

Computers and Challenges of Writing in Persian: A Personal History Spanning Five Decades

[Version of October 7, 2018]

Behrooz Parhami

Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106-9560, USA
parhami@ece.ucsb.edu

Abstract — The Persian script has presented some difficulties, ever since printing presses were introduced in Iran in the 1600s. The appearance of typewriters created additional problems and the introduction of digital computers added to the design challenges. These difficulties persisted, until high-resolution dot-matrix printers and display devices offered greater flexibility to font designers and the expansion of the computer market in the Middle East attracted investments on improving the Persian script for computers. Nevertheless, certain peculiarities of the Persian script have led to legibility and aesthetic quality issues to persist in many cases. In this paper, I enumerate some of the features of the modern Persian script that made it a poor match to implementation on modern technologies and review the challenges presented by, and some of the solutions proposed for, each new generation of computer printers and display devices. Interestingly, the same features that make legible and pleasant printing/displaying difficult also lead to challenges in automatic text recognition. I conclude with an overview of current state of the art and areas that still need further work.

Keywords — Aesthetics; Arabic script; Computer display; Computer printing; Dot-matrix font; Font design; Keyboard; Legibility; Line-printer; Movable type; Persian script; Printing; Technology transfer; Typewriter

Acknowledgements — This paper arose from a pair of lectures delivered on November 19 (Persian) and November 20 (English), 2017, as part of the Bilingual Lecture Series on Iran, under the auspices of UCLA's Iranian Studies Program. Staying true to the spirit of bilingualism, I have prepared two versions of this paper in Persian and English. My sincere thanks go to the UCLA lecture series' coordinator, Dr. Nayereh Tohidi, and the program's administrators and generous sponsors. While I played a central role in formulating and categorizing the problems discussed here and keeping track of advances and developments for reporting purposes, solving the said problems was most certainly not a one-person effort. Dr. Farhad Mavaddat and Eng. Armen Nahapetian were two of my colleagues in Iran who played key roles during the early years of studying and solving the script-related problems. Our graduate students, as well as researchers/developers at other universities, government organizations, and private companies also contributed to the solutions. As a final side note, my use of the word "Persian" covers a set of related languages that Iranians call "Farsi" and other speakers refer to as "Dari," "Pashto," or "Urdu," depending on their geographic regions and ethnic backgrounds.

In this paper, I summarize the results of nearly two decades of research and development, while I was working in Iran, and a few years thereafter, on adapting computer technology to the Persian script. The work was a team effort involving multiple colleagues at my university, researchers/developers at other organizations that were developing computer applications in Iran, and several computer vendors targeting the Iranian market. My motivation is to create a formal recorded history of these efforts and, more importantly, present a picture of the state of the art in Persian printing and display technologies that would serve as a baseline for charting a path to solving a myriad of tough problems that remain unsolved or poorly solved after five decades.

In fact, the original title for my UCLA talks was “Fifty Years of Poor Penmanship: How Computers Struggled to Learn the Persian Script.” The fifty-year period alluded to in this alternate title began in the late 1960s, when I was an undergraduate student at Tehran University’s College of Engineering. In those days, Tehran University had just installed an IBM 1620 at its computer center across the street from the main campus in an effort to mechanize administrative data processing tasks and also offer the students access to computing capabilities. At about the same time, major banks and large governmental organizations were developing customer databases, as well as payroll and personnel information systems, all of which needed to deal with information, such as employee/customer names and street addresses, in Persian.

A lot has happened in the five decades since then. My start in computing, when I was a junior engineering student in 1967, was learning to program in the Fortran language, preparing programs on punched cards and receiving the output on scrolls of fan-fold paper, a day or so later, iterating until program bugs had been removed. No Persian printing was involved at the time and there was no display unit for us student users. Organizations that did use the Persian script at the time, had to tolerate very low-quality output. The Persian script produced by computer printers and display devices has improved substantially, thanks to the diligence of the early researchers and the rapid development of digital computing in Iran (and Arab countries, using pretty much the same script), which attracted investments and technical know-how to solving the problems.

I hope I succeed in conveying the scope of the difficulties in the early years and the hard work and ingenuity that went into solving them in many iterations, while at the same time adapting to rapid changes in computer printing and display technologies.

1. Computers and the Persian Language

When I worked at Sharif (formerly Arya-Mehr) University of Technology in Tehran, Iran, during 1974-1988 (the last two years on sabbatical leave in Canada), one track of my research dealt with technology transfer. Specifically, my colleagues and I were interested in adapting the then-new computer technology to Iran's culture, economy, and language [Parh77b]. We preferred to use the term "technology transplantation" to better highlight the challenges of adaptation and acceptance. If not introduced in the right way, a society may reject a new technology, in much the same way that human body could reject a new organ that is ill-suited to its expectations, needs, and immune response. This is where the notion of "appropriate technology," that matches the social and economic conditions of a particular geographic area and promotes self-sufficiency [Dunn79], while also taking environmental factors into account, comes into account.

Considerations of culture and language issues in connection with computers spans many themes and disciplines, as depicted in Fig. 1. I won't have much to say in this paper about cultural adaptation (the leftmost branch in Fig. 1), except to point out that the design of suitable native educational programs is an important part of the process [Parh85] [Parh86]. In the late 1970s and early 1980s, I wrote a number of columns for Iranian newspapers and trade publications on the various aspects of cultural and linguistic adaptation (e.g., [Parh78c]). I also contributed a steady stream of articles to *Computer Report (Gozarash-e Computer)* the monthly magazine of the Informatics Society of Iran. I will endeavor to collect these thoughts into a separate article.

Adapting computers to the Persian language (the right-hand side of Fig. 1) is where I spent much of my effort while working in Iran. The first item of business in this domain is to make sure we can talk about computers, computer technology, and various application domains, using an appropriate terminology within the framework of the Persian language [Parh81b]. Coming up with new terms, while ensuring technical and linguistic correctness, as well as acceptability by the public, is a rather difficult process, which is still ongoing in Iran. Many precise and linguistically-sound proposed terms (e.g., "raayaaneh" for "computer") have been pretty much ignored, while others (such as "reez-pardaaazeshgar" for "microprocessor") have been widely accepted. Pragmatism and flexibility must be exercised in dealing with certain established terms (such as "electronic") that have found their ways into many languages.

Next, we need to communicate with computers to provide them with guidelines and instructions for doing what we want them to do. In this domain, language considerations are secondary. While there was some effort in the 1970s to “Persianize” programming languages (particularly the COBOL language, which was verbose and used a large number of English words in its original form), a consensus soon developed that what is difficult in programming isn’t the learning of a few English terms, such as knowing what “while $x > 0$ ” means, but rather the mastering of programming notions and constructs. This line of research was thus abandoned, although it is now returning in the form of Persian voice commands (discussed under input below).

The rightmost branch of the tree in Fig. 1, handling Persian language information, is the most-relevant one as far as this paper is concerned [Parh77a] [Parh78b] [Parh81a]. Processing of Persian information requires that information be input to the computer, using keyboards [Parh84b], natural-language audio, completed forms, or scanned documents fed to optical character recognition systems [Parh81c]. Then, we need information storage and interchange codes [IPBO80], and perhaps data-compression and other storage support mechanisms, to keep the information and send it between various storage and processing devices as efficiently as possible.

The actual processing of information inside the computer brings about the need for software and algorithms that may be language-dependent. A good example is provided by the difficulties in the alphabetical sorting of Persian names, given multiple versions of letters and spelling variations, particularly for composite names. Activities such as content analysis and style-matching need to be undertaken in collaboration with experts in other fields, such as anthropology, law, literature, and sociology. Finally, translation of texts to/from Persian falls in this domain.

Last, but not least, is the problem of output production by computers, as shown in the highlighted box on the bottom-right of Fig. 1. When the computer interacts with the outside world by directly controlling devices or systems, linguistic factors do not come into play, but when the output is printed, displayed, or spoken, language becomes a prominent consideration. We will present a detailed discussion on producing printed and displayed Persian texts in the following sections of the paper, focusing on pertinent considerations, such as the legibility and aesthetic quality of the resulting script. Spoken output has seen very limited research in the case of Persian, but perhaps this situation will change over the next few years.

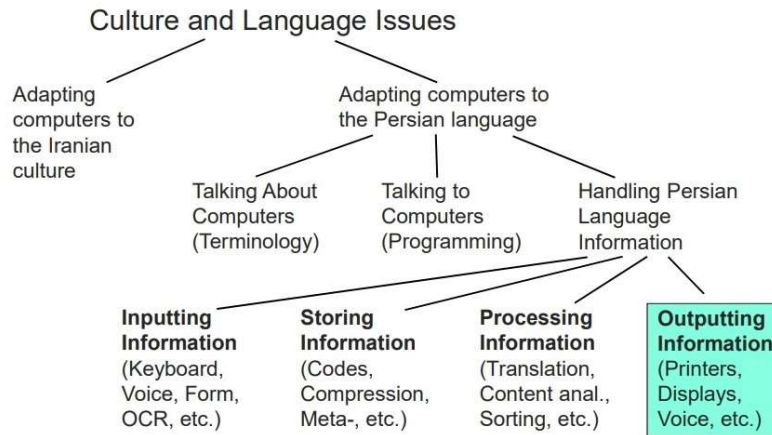


Fig. 1. Technology transplantation: Cultural and linguistic considerations for computer technology.

