Journal of Information Systems Engineering and Management
2024, 9(3)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Multi-Texture Fusion with ReliefF Subspace-Weighted Support Vector Machines for Lightweight Defect Classification of Electroluminescence Solar Cell Images

Priya Pal[1], Dr. Mukesh Yadav[2]

[1]*Research Scholar,* [2]*Professor*

[1,2]*Department of Electronics & Communication Engineering, SAGE University, Indore, India*

ARTICLE INFO

Received: 01 Aug 2024

Revised: 20 Sep 2024

Accepted: 29 Sept 2024

ABSTRACT

Electroluminescence (EL) imaging exposes micro-cracks, broken fingers and inactive regions in photovoltaic (PV) cells, but grading thousands of EL cell images by hand is slow and subjective. Most recent automated graders rely on deep convolutional neural networks (CNNs) or transformer hybrids that need graphics processing units (GPUs) and large labelled sets. This paper studies a deliberately lightweight alternative, MTF-RSWS, that runs on an ordinary central processing unit (CPU). Each cell is denoised with a median filter and enhanced with contrast-limited adaptive histogram equalisation (CLAHE); a 1,893-dimensional multi-texture fusion (MTF) vector is then built from local binary patterns (LBP), histograms of oriented gradients (HOG), grey-level co-occurrence matrix (GLCM) statistics, Gabor energies and intensity statistics. ReliefF ranks the features, and an ensemble of radial-basis-function support vector machines (RBF-SVMs) trained on ReliefF-weighted random subspaces (RSWS) makes the final decision by weighted soft voting. On a stratified 75/25 split of the public ELPV benchmark (2,624 cells, four defect-probability classes), MTF-RSWS reached 75.76% accuracy, 62.52% macro precision, 55.68% macro sensitivity and 57.99% macro F1-score, the best held-out macro F1-score among eight classical baselines; in the functional-versus-defective setting it reached 83.99% accuracy. Five-fold cross-validation, however, places its macro F1-score (55.53 ± 1.52%) within one standard deviation of a single SVM on the same features (56.14 ± 0.85%), so the held-out margin is not statistically reliable. Training took about 100 s on a two-core CPU and inference 5.8 ms per cell. The results sit below those of deep models on the same data, but they give a transparent, reproducible and hardware-light reference point, and the error analysis shows that the two uncertain classes remain the main obstacle.

Keywords: Electroluminescence imaging, photovoltaic cell defects, texture descriptors, ReliefF feature ranking, random subspace ensemble, support vector machine.

1. INTRODUCTION

Photovoltaic (PV) generation now carries a large share of new electricity capacity, and the economic value of a PV plant depends on its modules keeping their rated output for twenty years or more. Mechanical stress during manufacture, transport, installation and operation creates micro-cracks, finger interruptions and electrically inactive areas that are invisible to the eye but reduce power and grow over time [1], [2]. Electroluminescence (EL) imaging, in which a forward-biased module emits near-infrared light that is captured by a sensitive camera, makes these defects visible as dark lines and regions, and it has become the standard tool for quality control in factories and for audits in the field [3], [4].

The difficulty is scale. A single 60-cell module yields sixty cell images, a medium plant yields hundreds of thousands, and the visual judgement of an expert is both slow and inconsistent at the boundary between "probably fine" and "probably defective". The public ELPV benchmark makes this ambiguity explicit: each of its 2,624 cell images carries a defect probability of 0, 1/3, 2/3 or 1 assigned by experts [3], [4]. Automated grading on ELPV has followed the wider path of computer vision. Early work combined hand-crafted descriptors with support vector machines (SVMs) [4],

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

[6], [7], later studies trained convolutional neural networks (CNNs) [8]–[10], and recent ones use attention modules, deep feature fusion or very deep pre-trained backbones [11]–[13]. Reported accuracies have risen accordingly, but so has the hardware needed to train and deploy the models.

There is still a practical case for light, CPU-only graders. Many inspection stations in small factories and field audits run on embedded or office computers, labelled EL data are scarce, and maintenance engineers often want to know which image property drove a decision. Classical texture descriptors are cheap and interpretable, yet the ELPV literature rarely reports a careful, fully reproducible classical baseline in the four-class setting, and the gap between such a baseline and deep models is seldom measured under one protocol. This paper addresses that gap with a modest pipeline, MTF-RSWS, and reports its real performance including its weaknesses. The contributions are:

(i) a multi-texture fusion (MTF) descriptor that joins local binary patterns (LBP), histograms of oriented gradients (HOG), grey-level co-occurrence matrix (GLCM) statistics, Gabor energies and intensity statistics after median filtering and contrast-limited adaptive histogram equalisation (CLAHE);

(ii) a ReliefF subspace-weighted SVM (RSWS) ensemble in which ReliefF scores both guide the sampling of random feature subspaces and weight each member's soft vote;

(iii) an honest evaluation on ELPV with a held-out test set, five-fold cross-validation, eight baselines, an ablation study and per-class error analysis, all on a two-core CPU.

