Don't fear the unlabelled: Safe semi-supervised learning via simple debiasing

Anonymous Author(s)
Affiliation
Address
email

Abstract

Semi-supervised learning (SSL) provides an effective means of leveraging unlabelled data to improve a model's performance. Even though the domain has received a considerable amount of attention in the past years, most methods present the common drawback of lacking theoretical guarantees. Our starting point is to notice that the estimate of the risk that most discriminative SSL methods minimise is biased, even asymptotically. This bias impedes the use of standard statistical learning theory and can hurt empirical performance. We propose a simple way of removing the bias. Our debiasing approach is straightforward to implement and applicable to most deep SSL methods. We provide simple theoretical guarantees on the trustworthiness of these modified methods, without having to rely on the strong assumptions on the data distribution that SSL theory usually requires. In particular, we provide generalisation error bounds for the proposed methods. We evaluate debiased versions of different existing SSL methods, such as the Pseudo-label method and Fixmatch, and show that debiasing can compete with classic deep SSL techniques in various settings by providing better calibrated models. Additionally, we provide a theoretical explanation of the intuition of the popular SSL methods.

1 Introduction

The promise of semi-supervised learning (SSL) is to be able to learn powerful predictive models using partially labelled data. In turn, this would allow machine learning to be less dependent on the often costly and sometimes dangerously biased task of labelling data. Early SSL approaches—e.g. Scudder's (1965) untaught pattern recognition machine—simply replaced unknown labels by predictions made by some estimate of the predictive model and used the obtained *pseudo-labels* to refine their initial estimate. Other more complex branches of SSL have been explored since, notably using generative models (from McLachlan, 1977, to Kingma et al., 2014) or graphs (notably following Zhu et al., 2003). Deep neural networks, which are state-of-the art supervised predictors, have been trained successfully using SSL. Somewhat surprisingly, the main ingredient of their success is still the notion of pseudo-labels (or one of its variants), combined with a systematic use of data augmentation (e.g. Xie et al., 2019; Sohn et al., 2020; Rizve et al., 2021).

An obvious SSL baseline is to simply to throw away the unlabelled data. We will call such a baseline the *complete case*, following the missing data literature (e.g. Tsiatis, 2006). As reported in van Engelen & Hoos (2020), the main risk of SSL is the potential degradation caused by the introduction of unlabelled data. Indeed, semi-supervised learning outperforms the complete case baseline only in specific cases (Singh et al., 2008; Schölkopf et al., 2012; Li & Zhou, 2014). This degradation risk for generative models has been analysed in Chapelle et al. (2006, Chapter 4). To overcome this issue, previous works introduced the notion *safe* semi-supervised learning for techniques which never reduce predictive performance by introducing unlabelled data (Li & Zhou, 2014; Guo et al., 2020). Our loose

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

definition of safeness is as follows: *a SSL algorithm is safe if it has theoretical guarantees that are similar or stronger to the complete case baseline*. The "theoretical" part of the definition is motivated by the fact that any empirical assessment of generalisation performances of an SSL algorithm is jeopardised by the scarcity of labels. Unfortunately, popular deep SSL techniques generally does not benefit of theoretical guarantees without strong and essentially untestable assumptions on the data distribution (Mey & Loog, 2019) such the smoothness assumption (small perturbations on the features x do not cause large modification in the labels, $p(y|pert(x)) \approx p(y|x)$) or the cluster assumption (data points are distributed on discrete clusters and points in the same cluster are likely to share the same label).

Most semi-supervised methods rely on these distributional assumptions to ensure performance in entropy minimisation, pseudo-labelling and consistency-based methods. However, no proof is given that guarantees the effectiveness of state-of-the-art methods (Tarvainen & Valpola, 2017; Miyato et al., 2018; Sohn et al., 2020; Pham et al., 2021). To illustrate that SSL requires specific assumptions, we show in a toy example that pseudo-labelling fails at learning. To do so, we draw samples from two uniform distributions with a small overlap. Both supervised and semi-supervised neural networks are trained using the same labelled dataset. While the supervised algorithm learns perfectly the true distribution of $p(1|x)$, the semi-supervised learning methods (both entropy minimisation and pseudo-label) underestimate $p(1|x)$ for $x \in [1, 3]$ (see Figure 1). We also test our proposed method (DeSSL) on this dataset and show that the unbiased version of each SSL technique learns the true distribution accurately. See Appendix A for the results with Entropy Minimisation.

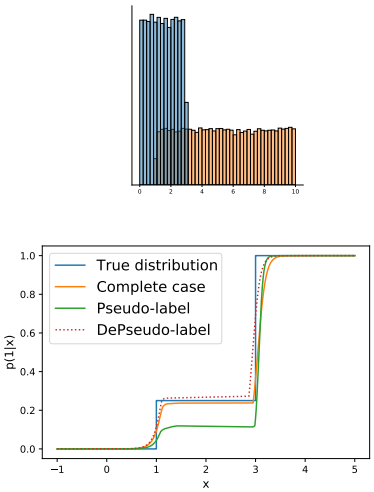

Figure 1: (Left) Data histogram. (Right) Posterior probabilities $p(1|x)$ of the same model trained following either complete case (only labelled data), Pseudo-label or our DePseudo-label. $n_l = 25,000, n_u = 25,000$.

1.1 Contributions

Rather than relying on the strong geometric assumptions usually used in SSL theory, we simply use the *missing completely at random (MCAR)* assumption, a standard assumption from the missing data literature (see e.g. Little & Rubin, 2019). With this only assumption on the data distribution, we propose a new safe SSL method derived from simply debiasing common SSL risk estimates. Our main contributions are:

- We introduce debiased SSL (DeSSL), a safe method that can be applied to most deep SSL algorithms without assumptions on the data distribution;
- We propose a theoretical explanation of the intuition of popular SSL methods. We provide theoretical guarantees on the safeness of using DeSSL both on consistency and calibration of the method. We also provide a generalisation error bound;
- We show how simple it is to apply DeSSL to the most popular methods such as Pseudo-label and Fixmatch, and show empirically that DeSSL leads to models that are never worse than their classical counterparts, generally better calibrated and sometimes much more accurate.

