

TorchCP: A Python Library for Conformal Prediction

Jianguo Huang^{1,2}, Jianqing Song^{1,3}, Xuanning Zhou^{1,4},
Bingyi Jing^{1,4}, Hongxin Wei^{1*}

¹Southern University of Science and Technology ²Nanyang Technological University

³Nanjing University ⁴The Chinese University of Hong Kong, Shenzhen

Editor: Zeyi Wen

Abstract

Conformal prediction (CP) is a powerful statistical framework that generates prediction intervals or sets with guaranteed coverage probability. While CP algorithms have evolved beyond traditional classifiers and regressors to sophisticated deep learning models like deep neural networks (DNNs), graph neural networks (GNNs), and large language models (LLMs), existing CP libraries often lack the model support and scalability for large-scale deep learning (DL) scenarios. This paper introduces TorchCP, a PyTorch-native library designed to integrate state-of-the-art CP algorithms into DL techniques, including DNN-based classifiers/regressors, GNNs, and LLMs. Released under the LGPL-3.0 license, TorchCP comprises about 16k lines of code, validated with 100% unit test coverage and detailed documentation. Notably, TorchCP enables CP-specific training algorithms, on-line prediction, and GPU-accelerated batch processing, achieving up to 90% reduction in inference time on large datasets. With its low-coupling design, comprehensive suite of advanced methods, and full GPU scalability, TorchCP empowers researchers and practitioners to enhance uncertainty quantification across cutting-edge applications.

Keywords: Conformal Prediction, Deep Learning

1. Introduction

Conformal prediction (CP) is a statistical framework generating prediction intervals or sets with guaranteed coverage probability (Vovk et al., 2022), offering a distribution-free solution to quantify predictive uncertainty. While CP algorithms have been developed beyond traditional classifiers and regressors to sophisticated deep learning (DL) models like graph neural networks (GNNs) and large language models (LLMs), existing CP libraries often lack the model support and scalability required for large-scale DL scenarios. This underscores the need for a comprehensive and scalable CP toolbox tailored for DL models, facilitating robust and accessible conformal prediction methods across cutting-edge applications.

This work presents TorchCP, a PyTorch-native library that advances conformal prediction for DL by implementing a wide range of state-of-the-art CP algorithms, including training and inference methods tailored for models such as DNN classifiers/regressors, GNNs, and LLMs. Its low-coupling design facilitates easy customization of components like trainers or scores, while full GPU support and batch processing ensure seamless scalability for large-scale tasks. TorchCP excels over existing libraries by delivering superior speed, with up to 90% faster inference times on large datasets compared to MAPIE and PUNCC (Appendix F),

*. Corresponding author: Hongxin Wei (weihx@sustech.edu.cn)

while also encompassing a broader algorithmic coverage than **TorchUQ**. This PyTorch-native integration and extensibility empower statisticians and DL practitioners to seamlessly apply CP across cutting-edge applications with enhanced DL compatibility.

2. Main contributions

Conformal prediction constructs a prediction set $C(X)$ for an input X , ensuring the true outcome Y satisfies $P(Y \in C(X)) \geq 1 - \alpha$, where α is a miscoverage rate. Existing CP libraries encounter significant challenges in DL, including limited native support for modern models (like DNNs, GNNs, and LLMs), poor computational efficiency when scaling to large datasets, and monolithic designs that restrict the flexible customization for diverse applications. **TorchCP** overcomes these shortcomings through a PyTorch-native framework that improves model compatibility, efficiency, and adaptability, thereby enhancing CP’s applicability to DL. Its low-coupling design and GPU-accelerated batch processing provide a robust foundation for tackling complex, real-world problems, detailed in the following parts.

2.1 Advantages for DL scenarios

Low-coupling design Low-coupling design is one of the most fundamental concepts in object-oriented design (McNatt and Bieman, 2001), which minimizes the dependency a class has on other classes. While many libraries (e.g., **PUNCC**) are implemented in a high-coupling design, **TorchCP** is structured into three core modules: **Trainer**, **Score**, and **Predictor**, each tailored to specific aspects of CP workflows. The **Trainer** module implements CP-specific training algorithms, such as **ConfTr** (Stutz et al., 2022), which optimizes deep models to improve efficiency. The **Score** module calculates non-conformity measures, e.g., LAC (Sadinle et al., 2019), to quantify how atypical a new sample is relative to the training data. The **Predictor** module implements CP workflows, including calibration and prediction, to generate outputs, e.g., prediction sets/intervals, p-values, and conformal predictive distributions.

This low-coupling design is visualized in Figure 1, illustrating the interplay among modules. The **Score** module is used to create various score objects that serve as inputs for both the **Trainer** and **Predictor** modules. The **Trainer** module optimizes DL models using CP-specific loss functions. The **Predictor** module then integrates the trained model with the score objects to generate conformal outputs, such as prediction sets/intervals, p-values, or predictive distributions. In this manner, **TorchCP** achieves strong reusability and extensibility by providing standardized module interfaces, making it straightforward to define new methods. Appendix B provides guidelines for adding customized modules using the provided interfaces, while Appendix E discusses the advantages of **TorchCP**’s low-coupling design.

Methodological scope Table 1 outlines the CP algorithms in **TorchCP** across 4 types of DL models¹: DNN classifiers/regressors, GNNs, and LLMs, along with evaluation metrics in Table 4. Current libraries are typically limited in algorithmic coverage, e.g., **MAPIE** only

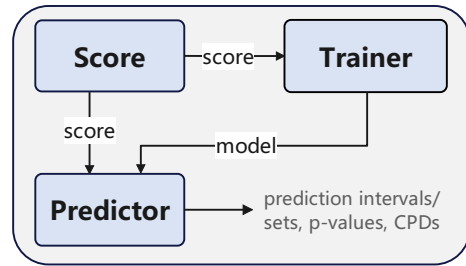


Figure 1: The modules of **TorchCP**.

1. Model types are defined by their input/output data differences, impacting the choice of CP algorithms.

Model types	Component	CP Algorithms
DNN Classifier	Trainer&Loss	ConfTr (Stutz et al., 2022), ConfTS (Xi et al., 2025), C-Adapter (Liu et al., 2025), UA Loss (Einbinder et al., 2022), SCPO (Bellotti, 2021)
	Score Function	LAC (Sadinle et al., 2019), APS (Romano et al., 2020), RAPS (Angelopoulos et al., 2020), SAPS (Huang et al., 2024), KNN (Gammerman and Vovk, 2007), Margin (Löfström et al., 2015), TOPK (Angelopoulos et al., 2020)
	Predictor	Split CP (Papadopoulos et al., 2002), Class-conditional CP (Vovk, 2012), Clustered CP (Ding et al., 2024), Weighted CP (Tibshirani et al., 2019), RC3P (Shi et al., 2024)
DNN Regressor	Score Function	ABS (Papadopoulos et al., 2002), NORABS (Papadopoulos et al., 2011), CQR (Romano et al., 2019), CQRR (Sesia and Candès, 2020), CQRM (Sesia and Candès, 2020), CQRFM (Kivaranovic et al., 2020), R2CCP (Guha et al., 2023)
	Predictor	Split CP, Ensemble CP (Xu and Xie, 2021), ACI (Gibbs and Candès, 2021), AgACI (Zaffran et al., 2022)
GNN	Trainer&Loss	CF-GNN (Huang et al., 2023)
	Score Function	DAPS (Zargarbashi et al., 2023), SNAPS (Song et al., 2024)
	Predictor	NAPS (Clarkson, 2023)
LLM	Predictor	Conformal LLM (Quach et al., 2024)

Table 1: The CP algorithms implemented in TorchCP for various deep learning models.

includes split CP. In TorchCP, we provide 6 training algorithms, 17 scores, and 10 predictors. A brief overview of all supported methods and their advantages is available in Appendix A.

Batch Processing and GPU Support TorchCP is a PyTorch-native library equipped with GPU acceleration and batch processing capabilities. It enables efficient computation on large datasets and scales seamlessly to DL applications, e.g., achieving up to 90% reduction in inference time on ImageNet (See Figure 6b). Appendix D.2 presents empirical evidence of its efficiency advantages on large-scale datasets.

