



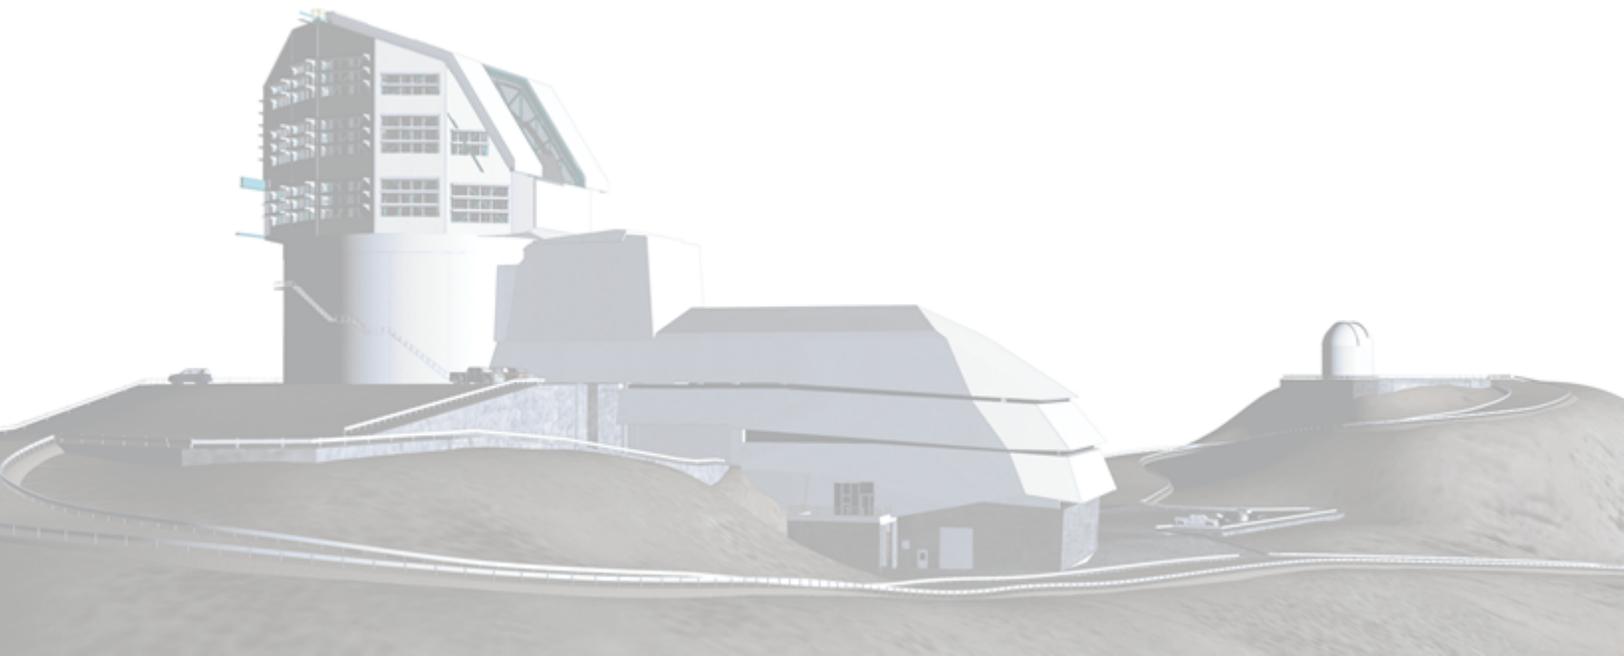
Vera C. Rubin Observatory  
Data Management

# Cell-Based Coaddition

**Jim Bosch**

**DMTN-075**

**Latest Revision: 2021-10-06**



## Abstract

Design sketch and mathematics for a new approach to building coadds and constraining their inputs.

