microplastics

MDPI

Article

A Domain-Adaptive Deep Learning Approach for Microplastic Classification

Max Barker [1], Tanmay Singha [2], Meg Willans [3], Mark Hackett [3] and Duc-Son Pham [1,*]

[1] School of Electrical Engineering, Computing and Mathematical Sciences, Curtin University, Bentley, WA 6102, Australia; barker.maxw@gmail.com
[2] The University of Notre Dame, Fremantle, WA 6160, Australia; tanmay.singha@nd.edu.au
[3] School of Molecular and Life Sciences, Curtin University, Bentley, WA 6102, Australia; meg.willans@postgrad.curtin.edu.au (M.W.); mark.j.hackett@curtin.edu.au (M.H.)
* Correspondence: dspham@ieee.org

Abstract

Microplastics pose a growing environmental concern, necessitating accurate and scalable methods for their detection and classification. This study presents a novel deep learning framework that integrates a transformer-based architecture with domain adaptation techniques to classify microplastics using reflectance micro-FTIR spectroscopy. A key challenge addressed in this work is the domain shift between laboratory-prepared reference spectra and environmentally sourced spectra, which can significantly degrade model performance. To overcome this, three domain-adaptation strategies—Domain Adversarial Neural Networks (DANN), Deep Subdomain-Adaptation Networks (DSAN), and Deep CORAL—were evaluated for their ability to enhance cross-domain generalization. Experimental results show that while DANN was unstable, DSAN and Deep CORAL improved target domain accuracy. Deep CORAL achieved 99% accuracy on the source and 94% on the target, offering balanced performance. DSAN reached 95% on the target but reduced source accuracy. Overall, statistical alignment methods outperformed adversarial approaches in transformer-based spectral adaptation. The proposed model was integrated into a reflectance micro-FTIR workflow, accurately identifying PE and PP microplastics from unlabelled spectra. Predictions closely matched expert-validated results, demonstrating practical applicability. This first use of a domain-adaptive transformer in microplastics spectroscopy sets a benchmark for high-throughput, cross-domain analysis. Future work will extend to more polymers and enhance model efficiency for field use.

Keywords: microplastics; transformer; reflectance micro-FTIR spectroscopy; domain adaptation; cross-domain learning; DANN; DSAN; Deep CORAL

Academic Editor: Juan A. Conesa

Received: 7 July 2025
Revised: 9 August 2025
Accepted: 14 August 2025
Published: 1 October 2025

Citation: Barker, M.; Singha, T.; Willans, M.; Hackett, M.; Pham, D.-S. A Domain-Adaptive Deep Learning Approach for Microplastic Classification. *Microplastics* **2025**, *4*, 69. https://doi.org/10.3390/microplastics4040069

1. Introduction

Microplastics have emerged as a pervasive and alarming contaminant, infiltrating water bodies and environments throughout the world [1–4]. Their widespread presence is no longer just an ecological concern [5]. Mounting scientific evidence points to both direct and indirect threats to human health. These microscopic plastic fragments, along with the chemicals they carry or leach, have been detected in our food, drinking water, and even in human tissues and organs [6,7]. This growing body of evidence underscores an urgent global imperative for action [8–10].

A wide range of techniques are available for microplastic analysis [11,12]. Visual inspection methods, while common, are often inaccurate [13]. In contrast, chemical charac-

terisation approaches, such as Fourier transform infrared (FTIR) spectroscopy and electron microscopy, offer accurate and non-destructive identification of microplastics, though they tend to be expensive and time-intensive. These methods are considered among the most dependable for microplastic characterisation [14]. Infrared spectroscopy, in particular, enables non-destructive analysis of microplastic particles. However, spectroscopic methods also present certain limitations. For instance, Attenuated Total Reflectance FTIR (ATR-FTIR) allows for detailed sample analysis but involves a laborious and costly process [13]. Conversely, micro-FTIR spectroscopy enables the detection of particle sizes and polymer types via a non-contact chemical characterisation process [13]. Its non-contact nature contributes to its suitability for high-throughput analysis. Unlike ATR-FTIR, which requires direct contact between the sample and the ATR crystal—thus impeding automation—micro-FTIR can process entire microplastic sample sets automatically, without human interaction.

Micro-FTIR analysis can be conducted using two primary modes: transmission and reflectance. Transmission mode is most commonly employed in microplastic studies, as it tends to produce higher-quality spectral data than reflectance mode. However, it comes with certain limitations, including the need for costly components such as infrared-transparent windows and specialised filter papers. Additionally, transmission mode is not suitable for analysing larger microplastic particles [13]. Reflectance mode offers a more cost-effective and high-throughput alternative for microplastic characterisation. The workflow for reflectance-based micro-FTIR analysis is illustrated in Figure 1. Although first demonstrated successfully as a proof of concept [15], this approach has not seen widespread adoption. One major barrier to its broader use is the challenge posed by refractive index variations and light dispersion effects, which complicate the interpretation of the resulting spectra [13].

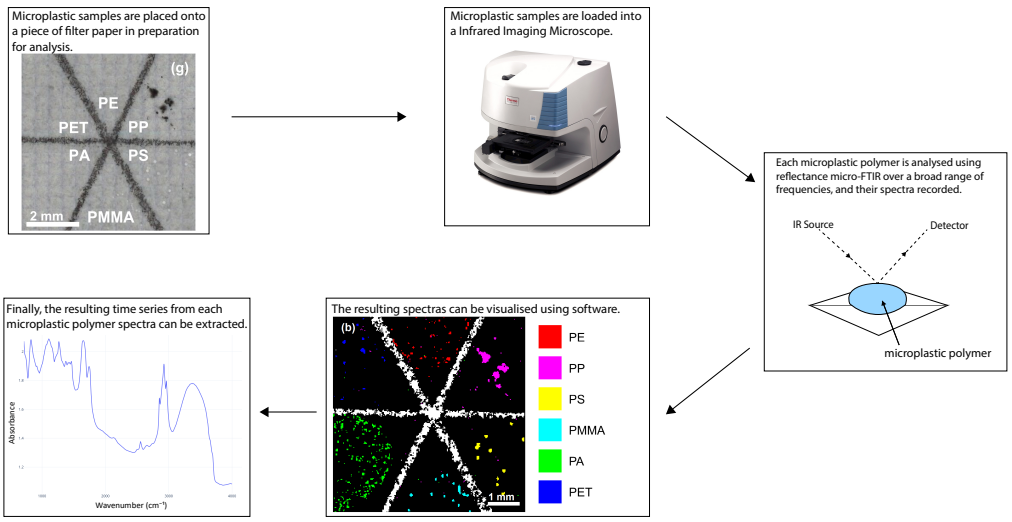

Figure 1. Reflectance micro-FTIR workflow—reproduced from [13].

The increasing threat posed by microplastics in aquatic and terrestrial ecosystems has motivated the development of automated and intelligent methods for microplastic detection and classification [16]. Modern deep learning approaches have been widely used in environmental monitoring tasks due to their ability to extract spatial features efficiently. For microplastic research, deep convolutional neural networks (DCNNs) have demonstrated high classification accuracy using electron micrographs and FTIR spectra. For example, CNN-based models have successfully identified polymer types and differentiated microbeads in wastewater samples, offering faster and more cost-effective solutions compared to traditional analytical techniques [17,18]. However, one major limitation of DCNNs is

their inherent inductive bias towards locality and translation invariance. Although these properties are useful for capturing local spatial features, they often fall short in capturing global context or long-range dependencies within the data, particularly when spectral or morphological characteristics vary over larger scales. This limitation becomes especially critical when analysing spectroscopic data where subtle patterns may span across different spectral domains.

In our previous work [19], we demonstrated the effectiveness of transformer-based models in accurately classifying microplastics, achieving competitive results on curated datasets. However, a critical challenge remains largely unresolved: the issue of *domain shift*. Most deep learning models, including transformers, are trained and evaluated on datasets that are sampled from the same or similar distributions. In real-world scenarios, however, microplastic samples can differ significantly across geographical locations, sample preparation protocols, imaging devices, and environmental conditions, leading to a distribution mismatch between the training (source) and deployment (target) domains. As a result, models trained in one domain often fail to generalise effectively when applied to a different domain.

To address this limitation, the field of domain adaptation has emerged as a promising approach [20]. Domain-adaptation is a subfield of transfer learning that focuses on adapting a model trained in one domain (source) to perform well in another domain (target), particularly when the target domain lacks labelled data. This is especially relevant for microplastic research, where obtaining labelled target data can be expensive and time-consuming. Domain-adaptation techniques are generally divided into shallow and deep approaches. Shallow methods align statistical distributions between domains and are lightweight but limited in complex settings. Deep methods, such as DANN [21], DSAN [22], and Deep CORAL [23], integrate adaptation into deep learning models to learn domain-invariant features at multiple levels of representation. Unsupervised domain adaptation (UDA), where only the source domain is labelled and the target domain remains unlabelled, is particularly appealing for microplastic classification due to the scarcity of labelled data in new domains. By incorporating UDA techniques into our transformer-based microplastic classification framework, we aim to improve model generalisation and robustness across diverse datasets and environmental conditions. In this work, we extend our previous study [19] to address the challenge of adapting microplastic classification models trained on reflectance micro-FTIR spectroscopy to new environments without relying on costly labelled datasets. We investigate several unsupervised domain-adaptation techniques, including DANN, DSAN, and Deep CORAL, integrated within a transformer-based architecture for microplastic classification using micro-FTIR spectroscopy. These methods are benchmarked on both synthetic and real-world datasets to assess their effectiveness in mitigating domain shift. Our objective is to develop a more robust and generalisable pipeline to support the large-scale deployment of microplastic classification systems across diverse environments. The remainder of this paper is structured as follows: Section 2 reviews the relevant literature on domain adaptation and examines key approaches in detail. Section 3 outlines our materials and methodology. Section 4 presents experimental results and discusses the effectiveness of our approach. Finally, Section 5 concludes with key findings and directions for future research.

