SelfClean: A Self-Supervised Data Cleaning Strategy

Fabian Gröger University of Basel Switzerland

Philippe Gottfrois University of Basel Switzerland

Ludovic Amruthalingam Lucerne University of Applied Sciences and Arts Switzerland Simone Lionetti Lucerne University of Applied Sciences and Arts Switzerland

Alvaro Gonzalez-Jimenez University of Basel Switzerland

Labelling Consortium* University Hospital of Basel Switzerland

Matthew Groh MIT Media Lab USA Alexander A. Navarini University of Basel Switzerland Marc Pouly Lucerne University of Applied Sciences and Arts Switzerland

Abstract

Most commonly used benchmark datasets for computer vision contain irrelevant images, near duplicates, and label errors. Consequently, model performance on these benchmarks may not be an accurate estimate of generalization ability. This is a particularly acute concern in computer vision for medicine where datasets are typically small, stakes are high, and annotation processes are expensive and error-prone. In this paper, we propose SELFCLEAN, a general procedure to clean up image datasets exploiting a latent space learned with self-supervision. By relying on self-supervised learning, our approach focuses on intrinsic properties of the data and avoids annotation biases. We formulate dataset cleaning as either a set of ranking problems, where human experts can make decisions with significantly reduced effort, or a set of scoring problems, where decisions can be fully automated based on score distributions. We compare SELFCLEAN against other algorithms on common computer vision benchmarks enhanced with synthetic noise and demonstrate stateof-the-art performance on detecting irrelevant images, near duplicates, and label errors. In addition, we apply our method to multiple image datasets and confirm an improvement in evaluation reliability.

1 Introduction

In traditional machine learning (ML), data cleaning was an essential part of the development process since minor contaminations in the dataset, such as irrelevant samples, near duplicates, label issues, and missing values, significantly impacted the model's performance and robustness [1]. However,

^{*}Valerie Amann, Elisabeth Goessinger, Hazem Juratli, Beda Muehleisen, Alina Mueller and Veronika Schmidt

with the rise of deep learning (DL) and large-scale datasets, data cleaning has become less crucial as large models have shown to work relatively well even if data quality is limited [2]. Validating and cleaning large datasets is challenging, especially for high-dimensional data, such as images, where manual verification is often infeasible. Thus, a lot of research has focussed on learning from noisy data [3] rather than fixing quality issues, as the overwhelming benefits of large-scale datasets are believed to exceed the drawback of diminished control. However, the rapid growth of image collections also comes with a sharp decrease in their quality and consistency especially if collection involves some degree of automation and thus is inherently prone to various errors [4]. For many domains, the size of publicly available datasets is still one of the main limiting factors to the progress of artificial intelligence (AI). In these low-data regimes, the influence and importance of "clean" data is more critical since even fractional amounts of poor-quality samples can substantially hamper performance [5]. This is especially important in high-stakes settings such as in the medical domain, where high-quality data is needed to train robust models and validate their performance. However, many practitioners rather focus on data quantity as a key performance driver and assume a high quality collection process by default [6]. Thus, even in sensitive domains such as dermatology, many existing datasets are known to contain varying noise levels, which can substantially undermine the progress of learning algorithms [7].

The difficulties in controlling quality traded off against the desire to report comparable results has led DL practitioners to leverage benchmark datasets despite of their underlying issues. For example, ten of the most used benchmark datasets in ML were evaluated and found to have an average label error rate of 3.4% in the evaluation set [8]. Errors in the evaluation and training sets undermine the framework by which scientific progress is measured. Contamination in evaluation sets corrupts scores, making it unclear which methods can successfully handle edge cases and obscuring how close performance is to its theoretical optimum. This is especially relevant since many popular benchmarks are saturating, i.e. only saw minor relative changes in performance over the last few years [9]. Data quality issues in training sets hinder optimization and may produce suboptimal results.

The three main types of data quality issues addressed in this paper illustrate these mechanisms well. Irrelevant samples, i.e. inputs that cannot be used for the task, add noise to evaluation metrics while slowing down and confusing training. Near duplicates, i.e. multiple views of the same object, produce arbitrary re-weighting in the evaluation set and reduce variability in the training set. Most importantly, they introduce potential leaks between training and evaluation sets that can lead to overoptimistic results. Label errors, i.e. wrongly annotated samples, result in incorrect evaluation and poison the training process. In the rest of this paper, we use the term "noise" for data quality issues including these three main categories. Notwithstanding the clear benefits of correcting noisy data, applying any automatic cleaning procedure on evaluation sets can be problematic. In fact, the model under evaluation may rely on similar mechanisms to the cleaning approach, invalidating the independence requirement of a performance estimate. An evaluation that ignores known data leaks is however also incorrect, leading to a conundrum.

This paper proposes a new method called SELFCLEAN to clean image datasets based on selfsupervised learning (SSL), illustrated in figure 1. We formulate dataset cleaning as a set of ranking problems, which massively reduces the efforts for manual corrections, or alternatively as a set of scoring problems, which can be used for fully-automatic decisions based on score distributions. We apply our approach to well-known benchmark datasets in computer vision and medical imaging, and discuss implications for reliability of results across these domains. The proposed method enables practitioners to audit data collections, increase evaluation reliability, and amend the training objective to potentially improve results. Our work contributes to data-centric ML [10] and aims to restore confidence in both newly-collected and existing datasets. In summary, our main contributions are:

- A new data cleaning procedure called SELFCLEAN, based on SSL, which can be used to find irrelevant samples, near-duplicates, and label errors by relying exclusively on inductive biases and the dataset itself. The procedure additionally yields a domain-specific pre-trained model that can be used for transfer learning.
- A performance analysis of the proposed cleaning method compared to competitive baselines in synthetic scenarios and expert validation in natural datasets.
- The application of SELFCLEAN to well-known benchmarks in computer vision, medical imaging, and dermatology, and open-sourced resources to clean and correct their errors.



Figure 1: SELFCLEAN first trains a self-supervised encoder on noisy data to obtain latent space representations for dataset samples. It then detects irrelevant samples with agglomerative clustering, near duplicates based on pairwise distances, and label errors using the intra-/extra- class distance ratio.

• A practical recommendation to clean development and evaluation splits of benchmark datasets as a reasonable trade-off to obtain more accurate performance estimates.

SELFCLEAN is a holistic framework based on self-supervised learning to clean multiple dataset errors in a joint pipeline that focuses on relatively low-data regimes typical for the medical domain.

2 Related work

Data cleaning is a core component of data analytics and a topic of interest in the data management community [11]. Recently, the data-centric AI initiative also brought it back to the attention of ML researchers, resulting in the development of various data cleaning tools. For instance, Vailoppilly et al. [12] proposed an all-in-one "data cleansing" tool based on dimensionality reduction, a DL noise classifier, and a denoising model. Open-source tools for data cleaning also started to appear, including CleanLab [13] and CleanVision [14], Lightly [15], and FastDup [16]. Most data cleaning approaches require dimensionality reduction to work with high-dimensional data such as images. This includes traditional approaches such as PCA [17] or t-SNE [18], and feature extraction with deep encoders, which are usually trained on natural image databases such as ImageNet [19]. Notably, in the last few years, SSL was shown to learn more representative latent spaces compared to supervised training [20–22]. Furthermore, Cao and Wu [23] demonstrated that SSL could learn meaningful latent spaces even with small datasets, low resolution, and small architectures. Inspired by these results, we use SSL as a basis to detect the three main types of data quality issues encountered in practice: irrelevant samples, near duplicates, and label errors [11]. Since these sub-problems are typically addressed separately in the literature, we briefly review them in turn.

Classification of irrelevant samples has significant overlaps with outlier detection, which can be addressed with supervised, unsupervised, and semi-supervised learning and was initially motivated by data cleansing, i.e. enabling models to fit the data more smoothly [24]. Unsupervised approaches assume outliers to be located in low-density regions and are based on either reconstruction errors [25, 26], classification [27], or probabilistic methods [28]. Supervised detectors, which require full ground truth labels, can identify known outlier types but risk missing unknown ones [29, 30]. Semi-supervised algorithms can capitalize on partial labels and retain the ability to detect unseen outlier types. Prominent approaches include training on in-distribution data and detecting anomalies that deviate from the dataset representation learned during training [31].

Near-duplicate detection is traditionally based on a representation-match strategy [32, 33]. Most DL approaches follow a similar method, where feature vectors are extracted by a deep network and utilized as representations for content-based matching [34]. Another possible solution is to use Siamese neural networks to learn a similarity metric between samples [35]. A more recent approach was proposed for copy detection, and uses a contrastive self-supervised objective along with entropy

regularization to promote consistent separation of image descriptions [36]. However, a caveat of this method is that it relies on a threshold that should be manually adapted for each dataset [37].

The identification of label errors is generally focused on prediction-label agreement via confusion matrices and proceeds by removing samples with low recognition rate [38] or parts of the minority classes [39]. There are exceptions, such as recent approaches based on supervised contrastive learning for label error correction [40, 41]. Another prominent method is confident learning, which identifies label errors based on noisy data pruning, using probabilistic thresholds to estimate noise and ranking examples to train with confidence [42].

3 Methodology

Let $\mathcal{X} = \{(\boldsymbol{x}_i, l_i) \mid i \in \mathcal{I}\}$ be an image classification dataset with index set $\mathcal{I} = \{1, \dots, N\}$, where $\boldsymbol{x}_i \in \mathbb{R}^{h \times w \times c}$ is the *i*-th sample with height *h*, width *w* and channels *c*, and $l_i \in \{1, \dots, L\}$ the *i*-th label. The feature extractor *f* with parameters θ maps a sample \boldsymbol{x}_i to the representation $\boldsymbol{e}_i = f(\boldsymbol{x}_i; \theta) \in \mathbb{R}^D$ where *D* denotes the latent dimension. The image of \mathcal{X} in the latent space is then $\mathcal{E} = \{\boldsymbol{e}_i \mid i \in \mathcal{I}\}$. The weights θ of the feature extractor *f* are pre-trained on \mathcal{X} using self-supervised learning (SSL). The assumption is that SSL produces meaningful latent spaces by training features to be invariant to augmentations. Furthermore, it does not suffer from semantic collapse like supervised learning, which removes any information unnecessary for class assignment [20].

In this work, we consider pre-trained features from SimCLR [43] and DINO [44], which have been shown to produce meaningful latent spaces [20, 21], but in principle any SSL method can be used. SimCLR is a popular discriminative SSL approach, that uses a contrastive loss to compare different views of the same image against other randomly sampled images and their transformations. DINO instead belongs to the distillation SSL family and trains a student network to match the outputs of a teacher network on different views of the same image. SELFCLEAN requires pre-training an encoder with SSL, and accordingly also produces a pre-trained model, which can be used to initialize learning for downstream tasks after cleaning. In appendix D.1, we compare the performance of the task-specific SSL pre-trained encoder against general-purpose features.

As SimCLR has feature normalization built into the training objective, it is natural to compare points in its latent space using cosine similarity $sim(\boldsymbol{u}, \boldsymbol{v}) = \boldsymbol{u}^{\top} \boldsymbol{v}/(||\boldsymbol{u}||_2 \cdot ||\boldsymbol{v}||_2)$ and the associated distance scaled to [0, 1], $dist(\boldsymbol{u}, \boldsymbol{v}) = (1 - sim(\boldsymbol{u}, \boldsymbol{v}))/2$. For DINO we explicitly include L_2 normalization during training and inference making the latent space a unit sphere of dimension D - 1, to facilitate comparison. In appendix D.2, we present an ablation study of the L_2 normalization and investigate the influence of the distance function.

Irrelevant samples We define samples to be irrelevant when they cannot contribute to the solution of the task at hand. For example, this could be an image from a different modality or one that does not contain any object of interest. Our irrelevant sample ranking is induced by agglomerative clustering with single linkage [45] based on the idea that the later a cluster is incorporated, the more it can be considered an outlier [46]. The ranking is given by sorting the resulting dendrogram such that, at each merge, the elements of the cluster with fewer leaves appear first. When a merge occurs among sub-clusters with the same number of elements, we sort the sub-cluster created at the larger distance first, and we allow for ties when only leaves are involved. Finally, for the automatic determination of irrelevant samples, we associate each sample with its rank and also a numeric score, which takes small values for abnormal instances, compatible with the described ranking. In appendix H, we construct such a score starting from the idea that merges, which happen at very different distances or between clusters of very different sizes, should produce large numerical variations.

Near duplicates We define near duplicates as pairs of images that contain two views of the same object. In this sense, exact duplicates are a special case of near duplicates. We rank potential near duplicates by sorting each pair of distinct samples $(i, j), i \neq j$ in ascending order according to the distance between their representations in the latent space, $dist(e_i, e_j)$. Thus, the first pair in the ranking has the smallest cosine distance compared to every other pair in the dataset.

Label errors Label problems subsume, for instance, label errors or ambiguous labels, whereas in this paper we only consider the former. Label errors are samples annotated with a wrong class label. We rank potential label errors by sorting samples in ascending order according to their intra-/extra-

class distance ratio, also known as N2 score [47]. This is a ratio that, for an anchor point e_i , compares the distances to the nearest representation of a different label $m_{\neq}(e_i)$ and the distance to the nearest representation of the same label $m_{=}(e_i)$:

$$m_{=}(\boldsymbol{e}_{i}) = \min_{j \in \mathcal{I}} \left[\delta_{l_{i}l_{j}} \cdot \operatorname{dist}(\boldsymbol{e}_{i}, \boldsymbol{e}_{j}) \right], \quad m_{\neq}(\boldsymbol{e}_{i}) = \min_{j \in \mathcal{I}} \left[(1 - \delta_{l_{i}l_{j}}) \cdot \operatorname{dist}(\boldsymbol{e}_{i}, \boldsymbol{e}_{j}) \right], \tag{1}$$

$$s_{\text{label}}(\boldsymbol{e}_i) = \frac{m_{\neq}^2(\boldsymbol{e}_i)}{m_{=}^2(\boldsymbol{e}_i) + m_{\neq}^2(\boldsymbol{e}_i)},\tag{2}$$

where $\delta_{l_i l_j}$ is the indicator function that evaluates to 1 if, and only if, $l_i = l_j$. The score s_{label} is bound between [0, 1] where label errors ought to be scored lower than true labels.