Section 2 reviews related work, Section 3 describes the method, Section 4 gives the experimental set-up, Section 5 presents and discusses the results, and Section 6 concludes.

2. RELATED WORK

2.1 Classical descriptors and machine learning

Deitsch *et al.* [4] introduced the ELPV data and compared an SVM trained on dense hand-crafted features with a fine-tuned VGG-19 network; the CNN was more accurate, but the SVM was much cheaper, and both struggled with the uncertain classes. The same group later proposed a robust cell-segmentation method for uncalibrated EL module images that makes such cell-level datasets possible [5]. Karimi *et al.* [6] built an automated pipeline that corrects, segments and classifies EL cell images with both hand-crafted features and CNNs, and Mantel *et al.* [7] extracted 25 statistical and geometric features per cell and showed that an RBF-SVM can reach high accuracy but low recall on rare defect types. These studies establish that texture statistics carry defect information, while leaving open how far a careful classical design can go in the four-class ELPV task.

2.2 Deep learning for EL and other PV images

Akram *et al.* [8] designed a light CNN for EL cells, Tang *et al.* [9] combined data augmentation with a CNN to identify several defect types, and Rahman and Chen [10] used a multi-attention network for polycrystalline cells. Su *et al.* [11] added complementary channel and spatial attention to a region-proposal detector for manufacturing defects. Demirci *et al.* [12] fused deep features from several pre-trained CNNs and classified them with an SVM, and Al-Waisy *et al.* [13] combined deep feature representations into a hybrid model; both report clear gains over earlier ELPV baselines in binary and four-class settings. Chen *et al.* [14] built an open pipeline that processed nearly 19,000 module images, and Verma *et al.* [15] compared VGG, ResNet and DenseNet models on a five-class EL set. Beyond EL, thermographic and visual images have been classified with multi-scale transfer-learning CNNs [16], EfficientNet features with neighbourhood component analysis and an SVM [17], ResNet-50 with Faster R-CNN localisation [18], simple CNNs [19], vision transformers [20], segmentation and transfer learning on aerial images [21], a real-time multi-variant CNN for RGB panel images [22] and unified detectors for several renewable assets [23]. Recent surveys stress that dataset bias, class imbalance and limited interpretability remain open problems [1], [2], [24], and generative augmentation has been proposed to ease data scarcity in crack imagery [25].

2.3 Research gap

Deep models dominate recent work, but their reported scores come from varied splits, augmentation orders and class definitions, and they demand GPU resources. A transparent, CPU-only baseline that uses feature relevance explicitly

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2024, 9(3)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

and is evaluated with a fixed held-out split and cross-validation is still useful as a reference point and as a deployable option where hardware is limited. MTF-RSWS is designed for that role.

3. PROPOSED METHOD

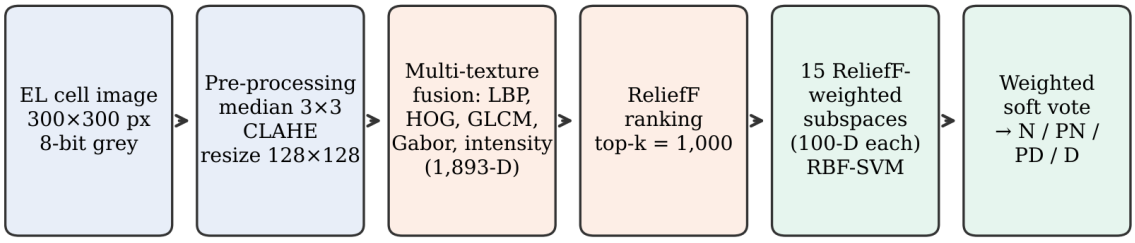


Figure 1. MTF-RSWS pipeline.

Figure 1 summarises the pipeline. An EL cell image of 300 × 300 pixels is first cleaned and contrast-enhanced, then described by five complementary texture families that are concatenated into one 1,893-dimensional MTF vector. ReliefF ranks every feature on the training data only, the 1,000 highest-ranked features form a candidate pool, and fifteen RBF-SVMs are trained on 100-feature subspaces drawn from this pool with probabilities proportional to their ReliefF weights. The members' class scores are combined by a vote in which each member counts in proportion to the total relevance of its subspace. The output is one of the four ELPV classes: normal (N), possibly normal (PN), possibly defective (PD) or defective (D).

3.1 Pre-processing

EL images contain sensor noise and strong brightness differences between mono- and polycrystalline cells. A 3 × 3 median filter suppresses isolated hot pixels while preserving thin crack edges, and CLAHE [26] with a clip limit of 2.0 on an 8 × 8 tile grid equalises local contrast so that faint cracks and dark inactive regions stand out. The enhanced image is resized to 128 × 128 pixels with area interpolation, which keeps feature extraction fast without removing busbar-scale structure.

