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.

CLEAR: Complex Learned Explicit Analytical Regularization for Ultra-Accelerated 4D Flow CMR Reconstruction

German Shâma Wache[0009−0001−1386−1702] and Sebastian Neumayer[0000−0002−4169−8447]

Chemnitz University of Technology, Germany
{german-shama.wache,sebastian.neumayer}@mathematik.tu-chemnitz.de

Abstract. While compressed-sensing regularizers enable interpretable reconstruction of 4D Flow CMR through transparent variational objectives, their hand-crafted nature is too restrictive under high acceleration. State-of-the-art learning-based approaches mitigate this, but typically encode regularization implicitly through unrolled network modules, which limits their interpretability. To address this limitation, we propose CLEAR, designed to combine the interpretability of compressed sensing with the flexibility of learned models. To the best of our knowledge, it is the first learned regularizer for a 4D reconstruction task. In the ultra-accelerated 10×–50× regime of the CMRx4DFlow2026 challenge, CLEAR outperforms compressed sensing locally low-rank (LLR) and the popular variational network FlowVN, while using less than 10k parameters and preserving an interpretable regularization structure.

Keywords: inverse problems · learned regularizer · medical imaging · 4D Flow CMR · variational reconstruction

1 Introduction

Four-dimensional flow cardiovascular magnetic resonance (4D Flow CMR) provides non-invasive, time-resolved three-dimensional blood velocity fields, ultimately supporting the assessment of cardiovascular diseases [1, 2]. Unlike conventional MRI, the object of interest is not only the complex image magnitude but also the blood velocity vector $\mathbf{v}(\mathbf{r}, t) = (\mathbf{v}_1(\mathbf{r}, t), \mathbf{v}_2(\mathbf{r}, t), \mathbf{v}_3(\mathbf{r}, t))^\top \in \mathbb{R}^3$, defined at spatial position $\mathbf{r}$ and cardiac phase $t \in \{1, \ldots, N_t\}$. This velocity is typically encoded into the phase of $N_e = 4$ complex images $\{\mathbf{x}_i\}_{i=0}^3$ according to

$$\mathbf{x}_i(\mathbf{r}, t) = \mathbf{x}_0(\mathbf{r}, t) \exp\left(\jmath\pi \frac{\mathbf{v}_i(\mathbf{r}, t)}{v_{\text{enc}}}\right), \qquad i = 0, 1, 2, 3, \tag{1}$$

where v_{enc} is the velocity value corresponding to a phase of $\pm\pi$, $\mathbf{x}_0$ the flow-compensated reference image with $\mathbf{v}_0 = 0$, and $\mathbf{x}_1, \mathbf{x}_2$, and $\mathbf{x}_3$ encode the three velocity directions. The velocity field is then obtained from the phase differences between the velocity-encoded images and the reference image.

The main limitation of 4D Flow CMR is its acquisition time [3], since four velocity encodings must be acquired over multiple cardiac phases. A common acceleration strategy is to acquire only a fraction of the k-space and to reconstruct the complex images before velocity extraction. For a single velocity encoding $i \in \{0, 1, 2, 3\}$, let $\mathbf{x}_i \in \mathbb{C}^{N_t \times N_x \times N_y \times N_z}$ denote the discretized image of interest on a $N_x \times N_y \times N_z$ spatial grid across N_t time frames. Further, we require the spatial Fourier transform $\mathcal{F}$ acting time-frame-wise, the scanner N_c-channel coil sensitivity operator $\mathbf{S}$, the undersampling mask $\mathbf{M}_R$ corresponding to an acceleration factor of R, and the measurement noise $\boldsymbol{\eta}$. Assuming Cartesian sampling, the corresponding accelerated k-space measurement $\mathbf{y}_i \in \mathbb{C}^{N_t \times N_c \times N_x \times N_y \times N_z}$ obeys

$$\mathbf{y}_i = \mathbf{M}_R \mathcal{F} \mathbf{S} \mathbf{x}_i + \boldsymbol{\eta}. \tag{2}$$

Reconstructing $\mathbf{x}_i$ from $\mathbf{y}_i$ is ill-posed since only a fraction $1/R$ of the k-space is acquired. The popular variational approach addresses this inverse problem by minimizing an objective consisting of a data-fidelity term and a carefully chosen regularizer $\mathcal{R}$ (promoting desired properties of $\mathbf{x_i}$), namely

$$\widehat{\mathbf{x}}_i = \arg\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{M}_R \mathcal{F} \mathbf{S} \mathbf{x} - \mathbf{y}_i\|_2^2 + \mathcal{R}(\mathbf{x}). \tag{3}$$