## Change Record

Version	Date	Description	Owner name
1	YYYY-MM-DD	Unreleased.	Jim Bosch

*Document source location:* <https://github.com/lst-dm/dmtn-075>

## Contents

<b>1 Likelihood Coadds</b>	<b>1</b>
1.1 Derivation and starting assumptions . . . . .	1
1.2 Approximate PSFs . . . . .	6
1.3 Full noise propagation . . . . .	9
<b>2 Implementation options</b>	<b>9</b>
2.1 Resampling in advance . . . . .	9
2.1.1 Native-resolution FFT convolution and deconvolution . . . . .	10
2.1.2 Single-kernel image-domain convolution . . . . .	11
2.2 Combining resampling and weights . . . . .	12
2.2.1 GALSIM-style Fourier-domain accumulation . . . . .	12
2.2.2 Analytic single-kernel image-domain resampling and weights . . . . .	12
<b>3 Extensions for relaxed assumptions</b>	<b>12</b>
3.1 Nonstationary noise . . . . .	12
3.2 Undersampled input images . . . . .	12
3.3 Missing pixels . . . . .	12
3.4 Chromatic PSFs . . . . .	12
<b>A References</b>	<b>12</b>
<b>B Acronyms</b>	<b>12</b>

# Cell-Based Coaddition

## 1 Likelihood Coadds

### 1.1 Derivation and starting assumptions

We start by writing the log likelihood of an arbitrary model of the local sky,  $\alpha$ , given multiple images  $z_j$  with differing PSFs  $\phi$ :

$$L \equiv \frac{1}{2} \sum_j \mathbf{R}_j^T \mathbf{C}_j^{-1} \mathbf{R}_j \quad (1)$$

where the residual vectors<sup>1</sup>  $\mathbf{R}_j$  are defined as

$$\mathbf{R}_j[\mathbf{x}_j] \equiv z_j[\mathbf{x}_j] - \int d^2\mathbf{r} \phi_j(\mathbf{A}_j\mathbf{x}_j + \mathbf{b}_j - \mathbf{r}) \alpha(\mathbf{r}) \quad (2)$$

where  $\mathbf{r}$  are coordinates in the projection we will coadd onto, and  $\mathbf{A}_j$  and  $\mathbf{b}_j$  are the linear transform and translation that map input-image coordinates  $\mathbf{x}_j$  to  $\mathbf{r}$ . Note that we define the PSF and model in coadd coordinates, not input-image coordinates. This likelihood is *local* because we have assumed the PSF is not spatially varying and the transformation is affine – a valid assumption in the neighborhood of any particular point, but not over larger scales. This is a common feature of the approaches to coaddition described here: we will explicitly subdivide the sky into small cells whose maximum extent is set by the scales on which these assumptions are valid.

This likelihood is quite general in many respects:

- the input pixel grid may or may not be well-sampled;
- pixels on the same image may have arbitrary noise correlations;
- pixels with no data or artifacts can simply be omitted from the sum, provided they can be identified.

---

<sup>1</sup>When we use matrix notation on image-like entities, we will invariably mean that the 2-d image is flattened into a 1-d vector.

We will need to place restrictions on these to make the method computationally tractable in later steps. In other respects, we have made powerful simplifying assumptions:

- all images have been background subtracted in advance;
- the true sky is completely static ( $\alpha$  is the same for every epoch);
- there is no wavelength dependence at all - so either the true sky is monochromatic, or all PSFs have the same throughput as a function of wavelength, so all wavelength dependency cancels.

We will relax each of these assumptions in later sections.

We now expand the quadratic log likelihood and group by powers of  $\alpha$ :

$$L = \frac{\kappa}{2} - \int d^2\mathbf{r} \Psi(\mathbf{r}) \alpha(\mathbf{r}) + \frac{1}{2} \int d^2\mathbf{r} \int d^2\mathbf{s} \Phi(\mathbf{r}, \mathbf{s}) \alpha(\mathbf{r}) \alpha(\mathbf{s}) \quad (3)$$

with

$$\kappa_c \equiv \sum_j \kappa_j \equiv \sum_j \sum_{\mathbf{x}_j} \sum_{\mathbf{y}_j} C_j^{-1}[\mathbf{x}_j, \mathbf{y}_j] z_j[\mathbf{x}_j] z_j[\mathbf{y}_j] \quad (4)$$

