

lowdown — simple markdown translator

lowdown — simple markdown translator

lowdown is a Markdown translator producing HTML5, *roff* documents in the **ms** and **man** formats, LaTeX, gemini, and terminal output. The [open source](#) C source code has no dependencies.

The tools are documented in [lowdown\(1\)](#) and [lowdown-diff\(1\)](#), the language in [lowdown\(5\)](#), and the library interface in [lowdown\(3\)](#).

To get and use *lowdown*, check if it's available from your system's package manager. If not, [download](#), [verify](#), and unpack the source. Then build:

```
% ./configure
% make
% make regress
# make install
```

lowdown is a [BSD.lv](#) project. Its portability to OpenBSD, NetBSD, FreeBSD, Mac OS X, Linux (glibc and musl), Solaris, and IllumOS is enabled by [oconfigure](#) and checked by BSD.lv's [build system](#).

One major difference between *lowdown* and other Markdown formatters is that it internally converts to an AST instead of directly formatting output. This enables some semantic analysis of the content such as with the [difference engine](#), which shows the difference between two markdown trees in markdown.

Output

lowdown produces HTML5 output in XML mode with **-Thtml**. It may produce either a fragment or standalone HTML5 document with **-s**.

It also produces simple LaTeX documents with **-Tlatex**. It uses the most basic packages possible.

The experimental **-Tgemini** outputs into the [Gemini](#) format.

PDFs may also be produced from *roff* documents via the **-Tms** and **-Tman**¹ outputs. These may be processed with *troff* system such as [groff](#) or (for **-Tman** only) [mandoc](#).

By way of example: this page, [index.md](#), renders as [index.latex.pdf](#) with LaTeX (via **-Tms**), [index.mandoc.pdf](#) with mandoc (via **-Tman**), or [index.nroff.pdf](#) with groff (via **-Tms**).

lowdown can output to ANSI-compatible UTF-8 terminals with **-Tterm**. This [glow](#)-inspired mode renders stylised Markdown-looking output for easy reading. (The traditional text output facilities of [groff](#) and [mandoc](#) may also be used for this.)

[see see see](#)

-Tman -Tterm -Tms

Only **-Thtml** and **-Tlatex** allow images and equations, though **-Tms** has limited image support with encapsulated postscript.

Input

Beyond traditional Markdown syntax support, *lowdown* supports the following Markdown features and extensions:

- autolinking
- fenced code
- tables

- superscripts
- footnotes
- disabled inline HTML
- “smart typography”
- metadata
- commonmark (**in progress**)
- definition lists
- extended image attributes

Examples

Want to quickly review your Markdown in a terminal window?

```
lowdown -Tterm README.md | less -R
```

I usually use *lowdown* when writing [sblg](#) articles when I’m too lazy to write in proper HTML5. ([sblg](#) is a simple tool for knitting together blog articles into a blog feed.) This basically means wrapping the output of *lowdown* in the elements indicating a blog article. I do this in my Makefiles:

```
.md.xml :
( echo "<?xml version=\"1.0\" encoding=\"UTF-8\" ?>" ; \
  echo "<article data-sblg-article=\"1\">" ; \
  echo "<header>" ; \
  echo "<h1>" ; \
  lowdown -X title $< ; \
  echo "</h1>" ; \
  echo "<aside>" ; \
  lowdown -X htmlaside $< ; \
  echo "</aside>" ; \
  echo "</header>" ; \
  lowdown $< ; \
  echo "</article>" ; ) >$@
```

If you just want a straight-up HTML5 file, use standalone mode:

```
lowdown -s -o README.html README.md
```

This can use the document’s meta-data to populate the title, CSS file, and so on.

The troff output modes work well to make PS or PDF files, although they will omit equations and only use local PS/EPS images in **-Tms** mode. The extra groff arguments in the following invocation are for UTF-8 processing (**-k**), tables (**-t**), and clickable links and a table of contents (**-mispdf**).

