LotteryCodec: Searching the Implicit Representation in a Random Networks for Low-Complexity Image Compression

Haotian Wu, Gongpu Chen, Pier Luigi Dragotti, Deniz Gündüz
15/06/2025

Outline

1. Preliminary
2. Lottery Codec Hypothesis
3. Verification of the Hypothesis
4. LotteryCodec scheme
5. Experimental Results
6. Conclusion and Future Work

1. Preliminary

From data representation to function representation

Parameterize a discrete signal as a **continuous** function.

Use neural networks to approximate the mapping from coordinates to signal intensities [1].

Shift the paradigm of data representation from feature-based to **function-based** representation.

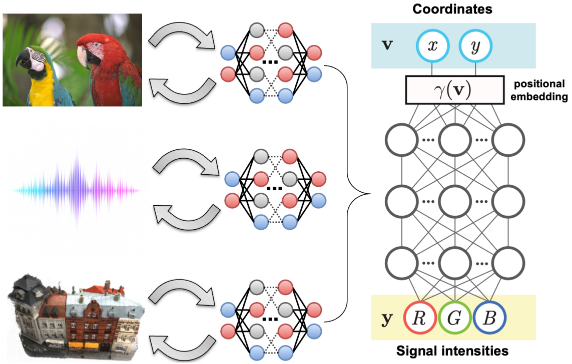

Fig. 1: Examples of INRs for various modalities.

1. Preliminary
From AE-based neural codec to overfitted codec

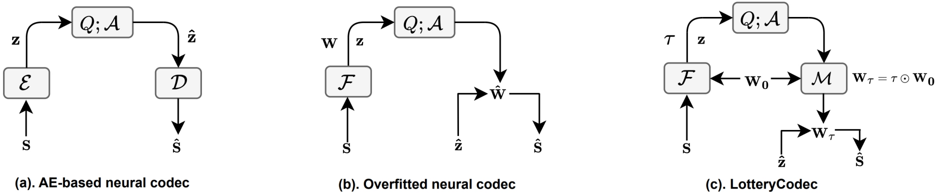

Fig. 2: Operational diagram of different compression models.

AE-based codecs leverage **advanced** architectures and **large-scale** datasets to achieve strong rate-distortion (RD) performance [2].

- An encoder-decoder pair maps the source to a quantized latent, which is entropy-coded to form the bitstream and decoded to reconstruct the signal.
- Developing **low-complexity**, **robust** codecs with **strong** RD performance remains an open challenge.

1. Preliminary
From AE-based neural codec to overfitted codec

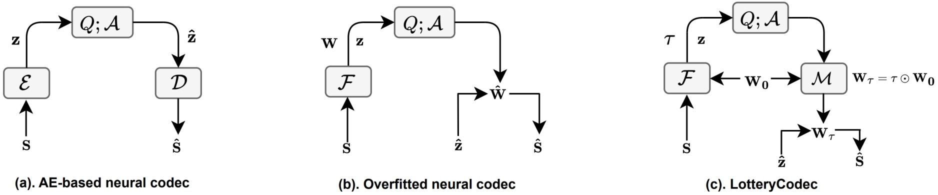

Fig. 2: Operational diagram of different compression models.

Overfitted codecs instead parameterize **each** data sample using **lightweight** neural functions, aiming for a **good, cheap, and fast** compression scheme [3], [4].

- Each source sample is represented by a neural function and an alternative latent vector.
- Decoding complexity is extremely low, as no data-generalization is required.
- Outperforming some existing codecs, such as BPG, HEVC, and BMS.

1. Preliminary

From overfitted codec to LotteryCodec

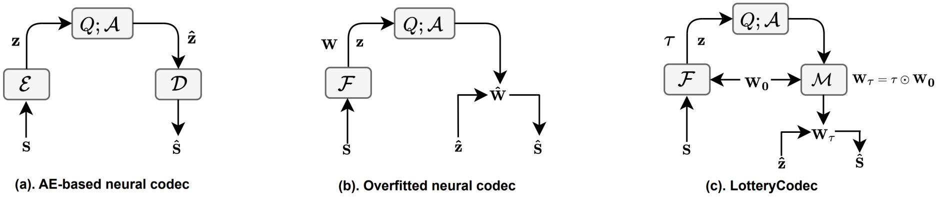

Fig. 2: Operational diagram of different compression models.

