

# Benefiting from quantum? A comparative study of Q-Seg, quantum-inspired techniques, and U-Net for crack segmentation

Akshaya Srinivasan<sup>1,4</sup>, Alexander Geng<sup>1</sup>, Antonio Macaluso<sup>2</sup>, Maximilian Kiefer-Emmanouilidis<sup>3,4</sup>, and Ali Moghiseh<sup>1</sup>

<sup>1</sup> Fraunhofer Institute for Industrial Mathematics (ITWM),  
Fraunhofer-Platz 1, 67663 Kaiserslautern

<sup>2</sup> Deutsches Forschungszentrum für Künstliche Intelligenz GmbH (DFKI),  
Campus D 3.2, 66123 Saarbrücken

<sup>3</sup> Deutsches Forschungszentrum für Künstliche Intelligenz GmbH (DFKI),  
Trippstadter Str. 122, 67663 Kaiserslautern

<sup>4</sup> Department of Computer Science and Research Initiative QC-AI, RPTU  
Kaiserslautern-Landau, Erwin-Schrödinger-Straße 48, 67663 Kaiserslautern

**Abstract** Exploring the potential of quantum hardware for enhancing classical and real-world applications is an ongoing challenge. This study evaluates the performance of quantum and quantum-inspired methods compared to classical models for crack segmentation. Using annotated gray-scale image patches of concrete samples, we benchmark a classical mean Gaussian mixture technique, a quantum-inspired fermion-based method, Q-Seg a quantum annealing-based method, and a U-Net deep learning architecture. Our results indicate that quantum-inspired and quantum methods offer a promising alternative for image segmentation, particularly for complex crack patterns, and could be applied in near-future applications.

**Keywords** Quantum computing, quantum image segmentation, quantum optimization, image processing, disordered systems

## 1 Introduction

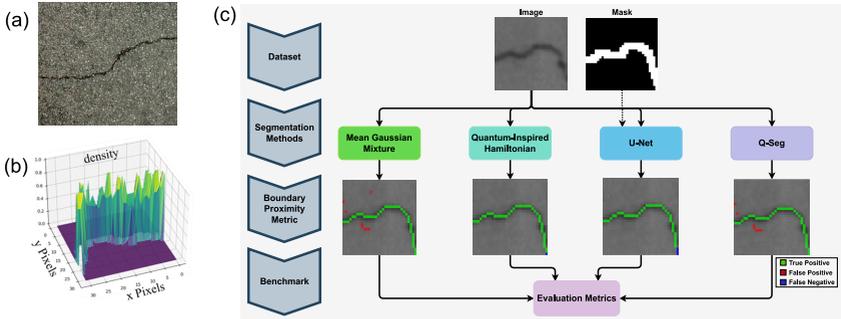
Quantum computing has emerged as one of the leading technologies to improve the efficiency and solvability of complex problems. Still,

the bridge between fundamental and applied research is very narrow and under construction. Unsupervised learning emerges as a particularly promising avenue for the adoption of quantum computing in machine learning. Classical algorithms often struggle to efficiently detect patterns in unlabeled data, a common scenario in many practical applications. Recent advancements have showcased the potential of quantum optimization techniques in addressing unsupervised segmentation tasks [1, 2]. Furthermore, combining quantum computing with classical methods have led to quantum-inspired (QI) and hybrid models like quantum transfer learning, which have been used for example for crack detection [3].

In this paper, we want to build on these developments, and furthermore evaluate how quantum effects in QI methods can be harnessed to advance classical algorithms as well as benchmarking current state-of-the-art approaches. As a use case we have chosen crack-segmentation, a real world problem, which we consider a tremendous important task to evaluate for example the quality of current roads and infrastructure, see Fig. 1 (a). By conducting a systematic comparison between four approaches, where two benefit from quantum, we seek to identify specific areas where non-classical approaches offer advantages. This research not only contributes to the understanding of quantum computing’s practical applications but also guides future developments in algorithm design and implementation within the field.

## 2 Segmentation Techniques

This section examines four methodologies for segmenting concrete cracks: Mean Gaussian Mixture (MGM), QI Hamiltonian, Q-Seg, and U-Net. Using a dataset of  $32 \times 32$  pixel images annotated with ground truth crack locations, each method processes input images to generate segmentation masks that delineate detected cracks. The approaches differ in complexity and computational demands, reflecting advancements in classical and quantum techniques. Figure 1 (c) provides a comparative overview of these workflows.



