



HAL
open science

On the design of text editors

Nicolas P. Rougier

► **To cite this version:**

| Nicolas P. Rougier. On the design of text editors. 2020. hal-03018466

HAL Id: hal-03018466

<https://inria.hal.science/hal-03018466>

Preprint submitted on 22 Nov 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On the design of text editors

Nicolas P. Rougier

Inria Bordeaux Sud-Ouest

Bordeaux, France

Nicolas.Rougier@gmail.com




Figure 1: GNU Emacs screenshots with hacked settings.

ABSTRACT

Code editors are written by and for developers. They come with a large set of default and implicit choices in terms of layout, typography, colorization and interaction that hardly change from one editor to the other. It is not clear if these implicit choices derive from the ignorance of alternatives or if they derive from developers' habits, reproducing what they are used to. The goal of this article is to characterize these implicit choices and to illustrate what are alternatives, without prescribing one or the other.

CCS CONCEPTS

- Human-centered computing → HCI theory, concepts and models; Interaction design theory, concepts and paradigms; Text input.

KEYWORDS

text, code, edition, design, color, typography, minimal

1 INTRODUCTION

Text (or code) editors are computer programs used for editing plain text. The first rudimentary text editor (QED) was released in 1965 [6, 16], soon followed by EDIT [4]. Since then, a plethora of text editors have been written, such as GNU Emacs [18] (1976) and vi/vim [19] (1976), for the oldest (and still actively developed), or Atom (2014) and Sublime Text (2008), for the more recent ones. Each editor offers some specific features (e.g. modal editing, syntax

colorization, code folding, plugins, etc.) that is supposed to differentiate it from the others at time of release. There exists however one common trait for virtually all of text editors: they are written by and for developers. Consequently, design is generally not the primary concern and the final product does not usually enforce best recommendations in terms of appearance, interface design or even user interaction. The most striking example is certainly the syntax colorization that seems to go against every good design principles in a majority of text editors and the motto guiding design could be summarized by "Let's add more colors" (using regex).

More generally, modern text editors comes with a large set of default and implicit choices that hardly change from one editor to the other. To take only one example, most editors (that I know) consider that there exists only two font weights (*regular* and *bold*): you can choose the *regular* font but rarely can you choose the boldness of the *bold* one. Consequently, if you choose a light or thin weight for the regular font, the difference with the bold font is dramatically accentuated. It is not clear to me if these implicit choices derive from the ignorance of alternatives or if they derive from developers' habits, reproducing what they are used to. The goal of this article is thus to characterized these implicit choices and to illustrate what are the alternatives. However, I do not recommend any specific alternative since this would require a user study that has yet to be done [17]. The goal is rather is to make the developer aware of the alternatives and to let her experiment.

Note: In this landscape, GNU Emacs and vi (or VIM) are very specific because they're highly hackable such that the advanced user




Figure 2: Default layout found in several modern editors

can configure them to their exact liking. The screenshots from the teaser figure have been designed using GNU Emacs.

2 LAYOUT

Beyond the actual text editing area where one can edit a file, most code editors comes with a number of additional features such as tabs, status bar, console, file browser, minimap, etc. This list may vary from one editor to the other, but it appears that there is a set of minimal features that a modern text editor is expected to offer. When these features are all active at once, this naturally and drastically reduces the size of the main editing area as illustrated on figure 2. More importantly, this clutters the space with secondary and peripheral information that a user doesn't look at or use very often when writing code. For example, the file (or project) browser is useful for giving an overview of a project or to select a specific file. But a developer spends most of her time in the text editing area and these peripheral information can be largely considered as a distractor or lost space. This might be the reason why a lot of editors allow to deactivate these components and some of them even offer a distraction free or zen mode (natively or via plug-in) where most of the peripheral information is actually hidden.




