Image Size Optimization in Light Sheet Microscopy

Strategies, Validation, and Best Practices for Managing Large Volumetric Datasets

MBF Bioscience | SLICE Systems | 2026

1. Introduction 2
2. Benefits of Image-Size Optimization 3
   2.1 Storage Requirements 3
   2.2 Efficient Data Transfer 3
   2.3 Analysis Processing Speed 3
3. Background: Understanding Image-Size Optimization Methods 5
   3.1 Pixel Binning 5
   3.2 Bit-Depth Reduction 5
   3.3 Z-Step Size 5
   3.4 Spatial Downsampling 6
   3.5 Image File Compression: Lossless Strategies and JPEG 2000 8
4. Optimizing Image Size in BrightSLICE Software 12
   4.1 Pixel Binning 12
   4.2 Bit-Depth Reduction 12
   4.3 Z-Step Size 13
   4.4 XY Save Resolution 13
   4.5 Image Resize 13
   4.6 File Format and Compression Ratio 14
   4.7 Summary: BrightSLICE Data Optimization Controls 15
5. Best Practices for Image-Size Optimization 17
6. Conclusion 19

© 2026 MBF Bioscience — MicroBrightField, LLC

1. Introduction

Modern light sheet fluorescence microscopy (LSFM) has revolutionized biological imaging by enabling rapid, high-resolution, three-dimensional imaging of intact specimens, including cleared organs and whole-brain preparations, with minimal photo-damage compared to conventional techniques. Whole-brain imaging, cleared-tissue datasets, and large organ surveys can now be acquired in a matter of hours. However, this breakthrough comes with a significant computing and data management challenge: a single whole-brain acquisition from an MBF Bioscience SLICE™ microscope system may occupy hundreds of gigabytes or more, and a typical experiment can generate terabytes of data in a single day, depending on the number of channels, bit depth, and spatial resolution employed.

The cumulative burden becomes acute when researchers work with datasets of ten, twenty, or more animals. In these scenarios, storage requirements can exceed practical limits of institutional infrastructure, and analysis pipelines may spend more time reading and writing data than performing actual computation. These bottlenecks are not merely inconvenient; they slow scientific discovery and increase operational costs.

Fortunately, the information density of biological images is rarely uniform across spatial scales. Many of the structures of greatest scientific interest, such as cell bodies, projection pathways, anatomical regions, are large relative to the imaging voxel size. This creates an opportunity: by reducing the spatial resolution of image data to match the scale of the structures being analyzed, researchers can achieve dramatic reductions in file size with minimal or no impact on analytical outcomes.

This white paper describes a suite of data-size optimization strategies, including spatial downsampling, bit-depth adjustment, and JPEG 2000 wavelet compression, that enable researchers to dramatically reduce dataset size without compromising scientific validity. Data-size reduction is not a compromise: when properly validated, optimally sized datasets can support the same statistical conclusions, while enabling faster analysis, reduced infrastructure costs, and more efficient collaboration. The guidance presented here is applicable to BrightSLICE™ and MBF Bioscience analysis software, including NeuroInfo®, Neurolucida® 360, and Stereo Investigator®, and is particularly relevant for researchers managing large experimental datasets.

---

2. Benefits of Image-Size Optimization

2.1 Storage Requirements

The most immediate benefit of data-size optimization is a proportional decrease in storage consumption, compared to raw data. For example, acquiring 10 whole-brain datasets with 2x downsampling in the X and Y dimensions and a 2x larger Z-step size reduces total archival storage requirements by a factor of approximately 8 relative to full resolution. This transforms a 6.4 TB archive into roughly 800 GB. This difference can mean the gap between a manageable archive and an infrastructure crisis.

Reduced storage translates into tangible cost savings, particularly when using cloud storage, high-performance storage arrays, or commercial archiving services. It also extends the useful life of existing storage infrastructure and enables researchers to archive larger datasets within fixed storage budgets.

2.2 Efficient Data Transfer

Smaller file sizes translate directly to reduced transfer times for all data movement operations. This is particularly relevant for transfer to cold storage (NAS or long-term archiving systems), retrieval from network storage for analysis, and sharing data with external collaborators, core facilities, or publication repositories.

2.3 Analysis Processing Speed

Analysis speed scales with the number of voxels that must be read, processed, and written. Downsampled datasets therefore process significantly faster across all software analysis pipelines, not because the algorithms themselves change, but because there is less data to move through the computational pipeline.

