SAND2020-7055C

Multimodal Data Fusion Via Entropy Minimization

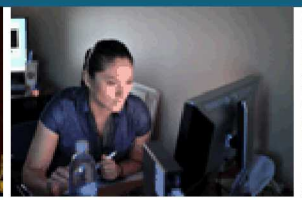

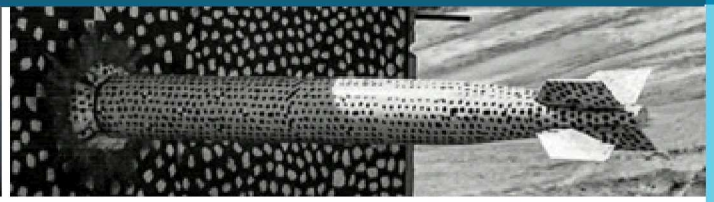

Joshua Michalenko, Lisa Linville, Dylan Z. Anderson

SAND #:

Unlimited Unclassified Release

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Multimodal Data Fusion

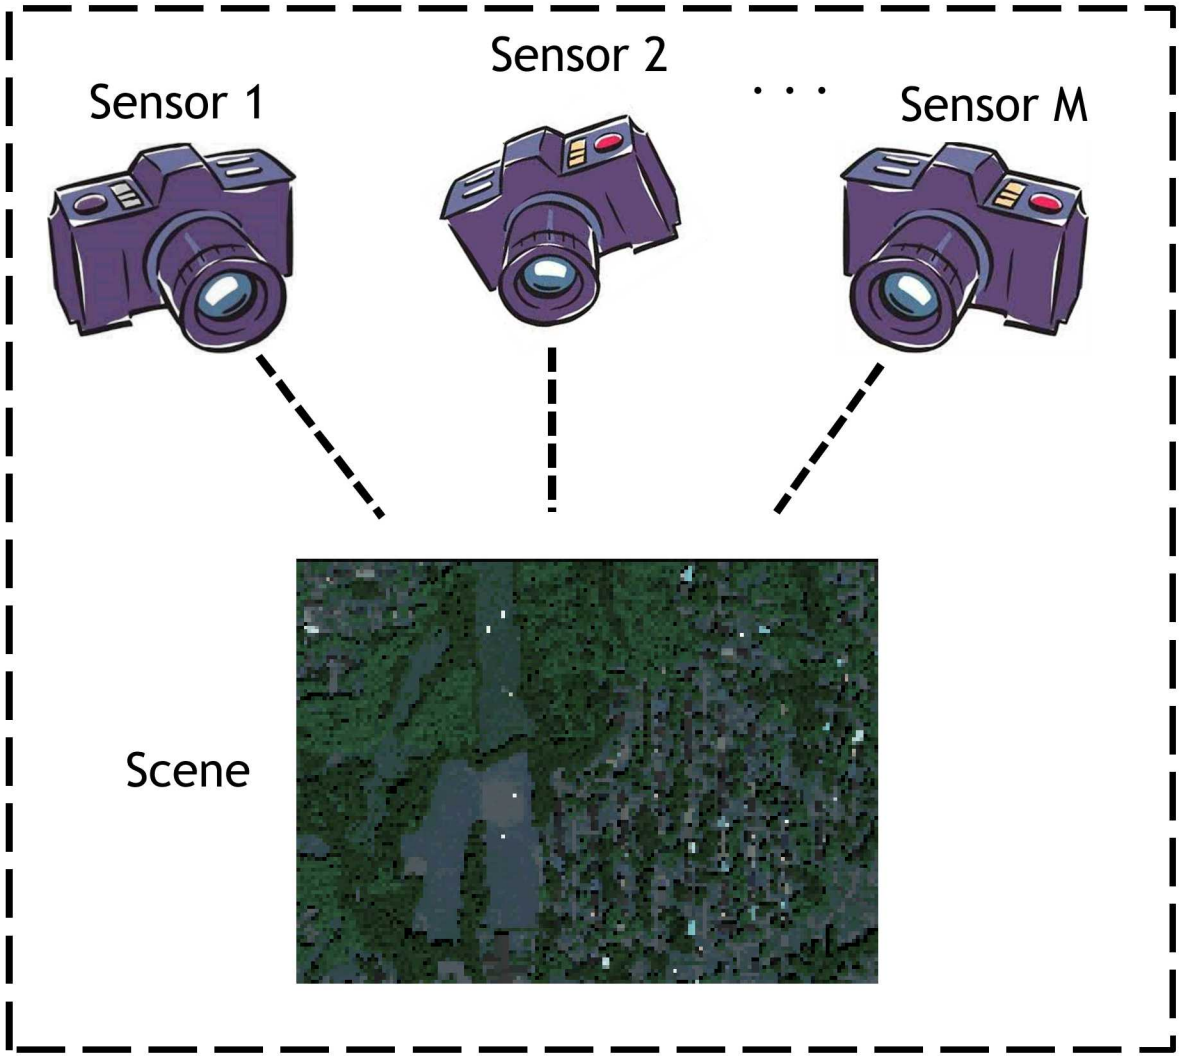

