

Enhancing Mass Transfer Coefficient Prediction from Field Emission Scanning Electron Microscope Images Through Convolutional Neural Networks and Data Augmentation

Original

Enhancing Mass Transfer Coefficient Prediction from Field Emission Scanning Electron Microscope Images Through Convolutional Neural Networks and Data Augmentation Techniques / Marcato, Agnese; Boccardo, Gianluca; Pisano, Roberto. - In: PROCESSES. - ISSN 2227-9717. - 13:2(2025). [10.3390/pr13020365]

Availability:

This version is available at: 11583/2999327 since: 2025-04-18T06:51:37Z

Publisher:

MDPI

Published

DOI:10.3390/pr13020365

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Article

Enhancing Mass Transfer Coefficient Prediction from Field Emission Scanning Electron Microscope Images Through Convolutional Neural Networks and Data Augmentation Techniques

Agnese Marcato , Gianluca Boccardo *  and Roberto Pisano 

Department of Applied Science and Technology, Politecnico di Torino, 24 Corso Duca degli Abruzzi, 10129 Torino, Italy; agnese.marcato@polito.it (A.M.); roberto.pisano@polito.it (R.P.)

* Correspondence: gianluca.boccardo@polito.it

Abstract: With the growing demand for drug products requiring lyophilization, it is essential to either expand aseptic drying capacity or improve the efficiency of existing capacity through process intensification, ensuring that resources are utilized to their full potential. In this regard, mathematical models are highly recommended to assist professionals in process optimization. To effectively utilise these models, it is also essential to develop robust techniques for determining key parameters, including the product resistance to vapour flow. Traditional experimental methods for evaluating this coefficient are time-intensive and/or require the insertion of probes into the product, which is not feasible at a manufacturing scale. This study addresses these challenges by introducing a novel deep learning framework designed to predict the mass transfer coefficient directly from Field Emission Scanning Electron Microscope images. This approach significantly streamlines the evaluation process, leveraging the high-resolution capabilities of Field Emission Scanning Electron Microscope for detailed analysis. In this work, we focus on advanced Field Emission Scanning Electron Microscope image processing, choice of strategic convolutional neural network configuration, and thorough model performance evaluation to predict the mass transfer coefficient. Given the frequent scarcity of datasets in this field, we have employed data augmentation techniques to enhance the robustness of our model. The results demonstrate good predictive accuracy (error on the interpolation test data lower than 5%), highlighting the potential of this framework to facilitate the assessment of mass transfer coefficients in freeze-dried products.

Keywords: deep learning; SEM; freeze-drying; convolutional neural networks; pharmaceuticals



Academic Editors: Weizhong Dai and Jie Zhang

Received: 22 August 2024

Revised: 23 December 2024

Accepted: 20 January 2025

Published: 28 January 2025

Citation: Marcato, A.; Boccardo, G.; Pisano, R. Enhancing Mass Transfer Coefficient Prediction from Field Emission Scanning Electron Microscope Images Through Convolutional Neural Networks and Data Augmentation Techniques.

Processes **2025**, *13*, 365. <https://doi.org/10.3390/pr13020365>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Freeze-drying, or lyophilization, is the gold standard for preserving many biopharmaceutical products [1]. This process involves removing water or other solvents under high vacuum and low temperatures, which are ideal for maintaining the native structure of biomolecules and, thus, preserving their bioactivity. The growing reliance on lyophilization is evident from the steady increase in the percentage of approved drug products that require lyophilization over the past decade [2]. However, despite its established effectiveness, lyophilization still faces several challenges. These include the limited use of mathematical modeling for monitoring, controlling, and optimizing the process, as well as the restrictions on product monitoring at the manufacturing scale due to the need to maintain sterility by avoiding the use of probes [3].

By integrating mathematical models with real-time data, manufacturers can enhance process control, reduce development time, and lower costs, ultimately improving product quality and ensuring regulatory compliance. Additionally, modeling facilitates scaling up the process from laboratory to industrial scale, ensuring consistency and efficiency. Simulating the various stages—freezing [4], primary drying [5], and secondary drying [6]—allows modeling to predict the behavior of the material being freeze-dried under different conditions. This predictive capability is crucial for determining optimal temperature and pressure settings, which are essential for maintaining the integrity of sensitive biological substances [7,8].

One significant challenge in applying mathematical models is obtaining accurate estimates of transport properties. While the heat transfer coefficient can be easily measured gravimetrically, determining the product resistance to vapor flow, which varies with each product, is more complex. This resistance is often inferred indirectly using the Product Temperature Response technique, based on the temperature profile during the primary drying process. However, this method is controversial because it requires inserting thermocouples into the product, potentially altering the freezing conditions in the monitored vial and resulting in a product morphology that may not represent the entire production batch. The Pressure Rise Test technique can solve this problem, but it is often affected by systematic errors, as discussed in [9].

