https://doi.org/10.2352/issn.2168-3204.2025.22.1.37

Automatic Restoration of Historical Stereoscopic Photographs for 3D Visualization at Scale

Dhruva Gowda-Storz and Sarah Kenderdine; Laboratory for Experimental Museology, EPFL; Lausanne, Switzerland

Abstract

Vast archives of stereographic photographs from the 19th and 20th centuries survive in collections worldwide. While extensively digitized, these artifacts remain largely inaccessible in their intended three-dimensional form. Contemporary stereoscopic displays offer ideal platforms for experiencing these historical media, yet a significant barrier persists: the labor-intensive process of restoring deteriorated stereographs for comfortable viewing. This paper addresses this challenge through two approaches: first, establishing a comprehensive framework for manual stereograph restoration that balances historical authenticity with viewing comfort; second, presenting our ongoing development of an automated pipeline that leverages recent advances in computer vision. Our approach aims to dramatically reduce the time and expertise required for restoration, potentially enabling unprecedented access to historical stereographic archives and facilitating their reintroduction to contemporary audiences through immersive technologies.

Stereoscopic Photography: A Rich Historical Medium Awaiting Rediscovery

Stereoscopic photography captures three-dimensional scenes by taking two photographs of a subject from slightly different viewpoints, mimicking human binocular vision. Unbeknownst to many today, stereoscopy was once a dominant form of optical entertainment in the 19th and 20th centuries, arguably even surpassing conventional photography in popularity for a period. Pioneered by Charles Wheatstone in the 1830s [1] and perfected commercially in the 1850s, stereography rapidly evolved into a global cultural phenomenon.

During its heyday, stereographs were produced and consumed in enormous quantities, used for education, entertainment, 'virtual' tourism, and scientific documentation. Conservative estimates suggest at least seven million distinct stereographic views were created worldwide during its decades of popularity [2], making it one of the most widely consumed visual media of its era. For many in the 19th and early 20th centuries, this three-dimensional medium was not a novelty but the norm for consuming photographic images.

These stereographs documented an expansive range of subjects: historical figures and events, world wars, cultural practices, now-lost monuments, scientific specimens, and scenes of everyday life. They offered viewers an unprecedented sense of presence and immersion in distant places and events—a quality that laid the groundwork for later developments in cinema and virtual reality. Although stereography eventually fell from prominence, its influence on visual culture and entertainment media persisted through the evolution of immersive technology.

Today, while stereography's cultural impact remains largely forgotten by the public, many of whom mistakenly view 3D media as a modern invention, vast collections of stereographs survive in archives worldwide. A significant portion of these collections have been digitized and made available through online portals. These three-dimensional archives represent an extraordinary resource of immersive historical windows holding immense value for historians, media archaeologists, and scholars of visual culture. However, the effective dissemination of this material to broader audiences in its intended three-dimensional form continues to pose significant technical and practical challenges.

Challenges and Potentials for Disseminating Historical Stereographs

Archives have undertaken significant efforts to digitize stereographs as photographic objects, creating digital surrogates that serve preservation, searchability, and basic accessibility needs. However, this conventional digitization approach necessarily omits the experiential dimension of stereography as a three-dimensional medium. With stereographs, there exists a critical distinction between digitization (creating a digital surrogate of the physical object) and what might be called "experiential digitization" (preserving the medium's essential viewing experience). A stereograph digitized as a flat scan (Figure 1) is analogous to preserving a vinyl record as a mere photograph, capturing its physical form while losing its primary function. This fundamental limitation of conventional digitization necessitates alternative approaches to make these historical media accessible in their intended form [3,4].

