Tabular-to-Image Encoding Methods for Melanoma Detection: A Proof-of-Concept

Vanesa Gómez-Martínez *, David Chushig-Muzo, Cristina Soguero-Ruiz *

Posted Date: 30 January 2026

doi: 10.20944/preprints202601.2351.v1

Keywords: melanoma detection; dermoscopic images; lesion-level statistical features; tabular-toimage; spatial encoding methods; tabular2image; LM-IGTD; convolutional neural networks

Article

Tabular-to-Image Encoding Methods for Melanoma Detection: A Proof-of-Concept

Vanesa Gómez-Martínez *, David Chushig-Muzo and Cristina Soguero-Ruiz *

Department of Signal Theory and Communications, Rey Juan Carlos University, Fuenlabrada, Madrid, Spain
* Correspondence: vanesa.gomez@urjc.es (V.G.-M.); cristina.soguero@urjc.es (C.S.-R.)

Abstract

Deep Learning (DL) models have demonstrated strong performance in dermatological applications, particularly when trained on dermoscopic images. In contrast, tabular clinical data—such as patient metadata and lesion-level descriptors—are difficult to integrate into DL-based pipelines due to their heterogeneous, non-spatial, and often low-dimensional nature. As a result, these data are commonly handled using separate classical machine learning (ML) models. In this work, we present a proof-of-concept study that investigates whether dermatological tabular data can be transformed into two-dimensional image representations to enable convolutional neural network (CNN)-based learning. To this end, we employ the Low Mixed-Image Generator for Tabular Data (LM-IGTD), a framework designed to transform low-dimensional and heterogeneous tabular data into two-dimensional image representations, through type-aware encoding and controlled feature augmentation. Using this approach, we encode low-dimensional clinical metadata, high-dimensional lesion-level statistical features extracted from dermoscopic images, as well as their feature-level fusion, into grayscale image representations. The resulting image representations serve as input to CNNs, and the performance is compared with ML models trained on tabular data. Experiments conducted on the Derm7pt and PH2 datasets show that traditional ML models generally achieve the highest Area Under the Curve values, while LM-IGTD-based representations provide comparable performance and enable the use of CNNs on structured clinical data used in dermatology.

Keywords: melanoma detection; dermoscopic images; lesion-level statistical features; tabular-to-image; spatial encoding methods; tabular2image; LM-IGTD; convolutional neural networks

1. Introduction

Skin cancer remains one of the most prevalent malignancies worldwide, with both melanoma and keratinocyte skin cancers showing a sustained increase, particularly among fair-skinned populations [1,2]. Early detection is essential, as timely diagnosis is strongly associated with improved treatment outcomes and prognosis, especially in melanoma [1,2].

Dermoscopy, a noninvasive imaging technique, has been extensively used to enhance the visual assessment of skin lesions and support clinical decision-making [3]. The availability of large, publicly accessible dermoscopic image datasets has further accelerated research in automated skin lesion analysis [4,5]. This has enabled the development of data-driven models with high predictive performance [6,7], with convolutional neural networks (CNNs) playing a central role to leverage the spatial structure of dermoscopic images for predictive tasks [8]. As a result, most recent advances in automated skin lesion assessment have primarily focused on image-based deep learning (DL) pipelines, often relying almost exclusively on dermoscopic images [9].

However, clinical decision-making in dermatology does not rely exclusively on images. In addition to dermoscopic patterns, clinicians consider patient-level clinical metadata, such as age, sex, and anatomical location, as well as lesion-level descriptors derived from image analysis, including

 doi:10.20944/preprints202601.2351.v1

asymmetry, border irregularity, color distribution, and texture-related characteristics [10]. While clinical metadata typically give rise to low-dimensional tabular representations, lesion-level descriptors extracted from dermoscopic images are often summarized through high-dimensional handcrafted or statistically derived feature sets [11,12]. Together, these complementary sources of tabular information capture clinically meaningful aspects of skin lesions. Such structured tabular information is clinically relevant but remains underutilized in modern DL–based dermatological pipelines [9,13]. This is largely due to the well-known challenges that tabular dermatological data pose for DL models, including their heterogeneity, lack of inherent spatial structure, and limited dimensionality. As a result, standard DL architectures—originally designed to model spatial or temporal correlations—often struggle to effectively learn from tabular representations [14–16].

To mitigate these limitations, recent work has explored transforming tabular data into image-like representations, thereby enabling the application of CNN-based models to non-visual data [17–22]. These approaches aim to impose a spatial structure on tabular variables, allowing convolutional architectures to exploit local patterns and relationships. However, such methods have rarely been investigated in the context of dermatological data [9], where heterogeneous clinical metadata and lesion-level descriptors are routinely available and clinically meaningful.

Within this line of research, the Low Mixed-Image Generator for Tabular Data (LM-IGTD) framework [23] introduced a permutation-based strategy to map numerical and categorical features onto two-dimensional images according to their statistical relationships. In addition, LM-IGTD incorporates a noise-based feature augmentation mechanism specifically designed to improve representations in low-dimensional settings, making it particularly suitable for clinical data scenarios commonly encountered in biomedical applications.

To the best of our knowledge, the application of tabular-to-image transformation frameworks to dermatological data has not yet been investigated. In this context, the present work is conceived as a proof-of-concept (PoC) study exploring the feasibility of LM-IGTD-based tabular-to-image representations for dermatological clinical data. Specifically, this study considers different sources of tabular information derived from dermoscopic datasets, including patient-level clinical metadata, lesion-level statistical features extracted from images, and their joint representation, in order to assess their suitability for CNN-based melanoma classification.

The remainder of this paper is organized as follows. Section 2 presents the publicly available PH2 and Derm7pt datasets, which provide dermoscopic images together with clinical metadata, as well as the feature extraction process and the LM-IGTD transformation method. Experimental results are reported in Section 3, followed by a discussion in Section 4 and concluding remarks in Section 5.

2. Materials and Methods

This section describes the datasets, feature extraction procedures, and tabular-to-image transformation pipeline employed in this study. An overview of the complete methodology is provided in Figure 1, which summarizes the main processing stages of the proposed pipeline, from tabular data construction to CNN-based melanoma classification.

