Visible Language

The journal of interface, experience, and communication design research

Volume 59 Number 2

August 2025

ISSN 0022-2224

Published continuously since 1967

inside

front

cover

Visible Language

The journal of interface, experience, and communication design research

Volume 59 Number 2

August 2025

ISSN 0022-2224

Published continuously since 1967

Visible Language

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:

Contents

EDITORIAL

Visible Language

Volume 59 Number 2

Reflecting on the August 2025 Issue — Considerations Nowadays and Implications For

Matthew Peterson

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

1. Considerations Nowadays: Identity, Expertise

What are we? As Maria dos Santos Lonsdale (2025) revealed in the April issue's introductory editorial, this issue takes as its consideration *the present,* following that issue's consideration of *the past.* The articles reveal various perspectives on our present. Well, "our" depends upon who is doing the reflecting. Ramanathan (2025), Zhang et al. (2025), and Medley and Haddad (2025), all in rather different ways, discuss cultural or personal identities that must be addressed in the things we design and the ways we go about designing them. An outsider might think that these external factors are what our field would struggle with — as we do — but that we would have internal clarity about our field's identity. Speaking for myself — as an editor of a journal in the field, and as a senior faculty member who instructs graduate students on the nature of the field and its scholarship — I have no such clarity.

Three practicing designers published in this issue make a point of acknowledging the presently dynamic nature of our field's identity: Hall (2025) outlines four new definitions for designers in their *evolving roles*: as advocates, curators, orchestrators, and emotion mediators. Shell (2025) focuses on *evolving responsibilities,* calling on designers to be the "human in the loop" in a world where it is no longer safe to assume that creative production is human production. Gupta (2025) also invokes *evolving responsibilities,* including not just the detail of expertise in collaboration but the defining

@: ISSUE〉 ARTICLE〉

Cite this article:

Peterson, M. (2025). Considerations nowadays and implications for. *Visible Language, 59*(2), iv–xii.

First published online August 15, 2025.

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

Visible Language Consortium:
University of Leeds (UK)
University of Cincinnati (USA)
North Carolina State University (USA)

title of "AI designer." All three of these designers have written in response to the rapid proliferation of commercial artificial intelligence technology, especially large language models and diffusion models.

But our field — again, "our" is not straightforward — was destabilized long before AI. When I began my studies of graphic design in the mid-1990s, the popular internet was just emerging, and the image element in HTML ("img") was recently invented (Andreessen, 1993) but not yet adopted in the web standards (Berners-Lee & Connolly, 1995). As the web standards were developed, some of us in graphic design began designing and coding websites. These included text and imagery, just like books and posters did, but they were displayed on monitors, in browser windows that varied in size and in their interpretation of code, and were responsive to scrolling inputs, mouse clicks, and hover states. Of course, even earlier graphic designers became involved in the design of personal computer interfaces, and before that other important changes occurred (see Bigelow & Holmes, 2025, in the April issue, for significant detail on such changes specifically in regard to font design). Mobile devices presented another set of screens and elevated the need for designers to create responsive typographic systems, with a wider range of display sizes and processing limitations.

The point is that *what we design* has been changing for decades. I think this is a more dramatic case of flux than other fields in design have faced. For instance, while architects have had to adjust to continuously changing tools and dependent building technologies, they are still designing buildings. Of course, some graphic designers are still designing books. So we have not coolly shifted along with technology. We are doing something else. It is just not clear to what degree it is expanding "our" expertise, or if "we" are fracturing.

I have used the term *graphic design,* but my colleagues at *Visible Language* at the Universities of Leeds and Cincinnati use *communication design.* I prefer the former, because I think it better incorporates user interface design, which relies on graphic devices but empowers people to do things instead of communicating to them. (That is, I think *communication* can be a limiting construct for conceptualizing user interfaces.) But that is not a robust argument, and *communication design* conversely has a level of abstraction that is probably more adaptable. The American Institute of Graphic Arts — *Arts* — now goes only by its once abbreviation, AIGA, which is paired with the more contemporary and noncommittal "professional association for design" (AIGA, n.d.). *Visible Language* was originally called *The Journal of Typographic Research,* but Merald E. Wrolstad (1971) decried "having to add a footnote every time the name [was] mentioned in order to explain its actual range of interests" (p. 5). Advances in reproduction technology likely had an influence on this decision, as one designer became capable of producing more kinds of things — typography *and* imagery — with existing technology and at an accept-

able level of competence. Returning to *graphic design,* my academic department recently changed our degree-granting programs from that to *graphic and experience design,* and in truth a presently more accurate title for the degrees would probably be *visual, interface, and experience design* (*visual design* is often used in software development contexts, especially when the term *graphic design* is still applied to other employees who focus on branding). And I will end this parade of titles by noting that as of 2024 *Visible Language* was described as covering *visual communication design,* but we now list it as *interface, experience, and communication design.*

What do we know? It gets worse if we consider what our *discipline* is, or what research and scholarship is and is not for us. I will not digress fully into those questions. But our research and our practice both rely on some base of knowledge, ability, or specialty — even if we often learn and enact our kind of design through more tacit pattern knowledge (Kolko, n.d.) and reflection-in-action (Schön, 1984). In my own consideration of these issues, I continually return to the *things* we can purport to be expert in that others cannot. I do not think generative AI is going to replace this expertise. I think generative AI crystallizes it. If everybody can make the same kind of bland derivative thing, then expertise in evaluating outputs — decision-making surrounding words, pictures, and their configurations in space and time — becomes all the more valuable.

Perhaps a pressing need for us is thus to make our tacit knowledge more explicit, and to not just assume, but to better understand how our immersive studio-based education and our highly flexible making practices can be leveraged. For instance, a design educator in our field may decide that in the past we learned about sets of relationships we can roughly call *semiotics* in large part — not exclusively — through long hours of making, seeing, questioning, and critiquing. And they may say that with production technology becoming more accessible and dramatically faster, to retain an expert "we," we must now be increasingly formal in how we learn such things. (I suppose I am that "educator" in this "instance.")

Academic journals are access points for knowledge generation, and this issue of *Visible Language* provides guidance for advancing our field through expertise retention and crystallization. Zhang et al. (2025) address the kind of nuanced understanding of typography that we must have if we are to be the experts, while Medley and Haddad (2025) consider the kinds of subtleties in imagery we need our students to invest in, if they are to have a say in the pictorial world we all inhabit. Armstrong et al. (2025) outline something they do not quite call a process, as a way to retain what we know about how we make — through visual exploration — while contributing to technology development.

It feels like an oversimplification. But we do make, select, and arrange typography and imagery, wherever and however they are instantiated, and expertly.

2. In this Issue

As reflected in the table of contents, this issue is divided into two parts:

1. Looking closely at type, text, and image (pages 109–175), and
2. Artificial intelligence and interface design (pages 176–251).

But there is another sequential organization related to our new publishing model. *Visible Language* is now diamond open access. The *diamond* distinction means that not only is it available free to readers, but authors do not incur any charges when their manuscripts are accepted, typeset, and published. This is in part enabled by the immediacy and lower cost of online publishing. (The other part is that the editorial board works to ensure financial support for the production costs that commitment of our own time cannot reduce, in service of the discipline.) Emphasizing online publication allows us to present some articles as "early views" in advance of indexing in one of the three issues per year. But *Visible Language* also has a history — beginning in 1967 — of printed volumes that are held in complete collections in some libraries (and by subscribing individuals as well). So at least for the time being, we are collecting issues in volume yearbooks that can be purchased as physical books at the end of the year. This means that the main articles of the April issue will be united with this August issue's articles (and the December issue's articles) in a continuous format. Knowing this, let me reveal the other structure:

- Invited article from the *past*-themed set in issue 59.1:
  - Rathna Ramanathan: "Research-Led Pluralist Typographic Practices: Case Studies from South Asia." Publication of this article was delayed, but by placing it first we will reunite it with the other invited articles in the yearbook.
- Author-submitted peer-reviewed articles — the lifeblood of academic journals:
  - Yuchan Zhang, Jeanne-Louise Moys, and Matthew Lickiss: "The Role of Text Alignment on Response Speed and Accuracy When Reading Chinese-English Bilingual Traffic Signs."
  - Stuart Medley and Hanadi Haddad: "Breaking Images: A Method for Improving Design Students' Visual Literacy."
  - Helen Armstrong, Ashley L. Anderson, Rebecca Planchart, Kweku Baidoo, and Matthew Peterson: "Addressing Uncertainty in LLM Outputs for Trust Calibration Through Visualization and User Interface Design."
- Curated articles from industry professionals as *Dispatches from Industry*:
  - Syashi Gupta: "A Seat at the Table: Designing for AI with Strategy, Vision, and Collaboration."

- Will Hall: “The Changing Definition of Designers in the Age of Generative AI.”
- Sierra Shell: “The Human Touch(point): Recommendations for Thoughtful AI Feature Design.”

In the volume yearbook, these articles will be followed by others authored by design students (see the call for papers in the April issue). The volume sequencing will thus be:

1. Invited articles from luminaries in the field,
2. Submitted articles from academic researchers subjecting their ongoing work to peer review,
3. Curated articles from industry professionals who would not otherwise have had publication in an academic journal in mind (but may now), and
4. Submitted and expert-reviewed articles from students who will soon be academic researchers or industry professionals.

In this way, the 2025 volume of *Visible Language* tells a story of scholarly engagement in reverse order.

3. Implications for: Research, Practice

In our last issue, Nigel Cross (2025) outlined the strides in design research and design understanding that have been made in the broad discipline. We are building on decades of progress. This is followed by Meredith Davis’s (2025) description of a stagnating educational paradigm in design that is faced with an ongoing paradigm shift in practice, largely under addressed for decades. Elsewhere Deborah Littlejohn (2023) notes a disconnect between the design profession and design education that holds design back (well, *graphic design* as she calls it there) from becoming a fully healthy *discipline*. To echo my own professed reflection and indecision earlier in this editorial, she notes that part of the “core problem” is that graphic design “lacks agreement from the field’s experts about what it is and what it is about” (Littlejohn, 2023, p. 55).

The editorial board of *Visible Language* is highly invested in the maturation of our field — whatever we call it, or whatever we call *them* if there is fracturing — into a proper discipline. There are two initiatives evident in this issue through which we hope to help develop our discipline.

First, we have established a column called *Dispatches from Industry* (pages 218–251), for which we have contacted industry professionals, engaged them in conversation, and when that conversation led to what we considered appropriate content for the journal, invited them to contribute. The initial *Dispatches* — Gupta (2025), Hall (2025), and Shell (2025) — all address AI, but that is a factor of curation in the moment (a moment for which AI has currency for both design and academic publishing).

We think *Dispatches from Industry* is highly valuable, as evidenced by the quality of these three entries. But beyond those individual pieces, the initiative itself has great potential because it surgically targets a weakness in our discipline (or field): an unhealthy relationship — really, too little of a relationship — between industry and academia. Industry works on short time scales, often measured in two week sprints and inflamed in reorgs and corresponding layoffs as companies constantly flow with changing conditions. Meanwhile in academia, concepts of contributions turn into manuscripts, which turn into articles, over one or two years. Funded projects last even longer, and dissemination of their results trails them further. Neither is a problematic model. There is a difference in user research and basic research reflected in their parent organizational structures. But user research and the practices tied to it are greatly enhanced if they are informed by basic knowledge. Furthermore, academic researchers cannot make all of the best decisions about what knowledge to pursue if they have only a cursory understanding of industry trends. We desperately need to bridge the gap that exists. We hope that *Dispatches from Industry* helps to do so.

Second, astute readers may have noticed in our last issue an additional prefatory element between the abstract and keywords in Van der Waarde and Thiessen (2025): an *implications for practice.* This prefatory element is not a straight summary, which would be redundant with the abstract. It is a guide for reading or *accessing* (as in a database) the article from a practice perspective. It seeks to proactively answer the practitioner's question: How is this academic article useful for me? I think it is helpful here to compare an abstract and *implications for practice.* So, for the same article:

> **Abstract:** The number of experiments that investigate the "readability" or "legibility" of texts is very substantial. Literature reviews of these studies appear regularly, and many publications refer to these experiments to suggest evidence for claims. Some of these claims have led to usable recommendations. However, most of these recommendations are often hard to apply and unhelpful. When we are teaching typography, we struggled to explain why the recommendations are difficult to use, why many reviews are uncritical, and why experiments rarely provide reliable evidence to support design decisions. A literature review, guided by experience in both commercial practice and university level education, lead to a list of themes and issues. There are at least 19 reasons why the results of many typographic experiments need to be questioned. This article provides 19 guidelines that could be used to evaluate experimental research into the ways in which texts are read. This list of reasons can be used as a checklist to assess and guide new typographic experiments. We hope to make sure experiments are worthwhile, future reviews are based on reliable sources, and recommendations are effective. (Van der Waarde & Thiessen, 2025, p. 77)

As with all proper abstracts, this embodies a form of problem statement and justification, it previews what the article will be like, and it outlines contributions. And while it does end with some guidance on how the article might be *used,* the following is far more specific and written with the practitioner's perspective in mind:

> **Implications for practice:** There are three practical applications of the findings of this review. Firstly, the 19 guidelines might help to critically review experimental findings and assess if they are relevant for practice — Table 1 is a handy checklist for this assessment. Secondly, the review shows that a typographic practice must be reader-focused. It is essential to involve readers throughout design processes, especially when the intention of information is to enable people to act. Performance criteria, evaluation methods and performance levels need to be relevant for readers. The result of this involvement is qualitative: a single remark from a single person can change the frame of a design project. And thirdly, the review shows that it is beneficial to look more intentionally at differences across readers and across reading activities. Involving people with different backgrounds and experiences will provide new insights into the ways visual information can be interpreted and applied. Listening and observing are fundamental design skills that need to be trained and honed. Even after decades of practice, it remains a humbling experience to find out how people really look at visual design. (pp. 77–78)

The *implications for practice* does not attempt to cover all of the content. For instance, this one directs the reader to Table 1 for quick access. In a proper reversal, and still in the interest of bridging academia and industry, the *Dispatches from Industry* columns in this issue include a counterpart *implications for research.* Here is an example from Shell (2025):

> **Implications for research:** This article raises issues that suggest fruitful research areas. The recommendation to compare a potential AI-driven design solution to a non-AI solution (Section 2) suggests the development of a conceptual framework for systematically assessing and justifying aspects of AI specification. Such a framework should be compact enough to integrate into R&D processes for ground-level impacts. The principle of thoughtful friction in AI product design (Section 3) is introduced in relation to frictive AI elements of granular typographic detail such as point size and positioning. Empirical research could determine which typographic and interface characteristics provide friction for contemporary users, to what degree different options do so, and what thresholds exist for users' emotional and cognitive responses. Further research could examine whether such frictive elements effectively facilitate trust calibration for potentially erroneous AI-generated content. Finally, the

> recommendation to integrate feedback mechanisms into AI-driven products (Section 4) suggests research aimed at understanding the current state of AI feature feedback collection in popular products, both in terms of mechanism and frequency. (pp. 241–242)

If you are a design researcher, above you can find clear direction for work that you know would be valued in industry. In this example, sections are explicitly referenced to guide the reader to areas of interest. But as an open access journal, *Visible Language* articles like these are available online and will be read not only by people, but also by machines. It is hard to predict how search engines will behave or continue to evolve (or how LLM summarization will do the same), but the *implications* elements are expected to aid in connecting potential readers to *Visible Language* content, even if they are unaware of the journal's existence. This could be impactful in bridging the industry–academia divide.

This issue's articles have many implications for research and practice, only some of which are collected under those prefatory headings. We hope our readers are able to find some that are particularly relevant to their own research, practice, or instruction, in whatever it is they call their areas of design.

4. References

AIGA (n.d.). Home | AIGA. *aiga.org*

Andreessen, M. (February 25, 1993). Proposed new tag: IMG. *World Wide Web History Center.* http://1997.webhistory.org/www.lists/www-talk.1993q1/0182.html

Armstrong, H., Anderson, A. L., Planchart, R., Baidoo, K., & Peterson, M. (2025). Addressing uncertainty in LLM outputs for trust calibration through visualization and user interface design. *Visible Language, 59*(2), 176–217. https://www.visible-language.org/journal/issue-59-2-addressing-uncertainty-in-llm-outputs-for-trust-calibration-through-visualization-and-user-interface-design

Berners-Lee, T., & Connolly, D. (September 22, 1995). Hypertext Markup Language – 2.0. *W3C.* https://www.w3.org/MarkUp/html-spec/html-spec_toc.html

Bigelow, C., & Holmes, K. (2025). Digital type challenges. *Visible Language, 59*(1), 23–54. https://www.visible-language.org/59-1/digital-type-challenges

Cross, N. (2025). Making design research visible. *Visible Language, 59*(1), 1–4. https://www.visible-language.org/59-1/making-design-research-visible

Davis, M. (2025). A shifting practice paradigm meets a persistent curriculum paradigm. *Visible Language, 59*(1), 5–22. https://www.visible-language.org/59-1/a-shifting-practice-paradigm-meets-a-persistent-curriculum-paradigm

Gupta, S. (2025). A seat at the table: Designing for AI with strategy, vision, and collaboration. *Visible Language, 59*(2), 218–231. https://www.visible-language.org/journal/issue-59-2-dfi-a-seat-at-the-table-designing-for-ai-with-strategy-vision-and-collaboration/

Hall, W. (2025). The changing definition of designers in the age of generative AI. *Visible Language, 59*(2), 232–240. https://www.visible-language.org/journal/issue-59-2-dfi-the-changing-definition-of-designers-in-the-age-of-generative-ai/

Kolko, J. (n.d.). *Teach slower to teach better.* https://www.jonkolko.com/writing/notes/teach-slower-to-teach-better

Littlejohn, D. (2023). Suspended between discipline and profession: A history of persistent immaturity and instability in the graphic design field. In G. Kaplan (Ed.), *After the Bauhaus, before the Internet: A history of graphic design pedagogy* (pp. 48–59).

Lonsdale, M. d. S. (2025). In with the new! *Visible Language, 59*(1), v–x. https://www.visible-language.org/59-1/in-with-the-new

Medley, S., & Haddad, H. (2025). Breaking images: A method for improving design students' visual literacy. *Visible Language, 59*(2), 153–175. https://www.visible-language.org/journal/issue-59-2-breaking-images-a-method-for-improving-design-students-visual-literacy/

Ramanathan, R. (2025). Research-led pluralist typographic practices: Case studies from South Asia. *Visible Language, 59*(2), 109–129. https://www.visible-language.org/journal/issue-59-2-research-led-pluralist-typographic-practices-case studies-from-south-asia

Schön, D. A. (1984). Problems, frames and perspectives on designing. *Design Studies, 5*(3), 132–136. https://doi.org/10.1016/0142-694X(84)90002-4

Shell, S. (2025). The human touch(point): Recommendations for thoughtful AI feature design. *Visible Language, 59*(2), 241–251. https://www.visible-language.org/journal/issue-59-2-dfi-the-human-touchpoint-recommendations-for-thoughtful-ai-feature-design/

van der Waarde, K., & Thiessen, M. (2025). Nineteen questions to evaluate typographic research: Chaff and wheat. *Visible Language, 59*(1), 77–99. https://www.visible-language.org/59-1/nineteen-questions-to-evaluate-typographic-research-chaff-and-wheat

Wrolstad, M. E. (1971). Visible Language: The journal for research on the visual media of language expression. *Visible Language, 5*(1), 5–12.

Zhang, Y., Moys, J.-L., & Lickiss, M. (2025). The role of text alignment on response speed and accuracy when reading Chinese-English bilingual traffic signs. *Visible Language, 59*(2), 130–152. https://www.visible-language.org/journal/issue-59-2-the-role-of-text-alignment-on-response-speed-and-accuracy-when-reading-chinese-english-bilingual-traffic-signs

INVITED ARTICLE

Research-Led Pluralist Typographic Practices: Case Studies from South Asia

Rathna Ramanathan

Central Saint Martins, University of the Arts, UK (r.ramanathan[at]csm.arts.ac.uk)

Abstract: This article is grounded in an exploration conducted by the author on publishing as a platform that brings intercultural communication, pluralism, graphic design and typography into productive dialogue with each other through engaged (in social and political issues; in creative, educational, and critical practice) and situated (local communities; international networks of editors, translators, designers, illustrators, publishers, and readers) design research frameworks and practices. This has resulted in an exploration of spaces in which new kinds of documents can be created, with, by and for marginalized publics, and, conversely, how the production of new texts and images creates spaces that enable emancipatory, temporary, or subversive practices to occur that suggest new directions for the practice of typography and typographic frameworks. This exploration through design research and practice, is framed by the author's own context, as that of a South Asian designer and researcher, working in the Global North.

Some of the initial thinking in this article was explored in a chapter for *The Routledge Companion to Design Research* — 2nd Edition. The article takes a holistic, post-disciplinary approach to graphic design and typography aiming to challenge notions of graphic design as purely aesthetic or craft-based, or as concerns of form and function. It calls for a shift in considering the wider politics and contributions of visual language — graphic design and typography specifically — to societal change. Additionally, it reframes research-led practices (and thereby visual language and typography), not as an elite activity but as a human practice that emerges as curiosity and intent. Such an approach is critical to undertake considering a global health crisis, climate emergency and with issues of conflict and social injustice where communication plays a pivotal role. The article concludes that how we approach design research and practice needs to be rethought so that it makes a meaningful contribution to planetary issues.

@: ISSUE〉 ARTICLE〉

Cite this article:

Ramanathan, R. (2025). Research-led pluralist typographic practices: Case studies from South Asia. *Visible Language, 59*(2), 109–129.

First published online August 15, 2025.

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

Visible Language Consortium:
University of Leeds (UK)
University of Cincinnati (USA)
North Carolina State University (USA)

Implications for practice: A holistic, post-disciplinary approach to graphic design and typographic research can challenge notions of graphic design as purely aesthetic, or as concerns of form and function, and speak to the shift needed in considering the wider politics and contributions of graphic design to societal change. Latin (Western) approaches to typography offer a singular view of typography as functional and rational. However, pluralistic approaches make more visible, through design and documentation, a broader approach to typography which acknowledges typography's link to language, as it is spoken, written, and read both culturally as well as materially. The history of the book which still looks primarily at the codex, needs to encompass the histories that are beyond the codex, to manuscripts, scrolls and other "book" traditions which are rarely documented or acknowledged. Where little evidence exists, historical practices can provide guidance for contemporary design frameworks and guidelines. With each of these contexts, research revealed approaches to similar design problems by designers, typographers and publishers in the past. Speaking to the contribution of practice, we must take an approach that suggests that we can design the means through which design happens, challenging the concepts, behaviors, and means of production as well as designing form.

Keywords: design research; Global South; graphic design; intercultural communication; publishing; typography

1. 'South Asia as a Site of Investigation'

In the article "Crisis in the Classics," Sheldon Pollock (2011) draws on a classic humanities dilemma. In a world of big problems, he notes, a small problem (in this case, the disappearance of classics — and classical textual knowledge — in the Indian subcontinent) can seem minuscule and seemingly irrelevant. Pollock, however, draws an important and critical connection that is relevant to our own practices and research: the small problems (or the perception of our problems as small) contain one very big question about what it means to be fully and richly (and diversely) human. This is surely the core of our task as visual communicators — past, present and future.

This article features case studies of three publishing projects — Harvard University Press' Murty Classical Library of India series (Figure 1); and Tara Books' *In the Land of Punctuation* (Figure 2) and *Liberté* (Figure 3) — which evidence a pluralistic and intercultural approach. As both examples are anchored in an Indian context of publishers and/or readers, one could question the relevance of this in the wider realm of visual language and typographic knowledge and practice. Yet this is precisely the point; rather than think of India as a national identity or a limited geographical space, the approach suggests using India as a framework in the manner suggested by Pinney (2013, p. 172). India thereby becomes a site of investigation in which you can develop a typographic framework or model that is relevant and potentially transportable to other models and

Figure 1. Front cover designs for the Murty Classical Library of India (MCLI) titles published 2015–2025.

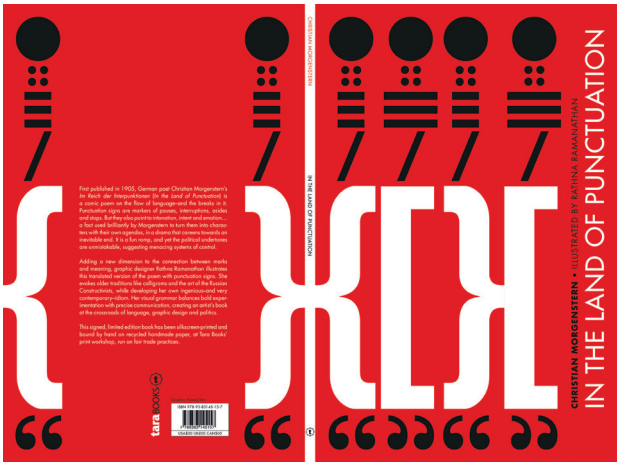

Figure 2. Cover design for Tara Books' *In the Land of Punctuation* (2017).

Figure 3. Cover design for Centre Pompidou and Tara Books' *Liberté* (2024).

contexts. This is particularly critical when establishing an equity and multiplicity in knowledge production in design research and practice.

To turn our attention to knowledge production in India, there has been coverage in media of "alternative facts" but as anyone from an oppressed or colonized society will note, alternative facts have existed as long as we have been writing history. This is often the basis of colonization (i.e., to present reality in a manner which suits one's own power, needs and contexts). One only needs to look at T.B. Macaulay, the British historian who oversaw introducing English concepts to education in India. When presenting on his findings, Macaulay (1835) dismissed Indian knowledge based on its difference. He refers to Indian history, astronomy, medicine, and religion as false, thereby dismissing hundreds of years of knowledge. Such an approach becomes critical and dangerous when employed by media platforms which are owned by and deployed for political gain all intent on a consistent "othering" of marginalized groups to achieve their purpose.

The colonial legacy is a painful legacy. Trivedi (2008) illustrates an example of how Indian knowledge was colonized using the *Hortus Malabaricus* ("Garden of Malabar"), a comprehensive treatise that documents the properties of the flora of the Western Ghats, a mountain range in India that crosses the states of Goa, Gujarat, Karnataka, Kerala, Maharashtra, and Tamil Nadu. Written in Latin and compiled over 30 years, the series was conceived by Hendrik van Rheede who was then Governor of Dutch Malabar and contains pen-and-ink-wash drawings of some 720 species which are accompanied by a detailed description in Latin. Apart from Latin, the plant names are included in Malayalam, Konkani, Urdu, and English. What is deeply troubling about this text is that whilst it was collated and compiled by "natives" as they are referred to — Indian experts in the field — it was available only in Latin until the 21st century. This text

has been largely inaccessible previously because it was not available in any Indian language. Knowledge about India, written with Indian knowledge has been inaccessible to Indians.

The origins of publishing and printing in India are entangled with colonial ambitions. Boaventura de Sousa Santos (2018) notes that these ambitions sought to discredit, erase, or appropriate the knowledges of the Global South with the aim of contributing to a dominant Global North knowledge and culture. For Tamil language speakers, the first time they read their language in printed form was to communicate a religious text that was alien to their culture (and beliefs). There is a tremendous power in this act of publishing — to use someone's language to represent back to them a culture and religion that is not their own. Who decides what is knowledge and who this knowledge is for? What is knowledge if language and the visual form prohibit people from accessing them? And what role do we play in this as researchers and designers who frame knowledge for reading and in addressing equity in knowledge production? As noted by Ansari (2020):

> Decolonization entails not only serious political commitments but epistemological ones: one has to engage with the colonial and precolonial past in order to arrive at a more nuanced and critical understanding of the present. (p. 8)

What of visual language and its importance? In a multipolar, global world of multiple experiences of being human, the way the world is presented back to us is decidedly narrow. Through social media, through packaging of global brands, through mainly consumer culture. To paraphrase Pollock (2011), we learn to know the world and ourselves through language — oral, written, and increasingly, sophisticated visual language. Language shapes the way we experience the world and way we use it as a tool influences what we nurture and care about. If we continue to normalize mainstream media approaches as the singular approach, we start to limit how we might enable others to describe and articulate their lived experiences and their cultures.

November, a design duo based in India who work on global visual language projects through typographic and type design interventions, noted in a conversation for the Walker Reader (Bhatt & November Studio, 2020), an aim to take a broader approach to their practice in order to be more representative of contextual needs of design and typography that rises from multiple cultures and modes of writing, reading and publishing:

> We emphasize the values of plurality, which are inherently democratic. It is about accepting and respecting the possibility of a multitude of experiences, opinions, and ways of life.

2. Pluralistic Approaches: Harvard University Press and the Murty Classical Library of India

Many classical Indian texts have never reached a global audience, and others are inaccessible even to Indian readers. The Murty Classical Library of India (MCLI) is a 100-year publishing project at Harvard University Press that aims to make available the great literary works of India from the past two millennia to redress this imbalance. The series provides modern English translations — many for the first time — alongside a vast number of Indian languages. The text in the appropriate regional script appears alongside the translation. Rohan Murty who envisioned MCLI was inspired by his own experience of education in India, and it is one that many middle-class, urban Indians, identify with. The texts that were studied in school were Shakespearean comedies and tragedies, poems from Wordsworth and Shelley, stories by Hardy and Kipling. However, missing from it was the same opportunity to partake of one's own classics and heritage.

There were several design challenges in this project. The first was at the time of the inception of MCLI, no typefaces existed that could set the range of characters in the texts in a manner that was readable, and accessible. Harvard Press commissioned a series of typefaces designed specifically for the library by Professor Fiona Ross (University of Reading) and John Hudson (Tiro Typeworks). The MCLI work featured here is not of typeface design but of the interior book design frameworks for 30 bilingual volumes, and typographic design for 19 bilingual volumes in Indian languages with English translations, as well as design and typographic guidelines in prose and poetry genres for several Indian languages including Apabhramsha, Avadhi, Bengali, Hindi, Kannada, Pali, Panjabi, Prakrit, Sanskrit and Telugu; and the Bangla, Devanagari, Gurmukhi, Kannada, and Telugu scripts. The combination of typesetting and design of bilingual Indic texts is unprecedented. These volumes were published in two editions, hardback for the scholarly market in the US and UK and paperback for the Indian popular market. The challenge of this project was to find contemporary design solutions to classical texts (pre-1800) whilst retaining their spirit and originality. The research was instrumental in supporting the expansion of readership in inclusive, decolonial and intercultural ways. This was achieved by creating a comprehensive typographic research framework for Indic scripts to preserve threatened narratives and to improve access and enhance reading for marginalized groups.

3. Tara Books and *In the Land of Punctuation* and *Liberté*

Tara Books is an Indian publisher founded in 1994 by a group of writers and designers committed to egalitarian principles. Tara was interested in changing the perspective from which stories are told which meant expanding the notion of authorship, the notion of the book and its content, and the role that design plays in the publishing process.

Publishing at Tara is reframed as a collaborative enterprise where the success of a book cannot be attributed to one individual because it is by nature, dialogic, collective, and heavily dependent on the work of others at every stage.

In an interview, publisher Gita Wolf (2021) refers to publishing as a cyclic conversation:

> We think of the book as a moment in time, a picture of a much longer process. There is a story of how the book was made, and then you have the book itself, and once the book is published you have the entire story of how it is received, and what else happens as a result of that reception.

The work with Tara Books is about giving a voice to marginalized people who do not normally get a voice, through the act of publishing. *The London Jungle Book* (2017) by Gond artist Bhajju Shyam is titled as such as both a homage and mirror-image counterpoint to Kipling's *The Jungle Book* (1894) and tells the story of Bhajju's journey to London from India. The book has a layer of historical significance: A century earlier, Bhajju's tribe had been studied by the British anthropologist Verrier Elwin, who married a Gond woman, and wrote several books about the tribe. Bhajju's grandfather had been Elwin's servant, so he had grown up with the writer's stories. Elwin had written in the preface to one of his books on the Gonds that he considered it a counterpart to Kipling's *Jungle Book. The London Jungle Book* was summarized by Bhajju (2017) with a decolonial statement of intent: "Elwin sahib wrote about my tribe, now it is my turn to write about his."

The other way of expanding reading that Tara Books explores, is through typography. Tara sees typography (to paraphrase Beirut, 2019) as a fundamental way to understand and engage with the world. Tara's approach to picture books challenges conventional separations of image and text and blurs the boundaries of what text or image should do. Research and expertise in new approaches to typography as well as nonstandard ways of designing and producing books informed a collaboration with Tara Books and the publication of experimental picture books *In the Land of Punctuation* and latterly, *Liberté.* These books employ typography as illustration. Research for the books drew from the understanding of how typography in children's books takes primarily a conventional Global North understanding, with text and image separated. This is counter to the understanding that we might experience word and image as equally visual, and particularly in India and Asian cultures where reading is a visual act.

In the Land of Punctuation as a project was motivating for three reasons. First, the work was out of print in the English language and available only in German, so it is mostly unknown to contemporary readers of English. The publisher felt the text and the context was still relevant and should be made available to a wider audience. Second, from a subject perspective, typography in the picture book context has, like much

of its content, tended to the safe and the cute. The text, due to its political content, called for research into typography and type play for more serious communication purposes. This was interesting within the context of a picture book as a literary but also a social, cultural, economic, and political product. And third, the project questioned the adherence to certain cultural norms. Building on the aim of equity in knowledge production, it was important to challenge the notion that a German poet should only be published in a Western context and only Europeans should work on European projects.

Publishers at the Centre Pompidou were taken with the approach in *Punctuation* and commissioned a special project using the same typographic framework, methods and research to celebrate the centennial of the Surrealists. Their focus was building from French Surrealist poetry but speaking to the international impact and reach of the Surrealist movement. There was a freedom in this commission, to choose an appropriate and relevant poem for our times. The final choice was Paul Eluard's *Liberté,* written in 1942 during the German occupation of France. Printed additionally in London by the Gaullist magazine *La France Libre*, thousands of copies were dropped by parachute by the Royal Air Force over occupied France.

4. Intercultural and Pluralistic Practices

This article is built on the premise that a contemporary and relevant approach to graphic design and typography necessitates a twofold understanding that i) design is not solely a craft, but a fundamental way to understand and engage with the world (Beirut, 2019), and ii) this requires the acknowledgement of non-mainstream, often marginalized approaches to the discipline, in particular, intercultural and pluralistic knowledge.

4.1. Typographic Research and Practice Beyond Global North Conventions and Understanding

Typography is visualizing language. As noted by several authors (Calvert, 2012; Gruendler, 2005; Lees-Maffei, 2019), since Beatrice Warde's proclamation in *The Crystal Goblet, or Printing Should Be Invisible,* typography in the Western tradition aims to establish a clear sense of "good" and "bad." Warde made several distinctions of the "good" which gave prominence to the form of typography over intention, and context. Warde's approach has framed modern typography thinking and is defined by Anglo-specific industrial, linguistic, and social contexts, i.e., the letterpress, which converts the page into a grid, Latin languages (predominantly English), and Western publishing, wherein the author (and thereby their words) is given primary importance. There is no acknowledgement nor understanding of other cultures, spoken language, or associative forms of typography, thereby creating a sense of hierarchy and marginalizing or othering other practices.

For example, in the context of the Indian subcontinent, where lithography preceded letterpress and letterpress was introduced with colonial intent, the form of the book was not the codex — the page was visual and spatial rather than linear and chronological, and the reader rather than the author was given prominence. Forms of typography that are associative with movement, sound, texture, particularly in relation to poetry produced by little presses, remains unrecognized beyond key figures such as Cobbing, Hamilton Finlay and Houedard. These (now marginalized) histories are rarely recognized as a part of design research, design history or practice. To extend Fry (2007), "[typography] is profoundly political. It either serves or subverts the status quo" (p. 8).

In the Land of Punctuation and *Liberté* explore the potential of a word-image visuality in typography. Design-led conversations and participatory reading sessions, and archival research which led to analysis of secondary and primary sources of ephemera from India (posters, murals, street signs) and from French and Russian archives (catalogs, publicity material, original artwork) inform the books. The research established visual examples of associative typography, wherein typography is concerned with the meaning and interpretation of the text and representing it using visual, verbal, and spatial aspects of typography. Typography in picture books takes primarily a conventional Global North understanding, with text and image separated. This is counter to the understanding that children in South Asian contexts grow up in an environment where word and image are equally visual; where reading is a visual act.

4.2. Intercultural Approaches to Typography and Book Design

With the MCLI series, the typesetting and design of bilingual Indic texts of such range and complexity is unprecedented in modern book design practice and posed multiple challenges that were addressed through three lines of enquiry. First, to establish a systematic bilingual book design for English translations of texts in ten different Indian languages and scripts grouped into four categories, namely, North Brahmic (Sanskrit, Hindi, Gujarati, Bengali), South Brahmic (Tamil, Telugu, Malayalam, and Kannada), Perso-Arabic (Urdu, Sindhi) and Prakrit (Pali). Second, to accommodate two genres — poetry and prose — in the template design. Third, as Indian texts do not use italics or bold, it was imperative to establish an Indic hierarchy and grammar through the application of typographic rules. There is a lack of attention to printing and typographic conventions in India as well as a lack of standards for typesetting modern Indian languages, as documented by Deshmukh (1958) and Ramakrishnan (2010). In addition, examples of bilingual design frameworks account for 3–4 different languages at most; here the task was to accommodate at the least the starting mission of 13 different languages and relevant scripts. The typographic and book interior designs aimed to recognize that some readers would be fluent in the language, while others might be

second-language or third-language speakers or not know English at all. It was essential that equity of access was provided for readers of all language fluencies.

With *In the Land of Punctuation* and *Liberté,* it is design and typography that situates universal narratives within a local context. The text was originally a 1928 German poem by Christian Morgenstern about politics, oppression and war that is recontextualized in a modern Indian setting and brought back to life. As noted above, Eluard's French poem spoke of freedom. The book becomes a research space to understand the politics that surrounds typography and language, where "politics" refers to the power that aesthetics that the visual and typography can carry as a voice and as a language in itself. Typography can be a tool which enables us to include rather than exclude, and to give those without a voice, an opportunity to have one.

4.3. History and Contemporary Practice

During the process of these projects, it was evident that precolonial and non-mainstream design histories are often unacknowledged and ignored in current design and historical research and practice. Yet cultural typographic histories can contribute and inform contemporary design practice. Western typography and book design have evolved without consideration for non-Western languages, typography, or design practices, so the challenge for MCLI was to incorporate Indic typographic traditions, design sensibilities and reader experiences into these bilingual editions, especially as the books are meant to be both for Western and Indian readership.

5. Research Methods

The projects employed several different research and design methods. Primary and secondary archival research was undertaken to focus on object research and establish an evidence-based understanding of practice and the sociocultural contexts in which book design and typographic design decisions were made. This included correspondence as well as original artwork. Extensive research was conducted on manuscripts, early printed books primarily in private and public collections in India and the UK; specialist archives including Bibilotheque Kandinsky (Centre Pompidou), St. Bride's Library, Roja Muthiah Library and SOAS (School of Oriental and African Studies) Library as well as Cooper Hewitt, Smithson Design Museum collections.

One of the design challenges of a 100-year publishing project was that it was essential that the system or standards that were being created survived the designers and researchers and provided whoever took this on in the future with a strong foundation to build on. The design act was to design texts whilst also simultaneously designing a system that would perpetuate. As noted by Farriss (1986), the key was to combine system (research) with process (design). Systems fit parts together in a synchronic

relationship explained by function; whilst process links them sequentially through cause and effect. The relationship of design is seen in motion, continually changing while remaining somewhat integrated.

With *Punctuation* and *Liberté,* research played a key role in building a sense of context of the time that the poems were written. A sense of authenticity within the design was embedded through material, narrative, and production. Visual research was conducted over three years using four sources. First, examples of "type in play" and "type as image" from a range of sources with the aim of analyzing the use of typography in these contexts. This research was limited to Morgenstern's lifetime with *Punctuation*. Second, investigations into the industrial production of typography and language (much of the context of Morgenstern's poem). Particular attention was paid to the way letterpress and typography as a medium could be used in communication of social and political themes. Third, photographic documentation of war in Germany, i.e., the visual imagery that stays in one's mind or in the popular imagination, even if one is unfamiliar with the firsthand experience of the war. Fourth, typographic testing and the investigation of use of red as a color in a variety of relevant contexts to draw attention for different reasons. With *Liberté,* the project was more challenging due to the relationship between the movement and typography. As noted by Poynor (2007):

> …one reason for Surrealism's relatively unexamined role in the history of graphic design is that it had no decisive impact on typographic methods and aesthetics. While graphic designers are still working today with typographic conventions that can be traced back to Modernism, Surrealism is not part of this narrative. (p. 50)

The visual research therefore had to consider the relationship between Surrealist concepts and the practice of typography and graphic design. Surrealists put the emphasis on automatic writing; it seemed pertinent to then consider an approach of automatic designing. To paraphrase Brian Schorn (as cited in Poynor, 2007), the process was fuelled by a desire to reach content not available through conventional typographic meanings, with each page an "individual universe without rules of logic."

With both projects, artifact analysis played a key role in establishing a relevant design approach. For MCLI, this focused on manuscripts and early printed books in Indian languages to provide both breadth and scope of knowledge and practice in pre-1800 Indian text design. This consisted of looking at objects whilst interrogating the contexts in which they were produced. The areas of research which fed into the practice were history of the book and printing in India; language, and scripts of India; reading and reader interactions with texts; and bilingual translations employing multiscript typography. For the *Punctuation* and *Liberté,* archival research and artifact analysis was conducted to investigate examples of typography in relation to poetry, particularly,

concrete, sound poetry, and nonsense verse. The research established visual examples of associative typography, wherein typography is concerned with the meaning and interpretation of the text and representing it using visual, verbal, and spatial aspects of typography.

Research through design practice methods were employed, using systematic analysis, typographic classification, iterative design, parallel prototyping and evaluation by expert editors and readers, often 4–5 per book. The design process functioned as a reflective research activity to enhance design practice through the examination of the tools and processes of design making, the critical act of recording and communicating steps, experiments, iterations of the design, and documentation to contextualize and communicate design actions through presentations.

The first consideration for the book design concept were the different languages and genres that the design had to accommodate. The MCLI task was to accommodate at the least the starting mission of 13 different languages and relevant scripts. The concept of "unity in diversity" is promoted strongly in India and is exemplified in the National Anthem written by Tagore. This became a guiding spirit for the interior design, i.e., to exemplify the best of the scripts and at same time, being relevant to the needs of the larger series.

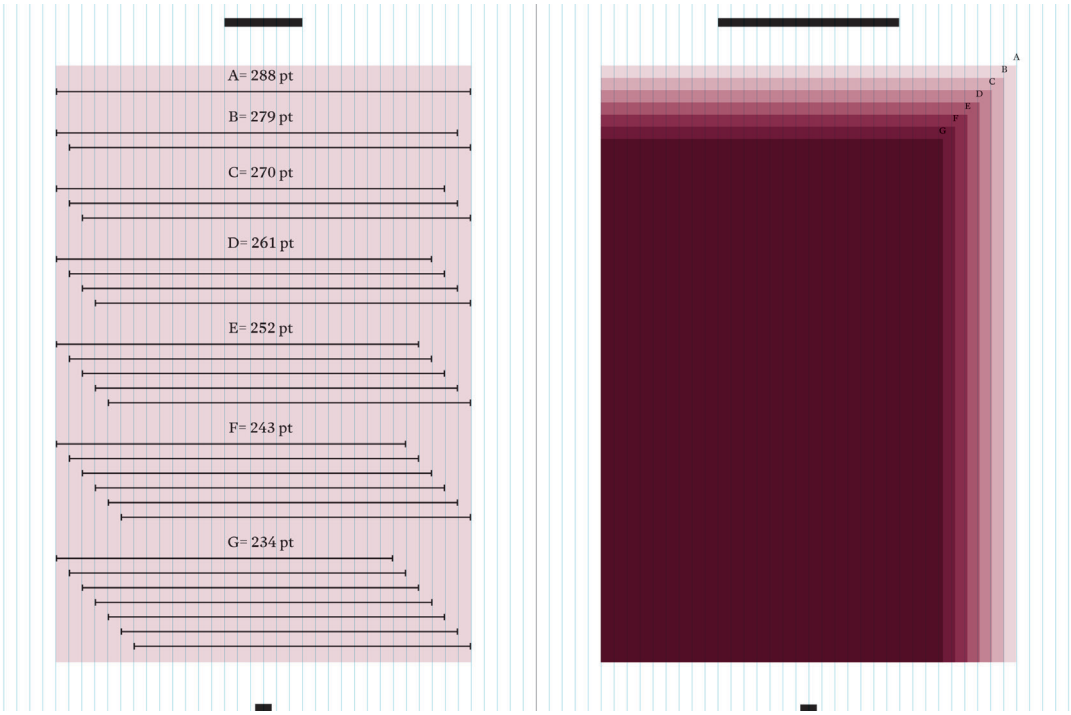

Figure 4. Diagram explaining the different lengths of text boxes in the polyphonic text framework (MCLI).

Figure 5. Diagram illustrating different permutations and combinations of text boxes in the polyphonic text framework.

It was important to acknowledge that the history of the book tradition in India is not the codex. It is the scroll or the manuscript. Textual content is shaped in part by the form (tools, materials and technology that produce form). With the introduction of printing in India tangled with colonial ambitions, this was something that also needed to be unraveled. Whilst conducting the research the aim was to pull out implicit understandings of how texts should be set as Farris (1986) noted. If it felt like something new was being built, this was not the intention. Instead, the research was reforming what existed for today's reader in a multilingual and intercultural context.

