Towards Maximizing Storage Efficiency in Pathological Whole Slide Imaging: ROI-Based Hybrid Image Compression

M. Omor Faruk
Ontario Centre for Talent Development
Western University Discovery Park
mdomor.faruk@uwaterloo.ca

Mahmud Hasan
Department of Computer Science
University of Western Ontario
mhasan62@uwo.ca

Abstract

Whole Slide Imaging (WSI) has revolutionized digital pathology but presents significant challenges in storing the high resolution images, given the massive file sizes. While standard image compression can reduce WSI file sizes, it applies blind compression on the whole WSI. This might lose potential bandwidth on non-diagnostic regions while compromising diagnostically relevant regions. To address this challenge, we propose a hybrid WSI compression technique that applies lossless compression to annotated clinically relevant Regions of Interest (ROIs) and lossy compression to the background. We implemented and evaluated a framework supporting multiple codecs, including lossless (JPEG-LS, PNG, ZSTD) and lossy (JPEG, JPEG-XL, JPEG2000) variants. Our results, evaluated on a representative set of WSI images, demonstrate compressed file sizes up to $30\times$ *smaller than the uncompressed originals, while preserving diagnostic quality in the annotated ROIs. The method efficiently stores ROI metadata for seamless reconstruction. This work provides a practical and efficient path for integrating task-aware compression into digital pathology workflows.*

1. Introduction

Digital pathology, driven by Whole Slide Imaging (WSI), is transforming the digitization of diagnostic results, clinical decision-making, pathological training, and data analysis. However, a major challenge in its widespread implementation is the massive data generated by high-resolution WSIs, which can range from hundreds of megabytes to gigabytes per slide, creating significant storage and transmission challenges [6, 10, 20].

While traditional image compression methods, such as JPEG/JPEG2000, can significantly reduce file sizes, their uniform application across entire pixel grid is inefficient in this case. This approach risks losing diagnostic details in relevant regions, while unnecessarily allocating bandwidth to compress non-informative background areas[11, 14].

This observation aligns with clinical workflows, where pathologists typically focus on specific Regions of Interest (ROIs) for diagnosis [2, 17]. Therefore, we propose an annotated ROI-based *hybrid* compression strategy that carefully applies lossless compression to diagnostically relevant ROIs and aggressively applies lossy compression to the background. This way, we can achieve substantial storage and bandwidth savings without compromising diagnostic accuracy. In this *proof-of-concept* work, we implement and evaluate such a hybrid compression framework for WSIs. Our key contributions are:

- A flexible methodology that encodes ROIs with state-of-the-art lossless codecs while applying configurable lossy compression to the background.
- Incorporation of ROI metadata directly into the compressed file, enabling seamless full-slide reconstruction.
- A comprehensive comparative analysis of eight different codec combinations on real-world WSI samples, demonstrating the effectiveness of our approach.

2. Background and Related Works

Early research demonstrated that diagnostic accuracy in WSIs can be preserved with lossy compression up to specific compression ratios, establishing a foundation for quality-controlled compression approaches [3, 4, 10]. This observation directly motivates *hybrid* compression frameworks, an active research direction that applies different strategies to diagnostically critical Regions of Interest (ROIs) and non-diagnostic background regions [15]. For example, some methods combine traditional transforms such as the Discrete Wavelet Transform (DWT) with deep learning modules to adaptively preserve critical features [7]. Building on this, deep learning–based compressors can surpass traditional codecs like JPEG-XL in perceptual quality by learning modality-specific features. However, these models often struggle to generalize across diverse datasets and are computationally intensive, limiting their practicality in clinical settings [7].

Beyond compression quality, computational overhead remains a key constraint in whole slide imaging workflows. WSIs are gigapixel-scale, and both encoding and decoding must scale efficiently to support clinical throughput. Methods relying on global optimization or deep neural inference across entire slides often incur substantial runtime and memory costs, hindering deployment [1, 8, 16]. In contrast, region-localized or tile-wise approaches better align with clinical pipelines [5, 23] by enabling parallel processing, bounded memory usage, and scalable handling of multiple ROIs [21, 22].