2.2 Tutorials

TorchCP provides a user-friendly framework for conformal prediction in deep learning with a detailed documentation². The library contains illustrative code showcasing functionality for common applications (See the “examples” folder³) with a practical guide in Appendix C. Moreover, in Appendix D, we provide reproduced results of classical CP methods, confirming the correctness of our implementation. Below, we present a generic example⁴ to show the usability of TorchCP. First, we utilize ConfTr (Stutz et al., 2022) to train a neural network, i.e., `init_model`, on a training dataset encapsulated in `train_loader`:

2. <https://torchcp.readthedocs.io/en/latest/index.html>

3. <https://github.com/ml-stat-Sustech/TorchCP/tree/master/examples>

4. The full source code is available at TorchCP CIFAR-100 Example.

Library	Model Support				Functionality Support						Coupling	Backend
	Classifier	Regressor	GNN	LLM	Training	Online	GPU	Batch	p -value	CPD		
nonconformist (Linusson et al., 2020)	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	low	Numpy
TorchUQ (Zhao et al., 2024)	✗	✓	✗	✗	✗	✓	✓	✓	✗	✗	low	PyTorch
Crepes (Boström, 2022)	✓	✓	✗	✗	✗	✓	✗	✗	✓	✓	low	Numpy
PUNCC (Mendil et al., 2023)	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	high	Numpy
MAPIE (Cordier et al., 2023)	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	high	Numpy
Fortuna (Detommaso et al., 2024)	✓	✓	✗	✗	✗	✗	✓	✓	✗	✗	high	JAX
TorchCP (ours)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	low	PyTorch

Table 2: Comparison of Open-Source CP Libraries⁵. “Classifier” and “Regressor” refer to DNN-based classifiers and regressors⁶. The term “Training” denotes support for CP-specific training, while “Online” refers to capabilities for semi-online conformal prediction. “ p -value” and “CPD” denote outputting p -values and conformal predictive distributions, respectively.

```

1 from torchcp.classification.trainer import ConfTrTrainer
2 trainer = ConfTrTrainer(init_model, alpha=0.1) # Initialize trainer with model
3 trained_model = trainer.train(train_loader, num_epochs=10) # Training 10 epochs
    
```

With the trained model, users can proceed with conformal prediction by selecting an appropriate score function and predictor. Below, we implement Split CP with a miscoverage rate of $\alpha = 0.1$, using the LAC score on calibration data encapsulated in `cal_loader`:

```

1 predictor = SplitPredictor(score_function=LAC(), model=trained_model, alpha=0.1)
2 predictor.calibrate(cal_loader) # Set threshold with calibration data
3 predict_set = predictor.predict(test_instance) # Predict set for test instance
    
```

3. Related library

Conformal prediction is a distribution-free framework for quantifying predictive uncertainty, yet existing libraries face notable limitations in DL applications. Libraries, such as **nonconformist**, **PUNCC**, and **MAPIE**, offer restricted native support for modern models like GNNs and LLMs, relying on NumPy without GPU acceleration, which hinders efficient scaling to large datasets. Moreover, tools like **TorchUQ** and **Fortuna** provide some GPU support but fall short in algorithmic coverage, e.g., **TorchUQ** supports only regressors. Additionally, many libraries like **PUNCC** are implemented in a high-coupling design, where predictors and score functions are tightly coupled within a single class. **TorchCP** addresses these gaps with its PyTorch-native integration, supporting a broader range of modern models and offering a modular, extensible architecture, as detailed in Table 2. We present a detailed comparison of the relevant libraries in Appendix F.

3. The definitions of model support and functionality support are presented in Appendix F.

4. **TorchCP** also supports traditional machine learning models like random forests, but is primarily designed for advanced DL models within the PyTorch ecosystem.

5. The definitions of model support and functionality support are presented in Appendix F.

6. **TorchCP** also supports traditional machine learning models like random forests, but is primarily designed for advanced DL models within the PyTorch ecosystem.

4. Conclusion and Future Work

In this work, we introduced TorchCP, a PyTorch-native library that seamlessly integrates conformal prediction algorithms into deep learning workflows. Supporting a wide range of state-of-the-art CP algorithms, TorchCP leverages PyTorch’s GPU acceleration and batch processing, offering a low-coupling design for easy customization and user-friendly documentation for accessibility. Looking ahead, we plan to expand TorchCP by incorporating emerging CP algorithms for evolving architectures like diffusion models and multimodal models, while optimizing for real-time performance. Through open-source collaboration on GitHub, we aim to foster community-driven innovation, empowering reliable AI in fields such as healthcare, finance, and autonomous systems.

Acknowledgments

Hongxin Wei is supported by the Shenzhen Fundamental Research Program (Grant No. JCYJ20230807091809020). Bingyi Jing’s research is partly supported by NSFC 12371290. We gratefully acknowledge the support of the Center for Computational Science and Engineering at the Southern University of Science and Technology for our research.

References

- Anastasios N Angelopoulos, Stephen Bates, et al. Conformal prediction: A gentle introduction. *Foundations and Trends® in Machine Learning*, 16(4):494–591, 2023.
- Anastasios N Angelopoulos, Rina Foygel Barber, and Stephen Bates. Theoretical foundations of conformal prediction. *arXiv preprint arXiv:2411.11824*, 2024.
- Anastasios Nikolas Angelopoulos, Stephen Bates, Michael Jordan, and Jitendra Malik. Uncertainty sets for image classifiers using conformal prediction. In *International Conference on Learning Representations*, 2020.
- Vineeth Balasubramanian, Shen-Shyang Ho, and Vladimir Vovk. *Conformal prediction for reliable machine learning: theory, adaptations and applications*. Newnes, 2014.
- Rina Foygel Barber, Emmanuel J Candes, Aaditya Ramdas, and Ryan J Tibshirani. Conformal prediction beyond exchangeability. *The Annals of Statistics*, 51(2):816–845, 2023.
- Stephen Bates, Anastasios Angelopoulos, Lihua Lei, Jitendra Malik, and Michael Jordan. Distribution-free, risk-controlling prediction sets. *Journal of the ACM (JACM)*, 68(6): 1–34, 2021.
- Anthony Bellotti. Optimized conformal classification using gradient descent approximation. *arXiv preprint arXiv:2105.11255*, 2021.
- Henrik Boström. crepes: a python package for generating conformal regressors and predictive systems. In Ulf Johansson, Henrik Boström, Khuong An Nguyen, Zhiyuan Luo, and Lars Carlsson, editors, *Proceedings of the Eleventh Symposium on Conformal and Probabilistic*

- Prediction and Applications*, volume 179 of *Proceedings of Machine Learning Research*. PMLR, 2022.
- Jase Clarkson. Distribution free prediction sets for node classification. In *International Conference on Machine Learning*, pages 6268–6278. PMLR, 2023.
- Thibault Cordier, Vincent Blot, Louis Lacombe, Thomas Morzadec, Arnaud Capitaine, and Nicolas Brunel. Flexible and Systematic Uncertainty Estimation with Conformal Prediction via the MAPIE library. In *Conformal and Probabilistic Prediction with Applications*, 2023.
- Gianluca Detommaso, Alberto Gasparin, Michele Donini, Matthias Seeger, Andrew Gordon Wilson, and Cedric Archambeau. Fortuna: A library for uncertainty quantification in deep learning. *Journal of Machine Learning Research*, 25(238):1–7, 2024.
- Tiffany Ding, Anastasios Angelopoulos, Stephen Bates, Michael Jordan, and Ryan J Tibshirani. Class-conditional conformal prediction with many classes. *Advances in Neural Information Processing Systems*, 36, 2024.
- Bat-Sheva Einbinder, Yaniv Romano, Matteo Sesia, and Yanfei Zhou. Training uncertainty-aware classifiers with conformalized deep learning. *Advances in Neural Information Processing Systems*, 35:22380–22395, 2022.
- Alexander Gammerman and Vladimir Vovk. Hedging predictions in machine learning. *The Computer Journal*, 50(2):151–163, 2007.
- Isaac Gibbs and Emmanuel Candes. Adaptive conformal inference under distribution shift. *Advances in Neural Information Processing Systems*, 34:1660–1672, 2021.
- Etash Guha, Shlok Natarajan, Thomas Möllenhoff, Mohammad Emtiyaz Khan, and Eugene Ndiaye. Conformal prediction via regression-as-classification. In *NeurIPS 2023 Workshop on Regulatable ML*, 2023.
- Jianguo Huang, HuaJun Xi, Linjun Zhang, Huaxiu Yao, Yue Qiu, and Hongxin Wei. Conformal prediction for deep classifier via label ranking. In *Forty-first International Conference on Machine Learning*, 2024.
- Kexin Huang, Ying Jin, Emmanuel Candes, and Jure Leskovec. Uncertainty quantification over graph with conformalized graph neural networks. *Advances in Neural Information Processing Systems*, 36, 2023.
- Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551*, 2017.
- Kevin Kasa and Graham W Taylor. Empirically validating conformal prediction on modern vision architectures under distribution shift and long-tailed data. In *ICML 2023 Workshop on Structured Probabilistic Inference $\{\backslash\mathcal{E}\}$ Generative Modeling*, 2023.

- Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *5th International Conference on Learning Representations*, 2017.
- Danijel Kivaranovic, Kory D Johnson, and Hannes Leeb. Adaptive, distribution-free prediction intervals for deep networks. In *International Conference on Artificial Intelligence and Statistics*, pages 4346–4356. PMLR, 2020.
- Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010.
- Henrik Linusson, Isak Samsten, Zygmunt Zając, and Martín Villanueva. nonconformist: Python implementation of the conformal prediction framework, 2020.
- Kangdao Liu, Hao Zeng, Jianguo Huang, Huiping Zhuang, Chi-Man Vong, and Hongxin Wei. C-adapter: Adapting deep classifiers for efficient conformal prediction sets. *European Conference on Artificial Intelligence*, 2025.
- Tuve Löfström, Henrik Boström, Henrik Linusson, and Ulf Johansson. Bias reduction through conditional conformal prediction. *Intelligent Data Analysis*, 2015.
- Valery Manokhin. Awesome conformal prediction, April 2022.
- Andrew Kachites McCallum, Kamal Nigam, Jason Rennie, and Kristie Seymore. Automating the construction of internet portals with machine learning. *Information Retrieval*, 3: 127–163, 2000.
- William B McNatt and James M Bieman. Coupling of design patterns: Common practices and their benefits. In *25th Annual International Computer Software and Applications Conference. COMPSAC 2001*, pages 574–579. IEEE, 2001.
- Mouhcine Mendil, Luca Mossina, and David Vigouroux. Puncc: a python library for predictive uncertainty calibration and conformalization. In *Conformal and Probabilistic Prediction with Applications*, pages 582–601. PMLR, 2023.
- Harris Papadopoulos. Guaranteed coverage prediction intervals with gaussian process regression. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12):9072–9083, 2024.
- Harris Papadopoulos and Haris Haralambous. Reliable prediction intervals with regression neural networks. *Neural Networks*, 24(8):842–851, 2011.
- Harris Papadopoulos, Kostas Proedrou, Volodya Vovk, and Alex Gammerman. Inductive confidence machines for regression. In *Machine Learning: ECML 2002: 13th European Conference on Machine Learning Helsinki, Finland, August 19–23, 2002 Proceedings 13*, pages 345–356. Springer, 2002.
- Harris Papadopoulos, Vladimir Vovk, and Alex Gammerman. Regression conformal prediction with nearest neighbours. *Journal of Artificial Intelligence Research*, 40:815–840, 2011.

- Victor Quach, Adam Fisch, Tal Schuster, Adam Yala, Jae Ho Sohn, Tommi S Jaakkola, and Regina Barzilay. Conformal language modeling. In *The Twelfth International Conference on Learning Representations*, 2024.
- Michael Redmond. Communities and Crime. UCI Machine Learning Repository, 2002.
- Yaniv Romano, Evan Patterson, and Emmanuel Candes. Conformalized quantile regression. *Advances in Neural Information Processing Systems*, 32, 2019.
- Yaniv Romano, Matteo Sesia, and Emmanuel Candes. Classification with valid and adaptive coverage. *Advances in Neural Information Processing Systems*, 33:3581–3591, 2020.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- Mauricio Sadinle, Jing Lei, and Larry Wasserman. Least ambiguous set-valued classifiers with bounded error levels. *Journal of the American Statistical Association*, 114(525):223–234, 2019.
- Matteo Sesia and Emmanuel J Candès. A comparison of some conformal quantile regression methods. *Stat*, 9(1):e261, 2020.
- Yuanjie Shi, Subhankar Ghosh, Taha Belkhouja, Jana Doppa, and Yan Yan. Conformal prediction for class-wise coverage via augmented label rank calibration. *Advances in Neural Information Processing Systems*, 37:132133–132178, 2024.
- Jianqing Song, Jianguo Huang, Wenyu Jiang, Baoming Zhang, Shuangjie Li, and Chongjun Wang. Similarity-navigated conformal prediction for graph neural networks. *Advances in Neural Information Processing Systems*, 2024.
- David Stutz, Ali Taylan Cemgil, Arnaud Doucet, et al. Learning optimal conformal classifiers. In *The Tenth International Conference on Learning Representations*, 2022.
- Herbert Susmann, Antoine Chambaz, and Julie Josse. Adaptiveconformal: An r package for adaptive conformal inference. *arXiv preprint arXiv:2312.00448*, 2023.
- Ryan J Tibshirani, Rina Foygel Barber, Emmanuel Candes, and Aaditya Ramdas. Conformal prediction under covariate shift. *Advances in Neural Information Processing Systems*, 32, 2019.
- Vladimir Vovk. Conditional validity of inductive conformal predictors. In *Asian Conference on Machine Learning*, pages 475–490. PMLR, 2012.
- Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random World, Second Edition*. Springer International Publishing, January 2022. Publisher Copyright: © Springer Verlag New York, Inc. 2005.

- HuaJun Xi, Jianguo Huang, Kangdao Liu, Lei Feng, and Hongxin Wei. Does confidence calibration improve conformal prediction? *Transactions on Machine Learning Research*, 2025. ISSN 2835-8856.
- Chen Xu and Yao Xie. Conformal prediction interval for dynamic time-series. In *International Conference on Machine Learning*, pages 11559–11569. PMLR, 2021.
- Margaux Zaffran, Olivier Féron, Yannig Goude, Julie Josse, and Aymeric Dieuleveut. Adaptive conformal predictions for time series. In *International Conference on Machine Learning*, pages 25834–25866. PMLR, 2022.
- Soroush H Zargarbashi, Simone Antonelli, and Aleksandar Bojchevski. Conformal prediction sets for graph neural networks. In *International Conference on Machine Learning*, pages 12292–12318. PMLR, 2023.
- Shengjia Zhao, Willie Neiswanger, and Rachel Luo. Torchuc: A library for uncertainty quantification based on pytorch, 2024.

Appendix A. Details of Conformal Prediction Methods in TorchCP

In this section, we provide brief introductions to the methods implemented in TorchCP. Furthermore, Table 3 offers practical guidelines to help users select appropriate methods based on their specific application needs.

Component	Method	Advantage/Use case
DNN Classifier		
Score	LAC	To achieve the smallest prediction sets
	APS	To approximate conditional coverage
	RAPS	For smaller, more stable sets than APS
	TOPK	A simple baseline with only logit ranking
	SAPS	Improving efficiency, conveying instance-wise uncertainty
	Margin	A non-conformity function based on prediction margin
	KNN	A score function based solely on feature-space nearest neighbors without a predictive model.
Predictor	Class-conditional CP	To achieve coverage for each class
	Clustered CP	For class-conditional coverage with limited data per class
	Weighted CP	Conformal prediction under covariate shift
Loss & Trainer	ConfTr	Efficient prediction sets with conformal-aware loss
	C-Adapter	Optimizing efficiency based on α -free loss
	ConfTS	Improving efficiency via temperature scaling
DNN Regressor		
Score	ABS	Simple and interpretable residual-based scores
	NORABS	Enhancing Prediction Interval Robustness through Normalized Scoring
	CQR	Flexible interval width via quantile regression
	R2CCP	Shorter prediction intervals than CQR via recasting regression as classification
Predictor	ACI	Reliable coverage under distributional shifts
	AgACI	Aggregates ACI predictions with adaptive weighting
	Ensemble CP	Stable prediction intervals without data splitting via model ensembling
GNN		
Score	DAPS	Improves efficiency of APS using structural neighbors
	SNAPS	Extends DAPS with feature neighbors for better efficiency
	NAPS	Constructing prediction sets for inductive node classification via Weighted CP
Loss & Trainer	CF-GNN	Reducing set size by attaching an adapter with a topology-aware inefficiency loss
LLM		
Predictor	Conformal LLM	To achieve statistically guaranteed prediction sets for language models through calibrated sampling

Table 3: Advantages or use cases of CP methods across different tasks.

