LLM-Driven Composite Neural Architecture Search for Multi-Source RL State Encoding

Yu Yu, Qian Xie, Junping Li, Aghamatlab Akbarzade, Nairen Cao, Li Jin

Introduction

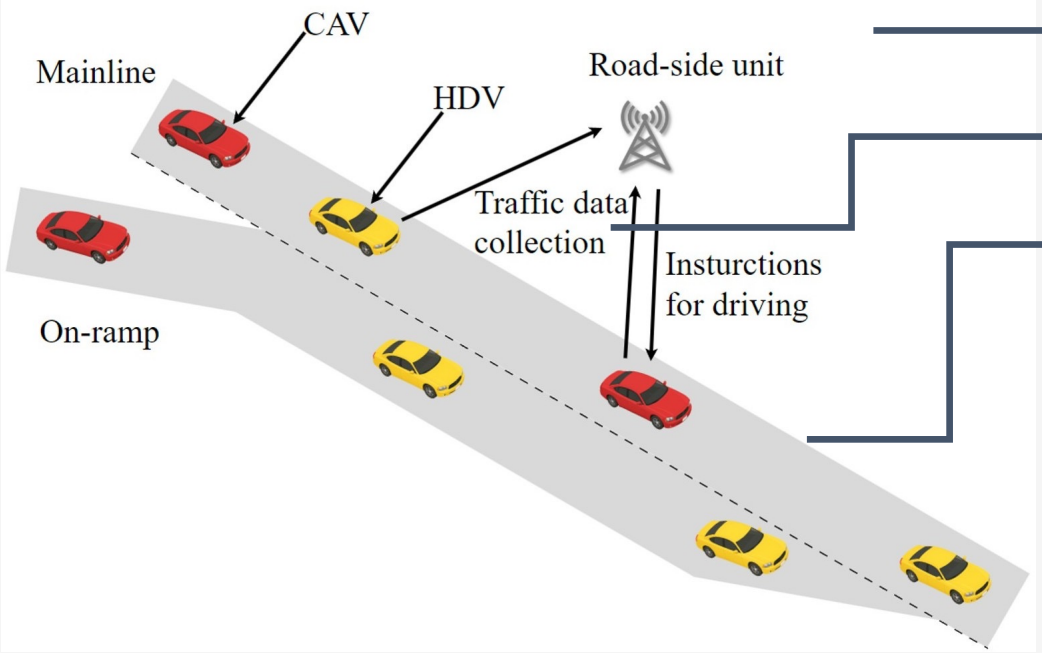

- Temporal Context
- Traffic State
- Sequence History

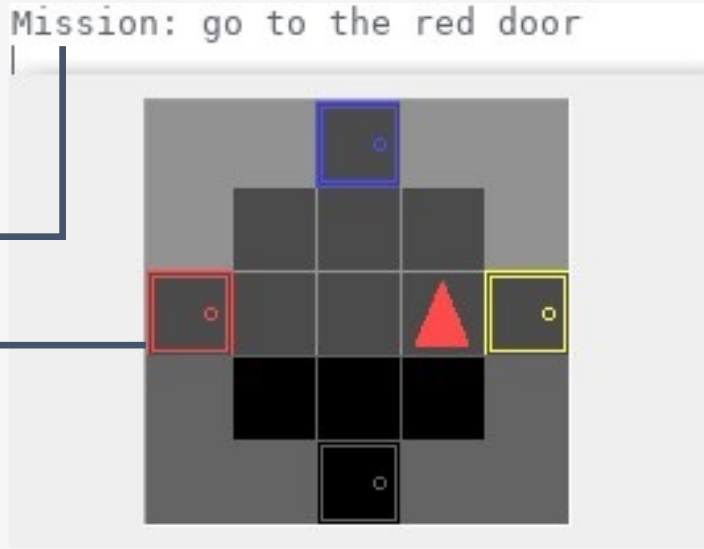

- Text Information
- Image Observation

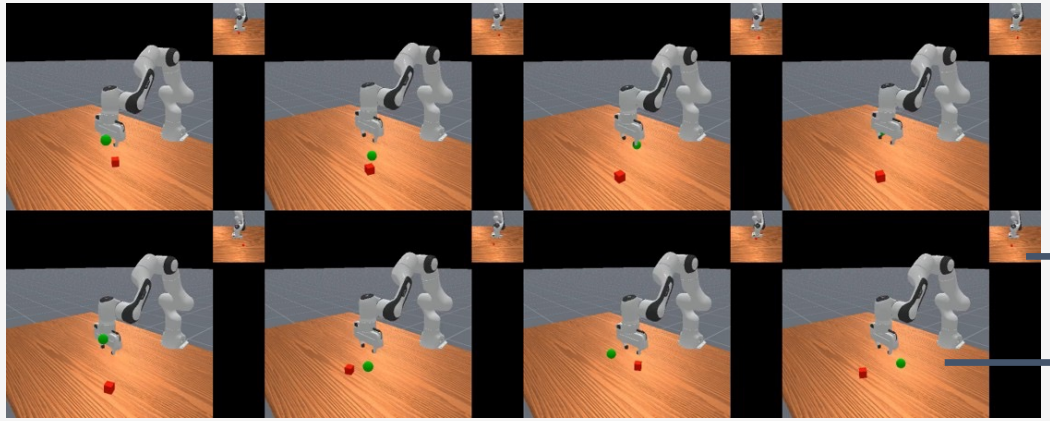

