System for monitoring road slippery based on CCTV cameras and convolutional neural networks

Dariusz Grabowski¹ · Andrzej Czyżewski¹

Received: 10 June 2020 / Revised: 22 August 2020 / Accepted: 24 August 2020 /

Published online: 9 September 2020

© The Author(s) 2020

Abstract

The slipperiness of the surface is essential for road safety. The growing number of CCTV cameras opens the possibility of using them to automatically detect the slippery surface and inform road users about it. This paper presents a system of developed intelligent road signs, including a detector based on convolutional neural networks (CNNs) and the transfer-learning method employed to the processing of images acquired with video cameras. Based on photos taken in different light conditions by CCTV cameras located at the roadsides in Poland, four network topologies have been trained and tested: Resnet50 v2, Resnet152 v2, Vgg19, and Densenet201. The last-mentioned network has proved to give the best result with 98.34% accuracy of classification dry, wet, and snowy roads.

Keywords Machine learning · Convolutional neural networks · Transfer learning · Road safety

1 Introduction

Proper assessment of the surface condition in terms of its slipperiness is crucial for road safety. For many drivers, this is a difficult task, which affects the problem of adjusting the speed of the vehicle to the conditions. According to a Polish Police report (Budzyński and Tubis 2019) from 2018, 24.1% of accidents caused by drivers took place for this very reason. Another research (Rama 2001) has shown that variable message signs based on weather conditions detection (especially in winter) have a positive effect on speed reduction among drivers. On this basis, it is reasonable to create a system that would accurately recognize the condition of the surface based on available data and inform road users about it. Also, it can be concluded that increasing the availability of road state detection technologies by reducing costs should positively affect safety. As described in the literature (Dey et al. 2015),

✉ Andrzej Czyżewski
ac@pg.gda.pl

Dariusz Grabowski
dariusz.grabowski1@gmail.com

¹ Multimedia Systems Department, Gdańsk University of Technology ETI Faculty,
Narutowicza 11/12, Gdańsk, Poland

easy access to such a system can also be a support for emergency management departments and facilitate better decision making. Information about the condition of the road allows emergency personnel to find the fastest route to the incident location. As it is demonstrated in this paper, this aim is possible to achieve by the use of existing infrastructure like roadside CCTV cameras and machine learning. Before the achieved results are discussed an overview of the system including developed intelligent road signs is presented. The data for neural network training were acquired from 27 CCTV cameras installed at the roadsides by the Polish General Directorate for National Roads and Highways (GDDKiA). The place of application of a trained neural network is an intelligent road sign developed as part of a conducted research and development project. The architecture of the system developed as a result of this project is briefly reviewed in the next chapter.

2 Information system architecture

At Gdańsk University of Technology, Faculty of Electronics, Telecommunications, and Informatics, in cooperation with AGH University of Science and Technology in Kraków and two companies a research project has been carried out under the title: “INZNAK – Intelligent road signs for adaptive vehicle traffic control, communicating in V2X technology” (Czyzewski et al. 2019). The main objective of the project is to prevent congestion and resulting vehicle collisions on motorways and expressways. In connection with the main objective defined above a scenario related to this objective is presented in Fig. 1. It is possible to use the developed intelligent signs also in a number of other traffic situations. An important feature of the intelligent road sign under development is to ensure the functionality of communication with vehicles equipped with the digital car-to-infrastructure communication interface. The designed system employing intelligent road signs communicates the speed calculated in relation to information received from a sequence of similar variable message signs along a section of motorway, connected to each other via a wireless network, optionally with the possibility of using remote management. A special feature of road signs is that they can operate autonomously, as the speed limit communicates through

Fig. 1 Scenario illustrating the use of the system of intelligent road signs

the signs is the result of their traffic measurements. The recommended speed is communicated by displaying it on a variable message sign and is also transmitted wireless to vehicles equipped with V2X interface (the interface to electronic communication between vehicle and infrastructure). A diagram illustrating how intelligent signs communicate in order for the system to make a decision is shown in Fig. 2. It is visible that multiple communication is in use, including:

- data exchange INZNAK with Management Centre
- data exchange road sign with neighbour road signs
- data exchange with vehicles having a communication unit operating in V2X standard

