

Article

Toward Automated Detection of Permanent Magnet Motors in WEEE Recycling Using Discriminative Transfer Learning

Niccolò Pezzati , Maurizio Guadagno , Lorenzo Berzi  and Massimo Delogu * 

Department of Industrial Engineering, University of Florence, Via di S. Marta 3, 50139 Florence, Italy; niccolo.pezzati@unifi.it (N.P.); maurizio.guadagno@unifi.it (M.G.); lorenzo.berzi@unifi.it (L.B.)

* Correspondence: massimo.delogu@unifi.it

Abstract

Rare Earth Elements (REEs) represent strategic and critical raw materials for the energy transition and must therefore be integrated into efficient and functional recycling processes. Their adoption in electric motors is rapidly expanding, raising significant challenges for end-of-life (EoL) management, starting from the collection phase. In this context, this work proposes the integration of an image-based classification framework within the Waste Electrical and Electronic Equipment (WEEE) recycling pipeline to selectively identify electric motors containing permanent magnets (PMs) and direct them toward dedicated recycling processes for rare earth recovery. The proposed methodology relies on a Discriminative Transfer Learning (DTL) approach based on a ResNeXt convolutional neural network (CNN), adapted to a proprietary and heterogeneous dataset of electric motors acquired in an industrial recycling facility. The objective is twofold: first, to identify motors containing PMs; second, to classify motors into construction categories according to their likelihood of incorporating PMs. Experimental results show promising performance in terms of PM-containing motor detection capability, establishing a robust foundation for the automated recovery of REEs at an industrial scale. Furthermore, the model's generalization capabilities can be further enhanced through the expansion of collaborative datasets and the integration of advanced scanning technologies.

Keywords: rare earth recovery; permanent magnets; functional recycling; end-of-life; image classification; discriminative transfer learning; convolutional neural networks

1. Introduction

The need to improve recycling and material recovery strategies for industrial products is increasingly emphasized by regulatory frameworks such as the European Green Deal [1] and the Critical Raw Materials Act (CRMA) [2]. These mandates explicitly call for optimized recovery processes, particularly for Critical Raw Materials such as Rare Earth Elements (REEs). The strategic importance of REEs stems not only from their geological scarcity but also from their ubiquitous applications across a wide range of sectors, from critical industrial infrastructure, such as automotive and energy, to widespread consumer goods, such as power tools. Electric motors represent a primary product containing these critical materials. They are significant consumers of copper, designated as a Strategic Raw Material, and, in particular, Permanent Magnets (PMs), which are predominantly composed of REEs, including Neodymium, Praseodymium, and Dysprosium. Due to their superior power density, which offers exceptional performance per unit of mass and volume, the adoption of rare earth PMs in electric motors has become widespread. This trend is no longer confined



Academic Editors: Jiangjian Xiao and Ahmed Abu-Siada

Received: 22 January 2026

Revised: 5 March 2026

Accepted: 13 March 2026

Published: 15 March 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

to high-power applications, such as electric vehicles or wind generators, but has permeated the consumer electronics market. Indeed, consumer devices, including commercially available power tools, now account for a significant share of demand; the home and consumer device sector currently holds the largest market share of PMs at 36.4% [3]. This proliferation is driven by a twofold need: to miniaturize products without compromising power or compliance with increasingly stringent energy-efficiency regulations. However, the upstream supply chain for these high-performance magnets poses significant challenges. The extraction and refining of REEs have a significant environmental impact and energy costs [4], as well as complex social impacts associated with their supply chains [5]. In contrast, comparative Life Cycle Assessment (LCA) studies show that recycling REEs results in a significant reduction in environmental impact compared to virgin material production [6,7]. Despite these advantages, there is a critical bottleneck in the End-of-Life (EoL) phase. Unlike large industrial motors, REE-containing components in household products are often not subject to detailed characterization and dismantling during their EoL stage. Instead, they risk being mixed with generic metal scrap and subjected to mass shredding, thus irreversibly destroying their intrinsic value. To mitigate this problem, it is essential to move towards “Functional Recycling”, a paradigm that aims to preserve the quality of virgin material so that, once recycled, it can perform its original function. However, the implementation of targeted recovery is impeded by the lack of standardized identification processes in large consumer recycling streams. Therefore, to ensure economic viability and compliance with CRMA targets, the development of efficient techniques for the classification and selective dismantling and recycling of PM-containing motors is essential.

Despite the extensive body of research on the subject, a gap remains between theoretical potential and industrial reality. Consequently, it is essential to develop more sophisticated separation and treatment techniques, together with dedicated collection infrastructure and specialized value chains [8], as well as implementing disassembly analysis tools from the very beginning of product design, in order to achieve effective industrial circularity [9].

In the industrial landscape, modern sorting and recycling plants are increasingly adopting automated recognition methods based on deep learning [10]. These approaches have proven effective in various sectors, including raw material separation [11] and specific applications such as electrochemical cell classification [12,13]. However, despite promising results in adjacent recycling sectors, the literature remains notably silent on established applications for the automatic classification of electric motors.

This study aims to fill this research gap. It evaluates the implementation of an automatic recognition method specifically designed to identify electric motors containing REEs, thereby contributing to the industrial enhancement of recycling processes for these critical materials. The samples adopted for the activity have been obtained by on-field investigations: dismantling yards and electric equipment repair workshops.

