

555 **A Notation**

556 For two vectors $p, q \in \mathbf{R}^d$ we denote by $p \cdot q = \sum_{i=1}^d p_i q_i$ their inner product. We use $p * q$ to denote
 557 their element-wise product, i.e., $(p * q)_i = p_i q_i$. We use the notation $\max_i p$ to denote the i -th largest
 558 element of the vector p . We use $\text{margin}(p)$ to denote the difference between the top-2 elements of p ,
 559 i.e., $\text{margin}(p) = \max_1 p - \max_2 p$. Moreover, we use $\text{margin}_k(p)$ to denote the top-k margin, i.e.,
 560 $\text{margin}_k(p) = \sum_{i=1}^k \max_i p - \max_{k+1} p$. Given a function $f(w) : \mathbf{R}^d \mapsto \mathbf{R}$ we denote by $\partial_w f(w)$
 561 the gradient of f with respect to the parameter w .

562 **B Detailed Description of SLaM**

563 **B.1 Estimating the Teacher’s Accuracy Parameters: $\alpha(x), k(x)$**

564 **Estimating the Teacher’s Accuracy $\alpha(x)$ via Isotonic Regression** We now turn our attention to
 565 the problem of estimating $\alpha(x)$ for each x of dataset B , i.e., the dataset labeled by the teacher model.
 566 In [27] the authors empirically observed that $\alpha(x)$ correlates with metrics of teacher’s confidence
 567 such as the “margin”, i.e., the difference between the probabilities assigned in the top-1 class and
 568 the second largest class according to the teacher’s soft label y_s . In particular, the larger the margin
 569 is the more likely is that the corresponding teacher label is correct. We exploit (and enforce) this
 570 monotonicity by employing isotonic regression on a small validation dataset to learn the mapping
 571 from the teacher’s margin at an example x to the corresponding teacher’s accuracy $\alpha(x)$.

572 To perform this regression task we use a small validation dataset V with correct labels that the
 573 teacher has not seen during training. For every example $x \in V$ we compute the corresponding
 574 soft-teacher label $y_s(x)$ and compute its margin $\text{margin}(x) = \max_1(y_s(x)) - \max_2(y_s(x))$. For
 575 every $x \in V$ we also compute the hard-prediction of the teacher and compare it with the ground-
 576 truth, i.e., for every x the covariate and response pair is $(\text{margin}(x), 1 - \text{err}(g(x), y(x)))$. We
 577 then use isotonic regression to fit a piecewise constant, increasing function to the data. Sorting
 578 the regression data $\{(\text{margin}(x), 1 - \text{err}(g(x), y(x)))x \in V\}$ by increasing margin to obtain a list
 579 $(c^{(1)}, \dots, r^{(1)}), \dots, (c^{(m)}, r^{(m)})$, isotonic regression solves the following task

$$\min_{\hat{r}^{(1)}, \dots, \hat{r}^{(m)}} \sum_{i=1}^m (r^{(i)} - \hat{r}^{(i)})^2$$

$$\text{subject to } \text{lb} \leq \hat{r}^{(i)} \leq \hat{r}^{(i+1)} \leq 1,$$

580 where the parameter lb is a lower bound on the values $\hat{r}^{(i)}$ and is a hyper-parameter that we tune.
 581 On the other hand, the upper bound for the values can be set to 1 since we know that the true value
 582 $\alpha(x)$ is at most 1 for every x (since it corresponds to the probability that the teacher-label is correct).
 583 After we compute the values $\hat{r}^{(1)}, \dots, \hat{r}^{(m)}$ for any given $c \in [0, 1]$ the output of the regressor is the
 584 value of $\hat{r}^{(i)}$ corresponding to the smallest $c^{(i)}$ that is larger-than or equal to c . This is going to be
 585 our estimate for $\alpha(x)$. We remark that finding the values $r^{(i)}$ can be done efficiently in $O(n)$ time
 586 after sorting the data (which has a runtime of $O(n \log n)$) so the whole isotonic regression task can
 587 be done very efficiently.

588 **Estimating $k(x)$.** We now describe our process for estimating the values of $\alpha(x)$ and $k(x)$ for
 589 every example of dataset B . Similarly to the binary classification setting, we estimate the accuracy
 590 probability $\alpha(x)$ using isotonic regression on a small validation dataset. The value of $k(x)$ can be set
 591 to be equal to a fixed value of k for all data, so that the top-k accuracy of the teacher on the validation
 592 data is reasonable (say above 60%). For example, in our ImageNet experiments, we used $k = 5$. We
 593 also provide a data-dependent method to find different values $k(x)$ for every example x . To do this
 594 we adapt the method for estimating the top-1 accuracy $\alpha(x)$ of the teacher from the validation dataset.
 595 For every value of $k = 2, \dots, L - 1$ we compute the top-k margin of the teacher’s predictions on the
 596 validation data which is equal to the sum of the top-k probabilities of the teacher soft-label minus the
 597 probability assigned to the $k + 1$ -th class, i.e.,

$$\text{margin}_k(y_s(x)) = \left(\sum_{i=1}^k \max_i y_s(x) \right) - \max_{k+1} y_s(x).$$

598 Using the top-k margin as the covariate and the top-k accuracy as the response we solve the cor-
 599 responding regression task using isotonic regression to obtain the value $\alpha_k(x)$ representing the
 600 probability that the true label belongs in the top-k predictions of the teacher soft-label. For some
 601 threshold, say 90%, for every x we set $k(x)$ to be the smallest value of k so that $\alpha_k(x) \geq 90\%$. We
 602 empirically observed that using larger thresholds for the top-k accuracy (e.g., 90% or 95%), is better.
 603 We remark that while using the top-k margin as the covariate in the regression task is reasonable, our
 604 method can be used with other “uncertainty metrics” of the teacher’s soft-labels, e.g., the entropy of
 605 the distribution of $y_s(x)$ after grouping together the top-k elements. The higher this entropy metric is
 606 the more likely that the top-k accuracy probability $\alpha(x)_k$ of the teacher is low.

607 B.2 SLaM for Distillation with Unlabeled Examples: Pseudocode

608 In this section we present pseudo-code describing the distillation with unlabeled examples setting
 609 and the SLaM method, Algorithm 1.

610 *Remark B.1.* We remark that in our experiments, we observed that not normalizing the mixing
 611 operation with $k(x) - 1$ resulted in better results overall. Therefore, the mixing operation used in our
 612 experimental evaluation of SLaM is $\text{mix}(f(x; w); \alpha(x), k(x)) = \alpha(x)f(x; w) + (1 - \alpha(x))(1 -$
 613 $f(x; w)) * \text{top}(y_s(x); k(x))$. For more details we refer the reader to the code provided in the
 614 supplementary material.

