
External Paired-Scanner Validation of Paired-Acquisition Neural Factorization on Canine Squamous Cell Carcinoma Histopathology

Matthew Vaishnav

Independent Researcher

Kitchener-Waterloo, Canada

`paired-acquisition-factorization-caninescc`

Abstract

Scanner-specific acquisition effects can confound representation learning in computational pathology, especially when models are evaluated across sites, devices, or staining/scanning pipelines. Paired-Acquisition Neural Factorization is a paired-acquisition projection objective designed to preserve biological region identity while reducing scanner-identifiable signal in the biological embedding. I evaluate whether a Paired-Acquisition Neural Factorization configuration selected on a human paired-scanner benchmark transfers to an independent public multi-scanner canine cutaneous squamous cell carcinoma dataset. After auditing annotation correspondence and applying deterministic geometry qualification, I retained 805 complete five-scanner regions from 44 biological samples. Using frozen DINOv2-Base features and locked hyperparameters, Paired-Acquisition Neural Factorization reduced sample-blocked scanner-probe accuracy by 0.380 absolute while improving cross-scanner pair cosine and preserving same-region retrieval. These results support Paired-Acquisition Neural Factorization as a paired-acquisition representation method that transfers beyond the development benchmark, while remaining a research-only validation rather than a clinical or diagnostic claim.

1 Introduction

Computational pathology representations can encode both tissue morphology and acquisition artifacts. When acquisition signal is correlated with scanner, site, or cohort membership, downstream performance may reflect technical shortcuts rather than robust biological representation. Paired-acquisition datasets provide a direct way to test whether a representation preserves tissue identity while reducing scanner-identifiable variation.

Paired-Acquisition Neural Factorization separates a frozen image representation into a biological embedding and an acquisition embedding by combining paired-region contrastive alignment, reconstruction, variance and covariance regularization, scanner-adversarial learning, and scanner-dependence

penalties. The present study asks whether the locked Paired-Acquisition Neural Factorization configuration transfers to an independent external paired-scanner benchmark.

2 Dataset and Geometry Qualification

I use the public Multi-Scanner Canine Cutaneous Squamous Cell Carcinoma histopathology dataset, which contains 44 biological samples scanned on five scanners: CS2, GT450, NZ20, NZ210, and P1000. The raw release contains 220 TIFF images and 6,215 annotation views corresponding to 1,243 candidate matched regions.

Before modeling, I audited the COCO annotations, SlideRunner SQLite database, TIFF dimensions, and scanner-to-scanner geometry. The P1000 view showed a coherent sample-specific orientation discrepancy rather than random annotation failure. I therefore applied deterministic P1000 orientation normalization and retained only regions for which every scanner view required at most 10% adaptive-crop padding. This produced a geometry-qualified paired-acquisition subset of 805 complete five-view regions, corresponding to 4,025 image views.

3 Frozen Encoder Triage

I evaluated frozen ResNet50 ImageNet, DINOv2-Base, and Phikon representations on the fold-0 validation split before running Paired-Acquisition Neural Factorization. DINOv2-Base was selected as the external substrate because it had the strongest cross-scanner cosine consistency and best worst-case retrieval while still retaining scanner-identifiable signal above chance.

Table 1: Frozen encoder triage on fold-0 validation.

Encoder	Probe	Avg cos	Worst cos
ResNet50	0.857	0.822	0.750
DINOv2-B	0.841	0.908	0.862
Phikon	0.980	0.810	0.779

Encoder	Avg R@1	Worst R@1	Rank
ResNet50	0.854	0.653	74.7
DINOv2-B	0.852	0.728	51.6
Phikon	0.776	0.640	59.1

4 Locked Paired-Acquisition Neural Factorization Protocol

The Paired-Acquisition Neural Factorization hyperparameters were locked before the external five-fold test. The validation run used seeds 901–905 on fold 0, fitting on train and evaluating on validation. The final external test used five sample-blocked folds and new seeds 911–915. For each fold, all non-test samples were used for fitting and the held-out fold was projected and evaluated exactly once.

