

FocusStack2XR: Provenance-Preserving Depth-Aware Visualization of Focal-Stacked Specimen Images

Bingxi Li

University of California, Berkeley, Berkeley, California, United States

ABSTRACT

We present FocusStack2XR, a provenance-preserving visualization pipeline and browser-based inspection prototype for focal-stack specimen imagery. Rather than producing metric 3D geometry, the system links displayed regions to their originating focal slices and exposes focus-index maps, confidence cues, conservative validity masks, and relative focus-depth proxies. We evaluate the pipeline on seven focal stacks spanning four specimen categories and 519 source images, and compare Laplacian and Tenengrad focus measures as a cross-measure consistency baseline. The prototype supports all-in-focus viewing, source-slice inspection, and uncertainty-aware browsing while preserving acquisition evidence.

Index Terms: focal stacking, scientific visualization, provenance, uncertainty, extended depth of field, XR.

1. INTRODUCTION

Focal stacking is widely used in macro and microscopy imaging to overcome limited depth of field. Multiple images are captured at different focus positions and fused into an all-in-focus result. Although useful for documentation, this result collapses focal-plane provenance: a viewer can no longer determine which source slice supports a displayed region or whether the focus assignment is ambiguous.

This limitation matters when focal-stack outputs are used in depth-aware or immersive interfaces. A plausible 2.5D presentation may imply geometric fidelity that the captured data do not support, especially for fine hairs, transparent wings, reflective shells, low-texture areas, support pins, and occlusions.

We present FocusStack2XR, a pipeline and early desktop browser prototype for evidence-aware inspection of focal-stack specimen imagery. The contribution is not a new metric reconstruction method. Instead, the system preserves source-slice provenance, exposes focus-derived uncertainty, and provides a relative focus-depth ordering that is explicitly treated as a visualization cue rather than physical geometry. The revised evaluation extends the original single-stack prototype to seven focal stacks from four specimen categories and adds both functional and algorithmic baselines.

Our contributions are:

- a provenance-preserving representation that links displayed regions to source focal slices;

- an uncertainty-aware focus-depth proxy with a conservative validity mask;
- a browser-based inspection prototype with overlay switching, slice browsing, and click-based source reporting; and
- a multi-specimen evaluation with a Tenengrad cross-measure consistency baseline.

2. RELATED WORK

Shape-from-focus methods infer local depth ordering from the focus position that maximizes an image sharpness measure [1]. Focus measures vary in robustness across texture, reflection, and transparency; Pertuz et al. review common operators including Laplacian- and gradient-based measures [2]. Multi-resolution fusion [3] and wavelet-based extended-depth-of-field methods [4] produce high-quality all-in-focus imagery, but conventional outputs do not necessarily expose source-slice provenance or uncertainty.

Provenance is important in scientific visualization because derived views depend on preprocessing, selection, and algorithmic choices [5]. FocusStack2XR applies this principle to focal-stack imagery by treating focus index and confidence as inspectable evidence rather than hidden intermediates. Immersive visualization can improve exploration but also amplify perceived realism [6]; therefore, our XR-facing design emphasizes traceability and explicit limits.

3. SYSTEM OVERVIEW

Figure 1 summarizes the pipeline. A focal stack is processed into an all-in-focus appearance baseline, a focus-index map, confidence values, a conservative validity mask, and a normalized relative focus-depth layer. These representations feed a browser viewer that supports focal-slice browsing, overlay switching, and click-based source-slice inspection.

4. METHOD AND IMPLEMENTATION

4.1 Acquisition and preprocessing

The original prototype used a 24-slice left-lateral stack captured at $1\times$ magnification, $f/4.0$, $1/200$ s, ISO 125, and 0.40 mm focus steps. The camera and pinned specimen were fixed on a motorized rail. The expanded evaluation contains seven stacks, all at 6240×4160 input resolution, totaling 519 source images.