In all three cases, SELFCLEAN leverages the local structure of the embedding space: cluster distances are computed only using the closest samples during agglomeration for irrelevant samples, nearduplicates are identified among sample pairs with the smallest distances, and finding label errors only considers the nearest examples of the same and a different class.

3.1 Operation modes

SELFCLEAN is deliberately built with two operation modes: human-in-the-loop and fully-automated. In the human-in-the-loop mode, samples are ranked based on the described criteria without specifying which ones should be actually corrected. The discovery of such candidates is essential as an exhaustive manual search is infeasible in practice, especially when considering pairwise relationships such as near duplicates. A human curator can then inspect issue-enriched data subsets, either confirming and correcting problems or looking for a specific rank threshold that gives the desired balance between precision and recall.

To perform automatic cleaning, specifying a fraction of data to drop or correct without a detailed inspection is suboptimal, as sensitivity to this parameter is expected to be very high. In this work, we make and empirically verify the hypothesis that (pairs of) samples can be scored with functions based on the embedding space distance, which produce a smooth distribution for clean samples, and relegate contaminated ones to significantly lower score values. Depending on the contaminated data distribution, it may then be possible to isolate problematic samples with statistical arguments based on two robust hyperparameters, namely the contamination rate guess α and the significance level q, as explained in appendix F.

4 Experimental setup

Datasets We experiment on a total of ten datasets, including three large-scale vision benchmarks (ImageNet [19], CelebA [48], Food-101N [49, 50]), two general medical datasets of X-rays (CheXpert [51]) and histopathological images (PatchCamelyon [52, 53]), and five dermatology datasets including dermatoscopy (HAM10000 [54], ISIC-2019 [54, 55]) and clinical images (Fitzpatrick17k [56], DDI [57], PAD-UFES-20 [58]). An overview of each dataset is presented in appendix B.

Training details We use a randomly initialized vision transformer (ViT) tiny with a patch size of 16×16 as encoder [59] for all experiments. Additionally, we compare against the same ViT supervised pre-trained on ImageNet. The latent representation is given by the class token output from the encoder's last layer, which has dimension 192 for a ViT tiny. The self-supervised pre-training follows the introductory papers of the respective methods [43, 44]. All SSL models were pre-trained for 100 epochs with only minor manual hyperparameter tuning to ensure proper convergence. Since we also consider small datasets, we increase the strength of augmentations to train with samples of sufficient variation. Details on hyperparameters can be found in table 9 of the appendix. We resize images to 224×224 pixels and normalize them using the mean and standard deviation of ImageNet [19]. The implementation is based on PyTorch 1.9 [60], and experiments are performed on an Nvidia DGX station, which features eight V100 GPUs, each with 32 GB of memory, 512 GB of system memory, and 40 CPU cores. Code and models for replicating our results can be found at https://github.com/***/***.

Table 1: Performance of various models on the detection of synthetic data quality issues. Evaluation is performed for each of the three considered issue types across two clean benchmark datasets, DDI and High-Quality Fitzpatrick17k (HQ-FST), augmented with two strategies for 10% synthetic contamination each (XR, CMED, MED, ARTE, LBL, and LBLC). Consult section 4 for more details.

s		DDI + 2	KR	DDI + CM	MED	HQ-FST	+ XR	HQ-FST +	CMED
mple	Method	AUROC (%)	AP (%)						
nt Sa	IN + IForest	89.5	37.0	89.6	52.3	88.1	37.2	96.3	74.8
evai	IN + HBOS	92.5	46.7	80.7	42.9	84.3	27.3	94.0	69.7
Irrel	SELFCLEAN (IN)	100.0	100.0	100.0	99.8	99.2	80.3	99.7	96.7
	SELFCLEAN (SimCLR)	90.8	32.8	77.8	31.6	95.5	48.1	81.4	26.0
	SELFCLEAN (DINO)	100.0	100.0	97.4	89.5	100.0	100.0	98.5	84.1
		DDI + M	ED	DDI + Al	RTE	HQ-FST +	MED	HQ-FST +	ARTE
cates	Method	AUROC (%)	AP (%)						
ilqu	pHashing	69.2	9.8	68.9	33.4	71.6	12.5	80.5	30.6
arD	SSIM	85.7	25.7	84.2	35.3	92.6	36.1	88.5	32.9
Ne	SELFCLEAN (IN)	96.7	6.0	94.9	29.2	97.8	27.9	93.3	27.2
	SELFCLEAN (SimCLR)	96.1	13.7	88.6	37.5	97.4	22.4	87.2	17.3
	SELFCLEAN (DINO)	99.6	63.0	99.6	65.4	99.3	61.2	95.3	38.9
		DDI + L	BL	DDI + LI	BLC	HQ-FST +	LBL	HQ-FST +	LBLC
OLS.	Method	AUROC (%)	AP (%)						
En	IN + Confident Learning	71.4	17.8	72.0	23.6	77.0	31.7	83.8	38.6
abel	FastDup	68.6	14.6	71.4	28.8	85.2	32.3	84.7	28.9
Г	SELFCLEAN (IN)	60.3	22.9	43.8	16.7	48.5	10.7	49.1	12.8
	SELFCLEAN (SimCLR)	39.6	8.0	42.8	16.0	56.3	17.7	47.2	13.2
	SELFCLEAN (DINO)	77.2	26.3	78.7	50.0	85.1	39.7	79.9	35.2

Evaluation metrics Our evaluation relies on ranking metrics as it constitutes the core of SELF-CLEAN independently of the operation mode. The proposed method and the baselines are therefore evaluated in terms of AUROC, AP, Recall@k, and Precision@k, following standard practice [61].

Synthetic experiment setup To compare SELFCLEAN against other methods, we create synthetic datasets by altering clean benchmarks for dermatology and start by considering each of the three data quality issues separately. Open-source dermatology datasets can be assumed of high quality since their size is relatively small, and curation often involves multiple experts annotating the full dataset. The two dermatology datasets, which we assume to be of high-quality, are High-Quality Fitzpatrick17k (HQ-FST) [56] and DDI [57]. The creation of synthetic datasets is inspired by typical noise present in the medical domain. For each of the three noise types, we define two distinct contamination strategies applied to both datasets, resulting in a total of twelve evaluations. These evaluations are then repeated for different percentages of contamination, i.e. 5% and 10%, mimicking real-world noise prevalence estimates [62]. For each noise category, we compare against other unsupervised methods that have performed well on the given task. A detailed description of these baselines can be found in appendix C. SELFCLEAN employs self-supervised pre-training, carried out on the contaminated dataset. Specifically, we train a separate model for every noise category, contamination level, and synthetic contamination strategy.

The first synthetic contamination strategy for irrelevant samples, *XR*, consists in adding images of a different modality, in our case lung X-rays of COVID-19 patients [63]. The second strategy for irrelevant samples, *CMED*, adds a number of samples from disparate datasets, consisting of surgical tools [64], X-ray images from arbitrary body parts [63], ImageNet samples [19], histopathological images [65], segmentation masks [66], and pictures of PowerPoint slides* [67] (see figure 12). Note that *XR* produces clustered outliers and *CMED* more isolated ones. The first contamination stategy for near-duplicates, called *MED*, adds samples from the original dataset after augmenting them with rotation, flipping, resizing, padding, and Gaussian blur (see figure 11a). The second approach for near-duplicates, *ARTE*, consists in adding samples from the original dataset after including artifacts such as watermarks, color bars, and rulers, followed by scaling and composition with other images to create a collage (see figure 11b). Detailed configurations for *MED* and *ARTE* can be found in table 10 of the appendix. For the first contamination strategy to evaluate label errors, named *LBL*, we take

^{*}Note that such contamination is common because medical datasets are often crawled from existing Power-Point slide collections.



Figure 2: Performance of the two best baselines compared to SELFCLEAN measured in terms of AP for a "mixed-contamination strategy" when varying the contamination rate. The artificial dataset is created from DDI by adding X-ray images (XR), then injecting augmented duplicates (MED), and finally changing labels at random (LBL).

a predefined percentage of samples of the original dataset and randomly change their labels, chosen uniformly from all others. The second, more challenging strategy to evaluate label errors, *LBLC*, consists of changing a predefined fraction to a label randomly extracted according to class prevalence in the original dataset.

Different contamination strategies can be applied sequentially to create a dataset with more challenging and realistic constellation of artificial data quality issues, resulting in a so-called "mixed-contamination strategy". We include such a scenario by choosing the first strategy for each of the three considered data problems. In order to consider all interactions, we start with adding irrelevant samples, proceed by creating near duplicates, and finally introduce label errors. Thus, there is the possibility, for example, that an irrelevant sample is further augmented, resulting in near-duplicates of an irrelevant sample. To preserve the meaning of the overall contamination rate C, each contamination in the sequence is added with prevalence C_S such that $(1 + C_S)^S = (1 + C)$, where S is the number of contamination steps.

Natural experiment setup To further validate the proposed approach, we evaluate it on data quality issues naturally found in benchmark datasets. To this end, we devise two different experiments. In the first experiment, we measure how well the ranking matches available metadata, e.g. if two images show the same person or if the label was obtained using gold standard tools in medical diagnosis. This experiment is, however, specific to each dataset depending on the applicability of the available metadata. Therefore, in a second experiment, we use SELFCLEAN to propose a ranking for some datasets and gather partial human annotations for validation. We collect annotations from medical experts for the medical datasets and rely on crowd workers for general image datasets. We then evaluate the proposed ranking using human annotations as described in appendix G.

5 Results

5.1 Synthetic contamination

Comparison on data quality issues Table 1 shows the results of the best two baseline methods compared to SELFCLEAN using supervised ImageNet (IN), SimCLR and DINO pre-training. Here performance is reported for the twelve synthetic datasets based on two dermatology benchmarks described in section 4 with a contamination rate of 10%; Table 4 in the appendix includes results for all baselines for both 5% and 10% contamination. SELFCLEAN with DINO pre-training performs roughly on par or better than the considered baselines for irrelevant-sample, near-duplicate, and label error detection. Some baselines for irrelevant sample detection perform better on clustered outliers and worse on isolated outliers or vice versa. In contrast, our method is able to detect both clustered and isolated ones. SELFCLEAN with supervised ImageNet features performs surprisingly well for irrelevant sample detection. On the other hand, SimCLR does not seem competitive compared to the other two pre-training strategies. This could be due to the small dataset size which may not be sufficient for SimCLR to be effective.

Influence of contamination rate Figure 2 illustrates the influence of the contamination rate for SELFCLEAN and the best two baseline models for each contamination category, i.e. IForest and HBOS for irrelevant samples, SSIM and pHASH for near duplicates, and FastDup and confident learning for label errors. Wherever appropriate we compare performance using both supervised ImageNet (IN) and self-supervised DINO features. Central value and error bars are obtained from training with three different random initializations and synthetic datasets. Note that this experiment is run on the mixed-contamination dataset. The dataset and augmentation strategies were chosen since they were challenging, as seen in the artificial evaluation from table 1, and include interaction effects. Our method outperforms most baselines for the three noise types over various contamination rates. The difference is apparent for irrelevant samples and near duplicates, where SELFCLEAN scores on average 84.0% and 19.3% higher in AP than the best baseline for the considered contamination rates. In the case of label errors, confident learning performs well for lower contamination rates, whereas FastDup works better for higher ones. SELFCLEAN is a good compromise throughout and results, on average, in 11.7% higher AP than FastDup.

5.2 Natural contamination

Comparison with metadata For some popular vision benchmarks, the literature provides annotations for label errors. For example Northcutt et al. [8] verified 5,440 samples of ImageNet's validation set and Lee et al. [50] 57,608 of Food-101N. Our label error ranking performs relatively poorly for ImageNet with AUROC = 67.3% and AP = 8.4%, and better for Food-101N where AUROC = 79.8% and AP = 47.8% indicate that verified label errors are identified well. Additionally, we evaluate near-duplicate detection against CelebA labels that indicate images of the same celebrity. Here we achieve an AUROC = 78.8% and AP = 30.9%, which is remarkably high considering that the model performs facial recognition without supervision. For medical datasets, we first check how well SELFCLEAN can find pairs of images showing the same lesion. We obtain good results for HAM10000 and ISIC-2019, with AUROC = 98.7%, AP = 28.4% and AUROC = 98.2%, AP = 26.6% respectively. On the other hand, for PAD-UFES-20 we only achieve AUROC = 71.0% and AP = 10.0%, which deserves further investigation to determine the mismatch. We also attempt to identify X-rays from the same patient within CheXpert and find some agreement with AUROC = 70.5%and AP = 7.5%, suggesting again that a case-by-case analysis should be performed. Furthermore, we investigate the correspondence between our label error ranking and the labelling method. For PAD-UFES-20, we consider gold-standard labelling and observe it does not correlate very well with our label error ranking (AUROC = 57.0%, AP = 61.3%). On the other hand, for the HAM10000 dataset, we observe that the labels obtained from follow-up, which are arguably the most likely to be correct since the considered nevi did not show any change during three follow-up visits or 1.5 years [54], are found by SELFCLEAN to be unlikely label errors (AUROC = 90.8%, AP = 86.3%). A table with results can be found in appendix E.2.

Comparison with human annotators We evaluate SELFCLEAN rankings based on human annotations across two medical and two common vision benchmarks as described in appendix G. The evaluation reveals that images ranked by SELFCLEAN as the most likely to contain data quality issues are also identified by human experts as containing data quality issues significantly more often than random images. As shown in table 8, we find statistically significant differences in nine of twelve evaluations for comparing the lowest 50 ranked images to a random selection and six of ten evaluations for comparing images ranked 1-25 to images ranked 26-50. In the second comparison, we exclude near duplicate evaluations since they only consist of data quality issues leading to no difference. Each evaluation indicates that the proposed ranking is reasonable and, to a large degree, coincides with human understanding of these three noise types. Therefore using SELFCLEAN can increase the efficiency when analyzing and fixing data quality issues.