Classical choices are compressed sensing (CS) regularizers, such as total variation (TV) [4] or the locally low-rank (LLR) [5]. However, their hand-crafted structure becomes restrictive at very high acceleration. Recent CMR reconstruction methods improve expressiveness by complex learned architectures. This includes region-focused transformer GAN reconstruction [6], personalized federated reconstruction [7], and unrolled low-rank and sparse networks [8]. Dynamic settings additionally incorporate recurrent models that exploit temporal dependencies [9], spatio-temporal attention [10], and deep separable spatio-temporal learning [11]. More recent advances include Mamba-based global modeling with uncertainty estimation [12], physics-informed synthetic data for improved generalization [13], and diffusion-prior implicit neural representations [14]. For 4D Flow CMR, FlowVN [15] introduced an unrolled variational network with learned spatio-temporal filters, while FlowMRI-Net [16] proposed a self-supervised unrolled network with complex recurrent modules. These methods are expressive, but their priors are typically implicit in network modules, limiting interpretability. To compromise between interpretability and expressivity, a middle-ground route consisting of learning an explicit regularizer for (3) emerged in 2D inverse problems; see [21] for an overview. A notable example is the learnable fields-of-experts regularizer [17], which was repopularized in [18, 19] with a recent extension to 3D in [20]. To the best of our knowledge, this route has not yet been explored for 4D inverse problems. Our contributions are thus the following:

- We propose a learned regularizer with analytical gradient, enabling complex-valued 4D Flow CMR reconstruction based on (3).
- We condition the regularizer on the acceleration factor, allowing it to handle a wide and continuous range of undersampling rates.

– We evaluate the method in an ultra-accelerated 10×–50× regime and compare it against CS and deep learning approaches.

2 Method

2.1 Regularizer architecture

Following the multi-subspace strategy in [15], the proposed regularizer replaces costly 4D spatio-temporal regularization with four branches operating in complementary 3D subspaces. Specifically, let $\mathbf{x} \in \mathbb{C}^{N_t \times N_x \times N_y \times N_z}$ be a 4D image. We order its axes as t, x, y, z and write $\prec$ for this order. For an axis $a \in \mathcal{A} := \{t, x, y, z\}$ and a slice index $b \in \{1, \dots, N_a\}$, let $a_1 \prec a_2 \prec a_3$ be the elements of $\mathcal{A} \setminus \{a\}$ and denote by $P_{a,b} \colon \mathbb{C}^{N_t \times N_x \times N_y \times N_z} \to \mathbb{C}^{N_{a_1} \times N_{a_2} \times N_{a_3}}$ the linear extraction operator that selects the b-th 3D hypersurface orthogonal to a. Then, we define the 4D regularizer

$$\mathcal{R}^{\mathrm{4D}}(\mathbf{x}) = \frac{\lambda}{4} \sum_{a \in \mathcal{A}} \sum_{b=1}^{N_a} \mathcal{R}^{\mathrm{3D}}_{a^\perp}\big(P_{a,b}\mathbf{x}\big), \quad \lambda > 0, \tag{4}$$

$\mathcal{R}^{\mathrm{3D}}_{a^\perp}$ being the regularizer branch acting in the 3D subspace $a^\perp$ orthogonal to axis a, and whose parametrization is given below.

Regularizer branch parametrization. We parametrize the regularizer branch acting in the 3D subspace $a^\perp$, $a \in \mathcal{A}$, with the fields-of-experts formulation

$$\mathcal{R}^{\mathrm{3D}}_{a^\perp}(\mathbf{z}) = \sum_{c=1}^{C} \sum_{n=1}^{N_{a_1}} \sum_{n'=1}^{N_{a_2}} \sum_{n''=1}^{N_{a_3}} \psi_{a,c}\big(|(W_a\mathbf{z})_{c,n,n',n''}|_{\varepsilon_0}\big), \tag{5}$$

where $W_a = [W_{a,1}, \dots, W_{a,C}]^\top$ is a learnable complex convolutional operator with C output channels applied to $\mathbf{z} \in \mathbb{C}^{N_{a_1} \times N_{a_2} \times N_{a_3}}$, $|\cdot|_{\varepsilon_0} := \sqrt{|\cdot|^2 + \varepsilon_0^2}$ with $\varepsilon_0 = 10^{-6}$ is applied component-wise to the complex filter responses, and the potential $\psi_{a,c}$ acts on the magnitude response of channel c. Following [21, 19], we build the potentials $\psi_{a,c}$ from the smoothed ℓ_1 family $s_\beta = \frac{\beta}{2}[|\cdot|^2 - (|\cdot| - \beta^{-1})_+^2]$ with $\beta > 0$: For each axis a, we use the 1-weakly convex base potential $\phi_{\beta_a} = s_{\beta_a} - s_1$, rescaled per output channel as $\psi_{a,c} = \alpha_{a,c}^{-2}\phi_{\beta_a}(\alpha_{a,c}\,\cdot)$, where $\beta_a > 1$ and $\alpha_{a,c} > 0$ are learnable.

