

Article

Machine Learning-Based Highway Truck Commodity Classification Using Logo Data

Pan He [†] , Aotian Wu [†] , Xiaohui Huang, Anand Rangarajan  and Sanjay Ranka ^{*}

Department of Computer and Information Science and Engineering, University of Florida, Gainesville, FL 32611, USA; pan.he@ufl.edu (P.H.); aotian.wu@ufl.edu (A.W.); xiaohuihuang@ufl.edu (X.H.); anand@cise.ufl.edu (A.R.)

^{*} Correspondence: ranka@cise.ufl.edu

[†] These authors contributed equally to this work.

Abstract: In this paper, we propose a novel approach to commodity classification from surveillance videos by utilizing logo data on trucks. Broadly, most logos can be classified as predominantly text or predominantly images. For the former, we leverage state-of-the-art deep-learning-based text recognition algorithms on images. For the latter, we develop a two-stage image retrieval algorithm consisting of a universal logo detection stage that outputs all potential logo positions, followed by a logo recognition stage designed to incorporate advanced image representations. We develop an integrated approach to combine predictions from both the text-based and image-based solutions, which can help determine the commodity type that is potentially being hauled by trucks. We evaluated these models on videos collected in collaboration with the state transportation entity and achieved promising performance. This, along with prior work on trailer classification, can be effectively used for automatically deriving commodity types for trucks moving on highways.

Keywords: freight analysis; scene text understanding; logo detection and recognition; commodity classification; deep learning; intelligent transportation system



Citation: He, P.; Wu, A.; Huang, X.; Rangarajan, A.; Ranka, S. Machine Learning-Based Highway Truck Commodity Classification Using Logo Data. *Appl. Sci.* **2022**, *12*, 2075. <https://doi.org/10.3390/app12042075>

Academic Editor: Lidia Jackowska-Strumillo

Received: 7 January 2022

Accepted: 6 February 2022

Published: 16 February 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Approximately 125 million households, nearly 7.7 million business establishments, and 90,000 governmental units in the U.S. require efficient and reliable movement of freight [1]. Freight transportation has become an indicator of economic growth and regional development, which makes freight analysis an increasingly important area. The main objective of freight analysis is to reduce freight transit time and transportation cost and improve the reliability of freight movement. In addition, it is beneficial in mitigating traffic congestion, better planning land use, and improving economic competitiveness [2].

Conventional freight data collection occurs through questionnaires filled manually by carriers, shippers, or receivers regarding the commodity type, origin, and destination [2]. However, survey-based methods have several apparent drawbacks, such as low response rates, unknown data reliability, and high time cost [3]. Although trucking companies are likely to keep detailed records of their truck and commodity information, most of them are unwilling to make these records public due to possible competition. As a consequence of the above limitations, current freight data have limited reliability, completeness, and efficiency, thereby reducing its applicability in downstream analysis and processing.

The video-based sensing technique has achieved consistent improvements in supporting a cost-effective and accurate traffic system. This has led to increasingly popular vision-based solutions in transportation applications aiming to improve efficiency and reduce costs. However, it is known that applying these solutions to freight classification is still in its infancy due to the challenges of the absence of a public dataset, real-time requirements of identifying objects, and large variations in environmental conditions.

Among various transportation modes, truck-based transportation acts as the major mode of commodity shipments in the U.S., carrying 62.7% of the total commodity tonnages

and 61.9% of the commodity values, according to [4]. Truck-based freight transportation is expected to grow in the next decade according to ATA's freight forecast. In response to this, the research community has developed various classification models for trucks and trailers, relying on the input data collected from traffic sensors such as weigh-in-motion (WIM), inductive loop detectors (ILD), and cameras [5–8]. However, the major limitation is that they fail to reveal the carried cargo from the limited cues identified from trucks.

Large-scale road-based freight data analysis is of great need to alleviate problems of traffic congestion, bottlenecks, and truck empty-mile wastage. In this paper, to the best of our knowledge, we present a fundamental video processing approach for freight analysis based on fine-grained visual information of truck images, e.g., logos and texts, collected in real-world environments. Logos provide important cues in identifying commodity types. Preliminary works [9,10] have shown effectiveness in detecting and recognizing license plates and predefined sets of vehicle brand logos. However, they cannot meet the requirement of freight analysis, where we are interested in reporting a broad range of logos carried by trucks with potential extensions in the future. In other words, the desired approach should be extensible, as it would be impractical to provide an exhaustive list of company logos, and new logos will likely show up. Therefore, an approach to bridging the gap between logo recognition and freight classification is desperately needed to supplement the existing data sources.

Prior work has successfully inferred the commodity type based on the trailer types, e.g., enclosed or tank, recognized from some truck images [6]. However, it fails to handle the majority of trucks with enclosed trailers. Fortunately, we might still infer the commodity types by leveraging those company logos potentially on truck bodies, which remains as the non-trivial task of detecting and recognizing logos on trucks. The challenges mainly lie in several factors such as uncontrolled illumination, occlusions, and background clutter. To address all mentioned problems, we have made the following contributions (A preliminary version of this manuscript has been presented in a conference [7], where we tackle the commodity classification by only utilizing text information of logos on trucks. In this paper, we leverage both text and image content information and develop an integrated approach to combine predictions from both text-based and image-based solutions):

- A coarse-to-fine universal logo detector that can estimate the locations of previously unseen logos. Since the detector is class-agnostic and not limited to a certain set of logos, it applies to a wide range of logos;
- An integrated approach to accurately link the detected logos to a company dataset customized for traffic scenarios. It proposes to leverage both text and image information from logos by a combination of texts generated from state-of-the-art solutions and logo types identified from our proposed logo matching method. The developed approach can be effectively extended to new logo classes and companies in which the traffic agency is highly interested;
- A novel end-to-end road video processing system to provide real-time dynamic commodity information by deploying sensors and edge devices in locations of interest. This utilizes the NAICS (North American Industry Classification System) taxonomy with searches aimed at commodity type inference based on the name of the company.

In addition, we have developed a new benchmark on commodity classification using logos. To the best of our knowledge, this is the first attempt at doing so, which could be beneficial in helping traffic engineers and researchers better evaluate their developed models.

The rest of the paper is organized as follows. Section 2 reviews previous studies that include relevant topics or techniques. Section 3 describes the overall commodity classification pipeline and details of all developed logo detection and recognition approaches. Section 4 describes the experimental settings and results of the different developed approaches. Finally, Section 5 concludes with an overall summary (while discussing limitations) and presents opportunities for future work.

