

Missing Modality Robustness in Semi-Supervised Multi-Modal Semantic Segmentation



Harsh Maheshwari



Yen-Cheng Liu



Zsolt Kira

Challenges with Multi-Modal Semantic Segmentation

Challenges with Multi-Modal **Semantic Segmentation**

Getting segmentation labels is laborious and costly

Challenges with **Multi-Modal Semantic Segmentation**

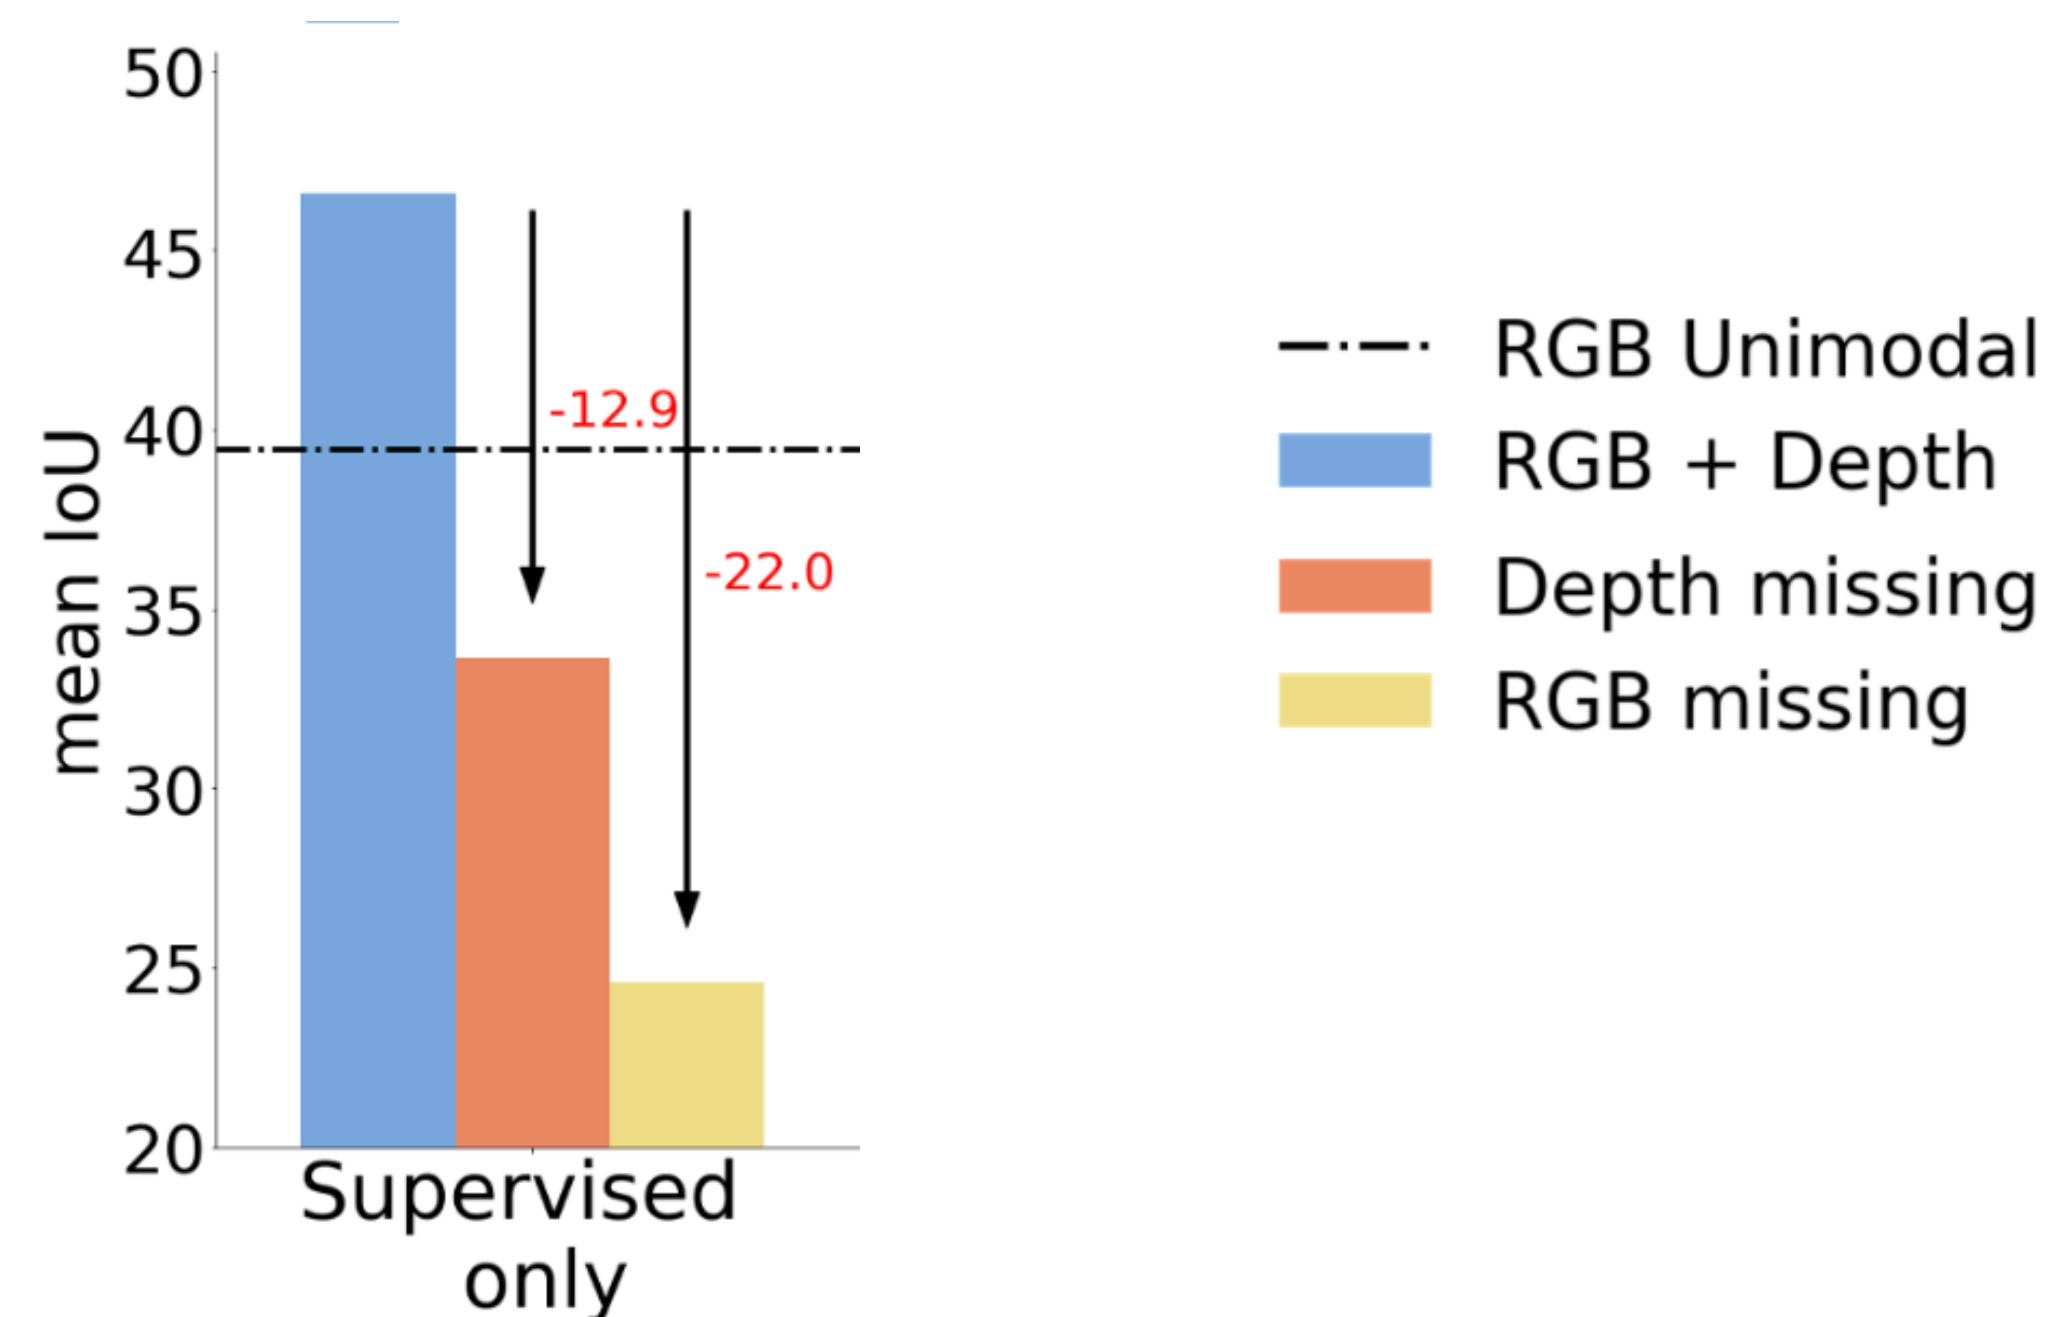
Getting segmentation labels is laborious and costly

Guaranteeing the presence of all modalities is difficult

Challenges with Multi-Modal Semantic Segmentation

Getting segmentation labels is laborious and costly

Guaranteeing the presence of all modalities is difficult



Solutions to these challenges

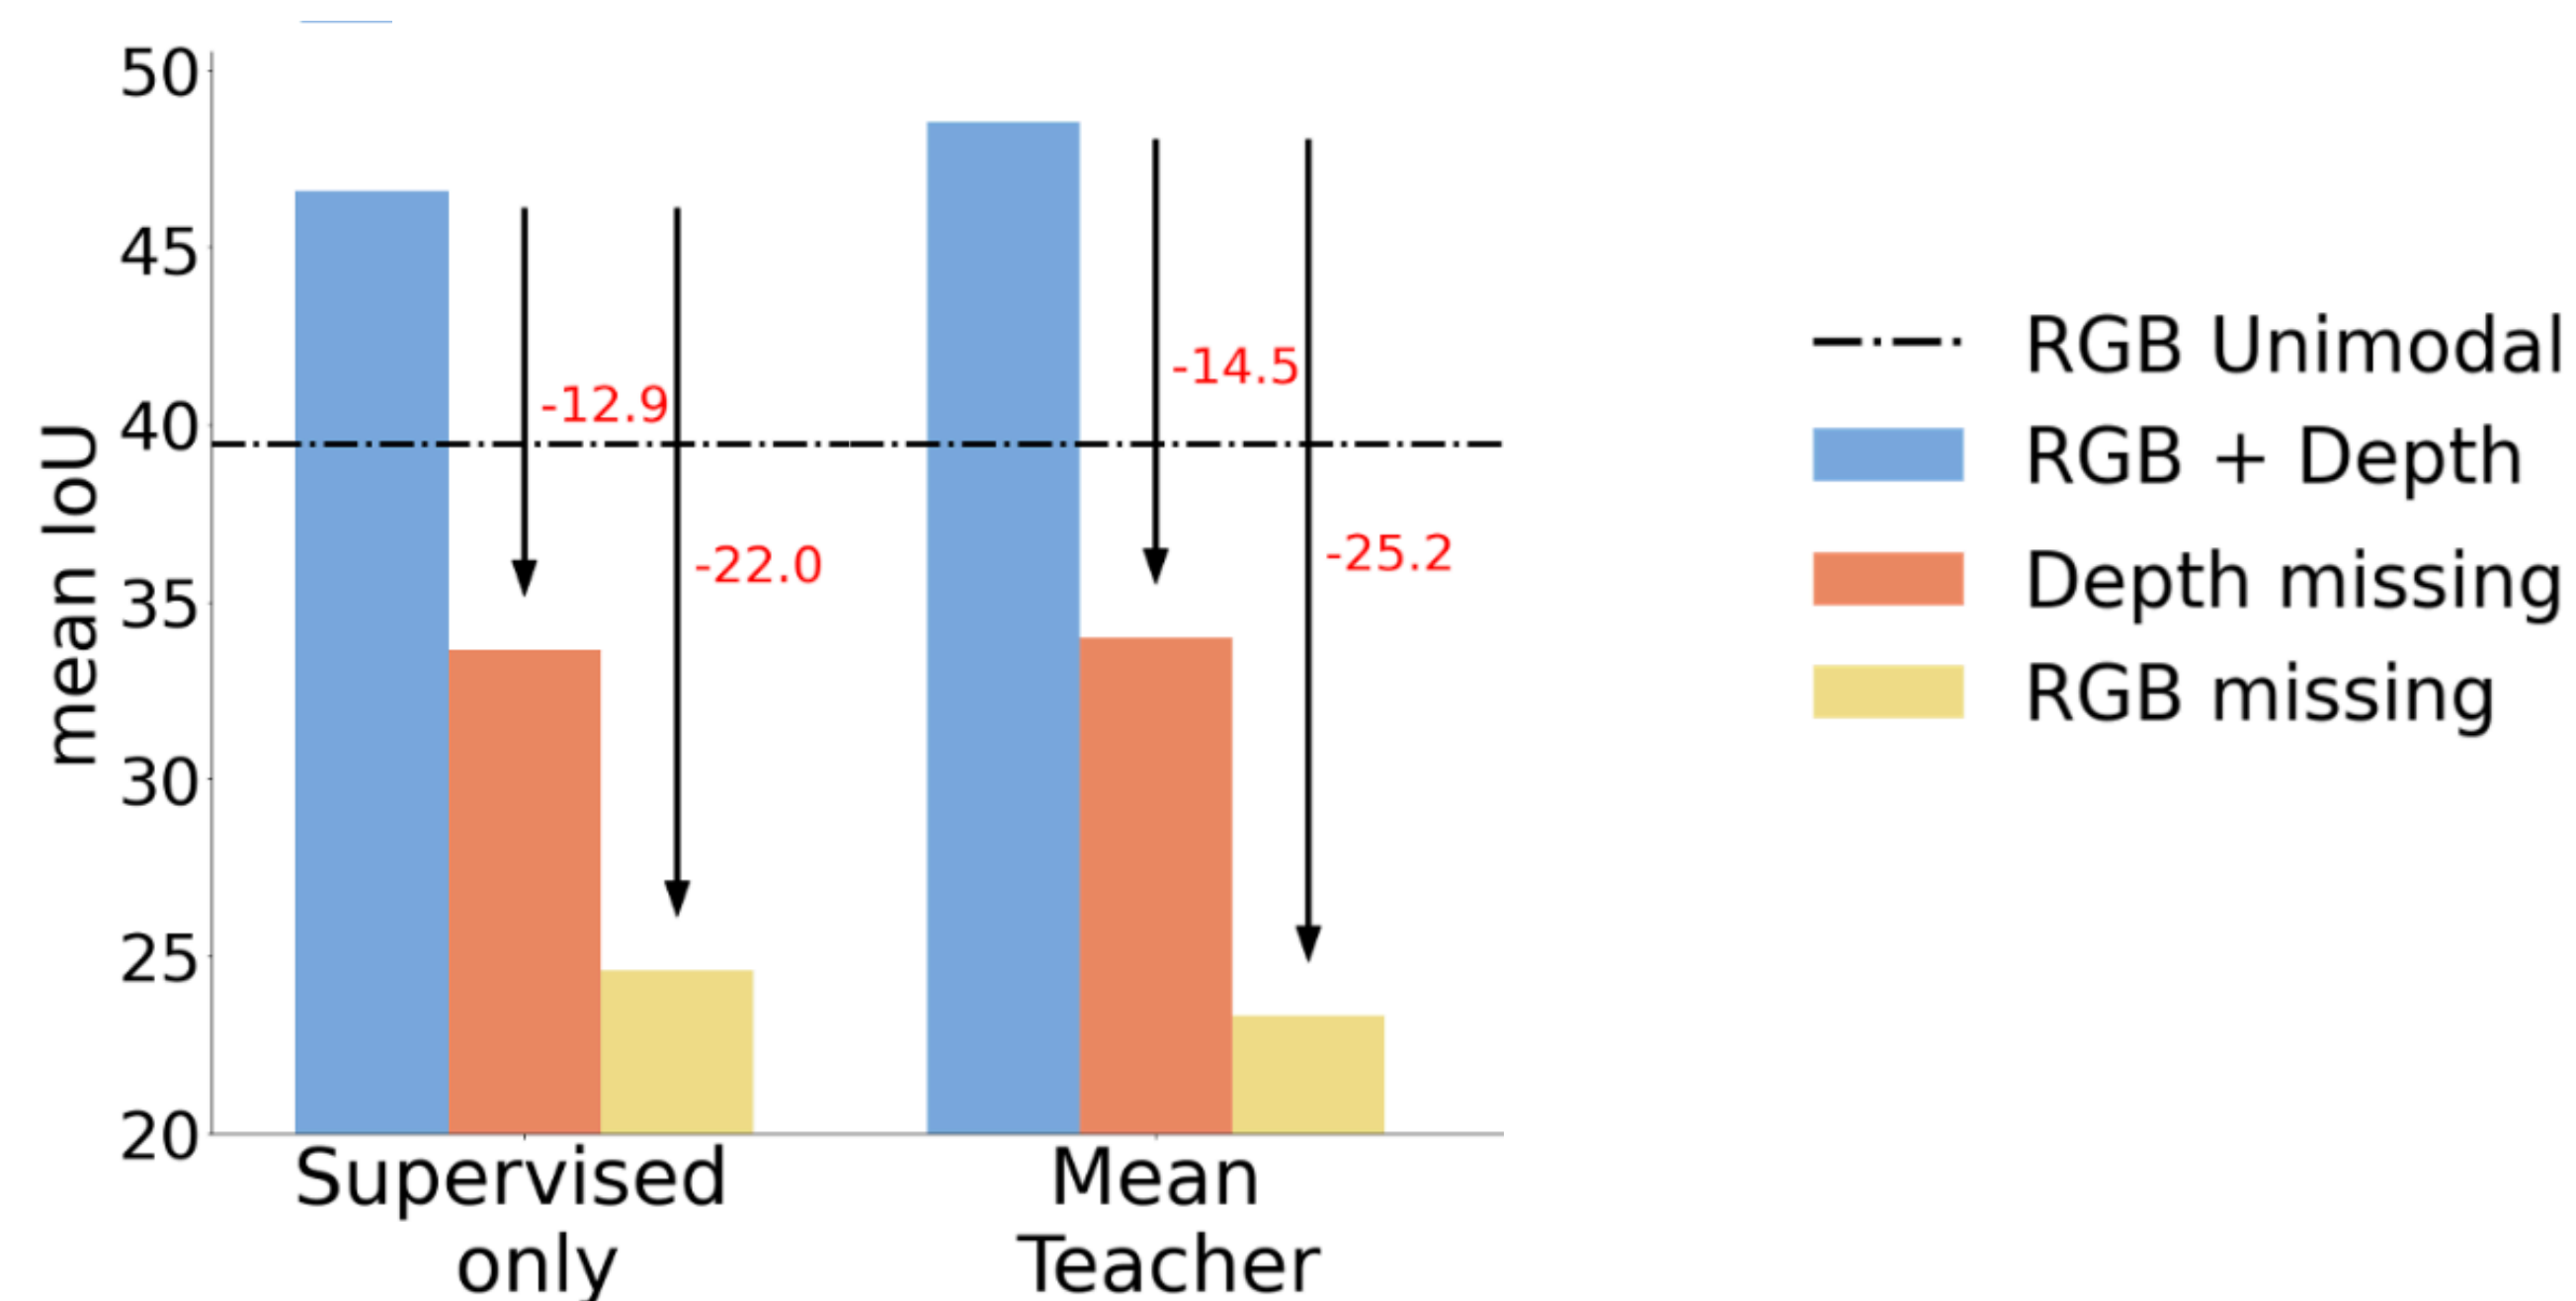
Getting segmentation labels is laborious and costly **Semi-Supervised**

