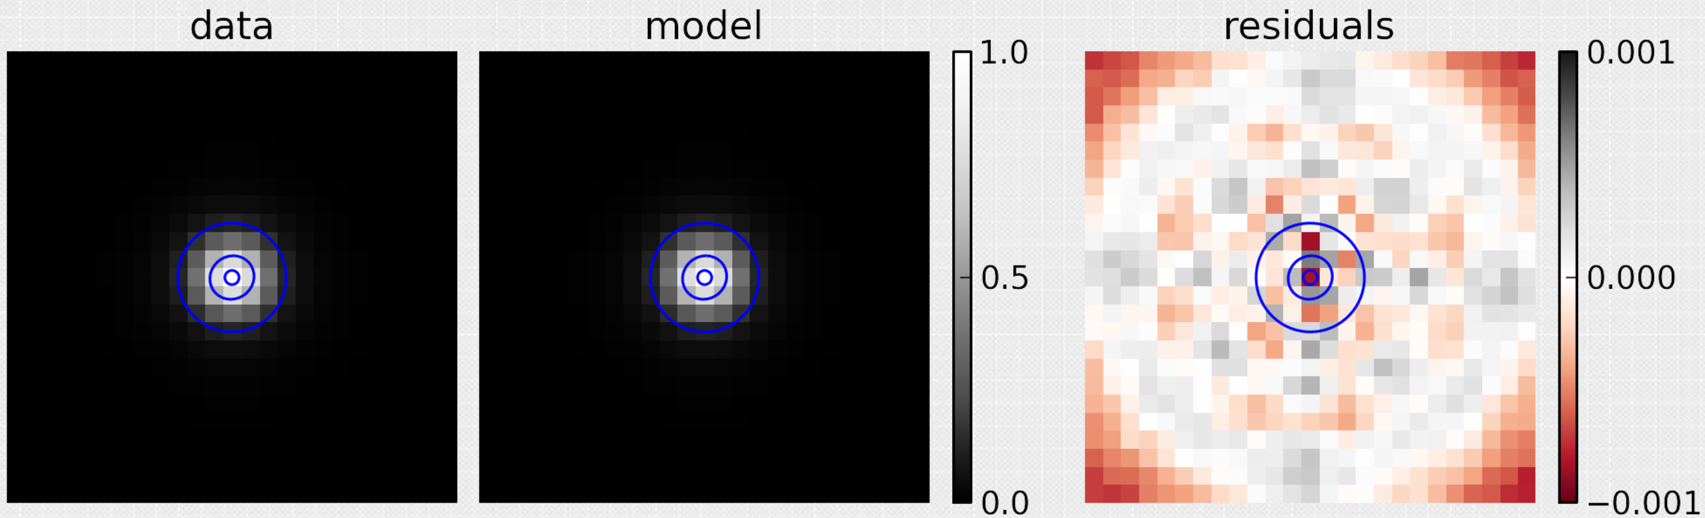

- Galaxy profiles are approximate...

  ...but in a straightforward and controllable way.

- Convolution is exact and *fast*:

  Model evaluation in 0.56ms; *lens*fit is *very approximately* 1.4ms (≠)

  Room for more improvement with custom CPU instruction sets, GPUs

Adaptive Importance Sampling

Given an unknown probability distribution $p(x)$ and a known proposal distribution $q(x)$ that approximates it, we can approximate the expectation value of $f(x)$ over $p(x)$ by drawing samples x_i from $q(x)$ and evaluating $p(x)$ at these points:

$$\int f(x)p(x)dx = \int f(x)\frac{p(x)}{q(x)}q(x)dx \approx \frac{1}{N}\sum_i^N f(x_i)\frac{p(x_i)}{q(x_i)}$$

- If $q(x)$ is nonzero everywhere that $p(x)$ is nonzero, this is an *unbiased estimator*, regardless of the number of samples.
- As $q(x)$ approaches $p(x)$, fewer samples are needed to get a good approximation.
- We can do this iteratively: parameterize $q(x)$, and update it to better match $p(x)$ based on the samples we draw.