6 Discussion

Influence of dataset cleaning To better understand the relevance of data cleaning, we examine the impact of using our method on performance estimates in table 2. We train linear and kNN classifiers based on our encoder on multiple classification benchmarks and measure the performance difference in F1 score when cleaning either the evaluation set or both splits. We use the automatic mode of SELFCLEAN with contamination rate guess $\alpha = 0.10$ and significance level q = 0.05. For most

Table 2: Influence of removing samples detected in the automatic cleaning mode with $\alpha = 0.10$ and q = 0.05 on downstream tasks. We report macro-averaged F1 scores for linear and kNN classifiers on DINO features over 100 random training/evaluation splits with 80% and 20% fractions respectively. We compute paired performance differences before and after cleaning the evaluation set, and before and after cleaning also the training set. We report the median and the intervals to the 5% (subscript) and 95% (superscript) percentiles. Additionally, we indicate significance of a paired permutation test on the difference sign with *p < 0.05, **p < 0.01, and ***p < 0.001.

	Differences Li	inear Classifier (%)	Differences kl	NN Classifier (%)
Dataset	Clean Eval	Clean Train	Clean Eval	Clean Train
DDI	$+1.0^{+11.1}_{-11.2}$	$-0.7^{+7.7}_{-10.8}$	$+1.2^{+1.9}_{-1.2}$ ***	$+0.0^{+1.7}_{-1.4}$ ***
HAM10000	$+0.1^{+3.2}_{-3.5}$	$-0.1^{+3.9}_{-3.6}$	$+0.2^{+0.5}_{-0.4}$ ***	$+0.2^{+1.3}_{-0.8}$ **
Fitzpatrick17k	$-0.6^{+2.9}_{-3.6}$ **	$+0.2^{+3.3}_{-3.9}$ *	$-4.1^{+1.2}_{-1.3}$ ***	$+0.1^{+2.0}_{-1.7}$
Food-101N	$+0.2^{+0.6}_{-0.5}$ ***	$+0.1^{+0.6}_{-0.5}$ **	$+0.1^{+0.1}_{-0.1}$ ***	$+0.1^{+0.2}_{-0.2}$ ***
ImageNet-1k	$-0.4^{+0.6}_{-0.6}$ ***	$-0.0^{+0.9}_{-0.5}$	$-0.4^{+0.1}_{-0.2}$ ***	$+0.4^{+0.3}_{-0.4}$ ***

benchmark datasets, cleaning the evaluation set and removing data leaks significantly alters scores. This is most apparent for datasets involving less curation, e.g. obtained from web crawling. Cleaning the training set has a significant positive impact for many benchmarks, indicating that noise in the training set hindered optimization. Detailed results are presented in appendix E.3.

Application to benchmark datasets We now apply the automatic mode of SELFCLEAN to wellknown image benchmark datasets and estimate the prevalence of data quality issues. The contamination rate guess $\alpha = 0.10$ and significance level q = 0.05 are conservative choices based on general noise level estimates and the study in appendix F. For highly-curated datasets involving extensive manual verification, such as DDI, PAD-UFES-20, HAM10000, CheXpert, and ImageNet-1k, we find noise levels below 1%. However, for ISIC-2019 and PatchCamelyon, we estimate 5.4% and 3.9% of near duplicates not accounted for in the metadata. Such errors, especially for highly curated data, should be addressed to regain confidence in the obtained scores. On the other hand, when considering datasets involving less manual curation, such as the Fitzpatrick17k, CelebA, and Food-101N, we find less than 1% of irrelevant samples, approximately 14.8%, 0.4%, and 1.4% near duplicates, and 0.6%, 0.5%, and 0.9% label errors, respectively. The abundance of near duplicates in these benchmarks can often be traced back to crawled data of different pages using the same illustration or thumbnail images. Since some datasets, such as Fitzpatrick17k, do not have fixed data splits considering these near duplicates, most random splits suffer from data leaks. Thus, all models evaluated on these benchmarks are optimistically biased. Detailed results are in appendix E.4.

Recommended use The conflict between resolving data quality issues and the veto against the examination of evaluation data, mentioned in the introduction, has no easy resolution. We suggest the following compromise, which avoids obvious pitfalls, as an improvement to the current situation. To refine performance estimates for a benchmark dataset, an SSL model can be trained on the training set. SELFCLEAN can then be used to clean both training and evaluation set, where for the latter the human-in-the-loop mode is required and label errors should not be considered. The number of problems found for each set separately and across them for near duplicates should be reported. This procedure further enables using the SSL pre-trained backbone as starting point for dataset tasks.

Limitations SELFCLEAN hinges on the considered dataset and inherits biases from its intrinsic composition. For example, given an image collection with a minority group that can be easily distinguished from the rest, the minority samples could be classified as irrelevant samples. Likewise, our method does not consider the dataset's context, so the candidates it provides for correction may represent desired features (e.g. images of the same patient), and suitability of samples for a specific task is not assessed. So far we only targeted three noise types and others such as blurred images or ambiguous labels will be ignored. The current formulation of SELFCLEAN does not scale well with dataset size, so additional work is required to apply it to large data collections. Finally, the SSL strategy needs to be adapted for other data modalities, which are outside the scope of this paper.

7 Conclusion and outlook

This paper proposed a novel data-cleaning strategy called SELFCLEAN, based on self-supervised learning. Our method can be used to detect irrelevant samples, near duplicates, and label errors either automatically or in collaboration with humans. We compared our proposal to state-of-the-art methods across multiple general and medical image benchmarks in both synthetic and natural contamination settings. Results showed that SELFCLEAN outperformed the considered baselines for synthetic data quality issues and demonstrated significant performance in natural settings when compared to metadata and expert annotations. Moreover, applying the cleaning strategy to highly-curated medical datasets and general vision benchmarks revealed multiple data quality issues. Removing problematic samples resulted in a significant difference for model scores. By correcting these datasets, confidence can be regained in reported benchmark performances. In the future, we plan to incorporate SELFCLEAN in annotation processes for higher data quality, and use it to detect problems during inference to enhance model robustness.

References

- Peng Li, Xi Rao, Jennifer Blase, Yue Zhang, Xu Chu, and Ce Zhang. CleanML: A Study for Evaluating the Impact of Data Cleaning on ML Classification Tasks, April 2021. URL http://arxiv.org/abs/ 1904.09483. arXiv:1904.09483 [cs].
- [2] David Rolnick, Andreas Veit, Serge Belongie, and Nir Shavit. Deep Learning is Robust to Massive Label Noise, February 2018. URL http://arxiv.org/abs/1705.10694. arXiv:1705.10694 [cs].
- [3] Nagarajan Natarajan, Inderjit S Dhillon, Pradeep K Ravikumar, and Ambuj Tewari. Learning with Noisy Labels. In Advances in Neural Information Processing Systems, volume 26. Curran Associates, Inc., 2013. URL https://proceedings.neurips.cc/paper/2013/hash/ 3871bd64012152bfb53fdf04b401193f-Abstract.html.
- [4] Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M Aroyo. "Everyone wants to do the model work, not the data work": Data Cascades in High-Stakes AI. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, pages 1–15, Yokohama Japan, May 2021. ACM. ISBN 978-1-4503-8096-6. doi: 10.1145/3411764.3445518. URL https://dl.acm.org/doi/10.1145/3411764.3445518.
- [5] Davood Karimi, Haoran Dou, Simon K. Warfield, and Ali Gholipour. Deep learning with noisy labels: exploring techniques and remedies in medical image analysis, March 2020. URL http://arxiv.org/ abs/1912.02911. arXiv:1912.02911 [cs, eess, stat].
- [6] Vasileios C. Pezoulas, Konstantina D. Kourou, Fanis Kalatzis, Themis P. Exarchos, Aliki Venetsanopoulou, Evi Zampeli, Saviana Gandolfo, Fotini Skopouli, Salvatore De Vita, Athanasios G. Tzioufas, and Dimitrios I. Fotiadis. Medical data quality assessment: On the development of an automated framework for medical data curation. *Computers in Biology and Medicine*, 107:270–283, April 2019. ISSN 0010-4825. doi: 10.1016/j.compbiomed.2019.03.001. URL https://www.sciencedirect.com/science/article/ pii/S0010482519300733.
- [7] Roxana Daneshjou, Mert Yuksekgonul, Zhuo Ran Cai, Roberto A. Novoa, and James Zou. SkinCon: A skin disease dataset densely annotated by domain experts for fine-grained debugging and analysis. January 2023. URL https://openreview.net/forum?id=gud0qopqJc4.
- [8] Curtis G. Northcutt, Anish Athalye, and Jonas Mueller. Pervasive Label Errors in Test Sets Destabilize Machine Learning Benchmarks, November 2021. URL http://arxiv.org/abs/2103.14749. arXiv:2103.14749 [cs, stat].
- [9] Simon Ott, Adriano Barbosa-Silva, Kathrin Blagec, Jan Brauner, and Matthias Samwald. Mapping global dynamics of benchmark creation and saturation in artificial intelligence. *Nature Communications*, 13 (1):6793, November 2022. ISSN 2041-1723. doi: 10.1038/s41467-022-34591-0. URL https://www.nature.com/articles/s41467-022-34591-0. Number: 1 Publisher: Nature Publishing Group.
- [10] Mohammad Hossein Jarrahi, Ali Memariani, and Shion Guha. The Principles of Data-Centric AI (DCAI), November 2022. URL http://arxiv.org/abs/2211.14611. arXiv:2211.14611 [cs].
- [11] Xu Chu, Ihab F. Ilyas, Sanjay Krishnan, and Jiannan Wang. Data Cleaning: Overview and Emerging Challenges. In *Proceedings of the 2016 International Conference on Management of Data*, pages 2201– 2206, San Francisco California USA, June 2016. ACM. ISBN 978-1-4503-3531-7. doi: 10.1145/2882903. 2912574. URL https://dl.acm.org/doi/10.1145/2882903.2912574.

- [12] Arun Prasad Vailoppilly, Ramkumar Sakthivel, and Resham Sundar Kumar. All-in-one Data Cleansing Tool. page 4, 2021.
- [13] Jonas Mueller, Anish Athalye, Caleb Chiam, Chris Mauck, and Curtis G. Northcutt et al. Cleanlab. GitHub. Note: https://github.com/cleanlab/cleanlab, 2017.
- [14] Jonas Mueller, Anish Athalye, Caleb Chiam, Chris Mauck, and Curtis G. Northcutt et al. Cleanvision. *GitHub. Note: https://github.com/cleanlab/cleanvision*, 2022.
- [15] Igor Susmelj, Matthias Heller, Philipp Wirth, Jeremy Prescott, and Malte Ebner et al. Lightly. *GitHub. Note: https://github.com/lightly-ai/lightly*, 2020.
- [16] Amir Alush, Dickson Neoh, and Danny Bickson et al. Fastdup. GitHub. Note: https://github.com/visuallayer/fastdup, 2022.
- [17] Andrzej Maćkiewicz and Waldemar Ratajczak. Principal components analysis (PCA). Computers & Geosciences, 19(3):303–342, March 1993. ISSN 0098-3004. doi: 10.1016/0098-3004(93)90090-R. URL https://www.sciencedirect.com/science/article/pii/009830049390090R.
- [18] Laurens van der Maaten and Geoffrey Hinton. Visualizing Data using t-SNE. Journal of Machine Learning Research, 9(86):2579-2605, 2008. ISSN 1533-7928. URL http://jmlr.org/papers/v9/ vandermaaten08a.html.
- [19] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. IEEE Conference on Computer Vision and Pattern Recognition, 2009.
- [20] Pierre Fernandez, Alexandre Sablayrolles, Teddy Furon, Hervé Jégou, and Matthijs Douze. Watermarking Images in Self-Supervised Latent Spaces, March 2022. URL http://arxiv.org/abs/2112.09581. arXiv:2112.09581 [cs].
- [21] Tongzhou Wang and Phillip Isola. Understanding Contrastive Representation Learning through Alignment and Uniformity on the Hypersphere. In *Proceedings of the 37th International Conference on Machine Learning*, pages 9929–9939. PMLR, November 2020. URL https://proceedings.mlr.press/v119/ wang20k.html. ISSN: 2640-3498.
- [22] Ben Sorscher, Robert Geirhos, Shashank Shekhar, Surya Ganguli, and Ari S. Morcos. Beyond neural scaling laws: beating power law scaling via data pruning, November 2022. URL http://arxiv.org/ abs/2206.14486. arXiv:2206.14486 [cs, stat].
- [23] Yun-Hao Cao and Jianxin Wu. Rethinking Self-Supervised Learning: Small is Beautiful, March 2021. URL http://arxiv.org/abs/2103.13559. arXiv:2103.13559 [cs].
- [24] Azzedine Boukerche, Lining Zheng, and Omar Alfandi. Outlier Detection: Methods, Models, and Classification. ACM Computing Surveys, 53(3):55:1–55:37, June 2020. ISSN 0360-0300. doi: 10.1145/3381028. URL https://doi.org/10.1145/3381028.
- [25] Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. Deep Anomaly Detection with Outlier Exposure, January 2019. URL http://arxiv.org/abs/1812.04606. arXiv:1812.04606 [cs, stat].
- [26] Davide Abati, Angelo Porrello, Simone Calderara, and Rita Cucchiara. Latent Space Autoregression for Novelty Detection. In 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 481–490, Long Beach, CA, USA, June 2019. IEEE. ISBN 978-1-72813-293-8. doi: 10.1109/CVPR.2019.00057. URL https://ieeexplore.ieee.org/document/8953541/.
- [27] Lukas Ruff, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander Binder, Emmanuel Müller, and Marius Kloft. Deep One-Class Classification. In Proceedings of the 35th International Conference on Machine Learning, pages 4393–4402. PMLR, July 2018. URL https: //proceedings.mlr.press/v80/ruff18a.html. ISSN: 2640-3498.
- [28] Jie Ren, Peter J. Liu, Emily Fertig, Jasper Snoek, Ryan Poplin, Mark Depristo, Joshua Dillon, and Balaji Lakshminarayanan. Likelihood Ratios for Out-of-Distribution Detection. In Advances in Neural Information Processing Systems, volume 32. Curran Associates, Inc., 2019. URL https://proceedings. neurips.cc/paper/2019/hash/1e79596878b2320cac26dd792a6c51c9-Abstract.html.
- [29] Raghavendra Chalapathy, Aditya Krishna Menon, and Sanjay Chawla. Anomaly Detection using One-Class Neural Networks, January 2019. URL http://arxiv.org/abs/1802.06360. arXiv:1802.06360 [cs, stat].