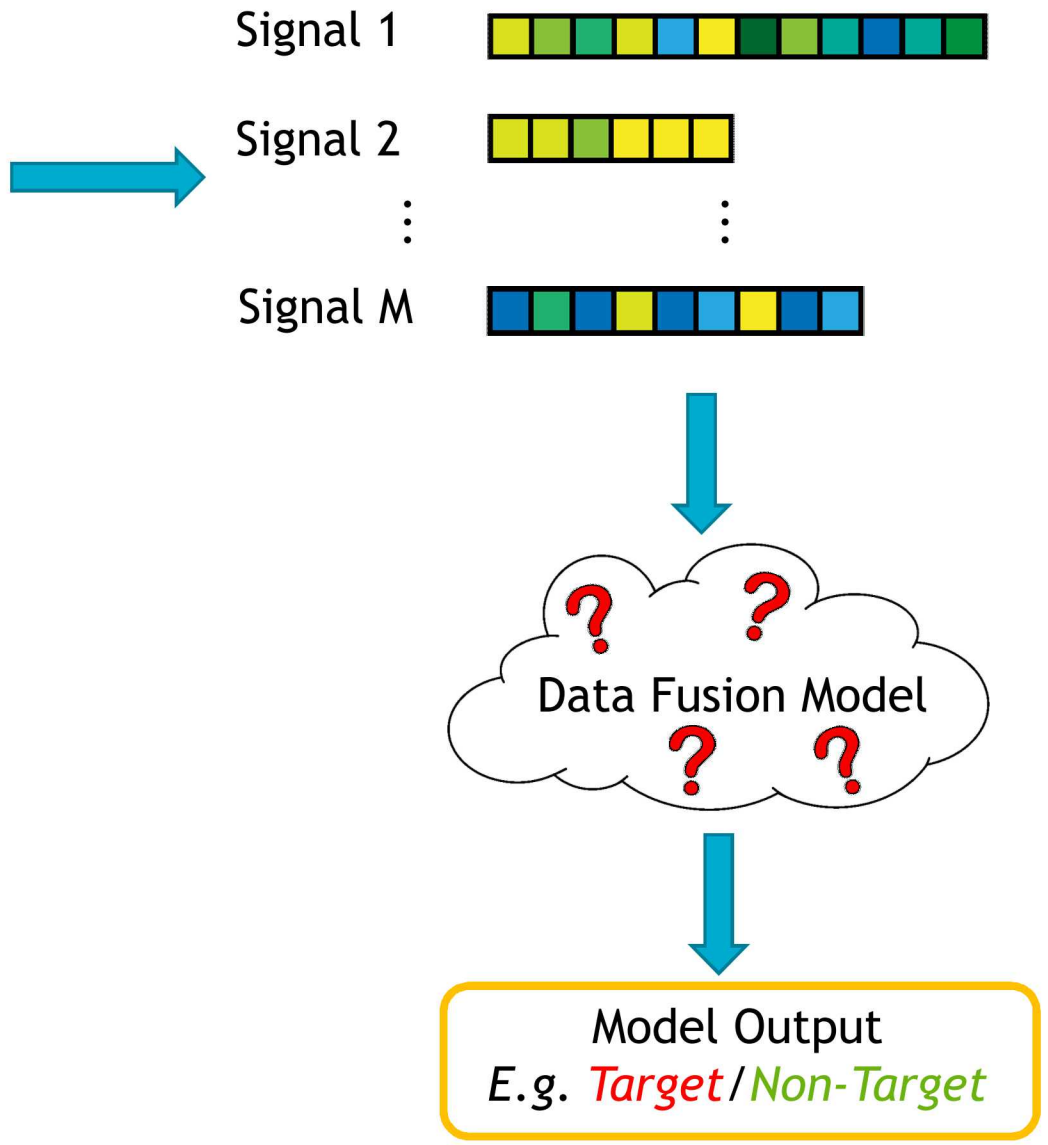

Data Fusion Models

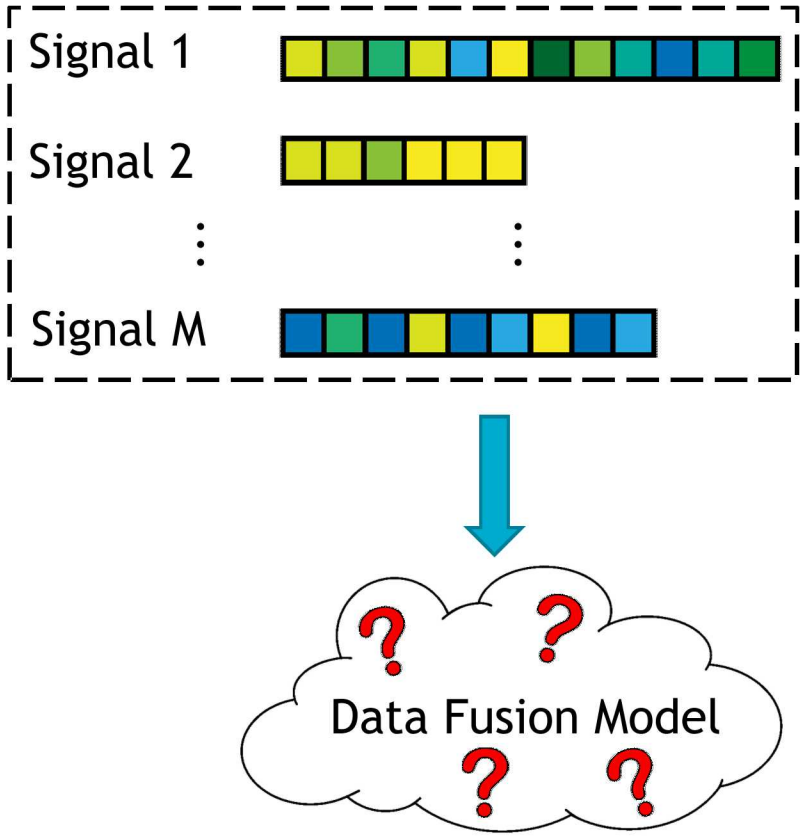

Competing Data Fusion Models

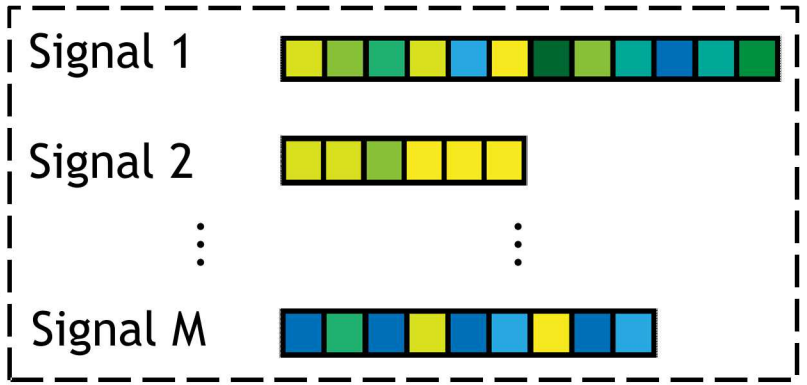

Fused Model Input (FMI)

- Concatenate Inputs

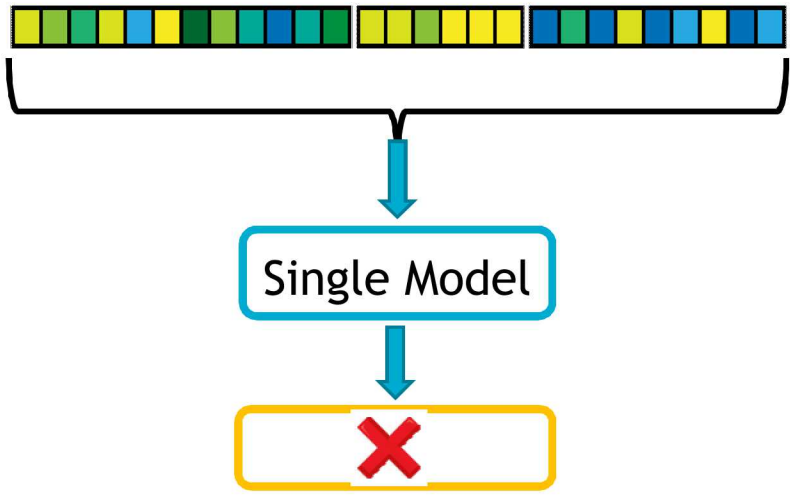

Co-information sharing

Unable to handle missing data

Fused Model Output (FMO)

