

Article

Fuzzy Overclustering: Semi-Supervised Classification of Fuzzy Labels with Overclustering and Inverse Cross-Entropy

Lars Schmarje ^{1,*}, Johannes Brünger ¹, Monty Santarossa ¹, Simon-Martin Schröder ¹, Rainer Kiko ²
and Reinhard Koch ¹

¹ Multimedia Information Processing Group, Kiel University, 24118 Kiel, Germany; jbr@informatik.uni-kiel.de (J.B.); msa@informatik.uni-kiel.de (M.S.); sms@informatik.uni-kiel.de (S.-M.S.); rk@informatik.uni-kiel.de (R.K.)

² Laboratoire d'Océanographie de Villefranche, Sorbonne Université, 06230 Villefranche-sur-Mer, France; rainer.kiko@obs-vlfr.fr

* Correspondence: las@informatik.uni-kiel.de

Abstract: Deep learning has been successfully applied to many classification problems including underwater challenges. However, a long-standing issue with deep learning is the need for large and consistently labeled datasets. Although current approaches in semi-supervised learning can decrease the required amount of annotated data by a factor of 10 or even more, this line of research still uses distinct classes. For underwater classification, and uncurated real-world datasets in general, clean class boundaries can often not be given due to a limited information content in the images and transitional stages of the depicted objects. This leads to different experts having different opinions and thus producing fuzzy labels which could also be considered ambiguous or divergent. We propose a novel framework for handling semi-supervised classifications of such fuzzy labels. It is based on the idea of overclustering to detect substructures in these fuzzy labels. We propose a novel loss to improve the overclustering capability of our framework and show the benefit of overclustering for fuzzy labels. We show that our framework is superior to previous state-of-the-art semi-supervised methods when applied to real-world plankton data with fuzzy labels. Moreover, we acquire 5 to 10% more consistent predictions of substructures.

Keywords: semi-supervised; fuzzy; deep learning; noisy; real-world; plankton; marine



Citation: Schmarje, L.; Brünger, J.; Santarossa, M.; Schröder, S.-M.; Kiko, R.; Koch, R. Fuzzy Overclustering: Semi-Supervised Classification of Fuzzy Labels with Overclustering and Inverse Cross-Entropy. *Sensors* **2021**, *21*, 6661. <https://doi.org/10.3390/s21196661>

Academic Editor: Hyungsik Nam

Received: 24 August 2021

Accepted: 2 October 2021

Published: 7 October 2021

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



Copyright: © 2021 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

Over the past years, we have seen the successful application of deep learning to many underwater computer vision problems [1–4]. Automatic analysis of underwater data allows us to monitor ecological changes by evaluating large amounts of for example plankton data [5,6]. While it is relatively easy to create a lot of underwater image data, its analysis is time-consuming and thus expensive because the annotation requires trained taxonomists. The possible reasons for this issue include the huge amounts of data, the high imbalance between classes and the variability of annotations [7].

In underwater classification, domain experts often differ in their annotations [7–9]. This issue arises due to the following reasons: Firstly, automatically captured underwater images often have a lower quality than images taken manually by humans. This difference in quality arises for example due to the underwater lighting conditions and no manual corrections to e.g. insufficient sharpness or not centering the target inside the focus. For example the analysis of benthic images can suffer from these issues [8,9]. Even in the best scenario, a single image generally does not contain most of the information needed for a clear identification (e.g., three-dimensional configuration, minute morphological details, fluorescence). Secondly, intermediate stages actually exist between classes [10]. For example, in Figure 1 we show two different physical appearances (puff & tuft) of trichodesmium, while the dataset also contains intermediate stages between these two classes.

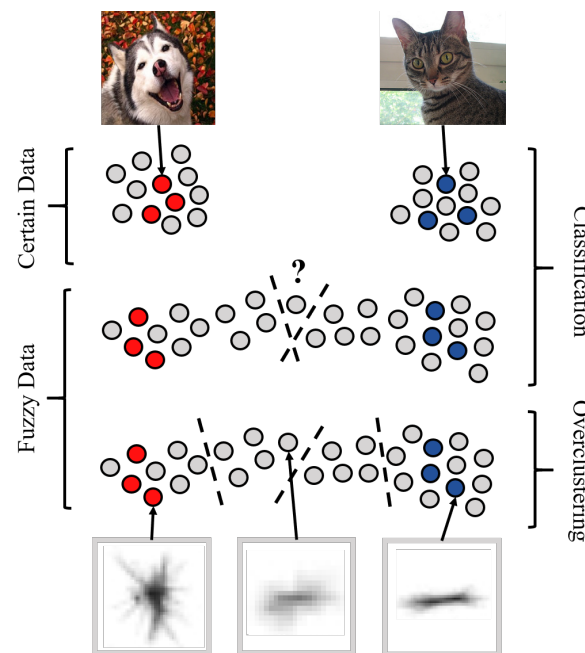


Figure 1. Illustration of fuzzy data and overclustering—The grey dots represent unlabeled data and the colored dots labeled data from different classes. The dashed lines represent decision boundaries. For certain data, a clear separation of the different classes with one decision boundary is possible and both classes contain the same amount of data (**top**). For fuzzy data determining a decision boundary is difficult because of intermediate datapoints between the classes (**middle**). These fuzzy datapoints can often not be easily sorted into one consistent class between annotators. If you overcluster the data, you get smaller but more consistent substructures in the fuzzy data (**bottom**). The images illustrate possible examples for certain data (cat & dog) and fuzzy plankton data (trichodesmium puff and tuft). The center plankton image was considered to be trichodesmium puff or tuft by around half of the annotators each. The left and right plankton image were consistently annotated.

This issue of different annotations is also known as *intra-* and *inter-observer* variability [11] and is common in many biological and medical application fields [8,9,12–17]. Even in a curated dataset [1], we quote Tarling et al. who state “there will very likely be inaccuracies, bias, and even inconsistencies in the labeling which will have affected the training capacity of the model and lead to discrepancies between predictions and ground truths” [18]. When aggregating multiple annotations per image, we call the resulting label fuzzy if we have different annotations between experts (non-zero variance), and certain if all annotations agree with each other. The mathematical formulation of a fuzzy label would be a unknown soft probability distribution l for k classes. The distribution $l \in (0, 1)^k$ can only be approximated with a high cost e.g., by averaging over multiple annotations.

Semi- and Self-Supervised Learning are promising approaches to decrease the needed amount of annotated data by a factor of 10 or even more [19–21]. These approaches leverage unlabeled data in addition to the normal labeled data to improve the training. A common strategy is to define a pretext task like image rotation prediction [22] or mutual information maximization [23] for pretraining. A broad overview of current trends, ideas and methods in semi-, self- and unsupervised learning is available in [24]. However, this research mainly focuses on established curated classification datasets such as STL-10 [25]. In these datasets, a clear distinction between classes such as cats and dogs are given. The hard partitioning of intermediate morphologies is not appropriate and does not allow the identification of substructures. We show that state-of-the-art semi-supervised algorithms are not well suited to handle fuzzy labels. These algorithms expect only certain labels as shown in the upper part of Figure 1. If we apply previous semi-supervised algorithms to fuzzy data which include fuzzy images, these algorithms arbitrarily assign undecidable images to one class (middle part of Figure 1).

