

HiClass: a Python Library for Local Hierarchical Classification Compatible with Scikit-learn

Fábio M. Miranda*

*Hasso Plattner Institute, Digital Engineering Faculty, University of Potsdam,
14482 Potsdam, Germany
Department of Mathematics and Computer Science, Free University of Berlin,
14195 Berlin, Germany*

FABIO.MALCHERMIRANDA@HPI.DE

Niklas Köhnecke*

*Hasso Plattner Institute, Digital Engineering Faculty, University of Potsdam,
14482 Potsdam, Germany*

NIKLAS.KOEHNECKE@STUDENT.HPI.UNI-POTSDAM.DE

Bernhard Y. Renard

*Hasso Plattner Institute, Digital Engineering Faculty, University of Potsdam,
14482 Potsdam, Germany*

BERNHARD.RENARD@HPI.DE

Editor: Alexandre Gramfort

Abstract

HiClass is an open-source Python library for local hierarchical classification entirely compatible with `scikit-learn`. It contains implementations of the most common design patterns for hierarchical machine learning models found in the literature, that is, the local classifiers per node, per parent node and per level. Additionally, the package contains implementations of hierarchical metrics, which are more appropriate for evaluating classification performance on hierarchical data. The documentation includes installation and usage instructions, examples within tutorials and interactive notebooks, and a complete description of the API. HiClass is released under the simplified BSD license, encouraging its use in both academic and commercial environments. Source code and documentation are available at <https://github.com/scikit-learn-contrib/hiclass>.

Keywords: Local Hierarchical Classification, Supervised Learning, Local Classifier per Node, Local Classifier per Parent Node, Local Classifier per Level

1. Introduction

Many classification problems across different application domains can be naturally modeled hierarchically (Figure 1 and appendix Figures 3-4), typically in the form of trees or directed acyclic graphs (Silla and Freitas, 2011). Examples of hierarchical classification problems are vast, ranging from musical genre classification (Ariyaratne and Zhang, 2012; Iloga et al., 2018) to text categorization (Javed et al., 2021; Ma et al., 2022), taxonomic classification of viral sequences in metagenomic data (Shang and Sun, 2021) and identification of COVID-19 in chest X-ray images (Pereira et al., 2020).

Nonetheless, many classifiers proposed in the literature are designed to completely ignore the existing hierarchy between classes by usually predicting only leaf nodes in a method-

*. These authors contributed equally to this work.

ology known as flat approach. Although easy to implement, the flat approach is incapable of dealing with problems where making a prediction for leaf nodes is not mandatory. Furthermore, since they consider the hierarchy during training, hierarchical models generally produce better results when compared with flat approaches, consistently achieving improvements in predictive performance (Silla and Freitas, 2011). Hence, in this manuscript we introduce **HiClass**, a Python library that implements the most common patterns for local hierarchical classifiers (see Figure 1 and Appendix C for more details), which can be employed in different application domains where the data can be hierarchically structured in the shape of trees or directed acyclic graphs. While the full hierarchical structure is ideally known, classification is also possible with only partially known structures and missing values in the outermost levels of the hierarchy.




Figure 1: Depiction of the flat approach and the local hierarchical classification models implemented in **HiClass**, for a two-level hierarchy. a) Flat approach, where a single multi-class classifier is trained to predict only the leaf nodes in hierarchical data, completely ignoring previous levels of the hierarchy. b) Local classifier per node, where several binary classifiers are trained for each existing node in the hierarchy, excluding the root node. c) Local classifier per parent node, where various multi-class classifiers are trained for each parent node in the hierarchy, with the goal of predicting the children nodes. d) Local classifier per level, where a multi-class classifier is trained for each level in the hierarchy.

2. Overview and Design

The API of `HiClass` is designed to be compatible with the `scikit-learn` API (Pedregosa et al., 2011), offering users a familiar API to train hierarchical models via `fit()`, store trained models with `pickle`, predict labels with the `predict()` method, and create machine learning pipelines. Moreover, each local classifier in `HiClass` is, by default, `scikit-learn`'s `BaseEstimator` class, enabling users to employ any of the classifiers already implemented in `scikit-learn`'s library. However, classes other than the `BaseEstimator` can also be used, as long as they implement the methods `fit()`, `predict()` and `predict_proba()`, thus enabling users to code their own local classifiers using third-party libraries, for example, `TensorFlow` (Abadi et al., 2016) or `PyTorch` (Paszke et al., 2019).

To allow a better integration with `scikit-learn` and faster training, `HiClass` uses for both features and labels `NumPy`'s structured arrays, whose data type is a composition of simpler data types (Harris et al., 2020). While features are exactly the same shape as expected for training flat models in `scikit-learn`, hierarchical training labels are represented as an array of shape $n_samples \times n_levels$, where each column must contain either a label for the respective level in the hierarchy or an empty string that indicates missing labels in the leaf nodes.

The directed acyclic graph (DAG) implemented in the library `NetworkX` (Hagberg et al., 2008) was the data structure chosen to make the models as generic as possible, since hierarchical data in the form of trees can also be represented with a DAG. Training can be performed in parallel for each local classifier by leveraging either the library `Ray`¹ or `Joblib`², while prediction, which is not a bottleneck, is performed from top to bottom, following the hierarchical structure to keep consistency among the several hierarchical levels.