2. Related Work

2.1. Logo Detection

The development of our universal logo detector benefits from advances in object detection techniques in recent years. The goal of object detection is to locate and classify certain categories of objects in given images and label the bounding box around the detected regions with a confidence score. Object detection techniques have also been used in logo detection. Several logo datasets have been released, such as FlickrLogos [11], LOGO-Net [12], WebLogo-2M [13] and QMUL-OpenLogo [14], which have catalyzed quite a few works on logo detection. Romberg et al. made use of quantized representations of basic spatial structures detected in logo images [11]. Pan et al. utilized CNN features for vehicle manufacturer brand recognition [15]. Su et al. trained one deep learning logo detection model on a synthesized logo dataset which leveraged extensive labelling costs [16]. Montserrat et al. adopted a two-stage approach and cropped regions of interest from the original image; this was followed by classification [17]. Our image-based recognition approach adopted a two-stage approach for the sake of better generalization ability.

2.2. Deep Metric Learning

Since building an exhaustive database for logos appeared on trucks is impractical, the extensibility logo recognition approach is of great importance. In other words, the approach should be able to quickly adapt to ever-growing logo classes. To address this challenge, we adopt deep metric learning, which essentially ‘learns to compare’ image pairs by shrinking intra-class distance while expanding inter-class distance. Several deep metric learning approaches have been proposed and been applied in various applications, including face identification [18], person re-identification [19,20], and image retrieval [21]. Metric learning is known to have a stronger discriminative power since the inter-class distance is directly maximized during training. Additionally, a learned metric can be generalized to unseen testing classes using the discriminative attributes and pairwise distances extracted in metric learning [22].

2.3. Content-Based Image Retrieval (CBIR)

CBIR seeks to solve the image retrieval problem using computer vision techniques. Two innovative works significantly advanced the efficiency and accuracy of CBIR. The first one is scale-invariant feature transform (SIFT) [23], which can robustly detect key points under changes in scale, rotation, noise, and illumination. The second one is the bag-of-visual-words (BoW) [24] model, which treats an image as a document and treats feature representations as words. The BoW model gives a compact representation of the whole image based on quantization of discriminative local features. In this paper, the BoW features were extracted to match with the features of samples in the gallery set, the effectiveness of which were measured through experiments.

3. Methodology

In this section, we describe the developed approach in detail. Logos, serving as the outward expression of brands, often consist of letters or texts with large variations in colors, font styles, and graphical figures. Logos can appear anywhere on truck bodies, making it hard to leverage any prior knowledge of context and placement. It becomes more challenging when moving to truck images of low resolutions, poor light and weather conditions, and diverse view angles. The existing work on logo detection relies on large datasets of sufficient fine-grained annotations on logos such as bounding boxes and logo types, which, however, are often unavailable in real-world traffic scenarios. In this paper, we have developed techniques following broad approaches for text-based and image-based logos as follows:

- For text-based logos, we employ state-of-the-art solutions for text detection [25,26] and recognition [27,28] to obtain raw text predictions with high accuracy. These predicted

- logo texts were matched to our built commodity database by comparing these texts with recorded company names via simple string matching algorithms;
- To identify image-based logos that do not contain text information, we developed a novel two-stage image-based approach: a universal logo detection stage that outputs all potential logo positions within images, followed by a logo recognition stage designed to incorporate various advanced image representations. If the universal logo detector returned bounding boxes with high confidence scores, we cropped out the corresponding image regions to obtain logo candidate images. Each logo image was processed by our developed model to obtain the image features. These features were matched with pre-computed image features stored in a database containing ground truth logo images of interest. It is considered as a correct match if the matching score surpasses a certain threshold; otherwise, we ignore this logo candidate;
 - For an automated system, it is not known a priori whether a given truck contains text- or image-based logos (or neither). We integrated the text-based approach and image-based approach using the pipeline shown in Figure 1. This pipeline naturally combines the advantages of both approaches in one integrated approach, and the later experiments in Section 4 validate its effectiveness in improving the performance of commodity classification.

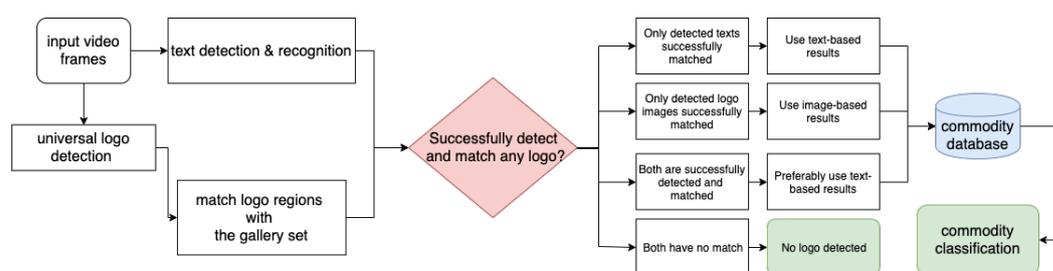


Figure 1. The overall pipeline of commodity classification.

3.1. Text-Based Logo Detection and Recognition

We implemented state-of-the-art scene-text solutions, extending the previous research work on EAST [25] for text detection and CRNN [28] for text recognition. The pipeline is shown in Figure 2. It consists of the following steps:

1. For a truck image acquired from a roadside traffic camera, we generate the text line/word map based on the features extracted from the multichannel FCN (fully convolutional network) model in EAST, which helps identify text regions of interest;
2. To remove detected texts largely overlapping each other, we apply the standard post-processing technique called NMS (non-maximum suppression), which results in oriented bounding boxes to indicate text line/word locations;
3. We then crop out the image region within each oriented bounding box and feed it into the CRNN model to translate the text image to a pure text string;
4. Finally, we match the predicted text string to a predefined logo class via additional techniques for word prediction and string matching.

We have achieved a high recall and a competitive recognition accuracy following the developed algorithms. The figures below show sample outputs. Though the texts might not be perfectly predicted, e.g., they might miss character prediction or wrongly recognize a few characters, this could be suitably fixed via the spelling correction methods that are publicly available, which we use to demonstrate in Section 4 that this text-based logo solution is beneficial in improving the performance of the overall pipeline.

It is worth mentioning that text-based logo detection and recognition has its limitations when solving the commodity classification. For those logos not containing any text or being too complicated to be successfully identified as texts, the text-based solution is no longer suitable. Instead, we have to seek a fundamentally different approach based on

the image content, which leads us to the image-based approach incorporating various advanced image representations for logo data described in the subsequent section.

