

DOMINO: Domain-Invariant Hyperdimensional Classification for Multi-Sensor Time Series Data

Junyao Wang, Luke Chen, Mohammad Abdullah Al Faruque
junyaow4@uci.edu, University of California, Irvine, United States

Motivation

Multi-Sensor Time Series Data:

With the emergence of IoT, heterogeneously connected sensors capture information over time, constituting multi-sensor time series data


Problem I: Sophisticated DNNs, e.g., RNNs have been proposed to capture spatial and temporal dependencies in these data. *(Too complicated for edge devices!)*

Problem II: Distribution Shift, a fundamental problem across data-driven ML

Distribution Shift:

The excellent relies heavily on the critical assumption that the training and inference data are from the same distribution.

Can be easily violated in real-world scenarios and can substantially degrade model performance in many embedded ML applications.



HDC Introduction

Cerebellum


Sparse high dimensional representations

Robustness against noise

Works well with multiple noisy input and computation

Efficient

The brain works at as low as 20W of energy



High-dimensional

Basic elements are hypervectors.

Holographic Encoding

Info of every feature is on all the dimensions

HDC Algebra

Simple and fast, very efficient computation.

- Binding (+):** Element-wise addition

$$\mathcal{H}_{bundle} = \mathcal{H}_1 + \mathcal{H}_2$$

$$\delta(\mathcal{H}_{bundle}, \mathcal{H}_1) \gg 0, \delta(\mathcal{H}_{bundle}, \mathcal{H}_2) \approx 0$$

- Bundling:** Element-wise multiplication

$$\mathcal{H}_{bind} = \mathcal{H}_1 * \mathcal{H}_2$$

$$\delta(\mathcal{H}_{bind}, \mathcal{H}_1) \approx 0, \delta(\mathcal{H}_{bind}, \mathcal{H}_2) \approx 0$$

- Reasoning:** measuring the similarity of hypervectors, e.g., cosine similarity is calculated as $\delta(\mathcal{H}_1, \mathcal{H}_2) = \frac{\mathcal{H}_1 \cdot \mathcal{H}_2}{\|\mathcal{H}_1\| \cdot \|\mathcal{H}_2\|}$

- A powerful learning solution for today's edge platforms

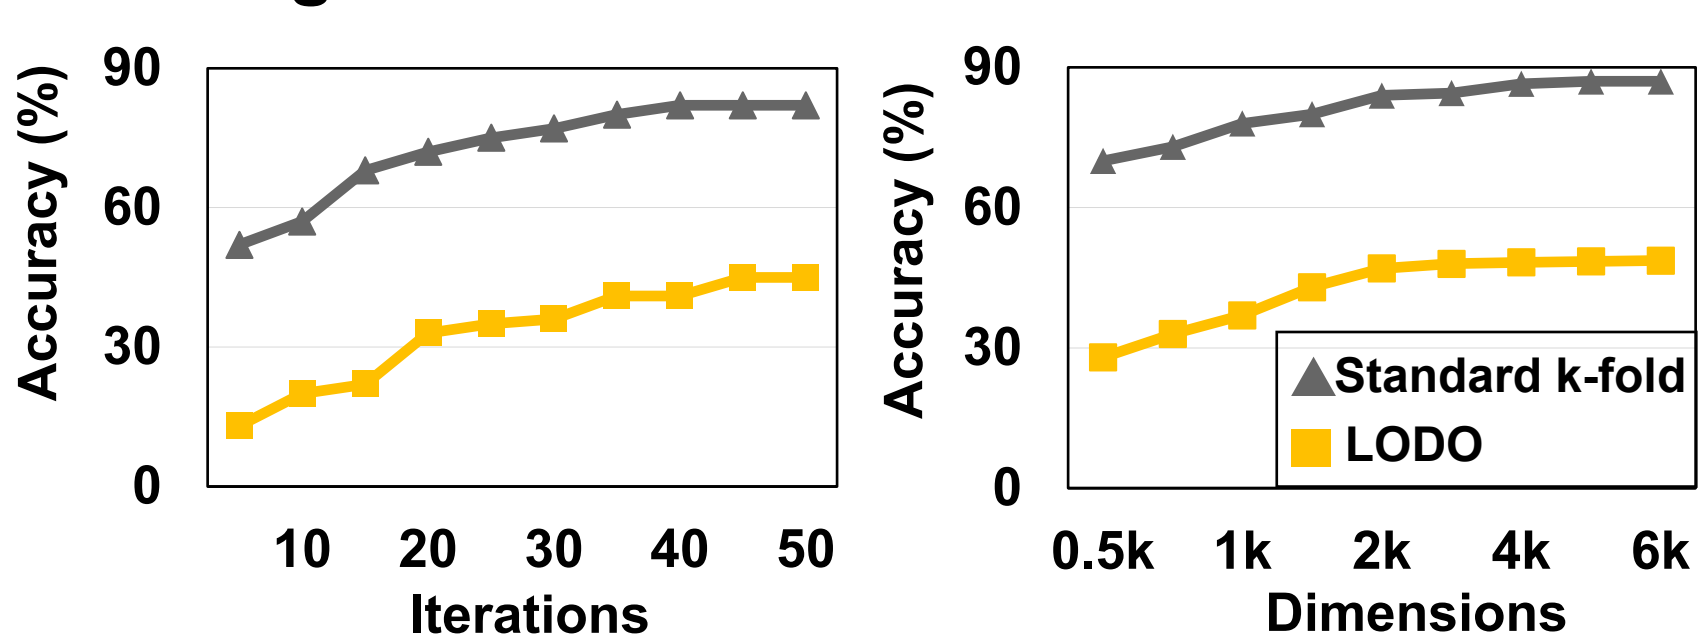
- Fast convergence, high computational efficiency, ultra-robustness against noise
- High-quality results comparable to SOTA DNNs