For clinical studies requiring absolute data integrity, lossless compression remains essential. While general-purpose lossless compressors such as LZMA can achieve high compression ratios, they are often computationally inefficient for very large images like WSIs. In response, specialized dictionary-based compressors [26] such as WISE exploit WSI-specific statistics, reporting average compression ratios of $36\times$ [18]. To mitigate the computational cost of compression and decompression, a parallel line of research explores analysis directly on compressed data. These adaptive decompression strategies can substantially accelerate analysis workflows while preserving diagnostic accuracy [28].

Despite these advances, current WSI compression techniques face limitations. Traditional codecs such as JPEG 2000 and cropping-based methods offer only modest storage savings [9], while hybrid ROI-based approaches provide adjustable fidelity but lack comprehensive benchmarking of modern codecs within clinically relevant workflows. Although deep learning methods achieve high perceptual quality, their computational demands and limited interpretability hinder clinical adoption. Specialized lossless compressors like WISE achieve high compression ratios but operate as monolithic processes without the region-specific flexibility of hybrid systems [12].

This work addresses these gaps through a systematic evaluation of a practical hybrid compression framework. By combining state-of-the-art lossless and lossy codecs, we provide a direct comparison of efficiency while guaranteeing lossless preservation of clinically significant ROIs, user-tunable background quality, and computational scalability suitable for real-world clinical deployment.

3. Methodology

To maximize storage efficiency, our proposed framework applies aggressive lossy compression to non-diagnostic background regions of the WSI. This approach is justified by the established principle that these areas can tolerate significant data reduction without impacting their utility for clinical or computational tasks [25]. The core challenge is to select codecs that provide an optimal trade-off between compression ratio, computational speed, and the perceptual quality of the reconstructed background.

We evaluated several state-of-the-art lossy codecs, including JPEG, JPEG-XL, and JPEG2000, for this work. These codecs were chosen for their widespread usage, efficient implementation, and ability to achieve high compression ratios. The quality of the compressed background is quantitatively assessed using standard metrics such as Peak Signal-to-Noise Ratio (PSNR) [13] and Multi-Scale Structural Similarity Index (MS-SSIM) [27] to ensure that the compression remains within acceptable, visually lossless bounds for the intended use case, even at high compression levels. This allows users to tune the compression level based on their specific storage versus quality requirements. Our proposed pipeline is illustrated in Figure 1.

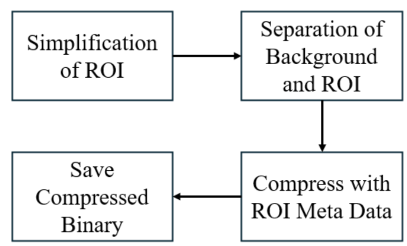

Figure 1. Block diagram of the proposed hybrid ROI-based compression and reconstruction pipeline.

3.1. Proposed Hybrid Compression Framework

The core methodology involves the following steps:

1. *ROI Simplification:* User-annotated ROIs, provided as polygons, are simplified to their minimum bounding rectangle (MBR) for efficient storage and processing.
2. *Spatial Separation:* The WSI is logically partitioned into the ROI region (I_{ROI}) and the background (I_{BG}).
3. *Independent Compression:* I_{ROI} is compressed using a lossless codec to guarantee diagnostic integrity, while I_{BG} is compressed with a lossy codec and a user-tunable quality parameter.
4. *Metadata and Packaging:* The compressed data, along with metadata (image dimensions, ROI bounding box, and segment sizes), is packaged into a single binary file.

3.2. Codec Configurations

We evaluated eight hybrid techniques, each pairing a *lossy background codec* with a *lossless ROI codec*, as defined in Table 1. The user-adjustable parameters for the lossy codecs control the quality-size trade-off for the background.