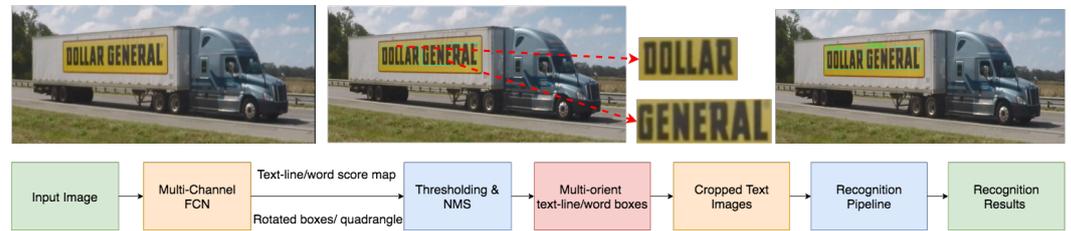


Figure 2. The detection and recognition pipeline of the text-based solution.

3.2. Image-Based Logo Detection and Recognition

Our image-based logo approach is general and works for images (and complex-font texts) using the following image-based approach: (i) a universal logo detector and (ii) a feature-matching-based logo recognizer. Figure 3 shows the training and inference pipeline of our staged approach.

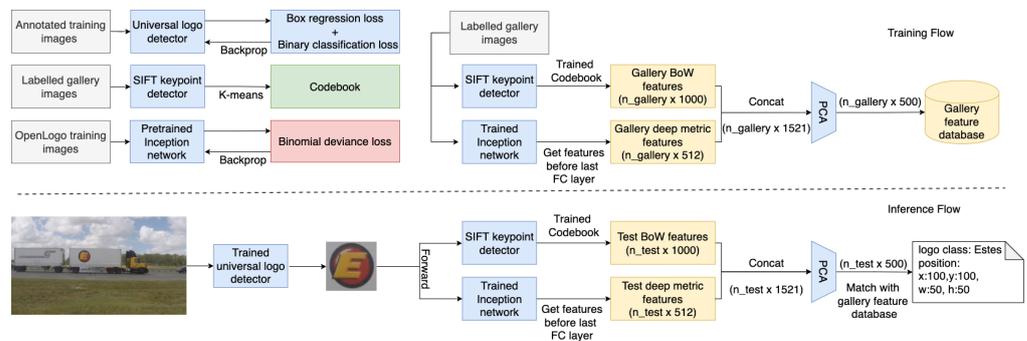


Figure 3. The flow diagram of image-based solution on both training and inference.

3.2.1. Universal Logo Detector

State-of-the-art deep learning methods have been discussed in [29] to train and evaluate object detectors on localizing and identifying logos in a closed set of classes. Logo detection is inherently a challenging task due to the presence of varying challenging factors in occlusions, uncontrolled illumination, and background clutter. Under these conditions, logo detectors tend to be susceptible to context changes. The authors presented a vivid example of a detector trained to localize a logo that only appears on shoes in all the training images. The detector fails to detect the same logo appearing on a coffee mug during the model inference.

In real traffic scenarios, we expect to encounter a large number of previously unseen logo classes. A logo detector trained with a fixed set of classes clearly cannot detect logos that are not in the fixed set. This requires retraining the logo detector and annotating more training data for new incoming logo classes, which makes it impractical because traffic agencies are interested in detecting all kinds of logo data to further identify the carried cargo within trucks. To overcome this limitation, a more promising direction is to develop a logo detector without fine-grained labeled training data for new logo classes. As shown in Figure 3, we designed the universal logo detector to localize any potential logos in contrast to popular logo detectors that detect and identify a fixed set of logos. Our proposed model alleviates the problem of collecting and annotating new training logo data for any future logos to be detected.

One-stage universal logo detector. The training was based on a large number of data images with bounding box annotations of logos. The model is designed to learn an abstract representation of all kinds of logos in the training stage. Thus, it can work with arbitrary logo data in the inference stage. In particular, the model is trained in a class-agnostic way:

all logo regions belong to one single category (the logo category). We consider other regions as the background category, thus creating a classifier in a binary fashion. By doing this, we intentionally remove specific logo class information and force the model to learn a generic representation of all logos. We refer to this model as the one-stage universal logo detector (one-stage ULD). We adapt the implementation of the popular Faster R-CNN [30] with a VGG16 backbone [31] for logo detection due to its simplicity and reasonably fast speed. The training follows the original papers [30]. We did not explore other advanced backbones (e.g., ResNet [32], DenseNet [33]) and detection approaches (e.g., MaskRCNN [34]) that could further improve the detection performance, as it is not the focus of our paper. We highlight the concept of universal logo detection for reporting the presence of any logo-like regions in this paper.

Coarse-to-fine universal logo detector. It is worth noting that logos on trucks (which are the vehicles of interest) usually appear on the truck bodies. To take advantage of this prior information, we further propose a coarse-to-fine universal logo detector by roughly estimating truck bounding boxes and conducting the universal logo detection only within each bounding box. The process takes raw images and determines the presence of truck objects within images by adopting the state-of-the-art detector called YOLO [35]. This design significantly improved the localization precision as shown in the later experiment section. After we obtained logo locations, we cropped all logo regions from the image and forwarded them to the logo recognition model to infer the logo/brand classes. The bottom part of Figure 4 shows sample outputs of our developed universal logo detector.

3.2.2. Reverse Image Search Logo Recognizer

In our preliminary experiments, we implemented a pipeline of logo detection and recognition using the developed universal logo detector and a commercial reverse image search. The reverse image search is a content-based image retrieval (CBIR) query approach [36] in which we provide the system with a sample image (search query) to search for related concepts about this image. We utilized the popular Google “Search by Image”, which allows us to search for related images just by uploading an image or image URL. It analyzes the submitted picture by constructing a mathematical model, comparing it to a large number of images in Google databases, and returning similar images and their annotations. The obtained results from this pipeline are not satisfactory as it usually reports relatively random and noisy predictions. This pipeline is difficult to customize to the task of freight analysis.

3.2.3. Feature-Matching Logo Recognizer

In this approach, we treated logo recognition as an image retrieval problem with a few sample images for each class. We collected the gallery images from the Internet, which included roughly 30 images per logo class. The gallery images are used as the templates against which all logo predictions are matched. Two sets of features were extracted to represent logo images, namely deep metric features and the BoW (Bag of Visual Words) features. The deep metric features draw upon recent advances in deep metric learning and have the capability of extracting high-level discriminative semantic information for similarity measurements. The BoW features instead extract low-level image information such as textures, corners, and edges for logo matching. The BoW features are commonly used in image retrieval for research problems and industrial applications. After extracting these two types of image feature representations, we combined them to obtain the fused features.

