

Last Modified: Sat May 14 16:33:16 UTC 2011

# PDFMiner

---

Python PDF parser and analyzer

[Homepage](#) [Recent Changes](#) [PDFMiner API](#)

- [What's It?](#)
- [Download](#)
- [Where to Ask](#)
- [How to Install](#)
  - [CJK languages support](#)
- [Command Line Tools](#)
  - [pdf2txt.py](#)
  - [dumppdf.py](#)
  - [PDFMiner API](#)
- [Changes](#)
- [TODO](#)
- [Related Projects](#)
- [Terms and Conditions](#)

## What's It?

---

PDFMiner is a tool for extracting information from PDF documents. Unlike other PDF-related tools, it focuses entirely on getting and analyzing text data. PDFMiner allows one to obtain the exact location of text in a page, as well as other information such as fonts or lines. It includes a PDF converter that can transform PDF files into other text formats (such as HTML). It has an extensible PDF parser that can be used for other purposes than text analysis.

## Features

- Written entirely in Python. (for version 2.4 or newer)
- Parse, analyze, and convert PDF documents.
- PDF-1.7 specification support. (well, almost)
- CJK languages and vertical writing scripts support.
- Various font types (Type1, TrueType, Type3, and CID) support.
- Basic encryption (RC4) support.
- PDF to HTML conversion (with a sample converter web app).
- Outline (TOC) extraction.
- Tagged contents extraction.
- Reconstruct the original layout by grouping text chunks.

PDFMiner is about 20 times slower than other C/C++-based counterparts such as XPDF.

**Online Demo:** (pdf -> html conversion webapp)

<http://pdf2html.tabesugi.net:8080/>

## Download

**Source distribution:**

<http://pypi.python.org/pypi/pdfminer/>

**github:**

<https://github.com/euske/pdfminer/>

## Where to Ask

**Questions and comments:**

<http://groups.google.com/group/pdfminer-users/>

## How to Install

---

1. Install [Python](#) 2.4 or newer. (**Python 3 is not supported.**)

2. Download the [PDFMiner source](#).
3. Unpack it.
4. Run `setup.py` to install:

```
# python setup.py install
```

5. Do the following test:

```
$ pdf2txt.py samples/simple1.pdf
Hello

World

Hello

World

H e l l o

W o r l d

H e l l o

W o r l d
```

6. Done!

## For CJK languages

In order to process CJK languages, you need an additional step to take during installation:

```
# make cmap
python tools/conv_cmap.py pdfminer/cmap Adobe-CNS1 cmapsrc/cid2code_Adobe_CNS1.txt
reading 'cmapsrc/cid2code_Adobe_CNS1.txt'...
writing 'CNS1_H.py'...
...
(this may take several minutes)

# python setup.py install
```

On Windows machines which don't have `make` command, paste the following commands on a command line prompt:

```
python tools\conv_cmap.py pdfminer\cmap Adobe-CNS1 cmapsrc\cid2code_Adobe_CNS1.txt
python tools\conv_cmap.py pdfminer\cmap Adobe-GB1 cmapsrc\cid2code_Adobe_GB1.txt c
python tools\conv_cmap.py pdfminer\cmap Adobe-Japan1 cmapsrc\cid2code_Adobe_Japan1
python tools\conv_cmap.py pdfminer\cmap Adobe-Korea1 cmapsrc\cid2code_Adobe_Korea1
python setup.py install
```

## Command Line Tools

---

PDFMiner comes with two handy tools: `pdf2txt.py` and `dumppdf.py`.

### pdf2txt.py

`pdf2txt.py` extracts text contents from a PDF file. It extracts all the text that are to be rendered programmatically, i.e. text represented as ASCII or Unicode strings. It cannot recognize text drawn as images that would require optical character recognition. It also extracts the corresponding locations, font names, font sizes, writing direction (horizontal or vertical) for each text portion. You need to provide a password for protected PDF documents when its access is restricted. You cannot extract any text from a PDF document which does not have extraction permission.

**Note:** Not all characters in a PDF can be safely converted to Unicode.

### Examples

```
$ pdf2txt.py -o output.html samples/naacl06-shinyama.pdf
```

(extract text as an HTML file whose filename is output.html)

```
$ pdf2txt.py -V -c euc-jp -o output.html samples/jo.pdf
```

(extract a Japanese HTML file in vertical writing, CMap is required)

```
$ pdf2txt.py -P mypassword -o output.txt secret.pdf
```

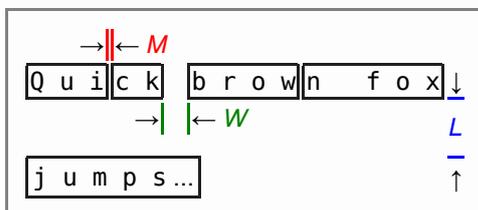
(extract a text from an encrypted PDF file)

