LORE: The Occam's Razor of Ordinal Embeddings

Vivek Anand[1], Alec Helbling[1], Mark A. Davenport[1], Gordon Berman[2], Sankaraleengam Alagapan[1] & Christopher J. Rozell[1]

[1]Georgia Institute of Technology [2]Emory University

Contact Me: vivek-anand.xyz

The Goal: Build Representations of Human Subjective Experience without making Assumptions

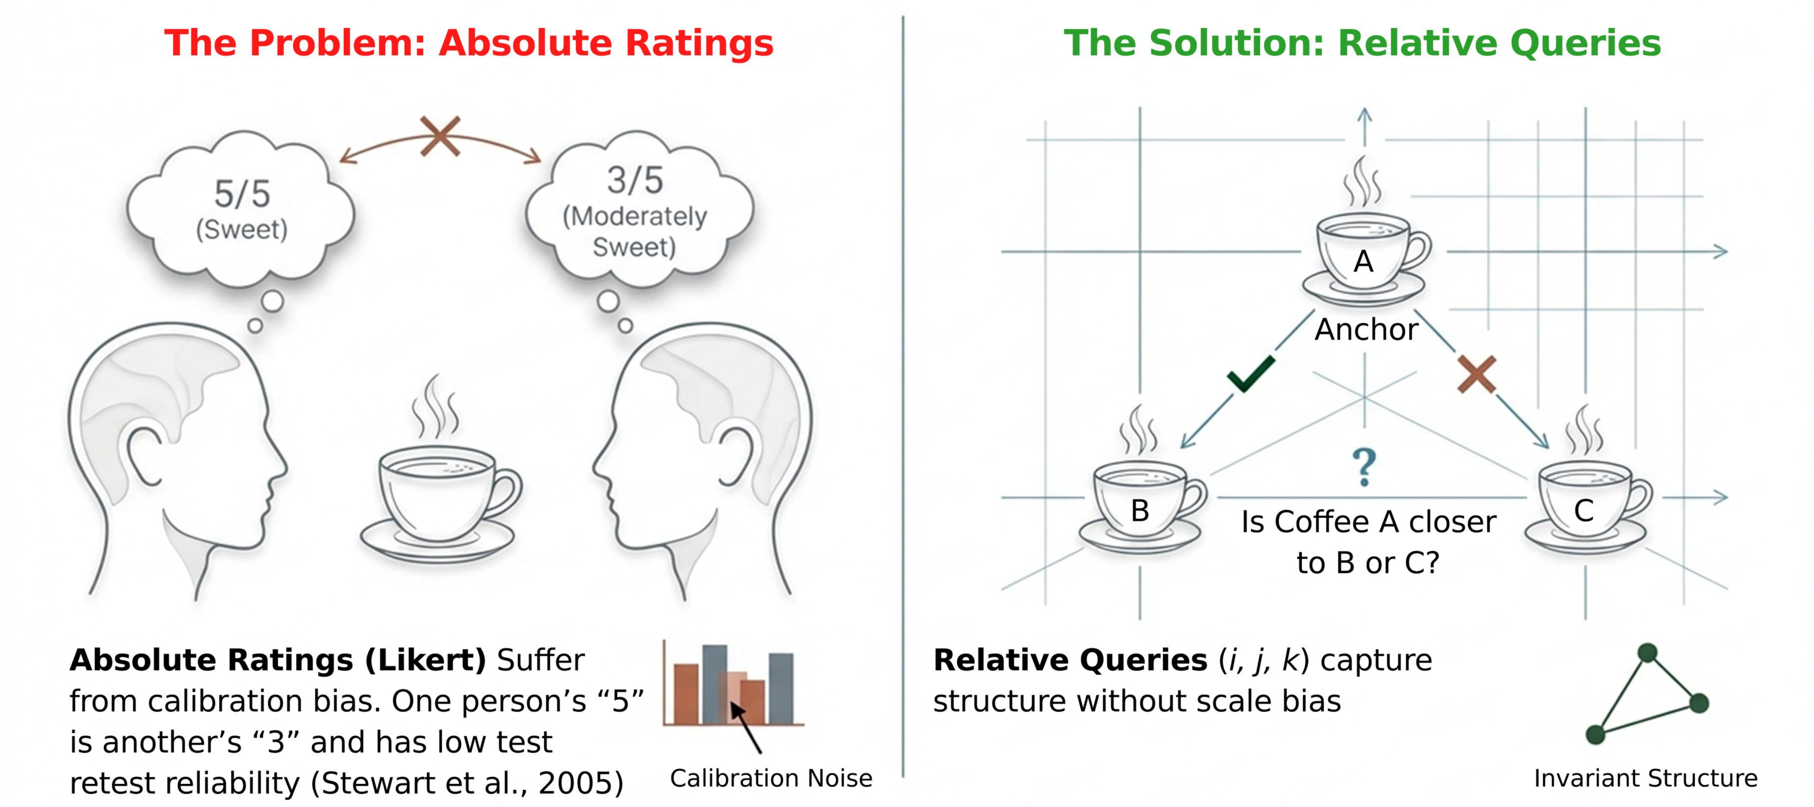