Figure 3: Influence of shape, margins and line spacing

If we now look closer at the single window mode (without any extra features), it is striking to observe that there are generally hardly

any margins as illustrated on figure 3A. While it is not conceivable to have a book or a PDF documents without margins, it is somehow considered perfectly normal to not have margins for a code editor (and for a terminal as well on that matter). It is even more surprising considering the recommendations that were made as early as 1986 by Van Nes [22]. Nonetheless, only a handful of modern editors offers the possibility to define margins. One possible explanation might be historical. Forty years ago, the standard vt100 terminal offered only a resolution of 80×24 characters and we could easily imagine that with such limited display, margins were out of the question. Today (2020), most screens have a much larger capability and it is thus difficult to understand why, if not habits, margins are still not an option. This would ease the reading (see 3B) and the same is true for line spacing that is generally set to its minimal value without the possibility to modify it while it would ease the reading even more.

Furthermore, it seems that the default shape of the window in single window mode has inherited the $\frac{3}{1}$ ratio from the vt100 console and the default window tends to be larger than tall. This is again a peculiar choice considering that code is mostly made of short lines (best practices recommend lines length between 72 and 90 characters, 80 being an heritage of IBM punch cards, see [12] for a review) that are mostly written on the left on the window. It would thus make sense to have taller windows as illustrated on figure 3C. Of course, this shape can be modified from the window manager or the settings, but there could be an option, for example, to enforce a fixed ratio of $\frac{1}{\sqrt{2}}$ (ISO 216, see figure 1).

3 TYPOGRAPHY

Typography is the poor relation of code edition [14]. It is as if typography recommendations had been frozen sometime during the eighties and nothing has ever changed since then. Typography in most code editors can be summarized as *Use only monospaced fonts with two weights (regular & bold) and two slants (normal & italic)*. However, digital typography has changed a lot since the eighties [2, 3] and, for example, most typefaces come in several weight variations, ranging from ultra thin to ultra black such that it is possible to define several couples. For example, in the case of Roboto Mono, we can use thin, light, regular, medium or bold:

```
def hello(name): print(f"Hello {name}!")
def hello(name): print(f"Hello {name}!")
def hello(name): print(f"Hello {name}!")
def hello(name): print(f"Hello {name}!")
```

Furthermore, with the advent variable fonts, we can consider several alternatives. Variable fonts have been introduced in version 1.8 of the OpenType font format specification. This new type of font includes (in a single file) multiple variations over one or several axis such as weight, width, optical size, slant etc and makes it possible to interpolate between these variations. For the end-user, this means she can precisely define the different typeface she wants to use. For a text editor, this means it is possible to have context dependent and dynamic font variations. For example, it is possible to have subtle weight variations of a text depending whether text is light text over a dark background or the opposite (think about

```

// Fast Inverse Square Root
float Q_rsqrt( float number )
{
    long i;
    float x2, y;
    const float threehalfs = 1.5F;

    x2 = number * 0.5F;
    y = number;
    i = * ( long * ) &y;           // Evil floating point bit level hacking
    i = 0x5f3759df - ( i >> 1 ); // !What the _?
    y = * ( float * ) &i;
    y = y * ( threehalfs - ( x2 * y * y ) ); // 1st iteration
    // y = y * ( threehalfs - ( x2 * y * y ) ); // 2nd iteration, this can be removed
    return y;
}

```

A

| | | |
|---|---|---|
| <p>Fast Inverse Square Root <small>master branch, last modified yesterday</small></p> <p>Evil floating point bit level hacking <small>What the _?</small></p> <p>1st iteration 2nd iteration, this can be removed</p> | <pre> float Q_rsqrt (float number) { long i; float x2, y; const float threehalfs = 1.5F; x2 = number * 0.5F; y = number; i = * (long *) &y; i = 0x5f3759df - (i >> 1); y = * (float *) &i; y = y * (threehalfs - (x2 * y * y)); // y = y * (threehalfs - (x2 * y * y)); return y; } </pre> | B |
|---|---|---|

Figure 4: A. Regular code display using syntax highlighting. B. Alternative code display taking advantage of typography.

selected text). Since any font axis can be made variable this offers a tremendous amount of possibility and probably a tremendous amount of ways to abuse it.

