AdaDCP: Learning an Adapter with Discrete Cosine Prior for Clear-to-Adverse Domain Generalization

Qi Bi • Yixian Shen • Jingjun Yi • Gui-Song Xia

1. University of Amsterdam, the Netherlands
2. Wuhan University, China

Background & Motivation

• Adverse weather (fog, rain, night, snow) severely blurs scenes and scarce annotations hinder training.
• Large-scale clear scenes are easy to collect but domain gaps remain for unseen adverse conditions.
• Vision foundation models (VFMs) offer inherent generalization but require adaptation to close the clear-to-adverse gap.

[2]

Contributions & Overview

- Introduce AdaDCP: adapter with discrete cosine prior enabling clear-to-adverse adaptation.
- Decouple frequency components using DCT and handle invariant (cA) and variant (cH/cV/cD) information separately.
- Design three modules:
  WIP (weather-invariant learning),
  WVP (weather-variant learning),
  WPA (cross-frequency alignment).
- Attain state-of-the-art performance across fog, night, rain and snow without adverse training data.

[7]

Observation

- DCT decomposes features into cA and cH/cV/cD components.
- cA captures weather-invariant structure; cH/cV/cD encode domain-specific details.
- DCT has superior energy compaction compared to FFT, DST, and wavelet transforms.
- Fine-grained frequency bands exhibit similar invariance/variance behaviours across domains.

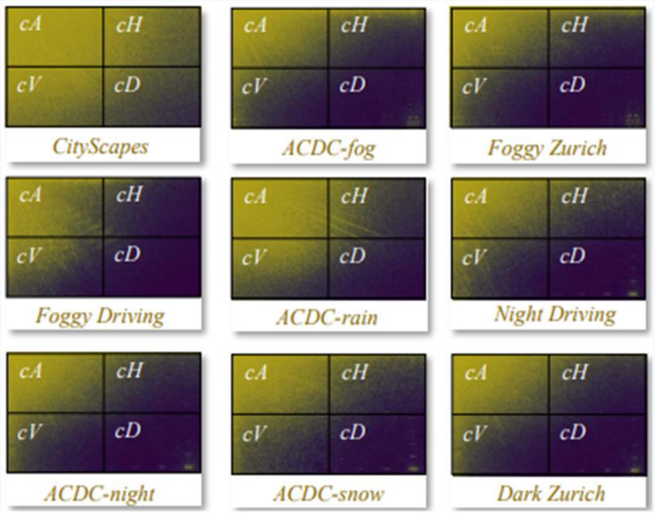

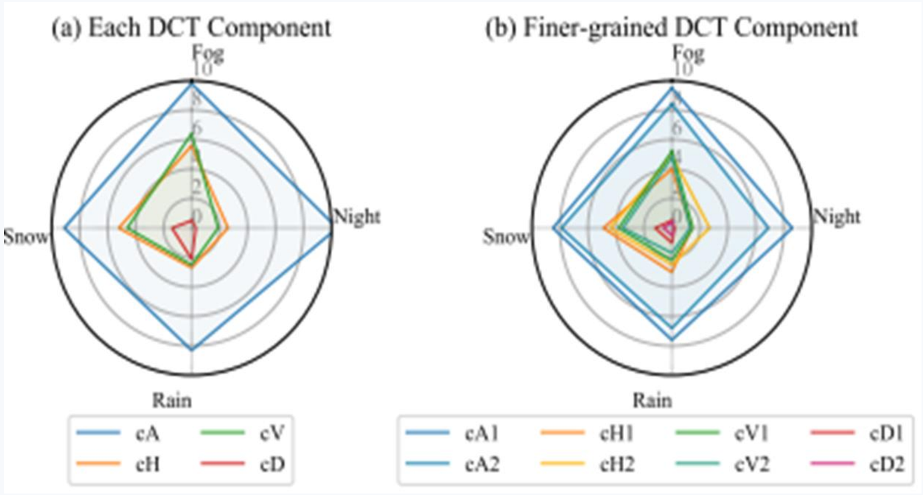

[3] [4]

