Semester Project

Enhancing Image Resolution through Optical System Deconvolution Using Neural Networks

Author
Jonas Schlör

Supervised by
Rasmus Ischebeck
Andreas Adelmann

December 22, 2025

Abstract

High-resolution optical diagnostics are critical for characterizing electron beam profiles in X-ray free electron laser (XFEL) facilities. However, diffraction limits and lens imperfections degrade image quality through blurring, which is quantitatively described by the system's point spread function (PSF). This study applies a neural-network-based image deconvolution method — the Richardson–Lucy Network (RLN) — to improve the spatial resolution of scintillator screen images acquired at the SwissFEL electron dump.

The PSF of the optical system was experimentally determined and used to generate synthetic training data. The RLN model was then applied to real beam images recorded simultaneously with two CMOS cameras, one of which was a high-resolution test sensor. Preprocessing included normalization and removal of radiation-induced sensor artifacts.

Quantitative analysis demonstrated significant resolution improvements. The transverse size of the electron bunch was reduced by 8.63 %, with a strong horizontal sharpening effect (31.29 %). In another beam-based screen resolution study, deconvolution improved resolution by up to 13.86 % depending on the optical filter configuration. These findings validate the RLN as a practical and effective tool for image enhancement in accelerator diagnostics.

Limitations include sensitivity to PSF mismatch and a lack of built-in uncertainty quantification. Future work will target deployment on upgraded high-resolution sensors and integration into SwissFEL's online diagnostics pipeline.

Contents

1 Introduction

Optical imaging systems are fundamentally constrained by both imperfections in their components and inherent physical limitations, particularly those arising from diffraction. Aberrations and diffraction effects introduce distortions and blurring, which collectively degrade the quality of the resulting image. These effects limit the achievable resolution of any imaging system, meaning that the true, undistorted image, or the truth of the ground, cannot be perfectly reconstructed.

To quantitatively describe these limitations, the point spread function (PSF) is used. The PSF characterizes how an optical system transfers spatial frequencies from the object to the image. Mathematically, the observed image can be modeled as a convolution of the true object with the PSF of the system. This function encapsulates all distortions introduced by the imaging process in a single matrix representation. Importantly, while the image is degraded, the information is not necessarily lost and, under certain conditions, can be recovered through a process known as deconvolution.

Deconvolution can be performed using classical numerical methods, such as the iterative Richardson-Lucy algorithm. More recently, neural network-based methods have been developed to improve classical deconvolution. One such method is the Richardson-Lucy Network (RLN), a machine learning model that uses algorithm unwrapping with the numerical approach. This neural network model learns the point spread function and is capable of producing higher-quality reconstructions compared to traditional techniques.

The Richardson-Lucy Network was previously implemented and evaluated using laboratory images in a bachelor's project by Guney Tekin. The present work focuses on adapting and applying this method to real-world accelerator data. Specifically, it involves imaging a scintillating crystal located at the electron dump of the SwissFEL, following the passage of electrons through the undulators. This application aims to demonstrate the method's effectiveness in practical, high-energy physics environments.

2 Theory

2.1 Point Spread Function

The point spread function (PSF), also referred to as the spatial impulse response function, is a fundamental descriptor of an optical imaging system. It characterizes the system's response to a point-like source of light and is an inherent property determined by the system's optical configuration, including lens imperfections and diffraction effects. In essence, the PSF describes how a single point in the object space is mapped to a distribution of light intensity in the image space due to these system imperfections. Mathematically, the image formation process can be modeled as a convolution between the true object and the PSF:

$$\text{Image}(x, y) = \text{Object}(x, y) \bigotimes PSF(x, y) \tag{1}$$

If the object consists of a single, ideal point source, the image recorded directly corresponds to the PSF itself.
This property allows experimental determination of the PSF by imaging a known point source. In practical setups, this can be achieved using a coherent light source such as a laser, passed through a sub-wavelength pinhole aperture to approximate a point emitter. The resulting intensity distribution on the image plane, recorded on a screen or sensor, provides a direct measurement of the PSF. This measured PSF can then be used to characterize the resolution limits of the system and serve as the basis for computational deconvolution methods aimed at reconstructing the original object.

