
pyconll Documentation

Release 2.2.1

Matias Grioni

Aug 07, 2020

Contents

| | | |
|----------|--------------------------------|-----------|
| 1 | pyconll | 3 |
| 1.1 | Links | 3 |
| 2 | CHANGELOG | 7 |
| 2.1 | [2.2.1] - 2019-11-17 | 7 |
| 2.2 | [2.2.0] - 2019-10-01 | 8 |
| 2.3 | [2.1.1] - 2019-09-04 | 8 |
| 2.4 | [2.1.0] - 2019-08-30 | 8 |
| 2.5 | [2.0.0] - 2019-05-09 | 8 |
| 2.6 | [1.1.4] - 2019-04-15 | 9 |
| 2.7 | [1.1.3] - 2019-01-03 | 9 |
| 2.8 | [1.1.2] - 2018-12-28 | 9 |
| 2.9 | [1.1.1] - 2018-12-10 | 10 |
| 2.10 | [1.1.0] - 2018-11-11 | 10 |
| 2.11 | [1.0.1] - 2018-09-14 | 10 |
| 2.12 | [1.0] - 2018-09-13 | 10 |
| 2.13 | [0.3.1] - 2018-08-08 | 11 |
| 2.14 | [0.3] - 2018-07-28 | 11 |
| 2.15 | [0.2.3] - 2018-07-23 | 11 |
| 2.16 | [0.2.2] - 2018-07-18 | 11 |
| 2.17 | [0.2.1] - 2018-07-18 | 12 |
| 2.18 | [0.2] - 2018-07-16 | 12 |
| 2.19 | [0.1.1] - 2018-07-15 | 12 |
| 2.20 | [0.1] - 2018-07-04 | 12 |
| 3 | Getting Started | 13 |
| 3.1 | Overview | 13 |
| 3.2 | Loading CoNLL-U | 13 |
| 3.3 | Traversing CoNLL-U | 13 |
| 3.4 | Outputting CoNLL-U | 14 |
| 3.5 | Conclusion | 14 |
| 4 | load | 15 |
| 4.1 | Example | 15 |
| 4.2 | API | 16 |
| 5 | token | 19 |

| | | |
|----------------------------|--------------------|-----------|
| 5.1 | Fields | 19 |
| 5.2 | API | 20 |
| 6 | sentence | 23 |
| 6.1 | Comments | 23 |
| 6.2 | Tokens | 23 |
| 6.3 | API | 24 |
| 7 | conll | 27 |
| 7.1 | API | 27 |
| 8 | tree | 29 |
| 8.1 | API | 29 |
| 9 | util | 31 |
| 9.1 | API | 31 |
| 10 | conllable | 33 |
| 10.1 | API | 33 |
| 11 | exception | 35 |
| 11.1 | API | 35 |
| Python Module Index | | 37 |
| Index | | 39 |

*Easily work with ****CoNLL*** files using the familiar syntax of **python**.*

1.1 Links

- [Homepage](#)
- [Documentation](#)

1.1.1 Installation

As with most python packages, simply use `pip` to install from PyPi.

```
pip install pyconll
```

`pyconll` is also available as a conda package on the `pyconll` channel. Only package 2.2.0 and newer are available for conda at the moment.

```
conda install -c pyconll pyconll
```

This package is designed for, and only tested with python 3.4 and up and will not be backported to python 2.x. File an issue if support is desired for a python version less than 3.4.

1.1.2 Motivation

This tool is intended to be a **minimal**, **low level**, and **functional** library in a widely used programming language. `pyconll` creates a thin API on top of raw CoNLL annotations that is simple and intuitive in a popular programming language.

In my work with the Universal Dependencies project, I saw a dissapointing lack of low level APIs for working with the CoNLL-U format. Most tooling focuses on graph transformations and DSLs for terse, automated changes. Tools

such as [Grew](#) and [Treex](#) are very powerful and productive, but have a learning curve and are limited the scope of their DSLs. [CL-CoNLLU](#) is simple and low level, but Common Lisp is not widely used in NLP, and difficult to pickup for beginners. [UDAPI](#) is in python but it is very large and has little guidance. `pyconll` attempts to fill the gaps between what other projects have accomplished.

Other useful tools can be found on the Universal Dependencies [website](#).

Hopefully, individual researchers find `pyconll` useful, and will use it as a building block for their tools and projects. `pyconll` affords a standardized and complete base for building larger projects without worrying about CoNLL annotation and output.

