

Terminal Lucidity: Envisioning the Future of the Terminal

Michael MacInnis
School of Computer Science
Carleton University, Ottawa, Canada

Olga Baysal
School of Computer Science
Carleton University, Ottawa, Canada

Michele Lanza
REVEAL @ Software Institute
USI, Lugano, Switzerland

Abstract—The Unix terminal, or just simply, the terminal, can be found being applied in almost every facet of computing. It is available across all major platforms and often integrated into other applications. Due to its ubiquity, even marginal improvements to the terminal have the potential to make massive improvements to productivity on a global scale. We believe that evolutionary improvements to the terminal, in its current incarnation as windowed terminal emulator, are possible and that developing a thorough understanding of issues that current terminal users face is fundamental to knowing how the terminal should evolve. In order to develop that understanding we have mined Unix and Linux Stack Exchange using a fully-reproducible method which was able to extract and categorize 91.0% of 1,489 terminal-related questions (from the full set of nearly 240,000 questions) without manual intervention.

We present an analysis, to our knowledge the first of its kind, of windowed terminal-related questions posted over a 15-year period and viewed, in aggregate, approximately 40 million times. As expected, given its longevity, we find the terminal’s many features being applied across a wide variety of use cases. We find evidence that the terminal, as windowed terminal emulator, has neither fully adapted to its now current graphical environment nor completely untangled itself from features more suited to incarnations in previous environments. We also find evidence of areas where we believe the terminal could be extended along with other areas where it could be simplified. Surprisingly, while many current efforts to improve the terminal include improving the terminal’s social and collaborative aspects, we find little evidence of this as a prominent pain point.

Index Terms—terminal, Stack Exchange, Unix, programming, development environment

I. INTRODUCTION

From humble beginnings as a teletypewriter, through video display terminals, to windowed terminal emulators, what we now think of as the Unix terminal, or just simply, the terminal, has become a fixture of the modern computing environment. The terminal is everywhere.

Before Windows Subsystem for Linux (WSL), the Microsoft POSIX subsystem, Cygwin, and Windows Services for UNIX provided a stopgap for those requiring access to Unix command-line tools on Windows [1]. WSL might have provided a similar anemic experience if not for the Windows Terminal which brought a full-featured, integrated, Unix terminal to Windows.¹ With the release of Windows 11, the Windows Terminal is now the default command-line

environment on Windows [2] making the terminal available, by default, across all major platforms.

The terminal is used so frequently when working with containers and the cloud that introductions to Docker or Kubernetes often contain inline discussions of the Unix shell and how to run commands from the terminal [3], [4]. While Python dominates the AI space (with C++ under the hood) and large language model (LLM) experimentation often assumes the presence of Jupyter Notebooks, the terminal is used to perform tasks that fall outside the purview of these tools [5]. Many modern text editors, such as Zed, feature an integrated terminal [6]. Installation instructions for languages like Go and Rust assume the presence of a terminal [7], [8]. The Integrated Development Environment (IDE), once the antithesis of the terminal-based approach, now often provides an integrated terminal. Microsoft’s VS Code IDE includes an integrated terminal [9]. KDevelop includes an integrated terminal [10]. JetBrains’ IDEs & Android Studio include an integrated terminal [11], [12] and the Android SDK includes a suite of command-line tools [13].

There are several command-line iOS development tools, and while Xcode does not currently provide an integrated terminal, iOS applications are developed on macOS which, due to its UNIX heritage, provides a terminal. Even before OS X the terminal was a prominent part of “the premier software development system for the Mac” [14] and when Mac OS X was introduced, the Macintosh became a “UNIX box”, and the terminal was front and center in early promotional material.²

We do not wish to overstate the importance of the terminal. We are not claiming anything close to “real developers use the terminal”. Using the terminal where it does not make sense is as strange as avoiding it where it does. We do, however, find that, given its ubiquity and utility in areas like software and systems development [3], [4], [15], [16], system administration [17], performance tuning [18], data science [19], computer and network security [20], etc. the terminal is often conspicuously absent from research discussions in these areas. We can only speculate as to why. It might be a case of familiarity breeding contempt. It could be that the terminal has become so common as to be invisible. We reject the idea that the terminal has achieved perfection with no potential for improvement.

¹<https://www.youtube.com/watch?v=8gw0rXPMPE>

²<https://www.brainmapping.org/MarkCohen/UNIXad.pdf>

The terminal deserves credit for its longevity and continued relevance over the last 50 years but it has not existed unchanged. The terminal is not even one entity but rather a large set of distinct applications with features that have converged to support popular terminal-based applications. We believe further evolution of the terminal is possible and that the key to understanding how the terminal should evolve lies in understanding the issues that current terminal users face.

The primary question we wish to address with this work is the possibility of developing a requirements-driven understanding of how the terminal should evolve based on a systematic analysis of a large body of windowed terminal-related questions posted over a 15-year period and viewed, in aggregate, approximately 40 million times.

II. RELATED WORK

While there are many papers that mine the Stack Exchange network for insights, such as [21]–[24], we know of no work that is directly related to the topic of this paper — investigating terminal-related issues to understand how the terminal should evolve. Other than [25], we know of no recent terminal-related papers. We feel that this absence combined with the popularity and longevity of the terminal, only serves to underscore the urgent need for more work in this area.

Popular terminals tend to be fairly conservative in the features they offer. There are, however, at any given time, often several efforts, at various stages in development, attempting to pull the terminal in new directions. Recent efforts include Xiki, Fig, sshx, and Warp. Xiki, a “shell console with GUI features”, promises a friendlier and more powerful command line [26], but development seems to be stalled. Fig, the “next-generation command line”, aimed to make “the command line easier and more collaborative” [27]. Unfortunately, Fig is sunsetting effective September 1st, 2024 [28]. Sshx, an application that lets users “share collaborative terminals over the web, with live cursors on an infinite canvas” [29] is being actively developed [30]. Warp, a “terminal built for the 21st century”, is currently in public beta [31]. A goal many of these efforts share appears to be improving the social and collaborative aspects of the terminal.

III. METHODOLOGY

To understand the issues that terminal users face we analyzed terminal-related questions extracted from Unix & Linux Stack Exchange (U&LSE).³ As of March 2023, U&LSE was one of “the three most actively-viewed sites” in the Stack Exchange network [32]. Topics discussed on U&LSE overlap with those discussed on Ask Ubuntu, Server Fault, Super User, and Stack Overflow but whereas those sites focus on a particular Linux distribution, system administrators, power users, and programmers, respectively, questions on U&LSE are of interest to an intersection of those audiences. There is no dataset that fits our purposes exactly. U&LSE works as a rough approximation and is more suited than other similar datasets.

³<https://unix.stackexchange.com/>

A. Selection and Filtering

Browsing the April 2024 U&LSE Data Dump, we observed that all of the potentially interesting questions we manually identified contained the word (or tag) *terminal*. We initially resisted filtering out closed questions and only performed minimal filtering based on score as we did not want to eliminate questions that might suggest more radical departures from the state of the art. What we found was that more radical questions often had fairly high scores. Low-scoring and/or closed questions, that were not completely unrelated, were often unclear and sometimes poorly-worded, near-duplicates of questions with higher scores. In line with other investigations [33], we found that eliminating closed questions significantly reduced noise and that a question score threshold could be used as a “slider” to reduce not only the noise but the size of the filtered data set.