A.1 Split CP

Conformal prediction (Papadopoulos et al., 2002; Balasubramanian et al., 2014; Manokhin, 2022; Angelopoulos et al., 2023, 2024) is a statistical framework that generates prediction sets/intervals containing ground-truth labels with a desired probability. *Split Conformal Prediction* (Vovk et al., 2022) (Split CP) is the most widely-used version of the conformal prediction procedure. Specifically, Split CP divides a given dataset into two disjoint subsets: one for training the base model and the other for conformal calibration. We next outline the main process of split conformal prediction:

1. Divide a given dataset into two disjoint subsets: a training fold \mathcal{D}_{tr} and a calibration fold \mathcal{D}_{cal} , with $|\mathcal{D}_{cal}|$ being n ;
2. Train a model on the training dataset \mathcal{D}_{tr} , and then define a non-conformity score function $V(\mathbf{x}, y)$;
3. Compute $\hat{Q}_{1-\alpha}$ as the $\frac{\lceil (n+1)(1-\alpha) \rceil}{n}$ quantile of the calibration scores $\{V(\mathbf{x}_i, y_i) : (\mathbf{x}_i, y_i) \in \mathcal{D}_{cal}\}$, where $\hat{Q}_{1-\alpha}$ is defined by

$$\hat{Q}_{1-\alpha} := \inf\{Q \in \mathbb{R} : \frac{|\{i : V(\mathbf{x}_i, y_i) \leq Q\}|}{n} \geq \frac{\lceil (n+1)(1-\alpha) \rceil}{n}\};$$

4. Use the threshold $\hat{Q}_{1-\alpha}$ to generate a prediction set for a new instance \mathbf{x}_{n+1} :

$$\mathcal{C}(\mathbf{x}_{n+1}) = \{y : V(\mathbf{x}_{n+1}, y) \leq \hat{Q}_{1-\alpha}\}.$$

A.2 DNN Classifier

Here, we provide a brief overview of conformal prediction methods implemented using a DNN classifier with the TorchCP library. Let $\mathcal{X} \subset \mathbb{R}^d$ be the input space and $\mathcal{Y} := \{1, \dots, K\}$ be the label space. We use $(X, Y) \sim \mathcal{P}_{\mathcal{X}\mathcal{Y}}$ to denote a random data pair satisfying a joint data distribution $\mathcal{P}_{\mathcal{X}\mathcal{Y}}$ and $f : \mathcal{X} \rightarrow \mathbb{R}^K$ to denote a classification neural network. Thus, the classifier $\pi : \mathcal{X} \rightarrow \Delta^{K-1}$ is defined as $\sigma \circ f$, where Δ^{K-1} is a $(K-1)$ -dimensional probability simplex and σ is a normalization function such as the softmax function. Let $\pi_y(x)$ be the probability of y conditional on x , i.e., $P(y|x)$.

LAC (Sadinle et al., 2019): The LAC measures the similarity between the example and the data space by the conditional probability $P(Y|X)$: $V(\mathbf{x}, y) := 1 - \pi_y(\mathbf{x})$.

APS (Romano et al., 2020): The score of Adaptive Prediction Sets (APS) calculates the non-conformity score using the cumulative sum of sorted softmax probabilities:

$$V(\mathbf{x}, y, u) := \sum_{y' \in \mathcal{Y}} \pi_{y'}(\mathbf{x}) \mathbb{1}(r_f(\mathbf{x}, y') < r_f(\mathbf{x}, y)) + u \cdot \pi_y(\mathbf{x}),$$

where $r_f(\mathbf{x}, y)$ denotes the rank of $\pi_y(\mathbf{x})$ among the descending softmax probabilities, and u is an independent random variable satisfying a uniform distribution on $[0, 1]$.

RAPS (Angelopoulos et al., 2020): RAPS adds a regularization to penalize noisy tail probabilities and regularize the number of samples in the uncertainty set:

$$V(\mathbf{x}, y, u; \pi) := \sum_{y' \in \mathcal{Y}} \pi_{y'}(\mathbf{x}) \mathbb{1}(r_f(\mathbf{x}, y') < r_f(\mathbf{x}, y)) + u \cdot \pi_y(\mathbf{x}) + \lambda \cdot (r_f(\mathbf{x}, y) - k)^+,$$

where λ represents the weight of regularization, $k \geq 0$ are regularization hyper-parameters and $(z)^+$ denotes the positive part of z .

TOPK (Angelopoulos et al., 2020): TOPK is a score function based solely on the ranking information: $V(\mathbf{x}, y, u) := \sum_{y' \in \mathcal{Y}} \mathbb{1}(r_f(\mathbf{x}, y') < r_f(\mathbf{x}, y)) + u$.

SAPS (Huang et al., 2024): The SAPS method mitigates the issue of probability miscalibration by only retaining the highest probability value and discarding all others:

$$V(\mathbf{x}, y, u) := \begin{cases} u \cdot \pi_{\max}(\mathbf{x}), & \text{if } r_f(\mathbf{x}, y) = 1, \\ \pi_{\max}(\mathbf{x}) + (r_f(\mathbf{x}, y) - 2 + u) \cdot \lambda, & \text{else,} \end{cases}$$

where λ is a hyperparameter representing the weight of ranking information and $\pi_{\max}(\mathbf{x})$ denotes the maximum softmax probability.

KNN (Gammerman and Vovk, 2007): KNN calculates the score based on the k closest training examples in the feature space. The score is calculated by dividing two sums: the numerator is the sum of the k shortest distances from an input point to other points that share its same label, while the denominator is the sum of the k shortest distances from that input point to points that have different labels than it does.

Margin (Löffström et al., 2015): The margin score is constructed on a well-known concept that a prediction with a larger margin is generally considered to be more conforming. Consequently, the margin score is defined by $V(\mathbf{x}, y) := \pi_y - \max_{j \in \mathcal{Y}: j \neq y} \pi_j$.

Class-conditional CP (Vovk, 2012): Class-conditional CP partitions the calibration data based on their labels y . The non-conformity scores' quantiles are then computed separately for each class using the corresponding calibration subset.

Clustered CP (Ding et al., 2024): Clustered CP extends Class-conditional CP by applying split conformal prediction to class clusters rather than individual classes. It achieves cluster-conditional coverage as an approximation of class-conditional guarantees.

Weighted CP (Tibshirani et al., 2019): Weighted CP addresses distribution shift in Conformal Prediction, where calibration and test data follow different distributions. The core idea is that data points closer to the test point in the feature space provide more valuable information about it.

RC3P (Shi et al., 2024): RC3P addresses the inefficiency of class-conditional conformal prediction, which often results in large prediction sets, especially for imbalanced classification tasks. The core idea is to selectively perform class-wise thresholding only for labels with low top- k error, by integrating a label rank calibration step into standard class-wise conformal calibration.

ConfTr (Stutz et al., 2022): ConfTr simulates the CP process based on end-to-end training by splitting mini-batches into calibration and prediction sets and designing differentiable conformal prediction steps: $\mathcal{L}_{\text{ConfTr}} = \max \left(0, \sum_{k=1}^K \mathcal{C}_{\theta, k}(\mathbf{x}) - \kappa \right)$, where θ is the trainable parameters of deep learning models and κ is the target set size. The definition of $\mathcal{C}_{\theta, y}(\mathbf{x})$ is given by $\sigma((\hat{Q}_{1-\alpha} - V(\mathbf{x}, y))/\tau)$ where τ is a temperature hyperparameter and σ is the sigmoid function.

ConfTS (Xi et al., 2025): ConfTS is a training method to optimize the T of temperature scaling for improved efficiency of prediction sets. The objective of ConfTS is to minimize the efficiency gap, i.e., $\hat{Q}_{1-\alpha} - V(\mathbf{x}, y)$. Formally, the loss function of ConfTS is:

$$\mathcal{L}_{\text{ConfTS}} = (\hat{Q}_{1-\alpha} - V(\mathbf{x}, y))^2$$

C-Adapter (Liu et al., 2025): C-Adapter proposes to optimize an adapter with a Conformal Discriminative (CD) Loss. The CD loss is constructed to improve the α -free efficiency of prediction sets: $\mathcal{L}_{\text{CD}} = \frac{1}{K} \sum_{j \in \mathcal{Y}} \{\sigma(V(\mathbf{x}, y) - V(\mathbf{x}, j))\}$.

SCPO (Bellotti, 2021): SCPO proposes an approach to train the deep learning model directly by optimizing for maximum predictive efficiency. To enable end-to-end training, the conformal predictor is approximated by a differentiable surrogate objective. Gradient descent is then used to minimize the following loss:

$$\mathcal{L}_{\text{SCPO}} = \sum_{k=1}^K \mathcal{C}_{\theta,k}(\mathbf{x}) + \lambda (\mathcal{C}_{\theta,y}(\mathbf{x}) - (1 - \alpha))^2,$$

where $\lambda > 0$ controls the relative importance of the two objective components.

Uncertainty-aware (UA) Loss (Einbinder et al., 2022): Uncertainty-aware Loss introduces a novel training objective that integrates conformal calibration directly into model training by encouraging the conformity scores to follow a uniform distribution. The novel uncertainty-aware component of this loss is:

$$\mathcal{L}_u = \sup_{w \in [0,1]} |\mathbb{E}[\mathbb{1}(V(\mathbf{x}, y, u) - w)]|,$$

where $V(\mathbf{x}, y, u) := \sum_{y' \in \mathcal{Y}} \pi_{y'}(\mathbf{x}) \mathbb{1}(r_f(\mathbf{x}, y') < r_f(\mathbf{x}, y)) + u \cdot \pi_y(\mathbf{x})$.

