

LORE: Jointly Learning the Intrinsic Dimensionality and Relative Similarity Structure From Ordinal Data

The Occam's Razor of Ordinal Embeddings

Vivek Anand, Alec Helbling, Mark A. Davenport, Gordon Berman, Sankaraleengam Alagapan, Christopher John Rozell

ICLR 2026 Presentation.



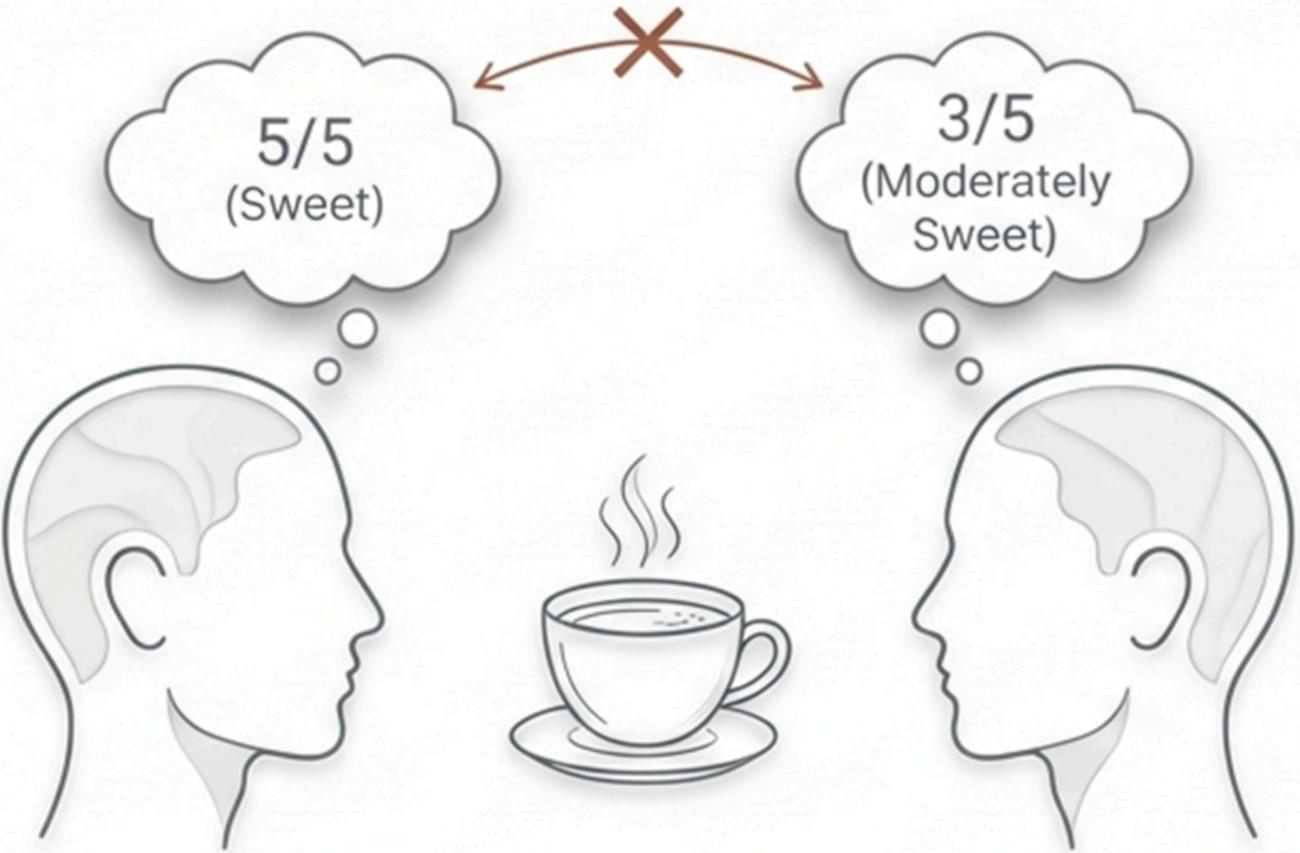
PAPER: arxiv.org/abs/2602.04192



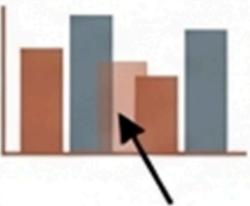
CODE: github.com/vivek2000anand/lore_iclr