To overcome this limitation, efforts have been made to directly assess product resistance by analyzing the morphology of freeze-dried products using Scanning Electron Microscopy (SEM) images [10]. SEM is a powerful imaging technique that uses a focused beam of electrons to generate detailed images of a sample's surface, while Field Emission Scanning Electron Microscopy (FESEM) offers even higher resolution and clarity by using a sharper electron beam. However, SEM image analysis is often complex and time-consuming, but it is a non-invasive form of analysis of the samples. A mathematical model, be it based on first-principles or on a data-driven methodology, would be able to bridge the gap between the necessity of invasive measures (substituted by SEM scans) and the accurate prediction of the relevant transport properties. Therefore, in this study, we aimed to develop an artificial intelligence-based system that can determine mass transport properties directly from SEM images, eliminating the need for a manual operator analysis of the structure. Predicting the mass transfer coefficient is crucial because it governs the rate at which moisture is removed from the product during lyophilization, directly affecting drying efficiency and the quality of the final product. By accurately estimating this parameter, manufacturers can optimize the lyophilization process and ensure the preservation of sensitive biopharmaceuticals. Recent advances in the analysis of mass transfer in dried porous products have utilized micro-CT tomography to provide valuable 3D structural insights [11–14]. While our current approach uses FESEM images, it could be extended to process 2D slices from micro-CT scans or further developed with three-dimensional convolutional layers to handle full 3D tomography images.

When dealing with images as a starting point for the construction of a data-driven model, convolutional neural networks (CNN) are the consolidated tool of choice. CNNs are a class of deep learning models specifically designed to process and analyze visual data by automatically detecting important features, such as edges and textures, through the application of convolutional filters. Freeze-dried products can be seen as porous media (see the visualization in Figure 1, and other published works in the literature [15]), and there are plenty of examples in the literature regarding CNN employed to extract macroscale representative parameters from porous media images [16]. For example, many works employed these networks for the prediction of the effect of the geometric structure on

fluid flow [17,18], but also filtration rates [19,20], electrical properties [21,22], and transport coefficients [23–26].

In this work, we built upon a previous study from the authors of [10], using the same dataset as a starting point. This dataset consists of FESEM images obtained from various formulations of freeze-dried products, generated through a systematic series of freeze-drying experiments. For detailed information about the experimental setup, readers are referred to the work by Grassini et al. [10]. In our previous approach, we used a series of numerical values coming from different FFT wavelet compression steps as input for the model. In the current work, instead, we have changed the paradigm for the construction of the model, employing the aforementioned CNN, which takes the whole image as a direct input. Moreover, we have expanded here upon the exploration of the performance of the model in the sense of its capability for generalization, which is all the more relevant for cases that suffer from a limited number of samples in the dataset (such as the current one).

Thus, in this work, we train a CNN to predict the mass transfer coefficient from the FESEM images of the dried products. In Figure 1, we show the workflow we implemented to achieve this. We processed the images in order to have an optimal input to the CNN; then, we trained it to accurately predict the mass transfer coefficient. The main issue we dealt with in this work was the limited number of images used for the neural network's training and testing, which is a common problem when dealing with real experimental images. To provide an accurate estimate of the accuracy, we present the results of a leave-one-out strategy, which allows us to showcase our network's performance clearly. Also, we exploit data augmentation techniques to maximize the extraction of available information from the dataset.