Motivations:

?

- Can each sample be parameterized under flexible complexity? $\rightarrow$**flexibility.**
- Can INRs encode the source into the network structure? $\rightarrow$ **lower rate.**
- Can we eliminate the neural function quantization for continuous parameters? $\rightarrow$ **better RD.**

Towards a **flexible** codec with **enhanced** RD performance $\rightarrow$ **LotteryCodec**.

1. Preliminary

Lottery ticket hypothesis

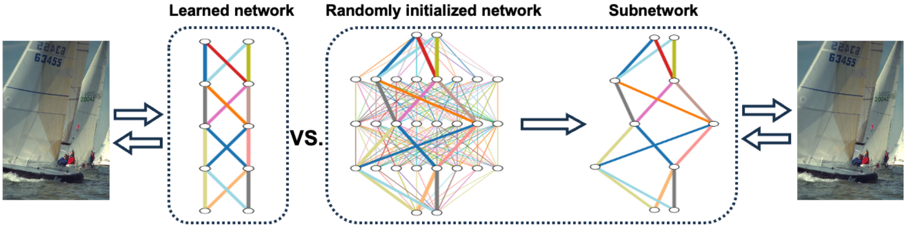

Fig. 3: Illustration of lottery ticket hypothesis.

- A randomly-initialized neural network can act as a handcrafted prior, encoding image statistics and prior information within its network structure [5].
- Over-parameterized neural networks contain high-performing untrained subnetworks [6].

2. Lottery Codec Hypothesis (LCH)

A New Paradigm for Overfitted Image Compression

Hypothesis Statement

Let $\mathsf{d}(\cdot)$ denote a distortion function and $\mathsf{H}(\cdot)$ the entropy function. For any overfitted image codec $g_{\mathbf{W}}(\mathbf{z})$, there exists an over-parameterized and randomly initialized network $g_{\mathbf{W}'}$ with $|\mathbf{W}'| > |\mathbf{W}|$ and a pair $(\tau', \mathbf{z}')$ as the 'winning tickets', such that

$$\mathsf{d}(\mathbf{S}, \mathbf{S}') \leq \mathsf{d}(\mathbf{S}, \mathbf{S}^*), \quad \mathsf{H}(\hat{\mathbf{z}}') = \mathsf{H}(\hat{\mathbf{z}})$$

Goal Achieve similar or better distortion under the same bit cost.

Empirical Support Extensive experiments were conducted to support its validity.

Theoretical Justification Any target network of width L_w and depth L_d can be approximated by pruning a random network that is a factor $O(\log(L_w L_d))$ wider and twice as deep [7], [8].

3. Verification of the hypothesis

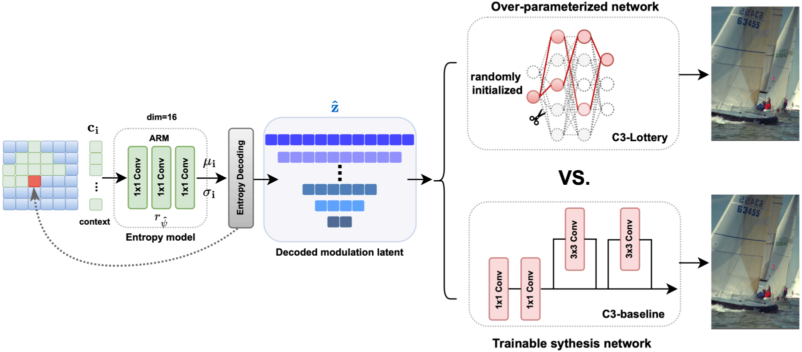

Fig. 4: Illustration of the experiment setup.

- The synthesis network is replaced with a randomly initialized over-parameterized network.
- Only a binary mask is learned, while all other components (from C3[9]) remain unchanged.
- We report the RD trade-off across varying network depths and widths.