Absolute Ratings (Likert) Suffer from calibration bias. One person's "5" is another's "3" and has low test retest reliability (Stewart et al., 2005)

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

The Ordinal Embedding Pipeline

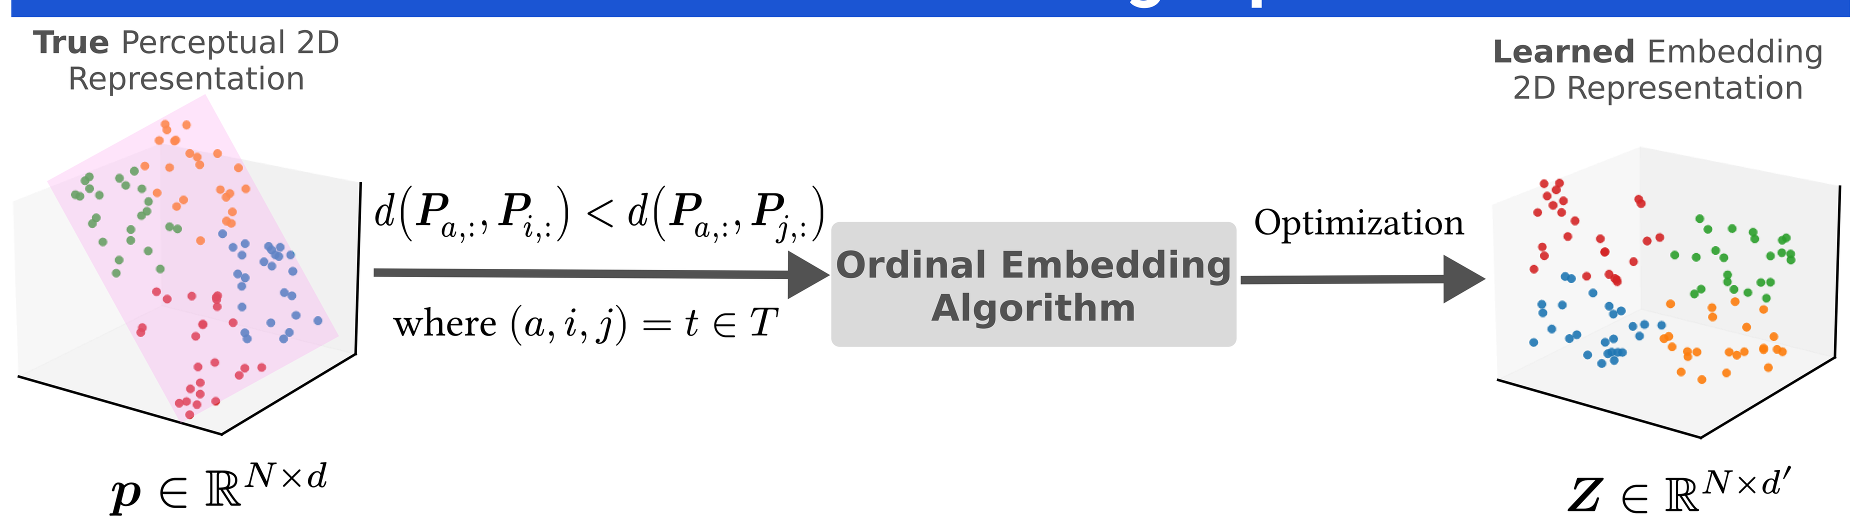

An Ideal Ordinal Embedding algorithm recovers both the relative similarity structure and dimensionality of the perceptual space from triplet responses