Importantly, many analysis algorithms operate at reduced resolution internally, regardless of the resolution at which data is provided. In these cases, providing full-resolution data incurs a loading and handling overhead without improving analytical accuracy. Key examples in MBF Bioscience software include:

- **Image filtering:** Some filters are computed at reduced resolution based on filter-specific settings. Using downsampled data as input reduces the overhead associated with internal downsampling prior to image filtering.
- **Atlas registration using Neuroinfo software:** The registration algorithm resamples image data to an appropriate resolution based on that of the reference atlas. Using downsampled input reduces loading time without affecting registration accuracy.
- **Cell detection:** Cell detection algorithms downsample image data before running the detection. You can eliminate this redundant step and substantially reduce processing time by using data that has already been downsampled; operations that take hours on full-resolution data may complete in minutes on appropriately downsampled images.

- **Image visualization:** Loading and rendering optimal size images onscreen requires less GPU and CPU memory, and processes more quickly, improving performance of real-time display.
- **Movie generation:** Creating 3D movies from optimal size images is substantially faster with minimal perceptible quality difference for most presentations.
- **Quantitative fluorescence measurements:** Downsampling has a low-to-moderate impact on fluorescence intensity measurements. Because pixel averaging during downsampling can alter intensity values, validate your downsampling settings carefully against full-resolution reference data before applying them to a full dataset.

---

3. Background: Understanding Image-Size Optimization Methods

Image-size optimization in microscopy encompasses several complementary approaches. Understanding each strategy, and when it is appropriate, is essential for effective data-size management.

3.1 Pixel Binning

Pixel binning is a true downsampling by the camera when acquiring images. By grouping adjacent camera pixels into superpixels, binning reduces spatial resolution while simultaneously increasing signal-to-noise ratio, because more photons are collected per effective pixel, and increasing image-acquisition speed. Binning is most appropriate when the target structures are substantially larger than the diffraction-limited resolution of the optical system, and when image-acquisition throughput is a priority.

Because pixel binned data are never acquired at full resolution, any detail that is lost cannot be recovered. Binning decisions should therefore be made thoughtfully during experimental design, informed by the structure-size considerations described in Section 3.4.

3.2 Bit-Depth Reduction

Bit-depth refers to the number of bits used to represent the intensity value of each pixel in an image. A 16-bit image can encode 65,536 discrete intensity levels, whereas an 8-bit image encodes only 256 intensity levels for the same dynamic range. The default setting for SLICE image-acquisitions is 16-bit which captures the full dynamic range of fluorescence signal; this is particularly important for imaging dim structures or for characterization of subtle differences in signal intensity that are biologically meaningful. While images can be acquired at 8-bit, which will reduce image file size by half, it comes at the cost of reduced intensity discrimination: fine gradations in signal that are distinguishable in 16-bit data may be collapsed into a single intensity value in 8-bit. For many morphological analyses, such as detecting the presence or absence of labeled cells, or tracing neuron-projection pathways, this loss is inconsequential and 8-bit data is entirely sufficient. For quantitative fluorescence intensity measurements however, 16-bit image data is recommended because the reduced dynamic range of 8-bit encoding can introduce rounding artifacts that compromise the accuracy of intensity-based comparisons. Determine the appropriate bit-depth based on the analytical requirements of the experiment and validate any reduction against full-bit-depth reference data before applying it to an entire experimental dataset.

3.3 Z-Step Size

The Z-step size, the axial distance between successive image planes, is not a downsampling operation in the strict sense, but increasing it is a practical and often underutilized method for reducing total image-acquisition time and dataset size. Because the total number of image planes is determined by the tissue depth divided by the Z-step, even modest increases in step size can meaningfully reduce file size and imaging time for thick specimens.

The appropriate Z-step is governed by the optical properties of the objective and the size of the structures being imaged. The axial depth of field of the objective sets a lower bound on the meaningful

Z resolution: sampling more finely than the depth of field does not recover additional structural information. For large structures, a Z-step that is two to five times the diffraction-limited axial resolution is often sufficient, while still ensuring that no significant portion of the structure is omitted between planes. The guiding principle is to image the structure with adequate axial sampling to support accurate detection and morphological analysis, without redundant oversampling that inflates file size without scientific benefit.

3.4 Spatial Downsampling

Spatial downsampling when saving an image reduces the number of pixels or voxels in an image by merging adjacent elements into a single, larger element whose value represents the combined region. Downsampling in the X and Y dimensions reduces the total pixel count by a factor of 4.

The most common downsampling approaches are averaging (where merged pixels take the mean intensity of their constituent elements) and subsampling (where one representative element is selected). Averaging generally preserves more signal information and is preferable for quantitative analysis.

The spatial resolution of the downsampled image is reduced proportionally. A dataset acquired at 1.8 µm pixel size becomes an approximately 3.6 µm pixel size dataset after 2x downsampling. Whether this resolution loss is scientifically meaningful depends entirely on the size of the structures being analyzed; this is addressed below.

