



World Conference on Transportation Research
Lectern Session A5-TS03A – AI/ML applications in Transportation:
From Infrastructure Sensing to User Trust
July 8, 2026

Weak supervision and fine-tuning with contrastive learning for multiclass lane marking segmentation

Julia Alves Porto,
PhD Candidate



Department of Transportation Planning and Engineering,
National Technical University of Athens

Agenda

➤ Background

- Motivation
- Related work

➤ Methodology

- Dataset preparation
- Models
- Training pipeline and performance monitoring

➤ Results and discussion

- End-to-End training
- Fine-Tuning
- Hyperparameter tuning: loss weight
- Hyperparameter tuning: learning rate
- Visual results
- Comparison with SOTA

➤ Conclusion, limitations and future work

- Conclusion
- Limitations and future work



Background



Motivation

- **Autonomous vehicle** navigation
- **Safety assessment:**
 - Lane quality
 - Permitted motions derived from lane type
- **Novelties:**
 - Self-supervised learning branch without negative examples
 - Discuss learning under different configurations with the proposed approach
 - Open-source multiclass lane marking dataset



Related work

➤ Lane marking segmentation


- Based on color and shape cues
- Deep-Learning
- Lane marking types

➤ Contrastive Learning and Weak Supervision

- Augmented pairs of positive examples in contrast with negative examples
- Bootstrap Your Own Latent – BYOL: eliminates the need for negative example by relying on architecture asymmetry (Grill et al., 2020)
- Applications: reduced datasets, noisy labels

➤ Lane marking segmentation with contrastive learning


- CLLD (Zoljodi et al., 2024): improved recall and F-Score on CuLane and accuracy on TuSimple
- LaneCorrect (Nie et al., 2025): LiDAR
- Lane detection model fine-tuning (Khan et al., 2025): better embedding for different datasets