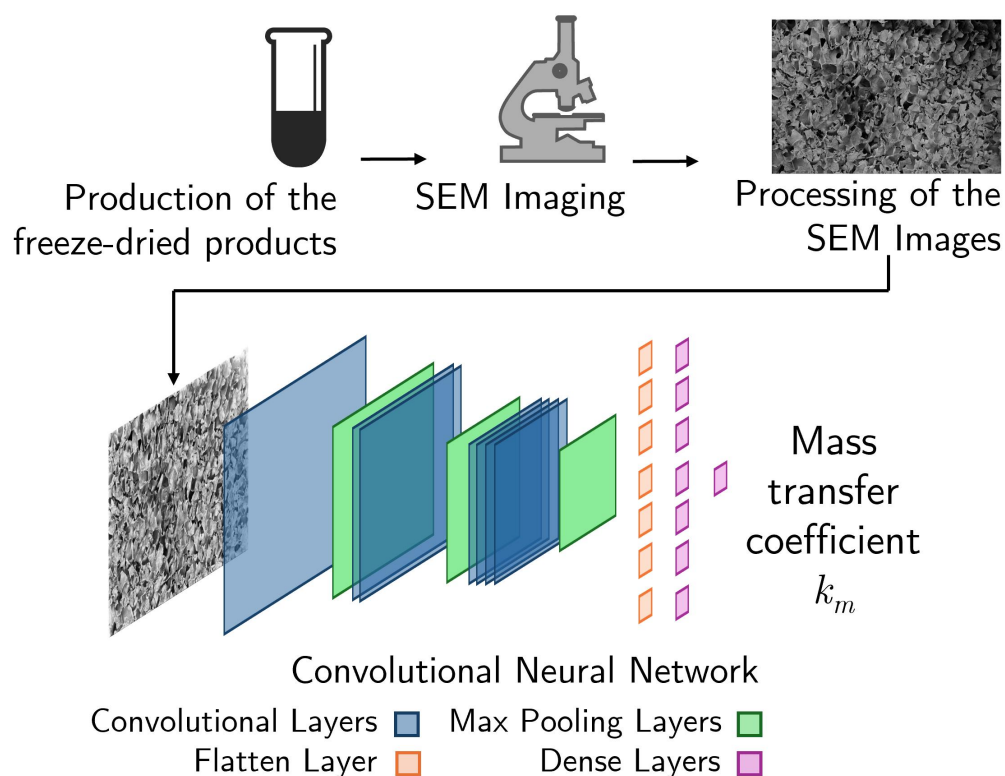


Figure 1. Workflow of this study: freeze-dried products are analyzed through FESEM; the images are processed, and then the CNN is trained to predict the mass transfer coefficient.

2. Dataset

In order to effectively train deep learning algorithms, we need a dataset that is composed of a suitable number of samples. Each sample needs to have two components: the input “feature”, which in our case is the FESEM image, and the output target, which is the value of the mass transfer coefficient associated with that image. The purpose of the trained CNN is to predict the mass transfer coefficient when given a new, unseen FESEM image as an input. This deep learning algorithm is thus a regression model. The most important metric to evaluate in a model of this kind is precisely the generalization performance on new samples not present in the original dataset.

As mentioned in the Introduction, the samples for this dataset were obtained experimentally by performing a number of freeze-drying cycles on different formulations. From these, both the input images and ground truth output targets are obtained, respectively, via FESEM imaging and a mass and heat transfer model based on temperature measurements obtained during the freeze-drying cycle.

The nine formulations employed are detailed in Table 1, representing different concentrations and combinations of dextran, mannitol, and polyvinylpyrrolidone (PVP).

All of the samples were produced by freezing the solutions at $-50\text{ }^{\circ}\text{C}$ for 4 h; the frozen samples were then dried at 10 Pa and then $-10\text{ }^{\circ}\text{C}$ until ice sublimation was complete. Afterward, the temperature was raised to $+20\text{ }^{\circ}\text{C}$ for 5 h to aid water desorption. For more details, the freeze-drying protocol was also described by Grassini et al. [10]. The necessary starting measurements for the evaluation of the mass transfer coefficient are the temperatures of the shelf T_f and the product T_p . The product temperatures are obtained via miniaturized thermocouples inserted into the vials. The estimation of the mass transfer coefficient k_m relies on the estimation of the overall heat transfer coefficient k_v . Then, the water vapour flow rate J_w associated with ice sublimation is estimated as follows:

$$J_w = \frac{k_v(T_f - T_p)}{\Delta H_s}, \quad (1)$$

where ΔH_s is the sublimation enthalpy. From this, the mass transfer coefficient k_m is obtained using this vapour flow estimation, via the following relation:

$$k_m = J_w \frac{RT_p L_d}{M_w} \frac{1}{P(T_p) - P_c}, \quad (2)$$

where M_w is the water molecular weight, $P(T_p)$ is the ice vapor pressure, P_c is the chamber pressure, L_d is the thickness of the freeze-dried layer calculated by integrating the vapor flow rate, and R is the ideal gas constant.

This estimated mass transfer is reported in the same Table 1 for each formulation. The standard deviation value associated with each k_m comes from having obtained a number of different measurements for each formulation, coming from each of the thermocouples inserted in the other vials. These are invasive measurements that inhibit further analysis of the sample, especially image analysis. Thus, some of the vials were left free of the thermocouples, and their product morphology was analyzed by means of FESEM imaging, using a central section of the dried cake.

Furthermore, the structure of the product, and thus its resistance to mass transport, can vary depending on the filling volume and the position of the vial during lyophilization. However, the proposed method is not directly affected by these parameters. If the product’s structure changes due to these factors, the corresponding resistance will also change, and the model will provide a new estimate of the transport parameter based on the updated

morphology. Therefore, as long as the samples can be prepared for SEM analysis without damaging their structure, the model remains applicable in a wide range of scenarios.

A Zeiss SUPRA 40 FESEM is employed with a magnification of $200\times$, resulting in a field of view of $2\text{ mm} \times 1.5\text{ mm}$; with a pixel resolution of about $1\text{ }\mu\text{m}$, the final image resolution, which needs to be the same over all the acquired samples, is 2048×1536 pixels. The acquisition protocol for these images is not fully automatic, as the operator needs to manually tune the contrast and brightness to obtain a reasonably sharp image of the sample structure. For the purposes of aiding reproducibility of the experiments and finding more information on the experimental protocol, the interested reader is referred to [10].

