RankSim: Ranking Similarity Regularization for Deep Imbalanced Regression

Yu Gong, Greg Mori, Frederick Tung

BOREALIS AI

Deep Imbalanced Classification

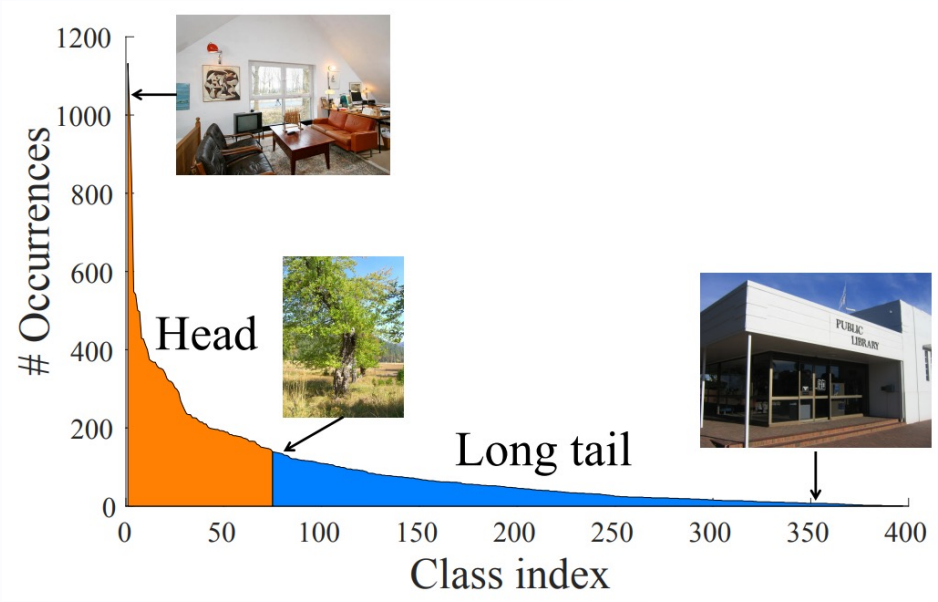

Figure source: Yu-Xiong Wang, et al. Learning to model the tail. *NeurIPS 2017*

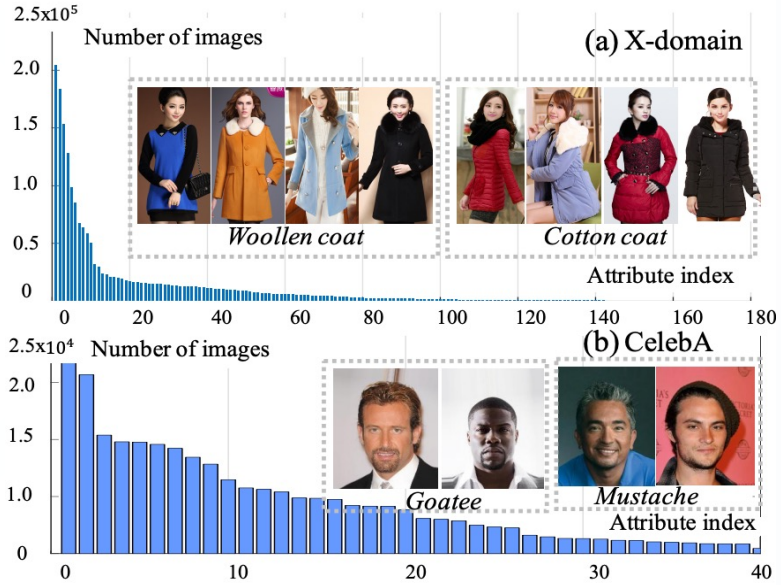

Figure source: Qi Dong, et al. Class Rectification Hard Mining for Imbalanced Deep Learning. *ICCV 2017*

- Most of the data belongs to a small number of classes (or long-tailed).

BOREALIS AI

Deep Imbalanced Regression

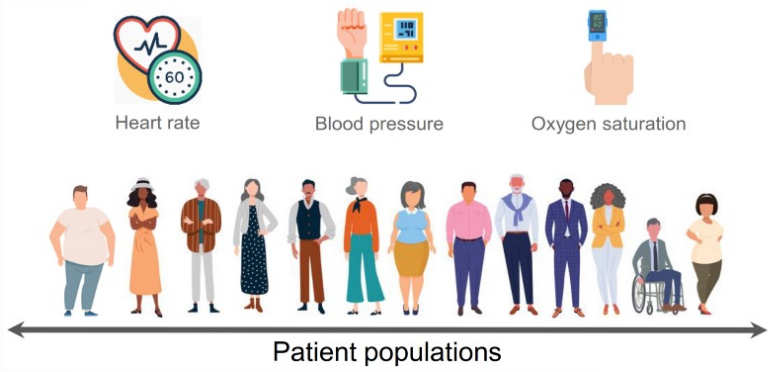

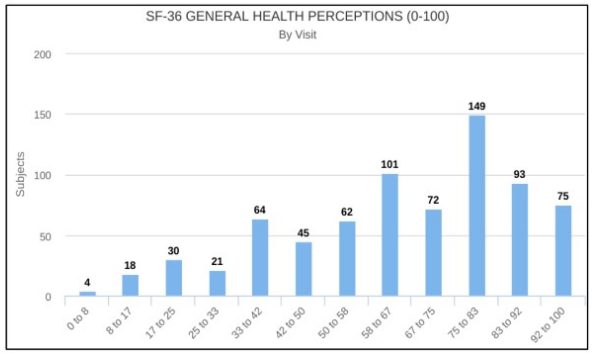

Health metrics

- Heart rate, blood pressure, etc.

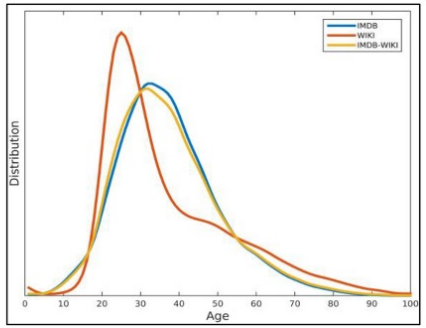

Age estimation

Figure source: Yuzhe Yang, et al. Delving into Deep Imbalanced Regression. *ICML 2021*