Determining the Appropriate Level of Downsampling

Structure Size as the Primary Determinant

The appropriate level of spatial downsampling is determined primarily by the size of the structures you want to analyze relative to the voxel size of the image data. The Nyquist-Shannon sampling theorem gives this a firm theoretical basis: a structure must span at least two voxels in each dimension to be reliably detected. Downsampling raises this minimum resolvable size, so the practical question is simply whether your structures of interest remain well above that threshold at the reduced resolution. Large structures such as neuronal cell bodies can typically accommodate aggressive downsampling, whereas small punctate structures such as synaptic puncta may require full or near-full resolution, since significant downsampling risks merging adjacent structures or causing small objects to fall below the detection threshold. When multiple structure types are present in the same dataset, use the downsampling level appropriate for the smallest structure of interest. Rather than thinking of downsampling as losing information, think of it as deliberately matching your resolution to the scale of the science being performed.

The Visual Appearance of Downsampled Data

When a downsampled image is zoomed to the same display magnification as a full-resolution image, it will appear more pixelated. This visual degradation is expected and is not, by itself, evidence that the data are unsuitable for analysis. The relevant question is not how the image appears when zoomed in, but whether structure-detection algorithms produce results that support the same scientific conclusions.

Downsampling Validation

The most reliable method for identifying the optimal downsampling level for a given experimental paradigm is systematic empirical validation. The goal of validation is not to produce a downsampled image that is indistinguishable from the original, it is to confirm that your experimental conclusions are preserved.

We recommend the following workflow:

1. **Select a representative sample.** Choose a specimen from your experimental dataset that is representative of typical signal quality, labeling density, and tissue condition.
2. **Process at multiple resolutions with identical parameters.** Run the complete analysis pipeline, including any image processing, registration, and cell detection (or other quantitative analysis), on the same specimen at full resolution, 1/2 resolution, and 1/4 resolution, using exactly the same parameter settings at each resolution. Do not re-optimize parameters for each resolution level.
3. **Compare results statistically.** Evaluate total counts and regional distributions across resolutions. Calculate the percentage difference between full-resolution and downsampled results.
4. **Apply the critical question.** Do the experimental and control groups in your dataset show the same statistically significant differences at downsampled resolutions as at full resolution? If yes, the downsampling level is validated for your experimental paradigm.

Interpreting Validation Results

Some variation in structure detection results between downsampled and native data is expected and acceptable. Detection algorithms are inherently imperfect: results can differ slightly between runs even at identical resolution due to stochastic variation. The relevant threshold is scientific equivalence, rather than zero variation.

A practical benchmark: When validating cell detection results, if total cell counts differ by less than a few percent between levels of downsampling, and if regional cell populations show the same patterns, the downsampling level is likely appropriate. The key criterion is whether the results lead to the same experimental conclusions or not. If validation reveals that 1/4 resolution produces equivalent results to full resolution, 1/2 resolution will also be acceptable. For subsequent samples in the same dataset, use the most aggressive validated downsampling level to maximize efficiency.

Experiment-Specific Validation

Validated downsampling levels are specific to a given experimental protocol. It is important to apply the same downsampling level to all samples in an experimental dataset; mixing resolutions introduces a confounding variable. Experiments that target different cell types or structural features typically require independent validations. In general, you should revalidate your established downsampling level when experimental methods, such as clearing chemistry, antibody concentration, or imaging parameters change between experiments.

3.5 Image File Compression: Lossless Strategies and JPEG 2000

Image compression reduces file sizes by encoding data more efficiently, either without any loss of information (lossless) or by selectively discarding data that has minimal impact on scientific conclusions (lossy). Both approaches have a role in light sheet microscopy data management.

For light sheet microscopy data, the MBF Bioscience JPX file format, based on the JPEG 2000 wavelet compression standard, is a particularly effective solution that goes beyond conventional lossless compression. JPEG 2000 offers several advantages over the traditional JPEG format that make it well-suited for scientific imaging: at high compression ratios it produces substantially fewer visual artifacts than conventional JPEG; it supports both lossless and lossy modes within the same standard; it allows images to be decoded at different resolutions and quality levels from a single compressed file; and it supports region of interest (ROI) coding, so that specific regions can be preserved at higher quality than the surrounding regions.