Restricting our analysis to questions containing the word (or tag) *terminal*, or, for good measure, the name of any one of 20 popular terminal emulators (*iterm*, *konsole*, *rxvt*, *xterm*, etc., see the replication package for the full list) — reduced the set of 238,249 questions to 20,805 questions. Due to pragmatic considerations, further reductions of the data set were necessary. We initially considered random sampling but ultimately decided against it so as not to hinder the reproducibility of our results. Working backwards, we found that questions with a score of less than ten provided “no new or relevant data” [34]. Excluding closed questions and questions with a single-digit score or less resulted in a set of 1,489 potentially interesting questions.

B. Discovering Local Regularity

Browsing this smaller set of questions, common tags, words, and phrases became more apparent. The language used by members of a community develops “a high level of local regularity” [35]. The U&LSE community has developed its own local dialect. The U&LSE community is also a subset and reflection of the wider (and much older) Unix and Linux community which has an already established jargon. In addition to common words and phrases, the U&LSE community has been developing a “folksonomy” [36] by applying tags to posts over the last 15 years.

C. Computer-Aided Qualitative Coding

Inspired by recent regex-based [37] and hybrid [38] qualitative coding efforts in quantitative ethnography and similar experiences in software engineering research [39], [40] we created simple patterns to extract and sort questions into categories. Where possible we preferred to base our patterns on existing tags so as to reuse Stack Exchange’s folksonomy. Next we based our patterns on words and short phrases that appeared in multiple related posts. Where the phrasing of a post was too idiosyncratic to be generalized we based our pattern on the exact phrase that we believed identified a post as being related to a category. Our automated process then sorted the data into their respective categories and the posts in those categories were manually reviewed.

Patterns were adjusted (which sometimes involved adding or removing categories) and the scripts run again. By using `git` (and because Stack Exchange posts are single-line records), we were easily able to see the effect of these changes across hundreds of iterations as we verified the categories we discovered “over and over again against actual data” [34]. Guiding each iteration was the goal of using the fewest number and smallest patterns possible while minimizing the number of manual corrections required - minimizing the “disagreements” between the human and the automated process.

Computer-aided qualitative coding allows one “rater” to compare and repeatedly refine their results against an automated process. In our case, the result of this refinement is a set of simple regular expressions able to sort 1,355 of 1,489 questions into our 26 discovered categories without manual intervention for an inter-rater reliability of 91%. By restricting ourselves to regular expressions we are able to present clear and unambiguous rules. While the patterns we created provide transparency they are overly simplistic in the face of unrestricted natural language. The categorization of 83 questions (5.6%) had to be manually corrected. A further 51 questions (3.4%) were assigned a correct category along with one or more incorrect categories and had to be manually adjusted. Manually adjusted questions are marked with an asterisk (*) when referenced in the results section. Manually corrected categorizations are detailed in the Appendix.

Other than the manually categorized questions, which are extracted before further processing there are no ordering restrictions on the patterns for categories other than the unrelated category. The unrelated category is a combination of questions that match a set of patterns as well as the unselected questions that fall through the pipeline. The set of patterns for the unrelated category is reliable enough that we use them as a pre-filter. The categorization process is depicted in Figure 1.




Fig. 1. Question categorization process.

D. Caveats

While we have made every effort to ensure the consistency of our categorization we make no claims regarding its applicability beyond the discussion in this paper or datasets other than U&LSE.

The primary advantage of our approach is that it is possible for it to be “independently replicated by other researchers” [41]. All of the scripts and patterns used in our analysis are available in the replication package along with instructions on their use.⁴

IV. RESULTS

The following is a summary of 1,489 terminal-related questions belonging to 25 of our 26 discovered categories. The only category not shown is the unrelated category. We have grouped and ordered categories in an effort to make the information more easily digestible (see Figure 2). Category names are given in each section along with the number of posts and aggregate number of views. 104 posts were assigned two categories, 11 posts were assigned three categories, and two posts were assigned four categories resulting in an inflated total of 923 posts for the 25 categories shown below and not 791 which is the number of unique posts not categorized as unrelated.




Fig. 2. Question views (in millions) per category.

A. Job Control, Multiplexing, and Remote Machines

The most frequently viewed, terminal-related, posts on U&LSE are related to job control, terminal multiplexers, and using the terminal to interact with remote systems.

⁴<https://doi.org/10.5281/zenodo.12788678>

1) **job-control (posts: 76, views: 5,061,222)**: First added to the 4BSD C shell in concert with 4BSD's new terminal driver, job control transformed the terminal into a multitasking interface.

Some people claim that a well-designed windowing system removes any need for job control. Some complain that the implementation of job control — requiring support from the kernel, the terminal driver, the shell, and some applications — is a hack. Some use job control with a windowing system, claiming a need for both [42].

The claim that “a well-designed windowing system removes any need for job control” is an exaggeration, though job control is less useful in the presence of a windowing system, with some users looking to “disable job control entirely” (137915). A major problem with job control and managing processes from the terminal in general, is that the mechanisms involved (controlling terminal, process groups, signals) are not widely understood, are under-specified, interact in complex and surprising ways, and allow applications to both override behavior and perform the same operations themselves. The shell itself complicates the situation by defaulting to delivering signals like SIGINT to the currently running process which can make tasks like “[t]erminating an infinite loop” (42287) more difficult than one might expect. All of this leaves users wondering how to perform relatively simple tasks like running “a command which will survive terminal close” (4004) or how to “detach a process from a bash script” (269805). While it is possible to “detach” a process from the terminal, the reverse is not easily or reliably possible (4034). There is the question of what to do with output from a program that is no longer connected to a terminal (17648), or worse, a program not connected to a terminal that attempts to read input (25024). The terminal's buffering and the speed at which commands can generate output combine to make job control features seem less effective (176917).

a) **io-routing (posts: 7, views: 1,269,311)** Just on the edge of what is possible, there is a desire to “attach a terminal to a detached process” (31824) without the use of a multiplexer, “view the output of a running process” (58550) from another terminal, and interact with processes in other terminals (365256).

b) **messaging (posts: 5, views: 462,386)** Overlapping to a certain degree with the desire to interact with processes in other terminals, the terminal provides ancient communication facilities used by the `msg`, `wall`, and `write` commands. With the exception of `wall`, these commands are now rarely used and often disabled on the receiving end (135108). This along with narrow scope of these tools typically leads to a search for alternatives (21334, 99460, 377768).

