



The journal of interface, experience, and
communication design research

Visible Language

ISSN 0022-2224

Published continuously since 1967

Volume

60

Number

1

April 2026

inside

front

cover



The journal of interface, experience, and
communication design research

Visible Language

ISSN 0022-2224

Published continuously since 1967

Volume **60** Number **1**

April 2026

The **Visible Language Consortium** is a collaboration among three institutions — through the School of Design at University of Leeds (UK), the Ullman School of Design at University of Cincinnati (USA), and the College of Design at North Carolina State University (USA) — from which the **editorial board** of *Visible Language* is composed:



Editor-in-Chief
Jeanne-Louise Moys

Associate Editor
Arjun Khara

Assistant Editor
Matthew Baxter

Editorial Manager
Rafiq Elmansy



Editor
Mike Zender

Associate Editor
Muhammad Rahman

Assistant Editor
D.J. Trischler



Editor
Matthew Peterson

Associate Editor
Deborah Littlejohn

Assistant Editor
Helen Armstrong

Advisory board: Naomi Baron Michael Bierut Charles Bigelow
Matthew Carter Keith Crutcher Mary Dyson Jorge Frascara
Ken Friedman Michael Golec Judith Gregory Kevin Larson
Aaron Marcus Tom Ockerse Sharon Poggenpohl Michael Renner
Stan Ruecker Katie Salen Peter Storkerson Karel van der Waarde
Mike Zender

iv–vi

The Editors Introduce the April 2026 Issue

Jeanne-Louise Moys, Mike Zender, and
Matthew Peterson

1–23

**The Terminological Development of
Graphic Design: Between Office Art and
Social Purpose**

David Preston

24–57

**Automation and Artificial Intelligence in the
Type Design Process: Insights from an
Industry Survey**

Alice Savoie, Kai Bernau, Wayne Daly, Raphaela
Haefliger, and Sebastian Baez-Lugo

58–83

**Evaluating Interactive Highlighting
Techniques in Digital Reading: An Empirical
Study of Hover-Based Line, Sentence, and
Paragraph Highlighting**

Szabolcs Vatány, Thi Huyen Nguyen, Anikó Illés, and
Ann Bessemans

84–89

Why Meta-Font Struck a Nerve

Mike Zender

90–119

**Metafont, Metamathematics, and
Metaphysics: Comments on Donald Knuth’s
Article “The Concept of a Meta-Font”**

Douglas R. Hofstadter

120–123

**A Profession Provoked: How Meta-Font
Struck a Nerve**

Deborah Littlejohn

124–125

Call for Papers: Typographic Landscapes

Volume 60 Number 1

The Editors Introduce the April 2026 Issue

Jeanne-Louise Moys ^a , Mike Zender ^b, and Matthew Peterson ^c 

^a School of Design, University of Leeds, UK; ^b College of Design, Architecture, Art, and Planning, University of Cincinnati, Cincinnati, OH, USA; ^c College of Design, North Carolina State University, Raleigh, NC, USA

Notes from Editor-in-Chief Jeanne-Louise Moys. Welcome to the first of three issues in volume 60. We would first like to pay tribute to all the editors, authors, reviewers, copyeditors, and designers who have contributed to *Visible Language* over the years and enabled us to meet this milestone of 60 years of publishing.

As a milestone year, it seems apt to open with David Preston's (2026) article, "The Terminological Development of Graphic Design." You, our readers, will have different views of the range of practices and sensibilities the term *graphic design* encompasses and whether it is still apt to describe how the profession and its artifacts have evolved. Many researchers, practitioners, and educators use alternative terms, such as *communication design* or *user-centered design*, instead of *graphic design*. Nevertheless, the term often seems to be more readily familiar to the general public and new undergraduates, even though their associations may lie more with advertising and branding applications than the range of visual practices and contexts designers and researchers engage with. Revisiting the roots of the term and how its usage has evolved is important to inform our current understanding of what we do as researchers, educators, and practitioners.

Notes from Editor Mike Zender. Alice Savoie, Kai Bernau, Wayne Daly, Raphaela Haefliger, and Sebastian Baez-Lugo's (2026) article, "Automation and Artificial Intelligence in the Type Design Process," reports on three research questions: how the type design industry is being impacted by AI, what risks and opportunities experts perceive,

@: [ISSUE](#) > [ARTICLE](#) >

Cite this article:

Moys, J.-L., Zender, M., & Peterson, M. (2026). The editors introduce the April 2026 issue. *Visible Language*, 60(1), iv–vi.

First published online April 26, 2026.

© 2026 Visible Language — this article is **open access**, published under the CC BY-NC-ND 4.0 license.

<https://visible-language.org/journal/>

Visible Language Consortium:

University of Leeds (UK)

University of Cincinnati (USA)

North Carolina State University (USA)

and how designers and computer scientists can collaboratively address ethics and quality (p. 26). In the body of the article, the authors add a fourth question: “what parts of their process” are designers “willing to delegate to a machine, and which parts are considered as meaningful, enjoyable, or highly dependent on human intervention?” (p. 33). They studied these questions through an online interview asking about *current practice and expectations* and *propensity to automate specific tasks*. The results they describe are interesting and offer a glimpse into the possibilities for AI in design beyond just typeface design.

Questioning by Mike Zender. As interesting as the author’s questions are, relevant questions not directly considered are equally interesting: to what extent are concerns about AI in type design the result of negativity bias; do we know enough about visual letterform concepts to create AI that can do more than mimic existing typefaces; what parameters define an excellent typeface design? That last question about quality is addressed by Savoie et al. (2026) obliquely, writing that some AI generated typefaces are “...rudimentary in their design, often encompassing only capital letters, and displaying obvious flaws” (p. 29). What precisely constitutes “rudimentary” and what the “flaws” are is not discussed, but these seem to be highly relevant questions for AI in type design or any kind of design, as are the metrics for making such judgments. This gets to a key issue about visual design: what exactly constitutes “good quality.” How will we know when an AI assisted or produced design is inadequate?

Through surveys and interviews, the authors expose the interesting quandary of machine learning from copyrighted material, as if human typeface designers generate their typefaces without having been influenced or “inspired” by existing letterforms, copyrighted or not. It is easy to fault a dumb machine, unable to defend itself, for doing what you have always done.

Intellectual property and creative control are two additional interesting issues explored in the article, along with the potential for AI to “function less as an autonomous designer and more as a brainstorming companion” (p. 45). This would parallel military applications of AI in such products as the “loyal wingman” of multiple unpiloted autonomous aircraft as companions for a single piloted warfighter.

Finally, the authors explore the idea that use of AI to assist in technical or repetitive tasks such as optimizing and kerning adoption also risks affecting “the very routines through which expertise is acquired” (p. 49). They quote one respondent saying,

...only doing something repeatedly makes you good at it and allows you to create novel and truly innovative things. Without that learning, it is difficult to imagine something amazing being developed. But if AI does all the generic work, how can anybody learn the skills that make a good designer? (p. 49)

Might machine learning for design make designers dumber? Will efficiency cause us to learn less? These issues and more are exposed and explored, in at least a provisional way, in this fascinating and informative article, making it well worth reading.

Notes from Editor Matthew Peterson. Szabolcs Vatány, Thi Huyen Nguyen, Anikó Illés, and Ann Bessemans (2026) address rather subtle differentiations of hover states for text highlighting in their article, “Evaluating Interactive Highlighting Techniques in Digital Reading.” Reading research has been an emphasis of *Visible Language* since its inception, and Vatány and colleagues continue that tradition. Text display technology changes rapidly and continuously, and so do readers, so we need scholars such as these to study reading, and to keep studying reading.

In addition to the above three new peer reviewed research articles, this first issue of Volume 60 includes a reprint of a 1982 article by Douglas Hofstadter, itself in reaction to another 1982 article by Donald Knuth, both in *Visible Language*'s 16th volume. I need not introduce it, as my colleague (and *Visible Language* editor) Mike Zender does so in an editorial preceding the reprint. I will simply note that it all centers on Knuth's Meta-Font, which raised issues about artificial intelligence long before the present moment. The reprint is followed by my colleague (and associate editor) Deborah Littlejohn's summary of letters published in Issue 60.4, which were solicited along with Hofstadter's original article, in reaction to Knuth's Meta-Font. I am not certain how many nested levels of *meta* this makes, but it is certainly a few.

References

- Hofstadter, D. R. (1982/2026). Metafont, metamathematics, and metaphysics: Comments on Donald Knuth's article “The Concept of a Meta-Font” [Reprint]. *Visible Language*, 60(1), 90–119. <https://www.visible-language.org/journal/issue-60-1-metafont-comments>
- Knuth, D. E. (1982). The concept of a meta-font. *Visible Language*, 16(1), 3–27. <https://journals.uc.edu/index.php/vl/article/view/5329/4193>
- Littlejohn, D. (2026). A profession provoked: How Meta-Font struck a nerve. *Visible Language*, 60(1), 120–123. <https://www.visible-language.org/journal/issue-60-1-how-metafont>
- Preston, D. (2026). The terminological development of graphic design: Between office art and social purpose. *Visible Language*, 60(1), 24–57. <https://www.visible-language.org/journal/issue-60-1-graphic-design-terminology>
- Savoie, A., Bernau, K., Daly, W., Haefliger, R., & Baez-Lugo, S. (2026). Automation and artificial intelligence in the type design process: Insights from an industry survey. *Visible Language*, 60(1), 24–57. <https://www.visible-language.org/journal/issue-60-1-automation-type-design>
- Vatány, S., Nguyen, T. H. T., Illés, A., & Bessemans, A. (2026). Evaluating interactive highlighting techniques in digital reading: An empirical study of hover-based line, sentence, and paragraph highlighting. *Visible Language*, 60(1), 58–83. <https://www.visible-language.org/journal/issue-60-1-highlighting-techniques>
- Zender, M. (2026). Why Meta-Font struck a nerve. *Visible Language*, 60(1), 84–89. <https://www.visible-language.org/journal/issue-60-1-why-metafont>



The Terminological Development of Graphic Design: Between Office Art and Social Purpose

David Preston 

Central Saint Martins, University of the Arts London, UK (d.preston[at]csm.arts.ac.uk)

Abstract: The origins of the term graphic design are typically framed through reductive, canonical narratives that trace its emergence to a singular event. This paper challenges such accounts by tracing a more complex and layered trajectory. Focusing on the neglected yet significant 476-page handbook *Graphic Design* by W. G. Raffé—the first book to feature the term in its title—it situates this text within a broader history of Anglophone design discourse. A close reading of this 1927 work reveals three key insights. First, Raffé conceived of graphic design as a socially engaged practice with a civic mission, foregrounding its communicative role in mobilizing the public—this orientation contrasts sharply with dominant, aesthetically driven definitions. Second, he articulated print reproduction as the essential enabling technology that distinguished graphic design from fine art and empowered visual communication to reach the masses at an accelerating speed and scale. Third, he attempted to codify graphic design as a professional discipline, using diagrams and schemas to reify the practice and articulate its principles. This was an important intervention prior to more formal attempts to professionalize graphic design in postwar Anglophone contexts. By recovering this overlooked text and locating it within a longer-term trajectory of development, the paper argues that the term graphic design did not emerge from a single moment or figure but evolved through decades of dispersed adoption. Revisiting Raffé’s foundational work offers valuable historical perspective on ongoing debates about the discipline’s identity and purpose in the post-digital era.

Implications for practice: This examination of disciplinary formation speaks directly to contemporary questions about specialization and expertise. Raffé’s 1927 text reveals how professional legitimacy was built through claims about specialized knowledge and social purpose. His emphasis on print as graphic design’s essential technology reminds us that disciplinary identity has always been tied to enabling media; today’s digital transformation represents continuity rather than rupture. Practitioners should recognize that the term graphic design emerged from specific Anglophone, colonial contexts. This awareness should inform more inclusive, culturally responsive approaches.

@: [ISSUE](#) > [ARTICLE](#) >

Cite this article:

Preston, D. (2026). The terminological development of graphic design: Between office art and social purpose. *Visible Language*, 60(1), 1–23.

First published online April 26, 2026.

© 2026 Visible Language — this article is **open access**, published under the CC BY-NC-ND 4.0 license.

<https://visible-language.org/journal/>

Visible Language Consortium:

University of Leeds (UK)

University of Cincinnati (USA)

North Carolina State University (USA)

As designers reposition themselves “upstream” into strategic roles or grapple with generative AI’s impact, Raffé’s core questions remain urgent: What is graphic design’s social function? What specialized expertise justifies professional status? How should practitioners relate to enabling technologies? Understanding this contested terminological history empowers designers to participate more consciously in ongoing debates about disciplinary boundaries, professional identity, and graphic design’s evolving global relevance.

Keywords: communication; design discourse; graphic design; professionalization; technology

1. Introduction

Despite its ubiquity, *graphic design* remains an underexamined term, with limited scholarship studying its adoption or development. Mainstream surveys posit neat narratives that borrow heavily from one another, eschewing depth and nuance to celebrate recognized milestones. This tendency has resulted in the underdevelopment and oversimplification of research in this area, with generalizations made about the significance of individual heroic figures. The eventual convergence of *graphic* (rooted in inscription and reproduction) and *design* (rooted in intentional form-giving and organization) signaled the crucial integration of artistic composition with purposeful communication; a development whose precise terminological origin remains contested.

The North American W. A. Dwiggins (1922) has been widely credited with coining the term “graphic design” in 1922 (Bierut et al., 1999; Mazur Thomson, 1997; Meggs, 1983), yet few have sought to understand the underlying meaning and relevance of his reference or how it fits into the longer-term adoption of the term. Some claim that Dwiggins’ role in the popularization of the term has been overstated when his usage was only inadvertent. For instance, Paul Shaw (2014a, para. 2) uses empirical evidence to argue that while Dwiggins may be a seminal figure in the history of graphic design, he is not responsible for coining the phrase and “never used the term ‘graphic designer’ in his life.”* Others have challenged the relevance of Dwiggins’ contribution, albeit indirectly, with British design historians arguing that the term was adopted gradually and came to prominence after World War II (Kinross, 1988, 1992; Stiff, 2009). In the US meanwhile, Gowan (1984) claims the term was only granted regulatory and economic recognition in the late 1970s to early 1980s.

* Dwiggins’ (1922) oft-cited article refers to “graphic design” just once, and only inadvertently. Shaw (2014a: para. 3) explains that Dwiggins was “not consciously coining a description for the ‘new design’ that he felt was needed” as he had not figured out what to call it yet; instead, he was likely using the words in place of “graphic art” to vary his writing. Dwiggins (1922) referred to “advertising artist” and “printing designer” when describing the new professional.

During the mid-1990s, *Visible Language* published a sequence of three special issues exploring “Critical Histories of Graphic Design.” Steve Baker (1994, p. 245) regretted “the restrictiveness and bogus neutrality of design history’s conventional linear narratives”; while Victor Margolin (1994, p. 233) called for a narrative approach that “probes more deeply into the way that graphic design has evolved.” This paper adopts a strategy to both zoom in and zoom out. Zooming inward, it draws close attention to one understudied British text significant as an early outlier: Walter George Raffé’s (1927b) *Graphic Design*, a book released five years after Dwiggin’s much-celebrated reference. Zooming outward enables a positioning of Raffé’s book within a broader 20th-century context and in relation to subsequent literature and debates.

Through a close reading of Raffé’s publication, this paper describes and interprets key characteristics to establish the significance of Raffé’s reference to graphic design. It examines his work’s role in the reification of graphic design and the later development of the discipline as a profession. The analysis reveals three threads of interest in relation to Raffé’s book. The first examines his attempts to position graphic design as a discipline with a wider social purpose and civic function to inform and mobilize the public. The second examines the role of print as an essential enabling technology of graphic design. A third considers Raffé’s attempts to frame graphic design as a burgeoning profession with a greater complexity than mere “office art.”

This paper examining the terminological development of graphic design is timely given the challenges facing the discipline. Within specialist communities there is an increasing pluralism to the ways that the term *graphic design* is interpreted and deployed. As the practice has evolved with technological advances and social mores it has become slippery to define. Paul Rodgers and Craig Bremner (2016, p. 22) have argued that design disciplines have fragmented and dissolved beyond recognition, with design today being characterized by “fluid, evolving patterns of practice that regularly traverse, transcend and transfigure disciplinary and conceptual boundaries.” There has also been much discourse around the democratization of design and innovation (von Hippel, 2005; Manzini, 2015), as well as questions around the relevance of professional status for graphic design more specifically (Rock, 1994). The print trade is no longer the nucleus of the profession that it was a hundred years ago (Bonsiepe, 1994), with educational institutions having adapted to reflect this, downplaying associations with print that the word “graphic” can connote and emphasizing “communication” instead.*

* Several British universities emphasize the communicative intent of the discipline with courses in “Graphic Communication” or “Graphic Communication Design,” others emphasize media, with courses in “Graphic and Digital Design,” “Graphic and Media Design” or “Graphics for Games.”

In a context where the boundaries of graphic design are increasingly contested and expanded, it is important to recognize that the term itself emerged within a specific Anglophone professional and industrial setting. Scholars have highlighted that the ways in which graphic design has been written about historically reflect the professional, institutional, and disciplinary frameworks of their authors (Mazur Thomson, 1993; Triggs, 2011). Visual communication practices have, of course, long existed across cultures, but the label *graphic design* reflects a particular epistemic and institutional lineage. Acknowledging this locality does not resolve the broader problem of Eurocentrism within design history, but it allows this study to clarify how one early British text contributed to the consolidation of a term that would later be applied globally. Its eventual international adoption owed more to the 20th-century dominance of English in commerce and professional discourse than to any inherent conceptual superiority over parallel terminologies. The aim here is not to universalize Raffé's perspective, but to understand how such texts helped form the conceptual vocabulary through which later histories—sometimes uncritically—have narrated the discipline.

2. Approach and Theoretical Basis

This paper employs case study methodology, using close textual analysis to examine Raffé's book as a significant but overlooked artifact in the terminological development of the discipline. The research design centers on systematic primary source analysis, treating the text as both a historical document and a conceptual intervention. The analytical approach combines content examination—analyzing the book's structure, arguments, visual elements, and rhetorical strategies. This is paired with contextual historical inquiry situating the text within broader patterns of professional emergence and terminological evolution.

The close reading methodology involves careful attention to language choices, conceptual approaches, and attempts to codify graphic design knowledge through diagrams and schemas. This analysis is triangulated with secondary sources including contemporary trade publications, educational records, and historical accounts to verify claims and establish broader significance. The historical method traces terminological development, examining how Raffé's 1927 usage relates to earlier commercial art discourse and later graphic design professionalization.

This approach allows both detailed understanding of the text's specific contributions to disciplinary formation and assessment of its significance within the trajectory of graphic design's emergence as a recognized professional field. By combining micro-level textual analysis with macro-level historical contextualization, the study illuminates how individual scholarly and professional interventions contribute to processes of disciplinary legitimation and terminological standardization.

The analysis draws on three theoretical lenses. First, professional sociology (Abbott, 1988) examines how occupational groups establish legitimacy by making knowledge claims and setting boundaries. Raffé's efforts to codify knowledge and distinguish it from "commercial art" are read as attempts to establish jurisdictional claims. Second, the paper engages with media theory. Raffé's positioning of print anticipates insights from McLuhan (1964) on media as extensions of human faculties and Benjamin (1935/2008) on the social implications of mechanical reproduction. Third, discourse analysis (Foucault, 1972) examines how terminological shifts reflect and constitute professional identity. The transition from *commercial art* to *graphic design* represents a discursive formation—a systematic way of organizing knowledge that shapes professional practice. These lenses position Raffé's text not merely as historical artifact, but as a site where professional, technological, and discursive forces converge, producing new forms of cultural authority and disciplinary legitimacy.

3. The Site of Study

3.1. Re-Evaluating Raffé

In his study of "commercial art," "graphic art," and "graphic design," Paul Shaw (2014a) identifies several texts published between 1880 and 1980 that warrant closer scrutiny. Shaw (2014b) shows that "graphic" and "design" appeared together in the 19th century, identifying several examples of incidental usage within US newspapers from 1842, 1856 and 1888. He also reports on early 20th-century usage within education. In 1922, California School of Arts and Crafts publicized a "Graphic Design and Lettering" short course; while in 1926, University of Illinois updated their curriculum from "Advertising Design" to "Graphic Design."

Among the examples Shaw draws attention to it is the entry from 1927—Englishman W. G. Raffé's *Graphic Design*, published by Chapman & Hall in London—that takes central focus here. Shaw notes, "this is the first book to have 'graphic design' in its title" (2014a, para. 60), a significant milestone that distinguishes Raffé's work from the other texts in his survey. While Shaw mistakenly writes off Raffé's book, believing that he never referred to graphic design within his text, Steve Baker (1990) identifies it as a book which has been neglected by design historians. This paper argues that the neglect for Raffé's work is misplaced; for despite being name-checked, it is yet to be scrutinized for its significance in the reification of graphic design as a discipline distinct from what had been called "commercial art."

Walter George Raffé, born in Wigan, England in 1881, was an artist, writer and teacher who followed the Art Nouveau style, working primarily with woodcut (Peppin & Micklethwait, 1984). He lived and worked in the Bradford district, studying at Leeds

College of Art whilst an art teacher at a local school. Gaining a scholarship to the Royal College of Art (RCA), he went on to travel, working in India as a teacher and serving as Principal of the Lucknow School of Art (Neville, n.d.). Raffé's appointment to this senior educational position in colonial India reflects the broader patterns of British imperial influence on design education during this period, demonstrating how Eurocentric ideals about design were transposed to other cultures through institutional channels.

By the end of the 1920s, he was editor of the trade journal *New Publicity* and had published three books on art and design: *Art and Labour* (1927a), *Graphic Design* (1927b), and *Poster Design* (1929). *Art and Labour* is a modest small-format edition of just 84 pages. Written and designed by Raffé, it set out his ideas across six ideological and philosophical essays. Raffé claims that art should not be an amusement but rather should be balanced with labor as part of daily activity—his position reflects his interest in the psychology of art and occultism. *Graphic Design* is highly pragmatic in tone, akin to a handbook, with Raffé's advice applied to professional contexts. It is a tome, with 476 pages and significantly larger in scale. The third book, *Poster Design*, is in a similarly large format, but half the extent, with 224 pages. Images take center stage, with full-page illustrations reproduced in full color—within *Graphic Design* images are predominantly black and white and rarely extend beyond half a page, supporting and illustrating the points made within the text.

As an educator, author and editor, Raffé was a figure of influence, but neither his work as a designer nor as editor and writer have been recognized beyond a name-check by historians. He was not a preeminent designer, unlike commercial artists like Edward McKnight Kauffer, Tom Purvis, or Austin Cooper, who were regularly celebrated in trade journals and annuals. Nevertheless, as editor of the journal *New Publicity* he brushed shoulders with the most respected practitioners. His key work *Graphic Design* (1927b) was well received in its time, with Chapman & Hall's British edition being supplemented by a US version from Bridgman the following year (Raffé, 1928) and a "Cheap Edition" being released four years later (Raffé, 1932).*

3.2. General Overview of the Book

In *Graphic Design*, Raffé (1927b) presents the discipline as a deeply technical endeavor, devoting large portions of the book to how images are reproduced through stylistic choices, printing methods and production techniques. His command of print's possibilities and limitations is evident throughout the examples he includes. For example, he shows how the same image can be rendered for print using different mark-making techniques such as outline, silhouette, stipple effect, or oblique line (Figure 1).

* These editions were likely identical in content given the closely matched extent of each volume.




Figure 1. Sample page showing a sower figure rendered in different ways. The five variations presented in the book are not given as stylistic alternates, but rather as forms suited to different means of reproduction. Image from: Raffé, 1927b, *Graphic Design*, p. 75. Public domain.

Section One on Technical Drawing outlines the purpose of drawing, the processes involved, and the tools that shape its outcomes, with guidance on proportion, composition, and adapting drawings to reproduction methods (Figure 2). Section Two on Technical Design broadens the focus to the psychological appeal of images and the aims of advertising, weaving together topics such as illustration, composition, tone, mood, legibility, scale, style, and typographic decision-making. Raffé uses these themes to show how graphic design operates across books, posters, stamps, labels, trademarks, fashion illustration, and cartoons. Section Three on Technical Methods covers reproduction processes, preparing artwork for print, handling photographs, selecting papers and boards, and the instruments and devices of production. The final section turns to the practicalities of professional life—studio organization, administrative routines, pricing and rights, client relations, agents, publicity, originality and plagiarism, and participation in competitions and exhibitions. The tone throughout is highly practical and professionally oriented, distinguishing this book from Raffé’s other publications.

While many British publications focused on national examples, Raffé’s text is international in scope, but remains within the Global North. Examples come from North America, Germany, France, Switzerland, Italy and Japan. From North America, the examples favor the direct commercialism of advertising, whereas the continental European references show artistic flair and a more poetic, less direct sales pitch.

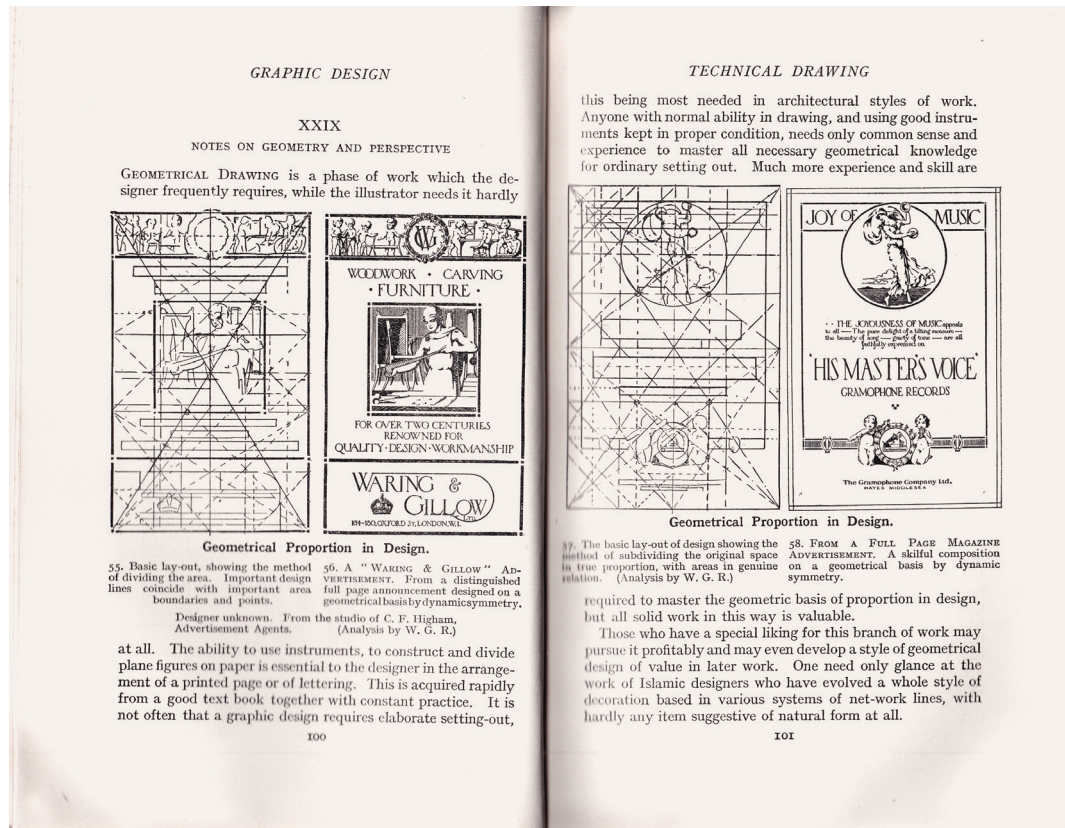


Figure 2. Sample spread showing design deconstructions focused on geometrical proportion in layouts that incorporate text and image content in harmony. Image from: Raffé, 1927b, *Graphic Design*, pp. 100–101. Public domain.

The book’s front cover (Figure 3) presents a white design on a navy-blue ground. The words “Graphic Design, W. G. Raffé” are set in all capitals, centered across four lines. Surrounding the text is a composite of pictures, including trademarks, design iconography and commercial illustrations. These represent typical “jobbing” work for commercial artists. The tone is traditional, channeling the Arts and Crafts movement, with hand-rendered serif lettering far removed from the perceived modernity of typographers of the era like Jan Tschichold or Eric Gill.

3.3. Training for a Profession

Raffé’s book presents itself as a how-to guide, “a good and straightforward book, written with the needs of students in mind” (1927b, p. viii), aiming to “help the young man or woman who desire to master graphic design and to work at it professionally” (p. vii). Several points stand out. Raffé emphasizes that the field is open to both men and women—a seemingly minor detail, yet significant given the era’s prevalent sexism. He also signals graphic design’s professional status from the outset, framing it as a serious occupation and later sharing his “professional secrets” (p. 422). This is surprising




Figure 3. Front cover design of W. G. Raffé's *Graphic Design*. Image from: Raffé, 1927b. Public domain.

considering that graphic design only came to be regarded as a serious professional activity after WWII. The Society of Industrial Artists (SIA) campaigned for design status from 1930 onwards, but limited headway was made until the society regrouped after the war (Armstrong, 2015).*

For entering the profession, Raffé believes that personal interaction between the novice and expert is key. He recommends tuition from an art school or a capable designer, “a weekly visit is often sufficient for the precise personal criticism and corrections” (1927b, p. vii). Raffé has several reservations about art school teaching, but recommends students seek formal education alongside experience on the job. He notes misalignments between the principles eschewed by educational institutions and industry demands, questioning the treatment of drawing within art schools as serving design and not the other way around. This highlights his awareness of disparities between professional methods and pedagogical principles prioritized by educational institutions.


* Outside Britain, relevant organizations include the AIGA (American Institute of Graphic Arts) founded in 1914 and the AGI (Alliance Graphique Internationale) which was incorporated in Paris in 1952.

Raffé's attempts to frame graphic design as a professional activity remain significant, establishing the discipline as distinct from commercial art. Commercial Art Limited distributed their bimonthly periodical from 1922 onwards, while Pitman released Verney Danvers' (1926) handbook, *Training in Commercial Art*. By comparison Raffé's book stood out in an increasingly competitive marketplace. Elsewhere publications flagged their ties with advertising, with Percy Bradshaw's (1925) *Art in Advertising* being key; others emphasized an editorial angle, like Vincent Steer's (1930) *Printing Design and Layout*, or the French printer Alfred Tolmer's (1931) *Mise en Page*.

Raffé (1927b) seeks to bestow integrity upon the discipline through elaborate diagrams and models. These demonstrate a rigorous attempt to codify and professionalize graphic design knowledge. He presents detailed schemas governing both the conceptual and formal aspects of work. These include charts based on human needs ("Advertising subject chart," Figure 4, p. 149) and dissecting images' graphic, dramatic, and decorative values ("Chart of illustration," p. 153). Further diagrams establish a theoretical basis for practice, including the relationship between conceptual and formal characteristics in advertisements, a method moving design beyond naturalistic imitation ("Relation of 'Nature' to Graphic Design," Figure 5, p. 161), and the psychological relations across all phases of the discipline ("The Eternal Triangle," Figure 6, p. 164).

Beyond these theoretical schemas, Raffé includes charts, lists, and itineraries covering printing methods, paper formats, poster standards, and historical printing techniques. Through this technical data, he seeks to codify graphic design into rules and laws following techno-scientific rationality. This establishes graphic design as more than art reframed for business ends, but as a practice with an esoteric knowledge base.

On a surface level, these diagrams appear authoritative, giving the discipline a professional sheen. Closer examination, however, reveals many of these schemas to be far-fetched, imposing artificial systematization onto practices that resist rigid codification. The elaborate frameworks, while visually compelling, attempt to link esoteric concepts in ways that overstate the systematic nature of graphic design. Figure 6, titled "The Eternal Triangle," exemplifies this tendency, presenting what Raffé describes as "the psychological relations recurrent in all phases of graphic design" (1927b, p. 164). Yet the diagram offers little practical guidance, and readers would struggle to apply its abstract schema. Its mystical overtones reflect Raffé's documented interest in occultism, infusing his professionalizing project with esoteric spirituality that sits uneasily alongside his claims to scientific rigor. This tension exposes the challenges faced by early attempts to theorize graphic design as a systematic discipline.



87. Advertising Subject Chart Showing Relation of Advertisable Items.

Figure 4. An example of one of Raffé’s many diagrams attempting to codify aspects of graphic design practice. It is notable that he incorporates both advertising and illustration within his conception of professional graphic design. Image from: Raffé, 1927b, *Graphic Design*, p. 149. Public domain.

4. Defining Graphic Design

4.1. Print as an Essential Enabling Technology of Graphic Design

The triumphant march of the printing press is blasting the trail for the designer, for, wherever there is print, there also graphic design is necessary. (Raffé, 1927b, p. 442)

For Raffé the power of graphic design lies in its potential for multiplicity. As he explains, fine art will no longer be carried from home to home, for images will be fixed all around, with art and design prevalent for the masses. Easel pictures are doomed to extinction, as “the ordinary single picture does not pull its social weight ... it represents too often only the hazy idea of one man” (Raffé, 1927b, p. 444). Raffé positions fine art as an inferior pursuit with limited social impact by comparison to graphic communication.

Raffé views print as an essential enabling technology for graphic design, through which the discipline finds its relevance, both socially and economically. Modern graphic design is “specifically designed for duplication by mechanical aid” (Raffé, 1927b, p. 8).




Figure 5. A further example of Raffé’s attempts to codify aspects of graphic design practice in diagrammatic form with a focus on the use of mass and line. Image from: Raffé, 1927b, *Graphic Design*, p. 161. Public domain.




Figure 6. A more holistic diagram on what Raffé calls “The Eternal Triangle” bringing together several of his concerns into one schema. Image from: Raffé, 1927b, *Graphic Design*, p. 164. Public domain.

As such, “the competent designer studies carefully his chosen process of duplication” (p. 9) and should tailor his graphics to the printing method and material in use. Raffé’s emphasis on print’s democratizing potential anticipates Benjamin’s (1935/2008) analysis of mechanical reproduction’s revolutionary implications. Where Benjamin (1935/2008, pp. 23–25) sees mechanical reproduction as a threat to art’s “aura,” Raffé champions multiplicity as graphic design’s social strength. This theoretical divergence reveals how professional discourse can oppose critical theory’s concerns, with Raffé seeing mechanical reproduction as a democratizing force and conduit for collective existence. This echoes Gitelman’s (2014) ideas about the social life of paper, as well as Innis (1951) and McLuhan’s (1964) arguments about the history of mass media being central to the history of civilization.