For processing, each image is converted to 8-bit grayscale, normalized to $[0, 1]$, and resized to a maximum width of 1600 pixels while preserving aspect ratio (1600×1067 for all evaluated stacks). No explicit inter-slice registration is applied in the current implementation; residual scale or alignment errors are therefore retained as a limitation.

⁰e-mail: bingxi.li@berkeley.edu

Figure 1. FocusStack2XR pipeline.

From focal-stack macro photography to relative depth and XR-ready specimen visualization.




Figure 1: FocusStack2XR processing pipeline. Focus-derived layers preserve source-slice evidence and uncertainty; relative focus-depth is not metric geometry.

4.2 Focus index, confidence, and validity

For slice I_i , the Laplacian focus response is

$$S_i(x, y) = G_\sigma * |\nabla^2 I_i(x, y)|, \quad (1)$$

using a 3×3 Laplacian kernel and Gaussian smoothing with $\sigma = 2.0$. The selected source slice is

$$F(x, y) = \arg \max_i S_i(x, y). \quad (2)$$

Confidence is the normalized margin between the largest and second-largest responses:

$$C(x, y) = \frac{S_{(1)}(x, y) - S_{(2)}(x, y)}{S_{(1)}(x, y) + 10^{-8}}. \quad (3)$$

A conservative validity mask is formed by thresholding the maximum Laplacian response at its 90th percentile, followed by 5×5 morphological opening and closing. This mask identifies high-evidence analysis regions; it is not claimed as complete specimen segmentation. The current mask may include support pins or acquisition clutter.

4.3 Relative focus-depth and viewer

The normalized focus-depth proxy is

$$D_r(x, y) = \frac{F(x, y)}{N - 1}, \quad (4)$$

where slice 1 is farthest and slice N is nearest. Pixels outside the validity mask are set to zero.

The viewer is implemented in HTML, CSS, and JavaScript. It displays a Zerene PMax all-in-focus baseline [7], raw slices,

and focus-derived overlays. Click coordinates are mapped from the rendered image to the analysis maps; the viewer reports the predicted source slice, normalized focus-depth, and confidence. The default overlay opacity is 0.65 with screen blending. The viewer is a desktop prototype, not a stereoscopic or head-mounted XR implementation.

5. EVALUATION

We evaluated seven stacks: bee front, side, and top views; a reflective beetle-shell top view; two elongated-insect views; and a transparent-wing specimen. Each stack had a corresponding PMax output. Figure 2 shows four representative cases spanning textured, reflective, elongated, and transparent structures.

To provide an algorithmic baseline, we recomputed focus indices with a 3×3 Sobel/Tenengrad magnitude followed by the same $\sigma = 2.0$ smoothing. Table 2 reports exact and ± 1 -slice agreement with the Laplacian result over valid pixels. These values measure *cross-measure consistency*, not ground-truth depth accuracy.

The three bee views show 96.3–97.9% agreement within one slice, and the reflective shell reaches 94.8%. The transparent-wing stack has the lowest median confidence (0.200) and the highest low-confidence fraction (50.0%), consistent with ambiguous evidence in transparent structures. The elongated view has the lowest cross-measure consistency (72.6%) and contains visible acquisition clutter, illustrating failure modes rather than a uniformly successful case. Per-stack processing time at 1600×1067 ranged from 4.4 to 22.1 s on the test laptop.

Table 1: Functional comparison of inspection baselines.

Workflow	All-in-focus	Source-slice link	Confidence	Relative focus-depth	Interaction
Raw focal-stack browsing	—	Direct slice viewing	—	—	Slice navigation
Zerene PMax	✓	—	—	—	Static output
FocusStack2XR	✓	✓	✓	✓ (proxy)	Slice, overlay, click inspection

Table 2: Multi-stack statistics. C_{med} is median confidence over valid pixels; Low- C is the fraction with $C < 0.2$. Agreement is Laplacian versus Tenengrad.