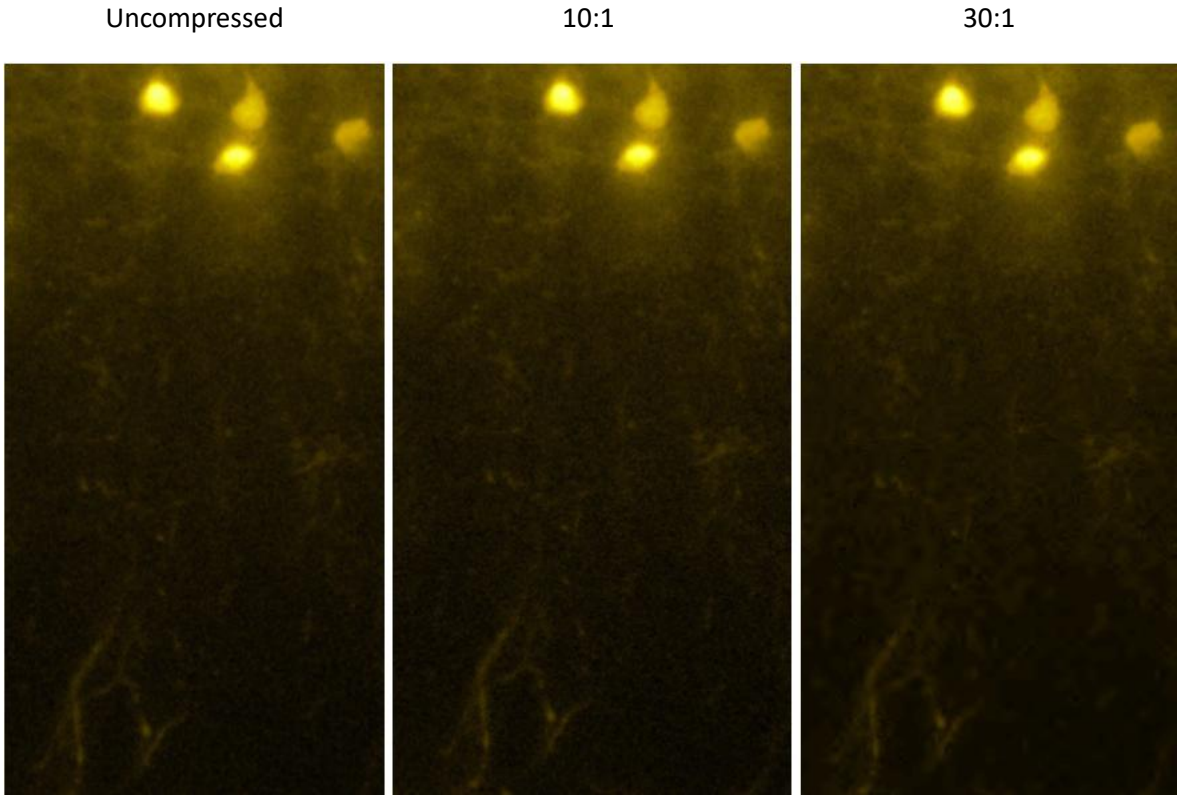
Lossless compression algorithms, such as LZW, DEFLATE, and Zstandard, reduce file sizes by efficiently encoding redundant data without discarding any information. The decompressed file is bit-for-bit identical to the original, making lossless compression an always-advisable baseline strategy with no analytical trade-off. File-size reduction is highly variable and depends on image characteristics: smooth, low-contrast regions compress well, whereas high-frequency detailed images compress poorly. For microscopy images, file-size reduction typically ranges from 20% to 60%. Lossless compression is most impactful when combined with spatial downsampling, which itself increases the proportion of spatially coherent regions amenable to compression.

Compression ratio values correspond to a bitrate setting that is the ratio between the total number of compressed bits and the product of the largest horizontal and vertical image component dimensions. The optimal rate setting is determined for each image, along with numerous other parameters including progression ordering, wavelet decomposition levels, and quality layers.

Generic lossy formats such as standard JPEG are not recommended for scientific microscopy data, because their block-based compression introduces artifacts that can compromise cell detection and fluorescence-intensity measurements. Importantly, compressed files should not be re-saved with additional compression as the quality degradation from chaining lossy compressions is multiplicative.

Figure 1: Comparison of Different Wavelet Based Compression Factors

The left panel shows an uncompressed image. The middle panel shows the same image with 10:1 compression with JP2000 wavelet and the right panel shows 30:1 JP2000 wavelet compression. The image with 10:1 compression exhibits virtually no compression artifacts. At 30:1, there are some minor artifacts and a slight reduction in dynamic range. What does this mean for data storage? If an uncompressed raw image was 1TB, a 10:1 compressed image would be 100GB, which makes storage and file handling far simpler and more rapid.