2) **multiplexers (posts: 107, views: 5,229,577)**: Terminal multiplexers like GNU `screen` [43] and `tmux` [44] enable launching multiple virtual terminals in one terminal session as well as detaching and reattaching to a session. They also allow users to “split the terminal into more than one 'view'”

(7453), group terminals into different sessions (24274), share “a terminal with multiple users” (122801), programmatically send input to a terminal (5305), and make a terminal read-only (7588). Not all users are “sold” on the concept of terminal multiplexers with some asking questions like, “Is screen useful?” (4116). The author of the `kitty` terminal has stated that “multiplexers add unnecessary overhead [and] suffer from a complexity cascade” [45]. Terminal multiplexers do add another layer which adds complexity and the features multiplexers provide sometimes overlap with those provided by other layers which can leave users wondering how to “make Screen scroll like a normal terminal” (43229) or how to copy and “paste text selections between `tmux` and the clipboard” (67673). Terminal multiplexers often require some tweaking (29907*, 38402, 38911, 348913, 511267). How to modify the default keybindings (5998, 453466) to be less intrusive (1636) or to avoid conflicts with other applications (12455) are frequent questions. Given their complexity, it is not always obvious which operations are required to achieve a desired effect or worse if the desired effect is possible at all. However, even with their additional complexity, terminal multiplexers are undeniably useful, particularly in a remote context.

3) **remote (posts: 66, views: 7,499,998)**: The terminal has always been a portal. Many applications and protocols have been developed over the years to extend the terminal's ability to communicate with a remote machine with SSH (and SCP and SFTP) now dominating. Within the confines of the terminal emulator window, SSH provides the illusion of interacting with a remote machine as if it were the local machine. This illusion breaks down when operations require access to both local and remote file systems with a common question being, “How to copy a file from a remote server to a local machine?” (188285) and vice-versa. The situation is a continuation of the Telnet/FTP divide (35538). While SSH can be used to make multiple “hops” from one remote machine to another, it is largely up to the user to remember how many “levels” there are in the current session if they wish “to exit everything except the 'root' shell” (3212) — this problem is not unique to SSH but is more apparent in a remote context. SSH can also be used to forward X connections, allowing remote access to graphical applications (7226) although not without issues (505908) and with no completely backward-compatible Wayland equivalent (776767).

SSH itself can be the cause of other issues, like the terminal appearing to hang when a connection is lost (196701) or interactive terminal-based programs not working correctly because no pseudo-terminal has been allocated (269474). SSH provides its own solutions to these issues, with mechanisms to break away from “an unresponsive ssh session” (2919) and force pseudo-terminal allocation (43496). It is somewhat remarkable that many terminal features, such as setting the window title (14113) or changing terminal colors, continue to work as expected over SSH connections. However, this does mean that remote systems may set the window title (40830) or change terminal colors (58982). There is also the question of how agnostic SSH should be (205567).

B. The Terminal as Windowed Terminal Emulator

The windowed terminal emulator brought the terminal into the modern computing environment and the world of the graphical user interface.

1) **integration (posts: 93, views: 4,321,881)**: Placing the terminal in a graphical environment alongside other applications creates an expectation that the terminal should integrate and be able to interact with both the environment and other applications. There is a desire to seamlessly “launch GUI applications from a terminal window” (150706) and similarly to close “various windows using the terminal” (85205). Users wish to configure not only the terminal itself “from within the terminal” (176175) but, in general, to be able to perform tasks that can be done outside the terminal from within the terminal. For example, being able to “set the audio volume using the terminal” (32206), “change the icon of an application’s window” (179174), or “[c]hange desktop wallpaper” (59653) from the terminal. Many terminals now support “clickable links” (112267) and while it is possible to make file names clickable (63417) this is not the default behavior.

Being a graphical application itself, common questions include how to “set the default size and position” (48984) for a new terminal window, configure a “shortcut to launch a terminal window” (47608), “launch an application with [the] default terminal emulator” (32547), and configure the “Open in Terminal” feature of other applications (336368, 547492). There are many looking for ways to launch a set of programs in multiple terminal windows or tabs (22682, 75902, 158434, 492365). With the amount of configuration and customization possible a common question is not only how to save the configuration of a single terminal (35748) but the configuration (including the size and placement) for a set of windows (76497, 274534, 593778).

2) **colors (posts: 63, views: 3,569,541)**: The terminal, as video display terminal, developed techniques for displaying color. Some of these survive in, and have subsequently been extended by, windowed terminal emulators. Even with widespread support for color, it is difficult to reliably detect the level of support, such as whether the “terminal supports 24-bit / true color” (450365). Configuring various commands to produce color can be similarly difficult with each having its own ad hoc method to enable the command “to display colored output” (94299). Further complications are caused by using commands in combination. Commands like `less` need to be configured to “properly display colours” (150328) (see also Section IV-D1). A consideration that is equally important in non-terminal contexts, there is a desire to choose colors that are not “impossible to read” (94498) in general as well as for those with “red green color blindness” (3692), for example. Choosing appropriate colors may involve detecting other colors that are already being used (1755, 172548). There are cases where “removing control chars (including console codes / colours)” (14684) or converting “colored output of [an] arbitrary program” (44956) to other representations is required.

3) **graphics (posts: 9, views: 1,641,946)**: Given that terminal emulators are graphical applications capable of displaying

color a natural next question is, “Are they capable of displaying graphics?” Some video display terminals had the ability to display graphics. In addition to character graphics [46], two techniques that survive in some terminal emulators are the Sixel bitmap graphics format [47] and the ReGIS vector graphics markup language [48]. Some terminal emulators have added further extensions [49]. Intermingling textual and graphical elements adds additional complexity which may be why many terminal emulators choose to avoid supporting graphics. Still, the desire remains to “view images from the terminal” (35333), “watch movies inside the terminal” (123674), and “draw images, shapes etc in a terminal” (146486).

4) **window-size (posts: 21, views: 1,073,338)**: Windowed terminals can have as many rows and columns as will fit on screen or be shrunk down to barely usable dimensions. Some commands need to be configured to use “the full width of the terminal” (9301). In some situations, users may wish to “set the maximum line width” (222528) in order to force a command to ignore the full size of the terminal. When lines are longer than the terminal is wide, should they be wrapped (20493, 70419)? Some terminal-based applications perform their own “long line wrapping” (12279*). Terminal windows may be resized while running with a corresponding expectation that the content will be dynamically reformatted “on terminal dimension changes” (11389).

5) **fonts (posts: 16, views: 444,840)**: From a terminal perspective, encodings like UTF-8 (see Section IV-E2) solve one half of a problem, defining which byte sequences correspond to which characters. The other half of the problem is correctly displaying the glyphs corresponding to those byte sequences. Many of these issues are terminal emulator specific. More universal issues include knowing what font the terminal is using (9696*, 154736), knowing how to change the terminal font (372447*, 407710), finding “the best font for rendering a codepoint” (162305), problems with non-monospace fonts (273059), and enabling or disabling more advanced font rendering features, such as ligatures (278175).

6) **mouse (posts: 11, views: 330,779)**: With graphical environments comes the mouse (or similar pointing device) and the expectation that the mouse should be able to interact with the text in the terminal window. Unlike other applications, many terminal-based applications do not allow “placing the cursor by clicking” (444601) on a position with the mouse. Some terminal-based applications do support the mouse but that support may need to be enabled (252995) or disabled by users who prefer that the terminal handle “mouse-select and copy-paste” (187695). When the terminal is handling the mouse, users may wish to configure the text that the terminal selects in response to mouse activity (45925, 91137, 290544) which may be complicated by the presence of other applications (318281).