Existing Methods Cannot Recover the True Dimensionality

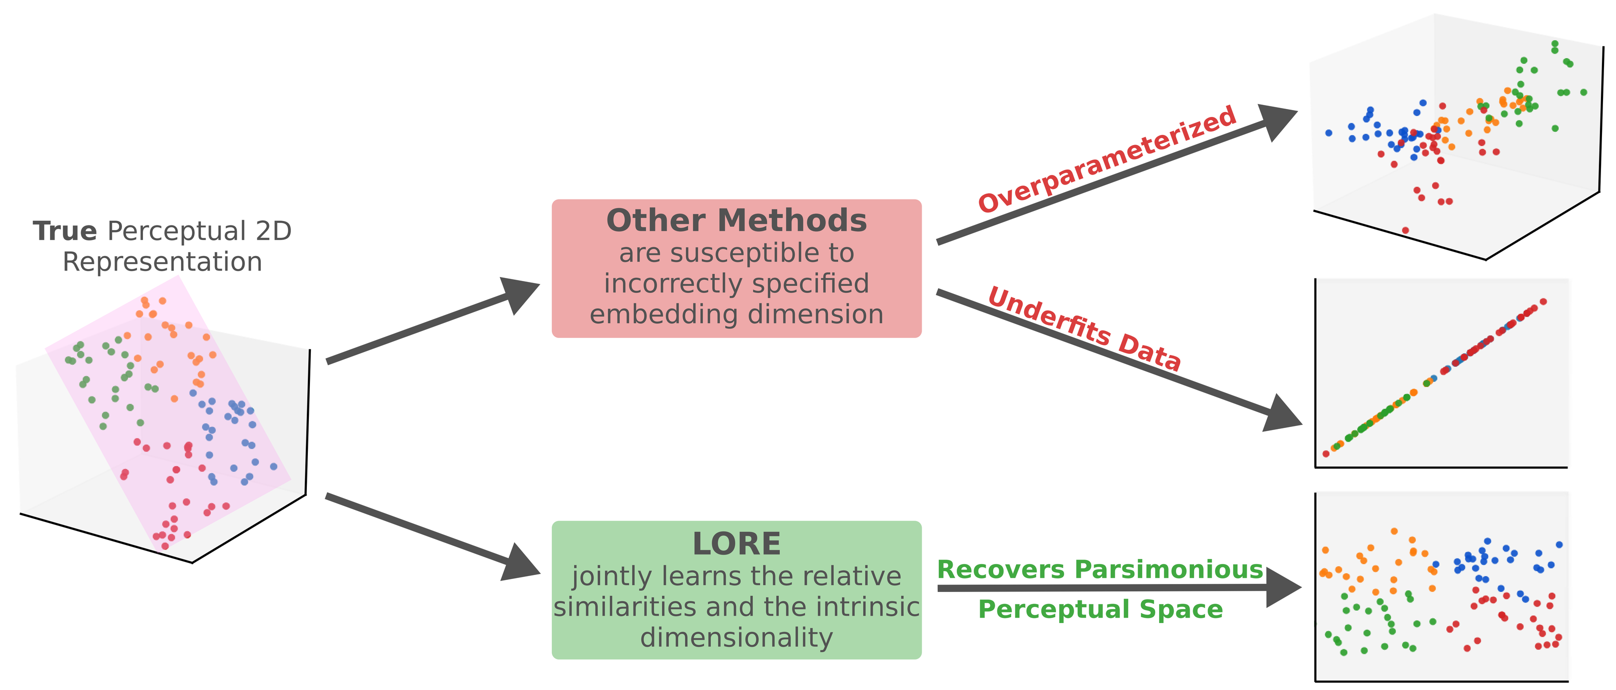

PAPER: arxiv.org/abs/2602.04192

CODE: github.com/vivek2000anand/lore_iclr

Key Insight: Penalize the Space the Embedding Uses

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

Fit the Triplet Responses

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

Penalize Space

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

The Schatten-p Quasi Norm aggressively crushes the small singular values (noise) but preserves larger ones better than the Nuclear Norm can.