Remark 1. The channel-wise rescaling $\psi_{a,c}$ preserves the 1-weak-convexity of the base potential. Hence, with spectrally normalized convolutional operators ($\|W_a\| = 1$, $a \in \mathcal{A}$), each $\mathcal{R}^{\mathrm{3D}}_{a^\perp}$ is 1-weakly convex, making the full regularizer $\mathcal{R}^{\mathrm{4D}}$ in (4) λ-weakly convex. The learnable regularizer is therefore allowed to be nonconvex and thus more powerful than purely convex penalties, while its negative curvature remains bounded via the learnable parameter λ, potentially also leading to more stable minimization of the objective in (3).

Remark 2. The smooth magnitude and smooth potentials make each regularizer branch $\mathcal{R}^{3\mathrm{D}}_{a^\perp}$ differentiable, with analytical gradient $\nabla\mathcal{R}^{3\mathrm{D}}_{a^\perp}(\mathbf{z}) = W_a^*(\psi_{a,c} \circ |\cdot|_{\varepsilon_0})'(W_a\mathbf{z})$, where the derivative acts component-wise on the complex response. Hence, $\nabla\mathcal{R}^{4\mathrm{D}}(\mathbf{x}) = \frac{\lambda}{4}\sum_{a\in\mathcal{A}}\sum_{b=1}^{N_a} P^*_{a,b}\nabla\mathcal{R}^{3\mathrm{D}}_{a^\perp}(P_{a,b}\mathbf{x})$, enabling gradient-based optimization of the objective in (3).

In practice, each convolutional operator W_a is normalized as required by Remark 1 and implemented as a cascade of three bias-free, unit-stride, unit-padding complex convolutions with $3\times3\times3$ kernels, giving an effective receptive field of $7\times7\times7$. The cascade uses 2, 4, and $C = 16$ output channels, respectively. The first layer of the cascade is constrained to have zero-mean filters, so that the regularizer responds to local variations rather than global intensity offsets, in the spirit of classical sparsifying filters such as finite differences and wavelets.

Acceleration-conditioned regularizer. We condition the regularizer $\mathcal{R}^{4\mathrm{D}}$ on the acceleration factor R by using the modified potential scales

$$\alpha_{a,c}(R) = R\exp\bigl(s_{c_{a,c}}(1/R)\bigr) \text{ for } R \in [R_{\min}, R_{\max}], \tag{6}$$

where $s_{c_{a,c}}$ is a linear spline with K equidistant knots on $[1/R_{\max}, 1/R_{\min}]$ and learnable parameters $c_{a,c}$. This enables $\mathcal{R}^{4\mathrm{D}}$ to handle a wide continuous range of undersampling rates and its effectiveness is demonstrated in Section 3. In practice, we used $R_{\min} = 9.0$, $R_{\max} = 51.0$ and $K = 5$.

2.2 Training and inference

Our goal is to learn a configuration of $\mathcal{R}^{4\mathrm{D}}$ that yields high-quality reconstructions through the variational formulation (3) across the entire range of R. We perform the underlying minimization with the *non-monotone accelerated proximal gradient (nmAPG)* algorithm [24, Supp. Thm. 4], guaranteeing convergence to a critical point. The algorithm is initialized with the zero-filled reconstruction and terminated once the relative change between iterates falls below 10^{-2}.

We train this reconstruction model end-to-end using a supervised ℓ_2 loss. This gives rise to a bilevel optimization problem: the reconstruction itself constitutes the lower-level problem (3), while the upper-level problem adjusts the configuration of $\mathcal{R}^{4D}$ to minimize the reconstruction loss. A central computational challenge is differentiating through the solution of (3). To make this backward pass tractable, we employ Jacobian-free backpropagation [23, 22]. We refer to [21] for details on the bilevel learning of (generic) regularization functionals.

At inference, given a four-point velocity encoded accelerated 4D Flow CMR acquisition, the reference image $\mathbf{x}_0$ and the velocity-encoded images $\mathbf{x}_1$, $\mathbf{x}_2$ and $\mathbf{x}_3$ are all reconstructed independently. Then, the velocity field is recovered as $\mathbf{v}_i = \frac{v_{\mathrm{enc}}}{\pi}\arg(\mathbf{x}_i\overline{\mathbf{x}_0})$, $i = 1, 2, 3$.

