
From Theory to Decision Rule: Calibrating the Noisy-Label Crossover for VLM Weak Supervision Across Three Medical-Imaging Benchmarks

Anonymous Authors¹

Abstract

Classical noisy-label theory predicts that downstream performance under weak supervision is bounded above by the labeler’s accuracy, implying a sharp *crossover*: once a gold-trained classifier matches the labeler, weak labels stop helping and start hurting. The prediction is theoretical; what is missing is a benchmark calibration that turns it into an instance-level statement for modern foundation-model labelers. We provide such a calibration for BiomedCLIP-generated weak labels on three medical-imaging benchmarks (PCAM, ISIC, NIH-CXR) and six downstream architectures spanning an $11\times$ parameter range. The crossover predicted by theory appears at $n_g \approx 100$ on PCAM, 20–50 on ISIC, and 250–500 on NIH-CXR; weak labels above the crossover degrade AUC by up to -0.10 . The location is architecture-invariant for four of five pretrained architectures, and a within-family DenseNet sweep ($2.5\times$ parameters, identical pre-training) confirms the labeler—not the student—is the binding constraint. The calibration in turn produces a decision rule operable from 10–20 gold labels: compare gold-only AUC to VLM accuracy on the user’s gold set. A structured-vs-random noise sign flip on NIH-CXR shows that the rate-only formulation of the bound is incomplete and identifies a concrete refinement (label-space projection) that future benchmarks can be designed to test.

1. Introduction

A common pattern for adapting foundation models (FMs) to specialized domains is to use the FM as a *labeler*: prompt

¹Anonymous Institution. Correspondence to: Anonymous <anonymous@example.com>.

Proceedings of the 43rd International Conference on Machine Learning, Seoul, South Korea. PMLR 306, 2026. Copyright 2026 by the author(s).

a large pretrained model for predictions on unlabeled data, then train a small task-specific network on the resulting pseudo-labels (Ratner et al., 2016; Xie et al., 2020; Sohn et al., 2020). In medical imaging, this turns a 400M-parameter vision–language model (VLM) such as BiomedCLIP (Zhang et al., 2024) into a free source of supervision for a deployable 8M-parameter classifier—one of the cheapest adaptation primitives available, and one that is increasingly drop-in because the FM and the downstream model share no parameters or gradients. The implicit assumption is that more weak labels are strictly helpful, with the only question being how to filter or weight them (Zhong et al., 2024; Patrini et al., 2017).

Classical noisy-label theory (Natarajan et al., 2013; Fréney & Verleysen, 2014; Song et al., 2022) predicts the opposite: downstream performance is bounded above by the labeler’s accuracy, so once a gold-only classifier matches the labeler, additional weak labels become contaminant. The bound implies a regime change—a *crossover*—in the gold-budget axis. What has been missing is a benchmark calibration of where this crossover sits for modern foundation-model labelers, how it depends on the downstream architecture (the axis a benchmark user actually varies), and how the bound translates into a decision rule operable from a small held-out set.

Relevance to CTB. The workshop calls for theory–benchmark loops in which theoretical predictions are sharpened by structured benchmarks and benchmarks are designed to test theory. The paper follows that loop end-to-end: a classical noisy-label upper bound predicts a phase transition; three benchmarks calibrate where it falls; an $11\times$ architecture sweep tests the bound’s invariance claim (the labeler, not the student, sets the ceiling); a structured-vs-random noise contrast probes which part of the bound is binding (rate vs. structure); and the calibration is then re-emitted as a decision rule with stated regimes of validity.

Contributions.

1. **Benchmark calibration of a theoretical bound.** We instantiate the noisy-label upper bound on three medical-

imaging benchmarks (PCAM, ISIC, NIH-CXR) under BiomedCLIP weak supervision and locate the predicted crossover at $n_g \approx 100, 20\text{--}50,$ and $250\text{--}500$ respectively.

