Comments for k2ub.

We acknowledge the contribution of Huang et al. and will cite their work explicitly, revising the abstract and introduction to better situate our paper with respect to prior work.

Discussion on Novelty. The two works share weight reuse, but *operate in different regimes and target different contribution* axes. (Implicit)^2 defines its representation as the fixed point of an input-injected layer:

$$z^\star = F_\theta(z^\star, x), \quad F_\theta(z, x) = \sin\big(W(z + \sin(Vx)) + Ux + b\big), \tag{1}$$

solved under equilibrium-oriented $\rho(\partial F/\partial z) < 1$ as a *stabilization* mechanism toward $z^\star$. Our model instead is a ***finite, non-contractive*** composition of a shared bias-free sinusoidal block, with no x-injection:

$$f_\psi(x) = W^{\text{out}} \circ \underbrace{(\sigma \circ W^{\text{rec}}) \circ \cdots \circ (\sigma \circ W^{\text{rec}})}_{R \text{ steps, } \textbf{no } x\text{-injection}} \circ\, \sigma \circ W^{\text{in}}(x), \tag{2}$$

whose objective is the integer-combination closure $\Omega^{(\ell+1)} \subseteq \text{span}_{\mathbb{Z}} \Omega^{(\ell)}$ (Eq. 9), making $|\Omega^{(R)}|$ **grow** with R. Tab. 3 confirms this empirically: iSIREN's spectrum saturates within one fixed-point iteration (33.5% across $\ell{=}1 \to \ell{=}8$, $\text{HF}_\ell{=}0.008$ throughout) while Ours grows monotonically 30.9%→94.6%. Three consequences: (a) no contractivity is assumed, hence **no convergence criterion is required (E.(ii))**; (b) bipolar code-space supervision and exact reconstruction (BER=0) are unreachable in the equilibrium formulation—iSIREN attains 47.16 dB vs. our ∞ (Tab. 3).

Table 1. **(Implicit)^2 vs. Ours comparison table**.

	(Implicit)^2	**Ours**
What is sought	$z^\star:\ z^\star{=}F_\theta(z^\star, x)$	$\Omega^{(R)}:\ \Omega^{(\ell+1)} \subseteq \text{span}_{\mathbb{Z}}\, \Omega^{(\ell)}$
Block	$\sin\big(W(z{+}\sin Vx){+}Ux{+}b\big)$	$\sin(\omega W^{\text{rec}} h)$, no x-inject.
Dynamical regime	fixed-point / equilibrium-oriented	**finite explicit unrolling**
Recurrence as	stabilization $\to z^\star$	**spectral enrichment** $\Omega^{(\ell)} \nearrow$
Iterations	$r{\to}\infty$ via solver	finite R (hyperparam.)
Backward	implicit diff. + Neumann	standard BPTT

Relation to Fourier-series views. We acknowledge an oversight and will cite Benbarka et al., (2022) explicitly. Concretely, we will revise Sec. 3.1 as:

> *This first-layer view is consistent with prior analyses connecting Fourier-feature mappings and SIREN-style layers (Benbarka et al., (2022)) . . .*

Evaluation on Non-sinusoidal models. Tab. 4 reports a matched-budget sweep across parameter budgets, with each architecture trained at its best LR identified in Tab. 2. Consistent with Tab. 3, *non-sinusoidal recurrence collapses at every budget*, confirming that sin-specific harmonic expansion (Tab. 3(c)) is the operative mechanism driving fidelity.

E.(i) refer the Tab. 2	**E.(ii)** no fixed-point solver convergence
E.(iii) refer the Tab. 3's PSNR col.	**E.(iv)** will specify in details

Table 2. **Learning rate sensitivity across architectures.**

Set5 (butterfly)	Rec.	lr=1e-3	lr=5e-4	lr=1.5e-4 (default)	lr=5e-5	best LR	PSNR @ best	Δ vs default
Non-Sinusoidal act.								
WIRE	✗	38.34	**40.13**	32.45	36.84	5e-4	40.13	+7.68
Gauss	✗	39.60	36.58	**60.31**	39.70	1.5e-4	60.31	0.00
PEMLP	✗	**34.20**	31.87	27.21	21.61	1e-3	34.20	+6.99
+ Rec.	✓	**31.04**	30.51	22.88	18.72	1e-3	31.04	+8.15
Sinusoidal act.								
SIREN	✗	**48.64**	42.25	35.97	32.49	1e-3	48.64	+12.67
+ Rec.	✓	5.82	15.15	**78.35**	50.96	1.5e-4	78.35	0.00
+ Bin. *(Ours)*	✓	8.53	30.34	77.55	$\infty_{\text{BER=0}}$	5e-5	$\infty_{\text{BER=0}}$	+N/A
FINER	✗	12.22	**48.36**	42.32	35.71	5e-4	48.36	+6.03
+ Rec.	✓	12.17	11.41	**85.62**	53.00	1.5e-4	85.62	0.00

Comments for tdVq.