**Figure 1:** Overview of crack segmentation motivation and methodology: (a) Cracks on roads illustrating real-world infrastructure challenges, (b) Results from the QI approach, accurately identifying crack locations using localized states tied to negative eigenvalues, and (c) Comparative pipeline of crack segmentation methods.

## 2.1 Mean Gaussian Mixture

The Gaussian mixture model is a fundamental image segmentation technique known for its simplicity and efficiency, especially when objects of interest, like pores, have distinct intensity levels. This classical method is computationally inexpensive and versatile, making it ideal for preliminary segmentation tasks. In our study, we adapt Otsu Thresholding [4] to segment cracks in concrete images. Otsu’s method determines the optimal threshold that maximizes between-class variance, effectively separating the foreground (cracks) from the background. To ensure consistent intensity across all samples, each image is normalized to a range of  $[0, 255]$ , addressing ambient lighting variations. Otsu method is applied to 30 images, calculating the optimal threshold for each. For consistency in our comparative analysis, we use the mean threshold across these images as a global threshold, allowing us to benchmark different segmentation methods. This approach balances individual image optimization and comparative consistency across the dataset.

## 2.2 Quantum-Inspired Hamiltonian

Due to the rise of quantum computing also QI-techniques have become more prominent in image processing [5]. In the original context, QI refers to the idea to evaluate classically how quantum effects like superposition, entanglement or wave function collapse (measurements) may change an algorithm of interest [6,7], and in the best case how to benefit from it. Simulating general many-body quantum systems and circuits becomes exponentially difficult as the number of particles or qubits increases. However, many problems can be reduced to polynomial complexity. For example, single-particle Hamiltonians allow each particle to be evaluated separately, with the combined dynamics described as presented here by Fermi-Dirac statistics [8]. We refer to [9] and [10] for a deeper discussion of the underlying physical effect, fitting in the context of this work. In this paper, we show in a proof-of-concept that the single-particle effect of Anderson Localization (AL) [11, 12] can be efficiently used for image- and especially crack-segmentation. Initially used to explain electron behavior in disordered lattices, this model introduces randomness into the potential energy landscape, leading to the localization of wave functions, see Fig.1 (b). Effects of AL and disorder have been found suitable for image and signal processing tasks such as image representation and denoising [10], augmentation [9, 13] and signal transfer in optical fibre [14]. Embedding an image in a Hamiltonian yield a simple matrix form  $N \times N$ , where  $N$  correspond to lattice sites or qubits. To formulate our Hamiltonian matrix, we slightly adjust the adaptive signal decomposition presented in [10]. The key QI idea is that the embedding of the images themselves will self-induce AL due to their disordered landscape, i.e. rough surfaces and cracks which seem like random disordered potentials. Thus, the eigenstates of the Hamiltonian will be close to a unit vector (hot encoded one), only in the strongly disordered areas, and rather extended in areas of weaker disorder. In our study, we embed the image on a 2D lattice. Here it is known that the localization length scales exponentially with disorder strength as well as energy. Thus leading to strong dependence on disorder effects, as particles will mostly localize in areas of the cracks and holes. The embedding works as follows. First, the  $m \times n$  images are flattened  $m \cdot n$ , such that a pixel value at  $A_{ij} \rightarrow a_l$  with  $l = j + ni$ . The corresponding single-particle Hamiltonian of  $N \times N$  size where

$N = m \cdot n$  reads

$$H_{i,j} = \begin{cases} a_i & \text{if } i = j \\ G(a_i, a_j) & \text{if } |i - j| = 1 \text{ and} \\ & i, j \bmod n \neq 0, \\ G(a_i, a_j) & \text{if } |i - j| = n \\ 0 & \text{otherwise} \end{cases} \quad G(a_i, a_j) = \exp\left(-\frac{(a_i - a_j)^2}{2\sigma^2}\right) \quad (1)$$