Guaranteeing the presence of all modalities is difficult

Challenges with Multi-Modal Semantic Segmentation

Getting segmentation labels is laborious and costly Semi-Supervised

Guaranteeing the presence of all modalities is difficult



Challenges with Multi-Modal Semantic Segmentation

Getting segmentation labels is laborious and costly

Semi-Supervised

Guaranteeing the presence of all modalities is difficult

Robustness to
Missing-Modalities

Linear Fusion

For low-label setting, a multi-modal model with
simple fusion mechanism

M3L: Multi-modal teacher for Masked Modality Learning

A **Semi-Supervised** framework to increase
robustness to **Missing-Modalities**

Linear Fusion

For low-label setting, a multi-modal model with
simple fusion mechanism

M3L: Multi-modal teacher for Masked Modality Learning

A Semi-Supervised framework to increase
robustness to Missing-Modalities

Linear Fusion: Multi-Modal Semantic Segmentation Model

For low-label setting, a multi-modal model with

simple fusion mechanism

and

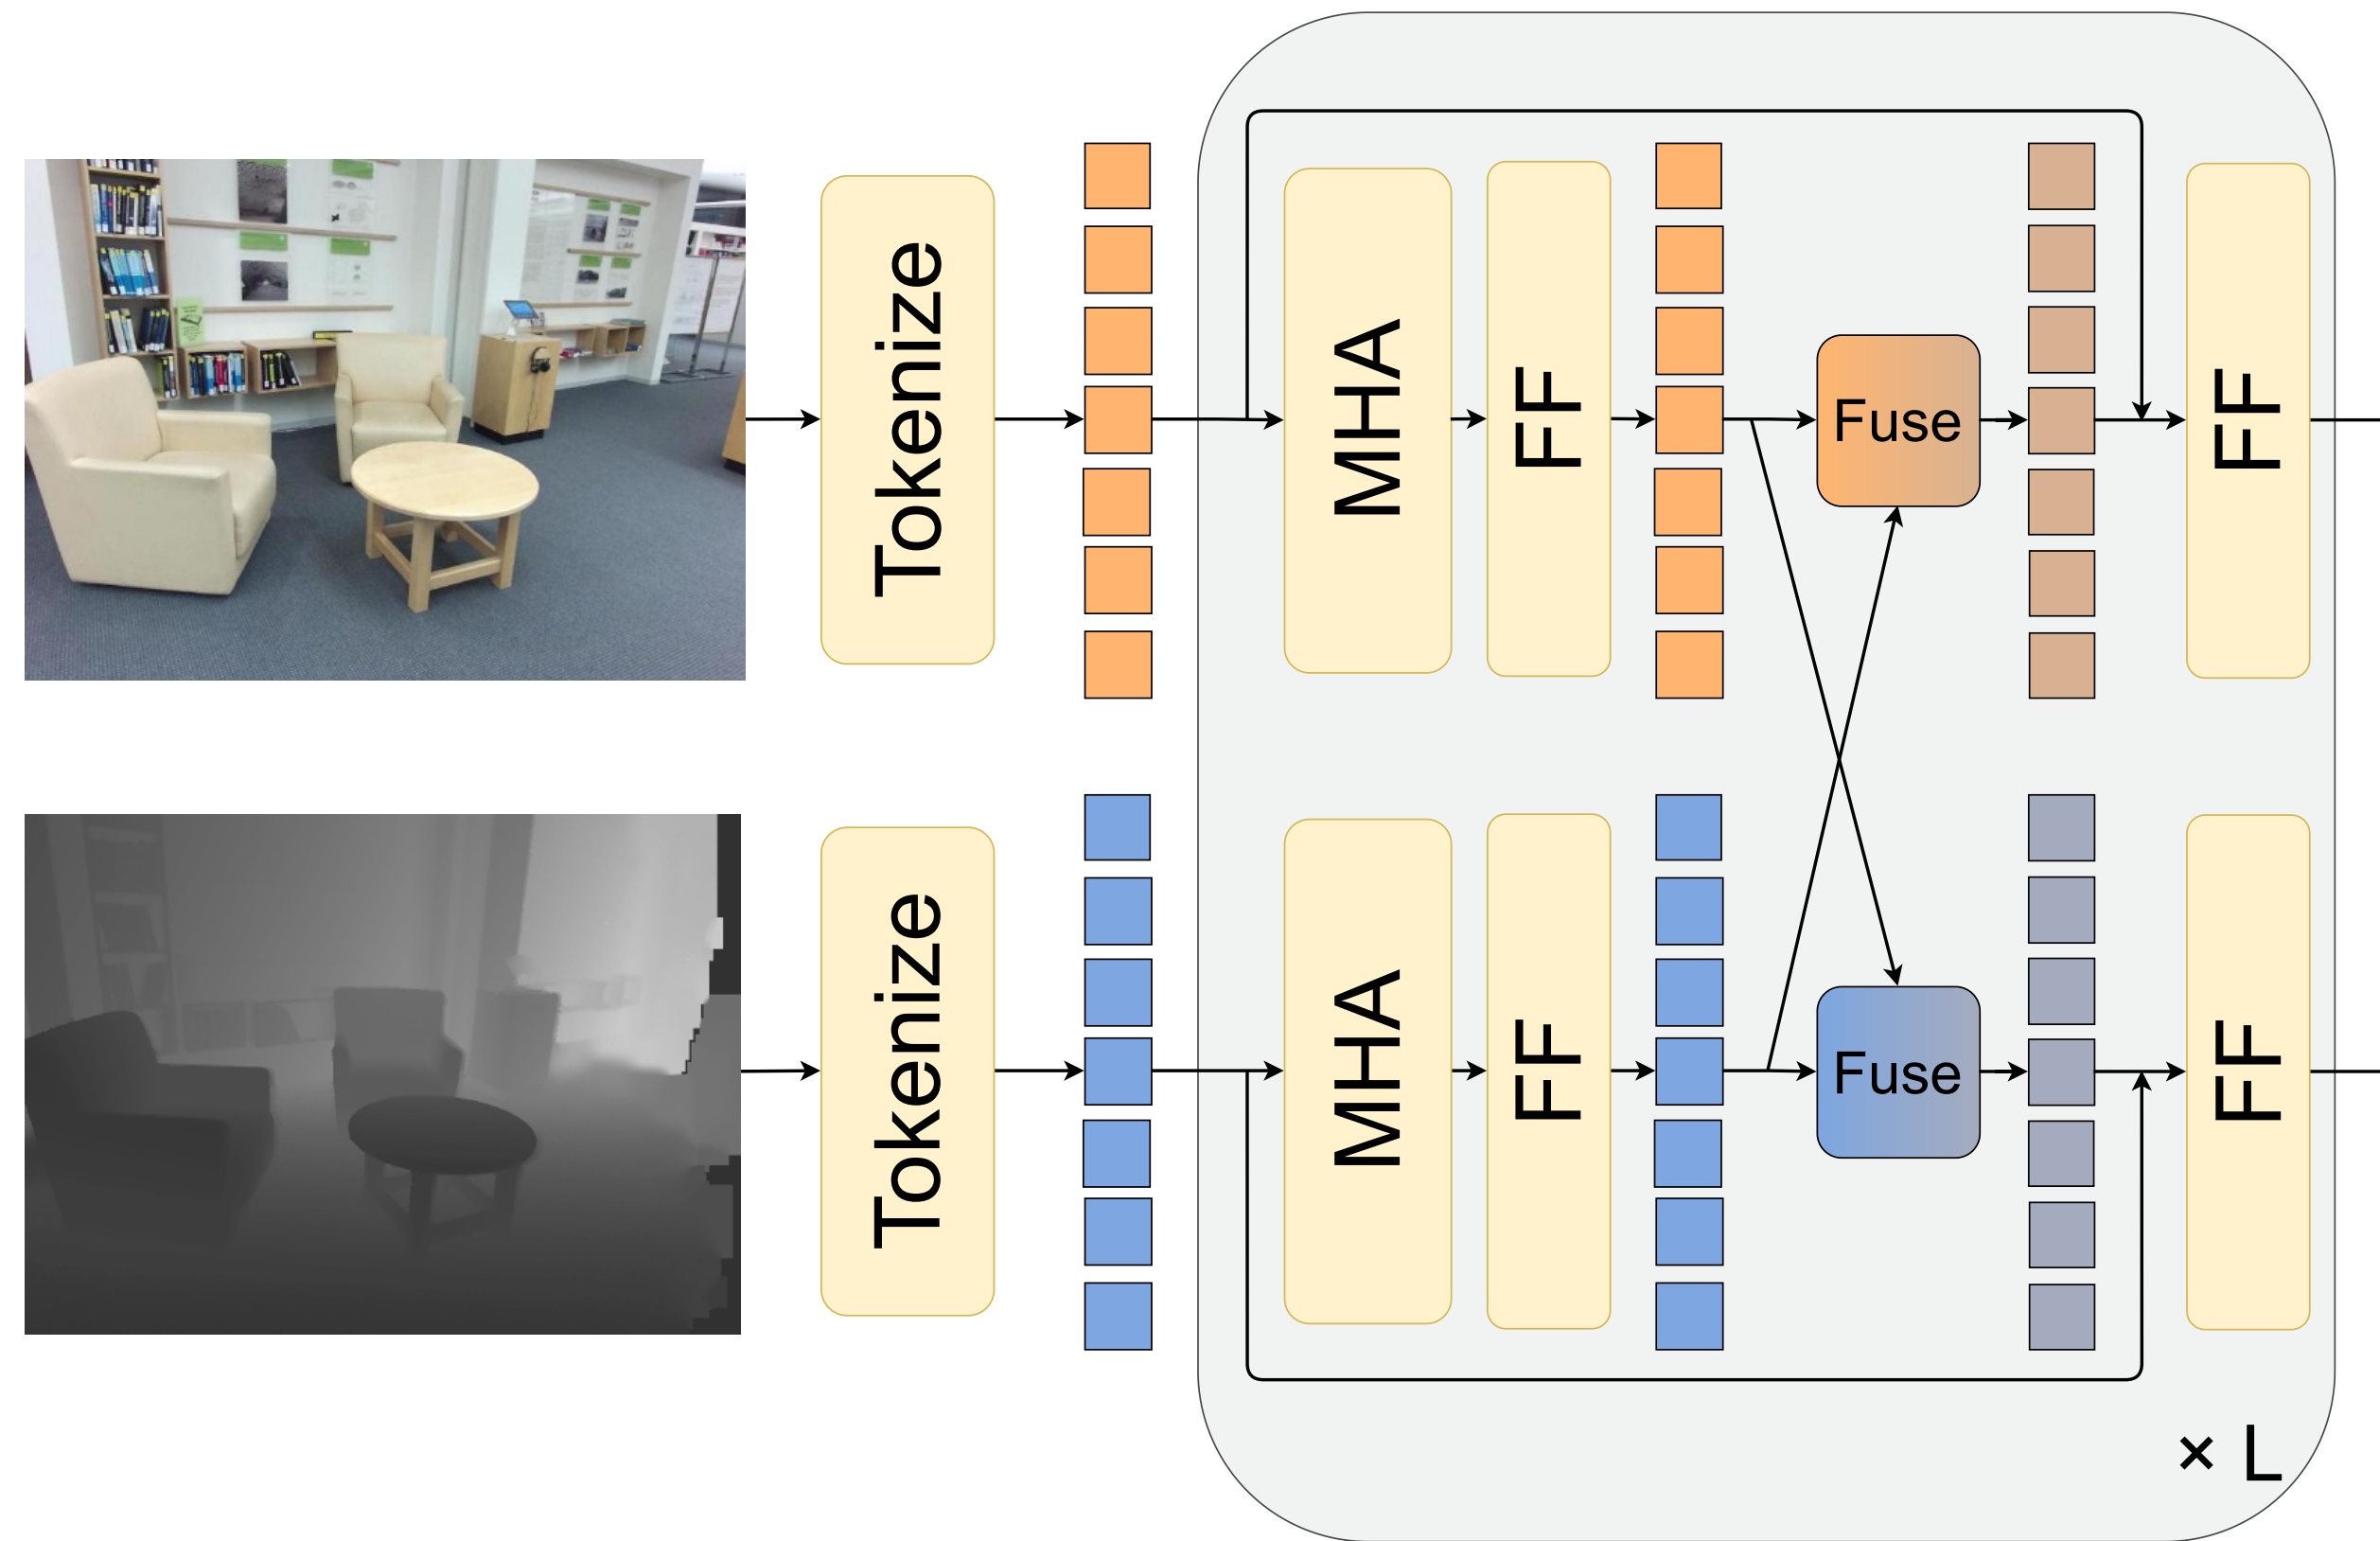
no extra trainable parameters

is desirable

Linear Fusion: Multi-Modal Semantic Segmentation Model

simple fusion mechanism

no extra trainable parameters

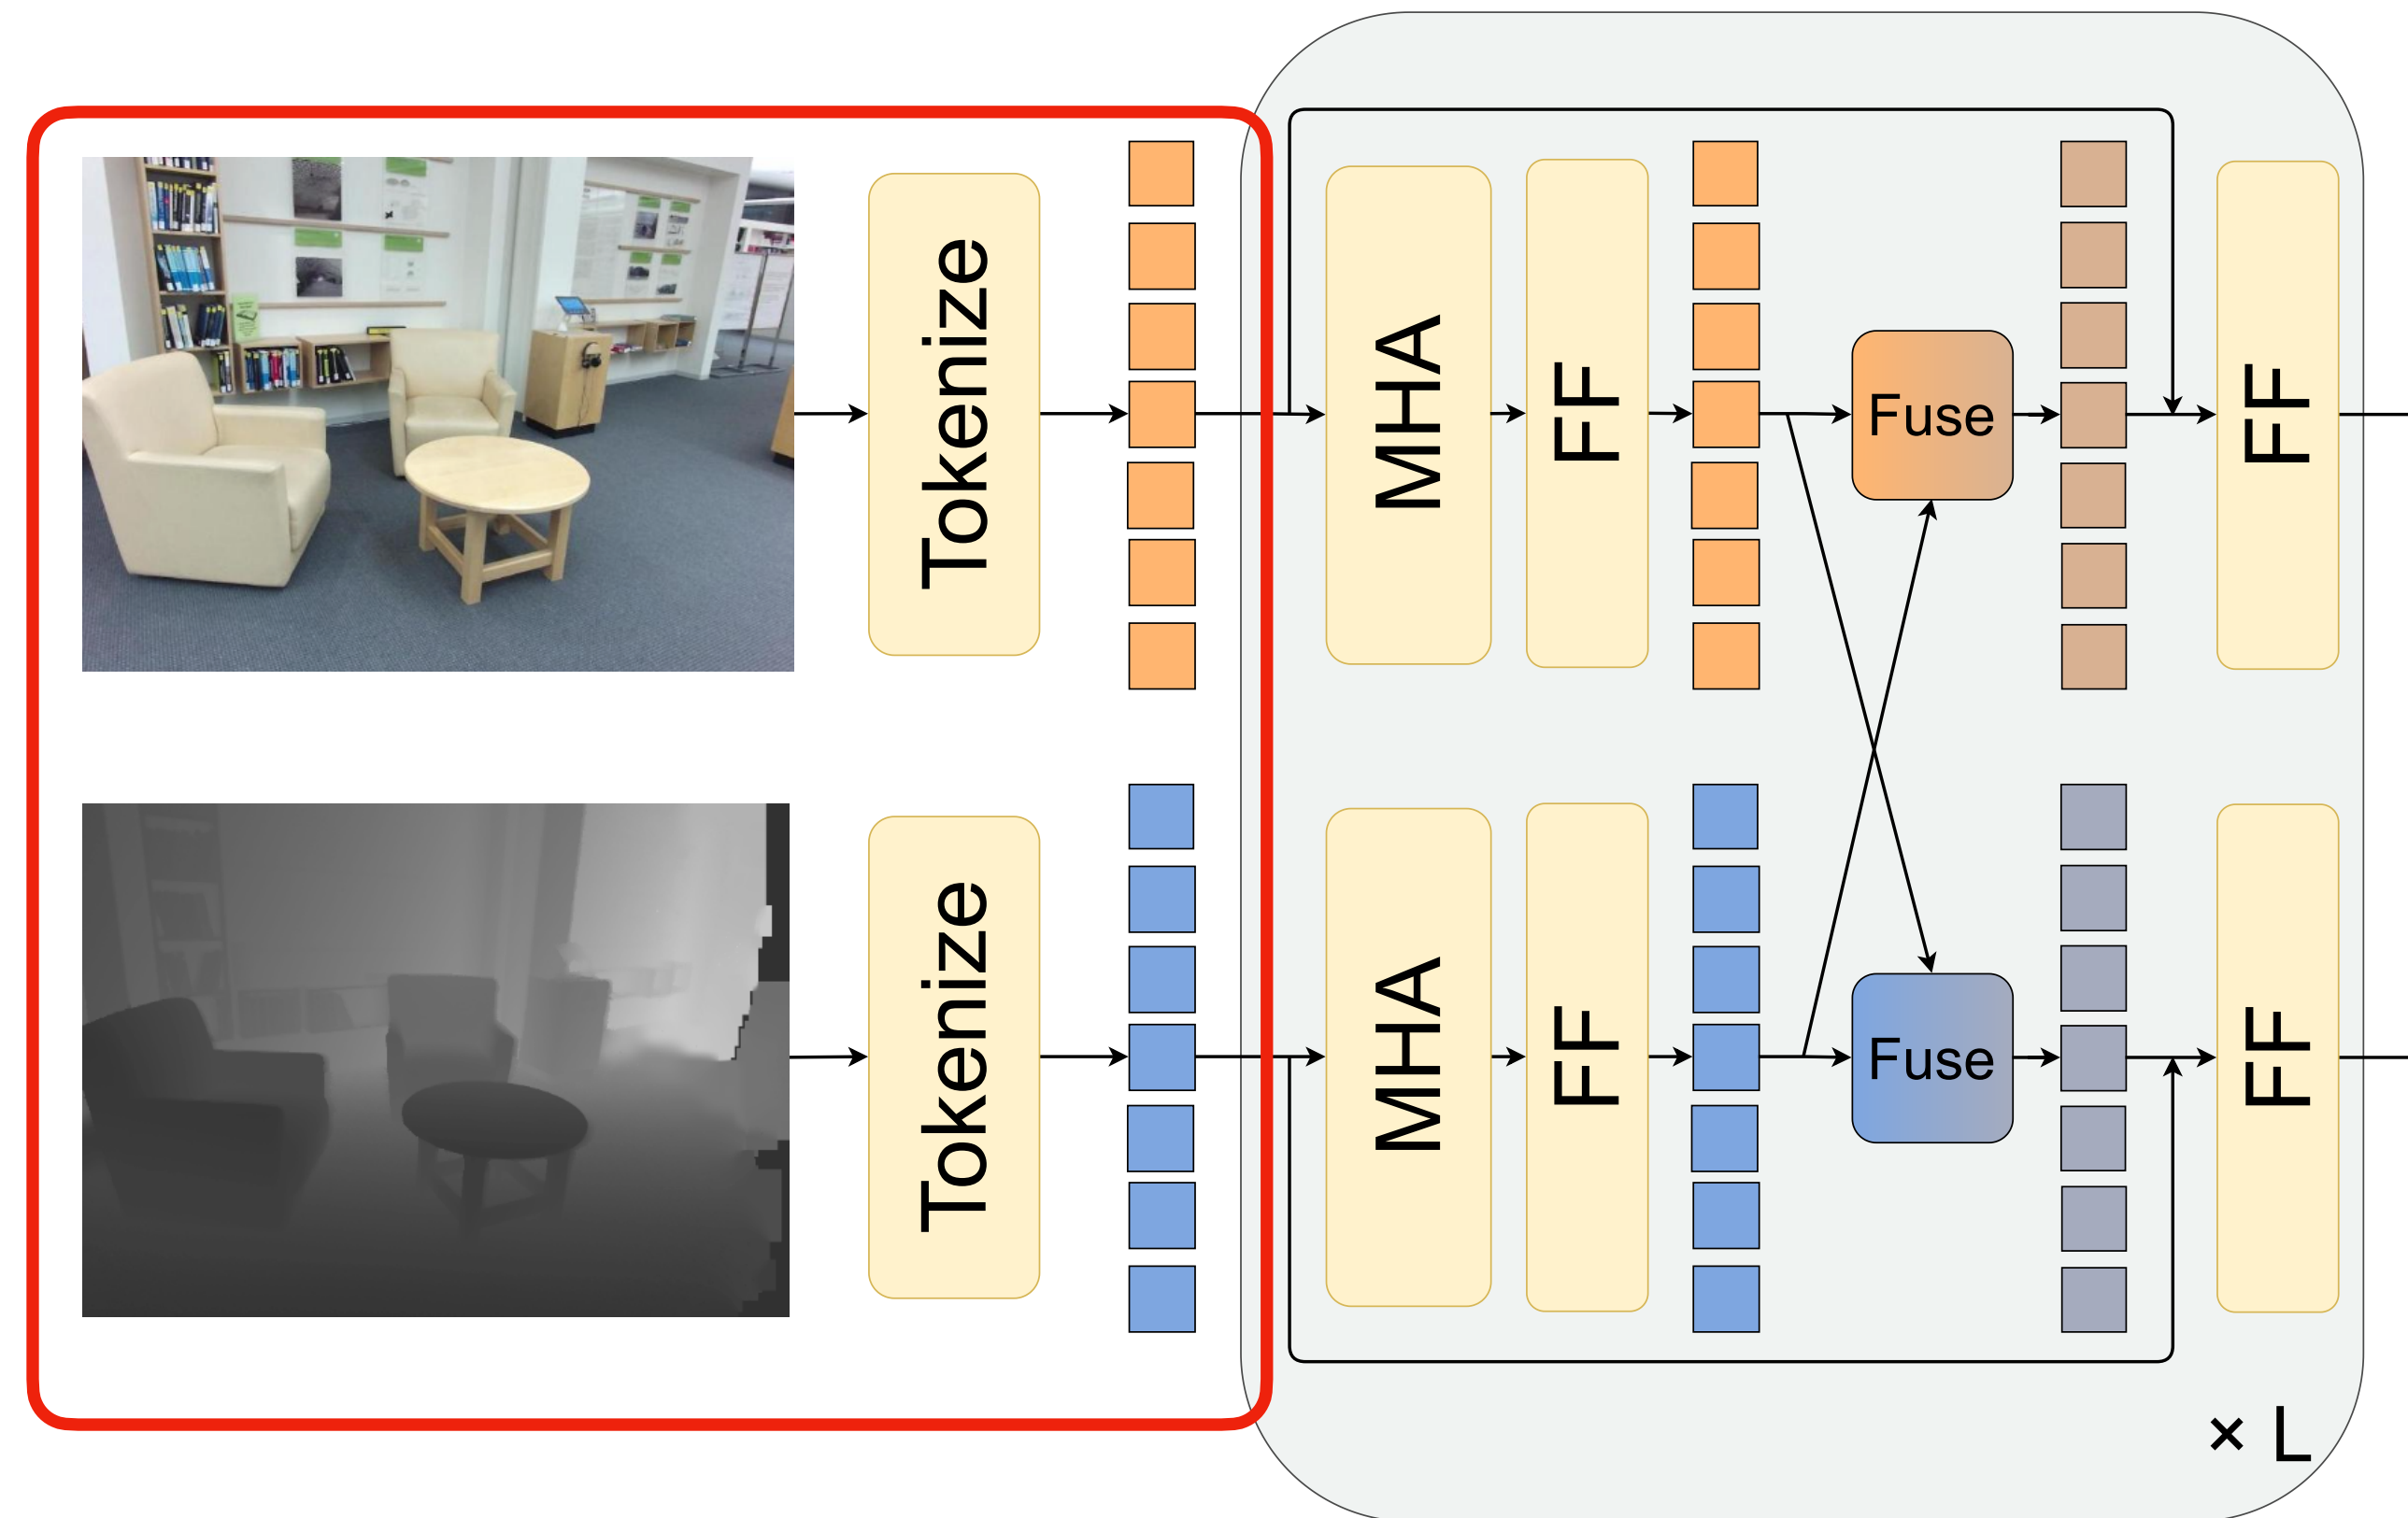


Linear Fusion

Linear Fusion: Multi-Modal Semantic Segmentation Model

simple fusion mechanism

no extra trainable parameters

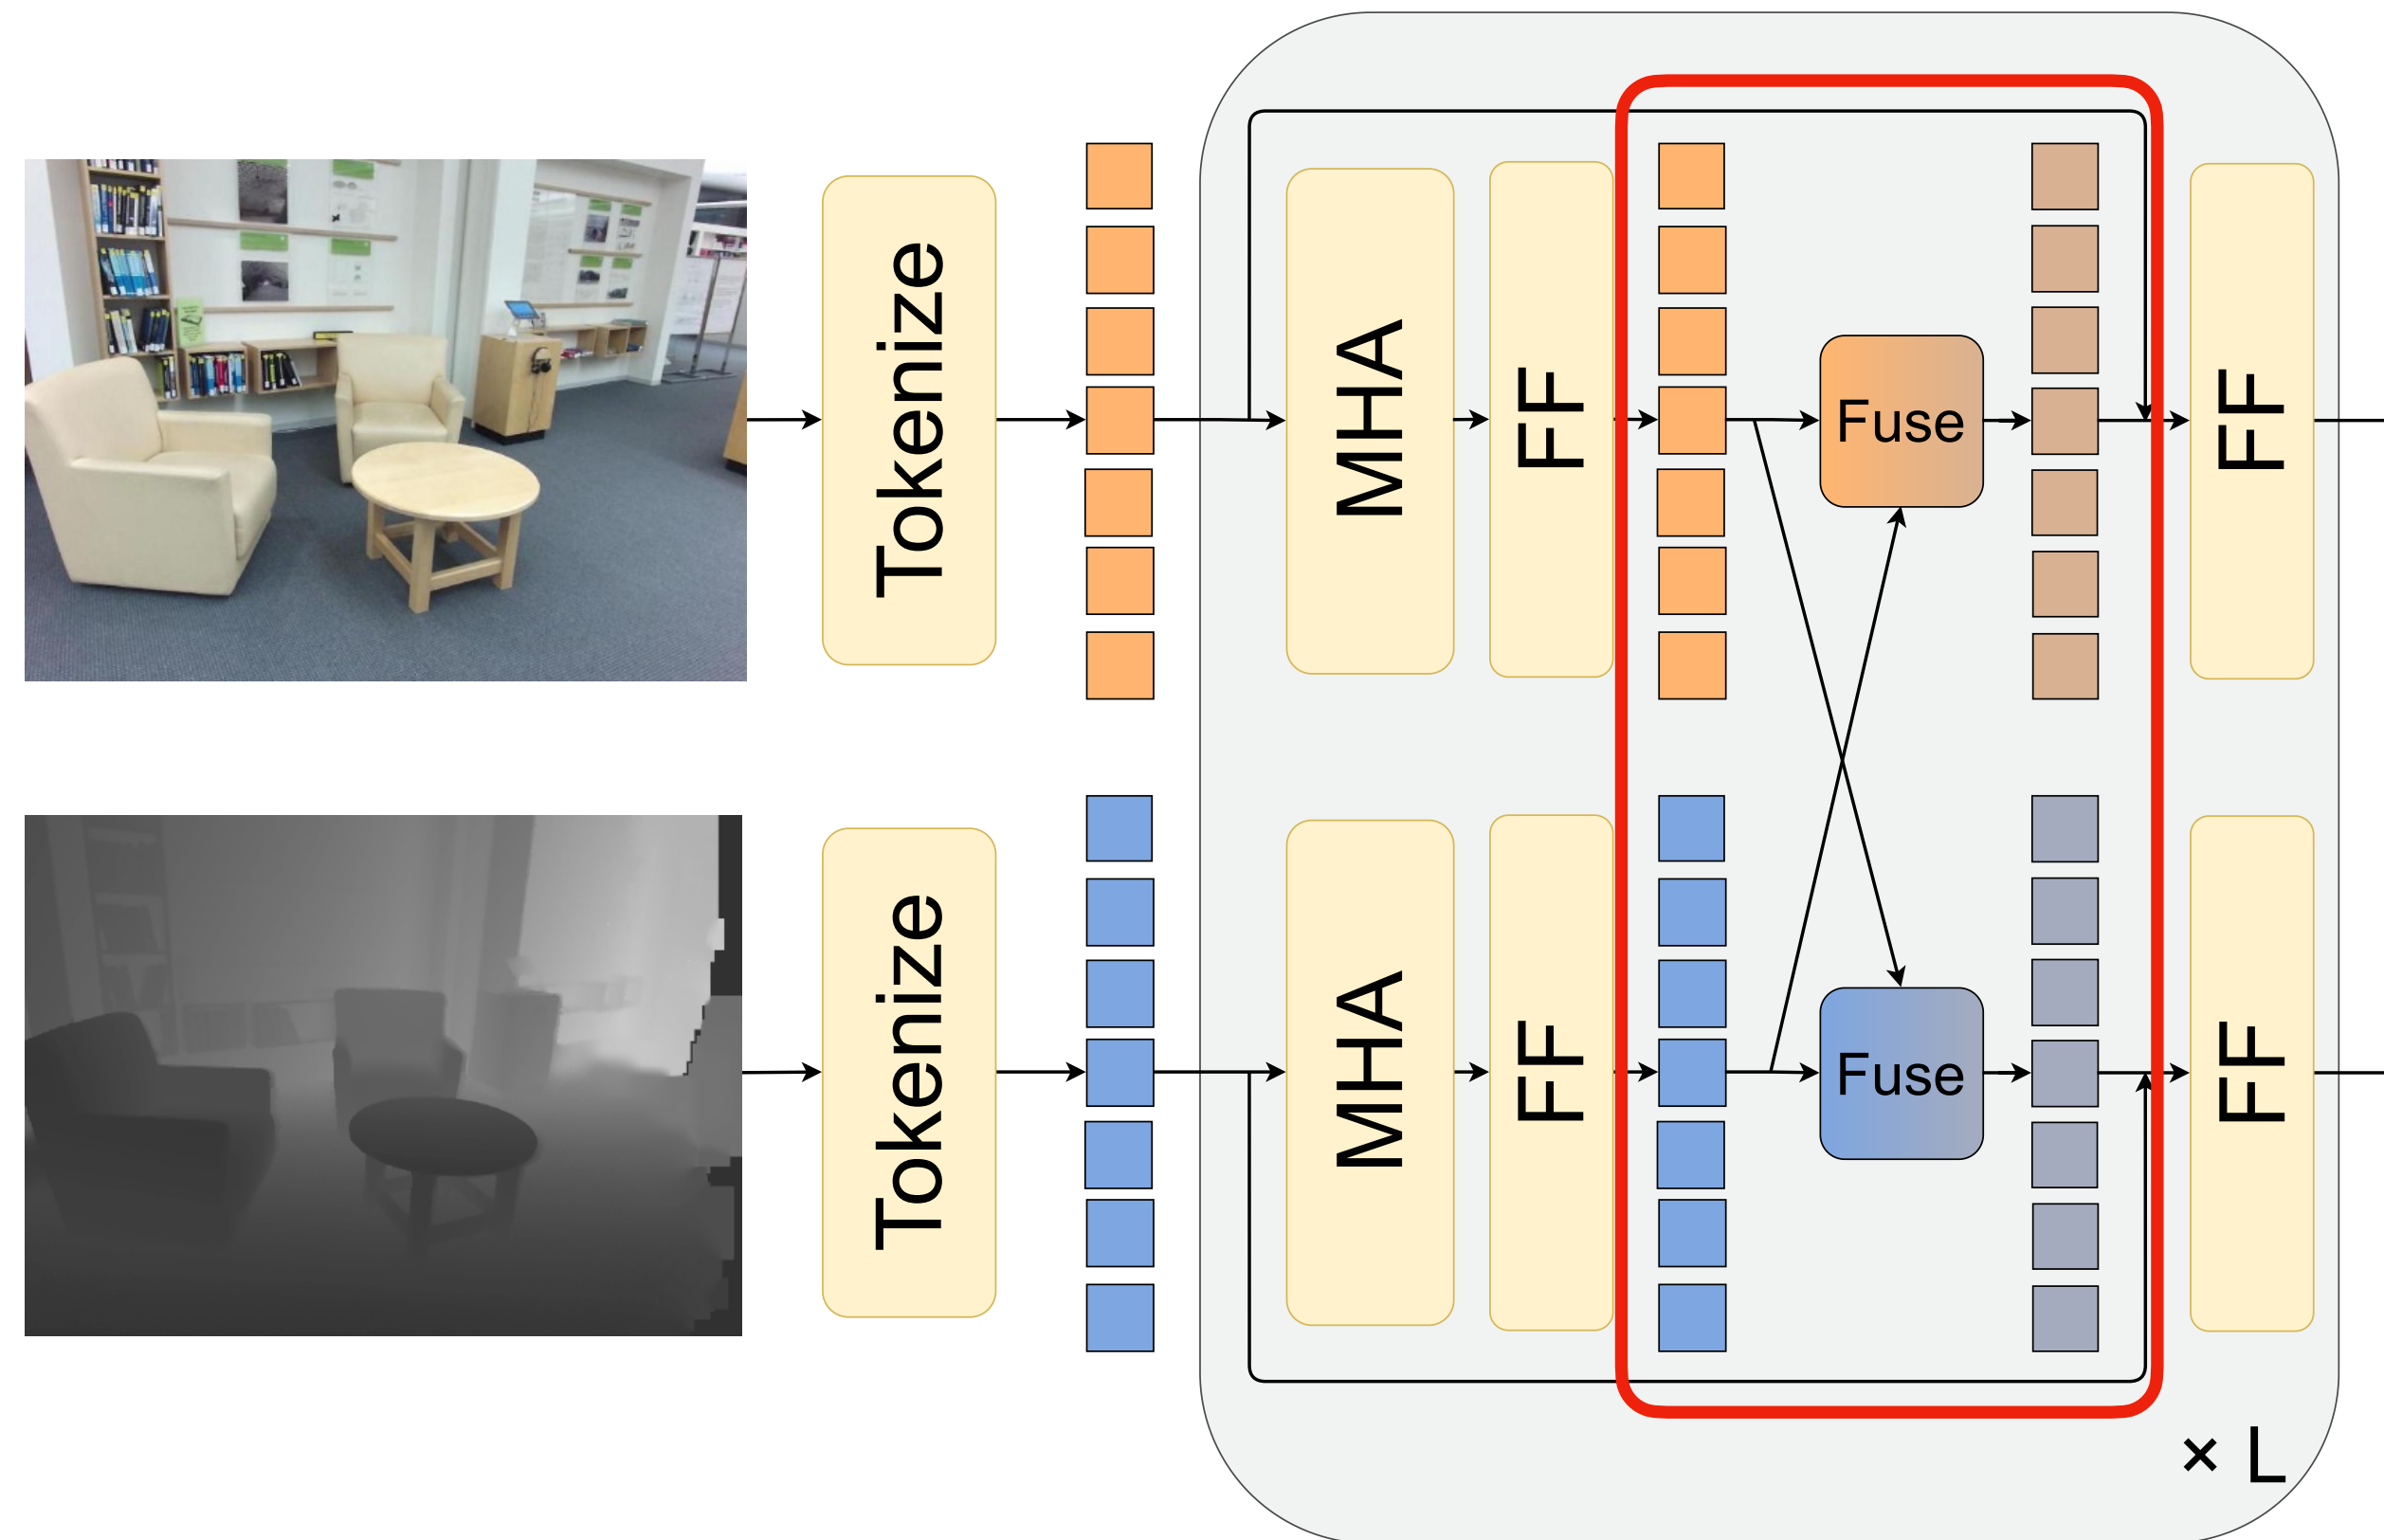


Linear Fusion

Linear Fusion: Multi-Modal Semantic Segmentation Model


simple fusion mechanism

no extra trainable parameters




