This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

# Supplementary Material for "Joint Multi-Velocity FlowVN Reconstruction with Flow-Aware Supervision for Accelerated 4D Flow MRI"

Veeranatesan S[1], Hariharan Devarajan[2], Pradeeba Sridar[3]

[1, 2, 3] Department of Medical Sciences and Technology, IIT Madras, India
md23b031@smail.iitm.ac.in, md23b011@smail.iitm.ac.in,
spradeeba@iitm.ac.in

## 1 Exploratory Development and Selection of the Flow-Aware Loss

We performed a sequential set of exploratory experiments to develop and select the flow-aware training objective for the joint multi-velocity FlowVN reconstruction framework. All exploratory experiments used the same 500-sample subset and a fixed training schedule of 20 epochs, with identical reconstruction architecture and optimization settings unless otherwise stated. Experiments were evaluated at both acceleration factors, R = 10 and R = 50.

The joint-reconstruction loss-development sequence started from the complex L1 reconstruction objective, in which the four velocity encodings were processed jointly but supervision was based only on the complex-valued image reconstruction error. Flow-related loss terms were then introduced to investigate whether explicit supervision of the reconstructed velocity field could improve flow-domain accuracy. The first modification added a velocity-magnitude loss with weight 0.1 while retaining the complex loss with weight 1.0. A subsequent experiment introduced both velocity-magnitude and directional losses using weights (1.0, 0.05, 0.2). Further experiments varied the angular contribution and introduced vessel masking to focus directional supervision on regions containing relevant flow information.

During this process, the angular-loss formulation was revised to an explicit vector-angle formulation based on the arccosine of the normalized dot product between reconstructed and reference velocity vectors. An auxiliary PhysicsCNN component present in some early exploratory implementations was subsequently removed before the final loss-refinement experiments so that the final objective consisted only of the explicitly defined complex-image, velocity-magnitude, and angular losses.

Following these modifications, several combinations of the revised velocity-magnitude and angular losses were evaluated. Configurations with angular weight 0.5 and velocity-magnitude weights of 0.05 and 0.1 were compared. The configuration with weights (wc, wv, wa) = (1.0, 0.1, 0.5) achieved the lowest angular error among the

evaluated configurations at both acceleration factors while maintaining competitive image- and velocity-domain metrics. It was therefore selected for final training. A relative-speed variant using the same nominal weights was also evaluated but did not improve performance.

The complete sequence of exploratory experiments is provided in Supplementary Table S1. The reported weights are the raw loss weights before normalization; during training, the weights were normalized by their sum. These experiments were intended for loss development and configuration selection rather than statistical comparison of independently optimized models. Each configuration was trained once on the same 500-sample subset for 20 epochs, without repeated random-seed experiments.

The selected (1.0, 0.1, 0.5) configuration was subsequently used for final training on the full available training dataset for 18 epochs and evaluation at R = 10 and R = 50.

**Table 1.** Complete sequence of exploratory loss-development experiments used to select the final flow-aware objective.

| **Configuration / experiment** | **Weights ($w_c$,$w_v$,$w_a$)** | **Key Modification** | **SSIM** | **nRMSE** | **RelErr** | **AngErr (°)** |
|---|---|---|---|---|---|---|
| Original FlowVN baseline, R10 | — | Original FlowVN; independent encoding reconstruction | 0.871 | 0.126 | 0.637 | 59.61 |
| Original FlowVN baseline, R50 | — | | 0.807 | 0.147 | 0.813 | 72.30 |
| Joint + Complex L1, R10 | (1, 0, 0) | Joint reconstruction of all 4 encodings; complex L1 objective | 0.889 | 0.107 | 0.684 | 58.97 |
| Joint + Complex L1, R50 | (1, 0, 0) | | 0.844 | 0.112 | 0.839 | 83.98 |
| Hybrid: velocity term, R10 | (1, 0.1, 0) | Added velocity-magnitude supervision | 0.891 | 0.100 | 0.749 | 64.83 |
| Hybrid: velocity term, R50 | (1, 0.1, 0) | | 0.807 | 0.153 | 0.881 | 85.77 |

| | | | | | | |
|---|---|---|---|---|---|---|
| Hybrid: velocity + angle, R10 | (1, 0.05, 0.2) | Reduced velocity weight and added directional supervision | 0.871 | 0.122 | 0.704 | 56.62 |
| Hybrid: velocity + angle, R50 | (1, 0.05, 0.2) | | 0.819 | 0.135 | 0.816 | 74.26 |
| Hybrid: angle only, R10 | (1, 0, 1) | Removed velocity term; angle supervision retained | 0.849 | 0.147 | 0.692 | 58.91 |
| Hybrid: angle only, R50 | (1, 0, 1) | | 0.827 | 0.130 | 0.781 | 73.62 |
| Hybrid: masked angle, R10 | (1, 0, 1) | Applied vessel mask to directional loss | 0.819 | 0.176 | 0.611 | 57.43 |
| Hybrid: masked angle, R50 | (1, 0, 1) | | 0.810 | 0.144 | 0.774 | 74.81 |
| Hybrid: masked + revised angle, R10 | (1, 0, 1) | Replaced angle-loss formulation with revised angular loss | 0.819 | 0.177 | 0.621 | 56.80 |
| Hybrid: masked + revised angle R50 | (1, 0, 1) | | 0.827 | 0.129 | 0.762 | 73.11 |
| Hybrid: masked + revised angle, no PhysicsCNN, R10 | (1, 0, 1) | Removed PhysicsCNN | 0.829 | 0.172 | 0.587 | 58.17 |
| Hybrid: masked + revised angle, no PhysicsCNN, R50 | (1, 0, 1) | | 0.830 | 0.126 | 0.734 | 74.15 |
| Hybrid: masked + revised angle, no PhysicsCNN, R10 | (1, 0, 0.5) | Reduced angle weight | 0.869 | 0.133 | 0.551 | 58.20 |
| Hybrid: masked + revised angle, no PhysicsCNN, R50 | (1, 0, 0.5) | | — | — | — | — |
| Hybrid: masked + revised angle + speed, no PhysicsCNN, R10 | (1, 0.05, 0.5) | Revised speed loss; revised angle loss; | 0.854 | 0.147 | 0.591 | 57.36 |

| | | | | | | |
|---|---|---|---|---|---|---|
| | | PhysicsCNN removed | | | | |
| Hybrid: masked + revised angle + speed, no PhysicsCNN, R50 | (1, 0.05, 0.5) | | 0.845 | 0.117 | 0.743 | 73.51 |
| Selected hybrid, R10 | (1, 0.1, 0.5) | Revised speed loss + revised angle loss; PhysicsCNN removed | 0.878 | 0.124 | 0.559 | 54.06 |
| Selected hybrid, R50 | (1, 0.1, 0.5) | | 0.849 | 0.109 | 0.738 | 71.89 |
| Selected hybrid: relative speed variant, R10 | (1, 0.1, 0.5) | Relative speed loss instead of final speed loss | 0.856 | 0.146 | 0.636 | 55.86 |
| Selected hybrid: relative speed variant, R50 | (1, 0.1, 0.5) | | 0.843 | 0.115 | 0.826 | 74.47 |