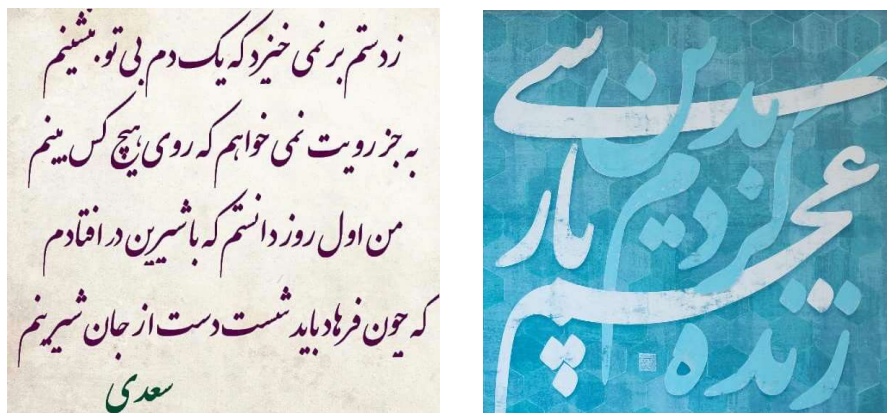


Fig. 2. Uncredited handwritten Nastaliq, left, and calligraphic art by Farokh Mahjoubi.

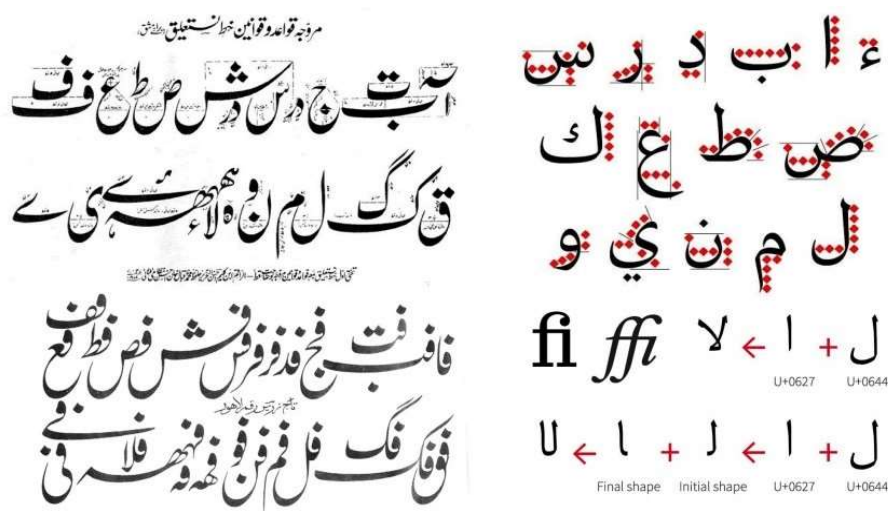


Fig. 3. Persian script: Rules for beautiful writing of letters and letter combinations.

2. A Brief History of the Persian Script

The modern Persian script is believed to be about 1200 years old, when it was adapted from Arabic [Laza08]. The Persian language itself is much older and, along with the script, came in previous versions of Old Persian and Middle Persian before then. For much of this 1200-year history, the script was written manually by scribes, who were hired by poets and others who wanted their writings archived and disseminated. The script used for this purpose was often Nastaliq, written with an eye toward legibility and beauty (Fig. 2, left panel). Each book produced in this way was one-of-a-kind, though a kind of primitive printing, consisting of etching the text on stone or wood and pressing paper or parchment against it after inking, was used on occasion to produce multiple copies.

The vast importance placed by Iranians on beauty in Persian script led to the appearance and quick spread of calligraphic art (Fig. 2, right panel). In calligraphy, the script's legibility is secondary, the prime focus being the proportions and interactions of textual elements and the color scheme. The perfection of this art form led to kings and architects becoming keen on incorporating calligraphic art on mosques, palaces, and other important buildings. In the process, artists developed special scripts that would make it easy to write Persian and Arabic texts, using tiles and other construction elements. In our subsequent discussion of dot-matrix displays and printers, we will return to the use of “pixel fonts,” where a black or white tile represents a pixel in forming the desired text with an array of tiles.

With regard to technological adaptation, the Persian script went through three phases of change, as the printing press, typewriters, and eventually computers entered the scene. The changes were cumulative, in the sense of the changes made in response to the printing press persisting through the adaptation to typewriters, and the latter changes affecting computer printing. Sections 3 and 4 outline the first two phases, 300-400 and 100-120 years ago, respectively, with the third phase (which is our main focus) beginning around 50 years ago, covered in the rest of the paper. Much of what I write about printing the Persian script also applies to generating it on display devices, but I will devote the latter part of Section 7 to the specific challenges and successes in displaying the Persian script. Bear in mind that nearly everything I say about Persian script is also applicable to Arabic.

Before the need arose for adapting the Persian script to modern technology, two major scripts Nastaliq and Naskh, were in common use, alongside many other less-popular ones. Rules for writing beautifully in various scripts were passed on from masters to students. For example, sizes and alternative forms for the letters of the alphabet were often specified in terms of the number of dots, horizontally and vertically (Fig. 3, right panel). Each master had a unique set of rules that set him apart in terms of style and reputation. Just as in English calligraphy, certain combination of letters, such as “fi” and “ffi” are specified as a unit, much like a single letter, Persian calligraphy has its rules for letter combinations, such as “laam-aleph” (the letter “laam” followed by the letter “aleph,” to form the combination “laa,” which looks different from “laam” and “aleph” letters put together), with such combinations being much more prevalent in the Nastaliq script (Fig. 3, left panel).

The Naskh script has become much more common than Nastaliq and has developed into a multitude of forms, because of its better match to the capabilities and limitations of modern display and printing technologies, as discussed in Sections 3 and 4. However, Nastaliq still holds a special place in the hearts of Persian-speaking communities worldwide. For example, the fanciest books of poetry are commonly produced in Nastaliq, which requires a calligrapher to write the entire book by hand, before it goes to offset printing. Flyers, posters, and brochures, which use Naskh-based print scripts for the most part, tend to include bits and pieces of Nastaliq for emphasis, such as in titles and headings, or for decorative purposes. Some progress has been made in producing the Nastaliq script on computer displays and printers, as discussed near the end of Section 7, but the script quality is not yet at a level to threaten the livelihood of master calligraphers.

It is worth noting that book production in Persian has a much longer history than the printing press or typewriters (subjects of Sections 3 and 4, respectively). In 2014, the US Library of Congress sponsored an exhibition, featuring an extensive lecture series, to celebrate 1000 years of the Persian book [LoC14]. An interesting history of publishing in the Persian language exists [Mora13], which despite packing much useful information, unfortunately lacks any photos or diagrams to convey the machines and mechanisms used for printing.

3. The Printing Press Comes to Iran

The printing press arrived in Iran some 400 years ago. Like many other symbols of the West, it is likely that the Shah at the time saw printing technology on one of his European trips and decided that he wanted to have it in Iran. The widespread use of printing machines did not start until much later, perhaps 300 years ago, in part because non-trivial adaptations were required.

The imported technology was based on the use of movable type, rectangular metal blocks each of which has the shape of one letter of the alphabet, digit, or symbols embossed on one surface. Combining the blocks horizontally formed the text to be imprinted on a line (Fig. 4). The typesetter would have in front of him large trays of types, with compartments for different letters and symbols. Multiple trays were needed to house different sets of fonts, such as those with different font sizes, italics, boldface, and so on. Several such lines were completed by the typesetter on a small hand-held tray and, when the small tray was filled, its contents were transferred to a larger tray and the setting of the next chunk of text began. This process continued, until 8 complete pages had been set onto a large tray. These 8 pages would be printed on one side of a large sheet of paper (lower-left panel of Fig. 4). Another 8 pages were set for what was to appear on the back side of the large sheet, with page numbers arranged as shown in the top-left panel of Fig. 5.

The sheet of paper with 16 pages of content would then be folded and put together with other folded sheets to form a complete book, whose number of pages was a multiple of 16 (sometimes necessitating blank pages at the end). These 16-page blocks were sewn together along the book's spine. The edges of the folded sheets were then cut with special guillotines to form the completed book, which then underwent the process of binding to put a hard or soft cover on it. While some books are still put together by sewing the spine, albeit mechanically, rather than by hand, modern book-making usually entails the use of glue, which is cheaper, faster, and results in flatter spines. Imperfect cutting in some older books led to the edges of some pages to remain connected to each other, necessitating manual separation by the reader.



