Spik-NeRF v2: Pushing the Limit of Spiking Neural Radiance Fields with ±I-LIF

Chen Cheng* · Qinlong Lan* · Gang Wan · Hao Guo† · Lei Liu · Zhanji Wei · Yitian Wu · Yufei Guo†

Taiyuan University of Technology · Space Engineering University
Intelligent Science & Technology Academy of CASIC
* Equal contribution † Corresponding authors

Motivation

- ANN-based NeRF delivers high-quality novel views, yet dense multiply–accumulate computation limits deployment on power-constrained devices.
- SNN-based NeRF reduces energy through event-driven inference, but spike quantization loses activation information.
- Higher spike fidelity often requires more timesteps, creating a quality–latency trade-off.

Key Contributions

- $\pm$I $-$ LIF neuron. Signed integer spikes in $\{-D,\ldots,D\}$ retain positive and negative activation information during training.
- Spike-driven inference. Signed integers become ternary events $\{-1,0,1\}$, replacing MACs with sparse additions and subtractions.
- Equivalent re-parameterization. A trained $\mathrm{I-LIF^{2D}}$ model become $\pm\mathrm{I-LIF^{D}}$ with inherited weights and compensated bias.
- Half the inference timesteps while preserving the next-layer input and rendering quality.

Methodology

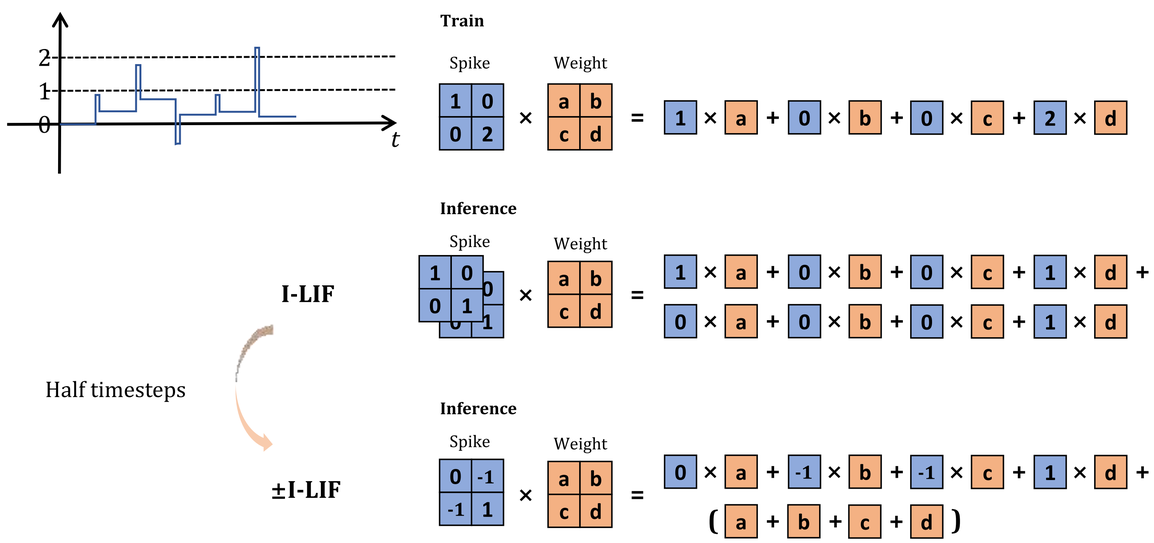


Illustration of the re-parameterization process for converting $\mathrm{I-LIF^{2D}}$ to $\pm\mathrm{I-LIF^{D}}$ during inference.

Key Results & Analysis

Metric	Method	Fern	Lego	Drums
PSNR↑	LIF (2 timesteps)	25.08	29.33	24.51
	LIF (4 timesteps)	25.59	30.35	25.07
	I-LIF^2	25.43	30.37	24.96
	±I-LIF^2	25.62	30.83	25.31
	I-LIF^4	25.82	31.11	25.35
SSIM↑	LIF (2 timesteps)	0.769	0.940	0.907
	LIF (4 timesteps)	0.795	0.951	0.918
	I-LIF^2	0.788	0.951	0.916
	±I-LIF^2	0.800	0.958	0.924
	I-LIF^4	0.809	0.960	0.925
LPIPS↓	LIF (2 timesteps)	0.207	0.039	0.086
	LIF (4 timesteps)	0.173	0.029	0.071
	I-LIF^2	0.185	0.029	0.071
	±I-LIF^2	0.160	0.023	0.061
	I-LIF^4	0.150	0.022	0.061

Per-scene quantitative comparison.

- I-LIF achieves slightly better performance yet doubles the inference latency.
- Under identical inference timesteps, ± I-LIF outperforms the original I-LIF neuron.
- Both ±I − LIF and I-LIF substantially surpass the vanilla LIF neuron.