1.1.3 Code Snippet

```
# This snippet finds what lemmas are marked as AUX which is a closed class POS in UD
import pyconll

UD_ENGLISH_TRAIN = './ud/train.conll'

train = pyconll.load_from_file(UD_ENGLISH_TRAIN)

aux_lemmas = set()
for sentence in train:
    for token in sentence:
        if token.upos == 'AUX':
            aux_lemmas.add(token.lemma)
```

1.1.4 Uses and Limitations

This package edits CoNLL-U annotations. This does not include the annotated text itself. Word forms on Tokens are not editable and Sentence Tokens cannot be reassigned or reordered. `pyconll` focuses on editing CoNLL-U annotation rather than creating it or changing the underlying text that is annotated. If there is interest in this functionality area, please create a github issue for more visibility.

This package also is only validated against the CoNLL-U format. The CoNLL and CoNLL-X format are not supported, but are very similar. I originally intended to support these formats as well, but their format is not as well defined as CoNLL-U so they are not included. Please create an issue for visibility if this feature interests you.

Lastly, linguistic data can often be very large and this package attempts to keep that in mind. `pyconll` provides methods for creating in memory conll objects along with an iterate only version in case a corpus is too large to store in memory (the size of the memory structure is several times larger than the actual corpus file). The iterate only version can parse upwards of 100,000 words per second on a 16gb ram machine, so for most datasets to be used on a dev machine, this package will perform well. The 2.2.0 release also improves parse time and memory footprint by about 25%!

1.1.5 Contributing

Contributions to this project are welcome and encouraged! If you are unsure how to contribute, here is a [guide](#) from Github explaining the basic workflow. After cloning this repo, please run `make hooks` and `pip install -r requirements.txt` to properly setup locally. `make hooks` setups up a pre-push hook to validate that code matches the default YAPF style. While this is technically optional, it is highly encouraged. `pip install -r requirements.txt` sets up environment dependencies like `yapf`, `twine`, `sphinx`, etc.

For packaging new versions, please use `setuptools` version 24.2.0 or greater for creating the appropriate packaging that recognizes the `python_requires` metadata.

README and CHANGELOG

When changing either of these files, please change the Markdown version and run `make docs` so that the other versions stay in sync.

Code Formatting

Code formatting is done automatically on push if githooks are setup properly. The code formatter is [YAPF](#), and using this ensures that coding style stays consistent over time and between authors.

CHAPTER 2

CHANGELOG

All notable changes to this project will be documented in this file.

The format is based on [Keep a Changelog](#) and this project adheres to [Semantic Versioning](#).

2.1 [2.2.1] - 2019-11-17

2.1.1 Fixed

- Branding information typo within setup.py
- Spurious command in Makefile recipe

2.1.2 Added

- Added `python_requires` clause to setup.py to prevent installation on unsupported platforms
- Include information in README about `setuptools` version needed to properly package within `python_requires` information
- Conda packaging support along with information in README about new installation method

2.1.3 Changed

- `pyconll` version is now housed in `.version` file so that this version only needs to be changed in one place before release.

2.2 [2.2.0] - 2019-10-01

2.2.1 Changed

- Use slots on Token and Sentence class for more efficient memory usage with large amounts of objects
- Remove source fields on Token and Sentence. These were not an explicit part of the public API so this is not considered a breaking change.

2.3 [2.1.1] - 2019-09-04

2.3.1 Fixed

- Solved `math.inf` issue with python 3.4 where it does not exist

2.4 [2.1.0] - 2019-08-30

2.4.1 Fixed

- The example `reannotate_ngrams.py` was out of sync with the function return type

2.4.2 Added

- ‘`find_nonprojective_deps`’ was added to look for non-projective dependencies within a sentence

2.5 [2.0.0] - 2019-05-09

2.5.1 Fixed

- `find_ngrams` in the `util` module did not properly match case insensitivity.
- `conllable` is now properly included in wildcard imports from `pyconll`.
- Issue when loading a CoNLL file over a network if the file contained UTF-8 characters. requests default assumes ASCII encoding on HTTP responses.
- The Token columns `deps` and `feats` were not properly sorted by attribute (either numeric index or case invariant lexicographic sort) on serialization

2.5.2 Changed

- Clearer and more concise documentation
- `find_ngrams` now returns the matched tokens as the last element of the yielded tuple.

2.5.3 Removed

- Document and paragraph ids on Sentences
- Line numbers on Tokens and Sentences
- Equality comparison on Tokens and Sentences. These types are mutable and implementing equality (with no hash overriding) causes issues for API clients.
- `SentenceTree` module. This functionality was moved to the `Sentence` class method `to_tree`.

2.5.4 Added

- `to_tree` method on `Sentence` that returns the `Tree` representing the `Sentence` dependency structure

2.5.5 Security

- Updates to `requirements.txt` to patch Jinja2 and requests

2.6 [1.1.4] - 2019-04-15

2.6.1 Fixed

- Parsing of underscore's for the `form` and `lemma` field, would automatically default to `None`, rather than the intended behavior.