The  $G(a_i, a_j)$  is the Gaussian difference only for nearest-neighboring pixels and  $\sigma^2$  is the Gaussian variance. The diagonal elements of the Hamiltonian matrix  $a_i$  correspond to the pixel values (potentials), while the off-diagonal  $G(a_i, a_j)$  elements represent the Gaussian weights between nearest-neighboring pixels (kinetic terms). The Hamiltonian in its diagonal form can already be considered as a thresholding technique, however, inferior to the MGM explained in Sec. 2.1. Only due to the kinetic terms we will get extended states which do not contribute significantly to the density of the crack or the whole picture at all. However, we have to be careful to construct kinetic terms for nearest neighbours, as otherwise we might generate extended states again [15]. Furthermore, we have found the Gaussian distance better than constant kinetic terms in [10]. The final mask shows the elementwise summation of magnitudes of all eigenstates (localized particles), tied only to negative eigenvalues. The sum of all negative eigenvalues corresponds to the many-body ground state energy and thus is the minimal energy of the system. This method shows surprising good results and efficiently finds the crack, see Fig.1 (b).

### 2.3 Q-Seg: Unsupervised Quantum Algorithm

Q-Seg is an innovative image segmentation method that utilizes quantum annealing [16, 17]. Initially developed for Earth observation images [2], Q-Seg adapts to detect cracks in concrete by efficiently solving the Maxcut problem using a D-Wave quantum annealer.

The segmentation procedure begins by converting the input image into a lattice graph where each pixel is a node, preserving spatial connectivity. Edges are weighted based on pixel similarity, calculated as squared differences in our case to enhance contrast to gray-scale crack images. The segmentation task becomes a graph cut problem, aiming

to find a maximum cut that best partitions the vertices based on edge weights. To overcome the computational challenges of finding maximum cuts, Q-Seg reformulates the problem into a Quadratic Unconstrained Binary Optimization (QUBO) formulation, suitable for quantum annealing [18]. The QUBO problem is mapped onto the *Pegasus* architecture [19] of the D-Wave quantum annealer, where the system starts in a superposition of all possible states and gradually evolves toward the lowest energy state that represents the optimal solution. The D-Wave quantum annealer iteratively adjusts system parameters and annealing cycles. This iterative adjustment increases the probability of reaching the global minimum.

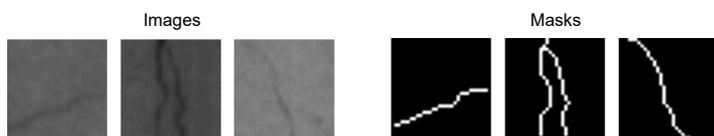
The final result of the quantum annealing process is a binary string corresponding to the segmented image, providing a direct solution to the image segmentation problem. This unsupervised segmentation approach proved effective beyond its original Earth observation application in adapted scenarios such as crack detection in concrete structures, demonstrating Q-Seg's versatility and robustness in various image segmentation tasks.

## 2.4 U-Net

U-Net is a deep-learning architecture designed for biomedical image segmentation [20], renowned for its performance in tasks with limited annotated data. Its versatility extends to various applications, including medical imaging, satellite imagery, and material defect detection. This study focuses on utilizing U-Net for crack detection in concrete, leveraging its strength in producing detailed segmentation masks.

The U-Net architecture features a U-shape consisting of an encoder and a decoder. The encoder reduces the spatial dimensions of the input image using convolutional and max-pooling layers, creating a lower-resolution representation. The decoder upsamples the image, restoring lost spatial dimensions. A notable advantage of U-Net is its skip connections between encoder and decoder layers, which allow access to high-resolution feature maps, improving segmentation accuracy. For binary segmentation tasks like crack detection, a sigmoid activation function classifies each pixel as a crack or background.

In this study, U-Net architecture is modified to handle  $32 \times 32$  pixel crack images, trained on a dataset of 456 labeled patches. The model,



**Figure 2:** Sample images of cracks with corresponding masks.

with approximately 21.7 million trainable parameters, generates a binary mask indicating crack presence. The training utilized a batch size of 16 over 50 epochs, completed in about 13 minutes on a local machine with an Intel Core i7 CPU and 16GB of RAM. U-Net’s ability to capture detailed features makes it a valuable tool for precise and reliable crack segmentation.