***Figure 1**. "A wonder to the primitive inhabitants—Santa Fe train crossing Canon Diablo, Arizona." Underwood and Underwood, 1903. From the author's private collection. This figure illustrates how digitized stereographs are most encountered today: as flat images stripped of their intended three-dimensional effect. When properly viewed, these paired images would create a striking illusion of depth that fundamentally transformed the viewing experience.*

To reconstitute the immersive experience of a stereograph, these images require translation into visualization frameworks that accommodate their binocular nature [3]. At the most basic level, digitized stereographs can be reprinted as physical stereo cards and viewed through traditional stereoscopes. Simple digital visualization methods include converting stereographs to animated GIFs (as implemented by some archives for online dissemination) or generating anaglyph representations viewable with inexpensive red/cyan glasses. More sophisticated approaches utilize digital stereoscopic displays, projectors, and virtual reality headsets, which show tremendous potential for high-fidelity dissemination of stereoscopic photographs.

The current cultural zeitgeist appears to prominently feature immersivity and stereoscopy through notions of virtual realities and metaverses, a movement that seems to echo stereoscopy's own popularity and influence more than a century ago. This presents a particularly opportune time for reviving this forgotten historical medium. As Peixoto et al. demonstrate, contemporary stereoscopic visualization technologies can effectively recreate many of the experiential qualities that made historical stereography compelling [3]. Such technologies are experiencing increasing adoption in both personal and institutional settings, creating technological infrastructure well-suited for (re)experiencing historical stereography. However, all these approaches still face a fundamental challenge: the need to translate flat, digitized representations into formats capable of reconstituting the stereoscopic illusion. This translation process requires significant technical intervention that presents substantial barriers to the widespread dissemination of stereographic archives.

Preparing Historical Stereographs for Stereoscopic Visualization through Restoration

The conversion of a digitized 2D scan of a stereograph into its intended three-dimensional format is best described as restoration. Following ICOM-CC's 2008 resolution on terminology, we define restoration as an action that facilitates an object's appreciation, understanding, or use when it has lost part of its function through alteration or deterioration [5]. Restoration of a stereo pair thus involves reinstating the immersive dimension removed by 2D digitization and correcting deteriorations that disrupt stereopsis, recovering what Muñoz-Viñas calls its 'legibility' [6].

While the restoration of stereoscopic photographs is common in 3D visualization projects, the knowledge remains largely tacit among specialists, with only limited published guidelines. Victor Flores [4] outlines a three-step preparation essential for comfortable stereoscopic visualization: first, left and right images must be cropped out of the stereo image, and transposed if pseudoscopic; second, images must be carefully cropped to ensure common visible areas in both views; and third, horizontal and vertical alignment must be achieved to enable comfortable stereopsis. Peixoto & Luz [3] recommend that the left and right stereo views be cropped out of the digitized image to their matching visual content, and dust and scratches be removed only when they interfere with the stereo experience (by creating ghostly parts of the image), otherwise leaving them as historical reminders.

Both approaches highlight that proper preparation extends beyond technical correction to consider the experiential qualities central to the original medium. However, these guidelines are sparse in detail and do not cover the entire gamut of the stereo restoration process in our experience. Flores does not address structural deteriorations to stereo pairs at all, and Peixoto & Luz do not discuss alignment as Flores does. Moreover, Peixoto & Luz's assertion of correcting damage only when necessary is only feasible in newer, less damaged corpora, and does not acknowledge the older material that is extremely difficult to view due to significant deterioration.

There is therefore a need to consolidate the tacit specialist stereo restoration knowledge, published archival guidelines, and 3D cinema postprocessing practices to better document the requirements of the historical stereographic medium.

The restoration of stereographic photographs can be broadly divided into three stages: stereo formatting, asymmetric defect removal, and aesthetic improvement, progressing from most essential to least essential, and from best preserving historical integrity to least preserving. This highlights a fundamental tension in stereograph restoration: balancing historical authenticity with viewing comfort, a tension that informs our three-stage approach.

Stereo Formatting

Stereo Formatting satisfies the core requirements for visualization in 3D, namely, cropping, transposition (if necessary), and alignment.

We recommend a two-pass cropping approach. The first pass separates the left and right stereo views from their frame (which may have various shapes like tombstones, circles, or squares) to create two digital image files. This initial rough crop should extract rectangles circumscribing each stereo image using software like Photoshop. These images can then be positivized if negative and transposed if pseudoscopic [4].

