ML4Seismic Partners Meeting 2023

Exploiting Structures of Data for Application Specific Representations

Kiran Kokilepersaud, Mohit Prabhushankar, and Ghassan AlRegib

OLIVES @GeorgiaTech

Georgia Tech

Introduction
Deep Learning is Trending Towards Large Generalized Models

Large Generalized Model – *One model* ***should*** *transfer to* ***any task***

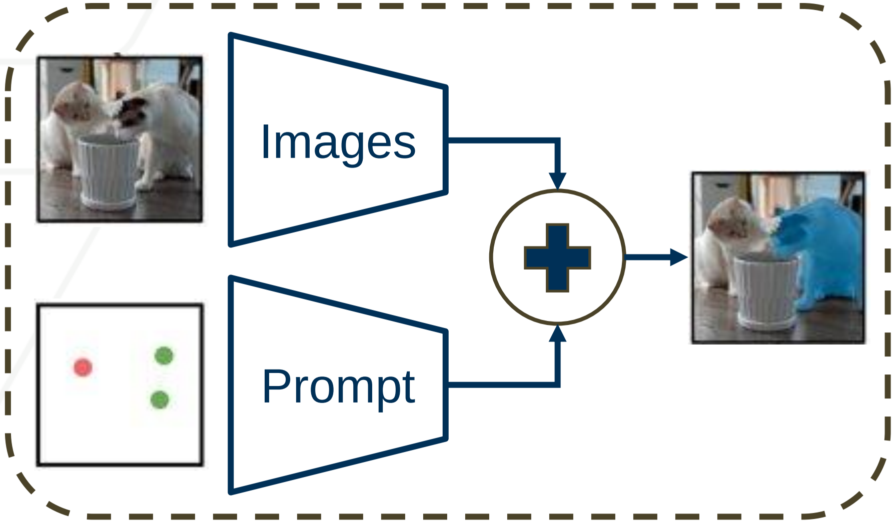

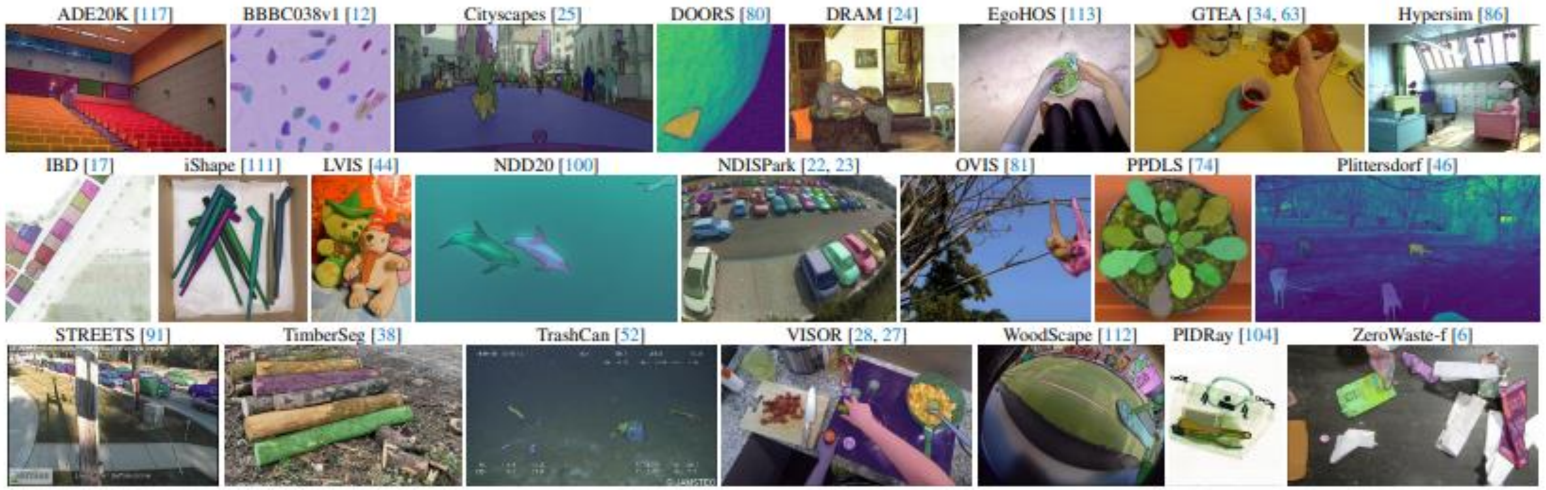

Figure 8: Samples from the 23 diverse segmentation datasets used to evaluate SAM's zero-shot transfer capabilities.

Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., ... & Girshick, R. (2023). Segment anything. *arXiv preprint arXiv:2304.02643*

Introduction

Generalized Models Rely on the Quality of the Produced Representation

Representation Learning – “Creating abstractions of data that enable ***extraction of useful information*** for a ***target downstream task***.”

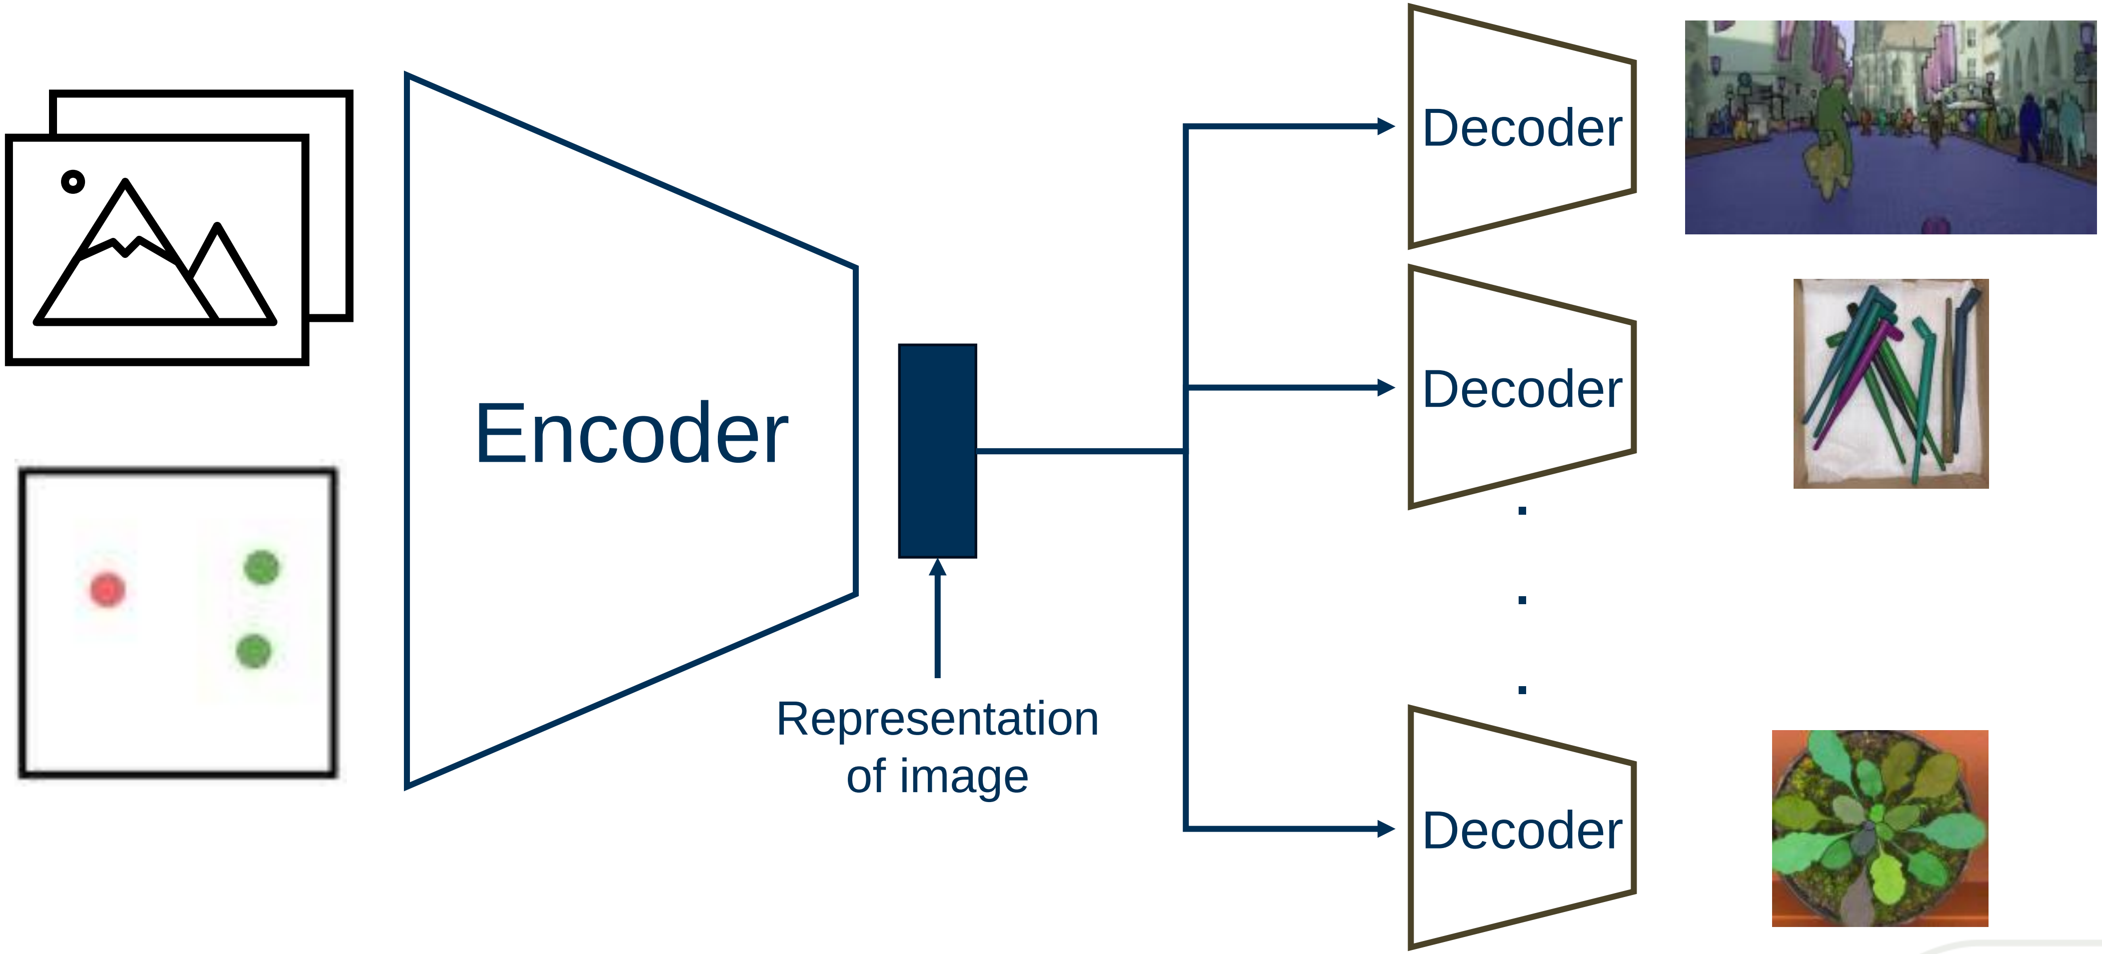

Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, *35*(8), 1798-1828.

Introduction
Generalized Representations Oftentimes do not Perform as Intended

Seismic Data

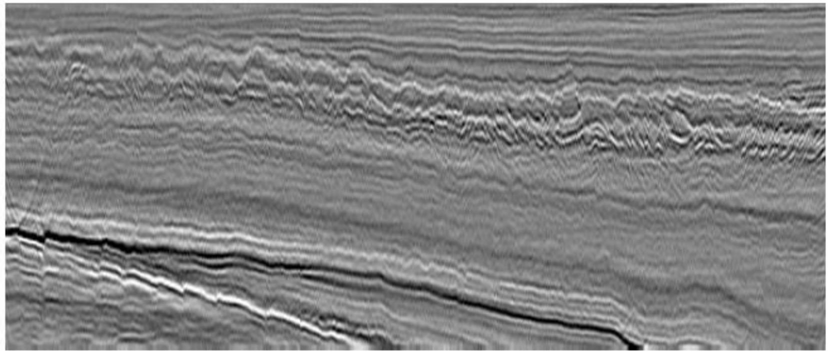

Fisheye Data

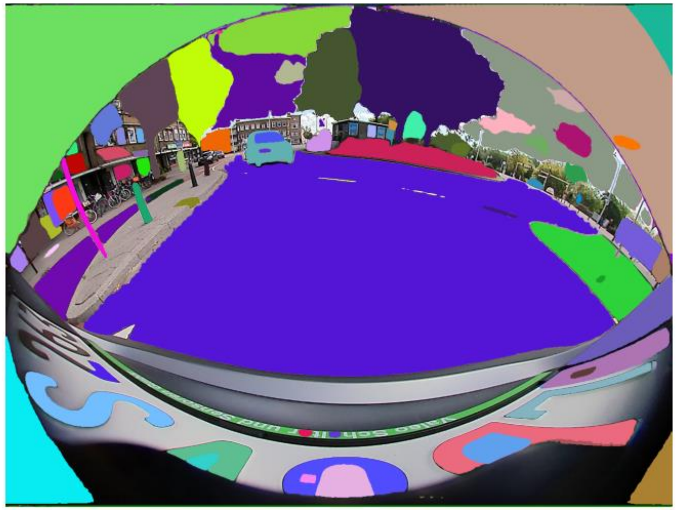

What is a Good Representation?