- Architecture-invariance test of the bound.** Six PCAM architectures spanning $11\times$ parameters and a within-family DenseNet ladder (DN-121/169/201) share the same crossover and a 0.024-AUC ceiling band, consistent with the theory’s prediction that the labeler, not the student, governs the ceiling.
- Probing which part of the theory binds.** Matched-rate structured VLM noise underperforms random noise by 0.09–0.11 AUC on PCAM but *outperforms* it by 0.03–0.04 AUC on NIH-CXR. The sign flip indicates that the rate-only formulation of the bound is incomplete on multi-label tasks where the FM projects onto a coarser label space, and identifies label-space projection as a concrete axis a refined statement should condition on.
- A theory-anchored decision rule.** The calibrated bound yields a one-line rule—compare gold-only AUC to VLM accuracy on the user’s gold set, both estimable from 10–20 labels—that is architecture-invariant on the pretrained students we tested and replicates across the three benchmarks at task-specific crossover counts.

2. Setup

Datasets. **PatchCamelyon (PCAM)** (327K 96×96 histopathology patches, 50/50 metastasis label); **ISIC 2019** (25K dermoscopy images, 8-class lesion classification); **NIH-CXR** (112K chest radiographs, 15-class multi-label). The three benchmarks span a wide range of FM–task alignments.

Foundation-model labeler. BiomedCLIP (Zhang et al., 2024) produces zero-shot predictions via contrastive matching to text prompts. Unfiltered accuracy / AUC: PCAM 72.8%/0.84, ISIC 64.7%/0.62 (binary malignancy score tiled across 8 classes), NIH-CXR 71.7%/0.62 (binary normal/abnormal tiled across 15 classes). Confidence filtering (≥ 0.70 positive, ≤ 0.03 negative) on PCAM retains $\sim 51\%$ of predictions at 83.6% accuracy. ISIC and NIH-CXR are intentionally cases where a general-purpose medical VLM lacks task-specific vocabulary.

Adaptation pipeline. We train downstream classifiers on mixtures of n_g gold and n_w weak labels with BCEWithLogitsLoss, Adam ($\text{lr}=10^{-4}$, weight decay 10^{-5}), batch size 64, ≤ 15 epochs with early stopping (patience 3 on val AUC). The PCAM grid sweeps $n_g \in \{10, 20, 50, 100, 500, 1\text{K}, 5\text{K}, \text{full}\} \times n_w \in \{0, 1\text{K}, 5\text{K}, 50\text{K}, \text{full}\}$ with 3 seeds.

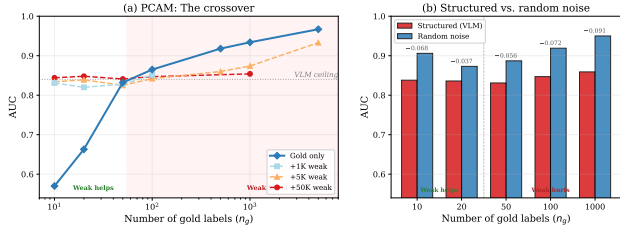


Figure 1. The VLM adaptation crossover on PCAM. Left: AUC vs. gold count at four weak-label doses; the gold-only curve crosses the BiomedCLIP ceiling (~ 0.84 , dashed) near $n_g=100$. Right: at matched noise rates, structured VLM errors consistently underperform random corruption.

Table 1. PCAM AUC by gold/weak configuration (DenseNet-121, 3-seed mean). The crossover sits at $n_g \approx 100$: below this, weak labels add up to $+0.27$; above, they cost up to ~ 0.10 .

$n_g \downarrow / n_w \rightarrow$	0	1K	5K	50K	Δ best
10	.579	.830	.835	.846	+0.267
20	.662	.829	.842	.851	+0.189
50	.832	.834	.830	.846	+0.014
100	.868	.855	.842	.851	−0.017
500	.922	.885	.866	.844	−0.037
1,000	.936	.908	.876	.857	−0.028
5,000	.970	.954	.937	.874	−0.106

Architecture sweep. On PCAM we replicate the core grid with six downstream architectures: DenseNet-121/169/201 (8/14/20M, ImageNet), DenseNet-264 (33M, random init—we flag and exclude this from monotonicity claims), ResNet-50 (25M, ImageNet), and ConvNeXt-Base (89M, ImageNet). The within-family DN-121/169/201 ladder holds stem, growth rate, normalization, and pretraining identical, varying only depth in blocks 3 and 4—the cleanest test of whether downstream capacity matters.