Algorithm 1 Student Label Mixing (SLaM) Distillation

Input: Labeled Dataset A, Labeled Validation dataset V, Unlabeled Dataset U

Output: A trained Student model $f(x; w)$

Train Teacher model on Labeled Dataset A
 Pre-train Student model on Labeled Dataset A

Label examples of Dataset U using the Teacher
 $B \leftarrow \emptyset$
for each $x \in U$ **do**
 Add $(x, y_s(x))$ to B # For hard-distillation use $y(x)$
end for

Learn Teacher Accuracy Statistics $\alpha(x), k(x)$ Algorithm 2
 $\hat{\alpha}(x), \hat{k}(x) \leftarrow \text{LearnAccuracyStatistics}(y(\cdot), V, B)$
 Train student $f(x; w)$ using the SLaM loss:

$$\sum_{(x,y) \in A \cup V} \ell(y, f(x; w)) + \sum_{(x,y) \in B} \ell(y, \text{mix}(f(x; w); \hat{\alpha}(x), \hat{k}(x)))$$

615 C SLaM Consistency

616 In the following proposition we show that any minimizer of the SLaM loss over the noisy teacher-data
 617 must agree with the ground-truth for all x (that have positive density). To keep the presentation
 618 simple and avoid measurability issues (e.g., considering measure zero sets under X) in the following
 619 we will assume that the example distribution X is supported on a finite set. We remark that one
 620 can easily adapt the proof to hold for any distribution X (but the result will hold after excluding
 621 measure-zero sets under X).

622 **Proposition C.1** (SLaM Consistency). *Let D be the distribution of the teacher-labeled examples*
 623 *of dataset B, i.e., we first draw $x \sim X$ and then label it using the noisy teacher of Definition 3.2*
 624 *Moreover, assume that there exists some parameter $w^* \in \mathcal{W}$ such that the ground-truth $g(x) =$*
 625 *$f(x; w^*)$. Denote by $\mathcal{L}^{\text{SLaM}}(w) = \mathbf{E}_{(x,y) \sim D}[\ell(y, \text{mix}(f(x; w); \alpha(x), k(x)))]$, the SLaM objective.*
 626 *The following hold true.*

627 1. w^* minimizes the SLaM objective.

Algorithm 2 Estimating Teacher’s Accuracy Statistics $\alpha(x), k(x)$

Input: (Noisy) Teacher Model $y_s(x)$, Labeled Validation dataset V , Isotonic-Regression lower-bound $lb \in [0, 1]$, and top-k accuracy threshold $t \in [0, 1]$.

Output: Estimates $\hat{\alpha}(x), \hat{k}(x)$ of the actual $\alpha(x), k(x)$.

Create Soft-labels for the Validation dataset using the teacher model $\{y_s(x) : x \in V\}$.

for $j = 1$ to $L - 1$ **do**

 # Map $y_s(x)$ to top-j margin and accuracy pairs on the Validation V

$$C \leftarrow \left\{ \left(\sum_{r=1}^j \max_r y_s(x) - \max_{j+1} y_s(x), 1 - \text{err}(y_s(x), z) \right) : (x, z) \in V \right\}.$$

Set $\hat{\alpha}_j(x)$ to be the output of Isotonic-Regression with lower-bound lb on the (covariate, response) pairs in C . # See Appendix [B.1](#)

end for

$\hat{\alpha}(x) \leftarrow \hat{\alpha}_1(x)$

$\hat{\alpha}_L(x) \leftarrow 1$ # The top-L accuracy is always (trivially) equal to 1

Given example x for some threshold t set $\hat{k}(x)$ to be the smallest integer $r \in \{1, \dots, L\}$ so that $\alpha_r(x) \geq t$.

628 2. Assuming further that for all x it holds that $\alpha(x)k(x) \neq 1$, we have that any minimizer w of
629 the SLaM objective satisfies: $f(x; w) = g(x)$ for all x .

630 *Proof.* Fix any example $x \in X$. By Definition [3.2](#) we have that the corresponding teacher label y is
631 correct with probability $\alpha(x)$ and a uniformly random incorrect label out of the top-k labels according
632 to the teacher soft-label $y_s(x)$. Recall for an L -dimension score vector p , by $\text{top}(p; k) \in \{0, 1\}^L$ we
633 denote the vector that has 1 on the positions of the top-k elements of p , e.g., $\text{top}((1, 2, 3, 4, 5); 2) =$
634 $(0, 0, 0, 1, 1)$. Conditional on x , the corresponding expected noisy teacher label is

$$\begin{aligned} \mathbf{E}[y \mid x] &= \mathbf{P}[y = g(x) \mid x]g(x) + \mathbf{P}[y \neq g(x) \mid x] \mathbf{E}[y \mid x, y \neq g(x)] \\ &= \alpha(x)g(x) + (1 - \alpha(x)) \mathbf{E}[y \mid y \neq g(x), x]. \end{aligned}$$

635 We know that the expected teacher label conditional on it being wrong $\mathbf{E}[y \mid y \neq g(x), x]$ is a
636 uniformly random incorrect label from the top-k labels of the corresponding $y_s(x)$.
637 Assume first that $k = L$, since the ground-truth is represented by a one-hot vector, the distribution
638 of uniformly random incorrect labels conditional on x can be written as $(1 - g(x))/(L - 1)$. For
639 example, if the ground-truth label is $g(x) = (1, 0, 0, 0, 0)$ then a uniformly random incorrect label
640 has probability distribution $(0, 1/4, 1/4, 1/4, 1/4)$. Assume now that $k(x) = 3$ and $\text{top}(y_s(x); 3) =$
641 $(1, 1, 1, 0, 0)$. Then the distribution of the (incorrect) teacher label becomes $(0, 1/2, 1/2, 0, 0)$. Using
642 $*$ to denote element-wise multiplication of two vectors, we have

$$\mathbf{E}[y \mid x, y \neq g(x)] = \frac{1 - g(x)}{k(x) - 1} * \text{top}(y_s(x); k(x))$$

643 Therefore, we obtain

$$\mathbf{E}[y \mid x] = \alpha(x)g(x) + (1 - \alpha(x)) \frac{1 - g(x)}{k(x) - 1} * \text{top}(y_s(x); k(x)) = \text{mix}(g(x); \alpha(x), k(x)).$$

644 Therefore, by using the fact that Cross-Entropy is linear in its first argument, we obtain that the
645 expected SLaM loss on some example x is

$$\begin{aligned} \mathbf{E}[\text{ce}(y, \text{mix}(f(x; w); \alpha(x), k(x))) \mid x] &= \text{ce}(\mathbf{E}[y \mid x], \text{mix}(f(x; w); \alpha(x), k(x))) \\ &= \text{ce}(\text{mix}(g(x; w); \alpha(x), k(x)), \text{mix}(f(x; w); \alpha(x), k(x))). \end{aligned}$$

646 We first have to show that there exist some parameter $w \in \mathcal{W}$ that matches the (expected) ob-
647 served labels $\mathbf{E}[y \mid x]$. Observe first that by using the realizability assumption, i.e., that there
648 exists w^* so that $f(x; w^*) = g(x)$ we obtain that, for every x , it holds $\text{mix}(g(x); \alpha(x), k(x)) =$
649 $\text{mix}(f(x; w^*); \alpha(x), k(x))$. In fact, by Gibb’s inequality (convexity of Cross-Entropy) we have that
650 w^* is a (global) minimizer of the SLaM objective.