Bengio et. al – “captures the posterior distribution of the ***underlying explanatory factors*** for the observed input.”

Explanatory Factors = Any component of the data distribution that **results in variation between samples**.

Low Level Factors =
Random Augmentation

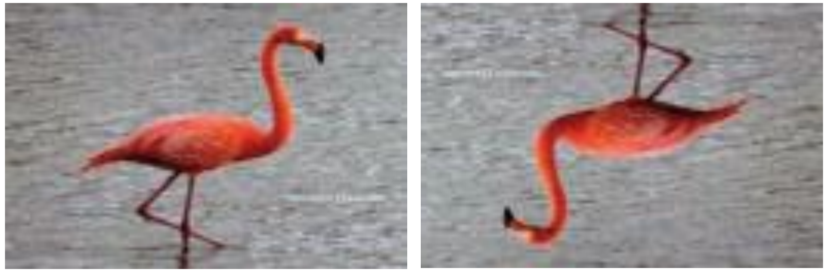

High Level Factors =
Taxonomical Classes

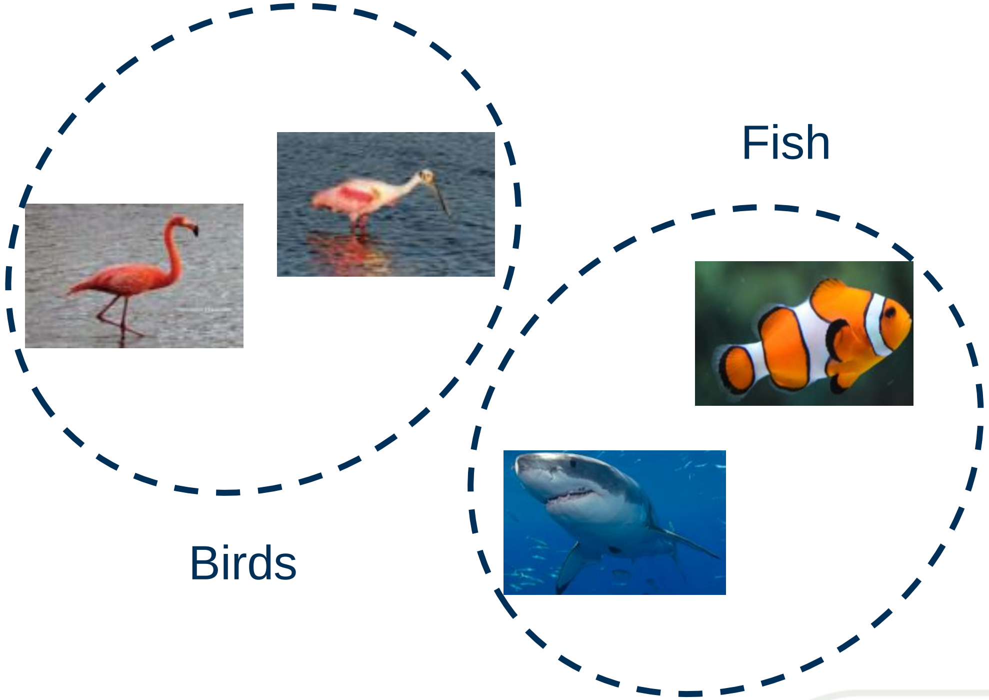

Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, *35*(8), 1798-1828.

Introduction

Traditional Representation Learning does not Incorporate Explanatory Factors

Contrastive learning is one popular representation learning approach.

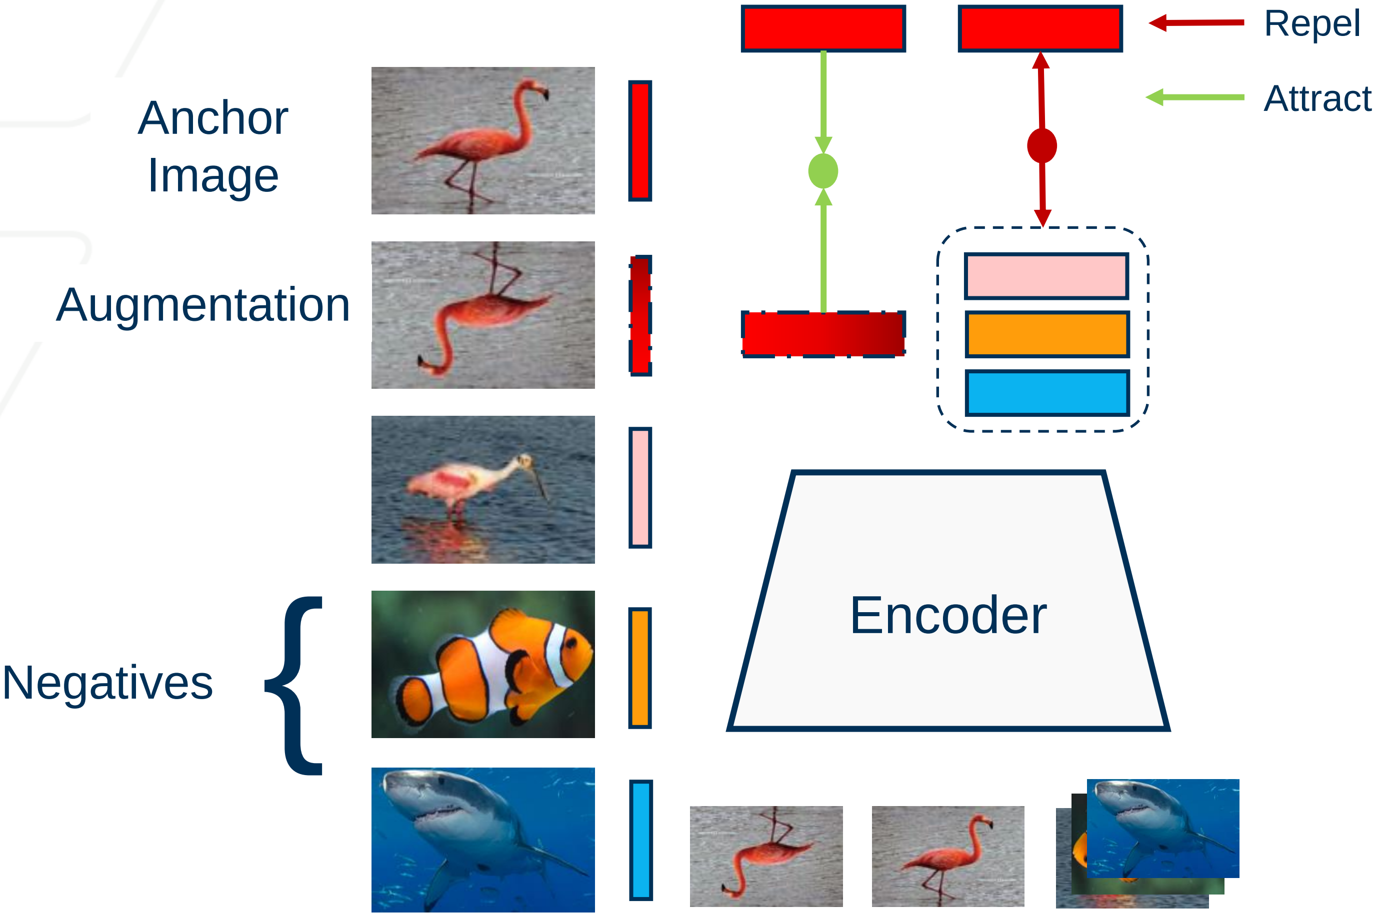

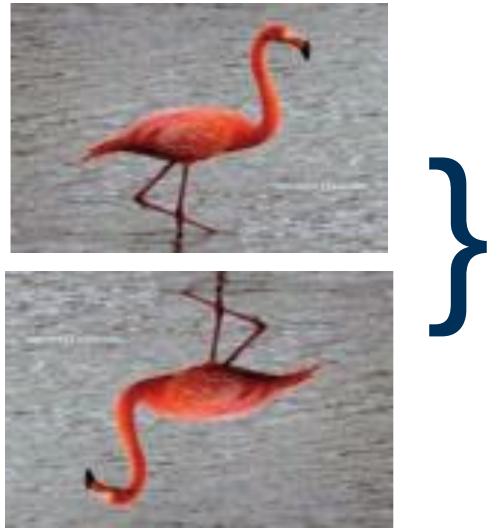

Invariance to augmentation

Higher level factors can potentially **better inform representations**.

Introduction

What are Higher Level Representational Factors?

Higher level explanatory factors exist in a variety of different domains.

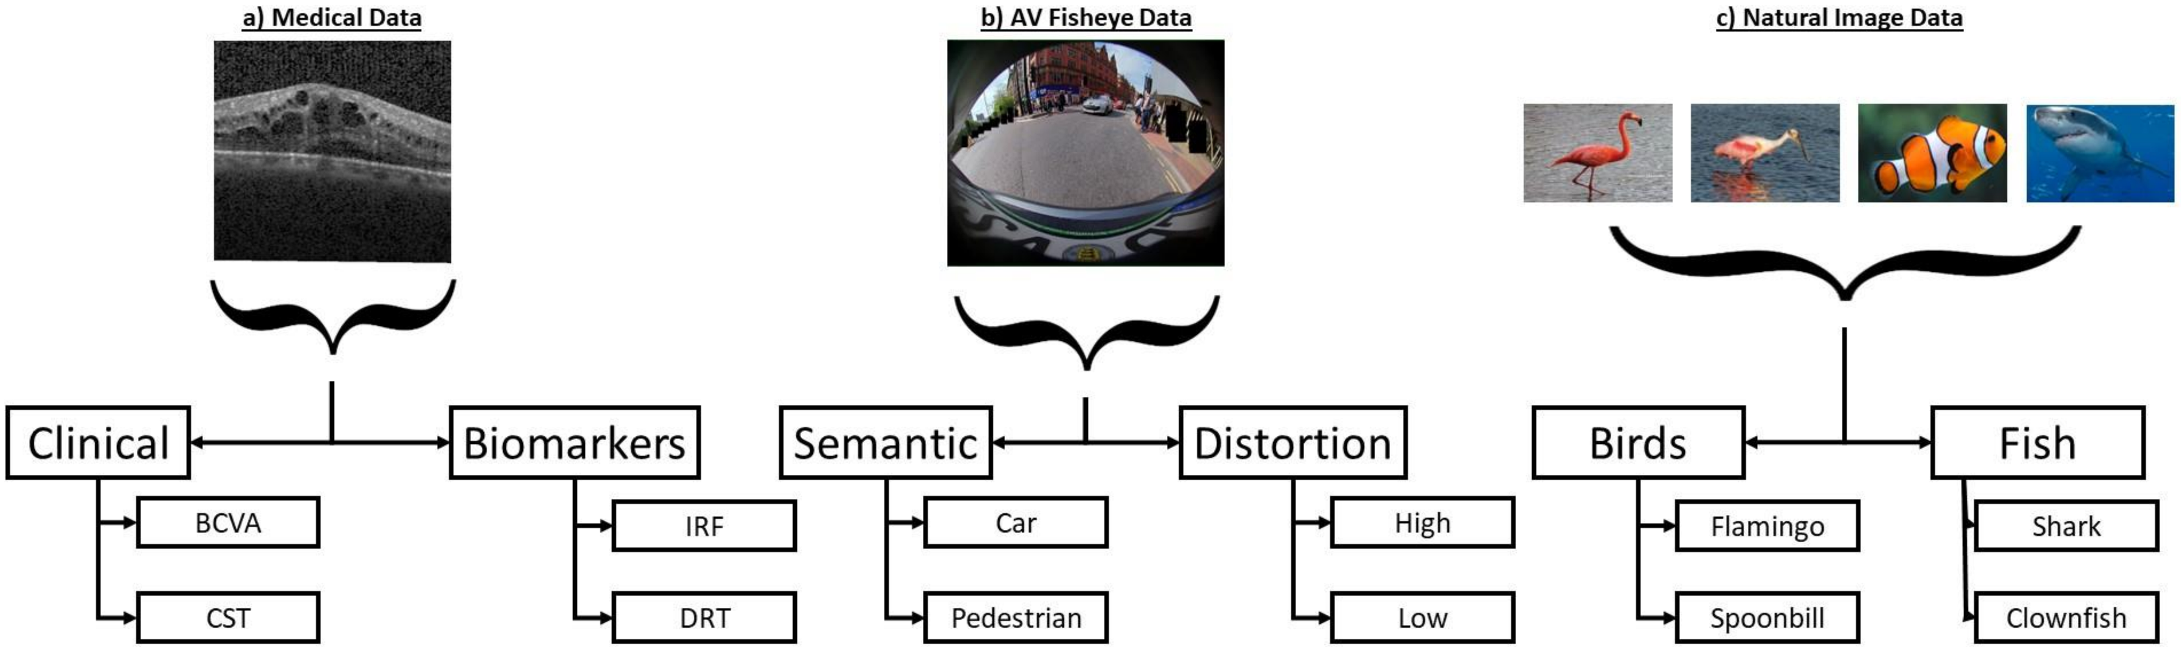

Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, *35*(8), 1798-1828.