$$\Psi_c(\mathbf{r}) \equiv \sum_j \Psi(\mathbf{r}) \equiv \sum_j \sum_{\mathbf{x}_j} \sum_{\mathbf{y}_j} C_j^{-1}[\mathbf{x}_j, \mathbf{y}_j] z_j[\mathbf{x}_j] \phi_j(\mathbf{A}_j \mathbf{y}_j + \mathbf{b}_j - \mathbf{r}) \quad (5)$$

$$\Phi_c(\mathbf{r}, \mathbf{s}) \equiv \sum_j \Phi(\mathbf{r}, \mathbf{s}) \equiv \sum_j \sum_{\mathbf{x}_j} \sum_{\mathbf{y}_j} C_j^{-1}[\mathbf{x}_j, \mathbf{y}_j] \phi_j(\mathbf{A}_j \mathbf{x}_j + \mathbf{b}_j - \mathbf{r}) \phi_j(\mathbf{A}_j \mathbf{y}_j + \mathbf{b}_j - \mathbf{s}) \quad (6)$$

$$(7)$$

$\Psi_c$  and  $\Phi_c$  are sufficient statistics for our model ( $\kappa_c$  does not constrain the model; it enters only if one is interested in the Bayesian evidence or other goodness-of-fit measures), and they include all sums over input images. But as continuous functions, we can't consider them a coadd, let alone a practical one.

If we choose our output grid such that all input PSFs are well sampled, however, we can replace

$$\phi_j(\mathbf{A}_j \mathbf{x}_j + \mathbf{b}_j - \mathbf{r}) = \sum_{\mathbf{p}} \phi_j(\mathbf{A}_j \mathbf{x} + \mathbf{b}_j - \mathbf{p}) \text{sinc}^2(\mathbf{p} - \mathbf{r}) \quad (8)$$

and obtain

$$\Psi_c(\mathbf{r}) = \sum_{\mathbf{p}} \Psi_c(\mathbf{p}) \text{sinc}^2(\mathbf{r} - \mathbf{p}) \quad (9)$$

$$\Phi_c(\mathbf{r}, \mathbf{s}) = \sum_{\mathbf{p}} \sum_{\mathbf{q}} \Phi_c(\mathbf{p}, \mathbf{q}) \text{sinc}^2(\mathbf{r} - \mathbf{p}) \text{sinc}^2(\mathbf{s} - \mathbf{q}) \quad (10)$$

i.e.  $\Psi_c$  and  $\Phi_c$  are guaranteed to be well-sampled as well, even if the input pixel grids are not. We are still quite far from a practical coadd, however; we want an image as well as a PSF and/or covariance matrix that enter into the likelihood in essentially the same way that a single input image does. That means we want to solve

$$\Psi(\mathbf{r}) = \sum_{\mathbf{p}} \sum_{\mathbf{q}} C_c^{-1}[\mathbf{p}, \mathbf{q}] \phi_c[\mathbf{r} - \mathbf{p}] z_c[\mathbf{q}] \quad (11)$$

$$\Phi(\mathbf{r}, \mathbf{s}) = \sum_{\mathbf{p}} \sum_{\mathbf{q}} C_c^{-1}[\mathbf{p}, \mathbf{q}] \phi_c[\mathbf{r} - \mathbf{p}] \phi_c[\mathbf{s} - \mathbf{q}] \quad (12)$$

for  $\phi_c$ ,  $C_c$ , and  $z_c$ . There is no single unique solution, and in fact no guarantee of any solution for completely general inputs. This approach is one limit of the IMCOM method of ?, which utilizes a weight function that minimizes S/N losses during coaddition along with errors in PSF reconstruction; when there is no S/N loss and the PSF reconstruction is exact, their solution satisfies [11] and [12], while providing diagnostics when no such solution is possible. IMCOM is also extremely computationally expensive, as is any method that relies on direct accumulation of a fully general  $\Phi$ . In this technical note, we will instead focus on methods that start with simplifying assumptions that are at least *mostly* true, and then investigate ways to iteratively and incrementally deal with those assumptions being weakly violated.

The first of these simplifying assumptions is that the input noise in each image is constant:<sup>2</sup>

$$C_j[\mathbf{x}_j, \mathbf{y}_j] \rightarrow \sigma_j^2 \delta[\mathbf{x}_j - \mathbf{y}_j] \quad (13)$$

<sup>2</sup>It should be possible to assume only that the noise is stationary, but the algebra is much more involved.