In India, reading is a public and social activity as well as a private activity. In India where there was and still is a sophisticated oral culture, there is a belief that oral communi-

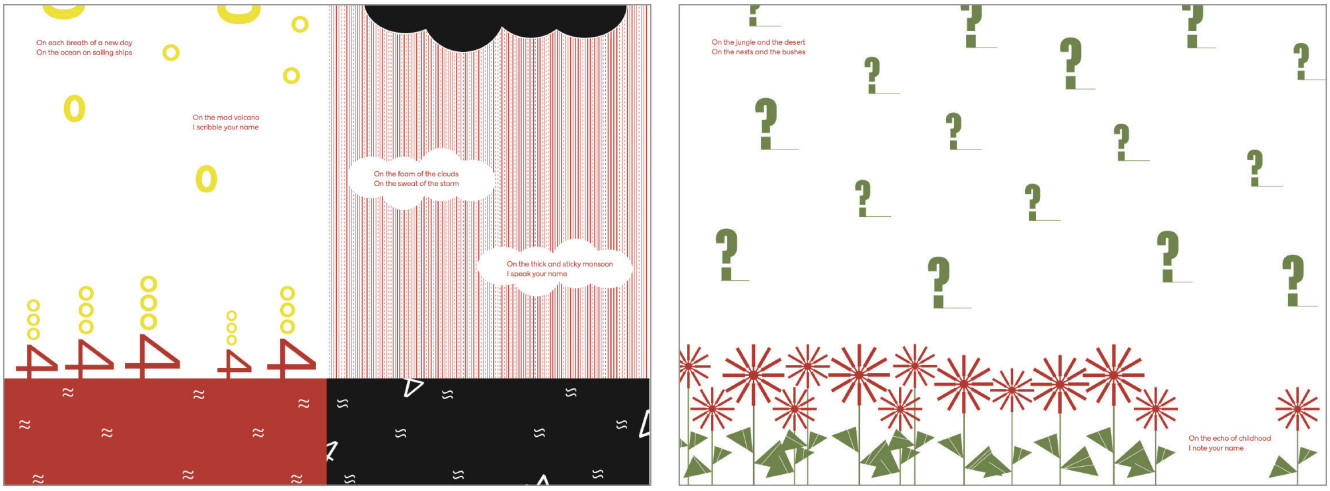

Figure 6. Two pages from *Liberté* (2024) which illustrate how meaning is created using typographic texture, movement, location and shape.

cation is still seen as an indication of one's ability as well as one's sincerity; it is also an affirmation of the belief that while what is written can always be read, what is meant to be heard must be spoken and lived.

In music, polyphony is a texture consisting of two or more independent melodic voices, as opposed to music with just one voice (monophony) or music with one dominant melodic voice accompanied by chords (homophony). Taking this as inspiration for facing-page translations in two different languages, the book design adjusts according to the languages and their relationship to each other (in terms of length of language). This system highlights the nature of each text and puts the languages directly in relation to each other on the spread that gives them equal emphasis. Following this concept, a grid was developed as a skeleton of the book, which allows different positioning of the elements on the page according to the length of languages (Figure 4). The template offers a systematic and flexible approach to the design of these classical texts in multiple languages. The width of the text box on the page adjusts according to the language in use and to the type of text (poetry, prose, etc.). The relationship between Indic and English text on the page results in a unique layout for each language/genre. The system aims to highlight the nature of the texts and put emphasis on hypertextuality.*

The page is formatted into a grid which divides the width of the page. While the top and base margins, placement of folios and running heads are set across the series, the side (inner and outer) margins of the template are flexible. The inner and outer margins allow the text block to contract and expand in relation to the language on the reflecting

* Hypertextuality refers to the way in which texts are interconnected through links, allowing readers to navigate from one piece of content to another in a nonlinear manner.

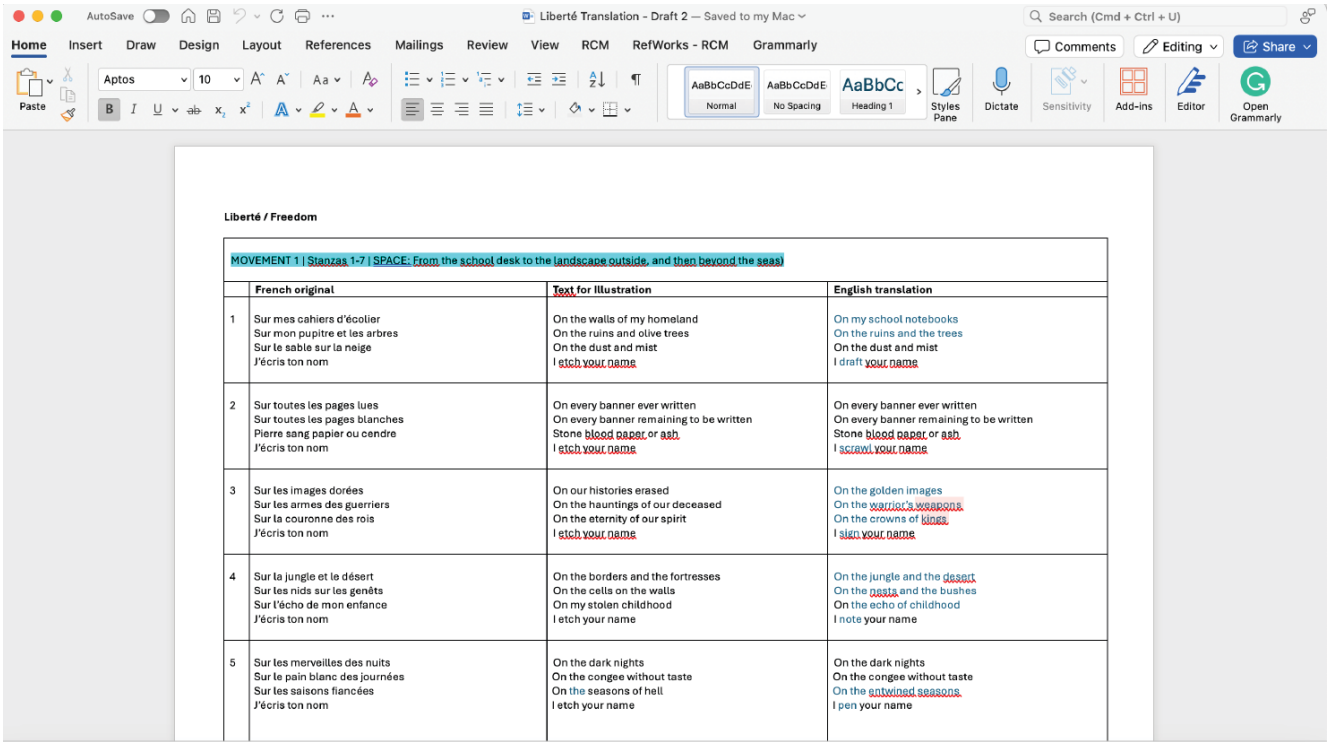

Liberté / Freedom

MOVEMENT 1 | Stanzas 1-7 | SPACE: From the school desk to the landscape outside, and then beyond the seas)

	French original	Text for Illustration	English translation
1	Sur mes cahiers d'écolier Sur mon pupitre et les arbres Sur le sable sur la neige J'écris ton nom	On the walls of my homeland On the ruins and olive trees On the dust and mist I etch your name	On my school notebooks On the ruins and the trees On the dust and mist I draft your name
2	Sur toutes les pages lues Sur toutes les pages blanches Pierre sang papier ou cendre J'écris ton nom	On every banner ever written On every banner remaining to be written Stone blood paper or ash I etch your name	On every banner ever written On every banner remaining to be written Stone blood paper or ash I scrawl your name
3	Sur les images dorées Sur les armes des guerriers Sur la couronne des rois J'écris ton nom	On our histories erased On the hauntings of our deceased On the eternity of our spirit I etch your name	On the golden images On the warrior's weapons On the crowns of kings I sign your name
4	Sur la jungle et le désert Sur les nids sur les genêts Sur l'écho de mon enfance J'écris ton nom	On the borders and the fortresses On the cells on the walls On my stolen childhood I etch your name	On the jungle and the desert On the nests and the bushes On the echo of childhood I note your name
5	Sur les merveilles des nuits Sur le pain blanc des journées Sur les saisons fiancées J'écris ton nom	On the dark nights On the congee without taste On the seasons of hell I etch your name	On the dark nights On the congee without taste On the entwined seasons I pen your name

Figure 7. Excerpt from the decolonial visual translation provided by Adeena May which assisted with the contemporary typographic interpretation.

page. The aim of the spread is to let each language reflect the other rather than letting one design decide the others. This means that the text boxes on the verso and the recto need not be of the same width which allows for the text to be placed on the page in 28 different ways (Figure 5). As the two languages do not have to be the same width, variations are possible. For *Punctuation* and *Liberté*, associative examples of typography were classified into different representational categories, forming a type palette and toolbox from which design drafts could be formed. They included texture, movement, location, shape, sound, and color (Figure 6).

A co-research and co-design process was undertaken with MCLI and *Punctuation* as well as *Liberté*. With MCLI, the process of establishing design frameworks involved iterations based on feedback from editors and translators working with Indic languages, as well as printers and binders. The book design and typography were iteratively designed with type designers, with the book design responding to the type design, and vice versa. The design was reviewed by language experts such as Rupert Snell (2018) as one that befitted the origins of the text as well a modern contemporary reading. With the *Punctuation*, readers tested early design drafts. Based on their understanding of the pages of typographic play, words and shapes were adjusted accordingly. This iterative process underpinned the aims of the book, i.e., to enable typography and language to expand

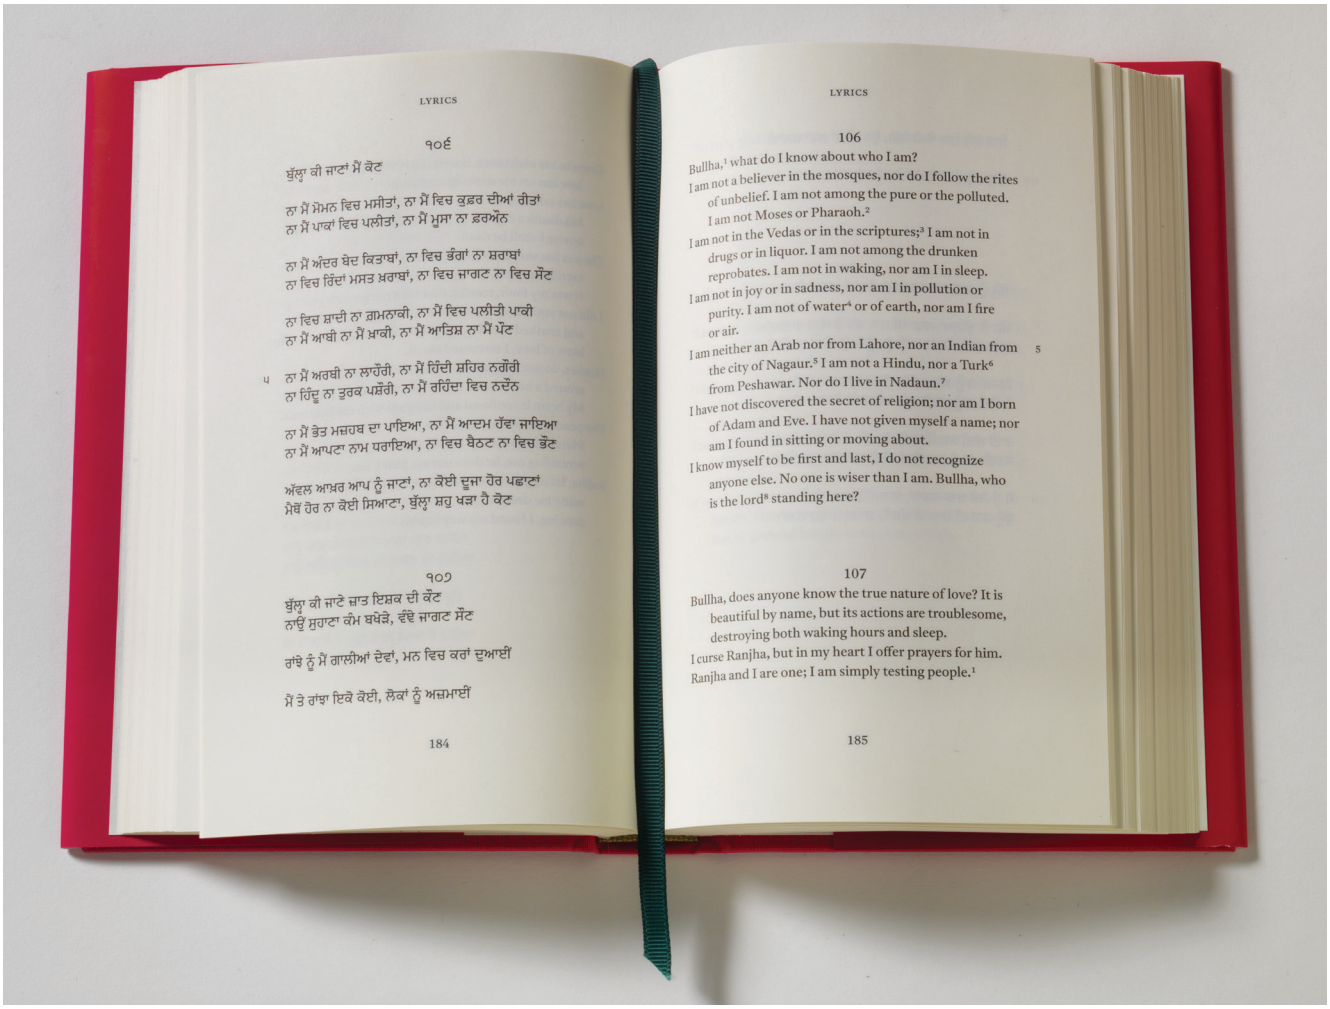

Figure 8. Open page spread illustrating how the book can be read. This title *Sufi Lyrics* has Gurmukhi script (Panjabi) on the left and Latin script (English) on the right.

aspects of reading (to incorporate sound, shape, texture, movement, color). With *Liberté*, the testing process had to work with two publishers — one who had grown up with the poem, and had previous context (a well-known poem taught in French schools), and the other without the national context for whom the work had to translate. Translation was an important aspect of the latter, with *Afterall* editor Adeena May providing a decolonial visual translation (Figure 7) for the typographic narrative.

In addition, contextual design methods were used to establish material-based narratives for MCLI. With MCLI, the books are produced in a hardback or library edition and a less expensive paperback, for the mass market in India. The hardback is bound to lie flat so the reader can make notes in the side margins and cross-reference the bilingual texts with ease (Figure 8).

6. Findings, Insights and Conclusions

The project aimed to establish the relevance of an approach not just to "non-Latin" typography but more broadly to the practice of typography, in relation to language. The aim here was to make more visible, through design and documentation, a broader approach to typography which acknowledges typography's link to language, as it is spoken, written, and read both culturally as well as materially. As noted by Pollock (2011), such approaches provide many occasions for learning something about our "shared humanity" from these works, but they also "give access to radically different forms of human consciousness for any given generation of readers, and thereby expands for them the range of possibilities of what it means to be a human being" (p. 36).

The history of the book which looks primarily at the codex, needs to encompass the histories that are beyond the codex, to manuscripts, scrolls and other "book" traditions which are rarely documented or acknowledged. Research revealed that there are no existing bilingual design frameworks for the presentation of Indian texts in Indian scripts and languages, nor as translations into English. In a letter written by Tim Jones, Director of Design and Production at Harvard University in 2018, he noted, "we had never applied the facing-page translation concept to such a wide array of languages and scripts." Jones goes on to say that the research "addressed the critical need for a unified design approach that could encompass a wide array of variation and many disparate requirements."

As discussed previously, existing approaches with Indian texts come from colonial roots of printing which have aimed to synthesize "non-Latin" scripts with a Latin page, rather than from the requirements of the scripts, languages, or texts themselves. This is a primary framework for research and design for multiple languages that can be applied to other world languages. The broader aim is to show the relevance of this approach not just to "India," not just to "non-Latin," but more broadly to the practice of design and typography and the relevance of research. Such a decolonial and intercultural typography acknowledges all periods of textual history, not just the dominant and the easily accessible.

Typographic guidelines for Indian texts that respond to Indic hierarchy and grammar in the application of typographic rules can enable contemporary reading and accommodate multiple (and new) readers. Indian typography borrows conventions from Western models of typography, converting typographic styles such as "bold," "underline," "italics," and "slanted" to contexts which do not use such styles. The research addressed the challenges of emphasis and hierarchy in texts by providing solutions more relevant to the roots of Indian scripts, for example, by employing color, size, and location (Figure 9). In relation to the layout, as noted, the design framework was based on the

MCLI STYLE GUIDE FOR TYPESETTING: HINDI TEXT

The following guide refers to treatment of text set in Hindi.
All text is set in MCLI Antwerp and Murty Hindi only.
Paragraph style language setting to be set in Hindi
** Refers to alternative treatments of the text adopted for the design of specific titles.*

PARAGRAPH STYLE CONVENTIONS

H4: HI Subsection Num 11.5pt
STYLE: Regular; SIZE: 11.5pt; LEADING: 17.5pt;
KERNING: Metrics; TRACKING: -25;
CASE: Normal; ALIGNMENT: Centre;
TO USE FOR: Main Text Poem Numbers;

H5: HI Side Number 9pt
STYLE: Regular; SIZE: 9pt; LEADING: 17.5pt;
KERNING: Metrics; TRACKING: -10;
CASE: Normal; ALIGNMENT: Left

H5: HI Side Number 9pt 80% K
STYLE: Regular; SIZE: 9pt; LEADING: 17.5pt;
KERNING: Metrics; TRACKING: -10;
CASE: Normal; ALIGNMENT: Left
TO USE FOR: Alternative Numbering System;
COLOUR: Normal

H5: HI Poetry Text 11.5pt
STYLE: Regular; SIZE: 11.5pt; LEADING: 17.5pt;
KERNING: Metrics; TRACKING: 0;
1ST LINE INDENT: -18pt; LEFT INDENT: 18pt

H5: HI Poetry Text 11.5pt 1st Line of Pg
STYLE: Regular; SIZE: 11.5pt; LEADING: 17.5pt;
KERNING: Metrics; TRACKING: 0
ALIGN TO GRID: First line only

H5: HI Poetry Text 11.5pt 1st Line
STYLE: Regular; SIZE: 11.5pt; LEADING: 17.5pt;
KERNING: Metrics; TRACKING: 0
1ST LINE INDENT: -18pt; LEFT INDENT: 75pt

H6: HI Notes Text 8pt
STYLE: Regular; SIZE: 8pt; LEADING: 10.5pt;
KERNING: Metrics; TRACKING: 0;
ALIGNMENT: Left Justify;
1ST LINE INDENT: -18pt; LEFT INDENT: 18pt

H7: HI Notes section Number 9pt
STYLE: Regular; SIZE: 9pt; LEADING: 10.5pt;
KERNING: Metrics; TRACKING: 0;
ALIGNMENT: Centre

H8: HI Notes Text 18pt Indent
STYLE: Regular; SIZE: 8pt; LEADING: 10.5pt;
KERNING: Metrics; TRACKING: 0;
ALIGNMENT: Left justify;
1ST LINE INDENT: -18pt; LEFT INDENT: 18pt;

CHARACTER STYLES CONVENTIONS

HI Note Ref H5
STYLE: Italic; SIZE: 11.5pt; LEADING: 17.5pt;
TRACKING: 0; CASE: Normal;
POSITION: Opentype Superior/Superscript

HI EN Note Text 9pt H8
STYLE: Regular; SIZE: 8pt; LEADING: 10.5pt;
TRACKING: 0; CASE: Normal; POSITION: Normal

HI EN Note Side Num 7pt
STYLE: Regular; SIZE: 7pt; LEADING: 10.5pt;
TRACKING: 0; CASE: Normal; POSITION: Normal

Figure 9. Typesetting guide for setting Hindi-English in the MCLI series.

relationship of scripts and languages to each other, on a facing page to enable reading for both scholarly means as well as for pleasure, and with readers of different fluencies.

The approach was commended by Walker (2017, p. 8) in a paper "Research in Graphic Design" as an example of good practice: "Rathna Ramanathan and Fiona Ross's work on book and typeface design for the Murty library is an excellent example drawing together cultural and historical precedent to inform contemporary graphic design." Schulze and Arnall (2011) proposed that we can design the means through which design happens, challenging the concepts, behaviors, and means of production as well as designing form. The project is not just about the spirit of the design process but about the impact of the project through design on everyday situations. The typefaces used in the project are available open-source to anyone working in the Indian context. The books are being brought back into universities, are available at an affordable price to the Indian public as well as accessible to an international audience. As noted by Pollock (2011), we need ways of describing the world that do not just belong to one tradition. MCLI, its purpose, design and production in all aspects are just one small step towards that.

With *Punctuation,* involving children and readers into the design process contributed to the final design of the text, and the approach changed the practice of the publishers themselves. In an interview conducted in 2021, Wolf noted in relation to the research,

> This has left a legacy that can be seen in terms of the strength brought to typography and design as a voice, to the process of the book understood as an ongoing conversation in which typography also has a voice.

With MCLI, the impact of this research has been twofold. Firstly, it has enabled the preservation of and access to Indic classical texts and Indic scripts by providing typographic frameworks and design guidelines for publication of bilingual books in Indic and Latin scripts by the Murty Classical Library of India. For Tara Books, the research has developed an approach to typography that empowers marginalized communities of readers as well as expanding readership in inclusive and decolonial ways. The success of *Punctuation* inspired the commission of another project — bringing that dialogue and process to a wider audience.

For us to address global challenges such as climate, health, or fake news, we need to acknowledge that communication is a fundamental right that needs to encompass culture and recognize context. It must understand that readers have multiple perspectives, reading fluencies, and bring their own contexts to the page or the screen. To build pluralistic and intercultural frameworks for typographic practice, we do this through the depth and interrogation of research not as an elite activity but as an everyday practice. This requires, primarily, a genuine need to know and understand that which

is not known or understood, rather than to solely pursue something that is "new" or "original" for practice or research in design.

7. References

Appadurai, A. (2006). The right to research. *Globalisation, Societies and Education, 4*(2): 167–177. https://doi.org/10.1080/14767720600750696

Ansari, A. (2021). Decolonisation, the history of design, and the designs of history. Paper presented at the Annual Design History Conference, online, September 1.

Beirut, M. (2019). [Blurb]. In D. Reinfurt, A. Micheals, and E. Lupton. *A *new* program for graphic design*. Inventory Press.

Bhatt, S., & November Studio (2020). Pluralism and power dynamics in Indian design: November Studio. *Walker Art Center.* https://walkerart.org/magazine/pluralism-indian-design-november-shiva-nallaperumal-juhi-vishnani/

Calvert, S. (2012). Materia Prima, text-as-image. *Journal of Writing in Creative Practice, 4*(3), 309–328.

Deshmukh, C. D. (1958). *The printing press in India*. Marathi Samshodhana Mandala.

De Sousa Santos, B. (2018). *The end of the cognitive empire: The coming of age of epistemologies of the south*. Duke University Press.

Eluard, P. (1942). *Liberté*. In *Poésie et verité.* Les Éditions de la main à plume.

Farriss, N. (1986). Foreword. In A. Appadurai (Ed.). *The social life of things: Commodities in cultural perspective*. Cambridge University Press.

Fry, T. (2007). Book review: The Archework Papers. *Design Issues, 23*(3): 88–92. http://doi.org/10.1162/desi.2007.23.3.88

Gruendler, S. (2005). *The life and work of Beatrice Warde* [PhD dissertation]. University of Reading.

Jones, T. (2018). *Letter to the Royal College of Art*. Harvard University.

Lees-Maffei, G., & Lees-Maffei, N. P. (2019). *Reading graphic design in a cultural context*. Bloomsbury Visual Arts.

Macaulay, T. B. (1835). Minute on education. Bureau of Education. In H. Sharp (Ed.), *Selections from educational records, part I (1781—1839). Superintendent, government printing, 1920. Reprint.* National Archives of India, 1965, 107–117.

Nallaperumal, S., & Vishani, J. (2020). In S. Bhatt (Ed.), *Pluralism and power dynamics in Indian design: November studio*. https://walkerart.org/magazine/pluralism-indian-design-november-shiva-nallaperumal-juhi-vishnani/

Pinney, C. (2013) More than local, less than global: Anthropology in the contemporary world. In C. Shore & S. Trnka (Eds.), *Up close and personal: On peripheral perspectives and the production of anthropological knowledge* (pp. 160–175). Berghahn Books.

Pollock, S. (2011). Crisis in the classics. *Social Research, 78*(1), 21–48.

Poynor, R. (2007). Dark tools of desire. *Eye, 63*(16), 54–61. https://www.eyemagazine.com/feature/article/dark-tools-of-desire

Ramakrishnan, C. (2010). Tamil typography. Email correspondence.

Ramanathan, R. (2023). Politics of publishing: Exploring decolonial and intercultural frameworks for marginalized publics. In J. Yee & P. A. Rodgers (Eds), *The Routledge companion to design research* (pp. 75–90). Routledge Art History and Visual Studies Companions. Routledge.

Schulze, J., & Arnall, T. (2011). Change through making. *Eye*, Summer. https://www.eyemagazine.com/feature/article/change-through-making

Shyam, B, & Wolf, G. (2017). *The London Jungle Book*. Tara Books.

Snell, R. (2018). Reviews: The Epic of Ram. *Bulletin of the School of Oriental and African Studies*, *80*(1), 165–167.

Till, J. (2021). Research after research. Paper presented at IoA Sliver Lecture Series 21/22: Research Cultures, online, December 2.

Trivedi, H. (2008). The 'book' in India: Orality, manuscript, print (post) colonialism. In R. Fraser & M. Hammond (Eds.), *Books without borders*, Volume 2 (pp. 12–33). Palgrave Macmillan.

Wolf, G. (2021). Unpublished interview [for REF impact case study at RCA].

Author

Rathna Ramanathan is a typographer, practice design researcher and academic known for her expertise in intercultural communication and alternative publishing practices. She is Provost, Central Saint Martins, Executive Dean for Global Affairs and Professor in Design and Intercultural Communication at the University of the Arts London. For the past thirty years, Rathna has led research-driven, intercultural, multi-platform graphic communication practice design research projects primarily in the Global South, all fuelled by a love for, and lifelong interest in typography and languages, and a belief in communication as a fundamental human right.

RESEARCH ARTICLE

The Role of Text Alignment on Response Speed and Accuracy When Reading Chinese-English Bilingual Traffic Signs

Yuchan Zhang, Jeanne-Louise Moys, and Matthew Lickiss

School of Design, University of Leeds, UK

Corresponding author: Yuchan Zhang (Y.Zhang15[at]leeds.ac.uk)

Abstract: This paper discusses the effect of text alignment on Chinese-English bilingual traffic signs on the speed and accuracy of navigational responses. Two text alignment conditions (centered and left settings) were tested in relation to sign complexity and the separating spacing between place names. Video materials were used to provide a safe way to simulate how and where road signs may appear in a driving scenario, while efficiently testing many variations. A total of 36 participants who read English but not Chinese engaged in this study. The results suggest that left alignment improves speed and accuracy in making decisions when responding to three-directional signs with narrower separating spacing, whereas centered alignment may be beneficial for one-and two-directional signs. These findings highlight the value of including text alignment specification in guidance for bilingual signage, especially when accounting for sign complexity. However, further studies are needed using methods with a higher ecological validity and a broader range of participants before robust recommendations can be devised.

Implications for practice: This study underscores the need for more nuanced typographic guidance in the design of Chinese-English bilingual road signs. The findings highlight the importance of accounting for the complex interactions between typographic and spatial attributes in sign composition for drivers to make quick and accurate decisions. In particular, practitioners and policymakers should specify text alignment, as it impacts drivers' speed and accuracy in navigation. These insights contribute to the development of safer and more efficient road navigation systems.

@: ISSUE〉 ARTICLE〉

Cite this article:

Zhang, Y., Moys, J.-L., & Lickiss, M. (2025). The role of text alignment on response speed and accuracy when reading Chinese-English bilingual traffic signs. *Visible Language, 59*(2), 130–152.

First published online August 15, 2025.

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

Visible Language Consortium:
University of Leeds (UK)
University of Cincinnati (USA)
North Carolina State University (USA)

Keywords: bilingual typography; Chinese-English road signs; sign legibility; text alignment; text spacing

1. Introduction

Traffic signage, though experienced as an everyday part of life, presents a number of design and layout challenges compared to more traditional texts. For example, traffic signs are typically highly restricted in their use of space and layout, contain comparatively short text fragments (potentially a single word, but with high degrees of variation), that may appear in combination with schematic elements (arrows and lines) and symbols. The positioning of elements within a sign can also convey navigational cues. Drivers typically have to make very rapid decisions in response to this information while processing other significant safety considerations. In this respect, conventions for traffic sign presentation are highly important and often established in legal guidance.

The use of guidelines ensures that individual signs function coherently within a signage system. Such a system builds a set of conventions that allow sign creators to consistently develop signage for a range of scenarios. Examples include the US *Manual on Uniform Traffic Control Devices (MUTCD)* (Federal Highway Administration, 2009) and *Standard Highway Signs (SHS)* (Federal Highway Administration, 2012) and the UK's *Traffic Signs Manual* (Department for Transport, 2016). When applying the ideas of guidelines to real world road infrastructure, however, huge variations between both junction topology and the naming of relevant roads and destinations present problems not only in designing consistent road signage, but in codifying the signage system to ensure consistency.

The majority of research on traffic signs tends to focus on signs in a single language and script and on factors like typeface choice and size (Beier, 2016; Dobres et al., 2017; Gálvez et al., 2016; Waller, 2007). While existing Chinese-English sign guidelines, such as the *Technical Guidelines for the Replacement of National Expressway Network Related Traffic Signs* (Research Institute of Highway Ministry of Transport & Beijing Communications Highway Survey and Design Institute, 2007), *Technical Guidelines for the Adjustment of National Highway Network Traffic Signs* (Research Institute of Highway Ministry of Transport et al., 2017), and *GB 5768.2-2022: Road Traffic Signs and Markings* (State Administration for Market Regulation and Standardization Administration of China, 2022) are based on research, such guidelines may not always lead to optimal outcomes in practice for bilingual signs (Figure 1), as the standards are typically based on monolingual design research. When the guidance is provided for bilingual design,

Figure 1. A Chinese-English bilingual traffic signage. Photographed in Beijing, China, 2024 by the first author.

it often overlooks the complex typographic nuances of different scripts (see detailed discussion in Section 1.1). In this context, there is scope for additional research that considers specifications for bilingual signs.

Studies have confirmed that drivers require more reading time for bilingual signs because of the doubled information provided (Rutley, 1972). To minimize reading time, the findings and solutions of studies focus on distinguishing the two languages to help users quickly locate the needed information (Rutley, 1974). The solution for distinguishing between two languages is more effective when the signs contains two similar scripts (e.g., English and Welsh). However, this approach is less useful for signs that feature distinctly different writing systems, like English and Arabic, where the scripts are inherently easy to differentiate. Therefore, bilingual sign legibility research requires additional considerations to address challenges unique to such contexts. Eid (2009) and Petretta (2014) suggest considering text spacing, script alignment, information sequence, and the role of pictorial elements when designing Arabic-English signs. Nevertheless, there is limited research on sign legibility that addresses the design challenges of combining the logographic Chinese script with the alphabetic Latin script.

In both Chinese and English contexts, there is research to support that text spacing can be arranged to assist in sign legibility (Chan et al., 2014; Garvey & Kuhn, 2004; Hsu & Huang, 2000; Tejero et al., 2018). More recently, the authors' previous studies have

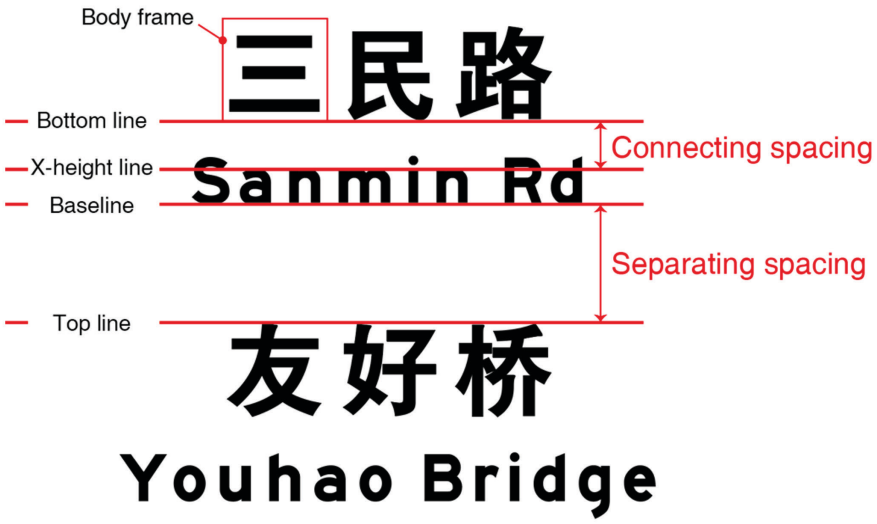

Figure 2. Connecting spacing between Chinese and English text and separating spacing of two bilingual place names. Connecting spacing is "line spacing" between the two languages and separating spacing is "line spacing" that separates two different place names.

demonstrated that the interline spacing* has impact on the legibility of Chinese-English bilingual traffic signs (CEBTS) (Zhang, 2021; Zhang & Moys, 2022). Specifically, the kinds of interline spacing analyzed include connecting spacing — the vertical spacing connecting a Chinese location name to its English translation (Zhang & Moys, 2022) — and separating spacing — the vertical distance allocated to distinguish between two bilingual names within a single direction (Zhang, 2021) (Figure 2).

However, there seems to be comparatively less research on the role of alignment of two scripts in CEBTS. There are four basic alignments: left-aligned, right-aligned, justified, and centered. Left-aligned text maintains uniform word spacing, aligning all lines to a common vertical point on the left while leaving the right edge uneven. Right-aligned text, on the other hand, aligns all lines to the right, resulting in an uneven left edge. Justified alignment adjusts the word and letter spacing so that all lines are of equal length, creating a uniform, block-like appearance. Centered alignment positions each

* The term *interline spacing* in Latin typography is similar to *line spacing,* but the definition of line spacing differs in the Chinese context. In Chinese typography, since Chinese characters do not contain descenders, line spacing refers to the vertical distance between either the bottom or top lines of the square characters to the corresponding bottom or top lines of the next row, rather than being based on the baseline as in Latin typography. To address the challenges of accurately defining spacing in bilingual text combining Chinese and English, the term "interline spacing" was adopted to ensure clarity and precision. This distinction is particularly important to avoid ambiguity when discussing spacing conventions across different typographic systems and will be used in the following sections to ensure consistent terminology.

line symmetrically, with equal spacing on both sides, resulting in a balanced but ragged appearance on both the left and right edges.

Studies of typography for print and screen indicate that text alignment can play an important role. For example, Ling and Van Schaik (2007) demonstrate that text should be left-aligned rather than justified when information is presented that needs to be scanned quickly, particularly on web browsing. Similarly, Hartley et al. (1975) advocate for the use of left-aligned texts in tables as centered alignment took longer to create and produced more errors. Although the contexts of these studies are different, both consider information that people are more likely to scan and take in at a glance.

For continuous reading, there are early studies that support the importance of left-aligned and uniform spacing (Gregory & Poulton, 1970; Hartley & Mills, 1973). However, earlier work examining the influence of text alignment on use of printed material found no advantage of left alignment (Fabrizio et al., 1967). Nevertheless, professional standards within information design tend to advocate for text that is left aligned rather than justified or centered for continuous reading (Luna, 2018). This premise is often attributed to an understanding that the text is easier to read because the spacing between words is more even. However, Dyson (2018) suggests that since the research is neither substantial nor conclusive this may be an aesthetic convention. There is an aesthetic rationale for using left-aligned text to ensure uneven gaps between words in justified text do not form "rivers" on the typeset page.

While this aesthetic consideration is less applicable to sign design than continuous text, there has been relatively little research into the role of alignment in signage. Barker and Fraser (2004) recommend left- and centered-alignment for alphabetic signage. In terms of bilingual sign scenarios involving Arabic and English, Eid (2009) and Petretta (2014) suggest vertically staggering the two scripts to improve legibility for shorter messages. Guidance in China such as *Technical Guidelines for the Adjustment of National Highway Network Traffic Signs* (Research Institute of Highway Ministry of Transport et al., 2017) and *5768.2-2022: Road Traffic Signs and Markings* (State Administration for Market Regulation and Standardization Administration of China, 2022) recommend vertically staggering Chinese and English on signs, with Chinese placed above. They also suggest using left- or centered-alignment on traffic signs.

1.1. Real World Examples

To ensure that the materials developed for the testing reported in this paper would be reasonably representative of real signs, a reference sample of current CEBTS in Chinese urban areas was collected. The sample comprises 143 signs from four cities in China: Beijing ($^{25}/_{143}$), Shanghai ($^{23}/_{143}$), Wuxi ($^{16}/_{143}$), and Dalian ($^{79}/_{143}$), photographed between 2017 and 2019. Therefore, the experimental materials are based on the

samples observed during this timeframe and the standards published from the 1990s up to the date of sample collection. While there have been some minor changes in the latest issued standard in 2022, the changes did not involve text spacing or alignment, and therefore had no impact on the eligibility and validity of the material design. In addition, although the samples were collected and observed a few years ago, since the current signs remain unchanged as observed in 2024, the material remains representative of what is seen in urban areas in China at the time of writing.

The sample was examined based on: a) the method by which bilingual location names are positioned in relation to one another, and b) the method by which the English translation aligns with its Chinese counterpart within a bilingual place name. The examination of the sample indicates that the two recommended alignment methods (left and centered alignment) in the standard *GB 5768.2-2022* (2022) have not been appropriately implemented, revealing inconsistencies in their application (Figure 3 and Figure 4). The most widely used method is centered (66/143) and there seems to be a tendency towards justified alignment (24/143). As depicted in Figure 3, all Chinese place names are formatted to uniform line lengths by adjusting (in places quite substantially) the horizontal spacing between words in both scripts, and the widths of letters are altered to achieve this goal.

However, the guidance in *GB 5768.2-2022* (2022) recommending left alignment was found rarely used to align English transliteration to its corresponding Chinese location name and to be applied only when aligning multiple bilingual location names (1/143) (e.g., the two bilingual location names are left-aligned on the right-hand side of the sign

Figure 3. Uniform line lengths in bilingual location names achieved through adjusted text spacing. Photographed by the first author in Dalian, China in 2018. The sign remains unchanged as observed in 2024.

Figure 4. The Latin texts are central-aligned with their corresponding Chinese texts, but the two bilingual place names are left-aligned on the right-hand side of the sign. Dalian sample. Photographed by the first author in 2018. The sign remains unchanged as observed in 2024.

shown in Figure 4). Notably, none of the 143 collected samples employed left alignment to arrange the two scripts within a bilingual location name. In contrast, the Latin text tends to be either justified (Figure 3) or centered (Figure 4) to align with its Chinese counterpart.

All the above leads to the main research question of this paper: does text alignment in bilingual traffic signs with two different scripts (and languages) affect how quickly and accurately people might make direction decisions?

Although the analysis of real-world examples has shown a prevalence of justified alignment, its current application compromises legibility, particularly in English texts, as evidenced in Figure 3. Considering the nature of the two scripts — where Chinese characters encode more information per glyph than Latin letters (often resulting in shorter Chinese place names compared to their English counterparts), the disparity in message lengths is not successfully addressed by the justified method in practice. Consequently, justified alignment is excluded from this study. As centered alignment is still widely applied and recommended by standards, this study focuses on comparing left alignment with centered alignment and seeks to determine:

1. whether there is a difference between the two alignments of the bilingual location names, centered- or left-aligned, in the legibility of CEBTS; and
2. if differences are found, which one could improve participants' speed and accuracy when identifying bilingual place names.

2. Method

2.1. Study Setup

The experimental design used 3D video simulations to replicate a driver's view of approaching CEBTS at a constant speed. Participants were shown signs that feature bilingual location names and directional arrows, becoming larger on screen until partic-

ipants can discern and report their observations. By employing a threshold method (Dyson, 2018), the experiment measured the earliest point at which participants can identify the signs and can use the information they gained to answer a question phrased in the format: "Which direction leads to destination xxx?" As soon as participants felt they can discern the answer from the video they indicated their directional choice by using the arrows on a keyboard. Response time and accuracy were recorded to assess the sign legibility. To emphasize response speed, video clips were displayed for no longer than seven seconds, concluding immediately upon the participant's response. See Figure 5 for further visual details. This experimental approach has parity with the methods used in the studies of separating and connecting space (Zhang, 2021; Zhang & Moys, 2022), enabling clear comparisons to be made across the interpretation of results.

The video stimuli, rendered at a resolution of 1280×1024 pixels, were displayed on a 75-inch monitor with a Full HD resolution and a refresh rate of 60 Hz. The monitor was set to ensure that the video content was scaled appropriately. The presentation, timing, data collection, and storage were managed using E-Prime 2.0 software.

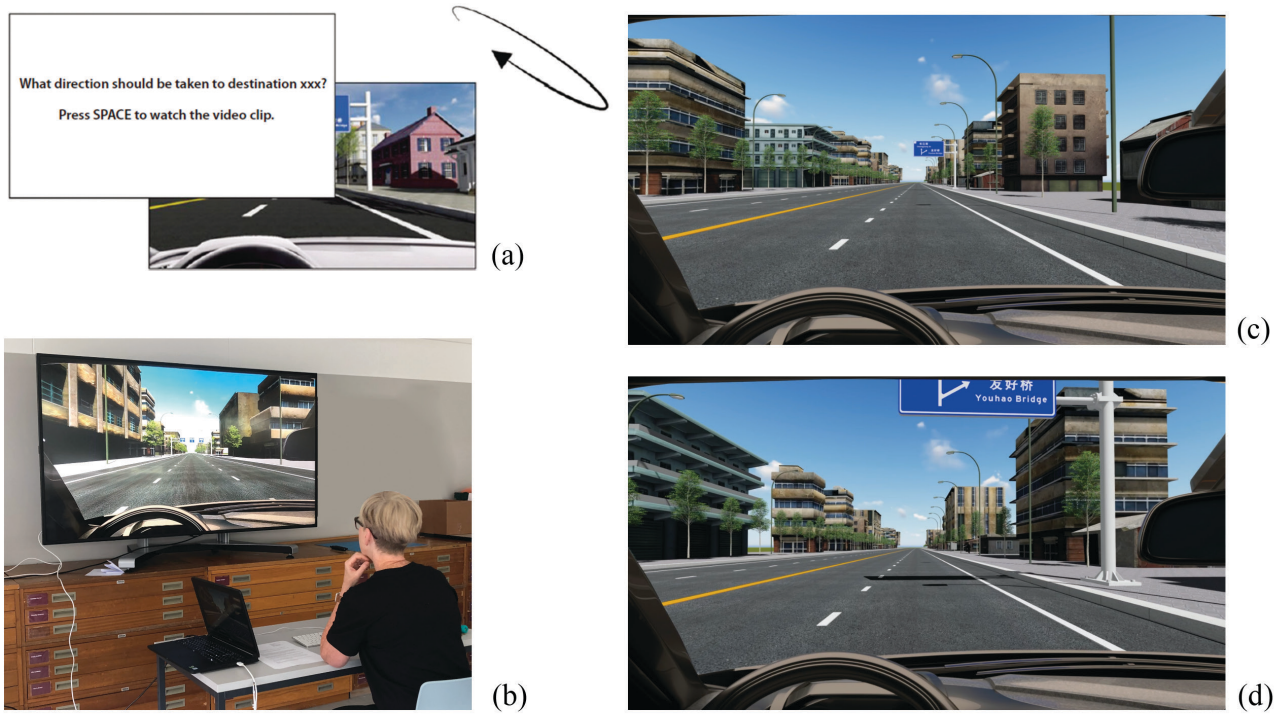

Figure 5. (a) The study procedure presented participants with a question to answer prior to each video stimulus. (b) A participant engaged in the test. (c) A screenshot showing the starting point of the video displayed to participants, where the road sign appears at its smallest. (d) A screenshot showing the ending point of the video, displayed at the seven-second mark.

The study was conducted in a university teaching room (Figure 5b). Each participant was tested individually in a quiet and controlled environment without external interruptions. Participants were instructed to turn off personal devices and other potential sources of distraction before beginning the session.

A pilot experiment was conducted to ensure the equipment and materials were set up appropriately and comfortably for the participants. This involved ensuring the table and chair heights were comfortable, positioning the monitor 1.6 meters away to provide clear visibility of the materials, and confirming that the modified computer keyboard with five directional arrows (Figure 6) was intuitive and easy for participants to understander and use. To maintain consistent conditions, the same room and equipment were used for both the pilot and main sessions, and the setup was aligned with the standards of Zhang (2021) and Zhang and Moys's (2022) previous studies.

2.2. Material Design

Parameter 1: Separating spacing. Previous research has indicated that reading speeds fluctuate across three levels of separating spacing (Figure 2): 0.5H, 0.75H, and 1H (where H represents the height of one Chinese character) (Zhang, 2021). Specifically, faster response times were observed at the 0.5H and 0.75H separation levels (with no significant difference between the two) compared to 1H (which shows a significant difference).

This raises an interesting question about how text alignment might further impact legibility at the 0.5H and 0.75H separation levels (Investigation 1) (Figure 7a,b). Additionally, for signs displaying only a single location name — where separating spacing is not applicable — it is important to investigate how alignment affects legibility independently of separating spacing (Investigation 2) (Figure 7c,d).

Parameter 2: Sign complexity. To ensure the study was reasonably representative of the varying degrees of sign complexity evident in the real-world sign sample (see section 1.1), the research investigated three levels of sign complexity: one-directional

Figure 6. A computer keyboard, adjusted to provide five directional arrows, enables participants to enter their responses.