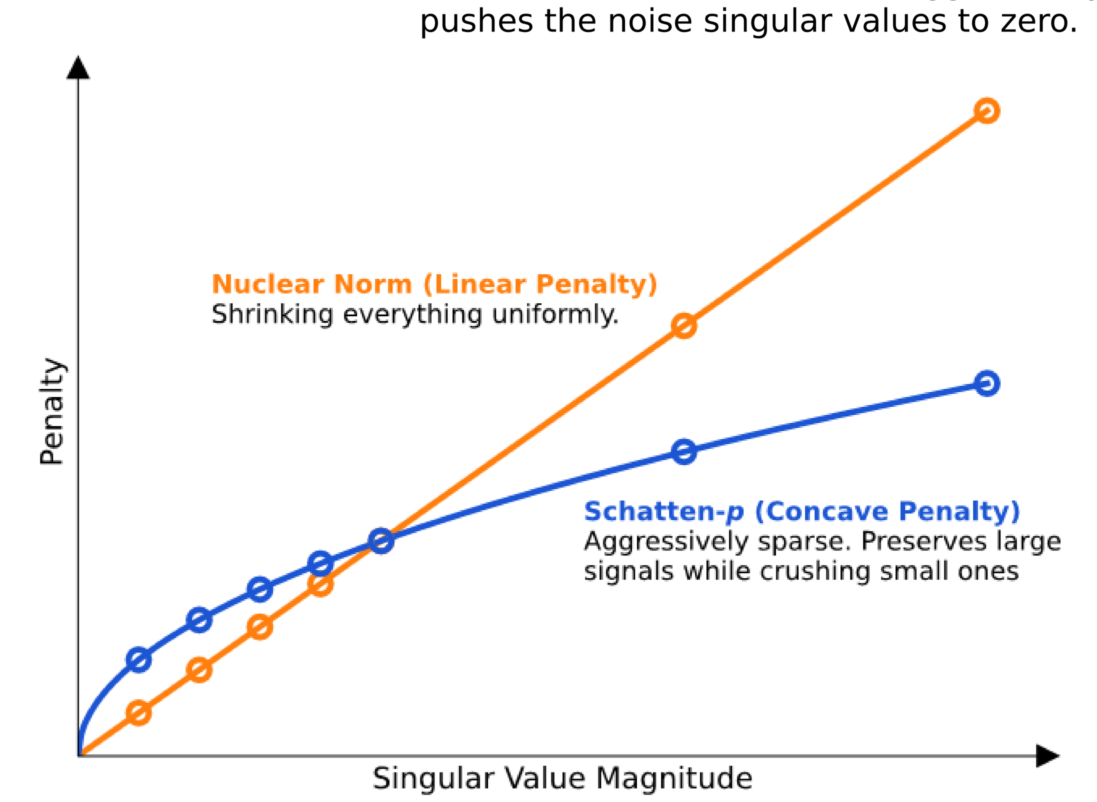

LORE Provably Converges to a Stationary Point

Theorem *(LORE converges to a stationary point)* The sequence of OEs generated by the LORE algorithm $\{Z^k\}_{k=1,2,3,...}$ converges. i.e.

$$\sum_{k=1}^{+\infty} \|Z^{k+1} - Z^k\|_F < +\infty$$

LORE's iteratively reweighted optimization algorithm ensures that the nonconvex objective provably reaches a stationary point. Local minima for ordinal embeddings are thought to be close to global minima. [2]