Deep Metric Features. Inspired by [14], we trained a DCNN (deep convolutional neural network) classifier on the QMUL-OpenLogo dataset and used the output features of the second-to-last layer as the feature representation for a given logo image. We did not train the classifier directly on the logo classes of interest, considering that new unseen logo classes might be added. Instead, we aim at projecting logo images into a feature space such

that logos from different classes are separable. Therefore, we seek to increase the inter-class distances of logo features, bringing us to the popular deep metric learning (DML).



Figure 4. Sample outputs from the developed approaches.

The core concept of DML is to find a good representation of images with a good metric for similarity measurement. To measure the similarity of two feature vectors, we simply choose to use the common cosine similarity because it is bounded and invariant to feature magnitude; its formula is shown as follows:

$$\cos(\mathbf{x}_i, \mathbf{x}_j) = \frac{\mathbf{x}_i \cdot \mathbf{x}_j}{\|\mathbf{x}_i\| \cdot \|\mathbf{x}_j\|}, \quad (1)$$

where \mathbf{x}_i and \mathbf{x}_j are the feature vectors extracted from image i and j using the DML model. If the score of cosine similarity is close to 1, it means these two feature vectors are likely to come from the same logo class, otherwise not. Formally, pairs are called positive pairs if they have the same label; otherwise, they are called negative pairs. To make logo features separable, we prefer a trained model that assigns higher similarity to positive pairs and lower similarity to negative pairs. We chose to use binomial deviance [37] as the loss function with the formula:

$$L = \sum_{i,j} \left[\frac{1}{P_i} \sum_{\mathbf{x}_i = \mathbf{x}_j} \log[1 + e^{\alpha(\lambda - s_{ij})}] + \frac{1}{N_i} \sum_{\mathbf{x}_i \neq \mathbf{x}_j} \log[1 + e^{\beta(s_{ij} - \lambda)}] \right] \quad (2)$$

where P_i and N_i denote the count of positive pairs and negative pairs related to \mathbf{x}_i , respectively. s_{ij} denotes the similarity of pair $(\mathbf{x}_i, \mathbf{x}_j)$. α , β , and λ are hyperparameters, which are chosen based on the best heuristic setting of [14]. The hyperparameters α and β were set to be 40 and 0, respectively. The hyperparameter λ was set to be 0.5. The loss function lays more emphasis on hard samples where positive pairs get low scores or negative pairs get high scores, which enforces the model to find more discriminative features. The deep metric learning model was implemented in PyTorch. We used the Inception network [38] as the backbone with a global pooling layer and a fully connected layer added on top of it. The network was trained in a pairwise way by gradient descent using the Adam optimizer [39].

Bag-of-words Features. Deep learning models are good at extracting high-level semantic features, whereas, for logos, low-level features (such as textures, corners, and edges) can be useful for recognition. To incorporate low-level features, we make use of the bag-of-words (BoW) features, which are the most commonly used image representations in image retrieval literature [40,41]. In the training stage, local features were extracted using the scale-invariant feature transform (SIFT) descriptor [42]. This resulted in a large number of features for each logo image. To further find a compact and fixed-length representation of each image, feature quantization is required. This was achieved by visual codebook learning, after which, each local feature can be assigned to a visual word in the codebook. In this way, an image could be discriminatively represented by a histogram of these visual words. The codebook was learned by clustering gallery local features using the k-means algorithm and regarding the centers of clusters as visual words. In the inference stage, local features were obtained by the same method following the training stage. For each feature, the nearest visual word in a trained codebook was found and assigned to the corresponding bin of histogram, which resulted in a fixed-length feature vector for each testing image.

Fused Features. Feature fusion has been proven to be effective when the classes of features to be fused are heterogeneous, which can result in better performance than the best single class of features. BoW features are essentially sparse histograms of low-level local features, and DML features are high-level semantic features with large receptive fields. Both of the extracted features are heterogeneous. Therefore, combining these two sets of features tends to give a better result, as they are complementary to each other. Following this intuition, we concatenated these two features and obtained 1,512-dimensional feature vectors. We then used principal component analysis (PCA) for dimensionality reduction [43]. This resulted in 500-dimensional reduced feature vectors as the final representation. The experiment results demonstrated the effectiveness of our fusion approach.

3.3. Integrated Logo Model

The main advantage of the text-based approach is its robustness when the logo is mostly text. Though the model outputs incorrect or incomplete text predictions in some challenging traffic scenarios, it works well in most cases. The errors from the partially correct predictions can be corrected by finding the most similar company name using approximate string matching algorithms. This approach does not work when the logo consists essentially of non-text images.

The image-based approach has the potential to cover scenarios where the text information is not available on truck bodies or the text-based approach fails to detect any texts from trucks. Due to the high recall of the universal logo detector, we are likely to get a good subset from all potential logos. However, the accuracy of this approach may not be high as compared to a text-based approach when only logos that consist of text are present.

We now discuss potential approaches that combine both of these approaches. Let the set of detected bounding boxes from the text-based approach be denoted as $\mathbf{B}_{\text{text}} = B_i^k$, $k = 1, 2, \dots, N_{\text{text}}$, where N_{text} denotes the number of detections from the text-based model. Each B_i^k is associated with two confidence scores, namely the detection score (denoted as S_i^0), which measures how likely the detected region contains text, and the matching score (denoted as S_i^1), which measures the similarity between the detected text with the matched text in defined logo classes. Similarly, for the image-based approach, the set of detected bounding boxes is denoted as $\mathbf{B}_{\text{image}} = B_i^k$, $k = 1, 2, \dots, N_{\text{image}}$; each B_i^k is associated with one matching score (denoted as S_i), which measures the similarity between the detected region and the matched one in the collected logo gallery. There are two possible approaches for integrating these models:

1. Text-focused Approach: If both text-based and image-based models detect a logo in the same location, (i.e., $\text{IoU}(B_i^m, B_i^n) > 0.3$), we rely on the label from text-based model;
2. Combined Approach: We train a decision tree classifier with four output classes indicating whether to use the text-based model's result, the image-based model's result, neither or both (corresponding to 1, 2, 0, and 3). Ambiguity arises at the testing

time when the image-based and text-based models give different predicted labels, while the classifier outputs 3. We resolve the issue by preferring the text-based model's result in this case.