BOREALIS AI

Deep Imbalanced Regression

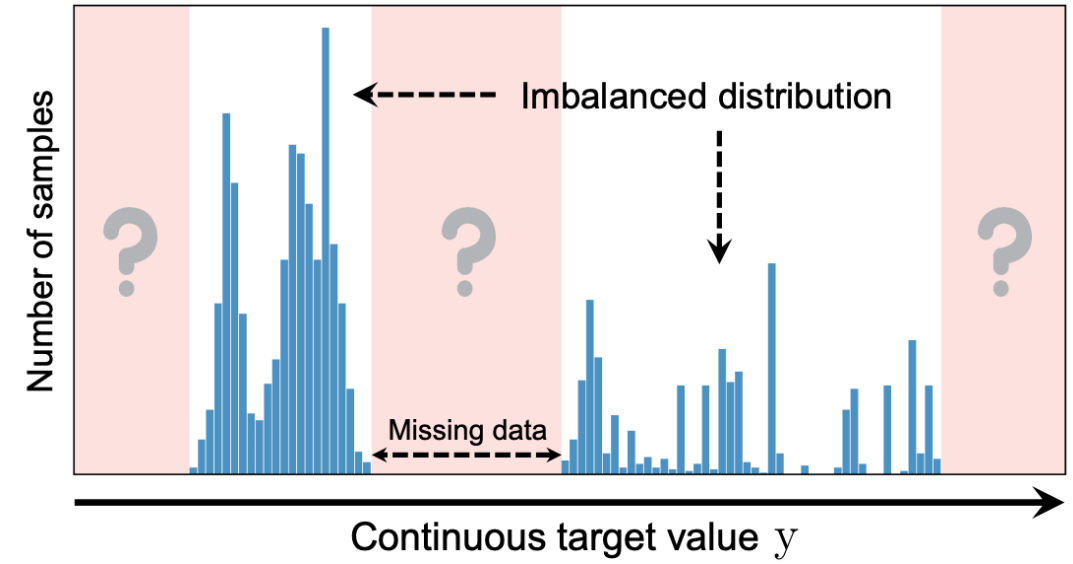

Figure source: Yuzhe Yang, et al. Delving into Deep Imbalanced Regression. *ICML 2021*

- Target y: *continuous* and potentially *infinite*
- Some targets may have *zero* samples
- Relationship between targets: *ordering*

BOREALIS AI

Deep Imbalanced Regression

- Training set: $\{(\mathbf{x}_i, y_i)\}_{i=1}^{N}$
- $\mathbf{x} \in \mathbb{R}^d$ is the *input*
- $y \in \mathbb{R}$ is the *continuous* target label
- $\mathbf{z} = f(\mathbf{x}; \theta)$ is the *feature* for $\mathbf{x}$, parameterized by deep neural networks

BOREALIS AI

Inductive Bias: Similarity Ordering

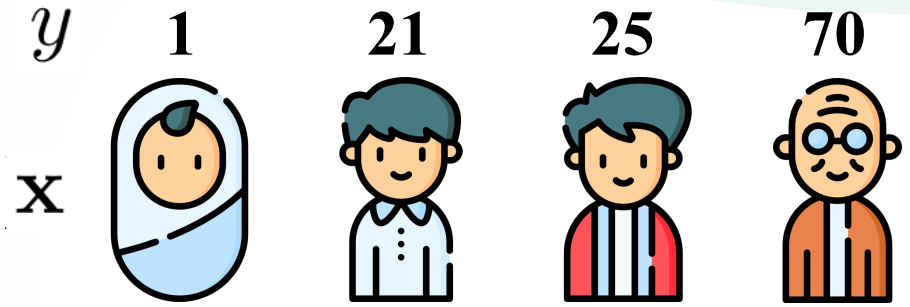

Similarity (,) > Similarity (,)

- **21**-year-old adult is more similar to **25**-year-old adult, than **70**-year-old man.

BOREALIS AI

Inductive Bias: Similarity Ordering

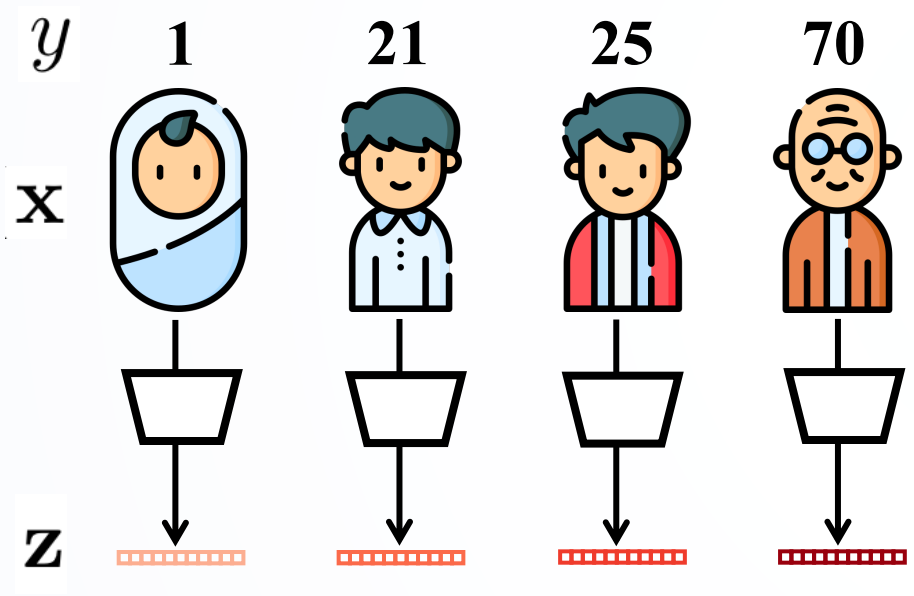

- $\sigma(*,*)$ denotes a similarity function
- $\sigma(\mathbf{z}_{21}, \mathbf{z}_{25}) > \sigma(\mathbf{z}_{21}, \mathbf{z}_{70})$

- Data that are closer in *label space* are also closer in *feature space*.