2.7 [1.1.3] - 2019-01-03

2.7.1 Fixed

- When used on Windows, the default encoding of Windows-1252 was used when loading CoNLL-U files, however, CoNLL-U is UTF-8. This is now fixed.

2.8 [1.1.2] - 2018-12-28

2.8.1 Added

- *Getting Started* page on the documentation to make easier for newcomers

2.8.2 Fixed

- Versioning on docs page which had not been properly updated
- Some documentation errors
- `requests` version used in `requirements.txt` was insecure and updated to newer version

2.9 [1.1.1] - 2018-12-10

2.9.1 Fixed

- The `pyconll.tree` module was not properly included before in `setup.py`

2.10 [1.1.0] - 2018-11-11

2.10.1 Added

- `pylint` to build process
- `Conllable` abstract base class to mark CoNLL serializable components
- Tree data type construction of a sentence

2.10.2 Changed

- Linting patches suggested by `pylint`.
- Removed `_end_line_number` from `Sentence` constructor. This is an internal patch, as this parameter was not meant to be used by callers.
- New, improved, and clearer documentation
- Update of `requests` dependency due to security flaw

2.11 [1.0.1] - 2018-09-14

2.11.1 Changed

- Removed test packages from final shipped package.

2.12 [1.0] - 2018-09-13

2.12.1 Added

- There is now a `FormatError` to help make debugging easier if the internal data of a `Token` is put into an invalid state. This error will be seen on running `Token#conll`.
- Certain token fields with empty values, were not output when calling `Token#conll` and were instead ignored. This situation now causes a `FormatError`.
- Stricter parsing and validation of general CoNLL guidelines.

2.12.2 Fixed

- `DEPS` parsing was broken before and assumed that there was less information than is actually possible in the UD format. This means that now `deps` is a tuple with cardinality 4.

2.13 [0.3.1] - 2018-08-08

2.13.1 Fixed

- Fixed issue with submodules not being packaged in build

2.14 [0.3] - 2018-07-28

2.14.1 Added

- Ability to easily load CoNLL files from a network path (url)
- Some parsing validation. Before the error was not caught up front so the error could unexpectedly later show up.
- Sentence slicing had an issue before if either the start or end was omitted.
- More documentation and examples.
- Conll is now a `MutableSequence`, so it handles methods beyond its implementation as well as defined by python.

2.14.2 Fixed

- Some small bug fixes with parsing the token dicts.

2.15 [0.2.3] - 2018-07-23

2.15.1 Fixed

- Issues with documentation since docstrings were not in RST. Fixed by using napoleon sphinx extension

2.15.2 Added

- A little more docs
- More README info
- Better examples

2.16 [0.2.2] - 2018-07-18

2.16.1 Fixed

- Installation issues again with wheel when using `pip`.

2.17 [0.2.1] - 2018-07-18

2.17.1 Fixed

- Installation issues when using `pip`

2.18 [0.2] - 2018-07-16

2.18.1 Added

- More documentation
- Util package for convenient and common logic

2.19 [0.1.1] - 2018-07-15

2.19.1 Added

- Documentation which can be found [here](#).
- Small documentation changes on methods.

2.20 [0.1] - 2018-07-04

2.20.1 Added

- Everything. This is the first release of this package. The most notable absence is documentation which will be coming in a near-future release.

3.1 Overview

`pyconll` is a low level wrapper around the CoNLL-U format. This document explains how to quickly get started loading and manipulating CoNLL-U files within `pyconll`, and will go through a typical end-to-end scenario.

To install the library, run `pip install pyconll` from your python enlistment.

3.2 Loading CoNLL-U

To start, a CoNLL-U resource must be loaded, and `pyconll` can load from files, urls, and strings. Specific API information can be found in the [load](#) module documentation. Below is a typical example of loading a file on the local computer.

```
import pyconll

my_conll_file_location = './ud/train.conll'
train = pyconll.load_from_file(my_conll_file_location)
```

Loading methods usually return a `Conll` object, but some methods return an iterator over `Sentences` and do not load the entire `Conll` object into memory at once.

3.3 Traversing CoNLL-U

After loading a CoNLL-U file, we can traverse the CoNLL-U structure; `Conll` objects wrap `Sentences` and `Sentences` wrap `Token`. Here is what traversal normally looks like.

```
for sentence in train:
    for token in sentence:
```

(continues on next page)

(continued from previous page)

```
# Do work within loops
pass
```

Statistics such as lemmas for a certain closed class POS or number of non-projective punctuation dependencies can be computed through these loops. As an abstract example, we have defined some predicate, `sentence_pred`, and some transformation of noun tokens, `noun_token_transformation`, and we wish to transform all nouns in sentences that match our predicate, we can write the following.

```
for sentence in train:
    if sentence_pred(sentence):
        for token in sentence:
            if token.pos == 'NOUN':
                noun_token_transformation(token)
```

Note that most objects in `pyconll` are mutable, except for a select few fields, so changing the `Token` object remains with the `Sentence`.