3. Verification of the hypothesis

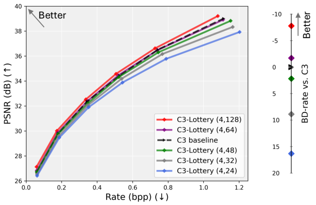

RD curve and BD rate for different configurations.

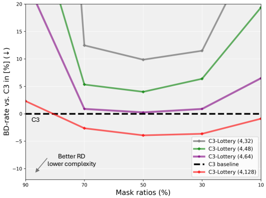

BD-rate versus mask ratios.

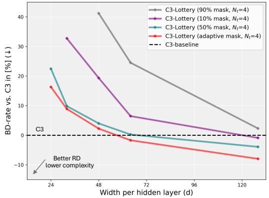

BD-rate over varying width and depth configurations.

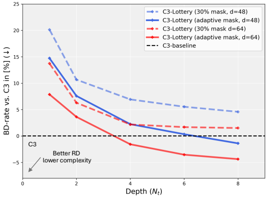

BD-rate over varying width and depth configurations.

Fig. 5: Experimental verification of the LCH across different over-parametrization configurations.

Experimental Observations

- When sufficiently over-parameterized, we can find a subnetwork and latent that match the original RD.
- RD performance (PSNR vs. the rate from $\hat{z}$) improves significantly with increasing width and depth.
- Perform best around 50% mask ratio, which maximizes structure entropy for coding richer information.

4. LotteryCodec scheme

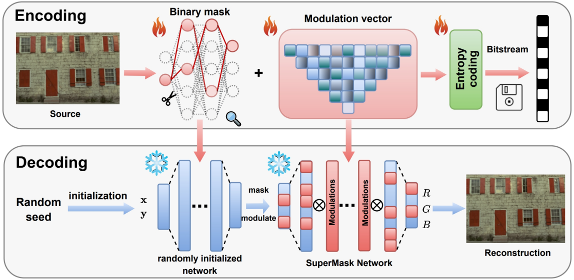

Fig. 6: Illustration of LotteryCodec.

- However, greater over-parameterization increases search complexity and bit cost.

We introduce a **rewind modulation mechanism** to enhance RD and simplify search.

- The source image is encoded into a binary mask and latent modulations.
- Preserved advantages: lower mask bit cost and flexible coding complexity

4. LotteryCodec Scheme

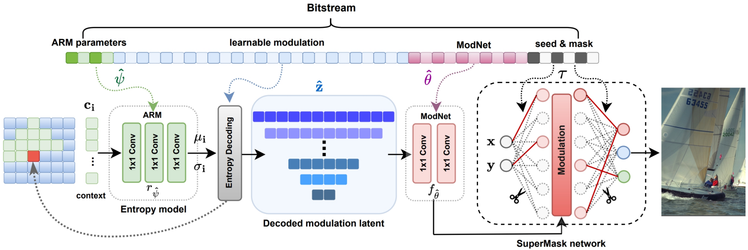

Fig. 7: Image decoding process in LotteryCodec.

- The receiver initializes a random network and uses a modulated subnetwork to reconstruct the source.

Rate and distortion expressions:

$$\mathsf{R} = \mathbb{E}_{\mathbf{S}\sim p_s}\left[-\log_2 p_{\hat{\psi}}(\hat{\mathbf{z}}) - \log_2 p(\tau) + \mathsf{R}_{\hat{\theta}} + \mathsf{R}_{\hat{\psi}}\right], \mathsf{D} = \mathbb{E}_{\mathbf{S}\sim p_s}\left[d(\mathbf{S}, g_{\tau\odot \mathbf{w}_0}(f_{\hat{\theta}}(\hat{\mathbf{z}}), \mathbf{x}))\right]$$

RD cost optimization: $\mathcal{L} = \mathsf{D} + \lambda \mathsf{R}(\hat{\mathbf{z}})$.

- This loss function excludes the rate-term from networks and mask for their minimal contributions.

