ICLR

Planckian Jitter: countering the color-crippling effects of color jitter on self-supervised training

Simone Zini [1] ✉️, Alex Gomez-Villa [2], Marco Buzzelli [1] 🗣️, Bartłomiej Twardowski [2,4], Andrew D. Bagdanov [3], Joost van de Weijer [2]

[1] Department of Informatics Systems and Communication, University of Milano – Bicocca, Milan, Italy
[2] Computer Vision Center, Universitat Autònoma de Barcelona, Barcelona, Spain
[3] Media Integration and Communication Center, University of Florence, Florence, Italy
[4] IDEAS NCBR, Warsaw, Poland

Self-supervised training and color jitter

- Self-supervised training:
  - Pretrain a model using only positive examples
    - Exploit self-consistency and data augmentation (jitter)
  - Use as starting point to train other downstream tasks
- The effects of color jitter:
  - Produces the best results in common downstream tasks
  - **How does it impact color-sensitive downstream tasks?**
    - Jitter induces invariance to the same transformations

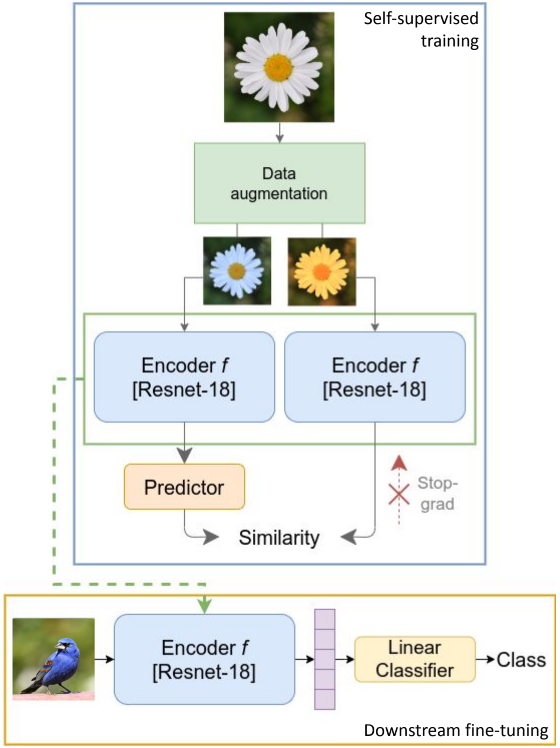

Chen, X., & He, K. (2021). Exploring simple siamese representation learning. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 15750-15758).

Default Color Jitter

- Transformations
  - Hue and saturation change
  - Random grayscale
  - Brightness and contrast change
- Experimental setup
  - Pretraining: SimSiam [Chen 2021]
  - Downstream tasks:
    - **CIFAR-100** (chromatic variations are irrelevant)

      → Default Color Jitter induces invariance
    - **FLOWERS-102** (color information is critical)

      → Default Color Jitter degrades the performance

AUGMENTATION	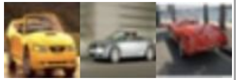 **CIFAR-100**	 **FLOWERS-102**
None	41.93%	36.47%
Default Color Jitter	↑ 59.93%	↓ 30.00%

Chen, X., & He, K. (2021). Exploring simple siamese representation learning. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 15750-15758).

Planckian Jitter

- Alternative color data augmentation based on the physical properties of light
  - Sample illuminant spectrum $\sigma_T(\lambda)$
    - from the distribution of a **black-body radiator**
  - Transform the sampled spectrum into sRGB $\rho_T \in \mathbb{R}^3$
  - Re-illuminate image via von-Kries-like transform
  - Add brightness and contrast variation
- Effects of Planckian Jitter:
  - Realistic color variations are effective on both tasks, compared to no augmentation
  - On the color-invariant CIFAR-100, Default Color Jitter still achieves better results

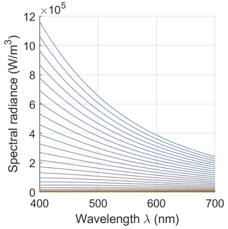

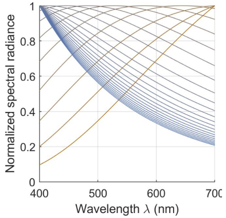

AUGMENTATION	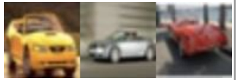 **CIFAR-100**	 **FLOWERS-102**
None	41.93%	36.47%
Default Color Jitter	↑ 59.93%	↓ 30.00%
Planckian Jitter	**↓ 47.31%**	**↑ 42.75%**

