

A data-centric approach for improving ambiguous labels with combined semi-supervised classification and clustering

Lars Schmarje¹[0000-0002-6945-5957], Monty Santarossa¹[0000-0002-4159-1367],
Simon-Martin Schröder¹[0000-0002-6603-9907], Claudius
Zelenka¹[0000-0002-9902-2212], Rainer Kiko²[0000-0002-7851-9107], Jenny
Stracke³[0000-0002-9986-9720], Nina Volkmann⁴[0000-0003-2870-9954], and
Reinhard Koch¹[0000-0003-4398-1569]

¹ MIP, Computer Science, Kiel University, Germany
{las,msa,sms,czw,rk}@informatik.uni-kiel.de

² LOV, Sorbonne Université, France rainer.kiko@obs-vlfr.fr

³ ITW, University of Bonn jenny.stracke@itw.uni-bonn.de

⁴ WING & ITTN, University of Veterinary Medicine Hannover
nina.volkmann@tiho-hannover.de

Abstract. Consistently high data quality is essential for the development of novel loss functions and architectures in the field of deep learning. The existence of such data and labels is usually presumed, while acquiring high-quality datasets is still a major issue in many cases. Subjective annotations by annotators often lead to ambiguous labels in real-world datasets. We propose a data-centric approach to relabel such ambiguous labels instead of implementing the handling of this issue in a neural network. A hard classification is by definition not enough to capture the real-world ambiguity of the data. Therefore, we propose our method "Data-Centric Classification & Clustering (DC3)" which combines semi-supervised classification and clustering. It automatically estimates the ambiguity of an image and performs a classification or clustering depending on that ambiguity. DC3 is general in nature so that it can be used in addition to many Semi-Supervised Learning (SSL) algorithms. On average, our approach yields a 7.6% better F1-Score for classifications and a 7.9% lower inner distance of clusters across multiple evaluated SSL algorithms and datasets. Most importantly, we give a proof-of-concept that the classifications and clusterings from DC3 are beneficial as proposals for the manual refinement of such ambiguous labels. Overall, a combination of SSL with our method DC3 can lead to better handling of ambiguous labels during the annotation process. ⁵

Keywords: Data-Centric, Clustering, Ambiguous Labels




Fig. 1: Benefit of data-centric classification and clustering (DC3) over Semi-Supervised Learning (SSL) - Real-world datasets often suffer from intra- or inter-observer variability (IIV) during the annotation and thus no clear separation of classes is given as in common benchmark datasets. Images with a high variability between the annotations therefore have a ambiguous label. These ambiguous labels can perturb SSL approaches (see lightning bolt) and result in inconsistent predictions. Our method DC3 can be used in combination with SSL to identify ambiguous images automatically and cluster them, while classifying the rest as usual. Therefore, we avoid label ambiguity during training and generate cluster proposals which can be used to create more consistent labels.

1 Introduction

In recent years, deep learning has been successfully applied to many computer vision problems [21, 49, 11, 42, 40, 15]. The availability of large high-quality datasets was a main reason for this success, as this enabled machine learning to incorporate a wide variety of real world patterns [30]. Many novel loss functions and architectures have been proposed including options to handle imperfect data [51, 58]. This model-centric view mostly tries to deal with issues like label bias [35], label noise [26] or ambiguous labels [17] instead of improving the dataset during the annotation process. Following recent data-centric literature [4, 43, 45], we therefore investigate in this paper an approach to improve the dataset during the annotation process.

Specifically, we study the impact of ambiguous labels due to *intra- or inter-observer variability* (IIV). Such variability may arise from variability / inconsistency of annotations over time or between annotators. This issue is common when annotating data [39, 45, 26, 43, 47, 24, 44, 37, 5, 46, 16, 25, 14, 19]. The literature names different possible reasons for this variability such as low resolution [39], bad quality [22, 47], subjective interpretations of classes [25, 37] or mistakes [26, 33].

We assume that this variability can be modeled for each image with an unknown soft probability distribution $l \in [0, 1]^k$ for a classification problem with k classes. Many previous methods use a hard label instead of a soft label for training and therefore can not model this issue by definition. We call a label and its corresponding image *certain* if all annotators would agree on the

⁵ Source code is available at <https://github.com/Emprime/dc3>

classification ($l \in \{0, 1\}$) and *ambiguous* if they would disagree ($l \in (0, 1)$). In other words, ambiguous images are likely to have different annotations due to IIV while certain images do not. It is problematic that the unknown distribution l can only be estimated with expensive operations such as the acquisition of multiple annotations. Real-world example images with certain and ambiguous labels are given in Figure 3 and detailed definitions are given in subsection 2.1.

The goal of this paper is to introduce a method which provides predictions which are beneficial for improving ambiguous labels via relabeling in a downstream task. The quality of ambiguous labels and thus the performance of trained models [4, 60] can easily be improved with more annotations. However, more annotations are associated with a higher cost in the form of human working hours. Semi-Supervised Learning (SSL) can reduce these costs because it has shown great potential in reducing the amount of required labeled data to 10% or even 1% while maintaining classification performance [49, 27, 11, 61, 8]. SSL can even boost performance further [59, 40] on already large labeled datasets like ImageNet [30].

Therefore, we propose Data-Centric Classification & Clustering (DC3) which can be used in combination with many SSL algorithms to perform a combined semi-supervised classification and clustering. It simultaneously distinguishes between ambiguous and certain images, classifies the certain images and clusters visually similar ambiguous images. A graphical summary is provided in Figure 1. We will show that this approach leads to better classifications and more compact clusters across multiple semi-supervised algorithms and non-curated datasets. Furthermore, we give a proof-of-concept that these improvements lead to a greater consistency of labels based on proposals from DC3.

The key contributions of this paper are: (1) DC3 allows an SSL algorithm to predict on average a 7.6% better F1-Score for classifications and a 7.9% lower inner distance of clustering across multiple algorithms and non-curated datasets. The hyperparameters of DC3 are fixed across all algorithms and datasets which illustrates the general applicability of our method. (2) We give a proof-of-concept that these improved predictions can be used to create labels on average 2.4-fold faster and 6.74% more consistent, in comparison to the non-extended algorithms and a consensus process. This leads to higher quality data for further evaluation or model training. (3) DC3 can be used in combination with many SSL algorithms without a noticeable trade-off in terms of run-time or memory consumption, which should enable many further applications.

1.1 Related Work

Our method is mainly related to Data-Centric Machine Learning, Semi-Supervised Learning and Classification & Clustering.

Data-Centric Machine Learning aims at improving the data quality rather than improving the model alone [45, 36]. The data issues like imperfect, ambiguous or erroneous labels [4, 52, 60, 14, 24, 13, 1] are often handled in a model-centric approach by detecting errors or making the models more robust [50, 9, 26, 2]. We

want to use the predictions of our model to improve the annotation process and therefore prevent or minimize the quality issues before they need to be handled in particular.