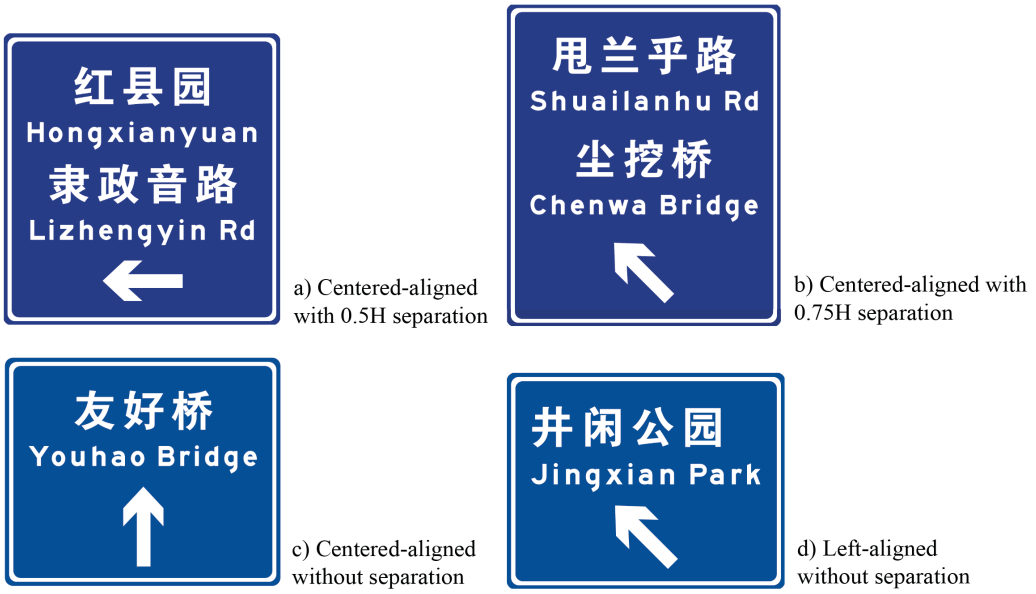

Figure 7. Four stimuli varying in text alignment and separating spacing used in the investigations.

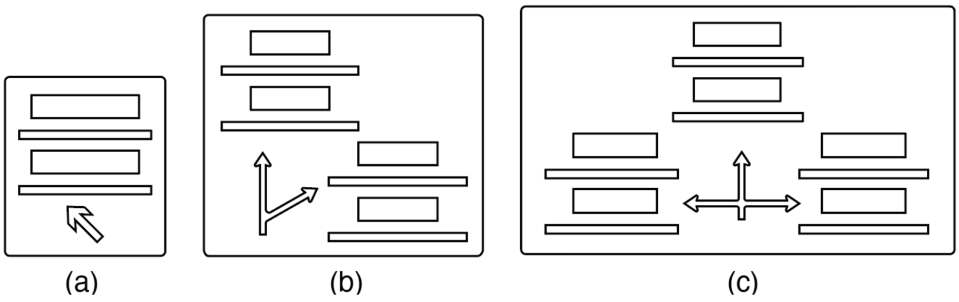

Figure 8. Three levels of sign complexity. (a) one-directional sign; (b) two-directional sign; (c) three-directional sign. Extracted from Zhang (2021).

(41 out of 143 samples), two-directional (25⁄143), and three-directional signs (54⁄143) (Figure 8).* The three levels of sign complexity adopted were the same as those used in Zhang (2021) and Zhang and Moys (2022) to enable continuity and comparison with our studies considering the impact of both connecting and separating spacings on the legibility of CEBTS.

* Sign complexity was categorized into three levels based on the number of directional indicators displayed. These three categories can encompass a broad spectrum of sign types used on CEBTS.

Parameter 3: Schematic elements. The schematic elements (including arrows, route arms, and orientation-direction patches) for the materials were designed to adhere to *Technical Guidelines for the Adjustment of National Highway Network Traffic Signs* (Research Institute of Highway Ministry of Transport et al., 2017) as it was the latest standard at the time of this study before the new 2022 standard was issued, ensuring the findings were relevant to the then-current regulation and practice (Figure 9). Although the new 2022 standard has been issued, it does not affect the validity of this study's materials, as the limited updates related to visual design are irrelevant to the schematic elements used in the two investigations. The schematic elements were controlled throughout the study. Although different forms of schematic elements were used in the stimuli (arrows in the one-directional signs and route arms in the two- and three-directional signs) and different sign layouts (asymmetric two-directional signs and symmetric

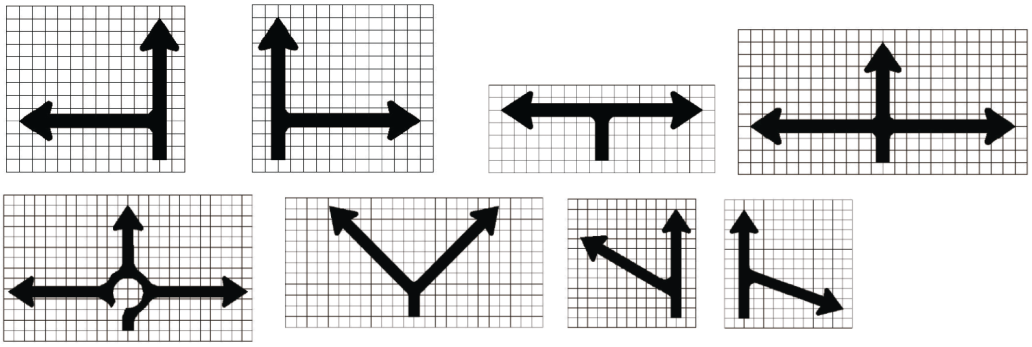

(a). Route arm guidance provided in the standard. Extracted from the Figure C.1.2, p.115. Route arms were applied to the two-and three-direction signs in this study.

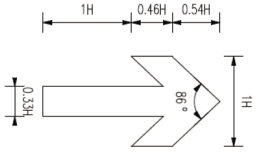

(b). Arrow guidance provided in the standard.Extracted from the Figure 3.5.3-3, p.16.Arrows were applied to the one-direction signs.

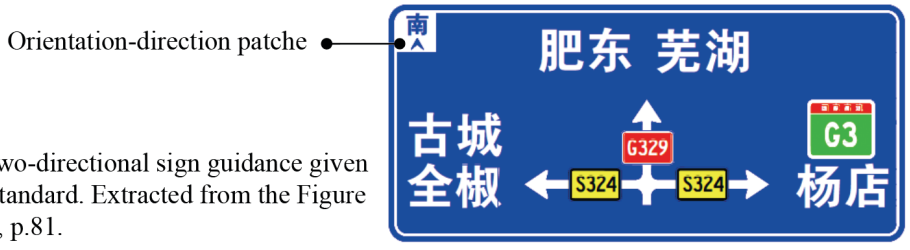

(c). A two-directional sign guidance given in the standard. Extracted from the Figure 7.3.4-1, p.81.

Figure 9. Some of the visual design guidelines extracted from *Technical Guidelines for the Adjustment of National Highway Network Traffic Signs* (2017), which the design of the stimuli in the Investigations adhered to.

three-directional signs) may affect participants' reading performance. All participants read all sign complexity levels and their performance was analyzed and compared according to each level, which enabled the isolation of the effect (see Section 2.3 for detailed testing methods). As with any legibility research involving multiple variables, there must be a balance between the isolation of all possible variables and real-world pragmatics in experimental design that accounts for the variation seen in the relevant materials — traffic signage being an area of such variation.

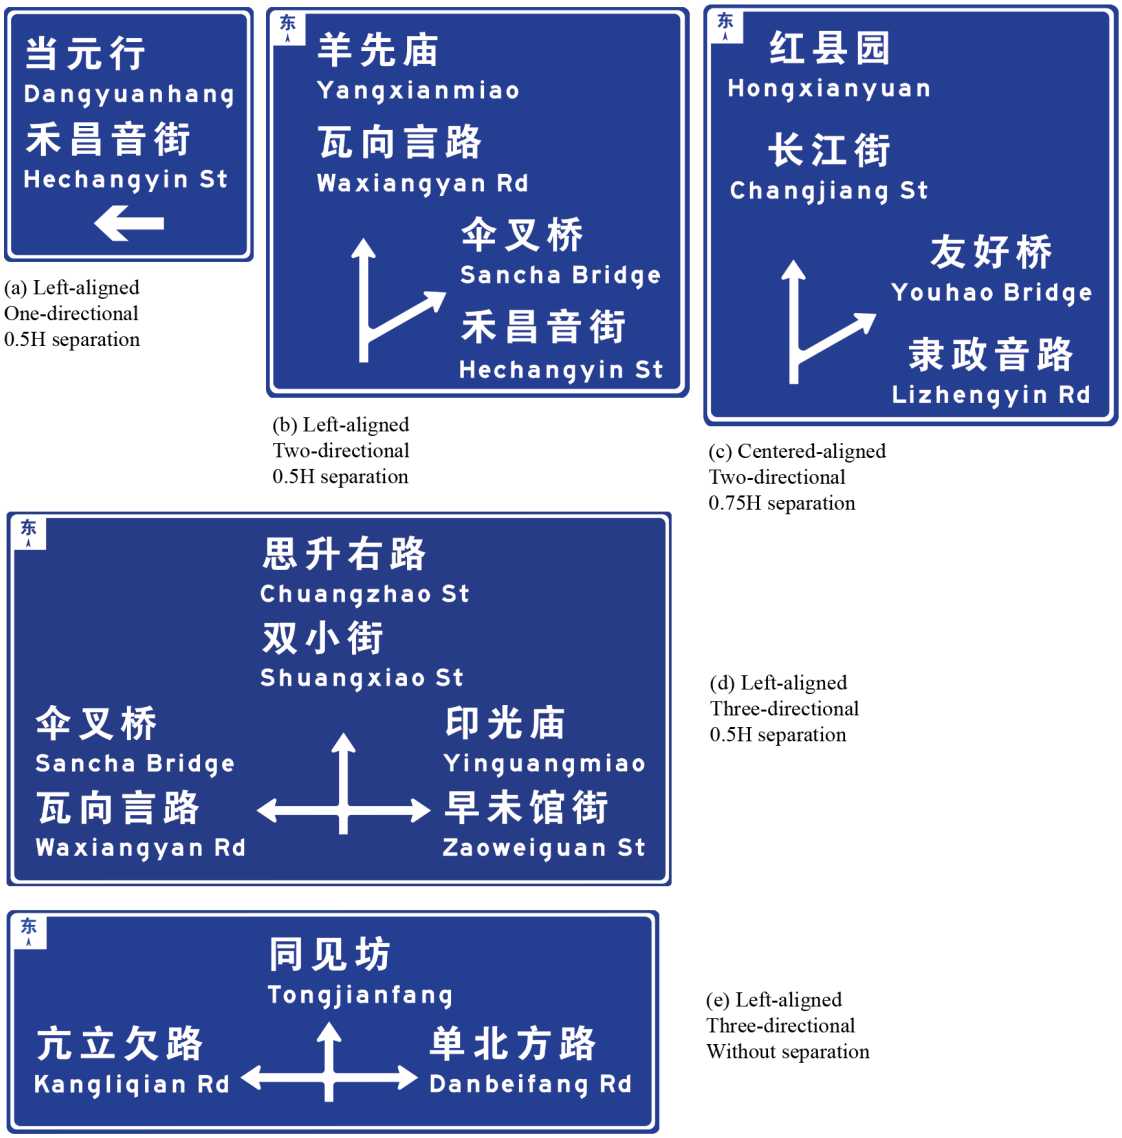

Figure 10. Five stimuli varying in text alignment, sign complexity, and separating spacing used in the study.

Other material design considerations. In addition to the above parameters, the material design also considered the scale of signs, positioning, lane width, and driving speed. Relevant elements were aligned with existing traffic standards and Chinese traffic guidelines to closely replicate drivers' interactions with real-world traffic signs. The material design deliberately excluded familiar location names, employing instead randomly combined Chinese characters and transliterated English counterparts to minimize the influence of prior knowledge, as many studies suggest that familiarity assists in reading signs (Lay, 2004; Zineddin et al., 2003).

Furthermore, the connecting spacing (Figure 2) was standardized to half the height of a Chinese character for both Investigations 1 and 2, because it was demonstrated as a generalizable spacing across sign complexities for the purpose of improving response speed (Zhang, 2021). The English location names in the stimuli were consistently set at 12 letters in length through transliteration. While in the real world this would vary, the limit set here allows for a direct comparison of outcomes and is crucial for distinguishing the effects of text alignment from those of the length of English names. In addressing the impact of the number of place names, for three-directional signs where the number of place names has been shown to influence reading performance (Zhang, 2021), this study maintained a consistent number of place names across all three-directional sign conditions. Figure 10 illustrates some of the stimuli tested in this study.

2.3. Participants and Testing Methods

In total, 36 participants were enrolled to participate in both Investigation 1 (with the separation parameter) and Investigation 2 (without separation parameter), with Investigation 1 first and followed by Investigation 2 for all individuals. Participants were recruited using specific screening questions to ensure they:

a. Had normal or corrected vision.
b. Had driving experience and age between 25 to 55 years old, as both driving experience and age play a role in reading road signs (Cantin et al., 2009). The screening criteria excluded gender due to limited evidence suggesting its influence on signage reading.
c. Did not read Chinese and used English as a first or second language.

Many experimental permutations would be possible in variations around point (c). According to Yang et al. (2020) the users of CEBTS can be grouped based on their language proficiency: Chinese drivers, bilingual drivers who are literate in both Chinese and English, and drivers who cannot read Chinese. In this study, only participants who were unable to read Chinese and relied solely on English information (and schematic cues, i.e., arrows) were included. This was because the distinct visual appearances of the two languages and the larger type size of Chinese text could aid Chinese and

bilingual groups in completing the task. According to National Bureau of Statistics of China (2020), the number of foreigners (who may not read Chinese) living in China has been steadily increasing, with the foreign population estimated to be around one million in 2020.

Another reason for focusing on this specific group is that the study was conducted in the UK, where it was challenging to find Chinese drivers with no English skills at all. Research has shown that bilingual individuals may experience additional cognitive processing when reading bilingual information, potentially leading to longer reading times due to interference from the second language (Zirnstein et al., 2018), potentially affecting the findings of this study (see Section 4 for further discussion about this criterion).

In Investigation 1, the participants' tasks were tested in three levels of sign complexity separately. In each level, the two alignments (centered and left) and the two levels of separating spacing (0.5H and 0.75H) were tested. Investigation 1 used a within-subject and between-subject mixed design. All 36 participants viewed both alignment groups: reading stimuli where the location names were centered and also reading stimuli where the location names were left aligned. The order in which participants received each stimulus was random, with the 36 participants being systematically split into two groups:

a. In one group 18 participants were shown both alignments under 0.5H separation.
b. In the other group 18 participants were shown both alignments under 0.75H separation.

Each stimulus was presented three times to each participant.

In Investigation 2, the same 36 participants performed a cross-over design by receiving six stimuli resulting of two alignments across three levels of sign complexity (2 alignments × 3 sign complexities). Each stimulus was presented only once to each participant, and the stimuli were displayed in random order.

3. Result & Analysis

3.1. Investigation 1: The Impact of Variations in Alignment in Interaction with Different Separating Spacings, Measured in Response Time

Investigation 1 examined the effect of the two alignments under 0.5H and 0.75H separating spacing. The mean and SD of response times for the centered and left alignments with both separating space under three levels of sign complexity were calculated and listed in Table 1.

Table 1. Mean and SD of response times (in seconds) for the centered- and left-alignments with both 0.5H and 0.75H separating space under three levels of sign complexity. The two alignment methods achieved a significant difference on response time when the three-directional signs using 0.5H separating spacing, in bold.

	One-directional sign		**Two-directional sign**		**Three-directional sign**	
	0.5H	**0.75H**	**0.5H**	**0.75H**	**0.5H**	**0.75H**
Centered-aligned	M: 5.139 SD: .749	M: 5.049 SD: .794	M: 4.711 SD: .874	M: 4.635 SD: .847	M: 5.433 SD: .592	M: 4.914 SD: .991
Left-aligned	M: 5.234 SD: 1.105	M: 5.482 SD: .912	M: 4.494 SD: .703	M: 4.797 SD: .679	M: 4.984 SD: 1.103	M: 5.281 SD: .764
Interaction effect	p = .339		p = .087		***p* = .012**	
Main effect	p = .191		p = .800		N/A	
Comparison	N/A		N/A		***p* = .033**	p = .140

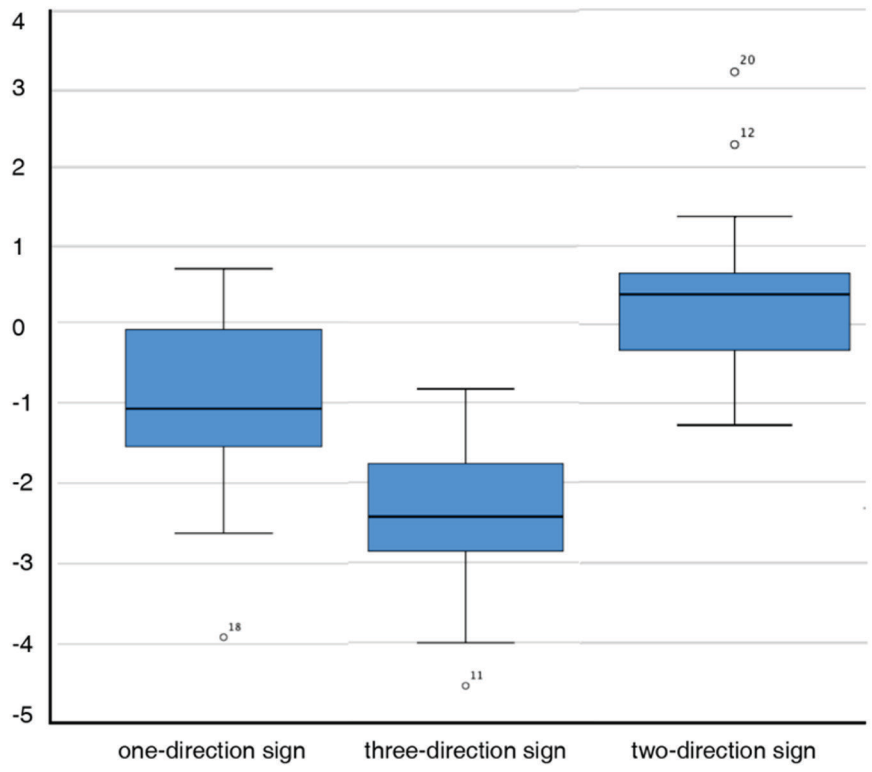

Figure 11. Outliers in three sign complexities shown in a boxplot conducted for a two-way ANOVA examining the effect of the two alignments with the two separating spacing levels on reading speed.

A two-way mixed ANOVA examined the effect of the two alignments with the two separating spacing levels on the participants' reading speed. Outliers were assessed by inspection of a boxplot (Figure 11). One outlier is detected that was more than 1.5 box-lengths from the edge of the box in the one-and three-directional sign conditions, and two outliers were detected in the two-directional sign condition. Inspection of their values did not reveal them to be extreme and they were kept in the analysis. In all three sign complexity conditions, the data was normally distributed, as assessed by Shapiro-Wilk's test of normality ($p > .05$). There was homogeneity of variances ($p > .05$) and covariances ($p > .001$), as assessed by Levene's test of homogeneity of variances and Box's M test respectively.

In both one-and two-direction sign conditions, there was no significant interaction between the separation levels and the two alignments on participants' response times:

one-directional sign condition: $F(1, 30) = .733$, $p = .399$, partial $\eta^2 = .024$;

two-directional sign condition: $F(1, 34) = 3.103$, $p = .087$, partial $\eta^2 = .084$.

The main effect analysis showed that there was no significant difference in mean response times between the two alignments regardless of the separating spacing:

one-directional sign condition: $F(1, 30) = 1.789$, $p = .191$, partial $\eta^2 = .056$;

two-directional sign condition: $F(1, 34) = 0.065$, $p = .800$, partial $\eta^2 = .002$.

In a three-direction sign condition, however, there was a significant interaction between the two independent variables on participants' response time, $F(1, 32) = 7.153$, $p = .012$, partial $\eta^2 = .183$. With 0.5h separation, the speed was significantly faster when the location names are left-aligned rather than centered ($M = .45$, $SE = .19$s, $p = .033$). However, the difference between the two alignments under 0.75h separation was not significant ($M = .37$, $SE = .24$s, $p = .140$).

3.2. Investigation 1: The Impact of Variations in Alignment in Interaction with Different Separating Spacings, Measured in Accuracy

An exact McNemar's test was included to determine if there was a significant difference in the accuracy between the two alignments for reading CEBTS. Table 2 lists the accuracy of the two alignments for each condition. It shows that there was a significant difference between the two alignments in accuracy in the three-direction sign condition with 0.5H separating spacing: $p = .039$ (bold in Table 2). With the location names left-aligned, the number of responses in which no error was made had increased to 16 (94.1%), with a concomitant reduction in the number of participants whose responses with errors to 1 (5.9%).

3.3. Investigation 2: The Impact of Alignment on Reading Performance, Measured in Response Time

Investigation 2 explored whether the two alignments may cause a significant difference in response speed and accuracy when participants reading CEBTS which includes only one location name within one direction. In other words, Investigation 2 examined how the two alignments affect legibility independently of separating spacing. The data was analyzed in terms of the three levels of sign complexity.

Table 2. Accuracy (without any errors) of two alignments on reading stimuli for each sign combination.

0.5H Separation	**Centered**	**Left-aligned**	**Exact Sig.**
One-direction sign	83.3%	72.2%	$p = .625$
Two-direction sign	88.9%	94.4%	$p = 1.000$
Three-direction sign	52.9%	94.1%	$\mathbf{p = .039}$

0.75H Separation	**Centered**	**Left-aligned**	**Exact Sig.**
One-direction sign	76.5%	58.8%	$p = .180$
Two-direction sign	83.3%	83.3%	$p = .928$
Three-direction sign	77.8%	77.8%	$p = 1.000$

Table 3. Mean and SD of response times (in seconds) for the centered and left alignments under three levels of sign complexity.

	One-direction sign	**Two-direction sign**	**Three-direction sign**
Centered	M: 2.486 SD: .679	M: 1.857 SD: .696	M: 4.910 SD: .890
Left-aligned	M: 3.419 SD: .929	M: 4.150 SD: .982	M: 4.787 SD: .896
Analysis	Centered achieved faster responses with a significant difference. 95% CI [.594, 1.270], t (35) = 5.598, d = 0.93. $\mathbf{p < .0005}$	Centered achieved faster responses with a significant difference. 95% CI [1.995, 2.590], t (35) = 15.634, d = 2.61. $\mathbf{p < .0005}$	No significant mean difference. 95% CI [−.371, .124], t (32) = −1.013 $p = .319$

Table 4. Accuracy (without any errors) of two alignments on reading stimuli in Investigation 2.

	Central	Left-aligned	Exact Sig.
One-directional sign	94.4%	100%	$p = .500$
Two-directional sign	94.4%	83.3%	$p = .180$
Three-directional sign	94.4%	97.2%	$p = 1.000$

The mean and SD of response times for the centered and left alignments in the three levels of sign complexity are listed in Table 3.

A paired-samples t-test was used to determine whether there was a significant mean difference between the response time when participants read a centered sign compared to a left-aligned sign. The three levels of sign complexity were tested separately.

In one- and two-direction sign conditions, one outlier was detected that was more than 1.5 box-lengths from the edge of the box in a boxplot. Inspection of their values did not reveal them to be extreme and they were kept in the analysis. There were no outliers as assessed by the boxplot in three-directional sign conditions. The assumption of normality was not violated, as assessed by Shapiro-Wilk's test ($p > .05$).

In both one- and two-direction sign conditions, the participants responded faster when reading the sign where the location names were centered-aligned as opposed to the location names were left-aligned. A statistically significant mean increased of .933s in the one-direction sign and 2.293s in the two-direction sign, respectively. However, there was no significant mean difference between the two alignments in a three-direction sign condition.

3.4. Investigation 2: The Impact of Alignment on Reading Performance, Measured in Accuracy

An exact McNemar's test was conducted to determine if there was a significant difference in the accuracy between two alignments when reading CEBTS. Table 4 lists the accuracy of the two alignments for each condition, from the Exact Sig. column, it shows that there was no significant difference between the two alignments in each condition.

4. Discussion

The results suggest that text alignment can have an effect on sign legibility. In Investigation 1, the results show that the participants performed at a faster speed and with higher accuracy when shown the left-alignment than the centered-alignment in a three-

direction sign condition with 0.5h separating spacing. However, this difference between the two alignments was not significant when using the 0.75h separation. Additionally, in one- and two-directional sign conditions, the two alignments did not achieve a significant difference under 0.5h or 0.75h separations. This indicates that either centered or left alignment can be used for one- and two-directional signs. Although, in a three-directional sign condition, participants responded faster and with higher accuracy with left alignment. Nevertheless, using a larger separating spacing (0.75h compared with using 0.5h separation) reduced the influence that was caused by using the two different alignments.

In one- and two-directional sign conditions, however, the two different alignments had a strong impact on reading speed in Investigation 2. This implies that the participants responded faster when they were shown the centered-alignment than left-alignment, when reading CEBTS which only indicated one place name per direction. However, this difference between the two alignments was not significant for three-directional signs.

The findings from Investigations 1 and 2 suggest that left alignment can potentially improve driver performance in scenarios where separating spacing is a factor — that is, when multiple bilingual location names are stacked vertically. However, in cases without the influence of separating spacing — where there is only one bilingual location name — centered text can enable drivers to respond more rapidly, particularly when reading one- and two-directional signs.

The findings also indicate that different combination of typographic elements and sign complexity may require different sign specifications for optimum performance. While there is substantial scope for further research, this nevertheless indicates that guidance for sign systems require much more precise typographic specifications for the spacing and positioning of different scripts. Zwaga et al. (1999, p. xvii) have suggested that information designers often "expect very detailed and narrow guidance" in comparison to the "general, broad guidelines" that emerge from research.

The findings of this study are potentially applicable to other writing systems, for example, Japanese-English and Korean-English signs. As both Japanese and Korean scripts incorporate Chinese characters (Kanji in Japanese and Hanja in Korean), the principles of alignment and spacing found effective for Chinese-English signs in this study might be transferable to Japanese-English and Korean-English signage. Future research could explore these variables in different linguistic contexts to develop legible and effective sign guidance. This would contribute to improved dual-script sign systems that support seamless navigation across various cultural and linguistic settings.

More broadly, however, the findings of the study demonstrate that it is important to consider text alignment in different everyday contexts and that different alignments

might be more effective in relation to different levels of information complexity. This study has focused on road signs. However, there are a range of everyday contexts in which people need to make quick decisions from signs. The research presented here could be extended to different sign contexts, including other transport contexts, health-care environments and urban spaces in which people may be moving on foot or with mobility aids. Signs in these contexts might be presented at different scales and heights. Thus, we propose that there is significant scope to consider the role of text alignment in bilingual signs for everyday decision-making. Such studies could help complement the existing research that focuses on the role of alignment in continuous reading from print or screen and tends to privilege a singular script or language.

In order to establish the influence of typographic factors, the study reported here only recruited participants who would only be able to read the English text and the findings had a direct impact on improving CEBTS for foreign drivers using these signs while driving in China. Nevertheless, the findings of this study might also benefit Chinese and bilingual drivers, as Yang et al. (2020) found that the use of English place names affects all three groups of drivers (see Section 2.3). The signs used in this study were from a Chinese context, but tested with participants who were could not read Chinese (to control for the inference between the two scripts). As such additional studies could establish the role of text alignment for Chinese and bilingual drivers, as this might be expected to increase cognitive load during decision making and response time. Ultimately, this study aims to improve bilingual sign design guidance for all drivers, not just a specific group. The insights gained here would be instrumental in informing future research exploring the design of more effective bilingual traffic signs, benefiting both international and local drivers.

Typographic research with multiple variables that are interrelated can be time consuming and expensive to conduct. Methodologically, the study demonstrates how video can be used to help narrow down the potential range of typographic variables and combinations that are appropriate to test. Larger scale studies using driving simulation would be essential to confirm the safest range of specifications to use in real-life signs with higher ecological validity. Nevertheless, the approach taken here can be extended to efficiently identify appropriate parameters to test as a preliminary approach. There is also potential to replicate this study using AR/VR instead of video, especially as the video material used in the presented study was generated through 3D modelling.

5. Conclusion

This study evaluates whether there is a difference between the centered and left alignment of the bilingual location name in the participants' speed and accuracy in decision making when encountering Chinese English bilingual traffic signs (CEBTS). It

also evaluates if the difference between the two alignments might relate to the changes in the separating spacing (the vertical distance to separate the two bilingual location names) and the sign complexity.

A total of 36 participants who understood English but could not read Chinese engaged in this study. Accuracy and response time were sometimes significantly better for certain text spacing conditions and certain sign complexities: the results suggest that left alignment improves speed and accuracy in three-directional signs with narrower separating spacing, whereas centered alignment may be beneficial when separating spacing is irrelevant and when used on one and two-directional signs.

There is a case to be made for more nuanced typographic guidance for road signs, given the importance of considering the multivariate interactions between typographic and spatial attributes in sign composition combined with the speed with which drivers need to be able to make accurate judgments. However, more research is necessary, which includes research with more ecological validity in real-life situations instead of research behind a computer screen.

6. References

Barker, P., & Fraser, J. (2004). *Sign design guide: A guide to inclusive signage*. JMU and the Sign Design Society.

Beier, S. (2016). Designing legible fonts for distance reading. In M. C. Dyson & C. Y. Suen (Eds.), *Digital fonts and reading* (pp. 79—93). World Scientific.

Cantin, V., Lavallière, M., Simoneau, M., & Teasdale, N. (2009). Mental workload when driving in a simulator: Effects of age and driving complexity. *Accident Analysis & Prevention, 41*(4), 763–771.

Chan, A. H., Tsang, S. N., & Ng, A. W. (2014). Effects of line length, line spacing, and line number on proofreading performance and scrolling of Chinese text. *Human Factors, 56*(3), 521–534.

Department for Transport. (2016). *Traffic signs manual*. https://www.gov.uk/government/publications/traffic-signs-manual

Dobres, J., Chrysler, S. T., Wolfe, B., Chahine, N., & Reimer, B. (2017). Empirical assessment of the legibility of the highway gothic and Clearview signage fonts. *Transportation Research Record, 2624*(1), 1–8.

Dyson, M. C. (2018). *Legibility: how and why typography affects ease of reading*. Centro de Estudios Avanzados de Diseño.

Eid, S. (2009). *Bilingual Information Design in Saudi Arabia* [unpublished master's dissertation]. University of Reading.

Fabrizio, R., Kaplan, I., & Teal, G. (1967). Readability as a function of the straightness of right-hand margins. *Visible Language, 1*(1), 90.

Federal Highway Administration. (2009). *Manual on uniform traffic control devices*.

Federal Highway Administration. (2012). *Standard highway signs*.

Gálvez, F., Ramírez, R., & Gallardo, V. (2016). The design of RutaCL: Developing and measuring performance for highway typeface. *Information Design Journal*, *22*(2), 127–146. https://doi.org/doi:10.1075/idj.22.2.06gal

Garvey, P. M., & Kuhn, B. T. (2004). Highway sign visibility. In M. Kutz (Ed.), *Handbook of transportation engineering* (pp. 11.11–11.19). McGraw-Hill Education.

Gregory, M., & Poulton, E. (1970). Even versus uneven right-hand margins and the rate of comprehension in reading. *Ergonomics*, *13*(4), 427–434.

Hartley, J., & Mills, R. L. (1973). Unjustified experiments in typographical research and instructional design. *British Journal of Educational Technology*, *4*(2), 120–131.

Hartley, J., Young, M., & Burnhill, P. (1975). On the typing of tables. *Applied Ergonomics*, *6*(1), 39–42.

Hsu, S.-H., & Huang, K.-C. (2000). Interword spacing in Chinese text layout. *Perceptual and Motor Skills*, *91*(2), 355–365.

Lay, M. G. (2004). History of traffic signs. In C. Castro & T. Horberry (Eds.), *The human factors of transport signs* (pp. 17–24). CRC Press, Taylor & Transport Signs.

Ling, J., & van Schaik, P. (2007). The influence of line spacing and text alignment on visual search of web pages. *Displays*, *28*(2), 60–67.

Luna, P. (2018). *Typography: A very short introduction*. Oxford University Press.

National Bureau of Statistics of China. (2020). *Statistical report on the foreign population in China*. https://www.stats.gov.cn

Petretta, J. (2014). Arabic sign design: Right to left and left to right. *Information Design Journal*, *21*(1), 18–23. https://doi.org/10.175/idj.21.1.04pet

Research Institute of Highway Ministry of Transport, & Beijing Communications Highway Survey and Design Institute. (2007). *Technical guidelines for the replacement of national expressway network related traffic Signs* (国家高速公路网相关标志更换工作实施技术指南). China Communications Press. http://www.jianbiaoku.com/webarbs/book/13627/676677.shtml

Research Institute of Highway Ministry of Transport, Highway Administration of Anhui Province, & Beijing Highway Survey and Design Institute. (2017). *Technical guidelines for the adjustment of national highway network traffic signs* (国家公路网交通标志调整工作技术指南). https://www.wdfxw.net/doc100277416.htm?fs=1

Rutley, K. (1972). *An investigation into bilingual (Welsh/English) traffic signs*. Transport and Road Research Laboratory.

Rutley, K. (1974). An second investigation into bilingual (Welsh/English) traffic signs. *TRRL supplementary report 34 UC*. Transport Research Laboratory.

State Administration for Market Regulation and Standardization Administration of China. (2022). *Mandatory national standard of the People's Republic of China GB 5768.2-2022, Road traffic signs and markings* (中华人民共和国强制国家标准 GB 5768.2-2022道路交通标志和标线). In Part 2: Road traffic signs (第二部分：道路交通标志). https://commons.wikimedia.org/wiki/File:GB_5768.2-2022.pdf

Tejero, P., Insa, B., & Roca, J. (2018). Increasing the default interletter spacing of words can help drivers to read traffic signs at longer distances. *Accident Analysis & Prevention*, *117*, 298–303.

Waller, R. (2007). Comparing typefaces for airport signs. *Information Design Journal*, *15*(1), 1–15.

Yang, Y., Chen, J., Easa, S. M., Zheng, X., Lin, W., & Peng, Y. (2020). Driving simulator study of the comparative effectiveness of monolingual and bilingual guide signs on Chinese highways. *Transportation research part F: Traffic psychology and behaviour*, *68*, 67–78.

Zhang, Y. (2021). Effects of text space of Chinese-English bilingual traffic sign on driver reading performance. *Displays, 67*, 102002. https://doi.org/https://doi.org/10.1016/j.displa.2021.102002

Zhang, Y., & Moys, J.-L. (2022). Effects of changes in spacing on dual-script sign legibility: The role of vertical connecting spacing in bilingual (Chinese-English) traffic signs. *Information Design Journal, 27*(2), 213–231. https://doi.org/10.1075/idj.21008.zha

Zineddin, A. Z., Garvey, P. M., Carlson, R. A., & Pietrucha, M. T. (2003). Effects of practice on font legibility. *Proceedings of the Human Factors and Ergonomics Society 47th annual meeting, 47*(13), 1717–1720. https://journals.sagepub.com/doi/pdf/10.1177/154193120304701326

Zirnstein, M., van Hell, J. G., & Kroll, J. F. (2018). Cognitive control ability mediates prediction costs in monolinguals and bilinguals. *Cognition, 176,* 87–106.

Zwaga, H., Boersema, T., & Hoonhout, H. (1999). *Visual information for everyday use: Design and research perspectives* (pp. xvii–xxxxiv).

Authors

Yuchan Zhang, Ph.D., is a lecturer in Graphic, Digital, and Communication Design at the School of Design, University of Leeds. She also serves as the Deputy Programme Leader of MA Design. Her research focuses on understanding how multilingual information can enhance the quality of the mobility experience in environments characterised by linguistic diversity. Additionally, her research explores the role of digital technologies, such as augmented reality (AR) and virtual reality (VR), in meeting users' physical and psychological needs during wayfinding processes.

Jeanne-Louise Moys, Ph.D., is an Associate Professor in the University of Leeds, School of Design. She is also the School Director of Postgraduate Research Studies. Her research typically engages with both themes of culture, identity and inclusion in design and human-centered design projects linked to typography and information design. Her pedagogic research profile, demonstrates a particular interest in partnerships and inclusion. She is a Senior Fellow of the Higher Education Academy and a UK Council for Graduate Education recognised research supervisor.

Matthew Lickiss, Ph.D., is Associate Professor (Graphic and Communication Design) at the School of Design, University of Leeds. His recent research and teaching centre on user centred information design, user experience, and design for communicating risk and uncertainty. He is particularly interested in how differing presentations of information using graphical means can elicit different decisions from users. Prior to this, his studies were focused on multimodal document theory, exploring how meaning changes or remains the same as the delivery of documents changes.

RESEARCH ARTICLE

Breaking Images: A Method for Improving Design Students' Visual Literacy

Stuart Medley [a,b] and **Hanadi Haddad** [b]

[a] UNIDCOM, IADE, Portugal; [b] Edith Cowan University, Australia

Corresponding author: Stuart Medley (s.medley[at]ecu.edu.au)

Abstract: In pursuit of competitive advantage, a growing number of organizations are adopting design-thinking strategies with a strong emphasis on visual methods. As a result, graphic design education must increase a focus on cultivating visual literacy as a thinking tool, in addition to its traditional processes for producing polished artifacts. This article proposes such a pedagogical approach; teaching students to deconstruct an image into pictures of differing levels of fidelity. The spell of realism broken, students can begin embarking on their own stylistic visual communication paths. Drawing on J.J. Gibson's distinction between image and picture, students explore how deliberate choices of pictorial form can advance specific communication goals. Classroom activities encourage them to imagine a world without text — where meaning must be carried solely by pictures — and to challenge the cultural assumption that photographic accuracy is synonymous with effective depiction.

Keywords: graphic design education; studio exercises; visual communication; visual literacy

1. Introduction

In this article, we take the position that visual literacy is a key competency for design students graduating in 2025. Over the past two decades digital cameras, smartphones and AI-driven image generators have radically lowered the cost of creating and sharing imagery. By 2011, more iPhones were "born" than human babies, and YouTube now receives more video in a single month than the three major US television networks

@: ISSUE〉 ARTICLE〉

Cite this article:

Medley, S., & Haddad, H. (2025). Breaking images: A method for improving design students' visual literacy. *Visible Language, 59*(2), 153–175.

First published online August 15, 2025.

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

Visible Language Consortium:
University of Leeds (UK)
University of Cincinnati (USA)
North Carolina State University (USA)

produced in their first 60 years (Apkon, 2013). Global organizations have responded by embracing design thinking, a mindset in which visualization and rapid prototyping are central (Sarooghi et al., 2019). Consulting firms such as Deloitte, KPMG and PwC increasingly recruit graduates who can sketch systems, storyboard user journeys and communicate strategy through diagrams rather than decks of text-heavy slides. The future looks likely to witness a further expansion of visual media and technologies. The corporate, government and not-for-profit sectors are already growing interested in developing the capabilities of their own staff. Looking for a competitive edge, a growing number of these organizations are adopting design thinking strategies, with a focus on visual methods.

How would one be educated to be a productive communicator in a world where communication happens more through pictures and less through the written word, and clients are looking for help with their visual strategies? In an economy that has moved away from the manufacture of things towards the provision of services, the future may be more about visual design as a thinking tool, and less as an end in itself. Graphic designers, illustrators and others working in fields of visual communication must start to concern themselves with how pictures can be put to use with deliberate intent. Visual literacy — defined here as the ability to read, write and reason with pictures — must therefore sit alongside conventional literacy in any contemporary design curriculum (Avgerinou, 2009; Kedra, 2018). This article responds to that need by (1) synthesizing the literature on visual literacy in design, (2) presenting a structured model for analyzing pictorial fidelity, and (3) describing studio projects that have improved students' confidence in choosing, making and manipulating images.

2. Literature Review

2.1. Defining Visual Literacy

Debates on textual literacy since the 1990s have moved away from relatively simplistic notions of what it means to be literate, to a more fine-grained understanding (Barro & Lee, 2001, p. 556). Questioning what purposes literacy skills are put towards raised the bar put in front of literacy: to read and fill out a form might be a necessary competency to participate in society, but to read and write reports is of a different level. It is an even more literate person who can write an affecting haiku or a play. Research into literacies in many disciplines is ongoing, with an urgent focus on problem solving in a technology-rich environment (Martin, 2018; Tang & Williams, 2019).

The same, dimensional attention can be given in the realm of visual literacy. As with the multifaceted approach to literacy, a dedicated program in visual literacy should ask what tasks pictures are being applied to. Looking at an instruction sheet and constructing an

IKEA chair might require different visual competencies than interpreting a commercial on YouTube, sketching a flow diagram or drawing a storyboard for an animated film.

Multiple authors portray the 21st century as a "bain d'images," in which pictorial representation dominates daily communication (Avgerinou, 2009; Elkins, 2008). Studies on problem-solving in technology-rich environments conclude that visual literacy is among the most critical graduate attributes (Martin, 2018; Tang & Williams, 2019). The convergence of ubiquitous image technology and service-oriented economies thus demands that designers master picture-based argument as fluently as text-based rhetoric.

Attempts to pin down "visual literacy," however, are made difficult by a span across multiple disciplines with, at times, conflicting agendas (Kędra, 2018). Most definitions converge on intentionality: the competent communicator can both decode and encode images with deliberate purpose (Haanstra & Wagner, 2019). *The Visual Literacy White Paper* (Bamford, 2003) and the *European Framework for Visual Competence* (Haanstra & Wagner, 2019) place awareness of why a particular pictorial strategy is chosen at the heart of the construct. This intentional dimension provides the conceptual bridge to design education, where appropriateness-to-task is already a central criterion in typography teaching. We will return to this purpose in the Concepts and Pedagogies section below.

2.2. The Need for Broader Visual Literacy in Graduates

Competency in visual literacy is necessary to prepare learners for living and working in a visually saturated environment (Doyle et al., 2018). Back in 2008 James Elkins began his *Visual Literacy* book with the notion that "a tremendous force of rhetoric has been brought to bear on the notion that ours is a predominantly visual culture" (Elkins, 2008, p. vii). Since then the rhetoric seems less like persuasion and more like stating the obvious. The future will likely witness deeper visual immersion in the "bain d'images." Pictorial representations are predominant modes of daily communication, and skills in visual "reading and writing" are therefore essential for full participation in any communication process. Ernesto Peña Alonso, author of *Visualizing Visual Literacy,* points out that among other drivers pushing the world in an increasingly visual direction, there is an increasing prevalence of contemporary visual technologies (Peña Alonso, 2018).

2.3. Industry Demand for Visual Thinking

As Meredith Davis (2018) pointed out in her Introduction to Design Futures, design for print (in the US) was in decline while the general trend for labor was showing 7% growth. At the same time, web design was growing at over twice the rate of general labor growth (15%), suggesting a shift to graphic design with a more visual interface

that affords interactions akin to conversations, as opposed to graphic design which delivers a perfected, finalized text heavy message from its source to a receiving reader. In addition, Davis describes a pace of technological change to a level where "there is too much technical knowledge for production to be a common threshold for formgiving responsibility" (p.5). While it may seem counterintuitive, design educators might think twice about trying to keep up with technology. Davis (2018) warns that:

> ...college faculty must be cautious not to overload curricula with content of temporary relevance at the expense of more enduring knowledge that transcends a rapidly changing context. At the same time, educators must rethink how to deliver lasting concepts and principles in light of a radically changed landscape for professional practice that bears little resemblance to the past. Curricula must be rethought from the ground up, not modified through endless additions to an industrial-age model. (p.5)

Davis's recommendations chime with the 21st Century business practices of some major corporations. Deloitte's "Centre for the Edge" produced a report about Toyota's "counterintuitive" approach of teaching its autoworkers to do by hand what the Toyota assembly-line robots have been doing much more quickly. An essential finding of the report is that the pace of change for technologies is far and fast outstripping the ability of staff to keep up:

> First, because of changing customer expectations and the pace at which technology is becoming able to replicate human skills, the number and variety of skills required to serve a profitable market is growing faster than the workforce can learn them. And two, skills themselves are becoming less central to creating the type of value that will differentiate a company and help build a deep, long-term relationship with customers. (Hagel et al., 2019, p. 4)

Toyota, rather than wishing to return to a time when the labor in its plants was all manual, is hoping its workers will:

> ...draw upon qualities such as imagination, creativity, problem-solving, and experimentation. The intent is to arm these workers with the right capabilities to enable them to continue to ask the right questions of unforeseen problems and develop new solutions. (p. 2)

Problem solving in a technology-rich environment is a pressing issue. The premise is that skills become obsolete (and quickly!) while capabilities endure.

Similarly, Hagel et al. (2019) identify a wider economic shift from discrete technical skills — often rapidly automated — to enduring human capabilities such as creativity, problem-framing and visual reasoning.

