Electronics and Telecommunications Research Institute

CVPR JUNE 3-7, 2026 DENVER COLORADO

ELiC: Efficient LiDAR Geometry Compression via Cross-Bit-depth Feature Propagation and Bag-of-Encoders

Junsik Kim, Gun Bang, Soowoong Kim

Electronics and Telecommunications Research Institute (ETRI), Daejeon, Korea

Challenges in LiDAR Geometry Compression

📦 Massive Data ➔ Storage & Bandwidth Bottleneck

🌐 High Bit-Depth Sparsity ➔ Broken Spatial Context

⚡ Real-Time Constraint ➔ Limited Compute Budget

Due to these challenges, achieving both high compression efficiency and real-time performance becomes difficult.

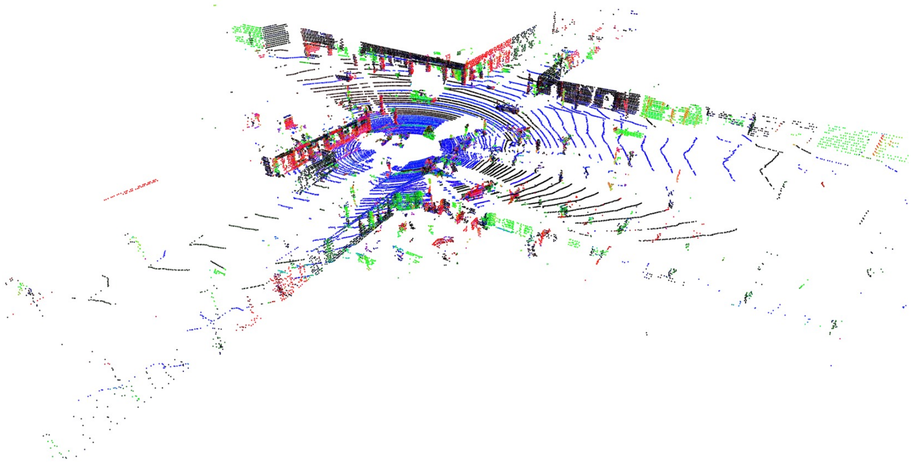

Our Approach: ELiC

🔁 Reuse Information ➔ Cross-bit-depth Feature Propagation

🎛️ Adapt to Distribution ➔ Bag-of-Encoders

⚡ Reduce Overhead ➔ Morton-order Hierarchy

Together, these components enable ELiC to achieve efficient compression while maintaining real-time performance.