Validation of the Jacobi–Anger expansion. On a regular $H{\times}H$ grid, we forward each architecture, capture $\mathbf{h}^{(\ell)}{=}\sigma(\omega \mathbf{W}^{(\text{rec})}\mathbf{h}^{(\ell-1)}) \in \mathbb{R}^{H^2 \times C}$, and take the per-channel 2-D DFT $\mathbf{F}_c^{(\ell)}{=}\text{DFT}(\mathbf{h}_{:,c}^{(\ell)})$, where $|\mathbf{F}_c^{(\ell)}[u, v]|$ is the strength of channel c at frequency (u, v). Channel-averaging and peak-normalising in dB:

$$M_\ell^{\text{dB}}[u, v] = 20 \log_{10} \frac{(1/C)\sum_c |\mathbf{F}_c^{(\ell)}[u, v]|}{\max_{u', v'}(\cdot)}, \tag{3}$$

we report two scalars per depth/recurrent—spectral support (breadth; S_ℓ^τ) and radial HF content (strength on the upper Nyquist half $\mathcal{U}{=}\{H/4, \ldots, H/2{-}1\}$; HF_ℓ):

$$S_\ell^\tau{=}H^{-2}\big|\{(u, v) : M_\ell^{\text{dB}}{>}\tau\}\big|, \quad \text{HF}_\ell{=}\tfrac{1}{|\mathcal{U}|}\sum_{k \in \mathcal{U}} \tfrac{R_\ell(k)}{R_\ell(0)}, \tag{4}$$

with $R_\ell(k)$ the radial profile of M_ℓ at radius k from DC. At random init, sinusoidal recurrence shows monotone growth in S_ℓ^τ (Tab. 3(a)) while non-sinusoidal recurrence collapses (Tab. 3(c)) and iSIREN saturates within one fixed-point iteration, separating sinusoidal growth from recurrent collapse.

Table 3. **Spectral support S_ℓ^τ and radial HF_ℓ.**

Architecture	Reccurent	Mode	S_ℓ^τ per depth/recurrent ℓ (%) 0	1	2	3	4	5	6	7	8	HF_ℓ at $\ell{=}0, 8$ 0	8	PSNR *for ref.*
(a) Sinusoidal recurrent, untrained														
SIREN	✓	$\omega{=}256/45$	30.9	26.5	34.3	43.7	52.0	61.3	72.3	85.0	**94.6**	0.022	0.101	—
FINER	✓	$\omega{=}256/45$	51.4	53.6	72.4	89.3	**100**	100	100	100	100	0.048	0.232	—
(b) Sinusoidal recurrent, trained on Set 5 (Bevilacqua et al., BMVC 2012)														
SIREN	✓	Float	31.0	17.1	17.2	19.1	20.5	12.9	16.4	28.1	61.7	0.020	0.011	78.35
+ Bin. *(Ours)*	✓	bipolar code	31.1	42.9	72.9	**99.8**	100	100	100	100	100	0.034	**0.160**	$\infty_{\text{BER=0}}$
FINER	✓	Float	53.9	41.4	46.6	46.0	37.4	16.4	41.6	100	100	0.044	0.016	85.62
iSIREN (Huang et al., 2021)	✓	DEQ	32.4	33.5	33.5	33.5	33.5	33.5	33.5	33.5	33.5	0.008	0.008	47.16
► *Sinusoidal-based Spectral methods — multiplicative (MFN) and bandlimited (BACON)*														
GaborNet (Fathony et al., 2021)	✗	$L{=}3$	40.6	58.2	73.0	35.2	—	—	—	—	—	0.038	0.008†	44.48
FourierNet (Fathony et al., 2021)	✗	$L{=}3$	41.4	58.3	66.6	36.0	—	—	—	—	—	0.050	0.008†	38.38
BACON‡ (Lindell et al., 2022)	✗	freq=128	0.9	6.3	14.0	15.1	—	—	—	—	—	0.059	3e-4†	27.88
*(c) Non-sinusoidal recurrent, untrained — **spectrum collapses***														
Gauss	✓	$\sigma{=}30$	1.0	1.9	14.0	1.2	**0.0**	0.0	0.0	0.0	0.0	2e-4	7e-3	12.02
PEMLP	✓	$N_f{=}10$	0.8	0.9	0.8	0.5	0.1	0.1	**0.0**	0.0	0.0	6e-4	2e-5	24.09

† HF reported at the architecture's final depth $\ell{=}3$ (followed official implementation). ‡ Single-scale BACON.

About overparameterization. Tab. 4 sweeps ∼100K–600K budgets across architectures. The gains are strongest in the moderate-to-high budget regime and align with the spectral diagnostic: Representatively, SIREN+Rec gains **+3.7 dB** at 100K. Consistent with Tab. 3(c), recurrence degrades non-sinusoidal models under spectral collapse, whereas sinusoidal recurrence improves fidelity.