BOREALIS AI

Inductive Bias: Similarity Ordering

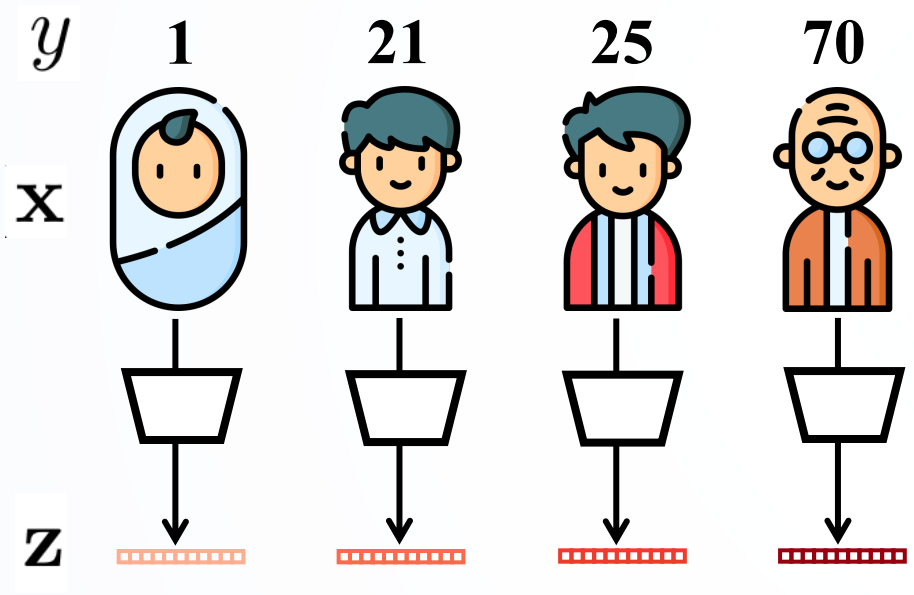

- $\sigma(\mathbf{z}_1, \mathbf{z}_{21}) > \sigma(\mathbf{z}_1, \mathbf{z}_{25}) > \sigma(\mathbf{z}_1, \mathbf{z}_{70})$
- $\sigma(\mathbf{z}_{21}, \mathbf{z}_{25}) > \sigma(\mathbf{z}_{21}, \mathbf{z}_1) > \sigma(\mathbf{z}_{21}, \mathbf{z}_{70})$
- $\sigma(\mathbf{z}_{25}, \mathbf{z}_{21}) > \sigma(\mathbf{z}_{25}, \mathbf{z}_1) > \sigma(\mathbf{z}_{25}, \mathbf{z}_{70})$
- $\sigma(\mathbf{z}_{70}, \mathbf{z}_{25}) > \sigma(\mathbf{z}_{70}, \mathbf{z}_{21}) > \sigma(\mathbf{z}_{70}, \mathbf{z}_1)$

BOREALIS AI

Method

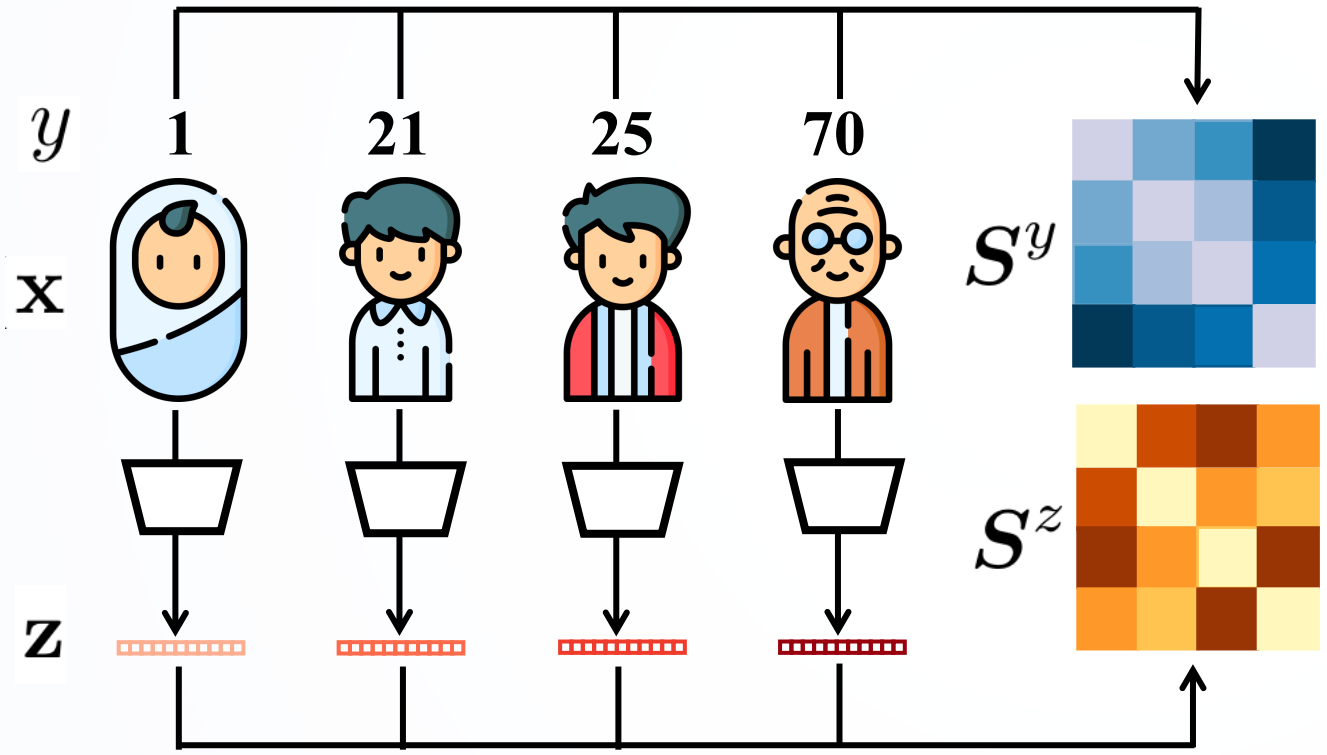

Pairwise Similarity Matrix

Label Similarity Matrix (*ground-truth*)

$$S^{y}_{i,j} = \sigma^{y}(y_i, y_j)$$

Feature Similarity Matrix (*learnable*)