651 We next show that *any (global) minimizer* of the SLaM objective must agree with the ground-truth
652 for every x . Since we have shown that w^* is able to match the (expected) labels $\mathbf{E}[y \mid x]$ any
653 other minimizer w must also satisfy $\text{mix}(g(x); \alpha(x), k(x)) = \text{mix}(f(x; w); \alpha(x), k(x))$. Assume
654 without loss of generality that $g_0 = 1$, i.e., the ground-truth label is 0. We observe that by using that
655 $\text{mix}(g(x; w); \alpha(x), k(x)) = \alpha(x)g(x) + (1 - \alpha(x))\frac{1-g(x)}{k(x)-1} * \text{top}(y_s(x); k(x))$ and the fact that the
656 ground-truth belongs in the top- $k(x)$ of the teacher’s predictions conditional that the teacher’s top-1
657 prediction is incorrect (thus $\text{top}(y_s(x))_0 = 1$), we obtain that

$$\alpha(x)g_0(x) + (1 - \alpha(x))(1 - g_0(x)) / (1 - k(x)) = \alpha(x)f(x; w)_0 + (1 - \alpha(x))(1 - f(x; w)_0) / (k(x) - 1).$$

658 Using the fact that $g_0 = 1$ we can simplify the above expression to

$$(1 - f(x; w)_0) \left(\alpha(x) - \frac{1 - \alpha(x)}{k(x) - 1} \right) = 0.$$

659 Using the assumption that $\alpha(x)k(x) \neq 1$ we obtain that the term $\left(\alpha(x) - \frac{1 - \alpha(x)}{k(x) - 1} \right)$ is not vanishing
660 and therefore it must hold that $f(x; w)_0 = 1 = g_0$, i.e., the student model must be equal to the
661 ground-truth.

662

□

663 D Extended Experimental Evaluation

664 We implemented all algorithms in Python and used the TensorFlow deep learning library [1]. We ran
665 our experiments on 64 Cloud TPU v4s each with two cores.

666 D.1 Implementation Details: Vision Datasets

667 Here we present the implementation details for the vision datasets we considered.

668 *Remark D.1.* We note that in all our experiments, “VID” corresponds to the implementation of the
669 loss described in equation (2), (4) and (6) of [2] (which requires appropriately modifying the student
670 model so that we have access to its embedding layer).

671 **Experiments on CIFAR-10/100 and CelebA** For the experiments on CIFAR-10/100 and CelebA
672 we use the Adam optimizer with initial learning rate $\text{lr} = 0.001$. We then proceed according to the
673 following learning rate schedule (see, e.g., [25]):

$$\text{lr} \leftarrow \begin{cases} \text{lr} \cdot 0.5 \cdot 10^{-3}, & \text{if } \#\text{epochs} > 180 \\ \text{lr} \cdot 10^{-3}, & \text{if } \#\text{epochs} > 160 \\ \text{lr} \cdot 10^{-2}, & \text{if } \#\text{epochs} > 120 \\ \text{lr} \cdot 10^{-1}, & \text{if } \#\text{epochs} > 80 \end{cases}$$

674 Finally, we use data-augmentation. In particular, we use random horizontal flipping and random
675 width and height translations with width and height factor, respectively, equal to 0.1.

676 The hyperparameters of each method are optimized as follows. For SLaM we always use 0.5 as
677 the lower bound for isotonic regression (i.e., the parameter lb in Algorithm 2). As CelebA is
678 a binary classification benchmark $k(x)$ is naturally set to 2 for all examples. For CIFAR-10/100
679 we used the data-dependent method for estimating $k(x)$ (see Algorithm 2) with threshold pa-
680 rameter $t = 0.9$. For weighted distillation we do a grid search over updating the weights ev-
681 ery $\{1, 25, 50, 100, 200\}$ epochs and we report the best average accuracy achieved. Finally, for
682 VID we search over $\{0.001, 0.1, 0.2, 0.5, 0.8, 1.0, 2.0, 10.0, 50.0, 100.0\}$ for the coefficient of the
683 VID-related term of the loss function, and for the PolyLoss we optimize its hyperparameter over
684 $\{-1.0, -0.8, -0.6, -0.4, -0.2, 0.5, 1.0, 2.0, 50.0, 100.0\}$.

685 **Experiments on ImageNet** For the ImageNet experiments we use SGD with momentum 0.9 as the
686 optimizer. For data-augmentation we use random horizontal flipping and random cropping. Finally,
687 the learning rate schedule is as follows. For the first 5 epochs the learning rate lr is increased from




Figure 3: Comparison of distillation methods on ImageNet. On the horizontal axis we plot the size of Dataset A as a percentage of the whole training dataset. On the vertical axis we plot the accuracy of the trained student-model on the test dataset.

Table 5: Experiments on CIFAR-10 (hard-distillation). See Section 4.2 for details.

Labeled Examples	5000	7500	10000	12500	15000	17500
Teacher	61.30	68.98	72.42	73.92	76.63	78.63
Vanilla	62.26 ± 0.45	69.07 ± 0.11	72.09 ± 0.11	73.43 ± 0.16	75.93 ± 0.25	77.43 ± 0.15
Taylor-CE [20]	63.14 ± 0.07	69.98 ± 0.11	72.72 ± 0.36	73.77 ± 0.28	76.26 ± 0.29	77.88 ± 0.20
UPS [48]	64.27 ± 0.08	70.93 ± 0.26	73.78 ± 0.16	74.66 ± 0.29	77.38 ± 0.37	78.95 ± 0.08
VID [3]	61.95 ± 0.22	66.91 ± 0.21	69.59 ± 0.24	72.16 ± 0.47	74.83 ± 0.11	75.55 ± 0.21
Weighted [27]	63.22 ± 0.45	71.04 ± 0.26	72.84 ± 0.12	74.20 ± 0.16	76.56 ± 0.24	78.23 ± 0.15
SLaM (Ours)	66.40 ± 0.31	72.44 ± 0.17	74.77 ± 0.13	75.64 ± 0.19	77.99 ± 0.36	79.26 ± 0.26

688 0.0 to 0.1 linearly. After that, the learning rate changes as follows:

$$lr = \begin{cases} 0.01, & \text{if } \#epochs > 30 \\ 0.001, & \text{if } \#epochs > 60 \\ 0.0001, & \text{if } \#epochs > 80. \end{cases}$$

689 The hyperparameters of each method are optimized as follows. For SLaM we do a hyperparameter
 690 search over $\{0.55, 0.60, 0.65, 0.70\}$ for the lower bound for isotonic regression, and we keep the best
 691 performing value for each potential size of dataset A . We used the fixed value 5 for $k(x)$, as the
 692 top-5 accuracy of the teacher model was satisfactory (much higher than its top-1 accuracy) on the
 693 validation dataset. For Taylor-CE we did a hyper-parameter search for the Taylor series truncation
 694 values in $\{1, 2, 3, 4, 5, 6, 10, 20, 50, 80, 100\}$. For weighted distillation we compute the weights in a
 695 one-shot fashion using the pre-trained student (as in the ImageNet experiments in [27]). For VID we
 696 search over $\{0.1, 0.3, 0.5\}$ for the coefficient of the VID-related term of the loss function, and for the
 697 PolyLoss we optimize its hyperparameter over $\{1.0, 2.0, 50.0, 100.0\}$.