Framework Overview & Method Design

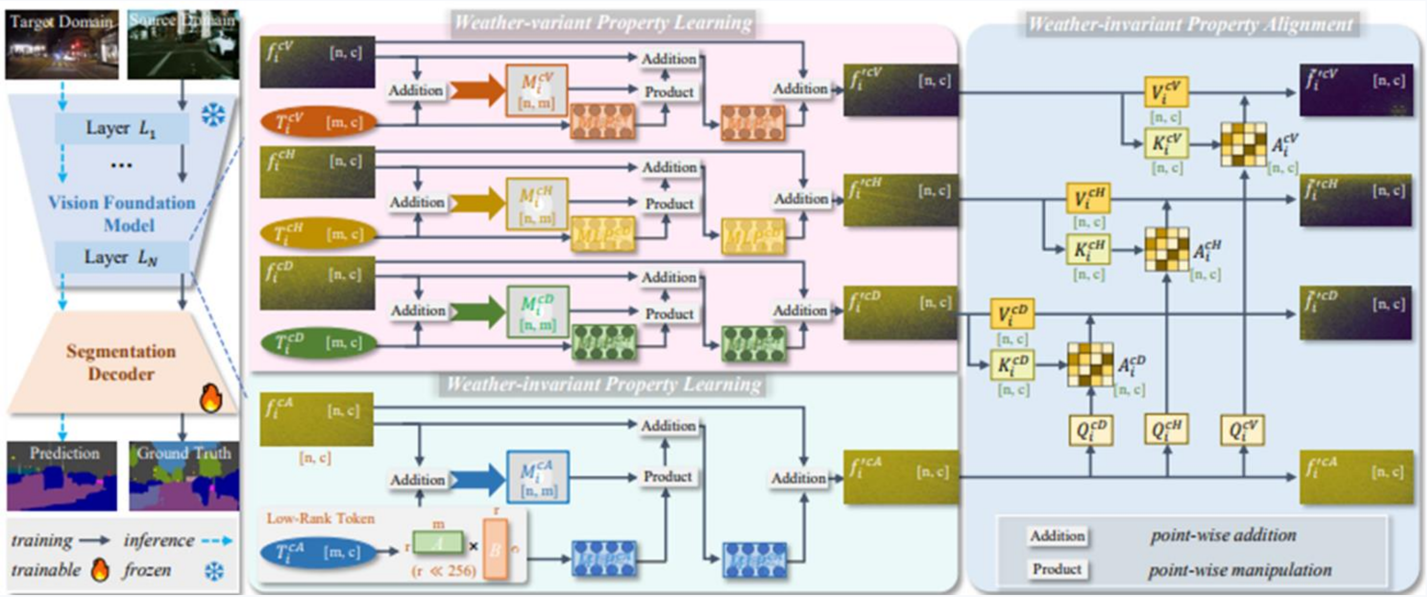

- Insert AdaDCP modules after each VFM layer to modulate features.
- Separate cA and cH/cV/cD channels and learn them via WIP and WVP.
- Align invariant and variant channels through cross-attention in WPA.

[8] [9]

