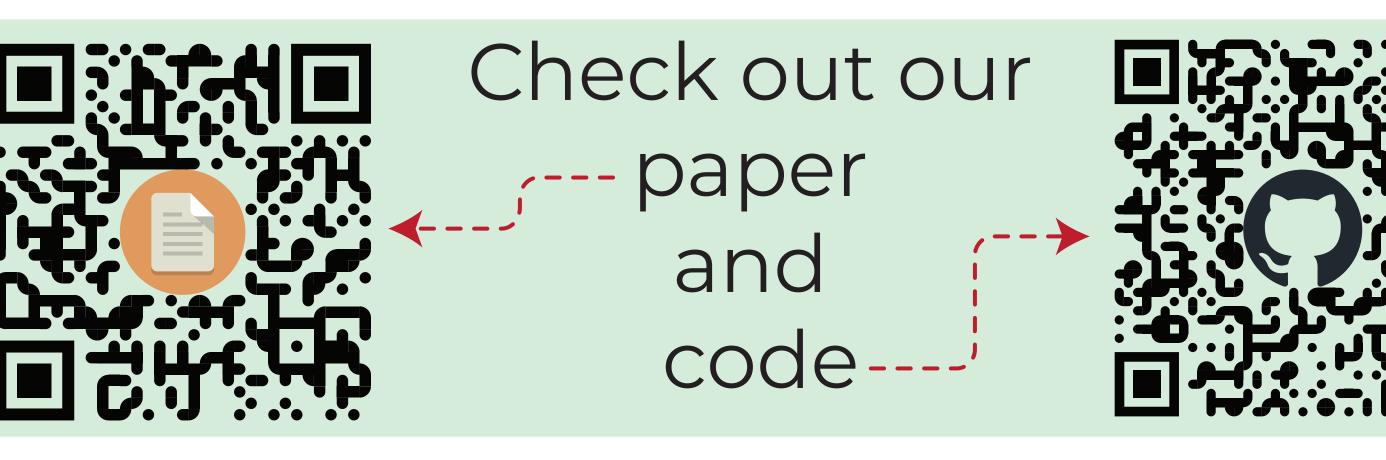
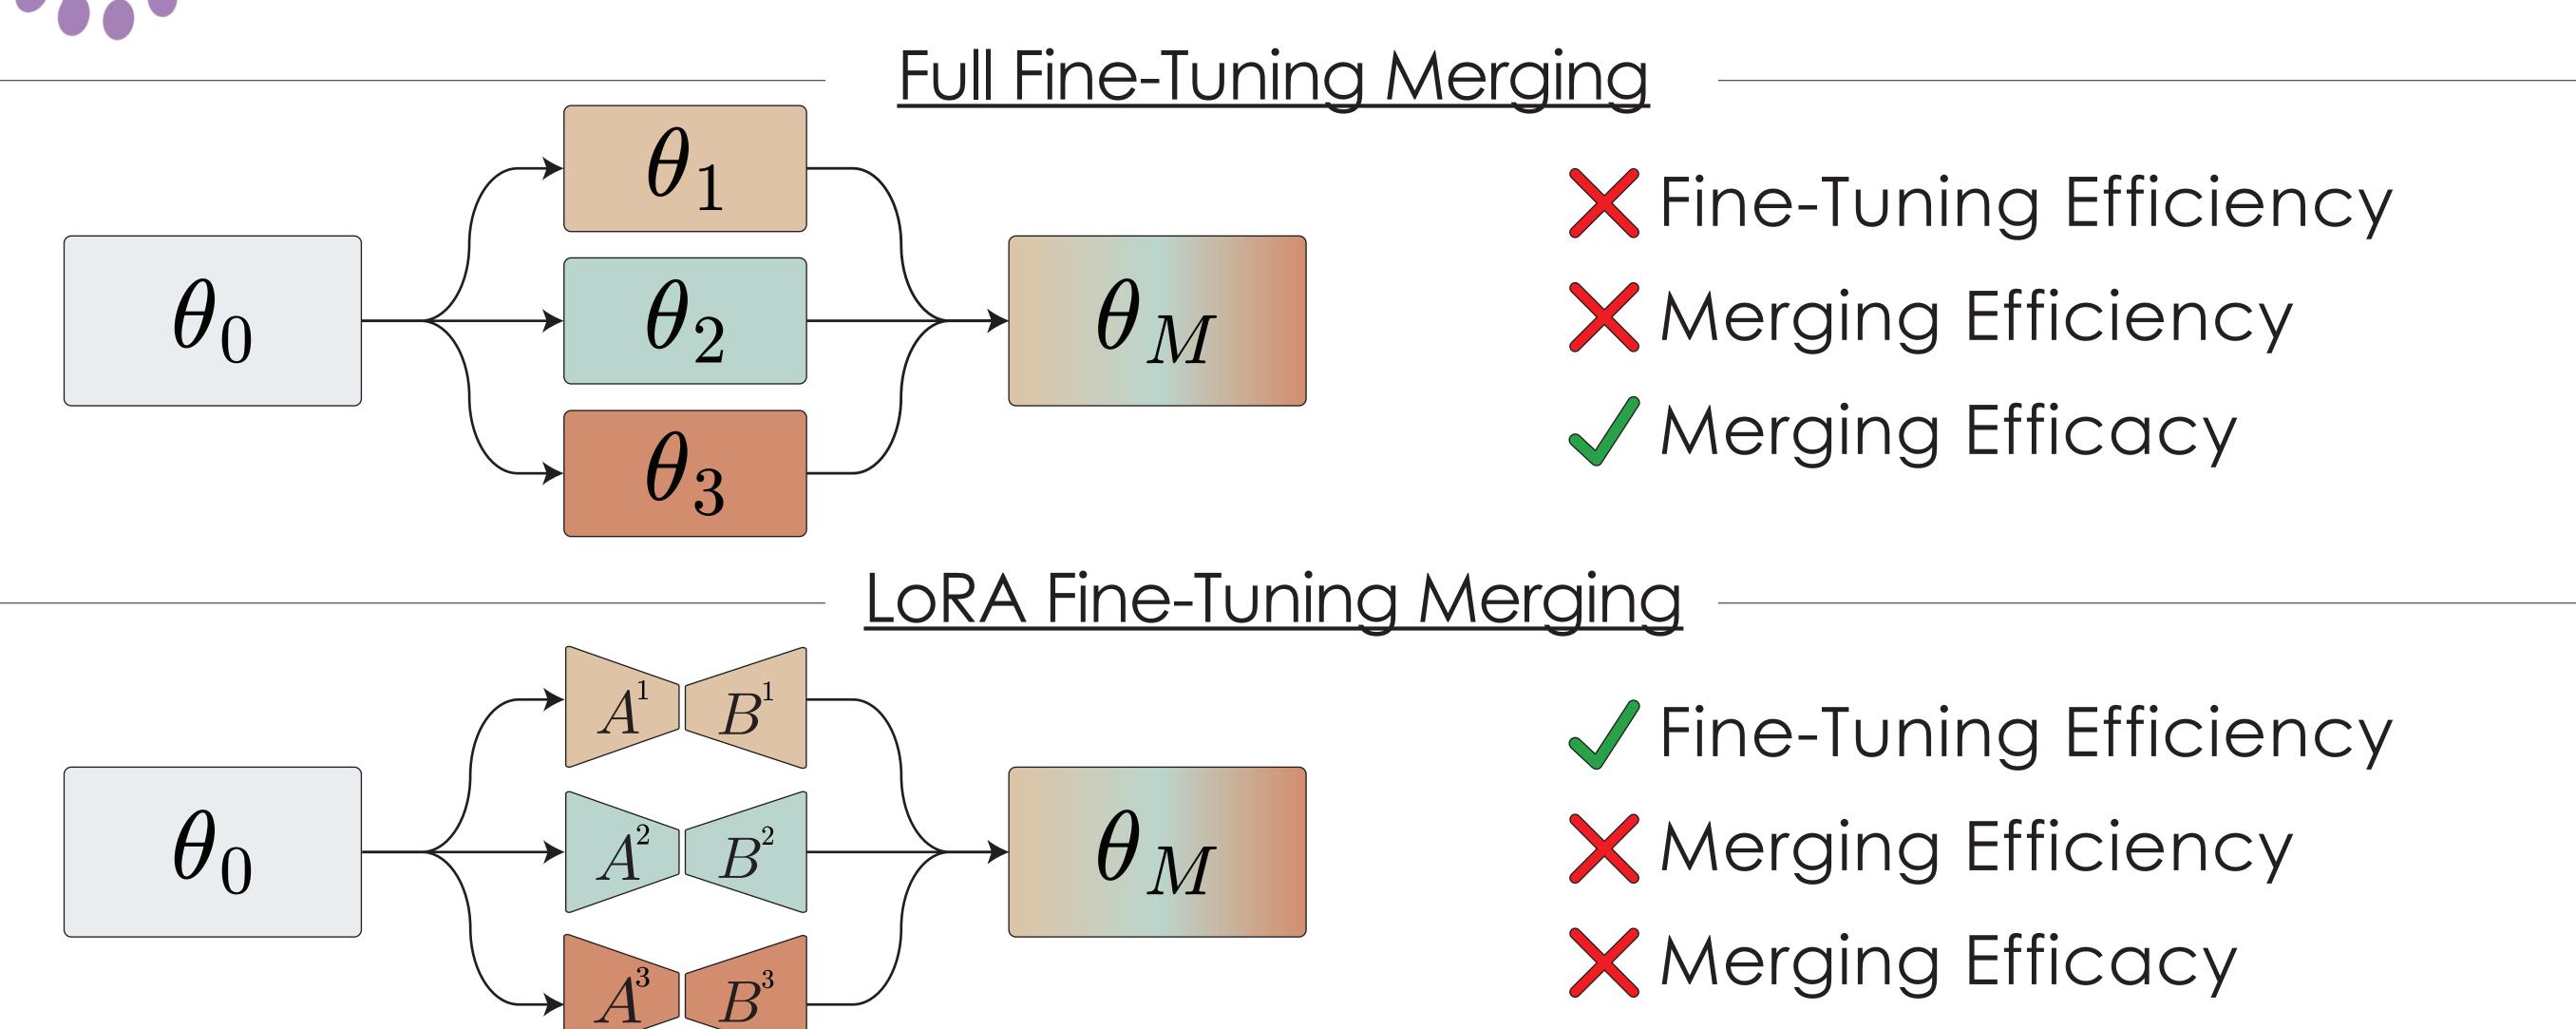