Remark 3. We refer to the proposed method as CLEAR: a **C**omplex, **L**earned, **E**xplicit, **A**nalytically differentiable **R**egularizer. The name also reflects the method's central idea: an interpretable learned regularizer (see Section 3.3) is inserted into an explicit/clear variational objective to be minimized.

3 Experiments and Results

3.1 Dataset and experimental setup

We use the CMRx4DFlow2026 challenge data from the CMRx reconstruction series, whose previous datasets and challenge summaries are described in [25–30]. The full collection contains over 400 multi-center, multi-vendor Cartesian 4D Flow CMR cases, including more than 300 aortic cases, acquired from more than 10 centers, 4 vendors (GE, Philips, Siemens, United Imaging), 3 field strengths (1.5T, 3T, 5T), and over 6 anatomical regions. The reported acquisition range is $200 \times 200 \times 40$ to $450 \times 450 \times 90$ mm^3 FOV, 0.9–3.0 mm spatial resolution, 23–120 ms temporal resolution, and $v_{\text{enc}} = 40$–100 cm/s for liver/kidney/brain and 100–200 cm/s for aorta/carotid. Each case provides $N_e = 4$ k-spaces of size $N_t \times N_c \times N_x \times N_y \times N_z$, corresponding to its four image velocity encodings, alongside coil sensitivity maps, segmentation masks, and v_{enc}. The dimensions vary between cases. Retrospective undersampling uses Gaussian k-t masks of size $N_t \times 1 \times N_x \times N_y \times 1$ at $R \in [10, 50]$, per case encoding. In our experiments, for each four-point velocity encoded case $\{\mathbf{x}_i\}_{i=0}^{3}$, the $R_i\times$-accelerated k-space $\mathbf{y}_i$ of encoding i is normalized as $\mathbf{y}_i \leftarrow (\|\mathbf{M}_{R_i}\|_2 / \|\mathbf{y}_i\|_2)\mathbf{y}_i$.

Training details. We train on 137 out of the 138 fully sampled aortic cases from the official training split and keep one for qualitative assessment. This is done on encoding patches of size $P_t \times N_c \times N_x \times N_y \times P_z$, with $P_t = 15$, $P_z = 7$ and a batch size of 5. In a given batch, each element consists of a randomly selected encoding from the same four-point velocity-encoded case and is retrospectively undersampled at a random $R \in [10, 50]$. We perform $50,000$ steps of the Adam optimizer with an initial learning rate of 5×10^{-3} decaying to 1×10^{-4} according to a cosine annealing schedule. On an NVIDIA RTX PRO 6000 Blackwell, the training takes around 96 hours.

Evaluation setup. We evaluate on all official validation sets, in which the cases are labeled as 10, 20, 30, 40 and 50×-accelerated. **Direct R1** evaluates the challenge's Regular Task 1 on 32 undersampled aortic cases, targeting accurate aorta reconstruction; **Generalizability S1** evaluates on the challenge's Special Task 1 on 40 cases from unseen sites and diseases; **Generalizability S2** evaluates the challenge's Special Task 2 on 40 unseen-anatomy cases: 10 cerebrovascular, 10 portal vein, 10 renal artery, and 10 carotid cases. Since validation ground truth is withheld, we report the official leaderboard metrics computed against hidden references within the segmentation mask: RelErr for vector-field magnitude error, AngErr for mean angular vector-field error, and nRMSE and SSIM for image magnitude, with SSIM averaged over the mask. We compare against LLR, representing compressed sensing, and FlowVN, representing deep learning. Both implementations are from the official challenge GitHub repository[1], and FlowVN uses the organizer-provided trained weights. The LLR-based reconstruction is computed using FISTA [31]. Comparisons against state-of-the-art methods are part of the official challenge leader-board and its final report.

[1] https://github.com/CmrxRecon/CMRx4DFlow2026

Table 1. Quantitative performance, computational complexity, and acceleration-conditioning ablation. **(a)** Leaderboard results on Direct R1 and the generalizability sets S1 and S2. **(b)** Model complexity and reconstruction time for a four-point-encoded, 30×-accelerated, 10-coil k-space case on a $28 \times 84 \times 108 \times 18$ spatio-temporal grid, measured on an NVIDIA RTX PRO 6000 Blackwell GPU. **(c)** Comparison of acceleration-specific CLEAR models (at $R = 10, 20, 30, 40, 50$) and the acceleration-conditioned CLEAR model on Direct R1. The best value in each comparison is shown in bold.