Source: Merugu and Adarsh, 2022

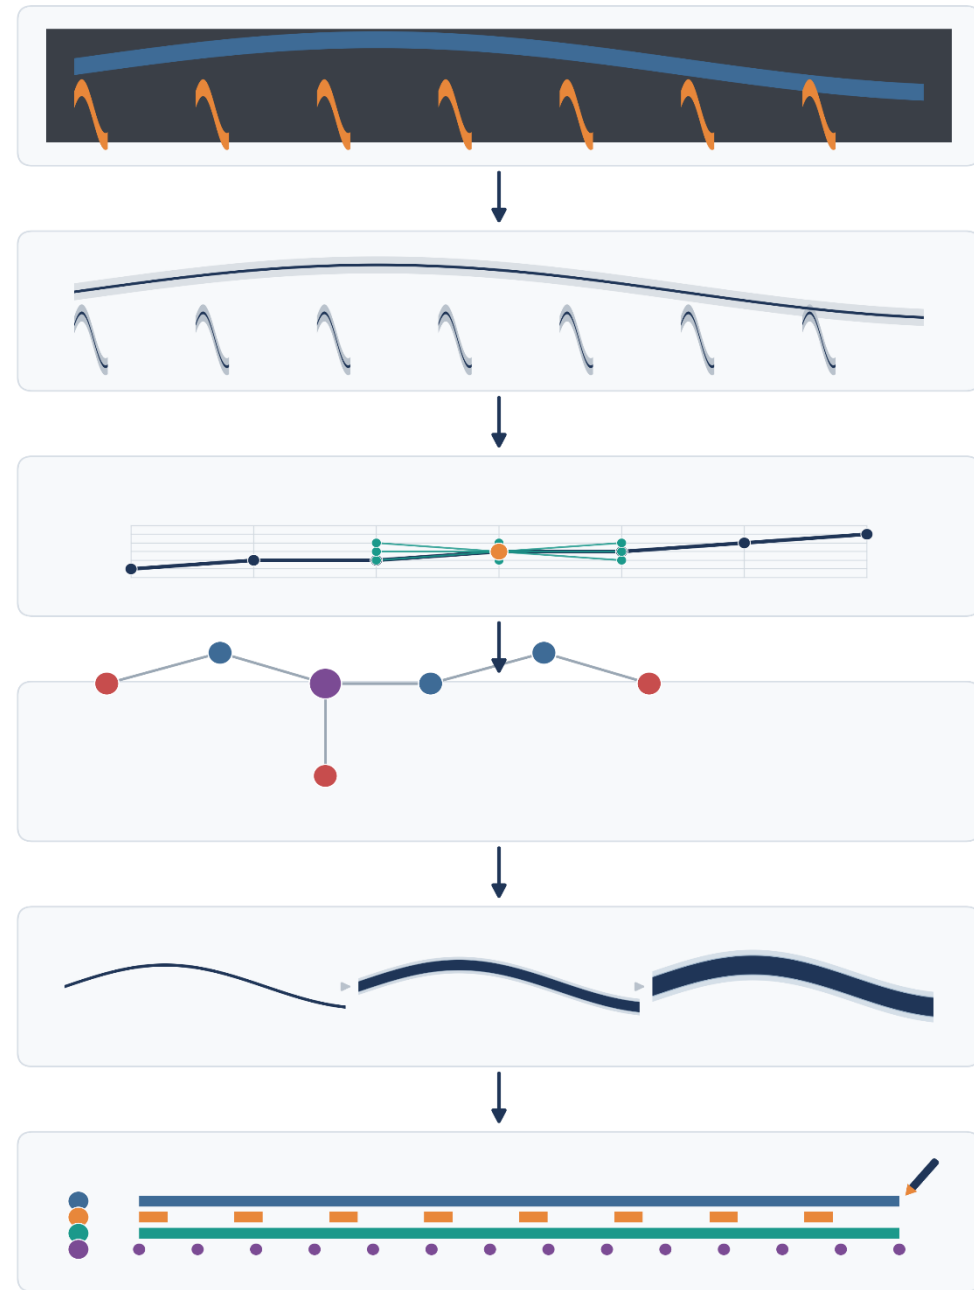


Methodology



Dataset Preparation

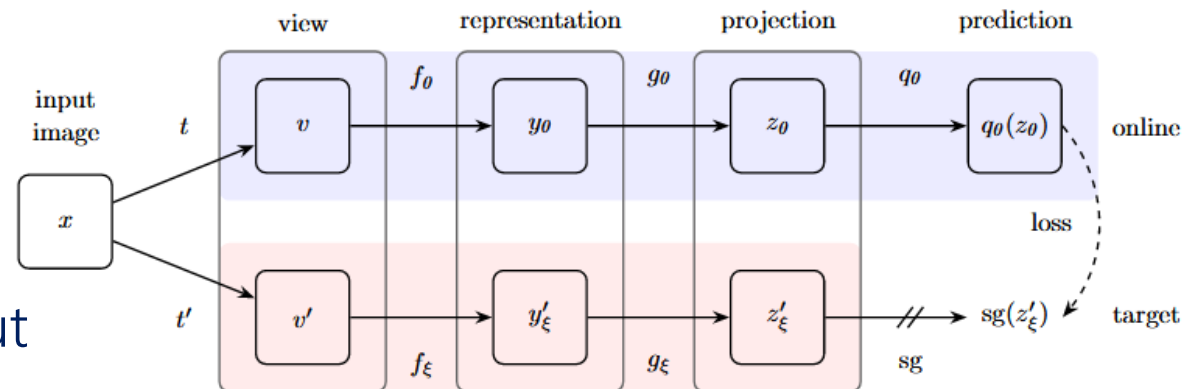
- **TuSimple dataset** (Yoo et al., 2020)
 - US Highways
 - 6,000 folders with at least 20 frames each
 - 1280 x 720 resolution size
- **Marking type classification**
 - Thinning
 - Graph structure from 8-level connectivity
 - Classification of separate entities by number of connections
 - Watershed to restore original thickness
 - Manual assignment of lane types



Models

➤ BYOL

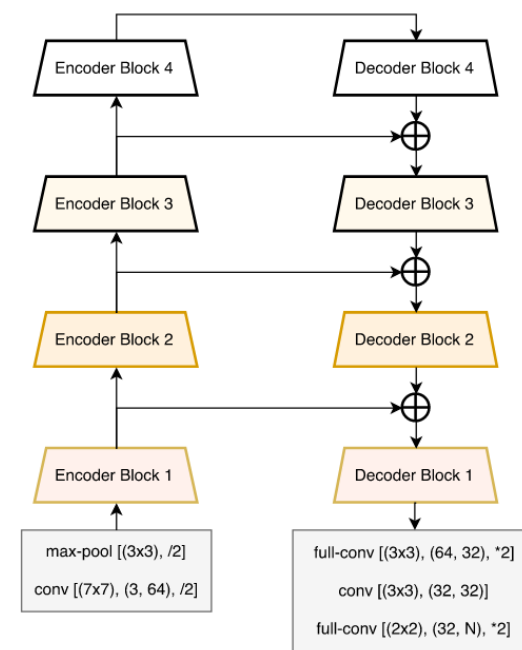
- Two differently augmented views per input
- Online branch (θ , training loss)
- Target branch (ξ , exponential moving average)
- Mean squared error between L2-normalized online prediction and target projection
- ResNet50 backbone
- 2-layer MLP with one-dimensional Batch Normalization for the projection and prediction layers