First, patient-level clinical metadata and lesion-level statistical features are extracted from the PH2 and Derm7pt datasets and represented as structured tabular data. These tabular representations are subsequently analyzed either independently or jointly, depending on the experimental setting. Next, the resulting tabular representations are transformed into two-dimensional grayscale images using the LM-IGTD framework, which incorporates feature ranking, pixel arrangement optimization, and noise-based feature augmentation when required. Finally, the generated images are used as input to CNN-based models for melanoma classification.

 doi:10.20944/preprints202601.2351.v1

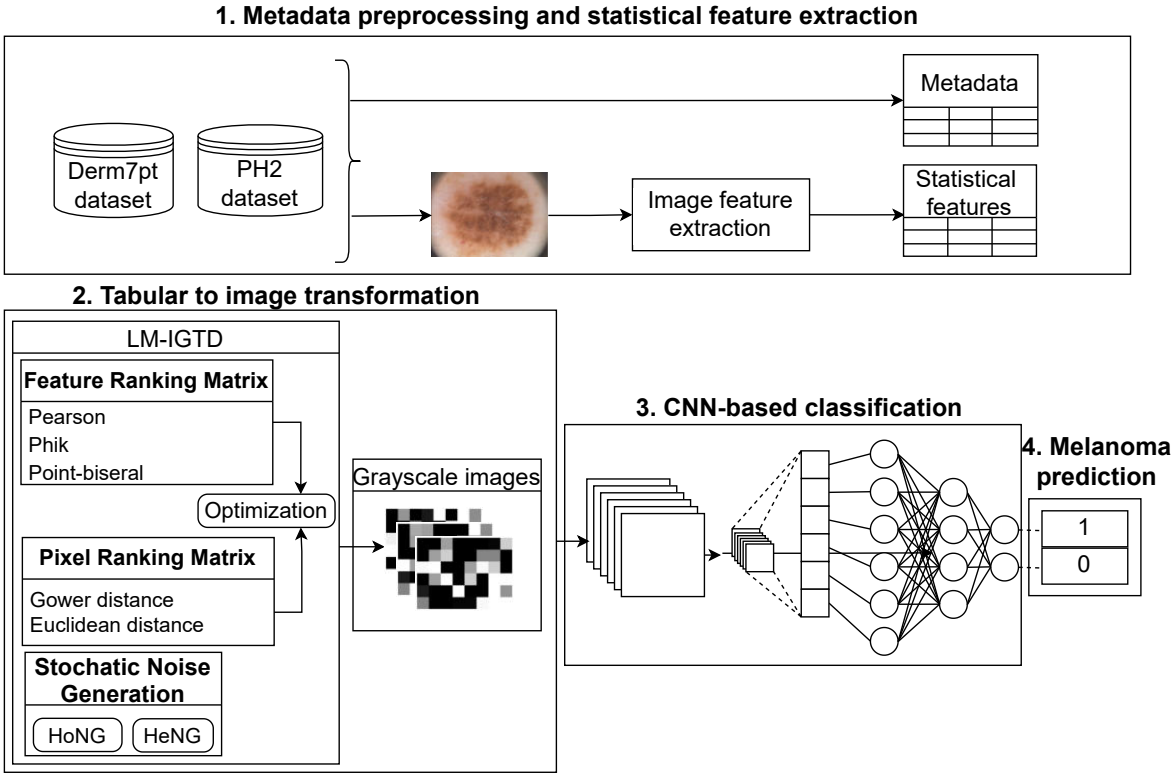

Figure 1. Overview of the proposed methodology.

2.1. Dataset Description

This study uses two publicly available dermoscopy datasets, PH2 [24] and Derm7pt [25], both of which provide dermoscopic images together with clinical metadata.

The PH2 dataset consists of 200 dermoscopic images acquired at the Dermatology Service of Hospital Pedro Hispano (Matosinhos, Portugal), including 160 benign lesions—80 common nevi and 80 atypical nevi—and 40 melanomas [24]. In this study, common and atypical nevi are grouped into a single *not melanoma* class, while all melanoma cases are assigned to the *melanoma* class, resulting in a binary classification scenario. PH2 provides metadata associated with dermoscopic criteria, including lesion colors, asymmetry, pigment network, dots and globules, streaks, regression areas, and blue-whitish veil.

The Derm7pt dataset contains over 2,000 dermoscopic and macroscopic images, accompanied by metadata derived from the seven-point checklist and additional clinical information [25]. This study focuses only on the 1,011 dermoscopic images available in the dataset. Lesions are grouped into a binary classification setting, with melanoma-related categories merged into a single *melanoma* class and the remaining lesion types assigned to the *not melanoma* class, resulting in 252 melanoma and 759 not melanoma samples. The associated metadata include dermoscopic criteria (pigment network, streaks, pigmentation, regression structures, dots and globules, blue-whitish veil, and vascular structures), as well as clinical and lesion-related variables such as diagnostic difficulty, elevation, anatomical location, sex, and management.

Table 1 summarizes the main characteristics of the tabular metadata used in this study, which consist of low-dimensional, clinically interpretable variables.

 doi:10.20944/preprints202601.2351.v1

Table 1. Overview of the tabular metadata used in the PH2 and Derm7pt datasets.

Dataset	Samples	Melanoma / Not melanoma	No. of variables	Metadata variables
PH2	200	40 / 160	7	Asymmetry; pigment network; dots and globules; streaks; regression areas; blue-whitish veil; lesion color
Derm7pt	1,011	252 / 759	12	Pigment network; streaks; pigmentation; regression structures; dots and globules; blue-whitish veil; vascular structures; diagnostic difficulty; elevation; anatomical location; sex; management

2.2. Image Feature Extraction

Lesion-level statistical features were extracted from dermoscopic skin lesions to obtain clinically and statistically meaningful descriptors represented in tabular form. These features aim to capture complementary aspects of lesion morphology, color distribution, and texture patterns that are routinely assessed by dermatologists during visual examination [26]. Following established practice in dermoscopic image analysis, we considered geometric features, color features, and a diverse set of local and global texture descriptors [27,28].

Geometric features were derived using the ABCD rule [29], which encodes asymmetry, border irregularity, color variation, and lesion diameter—criteria closely related to melanoma risk assessment. Color information was characterized using multiple color spaces, including RGB, HSV, CIE L*a*b, CIE L*u*v, and YCrCb. For each color channel, first-order statistical moments were computed to describe the distribution and variability of pigmentation within the lesion.

Texture features were extracted to characterize spatial variations and structural patterns within the lesion region. A broad range of statistical and signal-processing approaches was employed to characterize texture, including first- and second-order statistics, run-length and size-zone matrices, wavelet-based representations, and local pattern descriptors [30]. These methods quantify heterogeneity, regularity, and spatial organization of pixel intensities, which have been shown to be informative for melanoma detection.

Table 2 provides an overview of the extracted feature categories, the corresponding techniques, and their dimensionality, together with representative studies illustrating the use of these descriptors in dermoscopic image analysis and melanoma detection.

Table 2. Summary of handcrafted image features extracted from skin lesions.

