

# On the Effectiveness of Heterogeneous Ensemble Methods for Re-identification

Simon Klüttermann<sup>1,2</sup>, Jérôme Rutinowski<sup>1,2</sup>, Frederik Polachowski<sup>1</sup>, Anh Nguyen<sup>1</sup>  
Britta Grimme<sup>4</sup>, Moritz Roidl<sup>1,2</sup>, Emmanuel Müller<sup>1,2,3</sup>

<sup>1</sup>*TU Dortmund University, Dortmund, Germany*

<sup>2</sup>*Lamarr Institute for Machine Learning and Artificial Intelligence, Dortmund, Germany*

<sup>3</sup>*Research Center Trustworthy Data Science and Security, Dortmund, Germany*

<sup>4</sup>*Paderborn University, Paderborn, Germany*

simon.kluettermann@cs.tu-dortmund.de

**Abstract**—In this contribution, we introduce a novel ensemble method for the re-identification of industrial entities, using images of chipwood pallets and galvanized metal plates as dataset examples. Our algorithms replace commonly used, complex siamese neural networks with an ensemble of simplified, rudimentary models, providing wider applicability, especially in hardware-restricted scenarios. Each ensemble sub-model uses different types of extracted features of the given data as its input, allowing for the creation of effective ensembles in a fraction of the training duration needed for more complex state-of-the-art models. We reach state-of-the-art performance at our task, with a Rank-1 accuracy of over 77% and a Rank-10 accuracy of over 99%, and introduce five distinct feature extraction approaches, and study their combination using different ensemble methods.

## I. INTRODUCTION

Even though the use of deep learning models has become more prevalent for numerous applications, such as industrial computer vision tasks, more shallow approaches remain advantageous in their own right. This is because the training process of deep learning models requires larger amounts of time and data than their shallower counterparts, and allow less control over the functions learned. Shallow models, however, tend to yield a lower performance in their respective tasks, mainly because of their lower complexity. To reduce the overhead of training complex, monolithic deep learning models and to circumvent the poorer performance of shallow models, ensemble learning can be used. Ensemble models can combine the strengths of multiple simpler models by retaining their training simplicity while improving results [1]. The combination of such models can be realized in many different ways [2], [3], and the ensembles can also partially or entirely consist of deep learning models [4]. But while ensembles are a common and straightforward method in supervised machine learning, adapting them to re-identification tasks provides its own challenges.

In industrial settings, providing effective and efficient results is vital: Errors can imply the endangerment of human life or considerable risks and costs. However, at the same time, the methods used need to be able to work in real-time and often even in hardware-restricted settings, to be economically viable. Thus, ensembles can provide a very practical solution for

many applications, in contrast with demanding deep learning solutions [4].

One such industrial application is the re-identification of logistical entities, i.e., inanimate objects used in the context of warehousing. A prominent logistical entity is the Euro-pallet, which is a widely used type of load carrier, with hundreds of millions of them in constant circulation [5]. Despite their industrial relevance, Euro-pallets, like other types of standardized pallets, are not equipped with inherent identifiers (e.g., an ID or a barcode) and are instead only identifiable as part of a cluster by, e.g., their place of assembly [6]. This, in turn, does not lead to the individual identification of the pallet but rather its classification (i.e., knowing where a pallet was produced but not making it distinguishable from other pallets produced at the same site). Due to these limitations, Euro-pallets are usually only identified through documentation such as waybills.

Identifying individual pallets would provide the industry with further knowledge about operational processes, permitting their analysis and improvement [7], [8], beyond general detection and classification. This provides users with information on individual pallets and the goods they are carrying in a dispatch network of pallets, permitting process optimization that would otherwise be impossible to realize. While there have been some suggestions to use artificial markers (e.g., QR-codes), these provide their own challenges through the high number of pallet blocks in circulation and the effects of wear and tear. Thus, we are instead trying to recognize the pallet blocks themselves.

A first attempt to solve this challenge, based on the inherent visual characteristics of Euro-pallets, was presented in [9]. In this contribution, the authors proved the feasibility of identifying the chipwood pallet blocks used in Euro-pallets based on their unique surface structure. This work was further expanded upon in [10], in which a larger batch of images, recorded in the industry, was successfully used for re-identification. However, in both contributions, sophisticated deep learning models were employed. In this contribution, we will instead try to tackle the same challenge by employing ensemble models, in order to provide efficient predictions that still reach a comparable accuracy.