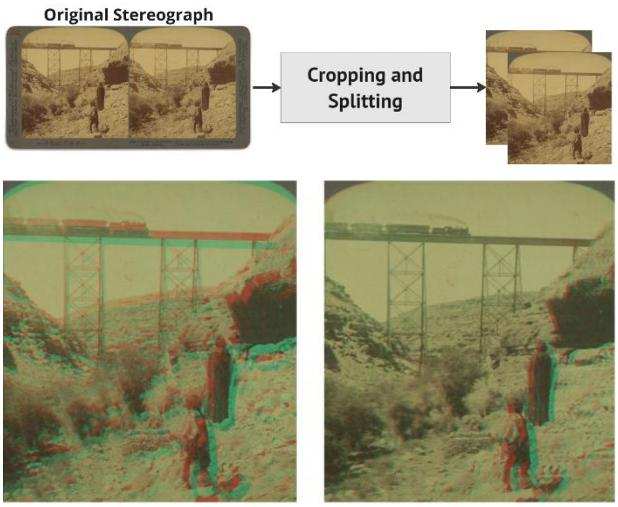

Figure 2*. Cropping and splitting (top) and alignment (bottom) are key preprocessing steps for formatting stereographs for comfortable visualization.*

Alignment follows, best accomplished with stereo-specific software like Stereo Photo Maker [7]. Two types of alignment are necessary: vertical alignment (rectification) eliminates vertical disparity by ensuring corresponding pixels lie on the same horizontal line, while horizontal alignment manipulates the depth budget and zero disparity point (ZDP). Stereo Photo Maker's auto-align function provides a reliable starting point that can be manually refined using anaglyph visualization and should ideally be verified on a 3D display.

After alignment, a second mutual cropping is recommended to remove inconsistencies such as boundary violations: visual features present in only one image [8], and frame-related misalignments that may manifest in older handcrafted stereo cards after our preceding alignment. This mutual cropping, best done in stereo tools like Stereo Photo Maker, should maintain identical aspect ratios and resolutions for the two output images to ensure consistency during visualization.

With these steps completed, stereographs are ready for basic stereopsis, though damage and aging may still require further restoration for comfortable viewing.

Asymmetric Defect Correction

Given the age of the stereographic medium, many surviving stereographs are over a century old and have consequently deteriorated significantly due to handling and aging during this time. Historical stereo pairs commonly exhibit defects such as dust, spots, scratches, cracks, mold (Figure 3), fading, and yellowing (Figure 4).

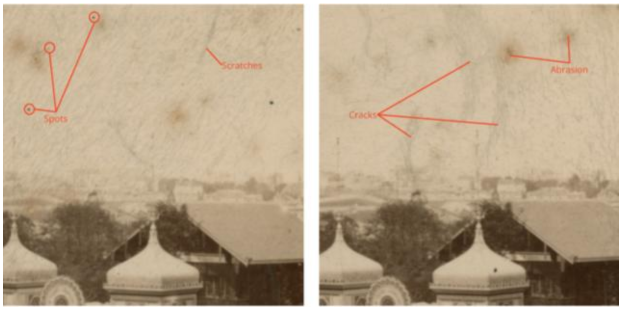

Figure 3. *Examples of damage and deterioration encountered in an albumen stereo card from 1867. Note that the damage in the left view is different from the right.*

In addition to localized structured damage, stereographs exhibit exposure differences between the two views (Figure 4(b)). These differences could stem from the two stereo views being taken at different times, different illuminations from the two viewpoints at the time of capture, limitations in the photographic technology of the time, or other factors.