Raffé plays up the role of the graphic designer as a communicator, acknowledging the distinct communicative aspect in graphic design:

Graphic Design will be required more and more, with increased means of transport, to touch the nations of the world more quickly than even the printed word can do. People move about from city to city; books and papers also move about; between these two modern facts graphic design will achieve its work, as a means of communication. (1927b, p. 444)

The word “communication” is unusual in early 20th-century commercial art—though the terms now commonly conjoin in Communication Design courses.* Communication was not typically associated with design at this time but referred to transportation or transmission, as evident in the history of city planning. Raffé’s idea of communication links to speed and movement as he refers to the transportation of ideas through society. This recalls Schivelbusch’s (2014) concept of the machine ensemble—linking progress to acceleration in the scale, scope and speed of machine integration. Raffé’s emphasis on the communication and transportation of ideas through society provides alternate focus versus the aesthetic and technical connotations of the word “graphic,” or the financial connotations of “commercial”—as in commercial art. This defining focus on what graphic design can do, as opposed to how it looks, or how it is made is unusual in the history of graphic design.

4.2. Office Artists and the Value of Original Design

According to Raffé, the “ordinary commercial man” (1927b, p. 3) is becoming increasingly aware of design, with opportunities emerging for freelancers to take commissions from publishers, advertising agencies and printers. Although enthusiastic about freelance life, he is condescending about in-house designers, claiming that commercial art in Britain is stunted by the limitations upon “office artists” required to churn out “maximum quantity at a known minimum quality” (1927b, p. 12). These artists are inhibited by their work conditions, being “cooped up in a dingy office all day” and unable to roam the streets in search of inspiration (1927b, p. 14). Graphic design is more than simply office art and should not be viewed as a formulaic practice to be conducted by rote.

Raffé (1927b, p. 5) questions the relationship between design and drawing, saying that although drawing is “a means of expression fundamental to graphic design ... it does not properly occupy the premier place so often wrongly claimed for it.” He criticizes how art schools concentrate on routine drawing competence at the expense of purposeful, applied design. As he understands it, students are incorrectly instructed on how to draw, before creating their designs; as in his mind, design is more sophisticated and

* Shaw (2014a) claims “visual communication” came to the fore in the late 1960s, as “graphic design” was on its way to common acceptance within the English-speaking world. He links these developments to the publication of John Cataldo’s (1966) book on the subject.

should be treated as primary. Raffé (1927b, p. 6) argues that “a really good designer can succeed, even if he cannot ‘draw,’” a surprising claim given that his opening chapter focuses on drawing. He downplays its importance, suggesting that drawing errors are easily corrected, whereas weak design is not. As he notes, drawing is of little significance without “ability in the much higher phase of original design” (1927b, p. 16). This underscores that, for Raffé, the designer’s real value lies in strategic judgement rather than technical execution.

Raffé emphasizes design as the critical element of the practice. Design, for him, infers a bridging force linking aesthetics with printing know-how. “The power of design is the synthesizing ability which links together craft knowledge and methods to the end desired” (1927b, p. 5). He understands design to be a purposive and communicative act linked to an audience; a strategic form of image-making allied with prowess in print production. He describes drawing as a “statement of information” or activity of “scientific analysis,” while design centers around “the deliberate selection and arrangement of material.” Raffé substantiates his views on drawing with various diagrams and talks through three stages of drawing: “preliminary,” “drawing,” and “design.” These diagrams elevate design above its associations with rudimentary drawing processes.

Graphic design is often conceptualized as a synthesis of text and image, with this becoming a defining characteristic of the discipline. For Raffé typography receives less attention than the pictorial or technical components. He identifies lettering as one of the three main areas of printed work: “lettering by type or hand,” “illustration” and “decoration.” Yet he sets aside just 31 pages for the subject, versus 127 for technical drawing. This reflects the technical conditions of the era, with commercial artists producing images on posters, handbills or books, but not typically being responsible for setting movable type. Although they may have rendered the lettering or titling on their artworks, body text would have been set by specialist typesetters for whom other handbooks were available.*

Raffé’s examples favor the pictorial expression and aesthetic flair of commercial art image-makers over modern typography. Jan Tschichold’s (1925, 1928) ideas on modern design began to spread through the continent from the ’20s, but decades passed before they became commonplace in British culture and education through the work of Herbert Spencer (1952), Anthony Froshaug and others. Raffé’s framing of graphic design can be seen to differ from the conception which became entrenched in the 1960s, where text and image were regarded as equal partners in a fully integrated relationship. For

* Many poster artists of the era relied on print specialists to prepare artworks for press and tidy up their cursory attempts at lettering.

Kinross (1992) the increasing prevalence of photographic technology helped to enable this pairing of text and image to prosper as the century progressed.

5. Raffé's Book in the Context of Broader Terminological Developments

5.1. Commercial Art and Graphic Design (1920–1960)

For Raffé, graphic design is largely interchangeable with commercial art, switching between the terms as he claims his book will not immediately make you a “successful commercial artist” (1927b, p. viii). Falling back on the established terminology implies a lack of confidence to commit, demonstrating how new ideas are often framed using recognizable concepts. These two terms he refers to frequently, but there is no clear pattern of usage, suggesting that he varied his language to stimulate the reader and avoid repetition. He also mentions “graphic art,” “commercial design,” “commercial graphic design,” “commercial graphic art” and “artistic publicity.”

Given the entanglement between graphic design and commercial art, it is essential to trace the historical development of these terms to explain how graphic design achieved prominence. Commercial art had been prevalent during the inter-war period, with a journal of the same name published between 1922 and 1936.* Richard Guyatt (1963, p. 21) claims:

The term ‘Graphic Design’ was unknown before the war. Those were the days of ‘Commercial Art,’ of the ‘Poster Artist,’ when art applied itself to commerce while still wearing a beard and beret as a reminder of noble origins.

According to Kinross (1988), “graphic design” only began to be established in Britain after 1948 when the RCA introduced the first course in the subject under Richard Guyatt. Frayling (2013, p. 468) claims that Guyatt had been “the first to use the phrase publicly in Britain”—but Raffé had published his text two decades earlier. The choice of course name is curious, with Guyatt explaining that they didn’t know what else to call it. During the restructuring of 1948, “School of Publicity Design” had been proposed, but this led to a backlash in *The Times*, with critics resistant to the vulgarity of the word “publicity.” As Guyatt (1963, pp. 21–22) recalls, “with a certain sense of relief, but not much conviction, the name ‘Graphic Design’ was chosen. No one was quite sure what it meant, but it had a purposeful ring.”

In the United States, Leon Friend and Joseph Hefter’s (1936) *Graphic Design: A Library of Old and New Masters in the Graphic Arts* is considered significant, being broadly similar

* *Commercial Art* became *Commercial Art and Industry* in 1932, then later *Art and Industry* in 1937.

in scope with Raffé's book. As previously discussed, short courses in graphic design had existed in the US in the 1920s, but Rob Roy Kelly (2001) argues that the period from 1950 to 1955 was most significant for the discipline's development. It was during this period that Yale established their influential graphic arts program under Department Chair Josef Albers. Notably, the events at Yale loosely echo developments at the RCA; however, it was not until the late 1950s, when Guyatt visited Yale with RCA Director Robin Darwin, that program head Alvin Eisenman was encouraged to change the course name from "graphic arts" to "graphic design"—a change that was implemented soon after (Kelly, 2001).

At Yale the importance of continental European ideas is clear, with the influential Albers having trained and taught at the Bauhaus. The lineage of influence at the RCA is less clear, as they were known to disregard continental modernism, virtually ignoring foreign developments until much later when designers such as Anthony Froshaug joined (Seago, 1995). While the RCA remained traditional, emphasizing steel and woodblock gravure, fine typography, and bookbinding, London's Central School embraced a progressive, avant-garde approach, drawing on continental modernism and prioritizing experimental typographic design over mere letter formation. By 1953 the Central School updated their Department of Book Production, reflecting developments at the RCA, renaming it the School of Book Production and Graphic Design (Johnstone, 1953).

It took decades for the term graphic design to supplant commercial art in Britain. Another milestone was the publication of John Lewis and John Brinkley's (1954) *Graphic Design with Special Reference to Lettering, Typography and Illustration*. Lewis and Brinkley were tutors at the RCA, with Department Head Guyatt responsible for the dust-jacket and introduction. The book reflects the tendencies of the department in which they taught, with illustrated examples focusing on British traditionalism and nostalgia. Photography is notably absent, and they prefer to focus on craft over the commercial or conceptual bases of the discipline. By separating the book into the fields of "typographic design" and "illustration design," they go against the conception of graphic design as a conjunction of text and image.

Together, these examples show the term *graphic design* coming into public consciousness through gradual usage, rather than a "big bang" moment. Despite the precedent for the term established by Raffé, commercial art remained dominant during the first half of the century in Britain. After WWII, the status and opportunities for design in Britain shifted significantly, as reconstruction efforts were guided by utopian ideals of social betterment (Maguire & Woodham, 1997). As such, design was seen as a tool which could contribute to the envisioning of a new and better world. From this context of social and economic reconstruction, a new generation of graphic designers emerged

from art school determined to make their mark and distinguish themselves from their predecessors.

5.2. Graphic Design from 1960

While the educational initiatives of the 1950s laid the foundations for graphic design to flourish, the '60s saw it become fully legitimized as a professional activity. By this time graduates from newly established courses at the RCA, Central and Yale could develop their careers, applying what they learnt to their own professional practices. In 1959 the Association of Graphic Designers London (AGDL) was formed by Derek Birdsall, with an exhibition occurring the following year under the title "Graphic design: London" (Braybon, 2018). This led in turn to the publication of *17 Graphic Designers London* (Commander, 1963a), a book showcasing the work of 11 design practices, including a trio of groups (Fletcher/Forbes/Gill, BDMW, and Kinnear Calvert). Others worked independently within publishing, art direction or editorial design. Of those featured, most were young and had graduated from London art schools after the war, many studying together. This collective positioned themselves as the doyens of a new age of graphic design in Britain.

Art Director of printing firm Balding + Mansell, John Commander was a key figure, serving as editor on *17 Graphic Designers London*, Chairman for the Designers & Art Directors' Association (D&AD), and helping to organize the first D&AD exhibition and annual (Commander, 1963b). The strong overlap between these events is significant, with many of those featured in *17 Graphic Designers London* serving as jury members for the D&AD show. Those involved competed with the established SIA, representing a more specialized and progressive vision for the discipline. They associated commercial art with the older generations and understood it to be fusty and unsophisticated. As Paul Stiff (2009) explains, commercial art was seen as merely picture making for business, whereas graphic design was to be viewed as sophisticated and rational. Collectively, the events of 1963 mark a distinct break from the painterly whimsy of British commercial art. Continental ideas came to the fore, with sans-serif typefaces increasingly ubiquitous. The commercial wit of the American approach combined confidently with the structural flamboyance of European typographic principles and ideas.

6. Conclusion

This research has demonstrated that the emergence of *graphic design* as a disciplinary term was more complex and geographically distributed than previously understood. Rather than originating with a single figure, or emerging suddenly, the term's development involved multiple practitioners across different national contexts working incrementally to establish new professional categories. Adoption should thus be

understood as involving several generations of practitioners spread across time periods and multinational terrain.

This paper examined W. G. Raffé's 1927 text *Graphic Design* as a significant but overlooked contribution to this terminological development. At a time when "commercial art" dominated professional discourse, Raffé's use of "graphic design" in his book title represented an anomalous outlier. Though the book received popular recognition in its time, it has been neglected by design historians, despite some recognizing its significance as the first book to carry the title.

6.1. Raffé's Foundational Contributions

Three core findings arise, challenging conventional narratives about graphic design's historical emergence and offering insights into the discipline's foundational concepts. Most significantly, Raffé positioned graphic design as a practice with vital social function, emphasizing its communicative capacity to mobilize public consciousness. He articulated how the multiplicity afforded by print processes empowered graphic design to transport ideas through society and generate collective awareness. Prioritizing graphic design's social benefits over its formal attributes or creative dimensions, Raffé articulated the discipline's wider purpose and value—contrasting sharply with typical aesthetic-focused conceptualizations. In linking the discipline to technological advancement and positioning speed, scale, and movement as markers of its future value, he anticipates later arguments by media scholars like Innis, McLuhan, and Gitelman regarding the transformative social power of visual culture.

Raffé's second major contribution was articulating graphic design as more than commercial picture-making. He positioned print reproduction as graphic design's essential enabling technology. Not merely a tool, but the fundamental basis for the discipline's social relevance and professional distinction from fine art. For Raffé, original design represented a "much higher phase" (1927b, p. 16) than drawing alone, requiring designers to move beyond personal expression toward a deep technical knowledge of print production techniques which informed how they designed. This understanding of the medium as deterministic foreshadowed McLuhan's (1964) thinking, anticipating actor-network theorists (Latour, 2005) insights about technology's agency in social formation.

Finally, Raffé's efforts to codify graphic design knowledge through extensive diagrams, schemas, and analytical frameworks represented an early attempt at disciplinary professionalization. By treating picture-making as a techno-science requiring specialized esoteric knowledge, he aimed to establish concrete foundations for professional legitimacy, paralleling developments in established professions like medicine and law. This was particularly significant as design was not widely regarded as a legitimate profession

in Britain until the post-war period, and art schools did not teach the subject until the late 1940s and '50s. His focus on education–industry relationships and practical career guidance for aspiring designers helped build momentum for the professionalization project, claiming jurisdiction (Abbott, 1988) over competing domains like advertising.

6.2. Contemporary Challenges

Revisiting Raffé's attempt to reify "graphic design" is relevant when considering challenges facing the discipline's identity a century later. While print was graphic design's essential enabling technology for Raffé, digital media has replaced it as the dominant communication platform. This shift has reshaped the professional frameworks associated with early 20th-century design practice.

The ideal of the professional designer has also been unsettled by the democratization of design tools and the rise of "design thinking" as a general problem-solving approach. Professional bodies such as Britain's Chartered Society of Designers have struggled to maintain relevance, while some practitioners have repositioned themselves "upstream" into strategic roles beyond project delivery and material craft. The emerging emphasis on issue-based rather than discipline-based practice raises fundamental questions about where specialized expertise sits in relation to generic, cross-sector design capabilities.

As Afonso Matos (2022) observes, definitions of graphic design have been broadening and diffusing, with technical proficiency overshadowed by notions of design as a mode of thinking or attitude. The development of generative artificial intelligence has added further pressure to creative labor markets, even as graphic design continues to shape visual culture. Image carousels and short-form videos have become central communication channels for younger audiences in the post-digital era, extending the accelerated visual communication dynamics that Raffé identified in the 1920s—though now screens have replaced printed pages.

Raffé's 1927 intervention articulated a professional and technical jurisdiction for graphic design, distinguishing it from "commercial art." Yet in grounding this jurisdiction in the technical and organizational conditions of Western print culture, his account also reflects the particular cultural and institutional setting from which the discipline emerged. Raffé's own career, including his time as Principal of the Lucknow School of Art, underscores how pedagogical and professional models developed in one context were transplanted to others, sometimes with limited sensitivity to local traditions. Recognizing this does not diminish his contribution; rather, it situates graphic design as a historically contingent construct shaped by the global conditions and assumptions of its time.

As graphic design grapples with questions of inclusivity and global relevance, these historical foundations become important for understanding how disciplinary

concepts have developed and what they may have obscured. The terminological shift from “commercial art” to “graphic design” involved claims to professional legitimacy that inevitably reflected the cultural contexts in which such claims were made. This process helped consolidate an Anglophone canon, often at the expense of broader histories of visual communication across Africa, Asia, and the Americas. Attending to these coexisting traditions—rooted in diverse materials, symbolic systems, and social functions—does not replace the history Raffé represents but expands the frame through which the field’s development can be understood.

Raffé’s contribution must also be understood within its linguistic and geographic particularity. While this study has traced how “graphic design” became established within Anglo-American professional discourse, the term’s subsequent global adoption reflected the geopolitical ascendancy of English in international professional spheres as much as any conceptual consensus. Other linguistic traditions maintained parallel terminologies—such as the German *Gebrauchsgraphik* (Aynsley, 1992), the Japanese *Shōgyō Bijutsu* (Weisenfeld, 2000), or the French and Spanish *Arts Graphiques* and *Artes Gráficas*—each embedding different emphases and cultural priorities within their professional frameworks.

By revisiting figures such as Raffé and tracing their influence, this study opens a pathway toward a more plural understanding of how graphic design emerged as a recognized field of practice and knowledge. This perspective supports a more informed engagement with contemporary debates about the discipline’s future direction and global relevance. As graphic design approaches its second century as a named form of labor, it remains unclear which technologies and forms of practice will define its next phase. Yet the questions Raffé engaged—regarding design’s social function, its relationship to enabling technologies, and its claims to professional jurisdiction—continue to resonate as the field navigates ongoing technological change, increasing cultural diversity, and evolving professional identities.

7. Acknowledgments

During the preparation of this work, the author used Claude AI to support sentence clarity and readability. All content, ideas, and scholarly arguments remain the author’s own. Images reproduced in this article derive from a work first published in the United Kingdom in 1927 and in the United States in 1928 and are reproduced by David Preston on a public-domain basis under US copyright law, which places works published prior to January 1, 1929, in the public domain.

8. References

- Abbott, A. (1988). *The system of professions: An essay on the division of expert labor*. University of Chicago Press.
- Armstrong, L. (2015). Steering a course between professionalism and commercialism: The Society of Industrial Artists and the code of conduct for the professional designer 1945–1975. *Journal of Design History*, 29(2), 161–179. <https://doi.org/10.1093/jdh/epv038>
- Aynsley, J. (1992). Gebrauchsgraphik as an early graphic design journal, 1924–1938. *Journal of Design History*, 5(1), 53–72. <https://doi.org/10.1093/jdh/5.1.53>
- Baker, S. (1990). The sign of the self in the metropolis. *Journal of Design History*, 3(4), 227–234. <https://doi.org/10.1093/jdh/3.4.227>
- Baker, S. (1994). A poetics of graphic design? *Visible Language*, 28(3), 245–260. <https://journals.uc.edu/index.php/vl/article/view/5607>
- Benjamin, W. (2008). *The work of art in the age of mechanical reproduction* (J. A. Underwood, Trans.). Penguin Books. (Original work published 1935)
- Bierut, M., Helfand, J., Heller, S., & Poynor, R. (Eds.). (1999). *Looking closer 3: Classic writings on graphic design*. Allworth.
- Bonsiepe, G. (1994). A step towards the reinvention of graphic design. *Design Issues*, 10(1), 47–52. <https://doi.org/10.2307/1511655>
- Bradshaw, P. V. (1925). *Art in advertising: A study of British and American pictorial publicity*. The Press Art School.
- Braybon, A. (2018). Town shaped the sixties. *Eye*, 24(96), 54–71. <https://www.eyemagazine.com/feature/article/town-shaped-the-sixties>
- Cataldo, J. (1966). *Graphic design and visual communication*. International Textbook Company.
- Commander, J. (Ed.). (1963a). *17 graphic designers London*. Balding and Mansell.
- Commander, J. (Ed.). (1963b). *Design & art direction '63*. Balding and Mansell.
- Danvers, V. L. (1926). *Training in commercial art*. Pitman & Sons.
- Dwiggins, W. A. (1922, August 29). New kind of printing calls for new design. *Boston Evening Transcript*.
- Foucault, M. (1972). *The archaeology of knowledge* (A. M. Sheridan Smith, Trans.). Pantheon Books.
- Frayling, C. (2013). Guyatt, Richard Gerald Talbot. In L. Goldman (Ed.), *Oxford Dictionary of National Biography 2005–2008*. Oxford University Press.
- Friend, L., & Hefter, J. (1936). *Graphic design: A library of old and new masters in the graphic arts*. Whittlesey House.
- Gitelman, L. (2014). *Paper knowledge: Toward a media history of documents*. Duke University Press.
- Gowan, A. (1984). Review of *A History of Graphic Design; History of Graphic Design and Communication: A Sourcebook*, by P. B. Meggs & C. Ashwin. *Design Issues*, 1(1), 87–89. <https://doi.org/10.2307/1511549>
- Guyatt, R. (1963). *GraphicsRCA: Fifteen years' work of the School of Graphic Design, Royal College of Art*. Royal College of Art.
- Innis, H. (1951). *The bias of communication*. University of Toronto Press.
- Johnstone, W. (1953). Graphic design at the Central School. *Penrose Annual*, 47, 58–60.
- Kelly, R. R. (2001). The early years of graphic design at Yale University. *Design Issues*, 17(3), 3–14. <https://doi.org/10.1162/074793601750357141>
- Kinross, R. (1988). From commercial art to plain commercial. *Blueprint*, 48, 29–36.

- Kinross, R. (1992). Conversation with Richard Hollis on graphic design history. *Journal of Design History*, 5(1), 73–90. <https://doi.org/10.1093/jdh/5.1.73>
- Latour, B. (2005). *Reassembling the social: An introduction to actor-network theory*. Oxford University Press.
- Lewis, J., & Brinkley, J. (1954). *Graphic design: With special reference to lettering, typography and illustration*. Routledge & Kegan Paul.
- Maguire, P., & Woodham, J. (1997). *Design and cultural politics in postwar Britain: The Britain Can Make It exhibition of 1946*. Leicester University Press.
- Manzini, E. (2015). *Design, when everybody designs: An introduction to design for social innovation*. MIT Press.
- Margolin, V. (1994). Narrative problems of graphic design history. *Visible Language*, 28(3), 233–244. <https://journals.uc.edu/index.php/vl/article/view/5606>
- Matos, A. (Ed.). (2022). *Who can afford to be critical?* Set Margins' publications.
- Mazur Thomson, E. (1997). *The origins of graphic design in America, 1870–1920*. Yale University Press.
- Mazur Thomson, E. (1993). The literature of graphic design. *Art Documentation: Journal of the Art Libraries Society of North America*, 12(1), 7–11.
- McLuhan, M. (1964). *Understanding media: The extensions of man*. Signet Books.
- Meggs, P. B. (1983). *Meggs' history of graphic design*. J. Wiley & Sons.
- Neville, C. (n.d.). RAFFE, Walter George. *Not Just Hockney*. <https://www.notjusthockney.info/raffe-walter-george/>
- Peppin, B., & Micklethwait, L. (1984). *Book illustrators of the twentieth century*. Arco Publishing.
- Raffé, W. G. (1927a). *Art and labour: Six essays*. C. W. Daniel.
- Raffé, W. G. (1927b). *Graphic design*. Chapman & Hall [UK Edition].
- Raffé, W. G. (1928). *Graphic design*. Bridgman Publishers [US Edition].
- Raffé, W. G. (1929). *Poster design*. Chapman & Hall.
- Raffé, W. G. (1932). *Graphic design*. Chapman & Hall ['Cheap Edition'].
- Rock, M. (1994). On unprofessionalism. 2x4. <https://2x4.org/ideas/1994/on-unprofessionalism>
- Rodgers, P. A., & Bremner, C. (2016). The concept of the design discipline. *Dialectic*, 1(1), 19–38. <https://doi.org/10.3998/dialectic.14932326.0001.104>
- Schivelbusch, W. (2014). *The railway journey: The industrialization of time and space in the nineteenth century*. University of California Press.
- Seago, A. (1995). *Burning the box of beautiful things: The development of a postmodern sensibility*. Oxford University Press.
- Shaw, P. (2014a, June 4). Graphic design: A brief terminological history. *Blue Pencil*. <https://www.paulshawletterdesign.com/2014/06/graphic-design-a-brief-terminological-history/>
- Shaw, P. (2014b, June 8). Graphic design: More on the terminology of a profession. *Blue Pencil*. <https://www.paulshawletterdesign.com/2014/06/graphic-design-more-on-the-terminology-of-a-profession/>
- Spencer, H. (1952). *Design in business printing*. Sylvan Press.
- Steer, V. (1930). *Printing design and layout*. Virtue.
- Stiff, P. (2009). Austerity, optimism: Modern typography in Britain after the war. In P. Stiff (Ed.), *Typography papers 8* (pp. 5–68). Hyphen Press.
- Tolmer, A. (1931). *Mise en page*. The Studio Ltd.
- Triggs, T. (2011). Graphic design history: Past, present and future. *Design Issues*, 27(1), 3–6.

- Tschichold, J. (1925). *Sonderheft elementare typographie* [Special issue elemental typography]. *Typographische Mitteilungen*, 22(10).
- Tschichold, J. (1928). *Die neue typographie* [The new typography]. Verlag.
- von Hippel, E. (2005). *Democratizing innovation*. MIT Press.
- Weisenfeld, G. (2000). Japanese modernism and consumerism: Forging the new artistic field of Shōgyō Bijutsu (Commercial Art). In E. K. Tipton & J. Clark (Eds.), *Being modern in Japan: Culture and society from the 1910s to the 1930s* (pp. 75–98). University of Hawai'i Press.

Author

David Preston is a senior lecturer and researcher in graphic communication design at Central Saint Martins, UAL. A senior fellow of Advance HE and SEDA-recognized supervisor, he leads the Strategy & Identity platform, exploring branding and communication design through a social purpose lens. His 20-year practice spans research, knowledge exchange, and pedagogy, including Innovate UK partnerships. David convenes the Graphic Communication Doctoral Node, supporting UAL-wide researcher development, and has presented research internationally through conferences and invited lectures. His doctorate examined how programmatic branding transformed British graphic design practices, research now published as the monograph *The Development of Corporate Design* (Bloomsbury, 2025).



Automation and Artificial Intelligence in the Type Design Process: Insights from an Industry Survey

Alice Savoie^a, Kai Bernau^a, Wayne Daly^a, Raphaela Haefliger^a, and Sebastian Baez-Lugo^b 

^a ECAL/University of Art and Design Lausanne (HES-SO); ^b EPFL+ECAL Lab, Swiss Federal Institute of Technology Lausanne/EPFL

Corresponding author: Alice Savoie (alice.savoie[at]ecal.ch)

Abstract: This article investigates how automation (both deterministic and artificial intelligence-based) is integrated into professional type design practice, a field with exacting standards of craft and which relies on specific and long-established working methods. Drawing on an online survey conducted in early 2025, alongside detailed follow-up interviews with select type practitioners, we map current practices and attitudes, as well as the perceived risks and opportunities in the field of automation for type design in general, and the implementation of artificial intelligence (AI) in particular. Data analysis established that deterministic, rule-based automation is near-ubiquitous in type designers' workflows, and is already used for a variety of tasks such as interpolation, glyphset expansion, and various font engineering tasks. In contrast, AI tools have currently only been adopted by a minority of practitioners, and are largely being used for adjacent tasks such as writing code, gathering project documentation, or generating proofing strings. The majority of respondents expressed strong resistance to automating what they identify as the creative core of their work (e.g., sketching, drafting a basic alphabet, marking proofs), but show willingness to delegate the most labor-intensive, technical operations to software, with kerning repeatedly identified as the leading candidate for further automation—provided that human oversight and decision making remain throughout the process. Ethical concerns (such as training data provenance, lack of transparency, and environmental costs) lead to a cautious attitude towards generative AI, a position also fueled by some expressed anxieties about corporate concentration. We argue that sustainable and worthwhile innovation in typeface design should prioritize assistive tools that are transparent and encourage human decision-making, in order to optimize routine work without compromising iterative practices through which designers acquire judgement. Such tools would ideally balance streamlined workflows with the acquisition and reinforcement of highly specific skills, which in turn enable designers to preserve qualitative typographic standards.

@: [ISSUE](#) > [ARTICLE](#) >

Cite this article:

Savoie, A., Bernau, K., Daly, W., Haefliger, R., & Baez-Lugo, S. (2026). Automation and artificial intelligence in the type design process: Insights from an industry survey. *Visible Language*, 60(1), 24–57.

First published online April 26, 2026.

© 2026 Visible Language — this article is **open access**, published under the CC BY-NC-ND 4.0 license.

<https://visible-language.org/journal/>

Visible Language Consortium:

University of Leeds (UK)

University of Cincinnati (USA)

North Carolina State University (USA)

Implications for practice: The integration of automation and artificial intelligence within type design should serve to augment designers' creative and critical agency. Automated processes can effectively support technical and repetitive tasks—such as spacing, proofing, and data handling—allowing practitioners to concentrate on conceptual and aesthetic decision-making. Transparency and user control are central to this relationship; systems must remain interpretable and open to designer intervention. Ethical considerations warrant continued professional attention as automation becomes pervasive. Collaboration among designers, educators, and developers will be essential to ensure that emerging tools are aligned with the discipline's values of craft, intentionality, and typographic quality. Sustained engagement with scripting and AI literacy will empower designers to critically shape automated systems within their evolving practice.

Keywords: artificial intelligence; creative practice; ethics; intellectual property; type design

1. Introduction

1.1. Context and Aim of the Research

Artificial intelligence (AI) products and assistants are being inserted into every creative field, from writing text to programming to generating images and video. While there are currently many reservations about such systems, a broad spectrum of users have grown accustomed to these new tools. Until now, typeface design seems to be one of the few creative disciplines not yet penetrated by AI, for a few unique reasons.

Typeface design requires a particular workflow and has specific quality criteria that make it challenging for AI and machine-learning processes to achieve a sufficient standard. Letterforms must adhere to strict typographic conventions and optical adjustments to ensure readability and aesthetic harmony at minute details (Unger, 2018, p. 63), so one of the primary challenges in using generative AI for type design is achieving this high level of precision. Generating accurate and coherent typographic systems involves complex issues such as shape consistency across whole character sets, as well as spacing, kerning (which also requires building kerning groups), developing OpenType features, hinting, i.e., many crucial steps that are difficult for AI to master without developing extensive and nuanced training data.*

* Questions of “training” and the creation of “training data” will arise several times in this article. There are no standardized methods for creating training data from font files or type designs, nor for training algorithms with such data. Different decisions can be made at every stage, from the types of data included, perhaps only shapes, or also font metadata, to the sources from which the data is drawn and the way it is prepared. Possible sources include scanned typeface specimens, rasterized images, and vector outline data in a range of forms and formats. Such variation only adds to the opacity of the process and would logically lead to different outcomes.

For the Latin alphabet specifically, typeface design is an industry that is concerned with creating original work in increasingly minute variations of the same basic themes established over the last 500 years. It operates within the contradiction that a new typeface must look unique and distinctive, yet simultaneously so in line with history that “only very few recognize its novelty,” as said by Stanley Morison (1930/1951, p. 7). Little wonder then that the industry’s culture is full of possessiveness, where type designers and foundries are watchful of both being plagiarized as well as plagiarizing others (Heller, 2015; Monotype, 2025).

There are also far fewer people who make typefaces—or even make conscious choices about typefaces in a professional way—than there are people generating images, writing computer code or school essays using AI tools. With training costs reportedly in the billions of dollars for the top-ranking models such as Anthropic’s Claude or OpenAI’s ChatGPT (Buchholz, 2024; Henshall, 2024), the economic viability of a font-generating large language model (LLM) is questionable. This concern sets aside a further unresolved issue in current AI experiments in type design, namely whether a sufficiently large body of typefaces exists to provide adequate and viable training data.

Given the rapid advances in AI and the specific challenges posed by typeface design, it seems to be an appropriate time to reflect on the current and future impact of AI on the type design industry. A team of designers and educators from the Master Type Design at École cantonale d’art de Lausanne (ECAL / HES-SO)* came together to investigate the specificities of type design as a creative industry, and to better grasp the challenges and potentialities that AI brings to the community.

We set out to explore the following questions: How are current advancements in AI already affecting the type design industry? What are the perceived risks and opportunities according to type designers, font engineers and foundry owners? And how can we develop a collaborative model between designers and computer scientists that ensures ethical and qualitative use of AI in type design?

1.2. Automation and Typeface Design: A Shared Legacy

The desire to reduce human intervention in the production process is the primary driving force that led to typography’s departure from handwriting. Type designers, or designers at large, have always been practitioners of automated craft. (Wang, 2024, p. 7)

* The “Type Tomorrow” research project ran from January to June 2025 and was funded by HES-SO as part of its “Appel à projets stratégiques 2024: l’intelligence artificielle au service de la société: opportunités, défis et risques.” The research was led by Alice Savoie, Kai Bernau, Raphaela Haefliger and Wayne Daly (ECAL / HES-SO), with support from Sebastian Baez-Lugo (EPFL+ECAL Lab).

Use of computation to systematize or automate letterform generation predates AI by several decades. A pivotal early example is Donald Knuth's Metafont (1979), a parametric system that allowed users to define letterforms through mathematical descriptions of letter parts and their relationships (Knuth, 1979, 1985; Figure 1). Although revolutionary, Metafont was limited in visual quality and adoption, partly due to its steep learning curve and the difficulty of capturing typographic nuance in code.

The availability of digital tools opened up new perspectives for a young generation of designers (Poynor, 2003, p. 96), especially when they became readily available on consumer computers. In the late 1980s digital font editors such as Fontastic and Fontographer, in combination with the Macintosh computers they ran on, became essential tools for type designers, allowing for precise control over typefaces. These programs automated only very few parts of the process, such as certain aspects of shape drawing and storage of those drawings in the correct locations inside a font file. In the early 1990s, Petr van Blokland (with Just van Rossum and Erik van Blokland) produced a "hacked" version of Fontographer called RoboFog (Ulrich, 2022, p. 325). It included an application programming interface (API) in the then-new Python programming language, which allowed an extension of the software to automate far more tasks through scripting, laying important groundwork for the culture and the perceived tasks of the profession.



Figure 1. Variations on capital letter A using different parameters from Metafont. Originally published in Knuth (1985), p. 52.

Towards the end of the 1990s, the dominant font editor was Fontlab, which came with a Python interpreter already built-in, on top of which Erik van Blokland, Just van Rossum and Tal Leming built RoboFab, “a pythonic API to FontLab’s native objects” (FontParts, n.d.), heavily based on RoboFog (RoboFont, n.d.-a). A lively scene of individual type designers with a “DIY automation ethos” sprang to life from the possibilities provided by RoboFab, from small helper scripts to whole applications like MetricsMachine for kerning, or Superpolator for interpolation.

While interpolation between letter shapes as part of the type design process has been practiced since the 1970s, beginning at the firm URW with its Ikarus software (Ulrich, 2022, p. 275), the role of computation in type design expanded only in the late 1990s with technologies such as Adobe’s Multiple Master and Apple’s TrueTypeGX formats. These formats first allowed typeface *users*, not just typeface *designers*, to interpolate between extremes. Apple’s GX Variations formed the basis of the OpenType variable font format published in 2016. These systems still required human authorship of the base designs (Figure 2).