- [30] Shiyu Liang, Yixuan Li, and R. Srikant. Enhancing The Reliability of Out-of-distribution Image Detection in Neural Networks, August 2020. URL http://arxiv.org/abs/1706.02690. arXiv:1706.02690 [cs, stat].
- [31] Samet Akcay, Amir Atapour-Abarghouei, and Toby P. Breckon. GANomaly: Semi-supervised Anomaly Detection via Adversarial Training. In C. V. Jawahar, Hongdong Li, Greg Mori, and Konrad Schindler, editors, *Computer Vision – ACCV 2018*, Lecture Notes in Computer Science, pages 622–637, Cham, 2019. Springer International Publishing. ISBN 978-3-030-20893-6. doi: 10.1007/978-3-030-20893-6_39.
- [32] David G. Lowe. Distinctive Image Features from Scale-Invariant Keypoints. International Journal of Computer Vision, 60(2):91–110, November 2004. ISSN 0920-5691. doi: 10.1023/B:VISI.0000029664. 99615.94. URL http://link.springer.com/10.1023/B:VISI.0000029664.99615.94.
- [33] Yan Ke, Rahul Sukthankar, and Larry Huston. Efficient Near-duplicate Detection and Sub-image Retrieval. 2004.
- [34] Artem Babenko, Anton Slesarev, Alexandr Chigorin, and Victor Lempitsky. Neural Codes for Image Retrieval, July 2014. URL http://arxiv.org/abs/1404.1777. arXiv:1404.1777 [cs].
- [35] Jure Žbontar and Yann LeCun. Computing the Stereo Matching Cost with a Convolutional Neural Network. In 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 1592–1599, June 2015. doi: 10.1109/CVPR.2015.7298767. URL http://arxiv.org/abs/1409.4326. arXiv:1409.4326 [cs].
- [36] Ed Pizzi, Sreya Dutta Roy, Sugosh Nagavara Ravindra, Priya Goyal, and Matthijs Douze. A Self-Supervised Descriptor for Image Copy Detection. In 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 14512–14522, New Orleans, LA, USA, June 2022. IEEE. ISBN 978-1-66546-946-3. doi: 10.1109/CVPR52688.2022.01413. URL https://ieeexplore.ieee.org/document/ 9880343/.
- [37] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mahmoud Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning Robust Visual Features without Supervision, April 2023. URL http://arxiv.org/abs/2304.07193. arXiv:2304.07193 [cs].
- [38] Pengfei Chen, Benben Liao, Guangyong Chen, and Shengyu Zhang. Understanding and Utilizing Deep Neural Networks Trained with Noisy Labels, May 2019. URL http://arxiv.org/abs/1905.05040. arXiv:1905.05040 [cs, stat].
- [39] Yun Zhang, Zongze Jin, Fan Liu, Weilin Zhu, Weimin Mu, and Weiping Wang. ImageDC: Image Data Cleaning Framework Based on Deep Learning. In 2020 IEEE International Conference on Artificial Intelligence and Information Systems (ICAIIS), pages 748–752, March 2020. doi: 10.1109/ICAIIS49377. 2020.9194803.
- [40] Bin Huang, Yaohai Lin, and Chaoyang Xu. Contrastive label correction for noisy label learning. *Information Sciences*, 611:173–184, September 2022. ISSN 0020-0255. doi: 10.1016/j.ins.2022.08.060. URL https://www.sciencedirect.com/science/article/pii/S002002552200946X.
- [41] Youngjune Lee, Oh Joon Kwon, Haeju Lee, Joonyoung Kim, Kangwook Lee, and Kee-Eung Kim. Augment & Valuate : A Data Enhancement Pipeline for Data-Centric AI. 2021.
- [42] Curtis G. Northcutt, Lu Jiang, and Isaac L. Chuang. Confident Learning: Estimating Uncertainty in Dataset Labels, August 2022. URL http://arxiv.org/abs/1911.00068. arXiv:1911.00068 [cs, stat].
- [43] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A Simple Framework for Contrastive Learning of Visual Representations. In *Proceedings of the 37th International Conference on Machine Learning*. PMLR, 2020.
- [44] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging Properties in Self-Supervised Vision Transformers. 2021.
- [45] J. C. Gower and G. J. S. Ross. Minimum Spanning Trees and Single Linkage Cluster Analysis. Journal of the Royal Statistical Society. Series C (Applied Statistics), 18(1):54–64, 1969. ISSN 0035-9254. doi: 10.2307/2346439. URL https://www.jstor.org/stable/2346439. Publisher: [Wiley, Royal Statistical Society].

- [46] M.F Jiang, S.S Tseng, and C.M Su. Two-phase clustering process for outliers detection. *Pattern Recognition Letters*, 22(6-7):691–700, May 2001. ISSN 01678655. doi: 10.1016/S0167-8655(00)00131-8. URL https://linkinghub.elsevier.com/retrieve/pii/S0167865500001318.
- [47] Tin Kam Ho and M. Basu. Complexity measures of supervised classification problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(3):289–300, March 2002. ISSN 1939-3539. doi: 10.1109/34.990132. Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
- [48] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- [49] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 mining discriminative components with random forests. In *European Conference on Computer Vision*, 2014.
- [50] Kuang-Huei Lee, Xiaodong He, Lei Zhang, and Linjun Yang. CleanNet: Transfer Learning for Scalable Image Classifier Training with Label Noise, March 2018. URL http://arxiv.org/abs/1711.07131. arXiv:1711.07131 [cs].
- [51] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silviana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, Jayne Seekins, David A. Mong, Safwan S. Halabi, Jesse K. Sandberg, Ricky Jones, David B. Larson, Curtis P. Langlotz, Bhavik N. Patel, Matthew P. Lungren, and Andrew Y. Ng. CheXpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison, January 2019. URL http://arxiv.org/abs/1901.07031. arXiv:1901.07031 [cs, eess].
- [52] Bastiaan S Veeling, Jasper Linmans, Jim Winkens, Taco Cohen, and Max Welling. Rotation equivariant CNNs for digital pathology. June 2018.
- [53] Babak Ehteshami Bejnordi, Mitko Veta, Paul Johannes van Diest, Bram van Ginneken, Nico Karssemeijer, Geert Litjens, Jeroen A. W. M. van der Laak, and and the CAMELYON16 Consortium. Diagnostic Assessment of Deep Learning Algorithms for Detection of Lymph Node Metastases in Women With Breast Cancer. JAMA, 318(22):2199–2210, December 2017. ISSN 0098-7484. doi: 10.1001/jama.2017.14585. URL https://doi.org/10.1001/jama.2017.14585.
- [54] Philipp Tschandl, Cliff Rosendahl, and Harald Kittler. The HAM10000 dataset, a large collection of multisource dermatoscopic images of common pigmented skin lesions. *Scientific Data*, 5(1):180161, August 2018. ISSN 2052-4463. doi: 10.1038/sdata.2018.161. URL https://www.nature.com/articles/ sdata2018161. Number: 1 Publisher: Nature Publishing Group.
- [55] Marc Combalia, Noel C. F. Codella, Veronica Rotemberg, Brian Helba, Veronica Vilaplana, Ofer Reiter, Cristina Carrera, Alicia Barreiro, Allan C. Halpern, Susana Puig, and Josep Malvehy. BCN20000: Dermoscopic Lesions in the Wild, August 2019. URL http://arxiv.org/abs/1908.02288. arXiv:1908.02288 [cs, eess] version: 2.
- [56] Matthew Groh, Caleb Harris, Luis Soenksen, Felix Lau, Rachel Han, Aerin Kim, Arash Koochek, and Omar Badri. Evaluating Deep Neural Networks Trained on Clinical Images in Dermatology with the Fitzpatrick 17k Dataset. IEEE Computer Society, 2021.
- [57] Roxana Daneshjou, Kailas Vodrahalli, Roberto A. Novoa, Melissa Jenkins, Weixin Liang, Veronica Rotemberg, Justin Ko, Susan M. Swetter, Elizabeth E. Bailey, Olivier Gevaert, Pritam Mukherjee, Michelle Phung, Kiana Yekrang, Bradley Fong, Rachna Sahasrabudhe, Johan A. C. Allerup, Utako Okata-Karigane, James Zou, and Albert S. Chiou. Disparities in dermatology AI performance on a diverse, curated clinical image set. *Science Advances*, 8(32):eabq6147, August 2022. doi: 10.1126/sciadv.abq6147. URL https://www.science.org/doi/10.1126/sciadv.abq6147. Publisher: American Association for the Advancement of Science.
- [58] Andre G. C. Pacheco, Gustavo R. Lima, Amanda S. Salomão, Breno Krohling, Igor P. Biral, Gabriel G. de Angelo, Fábio C. R. Alves Jr, José G. M. Esgario, Alana C. Simora, Pedro B. C. Castro, Felipe B. Rodrigues, Patricia H. L. Frasson, Renato A. Krohling, Helder Knidel, Maria C. S. Santos, Rachel B. do Espírito Santo, Telma L. S. G. Macedo, Tania R. P. Canuto, and Luíz F. S. de Barros. PAD-UFES-20: A skin lesion dataset composed of patient data and clinical images collected from smartphones. *Data in Brief*, 32:106221, October 2020. ISSN 2352-3409. doi: 10.1016/j.dib.2020.106221. URL https://www.sciencedirect.com/science/article/pii/S235234092031115X.
- [59] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. arXiv:2010.11929 [cs], June 2021. URL http://arxiv.org/abs/2010.11929. arXiv: 2010.11929.

- [60] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library, 2019.
- [61] Steffen Rendle. Evaluation Metrics for Item Recommendation under Sampling, December 2019. URL http://arxiv.org/abs/1912.02263. arXiv:1912.02263 [cs].
- [62] Songqiao Han, Xiyang Hu, Hailiang Huang, Minqi Jiang, and Yue Zhao. ADBench: Anomaly Detection Benchmark. page 45, 2022.
- [63] Tulin Ozturk, Muhammed Talo, Eylul Azra Yildirim, Ulas Baran Baloglu, Ozal Yildirim, and U. Rajendra Acharya. Automated detection of COVID-19 cases using deep neural networks with X-ray images. Computers in Biology and Medicine, 121:103792, June 2020. ISSN 0010-4825. doi: 10.1016/j.compbiomed.2020.103792. URL https://www.sciencedirect.com/science/article/ pii/S0010482520301621.
- [64] Diana Martins Lavado. Sorting Surgical Tools from a Cluttered Tray Object Detection and Occlusion Reasoning. 2018.
- [65] Fabio A. Spanhol, Luiz S. Oliveira, Caroline Petitjean, and Laurent Heutte. A Dataset for Breast Cancer Histopathological Image Classification. *IEEE Transactions on Biomedical Engineering*, 63(7):1455–1462, July 2016. ISSN 1558-2531. doi: 10.1109/TBME.2015.2496264. Conference Name: IEEE Transactions on Biomedical Engineering.
- [66] Noel C. F. Codella, Veronica Rotemberg, Philipp Tschandl, M. Emre Celebi, Stephen W. Dusza, David A. Gutman, Brian Helba, Aadi Kalloo, Konstantinos Liopyris, Michael A. Marchetti, Harald Kittler, and Allan Halpern. Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (ISIC). *CoRR*, 2019.
- [67] Andre Araujo, Jason Chaves, Haricharan Lakshman, Roland Angst, and Bernd Girod. Large-Scale Queryby-Image Video Retrieval Using Bloom Filters, July 2016. URL http://arxiv.org/abs/1604.07939. arXiv:1604.07939 [cs].
- [68] Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation Forest. In 2008 Eighth IEEE International Conference on Data Mining, pages 413–422, December 2008. doi: 10.1109/ICDM.2008.17. ISSN: 2374-8486.
- [69] Markus Goldstein and Andreas Dengel. Histogram-based Outlier Score (HBOS): A fast Unsupervised Anomaly Detection Algorithm. September 2012.
- [70] Zheng Li, Yue Zhao, Xiyang Hu, Nicola Botta, Cezar Ionescu, and George H. Chen. ECOD: Unsupervised Outlier Detection Using Empirical Cumulative Distribution Functions. *IEEE Transactions on Knowledge* and Data Engineering, pages 1–1, 2022. ISSN 1041-4347, 1558-2191, 2326-3865. doi: 10.1109/TKDE. 2022.3159580. URL http://arxiv.org/abs/2201.00382. arXiv:2201.00382 [cs, stat].
- [71] D. Marr, E. Hildreth, and Sydney Brenner. Theory of edge detection. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 207(1167):187–217, January 1997. doi: 10.1098/rspb.1980.0020. URL https://royalsocietypublishing.org/doi/10.1098/rspb.1980.0020. Publisher: Royal Society.
- [72] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, April 2004. ISSN 1941-0042. doi: 10.1109/TIP.2003.819861. Conference Name: IEEE Transactions on Image Processing.
- [73] Yoav Freund and Robert E Schapire. A Short Introduction to Boosting. 1999.
- [74] Yann Dubois, Stefano Ermon, Tatsunori B. Hashimoto, and Percy S. Liang. Improving Self-Supervised Learning by Characterizing Idealized Representations. Advances in Neural Information Processing Systems, 35:11279–11296, December 2022. URL https://proceedings.neurips.cc/paper_files/paper/ 2022/hash/494f876fad056843f310ad647274dd99-Abstract-Conference.html.
- [75] Weinan E and Stephan Wojtowytsch. On the emergence of simplex symmetry in the final and penultimate layers of neural network classifiers. In *Proceedings of the 2nd Mathematical and Scientific Machine Learning Conference*, pages 270–290. PMLR, April 2022. URL https://proceedings.mlr.press/ v145/e22b.html. ISSN: 2640-3498.
- [76] Eric K. Tokuda, Cesar H. Comin, and Luciano da F. Costa. Revisiting Agglomerative Clustering, June 2020. URL http://arxiv.org/abs/2005.07995. arXiv:2005.07995 [cs, stat].