To meet these challenges, V2X (Czyzewski et al. 2019) technology was implemented to communicate signs with road infrastructure and passing vehicles. Therefore, it is possible to establish a connection with the vehicle's computer, in which information about e.g., recommended speed is transmitted. The software implemented during the stage has been tested in road traffic with the use of cars that transmitted information to a receiver located near the road (see Fig. 3). The system developed includes the following components:

- Radar sensors: microwave, lidar, and acoustic
- Bluetooth device scanner
- weather station measuring temperature, pressure and humidity
- brightness sensor
- temperature sensor
- Wi-Fi router, LTE and LoRaWAN modem
- microcomputer with built-in graphics card for image processing (Jetson)
- microcomputer to manage the active road sign system, including the transmission of information to the server.

In the course of the project, a number of solutions comprising a prototype system of autonomous road signs were developed and experimentally tested. The idea of autonomous, communicating road signs was realized as a system demonstrator. Some of the solutions

Fig. 2 Road information system employing intelligent road signs

Fig. 3 Photo of the measuring station of the intelligent road sign with CCTV camera sensor attached to the bottom of its housing

have the character of original new technologies, such as acoustic radar plus improved Doppler sensor (Czyżewski et al. 2019), and the artificial intelligence application to detect the degree of slippery surfaces, which is obtained from a traditional video camera. The last-mentioned subject is discussed in this paper.

3 Methods of slippery detection

There are many existing techniques applicable to detect road conditions (see Fig. 4). They differ in the type and cost of the sensors used, the accuracy of the classification, and some limitations. Such limitations include the time of day (e.g., limited operation after dark) or particularly difficult weather conditions with it can operate (e.g., limited operation in rain or fog). Another issue is the place where the sensor is installed. Some of them need to be embedded into the surface; others placed close to it or placed on specified height. Another kind of limitation is the distance from which device is possible to collect measurements and its angular resolution (Bystrov et al. 2018). Moreover, the results presented in this paper may serve as a reference for comparison of the research results presented in this paper. Solutions can also be divided into fixed platforms installed next to the roadway and those mounted on vehicles. One of the unconventional sensor types, based on audio analyses

Fig. 4 Photos of experimental setups for determining road surface condition

(Kongrattanasert et al. 2010), was created to classify wet/dry pavement. The basis for the study was 785,000 sound samples taken at different vehicle speeds, with various environmental sounds acquired from diverse types of the road surface. The data was used in the process of learning a recurrent neural network with long short-term memory (LSTM) modules. Field tests, which were carried out on diverse road types and at different speeds, showed 93.2% sensitivity of the system (including the speed of 0 km/h, in that case, the system analyzed the sounds from passing cars). In another study (Bystrov et al. 2017), a hybrid approach was used: simultaneous use of 24 GHz radar and 40 kHz sonar. Reflected signal features were used to train the multi-layer neural network. Field tests of this solution carried out in various weather conditions and on different roads (asphalt, gravel, grass, sand), proved 95% of the average efficiency in the classification of wet/snowy/icy surfaces. The problem of pavement condition classification was also implemented using infrared radiation. In the (Jonsson 2011) study, it was shown that using near infra-red (NIR) in the 1000 nm to 2000 nm range; it is possible to determine if the road is dry/wet/snowy/icy. These conclusions were drawn based on tests carried out in laboratory conditions. For this purpose, 3 IR detectors were used with peak sensitivity at 960 nm, 1550 nm, and 1950 nm. It has been proved that individual types of pavement can be separated from each other based on sensor readings. The results of the continuation of research (Jonsson et al. 2015) were published in 2015. The goal set was to create a system that will be able to measure road conditions continuously without affecting its users. Therefore, a CCTV camera mounted next to the roadway, and an infrared camera equipped with a number of filters was used. Field studies confirmed the results from the laboratory. Several signal classifiers were checked, of which the best results were obtained for the SVM (Support Vector Machine) and KNN (k-nearest neighbours) classifier. The average efficacy was 95.75% and 99.5%, respectively. The authors point out that creating a commercial system based on a NIR camera would be very expensive due to the price of the equipment. The condition of the surface can also be assessed on the basis of its temperature and solar radiation (Lu Junhui and Wang Jian-qiang 2010). It was noticed that there is a relationship between dry/wet/icy road temperature and insolation estimated based on geographical location, season, air temperature, humidity, and time of day. Based on the factors given, a neural network model was trained, which