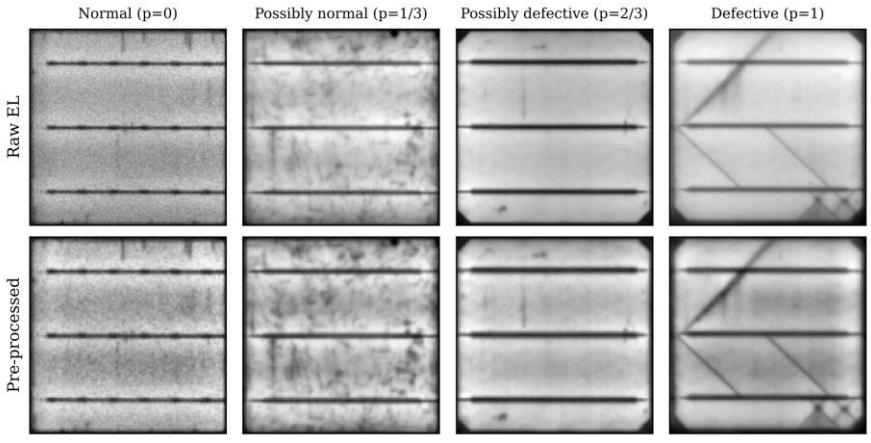


Figure 2. Example ELPV cells before and after pre-processing.

Figure 2 shows one randomly chosen cell from each class. The upper row is the raw EL image and the lower row the pre-processed version. CLAHE makes the diagonal cracks in the defective cell and the small vertical crack in the possibly defective cell much more visible, but it also amplifies the grain texture of polycrystalline cells, as in the possibly normal example. This texture resembles faint defects and is one reason why the PN and PD classes are hard to separate.

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2024, 9(3)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

3.2 Multi-texture fusion descriptor

Five descriptor families are computed on the 128 × 128 image. (a) Uniform rotation-invariant LBP histograms [27] at (P, R) = (8, 1) and (16, 2) give 28 features that capture micro-texture; the LBP code of a pixel with centre grey level g_c and P neighbours g_p on a circle of radius R is

$$LBP_{P,R} = \sum_{p=0}^{P-1} s(g_p - g_c)2^p \tag{1}$$

where $s(x) = 1$ for $x \geq 0$ and 0 otherwise. (b) HOG [28] with nine orientations, 16 × 16-pixel cells and 2 × 2-cell blocks gives 1,764 features that describe the direction of crack and finger edges. (c) GLCM statistics [29] (contrast, dissimilarity, homogeneity, energy, correlation and angular second moment) on 32 grey levels at distances 1, 2 and 4 pixels and four angles give 72 features; for a normalised co-occurrence matrix $p(i, j)$ the contrast is

$$Con = \sum_{i,j}(i-j)^2 p(i,j) \tag{2}$$

(d) The mean absolute response and standard deviation of twelve Gabor filters (four orientations × wavelengths of 4, 8 and 12 pixels) give 24 features sensitive to oriented line defects. (e) Five intensity statistics (mean, standard deviation, 5th percentile, median and the fraction of pixels darker than 0.25) describe inactive regions. All features are standardised to zero mean and unit variance using statistics from the training set only.

3.3 ReliefF feature ranking

ReliefF [30], [31] scores a feature by how well it separates each sample from its nearest neighbours of other classes compared with its nearest neighbours of the same class. For a randomly selected sample R with k nearest hits H_j and k nearest misses $M_j(C)$ in every other class C, the weight of feature f is updated over m iterations as

$$W_f \leftarrow W_f - \sum_{j=1}^{k} \frac{diff(f,R,H_j)}{mk} + \sum_{C \neq cl(R)} \frac{P(C)}{1-P(cl(R))} \sum_{j=1}^{k} \frac{diff\left(f,R,M_j(C)\right)}{mk} \tag{3}$$

where diff(·) is the normalised absolute difference of feature f between two samples and $P(C)$ is the prior of class C. Unlike filter scores such as variance or correlation, ReliefF accounts for feature interactions and handles more than two classes. We use $k = 10$ neighbours and a random sample of 900 training vectors to keep the computation light. The top $K = 1{,}000$ features form the pool T.