Introduction
Novelty of Our Work

Structures of application domain can reveal ***underlying components*** of data distribution.

We can exploit this for ***application-specific*** representation learning.

Optical Coherence Tomography (OCT)

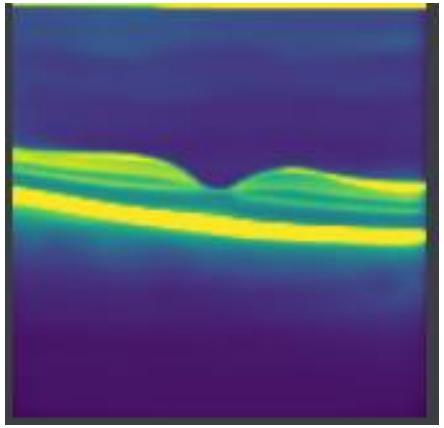

Fisheye

Seismic

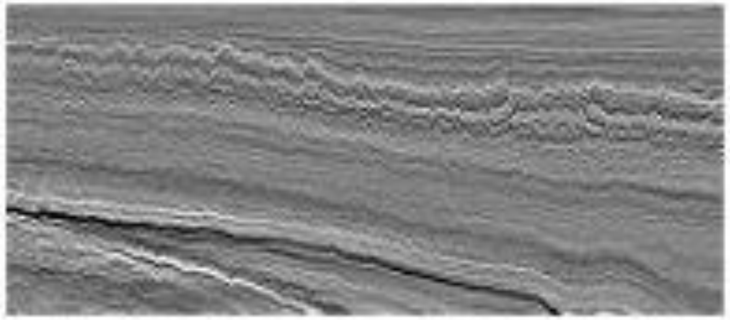

Example Application
Clinical Representations

Medical representations should reflect the interaction **clinical and biomarker factors**.

Optical Coherence Tomography (OCT)

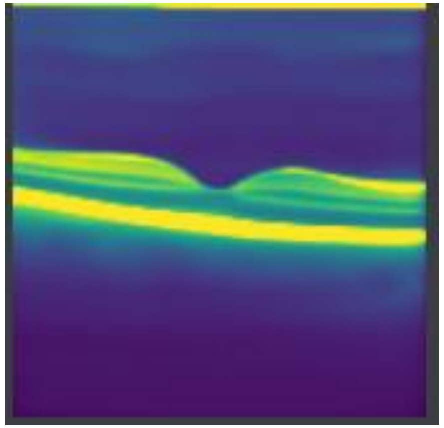

Fisheye

Seismic

Medical Example Application

Medical data is organized into biomarkers and clinical data

Clinical Data = Data collected naturally during ***routine clinical assessment***.

Biomarkers = Direct ***indicators of disease*** that have to be interpreted from OCT scans.

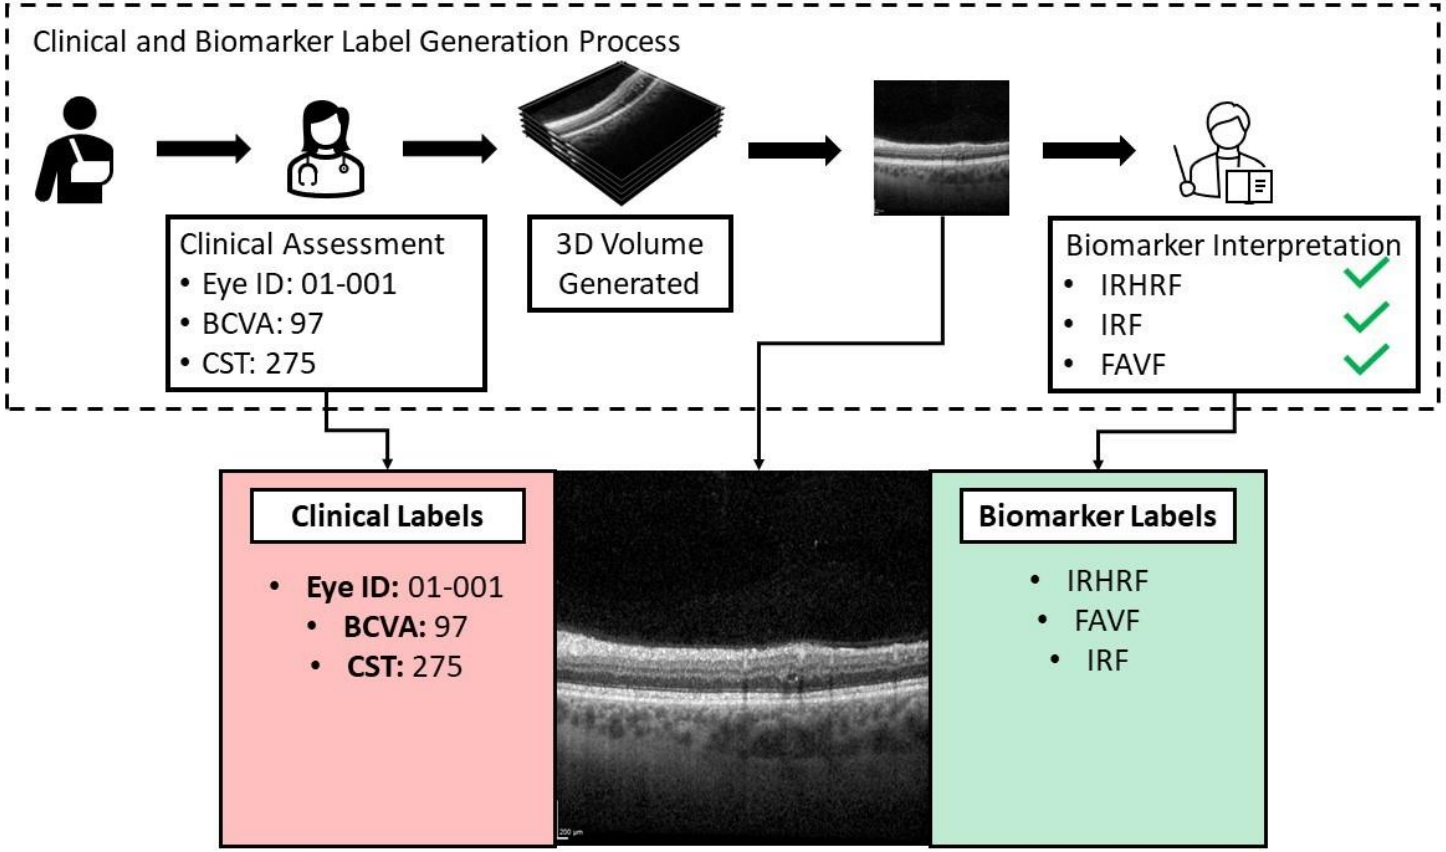

Kokilepersaud, K., Corona, S. T., Prabhushankar, M., AlRegib, G., & Wykoff, C. (2023). Clinically Labeled Contrastive Learning for OCT Biomarker Classification. *IEEE Journal of Biomedical and Health Informatics*.

Medical Example Application

There Exists Relationships between Biomarkers and Clinical Data

Can we shape ***representations on clinical data*** that will improve performance for ***biomarker detection***?.

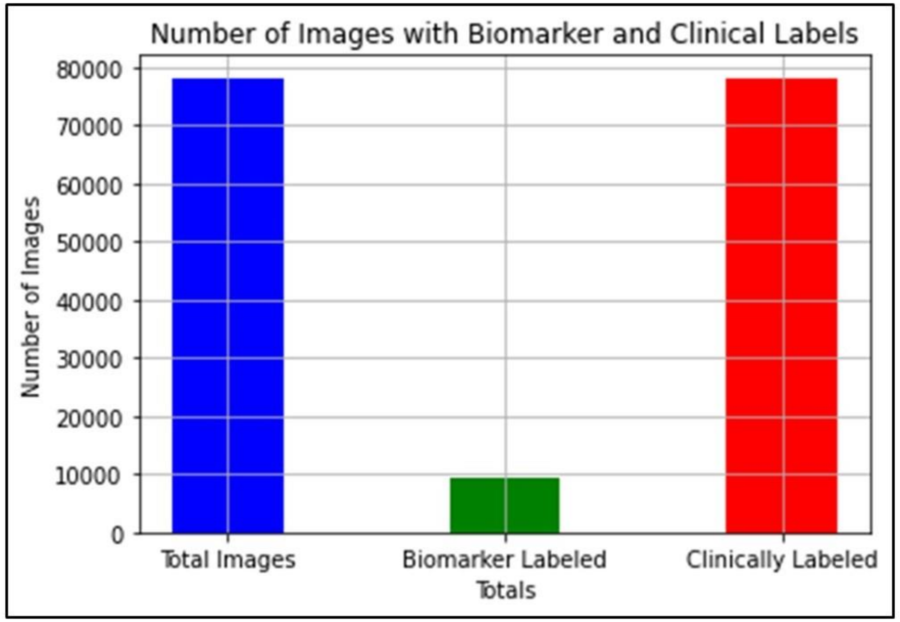

Clinical data is much **more prevalent** than biomarker data.

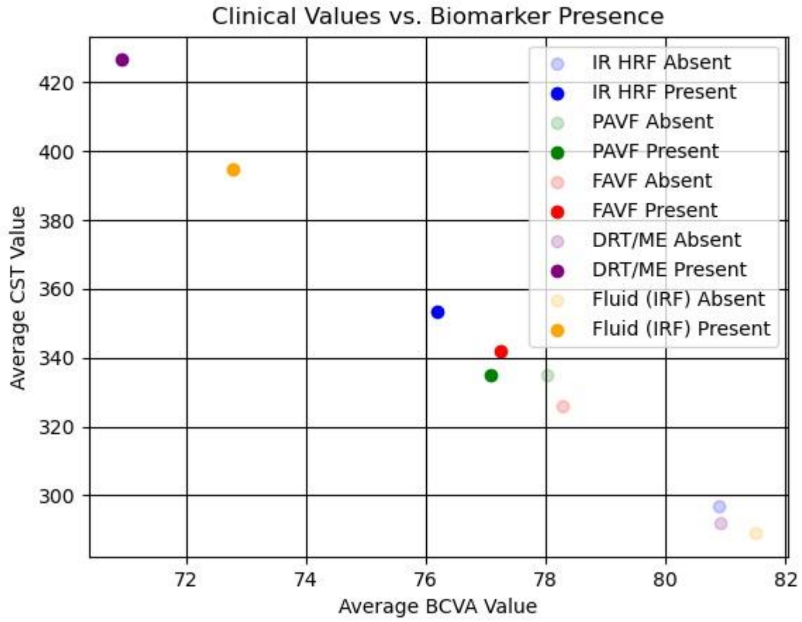

Clinical data **shares correlations** with biomarker data.