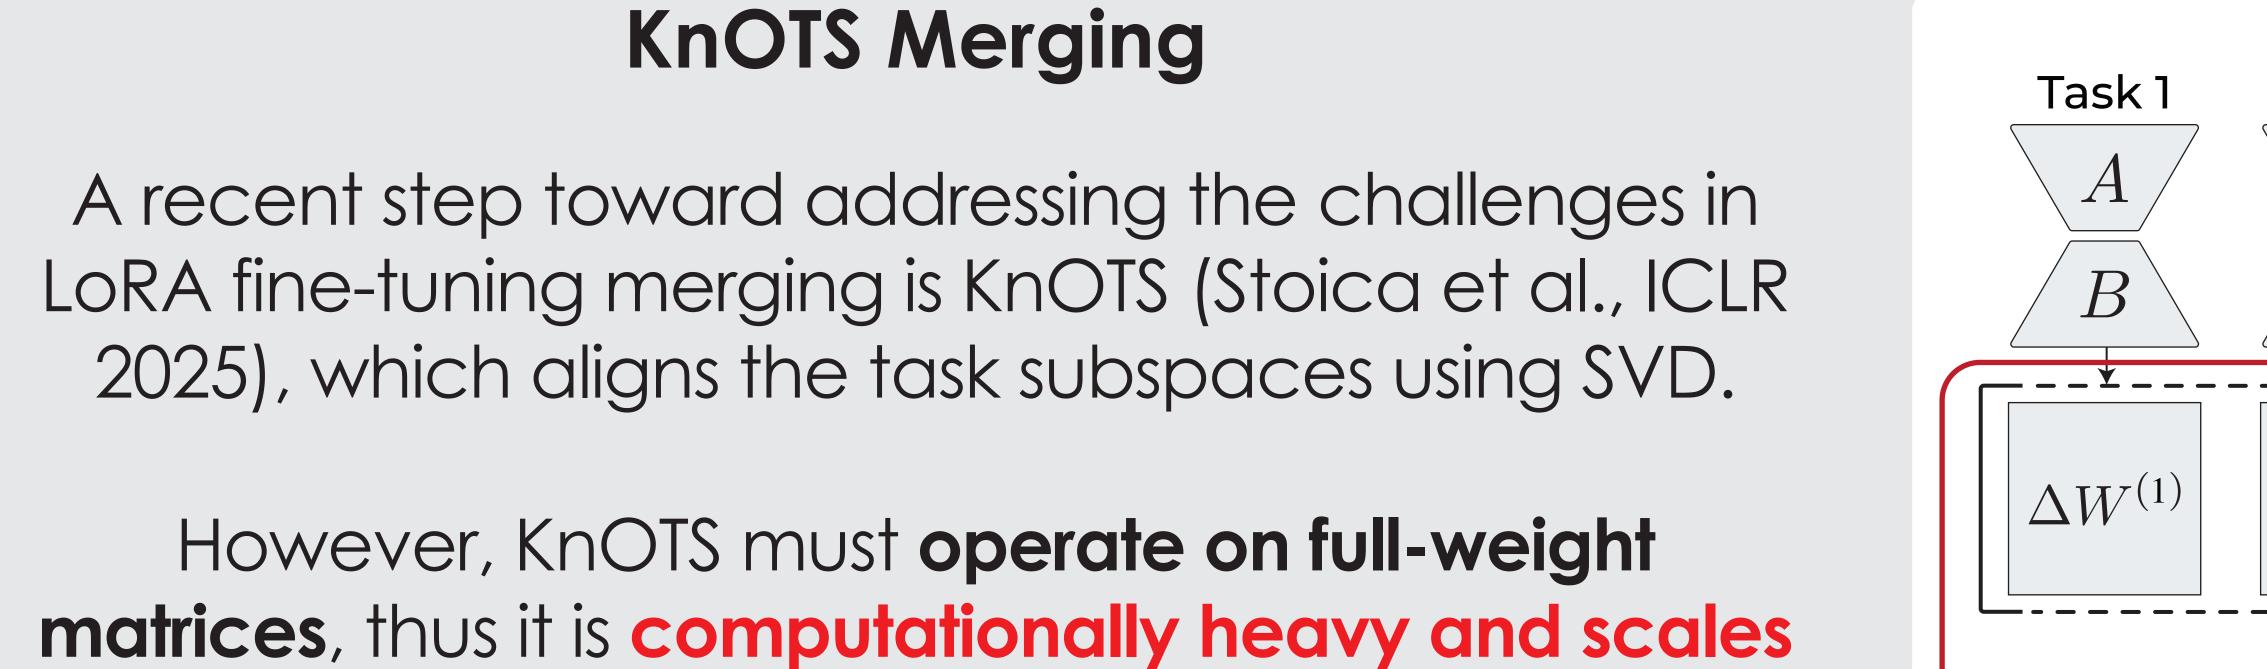