- Incorporates learning capability along with storing/loading information

- Unique advantages in dealing with time-series data

Challenges

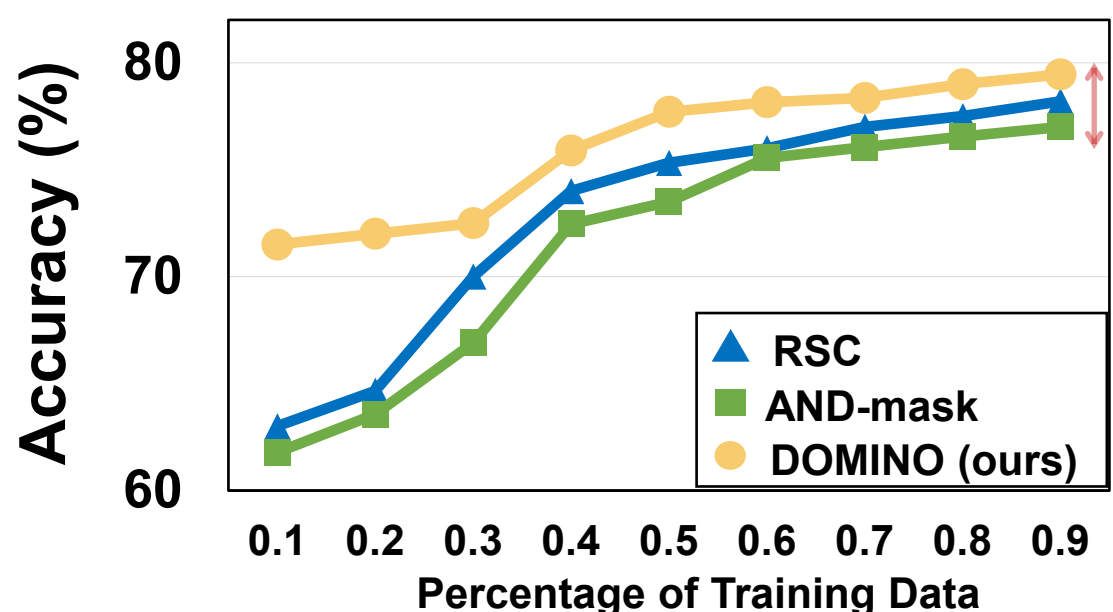
Existing HDCs are not immune to the distribution shift issue




The accuracy of leave-one-domain-out (LODO) cross-validation (CV) is considerably lower than the standard k-fold CV.

A very limited generalization capability of existing models.