3.4 Outputting CoNLL-U

Once you are done working with a `Conll` object, you may need to output your results. The object can be serialized back into the CoNLL-U format, through the `conll` method. `Conll`, `Sentence`, and `Token` objects are all `Conllable` which means they have a corresponding `conll` method which serializes the objects into the appropriate string representation.

3.5 Conclusion

`pyconll` allows for easy CoNLL-U loading, traversal, and serialization. Developers can define their own transformation or analysis of the loaded CoNLL-U data, and `pyconll` handles all the parsing and serialization logic. There are still some parts of the library that are not covered here such as the `Tree` data structure, loading files from network, and error handling, but the information on this page will get developers through the most important use cases.

This module defines the main interface to load CoNLL treebank resources. CoNLL treebanks can be loaded through a string, file, or network resource. CoNLL resources can be loaded and held in memory, or simply iterated through a sentence at a time which is useful in the case of very large files.

The fully qualified name of the module is `pyconll.load`, but these methods are imported at the `pyconll` namespace level.

4.1 Example

This example counts the number of times a token with a lemma of `linguistic` appeared in the treebank. If all the operations that will be done on the CoNLL file are readonly or are data aggregations, the `iter_from` alternatives are more memory efficient alternative as well. These methods will return an iterator over the sentences in the CoNLL resource rather than storing the CoNLL object in memory, which can be convenient when dealing with large files that do not need be completely loaded. This example uses the `load_from_file` method for illustration purposes.

```
import pyconll

example_treebank = '/home/myuser/englishdata.conll'
conll = pyconll.load_from_file(example_treebank)

count = 0
for sentence in conll:
    for word in sentence:
        if word.lemma == 'linguistic':
            count += 1

print(count)
```

4.2 API

A wrapper around the Conll class to easily load treebanks from multiple formats. This module can also load resources by iterating over treebank data without storing Conll objects in memory. This module is the main entrance to pyconll's functionalities.

`pyconll.load.iter_from_file(filename)`

Iterate over a CoNLL-U file's sentences.

Parameters `filename` – The name of the file whose sentences should be iterated over.

Yields The sentences that make up the CoNLL-U file.

Raises

- `IOError` if there is an error opening the file.
- `ParseError` – If there is an error parsing the input into a Conll object.

`pyconll.load.iter_from_string(source)`

Iterate over a CoNLL-U string's sentences.

Use this method if you only need to iterate over the CoNLL-U file once and do not need to create or store the Conll object.

Parameters `source` – The CoNLL-U string.

Yields The sentences that make up the CoNLL-U file.

Raises `ParseError` – If there is an error parsing the input into a Conll object.

`pyconll.load.iter_from_url(url)`

Iterate over a CoNLL-U file that is pointed to by a given URL.

Parameters `url` – The URL that points to the CoNLL-U file.

Yields The sentences that make up the CoNLL-U file.

Raises

- `requests.exceptions.RequestException` – If the url was unable to be properly retrieved.
- `ParseError` – If there is an error parsing the input into a Conll object.

`pyconll.load.load_from_file(filename)`

Load a CoNLL-U file given its location.

Parameters `filename` – The location of the file.

Returns A Conll object equivalent to the provided file.

Raises

- `IOError` – If there is an error opening the given filename.
- `ParseError` – If there is an error parsing the input into a Conll object.

`pyconll.load.load_from_string(source)`

Load the CoNLL-U source in a string into a Conll object.

Parameters `source` – The CoNLL-U formatted string.

Returns A Conll object equivalent to the provided source.

Raises `ParseError` – If there is an error parsing the input into a Conll object.

`pyconll.load.load_from_url(url)`

Load a CoNLL-U file at the provided URL.

Parameters `url` – The URL that points to the CoNLL-U file. This URL should be the actual CoNLL file and not an HTML page.

Returns A Conll object equivalent to the provided file.

Raises

- `requests.exceptions.RequestException` – If the url was unable to be properly retrieved and status was 4xx or 5xx.
- `ParseError` – If there is an error parsing the input into a Conll object.

The Token module represents a CoNLL token annotation. In a CoNLL file, this is a non-empty, non-comment line. Token members correspond directly with the [Universal Dependencies CoNLL definition](#) and all members are stored as strings. This means ids are strings as well. These fields are: `id`, `form`, `lemma`, `upos`, `xpos`, `feats`, `head`, `deprel`, `deps`, `misc`. More information on these is found below.

5.1 Fields

All fields are strings except for `feats`, `deps`, and `misc`, which are `dicts`. Each of these fields has specific semantics per the UDv2 guidelines. Since these fields are `dicts` these means modifying them uses python's natural syntax for dictionaries.