in field studies showed 90% efficiency in recognizing the aforementioned road conditions. Researchers indicate that this approach has low night-time efficiency. Another way to solve the problem of surface detection is to use a stereo camera (Kim et al. 2013). High efficiency was obtained using a camera equipped with filters that allow for obtaining 2 images at a given time: horizontally and vertically polarized. In addition, the test vehicle was equipped with a temperature and humidity sensor. Based on the data from these three sensors and with the help of the K-means classifier, the authors proposed an algorithm that allows detection of dry, wet, snowy, and icy surfaces. In the field tests, 97%, 95%, 87%, and 33% effectiveness respectively were obtained for each class. An approach based on image processing was also proposed in the literature (Nolte et al. 2018). The study grouped many sets of photos, mainly from studies on autonomous vehicles. The dataset was augmented by images mirroring, rotating, and scaling. Two CNN network architectures, InceptionNetV3 and ResNet50, were tested. In detecting the type of pavement (dry/wet asphalt, dirt, grass, cobblestone, snow), 90% and 92% average efficiency were obtained for given architectures, respectively. This paper presents an approach to detect road condition in categories which are crucial for pavement slipperiness assessment: dry, wet, and snowy. Subsequent chapters describe the process of collecting and processing data, conduct the learning process, and obtained the results of experiments.

4 Collecting dataset

A big challenge in methods based on machine learning algorithms is to obtain an extensive and adequately described data set. It is necessary to conduct an adequate learning process and to obtain satisfactory results. For the needs of this study, it was decided to use the network of measurement stations owned by the General Directorate of National Roads and Motorways, being the main road operator in Poland. Those devices, mounted along the main roads in the country, are equipped with a CCTV camera and a set of sensors for measuring weather conditions. The location of used cameras is shown on the map presented in Fig. 5. The features presented by the system are the temperature measured on various heights,

Fig. 5 The map of Poland with used cameras location

humidity, dew, road condition data (icy/snowy/dry/wet/moist), wind speed, and direction. To aggregate data on current road conditions, there was created an application for downloading and archiving road photos and corresponding weather conditions. The resulting set consists of over 3 300 000 records collected in the period from November 2018 to March 2020 all over Polish roads. Samples divided into categories are shown in Fig. 6. The collected dataset was made publicly available at Harvard Dataverse portal (Grabowski 2020).

Unfortunately, weather stations differ from each other by the set of available sensors and the complexity. Not all sensors distinguish between precipitation kinds like rain or snow (they detect the existence of precipitation, only). There is also a limited number of stations that can detect snow on the road. The connection between those categories is presented in Fig. 7.

Collected images were evaluated for quality using the BRISQUE method (Mittal et al. 2011). Those of poor quality have been rejected from further processing. As a result, due to the insufficient number of samples, the ice and moist categories were discarded. In the case of the moist and saline category, due to their poor discrimination between dry and wet

Fig. 6 Amount of data samples collected by category

Fig. 7 Amount of data samples collected by category

surfaces, it was decided not to include them in the classification algorithm. As a result, the dry, wet, and snow categories were chosen. A quite warm winter caused the amount of snowy road data compared to other categories, is rather low. For this reason, it was necessary to find photos of the snowy road, based on other weather parameters. It was achieved by finding graphics for which the precipitation was described as continuous/intense, and the temperature was below -2°C . Finally, for each category (dry, wet, snow), 711 images were selected from 27 roadside measurement stations. An equal number of photos for each category has been selected for each station. The dataset includes different lighting conditions like daylight, twilight, and street lightings switched on during the night.