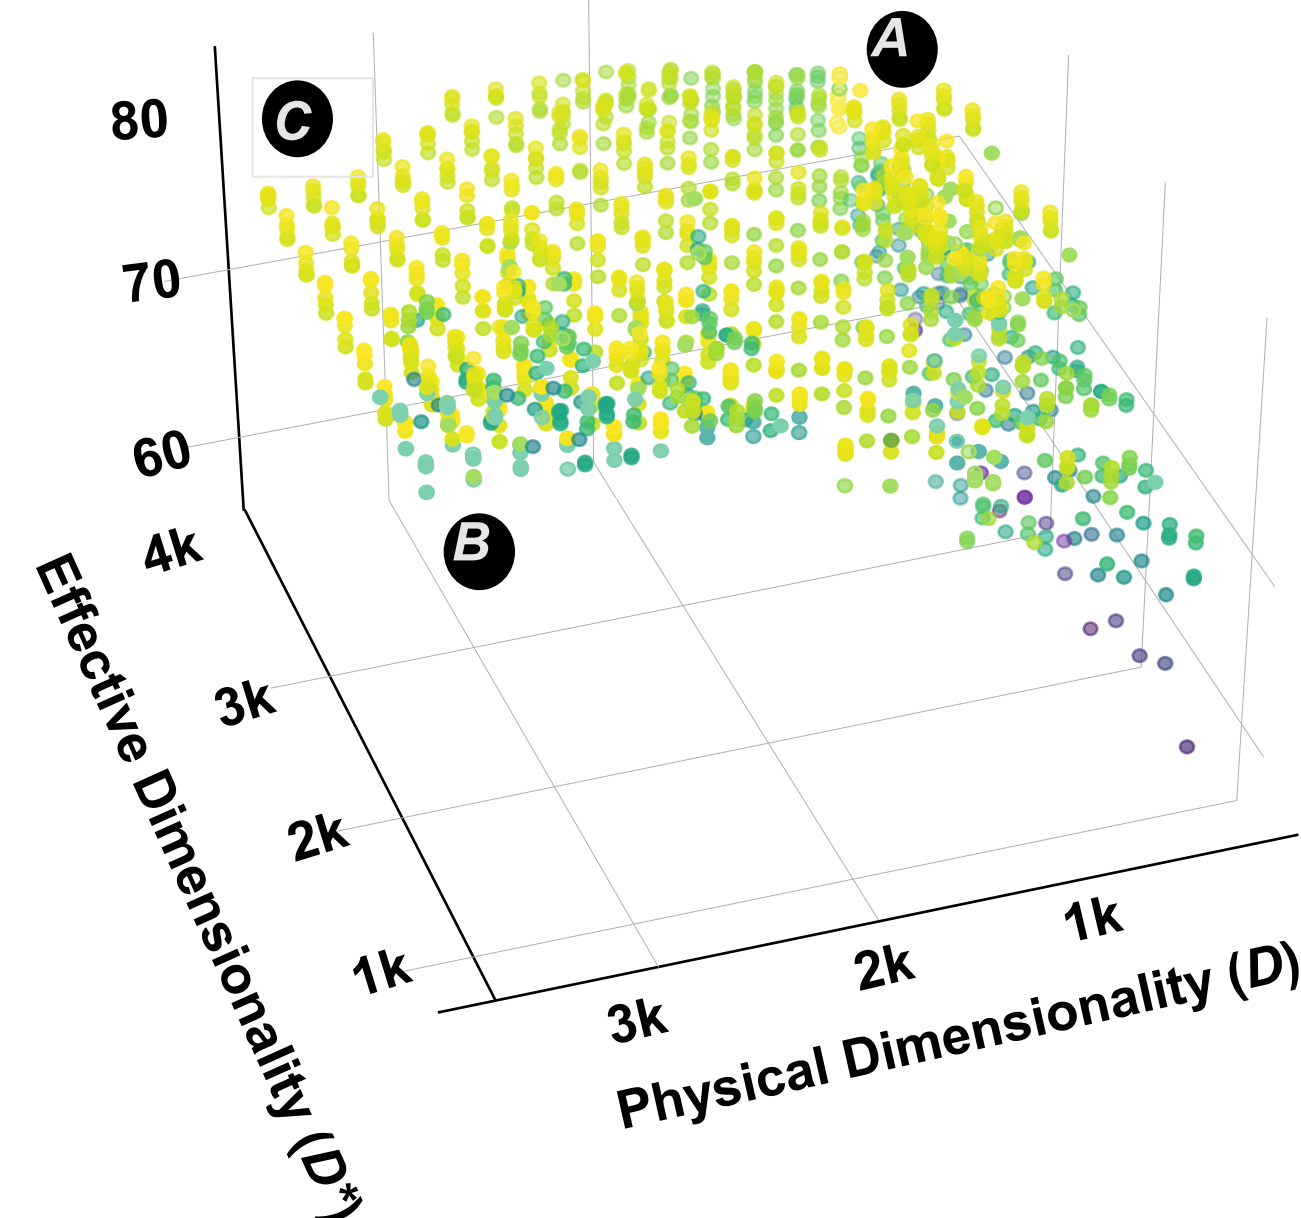
Evaluations



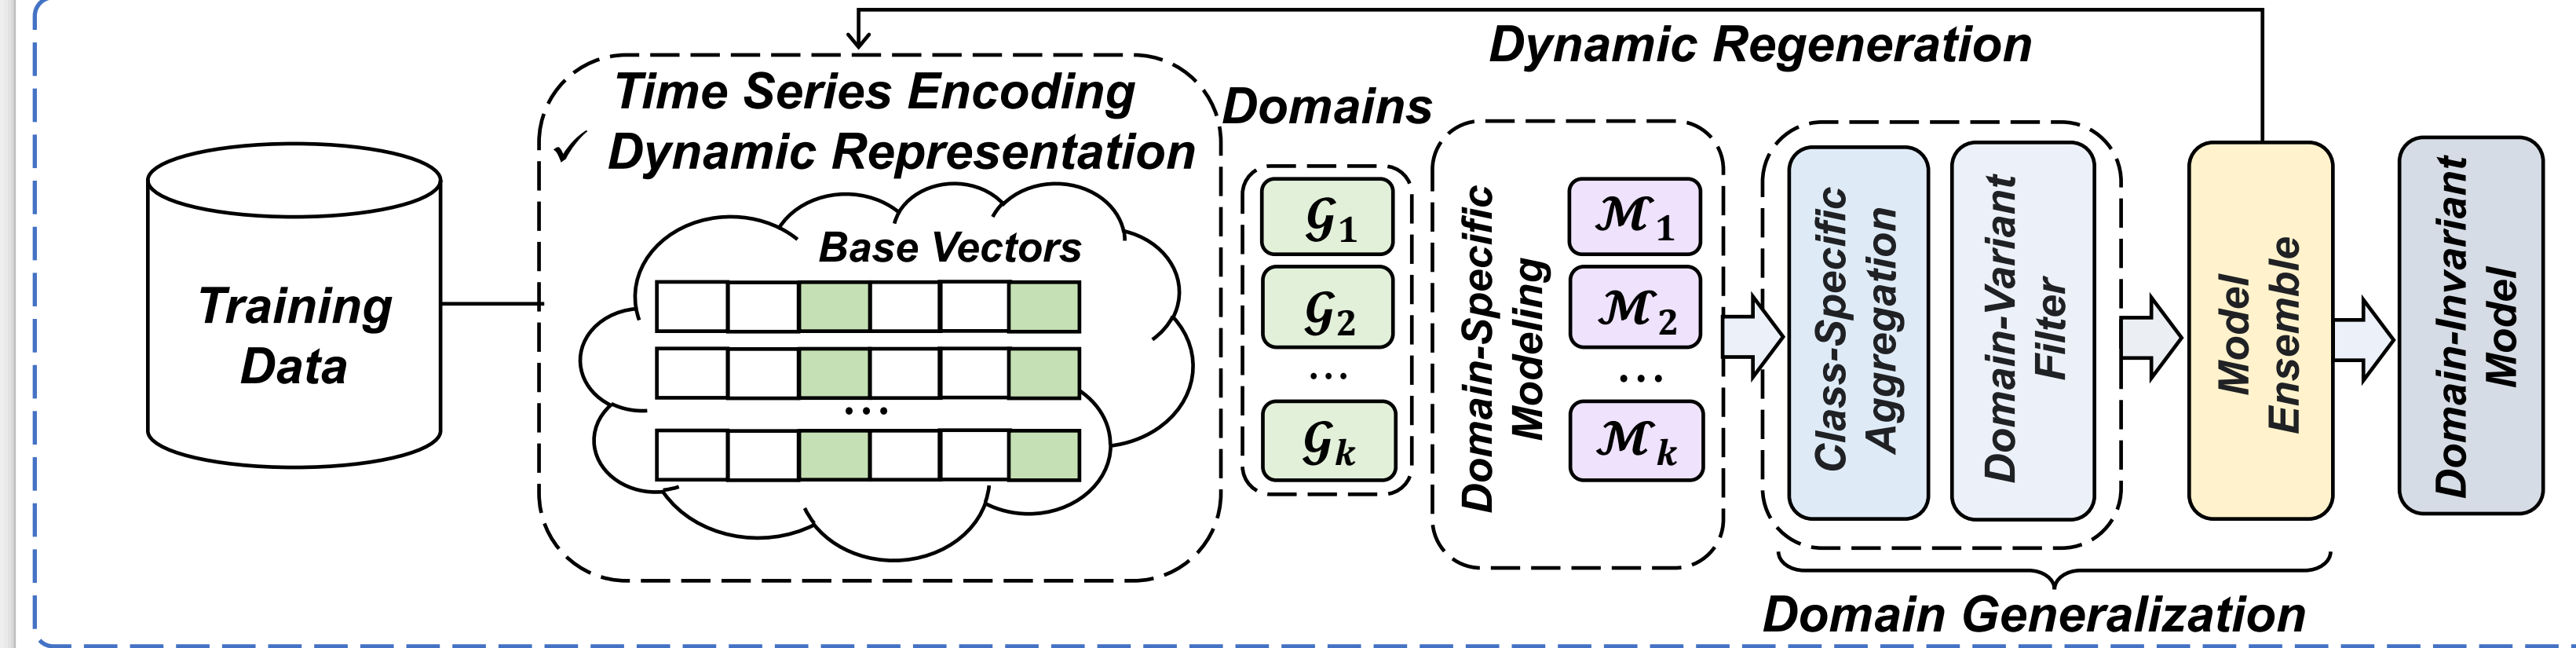
When using only 10% of the training data, DOMINO demonstrates 5.81% higher accuracy than AND-mask and 4.90% higher accuracy than RSC.