2.2 Richardson-Lucy Deconvolution

Richardson–Lucy deconvolution is an iterative method to recover an unknown image $\phi(y)$ from a measured image $\tilde{\varphi}(x)$ that has been blurred by a known point spread function (PSF) [1][2]. The image formation model is given by:

$$\tilde{\varphi}(x) = \int \phi(y) P(x|y)\, dy \tag{2}$$

where $P(x|y)$ is the conditional probability (i.e., the PSF), representing the probability that a photon emitted at position y is detected at position x. For a shift-invariant system:

$$P(x|y) = \text{PSF}(x - y) \tag{3}$$

Lucy introduces the inverse probability $Q(y|x)$, which satisfies:

$$Q(y|x) = \frac{\phi(y) P(x|y)}{\int \phi(y') P(x|y')\, dy'} \tag{4}$$

By Bayes' theorem and assuming Poisson noise statistics, the iterative update rule for estimating $\phi(y)$ is given by:

$$\phi^{(i+1)}(y) = \phi^{(i)}(y) \int \frac{\tilde{\varphi}(x)}{\varphi^{(i)}(x)} P(x|y)\, dx \tag{5}$$

where

$$\varphi^{(i)}(x) = \int \phi^{(i)}(y')P(x|y')\,dy' \tag{6}$$

Here, $\tilde{\varphi}(x)$ is the measured (blurred) image, $\phi^{(0)}(y)$ is the initial estimate of the true image, and $\phi^{(i)}(y)$ is updated at each iteration i.

This scheme converges under mild conditions to the maximum likelihood (ML) solution for Poisson-distributed measurement noise.

Algorithm 1 Richardson–Lucy Deconvolution for Imaging [2]

```
Require: Input image $I$, point spread function $P$
Ensure: Deconvolved estimate $\hat{E}^{L+1}$
Initialize estimate: $\hat{E}^{(0)} \leftarrow I$
Forward PSF: $f \leftarrow P$
Backward PSF: $b \leftarrow P^T$
for $k = 0$ to $L$ do
    $\hat{E}^{(k+1)} \leftarrow \hat{E}^{(k)} \cdot \left(\frac{I}{\hat{E}^{(k)} * f} * b\right)$
end for
return $\hat{E}^{(L+1)}$
```

2.3 Richardson Lucy Network

2.3.1 Algorithm Unwrapping

Algorithm unwrapping, also known as algorithm unrolling, is a technique that bridges classical iterative algorithms and neural networks by embedding the structure of a known algorithm directly into the architecture of a machine learning model. Instead of learning a function from scratch, the model mimics the step-by-step logic of an established algorithm, such as Richardson–Lucy deconvolution, while using learnable parameters at each iteration. This approach improves interpretability,since each operation corresponds to a meaningful step in the original algorithm, and typically requires fewer parameters than fully generic neural networks. These advantages make unrolling particularly well-suited for scientific and inverse problems, such as image deblurring or reconstruction, where the underlying physics is well understood. [2]

2.3.2 The Network

The Richardson–Lucy Network (RLN), proposed by Li et al. [3], is a neural network model based on the concept of algorithm unrolling. It leverages the iterative nature of the classical Richardson–Lucy Deconvolution (RLD) and embeds it into a trainable, end-to-end differentiable network architecture. This approach not only improves interpretability but also significantly reduces the number of parameters compared to fully generic deep learning models.

The core idea of RLN is to mimic the iterations of RLD:

$$E^{(k+1)} = E^{(k)} \cdot \left(\frac{I}{E^{(k)} \otimes f} \otimes b \right) \tag{7}$$

where I is the blurred input image, f and b are forward and backward kernels (approximating the PSF and its transpose), and $*$ denotes convolution. These kernels become trainable parameters, and the iteration is embedded into a neural network structure.

2.3.3 Network Structure

Li et al. enhance the basic unrolled model by introducing three conceptual stages: **Estimation**, **Update**, and **Output**.

- **Estimation:** Corresponds to the first RLD iteration. It processes a downsampled version of the input image through several 4-channel convolution layers with batch normalization and `softplus` activation:

  $$\texttt{softplus}(x) = \log(1 + e^x)$$

  It also uses a transposed convolution layer to upsample the intermediate result back to the original resolution.
- **Update:** Corresponds to the second RLD iteration, using the full-resolution image and fewer convolutional layers for computational efficiency. The output of the Estimation block is used as input, alongside the original image.
- **Output:** This stage deviates from the classical RLD scheme. It fuses outputs from Estimation and Update through several 8-channel convolutions to produce the final restored image.

