

# Parallel Discrete Convolutions on Adaptive Particle Representations of Images

Joel Jonsson<sup>1</sup>, Bevan L. Cheeseman, Suryanarayana Maddu, Krzysztof Gonciarz<sup>2</sup>, and Ivo F. Sbalzarini<sup>3</sup>

**Abstract**—We present data structures and algorithms for native implementations of discrete convolution operators over Adaptive Particle Representations (APR) of images on parallel computer architectures. The APR is a content-adaptive image representation that locally adapts the sampling resolution to the image signal. It has been developed as an alternative to pixel representations for large, sparse images as they typically occur in fluorescence microscopy. It has been shown to reduce the memory and runtime costs of storing, visualizing, and processing such images. This, however, requires that image processing natively operates on APRs, without immediately reverting to pixels. Designing efficient and scalable APR-native image processing primitives, however, is complicated by the APR's irregular memory structure. Here, we provide the algorithmic building blocks

required to efficiently and natively process APR images using a wide range of algorithms that can be formulated in terms of discrete convolutions. We show that APR convolution naturally leads to scale-adaptive algorithms that efficiently parallelize on multi-core CPU and GPU architectures. We quantify the speedups in comparison to pixel-based algorithms and convolutions on evenly sampled data. We achieve pixel-equivalent throughputs of up to 1TB/s on a single Nvidia GeForce RTX 2080 gaming GPU, requiring up to two orders of magnitude less memory than a pixel-based implementation.

**Index Terms**—Biomedical image processing, image filtering, image representation, image restoration, data compression, data structures / octrees, parallel processing, convolution, deconvolution.

Manuscript received August 10, 2021; revised March 28, 2022; accepted May 20, 2022. Date of publication June 14, 2022; date of current version June 20, 2022. This work was supported in part by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy through the Cluster of Excellence "Physics of Life" of TU Dresden under Grant EXC-2068-390729961; and in part by the Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI), Dresden/Leipzig, funded by the Bundesministerium für Bildung und Forschung (BMBF, Federal Ministry of Education and Research). The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Nelly Pustelnik. (Corresponding author: Ivo F. Sbalzarini.)

Joel Jonsson is with the Faculty of Computer Science, Technische Universität Dresden, 01069 Dresden, Germany, also with the Max Planck Institute of Molecular Cell Biology and Genetics, 01307 Dresden, Germany, also with the Center for Systems Biology Dresden, 01307 Dresden, Germany, and also with the Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI) Dresden/Leipzig, 01069 Dresden, Germany.

Bevan L. Cheeseman is with the Faculty of Computer Science, Technische Universität Dresden, 01069 Dresden, Germany, also with the Max Planck Institute of Molecular Cell Biology and Genetics, 01307 Dresden, Germany, also with the Center for Systems Biology Dresden, 01307 Dresden, Germany, and also with ONI Inc., Oxford OX2 8TA, U.K.

Suryanarayana Maddu is with the Faculty of Computer Science, Technische Universität Dresden, 01069 Dresden, Germany, also with the Max Planck Institute of Molecular Cell Biology and Genetics, 01307 Dresden, Germany, also with the Center for Systems Biology Dresden, 01307 Dresden, Germany, also with the Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI) Dresden/Leipzig, 01069 Dresden, Germany, also with the Flatiron Institute, New York, NY 10010 USA, and also with the NSF-Simons Center for Mathematical and Statistical Analysis of Biology, Faculty of Arts and Sciences, Harvard University, Cambridge, MA 02138 USA.

Krzysztof Gonciarz is with the Faculty of Computer Science, Technische Universität Dresden, 01069 Dresden, Germany, also with the Max Planck Institute of Molecular Cell Biology and Genetics, 01307 Dresden, Germany, and also with the Center for Systems Biology Dresden, 01307 Dresden, Germany.

Ivo F. Sbalzarini is with the Faculty of Computer Science and the Cluster of Excellence Physics of Life, Technische Universität Dresden, 01069 Dresden, Germany, also with the Max Planck Institute of Molecular Cell Biology and Genetics, 01307 Dresden, Germany, also with the Center for Systems Biology Dresden, 01307 Dresden, Germany, and also with the Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI) Dresden/Leipzig, 01069 Dresden, Germany (e-mail: ivos@mpi-cbg.de).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TIP.2022.3181487>, provided by the authors.

Digital Object Identifier 10.1109/TIP.2022.3181487

## I. INTRODUCTION

**F**LUORESCENCE microscopy enables long-term imaging of biological specimen at high spatial and temporal resolution with morphologically or biochemically specific labeling. This enables researchers to study biological structures and processes in living organisms [1]–[3], providing an invaluable source of information. However, handling and analyzing the large amounts of data produced causes significant computational demands, especially for three-dimensional (3D) images. Light-sheet microscopes, e.g., can acquire data at rates on the order of 10TB per hour [4], challenging image storage [5], visualization [6], and processing [7]. This “data bottleneck” often limits the throughput and scalability of fluorescence microscopy studies and leads to under-utilization of the information contained in the images.

To relax the data bottleneck, mainly three approaches are followed: (1) parallelization of image processing, (2) multi-resolution image representations, and (3) data compression. Parallelizing large images in distributed-memory computer clusters has enabled real-time segmentation of fluorescence microscopy images [8], and GPU acceleration has enabled interactive handling of large images [9], [10]. Multi-resolution representations and chunked hierarchical file formats aid analysis by allowing zoomable navigation of large volumes [11]. Tools like BigDataViewer [12] and TeraFly [13] have leveraged this to enable interactive visualization and annotation of Terabyte-sized volumetric images. Compression methods, such as the Keller Lab Block (KLB) [14] and B<sup>3</sup>D [15] formats, reduce the cost of storing and transferring data. However, compression does not reduce the runtime of image processing pipelines, as the data must be decompressed to its original size prior to processing.

Adaptive-resolution image representations, such as the Adaptive Particle Representation (APR) [16], provide an alternative to compression that can simultaneously reduce both

data size and processing times. In order to take full advantage of such representations, however, it is necessary to natively process the images in the data-reduced representation, ideally on parallel computer architectures. This enables efficient end-to-end pipelines that never have to revert to pixels, not even block-wise. For certain tasks, like graph-cut segmentation, this has been shown to reduce both memory consumption and runtime [16]. However, biological image analysis workflows typically require a wider range of algorithms. Adapting these algorithms, which have been designed for uniform pixel grids, to content-adaptive image representations, and efficiently implementing them using sparse data structures on parallel computers, is not trivial. Particularly challenging is the efficient implementation on GPUs, which are highly optimized for processing uniform data structures.

Here, we introduce the data structures and algorithms required for natively convolving APR images in parallel on both CPUs and GPUs. Previous formulations of APR convolution [16] were restricted to separable kernels and either required (locally) reconstructing the original pixels prior to convolving, or were limited to first-neighbor kernels. The present work extends APR convolution to dense kernels, proposes strategies for efficient parallel implementation, and provides scale-adaptive filters that operate on multiple resolution levels of an APR without having to reconstruct the finest level everywhere. We demonstrate the efficiency of this multi-scale approach, and we present an example of image deconvolution using the Richardson-Lucy algorithm natively implemented on the APR.

## II. BACKGROUND

For convenience, we recapitulate the fundamental concepts of the APR, but refer to Ref. [16] and its Supplementary Material for more details and for mathematical proofs.

### A. The Adaptive Particle Representation

The APR optimally adapts the local density of the sampling  $\mathcal{P} = \{(\mathbf{x}_p, f(\mathbf{x}_p))\}_{p=1}^{N_p}$  of some function, signal, or image,  $f : \mathbb{R}^n \rightarrow \mathbb{R}$ , defined over  $\Omega \subset \mathbb{R}^n$  so as to guarantee that the point-wise reconstruction error, for all of a wide-class of reconstruction methods  $\hat{f}$ , is bounded everywhere in  $\Omega$  by a user-set constant  $E$  relative to a (potentially spatially varying) local error scale. This places sampling points  $\mathbf{x}_p$  where they are required in order to approximate the unknown continuous  $f$  within the given error threshold at all original pixel locations. Figure 1 illustrates the result for a fluorescence microscopy image. In regions where the signal gradient is significant, the original (pixel) sampling density is retained. However, in regions of low signal gradient, such as the background or homogeneous areas in the interior of objects, the sampling density is reduced.

Mathematically, the APR bounds the infinity norm of the relative reconstruction error

$$\left\| \frac{f - \hat{f}}{\sigma} \right\|_{\infty} \leq E, \quad (1)$$

where  $f$  is the true signal,  $\hat{f}$  the reconstruction, and  $\sigma$  the local error scale. The reconstruction  $\hat{f}$  at location  $\mathbf{x}$  can be




Fig. 1. Illustration of a fluorescence microscopy pixel image and a corresponding Adaptive Particle Representation (APR). The image used is a crop of Hoechst-stained mouse blastocyst cells [17], available as image set BBBC032v1 from the Broad Bioimage Benchmark Collection [18]. The top half of each panel shows the original image (left) and the image reconstructed from the APR (right). The bottom halves show the sampling points (pixels or particles) as dots, colored by sample value and scaled according to their spatial extent. Pixels are uniform across the domain, whereas the APR particles adapt to the content of the image.

any positive weighted combination of sampled function values  $f(\mathbf{x}_p)$  within a certain radius  $R(\mathbf{x})$  of  $\mathbf{x}$ , where the function  $R : \Omega \rightarrow \mathbb{R}$  is called *resolution function*. Thus, for an APR, at any location  $\mathbf{x} \in \Omega$ , any reconstruction of the form

$$\hat{f}(\mathbf{x}) = \sum_{\mathbf{x}_p : |\mathbf{x} - \mathbf{x}_p| \leq R(\mathbf{x})} f(\mathbf{x}_p) w_p, \quad (2)$$

with  $\sum_p w_p = 1$  and  $w_p \geq 0^1$  fulfills the reconstruction condition in Eq. 1. The resolution function can intuitively be seen as a local length scale of the signal.

If we consider a signal originally evenly sampled at  $N$  points then, in general dimensions and for arbitrary  $R(\mathbf{x})$  and  $\mathcal{P}$ , only greedy locally optimal solutions can be found in quadratic runtime  $O(N^2)$ , becoming infeasible even for small problems [16]. In the APR both the resolution function  $R(\mathbf{x})$  and the sampling locations  $\mathbf{x}_p$  are therefore restricted to be power-of-two fractions of the image edge length  $|\Omega|$ . Under these restrictions, globally optimal sampling solutions can be found in linear time  $O(N)$  [16]. The use of power-of-two decompositions is common in image representations and image processing algorithms, such as image pyramids [19], tree-based methods [20], and wavelet decompositions [21].