Feature category	Technique	Representative studies	No. of features
Geometric	ABCD rule	[31–33]	4
Texture	DWT	[34,35]	34
	FDTA	[36]	4
	FOS	[32,37]	15
	GLCM	[38–40]	14
	GLDS	[41]	5
	GLRLM	[42,43]	16
	GLSZM	[44,45]	14
	HOS	[46]	2
	King	[47]	5
	LBP	[48,49]	6
	SFM	[41]	4
	WP	[50,51]	125
Color	RGB	[39,52]	4
	HSV	[39,52,53]	4
	CIE L*a*b	[39,52]	12
	CIE L*u*v	[54,55]	12
	YCrCb	[39,52,54]	12

Acronyms: Asymmetry, Border, Color, and Diameter (ABCD); discrete wavelet transform (DWT); fractal dimension texture analysis (FDTA); first-order statistics (FOS); gray-level co-occurrence matrix (GLCM); gray-level difference statistics (GLDS); gray-level run-length matrix (GLRLM); gray-level size zone matrix (GLSZM); higher-order spectra (HOS); local binary pattern (LBP); statistical feature matrix (SFM); wavelet packet (WP).

 doi:10.20944/preprints202601.2351.v1

Overall, the resulting feature set constitutes a high-dimensional tabular representation that captures complementary low-level characteristics of skin lesions, making it suitable for tabular-to-image transformation using the LM-IGTD framework, which supports both high- and low-dimensional heterogeneous data. A more detailed description of the feature extraction procedures and associated descriptors can be found in our previous work [41].

2.3. Tabular-to-Image Transformation Using LM-IGTD

LM-IGTD [23] is a tabular-to-image transformation framework designed to convert structured tabular data into two-dimensional grayscale image representations for image-based DL models. LM-IGTD builds upon the original IGTD [21] approach by extending its applicability to low-dimensional and mixed-type datasets, which are frequently encountered in clinical and biomedical domains. While IGTD focuses on permuting features to preserve their statistical relationships in a spatial layout, LM-IGTD introduces additional mechanisms—namely, type-aware similarity measures for mixed data and noise-based feature augmentation—to address the challenges posed by heterogeneous variable types and limited feature dimensionality.

A central limitation when applying tabular-to-image transformations to clinical metadata lies in the low dimensionality of many real-world datasets. When the number of available features is small, the resulting image representations may have limited spatial resolution, restricting the ability of CNN-based models to effectively exploit spatial patterns [14,56]. LM-IGTD explicitly addresses this issue by incorporating a stochastic noise-based feature augmentation strategy, which increases the dimensionality of the tabular input prior to image generation while preserving its underlying statistical structure.

The feature augmentation process is performed in an unsupervised manner and introduces synthetic features through controlled stochastic perturbations of the original variables. For numerical features, additive Gaussian noise is applied, with noise intensity governed by randomly sampled scaling parameters. For categorical variables, a swap-noise strategy is employed, whereby category values are randomly exchanged among samples with a given probability. LM-IGTD supports both Homogeneous Noise Generation (HoNG), in which a fixed number of noisy variants is generated per original feature, and Heterogeneous Noise Generation (HeNG), where a larger pool of noisy features is generated, and a subset is randomly selected. Among the generated augmented datasets, the one exhibiting the most compact and well-separated structure—assessed using unsupervised clustering criteria—is selected for subsequent image generation.

In the experimental evaluation, different augmentation levels were explored for low-dimensional clinical metadata by generating 3, 5, or 7 synthetic noisy features per original variable, in order to control the effective dimensionality of the resulting tabular representations.

In the present PoC study, noise-based feature augmentation is applied exclusively to low-dimensional clinical metadata, where the limited number of variables would otherwise lead to image representations with insufficient spatial resolution. In contrast, lesion-level statistical features extracted from dermoscopic images are inherently high-dimensional and naturally yield image representations with an adequate number of pixels. Consequently, these statistical feature sets are transformed directly into image form using LM-IGTD without the introduction of additional synthetic features.

Once the augmented tabular representation is obtained (when applicable), the image generation process proceeds as follows. In LM-IGTD, each tabular sample is mapped to a grayscale image in which each pixel corresponds to a single input feature, and the pixel intensity represents the normalized feature value for that sample. The spatial arrangement of features within the image is determined through an optimization process that seeks to preserve intrinsic relationships among features. To this end, LM-IGTD computes a feature ranking matrix encoding pairwise similarities or dissimilarities between features using type-aware statistical measures. Specifically, different correlation metrics are employed depending on the variable types, including Spearman [57] correlation for numerical–numerical pairs, point-biserial [58] correlation for numerical–binary pairs, and Phik [59] correlation for pairs involving

 doi:10.20944/preprints202601.2351.v1

categorical variables. This design allows numerical, binary, and categorical features to be jointly considered within a unified spatial encoding.

In parallel, a pixel ranking matrix is defined based on the pairwise distances between pixel locations on the two-dimensional image grid, computed using standard spatial distance measures (Euclidean for numerical data or Gower [60] for mixed types). An iterative optimization procedure is then employed to align the feature ranking matrix with the pixel ranking matrix, such that features exhibiting similar statistical relationships—according to the correlation measures described above—are assigned to nearby pixel locations, while dissimilar features are placed farther apart. This alignment ensures that spatial proximity in the generated image reflects statistical relatedness between variables.

3. Results

This section presents the experimental setting and the results obtained using different tabular data modalities and image representation strategies. Results are organized by data modality, including clinical metadata, lesion-level statistical features, and their fusion. For each modality, performance is reported for both traditional machine learning (ML) models trained directly on tabular data and CNN-based models trained on LM-IGTD-generated image representations, and results are shown for the Derm7pt and PH2 datasets.

3.1. Experimental Setting

For each dataset, samples were randomly split into training (80%) and test (20%) sets, with 15% of the training data further reserved for validation and used for hyperparameter tuning and early stopping. All experiments were repeated five times using different random seeds, applied to the training–test splitting, and results are reported as average performance across runs. Given the class imbalance present in both datasets, random undersampling was applied exclusively to the training data prior to model fitting. Model performance was evaluated using the area under the receiver operating characteristic curve (AUC), sensitivity, and specificity [61].

For tabular baselines, the evaluated ML models included Decision Tree (DT), Random Forest (RF), Support Vector Machine (SVM), and Least Absolute Shrinkage and Selection Operator (LASSO) [61,62]. The evaluated models and their corresponding hyperparameter ranges for both tabular baselines and CNN-based approaches are summarized in Table 3 and Table 4.

Table 3. Hyperparameter ranges considered for ML models trained on tabular data.

Model	Hyperparameter	Values considered
DT	Maximum depth	$[2, 8]$
	Minimum samples per split	$[15\%, 20\%]$ of training set size
RF	Number of trees	$\{10, 20, 30\}$
	Maximum depth	$[2, 8]$
	Minimum samples per split	$[2, 10]$
SVM	Regularization parameter	$[0.1, 10]$
	Kernel coefficient	$\{10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}\}$
LASSO	Regularization parameter	$[10^{-1.5}, 10^{0.4}]$

 doi:10.20944/preprints202601.2351.v1

Table 4. Hyperparameter ranges considered for CNN-based models trained on LM-IGTD-generated images.