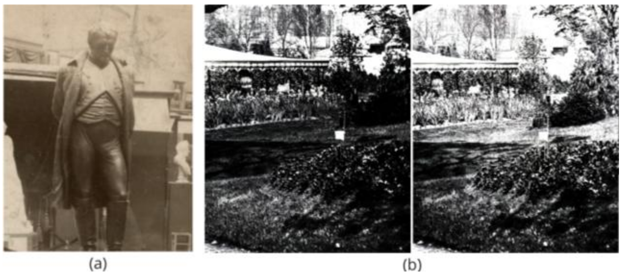

Figure 4*. (a) fading and yellowing of stereo views and (b) localized exposure differences.*

Unlike monoscopic photographs where defects are purely aesthetic, stereograph deterioration significantly impacts 3D visualization due to asymmetry between views. When viewing damaged stereographs stereoscopically, our visual system attempts to fuse the two views while simultaneously processing contradictory information where one eye sees a defect absent in the other view. This triggers binocular rivalry, a visual conflict that inhibits stereopsis, causing viewing discomfort and breaking immersion. Studies have shown that such stereoscopic discomfort can trigger negative emotional responses to the displayed content [9]. While contemporary stereoscopy research addresses various sources of visual fatigue [10], this kind of rivalry is unique to historical stereographs and requires targeted remediation, especially in older material where our visual systems are less tolerant of such inconsistencies.

The goal of defect correction is not to create a perfect image but specifically to target any defects or differences in the two images that may trigger rivalry. As pointed out by Peixoto & Luz [3], there is room for interpretation about how much to remove, since any removal is an alteration to the historical authenticity and lived experience of the stereographic object. Yet, in many cases, rivalry removes the 'legibility' or usability of the stereograph, which should be restored, especially in more extreme examples. This restoration typically targets structured defects and exposure differences.

To address these issues, we employ a systematic approach to defect correction that targets specific types of asymmetries. Structured damage can be corrected using an image processing tool such as Adobe Photoshop and requires careful manual intervention on the image. Minor structured defects, such as dust or specks, can be removed using the selective application of Photoshop's dust and scratches filter, while larger scratches, tears, or blemishes are manually retouched using the clone stamp, healing brush, and remove tools, identifying defects by eye. All restoration steps are cross-checked between the left and right views to ensure that only asymmetrical damage is corrected, and that corrections fully address any rivalry from the original defects.

Exposure differences, which arguably generate more rivalry than structured defects in our experience, can be corrected with the same tools as before. Aligning the global exposure of the two images through levels adjustment helps visually align them. Remaining local exposure differences are more difficult to address, but a common technique is the use of careful manual or AI-assisted object selection and then adjusting the exposures of individual objects or regions to align between the two views.

Validation of successful restoration is straightforward but essential: the corrected stereograph must be viewed in 3D. Remaining sources of rivalry or discomfort become immediately apparent during stereoscopic viewing, making this the definitive test for restoration quality. This testing should be performed periodically throughout the restoration process, not just at its conclusion, to ensure that corrections are effective and to identify any previously unnoticed asymmetries. Ideally, restoration would occur while viewing the image in 3D in real-time, which would immediately reveal remaining defects and verify correction quality. However, this capability isn't directly supported by existing software and requires custom solutions, presenting an opportunity for technical development in this field.

This manual process is highly involved, requiring specialized skill, careful attention, and considerable time investment. Some defects remain camouflaged when viewed monoscopically but appear as significant asymmetries when viewed stereoscopically. Conversely, features that initially appear to be defects may actually be symmetric between views and thus don't require correction. The process therefore relies heavily on expert judgment, and fully restoring an image is often neither feasible nor desirable in most cases. A practical approach is to address major sources of asymmetry until comfortable stereopsis is achieved without compromising historical authenticity. Depending on the severity of deterioration, this process can take anywhere from 30 minutes to several hours per image.

Aesthetic Improvement

Beyond the essential steps of stereo formatting and defect correction, additional enhancements can be applied to improve the viewing experience or restore certain aesthetic qualities.

These improvements, while optional, can significantly impact how viewers engage with historical stereographs.

Historical stereographs often exhibit grain patterns and loss of sharpness due to aging and digitization processes. Subtle denoising using tools like Topaz Photo AI or Photoshop's noise reduction filters can help reduce distracting grain while preserving image detail. Careful sharpening can then restore edge definition that may have been lost. These adjustments should be applied identically to both views to avoid introducing new asymmetries.