Source: Grill et al., 2020

➤ LinkNet

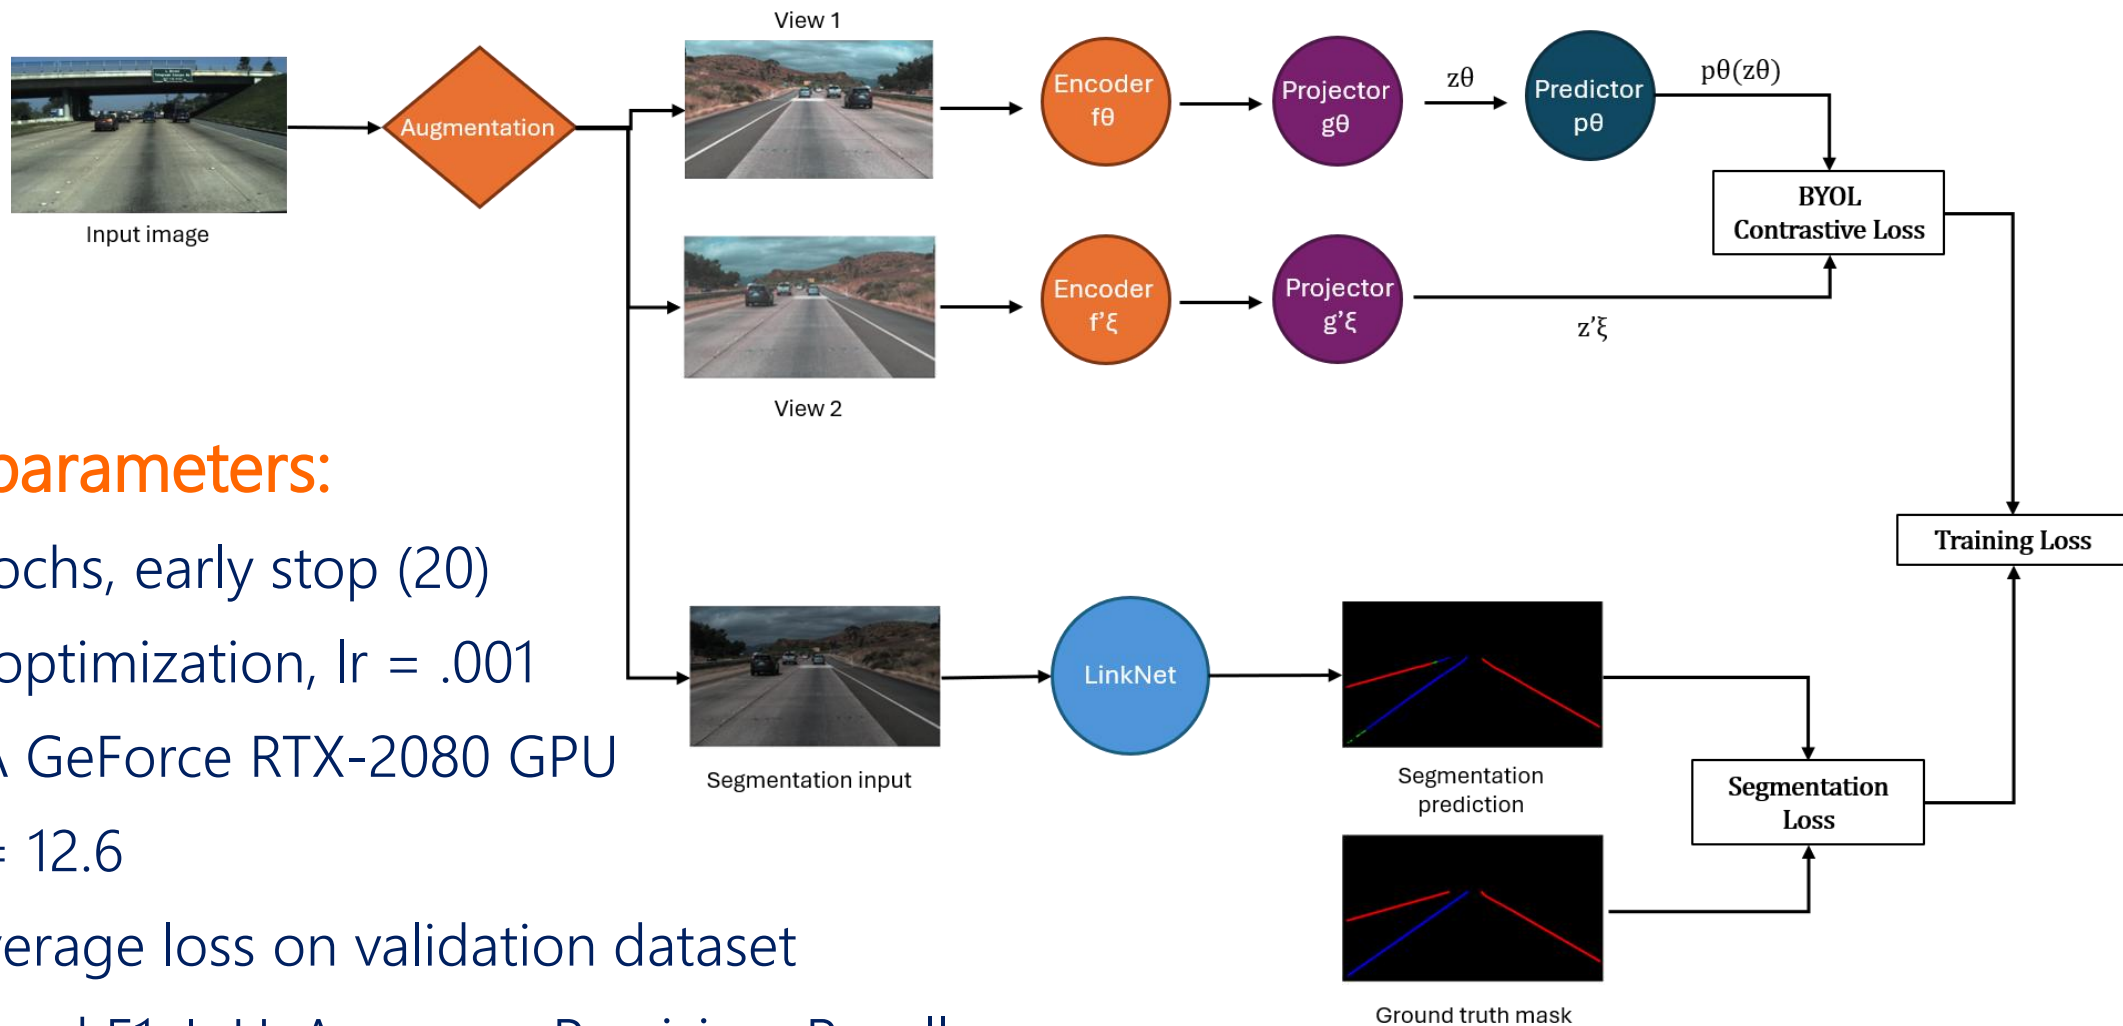
- ResNet50 backbone (He et al., 2016)
- Cross Entropy Loss