2 Semi-supervised learning

2.1 Learning with labelled data

The ultimate objective of most of the learning frameworks is to minimise a risk $\mathcal{R}$, defined as the expectation of a particular loss function L over a data distribution $p(x, y)$, on a set of models $f_\theta(x)$, parametrised by $\theta \in \Theta$. Thus, the learning task is finding θ^* that minimises the risk: $\mathcal{R}(\theta) = \mathbb{E}_{(X,Y)\sim p(x,y)}[L(\theta; X, Y)]$. The distribution $p(x, y)$ being unknown, we generally minimise

an approximation of the risk, the empirical risk $\hat{\mathcal{R}}(\theta)$ computed on a sample of n i.i.d points drawn from $p(x, y)$. $\hat{\mathcal{R}}(\theta)$ is an unbiased and consistent estimate of $\mathcal{R}(\theta)$ under mild assumptions. Its unbiased nature is one of the basic properties that is used for the development of traditional learning theory and asymptotic statistics (van der Vaart, 2000; Shalev-Shwartz & Ben-David, 2014).

2.2 Learning with both labelled and unlabelled data

Semi-supervised learning leverages both labelled and unlabelled data to improve the model's performance and generalisation. Further information on the distribution $p(x)$ provides a better understanding of the distributions $p(x, y)$ and also $p(y|x)$. Indeed, $p(x)$ may contain information on $p(y|x)$ (Schölkopf et al., 2012, Goodfellow et al., 2016, Chapter 7.6, van Engelen & Hoos, 2020).

In the following, we have access to n samples drawn from the distribution $p(x, y)$ where some of the labels are missing. We introduce a new random variable $r \in \{0, 1\}$ that governs whether or not a data point is labelled ($r = 0$ missing, $r = 1$ observed). The MCAR assumption states that the missingness of a label y is independent of its features and the value of the label: $p(x, y, r) = p(x, y)p(r)$. This is the case when nor features nor label carry information about the potential missingness of the labels. This description of semi-supervised learning as a missing data problem has already been done in multiple works –e.g. Seeger, 2000; Ahfock & McLachlan, 2019. Moreover, the MCAR assumption is implicitly made in most of the SSL works to design the experiments, indeed, missing labels are drawn completely as random in datasets such as MNIST, CIFAR or SVHN (Tarvainen & Valpola, 2017; Miyato et al., 2018; Xie et al., 2019; Sohn et al., 2020).

2.2.1 Complete case: throwing the unlabelled data away

In missing data theory, the complete case is the learning scheme that only uses fully observed instances, namely labelled data. The natural estimator of the risk is then simply the empirical risk computed on the labelled data. Fortunately, in the MCAR setting, the complete case risk estimate keeps the same good properties of the traditional supervised one: it is unbiased and converges pointwisely to $\mathcal{R}(\theta)$. Therefore, traditional learning theory holds for the complete case under MCAR. While these observations are hardly new (see e.g. Liu & Goldberg, 2020), they can be seen as particular cases of the theory that we develop below. The risk to minimise is

$$\hat{\mathcal{R}}_{CC}(\theta) = \frac{1}{n_l} \sum_{i=1}^{n_l} L(\theta; x_i, y_i). \tag{1}$$

2.2.2 Incorporating unlabelled data

A major drawback of the complete case framework is that a lot of data ends up not being exploited. A class of SSL approaches, mainly inductive methods with respect to the taxonomy of van Engelen & Hoos (2020), generally aim to minimise a modified estimator of the risk by including unlabelled data. Therefore, the optimisation problem generally becomes finding $\hat{\theta}$ that minimises the SSL risk,

$$\hat{\mathcal{R}}_{SSL}(\theta) = \frac{1}{n_l} \sum_{i=1}^{n_l} L(\theta; x_i, y_i) + \frac{\lambda}{n_u} \sum_{i=1}^{n_u} H(\theta; x_i). \tag{2}$$

where H is a term that does not depend on the labels and λ is a scalar weight which balances the labelled and unlabelled terms. In the literature, H can generally be seen as a surrogate of L. Indeed, it looks like the intuitive choices of H are equal or equivalent to a form of expectation of L on a distribution given by the model.

2.2.3 Some examples of surrogates

A recent overview of the recent SSL techniques has been proposed by van Engelen & Hoos (2020). In this work, we focus on methods suited for a discriminative probabilistic model $p_\theta(y|x)$ that approximates the conditional $p(y|x)$. We categorised methods into two distinct sections, the entropy and the consistency-based.

Entropy-based methods Entropy-based methods aim to minimise a term of entropy of the predictions computed on unlabelled data. Thus, they encourage the model to be confident on unlabelled data, implicitly using the cluster assumption. Entropy-based methods can all be described as an expectation of L under a distribution π_x computed at the datapoint x:

$$H(\theta; x) = \mathbb{E}_{\pi_x(\tilde{x},\tilde{y})}[L(\theta; \tilde{x}, \tilde{y})]. \tag{3}$$

For instance, Grandvalet & Bengio (2004) simply use the Shannon entropy as $H(\theta; x)$ which can be rewritten as equation (3) with $\pi_x(\tilde{x}, \tilde{y}) = \delta_x(\tilde{x})p_\theta(\tilde{y}|\tilde{x})$. Also, pseudo-label methods, which consist in picking the class with the maximum predicted probability as a pseudo-label for the unlabelled data (Scudder, 1965), can also be described as Equation 3. See Appendix B for complete description of the entropy-based literature (Berthelot et al., 2019; 2020; Xie et al., 2019; Sohn et al., 2020; Rizve et al., 2021; Zhang et al., 2021a) and further details.