2. Background

The increasing concern over the environmental impact of microplastics has underscored the need for accurate, scalable, and standardised classification methods. Traditional visual inspection techniques are labour-intensive and error-prone, while chemical characterisation methods such as Fourier-transform infrared (FTIR) spectroscopy provide

reliable and non-destructive analysis. Among these, micro-FTIR spectroscopy is particularly promising due to its ability to perform automated, high-throughput analysis of microplastic particles, determining both size and polymer composition. Unlike attenuated total reflectance FTIR (ATR-FTIR), which is contact-based and time-consuming, micro-FTIR operates in transmission and reflectance modes. While transmission mode offers high-quality spectra, it is expensive and unsuitable for larger particles. Reflectance mode is faster and more cost-effective but suffers from noisy and lower-quality spectra due to refractive and dispersion effects.

To mitigate these limitations, researchers have increasingly turned to machine learning approaches. Techniques such as partial least squares discriminant analysis (PLS-DA) have shown that polymers like polyethylene (PE) and polypropylene (PP) can be accurately classified using reflectance-mode FTIR spectra [13]. Building on these successes, deep learning approaches—particularly convolutional neural networks (CNNs)—have demonstrated impressive results. For instance, CNN-based models have achieved 98.33% accuracy on electron micrographs of microplastics, 89% accuracy in classifying microbeads in wastewater samples, and up to 99% accuracy in classifying FTIR spectra of PE and PP polymers.

Beyond CNNs, the emergence of transformer neural networks has revolutionised sequence modelling. Transformers outperform recurrent neural networks (RNNs) and long short-term memory (LSTM) models by processing input in parallel and leveraging multi-head self-attention mechanisms to capture long-range dependencies. Originally developed for natural language processing (NLP), transformers have been successfully adapted for time-series classification, forecasting, and anomaly detection. For classification tasks, the encoder is used to extract robust feature representations. Advanced transformer variants, such as two-tower architectures focusing on timestep-wise and channel-wise attention, have achieved state-of-the-art results in multivariate time-series tasks. These models have also outperformed both CNNs and RNNs in challenging domains, such as raw optical satellite time-series classification. More recently, a transformer model has been developed specifically for microplastic classification [19] to further improve performance across different datasets (see Table 1).

Table 1. Comparison between transformer, kNN and Random Forest models on the marine and standard testing sets—reproduced from [19].

Model	Marine Accuracy	Marine F1	Std. Accuracy	Std. F1
Transformer	98.7%	98.7%	92.3%	92.0%
CNN	98.7%	92.6%	91.2%	91.0%
KNN	94.8%	92.6%	88.0%	84.6%
Rand. Forest	97.4%	94.5%	68.0%	44.6%
Naive Bayes	98.7%	98.7%	80.0%	78.9%
LDA	97.4%	97.4%	58.8%	54.2%
Logistic Reg.	94.8%	94.7%	52.9%	19.9%

2.1. Domain Adaptation

To overcome these challenges, domain-adaptation techniques have gained attention. These approaches enable models trained on one data distribution (source domain) to generalise to a different distribution (target domain), even in the absence of target labels. The goal of many neural networks is to minimise an objective function by learning from a training dataset in such a way that the network can accurately and confidently make predictions on an unseen testing dataset [20]. Often, the testing data is not reflective of real-world unseen examples. This is because training and testing datasets typically come from the same distribution. When exposed to unseen data from a different distribution,

a machine learning model that once performed very well on the test set from its own distribution may not be able to generalise to unseen data.

Figure 2 illustrates an example of domain adaptation. In this figure, the dark and light stars represent the same semantic class but originate from different domains, resulting in distinct distributions. The same applies to the blue and yellow polygons. An optimal classifier trained on the source domain may perform well there but could make more errors on the target domain due to domain shift. However, domain adaptation, as shown in the bottom subplot, can optimally adjust the decision boundary, leading to improved classification performance.

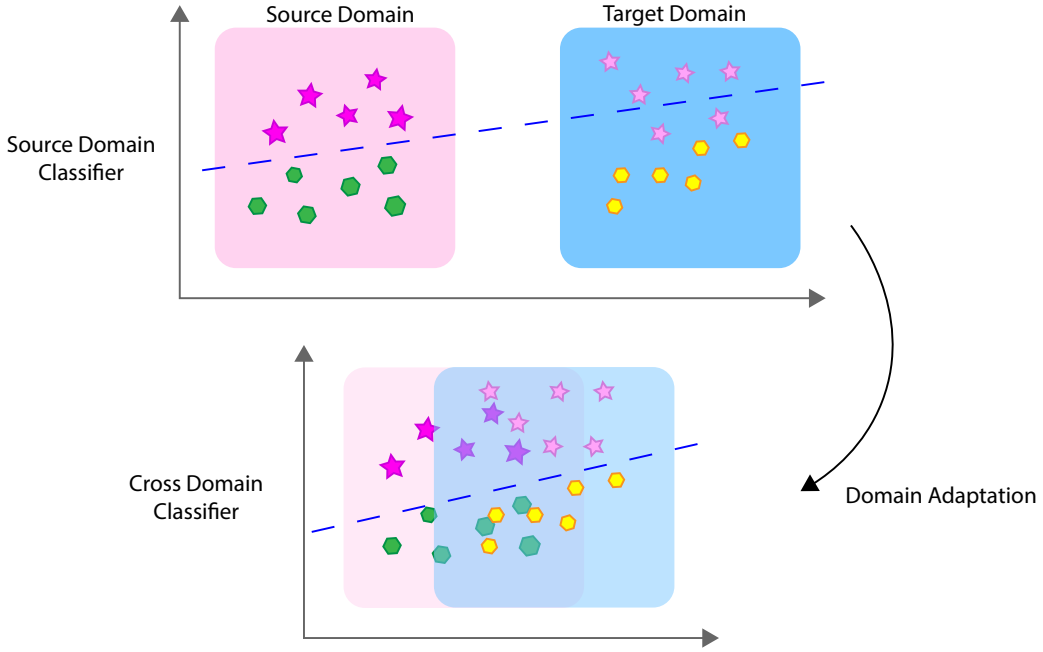

Figure 2. Illustrated example of domain adaptation. A source classifier may struggle to generalise to the target distribution, but by aligning the distributions generalisation may be achieved through a cross-domain classifier.

Domain adaptation is a field of transfer learning [24]. Domain-adaptation techniques can be broadly categorised into two main families:

- Shallow Domain-Adaptation Techniques: These methods operate by statistically aligning the distributions of the source and target domains. Two popular techniques in this category include:
  - Maximum Mean Discrepancy (MMD): A kernel-based method that measures the distance between distributions in a reproducing kernel Hilbert space (RKHS). It can be used to align the global or class-conditional distributions across domains.
  - CORrelation ALignment (CORAL): A method that aligns second-order statistics (covariance) of source and target features by minimising the Frobenius norm of their covariance difference. These techniques are computationally lightweight and can be incorporated into traditional machine learning pipelines but may struggle with complex feature spaces or nonlinear discrepancies.
- Deep Domain-Adaptation Techniques: These approaches integrate domain-adaptation objectives within deep learning models and can align feature representations at multiple abstraction levels. Key examples include:
  - Domain Adversarial Neural Networks (DANN): Employs a feature extractor shared between domains and a domain classifier trained adversarially using a

gradient reversal layer (GRL) [21]. This encourages the feature extractor to learn domain-invariant representations by confusing the domain classifier.

- Deep Subdomain-Adaptation Networks (DSAN): Instead of aligning entire domain distributions, DSAN [22] focuses on subdomain-level alignment, ensuring that class-specific distributions in the source and target domains are well matched using a modified version of MMD called Local MMD (LMMD).
- Deep CORAL: An extension of CORAL into the deep learning setting, Deep CORAL [23] aligns the covariance of source and target features at different layers of a deep network. It can be integrated with various architectures, including CNNs and transformers.

Next, we examine these three approaches in more detail.

2.2. Domain Adversarial Neural Networks

Domain adversarial neural networks (DANN) [21] are a type of unsupervised deep domain-adaptation architecture. The architecture is built in such a way that it can be implemented into any deep learning model. Figure 3 demonstrates the architecture. The green feature extractor can be represented as any neural network, be it a convolutional neural network, multilayer perceptron or transformer neural network. This architecture requires a labelled source set and an unlabelled target set. Each set is assigned a domain label representing the distribution it is part of. These labels are used to train the domain classifier.