Loss Function and Training

The RLN is trained in a supervised manner using pairs of blurred input images and corresponding ground truth images. Li et al. use a hybrid loss function combining the Mean Squared Error (MSE), the Structural Similarity Index (SSIM), and a regularization term:

$$L(I, GT, O, E) = \text{MSE}(GT, O) - \ln((1 + \text{SSIM}(GT, O))^2) + \text{MSE}(E, 0.8 \cdot GT + 0.2 \cdot I) \quad (8)$$

Modifications and Practical Considerations

Refinements by Tekin [2] replaced SSIM with the Multi-Scale SSIM (MSSIM) to better capture features across various spatial scales, resulting in a modified loss:

$$L(I, GT, O, E) = \text{MSE}(GT, O) - \ln((1 + \text{MSSIM}(GT, O))^2) + 0.1 \cdot \text{MSE}(E, GT) \quad (9)$$

To stabilize inference and improve generalization, the input images are normalized to have zero mean and unit variance:

$$x' = \frac{x - \mu_x}{\sigma_x} \quad (10)$$

$$y' = \text{RLN}(x') \quad (11)$$

$$y = y' \cdot \sigma_x + \mu_x \quad (12)$$

Furthermore, the output is clipped to the $[0, 1]$ range to match the expected intensity domain and reduce artifacts in the prediction.

This structured approach allows the RLN to learn interpretable, physically grounded image reconstructions with better generalization and fewer parameters compared to traditional convolutional networks.

2.4 The SwissFEL Accelerator

The information presented in this paragraph is based on the comprehensive overview of the SwissFEL X-ray Free Electron Laser (XFEL) facility by Milne et al. (2017) [4]. Figure 1 shows a schematic representation of the beamline.

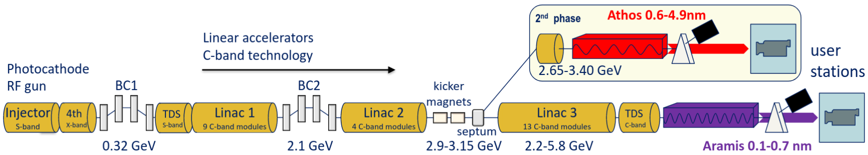

Figure 1: Schematic layout of SwissFEL. [4]

SwissFEL, located at the Paul Scherrer Institute, is a state-of-the-art X-ray free electron laser designed to generate spatially coherent, ultra-short X-ray pulses with extremely high peak brightness. It delivers photons in the 1–7 Å (1.77–12.4 keV) range with pulse durations of 0.2–20 fs (rms) and energies up to 1 mJ. The system is designed for high repetition rates (100 Hz) and can support advanced experimental techniques such as time-resolved spectroscopy and serial femtosecond crystallography.

2.4.1 Beam Dynamics and Free Electron Laser Principle

The fundamental goal was to achieve lasing at a wavelength of 1 Å using the lowest possible electron energy, thereby reducing overall size and cost. This is accomplished using an in-vacuum undulator with a short period $\lambda_u = 15$ mm and a magnetic field corresponding to $K = 1.2$. The FEL resonance condition is given by:

$$\lambda = \frac{\lambda_u}{2\gamma^2}\left(1 + \frac{K^2}{2}\right), \tag{13}$$

where γ is the Lorentz factor of the electron beam. To satisfy this condition at 1 Å, a beam energy up to 5.8 GeV is required. The normalized emittance is targeted at $\varepsilon_n = 430$ nm, and the peak current reaches up to 3 kA. The energy spread is approximately 350 keV (rms). SwissFEL operates flexibly across a wide range of bunch charges (10–200 pC) and supports special modes such as high-chirp or sub-femtosecond compression.

2.4.2 Injector and Linear Accelerator

The electron beam originates from a 2.5-cell S-band RF photoinjector gun, driven by a highly stable Yb:CaF$_2$ laser. Initial acceleration to 320 MeV is achieved using S-band linac modules, with further compression and emittance preservation ensured by a laser heater and X-band linearization cavities. The main acceleration to up to 5.8 GeV occurs in three C-band linac segments, composed of 26 individual modules.

2.4.3 Undulator System (Aramis Line)

The Aramis beamline features 13 in-vacuum undulator modules, each 4 m long, totaling an active length of 48 m. The undulators use NdFeB permanent magnets and are optimized for precision with a peak field of $B = 1.3$ T. The magnetic design supports a minimum gap of 3 mm and gap reproducibility better than 1 μm. The full undulator line is 60 m long, including intersection modules.