Conclusion

- Better representation. Signed integer training captures activation polarity and reduces quantization loss.
- Efficient inference. Ternary events retain spike-driven, multiplication-free computation.
- Lower latency. Re-parameterization halves virtual timesteps without changing the next-layer input.
- Strong rendering. With only 2–4 timesteps, Spik-NeRF v2 approaches ANN-based NeRF.

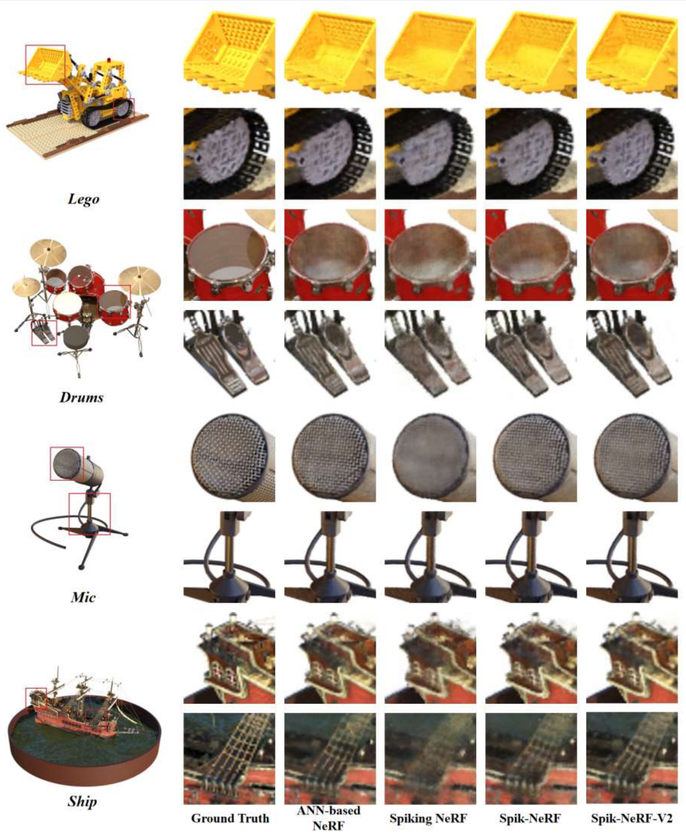
Rendering Results

Metric	Method	Chair	Drums	Ficus	Hotdog	Lego	Materials	Mic	Ship	Avg.
PSNR↑	ANN-based NeRF Mildenhall et al. (2021)	34.15	25.64	29.15	36.85	31.48	29.34	33.12	29.42	31.15
	Spiking-NeRF Li et al. (2025)	12.24	11.14	13.98	14.85	9.81	10.36	9.81	13.22	11.93
	Spiking NeRF	31.98	24.51	24.00	34.55	29.33	27.75	31.65	27.85	28.95
	Spik-NeRF Wan et al. (2025)	33.40	25.21	26.05	35.99	30.82	28.86	32.78	28.75	30.23
	Spik-NeRF v2	**33.71**	**25.35**	**27.60**	**36.40**	**31.11**	**29.07**	**33.04**	**29.12**	**30.68**
	Spik-NeRF v2*	**34.06**	**25.57**	**27.72**	**36.69**	**31.42**	**29.41**	**33.35**	**29.33**	**30.95**
SSIM↑	ANN-based NeRF Mildenhall et al. (2021)	0.979	0.929	0.966	0.979	0.965	0.958	0.978	0.874	0.953
	Spiking NeRF	0.963	0.907	0.891	0.969	0.940	0.937	0.969	0.838	0.927
	Spik-NeRF Wan et al. (2025)	0.973	0.921	0.929	0.978	0.957	0.950	0.975	0.856	0.942
	Spik-NeRF v2	**0.975**	**0.925**	**0.949**	**0.976**	**0.960**	**0.953**	**0.977**	**0.866**	**0.948**
	Spik-NeRF v2*	**0.977**	**0.928**	**0.950**	**0.979**	**0.964**	**0.957**	**0.979**	**0.872**	**0.951**
LPIPS↓	ANN-based NeRF Mildenhall et al. (2021)	0.014	0.053	0.022	0.015	0.020	0.024	0.023	0.086	0.032
	Spiking NeRF	0.034	0.086	0.119	0.032	0.039	0.047	0.043	0.126	0.066
	Spik-NeRF Wan et al. (2025)	0.021	0.065	0.063	0.021	0.027	0.033	0.028	0.105	0.045
	Spik-NeRF v2	**0.017**	**0.061**	**0.037**	**0.018**	**0.022**	**0.031**	**0.026**	**0.095**	**0.038**
	Spik-NeRF v2*	**0.016**	**0.056**	**0.036**	**0.016**	**0.020**	**0.026**	**0.023**	**0.088**	**0.035**