© 2026 MBF Bioscience — MicroBrightField, LLC

Evaluating File Compression Levels

The central question for any researcher considering lossy compression is: how much compression can be applied before it negatively impacts your ability to draw scientifically valid conclusions from your image data? The answer is application dependent. This section presents a systematic methodology for evaluating image-file compression levels and includes research findings conducted on SLICE microscopy data.

Image Quality Metrics

Use a combination of qualitative and quantitative metrics to assess the impact of compression on scientific data. The following approaches are recommended:

Qualitative Assessment (Visual Inspection)

- **Side-by-side comparisons:** Visually compare the original, uncompressed data with data compressed at various ratios.
- **Expert evaluation:** Have researchers familiar with the biological structures in the images assess the impact of compression on their ability to identify key features.

Quantitative Assessment (Objective Metrics)

- **Peak Signal-to-Noise Ratio (PSNR):** A classic metric that measures the ratio between the maximum possible signal power and the power of corrupting noise. Higher PSNR generally indicates better quality.
- **Structural Similarity Index (SSIM):** A perceptually based metric that considers changes in structural information, luminance, and contrast. A value closer to 1 indicates higher similarity (Wang et al., 2004).
- **Visual Information Fidelity (VIF):** An advanced metric based on information theory that quantifies the amount of visual information lost between the original and compressed image from the perspective of the human visual system. A value closer to 1 indicates less information has been lost (Sheikh & Bovik, 2006).
- **Learned Perceptual Image Patch Similarity (LPIPS):** A state-of-the-art metric that uses a deep neural network to compare images in a way that aligns with human perception. Lower scores indicate greater similarity (Zhang et al., 2018). For scientific images, the VGG network is often preferred over the default AlexNet, as it is more sensitive to the subtle textural changes and fine structural details that are critical in microscopy data.
- **Application-specific metrics:** The most critical evaluation is to measure the impact of compression on the specific quantitative analysis being performed, such as cell counting, segmentation, or fluorescence-intensity measurements.

File Compression Validation Protocol

The following protocol is recommended for determining the optimal file-compression level for a given dataset and analysis workflow:

1. **Select representative image datasets.** Choose a few datasets that encompass the range of image quality and biological structures present in your typical experiments.
2. **Define a range of compression ratios.** Start with a broad range (e.g., 5:1, 10:1, 15:1, 20:1, 50:1, 100:1).
3. **Compress the data.** Use a reliable JPEG 2000 implementation, such as the MBF Bioscience JPX file format, to compress the data at each compression ratio in your protocol.
4. **Perform qualitative and quantitative analysis.** Apply the metrics described in Section 5.1 to the compressed datasets.
5. **Analyze the results.** Plot the quantitative metrics as a function of the compression ratio to identify the point of diminishing returns.
6. **Determine the optimal compression level.** Select a compression level that provides significant file size reduction without unacceptably compromising the scientific integrity of the data.

Research Findings for SLICE Data

Research conducted on SLICE microscopy data identified a compression sweet spot of 15:1 to 20:1 using the MBF Bioscience JPX file format which implements JPEG 2000 wavelet-based compression. Within this range, expert observers confirmed that compression artifacts remain imperceptible, and automated quantitative measurements showed no significant impact on downstream analysis results. At a conservative 15:1 ratio, file sizes are reduced by approximately 93%, transforming a 500 GB dataset into roughly 33 GB, while preserving scientific validity.

These findings apply specifically to SLICE microscopy data. Researchers using other imaging systems should apply the same validation methodology to establish compression parameters appropriate for their specific instruments, sample types, and analysis workflows. SLICE users whose experiments involve highly quantitative measurements such as colocalization analysis, should independently confirm that their chosen compression level does not introduce artifacts that affect their specific analysis. Some workflows may require uncompressed data or a different compression setting than the standard recommendation. Finally, as with all parameters, it is essential to keep compression settings consistent throughout the course of a given experiment: mixing compression levels across samples introduces a systematic variable that can confound comparisons and compromise the integrity of quantitative results.

---

4. Optimizing Image Size in BrightSLICE Software

BrightSLICE, the SLICE microscope control software, provides multiple points in the imaging and image-processing workflow at which data size can be optimized. These controls span hardware-level camera configuration during image acquisition, z-step planning, Image Stitcher resolution settings, and file format and compression choices at export. Understanding where each control fits in the workflow, and its trade-offs, allows researchers to build a coherent, experiment-appropriate data management strategy from the outset.



For detailed information about implementing these controls see Optimizing Image-Data Size in the BrightSLICE user guide.

4.1 Pixel Binning

Pixel binning is the most upstream data reduction option in BrightSLICE; it is applied at the camera's sensor as image tiles are acquired, before data are written to disk. Binning combines adjacent camera pixels into a single superpixel, effectively reducing spatial resolution in the XY plane, increasing signal-to-noise ratio, improving dynamic range, and accelerating image-acquisition speed by reducing the amount of data that must be offloaded from the camera between frames.