According to Silla and Freitas (2011), the use of flat classification metrics might not be adequate to give enough insight on which algorithm is better at classifying hierarchical data. Hence, in `HiClass` we implemented the metrics of hierarchical precision (hP), hierarchical recall (hR) and hierarchical F-score (hF), which are extensions of the renowned metrics of precision, recall and F-score, but tailored to the hierarchical classification scenario. These hierarchical counterparts were initially proposed by Kiritchenko et al. (2006), and are defined as follows:

$$hP = \frac{\sum_i |\alpha_i \cap \beta_i|}{\sum_i |\alpha_i|}, hR = \frac{\sum_i |\alpha_i \cap \beta_i|}{\sum_i |\beta_i|}, hF = \frac{2 \times hP \times hR}{hP + hR}$$

, where α_i is the set consisting of the most specific classes predicted for test example i and all their ancestor classes, while β_i is the set containing the true most specific classes of test example i and all their ancestors, with summations computed over all test examples.

3. Code Quality Assurance

To ensure high code quality, *inheritance* is applied whenever possible to keep the code concise and easier to maintain, and all implementations adhere to the *PEP 8* code style (Van Rossum et al., 2001) enforced by `flake8` and the uncompromising code formatter `black`. API documentation is provided through *docstrings* (Goodger and van Rossum,

1. <https://www.ray.io/>

2. <https://joblib.readthedocs.io/en/latest/parallel.html>

2010), and the implementations are accompanied by *unit tests* that cover 98% of our code and are automatically executed by our *continuous integration workflow* upon commits.

4. Installation and Usage

HiClass is hosted on GitHub³, while tutorials and API documentation are available on the platform Read the Docs⁴. Packages for Python 3.7-3.9 are available for Linux, macOS and Windows and can be obtained with `pip install hiclass` or `conda install -c conda-forge hiclass`. Code 1 shows a basic example of fitting and evaluating a local hierarchical model with HiClass. More elaborate examples can be found in the tutorials.

```

from hiclass import LocalClassifierPerNode
from hiclass.metrics import f1
from sklearn.ensemble import RandomForestClassifier

# define mock data
X_train = X_test = [[1, 2], [3, 4]]
Y_train = Y_test = [
    ["Animal", "Mammal", "Cat"],
    ["Animal", "Reptile", "Turtle"],
]

# Use random forest classifiers for every node
rf = RandomForestClassifier()
lcpn = LocalClassifierPerNode(local_classifier=rf)
lcpn.fit(X_train, Y_train) # Train model
predictions = lcpn.predict(X_test) # Predict test data

# Print hierarchical F-score
print(f"f1: {f1(y_true=Y_test, y_pred=predictions)}")

```

Code 1: Example on how to use HiClass to train and evaluate a hierarchical classifier.

5. Comparison with Flat Classifiers

While HiClass focuses on ease of use and is fully written in a high level language, care has been taken to maximize computational efficiency. In Figure 2, we compare the hierarchical F-score, computational resources (measured with the command *time*) and disk usage. This comparison was performed between two flat classifiers from the library `scikit-learn` and Microsoft’s LightGBM (Ke et al., 2017) versus the local hierarchical classifiers implemented in HiClass. In order to avoid bias, cross-validation and hyperparameter tuning were performed on the local hierarchical classifiers and flat classifiers.

3. <https://github.com/scikit-learn-contrib/hiclass>

4. <https://hiclass.readthedocs.io/>

Comparison between the flat approach and the local hierarchical classifiers implemented in HiClass

Base classifier: LightGBM

| MODEL | F-SCORE | MEMORY USAGE (GB) | DISK USAGE (MB) | TRAINING TIME (S) |
|----------------------------------|---------|-------------------|-----------------|-------------------|
| Local Classifier per Parent Node | 0.7531 | 9.00 | 77 | 1,680 |
| Local Classifier per Node | 0.7901 | 31.92 | 412 | 3,355 |
| Local Classifier per Level | 0.6854 | 9.04 | 36 | 5,726 |
| Flat Classifier | 0.3710 | 4.54 | 30 | 4,284 |

Base classifier: Logistic Regression

| MODEL | F-SCORE | MEMORY USAGE (GB) | DISK USAGE (MB) | TRAINING TIME (S) |
|----------------------------------|---------|-------------------|-----------------|-------------------|
| Local Classifier per Parent Node | 0.7798 | 9.30 | 122 | 1,925 |
| Local Classifier per Node | 0.7763 | 21.01 | 123 | 245 |
| Local Classifier per Level | 0.7795 | 11.45 | 124 | 8,684 |
| Flat Classifier | 0.7748 | 3.15 | 108 | 3,436 |

Base classifier: Random Forest

| MODEL | F-SCORE | MEMORY USAGE (GB) | DISK USAGE (GB) | TRAINING TIME (S) |
|----------------------------------|---------|-------------------|-----------------|-------------------|
| Local Classifier per Parent Node | 0.7407 | 48.30 | 24 | 27,287 |
| Local Classifier per Node | 0.7668 | 55.19 | 27 | 24,617 |
| Local Classifier per Level | 0.7383 | 191.39 | 96 | 35,118 |
| Flat Classifier | 0.6672 | 162.40 | 81 | 5,215 |

Figure 2: Comparison between the flat approach and the local hierarchical classifiers implemented in HiClass, using the consumer complaints data set and Microsoft’s LightGBM (Ke et al., 2017), Logistic Regression and Random Forest as the base classifiers. For this benchmark, the metrics used were hierarchical F-score, memory usage in gigabyte, disk usage in megabyte or gigabyte, and training time in seconds. Overall, the hierarchical classifiers improved the F-score when compared with the flat approaches, while in some occasions the local hierarchical classifiers further reduced memory consumption, disk usage, and training time.