A Broader impact

SELFCLEAN is a new data-cleaning procedure that can be applied to any visual data collection. The procedure relies on SSL and therefore does not inherit annotation bias. Practitioners can choose if they want the cleaning process to happen fully automatically or with human intervention. Gaining insights into data collections of unknown quality can lead to the production of more reliable benchmarks and, in turn, results in performance measurements that are more meaningful and translate better to real-world performance. Moreover, reported benchmark results can be invalidated if they contain substantial contamination. SELFCLEAN is therefore a valuable tool to clarify which methods are the most valuable and to steer future research directions, as well as to accelerate AI applications.

The near-duplicate detection component of our approach could be used for person re-identification and data de-anonymization, even if it was not designed for this purpose. Although new in peer-reviewed publications and lacking extensive evaluation for data cleaning to the best of our knowledge, it can already be found in at least one publicly available tool [15]. We therefore conclude that increased awareness outweighs increased chances of malignant use.

For all conducted experiments we used about 5400 GPU hours which roughly corresponds to 600kg CO2.

B Datasets

In this study, we selected ten well-known open-source image datasets comprising three general-purpose vision benchmarks and seven medical datasets. The selection of datasets contains different modalities such as smartphone, high-resolution, X-ray, histopathology, dermatoscopy, and clinical images. The diversity of the datasets and domains should illustrate the proposed approach's versatility. Furthermore, some datasets were chosen because of their very high quality standards, and their curation involved large amounts of manual correction including validation by multiple domain experts. These high-quality datasets are mainly used to evaluate the proposed approach after injecting synthetic noise since they are assumed to be almost free of natural contamination.

Diverse Dermatology Images (DDI) is a public, deeply-curated, and pathologically-confirmed image dataset with diverse skin tones [57]. It contains 656 clinical images of 570 unique patients with 78 common and uncommon diseases originating from pathology reports of the Stanford Clinics.

PAD-UFES-20 is a public benchmark dataset composed of clinical images collected from smartphone devices including patient clinical data [58]. The dataset comprises 1,373 patients, 1,641 skin lesions, and 2,298 images for six different diagnoses: three skin diseases and three skin cancers.

HAM10000 is a public benchmark dataset consisting of 10,015 dermatoscopic images collected from different populations and institutions [54]. The collected cases include a representative sample of seven categories of pigmented lesions. Images were collected with the approval of the Ethics Review Committee of the University of Queensland (Protocol-No. 2017001223) and the Medical University of Vienna (Protocol-No. 1804/2017).

Fitzpatrick17k (**FST**) is a public benchmark dataset containing 16,577 clinical images with skin condition annotations and skin type labels based on the Fitzpatrick scoring system [56]. The images originate from two online dermatology atlases and thus are known to contain noise [7]. In this study, we used the middle granularity level, which partitions the labels into nine disease categories.

High-Quality Fitzpatrick17k (HQ-FST) is a subset of the Fitzpatrick17k dataset used in the paper [56] as a data quality check. It was obtained by randomly selecting 3% of the images (504 samples) and gathering annotations by two board-certified dermatologists to evaluate diagnostic accuracy. This subset is assumed to be of much higher quality than its original, larger counterpart.

ISIC-2019 is a public benchmark dataset of 25,331 dermoscopic images with metadata split into eight diagnostic categories. Additionally, the test set contains an additional outlier class not represented in the training data. The images originate from the HAM10000 [54] and the BCN_20000 [55] datasets.

PatchCamelyon consists of 327,680 color image patches extracted from histopathologic scans of lymph node sections [52] from the Camelyon16 dataset [53]. Each patch is annotated with a binary label indicating the presence of metastatic tissue. Camelyon16 contains 399 whole-slide images and corresponding glass slides of sentinel axillary lymph nodes, which were retrospectively sampled from 399 patients who underwent breast cancer surgery at two hospitals in the Netherlands. All metastases in the slides were annotated under the supervision of multiple expert pathologists.

CheXpert is a large public dataset for chest radiograph interpretation, consisting of 224,316 X-ray scans from 65,240 patients [51]. The authors retrospectively collected the chest radiographic examinations from Stanford Hospital, performed between October 2002 and July 2017 in both inpatient and outpatient centers, along with

their associated radiology reports. Labels were extracted from the free-text radiology reports with an automated rule-based system. The dataset further contains radiologist-labeled reference evaluation sets.

CelebFaces Attributes Dataset (CelebA) is a large-scale face attributes dataset with more than 200,000 celebrity images, each with 40 attribute annotations [48]. The images in this dataset cover large pose variations and mixed backgrounds. With 10,177 identities and 202,599 face images, CelebA includes many diverse samples and features rich annotations.

Food-101N is an image dataset that contains 310,009 images of food recipes divided into 101 classes [50]. Both Food-101N and the Food-101 [49] dataset share the same 101 classes. However, Food-101N has a significantly larger number of images and contains more noise. The pictures were scraped from Google, Bing, Yelp, and TripAdvisor, and 60,000 of them were manually verified for evaluation. The dataset includes information for each sample on whether or not it features a label problem.

ImageNet-1k (**IN**) is a well-known benchmark image database with 1,000 classes [19]. Images were scraped by querying words from WordNet's "synonym set" (synsets) on several image search engines. The images were labeled by Amazon Mechanical Turk workers, who were asked whether each image contained objects of a given synset.

C Baselines

We selected competitive and popular baseline approaches to detect each of the three data quality issue categories, i.e. irrelevant samples, near duplicates, and label errors. Some of these baselines require to encode images in a low-dimensional latent space. For this projection we used a ViT-tiny, the same architecture used for the proposed methodology, pre-trained on ImageNet or DINO, referred to as "IN + BASE" or "DINO + BASE" respectively, where BASE is the corresponding baseline. In this section, we briefly summarize each of the baselines used in this work.

C.1 Baselines for irrelevant samples

Isolation Forest (IForest) isolate observations by randomly selecting a feature and split value between the minimum and maximum of the selected feature, where the number of splits required to isolate a sample corresponds to the path length from the root node to the leaf node in a tree [68]. This path length averaged over a forest of random trees is a measure of normality, where noticeably shorter paths are produced for anomalies.

Histogram-based outlier detection (HBOS) is an efficient unsupervised method that creates a histogram of the feature vector for each dimension and then calculates a score based on how likely a particular data point is to fall within the histogram bins for each dimension [69]. The higher the score, the more likely the data point is an outlier, i.e. a feature vector coming from an anomaly will occupy unlikely bins in one or several of its dimensions and thus produce a higher anomaly score.

Deep One-Class Classifier with Auto Encoder (DeepSVDD) is a type of neural network for learning useful data representations in an unsupervised way [27]. DeepSVDD trains a neural network while minimizing the volume of a hypersphere that encloses the network representations of the data, forcing the network to extract the common factors of variation. DeepSVDD then detects outliers in the data by calculating the distance from the center.

Empirical Cumulative Distribution Functions (ECOD) is a parameter-free, highly-interpretable unsupervised outlier detection algorithm [70]. It estimates an empirical cumulative distribution function (ECDF) for each variable in the data separately. To generate an outlier score for an observation, it computes the tail probability for each variable using the univariate ECDFs and multiplies them together. This calculation is done in log space, accounting for each dimension's left and right tails.

C.2 Baselines for near duplicates

Perceptual Hash (pHashing) is a type of locality-sensitive hash, which is similar if features of the sample are similar [71]. It relies on the discrete cosine transform (DCT) for dimensionality reduction and produces hash bits depending on whether each DCT value is above or below the average value. In this paper, we use pHash with a hash size of 8.

Structural Similarity Index Measure (SSIM) is a type of similarity measure to compare two images with each other based on three features, namely luminance, contrast, and structure [72]. Instead of applying SSIM globally, i.e. all over the image at once, one usually applies the metrics regionally, i.e. in small sections of the image, and takes the mean overall. This variant of SSIM is often called Mean Structural Similarity Index. In this paper, we in fact apply SSIM locally to 8x8 windows but still refer to the method as SSIM for simplicity.

C.3 Baselines for label errors

Confident Learning (CL) is a data-centric approach that focuses on label quality by characterizing and identifying label errors in datasets based on the principles of pruning noisy data, counting with probabilistic thresholds to estimate noise, and ranking examples to train with confidence [42]. It builds upon the assumption of a class-conditional noise process to directly estimate the joint distribution between noisy (given) and uncorrupted (unknown) labels, resulting in a generalized CL that is provably consistent and experimentally performant. In this study, we use AdaBoost [73] as a classifier to estimate the probabilities. We did not observe any significant performance difference when using different classifiers similarly to Northcutt et al. [42].

C.4 Baselines for multiple noise types

FastDup is a powerful free and open-source tool designed to rapidly extract valuable insights from image and video datasets, aiming to increase the dataset quality and reduce data operations costs at an unparalleled scale [16]. It can detect and eliminate anomalies and outliers, identify duplicate and near-duplicate images and videos, and find wrongly-labeled samples.

D Self-supervised encoder

This section presents further investigation of the encoder pre-training and its influence on downstream tasks.

D.1 Performance of pre-trained encoder

Since SELFCLEAN relies on dataset-specific SSL to produce meaningful latent spaces, the learned features can further be used for transfer learning on the same dataset. In table 3, we consider DINO pre-training, as performed for SELFCLEAN, on different medical and vision benchmarks, i.e. DDI, HAM10000, Fitzpatrick17k, Food-101N, and ImageNet-1k. The features obtained are compared to a random and supervised ImageNet initialized ViT tiny in terms of linear performance on the respective classification task of the datasets, i.e. binary classification (benign/malignant) for DDI, multi-class classification for HAM10000, Fitzpatrick17k, Food-101N, and ImageNet-1k. The results show that ImageNet features perform better for the DDI and ImageNet-1k datasets. This is not unexpected since DDI is relatively small and the representation was obtained specifically for supervised ImageNet classification. In contrast, for HAM10000, Fitzpatrick17k, and Food-101N, the features produced by SELFCLEAN are lead to improved results. This indicates a further benefit of using SELFCLEAN, since the features used for cleaning can be recycled for the considered downstream task where they potentially improve performance.

Table 3: Macro-averaged test F1 scores for multiple vision benchmarks, obtained with linear classifiers on top of different pre-trained feature sets. Representations are extracted from a ViT tiny using random initialization weights as a baseline, transfer learning from supervised ImageNet, and DINO SSL pre-training as performed by SELFCLEAN. Results are obtained from runs with three different random seeds, of which we report the average performance and standard error.

Method	DDI (%)	HAM10000 (%)	Fitzpatrick17k (%)	Food-101N (%)	ImageNet-1k (%)
Random	$52.5~\pm~1.0$	$21.9~\pm~0.5$	$28.9~\pm~0.2$	5.1 ± 0.1	$0.8~\pm~0.1$
IN	63.6 ± 1.1	$57.5~\pm~0.4$	52.3 ± 0.1	$39.1~\pm~0.1$	49.1 ± 0.1
SELFCLEAN	$59.2~\pm~1.9$	$65.2~\pm~0.4$	53.8 ± 0.4	53.9 ± 0.2	$34.9~\pm~0.1$

D.2 Influence of *L*₂-normalization and distance functions

To compare DINO with SimCLR, we included explicit L_2 normalization in the latent space during both training and inference for DINO. A similar explicit L_2 normalization for representation layers is also used in theoretical works on SSL [74], where it was inherited from the neural collapse literature [75]. We investigate the influence of this L_2 normalization on the detection performance for the different dataset quality issues. Figure 3 shows the performance of SELFCLEAN with and without normalization for the mixed-contamination dataset constructed in 5.1. The results show that L_2 normalization increases robustness of irrelevant samples detection, improves performance in finding near duplicates, and it has little effect on label error detection. One possible explanation for the improved performance is that limiting the latent space to the unit hypersphere enforces a more direct relation between the training objective and the relative distances of encoded samples.

In a second experiment, we examined the influence of the choice of the distance function between cosine and Euclidean distance. Since the Euclidean and cosine distance on a L_2 normalized space are equivalent and produce the same ranking, we only show the results of different distance functions for the non-normalized latent



Figure 3: Performance of SELFCLEAN when changing the distance function used for comparison and removing the L_2 -normalization. The performance is measured in terms of average precision (AP) for a "mixed-contamination strategy" when varying the contamination rate. The artificial dataset is created from DDI by adding X-ray images (XR), then injecting augmented duplicates (MED), and finally changing labels at random (LBL).

space. Figure 3 shows that performance is almost unaffected by the choice of distance function, and suggests that normalization does not help by enforcing a specific distance, but rather by constraining the latent space.

D.3 Comparison of self-attention

In this experiment, we qualitatively compare on selected target datasets the last self-attention block of the encoders trained with DINO and supervised ImageNet classification. This experiment illustrates the difference between self-supervised training on the target domain and domain-agnostic pre-training. Figure 4 shows the self-attention masks for a random sample from ImageNet, CheXpert, and Fitzpatrick17k. A clear difference is visible for the self-attention masks obtained using the two pre-training strategies. For example, the masks show that encoders pre-trained with self-supervision attend to many more visual features than their supervised pre-trained counterparts, indicating that self-supervised representations contain more holistic information.

Supervised pre-training on ImageNet Mean Attention Overlay

















Top-4 Representative Patch



(a) Fitzpatrick17k [56] Supervised pre-training on ImageNet Attention Overla









Top-4 Representative Pa









Top-4 Representative Patch

(b) CheXpert [51] Supervised pre-training on ImageNet Mean Attention Overlay









DINO pre-training on ImageNet Mean Attention Overlay



(c) ImageNet [19]

Figure 4: Visualization of the last self-attention block of the encoder pre-trained with supervision on ImageNet and with DINO self-supervision on domain-specific datasets, for representative images of different datasets. The first column shows the input image, the second the focus of the model computed as the mean self-attention of the last block, the third is an overlay of the first two, and the fourth shows the four most representative patches.

