# Fish species classification in underwater video monitoring using Convolutional Neural Networks

**Frederik Kratzert**
Institute of Water Management, Hydrology and Hydraulic Engineering
University of Natural Resources and Life Sciences
Vienna, Austria
f.kratzert@gmail.com


**Helmut Mader**
Institute of Water Management, Hydrology and Hydraulic Engineering
University of Natural Resources and Life Sciences
Vienna, Austria

## Abstract

This report presents a case study for automatic fish species classification in underwater video monitoring of fish passes. Although the presented approach is based on the FishCam monitoring system, it can be used with any video-based monitoring system. The presented classification scheme in this study, is based on Convolutional Neural Networks that do not require the calculation of any hand-engineered image features. Instead, these networks use the raw video image as input. Additionally, this study investigates, if the classification accuracy can be increased by adding additional meta-information (date of migration and fish length) to the network. The approach is tested on a subset of 10 fish species (8099 individuals) occurring in Austrian river. On an independent test set, the presented approach achieves a classification accuracy of 93 %.

## 1   Introduction

This is a short technical report on about a recent study to classify fish from images of underwater video monitoring regarding their fish species. The study is part of the development of the FishCam monitoring system (Mader and Kratzert, 2016) that is designed for (semi-) automatic monitoring of fish migration in fish passes.

There exist already a variety of approaches for problem of classifying fish (Alsmadi et al., 2010; Badawi and Alsmadi, 2013; Lee et al., 2004; Matai et al., 2012; Rova et al., 2007; Spampinato et al., 2010) All of them have in common that the fish are classified according to engineered features, which are extracted from either the color/grayscale-image or the binary shape. This requires an exact delimitation of the fish from the background and the identification of certain key-points in an automated fashion. In real-world underwater video monitoring of e.g. fish passes in rivers with strong variations in the turbidity of the water this is an almost unfeasible task (Kratzert, 2016).

In recent years, Convolutional Neural Networks (CNNs) have gained a lot of attention and popularity regarding all kind of computer vision tasks (Farabet et al., 2013; Krizhevsky et al., 2012; Tompson et al., 2014). The advantage of these networks is that they take the entire image as input and learn to extract important features for the classification themselves during a training period. The FishNet software, which is part of the FishCam monitoring system, already uses a CNN in a first classification step, in which detected objects are divided into either being a fish or not being a fish.

Figure 1: Left: shows the construction of the detection tunnel in an underwater scenario. Right: shows an example of the recorded video imagery in clear water conditions.

This classification step achieves an overall accuracy of approx. 96 % correctly classified objects (Kratzert, 2016; Kratzert and Mader, 2016).

The main objective of this study is to investigate the use of CNNs for the second classification step, which classifies the detected fish into different fish species. Besides classifying fish only by images, also the use of additional meta-information or features (date of migration, length of the fish) is investigated. In addition, two different combination strategies are investigated, that can be used to combine the classification results of multiple images for one fish. Besides that, this report presents further information on has the following structure: first the FishCam detection tunnel is presented, which is used for recording the monitoring data, . Then the network architectures in more detail that is used in this study are presented in more detail, as well as the underlying data of this study. Then Finally, some preliminary results are presented coupled with a discussion and followed by an outlook for future studies.

## 2  Methods

### 2.1  Detection tunnel

To facilitate the automation of the video analysis, a detection tunnel was constructed for the FishCam monitoring system (Mader and Kratzert, 2016) that provides a static surrounding across different river (see Fig. 1). Through artificial illumination by LEDs with a light temperature of 6000 K, the lightning conditions are kept statically throughout day and night. Further, a white floor and background panel is installed to increase the contrasts in the videos. For the correct determination of the fish length, a mirror is installed in the upper part of the tunnel. For detailed description of all parts that are installed in the tunnel, see Mader and Kratzert (2016).

### 2.2  Adapted VGG16 Network

The underlying convolutional neural network architecture used in this study is the VGG16 (Simonyan and Zisserman, 2014). The convolutional part of this network is kept unchanged and the weights are used from a pre-trained network on the ImageNet dataset (Russakovsky et al., 2015). The only part that is adapted in this study is the full-connected top (see Fig. 2). The first adaption is the possible intergration of additional features (date of migration, fish length) at the beginning of the fully-connected top of the network. The hypothesis is that these additional features can help to increase the classification accuracy in general, since migration time and length can be a strong predictor/separator for fish classes. For example the time of the spawning migration is a type-specific property and naturally different fish have different length statistics (average length, maximum observed/possible length).