## Options

- o *filename*  
Specifies the output file name. By default, it prints the extracted contents to stdout in text format.
- p *pageno[,pageno,...]*  
Specifies the comma-separated list of the page numbers to be extracted. Page numbers start at one. By default, it extracts text from all the pages.
- c *codec*  
Specifies the output codec.
- t *type*  
Specifies the output format. The following formats are currently supported.
  - text : TEXT format. (Default)
  - html : HTML format. Not recommended for extraction purposes because the markup is messy.
  - xml : XML format. Provides the most information.
  - tag : "Tagged PDF" format. A tagged PDF has its own contents annotated with HTML-like tags. pdf2txt tries to extract its content streams rather than inferring its text locations. Tags used here are defined in the PDF specification (See §10.7 "Tagged PDF").
- I *image\_directory*  
Specifies the output directory for image extraction. Currently only JPEG images are supported.
- M *char\_margin*
- L *line\_margin*
- W *word\_margin*

These are the parameters used for layout analysis. In an actual PDF file, text portions might be split into several chunks in the middle of its running, depending on the authoring software. Therefore, text extraction needs to splice text chunks. In the figure below, two text chunks whose distance is closer than the *char\_margin* (shown as *M*) is considered continuous and get grouped into one. Also, two lines whose distance is closer than the *line\_margin* (*L*) is grouped as a text box, which is a rectangular area that contains a "cluster" of text portions. Furthermore, it may be required to insert blank characters (spaces) as necessary if the distance between two words is greater than the *word\_margin* (*W*), as a blank between words might not be represented as a space, but indicated by the positioning of each word.

Each value is specified not as an actual length, but as a proportion of the length to the size of each character in question. The default values are  $M = 1.0$ ,  $L = 0.3$ , and  $W = 0.2$ , respectively.



- C  
Suppress object caching. This will reduce the memory consumption but also slows down the process.
- n  
Suppress layout analysis.
- A  
Forces to perform layout analysis for all the text strings, including text contained in figures.
- V

Allows vertical writing detection.

**-Y *layout\_mode***

Specifies how the page layout should be preserved. (Currently only applies to HTML format.)

- `exact` : preserve the exact location of each individual character (a large and messy HTML).
- `normal` : preserve the location and line breaks in each text block. (Default)
- `loose` : preserve the overall location of each text block.

**-s *scale***

Specifies the output scale. Can be used in HTML format only.

**-m *maxpages***

Specifies the maximum number of pages to extract. By default, it extracts all the pages in a document.

**-P *password***

Provides the user password to access PDF contents.

**-d**

Increases the debug level.

---

## dumppdf.py

`dumppdf.py` dumps the internal contents of a PDF file in pseudo-XML format. This program is primarily for debugging purposes, but it's also possible to extract some meaningful contents (such as images).

### Examples

```
$ dumppdf.py -a foo.pdf
(dump all the headers and contents, except stream objects)
```

```
$ dumppdf.py -T foo.pdf
(dump the table of contents)
```

```
$ dumppdf.py -r -i6 foo.pdf > pic.jpeg
(extract a JPEG image)
```

### Options

**-a**

Instructs to dump all the objects. By default, it only prints the document trailer (like a header).

**-i *objno, objno, ...***

Specifies PDF object IDs to display. Comma-separated IDs, or multiple `-i` options are accepted.

**-p *pageno, pageno, ...***

Specifies the page number to be extracted. Comma-separated page numbers, or multiple `-p` options are accepted. Note that page numbers start at one, not zero.

**-r (raw)**

**-b (binary)**

**-t (text)**

Specifies the output format of stream contents. Because the contents of stream objects can be very large, they are omitted when none of the options above is specified.

With `-r` option, the "raw" stream contents are dumped without decompression. With `-b` option, the decompressed contents are dumped as a binary blob. With `-t` option, the decompressed contents are dumped in a text format, similar to `repr()` manner. When `-r` or `-b` option is given, no stream header is displayed for the ease of saving it to a file.

**-T**

Shows the table of contents.

**-P *password***

Provides the user password to access PDF contents.

-d

Increases the debug level.