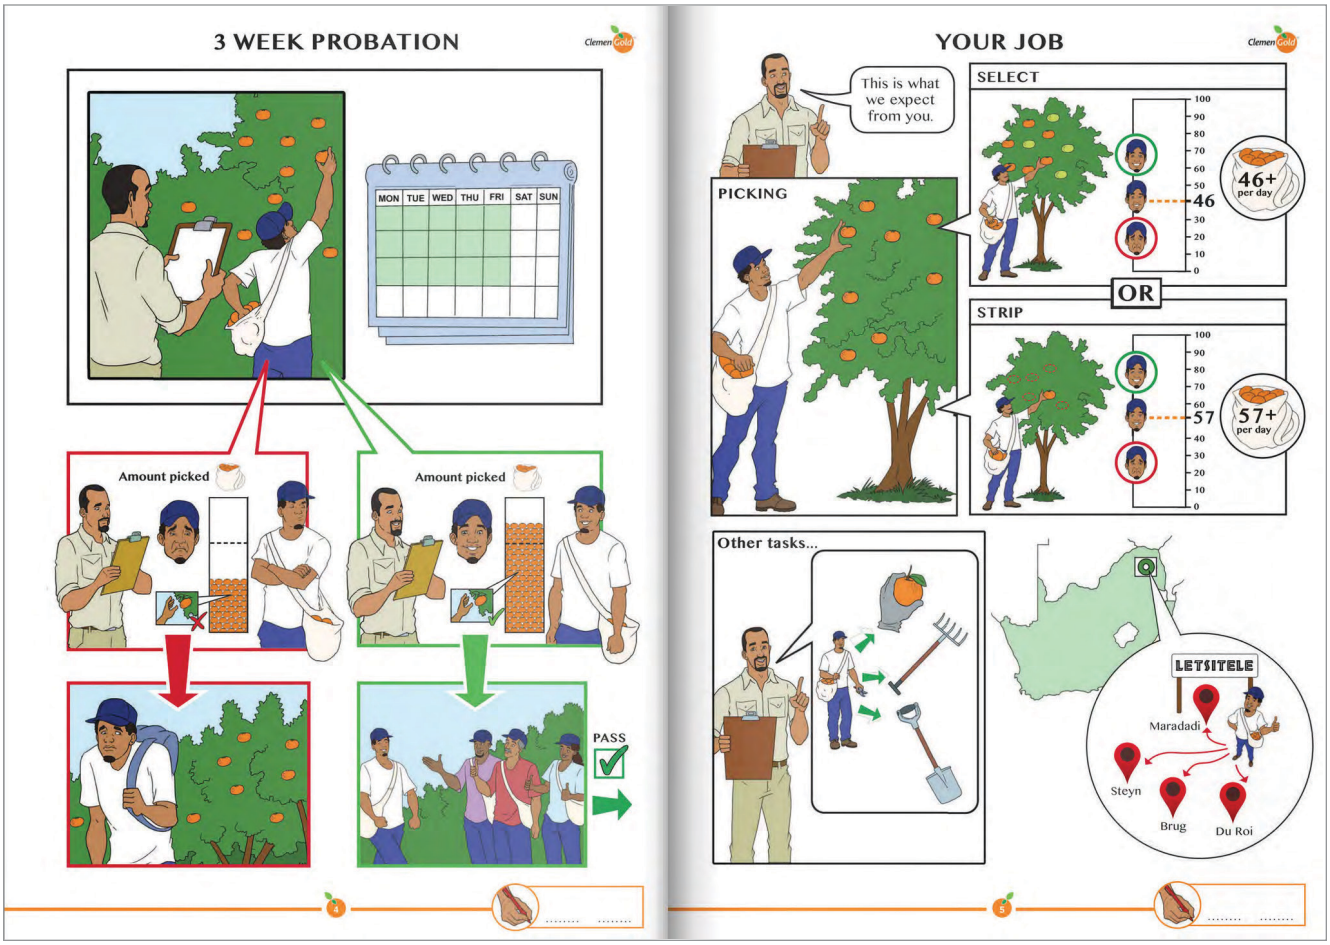

Figure 1. An example extract of a visual contract by Creative Contracts, Cape Town, South Africa, May 2016, for Indigo Fruit (Pty) Ltd, a farming business in Limpopo that grows and supplies Clemengold fruit.

Organizations describe visual thinking as an essential yet underutilized capability (Sarooghi et al., 2019). The Center for Visual Expertise now offers visual literacy training aimed at reducing workplace accidents (Thompson, 2020). More specifically, firms such as Indigo Fruit in South Africa (Figure 1), and Aurecon, a major engineering group in Australia, have begun using comics as employee contracts and induction materials. Aurecon quotes its Global Chief People Officer Liam Hayes as saying that "'The issue of engaging our talent and building their trust is becoming one of the biggest competitive differentiators across many industries and companies'" (Aurecon Group, 2018). Ranked Australia's 5th most innovative company in 2017 by the *Australian Financial Review,* Aurecon also decided to use the employment contract as an exemplar of thinking innovatively as part of its focus on shaping the future of work. This is a major leap. A threshold has surely been crossed when a practice as conservative as law can be made visual and the resulting visualizations must be legally binding.

At the same time, there is a move towards designers engaging with complexity. The complexity, in those nations which have moved away from handling the majority of their own manufacturing, comes from a shift from products to services:

> The economic basis of Western industrial nations has changed dramatically in the last three decades from manufacturing to the provision of information and services. Services now typically represent between sixty and seventy percent of the gross domestic product of developed nations. (Mager, 2008, p. 354)

Gültekin et al. (2016) explain that methods needed to guide designers in developing solutions that consider diverse stakeholder perspectives in these complex contexts are limited. Visually focused methods are high on this list of needs. Specific to the disciplines the authors teach, demand has grown locally in Australia for illustration, but not from the publishing sector. Service designers are working with illustrators and comics artists to help rapidly, and relatively cheaply, prototype service scenarios through sketches, storyboards and comic strips featuring predicted situations and personas. Illustrations in these contexts, apart from being inexpensive ways of producing prototypes, have several advantages. Illustrators can help visualize particular situations which do not yet exist, contribute their considered observations of local behaviors and appearances, amplify or completely avoid as necessary particular visual aspects pertinent to the communication, such as the intended service users' gender, ethnicity, age, and location in ways that are impossible to achieve through photography (Medley & Sercombe, 2019). Furthermore, the inclusion of pictures appropriate to the communication can increase engagement and empathy in the beholder (Grinstein et al., 2019).

In the conventional contexts of publishing, while there may be more outlets than ever for editorial illustration, the pay for illustrators has decreased. Illustrators can be commissioned from anywhere on the globe. For one-off visuals, where continuity or local sensitivities are deemed unimportant, the art director may easily choose the cheapest option. Stock illustration and more recently AI are choices that further undercut this market. In the service design context, by contrast, local knowledge is important, ongoing collaboration is paramount (projects may take years), and illustrators can be paid their worth. Roderick Mills, at the University of Brighton, suggests that illustration has value in this context because:

> Illustration has accessibility, and increasingly as it moves away from reproduction, has potential through performance and the potential of storytelling beyond the image [to] act as a pedagogic tool and for social engagement, [...] and to tell stories, narratives for varying communities. (Mills, 2017, p. 205)

In these service development contexts, the illustrations may never be published more widely than for the scrutiny of a small focus group.

2.4. Visual Literacy as a Key Competency for Graphic Design Students

The authors are lecturers and researchers in the field of visual communication; one from a professional background of graphic design and illustration. When we entered the academy in 2004 it was to teach these two disciplines. In searching for ways to make tacit, experiential knowledge explicit, we discovered that, in the educational institutions we had worked for in Australia and New Zealand, the emphasis in educating the next generation of graphic designers was placed almost entirely on typography and composition. This was exactly the way we had been taught more than a decade earlier. All the graphic design textbooks we could lay hands on as teachers were still dedicated to the application of type.

The study of graphic design is still typocentric. The word "typography" is still being used interchangeably with "graphic design," even in contemporary research that explores the history and future of education in the discipline (Vogel & Wang, 2019).

An outsider to the discipline which, according to most definitions is concerned with communication via the arrangement of type and image — Meggs even calls it a language made from these two elements — might expect both halves of that equation to be rich fields of experimentation and enquiry, but this has not been the case until late in the 20th Century (Meggs, 1992). While editor of *Eye Magazine*, Max Bruinsma commented on the "second class" status of pictures as compared with text in graphic design (Bruinsma, 1997). Some practitioners still report uncertainty in making visual meaning because of this disparity (Yates & Price, 2015).

A few texts have boldly and explicitly sought to address the relationship between text and image, such as Skolos and Wedell's (2006) *Type, Image, Message*, while being worthy developments from the typocentric texts of the 20th Century. However, they have limited their own scope by adopting a focus on photography as the principal choice visual communicators would make to embody the pictorial aspects of their designs: "designers must embrace the creative potential of photo-typographic space" (Skolos & Wedell, 2006, p. 10). Bo Bergström's very informative *Essentials of Visual Communication* nevertheless also suffers from this photographic bias (Bergström, 2009).

Graphic Design Theory (Davis, 2012) and *Visual Communication Design* (Davis & Hunt, 2017) do an excellent job of linking, in the former, the image aspects of graphic design's history to theories including Peirce's semiotics and hyperreality among others; and in the latter, focusing on user and audience interpretation of imagery through Gestalt principles — similar to Dondis' (1974) *A Primer of Visual Literacy* — and again Peircean semiotics (icon, index and symbol).

In short, while there are countless texts on typography for education in graphic design, there is only a relative handful dedicated to image for the same applications. Where

graphic design textbooks cover the image they tend to explore only the photographic. "Graphic design" is still a common term in the Anglosphere, especially regarding undergraduate design education. In the European context "graphic design" is related to the craft of giving form; skills which can be learned. Typography, in this sense, with its conventions around leading, tracking, line lengths and so forth, is well situated here: one can learn readability through rules. It is important for contemporary visual communication courses to focus on capabilities broader than those that have traditionally defined a graphic designer's competencies. "Visual communication" has a more open definition than "graphic design," reflecting Meredith Davis's (2018) recommendation to look to more enduring knowledge that transcends a rapidly changing context. This bigger picture should include visual literacy for designers.

Since 2000, illustration has dramatically re-emerged as a popular choice for graphic designers and art directors looking to embody the imagery of visual communication (Klanten & Hellige, 2005). Theorists too are seriously beginning to explore the potential of escaping the photographic image and the dominance of traditionally defined fine art discourse (Dowd, 2018). Carnegie Mellon's Susan Hagan (2019), arguing for the serious communicative application of illustration, says: "Illustrations are too often seen as shapes, colors, and textures that primarily trigger emotion, rather than as complex contributors to meaning that address the need to understand a problem from a fresh angle" (p. 163).

With the re-emergence of illustration in the new century, visual literacy in the graphic design context should encompass the hand-made aspects of picture making as well as photography and typography. To be visually literate in the context of deliberate visual communication should incorporate how pictures are made, rather than assuming they are made with a camera, and also include the uses to which pictures can be applied (Doyle et al., 2018).

2.5. The Need for Broader Visual Literacy in Graduates

Visual thinking is growing in importance. Companies seeking a competitive edge are shifting from having human resources departments that ensure the individual worker's compliance with the organization, to having human capital departments that are more aware of, and can leverage, their individual workers' skills and experiences, but also allow for individual learning styles so that staff can accrue new skills and capabilities (Llopis, 2019).

In the experience of one of the authors of this article — working with lawyers in the realm of comics contracts, or creating storyboards to help businesses design their services — it has been the client reaching out for help to think and work visually, rather than the illustrator trying to find new avenues for sales. The shift to service design

rather than products and communications ephemera means that illustrators, designers and art directors also need to change their own thinking about the venues where their work could appear; or not appear in the case of work that has no "public" beyond the clients' stakeholders in focus groups.

Visual literacy capabilities will allow illustrators and graphic designers a better view over the whole service process — where the conventional graphic design outputs such as web design, literature and signage fit in as touchpoints — but also allow them to contribute to the client's and user's understanding of processes. In the transition, illustrators and graphic designers will become consulting designers (Doyle et al., 2018, p. 891). In the undergraduate design course in our school, students experiment with visualization techniques, including storyboarding for service design. Prior to these sophisticated applications, design students can be introduced to a way of understanding these concepts that is not dissimilar to how many typography textbooks explain the application of type to convey intended meanings, as a choice focused on "appropriateness to task."

The Visual Literacy White Paper and the *European Framework for Visual Competencies* suggest that awareness of intentionality should be a key indicator of visual literacy (Bamford, 2003). This makes sense. Graphic design, illustration and visual literacy are allied here: all are concerned with visual forms created to express intended meanings to particular audiences. However, unlike words with their dictionary-searchable definitions and established vocabulary, pictures, even of a single subject can vary infinitely depending on how they are captured or made, from the specificity of a photograph of someone known to the viewer, to the very general, such as a pictogram of a person. So, how to understand and evaluate this intentionality of pictures?

While type education is built on tried and tested laws around readability of line lengths, and leading, and exacting numerical measures for type sizes, indents, gutter widths, and so on, for pictures the rules are less clear. Designers report enjoying working to rules: guidelines, a brief, a budget, a set of rules to play the game of design within. With limited choices, the work can begin sooner, the inevitably tight deadlines can be met. In our teaching around communicating using pictures, we explore one way to make images seem more evaluable.

As Robin Kinross (1984) observed, visual and verbal parallels can only go a short distance before they part ways. Is there a more nuanced way of looking for a parallel that could make sense to designers used to type application as a choice among options? Typography is about how the words are set, that is, how the message is put. Rather than the *content* of the message itself, the designer must get the *form* right. For the visual communication designer or art director, picture choice can also be about this focus

on the form more than the content; once the image is given, the focus can be put on selecting the *picture*: how the image is shown.

To develop visual literacy competences and dimensions such as the one described above, our teaching is based on some visual principles borrowed from J.J. Gibson and his intellectual heirs, such as W.J.T. Mitchell, who have productively separated the definitions of *image* and *picture* (in the fields of psychology and art history, respectively) (Gibson, 1971; Mitchell, 2009). Put simply, the *image* is the concept of what needs to be shown — either imagined or available from looking at the visible world — while the *picture* is the chosen way to apply the image. In the authors' teaching, this separation allows dealing with pictures in a way allied to how many typography textbooks explain the application of type: as a choice around "appropriateness to task." Students are asked to think about the tasks towards which they need to apply pictures.

3. Concepts and Pedagogies

3.1. The Development of Visual Literacy Capabilities

In the course the authors have designed, a question is put to students at the commencement of their studies: Why draw? Digital cameras exist on most people's phones (in the Australian experience) but what might be the advantages of *making* pictures rather than *taking* pictures? To borrow again from Gibson, what is to be gained from exploring the *chirographic*, or hand-made pictures, rather than only the *photographic*? For example, a bird in flight could be photographed or it could be drawn in a range of media and degrees of fidelity.

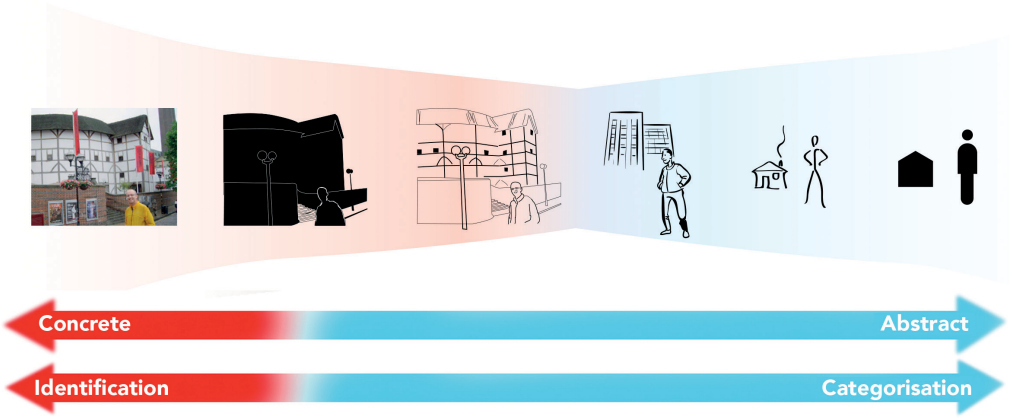

Figure 2. A visual realism continuum may show an array of pictures of iteratively reduced fidelity which bear a relationship to the same image. Pictures may range from "concrete to abstract," "specific to universal," from "identification to categorization," among other visual communication tasks.

To discuss these options with students, the concept of a visual realism continuum — a model used in different ways by various theorists in education and art history (Dwyer, 1972; Gropper, 1963; Knowlton, 1966; McCloud, 1993) — is introduced to help evaluate the communicative potential of pictures iteratively reduced in fidelity from their referent image. An example is given as Figure 2.

As students quickly familiarize themselves with this model, the photographic picture, prescribed by design practitioners and theorists as *the* way to embody the image (Müller-Brockmann, 1983), begins to look exactly like the narrow choice it is.

There are two major tasks of the visual system in establishing for a beholder what she or he is looking at in any moment. These tasks can be labelled *identification* and *categorization*. Towards the realistic or specific head of the visual realism continuum, the model helps in understanding identification. That is, which particular, individual example of the class of objects is being portrayed? Towards the abstract tail of the continuum, the model yields pictures that can help to communicate categorization. Pictures at this end of the continuum cannot be recognized as specific examples: they are more likely to have only the salient features that allow the beholder to quickly ascertain the broad category into which the drawn objects fit. The realism continuum model very quickly makes clear that there are appropriate choices to be made about depiction depending on whether the visual communication task at hand requires the viewer to be able to identify the subjects of the image or merely to be able to categorize them.

Furthermore, as the visual communication designer departs from visual realism and experiments with lower-fidelity options, they can impose their vision on to the image while making the picture. For example, color can be altered such that things that are conceptually alike (as related parts in a system) but visually dissimilar in reality, can be colored the same. Likewise, a uniform contour line may be added to the objects in a scene to more clearly relate them in the overall design. But texture and shape can also be manipulated, pushing the picture in the direction of eliciting synesthetic response from its beholders. For example, smooth textures and rounded corners can be imposed on the drawings to elicit a fondness from its beholders.

3.2. Some Teaching Methods to Improve Visual Literacy

Introductory identification and categorization sketches. We introduce these concepts to students at the commencement of the design degree in our school. To make these pictorial choices more concrete, we have devised a simple exercise that asks students to sketch two people in their classroom with just enough information in the drawings so that a third party can tell the difference between them. Students are asked to do these drawings as quickly as possible, and a stopwatch is started. Following this exercise,

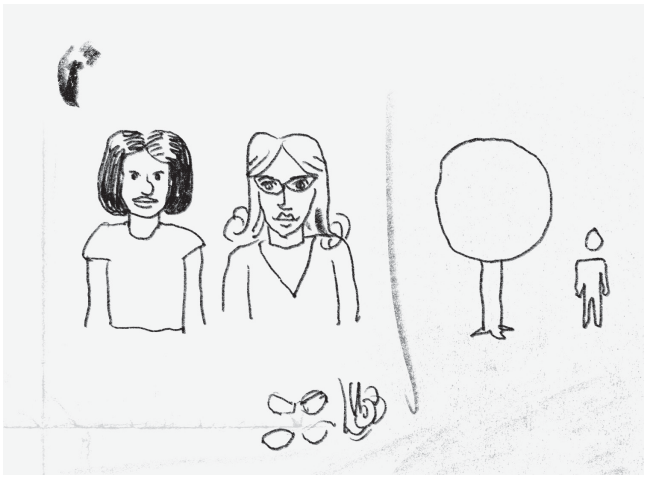

Figure 3. Identification and categorization sketch, courtesy Uriah Gray.

Figure 4. Two characters from the Breaking Images workshop, courtesy Nikola Kucharskoa. Students begin by distorting proportions rather than aiming for visual realism.

students are prompted to now draw a person and a tree so that a third party can tell the difference (Figure 3).

We have run this activity as a first class, first activity for five years. In total, more than 350 students have carried out the exercises. Generally, the first exercise takes between two and five minutes. The second takes only between 10 and 45 seconds. This exercise is a way of introducing the concept that different tasks call for different degrees of fidelity in drawing. The first exercise is about identification. As it deals with differentiating within the same class of objects (people), it necessarily requires more detail and more time and energy to produce. Specifically, more detail in the "short contours" of a drawing: those that sit within the long contours of an object's outline. The second, about categorization, requires only the long contours or outline of a subject's silhouette for the visual task of categorization to be successful. In this way, we introduce and make concrete the visual literacy concept of intentionality: to what tasks are we applying the pictures?

Comics theorist Neil Cohn (2014) laments that in the West students cannot equal the level of visual literacy exhibited in Japan, where children as young as six are able to produce complex visual narratives. Japanese culture, where drawn visual narratives are available for any age group, enables this approach. Anglophone cultures, which historically have encouraged children away from pictorial texts as they grow, by contrast celebrate individual discoveries and style (Doyle et al., 2018). This comparative weakness in visual instruction might be harnessed in a different way. If the escape from visual realism enables, among other things, quick categorization of subjects, and

increases appeal and engagement, encouraging the development of each individual's own drawing approaches or "production scripts" — how they put drawings together according to what they feel and know rather than what they see — might point to as yet undiscovered strengths of drawing and perhaps higher appeal and engagement with beholders of visual communication.

Breaking images: Accelerating towards finding a drawing style. In recent years, one of the authors of this article has conducted workshop as a "hothouse" for developing an individual's illustration style. The workshop questions the value of life drawing and realistic proportions and instead proposes steps to focus students on and accelerate towards their own approach to drawing (Figure 4). At the conclusion of the workshop two hours later, students have two characters of their own design (for the potential for tension or conflict in a narrative) in a two-panel comic strip that demonstrates aspects of the characters' traits (Figures 5 and 6).

In the first step, students are asked to draw a "normal" human face for their age group. Each student may use a remembered formula for this, or they can move the elements around in the face (e.g., erasing and re-drawing eyes, nose, mouth, face outline) until the face configuration begins to "disappear," by which, it is explained, none of the individual elements draws attention to itself as too large or small or placed in the wrong location.

Given that style becomes apparent where depictive choices have been made, in the next step students' choices are made overt: they are urged to "break" the face drawing by moving the elements around within the configuration. Students can begin by erasing

Figure 5 (left). Two-panel comic strip produced at the end of the workshop, courtesy Zuzanna Dominiak.

Figure 6 (below). Examples from the Breaking Images workshop, courtesy of (from left) Wenjie Zheng, Joanna Fung, Bruce Mutard, and Pete Corey.

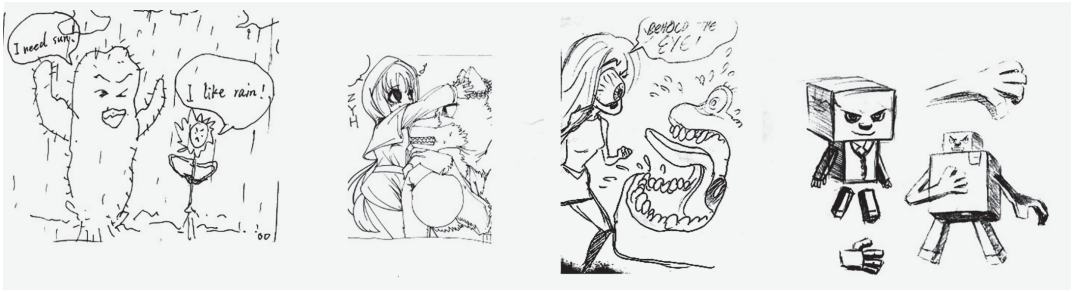

the eyes and re-drawing these elsewhere in the face, perhaps below the level of the mouth, or by moving the ears up to the top of the head. Paradoxically, the drawings have become more unique and memorable; more like useful character designs.

Students then should take a schematic approach to the entire figure, drawing and reflecting on the "right way" for a human figure to look. They are asked, "where do legs and arms bend?," and "are arms and legs thicker at the top or the bottom?" Again, choices are prompted. Students make multiple sketches, altering the relative drawn lengths of thigh and shin, and of upper arms and forearms, and inverting the thicknesses of limbs. Students learn that proportions matter much less than other perceived physical properties that comprise a visual syntax of the figure. For example, having the fold lines around the elbow working in the right direction is much more germane to the "reading" of the drawing than whether the elbow is halfway along the arm. Students have discovered that where elements are placed in relation to each other is more important for readability than actual proportion.

Students are asked to reflect upon the points at which they thought the tension between normality and uniqueness of expression "felt right" for them as picture-makers. In this way they begin to become aware of their own "production scripts": how a drawing should be put together from what the student knows and feels rather than from how things appear in reality. Escaping pictorial realism allows clear differentiation between character designs. As a result, the impression of consistency in repeated drawings strengthens, and the characters appear clearer to a beholder of the designs.

Caricature. Students are already seeing the outlandish results of deliberately "getting it wrong" with their schematic character anatomy. Caricaturing their own creation further exaggerates those differences that make the character unique. Students caricature their characters by comparing their unique design to a norm for that category of subjects, or, where no norm exists for the character created (for example, if their creation is a bizarre monster), a relatively simple composite may be derived through a combination of the common landmarks (for example, eyes, ears, legs) found in the two designs the participant has just made. Each of the two character designs can be compared to the composite and the differences from the composite should be exaggerated, as in Figure 7 for example.

The reasons for caricaturing the characters, rather than workshop participants accepting their initial character designs, is to push further the possibilities of the design, especially those visual aspects that make it unique, and to make clear to students the plasticity of any drawn creation. To bring the process back to visual communication with intentionality, caricaturing is an important visual capability to have and to understand. Any specific subject, including vehicles or a landscape, colors and shapes,

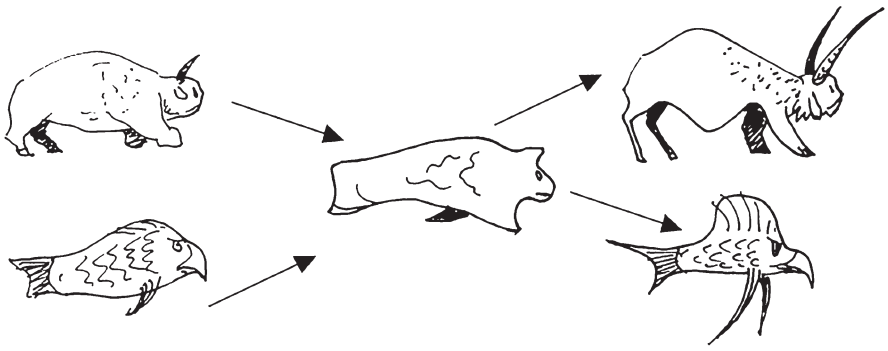

Figure 7. Common “landmarks” are identified in the original character designs (at left), from which a rough composite is made (at center). Landmarks in each character design are moved further away from the landmarks in the composite for the final, caricatured designs (at right). Illustration courtesy of Hanka Nazim.

can be caricatured. It is a fundamental human psychological faculty which improves recall of specific examples (Dror et al., 2008).

For students that progress to be graphic designers or art directors rather than illustrators, this workshop process gives them a compressed glimpse of what illustrators are doing over the course of their career: going deep within themselves to discover what feels right in the way of picture making. Illustrators can record the visible world, to be sure, but they combine recording with interpretation; bringing out their own emotions and intellect on to the page. They make visible as much as recording the already visible. They are building metaphors (Medley, 2019): making concrete on paper or screen concepts which otherwise are too abstract to grasp intellectually or emotionally. These metaphors work in illustrations but also in infographics.

As students progress through their course and gain a fuller understanding of the importance of visual fidelity, and the even greater importance of departures from visual fidelity, these pictorial projects become more complex. As the importance of visual communication only seems to be growing as the future becomes the present, and since these visual literacy capabilities are framed as a necessary skill for the future of work, these pictorial projects can themselves be future-focused. The content used in our pedagogy often comes from science fiction and science fact.

Future warning. The authors of this article prompt students to put themselves in the position of being the communication design experts consulted on the project. The first described here is based upon a US Government report from 1993, *Expert Judgment on Markers to Deter Inadvertent Human Intrusion into the Waste Isolation Pilot Plant,* or WIPP. As its full title suggests, the report describes the need to gather experts to make recommendations on how to stop people opening up a radioactive waste store (Trauth

et al., 1993). The added difficulty for this communication task is because the half-life of the radioactive waste it stores will remain deadly 10,000 years hence.

As with the original report, the students need to work on the assumption that none of the writing systems currently in use around the world may be in use in a world that will be 10,000 years older. This is a sound assumption to make: Egyptian hieroglyphics, though only around 5,000 years old, would have been incomprehensible without their comparison in the Rosetta Stone to known writing systems. Symbols require a common understanding between two or more parties. If the design problem at hand is based upon those future symbols being unknowable, what other strategies of communication can a design team adopt? Writing is out of the question.

After some discussion and testing, students often decide upon strategies using highly iconic (as opposed to symbolic) forms. A useful approach is to make a human figure, perhaps shown doubled over in apparent agony. But this is not the only strategy students have adopted. Where they choose an abstract approach, synesthetic forms, such as jagged edges and sharp points, are designed to repel any future discoverers. Interestingly, students often enquire if they can use typography. The answer is yes, but they are not allowed to use known letterforms. Again, the synesthetic approach, making repellent visual forms, can come into play. From this, students have observed that typography can conform to a kind of picture theory, but not vice versa. Picture trumps type in visual communication!

Student responses are made as prototype statues or bas-relief carvings (made from board) whose material considerations become paramount; what matter will stand the test of time? The medium of these messages matters very much. The prototypes are then tested for comprehensibility by students from outside the class. The feedback is positive, these generic human figures, shown in specific poses, and the repellent abstract forms help to communicate a sense of potential harm in the beholder. Students are learning through these exercises and associated class discussions and lectures that the usefulness of pictures in visual communication is only rarely allied with accurate drawing.

3.3. Challenges in Teaching to Improve Visual Communication Capabilities

Teaching deliberate communication based around images and pictures is not without its difficulties. Some of the greatest challenges are a function of the assumptions that circulate around pictures and drawing. As we explained above, graphic design theory itself, with its extreme bias towards typography, has not helped advance the cause of pictorial communication. The first problem we wish to discuss here is the assumption that pictures are vaguer than words when it comes to contexts of deliberate communication. The second assumption is around what constitutes good depiction. If

educators were clearer about these, they may in turn be able to overcome Kędra's (2018) identification of the marginalization of visual literacy across education curricula. Her observation reminds us that perhaps we have not progressed in education, at least in the Anglosphere, since Fransecky and Debes (1972) observed that we are "weaned away" from pictures as we progress as readers (p. 23).

To the first assumption, that pictures are too vague to be used deliberately without the addition of words to make clear the meaning of the communication: words, the argument runs, are needed to disambiguate pictures. Even some theorists in visual communication, such as, Timothy Donaldson, author of *Shapes for Sounds: Why Alphabets Look Like They Do,* believes this to be the case: "images always need to be explained with supporting text. More often the real communicating is being done by the words" (Donaldson, 2008, p. 9). On the one hand is the claim that pictures are vaguer than words. On the other is the reader who goes to see the movie of her favorite novel, and complains: "that's not how I pictured that setting or the main character." The moving picture has made concrete what the words allowed to remain relatively abstract.

Barnard and Johnson (2005) deftly countered the prevailing view that words outperform pictures when deliberate communication is a requirement. They demonstrated by adding pictures to verbal labels that pictures can disambiguate words just as effectively as the other way around. In the case of both pictures and words, it seems that each benefits from the context that sequence provides. The meaning of a single word in a sentence can be much clearer than that word by itself. Likewise, a single picture may be vaguer than the same picture in a sequence, for example, on an assembly instruction sheet or in a comic.

Another challenge is around the assumption of what constitutes good depiction. The WIPP student exercise, for example, is based on the premise of an illiterate audience for the messages. Communication is achieved, at least to begin with, through relatively iconic forms whose referent is clear. These projects present a science fiction of illiteracy, where communication must be achieved through figural forms. But what about those who cannot draw? Would such a future, requiring accurate depiction as a communication method, not be onerous and exclusionary? Gombrich (2002) declared that the "first prejudice teachers of art appreciation usually try to combat is the belief that artistic excellence is identical with photographic accuracy" (p. 158). As we have explained above however, high visual fidelity is really only necessary to help the beholder identify a specific subject within one class of similar subjects. Those who say they cannot draw often mean they cannot draw realistically (Dowd, 2018; Quito, 2018). For most graphic communication tasks, it is enough for the visual communicator to help the beholder *categorize*: to distinguish one class of subjects from a different class of subjects. In the case of the WIPP design problem, as long as the human figures the students design look

more like humans than they look like anything else, the communication is understood by the beholders tasked with decoding the warning.

Even early visual literacy theorists were fixated on visual realism. They recommended the camera as the means through which children could acquire visual literacy. The assumption that the visually real photographic or filmic image is the one to become expert in continues today with Stephen Apkon's (2009) *The Age of the Image* which, perhaps because the most prevalent kind of image on the Web is the photographic, makes no special mention of any other kind of picture. Cameras are miraculous tools, but provide communicators with an extremely narrow bandwidth of visual communication possibilities. They excel at recording that which is already visual, but not in making visual that which is not. We have intended to show above that the act of making visible is a necessary capability for future graduates of design.

Capabilities in both taking and making pictures are necessary to develop students' visual literacy. Research into style, pattern recognition and comprehension are necessary to further break down the historical assumptions about what constitutes good depiction. The spell of visual realism needs to be broken, so that, at most, students see it as one among many decisions for fixing the image into a picture.

3.4. Design Practices to Engage with Visual Thinking

Working in service design is one way for design educators to acquire this experience of applying pictures, rather than just type, in visual communication.

Most practicing visual communication designers will have by now come across a business that is looking for help in developing a service rather than a product and is looking for guidance through the relatively unfamiliar landscape of design thinking. Design instructors need to avail themselves of some of these practical experiences, in the process becoming familiar with visual design practices as part of the process to understand complexity, rather than the shape of the process's outcome. Visual communication needs to be appreciated as an input to design processes rather than (just) the resulting output. Discourse in the field discusses the benefits of this application towards designerly thinking, including idea generation, prototyping, visualization and its relationship to aesthetic style or aesthetic-lessness and meaning (Tonkinwise, 2011).

For this to happen, design educators will need to be allowed and encouraged to take time out from teaching for practice. Not only has work in service design strengthened the visual communication capabilities of one of this article's authors; it has alerted us to the possibility that clients may have stronger and clearer ideas about the application of the visual capabilities than designers and illustrators possess.

3.5. Course Programming to Enable the Acquisition of Appropriate Teaching Practices

One of the difficulties in focusing course programming on service design is around communicating to high school students the changing work environment for university design graduates. There is a disconnection, at least in Western Australia, between what high school students know about design and the reality of the changes in the design professions. Design in the academy is stuck between these points: how to market service design — working with intangible issues — to students who have in mind the design of physical artifacts. Design educators have always been located between school and industry. Building a bridge between these worlds is more complex than ever with the rapid changes in developed economies moving to the more complex contexts of service development rather than product manufacturing. Davis (2018) insists that "college programs must decide what they can and cannot promise students as professionally relevant outcomes of an undergraduate curriculum" (p. 5).

Our design team carries out "design roadshows"; visiting high schools to gauge, and increase, interest in our design courses. In our surveys of many high schools in Western Australia, the vast majority of students know of "graphic design" as a professional pursuit and have some notions of what graphic designers produce. Similarly, high school students understand at least narrow ideas of game design, because many of them play games in their leisure time, and of interaction design, since most are familiar with navigating the Web. In other words, students understand design as an output, but not as an input into business and organizational processes. Service design, systems or strategic design and other related terms are completely unfamiliar and yet are likely to be of even greater importance by the time the current crop of high school graduates has gone through higher education to become 2025's job seekers. Program directors need to work with university marketing staff to inform future students of these developments.

4. Conclusion

Shifts in technology, communication, and industry are demanding a wider and deeper understanding of the practice of graphic design in which picture generation is as equally valued as typographic form. The possibilities of more symbolic interactions using pictures as the communications medium could be explored and exploited if our society became more visually literate.

Meanwhile developments around thinking visually and more holistically are continuing apace in the corporate, government and not-for-profit sectors. As 21st Century businesses look for a competitive edge, visual thinking around their processes is seen as essential. Graphic designers and illustrators must quickly develop their visual literacy

capabilities. No other existing disciplines are better placed to adopt a position of visual expertise and take charge of deliberate visual communication.

One way of advancing the study of images and pictures for deliberate communication is to separate the definitions of image and picture. Students quickly grasp that type can be the inaudible "voice" through which the client's words are expressed to an audience. Likewise, they can understand that a picture, rendered in varying degrees of fidelity, can be the digital or printed embodiment of the image aspects of the client's message. The visual realism continuum is a useful conceptual model along which to place these pictorial choices. Students (and future designers) can then make informed, visually literate, decisions about whether the task at hand requires the communication's beholders to identify or merely categorize the subjects pictured.

In teaching these concepts, imagining communication scenarios where writing is disallowed as a system for symbolic interaction continues to be one effective way to build students' capabilities in visual literacy.

5. References

Abas, S. (2019). Reading the world – Teaching visual analysis in higher education. *Journal of Visual Literacy, 38*(1–2), 100–109. https://doi.org/10.1080/1051144X.2019.1574120

Apkon, S. (2013). *The age of the image: Redefining literacy in a world of screens.* Farrar, Straus and Giroux.

Ashwin, C. (2016). What is a drawing? *Drawing: Research, Theory, Practice, 1*(2), 197–209. https://doi.org/10.1386/drtp.1.2.197_1

Aurecon Group. (2018, May). *Australia's first visual employment contracts launched.* https://www.aurecongroup.com/about/latest-news/2018/may/visual-employment-contract

Avgerinou, M. D. (2009). Re-viewing visual literacy in the "bain d'images" era. *TechTrends, 53*(2), 28–34.

Balchin, W. G. V., & Coleman, A. M. (1966). Graphicacy should be the fourth ace in the pack. *Cartographica: The International Journal for Geographic Information and Geovisualization, 3*(1), 23–28.

Bamford, A. (2003). *The visual literacy white paper.* Adobe Systems.

Barnard, K., & Johnson, M. (2005). Word sense disambiguation with pictures. *Artificial Intelligence, 167*(1–2), 13–30.

Barro, R. J., & Lee, J. W. (2001). International data on educational attainment: Updates and implications. *Oxford Economic Papers, 53*(3), 541–563. https://doi.org/10.1093/oep/53.3.541

Bergström, B. (2009). *Essentials of visual communication.* Laurence King Publishing.

Bruinsma, M. (1997). Learning to read and write images. *Eye Magazine, 25*(7), 3–4.

Cohn, N. (2014). Framing "I can't draw": The influence of cultural frames on the development of drawing. *Culture & Psychology, 20*(1), 102–117. https://doi.org/10.1177/1354067X13515936

Davis, M. (2012). *Graphic design theory.* Thames & Hudson.

Davis, M., & Hunt, J. (2017). *Visual communication design.* Bloomsbury Visual Arts.

Davis, M. (2018). Introduction to design futures. AIGA Design Futures Trends. https://www.aiga.org/sites/default/files/2021-02/introduction-to-design-futures.pdf

Donaldson, T. (2008). *Shapes for sounds.* Mark Batty Publisher.

Dondis, D. A. (1974). *A primer of visual literacy.* MIT Press.

Dowd, D. B., & Plunkett, S. H. (2018). *Stick figures: Drawing as a human practice.* Spartan Holiday Books.

Doyle, S., Grove, J., & Sherman, W. (2018). *History of illustration.* Bloomsbury Publishing.

Dror, I. E., Stevenage, S. V., & Ashworth, A. R. S. (2008). Helping the cognitive system learn: Exaggerating distinctiveness and uniqueness. *Applied Cognitive Psychology, 22*(4), 573–584. https://doi.org/10.1002/acp.1383

Dwyer, F. M. (1972). *A guide for improving visualized instruction.* Learning Services.

Elkins, J. (Ed.). (2007). *Visual literacy (1st ed.).* Routledge. https://doi.org/10.4324/9780203939574

Fransecky, R. B., & Debes, J. L. (1972). *Visual Literacy: A Way to Learn—A Way to Teach.*

Gibson, J. J. (1971). The information available in pictures. *Leonardo, 4*(1), 27–35. https://doi.org/10.2307/1572223

Gombrich, E. H. (2002). *Art and illusion: A study in the psychology of pictorial representation (6th ed.).* Phaidon Press.

Grinstein, A., Hagtvedt, H., & Kronrod, A. (2019). Aesthetically (dis)pleasing visuals: A dual pathway to empathy and prosocial behavior. *International Journal of Research in Marketing, 36*(1), 83–99. https://doi.org/10.1016/j.ijresmar.2018.09.003

Gropper, G. L. (1963). Why is a picture worth a thousand words? *Educational Technology Research and Development, 11*(4), 75–95.

Gültekin, P., Bekker, T., Lu, Y., Brombacher, A., & Eggen, B. (2016). Combining user needs and stakeholder requirements: The value design method. In P. Markopoulos et al. (Eds.), *Collaboration in creative design* (pp. 97–119). Springer. https://doi.org/10.1007/978-3-319-29155-0_6

Haanstra, F., & Wagner, E. (2019). A European framework for visual competencies. In *The International Encyclopedia of Art and Design Education* (pp. 1–9). https://doi.org/10.1002/9781118978061.ead101

Hagan, S. (2019). Collaborative problem solvers in three environments. In A. Male (Ed.), *A companion to illustration* (pp. 163–181). Wiley. https://doi.org/10.1002/9781119185574.ch7

Hagel, J., Brown, J. S., & Wooll, M. (2019). Skills change, but capabilities endure. *Deloitte Insights.* https://www2.deloitte.com/content/dam/insights/us/articles/6332_From-skills-to-capabilities/6332_Skills-change-capabiliites-endure.pdf

Kędra, J. (2018). What does it mean to be visually literate? Examination of visual literacy definitions in a context of higher education. *Journal of Visual Literacy, 37*(2), 67–84. https://doi.org/10.1080/1051144X.2018.1492234

Kinross, R. (1984). Semiotics and designing. *Information Design Journal, 4*(3), 190–198. https://doi.org/10.1075/idj.4.3.02kin

Klanten, R., & Hellige, H. (Eds.). (2005). *Illusive: Contemporary illustration and its context.* Die Gestalten Verlag.

Knowlton, J. Q. (1966). On the definition of "picture." *AV Communication Review, 14*(2), 157–183.

Llopis, G. (2018, January 8). HR departments must urgently become human capital departments. *Forbes.* https://www.forbes.com/sites/glennllopis/2018/01/08/hr-departments-must-urgently-become-human-capital-departments/

Mager, B. (2008). Service design. In M. Erlhoff & T. Marshall (Eds.), *Design dictionary* (pp. 354–356). Birkhäuser.

Martin, J. P. (2018). *Skills for the 21st century: Findings and policy lessons from the OECD survey of adult skills. IZA Policy Paper No. 138.* Institute of Labor Economics. http://hdl.handle.net/10419/180633

McCloud, S. (1993). *Understanding comics: The invisible art.* Kitchen Sink Press.

Medley, S. (2019). Making visible: Illustration through identification, categorization, and metaphor. In A. Male (Ed.), *A companion to illustration* (pp. 19–46). Wiley. https://doi.org/10.1002/9781119185574.ch1

Medley, J., & Sercombe, J. (2019). 'The kids just lit up': Mode of depiction and number of panels matters in service design storyboards. In *CONFIA 2019 Proceedings* (pp. 340–349).

Meggs, P. B. (1992). *Type and image: The language of graphic design.* John Wiley & Sons.

Mills, R. (2017). Beyond image: Situated illustration, a pedagogic practice. In *5th CONFIA International Conference on Illustration and Animation Proceedings* (pp. 205–215). Edicao IPCA.

Mitchell, W. J. T. (2009). Visual literacy or literary visualcy? In J. Elkins (Ed.), *Visual literacy* (pp. 19–38). Routledge.

Müller-Brockmann, J. (1983). *The graphic designer and his design problems.* Hastings House.

Peña Alonso, E. J. (2018). *Visualizing visual literacy* (Master's thesis, University of British Columbia). https://doi.org/10.14288/1.0368982

Postema, B. (2013). *Narrative structure in comics: Making sense of fragments.* RIT Press.

Quito, A. (2018). Drawing is the best way to learn, even if you're not Leonardo da Vinci. *Quartz.* https://qz.com/quartzy/1381916/drawing-is-the-best-way-to-learn-even-if-youre-no-leonardo-da-vinci

Sarooghi, H., Sunny, S., Hornsby, J., & Fernhaber, S. (2019). Design thinking and entrepreneurship education: Where are we, and what are the possibilities? *Journal of Small Business Management, 57*(S1), 78–93. https://doi.org/10.1111/jsbm.12541

Skolos, N., & Wedell, T. (2006). *Type, image, message: A graphic design layout workshop.* Rockport Publishers.

Tang, K.-S., & Williams, P. (2019). STEM literacy or literacies? Examining the empirical basis of these constructs. *Review of Education, 7*(3), 675–697. https://doi.org/10.1002/rev3.3162

Thompson, D. S. (2020). *Picturing the state of visual literacy initiatives today. In J. Elkins (Ed.), Visual literacy (2nd ed.).* Routledge. https://doi.org/10.2307/j.ctv36xvs64.11

Tonkinwise, C. (2011). A taste for practices: Unrepressing style in design thinking. *Design Studies, 32*(6), 533–545. https://doi.org/10.1016/j.destud.2011.07.001

Trauth, K. M., Hora, S. C., & Guzowski, R. V. (1993). *Expert judgment on markers to deter inadvertent human intrusion into the Waste Isolation Pilot Plant (No. SAND--92-1382).* Sandia National Labs.

Vogel, C. M., & Wang, X. (2019). Observations on the state of design education: Past, present, future. *Design Management Review, 30*(1), 26–32. https://doi.org/10.1111/drev.12158

Yates, D., & Price, J. (2015). *Communication design: Insights from the creative industries.* Bloomsbury Publishing.

