Supplementary Material

S1 Data composition by training stream

Table S1 gives the per-stream composition of the 138-subject aortic cohort described in Sec. 3 of the main text. The three groups are disjoint at the subject level.

Table S1. Per-stream composition of the aortic cohort.

Stream	Center	Scanner	Field	Subjects
Supervised	Center007	GE Architect	3.0 T	67
	Center008	GE Premier	3.0 T	33
	Center010	Philips Ingenia	3.0 T	5
	Center012	Philips Ambition	1.5 T	9
	Center012	Philips Ingenia	3.0 T	5
	Supervised total			**119**
Self-supervised	Center011	UIH uMR890	3.0 T	**6**
Held out	Center007	GE Architect	3.0 T	8
	Center008	GE Premier	3.0 T	3
	Center010	Philips Ingenia	3.0 T	1
	Center012	Philips Ambition	1.5 T	1
	Held-out total			**13**
Cohort total				**138**

S2 Data-consistency residual on acquired samples

The data-consistency update is a learned reweighting of the k-space residual rather than a hard projection, so agreement with the measurements is not guaranteed by construction and has to be measured. Table S2 reports the relative residual on the acquired samples after the final unrolled stage, $\|\mathcal{A}\hat{x} - k\|_2/\|k\|_2$, for the proposed model on the local held-out cohort. The reconstruction stays within 9% of the measured k-space across the full acceleration range, consistent with the claim in the main text that relative velocity phase is set by per-encoding data consistency. Residuals on the out-of-distribution sets are larger, and are reported below.

Table S2. Relative residual on acquired samples after the final unrolled stage, proposed model, local held-out cohort.

R	10	20	30	40	50	mean
$\|\mathcal{A}\hat{x} - k\|_2/\|k\|_2$	0.056	0.063	0.070	0.079	0.088	0.071

Out-of-distribution residual. Figure S1 evaluates the same quantity on the two generalization sets, for the proposed model and for the FlowVN baseline trained on the same split. On the cross-site set the residual is 0.085 for the proposed model against 0.137 for the baseline, a 38% reduction, and the margin widens monotonically with acceleration, from 32% at R=10 to 43% at R=50. On the

cross-organ set the residuals are 0.059 and 0.171, a 66% reduction, and the proposed model has the lower residual in every one of the forty cases.

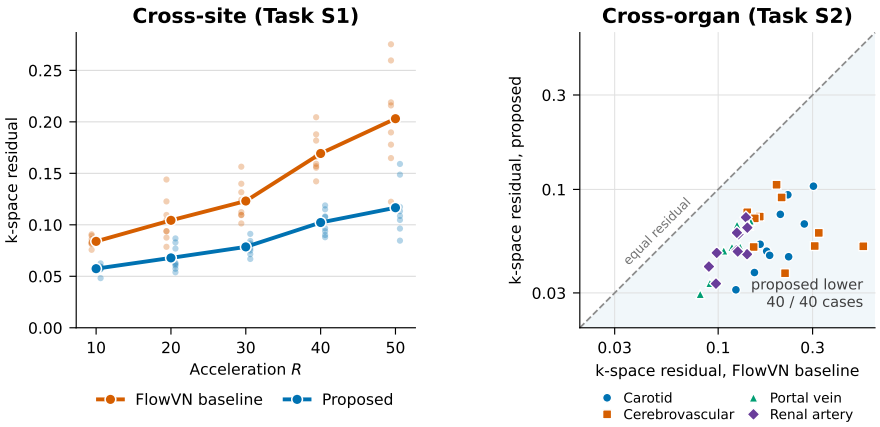

Fig. S1. K-space residual on the out-of-distribution sets. *Left, cross-site:* faint points are individual subjects, solid lines their per-R means. *Right, cross-organ:* each point is one case, reconstructed by both models; points below the dashed line are cases where the proposed model agrees more closely with the acquired samples. Axes are logarithmic.

S3 Reconstruction at the ends of the acceleration range

Figures S2 and S3 show the reconstructions and the associated error maps at R=10 and R=50 for the subject used in the main text (Center007, GE 3.0 T Architect, P048), at the same slice and the same peak-systole frame. At R=10 all three reconstruction methods recover the vessel lumen and separate mainly on velocity.