1.1. Literature Background

Since the breakthrough results achieved by AlexNet in 2012 [14] in image classification tasks, computer vision has undergone a radical paradigm shift. Indeed, there has been a transition from “feature-based” learning methods, in which recognition features were manually classified, to machine learning approaches in which deep neural networks autonomously learn relevant representations from data. In particular, the deployment of increasingly deep and complex Convolutional Neural Networks (CNNs) has enabled a significant performance increase across numerous computer vision tasks. A fundamental contribution in this domain was introduced by ResNet (Residual Network), which over-

came the difficulties associated with training very deep CNNs by using residual blocks. These blocks allow information to propagate via skip connections, thereby preventing performance degradation and facilitating the construction of significantly deeper networks compared to previous architectures [15]. However, this architecture has the limitation of requiring a large number of residual blocks to achieve optimal performance, resulting in increased model depth and a higher total number of parameters [16]. To address this complexity, the ResNeXt architecture was introduced; in this architecture, convolution operations within each residual block are split into parallel groups, allowing the network to retain increased expressive capacity while reducing its overall depth. The number of parameters per parallel block remains bounded thanks to a bottleneck layer, which reduces data dimensionality before filter application. This design enables greater representational capacity at a lower computational cost than traditional ResNet architectures; for this reason, ResNeXt is widely considered for specific learning objectives [17]. While architectures such as DenseNet may yield marginally superior performance, the associated increase in complexity renders them less suitable for the current study, particularly given the constraints imposed by data scarcity. Training a deep neural network from scratch generally requires very large datasets and substantial computational resources. For specific applications, this strategy is often unfeasible. A consolidated alternative is to use an established network pre-trained for a specific objective, namely, transfer learning. This approach is based on the following phases:

- Freezing the weights of deep convolutional layers. The weights of the initial convolutional layers are frozen. These layers act as generic feature extractors, capturing fundamental patterns such as edges, gradients, and textures that remain relevant across diverse image domains.
- Replacement of the final fully connected layers. The final fully connected layers are removed and replaced. The new classification layer is initialized and dimensionally adapted to match the specific number of classes in the target problem.
- Fine-tuning of the new layers. Initially, only the newly added layers are trained on the new dataset (Shallow Tuning). However, if the target domain differs significantly from the original source domain, it may be necessary to unfreeze and refine the convolutional layers adjacent to the fully connected block (Deep Tuning). This step allows the network to adapt its high-level feature representations to the specific characteristics of the new dataset [18].

Within modern detection pipelines, classification models are often integrated with object detection algorithms, such as Faster R-CNNs [19], which locate regions of interest and associate classes and coordinates with the identified objects. Alongside these so-called two-stage models, one-stage approaches have also been developed, such as RetinaNet [20] and YOLO [21]; these can simultaneously identify object positions and categories, offering lower inference times but generally lower accuracy.

Transfer learning methods find application in numerous sectors, including medical diagnostics [18,22], the food industry [23], and facial recognition systems [24]. In the context of recycling methods, such techniques are used for generic waste classification, benefiting from the availability of public datasets on macro-categories, such as TrashNet [25]. However, specific datasets for the classification and recycling of EoL electric motors containing PMs are currently unavailable, posing a significant limitation to the development and training of CNNs for this purpose. The most closely related application found in the literature concerns the recognition of electrochemical cells from electronic waste [12,13], often based on radiographic images with varying penetration capabilities. These enable the observation of internal components that would otherwise be occluded, thereby significantly improving recognition performance [26]. In this context, radiographic images are used not only to

identify cells but also to classify them by type, thereby directing them to specific recycling processes. With reference to the problem at hand, the objective of this paper is to lay the foundations for integrating computer vision techniques into the recognition and separation processes of electric motors containing PMs with high rare earth content using transfer learning methods.

Given the absence of public datasets specific to this domain, the results and subsequent analysis must be contextualized within the limits of the available data. The study relies on a dedicated image acquisition campaign conducted at an industrial Waste Electrical and Electronic Equipment (WEEE) recycling facility. This campaign yielded a heterogeneous dataset comprising images from 229 motor samples, captured from multiple angles.

1.2. Objectives and Work Structure

The objective of this work is therefore to implement and evaluate a semi-automatic recognition method for identifying electric motors with a high probability of containing rare earth PMs, using a transfer learning approach. Such motors can then be sent for dedicated treatment processes, such as a PM's selective dismantling or to a disintegration in a controlled hydrogen environment [8], which is particularly effective for recovering magnets in mechanically accessible and permeable systems [27]. Following a review of the literature presented above, the following work is developed in the following sections: Materials and Methods, Results, Discussion and Conclusion. Section 2 details the methodology employed for the image acquisition campaign and the subsequent motor recognition process. Specifically, a transfer learning approach was applied to the pre-trained ResNeXt50 architecture to address two distinct classification tasks:

- PM Probability Assessment. Identifying motors with a high probability of containing PMs.
- Motor-Type Classification. Classifying the motors into the four most commercially widespread categories.

Section 3 presents the performance evaluation of the trained CNN for both tasks using widely accepted accuracy metrics in the literature, highlighting the strengths and limitations of the model within the specific context of the case study. Finally, these results are discussed, drawing conclusions on the objectives achieved and possible future research in this field.


2. Materials and Methods

2.1. Dataset Acquisition and Motor Characterization

The image campaign acquisition has been conducted in a WEEE industrial facility, reporting a small dataset of heterogeneous images acquired from 229 electric motors. The images have been captured from multiple viewpoints to document diverse geometric and structural variations. The sample primarily originates from household appliances (e.g., washing machines, vacuum cleaners) and light mobility devices (e.g., electric scooter hub motors). Notably, the data reflects a market shift towards Brushless Direct Current (BLDC) Motor and Permanent Magnet Synchronous Motor (PMSM) architecture at the expense of traditional universal motors. This trend implies a growing demand for REEs, as these high-efficiency architectures are the most frequent adopters of PMs [28]. Table 1 reports the number of identified samples categorized by construction type, a task which is quite simple if done by a motor design expert examining the sample directly. The analyzed dataset primarily consists of small-to-medium-sized motors; representative examples of the collected samples are shown in Figure 1.

Table 1. Summary of identified motor categories, sample counts, and associated PM likelihood.

Type	Quantity	PM
AC synchronous (PMSM/BLDC)	94	Probably
Universal brushed	72	No
DC brushed	43	Unlikely
AC asynchronous	20	No

**Figure 1.** Representative examples of the collected motors: (a) Synchronous motor, (b) Universal motor, (c) DC brushed motor, (d) Asynchronous motor.