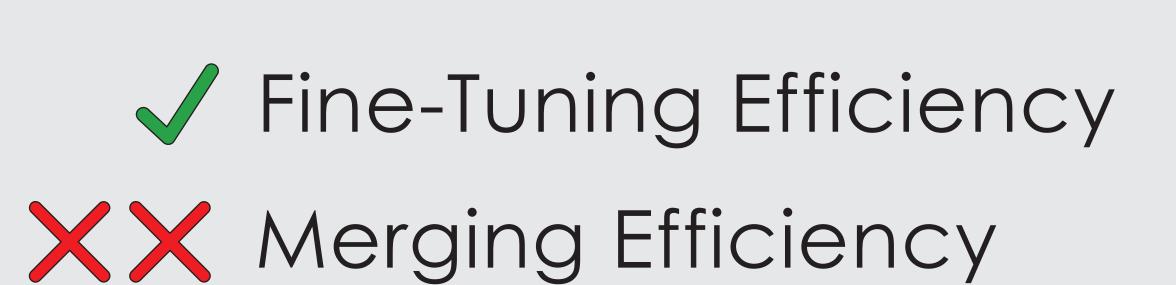
Accurate and Efficient Low-Rank Model Merging in Core Space

Aniello Panariello*, Daniel Marczak*, Simone Magistri, Angelo Porrello, Bartłomiej Twardowski, Andrew D. Bagdanov, Simone Calderara, Joost van de Weijer



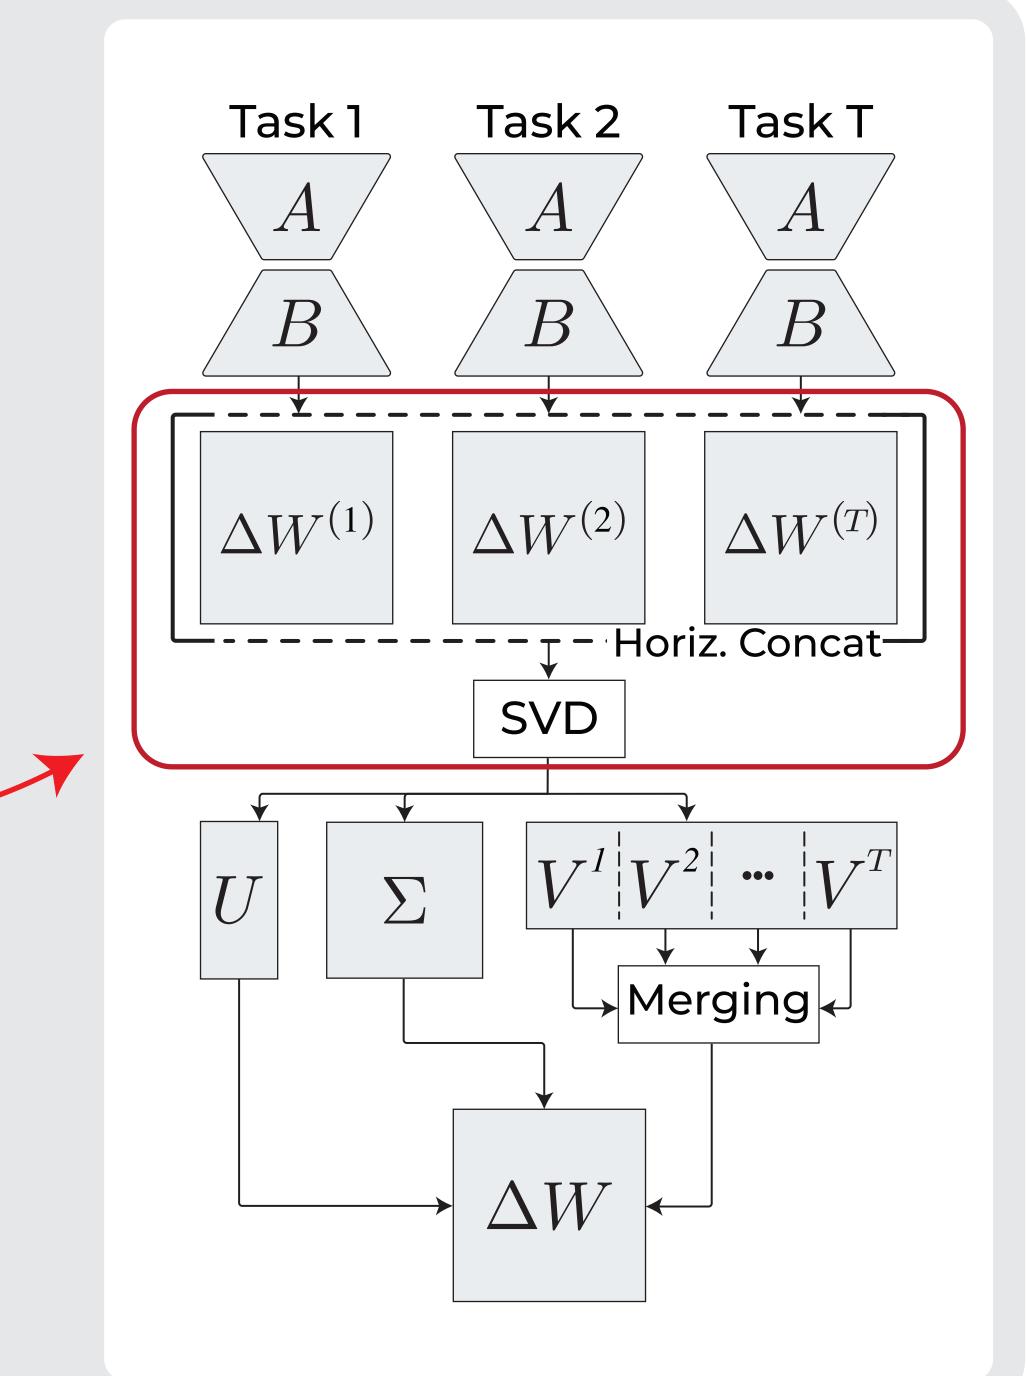






poorly with model size.

— Merging Efficacy



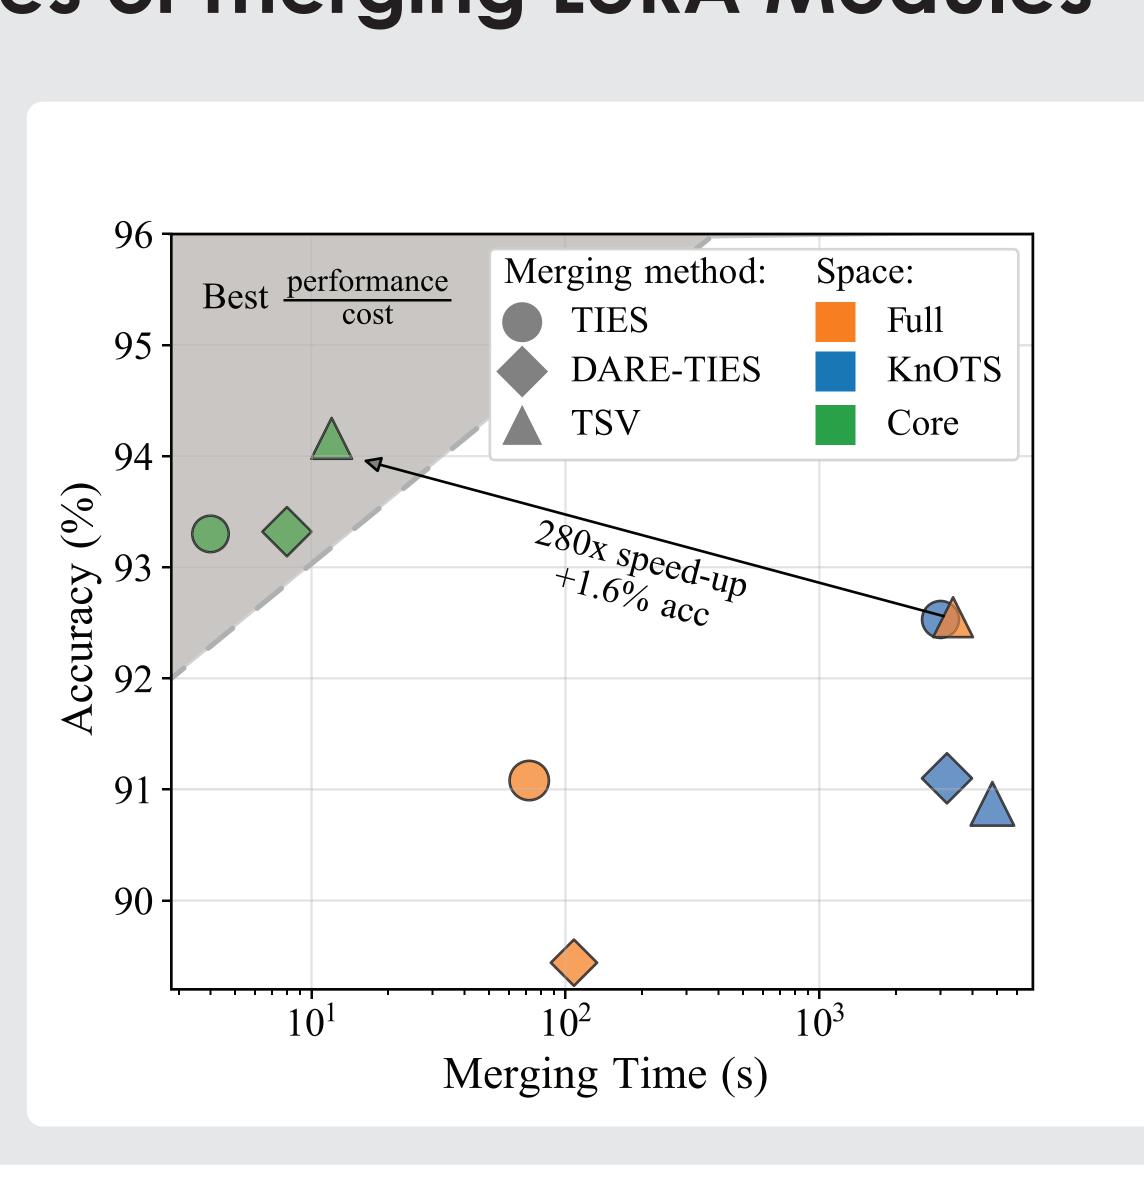
Our Core Space tackles all the challenges of merging LoRA Modules