To Map Perception, Measure Relationships not Ratings

The Problem: Absolute Ratings

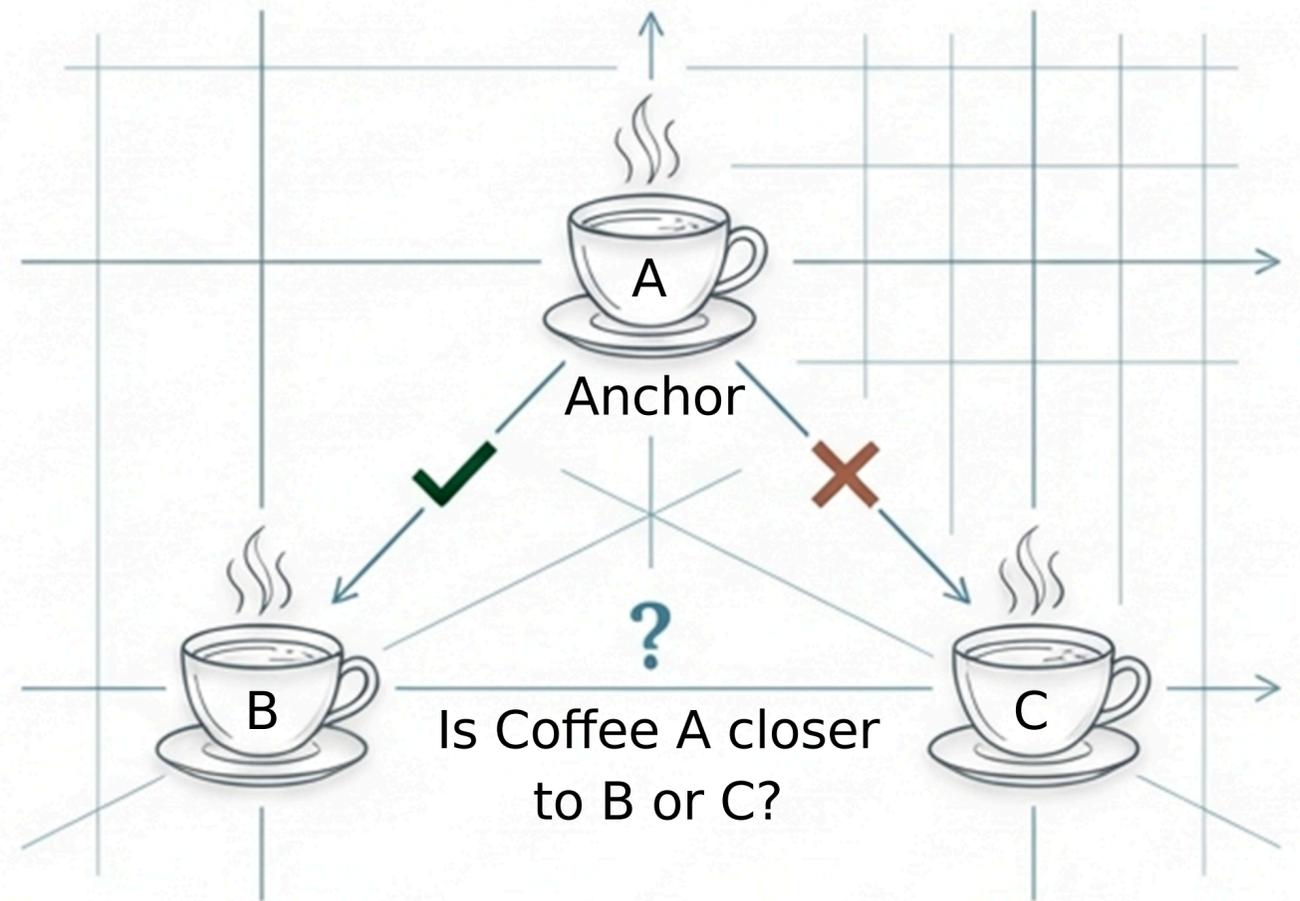


Absolute Ratings (Likert) Suffer from calibration bias. One person's "5" is another's "3"