5 Proposed approach

It was decided to process the collected dataset to use convolutional neural networks (CNN). It was chosen to conduct the learning process based on the transfer learning method. The method involved using a model trained on a large data set (the type of data does not have to be associated with road data). On top of that, there is a need to fine-tune only the last layers of the neural network. Research (Pan et al. 2018) (Shin et al. 2016) (Yosinski et al. 2014) shows that in many fields, the application of the described approach allows for gaining better results (higher precision) than when using only photos related to the model being built or building a model from scratch. The first step of the proposed solution is the unification of grouped into classes labelled photos. Each of them was cut to a square with dimensions of 224×224 . Then the photos of all categories were divided into 3 groups: 70% training set, 15% validation set for checking performance during training, and 15% test set for getting the accuracy of the resulted model. Each set is converted into a binary file in the Record IO format to efficiently operate on data. Before starting the learning process, individual photos of the training set are multiplied by random transformation such as horizontal flip, hue, and brightness adjustment. As described in the literature (Huang et al. 2017), this approach allows for increasing the precision of the created system and immunizing random changes in the data from cameras. The next step is to acquire a pre-trained CNN model. In this work, models Resnet50 v2 (He et al. 2016), Resnet152 v2 (He et al. 2016), Densenet201 (Huang et al. 2016), and Vgg19 (Simonyan and Zisserman 2014) were used, trained based on the Imagenet 1000 dataset. On top of that learning process is performed using train and validation sets to gain new skills as recognizing a dry, wet, and snowy road. The described process is presented in Fig. 8.

The results obtained in this way exceed the current state of the art results. For comparison, articles were selected in which predictions were made for dry, wet, and snowy roads.

Fig. 8 Neural network model with the transfer learning approach. An image after cropping is multiplied with random transformations and then used for the training

For comparison with the first of them (Nolte et al. 2018), we rejected the other recognized classes (cobblestone, grass, dirt). With this assumption, the authors obtained 93.33% accuracy.

6 Experiments

In order to find the best matching neural network topology, several tests were carried out to accomplish the assumed task. For this purpose, a machine with a GeForce RTX 2080Ti graphics card with 11 GB of RAM and with 18.8 TFLOPS computing power was used. The tests of each network were carried out with the same set of data. The learning rate parameter was set to 0.001 to prevent a quick overwriting of weights from the base model. The batch size for each topology was set to 32. Progress in the learning process for particular networks is shown in Fig. 9, and the learning results are shown in Table 1. Differences

Fig. 9 The learning process of selected neural networks for 30 iterations

Table 1 Performance comparison of selected models

Network	Accuracy	Precision	Recall	F1 Score
Resnet152_v2	96.05%	94.08%	94.18%	94.09%
Resnet50_v2	96.68%	95.02%	95.02%	95.02%
Densenet201	96.88%	95.33%	95.45%	95.35%
Vgg19	96.68%	95.02%	95.05%	95.03%

between the individual results are quite small. The worst is the Resnet152 network, with accuracy at 96.05%. The results of the Vgg19 and Resnet50 networks for the training set are almost the same. Densenet201 is the best in this ranking, with an accuracy of 96.68%. It was decided to go deeper into the Densenet201 network based on that results and try to find even better parameters to achieve the best possible score. As a result of several experiments, the learning rate parameter was increased to 0.01, and the number of iterations was reduced to 15. This resulted in a score of 1.46 percentage points better compared to the best solution described above, so the final result is 98,34%. The network, based on the given parameters, has obtained a better ability to generalize features from the inputs images. The total time of training the network was 5 minutes 38 seconds. Based on the matrix of results (Fig. 10) for particular classes of images, it can be concluded that the biggest problem is distinguishing between snow and wet pavements. Fig. 11 shows misclassified photographs from the test set. One of the reasons for incorrect classification is passing cars. The highly reflected headlight of cars is typical for wet surfaces. In poor lighting, especially at night, this effect may be similar even for dry surfaces. Passing trucks posed another problem. Their size makes them fill most of the frame, while the white tarpaulin effectively imitates the snowy surface. Another issue is the noise of the picture. This may be due to the camera lens being splashed with water or because of the appearance of fog/dust. Most of these problems can be solved by examining a sequence of images from the same location and making decisions based on summary observation. Some images were found that were difficult to classify automatically, for which the neural network gave particularly erroneous results at the output. They are presented in Fig. 12. On the other hand, many photos taken in severe conditions, such as poor lighting or intense rainfalls/snowfalls have been correctly classified. The corresponding examples are presented in Fig. 13. The results obtained in