- State Information
- RGB Observation

Multi-Source RL State Encoding

Mixed-autonomy traffic control

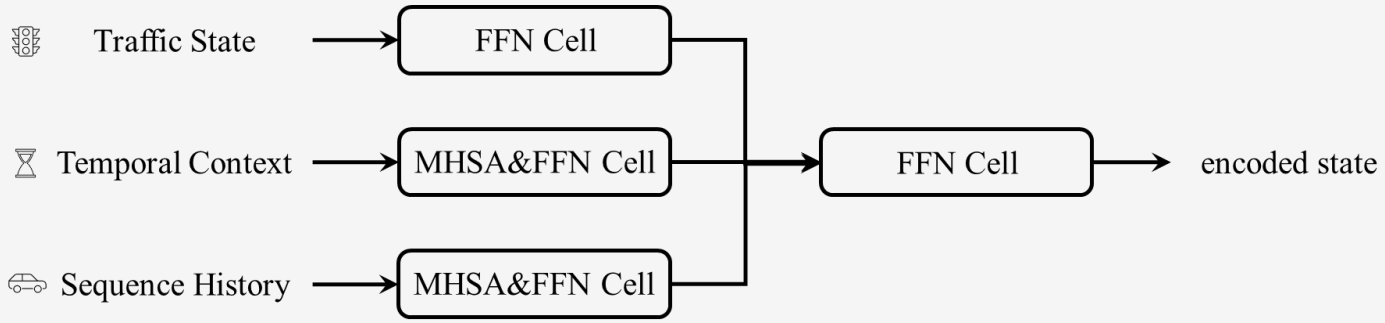

MiniGrid goal-oriented tasks

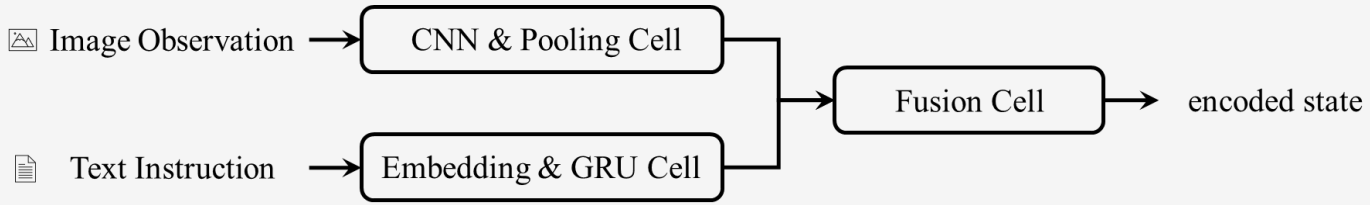

ManiSkill robotic control

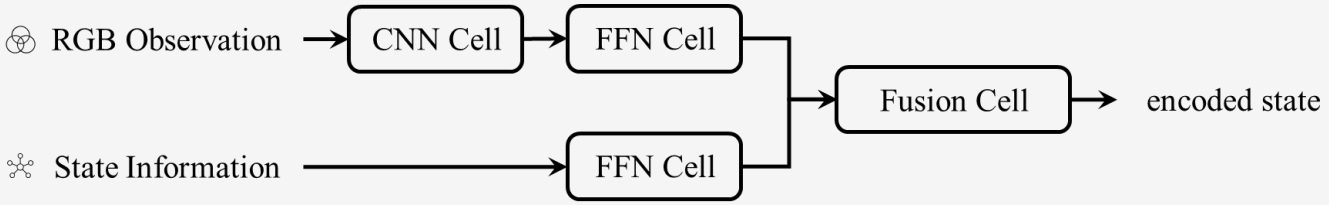

Designing state encoders for multi-source-RL remains underexplored.