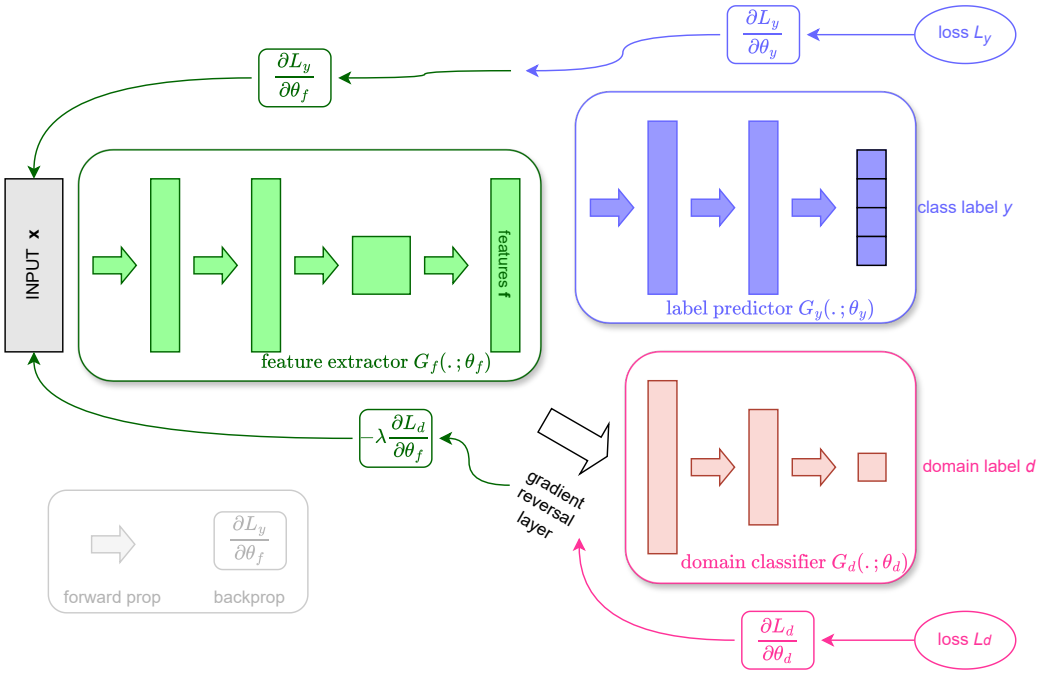

Figure 3. DANN architecture (re-drawn from [21]). The green feature extractor and blue label predictor create a standard feed-forward model. The domain classifier in pink represents a feed-forward neural network to learn domain-specific features.

The labelled source dataset will first flow through the network's feature representation stage, and then into the label predictor where the class labels will be predicted. During this stage, the features are also sent through the domain classifier. The domain classifier's role is to predict the domain of the features. So for the source set, the features are assigned a domain label of 0 and the target set a domain label of 1. The computed domain loss will then act as a confusion loss, as during backpropagation the loss will go through the gradient reversal layer which will flip the gradient. The target set goes through a similar process,

with the class label prediction stage being ignored. The idea of the gradient reversal is to cause confusion in the network, by forcing it to learn features that are not domain specific.

DANN is a promising architecture as it has successfully been applied in time-series domain adaptation using a 1D-CNN [25]. The architecture has improved classification accuracies for human activity recognition (HAR) sensor data by 13.8% (60.26–76.42%) and in sleep stage classification (SSG) by 4.9% (55.35–60.26%) [25]. DANN has also successfully been applied in computer-vision domain adaptation, improving classification accuracies on the MNIST-M dataset by 24.4% (52.25–76.66%).

2.3. Deep Subdomain Adaptation

Deep Subdomain Adaptation (DSAN) [22] is a statistical approach towards unsupervised domain-adaptation that utilises maximum mean discrepancy (MMD) as a means to align source and target domains. MMD is a kernel-based method that determines the distance between two distributions. In the case of DSAN, a modification of MMD is proposed that calculates the difference between the kernel mean embeddings of subdomains while considering the weight of the different samples. This is known as Local Maximum Mean Discrepancy (LMMD).

Figure 4 demonstrates the ultimate goal of DSAN and LMMD loss. Rather than simply aligning the domains, the goal of DSAN is to align the subdomains within each domain. This is carried out because global domain-adaptation may lose some of the fine-grained information contained within the subdomains [22]. DSAN has seen successful applications in computer-vision applications. One example is in the classification of the SVHN image dataset [26] when trained using the MNIST dataset [27]. Using DSAN, a classification performance increase of 30% was achieved, from 60.1% to 90.1% [22]. DSAN has also shown promising results in time-series datasets. In HAR classification, DSAN achieved a performance increase of 16.1% (60.26–76.42%). DSAN improved classification performance in SSD classification by 5.2% (55.35–60.57%).

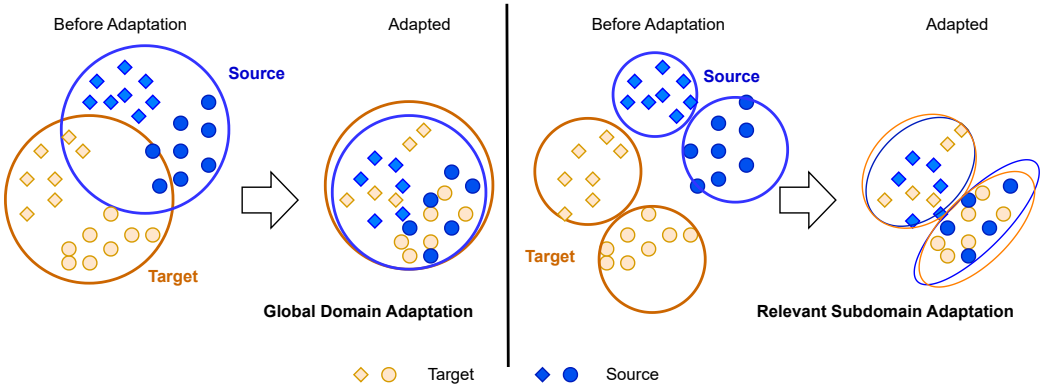

Figure 4. Demonstration of the differences between global domain adaptation (**left**) and subdomain adaptation (**right**) achieved by the DSAN architecture. The aim is to capture fine-grained information contained within the subdomains as opposed to losing this information through global domains (re-drawn from [22]).

2.4. Deep CORAL

A deep learning approach based on CORrelation ALignment (Deep CORAL) [23] is an unsupervised domain-adaptation method that attempts to align feature distributions between source and target domains. CORAL was initially created as a simple method of domain adaptation that can be integrated into any machine-learning model [28]. It works by using second-order statistics, aligning distributions by calculating the Frobenius norm

between the covariance of the source and target distributions (CORAL loss). Deep CORAL is an extension of this method for integration in deep learning models.

Figure 5 demonstrates the original implementation of the Deep CORAL architecture. The implementation uses a shared AlexNet architecture to compute classification and CORAL loss. CORAL loss is calculated with the feature representations from AlexNet when running the source and target datasets through the network. Deep CORAL is not bounded to the AlexNet architecture and is designed to integrate it into other deep learning architectures.

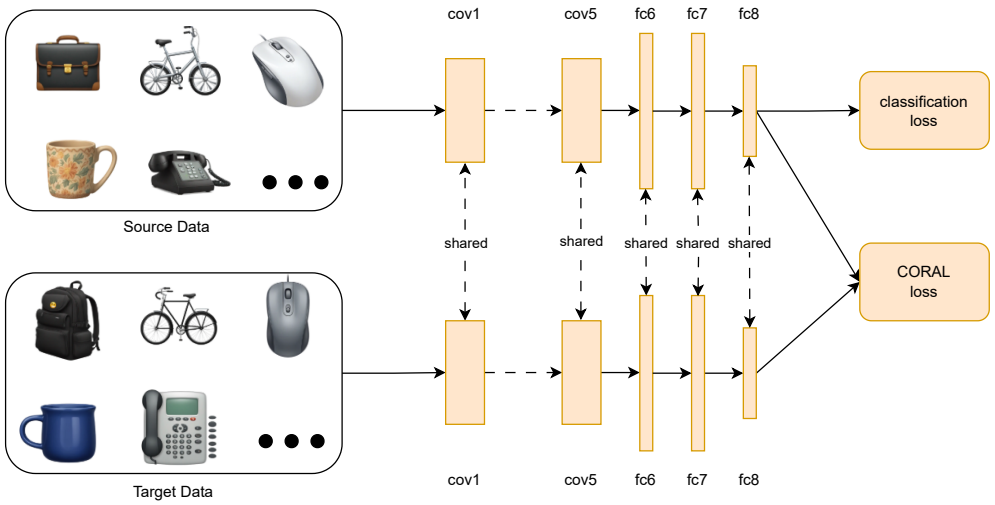

Figure 5. Deep CORAL architecture demonstrating the use of a shared feature extractor and classifier to calculate classification and CORAL loss when trained using labelled source and unlabelled target data (re-drawn from [23]).

While deep learning has shown high accuracy in microplastic classification, the challenge of domain generalisation remains underexplored. Combining transformer-based models with advanced domain-adaptation techniques such as DANN and DSAN offers a promising solution. Transformers excel at capturing long-range dependencies, while domain adaptation addresses distribution shifts between environments—together enabling scalable, standardised, and generalisable models for microplastic classification across diverse real-world datasets.

3. Materials and Methods

Three unsupervised domain-adaptation methods were implemented and tested to explore domain adaptation's impact on model generalisability. The domain-adaptation architectures utilised were DANN [21], Deep CORAL [23] and DSAN [22]. These three methods were selected as they represent two common domain-adaptation techniques, namely divergence-based and adversarial-based methods. The divergence-based methods include Deep CORAL and DSAN and take a statistical approach to align the distributions during training. Adversarial domain adaptation is the technique that DANN employs, where rather than a statistical approach, DANN attempts to align the distribution by forcing the network to learn domain invariant features through a separate domain prediction network. The goal is not only to improve the performance of the model but also to see whether divergence-based methods or adversarial-based methods are most appropriate. There is no one definitive best method of domain adaptation because each application will perform differently depending on the data and the architecture. The selection of these methods was also informed by a time-series domain-adaptation review in which they were all shown to have great performance increases on accelerometer data [25]. These methods

will be compared and contrasted with the output performance metrics. A portion of the standard polymer dataset will be reserved for final testing after training the model with domain adaptation to verify the performance.

