TriSCoV-Net: Cross-Scale Verified Virtual Immunomarker Proxy Generation from Diffusion Magnetic Resonance Imaging

Hongjie Zhu[1], Huayuan Zhou[1], Hao Tang[1], Hanyu Liu[1], and Chao Li[1,2⋆]

[1] University of Dundee, Dundee, UK
2619881@dundee.ac.uk, 2667057@dundee.ac.uk, 2485744@dundee.ac.uk, 2485644@dundee.ac.uk
[2] University of Cambridge, Cambridge, UK
cl647@cam.ac.uk

Abstract. Diffusion MRI (dMRI) provides white-matter microstructural information, but its associations with marker-specific histopathology are often assessed by global or regional scalar summaries, which may obscure focal, boundary-localized, or region-specific patterns. We propose a new dMRI-to-IHC spatial proxy mapping task under weak MRI–histology registration. Given dMRI and auxiliary structural priors, the goal is to generate inspectable tissue-scale virtual immunomarker proxy maps that preserve regional marker distributions and support fixed-protocol quantification. By retaining spatial heterogeneity and regional MRI–histology correspondence, this formulation moves beyond direct scalar readout regression, but remains challenging due to cross-domain gaps, weak alignment, and scale breaks in cascaded generation. To address these challenges, we propose TriSCoV-Net, a cross-scale verified framework that combines dMRI-conditioned proxy generation, PLI-assisted conditional super-resolution, support-domain constraint, and training-time cross-scale fine-tuning. Using a fixed ImageJ DAB workflow, we select 16.8 μm/pixel as the primary output and 4.2 μm/pixel as an auxiliary scale. On ODBB under LOSO evaluation, TriSCoV-Net improves image fidelity and fixed-protocol quantitative agreement over representative baselines.

Keywords: Diffusion MRI · Virtual Immunohistochemistry · Weak Registration · Conditional Super-Resolution · Polarization Light Imaging

1 Introduction

Neurological disorders are inherently cross-scale and spatially heterogeneous: demyelinating lesions reflect heterogeneous pathological mechanisms [15], astrocyte-related gliosis shows diverse cellular states and lesion-level spatial organization [14], and white-matter abnormalities exhibit spatial heterogeneity linked to risk factors, cognition, genetics, and atrophy [6]. Global marker or lesion

⋆ Corresponding author.

summaries may miss region-specific patterns, as different white-matter hyperintensity distributions have been associated with distinct cognitive phenotypes [11]. Yet clinical and cohort studies primarily rely on diffusion magnetic resonance imaging (dMRI) for macroscopic observation, whereas key tissue phenotypes arise at mesoscopic-to-microscopic scales. The Oxford Digital Brain Bank (ODBB) provides paired dMRI, microscopy, and tissue-section staining from the same specimens, offering a key testbed for spatial MRI–histology association [25].

In this work, we propose a new dMRI-to-immunohistochemistry (IHC) spatial proxy mapping task under weak MRI–histology registration. Given dMRI and auxiliary structural priors, the goal is to generate tissue-scale virtual immunomarker proxy maps that preserve regional marker distributions and support fixed-protocol quantification. The generated maps are not intended to replace physical IHC staining at the microscopic level. Instead, they serve as inspectable spatial proxy representations that bridge macroscopic dMRI measurements and tissue-scale marker distributions. Compared with direct scalar readout regression, this formulation retains spatial heterogeneity and regional MRI–histology correspondence before quantitative summaries are computed. It therefore enables a more transparent workflow: spatial proxy maps can first be inspected visually and then measured using the same downstream quantification protocol as real staining images.

Diffusion MRI indirectly reflects white-matter microstructure, whereas polarization light imaging (PLI) provides myelin-birefringence-derived cues of fiber orientation and tissue organization. Prior comparisons with histology and PLI support the biological relevance of dMRI-derived orientation dispersion [16], while biophysical models decompose dMRI signals into microstructural components such as neurite density and orientation dispersion [31]. We use slice-level PLI priors as auxiliary structural cues. This task involves three coupled challenges: i) *insufficient correspondence*, where cross-modal structural and biological links are incomplete; ii) *weak-registration geometry*, where deformation, missing regions, and partial overlap can induce structural drift or spurious staining; and iii) *cross-scale consistency*, where outputs at one scale may yield unstable statistics across scales. Existing image-to-image translation methods can learn cross-domain mappings but may be structurally unreliable under small-sample weak alignment [9,33]. GAN-based super-resolution sharpens textures [27], but the perception–distortion trade-off can reduce pixel-level fidelity [2]. MRI-to-microscopy transfer and synthesized myelin staining studies further indicate the need for structural priors and constrained cross-scale learning [20,19].