5.1.1 feats

`feats` is a key-value mapping from `str` to `set`. An example entry would be key `Gender` with value `set((Feminine,))`. More features could be added to an existing key by adding to its `set`, or a new feature could be added by adding to the dictionary. All features must have at least one value, so any keys with empty sets will throw an error on serialization back to text.

5.1.2 deps

`deps` is a key-value mapping from `str` to `tuple` of cardinality 4. This field represents enhanced dependencies. The key is the index of the token head, and the tuple elements define the enhanced dependency. Most Universal Dependencies treebanks, only use 2 of these 4 dimensions: the token index and the relation. See the Universal Dependencies guideline for more information on these 4 components. When adding new `deps`, the values must also be tuples of cardinality 4.

5.1.3 misc

For `misc`, the documentation only specifies that values be separated by a `'|'`, so not all keys have to have a value. So, values on `misc` are either `None`, or a set of `str`. A key with a value of `None` is output as a singleton, with no separating `'='`. A key with a corresponding set value will be handled like `feats`.

5.1.4 Examples

Below is an example of adding a new feature to a token, where the key must first be initialized:

```
token.feats['NewFeature'] = set(('No', ))
```

or alternatively as:

```
token.feats['NewFeature'] = set()
token.feats['NewFeature'].add('No')
```

On the miscellaneous column, adding a singleton field is done with the following line:

```
token.misc['SingletonFeature'] = None
```

5.2 API

Defines the Token type and parsing and output logic. A Token is the based unit in CoNLL-U and so the data and parsing in this module is central to the CoNLL-U format.

class `pyconll.unit.token.Token` (*source*, *empty=False*)

A token in a CoNLL-U file. This consists of 10 columns, each separated by a single tab character and ending in an LF (`'n'`) line break. Each of the 10 column values corresponds to a specific component of the token, such as id, word form, lemma, etc.

This class does not do any formatting validation on input or output. This means that invalid input may be properly processed and then output. Or that client changes to the token may result in invalid data that can then be output. Properly formatted CoNLL-U will always work on input and as long as all basic units are strings output will work as expected. The result may just not be proper CoNLL-U.

Also note that the word form for a token is immutable. This is because CoNLL-U is inherently interested in annotation schemes and not storing sentences.

__init__ (*source*, *empty=False*)

Construct a Token from the given source line.

A Token line ends in an LF line break according to the CoNLL-U specification. However, this method accepts a line with or without the LF line break.

On parsing, a `'_'` in the form and lemma is ambiguous and either refers to an empty value or to an actual underscore. The empty parameter flag controls how this situation should be handled.

This method also guarantees properly processing valid input, but invalid input may not be parsed properly. Some inputs that do not follow the CoNLL-U specification may still be parsed properly and as expected. So proper parsing is not an indication of validity.

Parameters

- **line** – The line that represents the Token in CoNLL-U format.

- **empty** – A flag to control if the word form and lemma can be assumed to be empty and not the token signifying empty. If both the form and lemma are underscores and empty is set to False (there is no empty assumption), then the form and lemma will be underscores rather than None.

Raises `ParseError` – On various parsing errors, such as not enough columns or improper column values.

conll()

Convert this Token to its CoNLL-U representation.

A Token's CoNLL-U representation is a line. Note that this method does not include a newline at the end.

Returns A string representing the Token in CoNLL-U format.

form

Provide the word form of this Token. This property is read only.

Returns The Token form.

is_multiword()

Checks if this Token is a multiword token.

Returns True if this token is a multiword token, and False otherwise.

The `Sentence` module represents an entire CoNLL sentence, which is composed of comments and tokens.

6.1 Comments

Comments are treated as key-value pairs, separated by the `=` character. A singleton comment has no `=` present. In this situation the key is the comment string, and the value is `None`. Methods for reading and writing comments on `Sentences` are prefixed with `meta_`, and are found below.

For convenience, the `id` and `text` comments are accessible through member properties on the `Sentence` in addition to metadata methods. So `sentence.id` and `sentence.meta_value('id')` are equivalent but the former is more concise and readable. Since this API does not support changing a token's form, the `text` comment cannot be changed. Text translations or transliterations can still be added just like any other comment.

6.1.1 Document and Paragraph ID

In previous versions of `pyconll`, the document and paragraph id of a `Sentence` were extracted similar to `text` and `id` information. This causes strange results and semantics when adding `Sentences` to a `Conll` object since the added sentence may have a `newpar` or `newdoc` comment which affects all subsequent `Sentence` ids. For simplicity's sake, this information is now only directly available as normal metadata information.

6.2 Tokens

This is the heart of the sentence. Tokens can be indexed on `Sentences` through their `id` value, as a string, or as a numeric index. So all of the following calls are valid, `sentence['5']`, `sentence['2-3']`, `sentence['2.1']`, and `sentence[2]`. Note that `sentence[x]` and `sentence[str(x)]` are not interchangeable. These calls are both valid but have different meanings.