2.2. Two-Stage Detection Pipeline and Data Preprocessing

Based on the samples collected, the study pursues a hierarchical dual objective:

- PM Detection. Primary, to identify motors with a high probability of containing PMs.
- Typology Classification. Subsequently, to classify motors according to their construction typology, as specific categories inherently possess a higher probability of housing PMs. Consequently, each image was assigned to a target label corresponding to one of the four categories listed in Table 1.

Since the quality, quantity, and specific arrangement of PMs are strictly dependent on the motor's construction typology, this dual approach provides a qualitative framework for the method. Rather than relying on a simple binary inference, the model contextualizes the presence of PMs within the architectural typology of the motor.

The training process was conducted in two distinct phases:

- Region Identification. Fine-tuning of a simplified Faster R-CNN to detect the specific Region of Interest (RoI) containing the motor.
- Classification. Fine-tuning of a ResNeXt network to classify the motors within the proposed RoI. The dimensions of the detected regions were expanded by 30% to prevent the loss of boundary details.

Given the preliminary nature of this study and the limited availability of industrial samples, this two-stage pipeline was preferred over a single-stage architecture. This choice aimed to simplify the learning task by reducing the number of parameters requiring simultaneous optimization, which is a critical consideration when dealing with small-scale

datasets. Single-stage detection methods may be evaluated in future work if a larger dataset becomes available.

To facilitate the manual annotation of RoIs and the subsequent training of the Faster R-CNN for bounding box detection, a custom annotation tool was developed. As illustrated in Figure 2, this tool records user inputs and exports the data in compliance with the COCO syntax [29], generating a dataset directly compatible with the training requirements of the region proposal network.




Figure 2. Interface of the annotation tool used for manually defining motor bounding boxes.

To mitigate error propagation within the sequential pipeline, a conservative 30% expansion was applied to the detected RoIs. This margin minimizes localization bias and ensures the motor remains fully visible, preserving the critical boundary details necessary for accurate classification.

Once the RoIs have been identified, the second stage of the classification pipeline is initiated. In this phase, to effectively expand the dataset size and enhance training robustness, data augmentation techniques were employed. This process involved applying pseudo-random transformations to the identified RoIs. By generating diverse training examples, this approach mitigates the risk of overfitting and prevents the network from learning spurious patterns or shared background noise. The training framework was implemented using PyTorch (v. 2.6), a leading framework for deep learning in Python (v. 3.11.9). This platform leverages GPU acceleration via CUDA, enabling efficient parallelization of computational tasks. The main Jupyter notebook and the Python functions developed are provided in the Supplementary Materials (Files S1–S4).

This phase began with the normalization of the entire image set, based on the dataset's mean and standard deviation, to standardize the distribution of pixel values. For each of the four motor categories, the samples were partitioned into 80% for training and 20% for testing. To ensure a realistic evaluation of the model's generalization capability and avoid data leakage, the split was performed by randomly selecting motor indices rather

than individual images. This strategy ensures that multiple images of the same motor are assigned exclusively to either the training or the testing set, preventing the network from learning spurious correlations based on specific motor instances.

Subsequently, Dataset objects were constructed containing the file paths of images processed by the identification network, the target classes, and the specific transformations for data augmentation. For the training set, data augmentation was performed using the torchvision.transforms library to create two augmented copies of each original image through pseudo-random transformations.

Note that image cropping was intentionally excluded from the augmentation pipeline, as the specific RoIs had been previously identified and expanded by 30% to preserve boundary details. The specific transformations applied are summarized in Table 2.

Table 2. Dataset composition and data augmentation parameters.

Feature	Train Set	Test Set
Number of Copies	2	1
Total Number of Images	1408	169
Random Rotation	[0; 15]°	No
Resize	(600, 600) px	(600, 600) px
Color Jitter (B, S, C, H)	Yes, [0; 0.2]	No
Random Flip	Yes	No
Normalization	Yes	Yes

Finally, two Data-Loader objects were instantiated to manage image loading in mini-batches of 64 samples for train and test sets, presenting them to the model in random order at each training iteration. This approach is a well-established practice that ensures greater stability in the convergence process, albeit at the cost of slightly reduced computational speed [30].

2.3. Network Optimization and Discriminative Transfer Learning

ResNeXt-50 [31], a 50-layer deep network, was adopted as the reference architecture. This network represents an optimal trade-off between complexity (25 million parameters) and performance, achieving 81% accuracy on the ImageNet dataset [32]. Following the convolutional backbone, a Global Average Pooling layer is employed to condense the spatial feature maps. This layer computes the mean value for each convolution channel, resulting in a distinct feature vector that serves as input for the subsequent layers. Consequently, the original fully connected block was modified by incorporating two parallel classification heads:

- **PM Probability Head.** This fully connected block terminates in a single-neuron layer to estimate the likelihood of PM presence. The raw network output is processed through a Sigmoid activation function, which maps the result to a probability value in the range [0; 1].
- **Multi-class Typology Head.** It terminates in a four-neuron layer configured to classify the motor's construction typology, in accordance with the categories outlined in Table 1. A Softmax activation function is applied to generate a probability distribution across the classes, where the highest probability identifies the predicted output class.

The schematic representation of the ResNeXt adaptation for transfer learning is illustrated in Figure 3.




Figure 3. Architecture of the multi-task model based on ResNeXt-50, designed to estimate PM likelihood and classify motor typology. A bottleneck block of the backbone with a cardinality of 32 is showcased; Parallel paths are aggregated and combined with the input through a residual skip connection.

The training phases were conducted using the Adam optimizer [33], employing a composite loss function obtained as the sum of the binary cross-entropy loss (for PM detection) and the multi-class cross-entropy loss (for classification). A higher weight was assigned to the latter to emphasize the relevance of this task within the application context.