Dataset	Slices	C_{med}	Low- C (%)	Exact (%)	± 1 slice (%)
Bee front	91	0.224	44.5	80.9	96.3
Bee side	75	0.250	39.2	83.8	97.0
Bee top	127	0.274	35.7	86.7	97.9
Beetle shell top	63	0.220	44.9	75.6	94.8
Elongated insect	62	0.266	36.9	54.1	72.6
Long insect left	24	0.230	44.5	62.5	83.2
Transparent wings	77	0.200	50.0	70.4	92.7

6. XR-FACING INSPECTION ROLE

Figure 2 shows the browser prototype. Its current role is to expose acquisition evidence: the user can compare the all-in-focus image with focus-index, confidence, validity, and relative-depth overlays; browse source slices; and click a region to retrieve its predicted source slice. The focus-depth layer is not used to create metric geometry, stereoscopic disparity, or free-viewpoint rendering. A future WebXR or head-mounted implementation could reuse the same evidence layers for spatial placement or layered inspection, but their scientific interpretation would remain a relative focus ordering.

7. DISCUSSION AND LIMITATIONS

The expanded evaluation demonstrates that the pipeline runs across multiple views and specimen types, but it does not establish metric depth accuracy. The Tenengrad comparison is a consistency check between focus measures, not validation against calibrated geometry. The validity mask is conservative and may include pins or clutter; no manual segmentation ground truth was available. No explicit registration is currently applied. In addition, the viewer remains a desktop browser interface and has not been evaluated in AR/MR/VR hardware or with users. These limitations define the next steps: calibrated focus stepping, registration, support-structure annotation, and an XR implementation that preserves the same provenance and uncertainty semantics.

8. CONCLUSION


FocusStack2XR preserves the evidence behind focal-stack-derived visualization rather than presenting an unsupported 3D model. Across seven stacks, the revised evaluation shows strong cross-measure consistency for several specimen/view combinations while revealing transparent and cluttered failure cases. The browser prototype provides a practical foundation for future XR inspection in which relative focus-depth, confidence, and source-slice provenance remain visible and scientifically qualified.

ACKNOWLEDGMENTS


The author thanks the specimen-imaging project staff for access to the setup and permission to use the specimens. OpenAI ChatGPT assisted with language editing, LaTeX formatting, schematic Figure 1, and analysis-script drafting. The author performed data collection, code execution, verification, interpretation, and final technical review. No human participants or external funding were involved. All figures were created by the author from project data. A browser demonstration video accompanies the paper.

REFERENCES

- [1] S. K. Nayar and Y. Nakagawa, “Shape from focus,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 16, no. 8, pp. 824–831, 1994.
- [2] S. Pertuz, D. Puig, and M. A. Garcia, “Analysis of focus measure operators for shape-from-focus,” *Pattern Recognition*, vol. 46, no. 5, pp. 1415–1432, 2013.
- [3] P. J. Burt and E. H. Adelson, “A multiresolution spline with application to image mosaics,” *ACM Trans. Graph.*, vol. 2, no. 4, pp. 217–236, 1983.
- [4] B. Forster, D. Van De Ville, J. Berent, D. Sage, and M. Unser, “Complex wavelets for extended depth-of-field,” *Microscopy Research and Technique*, vol. 65, no. 1–2, pp. 33–42, 2004.
- [5] E. D. Ragan, A. Endert, J. Sanyal, and J. Chen, “Characterizing provenance in visualization and data analysis,” *IEEE Trans. Vis. Comput. Graph.*, vol. 22, no. 1, pp. 31–40, 2016.
- [6] D. A. Bowman and R. P. McMahan, “Virtual reality: How much immersion is enough?” *Computer*, vol. 40, no. 7, pp. 36–43, 2007.
- [7] Zerene Systems LLC, “Zerene Stacker,” software documentation, 2026. [Online]. Available: <https://zerenesystems.com/cms/stacker>



(a) Representative PMax outputs (top) and Laplacian focus-index maps (bottom) across four specimen categories.



(b) Browser prototype showing PMax, a focus-index overlay, and click-based source-slice reporting.

Figure 2: Expanded evaluation and XR-facing inspection prototype. Transparent structures and acquisition clutter expose challenging focus evidence, while the viewer preserves source-slice and confidence information.