Framework Overview & Method Design

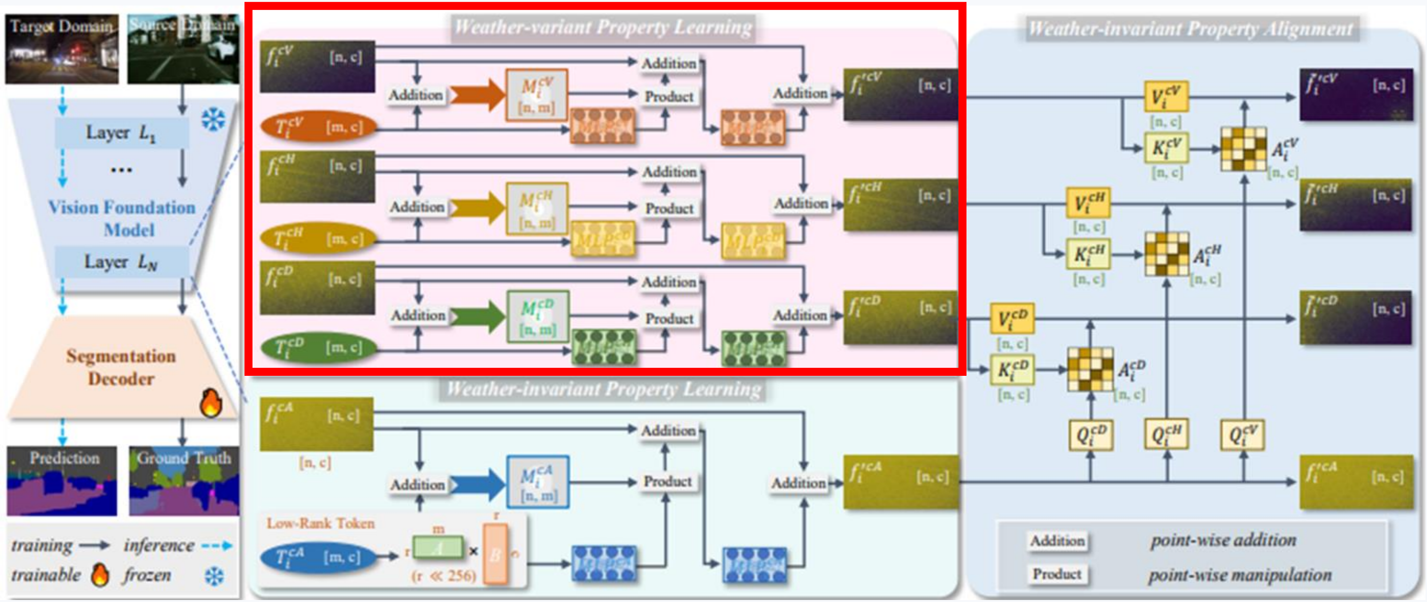

- Extract components cH/cV/cD via high-pass filtering.
- Learn separate token sets to capture weather-specific patterns.
- Fuse adapted features back into the network.

[8] [9]