Source: Chaurasia and Culurciello, 2017



Training pipeline and performance monitoring



➤ Training parameters:

- 100 epochs, early stop (20)
- Adam optimization, lr = .001
- NVIDIA GeForce RTX-2080 GPU
- Cuda = 12.6
- Best average loss on validation dataset
- Monitored F1, IoU, Accuracy, Precision, Recall

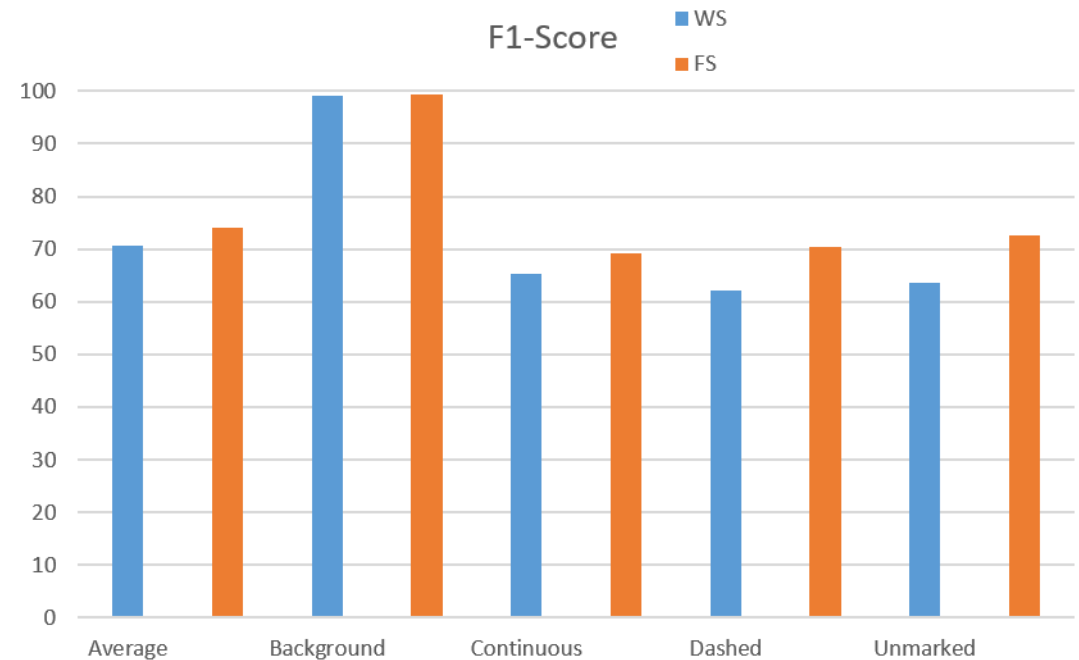



Results and discussion



End-to-End Training with Full Dataset

- 2h training per epoch
- Early stopping criteria after 87 epochs for WS, 20 epochs for FS
- Weak supervision underperforms Full Supervision



Fine-Tuning

- **Reduced dataset** – 0.4 of its original height and width
- **Augmentation:**
 - Horizontal Flip, Color Jitter, Convert to Gray, Gaussian Blur, Solarize
 - + Coarse Dropout (4 random squares of 16x16 pixels)
 - + Coarse Dropout (1 – 8 random squares, $p = 0.5$)

