Parameter selection in Leonardo

Both Leonardo-DeStripe and Leonardo-Fuse are distributed with well-tested default parameter settings that were optimized across diverse volumetric microscopy datasets and are sufficient for most practical applications. In typical use cases, no parameter adjustment is required.

In more challenging imaging conditions, however, selective tuning of a small number of user-adjustable parameters may further improve performance or adapt the method to specific acquisition characteristics. Importantly, these parameters primarily modulate the trade-off between artifact suppression and preservation of genuine biological structures, for Leonardo-DeStripe, as well as computational efficiency, for both Leonardo-DeStripe and Leonardo-Fuse.

Below, we describe the key user-adjustable parameters in each module, explain their functional roles, and outline recommended strategies for incremental adjustment. A complete technical description of all parameters is provided in the online documentation (https://leonardo-toolset.readthedocs.io/en/latest/api.html).

1 Leonardo-DeStripe: parameter discussion and practical tuning

Leonardo-DeStripe is designed to operate robustly with a single set of default parameters, which were used unchanged across all datasets presented in this work. In our experiments, we did not observe cases of catastrophic failure or instability of the destriping procedure. However, in particularly challenging datasets—such as those with irregular stripe orientations, drifting illumination, or biological structures strongly aligned with the stripe direction—the balance between stripe suppression and structure preservation may require minor adjustment. Such adjustment is typically not required to prevent failure of the method, but rather to refine performance when residual stripe patterns remain visible or when genuine structures appear slightly attenuated.

Key parameters controlling the suppression–preservation trade-off

- `illu_orient` **(str, optional): Activates the ill. prior in the post-processing, which was introduced to explicitly preserve genuine biological structures that align with the stripe direction.** This prior incorporates knowledge of the illumination orientation to prevent unintended attenuation of elongated features during stripe suppression. When illumination direction is known, enabling this parameter improves structural fidelity, particularly in datasets where biological structures are directionally aligned with illumination-induced stripe patterns.

- `allow_stripe_deviation` **(bool, default = False): Enables an additional penalty within ill. prior to account for non-ideal stripe patterns**, such as slightly tilted, wavy, or drifting stripes. When set to True, the method applies stronger suppression of irregular or unstable stripe artifacts. However, it may slightly reduce fine structural detail, particularly in regions containing weak or thin features. We recommend enabling this option only when stripe patterns deviate from a stable, straight orientation and residual artifacts persist after standard parameter adjustment. For datasets with well-aligned and stable stripe structures, keeping this parameter False better preserves structural fidelity. This parameter is effective only when the ill. prior is activated via illu_orient.
- `guided_upsample_kernel` **(int, default = 49): Controls the window size of the guided upsampling strategy in the post-processing module.** Larger kernel sizes promote sample structures being better preserved. However, residual stripes also became more apparent. The effect of tuning it on a zebrafish dataset is given below.