2.4.4 X-ray Optics and Photon Parameters

The X-ray beam produced has wavelengths from 1–7 Å, with photon energies between 1.77–12.4 keV. The pulse energy ranges from 0.01–1 mJ, with durations down to 0.2 fs (rms). SwissFEL incorporates sophisticated X-ray optics, including double crystal monochromators and Kirkpatrick-Baez focusing mirrors, optimized for both monochromatic and pink beam operation.

3 Methods

3.1 Setup

The imaging setup is designed to detect light emitted from a scintillating crystal when struck by the high-energy electron beam at SwissFEL. A 100 µm thick YAG:Ce (Yttrium Aluminum Garnet doped with Cerium) scintillator is used as the active material. Upon interaction with the electron beam, the scintillator emits visible light uniformly throughout its volume, producing a bright, extended light source. This light is then imaged onto a CMOS camera sensor using a macro lens with a focal length of $f = 200$ mm.

Figure 2 illustrates the optical layout. The electron beam (orange path) strikes the scintillator at an incident angle of 8.1°. The resulting scintillation light (green path) exits the crystal and is directed through the imaging system. Coherent optical transition radiation (COTR, shown in red) produced at the vacuum-crystal interface is deflected away from the imaging axis using a mirror, preventing unwanted signal contamination.

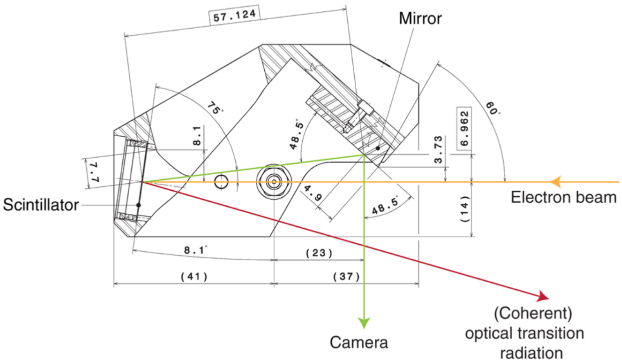

Figure 2: Optical setup. The orange line represents the electron beam, hitting the scintillator at an angle of 8.1°. The red line indicates coherent optical transition radiation, while the green line shows the desired scintillation light path. Adapted from [5].

The imaging geometry fulfills three main functions simultaneously:

1. **Efficient Imaging of the Scintillator:** The emission is refracted at the crystal surface according to Snell–Descartes law, ensuring proper propagation into the imaging system.
2. **Distortion-Free Focus via the Scheimpflug Condition:** Since the object plane (scintillator) is tilted with respect to the optical axis, the lens and sensor are aligned using the Scheimpflug principle. The criterion ensures that the entire tilted object plane is in sharp focus. The required angular alignment is given by:

$$\tan\phi = \left(1 + \frac{f}{s}\right)\tan\theta, \tag{14}$$

   where ϕ is the tilt angle of the object (here 8.1°), f the focal length of the lens, s the distance from lens to sensor, and θ the lens tilt angle (here 15°). This prevents distortion and focus variation across the image.
3. **Suppression of COTR:** Optical transition radiation, which is emitted when charged particles cross interfaces of differing refractive index, is strongly directional and can dominate over scintillation light. The setup deflects this radiation away from the optical axis using a mirror, avoiding interference with the scintillator signal [5].

Two CMOS cameras were used in the setup, both equipped with identical 200 mm macro lenses. A 50/50 beam splitter was used to send the same optical signal to both detectors simultaneously. One camera served as the legacy system (7.9 µm pixels), while the second was a new, high-resolution device with much smaller pixels (1.85 µm), characterized by a smaller field of view and significantly smaller pixels. Although this reduces the total imaged area, the higher resolution is crucial for directly resolving and analyzing the point spread function (PSF) of the lens system.

3.2 Modelling and Model Integration

The Richardson–Lucy Network (RLN) model used in this work builds on the implementation developed by Güney Tekin in her Bachelor's thesis, adapted to run with `Apptainer` containers and TensorFlow 2. The core approach is data-driven image deblurring, with training images generated synthetically based on known PSFs [2]. To ensure reproducibility, the full source code and explanatory README are provided in the RLN-at-SwissFEL repository [6].

The training image generation pipeline was preserved. It creates artificial images containing geometric shapes (e.g., lines, circles, edges), which are then convolved with the measured PSF to simulate the effect of optical blur. The PSF used in this study was obtained experimentally in the lab using the same 200 mm lens and represents the optical characteristics of the imaging system.