In the 2010s, new font editors such as Glyphs and RoboFont emerged, thanks to the dedication of a handful of type designers and computer programmers eager to develop contemporary tools. These two programs are diametrically opposed with regards to their view of integrating automation: Glyphs integrates many powerful functions and ways of optimizing and automating standard workflows (“Whose workflows?” goes one common critique of this approach) that offer much guidance and simplicity. RoboFont is a continuation of the lineage that began with RoboFog and RoboFab and abstains from bringing its own decisions to the user, instead requiring the user to think about and choose or create their own automations (RoboFont, n.d.-a). Famously, not even a tool to draw rectangles and ellipses is included with the core RoboFont application. Both font editors encourage, but RoboFont *requires* designers to add programming to their skill sets.

Communities that discuss and share plug-ins/extensions and scripts exist for both platforms as witnessed by public directories and dedicated extension managers (Robofont Mechanic and Glyphs Plugin Manager, respectively), and they are testimony to designers’ changing practices.

Contemporary type designers thus have at hand powerful tools that are seamlessly integrated into their design workflow. Many are embracing the use of programming and automation as part of their practice. Some designers with solid programming skills are marketing tools that can be easily integrated into a design process, such as Tim Ahrens’ Font Remix Tools (Remix Tools, n.d.) and Kern On (Kern On, n.d.) or Tal Leming’s MetricsMachine (Metrics Machine, n.d.).



Figure 2. Multiple Master axes implemented in the Myriad typeface family, designed by Robert Slimbach and Carol Twombly for Adobe Systems in 1992, and Minion, originally designed by Robert Slimbach for Adobe Systems in 1990. Originally published in Adobe Systems Inc. (1992, p. 7). Image: courtesy Musée de l'imprimerie et de la communication graphique Lyon.

1.3. AI, the Next Step in Type Design Workflow Automation?

The authors note that the term AI is often used rather loosely by themselves, the survey participants, the cited sources, and the broader audience. In this context, it may refer to systems that range from expert-operated specialized software trained on narrowly defined input data to produce narrowly defined outputs, to general large language models trained on vast internet-scale datasets that generate statistically probable responses to a wide range of queries. As many such systems are proprietary, it is not always possible to determine where a particular AI system lies along this spectrum.

A number of studies, articles and projects have explored the potential of AI for typography and typeface design, including experiments by computer scientists (Gao et al., 2008; Bataineh et al., 2012; Mohammed Javed et al., 2014; Murdock et al., 2015; Wang et al., 2015; Diem et al., 2017; Shinahara et al., 2019; Srivatsan et al., 2019; Wasim et al., 2024). Many of these experiments tackle isolated typographical aspects such as baseline detection (Murdock et al., 2015) or serif recognition (Wasim et al., 2024). Others attempt to generate entire alphabets based on fragmentary data—a practice that, if successful, opens up dizzying prospects for typeface generation. In 2018, Azadi and colleagues applied generative adversarial networks (GANs) to generate images of alphabets, presenting a model called “Multi-Content GAN” that can synthesize missing characters by learning from a few examples (Azadi et al., 2018). While their study demonstrated that GANs could produce stylistically coherent alphabets, a number of issues remained unsolved: the set of letters generated only included capital letters, featuring highly decorative styles unsuitable for continuous reading. Furthermore, the experiment was based on images rather than vector-based fonts. This is also the case for the Google Labs GenType experiment (Google Labs, n.d.; Figure 3), which generates illustrative alphabets through a user prompt.

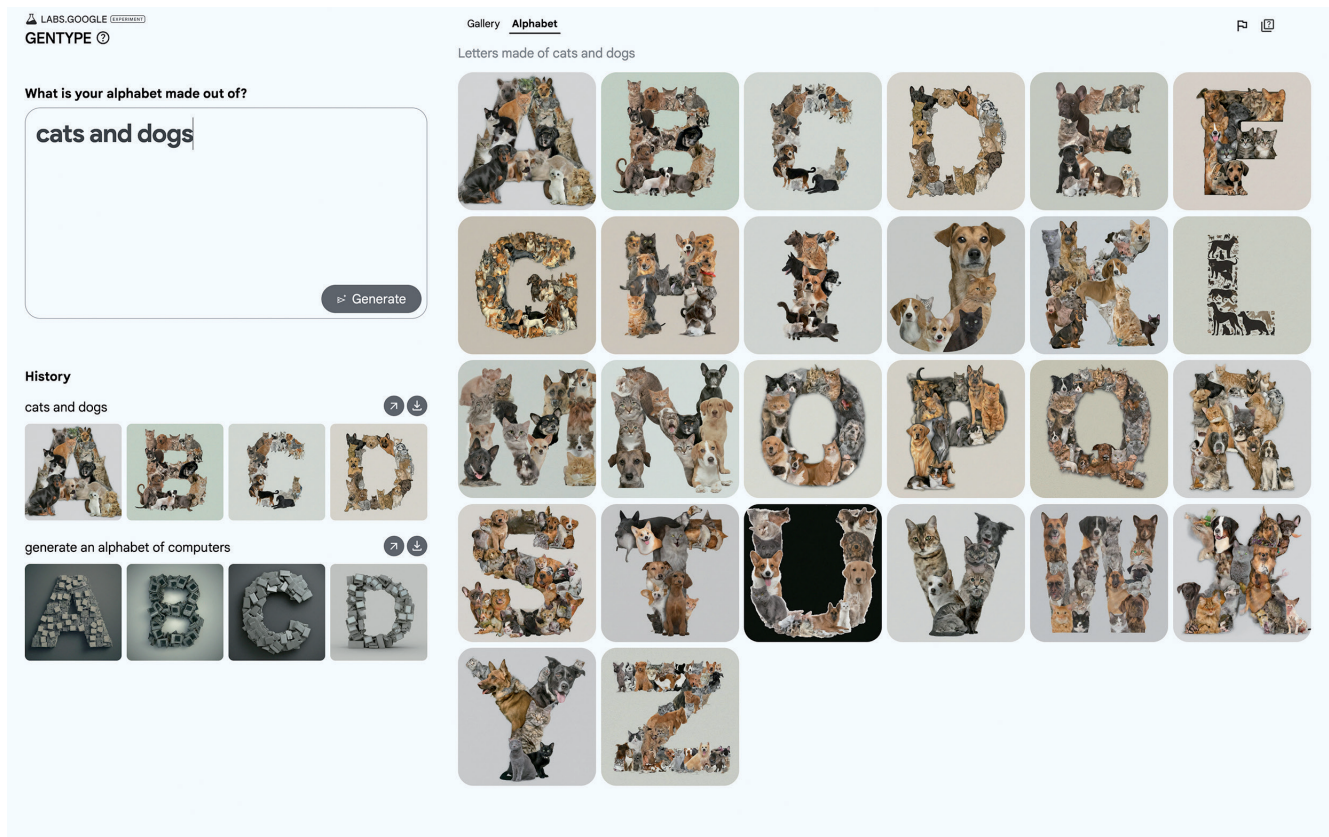


Figure 3. An alphabet made of cats and dogs, AI generated using the GenType model by Google Labs, October 2025. GenType is a trademark of Google LLC. Screenshot by the authors.



Figure 4. Letterforms generated with AI by Orlando Brunner. Published in Brunner (2024).

The last few years have seen initiatives which use generative AI and machine learning to output actual typefaces in vector format. Carlier et al. (2021) opened the way by generating and manipulating scalable vector graphics (SVGs) using deep learning techniques. Unlike raster images, vector graphics are defined by mathematical formulas, allowing them to scale without loss of quality. The team's research opened up new perspectives for generating character shapes, and potentially full typefaces. Initiatives to produce OpenType typefaces also recently emerged from independent type designers, programmers and type foundries (Brunner 2024; NaN Foundry, n.d.; Wentzel 2024; Figure 4). While these experiments are worthy of consideration, the resulting typefaces remain rudimentary in their design, often encompassing only capital letters, and displaying obvious flaws. The recently beta-launched Typograph service promises to “generate a typeface in seconds” from verbal prompts (Typograph, n.d.). However, co-founder Viktor Persson explains that the typefaces are actually created by more conventional means, with a chat-AI based interface sitting between the user input, and a variety of typeface models that are blended and interpolated conventionally, through

the interpretation of the chat prompts.* In recent years, programmers and typeface designers have joined forces and experimented with AI-assisted kerning.†

Most of these initiatives have raised fierce debates within the industry regarding the dataset used for training the algorithms, and type foundries have started updating their licenses to protect their designs from being exploited (Lineto, 2024). In theory such license addenda would on the one hand protect designers' creations from being incorporated into learning data and closely imitated by a theoretical AI font creation software; on the other hand, it would limit the quality of the output dramatically. In line with the wider AI industry, companies might just train their models on copyrighted and/or pirated data anyway, and claim that "it would be impossible to train today's leading AI models without using copyrighted materials," as OpenAI responded to an inquiry of the U.K. House of Lords' Communications and Digital Select Committee in relation to a copyright lawsuit brought against OpenAI by the *New York Times* (LLM0113, 2023).

It remains unclear whether relying only on Open Source typefaces can create quality results. Training outcomes depend not only on the amount of available data, but also on the quality and consistency of that data, as well as on the methods used to translate training material (such as outline descriptions or pixel images) into generated outline structures. These processes involve numerous technical choices that can strongly influence the quality of the output. Nonetheless, access to large and carefully curated typeface collections may still confer an advantage, which would mean that type foundries with the largest intellectual property (IP) resources are best positioned to assemble high-quality training datasets.

Beyond its own legacy library, Monotype Inc. has been buying type foundries and catalogs at a quick pace over the last 20 years,‡ and now claims to own over 150,000 fonts, completely dominating the market in many aspects including its ability to create AI training sets with this font data. At the moment, Monotype's efforts in the field of AI seem to be more in the field of categorization and recommendations of existing

* Viktor Persson interviewed by Kai Bernau, June 23, 2025.

† Some of these experiments were discussed in interviews carried out as part of this research project: see Cem Eskinazi interviewed by Raphaela Haefliger, June 26, 2025; Tal Leming interviewed by Kai Bernau, May 30, 2025.

‡ Notable acquisitions include: Linotype (which itself had previously absorbed font libraries from the likes of Deberny & Peignot, Haas and Stempel), 2006; China Type Limited, 2006; Ascender Corp., 2010; Bitstream and MyFonts, 2012; FontShop and the FontFont Library, 2014; URW Type Foundry, 2020; FontSmith, 2020; Hoefler & Co., 2021; Berthold type foundry, 2022; Milieu Grotesque foundry, 2023; Fontworks, 2023; 39 typeface families by David Berlow from the Font Bureau Library, 2023; Colophon Foundry, 2023; the rights to the catalog of SharpType, 2024; DSType font collection, 2024 (Monotype, n.d.-d; Luk, 2023; Fonts In Use, n.d.).

fonts (Monotype, n.d.-a) as well as assistants—a “creative partner, working alongside [designers] to add to [their] human creative capacities”—but also adds “Expect more and expect it soon” (Monotype, 2025). In 2025, Monotype announced a partnership with Blaze Type to conduct experiments in AI-automated completion of character sets (Monotype, n.d.-b), with improved formal results, in vector format. On the page describing the experiment, no mention is made as to the training data, but in an interview, Monotype’s Emilios Theofanous describes “a tool for typographic AI generation” where users provide “input of specific characters and [...] get back either an expanded character set [...] or alternates.” He notes that at the moment the tool is “only used internally,” “only for exploration and experimentation” and that for training data it uses “just Monotype’s IP,” not that from Monotype’s partners.*

2. Methodology

2.1. Aim and Context of the Survey

In response to this context, our team decided to run a survey to better understand how professionals—type designers, font developers, font engineers, etc.—perceive the evolving role of AI in type design. We also wanted to identify what parts of their process designers are willing to delegate to a machine, and which parts are considered as meaningful, enjoyable, or highly dependent on human intervention. The survey further shed light on the perceived risks and opportunities associated with AI in type design. We felt that understanding these perceptions was crucial for addressing concerns and identifying areas where AI can be most beneficial, or potentially harmful.

The survey was conducted through the online platform Survey King and was publicly accessible for six weeks. Respondents’ participation was entirely voluntary and responses remained anonymous, analyzed in aggregate form only.† The study was carried out in accordance with our institution’s research procedures and ethical guidelines, with internal approval secured before data collection commenced. To maximize ease of use for participants, all individual questions were optional, resulting in variable sample sizes for each analysis. When quoting open responses from the questionnaire in the

* Emilios Theofanous interviewed by Alice Savoie, July 7, 2025.

† The survey was conducted between February 18 and March 31, 2025 and was promoted through social media (posts on Instagram and Mastodon, through ECAL-affiliated accounts as well as personal accounts of team members, friends, and colleagues who kindly shared or republished the survey), a conference (an announcement during the Automatic Type Design 3 conference in Nancy) and email (via the ECAL newsletter system, targeting industry members, ECAL alumni, and others). Respondents had the option to provide their email address in the final stage of the survey, as a means to be kept informed about the publication of survey results. Recipients were also encouraged to share the survey with their own contacts.

sections that follow, we reproduce them verbatim and attribute an anonymized identifier to each respondent (e.g. [P35]). From a total sample of $N=157$ participants, $n=123$ completed the questionnaire, and $n=34$ submitted only partial responses.* Because no comprehensive demographic study of the type industry currently exists, it is not possible to establish how closely our sample reflects (or diverges from) the broader professional population.

Once the survey data had been analyzed, our team conducted a series of interviews to clarify and expand on some of the findings. Practitioners who collectively spanned a range of roles, geographic regions and levels of experience were invited to take part, selected purposely to discuss specific experiments referenced by survey respondents and/or to draw on their knowledge of ongoing automation and AI initiatives within their workplaces. Semi-structured interviews were either carried out through online meetings (video conference) or via e-mail. We asked interviewees to share their own views and recent experiences with automation and the use of AI in the type design process, and, where relevant, to reflect on their own or their employer's experiments. Interviewees include: Matthew Carter (founder, Carter & Cone), Cem Eskinazi (type designer, independent), Tal Leming (type designer and coder, Type Supply), Viktor Persson (founder, Typograph), Keitaro Sakamoto (type director, Morisawa), Emilios Theofanous (type director, Monotype). All conversations were transcribed, and illustrative quotations are woven into the subsequent discussion; interview quotations are not anonymized and are reproduced with the interviewees' informed consent.

2.2. Survey Structure and Analysis

Our survey was comprised of four sections:

1. **Introductory and ethical information:** respondents first encountered a brief statement outlining the study's aims, the intended use of the data, and the measures taken to ensure anonymity. They were then asked to confirm their informed consent before proceeding.
2. **Current practice and expectations:** participants detailed how they currently employ automation in their workflows, distinguishing between deterministic (i.e., rule-based) and non-deterministic (i.e., AI-driven) methods. A five-point Likert scale followed, gauging the extent to which they believe AI will reshape the type industry in the near future.
3. **Propensity to automate specific tasks:** using a five-point Likert scale, respondents indicated the likelihood that they would delegate each of 16 typical tasks involved in a type design process, from early sketching through final quality assurance, to

* For these 34 partial submissions, only their answers to the open-ended questions were considered, using them for qualitative sentiment analysis.

(a) deterministic scripts and (b) AI-powered tools. Open questions invited them to elaborate on which tasks they find especially tedious or enjoyable, and why they would—or would not—automate them.

4. Demographic and professional profile: the final section gathered background information: age, gender, country of residence, profession(s) and employment situation, years of industry experience, and the writing systems in which respondents routinely work. An optional comment box invited participants to add any further observations before submitting the questionnaire.

Survey analysis. The quantitative data from this online survey was analyzed using R (version 4.4.2) within the RStudio environment. Data wrangling and descriptive statistics were performed using packages from the *tidyverse* suite, primarily *dplyr* for data manipulation and *ggplot2* for visualization. Where appropriate, inferential statistical analyses were also conducted, including non-parametric pairwise comparisons, correlations, and linear regression. Qualitative data from the open-ended survey questions and follow-up interviews was analyzed using a thematic analysis approach. This process was facilitated by a rainbow spreadsheet to organize codes and identify emergent themes.

2.3. Profile of Respondents

Age and industry experience. 77% of participants fall within the 25 to 44 age bracket, with only one respondent aged over 65 and two under 24 (Figures 5 and 6). This skew almost certainly reflects the cohort most active in contemporary type design, enhancing the survey’s relevance to current workflows. Even so, the relative absence of older practitioners should be taken into account when interpreting the results.




Figure 5. Survey participants profile: age category distribution ($n=120$).




Figure 6. Survey participants profile: industry experience ($n=122$).

Professional experience is more evenly distributed. Just under a quarter of respondents (23%) have worked in the field for fewer than five years; 28.7% for five to ten years; 32.8% for ten to 20 years; and 15.6% for more than two decades.

Gender distribution. There is a clear gender imbalance in respondents’ profiles, as male respondents (61.5%) outnumber female respondents (29.5%) by approximately 2:1, with a small proportion of non-binary or undisclosed respondents (Figure 7). It is likely that these figures reflect long-observed industry imbalances.

Geographic distribution and scripts. A substantial majority of participants (nearly 80%) were based in either Europe or the United States, with a high concentration of respondents living in Switzerland and France (Figures 8 and 9). Although these figures cannot be taken as representative of the global type design community, this distribution reflects the research team’s institutional location and outreach channels.

Furthermore, nearly every participant reported designing for the Latin script (99.2%). Substantial minorities also worked with Cyrillic (36.4%), Greek (26.4%) and Arabic (11.6%), with smaller shares for Hebrew (5.8%). Additional scripts including Thai, Tamil (each 5%), Tai Viet and Chinese (each 4.1%) were represented by a handful of specialists. These figures highlight the heavy representation of industry professionals working in the West and on the Latin script, something that should be considered when interpreting the results, as findings may therefore not extrapolate easily to under-sampled regions.

Professional activities of respondents. 79% of respondents identified as typeface designers, but with multiple answers allowed, a substantial minority also claimed




Figure 7. Survey participants profile: gender distribution ($n=122$).




Figure 8. Survey participants profile: country distribution ($n=107$).




Figure 9. Survey participants profile: script distribution (multiple answers allowed, $n=121$).

technical titles such as “font developer” (34.7%) and “font engineer” (25.8%; Figures 10 and 11). Respondents also include 37.1% who identify as “Graphic designers who make fonts.” Another set of responses shows that 52% identify as freelancers and 37.4% as business owners, with only 12.2% of participants claiming to work for a medium or large company of 20 people or more. These figures underscore the fact that the community includes a large number of small studios and independent practitioners. It also highlights the fact that respondents’ jobs might include a diverse range of skills, and that creative and engineering roles frequently overlap.

3. Survey Findings

3.1. Current Use of Deterministic Automation in Typeface Design

Survey responses confirm that deterministic (rule-based or non-AI) automation has become an integral component of everyday practice. Three quarters of respondents (76.4%) reported employing deterministic tools in their workflow, most commonly via Python scripts or plug-ins for font editors such as Glyphs and RoboFont. Usage rates appear consistent (79–87%) across every age bracket. Whereas a higher percent-

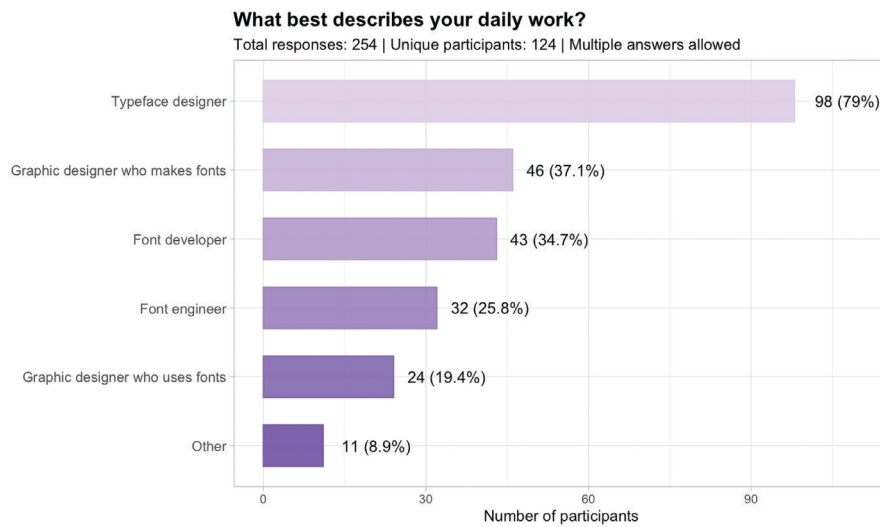


Figure 10. Survey participants profile: professional activities (multiple answers allowed, $n=124$).




Figure 11. Survey participants profile: position in industry (multiple answers allowed, $n=123$).

tage of male participants (89%) reported using non-AI automation compared to female participants (69%), professional identity (designer, developer, engineer, etc.) does not predict willingness to automate. As a complement to these figures, pairwise comparisons (Mann-Whitney U tests) revealed no statistically significant difference in automation likelihood between genders for either the AI condition or the non-AI condition. Predictive relationships between professional roles and willingness to automate were assessed with regression analyses. Perspectives on automation benefits and drawbacks are commonly experienced by all individuals, and are not influenced by their gender identity.

A minority of respondents (approximately one quarter) reported minimal or no use of automation, citing a preference for “hand-drawn” letterforms, skepticism about tool accuracy, or a desire to understand every step of production. Respondents’ comments suggest that, for many, the boundary between “manual” and “automated” work is now blurred as it could be argued that by default, anyone using a font editor and dealing with digital type relies on some kind of automation, one way or another. Deterministic automation has become foundational to digital type design, even if barely visible to some. Or as one participant asked, “Does anyone do type design without these kinds of tools?” (P134).

Deterministic automation is enlisted across a broad spectrum of tasks, including:

- ▶ Outline cleanup and error detection: scripts and applications for detecting drawing inconsistencies and issues in point placement
- ▶ Glyph scaling, glyphset expansion: designing small caps, superiors, inferiors and other derivative forms
- ▶ Interpolation and design space exploration: generating intermediate styles, supplementary weights and widths, optical sizes, etc.
- ▶ Spacing and kerning: comparing tables across masters, building kerning classes, suggesting pair values
- ▶ OpenType feature code writing
- ▶ Quality assurance: proofing, testing of exported fonts, etc.

In relation to these tasks, the survey identifies a number of utilities that have become de facto industry standards, as shown in Table 1.

Table 1. Table listing the applications and plug-ins that survey respondents reported using for specific tasks in the type design process.

Purpose	Tools and mentions by respondents
Outline clean-up and error detection	Speedpunk (3); Red Arrows (3); GlyphNanny (2)
Glyph scaling, glyphset expansion	RMX suite of tools (62); UFO Stretch (4)
Interpolation and design space explorations	Prepolator (21); Variable Font Preview (2)
Spacing and kerning	Kern On (30); Metrics Machine (5); HT LetterSpacer (4); iKern (2); LS Cadencer (1); DTL Foundry-Master (1)
Proofing and specimen generation	DrawBot (4); Word-O-Mat (1)
QA, mastering, writing OpenType features	FontTools (3); Diffenator 2 (2); FontBakery (2); Glyphs (1)

Respondents praised the RMX suite of tools in particular as a “quick sketching” environment for testing weights and widths, defining design space boundaries and producing scaled variants of base glyphs (such as superiors and inferiors or small caps) that are subsequently refined by hand. Likewise, Kern On is widely adopted to establish a first kerning pass, although many practitioners emphasized the need for manual oversight: “I always kern the basics manually and respect those values while the auto-kerning handles the rest” (P101). Prepolator is also frequently mentioned as an effective tool for cleaning up master styles prior to interpolation.

A further theme to emerge from the survey, and one that appears distinctive to the type sector, is that the majority of the software and apps mentioned by respondents have been conceived and coded by practitioners themselves. The principal font editors (Glyphs and RoboFont) as well as widely-adopted tools such as RMX Tools, Kern On, Metrics Machine and Prepolator, were all conceived, coded and released by practicing typeface designers with a fluency in software development. This dual identity collapses the conventional separation at play in most allied industries, nurturing a collegial culture in which design problems are addressed by writing and sharing code.

While deterministic automation is welcome for “boring and repetitive tasks,” many respondents are keen to retain creative control. Several interviewees stressed that automation should remain subservient to judgement: “While these automation tools are helpful for providing a rough visualization, I ultimately rely on manual correction to ensure accuracy and alignment with my standards” (P38). Others drew attention to the pragmatic benefits of scripting: “A simple rule was once explained by Petr van Blokland*: ‘if you have to do the same action more than ten times, then script it.’” (P224). Custom scripts enable type designers to tailor workflows to specific projects, further establishing automation as a craft resource: “I automate repetitive tasks as much as I can. I also use it to avoid human-made errors. Batch editing a large number of files or operating in a big design space can sometimes be challenging. I use custom tools to solve those problems” (P251).

3.2. Current Use of AI in the Type Design Industry

In contrast with the near-ubiquitous presence of deterministic scripting, AI tools remain a minority pursuit among professional type designers, irrespective of age. When asked “Do you currently use AI-powered tools in your type design work?” just under one third of respondents (31%) replied in the affirmative, and barely one in five (19%) reported having witnessed a “convincing” deployment of AI for type design tasks.

* Petr van Blokland is a Dutch graphic and type designer and software developer. He teaches in the Master Type and Media at the Royal Academy of Arts in The Hague.

Where AI is employed, it tends to support activities that are adjacent, but not central, to the drawing and engineering of typefaces, such as:

- ▶ Documentation and preliminary research
- ▶ Copywriting and translation assistance
- ▶ Writing or debugging code (e.g. Python scripts)
- ▶ Generating proofing documents, text samples and letter strings
- ▶ Studio administration and client communication
- ▶ Image generation and image processing

This pattern suggests a complementary, rather than competitive, relationship between AI and deterministic automation. Leading the way, OpenAI's ChatGPT dominates current usage (26 mentions), followed by LLM alternatives such as Claude.ai and DeepSeek. For code writing and debugging, respondents mentioned GitHub Copilot and Cursor, whereas tasks involving image generation prompted references to DALL·E, Hugging Face and Adobe Photoshop's generative-AI features. Overall, the data portrays AI as an emergent "studio assistant" that is valuable for ancillary work, but not yet embedded in a core type design workflow.

Of course this relationship is almost certainly shaped by current tool availability, since no commercially accessible AI applications address the critical stages of designing, developing or engineering typefaces. The handful of AI tools cited by respondents were at an exploratory stage, with mentions of initiatives specific to the development of Chinese, Japanese, and Korean (CJK) typefaces, experiments with AI-powered kerning tools, and the use of LLMs for generating typefaces. In the absence of market-ready specialist applications, machine-learning solutions remain confined to peripheral tasks, leaving deterministic automation unchallenged at the core of professional practice.

The prevailing sentiment across all participants' age groups (under 35, 35–45, and over 45) is that AI will likely transform the type design industry. When combining the "very likely" and "extremely likely" responses, a strong belief in transformation is evident (33%, 49% and 44%, respectively). Conversely, the proportion of those who felt a transformation was unlikely was consistently low (14%, 14%, and 12%, respectively). Despite these variations, a Kruskal-Wallis test showed no statistically significant difference in sentiment across the three age groups. This result indicates a widespread and consistent belief in AI's future impact, as professionals from all generations agree that the industry is more likely than not to be transformed by this technology. Whether a professional identifies as a typeface designer, a graphic designer, a font engineer, or a font developer, they all share the same fundamental perspective on automation tools.

3.3. Skepticism Expressed Towards AI: Ethical Issues and Corporate Dominance

When addressing the possible use of AI for core type design tasks, qualitative answers paint a picture of guarded skepticism with respondents commonly describing current AI tools as “imprecise,” “unpredictable” or “not ready for production.” Respondents further voiced three recurring concerns that echo wider debates in the creative industries:

- ▶ Intellectual property issues: the datasets used to train LLMs are likely to contain copyrighted typefaces used without permission from the original designers or type foundries.
- ▶ Energy consumption and carbon cost: the environmental footprint of large-scale inference was frequently cited as a deterrent.
- ▶ Opacity: designers mistrust what are seen as “black box” systems whose decision-making logic cannot be tracked or corrected. “Coding is the one area where I can see myself using AI automation, to help in building tools of my own design. I use a lot of non-AI automation in my workflow, but I can trust its output as I know what code I wrote and it is not going to hallucinate answers” (P263).

These (legitimate) concerns partly explain the community’s preference for deterministic tools, which offer much greater transparency at this stage, and require fewer resources.

A small but vocal group of respondents warn that the next wave of AI could entrench the power of corporations that already own extensive libraries of licensed fonts. They fear that companies such as Monotype or Google will “acquire the legal right to train an AI model on their library’s font data” and then offer “fully bespoke generated fonts at a fraction of the cost and time of a human designer” (P58). One respondent anticipates that the technology will simply magnify existing market imbalances: “It will be used by huge companies ... that are already changing (harming) the industry in many ways. Only the means will change” (P304).

Beyond outright market domination, respondents see a subtler threat to the values of their craft, with eight participants predicting that profit and speed will eclipse quality and originality once AI-driven generative tools become mainstream: “People will choose convenience over quality” (P30). One respondent expects that cost savings will prove “too hard to turn down for clients” even if the results are mediocre (P58). Several foresee a deluge of derivative work: “Someone will make a model that generates typefaces ... trained on fonts of poor quality and ethically dubious originality. Most people won’t know the difference and will pay £5 for a ‘custom font’” (P133). The likely outcome is that “mediocre work [will] pay even less” (P151), making it harder for independent designers to sustain their practice.

In that scenario, craft itself may become a mark of distinction. As one participant suggests, “There is always a connotation of premium or exclusivity when something is hand-made ... perhaps there will be such a distinction for digital works as well” (P93).

3.4. Retaining Creative Control

A clear hierarchy emerges from the Likert scale questions concerned with respondents’ propensity to automate specific tasks. Tasks perceived as being key to the creative process attract the strongest resistance to automation, whether AI-based or not (Figure 13). These tasks include:

- ▶ Gathering references
- ▶ Sketching ideas
- ▶ Drafting the basic alphabet
- ▶ Extending glyphsets
- ▶ Marking up proofs

Only 22% of respondents state they are “very likely” or “extremely likely” to entrust the drafting of the basic alphabet to deterministic tools, and a mere 9% to AI-powered tools. Qualitative comments reinforce this pattern, with many respondents emphasizing the pleasure they derive from drawing (35 mentions), sketching (26), and exploratory design work (23). The respondents explicitly identified “the design process” or “creative part” as the heart of their practice, with six framing this territory as “decision-making” that must not give way to an algorithm. As one respondent remarked: “I started drawing type with a fascination for its formal qualities and craft-related aspects and skills. My personal interest in type design is about exploring one’s mind and its connection to the hand and eye. I have no real interest in having something else make decisions for me within this process” (P205).

Furthermore, five respondents praised the inherent slowness of type design, including:

The slowness of type design is also what makes it enjoyable, therapeutic and gives a stronger sense of satisfaction once you reach a result you like (P31).

Type design is a slow process. In my professional practice, I try to incorporate aspects of it in every project, whether a couple of letters for an identity or an entire alphabet. It’s something I particularly enjoy and which gives sense to my entire practice. I don’t think the world is in need of another thousand fonts and those produced should be carefully crafted and thought through (P56).

Saving time is not key because the maturation of a font takes time—for good ones at least (P63).

In line with the above, respondents frequently highlighted the early stages of defining a typeface's identity as particularly enjoyable: finding references and gathering inspiration (4 mentions), defining the concept (6), mapping the family or design space (6), designing the core alphabet or basic glyphset (13), designing the extended or entire glyphset (7), applying optical corrections or refining shapes (6).

Interestingly, spacing was mentioned 13 times as an enjoyable task that should not be delegated to AI, indicating that respondents drew a clear distinction between spacing and kerning. Six respondents specifically named kerning as an enjoyable task that they would not delegate to AI, although, as we shall see below, this view is not reflected in the quantitative data.

Although most practitioners reserve the early, exploratory phase to human work, a small cohort of respondents (seven) view AI as a catalyst for creative ideation, an instrument for “achieving results we could not imagine ourselves” (P14). They endorse AI systems as rapid, low-stakes sketching partners, useful “for testing ... a weight or width for a new style” (P8), or “to explore different directions faster” (P145). Several respondents describe using machine output as raw material that can be subsequently refined: one appreciates AI's capacity “to explore the project boundaries and to produce shapes that can be re-incorporated ... after some edits to the ‘proper’ design” (P223). Others regard it as a provocation, for instance to feed LLMs with disparate historical samples “just to see how it [is] digested ... before starting a new drawing” (P208). One respondent notes that such tools let them “create things [they] would not otherwise create” (P279). Together, these voices suggest that, in limited hands, AI can function less as an autonomous designer and more as a brainstorming companion.

3.5. Potential for Automating Technical Tasks Using AI

By contrast, enthusiasm for automation grows steadily as tasks become more technical or repetitive (Figure 12), such as:

- ▶ Optimizing font rendering
- ▶ Generating proofing documents
- ▶ Kerning
- ▶ Programming
- ▶ Writing OpenType features
- ▶ Setting vertical metrics and font info

These results resonate with respondents' qualitative responses, which highlight throughout the survey their willingness to remain in control of key creative labor and decisions. Many use or wish for tools that streamline these routine jobs to speed up the process, while requiring human control and adjustments. Automation is seen as a way to free designers from “boring work,” allowing them to focus on the artistic side.