E Detailed dataset cleaning results

This section provides extended tables with performance results related to dataset cleaning. More precisely, section E.1 investigates synthetic contamination detection with different methods, metrics, and noise levels, expanding on section 5.1. Section E.2 presents in tabular form the comparison of SELFCLEAN with available metadata as discussed in section 5.2. Section E.3 extends table 2 in section 6 by including information on the performances used to compute paired differences. Finally, E.4 summarizes the number of data quality issues found in the considered benchmark datasets using the fully-automatic version of SELFCLEAN.

E.1 Detailed comparison on synthetic data quality issues

Table 4: Performance of various models on the detection of synthetic data quality issues. Evaluation is performed for each of the three considered issue types across two clean benchmark datasets, DDI and High-Quality Fitzpatrick17k (HQ-FST), augmented with two strategies for synthetic contamination each (XR, CMED, MED, ARTE, LBL, and LBLC). All scores are reported in percentages (%), P@k and R@k indicate precision and recall at k respectively. Please consult section 4 for more details.

	Method							Co	ntamin	ation 5%	é						
			DDI	+ XR			DDI +	CMED			HQ-FS	ST + XR		1	HQ-FST	+ CMED	
ples		P@20	R@20	AUROC	AP	P@20	R@20	AUROC	AP	P@20	R@20	AUROC	AP	P@20	R@20	AUROC	AP
Sam	IN + IForest	30.0	18.2	92.3	26.9	55.0	36.7	91.8	53.1	15.0	11.5	81.5	12.8	70.0	53.8	96.3	68.4
/ant	IN + DeepSVDD	0.0	50.5 0.0	93.4 71.3	47.4	10.0	55.5 6.7	69.4	55.5 11.9	20.0	3.8	85.8	19.2	10.0	42.5	94.5 60.6	11.4
rele	IN + ECOD	40.0	24.2	91.5	28.4	70.0	46.7	97.9	68.2	10.0	7.7	79.7	12.1	45.0	34.6	94.3	42.9
Ц	FastDup	0.0	0.0	9.7	2.7	60.0	40.0	/4.6	45.2	0.0	0.0	7.5	2.7	85.0	65.4	99.2	86.8
	SELFCLEAN (IN)	100.0 65.0	60.6 39.4	100.0 91.8	100.0 49.0	100.0 45.0	66.7 30.0	100.0 73.3	100.0 36.8	80.0 90.0	61.5 69.2	99.2 96.0	70.6	100.0 40.0	76.9 33 3	99.9 93.8	98.8 42.4
	SELFCLEAN (DINO)	100.0	60.6	100.0	100.0	100.0	66.6	98.9	92.8	100.0	76.9	100.0	100.0	85.0	70.8	99.1	83.2
			DDI	+ MED			DDI 4	ARTE			HQ-FS	Γ + MED			HQ-FST	+ ARTE	
icates		P@20	R@20	AUROC	AP	P@20	R@20	AUROC	AP	P@20	R@20	AUROC	AP	P@20	R@20	AUROC	AP
IquC	pHashing	20.0	12.1	65.1	9.3	40.0	24.2	62.7	24.3	5.0	3.8	71.0	1.1	40.0	30.8	80.9	30.5
ear I	SSIM FastDup	45.0 45.0	27.3	85.9 70.8	22.1 16.9	40.0 15.0	24.2 9.1	53.7	25.4 7.6	35.0 15.0	26.9	87.0 56.1	24.7	40.0 10.0	30.8 7.7	81.2 44.6	32.9 3.4
Z	SELFCLEAN (IN)	20.0	12.1	94.8	7.0	35.0	21.2	91.9	21.7	30.0	23.1	99.1	18.2	40.0	30.8	96.1	29.1
	SELFCLEAN (SimCLR)	35.0	21.2	98.5	21.6	30.0	18.2	95.5	18.1	35.0	26.9	95.7	18.5	15.0	11.5	93.6 96 7	9.1 44 1
	SELFCLEAN (DINO)	00.0		+ I RI	55.7	95.0	DDL	IBLC	/1.0	50.0	HOFE	T+IRI	56.2	50.0	HOFST	+1 BLC	44.1
S		P@20	R@20	AUROC	AP	P@20	R@20	AUROC	AP	P@20	R@20	AUROC	AP	P@20	R@20	AUROC	AP
Erro	IN + Confident Learning	15.0	0.1	74.6	16.0	10.0	61	75.2	12.5	30.0	24.0	76.5	18.0	20.0	12.0	60.8	16.3
abel]	FastDup	15.0	9.1	77.9	12.9	35.0	7.5	73.8	26.3	15.0	11.5	82.8	17.3	25.0	16.7	86.2	22.6
Ц	SELFCLEAN (IN)	10.0	6.1	60.4	13.3	20.0	4.3	44.8	13.9	20.0	15.4	53.8	13.5	15.0	9.7	56.6	8.9
	SELFCLEAN (SimCLR)	10.0	6.1	50.9	6.9	20.0	4.3	42.8	14.0	25.0	19.2	48.4	13.2	15.0	10.0	46.7	10.5
	SELFCLEAN (DINO)	15.0	9.1	12.2	13.0	75.0	10.5	/5.4	41.0	30.0	23.1	00.0	20.4	40.0	20.7	00.5	33.4
								Cor	itamina	ation 10%	//0						
			DDI	+ XR			DDI +	CMED			HQ-FS	ST + XR			HQ-FST	+ CMED	
nples		P@20	R@20	AUROC	AP	P@20	R@20	AUROC	AP	P@20	R@20	AUROC	AP	P@20	R@20	AUROC	AP
Sar	IN + IForest IN + HBOS	50.0 55.0	15.4 16.9	89.5 92.5	37.0 46.7	65.0 65.0	21.7	89.6 80.7	52.3 42.9	55.0 25.0	21.6 9.8	88.1 84.3	37.2 27.3	85.0 90.0	33.3 35.3	96.3 94.0	74.8 69.7
vant	IN + DeepSVDD	10.0	3.1	52.4	9.8	20.0	6.7	75.9	21.0	5.0	2.0	51.3	9.0	20.0	7.8	66.5	23.1
rrele	IN + ECOD FastDup	25.0	7.7	85.1	27.1	75.0	25.0	96.9 43.8	68.5 26.6	0.0	0.0	71.9	15.6	55.0	21.6	93.8 95.3	53.1
Ι	Ser ECLE (N)	100.0	20.0	100.0	4.2	100.0	20.0	100.0	20.0	80.0	21.4	4.5	4.7	100.0	20.2	95.5	06.7
	SELFCLEAN (IIN) SELFCLEAN (SimCLR)	5.0	1.5	90.8	32.8	45.0	33.3 15.0	77.8	31.6	80.0 0.0	0.0	99.2 95.5	48.1	30.0	12.5	81.4	26.0
	SELFCLEAN (DINO)	100.0	30.8	100.0	100.0	100.0	33.3	97.4	89.5	100.0	39.2	100.0	100.0	95.0	39.6	98.5	84.1
s			DDI	+ MED			DDI 4	ARTE			HQ-FS	Γ + MED			HQ-FST	+ ARTE	
licate		P@20	R@20	AUROC	AP	P@20	R@20	AUROC	AP	P@20	R@20	AUROC	AP	P@20	R@20	AUROC	AP
Dup	pHashing	30.0	9.2	69.2	9.8	95.0	29.2	68.9	33.4	35.0	13.7	71.6	12.5	75.0	29.4	80.5	30.6
lear]	FastDup	70.0 40.0	12.3	85.7 58.6	25.7 10.0	40.0	30.8 12.3	84.2 59.6	35.3 5.6	65.0 25.0	25.5 9.8	92.6 56.1	7.2	80.0 35.0	31.4 13.7	88.5 52.9	32.9 8.9
2	SELFCLEAN (IN)	15.0	4.6	96.7	6.0	75.0	23.1	94.9	29.2	60.0	23.5	97.8	27.9	60.0	23.5	93.3	27.2
					127	95.0	29.2	88.6	37.5	55.0	21.6	97.4 99.3	22.4	75.0	18.2	87.2	17.3 38.9
	SELFCLEAN (SimCLR) SELFCLEAN (DINO)	40.0 95.0	12.3 29.3	96.1 99.6	63.0	100.0	31.0	99.6	65.4	95.0	37.3	,,,,,	61.2	00.0	31.4	15.5	
	SELFCLEAN (SimCLR) SELFCLEAN (DINO)	40.0 95.0	12.3 29.3 DDI	96.1 99.6 + LBL	63.0	100.0	31.0 DDI 4	99.6 • LBLC	65.4	95.0	37.5 HQ-FS	T + LBL	61.2	00.0	51.4 HQ-FST	+ LBLC	
ors	SELFCLEAN (SimCLR) SELFCLEAN (DINO)	40.0 95.0 P@20	12.3 29.3 DDI R@20	96.1 99.6 + LBL AUROC	63.0 AP	100.0 P@20	31.0 DDI + R@20	99.6 • LBLC AUROC	65.4	95.0 P@20	37.3 HQ-FS R@20	T + LBL AUROC	61.2 AP	P@20	31.4 HQ-FST R@20	+ LBLC AUROC	AP
Errors	SELFCLEAN (SimCLR) SELFCLEAN (DINO) IN + Confident Learning	40.0 95.0 P@20 25.0	12.3 29.3 DDI R@20 7.9	96.1 99.6 + LBL AUROC 71.4	AP	100.0 P@20 30.0	31.0 DDI + R@20 9.7	99.6 • LBLC AUROC 72.0	65.4 AP 23.6	95.0 P@20 35.0	37.3 HQ-FS R@20 15.6	T + LBL AUROC 77.0	AP 31.7	P@20 40.0	31.4 HQ-FST R@20 15.1	+ LBLC AUROC 83.8	AP 38.6
Label Errors	SELFCLEAN (SimCLR) SELFCLEAN (DINO) IN + Confident Learning FastDup	40.0 95.0 P@20 25.0 15.0	12.3 29.3 DDI R@20 7.9 4.7	96.1 99.6 + LBL AUROC 71.4 68.6	AP 17.8 14.6	100.0 P@20 30.0 25.0	31.0 DDI 4 R@20 9.7 4.4	99.6 • LBLC AUROC 72.0 71.4	65.4 AP 23.6 28.8	95.0 P@20 35.0 40.0	37.3 HQ-FS R@20 15.6 16.0	T + LBL AUROC 77.0 85.2	AP 31.7 32.3	P@20 40.0 30.0	31.4 HQ-FST R@20 15.1 11.8	+ LBLC AUROC 83.8 84.7	AP 38.6 28.9
Label Errors	SELFCLEAN (SimCLR) SELFCLEAN (DINO) IN + Confident Learning FastDup SELFCLEAN (IN) SELFCLEAN (C) C C	40.0 95.0 P@20 25.0 15.0 40.0	12.3 29.3 DDI R@20 7.9 4.7 12.9	96.1 99.6 + LBL AUROC 71.4 68.6 60.3	AP 17.8 14.6 22.9	100.0 P@20 30.0 25.0 15.0	31.0 DDI + R@20 9.7 4.4 2.5	99.6 - LBLC AUROC 72.0 71.4 43.8 43.8	AP 23.6 28.8 16.7	95.0 P@20 35.0 40.0	37.3 HQ-FS R@20 15.6 16.0 6.1	JJ.J T + LBL AUROC 77.0 85.2 48.5	AP 31.7 32.3 10.7	P@20 40.0 30.0 20.0	31.4 HQ-FST R@20 15.1 11.8 7.8	+ LBLC AUROC 83.8 84.7 49.1	AP 38.6 28.9
Label Errors	SELFCLEAN (SimCLR) SELFCLEAN (DINO) IN + Confident Learning FastDup SELFCLEAN (IN) SELFCLEAN (SimCLR) SELFCLEAN (DINO)	40.0 95.0 P@20 25.0 15.0 40.0 10.0 30.0	12.3 29.3 DDI R@20 7.9 4.7 12.9 3.3 9.8	96.1 99.6 + LBL AUROC 71.4 68.6 60.3 39.6 77.2	AP 17.8 14.6 22.9 8.0 26.3	100.0 P@20 30.0 25.0 15.0 15.0 85.0	31.0 DDI 4 R@20 9.7 4.4 2.5 2.6 14.8	99.6 • LBLC AUROC 72.0 71.4 43.8 42.8 78.7	65.4 AP 23.6 28.8 16.7 16.0 50.0	95.0 P@20 35.0 40.0 15.0 10.0 50.0	37.3 HQ-FS R@20 15.6 16.0 6.1 4.1 20.4	T + LBL AUROC 77.0 85.2 48.5 56.3 85.1	AP 31.7 32.3 10.7 17.7 39.7	P@20 40.0 30.0 20.0 15.0 40.0	31.4 HQ-FST R@20 15.1 11.8 7.8 6.0 16.0	49.1 47.2 79.9	AP 38.6 28.9 12.8 13.2 35.2

E.2 Detailed comparison with metadata

Dataset	Metadata	Positive Samples (%)	AUROC (%)	AP (%)
PAD-UFES-20	Same Lesion	0.06	71.0	10.0
PAD-UFES-20	Not Biopsied	41.60	57.0	61.3
HAM10000	Same Lesion	0.01	98.8	28.4
HAM10000	Not Follow-Up	63.89	90.8	86.3
ISIC-2019	Same Lesion	0.01	98.2	26.6
CheXpert	Same Patient	0.01	70.5	7.5
ImageNet-1k [†]	Verified Label Errors	88.73	67.3	8.4
Food-101N [‡]	Verified Label Errors	18.51	79.8	47.8
CelebA	Same Person	0.02	78.8	30.9

Table 5: Comparison of the SELFCLEAN ranking with metadata from multiple benchmark datasets. For reference, we include the proportion of positive samples, also corresponding to the not-informed baseline performing best in terms of AP. Please consult section 5.2 for guidance on interpretation.