See Wraith et al (2013) for more information

Adaptive Importance Sampling

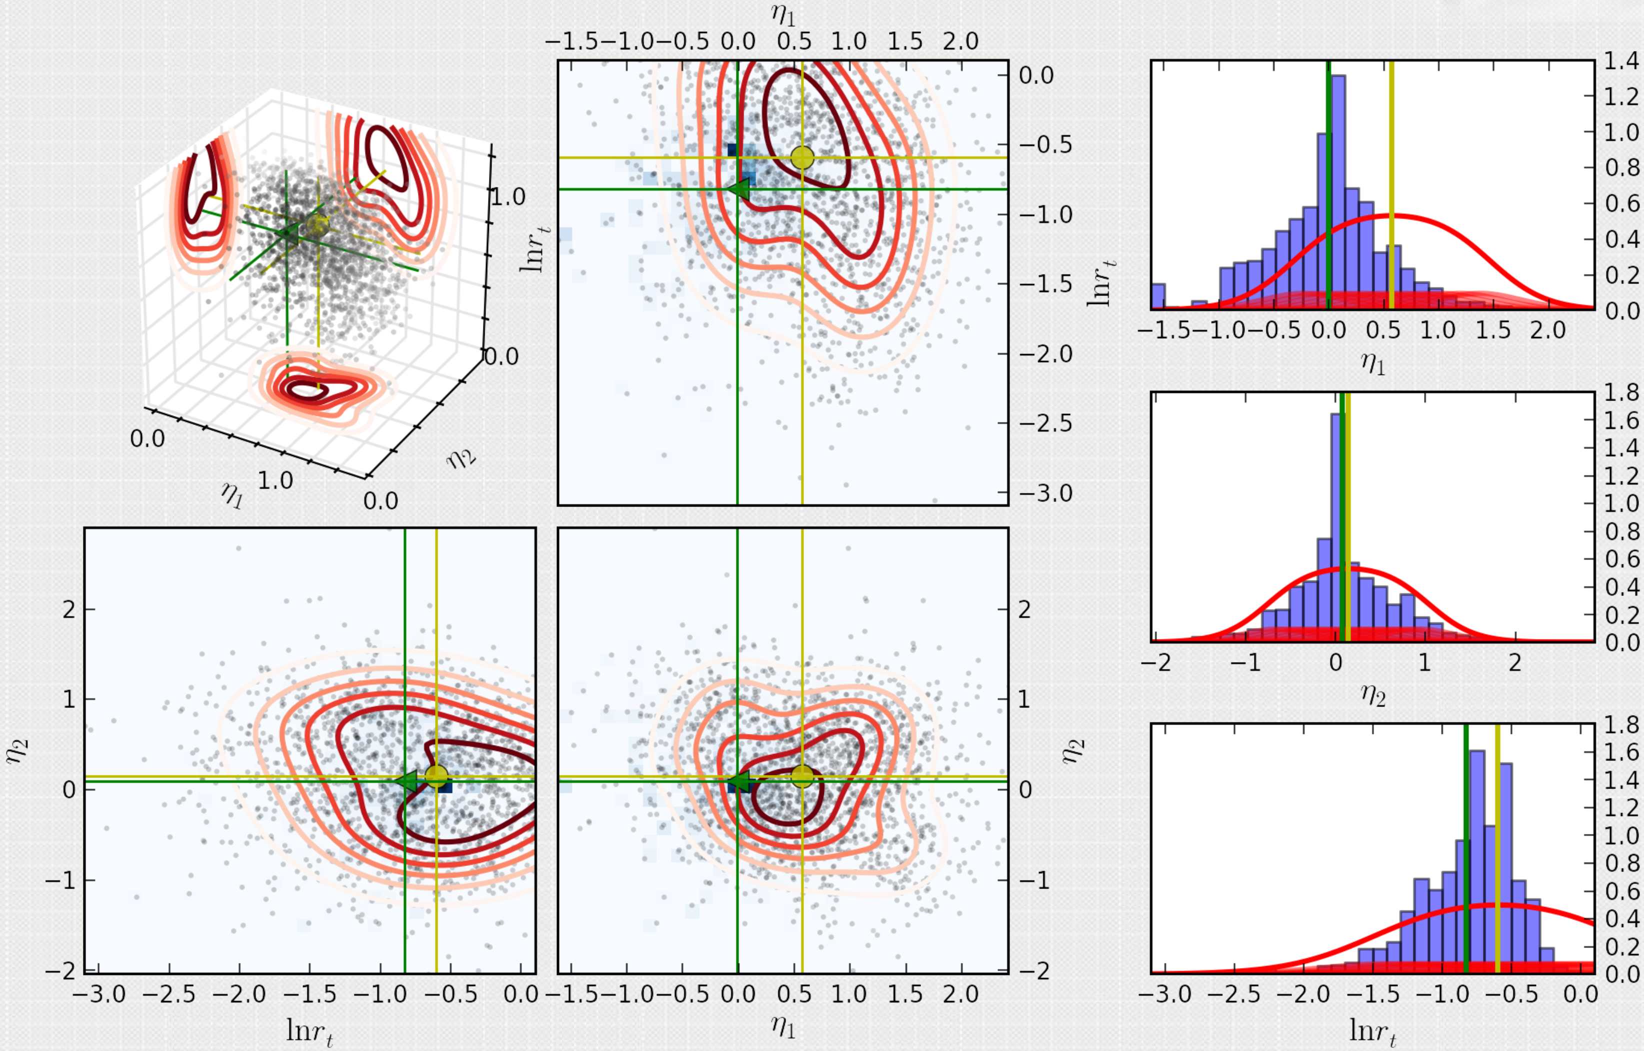

iteration=1, perplexity=0.263365, essf=0.115455

Adaptive Importance Sampling

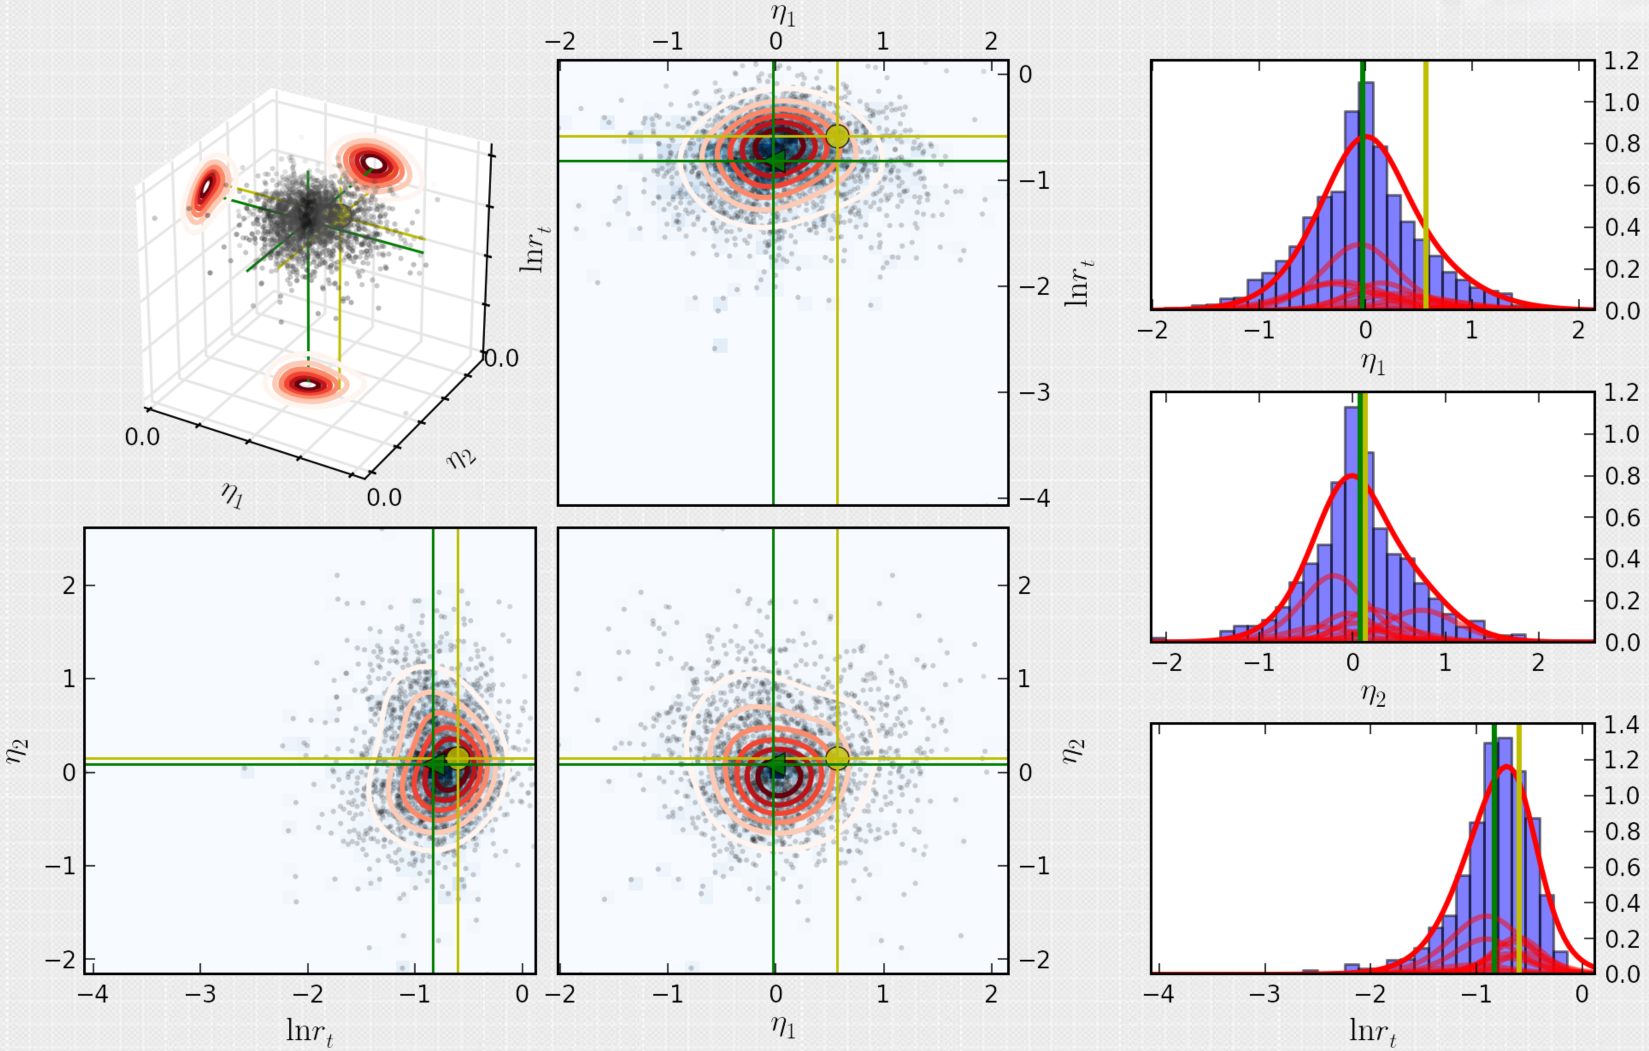

iteration=2, perplexity=0.872371, essf=0.737913

Adaptive Importance Sampling

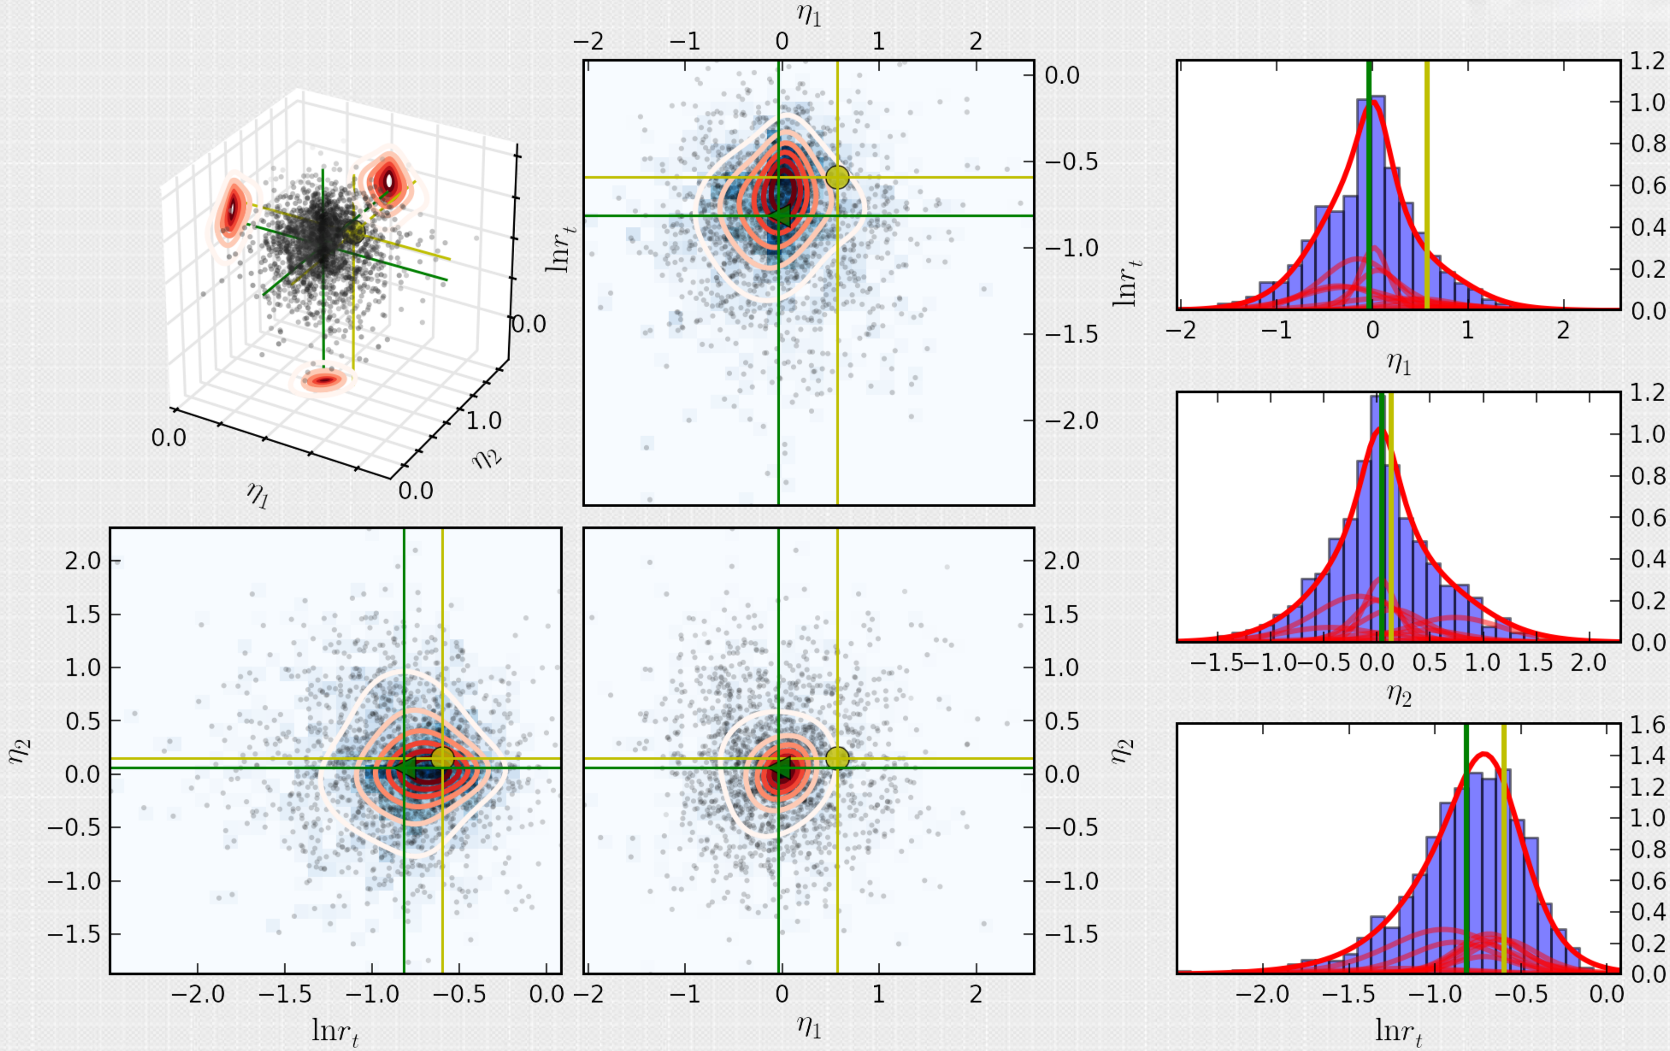

iteration=3, perplexity=0.939877, essf=0.893093

Adaptive Importance Sampling

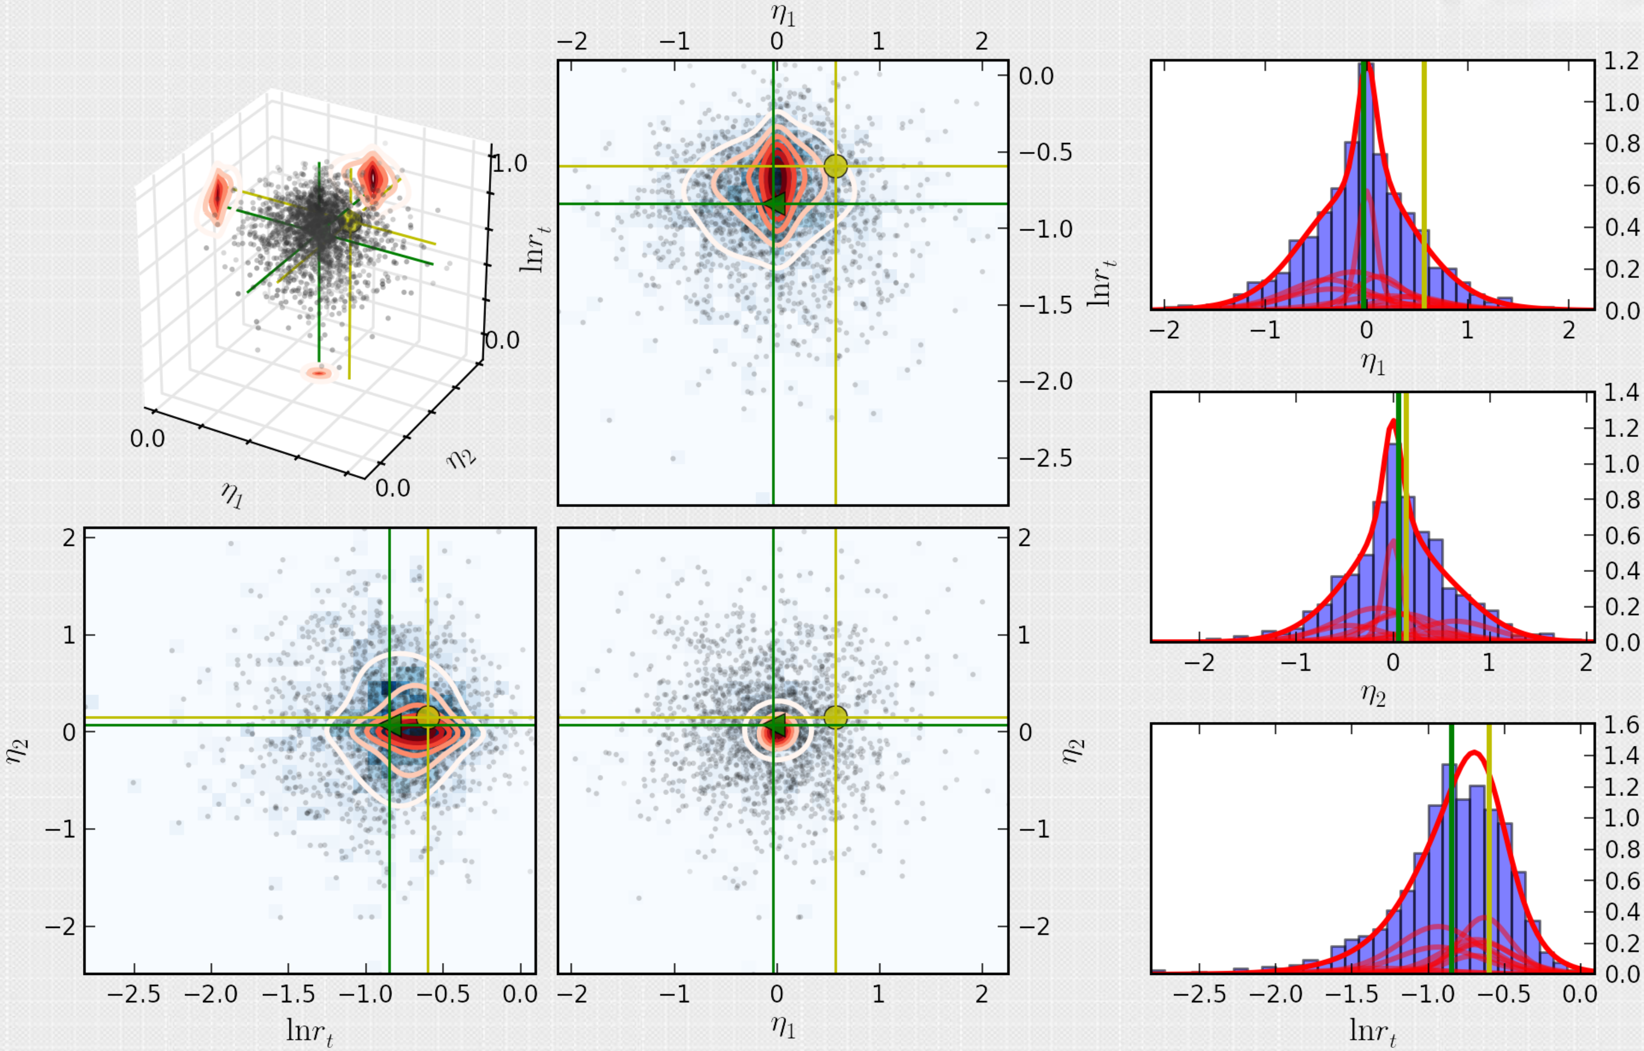

iteration=4, perplexity=0.964655, essf=0.934239

Using coadd measurements to feed MultiFit

- This is where adaptive importance sampling is really valuable: assuming the likelihood surface on the coadd is similar to the likelihood surface in MultiFit, we