The later experimental evaluation shows that the first integration approach performs slightly better on our collected testing dataset. This is potentially because it consists of more text-based logos. The second approach has a better generalization ability since it can automatically learn the parameters that can combine the outputs of the two approaches. It can also be modified accordingly to adapt to a dataset with different text versus non-text distributions. A user can choose either one of the two approaches, based on the characteristics of the target dataset.

3.4. Commodity Classification with Logo Data

The linkage between logo recognition and the commodity classification is provided by a commodity database we built. We utilized the North American Industry Classification System (NAICS), a standard that classifies business establishments with the aim of collecting and analyzing business-related statistical data. It is a comprehensive and well-structured system that classifies economic activities hierarchically into levels of groups, such as sectors, subsectors, and industry groups.

We built our commodity database by searching the NAICS code for each logo class and stored its corresponding commodity description. Samples of NAICS code and commodity description correspondences are shown in Table 1. With the database, the results from our logo detection and recognition pipeline are linked to their commodity description. This process completes our commodity classification solution. To the best of our knowledge, the proposed pipeline is the first attempt in this direction.

Table 1. Samples of the NAICS code searching.

NAICS Code	Description	NAICS Code	Description
311919	Other Snack Food Manufacturing	485119	Other Urban Transit Systems
337127	Institutional Furniture Manufacturing	488510	Freight Transportation Arrangement
424490	Other Grocery and Related Products Merchant Wholesalers	484230	Specialized Freight (except Used Goods) Trucking, Long-Distance
445110	Supermarkets and Other Grocery (except Convenience) Stores	532120	Truck Utility Trailer and RV Rental and Leasing
484121	General Freight Trucking Long-Distance Truckload	551112	Offices of Other Holding Companies

4. Experiments

In this section, we provide the details and statistics of the collected dataset and conduct extensive experiments to evaluate each module of our pipeline along with carefully designed ablation studies.

4.1. Dataset Collection and Processing

Benchmark Datasets. We evaluated our logo detection and recognition approaches on video frames captured by roadside cameras provided by the Florida Department of Transportation (FDOT). Among all logos shown up in the recorded videos, We picked 26 logo classes based on frequency of occurrence, which contains several top carrier companies in the US (<https://www.ttnews.com/top100/for-hire/2019> Accessed: 10 June 2020). When choosing logo classes, we also diversify the classes by including styled text logos, shape-based logos, and logos shown on different types of trailers. The chosen 26 classes do not represent full coverage of all logo classes of interest but are illustrative to evaluate our proposed approach. On one of the roadside videos, we annotated all logos that belong to the 26 chosen logo classes as a dataset, referred to as the Annotated Logo Dataset (ALD).

Each annotation consists of a bounding box around the target logo and the logo class. This dataset consists of 4486 images and 5020 logos and is used to evaluate logo detection and recognition performance. Detailed distributions of logo classes are shown in Table 2. In addition, we collected a gallery logo dataset (GLD) from the internet. For each logo class, we collected around 30 samples. This dataset is used in our feature-matching logo recognizer. Utilizing both datasets, we are able to execute and evaluate the logo detection and recognition pipeline for freight classification.

To provide a more in-depth evaluation and analysis of our logo recognition pipeline, we divided the 26 logo classes into 3 groups ('easy', 'medium', and 'difficult') according to the recognition difficulty. The detailed division can be found in Table 2 and samples of each group can be found in Figure 5. Most of the 'easy' logos tend to have a relatively clean background and high contrast between texts and the background. For example, we can easily separate the Dollar General text of the dark color from its smooth and single background with the yellow color. The model can extract discriminative features from these logos. We define logos containing multiple text lines as the 'medium' logos, which covers the typical cases such as Heartland Express and US Foods. The challenge mainly comes from the text arrangement. In addition, we have to handle logos of different colors and textures. The 'difficult' logos is very challenging due to their artistic fonts and figures (such as 'OD' and 'E'), extremely small text sizes, and low contrasts and reflective lighting conditions caused by logo and truck compartment materials.

Training Datasets. To train our universal logo detector and deep metric learning model, we exploited the existing large logo dataset called QMUL-OpenLogo [14] due to its rich annotations of logo instances with diverse appearances and background contexts. The QMUL-OpenLogo dataset consists of 27,083 images from 352 logo classes.

Table 2. Logo distributions of the Annotated Logo Dataset. We use (E), (M), and (D) to represent logo classes in groups 'easy', 'medium', and 'difficult'.

Logo Class & Group	Images	Logo Class & Group	Images	Logo Class & Group	Images	Logo Class & Group	Images
Ashley (E)	83	E (D)	248	Lays (M)	64	UPS (D)	236
Atlas (M)	52	FedEx (E)	1128	OD (D)	392	US Foods (M)	163
Budget (M)	47	HamburgSUD (M)	63	Opies (D)	51	Werner (D)	142
CarrollFulmer (M)	30	HeartlandExpress (M)	245	Prime (E)	48	XTRA (E)	489
Celadon (E)	107	heyl (M)	50	RBI (D)	281	YRC (M)	53
Davis (D)	95	JNJ (M)	168	SouthernAG (M)	174	Total	5,020
Dollar General (E)	199	Landstar (E)	362	Sunstate (E)	50		

Evaluation Protocols. We adopted a standard object detection evaluation protocols to evaluate our developed pipeline. Specifically, we adopt two commonly used metrics, recall, and precision, together with the average precision (AP) that measures the detection accuracy of a detector. AP computes the average precision for recall values ranging from 0 to 1. The general definition has the formula:

$$AP = \int_0^1 p(r) dr, \quad (3)$$

where $p(r)$ is the precision value at the recall value r . In practice, the equation is replaced with a finite sum over several recall values, such as the 11-point interpolated AP used in the Pascal VOC challenge [44] that is defined as the mean precision at a set of 11 equally spaced recall values ($[0, 0.1, 0.2, \dots, 1]$). We follow the new evaluation protocol of the Pascal VOC challenge where they use all data points, rather than interpolating only 11 equally spaced points [44] (we used the open-source evaluation tool Object-Detection-Metrics from: <https://github.com/rafaelpadilla/Object-Detection-Metrics#interpolating-all-points> Accessed 25 July 2019). The mean average precision (mAP) is the average of AP over all classes or or

categories. To decide whether a detection is correct, we calculate *IoU* (Intersection over Union) between predicted logos and ground truth logos. If the *IoU* is greater than a certain threshold (such as 0.5), the prediction is considered true; otherwise, it is considered false. The formula of *IoU* is provided as follows:

$$IoU = \frac{|B \cap B_{gt}|}{|B \cup B_{gt}|}, \quad (4)$$

where B and B_{gt} represents the predicted bounding box and its corresponding ground-truth, respectively.