Authors

Stuart Medley is a research fellow at UNIDCOM, Portugal, and an honorary associate professor of design at Edith Cowan University exploring the application of comics for legal contracts and in the health sector. He authored the book *The Picture in Design* and other publications,

including for Wiley's *Companion to Illustration.* He makes comics to improve employment and housing services, and has worked as an illustrator for 25 years with clients including Anglicare and Austal Ships in Australia, and the Imperial War Museums in the UK. He is a co-founder and art director for Hidden Shoal, a record label with a roster of international artists, with the BBC, Canadian Broadcasting, and Showtime as clients. He is a co-founder and funding director of the Perth Comic Arts Festival.

Dr. Hanadi Haddad has a background in interactive media, design and interactive television. She is a lecturer and supervisor at Edith Cowan University in design in the undergraduate and postgraduate programs. Her research interests include strategic visual communication, design thinking, service design and innovation, co-design/creation, usability, design theory and work integrated learning. She was previously a post-doctoral research associate at Murdoch University. A graduate of Curtin University, her PhD involved multidisciplinary research into visual modalities of character agents between the School of Design and School of Computer Science.

RESEARCH ARTICLE

Addressing Uncertainty in LLM Outputs for Trust Calibration Through Visualization and User Interface Design

Helen Armstrong [a], **Ashley L. Anderson** [a,b], **Rebecca Planchart** [a], **Kweku Baidoo** [a], **and Matthew Peterson** [a]

[a] Department of Graphic Design and Industrial Design, North Carolina State University, Raleigh, NC, USA; [b] School of Visual Arts, Virginia Polytechnic Institute and State University, Blacksburg, VA, USA

Corresponding author: Helen Armstrong (hsarmstr[at]ncsu.edu)

Abstract: Large language models (LLMs) are becoming ubiquitous in knowledge work. However, the uncertainty inherent to LLM summary generation limits the efficacy of human-machine teaming, especially when users are unable to properly calibrate their trust in automation. Visual conventions for signifying uncertainty and interface design strategies for engaging users are needed to realize the full potential of LLMs. We report on an exploratory interdisciplinary project that resulted in four main contributions to explainable artificial intelligence in and beyond an intelligence analysis context. First, we provide and evaluate eight potential visual conventions for representing uncertainty in LLM summaries. Second, we describe a framework for uncertainty specific to LLM technology. Third, we specify 10 features for a proposed LLM validation system — the Multiple Agent Validation System (MAVS) — that utilizes the visual conventions, the framework, and three virtual agents to aid in language analysis. Fourth, we provide and describe four MAVS prototypes, one as an interactive simulation interface and the others as narrative interface videos. All four utilize a language analysis scenario to educate users on the potential of LLM technology in human-machine teams. To demonstrate applicability of the contributions beyond intelligence analysis, we also consider LLM-derived uncertainty in clinical decision-making in medicine and in climate forecasting. Ultimately, this investigation makes a case for the importance of visual and interface design in shaping the development of LLM technology.

Implications for practice: This article provides guidance on explainability and transparency for AI interface design through the consideration of uncertainty in LLM summaries. Our Uncertainty Framework for Explainable Summaries (UFES) can guide system design and help users interpret

@: ISSUE〉 ARTICLE〉

Cite this article:

Armstrong, H., Anderson, A. L., Planchart, R., Baidoo, K., & Peterson, M. (2025). Addressing uncertainty in LLM outputs for trust calibration through visualization and user interface design. *Visible Language, 59*(2), 176–217.

First published online August 15, 2025.

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

Visible Language Consortium:
University of Leeds (UK)
University of Cincinnati (USA)
North Carolina State University (USA)

and act on LLM outputs. Our specifications for 10 features in a Multiple Agent Validation System (MAVS) can be implemented with current technology to aid user understanding, trust calibration, and decision-making. As an open resource, we provide eight visualization options with readable code that represent uncertainty within relevant passages of text. We also include four prototypes of MAVS to reference for implementation, or to educate stakeholders on the value of LLMs when carefully leveraged. While we situate this guidance in an intelligence analysis context, outcomes are relevant to any LLM systems that produce summaries of information.

Statement of applicability: This is a visual and interface design investigation in the context of intelligence analysis (specifically language analysis). It is relevant to any human-centered application of LLM technology that explicitly addresses uncertainty in outputs.

Application domains: intelligence analysis, knowledge work, computer science (LLM technology and RAG), psychology, and human-computer interaction; connections are also drawn to LLM use in medicine (diagnosis) and climate science (forecasting).

Keywords: explainable AI; human-machine teaming; intelligence analysis; large language models; trust calibration; uncertainty; user interface design; visual representation

1. Introduction

Researchers and practitioners increasingly rely on generative artificial intelligence (AI) systems as essential tools for navigating information-dense environments. When faced with vast quantities of information, humans can deploy large language models (LLMs) — AI systems trained to understand and generate human language — to efficiently summarize content. But these summaries are not foolproof. LLM summaries will always include a level of uncertainty. And we do not yet have a conventional means for understanding this uncertainty or for presenting it to users. Without this, users cannot properly calibrate their trust in AI systems, leaving the full potential of human-machine teaming unrealized.

This article addresses the need to visually signify the uncertainty specific to LLM summaries, to guide that signification with an uncertainty framework, and to situate the resulting signification and explanation using interface design strategies. We do this by reporting on a 12-month design investigation that examines uncertainty in LLM summaries for the intelligence community, funded by the United States Department of Defense and focused on language analysis — analysis where source information, or *intelligence traffic,* primarily takes the form of transcribed or written language. This investigation was a collaboration with the Laboratory for Analytic Sciences (LAS), which supports the advancement of technology and tradecraft relevant to the

mission of the United States intelligence community. Among other research areas, LAS focuses on human-machine teaming, particularly the optimal utilization of automated systems by intelligence analysts. This high-stakes, security-critical space demands exceptional precision, making it an ideal test environment for AI explainability and LLM uncertainty.

The investigation began as an exploration of uncertainty visualization and expanded to encompass a proposal for an interactive system with multiple LLM agents that assist language analysts in summary validation. The resulting validation process enables critical trust calibration between analysts and the AI system. It is directly relevant to LLM use by all knowledge workers beyond the intelligence context, and further, to any situation that is dependent upon LLM summarization for decision-making. The general contributions of this investigation, as articulated in this article, are:

1. An open resource of implementable visual conventions for representing uncertainty, with criteria for selecting among them (Section 4).
2. A framework for uncertainty in LLM summaries (Section 5).
3. Design specifications for a Multiple Agent Validation System (MAVS) to empower knowledge workers while helping them calibrate trust in AI, including a 10-element feature set (Section 6).
4. Rich prototypes for alternative versions of MAVS. The prototypes include a simulation interface that allows users to utilize MAVS in a scripted scenario (Section 7.1) and three narrative interfaces that enrich the MAVS concept by envisioning dynamically reconfigurable user interface implementations (Section 7.2).

Further special contributions for the design community are:

5. The investigation's overall process, which is a model that can readily be adapted for any design investigations that involve collaboration with non-design experts and that center visual exploration in their activities (Section 3).
6. Similarly adaptable processes for developing a framework (Section 5.1), producing a range of visual studies (Section 4), and designing speculative interfaces with a STEM focus on implementation rather than as a mode of criticism (Section 7).

Finally, special contributions for other communities of practice — medicine and climate science communities — are:

7. Explicit declarations regarding how this work may be applied in medical diagnosis and climate forecasting, as examples of application beyond the intelligence analysis space (Section 8).

2. Literature Review

Language analysts in the intelligence community have only recently begun exploring how they might utilize LLM summaries in their workflow, shifting from direct database queries to increasingly relying on more opaque LLM processes of retrieval, analysis, and summarization. This emergent human-machine teaming can augment human cognitive abilities and leverage human and AI capabilities. AI-assisted human decision-making has the potential to outperform full automation in the national security sector, as in other critical sectors such as medicine, law, financial services, and law enforcement (Tomsett et al., 2020; Zhang et al., 2020). However, the probabilistic nature of AI models — leading to fundamental levels of uncertainty — necessitates human oversight in critical scenarios. The uncertainty specific to LLMs makes it difficult for analysts to gauge information reliability. When users cannot fully grasp how automated systems work — particularly in complex scenarios where comprehensive technical knowledge is impractical — their willingness to rely on the systems depends heavily on trust. If uncertainty and vulnerability were not factors, trust would not be necessary (Lee & See, 2004).

Establishing and maintaining human trust of AI systems is quite challenging in a high-stakes environment. Inappropriate levels of trust between users and AI systems often lead to misuse (as overreliance) or disuse (as underreliance) of automation (Lee & See, 2004). Dietvorst et al. (2016) frame underreliance as *algorithm aversion*. Humans tend to trust algorithms until they detect that they are imperfect — and all algorithms are imperfect — at which point they may avoid the algorithms, bypassing them. Alternatively, if an interface gives users some control over an AI's predictions — and it can be very limited control — humans will have a greater tendency to use the AI, or to overcome algorithm aversion (Dietvorst et al., 2016). Whether due to overreliance or underreliance, miscalibrated trust diminishes the benefits of using an AI system. Overtrusting is particularly problematic when humans reinforce their own negative societal biases (Stevenson, 2018; Suresh et al., 2020). The AI system can become a convenient excuse for problematic recommendations rather than augmenting and improving human decision-making. Lee and See (2004) have noted that the diminishment of trust through system misuse and disuse is a closed-loop process: "If the system is not trusted, it is unlikely to be used; if it is not used, the operator may have limited information regarding its capabilities, and it is unlikely that trust will grow" (p. 68). So how might we interrupt the loop of diminishing trust?

User interface design is one possible answer to the question of trust calibration. Appropriate trust calibration occurs when a user's trust in an automated system corresponds accurately with the system's capabilities (Lee & See, 2004). Achieving appropriate trust calibration can produce superior human-machine performance

(Sorkin & Woods, 1985; Wickens et al., 2000). Specific interface features have been proposed that might support the appropriate calibration of trust (e.g., Corritore et al., 2003; Cummings, 2006). Borgo et al. (2024) synthesized 40 related papers to provide nine claims about the impact of interface design choices on perceived trustworthiness. Four of these claims are particularly relevant here:

1. "AI transparency, intelligibility, or explainability fosters trust,"
2. "Communicating uncertainty fosters trust,"
3. "Adding interactivity fosters trust," and
4. "Social factors influence trust" (Borgo et al., 2024, pp. 23–24).

AI transparency, intelligibility, or explainability fosters trust. Interface features that empower users to verify results through access to and investigation of original sources produce transparency for otherwise opaque AI models (Borgo et al., 2024; Sultanum et al., 2019), positively impacting trust (Dasgupta et al., 2017; Krueger et al., 2020; Sperrle et al., 2021). Sultanum et al. (2019) explain that linking back to original sources enables users to analyze and compare source material with the LLM output. The resulting analysis offers insight into the LLM's process, which helps the user more clearly delineate the boundaries of system capabilities and understand its outputs. Methods of delineation have frequently been addressed in the literature, although researchers disagree on the details (Bansal et al., 2019; Bellotti et al., 2001; Borgo et al., 2024; Doshi-Velez & Kim, 2017; Sultanum et al., 2019; Tintarev & Masthoff, 2007; Weisz et al., 2024).

Communicating uncertainty fosters trust. Many researchers point to the role of uncertainty awareness in trust calibration (Amershi et al., 2019; Bansal et al., 2019; Kocielnik et al., 2019; Tomsett et al., 2020). A user's trust in a system correlates directly with how well the user understands its underlying uncertainties (Sacha et al., p. 76). An important factor for such understanding is user comprehension of what a system does not or cannot know (Tomsett et al., 2020). Once aware of output uncertainties the user can more quickly form an accurate mental model of the system's true capabilities (Tomsett et al., 2020, p. 2). Borgo et al. (2024) suggest that the user interface should clearly display the uncertainties and limitations inherent in a system's data and results. Essential information about uncertainty should be prioritized to address human cognitive limitations (Alhadad et al., 2018; Baldassi et al., 2006), and the design of the interface should carefully direct users' attention (Shneiderman, 1996; Rosenholtz et al., 2007).

Padilla et al. (2018) consider practical strategies for reducing cognitive load during decision-making with visualizations. They recommend that designers focus on prioritizing and hierarchically structuring information: "Identify the critical information needed for a task and use a visual encoding technique that directs participants' attention to this information" (p. 22). Alhadad et al. (2018) suggest several strategies for directing attention and reducing cognitive strain through visualizations, including coherence,

chunking, contiguity, segmenting, and signaling. These recommendations point to the role visual and interaction design play in focusing the user on vital information needed to make decisions.

Dual process theory (Evans & Stanovich, 2013; Padilla et al., 2018) breaks decision-making into two types of processing: humans first make simple, lightweight decisions (Type 1 processing) before moving on to complex, demanding, laborious decisions (Type 2 processing). The dual process approach aligns with established user experience (UX) design principles, such as progressive disclosure through layered interfaces, the principle that UX should reveal increasingly complex data to users in stages or layers (Forsey et al., 2024; Joshi et al., 2017). Designers should leverage such strategies to support both Type 1 and Type 2 processing to impact trust calibration through sustained interaction. Simply visualizing uncertainty is not enough.

Adding interactivity fosters trust. Borgo et al. (2024) also emphasize that interactive features can build trust by enabling users to test and verify system behavior; to customize outputs to better serve their needs; and to contribute domain expertise to improve performance. Hands-on interaction helps users understand a system's capabilities and limitations, allowing them to better predict its behavior across different scenarios (p. 26). Accurately predicting system behavior across scenarios is key to successful trust calibration. A user's ability to predict such behavior affects their own tendencies to either engage or disengage with AI.

Social factors influence trust. Personifying an AI system as a virtual agent can foster trust (Weisz et al., 2024), particularly when the interface combines modalities such as speech, voice, and visual presence (Rheu et al., 2021). Nass and Brave (2005) argue that humans instinctively process artificial voices like human ones — a natural, social response that makes voice interfaces effective tools for building trust when designed to mimic human interaction patterns. Graaf and Malle (2017) showed that virtual agents provide an effective avenue for fostering user trust because users attribute human-like intentions and reasoning to AI systems. To fulfill their potential, virtual agents must be able to explain system actions, or else systems will remain opaque to users and mistrust is likely to develop (Williams et al., 2015). If virtual agents offer insightful explanations and exhibit human behavioral and linguistic patterns, users of AI systems can more easily form accurate mental models to guide system use and decision-making.

The literature thus provides guidance for addressing trust calibration. It supports these three key points:

1. Human-machine teaming can produce better results in critical decision-making spaces than human or AI working alone (Tomsett et al., 2020; Zhang et al., 2020; Zhao et al., 2023).

2. Appropriate trust calibration is key to successful human-machine teaming (Lee & See, 2004; Sorkin & Woods, 1985; Stevenson, 2018; Suresh et al., 2020; Wickens et al., 2000).
3. Interface design can be used to effectively communicate AI capabilities and thus support trust calibration (Borgo et al., 2024; Corritore et al., 2003; Cummings, 2006).

Furthermore, the literature suggests fundamental interface design strategies for facilitating trust calibration (TC):

- *TC1: Transparency.* Support trust calibration by enabling users to verify and interrogate LLM outputs with interface features such as direct source investigation (Bellotti et al, 2001; Borgo et al., 2024; Dasgupta et al., 2017; Doshi-Velez & Kim, 2017; Krueger et al., 2020; Sperrle et al., 2021; Sultanum et al., 2019; Tintarev & Masthoff, 2007; Weisz et al., 2024).
- *TC2: Visualization.* Support trust calibration by visualizing uncertainty to communicate limitations inherent to the LLM's data and results (Amershi et al., 2019; Banshal et al., 2019; Borgo et al., 2024; Kocielnik et al., 2019; Sacha et al., 2015; Tomsett et al., 2020).
- *TC3: Alignment.* Support trust calibration by enabling users to interact with visualizations in alignment with human decision-making processes (Alhadad et al., 2018; Baldassi et al., 2006; Evans & Stanovich, 2013; Kirschner et al., 2011; Padilla et al., 2018; Rosenholtz et al., 2007; Shneiderman, 1996).
- *TC4: Interactivity.* Support trust calibration by enabling users to affect system results through a range of explicit user interactions, including but not limited to user settings (Borgo et al., 2024; Dietvorst, 2016; Lee & See, 2004).
- *TC5: Virtual Agents.* Support trust calibration by enabling users to conceptualize model functionality and seek explanation through multiple AI agents (Borgo et al., 2024; Graaf & Malle, 2017; Nass & Brave, 2005; Rheu et al., 2021; Weisz et al., 2025; Williams et al., 2015).

We will revisit TC1–TC5 as we consider specific user interface features for representing and explaining uncertainty in LLM summaries.

3. Investigation Process

Early in this investigation we focused on developing a wide range of visual conventions for representing uncertainty in LLM summaries. This early exploratory work soon shifted to a more convergent, evaluative phase in which the extended multidisciplinary collaborative team — language analysts, computer scientists, and psychologists informing design researchers — narrowed the options down from approximately 150 to eight. In

parallel with these visual explorations, we began adapting an existing framework for uncertainty (Skeels et al., 2010) but ultimately developed a new framework for classifying uncertainty specific to LLM summaries.

In a parallel phase of the project, we considered how language analysts might interact with visualizations of uncertainty. We realized quickly that if analysts did not appropriately trust the information represented by visualizations, they would not use that information — in which case representational quality would be irrelevant. To address appropriate trust calibration, we drew from UX findings established in an earlier project with the Laboratory for Analytic Sciences (LAS, 2024), as well as the trust in automation literature. Through these efforts we gleaned five trust calibration interface strategies appropriate to intelligence analysis (Section 2).

Using a persona, scenarios, task flows, and interface strategies as a starting point, we decided to develop an interactive simulation interface, both as a concept generator for ourselves and as an educational tool for language analysts who use AI in their tradecraft. We pinpointed 10 core system features to address trust calibration within this simulation and cohered them into an LLM validation system concept. While the interactivity of the simulation enables analysts to directly experience the proposed system in full interaction fidelity, the requisite development time limited our own formative design exploration of the proposed system's potential. To overcome this limitation, we pivoted to additional scenario video prototyping. The three resulting narrative interfaces involved no backend development, resulting in a more nimble iterative process that reflected our evolving understanding of uncertainty and trust. These narrative interfaces envision future possibilities for trust calibration in dynamically reconfigurable user interfaces. This dual method of developing key interface features within both current and future interface structures allowed us to pivot in response to expert assessments, while permitting the lateral movements typical of design exploration.

The orderliness of this description and this section title's implicit suggestion of a singular process are both potentially misleading. As shown in Figure 1, we engaged in a 12-month discovery process that, through sustained interaction with our partners outside of design and through our own sensemaking-through-design, was neither orderly nor predictable. The project coordination strip in Figure 1 lists four key planning moments for our project. It was not until we were past the quarter mark that it became possible to create a "plan-out," the first plan that envisioned the project through completion. We knew from experience that early efforts would need to be fulfilled before later efforts could be specified. These plans are visualized in Figure 1 as ripples that affected the main investigation outcomes because each plan synthesized collective sensemaking and suggested adjustments to all ongoing tasks.

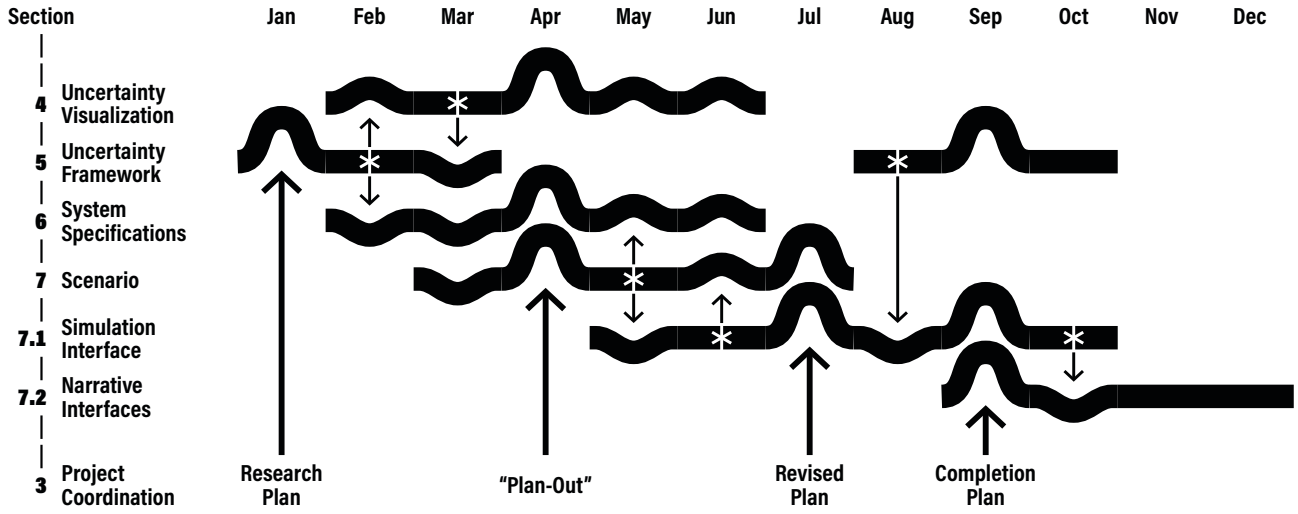

Figure 1. Investigation process for the project "Developing Visual Conventions for Explainable LLM Outputs in Intelligence Analysis Summaries," conducted January ("Jan") through December. Section numbers are references internal to this article. Ripples represent impacts spreading from one investigation component to others.

Likewise, our work toward individual outcomes frequently caused ripples across parallel tasks. Figure 1 depicts these ripples as smaller than those created by the comprehensive plans, but they were crucial nonetheless. As depicted, five formative investigation contributions caused ripples:

1. Uncertainty visualization (Section 4)
2. The uncertainty framework (Section 5)
3. LLM validation system specifications (Section 6)
4. A scenario encapsulated in the prototypes (Section 7)
5. The simulation interface (Section 7.1)

A sixth contribution was summative, informed by all previous work:

6. The narrative interfaces (Section 7.2)

It was thus through an immersive and messy process that this investigation took form. We now address the main investigation contributions in turn.

4. Uncertainty Visualization

To discover visual conventions for uncertainty in LLM summaries, we utilized design exploration and co-design, sharing progress in biweekly sessions with the extended multidisciplinary collaborative team. This meeting structure encouraged team participation regarding the formal qualities of the visual studies, their effectiveness in

conveying uncertainty, and the broader implications for intelligence analysis workflows and decision-making processes. We gathered feedback through interviews, surveys, and informal user testing. To ensure flexibility for incorporating visual conventions into summarized LLM outputs, we organized exploratory studies into three uncertainty cue locations: *inline,* embedded directly within the summary text itself; *interstitial,* positioned in the spaces between lines of text; and *adjacent,* appearing outside of the summary. To maximize variation, we did not initially concern ourselves with practicality, but we later removed all impractical studies from consideration. We created what we called "concept zero," a simplified interface prototype for situating visual conventions. In the midst of our exploration we adopted an accessible color palette from the IBM Design Language to address Web Content Accessibility Guidelines. Ultimately, we generated approximately 150 visual studies for conveying uncertainty in LLM summaries.

Discussions with our collaborators also produced criteria to guide future work that seeks to establish a singular visual convention for representing uncertainty. Using these criteria, researchers could stage empirical studies relating visual conventions to mental models of AI and to user preferences.

- *Experiential.* Visualizations of uncertainty should provide a sense of severity at a glance, with text subjectively *feeling* as uncertain as it has been deemed to be.
- *Reflective.* Visualizations of uncertainty should make sense upon reflection, ideally aligning with an accurate or at least useful mental model of uncertainty. This accuracy or utility will increase educational impact.
- *Legible.* Uncertain text should be readable so that it can be validated by users, though legibility can be dynamically variable if users are permitted to inspect text passages, in which case a responsive system can make text clearer.
- *Implementable.* The display of visualizations needs to be technically feasible. We have favored common web display technology as a gauge of implementation feasibility, including considerations of which visual conventions utilize prescribed display functions (e.g., blur in CSS) and which require workarounds (e.g., background images to approximate inline display functions that do not exist in CSS).

Using these four criteria, we derived 12 options from the collection of visual studies, each of which we implemented in CSS (in the simulation interface, Section 7.1). In consultation with our collaborators, we deactivated four of these options for a total of eight potential visual conventions for uncertainty. These are reproduced in Figure 2. We evaluated these visual conventions according to the semiotic modes by which they operate, the conceptual implications of those modes, and the patterns of criteria

fulfillment across the set. Semiotic modes and conceptual implications are documented in Table 1.

- *Semiotic modes.* Four visual conventions (VCs) utilize analogies exclusively (VC1) or primarily (VC6–CV8). This is potentially powerful because analogies reveal structural similarities between two domains that reflect true characteristics (Gentner & Smith, 2012). Three VCs utilize metaphors (VC2–VC4), which can be coherent and improve understanding, while being dependent upon familiarity with a suggested source domain (Johnson, 1987; Lakoff & Johnson, 1980). But VC5

Figure 2. Eight potential visual conventions for representing uncertainty severity in LLM summaries.

Table 1. Semiotic modes and conceptual implications for the eight uncertainty-signifying visual conventions shown in Figure 2.

Visual convention	Evaluation: mode and implication
VC1: Strikethrough	**Mode:** response analogy, uncertain passages should be editorially rejected (thin line) or redacted (thick line). **Implication:** uncertainty is a mistake.
VC2: Transparency	**Mode:** disappearance metaphor, uncertain passages are fading out of existence. **Implication:** certainty is tangibility.
VC3: Static	**Mode:** television-static metaphor, uncertainty makes passages difficult to resolve. **Implication:** uncertainty is interference in a passage-signal.
VC4: Fill	**Mode:** fluid-volume metaphor based on an up-is-more orientational metaphor. **Implication:** uncertainty is a quantity in passages.
VC5: Pattern	**Mode:** primarily arbitrary symbolic representation (pattern style); secondarily density analogy (pattern repetition interval). **Implication:** uncertainty is a pattern — this is meaning-poor — or uncertainty is a quantity in passages.
VC6: Text Blur	**Mode:** visual perceptual analogy, certain passages are focused on (attended to) instead of uncertain passages. **Implication:** certainty is in focus or is comfortably accessible.
VC7: Zig-Zag	**Mode:** primarily editorial markup analogy; secondarily wave frequency metaphor. **Implication:** uncertainty passages are unfinished, or uncertainty is a force.
VC8: Weight	**Mode:** primarily mass analogy; secondarily arbitrary relative representation (color). **Implication:** uncertainty is conspicuous — this is meaning-poor.

uses arbitrary distinctions of patterns, which likely limits its efficacy, requiring viewers to learn symbols.

- *Conceptual implications.* Two VCs conceptualize *certainty as a thing* and *uncertainty as a reduction of that thing* (VC2, VC6). We suspect that this is a healthy way to view uncertainty. The remainder conceptualize *uncertainty as a thing,* which may be useful because uncertainty is what language analysts must interrogate. Two VCs are listed as being "meaning-poor" in Table 1 because their semiotic modes as executed do not suggest an obvious mental model, in our estimation (VC5, VC8).

- *Experiential criterion.* The two VCs that use graphic mark variability to differentiate degrees of severity — background patterns (VC5) and wavy underlines (VC7) — do not appear to give an immediate impression of uncertainty. The VC that depends on font weight variation (VC8) is possibly too subtle for viewers without graphic design expertise. The remaining VCs all appear to have some claim to the experiential.
- *Reflective criterion.* Our difficulty in describing the conceptual implications of the VCs we labeled as "meaning-poor" (VC5, VC8) would likely equate with minimal contributions to viewers' understanding of uncertainty. The remaining VCs appear to have some claim to the reflective.
- *Legible criterion.* Legibility was largely assured through hover states that offer a relatively unobstructed view of otherwise obscured text. Three VCs render passages of "obvious" severity as entirely or nearly unreadable before hovering (VC1, VC3, VC6), while two VCs effectively leave all text unobscured at all times (VC7, VC8). The VCs tend to interact with qualities that may typically be adjustable through accessibility settings, which is a complicating factor.
- *Implementable criterion.* Three VCs utilize normal web display settings without the need for background images or pseudo-class workarounds that are not compatible across browsers (VC2, VC6, VC8), while the remainder require background images or workarounds to implement. VC1 is a special case. Though strikethrough is readily available in HTML and in more rudimentary applications, it is not a customizable property in CSS. VC1 would be the most implementable if only one level of severity was to be signaled, but indicating multiple levels of severity with strikethrough currently requires a workaround.

The eight visual conventions have been implemented in a web-based simulation with readable code. This open resource enables others to reevaluate and even modify them, thus discovering additional strengths and weaknesses in the visual conventions.

5. Uncertainty Framework

5.1. Uncertainty Framework Development

One of the first tasks we undertook in our investigation was to define uncertainty in relation to LLM summaries for intelligence analysis. In our search of existing literature in early 2024, we found no uncertainty frameworks specific to LLM summaries for intelligence analysis. However, we did locate relevant research on trust in AI and automated decision aids (Zerilli et al., 2022; Fell et al., 2020; Heger et al., 2016; Manzey et al., 2012; Okamura & Yamada, 2020; Prabhudesai et al., 2023; Vaswani et al., 2017), as well as on uncertainty information and explanation visualization in other contexts (Karran et al., 2005; Skeels et al., 2010; Thomson et al., 2005).

Among these, Skeels et al. (2010) provided the most useful initial framework, having structured classification of uncertainty with what we believed to be sufficient range and granularity for visual exploration guidance. The framework identifies five types of uncertainty: completeness, credibility, disagreement, inference, and measurement precision. Skeels et al. (2010) isolated these types through an analysis of uncertainty across many scientific fields, including ecology, computational biology, and medicine, with a focus on improving information visualization — as in diagrams, not in the annotation of text as is required in language analysis.

Our revision of, and ultimate departure from, Skeels et al.'s (2010) framework was informed through biweekly conversations with collaborating experts in language analysis, computer science, and psychology. We asked them to speculate on what kinds of uncertainty might be associated with LLMs, and we tried to map their suggestions to Skeels et al.'s (2010) types. We began to remove types, add types, and rename types. We repeatedly needed to pull back to determine *where* exactly we were attempting to identify uncertainty — out in the world, in the sources, or in the summary itself? We ultimately decided to limit our investigation to the uncertainty types that might appear in the summary itself due to the probabilistic nature of LLM technology.

5.2. Overview of Uncertainty Types

Our LLM-oriented uncertainty framework is a contribution to research in the context of intelligence analysis, knowledge work, and interface design. Though it is not the result of a systematic study like Skeels et al.'s (2010) general uncertainty framework, a provisional framework provides the requisite *a priori* structure for subsequent empirical studies that could validate it or suggest adjustments. It is a necessary first step.

We call our framework the Uncertainty Framework for Explainable Summaries (UFES). We do not specify the intelligence context or language analysis in the title as we have adopted definitions that we believe more fundamentally address LLM summaries. There are five types of uncertainty in UFES, and they are defined as follows.

1. *Meaning uncertainty:* misinterpreting word sense for technical, cultural, or uncommon terms, or for jargon.
2. *Reference uncertainty:* mistaking associations from demonstratives ("those"), adverbs ("there"), definite articles ("the"), or pronouns ("they").
3. *Conjecture uncertainty:* jumping to conclusions, incorrectly completing partial information, or making assumptions.
4. *Credibility uncertainty:* trusting a statement that was unserious, humorous, incongruous, a *non sequitur,* a manipulation, or an apparent lie.
5. *Evidence uncertainty:* making a claim without supporting evidence, either drawing from opaque training data or by hallucination.

Table 2. An imagined surreptitious recording between fictional characters André Silva and Baaba Owusu, and concerning fictional countries Avalon and Oceania. Line 4 includes a break in recording or transcription.

Line	Transcription
1	**Silva:** He didn't like having to give that speech as [indecipherable] those rural teachers. They aren't going to see the big picture
2	**Owusu:** There is no big picture. It's just what you say to them, and what you say to the big public schools.
3	**Silva:** Look, it was all about timing. And that's gonna be on the cycle for a couple of days, he gets to duck out. Look at what we have coming up. It's manufacturing on the border, customs, coordination, the union bosses on both sides. The Avalon secretary
4	[indecipherable audio, 5′23″]
5	[unidentified]: and he plans to make an ultimatum, quietly.
6	**Owusu:** But that secret won't keep.
7	**Silva:** Sure but
8	**Owusu:** But Oceania will know what it means. They're ready to act on it. It's only days

As suggested earlier, UFES primarily refers to uncertainty *in summaries themselves,* not uncertainty in source information, which is irreducible when that information is isolated. For instance, if a query addresses the possibility of life on Mars, and a summary states that there is no *evidence* of life on Mars, it is not a case of evidence uncertainty — the stated lack of evidence is an accurate accounting of a knowledge reality. Essentially, UFES is concerned with misleading claims in summaries, and it serves to identify the manner in which a claim may be misleading.

Our thinking here was crystalized in a summary example that utilizes three fictional countries — *Kobian* administration officials discuss *Avalon* and *Oceania.* This example is provided in Table 2. The following summary claim is an accurate representation of the source material displayed in that table, in which "[indecipherable audio, 5′23″]" separates a pronoun from its thus unknown referent. An imagined summary provided by an LLM that consults the Table 2 source includes:

> There is mention of an "ultimatum," but an apparent gap in the source recording makes it impossible to determine who is making an ultimatum that will serve as a communication to Oceania, or if this is a serious ultimatum and the degree to which it is mission relevant.

Though there is obviously irreducible uncertainty in the imagined source material — whatever was said in the missing five minutes — the summary's claim itself is not uncertain.

UFES could inform development of an LLM validation system by providing a starter language from which to articulate soft prompts. But the more immediate use of UFES is to improve human understanding of uncertainty in the LLM outputs. We now provide more detail on the five types of uncertainty we have defined here.

5.3. Additional Detail on Uncertainty Types in UFES

Meaning and reference uncertainty. Meaning and reference uncertainty emerged late in our development process, and only when we consciously confronted the mismatch between Skeels et al.'s (2010) numeric focus and our language focus. As such, these types have no corollary in their framework. We intend meaning uncertainty to be more localized, at the level of individual terms, and reference uncertainty to be distributed, as arising from relationships between statements. We considered "denotative" and "indexical" as alternative names for these types of uncertainty, respectively, because both deal with meaning derivation. But we opted for a more colloquial or less technical terminology.

Reference uncertainty begs further explanation than its definition alone. Returning to the example of Table 2, consider a different LLM summary claim to the one used above:

> …The Secretary of Labor appears to be preparing to make an ultimatum regarding the manufacturing conflict…

Cross-referencing with Table 2, the gap in transcription noted in line 4 occurs between mention of "the Avalon secretary" and the "he" who is making an ultimatum. It is a questionable assumption that "he" is the secretary when there is a gap of over five minutes. Even attribution of "labor" to the secretary is questionable. It appears to be an assumption following "union bosses" in the sentence immediately preceding "the Avalon secretary." These represent two degrees of reference uncertainty embedded in the claim that *the Secretary of Labor* is *making an ultimatum.*

Conjecture uncertainty. Conjecture uncertainty bears some similarity to Skeels et al.'s (2010) inference uncertainty. In early stages of developing UFES, we had a broader definition of inference uncertainty. The difference is apparent in our notes:

> An assertion is in some manner ambiguous, with more than one possible meaning available to complete it. An error occurs when a claim about the assertion relies on a misinterpretation (in cases of logically resolvable ambiguity), or otherwise fails to acknowledge the ambiguity inherent to the assertion (in cases of irreducible ambiguity). Misinterpretations and ambiguity can be rooted in: idiomatic

> expressions; cultural nuances; context-specific phrases; polysemy; unconventional sentence structure; conjugation; pronouns; rare or uncommon terms; technical jargon; tone; humor; or sarcasm.

This transitional definition is so general that it embodies our revised sense of three types of uncertainty: meaning ("rare or uncommon terms"), reference ("pronouns"), and credibility ("sarcasm"). Basically, the transitional definition is too broad to be useful, and it is emblematic of the verbal gymnastics that were necessary to conform linguistic considerations to Skeels et al.'s (2010) numeric considerations. Ultimately, in addition to refining our own definition of a related kind of uncertainty, we opted to designate it *conjecture* to avoid a mismatched comparison with Skeels et al.'s *inference.*

Nevertheless, our conjecture uncertainty is admittedly difficult to isolate. To retain the focus on LLM summary generation, we emphasize that it refers to completing partial information without adequately acknowledging the completion act. Conjecture uncertainty occurs when an alternative summary claim could have been drawn from the same partial information. It is an assumption. This does embody aspects of Skeels et al.'s (2010) completeness uncertainty, but missing numeric values in a data set are far more conspicuous, and far less ambiguous, than partial linguistic information. The following two examples may help to clarify conjecture uncertainty further.

First, imagine a scenario where there exists copious intelligence traffic about a series of meetings between an adversarial country's president and a group of legislators on a particular issue, yet none of that traffic offers specific details about the meetings themselves. Instead, what is available are conversations among those legislators that occurred after the meetings, in which they complain about interpersonal dynamics and personal agendas (e.g., getting the chief of staff to admit that he is wrong about *anything*). If a summary characterized these meetings as the president strategizing against the legislators based solely on such traffic, it would have conjecture uncertainty. It is an assumption that the expressed feelings following the meetings transfer fully to the meeting's agenda.

Second, imagine another scenario where available traffic inconsistently presents a fictional president's views on anti-ballistic missile deployment. In two sources, he appears strongly and moderately for increased deployment, and in two other sources he appears strongly and moderately against it. Taken at face value, it could appear that the president has no strategy, or that he is obfuscating, and thus a summary may characterize his views on the matter as "suggesting an absent or clandestine strategy." However, if a closer look at the sources reveals that his statements (or insider statements about his views) were made in confidence and with conviction, the "absent" or "clandestine" claim has conjecture uncertainty — it is a step too far in assumption. (We have played this scenario out further, where ordering the sources by date reveals that the president's

views evolved over time — as in, the sources were all accurately capturing moments in a sequence.)

Credibility uncertainty. For credibility uncertainty alone we retain Skeels et al.'s (2010) name for a type. We did so because, though the original type does indeed focus on numerics, its concept of credibility transfers to linguistics in ways we do not find distorting. While we do expand the type's definition, we do not feel the need to replace any major aspect of it.

The most straightforward interpretation of credibility uncertainty in LLM outputs is trusting the words of somebody inherently untrustworthy, and our definition does account for this. Skeels et al. (2010) note that a "human source may be considered untrustworthy based on past behavior or associations" (p. 76). Credibility uncertainty can also be more contextual. Some people may be inherently more credible than others based on their expertise and believing certain statements from inherently nonexpert sources will carry a degree of credibility uncertainty. Skeels et al. (2010) account for this as well: "…information from a specialist may lead to less uncertainty than information from a generalist…" (p. 76).

We additionally consider situational factors. The situation in which a person makes a statement is an everpresent complication, whether that person is generally trustworthy or untrustworthy. If an expert or an insider is making a joke, they are not necessarily leveraging their beliefs, and thus their access to relevant knowledge does not validate the joke's implications. And much speech is rhetorical, aimed at convincing others through argumentation more than through the explication of truth, even in casual conversation. In practice, there are not purely trustworthy or untrustworthy people. Instead there are situations in which individual statements may or may not be credible. Thus, credibility uncertainty refers to assertions themselves and not the people who make them. An assertion may be less credible due to its speaker's identity, but it is the assertion itself that has credibility value in our framework — the statement is the thing.

Evidence uncertainty. When objective truth and complete understanding are not realistic goals, what does it mean to have or to lack "evidence"? The other four types of uncertainty in UFES all concern the interpretation of available information. Clearly a summary can include claims that are more thoroughly unfounded than they are uncertain in meaning, reference, conjecture, or credibility.

In LLM summary generation powered by retrieval-augmented generation (RAG), an LLM has documentable access to its summary (of course) and the RAG sources, but not to the training data that constitutes its underlying foundational model. Evidence uncertainty occurs when there is a discrepancy between the *summary* and *documentable sources.* A claim is made in the summary for which no form of evidence is available.

It is impossible to know why an LLM made a claim if no evidence is given. If there is no basis for the claim (besides perhaps a biasing query), it is a case of hallucination. But this is indistinguishable from two other possibilities: that the claim came from the black-box training data; or that there was some kind of failure in documenting the normally documentable sources. What is most immediate in an intelligence analysis context is that a claim has no supporting evidence. This is what determines a course of action for a language analyst or for other users.

Pulling out from an intelligence analysis context, this view of evidence uncertainty still has efficacy. Whatever the mechanism is — hallucination, training data, source disclosure error, or something else — a claim in a summary that cannot in any way be supported cannot be validated.

6. Specifications for an LLM Validation System

This section outlines the design specifications for a Multiple Agent Validation System, contributing to the intelligence community and any other cases involving decision-making and sensemaking with LLM summaries. These specifications are the product of all threads of the investigation and of sustained interaction with our multidisciplinary collaborators. As our exploration of uncertainty visualization increasingly raised issues of interface design, we asked language analysts what actions they might take when they encountered uncertainty in an LLM summary. They identified seven specific actions: (1) viewing source files, (2) asking the LLM about its sources, (3) assessing relevance to the query, (4) asking the LLM about its summary, (5) submitting a new query, (6) modifying their existing query, and (7) searching for more information elsewhere. This and other feedback eventually coalesced into an LLM validation system concept.

6.1. The MAVS Concept

We propose a Multiple Agent Validation System (MAVS) to make knowledge workers more efficient while mitigating the limitations of LLM technology, and to facilitate healthy trust calibration by addressing common user struggles with automation. Our MAVS specifications include 10 discrete features, conceptually distributed among three virtual agents: a Query Agent, an Analytic Agent, and an Evaluative Agent. Whether or not these virtual agents are implemented in separate LLMs or as roles within a single LLM, they are instantiated in the feature set as distinct entities to aid the user in developing an accurate mental model — both of MAVS's underlying processes and of LLM technology more generally. Figure 3 is a process diagram of MAVS. Note that validation — the V in MAVS — is completed by the user.

We organize the feature set here according to the virtual agents. Table 3 lists how MAVS features address the trust calibration interface design strategies that emerged in the

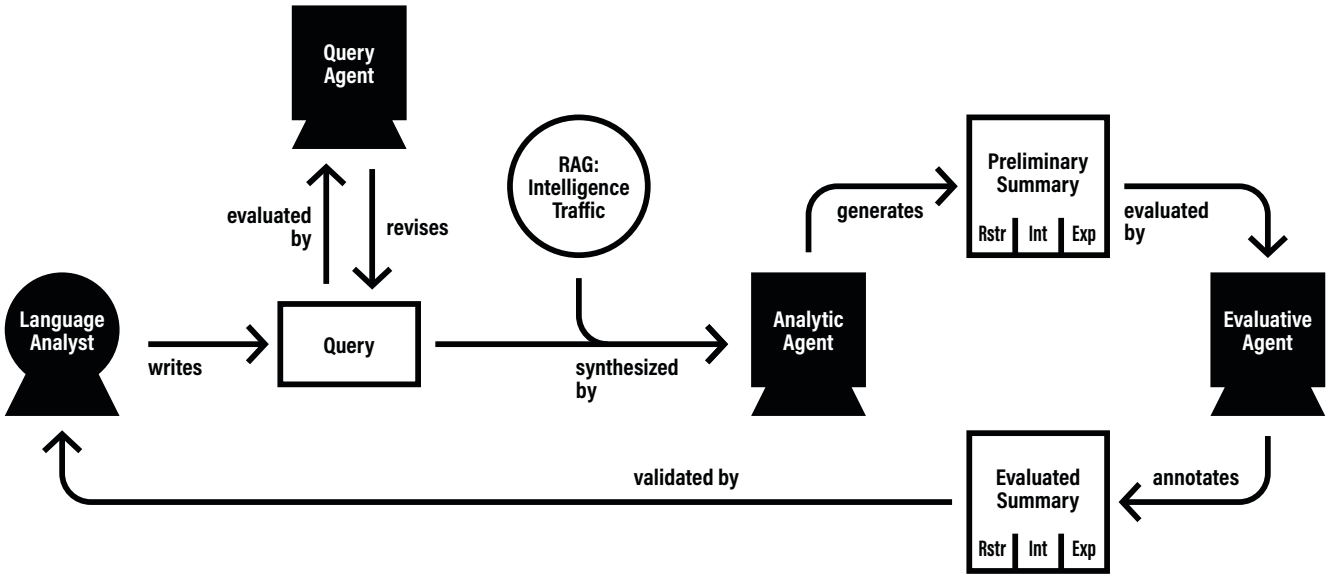

Figure 3. MAVS process diagram for the intelligence analysis context. Retrieval-augmented generation (RAG) focuses the system on intelligence traffic (i.e., collected sources). Summaries are available to users in three versions per the analytic sensitivity setting: restrictive ("rstr"), intermediate ("int"), and expansive ("exp").

Table 3. How the 10 core MAVS features address the trust calibration interface strategies (Section 2). The virtual agents are identified as QA (Query Agent), AA (Analytic Agent), and EA (Evaluative Agent).

Feature	Agent	Trust calibration interface strategies
Query History	QA	TC1: Transparency
Query Reshuffle	QA	TC4: Interactivity
Analytic Sensitivity	AA	TC4: Interactivity; TC5: Virtual Agents
Summary Sources	AA	TC1: Transparency
Visualization Sensitivity	EA	TC2: Visualization; TC3: Alignment; TC4: Interactivity; TC5: Virtual Agents
Uncertainty Visualization	EA	TC2: Visualization; TC3: Alignment
Uncertainty Alert Type Identification	EA	TC1: Transparency; TC2: Visualization; TC3: Alignment
Flagged Sources	EA	TC1: Transparency
Evaluative Agent Chat	EA	TC1: Transparency; TC4: Interactivity; TC5: Virtual Agents
Evaluation Export	EA	TC1: Transparency

literature review (Section 2, TC1–TC5). MAVS utilizes these key strategies to help users develop appropriate trust calibration with the goal of improving overall performance in human-machine collaboration. Figure 4 delineates areas within the simulation interface that are dedicated to the three virtual agents — the simulation interface is described in some detail in Section 7.1.

