TACO: Temporal Consensus Optimization for Continual Neural Mapping

Xunlan Zhou[1*], Hongrui Zhao[2*]
Negar Mehr[3]
* Equal Contribution

1: Nanjing University 2: UIUC 3: UC Berkeley

Reach Me TACO Our Lab

Robotics: Science and Systems 2026

Challenge in Neural Mapping

Current methods rely on replaying historical observations to preserve consistency and assume static scenes.

Replay-beased

Keeps outdated observations, leading to ghost artifacts.

Regularization-beased

Over-constrains updates, poor adaptation.

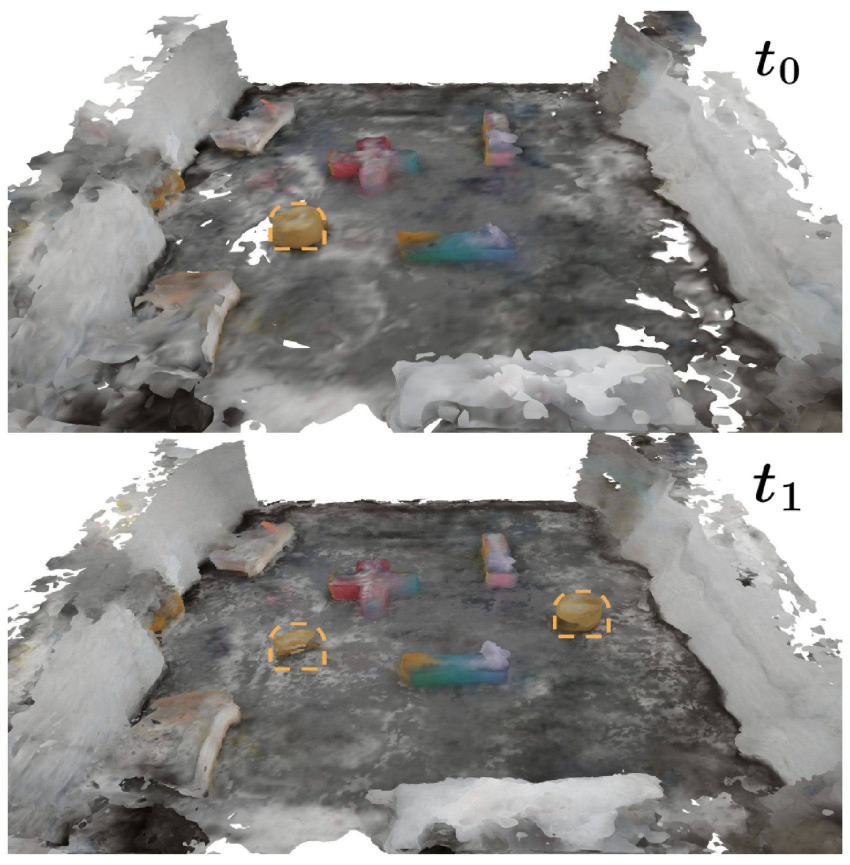

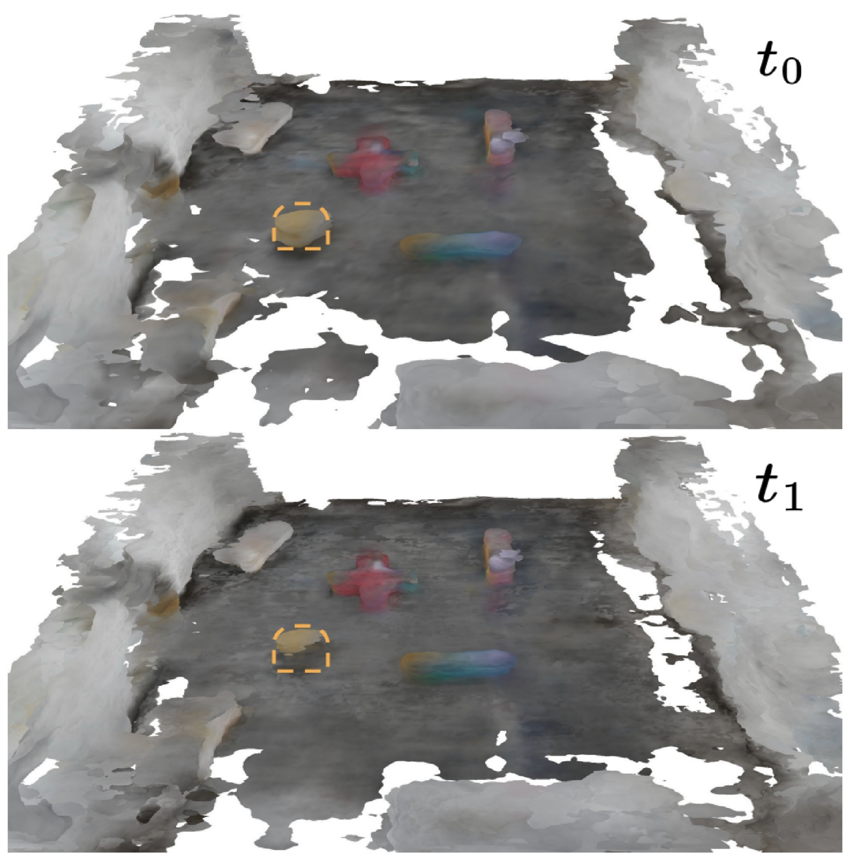

Can a neural map remember stable geometry while adapting to new changes?