A single-input contribution to  $\Phi_c$  is then

$$\Phi_j(\mathbf{r}, s) = \frac{1}{\sigma_j^2} \sum_{\mathbf{x}_j} \phi_j(\mathbf{A}_j \mathbf{x}_j + \mathbf{b}_j - \mathbf{r}) \phi_j(\mathbf{A}_j \mathbf{x}_j + \mathbf{b}_j - s) \quad (14)$$

with Fourier transform

$$\widetilde{\Phi}_j(\mathbf{u}, \mathbf{v}) = \frac{1}{\sigma_j^2} \widetilde{\phi}_j(-\mathbf{u}) \widetilde{\phi}_j(-\mathbf{v}) e^{-2\pi i \mathbf{b}^T (\mathbf{u} + \mathbf{v})} \sum_{\mathbf{x}_j} e^{-2\pi i \mathbf{x}_j^T \mathbf{A}_j^T (\mathbf{u} + \mathbf{v})} \quad (15)$$

To simplify further, we'd like to be able to apply the Poisson summation formula to the sum over  $\mathbf{x}_j$ , but we need to be careful about the summation limits; so far we've implicitly (and vaguely) summed over all input pixels relevant to a cell, excluding any completely missing pixels. To instead sum over  $\mathbb{Z}^2$ , we need the original sum in [2] to truncate outside some finite region by defining both  $z_j$  and  $\alpha$  to be zero outside that region, and we need to assume there are no missing pixels. Neighboring cells will necessarily define truncation inconsistently from each other, so we will also need to build the coadd for a cell with some padding, such that coadd pixels within the cell are (to good approximation) not affected by the truncation, and thus consistent with their neighbors. How much padding will depend on the details of the algorithm, and the degree to which we aim for compactness in the real domain vs. compactness in the Fourier domain.

With these qualifications,  $\widetilde{\Phi}_j$  becomes

$$\widetilde{\Phi}_j(\mathbf{u}, \mathbf{v}) = \frac{1}{\sigma_j^2} \widetilde{\phi}_j(-\mathbf{u}) \widetilde{\phi}_j(-\mathbf{v}) e^{-2\pi i \mathbf{b}^T (\mathbf{u} + \mathbf{v})} \sum_{\mathbf{k}_j \in \mathbb{Z}^2} \delta(\mathbf{k}_j - \mathbf{A}_j^T (\mathbf{u} + \mathbf{v})) \quad (16)$$

But  $\phi_j$  is band-limited on the coadd pixel grid, so  $\widetilde{\phi}_j$  is zero if  $|\mathbf{u}|_1 \geq \frac{1}{2}$  or  $|\mathbf{v}|_1 \geq \frac{1}{2}$ , and we can assume  $|\mathbf{u} + \mathbf{v}|_1 < 1$ . Considering the case where  $\mathbf{A}_j$  is a uniform scaling for simplicity:

- if  $|\mathbf{A}_j| < 1$ , the coadd pixels are larger than the input pixels; only the  $\mathbf{k}_j = 0$  term in the sum survives (note that the input image then must be well-sampled, because the coadd is).
- if  $|\mathbf{A}_j| > 1$ , the coadd pixels are smaller than the input pixels, and other terms may survive only if  $\mathbf{A}_j$  is greater than twice the actual band limit of the PSF – or, in other words, only if the input image is not well sampled.