6.3 API

Defines the Sentence type and the associated parsing and output logic.

class `pyconll.unit.sentence.Sentence` (*source*)

A sentence in a CoNLL-U file. A sentence consists of several components.

First, are comments. Each sentence must have two comments per UD v2 guidelines, which are `sent_id` and `text`. Comments are stored as a dict in the meta field. For singleton comments with no key-value structure, the value in the dict has a value of `None`.

Note the `sent_id` field is also assigned to the `id` property, and the `text` field is assigned to the `text` property for usability, and their importance as comments. The `text` property is read only along with the paragraph and document id. This is because the paragraph and document id are not defined per Sentence but across multiple sentences. Instead, these fields can be changed through changing the metadata of the Sentences.

Then comes the token annotations. Each sentence is made up of many token lines that provide annotation to the text provided. While a sentence usually means a collection of tokens, in this CoNLL-U sense, it is more useful to think of it as a collection of annotations with some associated metadata. Therefore the text of the sentence cannot be changed with this class, only the associated annotations can be changed.

__getitem__ (*key*)

Return the desired tokens from the Sentence.

Parameters *key* – The indicator for the tokens to return. Can either be an integer, a string, or a slice. For an integer, the numeric indexes of Tokens are used. For a string, the id of the Token is used. And for a slice the start and end must be the same data types, and can be both string and integer.

Returns If the key is a string then the appropriate Token. The key can also be a slice in which case a list of tokens is provided.

__init__ (*source*)

Construct a Sentence object from the provided CoNLL-U string.

Parameters *source* – The raw CoNLL-U string to parse. Comments must precede token lines.

Raises `ParseError` – If there is any token that was not valid.

__iter__ ()

Iterate through all the tokens in the Sentence including multiword tokens.

__len__ ()

Get the length of this sentence.

Returns The amount of tokens in this sentence. In the CoNLL-U sense, this includes both all the multiword tokens and their decompositions.

conll ()

Convert the sentence to a CoNLL-U representation.

Returns A string representing the Sentence in CoNLL-U format.

id

Get the sentence id.

Returns The sentence id. If there is none, then returns `None`.

meta_present (*key*)

Check if the key is present as a singleton or as a pair.

Parameters *key* – The value to check for in the comments.

Returns True if the key was provided as a singleton or as a key value pair. False otherwise.

meta_value (*key*)

Returns the value associated with the key in the metadata (comments).

Parameters **key** – The key whose value to look up.

Returns The value associated with the key as a string. If the key is a singleton then None is returned.

Raises `KeyError` – If the key is not present in the comments.

remove_meta (*key*)

Remove a metadata element associated with the Sentence.

Parameters **key** – The name of the metadata / comment.

Raises `KeyError` – If the key is not present in the Sentence metadata.

set_meta (*key, value=None*)

Set or add the metadata or comments associated with this Sentence.

Parameters

- **key** – The key for the comment.
- **value** – The value to associate with the key. If the comment is a singleton, this field can be ignored or set to None.

text

Get the continuous text for this sentence. Read-only.

Returns The continuous text of this sentence. If none is provided in comments, then None is returned.

to_tree ()

Creates a Tree data structure from the current sentence.

An empty sentence will create a Tree with no data and no children.

Returns A constructed Tree that represents the dependency graph of the sentence.

This module represents a CoNLL file: a collection of CoNLL annotated sentences. Users should use the `load` module to create CoNLL objects rather than directly using the class constructor. The `Conll` object is a wrapper around a list of sentences that can be serialized into a CoNLL format, i.e. it is `Conllable`.

`Conll` is a subclass of `MutableSequence`, so `append`, `reverse`, `extend`, `pop`, `remove`, and `__iadd__` are available free of charge, even though they are not defined below. This information can be found on the [collections](#) documentation.

7.1 API

Defines the `Conll` type and the associated parsing and output logic.

class `pyconll.unit.conll.Conll` (*it*)

The abstraction for a CoNLL-U file. A CoNLL-U file is more or less just a collection of sentences in order. These sentences are accessed by numeric index. Note that sentences must be separated by whitespace. CoNLL-U also specifies that the file must end in a new line but that requirement is relaxed here in parsing.

`__contains__` (*other*)

Check if the `Conll` object has this sentence.

Parameters *other* – The sentence to check for.

Returns True if this Sentence is exactly in the `Conll` object. False, otherwise.

`__delitem__` (*key*)

Delete the Sentence corresponding with the given key.

Parameters *key* – The info to get the Sentence to delete. Can be the integer position in the file, or a slice.

`__getitem__` (*key*)

Index a sentence by key value.

Parameters *key* – The key to index the sentence by. This key can either be a numeric key, or a slice.