6.2. Query Agent Features

The Query Agent assists the user in the querying process. It has the lowest instantiation profile of the MAVS virtual agents.

Feature 1: Query history. The formality of the querying process, in which the user's investigation is comprehensively represented as text inputs and text outputs, affords a remarkably complete and accurate record of that investigation. The query history feature documents all user queries in a listing to which the user can return. The Query Agent dynamically generates short titles for queries, as commercial LLM products currently do. This listing indicates adjustments to queries resulting in distinct summaries with index counts of two and greater.

Feature 2: Query reshuffle. Prompting — writing queries — is a special skill and it can be done poorly, reducing or reversing the effectiveness of LLM summaries. With the query reshuffle feature, the user can request that the Query Agent analyze and improve their query, producing a new summary. This is an established capability for AI given the right training. A frequent outcome of query revision is debiasing, removing elements of queries that can push results in an inappropriate direction. For instance, the following query is a directive, not a question, which could effectively coax an LLM into confirming the query premise irrespective of the evidence: "Explain Nicolau's plan for anti-ballistic missile development and expansion." (Rysz Nicolau is president of the fictional country Kobia.) The directive assumes that Nicolau indeed has plans for such development and expansion. A debiased version of this prompt might be: "Does Nicolau have any plans for anti-ballistic missiles?" Query reshuffle modifies the query, which causes the Analytic Agent to generate a new summary, adding to the index count of the pre-shuffled query's listing in the query history.

6.3. Analytic Agent Features

The Analytic Agent responds to the user's query by drawing from a vast quantity of information and returning a summary. This is probably the most familiar role for an LLM.

Feature 3: Analytic sensitivity. The analytic sensitivity feature permits the user to influence how the Analytic Agent generates summaries with settings of *expansive, intermediate,* and *restrictive.* The expansive setting increases discovery by presenting

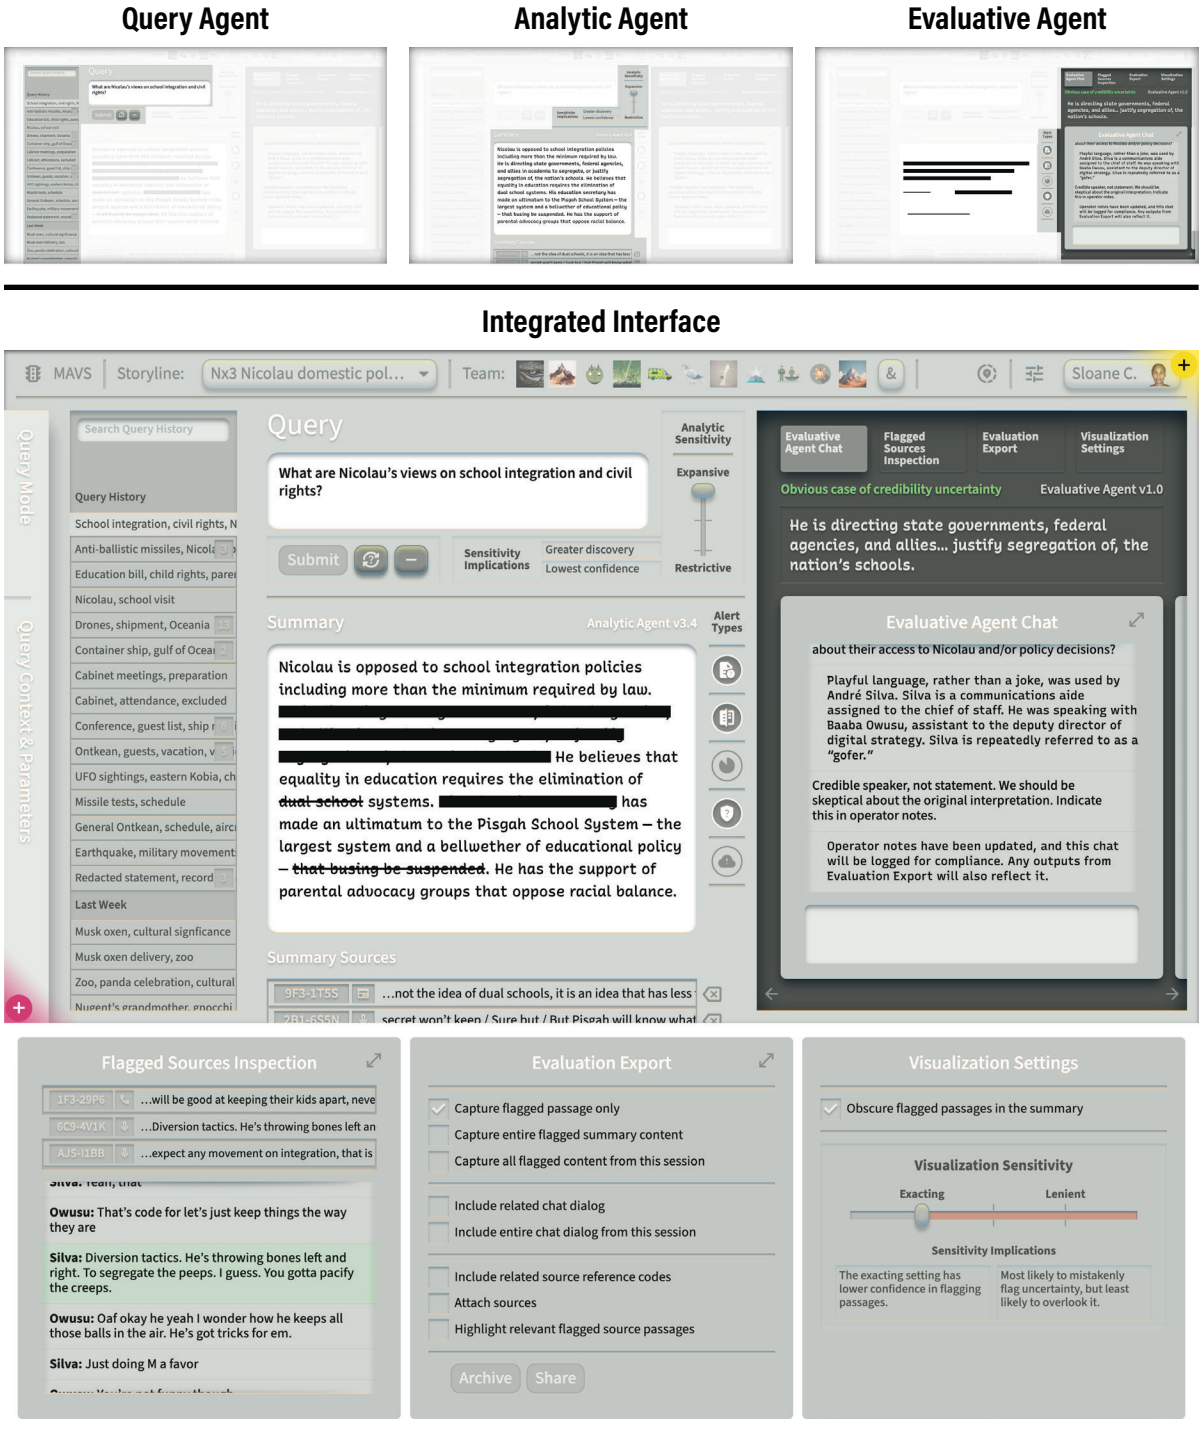

Figure 4. The three virtual agents in MAVS (top) as segmented in the full simulation interface (see Section 7.1). The three panels at bottom cycle through the area occupied by the Evaluative Agent chat at center right.

more possibilities to the user, but with a corresponding decrease in confidence. This may be appropriate in exploratory fact-finding. The restrictive setting results in greater confidence, but with lesser discovery that may overlook low-likelihood but potentially high-impact possibilities. This may be appropriate in crisis situations. As a setting that can be toggled, analytic sensitivity allows the user to see three versions of the Analytic Agent's summary for each query.

The analytic sensitivity feature could evolve iteratively through training and soft prompts. One possible method of implementation is for the Analytic Agent to independently generate *n* summaries, and then to compare those summaries. Claims that are shared across the highest proportion of independent summaries could be emphasized in a single common-claim summary — what is displayed with the restrictive setting. Reasonably strong claims that appear less frequently could be prioritized for a single uncommon-claim summary — some of what is displayed with the expansive setting. This method does beg the question as to how to keep the expansive summary to a reasonable length while still embodying some of the common claims, which should certainly not be ignored. Restrictive summaries may prove less useful due to the care taken to maximize reliability.

Feature 4: Summary sources. The ability to evaluate an LLM's claims is contingent upon access to the source material from which it reproduces patterns of language. Current LLM technology permits an accounting of this source material through retrieval-augmented generation (RAG). RAG enables an LLM to access a data set and to explicitly cite the sources of its claims in that data set — in contrast to the black-box behavior normally associated with LLMs. The summary sources feature discloses the specific sources the Analytic Agent used to generate a given summary. A relevant excerpt of the source is paired with metadata — in our case, a file code and an indication of recording medium — and the user can open the source directly to inspect it. The user can also remove a source from consideration, in which case the summary updates and the query history listing index count increases. Crucially in MAVS, the Evaluative Agent also has access to the disclosed summary sources, which is the basis of its validation process.

6.4. Evaluative Agent Features

The Evaluative Agent, the defining factor in MAVS, embodies an atypical role for an LLM agent. It utilizes specialized training and RAG to enact UFES (the uncertainty framework). The Evaluative Agent helps the user understand the Analytic Agent and validate its outputs. Ultimately, it evaluates congruence between the Analytic Agent's summary and its disclosed sources.

Feature 5: Visualization sensitivity. The visualization sensitivity feature permits the user to control when the Evaluative Agent bothers to provide markup, based on the

severity of uncertainty. Settings — *exacting, intermediate,* or *lenient* — characterize how the analyst asks the Evaluative Agent to behave. In attempting to recognize all instances of uncertainty in the summary, the exacting setting is the most likely to lead to mistakenly flagged passages. However, it is the least likely to overlook uncertainty — i.e., the most prone to false positives. The lenient setting results in more reliable uncertainty alerts. It is least likely to mistakenly flag instances of uncertainty, but most likely to overlook uncertainty — i.e., the most prone to false negatives. Unlike analytic sensitivity, toggling visualization sensitivity does not change the content of the summary. Instead, it flags more or fewer passages. The settings are operationalized according to plain language visible to the user: the lenient setting only flags *obvious* cases of uncertainty; the intermediate setting additionally flags *likely* cases; and the exacting setting additionally flags *conceivable* cases. There is no user option to bypass flagging obvious cases of uncertainty, as this would offer no initial means to validate the work of the Analytic Agent, hindering trust calibration.

Feature 6: Uncertainty visualization. There is an emotional component to uncertainty when there are professional stakes involved, and especially when there are security implications. In the flow of knowledge work, it is desirable that the user's emotional or gut sense of information is positively correlated with its certainty. The uncertainty visualization feature obscures passages in the summary commensurate with their assessed uncertainty severity. The levels of uncertainty controlled by visualization sensitivity — conceivable, likely, and obvious — are represented by increasing degrees of masking. A visual convention for achieving this may fully obscure obvious cases of uncertainty. However, cursor hover states for the summary itself permit the user to read flagged passages by responsively improving legibility.

Feature 7: Uncertainty alert type identification. To make accurate assessments when validating statements that have some degree of irreducible uncertainty, knowledge workers need to understand the basis of the uncertainty. We initially considered visualizing types of uncertainty instead of only severity level, but decided that this gives the user too much to learn and is distracting. Instead, the uncertainty alert type identification feature verbally identifies the type of uncertainty adjacent to the summary and only upon inspection. The Evaluative Agent identifies which of the five types of uncertainty is the reason a given passage was flagged, and concise definitions are provided for the types within the interface.

Feature 8: Flagged sources. An Evaluative Agent will be no more perfect than an Analytic Agent. Therefore a knowledge worker must be able to leverage their expertise to assess the Evaluative Agent's outcomes when validating the Analytic Agent's outcomes. The Evaluative Agent reports which of the Analytic Agent's disclosed sources led to an uncertainty alert. The flagged sources feature mimics the summary sources feature,

allowing the user to directly inspect sources in relation to an alert. An excerpt and metadata are immediately available, and the user can jump to a highlighted portion of the source information to begin the validation process, or they can open up the entire source.

Feature 9: Evaluative Agent chat. Providing users with a natural language mode of inquiry reduces the need to learn technically peculiar or unnatural interactions. A chat feature allows the user to engage in a conversation with the Evaluative Agent. When the user selects an individual flagged passage, the Evaluative Agent proactively explains why it was flagged. The user can engage the Evaluative Agent in conversation about the types of uncertainty, even if the agent does not tend to offer this information upon its initial description. This enables users to exert some control over LLM outputs, thus facilitating trust calibration.

Feature 10: Evaluation export. The more efficient augmented knowledge work becomes, the more difficult it will be for users to keep track of their investigations. The evaluation export feature documents the querying process for the user, along with all uncertainty alerts and interactions with the Evaluative Agent. In an intelligence analysis context, this is doubly important for compliance (i.e., aligning with strict regulations for reporting and documentation).

A robust implementation of MAVS would include these 10 features. Many of them can be experienced in the simulation interface (discussed in Section 7.1 and available in Peterson & Armstrong, 2024).

7. LLM Validation Prototypes

The LLM validation prototypes presented in this section contribute to knowledge work and the intelligence community in two important ways: they explicitly illustrate specifications that could be operationalized for practice (or for research in advance of practice); and they serve as educational tools that can help language analysts understand the potential of AI and set the stage for healthy trust calibration.

Scenarios situate design investigations within real-world contexts. In the scenario we developed for this investigation, a United States language analyst Sloane has been assigned to monitor and investigate the fictional country of Kobia and its president Rysz Nicolau. Our team employed this scenario within a simulation interface (in interactive demonstration form) and in three additional narrative interfaces (in video form). These prototypes are available in Peterson and Armstrong (2024). To most effectively demonstrate key MAVS functionality, we selected specific sections of the scenario to highlight within each prototype. The scenario content includes summaries, source excerpts, and chat scripts.

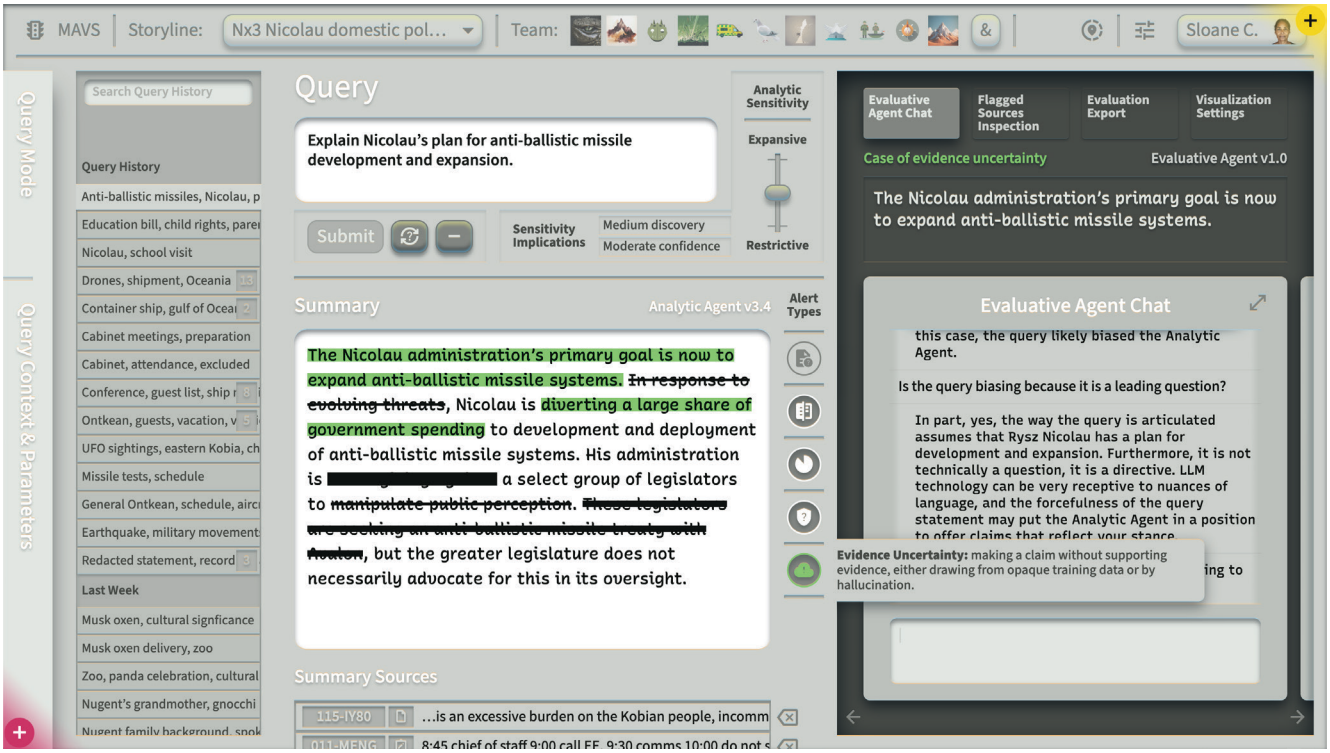

Figure 5. Simulation interface. In the pictured state, the user is hovering over the evidence uncertainty alert icon, which automatically highlights all evidence uncertainty passages in the summary.

7.1. Simulation Interface

The first Multiple Agent Validation System prototype is a simulation interface populated with scenario content. This prototype provides a realistic first-person experience of numerous MAVS features. Users play the part of Sloane as she engages in analysis of President Nicolau and his administration, with a narrative that permits significant lateral exploration through optional content. The web-based simulation interface was built in HTML, CSS, and JavaScript. While it does not incorporate actual AI — instead simulating AI — it is based on a tool developed by LAS (the collaborating lab) that utilizes an LLM and RAG with representative intelligence traffic.

We utilized a familiar control panel metaphor for the visual design of the interface to reduce cognitive load. Because simulation users are being asked to learn about an unfamiliar system (MAVS), they are not additionally asked to learn new interface conventions (Figure 5). The incomplete nature of the embodied scenario complicates the educational aspect of the prototype. Interface elements that are available at one time (i.e., that are scripted) are not available at others. To guide users through the scenario and to make sense of what is and is not interactable along the way, an instructional panel overlays one corner of the interface (Figure 6). The instructional panel suggests next steps with check boxes for completed tasks.

Another panel overlay permits users to select among eight uncertainty visualization styles to be utilized in the simulated LLM summaries (Figure 7). Like the instructional panel, the visualization panel would not be included in MAVS, for which a single visual convention for representing uncertainty would have been adopted. As such, the simulation interface allows practicing language analysts to experience different visual

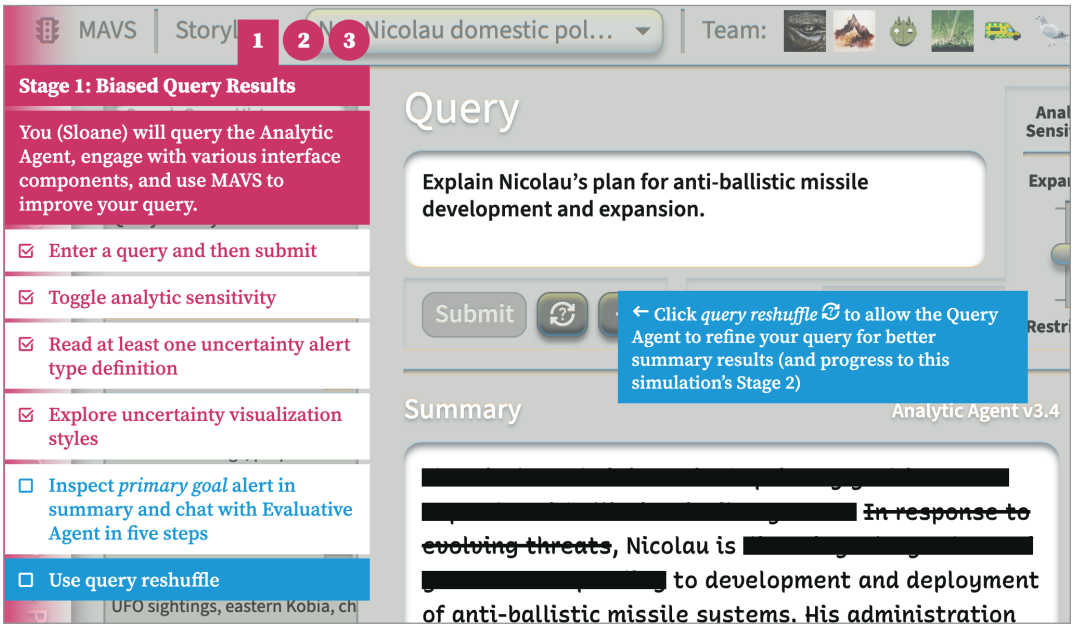

Figure 6. Simulation interface instructional panel with task hint. As listed tasks are completed, they are checked off. Hovering over incomplete tasks activates hints that point to interface elements.

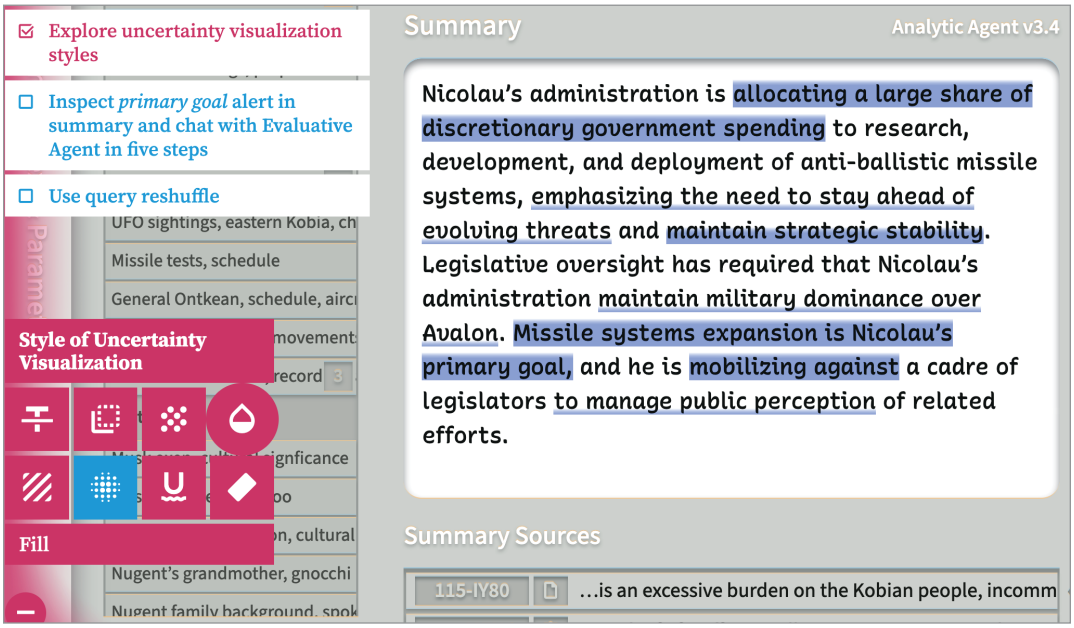

Figure 7. Simulation interface visualization panel with corresponding visual convention displayed in the summary.

conventions and consider their own preferences, leaving open the possibility of an end-user-informed determination of the ideal visual convention to adopt. In this way and others the simulation interface is a rich experimental stimulus that could be used to empirically test MAVS before costly development efforts are undertaken.

7.2. Narrative Interfaces

The simulation interface leverages familiarity by suggesting a control panel metaphor. However, this metaphor limits the potential capabilities of emergent technology as envisioned in MAVS. We thus also developed scenario videos for three distinct MAVS prototypes that are based on unconventional UX patterns, which may more naturally exemplify MAVS functionality. To focus the resultant narrative interfaces, we looked to the trust calibration (TC) literature and related interface design strategies TC1–TC5 (see Section 2). Based on these sources, we wrote three prompts to guide our interfaces.

1. *Transparency through interrogation and verification:* How might the interface utilize query recommendations, nudging, verification, and source inspection to calibrate trust between users and virtual agents? (Corresponds with TC1: Transparency and TC3: Alignment.)
2. *Multi-agent dialogue:* How might the interface use conversational AI to calibrate trust between users and virtual agents? (Corresponds with TC3: Alignment and TC5: Virtual Agents.)
3. *Context-driven:* How might the interface respond to the needs of specific users, customers, or storylines to calibrate trust between users and virtual agents? (Corresponds with TC3: Alignment and TC4: Interactivity.)

We utilized common UX methods to engage with our collaborators as we developed the narrative interfaces, including personas, scenarios, task flows, low- and high-fidelity sketching, and what-if prompts.

Narrative interface 1: Transparency through interrogation and verification. Language analysts want to leverage their own intricate understanding of human language to verify and interrogate data themselves. The first narrative interface provides an uncertainty alert report panel with a natural language explanation of identified uncertainty errors, along with key excerpts of flagged sources and quick access to the full sources themselves (Figure 8). This collects elements together in one space that the simulation interface distributes among separate zones. The uncertainty alert report panel reconfigures to accommodate the Evaluative Agent chat, and as with the simulation interface, the conversational interaction permits the analyst to submit operator notes — records for the chain of command and compliance — as they are suggested during sensemaking.

The first narrative interface deviates most dramatically from the simulation interface. It presents the user's workflow as branching diagrams of expanded and collapsed

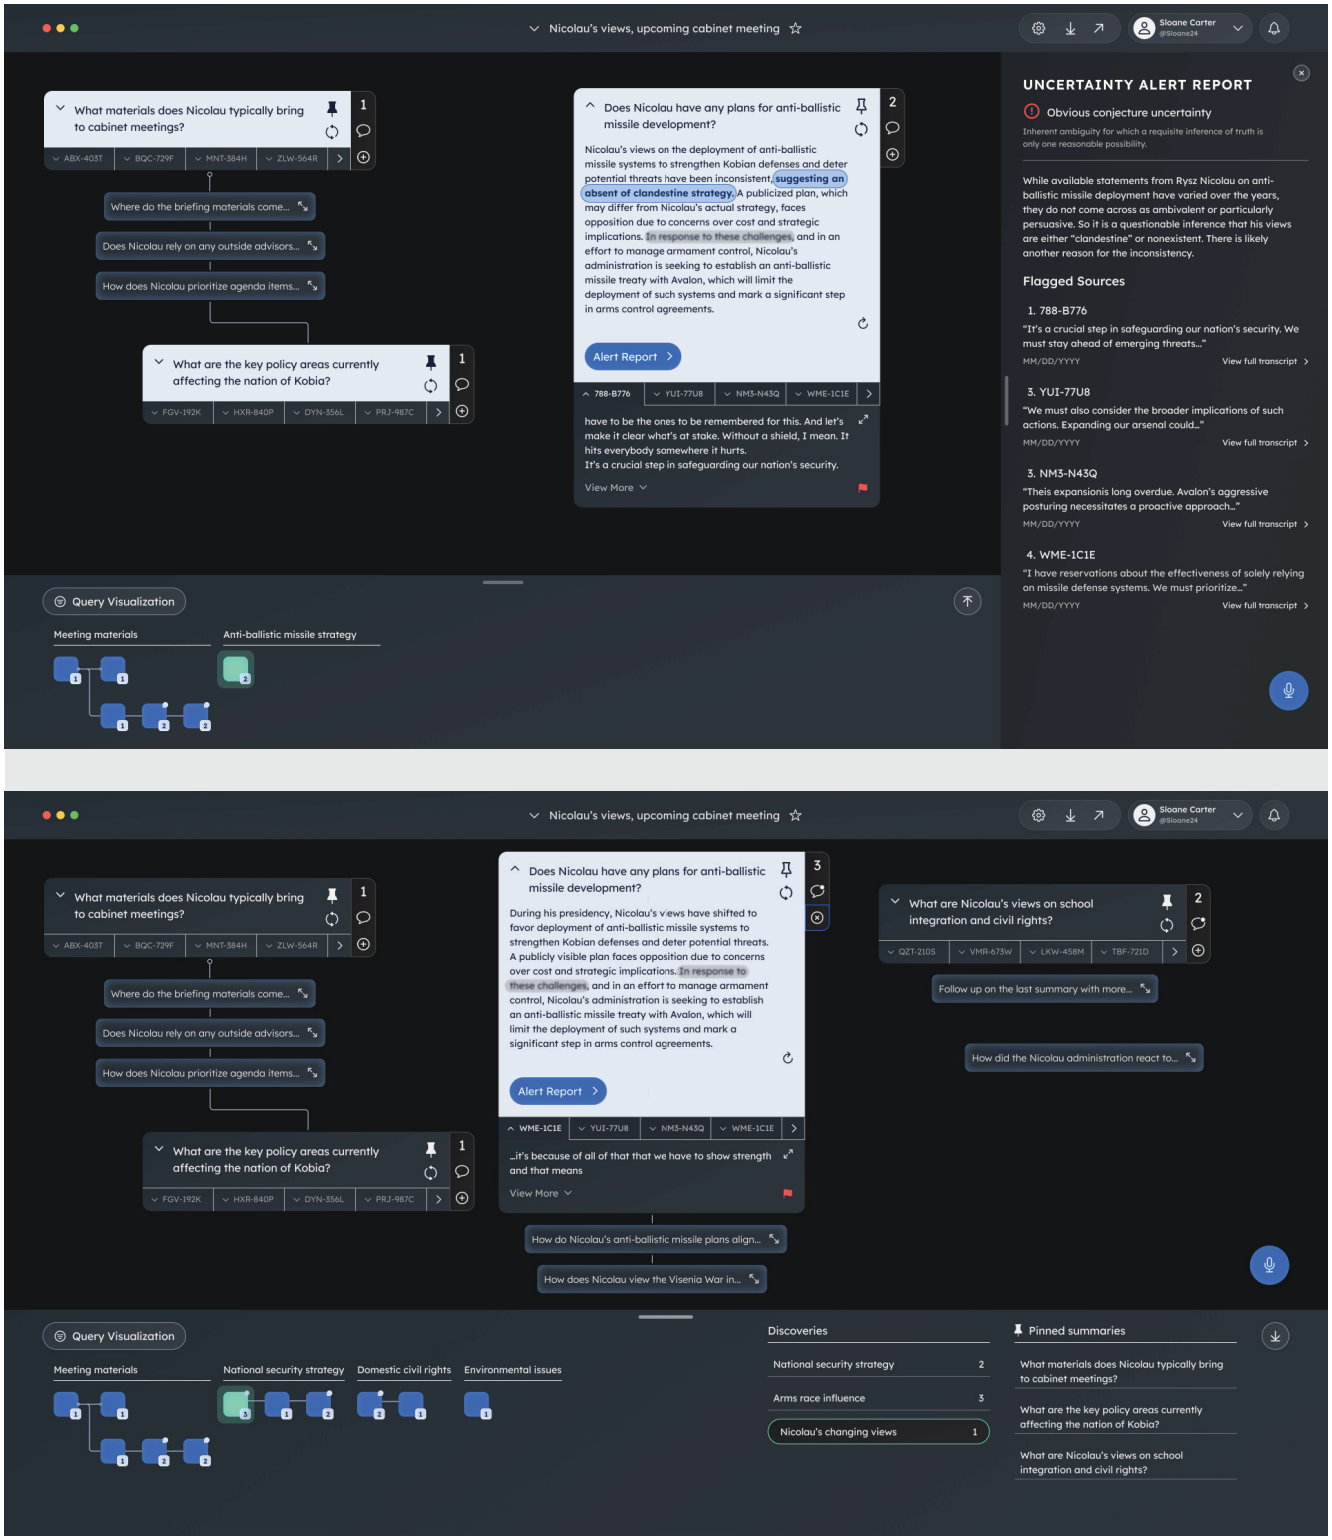

Figure 8. First narrative interface. Prompt-summary elements expand and collapse to varying degrees and their arrangement reflects the language analyst's investigatory process.

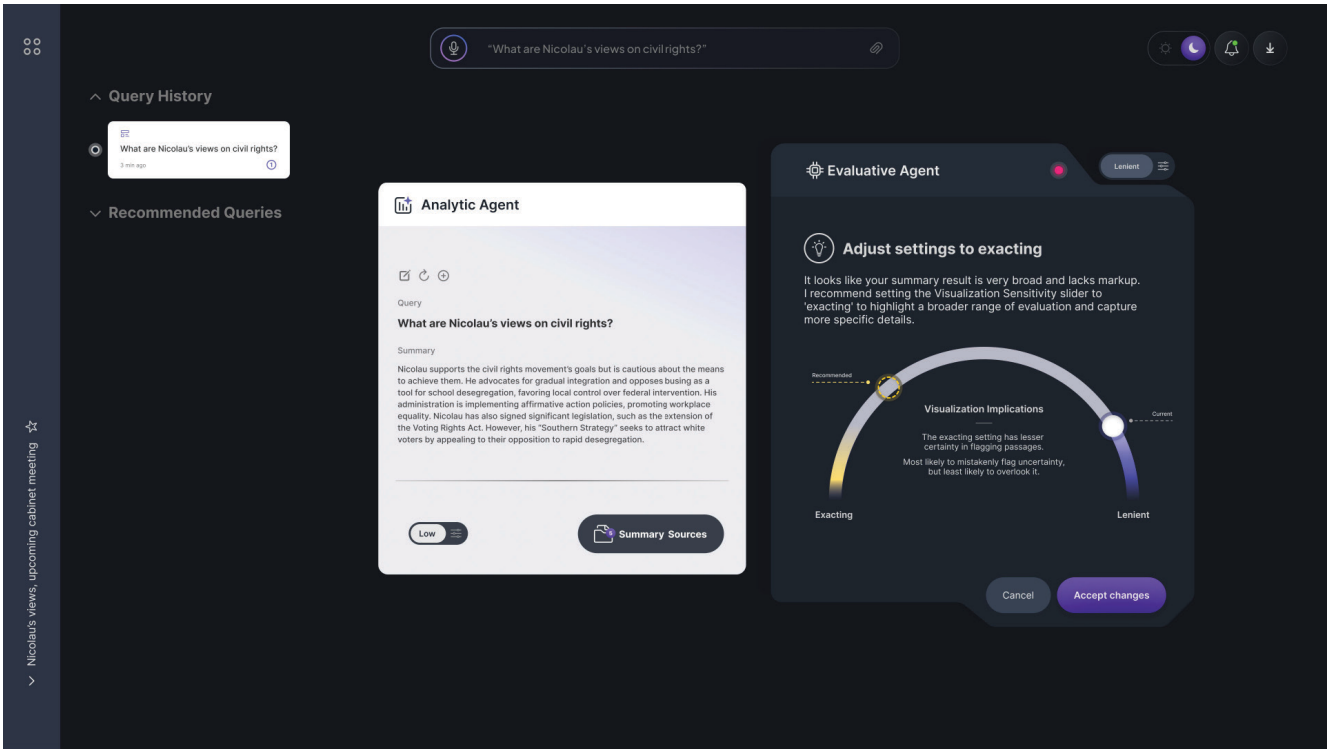

Figure 9. Second narrative interface with the Evaluative Agent active. The Analytic and Evaluative agents are presented as distinct entities that can respond to one another, and they make recommendations for the user.

prompt-summaries. In the central portion of the interface, prompts are chained together and break out when expanded to display summaries and other elements. In a lower query visualization strip, queries are minimized into icons and are organized into categories, while also reflecting investigatory pathways through chained connections. The interface constantly reconfigures as the user progresses. When appropriate, it nudges the user to revisit key points in the querying process and occasionally appends key insights to query icons. The interface enables the user to study how the Analytic and Evaluative Agents arrive at their conclusions, which helps the user better understand the system's capabilities.

Narrative interface 2: Multi-agent dialogue. The second narrative interface overtly presents the Analytic and Evaluative Agents as separate entities, embodied in adjacent floating panels. The user can converse with either virtual agent using the agent's panel, which grows slightly larger and includes a softly pulsing red light to reflect active engagement. Separating the virtual agents helps the user to conceptualize MAVS functionality by distinguishing analytic processes from evaluative processes. This facilitation is furthered with floating queries that can be dragged into a virtual agent's panel for a response, embodying the virtual agent with a recipient role. The pattern of virtual agents responding to each other, back-and-forth and through reciprocal panel

shrinking and growing behaviors, strengthens this embodiment. The greater the degree to which the interface distinguishes virtual agents, the easier it will be for the user to differentiate their system functions.

Narrative interface 3: Context-driven. The third narrative interface is an unconventional user interface that reconfigures itself as the user's investigation evolves (Figure 10). The user's investigatory process is structured as one long conversational flow. Both Analytic and Evaluative Agents converse with the user within this flow. The Analytic Agent auto-fills the analytic sensitivity setting according to the current storyline's dynamics. The user can adjust the setting, but recommendations are clearly indicated, possibly encouraging the user to experience settings they would not otherwise utilize. Recommended settings consider a variety of factors, such as the criticality of the storyline and past user behavior in similar situations. The Evaluative Agent also recognizes the criticality and greater context of the current storyline. For instance, it does not merely flag passages for contextually relevant uncertainty, it does so through a liquid panel that flows into the Analytic Agent's summary results. This violation of conventional panel integrity in interface design is a visual analogy for distinct virtual agent interactions and an incisive form of evaluation. Finally, the entire screen display also adjusts contextually, with the color scheme changing and element count reducing in critical high-stakes storyline periods to focus the user's attention.

These narrative interfaces provided an avenue for investigating trust calibration and interface design beyond the initial simulation interface. They embrace the potential of machine learning capabilities to build trust through transparency and conceptualization. The video format for presenting these interfaces guides viewers through a coherent workflow, making unconventional UX patterns sensible upon initial viewing. The novel element display and UX patterns within the narrative interfaces may potentially inform AI-based interface design beyond our investigation's focus on intelligence analysis.

8. Application and Transfer of Results

This section outlines aspects of our investigation that may contribute to communities of practice beyond intelligence analysis. We identify two application areas where investigation outcomes may be particularly relevant: LLM-assisted clinical decision-making in medicine and LLM-assisted climate forecasting. We examine relevant research in each area, viewing challenges related to uncertainty through the lens we have established. This suggests future work, but it also serves as a demonstration of how investigation outcomes can be adapted for additional areas not covered here.

Clinical decision-making in medicine. Researchers are actively exploring how LLMs can enhance clinical decision-making and provide diagnostic support in medicine

Figure 10. Third narrative interface. Chat elements flow into one another, and the system reconfigures itself in high-stakes moments.

(Nasarian et al., 2024; Panagoulias et al., 2023; Prabhod, 2023; Rajashekar et al., 2024; Savage et al., 2025). Some work has focused more narrowly on emergency care, digital pathology, and telehealth (Taylor et al., 2024; Kwan, 2024; Ullah et al., 2024). A common application of LLMs in clinical decision-making is assisting clinicians in prioritizing differential diagnoses (Prabhod, 2023; Taylor et al., 2024). Differential diagnosis is the systematic process used by clinicians to identify the most likely diagnoses from a set of competing possibilities (Cook & Décary, 2019). While established systems support this process, there is growing interest in expanding AI's role to mitigate diagnostic errors, improve information gathering, and facilitate diagnostic feedback (Taylor et al., 2024).

Despite the recognized potential of machine learning in this domain, a frequently cited challenge to integration is the lack of explainability in LLM-augmented systems, which has been shown to undermine user trust and hinder technology adoption (Panagoulias et al., 2023; Rajashekar et al., 2024; Savage et al., 2025; Ullah et al., 2024). This is particularly critical in high-stakes medical environments, where the urgency of decision-making, the fragmented nature of data, and the potential for cognitive overload leave little tolerance for uncertainty. Some researchers have focused on developing evaluation systems to measure or minimize uncertainty, but the literature does not fully address how to communicate uncertainty effectively to users in a medical context (e.g., Panagoulias et al., 2023; Savage et al., 2024). Nasarian et al. (2024) argue that many of the current explainable AI efforts are developer-centric, and that they tend to overlook the actual needs of end users. They further suggest that while machine learning professionals and developers tend to favor technical explanations, clinicians and patients would benefit from more intuitive visual formats. Both Prabhod (2023) and Kwan (2024) argue that future research should take a user-centered design approach, and should explore ways to provide training and meaningful engagement for clinicians. Kwan (2024) emphasizes that understanding user needs, defining user personas, and building prototypes are essential steps in developing AI-driven systems that can support clinical decision-making. Similarly, Taylor et al. (2024) emphasize the importance of designing AI tools that integrate seamlessly into clinician workflows and platforms without adding unnecessary complexity — namely in electronic health records (EHRs).

Climate forecasting. Similar issues arise in climate forecasting, where LLM-augmented systems must balance accuracy, interpretability, and usability to support not only decision-making, but also data analysis, communication to lay audiences, and generation of climate scenarios that can further inform decisions (Biswas, 2023). A preliminary search anecdotally suggests that the research in this area may not be as established as in medical diagnosis — most relevant work is in the form of unvetted uploads to preprint servers that we do not cover here.

Biswas (2023) explored how ChaptGPT could be leveraged to support climate research and policymaking. In this context LLMs are useful for generating instructive climate scenarios. However, the tendency of LLMs to hallucinate is significant due to limited context and expertise on climate data. For this reason, Biswas (2023) suggests that AI be used alongside traditional climate research methods, rather than as a replacement. Vaghefi et al. (2023) developed a specialized LLM, ChatClimate, to respond to queries related specifically to climate science. Their goal in creating this tailored model was to address challenges like hallucination and the presence of outdated information that might arise when using general purpose models. ChatClimate was not developed to replace the kinds of decision-making currently done by climate experts, but to increase the speed at which quality information on climate science can be accessed.

In contrast to broader applications of LLMs in climate science, Lawson et al. (2025) focused on weather forecasting at a more immediate ground level. They examined how well ChatGPT could analyze meteorological imagery and communicate hazard summaries in English and Spanish. ChatGPT struggled with the same challenges encountered in other climate applications, including hallucination and a lack of explainability and trustworthiness. Lawson et al.'s (2025) findings suggest that work remains to be done. Since these models are being leveraged for both long-term climate projections and real-time weather hazards, better representations of uncertainty could help facilitate trust calibration and generally improve system performance.

Implications. Our investigation offers several contributions that may be relevant to both clinical decision-making in medicine and climate forecasting, where LLM-augmented systems currently struggle with explainability and trust. The Uncertainty Framework for Explainable Summaries (Section 5) could help both clinicians and climate researchers interpret model outputs more effectively. The framework's pairing of evidence and credibility uncertainty would add nuance to understanding of uncertainty in both application areas, and the Multiple Agent Validation System's Evaluative Agent chat feature (Section 6) would help climate scientists vet the outdated information they frequently encounter. MAVS generally facilitates the rapid information gathering important in both application areas. In clinical settings, it could help mitigate the risks of misdiagnosis by ensuring AI-assisted insights are more transparent. In climate forecasting, it could improve trust in AI-generated climate scenarios by providing clearer explanations of LLM limitations. MAVS was conceptualized through user-focused design exploration, and thus does not suffer from the "developer-centric" emphasis of AI efforts in clinical decision-making (Nasarian et al., 2024).

A key investigation outcome for both application areas is the open resource of visual conventions for representing uncertainty in LLM-generated summaries (Section 4 and Section 7.1). The particulars of an application space are likely to influence valuation of

visualization efficacy, and the emergent criteria for selection can scaffold fresh valuation (Section 4). These contributions address core issues of accuracy and interpretability in domains where decisions have real-world consequences.

9. Discussion

Automation transparency has a positive impact on user task performance (van de Merwe et al., 2022). Knowledge workers need insight into the LLM systems they increasingly rely upon. And uncertainty is unavoidable with LLMs. While communicating this uncertainty benefits users, user outcomes differ based on the indicated degree of uncertainty — e.g., Kunze et al. (2019) noted participant behavioral changes at three levels of uncertainty, which correspond with our signification of conceivable, likely, and obvious levels.

Kunze et al. (2019) highlighted a "drawback" to displaying uncertainty: users need to look away from the task at hand to attend to visualizations (p. 355). But they were studying automated driving systems. When we explored uncertainty visualization, we did consider adjacent visualizations of uncertainty separate from the uncertain summaries, but our eight potential visual conventions are all inline, occurring directly within uncertain textual passages themselves (Section 4). The relevant literature on visualization tends to address standalone representations — the uncertainty framework we started with, Skeels et al. (2010), is a good example. The simple fact that our recommended forms of signification occur at the locus of uncertainty for LLMs — *in* and *as* written language — is possibly a powerful visual affordance for user interface design and transparent AI. It is possible for uncertain text to be — really, *to appear* — uncertain itself. This seems more desirable than providing an additional thing for overtaxed users to look at.

A result of interdisciplinary collaboration between language analysts, computer scientists, psychologists, and designers, this investigation provides human-centered recommendations that can guide LLM technology development. The expertise of the extended collaborative team ensures that the Multiple Agent Validation System, as described, is both implementable and relevant. This places our speculative design squarely in the present.

As a discovery-based process, this investigation suggested new research questions instead of answering preconceived ones. There are a variety of ways to continue this work. The most direct way would be completing the theory building and testing cycle through the empirical study of core project premises. The simulation interface (Section 7.1) could be used to test the MAVS summary validation process with intelligence analysts. The objective would be to identify uncertainty visualizations and

interface features that optimize intelligence analysts' ability to accurately validate LLM summaries with analytic and evaluative assistance from AI (as simulated, not implemented). This suggests two research questions:

- *Research Question 1 (RQ1):* What preferences do intelligence analysts have in the design of uncertainty communication (including visualization), and how do those preferences translate into trust attitudes and dependence behaviors?
- *RQ2:* When intelligence analysts are presented with potentially erroneous information, what actions do they take to validate and integrate the information with existing schemas, and how does this influence the performance of the analyst-automation team?