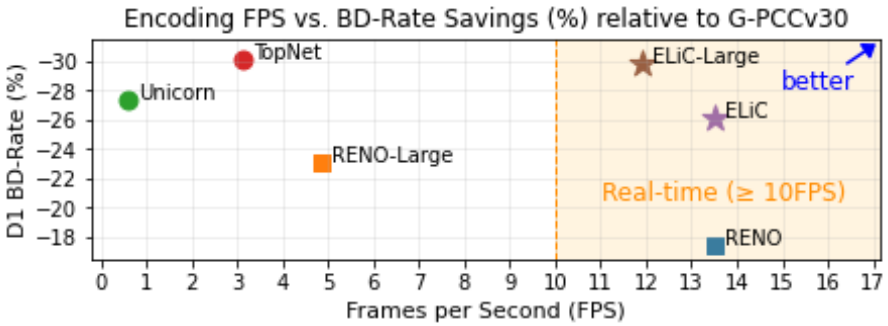

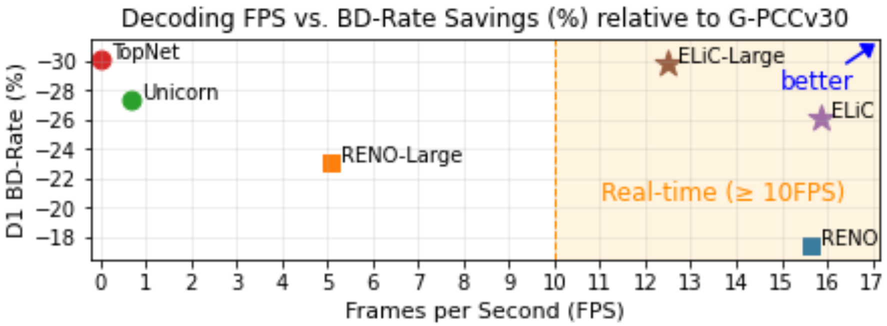

Cross-bit-depth Feature Propagation

Electronics and Telecommunications Research Institute

- Reuse dense low-level context for sparse high-level prediction

Conventional: Level-independent Processing

- Each depth processed independently
- Re-derives spatial context every time
- Fails in sparse high bit-depth

Proposed: Cross-bit-depth Feature Propagation

- Reuse features from lower depths
- Propagate rich contextual information
- Robust even in sparse regions

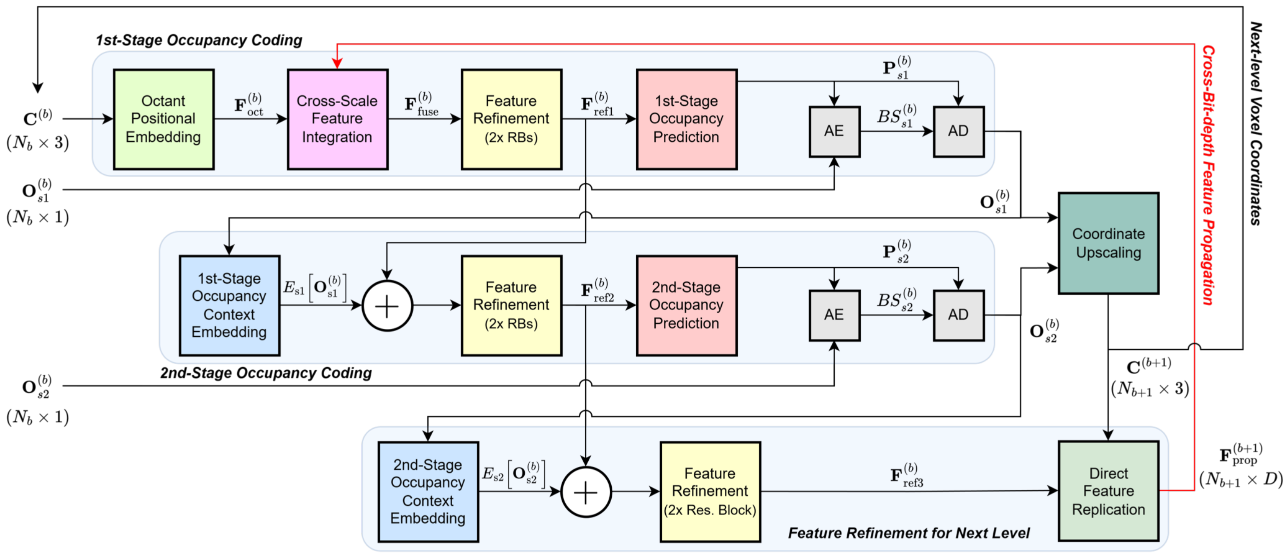

Bag-of-Encoders

- Different bit-depth levels require different model capacities

Distribution varies across bit-depths

- Different occupancy patterns per level
- Single encoder is suboptimal

Bag-of-Encoders (BoE)

- Multiple encoders with diverse capacities
- Select the best encoder per level
- Match model to data distribution