Both fully connected output networks were trained with a learning rate η of 1×10^{-3} and a weight decay λ of 1×10^{-5} . The weight update during training is defined by the following system:

$$L_{Adam}(\Theta_t) = 0.6 \cdot L_{Cross}(\Theta_t) + 0.4 \cdot L_{Binary}(\Theta_t) \quad (1)$$

$$L(\Theta_t) = L_{Adam}(\Theta_t) + \lambda \sum_{\theta_i \in \Theta} |\theta_i| \quad (2)$$

$$\Theta_{t+1} = \Theta_t - \eta \nabla L(\Theta_t) \quad (3)$$

where $L(\Theta_t)$ represents the overall loss function with respect to the set of parameters θ at iteration t . It consists of a weighted average of the two loss functions corresponding to the network's two tasks, calculated via the Adam optimization method, to which a regularization term on the parameters is added, scaled by the coefficient λ . Parameter updates occur gradually and are controlled by the learning rate η .

The training was executed in two sequential stages to ensure stable convergence and effective feature specialization, as summarized in Table 3:

- **Shallow Tuning.** In the initial phase, only the newly added classification heads were trained while the entire ResNeXt backbone remained frozen. This stage aimed to adapt the randomly initialized weights of the Multilayer Perceptron heads to the generic features extracted by the pre-trained backbone.
- **Deep Tuning.** During the subsequent phase, the final convolutional bottleneck block was unfrozen to allow for finer feature specialization. By employing a reduced learning rate, the network preserved the fundamental knowledge acquired from ImageNet while adapting its filters to the specific structural patterns and textures of the electric motor dataset. This practice of applying differentiated learning rates across network layers is defined as Discriminative Transfer Learning (DTL).

Table 3. Experimental settings for the training stages.

Feature	Shallow Tuning	Deep Tuning
Trainable Layers	PM head Multi-class Head	PM head Multi-class Head Last Bottleneck Block
Trainable Parameters	/	
	~780 thousand	~5.2 million
Learning Rates	Heads: 1×10^{-3}	Heads: 1×10^{-3}
	/	Last Bottleneck: 1×10^{-4}
Epochs	35	40

3. Results

The neural network was trained for 75 epochs, each consisting of 22 mini-batches of 64 images, for a total training time of approximately 92 min on an NVIDIA RTX 4060 GPU. The dataset was split into 80% for training and 20% for testing. The total number of images was increased compared to the initially available motors, thanks to photographs acquired from different angles and then augmented. To prevent data leakage, images of the same motor were not included in both the training and test sets, to avoid spurious correlations and ensure a realistic evaluation of the model's generalization capability. Given the limited dataset, the stability of the proposed method has been validated by repeating the experiment across 12 different random seeds for dataset partitioning. This process yielded an average accuracy of 0.80 for PM detection and 0.70 for typology classification, with performance fluctuations limited to approximately 6% from the mean value (detailed results are provided in the Supplementary Materials, Figure S2).

Most motors identified with a high PM probability were disassembled, and their Rare REE concentrations were later verified via Scanning Electron Microscopy (SEM). For the subset of samples not directly subjected to SEM analysis, the presence of PMs was inferred based on their constructive and architectural affinity with the validated units. This approach is consistent with the documented industrial trend toward adopting REE-based architectures in high-performance motors, as previously discussed in Section 1.

The network's performance was evaluated for both classification tasks and analysed in multiple sections as follows: first, the standard performance metrics widely adopted in the literature are introduced to establish a baseline for evaluation. Subsequently, the Receiver Operating Characteristic (ROC) and Precision–Recall (PR) curves are analysed. Following this, confusion matrices are examined to identify potential spurious correlations and to gain insight into class-specific misclassifications. Then, to further validate the model's decision-making process, an explainability analysis is provided, incorporating saliency heatmaps and quantitative metrics. Finally, a comparative benchmarking of the model's performance against baseline architectures is presented to validate the effectiveness of the proposed framework.

3.1. Model Performance Metrics

The classification performance of the model on the test dataset was evaluated using standard metrics commonly adopted in image classification, including the following.

- True Positive Rate (TPR) and False Positive Rate (FPR), defined as:

$$TPR = \frac{\text{true positives}}{\text{true positives} + \text{false negatives}} \quad (4)$$

$$FPR = \frac{\text{false positives}}{\text{false positives} + \text{true negatives}} \quad (5)$$

An image is considered positive when it is correctly assigned to its true class and negative otherwise. These metrics aim to evaluate the network's ability to detect samples with PM (TPR) and the probability of generating false alarms (FPR), respectively. The reciprocal of the FPR corresponds to the True Negative Rate (TNR), which expresses the model's ability to correctly identify samples without PM. In the case of binary classification, TPR and FPR values vary depending on the confidence threshold used to determine whether an image is classified as positive. A higher confidence threshold results in a cleaner sample of selected motors but increases the risk of excluding positive cases.

- Youden Index (Y). This parameter provides a criterion to identify the optimal trade-off between detecting positive (TPR) and negative ($1 - FPR$) samples. It is defined as:

$$Y = TPR + 1 - FPR \quad (6)$$

While this index can be used to determine the confidence threshold, it treats false positives and false negatives as equally weighted. In industrial applications, however, the costs of these errors are rarely symmetrical, especially when dealing with the recovery of high-value strategic materials.

- Precision (P). This metric measures the model's ability to correctly identify positive samples while minimizing false alarms:

$$P = \frac{\text{true positives}}{\text{true positives} + \text{false positives}} \quad (7)$$

This metric is particularly relevant in practical applications, as it indicates the purity of the identified positive samples, thereby reducing time and costs for downstream processing.