$$S^{\mathbf{z}}_{i,j} = \sigma^{\mathbf{z}}(\mathbf{z}_i, \mathbf{z}_j)$$

BOREALIS AI

Method

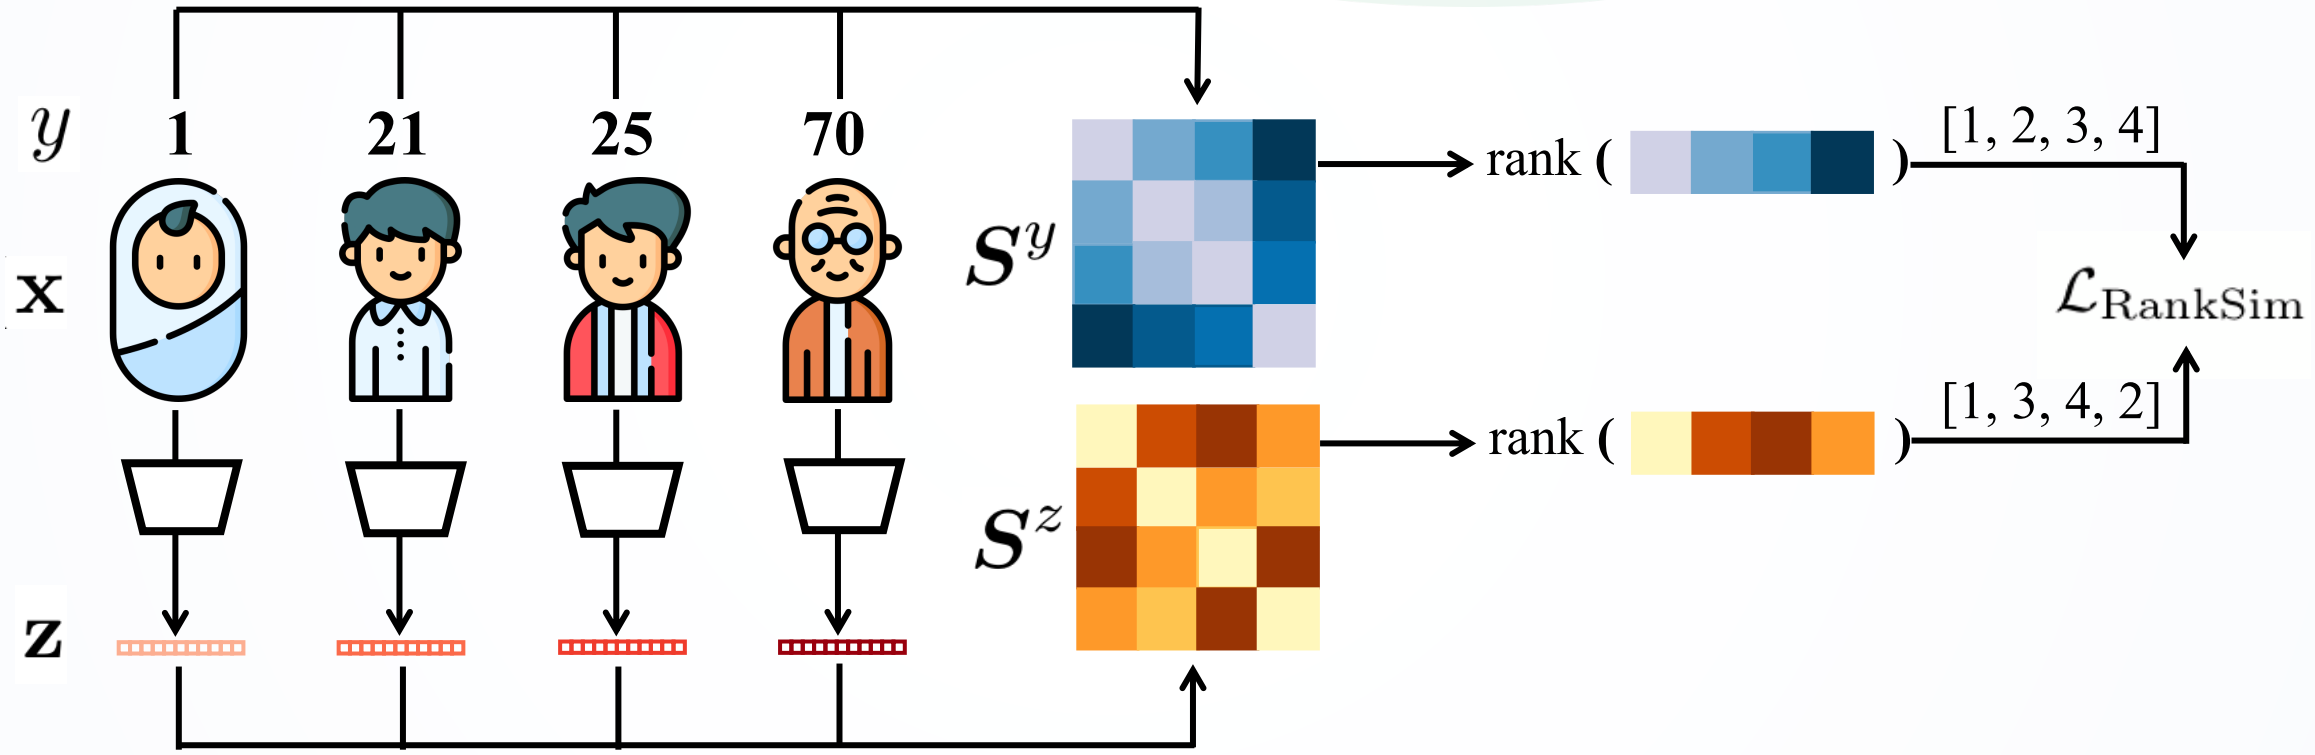

BOREALIS AI

Method

- $\mathcal{M} = \{(\mathbf{x}_i, y_i), i = 1, ..., M\}$
- Ranking function: $\mathbf{rk}(\mathbf{a})_i = 1 + |\{j : \mathbf{a}_j > \mathbf{a}_i\}|$.

$$\mathcal{L}_{\text{RankSim}} = \sum_{i=1}^{|\mathcal{M}|} \ell\left(\mathbf{rk}(\boldsymbol{S}^{y}_{[i,:]}),\ \mathbf{rk}(\boldsymbol{S}^{z}_{[i,:]})\right)$$

- For a given input sample, $\mathcal{L}_{\text{RankSim}}$ encourages the sorted list of its *neighbors in label space* to match the sorted list of its *neighbors in feature space*.

BOREALIS AI

Method

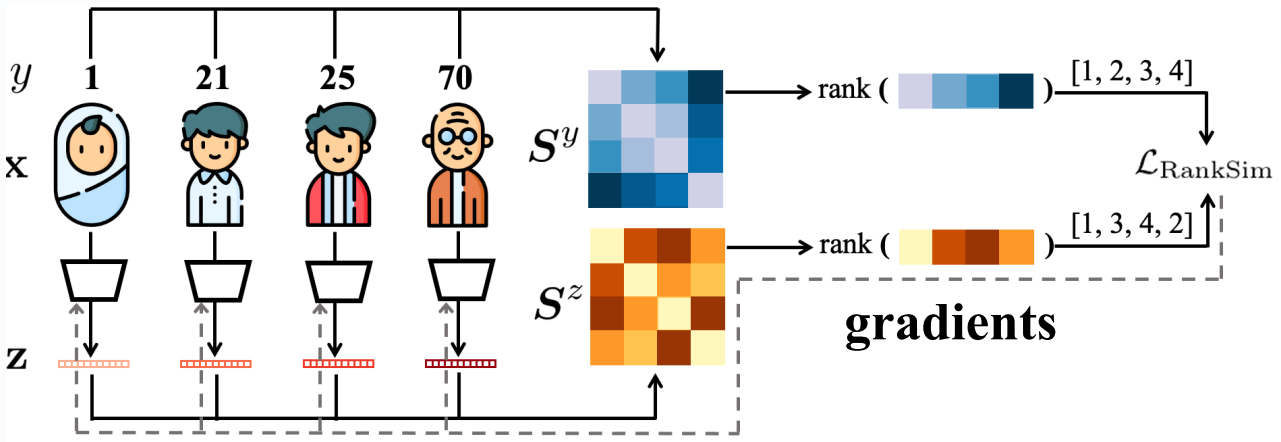

- $\mathcal{L}_{\text{objective}} = \mathcal{L}_{\text{regression}} + \gamma \mathcal{L}_{\text{RankSim}}$

- We recast the ranking as the minimizer of a linear combinatorial objective [1, 2].
- Two hyperparameters are introduced:
  - Interpolation strength of differential ranking λ
  - Balancing weight γ

[1] Rolínek, M., Musil, V., Paulus, A., Vlastelica, M., Michaelis, C., and Martius, G. Optimizing rank-based metrics with blackbox differentiation. CVPR 2020.
[2] Vlastelica, M., Paulus, A., Musil, V., Martius, G., and Rol´ınek, M. Differentiation of blackbox combinatorial solvers. ICLR 2020.