The blurring contribution of the scintillator crystal itself was not included in the model, based on prior studies indicating that its spatial blurring effect is below 1 μm, significantly less than a single camera pixel [7, 8].

For dataset acquisition, synchronized images were collected simultaneously from both CMOS cameras (DSCR050 and DSCR051) using locally merged data streams, tagged with common bunch IDs for later pairing and comparison.

Before feeding the images to the neural network, a custom preprocessing routine was applied. Raw 16-bit integer images were converted to 32-bit floats and normalized to the range $[0, 1]$. Damaged sensor regions, identifiable by abnormally bright pixels resulting from radiation exposure, were masked out to prevent artifacts during training and evaluation.

This modeling setup enables the network to learn a realistic mapping between degraded and deblurred images using physically consistent image degradations, with minimal synthetic assumptions.

4 Results

4.1 Image Enhancement

4.1.1 Single Pixel Artifacts

Radiation damage from scattered X-ray photons can lead to individual pixels on the camera sensor becoming persistently ultra-bright. These so-called "hot pixels" introduce significant artifacts in the network output, particularly during convolution operations that amplify localized features. Figure 3 illustrates the impact of such bright single-pixel peaks on the reconstructed image quality.

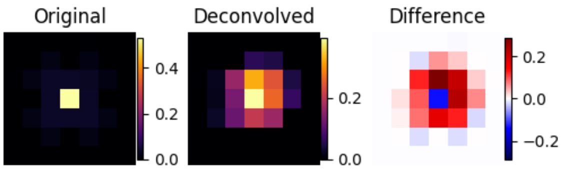

Figure 3: Effect of ultra-bright single-pixel artifacts on the model output. The presence of high-intensity pixels introduces artificial structures in the deblurred image.

4.1.2 Direct Output Comparison

A qualitative comparison of model outputs from the two camera systems is shown in Figure 4. Despite both cameras capturing the same scintillation event, their outputs differ significantly due to differences in optical quality and focus. An observable effect is contrast amplification: brighter regions become more intense, while dim regions are further suppressed. This is a known side-effect of Richardson–Lucy-like deconvolution and occurs because the model attempts to redistribute intensity to maximize sharpness. Background subtraction was applied uniformly across input and output images and dynamic contrast enhancement was mitigated with by a post-processing normalization step.

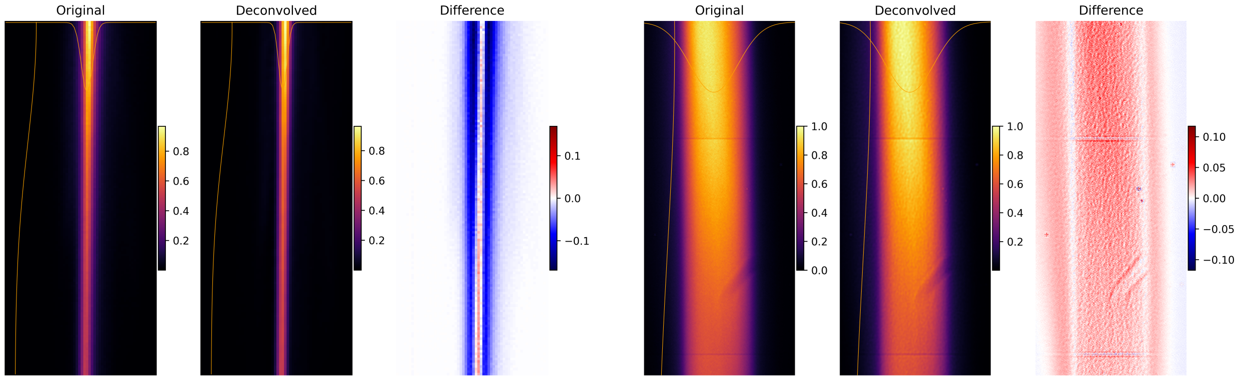

(a) DSCR050 camera

(b) DSCR051 camera

Figure 4: Comparison of RLN model outputs for both cameras. DSCR050 (left) demonstrates a significant sharpening effect while DSCR051 (right) shows reduced enhancement due to blur in the original image.