If outputting PDF, use the `pdfroff` script instead of **-Tpdf** output. This allows image generation to work properly. If not, a blank square will be output in places of your images.

```
lowdown -sTms README.md | groff -itk -mispdf > README.ps
lowdown -sTms README.md | pdfroff -itk -mispdf > README.pdf
```

The same can be effected with systems using [mandoc](#):

```
lowdown -sTman README.md | mandoc -Tps > README.ps
lowdown -sTman README.md | mandoc -Tpdf > README.pdf
```

More support for PDF (and other print formats) is available with the **-Tlatex** output.

```
lowdown -sTlatex README.md | pdflatex
```

For terminal output, troff or mandoc may be used in their respective **-Tutf8** or **-Tascii** modes. Alternatively, *lowdown* can render directly to ANSI terminals with UTF-8 support:

```
lowdown -Tterm README.md | less -R
```

Read [lowdown\(1\)](#) for details on running the system.

Library

lowdown is also available as a library, `lowdown(3)`. This is what's used internally by `lowdown(1)` and `lowdown-diff(1)`.

Testing

The canonical Markdown tests are available as part of a regression framework within the system. Just use `make regress` to run these tests.

I've extensively run `AFL` against the compiled sources with no failures—definitely a credit to the `hoedown` authors (and those from whom they forked their own sources). I'll also regularly run the system through `valgrind`, also without issue.

Code layout

The code is neatly layed out and heavily documented internally.

First, start in `library.c`. (The `main.c` file is just a caller to the library interface.) Both the renderer (which renders the parsed document contents in the output format) and the document (which generates the parse AST) are initialised.

The parse is started in `document.c`. It is preceded by meta-data parsing, if applicable, which occurs before document parsing but after the BOM. The document is parsed into an AST (abstract syntax tree) that describes the document as a tree of nodes, each node corresponding an input token. Once the entire tree has been generated, the AST is passed into the front-end renderers, which construct output depth-first.

There are a variety of renderers supported: `html.c` for HTML5 output, `nroff.c` for **-ms** and **-man** output, `latex.c` for LaTeX, `gemini.c` for Gemini, `term.c` for terminal output, and a debugging renderer `tree.c`.

Example

For example, consider the following:

```
## Hello world
```

First, the outer block (the subsection) would begin parsing. The parser would then step into the subcomponent: the header contents. It would then render the subcomponents in order: first the regular text “Hello”, then a bold section. The bold section would be its own subcomponent with its own regular text child, “world”.

When run through the **-Ttree** output, it would generate:

```
LOWDOWN_ROOT
  LOWDOWN_DOC_HEADER
  LOWDOWN_HEADER
    LOWDOWN_NORMAL_TEXT
      data: 6 Bytes: Hello
    LOWDOWN_DOUBLE_EMPHASIS
      LOWDOWN_NORMAL_TEXT
        data: 5 Bytes: world
  LOWDOWN_DOC_FOOTER
```

This tree would then be passed into a front-end, such as the HTML5 front-end with **-Thtml**. The nodes would be appended into a buffer, which would then be passed back into the subsection parser. It would paste the buffer into `<h2>` blocks (in HTML5) or a `.SH` block (troff outputs).

Finally, the subsection block would be fitted into whatever context it was invoked within.

Compatibility

lowdown is fully compatible with the original Markdown syntax as checked by the Markdown test suite, last version 1.0.3. This suite is available as part of the `make regress` functionality.

How Can You Help?

Want to hack on *lowdown*? Of course you do.

- There are bits and bobs remaining to be fixed or implemented. You can always just search for `TODO`, `XXX`, or `FIXME` in the source code. This is your best bet.
- Footnotes in **-Tms** with groff extensions should use `pdfmark` to link to and from the definition.
- If you want a larger project, a **-Tpdf** seems most interesting (and quite difficult given that UTF-8 need be present). Another project that has been implemented elsewhere is a parser for mathematics such that `eqn` or similar may be output.

¹ You may be tempted to write `manpages` in Markdown, but please don't: use `mdoc(7)`, instead — it's built for that purpose! The `man` output is for technical documentation only (section 7).