BOREALIS AI

Benchmarks

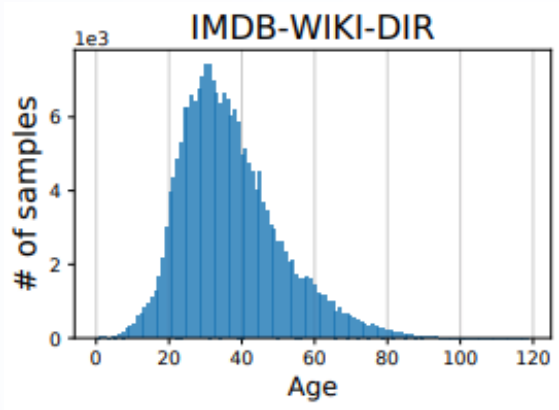

IMDB-WIKI-DIR

Age Estimation

191,509 imbalanced training samples

11022 balanced validation samples

11022 balanced test samples

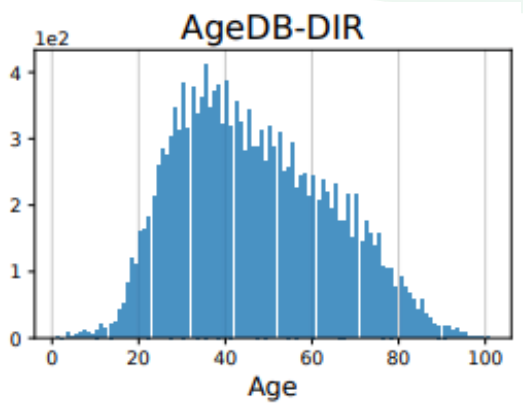

AgeDB-DIR

Age Estimation

12,208 imbalanced training samples

2140 balanced validation samples

2140 balanced test samples

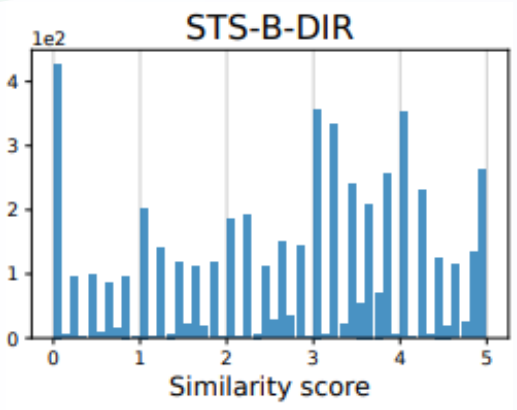

STS-B-DIR

Text Similarity Score Estimation

5249 imbalanced training samples

1000 balanced validation samples

1000 balanced test samples

Figure source: Yuzhe Yang, et al. Delving into Deep Imbalanced Regression. *ICML 2021*

BOREALIS AI

Baselines

- Focal-R [1]: regression variant of Focal loss
- RRT [2]: two-stage training
- SQINV / INV [3]: re-weighting
- LDS / FDS [3]: label and feature smoothing

[1] Lin, T.-Y., Goyal, P., Girshick, R., He, K., and Dollar, P. Focal loss for dense object detection. ICCV 2017.
[2] Yang, Y., Zha, K., Chen, Y.-C., Wang, H., and Katabi, D. Delving into deep imbalanced regression. ICML 2021.
[3] Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., and Kalantidis, Y. Decoupling representation and classifier for long-tailed recognition. ICLR 2020.

BOREALIS AI

Metrics

- Mean squared Error (MSE)
- Mean Absolute Error (MAE)
- Geometric Mean (GM): $(\prod_{i=1}^{n} e_i)^{\frac{1}{n}}$

- Many-shot: labels with > 100 training samples
- Medium-shot: labels with 20 to 100 training samples
- Few-shot: labels with < 20 training samples

BOREALIS AI

Results

IMDB-WIKI-DIR

- Age Estimation
- MSE, lower is better
- GM, lower is better

	MAE ↓				GM ↓			
	All	Man.	Med.	Few	All	Man.	Med.	Few
VANILLA	8.06	7.23	15.12	26.33	4.57	4.17	10.59	20.46
+ RANKSIM	7.72	**6.93**	14.48	25.38	4.27	**3.90**	10.02	17.84
+ LDS	7.83	7.31	12.43	22.51	4.42	4.19	7.00	13.94
+ LDS + RANKSIM	**7.57**	7.00	**12.16**	22.44	**4.23**	4.00	6.81	13.23
+ FDS	7.85	7.18	13.35	24.12	4.47	4.18	8.18	15.18
+ FDS + RANKSIM	7.74	**6.93**	14.71	24.91	4.34	3.96	10.35	16.85
+ LDS + FDS	7.78	7.20	12.61	22.19	4.37	4.12	7.39	12.61
+ LDS + FDS + RANKSIM	7.69	7.13	12.30	**21.43**	4.34	4.13	**6.72**	**12.48**
FOCAL-R	7.97	7.12	15.14	26.96	4.49	4.10	10.37	21.20
+ RANKSIM	7.77	6.99	14.23	26.01	4.38	4.03	9.25	20.16
+ LDS	7.90	7.10	14.72	25.84	4.47	4.09	10.11	19.14
+ LDS + RANKSIM	7.71	6.99	**13.65**	24.97	4.31	3.98	**8.72**	17.56
+ FDS	7.96	7.14	14.71	26.06	4.51	4.12	10.16	19.56
+ FDS + RANKSIM	7.75	7.01	14.06	**24.56**	4.33	3.99	9.04	**16.26**
+ LDS + FDS	7.88	7.10	14.08	25.75	4.47	4.11	9.32	18.67
+ LDS + FDS + RANKSIM	**7.67**	**6.91**	14.07	25.01	**4.28**	**3.93**	9.38	18.41
RRT	7.81	7.07	14.06	25.13	4.35	4.03	8.91	16.96
+ RANKSIM	7.55	6.83	13.47	24.72	4.17	3.86	8.66	15.54
+ LDS	7.79	7.08	13.76	24.64	4.34	4.02	8.72	16.92
+ LDS + RANKSIM	7.56	6.83	13.06	24.78	4.23	3.91	8.55	17.44
+ FDS	7.65	7.02	12.68	23.85	4.31	4.03	7.58	16.28
+ FDS + RANKSIM	**7.35**	6.81	**11.50**	**22.75**	**4.05**	3.85	**6.05**	**14.68**
+ LDS + FDS	7.65	7.06	12.41	23.51	4.31	4.07	7.17	15.44
+ LDS + FDS + RANKSIM	7.37	**6.80**	11.83	23.11	4.06	**3.84**	6.33	14.71
SQINV	7.87	7.24	12.44	22.76	4.47	4.22	7.25	15.10
+ RANKSIM	**7.42**	**6.84**	**12.12**	22.13	**4.10**	**3.87**	6.74	12.78
+ LDS	7.83	7.31	12.43	22.51	4.42	4.19	7.00	13.94
+ LDS + RANKSIM	7.57	7.00	12.16	22.44	4.23	4.00	6.81	13.23
+ FDS	7.83	7.23	12.60	22.37	4.42	4.20	6.93	13.48
+ FDS + RANKSIM	7.50	6.93	12.09	21.68	4.19	3.97	**6.65**	13.28
+ LDS + FDS	7.78	7.20	12.61	22.19	4.37	4.12	7.39	12.61
+ LDS + FDS + RANKSIM	7.69	7.13	12.30	**21.43**	4.34	4.13	6.72	**12.48**
OURS VS. VANILLA	+0.71	+0.43	+3.62	+4.90	+0.52	+0.33	+4.54	+7.98
OURS VS. Yang et al.	+0.30	+0.22	+0.91	+0.76	+0.26	+0.18	+0.88	+0.13