Before the date of migration can be used as an input, it must be transformed into a different format. In this case, the decision was to use only the week of the year as feature and encode this information into one-hot-encoding. One-hot encoding transforms a categorical feature (in this case the week of
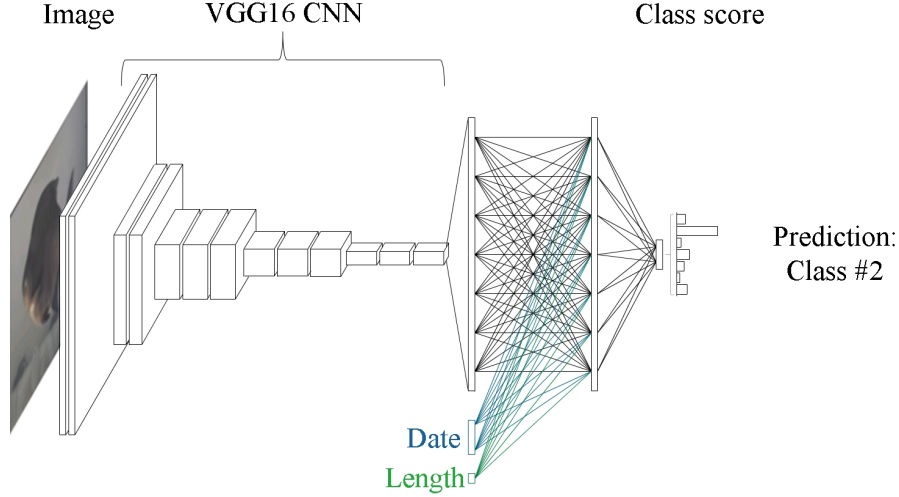
Figure 2: Adapted VGG16 network architecture used in this study. VGG16 CNN denotes the convolutional body of the original VGG16 network that is used unchanged. The class probabilities (depicted on the right) are calculated in the last layer of the network and from these, the precited class is derived.

the year) into a vector of *n* elements (where *n* is the maximum number of categories, which is 52 in this case), where each element is set to 0, except the one of the specific category (the migration week of the specific sample), sample, which is set to 1. Thus, the date input feature is transformed into another 52 dimensional input vector.

The second adaption made to the fully-connected top is that a batch-normalization layer (Ioffe and Szegedy, 2015) is added after each fully-connected layer, because a preliminary experiment has shown that this can help to generally increase the classification accuracy. A batch-normalization layer is also added to the length input feature, to learn an optimal input feature scaling during the training period for this feature.

The last adaption made is that the number of output neurons in the last layer is adjusted to the number of classes used in this experiment. Because the data set contains 10 different fish species, the number of neurons in the last layer is set to 10. The activations of the output neurons in the last layer is squashed through the softmax function, which is given by the following equation:

$$f(z_i) = \frac{e^{z_i}}{\sum_{n=0}^{N} e^{z_n}}, \tag{1}$$

where $z$ is the vector of activatoins of the output (last) layer of the network, $z_i$ is the *i-th* element of the vector $z$ and $N$ is the number of neurons (number of classes) in the output layer. The softmax function from Eq. 1 transforms a vector containing arbitrary real values into a vector of real values in the range of (0, 1) that add up to 1. Therefore, the resulting values of the softmax transformation can be interpreted as class probabilities.

To analyse the effect of the additional features used in this study, four different network configurations are tested:

1. **VGG16:** Classification based only on the images.

2. **VGG16 + length:** Classification based on the images, as well as the length of the fish.

3. **VGG16 + date:** Classification based on the images, as well as the date of the migration.

4. **VGG16 + both:** Classification based on the images, as well as the length and the date.

All of these network configurations are trained and evaluated on the exact same data, to make the results comparable.

3

| Species | Total number of fish | Total number of images |
|---|---|---|
| Barbel (*Barbus barbus*) | 510 | 3827 |
| Brown trout (*Salmo trutta fario*) | 475 | 5901 |
| Burbot (*Lota lota*) | 279 | 2296 |
| Common bream (*Abramis brama*) | 1122 | 9607 |
| Common chub (*Squalius cephalus*) | 1138 | 10881 |
| Common nase (*Chondrostoma nasus*) | 639 | 7580 |
| European bullhead (*Cottus gobio*) | 187 | 1333 |
| European perch (*Perca fluviatillis*) | 1149 | 10982 |
| Grayling (*Thymallus thymallus*) | 1192 | 18044 |
| Rainbow trout (*Oncorhynchus mykiss*) | 1408 | 12178 |

Table 1: Overview of the data used in this study.