7) **cut-copy-paste (posts: 67, views: 1,612,489)**: Terminal-based applications implementing copy&paste often do not use the clipboard. Due to independent implementations, copy&paste differs between terminal-based applications, falling back to the terminal to provide it (163715).

To further complicate matters, the X Window System not only has a clipboard but a primary *and* secondary selection [50] with Wayland also providing support for a primary selection in addition to the clipboard [51]. There is an obvious desire to “use just one unified clipboard” (13585) and have the “terminal selection automatically copy to the system clipboard” (172044) or, at a minimum, simply for consistent behavior from application to application (178070). These complications in addition to the peculiarities of some terminal-based applications create a situation where copying and pasting between the terminal and other applications “does not always work” (136229) as expected. Copying between the terminal and other applications also raises conversion questions (293472).

Specific to the terminal, it is not always easy or obvious how to “copy multi-page text from the terminal” (109935) or how to copy the text faithfully and not “as displayed” with characters that are a result of padding, line wrapping, or other display issues (126577, 218248). There is a desire to access the system clipboard from the terminal (125839) and easily “copy a file or directory location as a path and then paste it in a terminal window” (102551). When copying relative file paths from the terminal the working directory of the terminal-based program that produced the relative file path is needed “to construct a full path” (211418) (see also Section IV-D3). There is a desire to quickly copy the last command from the terminal (91863) and paste commands into the terminal and either have them automatically executed (364047, 633370, 689394*) or prevent their automatic execution (96357, 202732, 230678, 309786*). Bracketed paste was developed to address the problem of differentiating between typed and pasted text [52] and is what now prevents the automatic execution of pasted commands. While this distinction is important in some situations (355610) bracketed paste can cause unexpected behavior in terminal-based applications that do not support it (196098). When pasting text as input to a terminal-based application the pasted text “gets mixed with the final output” (693720) causing the two to be confusingly intermingled. An “ancient Unix limitation” also limits lines to 4,096 characters (131105, 643777). Given the heavy use of the keyboard some wish to “select text [...] via the keyboard only” (210768) and perform other operations such as switching “back and forth between [...] tabs using the keyboard” (67861⁵).

The terminal’s keyboard shortcuts differ from those in most other applications. This is perhaps not surprising when we consider that the terminal predates many of these other applications (16062). A frequent stumbling block, Ctrl-C, the keyboard shortcut used to copy in most applications, is the terminal’s default shortcut for interrupting a command. Ctrl-V, which users might expect to paste a previously cut or copied selection, is used by popular shells like `bash` to mean “insert the next character verbatim”. The serendipitous existence and use of the Command key on the Macintosh allows it to circumvent some of these issues.

However, the issue remains for the vast majority of users. Often there is a desire simply for consistency (52673*, 271150, 339160) and a “[s]ingle set of keyboard shortcuts for copy/paste/cut across all [...] applications” (33501). Previous attempts to address these issues have not been successful (18589⁵).

8) *key-bindings* (posts: 30, views: 677,455): Not limited to shortcuts for copy and paste, at one point even the default behavior of the backspace and delete keys when using the terminal was potentially problematic [53]. Some keys, like the meta key, that were present on keyboards of the day but are no longer typical, have to be remapped to other keys, like the alt or escape key (28993). Different terminals send different byte sequences for the home and end keys (20298), as well as other special keys (198711). In addition, not all key *combinations* have a standard representation (79374) (although that situation appears to be improving) while others overlap (226327). Different keyboard layouts change what is produced when a physical key is pressed. Opinions differ on whether shortcuts should be bound to the physical key (91355, 214273).

C. The Terminal as Dutiful Scribe

There are many “full-screen” terminal applications but interactions with the terminal are often much less complicated. The terminal is frequently used in a command and response scenario that some refer to as typescript-style interaction [54]. In this scenario, the role of the terminal is that of a dutiful scribe — preserving a record of commands issued and their output.

1) *scrollback* (posts: 73, views: 3,304,743): Scrollback, a record of commands issued and their output along with previous prompts, can become quite cluttered (107080) with users looking for techniques to quickly “scroll to the last command” (287541*), search scrollback (12848), save “terminal output to a file” (200637), distinguish normal output from errors (290696), and provide additional context such as “exit status code after each command” (703918). Sensitive information may be used when issuing a command or appear as part of a command’s output. The question then becomes how to properly purge this record (26975).

When it is known that a command will produce many *pages* of output a *pager* can be used. Ideally a pager would only be invoked when required (245064) — a similar effect can be achieved by configuring a pager to quit when it determines it is not required (107315). Some commands like `git` and `man` will pipe their output to a pager by default but there are those who wish they did not (23394). Pagers also come with their own set of issues (23304, 96635, 104499, 134385, 158583, 169886, 197199*, 257788, 412060, 429466).

Pagers and other “full-screen” terminal applications must implement and manage their own scrollback if required (145181). These applications often switch to the terminal’s “alternate screen” which has no scrollback. If the alternate screen is disabled or not used, the most recent screen contents for a full-screen application will become part of the scrollback buffer when the application terminates [55].

⁵key-binding question

In the more typical scenario, the screen appears to clear when a full-screen application switches away from the alternate screen as it terminates (27941, 60533). While enough seem to prefer the screen “clearing” behavior (60499) that it has largely become the default there are situations where restoring the “previous screen upon exit” is not desired (59203). Use of the alternate screen can be inconsistent across applications and systems resulting in the need to “configure the restore behavior consistently” (85398). Accidentally switching to the alternate screen may result in users wondering why “[s]crolling is disabled all of a sudden” (259922).

In a terminal, previous lines scroll up, mimicking paper on a teletype. As a result, the terminal prompt has a tendency to migrate to the bottom of its window. There are some who wish this orientation could be reversed “like a normal terminal only upside down” (324878) with the prompt “at the top of the terminal” (307177*). Even though terminals mostly deal with “just text”, the volume and the speed at which that text can be produced mean that even modern terminal emulators have performance issues to consider (41225) with many benchmarks focusing on scrolling speed [56]. One issue that remains unsolved is how to easily use “text from previous commands’ output” (385) with, or when constructing a new command (see the division of labor discussed in Section IV-C2).

a) **flow-control (posts: 7, views: 1,015,024)** An older solution, called flow control [57], allows users to “pause the output” (294625) of the terminal. Useful in some cases, this feature is the cause of posts asking how to “unfreeze [the terminal] after accidentally pressing Ctrl-S” (12107), wondering why the feature exists at all (137842), and asking how to “permanently disable” it (332791).

2) **command-history (posts: 35, views: 2,533,145):** The division of labor that has evolved between the terminal and the shell is most likely not what would be designed from scratch. The terminal’s canonical mode provides simple line editing but shells and other command-line interfaces often implement their own, more advanced, command history, line editing, and completion. A common desire is to “maintain the same history across multiple terminals” (131504) but opinions differ on when the synchronization should occur (212646). There are users who prefer to not share history across terminals (295165) and some who would prefer a “[p]er-directory history” (71189). Also common is the desire to save more than the default number of commands (20861). With command history being a rolling buffer there is sometimes a desire to “save the last command” (38072*) to a separate, more permanent location. Searching command history is possible but how “to cycle through” (73498) results or continue searching (47814) is not intuitive for some users, and we can see others looking for alternate behavior, such as searching “for a previous command with [what was] just typed” (231605). Different shells (as well as other command-line interfaces) have different default command history behavior, and customizing command history behavior can be difficult because of the amount of shell knowledge (265957) and in some case line editing library knowledge (420356) required.