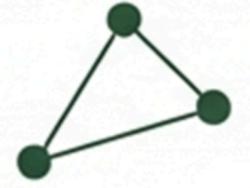


Calibration Noise

The Solution: Relative Queries

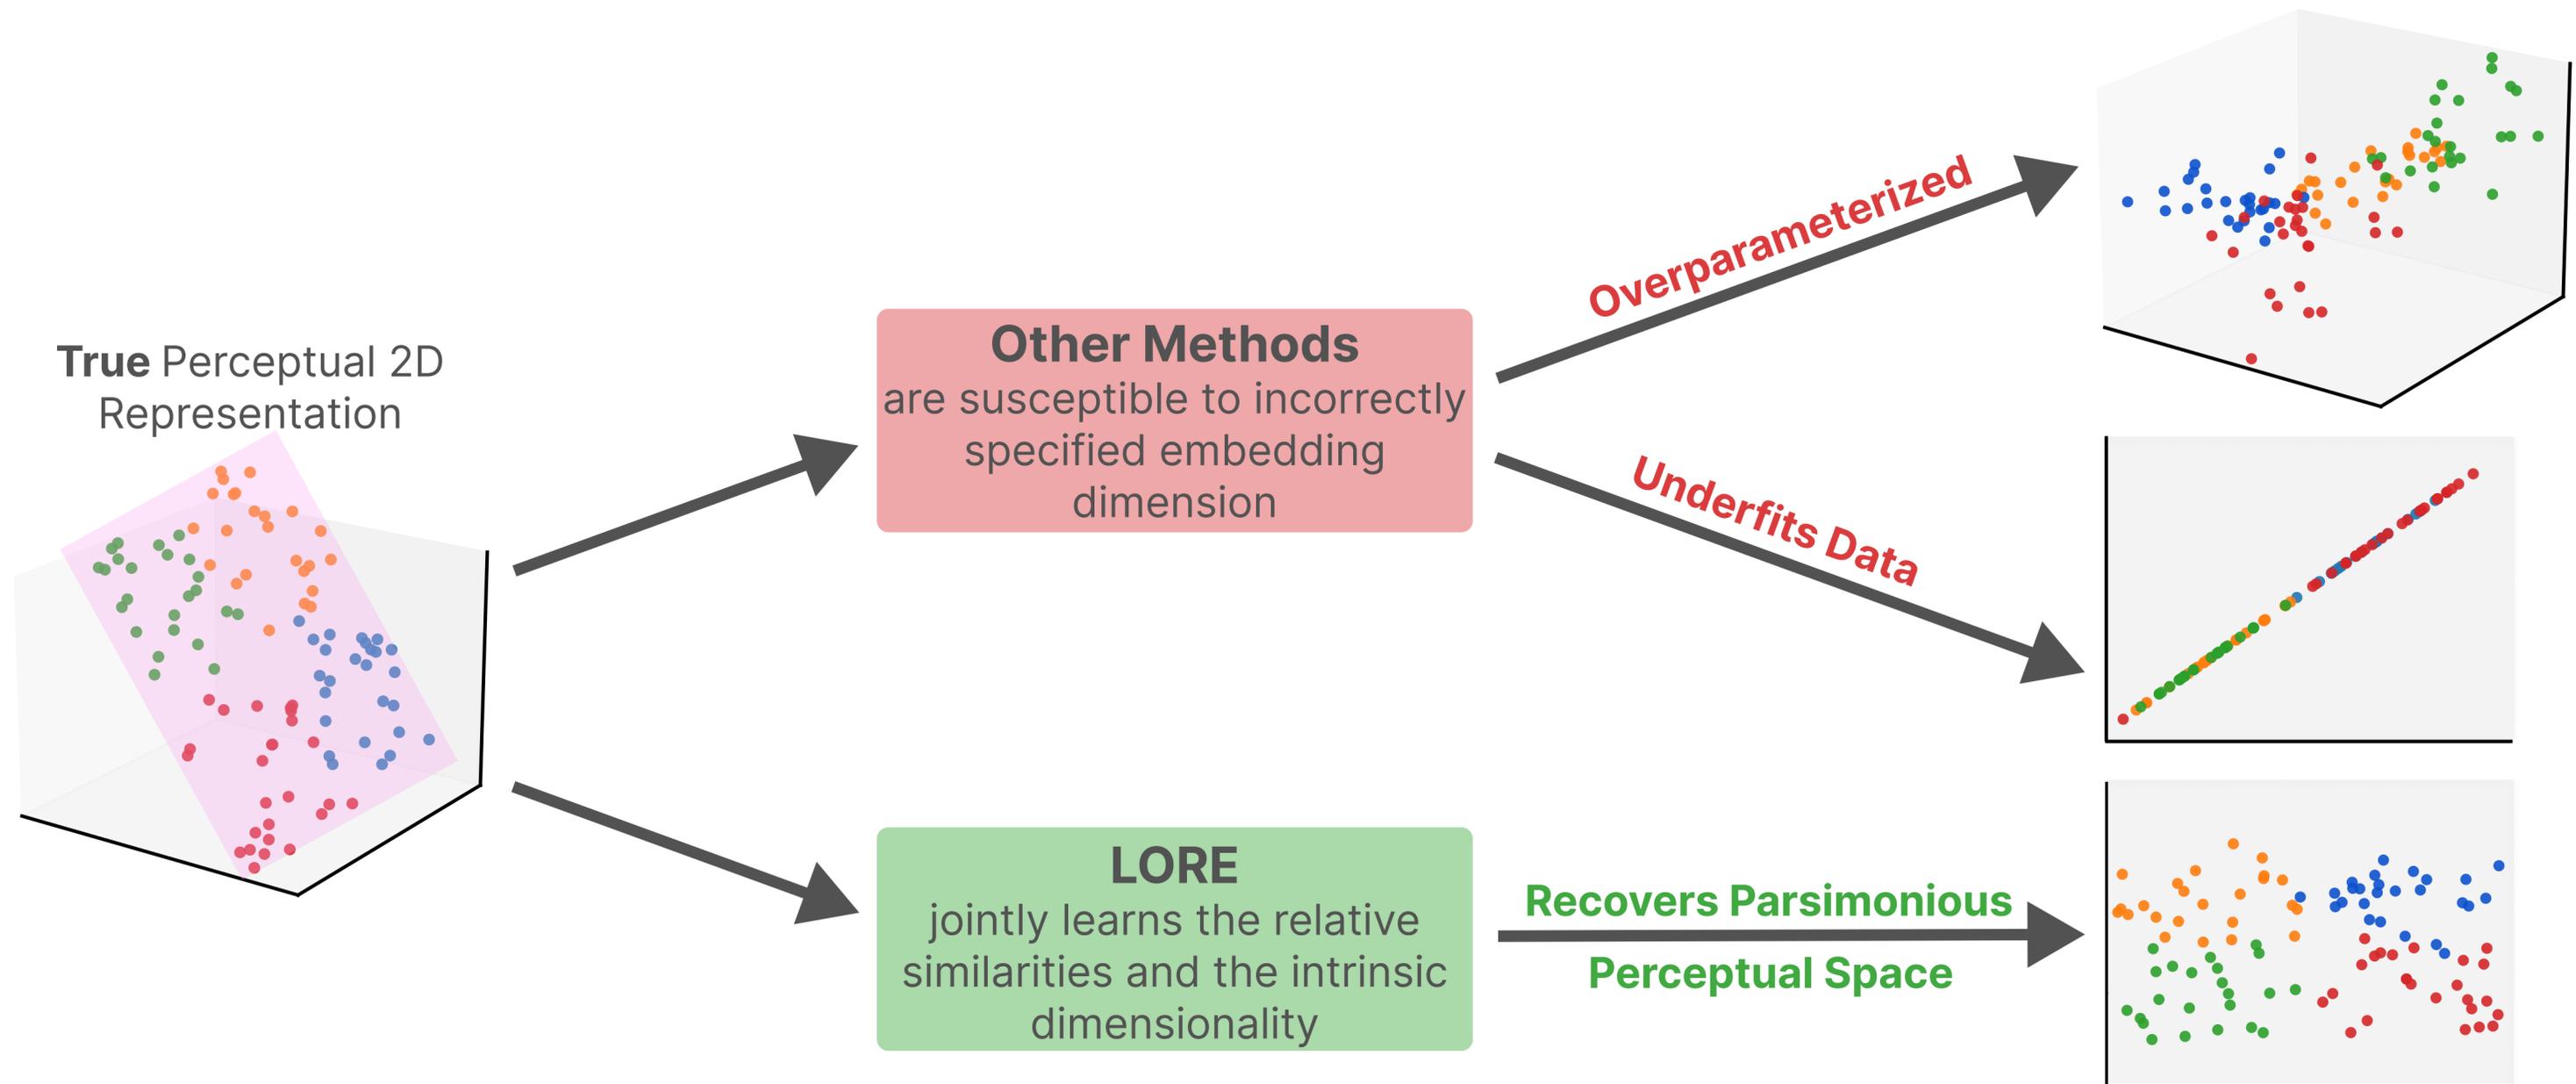


Relative Queries (i, j, k) capture structure without scale bias



Invariant Structure

The Conundrum: Needing to Guess Dimensionality



The LORE Objective Function

$$\min_{\mathbf{Z}} \Psi(\mathbf{Z}) = \sum_{(a,i,j) \in T} \log(1 + \exp(1 + d(\mathbf{Z}_{a,:}, \mathbf{Z}_{i,:}) - d(\mathbf{Z}_{a,:}, \mathbf{Z}_{j,:}))) + \lambda \sum_{i=1}^{\min\{N,d'\}} \sigma_i(\mathbf{Z})^p$$

