SimPer: Simple Self-Supervised Learning of Periodic Targets

Yuzhe Yang[1], Xin Liu[2], Jiang Wu[3], Silviu Borac[3], Dina Katabi[1], Ming-Zher Poh[3], Daniel McDuff[2,3]

[1]MIT CSAIL [2]University of Washington [3]Google

Background & Motivation

Periodicity is ubiquitous in nature!

1. *Healthcare:* Video-based vital signs measurement
2. *Environmental sensing:* Precipitation patterns / land surface temperature now- & future-casting
3. *Behavior analysis:* Temporal morphology in human motions for rehabilitation & neurological diseases

However... ***Labeling*** such data is typically challenging and resource intensive.

Idea: *Self-supervised learning*!

Pitfalls of SOTA SSL on Periodic Tasks

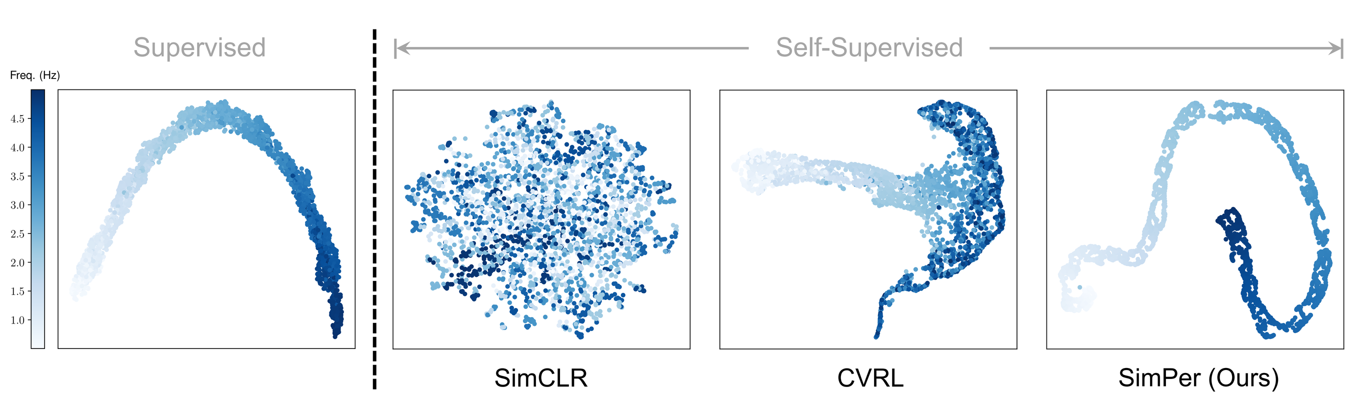

Observations:

- Existing SSL schemes fail to capture the underlying ***periodic*** or ***frequency*** information in data
- *Low frequency **resolution*** across learned features

Challenges:

1. How to define ***positive*** and ***negative*** samples in periodic learning?
2. How to measure ***feature similarity*** for periodic representations?
3. How to enable contrastive self-supervised learning over ***continuous*** targets?