OLIVES @GeorgiaTech

Georgia Tech

Medical Example Application

Shape Clinical Representations for Downstream Biomarker Detection

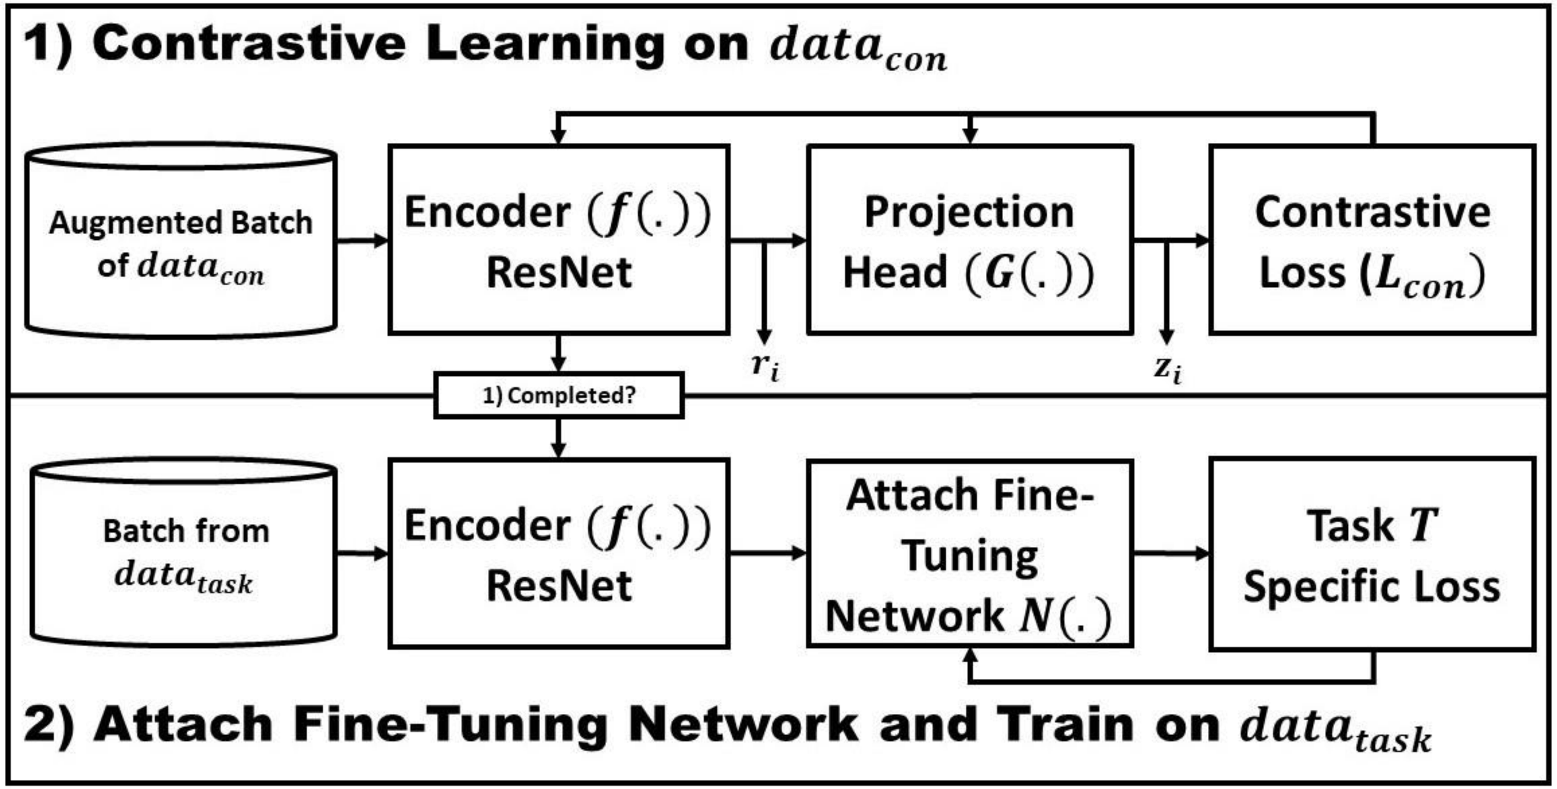

$$data_{con} = Data\ Labeled\ by\ Clinical\ Info$$

$$L_{con} = Clinical\ Contrastive\ Loss$$

$$data_{task} = Data\ Labeled\ by\ Biomarkers$$

$$N(.) = Linear\ Layer$$

Medical Example Application

Clinical Representations out-perform Standard Representations

Performance Metrics Averaged Across All Biomarkers				
Method	AUROC	Precision	Sensitivity	Specificity
PCL [10]	.676 ± .002	.676	.572	.681
SimCLR [8]	.761 ± .003	.748	.591	.779
Moco v2 [9]	.737 ± .002	.734	.597	.711
Eye ID	.802 ± .001	.769	.701	.741
CST	.793 ± .001	**.792**	.601	**.803**
BCVA	.801 ± .001	.785	.640	.792
BCVA + Eye ID	.804 ± .002	.756	**.723**	.708
BCVA + CST	.807 ± .001	.783	.643	.789
CST + Eye ID	**.819** ± .001	.756	.694	.732
BCVA + CST + Eye ID	.817 ± .001	.776	.677	.764

Clinical representations **out-perform SOTA** approaches.

Averaged Multi-Label AUROC with varying Biomarker Access				
Method	25%	50%	75%	100%
Supervised	.703 ± .002	.716 ± .003	.719 ± .002	.722 ± .005
PCL [10]	.675 ± .003	.681 ± .004	.683 ± .002	.681 ± .002
SimCLR [8]	.679 ± .004	.709 ± .006	.718 ± .003	.727 ± .002
Moco v2 [9]	.709 ± .006	.722 ± .002	.732 ± .001	.734 ± .002
Eye ID	.754 ± .005	.778 ± .003	.789 ± .001	.795 ± .001
CST	.694 ± .004	.721 ± .003	.739 ± .001	.749 ± .001
BCVA	.760 ± .009	.788 ± .001	.783 ± .001	.790 ± .001
BCVA + Eye ID	.761 ± .004	.786 ± .004	.794 ± .002	.795 ± .002
BCVA + CST	.712 ± .005	.751 ± .007	.773 ± .006	.782 ± .001
CST + Eye ID	**.766** ± .013	**.786** ± .003	**.803** ± .004	**.806** ± .003
BCVA + CST + Eye ID	.747 ± .005	.778 ± .003	.802 ± .004	.806 ± .002

Clinical representations formed from **multiple distributions are more robust** to lesser available data.

Example Application
FishEye Representations

Fisheye representations should reflect both **the semantic context and distortion**.

Optical Coherence Tomography (OCT)

Fisheye

Seismic

Fisheye Example Application

Fisheye data exhibits changes as a function of distortion

Semantic performance **worsens further from center** (higher distortion).
Can we constrain representations that **account** for this effect?

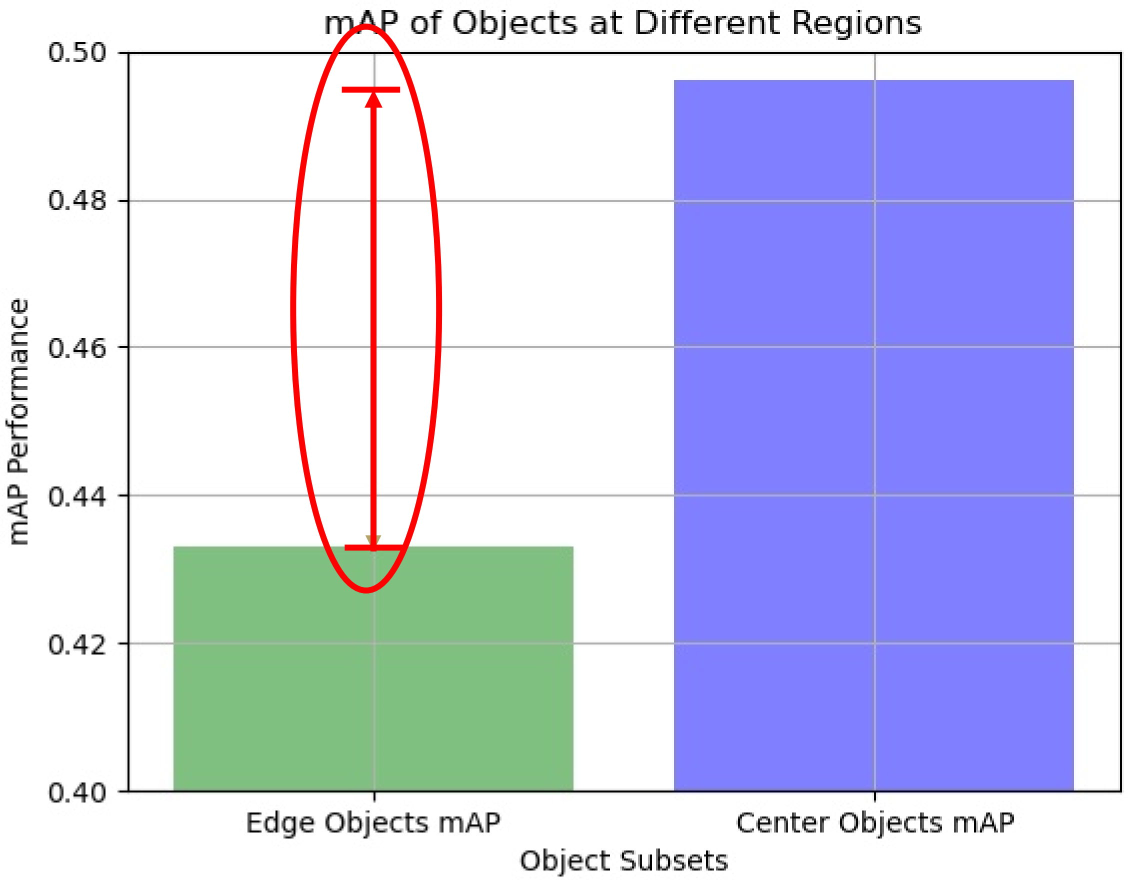

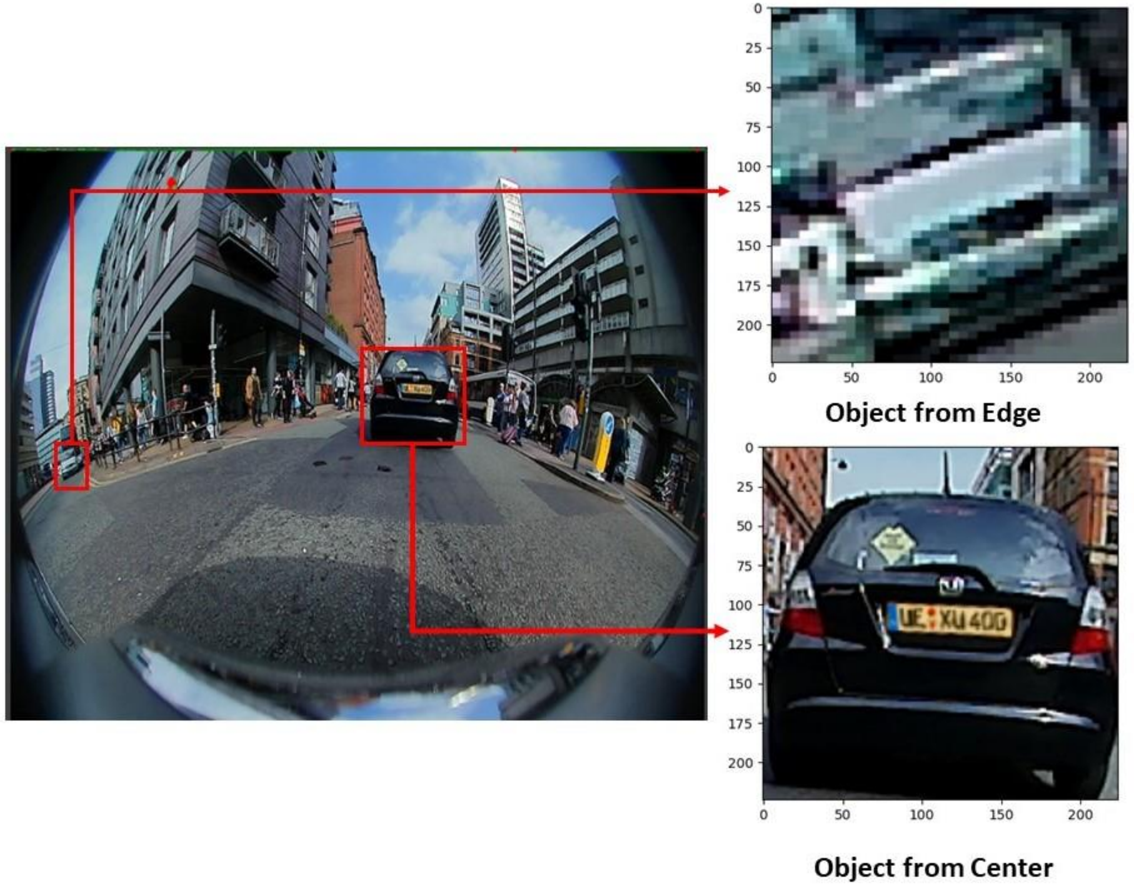

Kokilepersaud, K., Prabhushankar, M., Yarici, Y., AlRegib, G., & Parchami, A. (2023). Exploiting the Distortion-Semantic Interaction in Fisheye Data. *IEEE Open Journal of Signal Processing*.