3.4 ReliefF subspace-weighted SVM ensemble

The random subspace method [32] trains each ensemble member on a different subset of features, which reduces variance when features outnumber informative samples. In RSWS, member m receives a subspace S_m of 100 features drawn from T without replacement, where feature f is chosen with probability

$$\pi_f = \frac{W_f^+}{\sum_{g \in T} W_g^+} \tag{4}$$

and W^+ denotes $\max(W, \varepsilon)$. Each member is an RBF-SVM [33] with $C = 20$, γ set to the inverse of the feature variance times the subspace size, and class weights inversely proportional to class frequency to counter the strong imbalance of ELPV. Each member's weight is the normalised relevance of its subspace, and the final label is obtained by soft voting over the softmax of the one-versus-rest decision values $d_m(\mathbf{x})$:

$$\alpha_m = \frac{\sum_{f \in S_m} W_f^+}{\sum_{n=1}^{15} \sum_{f \in S_n} W_f^+} \tag{5}$$

$$\hat{y} = arg\ max_c \sum_{m=1}^{15} \alpha_m softmax\left(d_m(x)\right)_c \tag{6}$$

The ensemble size, pool size, subspace size and C were chosen from a small grid ($K \in \{300, 600, 1{,}000\}$, subspace $\in \{100, 200\}$, $C \in \{5, 20\}$) on a validation split carved out of the training partition; the test partition was not used for any choice. The procedure is summarised in Algorithm 1.

Algorithm 1 lists the five steps. Augmentation by flipping is applied only to training cells, after the split, so no flipped copy of a test cell can leak into training. Because every step is deterministic given the random seed, the whole pipeline can be rerun exactly on a laptop.

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2024, 9(3)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Algorithm 1. MTF-RSWS training and inference

Input: training cells $\{(I_i, y_i)\}$, pool size K, members M, subspace size s

1: for each I_i: median 3 × 3 → CLAHE → resize 128 × 128; add horizontal and vertical flips

2: $\mathbf{x}_i \leftarrow$ [LBP | HOG | GLCM | Gabor | intensity]; standardise with training statistics

3: $W \leftarrow$ ReliefF($\mathbf{X}$, $\mathbf{y}$, k = 10) on 900 sampled vectors; $T \leftarrow$ top-K features by W

4: for $m = 1 \ldots M$: draw $S_m \subset T$, $|S_m| = s$, with probabilities from Eq. (4); train RBF-SVM on $\mathbf{X}[:, S_m]$; compute α_m by Eq. (5)

5: **Inference:** pre-process and describe test cell $\mathbf{x}$; return $\hat{y}$ by the weighted soft vote of Eq. (6)

4. EXPERIMENTAL SET-UP

4.1 Dataset and protocol

Experiments use the ELPV benchmark [3]–[5]: 2,624 grey-scale EL cell images of 300 × 300 pixels extracted from 44 modules (1,074 monocrystalline and 1,550 polycrystalline cells). The expert defect probabilities define four classes, N (1,508 cells), PN (295), PD (106) and D (715), so the data are strongly imbalanced. A stratified 75/25 split with a fixed seed gives 1,968 training and 656 test cells (377 N, 74 PN, 26 PD, 179 D). For the binary setting, N and PN are merged as functional and PD and D as defective, following [4]. In addition, stratified five-fold cross-validation is run for the strongest models to estimate the spread of the scores.

4.2 Baselines, metrics and environment

Eight classical models are compared, all with the same pre-processing and training-only augmentation: k-nearest neighbours (k-NN, k = 7) on MTF, a decision tree (DT), a random forest (RF, 400 trees), an RBF-SVM on LBP only, an RBF-SVM on HOG only, an RBF-SVM on the full MTF vector, an RBF-SVM on 150 principal components of MTF (MTF+PCA+SVM) and the proposed MTF-RSWS. Accuracy, macro-averaged precision, macro-averaged sensitivity (recall) and macro F1-score are reported; macro averaging gives the rare PD and PN classes the same weight as N and D, so it is the main criterion here. For K classes the macro F1-score is

$$Macro - F1 = \frac{1}{K}\sum_{c=1}^{K}\frac{2P_cR_c}{P_c+R_c} \quad (7)$$

where P_c and R_c are the precision and recall of class c. All code was written in Python with scikit-learn, scikit-image, OpenCV and skrebate and run on a two-core CPU with 8 GB of memory and no GPU.

5. RESULTS AND DISCUSSION

5.1 Four-class comparison

Table 1. Four-class results on the ELPV test set (656 cells).

Method	ACC (%)	PRE (%)	SEN (%)	F1 (%)	Train (s)	Test (ms/cell)
k-NN (MTF)	66.16	50.38	38.82	39.8	0.2	0.28
Decision tree (MTF)	62.5	43.24	42.89	43.04	14.1	0.02
Random forest (MTF)	73.32	50.14	50.22	50.01	28.7	0.21

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2024, 9(3)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

LBP + SVM	60.21	44.61	47.38	44.96	1.4	0.27
HOG + SVM	73.48	58.15	51.1	53.07	26.3	4.42
MTF + PCA + SVM	73.48	57.74	54.96	56.06	5.1	0.6
MTF + SVM	75	61.44	53.22	55.62	24.1	4.64
MTF-RSWS (proposed)	**75.76**	**62.52**	**55.68**	**57.99**	**100.2**	**5.78**

Table 1 lists the held-out results. On this split MTF-RSWS gives the highest macro F1-score (57.99%), macro sensitivity (55.68%) and accuracy (75.76%), as well as the highest macro precision (62.52%). Fusing the five descriptor families helps clearly: the full MTF vector beats LBP alone by 10.7 points of macro F1-score and HOG alone by 2.6 points. Among classifiers on the same features, the SVM variants are well ahead of *k*-NN and the decision tree and also ahead of the random forest, which tends to absorb the minority classes into N. The cost of the ensemble is a training time of about 100 s, most of it spent on ReliefF, compared with 24 s for a single SVM; inference stays below 6 ms per cell, which is fast enough for in-line inspection on a CPU.