Semi-Supervised Learning [10] is mainly developed on curated benchmark datasets [30, 12, 29] where the issue of IIV is not considered. In contrast to other SSL research [11, 61, 8, 18, 49], we are not evaluating on these curated benchmarks but work with new real-world datasets for two reasons. Firstly, curated datasets do not suffer so much from IIV because they were already cleaned. Recent research indicates that even these datasets suffer from errors in the labels which negatively impact the performance [39, 4]. Secondly, if we want to evaluate the IIV issue, we need an approximation of the variability of the label for each image e.g. in the form of multiple annotations per image. However, this information is not provided for current state-of-the-art benchmarks except for datasets like [39, 4].

Classification&Clustering was investigated in detail [41, 38, 6, 7, 54]. However, classical low dimensional approaches are difficult to extend to real-world images [41, 38, 6]. and many deep-learning methods use the clustering only as a proxy task before the actual classification [55, 23, 43] or iterate between classifications and clusters [7, 54]. The work by Smieja et al. is a rare example where classification and clustering results are generated in parallel in each training step [48]. However, we want to automatically decide which data should be classified or clustered due to their underlying ambiguity.

2 Method

Our method Data-Centric Classification & Clustering (DC3) is not an individual method but an extension for SSL algorithms such as [3, 49, 53, 32, 31]. Any image classification model can be combined with DC3 as long as it is compatible with the definition of an arbitrary SSL algorithm below.

2.1 Definitions

We assume that every image $x \in X$ has an unknown soft probability distribution $l \in [0, 1]^k$ for a classification problem with k classes. This assumption is based on two main reasons. Firstly, inconsistent annotations exist due to subjective opinions from the annotators, e.g. the grading of an illness [25]. A hard label $l \in \{0, 1\}^k$ could not model such a difference over the complete annotator population. Secondly, if we consider biological processes, images of intermediate transition stages between two classes, such as the degeneration of a living underwater organism to dead biomass exist [43].

An image and its corresponding label l are *ambiguous* if $i, j \in \{1, \dots, k\}$ exist with $i \neq j, l_i > 0$ and $l_j > 0$. Otherwise the image and its label are *certain*. The ambiguity of a label is $1 - \max_{i \in \{1, \dots, k\}} l_i$. An image might be ambiguous because




Fig. 2: Our method DC3 and an extended arbitrary SSL method – The SSL algorithm passes an image x through the network Φ and outputs a classification $p_n(x)$. We add two additional outputs: an overclustering $p_o(x)$ and a ambiguity estimation $p_a(x)$. The ambiguity estimation $p_a(x)$ is used to determine if the classification or the overclustering output is used for our method DC3. Only some labels are available for the classification output and therefore most images have to be trained completely self-supervised on all outputs.

it is actually an intermediate or uncertain combination of different classes as stated above. For this reason, ambiguous images are not just wrongly assigned images.

An SSL algorithm uses a labeled dataset X_l and an unlabeled dataset X_u for the training of a neural network Φ with $X = X_l \cup X_u$. For all images $x \in X_l$ a hard label l is available while no label information is available for $x \in X_u$. The output $p_n(x) := \Phi(x)$ is a probability distribution over the k classes.

2.2 DC3

Our method DC3 extends an arbitrary SSL algorithm. The SSL algorithm passes an image x through the network Φ and predicts a classification $p_n(x) \in [0, 1]^k$. DC3 calculates two additional outputs without a noticeable impact on training time or memory consumption: a clustering assignment $p_o(x) \in [0, 1]^{k'}$ with $k' > k$ and an ambiguity estimation $p_a(x) \in [0, 1]$. The cluster assignment partitions visually similar images in more clusters than classes exist (overclustering with $k' > k$). The ambiguity estimation is used to determine if a classification ($p_n(x)$) or an (over)clustering ($p_o(x)$) should be used as the final output. The image is predicted as certain and the classification is used if $p_a(x) < 0.5$. Otherwise, the image is estimated as ambiguous and the clustering is used as output.

A key difference to previous literature [23, 43, 7] is that we do not create an additional or only a clustering of all samples. We create SSL classifications for certain images while ambiguous images are clustered without prescribed knowledge. Moreover, it is not feasible to determine ambiguous images before this classification/clustering and thus we have no ground-truth for this decision as well. These conditions led us to formulate three goals for our development: 1. The underlying SSL training must be possible and not negatively impacted while computing an additional overclustering. 2. A degeneration to one or random cluster assignments has to be avoided as no ground-truth is available for the

clustering. 3. A balance between certain and ambiguous images is needed as the same argument (no-ground truth) applies to the ambiguity estimation $p_a(x)$. For this purpose, the network is trained by minimizing the following loss function which benefits from SSL but avoids the described degenerations.

$$L(x) = L_{SSL}(x) \cdot [1 - p_a(x)] + \lambda_{CE^{-1}} L_{CE^{-1}}(x) \cdot [1 - p_a(x)] + \lambda_a L_A(x) + \lambda_s L_S(x) \cdot p_a(x) \quad (1)$$

The first three loss terms correspond to the outputs $p_n(x)$, $p_o(x)$ and $p_a(x)$ and the three goals described above, respectively. The last term (L_S) is optional and stabilizes the training. The λ values are weights to balance the impact of each term. The first loss L_{SSL} is the loss calculated by the original SSL algorithm and is only scaled with $[1 - p_a(x)]$ to prevent the original SSL training on images the network predicts as ambiguous.

The second loss $L_{CE^{-1}}$ incentives visually homogeneous clusters of the images by pushing images from different classes into different clusters. This loss is needed to prevent a degeneration of the clustering. A similar loss was used in [43] but could only be trained on labeled data, with pretrained networks and several inefficient stabilizing methods like repeating every sample 3-5 times per batch. We generalized the formula for two input images x, x' of the same mini-batch which should not be of the same class:

$$CE^{-1}(p_o(x), p_o(x')) = - \sum_{c=1}^k p_o(x)_c \cdot \ln(1 - p_o(x')_c). \quad (2)$$

For the selection of x, x' , we use either the ground-truth label l of x if it is available or the Pseudo-Label based on the network prediction $p_n(x)$. The loss is also scaled with $[1 - p_a(x)]$ because it uses an estimate of the class for an image which could be wrong / ill-suited for ambiguous images.