Component	Hyperparameter	Values considered
Convolutional layers	Number of filters	$\{8, 16, 32, 64\}$
	Kernel size	$\{3, 4, 5\}$
Pooling layers	Pool size	$\{2, 3\}$
Fully connected layers	Number of units	$\{32, 64, 128\}$
Regularization	Dropout rate	$\{0.1, 0.2, 0.4\}$
Optimization	Learning rate	$[10^{-4}, 10^{-2}]$
	Optimizer	{Adam, RMSprop}
LM-IGTD images	Minimum image size (pixels)	$\{25, 35, 45\}$
Training	Batch size	$\{10, 20, 30\}$
	Early stopping	based on validation loss

3.2. Results Using Clinical Metadata

This subsection reports the results obtained using clinical metadata variables described in Section 2.1 (Table 1). Due to the low dimensionality of metadata in both datasets, LM-IGTD was applied together with noise-based feature augmentation to generate image representations suitable for CNN-based learning. Both HoNG and HeNG noise generation strategies were evaluated using different augmentation levels, considering the addition of 3, 5, and 7 synthetic features per original variable. Table 5 summarizes the classification performance obtained on the PH2 and Derm7pt datasets using clinical metadata.

Table 5. Classification performance using clinical metadata from the PH2 and Derm7pt datasets. Results are reported for traditional ML models trained on tabular data and CNN-based models trained on LM-IGTD image representations with different noise augmentation configurations (3, 5, 7). Best results for each dataset are shown in bold.

Dataset	Data type	Model	AUC	Sensitivity	Specificity
PH2	Tabular	DT	**0.833 ± 0.085**	**0.875 ± 0.102**	0.792 ± 0.126
		LASSO	0.818 ± 0.077	0.875 ± 0.102	0.760 ± 0.126
		RF	0.828 ± 0.084	0.875 ± 0.102	0.781 ± 0.133
		SVM	0.771 ± 0.045	0.792 ± 0.156	0.750 ± 0.111
	Image - no noise	CNN	0.785 ± 0.036	0.667 ± 0.156	0.904 ± 0.115
	Image + HoNG (3)		0.743 ± 0.041	0.625 ± 0.102	0.861 ± 0.041
	Image + HeNG (3)		0.734 ± 0.046	0.500 ± 0.102	**0.968 ± 0.026**
	Image + HoNG (5)		0.712 ± 0.064	0.583 ± 0.059	0.840 ± 0.070
	Image + HeNG (5)		0.702 ± 0.033	0.458 ± 0.118	0.946 ± 0.055
	Image + HoNG (7)		0.681 ± 0.052	0.542 ± 0.102	0.823 ± 0.087
	Image + HeNG (7)		0.693 ± 0.040	0.417 ± 0.083	0.955 ± 0.044
Derm7pt	Tabular	DT	0.847 ± 0.029	0.902 ± 0.064	0.792 ± 0.040
		LASSO	0.848 ± 0.021	0.928 ± 0.033	0.768 ± 0.008
		RF	**0.863 ± 0.014**	**0.935 ± 0.040**	0.792 ± 0.017
		SVM	0.858 ± 0.034	0.908 ± 0.056	0.807 ± 0.016
	Image - no noise	CNN	0.803 ± 0.023	0.797 ± 0.051	**0.809 ± 0.009**
	Image + HoNG (3)		0.792 ± 0.006	0.778 ± 0.024	0.807 ± 0.020
	Image + HeNG (3)		0.810 ± 0.018	0.824 ± 0.028	0.796 ± 0.021
	Image + HoNG (5)		0.777 ± 0.012	0.778 ± 0.018	0.776 ± 0.019
	Image + HeNG (5)		0.812 ± 0.010	0.817 ± 0.021	0.807 ± 0.016
	Image + HoNG (7)		0.803 ± 0.020	0.824 ± 0.042	0.783 ± 0.019
	Image + HeNG (7)		0.798 ± 0.006	0.791 ± 0.040	0.805 ± 0.030

 doi:10.20944/preprints202601.2351.v1

Results on the PH2 dataset show increased variability across noise configurations, which can be attributed to the smaller dataset size and the limited number of available metadata variables. In this setting, CNN-based models display a noticeable trade-off between sensitivity and specificity, especially for HeNG, which often favors higher specificity at the expense of sensitivity. Nevertheless, the overall performance trends remain coherent across runs, indicating that LM-IGTD can be applied to low-dimensional clinical metadata even in small-sample scenarios.

For the Derm7pt dataset, CNN-based models trained on LM-IGTD-generated image representations achieve competitive and stable performance across all configurations. While their overall performance remains below that of tabular baselines, image-based models exhibit consistent behavior across noise levels. In particular, configurations employing HeNG tend to provide more stable results than image representations generated without noise, suggesting that controlled feature augmentation can help mitigate the limitations imposed by the low dimensionality of clinical metadata.

Overall, these results support the feasibility of transforming low-dimensional dermatological metadata into image representations using LM-IGTD, enabling the use of CNN-based models in settings where traditional tabular learning approaches remain dominant.

3.3. Results Using Lesion-Level Statistical Features

This subsection reports the results obtained using lesion-level statistical features extracted from dermoscopic images. In contrast to clinical metadata, these features form a high-dimensional tabular representation composed of numerous handcrafted descriptors capturing geometric, color, and texture-related characteristics of skin lesions. As the original feature dimensionality is sufficient to generate image representations with adequate spatial resolution, LM-IGTD was applied without noise-based feature augmentation. Table 6 summarizes the classification performance obtained on the PH2 and Derm7pt datasets using lesion-level statistical features. For both datasets, traditional ML models trained directly on tabular data provide strong baseline performance, particularly in terms of sensitivity. In contrast, CNN-based models trained on LM-IGTD-generated image representations consistently achieve higher specificity. Although AUC values are comparable across approaches, the two modeling paradigms favor different error profiles, reflecting a trade-off between false negatives and false positives rather than a single clearly superior approach.

Table 6. Classification performance using statistical features for PH2 and Derm7pt datasets. Results are reported for traditional ML models trained on tabular data and CNN-based models trained on image representations without noise augmentation. Best results for each dataset are shown in bold.

Dataset	Data type	Model	AUC	Sensitivity	Specificity
PH2	Tabular	DT	0.739 ± 0.085	0.833 ± 0.236	0.667 ± 0.167
		LASSO	0.781 ± 0.069	$\mathbf{0.833 \pm 0.236}$	0.729 ± 0.164
		RF	0.771 ± 0.079	0.833 ± 0.236	0.708 ± 0.191
		SVM	0.740 ± 0.066	0.667 ± 0.236	0.812 ± 0.125
	Image - no noise	CNN	$\mathbf{0.785 \pm 0.036}$	0.667 ± 0.156	$\mathbf{0.904 \pm 0.115}$
Derm7pt	Tabular	DT	0.648 ± 0.023	0.627 ± 0.111	0.669 ± 0.105
		LASSO	$\mathbf{0.709 \pm 0.017}$	$\mathbf{0.752 \pm 0.037}$	0.667 ± 0.017
		RF	0.682 ± 0.014	0.680 ± 0.033	0.684 ± 0.057
		SVM	0.696 ± 0.032	0.667 ± 0.064	0.726 ± 0.003
	Image - no noise	CNN	0.633 ± 0.023	0.503 ± 0.082	$\mathbf{0.763 \pm 0.057}$

On the PH2 dataset, CNN-based models achieve comparable or improved performance in terms of AUC and specificity relative to tabular baselines, despite the limited sample size. Notably, the increase in specificity is more pronounced than in Derm7pt, indicating that LM-IGTD-based representations may be particularly effective in small-sample settings when the original feature space is high-dimensional.

 doi:10.20944/preprints202601.2351.v1

For the Derm7pt dataset, CNN-based models trained on LM-IGTD-generated image representations exhibit an increase in specificity compared to tabular baselines, while traditional ML models achieve higher sensitivity. This behavior suggests that the spatial encoding induced by LM-IGTD leads CNNs to learn more conservative decision boundaries, favoring the reduction of false positive predictions when learning from high-dimensional statistical descriptors. Overall, these results demonstrate that LM-IGTD can effectively transform high-dimensional dermatological statistical features into image representations suitable for CNN-based learning. This supports its applicability beyond low-dimensional clinical metadata and motivates its use in subsequent multimodal fusion experiments.

