Rubin Observatory

Building and using coadds

Jim Bosch

Princeton / Rubin Observatory DM

with lots of help

Why coadds?

Rubin Observatory

- *Most* Objects are detected on coadds.
  - Some will come from matched difference-image detections
  - *None* will come from single-visit direct detections.
- Difference images will be built by subtracting single-visit images from coadd templates.
- Objects are *almost entirely* deblended on coadds (except in crowded fields; see Colin's talk on Thursday).
- Objects are *mostly* measured on coadds (some are derived from forced photometry on single visits).

Outline

Rubin Observatory

- Flavors of coadds and what we do now
  - Direct coaddition with PSF propagation ✓
  - PSF-matched coaddition ✓
  - Detection/likelihood coaddition **(?)**
  - Kaiser/decorrelated coaddition **(?)**
  - Multi-band and χ^2 coaddition **(!)**
- What's working well and what isn't
  - Artifact rejection on PSF-matched differences ✓
  - Correlated noise propagation **(!)**
  - PSFs and resampling kernels **(!)**
  - PSFs and discontinuities **(!)**
  - PSFs and computational performance **(!)**

✓ implemented
❌ not going to be implemented
(!) will be implemented
(?) under consideration
(!) current problem

Outline

Rubin Observatory

- Future plans (and anti-plans)
  - Cell, per-object, and other managed-discontinuity approaches (and the relationship between that and deblending) **(!)**
  - Likelihood, decorrelated, and multi-band coadds **(?)**
  - Multi-epoch fitting ❌

✓ implemented

❌ not going to be implemented

(!) will be implemented

(?) under consideration

(!) current problem

Flavors of coadds, and what we do now

Direct coaddition ✓

Just resample and combine images

- … with a linear statistic - [weighted] mean, *with no outlier rejection at all;*
- … with per-image (or at least large-spatial-scale) weights, *not per-pixel weights.*

Propagate the PSF model (at a point) by

- evaluating the single-visit PSF models at that point;
- resampling and combining those images with the same weights.

The PSFs *must* be defined to include the pixel responses - coadd pixels aren't rectangles.

The coadd PSF is discontinuous where the set of input images changes.

Direct coaddition ✓

Direct coaddition is not optimal, but it's not bad (it depends on seeing distribution).

- If you weigh by exposure time (there are usually better choices), it's just like a single longer exposure.
- HSC Wide loses 1.7±1.5% in S/N relative to optimal.

PSF-matched coaddition ✓

After warping, match all input images to a model target PSF.

- "Matching": convolve with a difference kernel.
- Target PSF must be larger than any input PSF, so throw away the worst data.

No need to propagate the PSF model: it's just the target PSF.

- No discontinuities in the coadd PSF.
- Can use the same target PSF across bands for consistent colors (see Thu talk).

Outlier rejection *is* [mostly] reasonable when building these.

PSF-matched coaddition ✓

Matching is *theoretically* fine for S/N, but bad in practice.

- The matching kernel moves noise from variance to covariance, and if we propagated and used the full noise covariance matrix, *there be would no loss in resolution.*
- No one ever does that (it's definitely impractical to do *exactly*).
- PSF-matched coadds still lose S/N relative to direct coaddition because you throw away more images.

Detection/likelihood coaddition **(?)**

- Correlating a background-dominated image by its own PSF yields a *detection map* - an image whose pixel values are proportional to the log likelihood of there being an isolated point source centered on that position.
- An optimal coadd detection map can be built by summing single-visit detection maps (adding log likelihoods = multiplying likelihoods).
- Optimal coadd detection for non-isolated and/or extended sources can also be defined in terms of this coadd.
- Actually, *any* operation that assumes background-dominated noise and a static sky can be done on this coadd (along with its propagated PSF) with no loss of information - it's a *sufficient statistic*.

Kaiser/decorrelated coaddition **(?)**

- Detection maps are an *inconvenient* sufficient statistic - they have such highly correlated noise that we can't really treat them as a regular image.
- To solve that problem, we can decorrelate after coaddition. If the input noise and PSFs are spatially constant, and there is no missing data or edges, this is actually very straightforward - in Fourier space.
  - These assumptions are locally fine, but Fourier transforms are not local.
  - This algorithm *is* feasible, but it's full of complex bookkeeping and careful handling of DFT vs continuous FT subtleties.