For comparison purposes, we used a snapshot from 02/11/2022 of the consumer complaints data set provided by the Consumer Financial Protection Bureau of the United States (Bureau and General, 2022), which after preprocessing contained 727,495 instances for cross-validation and hyperparameter tuning as well as training and 311,784 more for validation. Additional descriptions about this data set, preprocessing, feature extraction,

and experimental setup are available in Appendices A-B. The benchmark was computed on multiple cluster nodes running GNU/Linux with 512 GB physical memory and 128 cores provided by two AMD EPYC™ 7742 processors. A reproducible Snakemake pipeline (Köster and Rahmann, 2012) is available in our public repository⁵.

Our results reveal that the hierarchical F-score was enhanced by all local hierarchical approaches when compared with the flat classifiers, achieving a maximum improvement of $\approx 113\%$ when comparing the local classifier per node (LCPN) with the flat LightGBM approach. Regarding training time, the LCPN and local classifier per parent node (LCPPN) decreased computational time by $\approx 93\%$ and $\approx 44\%$, respectively, when compared with the flat Logistic Regression classifier. When compared with the flat Random Forest, the LCPN and LCPPN reduced memory usage by $\approx 66\%$ and $\approx 70\%$ and disk usage by $\approx 67\%$ and $\approx 70\%$, respectively.

6. Conclusion

HiClass is a Python package that provides implementations of popular machine learning models and evaluation metrics for local hierarchical classification. Thanks to its compatibility with the `scikit-learn` API, users can choose among existing classifiers to create powerful hierarchical classifiers for hierarchical data.

In future releases, we plan to implement multi-label hierarchical classification to expand the range of problems the library can solve. Additionally, we also plan to implement the global approach, which trains a single model that learns the entire hierarchical structure. Lastly, we also intend to add support for incremental learning in order to enable training of larger than memory data sets.

Acknowledgments

The authors gratefully acknowledge all users for requesting useful new features and reporting bugs. BYR gratefully acknowledges support by the BMBF-funded Computational Life Science initiative (project DeepPath, 031L0208, to B.Y.R.).

Appendix A. Hierarchical Data

In this appendix we delineate what constitutes hierarchical data and the data set used for evaluation in Section 5.

Numerous real-life problems can be naturally modeled hierarchically, i.e., they can be computationally represented as directed acyclic graphs or trees. Two notorious examples of hierarchical data are music genre and phylogeny, which are depicted in Figures 3-4, respectively.

HiClass facilitates training local hierarchical classifiers. Hierarchical labels can be simply defined in a $m \times n$ matrix, where each row is a training example and each column is a level in the hierarchy. Such matrix can be represented with Python lists, NumPy arrays or Pandas DataFrames. Python lists and Pandas DataFrames are automatically converted

5. https://github.com/scikit-learn-contrib/hiclass/tree/main/benchmarks/consumer_complaints

to NumPy arrays for efficient processing. Training features need to be numerical, hence feature extraction might be necessary depending on the data.




Figure 3: When labeling and retrieving musical information, the genre plays an important role, since having the musical genres structured in a class hierarchy simplifies how users browse and retrieve this information (Silla and Freitas, 2011). Image adapted from Silla and Freitas (2011).

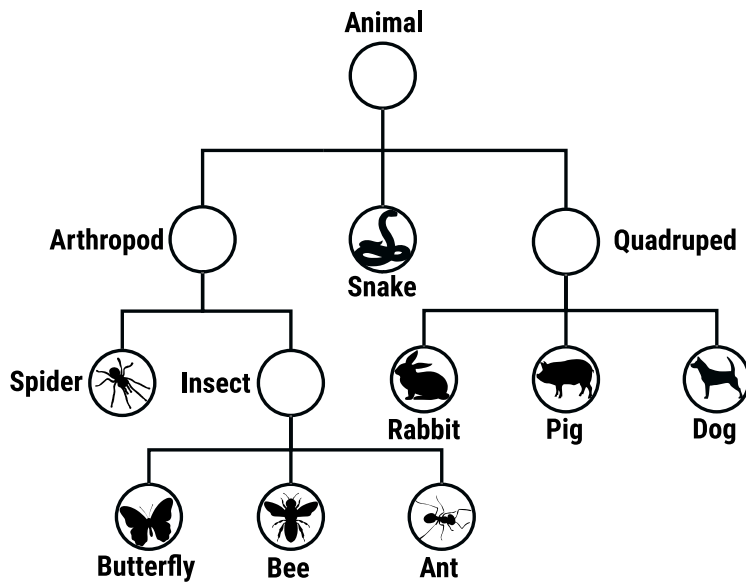


Figure 4: When grouping different species, scientists typically build a hierarchical structure to distinguish organisms according to evolutionary traits. Image inspired by Barutcuoglu and DeCoro (2006).