## Changes

---

- 2010/05/15: Speed improvements for layout analysis.
- 2010/05/15: API changes. `LText.get_text()` is added.
- 2010/04/20: API changes. `LTPolygon` class was renamed as `LTCurve`.
- 2010/04/20: `LLine` now represents horizontal/vertical lines only. Thanks to Koji Nakagawa.
- 2010/03/07: Documentation improvements by Jakub Wilk. Memory usage patch by Jonathan Hunt.
- 2010/02/27: Bugfixes and layout analysis improvements. Thanks to fujimoto.report.
- 2010/12/26: A couple of bugfixes and minor improvements. Thanks to Kevin Brubeck Unhammer and Daniel Gerber.
- 2010/10/17: A couple of bugfixes and minor improvements. Thanks to standardabweichung and Alastair Irving.
- 2010/09/07: A minor bugfix. Thanks to Alexander Garden.
- 2010/08/29: A couple of bugfixes. Thanks to Sahan Malagi, pk, and Humberto Pereira.
- 2010/07/06: Minor bugfixes. Thanks to Federico Brega.
- 2010/06/13: Bugfixes and improvements on CMap data compression. Thanks to Jakub Wilk.
- 2010/04/24: Bugfixes and improvements on TOC extraction. Thanks to Jose Maria.
- 2010/03/26: Bugfixes. Thanks to Brian Berry and Lubos Pintes.
- 2010/03/22: Improved layout analysis. Added regression tests.
- 2010/03/12: A couple of bugfixes. Thanks to Sean Manefield.
- 2010/02/27: Changed the way of internal layout handling. (`LTextItem` -> `LChar`)
- 2010/02/15: Several bugfixes. Thanks to Sean.
- 2010/02/13: Bugfix and enhancement. Thanks to André Auzi.
- 2010/02/07: Several bugfixes. Thanks to Hiroshi Manabe.
- 2010/01/31: JPEG image extraction supported. Page rotation bug fixed.
- 2010/01/04: Python 2.6 warning removal. More doctest conversion.
- 2010/01/01: CMap bug fix. Thanks to Winfried Plappert.
- 2009/12/24: `RunLengthDecode` filter added. Thanks to Troy Bollinger.
- 2009/12/20: Experimental polygon shape extraction added. Thanks to Yusuf Dewaswala for reporting.
- 2009/12/19: CMap resources are now the part of the package. Thanks to Adobe for open-sourcing them.
- 2009/11/29: Password encryption bug fixed. Thanks to Yannick Gingras.
- 2009/10/31: SGML output format is changed and renamed as XML.
- 2009/10/24: Charspace bug fixed. Adjusted for 4-space indentation.
- 2009/10/04: Another matrix operation bug fixed. Thanks to Vitaly Sedelnik.
- 2009/09/12: Fixed rectangle handling. Able to extract image boundaries.
- 2009/08/30: Fixed page rotation handling.
- 2009/08/26: Fixed zlib decoding bug. Thanks to Shon Urbas.
- 2009/08/24: Fixed a bug in character placing. Thanks to Pawan Jain.
- 2009/07/21: Improvement in layout analysis.
- 2009/07/11: Improvement in layout analysis. Thanks to Lubos Pintes.
- 2009/05/17: Bugfixes, massive code restructuring, and simple graphic element support added. `setup.py` is supported.
- 2009/03/30: Text output mode added.
- 2009/03/25: Encoding problems fixed. Word splitting option added.
- 2009/02/28: Robust handling of corrupted PDFs. Thanks to Troy Bollinger.
- 2009/02/01: Various bugfixes. Thanks to Hiroshi Manabe.
- 2009/01/17: Handling a trailer correctly that contains both `/XrefStm` and `/Prev` entries.
- 2009/01/10: Handling Type3 font metrics correctly.
- 2008/12/28: Better handling of word spacing. Thanks to Christian Nentwich.
- 2008/09/06: A sample pdf2html webapp added.
- 2008/08/30: ASCII85 encoding filter support.
- 2008/07/27: Tagged contents extraction support.
- 2008/07/10: Outline (TOC) extraction support.
- 2008/06/29: HTML output added. Reorganized the directory structure.
- 2008/04/29: Bugfix for Win32. Thanks to Chris Clark.
- 2008/04/27: Basic encryption and LZW decoding support added.
- 2008/01/07: Several bugfixes. Thanks to Nick Fabry for his vast contribution.
- 2007/12/31: Initial release.
- 2004/12/24: Start writing the code out of boredom...

## TODO

---

- [PEP-8](#) and [PEP-257](#) conformance.
- Better documentation.
- Better text extraction / layout analysis. (writing mode detection, Type1 font file analysis, etc.)
- Robust error handling.
- Crypt stream filter support. (More sample documents are needed!)
- CCITTFax stream filter support.

## Related Projects

---

- [pyPdf](#)
- [xpdf](#)
- [pdfbox](#)
- [mupdf](#)

## Terms and Conditions

---

(This is so-called MIT/X License)

Copyright (c) 2004-2010 Yusuke Shinyama <yusuke at cs dot nyu dot edu>

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

---

*Yusuke Shinyama (yusuke at cs dot nyu dot edu)*