This paper makes the following contributions: (1) an implementation of several domain-adaptation methods suitable for microplastic classification; and (2) an evaluation of these domain-adaptation methods on collected data for making the final recommendation on the integration with the reflectance micro-FTIR workflow. To the best of our knowledge, this work is the first to consider domain adaptation for microplastic classification.

3.1. Domain Adversarial Neural Networks

The DANN [21] architecture was implemented directly into the transformer architecture by adding a feed-forward neural network for domain classification. Two cross-entropy loss functions are used for class and domain classifications. The source domain dataset was allocated the source domain labels of 0.0, and the target domain was allocated the labels 1.0. During each training iteration, the gradient reversal intensity value λ was determined based on the current epoch, controlling the strength of the domain confusion. λ is calculated each epoch given by the following:

$$p = \frac{(batch\_index + current\_epoch \times len(dataloader))}{total\_epochs \times len(dataloader)}$$

and,

$$\lambda = \frac{2.0}{1.0 + e^{-10 \times p}} - 1.$$

Class labels, source and target domain labels are all predicted during training. Each domain prediction has its loss calculated and reversed via the gradient reversal layer. The final loss value is calculated as the sum of the source loss, domain loss and class prediction loss. This architecture's simple design meant minimal complexity was added to the transformer model; thus, no changes had to be made to the model's hyperparameters. Validation and standard dataset performance were recorded in terms of accuracy and loss. Finally, upon completion of model training, the marine and standard test sets were used to evaluate performance.

3.2. Deep Subdomain-Adaptation Networks

The implementation of DSAN utilises the transformer's encoder layer for feature extraction and a feed-forward network for classification. These are shared layers in that they are used for source and target learning and classification. An additional module was created for the calculation of the LMMD loss. During forward propagation, the model extracts the feature representations for both the marine (source) and standard (target) datasets. Once the features are extracted, the marine features are passed into the classifier, returning the resulting predictions. Then, the target features are fed into the classifier with the target predictions saved. Unlike DANN, the target and source domains are not given domain labels in this method. Similar to DANN, the model never receives access to the target labels. Then, with the predictions for both datasets and the marine dataset labels, the LMMD loss is computed as follows:

$$d_{\mathcal{H}}(p, q) \triangleq \mathbf{E}_c \left\| \mathbf{E}_{p^{(c)}}[\phi(\mathbf{x}^s)] - \mathbf{E}_{q^{(c)}}\left[\phi\left(\mathbf{x}^t\right)\right] \right\|_{\mathcal{H}}^2, \quad (1)$$

where $\mathbf{x}^s$, $\mathbf{x}^t$ are the instances in the source and target domains, $p^{(c)}$ and $q^{(c)}$ are the distributions in the source and target domains, respectively. This measures the discrepancy of local

distributions. By minimising the equation, the distributions of the relevant subdomains are drawn closer, capturing fine-grained information specific to each subdomain.

Once the network has computed the LMMD loss, the model will return both LMMD loss as well as the predictions the model made on the source/marine polymer dataset. LMMD loss was used in the training loop with the marine polymer prediction loss to train the model to learn domain invariant features during backpropagation. LMMD loss is gradually strengthened to prevent learning degenerate features. During training, the validation and standard dataset accuracies and loss are recorded. Upon completion of each training iteration, the model is given a final marine and standard polymer test sets to evaluate its final performance.

3.3. Deep CORAL

The implementation of Deep CORAL was of similar complexity to DSAN. It involved using a shared transformer for feature extraction and a linear layer for marine/source classification. The marine and standard datasets are passed through the feature extractor during forward propagation. The resulting features are then used to calculate the coral loss. First, the covariance for the source and target sets are calculated:

$$C_S = \frac{1}{n_S - 1}\left(D_S^\top D_S - \frac{1}{n_S}\left(\mathbf{1}^\top D_S\right)^\top \left(\mathbf{1}^\top D_S\right)\right) \tag{2}$$

and,

$$C_T = \frac{1}{n_T - 1}\left(D_T^\top D_T - \frac{1}{n_T}\left(\mathbf{1}^\top D_T\right)^\top \left(\mathbf{1}^\top D_T\right)\right), \tag{3}$$

where $\mathbf{1}$ represents a column vector with all elements equal to 1. Finally, CORAL loss can be calculated as follows

$$\ell_{CORAL} = \frac{1}{4d^2}\|C_S - C_T\|_F^2. \tag{4}$$

Here, $\|C_S - C_T\|_F^2$ represents the Frobenius norm between source and target convariances. The calculation of CORAL loss occurs during forward propagation and is returned along with the source/marine dataset predictions. Similar to DSAN and DANN, the influence that coral loss has on the total loss is increased over time via the computed λ for controlling the strength of coral loss:

$$\ell = \ell_{CLASS} + \lambda \ell_{CORAL}. \tag{5}$$

The accuracy and loss were recorded during model training for validation and standard datasets. Standard and marine test sets were employed to verify each model's performance after training.

3.4. Analysis on Spectral Data

Creating a robust deep learning model is part of a larger puzzle in the search for a unified protocol for microplastic analysis. To fully realise the potential of the transformer, the model will be tested with complete spectral files collected using reflectance micro-FTIR. This analysis will examine how well the transformer model integrates with spectral data and assists researchers in microplastic analysis. The process will emulate a real-world scenario where the model predicts polymer types from spectra with a different distribution than the training data, providing researchers with a visual guide to interpret the original spectral data alongside the model's predictions.

3.5. Datasets

The datasets used in this study were collected as part of [13] and comprise microplastic polymers polyethylene (PE) and polypropylene (PP) from two distinct domains. The first

is the marine polymer dataset, containing environmental microplastic spectra collected from marine saltwater, while the second is the standard polymer dataset, representing a non-environmental domain. The standard dataset was compiled from commercially sourced samples obtained from various suppliers and, in some cases, manually created by filing down larger plastic materials. Unlike environmentally derived microplastics, these were intentionally produced and thus represent a controlled reference domain. Both datasets were acquired using reflectance micro-FTIR spectroscopy, a high-throughput and non-destructive analytical technique that records the absorption of infrared radiation by polymers. This approach is consistent with the previous literature [17], as the collection of environmental microplastic samples can be costly and time-consuming, and publicly available datasets remain limited.

The marine polymer dataset contains 512 spectral samples of PE and PP polymers sourced from marine saltwater environments, along with additional polymer types including polystyrene (PS), polymethyl methacrylate (PMMA), polyamide (PA), and polyethylene terephthalate (PET). For this study, only the PE (130 samples) and PP (122 samples) spectra were used. The dataset was split into training, validation, and testing subsets, with 70% of the samples were allocated for training and the remaining 30% equally divided between validation and testing. This dataset serves as the source domain for the domain-adaptation experiments.

The standard polymer dataset comprises 169 spectral samples, consisting of 89 PE and 80 PP polymers. These samples were purchased as standard references and were not collected from the environment. The dataset is used exclusively as the target domain to evaluate the trained model's ability to generalise to samples from a different distribution. In this study, the marine dataset supports model training and in-domain testing, while the standard dataset is used to assess cross-domain generalisation performance.

4. Experimental Results

4.1. Domain Adversarial Neural Networks Implementation

Figure 6 demonstrates the best-performing model regarding validation accuracy and loss over 20 epochs of training. Throughout training, the gradient reversal layer intensity gradually increases, influencing the degree of domain confusion. By the 15th epoch, maximal domain confusion is applied with a gradient reversal λ above 0.99. As λ increases, the training accuracy decreases and loss increases, indicating that the model is being forced to learn domain-invariant features despite the hindering effect of domain confusion. Figure 7 shows similar trends in the standard dataset accuracies. The accuracy and loss patterns in the standard dataset closely follow those seen in the marine validation set, with both showing an initial decrease in performance followed by a sudden upturn as the model adapts to the domain confusion.

While this result is promising, not all training iterations were successful. Figure 8 demonstrates model training when domain confusion was unsuccessful. The hyperparameters for this experiment mimic the successful experiment, yet vastly different results are observed. In this iteration, the domain confusion has inhibited the model from learning from the training set due to degenerate features. The result is a model that cannot classify samples from either dataset. Table 2 displays the recorded results of the experiments thus far. For the best performing iterations seen in Figures 6 and 7, overall accuracies of 99% and 92% were achieved by the transformer model using DANN domain adaptation. These results are identical to the performance achieved during model tuning observed in [19], demonstrating that DANN could not make any meaningful improvements to the model. The failed training iteration observed in Figure 8 is also recorded. Precision and recall values of 0.00 for the classification of PE in both marine and standard test sets indicate that

the model is heavily biased towards making PP classifications, with this bias so extreme that no correct PE predictions were made. This could result from the model learning degenerate features, causing it to favour one class randomly. The failed iterations represent the majority of training iterations, as fewer successful iterations were observed. As a result, this method is not ideal for predicting unlabelled data as it is far too unreliable. Consistent results are necessary for domain adaptation to be successful.