Consumer Complaints Data Set

The consumer complaints data set is a database of complaints sent to companies in the United States (Bureau and General, 2022). This database, which is depicted on Table 1, is maintained by the Consumer Financial Protection Bureau (CFPB) and is updated regularly. For full reproducibility of the results, we took a snapshot from 02/11/2022 that we used for evaluation in our benchmark, and it is available upon request.

| Consumer Complaint Narrative | Product | Sub-product |
|------------------------------|--------------|-----------------------------|
| Loan | Student loan | Struggling to repay loan |
| Credit reporting | Reports | Unable to get annual report |

Table 1: Depiction of consumer complaint data set, obtained from the Consumer Financial Protection Bureau (Bureau and General, 2022).

According to the CFPB, complaints can give us insights into problems customers are experiencing and help regulate products and services, enforce laws, and educate and empower consumers. For our benchmark we used the complaint narratives as feature, which according to the CFPB are consumers’ descriptions of their experiences in their own words. The classification task here consists of labeling which product and sub-product the consumer is complaining about. Classifying missing labels is a recurrent problem in various application domains, since it is common for users not to use appropriate keywords or leave fields entirely empty. For instance, automatic classification of IT tickets has become quite popular in the last couple of years (Revina et al., 2020), but unfortunately we could not find any comprehensive data set publicly available for a benchmark.

Feature extraction was performed using the CountVectorizer and TfidfTransformer implementations available on scikit-learn (Pedregosa et al., 2011), in order to compute a matrix of token counts and term-frequency, respectively, with default parameters used for both. Rows with empty cells were discarded, and around 70% of the remaining data was used for hyperparameter tuning and training (727,495 examples), while 30% was held for validation (311,784 examples). For the flat classifiers, labels were concatenated before training and split back into their original shape after prediction.

Appendix B. Evaluation

This appendix provides additional description about the experimental setup mentioned in Section 5.

In order to avoid overfitting, our experiment began with splitting the consumer complaints data set into two subsets. The first subset contained 70% of the data for hyperparameter tuning and training, while the remaining 30% was held for testing (Figure 5). This first division was performed via scikit-learn’s `train_test_split` method. Afterwards, the training subset was further divided into 5 splits to perform 5-fold cross-validation for hyperparameter tuning. These splits were achieved with the help of the `KFold` class from scikit-learn.

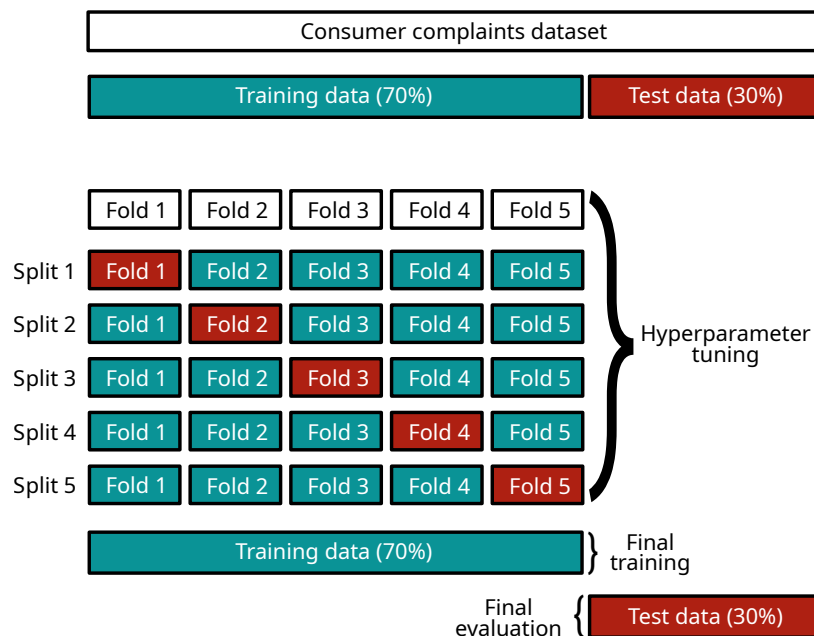


Figure 5: Visual representation of the experiments performed for hyperparameter tuning, training and evaluation. First the data set was split with 70% of the data being used for hyperparameter tuning and training, while 30% was held for a final evaluation. The subset with 70% of data held for training was further split into 5 subsets for 5-fold cross-validation and identification of best hyperparameter combination. Lastly, the models were trained on the whole 70% training data, using the best hyperparameters detected, and a final evaluation was performed with the 30% test data held in the beginning. Image adapted from Scikit-learn (2022).

For hyperparameter tuning, the models were trained using 4 folds as training data and validated on the remaining one. This process was repeated 5 times, with different folds combinations being used in each iteration, and the average hierarchical F-score was reported as the performance metric. The selection of the best hyperparameters was assisted by Hydra (Meta, 2022) and its plugin Optuna (Akiba et al., 2019), through a grid search using the combinations of hyperparameters described in Tables 2-4. After the best hyperparameter combinations were identified, the models were trained once more, but this time using the entire training data (70% of the full data set) and a final evaluation was carried out on the test data held in the beginning (30% of the entire data). All these steps were automatically executed in a Snakemake pipeline, which is available in our repository⁶.