## 2.3 Training data

The data used in this study arrives from over three years of monitoring fish migration with the FishCam system across a variety of rivers in Austria. In the latest experiments, a subset of 10 of all fish species occurring in Austrian rivers are used for testing this classification approach (see Table 2.3). This subset is chosen under the following consideration:

- The subset should contain fish species that look somewhat similar to each other (Rainbow trout (*Oncorhynchus mykiss*) and Brown trout (*Salmo trutta fario*) as well as Common nase (*Abramis brama*) and Common chub (*Squalius cephalus*)), to see if classification errors occur more frequently between this species.

- The subset should also contain fish species that look clearly different (European bullhead (*Cottus gobio*), European perch (*Perca fluviatillis*) and Burbot (*Lota lota*)) to the rest, to test how good these species can be separated in the classification procedure.

- Although more data is available, only samples of fish pass monitorings are chosen, where each of the 10 species occur to avoid the possibility of classification bias due to differences in the specific video imagery.

- The subset should contain more common species, as well as some rarely seen species, as the situation is naturally in the river systems. This leads to an inhomogeneous amount of available data per class and the focus on training a network with uneven sample sizes.

- Additionally, the dataset should represent the different image qualities caused by turbidity occurring in real-world monitorings in Austrian rivers, so that this approach can be tested for an all-year round use. The image qualities can range from crystal-clear to highly turbid water (see Fig 3).

Human experts classified all samples that are used in this experiment in a first step, so that truth data for training and testing is available. These experts also measured the length of each fish from the recorded videos. Then for each fish in each video image crops were extracted automatically using the FishNet software (Kratzert, 2016). The image crops were rotated, so that major-axis of each fish is horizontally and it was taken care, that each crop shows exactly a single fish in full extend. The amount of images per fish varies heavily, due to different dwell times in front of the camera. The total numbers of fish and image samples per species is shown in Table 2.3.

For training and evaluating, the dataset is randomly divided into a training set (80 % of all fish per species), a validation set (5 %) and a test set (15 %). Because the splitting of the dataset is made on the fish entity (with all its images) and not on the images itself, the final partitioning of training/validation/test set according to numbers of images is slightly different. This is, because the numbers of images per fish vary (see above). However, this approach is chosen, to make sure that the final test set only contains fish individuals, of which no image is contained in the training data. The training set is then used for the actual training of the classification network. The validation set is used to monitor the accuracy on an independent set of data during the training process and for early stopping to avoid the effect of overfitting. The test set is used after the training process is finished for the final evaluation (presented below in Section 3).
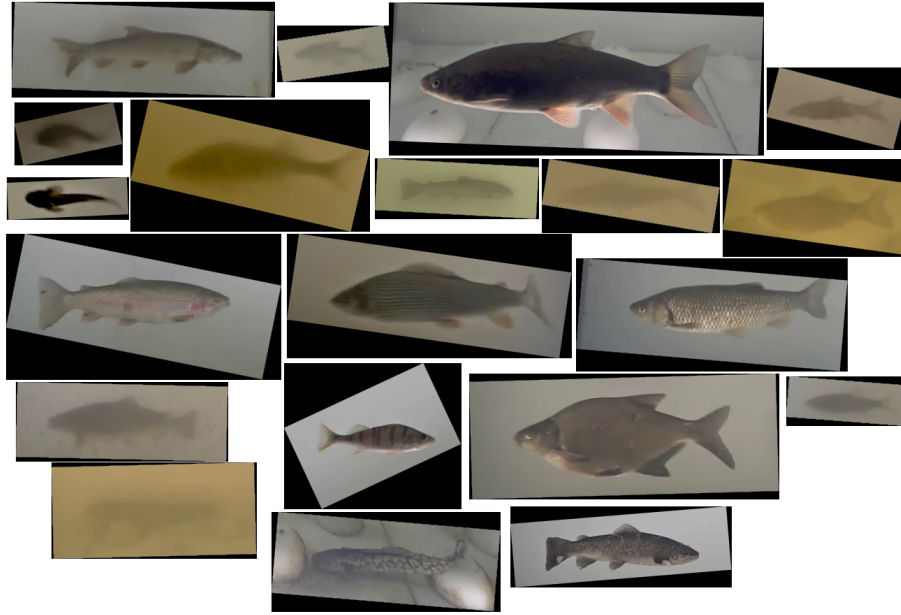
Figure 3: Random images of used data to present the difference in the image quality.