3.3. Mathematical Formulation

Let the original WSI be defined as a function $I : \Omega \to \mathbb{R}^c$, where $\Omega \subset \mathbb{Z}^2$ is the spatial domain and c is the number of

Table 1. Evaluated Hybrid Compression Techniques

Technique ID	Background (Param)	ROI (Lossless)
T1	JPEG 2000 (40)	JPEG-LS
T2	JPEG (50)	DEFLATE
T3	JPEG (50)	ZSTD
T4	JPEG (50)	LZMA
T5	JPEG (50)	PNG
T6	JPEG (50)	JPEG-LS
T7	JPEG-XL (0)	JPEG-LS
T8	JPEG-XL (0)	JPEG-XL-LS

color channels. The user-defined Region of Interest is given as a binary mask $M : \Omega \rightarrow \{0, 1\}$.

The image is partitioned into two disjoint regions:

$$I_{ROI} = I \odot M \quad \text{and} \quad I_{BG} = I \odot (1 - M) \tag{1}$$

where $\odot$ denotes element-wise multiplication.

We first define the set of compressed components as:

$$\mathcal{S}(I, M, q) = \{\mathcal{C}_{\text{lossless}}(I_{ROI}), \mathcal{C}_{\text{lossy}}(I_{BG}, q), \mathcal{M}\} \tag{2}$$

where $\mathcal{M}$ denotes the metadata required for reconstruction.

The final hybrid compressed representation is obtained by applying a packaging operator $\mathcal{P}(\cdot)$:

$$\mathcal{C}_{\text{hybrid}}(I, M, q) = \mathcal{P}(\mathcal{S}(I, M, q)) \tag{3}$$

which concatenates the compressed components and metadata into a single binary container.

3.4. ROI Simplification to Bounding Box

For each WSI, the diagnostically relevant ROI is initially provided as a set of N 2D coordinates, $S = \{(x_1, y_1), (x_2, y_2), \ldots, (x_N, y_N)\}$, often outlining a polygonal area. To standardize storage and facilitate later processing, we compute the minimum bounding rectangle (MBR) that encloses the ROI:

$$\begin{aligned} x_{\min} &= \min_i(x_i), \quad & x_{\max} &= \max_i(x_i), \\ y_{\min} &= \min_i(y_i), \quad & y_{\max} &= \max_i(y_i), \\ w &= x_{\max} - x_{\min}, \quad & h &= y_{\max} - y_{\min} \end{aligned} \tag{4}$$

where $(x_{\min}, y_{\min})$ are the top-left corner and (w, h) are the width and height of the bounding box. The ROI is thus reduced from an arbitrary polygon to a rectangle, optimizing both compression and metadata size.

3.5. Decompression and Reconstruction

The background and ROI are decompressed independently, and the original image is reconstructed by superimposing the lossless ROI onto the lossy background at the stored bounding box coordinates. A key implementation detail is the addition of an 8 pixel black margin around the ROI during JPEG background compression to prevent block artifacts at the boundary. This workflow ensures perfect reconstruction of diagnostically critical regions while achieving high compression ratios.

4. Results

Our hybrid compression framework was evaluated on a small subset (seven) of WSIs, demonstrating substantial file size reductions while guaranteeing lossless preservation of diagnostically critical Regions of Interest (ROIs). For this *proof-of-concept* study, the WSIs were spatially downsampled to manageable resolutions to facilitate systematic evaluation. The resulting image resolutions ranged approximately from $5{,}000 \times 3{,}000$ to $9{,}000 \times 9{,}000$ pixels, covering a representative spectrum of large-scale WSI dimensions. Perceptual quality was assessed using PSNR and SSIM computed over the entire reconstructed image.