E.3 Detailed influence of dataset cleaning

 $\begin{array}{r} 40.3^{+0.8}_{-0.9}\\ 31.2^{+0.8}_{-0.9}\end{array}$

Food-101N

ImageNet-1k

Table 6: Influence of removing samples detected in the automatic cleaning mode with $\alpha = 0.10$ and q = 0.05 on downstram tasks. We report macro-averaged F1 scores for linear and kNN classifiers on DINO features over 100 random training/evaluation splits with 80% and 20% fractions respectively. We compute paired performance differences before and after cleaning the evaluation set, and before and after cleaning also the training set. We report the median and the intervals to the 5% (subscript) and 95% (superscript) percentiles. Additionally, we indicate significance of a paired permutation test on the difference sign with *p < 0.05, **p < 0.01, and ***p < 0.001.

			Linear Classifier	ſ		
		Scores (%)		Differences (%)		
Dataset	Cont + Cont	Cont + Clean	Clean + Clean	Clean Eval	Clean Train	
DDI HAM10000 Fitzpatrick17k	$\begin{array}{c} 59.2^{+9.6}_{-10.2} \\ 62.6^{+4.2}_{-4.2} \\ 52.8^{+2.6}_{-3.1} \end{array}$	$\begin{array}{c} 59.6^{+12.0}_{-11.2} \\ 63.0^{+3.3}_{-4.0} \\ 52.5^{+2.5}_{-4.1} \end{array}$	$58.9^{+9.0}_{-9.7}$ $62.8^{+3.2}_{-3.8}$ $52.6^{+2.9}_{-2.8}$	$ \begin{array}{c} +1.0^{+11.1}_{-11.2} \\ +0.1^{+3.2}_{-3.5} \\ -0.6^{+2.9}_{-3.6} ** \end{array} $	$\begin{array}{r} -0.7^{+7.7}_{-10.8} \\ -0.1^{+3.9}_{-3.6} \\ +0.2^{+3.3}_{-3.9} \ * \end{array}$	
Food-101N ImageNet-1k	$50.0^{+0.9}_{-1.2}$ $42.4^{+0.7}_{-0.9}$	$50.1^{+1.1}_{-1.0} \\ 42.0^{+0.9}_{-0.9}$	$50.4^{+0.8}_{-1.2}$ $42.2^{+0.6}_{-1.0}$	$ \begin{array}{c} +0.2^{+0.6}_{-0.5} & *** \\ -0.4^{+0.6}_{-0.6} & *** \end{array} $	$^{+0.1^{+0.6}_{-0.5}}_{-0.0^{+0.9}_{-0.5}}$	
			kNN Classifier			
		Scores (%)		Differer	nces (%)	
	Cont + Cont	Cont + Clean	Clean + Clean	Clean Eval	Clean Train	
DDI HAM10000 Fitzpatrick17k	$58.2_{-8.3}^{+7.7}$ $58.3_{-4.9}^{+3.4}$ $60.2_{-1.9}^{+1.8}$	$\begin{array}{c} 59.2^{+7.5}_{-8.3}\\ 58.3^{+3.7}_{-4.7}\\ 56.1^{+1.9}_{-2.2}\end{array}$	$59.7^{+7.3}_{-8.8}$ $58.7^{+3.1}_{-4.6}$ $56.1^{+2.0}_{-2.3}$	$ \begin{array}{c} +1.2^{+1.9}_{-1.2} *** \\ +0.2^{+0.5}_{-0.4} \\ -4.1^{+1.2}_{-1.3} *** \end{array} $	$ \begin{array}{c} +0.0^{+1.7}_{-1.4} *** \\ +0.2^{+1.3}_{-0.8} ** \\ +0.1^{+2.0}_{-1.7} \end{array} $	

 $\begin{array}{r} 40.5^{+0.7}_{-1.1}\\ 31.1^{+0.8}_{-0.9} \end{array}$

 $+0.1^{+0.2}_{-0.2}_{+0.4}^{+0.3}_{+0.4}^{+0.3}_{+0.4}^{+0.3}$

 $+0.1^{+0.1}$ ***

 $+0.1_{-0.1}^{+0.1}$ $-0.4_{-0.2}^{+0.1}$ ***

 $\begin{array}{r} 40.4^{+0.7}_{-1.1}\\ 30.8^{+0.9}_{-0.9}\end{array}$

[†]Refers to the subset of ImageNet-1k validation set which was verified by Northcutt et al. [8].

^{*}Refers to the subset of Food-101N set which was verified by Lee et al. [50].

E.4 Details on the application to benchmark datasets

Table 7: Estimated percentage of data quality issues in multiple vision benchmarks estimated using SELFCLEAN's automatic mode with $\alpha = 0.10$ and q = 0.05. Images marked as originating from the same person, patient, or lesion were excluded from the near-duplicate count whenever this metadata was available. See appendix K for examples of the problems which were automatically found.

	Es	stimated Errors	
Dataset	Irrelevant Samples	Near Duplicates	Label Errors
DDI	1(0.2%)	4(0.6%)	5(0.8%)
PAD-UFES-20	0(0.0%)	0(0.0%)	5(0.4%)
HAM10000	0(0.0%)	1(0.6%)	17(0.2%)
Fitzpatrick17k	18(0.1%)	2,446 (14.8%)	$103\ (0.6\%)$
ISIC-2019	0(0.0%)	1,200(5.4%)	$97\ (0.4\%)$
PatchCamelyon	98~(0.3%)	12,845 (3.9%)	589(0.2%)
CheXpert [§]	6(0.0%)	0(0.0%)	303~(0.6%)
CelebA	2(0.0%)	810 (0.4%)	1,033(0.5%)
Food-101N	310(0.1%)	4,433(1.4%)	2,728(0.9%)
ImageNet-1k [¶]	0(0.0%)	36(0.1%)	262~(0.5%)

[§]Refers to the label errors in atelectasis.

[¶]Refers to the ImageNet-1k validation set, similar to [8].

F Automatic cleaning

In section 3 we constructed scores which take extreme values for candidate problematic data points. Isolating data issues using such a score is essentially a one-dimensional anomaly detection problem. Here we construct a procedure to detect outlier scores which works well in our case, and demonstrate that despite the introduction of two hyperparameters our results do not depend much on arbitrary choices.

F.1 Automatic cleaning procedure

We start with the intuition that detecting problematic samples is easy if the scores are smoothly distributed for normal data, but are far from the bulk for data with quality issues. However, all our scores range from 0 to 1 and increasingly extreme issues are simply closer and closer to zero without leaving large gaps on the score scale. For this reason, we expand the neighborhoods of 0 and 1 using a logit transformation, $\tilde{s} = \log[s/(1-s)]$. The transformed scores \tilde{s} then range over the whole real axis enabling a better separation between normal and problematic samples.

Since the logit transformation has jacobian $|d\tilde{s}/ds| = e^{-\tilde{s}}/(1 + e^{-\tilde{s}})^2$, under broad assumptions we expect the dominant behavior of the transformed score distribution to drop at least as quickly as a logistic probability density function for $\tilde{s} \to \pm \infty$. Note that this is the case even if the original score distribution is not just constant but presents an integrable power-law singularity for $s \to 0, 1$.

In order to identify the samples which are very unlikely to belong to normal data, we first attempt to isolate a region on the left tail of the distribution which is not affected by noise. To this end, we introduce a hyperparameter α , the "contamination rate guess", which represents a generous estimate of the fraction of dirty data in the dataset. For data quality issues where a score is associated to each sample, we simply drop the lowest $\lfloor \alpha_1 N \rfloor$ scores with $\alpha_1 = \alpha$, while when a score is associated to a pair of samples, we discard the lowest $\lfloor \alpha_1 N (N-1)/2 \rfloor$ scores with $\alpha_1 = \alpha^2$. Indeed, in the case where there are no interactions (e.g., only pairs of near-duplicates) we expect αN abnormally low scores, but in the worst case interaction scenario (e.g., all views of the very same sample) we await $\alpha N(\alpha N - 1)/2$ low out-of-distribution scores, which reduces to the above expression for α_1 when $\alpha N \gg 1$. Besides dropping the potentially problematic samples, we also select an upper score bound for the range of interest, since we aim at reproducing only the smooth *left* tail of the distribution. Reasonable choices are values between the lower score cutoff determined by α_1 and the median, paying attention that enough data is included for sufficient robustness to noise. For this reason, we choose the upper score cutoff to be the quantile corresponding to a fraction of data α_2 which is the geometric mean between α_1 and 1/2, i.e. $\alpha_2^2 = \alpha_1/2$. We observe that the range produces robust statistical information if the number of samples is sufficiently large and $\alpha \ll 1/2$, where in practice $\alpha \lesssim 1/4$ is already stable.

Following our heuristic argument, we approximate the smooth component of the left tail of the distribution using a logistic distribution with suitably chosen scale and location parameters, which has probability density function

$$pdf(\tilde{s};\mu,\sigma) = \frac{1}{\sigma} pdf\left(\frac{\tilde{s}-\mu}{\sigma}\right), \qquad pdf(\hat{s}) = \frac{e^{-\hat{s}}}{(1+e^{-\hat{s}})^2}.$$
(3)

Given the score cutoffs \bar{s}_1 and \bar{s}_2 corresponding to the quantiles α_1 and α_2 of the empirically observed distribution, the scale σ and location μ can be estimated as

$$\sigma = \frac{\bar{s}_2 - \bar{s}_1}{\bar{s}(\alpha_2) - \bar{s}(\alpha_1)}, \qquad \mu = \frac{\bar{s}_1 \bar{s}(\alpha_2) - \bar{s}_2 \bar{s}(\alpha_1)}{\bar{s}(\alpha_2) - \bar{s}(\alpha_1)}, \qquad \bar{s}(\alpha_m) = \log \frac{\alpha_m}{1 - \alpha_m} \quad \text{for} \quad m = 1, 2.$$
(4)

Here $\bar{s}(\alpha_m)$ indicates the percentage point function of the logistic distribution, i.e. the inverse of its cumulative distribution function. Note that the whole estimation procedure for the left tail of the distribution relies exclusively on quantiles and is therefore naturally robust to outliers.

With an estimate of the smooth score distribution for normal data, we can identify abnormal samples by requesting that they be very unlikely generated by the same random process. This is achieved by demanding that the probability of obtaining a score below an outlier cutoff s_{cut} be less than a significance level q times the expected fraction of outliers, which is $2\alpha/(N-1)$ in the case of pairs of samples and α otherwise. We set the hyperparameter q to 0.05 corresponding to a 95% one-sided confidence level and study the influence of this choice in section F.4. All samples with scores lower than the outlier cutoff will be then classified as problematic.

The value of the procedure described above lies in the fact that, despite requiring several engineering steps and introducing two additional hyperparameters, it detects a number of outliers which is largely independent of any reasonable choice for α and q. The remaining parts of appendix F are dedicated to showing that the procedure is intuitive and assumptions are empirically acceptable (F.2), and to demonstrating that detected outliers exhibit low sensitivity to the contamination rate guess α (F.3) and to the significance level q (F.4).



(d) Food-101N [50]

Figure 5: Score histogram (blue) and associated left-tail distribution fit (solid orange) with outlier cutoff (dashed orange) for all considered noise types and representative datasets. The green shaded area represents the range $[\bar{s}_1, \bar{s}_2]$ which is used to determine location and scale of the associated logistic distribution. The values $\alpha = 0.10$ and q = 0.05 are used throughout.

F.2 Automatic cleaning examples

In figure 5 we illustrate the fit to the left tails of distributions for representative datasets, together with the relevant range used to estimate scale and location, and the position of the outlier cutoff to classify data quality issues. We observe that the probability density function is a qualitatively good estimate of the density-normalized histograms in the expected range, i.e. for the smooth component of the histogram's left tail, within sampling uncertainties. The fit quality is somewhat lower for irrelevant samples, probably due to the score range which is all above $\tilde{s} = 0$. We also carried out experiments with a gaussian functional form for score distribution tails, and observed only minor changes which resulted in a slightly reduced number of detected problems.

F.3 Influence of the contamination rate guess α

In figure 6 we analyze the sensitivity of the number of detected data quality problems on the contamination rate guess α , for all noise types and representative datasets analyzed in this paper. In these plots, the significance level q is fixed to its default value of 0.05. We observe that the fraction of found problems does not depend much on α over several orders of magnitude, suggesting a sensitivity to this hyperparameter that is approximately vanishing or at most logarithmic. It is by virtue of this reduced dependence that we can fix $\alpha = 0.10$ throughout the rest of the paper, and that fully-automatic cleaning is able to produce stable results with limited prior knowledge of dataset quality.

F.4 Influence of the significance level q

In figure 7 we report the fraction of detected problematic samples as a function of the significance level q, for all considered noise types and for representative datasets. We can see that this hyperparameter essentially determines the number of outliers found, which is monotonically increasing with q. Moreover, the number of identified issues has, in most cases, a dependence on q which is less than linear. In some cases, especially when the number of detected outliers is below percent level or q approaches 1, we see more severe sensitivity to the specific value. This may be because the empirical score distribution changes more abruptly than estimated by the logistic fit, as happens for irrelevant samples, or because the region immediately below the lower score cutoff $\bar{s}(\alpha_1)$ (which corresponds to q = 1) is densely populated almost by construction. It is however clear that q regulates how extreme scores need to be for a sample to be considered problematic. A value of $q = 10^{-3}$ will only select very apparent data quality issues, q = 1/4 will almost certainly also include a significant fraction of valid samples, and our choice of q = 0.05 strikes a compromise between precision and recall.



Figure 6: Dependence of the fraction of detected data quality issues on the contamination rate guess α for all considered noise types and representative datasets, at a fixed significance level q = 0.05. The observed behavior is reported in blue. It is outside of the lower margin of the plots when no problems are found. The green solid line represents a fraction of detected issues which is equal to the contamination rate guess for reference. The vertical dotted red line indicates the default value $\alpha = 0.10$ used in the rest of the paper.



Figure 7: Impact of the choice of the significance level q on the fraction of detected data quality issues, across noise types and for representative datasets, for a fixed contamination rate guess $\alpha = 0.10$. The observed dependency on q is reported in blue, and it is below the lower margin of the plots when no problematic samples are found. The diagonal green solid line is just a reference to guide reading, and the dotted red line indicates the default choice q = 0.05.