The third loss L_A allows the ambiguity estimation. As stated above, the underlying distribution l is unknown and thus we do not know during training if x is ambiguous or certain. However, we can expect to know or be given a prior probability $p_A \in [0, 1]$ of the expected percentage of ambiguous images in the total dataset. We set p_A to a fixed value which balances certain and ambiguous images and the details are given in subsection 3.3. Based on this probability, we can estimate a Pseudo-Label of the ambiguity of each image in a batch during training. The loss L_A is the binary cross-entropy between the Pseudo-Label $h(x)$ and $p_a(x)$. The usage of hot-encoded Pseudo-Labels forces the network to make more confident predictions. The formulation is given below with i as the index of the image x inside the given batch, when all images inside the batch are sorted in ascending order based on p_a .

$$\begin{aligned} L_A(x) &= CE(h(x), p_a(x)) \\ &= -(1 - h(x)) \cdot \ln(p_a(x=0)) \\ &\quad - h(x) \cdot \ln(p_a(x=1)) \text{ with} \end{aligned} \quad (3)$$

$$h(x) = \begin{cases} 1 & i \leq \text{batch size} \cdot p_A \\ 0 & \text{else} \end{cases}$$

The fourth term L_S is the cross-entropy (CE) between $p_o(x)$ and $p_o(x')$ for two differently augmented versions x, x' of the same image. This loss is scaled with $p_a(x)$ and incentives that augmented versions of the same ambiguous image are in the same output cluster. We use CE because it indirectly minimizes also the entropy of $p_o(x)$ which leads to sharper predictions. Many SSL algorithms already use a differently augmented version x' of x as secondary input [3, 49, 53, 31, 23] which allows an easy computation. Otherwise, the fourth term is not calculated and treated as zero.

It is important to note that only the proposed combination of the individual parts leads to a successful training of all desired outputs. We show in section 4 that the combined clustering and classification (CC) based on $p_a(x)$ and the loss $L_{CE^{-1}}$ are the two essential parts to DC3.

3 Experiments

3.1 Datasets

That our method can be applied to many SSL algorithms across different real-world ambiguous datasets without major changes is a major advantage. While many datasets [39, 26, 43, 47, 24, 37, 5, 16, 25, 14] suffer from annotation variability, we do not know the unknown underlying distribution l to evaluate the ambiguity or any related metrics. We can approximate l with the average over multiple annotations from humans. An annotation is the hard coded guess $a = (a_1, \dots, a_k) \in \{0, 1\}^k$ of a class for an image from a human with exactly one $i' \in \{1, \dots, k\} : a_{i'} = 1$ and for all $j \in \{1, \dots, k\} \setminus \{i'\} : a_j = 0$. We assume that the approximation \hat{l} as the average of n annotations is identical to the unknown distribution l for $n \rightarrow \infty$. This leaves the issue that we need multiple annotations per image for a dataset with ambiguous labels which are often not available. However, all datasets summarized in Table 1 have multiple annotations and thus allow the approximation of \hat{l} . Nine visual examples for all datasets are given in Figure 3 and the datasets are shortly introduced below.

The *Plankton* dataset was introduced in [43]. The dataset contains 10 plankton classes and has multiple labels per image due to the help of citizen scientists. In contrast to [43], we include ambiguous images in the training and validation set and do not enforce a class balance which results in a slightly different data split as shown in Figure 3. Moreover, we processed the data by recentering the images and removing artifacts like scale bars.

The *Turkey* dataset was used in [56, 57]. The dataset contains cropped images of potential injuries of the birds which were separately annotated by three experts as not injured or injured.

The *Mice Bone* dataset is based on raw data which was published in [47]. The raw data are 3D scans from collagen fibers in mice bones. The three proposed classes are similar as well as dissimilar collagen fiber orientations and not relevant regions due to noise or background. We used the given segmentations to cut image regions from the original 2D image slices which mainly consist of one

Table 1: Overview of the used datasets – # is an abbreviation for number. The class imbalance is given as the percentage of the smallest and largest class with regard to the complete dataset. \hat{p}_A is the expected prior ambiguity probability of the dataset. n is the average of annotations per image.

Name	# classes	Input size [px]	# Images			Class Imbalance [%]		\hat{p}_A [%]	n
			Train	Val	Unlabeled	Smallest	Largest		
Plankton [43]	10	96x96	1964	2456	7860	4.16	30.37	44	24
Turkey [?]	2	96x96	1299	1542	5199	9.66	90.33	22	3
Mice Bone [47]	3	224x224	277	169	278	10.81	63.98	65	3
CIFAR-10H [39]	10	32x32	1600	2000	6400	9.88	10.16	32	51

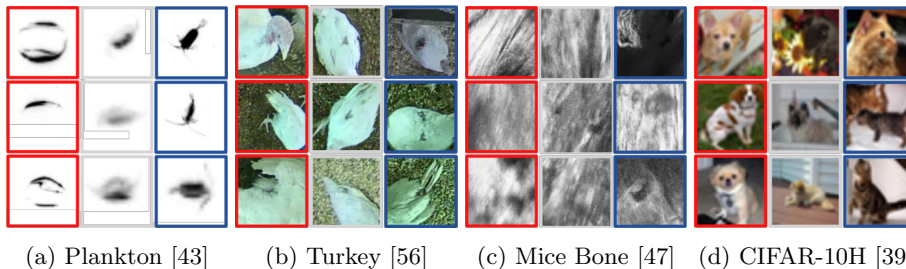


Fig. 3: Example images for the ambiguous real-world datasets – All datasets have certain images (red & blue) and ambiguous images between these classes (grey). The classes are Bubble & Copepod, Not Injured & Injured, Similar & Dissimilar Orientations and Dog & Cat respectively.

class. We generated ambiguous GT labels on 10% of the generated images by averaging over three classifications from an expert.

The *CIFAR-10H* [39] dataset provides multiple annotations for the test set of *CIFAR-10*[29]. This dataset is interesting because it illustrates that even the hard labels from benchmark datasets like *CIFAR-10* are based on soft labels due to IIV.

As stated above the approximation of \hat{l} is only possible with multiple annotations per image. For the *STL-10* dataset [12], only one annotation / label per image is given. We still include some results of this dataset to illustrate the performance on previous benchmarks.

For all datasets, we split our images X into a labeled X_l and an unlabeled X_u training set. We keep additional images as a validation subset. On X_l , we use for each image a random hard label sampled from the corresponding \hat{l} . This simulates the noisy approximation of the true ground truth label l . On X_u , we can only use the image information and not any label information. The validation data is used to compare the trained networks and to detect issues like overfitting.

3.2 Metrics

We want to measure the quality of classification and clusters over the certain and ambiguous data respectively which we assume are better proposals in the annotation or evaluation process. Based on this reasoning, we decided to use the weighted F1-Score on certain data and the mean inner distance on ambiguous data. The ambiguity is determined by the network output p_a . We define the metrics in detail below and give in subsection 3.5 a proof-of-concept for the higher consistency of labels based on proposals selected by the defined metrics. Common metrics like accuracy are not used as the class imbalance of several of our datasets would lead to misleading results.