SimPer: Simple SSL of Periodic Targets

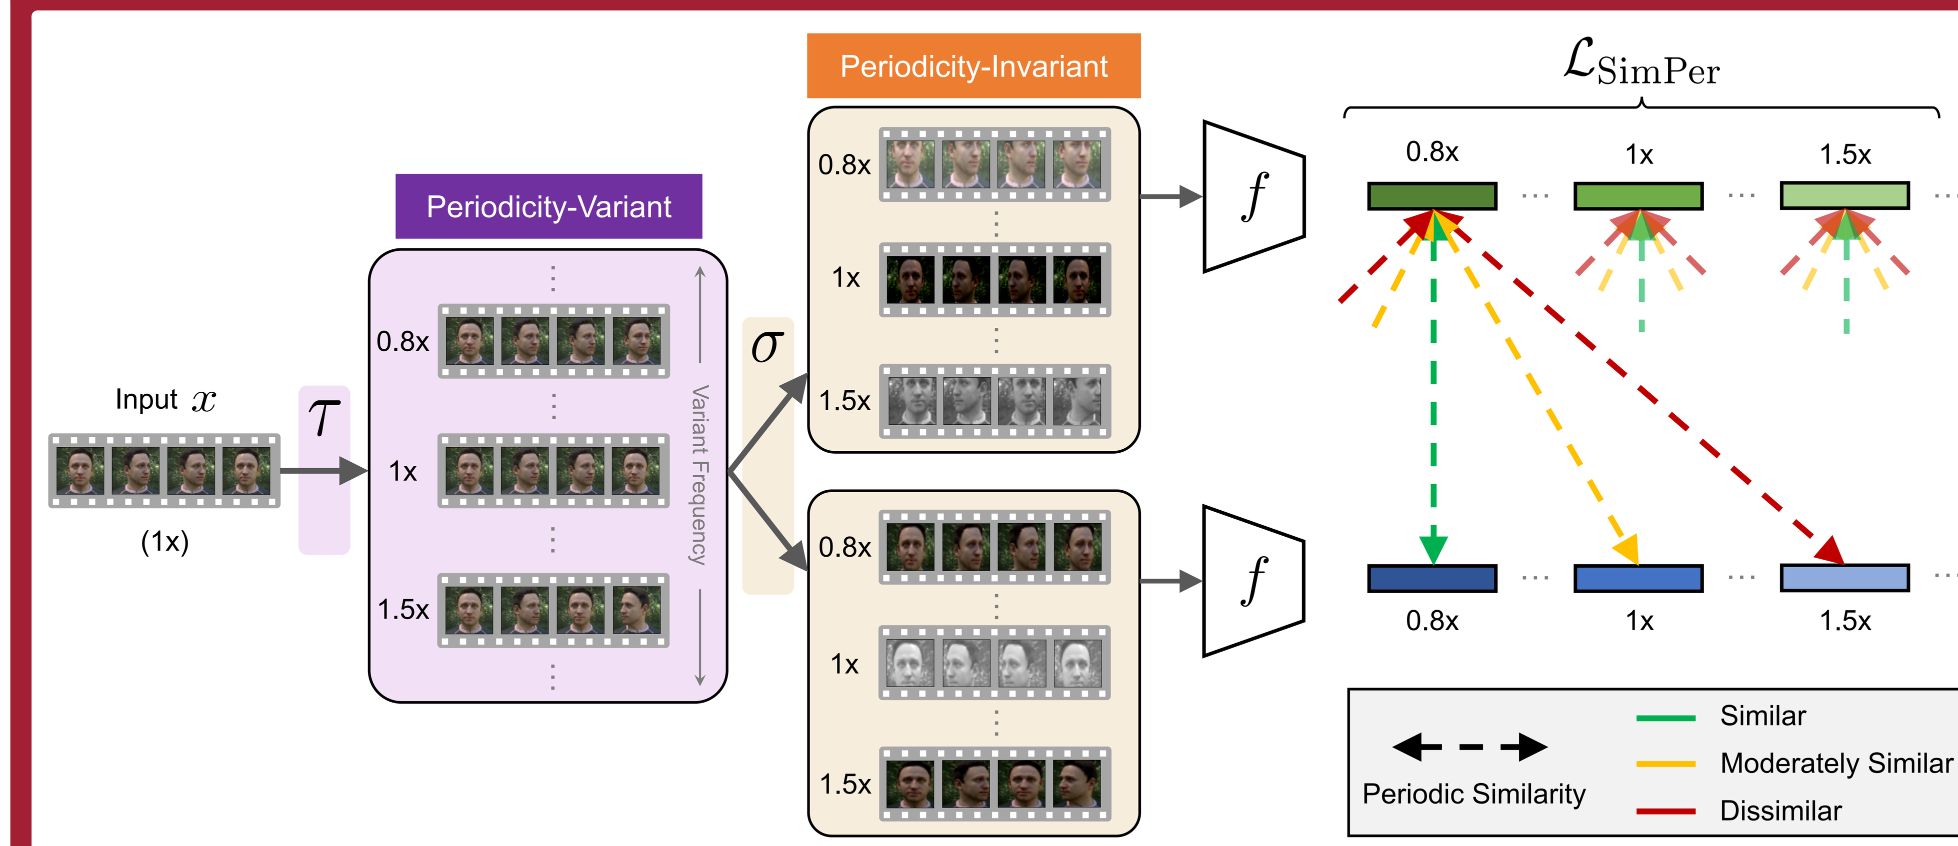

- ***Core Idea:*** Constructing & contrasting samples with different ***frequencies***
- ***Periodicity-Variant & Invariant Aug.:*** *Negative* views via ***frequency*** (speed) transformations from the ***same*** instance; *Positive* views via transforms that do not change the speed
- ***Temporal Self-Contrastive Learning:*** *Arbitrarily* large negative sample sizes under Nyquist sampling theorem; Natural *hard* negative samples & robust to spurious correlations