Fig. 10 Prediction results for Densenet201 network with $lr=0.01$ and after 15 iterations

	dry	snow	wet
dry	0.991	0.000	0.019
snow	0.009	1.000	0.000
wet	0.019	0.028	0.962

Predicted labels

True labels

Fig. 11 Misclassified images in different categories

this way exceed the current state of the art results. For comparison, articles were selected in which predictions were made for dry, wet, and snowy roads. For comparison with the first of them (Nolte et al. 2018), we rejected the other recognized classes (cobblestone, grass, dirt). With this assumption, the authors obtained 93.33% accuracy.

In another publication (Amthor et al. 2015), researchers for the classes mentioned above obtained 96.79% accuracy. The results of the research presented here exceed hitherto best-known results by 1.55 percentage points. Another essential factor of the surface slippery detector evaluation is the processing time for a particular frame. Performing calculations on a powerful cloud computing machine is a relatively easy task. The situation gets complicated when the detectors are mounted next to the road, in an environment of limited internet coverage. In this case, sending data to the computing server may be difficult and cause the device to malfunction. It is, therefore, vital for detectors to be autonomous. Assessment of the surface condition should be possible on a platform equipped with a CCTV camera and a relatively inexpensive GPU module such as Nvidia Jetson. Therefore, it is not only accuracy important but also the average processing time is essential. When choosing network topologies, the right balance should be found between these two factors. The graphic Fig. 14 shows the average image processing time for all described topologies. These results were generated on another platform: GeForce GTX 1080 Ti with 27.4 TFLOPS computing power. The shortest processing time was obtained for the Vgg19 network. Whereas the

Fig. 12 Difficult images in terms of classifying them to adequate categories

Fig. 13 Correctly classified images. From left to right: dry, snow, wet, snow

Fig. 14 Average processing time for all considered topologies

relatively small differences in performance between topologies (Table 1), the Vgg19 network seems to be the right candidate for use on embedded platforms.

7 Conclusions

The model created employing the transfer learning method allowed obtaining satisfying results at a low computing cost. It was indicated by experimental results, which brought 98.34% accuracy in recognition of dry, wet, and snowy roads. It has also been proven that detection can also be carried out in conditions of limited lighting (street lights case). The slipperiness detector solution proposed in this paper, utilizing the existing road measuring stations and other roadside cameras with public access, allows for covering a large area without incurring high costs. The efficiency and universality of the developed model can be increased further by the usage of data from cameras with various characteristics, i.e., with a different configuration, placed in diverse lighting conditions, locations, angles, and covering road sections of various categories (motorways/expressways/national roads).

Acknowledgments The project is financed by the Polish National Centre for Research and Development (NCBR) from the European Regional Development Fund No. POIR.04.01.04-00-0089/16 entitled: “INZNAK: Intelligent Road Signs with V2X Interface for Adaptive Traffic Controlling”.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Amthor, M., Hartmann, B., Denzler, J. (2015). Road Condition Estimation Based on Spatio-Temporal Reflection Models. In *Lecture Notes in Computer Science including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics*, 9358, Springer International Publishing (pp. 3–15).
- Budzyński, M., & Tubis, A. (2019). Assessing the Effects of the Road Surface and Weather Conditions on Road Safety. *Journal of Konbin*, 49(3), 323–349.