During training we do not enforce a balance between ambiguous and certain predictions to keep the required prior knowledge minimal. This can lead to uninformative metrics and therefore we call a training *degenerated* if no more than 10% of the validation data are either predicted as ambiguous or certain. We use the *weighted F1-Score* on certain images, based on the number of images per class to avoid instability due to classes with no or very few certain (predicted) images. For the ambiguous images, we use the mean inner euclidean distance (d) to the centroid on the soft / ambiguous Ground-Truth (GT) labels. The metric d is based on the soft GT and thus also minimal for classifications of the majority class which allows an evaluation also on classified data. The equation for a set of clusters of images X is given in Equation 4 with sets $C \in X$ as clusters and the corresponding approximated soft label distribution \hat{l}_x for each image $x \in C$. The centroid per cluster is given as μ_C .

$$\begin{aligned}
 d(X) &:= \frac{1}{|X|} \sum_{C \in X} \frac{1}{|C|} \sum_{x \in C} \|\hat{l}_x - \mu_C\|_2 \text{ with} \\
 \mu_C &:= \frac{1}{|C|} \sum_{x \in C} \hat{l}_x
 \end{aligned}
 \tag{4}$$

We use the vanilla (unchanged) SSL algorithms as baseline experiments. For these experiments and some ablation experiments, we have no ambiguity prediction $p_a(x)$. In these cases, we assume all images to be certain and use $p_n(x)$ as output. We often noticed that the classification improved while the clustering degenerated and the other way round. Therefore, we determine the best performance considering the difference (d -F1) between distance and F1-Score (smaller is better). It is important to note that this balancing is arbitrary, but we give a proof of concept that the proposals calculated by these metrics lead to more consistent annotations which justifies their definition. In general, we have 3 runs per setup but we exclude results that degenerate as described above. We report the best of these runs based on the (d -F1)-score over all non-degenerated runs. All scores are calculated on the validation data which is in general about 20% of all the data (see details in Table 1).

3.3 Implementation Details

All methods use the same code base and share major hyperparameters which is crucial for valid comparisons [28]. We use the prior ambiguity $p_A = 0.6$ and loss weights $\lambda_{CE^{-1}} = 10$, $\lambda_f = 0.1$ and $\lambda_s = 0.1$ across all experiments. It is important to note that we do not use the actual prior probability of ambiguous images p_A as given in Table 1 because the probability is unknown or would require multiple annotations per image. We use a constant approximation across all datasets and show in section 4 that this approximation is comparable or even better than \hat{p}_A . This parameter is essential for balancing the certain and ambiguous images. The batch size was 64 for all datasets except for the mice bone dataset with a batch size of 8. The additional losses L_A and L_S are only applied on the unlabeled data while $L_{CE^{-1}}$ is also calculated on the labeled data. These hyperparameters were determined heuristically on the Plankton dataset with Mean-Teacher and show strong results across different methods and datasets as shown in subsection 3.4. Most likely these parameters are not optimal for an individual combination of a method and a dataset but they show the general applicability across methods and datasets. We want to show that DC3 can be applied successfully to other datasets without hyperparameter optimization and thus did not investigate all combinations in detail. Nevertheless, we refer to the supplementary for more detailed insights about individual hyperparameters and the complete pseudo code for the loss calculation.

3.4 Evaluation

The comparison between different SSL algorithms and their extension with DC3 is given in Table 2. The best results were selected as described in subsection 3.2. The complete results and additional plots are given in the supplementary. We see that DC3 improves the classification and clustering performance across the majority of classes and methods by 5 to 10%. (d -F1) is improved by up to 40% for 16 out of 19 method-dataset-combinations. On average, we achieve a 7.6% higher F1-Score for certain classifications and a 7.9% lower inner distance for clusterings of ambiguous images if we look at all non excluded method-dataset-combinations. Even on STL-10 (without the possibility to evaluate ambiguous labels) DC3 creates up to 9% better classifications. Overall, we see the most benefit on the Mice Bone and Turkey dataset which we attribute to the worse initial approximation of \hat{l} . The different vanilla algorithms achieve quite similar results for each dataset. Only FixMatch achieves a more than 5% better F1-Score on the curated STL-10 and CIFAR-10H dataset. In general, we see that DC3 can be beneficially applied to a variety of datasets and methods and predicts better classifications and more compact clusters.

Additional results about the impact of ambiguous data, the unlabeled data ratio and the interpretability can be found in the supplementary.

Table 2: Performance across different methods and datasets – The vanilla algorithm is highlighted in light grey. Better results in comparison to the vanilla algorithm are marked bold. The definition of the metrics are given in subsection 3.2. CE stands for supervised Cross-Entropy training. All values are given in %. Reasons for exclusion: H - Hardware Restrictions

Methods	Plankton			Turkey			Mice Bone			CIFAR-10H			STL-10
	F1 \uparrow	$d \downarrow$	$(d-F1) \downarrow$	F1 \uparrow	$d \downarrow$	$(d-F1) \downarrow$	F1 \uparrow	$d \downarrow$	$(d-F1) \downarrow$	F1 \uparrow	$d \downarrow$	$(d-F1) \downarrow$	F1 \uparrow
CE	86.71	30.45	-56.26	83.84	42.98	-40.86	69.55	54.75	-14.80	67.71	55.80	-11.91	80.48
CE + DC3	78.24	23.41	-54.84	85.79	27.64	-58.14	93.88	36.58	57.30	78.27	54.52	-23.75	88.45
Mean-Teacher [53]	88.72	25.84	-62.88	81.82	45.12	-36.70	66.41	48.83	-17.58	73.53	46.93	-26.59	80.67
Mean-Teacher [53] + DC3	91.30	24.84	-66.46	86.45	33.92	-52.53	89.4	35.11	-54.73	85.13	52.44	-32.69	89.28
Pi-Model [31]	87.57	28.43	-59.14	82.11	39.46	-42.65	68.15	54.11	-14.04	71.53	49.13	-22.40	82.56
Pi-Model [31] + DC3	79.79	19.08	-60.71	87.43	23.33	-64.10	88.01	30.99	-57.02	83.05	43.40	-39.65	89.54
Pseudo-Label [32]	87.62	27.42	-60.20	82.37	44.88	-37.49	66.60	57.03	-9.57	69.70	53.30	-16.40	82.48
Pseudo-Label [32] + DC3	89.31	31.76	-57.55	83.44	35.04	-48.41	86.58	37.52	-49.06	83.74	51.32	-32.42	88.87
FixMatch [49]	85.81	30.29	-55.52	82.14	43.33	-38.81	H	H	H	78.09	41.99	-36.10	89.35
FixMatch [49] + DC3	87.20	31.28	-55.92	83.56	28.17	-55.39	H	H	H	83.09	49.49	-33.60	91.45

Table 3: Consistency comparison of generated labels from proposals – The first column describes the annotator selection and the used proposals. The Cohen’s kappa coefficient κ measures the agreement of between the used repetitions and Time gives annotation time in minutes. Results which are within one percent or minute of the best result per dataset and annotator selection are marked bold.