Framework Overview & Method Design

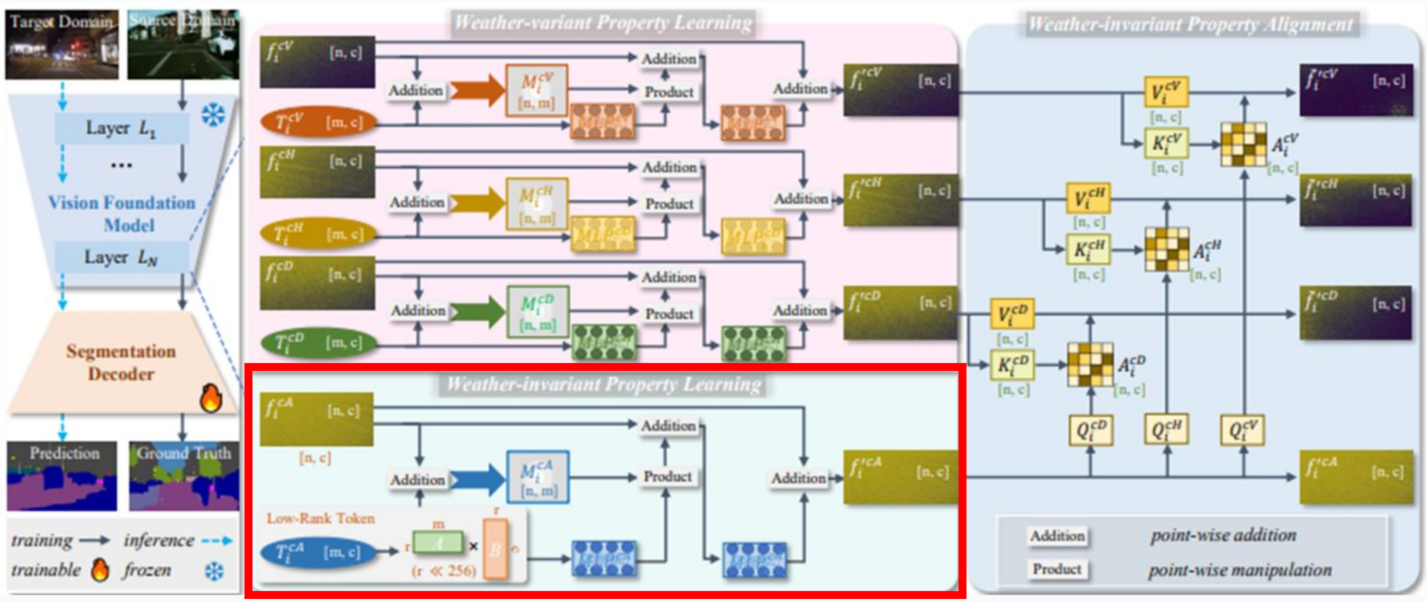

- Isolate component cA via DCT and low-pass filtering.
- Project cA into a compact token space and learn low-rank adaptation.
- Enhances semantic consistency across clear and adverse domains.

[8] [9]

Framework Overview & Method Design

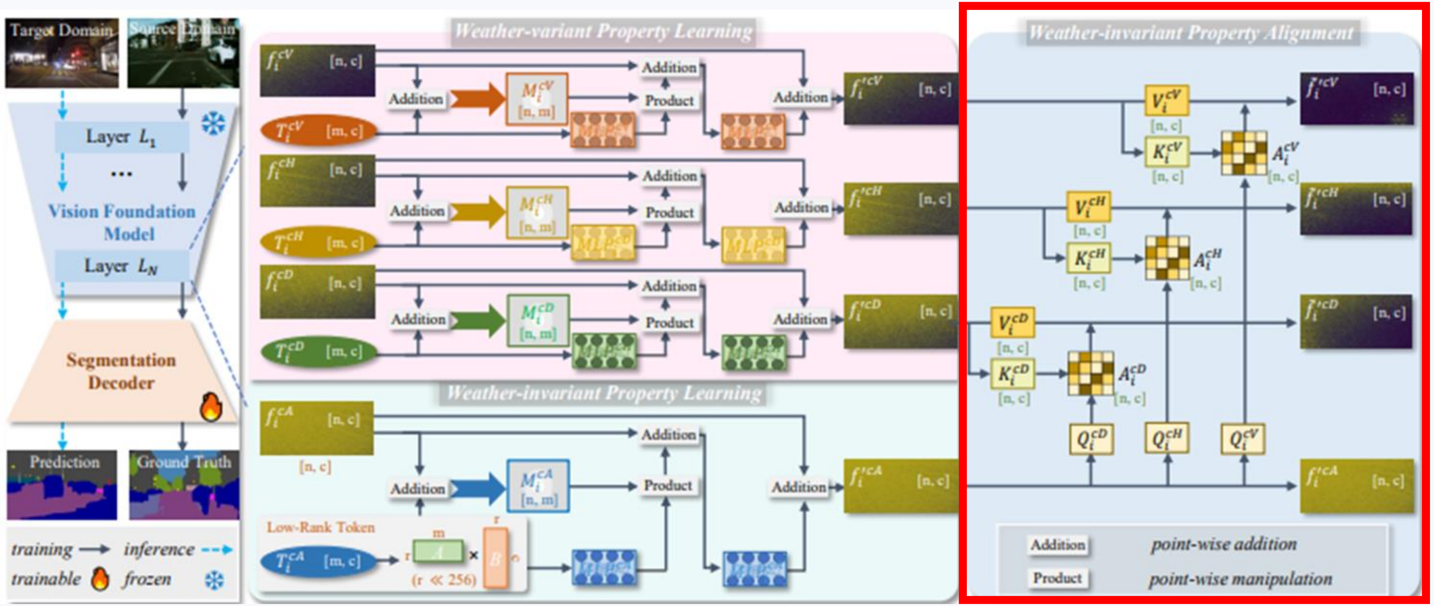

- Employ cross-attention between invariant and variant channels to highlight aligned features.
- Weighted aggregation suppresses adverse noise while retaining salient details.
- Projects aligned features back to the main branch for subsequent layers

[8] [9]