Figure 5. Samples of logo classes with different groups ('easy', 'medium', and 'difficult').

4.2. Experimental Results

In this section, we evaluated all three proposed approaches for the following: evaluation of the universal logo detector, evaluation of logo classification, and evaluations on end-to-end logo detection and recognition. We conducted ablation studies for each step to verify the effectiveness of our model design. These studies illustrated and detailed the advantages and disadvantages of model variant of each component, which sheds light on exploring freight classification using a particular approach.

4.2.1. Universal Logo Detector

We compared one-stage and coarse-to-fine universal logo detectors (ULD) on detection performance in terms of the recall (Rec), precision (Prec), and the average precision (AP) (Table 3). With $IoU = 0.3$, the one-stage ULD detector achieved a recall of 80.0%. The coarse-to-fine ULD detector achieves a higher recall of 85.7%, which is beneficial to the logo recognition stage, where the low precision can be further improved. As mentioned, the one-stage ULD detector directly predicts bounding boxes of logos within truck images, while it ignores the prior information that logos usually appear on the truck bodies. After we added this prior information in the coarse-to-fine ULD detector, we obtained consistently better performance with all different IoU threshold values in terms of the recall (+5%) and average precision (+5%).

Table 3. Comparisons between variants of the universal logo detectors.

	IoU = 0.1			IoU = 0.3			IoU = 0.5		
	Recall	Precision	Average Precision	Recall	Precision	Average Precision	Recall	Precision	Average Precision
One-stage ULD	83.7	68.9	74.9	80.0	65.9	68.3	69.9	57.6	56.4
Coarse-to-fine ULD	88.1	65.3	77.9	85.7	63.5	73.5	73.5	54.5	60.8

4.2.2. Feature Matching-Based Logo Recognizer

We evaluated the logo recognizer with different feature representations. To only evaluate the performance of the logo recognition component, we assumed that logo regions are already available to the recognizer, which was obtained by directly cropping out logo regions using ground truth annotations from the ALD dataset.

The results are reported in Table 4 with the top- k accuracy, which is the fraction of test images for which the correct label is among the k most probable model predictions. We chose different values of k for a comprehensive evaluation. The GLD dataset was used as template samples for feature matching.

Table 4. Evaluations on logo recognition models with different feature representations.

	Top-1 Accuracy	Top-3 Accuracy	Top-5 Accuracy
Deep Metric	90.0	96.3	97.1
Bag of Words	88.3	95.8	97.3
Fused	95.3	97.9	98.3

We were able to achieve the top-1 accuracies of 90.0% and 88.3% with deep metric features and bag-of-words features, respectively. We achieved the best performance if we fused them, which demonstrated that these two features are indeed partially complementary to each other.

4.2.3. Integrated Logo Model

We evaluated the developed logo models using only the text-based approach, only the image-based approach and the integrated approach. It is worth noting that certain logo classes (such as 'davis') are 'relatively unclear' and difficult for the image-based approach to recognize, as they are likely to be cluttered with scene context. In many cases, it may be worth considering them as texts rather than as image-based logos. This can be effectively addressed using the proposed text-based approach. For example, the 'SouthernAG' logo class can be handled by the text-based solution with high recall and precision, although it achieves poor performance with the image-based approach.

Image-based Logo Detection and Recognition. We evaluated the image-based logo solution using the mAP metric. We achieved an mAP of 66.2%. Given the relatively small number of samples for each logo in our datasets, these results are promising. We expect additional improvements as additional annotated images are added to the dataset.

Text-based Logo Detection and Recognition. The text-based solution has a high precision on matching texts to company names. In general, it achieved higher precision compared to the image-based solution. The recall is slightly worse, which can be partly attributed to the fact that the text detection tends to predict tighter bounding boxes around text regions. The logo region is usually larger than the text regions, as it usually consists of both text and figure regions. Because the ALD dataset is annotated as logos, bounding boxes predicted by our text solution are expected to be smaller than the ground truth annotations, which can worsen the recall.

Ablation Study on Integration Approaches. To evaluate the two integration approaches we presented in Section 3, we apportioned the ALD dataset into training and testing set with a 70–30 split. The evaluation result of the two approaches on the testing set of ALD is shown in Table 5. On the apportioned testing set, approach 1 (text-focused approach) yields better overall performance in terms of mAP and Acc, while approach 2 gives higher precision. The results may vary on different datasets. We choose to use approach 1. The subsequent experiments were conducted on the whole ALD dataset (as opposed to the apportioned testing set) because approach 1 does not require additional training.

Table 5. Ablation study of integration approaches.

Text-focused Approach				Combined Approach			
Recall	Precision	Average Precision	Accuracy	Recall	Precision	Average Precision	Accuracy
87.1	69.8	82.2	87.2	82.6	78.7	73.2	86.4

Integrated Logo Model. The end-to-end logo detection and recognition performance were measured on the ALD dataset. Besides the aforementioned detection metrics (Rec, Prec, and mAP), we also evaluated the classification accuracy (Acc), where a single predicted class is obtained by a majority voting of all bounding boxes' labels detected in one frame.

The experimental result is showed in Table 6. We obtained an mAP of 81.2%, surpassing both the text-based and image-based approach significantly. Given that there were 26 logo classes, these results are promising. As discussed earlier, additional annotation, in particular for the difficult classes, should further improve the overall accuracy.

Table 6. Evaluations of the text-based and image-based approaches on three logo groups.

Logo Groups	Image-Based			Text-Based			Integrated Approach		
	Precision	Recall	Average Precision	Precision	Recall	Average Precision	Precision	Recall	Average Precision
Easy	74.7	88.5	84.6	90.7	95.6	91.7	86.5	98.0	95.1
Medium	60.3	83.1	69.9	93.4	90.6	88.9	72.6	95.1	90.2
Difficult	45.4	50.8	39.5	39.9	24.5	23.4	56.2	59.5	51.4
Overall	60.7	76.1	66.2	71.5	67.3	65.6	72.5	86.4	81.2

As can be found in Table 6, the text-based approach performs well in the easy and medium categories. It fails to detect logo classes such as 'OD', 'Opies', and 'E', where 'E' and 'OD' logos are designed with artistic fonts and figures. The 'Opies' logo usually appears on the body of the tank truck, where the compartment is made of reflective materials. The lighting reflection causes the text-based approach to fail to detect 'Opies'. These studies can help traffic agencies customize their specific tasks by choosing a particular solution considering the characteristics of the data.