	Plankton		Turkey		Mice Bone		CIFAR-10H	
	κ [%] \uparrow	Time [min] \downarrow	κ [%] \uparrow	Time [min] \downarrow	κ [%] \uparrow	Time [min] \downarrow	κ [%] \uparrow	Tim [min] \downarrow
A1	73.00 \pm 1.51	51.09 \pm 2.36	88.08 \pm 3.43	14.56 \pm 0.84	71.35 \pm 2.56	13.94 \pm 2.25	92.70 \pm 1.69	40.58 \pm 1.93
A1 + SSL	85.00 \pm 2.52	12.69 \pm 3.37	85.63 \pm 3.66	10.70 \pm 0.44	72.00 \pm 2.87	6.59 \pm 1.65	94.85 \pm 0.91	14.33 \pm 1.48
A1 + DC3	90.29 \pm 1.41	11.32 \pm 1.43	91.95 \pm 1.22	11.57 \pm 0.64	81.36 \pm 2.17	6.74 \pm 1.05	94.70 \pm 0.52	14.65 \pm 0.60
A2	85.25 \pm 1.79	61.99 \pm 10.98	81.54 \pm 0.89	18.11 \pm 4.30	68.63 \pm 6.66	11.06 \pm 3.60	98.81 \pm 0.14	33.08 \pm 5.36
A2 + SSL	94.88 \pm 0.52	9.23 \pm 0.76	81.10 \pm 3.39	9.48 \pm 0.83	59.63 \pm 6.20	12.07 \pm 4.77	98.00 \pm 0.27	12.66 \pm 0.69
A2 + DC3	94.04 \pm 0.67	10.32 \pm 0.07	81.83 \pm 1.98	9.91 \pm 0.39	72.19 \pm 3.23	9.13 \pm 2.98	98.29 \pm 0.19	14.27 \pm 0.69
A3	84.74 \pm 1.02	21.54 \pm 1.54	78.27 \pm 1.08	19.35 \pm 1.16	56.27 \pm 4.03	10.15 \pm 2.12	93.22 \pm 1.01	21.96 \pm 1.10
A3 + SSL	88.59 \pm 0.84	9.02 \pm 0.20	88.44 \pm 1.74	13.24 \pm 0.32	72.32 \pm 0.61	8.02 \pm 1.23	92.37 \pm 1.78	9.79 \pm 0.52
A3 + DC3	88.57 \pm 0.62	7.76 \pm 0.27	91.94 \pm 1.04	14.05 \pm 0.51	72.77 \pm 2.74	9.56 \pm 1.71	94.81 \pm 0.96	9.50 \pm 0.74

3.5 Proof-of-concept improved data quality

We show above that DC3 can lead to better classifications and clusters than SSL alone. In accordance with previous literature [45, 43], we give a proof-of-concept in Table 3 that the annotation process can be improved with cluster-based proposals. As an SSL algorithm we used Mean-Teacher and for the datasets Plankton, Turkey and CIFAR-10H we used a random subsample of 10% for the evaluation. We conducted experiments with a pool of 6 annotators which consisted of domain experts and inexperienced hired workers which were paid a fixed wage per hour. We assigned 3 annotators from the pool per dataset. This means that annotator named e.g. A1 might be a different person between datasets in Table 3. We compare the annotations over time from each annotator. We investigated three different proposals for the annotation. The baseline is not using any proposals, the second is using the SSL predictions (classification) and the third is using the DC3 predictions (classification + clusters). For each cluster, a rough description was given as guidance during the annotation. After a training

Table 4: Ablation results averaged over different methods – The vanilla algorithms / baselines are highlighted in light grey. Each lower row extends this baseline individually with CE^{-1} [43], Clustering & Classification (CC) or both (DC3). CC can be interpreted as DC3 without CE^{-1} . The prior ambiguity estimate p_A is given in brackets if applicable. Results that improve over the baseline are marked in bold. The metrics are defined in subsection 3.2. The column ‘Ambiguous’ gives the percentage of predicted ambiguous data and the last column gives the number of non-degenerated runs over which we averaged

	F1		d		$(d-F1)$		Ambiguous		# Runs
	best	mean \pm std	best	mean \pm std	best	mean \pm std	best	mean \pm std	
CIFAR-10H									
Baseline	0.7809	0.7153 \pm 0.0359	0.4199	0.5027 \pm 0.0469	-0.3611	-0.2126 \pm 0.0827	-	-	15
+ CE^{-1}	0.7383	0.7191 \pm 0.0164	0.4692	0.4929 \pm 0.0243	-0.2691	-0.2262 \pm 0.0404	-	-	12
+ CC ($p_A = 0.6$)	0.8565	0.7471 \pm 0.1246	0.8657	0.8768 \pm 0.0129	0.0092	0.1297 \pm 0.1374	0.6145	0.5923 \pm 0.0322	12
+ DC3 ($p_A = 0.32$)	0.6656	0.6970 \pm 0.0469	0.2155	0.3684 \pm 0.1227	-0.4501	-0.3286 \pm 0.0836	0.2910	0.3115 \pm 0.0140	12
+ DC3 ($p_A = 0.6$)	0.8305	0.7457 \pm 0.1097	0.4340	0.4741 \pm 0.0584	-0.3965	-0.2716 \pm 0.0928	0.6125	0.5860 \pm 0.0290	15
Plankton									
Baseline	0.8872	0.8652 \pm 0.0212	0.2584	0.2915 \pm 0.0240	-0.6287	-0.5737 \pm 0.0444	-	-	15
+ CE^{-1}	0.8896	0.8803 \pm 0.0060	0.2540	0.2690 \pm 0.0098	-0.6356	-0.6113 \pm 0.0154	-	-	12
+ CC ($p_A = 0.6$)	0.8919	0.9128 \pm 0.0427	0.4085	0.7702 \pm 0.1630	-0.4833	-0.1426 \pm 0.1375	0.6242	0.5927 \pm 0.0127	12
+ DC3 ($p_A = 0.44$)	0.8625	0.9049 \pm 0.0340	0.2192	0.3269 \pm 0.0526	-0.6433	-0.5780 \pm 0.0305	0.4365	0.4451 \pm 0.0204	11
+ DC3 ($p_A = 0.6$)	0.9130	0.8768 \pm 0.0640	0.2484	0.3004 \pm 0.0750	-0.6646	-0.5764 \pm 0.0416	0.6164	0.5893 \pm 0.0202	14
Turkey									
Baseline	0.8211	0.8213 \pm 0.00s69	0.3946	0.4428 \pm 0.0209	-0.4265	-0.3786 \pm 0.0230	-	-	15
+ CE^{-1}	0.7998	0.7998 \pm 0.0000	0.3338	0.3338 \pm 0.0000	-0.4660	-0.4660 \pm 0.0000	-	-	12
+ CC ($p_A = 0.6$)	0.8527	0.8264 \pm 0.0469	0.3400	0.3435 \pm 0.0408	-0.5127	-0.4829 \pm 0.0128	0.5837	0.5646 \pm 0.0427	12
+ DC3 ($p_A = 0.22$)	0.7998	0.7998 \pm 0.0000	0.1675	0.2252 \pm 0.0646	-0.6322	-0.5746 \pm 0.0646	0.5000	0.3674 \pm 0.2054	4
+ DC3 ($p_A = 0.6$)	0.8743	0.8432 \pm 0.0350	0.2333	0.3270 \pm 0.0692	-0.6410	-0.5162 \pm 0.0643	0.8093	0.6387 \pm 0.2354	12

phase for the inexperienced annotators, we averaged across three repetitions for every annotator, proposal and dataset combination.