The camera used in the SLICE system has a very small native pixel size. This means that moving from 1×1 binning (no binning) to 2×2 binning produces only a modest reduction in optical resolution, considerably less than the equivalent step from 2×2 to 3×3. For experiments that target large structures, such as neuronal cell bodies or somas, 2×2 binning is recommended as a default starting point. For even larger structures, 3×3 binning may provide sufficient resolution and speed up image acquisition even more. Note that binning alters the pixel-to-micron scaling of the resulting images; downstream software must be informed of the correct voxel dimensions to ensure accurate spatial measurements.

Because pixel binning is applied while acquiring images, it cannot be reversed in post-processing. Make the decision to bin deliberately during experimental design, informed by the structure-size validation principles described in Section 3.4.

4.2 Bit-Depth Reduction

The default setting for SLICE acquisitions is 16-bit, which can represent 65,536 discrete fluorescence intensity values. For many biological samples and downstream analyses, this dynamic range exceeds what is scientifically necessary. You can acquire 8-bit images (256 intensity levels for the same dynamic range), reducing file size by 50% without affecting spatial resolution. The trade-off is a reduced ability to discriminate small differences in fluorescence intensity; this may be acceptable for visualization and morphological analysis but could affect quantitative fluorescence measurements. Validate bit-depth reduction in relation to the specific quantitative requirements of each experiment.

4.3 Z-Step Size

The Z-step size, the axial distance between successive image planes, directly determines the number of planes acquired per volume, and therefore has a proportional impact on both acquisition time and final file size. Selecting an appropriate Z-step size avoids unnecessary oversampling while preserving sufficient cross-sectional structural information for accurate experimental analysis. For most SLICE acquisitions, a Z-step size of 4–5 µm is recommended as a practical default, providing a reliable balance between axial sampling density and data volume. Where axial resolution is not critical to the experimental design, increasing the Z-step to 8–10 µm may be appropriate and will yield a corresponding reduction in image size.

4.4 XY Save Resolution

Spatial downsampling by reducing the XY save resolution combines signal from adjacent pixels to produce a lower-resolution stitched output image; it is the post image-acquisition version of pixel binning. The result is a smaller file that is well-suited for target structures that are large enough to be detected at reduced resolution.

As with all spatial downsampling, reducing XY resolution involves a trade-off between file size and image detail. The appropriate setting depends on the experimental objectives and should be validated against full-resolution data using the workflow described in Section 3.4. Once a resolution level is validated for a given experimental paradigm, it can be applied consistently to all subsequent samples in the dataset.

The BrightSLICE Image Stitcher and associated post-processing tools benefit significantly from working with downsampled data. Stitching is a computationally intensive operation, and processing at a reduced resolution can substantially decrease both compute time and peak memory usage, particularly for large, multi-tile whole-organ acquisitions. The Image Stitcher allows users to generate downsampled output images optimized for specific downstream analysis needs, rather than committing exclusively to full-resolution output.

In practice, 1/2 and 1/4 resolution are the most commonly used levels for analysis workflows targeting cell-body-scale structures. For visualization-only purposes or atlas registration, 1/8 or 1/16 resolution may be sufficient. Note that choosing 1/2 resolution is effectively the same as acquiring images with 2x2 pixel binning, but without the advantages of faster imaging and smaller raw data size.

4.5 Image Resize

The Image Resize feature in BrightSLICE performs a similar spatial reduction operation to the XY save resolution and can be applied to previously saved images or to subvolumes extracted from a larger dataset. This makes it possible to generate downsampled working copies from existing full-resolution archives without re-running the full Tissue Scan (image acquisition) workflow. When working with subvolumes (portions) of a larger image, we recommend that you extract subvolumes at full resolution and use the Image Resize feature to save subvolumes at the validated downsampled resolution for analysis.

4.6 File Format and Compression Ratio

File format selection and compression settings have a direct and significant impact on stored file size; optimize these settings and include them in your data-management strategy.

Compression Ratio

BrightSLICE enables you to adjust the compression ratio when saving images. Research on SLICE data indicates that 15:1 to 20:1 compression is optimal in most cases. It reduces file size ~93% but preserves scientific validity. This is the recommended default compression ratio for both working files and long-term archiving. More aggressive compression ratios are possible but should be validated for each data type and analysis workflow. Researchers performing quantitative fluorescence intensity measurements should validate their compression settings against uncompressed or lightly compressed reference data to confirm that intensity values are not meaningfully affected.