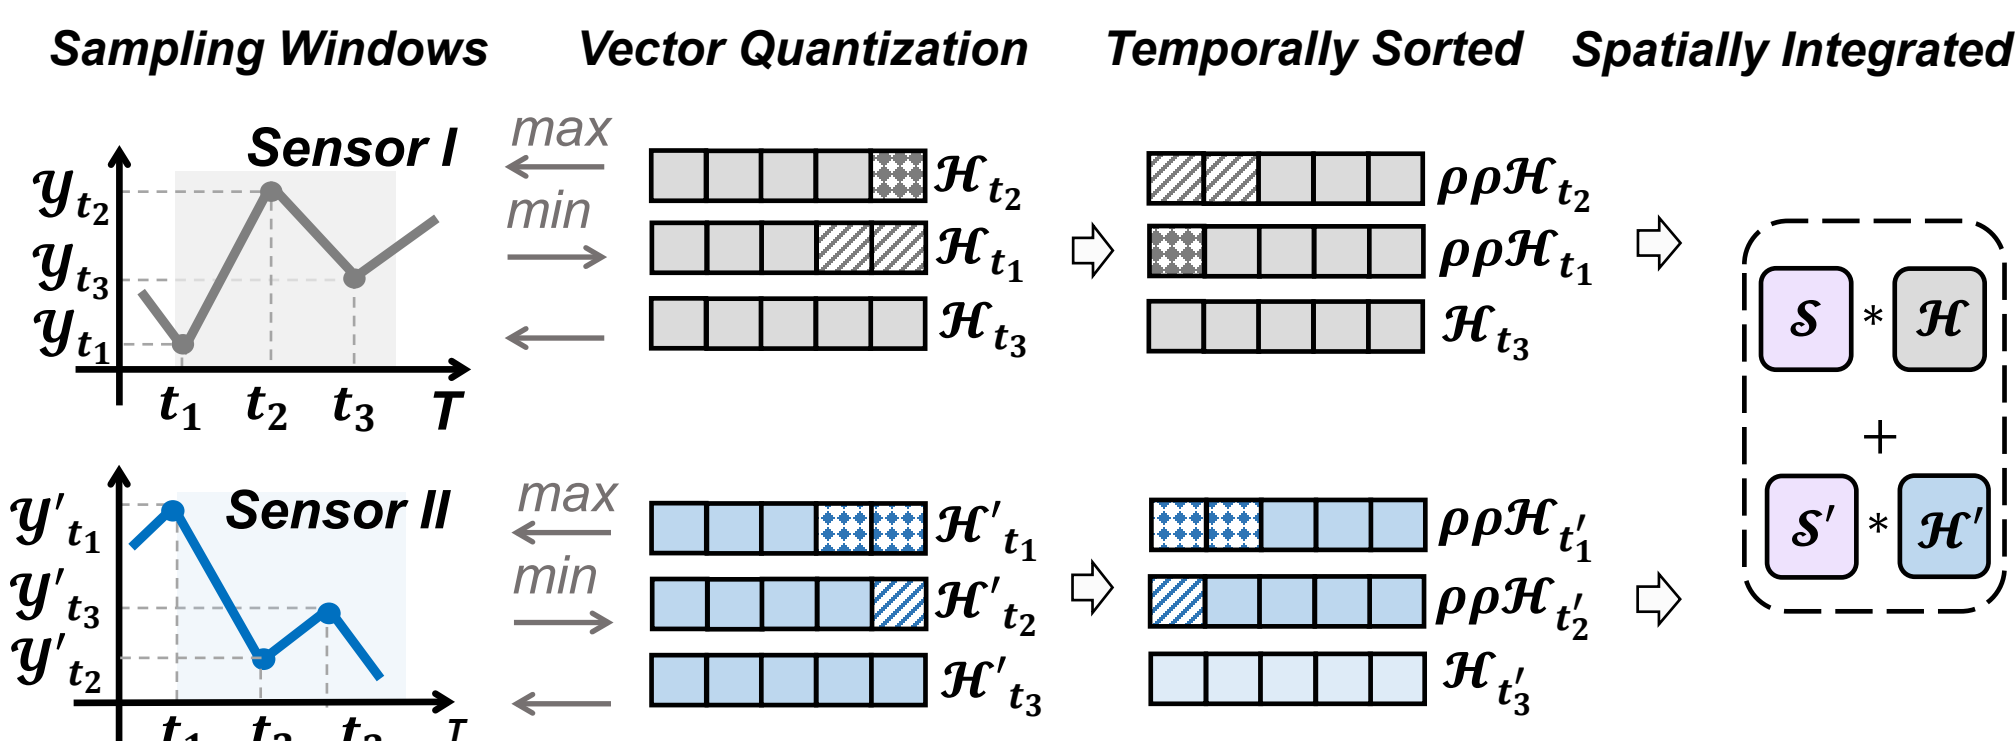
DOMINO outperforms SOTA CNN-based DG approaches by exhibiting on average 1.18% and 2.58% higher accuracy than RSC and AND-mask.