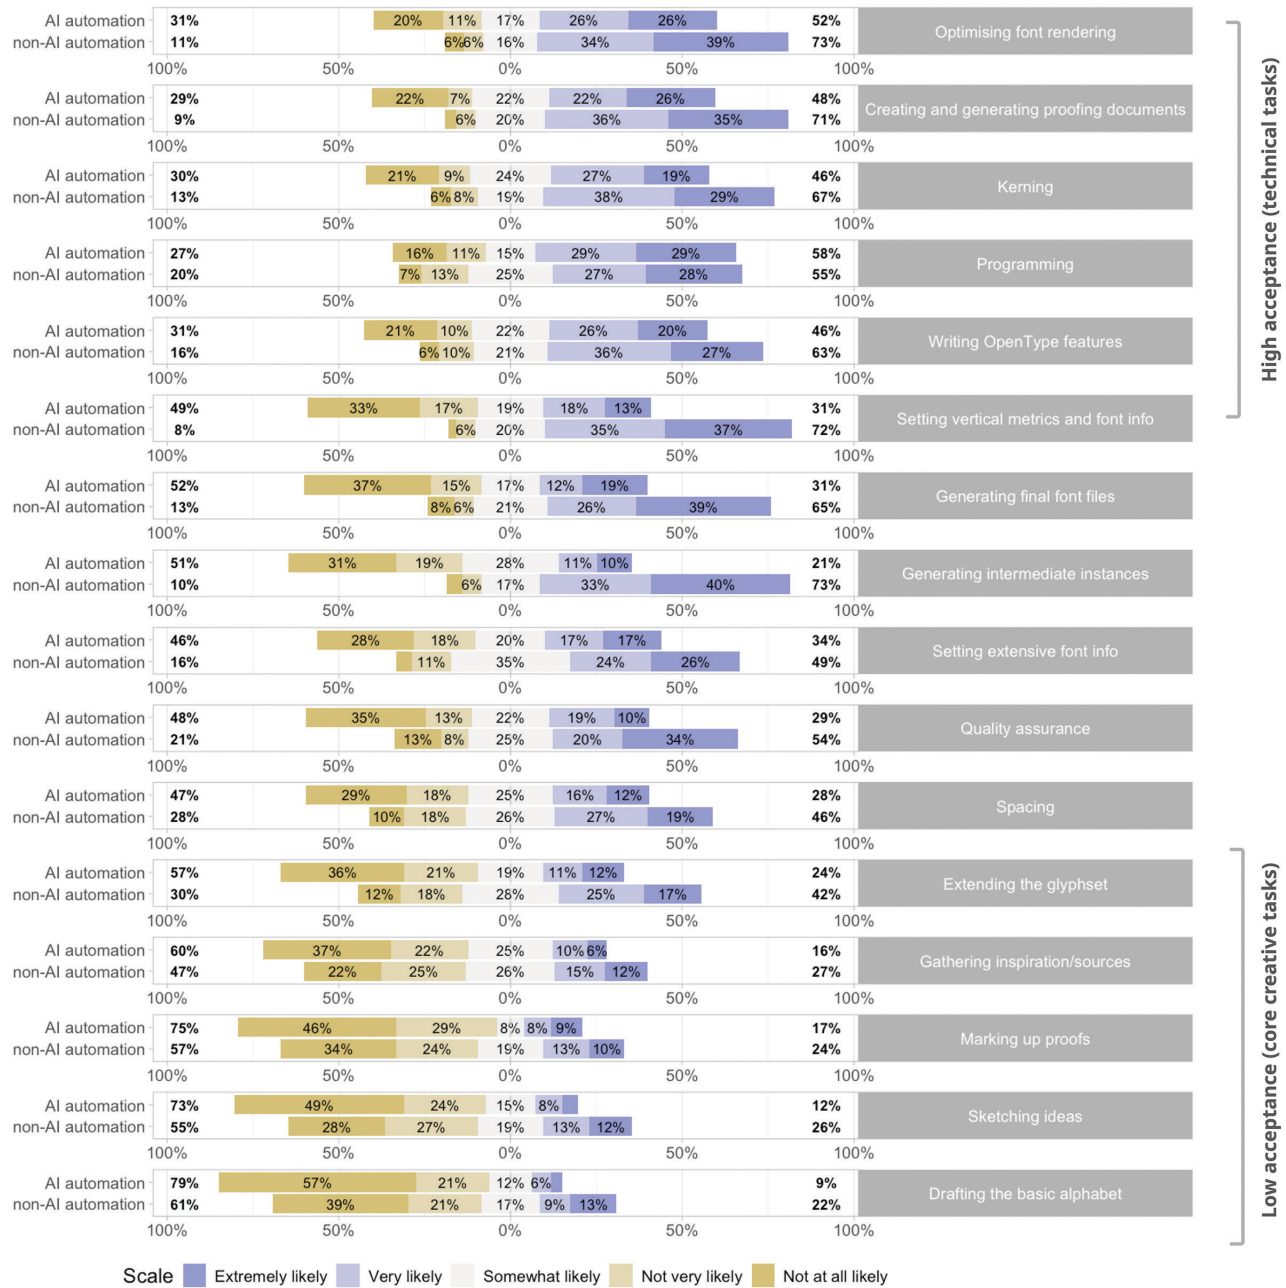


Figure 12. Likert scale on propensity to automate specific tasks using AI versus non-AI (deterministic) tools. Tasks are ranked from highest to lowest overall acceptance, separating high acceptance for technical tasks from low acceptance for creative tasks ($n=89$).

It's important that we look at AI and automation as assistants to a designer and not the replacement of the designer. For the sake of our industry, AI should not be the answer to lack of creativity but rather as a way to improve efficiency (P108).

I hope and believe that AI will make type design easier and largely eradicate tedious tasks such as setting vertical metrics and producing spacing/kerning. The type designer will still however rely on a good eye to quality assess the machine's quality assessments (P31).

I hope it will help font developers to make repetitive work quicker (see kerning, accents, etc.) but not affect the drawing part too much, where, in my opinion, it should play only an assistive role. My dream is to use AI in the testing environment, such as proofing and checking, where sometimes you can easily make mistakes just because you're too tired to look again at the same thing for so long (P218).

For tasks traditionally regarded as repetitive rather than creative, respondents expressed markedly greater confidence in deterministic automation than in AI-driven solutions. For example:

- ▶ Setting vertical metrics and font info tables: 72% of respondents were “very” or “extremely” likely to use deterministic tools, compared with 31% who would choose AI.
- ▶ Generating intermediate instances: 73% were “very” or “extremely” likely to rely on deterministic tools, versus 21% for AI.
- ▶ Spacing: only 10% were “not at all likely” to automate spacing with deterministic tools, in contrast to 29% for AI.
- ▶ Entering extensive font info data: 3% were “not at all likely” to automate this task via deterministic tools, in contrast to 28% for AI.
- ▶ Quality assurance: 54% were “very” or “extremely” likely to use deterministic tools, compared with 29% for AI.

Spearman correlation analyses indicate that attitudes toward AI automation are strongly inter-correlated, reflecting an “all-or-nothing” approach, suggesting that those receptive to AI tend to be receptive across the board. In contrast, attitudes toward deterministic automation tools are compartmentalized and task-specific, indicating a more pragmatic “task-by-task” adoption. Designers' longstanding familiarity with deterministic scripting likely underpins this asymmetry. AI tools, by contrast, are not yet widely integrated into everyday workflows and consequently attract greater caution.

3.6. Kerning as a Candidate for AI Automation

Kerning emerged from the survey as the foremost candidate for further automation, with 44 respondents identifying it as the task they would most readily delegate to AI (provided the output met satisfying quality standards). This preference likely reflects current practice: deterministic tools such as Kern On (Figure 13), MetricsMachine and the third party service iKern (iKern, n.d.) already handle much of the kerning workload for many foundries and type designers. The perceived benefits are foremost efficiency and consistency with one respondent observing that kerning “is quite tedious and time-consuming, and even when I do it myself, I don’t feel totally confident about the result, so I would gladly trust an AI to do a better job” (P24). Another described it as “repetitive, error prone” (P133), implying that assistance can raise, rather than lower, consistency and overall quality standard.

Nevertheless, this enthusiasm is tempered by a desire to retain final control, with numerous respondents calling for a hybrid solution in which the machine suggests kerning values, while the designer approves them:

Yes, I believe AI can easily handle kerning and production, but I would always double-check and never publish a font without ensuring every detail aligns with my intentions (P38).

Some argued that kerning serves a secondary function as a systematic review of a completed typeface. Interviewees Cem Eskinazi and Emilios Theofanous both stressed that the kerning stage acts as “a good QA moment”^{*} and “an opportunity to run the final checks,”[†] underlining its role in the broader cycle of quality assurance, rather than merely spacing adjustment.

Two ongoing experiments with AI-powered kerning exemplify this assistant model. The MILK prototype, developed by the independent collective type.tools (Type.tools, 2020), and an experimental plug-in by Tal Leming (with assistance from computer scientist Lars van Blokland), both employ machine-learning to suggest kerning pairs. Leming emphasizes that the aim is to build “an assistant, not a black box,” noting that kerning is a subjective process that ultimately embodies “opinion” rather than immutable rule. From a technical standpoint, both teams report that the core modelling challenges are largely resolved, but some crucial outstanding issues remain that relate in large part to economics and ethics: the commercial market for professional kerning tools is modest, and few industry actors can absorb the costs of bringing such software to

^{*} Cem Eskinazi interview.

[†] Emilios Theofanous interview.

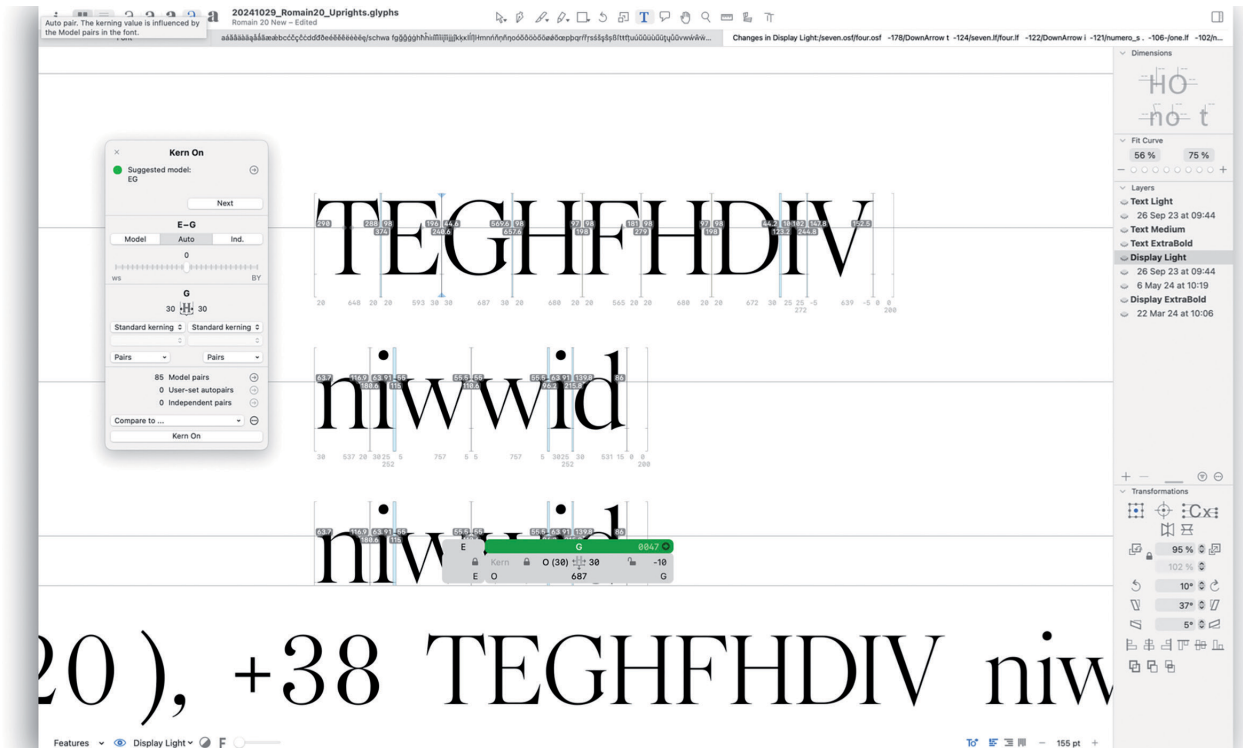


Figure 13. Screenshot of the Kern-On plug-in for partially-automated kerning of a typeface in the Glyphs font editor. Screenshot by the authors.

market; moreover, the energy required risks contradicting the efficiency gains that such an automation tool promises.

Taken together, these findings suggest that kerning should be considered not solely as a mechanical operation, but rather as an integral component of the design process. Designers seem to be willing to accept assistance as long as it optimizes their time, remains transparent about the process, and preserves their judgment and responsibility. The specific case of kerning therefore illustrates the broader argument articulated in our study, which is that automation is much better accepted when it supports and augments (rather than displaces) the designer’s decision-making process.

3.7. Learning by Doing: Acquiring Skills Through Repetition

Automation bears with it the promise to relieve designers of time-consuming chores, but its indiscriminate adoption also risks affecting the very routines through which expertise is acquired. One respondent to our survey articulated this dilemma:

...only doing something repeatedly makes you good at it and allows you to create novel and truly innovative things. Without that learning, it is difficult to imagine something amazing being developed. But if AI does all the generic work, how can anybody learn the skills that make a good designer? (P116)

Another linked the problem to tool making itself:

You can only write the tools you need if you fully understand what they need to do... Delegating things to AI is a magical thinking: “I cannot be bothered to figure this out, but I demand that it be done anyway” (P97).

Similar concerns have arisen in other fields; for instance, a recent article in *Le Monde* reported on French law firms, observing that LLMs already outperform junior solicitors for some administrative tasks and documentary research (Thomas, 2025). Yet it is precisely through this foundational, day-to-day work that young lawyers acquire the analytical discipline and knowledge on which their future careers depend.

Likewise, the cognitive process involved in designing typefaces appears to depend on iteration. Richard Sennett captures this point when he asks whether the machine is “a friendly tool or an enemy replacing the work of the human hand,” concluding that “making is thinking” (Sennett, 2008, pp. 65, 81). Type designer Gerard Unger echoed this sentiment from within the discipline, arguing in 1982 that he preferred tools “that *help* [him] think rather than those that *make* [him] think” and refusing to become “a parameterizer” (p. 354).

Sennett (2008) argues that “skill development depends on how repetition is organized” (p. 38), something that resonates with the type design activity. Each cycle of spacing, proofing or curve-correction feeds tacit knowledge back into the next design decision. If you remove this cycle, the feedback loop collapses. Some form of repetition is therefore necessary, whether executed with a pen, a Bézier handle or an AI assistant. Even tasks that might appear repetitive and boring may eventually carry satisfactions of their own—something described by British book designer David Pearson (2024) as “the pleasure of gently nibbling away at something” through slow, repetitive labor.

Tal Leming voices a related concern when reflecting on his own software:

Young designers at Font Bureau* used to get typefaces ready for interpolation... I wonder, did I fundamentally break something [when I wrote Prepolator]? I don't think that I did, I think economics played a bigger role than me figuring out how to calculate something. Kerning is a task given to assistants, who are still learning, in many foundries. Am I replacing it? I have deep ethical concerns about it.

* FontBureau was an influential early digital type foundry set up by type designer David Berlow (previously Mergenthaler, Linotype, Bitstream) and publication designer Roger Black in Boston in 1989. Its work influenced the industry through the designs it produced and the designers that came through and out of the firm (Fonts In Use, 2025).

The challenge therefore might lie in designing systems that support this learning process rather than supersede it, and in devising tools that invite experimentation and allow practitioners to grow from mechanical tasks to refined judgement. While automation may accelerate production and mask complexity at the same time, it could risk creating a generation of “parameterizers” (or one could say here, of “prompters”), removed from the skills that give the profession its distinctive value.

3.8. New Capabilities that Could Be Supported by AI

Survey participants suggested further areas for automation beyond kerning, and their responses list chores that are perceived as labor-intensive or technical:

Frequently cited were hinting (13 mentions), extending glyphsets (10), spacing (10), mastering/generating font files (7), vertical metrics (8), QA (5), generating text strings/proofing documents (6), adding font info (5), writing OpenType features (4), ensuring cross-platform compatibility (6), scripting/development (2), outline cleaning/refinement (3), and fixing interpolation issues (1). Drawing italics was mentioned twice, an activity that touches more directly on letterdrawing; however designers already have access to tools that derive an oblique from a slanted Roman style.

Interestingly, most of these operations are (at least in part) already handled by deterministic automation. The prevalence of such examples suggests that respondents are currently projecting familiar forms of automation onto AI, rather than envisioning new capabilities. But this tendency is likely to change as tools using machine learning mature.

Nevertheless, a handful of comments point beyond current routines: one respondent envisaged an assistant that would mediate between type designer and end user, “saving the user hours of type education” for instance by offering context-sensitive advice on stylistic sets (P8). Another wrote that he would welcome a tool that “keeps track of all the decisions made during the design process and checks if their application remains consistent” (P273) while another would like a tool that verifies technical compliance with specific environments such as Microsoft Office (P101).

The most ambitious proposal is a system able to extrapolate an entire font from a small set of key characters—an idea which, if successfully implemented, would be hailed by some as a breakthrough and by others as an existential threat:

It would be nice to design the key characters... and let a tool expand them to derivatives. But also be able to teach the tool how you would like the derivatives to look (P135).

I would like to automate drawing all those symbols which are not important for the design or look the same in most designs (like math symbols). Currency symbols are also no fun to draw and could be automated (P222).

Drafting an additional master based on an existing master (such as bold, slanted, etc.). Maybe AI can do a better job at extrapolating than conventional tools (P221).

Drawing based on a handmade sketch, finalizing letters, adding currency symbols, and other elements essential for a complete glyph set... I feel that the main characters are the most important, while the rest is more technical (P219).

Experimental platforms, including *Typograph.studio* and Monotype's *Human Types* and AI project (see footnote 3; Monotype, 2025; Monotype, n.d.-b), have already started to explore such functionality, but none yet is publicly available at the time of writing. Some respondents extended the concept to script expansion: one respondent envisaged rapid generation of CJK glyphs once the "design direction is set for key characters" (P93), while another would like to be able to add Arabic, Hebrew or Vietnamese ranges based on a Latin character set (P293).

The prospect of automated glyph expansion raises a number of concerns dealing with intellectual property, since the legal status of the libraries used for training remains unresolved. It also brings about questions of authorship, with the reservation of respondent P135 above echoing a more widespread reservation against computer-generated type design: that it would be faster to draw a typeface than to describe parametrically all detail decisions. The same notion, but within the field of programming itself, is expressed by programmer and consultant Kevin Henney (2009): "The act of describing a program in unambiguous detail and the act of programming are one and the same." And without an infinite amount of direction, who is really the author of the resulting shapes?

An AI typeface synthesizer can be imagined to create a possible shape of a letter to match a provided partial character set based on training data, that is to say: in a plausible, standard way. But would it be likely to exactly respond to the designers intent and every formal aspect specific to this unique typeface? Create something new on purpose? Or would it more likely simply propagate biases in standard maneuvers and recipes? The idea of autocompleting partial character sets therefore also raises the issue of script equity, as models trained predominantly on Latin data risk perpetuating existing cultural imbalances. As Borna Izadpanah wrote:

One of the crucial questions is whether these technological advancements will support the increasing diversity we are witnessing in type design. Will these tools continue the trend of disproportionately Latin-centric developments, reinforcing Euro-centric knowledge structures? Or will they place more emphasis on other world languages and writing systems? The risk is that AI-based tools will be

trained primarily on Latin script due to the abundance of available data, while other world scripts remain underdeveloped and underserved.*

4. Conclusion

In conclusion, our research indicates that deterministic automation is embedded in the type design industry, whereas AI-driven tools must still demonstrate their relevance, reliability and transparency, especially with respect to the provenance of training data. Practitioners are determined to retain control over the creative core of their workflow, particularly the drawing of letterforms, yet they are willing to delegate labor-intensive tasks such as kerning and font rendering optimization. Respondents stress that designers should keep final decision-making authority, and weigh any technological gain against ethical considerations. The absence of market-ready AI applications specific to typeface design undoubtedly underpins current skepticism. Nevertheless, general purpose LLMs, which are already competent at writing code and text, have already been adopted by a sizable minority of practitioners for peripheral activities such as script writing and proofing document generation.

Overall, the successful integration of automation into type design practice appears to depend on maintaining an appropriate balance between preserving creative authorship and key design decisions, while delegating routine technical operations to trustworthy software. Our findings also temper enthusiasm for wholesale automation, as iterative routines are not merely production steps but they are also the very means through which designers acquire tacit knowledge, and sharpen their judgement. If these cycles are removed entirely, the feedback loop that sustains expertise may wither. Therefore, the real balance could lie in developing tools that accelerate routine operations while encouraging experimentation and gradual mastery, rather than reducing designers to operators of opaque processes.

It should be remembered as well that these findings reflect the views of professionals whose expertise and livelihoods derive from the craft of type design. A wider population of font users—those who do not depend on designing and distributing typefaces for their income—may take a different view and positively welcome the prospect of generating complete fonts at the touch of a button. One participant succinctly captured the urgency of this shift, stating that:

...as an industry, we need to be prepared for a near future where AI starts to become actually useful for the kind of work we do, as well as a world where

* Borna Izadpanah, email correspondence with Alice Savoie, February 10, 2025.

anyone who can instruct an AI can generate a usable font. Figuring out what that means to us is tremendously relevant (P124).

5. References

- Adobe Systems (1992). *Myriad: A two-axis multiple master typeface*. Adobe Systems.
- Azadi, S., Fisher, M., Kim, V., Wang, Z., Shechtman, E., & Darrell, T. (2018). Multi-content GAN for few-shot font style transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 7564–7573). <https://doi.org/10.1109/CVPR.2018.00789>
- Bataineh, B., Sheikh Abdullah, S. N. H., & Omar, K. (2012). A novel statistical feature extraction method for textual images: Optical font recognition. *Expert Systems with Applications*, 39(5), 5470–5477. <https://www.sciencedirect.com/science/article/pii/S0957417411016241>
- Buchholz, K. (2024, August 23). The extreme cost of training AI models. *Forbes*. <https://www.forbes.com/sites/katharinabuchholz/2024/08/23/the-extreme-cost-of-training-ai-models/>
- Brunner, O. (2024). *Archeology of Consequent Forms*. Maxitype and VC-Edition.
- Carlier, A., Danelljan, M., Alahi, A., & Timofte, R. (2021). DeepSVG: A hierarchical generative network for vector graphics animation. In *Advances in Neural Information Processing Systems*, 33. Neural Information Processing Systems Foundation.
- Diem, M., Kleber, F., Fiel, S., Grüning, T., & Gatos, B. (2017). cBAD: ICDAR2017 Competition on baseline detection. In *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)* (Vol. 1, pp. 1355–1360). IEEE. <https://ieeexplore.ieee.org/document/8270153>
- FontParts (n.d.). FontParts. <https://fontparts.robotools.dev/en/stable/>
- Fonts In Use. (2025, June). Fonts in use: Font Bureau. <https://fontsinuse.com/foundry/33/font-bureau>
- Fonts In Use. (n.d.). Fonts in use: Monotype. <https://fontsinuse.com/foundry/61/monotype>
- Gao, L., Tang, Z., Lin, X., & Qiu, R. (2008). Comprehensive global typography extraction system for electronic book documents. In *Proceedings of the Eighth IAPR International Workshop on Document Analysis Systems* (pp. 615–621). IEEE Computer Society.
- Google Labs. (n.d.). GenType. <https://labs.google/gentype>
- Grebäck, M. (2022). Creating typefaces through machine learning. *Github*. <https://github.com/mansgreback/ai-typography?ssp=1&setlang=fr&cc=FR&safesearch=moderate>
- Heller, S. (2015, October 5). You wouldn't think it, but typeface piracy is a big problem. *Wired*. <https://www.wired.com/2015/10/you-wouldnt-think-it-but-typeface-piracy-is-a-big-problem>
- Henney, K. [@KevlinHenney]. (2009, August 17). Post on X (formerly Twitter). <https://x.com/kevinhenney/status/3361631527>
- Henshall, W. (2024, June 3). The billion-dollar price tag of building AI. *Time Magazine*. <https://time.com/6984292/cost-artificial-intelligence-compute-epoch-report/>
- iKern. (n.d.). iKern. <https://www.ikern.space>
- Kern On. (n.d.). Kern On. <https://kern-on.com>
- Knuth, D. E. (1979). Mathematical typography. *Bulletin (New Series) of the American Mathematical Society*, 1(2), 337–372.
- Knuth, D. E. (1985). Lessons learned from Metafont. *Visible Language*, 19(1), 35–53.
- Lineto. (2024). End user license agreement. <https://lineto.com/information/legal/legal/lineto-eula>
- LLM0113. (2023). House of Lords Communications and Digital Select Committee inquiry: Large language models. <https://committees.parliament.uk/writtenevidence/126981/pdf/>

- Luk, G., et al. (2023, December). Chart. https://typedrawers.com/discussion/comment/64717#Comment_64717
- Metrics Machine. (n.d.). Metrics Machine documentation. <https://extensionstore.robofont.com/extensions/metricsMachine/documentation/>
- Mohammed Javed, Nagabhushan, P., & Chaudhuri, B. B. (2014). Automatic detection of font size straight from run length compressed text documents. *International Journal of Computer Science and Information Technologies*, 5(1), 818–825. <https://www.ijcsit.com/docs/Volume%205/vol5issue01/ijcsit20140501178.pdf>
- Monotype. (n.d.-a). AI code of ethics. <https://www.monotype.com/monotype-ai-code-ethics>
- Monotype. (n.d.-b). BlazeType. <https://www.monotype.com/type-trends/human-types/blaze-type/>
- Monotype. (n.d.-c). Font plagiarism. <https://foundrysupport.monotype.com/hc/en-us/articles/360029957811-Font-Plagiarism>
- Monotype. (n.d.-d). Newsroom. <https://www.monotype.com/company/news-press>
- Monotype. (2025, May). Human types. <https://www.monotype.com/type-trends/human-types>
- Morison, S. (1930/1951). *First principles of typography*. Cambridge University Press. (Original work published 1930 in *The Fleuron*, Vol. 7.)
- Murdock, M., Reid, S., Hamilton, B., & Reese, J. (2015). ICDAR 2015 competition on text line detection in historical documents. In *2015 13th International Conference on Document Analysis and Recognition (ICDAR)* (pp. 1171–1175). IEEE.
- NaN Foundry. (n.d.). A machine learning font. <https://www.nan.xyz/projects>
- Pearson, D. [@davidpearson]. (2024, December 12). Instagram post. https://www.instagram.com/p/DDd6KOGCUTW/?hl=en&img_index=1
- Poynor, R. (2003). *No more rules: Graphic design and postmodernism*. Yale University Press.
- Remix Tools. (n.d.). <https://remix-tools.com>
- RoboFont. (n.d.-a). The pre-history of RoboFont. <https://doc.robofont.com/documentation/topics/robofont-pre-history/>
- RoboFont. (n.d.-b). RoboFont design principles. <https://robofont.com/documentation/topics/robofont-design-principles>
- Sennett, R. (2008). *The craftsman*. Allen Lane.
- Shinahara, Y., Karamatsu, T., Harada, D., Yamaguchi, K., & Uchida, S. (2019). Serif or sans: Visual font analytics on book covers and online advertisements [Unpublished manuscript]. *Arxiv*. <https://arxiv.org/abs/1906.10269>
- Srivatsan, N., Barron, J. T., Klein, D., & Berg-Kirkpatrick, T. (2020, May 16). A deep factorization of style and structure in fonts [Unpublished manuscript]. *Arxiv*. <https://arxiv.org/abs/1910.00748>
- Thomas, J. (2025, April 14). « C'est un métier ! » : chez les avocats, l'IA est déjà là. *Le Monde*. https://www.lemonde.fr/emploi/article/2025/04/14/c-est-un-metier-chez-les-avocats-l-ia-est-deja-la_6595713_1698637.html
- Tselovalnikov, S. (2023). Hey, computer, make me a font. <https://serce.me/posts/02-10-2023-hey-computer-make-me-a-font>
- Type.tools (Schumann, N., Eskinazi, C., & Otsuka, M.). (2020). MILK tool [Conference presentation]. *Typographics 2020*. <https://type.tools/> and <https://vimeo.com/431803660>
- Typograph (n.d.). Typograph. <https://typograph.studio/en>
- Ulrich, F. P. (2022). *The discourse in early digital type design technologies* (Doctoral thesis). University of Reading.
- Unger, G. (1982). Other replies to Donald E. Knuth's article, "The concept of a Meta-Font." *Visible Language*, 16(4), 354.

- Unger, G. (2018). *Theory of type design*. nai010 publishers.
- Wang, M. (2024). *Craft is dead, long live craft* (Master's thesis). ECAL.
- Wang, Z., Yang, J., Jin, H., Shechtman, E., Agarwala, A., Brandt, J., & Huang, T. S. (2015, July 12). DeepFont: Identify your font from an image [Unpublished manuscript]. *Arxiv*. <https://arxiv.org/abs/1507.03196>
- Wasim, S. T., Collaud, R., Défayes, L., Henchoz, N., Salzmann, M., & Ribes Lemay, D. (2024). Toward automatic typography analysis: Serif classification and font similarities. *Journal of Data Mining and Digital Humanities*. <https://doi.org/10.5281/zenodo.10623072>
- Wenzel, D. (n.d.). Automated type design. https://wenzeldaniel.com/w_atd.php

Authors

Alice Savoie is a typeface designer, researcher, and teacher. She holds an MA and a PhD from the University of Reading, where she is currently a Visiting Research Fellow. Her research examines technological change in typographic practice, the place of women in type history, and the role of type design in the digital humanities. She contributes to the research project *Reception of the New Typography in Francophone Graphic Scenes*, supervised by Professors Catherine de Smet (Université Paris 8) and Davide Fornari (ECAL/HES-SO). She teaches at the Atelier National de Recherche Typographique in Nancy and is a visiting professor on the Master Type Design programme at ECAL/University of Art and Design Lausanne (HES-SO). She lives and works in Lyon.

Kai Bernau has taught type design on the Master Type Design programme at ECAL/University of Art and Design Lausanne (HES-SO), and before that ECAL Master Art Direction, since 2011. His practice, Atelier Carvalho Bernau, designs reading experiences, from typefaces to interfaces, for culture and commerce. Commissions range from publications like the *New York Times* and *Esquire*, and publishers such as Walther König and Phaidon, to luxury brand Marsèll, cultural institutions Culturgest Lisbon and Villa Stuck in Munich, to academic clients like Brill publishers, Leiden and University of Warwick. The retail typefaces they created or collaborated on (Neutraface Slab, Publico, Lyon, Atlas, Neutral, Algebra) are used by leading brands and publications around the world. Previous achievements in generative, computer-aided design include a typeface system for Sandberg Institute, and art direction for Octavo publishers, Amsterdam. Kai Bernau holds a Masters degree from Type & Media in The Hague. He lives and works in Porto.

Wayne Daly is a graphic designer and co-director of the design practice Daly & Lyon. The studio focuses on collaborations with cultural and arts clients, producing books, exhibitions, identities and websites. Clients include Barbican Art Gallery, Design Museum, Royal Academy of Arts, Forensic Architecture, Hayy Jameel, Institute of Contemporary Arts, KW Institute for Contemporary Art, Sternberg Press, Tate Britain and Verlag der Buchhandlung Walther und Franz König. He holds an MA in Typographic Studies from the London College of Printing and is currently a visiting professor on the Master Type Design programme at ECAL/University of Art and Design Lausanne (HES-SO), where he co-directs the writing and thesis curriculum. He lives and works in London.

Raphaëla Haefliger is a graphic designer specialized in art direction, typography and type design. She holds a BA in Graphic Design from the University of Lucerne and an MA in Type Design from ECAL/University of Art and Design Lausanne (HES-SO). Since 2016 she has been working as an independent graphic designer in the commercial and cultural field. In addition to her self-employment, she works part-time at ECAL, Master Type Design and Research Department, ZHdK, MAS Design Direction and F+F: FA Visuelle Kommunikation.

Sebastian Baez-Lugo is a researcher in user experience (UX), cognitive and affective science. He earned his PhD in Neuroscience from the University of Geneva, conducting his doctoral work at the Swiss Center for Affective Sciences (CISA), where he investigated the mechanisms underlying human cognition and emotional processes. Driven by a passion for data analysis and visualization, Sebastian applies quantitative methods and insights from cognitive and affective science to explore how people think, feel, and interact with technology. At the EPFL+ECAL Lab, he leads the UX Research and Cognitive Science activities, overseeing projects that connect applied research with human-centred design in fields such as health innovation, artificial intelligence, and environmental technologies. He also mentors early-career design researchers, and builds collaborations between academic institutions and industry. He lives and works in the Lausanne area.



Evaluating Interactive Highlighting Techniques in Digital Reading: An Empirical Study of Hover-Based Line, Sentence, and Paragraph Highlighting

Szabolcs Vatány^a , Thi Huyen Nguyen^b , Anikó Illés^a , and Ann Bessemans^c 

^a Doctoral School, Moholy Nagy University of Art and Design; ^b Data Science Institute, Hasselt University; ^c Faculty of Architecture and Art, Hasselt University, READSEARCH research group

Corresponding author: Szabolcs Vatány (vatany.sz[at]gmail.com)

Abstract: Maintaining user engagement and supporting comprehension remain key challenges in digital reading environments. This study examines the impact of interactive hover-based text highlighting—on the line, sentence, and paragraph levels—on reading speed, comprehension, perceived attention, and user preferences during interlude reading. In a study with 80 participants, we compared these interactive techniques to static text presentations. While no statistically significant differences were observed in comprehension or reading speed across the tested highlighting methods, participants' subjective ratings showed significantly higher perceived attention sustainment with sentence- and paragraph-level highlighting compared to the static condition. These findings suggest that while such techniques may not enhance measurable reading performance, they can positively influence user experience. This work informs the design of digital reading interfaces by presenting the potential of user-preferred interaction mechanisms to support attentional engagement. Future research should investigate the long-term effects and adaptation to mobile contexts, as well as assess the relevance of these techniques for readers with attentional variability.

Implications for practice: Designers of digital reading interfaces should note that while interactive hover-based highlighting does not directly improve comprehension or reading speed, it can meaningfully enhance readers' perceived focus. Highlighting by grammatical units—sentences or paragraphs—was both preferred and rated as more attention-sustaining than line-by-line highlighting, which divided users. This distinction matters: aligning highlights with how readers naturally process meaning appears more effective than segmenting by visual lines. However, highlighting alone should not be expected to boost learning outcomes. Its practical value lies

@: [ISSUE](#) > [ARTICLE](#) >

Cite this article:

Vatány, S., Nguyen, T. H., Illés, A., & Bessemans, A. (2026). Evaluating interactive highlighting techniques in digital reading: An empirical study of hover-based line, sentence, and paragraph highlighting. *Visible Language*, 60(1), 58–83.

First published online April 26, 2026.

© 2026 Visible Language — this article is **open access**, published under the CC BY-NC-ND 4.0 license.