File Format: JPX vs. TIFF

BrightSLICE supports saving images in both TIFF and JPX (JPEG 2000 Extended) formats. For most workflows, JPX is strongly preferred over TIFF. JPX files are substantially smaller than their TIFF equivalents for comparable image content, due to the inherent efficiency of the JPEG 2000 wavelet-based compression architecture. For large volumetric datasets, the difference in file size between formats can be dramatic. JPX compression is also available in all other MBF Bioscience commercial software, as well as in the free MBF Bioscience MicroFile+® tool. Unless a downstream analysis tool requires TIFF input specifically, we recommend that you use JPX as your default output format.

4.7 Summary: BrightSLICE Data Optimization Controls

BrightSLICE software controls, described in the table below, provide multiple points in the imaging and post-processing workflow at which data size can be optimized.

For detailed information about implementing these controls see Optimizing Image-Data Size in the BrightSLICE user guide.

Strategy	Workflow Stage	Downsampling?	Recommended Use
Pixel Binning	Image acquisition	Yes (XY)	Pixel binning combines adjacent camera pixels into a single pixel in the image data, effectively reducing spatial resolution in the XY plane. • 10x detection objective: 2×2 pixel binning is recommended as default; 3×3 for large structures • 5x detection objective: 3x3 pixel binning is sufficient for most uses
Bit-Depth Adjustment	Image acquisition	No	Reducing bit depth to 8-bit may be useful for morphological analyses in which fine gradations in fluorescence intensity are not as critical as they are for quantitative analyses.
Z-Step Size	Image acquisition	No (axial sampling)	A Z-step size of 4–5 µm is recommended in most cases. This avoids unnecessary oversampling. Increasing the Z-step to 8–10 µm may be appropriate when the maximum axial resolution is not required.
XY Save Resolution	Post-image acquisition	Yes (XY)	After acquiring images, you can downsample the data by choosing a reduced "XY Save resolution" in the final step of image stitching. Choosing 1/2 resolution is effectively the same as acquiring images with 2x2 pixel binning, but without the advantages of faster imaging and smaller raw data size.
Image Resize	Post-image acquisition	Yes (XY)	Downsample existing image or subvolume files. This is equivalent to downsampling during image stitching using the XY save resolution dropdown menu, but it offers additional customization options.
Compression Ratio	Save/ Export	No	15–20:1 is recommended for SLICE data. This ratio offers ~93% file-size reduction compared to no compression. Validate for quantitative intensity measurements.
File Format (JP2 / JPX)	Save/ Export	No	Substantially smaller than TIFF with no loss of spatial resolution. Available in all MBF Bioscience software and our free MicroFile+ tool. • MBF JPEG 2000 is the recommended file format for use with MBF Bioscience software. • JPEG 2000 format is the recommended file format for use with software from other sources.

These data-reduction strategies are designed to be combined rather than chosen individually; most of them layer across the acquisition, post-acquisition, and export stages to compound their benefit without introducing resolution loss beyond any loss incurred at each step individually. A common example is pairing 2×2 pixel binning at image-acquisition with XY Save Resolution downsampling afterward: 2x2 binning halves the linear XY pixel dimension while acquiring images, and a subsequent 1/2 XY Save Resolution halves it again on the saved file, producing a combined 4× linear (16× total pixel count) reduction relative to native-resolution acquisition.

The same logic extends to most of the other controls in the table above. Bit-depth reduction to 8-bit and Z-step increases operate on entirely separate axes, intensity precision and axial sampling, respectively, and can be applied alongside any XY binning or downsampling combination without interacting with it. At the export stage, File Format and Compression Ratio work together by design: saving as JP2/JPX at the validated 15–20:1 compression ratio is itself the combination, layered on top of whatever upstream binning or downsampling choices were already made. The caveat that carries through every combination is the same one that applies to each technique individually: validate the combined reduction against full-resolution data for the specific measurement of interest before rolling it out across a full dataset, since morphological and visualization workflows tolerate considerably more cumulative reduction than quantitative intensity-based analyses do.

---

5. Best Practices for Image-Size Optimization

- **Plan Data Management at the Experimental Design Stage**

  Consider your storage requirements and analysis timelines before data are acquired. Decisions about camera binning, target resolution, and archiving strategy are most efficiently made during experimental design, when they can be aligned with analytical objectives and experiment size.

- **Match Resolution to the Scale of Structures of Interest**

  The resolution required for accurate detection is determined by the size of the target structures, not by the maximum achievable resolution of the imaging system. Begin validation at 1/2 resolution and extend to 1/4 resolution if initial results are consistent.