The DSCR051 camera (right) exhibits a generally softer image with visible blur. This degradation was attributed to imperfect optical alignment, which could not be corrected during the experiment due to restricted tunnel access. In contrast, the DSCR050 camera (left) captured

significantly sharper images and was therefore more suitable for model evaluation.
Interestingly, the pixel-wise difference between input and output images reveals opposite behavior for the two cameras. In the case of the sharper DSCR051 camera, the deconvoluted output often shows lower intensity values than the original, especially around the beam center, as indicated by the blue regions. This corresponds to successful suppression of halo-like blur from the point spread function (PSF), as the model redistributes intensity to sharpen the image. In contrast, the difference image for the blurrier DSCR050 camera exhibits mostly positive values (red), meaning the network has amplified certain features instead of deblurring them. This likely reflects the model's difficulty in handling unexpected blur from optical misalignment, which differs significantly from the PSF it was trained on. As a result, the RLN attempts compensation in a way that may lead to non-physical intensity shifts or artifacts.
Since the DSCR051 system was intended as a temporary test setup and a new high-resolution camera will replace it in the future, this limitation is noted but not critical to the overall findings. Consequently, most of the quantitative analysis focuses on results obtained from the DSCR050 camera.

To quantify the improvements achieved through deconvolution, the transverse sizes of the electron bunches were measured before and after processing. This was done by fitting Gaussian functions to the horizontal and vertical projections of each image and extracting the standard deviations σ_x and σ_y. The effective bunch size was then calculated as the geometric combination $\sigma = \sqrt{\sigma_x + \sigma_y}$.

The measurements were averaged over 42 image pairs. Table 1 summarizes the results. Notably, the deconvolution procedure reduced the effective bunch size by 8.63%, with a significant improvement in the horizontal direction (31.29%), while the vertical size remained almost unchanged.

Table 1: Average beam sizes before and after deconvolution, averaged over 42 images.

Parameter	**Input** [μm]	**Output** [μm]	**Improvement** [%]
Effective size σ	171.97 ± 0.21	157.12 ± 0.19	8.63
σ_x (horizontal)	39.48 ± 0.06	27.13 ± 0.04	31.29
σ_y (vertical)	723.16 ± 1.38	712.67 ± 1.35	1.45

4.2 Screen Resolution Measurement

To quantitatively assess screen resolution, we replicate a beam-based measurement from Juranić and Prat [9]. The method involves measuring the beam size at different beam energies while keeping the normalized emittance ϵ_n constant. As the energy γ changes, the expected beam size on the screen σ_{tot} follows:

$$\sigma_{\text{tot}}^2 = \sigma_{\text{scr}}^2 + \frac{\beta\epsilon_n}{\gamma}, \tag{15}$$

where σ_{scr} is the intrinsic resolution of the scintillator screen and optics, β is the beta function at the screen, and γ is the relativistic Lorentz factor. By fitting this equation to the energy-dependent measurements, one can extract both the resolution σ_{scr} and the transverse beam size $\beta\epsilon_n$. In the original paper, the resolution of the screen was determined to be 14.3 µm. Thus, the model was trained with a gaussian PSF with corresponding width.

The size of the beams were determined by fittings Gaussian peaks to the horizontal projections. The pixel size of the used camera was 6.5 µm.

A direct visual comparison of raw and deconvoluted images—including fitted beam profiles—is shown in figure 5.

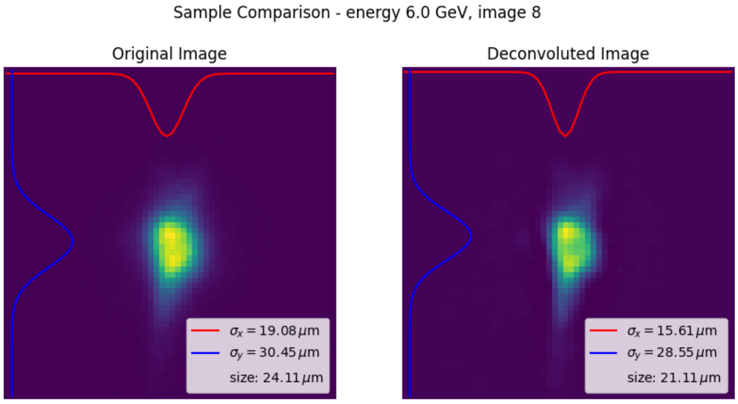

Figure 5: Comparison of raw and deconvoluted image including fitted beam sizes. Deconvolution reduces the apparent beam width from 24.11 μm to 21.11 μm (12.44% improvement). Image adapted from [9].

Figure 6 shows the screen resolution measurement described above. Although the exact numeric results from the original study could not be perfectly reproduced, the qualitative trends were replicated: (i) the new filter improves resolution over the old filter, and (ii) deconvolution introduces a measurable sharpening effect.

