

Supplementary Material: Selective Vision-Language Subspace Projection for Few-shot CLIP

Anonymous Author(s)

1 MORE ANALYSIS OF MODALITY GAPS .

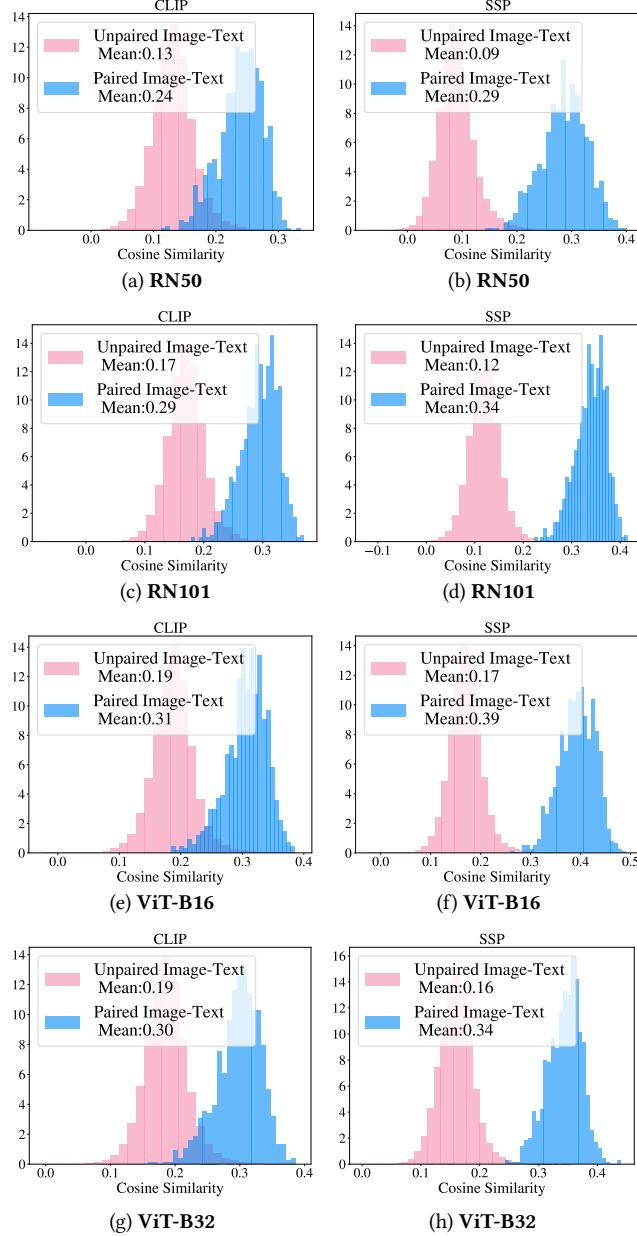


Figure A: Comparison of similarity distributions between text-image features of CLIP and SSP using different encoders on the ImageNet dataset.

To further analyze the effectiveness of our SSP working in narrowing modality gaps, we visualize the similarity distributions for both paired and unpaired image and text features. The comparisons are illustrated in Figure A. It is evident that in the original CLIP feature space, the similarities of paired image-text features only slightly exceed those of unpaired (the first column in Figure A). However, with the application of our SSP (the second column of Figure A), there is a noticeable improvement in reducing the similarities between unpaired image-text features (0.09 vs. 0.13 under RN50) and simultaneously enhancing the similarities between paired features (0.39 VS. 0.31 under ViT-B/16), which has effectively reduce the modality gaps.

2 DETAILS OF VON MISES-FISHER (VMF) DISTRIBUTION FORMULATION

As discussed in the previous works [4, 5], text and image features encoded from CLIP are located on a unit hypersphere. Consequently, these features follow a von Mises–Fisher (vMF) distribution [1]. The probability density function of the vMF distribution for the random d -dimensional unit vector f is expressed as:

$$p_d(f; \mu, \kappa) = \frac{1}{C_d(\kappa)} \exp \left(\kappa \mu^T f \right), \quad (1)$$

$$C_d(\kappa) = \frac{(2\pi)^{d/2} I_{(d/2-1)}(\kappa)}{\kappa^{d/2-1}},$$

where $\kappa \geq 0$, $\|\mu\|_2 = 1$ and $I_{(d/2-1)}$ denotes the modified Bessel function of the first kind at order $I_{(d/2-1)}$ as discussed in [1]. We utilize maximum likelihood estimation to calculate parameters μ and κ , where μ is the average of all features and κ is approximated as follows [3]:

$$\hat{\kappa} = \frac{\bar{R}(d - \bar{R}^2)}{1 - \bar{R}^2}. \quad (2)$$

Here \bar{R} is the length of sample mean, i.e., $\bar{R} = \|\bar{z}\|_2$, $\bar{z} = \frac{1}{N} \sum_{i=1}^N f_i$, and we calculate the $I_{(d/2-1)}(\kappa)$ based on the work [2]. Based on the above estimation, we can calculate metrics of distributions between the image and text features, as discussed in Figure 1 of the manuscript.