The code is publicly available at: <https://github.com/KDD-OpenSource/PalletEns>

## II. RELATED WORK

We define re-identification as the attempt to retrieve a previously recorded subject of interest over a network of cameras [11]. Thus, the predictions of a re-identification model are meant to match input data from a query dataset  $\mathcal{Q}$  with a given gallery dataset  $\mathcal{G}$ . For this purpose, the recorded subjects are assigned a descriptor during their first recording (generated, e.g., by feature extraction). Based on this descriptor, the similarity between subjects is compared to one another, and the subjects perceived as being the most similar are predicted to be the same [12], [13]. For this, a list of subjects, ranked by their similarity to the query data ( $\mathcal{Q}$ ), is generated from the gallery dataset ( $\mathcal{G}$ ). Prior to the matching task, training occurs using the training dataset ( $\mathcal{T}$ ). (These datasets are distinct ( $\mathcal{Q} \cap \mathcal{G} = \mathcal{Q} \cap \mathcal{T} = \mathcal{G} \cap \mathcal{T} = \emptyset$ ), in order to prevent information leakage and to limit overfitting.) After training, the task of re-identification can then formally be described as matching an image  $x_i \in \mathcal{Q}$  of a subject  $i$  to an image  $y_j \in \mathcal{G}$  of a subject  $j$  where  $i = j$  [14].

### A. Re-identification Use Cases

The most common use for re-identification is pedestrian surveillance [11], [15], [16], notably in the United Kingdom and the People’s Republic of China. Another application is vehicle surveillance [17], [18], which is analogous to pedestrian surveillance in its motivation. Even animals have been shown to be re-identifiable using their inherent visual features [19], [20]. A more industry-related application is the re-identification of materials [8]–[10], [21]–[23].

While different kinds of visual identification criteria can be used for the re-identification of pedestrians, such as their face [24] or their gait [25] and pedestrians can easily be segmented (e.g., into head, torso, legs), this does not apply to inanimate subjects, such as industrial entities. Still, first approaches to re-identification of industrial entities do exist [9], [10], [23], [26], offering the possibility to identify Euro-pallets by exploiting their unique surface patterns. In these contributions, PCB\_P4 (Part-based Convolutional Baseline) [27] and graph-based approaches have been deployed. The results of these studies demonstrate that the surface structure of chipwood and solid wood can be used for the re-identification of Euro-pallets. The datasets used for these works are publicly available [28], [29]. These works, however, do not employ ensemble methods and are only applied to a single logistical entity, namely Euro-pallets.

### B. Re-identification Approaches

An example of a re-identification approach is metric learning, which maps the subject’s similarity as feature embeddings, describing similarity as distances between two points in this embedding space [16], [30]. As loss functions, classification loss, contrastive loss and, most prominently, triplet loss are

often used to minimize intra-class distances and maximize inter-class distances [15], [16], [31].

As to evaluate re-identification results, the most common metric used is ranked accuracy, also known as top-k accuracy. This metric describes whether the query image is retrieved in the first  $k$  matches of the gallery data [11], [16]. This adds explainability and transparency to the model’s predictions, providing insight into the severity of a wrong retrieval. This insight can be especially useful, as there is a practical difference for many industrial use cases, between a prediction that is off by one rank or multiple ranks.

### C. Ensembles Usage for Re-identification

Using ensemble models is a common practice in machine learning, allowing the combination of different predictors to achieve a better performance [32]. There are many methods, ranging from the simple averaging of multiple models (bagging) [2] to training new models on the prediction of the ensemble sub-models (stacking) [33]. Additionally, ensembles can have further uses, like explainability [34]. However, there are surprisingly few applications to re-identification, as this requires us to adapt concepts like bagging and stacking. We want to summarize the existing approaches in this chapter.

A common approach to re-identification focuses on extracting general features representing aspects of a sample and using their resemblance as an indicator of similarity. This invites the use of ensembles to combine different features. A previous publication [35] employs ensembles by combining different hand-crafted color descriptors, which allows the extraction of information from more than one feature. Recent approaches [36]–[38] apply more advanced neural networks to find features using pre-trained models or to extract them directly.

The drawback of feature extraction is that many datasets contain a large amount of superfluous features, which tend to be included in well-studied, commonly extracted features. Additionally, there has not been enough research conducted yet to find specialized features for pallet re-identification and a solution emerging from such research would only be applicable to one use case or dataset, thus not leading to a general solution. Siamese neural networks [39], however, are able to effectively extract useful features automatically, providing a more generalizable and adaptable solution. Other approaches use ensembles for metric learning [40] or to better handle multiple data modalities [41].

However, to the best of our knowledge, ensembles have hardly been studied for siamese neural networks. Worth mentioning is [42], in which ensembles are used to enhance contrastive learning in the context of natural language processing, and [43], which studies the benefits of homogenous ensembles. While this ensemble method shares similarities with our approach, the application is fundamentally different, which limits the transferability of results.

## III. METHODOLOGY

This section aims to provide the reader with insights into our methodology. The section delves deeper into re-identification




Fig. 1. Excerpts of the used datasets. a) pallet block dataset [28], b) galvanized metal dataset [44]

task description for the given use case. It also describes the re-identification methods that are used for this task, as well as the ensembles that are comprised of the former.