<https://visible-language.org/journal/>

Visible Language Consortium:

University of Leeds (UK)

University of Cincinnati (USA)

North Carolina State University (USA)

in supporting sustained engagement, particularly when integrated alongside comprehension-oriented strategies in e-reading tools and digital textbooks.

Keywords: attention sustainment; digital reading; interactive highlighting; reading comprehension; typographic design; user experience

1. Introduction

1.1. Problem Statement and Motivation

Reading remains one of the primary ways through which people acquire information, learn, and find entertainment—in short, a core means of creating meaning and interpreting the world (Smith, 2004). The ability to read is fundamental to both social connections and personal development on intellectual and social levels. Reading enables participation in civic life, access to information and education, and the ability to communicate across distances and time through written correspondence and digital media (Bessemans, 2012). In a literate society, struggling with reading or being unable to read has significant consequences for quality of life, including reduced employment opportunities, limited access to healthcare information, difficulty navigating administrative systems, and social isolation, which is often used as an indicator of well-being (Poulsen et al., 2005). The rise of digital technologies and the World Wide Web has transformed reading habits, blurring the boundaries between personal, work, and study-related reading, particularly through screen-based formats. People now frequently switch between reading emails, social media, news articles, and work documents on the same device, often within the same session, fundamentally changing the nature of reading from discrete activities to a continuous, fragmented practice embedded throughout daily life (Mangen & van der Weel, 2016). The way text is displayed on screens significantly affects readability and the ability to maintain focus over time, compared to print-based media (Carr, 2020). According to prior studies (Hisgen & van der Weel, 2022), sustained reading over longer periods is essential for developing critical thinking skills. In today's highly media-saturated digital environment, interfaces are often visually busy and rich in stimuli. This creates a challenging setting for sustained concentration and can negatively affect readers' attention and comprehension.

While traditional approaches to improving reading focus on static text presentation, HCI research has explored interactive reading interfaces. For example, LiquidText introduced multitouch highlighting and flexible content manipulation to support active reading (Veras et al., 2014), and ScholarPhi augmented scientific papers with in-situ definitions to aid reader comprehension (Head et al., 2021). Intelligent skimming

tools like Scim even highlight salient content to guide readers' attention during quick reviews (Fok et al., 2022). Building on this work, our study examines a new hover-based highlighting technique in short-text reading, extending prior findings by focusing on sustained attention. Moreover, foundational eye-tracking research (e.g., Rayner's studies of reading) has long shown how managing visual attention can impact comprehension (Rayner, 1998), which our design-centric approach leverages and tests in an interactive setting. As the digital environment is highly interactive, we perform most of our activities by using the mouse or the keyboard. Integrating a small, supportive gesture into the process of digital reading might help sustain attention and thereby enhance comprehension.

Our review covers existing research on digital reading, attention sustainability, and comprehension enhancement. It points out the gaps in current methods and the uncovered potential for innovative approaches, such as hovering over text fragments and highlighting them, to strengthen and sustain reading attention and enhance comprehension. By our interpretation, previous studies indicate that reader interaction with text may foster engagement and thus improve comprehension, but empirical evidence on specific interactive methods remains limited.

We believe it is important to support the field of readability with valuable insights from research stemming from the design aspects of digital surfaces and reading, such as the stated problem of sustaining attention in digital contexts. The use of design practices in relation to typographic possibilities is still not explored deeply enough compared to their potential. Design patterns can help in defining strategies for readers to better access, grasp and understand information.

1.2. Research Questions and Hypothesis

This study addresses four research questions: (1) Does hover-based highlighting of text fragments improve reading comprehension? (2) Does hover-based highlighting affect reading speed? (3) Is using a participant's preferred highlighting method related to their performance on these measures? (4) How do readers evaluate each highlighting method regarding attention sustainability? These questions together aim to identify how the introduced interactive highlighting techniques influence both reading behavior and perceived attention.

We hypothesize that hover-based reading methods will enhance sustained reading attention, and thus comprehension, but make reading time longer because of the extra effort required for interaction with the cursor. We also hypothesize that the sentence method might be the most preferred due to its potential to help with line changes and maintain focus on the same small unit of content during reading. Connected to this assumption, we expect the line method to underperform in comprehension even if it

helps with line changing, as it is not a meaningful grammatical unit but only a visual one. We also expect that participants using their favorite method should perform better in the measurements of reading comprehension compared to other examined methods.

2. Relating Works on Reading

2.1. Interlude Reading

In readability studies, reading situations are often categorized based on the reading purpose, such as reading for enjoyment or for gaining knowledge (Kaakinen et al., 2018). Also, the capabilities of the reader can vary based on their age, reading skills, professional skills, or possible learning or reading disabilities. Another way these can be categorized is based on the length of the time it takes to read the text, which can specify the way of examination as a glanceable reading has to be investigated differently than a reading in some paragraphs or in the case of long-format texts (Beier et al., 2021; Sieghart & Gorbach, 2024).

We frame our investigation on readability within the context of short-term reading sessions. The term we are going to use is *interlude reading*. This term is established and defined by Wallace and colleagues as nestled between glanceable and long-form reading on the axis of time spent with reading a text, as the form of reading that occurs in short interludes (Wallace et al., 2020). Situations include, for example, reading short articles while waiting for public transport, or checking the news between two tasks at work.

Reading just for a short period is a very common situation nowadays. Since the advent of digital technologies, the amount of textual information has exponentially grown, creating an oversupply of materials to read online. As a result, readers tend to prefer shorter texts to long articles, which are faster to read, so it is possible to consume more information in a shorter and more compressed manner. This theory is further supported by the widespread appearance of short articles on online news portals, which are only a few paragraphs long (Barnhurst, 2013).

These short text formats fit well with the use of new devices, which offer the possibility of reading almost anywhere through phones or smartwatches, creating more opportunities for interlude reading during short breaks between two activities. As these situations happen usually in a busier environment, sustaining reading attention is more difficult, even for a short period of time.

We selected interlude reading as the primary context for our investigation due to its growing prevalence in digital environments, where users frequently engage with short texts during brief periods of availability (e.g., while commuting or multitasking).

This context provided a practical and controlled scope for evaluating hover-based highlighting techniques. However, our choice of interlude reading does not imply that such techniques are only suitable for short-form content. Rather, we used this scenario as a representative use case to examine how interactive highlighting affects attention and comprehension. The highlighting fragments tested—lines, sentences, and paragraphs—were chosen to match the typical structure of short texts, but we acknowledge that different fragment strategies might prove valuable in extended or deep reading contexts, which remain outside the scope of this study.

2.2. Sustaining Reading Attention

Reading itself is a complex cognitive process demanding a high level of attention to achieve good text comprehension. Attention can be disrupted by surroundings, noise, and visual elements, affecting the reading flow and thus reading comprehension. Previous studies of readability found that sustaining reading attention can be harmed by the appearance of hyperlinks and can pose an extra challenge to comprehension. Navigating through hypertexts while reading is likely to place additional demands on working memory compared to traditional linear reading (Cuddihy & Spyridakis, 2012; Salmerón et al., 2018; Wylie et al., 2018).

This distraction is even more accentuated because of the notification and pop-up systems of the devices we use for reading, especially during interlude reading. This results in poorer performance on concurrently performed tasks because limited attentional resources must be shared between tasks (Stothart et al., 2015).

As most interfaces and digital surfaces are visually oversaturated by multimodal content, they demand sustained attention from readers—a challenge intensified by call-to-action texts and buttons that constantly invite interaction. Compared to a physical book, on a digital surface, we not only tend to read, but also perform several other tasks like working, shopping, social interaction, etc. This multifunctionality enhances the spread of attention even when we are not performing any of these activities, making sustained reading attention even harder.

We hypothesize that active interaction, specifically hover-based highlighting of text fragments, may support sustained attention in highly interactive digital environments.

2.3. Supporting Reading with Interaction

Although printed reading always involved physical interaction, such as turning pages or pointing at text, digital reading introduced new forms of direct interface interaction, where input devices actively alter how text behaves and is presented. Since the digital paradigm shift, reading has undergone significant transformations in how, when, and why people engage with text, as on-screen texts displayed on electronic devices offer the

possibility to interact with the surfaces using different inputs like a mouse, keyboard, touchscreen, or voice command. Although this possibility has existed for some time, highly interactive reading approaches have not achieved widespread adoption, except in smaller phenomena such as digital and electronic literature.

The simplest and established interaction form with digital texts while reading is links on digital texts, usually used for navigation, word explanation or a digital glossary of the text content, with the use of clicking on links or hovering over them. Certain interactive features, such as digital glossaries and embedded questions with feedback, enhance comprehension by supporting the development of mental models (Clinton-Lisell et al., 2023). These features assist in bridging knowledge gaps and reinforce connections within the text, improving learning outcomes.

Interactive functions such as animations and story-related games can boost engagement and vocabulary development in younger readers. However, poorly aligned interactive features may distract from core reading tasks, highlighting the need for age-appropriate design (Hare et al., 2024). Contrary to earlier assumptions of a screen inferiority effect, no statistically significant difference in comprehension between digital and print mediums for narrative texts was detected when interactive features are not present. However, multimedia and interactive additions (dictionary, glossary, virtual tutor) positively affect comprehension by increasing engagement and promoting deeper processing of content (Schwabe et al., 2022).

3. Method

3.1. Participants

The study was conducted in two university-affiliated research laboratories located in Belgium and Hungary. Participants were recruited through university mailing lists and social media announcements targeting university students, staff, and members of the local community. To minimize bias in the assessment of attention-sustaining and preferred reading methods, participants were eligible only if they met specific criteria:

- ▶ Have at least intermediate English proficiency, meaning that they can read and understand written English texts fluently enough to follow study materials and instructions;
- ▶ Be a regular digital reader, defined as someone who habitually reads on-screen texts for information or leisure, engaging with digital reading interfaces several times per month or more often;
- ▶ Demonstrate an understanding of the study procedures before participation. The study information sheet and consent form were provided in written English

and explained verbally in English by the researcher. Each participant confirmed their understanding before starting the test;

- ▶ Be able and willing to complete the study;
- ▶ Be between 18 and 75 years of age.

Exclusion criteria included:

- ▶ Self-reported uncorrected visual impairments or severe reading difficulties;
- ▶ Professional expertise in typography, type design, or readability research that could bias subjective ratings;
- ▶ Difficulty understanding English instructions during the pre-test explanation;
- ▶ Insufficient reading fluency to comfortably read and comprehend standard informational texts.

In the study, 80 participants who met the inclusion criteria were enrolled. Thirty-nine (39) of the participants took part in the experiment in a research lab office room in Belgium, and 41 of the participants took part in the experiment at a collaborative space in Hungary. The study protocol was reviewed and approved by the Social and Ethical Committee of UHasselt Faculty of Science (reference no. REC/SMEC/2023-2024 40). All participants provided informed consent before participation and were free to withdraw at any time without penalty.

Most participants (55%) were women (44 participants). Participants in the age group 18–25 years accounted for 36.3% (29 participants), while the 25–40 years and 40–60 years represented 51.3% and 12.4% of the total number of participants, respectively. Of the 80 participants, 29 (36.3%) were Dutch native speakers. About their profession, 62 participants (77.5%) had no relation to typography, type design, graphic design or readability research, and 18 (22.5%) had some relation to these fields through coursework and interest, none held a degree above BA level in these fields (mostly Graphic Design BA students).

3.2. Materials and Tools

The materials used in the experiment consisted of four open-source text sections and their corresponding comprehension questions, obtained with permission from the study “Towards Individuated Reading Experiences: Different Fonts Increase Reading Speed for Different Individuals,” conducted by Wallace and colleagues (2022). The texts were downloaded from the Readability Consortium website (The Readability Consortium, n.d.). The four text levels were classified between 7.0 and 8.6 Flesch-Kincaid grade level English texts. After consulting with the researcher responsible for the measurement of the material, we decided to increase the number of multiple-choice options from three to four for each question, adding one distractor to reduce the chance of correct guessing. The text setting was based on the suggested general text setting for

desktop digital devices and websites, as the study primarily focuses on digital text reading. Each text consisted of two paragraphs and was between 897 and 962 characters long. Each text sample fitted entirely within one screen view (no scrolling required) to prevent visual discontinuity or navigation effects. This ensured that all participants viewed and interacted with the full text in a single frame. The letter size was set to 16 px, as this size is commonly used in similar reading studies and falls within the range recommended for body text in digital interfaces (Wallace et al., 2022). The width of the text block was 750 pixels to achieve around 70–80 characters (8–12 words) per line, a text setting recommended for digital texts by Web Content Accessibility Guidelines (WCAG) 2.1 (W3C, 2018). The text was aligned left (ragged-right) to avoid uneven word spacing that occurs with fully justified alignment, which could influence readability. The letter color was 100% black on a 100% white background. The text line height was set to 1.5 times the font size according to WCAG 2.1 suggestions. The font used in the experiment was a custom humanist-modernist sans-serif typeface (regular weight) designed by Szabolcs Vatóny, a member of the research team. Its use ensured full control over typographic consistency and legibility across all stimuli. The font's purpose in this study was not to examine typeface effects but to provide a neutral, reliable text presentation and to maintain complete control over text display and rendering. The test environment was designed as a single-window application opened in full-screen mode. The pages with the reading tasks contained only the texts and a button under the text to navigate to the next screen (Figure 1). These were separated by explanation and input pages.

All participants read from standardized 13- or 14-inch MacBook Pro devices provided by the research team, ensuring similar display resolution and luminance. Screen brightness was set to 100% on all devices to maintain consistent luminance levels. The font size and layout were locked in pixels, ensuring identical text scaling across both screen sizes. Participants sat approximately 50–60 cm from the screen, at a comfortable self-chosen distance, with the laptop placed on a table. They were instructed to adjust the screen angle for comfort but not to change the zoom level or window size. They could use either the built-in trackpad or a provided Bluetooth mouse. Although participants could choose between input devices, we did not record which option they used. This decision was made because the primary aim of the study was to examine the effects of the highlighting conditions in a context that reflected participants' natural reading habits and their typical interaction with digital devices. While the number and speed of cursor hover events were tracked, these data were not included in the present analysis, as they were not directly relevant to the research questions investigated in this study, which focused on reading comprehension and reading speed.

In each section, the text was designed with one of three interactive methods: line-by-line, sentence-by-sentence, or paragraph-by-paragraph or normal static reading (denoted as "Static"), which served as the reference condition and displayed all text at

100% black with no interactive highlighting. Specifically, for the line-by-line method, the participant reads a text by hovering over each line of the text with the cursor for succession from the start to the end of the text, while the hovered line is kept 100% black and the other lines are reduced to a 30% transparency. The other two interactive reading methods work the same way, but they involve hovering over each sentence (i.e., sentence-by-sentence, as shown as an example in Figure 1) and each paragraph (i.e., paragraph-by-paragraph) of the given text, whereas in the normal static reading condition hovering is not possible and no text is highlighted.

3.3. Procedure

We used a 4×4 Latin square model in which we had four sections (texts) with four reader groups compiled randomly.

Both testing environments were well-lit and moderately quiet shared workspaces. A modest level of background activity and ambient noise was intentionally maintained to reflect typical conditions in university or office study environments. This ensured that participants read within a realistic context that included minimal everyday distractions, allowing for a more ecologically valid assessment of attention-sustaining effects.



Figure 1. Screen used to display sections for participants reading with the four methods. In this example, the third sentence of the sample text is hovered over and highlighted using the sentence-by-sentence method.

The flow of the testing process started with a short explanation of the testing flow for the participants to have some idea of what they should do during the test. After starting the process on the laptop, they were faced with a letter of consent. After accepting it, the next screen was a preparation screen explaining which method they should use to read the next text and a “Start reading” button. After clicking the button, the participant was presented with the first task, which required them to read the text on the next screen using one of the reading methods. Participants were asked to move the cursor naturally across the text to activate the hover highlights, but they were free to proceed at their own pace. To finish the task, under the text displayed, they had to click the “Finish reading” button and were presented with two comprehension questions, each shown separately on consecutive pages after clicking the “Next” button. Each comprehension question offered four possible answers. They could select one of the answers and navigate further by clicking the “Next section” button to the next task screen.

After finishing the four sections, participants were presented with a post-reading survey and asked to indicate how much they felt the reading methods helped to maintain attention while reading compared to the static text, using a five-point scale, where two points on the left indicated “Harder than regular”, the middle point indicated “Same as regular static text” and two points on the right indicated “Easier than regular” (Figure 2). Afterwards, on the next page, participants selected one of the four reading methods (static, line-by-line, sentence-by-sentence, paragraph-by-paragraph) as their favorite, if they had one, and indicated if they felt motivated to read using that method.

During the study, we recorded each participant’s reading time for each method, their answers to the questions for each text provided and their ratings of the methods in terms of attention-sustaining ability, preferred method and motivation answers. The basic information on the demographic characteristics of the participants was also collected. All of these were tracked by our self-developed platform and exported at the end of each test.

3.4. Data Analysis

To explore the effects of the reading methods (static, line-by-line, sentence-by-sentence, paragraph-by-paragraph) on reading comprehension, we investigate the influence of the reading methods on the correctness of answering a question (i.e., the probability of having a “true” or “false” answer, compared to the correct answer) within each participant after reading the sections provided. We also want to investigate how the reading methods affect the time spent on each section (i.e., text) of a participant.

Because each participant answered multiple questions across different conditions, their responses were not statistically independent; the data, therefore, required a model accounting for within-participant dependency. It is, therefore, of paramount

Rate how the two hovering reading methods maintained your attention compared to regular static text reading. The middle circle is labeled 'Same as regular static text'.

Highlight Line by Line

four groups of people—housewives, chefs, doctors, and dietitians—were eager to adopt a product that would lead to more digestible, economical, and tastier foods. **Great foresight was shown in the development of Crisco.** The quality, as well as the quantity, of lard was steadily diminishing against the backdrop of a growing population. Prices were escalating. 'The high cost

Harder than regular Easier than regular

Highlight Sentence by Sentence

midwest of the United States. Soon after, this farmer contracted for a boat to be built that would be powered by steam. **Although advised by his builders to substitute the common four-bladed propellers from older versions of boats, he stuck to his original design.** With one propeller at either side of the rudder—called 'twin-propellers'—the boat was ready for use. This boat is now

Harder than regular Easier than regular

Highlight Paragraph by Paragraph

eager to adopt a product that would lead to more digestible, economical, and tastier foods. **Great foresight was shown in the development of Crisco.** The quality, as well as the quantity, of lard was steadily diminishing against the backdrop of a growing population. Prices were escalating. 'The high cost of living' became a commonly repeated phrase. Moreover, the country was

Harder than regular Easier than regular

Figure 2. Screenshot of the testing platform where participants rated whether the three interactive reading methods supported or hindered their ability to maintain reading attention compared to the static method.

importance that the statistical analysis accounts for the dependency in the data and the study design to ensure the accuracy of the statistical inference. To address our study objectives, we apply (generalized) linear mixed models, including both fixed effects and random effects (Demidenko, 2004; Faraway, 2014). In essence, a (generalized) linear mixed model is a statistical tool that is particularly useful to deal with repeated measures or clustered data where observations are not independent, which is the case in our study. Specifically,

- ▶ The fixed effects represent the relationships between the independent variables and the dependent variable. In our study, the independent variables are (i) the text provided in each section, (ii) the reading method of the text, (iii) the reading test group to which the participant belongs, and (iv) the participant’s indicated favorite reading method as an exploratory, post-hoc factor to examine whether self-reported preference related to performance, recognizing that this variable was not experimentally controlled and may reflect rather than cause perfor-

mance differences. The dependent variable is the correctness of the answer by a participant or the time spent on each section by a participant.

- ▶ The random effects, on the other hand, account for variability in the data that cannot be captured by the fixed effects. In our study, the random effects are attributed to participants and the questions they answered.

We analyzed comprehension (correct vs. incorrect answers) using a logistic mixed-effects model and analyzed reading time using a linear mixed-effects model on log-transformed times. These models included reading method, text section, group, and the participant's preferred method as fixed effects, with participant and question as random effects to account for repeated measures.

Analysis of the rating regarding attention-sustaining on reading methods. The participants were asked to give their evaluation on the difficulty, as shown in Figure 2, for each reading method (line-by-line, sentence-by-sentence, or paragraph-by-paragraph) in terms of how much it helped them maintain attention compared to the regular static method. They were also instructed to treat the midpoint of the scale as "Same as regular static text." To compare the proportions of participants' evaluations of the difficulty of each reading method relative to the static baseline, we applied a multinomial test followed by pairwise binomial tests, using the exact test formulation commonly associated with Read and Cressie's method (Read & Cressie, 1988) for categorical data. All the analyzes were performed in R, a language and environment for statistical computing. The statistical significance level was set at 5%.

4. Study Results

4.1. Descriptive Analysis

This part of the paper provides an overview of the study results through descriptive statistics, focusing on participants' number of correct answers, reading time, and rating regarding attention-sustaining related to the four reading methods: line-by-line, sentence-by-sentence, paragraph-by-paragraph, and static (control).

Number of correct answers. The reading comprehension of a participant was measured through two single-choice questions for each text. Overall, the paragraph-by-paragraph method had the highest percentage in the number of correct answers (84.4%), as shown in Table 1.

Mean reading time. In our study, the reading time is assumed to be the time participants spent on each section. As shown in Table 2, the sentence-by-sentence method generally resulted in the longest reading time with 78.32 seconds.

Table 1. Number of correct answers by reading groups, sections, methods and favorite methods.

Category	Group	Correct answers (160 total)	Correct answers (%)
Reading group	Group 1	131	81.88
	Group 2	122	76.25
	Group 3	127	79.38
	Group 4	129	80.62
Text number	Section 1	131	81.88
	Section 2	126	78.75
	Section 3	131	81.88
	Section 4	121	75.62
Reading method	Line-by-line	122	76.2
	Paragraph-by-paragraph	135	84.4
	Sentence-by-sentence	123	76.9
	Static	129	80.6
Favorite method	Line-by-line	84/112	75.0
	Paragraph-by-paragraph	135/168	80.4
	Sentence-by-sentence	197/240	82.1
	Static	93/120	77.5

Selection of the favorite method. The sentence-by-sentence method was the most preferred, chosen by over a third of the participants (37.5%), as shown in Table 3.

Rating regarding attention-sustaining by reading method. Participants rated each interactive reading method on a scale in which the two points on the left indicated that the method was considered harder than regular static reading, the two points on the right indicated that it was easier, and the midpoint represented the static reading. They generally rated the paragraph-by-paragraph and sentence-by-sentence methods easier, with many indicating a neutral or positive rating. Conversely, the line-by-line method was perceived as hardest, with higher percentages of negative ratings. The distribution of these ratings is illustrated in the associated density plot (Figure 3).

Table 2. Mean reading time in seconds by reading groups, texts, methods and favorite methods. “Meas.”: measurements.

Category	Group	Correct answers (160 total)	Mean reading time in seconds
Reading group	Group 1	131	66.98
	Group 2	122	82.38
	Group 3	127	74.92
	Group 4	129	74.64
Text number	Section 1	131	65.86
	Section 2	126	82.34
	Section 3	131	81.80
	Section 4	121	68.91
Reading method	Line-by-line	122	71.04
	Paragraph-by-paragraph	135	74.78
	Sentence-by-sentence	123	78.32
	Static	129	74.78
Favorite method	Line-by-line	84/112	77.22 (56 meas.)
	Paragraph-by-paragraph	135/168	76.26 (84)
	Sentence-by-sentence	197/240	74.56 (120)
	Static	93/120	70.58 (60)

Table 3. Number of participants selecting their favorite reading method.

Group	Participants	Percentage (%)
Line-by-line	14	17.5
Paragraph-by-paragraph	21	26.2
Sentence-by-sentence	30	37.5
Static	15	18.8




Figure 3. Distribution of participant ratings for each reading method on the attention-sustainment difficulty scale.

4.2. Results from Fitted Regression Models

To summarize the outcomes of the fitted models, none of the tested reading methods produced statistically significant effects on comprehension accuracy or reading time. While minor numerical variations were observed across sections and methods, these differences did not reach the threshold of significance.

The influence of reading methods on the correctness of the answer of a participant.

According to the fitted logistic Model 1, there were some numerical differences in comprehension accuracy across the four text sections, but none were statistically significant. Table 4 shows detailed estimates and *p*-values.

When comparing reading methods, Model 1 found no significant effect on correctness. All pairwise comparisons between methods showed $p > 0.25$ (see Table 4 for complete statistical details). The data showed no reliable differences in comprehension accuracy across highlighting conditions. If anything, the data trended in a direction where paragraph-level highlighting performed somewhat better than the others and line-level highlighting performed worse, but this pattern must be interpreted with caution as the results are not statistically significant. We also examined whether the order of exposure to methods (the reading test groups) influenced performance. No statistically significant differences in accuracy were found between the four groups.

Table 4. Estimated fixed-effect parameters, their 95% confidence intervals (CIs), and corresponding *p*-values for both models. The “Original estimates” column shows values on the log-odds scale for the logistic model (Model 1) and the original scale for the linear model (Model 2). The “Exponential scale” column shows odds ratios for the logistic model only (not applicable for the linear model). The 2.5% and 97.5% columns represent the lower and upper bounds of the 95% confidence interval, respectively. The numbers in bold indicate statistically significant results at 5% significance level... [Caption continues on the following page.]

Parameter*	Est.	Original estimates		Est.	Exponential scale		<i>p</i> -value	
		2.5% CI	97.5% CI		2.5% CI	97.5% CI		
Model 1 (correctness)	Intercept 0	1.809	0.191	3.427	6.105	1.211	30.779	0.025
	text_number 2	-0.367	-2.333	1.599	0.693	0.097	4.950	0.709
	text_number 3	0.425	-1.612	2.461	1.529	0.199	11.721	0.677
	text_number 4	-0.509	-2.473	1.456	0.601	0.084	4.289	0.605
	reading_method (L-by-L)	-0.329	-0.909	0.252	0.720	0.403	1.286	0.257
	reading_method (S-by-S)	-0.290	-0.871	0.292	0.749	0.418	1.339	0.319
	reading_method (P-by-P)	0.276	-0.342	0.894	1.317	0.710	2.444	0.372
	reading_test_group 1	0.187	-0.520	0.894	1.205	0.594	2.444	0.597
	reading_test_group 2	-0.412	-1.095	0.271	0.663	0.335	1.312	0.228
	reading_test_group 3	-0.042	-0.739	0.655	0.959	0.478	1.924	0.904
	favorite_method (L-by-L)	-0.275	-1.050	0.500	0.760	0.350	1.648	0.478
	favorite_method (S-by-S)	0.381	-0.301	1.062	1.464	0.740	2.893	0.264
	favorite_method (P-by-P)	0.227	-0.487	0.940	1.254	0.614	2.560	0.525
	Model 2 (time)	Intercept 0	4.086	3.836	4.336	59.511	46.352	76.405
text_number 2		0.198	0.122	0.274	1.219	1.130	1.315	0.000
text_number 3		0.184	0.108	0.260	1.202	1.114	1.296	0.000
text_number 4		0.017	-0.058	0.093	1.018	0.943	1.098	0.646
reading_method (L-by-L)		-0.049	-0.125	0.027	0.952	0.882	1.027	0.195
reading_method (S-by-S)		0.040	-0.036	0.116	1.041	0.965	1.123	0.292
reading_method (P-by-P)		-0.021	-0.097	0.055	0.980	0.908	1.057	0.586
reading_test_group 1		-0.099	-0.326	0.129	0.906	0.722	1.137	0.386
reading_test_group 2		0.081	-0.143	0.306	1.085	0.867	1.358	0.468
reading_test_group 3		-0.048	-0.275	0.178	0.953	0.760	1.195	0.671
favorite_method (L-by-L)		0.154	-0.109	0.418	1.167	0.897	1.519	0.241
favorite_method (S-by-S)		0.048	-0.177	0.273	1.050	0.838	1.314	0.667
favorite_method (P-by-P)		0.092	-0.145	0.328	1.096	0.865	1.388	0.438

Table 4 (caption continued from previous page). [* The parameters corresponding to the text_number take the text_number 1 as the reference text. The reference method for the parameters of the reading_method variable is the STATIC method. The reference group for the reading_test_group variable is the reading test group 4. The reference method for the favorite_method variable is the static method.] “Est.”: estimated; “L-by-L”: line-by-line; “S-by-S”: sentence-by-sentence; “P-by-P”: paragraph-by-paragraph.

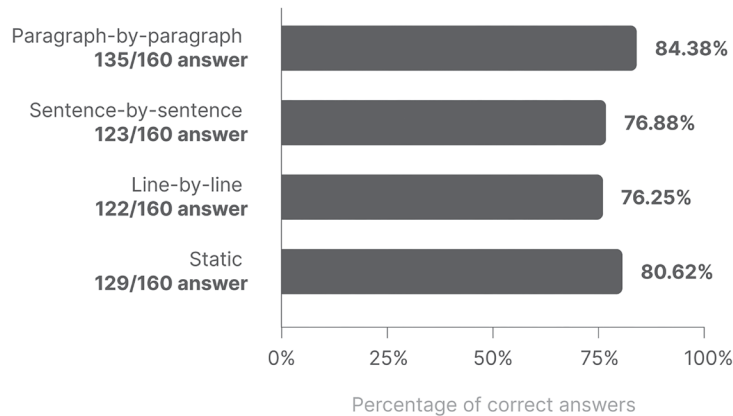


Figure 4. Correctness of the four methods.

Group 1 (order: line → sentence → static → paragraph) had higher odds of a correct answer than Group 4 (order: static → line → paragraph → sentence), whereas Group 2 (order: sentence → paragraph → line → static) and Group 3 (order: paragraph → static → sentence → line) had lower odds than Group 4. These variations were not significant (*p*-values all > 0.05), indicating that the sequence in which participants experienced the methods did not meaningfully affect their comprehension scores.

Finally, we examined whether participants’ stated preference for a particular highlighting method was associated with their overall comprehension performance. We found no significant relationship (all *p* > 0.26). These findings indicate that merely preferring a particular method was not reliably associated with better or worse comprehension performance. These results are summarized in Figure 4.

The effects of reading methods on the time spent on each section by the participant.

Model 2 analyzed the time each participant spent reading each section (Table 4). The results indicate that participants tended to spend more time on later sections compared to the first section. Specifically, Section 2 took about 21.9% longer to read than Section 1 (95% CI: 13.0–31.5%), and Section 3 took about 20.2% longer (95% CI: 11.4–29.6%). These increases were statistically significant, suggesting that participants became more engaged or required more time on the second and third texts. Time spent on Section 4 was only 1.8% higher than on Section 1, and this difference was not significant.

In contrast, the reading method had no significant impact on the time spent per section. Participants using the interactive highlighting methods read at nearly the same pace as in the static condition. Mean reading times were 71.04s for line-by-line, 78.32s for sentence-by-sentence, 74.78s for paragraph-by-paragraph, and 74.78s for static (Table 2). These small numerical differences were not statistically significant ($p > 0.05$ for all pairwise comparisons). Thus, none of the highlighting methods measurably accelerated or slowed down reading relative to static text. In practical terms, participants spent roughly the same amount of time reading each section regardless of whether hover-based highlighting was present.

Similarly, no reliable differences were observed in reading time across the different test groups (method orderings). For example, Group 2 (order: sentences \rightarrow paragraph \rightarrow line \rightarrow static) spent about 8.5% more time per section than Group 4 (order: static \rightarrow line \rightarrow paragraph \rightarrow sentence), while Group 1 (order: line \rightarrow sentence \rightarrow static \rightarrow paragraph) and Group 3 (order: paragraph \rightarrow static \rightarrow sentence \rightarrow line) spent approximately 9.6% and 4.7% less time per section than Group 4, respectively. These differences were not statistically significant, implying that the sequence in which participants encountered the methods did not systematically affect reading speed.

We also examined whether participants' self-reported favorite reading method was associated with their overall reading time. Participants who indicated a preference for one of the interactive highlighting techniques tended to spend slightly more time on sections overall than those who preferred static reading, but these differences were not significant. For instance, participants favoring line-by-line spent more time per section on average than those who favored static text ($p = 0.241$); those preferring sentence-by-sentence spent somewhat more time ($p = 0.667$), and those preferring paragraph-by-paragraph spent 9.6% more ($p = 0.438$). None of these effects reached statistical significance, indicating that preference was not reliably associated with faster or slower reading. As with comprehension, these patterns likely reflect subjective appeal rather than objective efficiency.

In summary, time-on-task remained consistent across all highlighting conditions and participant groups. Apart from the increased times for Sections 2 and 3 (possibly reflecting engagement or fatigue)—but notably not for Section 4, which was read at a similar pace to Section 1—hover-based highlighting neither significantly sped up nor slowed down reading. These results are summarized in Figure 5.

Analysis of the rating regarding attention-sustaining for reading methods. Beyond objective performance measures, participants rated each highlighting method on how it affected their ability to sustain attention, compared to normal static reading. These subjective evaluations (whether a method made it “harder,” “the same as regular,” or “easier” to maintain attention) differed significantly across the four methods (multi-




Figure 5. Time of the four methods.

nomial test, all $p < 0.05$). We conducted pairwise comparisons to identify which specific perceptions were driving these differences. The key findings for each method are as follows:

- ▶ **Line-by-line:** Participants' opinions on this method were polarized. A much larger proportion of participants rated the line-by-line highlighting as "harder" to sustain attention than regular reading, compared to those who said it was "about the same" (this difference was significant, $p < 0.001$). Likewise, significantly more participants rated line-by-line as "easier" than regular reading, compared to the "same" category ($p < 0.001$). In other words, although a slightly greater number found line-by-line highlighting more distracting than found it helpful, this difference was not statistically significant.
- ▶ **Sentence-by-sentence:** Participants generally perceived the sentence-by-sentence method as beneficial for sustaining attention. Significantly more participants rated it as "easier" compared to "harder" or "the same" ($p < 0.05$), suggesting that sentence-level highlighting was more often viewed as helpful than distracting or neutral. In summary, most readers found sentence-level highlighting helped them concentrate better, and a few found it detrimental; only the minority felt it was no different from normal reading.
- ▶ **Paragraph-by-paragraph:** Participants also viewed paragraph highlighting favorably. Significantly fewer participants rated the paragraph-by-paragraph method as "harder" compared to "easier" ($p < 0.05$). While the proportion who rated it "the same as regular" was lower than those who rated it "easier," this difference was not statistically significant ($p = 1.00$). Overall, this suggests that participants found paragraph highlighting either helpful or at least as effective as static reading, with relatively few perceiving it as more difficult.