Numerical Results

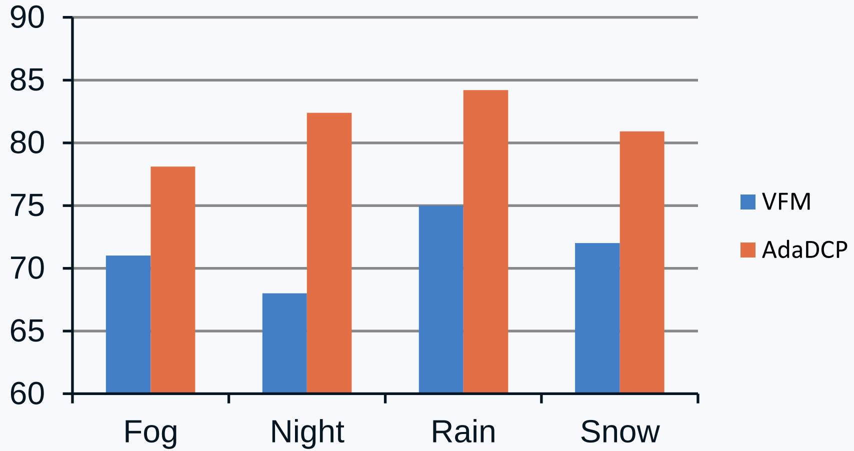

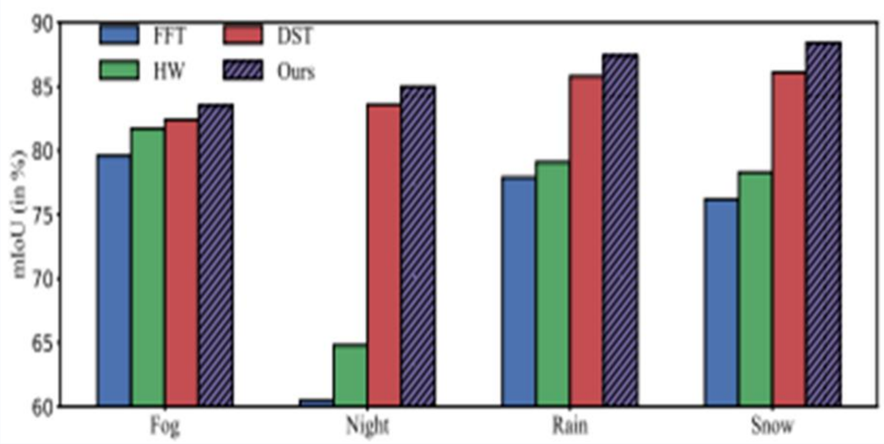

- AdaDCP yields mIoU gains of +7.1%, +14.4%, +9.2% and +8.9% on fog, night, rain and snow, respectively, over the bseline.
- Outperforms FFT, Haar Wavelet (HW) and Discrete Sine Transform (DST) across all conditions.
- Ablation studies confirm the effectiveness of each component.

[13]