3. The Crossover

The central finding on PCAM with the default DenseNet-121 student is a sharp regime change at $n_g=100$ (Figure 1, Table 1). At $n_g=10$, the gold-only classifier is near chance (0.579); adding 50K weak labels lifts AUC to 0.846—a $+0.27$ adaptation gain. At $n_g=100$ the gold-only baseline (0.868) already exceeds BiomedCLIP’s 0.84 ceiling; from there, every weak-label configuration *degrades* performance, and at $n_g=5\text{K}$ adding 50K weak labels costs -0.10 AUC. The locus of the crossover coincides exactly with the gold-only AUC crossing the VLM accuracy.

Cross-dataset replication. The crossover replicates on the other two benchmarks. On **ISIC**, BiomedCLIP outputs only a binary malignancy score on an 8-class task: the crossover comes earlier ($n_g \in [20, 50]$) and the post-crossover penalty is steeper, -0.14 AUC at $n_g=500$. On

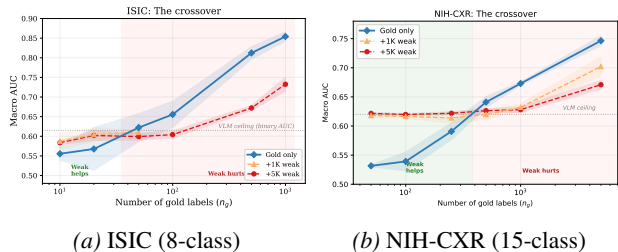


Figure 2. **The crossover replicates across modalities.** ISIC crosses at $n_g \in [20, 50]$; NIH-CXR at $n_g \in [250, 500]$. Crossover location tracks BiomedCLIP’s task-specific accuracy.

Table 2. PCAM AUC, BiomedCLIP (structured) vs. matched-rate random noise ($n_w=50K$, DN-121, 3-seed mean). Structured errors consistently cost ~ 0.10 AUC.

Gold	Structured	Random	$\Delta (S-R)$	Regime
10	.844	.946	-.102	Weak helps
20	.848	.938	-.090	Weak helps
50	.841	.940	-.100	Crossover
100	.847	.958	-.111	Weak hurts
1,000	.854	.958	-.104	Weak hurts

NIH-CXR, where the VLM is weaker (macro AUC ceiling ~ 0.62) and the task is 15-class multi-label, the crossover lies between 250 and 500 gold; the penalty grows with gold, reaching -0.076 AUC at $n_g=5K$ (Figure 2).

Structured noise depends on FM–task alignment. At a matched $\sim 27\%$ noise rate, structured BiomedCLIP errors underperform random label corruption on PCAM by $0.09\text{--}0.11$ AUC across the entire grid (Table 2). On NIH-CXR the sign reverses: structured labels *outperform* random by $+0.03\text{--}+0.04$ AUC, because BiomedCLIP’s binary normal/abnormal prediction carries partial signal across all 15 pathology classes that random corruption destroys. The lesson for adaptive FM pipelines is that the cost of structured FM errors depends on how directly the FM addresses the target task, not on the noise rate alone.

4. Architecture-Invariance: Testing the Theory’s Locus Claim

The noisy-label upper bound is a statement about the *labeler*: the downstream classifier inherits a ceiling that is a function of labeler accuracy, not of student capacity. This is a falsifiable prediction. We test it by sweeping the downstream architecture across $11\times$ parameters and asking whether the ceiling moves. It does not. The best weak-augmented AUC at $n_g \leq 50$ on PCAM falls in a tight band of 0.024 AUC across six architectures (Table 3). The within-family DN-121/169/201 ladder—which holds stem, growth rate, pre-training, and optimizer identical—shows no monotone trend in capacity at any gold count (Table 4). All four pretrained

Table 3. Weak-label ceiling on PCAM by architecture (best AUC at $n_g \leq 50$, any n_w , 3-seed mean). $11\times$ parameter spread maps to a 0.024 AUC band.