Fit the Data

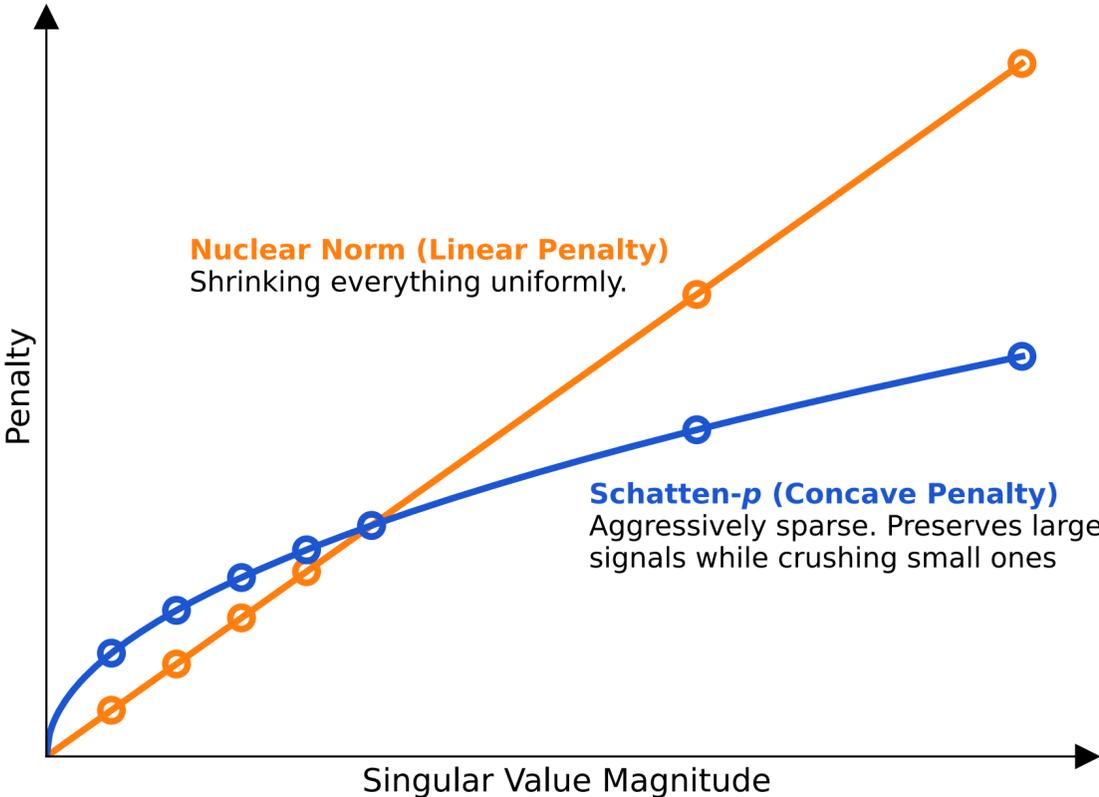
Penalize Complexity

Softplus Triplet Loss
Ensures the map respects the relative choices
(A is closer to i than j)

Schatten- p Quasi-Norm ($0 < p < 1$)
Unlike the Nuclear Norm, this aggressively pushes the noise singular values to zero.