Noisy labels are a common data quality issue and are discussed in the literature [11,26,27]. The fuzziness of labels is known as a special case of label noise that exist “due to subjectiveness of the task for human experts or the lack of experience in annotator[s]” [26]. In contrast to us, most methods [28–30] and literature surveys [11,26,27] interpret fuzzy labels as corrupted labels. We argue that fuzzy labels are valid signals derived from ambiguous images and that it is important to discover the substructures for real-world data handling [12–17].

Geng proposed to learn the label distribution to handle fuzzy data [31] and the idea was extended to the application of real-world images [32]. However, these methods are not semi-supervised and therefore depend on large labeled datasets. A variety of methods was proposed to handle fuzzy data in a semi-supervised learning approach [33–35]. These methods use lower-dimensional features spaces in contrast to images as input. Liu et al. proposed to use independent predictions of multiple networks as pseudo-labels for the estimation of the label distribution for photo shot-type classification [36]. We argue that the true label distribution is difficult to approximate and thus difficult to evaluate. We do not learn the label distribution but use clustering to identify substructures.

We propose *Fuzzy Overclustering* (FOC) which separates the fuzzy data into a larger number of visual homogeneous clusters (lower part, Figure 1) which can then be annotated very efficiently [10]. We will show on a Plankton dataset that state-of-the-art semi-supervised algorithms perform worse on fuzzy data in comparison to our method FOC which explicitly considers fuzzy images. Moreover, we will show that this leads to 5 to 10% more self-consistent predictions of plankton data.

One main idea is to rephrase the handling of fuzzy labels as a semi-supervised learning problem by using a small set of certain images and a large number of fuzzy images that are treated as unlabeled data. This approach allows us to use the idea of overclustering from semi-supervised literature [23,37] and apply it to fuzzy data. The difference to previous work is that we use overclustering not only to improve classification accuracy on the labeled data but improve the clustering and therefore the identification of substructures of fuzzy data. We show that overclustering allows us to cluster the fuzzy images in a more meaningful way by finding substructures and therefore allowing experts to analyze fuzzy images more consistently in the future.

We show the benefits of our method mainly on a plankton dataset which highlights the benefit for underwater classification. However, the issue of fuzzy labels is neither limited to plankton data nor to underwater classification. On a synthetic dataset, we show a proof-of-concept for the generalizability of our model to other datasets.

Our key contributions are:

- We identify an issue of semi-supervised algorithms that they do not work well with fuzzy labels. However, such fuzzy labels occur regularly in underwater image classification e.g due to high natural variation of depicted objects which leads to a high inter- and intraobserver variability.
- We propose a novel framework for handling fuzzy labels with a semi-supervised approach. This framework uses overclustering to find substructures in fuzzy data and outperforms common state-of-the-art semi-supervised methods like FixMatch [38] on fuzzy plankton data.
- We propose a novel loss, *Inverse Cross-entropy* (CE^{-1}), which improves the overclustering quality in semi-supervised learning.
- We achieve 5 to 10% more self-consistent predictions on fuzzy plankton data.

2. Method

Our framework Fuzzy Overclustering (FOC) aims at creating an overclustering for fuzzy labels by using an auxiliary classification and not the other way round like previous literature [23,37]. In this section, we describe our framework in general and explain important parts in detail in the following subsections. We use the following notation for the given semi-supervised classification task. Our training data consists of the two subsets

X_l and X_u . X_l is a labeled image dataset with images $x \in X_l$ and corresponding labels y . X_u is an unlabeled image dataset, i.e., there is/exists no label for images $x \in X_u$.

We generate three inputs x_1, x_2, x_3 based on one image $x \in X_l \cup X_u$ depending on the availability of the corresponding label y . If y is not available, the images x_1 and x_2 are augmented views of x and x_3 is an augmented version of a random image $x' \in X_l \cup X_u$. If y is available, x_1 is an augmented view of x , x_2 is a supervised augmentation (see Section 2.3) and x_3 an inverse example. For the inverse example, we choose an image $x' \in X_l$ with a different label y' ($y' \neq y$). We use an augmented version of this image as third input $x_3 = g_3(x')$ with augmentation g_3 . We constraint the ratio from unlabeled to labeled data to a fixed ratio r to improve the run time of the model (see Section 2.4). The inputs are processed by a neural network Φ which is composed of a backbone like ResNet50 [39] and linear output prediction layers. Following [23], we call this linear predictors *heads* and use them either as normal or overclustering heads. As output we use the soft-max classifications of these normal and overclustering heads. If k_{GT} is the number of ground-truth classes a normal head outputs a probability for each of the k_{GT} classes. The overclustering head has k output nodes with $k > k_{GT}$ and give probabilities for more clusters than ground-truth classes (overclustering). Both type of heads are therefore fully connected layers with softmax activation but of different output size. We can average the training over multiple independent heads per type as shown in [23]. We use the notation Φ_{n_i} or Φ_{o_i} for the i -th normal or overclustering head respectively. An overview about the general pseudo code of FOC including the loss calculation is given in Algorithm 1.

For both heads the loss is different but can be written as the weighted sum of an unsupervised and a supervised loss as follows:

$$\mathcal{L} = \lambda_s \mathcal{L}_s + \lambda_u \mathcal{L}_u \quad (1)$$

\mathcal{L}_s is cross-entropy (\mathcal{L}_{CE}) for the normal head and our novel CE^{-1} loss ($\mathcal{L}_{CE^{-1}}$) for the overclustering head (see Section 2.1). For both heads \mathcal{L}_u is the mutual information loss \mathcal{L}_{MI} (see Section 2.2). An illustration of the complete pipeline is given in Figure 2. We initialize our backbones with pretrained weights and can therefore directly use RGB images as input. For further implementation details see Section 3.2.

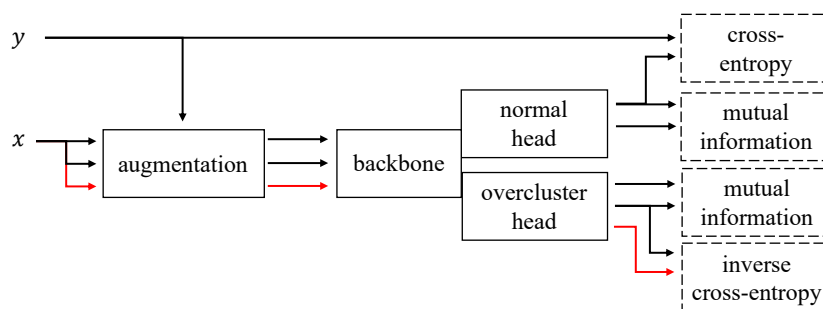


Figure 2. Overview of our framework FOC for semi-supervised classification—The input image is x and the corresponding label is y . The arrows indicate the usage of image or label information. Parallel arrows represent the independent copy of the information. The usage of the label for the augmentations is described in Section 2.3. The red arrow stands for an inverse example image x' with a different label than y . The output of the normal and the overclustering head have different dimensionalities. The normal head has as many outputs as ground-truth classes exist (k_{GT}) while the overclustering head has k outputs with $k > k_{GT}$. The dashed boxes on the right side show the used loss functions. More information about the losses inverse cross-entropy and mutual information can be found in Sections 2.1 and 2.2 respectively.

If we use FOC with $\lambda_s = 0$ and without supervised augmentations our model is comparable to the pretext task of Invariant Information Clustering (IIC) [23]. We can use this configuration as a warm-up to pretrain the weights. During the evaluation, we

will refer to using the pretext task for IIC and the warm-up of FOC synonymously. Our framework FOC can also be used to perform standard unsupervised clustering. The details about unsupervised clustering and a comparison to previous literature is given in the supplementary.