LORE has a wide and stable regularization setting

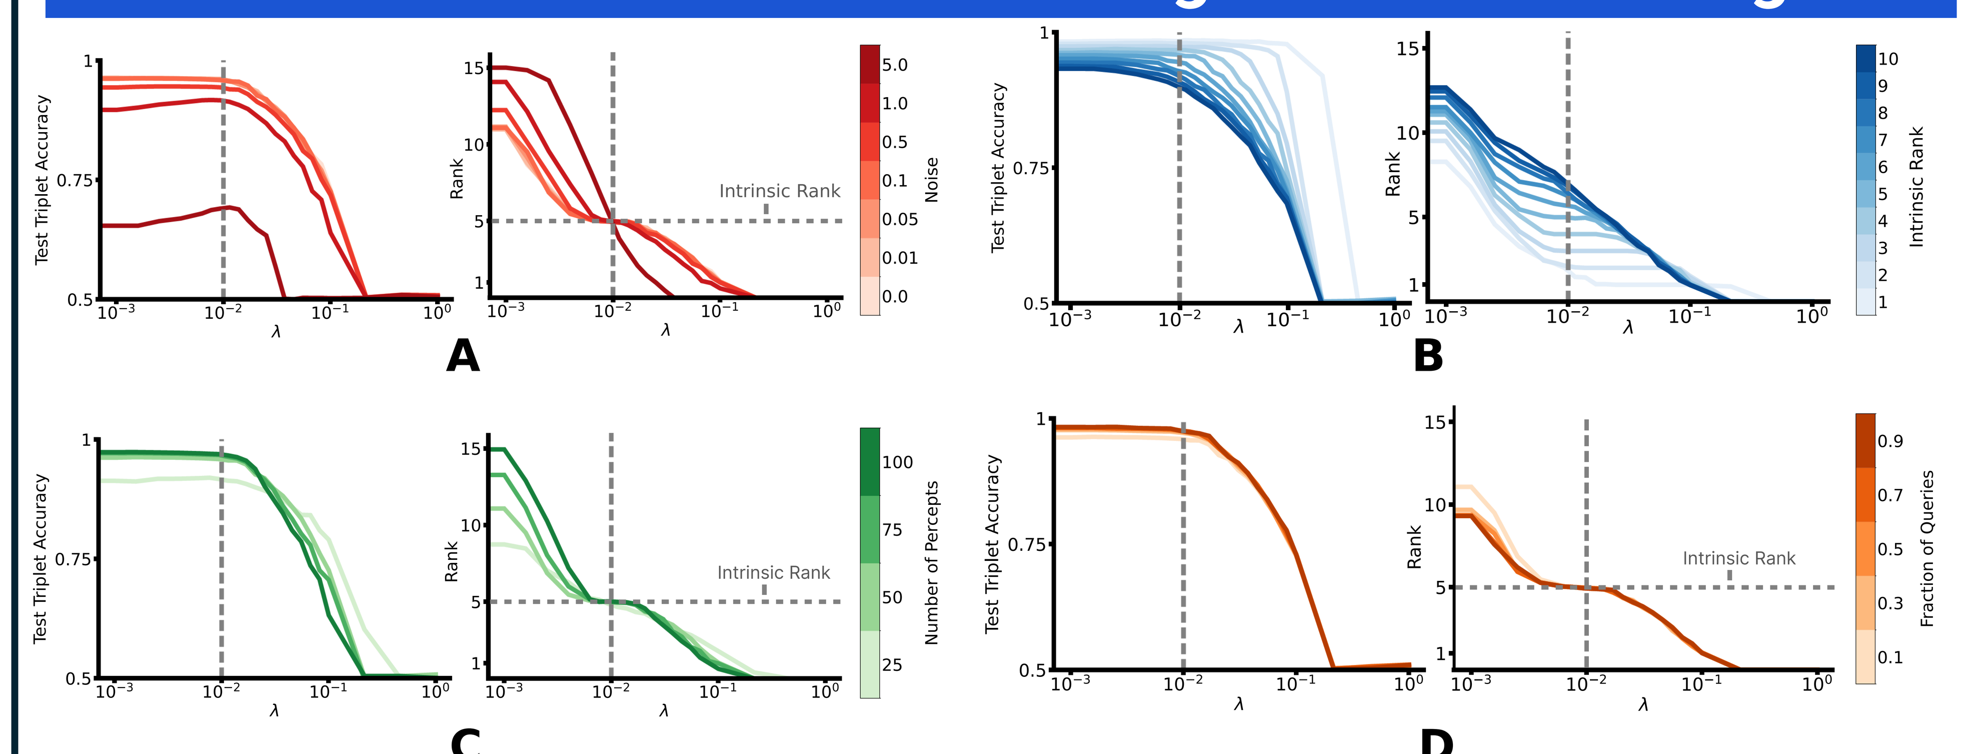

λ = 0.01 has stable intrinsic rank recovery as A) Response Noise, B) Intrinsic Rank, C) Number of Percepts and D) Fraction of Queries vary.