To address these challenges, we propose TriSCoV-Net, a Tri-Stage Scale-Consistent Verification Network. It injects dMRI spatial representations and slice-level PLI priors into the generator, providing fiber-orientation and myelin-related cues to reduce ill-posedness. We further use a support-domain constraint to suppress drift and spurious staining in evidence-insufficient regions. Finally, cross-scale joint optimization enforces a coherent scale chain for fixed-protocol quantification. Our main contributions are summarized as follows: i) we propose a new dMRI-to-IHC spatial proxy mapping task under weak MRI–histology

registration, aiming to generate inspectable virtual immunomarker proxy maps for tissue-scale quantification; ii) we develop TriSCoV-Net, which integrates dMRI-conditioned proxy generation, PLI-assisted conditional super-resolution, support-domain constraints, and cross-scale fine-tuning; iii) we evaluate the method on ODBB with LOSO testing and a fixed ImageJ DAB workflow, demonstrating improvements in image fidelity and quantitative agreement over representative baselines.

2 Method

2.1 Problem Definition

We study a dMRI-driven virtual immunomarker proxy-mapping task under weak registration, assisted by slice-level global PLI information. Because dMRI primarily encodes voxel-wise, population-level microstructural statistics rather than single-cell morphology or molecularly specific immune responses, reliable cell-level recovery from dMRI alone is not identifiable [10,17,29]. We therefore target tissue-scale immunomarker proxy maps for spatial inspection, robust quantification, and controlled comparison. Unlike direct scalar readout regression, this formulation preserves regional heterogeneity while avoiding non-interpretable fine textures when evidence is insufficient. To choose a verifiable output scale, we adopt a fixed ImageJ DAB quantification workflow [21] and measure myelin coverage (%Area) on real IHC across downsampling factors, observing inflection points near 16.8 μm/pix and 4.2 μm/pix. We set 16.8 μm/pix as the primary tissue-scale output and use 4.2 μm/pix only as an auxiliary training scale to improve cross-scale stability under weak registration and limited evidence.

2.2 Overview

Fig. 1 illustrates TriSCoV-Net, with inference and training-only fine-tuning separated for clarity. Given dMRI input $X \in \mathbb{R}^{B\times C\times H\times W}$, where B, C, H, and

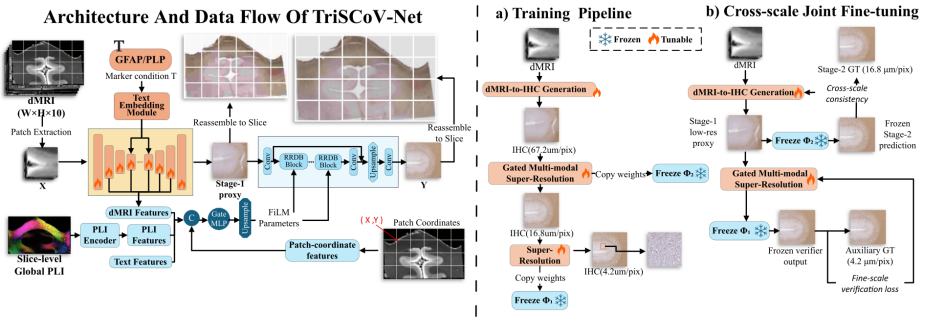

Fig. 1: Overview of TriSCoV-Net. **Left:** architecture and data flow of TriSCoV-Net. **Right:** training-only cross-scale verification, where frozen Φ_1 maps the Stage-2 output to the auxiliary finer scale for consistency supervision.

W denote batch size, dMRI channels, height, and width, and marker condition $T \in \mathbb{R}^{K\times D}$, where K is the number of markers and D is the embedding dimension, we jointly model glial fibrillary acidic protein (GFAP) and proteolipid protein (PLP). Stage 1 generates a low-resolution marker-conditioned proxy from dMRI. Stage 2 uses the Stage-1 proxy, shared dMRI encoder features, and slice-level global PLI features to produce the primary tissue-scale proxy. During training only, a frozen adversarial Residual-in-Residual Dense Block (RRDB) verifier Φ_1 maps the Stage-2 output to the auxiliary finer scale and provides cross-scale consistency signals for joint fine-tuning. Thus, Stage 3 denotes training-time cross-scale fine-tuning based on frozen verification.