Returns The corresponding sentence if the key is an int or the sentences if the key is a slice in the form of another Conll object.

Raises `TypeError` – If the key is not an integer or slice.

__init__ (*it*)

Create a CoNLL-U file collection of sentences.

Parameters *it* – An iterator of the lines of the CoNLL-U file.

Raises `ParseError` – If there is an error constructing the sentences in the iterator.

__iter__ ()

Allows for iteration over every sentence in the CoNLL-U file.

Yields An iterator over the sentences in this Conll object.

__len__ ()

Returns the number of sentences in the CoNLL-U file.

Returns The size of the CoNLL-U file in sentences.

__setitem__ (*key, sent*)

Set the given location to the Sentence.

Parameters *key* – The location in the Conll file to set to the given sentence. This only accepts integer keys and accepts negative indexing.

conll ()

Output the Conll object to a CoNLL-U formatted string.

Returns The CoNLL-U object as a string. This string will end in a newline.

insert (*index, value*)

Insert the given sentence into the given location.

This function behaves in the same way as python lists insert.

Parameters

- **index** – The numeric index to insert the sentence into.
- **value** – The sentence to insert.

write (*writable*)

Write the Conll object to something that is writable.

For simply writing, this method is more efficient than calling `conll` than writing since no string of the entire Conll object is created. The final output will include a final newline.

Parameters *writable* – The writable object such as a file. Must have a write method.

`Tree` is a very basic immutable tree class. A `Tree` can have multiple children and has one parent. The parent and child of a tree are established when a `Tree` is created. Accessing the data on a `Tree` can be done through the `data` member. A `Tree` is created through the `TreeBuilder` module which is an internal API in `pyconll`. A `Tree`'s use in `pyconll` is when creating a `Tree` structure from a `Sentence` object in the `sentencetree` module.

8.1 API

A general immutable tree module. This module is used when parsing a serial sentence into a `Tree` structure.

class `pyconll.tree.tree.Tree`

A tree node. This is the base representation for a tree, which can have many children which are accessible via child index. The tree's structure is immutable, so the data, parent, children cannot be changed once created.

As is this class is useless, and must be created with the `TreeBuilder` module which is a sort of friend class of `Tree` to maintain its immutable public contract.

__getitem__ (*key*)

Get specific children from the `Tree`. This can be an integer or slice.

Parameters *key* – The indexer for the item.

__init__ ()

Create a new empty tree. To create a useful `Tree`, use `TreeBuilder`.

__iter__ ()

Provides an iterator over the children.

__len__ ()

Provides the number of direct children on the tree.

Returns The number of direct children on the tree.

data

The data on the tree node. The property ensures it is readonly.

Returns The data stored on the Tree. No data is represented as None.

parent

Provides the parent of the Tree. The property ensures it is readonly.

Returns A pointer to the parent Tree reference. None if there is no parent.

This module provides additional, common methods that build off of the API layer. This module simply adds logic, rather than extending the API. Right now this module is pretty sparse, but will be extended as needed.

9.1 API

A set of utilities for dealing with pyconll defined types. This is simply a collection of functions.

`pyconll.util.find_ngrams(conll, ngram, case_sensitive=True)`

Find the occurrences of the ngram in the provided Conll collection.

This method returns every sentence along with the token position in the sentence that starts the ngram. The matching algorithm does not currently account for multiword tokens, so “don’t” should be separated into “do” and “not” in the input.

Parameters

- **sentence** – The sentence in which to search for the ngram.
- **ngram** – The ngram to search for. A random access iterator.
- **case_sensitive** – Flag to indicate if the ngram search should be case sensitive. The case insensitive comparison currently is locale insensitive lowercase comparison.

Returns An iterator of tuples over the ngrams in the Conll object. The first element is the sentence, the second element is the numeric token index, and the last element is the actual list of tokens references from the sentence. This list does not include any multiword token that were skipped over.

`pyconll.util.find_nonprojective_deps(sentence)`

Find the nonprojective dependency pairs in the provided sentence.

Dependencies are provided as a list of ordered pairs. Each ordered pair represents a non-projective dependency pair. Each element in the ordered pair is a token, that makes a dependency with its governor. So each token is the base of its dependency, and the two tokens’ dependencies cross in a non projective way.

Parameters **sentence** – The sentence to check for nonprojective dependency pairs.

Returns An iterable of pairs which represent the children of a nonprojective dependency pair.

`Conllable` marks a class that can be output as a CoNLL formatted string. `Conllable` classes implement a `conll` method.

10.1 API

Holds the `Conllable` interface, which is a marker interface to show that a class is a `Conll` object, such as a `treebank`, `sentence`, or `token`, and therefore has a `conll` method.

class `pyconll.conllable.Conllable`

A `Conllable` mixin to indicate that the component can be converted into a CoNLL representation.

conll ()

Provides a `conll` representation of the component.