Considering the analogy of an image pyramid for a function in one dimension, as illustrated in Fig. 2, the APR corresponds to a partition of the domain  $\Omega$  across resolution levels  $l$ , where at each location the coarsest element is selected under the condition that Eq. 1 holds. Intuitively, this can be thought of as a pruning operation of the full tree, where branches in areas of low signal gradient (relative to  $\sigma$ ) are cut, and the sampling values are pushed to coarser ascendant nodes. The APR particles (green dots in Fig. 2B) thus constitute the leaf nodes in a pruned tree structure. The locations of particles are taken to be the centers of the corresponding grid cells (blue intervals in Fig. 2B), which we refer to as *particle cells*. The particle cells partition the image domain, and correspond

<sup>1</sup>The non-negativity constraint can be relaxed, and classes of adaptation satisfying higher-order constraints can also be formulated [16].




Fig. 2. Schematic comparing a regular sampling (pixels) to an APR in one dimension. (A) shows the pixels as green dots, and successively downsampled (by factors of 2) super-pixels as yellow dots. This forms an image pyramid or, by connecting spatially overlapping elements, a full binary tree. The APR particles (green dots in B) define a partition of the image domain and coincide with nodes of this tree at different levels of resolution.

exactly to grid cells in the image pyramid. In this way, they can be described by their resolution level  $l_p$  and a multi-index  $\mathbf{i}_p = (i_{p,1}, \dots, i_{p,n})$  defining the location of the cell in the grid at level  $l_p$ . This allows the location  $\mathbf{x}_p$  of a particle cell to be written in terms of  $\mathbf{i}_p$  and  $l_p$  as

$$\mathbf{x}_p = \mathbf{i}_p \frac{|\Omega|}{2^{l_p}}. \quad (3)$$

For simplicity we assume the image domain  $\Omega$  to be a hypercube starting at  $\mathbf{0}$  with all edges of length  $|\Omega|$  a power of two. For anisotropic domains, or domains that are not a power of two,  $|\Omega|$  is the maximum edge length extended accordingly.

### B. Neighborhood and Reconstruction

The resolution function  $R(\mathbf{x})$  is taken to be a piecewise constant function, defined in terms of the APR particles as:

$$R(\mathbf{x}) = \sum_{p=1}^{N_p} \frac{|\Omega|}{2^{l_p}} \left[ \left\lfloor \frac{\mathbf{x}}{2^{l_p}} \right\rfloor = \mathbf{i}_p \right], \quad (4)$$

where  $[A] = 1$  if  $A$  is true and 0 otherwise. Thus, the value of the resolution function  $R(\mathbf{x})$  at any location  $\mathbf{x} \in \Omega$  is determined solely by the level of the particle cell which contains  $\mathbf{x}$ . This is guaranteed to be unique, since the particle cells partition the image domain  $\Omega$ . In the example of Fig. 2B, evaluating  $R(\mathbf{x})$  corresponds to finding the level  $l$  for which the blue interval containing  $\mathbf{x}$  holds a particle (green dot). The resolution function then takes the value  $R(\mathbf{x}) = \frac{|\Omega|}{2^l}$ , which defines the maximum radius of the reconstruction neighborhood in Eq. 2.

The reconstruction condition in Eq. 1 is valid for any positive weighted combination of particles within the radius  $R(\mathbf{x})$ . This affords a lot of flexibility in defining different reconstruction methods. However, in image processing algorithms we have found that it is often favorable to use the simplest possible method. That is, to reconstruct the signal at a location  $\mathbf{x}$ , we find the particle  $(\mathbf{x}_p^*, f(\mathbf{x}_p^*))$  whose particle cell contains  $\mathbf{x}$  and take  $\hat{f}(\mathbf{x}) = f(\mathbf{x}_p^*)$ . We refer to this as piecewise constant reconstruction. Compared to more general reconstructions, this method is computationally efficient to evaluate, as the reconstructed values do not depend on neighboring particles.

Moreover, this allows an intuitive view of particles and particle cells as pixels of different sizes. In the remainder of this text, unless otherwise stated, all reconstructions are assumed to be piecewise constant.

### C. Determining the APR

Computing the APR from a pixel image consists in finding the largest resolution function  $R(\mathbf{x})$  such that Eq. 1 holds everywhere. However, finding the optimal resolution function, in general, requires algorithms of quadratic time complexity  $O(N^2)$  in the number of pixels  $N$ . In order to solve the problem in linear time, for arbitrary images, the tighter bound

$$R(\mathbf{x}) \leq \min_{\mathbf{y}: |\mathbf{x}-\mathbf{y}| \leq R(\mathbf{x})} L(\mathbf{y}) \quad (5)$$

is considered, where

$$L(\mathbf{y}) = \frac{E\sigma(\mathbf{y})}{|\nabla f(\mathbf{y})|}. \quad (6)$$

Thus, determining the APR requires estimating the gradient magnitude  $|\nabla f|$  of the pixel intensity field, as well as the local error scale  $\sigma(\mathbf{x})$  at the original  $N$  pixels. There is some freedom in choosing how these quantities are computed. In the implementation of Ref. [16], which we also make use of here, the intensity gradient is computed using smoothing cubic B-splines for robust and tunable gradient estimation in the presence of noise [22]. The local error scale should be a sufficiently smooth function reflecting the local range of intensities. We here use a rescaled and smoothed estimate of the local intensity standard deviation, computed by recursive filters on the (downsampled by a factor of 2) B-spline image.

Once computed,  $|\nabla f|$  and  $\sigma$  are combined with the user-provided error bound  $E$  to form  $L$ , which is quantized into a tree structure, and the APR solution is computed using a linear-time algorithm called the *Pulling Scheme* [16]. This algorithm outputs the adaptive tree structure shown in Fig. 2B, which implicitly defines both the resolution function  $R(\mathbf{x})$  and the particle locations  $\{\mathbf{x}_p\}_{p=1}^{N_p}$ .

The final step requires determining the particle intensity values  $f(\mathbf{x}_p)$ . Since particles at the finest resolution coincide with the original pixels, those values remain unchanged. Intensities at coarser particle locations are resampled. This can be done in a number of ways. Here, as well as in [16], coarse particle values are determined by average downsampling the pixel values. This is simple and provides inherited denoising. Viewed as an operation on the tree structure, coarse particle values (i.e., leaf nodes of pruned branches in the tree) are thus obtained by propagating the values of original leaf nodes (i.e., pixels) upward, level by level, taking the average of the combined nodes at each step.

### D. APR-Based Image Processing

Similar to pixel images, processing algorithms can be defined for the APR. For example, the APR naturally lends itself to graph-cut segmentation where the particles constitute the nodes of the graph, and edges are drawn between neighbors. Ref. [16] also introduced two types of convolutions for



the APR, both limited to separable kernels. The first was an *adaptive filtering* algorithm that uses optimized data structures to fetch the neighbors of a particle along a given dimension. This allows the variable distance between neighboring particles at different resolution levels to be taken into account, but is limited to small neighborhoods (direct neighbors in Ref. [16]). Larger kernels were handled by the second method, which applies the separable convolution operation on 2D slices of reconstructed pixel values. The computational cost is reduced by only applying the convolution to pixels that intersect with a particle. The resulting algorithm, however, is not consistent with the pixel operation.

Here, we address all of these issues by providing a data structure that is more appropriate for the APR memory access patterns, especially on the GPU, and algorithms to evaluate convolutions at any resolution level (Section III). We also extend the notion of APR convolution to dense stencils, operating natively on multiple resolution levels to avoid pixel reconstruction (Section IV). Finally, we provide parallel implementations on both CPUs and GPUs (Section V).

### III. DATA STRUCTURES AND ALGORITHMS

We start by detailing data structures and algorithmic strategies that can be used to define and efficiently implement a wide range of image processing algorithms to natively work on the APR. This requires sparse tree data structures with the corresponding iterators, as well as local isotropic patch reconstruction.

#### A. Sparse APR Data Structure

As described in Eq. 3, the location  $\mathbf{x}_p$  of a particle (cell) is defined by its level  $l_p$  and multi-index  $\mathbf{i}_p$  encoding the cell coordinates in the grid at the corresponding resolution level. This definition offers a lot of flexibility in choosing data structures for storing the APR. Here, we base our design on the analogy with image pyramids, as illustrated in Fig. 2B. Since the particle cells partition the image domain, the APR corresponds to a pyramid of disjoint sparse images, where each location  $\mathbf{x} \in \Omega$  is covered by exactly one particle cell at some resolution level. From this perspective, the APR amounts to a set of sparse images at different resolutions, which can be encoded using any sparse array format.

The original APR data structures [16] store the particle values  $\{f(\mathbf{x}_p)\}_{p=1}^{N_p}$  as a single, contiguous vector, while the spatial coordinates  $\{\mathbf{x}_p\}_{p=1}^{N_p}$  are encoded as follows: The levels and all but one of the spatial dimensions are stored in a dense array with each “row” in the contiguous dimension<sup>2</sup> compressed. The sparse compression is done by storing the first and last index of each contiguous block of particles, along with a pointer to the particle values, in a red-black tree structure. This allows for efficient random access via red-black tree search in the sparse dimension.

While efficient random access is important in some applications, we argue that it is not required for most image

<sup>2</sup>In order to stay consistent with the software libraries, we assume that the image dimensions are ordered as  $(z, x, y)$ , where  $y$  is contiguous in memory, i.e., that values at locations  $(z, x, y)$  and  $(z, x, y+1)$  are adjacent in memory.