Consistency-based methods Another range of SSL methods minimise a consistency objective that encourages invariant prediction for perturbations either on the data either on the model in order to enforce stability on model predictions. These methods rely on the smoothness assumption. In this category, we cite Π-model from (Sajjadi et al., 2016), temporal ensembling from (Laine & Aila, 2017), Mean-teacher proposed by (Tarvainen & Valpola, 2017), virtual adversarial training (VAT) from (Miyato et al., 2018) and interpolation consistent training (ICT) from (Verma et al., 2019). We remark that these objectives H are equivalent to an expectation of L (see Appendix B). The general form of the unsupervised objective can be written as

$$C_1\mathbb{E}_{\pi_x(\tilde{x},\tilde{y})}[L(\theta; \tilde{x}, \tilde{y})] \leq H(\theta; x) = \textbf{Div}(f_{\hat{\theta}}(x, .), \text{pert}(f_\theta(x, .)) \leq C_2\mathbb{E}_{\pi_x(\tilde{x},\tilde{y})}[L(\theta; \tilde{x}, \tilde{y})], \tag{4}$$

where the **Div** is a non-negative function that measures the divergence between two distributions, $\hat{\theta}$ is a fixed copy of the current parameter θ (the gradient is not propagated throught $\hat{\theta}$) and $0 \leq C_1 \leq C_2$.

Previous works also remarked that H is an expectation of L for entropy-minisation and pseudo-label (Zhu et al., 2022; Aminian et al., 2022). We describe a more general framework covering further methods and provide with our theory an intuition on the choice of H.

2.3 Theoretical guarantees

The main risk of SSL is the potential degradation caused by the introduction of unlabelled data when distributional assumptions are not satisfied (Singh et al., 2008; Schölkopf et al., 2012; Li & Zhou, 2014), specifically in settings where the MCAR assumption does not hold anymore (Oliver et al., 2018; Guo et al., 2020). Additionally, in (Zhu et al., 2022), the authors show disparate impacts of pseudo-labelling on the different sub-classes of the population. To mitigate these problems, previous works introduced the notion *safe* semi-supervised learning for techniques which never reduce learning performance by introducing unlabelled data (Li & Zhou, 2014; Kawakita & Takeuchi, 2014; Li et al., 2016; Gan et al., 2017; Trapp et al., 2017; Guo et al., 2020). As remark by Oliver et al. (2018), SSL performances are enabled by leveraging large validation sets which is not suited for real-world applications. Then, theoretical guarantees are required to use safely SSL algorithms. For this reason, in our work, we consider as *safe* a SSL algorithm that has theoretical guarantees that are similar or stronger than those of the complete case baseline. Even though the methods presented above produce good performances in a variety of SSL benchmarks, they generally do not benefit from theoretical guarantees, even elementary. More over,Schölkopf et al. (2012) identify settings on the causal relation between the features x and the target y where SSL may systematically fail, even if classic SSL assumptions hold. Our example of Figure 1 also shows that classic SSL may fail to generalise in a very benign setting with a large number of labelled data.

Presented methods minimise a biased version of the risk under the MCAR assumption and therefore classical learning theory cannot be applied anymore, as we argue more precisely in Appendix C. Learning over a biased estimate of the risk is not necessarily unsafe but it is difficult to provide theoretical guarantees on such methods even if some works try to do so with strong assumptions on the data distribution (Mey & Loog 2019, Section 4 and 5). Additionally, we remark that the choice of H can be confusing as seen in the literature. For instance, Grandvalet & Bengio (2004) and Corduneanu & Jaakkola (2003) perform respectively entropy and mutual information *minimisation* whereas Pereyra et al. (2017) and Krause et al. (2010) perform *maximisation* of the same quantities.

2.4 Related works

Previous works already proposed safe SSL methods with theoretical guarantees. Unfortunately, so far these methods come with either strong assumptions or important computational burden. Li & Zhou (2014) introduced a safe semi-supervised SVM and showed that the accuracy of their method is never worse than SVMs trained with only labelled data with the assumption that the true model is accessible. However, if the distributional assumptions are not satisfied, no improvement or degeneration is expected. Sakai et al. (2017) proposed an unbiased estimate of the risk for binary classification by including unlabelled data. The key idea is to use unlabelled data to better evaluate on the one hand the risk of positive class samples and on the other the risk of negative samples. They provided theoretical guarantees on its variance and a generalisation error bound. The method is designed only for binary classification and has not been tested in a deep learning setting. It has been extended to ordinal regression in follow-up work (Tsuchiya et al., 2021). In the context of kernel machines, Liu & Goldberg (2020) used an unbiased estimate of risk, like ours, for a specific choice of H. Guo et al. (2020) proposed DS^3L, a safe method that needs to approximately solve a bi-level optimisation problem. In particular, the method is designed for a different setting, not under the MCAR assumptiom, where there is a class mismatch between labelled and unlabelled data. The resolution of the optimisation problem provides a solution not worse than the complete case but comes with approximations. They provide a generalisation error bound. Also, the method does not outperform classic SSL methods in the MCAR setting as it is designed for non-MCAR situations. Sokolovska et al. (2008) proposed a safe method with strong assumptions such that the feature space is finite and the marginal probability distribution of x is fully known. Fox-Roberts & Rosten (2014) proposed an unbiased estimator in the generative setting applicable to a large range of models and they prove that this estimator has a lower variance than the one of complete case.

3 DeSSL: Unbiased semi-supervised learning

In order to overcome the issues introduced by the second term in the approximation of the risk for the semi-supervised learning approach, we propose DeSSL, an unbiased version of the SSL estimator using labelled data to annul the bias. The idea here is to retrieve the properties of classical learning theory. Fortunately, we will see that the proposed method can eventually have better properties than the complete case, in particular with regards to the variance of the estimate. The proposed DeSSL objective is

$$\hat{\mathcal{R}}_{DeSSL}(\theta) = \frac{1}{n_l}\sum_{i=1}^{n_l} L(\theta; x_i, y_i) + \frac{\lambda}{n_u}\sum_{i=1}^{n_u} H(\theta; x_i) - \frac{\lambda}{n_l}\sum_{i=1}^{n_l} H(\theta; x_i). \tag{5}$$

Under the MCAR assumption, this estimator is unbiased for any value of the parameter λ. For a proof of this result see Appendix D.

Intuitively, for entropy-based methods H should be applied only on unlabelled data to enforce the confidence of the model only on unlabelled datapoints. Whereas, for consistency-based method, H can be applied to any subset of data points. Our theory and proposed method remain the same whether H is applied on all the available data or not (see Appendix I).