As a first step, we describe the herein used datasets (see Fig. 1). The first dataset [28], containing images of wooden pallet blocks, is an open-source dataset that is freely available online, providing insights into the dataset creation process and recording setup. It comprises 5,020 RGB images of pristine, unbranded EPAL pallet blocks. The images in the dataset have been recorded using two different cameras and five different perspectives: central, left-hand and right-hand side rotation, and left-hand and right-hand side shift. Half of those images (2,510), taken by camera 1 in different rotations, are used for the experiments in this contribution. Other parts of the dataset, as well as its entirety, could also be used for the experiments. With 10 images recorded per pallet block, the dataset contains images of a total of 502 different pallet blocks.

Additionally, we use a second dataset [44], which contains 5,088 images of galvanized metal plates. Galvanized metal, like wood, is a commonly used material (e.g., for pallet cages or metal crates) and also possesses unique surface structure that makes it identifiable. The plates used for this dataset are recorded using ambient light and photo light, and both at a 90° and a 75° angle. For our experiments, we used the 2,544 images of the dataset recorded with photo lighting under 75° and 90° angles.

### A. Task Description

The task of re-identification is often solved by learning a transformation  $f$  of a dataset such that the distances between the representations of the same entity are minimized and the distances between different ones are maximized. This goal can

be formalized as follows (with  $x^i$  being a sample following concept  $i$ ):

$$\|f(x_a^i) - f(x_b^j)\|_2 = \begin{cases} \text{small} & \text{if } i = j \\ \text{large} & \text{if } i \neq j \end{cases} \quad \forall x, y \quad (1)$$

This task can be achieved by minimizing a suitable loss function, e.g., triplet loss [45]:

$$L_{\text{triplet}} = \max(0, \|f(x_a^i) - f(x_b^i)\|_2 - \|f(x_a^i) - f(x_c^j)\|_2 + \alpha), \quad i \neq j \quad (2)$$

where  $x_a^i$ ,  $x_b^i$  and  $x_c^j$  represent three samples of two different subjects  $i$  and  $j$  and  $\alpha$  is a hyperparameter. This loss value is averaged over a large number of different triplets.

Our goal in this contribution is to study the case in which we have not only a single transformation  $f$  but also a second transformation  $g$ . Therefore, we define and investigate an ensemble  $e$  including both transformations, with an ensemble creation transformation  $T$  creating an ensemble.

$$e(x) = T(f(x) \oplus g(x)) \quad (3)$$

### B. Re-identification Methods

In order to study ensembles, we require sub-models. We chose heterogeneous sub-models, as this allows using pre-existing knowledge about our dataset. Additionally, this is a faster approach, as a lower number of sub-models is required. To the best of our knowledge, this is the first application of heterogeneous ensembling for re-identification. However, this also means that we might need to change the sub-models depending on the dataset, in order to maximize the ensemble's prediction accuracy. Here, we suggest five siamese models for each dataset, all based on different approaches:

**Image-based re-identification:** The most direct approach for training a siamese neural network on image-based data is to use convolutional neural networks. The benefit of this approach is its ability to extract information available in the entire image. However, this also represents the most time-consuming approach, and is the most susceptible to overfitting. Further information on the training process of this model can be found in the supplementary material.

**Graph-based re-identification:** For our graph-based approach, we manually extract the most essential information before training a siamese neural network on it. For this, we use the approach implemented in [23], where anomalous parts of an image are considered nodes in a graph. Notice that our performance differs from the one stated in [23], as we use a different fraction of query and gallery data. Alas, this approach is specialized to the pallet block dataset, and our attempts to generalize it to the metal dataset were unsuccessful. Thus, we only use it on the first dataset.

**Linear quantile re-identification:** To compensate for the lack of applicability of the graph based approach to the metal dataset, we developed a different model, exploiting the perspective-wise shearing in the galvanized metal dataset. For this purpose, we split the images into pixel columns and use

the 0.2, 0.5 and 0.8 quantile of each pixel value in this column as an input for a siamese neural network.

**Brightness, average color and color variance based re-identification:** The remaining three approaches split a given image into sub-images of size  $16 \times 16$  px each. We iterate over all sub-images of an image (with 50% overlap) and calculate 1 – 3 values for each, resulting in a representation with 768/2304 dimensions. For the *brightness*-based approach, we calculate the mean of each sub-image, while for the *average color*-based approach, we calculate a mean for each color and sub-image. Finally, for the *color variance*-based approach, we use the standard deviation of each color throughout the sub-image. The values for each sub-image are concatenated into a vector and used to train a siamese neural network. Each network uses three dense layers with 100 nodes each to generate a 50-dimensional representation. While these approaches ignore much of the available information given in an image, they are also straightforward and time-efficient, reducing the required training time by multiple orders of magnitude.