- F1 Score. This metric provides a balanced measure between precision and detection capability, evaluating the trade-off between highly precise but limited classification and a more complete but noisier classification. It is defined as:

$$F1 = 2 \frac{P \cdot TPR}{P + TPR} \quad (8)$$

In industrial applications, the score represents an optimal parameter for selecting the confidence detection threshold because it is directly aimed at balancing the need to minimize false positives while ensuring a high volume of elements is recovered. This metric aligns the model's performance with the economic and strategic requirements of REE recycling. Furthermore, the detection threshold can be strategically adjusted to favor either Precision or the quantity recovered, depending on specific industrial needs. Figure 4 shows the histogram of classification probabilities for the test sample. In this example, the confidence threshold was set to the optimal value of 0.38, corresponding to the maximum for both the F1 score and the Youden criterion.




Figure 4. Histogram of classification probabilities for the test image sample. The confidence threshold for assigning positives is set at the optimal value of 0.38.

The selection of 0.38 as the optimal threshold indicates that the model exhibits high confidence in identifying negative samples, while maintaining a more cautious approach toward positive identification.

3.2. Receiver Operating Characteristic and Precision–Recall Curves

The network performance and the determination of the optimal detection threshold were evaluated using ROC and PR curves. In the binary classification setting, the ROC curve represents the relationship between the *TPR* and the *FPR* as a function of the confidence threshold, with the optimal operating point identified as the one maximizing the Youden index Y . Conversely, the *PR* curve illustrates the trade-off between Precision and *TPR* (also referred to as Recall) as the same threshold varies. The ROC and PR curves corresponding to the binary classification task are reported in Figure 5. It can be observed that the optimal confidence threshold, determined consistently according to both the Youden index and the *F1* score, is equal to 0.38.




Figure 5. ROC and PR curves for the binary classification task of permanent magnet detection. The curves are obtained by varying the confidence threshold used for identification.

An analysis of the network behavior as a function of the confidence threshold shows that a significant number of images containing PMs are assigned intermediate, yet meaning-

ful, classification probabilities. Conversely, the number of false negatives remains limited, indicating a strong capability of the model to correctly exclude simple motor structures lacking PMs, which are generally easier to identify. For this reason, the adoption of a relatively low recognition threshold is justified, as it minimizes the risk of discarding motors that actually contain PMs, at the expense of a marginal reduction in overall precision.

ROC and PR curves are also reported for the multi-class classification task involving the four motor construction types (see Figure 6). In this configuration, the network assigns to each image a probability of belonging to each class and ultimately labels the image according to the class with the highest probability.




Figure 6. ROC and PR curves for the multi-class classification of motor construction types. The weighted average curve represents the mean values of the Youden index and F1 score across classes, weighted by the number of images in the validation dataset.

From the analysis of the results, the following observations can be drawn:

- Universal motors are easily identified due to the characteristic cylindrical shape of the stator and the presence of brushes and commutators.
- Synchronous motors exhibit good identification performance, owing to the presence of magnets with repeated geometries in the rotor and coils in the stator, which constitute distinctive visual features.
- Brushed DC motors are recognized with high precision but lower recall, likely due to partial occlusion of structural features (such as the commutator, brushes, and arc-shaped magnets) caused by casings that are difficult to remove during upstream shredding processes.
- Asynchronous motors show lower classification accuracy, plausibly due to their limited representation in the collected dataset. However, this limitation does not significantly affect the overall outcome, as this motor type is often misclassified by the network as a universal motor. Since both classes lack PMs and are therefore not subject to rare earth recovery processes, this confusion does not compromise the practical applicability of the proposed approach. This aspect is further discussed in the following section.

3.3. Confusion Matrices

The confusion matrix is a widely adopted indicator for assessing the model's ability to correctly classify images belonging to different categories. Column indices correspond to the predicted classes, while row indices represent the ground truth classes. The matrix values are normalized with respect to the total number of images in each class.

For the binary classification task, diagonal entries correspond to the TPR and the TNR, whereas off-diagonal elements represent their complementary quantities, namely the False Negative Rate (FNR) and the False Positive Rate (FPR), respectively. In the multi-class classification setting, diagonal elements indicate images correctly assigned to their corresponding class i , while off-diagonal values represent samples from class i that are erroneously classified as belonging to class j . In this way, the confusion matrix provides a clear visualization of the network's tendency to confuse specific classes.

From the results reported in Figure 7, it emerges that the model exhibits promising performance in the binary classification task, whereas more pronounced limitations are observed in the multi-class classification scenario.




Figure 7. Confusion matrices for the multi-class and binary classification tasks.

In particular, the recognition of brushed DC motors and asynchronous motors appears less accurate, likely due to the limited number of available samples for asynchronous motors and the difficulty in accessing distinctive structural features in brushed DC motors. Nevertheless, the network demonstrates high identification precision for the class AC sync grouping the two motor types that most frequently employ PMs (BLDC and PMSMs). Furthermore, the class of universal motors, which is highly represented in the analyzed sample, is easily recognizable and can therefore be reliably excluded during the selection process of motors containing PMs to be directed toward recycling operations. These results are reflected in the confusion matrix for the binary classification task, which shows encouraging values of 0.81 for both TPR and TNR. These findings do not yet guarantee full industrial robustness, thereby necessitating further validation through larger and more diversified datasets; however, for the case of waste motor processing, the ability to concentrate most machines supposed to have REE in a different stream from the other types could be sufficient. Examples of identification and classification outcomes are reported in Figure 8, highlighting that the network achieves better performance when motor features are more evident and clearly visible.

3.4. Model Explainability Analysis

The decision-making processes of deep neural networks often remain opaque due to their inherent black box nature. To ensure the model effectively recognizes distinctive morphological features, an explainability analysis was conducted. Gradient-weighted Class Activation Mapping (Grad-CAM) [34] provides qualitative insights by computing pixel-specific importance based on weighted feature activations. Figure 9 illustrates Grad-CAM saliency maps for four representative images.




Figure 8. Examples of motor region identification followed by binary and multi-class classification, compared against the ground truth. Green boxes represent correct evaluations for both tasks; orange boxes indicate an incorrect prediction in one of the two tasks.