Periodic Feature Similarity

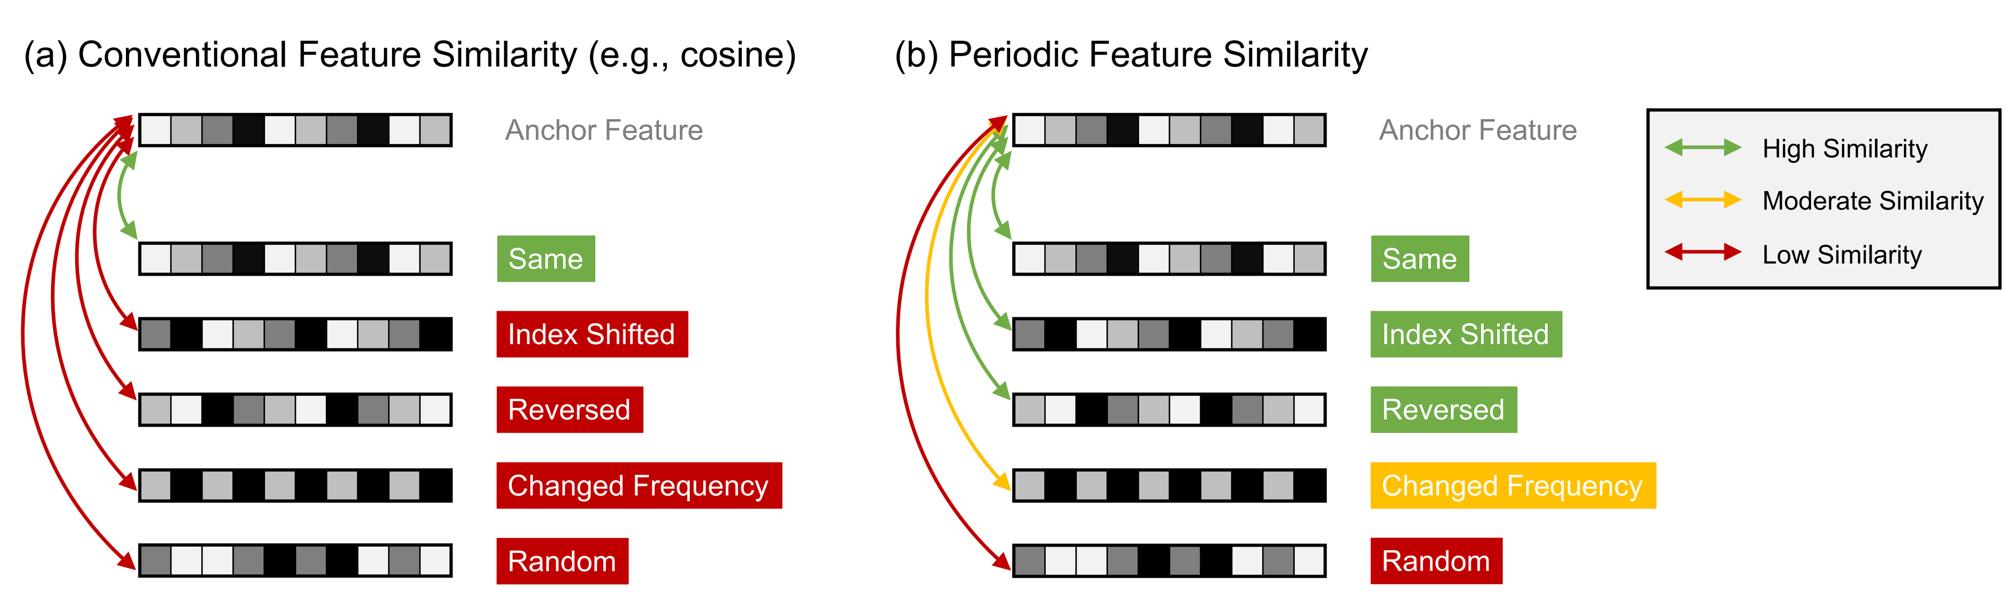

Concrete Instantiations:

Maximum cross-correlation / Normalized power spectrum density

Generalized Contrastive Loss for Continuous Targets

Classic InfoNCE Loss: hard classification task

$$\mathcal{L}_{\text{InfoNCE}} = -\log \frac{\exp(\text{sim}(\mathbf{z}, \tilde{\mathbf{z}})/\nu)}{\sum_{\mathbf{z}' \in \mathcal{Z}\setminus\{\mathbf{z}\}} \exp(\text{sim}(\mathbf{z}, \mathbf{z}')/\nu)}$$