698 D.2 Hard-Distillation

699 Here we present results on hard-distillation. The hyper-parameters of all methods are chosen the
 700 same way as in our soft-distillation experiments, see Appendix D.1. Tables 5, 6 and 7 contain
 701 our results on CIFAR-10, CIFAR-100 and CelebA, respectively. We observe that in almost all
 702 cases, SLaM consistently outperforms the other baselines. Moreover, for CIFAR-10 and CIFAR-100
 703 hard-distillation performs worse than soft-distillation (as it is typical the case) but in CelebA hard-
 704 distillation seems to be performing on par with (sometimes even outperforming) soft-distillation. A
 705 plausible explanation for the latter outcome is that in our CelebA experiments the teacher and student
 706 have different architectures (MobileNet and ResNet, respectively) so that soft-labels from the teacher
 707 are not so informative for the student. (This is also a binary classification task where the information
 708 passed from the teacher to the student through its soft-labels is limited.)

Table 6: Experiments on CIFAR-100 (**hard**-distillation). See Section 4.2 for details.

Labeled Examples	5000	7500	10000	12500	15000	17500
Teacher	35.97	44.65	49.62	55.68	59.19	62.05
Vanilla	36.36 ± 0.04	44.15 ± 0.10	50.22 ± 0.07	55.55 ± 0.24	58.85 ± 0.1	61.43 ± 0.19
Taylor-CE [20]	39.12 ± 0.14	46.87 ± 0.10	52.64 ± 0.22	57.19 ± 0.28	59.95 ± 0.11	62.36 ± 0.21
UPS [48]	39.49 ± 0.13	48.36 ± 0.44	53.95 ± 0.10	57.95 ± 0.10	60.59 ± 0.29	62.09 ± 0.28
VID [8]	37.19 ± 0.09	44.67 ± 0.16	50.63 ± 0.35	54.78 ± 0.07	59.27 ± 0.14	62.01 ± 0.05
Weighted [27]	38.04 ± 0.29	46.45 ± 0.22	52.33 ± 0.18	57.43 ± 0.13	60.81 ± 0.09	63.02 ± 0.06
SLaM (Ours)	42.01 ± 0.29	49.08 ± 0.14	54.49 ± 0.17	58.53 ± 0.04	61.12 ± 0.15	63.21 ± 0.18

Table 7: Experiments on CelebA (**hard**-distillation). See Section 4.2 for details.

Labeled Examples	2%	3%	4%	5%	6%	7%
Teacher	86.19	88.25	88.95	91.31	92.09	92.62
Vanilla	89.73 ± 0.08	91.61 ± 0.09	92.05 ± 0.11	93.41 ± 0.13	94.02 ± 0.15	94.05 ± 0.04
Taylor-CE [20]	90.62 ± 0.05	92.19 ± 0.02	92.66 ± 0.11	93.60 ± 0.14	94.00 ± 0.04	94.38 ± 0.10
UPS [48]	89.35 ± 0.04	91.30 ± 0.04	91.95 ± 0.12	93.18 ± 0.07	93.71 ± 0.04	94.18 ± 0.03
VID [8]	89.92 ± 0.21	91.60 ± 0.11	92.20 ± 0.12	93.51 ± 0.15	94.08 ± 0.15	94.27 ± 0.10
Weighted [27]	90.06 ± 0.06	91.97 ± 0.13	92.45 ± 0.10	93.60 ± 0.07	93.94 ± 0.12	94.25 ± 0.16
SLaM (Ours)	90.43 ± 0.05	92.25 ± 0.11	92.71 ± 0.08	93.96 ± 0.17	94.39 ± 0.21	94.52 ± 0.12




Figure 4: Comparison of distillation methods on CIFAR-10,100 and CelebA. On the horizontal axis we plot the size of Dataset A as a percentage of the whole training dataset. On the vertical axis we plot the accuracy of the trained student-model on the test dataset.

709 D.3 Large Movies Reviews Dataset Results

710 Here we present the results and the implementation details regarding the experiments on the Large
 711 Movies Reviews dataset. Recall that we use an ALBERT-large model as a teacher, and an ALBERT-
 712 base model as a student. We also use 2%, 4%, 8%, 40% percent (or 500, 1000, 2000, 10000 examples)
 713 from the training dataset and split the remaining data in a validation dataset of 500 examples and
 714 an unlabeled dataset U. We compare the methods on the soft-distillation. For each trial we train the
 715 student model for 40 epochs and keep the best test accuracy over all epochs. We perform 3 trials and
 716 report the average of each method and the variance of the achieved accuracies over the trials. The
 717 results of our experiments can be found in Table 8. We remark that we did not implement the UPS
 718 method for this dataset as the data-augmentation method for estimating the teacher’s accuracy could
 719 not be readily used for this NLP dataset. Moreover, using dropout and Monte Carlo estimation for
 720 the uncertainty was also not compatible with the Albert model used in this experiment.

721 Since we are dealing with ALBERT-models (which are already pre-trained), we do not pre-train the
 722 student model on dataset A except in the case of “weighted-distillation” [27], where we pre-train the
 723 student model on dataset A just for 1 epoch. The teacher model is trained using the Adam optimizer
 724 for 20 epochs with initial learning rate 10^{-6} . The student model is trained also using the Adam
 725 optimizer but for 40 epochs and with learning rate 10^{-7} .

726 The hyperparameters of each method are optimized as follows. For SLaM we do a hyperparameter
 727 search over $\{0.5, 0.6, 0.7, 0.8, 0.9\}$ for the lower bound for isotonic regression, and we keep the best
 728 performing value for each potential size of dataset A. As this is a binary classification benchmark we
 729 naturally set $k(x) = 2$ for all examples. For weighted distillation we do a grid search over updating
 730 the weights every $\{1, 10, 20, 40\}$ epochs and, similarly, we report the best average accuracy achieved.
 731 Finally, for VID (recall also Remark D.1) we search over $\{0.1, 0.5, 1.0, 2.0\}$ for the coefficient of

Table 8: Experiments on the Large Movies Reviews Dataset (**soft**-distillation). See Section [D.3](#) for details.

Labeled Examples	2%	4%	8%	40%
Teacher	77.52	84.04	85.44	88.3
Vanilla	80.93 ± 0.10	85.12 ± 0.29	85.99 ± 0.08	87.50 ± 0.6
Taylor-CE [20]	79.5 ± 0.38	85.14 ± 0.13	85.98 ± 0.14	87.57 ± 0.3
VID [3]	81.76 ± 0.32	85.33 ± 0.35	86.17 ± 0.06	87.71 ± 0.01
Weighted [27]	81.1 ± +0.1	85.2 ± 0.05	86.13 ± 0.17	87.8 ± 0.25
SLaM (Ours)	81.88 ± 0.23	85.5 ± 0.09	86.23 ± 0.13	87.73 ± 0.38




Figure 5: Composability the fidelity-based weighting scheme of [\[17\]](#). The x -axis shows the different values of the fidelity hyper-parameter β and the size of dataset A . From left to right we increase the size of dataset A from 10% to 35% and for each size we try different values of β . We observe that SLaM on its own (shown in green) is usually much better than the fidelity weighting scheme (shown in orange). Moreover, using SLaM on top of the fidelity weighting scheme (shown in blue) consistently improves its performance.

732 the VID-related term of the loss function, and for the PolyLoss we optimize its hyperparameter over
 733 $\{-1.0, -0.8, -0.6, -0.4, -0.2, 0.5, 1.0, 2.0\}$.