Answering these research questions would result in design principles that could guide work in human-machine teaming, explicit design implementations for effective communication about the uncertainty of LLMs, and an understanding of analyst information validation behavior based on trust.

Individual components of this investigation suggest other possibilities. Collaboration with experts and design exploration suggested deviating from Skeels et al.'s (2010) framework for uncertainty, but unlike that framework, our proposed Uncertainty Framework for Explainable Summaries has not been validated. A qualitative study of UFES with members of the intelligence community and LLM developers could utilize the framework's five types as *a priori* codes to refine it. Likewise, there has been no validation of the eight implemented visual conventions for representing uncertainty. Qualitative or quantitative research with intelligence analysts could tease out how impressions of the visualizations correspond with, and contribute to, mental models of uncertainty.

- *RQ3:* What types of uncertainty are present in LLM summaries, and how do intelligence analysts and LLM developers conceptualize these types?
- *RQ4:* How does the visualization of uncertainty in LLM summaries impact conceptualizations of uncertainty?

The results of studies like these could lead to modifications to MAVS specifications, such as altering the descriptors for analytic and visualization sensitivity.

As described in Section 8, this investigation has relevance for clinical design-making in medicine and climate forecasting. There are potential investigations in these transfer domains. For instance, there is interest in improving information gathering for differential medical diagnosis (Taylor et al., 2024), and in generating climate scenarios that are context-sensitive (Biswas, 2023).

- *RQ5:* How do clinicians conceptualize and utilize restrictive and expansive settings for differential diagnosis when an analytic sensitivity feature is available in MAVS?
- *RQ6:* How do climate scientists prompt an Evaluative Agent to vet an Analytic Agent's generated climate scenarios in climate policymaking?

In terms of application, the speculative interfaces explored in this investigation provide various pathways for developing MAVS according to context and constraints. Operationalizing the behaviors of the Query, Analytic, and Evaluative Agents would require significant iteration in training LLMs, but fundamentally new training methods need not be developed.

This investigation resulted in visual conventions for representing uncertainty in LLM summaries, a related framework for uncertainty, specifications for an LLM validation system, and situated prototypes of such a system. We have been explicit about the investigation's contributions throughout to promote application in and beyond the intelligence analysis context. Trust calibration plays a key role in successful human-machine teaming. We will never move beyond human limitations if we cannot trust AI to support and augment our abilities. User interface design plays a critical role in trust calibration because interfaces lie between humans and automated systems. Through thoughtful interface design that prioritizes transparency and human understanding, we can build the foundations of trust necessary for humans and AI to work together effectively, revealing new possibilities.

10. Acknowledgments

The authors were the investigation's principal investigators (Peterson and Armstrong) and research assistants (Planchart, Baidoo, and Anderson, along with Kayla Rondinelli). Collaborators at the Laboratory for Analytic Sciences included Christine Brugh, Patti K., Sue Mi K., and Mike G., among others. RQ1 and RQ2 were written with Colleen Patton. This material is based upon work done, in whole or in part, in coordination with the Department of Defense (DoD). Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the DoD and/or any agency or entity of the United States Government.

11. References

Alhadad, S. S. (2018). Visualizing data to support judgement, inference, and decision making in learning analytics: Insights from cognitive psychology and visualization science. *Journal of Learning Analytics, 5*(2), 60–85. https://doi.org/10.18608/jla.2018.52.5

Amershi, S., Inkpen, K., Teevan, J., Kikin-Gil, R., Horvitz, E., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., & Bennett, P. N. (2019). Guidelines

for human-AI interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (pp. 1–13). Association for Computing Machinery. https://doi.org/10.1145/3290605.3300233

Baldassi, S., Megna, N., & Burr, D. C. (2006). Visual clutter causes high-magnitude errors. *PLOS Biology, 4*(3), 387–394. https://doi.org/10.1371/journal.pbio.0040056

Bansal, G., Nushi, B., Kamar, E., Lasecki, W. S., Weld, D. S., & Horvitz, E. (2019). Beyond accuracy: The role of mental models in human-AI team performance. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing, 7*(1), 2–11. https://doi.org/10.1609/hcomp.v7i1.5285

Bellotti, V., & Edwards, K. (2001). Intelligibility and accountability: Human considerations in context-aware systems. *Human–Computer Interaction, 16*(2–4), 193–212. https://doi.org/10.1207/S15327051HCI16234_05

Biswas, S. S. (2023). Potential use of Chat GPT in global warming. *Annals of Biomedical Engineering, 51*(6), 1126–1127. https://doi.org/10.1007/s10439-023-03171-8

Borgo, R., Wall, E., Matzen, L., El-Assady, M., Masters, P., Hosseinpour, H., Endert, A., Chau, P., Perer, A., Schupp, H., Strobelt, H., & Padilla, L. (2024). Trust junk and evil knobs: Calibrating trust in AI visualization. *Proceedings of the IEEE 17th Pacific Visualization Conference (PacificVis)*, 22–31. https://doi.org/10.1109/pacificvis60374.2024.00012

Cook, C. E., & Décary, S. (2020). Higher order thinking about differential diagnosis. *Brazilian Journal of Physical Therapy, 24*(1), 1–7. https://doi.org/10.1016/j.bjpt.2019.01.010

Corritore, C. L., Kracher, B., & Wiedenbeck, S. (2003). Online trust: Concepts, evolving themes, and a model. *International Journal of Human-Computer Studies, 58*(6), 737–758. https://doi.org/10.1016/S1071-5819(03)00041-7

Cummings, M. L. (2006). Automation and accountability in decision support system interface design. *Journal of Technology Studies, 32*(1), 23–31. http://hdl.handle.net/1721.1/90321

Dasgupta, A., Lee, J.-Y., Wilson, R., Lafrance, R. A., Cramer, N., Cook, K., & Payne, S. (2017). Familiarity vs. trust: A comparative study of domain scientists' trust in visual analytics and conventional analysis methods. *IEEE Transactions on Visualization and Computer Graphics, 23*(1), 271–280. https://doi.org/10.1109/tvcg.2016.2598544

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2018). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science, 64*(3), 1155–1170. https://doi.org/10.1287/mnsc.2016.2643

Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv*. https://arxiv.org/abs/1702.08608

Evans, J. St. B. T., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science, 8*(3), 223–241. https://doi.org/10.1177/1745691612460685

Fell, L., Gibson, A., Bruza, P., & Hoyte, P. (2020). Human information interaction and the cognitive predicting theory of trust. *Proceedings of the 2020 Conference on Human Information Interaction and Retrieval*, 145–152. https://doi.org/10.1145/3343413.3377981

Forsey, H., Leahy, D., Fields, B., Minocha, S., Attfield, S., & Snell, T. (2024). Designing for learnability: Improvement through layered interfaces. *Ergonomics in Design: The Quarterly of Human Factors Applications*. https://doi.org/10.1177/10648046241273291.

Gentner, D., & Smith, L. (2012). Analogical reasoning. In V. S. Ramachandran (Ed.), *Encyclopedia of human behavior (2nd ed.),* pp. 130–136. Elsevier.

Graaf, M. D., & Malle, B. F. (2017). How people explain action (and autonomous intelligent systems should too). *Proceedings of the AAAI Fall Symposium*. https://api.semanticscholar.org/CorpusID:53292049

Heger, O., Kampling, H., & Niehaves, B. (2016). Towards a theory of trust-based acceptance of affective technology. *European Conference on Information Systems*. https://www.semanticscholar.org/paper/Towards-a-Theory-of-Trust-based-Acceptance-of-Heger-Kampling/4640516f7b9a2a693a80ce667c98dba1fc7eb8bb

Johnson, M. (1987). *The body in the mind: The bodily basis of meaning, imagination, and reason.* University of Chicago Press.

Joshi, S., Nistala, P. V., Jani, H., Sakhardande, P., & Dsouza, T. (2017). User-centered design journey for pattern development. In *Proceedings of the 22nd European Conference on Pattern Languages of Programs* (EuroPLoP '17) (Article 23, pp. 1–19). Association for Computing Machinery. https://doi.org/10.1145/3147704.3147730

Karran, A. J., Demazure, T., Hudon, A., Senecal, S., & Léger, P.-M. (2022). Designing for confidence: The impact of visualizing artificial intelligence decisions. *Frontiers in Neuroscience, 16*. https://doi.org/10.3389/fnins.2022.883385

Kirschner, P. A., Ayres, P., & Chandler, P. (2011). Contemporary cognitive load theory research: The good, the bad, and the ugly. *Computers in Human Behavior, 27*(1), 99–105. https://doi.org/10.1016/j.chb.2010.06.025

Kocielnik, R., Amershi, S., & Bennett, P. N. (2019). Will you accept an imperfect AI? Exploring designs for adjusting end-user expectations of AI systems. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (pp. 1–14). Association for Computing Machinery. https://doi.org/10.1145/3290605.3300641

Krueger, R. F., Beyer, J., Jang, W.-D., Kim, N., Sokolov, A., Sorger, P. K., & Pfister, H. (2020). Facetto: Combining unsupervised and supervised learning for hierarchical phenotype analysis in multi-channel image data. *IEEE Transactions on Visualization and Computer Graphics, 26*(1), 227–237. https://doi.org/10.1109/TVCG.2019.2934547

Kunze, A., Summerskill, S. J., Marshall, R., & Filtness, A. J. (2019). Automation transparency: Implications of uncertainty communication for human-automation interaction and interfaces. *Ergonomics, 62*(3), 345–360. https://doi.org/10.1080/00140139.2018.1547842

Kwan, H. Y. (2024). User-focused telehealth powered by LLMs: Bridging the gap between technology and human-centric care delivery. *2024 4th International Conference on Computer Communication and Artificial Intelligence (CCAI)*, 187–191. https://doi.org/10.1109/CCAI61966.2024.10603150

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by.* University of Chicago Press.

LAS [Laboratory for Analytic Sciences] (2024). Sight unseen, students design UX for intelligence analysts [webpage]. https://ncsu-las.org/2024/09/sight-unseen-students-design-ux-for-intelligence-analysts/

Lawson, J. R., Trujillo-Falcón, J. E., Schultz, D. M., Flora, M. L., Goebbert, K. H., Lyman, S. N., Potvin, C. K., & Stepanek, A. J. (2025). Pixels and predictions: Potential of GPT-4V in meteorological imagery analysis and forecast communication. *Artificial Intelligence for the Earth Systems, 4*(1), 240029. https://doi.org/10.1175/AIES-D-24-0029.1

Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors: The Journal of the Human Factors and Ergonomics Society, 46*(1), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392

Manzey, D., Reichenbach, J., & Onnasch, L. (2012). Human performance consequences of automated decision aids: The impact of degree of automation and system experience. *Journal of Cognitive Engineering and Decision Making, 6*(1), 57–87. https://doi.org/10.1177/1555343411433844

Nasarian, E., Alizadehsani, R., Acharya, U. R., & Tsui, K.-L. (2024). Designing interpretable ML system to enhance trust in healthcare: A systematic review to proposed responsible

clinician-AI-collaboration framework. *Information Fusion, 108*, 102412. https://doi.org/10.1016/j.inffus.2024.102412

Nass, C., & Brave, S. (2007). *Wired for speech: How voice activates and advances the human-computer relationship*. MIT Press.

Okamura, K., & Yamada, S. (2020). Empirical evaluations of framework for adaptive trust calibration in human-AI cooperation. *IEEE Access, 8*, 220335–220351. https://doi.org/10.1109/ACCESS.2020.3042556

Padilla, L. M., Creem-Regehr, S. H., Hegarty, M., & Stefanucci, J. K. (2018). Decision making with visualizations: A cognitive framework across disciplines. *Cognitive Research: Principles and Implications, 3*, 29. https://doi.org/10.1186/s41235-018-0120-9

Panagoulias, D. P., Palamidas, F. A., Virvou, M., & Tsihrintzis, G. A. (2023). Evaluating the potential of LLMs and ChatGPT on medical diagnosis and treatment. *2023 14th International Conference on Information, Intelligence, Systems & Applications (IISA)*, 1–9. https://doi.org/10.1109/IISA59645.2023.10345968

Peterson, M., & Armstrong, H. (2024). Developing visual conventions for explainable LLM outputs in intelligence analysis summaries [webpage]. https://textimage.org/uncertainty/

Prabhod, K. J. (2023). Integrating large language models for enhanced clinical decision support systems in modern healthcare. *Journal of Machine Learning for Healthcare Decision Support, 3*(1), 18–62. https://medlines.uk/index.php/JMLHDS/article/view/23

Prabhudesai, S., Yang, L., Asthana, S., Huan, X., Liao, Q. V., & Banovic, N. (2023). Understanding uncertainty: How lay decision-makers perceive and interpret uncertainty in human-AI decision making. *Proceedings of the 28th International Conference on Intelligent User Interfaces*, 379–396. https://doi.org/10.1145/3581641.3584033

Rajashekar, N. C., Shin, Y. E., Pu, Y., Chung, S., You, K., Giuffre, M., Chan, C. E., Saarinen, T., Hsiao, A., Sekhon, J., Wong, A. H., Evans, L. V., Kizilcec, R. F., Laine, L., McCall, T., & Shung, D. (2024). Human-algorithmic interaction using a large language model-augmented artificial intelligence clinical decision support system. *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–20. https://doi.org/10.1145/3613904.3642024

Rheu, M., Shin, J. Y., Peng, W., & Huh-Yoo, J. (2020). Systematic review: Trust-building factors and implications for conversational agent design. *International Journal of Human–Computer Interaction, 37*(1), 81–96. https://doi.org/10.1080/10447318.2020.1807710

Rosenholtz, R., Li, Y., & Nakano, L. (2007). Measuring visual clutter. *Journal of Vision, 7*(2), 17–22. https://doi.org/10.1167/7.2.17

Sacha, D., Senaratne, H., Kwon, B. C., Ellis, G., & Keim, D. A. (2016). The role of uncertainty, awareness, and trust in visual analytics. *IEEE Transactions on Visualization and Computer Graphics, 22*(1), 240–249. https://doi.org/10.1109/tvcg.2015.2467591

Savage, T., Wang, J., Gallo, R., Boukil, A., Patel, V., Safavi-Naini, S. A. A., Soroush, A., & Chen, J. H. (2025). Large language model uncertainty proxies: Discrimination and calibration for medical diagnosis and treatment. *Journal of the American Medical Informatics Association, 32*(1), 139–149. https://doi.org/10.1093/jamia/ocae254

Shneiderman, B. (1996). The eyes have it: A task by data type taxonomy for information visualizations. *Proceedings 1996 IEEE Symposium on Visual Languages*, 336–343. https://doi.org/10.1109/VL.1996.545307

Skeels, M., Lee, B., Smith, G., & Robertson, G. G. (2010). Revealing uncertainty for information visualization. *Information Visualization, 9*(1), 70–81. https://doi.org/10.1145/1385569.1385637

Sorkin, R. D., & Woods, D. D. (1985). Systems with human monitors: A signal detection analysis. *Human–Computer Interaction, 1*(1), 49–75. https://doi.org/10.1207/s15327051hci0101_2

Sperrle, F., El-Assady, M., Guo, G., Borgo, R., Chau, D. H., Endert, A., & Keim, D. (2021). A survey of human-centered evaluations in human-centered machine learning. *Computer Graphics Forum, 40*(3), 543–568. https://doi.org/10.1111/cgf.14329

Stevenson, M. T. (2017). Assessing risk assessment in action. *Minnesota Law Review, 103*, 303. https://doi.org/10.2139/ssrn.3016088

Sultanum, N., Singh, D., Brudno, M., & Chevalier, F. (2019). Doccurate: A curation-based approach for clinical text visualization. *IEEE Transactions on Visualization and Computer Graphics, 25*(1), 142–151. https://doi.org/10.1109/tvcg.2018.2864905

Suresh, H., Lao, N., & Liccardi, I. (2020). Misplaced trust: Measuring the interference of machine learning in human decision-making. *Proceedings of the 12th ACM Conference on Web Science* (pp. 315–324). https://doi.org/10.1145/3394231.3397922

Taylor, R. A., Sangal, R. B., Smith, M. E., Haimovich, A. D., Rodman, A., Iscoe, M. S., Pavuluri, S. K., Rose, C., Janke, A. T., Wright, D. S., Socrates, V., & Declan, A. (2024). Leveraging artificial intelligence to reduce diagnostic errors in emergency medicine: Challenges, opportunities, and future directions. *Academic Emergency Medicine*. https://doi.org/10.1111/acem.15066

Thomson, J., Hetzler, E., MacEachren, A., Gahegan, M., & Pavel, M. (2005). A typology for visualizing uncertainty. *Visualization and Data Analysis 2005, 5669*, 146–157. https://doi.org/10.1117/12.587254

Tintarev, N., & Masthoff, J. (2007). Effective explanations of recommendations. *Proceedings of the 2007 ACM Conference on Recommender Systems* (pp. 153–156). https://doi.org/10.1145/1297231.1297259

Tomsett, R., Preece, A., Braines, D., Cerutti, F., Chakraborty, S., Srivastava, M., Pearson, G., & Kaplan, L. (2020). Rapid trust calibration through interpretable and uncertainty-aware AI. *Patterns, 1*(4), 100049. https://doi.org/10.1016/j.patter.2020.100049

Ullah, E., Parwani, A., Baig, M. M., & Singh, R. (2024). Challenges and barriers of using large language models (LLM) such as ChatGPT for diagnostic medicine with a focus on digital pathology – a recent scoping review. *Diagnostic Pathology, 19*(1), 43. https://doi.org/10.1186/s13000-024-01464-7

Vaghefi, S. A., Stammbach, D., Muccione, V., Bingler, J., Ni, J., Kraus, M., Allen, S., Colesanti-Senni, C., Wekhof, T., Schimanski, T., Gostlow, G., Yu, T., Wang, Q., Webersinke, N., Huggel, C., & Leippold, M. (2023). ChatClimate: Grounding conversational AI in climate science. *Communications Earth & Environment, 4*(1), 1–13. https://doi.org/10.1038/s43247-023-01084-x

Van de Merwe, K., Mallam, S., & Nazir, S. (2024). Agent transparency, situation awareness, mental workload, and operator performance: A systematic literature review. *Human Factors, 66*(1), 180–208. https://doi.org/10.1177/00187208221077804

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems, 30*. https://proceedings.neurips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html

Weisz, J. D., He, J., Muller, M., Hoefer, G., Miles, R., & Geyer, W. (2024). Design principles for generative AI applications. *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Article 378, pp. 1–22). Association for Computing Machinery. https://doi.org/10.1145/3613904.3642466

Wickens, C. D., Gempler, K., & Morphew, M. E. (2000). Workload and reliability of predictor displays in aircraft traffic avoidance. *Transportation Human Factors, 2*(2), 99–126. https://doi.org/10.1207/STHF0202_01

Williams, T., Briggs, P., & Scheutz, M. (2015). Covert robot-robot communication: Human perceptions and implications for human-robot interaction. *Journal of Human-Robot Interaction, 4*(2), 24–49. https://doi.org/10.5898/JHRI.4.2.Williams

Zerilli, J., Bhatt, U., & Weller, A. (2022). How transparency modulates trust in artificial intelligence. *Patterns, 3*(4), 100455. https://doi.org/10.1016/j.patter.2022.100455

Zhang, Y., Liao, Q. V., & Bellamy, R. K. E. (2020). Effect of confidence and explanation on accuracy and trust calibration in AI-assisted decision making. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 295–305). https://doi.org/10.1145/3351095.3372852

Zhao, J., Wang, Y., Mancenido, M. V., Chiou, E. K., & Maciejewski, R. (2023). Evaluating the impact of uncertainty visualization on model reliance. *IEEE Transactions on Visualization and Computer Graphics, 1–15*. https://ieeexplore.ieee.org/document/10058545

Authors

Helen Armstrong is a professor of graphic and experience design and the director of the MGXD program at NC State University. Her research focuses on digital rights, human-machine teaming, and accessible design. Armstrong authored *Graphic Design Theory; Digital Design Theory*; and co-authored *Participate: Designing with User-Generated Content.* Her recent book, *Big Data, Big Design: Why Designers Should Care About Artificial Intelligence,* demystifies AI — specifically machine learning — while inspiring designers to harness this technology and establish leadership via thoughtful human-centered design. Armstrong is a past member of the AIGA National Board of Directors, the editorial board of *Design and Culture,* and a former chair of the AIGA Design Educators Community.

Ashley L. Anderson is an assistant professor of graphic design at Virginia Tech and a PhD in Design candidate at NC State University. Her research focuses on human-centered design and visual representation, particularly in the context of mental health and psychological intervention. She examines how design can shape and enhance the theories, processes, and methods used in psychological intervention.

Rebecca Planchart is a product designer at Pendo.io, a software experience management solution, where she supports enterprise platform and conversational AI initiatives. Her past research explored explainability and trust calibration in AI systems through UX and UI strategies. She is particularly interested in leveraging explainable AI to support users in high-stakes decision-making contexts.

Kweku Baidoo is a lecturer in graphic and experience design at NC State University. His work explores trust-centered design and visual strategies that support human understanding of complex AI systems. He is particularly interested in how AI-assisted decision-making can be designed to enhance appropriate user trust and performance in high-stakes domains such as healthcare.

Matthew Peterson is an associate professor of graphic and experience design at NC State University. His research focuses on visual representation in user interface design, recently including the facilitation of AI in intelligence analysis workflows through human-machine teaming, text-image integration in immersive user information systems, and the facilitation of scale cognition and numeracy in virtual environments.

DISPATCHES FROM INDUSTRY

A Seat at the Table: Designing for AI with Strategy, Vision, and Collaboration

Syashi Gupta

SAP, San Jose, California, USA (syashigupta[at]gmail.com)

Abstract: The speed at which artificial intelligence technology is integrated into products to ease user flows is redefining the role of designers, giving rise to specialized "AI designers" or "AI design specialists." In this article, I explore the evolving responsibilities of designers in the AI landscape, emphasizing the critical need for deep collaboration with engineering, legal, and product teams. Drawing from direct experience, I highlight the challenges of translating complex AI capabilities into user-centric, valuable product features, especially within established organizations grappling with legacy systems and lengthy development cycles. Key takeaways underscore designers' need to possess strong data literacy, continuously learn in a fast-paced field, and strategically advocate for AI applications that address genuine user needs. I outline the essential skills designers must cultivate, the opportunities presented by adaptive AI interfaces, the high stakes involved in responsible AI development, and the pressing questions the design community must address to shape a human-centered AI future.

Implications for research: This article focuses on the role and responsibilities of the emerging AI designer in modern product design and development. The distinction between AI for efficiency and AI for augmentation (Section 2.3) suggests a comprehensive framework that can help AI designers apply these categories and advocate for user and societal needs in the rush to incorporate AI functions into existing services. The discussion of user feedback loops (Section 2.6) characterizes good feedback systems as being granular, contextual, and actionable, with a palette of available UX patterns including inline corrections for refinement, transparent confidence scores, and feedback tagging. Empirical research is needed to provide AI designers with a generalized understanding of how these UI characteristics and UX patterns impact human understanding, and how they interact.

Keywords: AI design; AI implementation; AI strategy; collaborative practice; design industry

@: ISSUE〉 ARTICLE〉 DISPATCHES〉

Cite this article:

Gupta, S. (2025). A seat at the table: Designing for AI with strategy, vision, and collaboration. *Visible Language, 59*(2), 218–231.

First published online July 15, 2025.

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

Visible Language Consortium:
University of Leeds (UK)
University of Cincinnati (USA)
North Carolina State University (USA)

1. Evolution of Designers' Roles in the World of AI

The role of designers in artificial intelligence (AI) technology development has evolved rapidly over the past decade. The new role of "AI designer" or "AI design specialist" reflects the field's unique challenges and responsibilities. Throughout my experience designing AI-powered features, several recurring challenges have shaped and unfolded my role as a designer.*

I have seen AI retrofitted into legacy systems, which leads to feature bloat and redundant workflows, often overwhelming users instead of streamlining their experience. Long development cycles can result in an AI model becoming outdated or needing an update by the time a feature ships. A lack of direct collaboration between designers and engineers can cause teams to miss key edge cases, while inadequate feedback mechanisms in a product limit a designer's ability to understand the real user needs and the system's practical function. And I have seen AI introduced without a real use case, creating a fake aura of innovation without lending actual value to the user.

I have also witnessed experiences delivered by designers who lack a general understanding of how AI systems work, which makes it challenging to anticipate model behavior or system risks. These issues have pushed me to consider a new definition of design practice that is more data-literate, collaborative, user-centered, and strategy-driven. The necessary evolution of the AI designer's role is reflected in Figure 1.

2. Key Considerations for Designers Working with AI

2.1. Designers Are Strategists, Not Service Providers

Historically, technical teams and project stakeholders tend to perceive design as a downstream activity, incorporating it only after making core technical decisions. In the AI age, that model is outdated. Designing for AI is not just about creating sleek, beautiful interfaces, but also about defining how intelligent systems behave, adapt, and evolve.

Designers are increasingly responsible for:

- Mapping user pain points that AI can solve.
- Framing problems that guide model development.

---

* **Author note:** This article is based on firsthand experiences designing AI-powered features in creative tools used by millions, and in ongoing work in a cross-functional AI design systems team. It reflects real-world complexities, hard-learned lessons, and a deep belief that when design and development work in harmony, everyone wins, especially users.

- Visualizing system behavior to align teams across functions.
- Driving decisions about data input, transparency, and explainability.
- Advocating for human agency and ethical design.

Designers understand user behavior, human-computer interaction, and system feedback loops, so involving them from the start enables more meaningful use of AI.

2.2. From Visuals to Vision: Collaborating with Engineers and Product Teams

AI systems are not static. They are probabilistic, learning, and reactive. To design compelling experiences, designers must work closely with engineering teams to understand the AI system architecture, data pipelines, model behaviors, and potential edge cases.

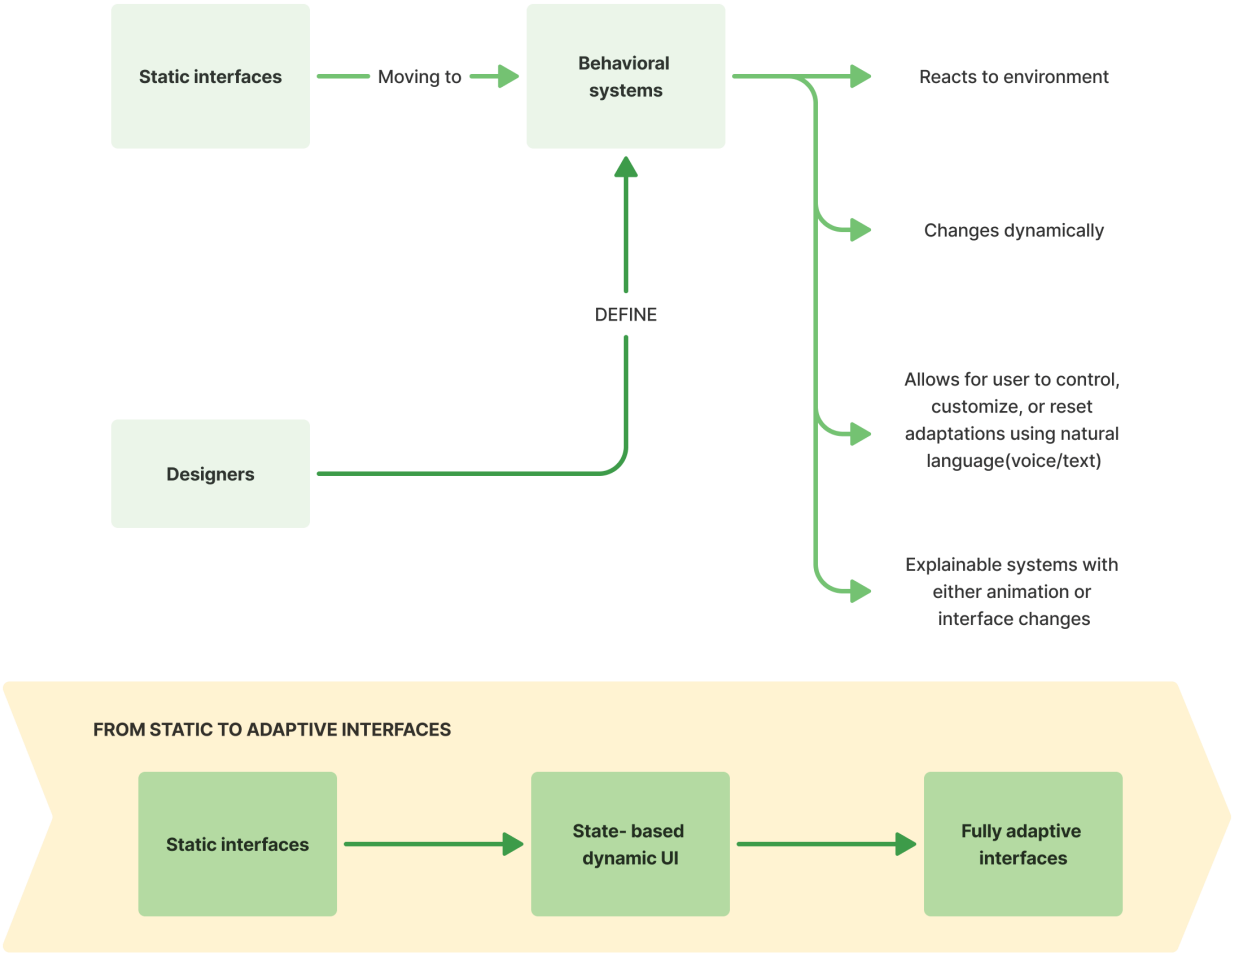

Figure 1. Evolution of AI-driven interfaces from static to fully adaptive interfaces that responds to the user's input.

Take, for example, Figma Make — a vibe coding tool that generates UI drafts using design system components (Levin, 2024). In 2024, Figma released an update following the discovery that the model inadvertently mimicked real-world app designs, so they held back its release to further train and reassess the model's performance. This emphasizes the importance of questioning how models are trained, what data is used, and where accountability lies (Figma, n.d).

Following are key questions that designers should ask in AI projects:

- What models power this experience?
- Is the model fine-tuned on user data? How is data collected and stored?
- What is the delivery mechanism (e.g., plugin, integrated system, API)?
- How will users give feedback, and how will the feedback be addressed? Does feedback require the user to share their data or their generation?
- Where is the data processing happening for the user? Is the user data anonymized? What is the default setting for enabling AI for the user? For example, Shakir (2024) outlines Apple's approach to keeping user data private.

Understanding AI system architecture is essential. Without this knowledge, designers risk shipping inconsistent experiences that fail to meet user expectations and raise data privacy concerns.

2.3. Designing AI: Where Capability Meets Context

AI's technical power means little if it does not meet real user needs. The most sophisticated algorithms are irrelevant when applied to the wrong problems, or worse, when they introduce unnecessary friction. Designers play a crucial role in grounding AI applications in real-world utility, ensuring that features address authentic pain points rather than showcasing novelty for novelty's sake.

In designing generative tools for image enhancement and object removal, I have witnessed firsthand how challenging it can be to bridge the gap between what AI can do and what users need. The gap between AI capability and users' needs is especially evident in creative tools, where users want precision and control, and considers those needs more essential than AI automation.

In general, AI use cases can be divided into two major value categories:

- **AI for efficiency:** streamlining repetitive, low-value tasks like smart cleanup, batch tagging, or background removal, to give users back their time.
- **AI for augmentation:** enhancing user creativity or judgment by delivering results better or faster than manual efforts, such as noise reduction, super-resolution, outcropping (extending the canvas), or improved stitching of frames in a video/movie.

For each implementation, designers must ask:

- Does this empower the user or obscure their control?
- Is it elevating human decision-making or replacing it entirely?

The distinction between assistive and autonomous systems is increasingly subtle. Designers are uniquely positioned to define this boundary — what gets handed off to the machine versus what remains in the user's hands? At Config 2025, speaker and roboticist Madeline Gannon stated that "automation is not inevitable; it is intentional" (2025, 02:35). As designers, we decide how an interaction works and what level of agency the user retains during and after an AI-driven process. Designing AI means shaping the conversation between human intent and machine intelligence.

2.4. Explainability: Building Trust, Not Just Functionality

Explainability is especially critical in generative AI systems, which are increasingly capable of learning and generating things that they were not explicitly trained on. These AI capabilities, even though quite excellent, require accountability, and through explaining the reasoning or logic, or data source, adoption and trust can be established with users (Stanford Institute for Human-Centered Artificial Intelligence, n.d.). Founders of Anthropic, which developed Claude AI, state that:

> ...generative AI systems are *grown* more than they are *built* — their internal mechanisms are "emergent" rather than directly designed. It is a bit like growing a plant or a bacterial colony: we set the high-level conditions that direct and shape growth, but the exact structure that emerges is unpredictable and difficult to understand or explain. (Amodei, 2025, paraphrasing Chris Olah)

This underscores the need for intentional design that makes AI systems *legible* to users. Whether through popovers, tooltips, or summary cards, designers must create affordances that explain how results are generated, what confidence levels are involved, and how users can override or refine outputs as reflected in Figure 2.

Explainability is not just an ethical concern, it is imperative for using any system with AI in it. It is how we build trust in systems that are otherwise black boxes.

2.5. Retrofitting Better: Using AI Tools to Present Vision

Startups often can integrate AI faster, not because they have better ideas, but because they operate with fewer constraints. Large organizations, in contrast, must navigate technical debt, complex review processes, legacy UX patterns, and rigorous legal, ethical, and accessibility checks.

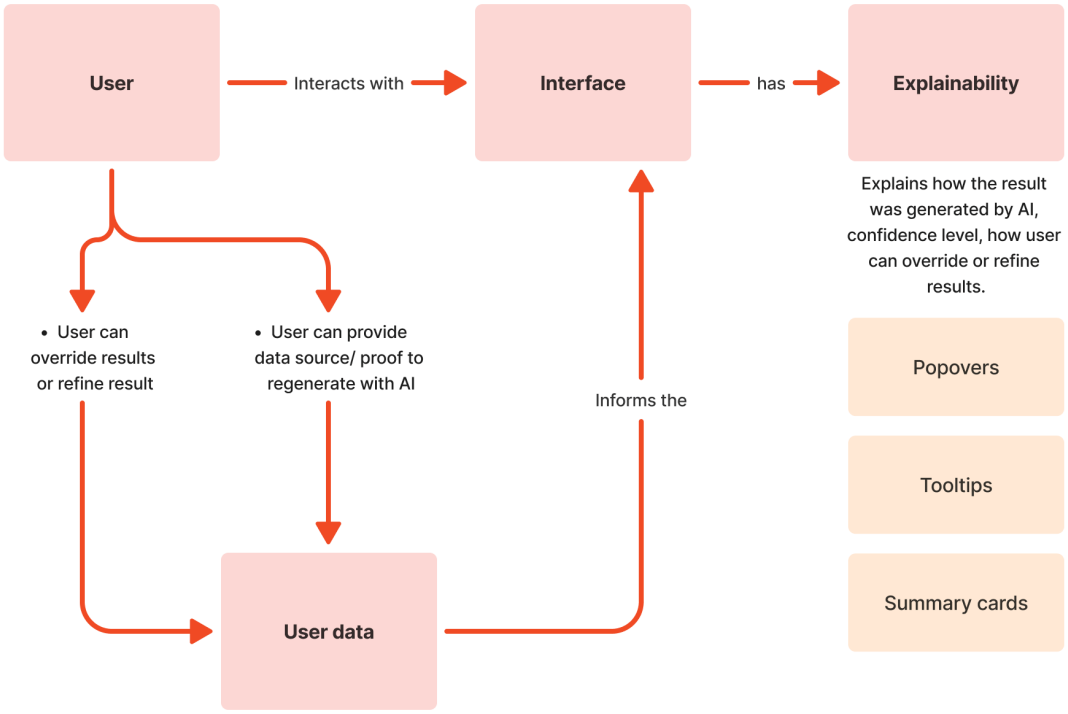

Figure 2. An explainability flow showing how users interact with an interface that should explain AI outputs and also help with refining, overriding, or regenerating results, through explanations using interface elements like popovers, tooltips, and summary cards.

Retrofitting AI into these mature systems often results in:

- **Feature bloat:** adding AI features without removing redundant ones.
- **Workflow redundancy:** offering multiple tools that solve the same or similar problems differently, without clearly indicating their distinction to the user. Photoshop's number of retouching tools reflects this (Figure 3).
- **Conflicting paradigms**: merging legacy interactions (like mouse select) with newer interactions for AI (e.g., prompts, voice, gestures), creating inconsistency and confusion.

The cumulative result is a bloated, disjointed experience that can overwhelm and confuse users rather than delight them with AI. Worse still, due to lengthy development cycles, the AI models powering these features are sometimes outdated by the time a feature ships.

In these scenarios, designers must advocate for modular systems, where AI features can evolve independently and plug into existing workflows without destabilizing them. However, the strategy also needs communication, which is where modern tools come in.

Tools that support *vibe coding* (low-code or no-code prototyping) can enable designers to prototype AI features faster and more accurately, even when developer collaboration is limited. These tools can turn sketches into working prototypes, making testing ideas easier, sharing vision, and aligning cross-functional teams early, an example of which is shown in Figures 4 and 5. I encourage students and professionals to use these tools to design and think through AI's behavior. Rapid prototyping is more than an execution tool — it is a strategy tool.

2.6. Feedback Loops: Where Did the Insights Go?

In AI-driven products, user feedback is not just important but essential for improvement. Nevertheless, most current systems treat feedback as a checkbox, using basic thumbs-up/down ratings that fail to capture nuance. What happens when a user's frustration does not fit into binary feedback?

Designers must build feedback mechanisms that reflect the complexity of AI interactions. Good feedback systems are:

- **Granular:** allow users to comment on why something worked or did not.
- **Contextual:** let users respond within the flow of interaction, not after.
- **Actionable:** tag feedback for tone, relevance, accuracy, or usability.

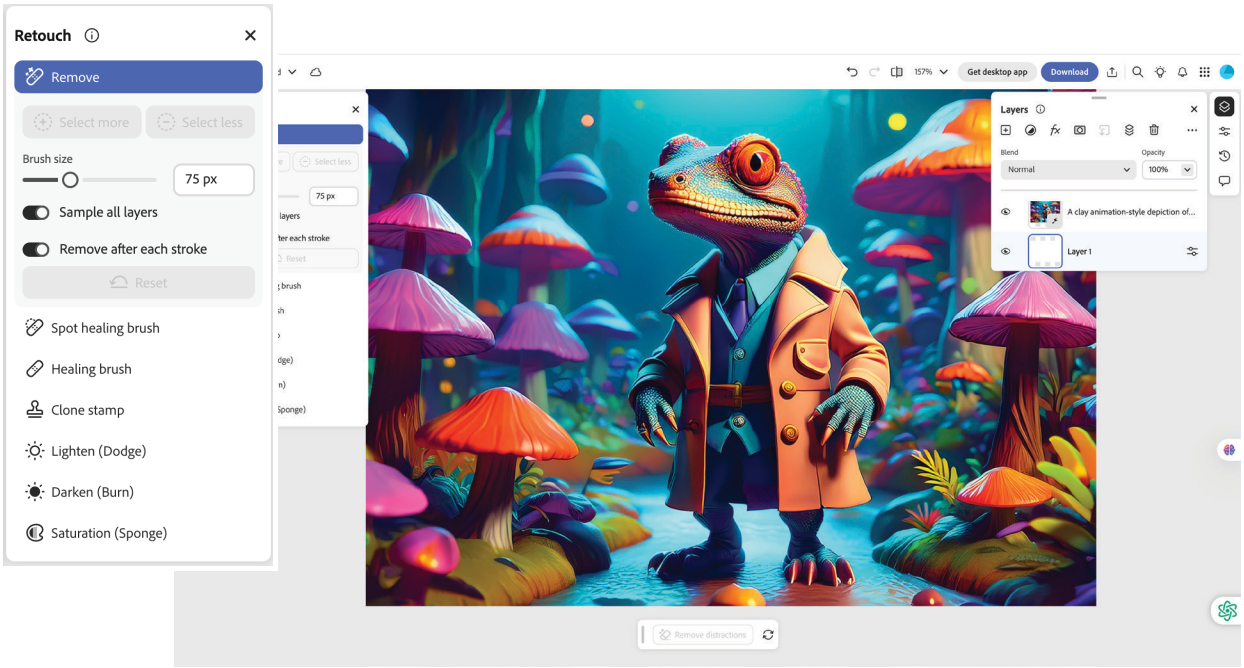

Figure 3. Beginners often struggle to choose between Photoshop's retouching tools (inset detail at left): the Healing Brush, Spot Healing Brush, Clone Stamp, and Remove Tool. Captured June 4, 2025.

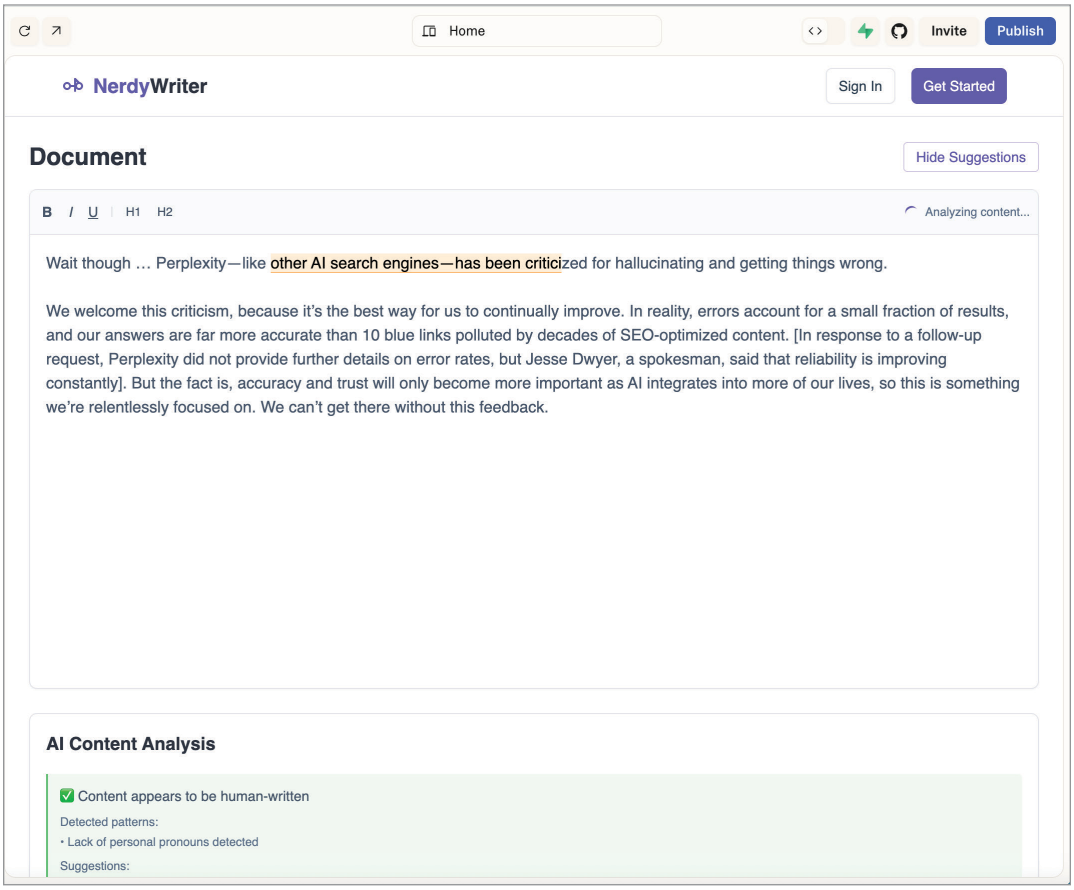

Figure 4 (above). An early prototype of NerdyWriter, a Grammarly-style writing tool built with Lovable.ai using "vibe coding" (Lovable, n.d.). Though AI models are not fully integrated, this mockup helps explore user flows and edge cases. The screenshot shows AI-generated content, inline grammar suggestions, and a scrollable horizontal layout for suggestion cards. It also raises critical UX questions: What if AI analysis fails? Can users regenerate or prompt AI? How do we distinguish between AI and user-edited or pasted content? This approach enables early thinking around product behavior and design, even before full AI integration.

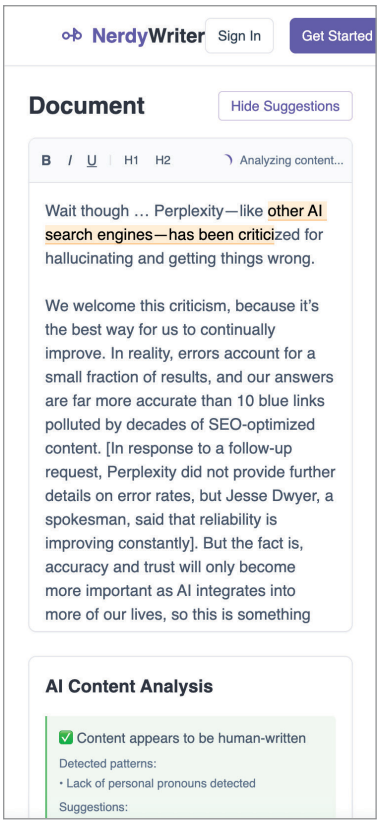

Figure 5 (left). Mobile version of the Nerdy writer. I can see some issues that are easily fixed.