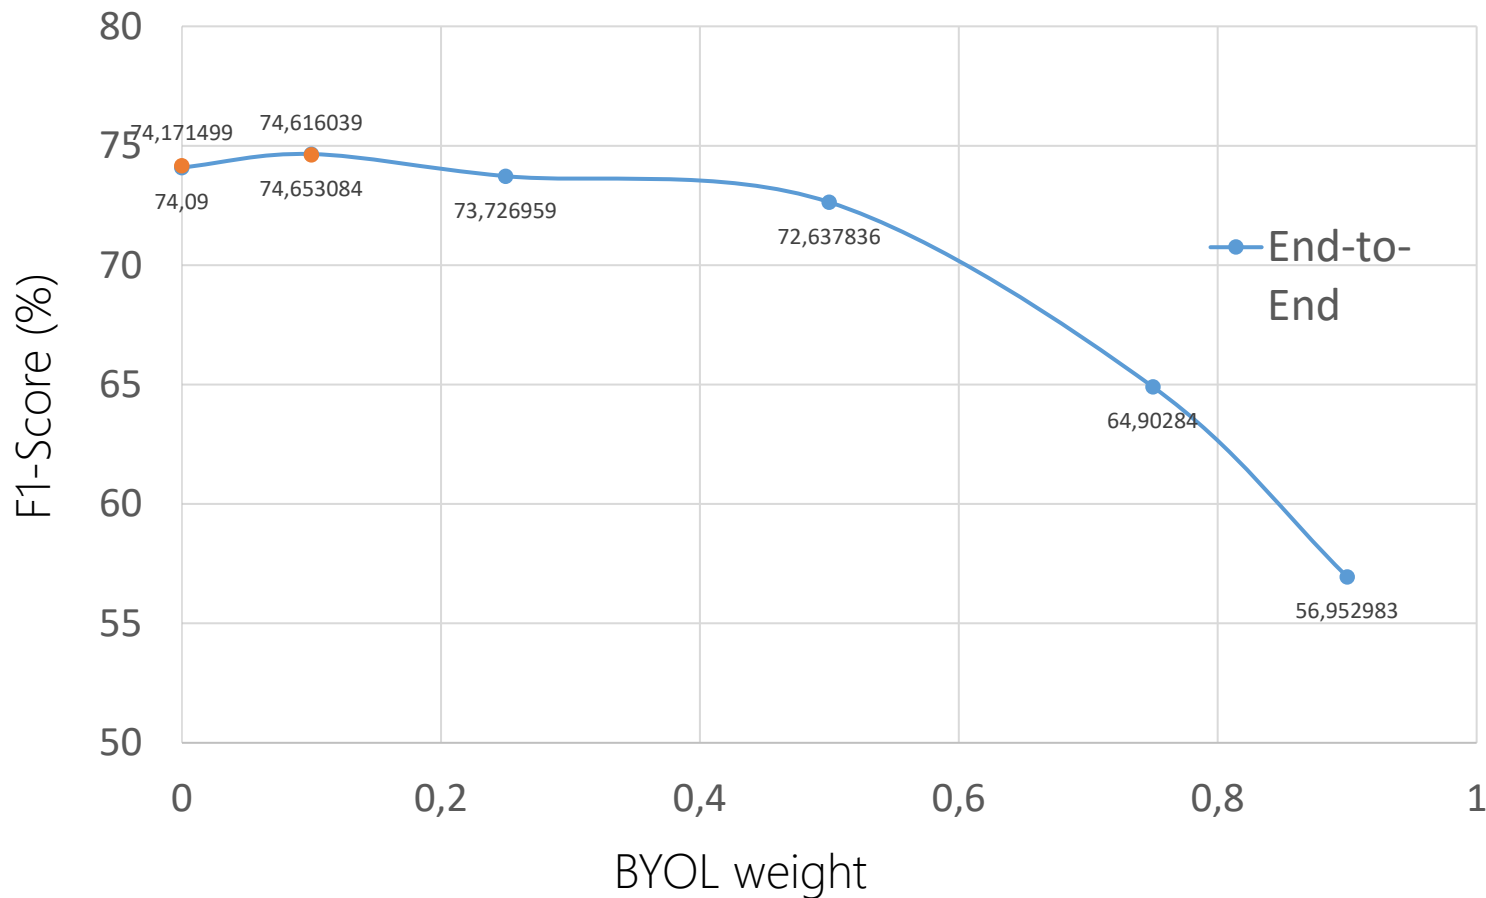
F1-Score per training configuration

Marking type	FS	WS	FT1	FT2	FT3
Average	70.8170	-22.14%	-0.14%	+1.98%	+2.03%
Background	99.1349	-0.21%	+0.00%	+0.03%	+0.03%
Continuous	63.2074	-22.36%	+2.13%	+3.58%	+3.16%
Dashed	63.0277	-41.38%	-1.57%	+5.34%	+4.63%
Unmarked	79.2775	-21.89%	+0.01%	+2.88%	+2.87%