Many stereographs show significant fading or yellowing that diminishes their original contrast. Converting yellowed images to black and white can often better represent their original appearance than attempting to preserve the discoloration. Judicious contrast enhancement through techniques like levels adjustment or curves can reveal details that have become obscured over time. When historically appropriate, these adjustments help recover the visual impact of the original stereograph as it would have appeared when new.

Even after the initial formatting stage, further cropping may be beneficial to eliminate problematic edge areas where one view contains details absent in the other (boundary violations). Additionally, cropping can improve composition or adjust the aspect ratio to better suit modern display methods. This refinement stage should aim to enhance the viewing experience without sacrificing significant image content.

Horizontal alignment can be further fine-tuned to manage the "depth budget" of the image—the range between the nearest and farthest points in the scene. Period stereographs were sometimes produced with exaggerated separation that can cause discomfort for modern viewers. Conversely, some images may benefit from slightly enhanced depth. Such adjustments should be approached cautiously, balancing historical accuracy with comfortable viewing, but at this stage, it is ultimately an artistic decision that impacts the viewing experience.

The most interventionist enhancement is AI-assisted colorization of monochrome stereographs. While this introduces elements not present in the original artifact and risks historical misrepresentation, it can enhance immersion and engagement for general audiences. If employed, colorization should be clearly identified as a modern interpretation rather than a historically accurate representation. The open-source DeOldify is the best model for this, in our opinion.

These aesthetic enhancements exist on a spectrum of intervention, from minimal (noise reduction) to substantial (colorization). Each step further from the original introduces both benefits and ethical concerns. The decision to apply such enhancements should consider the intended audience, display context, and purpose of the digitization project. In all cases, it is recommended to maintain both the original scans and unenhanced formatted stereo pairs alongside any aesthetically improved versions, preserving access to the less mediated historical record.

The High Cost of Manual Restoration as a Barrier to Wider Dissemination

We became aware of the cost of restoring stereo pairs while attempting to virtually reconstruct the Universal Exposition of 1867 on our panorama+ stereoscopic visualization system. This event was extensively documented with stereo photography, and our laboratory acquired more than 1400 stereo views that covered a large portion of the exposition. We quickly realized that the time and money required to prepare a stereo dataset for visualization, while manageable for smaller collections of stereo pairs, quickly became unfeasible when working with larger collections of images, especially ones as old and damaged as ours.

Splitting, transposing, and aligning stereographs can be done relatively quickly for individual images, but this processing time quickly accumulates when working with larger corpora. What took the most time, however, was the digital restoration of stereo images to reduce asymmetrical differences. Restoration is a labor-intensive process that requires high proficiency in image editing software and a keen eye for detail, as it is often challenging to separate defects from image features that look like defects. Great care must be taken not to introduce new sources of asymmetry when repairing defects. Depending on the resolution of the stereograph and the degree of damage, restoring a stereo pair could take anywhere from 30 minutes to two hours until completion, following which it needs to be verified on a stereo display and possibly corrected again. The work hours required for this workflow rapidly scale with larger and larger archives, costing large-scale projects such as ours a lot of time and money.

The high cost of preparing stereographs for visualization at scale is a significant factor that extends far beyond our project. The cost of restoring collections of stereographs in poor condition (often older material) makes them unfeasible to work with in projects with limited budgets, making it easier to skip this material and select stereographs in better condition. These factors can affect curatorial freedom and selectively suppress the dissemination of especially damaged material, regardless of its historical value. More importantly, this cost has far greater consequences for archives as their collections can number in the hundreds of thousands, requiring strong justification to process stereographs at this scale for dissemination. This cost is therefore a key barrier to the widespread dissemination of digitized historical stereographs.

The rapidly increasing adoption of virtual reality and stereoscopic visualization in personal, commercial and museological spheres, and their compatibility with visualizing the stereographic medium in high fidelity make them excellent candidates for disseminating stereoscopic archives. Finding a way to automate these preprocessing steps to make the process cheaper and more efficient would help facilitate unprecedented access to stereographic archives and provide new avenues to reintroduce them to the public.