The materials used in these experiments are insulating, so the resulting images are prone to suffer from charge effects, where the electron beam charge may pool on the sample surface due to the poor conductivity, resulting in a darker image area. To reduce these effects, the samples are coated with a 10 nm metallizing chromium layer, which increases surface conductivity. It has to be noted that this effect can only be reduced and not totally avoided. It leaves a noticeable artifact in the final image, impacting image-based analyses. Another effect impacting image quality comes from the heterogeneity of the surface: this leads to only a portion of the image being in focus (usually the center), and the borders are described with a lower definition.

These issues are addressed in the post-processing of the images as detailed in the next section.

Table 1. Dataset of experiments, from [10].

Formulation	Composition	k_m (m^2s^{-1})	Std %
A	5% dextran	1.9×10^{-3}	2.3
B	2.5% dextran + 2.5% mannitol	1.1×10^{-3}	9.3
C	4% dextran + 1% mannitol	1.4×10^{-3}	4.5
D	1% dextran + 4% mannitol	1.0×10^{-3}	4.4
E	2.5% mannitol + 2.5% PVP	9.3×10^{-4}	8.2
F	5% PVP	9.4×10^{-4}	8.9
G	4% mannitol + 1%PVP	9.6×10^{-4}	9.3
H	1% mannitol + 4% PVP	9.3×10^{-4}	3.5
I	10% mannitol	2.8×10^{-4}	8.7

3. Processing of the FESEM Images

The processing of the FESEM images is a crucial step in this workflow since they are the only input data used to predict the mass transfer coefficient. As such, their quality directly impacts the accuracy of the deep learning model. In particular, we need to address the presence of artifacts due to the scanning technique and also to improve the quality of the image. As described in the previous section, the former is due to the charging effect, as can be clearly seen in Figure 2a. The presence of this artifact is detrimental to the neural network accuracy as it alters the perceived morphological structure of the dried product. In contrast, the main issue in terms of image quality comes from low contrast, which impairs the neural network's interpretative capability.




Figure 2. Steps of processing of the FESEM images and relative histogram: (a) original FESEM image; (b) background subtraction output; and (c) histogram equalization output. Original image for formulation C.

The first step in the processing of the images consists of a background subtraction operation. This procedure aims to subtract from the pixel value the average local background intensity, calculated with a “rolling ball” algorithm [27,28]. The radius of this averaging ball was set to be 50 pixels, a value that needs to be chosen to be similar to the characteristic dimension of the largest feature present in the image. The result of this operation can be seen in Figure 2b. Qualitatively, it can be observed that this operation does, in fact, remove the dark shade in the original FESEM image. At the same time, the histogram of the pixels’ color distribution is compressed, skewing towards lighter values.

Thus, to enhance the contrast of the image, we applied a histogram equalization procedure. This process is a non-linear monotonic mapping that results in an image that leverages the entire pixel intensity spectrum. This is achieved by aiming to linearize the cumulative distribution function of the intensities [29]. The result of this operation and the relative stretched histogram can be seen in Figure 2c. These two operations in sequence were applied to all the FESEM images corresponding to the nine formulations of Table 1. In Figure 3, the final output images that are employed for the training of the CNN are shown.




Figure 3. Final processed FESEM images of the nine formulations samples. The labels (A–I) refer to the formulations of Table 1.

4. Deep Learning Framework

4.1. Convolutional Neural Network

As mentioned in this work, we are interested in predicting the values of mass transfer coefficients from images. CNNs represent a suitable approach for this type of regression problem. The literature proposes many different CNNs. The main distinguishing point between these is their architecture, which means both the hyperparameters of the layers and the operations that these layers perform.

The architecture we have chosen in this work is sketched in Figure 4: this network is very similar to another that has already been successfully employed in treating images of porous media [20,30]. As the Figure shows, in this case, we employed three repeating blocks, each consisting of a convolutional layer followed by a max pooling layer. The number of filters in the convolutional layers increases in each block: four in the first one, then eight, and then sixteen. The kernel size of this layer is 3-by-3, which is the most computationally effective for training on graphical processing units (GPU) [31]. We employed rectified linear units (ReLUs) as activation functions. Concerning instead the max pooling layers, we used a 2-by-2 pool size. A flatten layer is then used to change the dimensionality of the output from the two-dimensional data type of the processed image to a one-dimensional array. The subsequent dense layer consists of ten neurons, followed by a final dense layer with a single neuron, which outputs the prediction. Since this is a regression problem, the final layer contains only one neuron to generate the continuous output value. The convolutional neural network was trained using the Adam optimization algorithm with an initial learning rate of 0.0001, a batch size of 16. We employed dropout as a regularization technique to prevent overfitting, in particular setting a dropout rate of 50% [32]. Once the model is trained, it can be used to predict mass transfer coefficients from new FESEM images

almost instantaneously, making the prediction process highly efficient. While the training phase requires approximately 5 h on a GPU, the time complexity for making predictions with new data is minimal (order of seconds), allowing for rapid and scalable application. All the training has been performed on a single Nvidia Quadro GTX 4000 GPU (Nvidia Corporation, Santa Clara, CA, USA).

