
htmlmin Documentation

Release 0.1

Dave Mankoff

Sep 27, 2017

Contents

1	Quickstart	3
1.1	Command Line	4
2	Tutorial & Examples	7
3	API Reference	9
3.1	Main Functions	9
3.2	WSGI Middleware	10
3.3	Decorator	11
4	Install	13
5	Source Code	15
6	Features	17
7	Indices and tables	19

An HTML Minifier with Seatbelts

For single invocations, there is the `htmlmin.minify` method. It takes input html as a string for its first argument and returns minified html. It accepts multiple different options that allow you to tune the amount of minification being done, with the defaults being the safest available options:

```
>>> import htmlmin
>>> input_html = '''
    <body style="background-color: tomato;">
    <h1> htmlmin rocks</h1>
    <pre>
    and rolls
    </pre>
</body>'''
>>> htmlmin.minify(input_html)
u' <body style="background-color: tomato;" <h1> htmlmin rocks</h1> <pre>\n
↳and rolls\n      </pre> </body>'
>>> print htmlmin.minify(input_html)
<body style="background-color: tomato;" <h1> htmlmin rocks</h1> <pre>
    and rolls
</pre> </body>
```

If there is a chunk of html which you do not want minified, put a `pre` attribute on an HTML tag that wraps it. `htmlmin` will leave the contents of the tag alone and will remove the `pre` attribute before it is output:

```
>>> import htmlmin
>>> input_html = '''<span> minified </span><span pre> not minified </span>'''
>>> htmlmin.minify(input_html)
u'<span> minified </span><span> not minified </span>'
```

The `minify` function works well for one off minifications. However, if you are going to minify several pieces of HTML, the `Minifier` class is provided. It works similarly, but allows for persistence of options between invocations and recycles the internal data structures used for minification.

Command Line

htmlmin is invoked by running:

```
htmlmin input.html output.html
```

If no output file is specified, it will print to `stdout`. If not input specified, it reads from `stdin`. Help with options can be retrieved at any time by running `htmlmin -h`:

```
htmlmin -h
usage: htmlmin [-h] [-c] [-s] [--remove-all-empty-space] [-H] [-k] [-p [TAG [TAG ...
↳]]] [-e ENCODING]
           [INPUT] [OUTPUT]

Minify HTML

positional arguments:
  INPUT                File path to html file to minify. Defaults to stdin.
  OUTPUT              File path to output to. Defaults to stdout.

optional arguments:
  -h, --help          show this help message and exit
  -c, --remove-comments
↳individual basis    When set, comments will be removed. They can be kept on an
↳be removed from     by starting them with a '!': <!--! comment -->. The '!' will
↳of your comment,    the final output. If you want a '!' as the leading character
↳                    put two of them: <!--!! comment -->.

  -s, --remove-empty-space
↳cases.              When set, this removes empty space between tags in certain
↳a newline           Specifically, it will remove empty space if and only if there
↳such as             character occurs within the space. Thus, code like
                    '<span>x</span> <span>y</span>' will be left alone, but code
                    '
                    </head>
                    <body>
                    ...'
↳your                will become '...</head><body>...'. Note that this CAN break
↳caution.           html if you spread two inline tags over two lines. Use with

  --remove-all-empty-space
↳this can and        When set, this removes ALL empty space between tags. WARNING:
↳</i> <i>Y</i>'       likely will cause unintended consequences. For instance, '<i>X
↳other text will     will become '<i>X</i><i>Y</i>'. Putting whitespace along with
↳result. Whitespace is avoid this problem. Only use if you are confident in the
↳                    not removed from inside of tags, thus '<span> </span>' will
↳be left alone.
```



```
-H, --in-head          If you are parsing only a fragment of HTML, and the fragment
↳occurs in the        head of the document, setting this will remove some extra
↳whitespace.

-k, --keep-pre-attr    HTMLMin supports the proprietary attribute 'pre' that can be
↳added to elements    to prevent minification. This attribute is removed by default.
↳ Set this flag to    keep the 'pre' attributes in place.

-a PRE_ATTR, --pre-attr PRE_ATTR
↳it should not        The attribute htmlmin looks for to find blocks of HTML that
↳k' is                minify. This attribute will be removed from the HTML unless '-
                    specified. Defaults to 'pre'.

-p [TAG [TAG ...]], --pre-tags [TAG [TAG ...]]
↳left unminified.    By default, the contents of 'pre', and 'textarea' tags are
↳'script' and 'style' You can specify different tags using the --pre-tags option.
                    tags are always left unminified.

-e ENCODING, --encoding ENCODING
                    Encoding to read and write with. Default 'utf-8'.
```