(a) Leaderboard metrics

Method	Direct R1				Generalizability S1				Generalizability S2			
	RelErr↓	AngErr↓	nRMSE↓	SSIM↑	RelErr↓	AngErr↓	nRMSE↓	SSIM↑	RelErr↓	AngErr↓	nRMSE↓	SSIM↑
LLR	0.583	39.815	0.127	0.799	0.522	34.764	0.128	0.857	0.465	31.951	0.102	0.863
FlowVN	0.526	41.562	0.059	0.927	0.626	40.599	0.110	0.909	0.803	44.044	0.120	0.854
CLEAR (ours)	**0.396**	**30.214**	**0.042**	**0.959**	**0.354**	**23.211**	**0.046**	**0.983**	**0.293**	**24.800**	**0.033**	**0.989**

(b) Complexity and reconstruction time

Metric	LLR	FlowVN	CLEAR (ours)
Recon. Time ↓	69 s	**4 s**	64 s
# of param. ↓	–	64,860	**8,317**

(c) Acceleration-conditioning ablation

CLEAR variant	RelErr↓	AngErr↓	nRMSE↓	SSIM↑	Train.Time↓
Acceleration-specific	0.403	31.018	0.043	0.958	146 h
Acceleration-conditioned	**0.396**	**30.214**	**0.042**	**0.959**	**96 h**

3.2 Quantitative and qualitative results

Table 1 shows that CLEAR achieves the best performance across all validation settings and leaderboard metrics. On Direct R1, it improves both velocity accuracy and magnitude reconstruction, reducing RelErr, AngErr, and nRMSE while increasing SSIM compared with LLR and FlowVN. The gains are even larger on the generalizability sets S1 and S2, suggesting that the learned regularizer transfers well to unseen sites, diseases, and anatomical regions. CLEAR is also parameter-efficient, using about eight times fewer parameters than FlowVN (8,317 vs. 64,860), while reconstruction time is comparable to that of LLR, though much slower than FlowVN. Finally, the ablation shows that acceleration-conditioned CLEAR slightly outperforms the acceleration-specific trained with the same data, architecture, optimizer, and number of steps. Although the gain is small, conditioning is useful because the effective undersampling factors are not exactly the nominal values $R = 10, 20, 30, 40, 50$. For instance, two cases labeled as 30× may correspond to true accelerations of 30.22× and 30.07×. The continuous parametrization in (6) adapts the potential scales to each such case within a single model, whereas acceleration-specific models use one fixed set of scales per nominal factor and require a separate training run for each.

To assess sensitivity to the stopping criterion of nmAPG, we also monitored reconstruction quality and objective decrease for different tolerances, as shown in Figure 1. Relaxing the tolerance from 10^{-1} to 10^{-2} improved the reconstruction (RelErr 0.642 → 0.295, nRMSE 0.239 → 0.073), whereas tightening it to 10^{-3} gave only marginal gains (RelErr 0.289, nRMSE 0.065) while increasing runtime from 58 s to 150 s. We therefore used an nmAPG tolerance of 10^{-2} for all leaderboard evaluations, as it provided a favorable accuracy–runtime compromise. The stable objective decrease further showcases optimization stability.

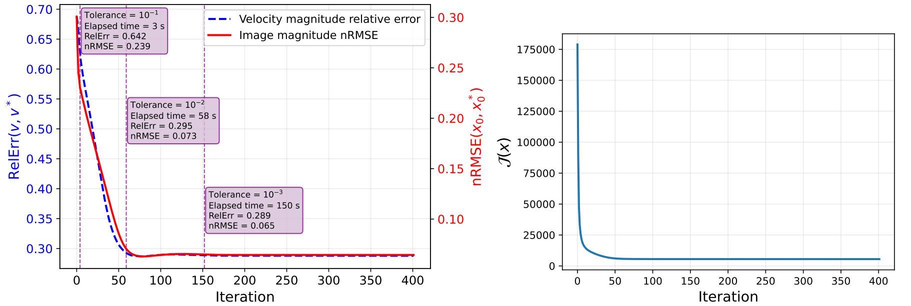

(a) **(b)**

Fig. 1. Convergence and stopping-tolerance sensitivity of CLEAR. (a) RelErr and nRMSE during batched reconstruction of the four velocity-encoded images in a 4D Flow CMR case, with markers indicating different nmAPG stopping tolerances. (b) Decrease of the objective (3) during reconstruction of a single velocity-encoded image.

