

Text rendering modes

The PDF operator behind this is Tr, and it says what showing a text does to the glyph outlines: fill them, stroke them, both, or neither. tclpdf writes it as -render, with a word for each mode.

Rechnung

-render fill
the default - the outlines are filled with -color

Rechnung

-render stroke
only the outline, in -stroke and -strokeWidth

Rechnung

-render fillStroke
filled first, then stroked - one call, two colours

-render invisible
drawn nowhere, and still in the file: select the line

The fourth line above is set with -render invisible. It is in the content stream, it advances the text matrix like any other (9.3.6), and pdftotext gives it back - which is exactly what a scanned page needs under its own picture.

-strokeWidth is a line width, not a font weight

The graphics state parameters of a stroking mode are read in user space rather than in text space (9.3.6). The same 0.3 mm outline is therefore the same 0.3 mm at every size - it does not grow with the face, which is what makes it a rule and not a weight.

Abc Abc Abc

The clipping modes are refused, by name

tclpdf: -render clip is a clipping mode (9.3.6, modes 4 to 7), and tclpdf does not write those: the glyph outlines become a clipping path at ET that stays in force until the next Q, this package sets one BT/ET per line inside its own q/Q, and a second line would be clipped to the intersection with the first, which is empty - available are fill, stroke, fillStroke and invisible

Check it yourself

```
qpdf --check 01.12-render-modes.pdf
qpdf --qdf --object-streams=disable 01.12-render-modes.pdf - | grep Tr
pdftotext 01.12-render-modes.pdf - | head
```