Composite Neural Architecture Search

NAS for supervised learning

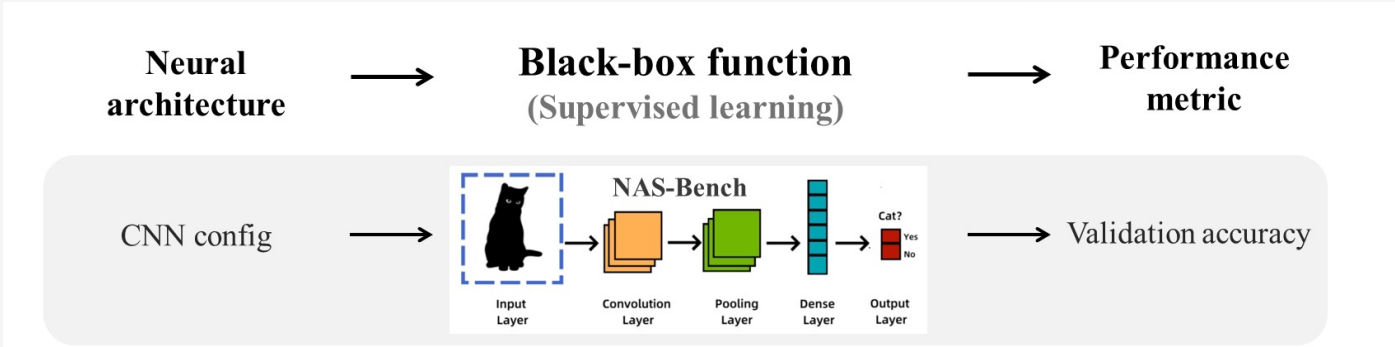

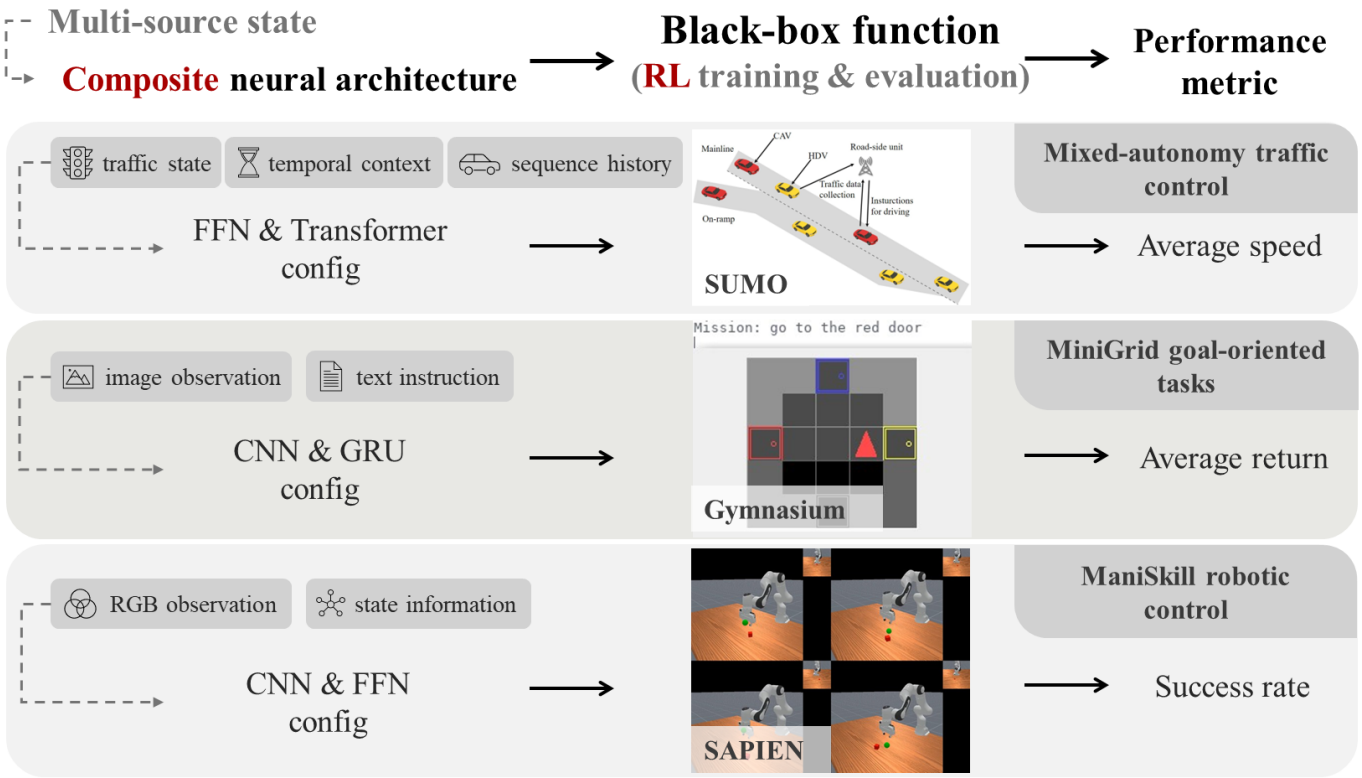

NAS for multi-source RL state encoding

- Unlike supervised learning, multi-source RL requires composite NAS.

LACER: An LLM → State Encoder → RL Pipeline

- **L**LM-driven Neural **A**rchitecture Search for **C**omposite State **E**ncoders in **RL**