At R=50 the zero-filled and CS-LLR reconstructions have lost the velocity information almost entirely; the baseline retains its gross structure, and the proposed model remains ahead of it on every metric.

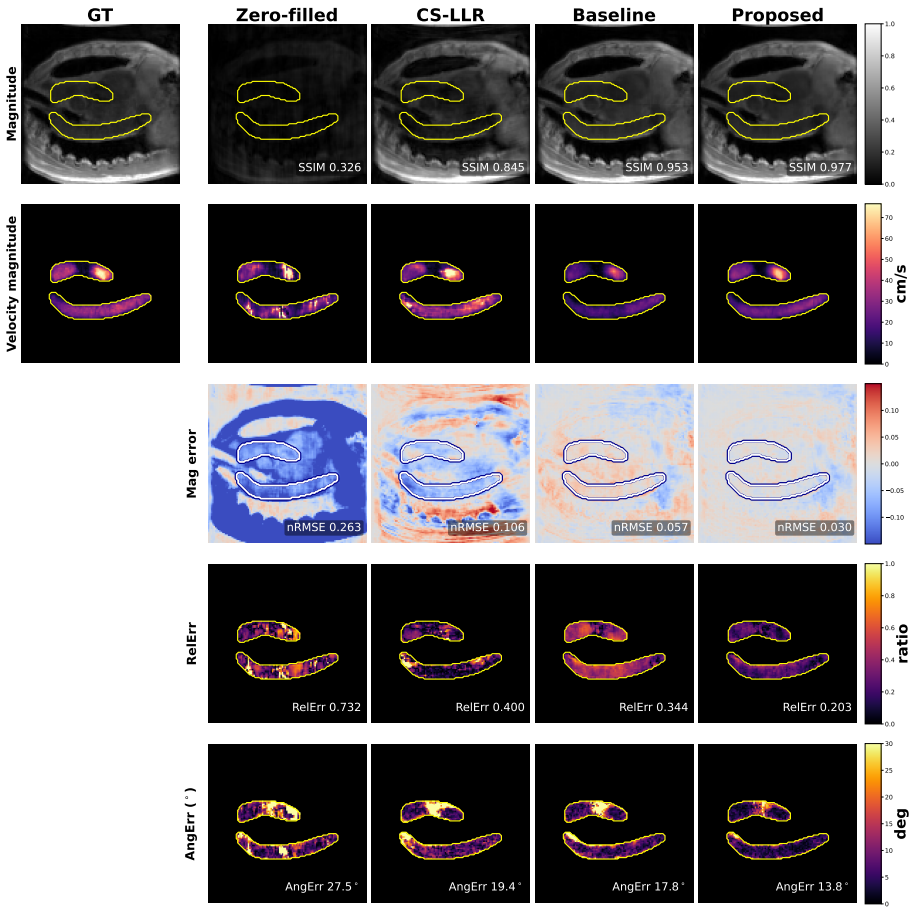

Fig. S2. Qualitative comparison of four approaches at R=10, at a peak-systole frame of a held-out subject. The segmentation ROI is outlined on each panel.