Linear Fusion


Linear Fusion: Multi-Modal Semantic Segmentation Model




Linear Fusion: Multi-Modal Semantic Segmentation Model



Linear Fusion: Multi-Modal Semantic Segmentation Model



Linear Fusion: Multi-Modal Semantic Segmentation Model



Linear Fusion

For low-label setting, a multi-modal model with
simple fusion mechanism

M3L: Multi-modal teacher for Masked Modality Learning

A Semi-Supervised framework to increase
robustness to Missing-Modalities

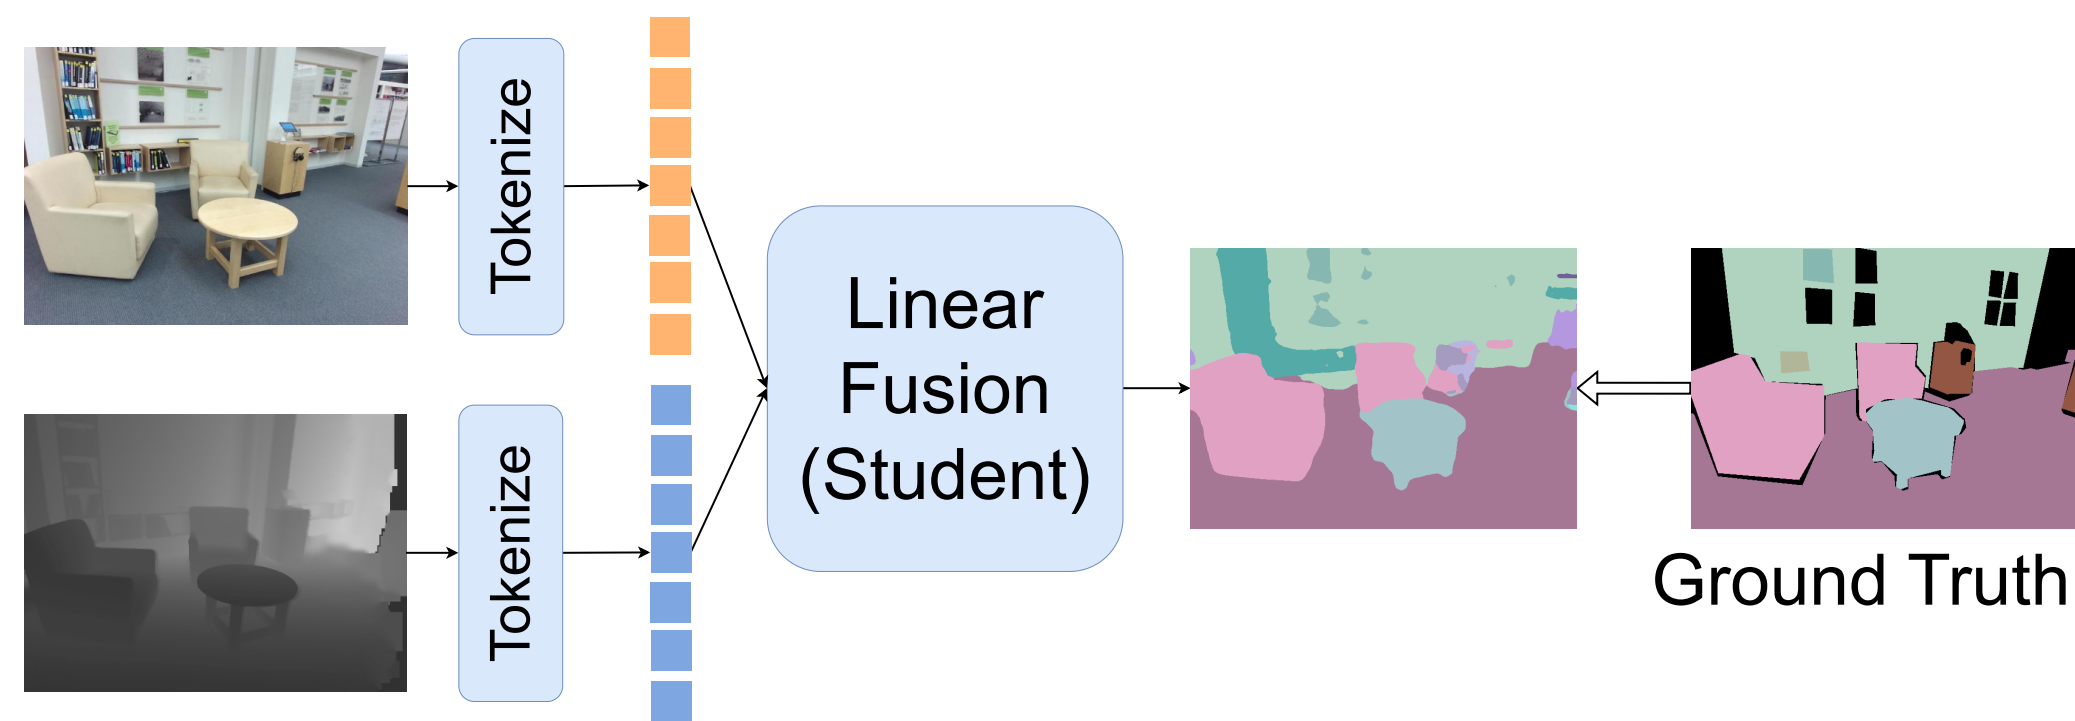
Linear Fusion

For low-label setting, a multi-modal model with simple fusion mechanism

M3L: Multi-modal teacher for Masked Modality Learning

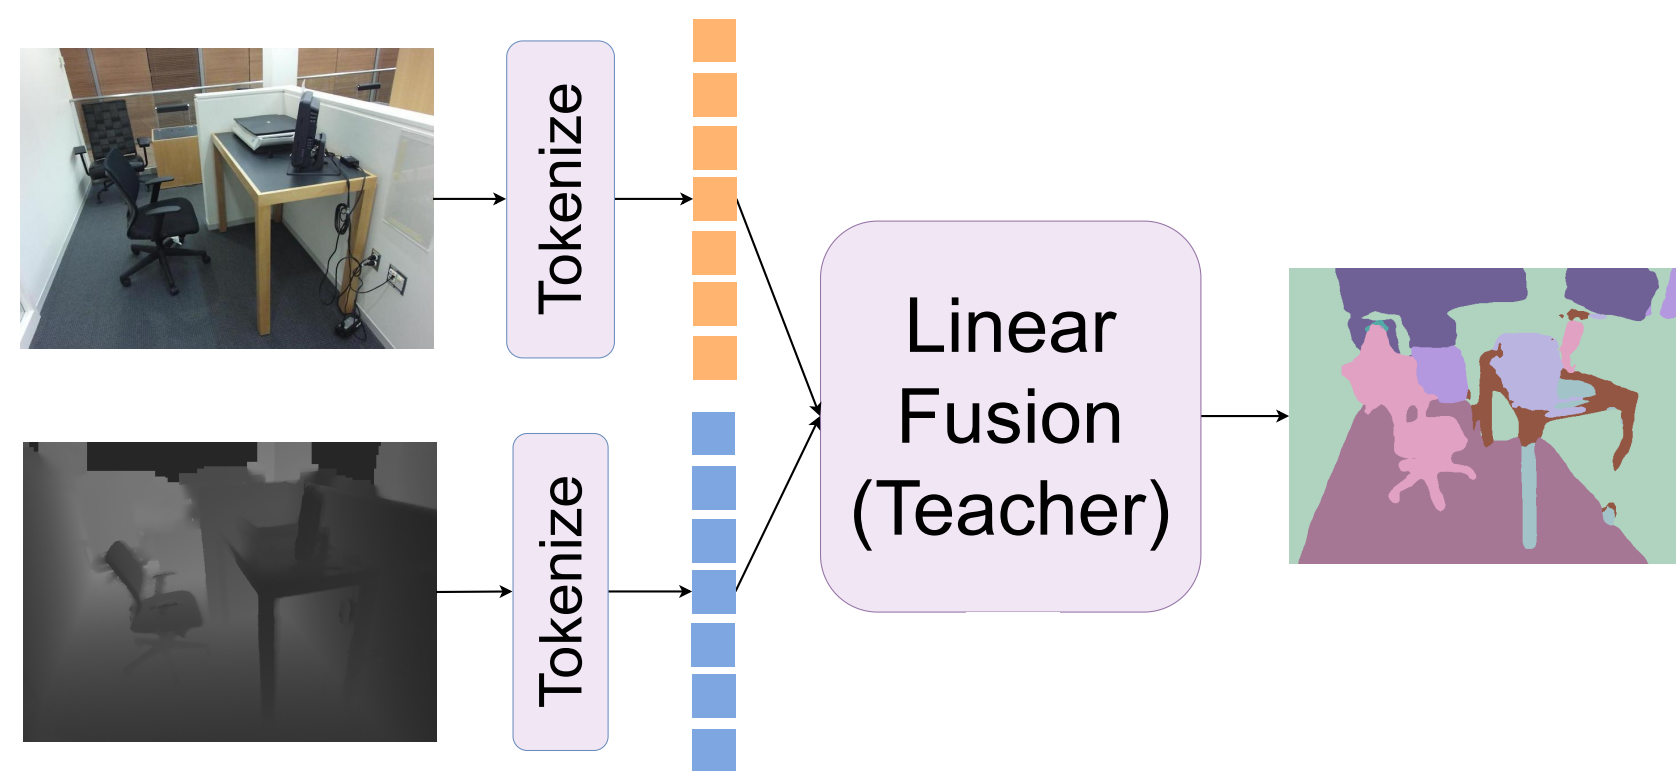
A **Semi-Supervised** framework to increase robustness to **Missing-Modalities**