4. LotteryCodec scheme

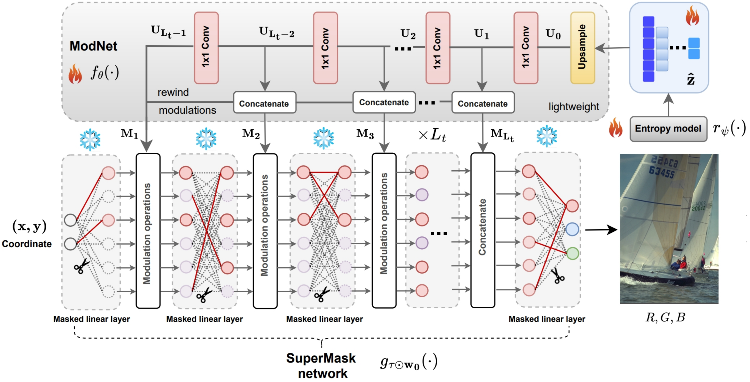

Fig. 8: Illustration of the rewind modulation mechanism in LotteryCodec.

- Global **edge-popup** algorithm [6] for mask learning, combined with **Fourier initialization** [10].
- Concatenating modulated neurons in a **rewind fashion** enriches the SuperMask with sign and magnitude cues, enabling deeper-layer feature reactivation while preserving high-level representations.

5. Experimental Results

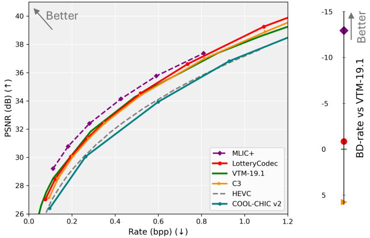

RD curve and BD rate on Kodak.

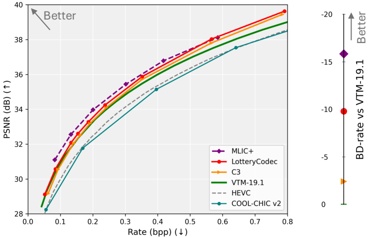

RD curve and BD rate on CLIC2020.

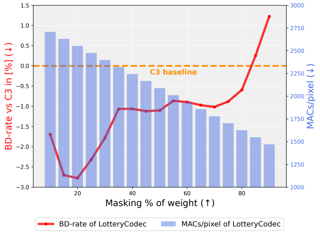

BD-rate and decoding complexity across different mask ratios on Kodak.

Fig. 9: Performance of LotteryCodec and other schemes.

Experimental Observations

- LotteryCodec achieves state-of-the-art RD performance among overfitted codecs, and outperforms VTM-19.1 on both Kodak and CLIC2020.
- Decoding complexity drops with increasing mask ratio, while RD performance peaks at a 20% mask ratio, lower than the 50% observed without rewind modulation (lower structure entropy for simplified coding).

5. Experimental Results

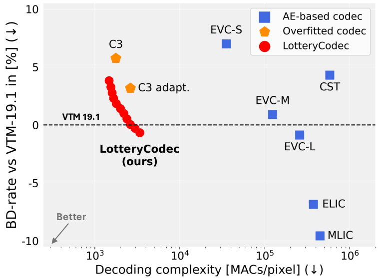

RD vs. decoding complexity on Kodak.

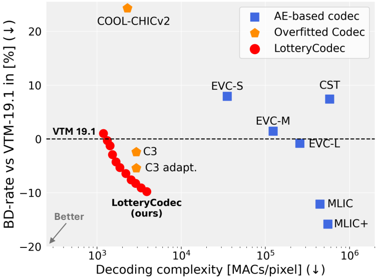

RD vs. decoding complexity on CLIC2020.

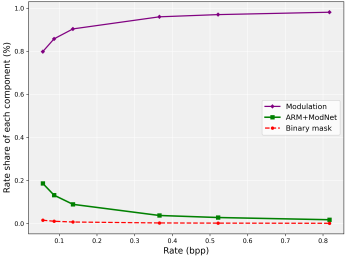

Compression cost distribution.

Fig. 9: Performance of LotteryCodec and other schemes (more ratio choices -> more flexible coding scheme).