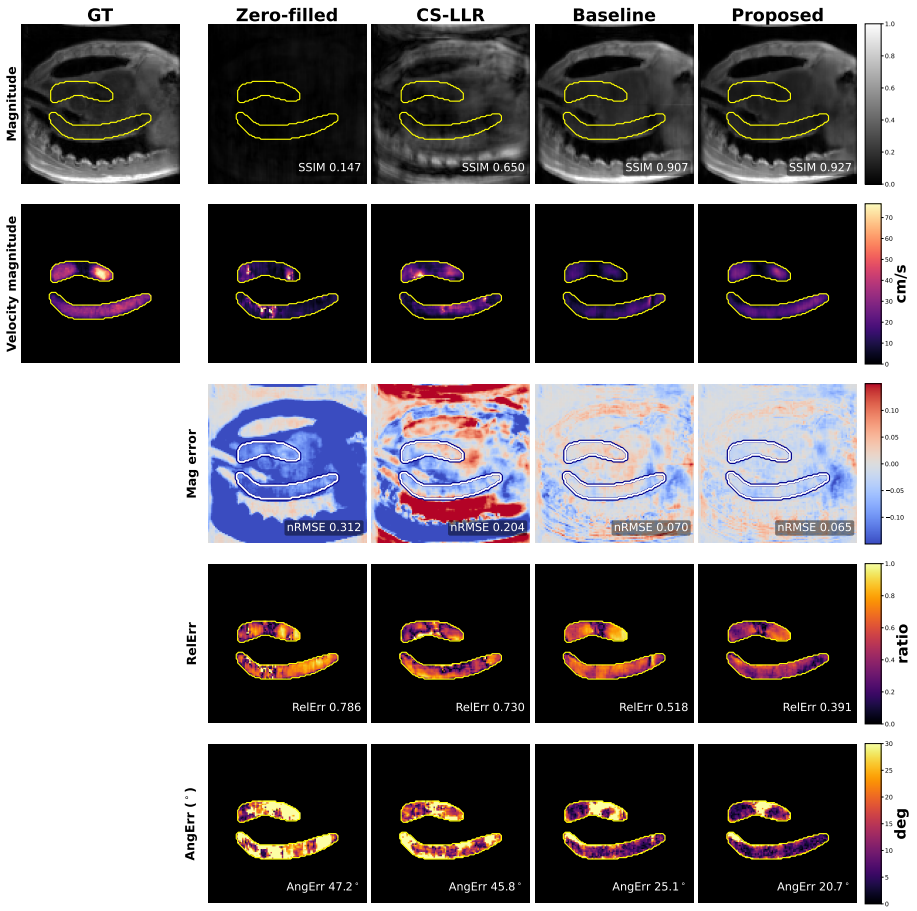

Fig. S3. Qualitative comparison of four approaches at R=50, at a peak-systole frame of a held-out subject. The segmentation ROI is outlined on each panel.

S4 Qualitative zero-shot generalization

Figures S4 and S5 show the frozen model on the cross-site and cross-organ sets scored in Table 1 of the main text, alongside the FlowVN baseline trained on the same split. Neither set has a released reference, so the panels carry no ground-truth column, no error maps and no per-case metric.

Each row of Fig. S4 is a different subject at a different acceleration, from R=10 to R=50; each row of Fig. S5 is a different non-aortic anatomy, again from a different subject. Every panel is cropped to the evaluation ROI of its case. Velocity is recovered from the phase difference between each of the three directional encodings x_j and the flow-compensated reference encoding x_0, scaled by the case's own velocity encoding sensitivity, $v_j = \angle(x_j x_0^*)\,\mathrm{venc}/\pi$, and the panels show $\|v\|_2$ in $\mathrm{cm\,s^{-1}}$.

Magnitude is near-identical between the two models at every acceleration, while the velocity fields separate progressively from R=30 upward, which locates the difference between the two training regimes in relative phase rather than in image magnitude. Without a reference, however, neither the agreement in magnitude nor the separation in velocity establishes which reconstruction is the more accurate.