Latent Space Combination (LSC)

- Planckian Jitter
  - **Pros**: the learned representation yields a high-quality color description of scene objects
  - **Cons**: the quality of shape/texture representation drops (color is used in self-supervision to solve cases where previously shape and texture were required)
- To exploit all information, we learn:
  - One representation with Default Color Jitter
  - One representation with Planckian Jitter
  - Concatenate the representations into a single vector
- Results:
  - Improved performance on both tasks

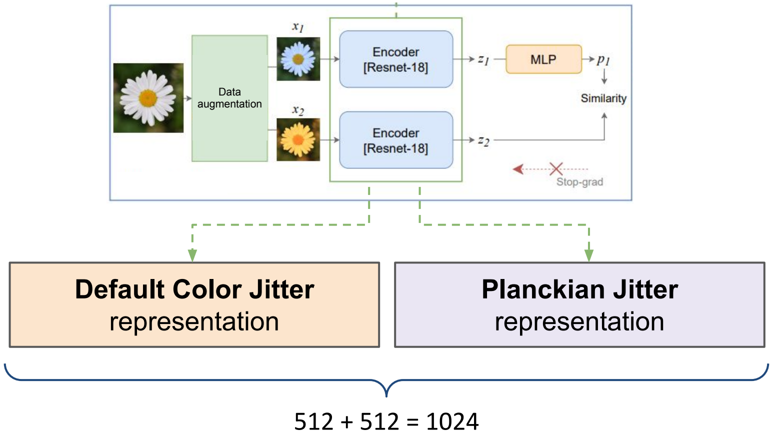

AUGMENTATION	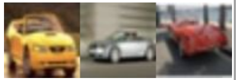 CIFAR-100	 FLOWERS-102
None	41.93%	36.47%
Default Color Jitter	↑ 59.93%	↓ 30.00%
Planckian Jitter	↓ 47.31%	↑ 42.75%
LSC	**↑ 63.54%**	**↑ 51.66%**

Planckian Jitter: countering the color-crippling effects of color jitter on self-supervised training

Simone Zini, Alex Gomez-Villa, Marco Buzzelli, Bartłomiej Twardowski, Andrew D. Bagdanov, Joost van de Weijer

- Overview:
  - Self-supervised training learns representations without labeled data, exploiting data augmentation (jitter)
  - Default Color Jitter deteriorates performance on color-sensitive downstream tasks
  - We propose a physics-based Planckian Jitter.
    - Beneficial when the intrinsic color of the objects is crucial for discrimination.
  - Exploiting both color and shape information (LSC) leads to overall superior results.
- See full paper for:
  - Applicability to other self-learning frameworks (SimCLR, BarlowTwins, VicReg)
  - Effects on other downstream tasks (Cub200, VegFru, T1K, USED)
  - Color sensitivity analysis

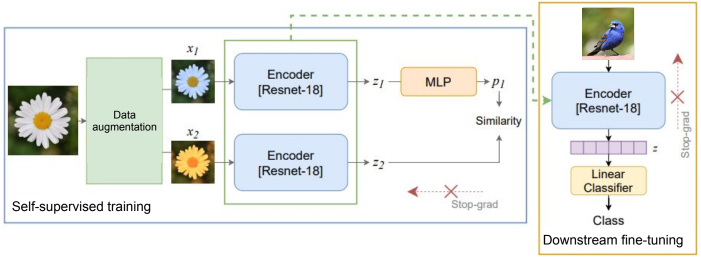

AUGMENTATION	CIFAR-100	FLOWERS-102
None	41.93%	36.47%
Default Color Jitter	↑ 59.93%	↓ 30.00%
Planckian Jitter	↓ 47.31%	↑ 42.75%
LSC	**↑ 63.54%**	**↑ 51.66%**

Download our code at:
https://github.com/TheZino/PlanckianJitter

Already available in Kornia