Architecture	Params	Pretrained	Weak ceiling
DenseNet-121	8M	yes	.857
DenseNet-169	14M	yes	.841
DenseNet-201	20M	yes	.845
ResNet-50	25M	yes	.843
DenseNet-264	33M	no	.849
ConvNeXt-Base	89M	yes	.865

Table 4. Within-family DenseNet sweep on PCAM, gold-only AUC (3-seed mean \pm std). The three curves lie within one std of each other from $n_g=50$ on, with no monotone trend in parameter count.

n_g	DN-121 (8M)	DN-169 (14M)	DN-201 (20M)
10	.570 \pm .094	.617 \pm .091	.549 \pm .066
50	.832 \pm .030	.825 \pm .032	.841 \pm .022
100	.865 \pm .003	.865 \pm .005	.876 \pm .002
500	.918 \pm .001	.900 \pm .003	.909 \pm .001
5,000	.967 \pm .001	.958 \pm .001	.961 \pm .001

architectures except ResNet-50 cross at $n_g=100$; ResNet-50 crosses later only because it has a lower gold-only scaling exponent ($\beta \approx 0.27$ vs. $0.5\text{--}0.75$), an inductive-bias effect rather than a capacity effect.

PCAM at 96×96 resolution appears capacity-saturated near 8M parameters within this family. The implication for benchmark design is that aggregate scores at large student capacity will overstate the ceiling’s tightness on a single labeler–task pair: the ceiling is a property of the benchmark–labeler joint, not of the model under test. A capability claim of the form “model M achieves AUC x under VLM weak supervision” should be reported as a labeler-conditioned upper bound, not a model property.

5. Confidence Filtering as a Regime-Dependent Knob

Confidence filtering at the FM (≥ 0.70 pos, ≤ 0.03 neg) keeps $\sim 51\%$ of predictions and raises labeler accuracy from 72.8% to 83.6% . On DN-121 the effect is regime-dependent: below the crossover (≤ 20 gold), the volume loss outweighs the accuracy gain ($\Delta \text{AUC} \in [-0.007, -0.005]$); above it, CF adds $+0.01\text{--}+0.03$ AUC. Across the architecture sweep, CF benefit tracks sample efficiency rather than parameter count: DN-121 and ConvNeXt-Base both benefit; ResNet-50 is neutral. The sign of the intervention’s effect therefore flips across the regime boundary, so any aggregate score for CF that does not condition on the regime will average two opposite effects.

6. From Calibration to Decision Rule

Rule. Estimate (i) FM accuracy on your gold set and (ii) gold-only AUC of any small ($\sim 8M$) pretrained student. *If gold-only AUC < FM accuracy, train with weak labels. Otherwise, additional weak labels degrade performance and the next budget unit is better spent on gold annotation.*

The rule is the empirical instantiation of the noisy-label upper bound at decision time: both inputs estimate the two quantities the bound compares, and both are estimable from 10–20 gold labels. The benchmark calibration converts a population-level statement into an instance-level prediction with a stated regime of validity (architecture-invariant on the four pretrained PCAM students we tested; replicates on ISIC and NIH-CXR with task-specific crossover counts; PubMedCLIP replication on ISIC consistent with the predicted ceiling shift).

What the rule does and does not predict. The rule predicts whether VLM weak supervision will help in expectation, and on the medical-imaging axis it has been verified across three benchmarks, six architectures, and three seeds. It does not predict the magnitude of the post-crossover penalty (which scales with n_w and with the structured-vs-random gap of Section 3); it does not predict per-class behavior on multi-label tasks where the FM projects onto a coarser label space; and it does not adjust for distribution shift between the gold set and the unlabeled pool. Each is a candidate axis for a sharper theoretical statement that future benchmarks could be designed to test.

7. Discussion

Theory to benchmark to refined theory. The noisy-label bound (Natarajan et al., 2013; Frénay & Verleysen, 2014) predicts *that* a crossover exists; the benchmark calibration locates *where* (three task-specific points) and supports what the bound says about *whom* (the labeler, not the student). The structured-vs-random sign flip on NIH-CXR (Section 3) cuts the other way: it shows the rate-only formulation of the bound is incomplete, because partial-signal projections of the FM’s label space can outperform random noise of the same nominal rate. A sharper statement would condition on a label-space-projection term—a candidate refinement that this benchmark surfaces and that future benchmarks (per-class crossovers, multi-label projections) could be designed to test directly.