3.1 Does the DeSSL risk estimator make sense?

The most intuitive interpretation is that by debiasing the risk estimator, we get back to the basics of learning theory. This way of debiasing is closely related to the method of control variates (Owen, 2013, Chapter 8) which is a common variance reduction technique. The idea is to add an additional term to a Monte-Carlo estimator with a null expectation in order to reduce the variance of the estimator without modifying the expectation. Here, DeSSL can also be interpreted as a control variate on the risk's gradient itself and should improve the optimisation scheme. This idea is close to the optimisation schemes introduced by Johnson & Zhang (2013) and Defazio et al. (2014) which reduce the variance of the gradients' estimate to improve optimisation performance.

Another interesting way to interpret DeSSL is as a constrained optimisation problem. Indeed, minimising $\hat{\mathcal{R}}_{DeSSL}$ is equivalent to minimising the Lagrangian of the following optimisation problem:

$$\begin{aligned} \min_{\theta} \quad & \hat{\mathcal{R}}_{CC}(\theta) \\ \text{s.t.} \quad & \frac{1}{n_u}\sum_{i=1}^{n_u} H(\theta; x_i) = \frac{1}{n_l}\sum_{i=1}^{n_l} H(\theta; x_i). \end{aligned} \tag{6}$$

The idea of this optimisation problem is to minimise the complete case risk estimator by assessing that some properties represented by H are on average equal for the labelled data and the unlabelled data. For example, if we consider entropy-minimisation, this program encourages the model to have the same confidence on the unlabelled examples as on the labelled ones.

The debiasing term of our objective will penalise the confidence of the model on the labelled data. Pereyra et al. (2017) actually show that penalising the entropy in a supervised context acts as a strong regulator for supervised model and improves on the state-of-the-art on common benchmarks. This comforts us in the idea of debiasing using labelled data in the case of entropy-minimisation. Similarly, the debiasing term in pseudo-label turns the problem into plausability inference as described by Barndorff-Nielsen (1976).

Our objective also resembles doubly-robust risk estimates used for SSL in the context of kernel machines by Liu & Goldberg (2020) and for deep learning in a recent preprint (Hu et al., 2022). In both cases, their focus is quite different, as they consider weaker conditions than MCAR, but very specific choices of H.

3.2 Is $\hat{\mathcal{R}}_{DeSSL}(\theta)$ an accurate risk estimate?

Because of the connections between our debiased estimate and variance reduction techniques, we have a natural interest in the variance of the estimate. Having a lower-variance estimate of the risk would mean estimating it more accurately, leading to better models. Similarly to traditional control variates (Owen, 2013), the variance can in fact be computed, and optimised in λ:

Theorem 3.1. *The function $\lambda \mapsto \mathbb{V}(\hat{\mathcal{R}}_{DeSSL}(\theta))$ reaches its minimum for:*

$$\lambda_{opt} = \frac{n_u}{n}\frac{\mathrm{Cov}(L(\theta; x, y), H(\theta; x))}{\mathbb{V}(H(\theta; x))}, \tag{7}$$

and at λ_{opt}:

$$\mathbb{V}(\hat{\mathcal{R}}_{DeSSL}(\theta))|_{\lambda_{opt}} = \left(1 - \frac{n_u}{n}\rho^2_{L,H}\right)\mathbb{V}(\hat{\mathcal{R}}_{CC}(\theta)) \leq \mathbb{V}(\hat{\mathcal{R}}_{CC}(\theta)), \tag{8}$$

where $\rho_{L,H} = \mathrm{Corr}(L(\theta; x, y), H(\theta; x))$.

A proof of this theorem is available as Appendix E. This theorem provides a formal justification to the heuristic idea that H should be a surrogate of L. Indeed, DeSSL is a more accurate risk estimate when H is strongly positively correlated with L, which is likely to be the case when H is equal or equivalent to an expectation of L. Then, choosing λ positive is a coherent choice. We also demonstrate in Appendix E that L and H are positively correlated when L is the negative likelihood and H is the entropy. Other SSL methods have variance reduction guarantees and already has shown great promises in SSL, see Fox-Roberts & Rosten (2014) and Sakai et al. (2017). In a purely supervised context, Chen et al. (2020) show that the effectiveness of data augmentation techniques lays partially on the variance reduction of the risk estimate. A natural application of this theorem would be to tune λ automatically by estimating λ_{opt}. In our case however, the estimation of $\mathrm{Cov}(L(\theta; x, y), H(\theta; x))$ with few labels led to extremely unstable unsatisfactory results.

3.3 Calibration

The calibration of a model is its capacity of predicting probability estimates that are representative of the true distribution. This property is determinant in real-world application when we need reliable predictions. A scoring rule $\mathcal{S}$ is a function assigning a score to the predictive distribution

$p_\theta(y|x)$ relative to the event $y|x \sim p(y|x)$, $\mathcal{S}(p_\theta, (x, y))$, where $p(x, y)$ is the true distribution (see e.g. Gneiting & Raftery, 2007). A scoring rule measures both the accuracy and the quality of predictive uncertainty, meaning that better calibration is rewarded. The expected scoring rule is defined as $\mathcal{S}(p_\theta, p) = \mathbb{E}_p[\mathcal{S}(p_\theta, (x, y))]$. A proper scoring rule is defined as a scoring rule such that $\mathcal{S}(p_\theta, p) \leq \mathcal{S}(p, p)$ (Gneiting & Raftery, 2007). The motivation behind having proper scoring rules comes from the following: suppose that the true data distribution p is accessible by our set of models. Then, the scoring rule encourages to predict $p_\theta = p$. The opposite of a proper scoring rule can then be used to train a model to encourage the calibration of predictive uncertainty: $L(\theta; x, y) = -\mathcal{S}(p_\theta, (x, y))$. Most common losses used to train models are proper scorings rule such as log-likelihood.