M3L: Multi-modal teacher for Masked Modality Learning



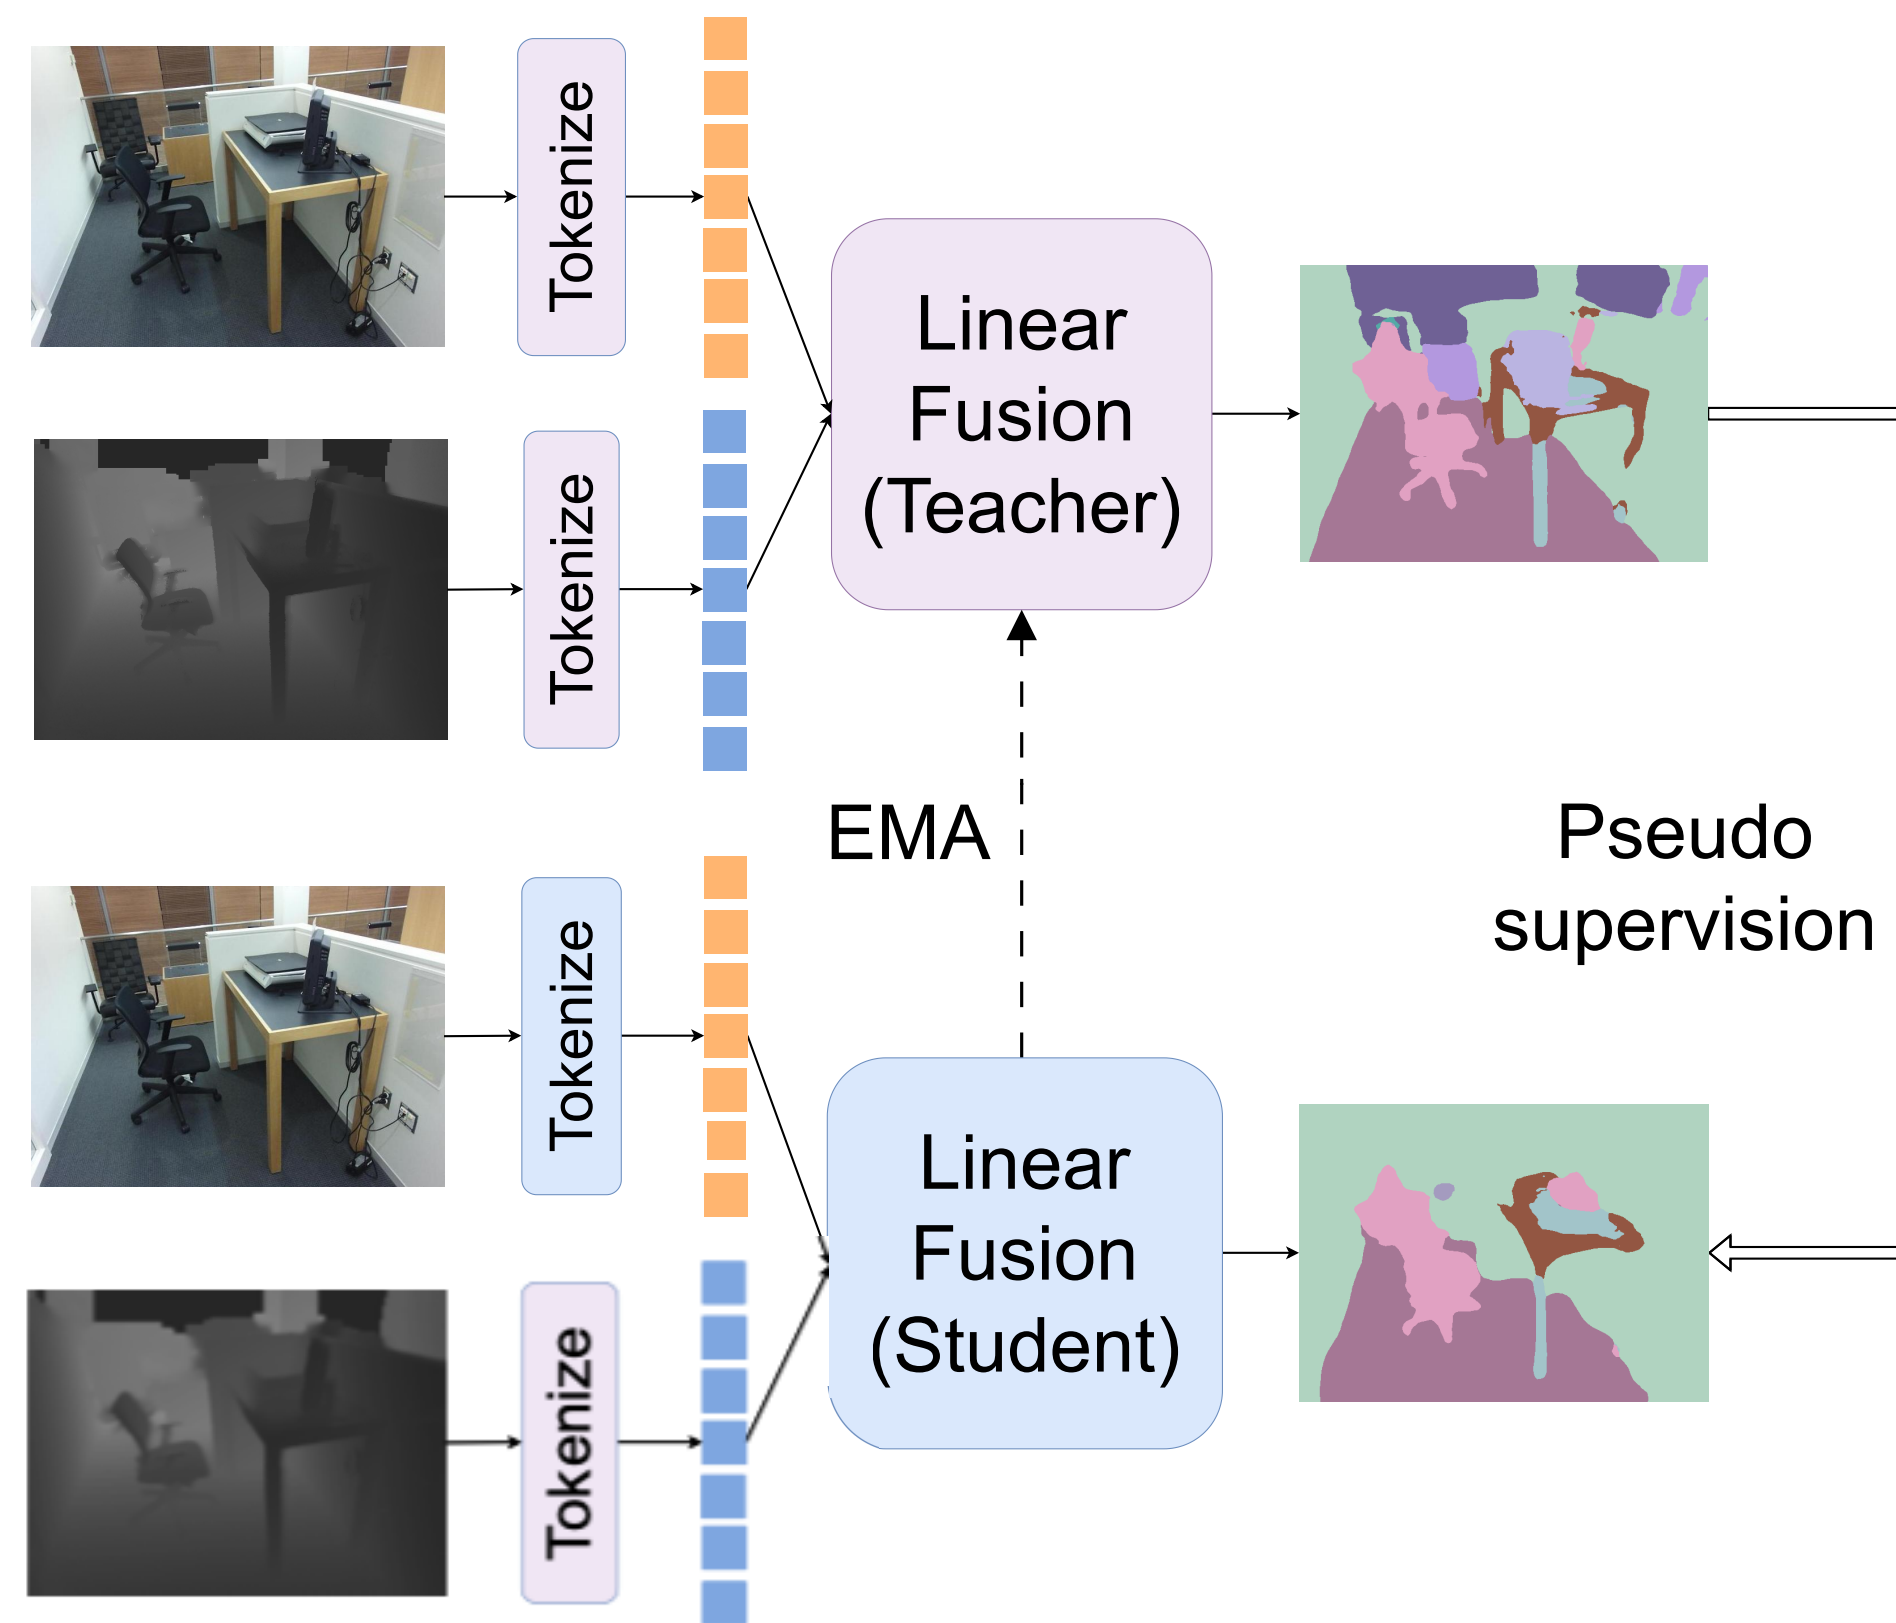
(a) Supervised

M3L: Multi-modal teacher for Masked Modality Learning



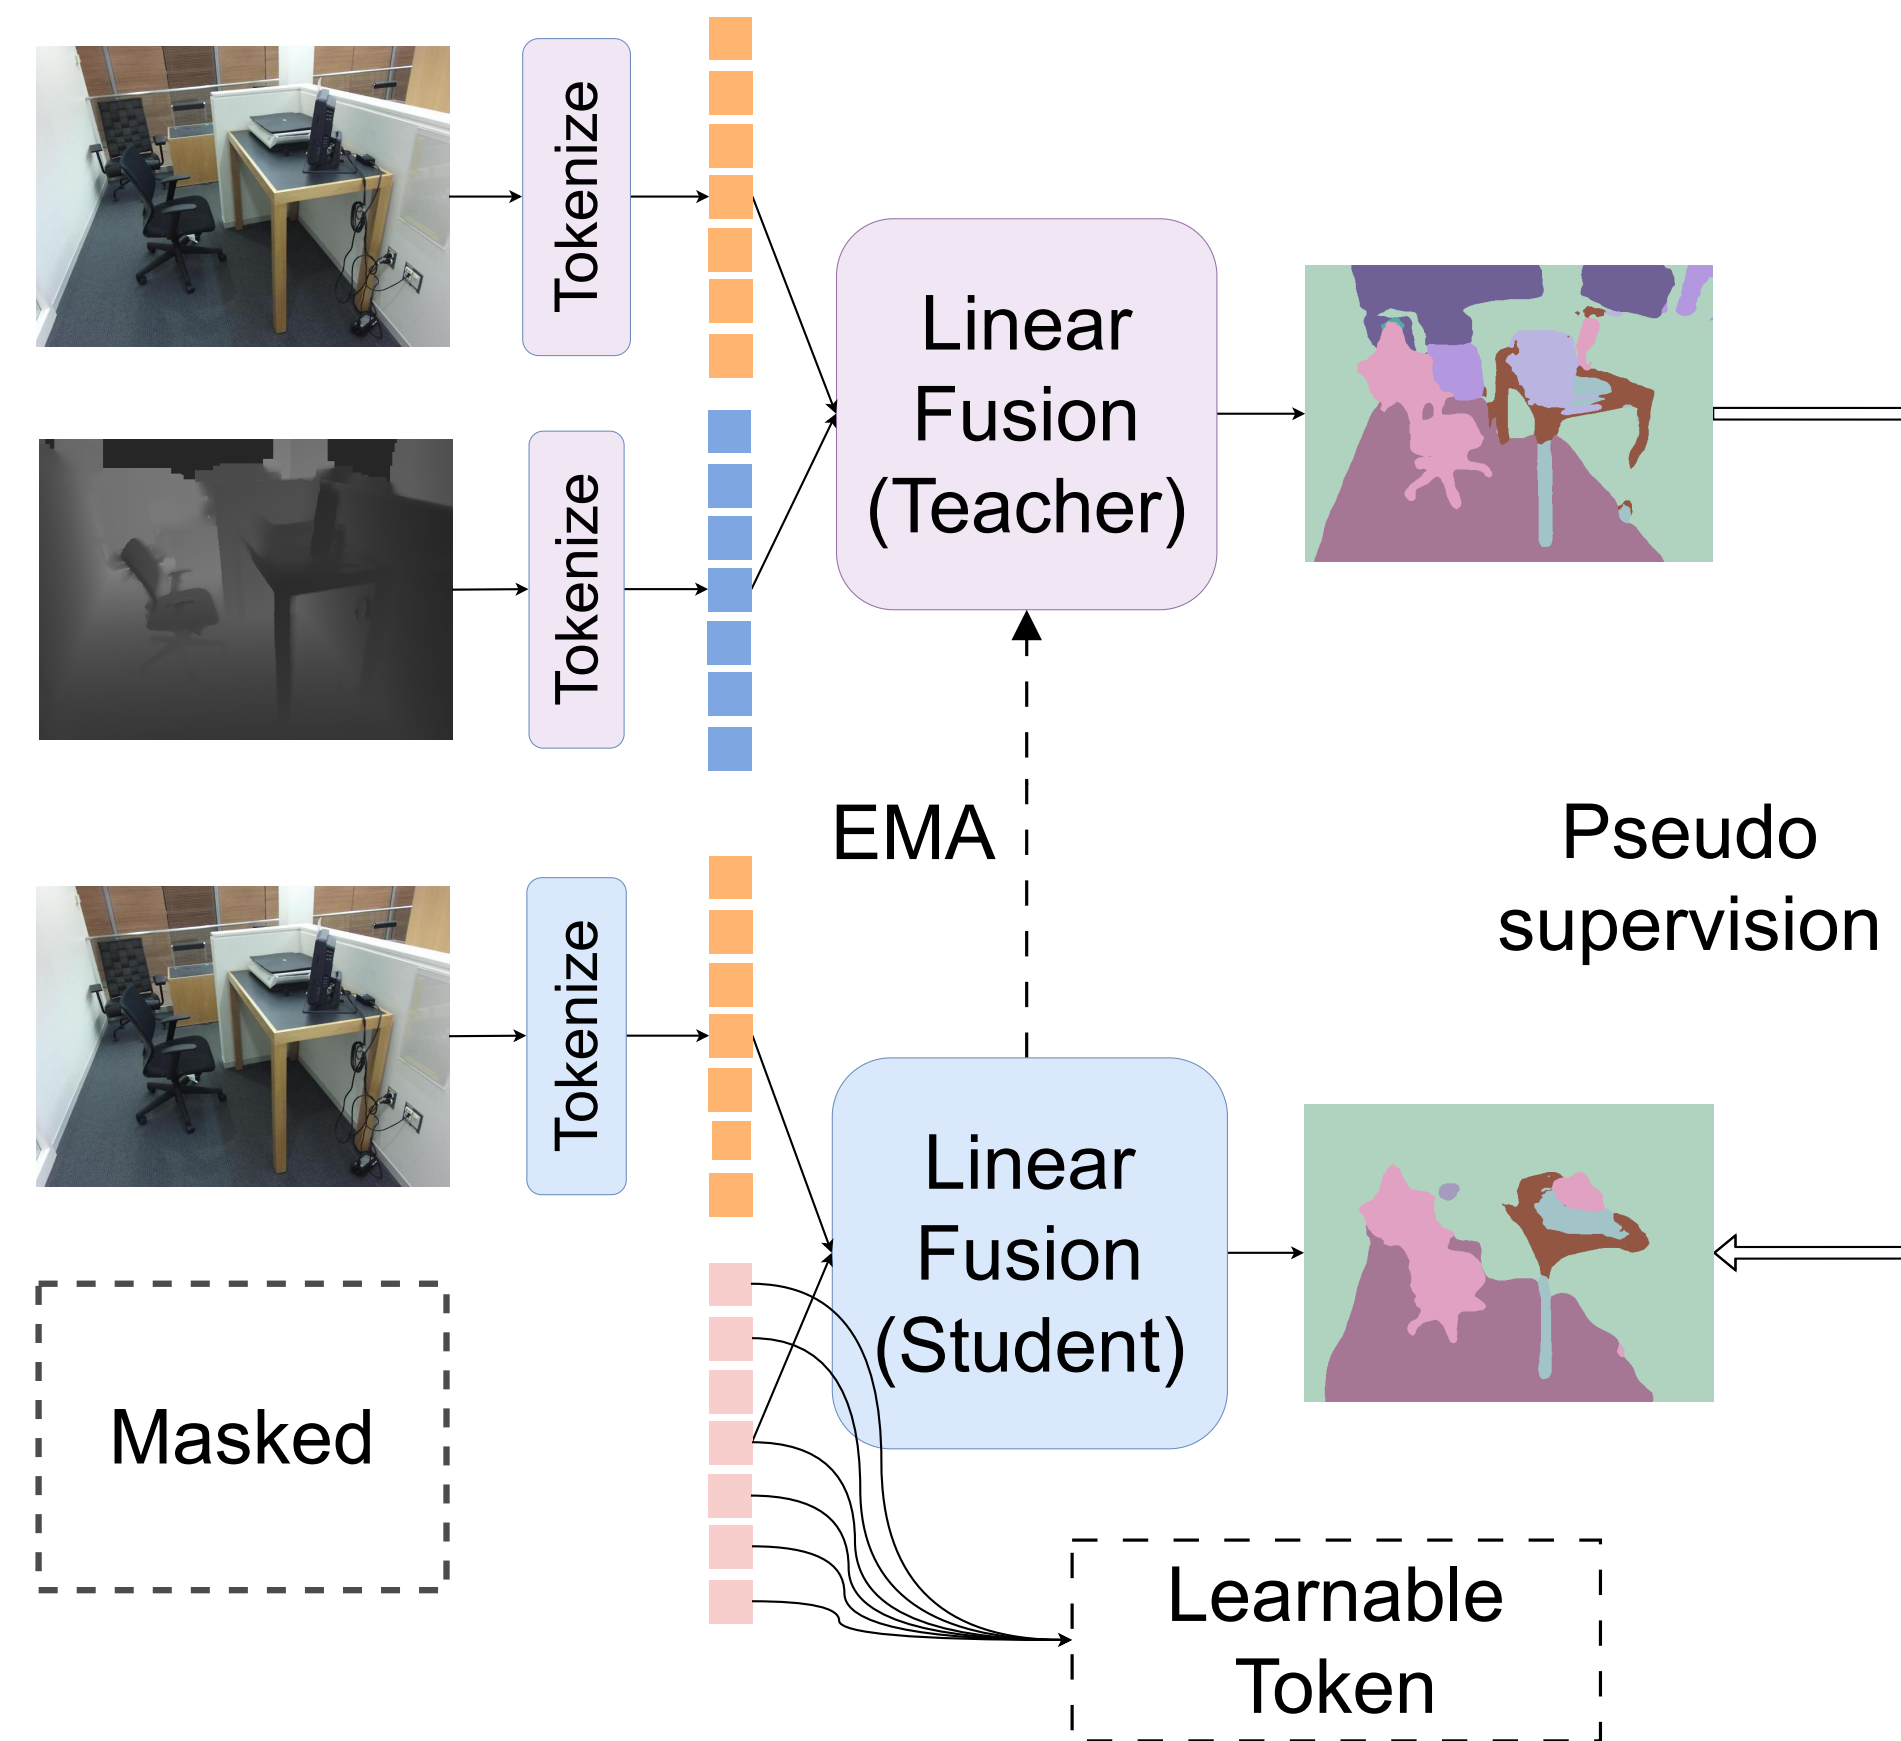
(b) Unsupervised