DOMINO: HDC-Based Domain Generalization



Multi-Sensor Time Series Encoding



N-gram Encoding:

- Assign random hypervectors \mathcal{H}_{max} and \mathcal{H}_{min} to represent the maximum and minimum signal values, i.e., ψ_{max} and ψ_{min} .

- Vector quantization to values between ψ_{max} and ψ_{min} . For instance,

$$\mathcal{H}_{t_3} = \mathcal{H}_{t_1} + \frac{y_{t_3} - y_{t_1}}{y_{t_2} - y_{t_1}} \cdot (\mathcal{H}_{t_2} - \mathcal{H}_{t_1});$$

$$\mathcal{H}'_{t_3} = \mathcal{H}'_{t_2} + \frac{y'_{t_3} - y'_{t_2}}{y'_{t_1} - y'_{t_2}} \cdot (\mathcal{H}'_{t_1} - \mathcal{H}'_{t_2})$$

- Temporally sorting by rotation shifts (ρ), e.g., $\mathcal{H} = \rho\rho\mathcal{H}_{t_1} * \rho\mathcal{H}_{t_2} * \mathcal{H}_{t_3}$
- Spatially integrating by binding, e.g., $\mathcal{H} = \mathcal{S}_1 * \mathcal{H}_1 + \dots + \mathcal{S}_n * \mathcal{H}_n$

Domain-Specific Modeling

Algorithm 1 Domain-Specific Modeling
Input: N encoded training samples $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_N$ with n classes and k domains, domains for training data $\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_k$, label of each data sample $\mathcal{L}_1, \mathcal{L}_2, \dots, \mathcal{L}_N$, learning rate η .
Output: k domain-specific models $\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_k$ after one training iteration.

- for each $\lambda \in [1, k]$ do
- Initialize a domain-specific model \mathcal{M}_λ consisting of one class hypervectors for each class $\mathcal{M}_\lambda = \{C_\lambda^1, C_\lambda^2, \dots, C_\lambda^c\}$
- for each $\mathcal{H}_i \in \{\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_N\}$ do
- if $\text{domain}(\mathcal{H}_i) = \mathcal{G}_\lambda$ then
- $C_{max} = \arg \max_{C_\lambda^x} \{\delta(\mathcal{H}_i, C_\lambda^1), \dots, \delta(\mathcal{H}_i, C_\lambda^c)\}$
- if $\mathcal{L}_i = C_{max}$ then
- continue
- else if $\mathcal{L}_i \neq C_{max} \wedge \mathcal{L}_i = C_\lambda^x$ then
- $C_{max} \leftarrow C_{max} - \eta \cdot [1 - \delta(\mathcal{H}_i, C_{max})] \times \mathcal{H}_i$
- $C_\lambda^x \leftarrow C_\lambda^x + \eta \cdot [1 - \delta(\mathcal{H}_i, C_\lambda^x)] \times \mathcal{H}_i$
- $\mathcal{M}_\lambda = \text{Normalize}(\mathcal{M}_\lambda)$
- return $\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_k$

Our algorithm provides a higher chance for non-common patterns to be properly included

$$\delta(\mathcal{H}, C_\lambda^t) = \frac{\mathcal{H} \cdot C_\lambda^t}{\|\mathcal{H}\| \cdot \|C_\lambda^t\|} = \frac{\mathcal{H}}{\|\mathcal{H}\|} \cdot \frac{C_\lambda^t}{\|C_\lambda^t\|}$$

- A large $\delta(\mathcal{H}, \cdot)$ indicates the input data points is marginally mismatches
- A small $\delta(\mathcal{H}, \cdot)$ indicates a noticeably new pattern

Domain Generalization (DG)

Algorithm 2 Domain Generalization

Input: k domain-specific models $\{\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_k\}$ each with size $n \times D$, regeneration rate \mathcal{R} .
Output: Domain-variant dimensions \mathcal{U} to drop.