Our final simplifying assumption is naturally that the input images are well-sampled; this yields

$$\widetilde{\Phi}_j(\mathbf{u}, \mathbf{v}) = \frac{1}{\sigma_j^2} \widetilde{\phi}_j(-\mathbf{u}) \widetilde{\phi}_j(-\mathbf{v}) e^{-2\pi i \mathbf{b}^T (\mathbf{u} + \mathbf{v})} \delta(-\mathbf{A}_j^T (\mathbf{u} + \mathbf{v})) \quad (17)$$

$$= \frac{1}{\sigma_j^2} \left| \widetilde{\phi}_j(\mathbf{u}) \right|^2 \delta(\mathbf{u} + \mathbf{v}) \quad (18)$$

If we then accumulate the contributions from each input image in the Fourier domain, the delta function factors out of the sum, and we can write

$$\widetilde{\Phi}_c(\mathbf{u}, \mathbf{v}) = \delta(\mathbf{u} + \mathbf{v}) \sum_j \frac{1}{\sigma_j^2} \left| \widetilde{\phi}_j(\mathbf{u}) \right|^2 = \delta(\mathbf{u} + \mathbf{v}) \frac{1}{\sigma_c^2} \left| \widetilde{\phi}_c(\mathbf{u}) \right|^2 \quad (19)$$

This is essentially the continuous Fourier version of [12], but here the solution is trivial, and in many respects ideal:

$$\widetilde{\phi}_c(\mathbf{u}) = \sigma_c \sqrt{\widetilde{\Phi}_c(\mathbf{u}, \mathbf{u})} = \sigma_c \sqrt{\sum_j \frac{1}{\sigma_j^2} \left| \widetilde{\phi}_j(\mathbf{u}) \right|^2} \quad (20)$$

$$\frac{1}{\sigma_c} = \int d^2 \mathbf{u} \sqrt{\sum_j \frac{1}{\sigma_j^2} \left| \widetilde{\phi}_j(\mathbf{u}) \right|^2} \quad (21)$$

Note that

- the coadd PSF  $\phi_c$  is the Fourier-domain square root of  $\Phi_c$ ;
- the noise in the coadd is also constant, and its standard deviation  $\sigma_c$  can be derived as the constant that normalizes the coadd PSF  $\phi_c$ .

With these in hand, we can compute the Fourier-domain coadd image  $z_c$  from the Fourier transform of  $\Psi_c$ :

$$\widetilde{z}_c(\mathbf{u}) = \frac{\widetilde{\Psi}_c(\mathbf{u}) \sigma_c^2}{\widetilde{\phi}_c(\mathbf{u})} \quad (22)$$

Obtaining  $\widetilde{\Psi}_c$  from  $z_j$  and  $\phi_j$  involves a continuous Fourier transform of either noisy, sampled data or something derived from it; a subtly difficult task. But it is also possible to reframe this

as a linear operation in the image domain:

$$z_c(\mathbf{r}) = \sum_j \sum_{\mathbf{x}_j} K_j(\mathbf{r} - \mathbf{A}_j \mathbf{x}_j - \mathbf{b}_j) z_j[\mathbf{x}_j] \quad (23)$$

$$\tilde{K}_j(\mathbf{u}) \equiv \frac{\sigma_c^2 \tilde{\phi}_j^*(\mathbf{u})}{\sigma_j^2 \tilde{\phi}_c(\mathbf{u})} \quad (24)$$

This simple result still hides a lot of practical complexity, because  $K$  is

- generally not compact in the image domain (in fact, it is band-limited, and hence *formally* must have infinite extent in the image domain);
- potentially expensive to evaluate at the many non-integer points implied by 23.

Overall, it is unclear whether we would be best served by approach that approximates the continuous Fourier transforms with DFTs (with all of the boundary-condition and folding issues that always entails) or one that works in the image domain, perhaps with analytic approximations or basis functions in the representation of  $\phi$  that have analytic Fourier transforms. We will explore several possibilities in Section 2.

## 1.2 Approximate PSFs

In many case we may not want to use or may not be able to use the best available model of the PSF when building a coadd. Using an approximation formally breaks the property that the coadd is a sufficient statistic for the static sky, but the actual information loss could nevertheless be quite small (much smaller than the loss due to building naive direct coadds in all cases). In many of the implementation options we discuss later in Section 2, using an approximation with an analytic Fourier transform is required. In other cases:

- The approximation may be more compact in the Fourier and/or image domain.
- The approximation may be less affected by measurement noise or numerical instability, both of which can cause trouble with deconvolution operations.
- The true PSF may be chromatic, so even our best model evaluated with some nominal SED is only approximately right for other SEDs.

- We may have already built the coadd with a PSF we now know to be subtly wrong, and want to understand systematic errors this may have introduced into later processing and/or update the coadd without starting over.