Finally, we provide the run-time for each module in the following. On average, the text-based logo recognizer takes 1.95 s, the coarse-to-fine ULD takes 2.15 s, the feature-matching based logo recognizer takes 1.61 s, and the final commodity classification through database lookup takes 0.04 s for each frame. The experiments are performed on an NVIDIA Titan V GPU. The speed can be boosted using batch processing.

5. Summary and Conclusions

A vision-based freight classification approach has been proposed in the present paper. Our proposed solution consisting of text-based and image-based branches is able to capture most existing logos. Both our text-based and image-based solutions are general and can easily be extended to new logo classes. Our text-based approach using advanced scene-text solutions produces highly accurate predictions when the logo is mostly text. To extend it to new logo classes, only the text strings of the logos are needed. Our image-based approach serves as a complement and deals with logos with little or no text. It first detects all potential logos, and then performs feature matching with samples of different logo classes. For potential new logo classes, only around 30 samples need to be collected per class. Furthermore, we have developed a new freight classification benchmark based on logo data. To our best knowledge, ours is the first dataset collected to evaluate freight classification based on logo data. It can be useful in providing traffic engineers and researchers with a dataset to systematically evaluate their developed freight classification models.

We showed through experiments that our overall accuracy of 80% for 26 chosen logos is very promising. However, the logos we found from recorded highway videos (from the state of Florida) are not exhaustive, and the current study is limited by the size of the collected dataset. To further improve the accuracy and make our approach feasible in

automatic commodity classification applications, we state that collective effort is needed to build a comprehensive on-truck logo database, a logo-to-commodity database, as well as popularization of commodity logos printed on trucks.

Author Contributions: Conceptualization, P.H.; Data curation, A.W. and X.H.; Investigation, P.H.; Methodology, P.H. and A.W.; Project administration, A.R. and S.R.; Software, P.H.; Supervision, A.R. and S.R.; Visualization, A.W.; Writing—original draft, P.H. and A.W.; Writing—review & editing, X.H. All authors have read and agreed to the published version of the manuscript.

Funding: This work is supported by NSF CNS 1922782, by the Florida Dept. of Transportation (FDOT) and FDOT District 5. The opinions, findings and conclusions expressed in this publication are those of the author(s) and not necessarily those of the Florida Department of Transportation or the National Science Foundation.

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

References

1. Nguyen, L.X.; Chambers, M.; Goworowska, J.; Rick, C.; Sedor, J.; Berg, J.T.; Ford, C.; Liu, M.; Mallik, A.K.; Menegus, D.; et al. *Freight Facts and Figures*; United States Department of Transportation, Bureau of Transportation Statistics: Washington, DC, USA, 2018.
2. Wilson, J. A Concept for a National Freight Data Program. *TR NEWS* **2004**, *233*, 30.
3. Madar, G. Micro-Data Collection and Development of Trip Generation Models of Commercial Vehicles: An Application for Windsor, Ontario. Master's Thesis, University of Windsor, Windsor, ON, Canada, 2014.
4. Bronzini, M.; Firestine, T.; Fletcher, W.; Greene, D.; McGuckin, N.; Meyer, M.; Moore, W.H.; Rick, C.; Sedor, J. *Transportation Statistics Annual Report*; United States Department of Transportation, Bureau of Transportation Statistics: Washington, DC, USA, 2018.
5. Hernandez, S.V.; Tok, A.; Ritchie, S.G. Integration of Weigh-In-Motion and Inductive Signature Data for Truck Body Classification. *Transp. Res. Part C Emerg. Technol.* **2016**, *68*, 1–21. [[CrossRef](#)]
6. He, P.; Wu, A.; Huang, X.; Scott, J.; Rangarajan, A.; Ranka, S. Deep Learning based Geometric Features for Effective Truck Selection and Classification from Highway Videos. In Proceedings of the International IEEE Conference on Intelligent Transportation Systems (ITSC), Auckland, New Zealand, 27–30 October 2019.
7. He, P.; Wu, A.; Huang, X.; Rangarajan, A.; Ranka, S. Video-based Machine Learning System for Commodity Classification. In Proceedings of the International Conference on Vehicle Technology and Intelligent Transport Systems (VEHITS), Prague, Czech Republic, 2–4 May 2020.
8. He, P.; Wu, A.; Huang, X.; Scott, J.; Rangarajan, A.; Ranka, S. Truck and Trailer Classification With Deep Learning Based Geometric Features. *IEEE Trans. Intell. Transp. Syst. (T-ITS)* **2020**, *22*, 7782–7791. [[CrossRef](#)]
9. Psyllos, A.P.; Anagnostopoulos, C.N.E.; Kayafas, E. Vehicle Logo Recognition Using a SIFT-Based Enhanced Matching Scheme. *IEEE Trans. Intell. Transp. Syst.* **2010**, *11*, 322–328. [[CrossRef](#)]
10. Llorca, D.F.; Arroyo, R.; Sotelo, M.A. Vehicle Logo Recognition in Traffic Images using HOG Features and SVM. In Proceedings of the 16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013), The Hague, The Netherlands, 6–9 October 2013; pp. 2229–2234.
11. Romberg, S.; Pueyo, L.G.; Lienhart, R.; Van Zwol, R. Scalable Logo Recognition in Real-world Images. In Proceedings of the 1st ACM International Conference on Multimedia Retrieval, Trento, Italy, 18–20 April 2011; p. 25.
12. Hoi, S.C.; Wu, X.; Liu, H.; Wu, Y.; Wang, H.; Xue, H.; Wu, Q. LOGO-Net: Large-scale Deep Logo Detection and Brand Recognition with Deep Region-based Convolutional Networks. *arXiv* **2015**, arXiv:1511.02462.
13. Su, H.; Gong, S.; Zhu, X. WebLogo-2M: Scalable Logo Detection by Deep Learning from the Web. In Proceedings of the IEEE International Conference on Computer Vision, Venice, Italy, 22–29 October 2017; pp. 270–279.
14. Su, H.; Zhu, X.; Gong, S. Open Logo Detection Challenge. *arXiv* **2018**, arXiv:1807.01964.
15. Pan, C.; Yan, Z.; Xu, X.; Sun, M.; Shao, J.; Wu, D. Vehicle logo recognition based on deep learning architecture in video surveillance for intelligent traffic system. In Proceedings of the IET International Conference on Smart and Sustainable City 2013 (ICSSC 2013), Shanghai, China, 19–20 August 2013.
16. Su, H.; Zhu, X.; Gong, S. Deep Learning Logo Detection with Data Expansion by Synthesising Context. In Proceedings of the 2017 IEEE Winter Conference on Applications of Computer Vision (WACV), Santa Rosa, CA, USA, 24–31 March 2017; pp. 530–539.
17. Montserrat, D.M.; Lin, Q.; Allebach, J.; Delp, E. Scalable Logo Detection and Recognition with Minimal Labeling. In Proceedings of the 2018 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR), Miami, FL, USA, 10–12 April 2018; pp. 152–157.
18. Schroff, F.; Kalenichenko, D.; Philbin, J. FaceNet: A Unified Embedding for Face Recognition and Clustering. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Boston, MA, USA, 7–12 June 2015; pp. 815–823.

