

Supplementary Materials: CP-Prompt: Composition-Based Cross-modal Prompting for Domain-Incremental Continual Learning

Anonymous Authors

1 DATASETS

We use public datasets to perform experiments and evaluate the performance of CP-Prompt. The detailed descriptions and statistics of these datasets are as follows:

- **CDDDB**¹ [11] is a dataset used for continuous deepfake detection, where the DIL objective involves recognizing authentic and fake images across different domains. We adopted the Hard Setting from [Wang et al., 2022a], requiring learning on 5 continuous deepfake detection domains: GauGAN, BigGAN, WildDeepfake, WhichFaceReal, and SAN. This entails approximately 27,000 images.
- **CORE50**² [12] is designed for continuous object recognition, consisting of 11 domains, each with 50 categories. In DIL, the goal is to sequentially learn from 8 domains for incremental training and assess performance on the remaining 3 domains (unseen). This encompasses about 160,000 images.
- **DomainNet**³ [14] is a domain adaptation dataset commonly used as a benchmark for DIL methods. It comprises 6 domains, each with 345 categories. The DIL setup aligns with CaSSLe [Fini et al., 2022]. This involves around 600,000 images in total.

2 BASELINES

In our experiments, the compared baselines mainly include replay-based methods (iCaRL, LUCIR and LRCIL), distillation method (BiC), regularization method (EWC), self-supervised (SimCLR, BYOL, Barlow Twins and Supervised Contrastive), other non-prompt methods (ER, GDumb, DER++, and Co²L), and prompt-related methods (L2P, DyTox and S-liPrompts).

- **iCaRL** [16] proposes a prioritized exemplar selection based on herding strategy to filter old samples.
- **LUCIR** [8] proposes three loss functions to constrain the bias problem caused by the imbalance of old and new samples in the Distillation-based method.
- **LRCIL** [13] proposes a method to store activations volumes at some intermediate layer instead of storing a portion of past data.
- **BiC** [21] proposes a method to introduce a bias correction layer after the fully connected layer to offset the domain shift phenomenon.
- **EWC** [10] proposes a method to constrain historical model parameters to alleviate forgetfulness through a quadratic penalty function.

- **SimCLR** [4] utilizes a self-supervised framework of contrastive learning architecture, converted to incremental learning loss under the CaSSLe[6] architecture.
- **BYOL** [7] uses a self-supervised framework of contrastive learning architecture without negative sample pairs to perform domain incremental learning under the CaSSLe architecture.
- **Barlow Twins** [22] uses high-dimensional embedding to improve self-supervised learning performance and is applied in the CaSSLe architecture.
- **Supervised Contrastive** [9] proposes a method to have multiple anchors and multiple positive sample adaptive versions in a minibatch, and is applied in the CaSSLe architecture.
- **ER** [3] proposes to use the current gradient and historical gradient weighting when updating the gradient.
- **GDumb** [15] proposes that when learning new knowledge, the sampler uniformly samples old samples to balance the sample distribution.
- **DER++** [1] proposes an experience replay method, which requires the model to learn the process of outputting approximations to old samples rather than the results.
- **Co²L** [2] proposes a decoupled representation-classifier scheme with the aim of continuously learning and preserving representations.
- **L2P** [20] proposes a query strategy based on key-value pairs to dynamically select the corresponding prompt for each input.
- **DyTox** [5] proposes a transformer architecture based on a dedicated codec framework. Encoders and decoders are shared between all tasks. Through dynamic expansion of special tokens, the decoder network can be dedicated to personalized domain task distribution.
- **Dual-Prompt** [19] is a class incremental learning method. It proposes General Prompt and Expert Prompt to be embedded in Attention in a prefix-tuning manner.
- **HiDe-Prompt** [17] proposes a hierarchical components such as within-task prediction, task-identity inference and task-adaptive prediction. It saves more fine-grained class features and optimizes prompts through contrast loss to better distinguish class features.
- **S-liPrompts** [18] proposes a win-win strategy to solve the domain increment problem. By learning cross-domain independent prompts, the model can get the best performance in each domain without any mutual interference, and storing the learned prompts to eliminate the catastrophic forgetting problem.

¹<https://github.com/Coral79/CDDDB>

²<https://vlomonaco.github.io/core50/index.html#dataset>

³<http://ai.bu.edu/M3SDA/>

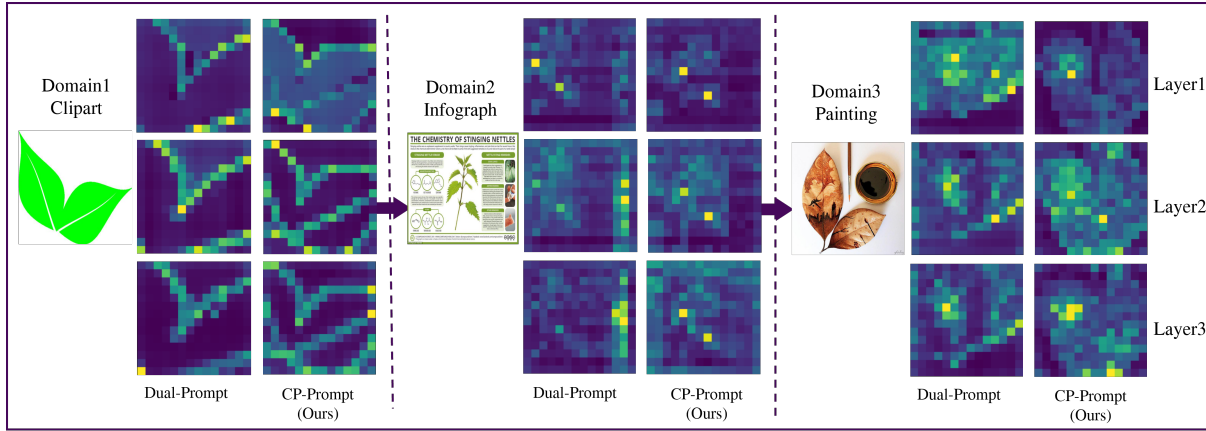


Figure 1: Performance variation of CP-Prompt by (a) using different prompt lengths; (b) adding new domain data.