### C. Ensembles

In this contribution, we study four different transformations  $T$  that create ensembles. While ensembles are used in many cases (see Section II), they are only rarely used for siamese network-based re-identification. We therefore aim to use ensemble models and study their applicability for the task at hand, since an ensemble usage would allow us to benefit from the combination of rudimentary methods instead of relying on a single, complex model, providing time-efficient and computationally less demanding results. We study here multiple methods, since these can have a large impact on the quality of the resulting ensemble [46].

The simplest ensemble transformation presented in this work, the *Concatenation* transformation, normalizes each ensemble component using the mean and the standard deviation (*std*) for all considered samples and concatenates them. It can be seen as an ensemble method following the concept of bagging [2].

$$T_{Concat}(x_0^i, \dots, x_n^i) = (z_0^i, \dots, z_n^i); z_j^i = \frac{x_j^i - \bar{x}_j}{s_{x_j}} \quad (4)$$

Here  $x_j^i$  is denoting the  $j$ -th component of the representation of sample  $x^i$ . Meanwhile, the *Neural Network Triplet* transformation is a stacking method [47], which uses a neural network to represent  $T_{NNtriplet}$  and to train it by minimizing Eq. (2) on the training set  $\mathcal{T}$ . The *Weighted Triplet* instead simply uses weighing factors:

$$T_{Wtriplet}(f(x) \oplus g(x)) = \alpha_f f(x) \oplus \alpha_g g(x) \quad (5)$$

These are found by minimizing Eq. (2), using gradient descent. Subsequently, the *Weighted Accuracy* transformation tries to maximize the probability that the closest sample  $y^j \in \mathcal{G}$  to a sample of the test set  $x^i \in \mathcal{Q}$  depicts the same subject  $i = j$ .

Since this probability is not continuous, gradient descent cannot be used for this task. Instead, the probability is optimized using the *flaml* library [48], which is designed to optimize non-continuous hyperparameter optimization problems, and thus does not rely on gradient descent. Additionally, a fifth ensemble, called the *Majority Vote* is created using a different approach to ensembling. For this ensemble, multiple sub-models do not calculate the distance to each gallery sample  $\in \mathcal{G}$ , but simply order the gallery samples based on their distance to the query sample. We choose the most common index of this order through all sub-models as an indicator of similarity between samples.

## IV. EXPERIMENTS

To test our novel approaches and to evaluate which ensemble method yields the highest performance, we train a siamese neural network on each of the data pre-processing steps described in Section III-B. This creates five different representations of each image, in which samples representing the same object are closer to each other than to different ones. While some pre-processing steps retain more information, others are able to express the information they contain in a more efficient manner.

To provide more meaningful results and to estimate prediction uncertainties, we employ cross-validation. We split our dataset into five different groups on the pallet block dataset and six different groups on the metal dataset. One of these folds is chosen as a *query* and *gallery* set, in which for each pallet block, one image is selected for the *query* set, while the rest become part of the *gallery* set. One fold is not used but held out for a possible future novelty detection task, while the remaining folds are used as a *training* set.

### A. Individual Results

First, we evaluate the individual sub-models and plot the resulting Rank-1 to Rank-10 accuracy, as shown in Fig. 2.

The two methods averaging color values (*brightness* and *average color*) provide subpar results on the pallet block dataset, with their overall performance reaching a Rank-1 accuracy of  $< 20\%$ . An entirely different effect can be observed for the metal dataset, with the simpler *brightness*-based method outperforming a method with image-wide information, demonstrating the effect that a careful selection of sub-models can provide.

Still, in both cases, using the *color variance*-based model provides a higher accuracy. We assume that this is due to high variances often representing unusual structures that can be used for re-identification, while in an average-based method these structures are averaged out. Even though the performance of these simple methods are not always competitive, they are an order of magnitude faster to train than other methods (i.e.,  $\sim 8s$  compared to about  $20min$  for the Inception model) and thus can be valuable ensemble sub-models.

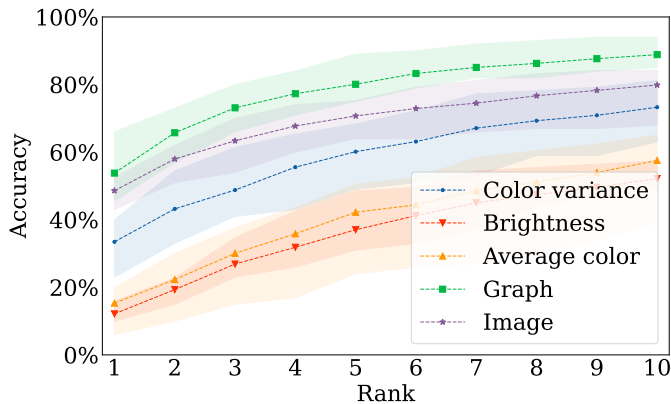
Notice that on both datasets, the best performance was not achieved using the highest number of available features. This implies the notion that performance can further be improved using an ensemble method.

## B. Ensemble Methods


In this section, the performance of the ensemble models is studied. Fig. 3 show the performance of the five different ensemble models discussed in Section III-C.