6. https://github.com/scikit-learn-contrib/hiclass/tree/main/benchmarks/consumer_complaints

| num_leaves | n_estimators | min_child_samples |
|-------------------|---------------------|--------------------------|
| 31 | 100 | 20 |
| 62 | 200 | 40 |

Table 2: Hyperparameters tested for LightGBM.

| solver | max_iter |
|---------------|-----------------|
| newton-cg | 10000 |
| lbfgs | |
| liblinear | |
| sag | |
| saga | |

Table 3: Hyperparameters tested for logistic regression.

| n_estimators | criterion |
|---------------------|------------------|
| 100 | gini |
| 200 | entropy |
| | log_loss |

Table 4: Hyperparameters tested for random forest.

For comparison purposes, in Figure 6 we show the results without hyperparameter tuning.

Appendix C. Algorithms Overview

This appendix provides rigorous descriptions for the algorithms mentioned in earlier sections.

HiClass provides implementations for the most popular machine learning models for local hierarchical classification, including the Local Classifier Per Node, the Local Classifier Per Parent Node and the Local Classifier Per Level. In the following subsections, we present in more details these different approaches for local hierarchical classification.

Local Classifier Per Node

One of the most popular approaches in the literature, the local classifier per node consists of training one binary classifier for each node of the class taxonomy, except for the root node. A visual representation of the local classifier per node is shown in Figure 7.

Each binary classifier can be trained in parallel using either the library Ray or Joblib. In order to avoid inconsistencies, prediction is performed in a top-down manner. For example, given a hypothetical test example, the local classifier per node firstly queries the binary classifiers at nodes “Reptile” and “Mammal”. Supposing that in this hypothetical situation the probability of the test example belonging to class “Reptile” is 0.8, and the probability of belonging to class “Mammal” is 0.5, then class “Reptile” is selected for the first level.

At the next level, only the classifiers at nodes “Snake” and “Lizard” are queried, and again the one with the highest probability is chosen.

Comparison between the flat approach and the local hierarchical classifiers implemented in HiClass

Base classifier: LightGBM

| MODEL | F-SCORE | MEMORY USAGE (GB) | DISK USAGE (MB) | TRAINING TIME (S) |
|----------------------------------|---------|-------------------|-----------------|-------------------|
| Local Classifier per Parent Node | 0.7077 | 8.97 | 79 | 1,533 |
| Local Classifier per Node | 0.7702 | 26.99 | 325 | 2,501 |
| Local Classifier per Level | 0.6383 | 9.17 | 23 | 4,769 |
| Flat Classifier | 0.3175 | 4.61 | 28 | 2,929 |

Base classifier: Logistic Regression

| MODEL | F-SCORE | MEMORY USAGE (GB) | DISK USAGE (MB) | TRAINING TIME (S) |
|----------------------------------|---------|-------------------|-----------------|-------------------|
| Local Classifier per Parent Node | 0.7799 | 9.17 | 122 | 1,899 |
| Local Classifier per Node | 0.7763 | 11.52 | 123 | 572 |
| Local Classifier per Level | 0.7795 | 11.45 | 124 | 8,639 |
| Flat Classifier | 0.7748 | 7.51 | 108 | 8,366 |

Base classifier: Random Forest

| MODEL | F-SCORE | MEMORY USAGE (GB) | DISK USAGE (GB) | TRAINING TIME (S) |
|----------------------------------|---------|-------------------|-----------------|-------------------|
| Local Classifier per Parent Node | 0.7359 | 24.91 | 12 | 13,803 |
| Local Classifier per Node | 0.7638 | 28.20 | 14 | 12,181 |
| Local Classifier per Level | 0.7330 | 96.31 | 48 | 17,462 |
| Flat Classifier | 0.6661 | 82.08 | 41 | 2,758 |

Figure 6: Comparison between the flat approach and the local hierarchical classifiers implemented in HiClass, using the consumer complaints data set and Microsoft’s LightGBM (Ke et al., 2017), Logistic Regression and Random Forest as the base classifiers. For this benchmark, the metrics used were hierarchical F-score, memory usage in gigabyte, disk usage in megabyte or gigabyte, and training time in seconds. No hyperparameter tuning was performed for these results. Overall, the hierarchical classifiers improved the F-score when compared with the flat approach, while in some occasions the local hierarchical classifiers further reduced memory consumption, disk usage, and training time.