Multi-band and χ^2 coaddition **(!)**

Rubin Observatory

- The optimal multi-band detection map is a weighted sum of per-band detection maps, where the weights depend on the SED of the object you want to detect.
- A χ^2 coadd (Szalay et al 1999) is a particular weighted sum that's optimal for SED of the sky (because it frames detection as rejecting a null hypothesis that a pixel is sky).
- We might want to make multi-band coadds for specific [idealized] SEDs; this may not be better than χ^2 in practice, but it might make our detection efficiency easier to characterize.
- All of these are easy; guessing what's best for science is hard.

Where we [will] use coadds

Rubin Observatory

	Today	Future
Object detection	direct	direct or likelihood, multi-band
Object deblending	direct	direct or decorrelated
Object measurement	direct	direct or decorrelated; +PSF-matched?
Subtraction templates for DIASources	direct	direct, PSF-matched, decorrelated, or DCR-corrected (Ian's talk Thur.)
Subtraction templates for artifact rejection	PSF-matched	PSF-matched

What's working well, and what isn't

CompareWarp: temporal artifact masking ✓

Rubin Observatory

Figure from Yusra AlSayyad

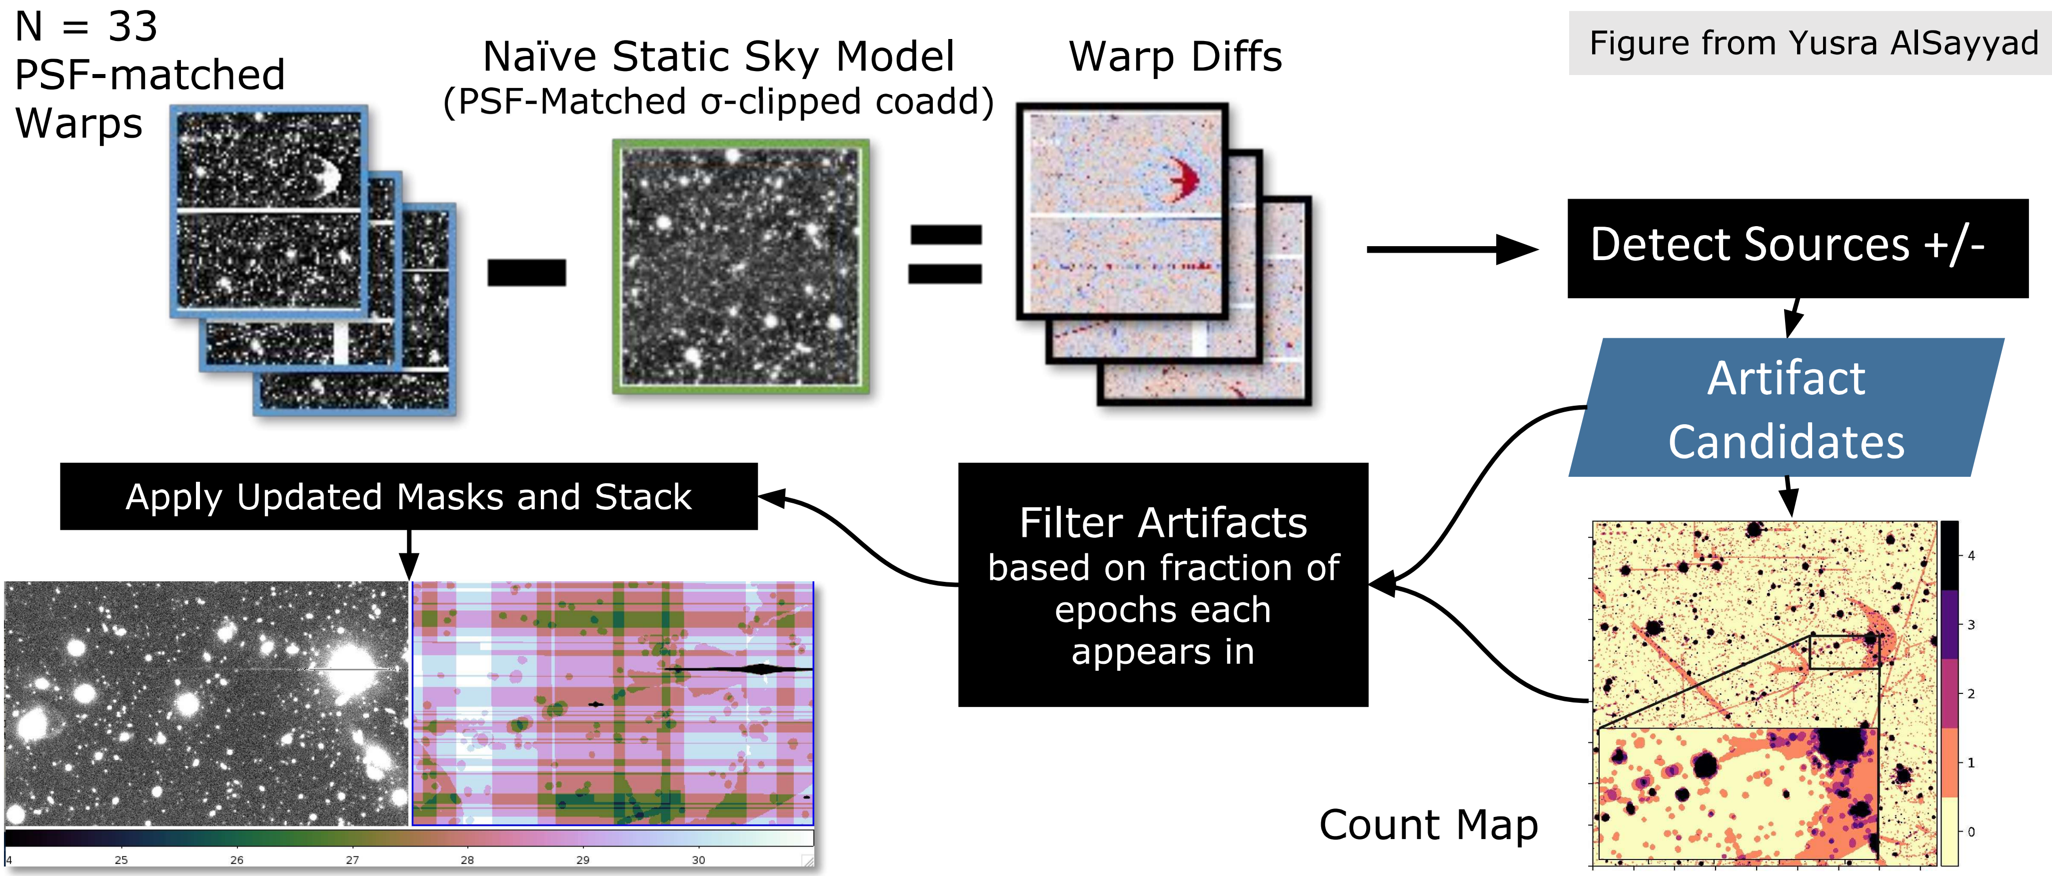

CompareWarp: temporal artifact masking ✓

Rubin Observatory