### 3 Dataset and Metrics for Segmentation Analysis

This study utilizes grayscale images of concrete for crack segmentation using four different methods. The original images measure approximately  $16,000 \times 32,000$  pixels, with cracks only 1 – 3 pixels wide, making detection challenging for the human eye and machine learning algorithms. To accommodate the limitations of the D-Wave quantum annealer, including the restricted number of qubits and limited runtime, we divide the images into smaller  $32 \times 32$  pixel patches. Our complete dataset consists of 456 manually annotated patches, split into 70% for training and 30% for validation of the U-Net model. We evaluate the performance of the segmentation methods—MGM, QI Hamiltonian, Q-Seg, and U-Net on an unseen test dataset of 30 patches. Figure 2 presents example patches with manually annotated masks highlighting the cracks, which are used as ground truth for performance comparison.

#### 3.1 Evaluation Metrics

To evaluate the segmentation methods, we use the confusion matrix, F1 score, and Intersection over Union (IoU). The confusion matrix provides four key metrics: True Positives (TP), False Positives (FP), False Negatives (FN), and True Negatives (TN), which assess the accuracy of

crack predictions. The F1 score combines precision and recall, while IoU measures the overlap between predicted and ground truth masks, providing a comprehensive evaluation of segmentation performance.

$$F1 \text{ Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad \text{IoU} = \frac{TP}{TP + FP + FN}. \quad (2)$$

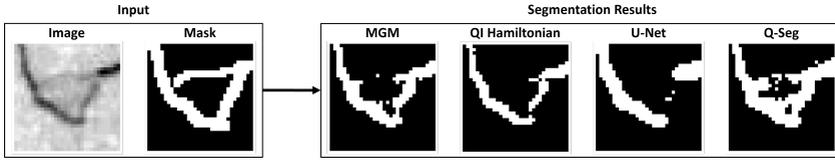
### 3.2 Boundary Proximity Metric

In traditional segmentation tasks, evaluation metrics such as the confusion matrix may not adequately reflect performance when slight deviations in boundary prediction occur. In crack segmentation tasks, where cracks are typically thin structures with irregular boundaries, these minor deviations should be tolerated to some extent. The Boundary Proximity Metric (BPM) addresses this by adjusting the boundary around the cracks in both the predicted segmentation  $I_P$  and the ground truth  $I_{GT}$ , allowing for more lenient evaluation in cases of minor misalignment. The process starts by skeletonizing both the ground truth  $S(I_{GT})$  and predicted segmentation results  $S(I_P)$ . Skeletonization reduces each crack to its core structure, which helps in focusing only on the most critical regions. After skeletonization, both the ground truth and the predicted results are dilated using flat disk structuring element  $B_r$  by a radius of  $r$  pixels. This dilation adjusts the boundary, expanding it to account for small deviations. Any predicted crack pixels that were previously identified as false positives or false negatives but fall within this dilated boundary (i.e., within  $r$  pixels of the ground truth) are then reassigned as true positives. This re-calibration of TP, TN, FN and FP are mathematically formulated as follows

$$\begin{aligned} I_{\widetilde{TP}} &= [S(I_{GT}) \oplus B_r] \cap S(I_P), & I_{\widetilde{FP}} &= [[S(I_{GT}) \oplus B_r] \cap S(I_P)] - S(I_P), \\ I_{\widetilde{FN}} &= [[S(I_P) \oplus B_r] \cap S(I_{GT})] - S(I_{GT}), & I_{\widetilde{TN}} &= I_{ones} - \sum (I_{\widetilde{TP}} + I_{\widetilde{FP}} + I_{\widetilde{FN}}), \end{aligned} \quad (3)$$

where  $I_{ones}$  is  $n \times n$  matrix with all entries equal to 1 and  $n$  is the size of the crack image. The new counts for true positives  $\widetilde{TP}$ , false positives  $\widetilde{FP}$ , false negatives  $\widetilde{FN}$ , and true negatives  $\widetilde{TN}$  are calculated by applying Eq. 3 and

$$\widetilde{X} = \|I_{\widetilde{X}}\|_1 = \sum_i \sum_j I_{\widetilde{X}}(i, j) \quad \text{where} \quad \widetilde{X} \in [\widetilde{TP}, \widetilde{FP}, \widetilde{FN}, \widetilde{TN}]. \quad (4)$$

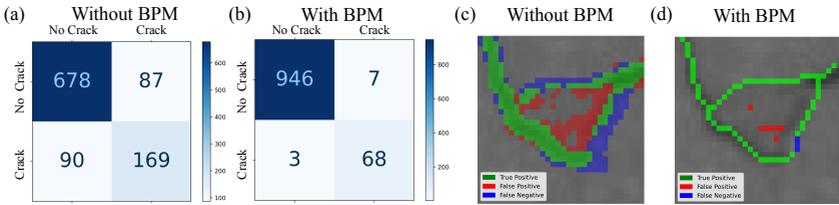


**Figure 3:** Crack segmentation results from four different techniques: MGM, the QI Hamiltonian method, U-Net, and Q-Seg.