2.3 dMRI and PLI Multi-modal Information Injection

To exploit dMRI-aligned structural cues and PLI-derived structural priors under weak registration, we introduce a globally gated spatial injection module. A shared Stage-1 encoder provides multi-scale dMRI skip features $\{s_i\}_{i=1}^{L}$ and a bottleneck feature F_{bot}, where L denotes the number of skip-feature levels. The marker condition T indexes GFAP or PLP and is encoded as a learned marker embedding e_T; together with patch-coordinate embedding e_{coord} and slice-level global PLI embedding e_{pli}, it provides auxiliary gate-conditioning signals. We form $c = [\text{GAP}(F_{\text{bot}}); e_T; e_{\text{coord}}; e_{\text{pli}}]$, where $\text{GAP}(\cdot)$ denotes global average pooling, and predict control variables by $(\alpha_{\text{dmri}}, \alpha_{\text{coord}}, \gamma, \beta, \ell) = \text{GateMLP}(c)$, where γ, β are FiLM parameters, $\alpha_{\text{dmri}}, \alpha_{\text{coord}}$ are spatial injection scales, and $\ell \in \mathbb{R}^L$ are skip-level logits. Smoothed scale weights are computed as $w = (1-\epsilon)\cdot \text{softmax}(\ell/\tau) + \epsilon/L$, where τ is a fixed softmax temperature and ϵ is a fixed uniform-mixing factor to avoid premature single-scale collapse. The dMRI spatial condition map is then

$$\mathbf{m}_{\text{dmri}} = \text{Conv}_{3\times3}\Big(\sum_{i=1}^{L} w_i \cdot \text{Up}\big(\text{Conv}_{1\times1}(\mathbf{s}_i)\big)\Big), \tag{1}$$

where $\text{Up}(\cdot)$ resizes features to a common target resolution (bilinear).

Given an intermediate RRDB feature $\mathbf{h}$, we first apply FiLM modulation $\mathbf{h} \leftarrow \mathbf{h} \odot (1+\boldsymbol{\gamma}) + \boldsymbol{\beta}$, where $\odot$ denotes element-wise multiplication, and then inject spatial priors in residual form:

$$\mathbf{h} \leftarrow \mathbf{h} + s_{\text{sp}}\big(\alpha_{\text{dmri}}\mathbf{m}_{\text{dmri}} + \alpha_{\text{coord}}\mathbf{m}_{\text{coord}}\big), \tag{2}$$

where $\mathbf{m}_{\text{coord}}$ is obtained by linear projection of $\mathbf{e}_{\text{coord}}$ followed by spatial broadcasting, and s_{sp} is a fixed scaling constant that bounds the injected update and keeps the RRDB trunk dominant.

2.4 Training Strategy

Weak registration with slice deformation, local missing regions, and incomplete field-of-view overlap weakens pixel-wise supervision and may induce spurious

staining in dMRI-unsupported regions. We therefore replace global pixel alignment with an evidence-aligned support-domain constraint: unsupported locations should revert to background. Concretely, we build a binary dMRI support mask M_{dmri} from the dMRI intensity map using a fixed threshold, and normalize the network output $\hat{Y}$ to $[-1, 1]$, where the white background corresponds to $+1$. The support-domain loss is

$$\mathcal{L}_{\mathrm{bg}} = \mathbb{E}_p\Big[(1 - M_{\mathrm{dmri}}) \cdot |\hat{Y} - 1|\Big], \tag{3}$$

where p indexes pixel locations. Because background is encoded as $+1$, this term penalizes non-background responses outside the dMRI support and suppresses drift or spurious staining under partial overlap. The constraint is applied at the supervision scale of each stage.

Cascaded generation is also prone to scale breaks: outputs that look plausible at one resolution can become statistically or structurally inconsistent after super-resolution, hurting controlled comparison and fixed-protocol quantification. To mitigate this, we perform cross-scale joint fine-tuning with frozen cross-scale links, enforcing consistency along the propagation path while avoiding writing non-identifiable high-frequency details back into the target output scale.