LORE outperforms on rank recovery; matches on accuracy on synthetic data

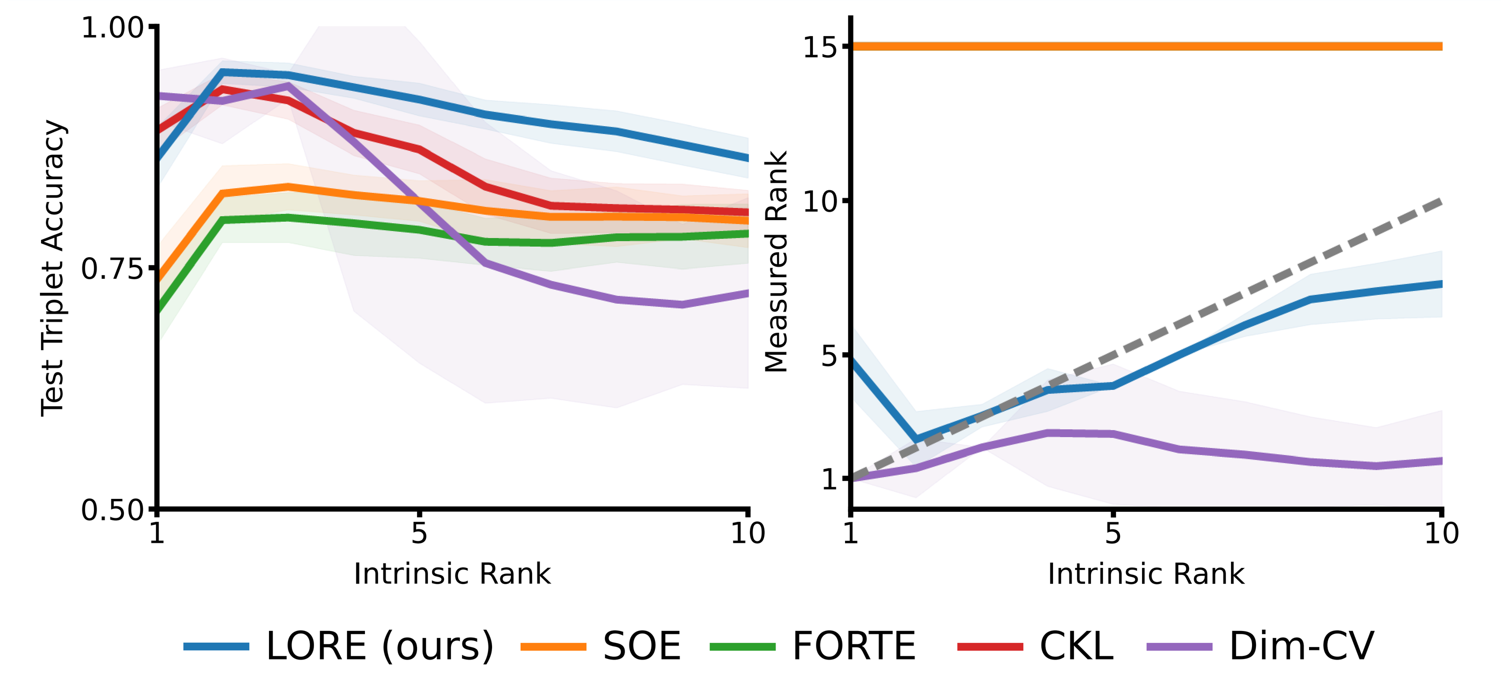

LORE Discovers Low Rank Structure

Method	Test Acc.	Rank	Time (s)
LORE	82.45 ± 0.27	$\mathbf{3.3 \pm 0.47}$	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

LORE learns low rank structure on the Food-100 dataset [3] without sacrificing test accuracy.

