MAST: Masked Augmentation Subspace Training for Generalizable Self-Supervised Priors

Chen Huang, Hanlin Goh, Jiatao Gu, Josh Susskind

ICLR 2023 | Apple Inc. | May 2023

Motivation

Limitations of existing SSL methods

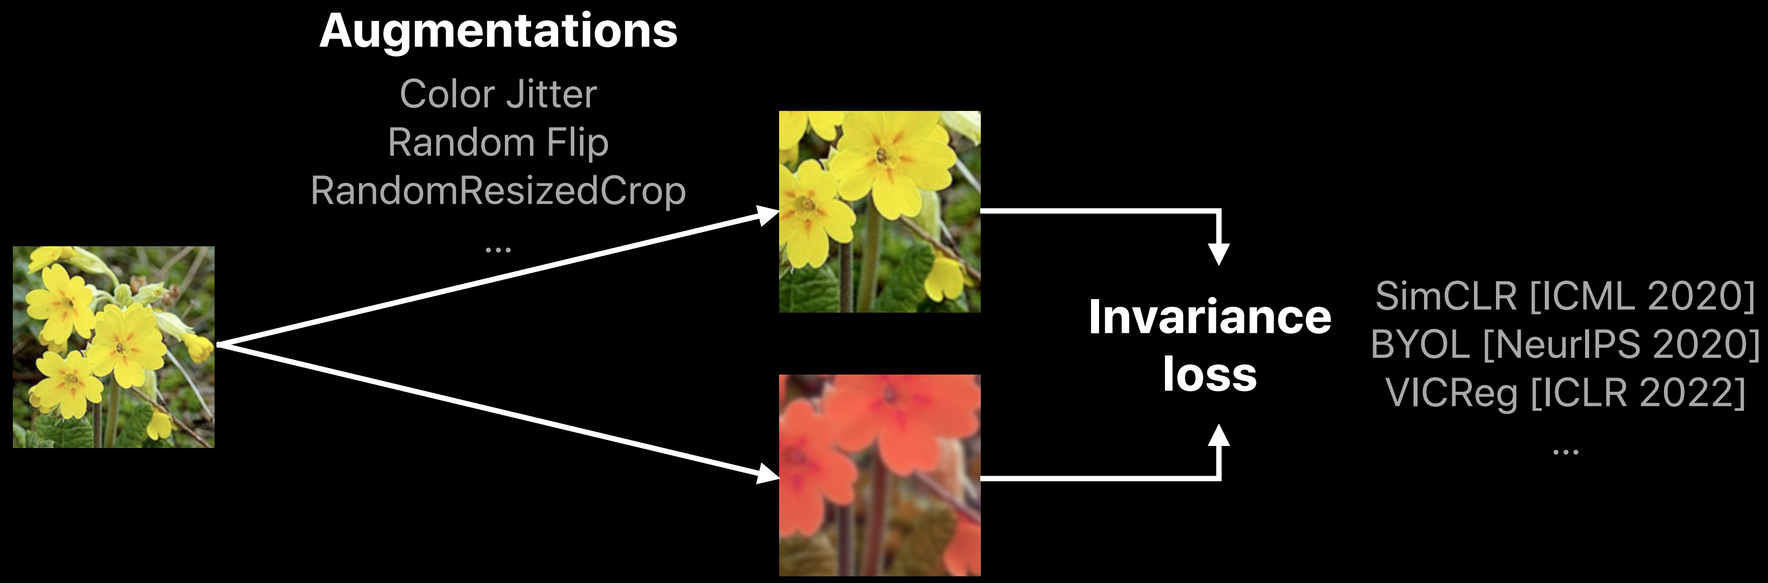

Which augmentations help generalization?