4.1. Compression Performance

State-of-the-art codecs, particularly JPEG-XL, achieved the highest compression efficiency. As illustrated in Fig. 2, the techniques $T8$ $\{$*JPEGXL-LS, JPEGXL(0)*$\}$ and $T7$ $\{$*JPEG-LS, JPEGXL(0)*$\}$ produced the smallest mean file sizes of 7.15 MB and 7.81 MB, respectively. This corresponds to a consistent size reduction of over 95% across all test images. Performance of other combinations of codecs are presented in Fig. 3.

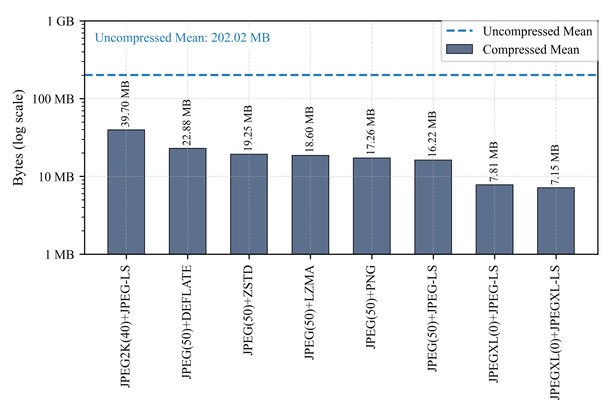

Figure 2. Compressed file sizes by technique. The dashed line shows the uncompressed mean (202.02 MB).

4.2. Background Quality and Trade-offs

The user-tunable property of proposed method allows for a flexible trade-off between background quality and file size. As quantified in Table 2, aggressive compression with JPEGXL(0) yields the smallest files but lower perceptual metrics (PSNR: 24.40 dB, SSIM: 0.66). In contrast,

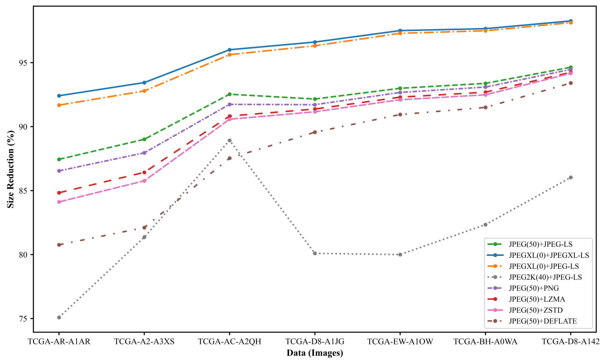

Figure 3. Size reduction across images and techniques. JPEG-XL methods (orange/blue) achieve the greatest and most consistent reduction.

JPEG(50) and JPEG2K(40) preserve background quality near visually lossless levels (PSNR ≥ 30 dB, SSIM ≥ 0.90) at the cost of a larger file size. Visual inspection of high-magnification crops (Fig. 4) confirms that increased compression introduces blurring and texture loss in the background, while the ROI remains perfectly preserved.

Table 2. Mean background perceptual quality for selected codecs.

Compression	Mean PSNR (dB)	Mean SSIM
JPEG2K (40)	34.74	0.96
JPEG (50)	30.60	0.91
JPEGXL (20)	27.79	0.82
JPEGXL (0)	24.40	0.66

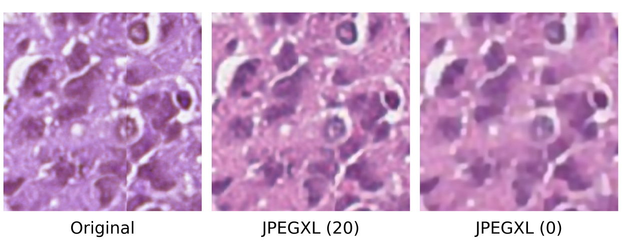

Figure 4. Visual comparison of background compression: original, moderate (JPEG-XL 20), and aggressive (JPEG-XL 0). ROI remains lossless throughout.