## 2.4 Training procedure

Because the dataset is highly unbalanced in the number of classes per sample, special care has to be taken during the training otherwise the network would be biased towards the most frequent class in the dataset. The approach used for this experiment is to provide batches during training, in which each fish species is similarly represented in numbers. A batch is a subset of the training data, used for updating the network parameters in one training set. This approach of oversampling samples of underrepresented classes logically implies that samples of these classes are shown more frequently to the network during training. To avoid the effect of overfitting various ways of data augmentation are used (Krizhevsky et al., 2012). Data augmentation is a process, in which the size of dataset is artificially increased through the application of random transformation on the data. For the images, these include small shifts in brightness, saturation, hue, contrast, as well as horizontal flips. In addition, the length feature is randomly changed during training by adding a random value in the range of [-20; 20] mm to the real value, measured by the human experts. This is done because:

1. Measured values are rounded by the human expert to even categories of 10 mm (e.g. 140 mm, 150 mm, 160 mm and so on).

2. A preliminary study in this research project has shown, that the accuracy of the measured length from the video imagery has an uncertainty range of approx. 2-3 cm.

It is worth noting at this point, that during the training is done on a per-sample base. This means, that the network tries to minimize the classification error made per sample and not per-fish (consisting of multiple samples).

## 2.5 Evaluation procedure

While during the training, the accuracy is calculate per training sample, for the final evaluation a more realistic approach is chosen. Because the main interest is to classify each fish correctly and of each fish, multiple images are available, the final prediction is made through combining the classification result of all images assigned to a fish. This can be made in different ways and in this study two different approaches are investigated:

1. **Mode:** For each sample (image + eventually additional features), the class with the highest class probability is chosen as the predicted class. The final prediction for a fish is made by selecting the class that got the most votes.

2. **Max-Prob:** The class probabilities of each sample are summed and as the final prediction, the class with the highest overall value is chosen.

The difference that might be important is that in the first approach (Mode) the information, if the classification network is uncertain about a class for a single sample (e.g. two classes with almost the same probability) is lost. In the second approach (Max-Prob) this information is conserved and might result in a difference overall score.

# 3 Results and discussion

All presented numbers derive from evaluating the trained networks on the independent test set, which consists of fish (and their corresponding images), the network has never seen before during the training period. In a first step, the overall performance of the four different network architectures is evaluated, combined with the two different aggregation strategies (see Section 2.5). In the next step, a more detail analysis of the best performing network is made.

## 3.1 Overall performance

Table 3.1 presents the overall classification accuracies of all network configurations.

| Network configuration | Accuracy per sample [%] | Accuracy per fish [%] | |
| --- | --- | --- | --- |
| | | Mode | Max-Prob |
| VGG16 | 85.0 | 89.7 | 90.3 |
| VGG16 + length | 85.7 | 89.8 | 90.2 |
| VGG16 + date | 88.5 | 92.2 | 92.3 |
| **VGG16 + both** | 89.4 | 93.0 | **93.3** |

Table 2: Overall classification accuracy of the different network architectures and aggregation types.

From the numbers presented in Table 3.1 it is visible that adding additional features has an overall positive effect on the classification accuracy. It is also visible that the length feature seems to be a less useful feature, compared to the date of migration. This is a somewhat pleasant result, since the date feature comes without any computational cost and can simply be read from the meta-data of the recorded video, while the length would have to be determined in a previous step. Studies have been made to automatize the length determination for the FishCam system, but it is a rather complicated process and only feasible in clear water conditions. As might be expected, it is also visible from Table 3.1 that deriving the final prediction from multiple images (samples) increases the overall prediction. Of the two different combination methods presented in Section 2.5, it can be seen in Table 3.1 that the Max-Prob approach performs marginally better, compared to the Mode approach. The two different approaches could potentially make a difference for cases, in which the classification network is rather uncertain between two classes for some images (e.g. two classes with almost equal probabilities) and more certain for some other images of the same fish (i.e. one class is a much higher probability then the others).

## 3.2 Analysis of the best performing network

In this subsection a more detailed analysis of the best performing architecture (VGG16 + both, see Section 2.2) is presented. The confusion matrix Figure 4 shows the predicted classes compared to the true classes of the test set data.