Fig. 4. Hand-typesetting with movable type, the tray holding the typeset text, and the printing process.



Fig. 5. Image transfer to large sheets of paper, folding the sheets, and book-binding.



Fig. 6. Movable Persian type in the compartments of the typesetter's tray.

The first challenge in Persian printing was to make the needed metal blocks holding the letters and other symbols (Fig. 6). This wasn't a straightforward process. Whereas the Latin alphabet consists of individual letters that are set next to each other, with a small gap between them, Persian letters needed to be connected to each other. Furthermore, the handwritten script was not readily decomposable into separate, non-overlapping symbols. Some examples are provided in the left panel of Fig. 7. The word "mojtame'e," in which the first two letters "meem" and "jeh" are almost completely overlapped horizontally, as are the next two letters "teh" and "meem," must be stretched out, so that letters appear side by side, rather than on top of one another.

So, calligraphers were recruited to design new fonts for Persian in which composability was placed ahead of aesthetic quality as a design consideration. The same word "mojtame'e" now appeared as shown in the center-top part of Fig. 7, with arrows pointing to its five constituent letters placed side by side. A similar problem arises for certain wide and/or tall Persian letters, whose handwritten forms might overlap with adjacent letters. In the word "changaal," for example, shown at the bottom-left of Fig. 7, the upper part of the letter "gue" extends to the left to overlap with the letter "noon" that precedes it. All such problems were solved by compromises in font design that traded off aesthetic quality, and in some cases readability, for composability.

Not all features of the Persian script cause difficulties in typesetting: there is also one helpful feature. Many Persian letters can be extended or elongated, as needed to fit the available space. This feature comes in handy when typesetting poems, which are traditionally formatted into two columns consisting of the first and second half-verses, with each column left- and right-adjusted. We see in the example supplied in the right panel of Fig. 7 that the letters "feh" and "yeh" in the circled word "aafaridam" have been elongated to ensure a justified left margin for the poem. The elongation mechanism is a dash-like symbol, that, like Latin dashes, comes in different lengths and is aligned with the horizontal line, the connections axis, on which all inter-letter connections take place. So, any letter to be connected to the preceding letter begins on the connection axis, and any letter to be connected to the following letter ends on the connection axis.

Besides connectivity of letters and horizontal overlap, certain other difficulties are inherent in the Persian script. Unlike the Latin script, in which letters have regular geometric shapes of comparable sizes (top half of Fig. 8), Persian letters are more curved and exhibit wider variations in width and height. Generally speaking, the Persian script is more compact horizontally, in the sense that the Persian translation of an English sentence tends to take up less space horizontally. This horizontal compactness is mainly due to short vowels not being written out in Persian. However, the Persian script takes up more space vertically, in part due to the need for more line spacing to keep the letters with tall lower parts on one line separated from letters on the following line that have tall upper parts. As a result, the metal type blocks for Persian have greater variations in size, making their manual handling by the typesetter somewhat more difficult.

This wider variation in letter sizes has the side effect of making fixed-width type much less legible than its Latin-script counterpart (more on this, when we introduce typewriter fonts). The bottom half of Fig. 8 contains a few verses of a poem, set with fixed-width and variable-width Persian fonts. The difference in legibility, aesthetic quality, and compactness are quite evident. The Persian alphabet has four additional letters compared with the Arabic alphabet, but these added letters have shapes that are very similar to existing Arabic letters, so font design and considerations of legibility and aesthetic quality aren't any different between the two languages, and we can pretty much ignore in our discussion specific mention of Arabic. One way to impart the difficulties discussed so far to non-Persian-speakers is to liken them to those encountered in printing Latin calligraphic script, which has both the problem of connectivity between adjacent letters and wide variations in letter sizes, particularly between initial-capital and regular letters.

The sample of newspaper front pages depicted in Fig. 9 provides a sense of modern Persian print scripts of various designs and sizes, collected in July 2017, upon the passing of Iranian-born mathematician and Fields-Medal winner Maryam Mirzakhani. Even though headline fonts are significantly larger than normal-text fonts in print media, legibility issues are pretty much the same, given that headlines must be readable from a much larger distance.

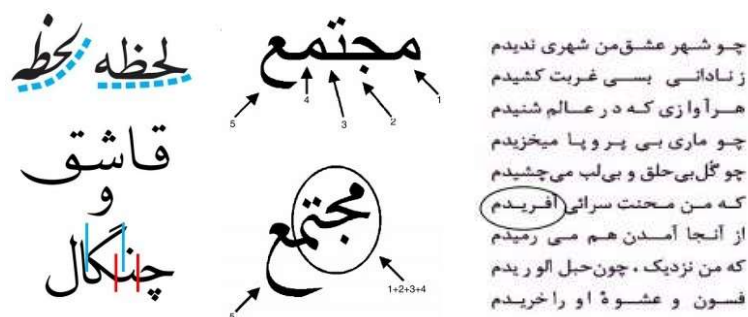


Fig. 7. Adaptation of the Persian script to movable-type print technology.

The Latin alphabet

abcdefghijklmnopqrstuvwxyz + / # % \$ & (
 ABCDEFGHIJKLMNOPQRSTUVWXYZ

The quick brown fox jumped over the fence. Arial 24
 The quick brown fox jumped over the fence. Times N R 24
 The quick brown fox jumped over the fence. Courier 20

The Persian alphabet

العربية فارسی
 Fixed-width characters
 کار پاکان را قیاس از خود مگیر
 آن یکی شیر است اندر بادیه
 و آن دگر شیر است کادم میخورد
 Variable-width characters
 کار پاکان را قیاس از خود مگیر
 آن یکی شیر است اندر بادیه
 و آن دگر شیر است کادم میخورد

Fig. 8. Comparing the Latin script and fonts with Persian/Arabic script.



Fig. 9. Modern Persian print fonts used in today's newspapers.

4. Persian Typewriters

Typewriters entered Iran around 120 years ago, but again it took a while for them to be used widely. Another set of adaptations were needed to allow Persian typewriting. In print type, each Persian letter has up to four shapes, corresponding to it appearing alone, at the beginning of a word, in the middle of a word, or at the end. The solo, initial, middle, and end forms of some letters, such as “ein,” shown at the top-left of Fig. 10 are quite different, requiring 4 typewriter keys (or two keys, without or with shift) to cover them. With 32 letters and 4 shapes, even ignoring digits and special symbols, 64 keys would be needed, along with the shift key. Fortunately, however, for most Persian letters, the initial and middle forms and the solo and end forms are sufficiently similar to allow their combining, with no great harm to the resulting script’s readability or aesthetic quality.

For example, the letter “meem” will not be damaged too much if it is given two forms rather than four. Ditto for the letters “beh” (bottom of Fig. 10) and the letter “heh” (left edge of Fig. 10), with slightly more compromise. This combining process allowed Persian typewriter keyboards to have only two forms for each letter, the only exceptions being the letters “ein” and “ghein.”

A breakthrough in typewriter technology occurred when smart electronic word-processors and associated software were introduced. Such smart input devices needed just one form of each letter on the keyboard, with the actual shape of the printed letter determined based on context (the previous and next letter). For example, if you want to type the word “kamtar,” composed of the four letters “keh,” “meem,” “teh,” and “reh” (recall that short vowels such as “a” are not printed, except perhaps as a diacritical mark, as shown near the top-right of Fig. 10), you begin by entering “keh.” At this point, the typewriter uses the solo shape, given that it has no information on what will come next. A display unit actually displays the solo “keh” and changes it later, if needed. A typewriter actually delays the printing of each symbol by one position, until it can determine the correct shape. Once you type the second letter “meem,” which connects to “keh,” the initial shape of “keh” is established, but the shape of “meem” depends on what comes next. This process continues until the end of a word (hitting of the space key) or line (return key) is encountered.

Mechanical Persian typewriters had keyboards such as the one shown in the left panel of Fig. 11. It was recognized early on that the arrangement of keys on the keyboard should be standardized, so that typing skills, once established, do not have to be re-learned. So, a standard layout was agreed upon and published [ISIR76]. Over the years, the print mechanism, which initially consisted of hammers hitting the ribbon and paper upon depression of a key, underwent improvements.

One of the highest-quality fonts for typewriters was offered by IBM in its Selectric line of word-processors, which used a golf-ball print mechanism (right panel of Fig. 11). When a key was pressed, the golf-ball mechanism would tilt and rotate to align the selected symbol with the print position before striking the golf-ball against the ribbon and paper. The golf-ball was easily removable for replacement with another golf-ball bearing a different font (italic, symbol, etc.), making it easy to typewrite technical manuscripts involving multiple typefaces and equations. Even multiple languages could be easily incorporated in the same document. In fact, I used such a typewriter to produce my first textbook, *Computer Appreciation* [Parh84a], sample pages of which appear in Fig. 12.