We see a general trend that the consistency improves and the annotation time decreases when proposals are used instead of None. Using DC3 proposals instead of SSL proposals, either leads to a similar or better consistency while the annotation time is often increased by one or two minutes. For this improvement, we credit the cleaner and more fine-grained outputs of the network. The additional verifications of the clusters could lead to the slightly increased annotation time. The individual benefits vary between the datasets and annotators. For example, the gains on the curated CIFAR-10H dataset are lower than on the uncurated Mice Bone dataset. On average across all annotators and datasets, we achieve an improved consistency of 6.74%, a relative speed-up of 2.4 and a maximum speed-up of 4.5 with DC3 proposals in comparison to the baseline.

4 Discussion

Ablation Study We pooled the runs between all methods to evaluate the impact of the individual components of our method DC3 and show the results in Table 4. The method FixMatch and the Mice Bone dataset are excluded from this ablation due to the up to 12 times higher required GPU hours and degenerated runs as before. Across the datasets, we see the best ($d-F1$)-scores are achieved

by DC3. The impact of the components varies between the datasets. We see that CE^{-1} positively impacts the clustering results which confirms the benefit of using CE^{-1} for overclustering [43]. CC often reaches a better F1-Score than the baseline and even surpasses DC3 sometimes. However, the inner distance (d) may increase as well. We conclude that CC and CE^{-1} on their own can lead to improvements but only the combination of both parts results in a stable algorithm across datasets and methods. Additionally, we see that the number of not degenerated runs is highest with the combination of CE^{-1} and CC. If we use an realistic amount of ambiguity \hat{p}_A in each dataset as p_A , we see that in general the F1-Score decreases and d -score improves. We attribute this difference to the lower prior ambiguity p_A because DC3 tries to predict more certain than ambiguous images. This leads to a lower inner distance but also includes more difficult images in the classification of the certain data. We believe this parameter is essential for balancing the improvements in the F1- and d -score for a specified usecase. We chose a p_A of 0.6 because we wanted to weight certain and ambiguous images almost equally but ensure very certain /fewer classifications.

Qualitative Analysis with t-SNE We investigated some t-SNE [34] visualizations in Figure 4. Comparing the predicted (DC3) classes and ambiguity with the ground truth (GT), we see more wrong classifications on ambiguous images. DC3 outputs higher ambiguity than expected due to the higher value of p_A , but the predicted ambiguous clusters are often located nearby of ambiguous regions in the GT. Additionally, the clusters in (c) partition the feature space in smaller regions which can be more easily verified by humans as shown in subsection 3.5. Overall, we see a better representation of the ambiguous feature space.

Limitations We showed that DC3 generalizes to different SSL algorithms and datasets without hyperparameter changes. However, the datasets only consist of up to several thousand images. Due to the required multiple annotations per image for the evaluation it is difficult to obtain datasets with millions of images. We focused on improving the classification and clustering and gave a proof-of-concept for the increased consistency of relabeled data. Due to the required human labor during the relabeling step, we could not investigate the consistency across more datasets and algorithms or investigate the usage of the improved data. We proposed to improve the annotation process based on human-validated network predictions. This could introduce a not-desired bias into the data. This might lead to a negative impact for humans or a group of humans for certain use-cases but we believe a small bias can be accepted in most applications because it is human controlled and systematically.

5 Conclusion

In real-world datasets, we often encounter ambiguous labels, due to intra- or interobserver variability, but also as intermediate classes might exist. We propose our method DC3 which is an extension to many SSL algorithms and allows to classify images with certain labels and cluster ambiguous ones. DC3

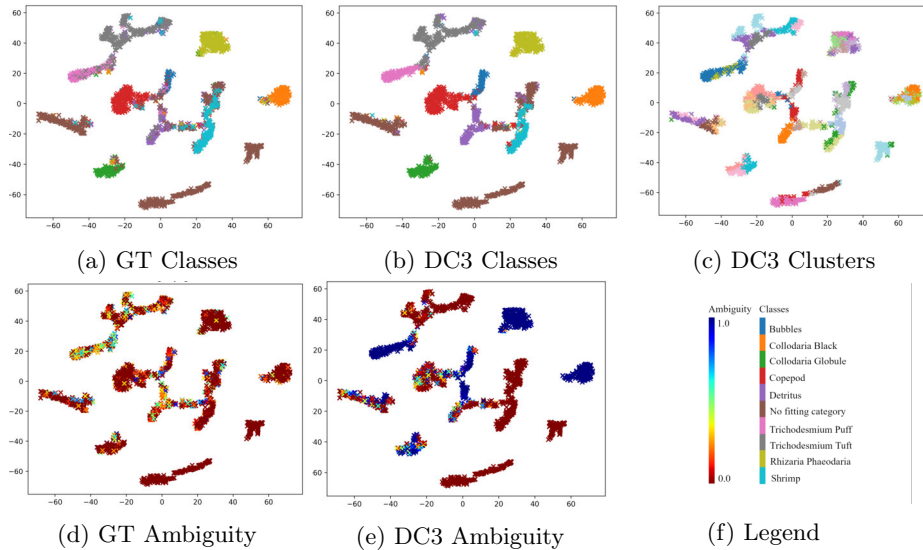


Fig. 4: t-SNE plots for Plankton dataset with Mean-Teacher – The same color was used 2–3 times for different clusters to ensure distinct colors.

also automatically determines which image should be treated as certain or ambiguous only based on a given prior probability p_A . On average, we achieve an increased F1-Score of 7.6% and a lower inner distance of clusters of 7.9% over all method-dataset-combinations. We give a proof-of-concept that these improved predictions can be used beneficially as proposals to create more consistent annotations. On average, we achieve an improved consistency of 6.74% and a relative speed-up of 2.4 when using DC3 proposals instead of no proposals. Therefore, SSL algorithms with DC3 are better suited to handle real-world datasets including ambiguous labeled images either by an improved classification / clustering or as a proposal during the annotation process with more insight.

Acknowledgements We acknowledge funding of LS by the ARTEMIS project (grant no. 01EC1908E) funded by the Federal Ministry of Education and Research (BMBF), Germany. SMS was funded by BMBF projects CUSCO (grant no. 03F0813D) and MOSAiC (grant no. 03F0917B). RKi was supported 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”. Funding for PlanktonID project were granted to RKi and RKo (CP1733) by the Cluster of Excellence 80 “Future Ocean” within the Excellence Initiative by the Deutsche Forschungsgemeinschaft on behalf of the German federal and state governments. Turkey data set was collected in the project “RedAlert – detection of pecking injuries in turkeys using neural networks” which was supported by the “Animal Welfare Innovation Award” of the “Initiative Tierwohl”.