Theorem 3.2. *If $\mathcal{S}(p_\theta, (x, y)) = -L(\theta; x, y)$ is a proper scoring rule, then $\mathcal{S}'(p_\theta, (x, y, r)) = -(\frac{rn}{n_l} L(\theta; x, y) + \lambda n(\frac{1-r}{n_u} - \frac{r}{n_l}) H(\theta; x))$ is also a proper scoring rule.*

The proof is available in Appendix F, and follows directly from unbiasedness and the MCAR assumption. The main interpretation of this theorem is that we can expect DeSSL to be as well-calibrated as the complete case.

3.4 Consistency

We say that $\hat{\theta}$ is consistent if $d(\hat{\theta}, \theta^*) \xrightarrow{p} 0$ when $n \to \infty$, where d is a distance on Θ. The asymptotic properties of $\hat{\theta}$ depend on the behaviours of the functions L and H. We will thus require the following standard assumptions.

Assumption 3.3. The minimum θ^* of $\mathcal{R}$ is well-separated: $\inf_{\theta: d(\theta^*, \theta) \geq \epsilon} \mathcal{R}(\theta) > \mathcal{R}(\theta^*)$.

Assumption 3.4. The uniform weak law of large number holds for both L and H.

Theorem 3.5. *Under the MCAR assumption, Assumption 3.3 and Assumption 3.4, $\hat{\theta} = \arg\min \hat{\mathcal{R}}_{DeSSL}$ is consistent.*

For a proof of this theorem see Appendix F. This theorem is a simple application of van der Vaart's (2000) Theorem 5.7 proving the consistency of a M-estimator. Also, this results holds for the complete case, with $\lambda = 0$ which prove that the complete case is a solid baseline under the MCAR assumption.

Coupling of n_l and n_u under the MCAR assumption Under the MCAR assumption, n_l and n_u are random variables. We have that $r \sim \mathcal{B}(\pi)$ (i.e. any x has the probability π of being labelled). Then, with n growing to infinity, we have $\frac{n_l}{n} = \frac{n_l}{n_l + n_u} \to \pi$. Therefore, both n_l and n_u grow to infinity and $\frac{n_l}{n_u} \to \frac{\pi - 1}{\pi}$. This implies $n_u = \mathcal{O}(n_l)$ and then when n goes to infinity, both n_u and n_l go to infinity too and even if $n_u >> n_l$.

3.5 Rademacher complexity and generalisation bounds

In this section, we prove an upper bound for the generalisation error of DeSSL. The unbiasedness of $\hat{\mathcal{R}}_{DeSSL}$ can directly be used to derive generalisation bounds based on the Rademacher complexity (Bartlett & Mendelson, 2002), defined in our case as

$$R_n = \mathbb{E}_{(\varepsilon_i)_{i \leq n}} \left[\sup_{\theta \in \Theta} \left(\frac{1}{n_l} \sum_{i=1}^{n_l} \varepsilon_i L(\theta; x_i, y_i) - \frac{\lambda}{n_l} \sum_{i=1}^{n_l} \varepsilon_i H(\theta; x_i) + \frac{\lambda}{n_u} \sum_{i=1}^{n_u} \varepsilon_i H(\theta; x_i) \right) \right], \quad (9)$$

where ε_i are i.i.d. Rademacher variables independent of the data. In the particular case of $\lambda = 0$, we recover the standard Rademacher complexity of the complete case. We can then now bound the generalisation error of a model trained using our new loss function.

Theorem 3.6. *We assume that labels are MCAR and that both L and H are bounded. Then, there exists a constant $\kappa > 0$, that depends on λ, L, H, and the ratio of observed labels, such that, with probability at least $1 - \delta$, for all $\theta \in \Theta$,*

$$\mathcal{R}(\theta) \leq \hat{\mathcal{R}}_{DeSSL}(\theta) + 2R_n + \kappa \sqrt{\frac{\log(4/\delta)}{n}}. \quad (10)$$

The proof follows Shalev-Shwartz & Ben-David (2014, Chapter 26), and is available in Appendix H.

4 Experiments

We evaluate the performance of DeSSL against different classic methods. The goal here is to compare DeSSL methods and their orginal counterparts. In particular, we perform experiments with simple SSL methods such as pseudo-label (PseudoLabel) and entropy minimisation (EntMIN) with varying λ on MNIST (LeCun & Cortes, 2010) and CIFAR-10 and CIFAR-100 (Krizhevsky, 2009) and compare them to the debiased method, respectively DeEntMin and DePseudoLabel. We also compare PseudoLabel and DePseudoLabel on five small datasets of MedMNIST (Yang et al., 2021a;b) with a fixed λ. The results of these experiments are reported below. In our figures, the error bars represent the size of the 95% confidence interval (CI). Finally, we modified the implementation of Fixmatch (Sohn et al., 2020) and compare it with its debiased version on CIFAR-10.

We also compare DeEntMin and DePseudoLabel to the biased version on a large range of tabular datasets commonly used in SSL benchmarks (Chapelle et al., 2006; Guo et al., 2010). We do not observe differences between the performance, see Appendix N. Finally, we show how simple it is to debias an existing implementation, by demonstrating it on the consistency-based models benchmarked by (Oliver et al., 2018), namely VAT, Π-model and MeanTeacher on CIFAR-10 and SVHN (Netzer et al., 2011). We observe similar performances between the debiased and biased version for the different methods, both in terms of cross-entropy and accuracy. Moreover, these results have been obtained using the hyperparameters finetuned for the biased versions. Therefore, it is likely that optimising the hyperparameters for DeSSL will yield even better with the right hyperparameters, see Appendix M.

4.1 MNIST

MNIST is an advantageous dataset for SSL since classes are well-separeted. We compare PseudoLabel and DePseudoLabel for a LeNet-like architecture using $n_l = 1000$ labelled data on 10 different splits of the training dataset into a labelled and unlabelled set. Models are then evaluated using the standard $10,000$ test samples. We used 10% of n_l as the validation set. We test the influence of the hyperparameter λ and report the accuracy, the cross-entropy and the expected calibration error (ECE, Guo et al., 2017) at the epoch of best validation accuracy, see Figure 2 and Appendix J. In this exemple SSL and DeSSL have the almost the same accuracy for all λ, however, DeSSL seems to be alway better calibrated. In order to break the cluster assumption, we reproduced the same experiment on a modified MNIST. Indeed, we had label noise by replacing the true label for 20% of the dataset by a randomly sampled label, see Appendix J. In this setting, DeSSL performs better for large λ in term of accuracy and also provides a better calibration.