A.3 DNN Regressor

For DNN Regressor, the goal is to predict a continuous target variable $y \in \mathbb{R}$ based on an input $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^d$, where $(X, Y) \sim \mathcal{P}_{\mathcal{X}\mathcal{Y}}$ follows a joint data distribution $\mathcal{P}_{\mathcal{X}\mathcal{Y}}$. Let $f : \mathcal{X} \rightarrow \mathcal{F}$ denote a regression model trained on a dataset $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$, where \mathcal{F} represents the space of model outputs, \mathbf{x}_i are the features and y_i are the observed target values.

ABS (Papadopoulos et al., 2002): The ABS score $V(\mathbf{x}, y) := |f(\mathbf{x}) - y|$ is designed to measure the absolute deviation of a prediction $f(\mathbf{x})$ from the observed target value y , where $f(\mathbf{x})$ is a point estimate of given $\{\mathbf{x}\}$. Then the prediction set for a new instance \mathbf{x}_{n+1} is constructed by $\mathcal{C}(\mathbf{x}_{n+1}) = [f(\mathbf{x}_{n+1}) - \hat{Q}_{1-\alpha}, f(\mathbf{x}_{n+1}) + \hat{Q}_{1-\alpha}]$.

NORABS (Papadopoulos et al., 2011; Papadopoulos and Haralambous, 2011): The NORABS score $V(\mathbf{x}, y) := \frac{|f(\mathbf{x}) - y|}{\gamma + \xi(\mathbf{x})}$ normalizes the regression error using difficulty measure deviation $\xi(\mathbf{x})$. The prediction set for \mathbf{x}_{n+1} is $\mathcal{C}(\mathbf{x}_{n+1}) = [f(\mathbf{x}_{n+1}) - \hat{Q}_{1-\alpha} \cdot (\gamma + \xi(\mathbf{x}_{n+1})), f(\mathbf{x}_{n+1}) + \hat{Q}_{1-\alpha} \cdot (\gamma + \xi(\mathbf{x}_{n+1}))]$. Additionally, we provide the normalized score for Gaussian Process Regression (Papadopoulos, 2024).

CQR (Romano et al., 2019): The CQR score $V(\mathbf{x}, y) := \max(f(\mathbf{x})_{\alpha/2} - y, y - f(\mathbf{x})_{1-\alpha/2})$ measures the maximum absolute deviation between the true value y and the predicted quantile interval $[f(\mathbf{x})_{\alpha/2}, f(\mathbf{x})_{1-\alpha/2}]$. Then the prediction set for a new instance \mathbf{x}_{n+1} is constructed by: $\mathcal{C}(\mathbf{x}_{n+1}) = [f(\mathbf{x})_{\alpha/2} - \hat{Q}_{1-\alpha}, f(\mathbf{x})_{1-\alpha/2} + \hat{Q}_{1-\alpha}]$.

CQRR (Sesia and Candès, 2020): The CQRR score extends the standard CQR method by introducing a scaling mechanism. This score normalizes the deviations between the true value y and the predicted quantile interval $[f(\mathbf{x})_{\alpha/2}, f(\mathbf{x})_{1-\alpha/2}]$ by the width of the

interval itself. Formally, for a data pair (\mathbf{x}, y) , the CQRR score is defined as $V(\mathbf{x}, y) := \max\left(\frac{f(\mathbf{x})_{\alpha/2} - y}{\hat{s}_\alpha(\mathbf{x})}, \frac{y - f(\mathbf{x})_{1-\alpha/2}}{\hat{s}_\alpha(\mathbf{x})}\right)$, where $\hat{s}_\alpha(\mathbf{x}) := f(\mathbf{x})_{1-\alpha/2} - f(\mathbf{x})_{\alpha/2} + \epsilon$ is the scaling factor, and $\epsilon > 0$ is a small constant added to prevent division by zero. The prediction set is given by $\mathcal{C}(\mathbf{x}_{n+1}) = \left[f(\mathbf{x}_{n+1})_{\alpha/2} - \hat{Q}_{1-\alpha} \cdot \hat{s}_\alpha(\mathbf{x}_{n+1}), f(\mathbf{x}_{n+1})_{1-\alpha/2} + \hat{Q}_{1-\alpha} \cdot \hat{s}_\alpha(\mathbf{x}_{n+1})\right]$.

CQRM (Sesia and Candès, 2020): The CQRM score extends the CQRR score by introducing two scaling factors to independently normalize the deviations for the lower and upper bounds of the prediction interval. This method requires the model to output predictions at three quantile levels: $[\alpha/2, 1/2, 1 - \alpha/2]$. The two scaling factors are defined as $\hat{s}_{\alpha, \text{lo}}(\mathbf{x}) := f(\mathbf{x})_{1/2} - f(\mathbf{x})_{\alpha/2} + \epsilon$ and $\hat{s}_{\alpha, \text{up}}(\mathbf{x}) := f(\mathbf{x})_{1-\alpha/2} - f(\mathbf{x})_{1/2} + \epsilon$. The non-conformity score for a data pair (\mathbf{x}, y) and prediction interval are then defined as:

$$V(\mathbf{x}, y) := \max\left(\frac{f(\mathbf{x})_{\alpha/2} - y}{\hat{s}_{\alpha, \text{lo}}(\mathbf{x})}, \frac{y - f(\mathbf{x})_{1-\alpha/2}}{\hat{s}_{\alpha, \text{up}}(\mathbf{x})}\right).$$

Then, the prediction set is given by

$$\mathcal{C}(\mathbf{x}_{n+1}) = \left[f(\mathbf{x}_{n+1})_{\alpha/2} - \hat{Q}_{1-\alpha} \cdot \hat{s}_{\alpha, \text{lo}}(\mathbf{x}_{n+1}), f(\mathbf{x}_{n+1})_{1-\alpha/2} + \hat{Q}_{1-\alpha} \cdot \hat{s}_{\alpha, \text{up}}(\mathbf{x}_{n+1})\right].$$

CQRFM (Kivaranovic et al., 2020): The CQRFM score is similar to CQRM, as it also requires the model to output predicted quantiles at $[\alpha/2, 1/2, 1 - \alpha/2]$. However, it normalizes the deviations from the median $f(\mathbf{x})_{1/2}$ by the fractional distance to the nearest quantile bounds. The two fractional scaling factors are defined in the same way as in CQRM. The non-conformity score and prediction interval are defined as:

$$V(\mathbf{x}, y) := \max\left(\frac{f(\mathbf{x})_{1/2} - y}{\hat{s}_{\alpha, \text{lo}}(\mathbf{x})}, \frac{y - f(\mathbf{x})_{1/2}}{\hat{s}_{\alpha, \text{up}}(\mathbf{x})}\right).$$

Then, the prediction set is given by

$$\mathcal{C}(\mathbf{x}_{n+1}) = \left[f(\mathbf{x}_{n+1})_{1/2} - \hat{Q}_{1-\alpha} \cdot \hat{s}_{\alpha, \text{lo}}(\mathbf{x}_{n+1}), f(\mathbf{x}_{n+1})_{1/2} + \hat{Q}_{1-\alpha} \cdot \hat{s}_{\alpha, \text{up}}(\mathbf{x}_{n+1})\right].$$

R2CCP (Guha et al., 2023): R2CCP converts a regression problem into a classification problem by dividing the range space into K bins. The model predicts the probability that each sample belongs to one of these K bins, similar to a standard classification task. To train the regression-to-classification model, R2CCP uses a custom loss function:

$$\mathcal{L}(\theta) = \sum_{i=1}^N \left(\sum_{k=1}^K (|y_i - m_k|^p \cdot \pi_{ik}) - \tau \cdot \sum_{k=1}^K \pi_{ik} \log(\pi_{ik}) \right),$$

where π_{ik} is the predicted probability for the k -th bin, m_k are the bin midpoints, τ is a regularization parameter, and p is the exponent applied to the absolute difference between the target value and the bin midpoint. Both τ and p are hyperparameters. After training the regression-to-classification model, R2CCP defines the non-conformity score as the linear interpolation of the model's softmax probabilities. Based on these scores, R2CCP outputs the interval where all non-conformity scores are smaller than $\hat{Q}_{1-\alpha}$.

Ensemble CP (Xu and Xie, 2021): Ensemble CP trains multiple models based on randomly sampled subsets of the data and aggregates their non-conformity scores using

a specific aggregation function, such as mean or median. It can update $\hat{Q}_{1-\alpha}$ over time, making it adaptable to dynamic time series.

ACI (Gibbs and Candes, 2021): ACI dynamically adjusts the confidence level $\alpha_t = \alpha_{t-1} + \gamma(\alpha - \text{err}_t)$ based on the historical error rates to adapt to distribution shifts, where α_t is the updated confidence level depending on time t , α is the target coverage, and $\text{err}_t := \sum_{i=1}^n w_i \cdot I(y_i \notin \hat{C}_i)$ is the weighted error rate at time step t . This dynamic updating enables ACI to maintain reliable prediction sets despite evolving data distributions.