Specifically, we first adversarially train an RRDB super-resolution network Φ_1 to map real IHC from the tissue scale (16.8 μm/pix) to the auxiliary finer scale (4.2 μm/pix). During joint fine-tuning, Φ_1 is frozen and used only as a training-time verifier. We also keep a frozen copy of the Stage-2 generator as Φ_2 to form a neighboring-scale consistency link, conditioned on the same inputs as Stage-2. Let the Stage-1 prediction be $\hat{Y}^{(1)}$ and the Stage-2 prediction be $\hat{Y}^{(2)}$, with ground truths $Y^{(2)}$ (tissue scale) and $Y^{(3)}$ (finer scale). The cross-scale consistency loss is $\mathcal{L}_{\mathrm{scale}} = \mathcal{L}_{1\to2} + \lambda_{23}\mathcal{L}_{2\to3}$, where

$$\mathcal{L}_{1\to2} = \left\|\Phi_2\Big(\hat{Y}^{(1)}\Big) - Y^{(2)}\right\|_1, \qquad \mathcal{L}_{2\to3} = \left\|\Phi_1\Big(\hat{Y}^{(2)}\Big) - Y^{(3)}\right\|_1. \tag{4}$$

Here λ_{23} balances the neighboring-scale term and the verifier-based term. Intuitively, $\mathcal{L}_{1\to2}$ constrains Stage-1→Stage-2 propagation via frozen Φ_2, while $\mathcal{L}_{2\to3}$ aligns Stage-2 outputs with finer-scale phenotypes via frozen verifier Φ_1, reducing scale breaks while avoiding forced non-identifiable cell-level details at 16.8 μm/pix. We first perform stage-wise pretraining. For each stage $k \in \{1, 2\}$, the model is optimized by:

$$\mathcal{L}^{(k)}_{\mathrm{base}} = \mathcal{L}_{\mathrm{adv}} + \lambda_{\mathrm{rec}}\mathcal{L}_{\mathrm{rec}} + \lambda_{\mathrm{bg}}\mathcal{L}_{\mathrm{bg}} + \lambda_{\mathrm{ssim}}\mathcal{L}_{\mathrm{ssim}} + \lambda_{\mathrm{grad}}\mathcal{L}_{\mathrm{grad}} + \lambda_{\mathrm{perc}}\mathcal{L}_{\mathrm{perc}}, \tag{5}$$

and then jointly fine-tune Stage-1 and Stage-2 with $\mathcal{L}_{\mathrm{joint}} = \mathcal{L}^{(1)}_{\mathrm{base}} + \mathcal{L}^{(2)}_{\mathrm{base}} + \lambda_S\mathcal{L}_{\mathrm{scale}}$, where L_{adv} is the standard GAN loss [5], $L_{\mathrm{rec}} = \|\hat{Y} - Y\|_1$, L_{bg} follows Eq. (3), $L_{\mathrm{ssim}} = 1 - \mathrm{SSIM}(\hat{Y}, Y)$, $L_{\mathrm{grad}} = \|\nabla\hat{Y} - \nabla Y\|_1$, and L_{perc} is a feature-space perceptual loss. The λ terms weight the corresponding losses, and λ_S controls cross-scale fine-tuning.

3 Experiments

3.1 Experimental Setup

We evaluate on the ODBB [25] following the standard MRI–to–tissue-slice weak-registration paradigm, where dMRI and stained tissue sections are approximately aligned and used as supervisory signals [8]. Cross-subject generalization is assessed using a subject-level leave-one-subject-out (LOSO) protocol. During training, we sample 1,536 (128×128) training samples per LOSO fold for stable optimization, sampled only from training slices, while evaluation and statistical aggregation are performed at the subject level. The model is implemented in PyTorch and trained using AdamW with an initial learning rate of 2×10^{-4}, decayed linearly after 100 epochs until convergence. We report PSNR [1], SSIM [28], and LPIPS [32] on the LOSO test set as mean±std. For fixed-protocol tissue-scale quantification, we use the same fixed ImageJ DAB quantification workflow [21] for all methods to compute %Area, Orientation, and Coherency, and report MAE against ground truth. These readouts are fixed-protocol tissue-scale summaries rather than standalone biological endpoints; statistical significance is assessed using two-sided Welch's t-tests between each baseline and Ours on subject-level aggregated scores. We evaluate the Stage-1 outputs and the final tissue-scale outputs after Stage-3 fine-tuning, with Stage-2 serving as the intermediate scale for cross-scale joint training. GFAP and PLP predictions are evaluated against their matched target stains, and metrics are aggregated over marker outputs at the subject level. All baselines are trained and evaluated under the same splits, sample size, and preprocessing, using public implementations. Stage-1 baselines include GAN-based translators (Pix2pix [9], CycleGAN [33], CUT [18]), diffusion-based synthesis (Palette [22]), and a small-sample virtual histology method [20], while final-output baselines include non-learning interpolation (Bicubic [12]), diffusion-based super-resolution (SR3 [23]), Transformer-based super-resolution (SwinIR [13]), and GAN-based super-resolution (ESRGAN [27]).