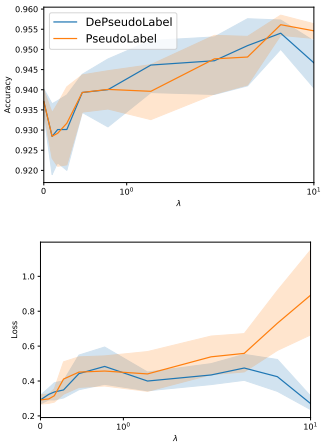

Figure 2: The influence of λ on Pseudo-label and DePseudo-label for a Lenet trained on MNIST with $n_l = 1000$: (Left) Mean test accuracy; (Right) Mean test cross-entropy, with 95% CI.

4.2 MedMNIST

We compare PseudoLabel and DePseudoLabel on different datasets of MedMNIST, a large-scale MNIST-like collection of biomedical images. We selected the five smallest 2D datasets of the collection, for these dataset it is likely that the cluster assumption no longer holds. We trained a 5-layer CNN with a fixed $\lambda = 1$ and n_l at 10% of the training data. We report in Table 1 the mean accuracy and cross-entropy on 5 different split of the labelled and unlabelled data and the number of labelled data used. We report the AUC in Appendix J. DePseudoLabel compete with PseudoLabel in terms of accuracy and even success when PseudoLabel's accuracy is less than the complete case. Moreover, DePseudoLabel is always better in term of cross-entropy, so calibration, whereas PseudoLabel is always worse than the complete case.

Table 1: Test accuracy and cross-entropy of Complete Case (CC), PseudoLabel (PL) and DePseudoLabel (DePL) on five datasets of MedMNIST.

DATASET	NL	CC		PL		DEPL	
		CROSS-ENTROPY	ACCURACY	CROSS-ENTROPY	ACCURACY	CROSS-ENTROPY	ACCURACY
DERMA	1000	1.95 ± 0.09	68.99 ± 1.20	2.51 ± 0.20	68.88 ± 1.03	$\mathbf{1.88 \pm 0.12}$	$\mathbf{69.30 \pm 0.85}$
PNEUMONIA	585	1.47 ± 0.04	83.94 ± 2.40	2.04 ± 0.04	$\mathbf{85.83 \pm 2.13}$	$\mathbf{1.40 \pm 0.06}$	84.36 ± 3.79
RETINA	160	1.68 ± 0.03	48.30 ± 3.06	1.80 ± 0.18	47.75 ± 2.50	$\mathbf{1.67 \pm 0.06}$	$\mathbf{49.40 \pm 2.62}$
BREAST	78	0.80 ± 0.04	76.15 ± 0.75	1.00 ± 0.26	74.74 ± 1.04	$\mathbf{0.70 \pm 0.03}$	$\mathbf{76.67 \pm 1.32}$
BLOOD	1700	$\mathbf{6.11 \pm 0.17}$	$\mathbf{84.13 \pm 0.83}$	6.61 ± 0.22	84.09 ± 1.17	6.53 ± 0.30	83.68 ± 0.59

4.3 CIFAR

We compare PseudoLabel and DePseudoLabel on CIFAR-10 and CIFAR-100. We trained a CNN-13 from Tarvainen & Valpola (2017) on 5 different splits. For this experiement, we use $n_l = 4000$ and use the rest of the dataset as unlabelled. Models are then evaluated using the standard $10,000$ test samples. For a more realistic validation set, we used 10% of n_l as the validation set. We test the influence of the hyperparameter λ and report the accuracy and the cross-entropy at the epoch of best validation accuracy, see Figure 3. We report the ECE in Appendix K. The performance of both methods on CIFAR-100 with $nl = 10000$ are reported in Appendix K. We observe DeSSL provides both a better cross-entropy and ECE with the same accuracy for small λ. For larger λ, DeSSL performs better in all the reported metrics. We performed a paired Student's t-test to ensure that our results are significant and reported the p-values in Appendix K. The p-values indicates that for λ close to 10, DeSSL is often significantly better in all the metrics. Moreover, DeSSL for large λ provides a better cross-entropy and ECE than the complete case whereas SSL never does.

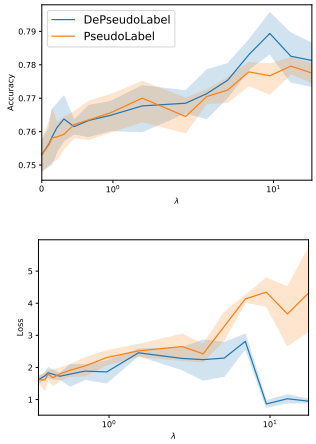

Figure 3: Influence of λ on Pseudo-label and DePseudo-label for a CNN trained on CIFAR with $n_l = 4000$: (Left) Mean test accuracy; (Right) Mean test cross-entropy, with 95% CI.

4.4 Fixmatch (Sohn et al., 2020)

We debiased a version of Fixmatch, see Appendix L for further details. For this experiment, we use $nl = 4000$ on 5 different folds.First, we report that a strong baseline using data augmentation reach 87.27% accuracy. Then, we observe that on the debiasing method improve both accuracy and cross-entropy of this modified version of Fixmatch. Inspired by Zhu et al. (2022), we show that our method improved performance on "poor" classes more equally than the biased version. Indeed, DeFixmatch improves Fixmatch by 1.57% overall but by 4.91% on the worst class. We report in Appendix L the accuracy per class of the different methods and the *benefit ratio* as defined by Zhu et al. (2022).

Table 2: 1st line: Accuracy, 2nd line: Worst class accuracy, 3rd line: Cross-entropy.