In addition to its advanced and high-quality print mechanism, IBM Selectric had a storage buffer for storing the content of an entire typed line. When the backspace key was pressed, the previous printed letter was read out from the buffer and was printed again, but with a white erasing ribbon instead of the black printing one. This caused the letter to be erased, providing a simple mechanism for correcting errors, without the mess of white-out liquid and its smudging problem.

There was a push in the mid-1970s to devise a unified standard layout for both typewriter and computer keyboards. An attempt in this direction, which was never approved as a standard or implemented, is depicted in the left panel of Fig. 13 [Parh84b]. Taking a cue from the “QWERTY” name for standard Latin keyboards, we called this keyboard layout “Zood-gozar,” based on the letters appearing in the middle section of the bottom row of keys. There was also a proposed standard information interchange code [IPBO80] [Parh84b], an extension of ISO standard code to the Persian language, which was also not adopted, as we hit the strikes and disruptions of the last couple of years before Iran’s Islamic Revolution, when government activities came to a stand-still.

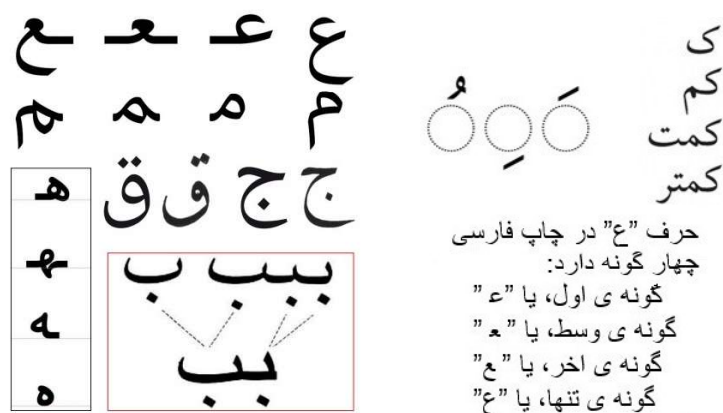


Fig. 10. Persian letters have four versions in print, two in typewriters, and one with smart text editors.



Fig. 11. An early Persian typewriter keyboard and IBM Selectric golf-ball print mechanism for Persian.

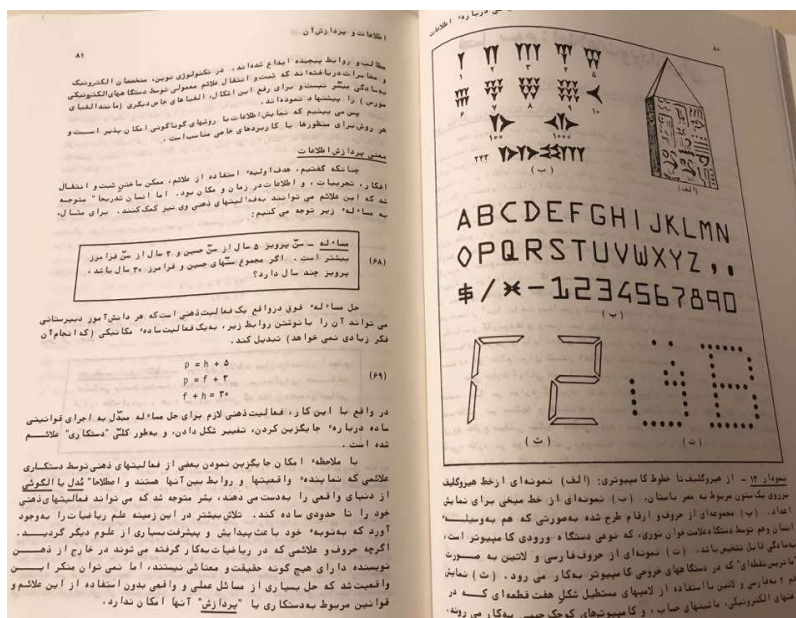


Fig. 12. IBM Selectric's Persian script, as used to typeset a book by the author.

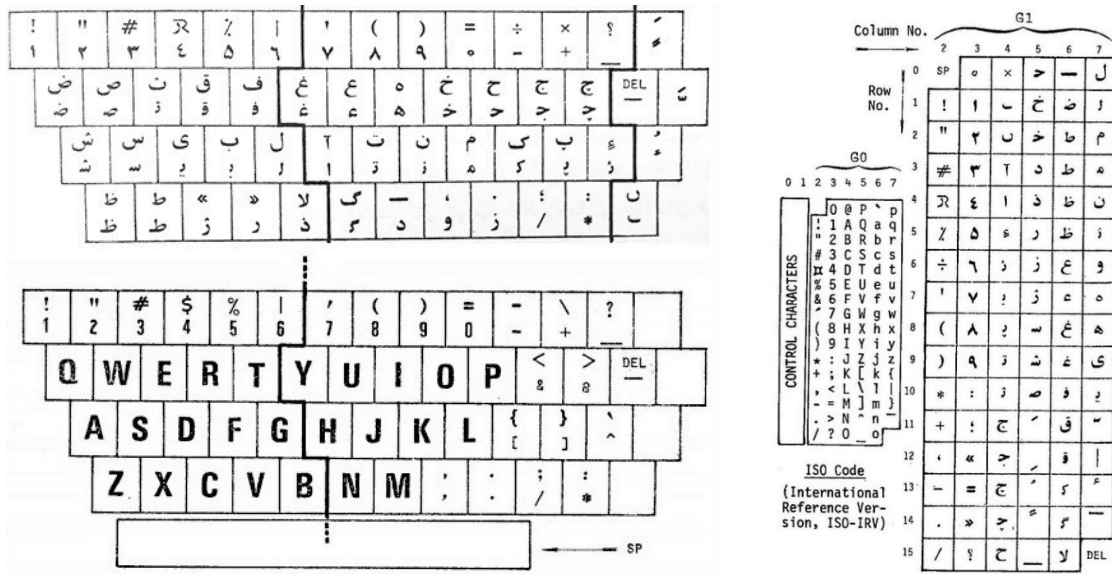


Fig. 13. Proposed standard Persian keyboard and information-interchange code..




Fig. 14. Drum printer mechanism, sample of the resulting Persian script, and an early line-printer.

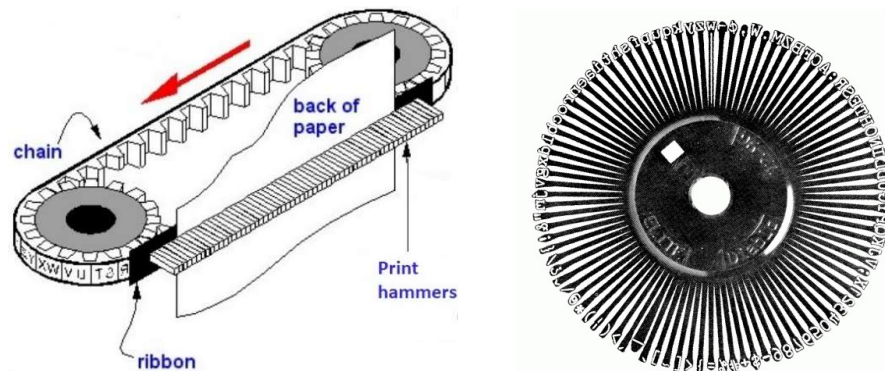


Fig. 15. Chain and daisy-wheel print mechanisms led to more compact printers.

5. Early Computer Printers and Displays

The earliest computer printers were known as line-printers, because they printed an entire line of text (typically on a wide paper, holding 132 columns) at once, advancing to the next line via moving of the continuous-feed fan-fold paper (right panel of Fig. 14). The print mechanism in the earliest line-printers was a drum (left panel of Fig. 14), with one ring of symbols for each of the print positions and each ring containing the entire printable character set. As the drum rotated at high speed, every letter eventually got aligned with the position of the line being printed. At that instant, a hammer would hit the paper and ribbon against the drum, creating the image of the desired symbol. In one drum rotation, taking a few milliseconds, one line would be printed. A drawback of drum printers for the Persian script was the gaps between adjacent symbols, necessitated by the placing of letters on the drum and also the spacing needed between hammers. The resulting fixed-width script appears near the center-top of Fig. 14. Slight errors in the timing of the hammer strikes would lead to the letters on the line to be mis-aligned, further deteriorating the script's legibility and aesthetic quality.

An innovation, that improved the print legibility and quality but was short-lived because of other problems it created, was switching the direction of lines and columns on a drum printer. With this variation, print lines would go along the perimeter of the drum rather than across the length of it. Consecutive Persian letters were now printed by the same hammer that could strike at any time, thus, at least in theory, removing the gap between symbols completely. However, errors in the timing of hammer strikes was still problematic, as was the need to print many times, even for pages that contained only 1 or 2 lines of text (the normal line-printer would skip the rest of the page once the few needed lines had been printed).