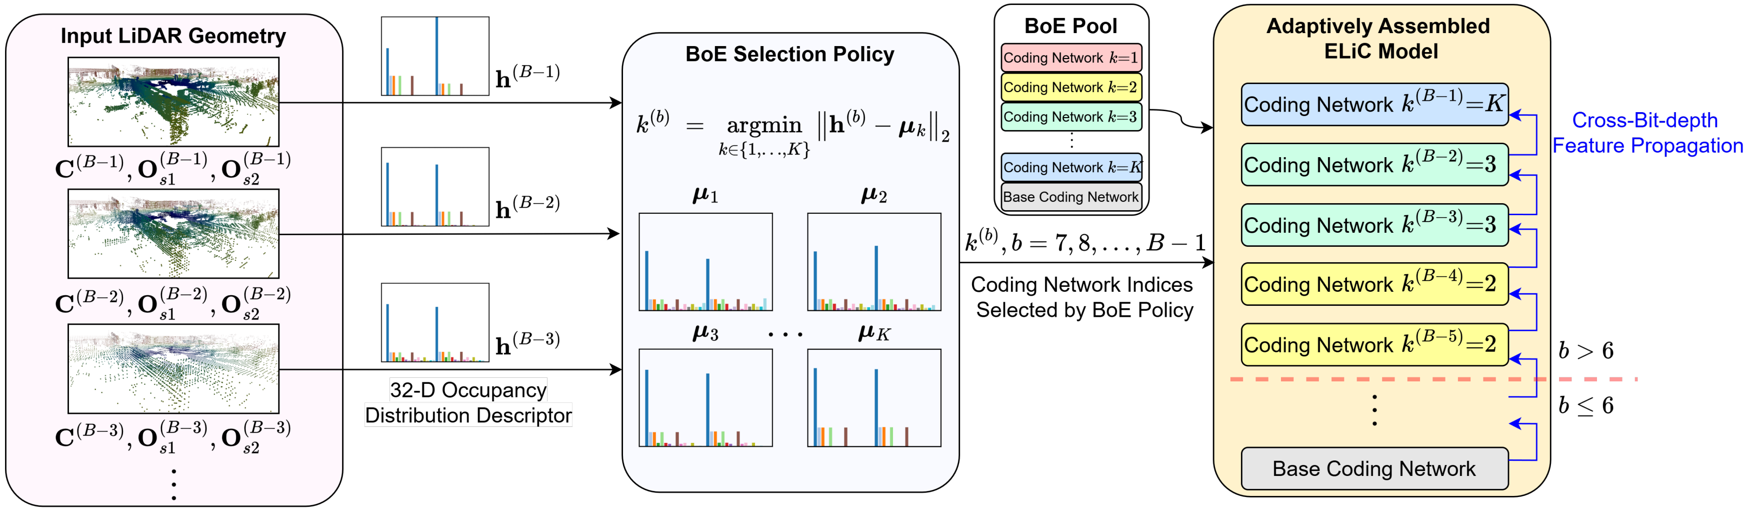

Morton-order Hierarchy

- Remove redundant sorting for faster processing

Repeated Sorting Overhead

- Sorting required at each bit-depth level
- High computational overhead
- Increased latency

Morton-order Hierarchy

- Maintain a global Morton order
- Eliminate per-level sorting
- Efficient hierarchical traversal

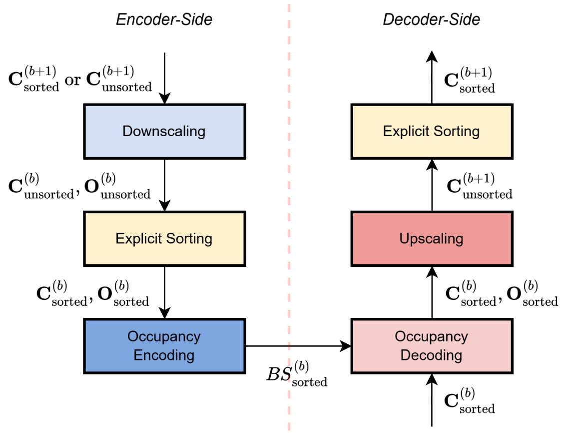

With explicit per-level sorting

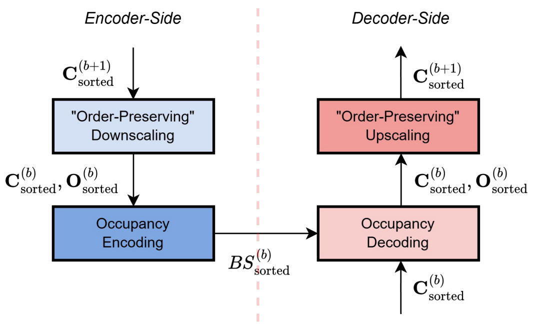

With Morton-order-preserving hierarchy

Results: Runtime Efficiency

- ELiC achieves the fastest runtime among all methods
  - Up to 4× faster than G-PCC [1]
  - Comparable to RENO [2], but with better compression

Table 1. Per-frame average encoding and decoding time measured on Ford and SemanticKITTI datasets. For TopNet, decoding time is reported over 10 frames due to its extremely long runtime.