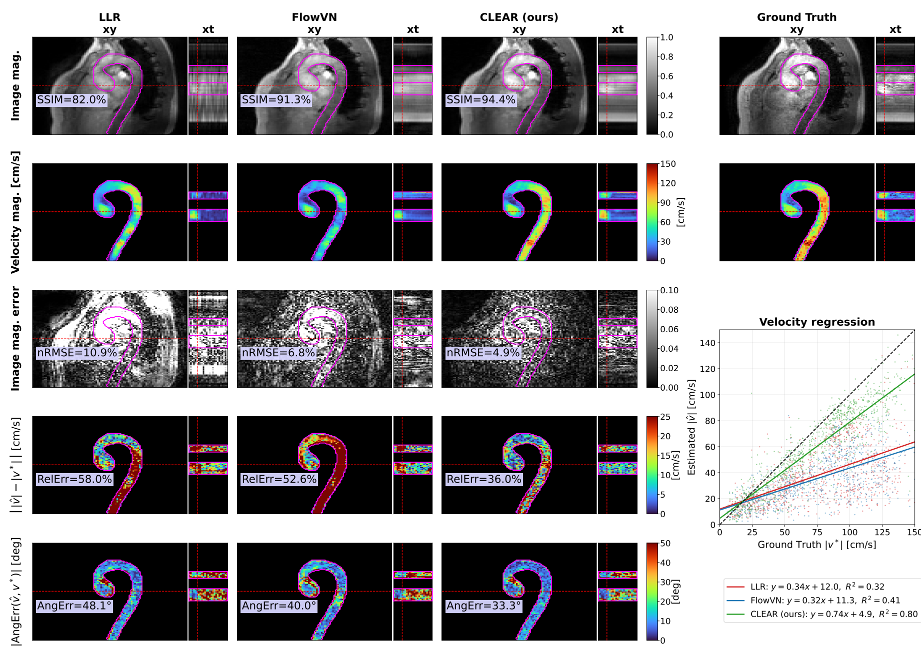

Fig. 2. Qualitative results on a 30× undersampled aortic case. Corresponding slice locations are illustrated with red dashed lines, indicating cross-section of the aorta and systolic peak. Scatter plot of velocity magnitude over manually segmented aorta (contour shown in magenta) is given together with correlation analysis ($y = ax + b$).

Figure 2 shows qualitative results on a representative 30× undersampled aortic case. Compared with LLR and FlowVN, CLEAR better preserves the

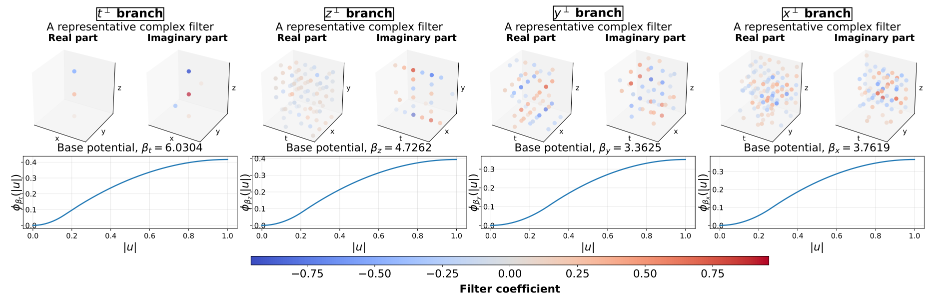

Fig. 3. Visualization of the learned components in CLEAR. Only filter coefficients with magnitudes larger than 15% of the maximum magnitude are displayed for clarity.

magnitude image and substantially reduces both magnitude and velocity errors within the segmented aorta. This is reflected by the improved image metrics (SSIM = 94.4%, nRMSE = 4.9%) and velocity metrics (RelErr = 36.0%, AngErr = 33.5°). The velocity scatter plot further shows that CLEAR provides the closest agreement with the reference velocities, with a slope closer to the identity and a markedly higher correlation ($R^2 = 0.80$) than LLR ($R^2 = 0.32$) and FlowVN ($R^2 = 0.41$). Although high velocities remain slightly underestimated, CLEAR produces the least biased and most spatially consistent reconstruction among the compared methods.

3.3 Interpretability

Figure 3 shows representative filters and potentials from the four 3D regularizer branches (5) in CLEAR. The $t^\perp$ branch learns a spatial high-pass filter, visually reminiscent of finite differences or local contrast detectors, whereas the $z^\perp$, $y^\perp$, and $x^\perp$ branches learn broader oscillatory spatio-temporal filters resembling higher-order derivatives and wavelet-like patterns. The learned potentials act on the magnitudes of the complex filter responses. Their smooth, saturating shape penalizes small responses strongly while treating large responses more mildly, encouraging incoherent features to vanish while preserving strong anatomical and flow-related structures. Thus, the filters indicate which local features are measured, and the potentials indicate how sparsity is imposed on them.