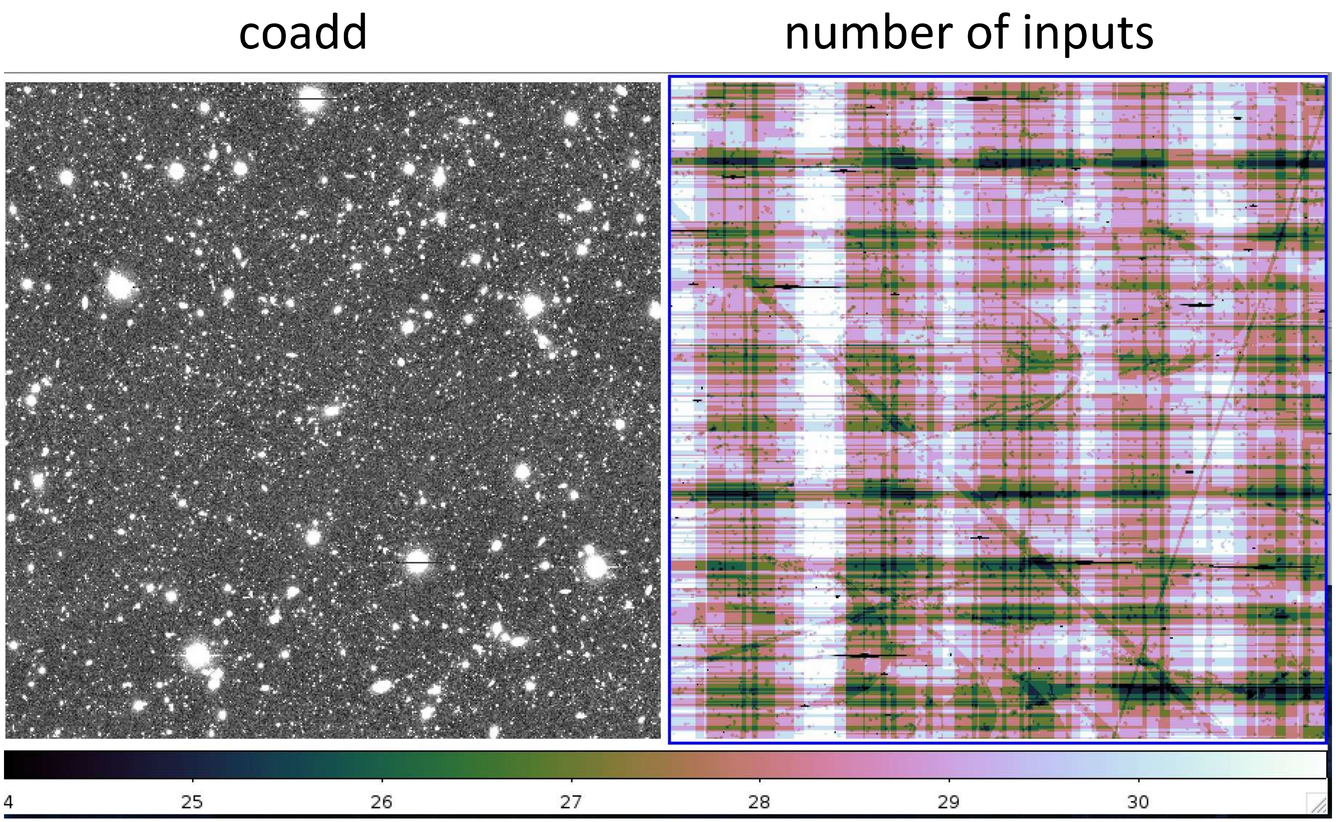

HSC-I
tract=9813 (COSMOS UD)
patch=6,4

Figure from Yusra AlSayyad

CompareWarp failure modes

Rubin Observatory

Insufficient dithers make ghosts land on same sky position and look persistent

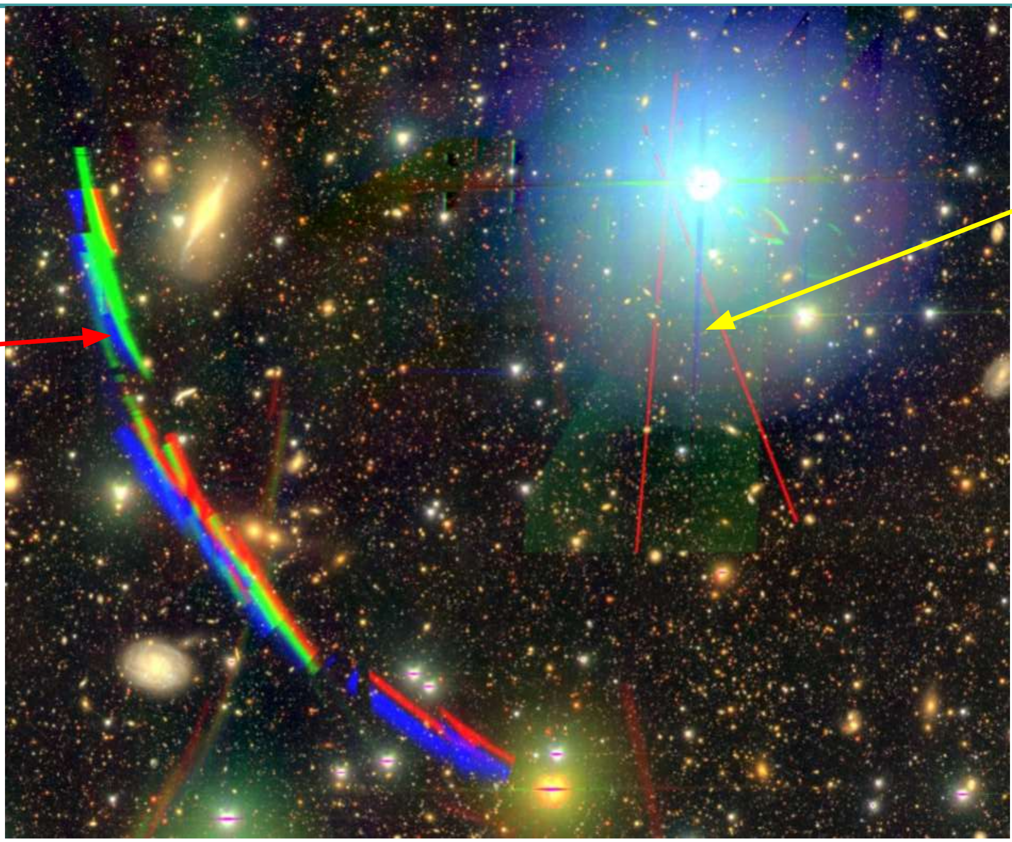

Bright stars are hard to subtract and look like variable sources, so they are often labeled persistent. If artifacts overlap these, they are also labeled persistent.

Remaining optical ghosts and satellite trails in HSC UD-COSMOS. The image is in gri and is approximately 23'x19'.
https://hsc-release.mtk.nao.ac.jp/doc/index.php/known-problems-2/

Figure from Yusra AlSayyad, NAOJ

Artifact rejection next steps: morphology

Rubin Observatory

CompareWarp is almost entirely about temporal detection.

- To improve significantly on it, we need to add morphological information.
- Morphological information will be even more necessary for image differencing - anything we remove from coadds with CompareWarp will by definition appear in difference images.
- We already do cosmic ray detection (earlier in the pipeline) morphologically; it's close, but it needs tuning and tweaking.
- Ghost prediction is in progress (for HSC, at NAOJ), and satellite trail detection is mostly working and nearly integrated.