Input	**G-PCC**		**RENO**		**RENO-Large**		**Unicorn**		**TopNet**		**ELiC**		**ELiC-Large**	
Bit-depth	**enc (s)**	**dec (s)**	**enc (s)**	**dec (s)**	**enc (s)**	**dec (s)**	**enc (s)**	**dec (s)**	**enc (s)**	**dec (s)**	**enc (s)**	**dec (s)**	**enc (s)**	**dec (s)**
16 bit	0.700	0.845	0.173	0.179	0.766	0.803	5.816	5.173	1.922	2357.856	0.172	0.157	0.198	0.168
15 bit	0.614	0.734	0.152	0.147	0.589	0.584	4.271	3.922	1.532	2091.421	0.145	0.135	0.162	0.143
14 bit	0.540	0.600	0.122	0.120	0.434	0.423	2.842	2.516	0.917	1136.171	0.121	0.112	0.14	0.119
13 bit	0.418	0.434	**0.098**	**0.090**	0.299	0.290	1.899	1.767	0.608	638.485	**0.095**	**0.086**	0.105	**0.096**
12 bit	0.302	0.266	**0.074**	**0.064**	0.207	0.197	1.661	1.473	0.322	308.066	**0.074**	**0.063**	**0.084**	**0.08**
Avg.	0.515	0.576	0.124	0.120	0.459	0.459	3.298	2.970	1.060	1306.400	0.121	0.111	0.138	0.121
vs. ELiC	4.24×	5.20×	1.02×	1.08×	3.78×	4.15×	27.16×	26.86×	8.73×	11,811×	1.00×	1.00×	1.14×	1.10×

[1] WG07 MPEG, "Enhanced G-PCC Test Model v30," ISO/IEC, 2025.
[2] K. You et al., "RENO: Real-Time Neural Compression for 3D LiDAR Point Clouds," CVPR, 2025.

Results: Compression Performance

ETRI Electronics and Telecommunications Research Institute

- ELiC achieves state-of-the-art compression efficiency
  - Ford [3]: Best BD-rate (-26.5%)
  - KITTI [4]: Near-SOTA (-33.3%)
- Better compression–runtime trade-off than all methods

Table 2. BD-Rate (%) on Ford and SemanticKITTI relative to the G-PCCv30. D1 = point-to-point, D2 = point-to-plane. Lower is better. Results are based on bit-depths {16, 15, 14, 13, 12}.

BD-Rate	**Ford**		**KITTI**	
vs. G-PCCv30	D1 (%)	D2 (%)	D1 (%)	D2 (%)
RENO	-14.02	-14.03	-20.90	-20.92
RENO-Large	-14.70	-14.70	-31.52	-31.53
Unicorn	-25.41	-25.42	-29.28	-29.30
TopNet	-26.26	-26.26	-34.10	-34.12
ELiC	-22.97	-22.97	-29.23	-29.24
ELiC-Large	-26.54	-26.53	-33.26	-33.27

[3] G. Pandey et al., "Ford Campus Vision and LiDAR Dataset," IJRR, 2011.
[4] J. Behley et al., "SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences," ICCV, 2019.

Ablation: Effect of Feature Propagation and BoE

- Feature propagation reduces bits even without BoE
- BoE further improves compression efficiency

Table 3. Bitrate savings and relative encoding/decoding runtime ratios of ELiC w/o BoE (1.28 MB) compared to RENO (1.10 MB).