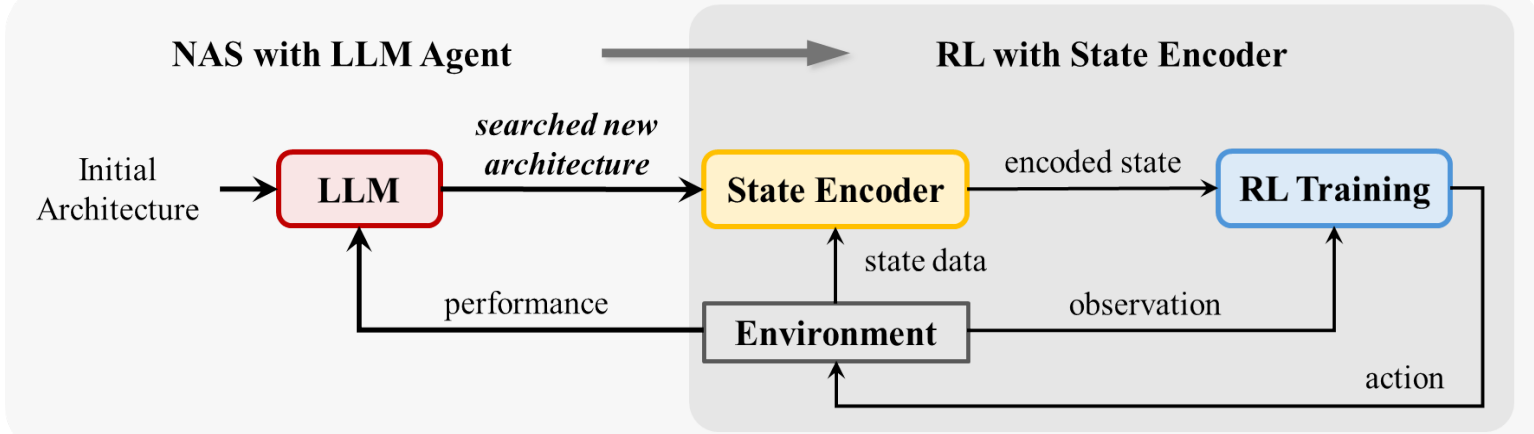

Module-specific search space

Table 1: Module-specific search spaces where bold values denote the configurations used by the Expert baseline.

Module	Operation	Heads / Activation	Dimension	Ratio	Depth
Time	MHSA, FFN	{**2**, 4, 8}	{**8**, 16, 32}	{1, **2**, 4}	{1, **2**, 3}
Traffic	FFN	{**relu**, gelu, swish}	{16, **32**, 64}	{**1**, 2, 4}	{**1**, 2, 3}
Sequence	MHSA, FFN	{2, **4**, 8}	{8, **16**, 32}	{1, **2**, 4}	{1, **2**, 3}
Fusion	FFN	{**relu**, gelu, swish}	{64, **128**, 256}	{**1**, 2, 4}	{**1**, 2, 3}

Pipeline comparison

GENIUS (existing LLM-based NAS)

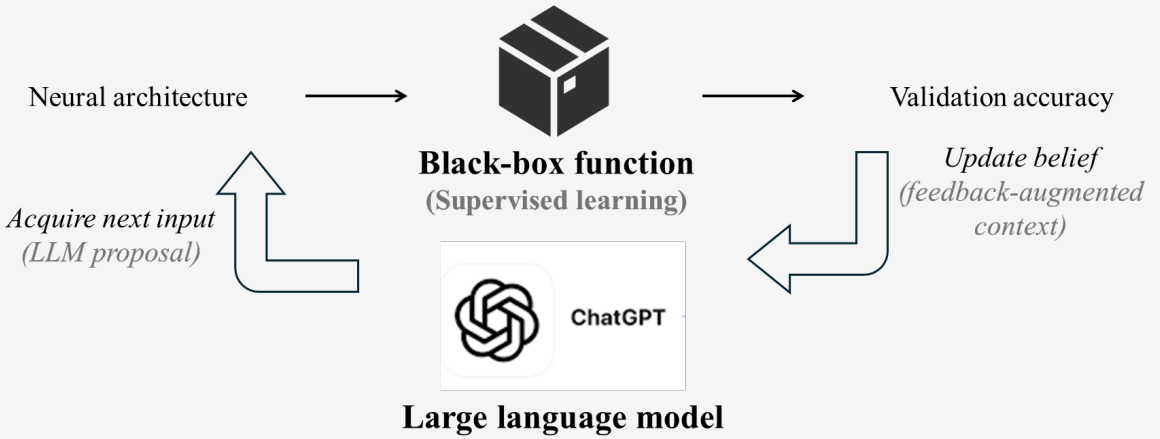

LACER (ours)