References

1. Addison, P.F.E.E., Collins, D.J., Trebilco, R., Howe, S., Bax, N., Hedge, P., Jones, G., Miloslavich, P., Roelfsema, C., Sams, M., Stuart-Smith, R.D., Scanes, P., Von Baumgarten, P., McQuatters-Gollop, A.: A new wave of marine evidence-based management: Emerging challenges and solutions to transform monitoring, evaluating, and reporting. *ICES Journal of Marine Science* **75**(3), 941–952 (2018). <https://doi.org/10.1093/icesjms/fsx216>
2. Algan, G., Ulusoy, I.: Image Classification with Deep Learning in the Presence of Noisy Labels: A Survey. *Knowledge-Based Systems* (2020). <https://doi.org/10.1016/j.knosys.2021.106771>
3. Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., Raffel, C.A.: Mixmatch: A holistic approach to semi-supervised learning. In: *Advances in Neural Information Processing Systems*. pp. 5050–5060 (2019)
4. Beyer, L., Hénaff, O.J., Kolesnikov, A., Zhai, X., van den Oord, A.: Are we done with ImageNet? arXiv preprint arXiv:2006.07159 (2020)
5. Brünger, J., Dippel, S., Koch, R., Veit, C.: ‘Tailception’: using neural networks for assessing tail lesions on pictures of pig carcasses. *Animal* **13**(5), 1030–1036 (2019). <https://doi.org/10.1017/S1751731118003038>
6. Cai, W., Chen, S., Zhang, D.: A simultaneous learning framework for clustering and classification. *Pattern Recognition* **42**(7), 1248–1259 (2009). <https://doi.org/10.1016/j.patcog.2008.11.029>
7. 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)*. pp. 132–149 (2018)
8. Caron, M., Goyal, P., Misra, I., Bojanowski, P., Mairal, J., Joulin, A.: Unsupervised Learning of Visual Features by Contrasting Cluster Assignments. *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)* (2020)
9. Cevikalp, H., Benligiray, B., Gerek, O.N.: Semi-supervised robust deep neural networks for multi-label image classification. *Pattern Recognition* **100**, 107164 (2020). <https://doi.org/https://doi.org/10.1016/j.patcog.2019.107164>
10. Chapelle, O., Scholkopf, B., Zien, A., Schölkopf, B., Zien, A.: Semi-supervised learning. *IEEE Transactions on Neural Networks* **20**(3), 542 (2006)
11. Chen, T., Kornblith, S., Swersky, K., Norouzi, M., Hinton, G.: Big Self-Supervised Models are Strong Semi-Supervised Learners. *Advances in Neural Information Processing Systems 33 pre-proceedings (NeurIPS 2020)* (2020)
12. 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*. pp. 215–223 (2011)
13. Crawford, K., Paglen, T.: Excavating AI: The Politics of Images in Machine Learning Training Sets. *AI and Society* pp. 1–12. <https://doi.org/10.1007/s00146-021-01162-8>
14. Culverhouse, P., Williams, R., Reguera, B., Herry, V., González-Gil, S.: Do experts make mistakes? A comparison of human and machine identification of dinoflagellates. *Marine Ecology Progress Series* **247**, 17–25 (2003). <https://doi.org/10.3354/meps247017>
15. Damm, T., Schmarje, L., Koser, N., Reinhold, S., Yilmaz, E., Krekieln, N., Lui, L.Y., Cummings, S.R., Koch, R., Glueer, C.C.: Artificial intelligence-driven hip fracture prediction based on pelvic radiographs exceeds performance of DXA: the “Study of Osteoporotic Fractures” (SOF). *Journal of Bone and Mineral Research* **37**, 193–193 (2021)

16. De Fauw, J., Ledsam, J.R., Romera-Paredes, B., Nikolov, S., Tomasev, N., Blackwell, S., Askham, H., Glorot, X., O'Donoghue, B., Visentin, D., Others, O'Donoghue, B., Visentin, D., Others: Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nature medicine* **24**(9), 1342–1350 (2018)
17. Gao, B.B., Xing, C., Xie, C.W., Wu, J., Geng, X.: Deep Label Distribution Learning With Label Ambiguity. *IEEE Transactions on Image Processing* **26**(6), 2825–2838 (2017)
18. Grill, J.B., Strub, F., Althé, F., Tallec, C., Richemond, P.H., Buchatskaya, E., Dohersch, C., Pires, B.A., Guo, Z.D., Azar, M.G., Piot, B., Kavukcuoglu, K., Munos, R., Valko, M.: Bootstrap your own latent: A new approach to self-supervised Learning. *Advances in Neural Information Processing Systems 33 pre-proceedings (NeurIPS 2020)* (2020)
19. Grossmann, V., Schmarje, L., Koch, R.: Beyond Hard Labels: Investigating data label distributions. *arXiv preprint arXiv:2207.06224* (2022)
20. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: *International Conference on Machine Learning*. pp. 1321–1330. PMLR (2017)
21. He, K., Gkioxari, G., Dollár, P., Girshick, R., R-cnn, M., Doll, P., Girshick, R.: Mask r-cnn. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2961–2969 (2017)
22. Jenckel, M., Parkala, S.S., Bukhari, S.S., Dengel, A.: Impact of Training LSTM-RNN with Fuzzy Ground Truth. In: *ICPRAM* (2018)
23. 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*. pp. 9865–9874. No. Iic (2019)
24. Jungo, A., Meier, R., Ermis, E., Blatti-Moreno, M., Herrmann, E., Wiest, R., Reyes, M.: On the effect of inter-observer variability for a reliable estimation of uncertainty of medical image segmentation. In: *Medical Image Computing and Computer Assisted Interventions, MICCAI*. pp. 682–690. Springer (2018)
25. 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 Journal of Biomedical and Health Informatics* **24**(5), 1413–1426 (2020). <https://doi.org/10.1109/JBHI.2019.2944643>
26. Karimi, D., Dou, H., Warfield, S.K., Gholipour, A.: Deep learning with noisy labels: exploring techniques and remedies in medical image analysis. *Medical Image Analysis* **65** (2020)
27. Kim, B., Choo, J., Kwon, Y.D., Joe, S., Min, S., Gwon, Y.: SelfMatch: Combining Contrastive Self-Supervision and Consistency for Semi-Supervised Learning (*NeurIPS*) (2021)
28. Kolesnikov, A., Zhai, X., Beyer, L.: Revisiting self-supervised visual representation learning. In: *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*. pp. 1920–1929 (2019)
29. Krizhevsky, A., Hinton, G., Others: Learning multiple layers of features from tiny images. *Tech. rep.* (2009)
30. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *Advances in neural information processing systems*. vol. 60, pp. 1097–1105. Association for Computing Machinery (2012). <https://doi.org/10.1145/3065386>
31. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. In: *International Conference on Learning Representations* (2017)