can do most of those iterations on the coadd, and only switch to MultiFit when we already have a good proposal distribution.

- If we can use a maximum likelihood approach instead, things are even easier: we can probably just do one or two Gauss-Newton steps in Multi-Fit to "tune up" the coadd fit result.

Work we can do on the coadd is ~200x (number of exposures) faster than work we do in MultiFit mode, so the more we can do there the better.

Implementation: Current Prototype

- Coaddition of PSF models fully implemented
- Prototype implemented this summer:
  - Multi-Shapelet model evaluation
    - one round of profiling: new convolution relation → 3.4x speedup
  - adaptive importance sampling
  - start from coadd (using StackFit PSF models), use to feed MultiFit
  - for more information, see the *MultiFit Prototyping Document (Summer 2013 Data Challenge Report)*
- No shear systematics tests using this code yet; just looking at likelihood surfaces, measuring computational performance.
- Next up: use greedy optimization to initialize proposal distributions for importance sampling.
- Currently fast enough that we have time for ~30-50 posterior samples; with a reasonable 6x future speedup from careful optimization on future hardware, we can do 200, which we think will be enough using importance sampling.

DESC's Role in Shape Measurement

Shape measurement is a significant fraction of the LSST Project's compute budget, and the LSST Project's compute budget is considerable. DESC shouldn't plan to re-do all of that with its own compute resources without a *very* good reason.

- If DESC thinks the Project's approach won't be good enough, and has an alternative that is computationally the same or cheaper, DESC should do the R&D necessary to make the case that the Project should change its baseline.
- If DESC thinks the Project's approach won't be good enough, and the alternatives are all much more computationally expensive, DESC should look for the necessary additional funding and work with the Project to make sure we leverage the work both parties will do to minimize the total work.

What can DESC do to help the Project?

- Help test our prototype (and its variants) and determine whether they will be good enough for LSST.
- Explore other shape measurement approaches, and examine them closely from a computational performance perspective.
- Keep working on PSF-related issues: the Project has plans for PSF estimation as well (though not as detailed as the shape measurement plans), and here it's even more crucial for the results of DESC's R&D work to be included in the DM processing, as we'll want to use improved PSF models for much more than just shear estimation.
- Help the Project get the right people. We'll be hiring many new DM developers as construction ramps up, so if you have any smart, software-savvy graduate students....

How the Project can help DESC

- Make it easier to install, use, and develop the DM Software Stack
  - DM work this year is focused on fixing bad designs, improving ease-of-use
  - DESC efforts for documentation are already helpful, but we need better coordination and organization.
  - We have a very useful Q/A forum (from the HSC project): http://hsca.ipmu.jp:8080/account/signin/
- Make the multifit prototype available in the stack!
  - It's not quite there yet. Sorry. Almost there.
- Want a tutorial? Host a DM member for a few days.
  - I can get someone up and running on the stack much more efficiently in person, and the same is true of many other DM developers – the turnaround time on questions and the ease of providing more detailed answers just makes a huge difference.
- What else?