- Initialize n empty matrices $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_n$, each with size $k \times D$.
- for each $\mathcal{T}_\gamma \in \{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_n\}$ do
- for each $\mathcal{M}_\lambda \in \{\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_k\}$ do
- $\mathcal{T}_\gamma[\lambda, :] = \mathcal{M}_\lambda[\gamma, :]$ ▷ Form n class-specific matrices
- $\sigma_\gamma = \text{Variance}(\mathcal{T}_\gamma, \text{columnwise})$ ▷ $\text{dim}(\sigma_\gamma) = 1 \times D$
- $\mathcal{V} = \sum_{\gamma=1}^n \sigma_\gamma$
- $\mathcal{U} = \text{argsort}(\mathcal{V}) \llbracket [(1 - \mathcal{R}) \cdot D] : D \rrbracket$ ▷ \mathcal{R} of dimensions with largest variance
- return \mathcal{U}

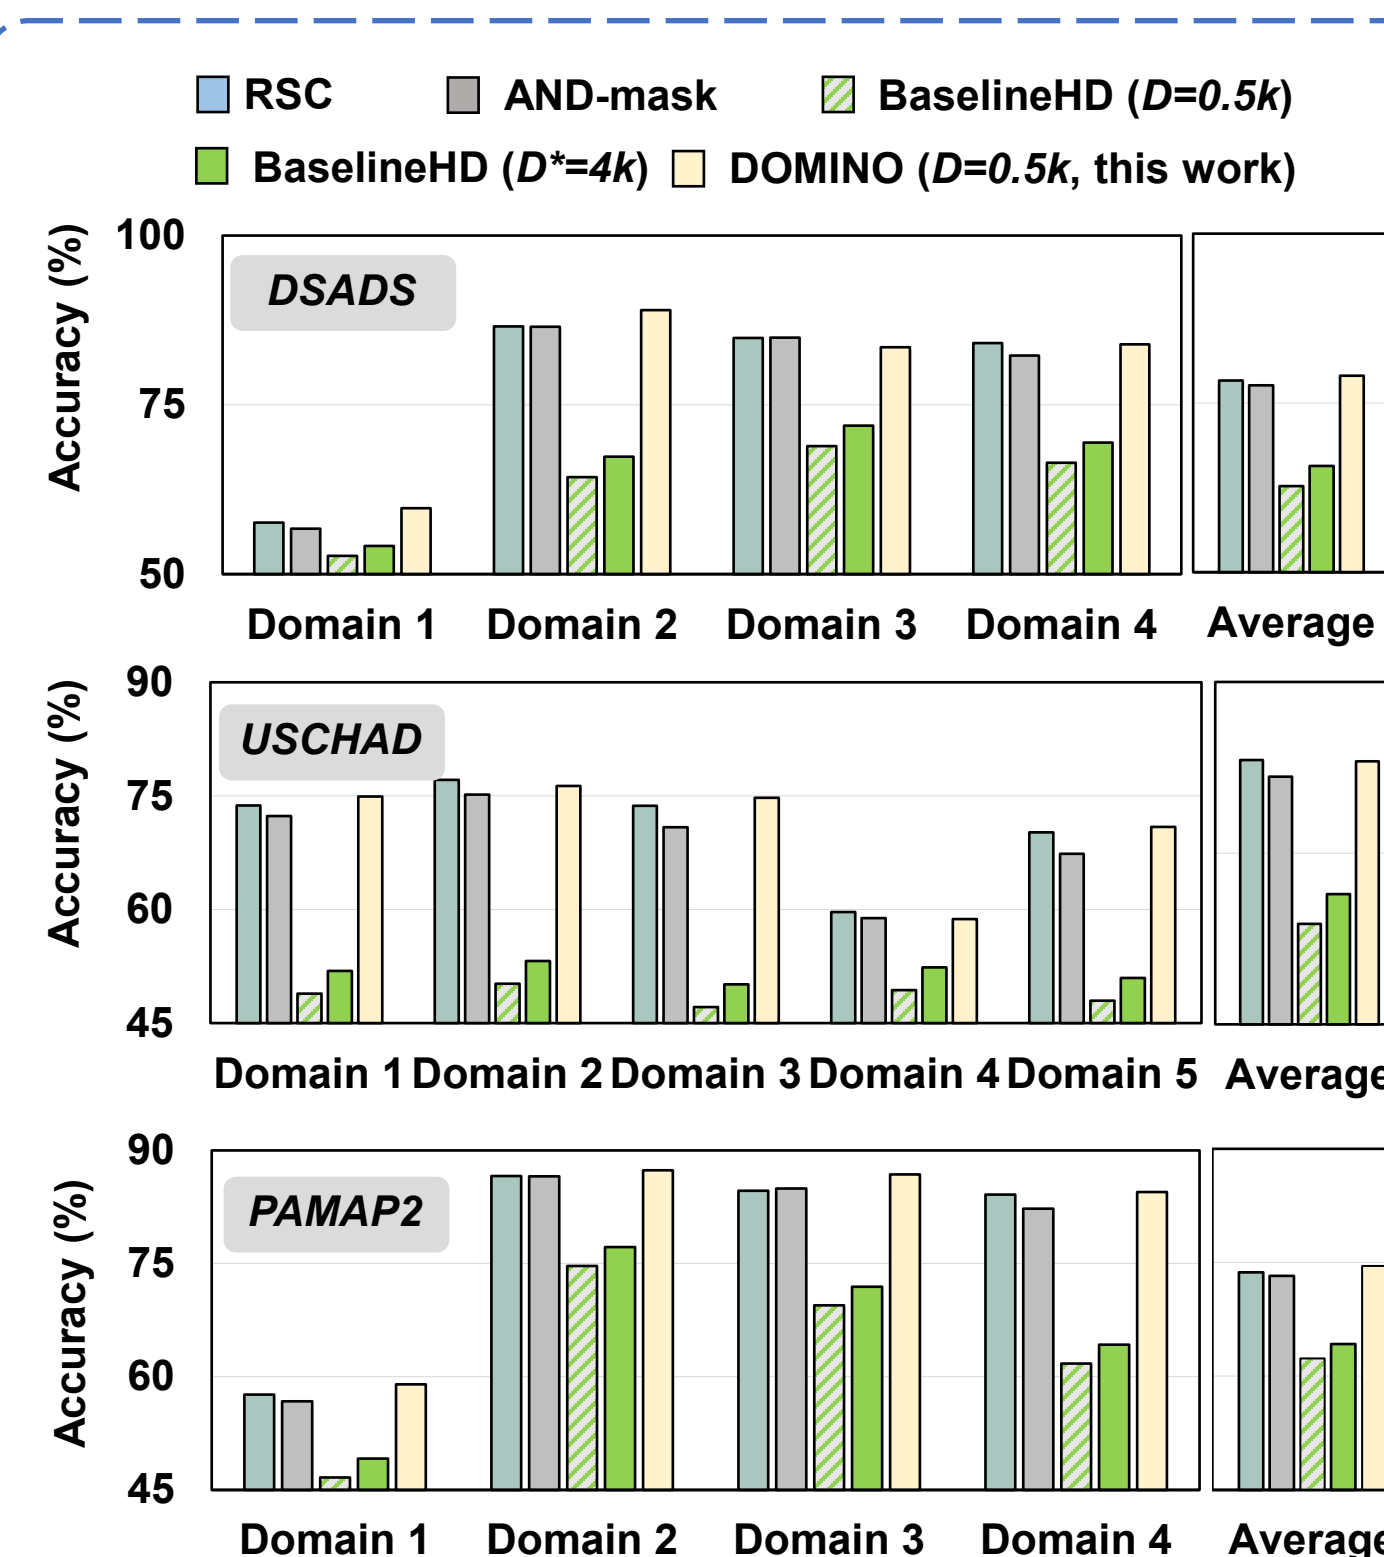
- Dimensions with large variance indicate, for the same class, these dimensions store very different information, and are hence considered domain-variant.