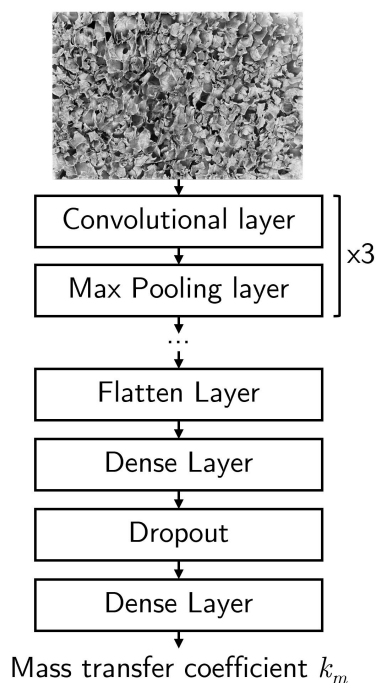


Figure 4. Architecture of the employed CNN.

4.2. Results

Generally, neural networks are trained on a portion of the dataset (the training set) and the remaining part is used to test their performance (the validation set). Since our dataset is comprised of just nine samples, it is hard to make a statistically significant choice on the repartition of the dataset. For this reason, we decided to employ the leave-one-out cross-validation (LOOCV) technique [33]. LOOCV addresses the limitations of the usual repartition strategies by utilizing a comprehensive approach for model validation that maximizes the use of the available data. The procedure for LOOCV involves the following steps:

- *Partitioning*—For a dataset containing N samples, LOOCV iterates N times, with each iteration i leaving out one sample from the dataset and using the remaining $N - 1$ samples for training the model. The left-out sample serves as the test set for that iteration.
- *Training and Evaluation*—In each iteration, the model is trained on $N - 1$ samples and then tested on the single left-out sample. This process yields an immediate measure of performance, which we expressed as an error between the real mass transfer coefficient value and the network prediction. This translates in the definition of the network's loss function as the mean square error between the ground truth (the experimental data) and the values predicted by the network.

The use of the LOOCV technique enables us to specifically demonstrate the generalization performance of the CNN. The usual analysis, which in our case would mean using eight data points for training and only one for validation, returns results that can be seen in Figure 5. The predictions on the training set are reasonably satisfactory, as expected. However, the prediction on the single validation point (in this case, formulation B) is higher,

with a relative percentage error of 12.9%. Nevertheless, this presentation of the algorithm performance, as mentioned, can be misleading as it is heavily influenced by the choice of the validation point: a different choice would result in quite a different evaluation of the model performance. For this reason, it is more appropriate to present an aggregate figure that takes into account all the possible choices.

The overall results of this procedure are found in Figure 6. All the results shown in this Figure are for the validation points: the same point (green color) for formulation B can be seen in both Figures.

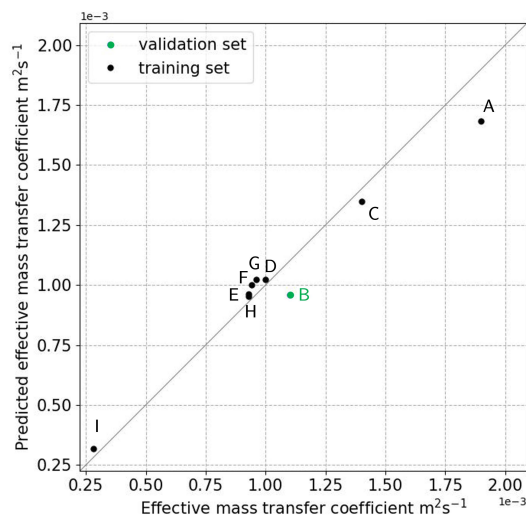


Figure 5. Parity diagram showing the distance between the neural network predictions and the corresponding experimentally obtained mass transfer coefficient. Points closer to the diagonal line mean lower error. The validation point (formulation B) is highlighted in green.




Figure 6. Parity diagram of the predicted vs. true mass transfer coefficient for the LOOCV trainings. All the points in the plot are validation samples.

Most of the validation results are satisfactory; in particular, the prediction errors for formulation B, D, E, F, G, and H are very low (see the first column of Table 2). As expected, the prediction errors for formulations A, C, and I are instead much higher, as they stand as outliers at the very limits of the explored mass transfer coefficients range. This is because when these points are validation points, we are, in fact, asking the network to extrapolate beyond the limits of the training set. When the available training set is limited in size, extrapolation performance worsens.

Table 2. Relative errors for the validation points of the leave-one-out technique: comparison of original dataset and augmented dataset.