3.2 Experimental Results and Analysis

Quantitative analysis. Under weak registration, we evaluate dMRI–to–IHC proxy prediction with auxiliary global PLI conditioning and report Stage-1 and final outputs in separate scale blocks in Tab. 1. Overall, our method yields a better balance between image consistency and biological quantification, improving perceptual quality while maintaining (or improving) reconstruction agreement and reducing deviations from GT-based quantitative readouts.

Qualitative analysis. In Fig. 2A, baselines under weak registration and cross-domain semantic gaps show background contamination, boundary spill-over, and structural instability. Our Stage-1 outputs instead show more stable contours and fewer spill-overs in low-support regions, consistent with conservative background reversion when evidence is insufficient. In Fig. 2B, our final outputs after Stage-3 fine-tuning better preserve weak granular expression patterns under a consistent background. The error maps show fewer high-error hotspots and weaker

Table 1: Quantitative comparison under weak registration. Stage-1 and final outputs are compared within separate scale blocks. P-values are baseline-vs-Ours.

Method	Image			Bio MAE			P-values (Image)			P-values (Bio)		
	PSNR↑	SSIM↑	LPIPS↓	Ori.↓	Coh.↓	Area↓	p-P	p-S	p-L	p-O	p-C	p-A
Stage-1: cross-domain synthesis												
Pix2pix [9]	15.87±1.07	0.626±0.014	0.603±0.016	41.02	0.245	14.29	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$	0.269	0.0004
CycleGAN [33]	18.83±1.08	0.652±0.026	0.407±0.033	33.26	0.315	12.65	0.018	0.0002	$< 10^{-4}$	$< 10^{-4}$	0.0070	0.0039
CUT [18]	15.04±1.17	0.549±0.032	0.627±0.030	32.16	0.364	13.47	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$	0.0003	$< 10^{-4}$	0.0005
Palette [22]	17.02±1.12	0.661±0.010	0.397±0.019	29.63	0.232	10.36	$< 10^{-4}$	0.002	$< 10^{-4}$	0.022	0.507	0.186
Pytlarz et al. [20]	17.86±1.15	0.671±0.021	0.574±0.019	30.27	0.347	13.50	$< 10^{-4}$	0.297	$< 10^{-4}$	0.0079	0.0002	0.0011
Ours	**19.46±0.91**	**0.677±0.024**	**0.369±0.028**	**25.55**	**0.203**	**8.75**	–	–	–	–	–	–
Final output after Stage-3 joint fine-tuning												
Bicubic [12]	29.11±1.70	0.735±0.046	0.365±0.010	13.18	0.156	9.69	0.0021	0.146	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$
SwinIR [13]	31.10±1.75	0.756±0.054	0.290±0.024	7.07	0.100	4.81	0.015	0.938	$< 10^{-4}$	0.0027	$< 10^{-4}$	0.460
ESRGAN [27]	30.97±1.64	0.745±0.052	0.215±0.010	7.71	0.124	6.59	0.028	0.429	$< 10^{-4}$	0.0008	$< 10^{-4}$	0.0094
SR3 [23]	26.34±0.82	0.704±0.026	0.257±0.006	18.00	0.096	6.59	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$	0.0019	0.018
Ours	**31.22±0.72**	**0.757±0.042**	**0.133±0.007**	**4.00**	**0.021**	**4.22**	–	–	–	–	–	–

Table 2: Ablation study. Settings are interpreted within their output scales. P-values compare each setting with the final frozen Tri-Scale model.