3.4. Results Using Fused Clinical Metadata and Statistical Features

This subsection reports the results obtained using the fusion of clinical metadata and lesion-level statistical features. Feature fusion was performed by concatenating clinical metadata and lesion-level statistical features at the tabular level, followed by either direct learning using traditional ML models or transformation into image representations using LM-IGTD for CNN-based classification. Given the high dimensionality of the fused feature space, noise-based feature augmentation was not applied. Table 7 summarizes the classification performance obtained on the PH2 and Derm7pt datasets using fused features. For both datasets, traditional ML models trained on the fused tabular representation provide strong baseline performance, achieving improved or comparable results relative to models trained on individual feature modalities.

Table 7. Classification performance using feature fusion for PH2 and Derm7pt datasets. Results are reported for traditional ML models trained on fused tabular features and CNN-based models trained on fused image representations without noise. Best results for each dataset are shown in bold.

Dataset	Data type	Model	AUC	Sensitivity	Specificity
PH2	Fusion (tabular)	DT	0.875 ± 0.062	0.875 ± 0.118	$\mathbf{0.875 \pm 0.118}$
		LASSO	$\mathbf{0.906 \pm 0.034}$	$\mathbf{0.958 \pm 0.059}$	0.854 ± 0.053
		RF	0.875 ± 0.036	0.833 ± 0.059	0.812 ± 0.026
		SVM	0.807 ± 0.048	0.792 ± 0.059	0.823 ± 0.039
	Fusion (image) – no noise	CNN	0.737 ± 0.114	0.625 ± 0.270	0.849 ± 0.055
Derm7pt	Fusion (tabular)	DT	0.797 ± 0.037	0.804 ± 0.058	$\mathbf{0.789 \pm 0.019}$
		LASSO	$\mathbf{0.848 \pm 0.008}$	$\mathbf{0.928 \pm 0.033}$	0.768 ± 0.031
		RF	0.801 ± 0.018	0.817 ± 0.046	0.785 ± 0.014
		SVM	0.814 ± 0.028	0.843 ± 0.042	0.785 ± 0.022
	Fusion (image) – no noise	CNN	0.775 ± 0.026	0.784 ± 0.061	0.765 ± 0.033

On the Derm7pt dataset, tabular fusion leads to improved performance compared to using lesion-level statistical features alone, while achieving performance comparable to that obtained using clinical metadata, particularly for LASSO-based models. CNN-based models trained on fused image representations achieve stable but lower performance, indicating that, in large datasets with rich tabular information, direct learning on fused tabular features remains more effective than image-based representations in this setting.

In contrast, results on the PH2 dataset highlight the benefit of feature fusion in small-sample settings. Tabular models trained on fused features achieve the highest AUC and sensitivity among all evaluated configurations. CNN-based models trained on fused image representations exhibit higher specificity but increased variability in sensitivity, reflecting the combined effect of limited sample size and high feature dimensionality.

Overall, these results indicate that feature fusion enhances classification performance when complementary sources of tabular information are combined. While traditional ML models remain strong baselines for fused tabular data, LM-IGTD-based image representations provide a coherent and

 doi:10.20944/preprints202601.2351.v1

stable alternative for integrating heterogeneous dermatological features within a unified image-based learning framework.

Figure 2 provides an overview of the comparative performance across data modalities, including clinical metadata, lesion-level statistical features, and their fusion.

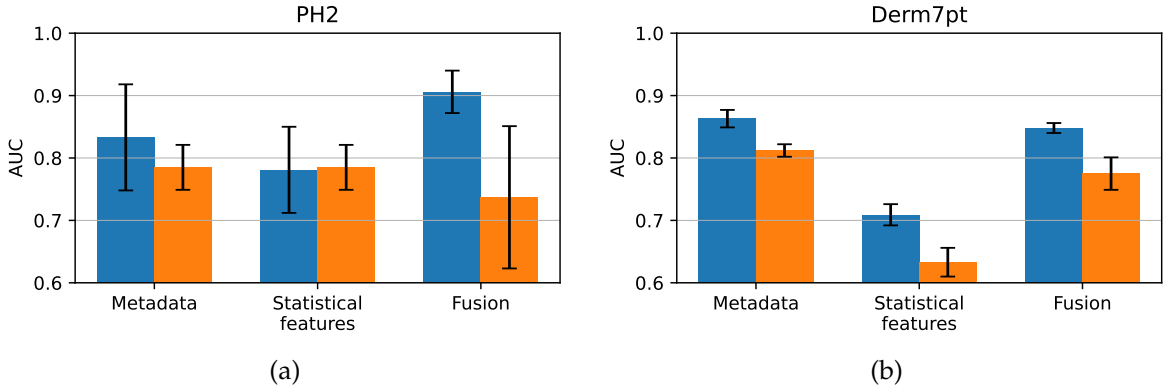

Figure 2. Summary of classification performance (AUC) obtained by the best-performing tabular ML models and LM-IGTD-based CNN models across different data modalities. Results are shown for (a) PH2 and (b) Derm7pt datasets. Blue bars: ML models trained on tabular data; orange bars: CNN-based models trained on LM-IGTD-generated image representations.

For each modality, the figure reports the AUC achieved by the best-performing tabular ML models and the best-performing LM-IGTD-based CNN models for the Derm7pt and PH2 datasets. Overall, the figure illustrates how classification performance varies across data modalities and modeling approaches. While traditional tabular models generally achieve strong AUC values, LM-IGTD-based CNN models provide competitive results in several configurations, highlighting their feasibility as an alternative representation strategy depending on the feature dimensionality and dataset size.

4. Discussion

This work presents a PoC study investigating the feasibility of representing clinically meaningful dermatological tabular data as two-dimensional image-like structures using the LM-IGTD framework, and subsequently applying CNNs for melanoma classification. Rather than aiming to outperform established image-based or tabular learning approaches, the primary objective of this study was to assess whether tabular-to-image transformations constitute a valid and coherent representation strategy for heterogeneous dermatological data sources.