M3L: Multi-modal teacher for Masked Modality Learning



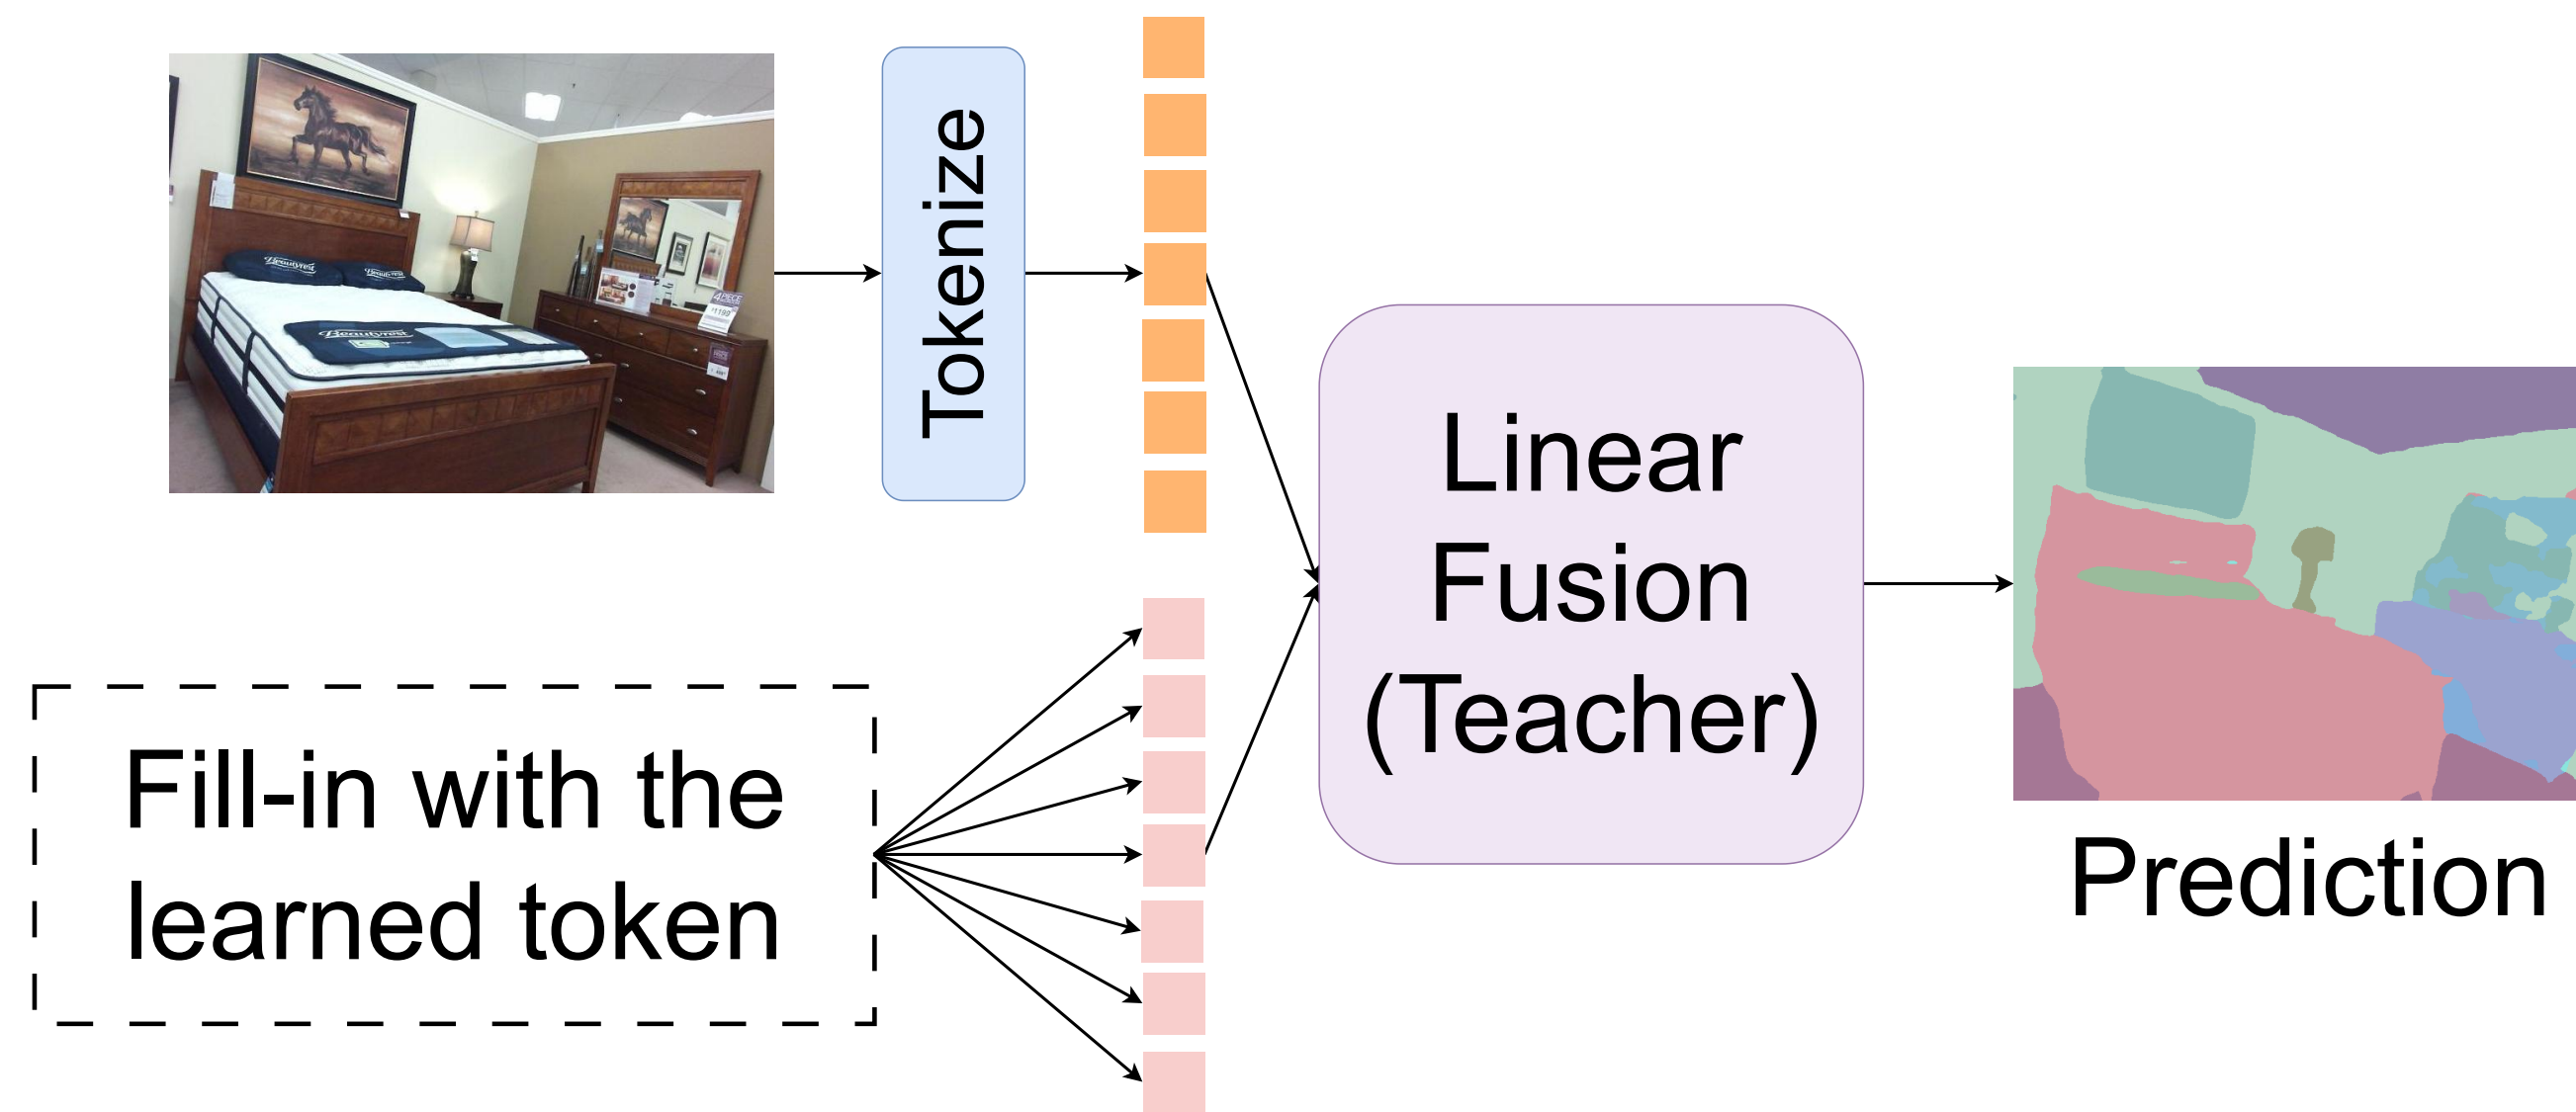
(b) Unsupervised

M3L: Multi-modal teacher for Masked Modality Learning



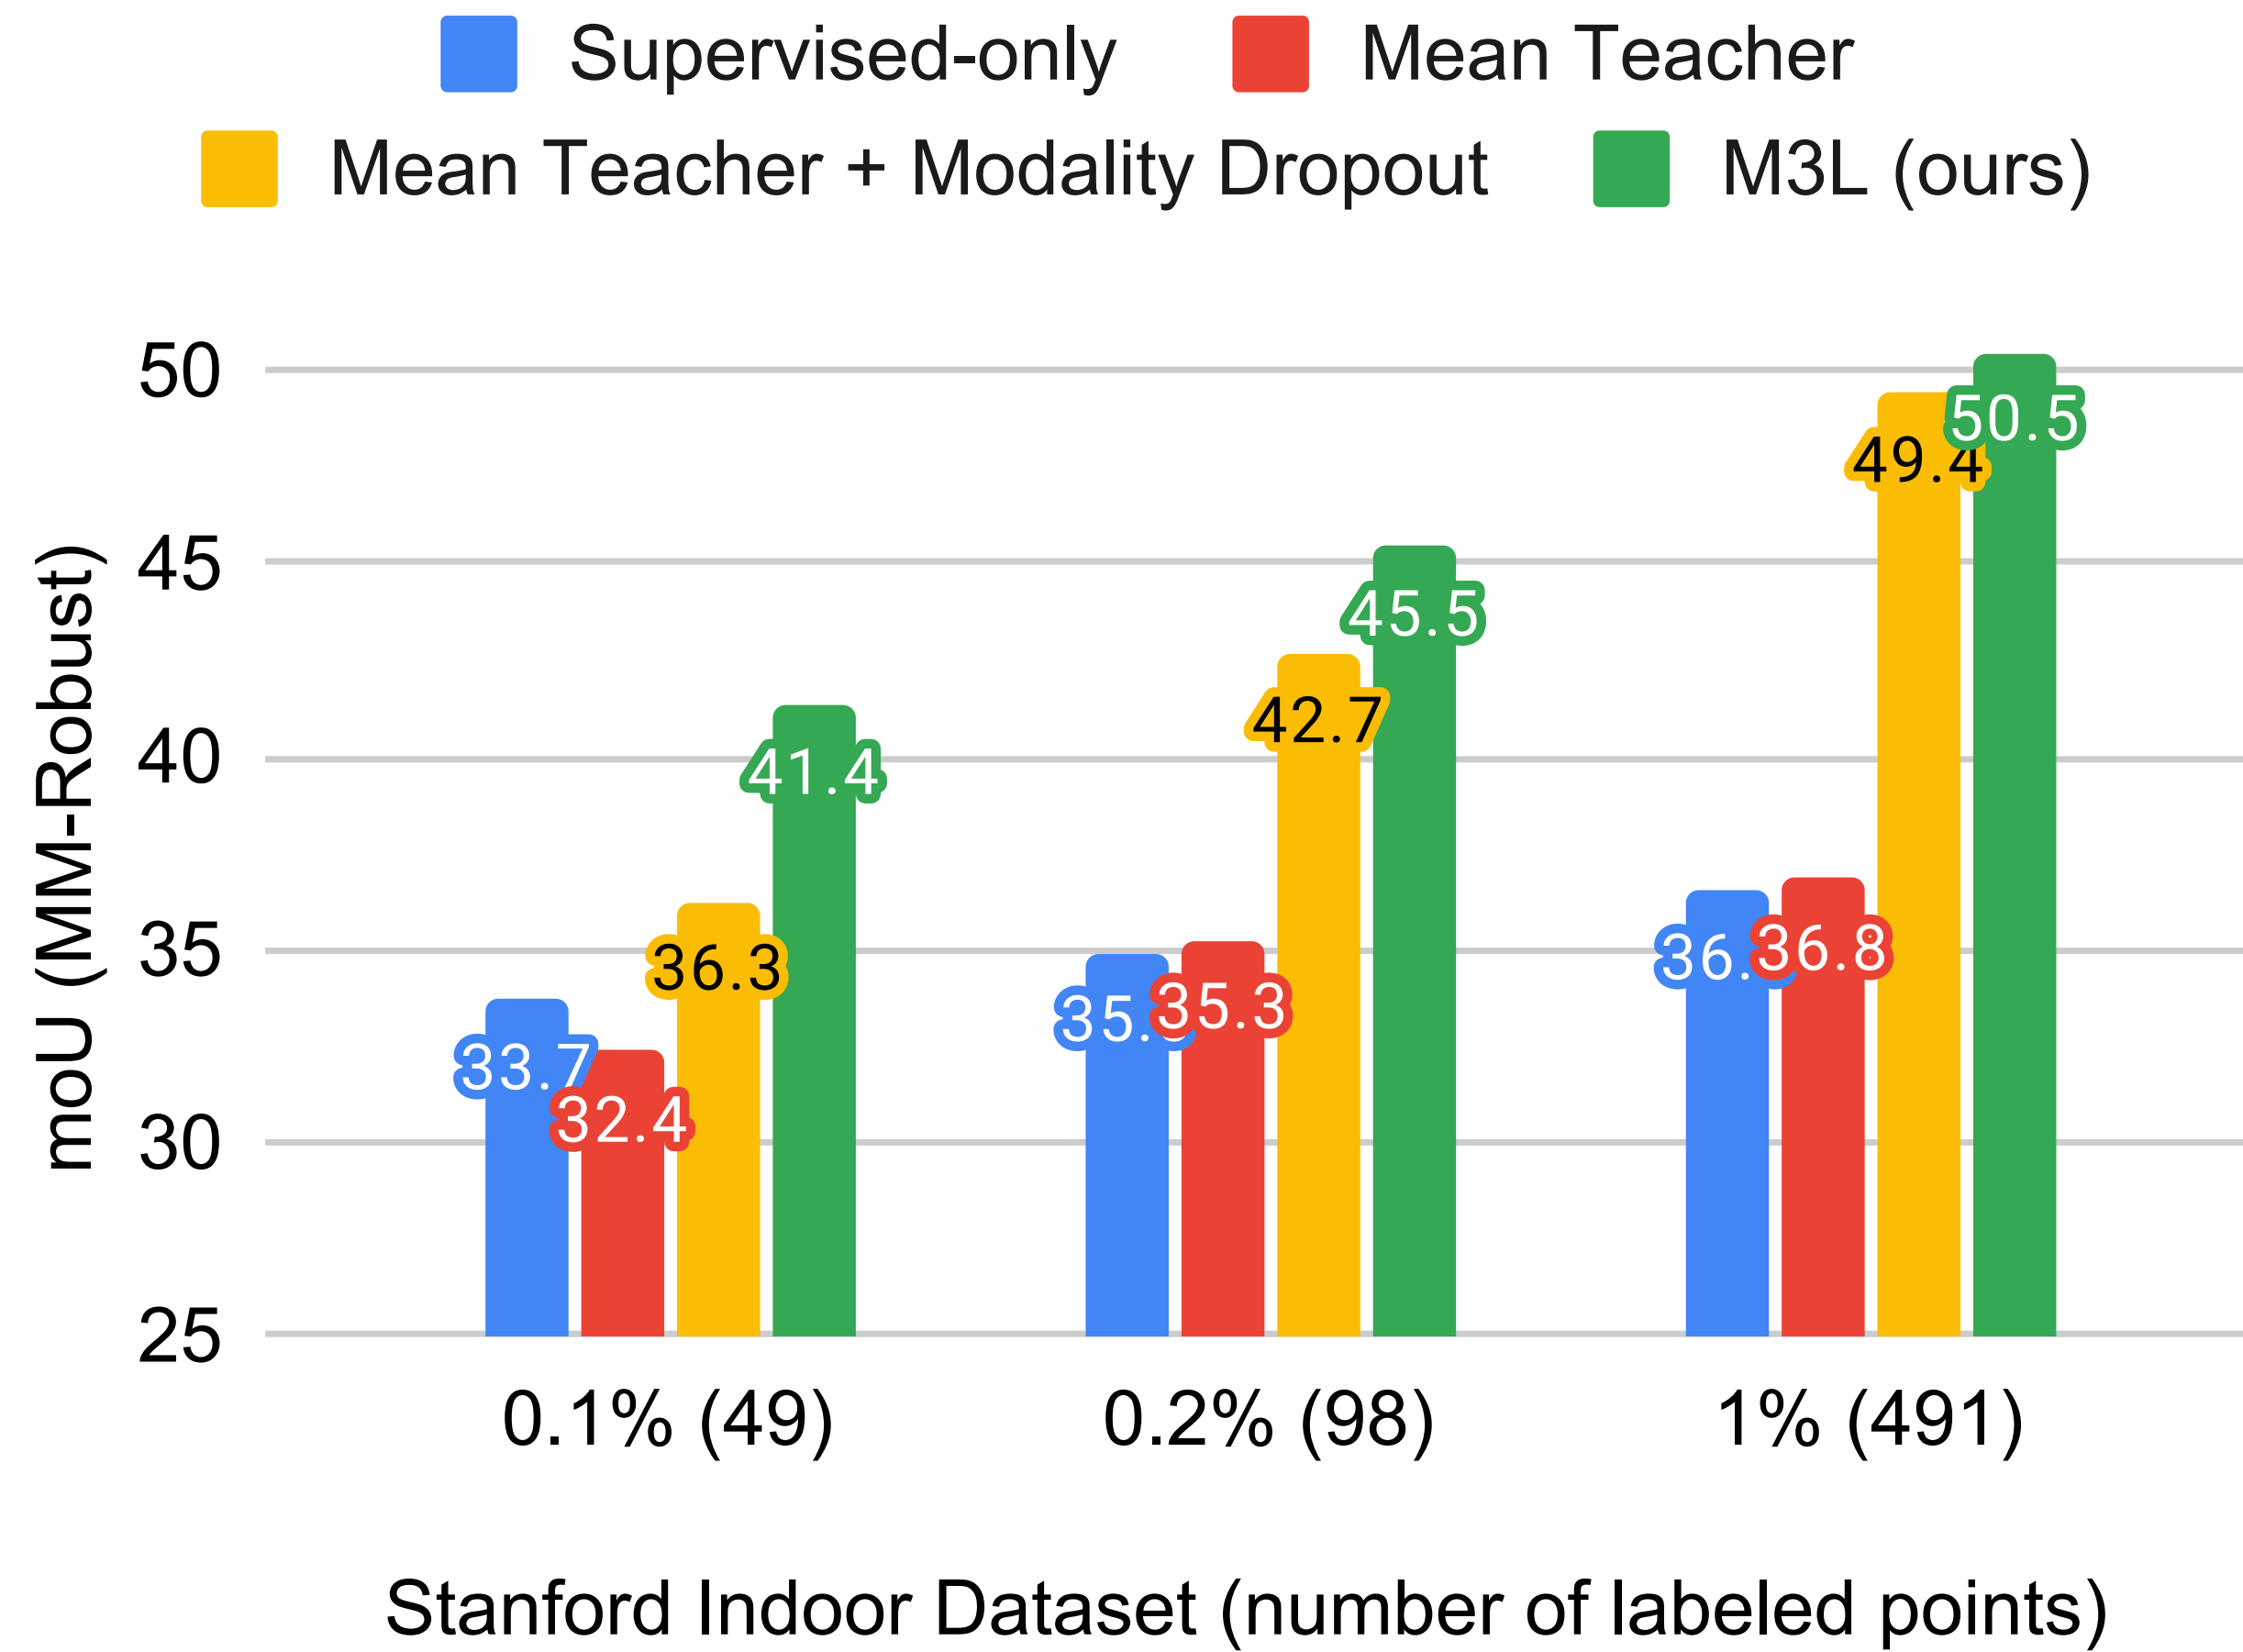
(b) Unsupervised

M3L: Multi-modal teacher for Masked Modality Learning