#### Algorithm 1: Linear APR Access Data Structure

```

class LinearAccess:
    /* minimum and maximum resolution levels */
    l_min, l_max
    /* image dimensions for each level */
    z_dim[], x_dim[], y_dim[]
    /* sparse access vectors */
    y_idx[] // y index of each particle
    xz_end[] // last particle of each row
    level_offset[] // offsets in xz_end per l
    /* helper function to access a given row */
    Function get_row(l, z, x):
        row ← level_offset[l] + z*x_dim[l] + x
        row_begin ← xz_end[row-1]
        row_end ← xz_end[row]
        return row_begin, row_end

```




Fig. 3. Illustration of the linear access data structure in 2D. The APR decomposes into a set of sparse grids, one for each resolution level. Each sparse grid is represented in compressed sparse row (CSR) format. The CSR data structures of each level are concatenated (yellow and green regions of the  $xz\_end$  and  $y\_idx$  vectors), with an additional vector  $level\_offset$  indicating the starting point of each level in the  $xz\_end$ .

processing tasks. Instead, these typically rely on the ability to iterate over neighborhoods of certain, fixed structure. Therefore, we here introduce a simplified data structure that explicitly stores the  $y$ -coordinates. This is equivalent to storing each resolution level  $l$  in compressed sparse row (CSR) format and concatenating the vectors of row offsets and  $y$ -indices for the different levels. Fig. 3 illustrates this data structure in 2D, and Algorithm 1 outlines how sparse rows of particles are accessed. The particles are indexed linearly in the order  $l \rightarrow z \rightarrow x \rightarrow y$ , with sparse compression along  $y$ . Coordinate indices in the sparse dimension are stored in the vector  $y\_idx$ , while the vector  $xz\_end$  encodes the last particle index in each sparse row, with one entry for each combination of  $(l, z, x)$ . An additional vector  $level\_offset$  stores the starting point of each level in the  $xz\_end$  vector. The metadata required to correctly access these vectors are the

**Algorithm 2:** Iteration Over Particles and Coordinates

---

```

Data: LinearAccess apr, intensity[]
for ( $l = \text{apr.l}_{\min}; l \leq \text{apr.l}_{\max}; l++$ ) do
  for ( $z = 0; z < \text{apr.z}_{\dim}[l]; z++$ ) do
    for ( $x = 0; x < \text{apr.x}_{\dim}[l]; x++$ ) do
       $r_{\text{begin}}, r_{\text{end}} \leftarrow \text{apr.get\_row}(l, z, x)$ 
      for ( $i = r_{\text{begin}}; i < r_{\text{end}}; i++$ ) do
         $y \leftarrow \text{apr.y\_idx}[i]$ 
         $p \leftarrow \text{intensity}[i]$ 
        /* do something */

```

---

minimum and maximum resolution levels, as well as the grid dimensions at each level.

This linear access data structure complements, rather than replaces, the random access data structure from Ref. [16]. Both of the data structures are available in LibAPR [23] (see Code Availability), and the similarity between them allows for efficient conversion from one to the other depending on the anticipated access pattern.

### B. Sequential Data Access and Iteration

Using the linear access data structure from Algorithm 1, querying particles at random locations is inefficient, but iterating sequentially over particles in the same sparse row is efficient, since both  $y$ -indices and particle properties are read from contiguous memory, promoting cache efficiency. Algorithm 2 illustrates how one would iterate over all particles of a given APR of class `LinearAccess`, accessing both spatial coordinates and intensity values.

Most algorithms require not only iteration over individual particles, but over pairs or neighborhoods of particles. In these cases, the optimal strategy depends on the size and structure of the neighborhood, as well as the type of processing hardware, e.g. CPU or GPU. In the following section we describe the main algorithmic contribution of this work, which includes two iteration strategies for accessing pairs and groups of particles.

### C. Neighborhood Data Access

Some of the most fundamental operations in image processing, such as spatial filtering and resampling operations, compute values for each pixel in the output image by accumulating values over fixed-size neighborhoods of the input image. These operations can be extended to the APR by instead considering neighborhoods of particles. However, due to the adaptive resolution, the structure of these neighborhoods varies across the domain. This complicates both the definition and implementation of such operations.

Rather than adapting the operations to anisotropic or non-uniform structures, the reconstruction condition in Eq. 1 enables the interpolation of information across resolution levels. This can be used to interpolate neighborhoods of particles to small local patches of uniform resolution, as shown in Fig. 4. Indeed, thanks to the reconstruction condition, coarse particles can be interpolated to the finest resolution




Fig. 4. Particle values can be interpolated between resolution levels to create locally isotropic neighborhoods with guaranteed error bounds.

level with a guaranteed point-wise bound on the interpolation error. This also implies a bounded error at all intermediate resolutions, compared to a similarly downsampled signal, i.e., the corresponding node in the image pyramid (or the full tree in Fig. 2A). Thus, since the APR particles and interior nodes coincide exactly with the corresponding tree nodes and the missing “pruned” nodes can be approximated, any image region can be reconstructed at any resolution with guaranteed error bounds.

While conceptually simple, this level interpolation constitutes an algorithmic challenge as it requires iterations over particles in multiple rows, across multiple resolution levels. For example, the  $3 \times 3$  patch in Fig. 4 intersects with 2 sparse rows at the coarser resolution level, 3 sparse rows at the target (medium) resolution, and 6 sparse rows at the finer resolution. Thus, reconstructing such a patch potentially requires iterations, or search, over 11 sparse rows across three levels. In 3D, a corresponding  $3 \times 3 \times 3$  patch uses values from 49 sparse rows. The number of rows required for local patch interpolation grows exponentially with dimension as additional finer resolution levels are included. In other words, reconstructing coarse isotropic patches on the fly, directly from the particles, is not feasible in practice. We address this by introducing an auxiliary data structure, which we call the *APR tree*.

### D. The APR Tree

We define the APR tree as the set of all interior nodes in the tree structure, i.e., all yellow nodes in Fig. 2B. Equivalently, this can be thought of as the set of elements in the image pyramid that can be obtained by downsampling the APR particles. We store and access these nodes using a separate instance of the `LinearAccess` APR data structure from Algorithm 1. The values of all APR tree particles (interior nodes) can be computed from the APR by recursive reduction over descendant particles. This requires iterating over all APR particles (leaf nodes), while simultaneously accessing the ascendant nodes. Algorithm 3 shows how this can be done, using what we call *synchronized iteration*. For each particle in the leading sparse row, the parent iterator is incremented until the  $y$ -indices align.<sup>3</sup>

For separable reduction operators, such as maximum or average reductions, synchronized iteration can be used to

<sup>3</sup>Note that Algorithm 3 depends on the assumption that the parent particle always exists, which for APRs is true by construction. The iteration strategy can be extended to pairs or groups of sparse rows with other spatial relations, but this likely requires additional conditions to be checked.

---

**Algorithm 3:** Synchronized Iteration Over a Sparse Row of Particles and Their Parent Nodes in the APR Tree

---

**Data:** LinearAccess apr, LinearAccess tree

**Input:**  $l, z, x$

```

/* APR particle row */
r_begin, r_end = apr.get_row(l, z, x)
/* Parent row in the APR tree */
p_begin, p_end = tree.get_row(l - 1, [z/2], [x/2])
j = p_begin // parent index
/* Iterate over APR particles */
for (i = r_begin; i < r_end; i++) do
    /* Increment j until y-indices align */
    while tree.y_idx[j] < [apr.y_idx[i]/2]
        do
            j++
    // tree node j is the parent of apr node i

```

---

compute the values of all particles in an APR tree in the following two steps:

- 1) iterate over all APR particles, reducing their values onto the parent nodes;
- 2) iterate over all APR tree particles, level by level from the finest resolution, to fill the tree.

This algorithm can be parallelized by distributing the sparse rows in each step across different threads or processes. However, care must be taken to avoid competing writes to the same address. We do this by distributing blocks of  $2 \times 2$  (in 3D) rows, such that each block corresponds to a single parent row.

The memory cost of storing the tree depends on the APR structure and hence the image contents. For a uniform APR in 3D, where all  $N_p$  particles are at level  $l_{\max}$ , the tree contains

$$N_t = \sum_{k=1}^{l_{\max}} \frac{N_p}{8^k} = \frac{N_p}{7} (1 - 8^{-l_{\max}}) \quad (7)$$

particles. This converges quickly to  $N_p/7$  with increasing  $l_{\max}$ . In general, the exact number of tree particles can be computed recursively by considering the number of APR and tree particles at each level. However, for sufficiently large volumes represented by sufficiently many particles,  $N_p/7$  remains a close approximation. In that case, the memory cost of the LinearAccess data structure is also approximately linear in the number of particles, such that the overall memory overhead is about  $1/7 \approx 14.3\%$ .

#### E. Isotropic Patch Reconstruction

By precomputing the values of all APR tree particles, locally isotropic patches can be reconstructed at any level without iterating over the sparse rows at finer resolutions. This reduces the complexity of this operation. In the following, we assume that all particle values are reduced by average downsampling, and that piecewise constant APR reconstruction is used.<sup>4</sup>

<sup>4</sup>For general linear reduction and reconstruction methods, the interpolation from level  $l_0$  to  $l_1$ , where  $l_0 < l_1 < l_{\max}$ , can be modified such that it corresponds to interpolation from  $l_0$  to  $l_{\max}$  followed by downsampling to  $l_1$ .

We denote the APR particles by  $\mathcal{P} = \{(\mathbf{x}_p, f(\mathbf{x}_p))\}_{p=1}^{N_p}$  and the APR tree by  $\mathcal{T} = \{(\mathbf{x}_t, f(\mathbf{x}_t))\}_{t=1}^{N_t}$ . Suppose we wish to reconstruct the value  $\hat{f}(\mathbf{x}_l)$  of a grid cell  $\mathbf{x}_l$  at resolution level  $l \leq l_{\max}$ . There are two possible cases:

- 1) There exists an APR particle  $(\mathbf{x}_p, f(\mathbf{x}_p))$  at resolution level  $l_p \leq l$  whose particle cell coincides with or contains  $\mathbf{x}_l$ .
- 2) The grid cell  $\mathbf{x}_l$  contains multiple APR particles at finer resolutions  $l_p > l$ . In this case, there exists a tree particle  $(\mathbf{x}_t, f(\mathbf{x}_t))$  at level  $l_t = l$  whose cell coincides with  $\mathbf{x}_l$ .

Since we use piecewise constant reconstruction, we obtain  $\hat{f}(\mathbf{x}_l) = f(\mathbf{x}_p)$  in the first case and  $\hat{f}(\mathbf{x}_l) = f(\mathbf{x}_t)$  in the second case. Thus, reconstructing an isotropic patch at level  $l$  requires iteration over APR particles within the patch at levels  $l' \leq l$  and, if  $l < l_{\max}$ , tree particles at level  $l$ . The LinearAccess data structure from Algorithm 1 allows patches of arbitrary shape and extent in the densely encoded dimensions. However, for performance reasons, it is best to let the patch span the entire sparse y-dimension in order to avoid linear searches due to the sequential access.

#### IV. APR PROCESSING

The ability to locally reconstruct isotropic neighborhoods allows for a wide range of image processing algorithms to be implemented natively for the APR. Arguably the most important low-level vision task is discrete convolution. Not only does this enable spatial filters, but it is also an essential component of high-level vision algorithms including convolutional neural networks. We therefore start by describing APR-native discrete convolution.