Satellite detection preview

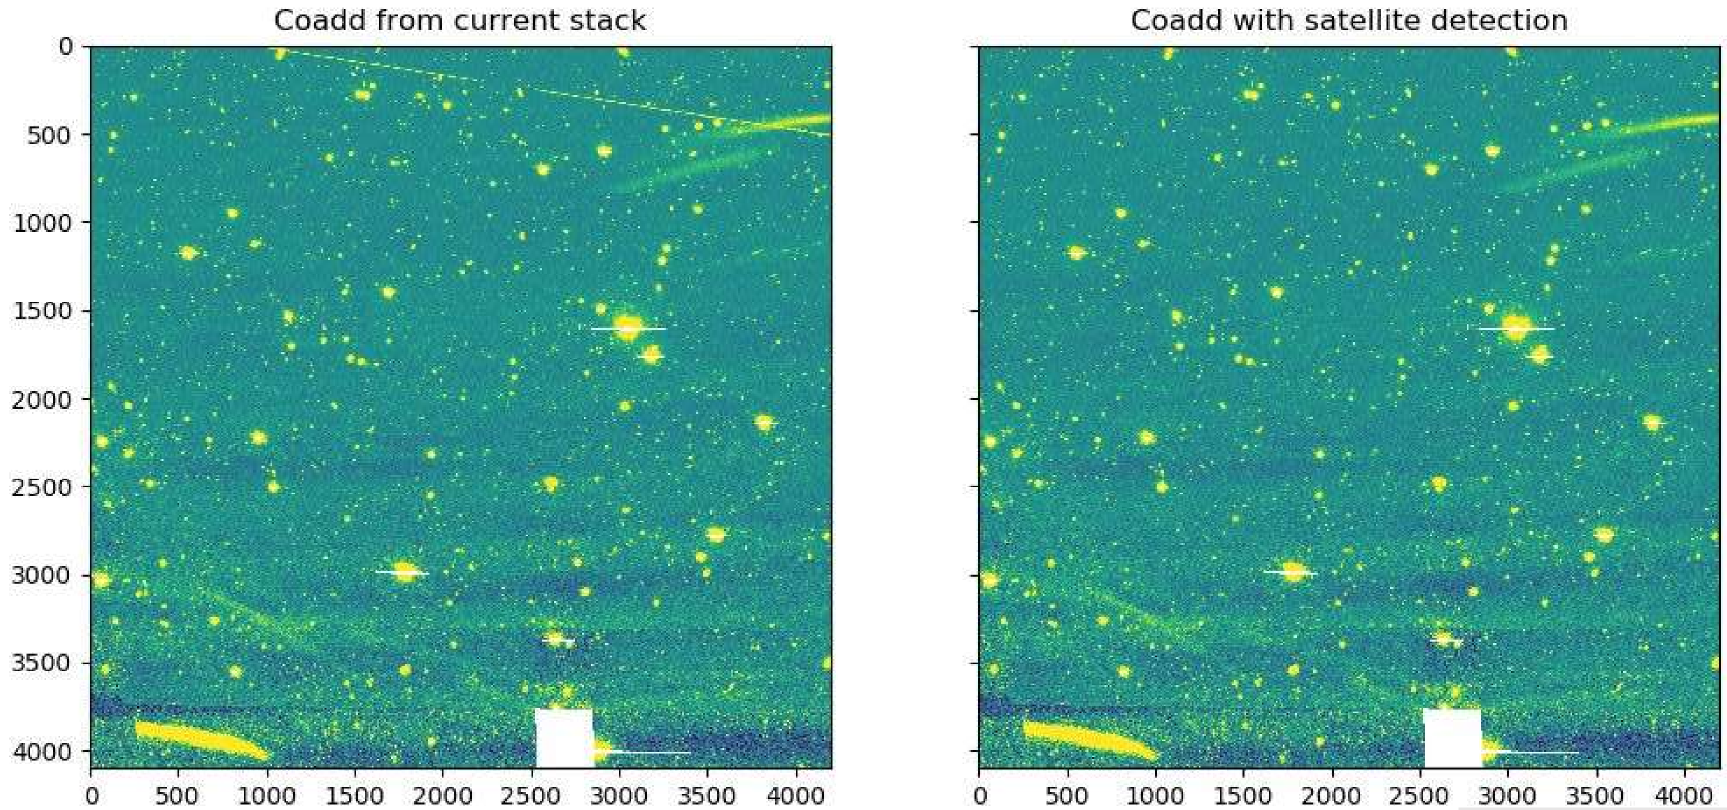

Figure from Clare Saunders

Correlated noise **(!)**

Noise is correlated in our coadds, but our algorithms and data structures only propagate per-pixel variance.

- Results in uncertainty underestimates for *all* coadd measurements.
- True sinc interpolation on data with background-dominated noise should not yield correlations, but (of course) we actually use Lanczos.
- We should be able to approximately propagate the covariance.
- We may also be able to use a better (for this purpose) resampling kernel (especially with Kaiser/decorrelated coaddition).

More on this topic in my galaxy photometry talk on Thursday.

PSFs and resampling **(!)**

Rubin Observatory

Our PSF metrics don't look as good on coadds as they do on individual images.

More on this in my shear estimation talk on Thursday.

PSFs and discontinuities **(!)**

Rubin Observatory

CCD boundaries and masking produce PSF discontinuities on the coadd. We know where those are, and we set a mask bit on any Object that lands on or near a discontinuity.

Many science cases won't care about these discontinuities at all, because they don't care about the PSF in excruciating detail; they can just ignore this flag.

PSFs and discontinuities **(!)**

Rubin Observatory

This works okay for the HSC Wide survey (5-6 epochs deep), where the number of Objects flagged this way is ~18%. It will not work for a survey ~20 (Y1) or ~200 (Y10) epochs deep.