Algorithm 1: Pseudocode for our method Fuzzy Overclustering

Data: Batch of images of size b from labeled image data X_l and unlabeled image data X_u

Result: calculate loss value for one given batch for a network Φ with n normal and overclustering heads

L : matrix of size $b \times 2n$;

```

/* iterate over batch */
for  $i \leftarrow 0$  to  $b$  do
   $x \leftarrow i$ -th image in batch;
  if label  $y$  for image  $x_i$  available then
     $x_1 \leftarrow g_1(x)$  with random augmentation  $g_1$ ;
    /* Supervised augmentation defined in Section 2.3 */
     $x_2 \leftarrow g_2(x)$  with supervised augmentation  $g_2$ ;
    /* Inverse example defined in Section 2 */
     $x_3 \leftarrow g_3(x')$  with random augmentation  $g_3$  and inverse example  $x'$ ;
  else
     $x_1 \leftarrow g_1(x)$  with random augmentation  $g_1$ ;
     $x_2 \leftarrow g_2(x)$  with random augmentation  $g_2$ ;
     $x_3 \leftarrow g_3(x')$  with random augmentation  $g_3$  and random image  $x'$ ;
  end
  /* iterate over heads */
  for  $j \leftarrow 0$  to  $n$  do
    calculate forward pass for outputs  $\Phi_{n_j}$  and  $\Phi_{o_j}$ ;
    /* CE loss for normal head */
     $L[i,j] \leftarrow \mathcal{L}_{CE}(\Phi_{n_j}(x_i), l_i)$  with  $l_i$ ;
    /*  $CE^{-1}$  loss for overclustering head */
     $L[i,j+n] \leftarrow \mathcal{L}_{CE^{-1}}(x_1, x_2, x_3)$  with Equation (2);
  end
end
/* calculate loss */
 $\mathcal{L}_s \leftarrow$  average supervised loss across heads and batch from  $L$ ;
 $\mathcal{L}_u \leftarrow$  unsupervised MI loss across batch with Equations (3) and (4);
 $\mathcal{L} \leftarrow \lambda_s \mathcal{L}_s + \lambda_u \mathcal{L}_u$ ;

```

2.1. Inverse Cross-Entropy (CE^{-1})

Inverse Cross-Entropy is a novel supervised loss for an overclustering head and one of the key contributions of this work. The loss is needed to use the label information for an overclustering head. For normal heads, we can use cross-entropy (CE) to penalize the divergence between our prediction and the label. We can not use CE directly for the overclustering heads since we have more clusters than labels and no predefined mapping between the two. However, we know that the inputs x_1/x_2 and x_3 should not belong to the same cluster. Therefore, our goal with CE^{-1} is to define a loss that pushes their output distributions (e.g., $\Phi(x_1)$ and $\Phi(x_3)$) apart from each other.

Let us assume we could define a distribution that $\Phi(x_3)$ should not be. In short, an inverse distribution $\Phi(x_3)^{-1}$. If we had such a distribution we could use CE to penalize the divergence for example between $\Phi(x_1)$ and $\Phi(x_3)^{-1}$.

One possible and easy solution for an inverse distribution is $\Phi(x_3)^{-1} = 1 - \Phi(x_3)$. For a binary classification problem, $\Phi(x_3)^{-1}$ can even be interpreted as a probability distribution again. This is not the case for a multi-class classification problem. We could use

a function like softmax to cast $\Phi(x_3)^{-1}$ into a probability distribution but decided against it for three reasons. Firstly, we would penalize correct behavior. For example in a three class problem with $\Phi_1(x_1) = 0.5 = \Phi_2(x_1)$ and $\Phi_3(x_3) = 1$ we only get $CE(\Phi(x_1), \Phi(x_3)^{-1}) = 0$ if $\Phi(x_3)^{-1}$ is not a probability distribution. Otherwise either $\Phi_1(x_3)^{-1}$ or $\Phi_2(x_3)^{-1}$ have to be real smaller than 1. Secondly, we are still minimizing the entropy of $\Phi(x_1)$ which leads to more confident predictions in semi-supervised learning [19,20,40–43]. The proof is given in the supplementary. Thirdly, it is easier and in practice, it is not needed. For the input $i = (x_1, x_2, x_3)$, we define the cross-entropy inverse loss $\mathcal{L}_{CE^{-1}}$ as shown in Equation (2).

$$\begin{aligned} \mathcal{L}_{CE^{-1}}(i) &= 0.5 \cdot CE^{-1}(\Phi(x_1), \Phi(x_3)) \\ &\quad + 0.5 \cdot CE^{-1}(\Phi(x_2), \Phi(x_3)), \text{ with} \\ CE^{-1}(p, q) &= - \sum_{c=1}^k p(c) \cdot \ln(1 - q(c)). \end{aligned} \quad (2)$$

2.2. Mutual Information (MI)

For the unlabeled data, we use the loss proposed by Ji et al. because it is calculated directly on the output clusters [23]. Therefore similar images are pulled to the same clusters while CE^{-1} pushes different images apart. For this purpose, we want to maximize the mutual information between two output predictions $\Phi(x_1), \Phi(x_2)$ with x_1, x_2 images which should belong to the same cluster and $\Phi : X \rightarrow [0, 1]^k$ a neural network with k output dimensions. We can interpret $\Phi(x)$ as the distribution of a discrete random variable z given by $P(z = c|x) = \Phi_c(x)$ for $c \in \{1, \dots, k\}$ with $\Phi_c(x)$ the c -th output of the neural network. With z, z' such random variables we need the joint probability distribution for $P_{cc'} = P(z = c, z' = c')$ for the calculation of the mutual information $I(z, z')$. Ji et al. propose to approximate the matrix P with the entry $P_{cc'}$ at row c and column c' by averaging over the multiplied output distributions in a batch of size n [23]. Symmetry of P is enforced as shown in Equation (3).

$$P = \frac{Q + Q^T}{2} \text{ with } Q = \frac{1}{n} \sum_{i=1}^n \Phi(x_i) \cdot \Phi(x'_i)^T \quad (3)$$

We can maximize our objective $I(z, z')$ with the marginals $P_c = P_{c'} = P(z = c)$ given as sums over the rows or columns as shown in Equation (4).

$$I(z, z') = \sum_{c=1}^k \sum_{c'=1}^k P_{cc'} \cdot \ln \frac{P_{cc'}}{P_c \cdot P_{c'}} \quad (4)$$

2.3. Supervised Augmentations

In the unsupervised pretraining, we use the same image x to create the two inputs $x_1 = g_1(x)$ and $x_2 = g_2(x)$ based on the augmentations g_1 and g_2 . Otherwise, without supervision, it is difficult to determine similar images. However, if we have the label y for x we can use a secondary image $x' \in X_l$ with the same label to mock an ideal image transformation to which the network should be invariant. In this case we can create $x_2 = g_2(x')$ based on the different image. We call this *supervised augmentation*.

2.4. Restricted Unsupervised Data

Unlabeled data has a small impact on the results but drastically increases the runtime in most cases. The increased runtime is caused by the facts that we often have much more unlabeled data than labeled data and that a neural network runtime is normally linear in the number of samples it needs to process. However, unlabeled data is essential for our proposed framework and we can not just leave it out. We propose to restrict the unlabeled data to a fixed upper-bound ratio r in every batch and therefore the unlabeled data per epoch. Detailed examples and experiments are given in the supplementary. It is important

to notice that we restrict only the unlabeled data per batch/epoch. While for one epoch the network will not process all unlabeled data, over time all unlabeled data will be seen by the network. We argue that the impact on training time negatively outweighs the small benefit gained from all unlabeled data per epoch.