Setting	Image			Bio MAE			P-values (Image)			P-values (Bio)		
	PSNR↑	SSIM↑	LPIPS↓	Ori.↓	Coh.↓	Area↓	p-P	p-S	p-L	p-O	p-C	p-A
Base	16.94±0.96	0.673±0.054	0.685±0.013	33.74	0.264	12.64	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$
+ $\mathcal{L}_{\text{bg}}$	19.46±0.91	0.677±0.024	0.369±0.028	25.55	0.203	8.75	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$
+ SR	25.03±0.71	0.699±0.035	0.300±0.014	12.26	0.187	6.15	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$	0.013
+ MRI-FiLM	26.45±0.75	0.705±0.048	0.273±0.024	8.64	0.153	5.96	$< 10^{-4}$	0.0003	$< 10^{-4}$	$< 10^{-4}$	$< 10^{-4}$	0.019
+ MRI&PLI-FiLM	28.34±0.77	0.725±0.048	0.252±0.008	6.26	0.124	5.44	$< 10^{-4}$	0.024	$< 10^{-4}$	0.0057	$< 10^{-4}$	0.101
+ Tri-Scale (unfrozen)	31.13±0.74	0.756±0.051	0.205±0.024	4.21	0.085	4.65	$< 10^{-4}$	0.129	$< 10^{-4}$	0.710	$< 10^{-4}$	0.541
+ Tri-Scale (frozen)	**31.22±0.72**	**0.757±0.042**	**0.133±0.007**	**4.00**	**0.021**	**4.22**	–	–	–	–	–	–

boundary error bands, indicating more concentrated errors and reduced structural drift. This agrees with the trends in LPIPS and fixed-protocol quantitative error metrics.

Ablation study. Tab. 2 shows that the proposed components are complementary and drive progressive gains. Adding the support-domain constraint, SR refinement, and multi-modal FiLM injection consistently improves image quality and fixed-protocol tissue-scale quantification errors in absolute terms. Frozen Tri-Scale joint fine-tuning is more stable than the unfrozen variant, supporting frozen cross-scale links for constrained propagation and reduced scale breaks.

4 Discussion

Reliable regions and failure modes. Identifiability in dMRI–to–IHC synthesis under weak registration is spatially non-uniform. From an evidence-support perspective, outputs fall into three regimes: i) high-support regions, where most methods recover coarse contours and differences mainly appear in weak granular patterns and texture statistics; ii) boundary regions, where geometric mismatch amplifies cross-domain ambiguity and baselines often trade consistency for over-smoothing, suppressing focal signals; and iii) low-support regions, where spurious staining and spill-over are most likely. In our small-sample, weak-registration setting, the tested diffusion-based generation/SR baselines are more prone to

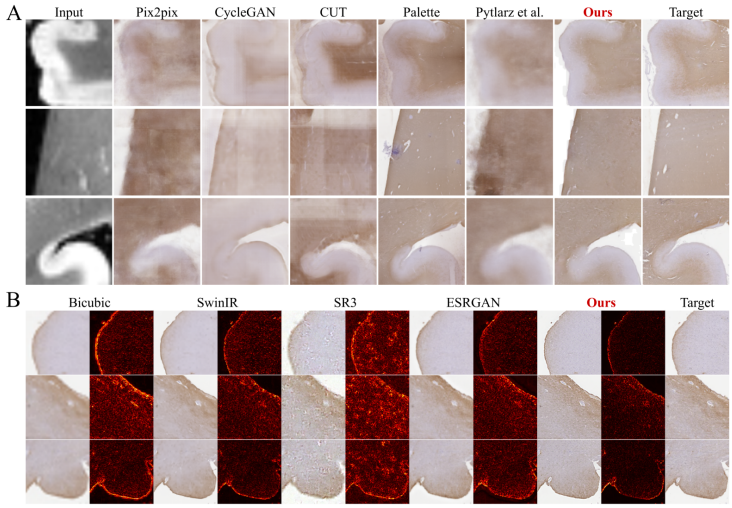

Fig. 2: Qualitative comparison of different models on 3 ROIs, covering i) tissue boundaries, ii) regions with weak granular patterns, and iii) white-matter transition areas. **A:** Stage-1 cross-domain synthesis examples. **B:** our final outputs after Stage-3 fine-tuning and corresponding error maps.

stochastic expressions and scale instability, whereas adversarial learning with conservative constraints yields more stable, controllable outputs. This empirical pattern supports our GAN-based core generator.