##### A. APR Filtering by Discrete Convolution

For a pixel image  $\mathbf{u} \in \mathbb{R}^s$  of size  $\mathbf{s} \in \mathbb{N}^d$  in dimension  $d$ , we refer to spatial filtering as the process of applying a dense discrete convolution operation with a stencil  $\mathbf{w} \in \mathbb{R}^f$ ,  $\mathbf{f} \in \mathbb{N}^d$ , of the same dimensionality but not necessarily the same size<sup>5</sup> as  $\mathbf{u}$ . The output  $\mathbf{o} = \mathbf{u} * \mathbf{w}$  produced by this operation is again an image, where each pixel is a linear combination of a neighborhood of pixels in  $\mathbf{u}$ , weighted by the values in  $\mathbf{w}$ . Now consider an APR  $\mathcal{P}$  representing  $\mathbf{u}$ . It is certainly possible to approximate  $\mathbf{o}$  by reconstructing the pixel image  $\hat{\mathbf{u}}$  from  $\mathcal{P}$  according to the reconstruction condition and computing  $\hat{\mathbf{o}} = \hat{\mathbf{u}} * \mathbf{w}$ . The output image  $\hat{\mathbf{o}}$  can then again be converted to a new APR  $\hat{\mathcal{P}}$  using the pulling scheme. When used in this way, the APR acts as a lossy compression technique, but provides no computational speedup over directly processing  $\mathbf{u}$ .

We therefore aim to define a native discrete convolution operation for the APR, which is consistent with convolving a fully reconstructed image, but only requires computational operations over the particles. Thus, rather than interpolating coarse particles to the finest resolution, we extend the convolution to operate directly on the coarser resolution levels where appropriate. This can be done by convolving each resolution

<sup>5</sup>Typically, the size of the stencil is on the order of a few pixels in each dimension.




Fig. 5. Illustration of APR convolution in one dimension. Filter stencils (dashed boxes) are applied centered on each APR particle (green dots), with neighboring particle values interpolated to the appropriate resolution, to compute the output value at the corresponding center particle. Interpolation from coarse to fine resolution is done on the fly, while values interpolated from fine to coarse resolution are precomputed and accessed via the auxiliary APR tree structure (yellow dots).

level separately, similar to how convolutions can be applied to each level of an image pyramid.

An APR corresponds to a sparse disjoint sampling of an image pyramid constructed from the original pixel image. Moreover, isotropic patch reconstruction (cf. Section III-E) allows us to approximate any region in this pyramid with a theoretically bounded reconstruction error according to Eq. 1. This affords some freedom in designing the convolution operation. Here, we define a convolution as an operator that acts only on the particle values  $f(\mathbf{x}_p)$ , keeping their positions  $\mathbf{x}_p$  fixed. More precisely, the filter is applied centered on each particle to compute the output value at that location. Neighboring particles are interpolated to an isotropic patch at the resolution of the center particle. The size of the patch must be at least the size of the stencil  $\mathbf{w}$ . This is illustrated in Fig. 5. In practice, it is typically advantageous to reconstruct larger patches such that the stencil can be applied to multiple locations in each patch.

Since reconstructed patches approximate the corresponding part of the pixel image pyramid, one can derive a bound on the error in the output values in terms of the filter  $\mathbf{w}$ , the error bound  $E$ , and the local error scale  $\sigma$  restricted to the resolution level of the patch:  $|\mathbf{o} - \hat{\mathbf{o}}| = |\mathbf{w} * (\mathbf{u} - \hat{\mathbf{u}})| \leq \|\mathbf{w}\| E \sigma$ . However, the sampling locations computed from the input signal may not be optimal to represent the convolved signal. Hence, while the APR convolution is consistent with the corresponding pixel convolution in the sense that output values can be interpreted as the corresponding values in the convolved image pyramid, the fixed sampling may result in additional interpolation or aliasing errors, which depend on the image.

Ideally, one would allow the convolution to adapt the sampling on the fly. This would require an error estimator that depends on the filtered signal. While this could trivially be achieved by evaluating the convolution on the full pixel grid, operating natively on APR would require: (i) finding the particles in the reconstruction neighborhood of each pixel location in the target particle cell; (ii) checking the reconstruction error for all combinations of these values; and (iii) ensuring that the resulting structure adheres to APR conventions (e.g., the level of neighboring particles cannot differ by more than one). This would incur significant additional computa-

tional cost. The main difficulty is point (i), as the particle locations in the neighborhood are subject to change, preventing local decisions. In the present work, we therefore consider the sampling locations fixed. More concretely, the result can be viewed as a bounded-error approximation of  $\mathbf{o}$  if the filter  $\mathbf{w}$  is such that it does not increase the required resolution at any particle location. This is a limiting assumption, which clearly does not hold for arbitrary filters. In particular, it excludes filters that result in a significant net translation of the signal. However, it intuitively holds, at least approximately, for important image filters such as gradient and sharpening, as well as many (near-symmetric) smoothing filters.

At the finest resolution level, the APR convolution operation is equivalent to convolving the reconstructed image  $\hat{\mathbf{u}}$ . With decreasing resolution the spatial extent of the filter grows, and the convolution is applied to (approximations of) downsampled versions of the image. Thus, applying the same stencil  $\mathbf{w}$  at coarser resolution levels is, in general, not consistent with the convolution at the finest resolution. We address this by modifying the stencil depending on the resolution level.

### B. Pixel-Consistent APR Convolutions by Stencil Restriction

In order to render APR convolution consistent with a convolution of the fully reconstructed image  $\hat{\mathbf{u}}$ , we use *operator restriction* as commonly encountered in the algebraic multigrid literature [24], [25]. Denote by  $R_l$  the restriction (downsampling) operator used to determine particle values at a coarse resolution level  $l < l_{\max}$ , and let  $P_l$  be the prolongation (upsampling) operator to reconstruct pixel values from particles at level  $l$ . Let  $\hat{\mathbf{u}}_l$  be the image reconstructed at level  $l$  and  $\hat{\mathbf{u}}$  the finest reconstruction at pixel resolution. We then have:

$$\hat{\mathbf{u}}_l = R_l \hat{\mathbf{u}}, \quad \hat{\mathbf{u}} \approx P_l \hat{\mathbf{u}}_l. \quad (8)$$

That is,  $\hat{\mathbf{u}}_l$  can be computed from  $\hat{\mathbf{u}}$  by downsampling via  $R_l$ , and  $\hat{\mathbf{u}}$  can be approximated by upsampling via  $P_l$ . Now consider the convolution of  $\hat{\mathbf{u}}$  with a stencil  $\mathbf{w}$ . This is a linear operation, which can be written as a matrix multiplication  $\mathbf{w} * \hat{\mathbf{u}} := K_w \hat{\mathbf{u}}$ . Using the interpolation operators  $R_l$  and  $P_l$ , we can now define a convolution operator  $K_{w_l}$  for the coarse image  $\hat{\mathbf{u}}_l$  on resolution level  $l$  as

$$K_{w_l} \hat{\mathbf{u}}_l := R_l K_w P_l \hat{\mathbf{u}}_l. \quad (9)$$

In this way,  $K_{w_l}$  is equivalent to interpolating  $\hat{\mathbf{u}}_l$  to the fine resolution, applying the convolution with  $\mathbf{w}$  there, and downsampling the result back to the coarse resolution. Similar to  $K_w$ , the coarse scale Toeplitz matrix  $K_{w_l}$  corresponds to a stencil  $\mathbf{w}_l$ , whose structure depends on the fine resolution stencil  $\mathbf{w}$ , as well as the restriction and prolongation operators. For stencils  $\mathbf{w}$  that are spatially invariant, the stencil  $\mathbf{w}_l$  can be computed without generating the matrix  $K_{w_l}$ , by evaluating only the  $|\mathbf{w}_l|$  unique non-zero elements.

The interpolation operators used in the standard APR formulation, i.e., block-wise average downsampling and piecewise constant upsampling, amount to particularly simple forms of  $R_l$  and  $P_l$ . This combination additionally has the property that  $R_l P_l = I$ , where  $I$  is the identity matrix. In this case,



the restricted stencil  $\mathbf{w}_l$  is computed directly by averaging the contributions of all elements of the fine stencil  $\mathbf{w}$  applied at all possible positions within the coarse grid cell.

This provides for internally consistent APR convolution by restricting the filter stencil to each resolution level, of which there are at most  $\lceil \log_2(|\Omega|) \rceil$ , where  $|\Omega|$  is the maximum image side length in pixels. Figure 6 shows an example of this approach for a Gaussian smoothing filter (middle row) applied directly on the APR levels after restriction (right panel), compared with applying the same filter on the raw pixels (middle panel) and on a low-noise ground-truth (GT) version of the pixel image (left panel). The respective input images are shown in the top row of the figure. The results of the smoothing filter computed on raw pixels and on the APR are visually indistinguishable and have similar peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) with respect to the blurred GT image. This confirms that the restricted APR filters provide pixel-consistent results.

### C. Resolution-Adaptive Filters

The APR convolution operation is inherently spatially adaptive, as guided by the structure of the APR. In the previous section, we adjusted the convolution stencils on each APR level such that the APR convolution is consistent with the full pixel operation. However, this may not always be the goal. Indeed, the spatial adaptivity of the APR can be exploited to provide scale-adaptive image filters. This can improve filter performance, e.g., in image gradient estimation. Since the smoothing length of the gradient estimator depends on the spacing between neighboring discretization points, it naturally varies with the APR resolution level. Thus, APR convolution with a level-dependent stencil  $\mathbf{w}_l$  of the form

$$\mathbf{w}_l = 2^{-(l_{\max}-l)} \mathbf{w}, \quad (10)$$

where  $\mathbf{w}$  is the pixel-level stencil, naturally adapts gradient estimation to varying length scales in the image contents. This is akin to the adaptive filters described in [16]. The convolution defined here, however, is not limited to direct face-connected neighbors, but can handle filter stencils of arbitrary size. The benefits are illustrated in Fig. 6 (bottom row) for gradients estimated using Sobel filters, which combine central finite differences with a smoothing filter orthogonal to the gradient direction. When used on the APR with level-dependent rescaling according to Eq. 10, the gradients along the edges of the objects are captured, while being less sensitive to imaging noise in low-gradient regions. This is reflected in the higher PSNR and SSIM values of the result computed on the APR compared to the result computed on pixels.