Figure 9. Grad-CAM generated saliency maps for a representative image for each class.

The network correctly prioritizes key structural components, such as the cylindrical stator in universal motors and the rotor coils in synchronous machines. While the model successfully captures critical rotor features, some focus on background elements is occasionally observed. Notably, the DC brushed category exhibits lower precision in feature localization; this is likely because internal structural markers are often occluded by casings that remain intact during the upstream shredding process. Conversely, in asynchronous motors, the model effectively targets the visible segments of the internal stator architecture.

A quantitative assessment of model explainability was conducted using the Deletion and Insertion Area Under Curve (DAUC and IAUC) metrics, as proposed by Petsiuk

et al. [35]. This procedure evaluates the model's prediction score as salient pixels are systematically masked in descending order of importance (Deletion) or progressively added (Insertion). A rapid drop in the prediction score during deletion indicates that the model relies on highly significant pixels for its classification, whereas a rapid rise during insertion signifies a robust explanation. Figure 10 reports the mean curves for both the binary and multi-class tasks across all samples.




Figure 10. Mean DAUC and IAUC curves for binary and multi-class tasks.

Regarding binary inference, the score drops sharply toward 0.5 as the most salient pixels are occluded. This indicates that when permanent magnet characteristics are clearly visible, the model is highly confident. For partially masked images, the model reverts to a more uncertain state, relying on the motor's broader macro-topological context. In the multi-class classification task, the model generally exhibits superior identification of key features, evidenced by a major probability drop following the removal of only a few salient pixels, with scores trending toward a 0.25 baseline for the four-class scenario. These findings suggest that model explanation could be enhanced by integrating radiographic images, providing deeper and systematic access to internal topological structures.

3.5. Baseline Performance Comparison

The performance of the proposed model was evaluated against other standard image classification architectures to validate the effectiveness of the Discriminative Transfer Learning approach. Each baseline model was trained for the same number of epochs (75) to ensure an objective comparison. The study compared the ResNeXt DTL performance with Shallow Tuning versions of AlexNet, ResNet18, DenseNet, and ResNeXt itself. Additionally, a DTL version of DenseNet was implemented to assess the impact of fine-tuning the convolutional backbone. The experimental settings and comparative results are summarized in Table 4.

Table 4. Comparison of accuracy metrics and training time between the proposed ResNeXt DTL model and selected baseline architectures.

Model	PM Acc. (%)	Classification Acc. (%)	Training Time (min)
ResNeXt DTL	80.2	73.2	92
DenseNet DTL	79.1	71.9	84
AlexNet ST	76.3	62.8	39
ResNet18 ST	80.3	55.1	52
DenseNet ST	77.8	58.6	68
ResNeXt ST	80.2	62.8	80

To maintain a standardized comparison across all models, the confidence threshold for the binary PM detection task was fixed at 0.5; this accounts for the slight decrease in accuracy compared to the results obtained using the optimized threshold of 0.38 discussed in the previous section. While the results for the binary task remained relatively consistent across the various architectures, significant performance variations were observed in the multi-class classification task. The proposed DTL approach provided a measurable improvement in the network's ability to extract domain-specific features compared to ST, which only updates the fully connected layers.

DenseNet, when implemented with a deep tuning approach, represents a highly effective alternative to the ResNeXt architecture. This performance is attributed to its dense connectivity and feature reuse capabilities, which effectively mitigate the vanishing gradient problem and allow for stable deep tuning. Model performance comparisons should be further investigated with a larger dataset once available to enhance the statistical robustness of these findings.

4. Discussion and Conclusions

The application of transfer learning for adapting CNNs to specific recognition tasks is a well-established practice and has been widely adopted in the field of waste classification and automated separation to support selective recycling processes.

In this context, the present work proposes a Discriminative Transfer Learning approach for the recognition and separation of electric motors originating from WEEE streams, with the aim of facilitating the recovery of PMs containing REEs. The ResNeXt50 network was adapted to a proprietary dataset acquired at an industrial WEEE recycling facility and trained to perform two distinct tasks: (i) the classification of electric motors into the four most prevalent construction types on the market, and (ii) the identification of motors exhibiting structural characteristics typical of PM-based machines.

The obtained results demonstrate the model's ability to identify motors containing PMs with good accuracy, enabling the selection of a suitable subset of samples for dedicated downstream treatment, albeit with a limited overall precision. However, it is essential to emphasize that the discussed metrics, while promising, are derived from a small dataset; therefore, they represent a preliminary proof-of-concept, and their industrial robustness must be demonstrated in future applications. The optimization of the confidence threshold allowed for the definition of an effective trade-off between detection capability and precision, tailored to the requirements of selective recycling applications.

Overall, the analysis confirms the validity of the proposed transfer learning approach as a supporting tool for the automated separation of electric motors in the context of recycling, particularly in processes involving hydrogen-based disaggregation techniques or selective dismantling. The supporting tool proposed in this paper will serve to improve industrial recycling processes for electric motors using rare earth permanent magnets, improving early treatment immediately after collection and the first phase of dismantling (i.e., removal of the motors from the plastic or metallic casings, even with destructive methods). Further development activities are nevertheless required to promote the maturation and scalability of the proposed framework to validate industrial feasibility. Future work should focus on:

- The expansion and sharing of an EoL electric motor dataset, representative of a broader and more diverse range of motor types and operating conditions.
- The integration of radiographic images acquired at different penetration depths, aimed at improving the recognition of internal geometric patterns in cases where external casings limit the visibility of structural components, in analogy with approaches already explored in related domains such as electrochemical cell classification.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/machines14030331/s1>, Table S1: Metrics/Detection Metrics, Figure S1: Metrics/Model Performance Comparison, Figure S2: Metrics/Model cross-validation, File S1: trainingMultitask.ipynb, File S2: Functions/baseline.py, File S3: Functions/motordataset.py, File S4: Functions/training.py, File S5: Cross Validation.ipynb, Data S1: models_training.json, Data S2: robustness_results.json.