Automatically Restoring Historical Stereographs at Scale

Building on the manual processes described above, we are developing an open-source pipeline to automate the restoration of historical stereographs at scale, addressing the prohibitive time and cost barriers currently limiting wider dissemination.

Several excellent methods for the automatic restoration of photographs [11,12] and film [13] have been proposed, but automatic stereograph restoration remains a niche area of research. To the best of our knowledge, only one paper attempts this [14], focusing more on visual improvements than correcting asymmetries for comfort optimization. This makes our proposed pipeline novel in scope, breadth, and depth.

Our pipeline is being developed in dialogue with archivists and stereo restoration experts to ensure it meets the needs of its end users. We've designed it around four key principles: modularity, allowing components to be used together or integrated into other workflows; non-destructive processing that preserves transparency and reversibility; historical integrity that minimizes interventions compromising authenticity; and scalability for efficiently processing large collections.

The pipeline's overall architecture follows the workflow illustrated in Figure 5, progressing from basic formatting through defect correction to final refinement. We are initially developing the system using our dataset of stereographs from the Universal Exposition of 1867 (containing albumen stereo cards, glass plates, and negatives from numerous archives) as a test case, with plans to expand compatibility to handle most varieties of stereographs available online.While still under active development, this architecture establishes the framework for our ongoing implementation efforts.

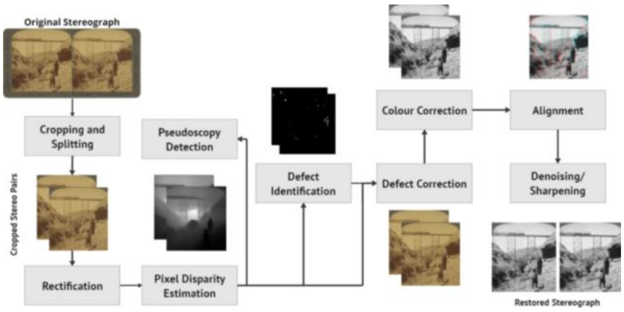

***Figure 5**. An overview of our proposed pipeline.*

The Pipeline

The pipeline begins with stereo formatting, which involves automatically cropping the stereo images from the scan using color and feature-based thresholding (such as the watershed method), and rectifying them (vertical alignment). After implementing various standard and specialized stereo rectification algorithms, including Luo et al.'s "single plate stereo rectification algorithm" [15], we found that the most robust and suitable model is Masuji Sato's stereo auto-align command-line tool [16], which we integrate into our pipeline as an automatic rectification engine.

Once we have our formatted stereo images, we move to the defect correction stage, where our approach leverages the 3D information encoded in the stereo pairs to identify asymmetries. This approach has only become viable in recent years due to advancements in deep learning-based optical flow and stereo matching models that can deliver highly accurate dense disparity maps.

Our implementation currently employs two such models: RAFT [17], a robust optical flow algorithm commonly used in various stereo projects working with real-world data, and more recently, Foundation Stereo [18], a foundation model for stereo matching that delivers highly accurate disparity maps with impressive generalization capabilities. Although Foundation Stereo outperforms RAFT significantly for our use case, it requires perfectly rectified images with the zero disparity point (ZDP) set at infinity, which is not always feasible with historical material. This remains an active area of improvement in our system.

The disparity maps, which encode the horizontal disparity between corresponding pixels in both stereo images, form the foundation for subsequent processing. One straightforward application is in identifying pseudoscopy, since we expect foreground pixels to move in a specific direction. An inverted disparity map therefore identifies pseudoscopy in many cases (Figure 6). When detected, the two images are automatically swapped. The other key use of the disparity map is in defect correction and exposure alignment.