I conclude my discussion of early computer printers by introducing two other common print mechanisms as examples. Chain and daisy-wheel printers, whose print mechanisms are depicted in Fig. 15, were cheaper and less bulky than drum printers, but they suffered from similar script-quality problems. In both cases, hammers would hit the embossed letters, on chain links or on the petal-like arms of the daisy-wheel, to print the corresponding letter on paper.

Before discussing how computer printing technology advanced beyond the early drum, chain, and daisy-wheel print mechanisms, let us discuss the technologies for outputting information on display units. One of the earliest and simplest display units for numerical data was the 7-segment display, shown at the top of Fig. 16. Seven small lights, later made of LEDs, would be turned off and on to illuminate the shapes of decimal digits, which we need on calculators, clocks, thermometers, and the like. This kind of line-segment display units were later extended to letters and other symbols (bottom of Fig. 16), using, for example, 9 segments, 14 segments, or other designs, with the results being acceptable for upper-case Latin letters, but rather deficient when used to display lower-case letters and certain special symbols.

We tried our hands at creating numeric and alphanumeric line-segment displays for Persian, with the proposals shown in Fig. 17. We found out that a reasonably readable set of numerals can be created using 7 curved segments (top of Fig. 17) and that no fewer than 18 segments were needed to approach a barely readable set of letters (bottom of Fig. 17). Both proposals remained at the research and prototype stage, as the approach was overtaken by the more flexible dot-matrix display technology. A 7-by-5 matrix of lights (7 rows, 5 columns), for example, was adequate for producing a readable Latin script for use on airport information boards, stadium scoreboards, and other situations where alphanumeric information was to be displayed in public places. As we developed the capability of making smaller dots/lights and to place them closer to each other, the quality of the resulting script improved through the use of 9-by-7 and even larger dot matrices.

Early CRT display units either generated desired textual and graphic elements by drawing lines on the screen (the way we write by drawing straight and curved strokes) or were equipped with special character-generator hardware that would shape the electron beam hitting the phosphorus screen to form the shape of a desired character. Eventually, CRT designs settled on the vector scheme, where a buffer would hold the dark/light or 0/1 pattern to be displayed on a dot matrix and the data in the display buffer was used to turn off or on the electron beam, as it scanned the surface of the display, according to the 0s and 1s in the display buffer. Color displays essentially held three copies of the said mechanism, one for each of the primary colors. Eventually, CRT displays would give way to flat-screen displays of various kinds, but the dot-matrix display method is here to stay.



Fig. 16. Line-segment display devices for numerals, Latin alphabet, and special symbols.

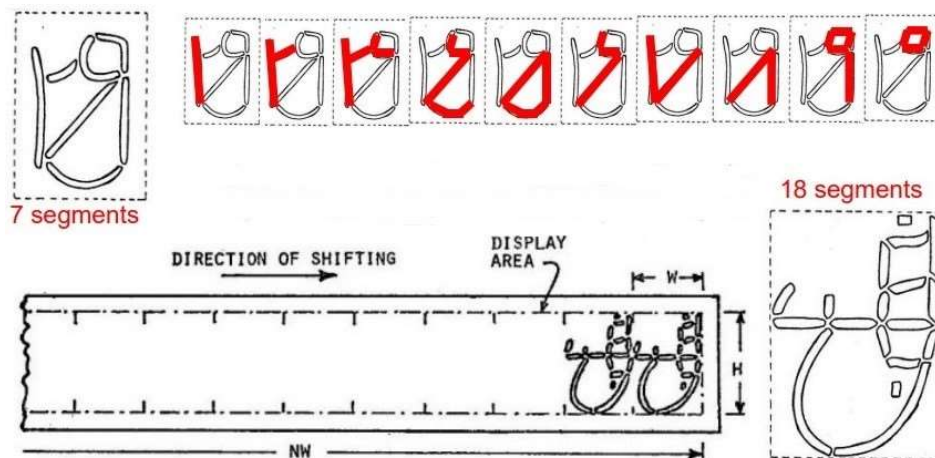


Fig. 17. Proposed line-segment display devices for Persian numerals and letters.

FLIGHT	DESTINATION	SCHED	BOARD	GATE	STATUS
QF 408	SYDNEY	7:00	6:40	23	
QF 610	BRISBANE	7:05	6:50	2	
QF 412	SYDNEY	7:30	7:10	24	
QF 414	SYDNEY	8:00	7:40	12	
DEPARTS FROM NON CITYFLYER GATE					
QF 612	BRISBANE	8:05	7:45	1	
QF 416	SYDNEY	8:30	8:10	24	
QF 418	SYDNEY	9:00	8:40	21	
QF 614	BRISBANE	9:05	8:45	1	
QF 420	SYDNEY	9:30	9:10	5	
DEPARTS FROM NON-CITYFLYER GATE					
QF 422	SYDNEY	10:00	9:40	24	



Fig. 18. Early dot-matrix displays were used at airports and on stadium scoreboards.

6. Progress in Computer Printing Technology

Several other printing technologies came and went, until improvements in dot-matrix printing made all other methods obsolete. The dot-matrix printing method has some elements in common with the square Kufic script, developed by artists to facilitate writing with tiles on the façade of mosques, palaces, and other buildings (left panel in Fig. 19). Just as a ceramic artist might have set all the tiles in one column before moving on to the next column, early dot-matrix printers had a column of 7 print needles that would be set to high and low, before being pushed forward to produce contact with a ribbon and paper, proceeding to the next column, either by moving the print mechanism or shifting the paper (right panel of Fig. 19). Eventually, more than 7 needles and more than one column of needles, aligned or skewed, would improve the print quality. The needles would eventually be replaced by tiny ink-jets that would shoot a drop of ink toward the paper to form a dot.

Dot-matrix printers offered extreme flexibility. Not only could we design fonts in a variety of ways, using our own preferences and trade-offs, we could also produce maps and geometric shapes, such as lines and circles (Fig. 20). As the dots became smaller and more of them could be placed in one square inch of space, print and image resolution improved, so much so that, today, we barely notice the tiny dots that form the letters and images on paper or on screen. Problems arising from the connectivity of Persian letters and their disparate sizes also disappeared. What was a hardware problem (designing print mechanisms and fitting them to the Persian script) is now a software notion: specifying the light and dark dots that form each Persian symbol.

We, therefore, got busy producing dot-matrix character sets for the Persian alphabet, as had been done previously for the Latin alphabet (Fig. 21). Once a dot-matrix alphabet had been created, we would copy the letters on small pieces of paper and would juxtapose them to write words and sentences (middle part of Fig. 21). If the quality wasn't to our satisfaction, we went back to the drawing board and made changes to letter designs. This process was eventually mechanized and we created the sample texts on a screen rather than by arranging pieces of paper. As dot-matrix size and resolution grew, more legible and aesthetically pleasing symbols could be created.

The impact of dot-matrix size on font quality is seen on the left half of Fig. 22, where the letter “R” is generated in two matrices: a 7-by-5 matrix and a 16-by-16 matrix with overlapped dots that leads to smoother curved edges. Interestingly, the legibility of a script is connected with the ease of recognizing it automatically by a computer program. In an early Persian character-recognition study using newspaper headlines (because, at the time, we did not have high-resolution scanners to use normal printed text), we solved the problem of how to decompose chunks of printed text into its constituent letters [Parh81c]. The result for the first text chunk in the word “Tehran” is depicted in the right panel of Fig. 22, where the first three letters have been separated before recognition. The results of this study not only showed the feasibility of automatic decomposition and recognition, but also helped us in the design of dot-matrix fonts with better legibility. Since our early work of the late 1970s, the field has advanced by leaps and bounds, so much so that a Google search for “Persian character recognition” produces more than 6 million hits.

Dot-matrix display and printing finally realized the goal of high-quality Persian output, while also allowing certain exotic fonts [Parh78a]. As seen in Figs. 23 and 24, the change in the Persian script quality was already significant from early line-printers and rudimentary dot-matrix printers of the 1970s to second-generation dot-matrix printers of the 1980s. This improvement trend has continued unabated in the three decades since, to the extent that today, we barely notice the collection of dots that forms each displayed or printed symbol. I will present the latest in Persian output-script production in Section 7, but the reader can preview the modern script quality in Figs. 29 and 30.

The production of high-quality print and display output has also quelled sporadic attempts at reforming the Persian script or changing the alphabet [Borj11], a la what was done in Turkey, to save the society from “backwardness.” I consider this outcome quite fortunate, as the right approach to technology transplantation is adapting technology to cultural factors, not the other way around. The extreme of insisting on language purity, which leads to resistance to new concepts and inflexibility in making trade-offs for the sake of producing legible output, for example, is just as bad as wholesale abandonment of centuries of traditions for short-term technological gain.