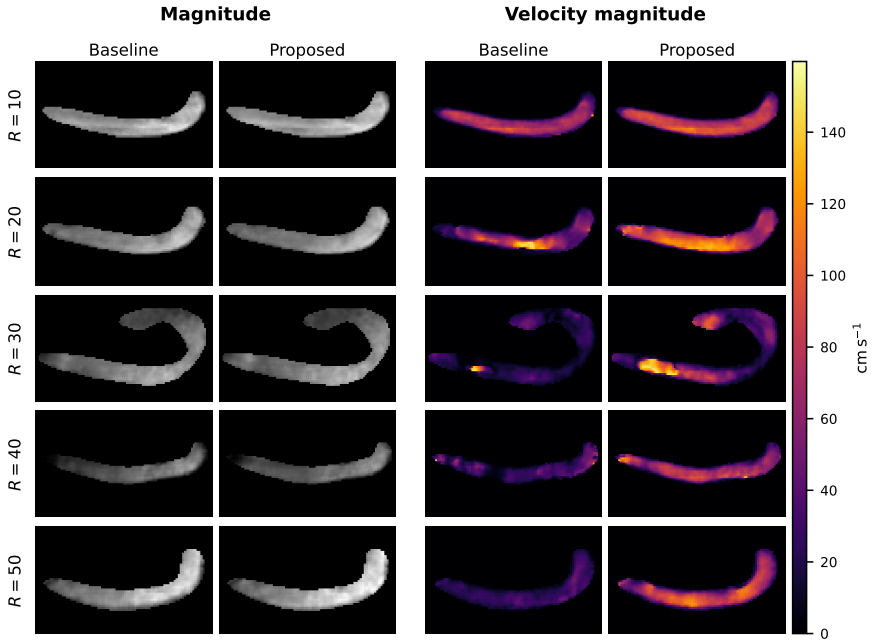

Fig. S4. Cross-site (Task S1). Aorta at an unseen site and scanner model (Center006, UIH 3.0 T uMR790). The challenge releases each validation subject at a single acceleration, so the rows sweep R across subjects: P027 (R=10), P034 (R=20), P012 (R=30), P028 (R=40), P025 (R=50).

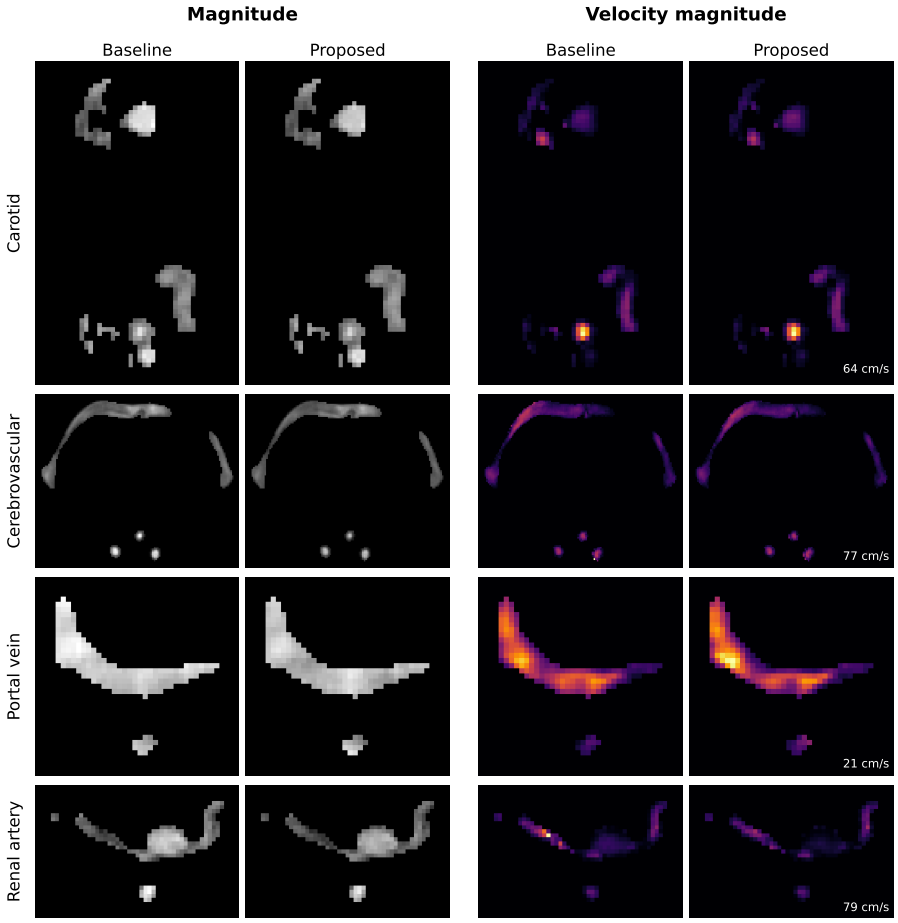

Fig. S5. Cross-organ (Task S2). Non-aortic anatomy at R=30, one subject per organ: carotid (Center008, GE 3.0 T Premier, P001), cerebrovascular (Center001, Siemens 3.0 T VIDA, P002), portal vein (Center002, UIH 5.0 T Jupiter, P001) and renal artery (Center014, Siemens 3.0 T VIDA, P007). The plane shown is the one with a near-largest in-ROI cross-section.