Returns A string `conll` representation of the base component.

Raises `NotImplementedError` – If the child class does not implement the method.

Custom exceptions for `pyconll`. These errors are a `ParseError` and a `FormatError`.

A `ParseError` occurs when the source input to a CoNLL component is invalid, and a `FormatError` occurs when the internal state of the component is invalid, and the component cannot be output to a CoNLL string.

11.1 API

Holds custom `pyconll` errors. These errors include parsing errors when reading treebanks, and errors when outputting CoNLL objects.

exception `pyconll.exception.FormatError`

Error that results from trying to format a CoNLL object to a string.

exception `pyconll.exception.ParseError`

Error that results from an improper value into a parsing routine.

Welcome to the `pyconll` documentation homepage. Module documentation, changelogs, and guidance pages are listed above in the table of contents. Those unsure where to start can see the [load](#), [conll](#), [sentence](#), and [token](#) pages which contain documentation for the base CoNLL-U data types. There is also the [Getting Started](#) page which goes through an end-to-end example of using `pyconll`.

The [github](#) homepage has a limited set of examples, tests, and the source.

p

- `pyconll.conllable`, 33
- `pyconll.exception`, 35
- `pyconll.load`, 16
- `pyconll.tree.tree`, 29
- `pyconll.unit.conll`, 27
- `pyconll.unit.sentence`, 24
- `pyconll.unit.token`, 20
- `pyconll.util`, 31

Symbols

`__contains__()` (pyconll.unit.conll.Conll method), 27
`__delitem__()` (pyconll.unit.conll.Conll method), 27
`__getitem__()` (pyconll.tree.tree.Tree method), 29
`__getitem__()` (pyconll.unit.conll.Conll method), 27
`__getitem__()` (pyconll.unit.sentence.Sentence method), 24

`__init__()` (pyconll.tree.tree.Tree method), 29
`__init__()` (pyconll.unit.conll.Conll method), 28
`__init__()` (pyconll.unit.sentence.Sentence method), 24
`__init__()` (pyconll.unit.token.Token method), 20
`__iter__()` (pyconll.tree.tree.Tree method), 29
`__iter__()` (pyconll.unit.conll.Conll method), 28
`__iter__()` (pyconll.unit.sentence.Sentence method), 24
`__len__()` (pyconll.tree.tree.Tree method), 29
`__len__()` (pyconll.unit.conll.Conll method), 28
`__len__()` (pyconll.unit.sentence.Sentence method), 24
`__setitem__()` (pyconll.unit.conll.Conll method), 28

C

Conll (class in pyconll.unit.conll), 27
`conll()` (pyconll.conllable.Conllable method), 33
`conll()` (pyconll.unit.conll.Conll method), 28
`conll()` (pyconll.unit.sentence.Sentence method), 24
`conll()` (pyconll.unit.token.Token method), 21
Conllable (class in pyconll.conllable), 33

D

data (pyconll.tree.tree.Tree attribute), 29

F

`find_ngrams()` (in module pyconll.util), 31
`find_nonprojective_deps()` (in module pyconll.util), 31
form (pyconll.unit.token.Token attribute), 21
FormatError, 35

I

id (pyconll.unit.sentence.Sentence attribute), 24
`insert()` (pyconll.unit.conll.Conll method), 28

`is_multiword()` (pyconll.unit.token.Token method), 21
`iter_from_file()` (in module pyconll.load), 16
`iter_from_string()` (in module pyconll.load), 16
`iter_from_url()` (in module pyconll.load), 16

L

`load_from_file()` (in module pyconll.load), 16
`load_from_string()` (in module pyconll.load), 16
`load_from_url()` (in module pyconll.load), 16

M

`meta_present()` (pyconll.unit.sentence.Sentence method), 24
`meta_value()` (pyconll.unit.sentence.Sentence method), 25

P

parent (pyconll.tree.tree.Tree attribute), 30
ParseError, 35
pyconll.conllable (module), 33
pyconll.exception (module), 35
pyconll.load (module), 16
pyconll.tree.tree (module), 29
pyconll.unit.conll (module), 27
pyconll.unit.sentence (module), 24
pyconll.unit.token (module), 20
pyconll.util (module), 31

R

`remove_meta()` (pyconll.unit.sentence.Sentence method), 25

S

Sentence (class in pyconll.unit.sentence), 24
`set_meta()` (pyconll.unit.sentence.Sentence method), 25

T

text (pyconll.unit.sentence.Sentence attribute), 25
`to_tree()` (pyconll.unit.sentence.Sentence method), 25

Token (class in `pyconll.unit.token`), [20](#)

Tree (class in `pyconll.tree.tree`), [29](#)

W

`write()` (`pyconll.unit.conll.Conll` method), [28](#)