Sensitive information may be used when issuing a command. A natural question is how to “temporarily suspend history tracking” (10922). The way this is typically implemented can cause confusion when encountered by accident with users wondering why “commands that start with spaces” (115917) are not being added to the command history. Similarly, users need a way to “clear the history for the current session” (544373) or “close a terminal without saving the history” (25049). Some commands are more dangerous than others, and simply recording them in command history makes it possible to inadvertently run them again (92742*). There are also those who would like a way to “save [a] command for later use” without executing it (569941*).

3) **prompt (posts: 24, views: 1,749,193):** Customizing the shell prompt is a gateway to many terminal issues (31695), through the peculiarities of the Unix shell (346924), and sometimes of a particular shell such as how `bash` (105926) or `zsh` (90772*) calculate the prompt’s display length. Some users are frustrated by the prompt’s tendency to move around and wish that the prompt was always “at the bottom” (153102) or “at the top” (218323) of the terminal (see also Section IV-C1). The line-oriented nature of the terminal and command-oriented nature of the shell can cause confusion with users wondering why they are being “repeatedly prompted with >” (158997). When the prompt is constructed dynamically, there are performance issues to consider (565905).

4) **line-editing (posts: 31, views: 529,774):** With the customization possible, it is interesting that most use the `emacs` or `vi` line-editing keybindings. Sometimes this is by default. Even when the possibilities are known (85390) there is the further step of knowing how to apply the required configuration (43003). These keybindings are internally consistent and have a long history but are inconsistent with most other, non-terminal, applications.

Line-editing libraries [58]–[61] allow for consistency across terminal applications (22740). However, customizations are then specific to each line-editing library. There are inevitable differences (252546) and a just as inevitable desire for “the same behavior” (264791) or at the very least the same functionality (42629). Not helping the situation is a division of labor that sees some sequences handled by the terminal and others by the terminal application leaving users wondering why “the arrow keys don’t work” (21733) or how to “erase a mistyped invisible password” (50493*). The command-oriented nature of the terminal makes it difficult to execute commands with fewer than two keystrokes (89622, 672295) particularly when substantially different applications may use the same line-editing library.

5) **completion (posts: 17, views: 326,981):** One of the most important features of modern line-editing libraries is completion, or tab-completion, as it is often called, after the key typically used to invoke it. Under the hood, there is a fair bit of machinery that makes completion possible which makes failures more difficult to debug (499195). Again, different applications and libraries implement this feature in different ways (24419*).

Interacting with the completion machinery can be difficult. For example, cancelling “completion, but only completion” (91161), knowing how to “redirect command completion output” (127506*), or how to “create a completion script” for a custom command (136466*). The division of labor that results in completion being implemented by terminal applications rather than the terminal itself also means that the completion machinery cannot (easily) access the contents of the scrollbar buffer (29878, 41886).

D. Context, Environment, and the Working Directory

1) **contextual-behavior** (*posts: 26, views: 589,235*): According to the Unix philosophy, we should expect “the output of every program to become the input to another, as yet unknown, program [and not] clutter output with extraneous information” [62]. In practice, a program may want to provide some “extraneous information” to make the output more human-readable (as well as potentially performing other operations) when it knows that it “is ‘connected’ to a terminal” (401934). This modal behavior adds complexity but is undeniably convenient and can even be used to stop common mistakes (638221) (see also Section IV-E1). It is also not without precedent — Unix has other short-circuits (29964) and contextual behavior baked in (515778). When a program defaults to the wrong mode users may wish to “trick a command into thinking its output is going to a terminal” (249723) or vice-versa (68574). Some commands side-step this by sending information more likely to be consumed by a human to the stream more likely to be connected to a terminal (166359). In addition to detecting the presence of a terminal, a program may wish to know the specific type of terminal (47037).

2) **environment-variables** (*posts: 36, views: 1,607,188*): The `TERM` environment variable can be inspected but, much like the HTTP User-Agent header, it can be set to anything and so cannot be completely trusted (23763). The `TERM` environment variable functions more as a hint to those terminal-based programs that choose to inspect it (43951). Other environment variables play important roles. The `EDITOR` and `VISUAL` environment variables allow users to “set the default [terminal-based] editor” (501862) and the `PATH` environment variable determines which commands can be invoked easily (45578). Many users have questions about the best “way to set environment variables” (88201) as well as where and how the initial values for environment variables are determined (198794, 246751, 277944*). Approaches like `path_helper` (22979) and `environment.d` (317282) have the potential to alleviate some of this confusion but in the short term they add additional complexity. Given the importance of environment variables in the terminal environment there is an understandable desire to open a new terminal with an environment “that is an exact clone” (6097) of the current environment or to revert to a last known good environment by clearing “out all variables without closing terminal” (172655).

3) **working-directory** (*posts: 16, views: 1,011,645*): Like all Unix processes, the terminal emulator process has a current working directory.

However, when asking if there is a way to open a “new terminal window with the same directory as the previous window” (327313), open a terminal in “a directory other than `$HOME`” (140602), change directories (81224), or be notified when entering a specific directory (18532) it is really the working directory of the application running in the terminal (usually a shell) that is being discussed.

E. Control Sequences, UTF-8, and the Shell

1) **control-sequences** (*posts: 47, views: 1,269,176*): Enabling much of the terminal’s functionality are the in-band control sequences created for the DEC VT series of terminals [63]. A subset of these were standardized [64] and are often referred to as ANSI escape codes or sequences [65]. The xterm terminal emulator supports a “kind of amalgam of ANSI and the VT-whatever standards” (5802) plus extensions [66], [67]. Later terminal emulators followed xterm’s lead. As with useful features in other applications, control sequences can, unfortunately, be abused [68] with a particular implementation exposing a vulnerability [69]. There is an obvious desire to “avoid escape sequence attacks” (15101). Thankfully, as with other products, security issues are often fixed soon after being found, and work to find additional vulnerabilities continues [70]. While not as severe as security-related issues, binary data may contain unintended control sequences which, if accidentally sent to the terminal, can wreak havoc on terminal settings and leave users wondering how to fix the “terminal after displaying a binary file” (79684*) or how to “prevent random console output from breaking the terminal” (247999). To avoid these problems some applications avoid outputting “binary” data unless explicitly instructed to do so (217936) (see also Section IV-D1). Even when the terminal is functioning as intended, misbehaving or misused applications may cause unexpected behavior (164944, 469770).

As the “escape” in “escape code” or “escape sequence” implies the vast majority of control sequences start with an escape character. Unfortunately, it is also possible to generate the escape character by pressing the escape key which exists on most keyboards. To distinguish between the escape key being pressed and the start of a control sequence many applications delay processing the escape character (318433). While the control sequences supported by terminal emulators have converged to support popular terminal-based applications, a terminal-based application may wish to “programmatically check” (228369) “what features a terminal supports” (390360).