***Figure 6**. Disparity maps of correctly transposed stereo pairs: left view with positive values (left) and right view with negative values (right), enabling automatic detection of pseudoscopy.*

To identify defects, we use the disparity maps to warp the two images to match each other, allowing them to be directly compared through subtraction or division to yield a difference map. The warping process leaves gaps due to occlusion, which we identify using a custom occlusion detection algorithm and exclude from the difference maps. These maps encode all features that are not common between the stereo pairs—in other words, features that are asymmetric.

The difference map is then separated into two components: high frequency and low frequency, by applying a Gaussian distribution to extract low-frequency differences (exposure variations) and then subtracting this from the original to isolate high-frequency differences (structured defects).

The high-frequency difference map localizes all high-frequency asymmetries, including defects of sufficient size and magnitude to potentially generate rivalry. However, the difference map also contains asymmetries generated from inaccuracies in the warping that stem from disparity estimation errors, and in some cases image features such as reflections. After appropriate filtering based on a method similar to that used in [12], the high-frequency map can generate a mask that guides inpainting.

Rather than using generative AI to inpaint, our methodology exploits the content redundancy between stereo pairs, repairing damaged regions of one image by sampling undamaged pixels in the other, similar to the approach used in [13]. This ensures that despite alterations to the historical image, we don't add new content but merely redistribute existing information, aligning with our non-destructive processing and historical integrity principles. If the same corresponding region in both images is damaged, we apply bilateral inpainting to maintain symmetry. Finally, the occluded regions we identified, which are unique to one stereograph and therefore don't contain asymmetric defects, are restored through monocular restoration methods such as [12].

The low-frequency component of our difference map measures the exposure differences between the two images, both local and global. By subtracting this difference from one image, we can completely and reliably align its exposure with the other, a task that is normally challenging to do manually.

Our pipeline generates multiple exposure alignments: the left image aligned to the right, the right image aligned to the left, and both images aligned to a middle exposure between them. The latter is computed by subtracting half of the low-frequency difference map from one image and adding half to the other, along with additional techniques for forward and backward warping and occlusion handling.

Finally, the pipeline concludes with optional aesthetic improvement steps: color correction to convert to black and white and improve contrast, horizontal alignment to optimize the depth budget for comfort, and application of denoising and sharpening algorithms to remove sensor noise and improve image quality for visualization.

With these automated approaches to stereo formatting, defect identification, correction, and enhancement, our pipeline aims to replicate the key benefits of manual restoration while dramatically reducing the time and expertise required, enabling processing at scales previously unattainable.

Conclusion and Potential Impact

This paper has introduced a comprehensive approach to restoring historical stereographs for contemporary visualization, establishing manual restoration best practices and presenting our work-in-progress pipeline for automating this process. By addressing the fundamental tension between historical authenticity and viewing comfort, we've developed a graduated framework that respects the original material while enabling effective visualization in modern contexts.

Our developing toolkit will facilitate low-cost processing of stereographs for stereoscopic displays, allowing users to determine their preferred balance between historical fidelity and viewing experience. This flexibility could democratize access to the vast corpus of historical stereoscopic material, potentially reigniting interest in this historically significant medium.

The dramatic reduction in processing costs could enable ambitious projects at previously unfeasible scales, as demonstrated by our Universal Exposition reconstruction. As similar endeavors become viable, we anticipate increased public engagement through immersive exhibitions. This work coincides with the growing presence of virtual reality displays in homes and educational settings, devices that often lack compelling historical content. Our pipeline could help unlock hundreds of thousands of historical stereographs as content for these platforms, providing archives with new channels to share their collections.

By prioritizing transparency, reversibility, and historical integrity in our automated processes, this research demonstrates how computational approaches can ethically extend access to historical visual media. As we continue development, we aim to establish open standards for stereograph restoration that can be adopted across cultural heritage institutions, ensuring this remarkable three-dimensional historical record becomes as accessible as possible for future generations.

Acknowledgements

We thank our collaborators, Dr. Edward Andò, Mallory Wittwer, and Dr. Florian Aymanns from the EPFL Center for Imaging, for their invaluable contributions towards the development of our pipeline.