- ***Cons:*** negative pairs possess meaningful **distance** in their *relative frequencies*

From Discrete Instance Discrimination to Continuous Contrast: soft regression

$$\mathcal{L}_{\text{SimPer}} = \sum_i -\sum_{j=1}^{M} \frac{\exp(w_{i,j})}{\sum_{k=1}^{M}\exp(w_{i,k})} \log \frac{\exp(\text{sim}(\mathbf{z}_i, \mathbf{z}'_j)/\nu)}{\sum_{k=1}^{M}\exp(\text{sim}(\mathbf{z}_i, \mathbf{z}'_k)/\nu)}, \quad w_{i,j} := \text{sim}_{\text{label}}(s_i, s_j)$$

Experiments & Results

4 different domains & tasks / 6 benchmark datasets

- ***Synthetic Data:*** RotatingDigits
- ***Action Counting:*** Countix
- ***Satellite Sensing:*** Land Surface Temperature (LST)
- ***Human Physiology:*** SCAMPS / UBFC / PURE

	RotatingDigits		SCAMPS		UBFC		PURE		Countix		LST	
Metrics	MAE↓	MAPE↓	MAE↓	MAPE↓	MAE↓	MAPE↓	MAE↓	MAPE↓	MAE↓	GM↓	MAE↓	ρ↑
SUPERVISED	0.72	28.96	3.61	5.33	5.13	4.72	4.25	4.93	1.50	0.73	1.54	**0.96**
SIMCLR	0.69	26.54	4.96	6.92	5.32	4.96	4.86	5.32	1.58	0.80	1.54	0.95
MOCO V2	0.64	24.73	5.33	7.24	5.05	4.64	4.97	5.60	1.54	0.79	1.53	0.95
BYOL	0.39	20.91	3.49	5.27	5.51	5.07	4.28	4.97	1.47	0.71	1.62	0.92
CVRL	0.34	18.82	5.52	7.34	5.07	4.70	4.19	4.71	1.48	0.71	1.49	**0.96**
SIMPER	**0.20**	**14.33**	**3.27**	**4.89**	**4.24**	**3.97**	**3.89**	**4.01**	**1.33**	**0.59**	**1.47**	**0.96**
GAINS	+0.52	+14.63	+0.34	+0.44	+0.89	+0.75	+0.36	+0.92	+0.17	+0.14	+0.07	+0.00