3 MORE PERFORMANCE COMPARISONS

Our SSP is a training-free method that can also be applied to training-required methods, such as Tip-F [4] and APE-T [5] mentioned in the manuscript. We evaluate our method against Tip-F and APE-T to demonstrate its efficacy. The summarized results of these applications are presented in Table A. We consistently observed that our SSP achieves the highest average accuracy across all benchmarks. However, in specific datasets like Aircraft and Cars under the 1-shot setting, it may not always outperform APE-T. When compared to Tip-F, our SSP demonstrates an average accuracy improvement of 0.68% in the 16-shot setting. Similarly, when

Table A: The classification accuracy (%) comparison on few-shot learning, i.e., 1-/2-/4-/6-/8-/16-shot, across 11 datasets. The results for LFA, Tip, and APE from our implementation by open public project, and the datasets include F102.(Flowers102), Euro(EuroSAT), F101.(Food101), SUN.(SUN397), C101.(Caltech101), UCF.(UFC101), and ImgN.(ImageNet).

Method	Pets	F102.	FGVC	DTD	Euro.	Cars	F101.	SUN.	C101.	UCF.	ImgN.	Avg.
CLIP	85.77	66.14	17.28	42.32	37.56	55.61	77.31	58.52	86.29	61.46	58.18	58.77
<i>16-shot</i>												
Tip-F (ECCV22)	89.70	94.40	35.28	67.71	84.95	76.41	79.33	71.28	92.84	76.74	65.48	75.83
Tip-F + SSP	91.17	95.21	35.58	68.62	85.62	76.70	79.50	71.89	93.27	78.46	65.63	76.51 ↑ 0.68
APE-T (ICCV23)	89.59	96.26	37.38	69.56	86.35	76.15	79.35	72.45	93.06	79.38	66.12	76.63
APE-T + SSP	90.11	96.43	37.65	70.45	86.75	76.40	79.49	72.65	93.57	79.57	66.12	77.20 ↑ 0.57
<i>8-shot</i>												
Tip-F (ECCV22)	88.20	91.88	30.30	63.12	77.72	69.72	78.59	69.02	92.06	74.62	63.93	72.65
Tip-F + SSP	90.35	93.71	30.57	63.56	79.23	71.68	78.92	69.80	92.48	74.89	64.14	73.58 ↑ 1.04
APE-T (ICCV23)	88.25	94.52	32.25	67.08	81.04	71.72	78.73	70.94	92.13	77.24	64.80	74.43
APE-T + SSP	88.88	94.84	32.52	67.49	81.16	71.19	78.64	71.08	92.45	77.77	64.83	74.62 ↑ 0.20
<i>4-shot</i>												
Tip-F (ECCV22)	87.50	89.40	25.86	57.80	73.43	64.85	78.06	66.32	91.81	71.87	62.53	69.95
Tip-F + SSP	87.77	89.77	26.13	58.45	74.01	66.18	78.36	67.20	92.18	73.13	63.09	70.57 ↑ 0.62
APE-T (ICCV23)	87.98	91.68	26.85	65.66	73.48	67.62	78.27	68.81	91.68	73.78	63.61	71.77
APE-T + SSP	88.23	91.92	27.12	65.78	74.35	68.28	78.30	68.88	91.68	74.76	63.65	72.09 ↑ 0.32
<i>2-shot</i>												
Tip-F (ECCV22)	87.23	83.19	22.92	54.31	66.26	63.00	77.71	63.92	90.30	68.28	61.76	67.17
Tip-F + SSP	87.51	84.41	24.78	54.91	67.64	64.53	77.76	65.94	90.72	70.07	62.53	68.25 ↑ 1.08
APE-T (ICCV23)	87.03	87.41	24.41	59.28	71.89	63.96	77.58	67.14	90.18	69.92	63.27	69.28
APE-T + SSP	87.57	88.75	24.93	59.04	72.37	63.89	77.64	67.00	90.87	69.97	63.28	69.57 ↑ 0.29
<i>1-shot</i>												
Tip-F (ECCV22)	86.41	80.02	20.82	48.52	58.95	58.65	77.50	62.56	90.30	65.01	61.17	64.54
Tip-F + SSP	86.67	81.93	21.15	49.29	61.19	59.91	77.62	62.85	90.56	65.44	61.95	65.32 ↑ 0.79
APE-T (ICCV23)	86.54	83.23	22.08	55.02	67.59	60.65	77.22	65.66	89.82	65.69	62.56	66.91
APE-T + SSP	86.73	84.21	21.63	55.61	68.63	60.55	77.53	65.71	90.43	66.56	62.58	67.29 ↑ 0.37

compared to APE-T, our SSP shows a 0.57% improvement. These consistent enhancements underscore the versatility and effectiveness of our SSP approach.

REFERENCES

- [1] Chaoqun Du, Yulin Wang, Shiji Song, and Gao Huang. 2024. Probabilistic Contrastive Learning for Long-Tailed Visual Recognition. *CoRR* abs/2403.06726 (2024).
- [2] Wittawat Jitkrittum, Aditya Krishna Menon, Ankit Singh Rawat, and Sanjiv Kumar. 2022. ELM: Embedding and Logit Margins for Long-Tail Learning. *CoRR* abs/2204.13208 (2022).
- [3] Suvrit Sra. 2012. A short note on parameter approximation for von Mises-Fisher distributions: and a fast implementation of I s (x). *Computational Statistics* 27 (2012), 177–190.
- [4] Renrui Zhang, Wei Zhang, Rongyao Fang, Peng Gao, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. 2022. Tip-Adapter: Training-Free Adaption of CLIP for Few-Shot Classification. In *ECCV* (35), Vol. 13695. 493–510.
- [5] Xiangyang Zhu, Renrui Zhang, Bowei He, Aojun Zhou, Dong Wang, Bin Zhao, and Peng Gao. 2023. Not All Features Matter: Enhancing Few-shot CLIP with Adaptive Prior Refinement. *CoRR* abs/2304.01195 (2023).