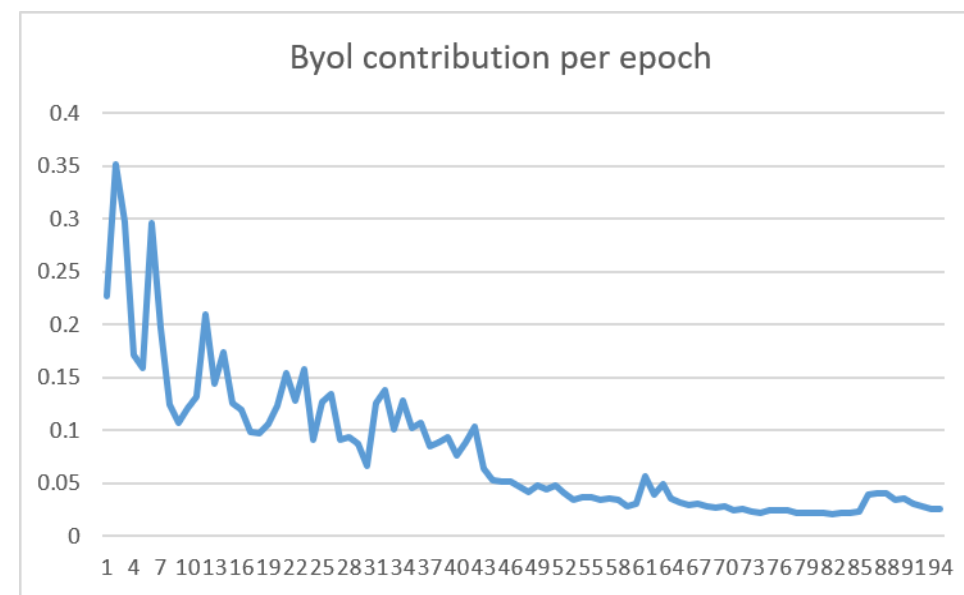


Hyperparameter tuning: loss weight

F1-Score per loss weight configuration

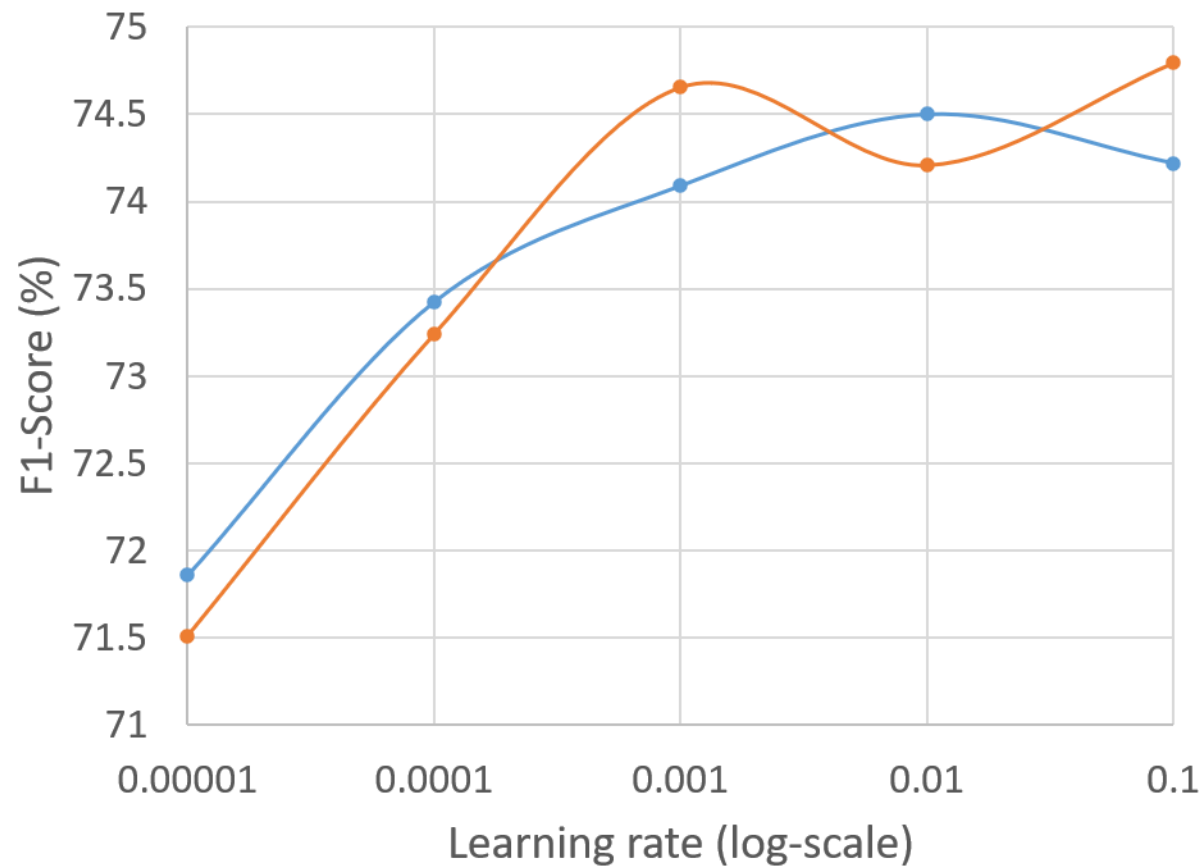


Byol contribution per epoch

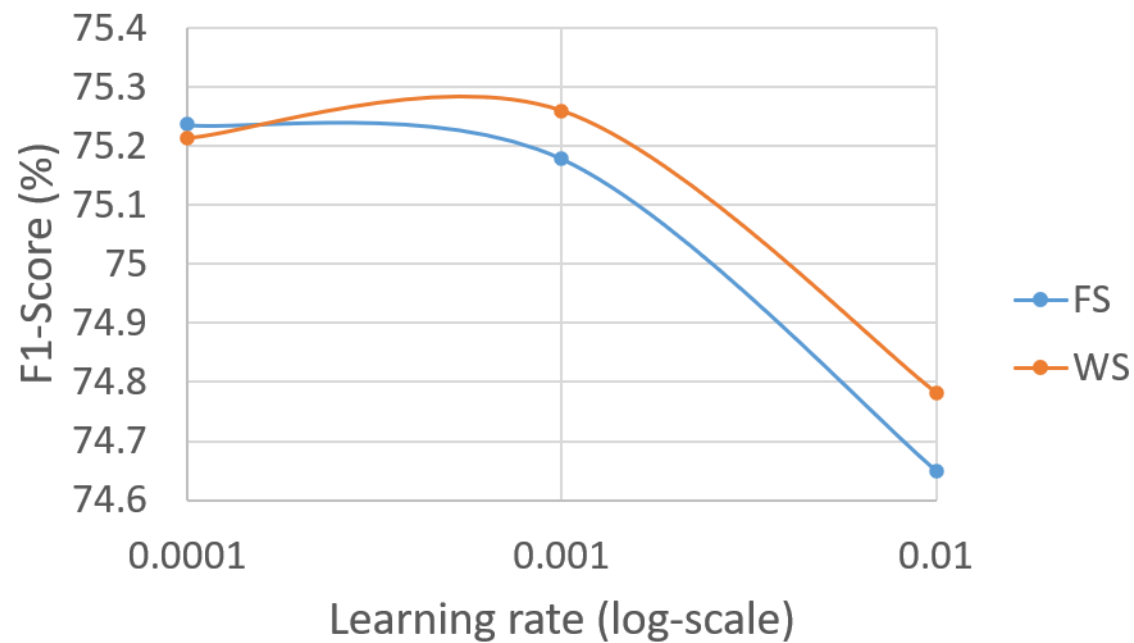


Hyperparameter tuning: learning rate

End-to-End Training

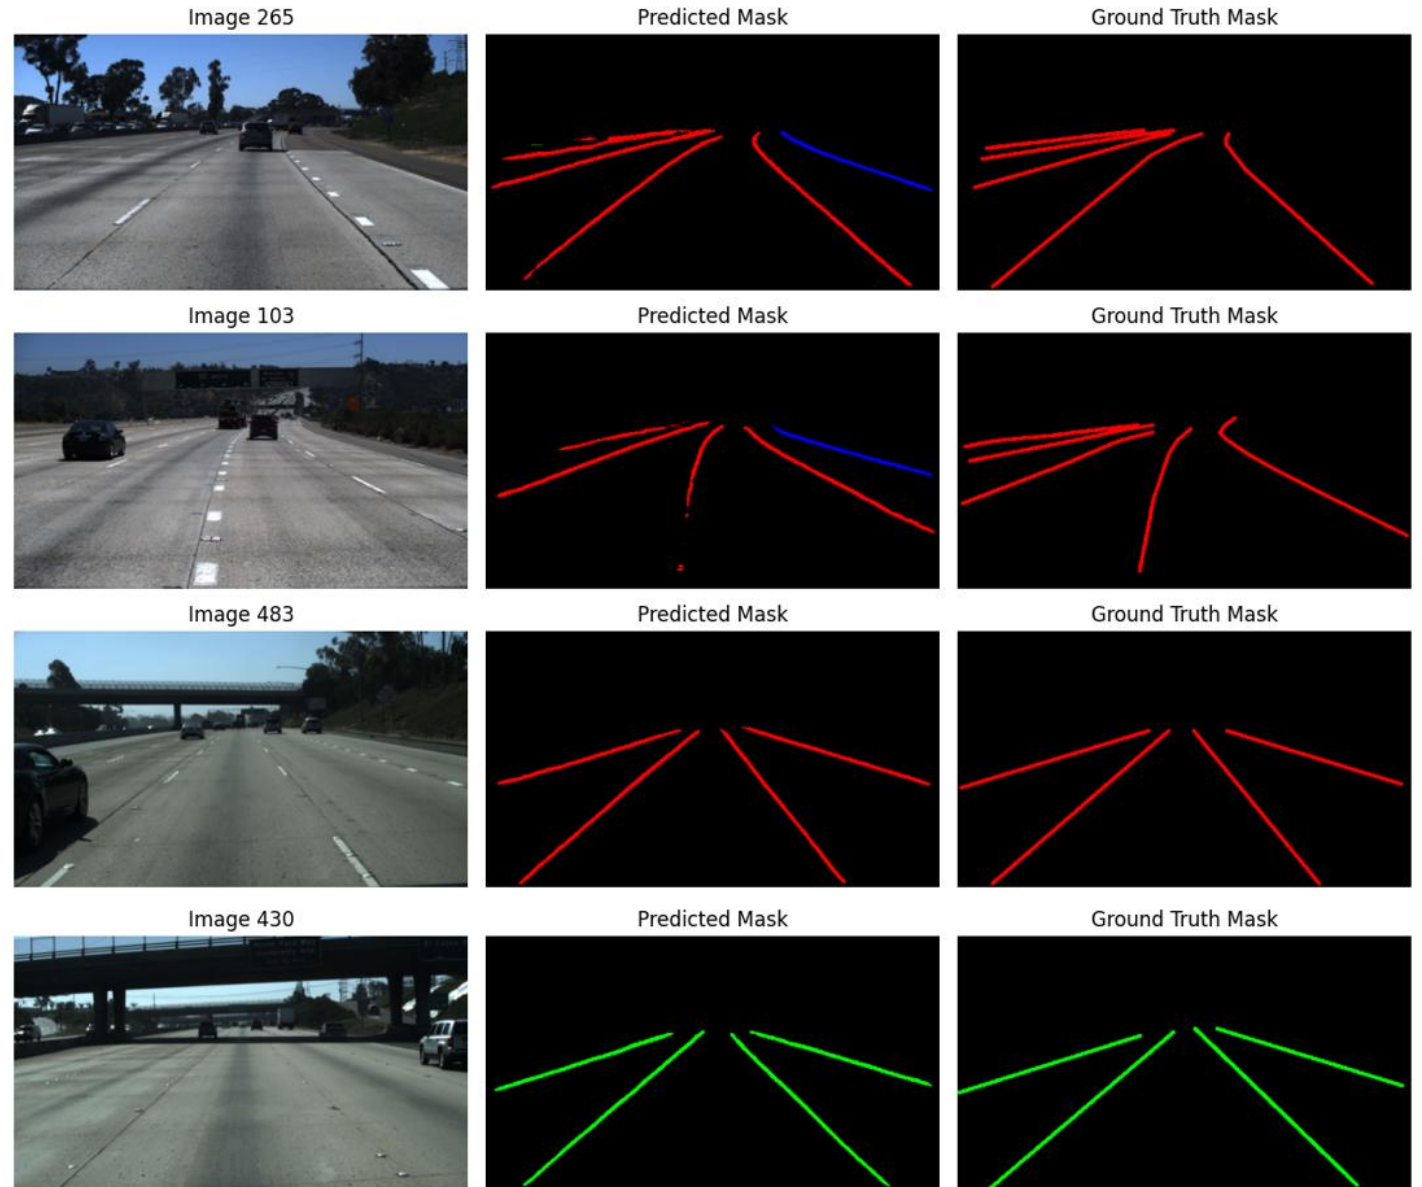


Fine Tuning




Visual results

- Results for **best performing model** (WS-FT)
- Optimal results under **good lighting conditions** and **linear shapes**
- **Labelling difficulties:**
 - Manual error
 - Lack of labels on original dataset



Visual results



- Fine-Tuning seems to improve both the lane detection and lane type identification



Comparison with State-of-the-Art