Another typographical features that was hardly used until very recently are ligatures, that is, the union of two or more glyphs into a single glyph. The Hasklig font by Ian Tuomi (based on Source Code Pro) is the first font to have taken advantage of ligatures and adapted them to source code. For example, the usual notation for *greater than or equal* in most programming languages is `>=` while the mathematical notation is \geq . Ligatures can be used to actually display the mathematical notation without changing the source code. Today font families such as Fira Code, Monoid, Iosevka, Inconsolata or JetBrains Mono all offers a various amount of such code-oriented ligatures [10]. Even if this feature does not entirely depend on the editor (editor must enforce ligature and the font must possess ligatures), it is an aspect to be considered when selecting the default font that is shipped with the editor. Note however that some typographers do not recommend the usage of such code ligature. Butterick [5] goes a bit further and explains that ligatures in programming fonts are a terrible idea because i) they contradict unicode and ii) they are guaranteed to be wrong sometimes.

Beyond eye-candy features, typography can also be exploited to introduce some radical changes in code layout as shown on figure 4. In this example, code and comments are separated in two distinct columns using two different fonts (monospaced font (Fira Code) for code and condensed font (Roboto Condensed) for comments). Comments are placed on the left and the short title is made prominent using larger and heavier font, enriched with contextual information (branch and last commit). It is interesting to note that this alternative occupies the exact same physical space as a regular display. Of

course, this mockup is only one possibilities among several alternatives and most probably, there are better ways to highlight what is deemed to be the important information in a source code. But I think this simple example illustrates the fact that typography can dramatically enhanced readability.

4 COLORIZATION

Syntax highlighting is pervasive and entrenched in virtually all code editors. It aims at visually distinguishing elements of the source code using different colors and/or typefaces. The origin can be traced back to the patent filed in 1982 by Klock and Chodak [9] about a syntax error correction method and apparatus (the first color terminal dates back to 1979 [8]). This has been further refined in 2009 with the concept of semantic highlighting [13] that uses knowledge of the underlying language to provide a finer control on how a given element should be highlighted. However, the advantages of syntax highlighting are far from being obvious. The most recent study [7] *could not find evidence in data that syntax highlighting as used in Eclipse has a beneficial effect on program comprehension for programming novices*. To make things worse, it is also not rare to have a Christmas tree effect resulting from an abuse of syntax highlighting [24]. Consider for example figure 4A that displays a code snippet using the default syntax highlighting of a recent editor. Even though the code is rather small, syntax highlighting results in six different colors being used all over the source code, making it difficult to assign a specific semantic to any given color. The question, is thus, does it help the developer? According to [1], *using more visual variety when rendering methods substantially reduces comprehension time of code features*. The problem with syntax highlighting is that it does not seem to be based on any specific principles and derives from the possibility of identifying code parts based on regular expressions, and the colorization of such expression. But there are no scientific recommendation on what to highlight or how to highlight. Only the *solarized* color palette crafted by Ethan Schoonover seems to enforce some design principles with *reduced brightness contrast while retaining contrasting hues for syntax highlighting readability* even though the author doesn't prescribe how to apply it such color scheme.

There exist however alternate use of colorization where the semantic of color is well defined. This is the case for multi-authored document where each author is identified with a unique color. This can be used during live editing such as notepads or post-edition using tools such as git blame (for example). Another possibility is to use colorization in order to show the modification history of a document, using light tint for old modifications and heavier tint for recent modifications. Wayne [23] goes a step further and denounces the use of syntax highlighting since it is a waste of an important information channel and suggest several alternative uses of color, among which, rainbow parenthesis, context highlighting, import highlighting, argument highlighting, type highlighting, etc. Instead of syntax or semantic colorization based on content, a simple alternative would be to adapt colorization to the reader, taking attentional constraints into account [11, 21]. For example, here is the color scheme that has been used to design the mockup on figure 1 and which is based on the perception rather than the content:

Critical face is for information that requires immediate action. It should be of high contrast when compared to other faces. This can be done (for example) by setting an intense background color, typically a shade of red. It must be used scarcely.