Author Contributions: Conceptualization, N.P., M.G. and L.B.; methodology, N.P.; software, N.P.; validation, N.P., M.G. and L.B.; formal analysis, N.P.; investigation, N.P., M.G. and L.B.; resources, N.P.; data curation, N.P.; writing—original draft preparation, N.P. and M.G.; writing—review and editing, N.P., M.G., L.B. and M.D.; visualization, N.P. and M.G.; supervision, L.B. and M.D.; project administration, L.B. and M.D.; funding acquisition, M.D. All authors have read and agreed to the published version of the manuscript.

Funding: Funded by the European Union under the HARMONY project (Grant agreement ID: 101138767 www.harmonyproject.eu/). However, views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or HADEA. Neither the European Union nor the granting authority can be held responsible for them. This work includes data obtained and processed as part of the “RAMITERA” project (Riciclare Motori Contenenti Terre Rare), funded by the Italian MASE—Ministry of the Environment and the Energy Security, under the 2021 Call for WEEE, Decree EC-DEC-85 of 7 September 2023 (Bando RAEE Edizione 2021).

Data Availability Statement: The original contributions presented in this study are included in the Supplementary Materials. Further inquiries can be directed to the corresponding author.

Acknowledgments: The authors would like to thank Giulia Cortina, Roberto Ciccarella, and Shaman Alwawi for their support in this research activity. The authors wish to thank also Seval srl (Colico, LC) for assistance in acquiring samples useful for this research.

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.

Abbreviations

The following abbreviations are used in this manuscript:

CRMA	Critical Raw Materials Act
REE	Rare Earth Element
PM	Permanent Magnet
LCA	Life Cycle Assessment
EoL	End-of-Life
CNN	Convolutional Neural Network
WEEE	Waste Electrical and Electronic Equipment
RoI	Region of Interest
BLDC	Brushless Direct Current Motor
PMSM	Permanent Magnet Synchronous Motor
DTL	Discriminative Transfer Learning
ST	Shallow Tuning
ROC	Receiver Operating Characteristic
PR	Precision–Recall
TPR	True Positive Rate
FPR	False Positive Rate
TNR	True Negative Rate
FNR	False Negative Rate