734 D.4 Combining with Teacher-Uncertainty-Based Reweighting Techniques

735 As we discussed in Section [2](#) our method can in principle be combined with teacher-uncertainty filtering
 736 and weighting schemes as these can be seen as preprocessing steps. To demonstrate this, we combine
 737 our method with the so-called fidelity-based weighting scheme of [\[17\]](#). The fidelity weighting
 738 scheme reweights examples using some uncertainty measure for teacher’s labels, e.g., by performing
 739 random data-augmentations and estimating the variance of the resulting teacher labels or using dropout
 740 and Monte Carlo estimation. More precisely, for every example x in the teacher-labeled dataset B , the
 741 fidelity-weighting scheme assigns the weight $w^{\text{Fid}}(x) = \exp(-\beta \text{uncertainty}^{\text{teacher}}(x))$ for some
 742 hyper-parameter $\beta > 0$. In our experiments we performed 10 random data augmentations (random
 743 crop and resize), estimated the coordinate-wise variance of the resulting teacher soft-labels, and
 744 finally computed the average of the variances of the k -classes, as proposed in [\[17\]](#). We normalized
 745 the above uncertainty of each example by the total uncertainty of the teacher over the whole dataset
 746 B . The weights of examples in dataset A are set to 1 and the reweighted objective is optimized over
 747 the combination of the datasets A, B .

$$\mathcal{L}^{\text{fid}}(w) = \frac{1}{|A \cup B|} \left(\sum_{(x,y) \in A} \ell(y, f(x; w)) + \sum_{(x,y) \in B} w^{\text{Fid}}(x) \ell(y, f(x; w)) \right). \quad (3)$$

748 To demonstrate the composability of our method with such uncertainty-based weighting schemes,
 749 we use CIFAR100 and the percentage of the labeled dataset A (as a fraction of the whole training
 750 set) is 10%, 15%, 20%, 25%, 30%, 35%, similar to the setting of Section [4.2](#). The teacher is a
 751 ResNet10 and the student is a ResNet56. We first train the student using only the fidelity weighting
 752 scheme, i.e., optimize the loss function of Equation [\(4\)](#) using different values for the hyperparameter
 753 $\beta \in \{0.1, 0.2, 1.0, 1.2, 2.0, 5.0, 10.0, 20.0\}$, i.e., ranging from mildly reweighting the examples of




Figure 6: CIFAR100: Temperature Ablation. On the x-axis we have the size of the labeled dataset (as a percentage of the whole training dataset) that the teacher model uses for training.

754 dataset B to more aggressively “removing” examples where the teacher’s entropy is large. For the
 755 same values of β we then train the student using the reweighted SLaM objective:

$$\mathcal{L}^{\text{Fid+SLaM}}(w) = \frac{1}{|A \cup B|} \left(\sum_{(x,y) \in A} \ell(y, f(x; w)) + \sum_{(x,y) \in B} w^{\text{fid}}(x) \ell(y, \text{mix}(f(x; w); \alpha(x), k(x))) \right). \quad (4)$$

756 For the combined SLaM + Fidelity method we did not perform hyper-parameter search and used
 757 the same parameters for the isotonic regression as we did in the “standard” SLaM experiment in
 758 CIFAR100 of Appendix D.1. We present our comprehensive results for all sizes of dataset A and
 759 values of the hyper-parameter β in Figure 5. Our results show that, regardless of the value of
 760 the hyperparameter β and the size of the labeled dataset A, using SLaM together with the fidelity
 761 weighting scheme provides consistent improvements. Moreover, in Figure 5, we observe that by
 762 using SLaM the achieved accuracy depends less on the hyper-parameter β : since SLaM takes into
 763 account the fact that some of the teacher’s predictions are incorrect, it is not crucial to down-weight
 764 them or filter them out.

765 D.5 Using Distillation Temperature

766 In this section we show that our approach can be effectively combined with temperature-scaling [26].
 767 Choosing the right distillation temperature often provides significant improvements. In our setting,
 768 the teacher provides much more confident predictions (e.g., soft-labels with high-margin) on dataset
 769 A (where the teacher was trained) compared to the teacher soft-labels of dataset B where the
 770 teacher is, on average, less confident. Given this observation, it is reasonable to use different
 771 distillation temperatures for dataset A and dataset B. We try different temperatures for dataset A
 772 and dataset B and perform vanilla distillation with temperature and also consider applying the
 773 temperature scaling before applying SLaM. For each size of dataset A we try pairs of temperatures
 774 $t_A, t_B \in \{0.01, 0.1, 0.5, 0.8, 1., 2., 5., 10., 100.\}$ and report the best accuracy achieved by vanilla
 775 distillation and the best achieved by first applying temperature scaling and then SLaM. In Figure 6
 776 we observe that SLaM with temperature scaling consistently improves over vanilla distillation with
 777 temperature.

778 D.6 Using SLaM with other loss functions beyond cross-entropy

779 In this section, we demonstrate that our method can be successfully applied when the student loss
 780 function comes from the families of losses introduced in [20] and [35]. We perform experiments on
 781 CIFAR-100 and ImageNet following the setting of Section 4.2. In particular, we compare vanilla
 782 distillation with unlabeled examples using the Taylor-CE loss of [20] and the PolyLoss of [35], with
 783 combining SLaM with these losses. For the Taylor-CE loss we set the “degree” hyperparameter to
 784 be 2 (as suggested in [20]) and we set the hyperparameter of the PolyLoss to be 2.0 (as suggested
 785 in [35]). The corresponding results can be found in Figure 7.




Figure 7: Using SLaM with PolyLoss [35] and Taylor CE [20]. On the x-axis we have the size of the labeled dataset (as a percentage of the whole training dataset) that the teacher model uses for training. See Appendix D.6 for more details.

786 E Distilling Linear Models and Learning Noisy Halfspaces

787 In this section we state and prove our convergence result for the SLaM method when applied
 788 to linear models. Our assumption is that the ground-truth $g(x)$ corresponds to a halfspace, i.e.,
 789 $g(x) = (\mathbf{1}\{w^* \cdot x > 0\}, \mathbf{1}\{w^* \cdot x \leq 0\})$ for some unknown weight vector w^* . We show that
 790 using SLaM with a linear model as the student will recover the ground truth classifier. We make the
 791 standard assumption that the ground-truth halfspace has γ -margin, i.e., that $\|w^*\|_2 = 1$ and that it
 792 holds $|w^* \cdot x| \geq \gamma$ for all examples x . For a fixed example x , the observed noisy teacher-label y
 793 satisfies Definition 3.2, i.e., $y = g(x)$ w.p. $\alpha(x)$ and $y = 1 - g(x)$ w.p. $1 - \alpha(x)$ (since $k = 2$ for
 794 binary classification). Our approach consists of using the standard cross-entropy loss $\text{ce}(p, q)$ and
 795 training a student-model consisting of a linear layer plus a soft-max activation, i.e.,

$$f(x; w) = (f_0(x; w), f_1(x; w)) = \left(\frac{1}{1 + e^{-w \cdot x}}, \frac{e^{-w \cdot x}}{1 + e^{-w \cdot x}} \right).$$

796 Recall, that for binary classification, we define the mixing operation as

$$\text{mix}(f(x; w); \alpha(x)) = \alpha(x)f(x; w) + (1 - \alpha(x))(1 - f(x; w)).$$

Algorithm 3 SLaM for Linear Models

Initialize weight vector of student $w^{(0)} \leftarrow 0$

for $t = 1, \dots, T$ **do**

 Draw example $x^{(t)} \sim X$.