3. Experiments

We conducted our experiments mainly on a real-world plankton dataset. We used the common image classification dataset STL-10 as a comparison with only certain labels and a synthetic dataset for a proof-of-concept for the generalizability to other datasets. We compare ourselves to previous work and make several ablations. Additional results like unsupervised clustering, more detailed ablations and further details are given in the supplementary material.

3.1. Datasets

While the issue of fuzzy labels is present in multiple datasets [12–17], they are not well suited for evaluations. If we want to quantify the performance on fuzzy labels, we need a dataset with very good fuzzy ground-truth. This can only be achieved with a high cost e.g. by multiple annotations and thus is often not feasible. For all used datasets, we ensure that the labeled training data only consists of certain images and that the fuzzy images are used as unlabeled data. If we include fuzzy labels in the labeled data which is used as guidance during training, this will lead to worse performance as illustrated in the ablations (Table 3).

3.1.1. Plankton

The plankton dataset contains diverse grey-level images of marine planktonic organisms. The images were captured with an Underwater Vision Profiler 5 [44] and are hosted on EcoTaxa [45]. In the citizen science project PlanktonID (<https://planktonid.geomar.de/en> (accessed on 6 October 2021)), each sample was classified multiple times by citizen scientists. The data for the PlanktonID project is a subset of the data available on EcoTaxa [45]. It was presorted to contain a more balanced representation of the available classes. The dataset consists of 12,280 images in originally 26 classes. We merged minor and similar classes so that we ended up with 10 classes. The class no-fit represents a mixture of left-over classes. The merging was necessary because some classes had too few images for current state-of-the-art semi-supervised approaches. After this process, a class imbalance is still present with the smallest class containing about 4.16% and the largest class 30.37% of all samples. We use the mean over all annotations as the fuzzy label. The citizen scientists agree on most images completely. We call these images and their label *certain*. However, about 30% of the data has at least one disagreeing annotation. We call these images and their label *fuzzy* and use the most likely class as ground-truth if we need a hard label for evaluation. The fuzzy labeled images are used only as unlabeled data. More details about the mapping process, the number of used samples and graphical illustrations are given in the supplementary.

3.1.2. STL-10

STL-10 is a common semi-supervised image classification dataset [25] and a subset of ImageNet [46]. It consists of 5000 training samples and 8000 validation samples depicting everyday objects. Additionally, 100,000 unlabeled images are provided that may belong to the same or different classes than the training images. In contrast to the plankton and synthetic dataset, no labels are provided for the unlabeled data and no fuzzy datapoints exist. We use this dataset only to illustrate the difference in the performance of FOC to previous semi-supervised methods.

3.1.3. Synthetic Circles and Ellipses (SYN-CE)

This dataset is a mixture of circles and ellipses (bubbles) on a black background with different colors. The 6 ground-truth classes are blue, red and green circles or ellipses. An image is defined as *certain* if the hue of the color is 0 (red), 120 (green) or 240 (blue) and the main axis ratio of the bubble is 1 (circle) or 2 (ellipse). Every other datapoint is considered *fuzzy* and the ground-truth label l is calculated as the product of the interpolation of the color p_c and the geometry p_g distribution. More details are in the supplementary. The dataset consists of 1800 *certain* and 1000 *fuzzy* labeled images for train, validation and unlabeled data split. We look at three subsets: *Ideal*, *Real* and *Fuzzy*. The *Ideal* subset uses the maximal class of the fuzzy label l as a ground-truth class and represents the ideal case that we certainly know the most likely label to each image. For the *Real* subset, the ground-truth classes are randomly picked with the distribution of the fuzzy label l and represent the real or common case. For example due to only one annotation, the percentage that the label corresponds to the actual most likely class is linear to the fuzzy label. The *Fuzzy* subset only uses *certain* labeled images as training data and represent a cleaned training dataset. We will show that this handling of fuzzy labels leads to a higher classification performance in comparison to the *Real* dataset in Section 3.5.1. The *Ideal* and the *Real* subset can be evaluated on the unlabeled data of the *Fuzzy* subset with some overlap in the images.

3.2. Implementation Details

As a backbone for our framework, we used either a ResNet34 variant [23] or a standard ResNet50v2 [39]. The heads are single fully connected layers with a softmax activation function. Following [23], we use five randomly initialized copies for each type of head and repeat images per batch three times for more stable training. We alternated between training the different types of heads. The inputs are either sobel-filtered images or color images for pretrained networks. For the ResNet34 backbone, we use CIFAR20 (20 superclasses in CIFAR-100 [47]) weights and for the ResNet50v2 backbone ImageNet [46] weights. We use in general $\lambda_s = 1 = \lambda_u$ and an unlabeled data restriction of $r = 0.5$. We call our Framework FOC-Light if we use $\lambda_u = 0$ and no warm-up. This means we do not use the loss introduced by [23] and therefore also do not have to use their stabilization methods like repetitions. During the pretext task or warm-up and the main training, we train the framework with Adam and an initial learning rate of 1×10^{-4} for 500 epochs. When switching from the pretext task to fine-tuning, we train only the heads for 100 epochs with a learning rate of 1×10^{-3} before switching to the lower learning rate of 1×10^{-4} . The number of outputs for the overclustering head should be about 5 to 10 times the number of classes. The exact number is not crucial because it is only an upper bound for the framework. We use 70 for STL-10 and 60 for the plankton dataset. We selected all hyperparameters heuristically based on the STL-10 dataset and did not change them for the plankton dataset. We used the recommended hyperparameters by the original authors for the previous methods. We compared with the following methods Semantic Clustering by Adopting Nearest neighbors (SCAN) [48], Information Invariant Clustering (IIC) [23], Mean-Teacher [49], Pi(-Model) [29], Pseudo-label [50] and FixMatch [38]. More detailed descriptions are given in the supplementary.

3.3. Metrics

The evaluation protocols vary slightly depending on the used output and dataset. The used data splits training, validation and unlabeled are defined above in Section 3.1.

On STL-10, we calculate accuracy of the validation data. Accuracy is the portion of true positive and true negatives from the complete dataset.

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \quad (5)$$

TP, TN, FP and FN are the true positive, true negative, false positive and false negative respectively. We calculate these values per class and then sum the up before calculating

the accuracy (micro averaging). For the overclustering head, we need to find a mapping between the output clusters and the given classes. We calculate this mapping based on the majority class in each cluster on the training data as in [23].

On the fuzzy plankton and synthetic datasets, we evaluate the macro-f1 score on the unlabeled data. We calculate the macro F1-Score i.e., the average of the F1-scores per class due to the skewed class distribution.

$$\text{F1-Score} = \frac{2\text{TP}}{2\text{TP} + \text{FP} + \text{FN}} \quad (6)$$

Mind that a micro averaged F1-Score would be in our case the same as the above defined accuracy. We use the unlabeled data as evaluation dataset because the fuzzy images, in which we are interested, are only included in the unlabeled data split by definition. The mapping for the overclustering head is calculated based on the unlabeled data split because we expect human experts to be involved in this process for the identification of substructures. The best unlabeled results of the fuzzy Plankton and Synthetic dataset are reported based on the validation metrics.

If not stated otherwise, we report the maximum score for the overclustering and the normal head and the average and standard deviation over 3 independent repetitions.