Overall, the classes of rather similar looking fish (Chub, Nase, Rainbow trout and Brown trout) have the highest rate of misclassifications. The wrong predicted classes are mainly made in between these classes as well. For example, Chubs get misclassified as Nase or Rainbow trout and Rainbow trouts are wrongly predicted as either Chub, Brown trout or Nase. The remaining classes have a rather low rate of misclassifications from zero misclassified samples (Burbot) to 3.6 % misclassified fish of the European perch class.

Predicted Class

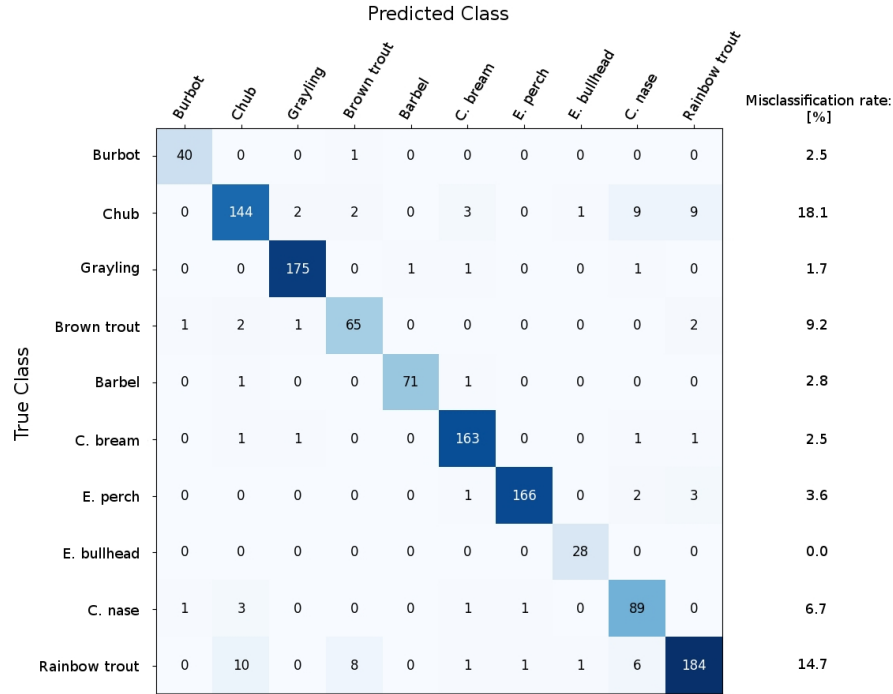| True Class | Burbot | Chub | Grayling | Brown trout | Barbel | C. bream | E. perch | E. bullhead | C. nase | Rainbow trout | Misclassification rate: [%] |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Burbot | 40 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 2.5 |
| Chub | 0 | 144 | 2 | 2 | 0 | 3 | 0 | 1 | 9 | 9 | 18.1 |
| Grayling | 0 | 0 | 175 | 0 | 1 | 1 | 0 | 0 | 1 | 0 | 1.7 |
| Brown trout | 1 | 2 | 1 | 65 | 0 | 0 | 0 | 0 | 0 | 2 | 9.2 |
| Barbel | 0 | 1 | 0 | 0 | 71 | 1 | 0 | 0 | 0 | 0 | 2.8 |
| C. bream | 0 | 1 | 1 | 0 | 0 | 163 | 0 | 0 | 1 | 1 | 2.5 |
| E. perch | 0 | 0 | 0 | 0 | 0 | 1 | 166 | 0 | 2 | 3 | 3.6 |
| E. bullhead | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 28 | 0 | 0 | 0.0 |
| C. nase | 1 | 3 | 0 | 0 | 0 | 1 | 1 | 0 | 89 | 0 | 6.7 |
| Rainbow trout | 0 | 10 | 0 | 8 | 0 | 1 | 1 | 1 | 6 | 184 | 14.7 |

Figure 4: Confusion matrix of the classification result of the test set. On the right side, the misclassification rate per fish species is provided.