2) **encodings** (*posts: 29, views: 615,415*): Many see 2008 as the year when UTF-8 became the rule rather than the exception [71]. This shift seems to have solved many previous terminal encoding issues. However, there are still some issues that remain. Among the remaining issues is the question of how to easily “type Unicode characters” (12244) beyond the limited set that are easily possible from the keyboard. A related issue is understanding the byte sequences that correspond to Unicode characters and how to represent these in other contexts (530788*). Knowing “the display width of a string of characters” (245013) is not simply about counting bytes.

Not all languages are left to right (100811, 167856). Even with UTF-8 as the default, there are issues with non-printing characters and legacy control codes inherited from ASCII that are “invisible in terminal output” (128019, 723082) as well as files in legacy formats (184538). Although support for UTF-8 is now nearly universal there is an understandable desire to be able to detect a terminal’s level of Unicode support (184345).

3) *quoting (posts: 11, views: 475,872)*: As the most common terminal application, the shell determines how characters are interpreted, which are best avoided (148043), and what is required to “use a special character” (296141). Misunderstanding how the shell operates can easily cause confusion (284638, 446193). In an effort to avoid some of this confusion, some terminal-based programs have started producing shell-friendly output (258679) (see also Section IV-D1).

V. DISCUSSION

Despite its deceptively simple appearance, the terminal is a large and complex application. Using the core Wayland window system code [72] as a point of reference, at just over 100,000 lines of code the venerable xterm [73] is approximately 2.5x larger than the core Wayland window system code. Even `foot`, a relatively new, “lightweight and minimalistic Wayland terminal emulator” [74] is larger than the core Wayland window system code itself. The terminal’s size, complexity, and long history result in what can seem like a bewildering array of issues. We have grouped and arranged these issues in an attempt to make them more digestible. We also believe that several themes emerge.

With the creation of the windowed terminal emulator the environment of the terminal changed. We claim that the terminal, as windowed terminal emulator, and the applications it supports, have yet to fully **adapt** to this new environment. We see evidence for this in posts asking why the mouse cannot be used to position the cursor when editing a command or why the terminal’s keyboard shortcuts for copy and paste are different from most other applications. The long history of the terminal includes the accumulation of quirks and the evolution of conventions that often differ from conventions established outside the terminal. While users can and do adapt, this raises the question of how much of this is inherent and how much is accidental and avoidable. We see further evidence in posts asking if there is some way to make file names in terminal output clickable by default, as well as in the many posts looking for ways to integrate the terminal with other applications or asking for ways to synchronize command history or environment variable changes across multiple terminals.

While the terminal has not yet fully adapted to its new graphical environment neither, we claim, has it managed to completely **untangle** itself from the video-display terminal era. Instead of a video-display terminal with one login shell, there are now often multiple terminal windows for the same user. The `jobs` command (even though it is provided by the shell) is particularly revealing here, in that it only shows jobs for the current shell in the current terminal and not all jobs on the current system for the current user (509147).

The line-oriented nature of the terminal and the command-oriented nature of interpreters like the shell forces each of these applications to re-implement much of the same functionality. The division of labor that has evolved between the terminal and the shell also makes it more cumbersome than one might expect to use the output of previous commands when constructing new commands. In general, it seems strange that terminal pagers are still required and that terminal features for interacting with scrollbar are so minimal.

Terminal scrollbar and command history are areas that seem underdeveloped and we believe more could be done to **extend** these features. This area has seen some activity recently: an improvement pioneered by FinalTerm and now adopted by many terminals, the `OSC 133` control sequences [75] allow command interpreters to mark the beginning of the prompt, the beginning of a command, the start of the command output, and the end of the command output. This enables terminals with support for this feature to quickly scroll to these marks. In addition to better navigation and searching of scrollbar and command history, the terminal could also provide more context in scrollbar as well as do more to distinguish regular output from errors. We believe the terminal should also do more to prevent errors caused by misbehaving applications or unintentional control sequences.

More terminal application-related than terminal-related, we see a number of opportunities to **simplify** the terminal experience. These range from promoting better defaults to investigating how to standardize feature detection to developing tools to translate configuration into the myriad ad hoc settings required by various terminal-based applications.

Based on aggregate views it is clear that integration with other applications, job control, terminal multiplexing, and accessing remote systems are important use cases. Attempts to improve the terminal ignore these at their peril. More than that, any attempt to improve the terminal must take great care to preserve the vast majority of the terminal’s functionality. We know from Hyrum’s Law that “all observable behaviors of [a] system will be depended on by somebody” [76]. Terminal-based applications have had half a century to develop a dependence on the terminal’s behavior (105325, 278884) - both documented and undocumented.

Finally, we see that some terminal-related issues are a result of not knowing where to look for more information. We hope that this paper may play a small part in rectifying that. We would also like to draw attention to resources like Thomas Dickey’s *XTerm Control Sequences* [66] and articles like Linus Åkesson’s “The TTY demystified” [57]. While terminal-related documentation does exist we believe what makes understanding the terminal more difficult is how this documentation often has to be pieced together as part of a historical record and that what is missing is comprehensive documentation of the terminal as it now exists.

In Section II we mentioned that a goal many current efforts to improve the terminal seem to share is improving the social and collaborative aspects of the terminal. We see little evidence of that as a prominent pain point.

VI. LIMITATIONS

This paper is based on questions posted on U&LSE and so inherits the biases of that community. As discussed previously, there are other similar sources of information but we believe U&LSE to be the most relevant for the discussion in this paper. Our approach could be applied to other datasets to investigate how much the choice of U&LSE affects our results.

Our findings are based on a subset of questions from U&LSE. It is possible that this subset omits relevant questions. Pragmatically, it is not possible to evaluate every question. To mitigate against this threat our goal was to make our selection and categorization rules as clear as possible. In addition, our selection and categorization process can be reproduced using the scripts in our replication package.

Using regular expressions we were able to categorize 91.0% of the posts identified by our selection process without manual intervention. Where possible we based our patterns on existing tags or common words or phrases occurring in more than one post. In cases where this was not possible, due to tags lacking in discriminatory power or idiosyncratic phrasing, we based our patterns on unique and identifiable passages still recognizable as being related to the category in question. Additionally, we found that the patterns we developed to identify a subset of all unrelated posts were sufficiently accurate to be used as a pre-filter. We exploited this characteristic to simplify the patterns needed for all other categories. The goal throughout the process was patterns that are easily understandable and stable. We do not claim that the patterns used are an optimal or minimal set.

It is possible for a question to belong to more than one category or for a post to contain more than one question. We found these threats to be more theoretical than practical. Only 7.9% of selected posts were assigned more than one category. In the vast majority of cases, there was a clearly dominant category with additional categories aiding rather than hindering understanding of the issue. On the rare occasion where a post contained more than one question (94299, for example) we did not find our analysis affected.

It is possible to disagree with the categories we have chosen as well as our interpretation of the results. While we cannot anticipate any and all objections we have made our replication package available to foster transparency and enable others to reproduce, examine, and experiment with our approach and our results.