4.3. Computational Cost and Scalability

To quantify the practical impact of the proposed framework on clinical workflows, we measured total execution time and compression ratio across all evaluated hybrid configurations on representative CPU hardware. Table 3 summarizes results for seven test cases (TC1–TC7), spanning a range of downsampled WSI resolutions.

Across configurations, execution time scaled with image size and codec complexity. Hybrid schemes using fast lossless ROI codecs (e.g., JPEG-LS) with conventional JPEG background compression achieved sub-second runtimes while providing substantial size reduction. In contrast, configurations employing more complex codecs such as JPEG-XL or LZMA achieved higher compression ratios at the cost of increased execution time, reflecting greater computational overhead.

Peak memory usage remained bounded and scaled with tile size rather than full slide resolution, enabling efficient processing of large WSIs. The framework supports tile-wise and parallel processing, allowing multiple ROIs to be handled independently without global recompression. As ROI coverage increases, both runtime and compression behavior converge toward those of whole-slide compression, consistent with the trends in Table 3.

Table 3. Mean compression ratio and mean total execution time across hybrid compression configurations, averaged over seven test cases (TC1–TC7).

ROI / Background Codec	Mean CR ($\times$)	Mean Time (s)
JPEG-LS / JPEG	13.35	0.50
JPEG-XL-LS / JPEG-XL	49.09	19.61
JPEG-LS / JPEG-XL	38.58	13.17
JPEG-LS / JPEG2000	5.74	18.16
PNG / JPEG	8.56	2.38
LZMA / JPEG	8.26	3.16
ZSTD / JPEG	8.14	3.38
DEFLATE / JPEG	5.59	12.85

Overall, these results show that the proposed hybrid framework enables controlled trade-offs between compression efficiency and computational cost, supporting practical deployment across diverse clinical scenarios.

5. Conclusion

This work demonstrates that a hybrid, ROI-based compression framework can substantially reduce Whole Slide Image (WSI) storage requirements up to 97% while guaranteeing lossless preservation of diagnostically critical regions. Our results show that modern codecs, particularly *JPEG-XL*, provide strong compression efficiency and robustness across diverse tissue types, while allowing user-controlled background quality without compromising ROI integrity. These findings highlight the practicality of ROI-aware compression for addressing the growing data management challenges in digital pathology.

Future work will focus on clinical validation using larger and more diverse datasets such as HISTAI [19], integration of automatic and AI-driven ROI detection, and compatibility with DICOM-WSI and pyramidal storage formats to support efficient random access and broader adoption. We also plan to further benchmark state-of-the-art codecs, investigate alternative ROI representations, evaluate scalability with increasing ROI complexity, and extend analysis to computational cost, downstream task robustness [24], and reproducible open-source deployment.

References

[1] Areej Alsaafin et al. Splice: Streamlining digital pathology image processing. In *International Symposium on Biomedical Imaging*, 2024. Explores computation cost reduction via patch selection. 2

[2] Kimberly Ashman et al. Whole slide image data utilization informed by digital diagnosis patterns. *Journal of Pathology Informatics*, 13:100113, 2022. 1

[3] M. Balkenhol et al. Impact of image compression on histopathological diagnosis. *Virchows Archiv*, 471:107–115, 2017. 1

[4] Yijiang Chen et al. Quantitative assessment of the effects of compression on deep learning in digital pathology image analysis. *JCO Clinical Cancer Informatics*, 4:221–233, 2020. 1

[5] David Clunie. Whole slide imaging dicom supplement 145, 2014. 2

[6] David Clunie et al. Implementing the DICOM standard for digital pathology. *Journal of Pathology Informatics*, 9(1):37, 2018. 1

[7] F. Dai. Deep learning based medical image compression using cross attention learning and wavelet transform. *Scientific Reports*, 15:40008, 2025. 1, 2