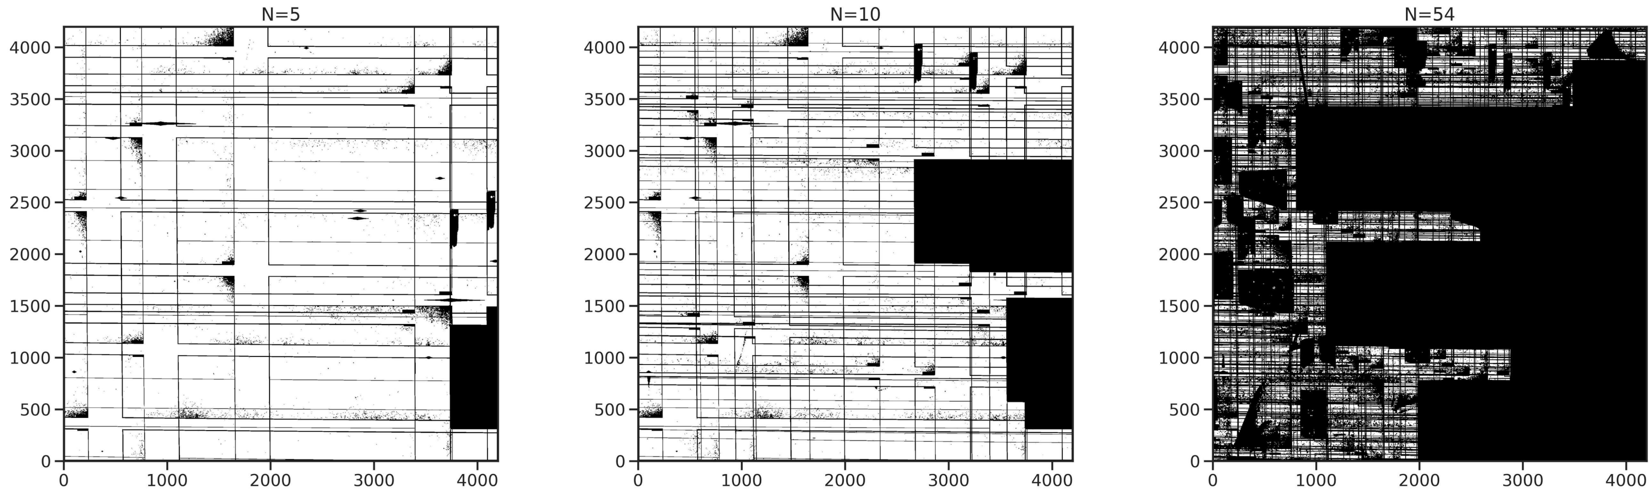

Figure by Yusra AlSayyad

PSFs and performance optimization **(!)**

Rubin Observatory

Every time we evaluate the PSF at a point on the coadd, we:

1. transform the coadd position to pixel coordinates on each input image;
2. evaluate the input-image PSF at that point;
3. resample that PSF image to the coadd coordinate system;
4. compute a weighted average of the input PSF images.

This makes coadd processing time scale with the number of input images, and when the number of inputs is large, it's the single slowest part of our pipeline.

Future plans and anti-plans

Managing discontinuities: per-Object coadds

1. Use a regular likelihood or direct coadd for detection.
2. Use a regular decorrelated or direct coadd for preliminary deblending and measurement.
3. Make a postage-stamp coadd for each Object. If any input image has a discontinuity or masked pixels *on that postage stamp*, do not include it at all.
4. Do final [deblending and] measurement on the postage-stamp coadd.

Managing discontinuities: cell-based coadds

Rubin Observatory

1. Build coadds in very small (~few arcsecond) cells:
    - Could be direct or likelihood→decorrelated.
    - Reject any input images that have discontinuities or masked pixels in that cell.
    - Propagate PSF by assuming it is constant within a cell.
    - Include PSF-width overlaps for detection smoothing or decorrelation *only*.
2. Run detection, deblending, and measurement on the cells. All algorithms must be able to utilize different PSFs for different cells.
    - Straightforward for detection (use overlaps for smoothing, then threshold).
    - Straightforward for forward modeling (just piecewise model evaluation).

Noise correlations at boundaries are hard to propagate, but easy to Monte Carlo.

Kaiser/decorrelated coaddition **(?)**

Is the extra complexity worth it?

- Probably not worth it just for S/N, unless the seeing distribution is *very* broad.
- Not having to think about which images to reject is very attractive (if Kaiser coaddition can deliver on that).
- Not having to think about correlated noise is even more attractive (again, if Kaiser coaddition can deliver).

Conspicuously absent: multifit

Simultaneous multi-epoch fitting ("multifit") has long been described as the way DM will perform its most systematics-limited Object measurements:

- stellar motions
- galaxy models (and shear)

This is changing. DM no longer believes multifit is necessary for these measurements, and it's a very inefficient way to use scarce compute resources.

We are working on changing baselined documents that (still) refer to it.

I'll speak more on this subject in later talks.

Extra slides

Comparing Flavors of Coadds

	Direct		PSF-Matched	Kaiser	Multi-Epoch Fitting
	per-pixel weights	per-epoch weights			
PSF	per-pixel / none	discontinuous	trivial (given)	discontinuous	multiple
Sufficient Statistic	slightly lossy	slightly lossier	very lossy	yes	yes
Noise Correlations	resampling only	resampling only	resampling and matching	maybe none?	none
Sampling Requirements	Nyquist	Nyquist	Nyquist	Nyquist (or IMCOM)	none
Edges Problematic?	every pixel is an edge	yes	no	yes	no