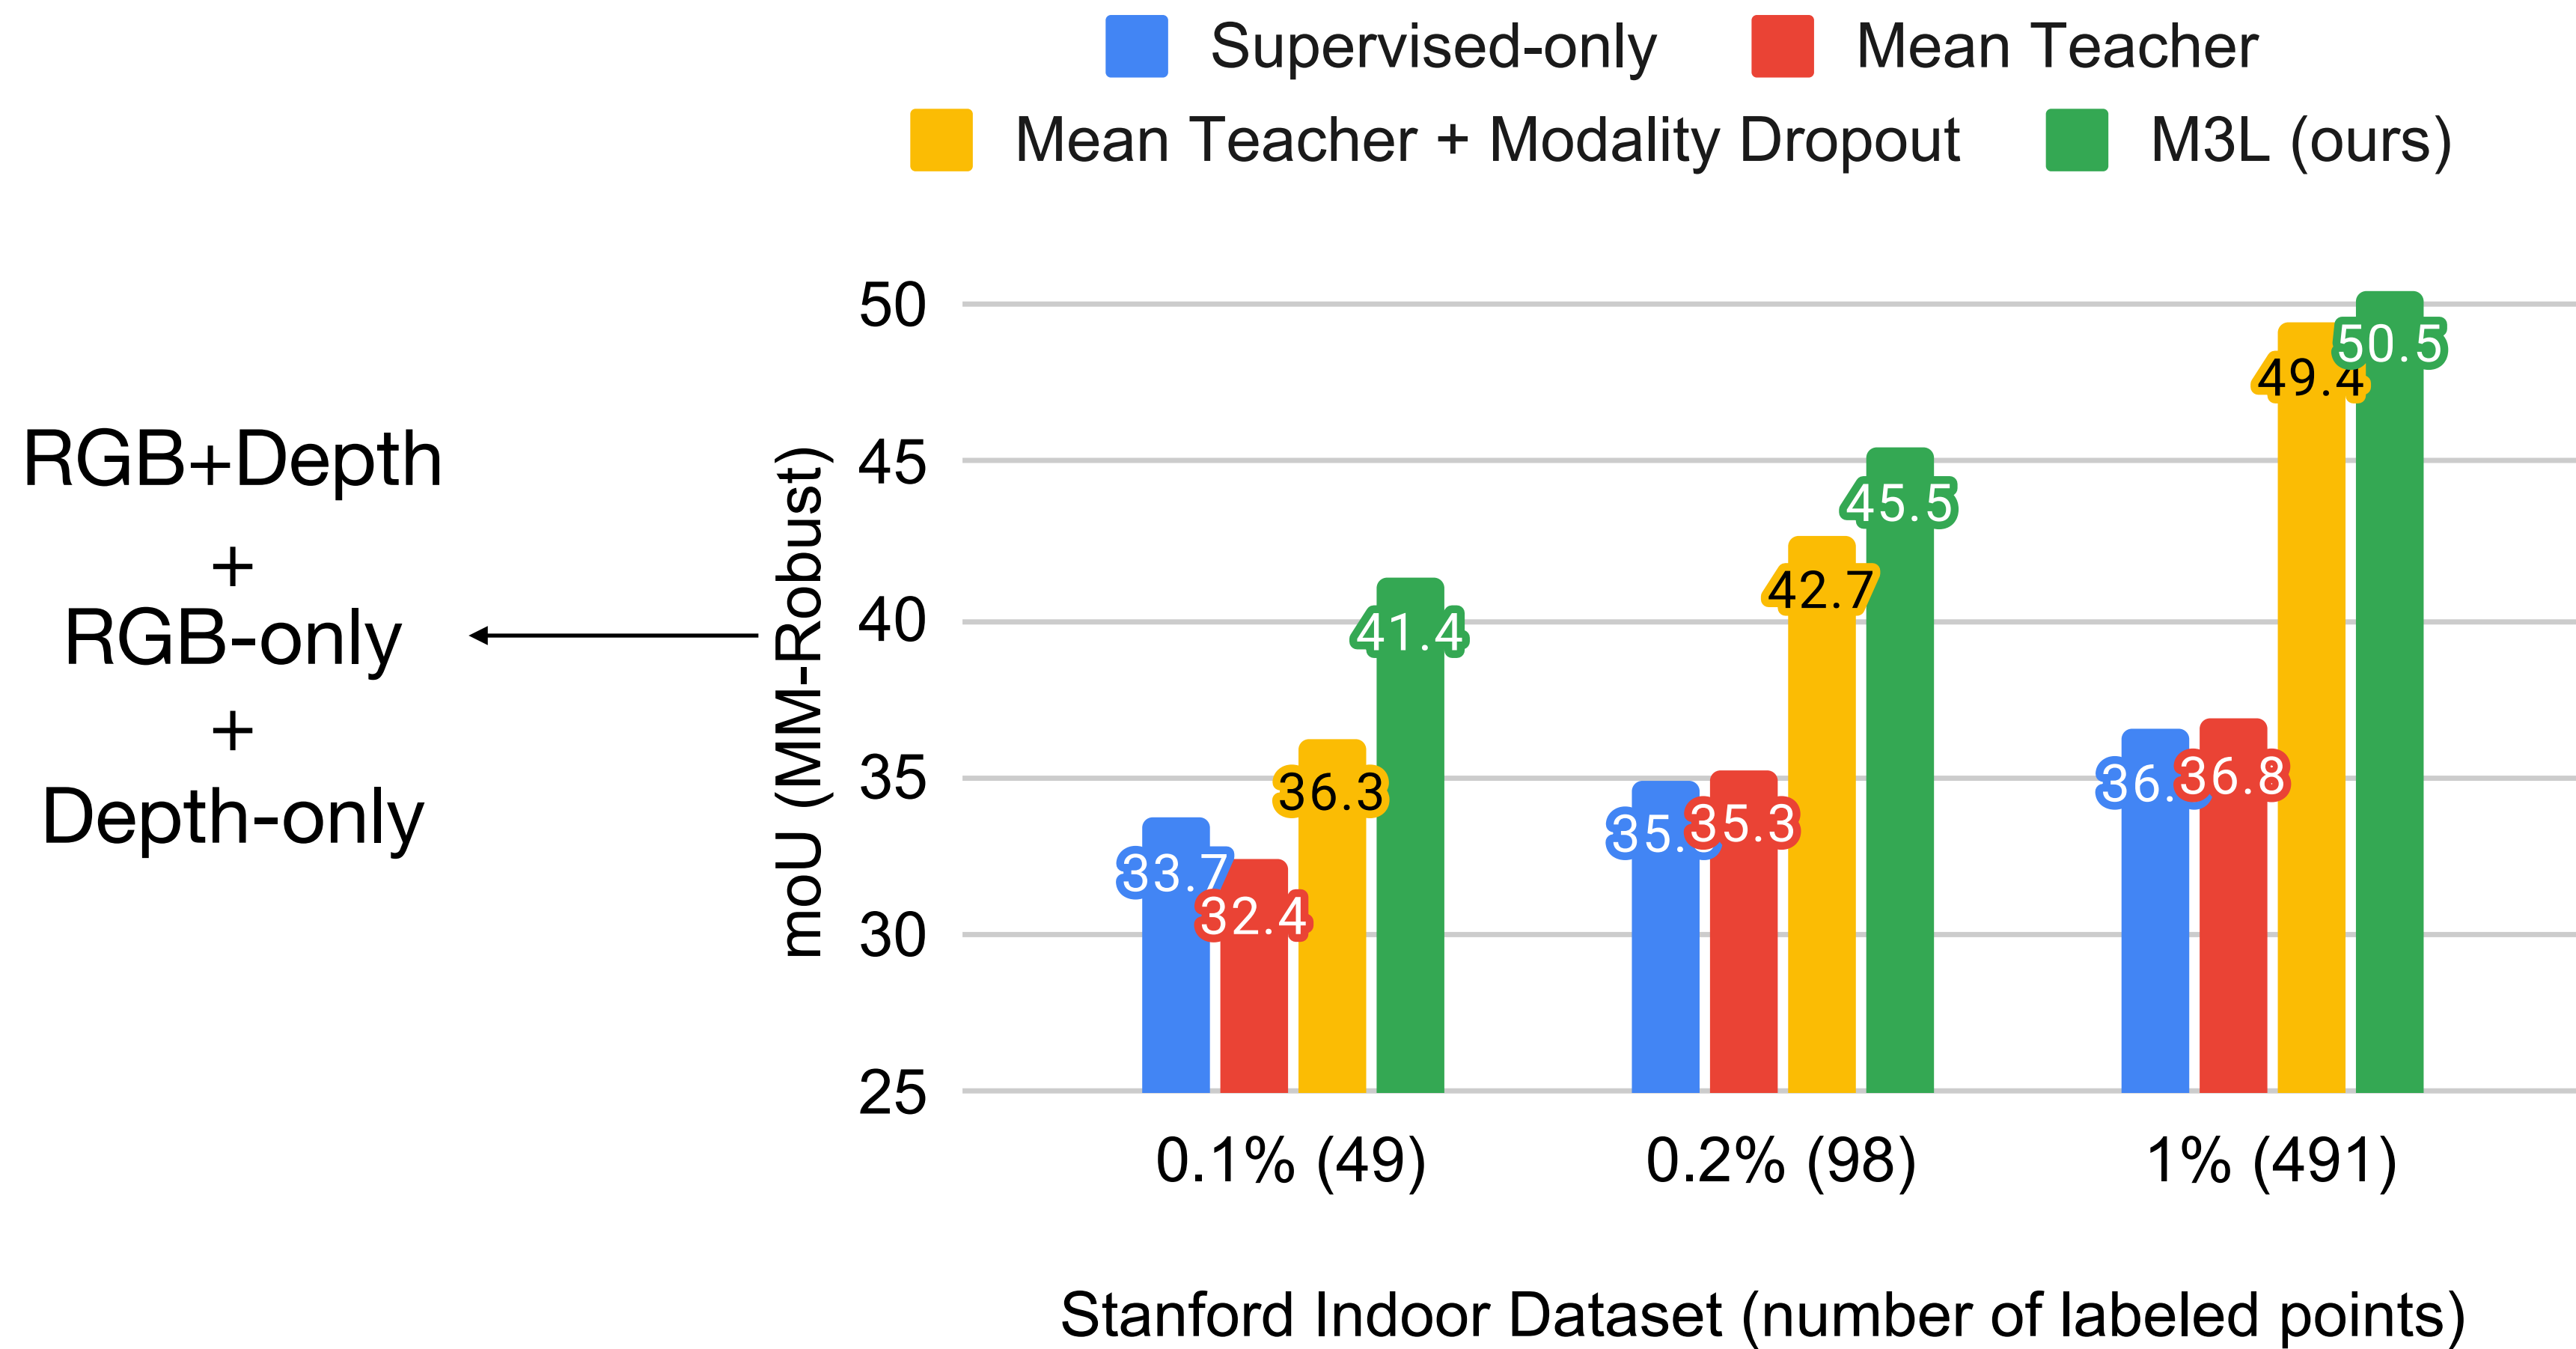


(c) Inference with missing modality


M3L: Multi-modal teacher for Masked Modality Learning



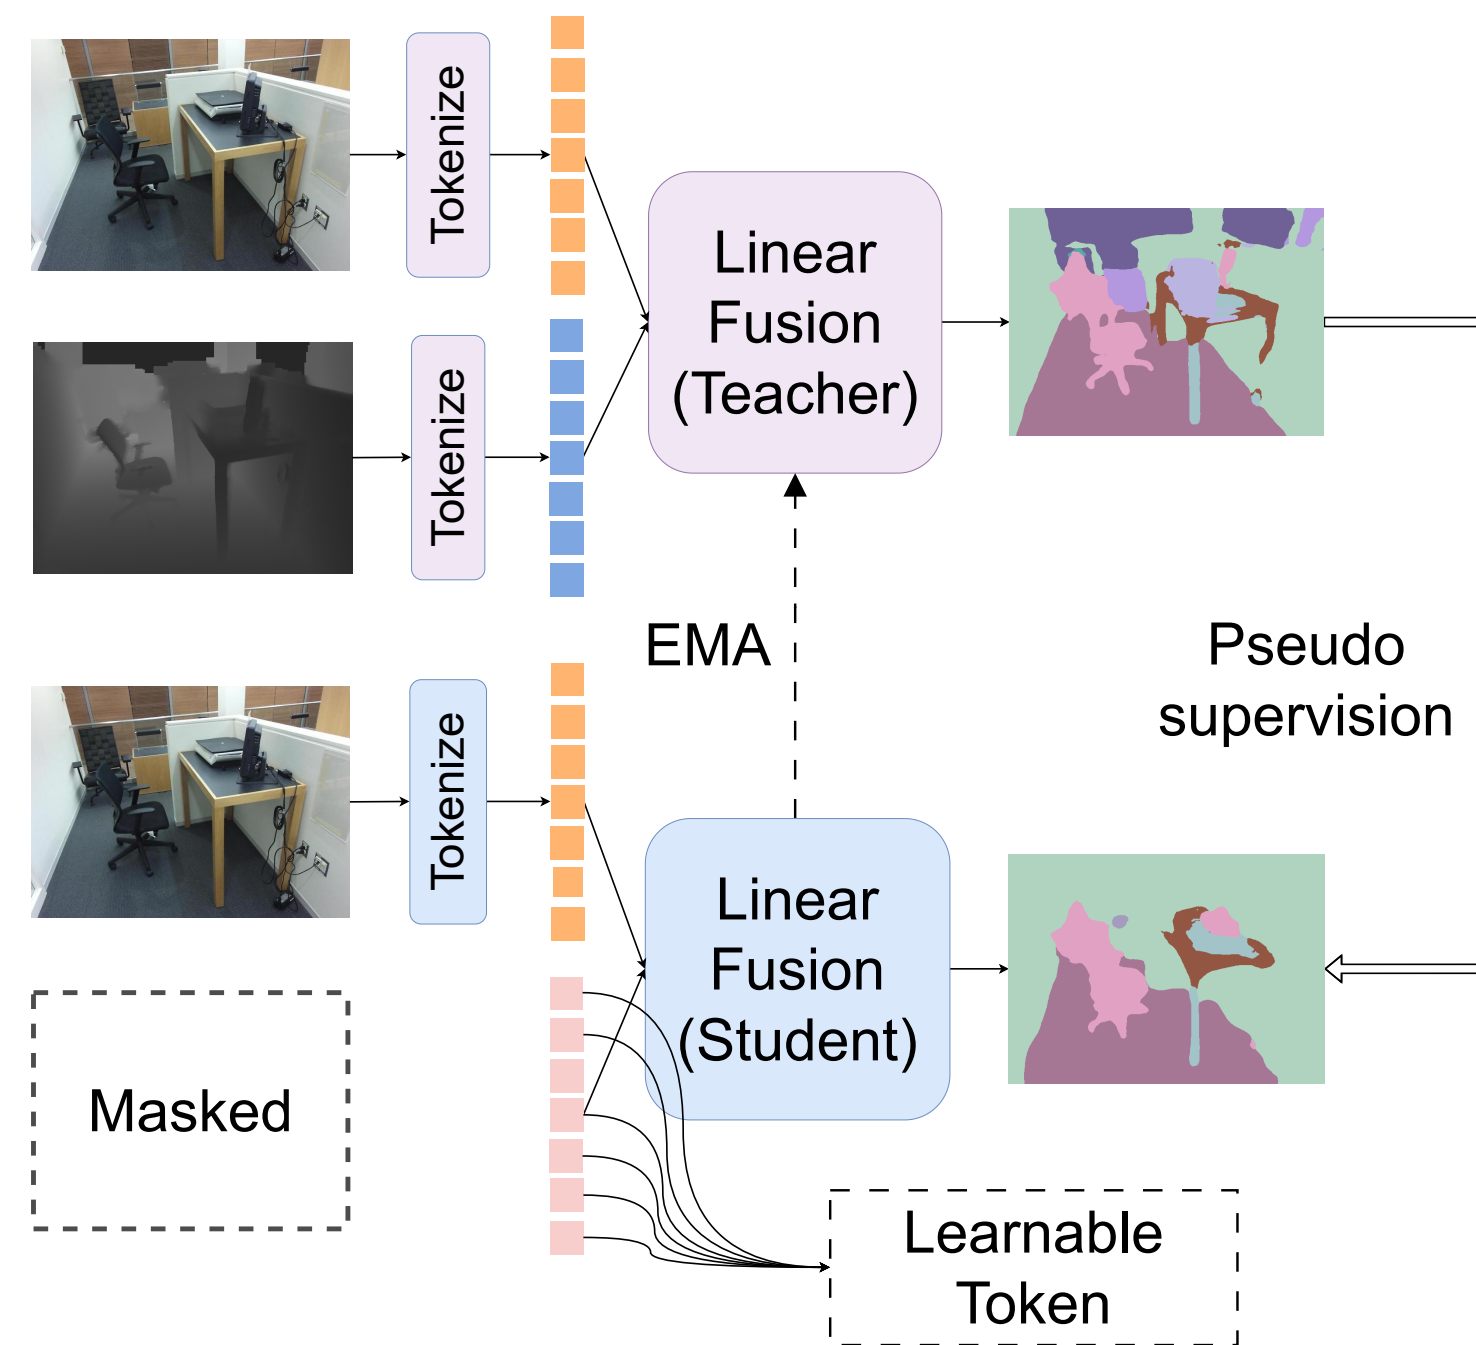
M3L: Multi-modal teacher for Masked Modality Learning