Limitations and future directions. Our goal is a quantification-oriented tissue-scale immunomarker proxy, not single-cell recovery. Conservative constraints improve verifiability but may be overly conservative for extremely weak, small focal signals. Evidence remains limited by open-dataset size and approximate registration quality; although LOSO and subject-level aggregation reduce overfitting risks, conclusions should be interpreted within the current data distribution and alignment quality. We use PLI as a microstructural prior. A key next step is to replace this input with more deployable fiber-based models [24,7,30] that introduce precise myelin-related in-vivo structural priors, preserving global microstructural conditioning while moving toward a fully in-vivo pipeline [26,4,3].

5 Conclusion

We present a new proxy-mapping task that moves beyond direct scalar readout regression by preserving regional marker distributions and spatial MRI–histology correspondence. TriSCoV-Net integrates dMRI-conditioned proxy generation, PLI-assisted conditional super-resolution, support-domain constraint, and training-time cross-scale fine-tuning to address cross-domain gaps, weak alignment, and scale inconsistency. Using a fixed ImageJ DAB workflow, we evaluate both image fidelity and quantitative agreement. Overall, this work provides

clearer scale criteria, evaluation protocols, and comparison baselines for dMRI-driven virtual immunomarker proxy mapping, while clarifying its limitations for tissue-scale MRI–histology association.

References

1. Recommendation itu-t j.340: Reference algorithm for computing peak signal-to-noise ratio (psnr) of a processed video sequence with compensation for constant spatial and temporal shifts. Tech. Rep. J.340, International Telecommunication Union (ITU-T) (Jun 2010)
2. Blau, Y., Michaeli, T.: The perception-distortion tradeoff. In: CVPR (2018)
3. de Boer, J.F., Hitzenberger, C.K., Yasuno, Y.: Polarization sensitive optical coherence tomography – a review [invited]. Biomedical Optics Express **8**(3), 1838–1873 (2017). https://doi.org/10.1364/BOE.8.001838
4. DePaoli, D., Côté, D.C.C., Bouma, B.E., Villiger, M.: Endoscopic imaging of white matter fiber tracts using polarization-sensitive optical coherence tomography. NeuroImage **264**, 119755 (2022). https://doi.org/10.1016/j.neuroimage.2022.119755
5. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances in Neural Information Processing Systems. vol. 27, pp. 2672–2680 (2014)
6. Habes, M., Sotiras, A., Erus, G., Toledo, J.B., Janowitz, D., Wolk, D.A., Shou, H., Bryan, N.R., Doshi, J., Völzke, H., Schminke, U., Hoffmann, W., Resnick, S.M., Grabe, H.J., Davatzikos, C.: White matter lesions: Spatial heterogeneity, links to risk factors, cognition, genetics, and atrophy. Neurology **91**(10), e964–e975 (2018). https://doi.org/10.1212/WNL.0000000000006116
7. Howard, A.F., Mollink, J., Kleinnijenhuis, M., Pallebage-Gamarallage, M., Bastiani, M., Cottaar, M., Miller, K.L., Jbabdi, S.: Joint modelling of diffusion mri and microscopy. NeuroImage **201**, 116014 (2019)
8. Huszár, I.N., Dudzik, M., Hess, M., et al.: Template-based imaging registration library (tirl). NeuroImage **265**, 119791 (2023)
9. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: CVPR (2017)
10. Jelescu, I.O., Budde, M.D.: Design and validation of diffusion mri models of white matter. Frontiers in Physics **5**, 61 (2017). https://doi.org/10.3389/fphy.2017.00061
11. Jiménez-Balado, J., Corlier, F., Habeck, C., Stern, Y., Eich, T.: Effects of white matter hyperintensities distribution and clustering on late-life cognitive impairment. Scientific Reports **12**, 1955 (2022). https://doi.org/10.1038/s41598-022-06019-8
12. Keys, R.G.: Cubic convolution interpolation for digital image processing. IEEE Transactions on Acoustics, Speech, and Signal Processing **29**(6), 1153–1160 (1981)
13. Liang, J., Cao, J., Sun, G., Zhang, K., Van Gool, L., Timofte, R.: Swinir: Image restoration using swin transformer. In: ICCV Workshops (2021)
14. Lo, C.H., Skarica, M., Mansoor, M., Bhandarkar, S., Toro, S., Pitt, D.: Astrocyte heterogeneity in multiple sclerosis: Current understanding and technical challenges. Frontiers in Cellular Neuroscience **15**, 726479 (2021). https://doi.org/10.3389/fncel.2021.726479
15. Lucchinetti, C., Brück, W., Parisi, J., Scheithauer, B., Rodriguez, M., Lassmann, H.: Heterogeneity of multiple sclerosis lesions: Implications for the pathogenesis of demyelination. Annals of Neurology **47**(6), 707–717 (2000). https://doi.org/10.1002/1531-8249(200006)47:6<707::AID-ANA3>3.0.CO;2-Q