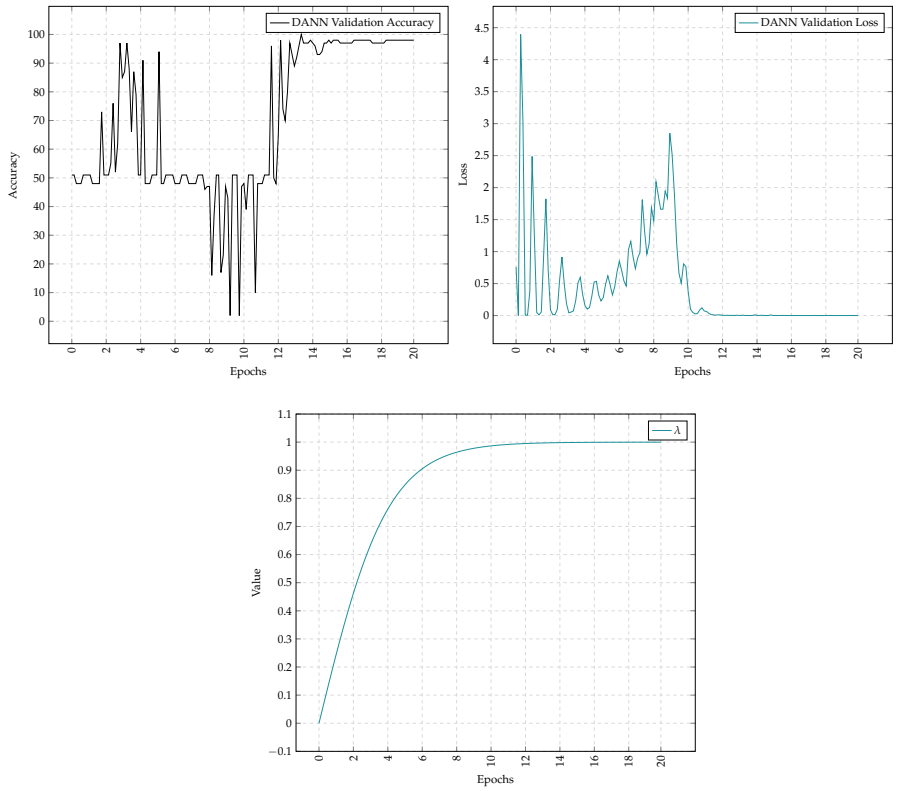

Figure 6. Validation accuracy, loss, and gradient reversal intensity λ over the training process. Values are logged per batch (150 points), with the x-axis scaled to represent a 20-epoch training duration.

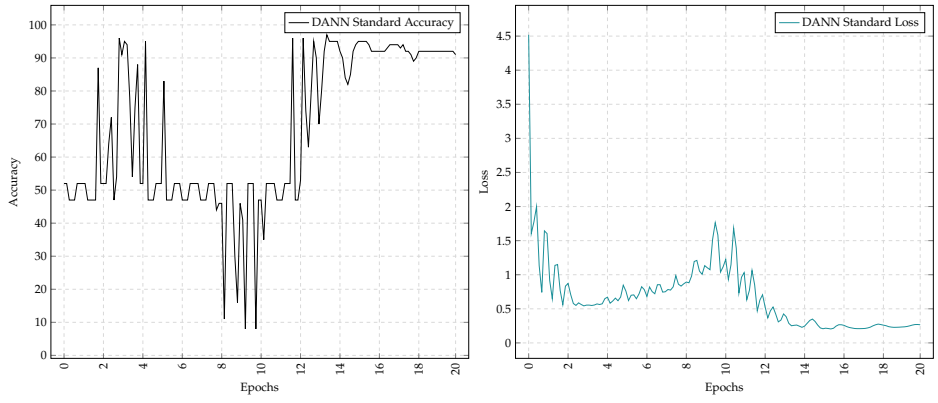

Figure 7. Standard accuracy and loss using DANN domain adaptation.

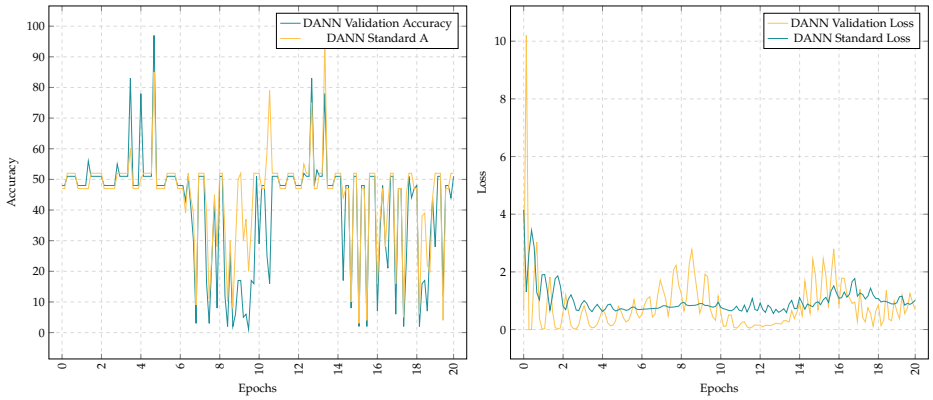

Figure 8. Accuracy and loss for both validation and standard sets when influenced by degenerated features.

Table 2. Marine and standard polymer test set accuracies when using DANN domain adaptation with respect to the best performing model and a failed model.

		Testing Set				Standard Testing Set			
	Polymer	**Prec.**	**Rec.**	**F1**	**Acc.**	**Prec.**	**Rec.**	**F1**	**Acc.**
Best	PE	97%	100%	99%	99%	87%	99%	92%	92%
	PP	100%	97%	99%	**99%**	99%	87%	92%	92%
Failed	PE	0%	0%	0%	52%	0%	0%	0%	53%
	PP	52%	100%	34%	52%	53%	100%	69%	53%

4.2. Deep Subdomain-Adaptation Networks Implementation

A significant change had to occur in the model's hyperparameters for the DSAN experiments to be successful. The DSAN model requires more GPU memory due to the space required when calculating the LMMD loss. Additional hyperparameter tuning was conducted to compensate for the memory requirements. The new parameters are as follows:

- Epochs: 5
- Batch size: 32
- Learning rate: 10^{-4}
- Encoding layers: 2
- Multi-headed attention layers: 4

The reduction in the epochs and the batch size meant that the model took shallower steps during gradient descent. Memory could have been conserved by reducing the encoding and multi-headed attention layers, but this would mean that the transformer is not utilised to its full capacity.

The resulting marine validation and loss values are shown in Figure 9. Unlike the results from DANN observed in Figure 6, this figure displays a much more conventional training curve. The adjusted hyperparameters are acceptable as the model has achieved satisfactory performance on the validation set with a final validation accuracy and loss of 98.7% and 0.07, respectively. The successful performance on the validation set also extends to the performance of the standard dataset, as seen in Figure 10. The loss and accuracy curves follow similar trends to those shown in Figure 6. These results show that statistical approaches to distribution alignment may be more appropriate for domain adaptation than adversarial approaches.

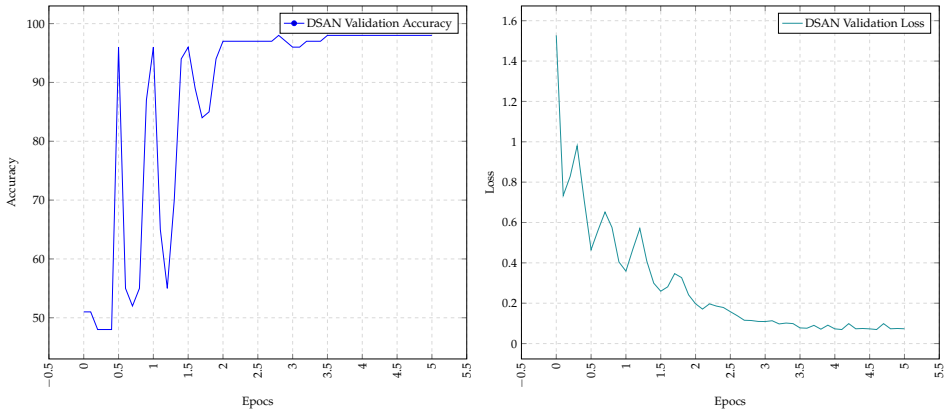

Figure 9. Validation accuracy and loss using DSAN domain adaptation.

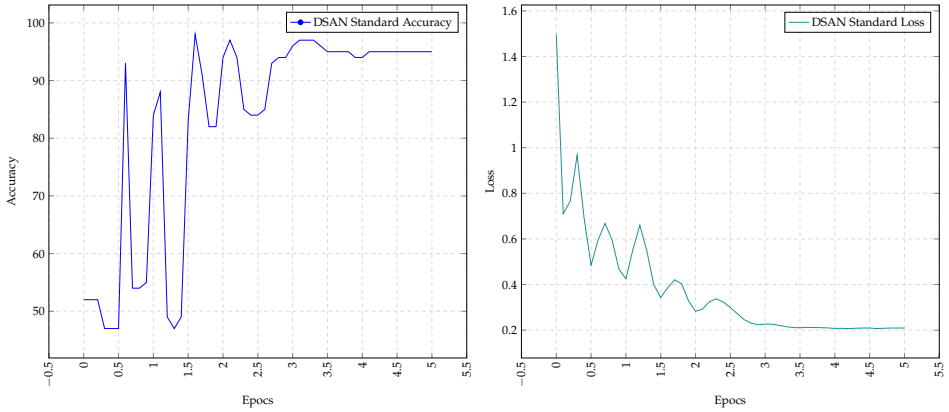

Figure 10. Standard accuracy and loss using DSAN domain adaptation.