VII. CONCLUSION

At one point in time, it was a commonly held belief that the terminal would go away — replaced by sleeker interfaces. There are some who continue to hold this belief. Exceeding the strength of this belief is the terminal’s stubborn refusal to do so. We believe the terminal has persisted because it provides a uniquely productive interface for many tasks. We also believe, and believe that this paper demonstrates, that a requirements-based understanding of how the terminal should evolve is possible.

Terminal lucidity “is an unexpected return of consciousness, mental clarity or memory shortly before death” [77]. We are not referring here to the death of the terminal. Far from it. The terminal’s popularity seems, if anything, to be increasing. We are referring to the death of prejudices and preconceived notions about the terminal that are increasingly hard to maintain. As these fall away, we hope it will be possible to see the true potential of the terminal clearly.

Acknowledgements

The authors thank Thomas Dickey for reviewing an early draft and suggesting ways to improve readability. The authors also gratefully acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC, RGPIN-2021-03809) and the Swiss National Science Foundation (SNSF, Project “FORCE”, SNF Project Number 232141).

VIII. APPENDIX

The terminal is often mentioned but not the focus in questions about popular terminal-based applications like `ssh` (48863, 71481, 85920, 108679, 207365, 317491, 321968, 323925, 352298, 371901, 758893) or `vim` (11369, 25965, 43033, 186166, 345494) or applications like `emacs` which may run in a terminal (8424, 174413). There are also questions about the Linux (or FreeBSD, etc.) console (16861, 23363, 45228, 49779, 485156). or windowed terminal-emulator adjacent applications like PuTTY (20257, 143684) that are not relevant to the discussion in this paper. Sometimes the terminal is simply mentioned incidentally (3022, 22561, 72657, 150501, 164262, 162979, 233459, 322814, 369097, 439486, 569993) or it, or a terminal-based application, is part of the situation but has no bearing on the problem or solution (14629, 23360, 25998, 34729, 37539, 52055, 146193, 166936). The terminal is mentioned in many questions that are really Unix, scripting, or shell-related questions (10698, 40708, 70615, 71492, 72962, 113556, 134734, 151757, 153357, 153508, 155838, 156084, 157689, 166853, 189254, 199085, 245746, 251691, 283586, 297729, 377618, 382191, 411688, 494378, 498685, 599013).

In addition to the above questions, manually categorized as unrelated, 13 questions were manually corrected and assigned to the following categories: control-sequences (105325), encodings (723082), environment-variables (23763), integration (152310, 234136, 268451, 316401), key-bindings (214273, 226327, 262129), prompt (262129), remote (3212), and scroll-back (12439).

REFERENCES

- [1] “Wikipedia: Windows Services for UNIX,” https://en.wikipedia.org/wiki/Windows_Services_for_UNIX, Dec. 2023.
- [2] K. Cinnamon, “Windows terminal is now the default in windows 11,” <https://devblogs.microsoft.com/commandline/windows-terminal-is-now-the-default-in-windows-11>, Oct. 2022.
- [3] J. Nickoloff and S. Kuenzli, *Docker in Action*, 2nd ed. Manning, 2019.
- [4] M. Lukša, *Kubernetes in Action*. Manning, 2017.
- [5] V. Nguyen, “How to create a custom language model,” <https://developer.nvidia.com/blog/how-to-create-a-custom-language-model/>, Mar. 2023.
- [6] “Zed,” <https://zed.dev/>, Mar. 2024.
- [7] “Installing Go,” <https://go.dev/doc/install>, Mar. 2024.