16. Mollink, J., Kleinnijenhuis, M., van Cappellen van Walsum, M.M.A., et al.: Evaluating fibre orientation dispersion in white matter: Comparison of diffusion mri, histology and polarized light imaging. NeuroImage **157**, 561–574 (2017)
17. Novikov, D.S., Fieremans, E., Jespersen, S.N., Kiselev, V.G.: Quantifying brain microstructure with diffusion MRI: Theory and parameter estimation. NMR in Biomedicine **32**(4), e3998 (2019). https://doi.org/10.1002/nbm.3998
18. Park, T., Efros, A.A., Zhang, R., Zhu, J.Y.: Contrastive learning for unpaired image-to-image translation (cut). In: ECCV (2020)
19. Pittayapong, S., Hametner, S., Bachratá, B., et al.: Synthesized myelin and iron stainings from 7t multi-contrast mri via deep learning. NeuroImage **318**, 121417 (2025)
20. Pytlarz, M., Onicas, A., Crimi, A.: Style transfer between microscopy and magnetic resonance imaging via generative adversarial network in small sample size settings. In: ICIP (2023)
21. Ruifrok, A.C., Johnston, D.A.: Quantification of histochemical staining by color deconvolution. Analytical and Quantitative Cytology and Histology **23**(4), 291–299 (2001)
22. Saharia, C., Chan, W., Saxena, S., et al.: Palette: Image-to-image diffusion models. ACM Transactions on Graphics (TOG) (2022), sIGGRAPH
23. Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D.J., Norouzi, M.: Sr3: Image super-resolution via iterative refinement. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) **44**(12), 121–132 (2022)
24. Tang, H., Chen, H., Li, C.: Generalized spherical neural operators: Green's function formulation. In: The Fourteenth International Conference on Learning Representations (2026), https://openreview.net/forum?id=XkGjzSDTnm
25. Tendler, B.C., Hanayik, T., Ansorge, O., et al.: The digital brain bank, an open access platform for post-mortem imaging datasets. eLife **11**, e73153 (2022)
26. Wang, H., Akkin, T., Magnain, C., Wang, R., Dubb, J., Kostis, W.J., Yaseen, M.A., Cramer, A., Sakadzic, S., Boas, D.: Polarization sensitive optical coherence microscopy for brain imaging. Optics Letters **41**(10), 2213–2216 (2016). https://doi.org/10.1364/OL.41.002213
27. Wang, X., Yu, K., Wu, S., et al.: Esrgan: Enhanced super-resolution generative adversarial networks. In: ECCV Workshops (2018)
28. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: From error visibility to structural similarity. IEEE Transactions on Image Processing **13**(4), 600–612 (2004). https://doi.org/10.1109/TIP.2003.819861
29. West, K.L., Kelm, N.D., Carson, R.P., Does, M.D.: A revised model for estimating g-ratio from MRI. NeuroImage **125**, 1155–1158 (2016). https://doi.org/10.1016/j.neuroimage.2015.08.017
30. Zeng, R., Lv, J., Wang, H., Zhou, L., Barnett, M., Calamante, F., Wang, C.: Fod-net: A deep learning method for fiber orientation distribution angular super resolution. Medical Image Analysis **79**, 102431 (2022)
31. Zhang, H., Schneider, T., Wheeler-Kingshott, C.A., Alexander, D.C.: Noddi: practical in vivo neurite orientation dispersion and density imaging of the human brain. NeuroImage **61**(4), 1000–1016 (2012)
32. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 586–595 (2018)
33. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: ICCV (2017)