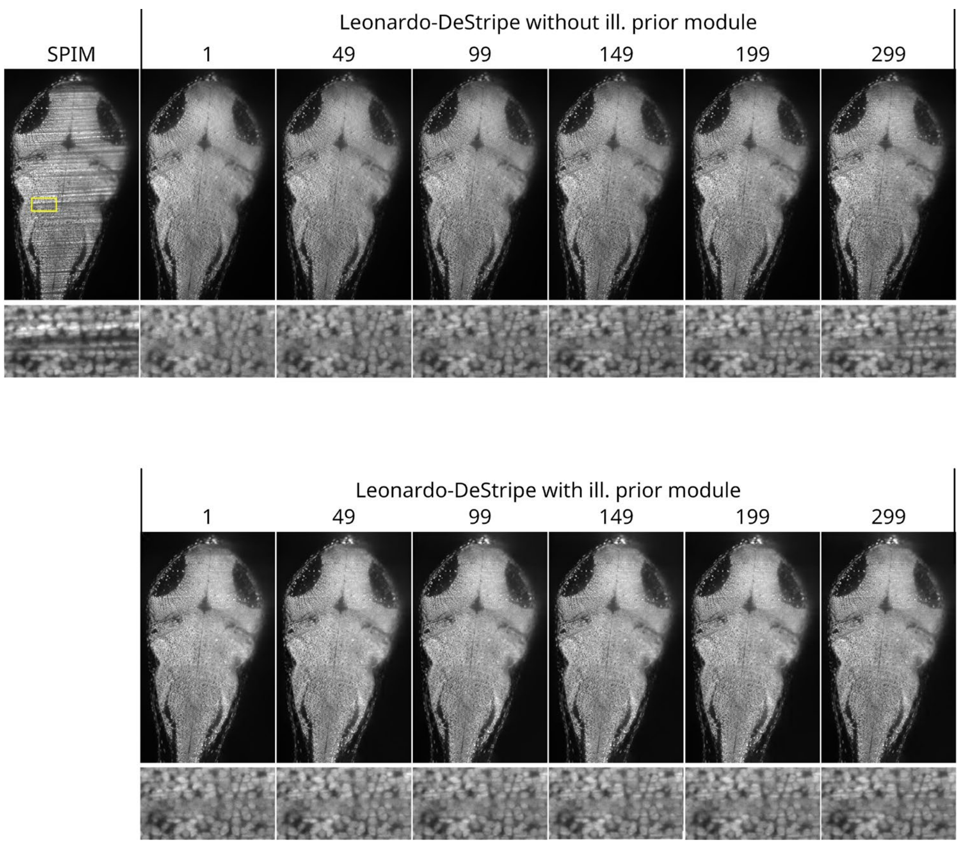

Impact of `guided_upsample_kernel`. To better see the effect, we first applied guided upsampling with different window sizes without ill. prior module (top row). As the window size increased, sample structures could be better preserved; however, residual stripes also became more apparent (see zoomed-in region). Notably, the ill. prior module (bottom row) enhanced robustness and effectively suppressed residual stripes across all window sizes.

- `resample_ratio` **(int, default = 3): Downsampling factor along the stripe direction used when training the graph neural network (GNN).** The default value (3) provides a good balance between computational efficiency and modeling accuracy and was used unchanged in all experiments presented in this study. In most applications, modification of this parameter is not required, particularly when the ill. prior is enabled, as structural preservation is then primarily governed by the post-processing module. If the ill. prior is not activated, modest increases in resample_ratio (up to 5) may, in some cases, help preserve structural integrity. However, excessively large values may impair the network's ability to model fine stripe patterns and are therefore not recommended.

For clarity, other parameters such as the Fourier wedge angle are generally kept at their default values; in practice, the user-facing tuning relevant to stripe-suppression/structure-preservation trade-offs is primarily controlled by illu_orient, allow_stripe_deviation, and guided_upsample_kernel (see parameter definitions above).

Recommended practical adjustment procedure

When parameter refinement is desired, we recommend the following incremental approach:

1. Start from the default parameter configuration.
2. Inspect Leonardo-DeStripe's default result that contains both stripe patterns and biologically relevant structures.
3. If illumination orientation is known and preservation of directionally aligned structures is important, enable the ill. prior via illu_orient.
4. If residual stripe patterns remain visible:
   a) When stripes are irregular (e.g., tilted, wavy, or drifting), enable allow_stripe_deviation.
   b) If residual artifacts persist, modestly decrease guided_upsample_kernel to stronger stripe suppression.
5. After each adjustment, visually assess whether fine structural features remain continuous and well contrasted.
6. In rare cases where the ill. prior leads to undesired stripe removal, users may disable it.

Adjustment should stop once stripe artifacts are sufficiently suppressed without noticeable distortion of genuine biological structures.

2 Leonardo-Fuse: robustness and practical configuration

Leonardo-Fuse was designed to operate reliably under a fixed default configuration and, in typical applications, does not require parameter tuning. All fusion results presented in this work were obtained using default settings. Parameter modification is generally considered only for improving computational efficiency or for accommodating specific dataset characteristics, such as highly sparse samples or extremely large volumetric datasets.

Segmentation-related configuration

- `require_segmentation` **(bool, default = True)**. Segmentation is primarily used to exclude background regions and suppress ghost artifacts near sample boundaries during fusion. Thus, when ghost artifacts are not severe, segmentation can be safely disabled to reduce runtime. Also, if the data is highly sparse, no segmentation is advised.

Fusion boundary estimation

- `resample_ratio` **(int, default = 2): Downsampling factor for estimating the fusion boundary.** Increasing this value reduces memory usage and computational cost and generally works well for datasets with smooth or slowly varying structures. In routine applications, modification is not required. For large volumetric datasets or resource-limited systems, moderate increases may accelerate processing while maintaining stable fusion behavior.

Registration-related configuration (dual-sided detection)

When dual-sided detection requires registration, Leonardo-Fuse performs internal alignment prior to fusion. Default registration settings are sufficient in most cases. Registration-related parameters primarily control computational cost and refinement behavior. Adjustment may be considered when:

- reliable pre-computed registration results are available (e.g., multi-channel scenarios),
- faster execution is desired for large datasets (by increasing axial or lateral downsampling), or
- fine-tuning refinement is unnecessary.

Appropriate settings should maintain accurate spatial alignment between views without introducing visible misregistration artifacts in the fused volume. Visual inspection of anatomical continuity across the fusion boundary provides a practical means of assessing alignment quality.