AgACI (Zaffran et al., 2022): AgACI is a parameter-free conformal prediction method for time series that builds on ACI by using online expert aggregation. It runs multiple models in parallel with different learning rates γ_k and adaptively aggregates their prediction intervals using weighted averages, where weights are based on the historical pin-ball loss. This approach eliminates the need to tune γ manually and ensures valid coverage while improving interval efficiency, especially under temporal dependencies.

A.4 GNN

In this section, we discuss the applications of conformal prediction algorithms on graph data. The Graph is defined as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} := \{v_i\}_{i=1}^N$ denotes the node set and \mathcal{E} denotes the edge set with $|\mathcal{E}| = E$. Let $\mathbf{X} := [\mathbf{x}_1, \dots, \mathbf{x}_N]^T$ be the node feature matrix, where $\mathbf{x}_i \in \mathbb{R}^d$ is a d -dimensional feature vector for node v_i . The label of v_i is $y_i \in \mathcal{Y}$, where $\mathcal{Y} := \{1, 2, \dots, K\}$ denotes the label space.

DAPS (Zargarbashi et al., 2023): DAPS proposes a diffusion-based method that incorporates neighborhood information by leveraging the network homophily. Specifically, the score function is defined as: $V(\mathbf{x}_i, y, \lambda) := (1 - \lambda)\hat{V}(\mathbf{x}_i, y) + \frac{\lambda}{|\mathcal{N}_i|} \sum_{v_j \in \mathcal{N}_i} \hat{V}(\mathbf{x}_j, y)$, where $\hat{V}(\cdot, \cdot)$ is the basic non-conformity score function, e.g., APS, and λ is a diffusion parameter.

SNAPS (Song et al., 2024): SNAPS, which is essentially an enhanced version of the DAPS, aggregates non-conformity scores of nodes with high feature similarity to ego node and one-hop graph structural neighbors. Formally, the score function is shown as: $V(\mathbf{x}_i, y, \lambda, \mu) := (1 - \lambda - \mu)\hat{V}(\mathbf{x}_i, y) + \frac{\lambda}{\mathbf{D}_s(i, i)} \sum_{j=1}^M \mathbf{A}_s(i, j)\hat{V}(\mathbf{x}_j, y) + \frac{\mu}{|\mathcal{N}_i|} \sum_{v_j \in \mathcal{N}_i} \hat{V}(\mathbf{x}_j, y)$, where \mathbf{A}_s is the adjacency matrix of k -NN graph, which measures the similarity between each node pair, and \mathbf{D}_s is the degree matrix of \mathbf{A}_s . M denotes the number of randomly selected nodes as targets used to construct k -NN graph. Both λ and μ are hyperparameters, which are used to measure the importance of three parts of non-conformity scores.

NAPS (Clarkson, 2023): NAPS, which adapts the weighted variant of CP from (Barber et al., 2023), assigns a weight to adjacent nodes, and a weight of zero otherwise. NAPS then applies nonexchangeable conformal prediction with the APS scoring function in the following equation: $\mathcal{C}(\mathbf{x}_{n+1}) = \{y \in \mathcal{Y} : V(\mathbf{x}_{n+1}, y) \leq Q_{1-\alpha}(\sum_{i=1}^n w_i \cdot \delta_{s_i} + w_{n+1} \cdot \delta_{+\infty})\}$, where $Q_\tau(\cdot)$ denotes the τ -quantile of a distribution and $s_i = V(\mathbf{x}_i, y_i)$. Moreover, w_i denotes a weight of s_i and δ_{s_i} denotes a point mass at s_i .

CF-GNN (Huang et al., 2023): CF-GNN introduces a topology-aware output correction model, akin to Graph Convolutional Network, which employs a conformal-aware inefficiency loss to refine predictions and improve the efficiency of post-hoc CP. For the loss, CF-GNN adopts the optimizing efficiency version of ConfTr.

Deep Models	Metrics
DNN Classifier	Coverage, Size, CovGap (Ding et al., 2024), WSC (Romano et al., 2020), SSCV (Angelopoulos et al., 2020), Violated classes (Kasa and Taylor, 2023), DiffViolation (Angelopoulos et al., 2020), S Criterion, N Criterion, U Criterion, F Criterion, M Criterion, F Criterion, OU Criterion, OF Criterion, OM Criterion, OE Criterion (Vovk et al., 2022), etc.
DNN Regressor	Coverage, Size
GNN	Coverage, Size, Singleton ratio (Zargarbashi et al., 2023)
LLM	SetLoss, SSCL (Quach et al., 2024)

Table 4: Implemented evaluation metrics in TorchCP.

A.5 LLM

Conformal LLM (Quach et al., 2024): This paper proposes a novel conformal procedure based on a risk control framework (Bates et al., 2021) for language models by reformulating the generation task as a sampling problem rather than enumerating all possible outputs. The core algorithm samples candidate outputs until a calibrated stopping rule is met, while using a rejection rule to filter out low-quality samples. The method guarantees that the output set contains at least one acceptable answer with high probability while maintaining the statistical guarantees of traditional conformal prediction.

A.6 Metrics

In this section, we provide a brief overview of the evaluation metrics supported by TorchCP for diverse modern models, as summarized in Table 4. These metrics facilitate rigorous and fine-grained assessment of prediction sets, offering deep insights into model uncertainty.

Appendix B. Extensibility

Each module in TorchCP can be easily extended. In the following, we provide basic guidelines for customizing your own scores, predictors, trainers, and metrics.

B.1 Add new scores

Adding new scores involves two steps:

1. Implementing a New Score Class: new Scores should be implemented in the `score` module, inheriting from the `BaseScore` class. Using classification as an example, one could alternatively choose to override the functions: `__init__()`, `__call__()`, `_calculate_all_label()`, and `_calculate_single_label()`.
2. Adding an Interface: After customizing the score class, register it in `score/__init__.py`.

B.2 Add new predictors

Similar to adding new scores, the addition of new predictors also consists of two steps:

1. Implementing a New Predictor Class: new Predictors should be implemented in the `predictor` module, inheriting from the `BasePredictor` class. Using classification as an example, one could alternatively choose to override the functions: `__init__()`, `calibrate()`, and `predict()`.
2. Adding an Interface: After customizing the predictor class, register it in `predictor/__init__.py`.

B.3 Add new trainers

Adding new trainers is similar to steps of B.1 and B.2:

1. Implementing a New Trainer Class: new Trainers should be implemented in the `trainer` module, inheriting from the `BaseTrainer` class. Using classification as an example, one could alternatively choose to override the functions: `__init__()`, `train()` and `validate()`. One could pass different loss functions in the `__init__()` method.
2. Adding an Interface: After customizing the trainer class, register it in `trainer/__init__.py`.

B.4 Add new metrics

Adding new metrics involves two steps:

1. Implementing a New Metric Function: New metrics should be implemented in the `utils/metrics.py`.
2. Adding an Interface: After customizing the metric function, register it by `Registry('METRICS')` in the `utils/metrics.py`.

Appendix C. Example Codes for TorchCP

TorchCP provides a collection of example codes in the “examples” folder ⁷ of the official codebase, designed to help users quickly get started with various conformal prediction tasks. In this section, we briefly introduce representative examples for the four supported tasks.

C.1 Example codes for offline conformal prediction

This part provides brief descriptions of example codes under the offline setting:

1. **DNN Classifier.** This example code⁸ applies Split CP with the LAC score on CIFAR-100 using a pretrained ResNet-20. It evaluates the coverage and average set size of the resulting prediction sets.

7. <https://github.com/ml-stat-Sustech/TorchCP/tree/master/examples>

8. https://github.com/ml-stat-Sustech/TorchCP/blob/master/examples/classification_splitcp_cifar100.py

2. **DNN Regressor.** The example code of regression⁹ shows the complete workflow of applying Split CP with the Conformalized Quantile Regression score (Romano et al., 2019) on a synthetic dataset in a regression task, including model training, calibration, prediction interval generation, and evaluation.
3. **GNN.** The GNN example¹⁰ implements a Graph Convolutional Network (GCN) (Kipf and Welling, 2017) for node classification on the CoraML dataset (McCallum et al., 2000) under a transductive setting. Conformal prediction is then applied to generate prediction sets with guaranteed coverage.
4. **LLM.** The LLM example¹¹ demonstrates how to deploy conformal prediction in large language model generation tasks, i.e., TriviaQA (Joshi et al., 2017).