 Label $x^{(t)}$ with (noisy) teacher to obtain $y^{(t)}$

 Compute the gradient of the SLaM loss at $(x^{(t)}, y^{(t)})$:

$$g^{(t)} \leftarrow \partial_w \text{ce}(y^{(t)}, \text{mix}(f(x^{(t)}; w^{(t-1)}), \alpha(x^{(t)}))) \big|_{w=w^{(t-1)}}$$

 Compute step size: $\lambda^{(t)} \leftarrow 1/r(f(x^{(t)}; w^{(t-1)}), \alpha(x^{(t)}))$ (see Lemma E.3 for the definition of $r(\cdot, \cdot)$).

 Update the student model: $w^{(t)} \leftarrow w^{(t-1)} - \lambda^{(t)} g^{(t)}$

end for

797 **Theorem E.1** (Student Label Mixing Convergence). *Let X be a distribution on \mathbf{R}^d and $g(x)$ be*
 798 *the ground-truth halfspace with normal vector $w^* \in \mathbf{R}^d$. Let D be the distribution over (noisy)*
 799 *teacher-labeled examples (x, y) whose x -marginal is X . We denote by $\alpha(x)$ the probability that the*
 800 *teacher label $y \in [0, 1]^2$ is correct, i.e., $\alpha(x) = \mathbf{P}_{(x, y) \sim D}[\arg\max(y) = g(x) \mid x]$. Assume that*
 801 *there exist $\beta, \gamma > 0$ such that for all examples x in the support of X it holds that $|w^* \cdot x| \geq \gamma$ and*
 802 *$|1/2 - \alpha(x)| \leq \beta$. Let $\epsilon > 0$. After $T = O(1/(\beta^2 \gamma^2 \epsilon^2))$ iterations of SLaM (Algorithm 3), with*
 803 *probability at least 99%, there exists an iteration $t \leq T$ where $\mathbf{P}_{x \sim X}[\text{err}(f(x; w^{(t)}), g(x))] \leq \epsilon$.*

804 **Remark E.2** (High-Probability Result). We remark that even though our learner succeeds with
 805 constant probability (at least %99) we can amplify its success probability to $1 - \delta$ by standard
 806 amplification techniques (i.e., by repeating the algorithm $O(\log(1/\delta))$ times and keeping the best
 807 result). To achieve success probability $1 - \delta$ the total sample complexity is $O(\log(1/\delta)/(\epsilon^2 \gamma^2 \beta^2))$.

808 *Proof.* We first provide simplified expressions for the gradient of the SLaM objective and the update
809 vectors $\lambda^{(t)} g^{(t)}$ used in Algorithm 3. In what follows we remark that for any binary classification
810 model $f(x; w) = (f_0(x; w), f_1(x; w))$ we have the following identities: (i) $(\text{mix}(f(x; w); \alpha(x)))_0 =$
811 $\text{mix}(f_0(x; w); \alpha(x))$, where to simplify notation we overload the mixing operation to also act on
812 the scalar $f_0(x; w)$, i.e., $\text{mix}(f_0(x; w); \alpha(x)) = \alpha(x)f_0(x; w) + (1 - \alpha(x))(1 - f_0(x; w))$; and (ii)
813 $f_1(x; w) = 1 - f_0(x; w)$.

814 **Lemma E.3** (SLaM Gradient). *The gradient of the SLaM objective is equal to*

$$\partial_w \text{ce}(y, \text{mix}(f(x; w); \alpha(x))) = r(f_0(x; w); \alpha(x)) \text{sgn}(2\alpha(x) - 1) ((\text{mix}(f_0(x; w); \alpha(x)) - y_0)x,$$

815 *where*

$$r(f(x; w); \alpha(x)) = \frac{f_0(x; w)(1 - f_0(x; w))}{\text{mix}(f_0(x; w); \alpha(x))(1 - \text{mix}(f_0(x; w); \alpha(x)))} |2\alpha(x) - 1|$$

816 *Let $L(x; w) = \mathbf{E}_{(x, y) \sim D}[\text{ce}(y, \text{mix}(f(x; w); \alpha(x))) \mid x]$ be the expected student label mixing loss con-*
817 *ditional on some example $x \in \mathbf{R}^d$. It holds $\partial_w L(x; w) = r(f(x; w); \alpha(x)) |2\alpha(x) - 1| (f_0(x; w) -$
818 $g_0(x)) x$.*

819 *Proof.* We first show the formula

$$\partial_w \text{ce}(y, \text{mix}(f(x; w); \alpha(x))) = r(f_0(x; w); \alpha(x)) \text{sgn}(2\alpha(x) - 1) ((\text{mix}(f_0(x; w); \alpha(x)) - y_0)x. \quad (5)$$

820 Using the chain rule, we obtain

$$\begin{aligned} \partial_w \text{ce}(y, \text{mix}(f(x; w); \alpha(x))) &= \\ &= - \frac{y_0}{\text{mix}(f_0(x; w); \alpha(x))} \partial_w (\text{mix}(f_0(x; w); \alpha(x))) \\ &= - \frac{y_1}{\text{mix}(f_1(x; w); \alpha(x))} \partial_w (\text{mix}(f_1(x; w); \alpha(x))). \end{aligned}$$

821 Now we observe that that for binary classification, it holds that $y_1 = 1 - y_0$, $\text{mix}(f_1(x; w); \alpha(x)) =$
822 $1 - \text{mix}(f_0(x; w); \alpha(x))$, and therefore, also $\partial_w \text{mix}(f(x; w); \alpha(x))_1 = -\partial_w \text{mix}(f(x; w); \alpha(x))_0$
823 to obtain the simplified expression:

$$\begin{aligned} \partial_w \text{ce}(y, \text{mix}(f(x; w); \alpha(x))) &= \\ &= - \frac{y_0}{\text{mix}(f_0(x; w); \alpha(x))} \partial_w (\text{mix}(f_0(x; w); \alpha(x))) \\ &= + \frac{1 - y_0}{1 - \text{mix}(f_0(x; w); \alpha(x))} \partial_w (\text{mix}(f_0(x; w); \alpha(x))). \end{aligned}$$

824 Further simplifying the above expression, we obtain:

$$\begin{aligned} \partial_w \text{ce}(y, \text{mix}(f(x; w); \alpha(x))) &= \\ &= \frac{\text{mix}(f_0(x; w); \alpha(x)) - y_0}{\text{mix}(f_0(x; w); \alpha(x)) (1 - \text{mix}(f_0(x; w); \alpha(x)))} \partial_w (\text{mix}(f_0(x; w); \alpha(x))). \end{aligned}$$

825 Using again the chain rule we obtain that

$$\partial_w (\text{mix}(f_0(x; w); \alpha(x))) = \alpha(x) \partial_w (f_0(x; w)) + (1 - \alpha(x)) \partial_w (1 - f_0(x; w)) = (2\alpha(x) - 1) \partial_w f_0(x; w).$$

826 Using the fact that the derivative of the sigmoid function $r(t) = 1/(1 + e^{-t})$, is $r'(t) = e^{-t}/(1 -$
827 $e^{-t})^2 = r(t)(1 - r(t))$, and the chain rule, we obtain that $\partial_w f_0(x; w) = f_0(x; w)(1 - f_0(x; w))x$.
828 Putting everything together we obtain the claimed formula for $\partial_w \text{ce}(y, \text{mix}(f(x; w); \alpha(x)))$.