BOREALIS AI

Results

AgeDB-DIR

- Age Estimation
- MSE, lower is better
- GM, lower is better

	MAE ↓				GM ↓			
	All	Man.	Med.	Few	All	Man.	Med.	Few
VANILLA	7.77	6.62	9.55	13.67	5.05	4.23	7.01	10.75
+ RANKSIM	7.13	6.51	8.17	10.12	4.48	4.01	5.27	6.79
+ LDS	7.67	6.98	8.86	10.89	4.85	4.39	5.80	7.45
+ LDS + RANKSIM	**6.99**	**6.38**	7.88	10.23	**4.40**	**3.97**	5.30	6.93
+ FDS	7.55	6.50	8.97	13.01	4.75	4.03	6.42	9.93
+ FDS + RANKSIM	7.33	6.49	8.53	11.98	4.82	4.19	6.16	8.99
+ LDS + FDS	7.55	7.01	8.24	10.79	4.72	4.36	5.45	6.79
+ LDS + FDS + RANKSIM	7.03	6.54	**7.68**	**9.92**	4.45	4.07	**5.23**	**6.35**
FOCAL-R	7.64	6.68	9.22	13.00	4.90	4.26	6.39	9.52
+ RANKSIM	7.15	6.45	7.97	11.50	4.53	4.10	5.10	8.50
+ LDS	7.56	6.67	8.82	12.40	4.82	4.27	5.87	8.83
+ LDS + RANKSIM	7.25	6.40	8.71	11.24	4.58	4.02	5.99	**7.52**
+ FDS	7.65	6.89	8.70	11.92	4.83	4.32	5.89	8.04
+ FDS + RANKSIM	7.25	6.72	**7.86**	**10.58**	4.54	4.22	**4.84**	7.57
+ LDS + FDS	7.47	6.69	8.30	12.55	4.71	4.25	5.36	8.59
+ LDS + FDS +RANKSIM	**7.09**	**6.17**	8.71	11.68	**4.46**	**3.85**	5.76	8.78
RRT	7.74	6.98	8.79	11.99	5.00	4.50	5.88	8.63
+ RANKSIM	7.11	6.53	8.00	10.04	4.52	4.19	**5.05**	**6.77**
+ LDS	7.72	7.00	8.75	11.62	4.98	4.54	5.71	8.27
+ LDS + RANKSIM	**6.94**	**6.43**	**7.54**	10.10	**4.37**	**3.97**	5.11	7.05
+ FDS	7.70	6.95	8.76	11.86	4.82	4.32	5.83	8.08
+ FDS + RANKSIM	7.11	6.55	7.99	**10.02**	4.49	4.13	5.13	6.85
+ LDS + FDS	7.66	6.99	8.60	11.32	4.80	4.42	5.53	6.99
+ LDS + FDS + RANKSIM	7.13	6.54	8.07	10.12	4.55	4.18	5.20	6.87
SQINV	7.81	7.16	8.80	11.20	4.99	4.57	5.73	7.77
+ RANKSIM	**6.91**	**6.34**	**7.79**	9.89	**4.28**	**3.92**	**4.88**	6.89
+ LDS	7.67	6.98	8.86	10.89	4.85	4.39	5.80	7.45
+ LDS + RANKSIM	6.99	6.38	7.88	10.23	4.40	3.97	5.30	6.90
+ FDS	7.69	7.10	8.86	9.98	4.83	4.41	5.97	**6.29**
+ FDS + RANKSIM	7.02	6.49	7.84	**9.68**	4.53	4.13	5.37	6.89
+ LDS + FDS	7.55	7.01	8.24	10.79	4.72	4.36	5.45	6.79
+ LDS + FDS + RANKSIM	7.03	6.54	**7.68**	9.92	4.45	4.07	5.23	6.35
OURS VS. VANILLA	+0.86	+0.45	+2.01	+3.99	+0.77	+0.38	+2.17	+4.40
OURS VS. Yang et al.	+0.56	+0.33	+0.70	+0.30	+0.43	+0.18	+0.52	-0.06

BOREALIS AI

Results

STS-B-DIR

- Text Similarity Score Estimation
- MSE, lower is better
- Pearson correlation, higher is better