LORE's Axes are Semantically Interpretable

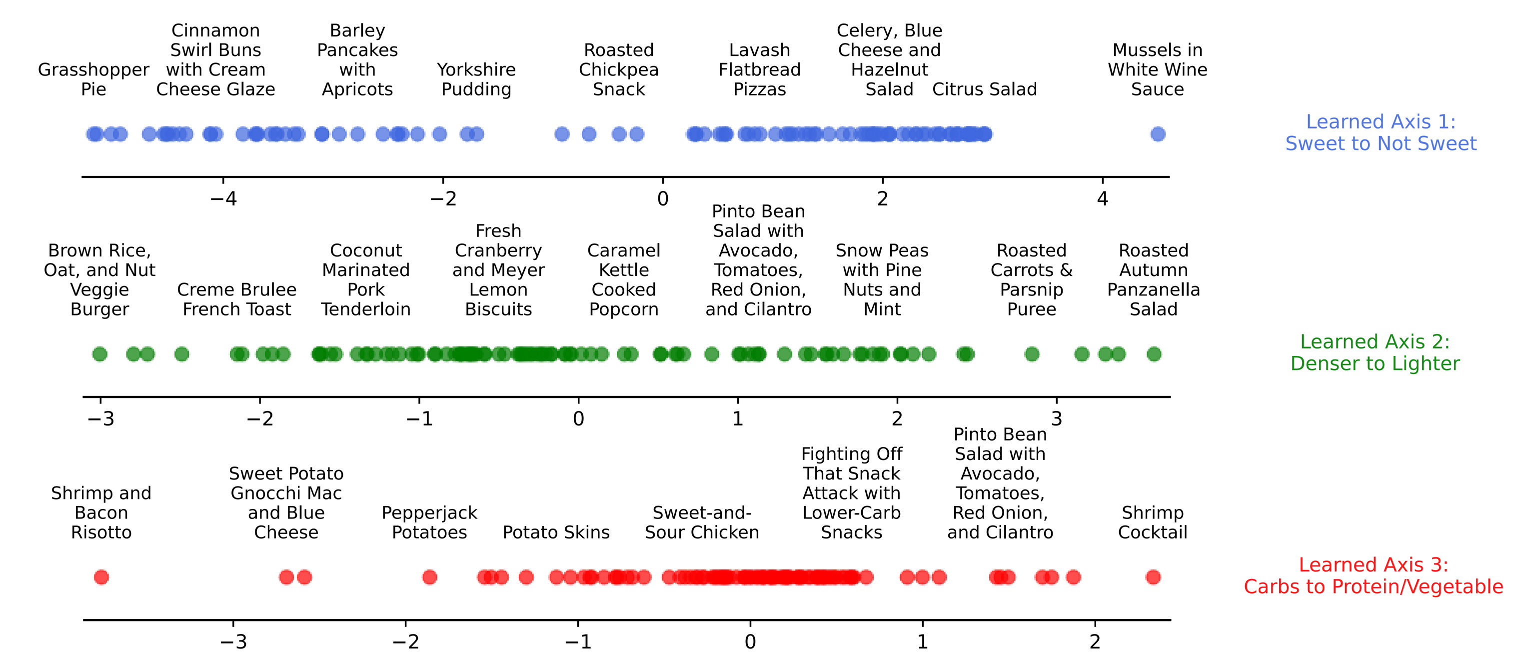

References

[1] N. Stewart, G. D. A. Brown, and N. Chater, "Absolute identification by relative judgment," *Psychological Review*, vol. 112, no. 4, p. 881, 2005.
[2] A. Bower, L. Jain, and L. Balzano, "The landscape of non-convex quadratic feasibility," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 3974–3978.
[3] M. Wilber, I. Kwak, and S. Belongie, "Cost-effective hits for relative similarity comparisons," in *Proc. AAAI Conf. Human Computation and Crowdsourcing*, vol. 2, 2014, pp. 227–233.

Acknowledgements

This work was supported by the National Science Foundation under Grant No. CCF-2107455. Additionally, supported in part by the NIH BRAIN Initiative through NIMH grant R61MH138966, the National Center for Advancing Translational Sciences of the National Institutes of Health under Award Number UL1TR002378 and KL2TR002381, and the Julian T. Hightower Chair at Georgia Tech. Vivek Anand was supported by the National Institutes of Health under the Data Science in Computational Neural Engineering Training Grant T32EB025816 and the National Defense Science and Engineering Graduate (NDSEG) Fellowship Program.