- Depends on downstream task
- Learned invariances may contradict each other

Similarity mismatch

- Between ambiguous views with strong augmentations

MAST: Masked Augmentation Subspace Training

Disentangled learning of augmentation invariances

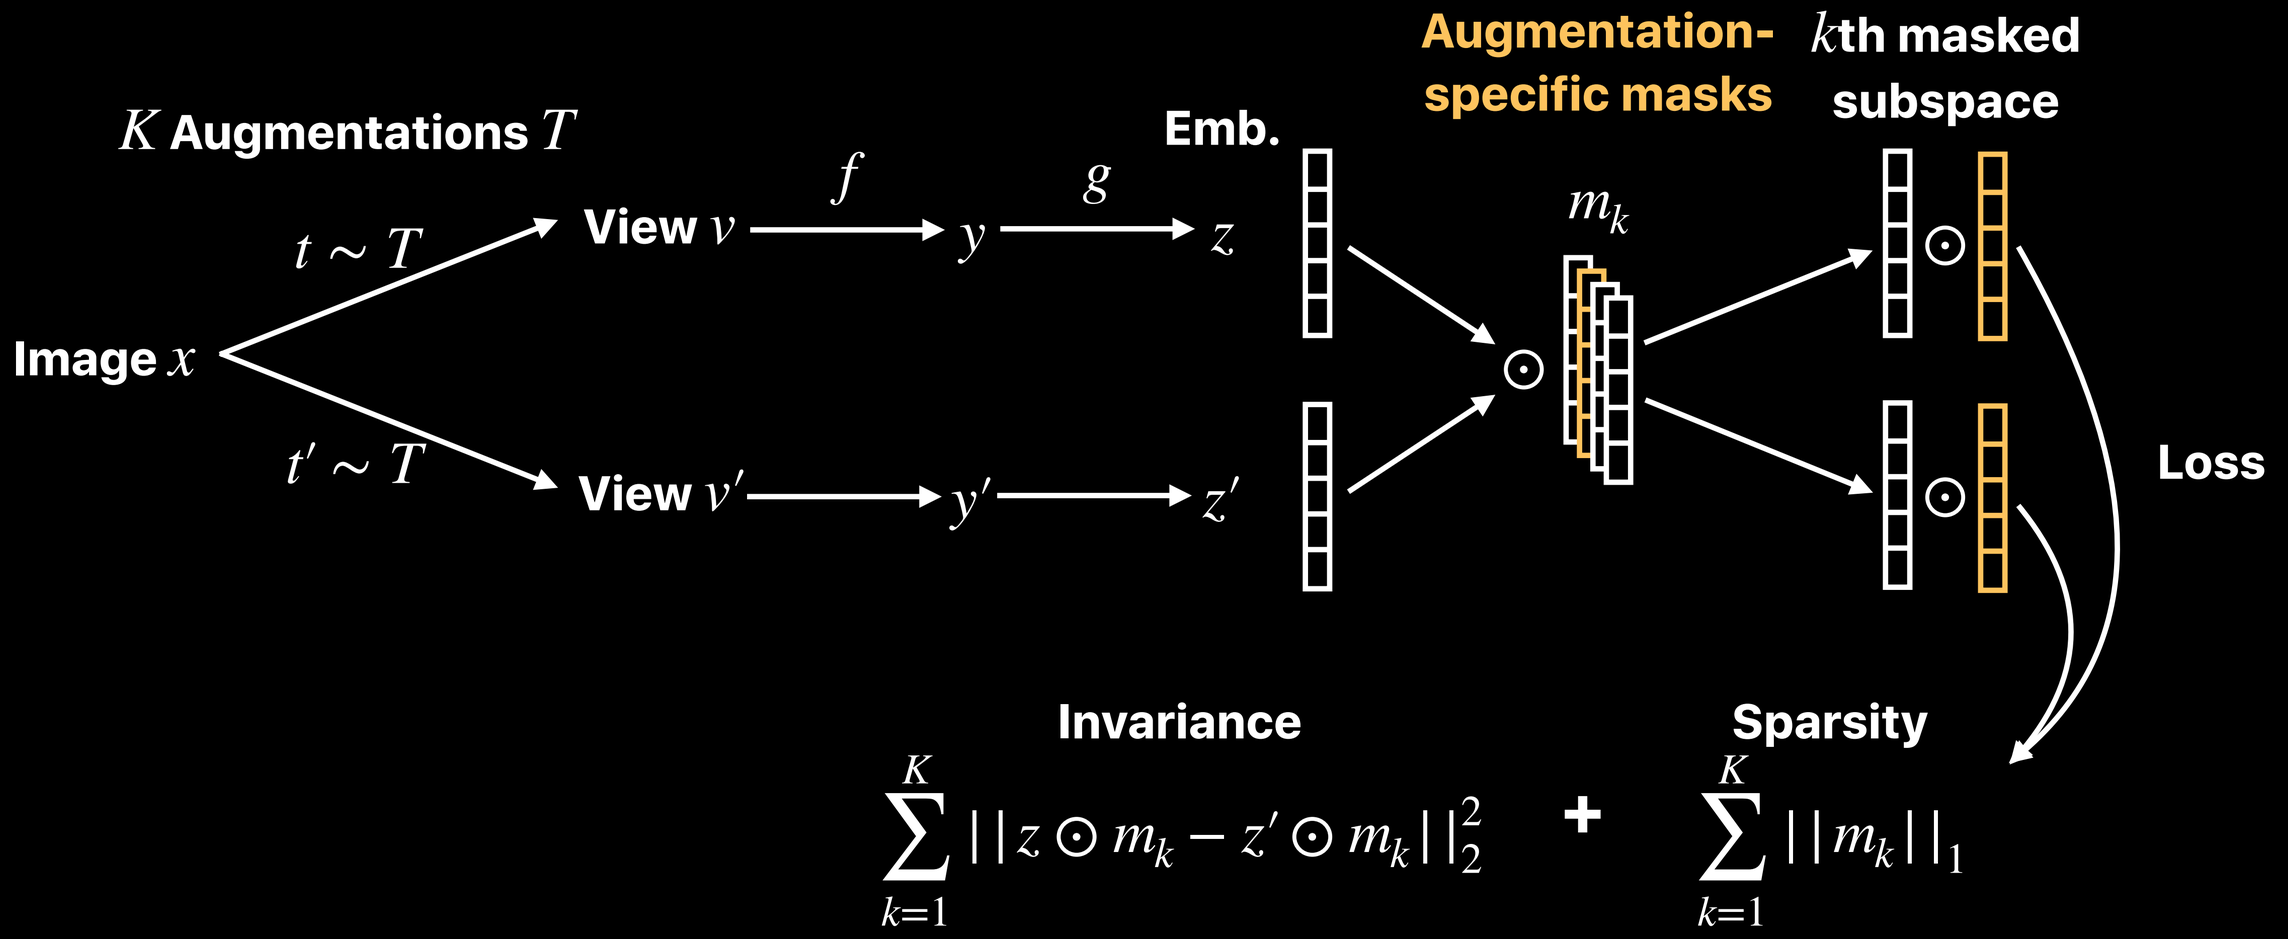

MAST

Visualization and analysis

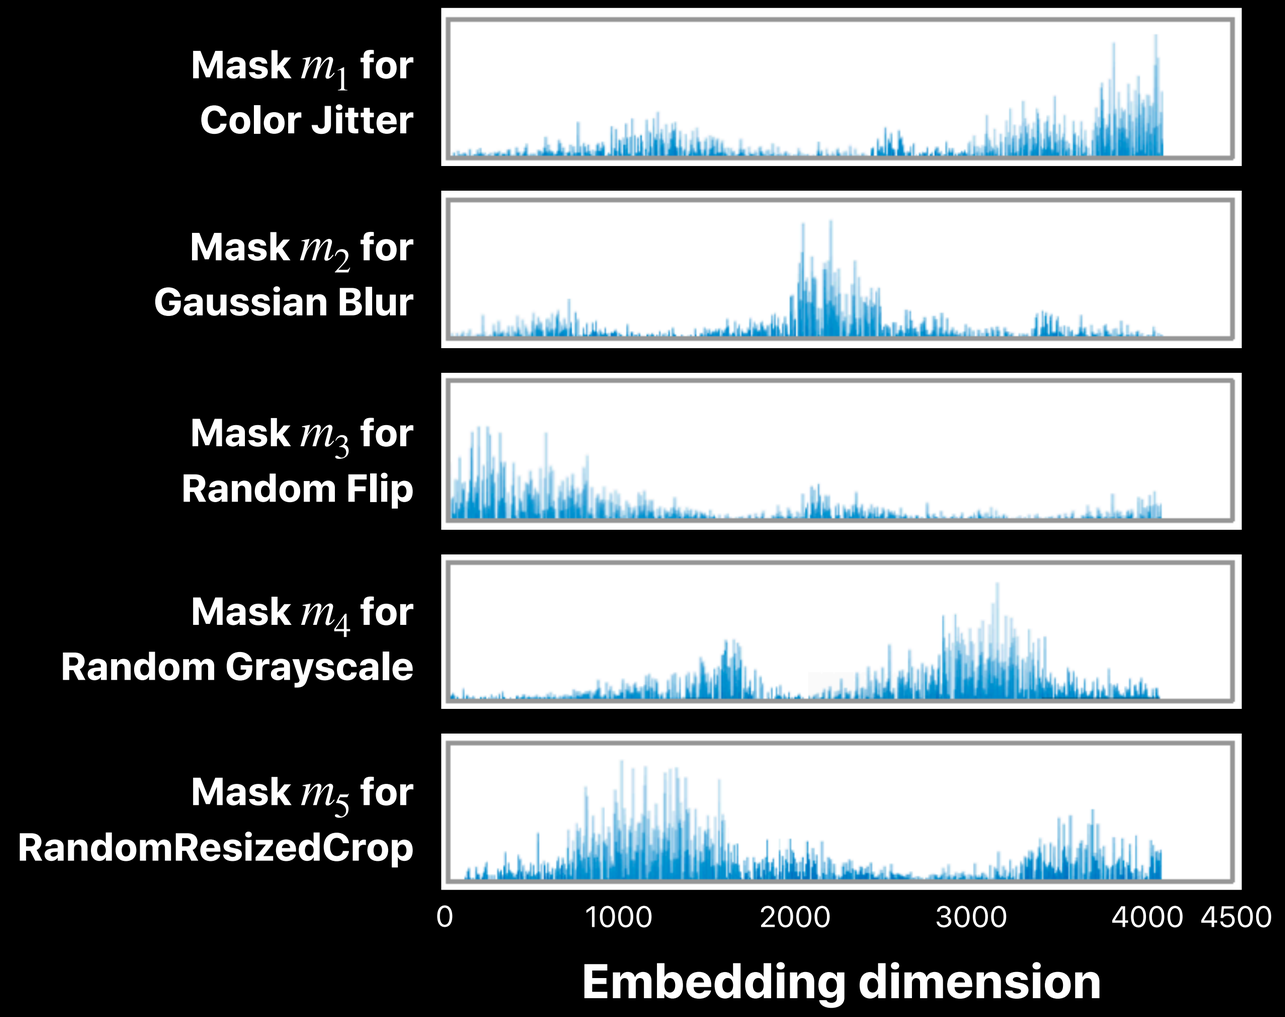

Benefits

- Masks can measure correlation/ redundancy of augmentations.
- Masked subspaces can encode contradicting invariances
- Preserves diverse priors for downstream generalization