Using this boundary proximity metric makes the evaluation more forgiving towards minor misalignment that would otherwise result in a higher count of false positives and false negatives. This approach is especially beneficial in crack segmentation, where small discrepancies in boundary prediction are often unavoidable due to the irregular shapes of cracks.

## 4 Results and Discussion

We have benchmarked MGM, QI Hamiltonian, Q-Seg, and U-Net using standard evaluation metrics and prediction time for segmenting 30 images. Additionally, we employ the BPM to refine the evaluation by considering slight deviations in the predicted crack boundaries compared to the ground truth. Each segmentation method shows distinct results in detecting cracks, as illustrated in Figure 3. This figure underscores the strengths and limitations of each approach in capturing fine details and improving prediction accuracy. Furthermore, Figure 4 demonstrates the visual comparison of the segmentation results, showing both the standard confusion matrix and the one after applying BPM. An overlay diagram illustrates the alignment between the predicted crack masks and the actual cracks. This comparison emphasizes the impact of BPM in improving segmentation accuracy, particularly in challenging cases where the cracks are faint or unclear. Table 1 provides a detailed comparison of the segmentation methods, both with and without BPM adjustments, highlighting their effectiveness in crack detection. The table includes average F1 scores and IoU values for each segmentation technique, allowing for a comprehensive performance assessment. We also present the corresponding prediction times to provide insight



**Figure 4:** Visual comparison of segmentation results, including the standard confusion matrix (a), the confusion matrix post-BPM application (b), and an overlay of predicted crack masks against actual cracks before (c) and after BPM (d).

**Table 1:** Performance comparison of four crack segmentation techniques with and without boundary proximity metric (BPM) using standard evaluation metrics.

Segmentation Methods	Metrics without BPM		Metrics with BPM		Prediction Time (s)
	Avg IoU	Avg F1 Score	Avg IoU	Avg F1 Score	
MGM	$0.5783 \pm 0.1611$	$0.7197 \pm 0.1449$	$0.7454 \pm 0.1836$	$0.8439 \pm 0.1781$	$0.032 \pm 0.007$
QI Hamiltonian	$0.6218 \pm 0.178$	$0.7478 \pm 0.1766$	$0.9447 \pm 0.1241$	$0.9693 \pm 0.1016$	$156 \pm 10$
U-Net	$0.6159 \pm 0.1440$	$0.7522 \pm 0.1145$	$0.8945 \pm 0.1834$	$0.9395 \pm 0.1697$	$2.292$
Q-Seg	$0.5728 \pm 0.1687$	$0.7079 \pm 0.1735$	$0.8014 \pm 0.1431$	$0.8753 \pm 0.1357$	$2.277 \pm 0.250$

into the computational efficiency of each segmentation task. From Table 1, we observe that the QI Hamiltonian method delivers the best overall segmentation performance, with the highest average IoU and F1 score using BPM. Without BPM, it performs within the same error bounds as U-Net, showcasing its robustness. However, its prediction time is significantly longer, which limits its efficiency for real-time applications, especially on larger datasets. U-Net, while slightly behind the QI Hamiltonian method in segmentation accuracy with BPM, is still highly competitive, especially without BPM. However, it demands considerable computational resources, and its training time of 13 minutes for 456 ( $32 \times 32$ ) samples is not included in the prediction time. Q-Seg has a prediction time similar to U-Net, which includes only the Quantum Processing Unit (QPU) access and qubit embedding time. Though it is not as accurate as QI Hamiltonian or U-Net, presents a competitive alternative with balanced performance and does not require labeled data for training, making it practical for scenarios where training data is limited. The MGM model performs comparably to Q-Seg but slightly worse with BPM. However, it is the fastest method, avoiding any training phase like U-Net. Despite its speed, this method lacks

segmentation accuracy and is unlikely to perform well with a single mean threshold on larger, more complex datasets. Overall, the outcomes suggest that while U-Net and the Hamiltonian method offer the highest accuracy, Q-Seg provides a balanced alternative with moderate performance and no training requirements.