The experimental results demonstrate that LM-IGTD can effectively encode both low-dimensional clinical metadata and high-dimensional lesion-level statistical features into structured image representations suitable for CNN-based learning. Across both the Derm7pt and PH2 datasets, CNN models trained on LM-IGTD-generated images exhibited stable and consistent performance trends, supporting the hypothesis that tabular dermatological data can be mapped to a spatial domain without destroying discriminative information. This observation is particularly relevant in dermatology, where structured clinical descriptors and derived lesion features are commonly used alongside visual assessment.

For clinical metadata, which typically consist of a limited number of heterogeneous variables, the incorporation of noise-based feature augmentation proved useful in stabilizing the generated image representations and facilitating CNN learning. In particular, heterogeneous noise generation led to more consistent performance in the larger Derm7pt dataset, suggesting that controlled stochastic augmentation can mitigate the limitations imposed by low feature dimensionality. These findings align with the design principles of LM-IGTD and confirm its applicability to clinical metadata scenarios frequently encountered in real-world dermatological settings. In contrast, lesion-level statistical features extracted from dermoscopic images are inherently high-dimensional and naturally support the construction of image representations without additional augmentation. In this context, LM-IGTD-based CNN models achieved performance levels comparable to traditional tabular baselines, while

 doi:10.20944/preprints202601.2351.v1

exhibiting a systematic increase in specificity in several experimental configurations. This behavior suggests that the spatial organization induced by LM-IGTD enables CNNs to learn more conservative decision boundaries when operating on high-dimensional statistical descriptors, potentially reducing false positive predictions.

Consistent with prior work [23,63], traditional ML models trained directly on tabular data remain strong baselines across all evaluated scenarios, particularly for low-dimensional clinical metadata. Importantly, the goal of this study is not to replace such models, but to explore an alternative representation paradigm that enables the use of DL architectures in contexts dominated by structured data. The fact that LM-IGTD-based CNN models achieve competitive and coherent performance—despite the additional representation step—constitutes a meaningful outcome, demonstrating that the tabular-to-image transformation preserves relevant information and supports stable learning behavior.

An additional contribution of this work lies in the exploration of feature-level fusion through tabular-to-image transformation. By concatenating clinical metadata and lesion-level statistical features prior to image generation, this study introduces a unified representation in which heterogeneous sources of information are jointly embedded within a single spatial structure. This approach can be interpreted as an implicit fusion mechanism operating entirely within the tabular domain, allowing CNNs to model interactions between clinical and lesion-level attributes without requiring raw image data at inference time. The fusion experiments indicate that this unified encoding is feasible and yields coherent performance trends across datasets, even though direct tabular learning remains highly effective.

Beyond the specific CNN architectures evaluated in this study, an important implication of the proposed tabular-to-image representation is that it enables the use of a broader class of DL models originally developed for visual data. Once tabular clinical information is encoded as a structured image, more advanced architectures such as residual [64] and densely connected networks [65], as well as more recent attention-based and transformer-based models [66,67], can be directly applied without requiring fundamental changes to the input representation. From a PoC perspective, the competitive performance observed with relatively simple CNN architectures suggests that further performance gains may be achievable by leveraging more expressive models capable of capturing long-range spatial dependencies and higher-order feature interactions [68].

Another promising direction for future research lies in the integration of LM-IGTD-generated images with raw dermoscopic images within multimodal learning frameworks. In such a setting, tabular-derived image representations could act as a complementary visual modality, encoding clinical metadata and lesion-level descriptors alongside dermoscopic appearance. This opens the possibility of jointly modeling structured clinical information and visual lesion characteristics using unified image-based architectures, potentially enhancing robustness and interpretability in dermatological decision-support systems. While the present study focuses on establishing feasibility, the proposed framework naturally lends itself to multimodal extensions that combine generated and acquired images.

The present study is subject to certain constraints that are inherent to its PoC nature. The experimental evaluation was conducted on two publicly available datasets, with relatively limited sample sizes—particularly in the case of PH2—and no external validation cohort was considered. In addition, the CNN architectures explored were kept relatively simple, as the emphasis of this work was placed on assessing the feasibility of the proposed representation rather than on architectural optimization. Finally, the analysis was limited to binary melanoma classification using tabular features derived from dermoscopic datasets. These aspects define the scope of the present study and naturally motivate future research involving larger and more diverse cohorts, external validation, more advanced network architectures, and extended clinical scenarios.

Overall, this PoC study demonstrates that LM-IGTD-based tabular-to-image representations provide a feasible and conceptually sound framework for encoding heterogeneous dermatological data. By enabling CNN-based learning on clinical metadata, lesion-level statistical features, and their

 doi:10.20944/preprints202601.2351.v1

fusion, this work lays the groundwork for future research exploring more advanced architectures and multimodal strategies for dermatological decision support.

5. Conclusions

This work presents a PoC study on the feasibility of representing heterogeneous dermatological tabular data as two-dimensional image-like structures using the LM-IGTD framework for CNN-based melanoma classification. The study focuses on assessing whether tabular-to-image transformations can serve as a coherent and stable representation strategy for clinically relevant structured dermatological data.

The experimental results show that LM-IGTD can effectively encode both low-dimensional clinical metadata and high-dimensional lesion-level statistical features into structured image representations without degrading their discriminative content. CNN models trained on these representations exhibited consistent and reproducible performance across datasets and data modalities, indicating that the proposed transformation supports stable learning behavior. While traditional ML models trained directly on tabular data remain strong baselines, the competitive performance of LM-IGTD-based CNNs confirms the feasibility of the proposed representation paradigm.

Beyond individual data modalities, the results also demonstrate that feature-level fusion through tabular-to-image transformation is feasible, enabling the joint representation of clinical metadata and lesion-level statistical features within a unified spatial structure. This highlights the flexibility of the LM-IGTD framework for integrating heterogeneous dermatological data sources within a single representation suitable for CNN-based learning.

Overall, this PoC study establishes LM-IGTD-based tabular-to-image representations as a conceptually sound and practically viable approach for encoding structured dermatological data. By demonstrating feasibility across multiple data types and experimental settings, this work provides a foundation for future research involving larger cohorts, external validation, and the exploration of more advanced architectures and multimodal learning strategies for dermatological decision support.