MAST

Uncertainty modeling

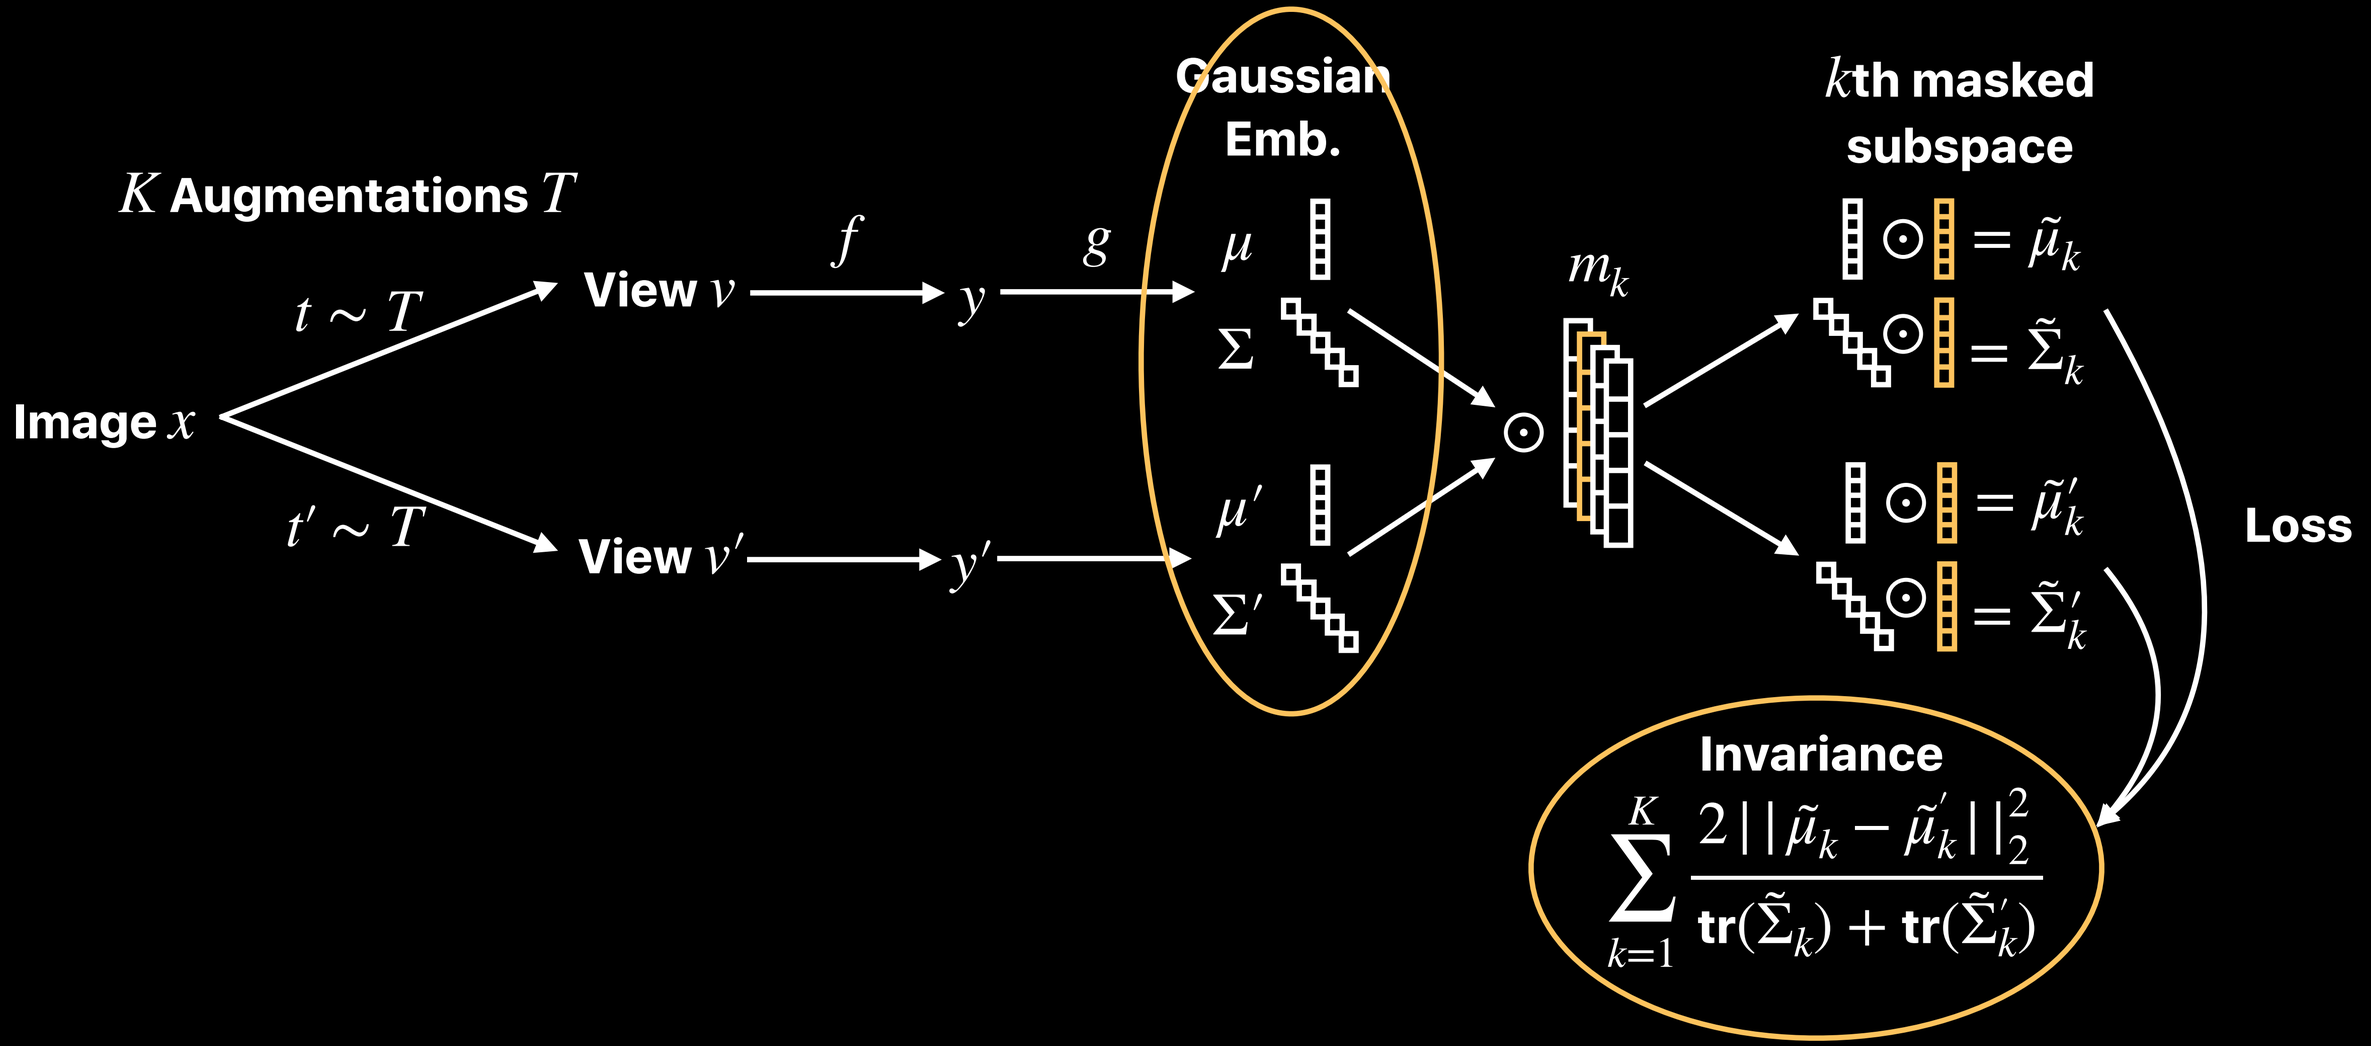

Results

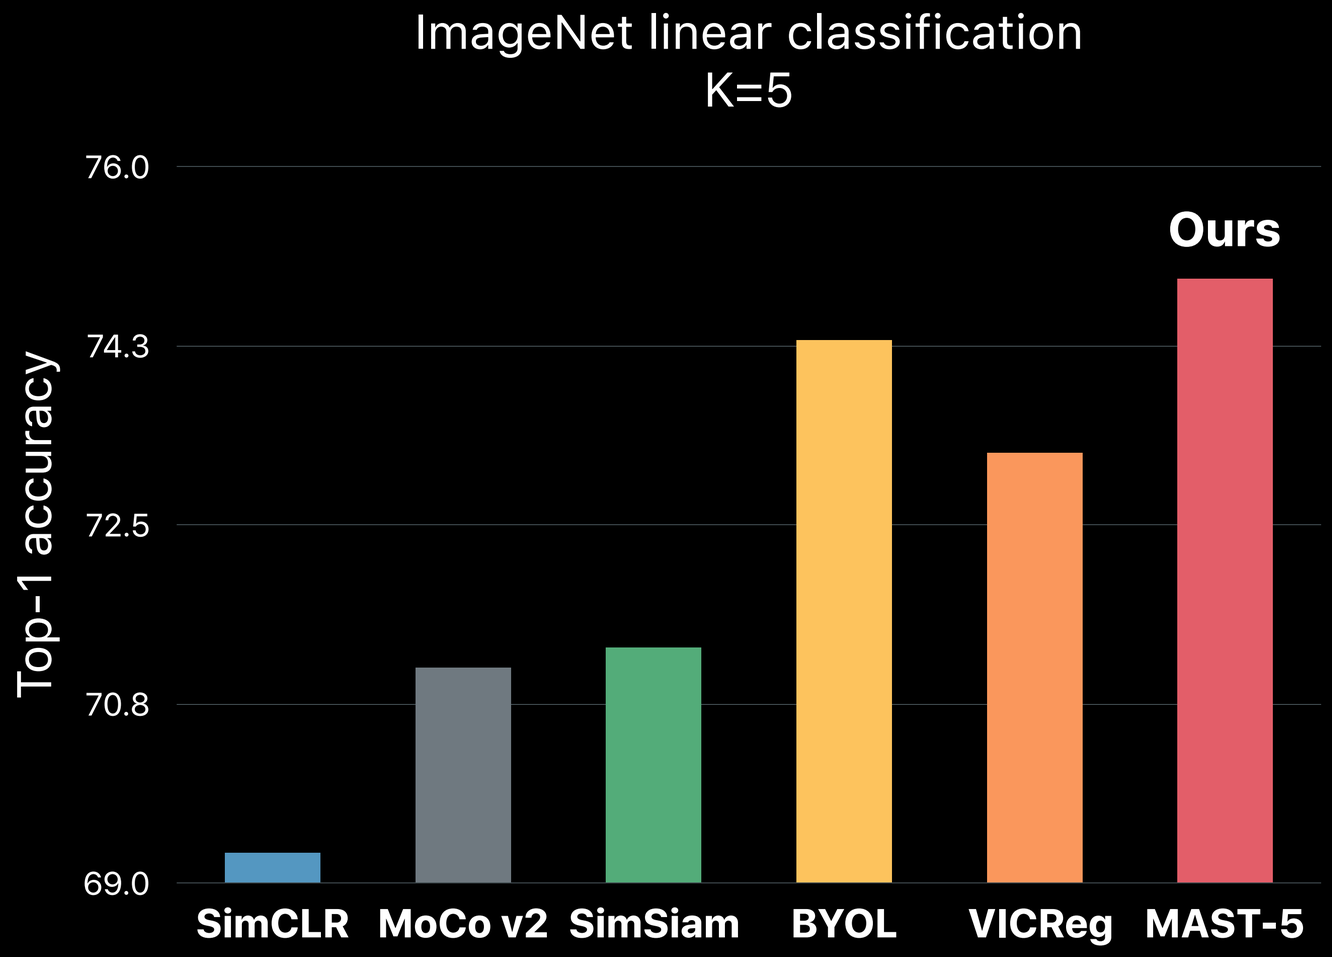

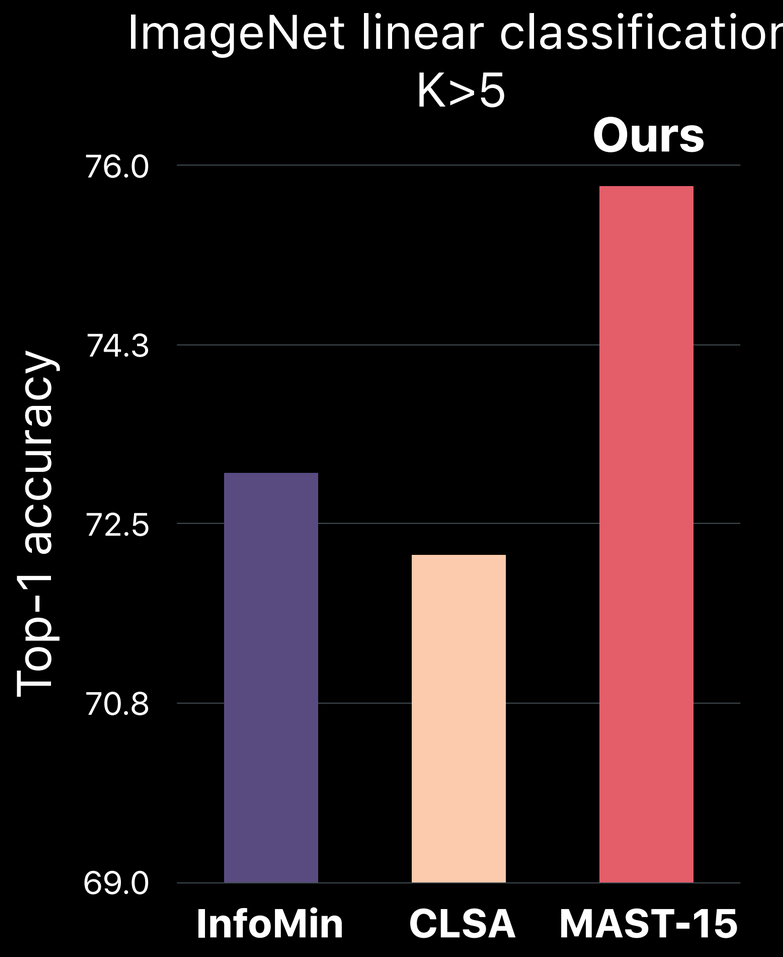

MAST remains competitive on various other downstream tasks

Conclusions

- MAST improves generalization of SSL via disentangled and uncertainty-aware learning of augmentation invariances
- Competitive on diverse downstream tasks without presuming any task information

Paper ID 4018