



BEST PAPER AWARD

This is presented to

Di Liu, Meng Chen, Baotong Lu, Huiqiang Jiang, Zhenhua Han, Qianxi Zhang, Qi Chen, Chengruidong Zhang, Bailu Ding, Kai Zhang, Chen Chen, Fan Yang, Yuqing Yang, Lili Qiu

for the paper entitled as

Retrieval Attention: Accelerating Long-Context LLM Inference via Vector Retrieval

Presented at the Fourth NeurIPS Efficient Natural Language and Speech Processing Workshop 2024 (ENLSP-IV)
December 14th, 2024



VAHID PARTOVI NIA

ENLSP
Award Committee Chair



PEYMAN PASSBAN

ENLSP
Award Committee Chair



LILI MOU

ENLSP
Award Committee Chair



Sponsored by