3.4. Results

3.4.1. State-of-the-Art Comparison

We compare the state-of-the-art methods on certain and fuzzy data in Table 1.

Table 1. Comparison of state-of-the-art on certain and fuzzy data—We use STL-10 as a certain dataset and the Plankton data as a fuzzy dataset. We report the Accuracy for STL-10 and the F1-Score for the Plankton data due to class imbalance. It is important to notice that STL-10 is a curated dataset while the Plankton dataset still contains the fuzzy images. For more details about the metrics see Section 3.3. The results of previous methods are reported in the original paper or the original authors code was used to replicate the results. The best results are marked bold. Legend: † A MLP used for fine-tuning. ‡ Used only 1000 labels instead of 5000. * Unsupervised method.

Method	Network	Type of Data	
		Certain	Fuzzy
SCAN * [48]	ResNet18	76.80 ± 1.10	37.64 ± 3.56
IIC [23]	ResNet34	85.76 ± 1.36	65.47 ± 1.86
IIC † [23]	ResNet34	88.8	66.81 ± 1.85
Mean-Teacher [49]	Wide ResNet28	78.577 ± 2.39 ‡	72.85 ± 0.46
Pi [29]	Wide ResNet28	73.77 ± 0.82 ‡	74.34 ± 0.58
Pseudo-label [50]	Wide ResNet28	72.01 ± 0.83 ‡	75.04 ± 0.52
FixMatch [38]	Wide ResNet28	94.83 ± 0.63 ‡	76.28 ± 0.27
FOC-Light (Ours)	ResNet50	–	72.79 ± 2.99
FOC (Ours)	ResNet50	86.12 ± 1.22	76.79 ± 1.18

We see that FOC reaches a performance of about 86% on certain data but is not able to reach the performance of FixMatch. FixMatch outperforms FOC by a clear margin of nearly 8% while using a fifth of the labels. This performance is expected as FOC does not focus like the others on classifying certain but fuzzy data. If we look at the less curated fuzzy Plankton dataset, we see that FOC outperforms all all methods by a small margin. All previous methods focus on certain and curated data and we see this leads to a huge performance degeneration if they are applied to fuzzy data. FixMatch reaches in both datasets the best performance except for our method FOC. We conclude that the overclustering from FOC is the key for handling fuzzy data because it allows more flexibility during training. Previous semi-supervised methods did not consider the issue of inter- and intraobserver variability and thus are worse than FOC in classifying fuzzy data.

If we use FOC-Light without the loss and stabilization of [23] the F1-Score drops slightly to 75% but the used GPU hours can be decreased from 58 to 4 h. We conclude that the overclustering head is more suitable for handling fuzzy real-world data as we assumed at the beginning. Moreover, we see that the combination of cross-entropy and our novel loss CE^{-1} can also successfully train an overclustering head.

3.4.2. Consistency

Up to this point, we analyzed classification metrics based on the 10 ground-truth classes but the quality of substructures was not evaluated. We can judge the consistency of each image within its cluster with the help of experts as a quality measure. An image is consistent if an expert views it as visually similar to the majority of the cluster. The consistency is calculated by dividing the number of consistent images by all images. The consistency over all classes or per class for FOC and FixMatch is given in Table 2 and raw numbers are provided in the supplementary. We provide a comparison based on all data and without the no-fit class because this class contains a mixture of different plankton entities. Visual similarity is therefore difficult to judge because it can only be defined by not being similar to the other nine classes. Based on the F1-Score, FixMatch and FOC perform similarly but if we look at the consistency we see that FOC is more than 5% more consistent than FixMatch. If we exclude the class no-fit from the analysis, FOC reaches a consistency of around 86% in comparison to 77% from FixMatch. For both sets, our method FOC reaches a higher average consistency per cluster and lower standard deviation. This means the clusters produced by FOC are more relevant in practice because there are fewer low-quality clusters which can not be used. Overall, this higher consistency can lead to faster and more reliable annotations.

Table 2. Consistency comparison on plankton dataset—The consistency is rated by experts over the complete data and a subset without the class no-fit. The score is given overall as an average per cluster with standard deviation and is described in Section 3.4.2. The best results are marked bold.

Method	All Data		Ignore Class No-Fit	
	Overall	Per Cluster	Overall	Per Cluster
FixMatch [38]	82.56	78.78 ± 28.22	77.11	69.61 ± 29.41
FOC (Ours)	87.80	79.66 ± 18.88	86.31	86.41 ± 13.68

3.4.3. Qualitative Results

We illustrate some qualitative results of FOC in Figure 3. All images in a cluster are visually similar, even the probably wrongly assigned images (red box). For the images in the first row, the annotators are certain that the images belong to the same class. In the second row, annotators show a high uncertainty of assignment between the two variants of the same biological object. This illustrates the benefit of overclustering since visual similar items are in the same cluster even for uncertain annotations. In a consensus process for the second row, experts could decide if the cluster should be the puff, tuft or a new borderline class. Moreover, this clustering could be beneficial for monitoring the current imaging process. We provide more randomly selected results in the supplementary.

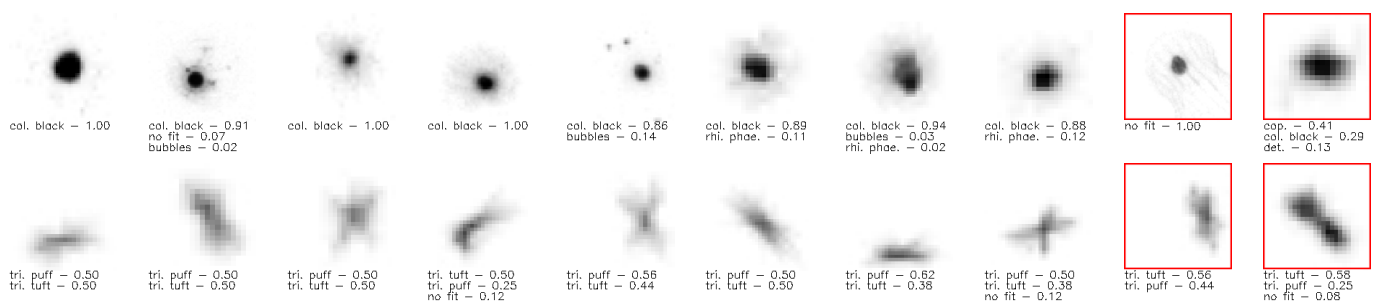


Figure 3. Qualitative results for unlabeled data—The results in each row are from the same predicted cluster. The three most important fuzzy labels based on the citizen scientists’ annotations are given below the image. The last two items with the red box in each row show examples not matching the majority of the cluster.

3.5. Ablation Studies

3.5.1. SYN-CE

We compare our framework with some previous methods on the three subsets of SYN-CE in Table 3. All semi-supervised methods reach almost a F1-Score of 100% on the unlabeled fuzzy data for the subset *Ideal*. In real-world data, it is unlikely that we have the real fuzzy ground-truth labels. It is more likely that we have uncertain/wrong labels for training and validation or no labels at all for fuzzy data like in the subsets *Real* or *Fuzzy*. In both cases, we see that our method reaches a superior performance with up to 10% higher F1-Score. While FOC-Light is only slightly better in comparison to the other semi-supervised methods on the *Real* subset it is comparable to the complete framework on the *Fuzzy* dataset. This is one indication that CE^{-1} is one of the key components for successfully training the overclustering heads. We see the F1-Score on the *Fuzzy* subset is around 10% higher than on the *Real* subset. We conclude that FOC can also generalize to other datasets. We conclude that these results support our idea of separating certain and fuzzy data during training because we do not need to potentially falsely approximate the real fuzzy ground-truth label like in the *Real* subset.