19. Yi, D.; Lei, Z.; Liao, S.; Li, S.Z. Deep Metric Learning for Person Re-Identification. In Proceedings of the 2014 22nd International Conference on Pattern Recognition, Stockholm, Sweden, 24–28 August 2014; pp. 34–39.
20. Chen, W.; Chen, X.; Zhang, J.; Huang, K. Beyond Triplet Loss: A Deep Quadruplet Network for Person Re-identification. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 403–412.
21. Chen, B.; Deng, W. Hybrid-Attention based Decoupled Metric Learning for Zero-Shot Image Retrieval. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 2750–2759.
22. Li, D.; Tian, Y. Survey and Experimental Study on Metric Learning Methods. *Neural Netw.* **2018**, *105*, 447–462. [[CrossRef](#)] [[PubMed](#)]
23. Lowe, D.G. Distinctive Image Features from Scale-Invariant Keypoints. *Int. J. Comput. Vis.* **2004**, *60*, 91–110. [[CrossRef](#)]
24. Sivic, J.; Zisserman, A. Video Google: A Text Retrieval Approach to Object Matching in Videos. In Proceedings of the IEEE International Conference on Computer Vision, Nice, France, 13–16 October 2003; p. 1470.
25. Zhou, X.; Yao, C.; Wen, H.; Wang, Y.; Zhou, S.; He, W.; Liang, J. EAST: An Efficient and Accurate Scene Text Detector. In Proceedings of the IEEE conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 5551–5560.
26. Tian, Z.; Huang, W.; He, T.; He, P.; Qiao, Y. Detecting Text in Natural Image with Connectionist Text Proposal Network. In Proceedings of the European Conference on Computer Vision, Amsterdam, The Netherlands, 8–16 October 2016; Springer: Berlin/Heidelberg, Germany, 2016; pp. 56–72.
27. He, P.; Huang, W.; Qiao, Y.; Loy, C.C.; Tang, X. Reading Scene Text in Deep Convolutional Sequences. In Proceedings of the AAAI Conference on Artificial Intelligence, Phoenix, AZ, USA, 12–17 February 2016.
28. Shi, B.; Bai, X.; Yao, C. An End-to-End Trainable Neural Network for Image-based Sequence Recognition and Its Application to Scene Text Recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* **2016**, *39*, 2298–2304. [[CrossRef](#)] [[PubMed](#)]
29. Fehérvári, I.; Appalaraju, S. Scalable Logo Recognition using Proxies. In Proceedings of the 2019 IEEE Winter Conference on Applications of Computer Vision (WACV), Waikoloa Village, HI, USA, 7–11 January 2019; pp. 715–725.
30. Ren, S.; He, K.; Girshick, R.; Sun, J. Faster r-cnn: Towards real-time object detection with region proposal networks. In Proceedings of the Advances in Neural Information Processing Systems, Montreal, QC, USA, 7–12 December 2015; pp. 91–99.
31. Simonyan, K.; Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv* **2014**, arXiv:1409.1556.
32. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 27–30 June 2016; pp. 770–778.
33. Huang, G.; Liu, Z.; Van Der Maaten, L.; Weinberger, K.Q. Densely connected convolutional networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 4700–4708.
34. He, K.; Gkioxari, G.; Dollár, P.; Girshick, R. Mask r-cnn. In Proceedings of the IEEE International Conference on Computer Vision, Venice, Italy, 22–29 October 2017; pp. 2961–2969.
35. Redmon, J.; Divvala, S.; Girshick, R.; Farhadi, A. You only look once: Unified, real-time object detection. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 27–30 June 2016; pp. 779–788.
36. Lew, M.S.; Sebe, N.; Djeraba, C.; Jain, R. Content-based multimedia information retrieval: State of the art and challenges. *ACM Trans. Multimed. Comput. Commun. Appl. TOMM* **2006**, *2*, 1–19. [[CrossRef](#)]
37. Wang, X.; Han, X.; Huang, W.; Dong, D.; Scott, M.R. Multi-Similarity Loss with General Pair Weighting for Deep Metric Learning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 5022–5030.
38. Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; Rabinovich, A. Going deeper with convolutions. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Boston, MA, USA, 7–12 June 2015; pp. 1–9.
39. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. *arXiv* **2014**, arXiv:1412.6980.
40. Sivic, J.; Zisserman, A. Efficient visual search of videos cast as text retrieval. *IEEE Trans. Pattern Anal. Mach. Intell.* **2008**, *31*, 591–606. [[CrossRef](#)] [[PubMed](#)]
41. Li, F.F.; Perona, P. A bayesian hierarchical model for learning natural scene categories. In Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05), San Diego, CA, USA, 20–26 June 2005; Volume 2, pp. 524–531.
42. Lowe, D.G. Object recognition from local scale-invariant features. In Proceedings of the Seventh IEEE International Conference on Computer Vision, Kerkyra, Greece, 20–27 September 1999; Volume 2, pp. 1150–1157.
43. Jolliffe, I.T. Principal components in regression analysis. In *Principal Component Analysis*; Springer: Berlin/Heidelberg, Germany, 1986; pp. 129–155.
44. Everingham, M.; Van Gool, L.; Williams, C.K.; Winn, J.; Zisserman, A. The PASCAL Visual Object Classes (VOC) Challenge. *Int. J. Comput. Vis.* **2010**, *88*, 303–338. [[CrossRef](#)]