Fisheye Example Application

Constrain Representations based on both Distortion and Semantic Information

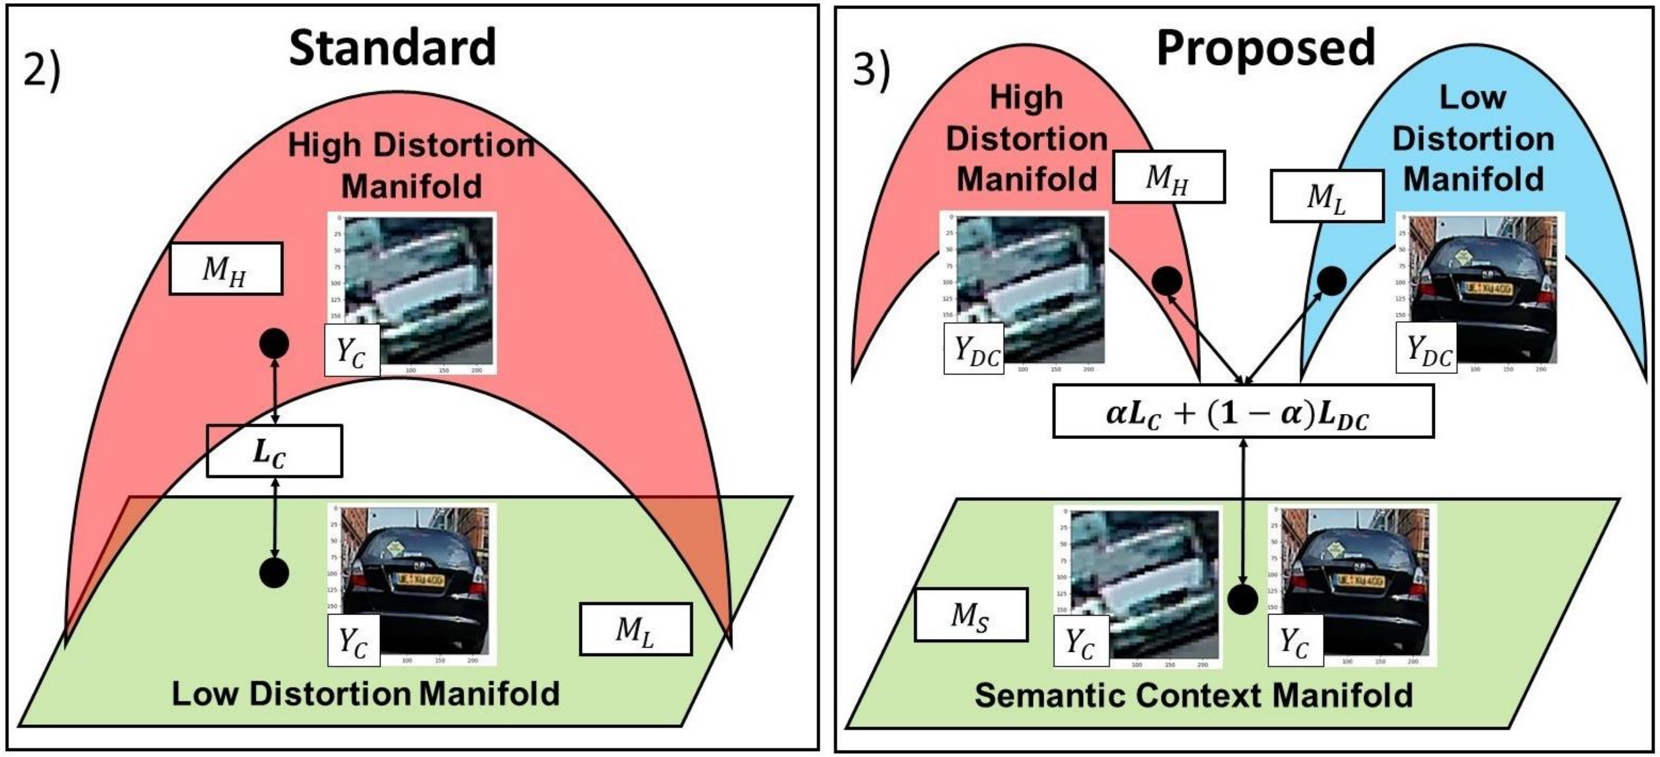

Standard

- Semantic **class loss** for training.

Proposed

- Introduce loss that balances both **semantic and distortion** information.

$$M_H = High\ Distortion\ Manifold, M_L = Low\ Distortion\ Manifold$$
$$M_S = Semantic\ Context\ Manifold$$
$$Y_C = Class\ Label\ , Y_{DC} = Distortion\ Class\ Label$$
$$L_C = Class\ Loss, L_{DC} = Distortion\ Class\ Loss$$

Fisheye Example Application

Extract Labels with Respect to Semantic Class and Distortion Class

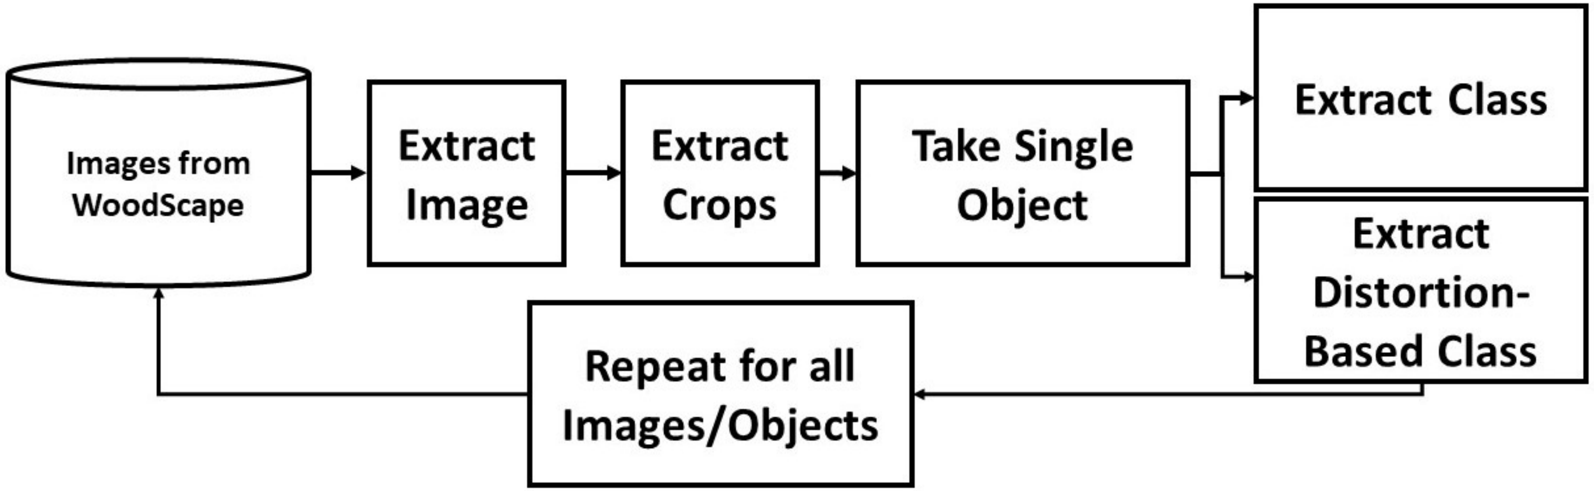

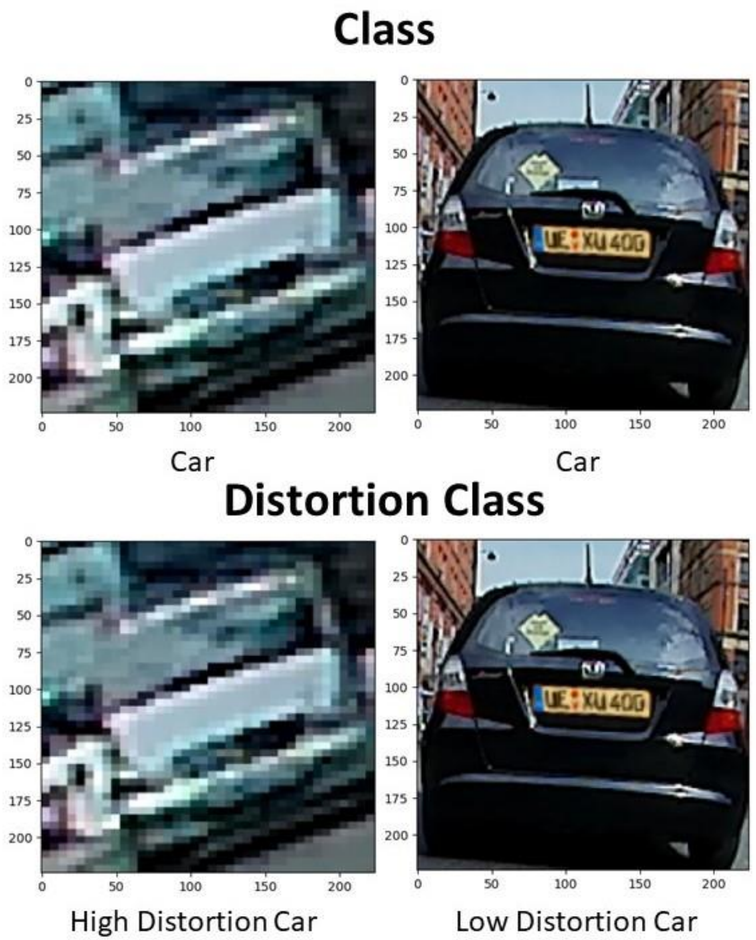

Fisheye Example Application

Use Combined Contrastive Loss Across Generated Labels

$$L_{DC} = Supervised\ Contrastive\ Loss\ on\ Distortion\ Labels$$
$$L_C = Supervised\ Contrastive\ Loss\ on\ Class\ Labels$$