- Train Ind. Models, Fuse outputs

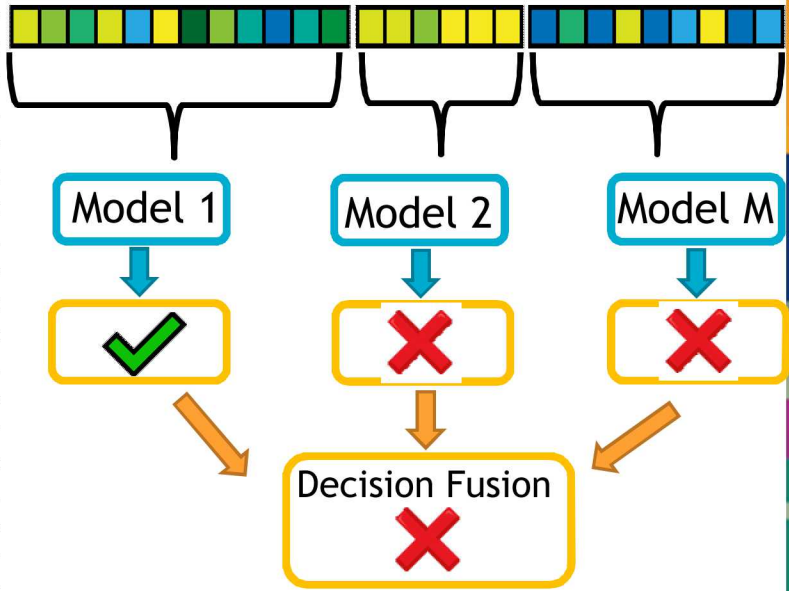

Can handle missing data

No co-information sharing

Competing Data Fusion Models

Fused Model Input (FMI)

- Concatenate Inputs

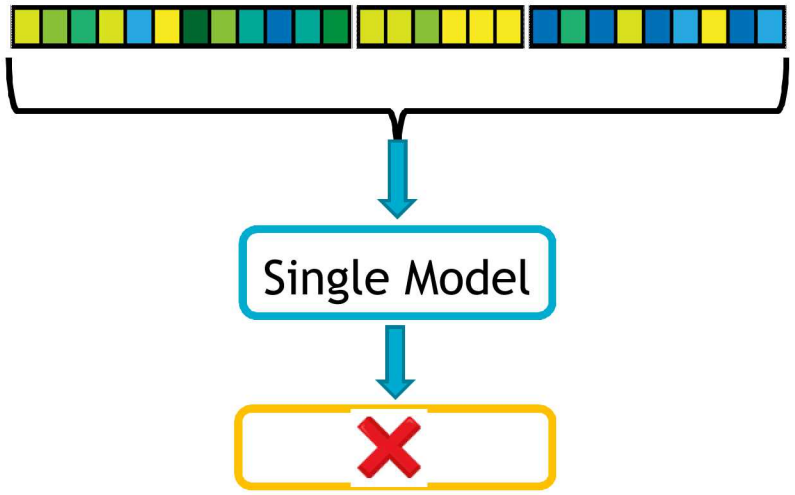

Co-information sharing

Unable to handle missing data

Fused Model Output (FMO)

- Train Ind. Models, Fuse outputs

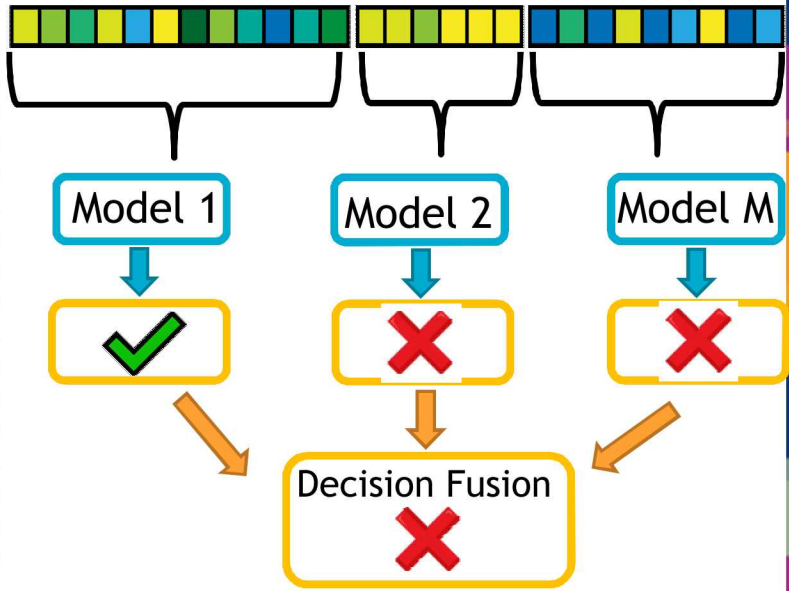

Can handle missing data

No co-information sharing

Competing Data Fusion Models

Fused Model Input (FMI)

- Concatenate Inputs

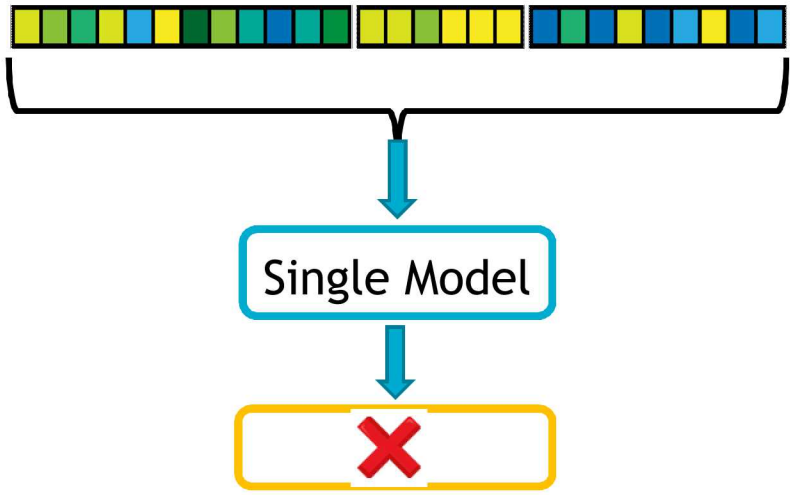

Co-information sharing

Unable to handle missing data

Entropy Minimization (EMIN)