- Bystrov, A., Hoare, E., Tran, T.Y., Clarke, N., Gashinova, M., Cherniakov, M. (2017). Automotive system for remote surface classification. *Sensors Switzerland*, 17(4), 745.
- Bystrov, A., Hoare, E., Tran, T.-Y., Clarke, N., Gashinova, M., Cherniakov, M. (2018). Sensors for Automotive Remote Road Surface Classification. In *2018 IEEE International Conference on Vehicular Electronics and Safety ICVES IEEE*, 1–6.
- Czyzewski, A., Kotus, J., Szwoch, G. (2019). Estimating Traffic Intensity Employing Passive Acoustic Radar and Enhanced Microwave Doppler Radar Sensor. *Remote Sensing*, 12(1), 110.
- Czyzewski, A., Sroczynski, A., Smialkowski, T., Hoffmann, P. (2019). Development of Intelligent Road Signs with V2X Interface for Adaptive Traffic Controlling. In *2019 6th International Conference on Models and Technologies for Intelligent Transportation Systems MT-ITS, IEEE*, 1–7.
- Dey, K.C., Mishra, A., Chowdhury, M. (2015). Potential of intelligent transportation systems in mitigating adverse weather impacts on road mobility: A review. *IEEE Transactions on Intelligent Transportation Systems*, 16(3), 1107–1119.
- Grabowski, D. (2020). Road CCTV images with associated weather data Harvard Dataverse, <https://doi.org/10.7910/DVN/SV9N9F>.
- He, K., Zhang, X., Ren, S., Sun, J. (2016). Identity Mappings in Deep Residual Networks. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 9908 LNCS, 630–645. arXiv:1603.05027.
- Huang, G., Liu, Z., van der Maaten, L., Weinberger, K.Q. (2016). Densely Connected Convolutional Networks. *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, 2017-Janua*, 2261–2269. arXiv:1608.06993.
- Huang, Z., Pan, Z., Lei, B. (2017). Transfer learning with deep convolutional neural network for SAR target classification with limited labeled data. *Remote Sensing*, 9(9), 907.
- Jonsson, P. (2011). Remote sensor for winter road surface status detection. In *Proceedings of IEEE Sensors, IEEE*, 1285–1288.
- Jonsson, P., Casselgren, J., Thornberg, B. (2015). Road surface status classification using spectral analysis of NIR camera images. *IEEE Sensors Journal*, 15(3), 1641–1656.
- Kim, Y., Baik, N., Kim, J. (2013). A Study on Development of Mobile Road Surface Condition Detection System Utilizing Probe Car. *Journal of Emerging Trends in Computing and Information Sciences*, 4(10), 742–750.
- Kongrattanasert, W., Nomura, H., Kamakura, T., Ueda, K. (2010). Automatic Detection of Road Surface States from Tire Noise Using Neural Network Analysis. *Proceedings of 20th International Congress on Acoustics, ICA 2010*, 130(7), 920–925.
- Lu Junhui, & Wang Jianqiang (2010). Road surface condition detection based on road surface temperature and solar radiation. In *2010 International Conference on Computer, Mechatronics, Control and Electronic Engineering: IEEE*.
- Mittal, A., Moorthy, A. K., Bovik, A. C. (2011). Blind/Referenceless Image Spatial Quality Evaluator. In *2011 Conference Record of the Forty Fifth Asilomar Conference on Signals, Systems and Computers ASILOMAR, IEEE*, 723–727.
- Nolte, M., Kister, N., Maurer, M. (2018). Assessment of deep convolutional neural networks for road surface classification. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. <https://doi.org/10.1109/itsc.2018.8569396>, <http://dx.doi.org/10.1109/ITSC.2018.8569396>.
- Pan, G., Fu, L., Yu, R., Muresan Ph Student, M. D. (2018). Winter Road Surface Condition Recognition Using a Pre-trained Deep Convolutional Neural Network. In *Transportation Research Board 97th Annual Meeting*.
- Rama, P. (2001). Effects of weather-controlled variable message signing on driver behaviour, VTT Publications, 1, 447.
- Shin, H., Roth, H., Gao, M., Lu, L., Xu, Z., Nogues, I., Yao, J., Mollura, D., Summers, R. (2016). Deep Convolutional Neural Networks for Computer-Aided Detection: CNN Architectures, Dataset Characteristics and Transfer Learning and Daniel Mollura are with Center for Infectious Disease Imaging HHS Public Access. *IEEE Transactions on Medical Imaging*, 35(5), 1285–1298.
- Simonyan, K., & Zisserman, A. (2014). Very Deep Convolutional Networks for Large-Scale Image Recognition ,1–14, arXiv:1409.1556.
- Yosinski, J., Clune, J., Bengio, Y., Lipson, H. (2014). How transferable are features in deep neural networks? In *27th International Conference on Neural Information Processing Systems, 2, MIT Press, Montreal*, 3320–3328.