829 To obtain the gradient formula for the expected loss conditional on some fixed example x , we can
830 use the fact that $\partial_w \mathbf{E}[\text{ce}(y, \text{mix}(f(x; w); \alpha(x))) \mid x] = \mathbf{E}[\partial_w \text{ce}(y, \text{mix}(f(x; w); \alpha(x))) \mid x]$. Now
831 using the formula of Equation (5) and the fact that $\mathbf{E}[y_0 \mid x] = \text{mix}(g_0(x); \alpha(x))$ by the definition of
832 our noise model, we obtain that

$$\begin{aligned} \partial_w L(x; w) &= r(f_0(x; w); \alpha(x)) \text{sgn}(2\alpha(x) - 1) (\text{mix}(f_0(x; w); \alpha(x)) - \text{mix}(g_0(x); \alpha(x))) \\ &= r(f_0(x; w); \alpha(x)) (2\alpha(x) - 1) (f_0(x; w) - g_0(x)) \end{aligned}$$

833

□

834 We first show the following claim proving that after roughly $T = 1/(\beta^2\gamma^2\epsilon^2)$ gradient iterations the
 835 student parameter vector $w^{(t)}$ will have good correlation with the ground-truth vector w^* .

836 *Claim 1.* Fix any T larger than a sufficiently large constant multiple of $\log(1/\delta)/(\epsilon^2\gamma^2\beta^2)$, and
 837 assume that for all $t \leq T$ it holds that $\mathbf{P}_{x \sim X}[\text{err}(f(x; w^{(t)}), g(x))] > \epsilon$. Then, we have $w^{(T)} \cdot w^* =$
 838 $\Omega(\beta\gamma\epsilon) T$, with probability at least $1 - \delta$.

839 *Proof.* Denote by $u^{(t)} = -\lambda^{(t)}g^{(t)}$ the update vector used in Algorithm 3. We observe that the
 840 weight vector at round T is equal to $w^{(T)} = \sum_{t=1}^T u^{(t)}$. In what follows we denote by $\mathcal{F}^{(t)}$ the
 841 filtration corresponding to the randomness of the updates of Algorithm 3. We define the martingale
 842 $q^{(T)} = \sum_{t=1}^T (\mathbf{E}[u^{(t)} | \mathcal{F}^{(t-1)}] - u^{(t)})$ with $q^{(0)} = 0$. We first show that under the assumption that
 843 $\mathbf{P}_{x \sim X}[\text{argmax}(f(x; w^{(t)})) \neq g(x)] > \epsilon$, for all $t \leq T$, it holds that $\sum_{t=1}^T \mathbf{E}[u^{(t)} | \mathcal{F}^{(t-1)}] \cdot w^* \geq$
 844 $(\epsilon\gamma\beta/2) T$. Using the SLam gradient expression of Lemma E.3 and the definition of the step size $\lambda^{(t)}$
 845 we obtain that $\mathbf{E}[u^{(t)} | \mathcal{F}^{(t-1)}] = \mathbf{E}_{x \sim X}[|2\alpha(x) - 1| (g_0(x) - f_0(x; w^{(t-1)})) x]$. Take any step t .
 846 We have that

$$\begin{aligned} \mathbf{E}[u^{(t)} | \mathcal{F}^{(t-1)}] \cdot w^* &= \mathbf{E}_{x \sim X}[|2\alpha(x) - 1| (g_0(x) - f_0(x; w^{(t-1)})) (x \cdot w^*)] \\ &= \mathbf{E}_{x \sim X}[|2\alpha(x) - 1| |g_0(x) - f_0(x; w^{(t-1)})| |x \cdot w^*|], \end{aligned}$$

847 where we used the fact that $(g_0(x) - f_0(x; w^{(t-1)})) \text{sgn}(x \cdot w^*) = |g_0(x) - f_0(x; w^{(t-1)})|$. Now,
 848 using the γ -margin assumption of the distribution D and the fact that $|2\alpha(x) - 1| \geq \beta$ we obtain

$$\begin{aligned} \mathbf{E}[u^{(t)} | \mathcal{F}^{(t-1)}] \cdot w^* &\geq \beta\gamma \mathbf{E}_{x \sim X}[|g_0(x) - f_0(x; w^{(t-1)})|] \\ &\geq \beta\gamma \mathbf{E}_{x \sim X}[|g_0(x) - f_0(x; w^{(t-1)})| \text{err}(g(x), f(x; w^{(t-1)}))] \\ &\geq (\beta\gamma/2) \mathbf{P}_{x \sim X}[\text{err}(g(x), f(x; w^{(t-1)}))] \geq \beta\gamma\epsilon/2, \end{aligned}$$

849 where for the penultimate inequality we used the fact that when $g(x)$ and $f(x; w^{(t-1)})$ disagree it
 850 holds that $|g_0(x) - f_0(x; w^{(t-1)})| \geq 1/2$. Take, for example, the case where $g_0(x) = 1$. Then
 851 $f_0(x; w^{(t-1)})$ must be smaller than $1/2$ otherwise the prediction of the model $\text{argmax} f(x; w^{(t-1)})$
 852 would also be 0 (and would agree with the prediction of $g(x)$). Finally, for the last inequality we used
 853 the fact that, by our assumption, it holds that $\mathbf{P}_{x \sim X}[\text{err}(g(x), f(x; w^{(t-1)}))] \geq \epsilon$. Therefore, we
 854 conclude that $\sum_{t=1}^T \mathbf{E}[u^{(t)} | \mathcal{F}^{(t-1)}] \cdot w^* \geq (\epsilon\gamma\beta/2) T$. Next, we shall show that $w^{(T)}$ also achieves
 855 good correlation with the optimal direction w^* with high probability. We will use the fact that $q^{(t)}$ is
 856 a martingale and the Azuma-Hoeffding inequality to show that $w^{(T)} \cdot w^*$ will not be very far from its
 857 expectation.

858 **Lemma E.4** (Azuma-Hoeffding). *Let $\xi^{(t)}$ be a martingale with bounded increments, i.e., $|\xi^{(t)} -$
 859 $\xi^{(t-1)}| \leq M$. It holds that $\mathbf{P}[\xi^{(T)} \geq \xi^{(0)} + \lambda] \leq e^{-\lambda^2/(2M^2T)}$.*

860 Recall that from Lemma E.3 we have that $\mathbf{E}[u^{(t)} | \mathcal{F}^{(t-1)}] = \mathbf{E}_{x \sim X}[|2\alpha(x) - 1| (g_0(x) -$
 861 $f_0(x; w^{(t-1)})) x]$ and

$$u^{(t)} = \text{sgn}(2\alpha(x^{(t)}) - 1) (y_0^{(t)} - \text{mix}(f_0(x^{(t)}; w^{(t-1)}), \alpha(x^{(t)})) x^{(t)}).$$

862 Observe that since $\|x\|_2 \leq 1$ for all x it holds that $\|u^{(t)}\|_2 \leq 1$. Therefore, the difference $\|\mathbf{E}[u^{(t)} |$
 863 $\mathcal{F}^{(t-1)}] - u^{(t)}\|_2 \leq 2$ with probability 1. Since $\|w^*\|_2 = 1$, using Cauchy-Schwarz, we also obtain
 864 that $|\mathbf{E}[u^{(t)} \cdot w^* | \mathcal{F}^{(t-1)}] - u^{(t)} \cdot w^*| \leq 2$.