- Train Ind. Models, Fuse outputs

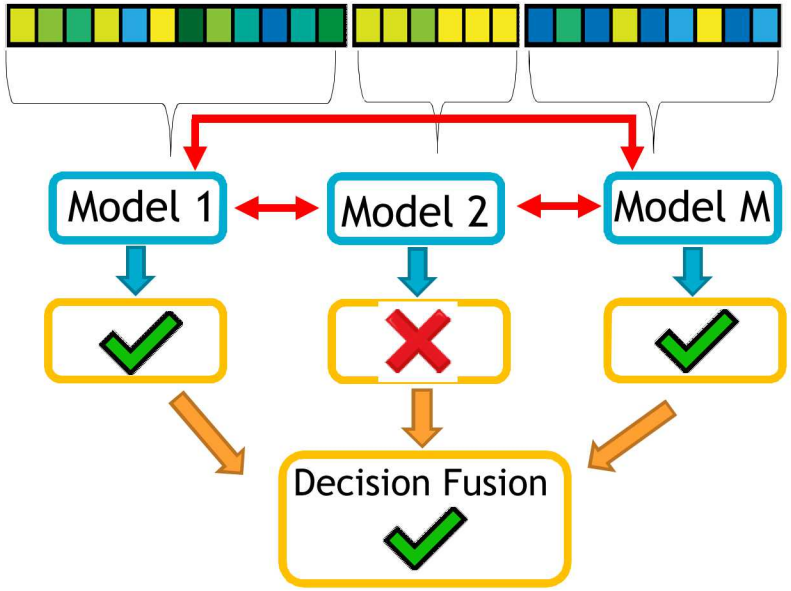

Co-information sharing

Can handle missing data

Fused Model Output (FMO)

- Train Ind. Models, Fuse outputs

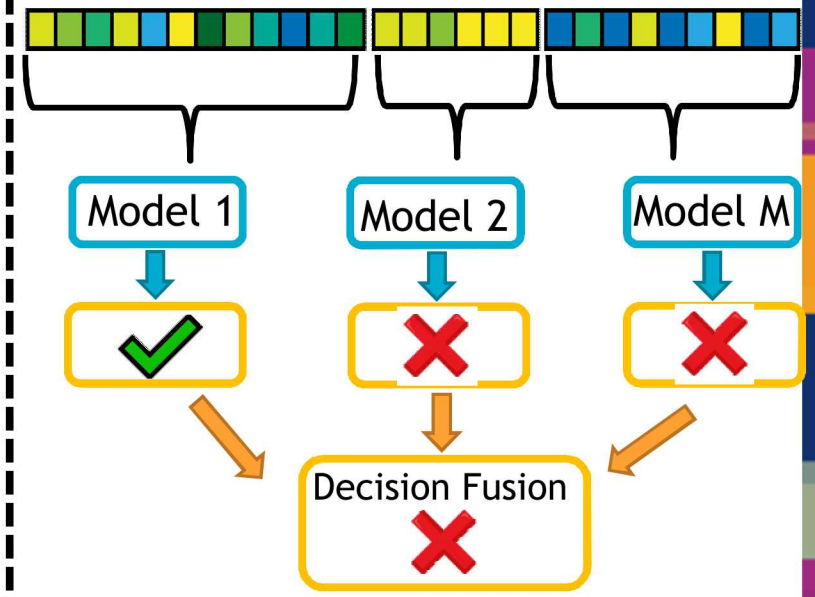

Can handle missing data

No co-information sharing

Entropy Minimization

- Entropy is a measure of the ‘uncertainty’ in a random variable

$$H(\mathbf{X}) = -\sum_{i \in \mathscr{C}} p_i \log(p_i)$$

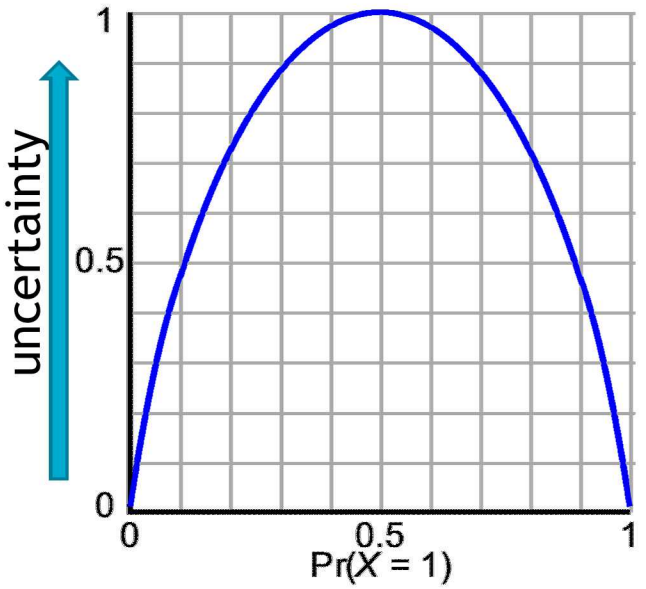

- Link individual models together during training by minimizing the entropy of their outputs for a given input

$$p(\mathbf{X}) = \frac{\sum_{m \in [M]} p(y|x^m)}{M}$$

Use Entropy of the averaged distribution as a conjugate loss term

$$\mathscr{L}(y, \hat{y}) = \left(\sum_{m \in [M]} \mathrm{CE}(y, \hat{y_m}) \right) + \gamma H(\mathbf{X})$$

Megascene Testbed

Large simulated hyperspectral scene using DIRSIG model

2 Sensors with ‘varying’ spectral resolutions

- AVIRIS-like sensor synthetically split
- Very Near InfraRed (VNIR)
  - 0.4-0.9 um 10nm resolution
- Short Wave InfraRed (SWIR)
  - 0.9-2.5 um 10 nm resolution