	Error, Original Dataset	Error, Augmented Dataset
A	5.99×10^{-1}	5.66×10^{-1}
B	1.29×10^{-1}	7.92×10^{-4}
C	5.36×10^{-1}	2.12×10^{-1}
D	2.2×10^{-3}	1.55×10^{-3}
E	5.66×10^{-3}	9.00×10^{-4}
F	3.78×10^{-5}	1.93×10^{-4}
G	5.80×10^{-4}	4.52×10^{-4}
H	6.92×10^{-2}	4.490×10^{-2}
I	2.67	1.92

5. Data Augmentation

5.1. Dataset Manipulation

This performance issue, stemming from the limited number of data points, can be remediated by enlarging the available dataset: this can be an expensive endeavour. In this work, we have chosen to artificially augment our dataset in order to tackle this issue without incurring additional experimental costs.

The most common way to do this is by both rotating and flipping the original images, as shown in Figure 7. In the Figure, both horizontal and vertical flip operations are shown, together with a 180° rotation. The latter is the only available rotation angle we can use while keeping the same image shape (i.e., the exact dimensions of the array providing the numerical representation of the image to the neural network).

This is effective for CNN models as these models are equivariant to translation but not to the flipping or rotation of the inputs. Hence, images manipulated in this way effectively act as new training data. Naturally, the four (original plus three manipulations) images for each formulation are labeled with the same mass transfer coefficient. In this way, the numerosity of the samples increases four-fold, from 9 to 36.




Figure 7. Magnified FESEM image for formulation C (only a limited portion is shown, for clarity). (Top left): original. (Top right): vertical flip. (Bottom left): horizontal flip. (Bottom right): 180° rotation.

5.2. Results

Having a larger dataset available would usually result in a correspondingly better generalization performance of the trained model. As can be seen in Table 2 and the new version of parity diagram of Figure 8, this is indeed the case. In the Table, we report

the relative errors between the predicted mass transfer coefficients and the experimental ones, for the original dataset in comparison with the augmented dataset. Error metrics generally decrease, showing an increased performance. It has to be noted, however, that extrapolation performances (formulations A and I) are still lacking: the errors decrease, but not in a way that would be satisfactory in actual operation.

The morphology of the pore structure of formulation I is the most distinct among those in the dataset. This is due to the higher solid content (10%) compared to the rest of the formulations (5%). This results in a denser pore structure with a smaller mean pore size: this difference can be clearly seen from Figure 3I. Having a single sample visually distinct from the rest of the dataset means that the neural network will have difficulties in predicting the mass transfer coefficient. Among other formulations with equal solid content (5%), the formulations lower in mannitol will show wider pores and, thus, higher mass transfer coefficients. Higher substitution rates of mannitol with dextran (as in formulations A and C) also increase the interconnectedness of the pores in the matrix, even with similar pore sizes, which is another factor positively impacting the mass transfer coefficient.

Formulation A, in particular, does not contain mannitol in its recipe, which results in the highest mass transfer values. A decreasing mass transfer coefficient with increasing mannitol content was already observed in [8].




Figure 8. Parity diagram of the predicted vs. true mass transfer coefficient for the LOOCV trainings using the artificially augmented dataset for training. All the points in the plot are validation samples.

The results of our study demonstrate that the proposed CNN model outperforms traditional methods for predicting the mass transfer coefficient, particularly in terms of speed and scalability. While classical approaches require labor-intensive measurements, such as inserting thermocouples and performing time-consuming analyses, our image-based model eliminates these steps and allows for more efficient, automated predictions. Moreover, data augmentation techniques were employed to enhance the model's generalization performance, especially when dealing with limited datasets, which is a common challenge in the field. Although the prediction accuracy is satisfactory, especially for formulations with similar morphological characteristics, there are some limitations in cases where the formulations are outliers or significantly different in composition. Future work should focus on expanding the dataset and refining the model for better extrapolation capabilities.

Overall, the integration of deep learning in the prediction of mass transfer coefficients marks a significant advancement over existing experimental approaches.

6. Conclusions

We showed that deep learning techniques, particularly CNNs, can enhance the prediction of the mass transfer coefficient from an experimental dataset of FESEM images. In this work, we focused on the neural network's generalization capability and assessed the accuracy of the predictions from new images by employing a leave-one-out technique. This allowed us to conduct a comprehensive analysis of the knowledge extracted from the dataset. The original dataset is comprised of nine samples, so the size of the dataset limits the generalization capability. We showed that artificial augmentation of the dataset can be useful for improving the performance of the neural network in the prediction of the mass transfer coefficient from new images.

New experimental data are necessary to further improve the model's performance, particularly the extrapolation outside the mass transfer coefficients' training range. Nevertheless, coupling with data augmentation techniques can still be a valid way to extract features from the dataset efficiently.

The limited proof of concept presented in this work shows that these deep learning models, appropriately trained and tuned, can greatly speed up the evaluation of the effectiveness of different formulations by surrogating the temperature measure from which the mass transfer coefficient can be computed and thus reducing the experimental load.