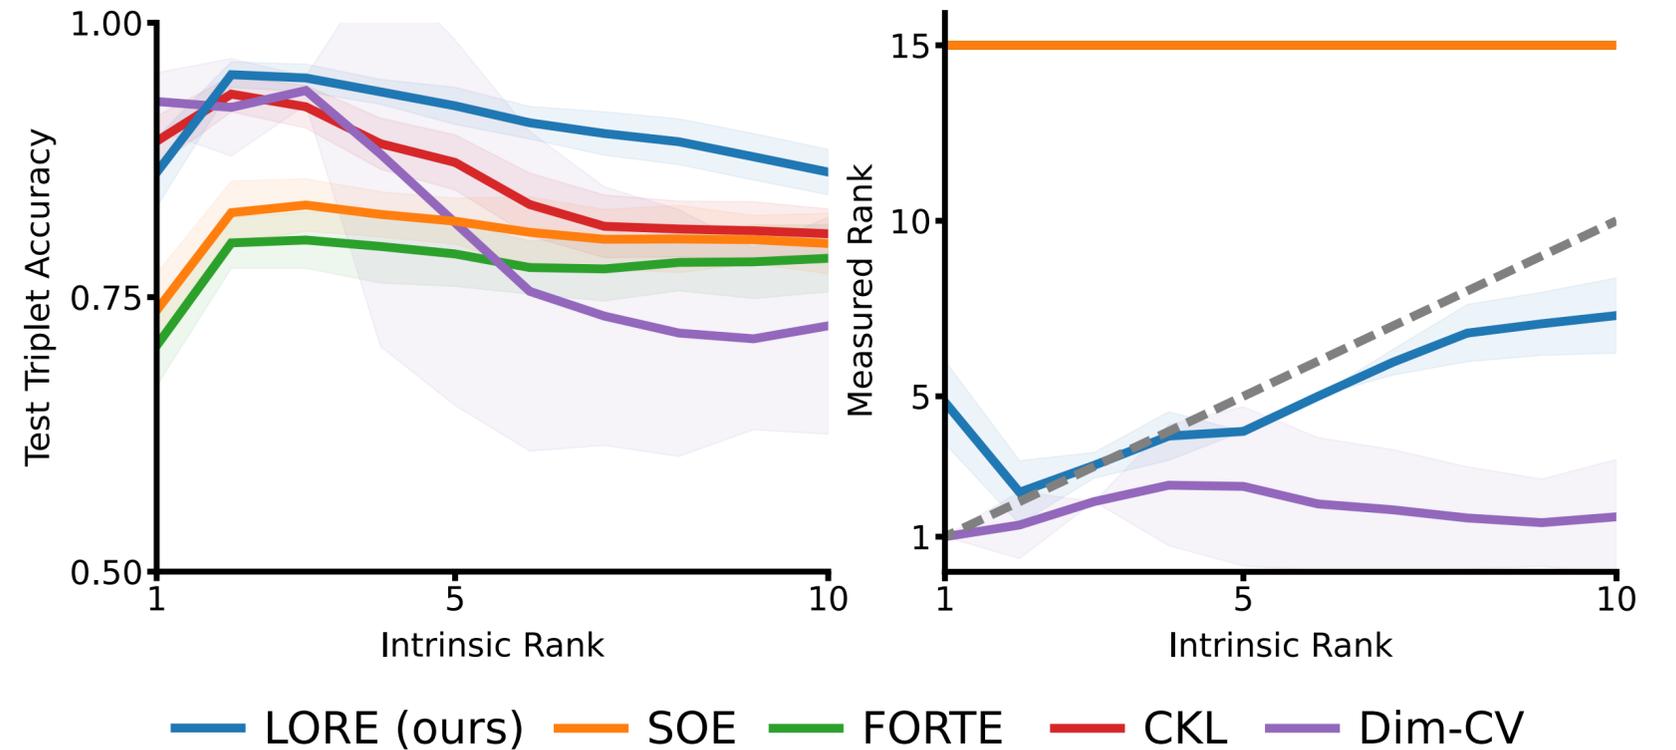
Theorem 1: Convergence Guaranteed

Despite the non-convex objective, LORE's iteratively reweighted algorithm is guaranteed to converge to a stationary point.



Results: Recovering the Perceptual space

Only LORE can recover the intrinsic dimensionality with high accuracy in a simulated perceptual experiment