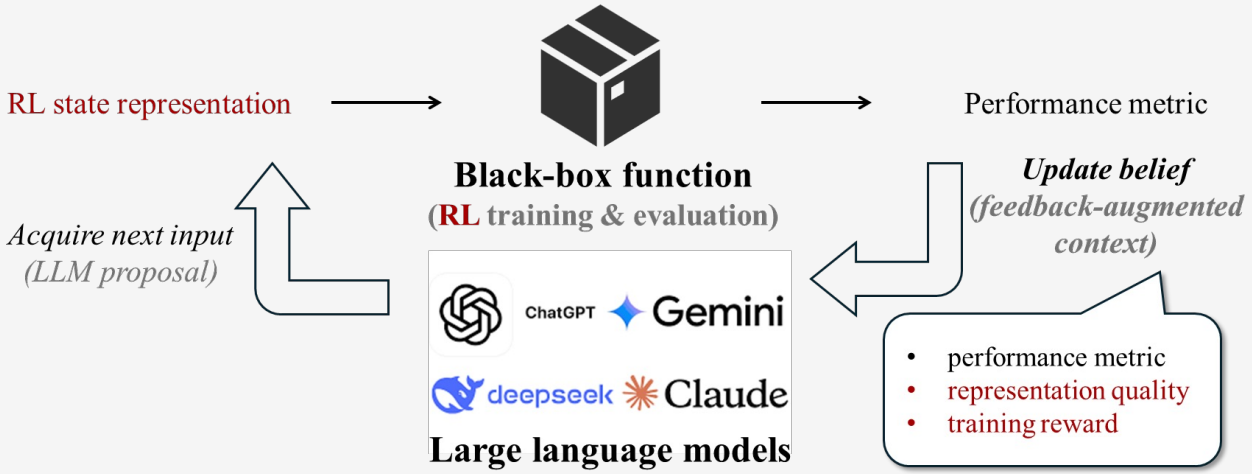

- Exploit side info on source-specific encoders beyond performance metric.