The final test set accuracies for both marine and standard polymer test sets are shown in Table 3. This table represents the best-performing iteration using DSAN domain adaptation for both test sets. It can be seen in the table that a significant performance increase has been observed for the standard test set. Overall accuracy and F1-scores of 95% represent a satisfactory boost in performance. A decrease in marine test set accuracy has also been observed. This likely results from the domain alignment, causing the model to perform slightly worse on the source dataset but with the added benefit of increased performance on the target dataset.

Table 3. DSAN performance metrics for both the marine and standard polymer testing sets.

	Marine Testing Set				**Standard Testing Set**			
	Prec.	**Rec.**	**F1**	**Acc.**	**Prec.**	**Rec.**	**F1**	**Acc.**
PE	92%	97%	95%	95%	93%	97%	95%	95%
PP	97%	93%	95%	95%	98%	93%	95%	95%
Avg.	94.5%	95%	95%	95%	95.5%	95%	95%	95%

4.3. Deep CORAL Implementation

Due to the covariance calculations, the hyperparameters for Deep CORAL remain unchanged following the changes made in DSAN's implementation. The covariance calculations demand a large amount of memory due to the size of the features computed by the transformer. Figure 11 displays the validation performance of an example iteration using Deep CORAL. The training accuracy on the validation set has a similar trend to

that of DSAN's results in Figure 9, following a desirable convergence curve. This further shows that statistical domain-adaptation methods may be more suitable for time-series data. Figure 12 shows the performance for the standard dataset. The accuracy and loss over the five training epochs closely match the trends observed in Figure 11. These similar trends indicate that the model is learning to classify both datasets purely from the unsupervised domain-adaptation method.

Digging deeper, we can observe the best-performing Deep CORAL iteration in Table 4. Metrics for the marine dataset show that Deep CORAL was able to preserve its ability to classify the source dataset. It achieved an overall accuracy on the marine dataset of 99%. On the standard dataset, it achieved an accuracy of 94%. This result is 1% less than what was achieved using DSAN. However, the benefit here is that Deep CORAL was able to maintain its performance on accurate classifications for the marine dataset.

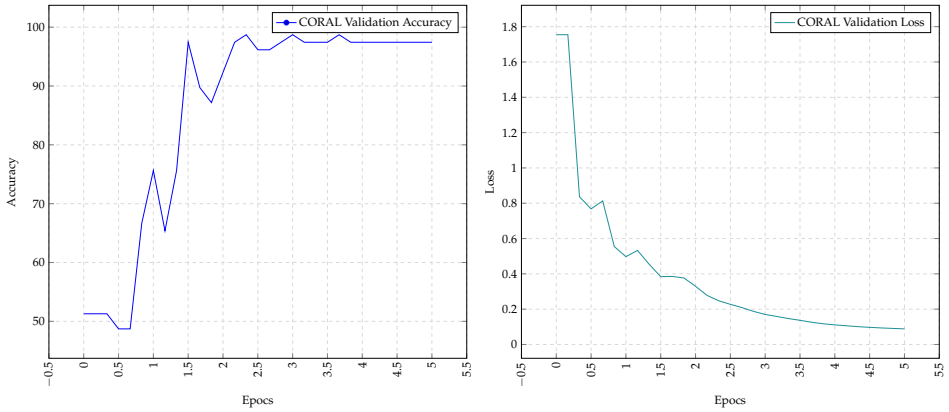

Figure 11. Validation accuracy and loss using CORAL domain adaptation.

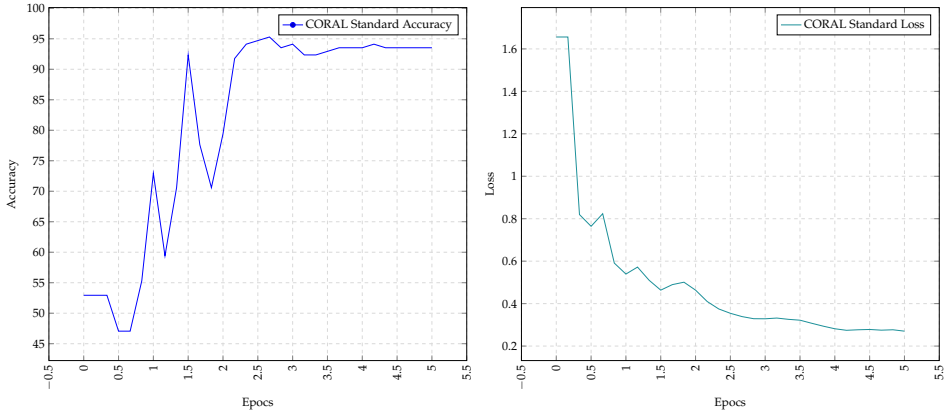

Figure 12. Standard accuracy and loss using Deep CORAL domain adaptation.

Table 4. Deep CORAL performance metrics for both the marine and standard polymer testing sets.

	Marine Testing Set				**Standard Testing Set**			
	Prec.	**Rec.**	**F1**	**Acc.**	**Prec.**	**Rec.**	**F1**	**Acc.**
PE	97%	100%	99%	99%	99%	88%	93%	94%
PP	99%	99%	99%	99%	95%	94%	94%	94%
Avg.	98%	99.5%	99%	99%	97.5%	93.5%	93.5%	94%

4.4. Comparison of Domain-Adaptation Methods

Table 5 offers a comparison of each domain-adaptation method. Of the three DANN was the worst-performing model due to its highly variable outputs. During successful

training iterations, it was unable to perform better than the vanilla transformer. Overall, the method's main proponent is its simplicity. To reiterate, the desired outcome is a model that can make consistent predictions on the target data, which the DANN method was unable to achieve. DSAN was the best performer on the standard dataset, with an accuracy of 95%. This may be due to the fact that it learns to align the local distributions, closing not only the domain shift but also the subdomain shift. The method was the worst performer in regard to the marine polymer set, with an accuracy of only 95%. While it may learn fine-grained information relevant to the subdomains, it could also be limited in its overall ability to learn hence why the marine polymer classifications are lower. Deep CORAL is the all-around best performer. While it scored 1% less in its accuracy on the standard polymer set, it was able to maintain its ability to accurately classify the marine polymer set. This could be due to its simplistic design, where only a few (yet large) calculations are made. Deep CORAL and DSAN are flawed in that they may not be the most appropriate architecture to be used with the transformer due to the large amount of memory required.

Table 5. Comparison of DANN, DSAN, and Deep CORAL domain-adaptation methods' performance on the marine and standard polymer testing sets.

	Marine Testing Set				**Standard Testing Set**			
Method	**Prec.**	**Rec.**	**F1**	**Acc.**	**Prec.**	**Rec.**	**F1**	**Acc.**
DANN	98.5%	98.5%	99%	99%	93%	93%	92%	92%
DSAN	94.5%	95%	95%	95%	95.5%	95%	95%	95%
CORAL	98.5%	99.5%	99%	99%	97.5%	93.5%	93.5%	94%

4.5. Integration with the Reflectance Micro-FTIR Workflow

To assess the model's practical application, the Deep CORAL model was integrated into the reflectance micro-FTIR analytical workflow. For testing, we used the original `.map` file from the standard polymer dataset, which contains complete spectral recordings for each pixel of the 241 × 244 IR image (58,804 spectra in total). This sample includes spectra for all standard polymer types as well as recordings from the filter paper on which they were placed.

Figure 13 shows the optical image of the standard polymer samples used in this experiment. Although the samples contain many polymer types, analysis was restricted to PE and PP, as these are the only polymers found in the marine dataset. Each spectra from the `.map` file was converted into CSV format, classified by the transformer, and used to visualise the polymer distribution in the image.

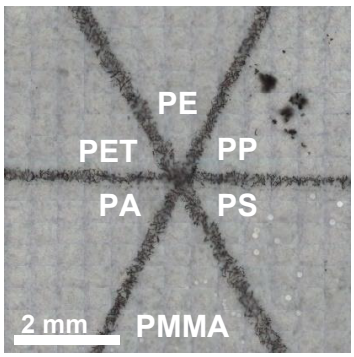

Figure 13. Optical image of the standard polymer sample, consisting of numerous microplastic polymers placed on filter paper—reproduced from [13].

The marine and standard polymer datasets were also updated to include the recorded spectra of the filter paper (FP) that the samples rest upon. This step was necessary for the model to identify the microplastic polymers from the background. Filter paper can change from sample to sample, so domain adaptation was employed on both the polymer and filter paper spectra during this step. To evaluate the validity of the predictions, the results will be compared to partial least-squares discriminant analysis (PLS-DA), which was trained and tested exclusively on this standard polymer dataset in the traditional supervised setting (no domain adaptation) considered in [13]. In contrast, the transformer was trained exclusively on the marine dataset and adapted to the standard polymer dataset using the Deep CORAL method, without access to the standard polymer training labels.

PLS-DA is a supervised classification technique that combines dimensionality reduction with discriminant analysis. In this approach, principal component analysis (PCA) is first applied to the spectral dataset to reduce dimensionality by transforming the large set of correlated wavenumber variables into a smaller set of orthogonal principal components (PCs) that capture the majority of the variance. This step facilitates visualisation, reduces noise, and addresses the issue of having more predictors than observations. The reduced feature set is then analysed using PLS-DA, which leverages known class labels to find latent variables that best separate predefined categories (e.g., different polymer types). The resulting model can classify unknown spectra by projecting them into the latent variable space and assigning them to the most likely class. Performance is typically evaluated using a hold-out test set or through cross-validation approaches such as leave-one-out, which iteratively assesses the model's predictive accuracy on unseen samples.