COMPLETE CASE	FIXMATCH	DEFIXMATCH
87.27 ± 0.25	93.87 ± 0.13	$\mathbf{95.44 \pm 0.10}$
70.08 ± 0.93	82.25 ± 2.27	$\mathbf{87.16 \pm 0.46}$
0.60 ± 0.01	0.27 ± 0.01	$\mathbf{0.20 \pm 0.01}$

5 Conclusion

Motivated by the remarks of van Engelen & Hoos (2020) and Oliver et al. (2018) on the missingness of theoretical guarantees in SSL, we proposed a simple modification of SSL frameworks. We consider frameworks based on the inclusion of unlabelled data in the computation of the risk estimator and debias them using labelled data. We show theoretically that this debiasing comes with several theoretical guarantees. We demonstrate these theoretical results experimentally on several common SSL datasets and some more challenging ones such as MNIST with label noise. DeSSL shows competitive performance in term of accuracy compared to its biased version but improves significantly the calibration. There are several future directions open to us. We showed that λ_{opt} exists (Theorem 3.1) and therefore our formula provides guidelines for the optimisation of λ. Finally, an interesting improvement would be to go beyond the MCAR assumption by considering settings with a distribution mismatch between labelled and unlabelled data (Guo et al., 2020; Cao et al., 2021; Hu et al., 2022).

References

Ahfock, D. and McLachlan, G. J. On missing label patterns in semi-supervised learning. *arXiv preprint arXiv:1904.02883*, 2019.

Aminian, G., Abroshan, M., Khalili, M. M., Toni, L., and Rodrigues, M. An information-theoretical approach to semi-supervised learning under covariate-shift. In *International Conference on Artificial Intelligence and Statistics*, pp. 7433–7449. PMLR, 2022.

Barndorff-Nielsen, O. Plausibility inference. *Journal of the Royal Statistical Society: Series B (Methodological)*, 38(2):103–123, 1976.

Bartlett, P. L. and Mendelson, S. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.

Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., and Raffel, C. A. Mixmatch: A holistic approach to semi-supervised learning. *Advances in Neural Information Processing Systems*, 2019.

Berthelot, D., Carlini, N., Cubuk, E. D., Kurakin, A., Sohn, K., Zhang, H., and Raffel, C. ReMix-Match: Semi-supervised learning with distribution matching and augmentation anchoring. *International conference on Learning Representations*, 2020.

Cao, K., Brbic, M., and Leskovec, J. Open-world semi-supervised learning, 2021.

Chapelle, O., Scholkopf, B., and Zien, A. Semi-supervised learning. *MIT Press*, 2006.

Chen, S., Dobriban, E., and Lee, J. H. A group-theoretic framework for data augmentation. *Journal of Machine Learning Research*, 21(245):1–71, 2020.

Corduneanu, A. and Jaakkola, T. On information regularization. In *UAI*. UAI, 2003.

Defazio, A., Bach, F., and Lacoste-Julien, S. SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives. *Advances in Neural Information Processing Systems*, 2014.

Fox-Roberts, P. and Rosten, E. Unbiased generative semi-supervised learning. *The Journal of Machine Learning Research*, 15(1):367–443, 2014.

Gan, H., Li, Z., Fan, Y., and Luo, Z. Dual learning-based safe semi-supervised learning. *IEEE Access*, 6:2615–2621, 2017.

Gneiting, T. and Raftery, A. E. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. *Advances in Neural Information Processing Systems*, 2014.

Goodfellow, I., Bengio, Y., and Courville, A. *Deep Learning*. MIT Press, 2016.

Grandvalet, Y. and Bengio, Y. Semi-supervised learning by entropy minimization. *Advances in Neural Information Processing Systems*, 2004.

Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. On calibration of modern neural networks. *International Conference on Machine Learning*, 2017.

Guo, L.-Z., Zhang, Z.-Y., Jiang, Y., Li, Y.-F., and Zhou, Z.-H. Safe deep semi-supervised learning for unseen-class unlabeled data. *International Conference on Machine Learning*, 2020.

Guo, Y., Niu, X., and Zhang, H. An extensive empirical study on semi-supervised learning. *IEEE International Conference on Data Mining*, 2010.

Hu, X., Niu, Y., Miao, C., Hua, X.-S., and Zhang, H. On non-random missing labels in semi-supervised learning. In *International Conference on Learning Representations*, 2022.

Johnson, R. and Zhang, T. Accelerating stochastic gradient descent using predictive variance reduction. *Advances in Neural Information Processing Systems*, 2013.

Kawakita, M. and Takeuchi, J. Safe semi-supervised learning based on weighted likelihood. *Neural Networks*, 53:146–164, 2014.

Kingma, D. P., Mohamed, S., Rezende, D. J., and Welling, M. Semi-supervised learning with deep generative models. In *Advances in Neural Information Processing Systems*, pp. 3581–3589, 2014.

Krause, A., Perona, P., and Gomes, R. Discriminative clustering by regularized information maximization. *Advances in neural information processing systems*, 23, 2010.

Krizhevsky, A. Learning multiple layers of features from tiny images. Technical report, MIT, NYU, 2009.

Laine, S. and Aila, T. Temporal ensembling for semi-supervised learning. *International Conference on Learning Representations*, 2017.

LeCun, Y. and Cortes, C. MNIST handwritten digit database. 2010. URL `http://yann.lecun.com/exdb/mnist/`.

Lee, D.-H. Pseudo-Label : The simple and efficient semi-supervised learning method for deep neural networks. *Workshop on challenges in representation learning, International conference on machine learning*, 2013.

Li, Y.-F. and Zhou, Z.-H. Towards making unlabeled data never hurt. *IEEE transactions on pattern analysis and machine intelligence*, 37:175–188, 2014.

Li, Y.-F., Kwok, J. T., and Zhou, Z.-H. Towards safe semi-supervised learning for multivariate performance measures. *AAAI Conference on Artificial Intelligence*, 2016.

Little, R. J. and Rubin, D. B. *Statistical Analysis with Missing Data*. John Wiley & Sons, 2019.

Liu, T. and Goldberg, Y. Kernel machines with missing responses. *Electronic Journal of Statistics*, 14:3766–3820, 2020.