The ensemble models provide much higher performances than the individual models. Given the results in Fig. 3, the ensembling approach yielding the highest accuracy on the pallet block dataset seems to be the concatenation, which allows us to improve the Rank-1 accuracy from 54% to over 70%, compared to the graph-based sub-model. On the metal dataset, the concatenation also performs well, but here, learning weights seems valuable, increasing the Rank-1 accuracy from 55% to almost 78%, compared to the best sub-model. It is quite counter-intuitive that the simplest (and unsupervised) approach to ensembling also yields one of the highest performances.

We studied three ways of weighing individual sub-models and expected the weighing approaches to improve results, given that Fig. 2 showed significant differences of the sub-models' performances. We would also have expected similarly functioning approaches (like our *brightness* and *average color*-based approaches) to potentially cancel one another

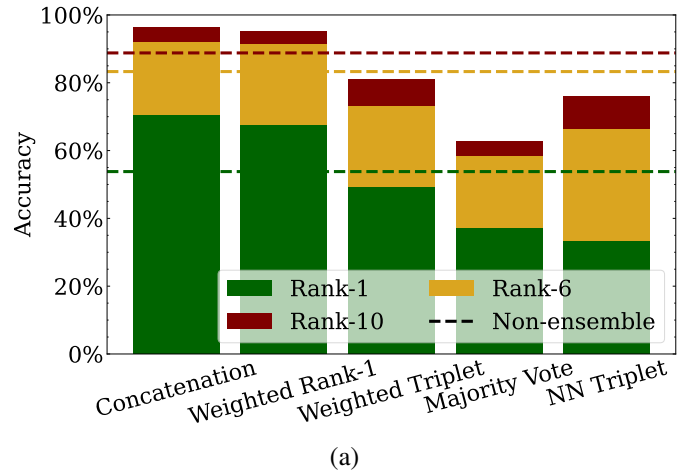


(a)

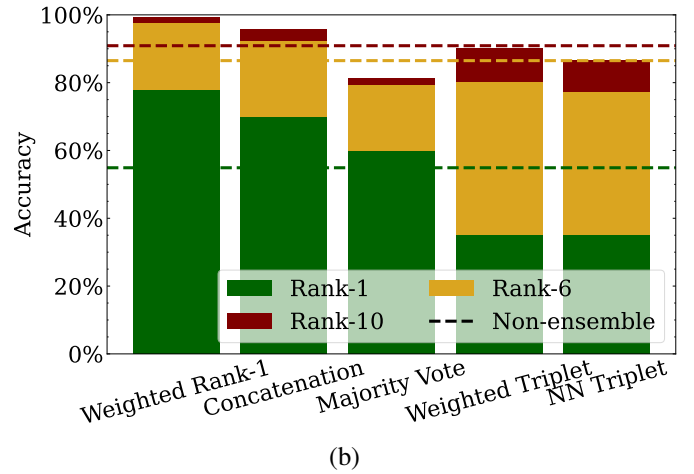


(b)

Fig. 2. Rank-k accuracy for the ensemble sub-models. The highlighted areas represent the highest and lowest scores per individual fold. a) pallet block dataset, b) metal dataset



(a)



(b)

Fig. 3. Rank-1, Rank-6, and Rank-10 accuracy for the five different ensemble approaches compared to our most effective non-ensemble model (the *graph*-based model). a) pallet block dataset, b) metal dataset

out. However, while weighing seems to increase the resulting performance in some cases, it seems not always to be an effective solution. One reason for this phenomenon might be the weighing approaches' susceptibility to overfitting.

The lowest (Rank-1) accuracy is obtained by the neural network approach, for which we use another neural network trained on the output of each siamese neural network to minimize the triplet loss (Eq. (2)) and generate a new common representation. This matches our expectation, since the neural network also contains the highest amount of parameters and thus has the highest potential of overfitting. Slightly better results are provided when a weighting factor is added and optimized to the output of each sub-model (*Weighted Triplet*). Still, the performance is inferior to that of the best-performing sub-model. We expect this to be a result of a local minimum of the optimization task because when changing the function to be optimized from a triplet loss to a Rank-1 accuracy on the training set (*Weighted Accuracy*) the performance can be improved, up to a level comparable to the performance of the *Concatenation* approach. In any case, this implies that there remains untapped potential for the application of more

sophisticated ensemble methods.

We tried another approach, in which the median prediction of each sub-model is chosen (*Majority Vote*). Still, this method does not perform well either. Through this ensemble method usage, we effectively perform an abstract segmentation of the data, analogous to the aforementioned pedestrian segmentation approaches, which however is restricted to this very use case, unlike our approach. Therefore, while other high-performance methods, especially in pedestrian re-identification are highly specialized for their specific domain of application, our approach yields a high degree of transferability, as we also demonstrate by applying it to two distinct datasets.