[8] N. Dimitriou et al. Deep learning for whole slide image analysis: An overview. *Frontiers in Medicine*, 6:264, 2019. 2

[9] Shahriar Faghani et al. Optimizing storage and computational efficiency: An efficient algorithm for whole slide image size reduction. *Mayo Clinic Proceedings: Digital Health*, 1(3):419–424, 2023. 2

[10] Maximilian Fischer et al. Unlocking the potential of digital pathology: Novel baselines for compression. *Journal of Pathology Informatics*, 2024. 1

[11] Henrik Helin et al. Optimized jpeg 2000 compression for efficient storage of histopathological whole-slide images. *Journal of Pathology Informatics*, 9:20, 2018. 1

[12] Miguel Hernández-Cabronero, I. Blanes, and Joan Serra-Sagristà. Transform optimization for the lossy coding of pathology whole slide images. In *Data Compression Conference (DCC)*. IEEE, 2016. IEEE DCC paper; JPEG2000-oriented WSI coding. 2

[13] A. Hore and D. Ziou. Image quality metrics: Psnr vs. ssim. In *International Conference on Pattern Recognition*, 2010. 2

[14] Elizabeth A. Krupinski et al. Compressing pathology whole-slide images using a human and model observer evaluation. *Journal of Pathology Informatics*, 3:17, 2012. 1

[15] Jonghyun Lee et al. Adaptive compression framework for giga-pixel whole slide images. *Nature Communications*, 2025. 1

[16] Z. Li et al. Computationally efficient adaptive decompression for whole slide image classification. *Journal of Biomedical Optics*, 2023. Discusses computation and memory costs of WSI processing. 2

[17] Zhi Liu et al. A whole-slide foundation model for digital pathology from real-world data. *Nature*, 630:181–188, 2024. 1

[18] Zhenyu Liu et al. WISE: A framework for gigapixel whole-slide-image lossless compression, 2025. 2

[19] Dmitry Nechaev, Alexey Pchelnikov, and Ekaterina Ivanova. Histai: An open-source, large-scale whole slide image dataset for computational pathology, 2025. 4

[20] NEMA DICOM. DICOM Whole Slide Imaging (WSI). https://dicom.nema.org/dicom/dicomwsi/, 2025. Web page describing DICOM WSI tiled multi-frame model. 1

[21] Liron Pantanowitz et al. Guidelines for validation of whole slide imaging systems. *Archives of Pathology & Laboratory Medicine*, 135:171–179, 2011. 2

[22] R. Randell et al. Diagnostic work practices and the adoption of digital pathology. *Journal of Pathology Informatics*, 4:28, 2013. 2

[23] David Taubman and Michael Marcellin. *JPEG2000: Image Compression Fundamentals, Standards and Practice*. Springer, 2002. 2

[24] David Tellez et al. Quantifying the effects of data augmentation and compression on convolutional neural networks in digital pathology. *Medical Image Analysis*, 58:101544, 2019. 5

[25] Vilppu J. Tuominen and Jorma Isola. Linking whole-slide microscope images with dicom by using jpeg2000 interactive protocol. *Journal of Digital Imaging*, 23(4):454–462, 2010. 2

[26] Feng Wang, Tianqi Pan, Jinhui Chen, Line Pouchard, and Chi-Ren Shyu. Managing and querying whole slide images. In *AMIA Annual Symposium Proceedings*, 2012. Standard-oriented WSI storage/querying; tile/region access. 2

[27] Zhou Wang et al. Image quality assessment: From error visibility to structural similarity. *IEEE TIP*, 13(4):600–612, 2004. 2

[28] Farhad Ghazvinian Zanjani, Svitlana Zinger, Bastian Piepers, Saeed Mahmoudpour, Peter Schelkens, and Peter H.N. de With. Impact of jpeg 2000 compression on deep convolutional neural networks for metastatic cancer detection in histopathological images. *Journal of Pathology Informatics*, 10:18, 2019. 2