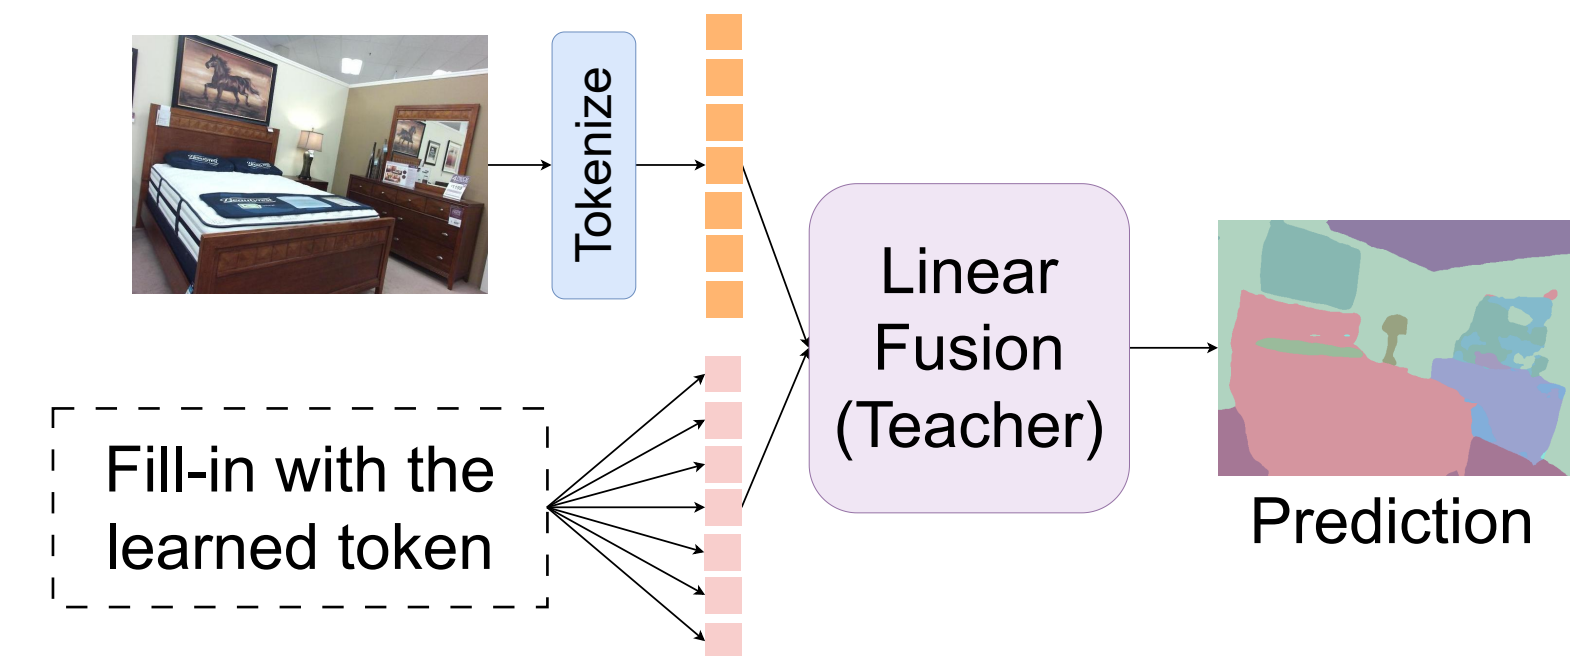
M3L: Multi-modal teacher for Masked Modality Learning



M3L: as unimodal semi-supervised segmentation framework




Training with both modalities



Inference with RGB only

M3L: as unimodal semi-supervised segmentation framework

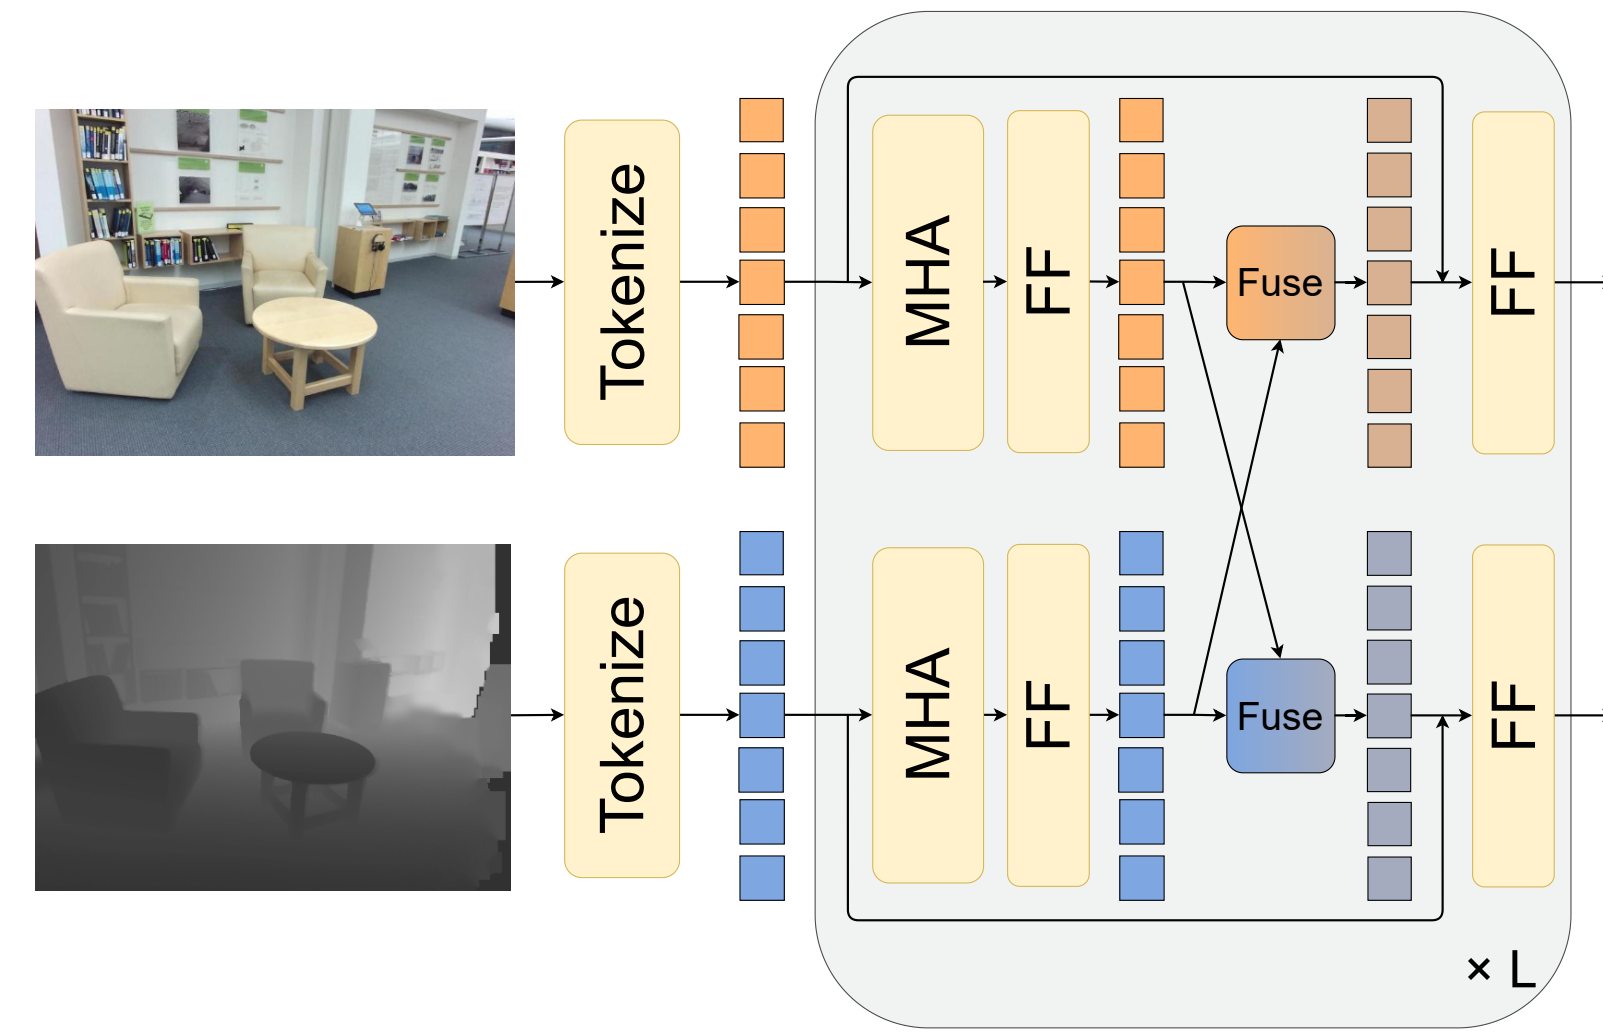


Contributions

Linear Fusion

Simple fusion mechanism

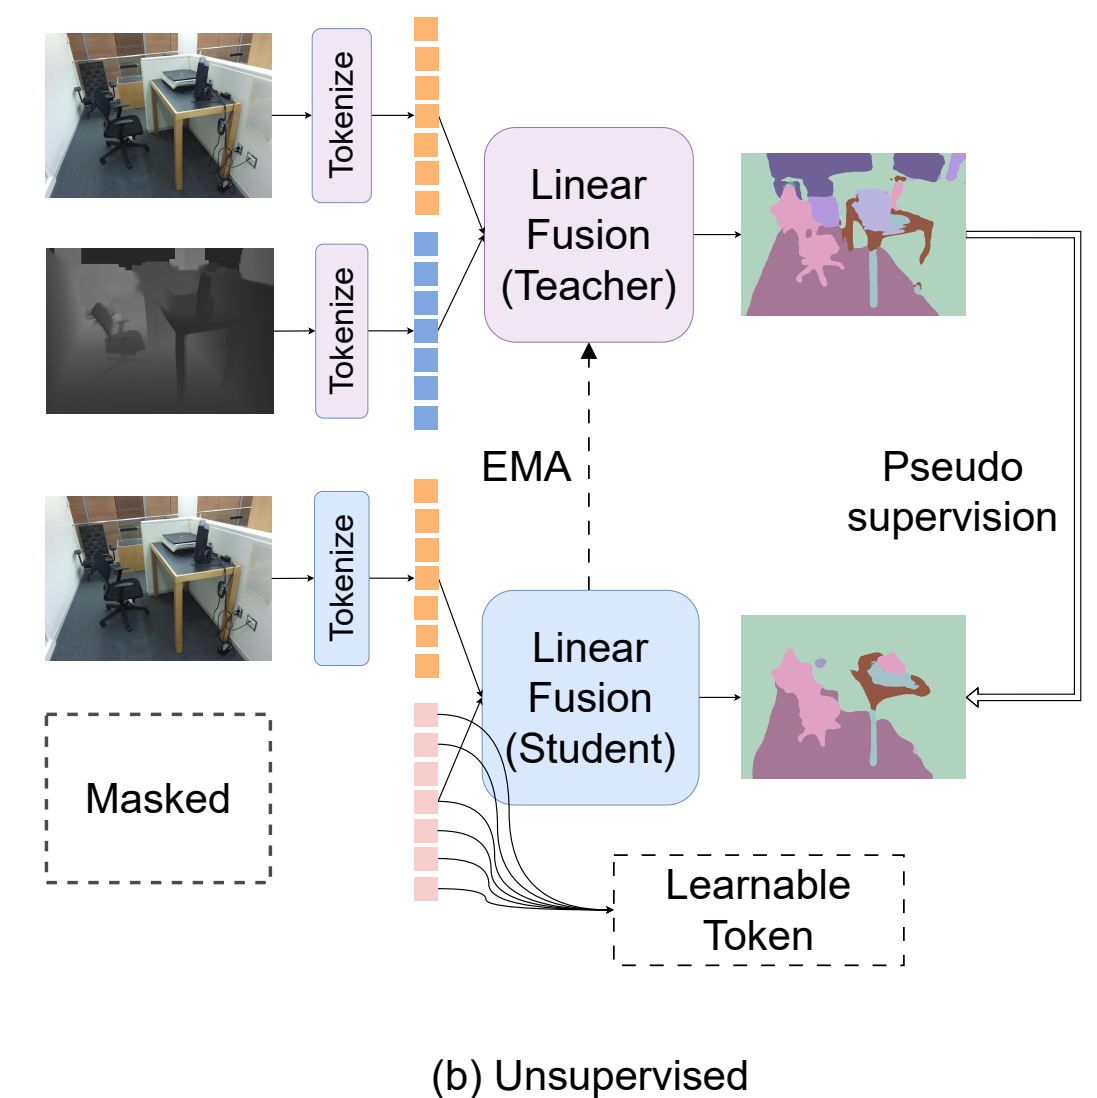
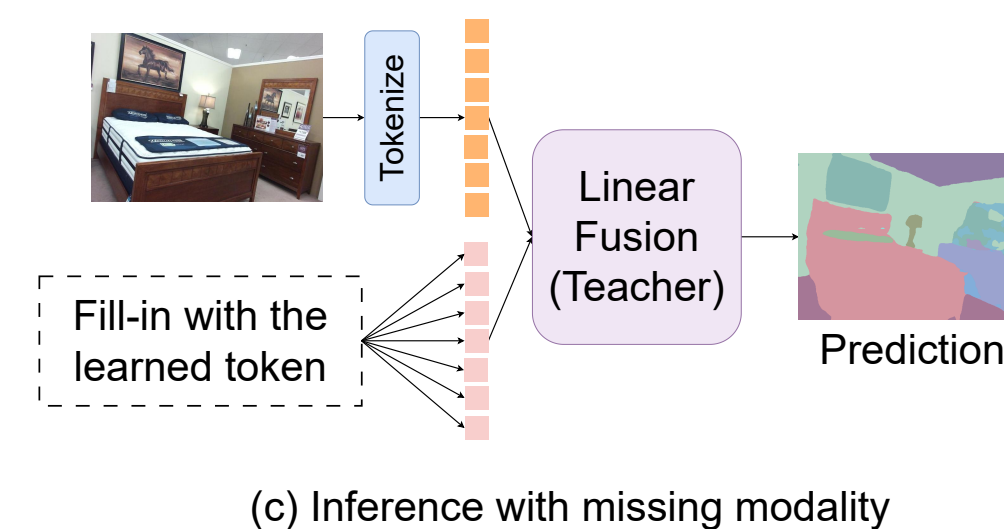
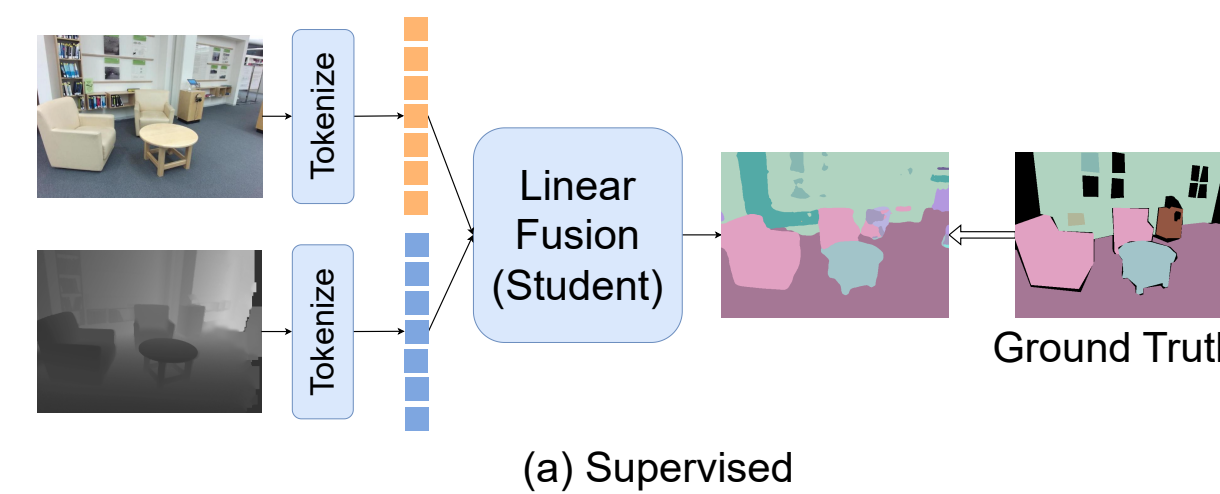
No extra trainable parameters



M3L

Semi-Supervised

Robustness to
Missing-Modalities



Thank you for listening!

Email: harsh.maheshwari@gatech.edu

Code: <https://github.com/harshm121/M3L>

Project Page: <https://harshm121.github.io/projects/m3l>