Table 3. Comparison to state-of-the-art on SYN-CE datasets—Each column represents a subset of the dataset SYN-CE. The results are F1-Scores which were calculated on the unlabeled data which include the fuzzy labels. All results within a one percent margin of the best result are marked bold.

Method	Ideal	Real	Fuzzy
Mean-Teacher [49]	97.11 ± 0.78	73.23 ± 2.49	66.57 ± 16.27
Pi [29]	98.44 ± 0.28	72.74 ± 2.43	77.69 ± 5.02
Pseudo-label [50]	98.17 ± 0.30	75.70 ± 1.98	89.48 ± 1.94
FixMatch [38]	98.32 ± 0.01	71.81 ± 1.06	93.82 ± 1.83
FOC-Light (Ours)	97.46 ± 4.39	78.77 ± 7.83	94.29 ± 0.87
FOC (Ours)	97.72 ± 4.52	83.86 ± 4.21	94.15 ± 0.29

3.5.2. Loss & Network

In Table 4 multiple ablations for STL-10 and the plankton dataset are given. The scores are averaged across the different output heads of our framework. Based on these tables, we illustrate the impact of the warm-up, the initialization and the usage of the MI and CE^{-1} loss for our framework. The normal accuracy can be improved by about 10% when using the unsupervised warm-up on the STL-10 dataset. On the plankton dataset, the impact is less but tends to give better results of some percent. Warm-up in combination with the MI loss leads to a performance which is not more than 10% worse than the full setup for all ablations except for one. For this exception, CE^{-1} is needed to stabilize the overclustering performance due to the poor initialization with CIFAR-20 weights. We attribute this worse performance to the initialization and not the different backbone because on STL-10 the CIFAR-20 initializations of the ResNet34 backbone outperform the ImageNet

weights of the ResNet50v2 backbone. We believe the positive effects of ImageNet weights for its subset STL-10 and the better network are negated by the different loss.

IIC is similar to FOC with warm-up and no additional losses but we also train an overclustering head for handling fuzzy data. Taking this into consideration, we achieve an 8 to 11% better F1-Score than IIC. A special case is FOC-light which does only use the CE^{-1} loss and therefore no stabilization method proposed in [23]. This decreases gpu memory usage and runtime and results in a total decrease of the GPU hours from 58 to 4 h. Overall, our novel loss CE^{-1} improves the overclustering performance regardless of the dataset and the weight initialization by 10% on STL-10 and up to 7% on the plankton dataset. We see that CE^{-1} is a key component for training an overclustering head and can even be trained without the stabilization of the warm-up and the MI loss.

Table 4. Ablation study—The second to fourth column indicates if a warm-up, the MI loss or our CE^{-1} loss were used respectively. The fifth column indicates if CIFAR-20 (C), ImageNet (I) or no (–) weights were used. Sobel filtered images are used as input for no weights. The Top1 and Top3 results are marked bold respectively. * Original authors code. † A MLP used for fine-tuning.

Method	Warm	MI	CE^{-1}	Weight	Accuracy	
					Overcluster	Normal
FOC		X		–	70.92 ± 2.42	76.39 ± 0.05
IIC * [23]	X			–		85.76
FOC	X	X		–	73.88 ± 0.21	82.01 ± 5.31
FOC	X	X	X	–	82.59 ± 0.06	86.49 ± 0.01
FOC	X	X	X	C	84.36 ± 0.64	78.59 ± 7.40
FOC	X	X	X	I	83.57 ± 0.10	85.21 ± 0.03
(a) STL-10						
Method	Warm	MI	CE^{-1}	Weight	F1-Score	
					Overcluster	Normal
IIC [23]	X			–	–	66.63
IIC † [23]	X			–	–	69.92
FOC				C	31.45 ± 6.02	39.35 ± 1.30
FOC		X		C	29.82 ± 2.98	60.65 ± 0.02
FOC		X	X	C	70.11 ± 1.99	64.10 ± 0.13
FOC	X			C	23.95 ± 2.63	58.71 ± 2.07
FOC	X	X		C	69.36 ± 0.05	56.59 ± 0.04
FOC	X	X	X	C	70.68 ± 0.10	58.09 ± 0.03
FOC				I	29.88 ± 2.75	54.92 ± 0.03
FOC-Light			X	I	74.93 ± 0.22	73.64 ± 0.06
FOC		X		I	72.70 ± 0.36	64.78 ± 0.04
FOC		X	X	I	73.93 ± 0.29	64.84 ± 0.03
FOC	X			I	73.93 ± 0.29	64.84 ± 0.03
FOC	X	X		I	69.64 ± 1.04	66.56 ± 0.08
FOC	X	X	X	I	74.01 ± 3.17	65.17 ± 0.18
(b) plankton dataset						

4. Conclusions

In this paper, we take the first steps to address real-world underwater issues with semi-supervised learning. Our presented novel framework FOC can handle fuzzy labels via overclustering. We showed that overclustering can achieve better results than previous state-of-the-art semi-supervised methods on fuzzy plankton data. The additional overclustering output is a key difference to previous work to achieve this superior performance.

While on certain data FOC is not state-of-the-art by a clear margin of over 10%, it slightly outperforms all other methods on the fuzzy plankton data. These beneficial effects have to be verified on other fuzzy datasets and with more semi-supervised algorithms in the future. Due to better performance of FOC on fuzzy data, we expect a similar outcome. We illustrated the visual similarity on qualitative results from these predictions and results in 5 to 10% more consistent predictions. We showed that CE^{-1} is the key component for training the overclustering head.

Supplementary Materials: The following are available at <https://www.mdpi.com/article/10.3390/s21196661/s1>, The details about unsupervised clustering and a comparison to previous literature.

Author Contributions: Conceptualization, L.S., J.B., M.S., S.-M.S., R.K. (Rainer Kiko) and R.K. (Reinhard Koch); methodology, L.S., J.B., M.S. and S.-M.S.; software, L.S.; validation, L.S.; formal analysis, L.S.; investigation, L.S., J.B., M.S., S.-M.S. and R.K. (Rainer Kiko); resources, L.S. and R.K. (Rainer Kiko); data curation, L.S., S.-M.S. and R.K. (Rainer Kiko); writing—original draft preparation, L.S., J.B., M.S., S.-M.S., R.K. (Rainer Kiko) and R.K. (Reinhard Koch); writing—review and editing, L.S., J.B., M.S., S.-M.S., R.K. (Rainer Kiko) and R.K. (Reinhard Koch); visualization, L.S., J.B., M.S., S.-M.S., R.K. (Rainer Kiko) and R.K. (Reinhard Koch); supervision, R.K. (Reinhard Koch); project administration, Not applicable; funding acquisition, Not applicable. All authors have read and agreed to the published version of the manuscript.