	MSE ↓				Pearson cor. (%) ↑			
	All	Man.	Med.	Few	All	Man.	Med.	Few
VANILLA	0.974	0.851	1.520	0.984	74.2	72.0	62.7	75.2
+ RANKSIM	**0.873**	0.908	**0.767**	0.705	**76.8**	71.0	**72.9**	85.2
+ LDS	0.914	0.819	1.319	0.955	75.6	73.4	63.8	76.2
+ LDS + RANKSIM	0.889	0.911	0.849	0.690	76.2	70.7	70.0	85.6
+ FDS	0.916	0.875	1.027	1.086	75.5	73.0	67.0	72.8
+ FDS + RANKSIM	0.884	0.924	**0.767**	**0.685**	76.5	70.4	72.5	**85.7**
+ LDS + FDS	0.907	**0.802**	1.363	0.942	76.0	**74.0**	65.2	76.6
+ LDS + FDS + RANKSIM	0.903	0.908	0.911	0.804	75.8	70.6	69.0	82.7
FOCAL-R	0.951	0.843	1.425	0.957	74.6	72.3	61.8	76.4
+ RANKSIM	0.887	0.889	0.918	0.745	76.2	70.8	70.4	84.6
+ LDS	0.930	**0.807**	1.449	0.993	75.7	**73.9**	62.4	75.4
+ LDS + RANKSIM	**0.872**	0.887	0.847	**0.718**	**76.7**	71.2	70.3	**85.1**
+ FDS	0.920	0.855	1.169	1.008	75.1	72.6	66.4	74.7
+ FDS + RANKSIM	0.913	0.952	0.793	0.723	75.6	69.6	**71.5**	84.8
+ LDS + FDS	0.940	0.849	1.358	0.916	74.9	72.2	66.3	77.3
+ LDS + FDS + RANKSIM	0.911	0.943	**0.779**	0.866	75.7	69.9	71.4	81.2
RRT	0.964	0.842	1.503	0.978	74.5	72.4	62.3	75.4
+ RANKSIM	**0.865**	0.876	0.867	**0.670**	**77.1**	72.2	68.3	**86.1**
+ LDS	0.916	0.817	1.344	0.945	75.7	73.5	64.1	76.6
+ LDS + RANKSIM	0.874	0.893	**0.833**	0.722	77.0	72.3	68.3	84.8
+ FDS	0.929	0.857	1.209	1.025	74.9	72.1	67.2	74.0
+ FDS + RANKSIM	0.871	0.874	0.898	0.734	76.8	72.0	**68.7**	84.5
+ LDS + FDS	0.903	**0.806**	1.323	0.936	76.0	**73.8**	65.2	76.7
+ LDS + FDS + RANKSIM	0.882	0.892	0.887	0.702	76.6	71.7	68.0	85.5
INV	1.005	0.894	1.482	1.046	72.8	70.3	62.5	73.2
+ RANKSIM	1.091	1.056	1.240	1.118	69.9	65.2	60.1	76.0
+ LDS	0.914	0.819	1.319	0.955	75.6	73.4	63.8	76.2
+ LDS + RANKSIM	**0.889**	0.911	**0.849**	**0.690**	**76.2**	70.7	**70.0**	**85.6**
+ FDS	0.927	0.851	1.225	1.012	75.0	72.4	66.6	74.2
+ FDS + RANKSIM	1.083	1.035	1.301	1.063	70.0	64.8	68.9	76.7
+ LDS + FDS	0.907	**0.802**	1.363	0.942	76.0	**74.0**	65.2	76.6
+ LDS + FDS + RANKSIM	0.903	0.908	0.911	0.804	75.8	70.6	69.0	82.7
OURS VS. VANILLA	+0.109	-0.023	+0.753	+0.314	+2.9	+0.3	+10.2	+10.9
OURS VS. Yang et al.	+0.038	-0.072	+0.260	+0.246	+1.1	-1.7	+5.7	+8.8