Implications for benchmark design. The architecture-invariance result implies that single-aggregate benchmark scores under VLM weak supervision are confounded with the labeler’s accuracy. A more informative benchmark

would (i) report gold-only and weak-augmented curves jointly, (ii) tag each operating point with the regime it sits in (below / at / above the crossover), and (iii) decompose the noise contribution into rate and structure terms. None of this requires new datasets; it requires reporting protocol changes that align scores with the theoretical quantities that govern them.

Failure modes and refinements the rule does not capture.

(i) *Distribution shift* between the FM’s pretraining corpus and the deployment domain can decouple FM accuracy on a small gold set from FM accuracy on the unlabeled pool. (ii) *Multi-VLM ensembles* (Ratner et al., 2016) restore the disagreement signal that single-FM weak labeling lacks and may shift the crossover. (iii) Tasks with extreme class imbalance or where the FM’s binary score is the wrong projection require a per-class crossover analysis instead of a scalar comparison. Each is a target for a structured benchmark designed against the refined theoretical statement.

Limitations. Specific crossover counts are dataset- and FM-specific; we have sampled three medical-imaging benchmarks and one primary VLM (plus a PubMedCLIP replication on ISIC consistent with the predicted ceiling). All experiments use naive label merging; noise-aware losses or sample reweighting could soften the post-crossover penalty. The within-family ladder is three points (DN-121/169/201) plus a confounded random-init top point; a custom-pretrained DN-264 or a width-varying ladder would strengthen the capacity-null claim.

Conclusion. A classical noisy-label upper bound predicts a regime change in VLM weak supervision; three medical-imaging benchmarks under an $11\times$ architecture sweep calibrate where it sits and support the bound’s invariance claim; the calibration re-emits as a decision rule operable from 10–20 labels; and the structured-vs-random sign flip identifies the next refinement the theory needs.

References

- Frénay, B. and Verleysen, M. Classification in the presence of label noise: a survey. *IEEE Transactions on Neural Networks and Learning Systems*, 25(5):845–869, 2014.
- Natarajan, N., Dhillon, I. S., Ravikumar, P., and Tewari, A. Learning with noisy labels. In *Advances in Neural Information Processing Systems*, volume 26, 2013.
- Patrini, G., Rozza, A., Menon, A. K., Nock, R., and Qu, L. Making deep neural networks robust to label noise: A loss correction approach. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2233–2241, 2017.

- Ratner, A. J., De Sa, C. M., Wu, S., Selsam, D., and Ré, C. Data programming: Creating large training sets, quickly. In *Advances in Neural Information Processing Systems*, volume 29, 2016.
- 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. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Song, H., Kim, M., Park, D., Shin, Y., and Lee, J.-G. Learning from noisy labels with deep neural networks: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, 34(11):8135–8153, 2022. doi: 10.1109/TNNLS.2022.3152527.
- Xie, Q., Luong, M.-T., Hovy, E., and Le, Q. V. Self-training with Noisy Student improves ImageNet classification. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10687–10698, 2020.
- Zhang, S., Xu, Y., Usuyama, N., Xu, H., Bagga, J., Tinn, R., Preston, S., Rao, R., Wei, M., Valluri, N., Wong, C., Tupini, A., Wang, Y., Mazzola, M., Shukla, S., Liden, L., Gao, J., Lungren, M. P., Naumann, T., Wang, S., and Poon, H. BiomedCLIP: A multimodal biomedical foundation model pretrained from fifteen million scientific image-text pairs. 2024.
- Zhong, L., Huang, Z., Liu, Y., Liao, W., Zhang, S., Wang, G., and Zhang, S. VLM-CPL: Consensus pseudo labels from vision-language models for annotation-free pathological image classification. *IEEE Transactions on Medical Imaging*, 2024.