Future work could focus on testing these approaches on larger datasets to assess their effectiveness in realistic scenarios. Optimizing the Hamiltonian method for GPU parallel processing could yield a 20x speedup [21] and exploring quantum simulations using ultra-cold gas setups [22, 23]. Additionally, exploring Q-Seg on gate-based quantum computing [24] and identifying other challenging domains for annotated data can enhance the application of quantum methods.

## Acknowledgements

We'd like to thank the Quantum Initiative Rhineland-Palatinate (QUIP) for their support. This work was also partially funded by the Research Initiative 'Quantum Computing for Artificial Intelligence' (QC-AI) and the Federal Ministry for Economic Affairs and Climate Action through the EniQmA project (funding number 01MQ22007A).

## References

1. T. Presles, C. Enderli, G. Burel, and E. H. Baghious, "Synthetic aperture radar image segmentation with quantum annealing," *IET Radar, Sonar & Navigation*, 2024.
2. S. M. Venkatesh, A. Macaluso, M. Nuske, M. Klusch, and A. Dengel, "Q-seg: Quantum annealing-based unsupervised image segmentation." *IEEE Computer Graphics and Applications*, 2024.
3. A. Geng, A. Moghiseh, C. Redenbach, and K. Schladitz, *arXiv:2307.16723*, 2023.
4. N. Otsu *et al.*, *Automatica*, vol. 11, no. 285-296, pp. 23–27, 1975.
5. B. Singh, S. Majumdar, and S. Indu, *Quantum Studies: Mathematics and Foundations*, vol. 11, no. 3, pp. 427–458, Oct 2024.
6. M. Moore and A. Narayanan, "Quantum-inspired computing," *Dept. Comput. Sci., Univ. Exeter, Exeter, UK*, 1995.

7. L. Huynh, J. Hong, A. Mian, H. Suzuki, Y. Wu, and S. Camtepe, *arXiv:2308.11269*, 2023.
8. I. Peschel, *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2004, no. 06, p. P06004, jun 2004.
9. N. Palaiodimopoulos, V. F. Rey, M. Tschöpe, C. Jörg, P. Lukowicz, and M. Kiefer-Emmanouilidis, in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 116–120.
10. S. Dutta, A. Basarab, B. Georgeot, and D. Kouamé, *IEEE Open Journal of Signal Processing*, vol. 2, pp. 190–206, 2021.
11. E. Abrahams, *50 years of Anderson Localization*. world scientific, 2010, vol. 24.
12. P. W. Anderson, *Phys. Rev.*, vol. 109, pp. 1492–1505, Mar 1958.
13. N. Palaiodimopoulos, J. Frkatovic, V. F. Rey, M. Tschöpe, S. Suh, P. Lukowicz, and M. Kiefer-Emmanouilidis, *arXiv:2409.16180*, 2024.
14. A. Mafi and J. Ballato, “Review of a decade of research on disordered anderson localizing optical fibers,” *Frontiers in Physics*, vol. 9, 2021.
15. N. E. Palaiodimopoulos, M. Kiefer-Emmanouilidis, G. Kurizki, and D. Petrosyan, *SciPost Phys. Core*, vol. 6, p. 017, 2023.
16. A. Das and B. K. Chakrabarti, *Rev. Mod. Phys.*, vol. 80, pp. 1061–1081, Sep 2008.
17. E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, *arXiv:0001106*, 2000.
18. H. Neven, G. Rose, and W. G. Macready, *arXiv:0804.4457*, 2008.
19. N. Dattani, S. Szalay, and N. Chancellor, *arXiv:1901.07636*, 2019.
20. O. Ronneberger, P. Fischer, and T. Brox, in *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*. Springer, 2015, pp. 234–241.
21. R. Okuta, Y. Unno, D. Nishino, S. Hido, and C. Loomis, in *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)*, 2017.
22. S. Barbosa, M. Kiefer-Emmanouilidis, F. Lang, J. Koch, and A. Widera, *Phys. Rev. Res.*, vol. 6, p. 033039, Jul 2024.
23. W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner, *Nature*, vol. 462, no. 7269, pp. 74–77, 2009.
24. S. M. Venkatesh, A. Macaluso, M. Nuske, M. Klusch, and A. Dengel, *arXiv:2405.14405*, 2024.