TACO Framework

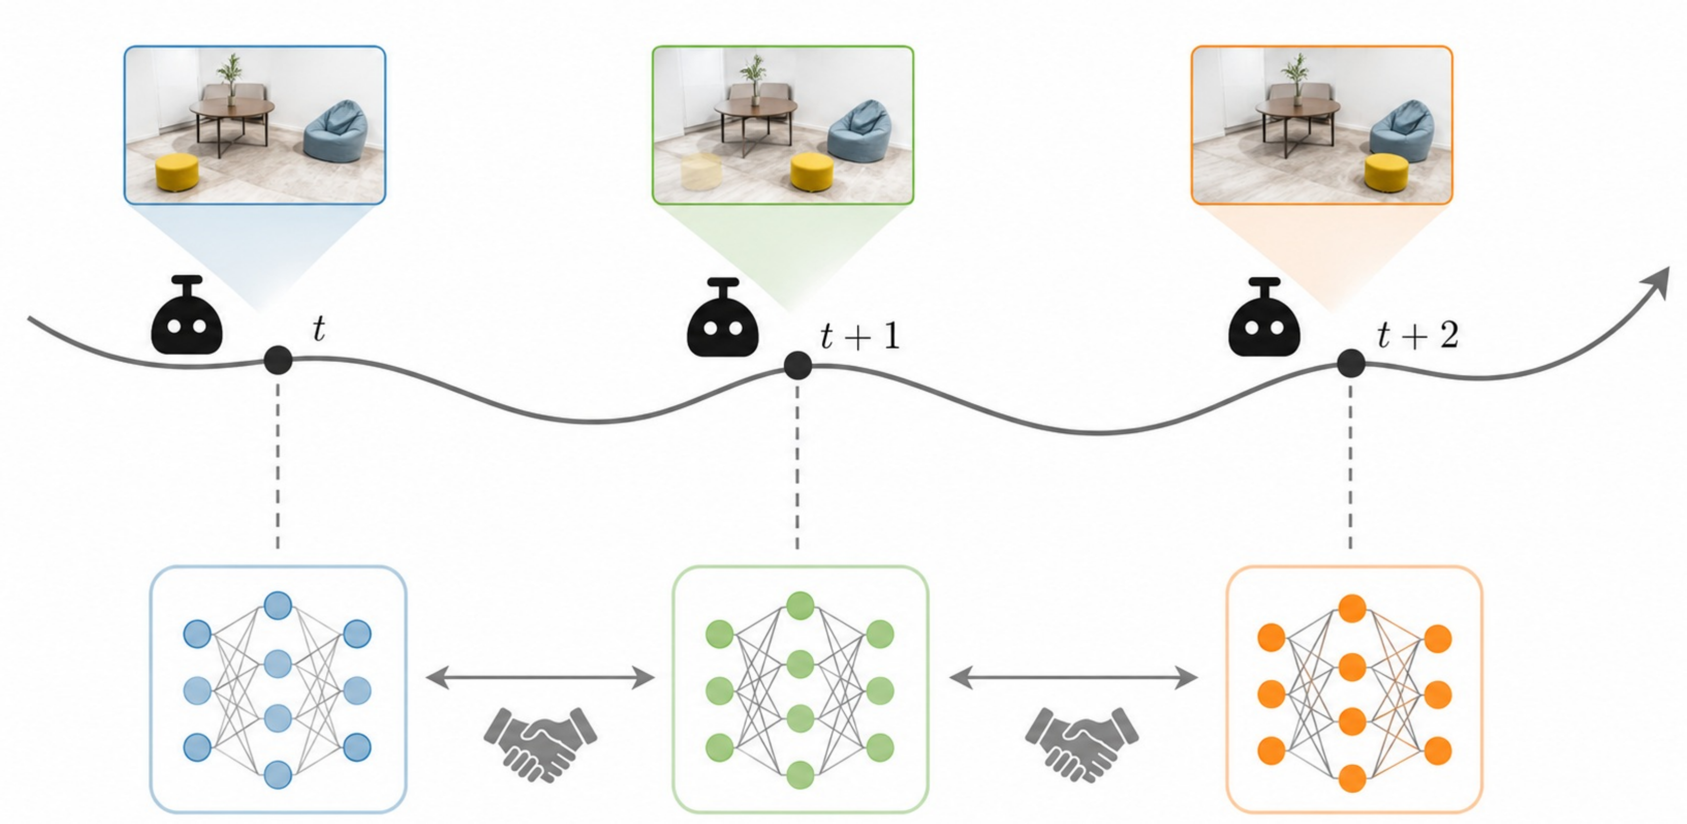

TACO treats past neural maps as historical snapshots. Instead of replaying old RGB-D frames, the current map consults its past states through an importance-weighted temporal consensus constraint.

Problem Formation

$$\min_{\Theta_t} L^{obj}(\Theta_t, R_t), \text{ s.t. } \Theta_t = z_{t,t-1}(\Theta_t).$$

We define the importance-weighted consensus target as:

$$z_{t,t-1}(\Theta_t) = \left(W_t(\Theta_t) + W_{t-1}(\Theta_t)\right)^{-1}\left(W_t(\Theta_t)\Theta_t + W_{t-1}(\Theta_t)\Theta_{t-1}\right)$$

The augmented Lagrangian:

$$\mathcal{L}^a = L^{obj}(\Theta_t, R_t) + \left(\Theta_t - z_{t,t-1}(\Theta_t)\right)^\top p + \frac{\rho}{2}\|\Theta_t - z_{t,t-1}(\Theta_t)\|_2^2.$$

Primal update:

$$\Theta_t^k = \arg\min_{\Theta_t} L^{obj}(\Theta_t, R_t) + \left(\Theta_t - z_{t,t-1}^{k-1}(\Theta_t^{k-1})\right)^\top p^{k-1} + \frac{\rho}{2}\|\Theta_t - z_{t,t-1}^{k-1}(\Theta_t^{k-1})\|_2^2.$$

Dual update: $p^k = p^{k-1} + \rho(\Theta_t^k - z_{t,t-1}^{k-1}(\Theta_t^{k-1})).$

The parameter importance:

$$\mathcal{L}_s^{\text{proxy}} = \frac{1}{|\mathcal{S}_s|}\sum_{r\in\mathcal{S}_s}\left(\|\hat{c}_{s,r}\|_2^2 + \|\hat{d}_{s,r}\|_2^2\right), \quad u_t^k(\Theta_t^k) = \sum_{s=1}^{k}\left|\frac{\partial \mathcal{L}_s^{\text{proxy}}}{\partial \Theta_t^s}\right|$$

Experiments & Results

Static preservation TACO matches replay-based Co-SLAM and outperforms continual learning baselines.

Dynamic adaptation TACO updates changed regions without ghosting or tearing.