Popout face is used for information that needs attention. To achieve such effect, the hue of the face has to be sufficiently different from other faces such that it attracts attention through the popout effect.

Strong face is used for information of a structural nature. It has to be the same color as the default color and only the weight differs by one level (e.g., light/regular or regular/bold). It is generally used for titles, keywords, directory, etc.

Salient face is used for information that are important. To suggest the information is of the same nature but important, the face uses a different hue with approximately the same intensity as the default face. This is typically used for links.

Faded face is for secondary information that is less important. It is made by using the same hue as the default but with a lesser intensity than the default. It can be used for comments, secondary information and also replace italic (which is generally abused anyway).

Subtle face is used to suggest a physical area on the screen. It is important to not disturb too strongly the reading of information and this can be made by setting a very light background color that is barely perceptible.

In fact, such cognitive colorization does not require any change in syntax in highlighting engines. It only requires a restricted set of colors and a careful selection of what information needs to be salient, faded or strong.

5 USER INTERACTION

One feature shared by all text editors is the linear representation of a text file and the use of scrolling to navigate forward or backward. This has become a seemingly ubiquitous part of our user experience and it seems difficult to imagine any alternative that nonetheless exist. It is, for example, quite common to have a dedicated navigation panel allowing to jump to a specific part of the code. This panel can be built automatically (e.g. function, method, class, definition) or manually using bookmark inside a text file. However, these navigation panels usually follows the linear structure of the code (even though some may propose a sorted list) and do not process semantically the content (when built automatically). These navigation panels are actually comparable to a table of content and this suggests a code source could be considered as a book as it has been proposed by Oman and Cook [14]. The author proposed to format source code as if it was printed on a book, taking advantage of typography, logical blocks separation and using the page paradigm (instead of scroll). As explained by the author, *the components of a book (preface, table of contents, indices and pagination, chapters, sections, paragraphs, sentences, punctuation, type style, and character case) are all designed to facilitate rapid information access and transfer*. Such printed book paradigm has been hardly used in any text editor but it is certainly a direction to explore further.

GNU Emacs – Elegance

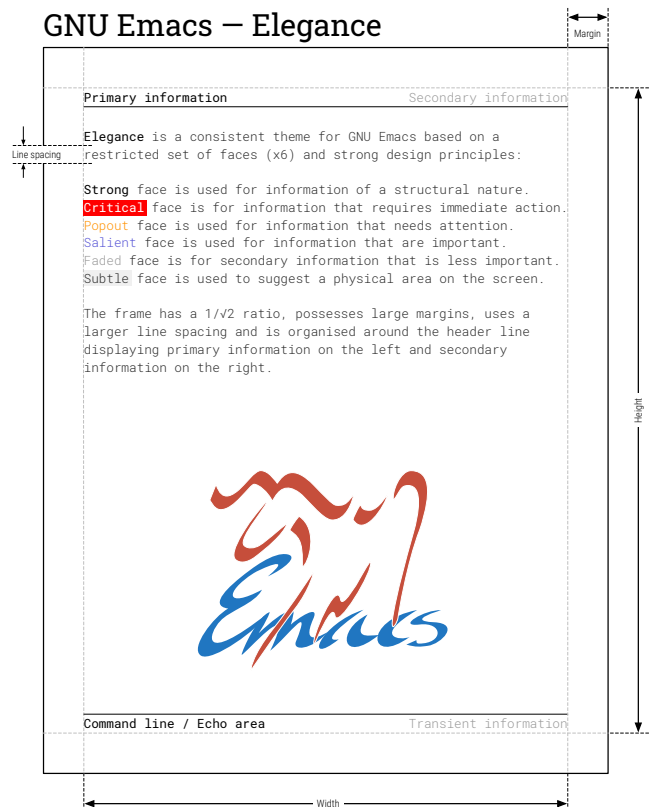


Figure 5: GNU Emacs with heavily modified settings enforcing some of the alternatives presented in this article.