- **Always Validate Before Committing to a Downsampling Level**

  Do not assume that a given downsampling level is appropriate without empirical testing. The time investment in validation is negligible relative to the cumulative time and storage savings achieved across a full experimental dataset.

- **Apply Validated Parameters Consistently**

  Once a downsampling level is validated for a given experimental paradigm, apply it consistently to all samples in the dataset. Inconsistent resolution across samples is a methodological confounding factor that complicates interpretation and may invalidate statistical comparisons.

- **Prioritize Results Over Appearance**

  The visual appearance of downsampled images at high magnification is a poor proxy for analytical quality. Focus on whether quantitative results, such as cell counts, regional distributions, statistical differences between groups, are preserved at the reduced resolution.

- **Leverage Internal Algorithm Optimization**

  Many analysis algorithms already operate at reduced resolution internally. Using downsampled data in these algorithms does not degrade accuracy; it simply eliminates the redundant step of loading full-resolution data then subsampling it as the first step of the analysis.

- **Document Your Approach**

  Record which downsampling strategy was applied to each dataset, at which stage of the workflow, and how it was validated. This documentation is essential for methodological reproducibility and for reporting your results to the scientific community in publications and seminars. It also provides a reference for future experiments using similar protocols.

- **Perform Deconvolution Before Compression**

  If deconvolution is part of your processing workflow, deconvolve raw, uncompressed image data before compressing it. Deconvolution algorithms are sensitive to the noise characteristics of the input data, and compression, even at modest ratios, can alter those characteristics in ways that reduce the accuracy of the deconvolution result. Establishing a workflow in which raw tiles are

deconvolved first and then compressed for storage and analysis ensures that both operations perform optimally.

- **Implement a Tiered Storage Strategy**

  A tiered storage system uses compressed and downsampled files strategically across different storage tiers to balance cost, accessibility, and data integrity. A practical framework is to maintain compressed JPX-format working copies on fast, accessible storage for active analysis; archive full-resolution or minimally downsampled datasets on lower-cost, higher-capacity storage such as NAS or cold storage for long-term retention; and use aggressively downsampled copies for sharing with collaborators and for initial exploratory analysis. This approach ensures that the right version of the data is available for each use case without incurring the cost of storing full-resolution copies everywhere.

- **Consider Archiving Raw Data for Critical Experiments**

  For experiments of particular scientific significance, consider archiving the original raw, uncompressed data with lossless compression if storage resources permit. While the 15:1 to 20:1 JPX compression range has been validated for SLICE data, archiving raw data provides an additional safeguard for future analysis using methods or tools that may have different sensitivities to compression. This practice is especially important if the experiment involves extremely precious samples, novel biological phenomena that have not been previously characterized, or if the data may be subject to future regulatory or publication scrutiny.

- **Educate Your Team on Compression Trade-offs**

  Ensure that all members of your research group understand the implications of working with lossy compressed data. While 15:1 to 20:1 JPX compression has been validated as scientifically robust for SLICE data, team members should understand what this means in practice: particularly that compressed files should not be re-compressed, that quantitative intensity measurements may require separate validation, and that the compression parameters used for each dataset should be recorded as part of standard data documentation.

---

6. Conclusion

Data-size reduction is not a concession to computational or infrastructure limitations; it is a principled optimization strategy that, when properly implemented, strengthens research workflows without compromising scientific rigor. The key insight is that the information required to answer a scientific question is rarely the same as the maximum information content of the raw imaging data. By aligning data storage and processing resolution with the actual analytical requirements of each experiment, you can achieve order-of-magnitude improvements in efficiency.

Spatial downsampling is the most impactful single strategy available to SLICE users for managing large volumetric datasets. Combined with complementary approaches such as camera binning and JPX compression at the validated 15:1 to 20:1 range, it provides a comprehensive framework for sustainable data management across large experimental datasets.

The path to effective data size reduction is systematic: validate carefully against full-resolution data, apply validated strategies consistently within each experimental dataset, and focus on whether scientific conclusions are preserved. With this framework in place, the time and storage savings compound with every additional sample acquired.

Downsampling and compression, properly validated and consistently applied, are among the highest-return optimizations available to researchers working with large, light-sheet microscopy datasets.

For questions about implementing these strategies with specific MBF Bioscience software, contact MBF Bioscience support or consult the online documentation at mbfbioscience.com/product-resources.

References

Wang, Z., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. (2004). Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing, 13*(4), 600–612.

Sheikh, H. R., & Bovik, A. C. (2006). Image information and visual quality. *IEEE Transactions on Image Processing, 15*(2), 430–444.

Zhang, R., Isola, P., Efros, A. A., Shechtman, E., & Wang, O. (2018). The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 586–595).