The two variants were a paired-consistency reference projection and the locked Paired-Acquisition Neural Factorization dependence-penalty configuration, denoted Paired-Acquisition Neural Factorization dep20. The external test reused the same objective structure selected before this benchmark was inspected as a final result source.

5 Results

Across five folds and five seeds, Paired-Acquisition Neural Factorization reduced scanner identifiability while preserving same-region retrieval. Descriptive run means are shown in Table 2. Sample-blocked paired contrasts over 44 biological samples are shown in Table 3.

Table 2: Locked five-fold descriptive run means.

Variant	Probe	Avg cos	Worst cos	Avg R@1
Reference	0.752868	0.696022	0.627300	0.930637
Paired-Acquisition Neural Factorization	0.361408	0.729961	0.656736	0.933392

Table 3: Sample-blocked Paired-Acquisition Neural Factorization minus paired-reference contrasts.

Metric	Diff.	CI low	CI high	<i>p</i>
Probe	-0.380347	-0.399616	-0.361107	0.000004
Avg cos	+0.033256	+0.030935	+0.035530	0.000004
Worst cos	+0.033104	+0.028789	+0.037550	0.000004
Avg R@1	+0.002326	+0.000005	+0.004789	0.064600
Worst R@1	+0.001731	-0.005823	+0.009250	0.660573

The main scanner-probe contrast was -0.380 absolute with a bootstrap 95% confidence interval of $[-0.400, -0.361]$ and Monte Carlo sign-flip $p = 0.000004$. All 44 sample blocks favored Paired-Acquisition Neural Factorization for scanner-probe reduction. Pair-cosine average and worst-pair cosine also improved for all 44 samples. Retrieval top-1 average and worst-pair retrieval remained non-inferior within the predefined 0.02 margin.

The acquisition factor retained scanner information, with mean acquisition scanner-probe accuracy of 0.865, while acquisition-space tissue retrieval was low at 0.181. The mean biological-acquisition cross-covariance RMS was 0.0898. This pattern is consistent with scanner signal being preferentially represented in the acquisition branch rather than the biological embedding.

6 Discussion

This external study supports the transferability of Paired-Acquisition Neural Factorization’s paired-acquisition objective. The method reduced scanner-identifiable signal by a large absolute margin on a dataset not used for hyperparameter selection, while improving cross-scanner cosine consistency and preserving same-region retrieval. The result is important because it demonstrates that the effect is not confined to the original human paired-scanner benchmark.

The study has several limitations. The dataset is canine SCC rather than human pathology, and the experiment evaluates representation identifiability rather than clinical performance. The geometry-qualified subset intentionally excludes regions that required substantial padding or unreliable crop support. Finally, Paired-Acquisition Neural Factorization is evaluated here as a projection on frozen features, not as an end-to-end diagnostic model.

7 Reproducibility and Claim Boundary

The accompanying repository contains the scripts, split metadata, compact result tables, and paper source needed to inspect the external validation package. Raw whole-slide TIFF files, extracted JPEG patches, frozen feature archives, and checkpoints are intentionally excluded from Git.

The supported claim is narrow: on an independent paired-scanner canine SCC benchmark, a Paired-Acquisition Neural Factorization projection with hyperparameters locked from the human SCORPION development benchmark reduced scanner identifiability while preserving same-region retrieval and improving cross-scanner cosine consistency. This is not clinical validation, diagnostic software, or evidence of patient-care utility.

8 Conclusion

A Paired-Acquisition Neural Factorization configuration locked from prior development transferred to an independent external paired-scanner canine SCC benchmark. On 44 sample-blocked biological units, Paired-Acquisition Neural Factorization reduced scanner-probe accuracy by approximately 0.38 absolute, improved cross-scanner cosine consistency, and preserved same-region retrieval.

This provides external paired-acquisition evidence for Paired-Acquisition Neural Factorization as a research method for scanner-aware representation identifiability in computational pathology.