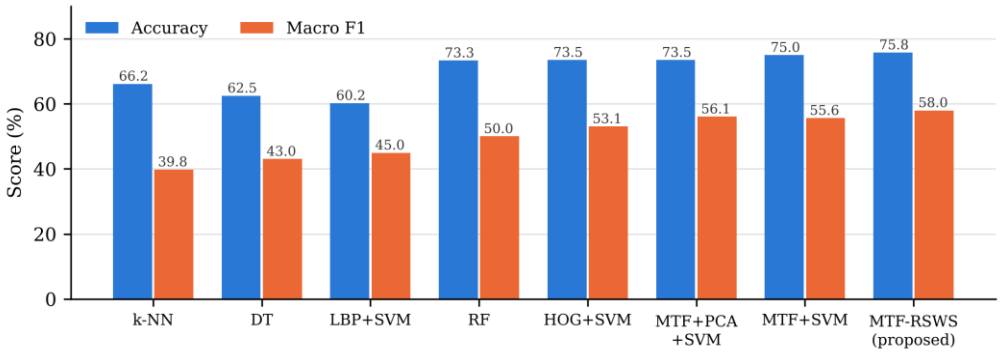


Figure 3. Accuracy and macro F1-score of all methods.

Figure 3 plots accuracy and macro F1-score side by side. The gap between the two bars is large for every model, between 15 and 27 points, because accuracy is dominated by the large N class whereas macro F1-score is pulled down by the small and ambiguous PN and PD classes. The figure also shows that the improvement of MTF-RSWS over the single-SVM baselines is small, 1.9 to 2.4 points of macro F1-score on this split, which motivated the cross-validation below.

5.2 Cross-validation and ablation

Table 2 repeats the comparison with stratified five-fold cross-validation, where each fold re-runs ReliefF and all training steps on its own training part. The picture is more sober than on the single split. MTF-RSWS has the highest mean accuracy (74.89 ± 1.02%), but its mean macro F1-score (55.53 ± 1.52%) is slightly below that of the single MTF + SVM (56.14 ± 0.85%) and the uniform subspace ensemble (56.08 ± 1.59%), and all three lie within one standard deviation of each other. The SVM-based models remain clearly ahead of the random forest (51.75 ± 1.74%). We therefore do not claim that RSWS is significantly better than a well-tuned single SVM; its held-out advantage in Table 1 is within the fold-to-fold variation, and the main value of the ensemble is the feature-relevance information it exposes.

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2024, 9(3)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Table 2. Five-fold cross-validation (mean ± standard deviation).

Method	ACC (%)	PRE (%)	SEN (%)	F1 (%)
Random forest (MTF)	72.45 ± 1.04	57.01 ± 4.31	51.19 ± 1.25	51.75 ± 1.74
MTF + PCA + SVM	72.22 ± 1.35	56.84 ± 1.79	54.38 ± 1.32	55.30 ± 1.53
MTF + SVM	74.77 ± 1.12	64.70 ± 6.22	53.59 ± 0.96	56.14 ± 0.85
Uniform subspace SVM (no ReliefF)	74.73 ± 0.61	66.97 ± 1.94	52.85 ± 1.27	56.08 ± 1.59
MTF-RSWS (proposed)	**74.89 ± 1.02**	**63.19 ± 3.18**	**53.08 ± 1.07**	**55.53 ± 1.52**

Table 3. Ablation of the RSWS components (held-out test set).

Variant	ACC (%)	F1 (%)	Binary ACC (%)	Binary F1 (%)
Single SVM, all 1,893 features	75.00	55.62	82.93	78.99
Single SVM, ReliefF top-1,000	76.22	56.45	83.23	79.43
15 uniform random subspaces, equal votes	76.52	57.40	85.21	81.77
MTF-RSWS (ReliefF-weighted subspaces and votes)	75.76	57.99	83.99	80.98

Table 3 isolates the two ingredients of RSWS. Restricting a single SVM to the ReliefF top-1,000 features raises macro F1-score from 55.62% to 56.45%, which confirms that ReliefF removes some uninformative HOG cells. A plain random-subspace ensemble with equal votes gives the highest accuracy (76.52%) and the best binary scores (85.21% accuracy), while ReliefF-weighted sampling and voting give the best four-class macro F1-score (57.99% versus 57.40%) and sensitivity. In other words, relevance weighting shifts the ensemble towards the minority classes at a small cost in overall accuracy; it does not dominate the simpler ensemble on every metric, and we report it as a trade-off rather than a uniform gain.