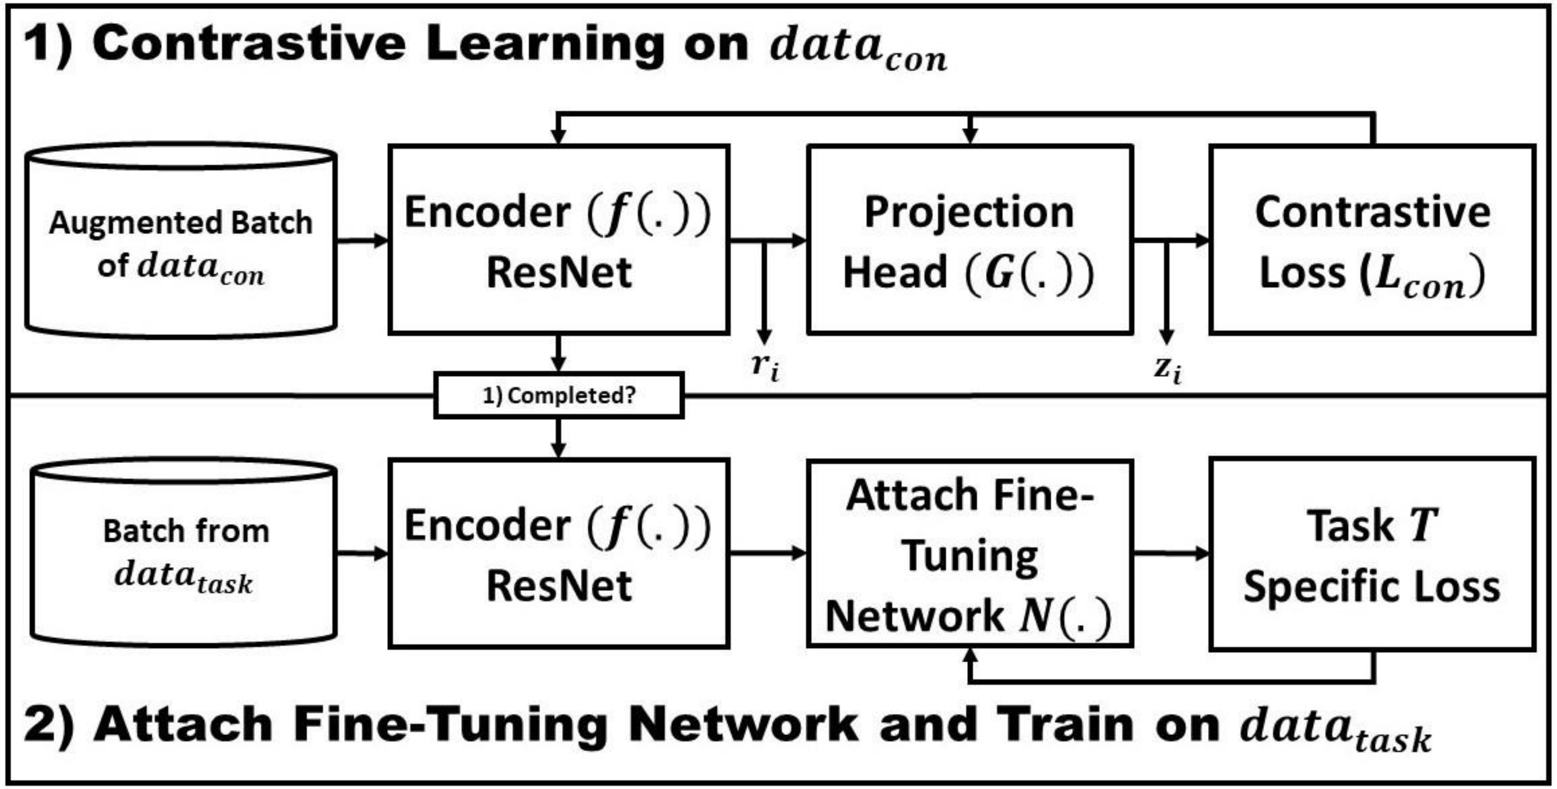

$$data_{con} = Patches\ with\ Distortion\ and\ Class\ Labels$$

$$L_{con} = \alpha L_c + (1 - \alpha) L_{DC}$$

$$data_{task} = WoodScape\ Object\ Detection$$

$$N(.) = Yolo\ v5$$

Fisheye Example Application

Performance Improves with Representations that Reflect Both Fisheye Components

- Alpha **controls balance between semantic and distortion** information in loss: $\alpha L_C + (1-\alpha)L_{DC}$
- **Equal weight** on both losses performed best
- Both semantic and distortion information important for fisheye representations

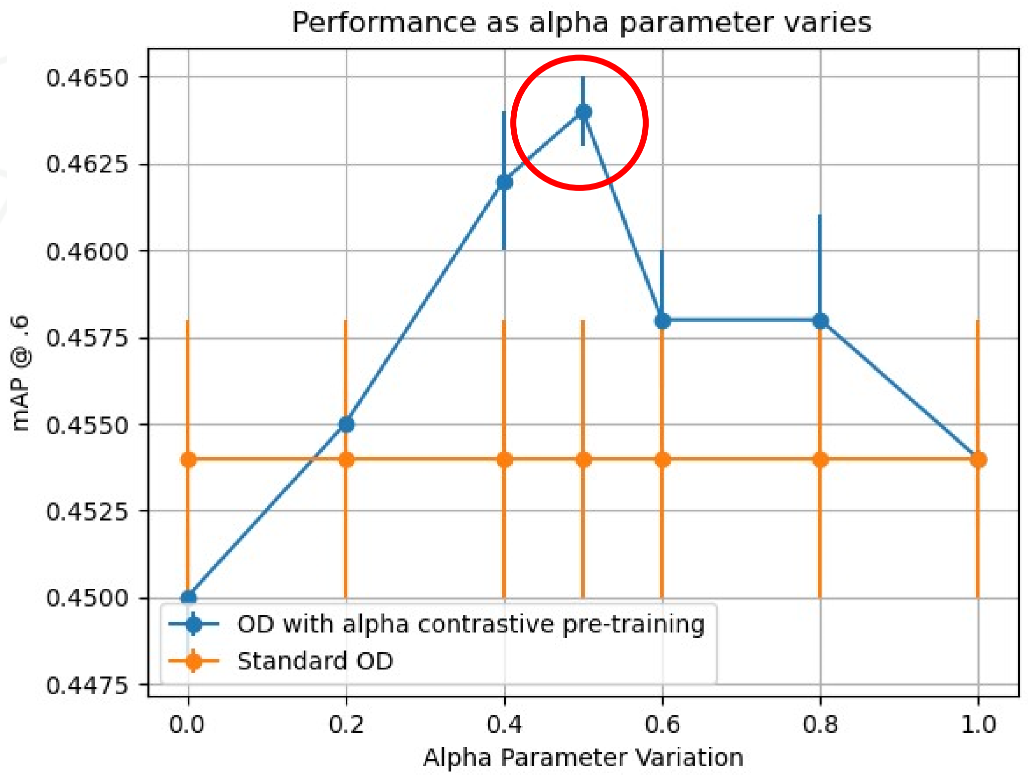

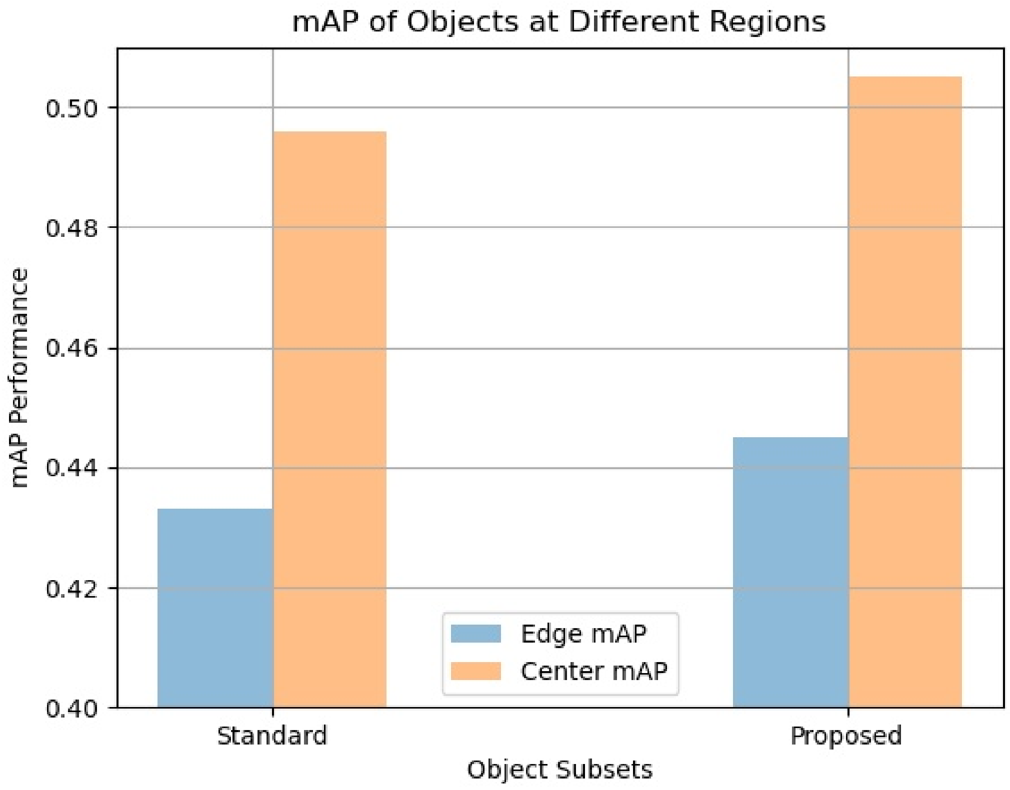

Example Application
Seismic Representations

What are good application specific representations for Seismic?

Optical Coherence Tomography (OCT)

Fisheye

Seismic

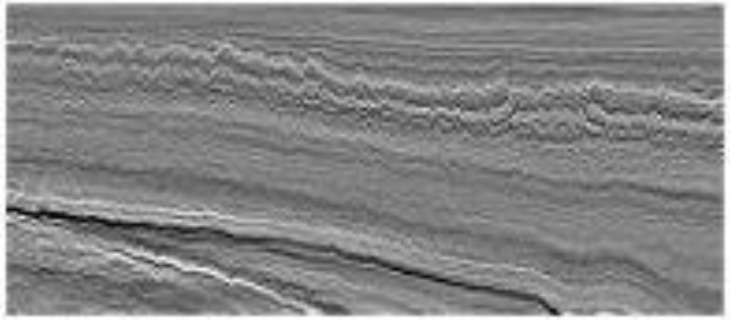

Connection to Seismic

Seismic can benefit From Representations based on Seismic-Specific Considerations

What are considerations for a good seismic representation space?

Multi-Modality

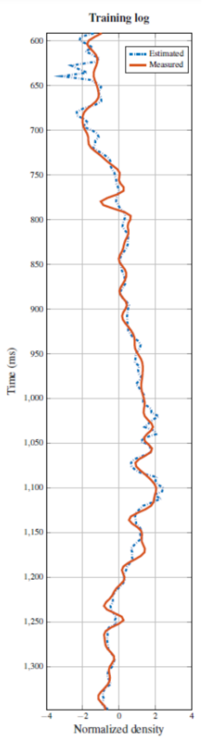

Physics Interactions

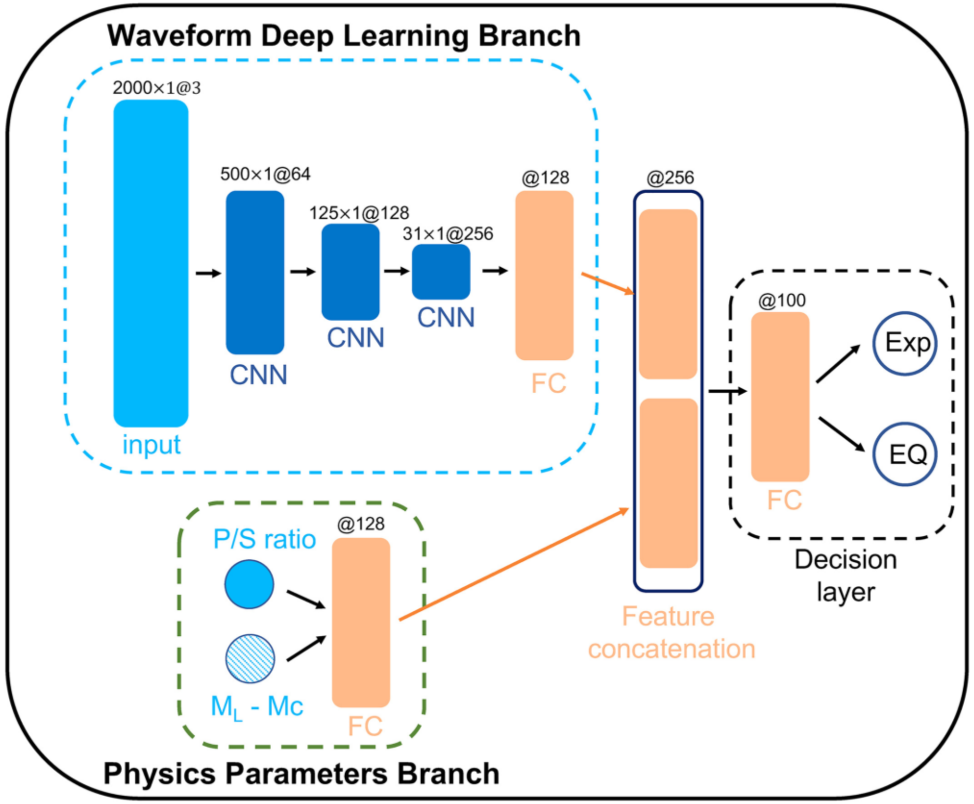

Human Interaction

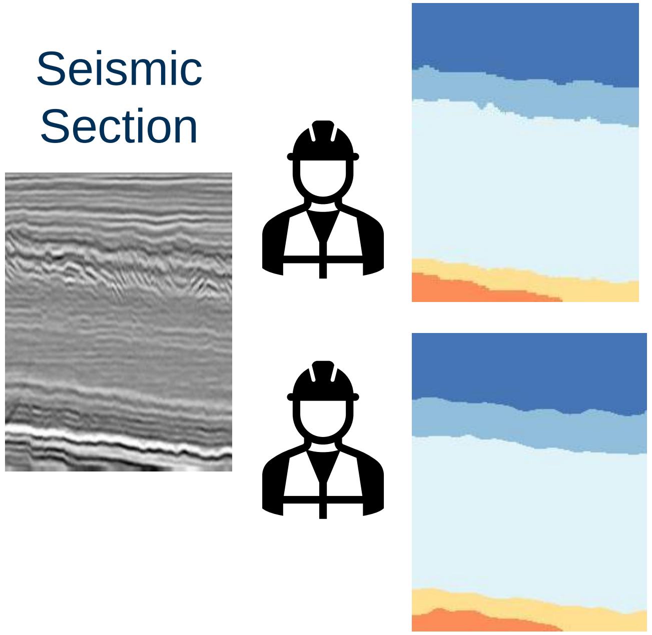

Kong, Q., Wang, R., Walter, W. R., Pyle, M., Koper, K., & Schmandt, B. (2022). Combining Deep Learning With Physics Based Features in Explosion-Earthquake Discrimination. *Geophysical Research Letters*, *49*(13), e2022GL098645.