Qualitative Results

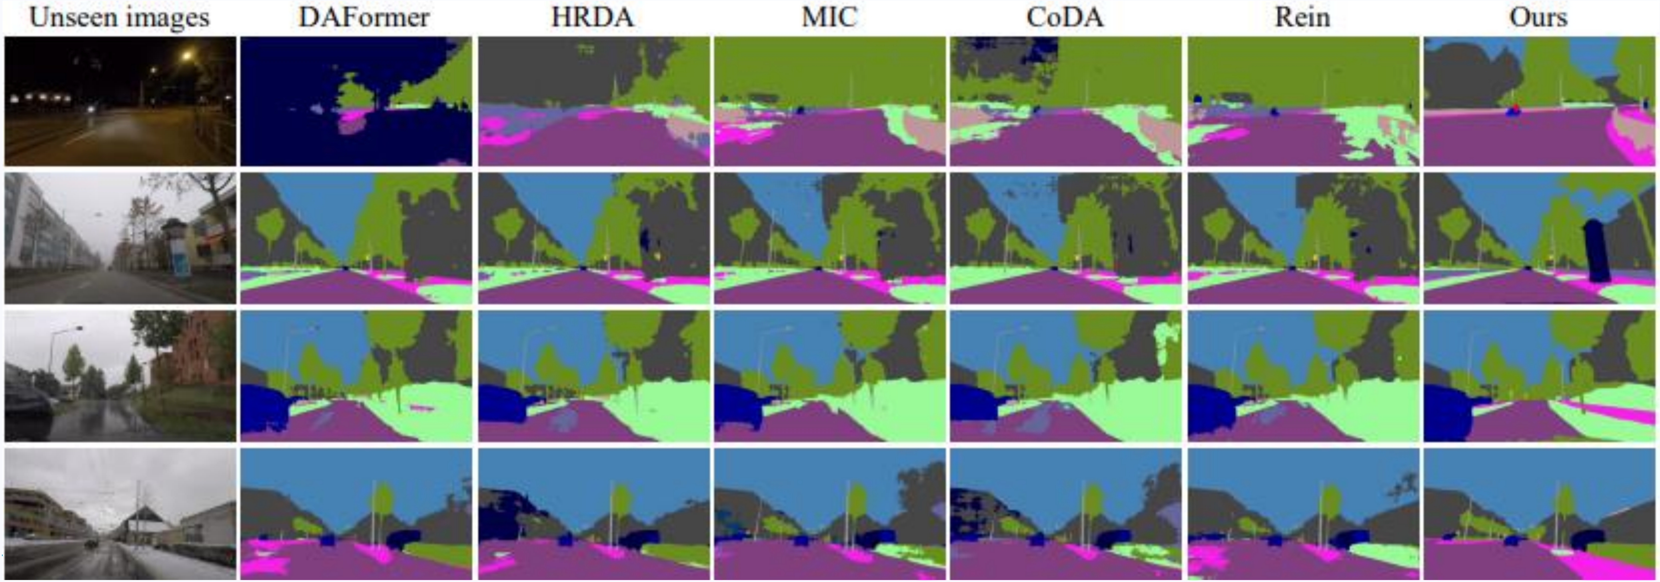

[14]

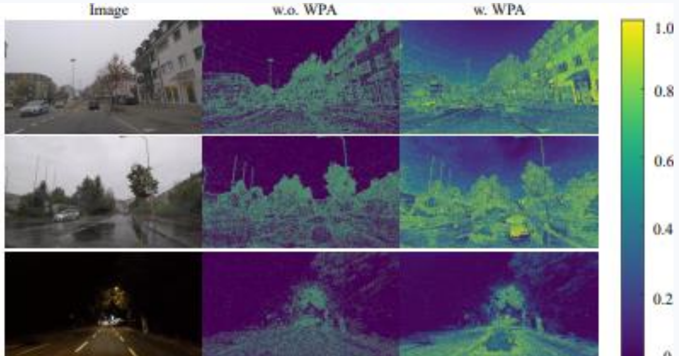

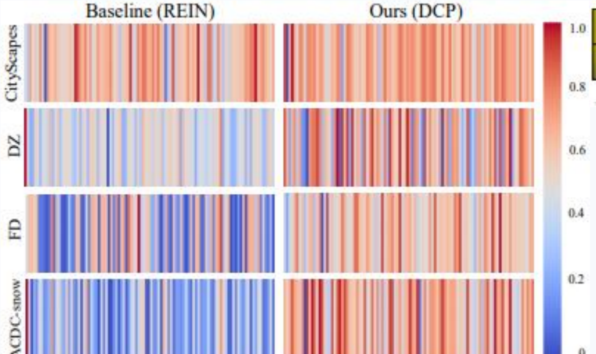

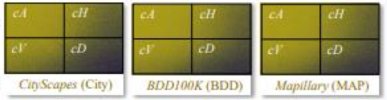

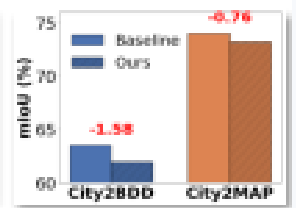

Conclusion

- AdaDCP produces sharper segmentation boundaries and preserves small objects under adverse conditions.
- Improves pixel-level predictions for both depth and dehaze tasks (not shown here) demonstrating cross-task generality.
- Future work: extend frequency-prior adaptation to other modalities and tasks.

[14] [15]

Thanks for your attention!

[1]