Funding: We acknowledge funding of L. Schmarje by the ARTEMIS project (Grant number 01EC1908E) funded by the Federal Ministry of Education and Research (BMBF, Germany). We acknowledge funding of M. Santarossa by the KI-SIGS project (Grant number FKZ 01MK20012E) funded by the Federal Ministry for Economic Affairs and Energy (BMWi, Germany). S-M Schöder was supported by the “CUSCO—Coastal Upwelling System in a Changing Ocean” project (Grant number 03F0813) funded by the Federal Ministry of Education and Research (Germany). R Kiko also acknowledges support via a “Make Our Planet Great Again” grant of the French National Research Agency within the “Programme d’Investissements d’Avenir”; reference “ANR-19-MPGA-0012”. Funds to conduct the PlanktonID project were granted to R Kiko and R Koch (CP1733) by the Cluster of Excellence 80 “Future Ocean” within the framework of the Excellence Initiative by the Deutsche Forschungsgemeinschaft (DFG) on behalf of the German federal and state governments. This work was supported by Land Schleswig-Holstein through the Open Access Publikationsfonds Funding Program.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The used STL-10 dataset is introduced in [25]. The raw image plankton data is hosted on EcoTaxa [45] and the annotations were created in the project PlanktonID <https://planktonid.geomar.de/de> (accessed on 6 October 2021). The annotations can be requested from the original data owners. The source code is available at <https://github.com/Emprime/FuzzyOverclustering> (accessed on 6 October 2021). The used data is available at <https://doi.org/10.5281/zenodo.5550918> (accessed on 6 October 2021).

Acknowledgments: We thank our colleagues, especially Claudius Zelenka, for their helpful feedback and recommendations on improving the paper. Moreover, we are grateful for all citizen scientist which participated in PlanktonID and the team of PlanktonID for providing us with their data. We thank Xu Ji, Ting Chen, Kihyuk Sohn and Wouter Van Gansbeke for answering our questions regarding their respective work.

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

References

1. Saleh, A.; Laradji, I.H.; Konovalov, D.A.; Bradley, M.; Vazquez, D.; Sheaves, M. A realistic fish-habitat dataset to evaluate algorithms for underwater visual analysis. *Sci. Rep.* **2020**, *10*, 14671. [[CrossRef](#)] [[PubMed](#)]
2. Gómez-Ríos, A.; Tabik, S.; Luengo, J.; Shihavuddin, A.S.M.; Krawczyk, B.; Herrera, F. Towards highly accurate coral texture images classification using deep convolutional neural networks and data augmentation. *Expert Syst. Appl.* **2019**, *118*, 315–328. [[CrossRef](#)]
3. Thum, G.W.; Tang, S.H.; Ahmad, S.A.; Alrifayy, M. Toward a highly accurate classification of underwater cable images via deep convolutional neural network. *J. Mar. Sci. Eng.* **2020**, *8*, 924. [[CrossRef](#)]

4. Knausgård, K.M.; Wiklund, A.; Sjørdalen, T.K.; Halvorsen, K.T.; Kleiven, A.R.; Jiao, L.; Goodwin, M. Temperate fish detection and classification: A deep learning based approach. *Appl. Intell.* **2021**. [[CrossRef](#)]
5. Lombard, F.; Boss, E.; Waite, A.M.; Uitz, J.; Stemann, L.; Sosik, H.M.; Schulz, J.; Romagnan, J.B.; Picheral, M.; Pearlman, J.; et al. Globally consistent quantitative observations of planktonic ecosystems. *Front. Mar. Sci.* **2019**, *6*, 196. [[CrossRef](#)]
6. Giering, S.L.C.; Cavan, E.L.; Basedow, S.L.; Briggs, N.; Burd, A.B.; Darroch, L.J.; Guidi, L.; Irsson, J.O.; Iversen, M.H.; Kiko, R.; et al. Sinking Organic Particles in the Ocean—Flux Estimates From in situ Optical Devices. *Front. Mar. Sci.* **2020**, *6*, 834. [[CrossRef](#)]
7. Addison, P.F.E.; Collins, D.J.; Trebilco, R.; Howe, S.; Bax, N.; Hedge, P.; Jones, G.; Miloslavich, P.; Roelfsema, C.; Sams, M.; et al. A new wave of marine evidence-based management: Emerging challenges and solutions to transform monitoring, evaluating, and reporting. *ICES J. Mar. Sci.* **2018**, *75*, 941–952. [[CrossRef](#)]
8. Durden, J.M.; Bett, B.J.; Schoening, T.; Morris, K.J.; Nattkemper, T.W.; Ruhl, H.A. Comparison of image annotation data generated by multiple investigators for benthic ecology. *Mar. Ecol. Prog. Ser.* **2016**, *552*, 61–70. [[CrossRef](#)]
9. Schoening, T.; Bergmann, M.; Ontrup, J.; Taylor, J.; Dannheim, J.; Gutt, J.; Purser, A.; Nattkemper, T.W. Semi-automated image analysis for the assessment of megafaunal densities at the Arctic deep-sea observatory HAUSGARTEN. *PLoS ONE* **2012**, *7*, e38179. [[CrossRef](#)]
10. Schröder, S.M.; Kiko, R.; Koch, R. MorphoCluster: Efficient Annotation of Plankton images by Clustering. *Sensors* **2020**, *20*, 3060. [[CrossRef](#)]
11. Karimi, D.; Dou, H.; Warfield, S.K.; Gholipour, A. Deep learning with noisy labels: Exploring techniques and remedies in medical image analysis. *Med. Image Anal.* **2020**, *65*, 101759. [[CrossRef](#)] [[PubMed](#)]
12. Brünger, J.; Dippel, S.; Koch, R.; Veit, C. ‘Tailception’: Using neural networks for assessing tail lesions on pictures of pig carcasses. *Animal* **2019**, *13*, 1030–1036. [[CrossRef](#)] [[PubMed](#)]
13. Schmarje, L.; Zelenka, C.; Geisen, U.; Glüer, C.C.; Koch, R. 2D and 3D Segmentation of Uncertain Local Collagen Fiber Orientations in SHG Microscopy. In *DAGM German Conference of Pattern Recognition*; Springer: Berlin/Heidelberg, Germany, 2019; Volume 11824 LNCS, pp. 374–386. [[CrossRef](#)]
14. De Fauw, J.; Ledsam, J.R.; Romera-Paredes, B.; Nikolov, S.; Tomasev, N.; Blackwell, S.; Askham, H.; Glorot, X.; O’Donoghue, B.; Visentin, D.; et al. Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nat. Med.* **2018**, *24*, 1342–1350. [[CrossRef](#)] [[PubMed](#)]
15. Karimi, D.; Nir, G.; Fazli, L.; Black, P.C.; Goldenberg, L.; Salcudean, S.E. Deep Learning-Based Gleason Grading of Prostate Cancer From Histopathology Images—Role of Multiscale Decision Aggregation and Data Augmentation. *IEEE J. Biomed. Health Inform.* **2020**, *24*, 1413–1426. [[CrossRef](#)]
16. dos Reis, F.J.C.; Lynn, S.; Ali, H.R.; Eccles, D.; Hanby, A.; Provenzano, E.; Caldas, C.; Howat, W.J.; McDuffus, L.A.; Liu, B.; et al. Crowdsourcing the general public for large scale molecular pathology studies in cancer. *EBioMedicine* **2015**, *2*, 681–689. [[CrossRef](#)]
17. Culverhouse, P.; Williams, R.; Reguera, B.; Herry, V.; González-Gil, S. Do experts make mistakes? A comparison of human and machine identification of dinoflagellates. *Mar. Ecol. Prog. Ser.* **2003**, *247*, 17–25. [[CrossRef](#)]
18. Tarling, P.; Cantor, M.; Clapés, A.; Escalera, S. Deep learning with self-supervision and uncertainty regularization to count fish in underwater images. *arXiv* **2021**, arXiv:2104.14964.
19. Berthelot, D.; Carlini, N.; Cubuk, E.D.; Kurakin, A.; Sohn, K.; Zhang, H.; Raffel, C. ReMixMatch: Semi-Supervised Learning with Distribution Alignment and Augmentation Anchoring. *arXiv* **2019**, arXiv:1911.09785.
20. Zhai, X.; Oliver, A.; Kolesnikov, A.; Beyer, L. S4L: Self-Supervised Semi-Supervised Learning. In Proceedings of the IEEE International Conference on Computer Vision, Seoul, Korea, 27 October–2 November 2019; pp. 1476–1485.
21. Chen, T.; Kornblith, S.; Norouzi, M.; Hinton, G. A Simple Framework for Contrastive Learning of Visual Representations. *arXiv* **2020**, arXiv:2002.05709.
22. Gidaris, S.; Singh, P.; Komodakis, N. Unsupervised Representation Learning by Predicting Image Rotations. *arXiv* **2018**, arXiv:1803.07728.
23. Ji, X.; Henriques, J.F.; Vedaldi, A.; Ji, X.; Henriques, J.F.; Vedaldi, A. Invariant information clustering for unsupervised image classification and segmentation. In Proceedings of the IEEE International Conference on Computer Vision, Seoul, Korea, 27 October–2 November 2019; pp. 9865–9874.
24. Schmarje, L.; Santarossa, M.; Schroder, S.M.; Koch, R. A Survey on Semi-, Self-and Unsupervised Learning for Image Classification. *IEEE Access* **2021**, *9*, 82146–82168. [[CrossRef](#)]
25. Coates, A.; Ng, A.; Lee, H. An analysis of single-layer networks in unsupervised feature learning. In Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, Fort Lauderdale, FL, USA 11–13 April 2011; pp. 215–223.
26. Algan, G.; Ulusoy, I. Image Classification with Deep Learning in the Presence of Noisy Labels: A Survey. *Knowl.-Based Syst.* **2021**, *215*, 106771. [[CrossRef](#)]
27. Song, H.; Kim, M.; Park, D.; Lee, J. Learning from Noisy Labels with Deep Neural Networks: A Survey. *arXiv* **2020**, arXiv:1406.2080.
28. Nguyen, D.T.; Mummadi, C.K.; Ngo, T.P.N.; Nguyen, T.H.P.; Beggel, L.; Brox, T. SELF: Learning to Filter Noisy Labels with Self-Ensembling. *arXiv* **2019**, arXiv:1910.01842.
29. Laine, S.; Aila, T. Temporal ensembling for semi-supervised learning. *arXiv* **2016**, arXiv:1610.02242.
30. Li, J.; Socher, R.; Hoi, S.C.H. DivideMix: Learning with Noisy Labels as Semi-supervised Learning. *arXiv* **2020**, arXiv:2002.07394.
31. Geng, X. Label distribution learning. *IEEE Trans. Knowl. Data Eng.* **2016**, *28*, 1734–1748. [[CrossRef](#)]