G Validating algorithmic rankings with human experts

In this section, we describe the procedure used to confirm that SELFCLEAN assigns low ranks to problematic samples and high ranks to normal data, as discussed in the second part of section 5.2. To this end, for each data quality issue type, we collect human annotations for the first 50 images and for 50 images randomly sampled from the dataset. Annotators use a custom tool which is shown in figure 8. The annotation process starts with the selection of a dataset and data quality issue (e.g. the Fitzpatrick17k dataset and irrelevant samples) and then proceeds with binary questions about single images or pairs thereof depending on the task. Section G.1 shows the task descriptions for each quality issue. Note that samples ranks are not displayed to avoid potential bias.

We paid crowd workers 0.03 US dollars per annotation for images from ImageNet and Food-101N, which roughly corresponds to 9 US dollars per hour. Medical experts were not compensated financially but were instead acknowledged with co-authorship in a labeling consortium.



Figure 8: The verification tool used by the medical experts for annotating the ranking of dataset quality issues by SELFCLEAN.

G.1 Task descriptions

This section gives detailed information about each task description shown to the annotators:

- Irrelevant samples: You will see one image, and your task is to judge whether the image reflects an irrelevant sample. We define samples to be irrelevant when they are not expected to be valid inputs in the context of a specific task.
- Near duplicates: You will see two images, and your task is to judge whether the images reflect a near-duplicate. We define near-duplicates as pairs of images that contain two views of the same object. In this sense, exact duplicates are an easy special case of near duplicates.
- Label errors: You will see one image along with the corresponding label, and your task is to judge whether the label of the image is correct. The goal is to detect mismatches between the label and the image. Label errors are simply samples annotated with a wrong class label.

G.2 Detailed results

In order to verify that problematic samples tend to appear first in the ranking provided by SELFCLEAN, for each noise type, we first consider all 100 annotated images and then the first 50 in the ranking only. We conduct Mann-Whitney U statistical tests to verify that humans are more likely to identify data quality issues in samples that appear first in the SELFCLEAN ranking. In order to gain more intuitive understanding, we also report the fraction of samples that were found to be problematic within the first 50 and the 50 random samples, and within samples ranked 1 through 25 and 26 through 50. Finally, we visualize the distribution of human-confirmed problems through the ranking by plotting the fraction of confirmed problems in a rolling window of ten ranks in figure 9.

We observe very good performance for near-duplicate detection throughout the considered datasets. Label-error identification is very convincing in all cases but for DDI. The different concentration of problems is mostly

observed between images with low ranking and random samples, while the difference between samples 1-25 and 26-50 is less pronounced. We observe that identifying label errors in a highly-curated dataset such as DDI is a highly-nontrivial task which might exceed the design of our experiment. Finally, the detection of irrelevant samples is the task where SELFCLEAN achieves the lowest agreement with human annotators. Nevertheless, our results still suggest a significant separation of irrelevant samples within the ranking in at least half of the cases.

Table 8: Comparison of the percentage of errors found by experts in the 50 lowest-ranked samples with 50 random samples, and in samples 1 to 25 with samples 26 through 50. We report the percentage of errors in each sample and the corresponding p-value of a Mann–Whitney U test, which represents the probability for the ranking to be unrelated to the position of positive samples.

			Percentage	e of Human-Co	nfirmed Problems		
Dataset	Data Quality Issue	Lowest 1-50 (%)	Random Sample (%)	p-value	Lowest 1-25 (%)	Lowest 26-50 (%)	p-value
DDI	Irrelevant Samples	12	8	0.25	20	4	0.04
DDI	Near Duplicates	12	0	0.006	24	0	0.005
DDI	Label Errors	22	32	0.86	20	24	0.63
Fitzpatrick17k	Irrelevant Samples	14	4	0.04	12	16	0.65
Fitzpatrick17k	Near Duplicates	100	0	$1.3 imes10^{-23}$	100	100	undef
Fitzpatrick17k	Label Errors	54	12	$4.4 imes 10^{-6}$	52	56	0.61
ImageNet	Irrelevant Samples	62	48	0.08	56	68	0.80
ImageNet	Near Duplicates	92	0	$2.1 imes \mathbf{10^{-20}}$	100	84	0.02
ImageNet	Label Errors	36	0	$1.6 imes \mathbf{10^{-6}}$	48	24	0.04
Food-101N	Irrelevant Samples	24	4	0.002	36	12	0.02
Food-101N	Near Duplicates	100	0	$f 1.3 imes 10^{-23}$	100	100	undef
Food-101N	Label Errors	72	34	$7.6 imes 10^{-5}$	80	64	0.61



Figure 9: Visualization of the percentage of quality issues found across the first 50 samples in the SELFCLEAN ranking and in 50 random samples, using a rolling window of size 10. Results are reported across four datasets and for each noise type.

H Scoring for irrelevant samples

This section describes how to construct a score based on hierarchical clustering, such that samples with a high probability of being irrelevant have significantly lower values compared to the bulk of the data. Note that, although in practice we use single-linkage agglomerative clustering, our heuristic construction can be applied to any hierarchical clustering and is formulated accordingly.

Hierarchical clustering over a set of leaves $\{1, \ldots, N\}$ is defined by a hierarchy of sets $C_n = \{C_{1n}, \ldots, C_{nn}\}$ for *n* going from *N* to 1. Without loss of generality, we assume indices of merged sets are always consecutive, and the other sets do not change their relative order

$$C_{in} = \begin{cases} C_{i(n+1)} & \text{if } i < i_n, \\ C_{i(n+1)} \cup C_{(i+1)(n+1)} & \text{if } i = i_n, \\ C_{(i+1)(n+1)} & \text{if } i > i_n, \end{cases} \quad \text{for } i = 1, \dots, n \text{ and } n = 1, \dots, N.$$
(5)

The hierarchy of sets induces a dendrogram, i.e. a tree graph where each cluster is a node connected to its direct parent and children. Each element of the hierarchy is also associated with a merge distance $d_{n-1} = d(C_{i_nn}, C_{(i_n+1)n})$. To define a ranking, we sort the dendrogram such that $|C_{i(n+1)}| \leq |C_{(i+1)(n+1)}|$, i.e. the cluster which contains the least leaves comes first, based on the idea that outliers are associated with merges containing fewer leaves [46]. In case of ties, the cluster created at the largest distance precedes the other.

To produce a scoring for each sample from the dendrogram, natural building blocks are the merge distance, the sizes of the merged clusters, and their interactions [76]. Accordingly, we define scores by drawing the dendrogram in a $[0,1] \times [0,1]$ square where the horizontal axis is one minus the distance and the vertical axis is the weight w_{in} of cluster *in* which is defined recursively. This graphical construction is illustrated in the right panel of figure 10. The score of each leaf s_i is determined as the area under the curve $f_i(d)$ where

$$f_i(d) = w_{jn}$$
 if $d_n \le d < d_{n-1}$ and $i \in C_{jn}$, $n = 1, \dots, N$, (6)

with $d_N = 0$ and $d_0 = 1$. For convenience, we define $p_{in} = |C_{in}|/N$ to be the probability of cluster *in* and set $w_{0n} = 0$ and $w_{11} = 1$. To define the weights, we propose a rule which we call leaves and distances (LAD) and reads

$$w_{i(n+1)} = \begin{cases} w_{in} & \text{if } i < i_n, \\ w_{(i_n-1)n} + (w_{i_nn} - w_{(i_n-1)n}) p_{i(n+1)} / p_{i_nn} & \text{if } i_n \le i \le i_n + 1, \\ w_{(i-1)n} & \text{if } i > i_n + 1. \end{cases}$$
(7)

This completes our construction of the scores s_i for irrelevant samples.

We note that although the formulation given here is limited to the case of binary merges and no ties among distances, generalization to these cases is straightforward except for resolving how to sort the clusters.



Figure 10: The left plot shows an example of a dendrogram for hierarchical clustering, and the right plot an illustration of the LAD scoring. In the left plot, the x-axis shows the ranking of the different points in brackets and the corresponding identification number. In the right plot, the right side of the y-axis shows the weights w_{in} corresponding to equation 7.

I Hyperparameters

In this section we specify parameter values used throughout the paper for both pre-training and evaluation.

Table 9 lists the hyperparameters used for pre-training for both DINO and SimCLR. The used parameters are very similar to the introductory papers of these approaches [43, 44] with the exception that the global crop scale is larger and there are more local crops, which we have found to be beneficial for smaller datasets (< 20,000). However, we observed no benefit or harm for larger datasets with these minor hyperparameter changes.

Table 10 lists the configuration for producing near-duplicate images using the MedAugment and ArtefactAugment strategies. The configuration was chosen to mimic the natural contamination of near duplicates in benchmark datasets.

Table 9: Hyperparameters used for pre-training using SimCLR and DINO on the dataset to clean. Here "-" indicates that the respective parameter is not used for the corresponding pre-training strategy. Parameters not given in this list follow the introductory paper. More detailed information about the hyperparameters can be found in the open-sourced codebase.

Hyperparameter	SimCLR	DINO
Batch size	90	64
Epochs	100	100
Optimizer	Adam	AdamW
Learning rate	0.001	0.0005
Min. learning rate	1e-6	1e-6
Weight decay	0.04	0.04
Weight decay end	0.4	0.4
Warmup epochs	10	10
Momentum teacher	-	0.996
Clip grad.	3.0	3.0
Base model	ViT tiny	ViT tiny
Model embedding dim.	192	192
Model output dim.	128	4096
Model patch size	16	16
Model drop path rate	0.1	0.1
Norm last layer	-	True
Global crops scale	-	(0.7, 1.)
Local crops scale	-	(0.05, 0.4)
Global crops number	-	2
Local crops number	-	12
Warmup teacher temp.	-	0.04
Teacher temp.	-	0.04
Warmup teacher temp. epochs	-	30
Contrastive temp.	0.5	-

Table 10: Parameter values used for the MedAugment (MED) (10a) and ArtefactAugment (ARTE) (10b) strategies. More detailed configurations can be found in the open-sourced codebase.

(a) MedAugment (l	MED)	
Hyperparameter	MED	(b) ArteractAugn
Rotation proba.	0.5	Hyperparameter
Padding proba.	0.5	Watermark proba.
Blur proba.	0.5	Colorbar proba.
Rotation degree range	(0, 180)	Collage proba.
Scale range	(0.5, 0.9)	Watermark max. s
Padding	3	Collage max. scale
Gaussian kernel size	5	

tAugment (ARTE)

ARTE 0.5 0.5 0.5 0.5 0.5

Watermark proba.
Colorbar proba.
Collage proba.
Watermark max. scale
Collage max. scale

J Visualization of artificial augmentation strategies

This section visualizes the different synthetic contamination strategies used for the evaluation of 5.1. Figure 11 visualizes both near-duplicate augmentation strategies, namely MedAugment (MED) and ArtefactAugment (ARTE). Figure 12 shows a random sample of the irrelevant images from the combined medical images (CMED) strategy added to a dataset for contamination. The gallery further shows some irrelevant X-ray images from the XR augmentation strategy.





(b) ArtefactAugment (ARTE)

Figure 11: Random samples of near duplicates produced by synthetic augmentation strategies and used to evaluate near-duplicate detection. Figure 11a shows examples of the MedAugment strategy (MED), consisting of random rotation, flipping, resizing, applying padding, and Gaussian blur. Figure 11b shows examples of the ArtefactAugment strategy (ARTE), which consists of adding typical artifacts found in medical image collections. This includes adding watermarks, color bars, and rulers and then scaling the image and randomly adding additional images to create a collage. Images one, two, four, and five show examples of collages.



Figure 12: A random sample of irrelevant samples from the *CMED* contamination strategy. The images originate from multiple datasets, consisting of surgical tools [64], X-ray images [63], ImageNet samples [19], histopathological images [65], segmentation masks [66] and pictures of PowerPoint slides [67].

K Inspection of benchmark datasets

This section contains illustrations of the rankings produced by SELFCLEAN for multiple vision benchmarks, divided into the three main noise categories of irrelevant samples, near duplicates, and label errors.

K.1 ImageNet-1k



Figure 13: Ranking produced by SELFCLEAN for near duplicates in the ImageNet-1k validation set, of which the top-15 are shown along with the respective rank and index.



Figure 14: Ranking produced by SELFCLEAN for irrelevant samples in the ImageNet-1k validation set, of which the top-15 are shown along with the respective rank and index.



Figure 15: Ranking produced by SELFCLEAN for label errors in the ImageNet-1k validation set, of which the top-15 are shown along with the respective rank, index, and original label.



Figure 16: Ranking produced by SELFCLEAN for near duplicates in CheXpert, of which the top-15 are shown along with the respective rank and index.



Figure 17: Ranking produced by SELFCLEAN for irrelevant samples in CheXpert, of which the top-15 are shown along with the respective rank and index.



Figure 18: Ranking produced by SELFCLEAN for atelectasis label errors in CheXpert, of which the top-15 are shown along with the respective rank, index, and original label.

K.3 PatchCamelyon



Figure 19: Ranking produced by SELFCLEAN for near duplicates in PatchCamelyon, of which the top-15 are shown along with the respective rank and index.



Figure 20: Ranking produced by SELFCLEAN for irrelevant samples in PatchCamelyon, of which the top-15 are shown along with the respective rank and index.



Figure 21: Ranking produced by SELFCLEAN for label errors in PatchCamelyon, of which the top-15 are shown along with the respective rank, index, and original label, i.e. if the patch is tumerous.

K.4 Fitzpatrick17k



Figure 22: Ranking produced by SELFCLEAN for near duplicates in the Fitzpatrick17k, of which the top-15 are shown along with the respective rank and index.



Figure 23: Ranking produced by SELFCLEAN for irrelevant samples in the Fitzpatrick17k, of which the top-15 are shown along with the respective rank and index.



Figure 24: Ranking produced by SELFCLEAN for label errors in the Fitzpatrick17k, of which the top-15 are shown along with the respective rank, index, and original label.



Figure 25: The figure shows multiple near-duplicates found in the Fitzpatrick17k dataset [56]. Figures 25a and 25b show that SELFCLEAN can find near-duplicates where the images were anonymized. Further, figures 25c and 25d show extreme cases of near-duplicates where the method found crops of the same image. Finally, figures 25e and 25f show examples where the near-duplicate detection failed.