Because the coadd likelihood will no longer be exactly equivalent to the original joint single-epoch likelihood with this approximate, we will instead start with [23] as a prescription for building a coadd, but replace the true (or best-estimate) per-epoch PSFs  $\phi_j$  with deliberate approximations  $p_j$ , and then derive the new effective coadd PSF.

We start by expanding  $\tilde{K}$  with that substitution:

$$\tilde{K}_j(\mathbf{u}) \equiv \frac{\sigma_c^2 \tilde{p}_j^*(\mathbf{u})}{\sigma_j^2 \sqrt{\sum_n \frac{1}{\sigma_n^2} |\tilde{p}_n(\mathbf{u})|^2}} \quad (25)$$

Note that  $\sigma_c$  is at this stage just an unknown normalization constant; we will show later that it is equal to the noise on the coadd when the effective coadd PSF is normalized.

To compute the effective PSF of the coadd, we coadd the image of a point source centered at  $\mathbf{r} = 0$  using [23]:

$$\phi_c(\mathbf{r}) = \sum_j \sum_{\mathbf{x}_j} K_j(\mathbf{r} - \mathbf{A}_j \mathbf{x}_j - \mathbf{b}_j) \phi_j(\mathbf{A}_j \mathbf{x}_j + \mathbf{b}_j) \quad (26)$$

It is important here that we use the true PSFs  $\phi_j$  (or our best-available model). We are also continuing to rely on our assumption that all input images are well-sampled - without that,  $K$  would need to be a position-dependent kernel to have a well-defined PSF on the coadd. That assumption also lets us perform this convolution in coadd coordinates<sup>3</sup>:

$$\phi_c(\mathbf{r}) = \sum_j \frac{1}{|A_j|} \sum_s K_j(\mathbf{r} - \mathbf{s}) \phi_j(\mathbf{s}) \quad (27)$$

We can now solve for  $\sigma_c$  by requiring

$$\int \phi_c(\mathbf{r}) d^2\mathbf{r} = 1 \quad (28)$$

---

<sup>3</sup>TODO: prove this by a Fourier round-trip; we get another sum where only  $\mathbf{k} = 0$  survives.

The noise covariance of the coadd can be computed via straightforward uncertainty propagation:

$$C_c(\mathbf{r}, \mathbf{s}) = \sum_j \sigma_j^2 \sum_{\mathbf{x}_j} K_j(\mathbf{r} - \mathbf{A}_j \mathbf{x}_j - \mathbf{b}_j) K_j(\mathbf{s} - \mathbf{A}_j \mathbf{x}_j - \mathbf{b}_j) \quad (29)$$

Its Fourier transform is

$$\widetilde{C}_c(\mathbf{u}, \mathbf{v}) = \sum_j \sigma_j^2 \widetilde{K}_j(\mathbf{u}) \widetilde{K}_j(\mathbf{v}) e^{-2\pi i \mathbf{b}_j^T (\mathbf{u} + \mathbf{v})} \sum_{\mathbf{x}_j} e^{-2\pi i \mathbf{x}_j^T \mathbf{A}_j^T (\mathbf{u} + \mathbf{v})} \quad (30)$$

$$= \sum_j \sigma_j^2 \left| \widetilde{K}_j(\mathbf{u}) \right|^2 \delta(\mathbf{u} + \mathbf{v}) \quad (31)$$

where the simplification comes from using the Poisson summation formula and rejecting modes other than  $\mathbf{k} = 0$ , just as in [18]. Substituting the definition of  $\widetilde{K}$  from [25], this simplifies further to

$$\widetilde{C}_c(\mathbf{u}, \mathbf{v}) = \left[ \sum_j \frac{\sigma_c^2}{\sigma_j^2} |\widetilde{p}_j(\mathbf{u})|^2 \right] \left[ \sum_n \frac{1}{\sigma_n^2} |\widetilde{p}_n(\mathbf{u})|^2 \right]^{-1} = \sigma_c^2 \quad (32)$$

Even with the PSF approximated, the noise in the coadd is uncorrelated and stationary.