Author Contributions: Conceptualization, V.G.-M., D.C.-M. and C.S.-R.; Methodology, V.G.-M., D.C.-M. and C.S.-R.; Software, V.G.-M.; Validation, V.G.-M., D.C.-M. and C.S.-R.; Formal analysis, V.G.-M.; Investigation, V.G.-M.; Resources, C.S.-R.; Data curation, V.G.-M.; Writing—original draft preparation, V.G.-M.; Writing—review and editing, V.G.-M., D.C.-M. and C.S.-R.; Visualization, V.G.-M.; Supervision, D.C.-M. and C.S.-R.; Project administration, C.S.-R.; Funding acquisition, C.S.-R. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the European Commission through the H2020-EU.3.1.4.2, European Project WARIFA (Watching the risk factors: Artificial intelligence and the prevention of chronic conditions) under the Grant Agreement 101017385; and by the Spanish Government through the grants AAVis-BMR PID2019-107768RA and PID2023-149457OB (both funded by the AEI agency, AEI/10.13039/501100011033). The study sponsors have not been involved in any stage of the study.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data used in this study are publicly available. The PH2 dataset is available at https://www.fc.up.pt/addi/ph2%20database.html, and the Derm7pt dataset is available at https://derm.cs.sfu.ca/Welcome.html.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

 doi:10.20944/preprints202601.2351.v1

Abbreviations

The following abbreviations are used in this manuscript:

AUC	Area Under the Receiver Operating Characteristic Curve
CNN	Convolutional Neural Network
DL	Deep Learning
DT	Decision Tree
HeNG	Heterogeneous Noise Generation
HoNG	Homogeneous Noise Generation
IGTD	Image Generator for Tabular Data
LASSO	Least Absolute Shrinkage and Selection Operator
LM-IGTD	Low Mixed-Image Generator for Tabular Data
ML	Machine Learning
PoC	Proof-of-Concept
RF	Random Forest
SVM	Support Vector Machine

References

1. Sun, Y.; et al. Global trends in melanoma burden: a comprehensive analysis from the global burden of disease study, 1990-2021. *Journal of the American Academy of Dermatology* **2025**, *92*, 100–107.
2. Garbe, C.; et al. Epidemiology of cutaneous melanoma and keratinocyte cancer in white populations 1943–2036. *European Journal of Cancer* **2021**, *152*, 18–25.
3. Rojas, K.D.; et al. Skin cancer: Primary, secondary, and tertiary prevention. Part II. *Journal of the American Academy of Dermatology* **2022**, *87*, 271–288.
4. Goyal, M.; et al. Artificial intelligence-based image classification methods for diagnosis of skin cancer: Challenges and opportunities. *Computers in Biology and Medicine* **2020**, *127*, 104065.
5. Wen, D.; et al. Characteristics of publicly available skin cancer image datasets: a systematic review. *The Lancet Digital Health* **2022**, *4*, e64–e74.
6. Li, H.; et al. Skin disease diagnosis with deep learning: A review. *Neurocomputing* **2021**, *464*, 364–393.
7. Gómez-Martínez, V.; et al. A multimodal and interpretable-based approach for improving melanoma detection using dermoscopy images, 2023. Available at https://doi.org/10.21203/rs.3.rs-3781517/v1.
8. Hasan, M.K.; et al. A survey, review, and future trends of skin lesion segmentation and classification. *Computers in Biology and Medicine* **2023**, *155*, 106624.
9. Wu, Y.; et al. Skin cancer classification with deep learning: a systematic review. *Frontiers in Oncology* **2022**, *12*, 893972.
10. Khan, M.A.; et al. An implementation of normal distribution based segmentation and entropy controlled features selection for skin lesion detection and classification. *BMC Cancer* **2018**, *18*, 1–20.
11. Venkat, N. The curse of dimensionality: Inside out. *Pilani (IN): Birla Institute of Technology and Science, Pilani, Department of Computer Science and Information Systems* **2018**, *10*.
12. Lee, H.D.; et al. Dermoscopic assisted diagnosis in melanoma: Reviewing results, optimizing methodologies and quantifying empirical guidelines. *Knowledge-Based Systems* **2018**, *158*, 9–24.
13. Park, A.J.; Weintraub, G.S.; Asgari, M.M. Leveraging the electronic health record to improve dermatologic care delivery: The importance of finding structure in data, 2020.
14. Shwartz-Ziv, R.; Armon, A. Tabular data: Deep learning is not all you need. *Information Fusion* **2022**, *81*, 84–90.
15. Arik, S.Ö.; Pfister, T. Tabnet: Attentive interpretable tabular learning. In Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence, USA, May 2021; Vol. 35, pp. 6679–6687.
16. Medeiros Neto, L.; Rogerio da Silva Neto, S.; Endo, P.T. A comparative analysis of converters of tabular data into image for the classification of Arboviruses using Convolutional Neural Networks. *Plos One* **2023**, *18*, e0295598.
17. Wolf, T.N.; et al. DAFT: a universal module to interweave tabular data and 3D images in CNNs. *NeuroImage* **2022**, *260*, 119505.

 doi:10.20944/preprints202601.2351.v1

18. Lee, E.; Nam, M.; Lee, H. Tab2vox: CNN-based multivariate multilevel demand forecasting framework by tabular-to-voxel image conversion. *Sustainability* **2022**, *14*, 11745.
19. Sharma, A.; et al. DeepInsight: A methodology to transform a non-image data to an image for convolution neural network architecture. *Scientific Reports* **2019**, *9*, 11399.
20. Bazgir, O.; et al. Representation of features as images with neighborhood dependencies for compatibility with convolutional neural networks. *Nature Communications* **2020**, *11*, 4391.
21. Zhu, Y.; et al. Converting tabular data into images for deep learning with convolutional neural networks. *Scientific Reports* **2021**, *11*, 11325.
22. Castillo-Cara, M.; et al. TINTO: Converting Tidy Data into image for classification with 2-Dimensional Convolutional Neural Networks. *SoftwareX* **2023**, *22*, 101391.
23. Gómez-Martínez, V.; et al. LM-IGTD: a 2D image generator for low-dimensional and mixed-type tabular data to leverage the potential of convolutional neural networks. *arXiv preprint arXiv:2406.14566* **2024**.
24. Mendonça, T.; et al. Ph2: A public database for the analysis of dermoscopic images. *Dermoscopy Image Analysis* **2015**.
25. Kawahara, J.; et al. Seven-point checklist and skin lesion classification using multitask multimodal neural nets. *Journal of Biomedical and Health Informatics* **2018**, *23*, 538–546.
26. Liu, J.; Zou, X.b. *Practical Dermoscopy*; Springer Nature, 2022.
27. Camacho-Gutiérrez.; et al. Multi-class skin lesion classification using prism-and segmentation-based fractal signatures. *Expert Systems with Applications* **2022**, *197*, 116671.
28. Singh, L.; Janghel, R.R.; Sahu, S.P. A hybrid feature fusion strategy for early fusion and majority voting for late fusion towards melanocytic skin lesion detection. *International Journal of Imaging Systems and Technology* **2022**, *32*, 1231–1250.
29. Friedman, R.J.; et al. Early detection of malignant melanoma: the role of physician examination and self-examination of the skin. *A Cancer Journal for Clinicians* **1985**, *35*, 130–151.
30. De Siqueira, F.R.; et al. Multi-scale gray level co-occurrence matrices for texture description. *Neurocomputing* **2013**, *120*, 336–345.
31. Ali, A.R.H.; et al. Automating the ABCD rule for melanoma detection: a survey. *IEEE Access* **2020**, *8*, 83333–83346.
32. Kingsly, A.; Sankaragomathi, B. Performance Analysis of Machine Learning Based Classifiers for the Diagnosis of Melanoma Cancer and Comparison. *Journal of Computational and Theoretical Nanoscience* **2018**, *15*, 558–575.
33. Thanh, D.N.; et al. Melanoma skin cancer detection method based on adaptive principal curvature, colour normalisation and feature extraction with the ABCD rule. *Journal of Digital Imaging* **2020**, *33*, 574–585.
34. Karuppiah, S.; et al. An Efficient Galactic Swarm Optimization Based Fractal Neural Network Model with DWT for Malignant Melanoma Prediction. *Neural Processing Letters* **2022**, *54*, 5043–5062.
35. Waladi, A.; et al. Melanoma classification using texture and wavelet analysis. In Proceedings of the 2019 International Conference of Artificial Intelligence and Information Technology (ICAIIT), IEEE, Yogyakarta, Indonesia, March 2019; pp. 336–343.
36. Kumar, T.K.; Himanshu, I.N. Artificial intelligence based real-time skin cancer detection. In Proceedings of the 2023 15th International Conference on Computer and Automation Engineering (ICCAE), IEEE, Sydney, Australia, March 2023; pp. 215–219.
37. Arif, M.; et al. Automated detection of nonmelanoma skin cancer based on deep convolutional neural network. *Journal of Healthcare Engineering* **2022**, *2022*.
38. Kumar, M.; et al. A de-ann inspired skin cancer detection approach using fuzzy c-means clustering. *Mobile Networks and Applications* **2020**, *25*, 1319–1329.
39. Bansal, P.; et al. Detection of melanoma in dermoscopic images by integrating features extracted using handcrafted and deep learning models. *Computers & Industrial Engineering* **2022**, *168*, 108060.
40. Chatterjee, S.; Dey, D.; Munshi, S. Integration of morphological preprocessing and fractal based feature extraction with recursive feature elimination for skin lesion types classification. *Computer Methods and Programs in Biomedicine* **2019**, *178*, 201–218.
41. Gómez-Martínez, V.; et al. Ensemble feature selection and tabular data augmentation with generative adversarial networks to enhance cutaneous melanoma identification and interpretability. *BioData Mining* **2024**, *17*, 46.