5.3 Error analysis and binary setting

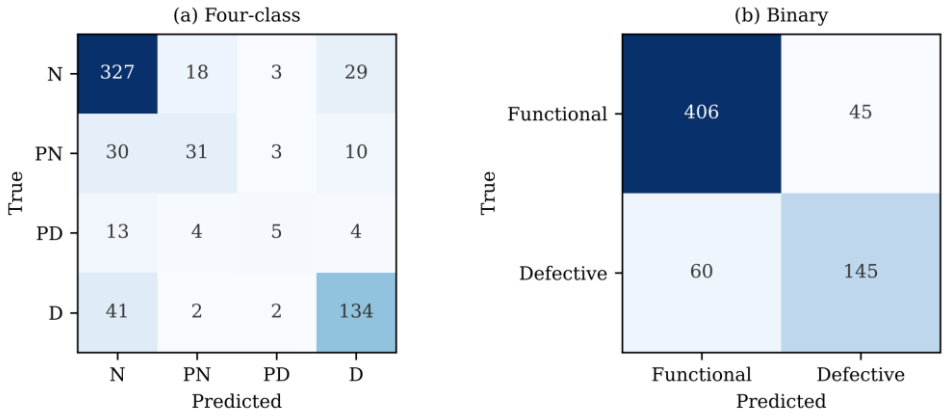


Figure 4. Confusion matrices of MTF-RSWS on the test set.

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2024, 9(3)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Figure 4(a) shows the four-class confusion matrix. Recall is 86.7% for N and 74.9% for D, but only 41.9% for PN and 19.2% for PD; 30 of the 74 PN cells are labelled N and 13 of the 26 PD cells are labelled N. The main error between the confident classes is also towards N: 41 defective cells, mostly with short or low-contrast cracks, are called normal. Figure 4(b) shows the binary result, 83.99% accuracy, 81.72% precision, 80.38% sensitivity and 80.98% F1-score, with 60 of 205 defective cells missed. These errors match the nature of the labels: PN and PD encode expert doubt, and texture statistics computed over the whole cell dilute a single short crack among large defect-free areas.

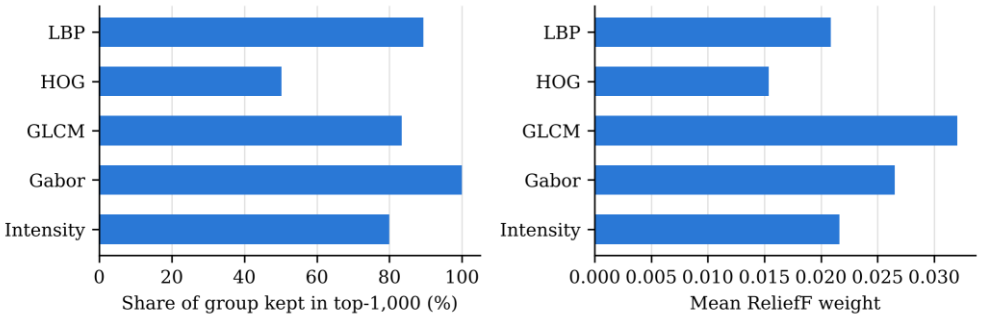


Figure 5. ReliefF relevance of the descriptor families.

Figure 5 shows how ReliefF judged each family on the training data. On the left is the share of each family that survives into the top-1,000 pool and on the right the mean ReliefF weight. GLCM, Gabor, LBP and intensity features are short, dense and almost all retained, and GLCM has the highest mean weight. HOG supplies most of the pool in absolute numbers, but only half of its 1,764 features are retained and its mean weight is the lowest, because many HOG cells describe busbars and cell borders that are common to all classes. This supports the use of relevance-guided sampling for the HOG-heavy MTF vector.

5.4 Comparison with published ELPV results

Table 4. Accuracy reported on ELPV by earlier studies.

Study	Model type	Classes	Reported ACC (%)
Demirci et al. [12]	Light CNN (L-CNN)	4	82.58
Demirci et al. [12]	Deep feature fusion + SVM (DFB-SVM)	4	89.63
Al-Waisy et al. [13]	ResNet-50 features	4	91.60
Al-Waisy et al. [13]	Hybrid deep features	4	95.35
This work	MTF-RSWS (hand-crafted, CPU only)	4	75.76
Deitsch et al. [4]	Dense feature descriptors + linear SVM	2	82.44
Deitsch et al. [4]	Fine-tuned VGG-19 CNN	2	88.42
This work	MTF-RSWS (hand-crafted, CPU only)	2	83.99

Values for [4] and [12] are as summarised in [13], and values for [13] are its own; their splits, augmentation and image sizes differ from ours, so the comparison is indicative only.