4 Discussion and Conclusion

We introduced CLEAR, a complex, learned, explicit regularizer for 4D Flow CMR reconstruction. The method combines multi-subspace 3D complex filters, smooth sparsity-promoting potentials, and acceleration conditioning within a variational formulation. Across the CMRx4DFlow2026 validation settings, CLEAR consistently improved magnitude and velocity metrics over the baselines with less than 10k parameters. The results suggest that learned regularizers can bridge part of the gap between classical compressed sensing and deep reconstruction networks. This is particularly relevant for clinical reconstruction, where robust and interpretable reconstruction methods are desirable.

The current design, however, has two limitations. First, the regularizer uses a smooth magnitude $|\cdot|_{\varepsilon_0}$ to map each complex filter response to a real scalar before applying the potential. This radial construction is simple and smooth but invariant to the phase of the complex response and thus may discard useful phase-dependent information in the learned feature space. Future work could therefore investigate more expressive and phase-aware mappings. Second, the regularizer processes each image encoding within a four-point velocity-encoded 4D Flow CMR case independently, which leads to independent reconstructions of each image encoding. A joint processing of the four encodings within a case, as done in [16], would exploit the correlations between them and potentially lead to improved reconstructions. This is also left for future work, as well as prospective validation, uncertainty estimation, and clinical reader studies.

Code availability. The source code and trained weights for CLEAR are publicly available on GitHub: https://github.com/Shamachrist7/CLEAR.

Acknowledgments. This study was funded by the DFG within the SPP2298 under the Project Number 543939932.

References

1. Markl, M., Frydrychowicz, A., Kozerke, S., Hope, M., Wieben, O.: 4D flow MRI. J. Magn. Reson. Imaging 36(5), 1015–1036 (2012). https://doi.org/10.1002/jmri.23632
2. Dyverfeldt, P., Bissell, M., Barker, A.J., et al.: 4D flow cardiovascular magnetic resonance consensus statement. J. Cardiovasc. Magn. Reson. 17, 72 (2015). https://doi.org/10.1186/s12968-015-0174-5
3. Bissell, M.M., Raimondi, F., Ait Ali, L., et al.: 4D Flow cardiovascular magnetic resonance consensus statement: 2023 update. Journal of Cardiovascular Magnetic Resonance 25, 40 (2023). https://doi.org/10.1186/s12968-023-00942-z
4. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. Physica D 60(1–4), 259–268 (1992). https://doi.org/10.1016/0167-2789(92)90242-F
5. Trzasko, J., Manduca, A.: Local versus global low-rank promotion in dynamic MRI series reconstruction. In: Proc. ISMRM, p. 4371 (2011)
6. Lyu, J., Li, G., Wang, C., et al.: Region-focused multi-view transformer GAN for cardiac cine MRI reconstruction. Med. Image Anal. 85, 102760 (2023). https://doi.org/10.1016/j.media.2023.102760
7. Lyu, J., Tian, Y., Cai, Q., et al.: Adaptive channel-modulated personalized federated learning for MRI reconstruction. Comput. Biol. Med. 165, 107330 (2023). https://doi.org/10.1016/j.compbiomed.2023.107330
8. Wang, Z., Qian, C., Guo, D., et al.: One-dimensional deep low-rank and sparse network for accelerated MRI. IEEE Trans. Med. Imaging 42, 79–90 (2023). https://doi.org/10.1109/TMI.2022.3203312