32. Gao, B.B.; Xing, C.; Xie, C.W.; Wu, J.; Geng, X. Deep Label Distribution Learning With Label Ambiguity. *IEEE Trans. Image Process.* **2017**, *26*, 2825–2838. [[CrossRef](#)] [[PubMed](#)]
33. Liu, J.; Ma, Y.; Qu, F.; Zang, D. Semi-supervised Fuzzy Min–Max Neural Network for Data Classification. *Neural Process. Lett.* **2020**, *51*, 1445–1464. [[CrossRef](#)]
34. Kowsari, K.; Bari, N.; Vichr, R.; Goodarzi, F.A. FSL-BM: Fuzzy Supervised Learning with Binary Meta-Feature for Classification. In *Future of Information and Communication Conference*; Springer: Cham, Switzerland, 2018; pp. 655–670.
35. El-Zahhar, M.M.; El-Gayar, N.F. A semi-supervised learning approach for soft labeled data. In *Proceedings of the 2010 10th International Conference on Intelligent Systems Design and Applications*, Cairo, Egypt, 29 November–1 December 2010; pp. 1136–1141.
36. Liu, Y.; Liang, X.; Tong, S.; Kumada, T. Photo Shot-Type Disambiguation by Multi-Classifer Semi-Supervised Learning. In *Proceedings of the 2019 IEEE International Conference on Image Processing (ICIP)*, Taipei, Taiwan, 22–25 September 2019; pp. 2466–2470.
37. Caron, M.; Bojanowski, P.; Joulin, A.; Douze, M. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, Munich, Germany, 8–14 September 2018; pp. 132–149.
38. Sohn, K.; Berthelot, D.; Li, C.L.; Zhang, Z.; Carlini, N.; Cubuk, E.D.; Kurakin, A.; Zhang, H.; Raffel, C. FixMatch: Simplifying Semi-Supervised Learning with Consistency and Confidence. *arXiv* **2020**, arXiv:2001.07685.
39. He, K.; Zhang, X.; Ren, S.; Sun, J. Identity Mappings in Deep Residual Networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, Amsterdam, the Netherlands, 8–16 October 2016; pp. 630–645.
40. Grandvalet, Y.; Bengio, Y. Semi-supervised learning by entropy minimization. *Adv. Neural Inf. Process. Syst.* **2005**, *367*, 529–536.
41. Xie, Q.; Dai, Z.; Hovy, E.; Luong, M.T.; Le, Q.V. Unsupervised Data Augmentation for Consistency Training. *arXiv* **2019**, arXiv:1904.12848.
42. Miyato, T.; Maeda, S.I.; Koyama, M.; Ishii, S. Virtual adversarial training: A regularization method for supervised and semi-supervised learning. *IEEE Trans. Pattern Anal. Mach. Intell.* **2018**, *41*, 1979–1993.
43. Berthelot, D.; Carlini, N.; Goodfellow, I.; Papernot, N.; Oliver, A.; Raffel, C.A. Mixmatch: A holistic approach to semi-supervised learning. *arXiv* **2019**, arXiv:1905.02249.
44. Picheral, M.; Guidi, L.; Stemann, L.; Karl, D.M.; Iddaoud, G.; Gorsky, G. The Underwater Vision Profiler 5: An advanced instrument for high spatial resolution studies of particle size spectra and zooplankton. *Limnol. Oceanogr. Methods* **2010**, *8*, 462–473. [[CrossRef](#)]
45. Picheral, M.; Colin, S.; Irisson, J.O. EcoTaxa, a Tool for the Taxonomic Classification of Images. 2017. Available online: <https://ecotaxa.obs-vlfr.fr/> (accessed on 6 October 2021).
46. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. Imagenet classification with deep convolutional neural networks. *Adv. Neural Inf. Process. Syst.* **2012**, *60*, 1097–1105. [[CrossRef](#)]
47. Krizhevsky, A.; Hinton, G. Learning Multiple Layers of Features from Tiny Images. Technical Report. 2009. Available online: <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf> (accessed on 6 October 2021).
48. Van Gansbeke, W.; Vandenhende, S.; Georgoulis, S.; Proesmans, M.; Van Gool, L. Scan: Learning to classify images without labels. In *European Conference on Computer Vision*; Springer: Cham, Switzerland, 2020; pp. 268–285.
49. Tarvainen, A.; Valpola, H. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *arXiv* **2017**, arXiv:1703.01780.
50. Lee, D.H. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on Challenges in Representation Learning*; ICML: Atlanta, GA, USA, 2013; Volume 3, p. 2.