## V. PARALLEL IMPLEMENTATION

The APR convolution operation described so far essentially consists of three algorithmic components:

- 1) computing the values of (interior) APR tree particles,
- 2) reconstructing isotropic patches, and
- 3) accumulating patch values to form the convolution output.

Filling the APR tree (step 1) is algorithmically independent of the convolution operation, while steps 2 and 3 are performed alternately for each location in the image domain. In order to take advantage of modern computing hardware, we have implemented these algorithmic components to use thread parallelism on multi-core CPUs and on GPUs. We describe the design principles used to achieve this and discuss their advantages and limitations.

### A. CPU Parallelization Using OpenMP

Modern CPUs typically comprise 6 to 48 cores, each optimized to execute a pipelined sequence of instructions called a *thread*. Parallelization on multi-core CPUs thus amounts to dividing a program into a small number of threads that can be distributed across the cores and executed concurrently. We provide a thread-parallel CPU implementation of the present APR convolution algorithm based on the Open Multi-Processing (OpenMP) [27] API, which provides a set of high-level directives and work-sharing constructs for SPMD (Single Program Multiple Data) parallelism, where the same set of instructions is executed concurrently at multiple data points.

In order to distribute the data points, in our case APR particles, across threads, we follow a domain-decomposition approach as is classic for pixel images. Execution efficiency demands that the threads iterate over the sparse  $y$ -dimension, which is accessed sequentially according to Algorithm 2. Therefore, we partition the APR by assigning to each CPU core a subset of the sparse rows, rather than a subset of the particles. However, the number of particles in each sparse row can vary greatly. This may lead to large differences in the amount of work done by each core, causing slowdown due to load imbalance. We address this by runtime dynamic load balancing, dynamically re-distributing sparse rows between threads at runtime.

Consider the convolution with a stencil  $\mathbf{w}$  of size  $k^3$  at some resolution level  $l$  in the APR of an image of size  $N_z \times N_x \times N_y$  pixels. The data are distributed such that each thread computes the output for the APR particles in a given  $z$ -slice. Each thread allocates a buffer of size  $k \times k \times (N_y + k - 1)$  for the isotropic patches, where the  $k - 1$  additional points in the  $y$ -dimension allow for padding to handle boundary conditions at the edges of the image. For each  $z$ -slice, the buffer is iterated over along the  $x$ -dimension with the center of the buffer at the given  $z$ -index. At each location, the thread iterates over the necessary<sup>6</sup> sparse rows to fill the  $k$  new rows at the beginning of the buffer (cf. Section III-E). The output values are then aggregated for each APR particle in the sparse row corresponding to the center location of the buffer.

### B. GPU Parallelization Using CUDA

In addition to the multi-threaded CPU implementation, we also provide an optimized implementation of the present

<sup>6</sup>Isotropic patch reconstruction does not have to include the particles from all coarser levels, but enough to ensure that neighborhoods of size  $k^3$  around each output particle are filled. Since the resolution levels of neighboring particles cannot differ by more than one [16], the number of coarser levels that must be considered scales logarithmically with the stencil radius  $(k - 1)/2$ .






Fig. 6. Illustration of APR filtering for smoothing and gradient estimation. The images are taken from the FMD dataset [26]. Top row: ground truth (GT) image obtained by averaging 50 acquisitions, raw image (a single acquisition), and the APR of the raw image. Middle row: Gaussian smoothing with standard deviation 2 pixels applied to each image from the top row, using restricted stencils for coarser APR particles. Bottom row: gradient magnitudes computed from the respective image of the top row. The ground truth gradients were computed using central finite differences, whereas the gradients on noisy raw image were computed using Sobel filters. For the APR, the filters were rescaled according to the particle distance at each resolution level (see Eq. 10). All peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) values are given with respect to the GT image in the same row.

APR convolution algorithm for GPUs. Our implementation uses the Compute Unified Device Architecture (CUDA) API for GPU programming and is hence limited to Nvidia GPUs. Recent Nvidia GPUs boast thousands of cores organized into arrays of *Streaming Multiprocessors* (SMs), each designed to concurrently execute hundreds of threads. In order to parallelize a program using CUDA, it is divided into *thread blocks* (TBs). Each TB executes a program, called a *CUDA kernel*, over a given block of data. The threads within a TB are executed on the same SM in groups of 32 threads, called *warps*. Threads within the same TB communicate via shared memory and barrier synchronization, with recent versions of CUDA additionally supporting warp-level communication and synchronization.

CUDA programs that make good use of the hardware resources of a GPU need to adhere to multiple design principles: First, all threads within a warp should simultaneously perform the same operations. This maximizes efficiency because the threads in a warp execute instructions in lockstep following to the SIMT (Single Instruction Multiple Threads) model of parallelism. Second, diverging control flow paths within a warp are to be avoided. Otherwise, the alternative control flow paths are executed sequentially, with masking employed to enable and disable the appropriate threads. Third, access to the global memory has to be minimized and done in stride, as those transactions carry a significant overhead. This is mainly achieved by *latency hiding*, where execution switches to another warp while one warp waits for the memory. Latency hiding requires that enough scheduled threads are

available, providing a fourth design requirement. Fifth, memory latency can be further reduced by: (1) loading frequently accessed data into shared memory or registers and (2) ensuring that the necessary global memory accesses are coalesced, such that each transaction serves as many threads as possible at once.

These programming requirements are not easily reconciled with a dynamically adaptive tree data structure like the APR. Considering the APR data structures proposed in Algorithm 1, particle values and y-indices are contiguous in memory along the sparse y-dimension. In order to promote coalesced memory access to these vectors, adjacent threads in a warp should therefore access adjacent particles within the same sparse row. We achieve this by distributing the sparse rows across warps. The size of each thread block is limited to 1024 threads on most GPU devices. Since a warp consists of 32 threads, one warp must handle multiple sparse rows if an operation requires accessing more than 32 different positions in  $z$  and  $x$ . This may lead to reduced concurrency due to increased register usage. By instead assigning a half-warp (16 threads) to each sparse row, this limitation can be relaxed to allow up to 64 different positions. Therefore, our design imposes an upper limit to the size of the stencil neighborhood that can be accessed during convolution. Assuming a cube-shaped stencil in 3D, with odd side length, the size is limited to  $7^3$  pixels.

Using this design, we provide optimized implementations of 3D APR convolutions in CUDA for filter stencils of size  $3^3$  and  $5^3$  pixels. For the sake of example, we describe the  $3^3$  algorithm, assuming that the tree particle values  $f(\mathbf{x}_i)$




Fig. 7. Illustration of our CUDA kernel operations for APR convolution at an arbitrary APR level  $l$  in 2D. Each warp corresponds to a fixed position  $(z, x)$  at level  $l$ . The threads load particles from up to three sparse rows at the corresponding position: the APR at levels  $l-1$  (top-left panel) and  $l$  (middle-left), and the APR tree at level  $l$  (bottom-left). Particle values are written to a uniform shared-memory buffer at level  $l$  to reconstruct local isotropic neighborhoods around target particle locations (top-right, loaded particle origins identified by color). Finally, the convolution stencil  $w_l$  (middle-right) is applied to interior positions in the buffer, computing the output at the level  $l$  APR particle locations (bottom-right).

of  $\mathcal{T}$  have previously been computed. For this stencil size, we choose a thread block size of  $32 \times 4 \times 4$ , resulting in 16 warps at different positions in a  $4 \times 4$  patch in the  $z-x$  plane. Isotropic patches are reconstructed in a shared memory buffer of size  $4 \times 4 \times 32$ . Reconstructing neighborhoods of size  $3^3$  around APR particles at level  $l$  requires values from the APR at levels  $l-1$  (if  $l > l_{\min}$ ) and  $l$ , as well as the APR tree at level  $l$  (if  $l < l_{\max}$ ). Thus, each thread accesses particles from up to three sparse rows, one from each of these structures, corresponding to the fixed  $(z, x)$  position of the warp. The shared memory buffer is iterated across the  $y$ -dimension in a synchronized manner. For each position of the buffer, the threads:

- 1) Update their particles by comparing the  $y$ -indices to the current location of the buffer. If a particle is behind the buffer (i.e. has a  $y$ -index smaller than the beginning of the buffer), the thread loads the next particle in the corresponding row. This is repeated until no particles are behind the buffer any more.
- 2) Reconstruct local isotropic patches by writing particle values within the range of the patch to the shared memory buffer.
- 3) Aggregate the convolution output for the level  $l$  APR particles in the  $2 \times 2$  center rows of the buffer.

Block-level barrier synchronization is required before steps 2 and 3 to ensure correctness of the output. Figure 7 illustrates this algorithm in 2D. The implementation for stencils of size  $5^3$  is analogous, with the main difference being the size of the thread block, which we then choose to be  $16 \times 8 \times 8$ .

We additionally implement high-level optimization strategies to reduce unnecessary work for sparse images. In this case, the sparse rows at fine resolution levels tend to be empty or contain only few particles. Therefore, thread blocks may needlessly reconstruct patches where there are no target particles. We avoid this by using a pre-processing CUDA kernel to detect the positions in the APR where the sparse rows contain at least one particle. The convolution stencils are then launched only at those locations. In addition, the thread blocks use a reduction operation to determine the minimum

and maximum  $y$ -indices of the target particles. Positions of the uniform buffer outside of this range are skipped.

## VI. RESULTS

We present benchmark results of computational performance and memory usage of our APR convolution implementations using both OpenMP on CPUs and CUDA on GPUs. We also present an example of image deconvolution using the Richardson-Lucy algorithm re-implemented for APRs and compare the results with those obtained by the classic algorithm on pixels.

Since the APR adapts to the contents of an image, the performance of APR processing algorithms depends on image contents through the number of particles required by the APR to reach the reconstruction error bound. We quantify this using the *computational ratio* (CR), defined as

$$\text{CR} = \frac{\text{Number of pixels in the original image}}{\text{Number of particles in the APR}}. \quad (11)$$

To put this metric into perspective, typical CR values for real-world fluorescence microscopy datasets range from 3.6 to 372.2 with an average of 51.1 and a median of 22.7 [16]. The CR value one can obtain for any given image depends mainly on the sparsity of the image content. The APR will represent regions where the signal gradient is significant (relative to the local error scale  $\sigma$ ) at high resolution. Thus, if the gradients are distributed densely across most of the image, high CR values cannot be obtained without significant loss of information.