Connection Seismic
Factors of Variation Exist within Seismic Volumes

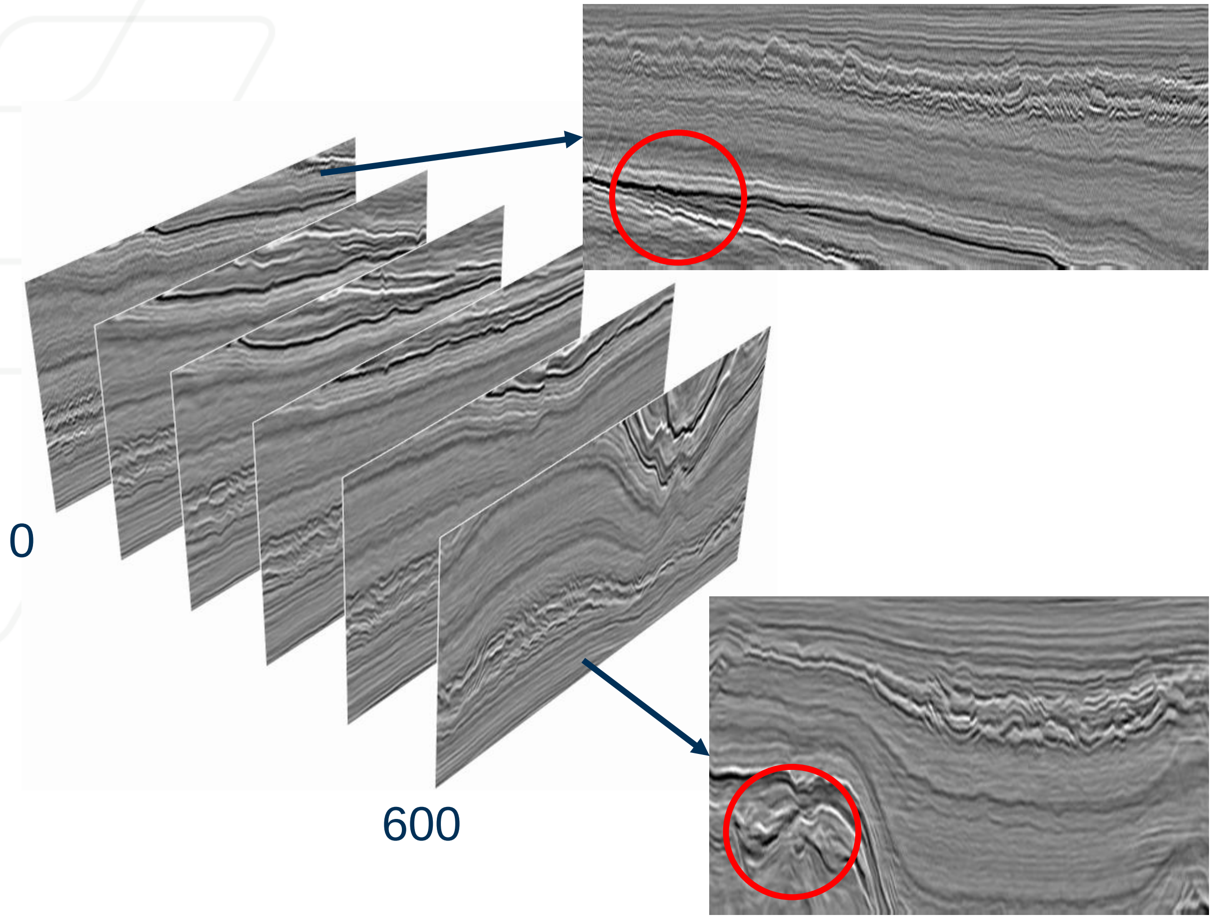

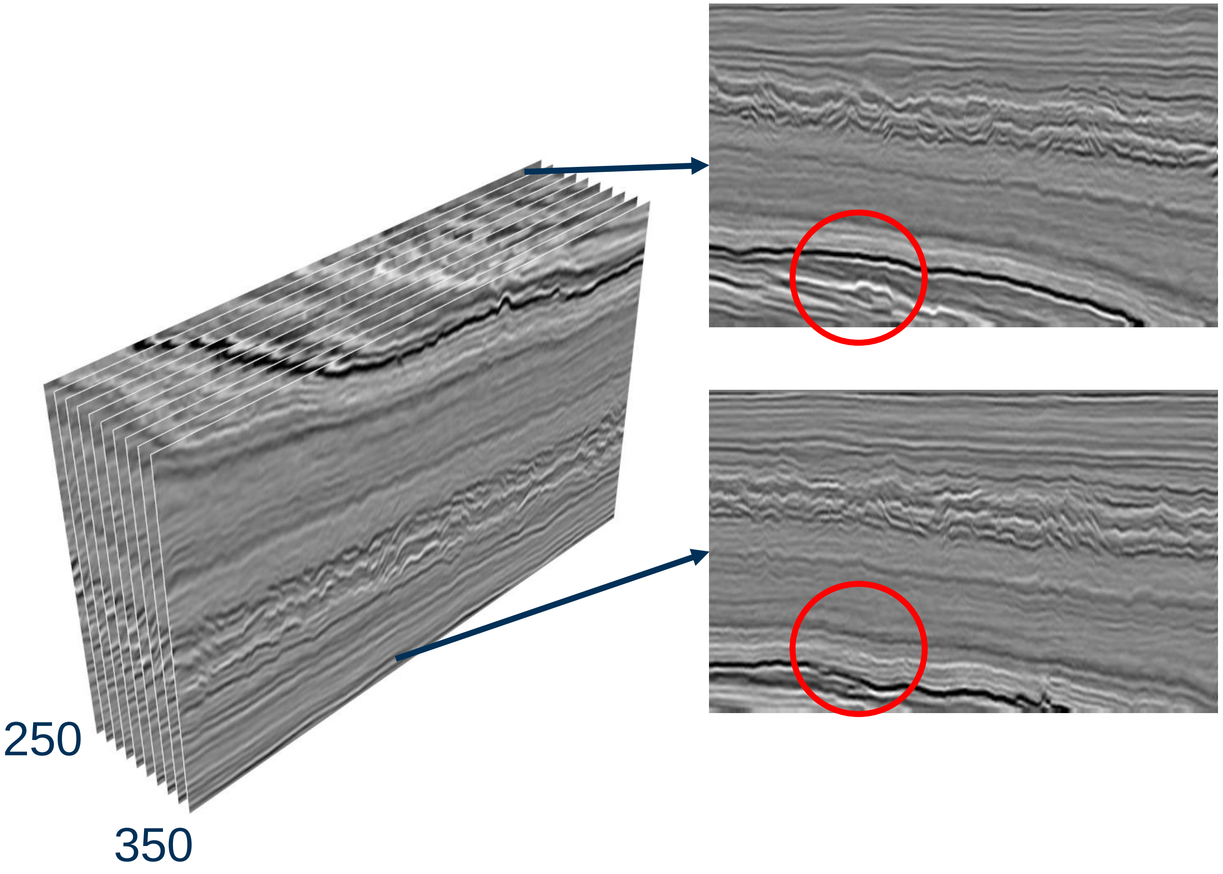

Different structures entirely

Similar structures, but still fine-grained differences

Connection to Seismic
Seismic Structures Exist within a Single Modality

Ideally want model to:
1) Associate **close slices** together
2) Learn **fine-grained structural** differences