### C. State of the Art Comparison

Finally, we compare our method to commonly used approaches (see Tab. I). In addition to the previously introduced models, we compare our models to a method similar to one used for Person Re-identification [15], [27]. The approach here is to use a pre-trained model, to add three Dense layers for the output to match, and to continue training the model. The main difference to common person re-identification models, is that we employ the InceptionV3 object detection model [49], instead of a model trained to re-identify pedestrians. We chose this model over, e.g., other state-of-the-art models like ResNet50 [50] and MobileNetV2 [51] as it seems to provide significantly higher performance in our experiments. Still, this approach is not able to outperform our ensemble methods while requiring significantly more time to train.

TABLE I

COMPARISON OF OUR METHOD WITH ALTERNATIVE APPROACHES ON THE PALLET BLOCK AND METAL DATASETS.

Dataset	Model	Rank-1	Rank-10	Runtime [s]
Pallet	Ensemble	<b>0.703</b> ± 0.079	<b>0.964</b> ± 0.020	287
	Graph [23]	0.526 ± 0.052	0.904 ± 0.045	50
	Image [9]	0.486 ± 0.032	0.799 ± 0.060	213
	Inception [49]	0.518 ± 0.05	0.918 ± 0.014	1, 116
Metal	Ensemble	<b>0.777</b> ± 0.054	<b>0.992</b> ± 0.007	257
	Color Variance	0.549 ± 0.057	0.909 ± 0.022	8
	Image [9]	0.360 ± 0.033	0.898 ± 0.047	224
	Inception [49]	0.681 ± 0.059	0.951 ± 0.013	1, 283

It is also again apparent, that our best individual sub-models are significantly outperformed by the ensemble model. Focusing on Rank-10 accuracy, the results of the four models in question converge between 79.9% and 99.2%, with the image-based model consistently providing the lowest re-identification accuracy. The accuracies obtained using the metal dataset were consistently higher than the ones using the pallet block dataset, implying a potential difference in difficulties between the two re-identification tasks. Further analysis of our ensemble model is found in the supplementary materials provided.

## V. CONCLUSION

In this contribution, the first study of heterogeneous siamese neural network ensembles for re-identification purposes has

been presented. Using such ensembles, we reached state-of-the-art performance on our datasets. With our novel approach, we obtain Rank-1 re-identification accuracies of 70% and 77% on the pallet block and metal datasets, respectively. One of our employed ensembling methods, the *Concatenation* method, a fairly simple approach, performs better than most of the more sophisticated ones evaluated in this work. In subsequent publications, we aim to study the use of a higher number of possibly even simpler models, while still aiming to maintain reliable re-identification results. We also show, how our ensemble approach can be applied and specified to different datasets, implying that it might be possible to apply our methods to further contrastive learning tasks. Since we restricted ourselves to the application for the re-identification of logistical entities, we invite the research community to apply similar approaches to other datasets. We also invite researchers to apply their approaches to the same datasets, to provide a comparison. Finally, we believe the further investigation of the explainability and trustworthiness of re-identification methods to be a promising research direction. For this reason, we expand upon this topic in our supplementary material in great detail.

## ACKNOWLEDGEMENTS

This research was funded by the Research Center Trustworthy Data Science and Security and the Lamarr Institute for Machine Learning and Artificial Intelligence.

This work is part of the project “Silicon Economy Logistics Ecosystem” which is funded by the German Federal Ministry of Transport and Digital Infrastructure.

## REFERENCES

- [1] T. Hastie, R. Tibshirani, J. Friedman, T. Hastie, R. Tibshirani, and J. Friedman, “Ensemble learning,” *The elements of statistical learning: data mining, inference, and prediction*, pp. 605–624, 2009. 1
- [2] L. Breiman, “Bagging predictors,” *Machine Learning*, vol. 24, pp. 123–140, Aug 1996. 1, 2, 4
- [3] Y. Freund and R. E. Schapire, “A short introduction to boosting,” in *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, pp. 1401–1406, Morgan Kaufmann, 1999. 1
- [4] M. Ganaie, M. Hu, A. Malik, M. Tanveer, and P. Suganthan, “Ensemble deep learning: A review,” *Engineering Applications of Artificial Intelligence*, vol. 115, 2022. 1
- [5] I. Deviatkin and M. Horttanainen, “Carbon Footprint of an EUR-sized Wooden and a Plastic Pallet,” in *E3S Web of Conferences*, vol. 158, EDP Sciences, 2020. 1
- [6] D. S. C. Packaging, “Pallet Production Specification - Part 1: Construction Specification for 800 mm × 1200 mm Flat Wooden Pallets; German version EN 13698-1:2003,” tech. rep., Beuth Verlag, 2004. 1
- [7] L. Krämer, J. Rutinowski, J. Endendyk, R. Vaut, and M. Roidl, “On blockchain-based token usage in cyber-physical production systems,” *Logistics Journal: Proceedings*, vol. 2022, no. 18, 2022. 1
- [8] J. Rutinowski, H. Youssef, A. Gouda, C. Reining, and M. Roidl, “The potential of deep learning based computer vision in warehousing logistics,” *Logistics Journal: Proceedings*, vol. 2022, no. 18, 2022. 1, 2
- [9] J. Rutinowski, T. Chilla, C. Pionzewski, C. Reining, and M. ten Hompel, “Towards Re-Identification for Warehousing Entities – A Work-in-Progress Study,” in *Proceedings of the IEEE Conference on Emerging Technologies In Factory Automation (ETFA)*, pp. 501–504, 2021. 1, 2, 6
- [10] J. Rutinowski, C. Pionzewski, T. Chilla, C. Reining, and M. Ten Hompel, “Deep learning based re-identification of wooden euro-pallets,” in *2022 21st IEEE International Conference on Machine Learning and Applications (ICMLA)*, pp. 113–117, IEEE, 2022. 1, 2