The left subplot of Figure 14 shows the pixel level classification of the polymer identities using PLS-DA. The results were confirmed to be highly accurate in [13]. As such, these results provide a strong basis for comparison with the cross-domain transformer model.

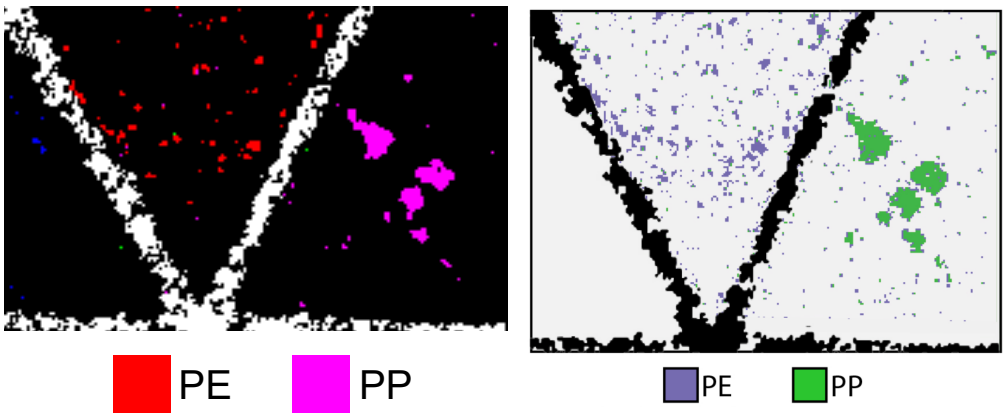

Figure 14. (**Left**): False colour map showing the prediction of the standard polymer identities using PLS-DA (Extracted from [13]), which approximates the upper bound of achievable classification performance on the target set. (**Right**): Pixel map showing the transformer's predictions of the standard polymer identities.

The right subplot of Figure 14 shows the pixel-level classifications of the polymer identities using the transformer model. For clarity, a graphite border has been overlaid to mark the PE and PP regions identified in Figure 13, while other polymer regions have been filled with the background colour. The predictions are largely successful, with both regions containing mostly correct polymer identities, aside from a few errors. It is important to reiterate that these predictions were made by a model with no access to the training labels of the standard polymer set, indicating that the underlying transformer model and the Deep CORAL domain-adaptation method have been largely successful in classification. However, when compared to the PLS-DA results in the left subplot, the transformer model

exhibits more classification errors, particularly in the form of false positives. The presence of PE classifications in the PP region, and vice versa, indicates that some misclassifications are present.

Figure 15 provides a closer view of the classification errors in both regions. In the PE region, most misclassifications occur along the edges of larger PE particles—a pattern also observed in the PP region, where PE misclassifications surround large PP microplastics. These errors are likely attributable to limitations of the reflectance micro-FTIR method. Because the IR imaging microscope has a fixed spatial resolution, smaller particles may not fully cover the detection area. Consequently, the recorded spectra can contain contributions from both the polymer and the underlying filter paper, producing mixed spectra that reduce classification accuracy. This effect likely explains the concentration of errors at polymer boundaries in both regions.

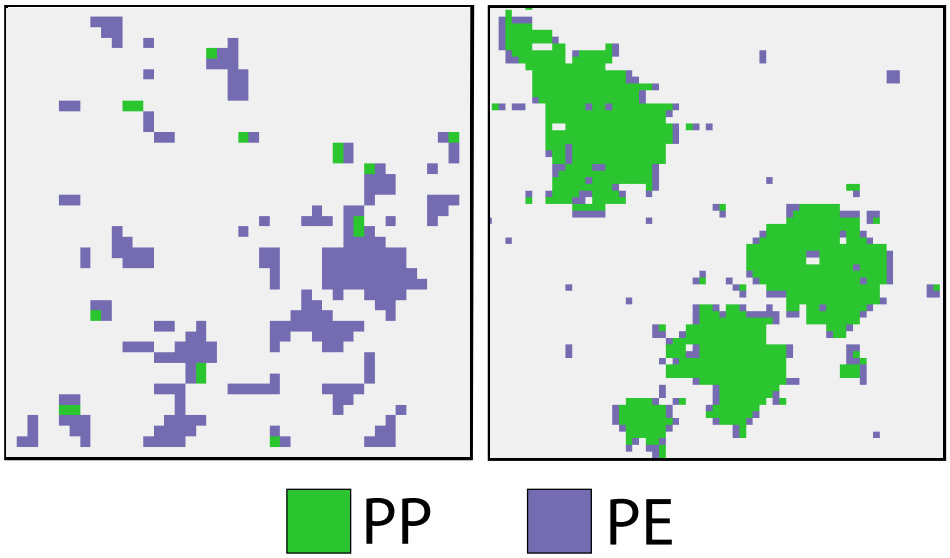

Figure 15. Exploded view of the polymer identity predictions.

For larger microplastics, polymer identity can often be inferred by the majority classification within a contiguous cluster. For example, in Figure 15, large green clusters are likely polyethylene simply because it is the dominant classification in that region. However, this assumption becomes less reliable for smaller clusters, where limited coverage and spectral mixing make the true identity less certain.

4.6. Discussion

The comparative evaluation of domain-adaptation methods aligns well with the existing literature on their respective strengths and limitations. Consistent with prior findings, the Domain-Adversarial Neural Network (DANN) showed limited robustness in handling complex domain shifts, which may be due to its sensitivity to noisy or heterogeneous data [21,29]. This outcome confirms the working hypothesis that simplistic adversarial alignment alone may be insufficient for the intricate spectral variations inherent in environmental microplastic datasets.

Deep Subdomain-Adaptation Network (DSAN), which explicitly aligns local subdomain distributions [30], performed strongly on the standard polymer set, supporting the premise that fine-grained domain alignment can improve classification when subdomain structures are present. However, its comparatively weaker performance on the marine dataset suggests challenges in scaling this approach to broader or more variable domain shifts, underscoring the need to tailor adaptation strategies to domain complexity.

Deep CORAL's consistently strong performance across datasets reinforces the value of covariance-based alignment methods [23], particularly in environmental applications characterized by diverse data sources and limited samples. Its relative simplicity and computational efficiency position it as a practical choice for real-world deployments involving resource-intensive architectures like transformers.

Broadly, these results highlight that effective domain adaptation in microplastic classification requires balancing methodological complexity with data variability and computational constraints. Given the variability of environmental microplastic samples and the scarcity of large, diverse datasets, methods like Deep CORAL that generalize well with modest resources are especially promising.

Looking forward, future research should explore hybrid adaptation frameworks that combine global and local alignment techniques to better capture domain heterogeneity. Additionally, developing domain-adaptation methods specifically optimized for transformer models could improve scalability. Augmenting training data through synthetic generation or targeted sampling may further enhance model robustness. Finally, integrating explainability and uncertainty quantification will be critical for fostering trust and practical application in environmental monitoring contexts.

5. Conclusions and Future Work

From the observations and results presented in this study, it is evident that integrating deep learning with reflectance micro-FTIR spectroscopy can lead to highly accurate and scalable methods for microplastics detection and classification. This work is the first to explore the application of deep neural networks—particularly a modified transformer-based model—for processing reflectance micro-FTIR spectra, establishing a new benchmark in automated microplastics analysis.

A key contribution of this study lies in identifying and addressing the domain-shift problem—a critical challenge that arises when polymer spectra differ significantly across datasets due to variations in environmental sources, instrument settings, or sample preparation. Such discrepancies can severely degrade model performance if not properly accounted for. To mitigate this, the study explored several domain-adaptation techniques, including both shallow and deep domain-adaptation strategies, to align the distributions of source and target domains.

The results demonstrate that these domain-adaptation methods substantially enhance model robustness. The proposed transformer model, when integrated with domain-adaptation strategies, achieved remarkable classification accuracies of 99% on the source domain and 94.4% on the target domain, clearly outperforming baseline approaches. This validates the potential of domain adaptation to bridge the domain gap and improve generalisability in real-world scenarios where labelled target data is limited or unavailable.

Moreover, the study highlights the feasibility of conducting end-to-end analysis on unknown samples without requiring prior polymer identity information—thanks to the strong generalization capabilities enabled by domain adaptation. Collectively, these findings establish a foundation for domain-adaptive deep learning frameworks in microplastic research and further underscore reflectance micro-FTIR spectroscopy as a powerful and practical tool for high-throughput microplastic identification across diverse environmental contexts.

A major shortcoming of this study is the inherent lack of data. Due to its limitations, reflectance micro-FTIR has not been widely adopted by the scientific community. This has resulted in a significant scarcity of available data, limiting the validation of the deep learning model. Access to two datasets has served this study well; however, additional datasets could better highlight the extent of domain shifts across different microplastic distributions. It may be that environmental microplastics are very similar regardless of location, which

would imply that domain shift is not a major concern. Conversely, the opposite could be true, suggesting a multitude of distinct domains to consider. If reflectance micro-FTIR becomes more widely used, large public datasets containing microplastic polymer spectra would enable future deep learning models to classify microplastics accurately across domains, thereby supporting serious practical applications.

The transformer neural network employed in this study was overall successful, but several considerations remain for future research. Firstly, transformers are computationally expensive deep learning models, requiring a capable GPU to train effectively. Other models achieving similar performance—such as 1D CNNs—have an advantage over transformers because they are less memory-intensive, making them more accessible to researchers in chemistry and related fields. Incorporating domain-adaptation methods significantly increases the complexity of the transformer model; domain adaptation might therefore be more suitable for other architectures like CNNs. Additionally, transformers generally require large amounts of training data to perform optimally. The full potential of this model may not have been realised here, as additional data could further improve its generalisation ability. Future research could explore replacing the transformer with other supervised classifiers drawn from different domains, such as medical spectral imaging or time-series anomaly detection, to further evaluate the generalisability of domain-adaptation methods for microplastic classification.