In addition to the CR, also other factors may directly or indirectly affect memory usage and computational performance. For example, the size of the APR data structure explicitly depends on the original image size in the  $x$  and  $z$  dimensions. Moreover, performance may vary between different APRs, even if both the image dimensions and the CR are equal, since the particle layout can affect, e.g., cache hit rates and conditional branches.

### A. APR Convolution Benchmarks

Here we present results on the memory usage and computational performance of our APR convolution implementations. The benchmarks are performed on APRs computed from a set of 10 synthetic images, containing spherical objects at varying densities, yielding a range of CR values from 1.04 to 1019.8. The files, as well as the code used to produce all data presented here, are available in LibAPR.

1) *Memory Usage:* In order to convolve an image of  $N$  pixels, the simplest approach is to allocate enough memory to hold the entire image plus an output buffer of the same size. If the input and output data types require  $s_i$  and  $s_o$  bytes of storage, respectively, this requires  $N(s_i + s_o)$  bytes of memory. Similarly, convolving an APR requires two buffers for the input and output particle intensities, but in addition also the access data structures of the APR and the APR tree, as well as the interior particle intensity values.

We benchmark this using synthetic images of size  $1024^3$  pixels and of different CR. The benchmark images are described and examples shown in Supplementary Material I-A. Figure 8 shows a breakdown of the memory




Fig. 8. Memory required to perform one convolution operation on a cubic image of  $1024^3$  pixels and the corresponding APRs at different computational ratios (CR). Pixel and particle intensities are stored using 32-bit data types. We neglect the memory required to store the weights of the filter stencil. The APR memory usage additionally includes the APR access data structure, as well as the access data structure and values of the APR tree (interior nodes).

requirements for APR convolution when all pixel and particle intensity values are stored as 32-bit data types. The straightforward convolution on pixels requires 8.59 GB of memory,<sup>7</sup> exceeding the available VRAM of most modern GPUs. Tiling strategies must then be implemented to enable processing, but this is only efficient if image tiles can be transferred to and from the GPU device concurrently with computation. This adds complexity to the implementation and most likely results in a performance penalty. Computing the convolution natively on an APR with a CR of 20.8 (which is smaller than the median of real-world microscopy datasets [16]) requires 0.58 GB of memory: 14.8 times less than the equivalent pixel operation. This makes straightforward processing of the image possible on almost any currently available GPU.

The memory required for APR convolution decreases monotonically with increasing CR, as shown in Fig. 8. For  $CR \approx 1$ , we observe that the memory overhead from the APR data structures is about 30%. However, this overhead is quickly amortized for higher CRs, since the size of the `xz_end` vector in the `LinearAccess` data structure is constant (see Algorithm 1). APR-native convolution at  $CR \approx 1020$ , a value not uncommon in fluorescence microscopy, requires 25.5 MB of memory, or 337 times less than the classic pixel-based implementation.

For GPU implementations, the benefit of reducing the memory requirement is two-fold: First, it enables larger image regions to be kept in memory and processed without tiling strategies. Second, the total amount of data that needs to be transferred to and from the GPU is reduced proportional to the CR, alleviating the performance bottleneck from the host-device transfer bandwidth.

2) *Computational Performance:* We present benchmark results for APR convolution using filter stencils of  $3^3$  and  $5^3$  pixels, respectively. All benchmarks are performed on an Alienware m15 laptop equipped with a GeForce RTX 2080 Max-Q GPU and an Intel Core i9-9980HK CPU running Ubuntu 18.04 and CUDA Toolkit 11.0 RC.

The APR convolution implementation in CUDA consists of three steps: 1) finding the locations of non-empty sparse rows,<sup>8</sup> 2) filling the APR tree, and 3) performing the convolution operation. All of these steps are implemented in CUDA and performed in sequence. In the benchmark results below, all three steps are included in the reported times. However, in practical situations requiring repeated convolutions, step 1 can be reused, and for convolution of the same input with multiple filters, the tree data can also be reused, leading to runtimes smaller than those reported here.

Rather than reporting absolute wall-clock times, we put the results on a more intuitive scale by computing the *effective throughput*, which we define as

$$\text{Effective throughput} = \frac{\text{size of pixel image in Bytes}}{\text{total processing time}}. \quad (12)$$

For a pixel algorithm, this is the classic data throughput, i.e., the number of Bytes processed per second. For an APR, the effective throughput states the throughput a pixel algorithm would need to have in order to achieve the same processing time. Figure 9 shows the scaling of the effective throughput for the same synthetic images of different CRs as already used in Fig. 8, subsampled to size  $512^3$ . All data are represented as 32-bit floating point numbers. Hence, the numerator in Eq. 12 is 537 MB. The GPU timings include all computational steps, but exclude the times for host-device data transfers. The runtimes for the corresponding pixel convolutions are shown as horizontal lines, as they do not depend on the CR. The CUDA pixel benchmarks were performed using ArrayFire v3.8.0 [28] in C++, while the OpenMP benchmarks use our own implementation.

The effective throughput of the ArrayFire CUDA implementation is 26.6 GB/s for the  $3^3$  filter and 10.7 GB/s for the larger  $5^5$  filter. Our CUDA APR convolutions break even with the performance of ArrayFire at CR values  $\approx 2$ . For higher computational intensity of the operation (i.e., for larger stencils), APR convolution breaks even earlier. At a CR of 20.8, the  $3^3$  and  $5^5$  APR convolution speeds correspond to pixel throughputs of 102.2 and 64.5 GB/s, respectively. That is 3.8 and 6.0 times faster than the ArrayFire implementation. For the benchmark image with  $CR=124$ , which is well within the range typical of real-world microscopy images [16], the effective throughput of APR convolution increases to 455.7 and 299.1 GB/s, respectively. Compared to ArrayFire, this corresponds to speedup factors of 17 and 28. Considering the theoretical peak performance of 6.447 TFLOPS (trillion floating point operations per second) of the benchmark GPU, the theoretical performance limit for the pixel convolution is 486 GB/s for the  $3^3$  filter and 103 GB/s for  $5^5$ . The effective

<sup>7</sup>Throughout this manuscript, we use decimal SI prefixes for Byte multiples, so for example 1 GB =  $10^9$  Byte.

<sup>8</sup>This is not strictly necessary, but we have found that it yields a speedup of up to a factor of 5 for high CR values, while the additional overhead is negligible for low CR values.



TABLE I

PERFORMANCE AND QUALITY METRICS OF THE DECONVOLUTION RESULTS FROM FIG. 10. THE INPUT IMAGE HAS CR=5.3. THE PSF COVERAGE IS THE FRACTION OF THE TOTAL INTENSITY OF THE PSF INCLUDED IN THE TRUNCATED ARRAY. REPORTED RUNTIMES FOR PIXEL IMAGES USE DECONVOLUTIONLAB2 [31] (SINGLE-THREADED), WHEREAS APR RUNTIMES USE MULTI-THREADING ON 8 CPU CORES. THE PSNR AND SSIM VALUES ARE COMPUTED WITH RESPECT TO THE GROUND-TRUTH IMAGE AFTER NORMALIZATION TO ACCOUNT FOR INTENSITY RANGE DIFFERENCES DUE TO TRUNCATION AND NORMALIZATION OF THE PSF. THE FULL PSF APR RESULT (MARKED BY \*) IS OBTAINED BY DECONVOLVING THE FULL-RESOLUTION RECONSTRUCTED VOLUME AND RESAMPLING THE RESULT ONTO THE PARTICLES

Data	PSF	PSF coverage [%]	Time per iteration [s]	Memory usage [MB]	PSNR [dB]	SSIM
Input	-	-	-	-	26.51	0.504
Pixels	Full	100	5.11	2038	29.47	0.888
Pixels	$21^3$	12.0	5.11	2038	27.45	0.458
APR	$11^3$	6.3	0.69	194	27.11	0.459
APR	$21^3$	12.0	3.49	206	27.88	0.469
APR	$41^3$	23.5	33.1	250	28.92	0.545
APR*	Full	100	5.11	2038	29.48	0.886

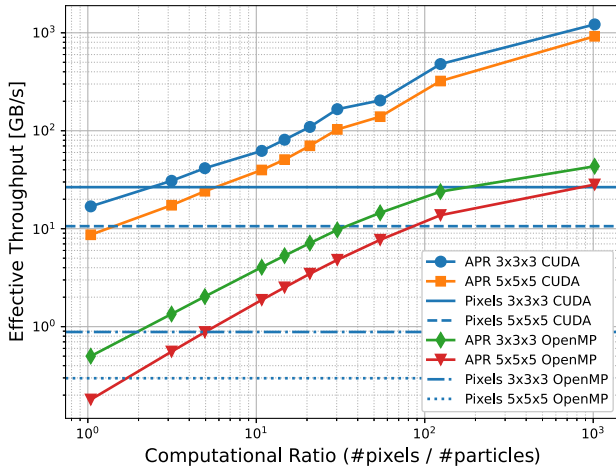


Fig. 9. Computational performance of APR convolutions using GPU (CUDA) and CPU (OpenMP) parallelization for 10 synthetic benchmark datasets of different computational ratios (CR). The effective throughput is calculated as the size of the original image in GB divided by the processing time in seconds. GPU processing times include all computations, but exclude data transfers. The horizontal lines show measurements for pixel convolutions using CUDA ArrayFire as well as OpenMP. We show results for stencils of size  $3^3$  and  $5^3$  voxels.

throughput of the  $5^3$  APR convolution exceeds this limit for CR values larger than 30.

The performance of the APR convolution on the multi-core CPU using OpenMP also scales with CR. Similar to the CUDA results, the APR performance breaks even with the corresponding pixel algorithm at CR values  $\approx 2$ . For large CR values  $> 100 \dots 500$ , APR performance on the CPU even exceeds pixel performance on the GPU. If data transfer times to and from the GPU are taken into account, the CPU breaks even at lower CR values. Comparing the performance of the APR convolution on the GPU and CPU, we observe that the GPU implementation (without data transfer times) is 15 to 45 times faster across benchmark datasets and filter sizes. This is discussed in more detail in Supplementary Material I-B.

In summary, these benchmarks suggest that native APR convolutions are suited to overcome some of the difficulties associated with processing large images on the GPU, providing real-time convolution implementations for edge computing, or accelerating CPU implementations to achieve GPU-like performance.