BOREALIS AI

Feature Visualization

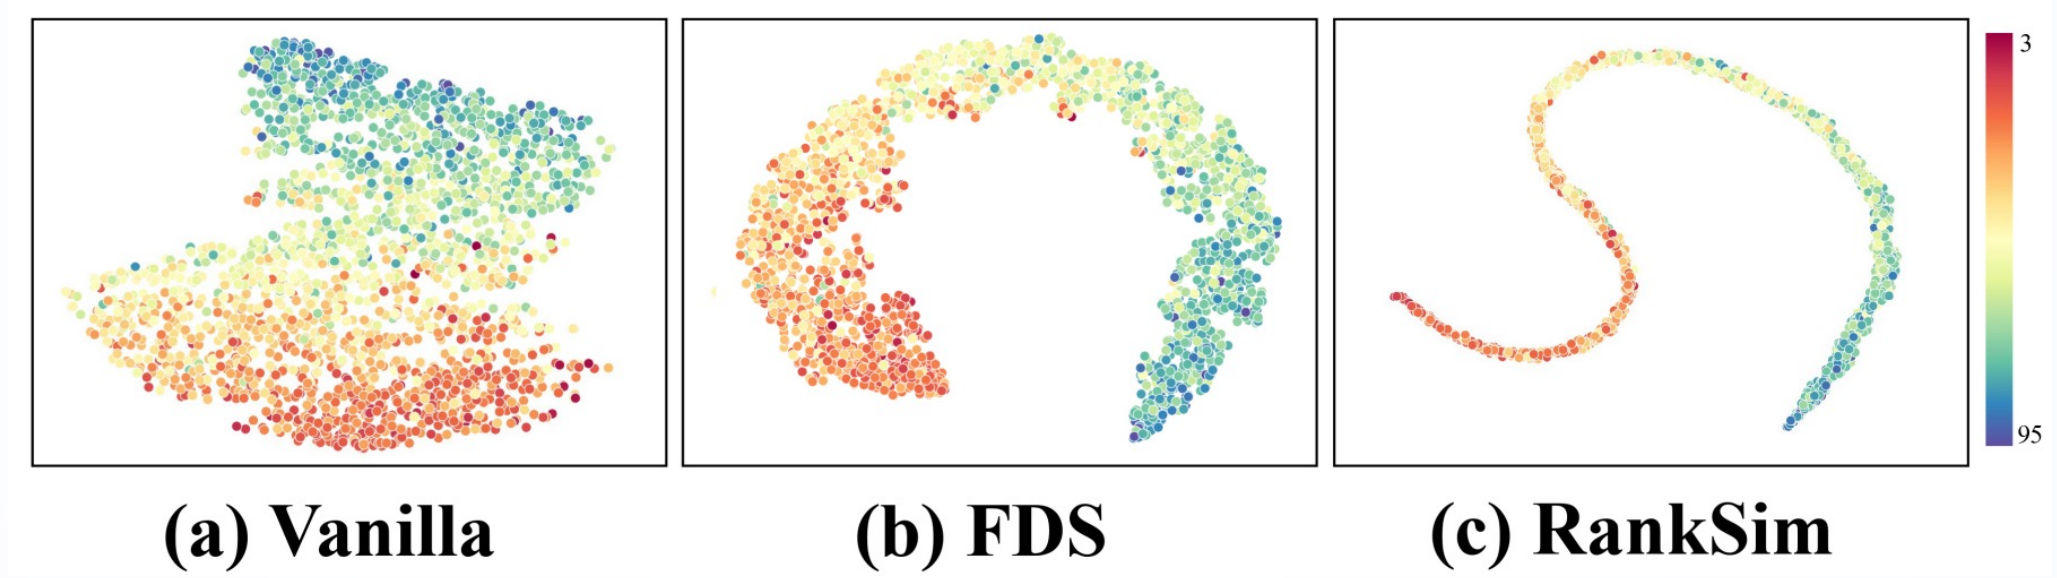

BOREALIS AI

Average feature-space ranking matrices

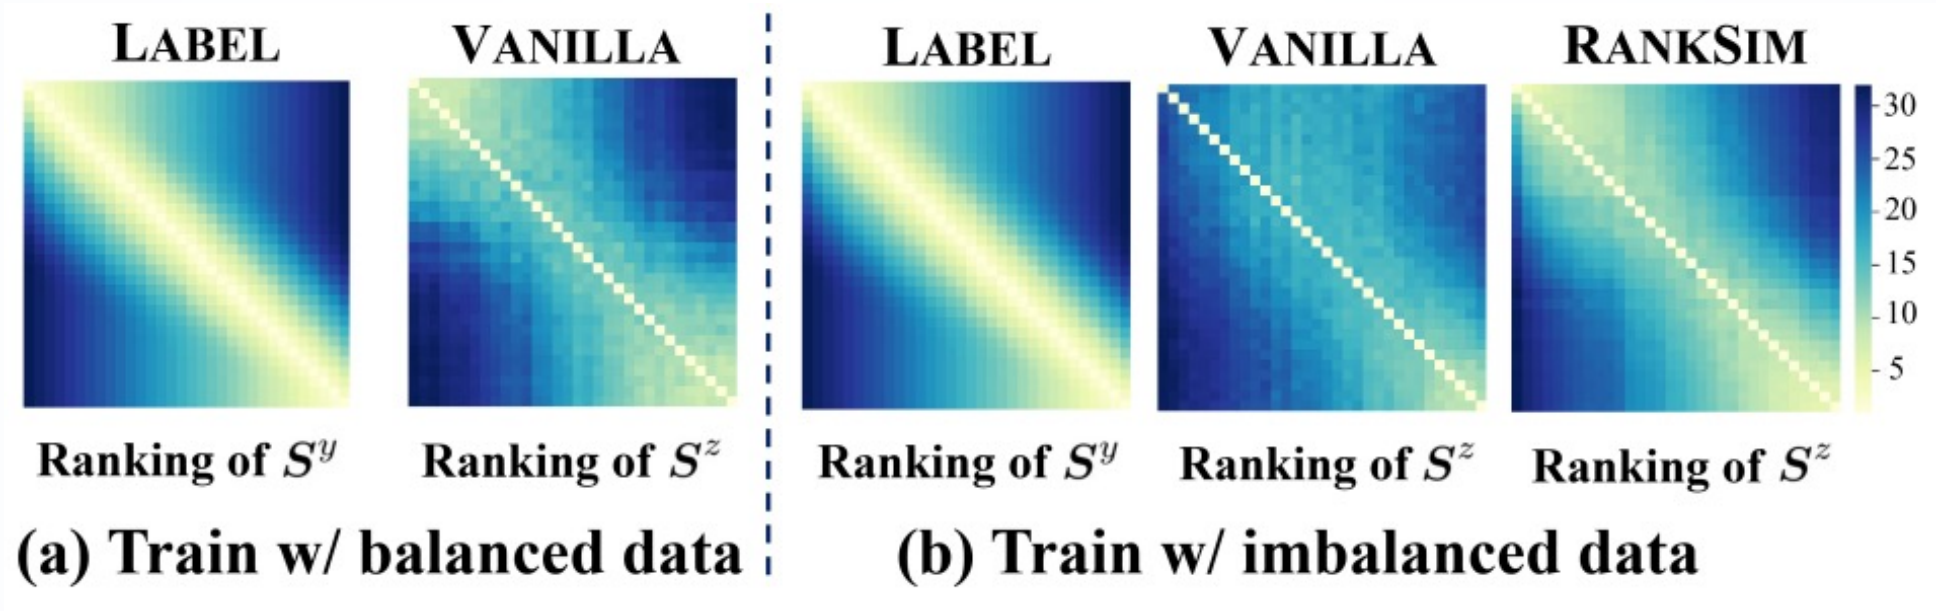

- Label-space and feature-space rankings tend to be consistent when training with balanced data

BOREALIS AI

Zero-shot Regression

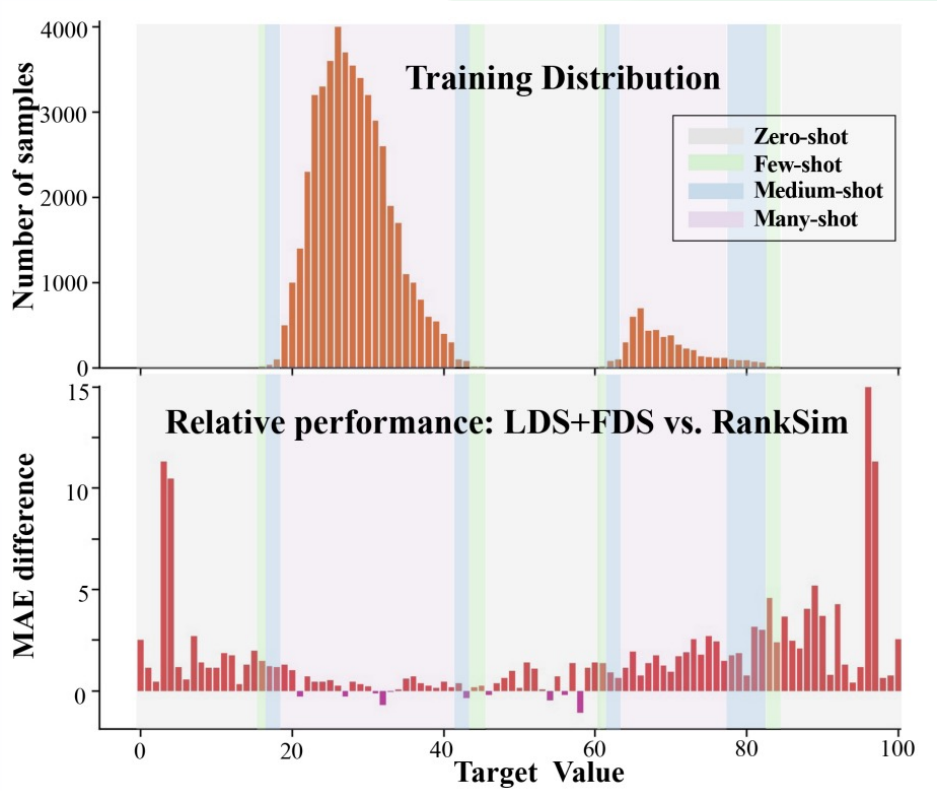

- **RankSim** significantly outperforms state-of-the-art in the zero-shot region.

BOREALIS AI

Conclusion

- A novel regularizer **RankSim** (data that are closer in label space also closer in the feature space) capturing both nearby & distant relationships.
- New state-of-the-art results on three benchmarks.
- Complementary to conventional imbalanced learning methods.

BOREALIS AI

Thank you

Code and pretrained weights

https://arxiv.org/pdf/2205.15236.pdf