Efficient: up to 600x speed improvent.

Accurate: consistent performance boost across vision and language benchmarks.

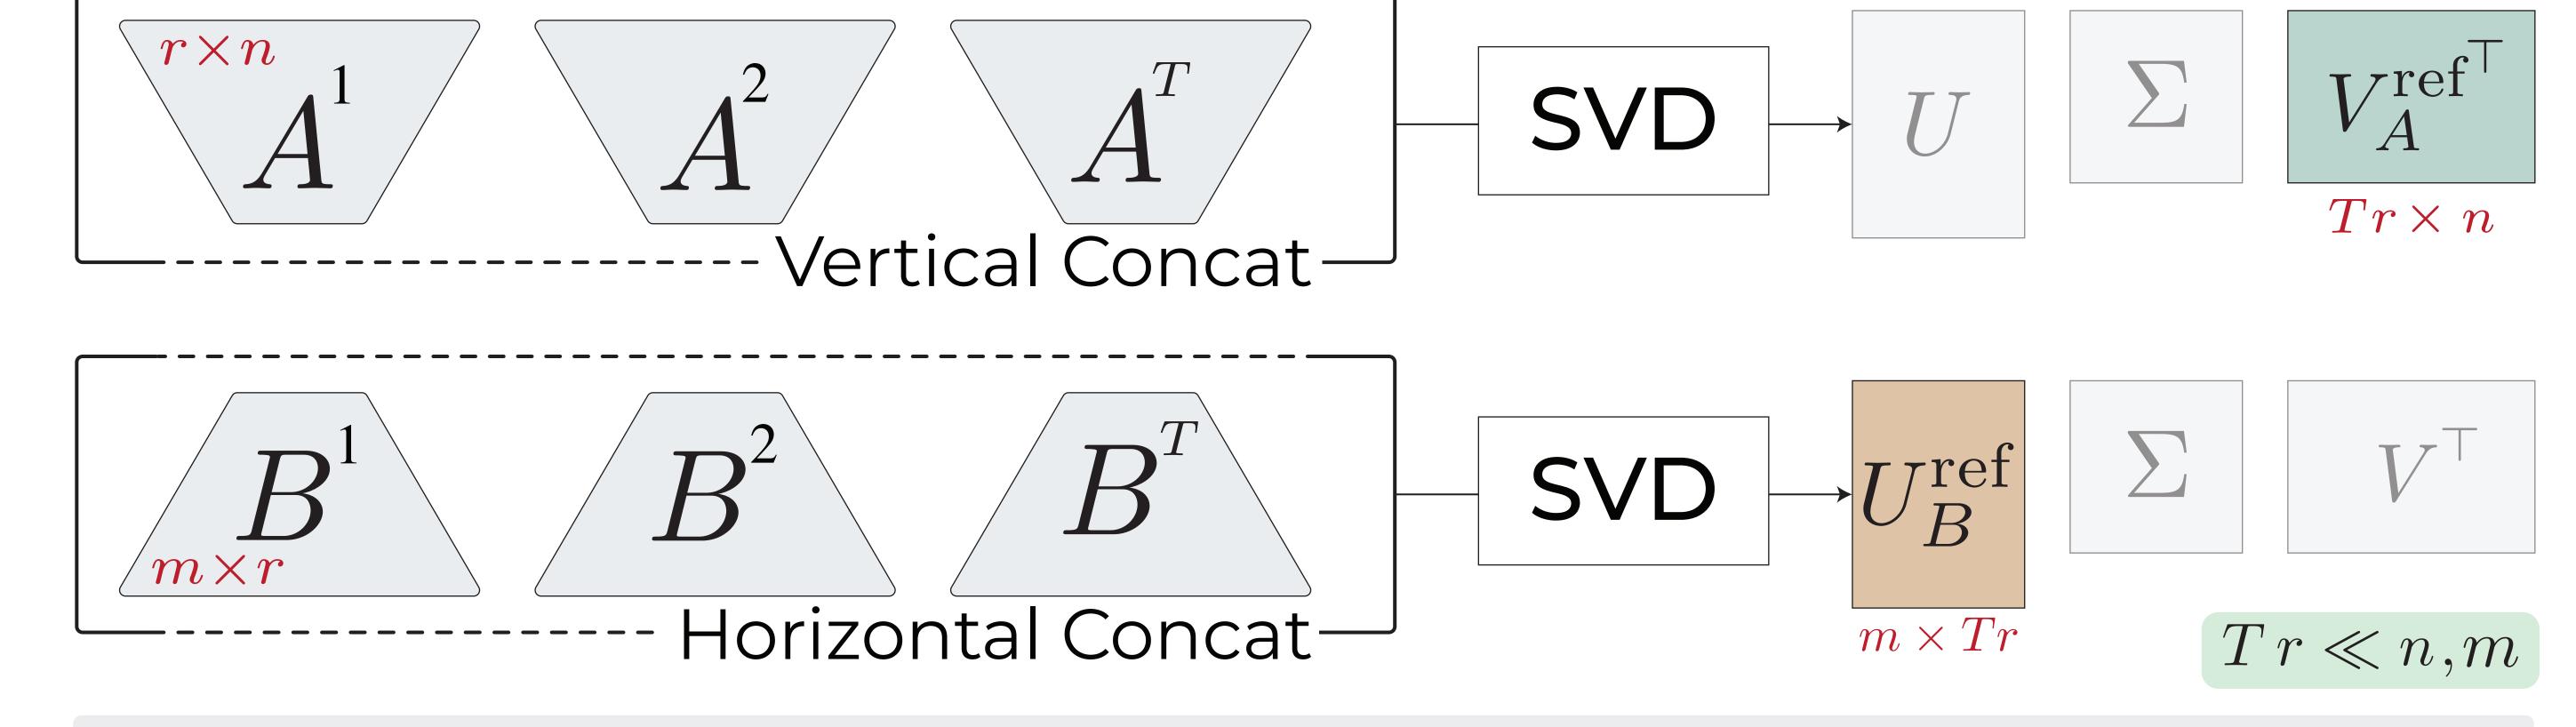
Universal: works with any merging technique (e.g., TIES, TSV, ISO, DARE).

In the Core Space we have a denser representation with no loss of information when projecting to it.

