【Accepted Paper】

Title

Indoor Multi-View Radar Object Detection via 3D Bounding Box Diffusion

Authors

Ryoma Yataka（Melco）, Pu (Perry) Wang（MERL）, Petros Boufounos（MERL）, Ryuhei Takahashi（Melco）

Abstract

Indoor perception using radar (*1) has attracted increasing attention for applications such as indoor monitoring and surveillance, due to its ability to preserve privacy and its high reliability even under hazardous conditions such as fires. However, radar sensors suffer from low spatial resolution, making it difficult to achieve highly accurate predictions while maintaining strong generalization performance.

In this paper, we propose Radar Object dEtection with 3D Bounding boX Diffusion (REXO) (*2) to improve the accuracy of radar-based perception. REXO explicitly associates 2D radar features obtained from multiple viewpoints within a shared 3D space common to all viewpoints, enabling the construction of 3D bounding boxes (BBox) in the radar coordinate system with strong generalization capability. Furthermore, by progressively applying denoising processes to the BBoxes in this 3D space, REXO achieves state-of-the-art BBox estimation accuracy, significantly outperforming existing methods.

(*1) A technology that detects and recognizes surrounding environments and objects using radar. Since radar uses radio waves to measure object position, velocity, and shape, it provides high reliability even in low-visibility or dark environments.
(*2) A deep learning model for object detection using radar data.

【Accepted Paper】

Title

LatentLLM: Activation-Aware Transform to Multi-Head Latent Attention

Authors

Toshiaki Koike-Akino(MERL), Xiangyu Chen(MERL), Jing Liu(MERL), Ye Wang(MERL), Pu (Perry) Wang(MERL),Matthew Brand(MERL)

Abstract

Large language models (LLMs) and multi-modal models (LMMs) have shown high performance but face challenges in practical deployment due to their high computational and memory requirements.
We have developed LatentLLM, an activation- and attention-aware model compression technology.
The method enables high-accuracy, reduced-dimension models without additional training by jointly compressing multiple weights.
This technology is expected to contribute to energy-efficient and scalable real-world deployment of advanced AI systems.

【Accepted Paper】

Title

Chain-of-Thought Driven Adversarial Scenario Extrapolation for Robust Language Models

Authors

Md Rafi Ur Rashid(Pennsylvania State University), Vishnu Asutosh Dasu (Pennsylvania State University),Ye Wang(MERL), Gang Tan (Pennsylvania State University), Shagufta Mehnaz (Pennsylvania State University)

Abstract

Large language models (LLMs) are widely used across many domains but still face safety challenges, including jailbreaks, toxic content, hallucinations, and social bias. Existing defenses often address only a single threat type or resort to rigid outright rejection, sacrificing user experience and failing to generalize across diverse and novel attacks. This paper introduces Adversarial Scenario Extrapolation (ASE), an inference-time framework that enables LLMs to internally anticipate adversarial scenarios and construct defensive responses without additional training. Evaluations show that ASE achieves near-zero jailbreak attack success rates and minimal toxicity, significantly reduces hallucinations and bias scores, and substantially lowers outright rejections. These results are expected to significantly contribute to improving the reliability of generative AI systems in real-world deployments, where both safety and usability are critically required.