### B. Richardson-Lucy Deconvolution

We demonstrate an application of native APR convolution for image restoration using the iterative Richardson-Lucy (RL) deconvolution algorithm [29], [30]. Deconvolution is a frequent example of an ill-posed inverse problem in microscopy image processing. Suppose we have acquired a blurred (from light diffraction) and noisy (detector shot noise and electronics noise) image

$$\mathbf{u} = \eta(\mathbf{i} * \mathbf{w}), \quad (13)$$

where  $\mathbf{i}$  is the imaged sample (e.g., the spatial distribution of fluorophores in the specimen),  $\mathbf{w}$  is the point-spread function (PSF) of the microscope optics, and  $\eta$  models the noise distribution. Under the assumption of Poisson-distributed noise, the RL algorithm attempts to recover  $\mathbf{i}$  from  $\mathbf{u}$  via the iterative updates

$$\mathbf{i}_{k+1} = \mathbf{i}_k \left( \frac{\mathbf{u}}{\mathbf{i}_k * \mathbf{w}} * \mathbf{w}^\dagger \right), \quad (14)$$

which amounts to a fixed-point iteration for maximizing the likelihood of observing  $\mathbf{u}$ . Division and multiplication are element-wise,  $*$  denotes discrete convolution, and  $\mathbf{w}^\dagger$  is the flipped PSF with the order of elements reversed in each dimension. We adapt this algorithm to the APR by replacing the convolution operations with their APR-native counterparts and restricting the PSF stencil to the different APR levels as described in Eq. 9. In microscopy, the PSF is typically measured experimentally on a grid of the same size as the image, and convolutions are applied in the Fourier domain. Since APR convolution is done in the spatial domain, using the full PSF quickly becomes computationally infeasible. Hence, APR deconvolution using the present discrete multi-level operators requires truncating the PSF.

We evaluate the restoration quality for different truncations using a synthetic microtubule dataset with known ground truth [31]. The 3D image is of size  $128 \times 256 \times 512$  voxels (64 MiB, 32-bit values). Figure 10 shows the ground truth and input images, as well as restorations on pixels and on APR (CR=5.3) for different PSF truncations. Quantitative metrics are given in Table I. As a baseline we consider the result of RL deconvolution using the full PSF on both the original input pixel image and its APR reconstruction. The result on the reconstructed volume is further re-sampled onto




Fig. 10. Maximum-intensity z-projections of the 3D ground truth and corrupted (blurred and noisy) synthetic microtubules dataset (see [31] for details), as well as APR and pixel restorations obtained by Richardson-Lucy deconvolution using the full-sized PSF and different truncations for different numbers of iterations. The reference and input data, as well as the full-sized pixel PSF, are taken from <http://bigwww.epfl.ch/deconvolution/data/microtubules>.

the particles to simulate APR-native processing. This suggests that there is no significant degradation due to the adaptive sampling.

Truncation of the PSF has the largest effect on the restoration quality. This is expected, as the chosen truncation sizes of  $11^3$ ,  $21^3$ , and  $41^3$  pixels only capture 6.3%, 12%, and 23.5% of the total PSF intensity, respectively. However, even with the PSF truncated to  $11^3$  pixels, the APR restoration (APR-RL) appears to have improved contrast and clearer edges, which may aid downstream processing tasks. Increasing the PSF

size leads to higher-quality restorations, at the expense of a proportional (to the number of PSF voxels) increase in the computational load. The restorations using a PSF of size  $21^3$  on both the pixel image and its APR show comparable quality, both visually and quantitatively, but the APR-native algorithm requires  $10\times$  less memory.

While these results are promising, there are a number of limitations to APR-native image restoration. Deconvolution in the Fourier domain using the full PSF is based on a physically accurate problem formulation, which is crucial in

applications requiring interpretability of the data. In addition to truncation of the PSF, APR-RL induces errors due to the lossy representation of the input data, and due to deviations of the APR convolution from the underlying forward process. Despite these limitations, however, the example presented here indicates that APR-based image restoration may be practical in realistic scenarios and leads to significant computational savings. This may be particularly useful in applications involving large volume images, where memory usage is the bottleneck. Higher restoration quality can also be expected in applications where the PSF has smaller support, such that truncation is not required. We present two additional examples of APR-RL in Supplementary Material II to corroborate these findings. One using real microscopy data, the other one using a small synthetic PSF.

## VII. DISCUSSION AND CONCLUSION

We have introduced the data structures and algorithms that enable discrete convolutions to be directly evaluated on Adaptive Particle Representations (APRs) of images using multi-core CPU or GPU acceleration. We have shown how spatial convolution can be equivalently defined on an APR by leveraging the reconstruction condition to interpolate information across resolution levels. We benchmarked the computational efficiency of our APR convolution using a set of synthetic images of varying content densities.

For fixed image sizes, we showed that the computational performance of our CPU and GPU implementation scales with image sparsity as measured by the computational ratio (CR). APR convolution on the GPU achieved speedups of 3.8- to 28-fold compared to CUDA pixel implementations for CR values corresponding to real-world fluorescence microscopy datasets [16]. At the same time, APR convolutions required orders of magnitude less memory, enabling tiling-free GPU convolution of images of tens to hundreds of Gigabytes. For the largest  $CR \approx 1020$  tested, typical of super-resolution or expansion microscopy images, we obtained pixel-equivalent processing throughputs of up to 1 TB/s on an inexpensive gaming GPU, which far exceeds the acquisition rate of current microscopes [4]. This enables complex workflows, comprising tens to hundreds of convolution operations, to be performed in real time during acquisition.

The APR convolution operation is naturally decomposed by resolution level, which offers the ability to modify the filter stencil across resolution levels. We presented two methods to do so: First, inspired by operator restriction in multi-grid methods, the filter can be downsampled to coarser resolutions, such that applying the downsampled filter on the coarser level is consistent with applying the original filter on the finest resolution. Second, the filter weights can be rescaled according to the distance between particles at each level, which is useful, e.g., in finite-difference computations. In general, the filters applied at each level can be independent of one another, providing additional design freedom for content-adaptive filters.

While the parallel CPU implementation of our algorithm accepts convolution kernels of arbitrary size in 1D, 2D, and 3D, our CUDA implementation is currently limited to filter

stencils of size  $3^3$  and  $5^3$ . The main difficulties in extending this to larger filters are the limitations to thread block size and register usage per thread. The efficiency of our implementation hinges on the ability to distribute the sparse rows corresponding to fixed positions in the  $z$ - $x$  plane across warps, in order to promote coalesced access to global memory. With increasing filter size, the amount of such positions quickly outgrows the maximum thread block size, such that each thread must iterate over additional rows to reconstruct isotropic patches. This adds multiples to the workload of each thread and leads to increasingly complex kernels that result in excessive register pressure, requiring compromise to avoid spillage. A possible solution is to decompose convolutions with larger stencils into sums of 1D or 2D convolutions applied line- or plane-wise [32].

In our benchmarks, we have neglected the time required to transfer the image data between the host computer and the GPU. Since data transfer times are proportional to the amount of data being sent, these can be expected to show similar scaling with CR as the memory usage. This means that in practical applications, where transfer times are significant, the overall benefit of the APR is expected to be larger than in the presented benchmarks. This is because the APR not only accelerates image processing, but also reduces the amount of data that needs to be transferred.

Finally, we have only considered discrete convolutions. While the presented approach can in principle be extended to nonlinear filters, morphological operations, and resampling techniques, discrete convolutions remain an essential component of high-level image processing methods. For example, image features computed by convolutions are routinely used in machine learning algorithms, such as random forests and support vector machines, to perform high-level vision tasks such as object detection or instance segmentation. The ability to efficiently convolve an APR enables us to extract per-particle features and perform these tasks directly on the APR, without reverting to pixels. This not only leverages the computational ratio of the APR, but it also benefits the machine learning algorithms, as the number of learning dimensions is reduced.

In the future, this could be exploited to design convolutional neural networks (CNNs) that directly operate on APR images using layers that implement APR convolutions. The filter weights can be learned by back-propagating the loss gradients through the interpolation across resolutions, as presented here. The filter weights for different resolution levels can be learned independently, leading to an inherent spatial scale-adaptivity of the network. The training process itself would then be guided by the spatial structure of the APR. This is in contrast to pixel CNNs, where each filter is applied uniformly across the entire image. We expect that the APR focusing on informative image regions may lead to faster convergence during training, and it may enable smaller, shallower networks. Using our design with fixed sampling points, one would have to consider the structural match between the network inputs and outputs. This restricts the range of possible applications, but should not pose an issue in many important tasks, such as object detection and segmentation.



Taken together, APR-native discrete convolution demonstrates the potential for scalable image processing on CPUs and GPUs. The presented data structures and algorithms provide a basis for APR-native image processing and machine learning. Combined with the results from Ref. [16], which showed the practical applicability of the APR to microscopy datasets, as well as existing software for generating, visualizing, and compressing APRs [23], [33], this brings within reach complete APR-based workflows for storing, visualizing, and processing large sparse image datasets without reverting to pixels. This will allow leveraging the full potential of the APR in big bio-image projects.

#### CODE AVAILABILITY

Algorithms for the generation and manipulation of APRs are implemented in the open-source C++ software library LibAPR [23] (available at <https://github.com/AdaptiveParticles/LibAPR>). The data structures and algorithms described in this work are also implemented and publicly available there. In addition, most APR functionality is available from Python through the package pyapr [33] (<https://github.com/AdaptiveParticles/pyapr>), built on top of LibAPR. The Python package additionally offers interactive visualization tools for APRs (by *z*-slice and by maximum intensity projection raycast), as well as visual aids for observing the conversion of pixel images to APRs.