Green Discs inserted as targets

- ~3.5 million pixels
- <.1% targets

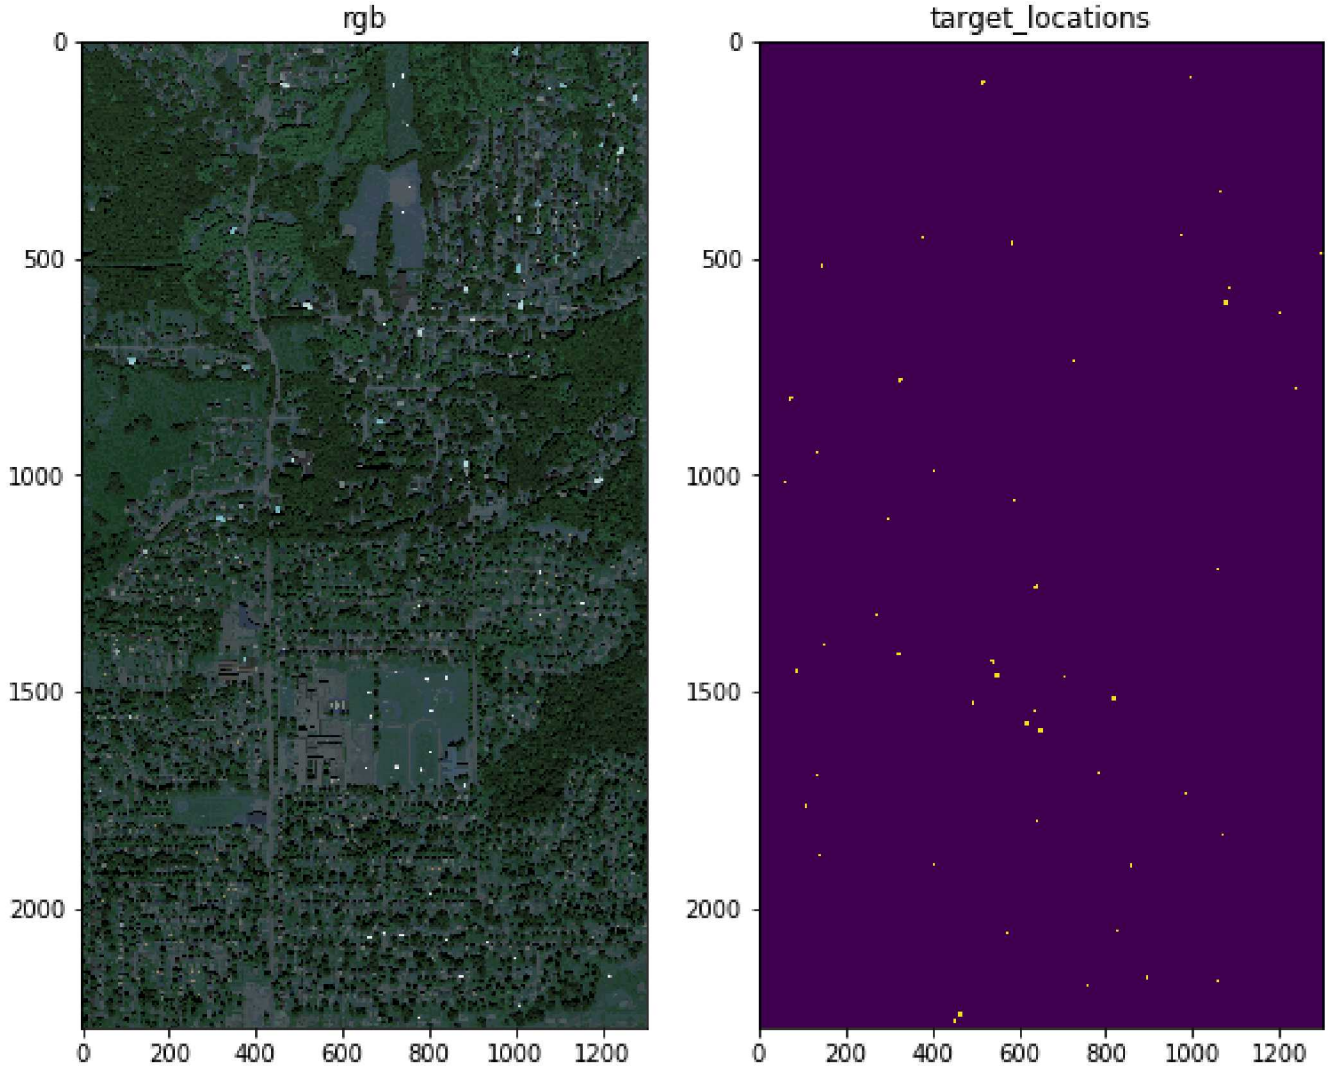

Megascene Testbed

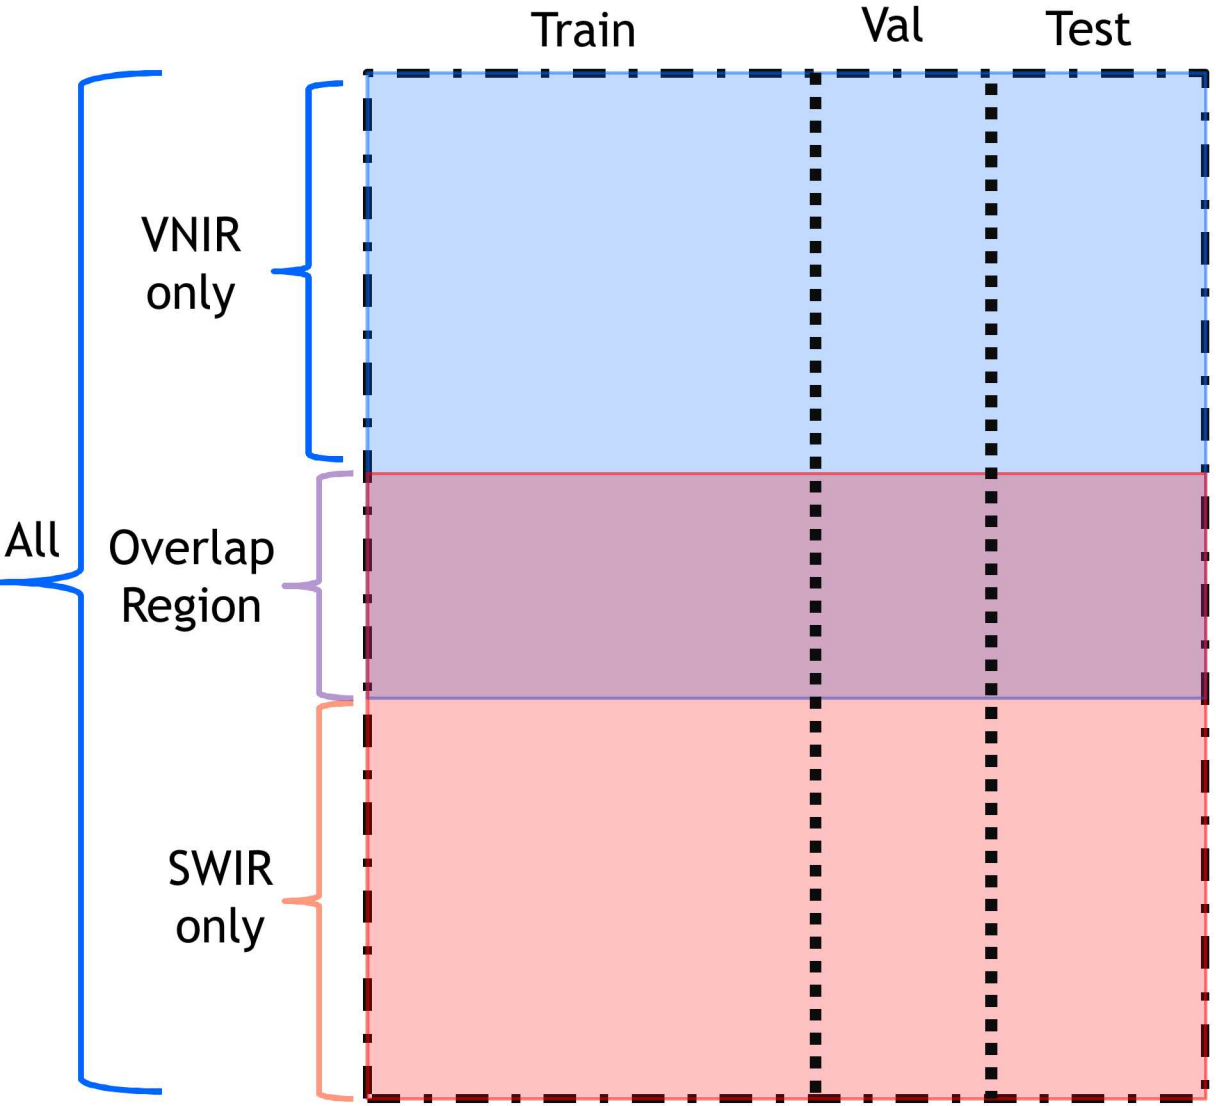

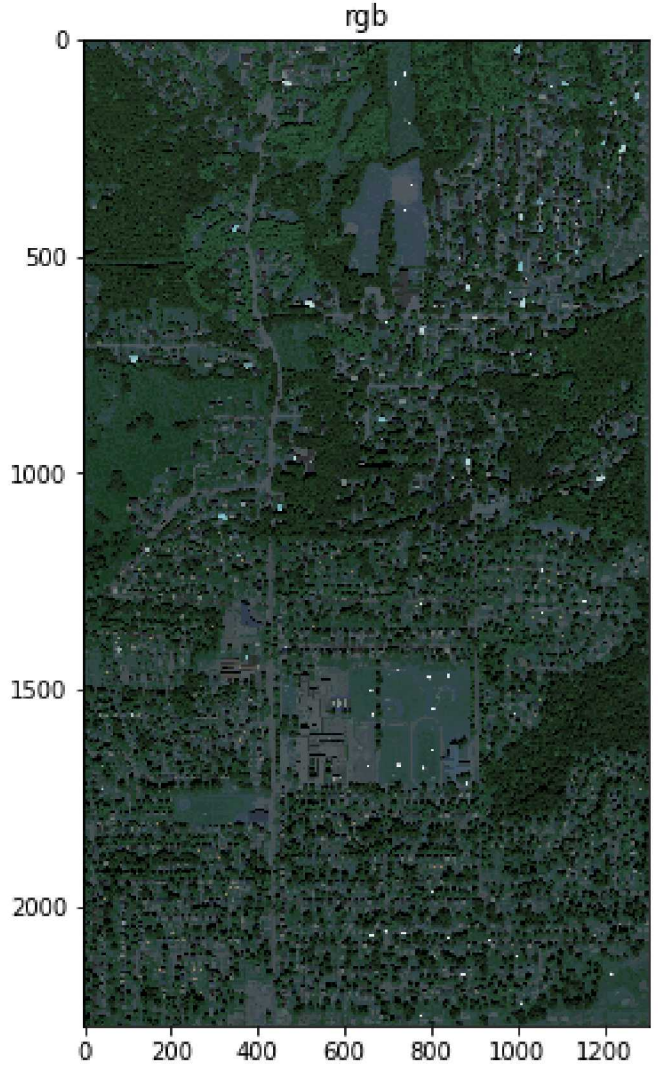

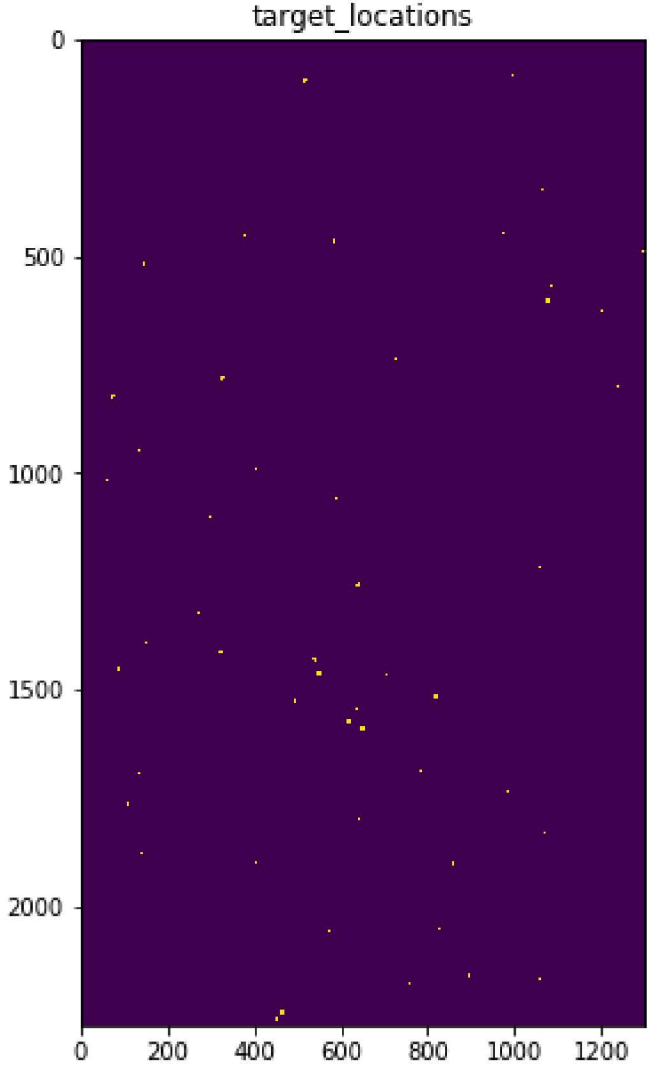

EMIN performance W.R.T. Data Availability