- [11] M. Ye, J. Shen, G. Lin, T. Xiang, L. Shao, and S. C. Hoi, "Deep Learning for Person Re-Identification: A Survey and Outlook," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 6, pp. 2872–2893, 2021. [2](#)
- [12] D. Yi, Z. Lei, S. Liao, and S. Z. Li, "Deep Metric Learning for Person Re-Identification," in *Proceedings of the International Conference on Pattern Recognition*, pp. 34–39, 2014. [2](#)
- [13] H.-X. Yu, A. Wu, and W.-S. Zheng, "Cross-View Asymmetric Metric Learning for Unsupervised Person Re-Identification," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pp. 994–1002, 2017. [2](#)
- [14] J. Rutinowski, B. Vankayalapati, N. Schwenzfeier, M. Acosta, and C. Reining, "On the applicability of synthetic data for re-identification," *arXiv preprint arXiv:2212.10105*, 2022. [2](#)
- [15] K. Islam, "Deep learning for video-based person re-identification: A survey," *arXiv preprint arXiv:2303.11332*, 2023. [2](#), [6](#)
- [16] Z. Ming, M. Zhu, X. Wang, J. Zhu, J. Cheng, C. Gao, Y. Yang, and X. Wei, "Deep Learning-Based Person Re-Identification Methods: A Survey and Outlook of Recent Works," *Image and Vision Computing*, vol. 119, p. 104394, 2022. [2](#)
- [17] S. D. Khan and H. Ullah, "A Survey of Advances in Vision-Based Vehicle Re-Identification," *Computer Vision and Image Understanding*, vol. 182, pp. 50–63, 2019. [2](#)
- [18] L. Wei, X. Liu, J. Li, and S. Zhang, "VP-ReID: Vehicle and Person Re-Identification System," in *Proceedings of the 2018 ACM on International Conference on Multimedia Retrieval*, pp. 501–504, 2018. [2](#)
- [19] Z. Jewell, S. Alibhai, F. Weise, S. Munro, M. Vuuren, and R. Vuuren, "Spotting Cheetahs: Identifying Individuals by their Footprints," *Journal of Visualized Experiments*, vol. 2016, 2016. [2](#)
- [20] M. Zuerl, R. Dirauf, F. Koeferl, N. Steinlein, J. Sueskind, D. Zanca, I. Brehm, L. v. Fersen, and B. Eskofier, "Polarbearvidid: A video-based re-identification benchmark dataset for polar bears," *Animals*, vol. 13, no. 5, 2023. [2](#)
- [21] J. C. Hermanson and A. C. Wiedenhoeft, "A Brief Review of Machine Vision in the Context of Automated Wood Identification Systems," *IAWA Journal*, vol. 32, no. 2, pp. 233–250, 2011. [2](#)
- [22] A. Klar, *Identifizierung von Ladehilfsmitteln auf der Basis Stochastischer Oberflächeninformationen*. PhD thesis, TU Dortmund University, Germany, 1995. [2](#)
- [23] S. Klüttermann, J. Rutinowski, C. Reining, M. Roidl, and E. Müller, "Towards graph representation based re-identification of chipwood pallet blocks," in *2022 21st IEEE International Conference on Machine Learning and Applications (ICMLA)*, pp. 1543–1550, 2022. [2](#), [3](#), [6](#)
- [24] Y. Kortli, M. Jridi, A. Al Falou, and M. Atri, "Face Recognition Systems: A Survey," *Sensors*, vol. 20, no. 2, 2020. [2](#)
- [25] A. Nambiar, A. Bernardino, and J. C. Nascimento, "Gait-Based Person Re-Identification: A Survey," *ACM Computing Surveys (CSUR)*, vol. 52, no. 2, pp. 1–34, 2019. [2](#)
- [26] N. Schwenzfeier, J. Rutinowski, M. Hesenius, C. Reining, and M. Acosta, "Generating embedding spaces for re-identifying pallets from their chipwood patterns," *Engineering Applications of Artificial Intelligence*, vol. 126, 2023. [2](#)
- [27] Y. Sun, L. Zheng, Y. Yang, Q. Tian, and S. Wang, "Beyond Part Models: Person Retrieval with Refined Part Pooling (and a Strong Convolutional Baseline)," in *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 480–496, 2018. [2](#), [6](#)
- [28] J. Rutinowski, T. Chilla, C. Pionzewski, C. Reining, and M. ten Hompel, "pallet-block-502 – A chipwood re-identification dataset," Sept. 2021. [2](#), [3](#)