#### REFERENCES

- [1] E. G. Reynaud, J. Peychl, J. Huisken, and P. Tomancak, "Guide to light-sheet microscopy for adventurous biologists," *Nature Methods*, vol. 12, no. 1, p. 30, 2014.
- [2] K. McDole *et al.*, "In toto imaging and reconstruction of post-implantation mouse development at the single-cell level," *Cell*, vol. 175, no. 3, pp. 859–876, 2018.
- [3] J. Huisken, "Slicing embryos gently with laser light sheets," *BioEssays*, vol. 34, no. 5, pp. 406–411, May 2012.
- [4] R. K. Chhetri, F. Amat, Y. Wan, B. Höckendorf, W. C. Lemon, and P. J. Keller, "Whole-animal functional and developmental imaging with isotropic spatial resolution," *Nature Methods*, vol. 12, no. 12, pp. 1171–1178, Dec. 2015.
- [5] C. Allan *et al.*, "OMERO: Flexible, model-driven data management for experimental biology," *Nature Methods*, vol. 9, no. 3, pp. 245–253, Mar. 2012.
- [6] U. Gunther *et al.*, "Scenery: Flexible virtual reality visualization on the Java VM," in *Proc. IEEE Vis. Conf. (VIS)*, Oct. 2019, pp. 1–5.
- [7] J. Schindelin *et al.*, "Fiji: An open-source platform for biological-image analysis," *Nature Methods*, vol. 9, no. 7, pp. 676–682, Jul. 2012.
- [8] Y. Afshar and I. F. Sbalzarini, "A parallel distributed-memory particle method enables acquisition-rate segmentation of large fluorescence microscopy images," *PLoS ONE*, vol. 11, no. 4, Apr. 2016, Art. no. e0152528.
- [9] L. A. Royer *et al.*, "ClearVolume: Open-source live 3D visualization for light-sheet microscopy," *Nature Methods*, vol. 12, no. 6, pp. 480–481, Jun. 2015.
- [10] R. Haase *et al.*, "CLIJ: GPU-accelerated image processing for everyone," *Nature Methods*, vol. 17, no. 1, pp. 5–6, Jan. 2020.
- [11] J. Moore *et al.*, "OME-NGFF: A next-generation file format for expanding bioimaging data-access strategies," *Nature Methods*, vol. 18, no. 12, pp. 1496–1498, 2021.
- [12] T. Pietzsch, S. Saalfeld, S. Preibisch, and P. Tomancak, "BigDataViewer: Visualization and processing for large image data sets," *Nature Methods*, vol. 12, no. 6, pp. 481–483, Jun. 2015.
- [13] A. Bria, G. Iannello, L. Onofri, and H. Peng, "TeraFly: Real-time three-dimensional visualization and annotation of terabytes of multidimensional volumetric images," *Nature Methods*, vol. 13, no. 3, pp. 192–194, Mar. 2016.
- [14] F. Amat, B. Höckendorf, Y. Wan, W. C. Lemon, K. McDole, and P. J. Keller, "Efficient processing and analysis of large-scale light-sheet microscopy data," *Nature Protocols*, vol. 10, no. 11, p. 1679, 2015.
- [15] B. Balázs, J. Deschamps, M. Albert, J. Ries, and L. Hufnagel, "A real-time compression library for microscopy images," *bioRxiv*, Jan. 2017, Art. no. 164624. [Online]. Available: <https://www.biorxiv.org/content/10.1101/164624v1>
- [16] B. L. Cheeseman, U. Günther, K. Gonciarz, M. Susik, and I. F. Sbalzarini, "Adaptive particle representation of fluorescence microscopy images," *Nature Commun.*, vol. 9, no. 1, p. 5160, Dec. 2018.
- [17] N. C. Rivron *et al.*, "Blastocyst-like structures generated solely from stem cells," *Nature*, vol. 557, no. 7703, pp. 106–111, May 2018.
- [18] V. Ljosa, K. L. Sokolnicki, and A. E. Carpenter, "Annotated high-throughput microscopy image sets for validation," *Nature Methods*, vol. 9, no. 7, p. 637, 2012.
- [19] E. H. Adelson, C. H. Anderson, J. R. Bergen, P. J. Burt, and J. M. Ogden, "Pyramid methods in image processing," *RCA Eng.*, vol. 29, no. 6, pp. 33–41, 1984.
- [20] D. Meagher, "Geometric modeling using octree encoding," *Comput. Graph. Image Process.*, vol. 19, no. 2, pp. 129–147, Jun. 1982.
- [21] P. Porwik and A. Lisowska, "The Haar-wavelet transform in digital image processing: Its status and achievements," *Mach. Graph. Learn.*, vol. 13, nos. 1–2, pp. 79–98, 2004.
- [22] M. Unser, A. Aldroubi, and M. Eden, "B-spline signal processing. II. Efficiency design and applications," *IEEE Trans. Signal Process.*, vol. 41, no. 2, pp. 834–848, Feb. 1993.
- [23] B. Cheeseman, K. Gonciarz, U. Günther, J. Jonsson, M. Emmenlauer, and M. Susik, "Cheesema/LibAPR: Initial release v1.1," Sep. 2018, doi: [10.5281/zenodo.1423158](https://doi.org/10.5281/zenodo.1423158).
- [24] E. Haber, L. Ruthotto, E. Holtham, and S.-H. Jun, "Learning across scales—Multiscale methods for convolution neural networks," in *Proc. AAAI Conf. Artif. Intell.*, vol. 32, 2018, pp. 3142–3148. [Online]. Available: <https://aaai.org/Library/AAAI/aaai18contents.php>
- [25] U. Trottenberg, C. W. Oosterlee, and A. Schuller, *Multigrid*. Amsterdam, The Netherlands: Elsevier, 2000.
- [26] Y. Zhang *et al.*, "A Poisson–Gaussian denoising dataset with real fluorescence microscopy images," 2018, *arXiv:1812.10366*.
- [27] L. Dagum and R. Menon, "OpenMP: An industry standard API for shared-memory programming," *IEEE Comput. Sci. Eng.*, vol. 5, no. 1, pp. 46–55, Mar. 1998.
- [28] J. Malcolm, P. Yalamanchili, C. McClanahan, V. Venugopalakrishnan, K. Patel, and J. Melonakos, "ArrayFire: A GPU acceleration platform," in *Modeling and Simulation for Defense Systems and Applications VII*, vol. 8403. Bellingham, WA, USA: International Society for Optics and Photonics, 2012, Art. no. 84030A.
- [29] W. H. Richardson, "Bayesian-based iterative method of image restoration," *J. Opt. Soc. Amer.*, vol. 62, no. 1, pp. 55–59, Jan. 1972.
- [30] L. B. Lucy, "An iterative technique for the rectification of observed distributions," *Astron. J.*, vol. 79, no. 6, p. 745, 1974.
- [31] D. Sage *et al.*, "DeconvolutionLab2: An open-source software for deconvolution microscopy," *Methods*, vol. 115, pp. 28–41, Feb. 2017.
- [32] A. Eklund, M. Andersson, and H. Knutsson, "True 4D image denoising on the GPU," *Int. J. Biomed. Imag.*, vol. 2011, pp. 1–16, Jan. 2011.
- [33] J. Jonsson, B. Cheeseman, and J. Scholler, "Adaptiveparticles/pyapr: V1.0.0-rc1," May 2022, doi: [10.5281/zenodo.6578053](https://doi.org/10.5281/zenodo.6578053).



**Joel Jonsson** received the M.Sc. degree in engineering mathematics and computational science from the Chalmers University of Technology, Gothenburg, Sweden. He is currently pursuing the Ph.D. degree in computer science with Technische Universität Dresden under the supervision of Prof. Ivo F. Sbalzarini. He is also affiliated with the Center for Systems Biology Dresden (CSBD), the Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI), and the Max Planck Institute of Molecular Cell Biology and Genetics (MPI-CBG). His research interests include machine learning, parallel computing, and biological image analysis.



**Bevan L. Cheeseman** received the Bachelor of Science degree in applied mathematics and the Bachelor of Commerce degree (Hons.) in finance from the University of Auckland, New Zealand, the M.Sc. degree in mathematics and statistics from the University of Melbourne, Australia, and the Ph.D. degree in mathematics from Technische Universität Dresden under the supervision of Prof. Ivo F. Sbalzarini in 2017. During this time, he was affiliated with the Center for Systems Biology Dresden (CSBD) and the Max Planck Institute of Molecular Cell Biology and Genetics (MPI-CBG). Following his Ph.D. degree, he was awarded a HFSP Fellowship to pursue work with the Department of Genetics, University of Cambridge. Currently, he is working at Oxford Nanoimaging (ONI), an Oxford and San Diego-based microscopy start-up, where he leads the Data Science and Cloud Teams. His primary research interests include adaptive resolution representations, data driven discovery, and scalable high-throughput approaches to image processing.



**Suryanarayana Maddu** received the Bachelor of Technology (B.Tech.) degree in mechanical engineering from NIT Karnataka, India, the M.Sc. degree in computational science from Ruhr University Bochum, Germany, and the Ph.D. degree in computer science from Technische Universität Dresden under the supervision of Prof. Ivo F. Sbalzarini in 2021. During this time, he was affiliated with the Center for Systems Biology Dresden (CSBD), the Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI), and the Max Planck Institute of Molecular Cell Biology and Genetics (MPI-CBG). Currently, he is holding an appointment as a Flatiron Research Fellow of the Center for Computational Biology, Simons Foundation, New York City, USA. He also holds a Visiting Scholar appointment with the NSF-Simons Center for Mathematical & Statistical Analysis of Biology, Harvard University. His primary research interests include statistical learning theory, physics-informed machine learning, and scientific numerical computing.



**Krzysztof Gonciarz** received the Master of Science and Engineering degree in computer science with major in information networks and systems from the Lodz University of Technology, Poland. He is currently employed as a Research Software Engineer in the group of Prof. Ivo F. Sbalzarini, with affiliations to the Center for Systems Biology Dresden (CSBD), the Max Planck Institute of Molecular Cell Biology and Genetics (MPI-CBG), and Technische Universität Dresden. His current research interests include image processing, computer vision, and applied deep learning.



**Ivo F. Sbalzarini** received the Graduate degree (Hons.) in mechanical engineering from ETH Zürich, Switzerland, with majors in computational fluid dynamics and control theory and the Doctorate (Dr. sc. techn.) degree (Hons.) in computer science from ETH Zürich under the supervision of Prof. Petros Koumoutsakos. In 2006, he was named as an Assistant Professor in computational science with the Department of Computer Science, ETH Zürich. In 2012, he and his group moved to Dresden, where he became one of the founding members of the Center for Systems Biology and the Chair of Scientific Computing for Systems Biology. Since 2021, he has been the Dean of the Faculty of Computer Science. He is currently a Professor with the Faculty of Computer Science, Institute of Artificial Intelligence, Technische Universität Dresden, and a Director of the Center for Systems Biology Dresden (CSBD). He is also a tenured Senior Research Group Leader with the Max Planck Institute of Molecular Cell Biology and Genetics (MPI-CBG), Dresden, Germany. He also works as an Area Lead for Applied Data Science and AI with the Federal Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI) and the Research Avenue Leader of Scientific Computing and Smart Microscopy with the Federal Cluster of Excellence "Physics of Life." His research interests include biological image analysis, adaptive discretization schemes, data-driven modeling, and parallel computing, all with applications in multi-scale problems of biological systems and the spatio-temporal systems biology of living matter. He received the Willi Studer Award from ETH Zürich for the Graduate degree and the Choras Award from the Weizmann Institute of Science, Israel, for the Doctorate degree.