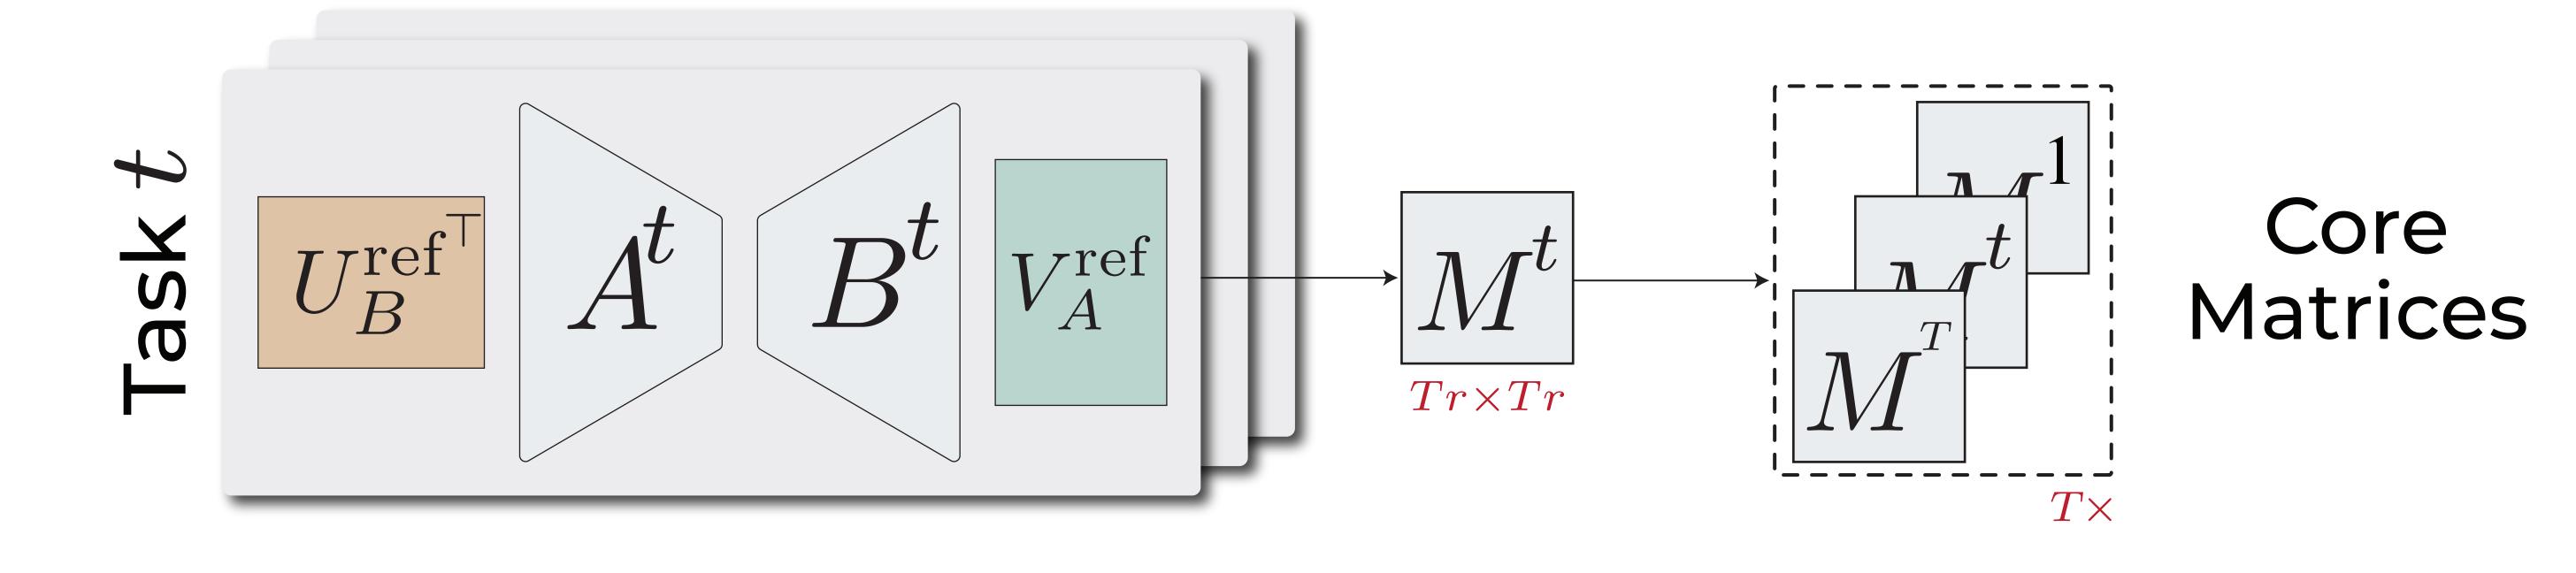


Core Space Goal: merging LoRA fine-tuned models efficiently and accurately while remaining in low-rank space.