With this formalism in hand, we can now look more closely at how this coadd algorithm compares to those in broad use in astronomy. A “direct” coadd is built by simply resampling all input images to the same grid, and combining them with a weighted mean. The weights must be very slowly-varying (often a per-epoch constant), and the combination must use a linear operator (no means or sigma-clipping) for the coadd to have a well-defined effective PSF. When the weights are proportional to the exposure time, this approach is statistically equivalent to taking a single longer exposure (at least for the ideal, static sky), but it is not optimal when the input images have different PSFs. In fact, we can see here that it is equivalent to using a delta function for all PSF approximations  $p_j$  (and kernels  $K_j$ ), or a constant weight in the Fourier domain; the optimal coadd with  $p_j \rightarrow \phi_j$  instead gives more weight to high spatial frequency contributions from the best-seeing images.

PSF-matched coadds, in which each input image is degraded to a target PSF that is approximately the same as the PSF of the worst-seeing image, are naturally even worse from a signal-to-noise standpoint. This makes sense in our formalism: these define  $\widetilde{K}_j$  such that  $\widetilde{p}_j$  appears in the denominator, not the numerator, and hence they downweight the high-frequency

modes from the better-seeing images even more than the constant Fourier weighting of direct coadds.

### 1.3 Full noise propagation

While the algorithm we have derived assumes stationary, uncorrelated noise, we can still use it coadd images with more complex noise if we are willing to ignore the structure of that noise when determining the weights. Just like PSF approximations, this means our coadds will no longer be a sufficient statistic for the static sky, but the loss of information should be quite small in practice.

To fully propagate the noise, we can use the same formula as [29], adjusted to use a full covariance matrix  $C_j$  for each input image:

$$C_c[\mathbf{r}, \mathbf{s}] = \sum_j \sum_{\mathbf{x}_j} \sum_{\mathbf{y}_j} K_j(\mathbf{r} - \mathbf{A}_j \mathbf{x}_j - \mathbf{b}_j) K_j(\mathbf{s} - \mathbf{A}_j \mathbf{y}_j - \mathbf{b}_j) C_j[\mathbf{x}_j, \mathbf{y}_j] \quad (33)$$

Unfortunately, even if the noise is just nonstationary (i.e. still uncorrelated), as is common in astronomical images, the noise on the coadd ends up both nonstationary and correlated, leaving [33] quite expensive in both compute and storage.

In practice, it will often be sufficient to propagate only the diagonal  $C_c[\mathbf{r}, \mathbf{r}]$  (or even just the original constant approximation  $\sigma_c$ ), and possibly augment this with a few Monte Carlo noise fields: noise images that are coadded in exactly the same way as the true input images. Even extremely demanding (in terms of systematic error control) science applications such as shear estimation for weak lensing can often get by with just one Monte Carlo noise realization (TODO citation).

## 2 Implementation options

### 2.1 Resampling in advance

In this section, we consider implementation options that assume all input images have already been resampled to the coadd grid in advance, presumably via a compact image-domain interpolant that approximates sinc (e.g. Lanczos). This inevitably introduces aliasing, which is

traditionally ignored; we will ignore it here as well, but do consider it to be one of the disadvantages of these approaches.

### 2.1.1 Native-resolution FFT convolution and deconvolution

In this approach, input data images  $z_j$  are FFT'd, which means we must first zero-pad them and apply some sort of multiplicative ramp to a buffer region between the output data region and the zero region, because the presence of sharp edges would cause problems. Best-available PSF models  $\phi_j$  and approximations  $p_j$  are either rendered directly in the Fourier domain or rendered as images and also FFT'd onto the same (discrete) Fourier grid. We then accumulate the following sums:

$$\widetilde{\Psi}_c[\mathbf{u}] = \sum_j \frac{\widetilde{p}_j^*[\mathbf{u}] \widetilde{z}_j[\mathbf{u}]}{\sigma_j^2} \quad (34)$$

$$\widetilde{\Phi}_c^{(1)}[\mathbf{u}] = \sum_j \frac{\widetilde{p}_j^*[\mathbf{u}] \widetilde{\phi}_j[\mathbf{u}]}{\sigma_j^2} \quad (35)$$

$$\widetilde{\Phi}_c^{(2)}[\mathbf{u}] = \sum_j \frac{\widetilde{p}_j^*[\mathbf{u}] \widetilde{p}_j[\mathbf{u}]}{\sigma_j^2} \quad (36)$$