Table 4 places MTF-RSWS among published ELPV results. In the binary setting its accuracy (83.99%) is above the hand-crafted SVM of Deitsch *et al.* [4] and slightly below their fine-tuned VGG-19. In the four-class setting it is clearly below the deep models of [12] and [13], by about 7 to 20 points. This is the expected price of using hand-crafted texture instead of learned features, and it defines the boundary of the present approach: MTF-RSWS is not a replacement for deep models when GPU resources and accuracy are paramount, but a transparent, light and

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

reproducible baseline, suitable for low-cost screening where defective and borderline cells are passed to a human expert or a heavier model.

5.5 Limitations

Four limitations should be noted. First, the minority classes contain few samples (26 PD cells in the test set), so their per-class scores have wide uncertainty. Second, only one public dataset was used, and generalisation to other EL cameras and cell technologies is untested. Third, whole-cell descriptors cannot localise a defect, which limits interpretability at the pixel level. Fourth, ReliefF dominates the training time; a faster relevance estimator would make retraining on new production lines easier.

6. CONCLUSION

This paper presented MTF-RSWS, a CPU-only pipeline for classifying EL solar cell images that combines median filtering and CLAHE, a 1,893-dimensional multi-texture fusion descriptor and a ReliefF subspace-weighted SVM ensemble. On the four-class ELPV test split it reached 75.76% accuracy and 57.99% macro F1-score on a held-out test set, the best macro F1-score among eight classical baselines, and 83.99% accuracy in the binary setting, with inference under 6 ms per cell. Cross-validation showed that this held-out gain over a single SVM is not statistically reliable, since the mean macro F1-scores lie within one standard deviation of each other, and the ablation showed that relevance weighting mainly shifts decisions towards the minority classes. The model remains well below deep networks in four-class accuracy, and most of its errors come from the possibly normal and possibly defective classes. Future work will add crack-localised descriptors on cell sub-regions, cost-sensitive thresholds for the uncertain classes, evaluation on further EL datasets and a comparison with compact CNNs under the same protocol.

REFERENCES

[1] A. Mellit and S. Kalogirou, "Artificial intelligence and internet of things to improve efficacy of diagnosis and remote sensing of solar photovoltaic systems: Challenges, recommendations and future directions," *Renewable and Sustainable Energy Reviews*, vol. 143, Art. no. 110889, 2021.

[2] Y. Hammoudi, I. Idrissi, M. Boukabous, Y. Zerguit, and H. Bouali, "Review on maintenance of photovoltaic systems based on deep learning and internet of things," *Indonesian Journal of Electrical Engineering and Computer Science*, vol. 26, no. 2, pp. 1060–1072, 2022.

[3] C. Buerhop-Lutz, S. Deitsch, A. Maier, F. Gallwitz, S. Berger, B. Doll, J. Hauch, C. Camus, and C. J. Brabec, "A benchmark for visual identification of defective solar cells in electroluminescence imagery," in *Proc. 35th European PV Solar Energy Conference and Exhibition (EU PVSEC)*, 2018, doi: 10.4229/35thEUPVSEC20182018-5CV.3.15.

[4] S. Deitsch, V. Christlein, S. Berger, C. Buerhop-Lutz, A. Maier, F. Gallwitz, and C. Riess, "Automatic classification of defective photovoltaic module cells in electroluminescence images," *Solar Energy*, vol. 185, pp. 455–468, 2019.

[5] S. Deitsch, C. Buerhop-Lutz, E. Sovetkin, A. Steland, A. Maier, F. Gallwitz, and C. Riess, "Segmentation of photovoltaic module cells in uncalibrated electroluminescence images," *Machine Vision and Applications*, vol. 32, no. 4, Art. no. 84, 2021.

[6] A. M. Karimi, J. S. Fada, M. A. Hossain, S. Yang, T. J. Peshek, J. L. Braid, and R. H. French, "Automated pipeline for photovoltaic module electroluminescence image processing and degradation feature classification," *IEEE Journal of Photovoltaics*, vol. 9, no. 5, pp. 1324–1335, 2019.

[7] C. Mantel, F. Villebro, G. A. dos Reis Benatto, H. R. Parikh, S. Wendlandt, K. Hossain, P. Poulsen, S. Spataru, D. Sera, and S. Forchhammer, "Machine learning prediction of defect types for electroluminescence images of photovoltaic panels," in *Proc. SPIE 11139, Applications of Machine Learning*, 2019, Art. no. 1113904.

[8] M. W. Akram, G. Li, Y. Jin, X. Chen, C. Zhu, X. Zhao, A. Khaliq, M. Faheem, and A. Ahmad, "CNN based automatic detection of photovoltaic cell defects in electroluminescence images," *Energy*, vol. 189, Art. no. 116319, 2019.

[9] W. Tang, Q. Yang, K. Xiong, and W. Yan, "Deep learning based automatic defect identification of photovoltaic module using electroluminescence images," *Solar Energy*, vol. 201, pp. 453–460, 2020.