- We sum up the variance vector of each class-specific matrix to obtain a vector representing the overall relevance of dimension to domains.

- We select the top \mathcal{R} portion of dimensions with the highest variance and regenerate each of them with a new randomly generated vector.

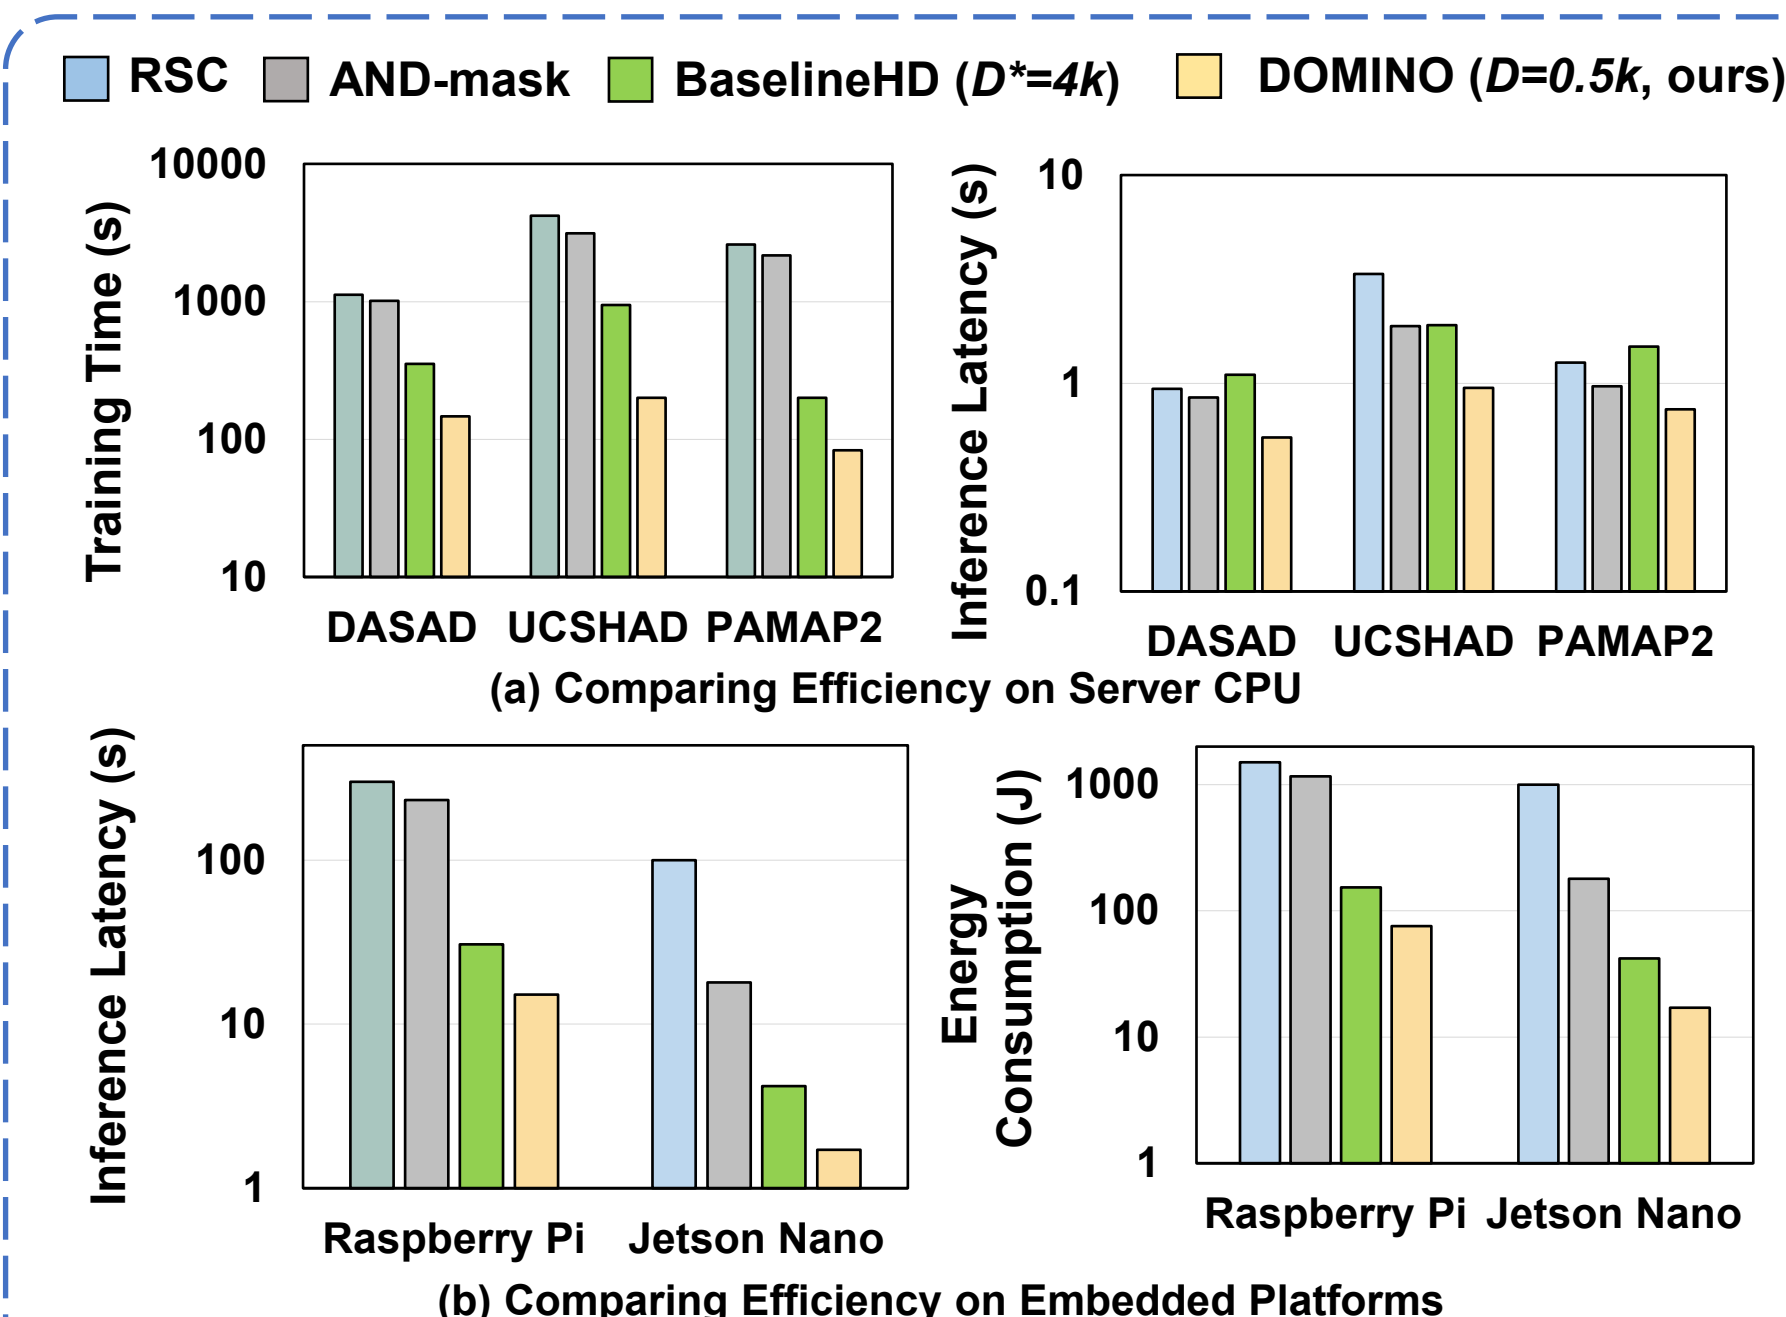
Model Ensemble

Evaluations



LODO Accuracy:

0.96% higher than RSC, 2.04% higher than AND-mask
11.70% higher than BaselineHD ($D^* = 4k$), 15.93% higher than BaselineHD ($D = 0.5k$)



On Server CPU, DOMINO exhibits:

- 16.34x faster training than RSC, 14.17x faster training than AND-mask
- 2.89x faster inference than RSC, 1.97x faster inference than AND-mask

On embedded devices:

- Raspberry Pi:
 - 19.79x faster than RSC, 15.31x faster than AND-mask
- Jetson Nano
 - 58.44x faster than RSC, 10.49x faster than AND-mask

Conclusion:

We propose DOMINO, an HDC-based domain generalization algorithm that provides significantly higher efficiency and better performance than SOTA DNN-based techniques. Our solution provides a resource-efficient and hardware-friendly solution, especially for today's edge devices, to mitigate the distribution shift challenge.