Ford		KITTI		Relative Runtime	
D1 (%)	D2 (%)	D1 (%)	D2 (%)	enc (s)	dec (s)
-8.60	-8.59	-6.61	-6.60	0.90×	0.85×

Table 4. Bitrate savings (D1 BD-rates) of ELiC according to BoE pool size K, relative to ELiC w/o BoE

BoE Pool Size	Ford D1 (%)	KITTI D1 (%)	BoE Pool Size	Ford D1 (%)	KITTI D1 (%)
$K = 3$	-1.31	-2.89	$K = 4$	-1.55	-3.06
$K = 5$	-1.98	-3.54	$K = 6$	-1.56	-3.53
$K = 7$	-1.43	-3.44	Per-Level	-2.52	-4.32

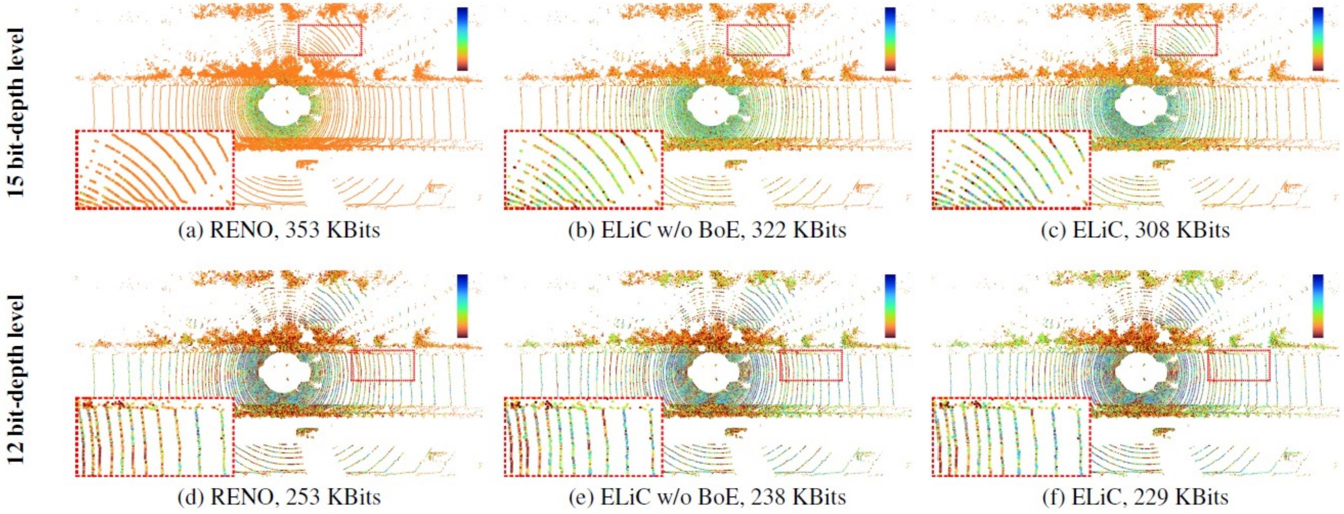

Figure 5. Per-point bit allocation on SemanticKITTI frame at the 15 and 12 bit-depth levels for RENO, ELiC w/o BoE, and ELiC (K=5). The colormap encodes lower (blue) to higher (dark orange) predicted bits per point.

Ablation: Effect of Morton-order Hierarchy

- Eliminates per-level sorting overhead
- ~15% latency reduction by removing sorting
  - Encoding: 0.142s → 0.121s
  - Decoding: 0.128s → 0.111s

Table 5. Runtime comparison between ELiC with explicit sorting and ELiC with Morton-order-preserving hierarchy.

ELiC with Explicit Sorting		ELiC with Morton-Order		Latency Reduction	
enc (s)	dec (s)	enc (s)	dec (s)	enc (s)	dec (s)
0.142	0.128	0.121	0.111	-0.021	-0.017

Conclusion

ETRI
Electronics and Telecommunications Research Institute

🔁 Reuse ➔ Preserve spatial context

🎛️ Adapt ➔ Match data distribution

⚡ Reduce ➔ Eliminate overhead

- State-of-the-art compression performance
- Real-time runtime (≥10 FPS)
- Best compression–runtime trade-off
- ELiC bridges the gap between efficiency and real-time LiDAR compression

ELiC