C.2 Example codes for semi-online conformal prediction

Semi-online conformal prediction means that after predicting a test point, its true label is immediately revealed, and the test sample is incorporated into the calibration set. TorchCP supports this setting as well, with the following examples:

1. **DNN Regressor.** In TorchCP, the semi-online setup of the regression task is integrated into `ACIPredictor` (Gibbs and Candes, 2021). Therefore, we added an example code¹² for the regression task in TorchCP.
2. **DNN Classifier.** For DNN classification tasks, this semi-online calibration approach can be naturally implemented in TorchCP by iteratively invoking the “calibrate()” function with new incoming data. To clarify this, we have added a semi-online conformal prediction example¹³ for a classification task in TorchCP.

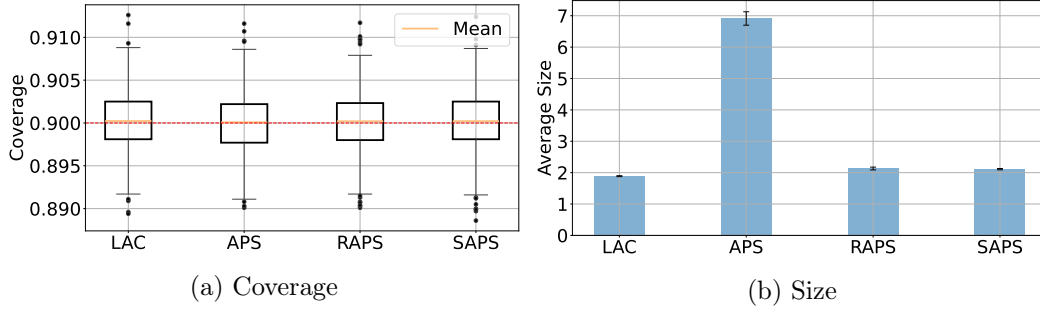
Appendix D. Empirical Results

This section presents empirical results of TorchCP. We first reproduce several classical conformal prediction methods across various tasks. Then, we benchmark the computational efficiency of TorchCP against several representative conformal prediction libraries.


D.1 Evaluation results

DNN Classifier. In Figure 2, we present the experimental results of Split CP on ImageNet, using LAC, APS, RAPS ($\lambda=0.1$, $k=0$), and SAPS ($\lambda=0.2$). Figure 3 compares Split

9. https://github.com/ml-stat-Sustech/TorchCP/blob/master/examples/regression_cqr_synthetic.py
10. https://github.com/ml-stat-Sustech/TorchCP/blob/master/examples/gnn_transductive_coraml.py
11. https://github.com/ml-stat-Sustech/TorchCP/blob/master/examples/llm_ConformalLM_TriviaQA.py
12. https://github.com/ml-stat-Sustech/TorchCP/blob/master/examples/timeseries_aci_synthetic.py
13. https://github.com/ml-stat-Sustech/TorchCP/blob/master/examples/classification_splitcp_cifar100_online.py

Figure 2: Results of image classification with $\alpha = 0.1$ on ImageNet dataset.

CP, Class-conditional CP, and Clustered CP on the ImageNet dataset, each employing the APS scoring rule. Results are averaged over 1,000 independent trials.

Figure 3: Conditional coverage results on ImageNet image classification with $\alpha = 0.1$.

DNN Regressor. In Figure 4, we present an example of experimental results on the Community dataset (Redmond, 2002), using ABS, CQR, CQRR, CQRM, CQRFM, R2CCP with Split CP. The number of trials is 1000.


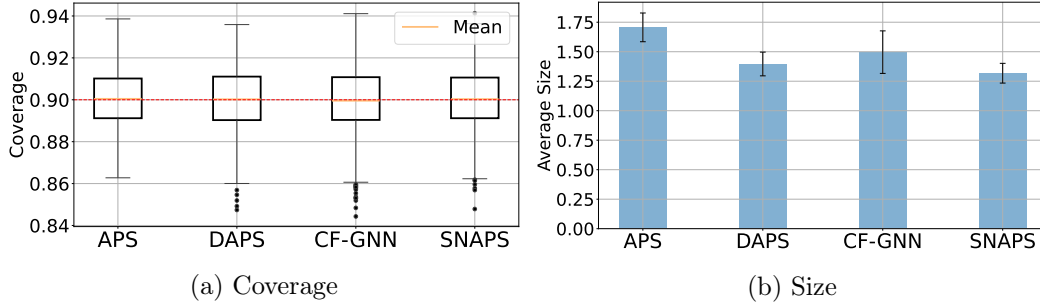
GNN. In Figure 5, we present experimental results on the CoraML dataset (McCallum et al., 2000) for node classification under a transductive learning setting. We evaluate four methods: APS, DAPS, CF-GNN, and SNAPS, each conducted over 1000 independent trials. For DAPS, the diffusion parameter is set to 0.5. In CF-GNN, the topology-aware correction model is implemented using a two-layer GCN. For SNAPS, the hyperparameters are set as follows: $\lambda = \frac{1}{3}$, $\mu = \frac{1}{3}$, and $k = 20$. These hyperparameter choices follow the default or recommended configurations reported in the respective original papers.

LLM. In Table 5, we present an example prediction set, generated by conformal LLM, of the TriviaQA dataset (Joshi et al., 2017). The generation process terminates after the fifth sample, as the cumulative score surpasses the conformal threshold.

D.2 Comparison of computational efficiency

D.2.1 CLASSIFICATION TASK

We evaluated the computational efficiency of TorchCP in image classification tasks by comparing its calibration and inference time with several existing CP libraries. Specifically, we use two standard benchmarks: MNIST (LeCun et al., 2010) and ImageNet (Russakovsky et al., 2015). Throughout all experiments, we used Split CP with APS as the conformal


 Figure 4: Results of regression with $\alpha = 0.1$ on Community dataset.

 Figure 5: Results of Graph node classification with $\alpha = 0.1$ on CoraML dataset.

Question	Men Against the Sea and Pitcairn’s Island were two sequels to what famous novel?		
	Answer	Score	Label
Answer 1	Mutiny On The Bounty	0.5347	True
Answer 2	Treasure Island	0.5087	False
Answer 3	Treasure Island	0.5087	False
Answer 4	Mutiny on the Bounty	0.1889	True
Answer 5	Robinson Crusoe	0.1384	False

Table 5: An Example of prediction sets from Conformal LLM on TriviaQA.

prediction method. Moreover, **TorchCP** supports two data processing modes: full-batch (processing the entire dataset at once), and batch (processing data in mini-batches). In our experiments, we set the batch size to 128 for batch processing. We use NVIDIA H100 for GPU computation and AMD EPYC 9654 for CPU computation.

MNIST. We compared the computational efficiency of five libraries: **PUNCC**, **MAPIE**, **Fortuna**, **Crepes**, and **TorchCP** (GPU-based full-batch processing). For the dataset, we use 48,000 samples for training, 6,000 for calibration, and 6,000 for testing. For **TorchCP**, we use a LeNet-5 model for training and inference using PyTorch on the GPU. For **PUNCC**, **MAPIE**, and **Crepes**, we obtain model predictions using PyTorch on a GPU, and then convert output probabilities to NumPy arrays for conformal calibration and inference using CPU. For **Fortuna**, we train the LeNet-5 model using the built-in training algorithm and perform model predictions with JAX on a GPU. We only report the total runtime for **Fortuna** as




Figure 6: Runtime comparison of different conformal prediction libraries.

it ties the conformal calibration and inference processes together. Figure 6a presents the running times of the five libraries on the MNIST dataset. The results show that TorchCP achieves the lowest calibration and inference time among the five libraries. Specifically, TorchCP completes calibration and inference for the entire dataset in just 0.07 seconds, using only 3.6% of the runtime needed by PUNCC (1.93 seconds). Moreover, the Crepes library exhibits the longest inference time on the MNIST dataset, likely due to its computation of p-values for each test sample’s non-conformity scores by comparing them against the entire calibration set. This substantial improvement highlights the efficiency advantage of TorchCP on small-scale datasets.

ImageNet. We compared the computational efficiency of four libraries: PUNCC, MAPIE, Crepes, and TorchCP (GPU-based batch processing). In the dataset, we use a pre-trained ResNet-50 model from torchvision and select 25,000 images from the validation set for calibration and another 25,000 for testing. For TorchCP, we leverage GPU-based batch processing to efficiently run the conformal pipeline. For PUNCC, MAPIE, and Crepes, we conducted experiments on the CPU, as they lack batch processing support, and their GPU-based full-batch processing causes out-of-memory errors. Figure 6b shows the results of running times of three libraries on the ImageNet dataset. The results demonstrate that TorchCP achieves the shortest calibration and inference times, cutting approximately 90% of the time required by PUNCC (104s vs. 995s). This significant enhancement underscores the scalability of TorchCP on large-scale datasets and models.