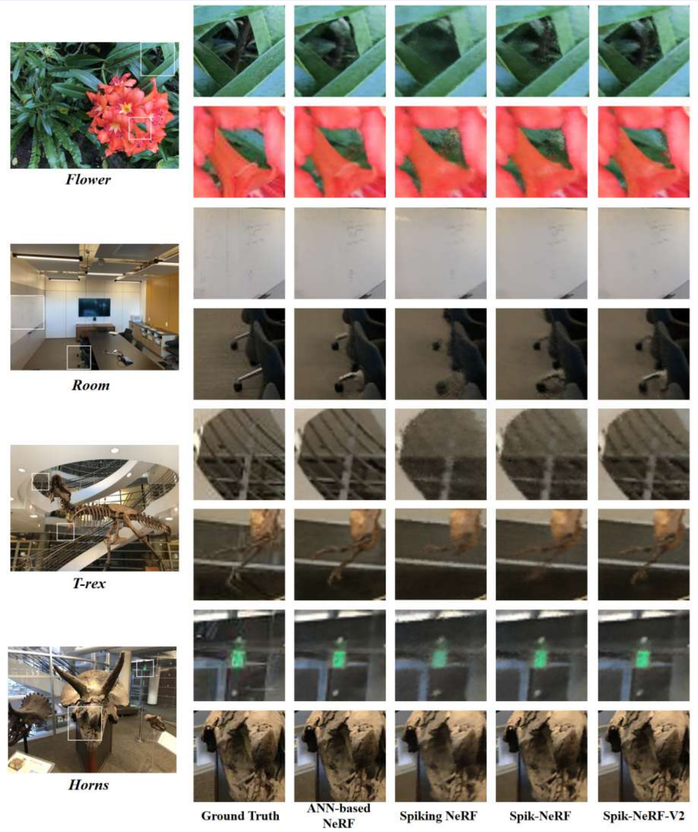
Per-scene quantitative results on the synthetic dataset

Metric	Method	Room	Fern	Leaves	Fortress	Orchids	Flower	T-Rex	Horns	Avg.
PSNR↑	ANN-based NeRF Mildenhall et al. (2021)	31.38	26.25	21.98	31.35	21.20	27.51	27.27	28.10	26.88
	Spiking-NeRF Li et al. (2025)	17.07	15.91	9.68	14.51	9.12	10.35	15.05	13.05	13.09
	Spiking NeRF	30.12	25.08	20.75	30.07	20.51	26.19	25.26	26.34	25.54
	Spik-NeRF Wan et al. (2025)	30.90	25.70	21.46	30.79	20.94	26.94	26.15	27.16	26.26
	Spik-NeRF v2	**31.07**	**25.82**	**21.57**	**30.91**	**21.05**	**27.10**	**26.43**	**27.37**	**26.42**
	Spik-NeRF v2*	**31.11**	**25.99**	**21.75**	**31.02**	**21.12**	**27.29**	**26.76**	**27.60**	**26.58**
SSIM↑	ANN-based NeRF Mildenhall et al. (2021)	0.931	0.836	0.790	0.896	0.734	0.853	0.896	0.877	0.851
	Spiking NeRF	0.904	0.769	0.691	0.837	0.644	0.782	0.825	0.800	0.781
	Spik-NeRF Wan et al. (2025)	0.920	0.802	0.744	0.869	0.688	0.820	0.862	0.835	0.817
	Spik-NeRF v2	**0.923**	**0.809**	**0.753**	**0.874**	**0.698**	**0.828**	**0.870**	**0.846**	**0.825**
	Spik-NeRF v2*	**0.925**	**0.819**	**0.770**	**0.883**	**0.714**	**0.839**	**0.881**	**0.857**	**0.836**
LPIPS↓	ANN-based NeRF Mildenhall et al. (2021)	0.049	0.101	0.119	0.059	0.122	0.075	0.062	0.078	0.083
	Spiking NeRF	0.098	0.207	0.203	0.130	0.223	0.134	0.135	0.166	0.162
	Spik-NeRF Wan et al. (2025)	0.066	0.164	0.157	0.092	0.170	0.101	0.095	0.123	0.121
	Spik-NeRF v2	**0.062**	**0.150**	**0.150**	**0.084**	**0.157**	**0.093**	**0.086**	**0.111**	**0.111**
	Spik-NeRF v2*	**0.058**	**0.133**	**0.135**	**0.073**	**0.141**	**0.083**	**0.076**	**0.098**	**0.099**

Per-scene quantitative results on the realistic dataset



The synthetic dataset



The realistic dataset