An important and critical aspect of interaction in a code editor is (of course) the actual input of text and commands since coding activity encompasses actual code writing but also code navigation. In most modern editors, such inputs are carried out via the combination of keyboard and mouse while, in older editors such as Emacs and vi, it is possible to issue text and commands from keyboard only. For Emacs, this is made possible using the command-line that is an integral part of Emacs and allow to type literal commands. This has probably inspired the command palette that is now found in modern editors. This command line is complemented by several keybindings that may be global or specific to a given mode (e.g. Python mode, Lisp mode, etc.) and any command can also be bound to any key sequence (even very complex and long ones). The philosophy of vi is quite different because it has adopted a modal approach and operates in two modes: the insert mode where keystrokes modify the document and the command mode where keystrokes are interpreted as commands. Such modal feature has been somehow killed by Tesler [20] who transformed the modal Bravo editor into the modeless Gypsy editor during the seventies and whose efficiency has been demonstrated [15]. This may very well had a profound impact on the community because the only surviving modal editor is actually vi and only the new kakoune editor (kakoune.org) adopted a similar modal approach (using an object + verb approach while vi uses verb + object). However, four

decades later, this efficiency might need to be re-evaluated, especially in light of challenges such as vimgolf (www.vimgolf.com) that demonstrates how to perform complex and real-world code transformation with only a few keystrokes. For example, let's consider the following text: "The quick brown fox jumps over the lazy dog." that we want to convert to "The quick lazy dog jumps over the brown fox." The actual sequence of vi keystrokes to transform the first text into the second is 2wd2w3wPd3w6bep (2w: Move forward two words – d2w: Delete two words – 3w: Move forward three words – P: Paste the previously deleted text before cursor position – d3w: Delete three words – 6b: Move back six words – e: Move to the end of the current word – p: Paste the previously deleted text after the cursor position). This keystroke sequence is of course a bit cryptic for those not familiar with vi but it is nonetheless quite efficient. No doubt that the casual user would not like to have to learn such meta-language before being able to use the editor.

But let me remind you that users of code editors are not casual users and such modal interface might need to be further exploited in modern editors. Similarly, there are plenty of "old" concepts that might be worth to be re-considered such as kill ring (extended clipboard), recursive undo (possibility to undo an undo command), rectangle selection, etc.

6 CONCLUSION

I've highlighted several implicit choices that are present in a number of both old and modern text editors and introduced several alternatives that, I think, are worth to be explored and exploited by future developers and designers. Several of these alternatives have been actually introduced 30 years ago by Oman and Cook [14] and I've mostly updated them in light of available technologies in 2020 and added new ones that were hardly imaginable in 1990. In the meantime, there are viable alternatives that are already implemented in historic editors (vi and GNU Emacs) that may be also worth to be reconsidered in the design of future code editors. Overall, the sum of all these alternatives offer great possibilities to enhance the user experience as illustrated in figure 5. This is not a mockup but my actual GNU Emacs configuration that I use for code, text, mail and agenda. Note that I'm not advocating for people to use such setup, the goal is merely to highlight an example of a radical design for a text editor that is perfectly usable according to my daily and personal experience.

REFERENCES

- [1] Dimitar Asenov, Otmar Hilliges, and Peter Müller. 2016. The Effect of Richer Visualizations on Code Comprehension. In *Proceedings of the 2016 CHI Conference*