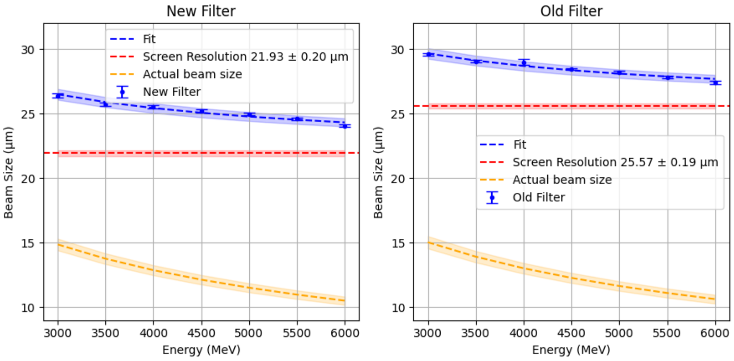

(a) Original data

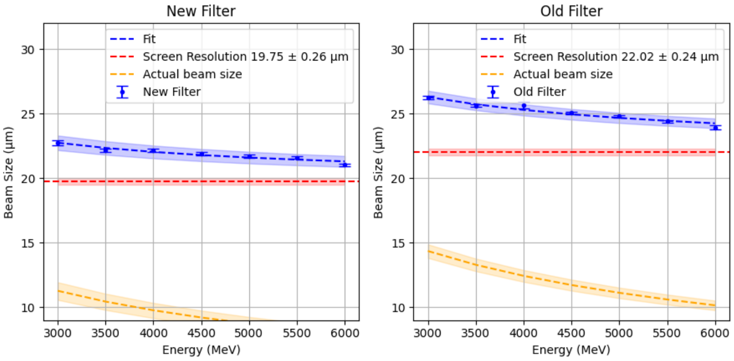

(b) Enhanced (deconvoluted) data

Figure 6: Comparison of beam size fits with and without deconvolution, based on the method in Juranić and Prat [9]. Fit model includes beam optics and screen resolution.

Extracted Parameters

The extracted resolution σ_{scr} and geometric product $\beta\epsilon_n$ for each configuration are summarized in Table 2. In both filter configurations, the screen resolution improves after deconvolution, with relative improvements of 9.94% (new filter) and 11.86% (old filter).

The geometric term $\beta\epsilon_n$, which reflects the intrinsic beam properties and should remain unaffected by post-processing, appears increased in the deconvoluted data. This is likely not a physical effect, but rather an artifact of the nonlinear least-squares fit used to extract the parameters. Since σ_{scr} and $\beta\epsilon_n$ are coupled in the fit model, changes in one can artificially shift the other to maintain agreement with the observed beam size. The result suggests a compensation effect due to parameter correlation. A more robust error analysis or constrained fitting strategy could be employed in future work to reduce this ambiguity.

Table 2: Comparison of extracted screen resolution σ_{scr} and geometric emittance-product $\beta\epsilon_n$ before and after deconvolution.

Filter Type	**Raw**	**Deconvoluted**	**Improvement [%]**
Screen Resolution $\boldsymbol{\sigma_{\text{scr}}}$ $[\boldsymbol{\mu m}]$			
New	21.93 ± 0.20	19.75 ± 0.26	9.94
Old	25.57 ± 0.19	22.02 ± 0.24	13.86
Geometric Term $\boldsymbol{\beta\epsilon_n}$ $[\boldsymbol{m^2}]$			
New	$(1.30 \pm 0.08) \times 10^{-6}$	$(0.75 \pm 0.09) \times 10^{-6}$	42.31
Old	$(1.33 \pm 0.08) \times 10^{-6}$	$(1.21 \pm 0.09) \times 10^{-6}$	9.77

5 Conclusion

The aim of this project was to improve the optical resolution of beam profile measurements at SwissFEL by applying neural-network-based deconvolution using the Richardson–Lucy Network (RLN). This entailed four primary objectives: (1) measuring the point spread function (PSF) of the optical system, (2) acquiring real accelerator data from the SwissFEL beamline, (3) applying the RLN to deconvolve the images, and (4) quantifying the improvement in spatial resolution.

Using experimental data collected from two CMOS cameras — one standard-resolution and one high-resolution — the PSF of the optical setup was determined and used to generate synthetic training data for the RLN. The model was successfully integrated and deployed to process real-world data, with dedicated preprocessing steps to normalize intensity values and remove sensor artifacts caused by radiation damage.