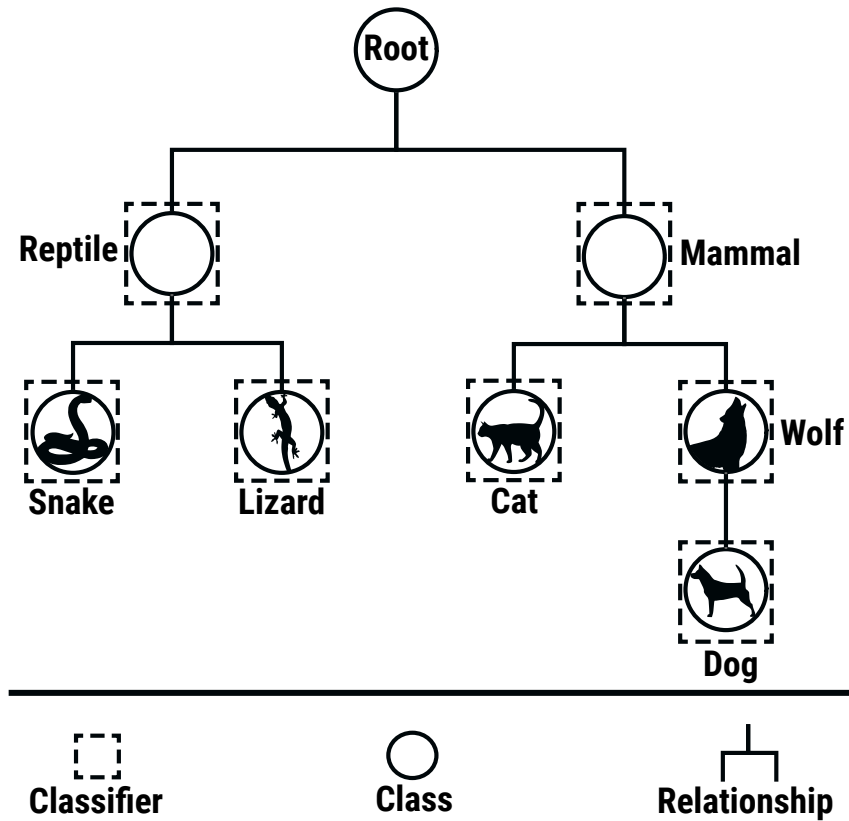


Figure 7: Visual representation of the local classifier per node approach, where binary classifiers (squares) are trained for each class (circles) of the hierarchy, excluding the root node.

TRAINING POLICIES

There are multiple ways to define the set of positive and negative examples for training the binary classifiers. In HiClass we implemented 6 policies described by Silla and Freitas (2011), which were based on previous work from Eisner et al. (2005) and Fagni and Sebastiani (2007). The notation used to define the sets of positive and negative examples is presented in Table 5, as described by Silla and Freitas (2011).

| Symbol | Meaning |
|-----------------------------|---|
| Tr | The set of all training examples |
| $Tr^+(c_i)$ | The set of positive training examples of c_i |
| $Tr^-(c_i)$ | The set of negative training examples of c_i |
| $\uparrow(c_i)$ | The parent category of c_i |
| $\downarrow(c_i)$ | The set of children categories of c_i |
| $\uparrow\uparrow(c_i)$ | The set of ancestor categories of c_i |
| $\downarrow\downarrow(c_i)$ | The set of descendant categories of c_i |
| $\leftrightarrow(c_i)$ | The set of sibling categories of c_i |
| $*(c_i)$ | Denotes examples whose most specific known class is c_i |

Table 5: Notation used to define the sets of positive and negative examples.

Based on this notation, we can define the different policies and their sets of positive and negative examples as follows:

| Policy | Positive examples | Negative examples |
|---------------------------|---|---|
| Exclusive | $Tr^+(c_i) = *(c_i)$ | $Tr^-(c_i) = Tr \setminus *(c_i)$ |
| Less exclusive | $Tr^+(c_i) = *(c_i)$ | $Tr^-(c_i) = Tr \setminus *(c_i) \cup \downarrow(c_i)$ |
| Less inclusive | $Tr^+(c_i) = *(c_i) \cup \downarrow(c_i)$ | $Tr^-(c_i) = Tr \setminus *(c_i) \cup \downarrow(c_i)$ |
| Inclusive | $Tr^+(c_i) = *(c_i) \cup \downarrow(c_i)$ | $Tr^-(c_i) = Tr \setminus *(c_i) \cup \downarrow(c_i) \cup \uparrow(c_i)$ |
| Siblings | $Tr^+(c_i) = *(c_i) \cup \downarrow(c_i)$ | $Tr^-(c_i) = \leftrightarrow(c_i) \cup \downarrow(\leftrightarrow(c_i))$ |
| Exclusive siblings | $Tr^+(c_i) = *(c_i)$ | $Tr^-(c_i) = \leftrightarrow(c_i)$ |

Table 6: Policies used to define the sets of positive and negative examples.

Using as example the class “Wolf” from the hierarchy represented in Figure 7, we have the following sets of positive and negative examples for each policy:

| Policy | $Tr^+(c_{\text{Wolf}})$ | $Tr^-(c_{\text{Wolf}})$ |
|---------------------------|-------------------------|--|
| Exclusive | Wolf | Reptile, Snake, Lizard, Mammal, Cat, Dog |
| Less exclusive | Wolf | Reptile, Snake, Lizard, Mammal, Cat |
| Less inclusive | Wolf, Dog | Reptile, Snake, Lizard, Mammal, Cat |
| Inclusive | Wolf, Dog | Reptile, Snake, Lizard, Cat |
| Siblings | Wolf, Dog | Cat |
| Exclusive siblings | Wolf | Cat |

Table 7: Sets of positive and negative examples for each policy, given the class “Wolf”.