The unnormalized coadd PSF model  $\widetilde{\phi}_c$  is then

$$\frac{\widetilde{\phi}_c[\mathbf{u}]}{\sigma_c^2} = \frac{\widetilde{\Phi}_c^{(1)}[\mathbf{u}]}{\sqrt{\widetilde{\Phi}_c^{(2)}[\mathbf{u}]}}; \quad (37)$$

we can then inverse FFT and solve for  $\sigma_c$  via

$$\sum_{\mathbf{r}} \phi_c[\mathbf{r}] = 1. \quad (38)$$

Finally, the Fourier-domain coadd image is

$$\widetilde{z}_c[\mathbf{u}] = \frac{\widetilde{\Psi}_c[\mathbf{u}] \sigma_c^2}{\sqrt{\widetilde{\Phi}_c^{(2)}[\mathbf{u}]}}; \quad (39)$$

which we also inverse FFT to form the coadd image.

This approach is in many respects the simplest possible implementation, especially if one already has an optimized pipeline for image resampling. It is also probably among the fastest: there are no expensive image-domain convolutions, and the padding necessary for the FFT is *probably* only proportional to the PSF size, not the size of the coadd cell.

The main downside of this approach is its nonlocality in the image domain. The FFT division by  $\sqrt{\Phi_c^{(2)}[\mathbf{u}]}$  in particular combines information from all pixels when computing the value of any other pixel, so even output pixels far from the cell edge will be affected by our artificial multiplicative ramp to some degree. The same is true of bad or missing pixels that are interpolated in the input images; their effect cannot be confined to some region around the interpolation.

It may be that this is not a serious problem in practice; the operations may be “local enough” in the image domain for any systematic errors introduced to be negligible. But this is largely an experimental question, and it certainly depends on the profiles of the input PSFs and their approximations; the challenges come from the kinds of messy situations that are often hard to simulate, let alone analyze mathematically.

Because this method never actually constructs the weighting operators  $K_j$  directly, there is no direct way to propagate a full noise covariance matrix with it via [33]; if that kind of noise propagation is desired one would instead have to construct a compact  $K_j$  using a different approach (e.g. Section 2.1.2) that is not quite the same as the one effectively used here. The same is true of the extensions in Section 3 that make use of an image-domain  $K_j$ .

## 2.1.2 Single-kernel image-domain convolution

This method is in many respects the opposite of the previous one, though it also assumes the input data has been resampled. We work entirely in the image domain; we compute approximations to  $K_j$  and then use [23] and [27] to directly accumulate the coadd image and effective PSF.

In order to make these discrete convolution operations efficient, we need to define  $K_j$  approximations with strictly finite support in the image domain (even if it is derived from an analytic form that does not truncate, we need to render images from it that do). This directly addresses the locality problems with FFT methods, like the one described in Section 2.1.1;

zero-padding and artificial ramps are unnecessary. We do need to expand the region we convolve by a buffer the size of  $K_j$ , but after doing so we can be confident that the original target region is fully unaffected by artifacts or missing pixels beyond that padding.

The challenge with this approach is in constructing the  $K_j$  approximations; finite support in the image domain makes it impossible to formally also make  $K_j$  band-limited, but we need to keep aliasing due to high-frequency power in  $K_j$  to a minimum.

## 2.2 Combining resampling and weights

### 2.2.1 GALSIM-style Fourier-domain accumulation

### 2.2.2 Analytic single-kernel image-domain resampling and weights

## 3 Extensions for relaxed assumptions

### 3.1 Nonstationary noise

### 3.2 Undersampled input images

### 3.3 Missing pixels

### 3.4 Chromatic PSFs

## A References

## B Acronyms

---

Acronym	Description
DM	Data Management
DMTN	DM Technical Note

---

---

FFT	Fast Fourier Transform
PSF	Point Spread Function
SED	Spectral Energy Distribution

---