CHAPTER 2

Tutorial & Examples

Coming soon...

Main Functions

`htmlmin.minify` (*input*, *remove_comments=False*, *remove_empty_space=False*, *remove_all_empty_space=False*, *reduce_boolean_attributes=False*, *keep_pre=False*, *pre_tags=(u'pre', u'textarea')*, *pre_attr='pre'*)

Minifies HTML in one shot.

Parameters

- **input** – A string containing the HTML to be minified.
- **remove_comments** – Remove comments found in HTML. Individual comments can be maintained by putting a ! as the first character inside the comment. Thus:

```
<!-- FOO --> <!--! BAR -->
```

Will become simply:

```
<!-- BAR -->
```

The added exclamation is removed.

- **remove_empty_space** – Remove empty space found in HTML between an opening and a closing tag and when it contains a newline or carriage return. If whitespace is found that is only spaces and/or tabs, it will be turned into a single space. Be careful, this can have unintended consequences.
- **remove_all_empty_space** – A more extreme version of `remove_empty_space`, this removes all empty whitespace found between tags. This is almost guaranteed to break your HTML unless you are very careful. nothing
- **reduce_boolean_attributes** – Where allowed by the HTML5 specification, attributes such as 'disabled' and 'readonly' will have their value removed, so 'disabled="true"' will simply become 'disabled'. This is generally a good option to turn on except when JavaScript relies on the values.

- **keep_pre** – By default, htmlmin uses the special attribute `pre` to allow you to demarcate areas of HTML that should not be minified. It removes this attribute as it finds it. Setting this value to `True` tells htmlmin to leave the attribute in the output.
- **pre_tags** – A list of tag names that should never be minified. You are free to change this list as you see fit, but you will probably want to include `pre` and `textarea` if you make any changes to the list. Note that `<script>` and `<style>` tags are never minimized.
- **pre_attr** – Specifies the attribute that, when found in an HTML tag, indicates that the content of the tag should not be minified. Defaults to `pre`.

Returns A string containing the minified HTML.

If you are going to be minifying multiple HTML documents, each with the same settings, consider using `Minifier`.

```
class htmlmin.Minifier(remove_comments=False, remove_empty_space=False, re-
                      move_all_empty_space=False, reduce_boolean_attributes=False,
                      keep_pre=False, pre_tags=(u'pre', u'textarea'), pre_attr='pre')
```

An object that supports HTML Minification.

Options are passed into this class at initialization time and are then persisted across each use of the instance. If you are going to be minifying multiple peices of HTML, this will be more efficient than using `htmlmin.minify`.

See `htmlmin.minify` for an explanation of options.

minify (*input)

Runs HTML through the minifier in one pass.

Parameters **input** – HTML to be fed into the minimizer. Multiple chunks of HTML can be provided, and they are fed in sequentially as if they were concatenated.

Returns A string containing the minified HTML.

This is the simplest way to use an existing `Minifier` instance. This method takes in HTML and minifies it, returning the result. Note that this method resets the internal state of the parser before it does any work. If there is pending HTML in the buffers, it will be lost.

input (*input)

Feed more HTML into the input stream

Parameters **input** – HTML to be fed into the minimizer. Multiple chunks of HTML can be provided, and they are fed in sequentially as if they were concatenated. You can also call this method multiple times to achieve the same effect.

output

Retrieve the minified output generated thus far.

finalize ()

Finishes current input HTML and returns mininified result.