Results

Comparison between LACER and baselines

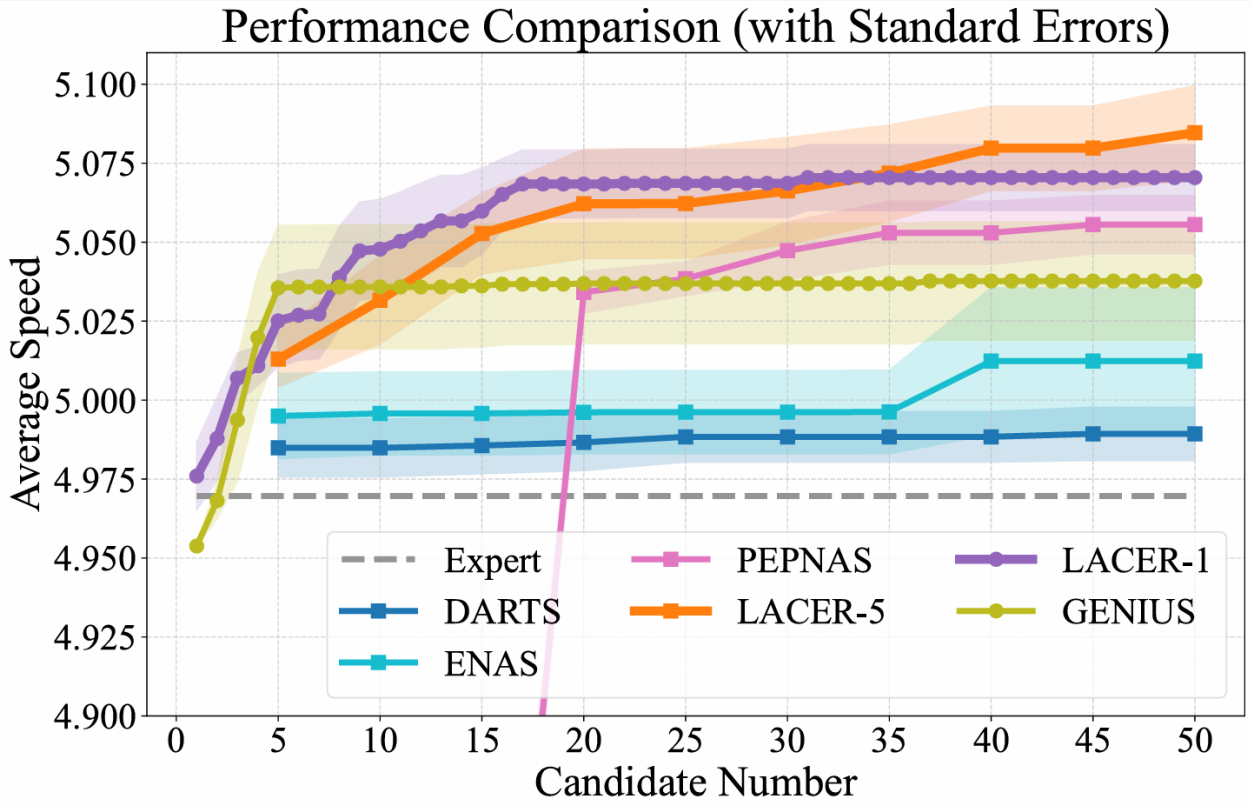

Ablation studies

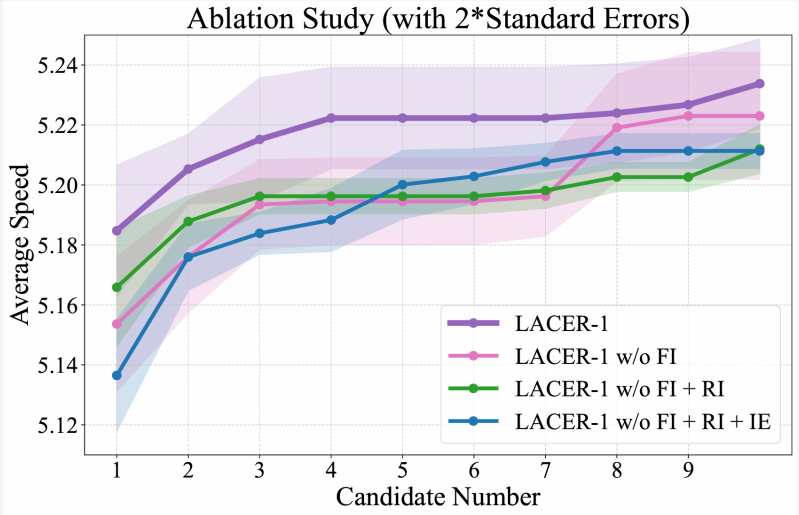

Computation time

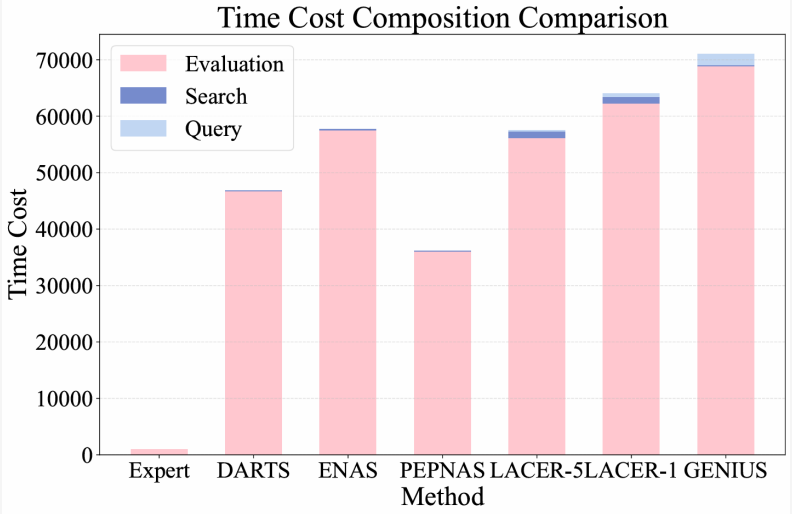

- LACER-5 denotes the batch variant (5 candidates per iteration).

LLM-Driven Composite Neural Architecture Search for Multi-Source RL State Encoding

Yu Yu, Qian Xie, Junping Li,
Aghamatlab Akbarzade,
Nairen Cao, Li Jin