## 4  Conclusion and outlook

In this study, a new approach for fish species classification in video-based monitoring is presented using a convolutional neural network. The classification is mainly based on the raw images from the video but it is shown how additional information (the date of the migration and the length of the fish) can help to enhance the classification accuracy. The study is conducted on a dataset containing 10 different species, occurring in Austrian rivers, coming from monitoring data of the FishCam monitoring system. Derived from multiple images per fish, the presented approach achieves a classification accuracy of 93.3 %, when using both additional informations (migration date and fish length). However, before this system could be used for operational monitoring of fish migration, more species have to be added to the dataset. Because the training data should come from the same system (in this case from monitoring using the FishCam detection tunnel), samples of some rarely seen species might be scarce. A different approach on the way to operational use could be to test, if prediction probabilities can be used to identify fish individuals that are not covered in the learned classes. These individuals could then be tagged for manual review. However, since this would be only the case for species that occur rarely, the approach presented in this study could possibly be applied in the future in operation monitoring.

## References

Alsmadi, M. K., Omar, K. B., Noah, S. A., and Almarashdah, I.: Fish Recognition Based on Robust Features Extraction from Size and Shape Measurements Using Neural Network, Journal of Computer Science, 6, 1060–1066, 2010.

Badawi, U. A. and Alsmadi, M. K. S.: A hybrid memetic algorithm (Genetic Algorithm and Tabu Local Search) with back-propagation classifier for fish recognition, IJCSI International Journal of Computer Sciences, 10, 384–356, 2013.

Farabet, C., Couprie, C., Najman, L., and LeCun, Y.: Learning hierarchical features for scene labeling, IEEE transactions on pattern analysis and machine intelligence, 35, 1915–1929, 2013.

Ioffe, S. and Szegedy, C.: Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift, in: Proceedings of the 32nd International Conference on Machine Learning, edited by Bach, F. and Blei, D., vol. 37, pp. 448–456, PMLR, Lille, France, URL http://proceedings.mlr.press/v37/ioffe15.html, 2015.

Kratzert, F.: Entwicklung einer Software zur automatisierten Erkennung von Fischen in videoüberwachten Fischaufstiegshilfen, Masterarbeit, Universität für Bodenkultur Wien, 2016.

Kratzert, F. and Mader, H.: FishCam – A semi-automatic video-based monitoring system of fish migration, EGU General Assembly Conf., 18, 2016.

Krizhevsky, A., Sutskever, I., and Hinton, G. E.: ImageNet Classification with Deep Convolutional Neural Networks, Advances In Neural Information Processing Systems, pp. 1097–1105, 2012.

Lee, D.-J., Schoenberger, R. B., Shiozawa, D., Xu, X., and Zhan, P.: Contour matching for a fish recognition and migration-monitoring system, Optics East, 5606, 37–48, https://doi.org/10.1117/12.571789, URL http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=852623, 2004.

Mader, H. and Kratzert, F.: The fishcam migration monitoring system for fish passes, in: Proceedings of the 11th International Symposium on Ecohydraulics, edited by Webb, J., Costelloe, J., Casas-Mulet, R., Lyon, J., and Stewardson, M., University of Melbourne, Melbourne, Australia, 2016.

Matai, J., Kastner, R., and Cutter, G. R.: Automated Techniques for Detection and Recognition of Fishes using Computer Vision Algorithms, NOAA Technical Memorandum, pp. 35–37, 2012.

Rova, A., Mori, G., and Dill, L. M.: One fish, two fish, butterfish, trumpeter: Recognizing fish in underwater video, MVA2007 IAPR Conference on Machine Vision Applications, May 16-18, 2007, Tokyo, JAPAN, pp. 404–407, URL http://www.cvl.iis.u-tokyo.ac.jp/mva/proceedings/2007CD/papers/11-02.pdf, 2007.

Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge, International Journal of Computer Vision, 115, 211–252, 2015.

Simonyan, K. and Zisserman, A.: Very Deep Convolutional Networks for Large-Scale Image Recognition, pp. 1–14, https://doi.org/10.1016/j.infsof.2008.09.005, URL http://arxiv.org/abs/1409.1556, 2014.

Spampinato, C., Giordano, D., Salvo, R. D., Fisher, R. B., and Nadarajan, G.: Automatic Fish Classification for Underwater Species Behavior Understanding, in: Proceedings of the first ACM international workshop on Analysis and retrieval of tracked events and motion in imagery streams, pp. 45–50, 2010.

Tompson, J., Jain, A., LeCun, Y., and Bregler, C.: Joint Training of a Convolutional Network and a Graphical Model for Human Pose Estimation, in: Proceedings of Advances in Neural Information Processing Systems 27, pp. 1799–1807, 2014.