Future research should focus on conducting additional freeze-drying experiments to include a broader range of formulations, incorporating datasets from imaging techniques like micro-CT for complementary 3D structural insights, and exploring advanced neural network architectures such as transfer learning with pre-trained models to enhance predictive performance.

Even if this workflow has been proposed for this specific application, we think that both dataset scarcity and the subsequent processing to obtain a suitable input to the neural network is a common issue when dealing with experimental images. The workflow we implemented can thus be useful for different applications in the chemical engineering field.

Author Contributions: Conceptualization, R.P.; methodology, R.P., G.B. and A.M.; software, A.M. and G.B.; validation, R.P., A.M. and G.B.; formal analysis, A.M., G.B. and R.P.; resources, R.P. and G.B.; data curation, R.P.; writing—original draft preparation, R.P., A.M. and G.B.; writing—review and editing, R.P., A.M. and G.B.; and supervision, R.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: Data are available upon request.

Acknowledgments: Computational resources were provided by hpc@polito, which is a project of Academic Computing within the Department of Control and Computer Engineering at the Politecnico di Torino (<http://www.hpc.polito.it>, accessed on 22 August 2024). We acknowledge the CINECA award under the ISCRA initiative for the availability of high-performance computing resources and support.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

FESEM	Field Emission Scanning Electron Microscopy
CNN	Convolutional neural network
PVP	Polyvinylpyrrolidone

References

1. Bhatnagar, B.; Tchessalov, S. Advances in freeze drying of biologics and future challenges and opportunities. In *Drying Technologies for Biotechnology and Pharmaceutical Applications*; Wiley: Hoboken, NJ, USA, 2020; pp. 137–177.
2. Food and Drug Administration. Novel Drug Approvals at FDA. 2024. Available online: <https://www.fda.gov/drugs/development-approval-process-drugs/novel-drug-approvals-fda> (accessed on 22 August 2024).
3. National Academies of Sciences, Engineering, and Medicine. *Innovations in Pharmaceutical Manufacturing on the Horizon: Technical Challenges, Regulatory Issues, and Recommendations*; The National Academies Press: Washington, DC, USA, 2021.
4. Arsiccio, A.; Pisano, R. Application of the quality by design approach to the freezing step of freeze-drying: Building the design space. *J. Pharm. Sci.* **2018**, *107*, 1586–1596. [[CrossRef](#)] [[PubMed](#)]
5. Pisano, R.; Fissore, D.; Barresi, A.A. In-line and off-line optimization of freeze-drying cycles for pharmaceutical products. *Dry. Technol.* **2013**, *31*, 905–919. [[CrossRef](#)]
6. Pisano, R.; Fissore, D.; Barresi, A.A. Quality by design in the secondary drying step of a freeze-drying process. *Dry. Technol.* **2012**, *30*, 1307–1316. [[CrossRef](#)]
7. Fissore, D.; Pisano, R.; Barresi, A.A. Using mathematical modeling and prior knowledge for QbD in freeze-drying processes. In *Quality by Design for Biopharmaceutical Drug Product Development*; Springer: New York, NY, USA, 2015; pp. 565–593.
8. Fissore, D.; Pisano, R.; Barresi, A.A. A model-based framework to optimize pharmaceuticals freeze drying. *Dry. Technol.* **2012**, *30*, 946–958. [[CrossRef](#)]
9. Tang, X.C.; Nail, S.L.; Pikal, M.J. Evaluation of manometric temperature measurement, a process analytical technology tool for freeze-drying: Part II measurement of dry-layer resistance. *Aaps Pharmscitech* **2006**, *7*, E77–E84. [[CrossRef](#)]
10. Grassini, S.; Pisano, R.; Barresi, A.A.; Angelini, E.; Parvis, M. Frequency domain image analysis for the characterization of porous products. *Measurement* **2016**, *94*, 515–522. [[CrossRef](#)]
11. Gruber, S.; Greiner, J.; Eppink, A.; Thomik, M.; Coppens, F.; Vorhauer-Huget, N.; Tsotsas, E.; Foerst, P. Pore shape matters—In-situ investigation of freeze-drying kinetics by 4D XCT methods. *Food Res. Int.* **2024**, *193*, 114837. [[CrossRef](#)] [[PubMed](#)]
12. Thomik, M.; Gruber, S.; Kaestner, A.; Foerst, P.; Tsotsas, E.; Vorhauer-Huget, N. Experimental study of the impact of pore structure on drying kinetics and sublimation front patterns. *Pharmaceutics* **2022**, *14*, 1538. [[CrossRef](#)]
13. Foerst, P.; Gruber, S.; Schulz, M.; Vorhauer, N.; Tsotsas, E. Characterization of lyophilization of frozen bulky solids. *Chem. Eng. Technol.* **2020**, *43*, 789–796. [[CrossRef](#)]
14. Pisano, R.; Barresi, A.A.; Capozzi, L.C.; Novajra, G.; Oddone, I.; Vitale-Brovarone, C. Characterization of the mass transfer of lyophilized products based on X-ray micro-computed tomography images. *Dry. Technol.* **2017**, *35*, 933–938. [[CrossRef](#)]
15. Stratta, L.; Adali, M.B.; Barresi, A.A.; Boccardo, G.; Marcato, A.; Tuccinardi, R.; Pisano, R. A diffused-interface model for the lyophilization of a packed bed of spray-frozen particles. *Chem. Eng. Sci.* **2023**, *275*, 118726. [[CrossRef](#)]
16. Pal, M.; Makauskas, P.; Malik, S. Upscaling porous media using neural networks: A deep learning approach to homogenization and averaging. *Processes* **2023**, *11*, 601. [[CrossRef](#)]
17. Wu, J.; Yin, X.; Xiao, H. Seeing permeability from images: Fast prediction with convolutional neural networks. *Sci. Bull.* **2018**, *63*, 1215–1222. [[CrossRef](#)] [[PubMed](#)]
18. Abdel Azim, R.; Aljehani, A. Neural network model for permeability prediction from reservoir well logs. *Processes* **2022**, *10*, 2587. [[CrossRef](#)]
19. Marcato, A.; Boccardo, G.; Marchisio, D.L. A computational workflow to study particle transport in porous media: Coupling CFD and deep learning. In *Computer Aided Chemical Engineering*; Elsevier: Amsterdam, The Netherlands, 2020; Volume 48, pp. 1759–1764.
20. Marcato, A.; Boccardo, G.; Marchisio, D. From computational fluid dynamics to structure interpretation via neural networks: An application to flow and transport in porous media. *Ind. Eng. Chem. Res.* **2022**, *61*, 8530–8541. [[CrossRef](#)]
21. Marcato, A.; Santos, J.E.; Liu, C.; Boccardo, G.; Marchisio, D.; Franco, A.A. Modeling the 4D discharge of lithium-ion batteries with a multiscale time-dependent deep learning framework. *Energy Storage Mater.* **2023**, *63*, 102927. [[CrossRef](#)]
22. Weber, R.M.; Korneev, S.; Battiato, I. Homogenization-informed convolutional neural networks for estimation of li-ion battery effective properties. *Transp. Porous Media* **2022**, *145*, 527–548. [[CrossRef](#)]