In summary, these self-reported ratings indicate that sentence-by-sentence and paragraph-by-paragraph highlighting were generally seen as making it easier to sustain attention compared to regular static text. The line-by-line method, on the other hand, received mixed feedback—readers either loved it or disliked it in terms of attention, and a significantly higher number leaned toward it being more challenging rather than neutral. Overall, the interactive highlighting techniques (especially at the sentence and paragraph levels) were perceived to aid concentration for many users, even though our objective measures did not show a clear performance benefit. These results are presented in Table 5.

5. Discussion

This study addressed four research questions about hover-based highlighting in digital reading, examining effects on comprehension (RQ1), reading speed (RQ2), the relationship between user preferences and performance (RQ3), and readers’ evaluations of attention sustainability (RQ4). Our evaluation revealed no statistically significant differences in reading comprehension or reading speed between the interactive highlighting techniques and the control condition. In other words, no highlighting approach measurably improved how well participants understood the text or how quickly they read it. There were slight performance differences in raw scores—for example, one

Table 5. Results of the rating regarding attention-sustaining for reading methods using pairwise binomial test.

Reading method	Rating	Number of participants	Pairwise binomial test results: comparison between...		p-value
Line-by-line	Harder	41 (51.3%)	Harder	Same as regular	<0.001
	Same as regular	5 (6.2%)	Harder	Easier	1.00
	Easier	34 (42.5%)	Same as regular	Easier	<0.001
Sentence-by-sentence	Harder	19 (23.8%)	Harder	Same as regular	1.00
	Same as regular	17 (21.2%)	Harder	Easier	<0.05
	Easier	44 (55%)	Same as regular	Easier	<0.05
Paragraph-by-paragraph	Harder	14 (17.5%)	Harder	Same as regular	0.068
	Same as regular	30 (37.5%)	Harder	Easier	<0.05
	Easier	36 (45%)	Same as regular	Easier	1.00

highlighting technique yielded nominally higher comprehension accuracy and the other showed a minor difference in reading time—but these trends did not reach significance. These variations in comprehension or speed across conditions may reflect small and unreliable effects, though the lack of statistical significance could also indicate no genuine effect, potentially attributable to the sample size and the short length of the texts used in this study.

However, our results did show a clear effect on readers' subjective focus. Participants reported significantly higher attention-sustainment specifically when using sentence-by-sentence and paragraph-by-paragraph highlighting techniques, compared to the static condition. Participants had mixed perceptions of line-by-line highlighting, with some rating it as more distracting while others found it helpful. In particular, the highlighting condition that segmented text by grammatical units (such as sentences or paragraphs) received higher ratings for sustaining attention than the condition highlighting equal-length visual chunks (e.g., one line of text at a time). This suggests that while comprehension scores remained similar, the way text was highlighted had a meaningful impact on how well readers felt they could maintain focus. The grammatical segmentation highlight may have helped readers concentrate by presenting ideas in coherent units, leading to a statistically significant improvement in self-reported attention compared to the visually segmented highlights.

Participants' preferences further reflected these attention findings. In post-reading surveys, most readers preferred the grammatical segmentation highlighting technique over the purely visual line-by-line highlighting. Many commented that the grammatically segmented highlights made it easier to follow the narrative and stay engaged, whereas the line-oriented highlights sometimes felt arbitrary or distracting when a sentence spanned multiple lines. This clear user preference is consistent with the higher reported engagement and perceived focus associated with the grammatical highlighting conditions. In essence, readers tended to favor the technique that helped them feel more attentive and immersed in the content. Notably, participants' preferences for particular highlighting methods were not reliably associated with better performance, suggesting that subjective appeal and objective effectiveness may be distinct dimensions of reading support. Even methods that did not improve measurable performance were valued by some readers for keeping their reading on track, indicating an important benefit in terms of user experience.

Our results hint that highlighting by meaningful linguistic units (grammatical segments like a sentence or a paragraph) may better support readers' natural processing of text than highlighting by line. In contrast, visual segmentation (highlighting one line at a time) provided some attentional benefit for certain readers but was counterproductive for nearly as many—reflecting its polarized reception. The absence of signifi-

cant comprehension differences suggests that no highlighting method automatically improves understanding in a short reading task, but the improved perceived attentional engagement with grammatical highlights could be practically important in scenarios requiring prolonged concentration.

5.1. Future Research

Building on this work, future studies might employ more sensitive, robust or longer-term experimental designs to determine whether performance differences emerge under extended reading conditions or among specific populations such as language learners or readers with attentional challenges. Further experiments could also compare alternative interaction types (e.g., tapping or keyboard control) to assess usability across devices. These directions would provide a more robust basis for evaluating when and how interactive highlighting benefits reading performance. Additionally, objective measures like eye-tracking could be used in future experiments to validate whether highlighted segments truly guide visual attention and reduce mind-wandering. By addressing these questions, future research can deepen our understanding of how and when interactive highlighting can most effectively support digital reading.

5.2. Design Implications

Our findings carry implications for the design of digital reading applications and study tools. Designers might consider incorporating interactive highlighting features that align with readers' cognitive processing of text. In practice, highlighting entire sentences or paragraphs—rather than arbitrary visual blocks—could potentially support comprehension and sustained attention by presenting information in semantically coherent chunks. User preference for grammatical segmentation highlights in our study suggests that many readers appreciate highlights that follow the natural flow of language, finding them more engaging and less disruptive. Therefore, giving users some control over how text is highlighted—or intelligently choosing the method based on content type—could improve overall satisfaction and engagement with textual platforms. It is also important to remember that highlighting alone did not boost test performance in our study, so such features should be seen as aids to engagement and focus rather than guaranteed learning enhancers. Effective UX design for educational reading tools might combine interactive highlighting with other support (e.g., annotations, summaries, self-quizzing prompts) to convert the moment-to-moment focus that highlights provide into deeper processing of the material. In summary, digital reading systems stand to benefit from thoughtful highlight design: emphasizing meaningful text units, accommodating individual preferences, and ultimately using highlighting not as an end, but to foster active reading and sustained engagement.

6. Conclusion

This study investigated three interactive hover-based highlighting techniques—two aligned with grammatical text segmentation (sentence and paragraph) and one based on visual segmentation (line-by-line)—to examine how they influence reading comprehension (1) and reading speed (2) in a digital reading context. The results showed no significant differences in comprehension or reading speed across highlighting conditions, suggesting that hover-based highlighting does not directly enhance these measurable aspects of reading performance. This finding does not support the first two hypotheses predicting improved comprehension and slower reading due to the added interaction effort. Moreover, readers expressed a clear preference for the grammatical highlighting methods, indicating that they found it more helpful and less distracting than the line-by-line highlighter. These results suggest that while interactive highlighting may not directly boost immediate comprehension outcomes, it can enhance user engagement and focus during reading.

Implications of our findings point toward the value of aligning interface design with cognitive reading processes. Highlighting text in meaningful units (like sentences or paragraphs) appears to resonate better with readers, potentially aiding concentration without imposing a penalty on reading efficiency. Designers of e-reading tools and digital textbooks might explore integrating smart highlighting features that maintain or increase reader attention, as our findings suggest this could offer experiential benefits even without measurable performance gains. At the same time, educators and learners may consider that highlighting alone is unlikely to enhance learning, but it may support comprehension when used alongside active reading strategies. Future work will further clarify the role of interactive highlights in learning: for example, examining long-term retention, adapting highlighting to individual needs, and testing these techniques in real-world educational settings. By advancing our understanding of how highlighting format influences reading, this research lays the groundwork for developing more effective digital reading experiences that keep readers focused and engaged, ultimately supporting better comprehension over time.

7. References

- Barnhurst, K. G. (2013). Newspapers experiment online: Story content after a decade on the web. *Journalism*, 14(1), 3–21. <https://doi.org/10.1177/1464884912448898>
- Beier, S., Berlow, S., Boucaud, E., Bylinskii, Z., Cai, T., Cohn, J., Crowley, K., Day, S. L., Dingler, T., Dobres, J., Healey, J., Jain, R., Jordan, M., Kerr, B., Li, Q., Miller, D. B., Nobles, S., Papoutsaki, A., Qian, J., ... Wolfe, B. (2021). Readability research: An interdisciplinary approach (Version 1). *arXiv*. <https://doi.org/10.48550/ARXIV.2107.09615>
- Bessemans, A. M. M. (2012). *Letterontwerp voor kinderen met een visuele functiebeperking* [PhD thesis]. Leiden University. <https://hdl.handle.net/1887/20032>

- Carr, N. G. (2020). *The shallows: What the Internet is doing to our brains*. W.W. Norton & Company.
- Clinton-Lisell, V., Seipel, B., Gilpin, S., & Litzinger, C. (2023). Interactive features of E-texts' effects on learning: A systematic review and meta-analysis. *Interactive Learning Environments*, 31(6), 3728–3743. <https://doi.org/10.1080/10494820.2021.1943453>
- Cuddihy, E., & Spyridakis, J. H. (2012). The effect of visual design and placement of intra-article navigation schemes on reading comprehension and website user perceptions. *Computers in Human Behavior*, 28(4), 1399–1409. <https://doi.org/10.1016/j.chb.2012.03.002>
- Demidenko, E. (2004). *Mixed models: Theory and applications*. Wiley. <https://doi.org/10.1002/0471728438>
- Faraway, J. J. (2014). *Linear models with R* (Second edition). CRC Press, Taylor & Francis.
- Fok, R., Kambhamettu, H., Soldaini, L., Bragg, J., Lo, K., Head, A., Hearst, M. A., & Weld, D. S. (2022). Scim: Intelligent skimming support for scientific papers. *Arxiv*. <https://doi.org/10.48550/ARXIV.2205.04561>
- Hare, C., Johnson, B., Vlahiotis, M., Panda, E. J., Tekok-Kilic, A., & Curtin, S. (2024). Children's reading outcomes in digital and print mediums: A systematic review. *Journal of Research in Reading*, 47(3), 309–329. <https://doi.org/10.1111/1467-9817.12461>
- Head, A., Lo, K., Kang, D., Fok, R., Skjonsberg, S., Weld, D. S., & Hearst, M. A. (2021). Augmenting scientific papers with just-in-time, position-sensitive definitions of terms and symbols. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–18. <https://doi.org/10.1145/3411764.3445648>
- Hisgen, R. G. W., & van der Weel, A. (2022). *De lezende mens: De betekenis van het boek voor ons bestaan*. Uitgeverij Atlas Contact.
- Kaakinen, J. K., Papp-Zipernovszky, O., Werlen, E., Castells, N., Bergamin, P., Baccino, T., & Jacobs, A. M. (2018). Emotional and motivational aspects of digital reading. In M. Barzillai, J. Thomson, S. Schroeder, & P. van den Broek (Eds.), *Studies in written language and literacy* (Vol. 17, pp. 141–164). John Benjamins Publishing Company. <https://doi.org/10.1075/swll.17.06kaa>
- Mangen, A., & Van der Weel, A. (2016). The evolution of reading in the age of digitization: An integrative framework for reading research. *Literacy*, 50(3), 116–124. <https://doi.org/10.1111/lit.12086>
- Poulsen, P. B., Buchholz, P., Walt, J. G., Christensen, T. L., & Thygesen, J. (2005). Cost analysis of glaucoma-related-blindness in Europe. *International Congress series*, 1282, 262–266. <https://doi.org/10.1016/j.ics.2005.05.091>
- Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, 124(3), 372–422. <https://doi.org/10.1037/0033-2909.124.3.372>
- Salmerón, L., Strømso, H. I., Kammerer, Y., Stadler, M., & Van Den Broek, P. (2018). Chapter 4. Comprehension processes in digital reading. In M. Barzillai, J. Thomson, S. Schroeder, & P. van den Broek (Eds.), *Studies in written language and literacy* (Vol. 17, pp. 91–120). John Benjamins Publishing Company. <https://doi.org/10.1075/swll.17.04sal>
- Schwabe, A., Lind, F., Kosch, L., & Boomgaarden, H. G. (2022). No negative effects of reading on screen on comprehension of narrative texts compared to print: A meta-analysis. *Media Psychology*, 25(6), 779–796. <https://doi.org/10.1080/15213269.2022.2070216>
- Sieghart, S., & Gorbach, R. P. (Eds.). (2024). *Gutes design für leichte sprache: Theorie und praxis zur DIN SPEC 33429* (1st ed.). utb GmbH. <https://doi.org/10.36198/9783838563077>
- Smith, F. (2004). *Understanding reading* (6th ed.). L. Erlbaum Associates.
- Stothart, C., Mitchum, A., & Yehnert, C. (2015). The attentional cost of receiving a cell phone notification. *Journal of Experimental Psychology: Human Perception and Performance*, 41(4), 893–897. <https://doi.org/10.1037/xhp0000100>

- The Readability Consortium. (n.d.). Resources. <https://thereadabilityconsortium.org/resources/>
- Veras, R., Paluka, E., Chang, M.-W., Tsang, V., Shein, F., & Collins, C. (2014). Interaction for reading comprehension on mobile devices. *Proceedings of the 16th International Conference on Human-Computer Interaction with Mobile Devices & Services*, 157–161. <https://doi.org/10.1145/2628363.2628387>
- Wallace, S., Bylinskii, Z., Dobres, J., Kerr, B., Berlow, S., Treitman, R., Kumawat, N., Arpin, K., Miller, D. B., Huang, J., & Sawyer, B. D. (2022). Towards individuated reading experiences: Different fonts increase reading speed for different individuals. *ACM Transactions on Computer-Human Interaction*, 29(4), 1–56. <https://doi.org/10.1145/3502222>
- Wallace, S., Treitman, R., Kumawat, N., Arpin, K., Huang, J., Sawyer, B., & Bylinskii, Z. (2020). Individual differences in font preference & effectiveness as applied to interlude reading in the digital age. *Journal of Vision*, 20(11), 412. <https://doi.org/10.1167/jov.20.11.412>
- World Wide Web Consortium (W3C). (2018). Web content accessibility guidelines (WCAG) 2.1. <https://www.w3.org/TR/WCAG21/>
- Wylie, J., Thomson, J., Leppänen, P. H. T., Ackerman, R., Kannianen, L., & Prieler, T. (2018). Chapter 3. Cognitive processes and digital reading. In M. Barzillai, J. Thomson, S. Schroeder, & P. van den Broek (Eds.), *Studies in written language and literacy* (Vol. 17, pp. 57–90). John Benjamins Publishing Company. <https://doi.org/10.1075/swll.17.03wyl>

Authors

Szabolcs Vatány is a doctoral student at Moholy-Nagy University of Art and Design Budapest, where his research focuses on digital readability and typography. He holds an MA in Graphic Design and has several years of professional experience as a UX designer. His background spans interface design, type design, and multiple areas of graphic communication. His work often explores the intersection of design, disability, and education, including projects aimed at improving accessibility and developing educational technologies. He has also led university courses on related topics, integrating design practice with inclusive and research-based approaches.

Thi Huyen Nguyen is a Doctor of Preventive Medicine with an MSc in Epidemiology and a PhD candidate in Statistics. She has over eight years of combined experience in public health, infectious disease modeling, and clinical research. She specializes in statistical modeling, data analysis, and reproducible research workflows, with applications spanning public health and non-health domains such as cognitive accessibility and design perception. Her work bridges quantitative methodology with real-world problem solving, fostering interdisciplinary collaboration across academic, governmental, and international partners.

Anikó Illés, PhD, is a psychologist and aesthete, professor of Moholy-Nagy University of Art and Design Budapest. Her main research interest is about issues concerning visual culture, such as psychology of art (appreciation of artworks, creativity, analyzing artworks), pedagogical aspects of art and museums, analyzing visual images as a research method for several topics in the fields of education and social psychology (schooling, educational gap, perceived social mobility of children, national identity). She was the Secretary General of International Association of Empirical Aesthetics (2020–2025). She was a Fulbright researcher at Montclair State University, NJ, USA (2020).

Ann Bessemans (prof. dr.) is a leading expert in legibility and an award-winning typographic and type designer. She is the founder of READSEARCH, the legibility research group at PXL-MAD School of Arts and Hasselt University (Belgium), where she is affiliated with the Faculty of Architecture and Arts and the Data Science Institute (DSI). In addition to her academic work, she

also operates as an independent designer and consultant. At PXL she teaches typography and type design. Bessemans is the program director of the international Master program “Reading Type & Typography.” With a PhD (under the supervision of Gerard Unger) from Leiden and Hasselt Universities, her data-driven legibility research, supported by multiple Microsoft grants, aims to improve typeface design for diverse reading needs, including those with impairments. She actively engages in various academic and research roles, actively participating in the board of directors at ATypI.



Why Meta-Font Struck a Nerve

Mike Zender

College of Design, Architecture, Art, and Planning, University of Cincinnati, Cincinnati, OH, USA
(zenderpm[at]ucmail.uc.edu)

What's old is old. But what's new is also old. One could argue that there are a limited number of concepts in our world and that time simply rehashes new iterations of existing concepts. And while one might argue that's not proof that it's true, Douglas R. Hofstadter's article reprinted following this introduction is one data point of evidence to support that what's new really is old. It is 44 years old but explains something important unfolding today: AI.

Reprint of this article was instigated by a routine figure reproduction request from an old article, in this case, Hofstadter's 1982 article "Metafont, Metamathematics, and Metaphysics: Comments on Donald Knuth's 1982 Article "The Concept of a Meta-Font"" from *Visible Language* 16.4 (see Figure 1). My memory was not good enough to recall much of Hofstadter's article at first, so I re-read it. When I did, I immediately asked him for permission to reprint it because even though it was from years ago it so directly addressed essential issues concerning AI today. To put that in perspective, Apple computer was founded just six years prior to Hofstadter's insightful article and is today the company with either the world's largest or second largest market cap depending on NVIDIA's, and NVIDIA is an AI company. To say that 44 years ago Hofstadter foresaw issues of at least economic significance, is a vast understatement.

Hofstadter (1982/2026) articulated in understandable language AI challenges that most of us only feel intuitively. He makes clear, in a non-trivial way, that machines are limited by being... machines. He explains that starting with a machine is to start with a very limited set of assumptions, mechanisms, and relationships with one target in mind and because one target is the focus other targets are of necessity excluded. But what's excluded may be critical. He said, "No matter how many new knobs—or even families

@: [ISSUE](#) > [ARTICLE](#) >

Cite this article:

Zender, M. (2026). Why Meta-Font struck a nerve. *Visible Language*, 60(1), 84–89.

First published online April 26, 2026.

© 2026 Visible Language — this article is **open access**, published under the CC BY-NC-ND 4.0 license.

<https://visible-language.org/journal/>

Visible Language Consortium:

University of Leeds (UK)

University of Cincinnati (USA)

North Carolina State University (USA)

extremely extended or extremely condensed without changing the heights or widths of the strokes. One can also imitate a typewriter by extending or condensing the individual characters so that each one has the same width. Note that the length of serifs is proportional to the width, so that an i has much longer serifs than an m in the typewriter style.

Of course we get a much better imitation of a typewriter when the distinction between thick and thin strokes disappears. Such a font looks typewriter-like even when its letters do not all have the same width.

The letters of Computer Modern are all drawn by pens having an elliptical nib; for example, the thick strokes of the h's in this sentence were made by a pen that would look like '—' if enlarged ten times. The ellipses have perfectly horizontal axes, not tipped as ' / ', because the letters are intended to have vertical stress. Different pens are used to draw different parts of the letters.

Five parameters control the dimensions of these elliptical pens: One for the thin hairlines, another for thick stem lines that are straight, another for thick stem lines that are curved, another for the bulbs on letters like aef..y, and another that gives an aspect ratio between horizontal and vertical dimensions. The height of the hairline pen is used also as the height of the pens that draw the thick vertical stem lines. If the first four of these pen-width parameters are equal and if the aspect ratio is 1/1, the pens will be perfect circles.

9 Knuth / Concept of a Meta-Font

An ellipse like '—' has an aspect ratio of 1/3, while the aspect ratio of ' / ' is 3/1. It is interesting to see what happens when sans-serif letters are drawn with pens of different aspect ratios:

- A pen of aspect 1/3 generated these letters.
- A pen of aspect 2/3 generated these letters.
- A pen of aspect 1/1 generated these letters.
- A pen of aspect 3/2 generated these letters.
- A pen of aspect 3/1 generated these letters.

The aspect ratio can also be varied when the pens have different widths and serifs are present; in this case the aspect affects the darkness of letters like g and s that have thick horizontal strokes:

- A pen of aspect 1/3 generated these letters.
- A pen of aspect 2/3 generated these letters.
- A pen of aspect 1/1 generated these letters.
- A pen of aspect 3/2 generated these letters.
- A pen of aspect 3/1 generated these letters.

(In the examples above, the widths of thick vertical stems for aspect ratios less than 1 are equal to the heights of thick horizontal stems for aspect ratios greater than 1.)

Special care is needed in the choices of the pen-width parameters. For example, undesirable blotches appear when the bulbs are too large for the stems; and the type has a disturbing inconsistency when the curved stems are substantially wider than the straight ones. **A font cannot get too bold without having portions of the letters run into each other.** Perhaps future meta-fonts will be

10 Visible Language XVI 1 1982

Serif details can be varied in several ways. For example, there are no 'sheared' serifs on the letters in this sentence. And the letters you are now reading have thrice as much shear as usual, just to make sure that the concept of shear is clear. Another serif-oriented concept is the amount of 'bracketing'; the serifs in this sentence have no brackets. But the brackets are exaggerated in this sentence, so the serifs appear darker. The difference can be understood most easily if we enlarge the letters:

- n**o bracketing;
- N**ormal bracketing;
- n**oticeable bracketing.

A curve that starts at the edge of the serif will be tangent to the stem at some distance above or below the serif; this vertical distance is the 'bracketing' parameter.

A third parameter affecting serifs is called the 'crispness': The example serifs above have been crisply squared off, using a special rectangular pen instead of an ellipse, but one can also specify

no crispness,

in which case only the elliptical pens are used. The typewriter-like font examples above are non-crisp.

The length of serifs is, of course, controllable too. The letters in this sentence have serifs that are 50% shorter than before. And in this sentence they are 50% longer than before—so long

12 Visible Language XVI 1 1982

change is quite dramatic—it is something like the gradual changes in our own faces as we grow older, except that this typeface is getting younger.

Hundreds of typefaces have appeared in this article, yet all of them belong to the Computer Modern Roman and Italic meta-fonts. Each letter has been specified by a computer program written in the METAFONT language, and the computer can draw any desired variant of that letter when the parameter values have been supplied. It is important to remember that none of these conventions and parameters are built into METAFONT itself; METAFONT is a general-purpose language intended to facilitate the design of meta-fonts, and Computer Modern is but one approach to font design using such a language.

Let us take a brief look at the program for the letter h, since this will give some insight into the way a meta-font can

17 Knuth / Concept of a Meta-Font

Figure 1. Pages 9, 10, 12, and 17 from Knuth (1982).

of knobs—you add to your... machine, you will have left out some possibilities” (p. 104). In the back of our minds, we fear that non-machine possibilities not only matter, they may be the most significant parts. The beauty of Hofstadter’s article for designers facing an AI world is that it’s centered on letterforms:

Clearly there is much more going on in typefaces than meets the eye—literally. The shape of letterforms is a surface manifestation of deep mental abstractions. It is determined by conceptual considerations and balances that no set of merely geometric knobs could capture. Underneath or behind each instance of **A** there lurks a concept... (p. 109)

He moves on to say that the spirit of a letter is defined not by geometry but by roles that visual forms play and he doesn’t stop at that, he describes that roles are modular and that they “overlap and mingle in a subtle way” within a context (p. 109).

Hofstadter’s thoughts about “deep mental abstractions” reinforce several lines of thinking about typefaces and design that have developed in the past 20 years that I am personally familiar with. One is the concept of a letterform skeleton that embodies the essential qualities of each letter. The starting point for Donald Knuth’s Meta-Font to which Hofstadter was responding, was a pen stroke. By providing his parameterized letterform generating system with a pen stroke Knuth bypassed the important question of what exactly an **A** is. At the heart of a typeface generator lies some definition of the features that constitute **A**-ness or, as Hofstadter might say, the essential visual “roles” needed for an **A** to be recognized as an **A**. While many would argue that we still don’t know exactly what defines an **A** (a good opportunity for research!), the last 100 years have seen some progress toward such a definition.

Edward Johnson’s 1906 book *Writing, Illuminating, and Lettering* described a letterform “skeleton or structural plan” as the essence of a letterform and illustrated this using single line pencil strokes as capturing letters “essential forms” (Johnson, 1906/1948, p. 240). In a sense, Knuth’s “pen-stroke” was simply putting Johnson into a computer. Several decades later, Adrian Frutiger, the designer of many prominent 20th century typefaces (Univers, Frutiger, and Avenir among them), defined the most elemental shapes of letterforms by overlapping the letterforms of “the most widely used typefaces in the world” in his 1979 book *Type Sign Symbol*. Frutiger said that the areas of overlap in all the different typefaces defined “a kind of basic skeleton” that had engraved themselves “in the subconscious of the reader as a kind of elemental form” (Frutiger, 1979, p. 64). Building upon Johnson and Frutiger in my 2019 *Visible Language* article “Letterform Legibility and Visual Perception: A Speculation” (see Figure 2), I speculated that the essential features of letterforms, that is their “skeletons,” mapped nicely onto distinct combinations of early features detected in the process of visual perception. Those early features are: horizontal, vertical, angular, diagonal, circle, open-form,

Visible Language 53 . 3 .

Letterform Legibility and Visual Perception: a speculation

Mike Zender

"Roman capital letters first achieved the forms we know today about AD 100. ... At their most formal they are based on very simple geometric shapes, symbols for the sounds in a language. And each letter is successful as a symbol because its shape is hard to confuse with the others and is easy to memorize."
(Sutton & Bartram, 1968, p. 6)

This short paper explores a straightforward insight: that the basic features of visual perception map instructively onto the letterform skeletons of the Latin alphabet. Linking findings from visual perception with knowledge about typography and reading might advance our knowledge of how letterforms function visually. This knowledge could be used to develop a formal measure of letterform legibility, to provide means to distinguish between a text and a display typeface, and to provide guidance for typeface design.

Keywords: letterform, visual perception, letterform skeleton, legibility

Zender
Letterform Legibility and Visual Perception

LETTERFORMS

10 popular typefaces overlapped, to reveal the cap "E" letterform skeleton.

PERCEPTUAL FEATURES

diagonal curve vertical horizontal

Four basic perceptual features: diagonal, curve, vertical, horizontal.

70

Visible Language 53 . 3 .

FIGURE 2

The process of object recognition: an impression falls on the retina, basic features are sorted, patterns are assembled, features are compared to stored activation patterns, then an object is recognized.

76

Zender
Letterform Legibility and Visual Perception

Figure 2. Pages 70/71 and 76/77 from Zender (2019).



Figure 3. Pages 82/83 from Zender (2019).

closed-form (Figure 3). I guessed that each letterform activated the base mechanisms of visual perception and that various combinations of these essential features may be part of what makes an A an A versus an O.

Hofstadter’s article invokes the suggestion that “roles,” perhaps in the sense of the interaction of basic visual forms like key perceptual features, are what define a letterform more than mere mechanical geometry. Abstracted further, it may be that verbs, actions of interaction, are more important conceptually than physical geometry to what makes an A an A. If true, then one might create an AI typeface generator based on the interaction of key features rather than on Knuth’s pen strokes or geometric mechanisms. Why someone would want to do this is another question, as is the part the designer plays in this whole AI generated typeface enterprise.

Because Hofstadter’s article is brilliant, he also grounds this with Gödel’s theorem that any system cannot be both complete and provable within itself and links that to the concepts of *completeness*—a system which presents every true positive—and *consistency*—a system with no false positives. These concepts are so critical to defining current AI system performance. If you’re not an AI person, forgive the lingo. We do not wish to over-simplify a very deep article that is none the less highly readable. After all,

Hofstadter won the Pulitzer prize for his 1979 book on Gödel and the Los Angeles Times Book Prize for his 2007 *I Am a Strange Loop*.

In addition to being pleased to have a track record of publishing concepts that stand the test of time, *Visible Language* is very honored to have had people the stature of Douglas R. Hofstadter and Donald E. Knuth contribute as authors in the past (though we do not include a reprint of Knuth's article here, it's available in our archive). We're also so honored to have had letter writers of such high stature as those who sent in additional letters in response to Knuth's article: Bigelow, Karow, Unger, and Zapf (Baudin et al., 1982; summarized in Littlejohn, 2026). They are all truly world-class thinkers as their contributions demonstrate. We re-present their work here not in the spirit of resting on our laurels but as intellectuals ourselves who believe that connecting concepts across time is something every scholar, and scholarly journal, should strive to do.

References

- Baudin, F., Bigelow, C., Bronsard, H.-P., Fisher, E., Jr., Ford, D., Gore, G., Jaspert, W. P., Kapr, A., Karow, P., Nesbitt, A., Rondthaler, E., Schappler, J., Tracy, W., Ungar, G., Zapf, H., & Knuth, D. E. (1982). Other replies to Donald E. Knuth's article, "The Concept of a Meta-Font" [V. Mollar & K. Schmidt, Trans.]. *Visible Language*, 16(4), 339–359. <https://journals.uc.edu/index.php/vl/article/view/5350>
- Frutiger, A. (1979). *Type sign symbol*. ABC Verlag.
- Hofstadter, D. R. (1979). *Gödel, Escher, Bach: An eternal golden braid*. Basic Books.
- Hofstadter, D. R. (1982/2026). Metafont, metamathematics, and metaphysics: Comments on Donald Knuth's article "The Concept of a Meta-Font" [Reprint]. *Visible Language*, 60(1), 90–119. <https://www.visible-language.org/journal/issue-60-1-metafont-comments>
- Hofstadter, D. R. (2007). *I am a strange loop*. Basic Books.
- Johnson, E. (1906/1948). *Writing, illuminating, and lettering*. Sir Isaac Pitman & Sons, Ltd.
- Knuth, D. E. (1982). The concept of a meta-font. *Visible Language*, 16(1), 3–27. <https://journals.uc.edu/index.php/vl/article/view/5329/4193>
- Littlejohn, D. (2026). A profession provoked: How Meta-Font struck a nerve. *Visible Language*, 60(1), 120–123. <https://www.visible-language.org/journal/issue-60-1-how-metafont>
- Zender, M. (2019). Letterform legibility and visual perception: A speculation. *Visible Language*, 53(3), 70–91. <https://journals.uc.edu/index.php/vl/article/view/4655/3511>



Metafont, Metamathematics, and Metaphysics: Comments on Donald Knuth’s Article “The Concept of a Meta-Font”

Douglas R. Hofstadter

Department of Philosophy, University of Indiana, Bloomington, IN, USA (dughof[at]iu.edu)

Abstract: It is argued that readers are likely to carry away from Donald Knuth’s article “The Concept of a Meta-Font” a falsely optimistic view of the extent to which the design of typefaces and letterforms can be mechanized through an approach depending on describing letterforms by specifying the settings of a large number of parameters. Through a comparison to mathematical logic, it is argued that no such set of parameters can capture the essence of any semantic category. Some different ways of thinking about the problem of the “spirit” residing behind any letterform are suggested, connecting to current [c. 1982] research issues in the field of artificial intelligence.

Keywords: artificial intelligence; font design; letter spirit; letterforms; meta-font; parametric design; typographic parameters

Prefatory note to the reprint. Following is a reprint of an article by Douglas Hofstadter that first appeared in 1982 in *Visible Language* issue 16.4, in response to another 1982 article by Donald Knuth from issue 16.1. It is presented here with minimal edits that integrate it into the current *Visible Language* typesetting, with most adjustments made for citation style. The reprint is preceded in this issue’s layout by an introductory editorial by Mike Zender (2026). It is followed by editorial by Deborah Littlejohn (2026), which summarizes additional responses to Knuth by other authors (Baudin et al., 1982). These responses (all as letters to the editor) appeared in issue 16.4 alongside Hofstadter’s response.

@: [ISSUE](#) > [REPRINT](#) > [ORIGINAL](#) >

Cite this article:

Hofstadter, D. R. (1982/2026). Metafont, metamathematics, and metaphysics: Comments on Donald Knuth’s article “The Concept of a Meta-Font” [Reprint]. *Visible Language*, 60(1), 90–119.

First published online April 26, 2026.

© 1982/2026 Visible Language — this article is **open access**, published under the CC BY-NC-ND 4.0 license.

<https://visible-language.org/journal/>

Visible Language Consortium:

University of Leeds (UK)

University of Cincinnati (USA)

North Carolina State University (USA)

1. The “Mathematization of Categories” and Metamathematics

Donald Knuth has spent the past several years working on a system allowing him to control many aspects of the design of his forthcoming books—from the typesetting and layout down to the very shapes of the letters! Never has an author had anything remotely like this power to control the final appearance of his or her work. Knuth’s TEX typesetting system has become well known and available in many countries around the world. By contrast, his Metafont system for designing families of typefaces has not become as well known or as available.

In his article “The Concept of a Meta-Font” (Knuth, 1982), Knuth sets forth for the first time the underlying philosophy of Metafont, as well as some of its products. Not only is the concept exciting and clearly well executed, but in my opinion the article is charmingly written as well. However, despite my overall enthusiasm for Knuth’s idea and article, there are some points in it that I feel might be taken wrongly by many readers, and since they are points that touch close to my deepest interests in artificial intelligence and esthetic theory, I felt compelled to make some comments to clarify certain important issues raised by “The Concept of a Meta-Font.”

Although his article is primarily about letterforms, not philosophy, Knuth holds out in it a philosophically tantalizing prospect for us: that with the arrival of computers, we can now approach the vision of a unification of all typefaces. This can be broken down into two ideas: (1) that underneath all A’s there is just one grand, ultimate abstraction that can be captured in a finitely parametrizable computational structure—a “software machine” with a finite number of “tunable knobs” (we could say “degrees of freedom” or “parameters,” if we wished to be more dignified); and (2) that every conceivable particular A is just a product of this machine with its knobs set at specific values.

Beyond the world of letterforms, Knuth’s vision extends to what I shall call the “mathematization of categories”: the idea that any abstraction or Platonic concept can be so captured—i.e., as a software machine with a finite number of knobs. (For more on this notion, see Hofstadter, 1982b.) Knuth gives only a couple of examples—those of the “meta-waltz” and the “meta-shoe”—but by implication one can imagine a “meta-chair,” a “meta-person,” and so forth.