Local Classifier Per Parent Node

The local classifier per parent node approach consists of training a multi-class classifier for each parent node existing in the hierarchy, as shown in Figure 8.

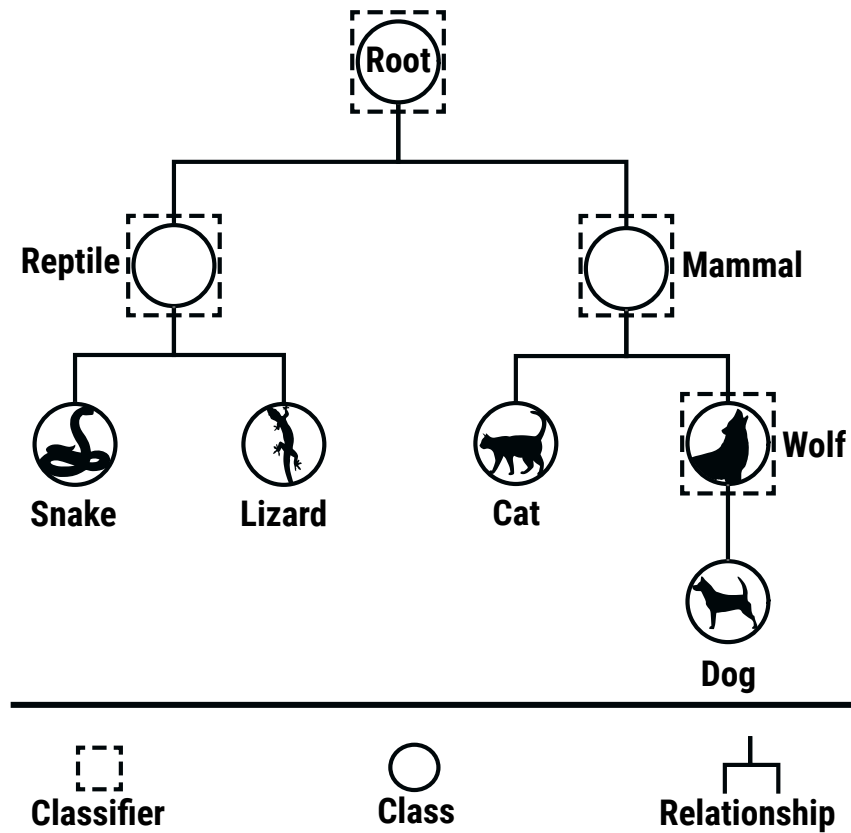


Figure 8: Visual representation of the local classifier per parent node approach, where multi-class classifiers (squares) are trained for each parent node existing in the class hierarchy (circles).

While training can be executed in parallel, prediction is always performed in a top-down style in order to avoid inconsistencies. For example, assuming that the classifier located at the root node decides that a test example belongs to class “Mammal”, then the next level can only be predicted by the classifier located at node “Mammal”, which in turn will ultimately decide if the test example belongs either to the class “Cat” or “Wolf/Dog”.

Local Classifier Per Level

The local classifier per level approach consists of training a multi-class classifier for each level of the class taxonomy. An example is displayed on Figure 9.

Similar to the other hierarchical classifiers, the local classifier per level can also be trained in parallel, and prediction is performed in a top-down mode to avoid inconsistencies. For example, supposing that for a given test example the classifier at the first level returns the probabilities 0.91 and 0.7 for classes “Reptile” and “Mammal”, respectively, then the one with the highest probability is considered as the correct prediction, which in this case is

class “Reptile”. For the second level, only the probabilities for classes “Snake” and “Lizard” are considered and the one with the highest probability is the final prediction.