9. Qin, C., Schlemper, J., Caballero, J., et al.: Convolutional recurrent neural networks for dynamic MR image reconstruction. IEEE Trans. Med. Imaging 38, 280–290 (2019). https://doi.org/10.1109/TMI.2018.2863670
10. Lyu, J., Wang, S., Tian, Y., et al.: STADNet: spatial-temporal attention-guided dual-path network for cardiac cine MRI super-resolution. Med. Image Anal. 94, 103142 (2024). https://doi.org/10.1016/j.media.2024.103142
11. Wang, Z., Xiao, M., Zhou, Y., et al.: Deep separable spatiotemporal learning for fast dynamic cardiac MRI. IEEE Trans. Biomed. Eng. 72, 3642–3654 (2025). https://doi.org/10.1109/TBME.2025.3574090
12. Huang, J., Yang, L., Wang, F., et al.: Enhancing global sensitivity and uncertainty quantification in medical image reconstruction with Monte Carlo arbitrary-masked Mamba. Med. Image Anal. 99, 103334 (2025). https://doi.org/10.1016/j.media.2024.103334
13. Wang, Z., Yu, X., Wang, C., et al.: One for multiple: physics-informed synthetic data boosts generalizable deep learning for fast MRI reconstruction. Med. Image Anal. 103, 103616 (2025). https://doi.org/10.1016/j.media.2025.103616
14. Lyu, J., Wang, G., Wang, Z., et al.: Diffusion-prior implicit neural representation for arbitrary-scale cardiac cine MRI super-resolution. Inf. Fusion 126, 103510 (2026). https://doi.org/10.1016/j.inffus.2025.103510
15. Vishnevskiy, V., Walheim, J., Kozerke, S.: Deep variational network for rapid 4D flow MRI reconstruction. Nat. Mach. Intell. 2, 228–235 (2020). https://doi.org/10.1038/s42256-020-0165-6
16. Jacobs, L., Piccirelli, M., Vishnevskiy, V., Kozerke, S.: FlowMRI-Net: a generalizable self-supervised 4D flow MRI reconstruction network. J. Cardiovasc. Magn. Reson. 27, 101913 (2025). https://doi.org/10.1016/j.jocmr.2025.101913
17. Roth, S., Black, M.J.: Fields of experts. Int. J. Comput. Vis. 82(2), 205–229 (2009). https://doi.org/10.1007/s11263-008-0197-6
18. Goujon, A., Neumayer, S., Bohra, P., Ducotterd, S., Unser, M.: A neural-network-based convex regularizer for inverse problems. IEEE Trans. Comput. Imaging 9, 781–795 (2023). https://doi.org/10.1109/TCI.2023.3306100
19. Goujon, A., Neumayer, S., Unser, M.: Learning weakly convex regularizers for convergent image-reconstruction algorithms. SIAM J. Imaging Sci. 17(1), 91–115 (2024). https://doi.org/10.1137/23M1565243
20. Wache, G.S., Chaithya, G.R., Tanabene, A., Neumayer, S.: Weakly convex ridge regularization for 3D non-Cartesian MRI reconstruction. arXiv:2603.27158 (2026)
21. Hertrich, J., Wong, H.S., Denker, A., et al.: Learning regularization functionals for inverse problems: a comparative study. Handbook of Numerical Analysis, Elsevier (2026). https://doi.org/10.1016/bs.hna.2026.04.001
22. Bolte, J., Pauwels, E., Vaiter, S.: One-step differentiation of iterative algorithms. In: NeurIPS, vol. 36, pp. 77089–77103 (2023)
23. Fung, S.W., Heaton, H., Li, Q., McKenzie, D., Osher, S., Yin, W.: JFB: Jacobian-free backpropagation for implicit networks. Proc. AAAI 36(6), 6648–6656 (2022). https://doi.org/10.1609/aaai.v36i6.20619
24. Li, H., Lin, Z.: Accelerated proximal gradient methods for nonconvex programming. In: NeurIPS, vol. 28, pp. 379–387 (2015)
25. Wang, C., Lyu, J., Wang, S., et al.: CMRxRecon: a publicly available k-space dataset and benchmark for cardiac MRI. Sci. Data 11, 687 (2024). https://doi.org/10.1038/s41597-024-03525-4
26. Wang, Z., Wang, F., Qin, C., et al.: CMRxRecon2024: a multimodality, multiview k-space dataset for accelerated cardiac MRI. Radiol. Artif. Intell. 7, e240443 (2025). https://doi.org/10.1148/ryai.240443

27. Wang, Z., Huang, M., Shi, Z., et al.: Enabling ultra-fast cardiovascular imaging across heterogeneous clinical environments with a generalist foundation model and multimodal database. arXiv:2512.21652 (2025)
28. Lyu, J., Qin, C., Wang, S., et al.: State-of-the-art cardiac MRI reconstruction: results of the CMRxRecon challenge. Med. Image Anal. 101, 103485 (2025). https://doi.org/10.1016/j.media.2025.103485
29. Wang, K., Qin, C., Shi, Z., et al.: Extreme cardiac MRI analysis under respiratory motion: results of the CMRxMotion challenge. Med. Image Anal., 103883 (2025). https://doi.org/10.1016/j.media.2025.103883
30. Wang, F., Wang, Z., Li, Y., et al.: Toward modality- and sampling-universal learning strategies for accelerating cardiovascular imaging. IEEE Trans. Med. Imaging 45, 1872–1887 (2026). https://doi.org/10.1109/TMI.2025.3641610
31. Beck, A., Teboulle, M.: A fast iterative shrinkage-thresholding algorithm for linear inverse problems. SIAM Journal on Imaging Sciences 2(1), 183–202 (2009). https://doi.org/10.1137/080716542
32. Wang, C., Li, Y., Lv, J., et al.: Recommendation for CMR imaging-based phenotypic study: imaging part. Phenomics 1, 151–170 (2021). https://doi.org/10.1007/s43657-021-00018-x