This is perhaps carrying Knuth’s vision further than he ever intended. Indeed, I suspect so; I doubt that Knuth believes in the feasibility of such a “mathematization of categories” opened up by computers. Yet any imaginative reader would be likely to draw hints of such a notion out of Knuth’s article, whether Knuth intended it that way or not. It is my purpose in this article to argue that such a vision is exceedingly unlikely to come about, and that such intriguingly flexible tools as meta-shoes, meta-fonts, modern electronic organs (with their “oom-pah-pah” and “cha-cha-cha” rhythms and

their canned harmonic patterns), and other many-knobbed devices will only help us see more clearly why this is so. The essential reason for this I can state in a very short way: I feel that to fill out the full “space” defined by a category such as “chair” or **A** or “waltz” is an act of infinite creativity, and that no finite entity (inanimate mechanism or animate organism) will ever be capable of producing all possible **A**'s and nothing but **A**'s (the same could be said for chairs, waltzes, etc.).

I am not making the trivial claim that, because life is finite, nobody can make an infinite number of creations; I am making the nontrivial claim that nobody can possess the “secret recipe” from which all the (infinitely many) members of a category such as **A** can in theory be generated. In fact, my claim is that no such recipe exists. Another way of saying this is that even if you were granted an infinite lifetime in which to draw all the **A**'s you could think up, thus realizing the full potential of any recipe you had, no matter how great it might be, you would still miss vast portions of the space of **A**'s.

In metamathematical terms this amounts to positing that any conceptual (or “semantic”) category is a “productive” set, a precise notion whose characterization is a formal counterpart to the description in the previous paragraphs (namely, a set whose elements cannot be totally enumerated by any effective procedure without overstepping the bounds of that set, but which can be approximated more and more fully by a sequence of increasingly complex effective procedures). The existence and properties of such sets first became known as a result of Gödel's Incompleteness Theorem of 1931 (Gödel, 1931/1962). It is certainly not my purpose here to explain this famous result, but a short synopsis might be of help. (Other useful references are: Chaitin, 1975; DeLong, 1970; Hofstadter, 1979; Nagel & Newman, 1958; Rucker, 1982; Smullyan, 1961, 1978.)

2. An Intuitive Picture of Gödel's Theorem

Gödel was investigating the properties of purely formal deductive systems in the sphere of mathematics, and he discovered that such systems—even if their ostensible domain of discourse was limited to one topic—could be viewed as talking “in code” about themselves. Thus a deductive system could express, in its own formal language, statements about its own capabilities and weaknesses. In particular, System X could say of itself through the Gödelian code, “System X is not powerful enough to demonstrate the truth of Sentence S.” It sounds a little bit like a science-fiction robot called “Robot 15” droning in a telegraphic monotone, “Robot-15 unfortunately unable to complete Task T-12—very sorry.” Now what happens if Task T-12 happens, by some crazy coincidence, to be not the assembly of some strange cosmic device but merely the act of uttering the preceding telegraphic monotone? (I say “merely” but of course that is a bit ironic.) Then Robot-15 could get only partway through the sentence before choking: “Robot-15 unfortunately unable to comp—.”

Now in the case of a formal system, System X, talking about its powers, suppose that Sentence G, by an equally crazy coincidence, is the one that says, “System X is regrettably not powerful enough to demonstrate the truth of Sentence G.” In such a case, Sentence G is seen to be an assertion of its own unprovability within System X. In fact we do not have to rely on crazy coincidences, for Gödel showed that given any reasonable formal system, a G-type sentence for that system actually exists. (The only exaggeration in my English-language version of G is that in formal systems there is no way to say “regrettably.”) In formal deductive systems this foldback takes place of necessity by means of a Gödelian code, but in English no Gödelian code is needed and the peculiar quality of such a loop is immediately visible.

If you think carefully about Sentence G, you will discover some amazing things. Could Sentence G be provable in System X? If it were, then System X would contain a proof for Sentence G, which asserts that System X contains no proof for Sentence G. Only if System X is blatantly self-contradictory could this happen—and a formal reasoning system that is self-contradictory is no more useful than a submarine with screen doors. So, provided we are dealing with a consistent formal system (one with no self-contradictions), then Sentence G is not provable inside System X. And since this is precisely the claim of Sentence G itself, we conclude that Sentence G is true—true but unprovable inside System X.

One last way to understand this curious state of affairs is afforded the reader by this small puzzle. Choose the more accurate of the following pair of sentences:

1. Sentence G is true *despite* being unprovable.
2. Sentence G is true *because* it is unprovable.

You’ll know you’ve really caught on to Gödel when both versions ring equally true to your ears, when you flip back and forth between them, savoring that exceedingly close approach to paradox that G affords. That’s how twisted back on itself Sentence G is!

The main consequence of G’s existence within each System X is that there are truths unattainable within System X, no matter how powerful and flexible System X is, as long as System X is not self-contradictory. Thus, if we look at truths as objects of desire, no formal system can have them all; in fact, given any formal system we can produce on demand a truth that it cannot have, and flaunt that truth in front of it with taunting cries of “Nah, nah!” The set of truths has this peculiar and infuriating quality of being uncapturable by any finite system, and worse, given any candidate system, we can use what we know about that system to come up with a specific Gödelian truth that eludes provability inside that system.

By adding that truth to the given system, we come up with an enlarged and slightly more powerful system—yet this system will be no less vulnerable to the Gödelian devilry than

its predecessor was. Imagine a dike that springs a new leak each time the proverbial Dutch boy plugs up a hole with his finger. Even if he had an infinite number of fingers, that leaky dike would find a spot he hadn't covered. A system that contains at least one unprovable truth is said to be "incomplete," and a system that not only contains such truths but that cannot be rescued in any way from the fate of incompleteness is said to be "essentially incomplete." Another name for sets with this wonderfully perverse property is "productive" (Rogers, 1967).

My claim—that semantic categories are productive sets—is, to be sure, not a mathematically provable fact but a metaphor. This metaphor has been used by others before me—notably, the logicians Emil Post (1944) and John Myhill (1952)—and I have written of it myself before (Hofstadter, 1979, 1982a).

3. Completeness and Consistency

Note that it is important to have the potential to fill out the full (infinite) space, and equally important not to overstep it. However, merely having infinite potential is not by any means equivalent to filling out the full space. After all, any existing Metafont **A**-schema—even one having just one degree of freedom!—will obviously give us infinitely many distinct **A**'s as we sweep its knob (or knobs) from one end of the spectrum to the other. Thus to have an **A**-making machine with infinite variety of potential output is not in itself difficult; the trick is to achieve *completeness*: to fill the space.

And yet, isn't it easy to fill the space? Can't one easily make a program that will produce all possible **A**'s? After all, any **A** can be represented as a pattern of pixels (dots that are either off or on) in an $m \times n$ matrix—hence a program that merely prints out all possible combinations of pixels in matrices of all sizes (starting with 1×1 and moving upwards to 2×1 , 1×2 , 3×1 , 2×2 , 1×3 , etc., as in Georg Cantor's famous enumeration of the rational numbers) will certainly cover any given **A** eventually. This is quite true. So what's the catch?

Well, unfortunately, it is hard—very hard—to write a screening program that will retain all the **A**'s in the output of this pixel-pattern program, and at the same time will reject all **K**'s, pictures of frogs, octopi, grandmothers, and precognitive photographs of traffic accidents in the twenty-fifth century (to mention just a few of the potential outputs of the generation program). The requirement that one must stay within the bounds of a conceptual category could be called *consistency*—a constraint complementary to that of completeness.

In summary, what might seem desirable from a knobbed category-machine is the joint attainment of two properties—namely, (1) completeness: that all true members of a category (such as the category of **A**'s [Figure 1] or the category of human faces [Figure 2])



Figure 1. The category of A's (drawn from Letraset, 1981). [Editor's note: This figure was used for the cover of *Visible Language* issue 16.4. All figures in this reprint were captured from a printed copy of the original publication, resulting in considerable imperfections.]



Figure 2. The category of human faces (drawn from Strich, 1981). [Editor’s note: The faces shown here are from Italian director Federico Fellini’s film archive photographs, and thus represent a narrower range of human faces than the range of typefaces in Figure 1.]

should be potentially producible eventually as output; and (2) consistency: that no false members of the category (“impostors”) should ever be potentially producible. In short, that the set of outputs of the machine should coincide exactly with the set of members of the intuitive category.

The twin requirements of consistency and completeness are metaphorical equivalents of well-known notions by the same names in metamathematics, denoting desirable properties of formal systems (theorem-producing machines)—namely, (1) completeness: that all true statements of a theory (such as the theory of numbers or the theory of sets) should be potentially producible eventually as theorems; and (2) consistency:

that no false statements of the theory should ever be potentially producible. In short, that the set of theorems of the formal system should coincide exactly with the set of truths of the informal theory.

The import of Gödel's Incompleteness Theorem is that these two idealized goals are unreachable simultaneously for any "interesting" theory (where "interesting" really means "sufficiently complex"); nonetheless, one can approach the set of truths by stages, using increasingly powerful formal systems to make increasingly accurate approximations. The goal of total and pure truth is, however, as unreachable by formal methods as is the speed of light by any material object. I suggest that a parallel statement holds for any "interesting" category (where again, "interesting" means something like "sufficiently complex," although it is a little harder to pin down): namely, one can do no better than approach the set of its members by stages, using increasingly powerful knobbed machines to make increasingly accurate approximations.

Intuition at first suggests that there is a crucial difference between the (metamathematical) result about the nonformalizability of truth and the (metaphorical) claim about the nonmechanizability of semantic categories; this difference would be that the set of all truths in a mathematical domain such as set theory or number theory is objective and eternal, whereas the set of all **A**'s is subjective and ephemeral. However, on closer examination, this distinction begins to blur quite a bit. The very fact of Gödel's proven nonformalizability of mathematical truth casts serious doubt on the objective nature of such truth. Just as one can find all sorts of borderline examples of **A**-ness, examples that make one sense the hopelessness of trying to draw the concept's exact boundaries, so one can find all sorts of borderline mathematical statements that are formally undecidable in standard systems and which, even to a keen mathematical intuition, hover between truth and falsity. And it is a well-known fact that different mathematicians hold different opinions about the truth or falsity of various famous formally undecidable propositions (the axiom of choice in set theory is a classic example). Thus, somewhat counterintuitively, it turns out that mathematical truth has no fixed and eternal boundaries, either. And this suggests that perhaps my metaphor is not so much off the mark.

4. A Misleading Claim for Metafont

Whatever the validity and usefulness of this metaphor, I shall now try to show some evidence for the viewpoint that leads to it, using Metafont as a prime example of a "knobbed category machine." In his article, Knuth comes perilously close, in one throwaway sentence, to suggesting that he sees Metafont as providing us with a mathematization of categories. I doubt he suspected that anyone would focus in on

that sentence as if it were the key sentence of the article—but as he did write it, it's fair game! That sentence ran:

The ability to manipulate lots of parameters may be interesting and fun, but does anybody really need a 6 1/7-point font that is one fourth of the way between Baskerville and Helvetica? (Knuth, 1982, p. 19)

This rhetorical question is fraught with unspoken implications. It suggests that Metafont as it now stands (or in some soon-available or slightly modified version) is ready to carry out, on demand, for any user, such an interpolation between two given typefaces. There is something very tricky about this proposition that I suspect most readers will not notice: it is the idea that jointly parametrizing two typefaces is no harder, no different in principle, from just parametrizing one typeface in isolation.

Indeed, to many readers, it would appear that Knuth has actually carried out such a joint parametrization. After all, in printing Psalm 23 (Figure 3), didn't he move from an old-fashioned, compact, serified face with relatively tall ascenders and descenders

The LORD is my shepherd;
 I shall not want.
 He maketh me to lie down
 in green pastures:
 he leadeth me
 beside the still waters.
 He restoreth my soul:
 he leadeth me
 in the paths of righteousness
 for his name's sake.
 Yea, though I walk through the valley
 of the shadow of death,
 I will fear no evil:
 for thou art with me;
 thy rod and thy staff
 they comfort me.
 Thou preparest a table before me
 in the presence of mine enemies:
 thou anointest my head with oil,
 my cup runneth over.
 Surely goodness and mercy
 shall follow me
 all the days of my life:
 and I will dwell
 in the house of the LORD
 for ever.

Figure 3. Donald Knuth's virtuoso Metafont rendition of Psalm 23, in which the font for each character is determined by the settings of 28 knobs, all of which change slowly but steadily as the psalm progresses.

and small x-height all the way to the other end of the spectrum: a modern-looking, extended, sans-serif face with relatively short ascenders and descenders and large x-height? Yes, of course—but the critical omitted point here is that these two ends of the spectrum were not pre-existing, prespecified targets; they just happened to emerge as the extreme products of a knobbed machine designed so that one more or less intermediate setting of its knobs would yield a particular target typeface (Monotype Modern Extended 8A).

In other words, this particular set of knobs was inspired solely and directly by an attempt to parametrize one typeface (Monotype Modern). The two extremes shown in the psalm are both variations on that single theme; the same can be said of every intermediate stage as well. There is only one underlying theme here (Monotype Modern), and a cluster of several hundred variants of it, each one of which is represented by a single character. The psalm does not represent the marriage of two unrelated families, but simply exhibits many members of one large family.

5. Joint Parametrization of Two Typefaces: A Far Cry from Parametrizing One Typeface

You can envision all the variants of Monotype Modern produced by twiddling the knobs on this particular machine as constituting an “electron cloud” surrounding a single “nucleus” (Figure 4a). Now by contrast, joint parametrization of two pre-existent, known typefaces (say, Baskerville and Helvetica as Knuth suggests [Figure 5]), would be like a cloud of electrons swarming around two nuclei, like a chemical bond (Figure 4b).

In order to jointly parametrize two typefaces in Metafont, you would need to find, for each pair of corresponding letters (say Baskerville **a** and Helvetica **a**) a set of discrete geometric features (line segments, serifs, extremal points, points of curvature shift, etc.) that they share and that totally characterize them. Each such feature must be equated with one or more parameters (knobs), so that the two letterforms are seen as




Figure 4. (a) An electron cloud surrounding a single nucleus; (b) A cloud of electrons around two nuclei, like a chemical bond.



Figure 5. Baskerville (above) and Helvetica.

produced by specific settings of their shared set of knobs. Moreover, all intermediate settings must also yield valid instances of the letter **a**. That is the very essence of the notion of a knobbed machine, and it is also the gist of the quote, of course: that we should now (or soon) be able to interpolate between any familiar typefaces merely by knob-twiddling.

Now I will admit that I think it is perhaps feasible—though much more difficult than parametrizing a single typeface—to jointly parametrize two typefaces that are not radically different. It is not trivial, to cite just one sample difficulty, to move between Baskerville’s round dot over the **i** to Helvetica’s square dot—but it is certainly not inconceivable. Conversely, it is not inconceivable to move between the elegant swash tail of the Baskerville **Q** and the stubby straight tail of the Helvetica **Q**—but it is certainly not trivial (Figure 6).

Moving from letter to letter and comparing them will reveal that each of these two typefaces has features that the other totally lacks. (You should disregard lowercase **g**, since the **g**’s of our two typefaces are as different from each other as Baskerville **B** is from Helvetica **H**; in both cases the two letterforms being compared derive from entirely different underlying “Platonic essences.” (It is Metafont’s purpose to mediate between different stylistic renditions of a single Platonic essence, not between distinct Platonic essences.) Presumably, in a case where one typeface possesses some distinct feature that the other totally lacks, there is a way to fiddle with the knobs that will make the feature nonexistent in one but present in the other. For instance, a knob



Figure 6. Blow-ups of Baskerville and Helvetica **i** and **Q**.

setting of zero might make some feature totally vanish. Sometimes it will be harder to make features disappear—it might require several knobs to have coordinated settings. Nonetheless, despite all the complex ways that Baskerville and Helvetica differ, I repeat, it is conceivable that somebody with great patience and ingenuity could jointly parametrize Helvetica and Baskerville. But the real question is this: Would such a joint parametrization easily emerge out of two separate, independently carried-out parametrizations of these typefaces? Hardly!

The Baskerville knobs do not contain even a hint of the Helvetica qualities—or the reverse. How can I convince you of this? Well, just imagine how great the genius of John Baskerville would have had to be for his design to have implicitly defined another typeface—and a typeface only discovered (or invented) two centuries later! To see this more concretely, imagine that someone who had never seen Helvetica naively created a Metafont rendition of Baskerville (that is, a metafont centered on Baskerville in the same sense as Knuth's sample metafont is centered on Monotype Modern). Now imagine that someone else who does know Helvetica comes along, twiddles the knobs of this Baskerville metafont, and actually produces a perfect Helvetica! It would be nearly as strange as having a marvelous music-composing program based exclusively on the style of G. F. Handel (who composed in England in a baroque, elegant 18th-century style) that was later discovered, totally unexpectedly, to produce many pieces indistinguishable in style from the music of Ernest Bloch (who composed in Switzerland in a sparse, crisp 20th-century style) when various melodic, harmonic, and rhythmic parameters were twiddled. To me, this is simply inconceivable.

6. Interpolating Between an Arbitrary Pair of Typefaces

The worst is yet to come, however. Presumably Knuth did not wish us to take his rhetorical question in such a limited way as to imply that the numbers $6\frac{1}{7}$ and $\frac{1}{4}$ were important. Pretty obviously, they were just examples of arbitrary parameter settings. Presumably, if Metafont could easily give you a $6\frac{1}{7}$ -point font that is $\frac{1}{4}$ of the way between Baskerville and Helvetica, it could as easily give you an $11\frac{2}{3}$ -point font that is $\frac{5}{17}$ of the way between Baskerville and Helvetica—and so on. And why need it be restricted to Baskerville and Helvetica? Surely those numbers weren't the only "soft" parts of the rhetorical question! Common sense tells us that Helvetica and Baskerville were also merely arbitrary choices of typeface. Thus the hidden implication is that, as easily as one can twiddle a dial to change point size, so one can twiddle another dial (or set of dials) and arrive at any desired typeface, be it Helvetica, Baskerville, or whatever. Knuth might just as easily have put it this way: "The ability to manipulate lots of parameters may be interesting and fun, but does anybody really need an X-point

font that is Y percent of the way between typeface T1 and typeface T2?” For instance, we might have set the four knobs to the following settings:

- X: 36
- Y: 50%
- T1: Magnificat
- T2: Stop

Each of these two typefaces (Figure 7a,b) is ingenious, idiosyncratic, and visually intriguing. I challenge any reader to even imagine a blend halfway between them, let alone draw it! And to emphasize the flexibility implied by the question, how about trying to imagine a typeface that is (say) one third of the way between Cirkulus and Block Up (Figure 7c,d)? Or one that is somewhere between Explosion and Shatter (Figure 7e,f)?

7. “A Posteriori” Knobs and the Frame Problem of AI

Shatter, incidentally, provides an excellent example of the trouble with viewing everything as coming from parameter settings. If you look carefully, you will see that



Figure 7. (a) Magnificat, (b) Stop, (c) Cirkulus, (d) Block Up, (e) Explosion, (f) Shatter, and (g) Helvetica Medium Italic.

Shatter is indeed a “variation on a theme,” the theme being Helvetica Medium Italic (Figure 7g). But does that imply that any meticulous parametrization of Helvetica would automatically yield Shatter as one of its knob-settings? Of course not. That is absurd. No one in their right mind would anticipate such a variation while parametrizing Helvetica, just as no one in their right mind when delivering their Nobel Lecture would say, “Thank you for awarding me my first Nobel Prize.” When someone wins a Nobel Prize, they do not immediately begin counting how many they have won. Of course, if they win two, then a knob will spontaneously appear in most people’s minds, and friends will very likely make jokes about the next few Nobel Prizes. Before the second prize, however, the “just-one” quality would have been an unperceived fact.

This is closely related to a famous problem in cognitive science (the study of formal models of mental processes, especially computer models) called the “frame problem” (Dennett, 1981), which can be epitomized this way: How do I know, when telling you I’ll meet you at seven at the train station, that it makes no sense to tack on the proviso, “as long as no volcano erupts along the way, burying me and my car on the way to the station,” but that it does make reasonable sense to tack on the proviso, “as long as no traffic jam holds me up”? And, of course, there are many intermediate cases between these two. The frame problem is about the question, “What variables (knobs) is it within the bounds of normalcy to perceive?” Clearly, no one can conceivably anticipate all the factors that might somehow be relevant to a given situation; one simply blindly hopes that the species’ evolution and the individual’s life experiences have added up to a suitably rich combination to make for satisfactory behavior most of the time. There are too many contingencies, however, to try to anticipate them all, even given the most powerful computer. One reason for the extreme difficulty in trying to make machines able to learn is that we find it very hard to articulate a set of rules defining when it makes sense and when it makes no sense to perceive a knob.

This brings us back to Shatter, seen as a variation on Helvetica. Obviously, once you’ve seen such a variation, you can add a knob (or a few) to your Metafont “Helvetica machine,” enabling Shatter to come out. (Indeed, you could add similar “Shatterizing” knobs to your “Baskerville machine,” for that matter!) But this would all be *a posteriori*: after the fact. The most telling proof of the artificiality of such a scheme is, of course, that no matter how many variations have been made on (say) Helvetica, people can still come up with many new and unanticipated varieties, such as: Helvetica Rounded, Helvetica Rounded Deco, Helvetican Flair, and so on (Figure 8a,b,c; Graphic Products Corporation, 1981).

Incidentally, it is important that I make it clear that although I find it easier to make my points with somewhat extreme or exotic versions of letters, these points hold just as



Figure 8. Top to bottom: (a) Helvetica Rounded, (b) Helvetica Rounded Deco, and (c) Helvetican Flair (Graphic Products Corporation, 1981).

strongly for more conservative letters. One simply has to look at a finer grain size, and all the same kinds of issues reappear.

No matter how many new knobs—or even new families of knobs—you add to your Helvetica machine, you will have left out some possibilities. People will forever be able to invent novel variations on Helvetica that haven’t been foreseen by a finite parametrization, just as musicians will forever be able to devise novel ways of playing “Begin the Beguine” that the electronic-organ builders haven’t yet built into their elaborate repertoire of canned rhythms, harmonies, and so forth. To be sure, the organ builders can always build in extra possibilities after they have been revealed, but by then a creative musician will have long since moved on to other styles. One can imagine Helvetica modified in many novel ways inspired by various extant typefaces such as Sunrise, Buster, Stack, Double, and so on (Figure 9). I leave it to readers to try to imagine such variants.

8. A Total Unification of All Typefaces?

The worst is still yet to come! Knuth’s throwaway sentence unspokenly implies that we should be able to interpolate any fraction of the way between any two arbitrary typefaces. For this to be possible, any pair of typefaces would have to share the exact same set of knobs (otherwise, how could you set each knob to an intermediate setting?). And since all pairs of typefaces have the same set of knobs, transitivity implies that all



Figure 9. Top to bottom: (a) Sunrise, (b) Buster, (c) Stack, and (d) Double.

typefaces would have to share a single, grand, universal, all-inclusive, ultimate set of knobs. (The argument is parallel to the following one: If, given any two people, they have the same number of legs, then all people have the same number of legs.)

Thus we realize that Knuth’s sentence casually implies the existence of a “universal A-machine”—a single Metafont program with a finite set of parameters, such that any combination of settings of them will yield a valid A, and conversely, such that any valid A will be yielded by some combination of settings of them. Now how can you possibly incorporate all of the previously shown typefaces into one universal schema?

Or look again at the 56 capital A’s of Figure 1. Can you find in them a set of specific, quantifiable features? (For a comparable collection for each letter of the alphabet, see Kuwayama, 1973.) Imagine trying to pinpoint a few dozen discrete features of the Magnificat A (Figure 1, A7) and simultaneously finding their “counterparts” in the Univers A (D3). Suppose you have found enough to characterize both completely. Now remember that every intermediate setting also must yield an A. This means we will have every shade of “cross” between the two typefaces.

This intuitive sense of a “cross” between two typefaces is common and natural, and occurs often to typeface lovers when they encounter an unfamiliar typeface. They may characterize the new face as a cross between two familiar typefaces (“Vivaldi is a cross between Magnificat and Palatino Italic Swash”) or they may see it as an exaggerated rendition of a familiar typeface (“Magnificat is Vivaldi squared”; Figure 10). What truth is there to such a statement? All one can really say is that each Magnificat letter looks “sort of like” its Vivaldi counterpart, only about “twice as fancy” or “twice as curly” or something vague along these lines. But how could a single “curliness” knob account for the mysteriously beautiful meanderings, organic and capricious, in each Magnificat letter?



Figure 10. Top to bottom: (a) Palatino Italic Swash caps, (b) Vivaldi caps, and (c) Magnificat caps.

Can you imagine twisting one knob and watching thin, slithery tentacles begin to grow out of the Palatino Italic **A**, snaking outwards eventually to form the Vivaldi **A**, then continuing to twist and undulate into ever more sinuous forms, yielding the Magnificat **A** in the end? And who says that that is the ultimate destination? If Magnificat is Vivaldi squared, then what is Magnificat squared?

Specialists in computer animation have had to deal with the problem of interpolation of different forms. For example, in a television series about evolution, there was a sequence showing the outline of one animal slowly transforming into another one. But one cannot simply tell the computer, “Interpolate between this shape and that one!” To each point in one there must be explicitly specified a corresponding point in the other. Then one lets the computer draw some intermediate positions on one’s screen, to see if the choice works. A lot of careful “tuning” of the correspondences between figures must be done before the interpolation looks good. There is no recipe that works in general.

9. The Essence of A-Ness Is Not Geometrical

Despite all the difficulties described above, some people, even after scrutinizing the wide diversity of realizations of the abstract **A**-concept, still maintain that they all do share a common geometric quality. They sometimes verbalize it by saying that all **A**’s have “the same shape” or are “produced from one template.” Some mathematicians are inclined to search for a topological or group-theoretical invariant. A typical suggestion might be: “All instances of **A** are open at the bottom and closed at the top.” Well, in Figure 1, sample A8 (Stop) seems to violate both of those criteria. And many others of the sample letters violate at least one of them. In several examples, such concepts as “open” or “closed” or “top” or “bottom” apply only with difficulty. For instance, is G7 (Sinaloa) open at the bottom? Is F6 (Calypso) closed at the top? What about A4 (Astra)?

The problem with the Metafont “knobs” approach to the **A** category is that each knob stands for the presence or absence (or size or angle, etc.) of some specifically geometric feature of a letter: the width of its serifs, the height of its crossbar, the lowest point on its

left arm, the highest point along some extravagant curlicue, the amount of broadening of a pen, the average slope of the ascenders, and so forth and so on. But in many A's such notions are not even applicable. There may be no crossbar, or there may be two or three or more. There may be no curlicue, or there may be a few curlicues.

10. Chauvinism Versus Open-Mindedness: Fixed Questionnaires Versus Fluid Roles

A Metafont joint parametrization of two A's presumes that they share the same features, or what might be called "loci of variability." It is a bold (and, I maintain, absurd) assumption that one could get any A by filling out an eternal and fixed questionnaire: "How wide is its crossbar? What angle do the two arms make with the vertical? How wide are its serifs?" (and so forth). There may be no identifiable part that plays the crossbar role, or the left-arm role; or some role may be split among two or more parts. You can easily find examples of these phenomena among the 56 A's in Figure 1. Some other examples of what I call role splitting, role sharing, role transferral, role redundancy, and role elimination are shown in Figure 11. These terms describe the ways that conceptual

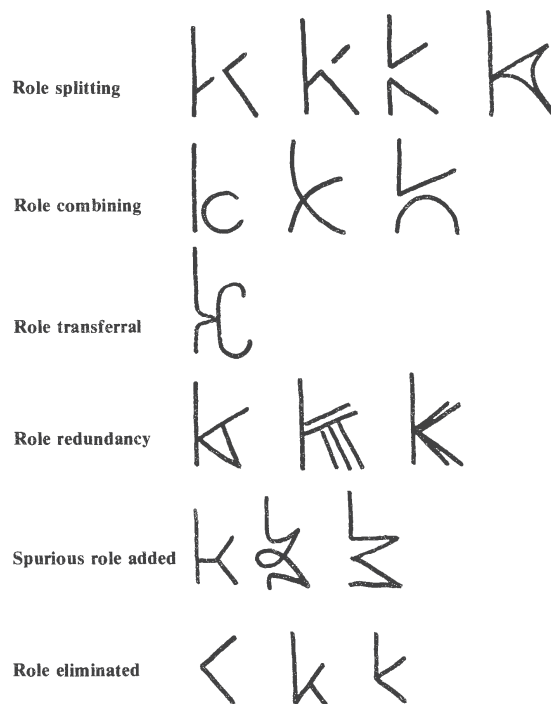


Figure 11. Examples of role splitting, combining, transferral, redundancy, added (spuriously), and eliminated in typographic features.

roles are apportioned among various geometric entities, which are readily recognized by their connectedness and gentle curvatures.

When I was 12, my family was about to leave for Geneva, Switzerland, for a year, so I tried to anticipate what my school would be like. The furthest my imagination could stretch was to envision a school that looked exactly like my one-story Californian stucco junior high school, only with classes in French (twiddling the “language” knob), and with the schoolbus that would pick me up each morning perhaps pink instead of yellow (twiddling the “schoolbus color” knob). I was utterly incapable of anticipating the vast difference that there actually turned out to be between the Geneva school and my California school.

Likewise, there are many “exobiologists” who have tried to anticipate the features of extraterrestrial life, if it is ever detected. Many of them have made assumptions that to others appear strikingly naive. Such assumptions have been dubbed “chauvinisms” by Carl Sagan (1973). There is, for instance, “liquid chauvinism,” which refers to the phase of the medium in which the chemistry of life is presumed to take place. There is “temperature chauvinism,” which assumes that life is restricted to a temperature range not too different from that here on the planet Earth. In fact, there is planetary chauvinism—the idea that all life must exist on the surface of a planet orbiting a certain type of star. There is carbon chauvinism, assuming that carbon must form the keystone of the chemistry of any sort of life. There is speed chauvinism, assuming that there is only one “reasonable” rate for life to proceed at. And so it goes.

If a Londoner arrived in New York, we might find it naive (or perhaps pathetic) if he or she asked: “Where is your Big Ben? Where are your Houses of Parliament? Where does your Queen live? When is your teatime?” The idea that the biggest city in the land need not be the capital, need not have a famous bell tower in it, and so on, seems totally obvious after the fact, but to the naive tourist it can come as a surprise.

The point here is that when it comes to fluid semantic categories such as **A**, it is equally naive to presume that it makes sense to refer to “the crossbar” or “the top” or to any constant feature. It is quite like expecting to find “the same spot” in any two pieces of music by the same composer. The problem, I have found, is that most people continue to insist that any two instances of **A** have “the same shape,” even when confronted with such pictures as Figure 1.

The analogy between Britain and the United States is a useful one to continue for a moment. The role that London plays in England is certainly multifaceted, but two of its main roles are “chief commercial city” and “capital.” These two roles are played by different cities in the U.S. On the other hand, the role that the American President plays in the U.S. is split into pieces in Britain, part being carried by the Queen (or King), and

part by the Prime Minister. Then there is a subsidiary role played by the President's wife—the “First Lady.” Her counterpart in Britain is also split, and moreover, these days “wife” has to be replaced by “husband,” whether one is thinking that the “President of England” is the Queen or the Prime Minister. (See Hofstadter, 1981, for an extended discussion of such analogy problems and their relation to machine intelligence.)

To think one can anticipate the complete structure of one country or language purely on the basis of being intimately familiar with another one is presumptuous and, in the end, preposterous. Even if you have seen dozens, you have not exhausted the potential richness and novelty in such domains. In fact, the more instances you have seen, the more circumspect you are about making unwarranted presumptions about unseen instances, although certainly your ability to anticipate the unanticipated (or unanticipable) improves! The same holds for instances of any letter of the alphabet or other semantic category.

11. The A Spirit

Clearly there is much more going on in typefaces than meets the eye—literally. The shape of a letterform is a surface manifestation of deep mental abstractions. It is determined by conceptual considerations and balances that no finite set of merely geometric knobs could capture. Underneath or behind each instance of **A** there lurks a concept, a Platonic entity, a spirit. This Platonic entity is not an elegant shape such as the Univers **A**, not a template with a finite number of knobs, not a topological or group-theoretical invariant in some mathematical heaven, but a mental abstraction—a different sort of beast. Each instance of the **A** spirit reveals something new about the spirit without ever exhausting it. The mathematization of such a spirit would be a machine with a specific set of knobs on it, defining all its “loci of variability” for once and for all. I have tried to show that to expect this is simply not reasonable. In fact, I made the following claim, above: “No matter how many new knobs—or even new families of knobs—you add to your ... machine, you will have left out some possibilities. People will forever be able to invent novel variations ... that haven't been foreseen by a finite parametrization...”

Of what, then, is such an abstract “spirit” composed? Or is it simply a mystically elusive, noncapturable essence that defies the computational—indeed, the scientific—approach totally? Not at all, in my opinion. I simply think that a key idea is missing in what I have described so far. And what is this key idea? I shall first describe the key misconception. It is to try to capture the essence of each separate concept in a separate “knobbed machine”—that is, to isolate the various Platonic spirits. The key insight is that those spirits overlap and mingle in a subtle way.

12. Happy Roles, Unhappy Roles, and Quirk-Notes

The way I see it, the Platonic essence lurking behind any concrete letterform is composed of conceptual “roles” rather than geometric parts. (A related though not identical notion called “functional attributes” was discussed by Barry Blesser and co-workers nearly ten years ago in *Visible Language*; Blesser et al., 1973.) A role, in my sense of the term, does not have a fixed set of parameters defining the extent of its variability, but it has instead a set of tests or criteria to be applied to candidates that might be instances of it. For a candidate to be accepted as an instance of the role, not all the tests have to be passed; not all the criteria have to be present. Instead, the candidate receives a score computed from the tests and criteria, and there is a threshold point above which the role is “happy,” and below which it is “unhappy.” Then below that, there is a cut-off point below which the role is totally dissatisfied, and rejects the candidate outright.

An example of such a role is that of “crossbar.” Note that I am not saying “crossbar in capital A,” but merely “crossbar.” Roles are modular: they jump across letter boundaries. The same role can exist in many different letters. This is, of course, reminiscent of the fact that in Metafont a serif (or generally, any geometric feature shared by several letters) can be covered by a single set of parameters for all letters, so that all the letters of the typeface will alter consistently as a single knob is turned. The difference is that my notion of “role” doesn’t have the generative power that a set of specific knobs does. From the fact that a given role is “happy” with a specific geometric filler, one cannot deduce exactly how that filler looks. There is, of course, more to a role’s “feelings” about its filler than simply happiness or unhappiness; there are a number of expectations about how the role should be filled, and the fulfillment (or lack thereof) can be described in “quirk-notes.” Thus, quirk-notes can describe the unusual slant of a crossbar (Figure 1, E1, Arnold Bocklin), the fact that it is filled by two strokes rather than one (E3, Airkraft), or the fact that it fails to meet (or has an unusual way of meeting) its vertical mate (A2, Eckmann Schrift; F5, Le Golf; and many others).