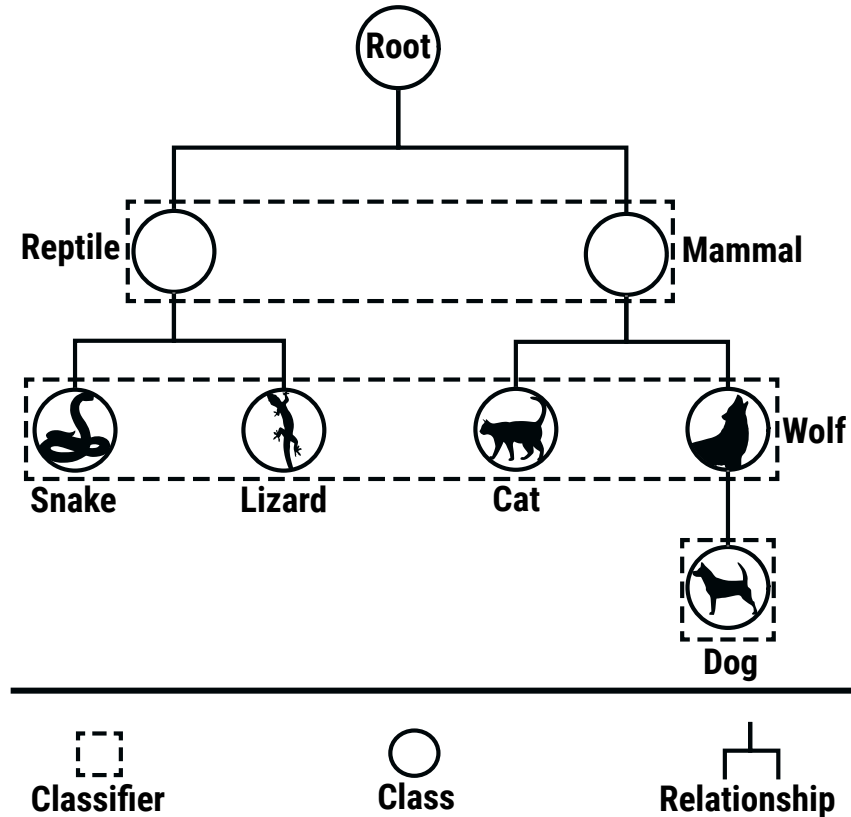


Figure 9: Visual representation of the local classifier per level approach.

References

- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: a system for large-scale machine learning. In *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 265–283, 2016.
- Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: a next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2623–2631, 2019.
- Hasitha B Ariyaratne and Dengsheng Zhang. A novel automatic hierachical approach to music genre classification. In *Proceedings of the IEEE International Conference on Multimedia and Expo Workshops*, pages 564–569. IEEE, 2012.

- Zafer Barutcuoglu and Christopher DeCoro. Hierarchical shape classification using bayesian aggregation. In *Proceedings of the IEEE International Conference on Shape Modeling and Applications 2006 (SMI'06)*, pages 44–44. IEEE, 2006.
- Consumer Financial Protection Bureau and Illinois Attorney General. Consumer complaints. <https://www.consumerfinance.gov/data-research/consumer-complaints/>, 2022. [Online; accessed 19-May-2022].
- Roman Eisner, Brett Poulin, Duane Szafron, Paul Lu, and Russell Greiner. Improving protein function prediction using the hierarchical structure of the gene ontology. In *Proceedings of the IEEE Symposium on Computational Intelligence in Bioinformatics and Computational Biology*, pages 1–10. IEEE, 2005.
- Tiziano Fagni and Fabrizio Sebastiani. On the selection of negative examples for hierarchical text categorization. In *Proceedings of the 3rd Language Technology Conference*, pages 24–28, 2007.
- David Goodger and Guido van Rossum. Docstring conventions. In *Pro Python*, pages 303–307. Springer, 2010.
- Aric Hagberg, Pieter Swart, and Daniel S Chult. Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2008.
- Charles R Harris, K Jarrod Millman, Stéfan J van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. Array programming with numpy. *Nature*, 585(7825):357–362, 2020.
- Sylvain Iloga, Olivier Romain, and Maurice Tchuente. A sequential pattern mining approach to design taxonomies for hierarchical music genre recognition. *Pattern Analysis and Applications*, 21(2):363–380, 2018.
- Taimoor Ahmed Javed, Waseem Shahzad, and Umair Arshad. Hierarchical text classification of urdu news using deep neural network. *arXiv preprint arXiv:2107.03141*, 2021.
- Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in Neural Information Processing Systems*, 30, 2017.
- Svetlana Kiritchenko, Stan Matwin, Richard Nock, and A Fazel Famili. Learning and evaluation in the presence of class hierarchies: application to text categorization. In *Proceedings of the Conference of the Canadian Society for Computational Studies of Intelligence*, pages 395–406. Springer, 2006.
- Johannes Köster and Sven Rahmann. Snakemake—a scalable bioinformatics workflow engine. *Bioinformatics*, 28(19):2520–2522, 2012.
- Yinglong Ma, Xiaofeng Liu, Lijiao Zhao, Yue Liang, Peng Zhang, and Beihong Jin. Hybrid embedding-based text representation for hierarchical multi-label text classification. *Expert Systems with Applications*, 187:115905, 2022.

- Meta. Hydra: a framework for elegantly configuring complex applications. <https://hydra.cc/>, 2022. [Online; accessed 14-Nov-2022].
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: an imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32:8026–8037, 2019.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: machine learning in python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Rodolfo M Pereira, Diego Bertolini, Lucas O Teixeira, Carlos N Silla Jr, and Yandre MG Costa. Covid-19 identification in chest x-ray images on flat and hierarchical classification scenarios. *Computer Methods and Programs in Biomedicine*, 194:105532, 2020.
- Aleksandra Revina, Krisztian Buza, and Vera G Meister. It ticket classification: the simpler, the better. *IEEE Access*, 8:193380–193395, 2020.
- Scikit-learn. Cross-validation: evaluating estimator performance. https://scikit-learn.org/stable/modules/cross_validation.html, 2022. [Online; accessed 14-Nov-2022].
- Jiayu Shang and Yanni Sun. Cheer: hierarchical taxonomic classification for viral metagenomic data via deep learning. *Methods*, 189:95–103, 2021.
- Carlos N Silla and Alex A Freitas. A survey of hierarchical classification across different application domains. *Data Mining and Knowledge Discovery*, 22(1):31–72, 2011.
- Guido Van Rossum, Barry Warsaw, and Nick Coghlan. Pep 8: style guide for python code. *Python.org*, 1565, 2001.