Quantitative analysis confirmed a substantial improvement in image sharpness after deconvolution. For electron bunch images, the effective bunch size was reduced by 8.63 %, with a strong improvement in the horizontal resolution (31.29 %) and minimal effect on the vertical axis (1.45 %), consistent with the expected streaked bunch geometry. These results demonstrate that the RLN is capable of resolving fine spatial structures that are otherwise blurred in raw images.

Additionally, in a beam-based screen resolution study inspired by Juranić and Prat [9], the extracted screen resolution values decreased from 21.93 µm to 19.78 µm for the new filter and from 25.57 µm to 22.82 µm for the old filter, confirming the efficacy of the deblurring model.

Limitations. Despite the promising results, several limitations must be acknowledged. First, the RLN is trained on synthetically generated data using a single measured PSF, which does not account for potential variation in the actual optical blur across the field of view or changes over time. In particular, blur arising from misalignment (as seen in the DSCR051 camera) could not be compensated effectively, as it deviated significantly from the training PSF. Second, while the scintillator was assumed to contribute negligible blurring based on literature values, any unexpected degradation from radiation damage or manufacturing tolerances would not be captured in the model.

Outlook. Future work should focus on deploying a high-resolution camera with a larger field of view to fully exploit the model's capabilities. Moreover, integrating RLN-based deconvolution into the standard SwissFEL data processing pipeline would enable routine high-precision beam size measurements. Retraining the model with PSFs measured across the full sensor area and incorporating synthetic examples of coherent optical transition radiation (COTR) and scintillator blurring would improve robustness.

References

[1] L. B. Lucy. “An iterative technique for the rectification of observed distributions”. In: *The Astronomical Journal* 79 (June 1974). Publisher: IOP ADS Bibcode: 1974AJ.....79..745L, p. 745. ISSN: 0004-6256. DOI: 10.1086/111605. URL: https://ui.adsabs.harvard.edu/abs/1974AJ.....79..745L (visited on 02/23/2025).

[2] Güney Tekin. *Deblurring of Images with the Richardson-Lucy-Network*. Bachelor Project. Zürich, Switzerland, June 24, 2024.

[3] Yue Li et al. “Incorporating the image formation process into deep learning improves network performance”. en. In: *Nature Methods* 19.11 (Nov. 2022). Publisher: Nature Publishing Group, pp. 1427–1437. ISSN: 1548-7105. DOI: 10.1038/s41592-022-01652-7. URL: https://www.nature.com/articles/s41592-022-01652-7 (visited on 02/23/2025).

[4] *SwissFEL: The Swiss X-ray Free Electron Laser*. URL: https://www.mdpi.com/2076-3417/7/7/720 (visited on 02/23/2025).

[5] Rasmus Ischebeck et al. “Transverse profile imager for ultrabright electron beams”. In: *Physical Review Special Topics - Accelerators and Beams* 18.8 (Aug. 2015). Publisher: American Physical Society, p. 082802. DOI: 10.1103/PhysRevSTAB.18.082802. URL: https://link.aps.org/doi/10.1103/PhysRevSTAB.18.082802 (visited on 03/21/2025).

[6] Jonas Schlöe. *RLN-at-SwissFEL*. https://gitea.psi.ch/schloe_j/RLN-at-SwissFEL. 2025.

[7] Andreas Koch et al. “X-ray imaging with submicrometer resolution employing transparent luminescent screens”. In: *Journal of the Optical Society of America A* 15.7 (July 1998), pp. 1940–1951. DOI: 10.1364/JOSAA.15.001940.

[8] A. R. Najafiyan et al. “BEAM PROFILE MONITOR DESIGN FOR A MULTIPURPOSE BEAM DIAGNOSTICS SYSTEM”. In: *Proceedings of the 10th International Beam Instrumentation Conference (IBIC’21)*. Pohang, Rep. of Korea: JACOW Publishing, 2021, pp. 218–222. ISBN: 978-3-95450-230-1. DOI: 10.18429/JACOW-IBIC2021-TUPP13. URL: https://doi.org/10.18429/JACOW-IBIC2021-TUPP13.

[9] Pavle Juranić and Eduard Prat. “Testing high-resolution transverse profile monitors by measuring the dependence of the electron beam size on the beam energy at SwissFEL”. In: *Review of Scientific Instruments* 94.7 (July 2023), p. 073301. ISSN: 0034-6748. DOI: 10.1063/5.0155444. URL: https://doi.org/10.1063/5.0155444 (visited on 05/20/2025).