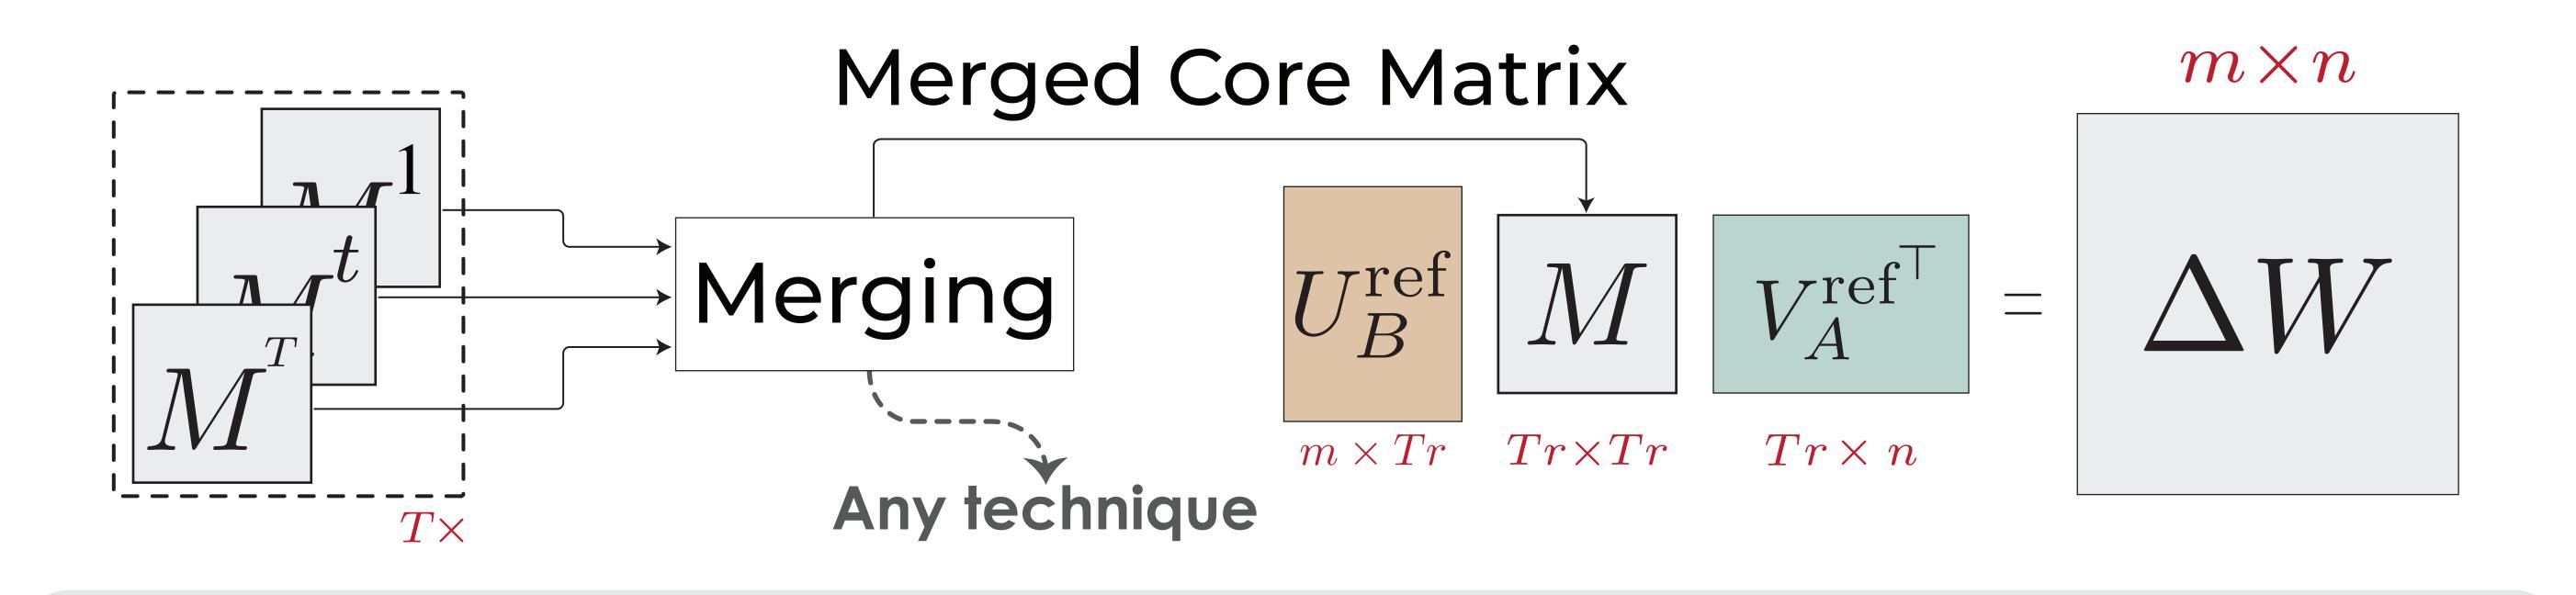
1 Compute reference bases



2 Project into Core Space



3 Merge & Reconstruct







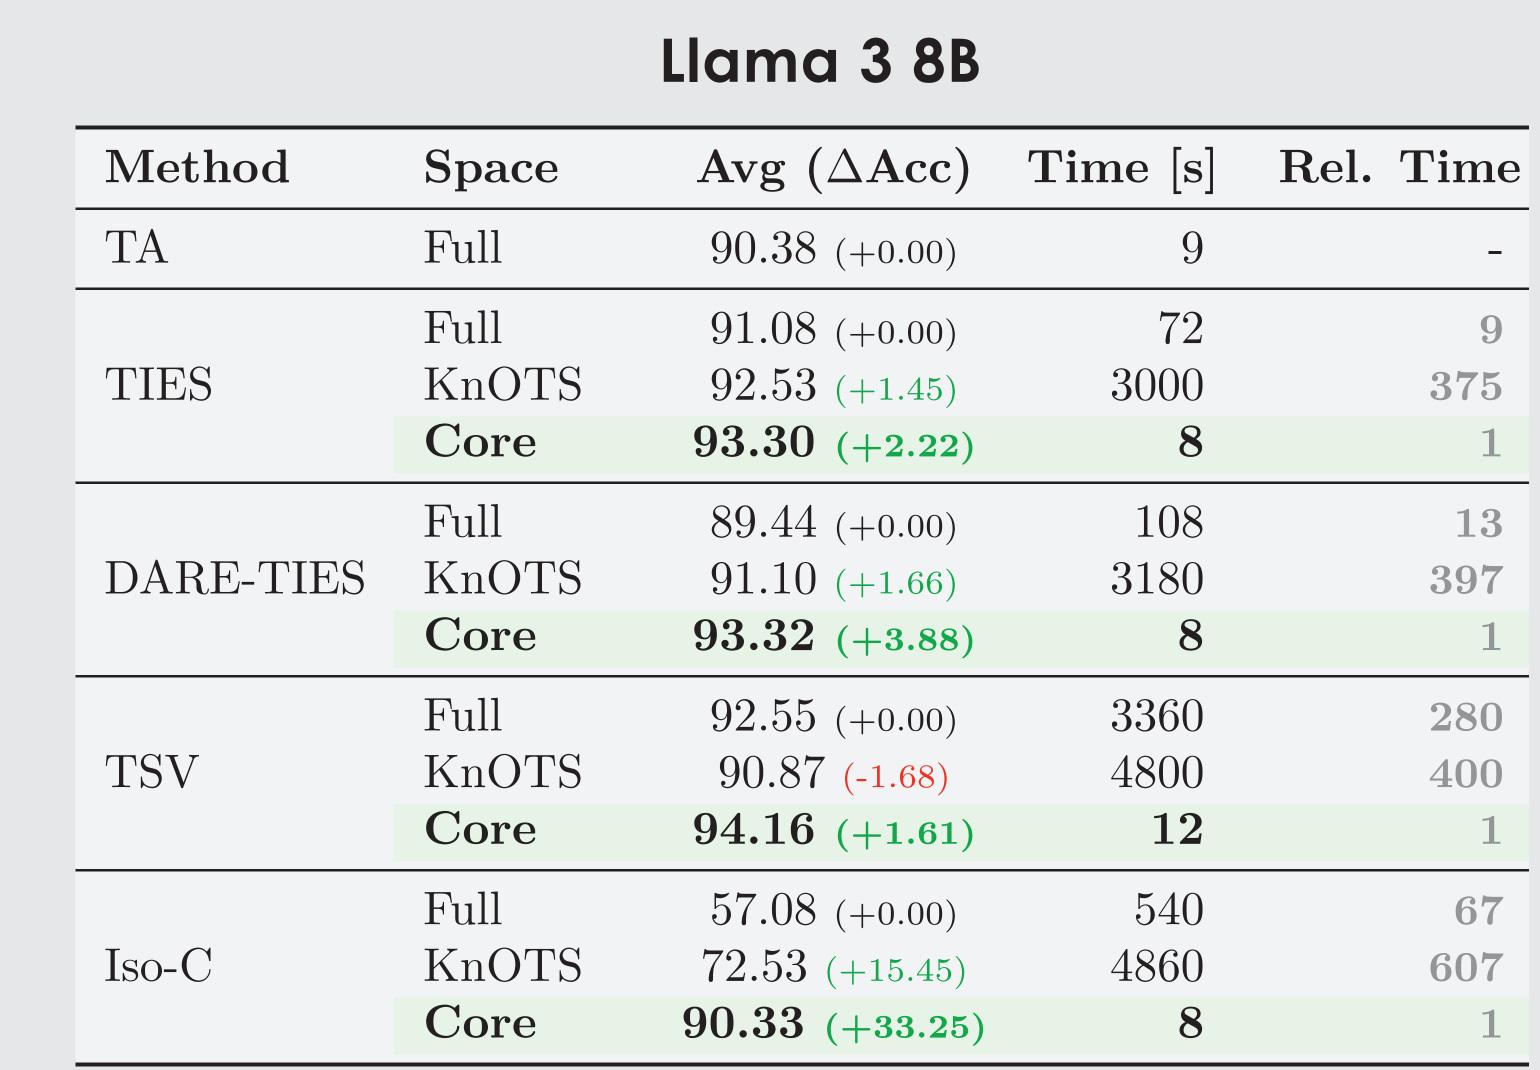




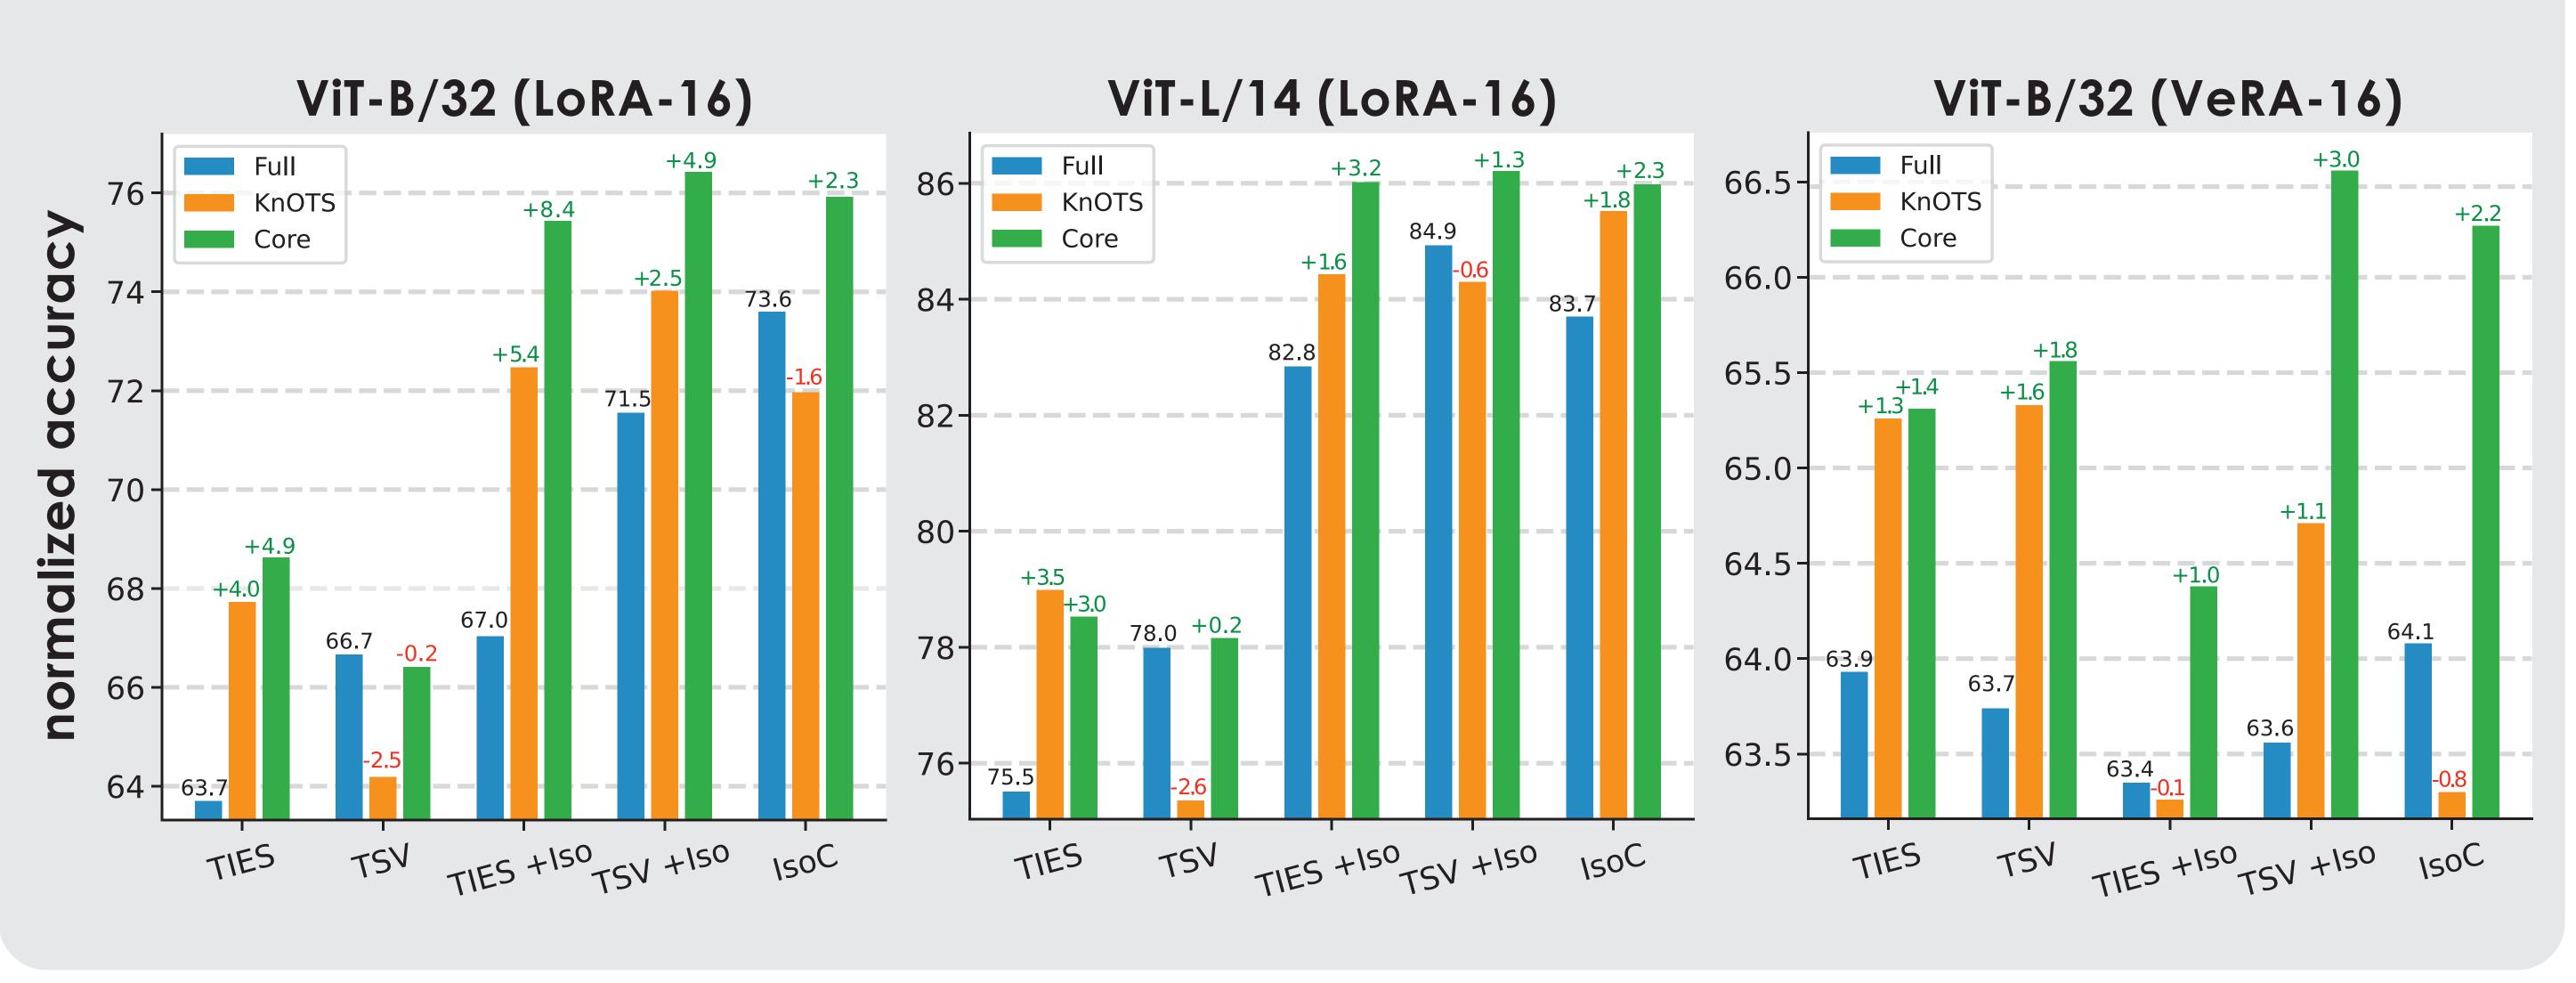


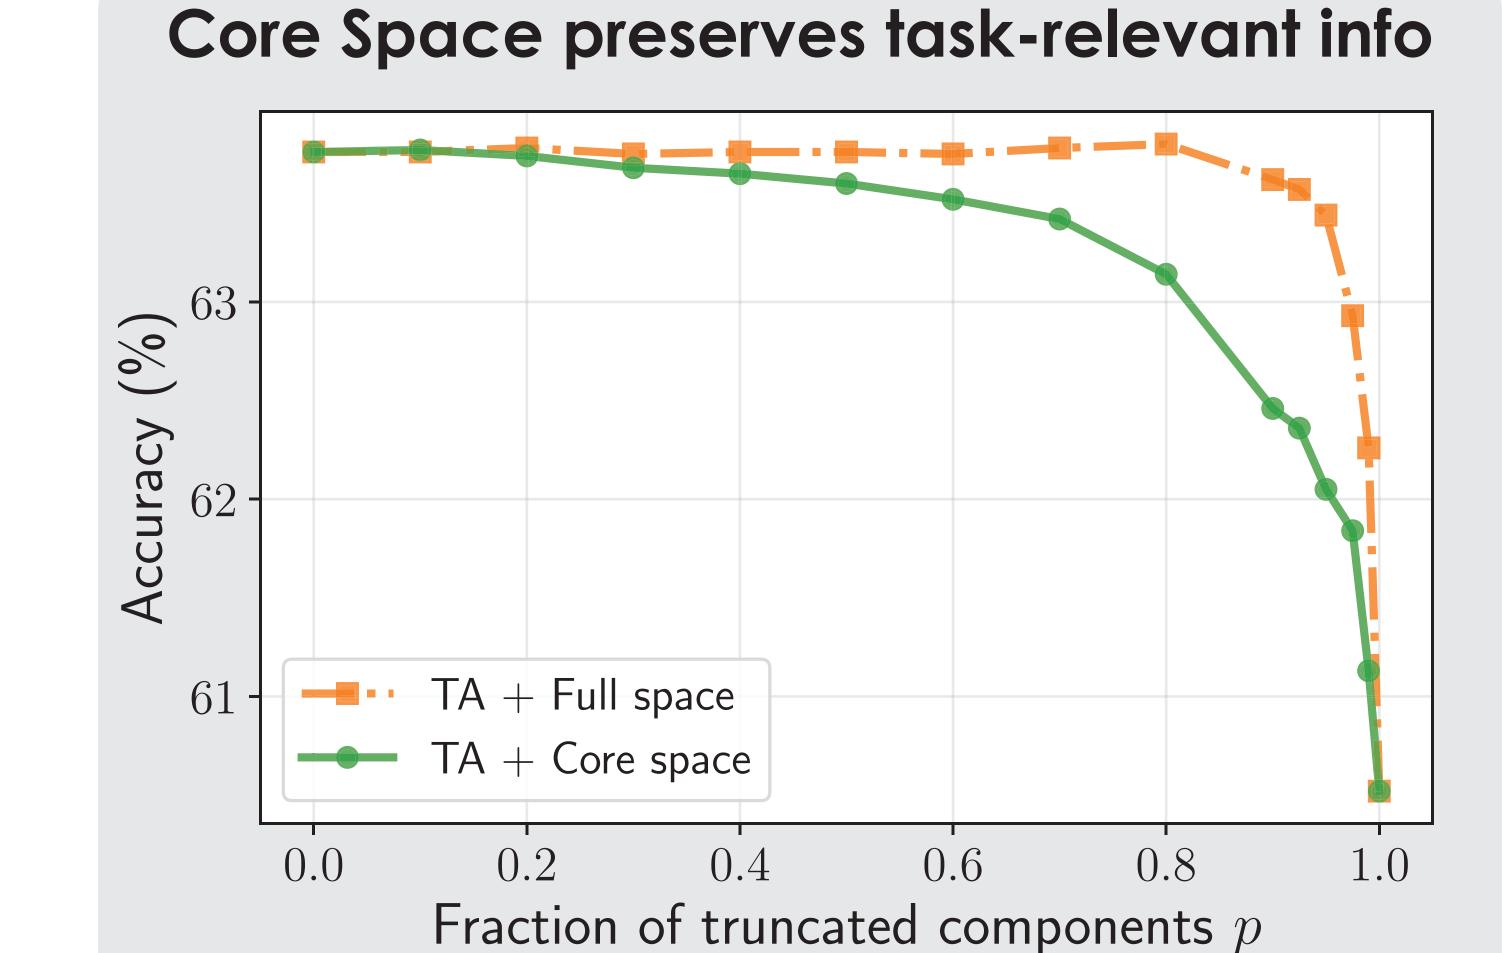
DEAS

NCBR O •



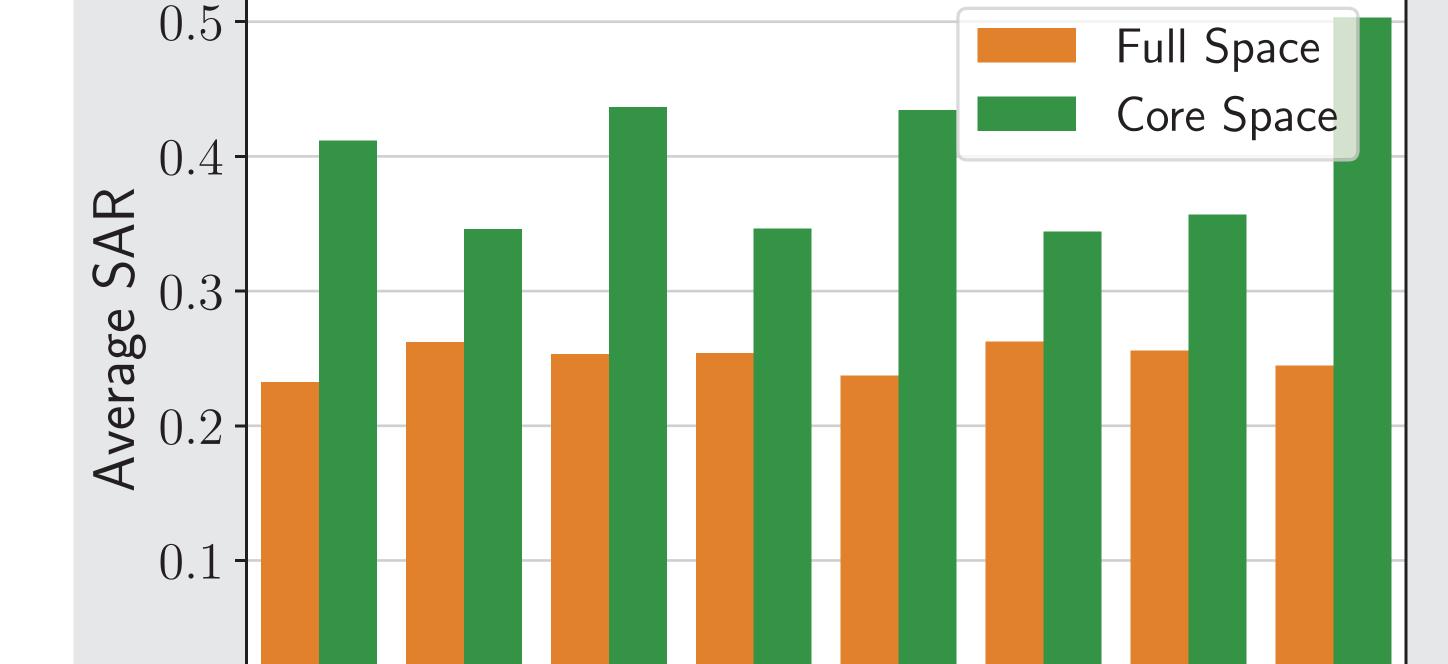
- Reliable both in Language and Vision tasks.
- Up to 600× faster for Llama 38B.
- Works reliably across ViT-B/32,
 ViT-L/14.
- Handles VeRA and different ranks.





Core Space is **compact and dense**, truncating any component reduces accuracy, unlike full space, which contains many redundant directions.

This compactness facilitates model merging as it extracts only the relevant components.



Core Space improves subspace alignment

Core Space enforces a shared basis, resulting in higher alignment and lower task interference.

Cars DTD EUROSATTSRB MNIST RESISCUN397 SVHN

Thus, task subspaces overlap more effectively, which reduces interference during merging and leads to better generalization.