865 Using Lemma E.4 and the fact that $q^{(0)} = 0$ we obtain that $\mathbf{P}[q^{(T)} \cdot w^* \geq (\beta\gamma\epsilon/4) T] \leq$
 866 $e^{-\beta^2\gamma^2\epsilon^2 T/128}$. Therefore we conclude that for any T larger than $128 \log(1/\delta)/(\beta^2\gamma^2\epsilon^2)$, with
 867 probability at least $1 - \delta$, it holds that $q^{(T)} \cdot w^* \geq (\beta\gamma\epsilon/4) T$ or equivalently $w^{(T)} \cdot w^* \geq (\beta\gamma\epsilon/4) T$,
 868 where we used our previously obtained bound for the expected updates $\sum_{t=1}^T \mathbf{E}[u^{(t)} | \mathcal{F}^{(t-1)}] \cdot w^* \geq$
 869 $(\beta\gamma\epsilon/2) T$.

870 □

871 *Claim 2.* Fix any $T \geq 1$. Then, we have $\|w^{(T)}\|_2 = O(\sqrt{T})$, with probability at least 99%.

872 *Proof.* We have that $\|w^{(T)}\|_2^2 = \|w^{(T-1)}\|_2^2 + 2u^{(T)} \cdot w^{(T-1)} + \|u^{(T)}\|_2^2$. Unrolling the iteration,
 873 we obtain that

$$\|w^{(T)}\|_2^2 = 2 \sum_{t=1}^T u^{(t)} \cdot w^{(t-1)} + \sum_{t=1}^T \|u^{(t)}\|_2^2 \leq 2 \sum_{t=1}^T u^{(t)} \cdot w^{(t-1)} + T, \quad (6)$$

874 where we used the fact that, since $\|x^{(t)}\|_2 \leq 1$, it holds that $\|u^{(t)}\|_2 \leq 1$ (see the proof of Claim [1](#)).
 875 We first show that $\sum_{t=1}^T \mathbf{E}[u^{(t)} \mid \mathcal{F}^{(t-1)}] \cdot w^{(t-1)} = O(T)$. We have

$$\begin{aligned} \mathbf{E}[u^{(t)} \mid \mathcal{F}^{(t-1)}] \cdot w^{(t-1)} &= \mathbf{E}_{x \sim X} [2\alpha(x) - 1 \mid (g_0(x) - f_0(x; w^{(t-1)})) (x \cdot w^{(t-1)})] \\ &\leq \mathbf{E}_{x \sim X} [(g_0(x) - f_0(x; w^{(t-1)})) (x \cdot w^{(t-1)})]. \end{aligned}$$

876 We will show that for x it holds that

$$g_0(x) - f(x; w^{(t-1)})(x \cdot w^{(t-1)}) \leq \frac{1}{e}.$$

877 Fix some x and let $s = w^{(t-1)} \cdot x$. Assume first that $g_0(x) = 1$. Then, we have

$$g_0(x) - f(x; w^{(t-1)})(x \cdot w^{(t-1)}) = \left(1 - \frac{1}{1 + e^{-s}}\right) s = s \frac{e^{-s}}{1 + e^{-s}} \leq \frac{1}{e},$$

878 where we used the fact that $s \frac{e^{-s}}{1 + e^{-s}} \leq 0$ for $s \leq 0$ and $s \frac{e^{-s}}{1 + e^{-s}} \leq se^{-s} \leq 1/e$ for $s \geq 0$ (using the
 879 elementary inequality $ze^{-z} \leq 1/e$ for all $z \in \mathbf{R}$). When $g_0(x) = 0$ we similarly have that

$$g_0(x) - f(x; w^{(t-1)})(x \cdot w^{(t-1)}) = -\frac{s}{1 + e^{-s}} \leq \frac{1}{e},$$

880 where we used the fact that when $s \geq 0$ it holds that $-\frac{s}{1 + e^{-s}} \leq 0$ and when $s \leq 0$, $-\frac{s}{1 + e^{-s}} \leq$
 881 $-s/e^{-s} = -se^s$. For the final inequality, we used again the inequality $ze^{-z} \leq 1/e$ for all $z \in \mathbf{R}$
 882 (where we replaced z with $-s$).

883 Therefore, we obtain that $\mathbf{E}[u^{(t)} \mid \mathcal{F}^{(t-1)}] \cdot w^{(t-1)} \leq 1/e$ and $\sum_{t=1}^T \mathbf{E}[u^{(t)} \mid \mathcal{F}^{(t-1)}] \cdot w^{(t-1)} \leq T/e$.
 884 Using the decomposition of Equation [\(6\)](#), linearity of expectation, and the tower rule for conditional
 885 expectations, we conclude that $\mathbf{E}[\|w^{(T)}\|_2^2] \leq (2/e + 1)T$. Using Markov's inequality we obtain that
 886 with probability at least 99% it holds that $\|w^{(T)}\|_2^2 = O(T)$ or equivalently $\|w^{(T)}\|_2 = O(\sqrt{T})$.

887 \square

888 We can now finish the proof of Theorem [5.1](#). Assume, in order to reach a contradiction, that
 889 for all $t \leq T$ it holds that $\mathbf{P}_{x \sim X}[\text{err}(f(x; w^{(t)}), g(x))] > \epsilon$. Now picking T to be larger than a
 890 sufficiently large constant multiple of $1/(\epsilon^2 \gamma^2 \beta^2)$ and using Claim [1](#) and Claim [2](#) we obtain that,
 891 with probability at least 99%, it holds that $w^{(T)} \cdot w^* / \|w^{(T)}\|_2 \geq \Omega(\beta \gamma \epsilon \sqrt{T})$, which can be made to
 892 be larger than 1 by our choice of T . However, this is a contradiction as by Cauchy-Schwarz we have
 893 $w^{(T)} \cdot w^* / \|w^{(T)}\|_2 \leq \|w^*\|_2 \leq 1$. Therefore, with probability at least 99%, it must be that for some
 894 $t \leq T$ it holds that $\mathbf{P}_{x \sim X}[\text{err}(f(x; w^{(t)}), g(x))] \leq \epsilon$.

895 \square




Figure 8: The landscape and gradient field of the population student label mixing loss for a simple 2 dimensional feature problem with a ground truth corresponding to a halfspace. We observe that the landscape is non-convex; however we can see that the corresponding gradient field “points towards the optimal direction” and therefore gradient descent converges to the global minimizer. A potential issue is the fact that the landscape contains regions where the gradients may almost vanish and this could lead to the gradient iteration of the student getting trapped there. To handle this issue, in Algorithm 3 we multiply the gradient of SLAM with an appropriate step-size.