 doi:10.20944/preprints202601.2351.v1

42. Tan, T.Y.; et al. Intelligent skin cancer detection using enhanced particle swarm optimization. *Knowledge-based Systems* **2018**, *158*, 118–135.
43. Murugan, A.; et al. Diagnosis of skin cancer using machine learning techniques. *Microprocessors and Microsystems* **2021**, *81*, 103727.
44. Peng, Y.; et al. Combining texture features of whole slide images improves prognostic prediction of recurrence-free survival for cutaneous melanoma patients. *World Journal of Surgical Oncology* **2020**, *18*, 1–8.
45. de Moura, L.V.; et al. Skin lesions classification using multichannel dermoscopic Images. *ANAIS DO XII SIMPÓSIO DE ENGENHARIA BIOMÉDICA-IX SIMPÓSIO DE INSTRUMENTAÇÃO E IMAGENS MÉDICAS, 2019, Brasil.* **2019**.
46. Singh, S.; Urooj, S. A Methodological Approach for Analysis of Melanoma Images. *Madridge Journal of Dermatology and Research* **2018**, *3*, 83–87.
47. Shrestha, B.; et al. Detection of atypical texture features in early malignant melanoma. *Skin Research and Technology* **2010**, *16*, 60–65.
48. Adjed, F.; et al. Fusion of structural and textural features for melanoma recognition. *IET Computer Vision* **2018**, *12*, 185–195.
49. Majtner, T.; et al. Combining deep learning and hand-crafted features for skin lesion classification. In Proceedings of the 2016 Sixth International Conference on Image Processing Theory, Tools and Applications, IEEE, Oulu, Finland, December 2016; pp. 1–6.
50. Jayaraman, P.; et al. Wavelet-based classification of enhanced melanoma skin lesions through deep neural architectures. *Information* **2022**, *13*, 583.
51. Narasimhan, K.; Elamaran, V. Wavelet-based energy features for diagnosis of melanoma from dermoscopic images. *International Journal of Biomedical Engineering and Technology* **2016**, *20*, 243–252.
52. Bansal, P.; et al. Improving the classification accuracy of melanoma detection by performing feature selection using binary Harris hawks optimization algorithm. *Soft Computing* **2022**, *26*, 8163–8181.
53. Barata, C.; et al. Melanoma detection algorithm based on feature fusion. In Proceedings of the 2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), IEEE, Milan, Italy, August 2015; pp. 2653–2656.
54. Shalu.; et al. Automated melanoma skin cancer detection from digital images. *International Journal of Biomedical Engineering and Technology* **2021**, *37*, 275–289.
55. Celebi, M.E.; et al. A methodological approach to the classification of dermoscopy images. *Computerized Medical Imaging and Graphics* **2007**, *31*, 362–373.
56. Hashemi, M. Enlarging smaller images before inputting into convolutional neural network: zero-padding vs. interpolation. *Journal of Big Data* **2019**, *6*, 1–13.
57. Yu, H.; Hutson, A.D. A robust Spearman correlation coefficient permutation test. *Communications in Statistics-Theory and Methods* **2022**, pp. 1–13.
58. Kornbrot, D. Point biserial correlation. *Wiley StatsRef: Statistics Reference Online* **2014**.
59. Baak, M.; et al. A new correlation coefficient between categorical, ordinal and interval variables with Pearson characteristics. *Computational Statistics & Data Analysis* **2020**, *152*, 107043.
60. Liu, P.; et al. A modified and weighted Gower distance-based clustering analysis for mixed type data: a simulation and empirical analyses. *BMC Medical Research Methodology* **2024**, *24*, 305.
61. Shalev-Shwartz, S.; Ben-David, S. *Understanding machine learning: From theory to algorithms*; Cambridge University Press, 2014.
62. Bishop, C.M. *Pattern Recognition and Machine Learning (Information Science and Statistics)*; Springer-Verlag: Berlin, Heidelberg, 2006.
63. Lara-Abelenda, F.J.; et al. Transfer learning for a tabular-to-image approach: A case study for cardiovascular disease prediction. *Journal of Biomedical Informatics* **2025**, *165*, 104821.
64. He, K.; et al. Deep residual learning for image recognition. In Proceedings of the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, June 2016; pp. 770–778.
65. Huang, G.; et al. Densely connected convolutional networks. In Proceedings of the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, July 2017; pp. 4700–4708.
66. Dosovitskiy, A. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* **2020**.

 doi:10.20944/preprints202601.2351.v1

67. Liu, Z.; et al. Swin transformer: Hierarchical vision transformer using shifted windows. In Proceedings of the Proceedings of the IEEE/CVF International Conference on Computer Vision, Montreal, QC, Canada, October 2021; pp. 10012–10022.
68. Khan, S.; et al. Transformers in vision: A survey. *ACM Computing Surveys (CSUR)* **2022**, *54*, 1–41.