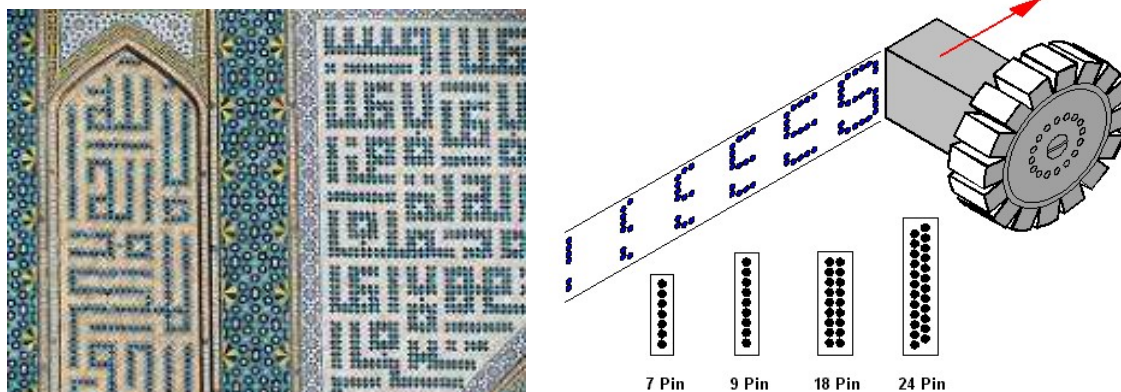


Fig. 19. The square Kufic script for writing with tiles, and early dot-matrix printing technology.

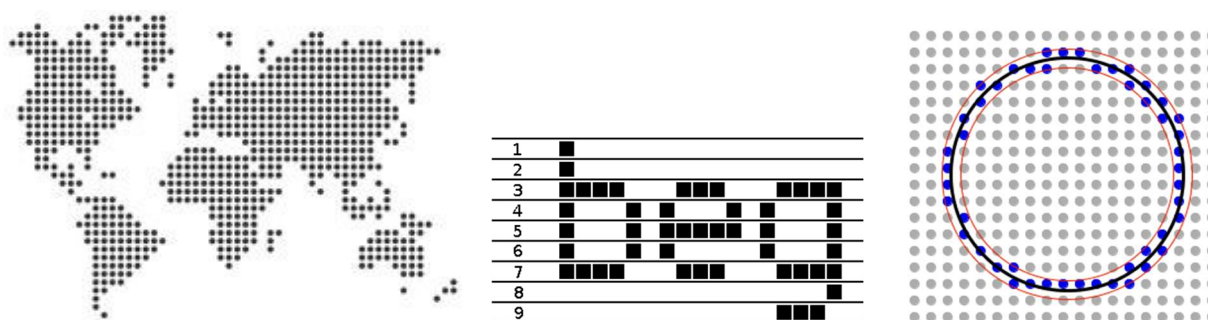


Fig. 20. Dot-matrix style of printing and display is quite flexible, but it isn't a cure-all.

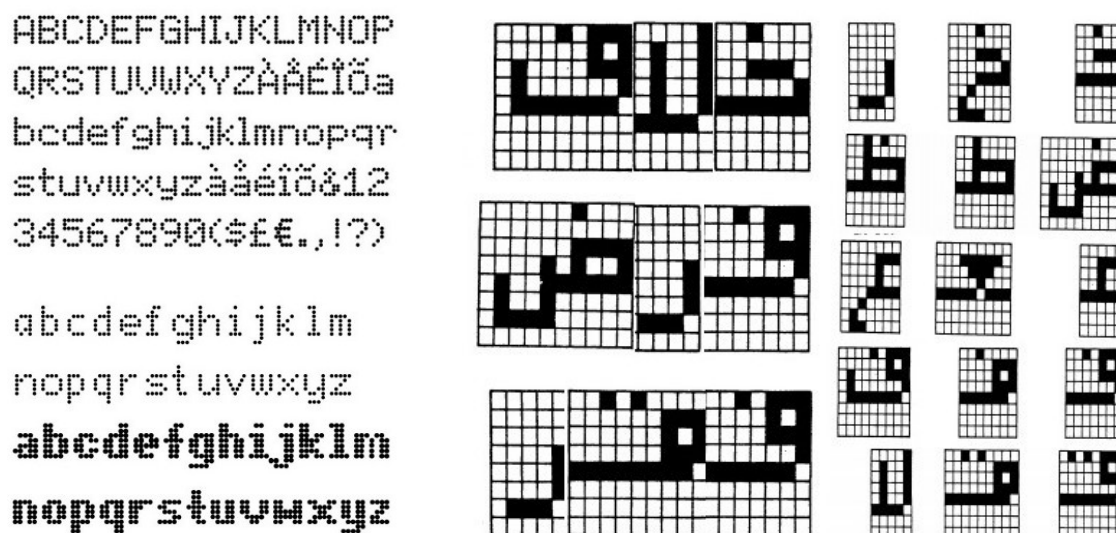


Fig. 21. Dot-matrix designs for some Persian letters, compared with those for the Latin alphabet.

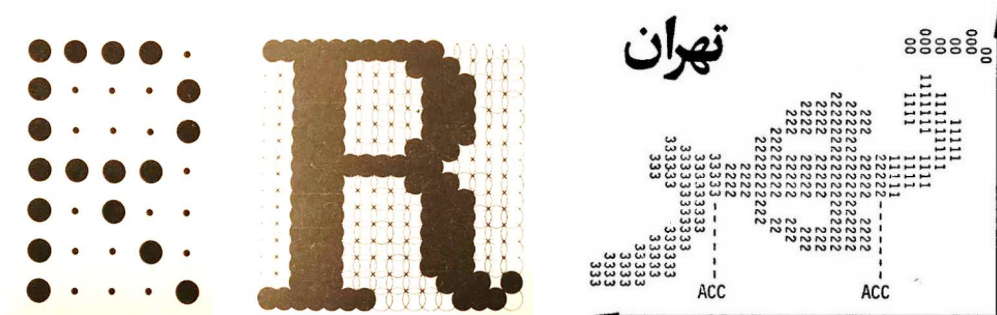


Fig. 22. Small 7-by-5 vs. large 16-by-16 (overlapped) dot matrices and example of connected letters.



Fig. 23. Samples of early computer printouts (1970s).

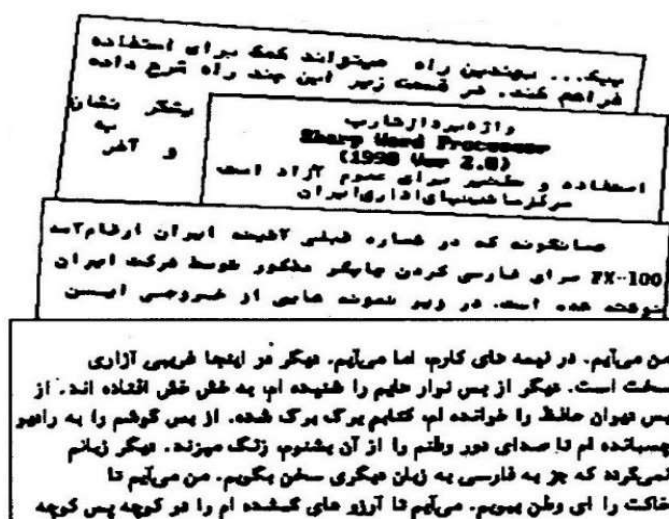


Fig. 24. Samples of later computer printouts (1980s).

7. Persian/Arabic Scripts on Today's Computers

Much progress has been made in improving the script quality in Persian printers and displays, all of which now use the dot-matrix scheme to form the symbols. To cite an example, the Persian font quality on Facebook shows continued improvement. The improvement does not always come on all platforms at once. In 2017, I made several suggestions for font quality improvement to Facebook and, in late 2017, a new font was rolled out on the mobile Facebook platform, but unfortunately, the desktop platform font (left panel in Fig. 25) is still illegible, unless the image is enlarged significantly. The right panel in Fig. 25 shows the font used on an on-line news site, which is significantly better. The lesson here is that users of Persian script should complain loudly at every opportunity in order to get better print and display quality. Now that we have huge data sets of Persian texts, such as posts on social media, there is no excuse not to use the data set to find out about user tastes and preferences, perhaps in a crowd-sourcing format.

As mentioned earlier, legibility problems of Persian fonts lead directly to difficulties in automatic script recognition. The newspaper headline in Fig. 26, from a three-decades-old study [Parh81c], depicts some of the problems in both domains: discerning connection points (a), letters that are very similar in shape (b_1 , b_2), wide variations in letter widths (c_1 , c_2), horizontal overlap between letters on a line (d), and vertical overlap between lines (e), all of which make segmentation of text into its constituent letters difficult. A byproduct of the study was the conclusion that for accurate printed-text recognition, a “pen-width” of 4 dots (left panel in Fig. 27) seems to be necessary. This is also a good guideline for highly readable printed and displayed Persian script. With a pen-width of 4 or more, letter features, such as curvatures and holes, become easily noticeable (Fig. 27).