Analysis: Data Efficiency

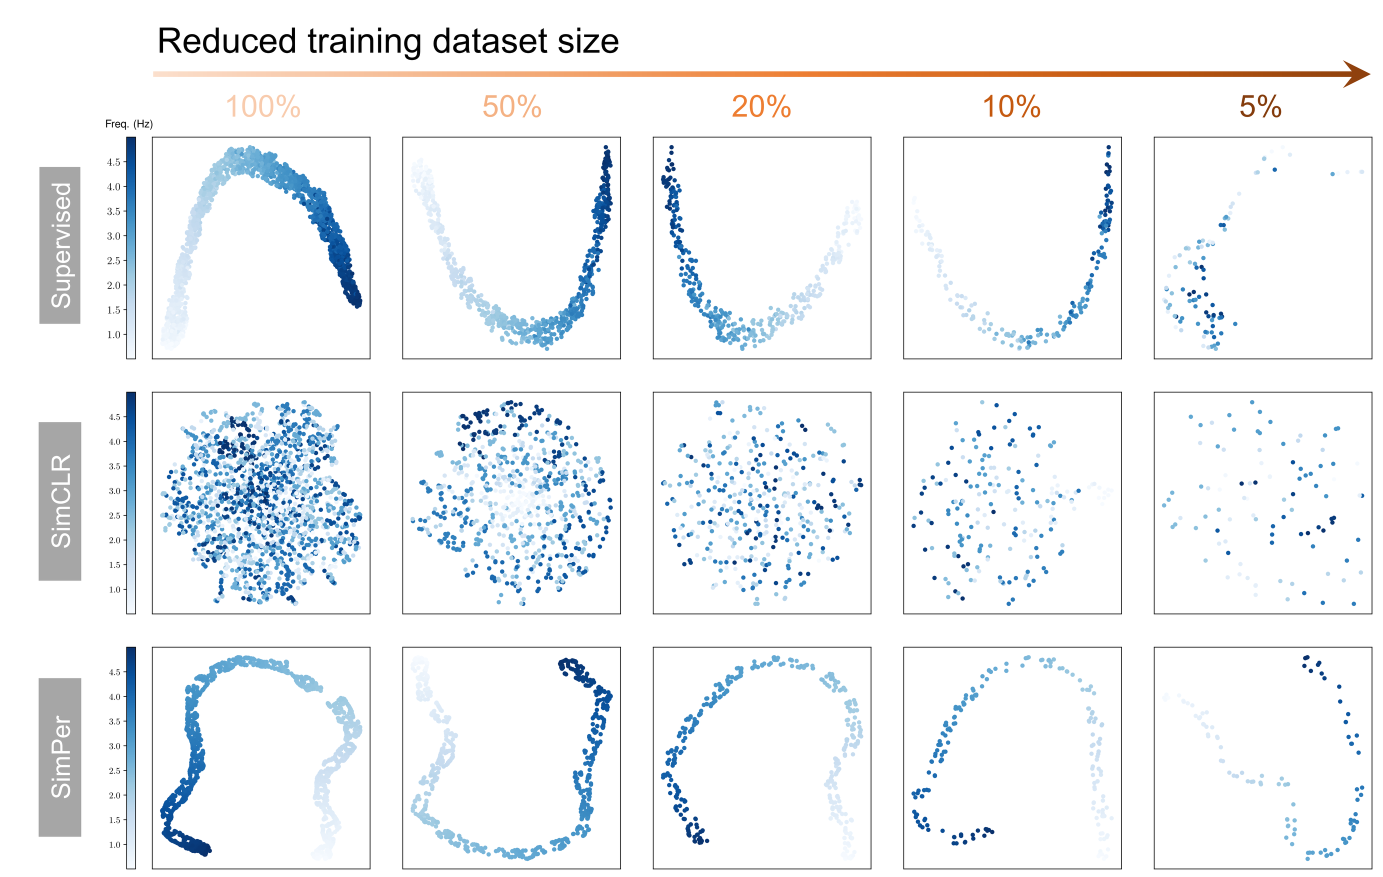

Analysis: Zero-Shot Generalization

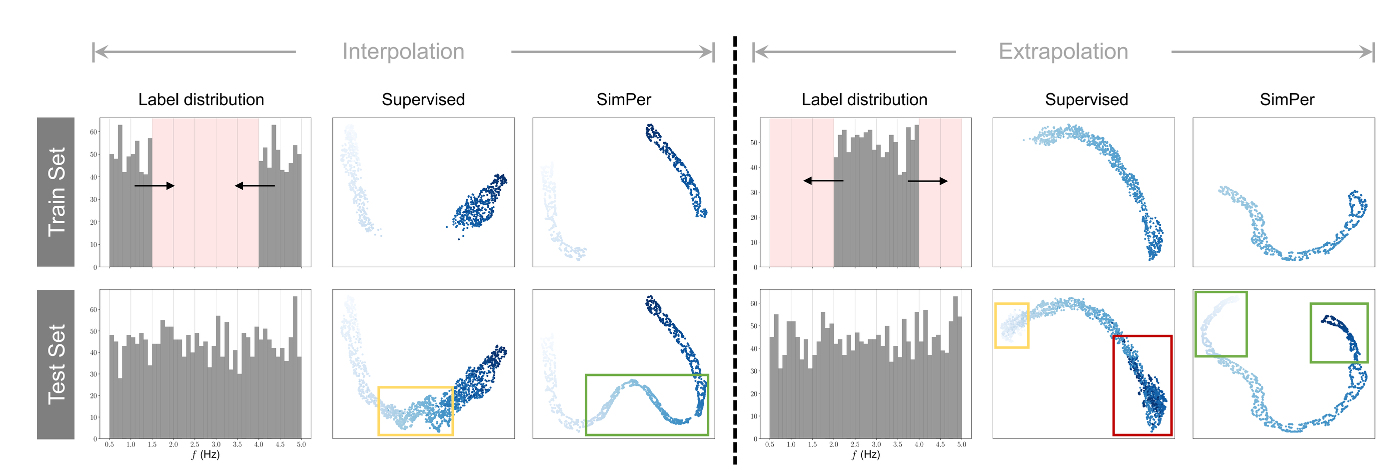

More Information

Code: `https://github.com/YyzHarry/SimPer`

Talk: `https://youtu.be/uEezGU3P_-I`

Project page: `https://simper.csail.mit.edu/`