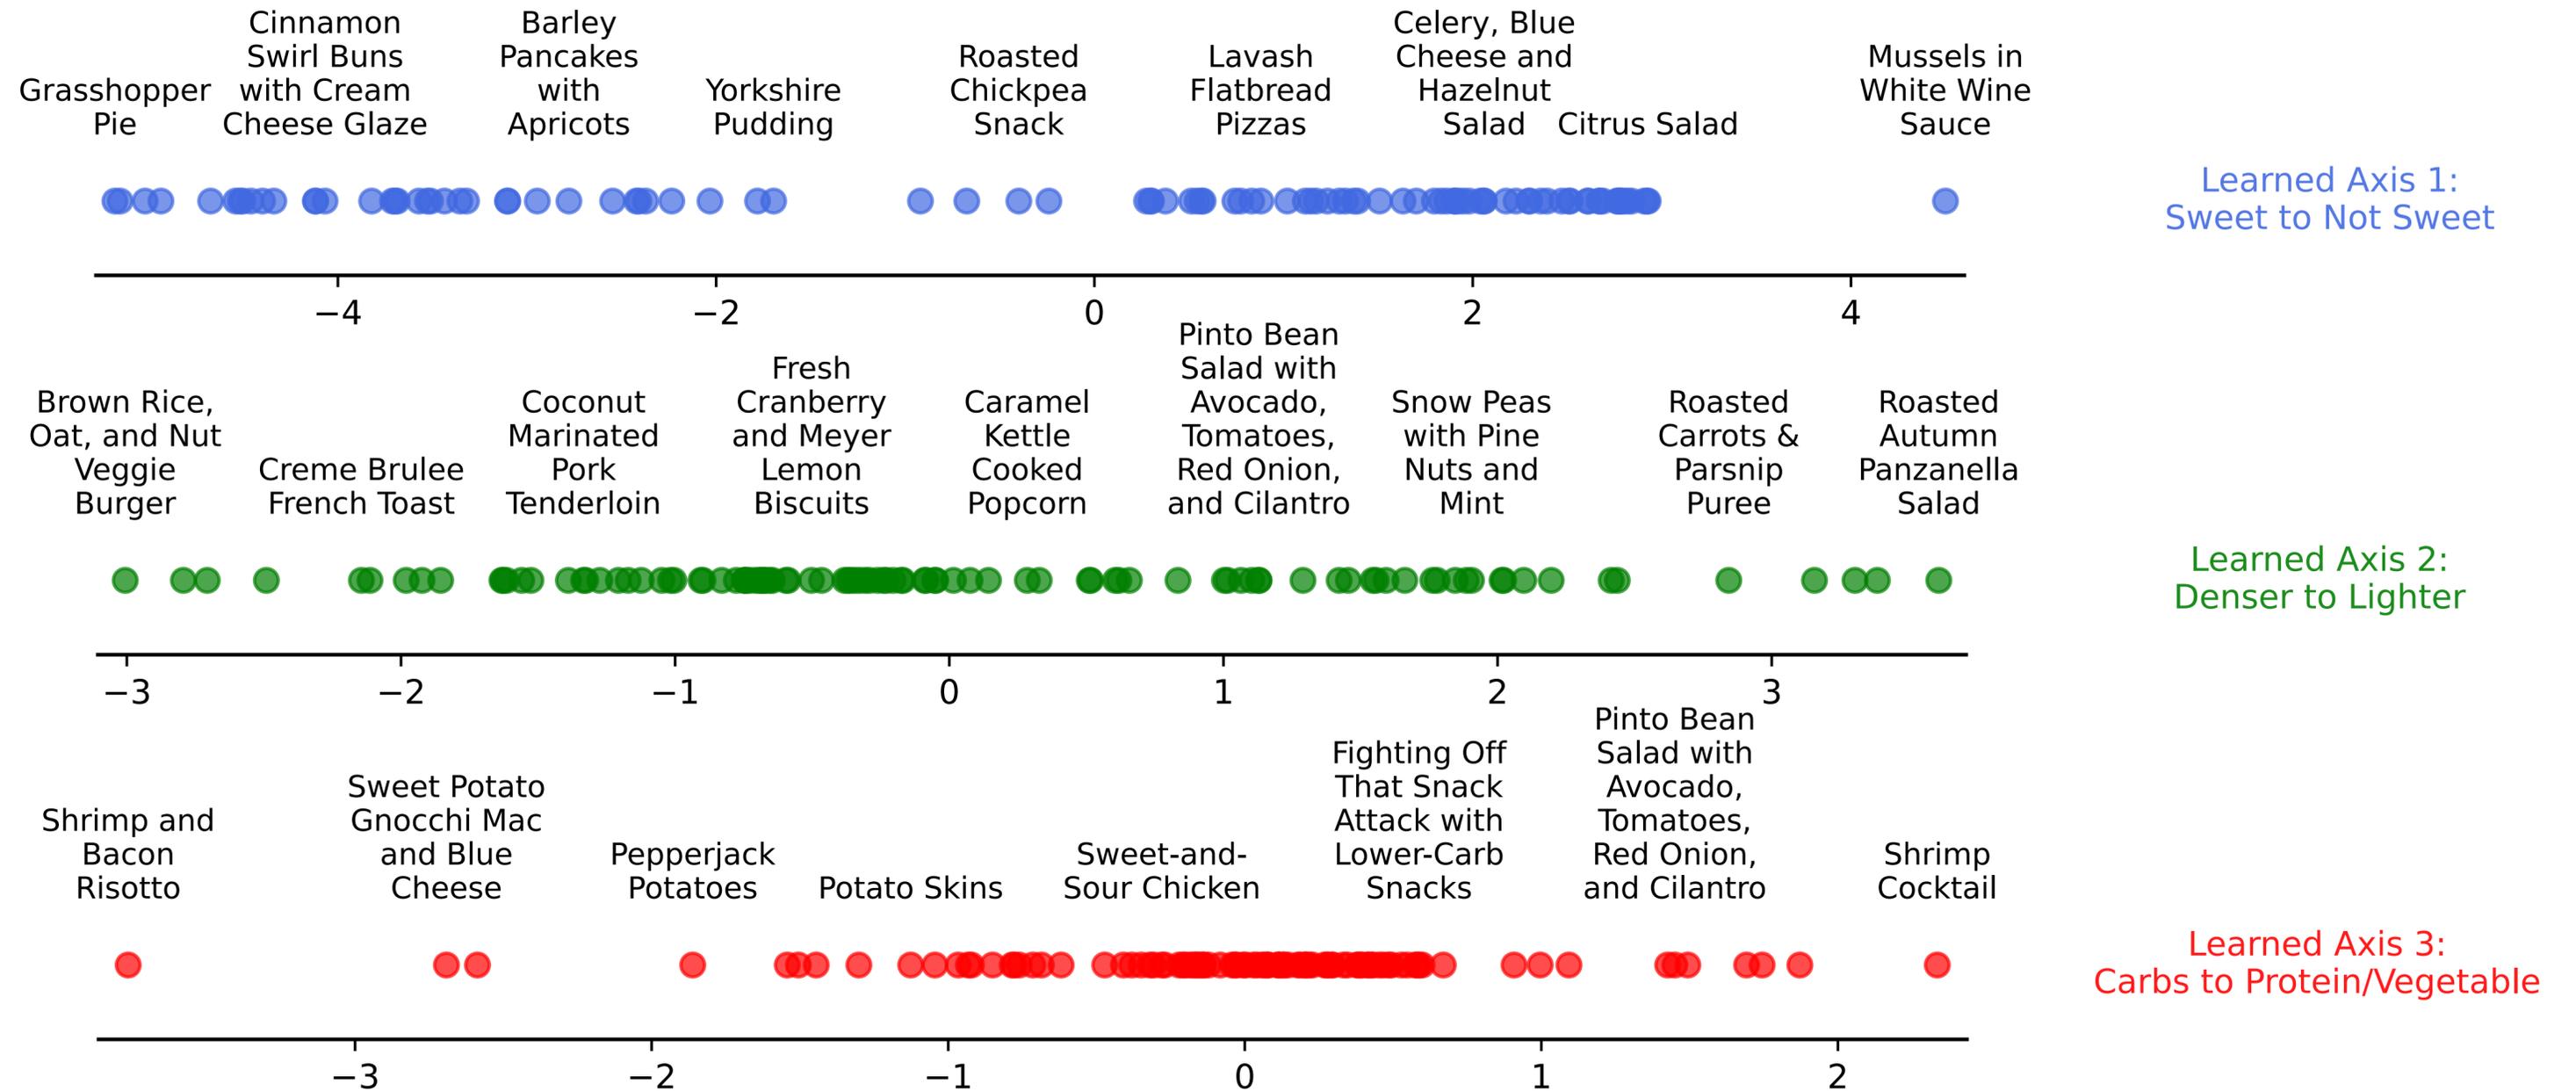


Only LORE learns low rank solutions while maintaining high accuracy on Food-100 dataset

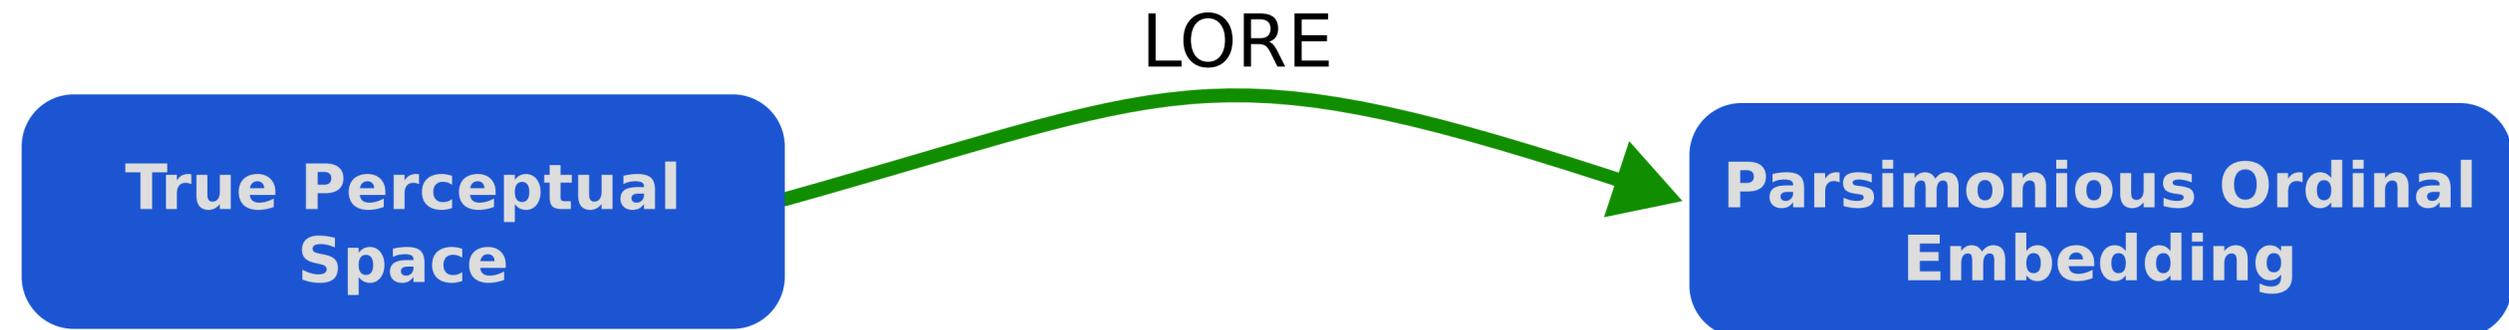
Method	Test Acc.	Rank	Time (s)
LORE	82.45 ± 0.27	3.3 ± 0.47	6.64 ± 3.90
SOE	82.34 ± 0.32	15 ± 0.00	27.09 ± 1.38
FORTE	81.73 ± 0.46	15 ± 0.00	6.34 ± 0.52
t-STE	82.79 ± 0.24	15 ± 0.00	40.93 ± 20.14
CKL	82.75 ± 0.20	15 ± 0.00	18.41 ± 7.89
Dim-CV	77.67 ± 0.02	1.47 ± 0.51	1721.9 ± 26.71

Semantically Interpretable Axes

Discovered without supervision



LORE: The Occam's Razor for Ordinal Embeddings



LORE removes the guesswork of estimating the intrinsic dimensionality of the perceptual space for ordinal embeddings. By balancing accuracy with parsimony, LORE is able to jointly learn both the intrinsic dimensionality and the relative similarities.