- [8] "Installing Rust," <https://www.rust-lang.org/tools/install>, Mar. 2024.
- [9] "Integrated terminal in visual studio code," <https://code.visualstudio.com/docs/terminal/basics>, Feb. 2024.
- [10] "Konsole," <https://konsole.kde.org/>, Mar. 2024.
- [11] "Get to know the android studio ui," <https://developer.android.com/studio/intro/user-interface>, Apr. 2023.
- [12] "The new terminal (beta) is now in jetbrains ides," <https://blog.jetbrains.com/idea/2024/02/the-new-terminal-beta-is-now-in-jetbrains-ides/>, Feb. 2024.
- [13] "Android studio command-line tools," <https://developer.android.com/tools>, Apr. 2023.
- [14] N. Stephenson, "In the beginning was the command line," <https://www.cryptonomicon.com/beginning.html>, 1999.
- [15] A. A. A. Donovan and B. W. Kernighan, *The Go Programming Language*. Addison-Wesley, 2015.
- [16] J. Edelman, S. S. Lowe, and M. Oswalt, *Network Programmability and Automation*. O'Reilly Media, 2018.
- [17] T. R. Hein, E. Nemeth, G. Snyder, B. Whaley, and D. Mackin, *UNIX and Linux System Administration Handbook*, 5th ed. Addison-Wesley, 2017.
- [18] B. Gregg, *Systems Performance*, 3rd ed. Pearson Education, 2020.
- [19] J. Janssens, *Data Science at the Command Line*, 2nd ed. O'Reilly Media, 2021.
- [20] S. Suehring, *Learning DevSecOps*. O'Reilly Media, 2024.
- [21] S. Freire, F. Gomes, L. Barbosa, T. S. Mendes, G. Reges, R. S. Maciel, M. Mendonça, and R. Spínola, "Requirements engineering issues experienced by software practitioners: A study on stack exchange," in *International Working Conference on Requirements Engineering: Foundation for Software Quality*. Springer, 2023, pp. 3–20.
- [22] M. Papoutsoglou, G. M. Kapitsaki, D. German, and L. Angelis, "An analysis of open source software licensing questions in stack exchange sites," *Journal of Systems and Software*, vol. 183, p. 111113, 2022.
- [23] H. Gujral, S. Lal, and H. Li, "An exploratory semantic analysis of logging questions," *Journal of Software: Evolution and Process*, vol. 33, no. 7, p. e2361, 2021.
- [24] A. Tahir, J. Dietrich, S. Counsell, S. Licorish, and A. Yamashita, "A large scale study on how developers discuss code smells and anti-pattern in stack exchange sites," *Information and Software Technology*, vol. 125, p. 106333, 2020.
- [25] M. MacInnis, O. Baysal, and M. Lanza, "Terminals all the way down," in *Proceedings of the ACM/IEEE 44th International Conference on Software Engineering: New Ideas and Emerging Results*, 2022, pp. 86–90.
- [26] C. Muth, "Xiki," <https://github.com/trogdoro/xiki>, Jan. 2020.
- [27] "Fig user manual," <https://fig.io/manual>, Mar. 2024.
- [28] "Fig," <https://fig.io/>, Mar. 2024.
- [29] E. Zhang, "sshx on /t/rust/," https://www.reddit.com/r/rust/comments/17oejp6/i_made_sshx_an_app_that_lets_you_share, Nov. 2023.
- [30] —, "sshx," <https://github.com/ekzhang>, Apr. 2024.
- [31] "Warp," <https://www.warp.dev/>, Feb. 2024.
- [32] "Wikipedia: Stack Exchange," https://en.wikipedia.org/wiki/Stack_Exchange, Feb. 2024.
- [33] H. Alharthi, D. Outioua, and O. Baysal, "Predicting questions' scores on stack overflow," in *Proceedings of the 3rd International Workshop on CrowdSourcing in Software Engineering*, ser. CSI-SE '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 1–7.
- [34] A. Strauss and J. Corbin, *Basics of Qualitative Research: Grounded Theory Procedures and Techniques*. Sage Publications, Jan. 1990.
- [35] A. Hindle, E. T. Barr, Z. Su, M. Gabel, and P. Devanbu, "On the naturalness of software," in *Proceedings of the 34th ACM/IEEE International Conference on Software Engineering*, 2012, pp. 837–847.
- [36] "Wikipedia: Folksonomy," <https://en.wikipedia.org/wiki/Folksonomy>, May 2024.
- [37] Z. Cai, C. Marquart, B. Eagan, Y. Xiao, and D. W. Shaffer, "A lightweight interactive regular expression generator for qualitative coding in quantitative ethnography," in *Advances in Quantitative Ethnography*, G. A. Irgens and S. Knight, Eds. Cham: Springer Nature Switzerland, 2023, pp. 455–469.
- [38] Z. Cai, B. Eagan, and D. W. Shaffer, "Negative reversion: Toward intelligent co-raters for coding qualitative data in quantitative ethnography," in *Advances in Quantitative Ethnography*, G. A. Irgens and S. Knight, Eds. Cham: Springer Nature Switzerland, 2023, pp. 426–437.
- [39] A. Bacchelli, M. D'Ambros, M. Lanza, and R. Robbes, "Benchmarking lightweight techniques to link e-mails and source code," in *16th Working Conference on Reverse Engineering*. IEEE, 2009, pp. 205–214.
- [40] S. Beyer, C. Macho, M. D. Penta, and M. Pinzger, "What kind of questions do developers ask on stack overflow? a comparison of automated approaches to classify posts into question categories," *Empirical Software Engineering*, vol. 25, pp. 2258–2301, 2020.
- [41] J. N. Wulff, G. B. Sajons, G. Pogrebna, S. Lonati, N. Bastardoz, G. C. Banks, and J. Antonakis, "Common methodological mistakes," *The Leadership Quarterly*, vol. 34, no. 1, p. 101677, 2023.
- [42] W. R. Stevens and S. A. Rago, *Advanced Programming in the UNIX Environment*, 3rd ed. Addison-Wesley, 2013.
- [43] "Gnu screen," <https://www.gnu.org/software/screen/>, Apr. 2016.
- [44] "tmux," <https://github.com/tmux/tmux>, Oct. 2021.
- [45] K. Goyal, "Kitty tmux-like daemon response," <https://github.com/kovidgoyal/kitty/issues/391#issuecomment-638320745>, Jun. 2020.
- [46] "Wikipedia: Semigraphics," <https://en.wikipedia.org/wiki/Semigraphics>, Jan. 2024.
- [47] "Wikipedia: Sixel," <https://en.wikipedia.org/wiki/Sixel>, Jan. 2024.
- [48] "Wikipedia: ReGIS," <https://en.wikipedia.org/wiki/ReGIS>, Jun. 2023.
- [49] K. Goyal, "Terminal graphics protocol," <https://sw.kovidgoyal.net/kitty/graphics-protocol/>, Jul. 2021.
- [50] C. Lindsey, "The secondary selection," <https://www.cs.man.ac.uk/~lin/dsec/secondary-selection.html>, Oct. 2018.
- [51] D. DeVault, "Wayland misconceptions debunked," <https://drewdevault.com/2019/02/10/Wayland-misconceptions-debunked.html>, Feb. 2019.
- [52] "Wikipedia: Bracketed Paste," https://en.wikipedia.org/wiki/Bracketed_Paste, Jan. 2023.
- [53] A. Baretta, "Consistent backspace and delete configuration," <https://web.archive.org/web/20181021233554/http://www.ibb.net/~anne/keyboard/keyboard.html>, Oct. 2018.
- [54] R. Pike, "Acme: A user interface for programmers," in *USENIX Winter*, 1994, p. 223–234.
- [55] T. Dickey, "Xterm faq - why doesn't the screen clear when running vi?" https://invisible-island.net/xterm/xterm.faq.html#xterm_tite, Dec. 2022.
- [56] "Performance - warp documentation," <https://docs.warp.dev/how-does-warp-compare/performance>, Feb. 2024.
- [57] L. Åkesson, "The tty demystified," <https://www.linusakesson.net/programming/tty/>, Jul. 2008.
- [58] C. Ramey and B. Fox, "Gnu readline," <https://tiswww.case.edu/php/chet/readline/rltop.html>, Jan. 2024.
- [59] J. Wiberg, "Editline," <https://github.com/troglobit/editline>, Dec. 2023.
- [60] S. Sanfilippo, "Linenoise," <https://github.com/antirez/linenoise>, Feb. 2024.
- [61] P. Harris, "Liner," <https://github.com/peterh/liner>, Jan. 2022.
- [62] M. D. McIlroy, E. N. Pinson, and B. A. Tague, "Unix time-sharing system: Foreword," *The Bell System Technical Journal*, vol. 57, pp. 1899–1904, 1978.
- [63] "Wikipedia: VT100," <https://en.wikipedia.org/wiki/VT100>, Dec. 2023.
- [64] "Ecma-48: Control functions for coded character sets," <https://www.ecma-international.org/publications-and-standards/standards/ecma-48/>, Jun. 1991.
- [65] "Wikipedia: ANSI Escape Code," https://en.wikipedia.org/wiki/ANSI_escape_code, Jan. 2024.
- [66] T. Dickey, "Xterm control sequences," <https://invisible-island.net/xterm/ctlseqs/ctlseqs.html>, Feb. 2024.
- [67] "Escape sequences," <https://bjh21.me.uk/all-escapes/all-escapes.xhtml>, Oct. 2020.
- [68] H. Moore, "Terminal emulator security issues," <https://hdm.io/writing/termulation.txt>, Feb. 2003.
- [69] E. Gerzi, "Don't trust this title: Abusing terminal emulators with ansi escape characters," <https://www.cyberark.com/resources/threat-research-blog/dont-trust-this-title-abusing-terminal-emulators-with-ansi-escape-characters>, Jan. 2022.
- [70] D. Leadbeater, "ESC[31m"?! ANSI terminal security in 2023," <https://dgl.cx/2023/09/ansi-terminal-security>, Oct. 2023.
- [71] "Wikipedia: UTF-8," <https://en.wikipedia.org/wiki/UTF-8>, Feb. 2024.
- [72] "Wayland," <https://gitlab.freedesktop.org/wayland/wayland>, Feb. 2024.
- [73] T. Dickey, "xterm," <https://invisible-island.net/archives/xterm/>, Feb. 2024.
- [74] D. Eklöf, "foot," <https://codeberg.org/dnkl/foot>, Feb. 2024.
- [75] "iterm2: Proprietary escape codes," <https://iterm2.com/documentation-escape-codes.html>, Jun. 2024.
- [76] H. Wright, "Hyrum's law," <https://www.hyrumslaw.com/>, Sep. 2017.
- [77] "Wikipedia: Terminal Lucidity," https://en.wikipedia.org/wiki/Terminal_Lucidity, Feb. 2024.