Experimental Observations

- LotteryCodec enables flexible RD–complexity trade-offs by adjusting the mask ratio, where more detailed results are updated on our project page.
- LotteryCodec achieves comparable or better RD performance than most AE-based codecs with at least 10× fewer MACs, and significantly outperforms other overfitted codecs at lower complexity.

5. Experimental Results
Coding efficiency

Table: Coding time for Kodak images on NVIDIA L40S (GPU) and Intel Xeon Platinum 8358 (CPU) with a masking ratio of 0.8 under structured pruning. Orange indicates GPU computation; blue indicates CPU computation.

Models	Encoding time	Decoding time
VTM 19.1	85.53 (s)	352.52 (ms)
EVC (S/M/L)	$20.23/32.21/51.35$ (ms)	$18.82/23.73/32.56$ (ms)
MLIC+	205.60 (ms)	271.31 (ms)
LotteryCodec (d $= 8/16/24$)	$13.86/14.64/14.92$ (sec/1k steps)	$261.33/267.58/278.31$ (ms)
C3 (d $= 12/18/24$)	$13.10/13.98/14.32$ (sec/1k steps)	$272.15/284.67/295.03$ (ms)

Experimental Observations

- LotteryCodec offers faster decoding with slightly higher encoding time than other overfitted codecs.
- While real-world latency depends on many factors and can benefit from engineering optimizations (e.g., C APIs), all reported results use the same unoptimized setup for fairness.

6. Conclusion and Future Work

- We introduced and validated the **lottery codec hypothesis**, proposing **LotteryCodec**, that compresses images into modulation vectors and a binary mask for a randomly initialized network.
- LotteryCodec achieves **state-of-the-art** RD performance among overfitted codecs while maintaining **low**, **adaptive** complexity.
- It can be extended as an alternative for **video coding**, offering better RD performance and control over complexity and rate.
- Project page: `https://eedavidwu.github.io/LotteryCodec/`

Scan for code and resources

IMPERIAL

Thank you.
Questions?

LotteryCodec: Searching the Implicit Representation in a Random Networks for Low-Complexity Image Compression
15/06/2025

References I

[1] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein, “Implicit neural representations with periodic activation functions,” Advances in neural information processing systems, vol. 33, pp. 7462–7473, 2020.

[2] J. Ballé, D. Minnen, S. Singh, S. J. Hwang, and N. Johnston, “Variational image compression with a scale hyperprior,” ICLR 2018-6th International Conference on Learning Representations, 2018.

[3] E. Dupont, A. Goliński, M. Alizadeh, Y. W. Teh, and A. Doucet, “Coin: Compression with implicit neural representations,” ICLR 2021-International Conference on Learning Representations Workshop Neural Compression 2021, 2021.

[4] T. Ladune, P. Philippe, F. Henry, G. Clare, and T. Leguay, “Cool-chic: Coordinate-based low complexity hierarchical image codec,” in Proceedings of the IEEE/CVF International Conference on Computer Vision, 2023, pp. 13 515–13 522.

[5] D. Ulyanov, A. Vedaldi, and V. Lempitsky, “Deep image prior,” in Proceedings of the IEEE/CVF International Conference on Computer Vision, 2018, pp. 9446–9454.

References II

[6] V. Ramanujan, M. Wortsman, A. Kembhavi, A. Farhadi, and M. Rastegari, "What's hidden in a randomly weighted neural network?" In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2020, pp. 11 893–11 902.

[7] A. Pensia, S. Rajput, A. Nagle, H. Vishwakarma, and D. Papailiopoulos, "Optimal lottery tickets via subset sum: Logarithmic over-parameterization is sufficient," Advances in neural information processing systems, vol. 33, pp. 2599–2610, 2020.

[8] A. da Cunha, E. Natale, and L. Viennot, "Proving the strong lottery ticket hypothesis for convolutional neural networks," in ICLR 2022-10th International Conference on Learning Representations, 2022.

[9] H. Kim, M. Bauer, L. Theis, J. R. Schwarz, and E. Dupont, "C3: High-performance and low-complexity neural compression from a single image or video," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 9347–9358.

[10] K. Shi, X. Zhou, and S. Gu, "Improved implicit neural representation with fourier reparameterized training," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 25 985–25 994.