References

[1] C. Wheatstone, "The Bakerian Lecture.—Contributions to the physiology of vision.—Part the second. On some remarkable, and hitherto unobserved, phenomena of binocular vision (continued)," Phil. Trans. R. Soc. Lond., 142, 1-17 (1852).

[2] W. C. Darrah, The World of Stereographs (Times and News Publishing, Gettysburg, PA, 1977).

[3] R. Peixoto, F. C. Luz, and J. Oliveira, The digitalization of analogue stereo photographs and the creation of the digital stereo archive, Proc. Archiving Conference, Vol. 16, Society for Imaging Science and Technology (2019).

[4] V. Flores, Involuntary Pseudos. The untransposed stereos in the digital archives, Proc. Archiving Conference, Vol. 16, Society for Imaging Science and Technology (2019).

[5] International Council of Museums Committee for Conservation, Resolution on terminology to characterize the conservation of tangible cultural heritage, Proc. 15th ICOM-CC Triennial Conference, New Delhi, 22–26 September 2008. Available: https://www.icom-cc.org/en/terminology-for-conservation.

[6] S. Muñoz-Viñas, Contemporary Theory of Conservation (Routledge, London, 2005).

[7] M. Sato, Stereo Photo Maker [software], available: https://stereo.jpn.org/eng/stphmkr.

[8] R. Baumgartner, 3D Postproduction: Stereoscopic Workflows and Techniques (Routledge, New York, 2014).

[9] M. Barreda-Ángeles et al., Exploring the effects of 3D visual discomfort on viewers' emotions, Proc. IEEE Int. Conf. Image Processing (ICIP) (2014).

[10] K. Terzić and M. Hansard, "Methods for reducing visual discomfort in stereoscopic 3D: A review," Signal Process. Image Commun., 47, 402–416 (2016).

[11] Z. Wan et al., Bringing old photos back to life, Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (2020).

[12] P. Kisilev, Automatic context-aware dust and scratch removal in scanned images, Proc. 15th IEEE Int. Conf. Image Processing (2008).

[13] P. Schallauer, A. Pinz, and W. Haas, "Automatic restoration algorithms for 35 mm film," Videre: J. Comput. Vision Res., 1(3), 59–85 (1999).

[14] M. Fanfani, C. Colombo, and F. Bellavia, "Restoration and enhancement of historical stereo photos through optical flow," Vision (Journal) (2021).

[15] X. Luo, Y. Kong, J. Lawrence, R. Martin-Brualla, and S. M. Seitz, KeystoneDepth: History in 3D, Proc. 2020 Int. Conf. on 3D Vision (3DV), 463–472 (2020).

[16] M. Sato, StereoAutoAlign for Win64 (Ver. 0.30—10 January 2022) [software], available: https://stereo.jpn.org/stereoautoalign/index_wine.html.

[17] Z. Teed and J. Deng, RAFT: Recurrent all-pairs field transforms for optical flow, Proc. ECCV 2020, 402–419 (2020).

[18] B. Wen, M. Trepte, J. Aribido, J. Kautz, O. Gallo, and S. Birchfield, "FoundationStereo: Zero-Shot Stereo Matching," arXiv preprint arXiv:2501.09898 (2025).

Author Biography

Dhruva Gowda Storz is a PhD Candidate at the Laboratory for Experimental Museology at EPFL in Switzerland. Dhruva holds an MSc in artificial intelligence from Imperial College London and a BSc in physics and neuroscience with a focus on computational neuroscience from Maastricht University. Drawing on his background in both the arts and the sciences, Dhruva's research interests lie at the intersection of culture and computing.

Professor Sarah Kenderdine is the director of the laboratory for experimental museology at EPFL, Switzerland, where she leads a team at the forefront of interactive and immersive experiences for galleries, libraries, archives and museums. In widely exhibited installation works, she has amalgamated tangible and intangible cultural heritage with new media art practice, especially in the realms of interactive cinema, augmented reality and embodied narrative. Sarah has produced over 110 exhibitions and installations for museums worldwide.