Reference	Dataset	Epochs	Lane Marking Type	F1-Score (%)	IoU (%)	Precision (%)	Recall (%)	Accuracy (%)
Ours (LinkNet with BYOL Fine-Tuning)	TuSimple	100	✓	74.74	62.44	77.73	72.40	99.20
LinkNet	TuSimple	100	✓	74.45	62.15	77.56	72.07	99.19
LaneCorrect (Nie et al., 2025)	TuSimple	?	×	92.91				91.95
CLLD (Zoljodi et al., 2024)	TuSimple	100	×					94.25
CLLD (Zoljodi et al., 2024)	CuLane	100	×	79.27		88.59	71.73	

➤ Comparable to CuLane-level results despite being **multiclass**



Conclusions, limitations and future work



Conclusion

- Dual-branch weakly supervised lane marking segmentation model
- Three class types:
 - Continuous
 - Dashed
 - Unmarked
- Comparable metrics with challenging benchmarks
 - **F1-Score = 0.75**




Limitations and Future Work

➤ **Limitations:**

- Dataset size
- Single training seed
- Pre-trained weights
- Different datasets

➤ **Future work:**

- Applications in transportation-related tasks
- Using the prediction to add semantic context for other attribute extraction tasks



World Conference on Transportation Research
Lectern Session A5-TS03A – AI/ML applications in Transportation: From
Infrastructure Sensing to User Trust
July 8, 2026

Weak supervision and fine-tuning with contrastive learning for multiclass lane marking segmentation

Julia Alves Porto,
PhD Candidate



Department of Transportation Planning and Engineering,
National Technical University of Athens