23. Wei, H.; Zhao, S.; Rong, Q.; Bao, H. Predicting the effective thermal conductivities of composite materials and porous media by machine learning methods. *Int. J. Heat Mass Transf.* **2018**, *127*, 908–916. [CrossRef]
24. Rani, K.S.; Jayadurga, R.; Raja, V.; Kumar, M.S.; Swathi, R.S.V.R.; Kumar, P. Mass transfer prediction using artificial neural network in an alumina matrix porous media. *Eur. Chem. Bull.* **2022**, *11*, 113–120.
25. Li, P.; Ma, C.; Chen, Z.; Wang, H.; Wang, Y.; Bai, H. A Review: Study on the Enhancement Mechanism of Heat and Moisture Transfer in Deformable Porous Media. *Processes* **2023**, *11*, 2699. [CrossRef]
26. Marcato, A.; Marchisio, D.; Boccardo, G. Reconciling deep learning and first-principle modelling for the investigation of transport phenomena in chemical engineering. *Can. J. Chem. Eng.* **2023**, *101*, 3013–3018. [CrossRef]
27. Sternberg, S.R. Biomedical image processing. *Computer* **1983**, *16*, 22–34. [CrossRef]
28. WS, R. Imagej, Us National Institutes of Health, Bethesda, Maryland, USA. 2011. Available online: <http://imagej.nih.gov/ij/> (accessed on 22 August 2024).
29. Pizer, S.M.; Amburn, E.P.; Austin, J.D.; Cromartie, R.; Geselowitz, A.; Greer, T.; ter Haar Romeny, B.; Zimmerman, J.B.; Zuiderveld, K. Adaptive histogram equalization and its variations. *Comput. Vision Graph. Image Process.* **1987**, *39*, 355–368. [CrossRef]
30. Alqahtani, N.; Alzubaidi, F.; Armstrong, R.T.; Swietojanski, P.; Mostaghimi, P. Machine learning for predicting properties of porous media from 2d X-ray images. *J. Pet. Sci. Eng.* **2020**, *184*, 106514. [CrossRef]
31. Ding, X.; Zhang, X.; Ma, N.; Han, J.; Ding, G.; Sun, J. Repvgg: Making vgg-style convnets great again. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Nashville, TN, USA, 20–25 June 2021; pp. 13733–13742.
32. Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; Salakhutdinov, R. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **2014**, *15*, 1929–1958.
33. Arlot, S.; Celisse, A. A survey of cross-validation procedures for model selection. *Stat. Surv.* **2010**, *4*, 40–79. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.