Now that we have an idea about the desirable dot-matrix size to accommodate a pen-width of 4 for high-quality fonts, it is also instructive to probe the lower boundary on matrix size to generate a script of adequate quality for certain mobile and embedded applications, where cost factors may inhibit the use of large dot-matrices [Parh95]. The simulation results for fonts in 7-by-5, 7-by-9/2, and 9-by-9/2 are shown in Fig. 28 (middle panel, top to bottom). A matrix dimension $m/2$ implies the presence of m rows/columns of dots in skewed format, so that the physical dimension of the matrix is roughly $m/2$, despite the fact that there are m elements. This kind of skewed arrangement of dots helps with generating fonts of higher quality, when the letters have curved or slanted strokes.



Behrooz Parhami

1 hr · 🌐 ▼

دانش و فن آوری

حذف نرم افزارهای ایرانی از فروشگاه مجازی گوگل

کار یاکان را قیاس از خود مگیر
گر چه باشد در نوشتن شیر، شیر
آن یکی شیر است اندر بادیه
آن دگر شیر است اندر بادیه
آن یکی شیر است کادم می خورد
و آن دگر شیر است کادم می خورد

ویدئو: سلول های بنیادی
محدودیت اینترنت: آزادی های آنلاین پس از انتخابات ۹۶
چهل سال پس از وویجرها: دو نامه از یک تمدن
محمدجواد آذری جهرمی و همزیستی شغلی با فیلترینگ و شنود

گوگل نیز پس از اپل
نرم افزارهای ایرانی را از
فروشگاه مجازی خود حذف
کرد. دلیل آن تبعیت از قانون
تحریم های ۲۲ ساله آمریکا
علیه ایران است.



Fig. 25. Sample dot-matrix display typefaces from Facebook (left) and an on-line site.

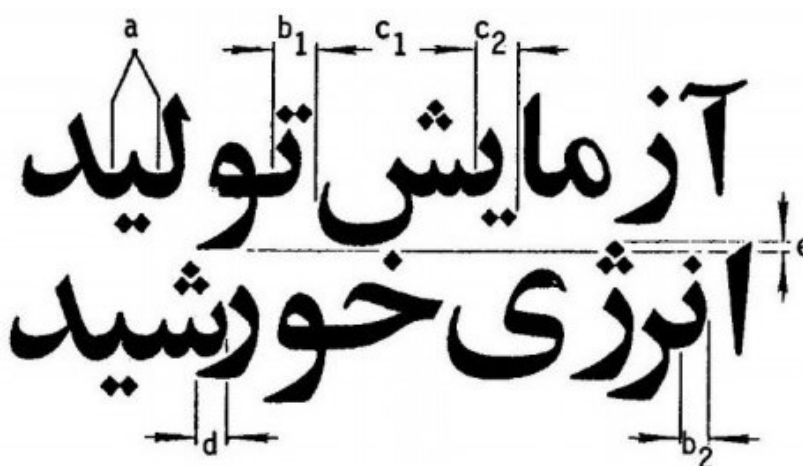


Fig. 26. Some of the challenges in the automatic recognition of Persian printed texts.

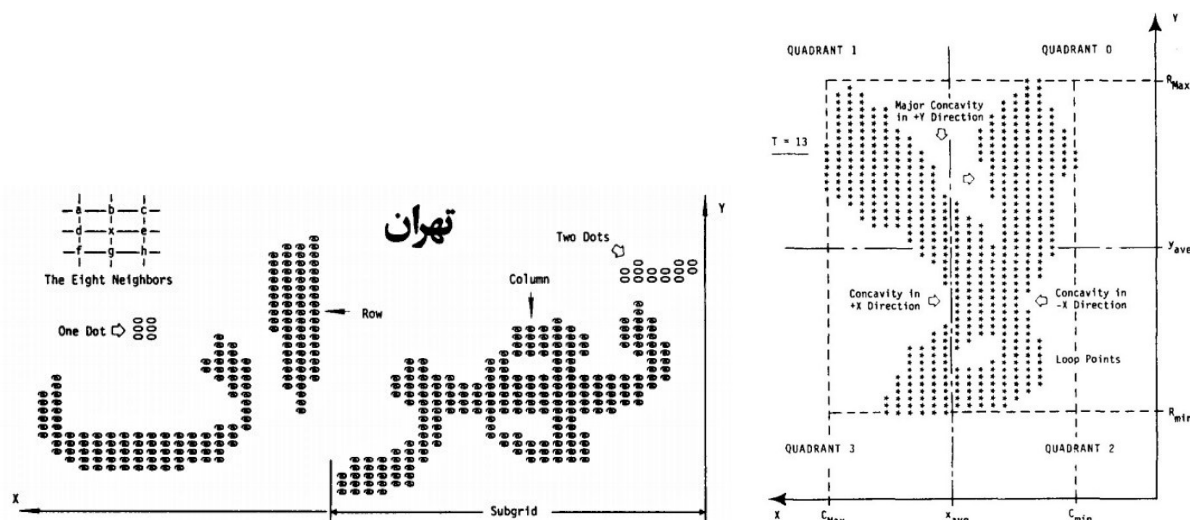


Fig. 27. Techniques and results from the automatic recognition of printed Persian texts.

Given that certain tall Persian letters are particularly challenging for displaying or printing with small matrices, a separate simulation was conducted with sample texts featuring many tall letters. Partial results are shown in the top-right of Fig. 28 for the same three dot-matrix dimensions, along with the sample text represented with variable-width dot-matrices. The study's conclusion was that a 9-by-9/2 dot-matrix is the absolute minimum for minimally legible Persian script.

In modern computer applications, scripts need to be displayed on devices of varying sizes, from large desktop screens to tiny displays on mobile devices. So, each font, is designed in various point-sizes (top of Fig. 29). Legibility considerations dictate that a 20-point font, say, not be a linear reduction by a factor of 2 of a 40-point font. In some cases, such as viewing a map on a small screen, the fonts and other image elements are actually shrunk linearly to fit the available space. It is thus important to conduct experiments on font legibility and aesthetic quality when such shrinkage takes place. The bottom half of Fig. 29 shows the impact of font-size adjustment/reduction as well as image resizing/shrinkage on the resulting script.

I conclude this section with actual modern Persian fonts, as implemented by Microsoft. The font menu of Microsoft Word, from which the samples in Fig. 30 have been drawn, offers multiple font options for the Persian script. The first four options (Arial, Cambria, Helvetica, Times New Roman) generate the same script on my version of Microsoft Word. Perhaps they will be replaced with distinct fonts in future. The other options (Calibri, Courier, Dubai, Dubai Light, Tahoma, Traditional Arabic) exhibit different levels of readability and aesthetic quality. The Calibri and Dubai fonts perhaps strike the best balance between readability and aesthetic quality, as evident from the enlarged samples at the left edge of Fig. 30. The Courier font suffers from typical problems of fixed-width fonts for Persian, but if having a constant symbol width is a requirement, it is more than bearable.

Beyond run-of-the-mill printing and display, decorative and artistic writing, using a variety of handwriting styles can be and has been mechanized. Some examples are provided by the on-line site NastaliqOnline.ir, which features an input box for entering a text and a menu to choose the style of writing. Then, upon issuing the “write” command, the site's software algorithm creates the resulting script and presents it in jpg format for copying and pasting into other documents. There are still problems with the script quality, but the site's designers are making improvements on a regular basis. Examples of the resulting scripts are provided in Figs. 31 and 32.



Fig. 28. Some results of a study to determine the minimum dot-matrix size for Persian script..



Fig. 29. Persian fonts of various point-sizes and the effect of scaling on script readability and quality.

8. Conclusion and Remaining Challenges

After facing five decades of challenges and responding with solutions and improvements, we are still nowhere near the end of the line. Display quality and resolution is continually improving, which is a big help. High-definition (HD) and ultra-high-definition (UHD) displays offer 1920-by-1080 and 3840-by-2160 matrices, respectively, with the latter doubling the resolution in each dimension and, thus, quadrupling the areal resolution. In printers too, even the cheapest models now offer a resolution of 600 dots per inch (DPI) at a minimum. So, hardware capabilities are no longer our main problems. The problems lie in font design and text-processing algorithms for information display and printing. There is much accumulated know-how about fonts and algorithms, which must be put to use in a systematic way, with empirical evaluation and extensive feedback.