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

[10] M. R. U. Rahman and H. Chen, "Defects inspection in polycrystalline solar cells electroluminescence images using deep learning," *IEEE Access*, vol. 8, pp. 40547–40558, 2020.
[11] B. Su, H. Chen, P. Chen, G. Bian, K. Liu, and W. Liu, "Deep learning-based solar-cell manufacturing defect detection with complementary attention network," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 6, pp. 4084–4095, 2021.
[12] M. Y. Demirci, N. Beşli, and A. Gümüşçü, "Efficient deep feature extraction and classification for identifying defective photovoltaic module cells in electroluminescence images," *Expert Systems with Applications*, vol. 175, Art. no. 114810, 2021.
[13] A. S. Al-Waisy, D. Ibrahim, D. A. Zebari, S. Hammadi, H. Mohammed, M. A. Mohammed, and R. Damaševičius, "Identifying defective solar cells in electroluminescence images using deep feature representations," *PeerJ Computer Science*, vol. 8, Art. no. e992, 2022.
[14] X. Chen, T. Karin, and A. Jain, "Automated defect identification in electroluminescence images of solar modules," *Solar Energy*, vol. 242, pp. 20–29, 2022.
[15] H. Verma, S. V. Siruvuri, and P. R. Budarapu, "A machine learning-based image classification of silicon solar cells," *International Journal of Hydromechatronics*, vol. 7, no. 1, pp. 49–66, 2024.
[16] D. Korkmaz and H. Acikgoz, "An efficient fault classification method in solar photovoltaic modules using transfer learning and multi-scale convolutional neural network," *Engineering Applications of Artificial Intelligence*, vol. 113, Art. no. 104959, 2022.
[17] Z. B. Duranay, "Fault detection in solar energy systems: A deep learning approach," *Electronics*, vol. 12, no. 21, Art. no. 4397, 2023.
[18] S. P. Pathak, S. Patil, and S. Patel, "Solar panel hotspot localization and fault classification using deep learning approach," *Procedia Computer Science*, vol. 204, pp. 698–705, 2022.
[19] M. Pa, M. N. Uddin, and A. Kazemi, "A fault detection scheme utilizing convolutional neural network for PV solar panels with high accuracy," in *Proc. IEEE 1st Industrial Electronics Society Annual On-Line Conference (ONCON)*, 2022, pp. 1–5.
[20] D. Dwivedi, K. V. S. M. Babu, P. K. Yemula, P. Chakraborty, and M. Pal, "Identification of surface defects on solar PV panels and wind turbine blades using attention based deep learning model," arXiv:2211.15374, 2022.
[21] A. Shaik, A. Balasundaram, L. S. Kakarla, and N. Murugan, "Deep learning-based detection and segmentation of damage in solar panels," *Automation*, vol. 5, no. 2, pp. 128–150, 2024.
[22] S. Prabhakaran, R. A. Uthra, and J. Preetharoselyn, "Feature extraction and classification of photovoltaic panels based on convolutional neural network," *Computers, Materials & Continua*, vol. 75, no. 1, 2023.
[23] A. S. M. Shihavuddin, M. R. A. Rashid, M. H. Maruf, M. A. Hasan, M. A. ul Haq, R. H. Ashique, and A. Al Mansur, "Image based surface damage detection of renewable energy installations using a unified deep learning approach," *Energy Reports*, vol. 7, pp. 4566–4576, 2021.
[24] S. R. Mohanty, M. U. Maruf, V. Singh, and Z. Ahmad, "Machine learning approaches for automatic defect detection in photovoltaic systems," arXiv:2409.16069, 2024.
[25] E. Branikas, P. Murray, and G. West, "A novel data augmentation method for improved visual crack detection using generative adversarial networks," *IEEE Access*, vol. 11, pp. 22051–22059, 2023.
[26] K. Zuiderveld, "Contrast limited adaptive histogram equalization," in *Graphics Gems IV*, P. S. Heckbert, Ed. San Diego, CA, USA: Academic Press, 1994, pp. 474–485.
[27] T. Ojala, M. Pietikäinen, and T. Mäenpää, "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 7, pp. 971–987, 2002.
[28] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005, vol. 1, pp. 886–893.
[29] R. M. Haralick, K. Shanmugam, and I. Dinstein, "Textural features for image classification," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-3, no. 6, pp. 610–621, 1973.
[30] I. Kononenko, "Estimating attributes: Analysis and extensions of RELIEF," in *Proc. European Conference on Machine Learning (ECML)*, LNCS 784, 1994, pp. 171–182.
[31] R. J. Urbanowicz, M. Meeker, W. La Cava, R. S. Olson, and J. H. Moore, "Relief-based feature selection: Introduction and review," *Journal of Biomedical Informatics*, vol. 85, pp. 189–203, 2018.

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2024, 9(3)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

[32] T. K. Ho, "The random subspace method for constructing decision forests," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, no. 8, pp. 832–844, 1998.
[33] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.