References

1. Communication from The Commission to The European Parliament, The European Council, The Council, The European Economic and Social Committee and The Committee of the Regions The European Green Deal 2019. Available online: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM%3A2019%3A640%3AFIN> (accessed on 20 November 2025).
2. Hool, A.; Helbig, C.; Wierink, G. Challenges and Opportunities of the European Critical Raw Materials Act. *Miner. Econ.* **2024**, *37*, 661–668. [\[CrossRef\]](#)
3. Permanent Magnets Market. Market.us. Available online: <https://market.us/report/permanent-magnets-market/> (accessed on 19 November 2025).
4. Mugion, R.G.; Elmo, G.C.; Ungaro, V.; Pietro, L.D.; Martucci, O.; Mugion, R.G.; Elmo, G.C.; Ungaro, V.; Pietro, L.D.; Martucci, O. A Systematic Literature Review of Selected Aspects of Life Cycle Assessment of Rare Earth Elements: Integration of Digital Technologies for Sustainable Production and Recycling. *Sustainability* **2025**, *17*, 5825. [\[CrossRef\]](#)
5. Werker, J.; Wulf, C.; Zapp, P.; Schreiber, A.; Marx, J. Social LCA for Rare Earth NdFeB Permanent Magnets. *Sustain. Prod. Consum.* **2019**, *19*, 257–269. [\[CrossRef\]](#)
6. Jin, H.; Afiuny, P.; McIntyre, T.; Yih, Y.; Sutherland, J.W. Comparative Life Cycle Assessment of NdFeB Magnets: Virgin Production versus Magnet-to-Magnet Recycling. *Procedia CIRP* **2016**, *48*, 45–50. [\[CrossRef\]](#)
7. Jin, H.; Afiuny, P.; Dove, S.; Furlan, G.; Zakotnik, M.; Yih, Y.; Sutherland, J.W. Life Cycle Assessment of Neodymium-Iron-Boron Magnet-to-Magnet Recycling for Electric Vehicle Motors. *Environ. Sci. Technol.* **2018**, *52*, 3796–3802. [\[CrossRef\]](#)
8. Cortina, G.; Guadagno, M.; Berzi, L.; Delogu, M. Environmental Impact and Recycling Routes of Rare Earth Elements in Permanent Magnets of Electric Machines for Industrial and Automotive Applications: A Systematic Review. *Eng. Proc.* **2026**, *Under Review*.
9. Guadagno, M.; Innocenti, E.; Berzi, L.; Corsi, S.; Delogu, M. Development of Procedures for Disassembly of Industrial Products in Python Environment. *Eng. Proc.* **2025**, *85*, 6. [\[CrossRef\]](#)
10. Fotovvatikhah, F.; Ahmedy, I.; Noor, R.M.; Munir, M.U. A Systematic Review of AI-Based Techniques for Automated Waste Classification. *Sensors* **2025**, *25*, 3181. [\[CrossRef\]](#) [\[PubMed\]](#)
11. Zhang, Q.; Yang, Q.; Zhang, X.; Bao, Q.; Su, J.; Liu, X. Waste Image Classification Based on Transfer Learning and Convolutional Neural Network. *Waste Manag.* **2021**, *135*, 150–157. [\[CrossRef\]](#) [\[PubMed\]](#)
12. Sterkens, W.; Diaz-Romero, D.; Goedemé, T.; Dewulf, W.; Peeters, J.R. Detection and Recognition of Batteries on X-Ray Images of Waste Electrical and Electronic Equipment Using Deep Learning. *Resour. Conserv. Recycl.* **2021**, *168*, 105246. [\[CrossRef\]](#)
13. Ueda, T.; Koyanaka, S.; Oki, T. In-Line Sorting System with Battery Detection Capabilities in e-Waste Using Combination of X-Ray Transmission Scanning and Deep Learning. *Resour. Conserv. Recycl.* **2024**, *201*, 107345. [\[CrossRef\]](#)
14. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. ImageNet Classification with Deep Convolutional Neural Networks. In *Proceedings of the 26th International Conference on Neural Information Processing Systems, Lake Tahoe, NV, USA, 3–8 December 2012*; Curran Associates Inc.: Red Hook, NY, USA, 2012; Volume 1, pp. 1097–1105.
15. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep Residual Learning for Image Recognition. *arXiv* **2015**, arXiv:1512.03385. [\[CrossRef\]](#)
16. Xie, S.; Girshick, R.; Dollár, P.; Tu, Z.; He, K. Aggregated Residual Transformations for Deep Neural Networks. *arXiv* **2017**, arXiv:1611.05431. [\[CrossRef\]](#)
17. Baker, N.A.; Zengeler, N.; Handmann, U.; Baker, N.A.; Zengeler, N.; Handmann, U. A Transfer Learning Evaluation of Deep Neural Networks for Image Classification. *Mach. Learn. Knowl. Extr.* **2022**, *4*, 22–41. [\[CrossRef\]](#)
18. Tajbakhsh, N.; Shin, J.Y.; Gurudu, S.R.; Hurst, R.T.; Kendall, C.B.; Gotway, M.B.; Liang, J. Convolutional Neural Networks for Medical Image Analysis: Full Training or Fine Tuning? *IEEE Trans. Med. Imaging* **2016**, *35*, 1299–1312. [\[CrossRef\]](#)
19. Ren, S.; He, K.; Girshick, R.; Sun, J. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *arXiv* **2016**, arXiv:1506.01497. [\[CrossRef\]](#)
20. Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; Dollár, P. Focal Loss for Dense Object Detection. *arXiv* **2018**, arXiv:1708.02002. [\[CrossRef\]](#)
21. Pandilova, E.; Petrov, M.; Spasev, V.; Dimitrovski, I.; Kitanovski, I. Transfer Learning with Yolo for Object Detection in Remote Sensing. In *ICT Innovations 2024. TechConvergence: AI, Business, and Startup Synergy*; Risteska Stojkoska, B., Janeska Sarkanjac, S., Eds.; Communications in Computer and Information Science; Springer Nature: Cham, Switzerland, 2025; Volume 2436, pp. 121–135, ISBN 978-3-031-86161-1. [\[CrossRef\]](#)
22. Madaan, V.; Roy, A.; Gupta, C.; Agrawal, P.; Sharma, A.; Bologna, C.; Prodan, R. XCOVNet: Chest X-Ray Image Classification for COVID-19 Early Detection Using Convolutional Neural Networks. *New Gener. Comput.* **2021**, *39*, 583–597. [\[CrossRef\]](#) [\[PubMed\]](#)
23. Xiao, Z.; Wang, J.; Han, L.; Guo, S.; Cui, Q. Application of Machine Vision System in Food Detection. *Front. Nutr.* **2022**, *9*, 888245. [\[CrossRef\]](#)
24. Kopalidis, T.; Solachidis, V.; Vretos, N.; Daras, P. Advances in Facial Expression Recognition: A Survey of Methods, Benchmarks, Models, and Datasets. *Information* **2024**, *15*, 135. [\[CrossRef\]](#)
25. TrashNet. Available online: <https://www.kaggle.com/datasets/feyzazkefe/trashnet> (accessed on 19 November 2025).

26. Li, Y.; Zhang, X. Intelligent X-Ray Waste Detection and Classification via X-Ray Characteristic Enhancement and Deep Learning. *J. Clean. Prod.* **2024**, *435*, 140573. [[CrossRef](#)]
27. Burkhardt, C.; van Nielen, S.; Awais, M.; Bartolozzi, F.; Blomgren, J.; Ortiz, P.; Xicotencatl, M.B.; Degri, M.; Nayeboossadri, S.; Walton, A. An Overview of Hydrogen Assisted (Direct) Recycling of Rare Earth Permanent Magnets. *J. Magn. Magn. Mater.* **2023**, *588*, 171475. [[CrossRef](#)]
28. Podmiljšak, B.; Saje, B.; Jenuš, P.; Tomše, T.; Kobe, S.; Žužek, K.; Šturm, S. The Future of Permanent-Magnet-Based Electric Motors: How Will Rare Earths Affect Electrification? *Materials* **2024**, *17*, 848. [[CrossRef](#)] [[PubMed](#)]
29. COCO—Common Objects in Context. Available online: <https://cocodataset.org/#format-data> (accessed on 15 December 2025).
30. Tian, Y.; Zhang, Y.; Zhang, H. Recent Advances in Stochastic Gradient Descent in Deep Learning. *Mathematics* **2023**, *11*, 682. [[CrossRef](#)]
31. Resnext50_32x4d—Torchvision 0.24 Documentation. Available online: https://docs.pytorch.org/vision/stable/models/generate_d/torchvision.models.resnext50_32x4d.html#torchvision.models.ResNeXt50_32X4D_Weights (accessed on 19 November 2025).
32. ImageNet. Available online: <https://www.image-net.org/index.php> (accessed on 19 November 2025).
33. Kingma, D.P.; Ba, J. Adam: A Method for Stochastic Optimization. *arXiv* **2014**, arXiv:1412.6980. [[CrossRef](#)]
34. Selvaraju, R.R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; Batra, D. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. *Int. J. Comput. Vis.* **2020**, *128*, 336–359. [[CrossRef](#)]
35. Petsiuk, V.; Das, A.; Saenko, K. RISE: Randomized Input Sampling for Explanation of Black-Box Models. *arXiv* **2018**, arXiv:1806.07421. [[CrossRef](#)]

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