- on *Human Factors in Computing Systems*. ACM. <https://doi.org/10.1145/2858036.2858372>
- [2] Charles Bigelow. 2020. The Font Wars, Part 1. *IEEE Annals of the History of Computing* 42, 1 (1 2020), 7–24. <https://doi.org/10.1109/mahc.2020.2971202>
- [3] Charles Bigelow. 2020. The Font Wars, Part 2. *IEEE Annals of the History of Computing* 42, 1 (1 2020), 25–40. <https://doi.org/10.1109/mahc.2020.2971745>
- [4] S. R. Bourne. 1971. A design for a text editor. *Software: Practice and Experience* 1, 1 (1 1971), 73–81. <https://doi.org/10.1002/spe.4380010107>
- [5] Matthew Butterick. 2019. *Practical Typography, second edition*. online. <https://practicaltypography.com/>
- [6] L. Peter Deutsch and Butler W. Lampson. 1967. An online editor. *Commun. ACM* 10, 12 (12 1967), 793–799. <https://doi.org/10.1145/363848.363863>
- [7] Christoph Hannebauer, Marc Hesenius, and Volker Gruhn. 2018. Does syntax highlighting help programming novices? *Empirical Software Engineering* 23, 5 (2 2018), 2795–2828. <https://doi.org/10.1007/s10664-017-9579-0>
- [8] IBM 1979. *IBM 3270 Information and Display System: Color and Programmed Symbol*. IBM. GA33-3056-0 File No. 5360/5370/53/4300/8100-09.
- [9] Anita H. Klock and Jan B. Chodak. 1982. Syntax error correction method and apparatus. <https://patents.google.com/patent/US4617643A/en> US Patent US4617643A.
- [10] Matej Latin. 2020. 5 monospaced fonts with cool coding ligatures. Blog post. <https://betterwebtype.com/articles/2020/02/13/5-monospaced-fonts-with-cool-coding-ligatures/>
- [11] Lori McCay-Peet, Mounia Lalmas, and Vidhya Navalpakkam. 2012. On saliency, affect and focused attention. In *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems - CHI '12*. ACM Press. <https://doi.org/10.1145/2207676.2207751>
- [12] Anuj Nanavati and Randolph Bias. 2005. Optimal Line Length in Reading—A Literature Review. *Visible Language* (01 2005).
- [13] David Nolden. 2009. From Syntax Highlighting to Semantic Highlighting. Blog post. <https://zwabel.wordpress.com/2009/01/08/c-ide-evolution-from-syntax-highlighting-to-semantic-highlighting/>
- [14] Paul W. Oman and Curtis R. Cook. 1990. Typographic style is more than cosmetic. *Commun. ACM* 33, 5 (5 1990), 506–520. <https://doi.org/10.1145/78607.78611>
- [15] Merle F. Poller and Susan K. Garter. 1983. A comparative study of moded and modeless text editing by experienced editor users. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems - CHI '83*. ACM Press. <https://doi.org/10.1145/800045.801603>
- [16] Dennis Ritchie. 2004. An incomplete history of the QED Text Editor. Blog post. <https://www.bell-labs.com/usr/dmr/www/qed.html>
- [17] Teresa L. Roberts and Thomas P. Moran. 1983. The evaluation of text editors: methodology and empirical results. *Commun. ACM* 26, 4 (4 1983), 265–283. <https://doi.org/10.1145/2163.2164>
- [18] Richard M. Stallman. 1981. EMACS the extensible, customizable self-documenting display editor. *ACM SIGPLAN Notices* 16, 6 (6 1981), 147–156. <https://doi.org/10.1145/872730.806466>
- [19] Sinclair Target. 2018. Where Vim Came From. Blog post. <https://twobithistory.org/2018/08/05/where-vim-came-from.html>
- [20] Larry Tesler. 2012. A personal history of modeless text editing and cut/copy-paste. *interactions* 19, 4 (7 2012), 70. <https://doi.org/10.1145/2212877.2212896>
- [21] Anne M. Treisman and Garry Gelade. 1980. A feature-integration theory of attention. *Cognitive Psychology* 12, 1 (1 1980), 97–136. [https://doi.org/10.1016/0010-0285\(80\)90005-5](https://doi.org/10.1016/0010-0285(80)90005-5)
- [22] F. L. Van Nes. 1986. Space, colour and typography on visual display terminals. *Behaviour & Information Technology* 5, 2 (April 1986), 99–118. <https://doi.org/10.1080/01449298608914504>
- [23] Hillel Wayne. 2020. Syntax highlighting is a waste of an information channel. Blog post. <https://buttondown.email/hillelwayne/archive/46a17205-e38c-44f9-b048-793c180c6e84>
- [24] Linus Åkesson. 2007. A case against syntax highlighting. Blog post. <http://www.linusakesson.net/programming/syntaxhighlighting/>