This method flushes any remaining input HTML and returns the minified result. It resets the state of the internal parser in the process so that new HTML can be minified. Be sure to call this method before you reuse the `Minifier` instance on a new HTML document.

WSGI Middleware

```
class htmlmin.middleware.HTMLMinMiddleware(app, by_default=True, keep_header=False, de-
                                         bug=False, **kwargs)
```

WSGI Middleware that minifies html on the way out.

Parameters

- **by_default** – Specifies if minification should be turned on or off by default. Defaults to `True`.
- **keep_header** – The middleware recognizes one custom HTTP header that can be used to turn minification on or off on a per-request basis: `X-HTML-Min-Enable`. Setting the header to `true` will turn minification on; anything else will turn minification off. If `by_default` is set to `False`, this header is how you would turn minification back on. The middleware, by default, removes the header from the output. Setting this to `True` leaves the header in tact.
- **debug** – A quick setting to turn all minification off. The middleware is effectively bypassed.

This simple middleware minifies any HTML content that passes through it. Any additional keyword arguments beyond the three settings the middleware has are passed on to the internal minifier. The documentation for the options can be found under [htmlmin.minify](#).

Decorator

`htmlmin.decorator.htmlmin(*args, **kwargs)`

Minifies HTML that is returned by a function.

A simple decorator that minifies the HTML output of any function that it decorates. It supports all the same options that [htmlmin.minify](#) has. With no options, it uses `minify`'s default settings:

```
@htmlmin
def foobar():
    return '    minify me!    '
```

or:

```
@htmlmin(remove_comments=True)
def foobar():
    return '    minify me! <!-- and remove me! -->'
```

`htmlmin` is an HTML minifier that just works. It comes with safe defaults and an easily configurable set options. It can turn this:

```
<html>
  <head>
    <title> Hello, World! </title>
  </head>
  <body>
    <p> How are <em>you</em> doing? </p>
  </body>
</html>
```

Into this:

```
<html><head><title>Hello, World!</title><body><p> How are <em>you</em> doing? </p></
↳body></html>
```

When we say that `htmlmin` has 'seatbelts', what we mean is that it comes with features that you can use to safely minify beyond the defaults, but you have to put them in yourself. For instance, by default, `htmlmin` will never minimize the content between `<pre>`, `<textarea>`, `<script>`, and `<style>` tags. You can also explicitly tell it to not minify

additional tags either globally name or by adding the custom `pre` attribute to a tag in your HTML. `htmlmin` will remove the `pre` attributes as it parses your HTML automatically.

It also includes a command-line tool for easy invocation and integration with existing workflows.

CHAPTER 4

Install

To install via pip:

```
pip install htmlmin
```


CHAPTER 5

Source Code

Source code is available on github at <https://github.com/manky/htmlmin>:

```
git clone git://github.com/manky/htmlmin.git
```


CHAPTER 6

Features

- Safely minify HTML with either a *function call* or from the *command line*.
- Extend what elements can and cannot be minified.
- Intelligently remove whitespace completely or reduce to single spaces.
- Properly handles unclosed HTML5 tags.
- Optionally remove comments while marking some comments to keep.
- Simple function *decorator* to minify all function output.
- Simple *WSGI middleware* to minify web app output.
- Tested in both Python 2.7 and 3.2:

CHAPTER 7

Indices and tables

- `genindex`
- `search`

F

finalize() (htmlmin.Minifier method), 10

H

htmlmin() (in module htmlmin.decorator), 11

HTMLMinMiddleware (class in htmlmin.middleware),
10

I

input() (htmlmin.Minifier method), 10

M

Minifier (class in htmlmin), 10

minify() (htmlmin.Minifier method), 10

minify() (in module htmlmin), 9

O

output (htmlmin.Minifier attribute), 10