> ***Minor:*** *We will also fix the requested clarity issues. Fig. 3 wording, Fig. 1 reference, bias-free equations, GT visuals, BER spelling, Table 7 bolding, Gray-code citation, output-quantization wording, and Fig. 7 caption.*

Table 4. **Effectiveness of Recurrent on Capacity-limited regime**

Architecture (best LR)	Rec.	PSNR (dB) ↑ ∼100K	∼200K	∼400K	∼600K	SSIM ↑ ∼100K	∼200K	∼400K	∼600K	LPIPS ↓ ∼100K	∼200K	∼400K	∼600K
(a) Non-sinusoidal baselines													
PEMLP (1e-3)	✗	**28.89**	**31.04**	31.80	**34.20**	**0.879**	**0.911**	0.920	**0.943**	**0.054**	**0.034**	0.025	**0.014**
+ Rec. (1e-3)	✓	27.48	29.64	**33.82**	31.04	0.847	0.886	**0.943**	0.903	0.072	0.049	**0.015**	0.028
Gauss (1.5e-4)	✗	**20.48**	**26.49**	**35.33**	**60.31**	**0.597**	**0.770**	**0.937**	**0.9997**	**0.290**	**0.066**	**0.007**	**0.0000**
+ Rec. (1.5e-4)	✓	11.74	11.77	11.92	12.02	0.074	0.069	0.098	0.109	1.282	1.268	1.247	1.187
(b) Sinusoidal variants (Ours)													
SIREN (1e-3)	✗	33.60	39.57	46.23	48.64	0.958	0.985	0.995	0.997	0.015	0.0010	0.0001	0.0000
+ Rec. (1.5e-4)	✓	**37.26**	**46.26**	**61.36**	78.35	**0.960**	**0.994**	**0.9998**	1.0000	**0.001**	**0.0001**	**0.0000**	0.0000
+ Bin. *(Ours)* (5e-5)	✓	23.26	24.83	38.97	$\infty_{\text{BER=0}}$	0.808	0.877	0.995	**1**	0.080	0.039	0.001	**0**
+ Weighted Bin. *(Ours$_W$)* (1.5e-4)	✓	32.29	40.01	53.91	73.38	0.941	0.983	0.999	1.0000	0.0074	0.0009	0.0000	0.0000
FINER (5e-4)	✗	**38.01**	40.59	45.56	48.36	0.979	0.986	0.994	0.996	0.004	0.002	0.0002	0.0001
+ Rec (1.5e-4)	✓	34.21	**44.45**	**68.06**	**85.62**	0.926	**0.991**	**1.0000**	**1.0000**	0.005	**0.0001**	**0.0000**	**0.0000**

Comments for Vw2E.

Comparison to domain SOTA & practicality. For NeRF with iNGP, we did not observe a meaningful gain, likely because the HashGrid provides a strong inductive bias. *In contrast*, *Ours* is competitive with LIIF in image SR, and also improves NeRV-style video INR by replacing only the stem at nearly identical model size (Tab. 5). This low-parameter advantage is relevant when many instance-specific decoders must be stored on-device.

Table 5. **Comparison across SR (LIIF) and Video (NeRV).**

Dataset	Scale	LIIF MLP (≈ 50K) PSNR↑	SSIM↑	LPIPS↓	*Ours* (≈ 50K) PSNR↑	SSIM↑	LPIPS↓
Benchmark (Model used: EDSR-baseline)							
Set5	×2	37.673	0.9485	0.0625	**37.721**	**0.9489**	**0.0574**
	×3	34.036	0.9104	0.1292	**34.103**	**0.9111**	**0.1275**
	×4	31.820	0.8711	**0.1797**	**31.850**	**0.8718**	0.1867
Set14	×2	33.325	0.9005	0.1022	**33.401**	**0.9016**	**0.0941**
	×3	30.116	0.8243	0.2121	**30.126**	**0.8249**	**0.2103**
	×4	28.377	0.7582	**0.2912**	**28.389**	**0.7598**	0.3010
B100	×2	31.968	0.8979	0.1624	**32.031**	**0.8996**	**0.1524**
	×3	28.926	0.8046	**0.2966**	**28.976**	**0.8060**	0.2969
	×4	**27.427**	**0.7294**	**0.3856**	**27.450**	**0.7308**	**0.3981**
Mean							
		31.519	0.8494	**0.2024**	**31.561**	**0.8505**	0.2027

Variant	Best config	PSNR ↑	SSIM ↑	LPIPS ↓
NeRV	H=128 lr=5e-4	38.414	0.9708	0.0170
+Rec	R=3 lr=5e-4	39.576	0.9768	0.0116
+Rec (sinusoidal)	R=5 ω=3 lr=1e-3 gc=1	**41.315**	**0.9828**	**0.0072**