McLachlan, G. J. Estimating the linear discriminant function from initial samples containing a small number of unclassified observations. *Journal of the American statistical association*, 72:403–406, 1977.

Mey, A. and Loog, M. Improvability through semi-supervised learning: A survey of theoretical results. *arXiv preprint arXiv:1908.09574*, 2019.

Miyato, T., Maeda, S.-i., Koyama, M., and Ishii, S. Virtual adversarial training: A regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41:1979–1993, 2018.

Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. Reading digits in natural images with unsupervised feature learning. 2011.

Newey, W. K. and McFadden, D. Large sample estimation and hypothesis testing. *Handbook of econometrics*, 4:2111–2245, 1994.

Oliver, A., Odena, A., Raffel, C., Cubuk, E. D., and Goodfellow, I. J. Realistic evaluation of deep semi-supervised learning algorithms. *Advances in Neural Information Processing Systems*, 2018.

Owen, A. B. *Monte Carlo theory, methods and examples*. 2013.

Pereyra, G., Tucker, G., Chorowski, J., Kaiser, Ł., and Hinton, G. Regularizing neural networks by penalizing confident output distributions. *Workshop track, International Conference on Learning Representations*, 2017.

Pham, H., Dai, Z., Xie, Q., and Le, Q. V. Meta pseudo labels. *Conference on Computer Vision and Pattern Recognition*, 2021.

Rizve, M. N., Duarte, K., Rawat, Y. S., and Shah, M. In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. *International Conference on Learning Representations*, 2021.

Sajjadi, M., Javanmardi, M., and Tasdizen, T. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. *Advances in Neural Information Processing Systems*, 2016.

Sakai, T., Plessis, M. C., Niu, G., and Sugiyama, M. Semi-supervised classification based on classification from positive and unlabeled data. *International conference on machine learning*, 2017.

Schölkopf, B., Janzing, D., Peters, J., Sgouritsa, E., Zhang, K., and Mooij, J. On causal and anticausal learning. *Internation conference on machine learning*, 2012.

Scudder, H. Probability of error of some adaptive pattern-recognition machines. *IEEE Transactions on Information Theory*, 11:363–371, 1965.

Seeger, M. Learning with labeled and unlabeled data. *Technical report*, 2000.

Shalev-Shwartz, S. and Ben-David, S. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge university press, 2014.

Singh, A., Nowak, R., and Zhu, J. Unlabeled data: Now it helps, now it doesn't. *Advances in Neural Information Processing Systems*, 2008.

Sohn, K., Berthelot, D., Li, C.-L., Zhang, Z., Carlini, N., Cubuk, E. D., Kurakin, A., Zhang, H., and Raffel, C. FixMatch: Simplifying semi-supervised learning with consistency and confidence. *Avances in Neural Information Processing Systems*, 2020.

Sokolovska, N., Cappé, O., and Yvon, F. The asymptotics of semi-supervised learning in discriminative probabilistic models. In *International Conference on Machine Learning*, 2008.

Tarvainen, A. and Valpola, H. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advancer in Neural Information Processing Systems*, 2017.

Trapp, M., Madl, T., Peharz, R., Pernkopf, F., and Trappl, R. Safe semi-supervised learning of sum-product networks. *Conference on Uncertainty in Artificial Intelligence*, 2017.

Tsiatis, A. A. *Semiparametric theory and missing data*. Springer, 2006.

Tsuchiya, T., Charoenphakdee, N., Sato, I., and Sugiyama, M. Semisupervised ordinal regression based on empirical risk minimization. *Neural Computation*, 33:3361–3412, 2021.

van der Vaart, A. W. *Asymptotic statistics*. Cambridge university press, 2000.

van Engelen, J. E. and Hoos, H. H. A survey on semi-supervised learning. *Machine Learning*, 109: 373–440, 2020.

Verma, V., Kawaguchi, K., Lamb, A., Kannala, J., Bengio, Y., and Lopez-Paz, D. Interpolation consistency training for semi-supervised learning. *International Joint Conference on Artificial Intelligence*, 2019.

Wei, C., Shen, K., Chen, Y., and Ma, T. Theoretical analysis of self-training with deep networks on unlabeled data. In *International Conference on Learning Representations*, 2021.

Xie, Q., Dai, Z., Hovy, E., Luong, M.-T., and Le, Q. V. Unsupervised data augmentation for consistency training. *Advances in Neural Information Processing Systems*, 2019.

Yang, J., Shi, R., and Ni, B. MedMNIST classification decathlon: A lightweight AutoML benchmark for medical image analysis. In *IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pp. 191–195, 2021a.

Yang, J., Shi, R., Wei, D., Liu, Z., Zhao, L., Ke, B., Pfister, H., and Ni, B. MedMNIST v2: A large-scale lightweight benchmark for 2D and 3D biomedical image classification. *arXiv preprint arXiv:2110.14795*, 2021b.

Zhang, B., Wang, Y., Hou, W., Wu, H., Wang, J., Okumura, M., and Shinozaki, T. FlexMatch: Boosting semi-supervised learning with curriculum pseudo labeling. *Advances in Neural Information Processing Systems*, 2021a.

Zhang, H., Cisse, M., Dauphin, Y. N., and Lopez-Paz, D. mixup: Beyond empirical risk minimization. *Internation Conference on Learning Representations*, 2017.

Zhang, S., Wang, M., Liu, S., Chen, P.-Y., and Xiong, J. How unlabeled data improve generalization in self-training? a one-hidden-layer theoretical analysis. In *International Conference on Learning Representations*, 2021b.

Zhu, X., Ghahramani, Z., and Lafferty, J. D. Semi-supervised learning using Gaussian fields and harmonic functions. *International conference on machine learning*, 2003.

Zhu, Z., Luo, T., and Liu, Y. The rich get richer: Disparate impact of semi-supervised learning. In *International Conference on Learning Representations*, 2022.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See section 5
   (c) Did you discuss any potential negative societal impacts of your work? [N/A] Theoretical work
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes] See Appendices
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] code and instructions to run Fixmatch.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Appendices
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] Confidence intervals on all figures.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] An estimation in Appendix
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] See Appendix
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] Code
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]