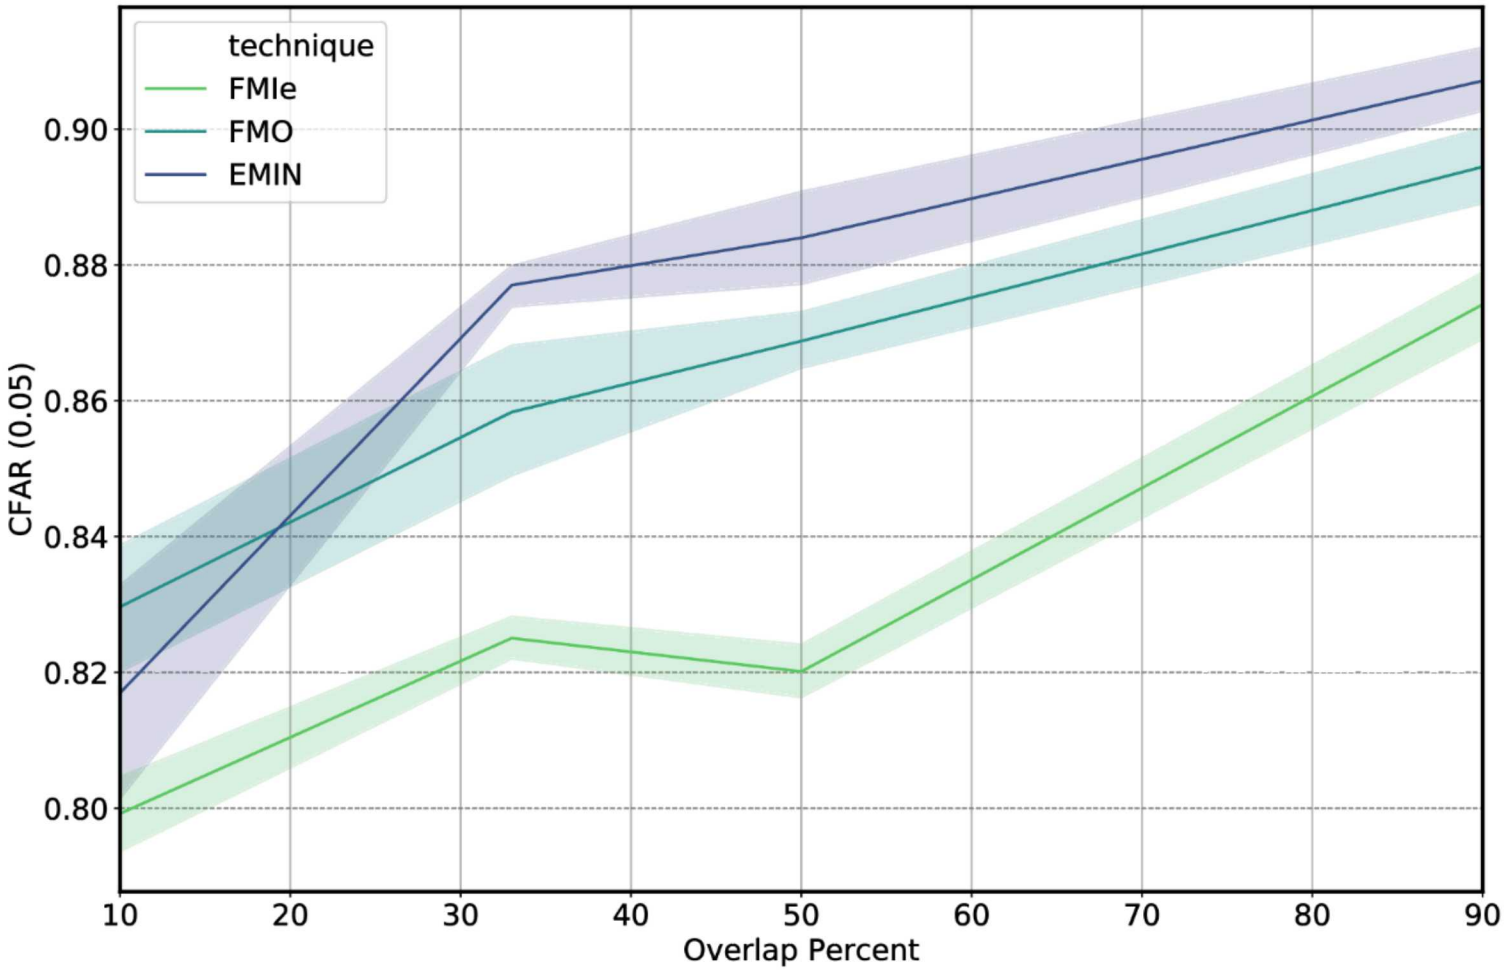

EMIN performance W.R.T. Data Availability

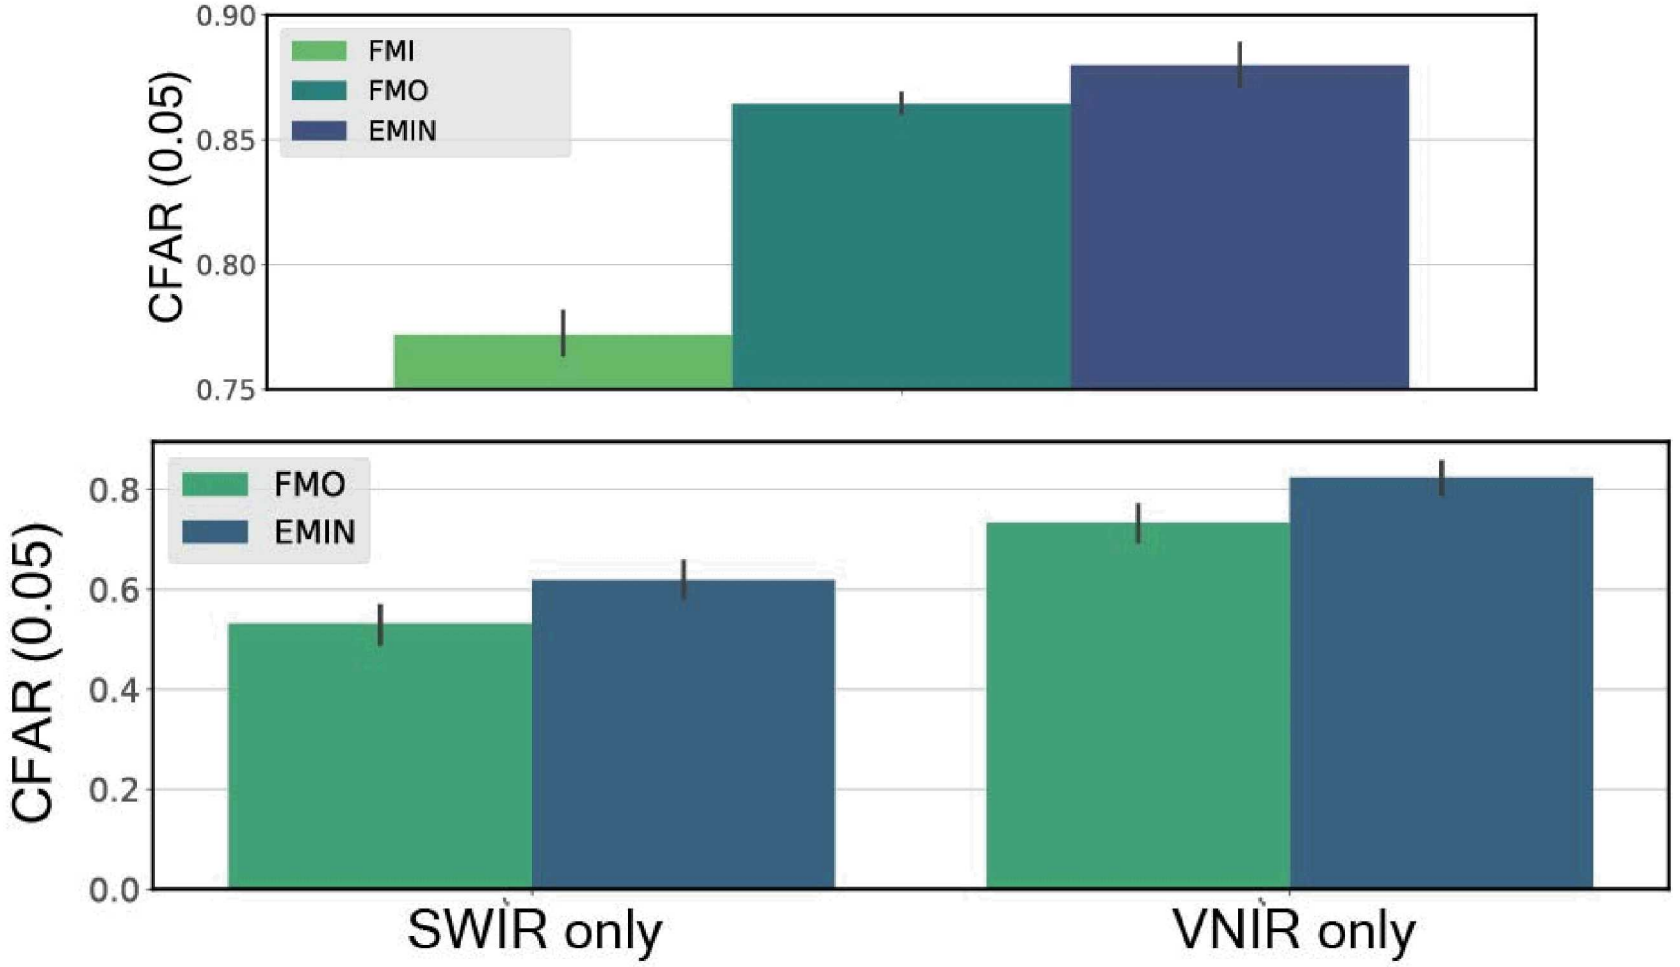

EMIN shows increases in performance for individual models

EMIN Performance W.R.T. Sensor Resolution

Decreasing SWIR Sensor Resolution

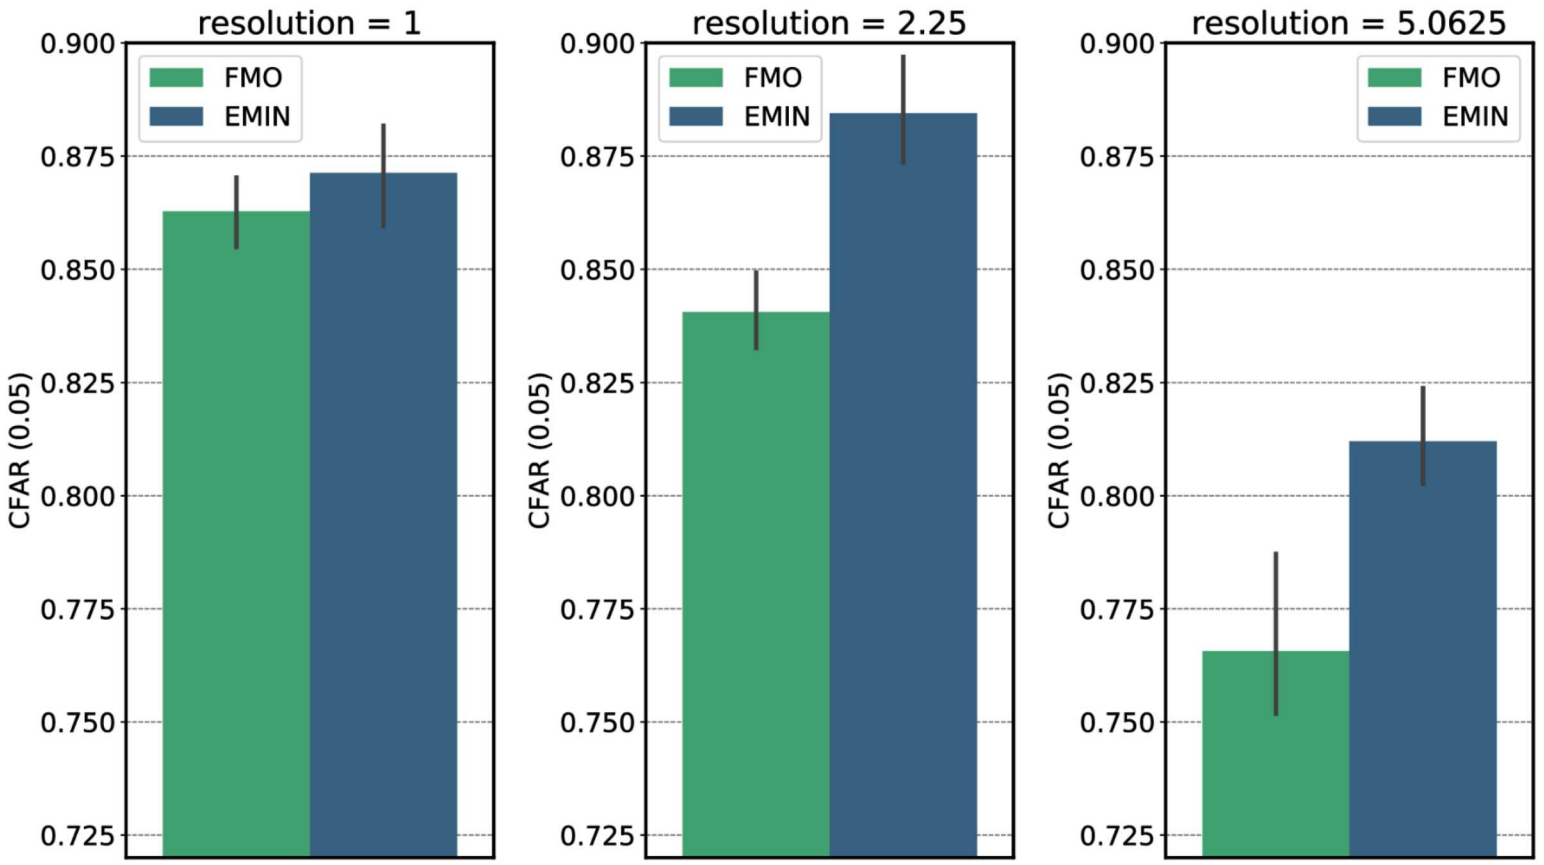

Performance increase from EMIN becomes larger with more realistic operating conditions

Conclusions

- We outline a new method for multimodal data fusion based on minimizing entropy
  - Method is flexible to any model with a gradient based update scheme
  - Hypothesize co-information sharing allows for higher performing individual models
- Using a synthetic target detection dataset, EMIN outperforms FMO and FMI while maintaining flexible inference
- Next steps:
  - More extensive characterization of EMIN effects
  - Real world datasets
  - Sensor networks of scale