3 IMPLEMENTATION DETAILS

Running Environment. We conduct all experiments on Ubuntu 18.04.2 LTS server with Intel(R) Xeon(R) CPU E5-2680 v4 @ 2.40GHz, 256G RAM and 8 NVIDIA GeForce RTX 3090-24GB. We implement CP-Prompt with Python 3.7.6 and PyTorch 1.7.0.

Hyper-parameter Settings. We implement our CP-Prompt in PyTorch. To ensure a fair comparison of experimental results, the personalized prompt length is set to 16, and the common prompt length is set to 6. The number of selected layers for personalized prompts ranges from 0 to 10. Pre-trained models are uniformly chosen as CLIP's ViT-B/16 version. The maximum number of epochs for CDDB-Hard dataset is set to 50, for CORE50 is 20, and for DomainNet is 30. The learning rate is set to 0.01, and the batch size is fixed at 128. For other details, please refer to https://anonymous.4open.science/r/CP_Prompt-C126.

4 MORE INFORMATION ABOUT COMMON PROMPT ATTENTION WEIGHTS

The common prompt architecture we designed can transfer knowledge between domains more effectively. This design can pay attention to effective information faster. For example in Figure 1, the CP-Prompt we propose can focus on core and effective information faster.

REFERENCES

- [1] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. 2020. Dark Experience for General Continual Learning: a Strong, Simple Baseline. In *NeurIPS 2020*, Hugo Larochelle, Marc' Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.).
- [2] Hyuntak Cha, Jaeho Lee, and Jinwoo Shin. 2021. Co²L: Contrastive Continual Learning. In *ICCV 2021*. 9496–9505.
- [3] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc' Aurelio Ranzato. 2019. On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486* (2019).
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In *ICML 2020*, Vol. 119. 1597–1607.
- [5] Arthur Douillard, Alexandre Ramé, Guillaume Couairon, and Matthieu Cord. 2022. DyTox: Transformers for Continual Learning with DYnamic Token eXpansion. In *CVPR 2022*. 9275–9285.
- [6] Enrico Fini, Victor G. Turrissi da Costa, Xavier Alameda-Pineda, Elisa Ricci, Karteek Alahari, and Julien Mairal. 2022. Self-Supervised Models are Continual Learners. In *CVPR 2022*. 9611–9620.
- [7] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Ávila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. 2020. Bootstrap Your Own Latent - A New Approach to Self-Supervised Learning. In *NeurIPS 2020*.
- [8] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. 2019. Learning a Unified Classifier Incrementally via Rebalancing. In *CVPR 2019*. 831–839.
- [9] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. 2020. Supervised Contrastive Learning. In *NeurIPS 2020*.
- [10] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences* 114, 13 (2017), 3521–3526.
- [11] Chuqiao Li, Zhiwu Huang, Danda Pani Paudel, Yabin Wang, Mohamad Shahbazi, Xiaopeng Hong, and Luc Van Gool. 2023. A Continual Deepfake Detection Benchmark: Dataset, Methods, and Essentials. In *WACV 2023*. 1339–1349.
- [12] Vincenzo Lomonaco and Davide Maltoni. 2017. CORE50: a New Dataset and Benchmark for Continuous Object Recognition. In *CoRL 2017*, Vol. 78. 17–26.
- [13] Lorenzo Pellegrini, Gabriele Graffieti, Vincenzo Lomonaco, and Davide Maltoni. 2020. Latent Replay for Real-Time Continual Learning. In *IROS 2020*. 10203–10209.
- [14] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. 2019. Moment Matching for Multi-Source Domain Adaptation. In *ICCV 2019*. 1406–1415.
- [15] Ameya Prabhu, Philip H. S. Torr, and Puneet K. Dokania. 2020. GDumb: A Simple Approach that Questions Our Progress in Continual Learning. In *ECCV 2020*, Vol. 12347. 524–540.
- [16] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H. Lampert. 2017. iCaRL: Incremental Classifier and Representation Learning. In *CVPR 2017*. 5533–5542.
- [17] Liyuan Wang, Jingyi Xie, and Xingxing Zhang et al. 2023. Hierarchical Decomposition of Prompt-Based Continual Learning: Rethinking Obscured Sub-optimality. In *NeurIPS 2023*.
- [18] Yabin Wang, Zhiwu Huang, and Xiaopeng Hong. 2022. S-Prompts Learning with Pre-trained Transformers: An Occam's Razor for Domain Incremental Learning. In *NeurIPS 2022*.
- [19] Zifeng Wang, Zizhao Zhang, and Sayna Ebrahimi et al. 2022. DualPrompt: Complementary Prompting for Rehearsal-Free Continual Learning. In *ECCV*.
- [20] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer G. Dy, and Tomas Pfister. 2022. Learning to Prompt for Continual Learning. In *CVPR 2022*. 139–149.
- [21] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. 2019. Large Scale Incremental Learning. In *CVPR 2019*. 374–382.
- [22] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. 2021. Barlow Twins: Self-Supervised Learning via Redundancy Reduction. In *ICML 2021*, Vol. 139. 12310–12320.