These quirk-notes are characterizations of stylistic traits of a perceived letterform. They do not contain enough information, however, to allow a full reconstruction of that letterform, whereas a Metafont program does contain enough information for that. However, they do contain enough information to guide the creation of many specific letterforms that have the given stylistic traits. All of them would be, in some sense, “in the same style.”

13. Modularity of Roles

The important thing is that this modularity of roles allows them to be exported to other letters, so that a quirk-note attached to a particular role in **A** could have relevance to

E, L, or T. Thus stylistic consistency among different letters is a by-product of the modularity of roles, just as the notion of letter-spanning parameters in Metafont gives rise to internal consistency of any typeface it might generate.

Furthermore, there are connections among roles so that, for instance, the way in which the “crossbar” role is filled in one letter could influence the way that the “post” or “bowl” or “tail” role is filled in other letters. This is to avoid the problem of overly simplistic mappings of one letter onto another, analogous to the Londoner asking an American where the American Houses of Parliament are. Just as one must interpret “Houses of Parliament” liberally rather than literally when “translating” from England to the U.S., so one may have to convert “crossbar” into some other role when looking for something analogous in the structure of another letter than **A**, such as **N**. In certain typefaces the diagonal stroke in **N** could well be the counterpart of the crossbar in **A**. But it is important to emphasize that no fixed (i.e., typeface-independent) mapping of roles in **A** onto roles in **N** will work; only the specific letterforms themselves (via their quirk-notes) can determine what roles (if any) should be mapped onto each other. Such cross-letter mappings must be mediated by a considerable degree of understanding of what functions are fulfilled by all the roles in the two particular letters concerned. (This fluid mapping of roles is discussed in more detail in Hofstadter, 1982b.)

14. Typographical Niches and Rival Categories

So far I have sketched very quickly a theory of “Platonic essences” or “letter spirits” involving modular roles—roles shared among several letters. This sharing of roles is one aspect of the overlapping and mingling that I spoke of above. There is a second aspect, which is suggested by the phrase “typographical niche.” The notion is analogous to that of “ecological niche.” When, in the course of perception of a letterform, a group of roles have been activated and have decided that they are present (whether happily or unhappily), their joint presence constitutes evidence that one of a set of possible letters is present. (Remember that since a role is not the property of any specific letter, its presence does not signal that any specific letter is in view.)

For instance, the presence of a “post” role and a “bowl” role in certain relative positions would suggest very strongly that there is a **b** present. Sometimes there may be evidence for more than one letter. The eye-mind combination is not happy with any such unstable state for long, and strains to make a decision. It is as if there is a very steep and slippery ridge between valleys, and a ball dropped from above is very unlikely to come to settle on top of the ridge. It will tumble to one side or the other. The valleys are the typographical niches.



Figure 12. Versions of **h** and **k** as rivals for the same typographical niche.

Now the overlapping of letters comes about because each letter is aware of its typographical rivals, its next-door neighbors, just over the various ridges that surround its space. The letter **h**, for instance, is acutely sensitive to the fact that it has a close rival in **k**, and vice versa (Figure 12). The letter **T** is very touchy about having its crossbar penetrated by the post below, since even the slightest penetration is enough to destroy its **T**-ness and to slip it over into **T**'s arch-rival niche, **t**. It's a low ridge, and for that reason, **T** guards it extra-carefully.

15. The Intermingling of Platonic Essences

This image is, I hope, sufficiently strong to convey the second sense of overlapping and intermingling of Platonic essences. "No letter is an island," one might say. There has to be much mutual knowledge spread about among all the letters. Letters mutually define each others' essences, and this is why an isolated structure supposedly representing a single letter in all its glory is doomed to failure.

A letterform-designing computer program based on the above-sketched notions of typographical roles and niches would look very different from one that tried to be a full "mathematization of categories." It would involve an integration of perception with generation, and moreover an ability to generalize from a few letterforms (possibly as few as one) to an entire typeface in the style of the first few. It would not do so infallibly; but of course it is not reasonable to expect "infallible" performance, since stylistic consistency is not an objectively specifiable quality.

In other words, a computer program to design typefaces (or anything else with an esthetic or subjective dimension) is not an impossibility; but one should realize that, no less than a human, any such program will necessarily have a “personal” taste—and it will almost certainly not be the same as its designers’ taste. In fact, to the contrary, the program’s taste will quite likely be full of unanticipated surprises to its programmers (as well as to everyone else), since that taste will emerge as an implicit and remote consequence of the interaction of a myriad features and factors in the architecture of the program. Taste itself is not directly programmable. Thus, although any esthetically programmed computer will be “merely doing what it was programmed to do,” its behavior will nonetheless often appear idiosyncratic and even inscrutable to its programmers, reflecting the fact—well known to programmers—that often one has no clear idea (and sometimes no idea at all) just what it is that one has programmed the machine to do!

16. The “Vertical” and “Horizontal” Problems: Two Equally Important Facets of One Problem

I have made a broad kind of claim: that true understanding of letterforms depends on more than understanding something about each Platonic letter in isolation; it depends equally much on taking into account the ways that letters and their pieces are interrelated, on the ways that letters depend on each other to define a total style. In other words, any approach to the impossible dream of the “secret recipe” for “A-ness” requires a simultaneous solution to two problems, which I call the “vertical” and the “horizontal” problems (Figure 13). The former is the question, “What do all the items in any column have in common?” The latter is the question, “What do all the items in any row have in common?”




Figure 13. The vertical and horizontal problems.




Figure 14. A few faces designed by Hermann Zapf. Top to bottom: (a) Optima, (b) Palatino, (c) Melior, (d) Zapf Book, (e) Zapf International, and (f) Zapf Chancery.

Actually, there is no reason to stop with two dimensions; the problem seems to exist at higher degrees of abstraction. We could lay out our table of comparative typefaces more carefully; in particular, we could make it consist of many layers stacked on top of each other, as in a cake. On each layer would be aligned many typefaces made by a single designer. This idea is illustrated in Figure 14, showing a few faces designed by Hermann Zapf (Optima, Palatino, Melior, Zapf Book, Zapf International, Zapf Chancery; see Zapf, 1960). Along with the Zapf layer, one can imagine a Frutiger layer, a Lubalin layer, a Goudy layer, and so on. One could try to arrange the typefaces in such a way that “corresponding” typefaces by various designers are aligned.

Now in this three-dimensional cake, the two earlier one-dimensional questions still apply, but there is also a new two-dimensional question: “What do all the items in a given layer have in common?” The third dimension can be explored as one moves from one layer to another, asking what all the typefaces in a given “shaft” have in common.

Moreover, a fourth dimension can be added if you imagine many such “layercakes,” one for each distinguishable period of typographical design. Thus our fourth dimension, like Einstein’s, corresponds to time. Now one can ask about each layercake, “What do all the items herein have in common?” This is a three-dimensional question. Presumably, one could carry this exercise even further.

If we go back to the “simplest” of these questions, the original “vertical” question of Figure 13, a naive answer to it could be stated in one word: “letter.” And likewise, a naive answer to the “horizontal” question of Figure 13 is also storable in one word:

“spirit.” In fact, the word “spirit” is applicable, in various senses of the term, to all the higher-dimensional questions, such as “What do all the typefaces produced in the Art Deco era have in common?” There is such a thing, ephemeral though it may be, as “Art Deco spirit,” just as there is undeniably such a thing as “French spirit” in music or “impressionistic spirit” in art (see Loeb, 1975).

Stylistic moods permeate whole periods and cultures and indirectly determine the kinds of creations that people in them come up with. They exert gentle but definite “downward” pressures. As a consequence, not only are the alphabets of a given period and area distinctive, but one can even recognize “the same spirit” in such things as teapots, coffee cups, furniture, automobiles, architecture, and so on (Bush, 1975). And one can also be inspired by a given typeface to carry its ephemeral spirit over into another alphabet, such as Greek, Hebrew, Cyrillic, or Japanese. In fact, this has been done in many instances (Figure 15). The problem I am most concerned with in my research is whether (or rather, how) susceptibility to such a “spirit” can be implanted in a computer program.

17. Letter and Spirit

These words “letter” and “spirit,” of course, recall the contrast between the “letter of the law” and the “spirit of the law,” and the way in which our legal system is constructed so that judges and juries will base their decisions on precedents. This means that any case must be “mapped,” in a remarkably fluid way, by members of a jury, onto previous cases. It is up to the opposing lawyers, then, to be advocates of particular mappings; to try to channel the jury members’ perceptions so that one mapping dominates over another. It is quite interesting that jury decisions are supposed to be unanimous, so that in a metaphorical sense, a “phase transition” or “crystallization” of opinion must take place. The decision must be solidly locked in, so that it reflects not simply a majority or even a consensus, but a totality, a unanimity (which, etymologically, means “one-souledness”). (For discussions of such “phase transitions,” see Hofstadter, 1982c and 1983, and for descriptions of computer models of perception in which a form of collective decision making is carried out, see Reddy, 1976, and Winston, 1975.)

In law, extant rules, statutes, and so on are never enough to cover all possible cases (reminding us once again of the fact that no fixed and rigid set of A-defining rules can anticipate all A’s). The legal system depends on the notion that people, whose experience covers much more than the specific case and rules at hand, will bring to bear their full range of experience not only with many categories but also with the whole process of categorization and mapping. This allows them to transcend the specific, rigid, limited rules, and to operate according to more fluid, imprecise, yet more powerful principles. Or, to revert to the other vocabulary, this ability is what allows people to transcend the



Figure 15. The “spirit” of some Roman typefaces carried over into Cyrillic and Greek typefaces (Compugraphic Corporation, 1982) and into Hebrew and Japanese typefaces (Biggs, 1977). The related Kana and Latin letters were designed by Yasaburo Kuwayama for the Nissan Company.

letter of the law and to apply its spirit. It is this tension between rules and principles, between the letter and the spirit, that is so admirably epitomized for us by the work of Knuth and others exploring the relationship between artistic design and mechanizability. We are entering a very exciting and important phase of our attempts to realize the full potential of computers, and Knuth’s article points to many of the significant issues that must be thought through very carefully.

To conclude, then, I wish to state that the mathematization of categories is an elegant goal, a wonderful beckoning mirage before us, and the computer is the obvious medium to exploit to try to realize this goal. Donald Knuth, whether he has been pulled by a

distant mirage or by an attainable middle-range goal, has contributed immensely, in his work on Metafont, to our ability to deal with letterforms flexibly, and has cast the whole problem of letters and fonts in a much clearer perspective than ever before. Readers, however, should not pull a false message out of his article: they should not confuse the chimera of the mathematization of categories with the quest after a more modest but still fascinating goal. In my opinion, one of the best things Metafont could do is to inspire readers to chase after what Knuth has rightly termed “the intelligence” of a letter, making use of the explicit medium of the computer to yield new insights into the elusive “spirits” that flit about so tantalizingly, hidden just behind those lovely shapes we call “letters.”

18. Acknowledgments

Editor’s note. Special thanks to Douglas Hofstadter, who encouraged this reprint upon request in communication with Mike Zender. Matthew Peterson handled editorial conversion of the original publication and wrote the prefatory remarks, and Muhammad Rahman captured the figures from a printed copy of issue 16.4. Mentions of letters such as **A** are in monospaced semibold in this reprint, though quotes alone were utilized in the original.

19. References Exclusive to the Reprint

- Baudin, F., Bigelow, C., Bronsard, H.-P., Fisher, E., Jr., Ford, D., Gore, G., Jaspert, W. P., Kapr, A., Karow, P., Nesbitt, A., Rondthaler, E., Schappler, J., Tracy, W., Ungar, G., Zapf, H., & Knuth, D. E. (1982). Other replies to Donald E. Knuth’s article, “The Concept of a Meta-Font” [V. Mollar & K. Schmidt, Trans.]. *Visible Language*, 16(4), 339–359. <https://journals.uc.edu/index.php/vl/article/view/5350>
- Hofstadter, D. R. (1982). Metafont, metamathematics, and metaphysics: Comments on Donald Knuth’s article “The Concept of a Meta-Font.” *Visible Language*, 16(4), 309–338. <https://journals.uc.edu/index.php/vl/article/view/5350/4214>
- Hofstadter, D. R. (2007). *I am a strange loop*. Basic Books.
- Hofstadter, D. R. (2025). *Ambigrammia: Between creation and discovery*. Yale University Press.
- Littlejohn, D. (2026). A profession provoked: How Meta-Font struck a nerve. *Visible Language*, 60(1), 120–123.
- Zender, M. (2026). Why Meta-Font struck a nerve. *Visible Language*, 60(1), 84–89.

20. References

- Biggs, J. R. (1977). *Letterforms and lettering*. Blandford Press.
- Blessner, B., Shillman, R., Cox, C., Kuklinski, T., Ventura, J., & Eden, M. (1973). Character recognition based on phenomenological attributes. *Visible Language*, 7(3), 209–223. <https://journals.uc.edu/index.php/vl/article/view/5146>

- Bush, D. J. (1975). *The streamlined decade*. George Braziller.
- Chaitin, G. (1975). Randomness and mathematical proof. *Scientific American*, 232(5), 47–52. <https://www.jstor.org/stable/24949798>
- Compugraphic Corporation (1982). *A portfolio of text and display type*. Compugraphic Corporation.
- DeLong, H. (1970). *A profile of mathematical logic*. Addison-Wesley.
- Dennett, D. C. (1964). Cognitive wheels: The frame problem of AI. In C. Hookway (Ed.), *Minds, machines and evolution* (pp. 129–152). Cambridge University Press.
- Gödel, K. (1931/1962). *On formally undecidable propositions of Principia Mathematica and related systems* (B. Meltzer, Trans.). Basic Books.
- Graphic Products Corporation (1981). *Format catalogue no. 7*. Graphic Products Corporation.
- Hofstadter, D. R. (1979). *Gödel, Escher, Bach: An eternal golden braid*. Basic Books.
- Hofstadter, D. R. (1981). Metamagical themas: How might analogy, the core of human thinking, be understood by computers? *Scientific American*, 245(3), 18–30. <https://www.jstor.org/stable/24964552>
- Hofstadter, D. R. (1982a). Metamagical themas: Can inspiration be mechanized? *Scientific American*, 247(3), 18–34. <https://www.jstor.org/stable/24966674>
- Hofstadter, D. R. (1982b). Metamagical themas: Variations on a theme as the essence of imagination. *Scientific American*, 247(4), 20–29. <https://www.jstor.org/stable/24966697>
- Hofstadter, D. R. (1982c). Who shoves whom around inside the Careenium? or what is the meaning of the word “I”? *Synthese*, 53, 189–218. <https://link.springer.com/article/10.1007/BF00484897>
- Hofstadter, D. R. (1983). Artificial intelligence: Subcognition as computation. In F. Machlup & U. Mansfield (Eds.), *The study of information: Interdisciplinary messages* (pp. 263–287). John Wiley and Sons.
- Kim, S. E. (1981). *Inversions: A catalogue of calligraphic cartwheels*. Byte Books.
- Knuth, D. E. (1982). The concept of a meta-font. *Visible Language*, 16(1), 3–27. <https://journals.uc.edu/index.php/vl/article/view/5329/4193>
- Kuwayama, Y. (1973). *Trademarks and symbols, volume 1: Alphabetical designs*. Van Nostrand Reinhold.
- Larcher, J. (1976). *Fantastic alphabets*. Dover Press.
- Letraset (1981). *Graphic art materials reference manual*. Letraset.
- Loeb, M. (1975). *New art deco alphabets*. Dover Press.
- Myhill, J. (1952). Some philosophical implications of mathematical logic: I. Three classes of ideas. *Review of Metaphysics*, 6(2), 165–198. <http://www.jstor.com/stable/20123317>
- Nagel, E., & Newman, J. R. (1958). *Gödel's proof*. New York University Press.
- Post, E. (1965). Recursively enumerable sets of positive integers and their decision problems. In M. Davis (Ed.), *The undecidable: Basic papers on undecidable propositions, unsolvable problems and computational functions* (pp. 304–337). Raven Press.
- Reddy, R. (1976). *Working papers in speech recognition, IV: The Hearsay II system* [Technical report]. Carnegie-Mellon University Computer Science Department. <https://apps.dtic.mil/sti/tr/pdf/ADA027844.pdf>
- Rogers, H. (1967). *Theory of recursive functions and effective computability*. McGraw-Hill.
- Rucker, R. (1982). *Infinity and the mind*. Birkhauser.
- Sagan, C. (1973). *Communication with extraterrestrial intelligence*. MIT Press.
- Smullyan, R. (1961). *Theory of formal systems*. Princeton University Press.

Smullyan, R. (1978). *What is the name of this book?* Prentice-Hall.
 Strich, C. (1981). *Fellini's faces: 418 photographs from the archives of Federico Fellini*. Holt, Rinehart, and Winston.
 Winston, P. H. (1975). *The psychology of computer vision*. McGraw-Hill.
 Zapf, H. (1960). *About alphabets*. MIT Press.

Author

[Original bio from 1982.] **Douglas R. Hofstadter** is an associate professor in the Computer Science Department at Indiana University (Bloomington, IN, USA). Immediately after receiving his Ph.D. in physics, he wrote *Gödel, Escher, Bach: an Eternal Golden Braid* (Basic Books, 1979), a book about mentality and consciousness, and their relation to abstract structures. He conducts research in artificial intelligence, concentrating on the higher-level aspects of perception, and how they are related to generalization, abstraction, and induction. His project Letter Spirit in particular is concerned with the notion of style in alphabets; his goal is to have it be able to extrapolate an entire typeface from a few sample letters given to it. He has co-edited, with philosopher Daniel Dennett, an anthology of fantasies and reflections on self and soul: *The Mind's I* (Basic Books, 1981). He also writes a column entitled “Metamagical Themas” each month in *Scientific American*.

[Bio for 2026.] **Douglas R. Hofstadter** is a distinguished professor of cognitive science at Indiana University (Bloomington, IN, USA). He recently added *Ambigrammia: Between Creation and Discovery (ABCD)* (Yale University Press) to his list of authored books. *ABCD* contains hundreds of ambigrams, defined as “a piece of writing expressly designed to squeeze in more than one reading” (Hofstadter, 2025, p. 1). In the example below, each color name can be read forward and also in a mirror. The final instance actually has *four* readings, with DOUG (the ambigram’s creator) and 2006 (its year of creation) simultaneously readable both forward and backward. In his *I Am a Strange Loop* (Basic Books), published earlier in 2007, Hofstadter offers one of his “firmest conclusions”: “we always think by seeking and drawing parallels to things we know from our past, and... we therefore communicate best when we exploit examples, analogies, and metaphors galore, when we avoid abstract generalities, when we use very down-to-earth, concrete, and simple language, and when we talk about our own experiences” (Hofstadter, 2007, p. xv). Among other down-to-earth intellectual contributions in simple language, he is known for Hofstadter’s Law (p. xv): “It always takes longer than you think it will take, even when you take into account Hofstadter’s Law.”





A Profession Provoked: How Meta-Font Struck a Nerve

Deborah Littlejohn 

College of Design, North Carolina State University, Raleigh, NC, USA (dklittle[at]ncsu.edu)

Knuth's Meta-Font appeared at a pivotal moment in late twentieth century design discourse, at a time when the field was beginning to grapple with computation in the context of wider (and longer, more entrenched) debates around intuition and craft versus system and rule. Upon publishing Knuth's (1982a) paper in issue 16.1, *Visible Language* solicited responses from notable design luminaries (extending to Knuth himself, a computer scientist at Stanford University), that would appear in issue 16.4 (Baudin et al., 1982). Their responses—16 in total—reveal a rich snapshot of a field deeply concerned with maintaining its aesthetic authority under conditions of technological inevitability, even as its practices were increasingly becoming mediated by code. Together, these letters depict a profession that was simultaneously fascinated, alarmed, and introspective.

The Meta-Font debate comes across as strikingly familiar today, given the myriad discussions around artificial intelligence in design. The same anxieties persist about technological advances displacing the designer or replacing design itself. On the flip side, we discern a familiar optimism—AI, like Meta-Font, is less a surrogate designer than an amplifier of the designer's intent; it can accelerate iteration and support new streams of design research. After forty years of technological change, design continues to confront the same question of how to ensure computational systems serve design judgment rather than overwrite it. In this sense, arguments about AI in design echo the Meta-Font debate by reaffirming that each technological advance prompts new articulations of the designer's role as the final arbiter of meaning, legibility, and professional responsibility.

@: [ISSUE](#) > [ARTICLE](#) >

Cite this article:

Littlejohn, D. (2026). A profession provoked: How Meta-Font struck a nerve. *Visible Language*, 60(1), 120–123.

First published online April 26, 2026.

© 2026 Visible Language — this article is **open access**, published under the CC BY-NC-ND 4.0 license.

<https://visible-language.org/journal/>

Visible Language Consortium:

University of Leeds (UK)

University of Cincinnati (USA)

North Carolina State University (USA)

The Case for Meta-Font: Analysis & New Tools

Respondents supportive of Meta-Font viewed it as a continuation of typographic tradition that promised unprecedented analytical power and research potential. Positive comments noted that Meta-Font should be interpreted as a conceptual proposal, rather than an aesthetic achievement. Hermann Zapf (1982) was representative of this view, stating that the system “shows the endless possibilities of this computer aided approach to type design and should not be examined or analyzed merely for its aesthetic values” (p. 356). Several contributors insisted Meta-Font was no radical departure from design practice at all, arguing instead that it was a continuation of typographic history. As Charles Bigelow (1982) explained, “the fundamental idea of a meta-font has been a common theme in the history of typography” (p. 339). He went on to position Knuth’s parametrization approach alongside centuries of expanding type families and systematic variation, from Garamond and Granjon to Frutiger’s Univers. The only thing original in Knuth, said Bigelow, is “the explicit implementation of the design ideas in a computer system” (p. 342). Gerard Unger (1982) also noted that “the urge to parameterize is, like Diderot’s and d’Alembert’s wish to describe and catalogue, a rational aim. And it is no coincidence that in the age of Rationalism the first Meta-Font—or rather, type family—was created by Fournier” (p. 356).

Supportive writers also argued that Meta-Font would not replace designers. As “a very human concept,” Peter Karow (1982) described it as a means by which designers could analyze ideas, test variants rapidly, and refine legibility with unprecedented speed (p. 347). Hermann Zapf (1982) similarly framed Meta-Font as “an ingenious computer based tool worked out for those individuals with less manual design experience,” not a mechanism for automating creativity away (p. 358). Other letters noted Meta-Font’s greatest value was in research, not style. Albert Kapr (1982) praised Knuth for lifting typographic questions “out of the area of graphic feeling into the limelight of scientific knowledge” and even proposed seven concrete optical questions—about stroke width, counters, and centers—that computers might help answer (pp. 348–349). Likewise, Gerard Unger (1982) hoped parametrization could refine legibility studies and lead to more precise design briefs (p. 356).

The Case Against Meta-Font: Aesthetics & the Sovereign Eye

Critical letters focused on aesthetics, philosophy, and the erosion of professional authority. A common criticism was blunt: results were ugly, “engineery,” and distressing. In addition to calling them “ugly,” David Ford (1982) found Knuth’s typefaces “amateurishly rendered” and declared that without convincing forms, Meta-Font remained only “potential abilities—as opposed to reality” (p. 344). Edward Rondthaler (1982) echoed this sentiment, noting that none of the Meta-Font variations shown were typographically better than the originals (p. 351). Skeptics strongly contended that typography

cannot be reduced to metrics without losing its soul. Gerard Unger (1982) insisted, “the gist of a type design cannot be found in its parameters” (p. 354–355), pointing out that the heart of a design often lies in its verbal, cultural intention rather than its measurable traits. Alexander Nesbitt (1982) delivered the harshest assessment: Meta-Font would “churn out an infinite number of ‘designs’ but nothing beautiful” (p. 351) and warned that technology leads designers away from “the eye being ‘the sovereign ruler of taste’” (p. 350).

Several respondents were anxious that Meta-Font would empower inexperienced or poorly trained designers. In Gary Gore’s (1982) most negative leaning letter (and he wrote three; one “angrily,” one “progressively,” and one “cordially,” but all “from the heart”), he feared Knuth’s tool would be “available to amateurs” and set typography “back to a new dark age” (p. 345). Henri-Paul Bronsard (1982) was primarily concerned with how Meta-Font would affect teaching. He recounted how his students asked why they should learn calligraphy at all if a program can design letters, raising for him unresolved questions about education and visual sensitivity (p. 342).

A Case for Caution without Rejection: Consensus & Shared Middle Ground

The letters submitted in response to Knuth (1982a) cannot be neatly categorized into “pro” versus “con” piles. For the most part, respondents conditionally accepted Knuth’s idea for automating typography, if it remained subordinate to human judgment and tradition. Even the stronger critics rarely called for a complete abandonment of the Meta-Font idea. Instead, they urged restraint, hierarchy, and design leadership. Hermann Zapf (1982) called for “brakes” to prevent indiscriminate deformations (p. 357), while Peter Karow (1982) cautioned that variation must be an option, not a mandate (p. 348). Walter Tracy (1982) presented his argument with a touch of humor when he stressed that Meta-Font should serve a designer’s prior intentions, not replace them with “someone fiddling with the parameters until...Eureka! the immaculate conception of a new typeface!” (p. 354).

This “middle ground” position suggests a common thread among all these replies: Meta-Font could be valuable as long as it remains subordinate to human judgment. The tension is between the new technology and how knowledge, skill, taste, and tradition survive technological change, a point that Bigelow (1982) stressed: “The computer requires rational, logical, and algorithmic descriptions, whereas the history of typeface evolution has been replete with accident, idiosyncrasy, serendipity, virtuosity, fortuity, and all of the other irrational, illogical, and intuitive forces to which art is subject” (pp. 342–343). As Knuth himself concluded, the hope is that forcing such principles into explicit form may reveal “how little we really know about letterforms”—and thereby help raise the art to an even higher level” (Knuth, 1982, p. 359).

References

- Bigelow, C. (1982). To the editor. *Visible Language*, 16(4), 339–343.
- Bronsard, H.-P. (1982). To the editor. *Visible Language*, 16(4), 341–344.
- Ford, D. (1982). To the editor. *Visible Language*, 16(4), 344–345.
- Gore, G. (1982). Here are three letters on the Meta-Font, all from the heart. *Visible Language*, 16(4), 345–346.
- Baudin, F., Bigelow, C., Bronsard, H.-P., Fisher, E., Jr., Ford, D., Gore, G., Jaspert, W. P., Kapr, A., Karow, P., Nesbitt, A., Rondthaler, E., Schappler, J., Tracy, W., Ungar, G., Zapf, H., & Knuth, D. E. (1982). Other replies to Donald E. Knuth's article, "The Concept of a Meta-Font" [V. Mollar & K. Schmidt, Trans.]. *Visible Language*, 16(4), 339–359. <https://journals.uc.edu/index.php/vl/article/view/5350>
- Kapr, A. (1982). To the editor. *Visible Language*, 16(4), 346–348.
- Karow, P. (1982). To the editor. *Visible Language*, 16(4), 346–350.
- Knuth, D. E. (1982a). The concept of a meta-font. *Visible Language*, 16(1), 3–27. <https://journals.uc.edu/index.php/vl/article/view/5329/4193>
- Knuth, D. E. (1982b). A reply from the author. *Visible Language*, 16(4), 358–359.
- Nesbitt, A. (1982). To the editor. *Visible Language*, 16(4), 349–352.
- Rondthaler, E. (1982). To the editor. *Visible Language*, 16(4), 350–353.
- Tracy, W. (1982). To the editor. *Visible Language*, 16(4), 353–355.
- Unger, G. (1982). To the editor. *Visible Language*, 16(4), 353–356.
- Zapf, H. (1982). To the editor. *Visible Language*, 16(4), 356–358.

Call for Papers: Special Issue, Early 2027

Typographic Landscapes: Migrating Types—Typographic Meaning-Making Across Boundaries

Visible Language invites submissions for a special issue on the emerging field of typographic landscape research, which connects scholarship in typography, graphic communication, and sociolinguistics to investigate typographic activities as social practice in public spaces.

Typographic landscape research studies all forms of sign-making involving texts in urban, rural, and virtual contexts, ranging from commercial shop signs to self-authorized stickers, artistic murals, and posters, from commemorative placards to regulatory and infrastructural signage, created by professionals and laypeople alike.

The field recognizes that the form of written language, the materiality of letters and signs, and their placement in architectural and other settings are part of the constituted meaning of the messages on display. Typography is not merely a vehicle for linguistic content but communicates meaning through graphic form and materiality itself.

Typographic landscape research considers typographic work as social and communicative practice that is layered and situated across time and space, investigating the use of typographic resources as a means for social actors to perform identities, debate power relationships, negotiate spaces of inclusion or exclusion, signify belonging, challenge or manifest cultural hegemony, preserve local heritage, reference distant geographies, and transform shared places.

This issue will be of interest to anyone designing with type, text and language, particularly those interested in the social dimensions of typography and graphic design, as well as those working with graphic communication in built environments, architecture, urban, and rural spaces.

Submission deadline and contacts

The deadline for submissions is Wednesday, September 2, 2026. Earlier submissions are welcome. When creating your submission, please select the option “special issue article” from the options provided in the “section” field.

All submissions will undergo an initial desk review and those that are to be considered for publication will be subject to double-blind peer review. Final acceptance will require approval from both the guest editors and editor-in-chief.

The special issue is guest-edited by Irmi Wachendorff, University of Reading, UK, and Yu Li, Loyola Marymount University, USA. Inquiries may be sent to the guest editors or to Assistant Editor Matthew Baxter.

- ▶ Irmi Wachendorff (irmi.wachendorff[at]reading.ac.uk)
- ▶ Yu Li (yu.li[at]lmu.edu)
- ▶ Matthew Baxter (M.G.Baxter[at]leeds.ac.uk)

More details are provided on the following page.

Scope and contributions sought

The title “Migrating Types—Typographic Meaning-Making across Boundaries” reflects the focus on typographic meaning-making between *here* and *there*, *now* and *then*, the *self* and the *other*, where such processes are especially active, productive, and consequential.

For this special issue, we are looking for submissions that focus specifically on typographic meaning-making in public (semiotic) spaces referencing across cultural, geographical, territorial, and temporal boundaries. Contributions might engage with (but are not limited to) the role of typography in:

- ▶ Spaces of exclusion, inclusion, migration, and cohabitation
- ▶ The depiction of social hierarchies, power structures, and social transformation
- ▶ The creation of belonging, performance of identity, and negotiation of ideologies

- ▶ Visual and cultural stereotypes, representation, and hegemony
- ▶ Place-making, preservation of local heritage and identities, commodification, and gentrification

We welcome submissions that:

- ▶ Draw on interdisciplinary scholarly foundations combining theoretical perspectives from graphic communication, typography, and sociolinguistics, as well as potentially social semiotics, anthropology, cultural, visual, and communication studies.
- ▶ Employ interdisciplinary methodological approaches (such as quantitative and qualitative empirical case studies, comparative cross-regional or historical analyses) and offer methodological innovations in data collection, mapping, coding, and frameworks for visual and multimodal analysis.

Criteria for inclusion

We expect robust, scholarly, analytical, and critical research. All submissions must meet the following criteria:

- ▶ Clearly stated research question
- ▶ Strong scholarly foundations and engagement with relevant theoretical frameworks
- ▶ Appropriate research design and transparent documentation of systematic data collection processes and parameters (such as sampling strategies, sample size, geographic and temporal scope)
- ▶ Use of visual data as essential evidence to construct and support the analytical argument
- ▶ Critical analysis that moves beyond describing typographic phenomena to analyzing their social, cultural, and ideological significance
- ▶ Balanced interpretations that, rather than isolated observations, connect findings to broader contexts and theoretical constructs and critically discuss limitations
- ▶ Articulation of significance and implications for the design discipline, design practice, and society
- ▶ Accessible writing style that combines scholarly rigor with clear prose and rich visual examples, serving the journal’s broad readership (practitioners, academics, and students)
- ▶ Stated ethical approval for the study and/or copyright permissions to reproduce images, as relevant to the research
- ▶ Submission of high-quality images following *Visible Language* guidelines

Structure and style

Submissions must follow the journal guidelines (with APA citations) and house style and include:

- ▶ Abstract (100–200 words)
- ▶ Introduction (with clearly stated research question)
- ▶ Research context (literature review, rationale, objectives)
- ▶ Methods
- ▶ Findings (including examples of visual material/data, as appropriate)
- ▶ Discussion (interpretation of results and implications)
- ▶ References (APA)

<https://www.visible-language.org/journal/calls-and-submissions/>

Interpolation. Lowercase z's from Recursive, a variable font designed by Stephen Nixon (ArrowType). Interpolation is shown for two variable axes: weight (300...1000, from top) and casual (0...1, from left). Additional axes are fixed across the samples: monospace (0), slant (-12), and cursive (1).



inside

back

cover

iv–vi

The Editors Introduce the April 2026 Issue

Jeanne-Louise Moys, Mike Zender, and
Matthew Peterson

1–23

**The Terminological Development of
Graphic Design: Between Office Art and
Social Purpose**

David Preston

24–57

**Automation and Artificial Intelligence in the
Type Design Process: Insights from an
Industry Survey**

Alice Savoie, Kai Bernau, Wayne Daly, Raphaela
Haefliger, and Sebastian Baez-Lugo

58–83

**Evaluating Interactive Highlighting
Techniques in Digital Reading: An Empirical
Study of Hover-Based Line, Sentence, and
Paragraph Highlighting**

Szabolcs Vatány, Thi Huyen Nguyen, Anikó Illés, and
Ann Bessemans

84–89

Why Meta-Font Struck a Nerve

Mike Zender

90–119

**Metafont, Metamathematics, and
Metaphysics: Comments on Donald Knuth’s
Article “The Concept of a Meta-Font”**

Douglas R. Hofstadter

120–123

**A Profession Provoked: How Meta-Font
Struck a Nerve**

Deborah Littlejohn

124–125

Call for Papers: Typographic Landscapes