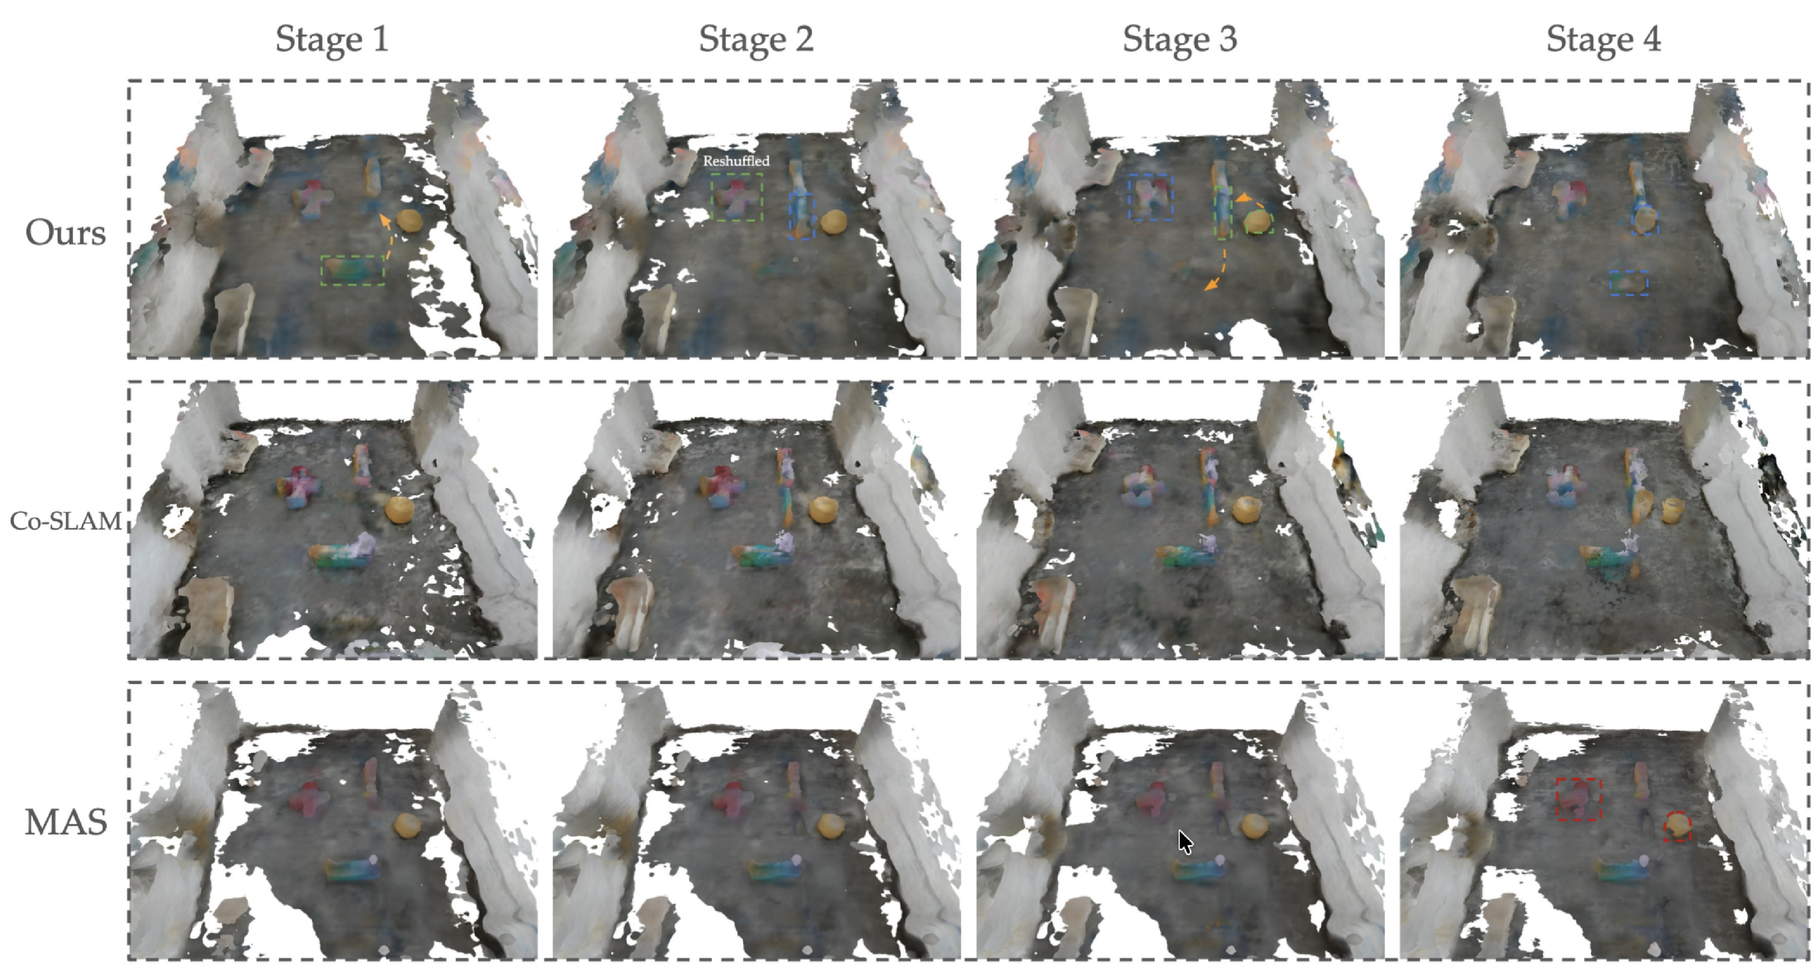

Memory efficiency

Method	CPU RSS (MB) ↓
Co-SLAM	4165.1
KR (Replay-based)	1912.4
UNIKD (Knowledge Distillation)	1945.1
Online MAS (Regularization)	1800.6
Ours	1864.5

TACO enables neural maps to remember stable geometry while adapting to new scene changes!