32. 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. vol. 3, p. 2 (2013)
33. Li, J., Socher, R., Hoi, S.C.H.: DivideMix: Learning with Noisy Labels as Semi-supervised Learning. In: International Conference on Learning Representations. pp. 1–14 (2020)
34. der Maaten, L., Hinton, G.: Visualizing data using t-SNE. *Journal of machine learning research* **9**(11) (2008)
35. Menon, A.K., Rawat, A.S., Yu, F., Jayasumana, S., Kumar, S., Reddi, S., Jitkrittum, W.: Disentangling sampling and labeling bias for learning in large-output spaces. In: International Conference on Machine Learning (2021)
36. Motamedi, M., Sakharykh, N., Kaldewey, T.: A Data-Centric Approach for Training Deep Neural Networks with Less Data. *NeurIPS 2021 Data-centric AI workshop* (2021)
37. Ooms, E.A., Zonderland, H.M., Eijkemans, M.J.C., Kriege, M., Mahdavian Delavary, B., Burger, C.W., Ansink, A.C.: Mammography: Interobserver variability in breast density assessment. *The Breast* **16**(6), 568–576 (2007). <https://doi.org/10.1016/j.breast.2007.04.007>
38. Peikari, M., Salama, S., Nofech-mozes, S., Martel, A.L.: A Cluster-then-label Semi-supervised Learning Approach for Pathology Image Classification. *Scientific Reports* (April), 1–13 (2018). <https://doi.org/10.1038/s41598-018-24876-0>
39. Peterson, J., Battleday, R., Griffiths, T., Russakovsky, O.: Human uncertainty makes classification more robust. *Proceedings of the IEEE International Conference on Computer Vision 2019-October*, 9616–9625 (2019). <https://doi.org/10.1109/ICCV.2019.00971>
40. Pham, H., Dai, Z., Xie, Q., Luong, M.T., Le, Q.V.: Meta Pseudo Labels (2020)
41. Qian, Q., Chen, S., Cai, W.: Simultaneous clustering and classification over cluster structure representation. *Pattern Recognition* **45**(6), 2227–2236 (2012). <https://doi.org/10.1016/j.patcog.2011.11.027>
42. Santarossa, M., Kilic, A., von der Burchard, C., Schmarje, L., Zelenka, C., Reinhold, S., Koch, R., Roider, J.: MedRegNet: unsupervised multimodal retinal-image registration with GANs and ranking loss. In: *Medical Imaging 2022: Image Processing*. vol. 12032, pp. 321–333. SPIE (2022)
43. 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* **21**(19), 6661 (2021). <https://doi.org/10.3390/s21196661>
44. Schmarje, L., Grossmann, V., Zelenka, C., Dippel, S., Kiko, R., Oszust, M., Pastell, M., Stracke, J., Valros, A., Volkmann, N., Koch, R.: Is one annotation enough? A data-centric image classification benchmark for noisy and ambiguous label estimation. *arXiv preprint arXiv:2207.06214* (2022)
45. Schmarje, L., Koch, R.: Life is not black and white - Combining Semi-Supervised Learning with fuzzy labels. *Proceedings of the Conference "Lernen, Wissen, Daten, Analysen"*, (2021)
46. Schmarje, L., Liao, Y.H., Koch, R.: A Data-Centric Image Classification Benchmark. *NeurIPS 2021 Data-centric AI workshop* (2021)
47. Schmarje, L., Zelenka, C., Geisen, U., Glüer, C.C., Koch, R.: 2D and 3D Segmentation of uncertain local collagen fiber orientations in SHG microscopy. *DAGM German Conference of Pattern Recognition 11824 LNCS*(November), 374–386 (2019). https://doi.org/10.1007/978-3-030-33676-9_26

48. Śmieja, M., Struski, L., Figueiredo, M.A.T.: A Classification-Based Approach to Semi-Supervised Clustering with Pairwise Constraints (2020)
49. 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. *Advances in Neural Information Processing Systems* 33 pre-proceedings (NeurIPS 2020) (2020)
50. Song, H., Kim, M., Park, D., Lee, J.G., Shin, Y., Lee, J.G.: Learning From Noisy Labels With Deep Neural Networks: A Survey. *IEEE Transactions on Neural Networks and Learning Systems* pp. 1–19 (2022). <https://doi.org/10.1109/TNNLS.2022.3152527>
51. Tajbakhsh, N., Jeyaseelan, L., Li, Q., Chiang, J.N., Wu, Z., Ding, X.: Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation. *Medical Image Analysis* **63**, 101693 (2020). <https://doi.org/https://doi.org/10.1016/j.media.2020.101693>
52. Tarling, P., Cantor, M., Clapés, A., Escalera, S.: Deep learning with self-supervision and uncertainty regularization to count fish in underwater images pp. 1–22 (2021)
53. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: *ICLR* (2017)
54. Tian, Y., Henaff, O.J., van den Oord, A.: Divide and Contrast: Self-supervised Learning from Uncurated Data (2021)
55. Van Gansbeke, W., Vandenhende, S., Georgoulis, S., Proesmans, M., Van Gool, L.: Scan: Learning to classify images without labels. In: *Proceedings of the European Conference on Computer Vision*. pp. 268–285 (2020)
56. Volkmann, N., Brünger, J., Stracke, J., Zelenka, C., Koch, R., Kemper, N., Spindler, B.: Learn to train: Improving training data for a neural network to detect pecking injuries in turkeys. *Animals* 2021 **11**, 1–13 (2021). <https://doi.org/10.3390/ani11092655>
57. Volkmann, N., Zelenka, C., Devaraju, A.M., Brünger, J., Stracke, J., Spindler, B., Kemper, N., Koch, R.: Keypoint Detection for Injury Identification during Turkey Husbandry Using Neural Networks. *Sensors* **22**(14), 5188 (2022). <https://doi.org/10.3390/s22145188>
58. Wei, Y., Feng, J., Liang, X., Cheng, M.m.: Object Region Mining with Adversarial Erasing : A Simple Classification to Object Region Mining with Adversarial. *CVPR* (March), 1568–1576 (2017)
59. Xie, Q., Luong, M.T., Hovy, E., Le, Q.V., Luong, M.T., Le, Q.V., Hovy, E., Le, Q.V.: Self-Training With Noisy Student Improves ImageNet Classification. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 10684–10695. IEEE (2020). <https://doi.org/10.1109/CVPR42600.2020.01070>
60. Yun, S., Oh, S.J., Heo, B., Han, D., Choe, J., Chun, S.: Re-Labeling ImageNet: From Single to Multi-Labels, From Global to Localized Labels. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2340–2350 (2021)
61. Zbontar, J., Jing, L., Misra, I., LeCun, Y., Deny, S.: Barlow Twins: Self-Supervised Learning via Redundancy Reduction (2021)