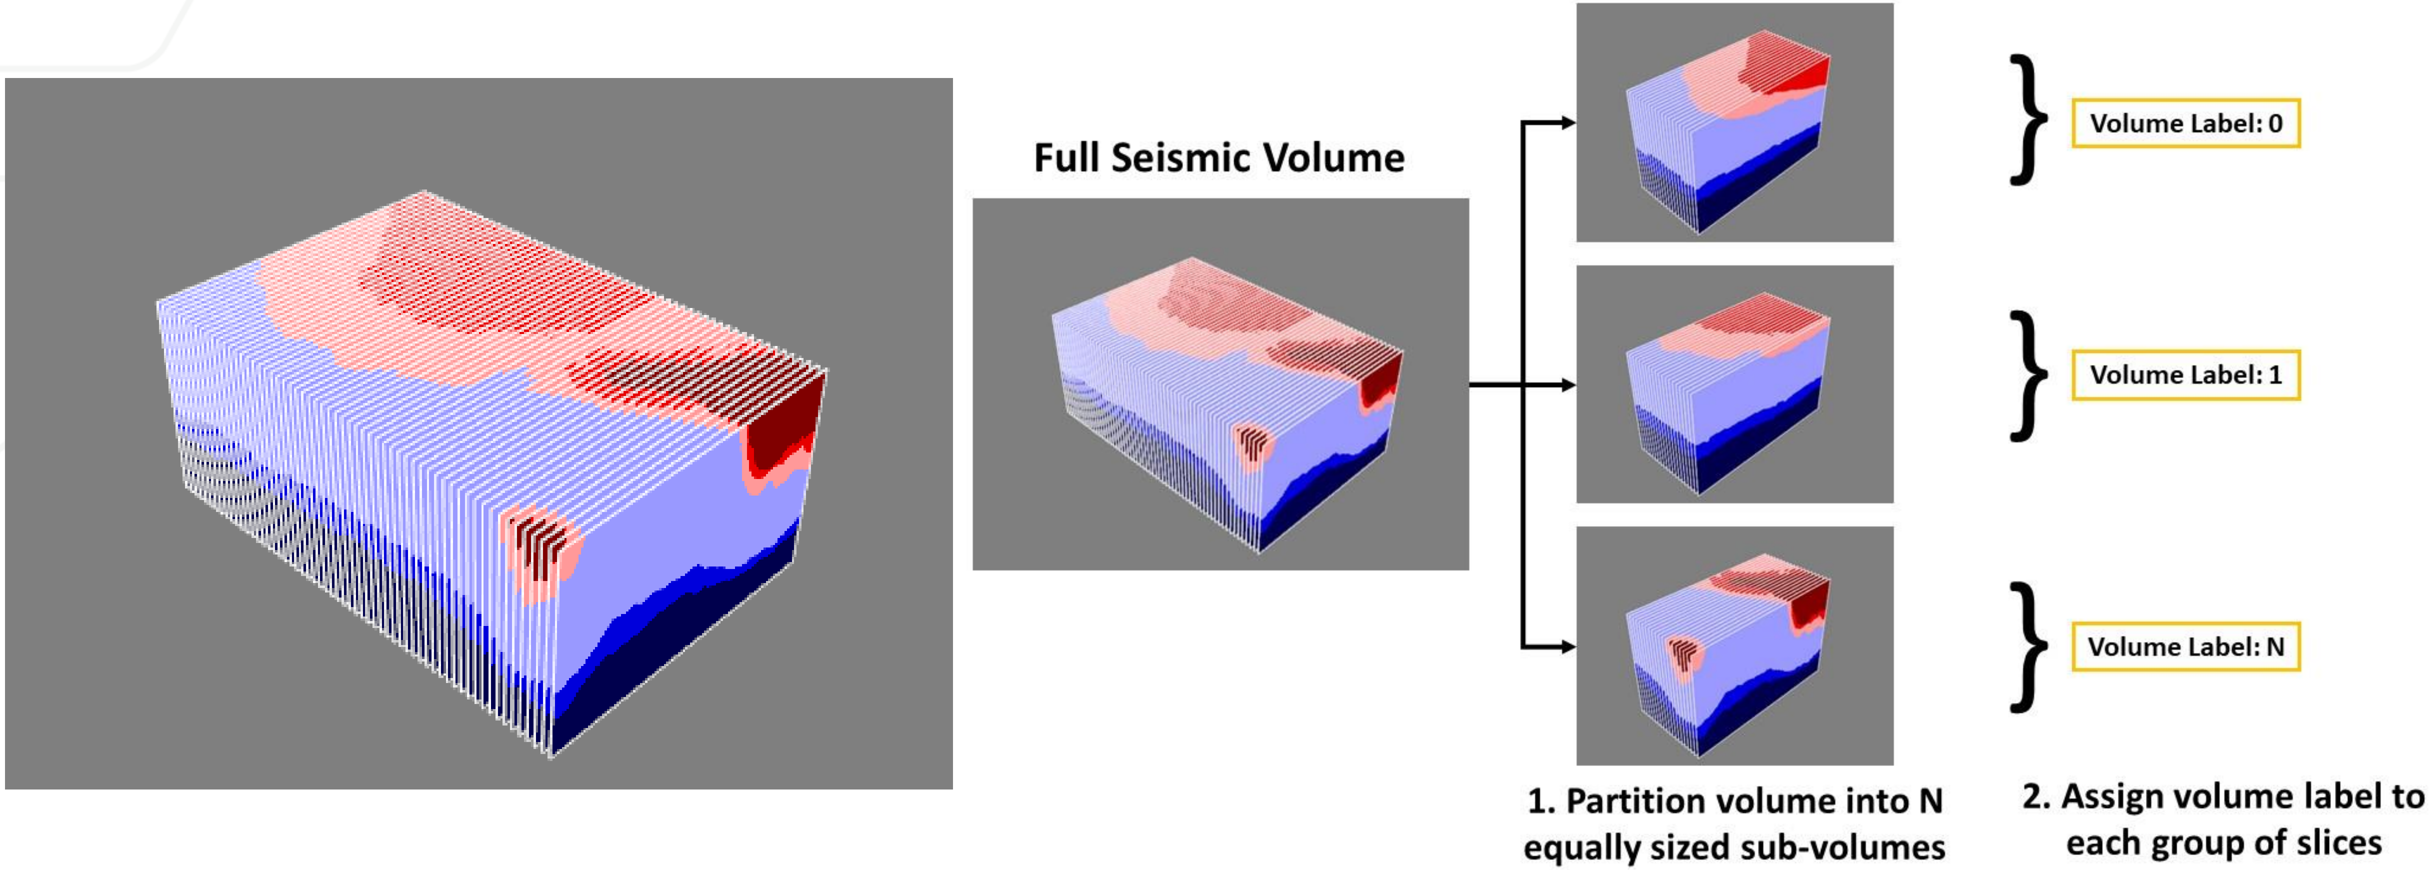

Kokilepersaud, K., Prabhushankar, M., & AlRegib, G. (2022, August). Volumetric supervised contrastive learning for seismic semantic segmentation. In *Second International Meeting for Applied Geoscience & Energy* (pp. 1699-1703). Society of Exploration Geophysicists and American Association of Petroleum Geologists.

Seismic Example Application

Use Volumetric Contrastive Loss

L_{VH} = *Volume Hard encourages* ***attention to fine − grained differences*** *between similar structures*
L_V = *Volumetric Contrastive Loss encourages* ***close slices in the volume*** *have similar embeddings*

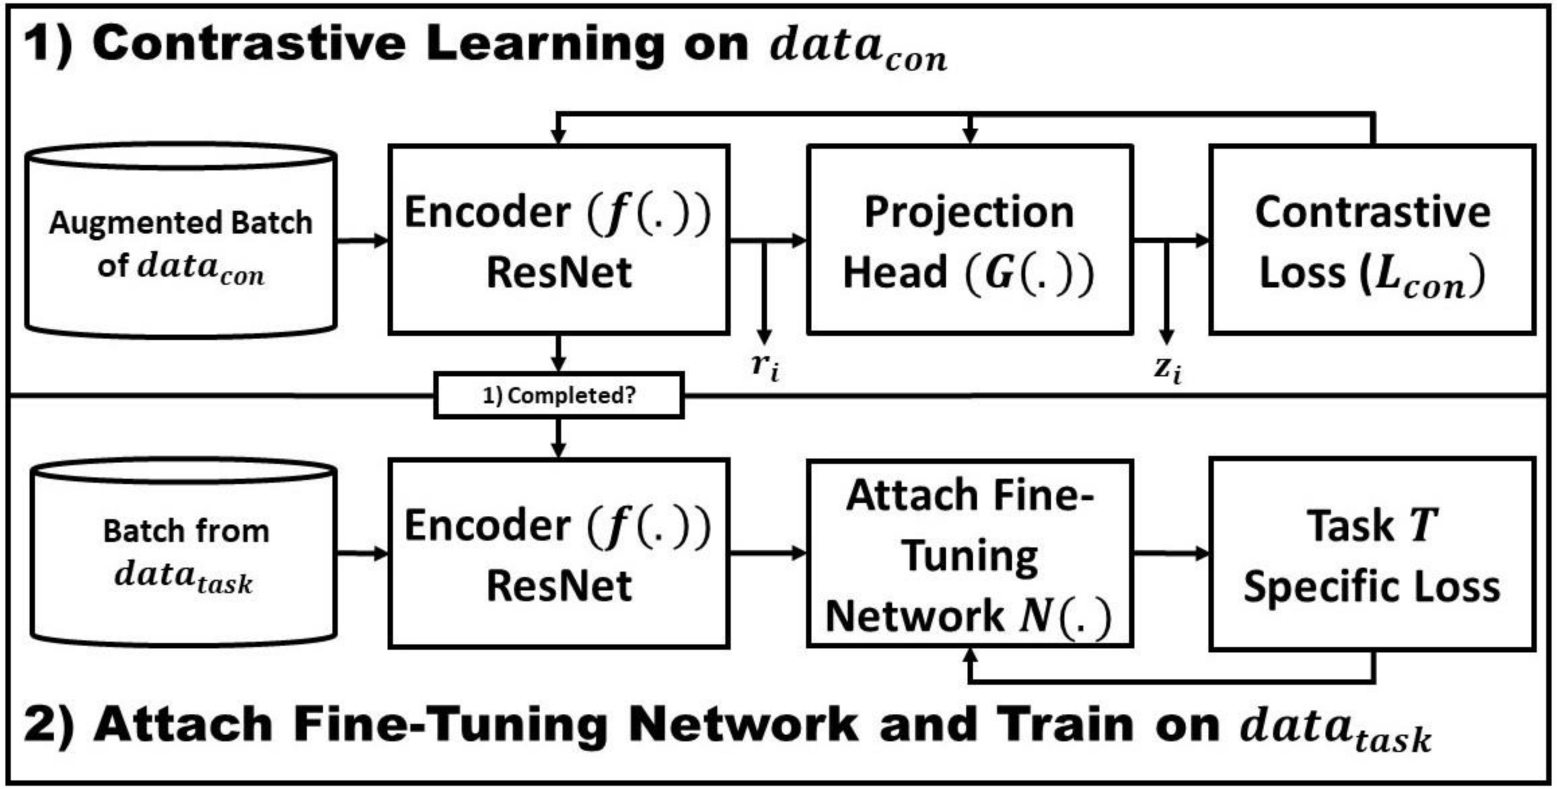

$data_{con}$ = *F3 Block with Volume Labels*

$$L_{con} = L_{VH} + L_V$$

$data_{task}$ = *F3 Block with Labels*

$N(.)$ = *DeepLab v3*

Seismic Example Application
Volumetric Labels Improve Performance

Observations

Using **volumetric contrastive learning** led to improvement in performance
Further **opportunities exist for understanding seismic-specific** representations

Method	MIOU
SimCLR	.6913
Volumetric Loss	**.6980**

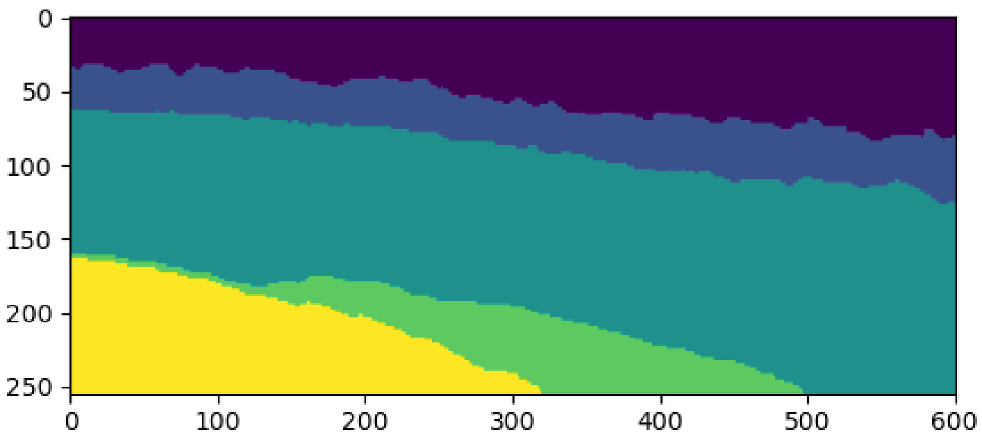

Ground Truth

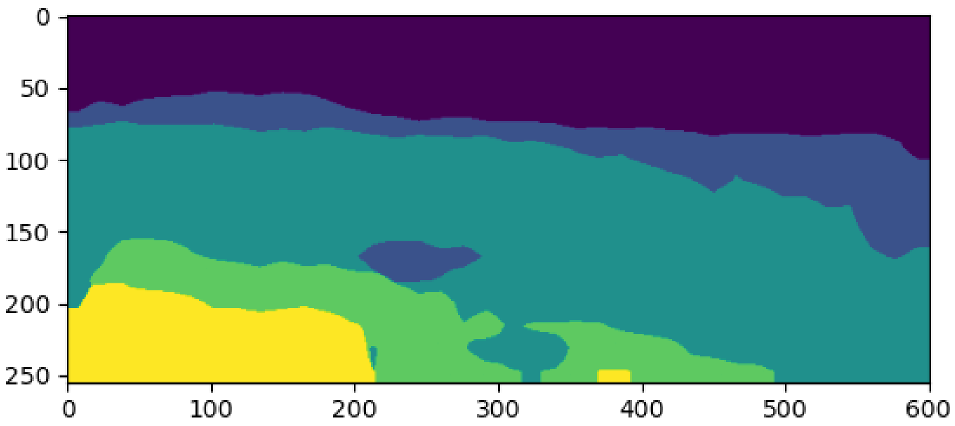

Volumetric

Conclusion

Application Specific

- Medical representations can be better shaped by **components relating to clinical information**
- Fisheye-specific representations can be better shaped by the interaction **of both semantic and distortion** based information
- Seismic representations have potential opportunities as well as volumetric positional information

Overall

- Every domain of data potentially has ***easily accessible distributions*** that can shape representations to better reflect the ***underlying distribution*** of data in the domain.

Publications and Code

1. Kokilepersaud, K., Prabhushankar, M., & AlRegib, G. (2022, August). Volumetric supervised contrastive learning for seismic semantic segmentation. In *Second International Meeting for Applied Geoscience & Energy* (pp. 1699-1703). Society of Exploration Geophysicists and American Association of Petroleum Geologists.

2. Kokilepersaud, K., Prabhushankar, M., AlRegib, G., Corona, S. T., & Wykoff, C. (2022, October). Gradient-based severity labeling for biomarker classification in oct. In *2022 IEEE International Conference on Image Processing (ICIP)* (pp. 3416-3420). IEEE.

3. Kokilepersaud, K., Corona, S. T., Prabhushankar, M., AlRegib, G., & Wykoff, C. (2023). Clinically Labeled Contrastive Learning for OCT Biomarker Classification. *IEEE Journal of Biomedical and Health Informatics*.

4. Kokilepersaud, K., Prabhushankar, M., Yarici, Y., AlRegib, G., & Parchami, A. (2023). Exploiting the Distortion-Semantic Interaction in Fisheye Data. *IEEE Open Journal of Signal Processing*.

5. K. Kokilepersaud, M. Prabhushankar, and G. AlRegib, "Leveraging Image Perturbations for Medical Contrastive Learning," *Journal of Biomedical and Health Informatics Special Issue on Machine Learning Technologies for Biomedical Signals Processing*, submitted on Sept. 30, 2023.

For more OLIVES content, please visit:

GitHub

SCAN ME

Publications

SCAN ME