Finally, the model operates as a black box, meaning the rationale behind its predictions is not transparent. Future research could benefit from integrating explainability features—such as providing users with reasons for specific classifications and confidence scores. This information could assist users in assessing the reliability of the model's predictions.

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

Funding: This research received no external funding.

Data Availability Statement: For research data used in this work and implementation of the baseline transformer classification model, please visit https://github.com/br3nr/microplastic-transformer (accessed on 13 August 2025). For the domain-adaptation methods implemented in this paper, please visit https://github.com/br3nr/microplastic-da (accessed on 13 August 2025).

Acknowledgments: The authors would like to acknowledge Aneesh Krishna for their administrative support throughout the project.

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

References

1. Lillo, P.M.; Udagedara, S.; Williamson, R.; Gorman, D. Cross-Shore Microplastic Accumulation on Sri Lanka's West Coast One Year After the Catastrophic X-Press Pearl Pollution Event. *Microplastics* **2025**, *4*, 37. [CrossRef]
2. Piskuła, P.; Astel, A.M. Microplastics occurrence in two mountainous rivers in the Lowland area—A case study of the Central Pomeranian Region, Poland. *Microplastics* **2022**, *1*, 167–185. [CrossRef]
3. Vethaak, A.D.; Legler, J. Microplastics and human health. *Science* **2021**, *371*, 672–674. [CrossRef] [PubMed]
4. Wu, Z.; Wang, J.; Yu, S.; Sun, Q.; Han, Y. Microplastic Pollution in China's Aquatic Systems: Spatial Distribution, Transport Pathways, and Controlling Strategies. *Microplastics* **2025**, *4*, 41. [CrossRef]
5. Barboza, L.G.A.; Frias, J.P.; Booth, A.M.; Vieira, L.R.; Masura, J.; Baker, J.; Foster, G.; Guilhermino, L. Microplastics pollution in the marine environment. In *World Seas: An Environmental Evaluation Volume III: Ecological Issues and Environmental Impacts*; Academic Press: Amsterdam, The Netherlands, 2018. [CrossRef]

6. Jenner, L.C.; Rotchell, J.M.; Bennett, R.T.; Cowen, M.; Tentzeris, V.; Sadofsky, L.R. Detection of microplastics in human lung tissue using μFTIR spectroscopy. *Sci. Total Environ.* **2022**, *831*, 154907. [CrossRef] [PubMed]
7. Leslie, H.A.; van Velzen, M.J.; Brandsma, S.H.; Vethaak, A.D.; Garcia-Vallejo, J.J.; Lamoree, M.H. Discovery and quantification of plastic particle pollution in human blood. *Environ. Int.* **2022**, *163*, 107199. [CrossRef]
8. Greeves, S.; McGovern, R.; McKinney, M. Plastic Pollution of the Tennessee River: Comparing Risk Perceptions and Preferred Policy Solutions Between Stakeholders and the Public. *Microplastics* **2025**, *4*, 40. [CrossRef]
9. Igalavithana, A.D.; Mahagamage, M.G.Y.; Gajanayake, P.; Abeynayaka, A.; Gamaralalage, P.J.D.; Ohgaki, M.; Takenaka, M.; Fukai, T.; Itsubo, N. Microplastics and potentially toxic elements: Potential human exposure pathways through agricultural lands and policy based countermeasures. *Microplastics* **2022**, *1*, 102–120. [CrossRef]
10. Li, D.; Li, P.; Shi, Y.; Sheerin, E.D.; Zhang, Z.; Yang, L.; Xiao, L.; Hill, C.; Gordon, C.; Ruether, M.; et al. Stress-induced phase separation in plastics drives the release of amorphous polymer micropollutants into water. *Nat. Commun.* **2025**, *16*, 3814. [CrossRef]
11. Adelugba, A.; Emenike, C. Comparative review of instrumental techniques and methods for the analysis of microplastics in agricultural matrices. *Microplastics* **2023**, *3*, 1–21. [CrossRef]
12. Li, D.; Sheerin, E.D.; Shi, Y.; Xiao, L.; Yang, L.; Boland, J.J.; Wang, J.J. Alcohol pretreatment to eliminate the interference of micro additive particles in the identification of microplastics using Raman spectroscopy. *Environ. Sci. Technol.* **2022**, *56*, 12158–12168. [CrossRef]
13. Willans, M. Developing a Protocol for Microplastic Detection Using Micro-Reflectance FTIR Spectroscopy. Bachelor's Honours Thesis, Curtin University, Bentley, WA, Australia, 2021.
14. Hidalgo-Ruz, V.; Gutow, L.; Thompson, R.C.; Thiel, M. Microplastics in the marine environment: A review of the methods used for identification and quantification. *Environ. Sci. Technol.* **2012**, *46*, 3060–3075. [CrossRef]
15. Harrison, J.P.; Ojeda, J.J.; Romero-González, M.E. The applicability of reflectance micro-Fourier-transform infrared spectroscopy for the detection of synthetic microplastics in marine sediments. *Sci. Total Environ.* **2012**, *416*, 455–463. [CrossRef]
16. Lee, G.; Jung, J.; Moon, S.; Jung, J.; Jhang, K. Microscopic Image Dataset with Segmentation and Detection Labels for Microplastic Analysis in Sewage: Enhancing Research and Environmental Monitoring. *Microplastics* **2024**, *3*, 264–275. [CrossRef]
17. Shi, B.; Patel, M.; Yu, D.; Yan, J.; Li, Z.; Petriw, D.; Pruyn, T.; Smyth, K.; Passeport, E.; Miller, R.J.; et al. Automatic quantification and classification of microplastics in scanning electron micrographs via deep learning. *Sci. Total Environ.* **2022**, *825*, 153903. [CrossRef]
18. Zhu, Z.; Parker, W.; Wong, A. PlasticNet: Deep Learning for Automatic Microplastic Recognition via FT-IR Spectroscopy. *J. Comput. Vis. Imaging Syst.* **2021**, *6*, 1–3. [CrossRef]
19. Barker, M.; Willans, M.; Pham, D.S.; Krishna, A.; Hackett, M. Explainable detection of microplastics using transformer neural networks. In Proceedings of the Australasian Joint Conference on Artificial Intelligence, Perth, WA, Australia, 5–8 December 2022; Springer: Cham, Switzerland, 2022; pp. 102–115. [CrossRef]
20. Farahani, A.; Voghoei, S.; Rasheed, K.; Arabnia, H.R. A Brief Review of Domain Adaptation. In *Advances in Data Science and Information Engineering*; Springer International Publishing: Cham, Switzerland, 2021; pp. 877–894. [CrossRef]
21. Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; Lempitsky, V. Domain-adversarial training of neural networks. *J. Mach. Learn. Res.* **2016**, *17*, 1–35.
22. Zhu, Y.; Zhuang, F.; Wang, J.; Ke, G.; Chen, J.; Bian, J.; Xiong, H.; He, Q. Deep subdomain adaptation network for image classification. *IEEE Trans. Neural Netw. Learn. Syst.* **2021**, *32*, 1713–1722. [CrossRef]
23. Sun, B.; Saenko, K. Deep CORAL: Correlation Alignment for Deep Domain Adaptation. In *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*; Springer: Berlin/Heidelberg, Germany, 2016; Volume 9915 LNCS, pp. 443–450. [CrossRef]
24. Pan, S.J.; Yang, Q. A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.* **2010**, *22*, 1345–1359. [CrossRef]
25. Ragab, M.; Eldele, E.; Tan, W.L.; Foo, C.S.; Chen, Z.; Wu, M.; Kwoh, C.K.; Li, X. ADATIME: A Benchmarking Suite for Domain Adaptation on Time Series Data. *arXiv* **2022**, arXiv:2203.08321. [CrossRef]
26. Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; Ng, A.Y. Reading Digits in Natural Images with Unsupervised Feature Learning. In Proceedings of the NIPS Workshop on Deep Learning and Unsupervised Feature Learning, 2011. Available online: http://ufldl.stanford.edu/housenumbers/nips2011_housenumbers.pdf (accessed on 13 August 2025).
27. Deng, L. The MNIST Database of Handwritten Digit Images for Machine Learning Research [Best of the Web]. *IEEE Signal Process. Mag.* **2012**, *29*, 141–142. [CrossRef]
28. Sun, B.; Feng, J.; Saenko, K. Return of Frustratingly Easy Domain Adaptation. In Proceedings of the AAAI Conference on Artificial Intelligence, Phoenix, AZ, USA, 12–17 February 2016; AAAI Press: Palo Alto, CA, USA, 2016; Volume 30, pp. 2058–2065. [CrossRef]

29. Tzeng, E.; Hoffman, J.; Saenko, K.; Darrell, T. Adversarial Discriminative Domain Adaptation. In Proceedings of the 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, 21–26 July 2017; pp. 2962–2971. [CrossRef]
30. Pei, Z.; Cao, Z.; Long, M.; Wang, J. Multi-adversarial domain adaptation. In Proceedings of the AAAI Conference on Artificial Intelligence, New Orleans, LA, USA, 2–7 February 2018; Volume 32. Available online: https://dl.acm.org/doi/10.5555/3504035.3504517 (accessed on 13 August 2025).