What remains to be done includes a better understanding of the trade-offs between legibility and aesthetic quality and in-depth studies of formatting issues, especially when bilingual and multilingual output is involved. The first problem can be investigated through the use of crowd-sourcing. The vast data sets of available Persian script (e.g., from Facebook and Twitter posts) can be used for occasional polls, putting two different scripts side-by-side and asking users to rate them. In the area of formatting, the opposite directions of writing in Persian (right to left) and Latin-based (left to right) scripts create problems. If I were to insert a Persian word here, say, **فارسی**, the formatting might be messed up if the Persian word appears close to a line-break so that it has to be moved to the next line. Punctuation marks, particularly parentheses and brackets, also create similar problems, sometimes moving from their proper places to different positions in the text.

The expanding user base demanding the capability to print and display the Persian script is helping with improvements. However, it is important that users stay active and engaged, providing feedback to the fullest possible extent in order to motivate investment in improving the hardware technologies and software algorithms. An integrated, easy-to-use bilingual keyboard and improved optical character recognition would be important first steps in solving the remaining input problems. Alongside improving the existing facilities for Persian input and output, we must keep an eye on what's coming in the near future: voice input and output. Much work has been done in English voice input, but we are at the beginning of the process of applying such techniques to Persian [Same08]. Ditto for Persian voice output, from a textual document [Boka10].

Display and printing technologies aside, studies are needed to gauge the Persian script itself, in order to develop a better understanding of legibility and aesthetic quality considerations. I will cite just two examples of corresponding studies for English hand-printed and printed scripts, as examples of how the concepts and results can be extended to Persian.

Appendix A in a book [Bail83] that I have used in a graduate-level course on fault-tolerant computing, where the impact of errors of various kinds on the correct functioning of computer systems is studied, contains an interesting table which provides empirical error rates, when manually-written texts and filled-out forms were keyed-in by operators of early computer input devices. According to that table, the least error-prone hand-printed English alphanumeric symbols, with error rates of less than 1%, are these 8 symbols: W, M, 3, 7, A, 9, E, C. Near the high end of the error spectrum, with error rates of around 5%, we find these 6 symbols: N, 0, 5, J, V, G. Particularly error-prone are the letters Z (13%) and I (25%). Error-proneness is intimately related to legibility, but I am unaware of any study of error-proneness for handwritten or printed Persian scripts.

In studies of reader comprehension of texts, the notion of readability [Zama12] has been widely studied. Readability, which is different from legibility, is affected by word usage (short vs. long, common vs. rare, etc.), sentence structures (short/simple vs. long/complex), and the physical parameters of the text in terms of font, size, color, and the like. Legibility is a more focused notion that pertains only to the letter/symbol shapes, their spacing, and combined appearance in words, the so-called word “shapes.” Again, I don’t know of any studies in this area specific to Persian, but efforts to characterize the legibility of Arabic script have begun [Chah12].

References

- [Bail83] R. W. Bailey, *Human Errors in Computer Systems*, Prentice-Hall, 1983.
- [Boka10] M. H. Bokaei, H. Sameti, H. Eghbal-zadeh, B. BabaAli, K. H. Hosseinzadeh, M. Bahrani, H. Veisi, and A. Sanian, "Niusha, the First Persian Speech-Enabled IVR Platform," *Proc. 5th IEEE Int'l Symp. Telecommunications*, 2010, pp. 591-595.
- [Borj11] M. Borjian and H. Borjian, "Plights of Persian in the Modernization Era," *Handbook of Language and Ethnic Identity: The Success-Failure Continuum in Language and Ethnic Identity Efforts*, Vol. 2, pp. 254-267, 2011.
- [Chah12] N. Chahine, "Reading Arabic: Legibility Studies for the Arabic Script," Doctoral Dissertation, Faculty of the Humanities, Leiden University, 2012.
- [Dunn79] P. D. Dunn, *Appropriate Technology. Technology with a Human Face*, Schocken, 1979.
- [IPBO80] Iranian Plan and Budget Organization, *Final Proposal for the Iranian National Standard Information Code (INSIC)*, Persian and English versions, 1980.
- [ISIR76] Institute of Standards and Industrial Research of Iran, *Character Arrangement on Keyboards of Persian Typewriters* (in Persian), ISIRI 820, 1976.
- [Laza08] G. Lazard, "The Rise of the New Persian Language," *The Cambridge History of Iran*, Vol. 4 (Period from the Arab Invasion to the Saljuqs), 2008, pp. 566-594.
- [LoC14] Library of Congress, "A Thousand Years of the Persian Book," Exhibition and Lecture Series, 2014.
- [Mora13] F. Moradi, L. Nanquette, M. Hosseinipour, A. Amiri, D. Rakhimov, and B. Biniiaz, *Publishing in Persian Language: In Iran, Afghanistan, Tajikistan, Uzbekistan, Europe and United States*, Int'l Alliance of Independent Publishers, 2013.
- [Parh77a] B. Parhami and F. Mavaddat, "Computers and the Farsi Language: A Survey of Problem Areas," *Information Processing 77* (Proc. IFIP World Congress), North Holland, 1977, pp. 673-676.
- [Parh77b] F. Mavaddat and B. Parhami, "Informatics in Iran: Problems and Prospects," *Proc. Int'l Conf. Computer Applications in Developing Countries*, Bangkok, Thailand, August 1977, pp. 121-133.
- [Parh78a] B. Parhami, "Optically Weighted Dot-Matrix Farsi and Arabic Numerals," *Information Technology 78* (Proc. 3rd Jerusalem Conf.), 1978, North-Holland, pp. 207-210.
- [Parh78b] B. Parhami, "On the Use of Farsi and Arabic Languages in Computer-Based Information Systems," *Proc. Symp. Linguistic Implications of Computer-Based Information Systems*, New Delhi, India, November 1978.

- [Parh78c] B. Prahmi, "Impact of Farsi Language on Computing in Iran," *Mideast Computer*, Vol. 1, pp. 6-7, September 1978.
- [Parh81a] B. Parhami, "Language-Dependent Considerations for Computer Applications in Farsi and Arabic Speaking Countries," *System Approach for Development* (Proc. IFAC Conf.), North-Holland, 1981, pp. 507-513.
- [Parh81b] B. Parhami and V. Daie, *Glossary of Computers and Informatics: English/Persian*, Informatics Society of Iran, Tehran, May 1981. Reprinted many times up to 1991; revised and reissued by an ISI committee in 1994.
- [Parh81c] B. Parhami and M. Taraghi, "Automatic Recognition of Printed Farsi Texts," *Pattern Recognition*, Vol. 14, Nos. 1-6, pp. 395-403, 1981.
- [Parh84a] B. Parhami, *Computer Appreciation (Aashnaaee baa Computer*, in Persian), Tehran, 1984. Reprinted in 1985, 1987, and yearly thereafter. As of 2001, was still being used as a textbook.
- [Parh84b] B. Parhami, "Standard Farsi Information Interchange Code and Keyboard Layout: A Unified Proposal," *J. Institution of Electrical and Telecommunications Engineers*, Vol. 30, No. 6, pp. 179-183, 1984.
- [Parh85] B. Parhami, "University Education in Computer Science and Technology: The New Iranian Plan," *Proc. IFIP 4th World Conf. on Computers in Education*, Norfolk, VA, August 1985, pp. 923-930.
- [Parh86] B. Parhami, "Computer Science and Engineering Education in a Developing Country: The Case of Iran," *Education and Computing*, Vol. 2, No. 4, pp. 231-242, 1986.
- [Parh95] B. Parhami, "On Lower Bounds for the Dimensions of Dot-Matrix Characters to Represent Farsi and Arabic Scripts," *Proc. 1st Annual CSI Computer Conf.*, Tehran, Iran, December 1995, pp. 125-130.
- [Same08] H. Sameti, H. Veisi, M. Bahrani, B. Babaali, and K. Hosseinzadeh, "Nevisa, a Persian Continuous Speech Recognition System," *Proc. 13th Int'l CSI Computer Conf.*, Springer, 2008, pp. 485-492.
- [Zama12] M. Zamanian and P. Heydari, "Readability of Texts: State of the Art," *Theory & Practice in Language Studies*, Vol. 2, No. 1, 2012.

About the Author

Behrooz Parhami (PhD, UCLA 1973) is Professor of Electrical and Computer Engineering, and former Associate Dean for Academic Personnel, College of Engineering, at University of California, Santa Barbara, where he teaches and does research in computer hardware architecture. A Life Fellow of IEEE, a Fellow of IET, and British Computer Society, and recipient of several other awards (including a most-cited paper award from *J. Parallel & Distributed Computing*), he has written six textbooks and more than 300 peer-reviewed technical papers. Most relevant to this paper, is his 12 years of experience at Iran's Sharif (formerly Arya-Mehr) University of Technology, spanning 1974 to 1986, where he studied the Persian language in connection with computers and their applications, participated in technology transfer and standardization projects, helped establish the Informatics Society of Iran (and its technical journal, *Computer Report*), and was a leader in developing computer engineering educational programs and associated curricula. Professionally, he serves on journal editorial boards and conference program committees and is also active in technical consulting.