Choices of batch processing. To assess the efficiency of two processing modes (Full-batch and batch), we conduct experiments with TorchCP on both MNIST and ImageNet, using the same data splits and models as in prior experiments. We utilize GPU-accelerated computing for all experiments except the full-batch mode on ImageNet, where the dataset’s size exceeds GPU memory capacity. Figures 7a and 7b present the running times of different data processing modes on MNIST and ImageNet, respectively. On the small-scale dataset, i.e., MNIST, using full-batch processing on GPU achieves the fastest implementation among the three choices, requiring only 26% computational time of that with CPU. When it comes to the ImageNet dataset, GPU-based batch processing completes calibration and inference in approximately 104.29 seconds, compared to 1022.55 seconds for CPU-based processing, achieving a 9.8 \times speedup. These findings suggest that GPU-based full-batch processing is optimal for small-scale datasets where GPU memory is sufficient to load the whole dataset.

Conversely, GPU-based batch processing is preferable for large-scale datasets like ImageNet, offering significant runtime advantages and enabling scalable, efficient deployment.




Figure 7: Runtime comparison of different data processing modes.




Figure 8: Comparison of time consumption and prediction intervals between AdaptiveConformal and TorchCP on the France electricity dataset.

D.2.2 REGRESSION TASK

We evaluate the computational efficiency and predictive performance of **TorchCP** and an R-based CP package, **AdaptiveConformal** (Susmann et al., 2023) on a time series regression task, using the France electricity forecasting dataset¹⁴. Specifically, we use the first 350 time steps to train a SARIMA model with order (1, 0, 0) and seasonal order (1, 1, 0), and construct prediction intervals for the remaining time steps. To assess computational efficiency, we apply the ACI algorithm (Gibbs and Candes, 2021) using both **TorchCP** and **AdaptiveConformal**. Both **AdaptiveConformal** and **TorchCP** are executed on CPU in this experiment to ensure a fair comparison. Moreover, since ACI performs recalibration for each test instance, the reported time includes both calibration and inference.

Figure 8a reports the runtime results of these two packages. The results show that **TorchCP** achieves significantly higher computational efficiency compared to the R-based package, reducing running time from 6.1 seconds to around 0.06 seconds. Figures 8b and 8c illustrate the prediction intervals produced by **AdaptiveConformal** and **TorchCP**, respectively. The two sets of intervals are nearly identical, demonstrating the correctness and consistency of the **TorchCP** implementation. This substantial improvement highlights the efficiency advantage of **TorchCP** for time series regression.

14. <https://cran.r-project.org/web/packages/opera/index.html>

Appendix E. The Advantage of our low-coupling design

Low-coupling design is one of the most fundamental concepts in object-oriented design, which minimizes the dependency a class has on other classes. The low-coupling design of TorchCP ensures flexibility in integrating diverse training algorithms, non-conformity score functions, and prediction workflows, accommodating the wide variety of CP methods and tasks. In contrast, most current packages focus on a high-coupling design where predictors and scoring functions are often predefined within a specific predictor (e.g., Split CP). This limits extensibility when new CP predictors or user-defined score functions are needed. TorchCP’s decoupled design allows users to independently select or define new trainers, scores, or predictors, enabling flexible integration of different components. In what follows, we elaborate on the concrete benefits of our low-coupling design:

- **Flexible combination:** users can freely combine Trainer, Score, and Predictor modules to customize conformal prediction workflows. For instance, a practitioner can use different non-conformity scores (e.g., LAC or APS) within any predictor or trainer, while other packages only implement specified scores for predictors. Moreover, some libraries (e.g., TorchUQ) implemented various non-conformity scores as functions, but this style suffers from extending to those scores with extra hyperparameters (e.g., the k and λ of RAPS).
- **Modification isolation:** our low-coupling design ensures that changes to one module (e.g., refining the APS score) do not affect other modules. This isolation minimizes the risk of unintended side effects, reducing maintenance complexity and preventing errors from propagating across the system. For instance, tightly coupled systems like PUNCC require redefining the predictor to modify a scoring function, while TorchCP confines changes to the target module, streamlining updates and preserving stability.
- **Research extensibility:** by adhering to standardized interfaces, new functionality (e.g., a novel non-conformity score) can be integrated into TorchCP without altering existing modules. Differently, MAPIE follows a high-coupling design, making it challenging to introduce new predictors, like class-conditional CP. Extensibility supports rapid iteration and long-term evolvability, enabling efficient adaptation to new research or applications.
- **Ease of testing and debugging:** decoupled modules can be tested and debugged independently, as their interactions are governed by well-defined interfaces. This modularity simplifies identifying and resolving issues within a single module (e.g., debugging a Trainer algorithm) without navigating complex dependencies. In TorchCP, this reduces testing overhead and accelerates development cycles compared to tightly coupled systems, where debugging often requires analyzing the entire pipeline.

Appendix F. Comparison with related libraries

Numerous libraries implement conformal prediction (CP) algorithms, yet their suitability for deep learning scenarios differs significantly, particularly in supporting modern deep learning models like GNNs and LLMs, as well as their scalability for large-scale datasets. In Table 2,

we present a concise comparison of open-source CP libraries, focusing on supported model types and functionalities:

- **Model type** refers to the types of machine learning models the library can work with. They are classified based on differences in input and output data, which influence the selection of applicable CP algorithms.
- **Functionality** describes the operational capabilities of the library, such as how it processes data, computational optimization, or prediction forms.

In the following, we compare several popular CP libraries, emphasizing TorchCP’s superior compatibility and performance with deep learning applications.

nonconformist (Linusson et al., 2020): A NumPy-based library supporting basic CP algorithms for classification and regression tasks. However, it lacks GPU acceleration, batch processing, and compatibility with deep learning models like DNN Classifiers, DNN Regressors, GNNs, or LLMs, making it unsuitable for complex deep learning scenarios. While **nonconformist** implements predictors and scores as different classes following a low-coupling design, its algorithmic coverage is limited, implementing only inductive CP and aggregated CP with 5 scores.

TorchUQ (Zhao et al., 2024): A PyTorch library focused on uncertainty quantification, including some CP methods for regression, with GPU support. However, it lacks comprehensive algorithm implementation, omitting CP algorithms for classifiers and CP-specific training, and does not support advanced deep learning models like GNNs or LLMs, limiting its applicability for complex deep learning Scenarios. **TorchUQ** uses a dictionary to map names to score functions in predictors, enabling developers to define new scores without modifying predictor logic. However, this style struggles to accommodate scores with additional hyperparameters (e.g., k and λ in RAPS), requiring extra logic to manage these variations, which increases coupling and complexity.

Crepes (Boström, 2022): a NumPy-based library with extensive prediction forms that supports the classifiers, regressors, and predictive systems with semi-online CP, excelling in statistical classification and regression tasks. However, it lacks GPU acceleration, batch processing, and seamless extensions to modern models like GNNs or LLMs, limiting scalability for engineers tackling large-scale or complex deep learning problems. Similar to **TorchUQ**, **Crepes** implements non-conformity scores as functions passed to predictors, so it also suffers from supporting scores with additional hyperparameters (e.g., k and λ in RAPS).

PUNCC (Mendil et al., 2023): A scikit-learn-compatible library delivers Conformal Prediction for tasks like regression, classification, and anomaly detection, but struggles with modern deep learning applications. It does not support advanced deep learning models such as DNNs, GNNs, or LLMs, lacks GPU acceleration and batch processing, and performs significantly slower on large datasets (up to 10x slower than **TorchCP** in ImageNet, see Figure 6b). **PUNCC** defines non-conformity scores as functions, but ties each score to a dedicated predictor, resulting in high coupling.

MAPIE (Cordier et al., 2023): A scikit-learn-based library providing CP sets/intervals for classical classification/regression tasks. It excels in traditional machine learning workflows but lacks support for deep learning models, such as DNNs, GNNs, or LLMs. It also does not offer online CP, GPU acceleration, and batch processing. **MAPIE** implements

scores as different classes with a `predict` function, making it challenging to support other predictors (like class-conditional CP (Vovk, 2012)).

Fortuna (Detommaso et al., 2024): a JAX-based library, supports CP methods in deep learning with GPU acceleration and select conformal prediction methods. It enhances scalability for deep learning but has limited algorithmic coverage, lacking training algorithms and online CP, and does not support GNNs or LLMs. Additionally, JAX’s steeper learning curve may challenge statisticians familiar with PyTorch. Like **PUNCC**, **Fortuna** defines a dedicated classifier for each score, leading to high-coupling design.

In contrast, **TorchCP**, a PyTorch-native library, significantly enhances conformal prediction for deep learning by providing a comprehensive suite of state-of-the-art CP algorithms. It offers exceptional extensibility through low-coupling design with three modules (See Appendix E) and unified interfaces across various algorithms (See Appendix B). Notably, **TorchCP** seamlessly integrates CP algorithms with advanced deep learning models, such as GNNs and LLMs, and leverages GPU-accelerated batch processing to ensure scalability for large-scale deep learning scenarios (See Appendix D.2).