- [29] J. Rutinowski, T. Chilla, and C. Pionzewski, "pallet-block-32965 – A chipwood re-identification dataset," Mar. 2022. [2](#)
- [30] S. Paisitkriangkrai, C. Shen, and A. van den Hengel, "Learning to rank in person re-identification with metric ensembles," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. [2](#)
- [31] Y. Zhai, X. Guo, Y. Lu, and H. Li, "In Defense of the Classification Loss for Person Re-Identification," *arXiv preprint: 1809.05864*, 2018. [2](#)
- [32] A. Zimek, R. J. Campello, and J. Sander, "Ensembles for unsupervised outlier detection: Challenges and research questions a position paper," *SIGKDD Explor. Newsl.*, vol. 15, p. 11–22, mar 2014. [2](#)
- [33] M. O. Sandim, *Using Stacked Generalization for Anomaly Detection*. PhD thesis, 2017. [2](#)
- [34] S. Klüttermann, C. Balestra, and E. Müller, "On the efficient explanation of outlier detection ensembles through shapley values," in *Advances in Knowledge Discovery and Data Mining* (D.-N. Yang, X. Xie, V. S. Tseng, J. Pei, J.-W. Huang, and J. C.-W. Lin, eds.), (Singapore), pp. 43–55, Springer Nature Singapore, 2024. [2](#)
- [35] X. Liu, H. Wang, Y. Wu, J. Yang, and M.-H. Yang, "An ensemble color model for human re-identification," in *2015 IEEE Winter Conference on Applications of Computer Vision*, pp. 868–875, 2015. [2](#)
- [36] A. Serbetci and Y. S. Akgul, "End-to-end training of cnn ensembles for person re-identification," *Pattern Recognition*, vol. 104, p. 107319, 2020. [2](#)
- [37] J. Wang, Y. Li, and Z. Miao, "Ensemble feature for person re-identification," 01 2019. [2](#)
- [38] Y. Yang, X. Liu, Q. Ye, and D. Tao, "Ensemble learning-based person re-identification with multiple feature representations," *Complexity*, vol. 2018, p. 5940181, Sep 2018. [2](#)
- [39] G. R. Koch, "Siamese neural networks for one-shot image recognition," 2015. [2](#)
- [40] S. Paisitkriangkrai, C. Shen, and A. van den Hengel, "Learning to rank in person re-identification with metric ensembles," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. [2](#)
- [41] M. Ye, X. Lan, Q. Leng, and J. Shen, "Cross-modality person re-identification via modality-aware collaborative ensemble learning," *IEEE Transactions on Image Processing*, vol. 29, pp. 9387–9399, 2020. [2](#)
- [42] S. Xu, E. Shijia, and Y. Xiang, "An ensemble model based on siamese neural networks for the question pairs matching task," in *CKKS Tasks*, 2018. [2](#)
- [43] S. Klüttermann, J. Rutinowski, and E. Müller, "The phenomenon of correlated representations in contrastive learning," in *International Joint Conference on Neural Networks (IJCNN)*. [2](#)
- [44] J. Rutinowski, J. Endendyk, C. Reining, and M. Roidl, "galvanized-636 – A galvanized steel re-identification dataset," Dec. 2022. [3](#)
- [45] E. Hoffer and N. Ailon, "Deep metric learning using triplet network," in *Similarity-Based Pattern Recognition* (A. Feragen, M. Pelillo, and M. Loog, eds.), (Cham), pp. 84–92, Springer International Publishing, 2015. [3](#)
- [46] S. Klüttermann and E. Müller, "Evaluating and comparing heterogeneous ensemble methods for unsupervised anomaly detection," pp. 1–8, 06 2023. [4](#)
- [47] D. Mienye and Y. Sun, "A survey of ensemble learning: Concepts, algorithms, applications, and prospects," *IEEE Access*, vol. PP, pp. 1–1, 09 2022. [4](#)
- [48] C. Wang, Q. Wu, M. Weimer, and E. E. Zhu, "Flaml: A fast and lightweight autml library," in *Fourth Conference on Machine Learning and Systems (MLSys 2021)*, April 2021. [4](#)
- [49] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," 06 2016. [6](#)
- [50] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," pp. 770–778, 06 2016. [6](#)
- [51] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," pp. 4510–4520, 06 2018. [6](#)