Some effective UX patterns include:

- Inline corrections or quick edits to refine the AI's result.
- Confidence scores with transparency and override options.
- Feedback tagging (e.g., "not accurate," "not useful," "inappropriate tone").

More importantly, this data must return to designers and machine learning (ML) teams, not just sit in dashboards. Feedback is not just about improving AI or achieving key product metrics, but also about the relationship between humans and the systems they rely on (Figure 6). Google addresses this with a "Feedback + Control" resource for designers (Google, n.d.). In addition, there is a tightly coupled relation between explainability, control, feedback, and trust, which designers can keep in mind when designing for human-AI interaction (Figure 7).

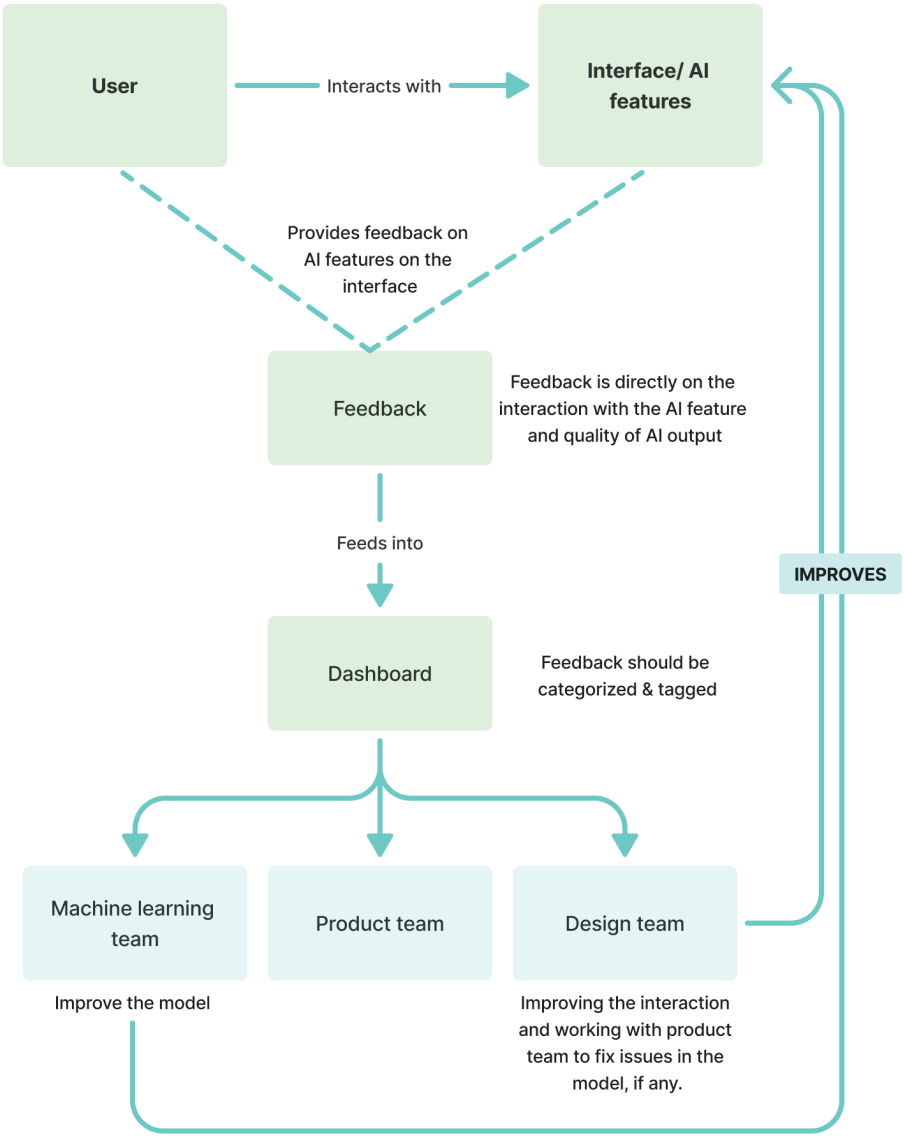

Figure 6. A feedback driven loop where user input on AI features is captured via the interface, categorized in a dashboard, and routed to machine learning, product, and design teams, ultimately improving both the model and user experience.

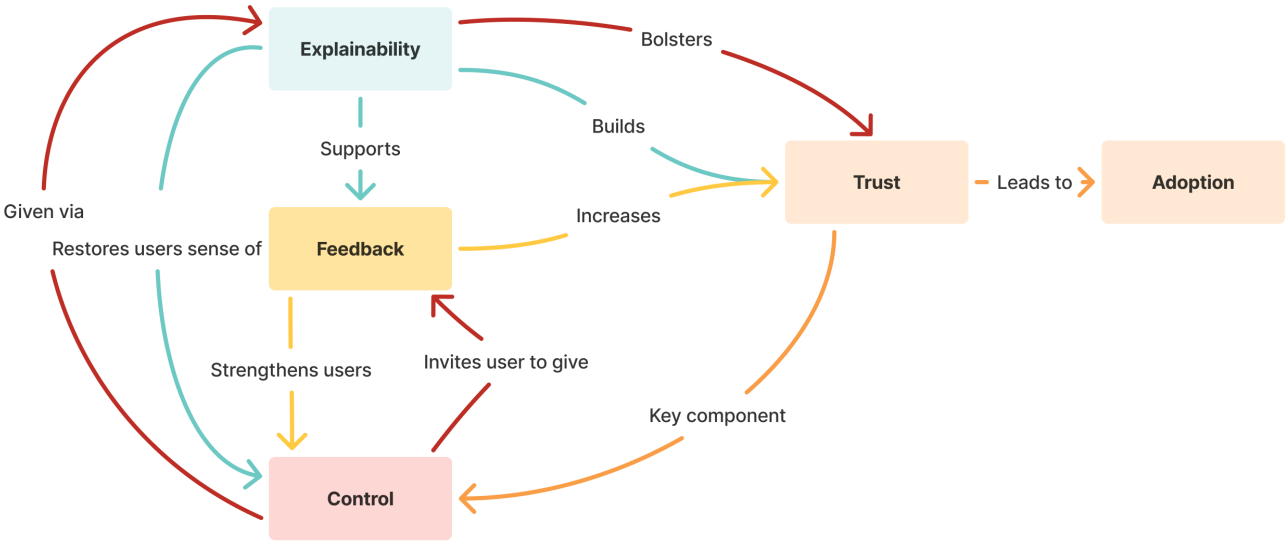

Figure 7. A concept map of the tightly coupled relationship between feedback, user control, explainability, and trust in the human-AI interaction space.

2.7. Raising Data Literacy: Essential for a Seat at the Table

One cannot truly design for something (like AI) until they understand it, at least conversationally. Designers need data fluency to move beyond just skinning the interface of an AI feature and shape how it behaves. It is not about becoming an ML engineer or creating own models; it is about asking the right questions, challenging assumptions, and collaborating as a peer with engineering and data science partners. In my experience, this is fundamental to moving from being a service provider to a strategist.

To collaborate effectively, designers must increase their data fluency in these key areas:

- **Understanding the machine's mindset:** designers do not need to build the models themselves, but understanding the fundamentals of *their training and performance evaluation* is crucial. What kind of data went in? What metrics define "success" or "failure" for the model? Knowing this helps anticipate where the AI might struggle and why it behaves in unexpected ways, as well as design interactions that gracefully handle uncertainty or errors. It takes the AI from a magic black box to something one can reason about. Additionally, I have encountered situations where models cannot be questioned because they are third-party. Even in such cases, it is still valid to ask engineers how they are developing over these models to deliver value to users and what the potentially worst-case scenario is, and then design backward for failing gracefully.
- **Decoding the AI lexicon:** Understand terms like bias, hallucination, fine-tuning, prompt engineering, data preprocessing, feature engineering, and model evaluation (e.g., Sponheim, 2024). Knowing these terms will enable designers

to have more informed discussions with data scientists and engineers, fostering effective collaboration.

- **Using Design tools to your advantage:** Tools like vibe coding or Figma Dev Mode can speed up prototyping, developer handoff, and user testing concepts, which would have been harder before just by using visual sketches. A plugged-in interaction or prototype helps designers while designing, or helps stakeholders understand the depth of user interactions and resolve edge cases faster by simulating actual system behavior.
- **Making sense of multimodal inputs and outputs:** AI is not about text boxes and clicks anymore. It is speaking, seeing, moving, and even feeling. From voice and gesture to augmented reality, mixed reality, and robotics, AI is stepping into the physical world. The more fluent designers are in multimodal interactions, the more ways they can pick up on what users actually want and need. Multimodal design can dramatically improve accessibility and lead to more intuitive, inclusive experiences.
- **Navigating legal and ethical tracks:** Design systems must be informed by AI accessibility standards and legal frameworks like the General Data Protection Regulation (GDPR), the AI Act (in the EU), or Section 508 compliance (in the US). These regulations directly impact multiple aspects of design, such as designing for an AI notice — how much of the AI notice should be shown, how many times it should be shown, and whether a user can remove the notice. Resources are available for designers to remain informed, such as Adobe's Content Authenticity Initiative (2024) and Grammarly's (n.d.) Authorship — including when writing an article like this.

As AI tools become more sophisticated, our ability to question them critically must be enhanced. Continuously learning and sharing is the way to remain relevant.

2.8. Creating a Culture of Co-Creation

A culture that values real collaboration between designers, developers, researchers, and product managers is the foundation for successful AI experiences. Some best practices that design teams can advocate for are:

- Integrated standups and sprint reviews that include design and development.
- Co-writing problem statements at the start of the AI feature journey.
- Designers embedded in ML workflows to explore possibilities early.
- Use of design artifacts (like user flows or journey maps) to align stakeholders across disciplines.

Tools like collaborative whiteboarding (e.g., FigJam, Miro), live design-developer environments (e.g., Figma Dev Mode), and shared metrics dashboards can make alignment visible and trackable, as shown in Figure 8.

3. Conclusion: AI Needs Designers Who Think Bigger

Design practice should be redefined as more data-literate, collaborative, user-centered, and strategy-driven. I have outlined this redefinition as follows (reflected in Figure 9).

1. **Designers are strategists:** AI demands early designer involvement to shape behavior, ethics, and system feedback, not just visuals.
2. **Cross-functional collaboration is essential:** Working closely with engineers, legal, and product teams helps define model behavior, edge cases, and accountability.
3. **Build AI with purpose:** Features must solve real user problems, not just showcase tech. Utility should be prioritized over hype.
4. **Balance assistive and autonomous AI:** Designers must define the line between user empowerment and automation, preserving human agency.
5. **Explainability builds trust:** It is crucial for usability and trust that system behavior be made transparent with tooltips, summaries, and confidence indicators.
6. **Retrofitting AI needs vision and modularity:** Avoid bloated, conflicting experiences by designing flexible, pluggable AI systems.
7. **Feedback loops must evolve:** We must move beyond basic ratings to granular, contextual, and actionable feedback that informs future iterations.
8. **Data fluency is critical:** Designers must understand AI terms, model behavior, legal standards, and multimodal interfaces if they are to design responsibly and collaboratively.
9. **Foster a culture of co-creation:** Embed design in ML/AI workflows, use shared artifacts, and prototype early to align teams and shape outcomes.

Designers are no longer simply producers of interfaces and interactions; they are shapers of behavior, authors of strategy, and advocates for ethics. Having a seat at the table with developers is not about status; it is about influence. A seat at the table ensures that AI products resonate, empower, and adapt.

As we continue to explore generative design, adaptive UIs, and artificial general intelligence (AGI)-informed products, one thing is clear: the future of AI will be *designed*. Moreover, the designers shaping it must be fluent in code, culture, and conscience.

4. References

Amodei, D. (2025, April 1). *The urgency of interpretability*. Darioamodei.com. https://www.darioamodei.com/post/the-urgency-of-interpretability

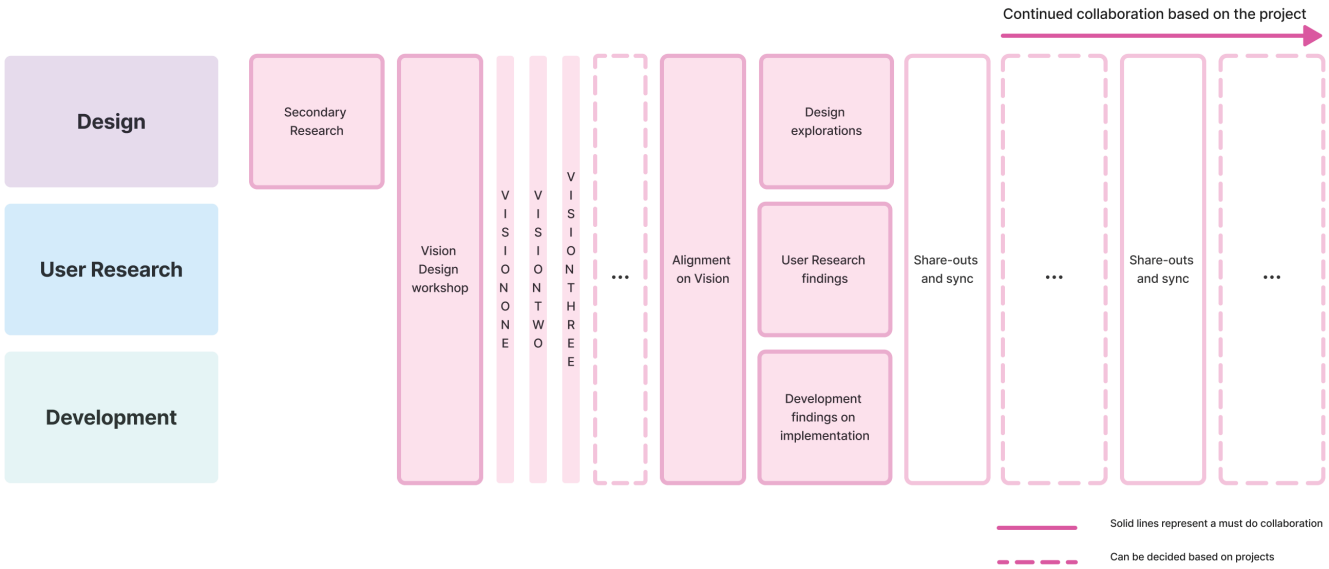

Figure 8. Visualizing a culture of co-creation between designers, user researchers, and development teams. The diagram distinguishes mandatory collaborations (solid lines) from those decided per project (dashed lines), promoting the collaboration between all stakeholders.

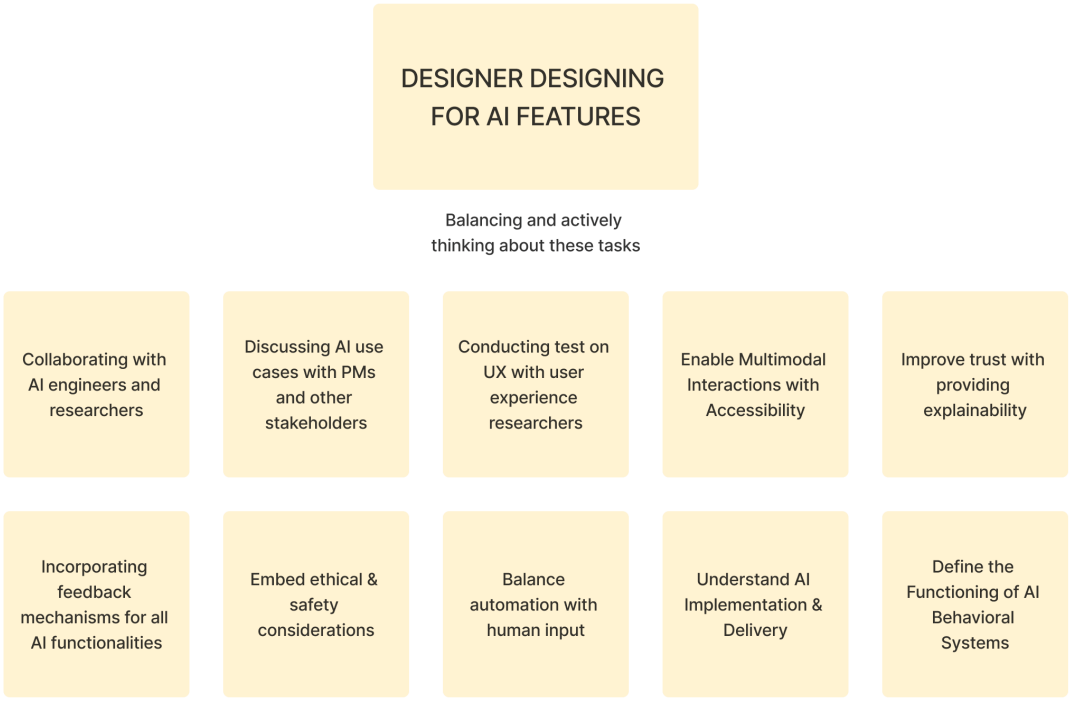

Figure 9. Key tasks and considerations for a designer working on AI features.

Content Authenticity Initiative (2024). *How it works*. https://contentauthenticity.org/how-it-works

Figma (n.d). *Building Figma AI*. https://www.figma.com/ai/our-approach/

Gannon, M. (2025, May 20). *Config 2025*. Figma. https://www.youtube.com/watch?v=3_3wGrLkTCA

Google (n.d.). Feedback + control. *People + AI guidebook*. Google. https://pair.withgoogle.com/chapter/feedback-controls/

Grammarly (n.d.). *Grammarly authorship*. https://www.grammarly.com/authorship

Levin, N. (2024, February 29). *An update on our Make Designs feature*. Figma Blog. https://www.figma.com/blog/inside-figma-a-retrospective-on-make-designs/

Lovable (n.d.). *NerdyWriter project*. https://nerdy-scribe-reveal.lovable.app/

Shakir, U. (2024, June 13). *Here's how Apple's AI model tries to keep your data private*. The Verge. https://www.theverge.com/2024/6/13/24175985/apple-intelligence-ai-model-local-cloud-privacy-how-it-works

Sponheim, S. (2024, January 21). *AI jargon: A glossary of key terms for UX professionals*. Nielsen Norman Group. https://www.nngroup.com/articles/ai-jargon-glossary/

Stanford Institute for Human-Centered Artificial Intelligence (n.d.). *AI index*. Stanford University. https://hai.stanford.edu/ai-index

Author

Syashi Gupta is a product designer passionate about AI design systems, responsible generative AI, and creative cross-functional collaboration. With over a decade of experience designing complex tools across the creative and enterprise space, she currently works at SAP, shaping ethical and scalable AI experiences. Previously, she led design for flagship generative features at Adobe, including background generation and AI-powered object removal in Photoshop and Lightroom. Syashi holds degrees in Information and Communication Technology and Graphic Design, and her master's thesis explored agency, trust, and interpretability in generative adversarial networks (GANs). She thrives at the intersection of usability, technical fluency, and ethical design. Syashi is always curious, and is often seen peering through her optimistically-pessimist glasses to judge new AI releases before they (inevitably) try to take her job.

DISPATCHES FROM INDUSTRY is a *Visible Language* initiative to strengthen the design discipline by bridging industry and academia in the field of interface, experience, and communication design. The column gathers insights from industry practitioners that might inform and guide design practice, while suggesting high-leverage areas for research and scholarship in design. Authors are invited by the editorial board.

DISPATCHES FROM INDUSTRY

The Changing Definition of Designers in the Age of Generative AI

Will Hall

RAIN, New York City, NY, USA and Abu Dhabi, UAE (hello[at]willhall.co)

Abstract: As generative AI transforms the boundaries of creativity and intelligence, the role of the designer is undergoing a profound redefinition. This article explores how design practice must evolve in response — not by resisting AI, but by reshaping how it operates within human systems. Drawing on two decades of fieldwork, product development, and leadership in conversational and multimodal AI, the author proposes four emerging identities for designers: advocate, curator, orchestrator, and mediator of emotion. Each represents a distinct but interdependent response to AI's strengths — and its blind spots. Designers must now move beyond aesthetics and usability to safeguard meaning, ensure ethical alignment, and preserve emotional resonance in systems that otherwise optimize for efficiency alone. The author argues that design's most vital role is to act as a counterforce to algorithmic reduction. In a moment defined by speed, scale, and automation, we must ask not just what AI can do — but what it should do, and for whom. The future will be automated. But it must also be human.

Implications for research: As AI decision-making increasingly requires designers to advocate for human values (Section 3), those designers must recognize what kinds of UI elements and properties impact values in AI-based interfaces, and how. A validated framework could focus designers' attention on crucial aspects. The assertion that AI cannot curate the quality of generated material as well as humans (Section 4) begs investigation, and an understanding of how human curation differs from AI curation might crystallize human capabilities that must be preserved. Systems that coordinate multiple models (Section 5) might compound the uncertainty inherent to AI, and standards for such systems must be further developed. And most directly, the discussion of the emergent area of neuroaesthetic design (Section 6) culminates in a call for frameworks that can strengthen arguments for system transparency over engagement metrics.

Keywords: artificial intelligence; design advocacy; generative AI

@: ISSUE〉 ARTICLE〉 DISPATCHES〉

Cite this article:

Hall, W. (2025). The changing definition of designers in the age of generative AI. *Visible Language, 59*(2), 232–240.

First published online June 25, 2025.

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

Visible Language Consortium:
University of Leeds (UK)
University of Cincinnati (USA)
North Carolina State University (USA)

1. Introduction

I remember it like it was yesterday — the first time I spoke to Alexa.

In the early days of the current AI boom, I often found myself in rooms where I was the only designer. Engineers marveled at the technical breakthroughs, but I was thinking beyond the tech. One day, in a locked-down conference room, we gathered around a very early version of an Echo speaker. For security reasons, the device was encased in a ridiculous shell, bolted shut to hide its true form. I had to sign my life away to be in the room with it. The only thing exposed was the now-iconic light ring.

Curious, I asked my first question:

> "Alexa, what's the best movie of all time?"
>
> Alexa: "Citizen Kane."

That moment stuck with me — not because of the answer itself (after all, *Citizen Kane* is a reasonable option), but because of its apparent certainty. Understanding natural language was one thing, but where was that answer coming from? Who decided *Citizen Kane* was the greatest movie ever? Was it a programmed truth, an algorithmic consensus, or something else entirely?

The answer did not come from a listicle or a panel of critics — the AI conjured it from a multitude of sources. That clunky little smart speaker was the first tremor of a seismic shift, a quiet signal of the tidal wave that would soon reshape everything: generative AI.

That question — where AI gets its answers — has only grown more urgent over time. The typical response is, "Well, it's a black box." But the reality is far more complicated.

At RAIN, where I spent over a decade as chief creative officer, my team built countless voice experiences. Every month, we ran headfirst into foundational questions no one had asked before — questions about intent, privacy, how children interact with AI (we built and launched Alexa for Kids), and the delicate trade-offs between convenience and security. Unlike other areas of design, where best practices and design patterns guide the way, we had no one to follow. We were making it up as we went every single day.

2. The Changing Landscape: From Designer as Definer to Design as Counterforce

When I graduated from college with a Bachelor of Graphic Design in the early 2000s, the dean handed each graduate a "Capital D Design" pin. This notion of a "Capital D Designer" is one that I love. In part, it stands for a definition of design that is bigger

than aesthetics or cold functionality. It calls for a more significant design notion that operates strategically and embraces interdisciplinary approaches. What's not to love?

But I also think it meant "designer as definer." An architectural, "top-down" notion is that we, as designers, define a given product, building, or experience and then allow our work to curate and guide the user accordingly. We *define*.

But in this new era of data, where insights and businesses emerge from the ether stochastically (seemingly randomly from complex probability distributions), the definition of "designer as definer" must evolve. Design can be a counterforce to AI's reductionism: preserving human context, emotional resonance, and ethical considerations in a landscape increasingly dominated by algorithmic efficiency. Through my field research and professional practice over the past two decades, I have identified four emerging definitions of the designer's role in the age of AI, each representing a vital dimension of our evolving discipline.

3. Definition One: Designer as Advocate

Advocacy is an overused word these days, but it takes on new weight in the context of AI. When we think of advocacy in design, we often picture standing up for the user — ensuring accessibility, clarity, and a frictionless experience. But in the age of generative AI, designers must advocate for users and the integrity of decision-making.

We are no longer just crafting experiences; we are shaping how intelligence itself interacts with the world.

One particularly illuminating experience came during my work developing AI assistants in the UAE. During this project, I encountered an AI loan approval system that made decisions based on an unexpected data point: the applicant's phone battery level at the time of application. This system achieved an unprecedented success rate in predicting loan defaults, yet no one could explain the causal relationship.

As the chief data scientist told me: "If you go in with a hypothesis, if you try to force meaning onto the data, you'll taint the system. You'll introduce bias. And the AI will fail."

This counterintuitive insight reveals a critical tension in AI development: the gap between correlation and meaning. AI systems do not "understand" the consequences of their outputs. They do not question whether their recommendations are fair, ethical, or even remotely logical to a human mind. They optimize for outcomes, following correlations detected in data — no matter how absurd, biased, or potentially harmful they might be.

This is precisely where designers must now become advocates. In a world of algorithmic decision-making, the designer becomes the last guardian of human values. Designers are not just tasked with making AI-generated decisions look good or feel seamless. We must interrogate those decisions — ask what is behind them, who they impact, and whether they align with human values.

Are we making products that empower or manipulate? Are AI-driven interactions transparent, or are we just making the black box more palatable? Are we designing with intention, or are we merely refining the illusion of control?

AI will not ask these questions. Engineers will optimize for performance. Businesses will optimize for revenue. However, designers must advocate for the human layer — not by resisting AI but by shaping how it interacts with people while aligning it with business and organizational goals.

This advocacy means making invisible systems visible — ensuring users know when AI is making a decision, providing ways to challenge or override it, and creating feedback loops that keep technology accountable. It also means advocating for interpretability. Just because an AI-generated insight is statistically valid does not mean we should blindly trust it.

In traditional models, we asked technology for answers. Now, as advocates, we must question those answers. Designers are no longer just creators of experiences — we are the last line of defense against systems that optimize without care.

4. Definition Two: Designer as Curator

If advocacy is about questioning AI's decisions, curation is about shaping them into something meaningful. Generative AI is a machine of endless possibility. In seconds, it can generate thousands of images, architectural layouts, UX flows — anything — far beyond what any human could create alone. But raw output is not design. AI does not understand quality. It has no instinct for what feels intuitive, compelling, or meaningful.

Like a museum curator selecting pieces for an exhibition, we do not just arrange what AI produces — we decide what belongs and why. This is part of what Maeda (2019) calls "computational design" — the ability to navigate vast possibility spaces and extract meaningful patterns that align with human values and intentions. Earlier this year at Adobe MAX, I watched a generative AI demo create hundreds of logo variations in seconds. While the tech was impressive, the real skill was not in prompt engineering — it was in the designer's ability to sift through the flood of outputs and identify the five with true creative promise worth developing further.

The curation process involves several layers of judgment that AI cannot replicate:

- Cultural context: understanding which AI-generated solutions will resonate within specific cultural frameworks.
- Intentional selection: identifying which options align with strategic objectives beyond surface-level aesthetics.
- Ethical filtration: removing options that may be technically impressive but ethically problematic.
- Coherence creation: assembling individual elements into holistic experiences that make sense to humans.

This shift redefines our role. It is no longer about dictating a singular vision but about guiding an emergent one. It demands a designer who thrives in uncertainty, embraces iteration, and knows how to elevate the right ideas while discarding the rest.

Because left unchecked, AI does not create meaning — it generates chaos. Patterns without purpose. Options without insight. Randomness that only feels intelligent. Designers must be the human layer that transforms AI's brute force into something intentional, beautiful (if that is a goal), and useful.

Some might argue that AI could eventually learn to curate effectively through reinforcement learning or analyzing human preferences. However, this perspective misunderstands the fundamental nature of curation as a culturally embedded practice. Curatorial judgment relies on contextual understanding that transcends statistical patterns — it requires cultural literacy, ethical reasoning, and an intuitive grasp of emotional resonance that remains uniquely human.

In this new era, we are not just creators but curators. Not defining from above, but shaping and co-creating from within.

This curation role connects directly to our advocacy function — while advocacy ensures AI systems respect human agency, curation ensures that they produce outcomes worthy of human attention. Together, they form a foundation upon which our next roles build.

5. Definition Three: Designer as Orchestrator

Before I unpack this definition, I need to establish two critical premises.

Premise one: technology is moving toward "context-first" experiences. Over the past 75 years, computing has transformed across three distinct waves:

1. Desktop-first (legacy era): Technology was static, tethered to a single location, with hardcoded, linear interactions. Users adapted to machines, navigating with keyboards and mice.

2. Mobile-first (transition era): Technology became portable and responsive, adapting to different screens and touch-based inputs while still requiring explicit user commands.
3. Context-first (emerging era): Technology is now ambient — unbound from devices, seamlessly aware of its environment, and powered by AI-driven personalization. Interfaces are human-first, engaging through natural language, gesture, sight, and motion.

This third wave is not just an advancement — it is a paradigm shift. When technology can see, hear, speak, and predict, the traditional interface becomes just one note in a broader symphony of interaction.

Premise two: the rise of multimodal AI. AI is no longer confined to a single mode of interaction. It is multimodal, processing and generating across multiple dimensions simultaneously:

- Converting spoken language into structured data,
- Generating visuals that align with a user's aesthetic preferences,
- Adapting its tone based on emotional cues,
- Writing code that materializes ideas in real-time.

As technology becomes more ambient and interactions shift to context-first models, AI-driven systems no longer operate in silos. Every touchpoint is infused with multimodal intelligence, requiring designers to embrace complexity and think in systems.

This shift is not just a technical evolution — it is a fundamental redesign of how we shape digital experiences. Designers are no longer just crafting interfaces but orchestrating ecosystems of intelligent systems.

Consider the challenge of a "simple" AI-driven experience:

- One model analyzes user behavior,
- Another generates responses,
- A third ensures brand consistency,
- A fourth handles visual composition,
- A fifth monitors for ethical considerations.

Left uncoordinated, these systems optimize for individual objectives, leading to disjointed, inconsistent, or conflicting experiences. Without deliberate orchestration, AI models operate like musicians playing different tunes — each technically proficient yet collectively chaotic.

In an era where AI makes autonomous decisions, someone must safeguard coherence, human-centeredness, and intent. The designer's role is no longer just about building screens — it is about engineering relationships between AI, interfaces, and users.

True design in this space is systemic, not static. It is about crafting environments where complexity is understandable to the user — where despite the immense intelligence behind the scenes, everything simply works.

Orchestration is a higher-order design skill and requires understanding the relationships between AI, interfaces, and users — and ensuring that what emerges is greater than the sum of its automated parts.

6. Definition Four: Designer as Mediator of Emotion

If designers act as advocates, curators, and orchestrators, then a new challenge emerges: How do we ensure that the products, experiences, and emotions we design remain deeply human? AI is reducing decision-making to cold optimization, but meaning is more than mere efficiency.

During a European research tour with a major luxury house, I encountered what I believe represents the next frontier of AI: neuroaesthetics, the science of designing directly for the brain's emotional circuitry. This emerging discipline moves beyond subjective taste, using brainwave analysis and cognitive modeling to scientifically decode how humans process beauty, art, and design at a neurological level.

Consider my recent project developing a digital fragrance experience. Through neuroaesthetic testing, we discovered specific color combinations and motion patterns that triggered olfactory responses in viewers even without physical scent present.

By mapping these brain-sensory connections, we created digital interfaces that could mimic the sensation of smelling citrus notes or woody undertones purely through visual and auditory cues — quantifiably more effective than traditional marketing approaches. Our AI generated these designed elements based on a neuroaesthetic understanding of the data.

Bentley exemplifies this science-driven aesthetic at scale. Their design studio employs neurological feedback to validate every element in their vehicles — from the precise curvature of dashboard lines to the rhythm of seat stitching patterns — ensuring each triggers the exact emotional cascade that defines their brand experience. What was once intuitive artistry has evolved into rigorously measured emotional engineering.

The implications go far beyond luxury. In healthcare, neuroaesthetic principles could inform AI-powered patient interfaces to reduce anxiety and boost treatment

adherence — using precise visual and auditory cues. In the workplace, environments could be engineered for focus and collaboration through intentional use of form, texture, temperature, and sound. Companies like Kinda Studios are already applying these ideas across industries, with encouraging early results.

Some critics argue that designing for emotional response is fundamentally manipulative. However, I would counter that all design influences emotion — the question is whether we do so consciously, ethically, and in service of human flourishing. As AI systems increasingly mediate our interactions with the world, designers must develop and apply rigorous ethical frameworks for emotional design that prioritize transparency, agency, and genuine well-being over mere engagement.

This new capability carries unprecedented ethical weight. If AI optimizes for efficiency *and* emotion, designers must now answer: Whose emotions are we optimizing for? What experiences are truly meaningful versus merely addictive? Which neural patterns are ethical for us to trigger? The greatest risk is not that AI will fail to understand human emotion, but that it will understand and manipulate it too well, without the ethical constraints that guide human designers. We need new frameworks for responsible neuroaesthetic design that prioritize transparency and human flourishing over pure engagement metrics.

Each of our four definitions progressively builds a more complete picture of design's evolving role:

- As advocates, we protect human autonomy within AI systems.
- As curators, we extract meaning from algorithmic abundance.
- As orchestrators, we coordinate complex systems into coherent experiences.
- As emotion mediators, we ensure these experiences resonate at a human level.

Together, these roles represent a unified framework for design practice that balances technological capability with enduring human needs.

7. Conclusion: Design as a Counterforce to Algorithmic Reduction

The four definitions I have outlined — designer as advocate, curator, orchestrator, and mediator of emotion — collectively point toward a new comprehensive vision of design in the age of AI. They represent not a rejection of technology but a more nuanced integration of it into our practice, one that preserves and amplifies distinctly human values amid increasing automation.

The "Capital D Design" pin I received at graduation carries new weight today. Design is indeed larger than aesthetics or functionality — it has become the essential counter-

force to algorithmic reduction, preserving the nuance, context, and emotional depth that make us human.

In a world where efficiency and optimization have become our default values, designers must stand for something different. We must be the ones who ask not just "how fast?" or "how accurate?," but "how meaningful?" and "for whom?" As AI threatens to flatten experience into whatever can easily be quantified, designers must be the champions of everything that resists such reduction — the ambiguous, the emotional, the cultural, the ethical.

The machines will refine, predict, and accelerate — but only we can ensure that they serve something greater than their own logic. Our job is not to compete with AI, nor to passively accept its influence, but to wield it with intent — to design not just for usability, but for humanity.

Because in the end, AI will not decide what kind of world we live in. We will.

8. Reference

Maeda, J. (2019). *How to speak machine: Computational thinking for the rest of us.* Portfolio.

Author

Will Hall is a serial entrepreneur, creative executive, and presenter on *America by Design* on CBS. A pioneer in conversational AI, he developed the first voice application for Alexa. More recently, he secured over $40 million in funding — backed by sovereign wealth in Abu Dhabi — to build and scale an AI company across the UAE and Asia. Over the past two decades, Will has led creative, product, and innovation work at companies including RAIN, Adult Swim (Pop), Honest, MRY, and Rockwell LAB, with a focus spanning media, tech, AI, and design. He has collaborated with 23 of the Fortune 100 and served on influential boards and advisory councils, including Google's Agency Council, NYCxDESIGN, and NC State's College of Design Leaders Council. He is also a member of the MIT Media Lab Consortium. Earlier in his career, he played a role in the record-breaking Alibaba IPO — the largest in market history. His mom still has no idea what he does.

DISPATCHES FROM INDUSTRY is a *Visible Language* initiative to strengthen the design discipline by bridging industry and academia in the field of interface, experience, and communication design. The column gathers insights from industry practitioners that might inform and guide design practice, while suggesting high-leverage areas for research and scholarship in design. Authors are invited by the editorial board.

DISPATCHES FROM INDUSTRY

The Human Touch(point): Recommendations for Thoughtful AI Feature Design

Sierra Shell

Product Design, SAS Institute Inc., Cary, NC, USA (sierrashell1991[at]gmail.com)

Abstract: This article examines the evolving responsibilities of designers in an era of explosive AI growth. AI is a horizontal technology affecting nearly all industries, and designers must position themselves as stewards of the "human in the loop" to balance technological capabilities with human needs. Three core principles can guide ethical AI design: First, AI features should solve specific user problems rather than being implemented for novelty or marketing purposes. Second, strategic friction can serve as a beneficial design element when deployed at consequential decision points, encouraging users to engage thoughtfully with AI-generated content. Third, robust user feedback mechanisms should be prioritized to ensure continuous improvement based on real-world usage. Generative AI should facilitate — never replace — human expertise to avoid the centralization of ideas and displacement of creativity. Drawing from industry examples, the article demonstrates that successful human-AI collaboration depends not on technological sophistication alone, but on thoughtful design that empowers users as active participants rather than passive consumers of AI outputs. Eight recommendations are provided to ensure that the three core principles discussed are incorporated into a product's design.

Implications for research: This article raises issues that suggest fruitful research areas. The recommendation to compare a potential AI-driven design solution to a non-AI solution (Section 2) suggests the development of a conceptual framework for systematically assessing and justifying aspects of AI specification. Such a framework should be compact enough to integrate into R&D processes for ground-level impacts. The principle of thoughtful friction in AI product design (Section 3) is introduced in relation to frictive AI elements of granular typographic detail such as point size and positioning. Empirical research could determine which typographic and interface characteristics provide friction for contemporary users, to what degree different options do so, and what thresholds exist for users' emotional and cognitive responses. Further research could examine whether such frictive elements effectively facilitate trust calibration for potentially

@: ISSUE〉 ARTICLE〉 DISPATCHES〉

Cite this article:

Shell, S. (2025). The human touch(point): Recommendations for thoughtful AI feature design. *Visible Language, 59*(2), 241–251.

First published online May 22, 2025.

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

Visible Language Consortium:
University of Leeds (UK)
University of Cincinnati (USA)
North Carolina State University (USA)

erroneous AI-generated content. Finally, the recommendation to integrate feedback mechanisms into AI-driven products (Section 4) suggests research aimed at understanding the current state of AI feature feedback collection in popular products, both in terms of mechanism and frequency.

Keywords: AI design; design industry; ethical AI; friction; product design; responsible AI

1. Introduction

In their viral lecture *The AI Dilemma,* Center for Humane Technology co-founders Tristan Harris and Aza Raskin (2023) describe how the introduction of a new technology necessitates a new class of responsibilities. They are not the first to suggest as much, and digital designers are familiar with the emergence of novel and disruptive tech. As a horizontal technology, AI is already influencing almost every industry and digital service. Companies are quick to evangelize their "GenAI" (generative AI) strategies, sometimes incorrectly labeling broadly defined AI features as generative AI to leverage its current popularity. Disentangling buzzwords and hype from the potential substantive improvement AI offers is the job of technologists collectively. However, for the designer, AI escalates both the power and consequences of design choices.

The role of designers in the creation of AI features, products, and services is to be stewards of the "human in the loop." To that end, product designers, visual designers, user researchers, and others in the field of human-computer interaction must thoroughly understand a user's problem and how AI features specifically and uniquely address it. They should use friction intentionally to highlight important decisions and other moments that require careful human consideration. Finally, to minimize organizational risk, designers should include a user feedback mechanism alongside GenAI features to ensure they understand users' experiences. Most importantly, user feedback should be prioritized over marketable but superficial new features.

2. AI Features Should Solve a User's Problem

Problem definition is nothing new to the world of design. Experts like those in the Nielsen/Norman Group have shared guidance on using "how might we" statements to ideate on the right problem (Rosala, 2021) and crafting user need statements when defining the problem using design thinking methodology (Gibbons, 2019). A designer should seek, above all, to help a user solve their problem in the simplest, most intuitive way possible. The first step in doing so successfully is accurate problem definition.

Companies often seek to increase the marketability of their products and services by implementing the novel technology *du jour.* Of 8,000 business leaders surveyed for Cisco's AI Readiness Index, 98% feel increased urgency to use AI within their organization ("Cisco 2024 AI Readiness Index," 2024). However, feeling pressure to utilize AI does not mean its implementation will inevitably be helpful. If used unnecessarily, it can introduce needless risk, requiring compliance with complex laws and standards. By understanding a user's problem, we can ensure an AI-driven solution is thoughtfully crafted and demonstrably better than one that does not employ AI.

A simple exercise can allow designers to feel confident that their AI use is thoughtfully addressing a user's problem: brainstorm a solution or task flow that does not incorporate AI at all. What might have been created 10 years ago without today's machine intelligence? Answering this question can help ensure that designers are able to clearly, quickly, and easily justify an AI-driven solution. Note that a designer's justification should be tied to the problem space itself, not to general characteristics of AI. In other words, the description of AI's advantage should not consist only of broad characteristics of AI.

- **Recommendation 1:** Make sure an AI-driven solution is clearly and demonstrably better than one without AI by creating and comparing to a non-AI design solution.

There is no doubt that AI has the power to improve product or service experience when used carefully. Such features can help a human user's expertise shine, or remove burdensome barriers. Consider GenAI. Generative AI is a subcategory of broad AI that creates new content, whether text, images, videos, reports, data, or other assets. GenAI is often successful when implemented in two contexts.

First, GenAI is a great tool for introducing productivity improvements like summarization, automation, and basic analysis. AI can accomplish mundane organizational tasks and tackle low-hanging analytical fruit. Thus, a user with deep expertise has more time to use their special skillset. GenAI can provide a shortcut that helps a user take advantage of their own expertise earlier in their process, removing the busy work.

Second, GenAI can grant access to skillsets a user may not otherwise have. For example, a small business may use Midjourney to create images or content for a marketing campaign. Without time or expertise in visual design or illustration, the business may have to rely on stock photos to accompany their product or service, getting lost among their competitors. But the specificity with which an image can be generated could empower them to stand out when putting their product, message, and brand out in the world.

- **Recommendation 2:** Use GenAI features to facilitate and complete mundane tasks, allowing users to leapfrog to tasks requiring their own expertise.

- **Recommendation 3:** Use GenAI features to provide access to expertise users would not otherwise have.

In these examples, GenAI *facilitates* a user's expertise. It never replaces it. AI features that override a user's talent threaten the human in the loop, displace creativity, and lead to a centralization of ideas. The use of AI as a means of removing human ingenuity in favor of a quick and uninspired version of a human-created asset should be avoided whenever possible. The centralization of output is noticeable. Education researchers who examined student work noted telltale signs of the use of chatbot assistance, including increased use of Latin terms and consistent, unnecessary vocabulary (Morrison, 2024). The prevalence of its use has created a distinct essay style detectable by many teachers and professors.

Furthermore, humans are still much better at producing creative assets for a given context. AI does not have the intuition to understand how an asset fits into a larger project or how it would be received by diverse end users. Lived human experience and a nuanced understanding of the problem area give human creators an advantage over machine intelligence when context matters most.

The use of AI systems comes with real-world costs for organizations (and for the planet), so such features should be introduced carefully. Especially with regulations like the EU's AI Act in play, using AI necessitates a thorough understanding of its risks, potential harms, and compliance burden. AI should not be thrown at every problem.

However, sometimes designers have no ability to influence a product or service's requirements. Incorporating AI may be a company's goal. When required to include an AI feature where there is no distinct need, one should consider how it can fit seamlessly into the interface. The feature should be placed appropriately in the visual hierarchy and treated as a help mechanism, not something created for a demo or marketing materials. Thoughtfully considering and designing AI features, and clearly communicating their specific value-add, will serve any design professional well.

- **Recommendation 4:** When encouraged to incorporate an AI feature for its own sake, think of it as a help mechanism tucked neatly into an appropriate place in the visual hierarchy. Features need not be displayed at the top level for marketing purposes.

3. Friction Can Be a Friend

For many designers, friction is a dirty word. Friction generally refers to anything within a user's experience that makes it more difficult to complete a task. For example, consider a website whose call-to-action button lies just below the scroll line. A user must scroll

slightly in order to click the button as desired by the site's owner. This is annoying friction — the additional requirement of scrolling likely reduces the frequency with which users complete the task (here, clicking a button).

Sometimes friction is employed intentionally because it works. Consider how frequently a user is bombarded with a dialog upon arrival to a retail site, requesting personal information in exchange for a discount. Exiting from these windows is arduous, as they generally use a tiny exit icon that is difficult to click. The window stops a user in their tracks, but can result in a higher number of individuals providing their information. While not exactly a dark pattern, the frictive UI element is designed to stop a user from accomplishing their task for the sake of the site's owner, not their user.

In the world of AI tools and features, friction finds a fundamental purpose. Helping users stop and consider their actions when using AI tools can support a healthier relationship with AI. This is especially true when it comes to decisioning tools. For companies that create, deploy, and manage their own AI models to enhance decision making, friction is key for safety and risk reduction.

As an example, SAS Institute Inc. recently published an open-source guide for trustworthy AI governance called the Trustworthy AI Life Cycle Workflow (hosted on the SAS GitHub page; SAS Institute Inc., 2025). The workflow (Figure 1) outlines steps for evaluating and deploying a more trustworthy AI system, based on the U.S. National Institute of Standards and Technology's (NIST) recommendations, standards, and best practices. Using the guide, teams can produce documentation outlining how the organization has done its due diligence to ensure a model is fair and minimizes harm.

After its initial release, internal reviewers in the Fraud domain at SAS were consulted for feedback. While the workflow was generally well-received, several comments described "barriers" to a seamless operation, and remarked how "extensive" and "big" the workflow was. Reviewers understandably wanted to streamline and shorten the process, making certain steps easier for users and decreasing the overall time to model deployment.

The team spent significant time discussing this feedback. However, many of the included steps erected intentional barriers; for example, rather than simply asking if a model uses personally identifiable information (PII), the workflow outlines steps to identify such variables, asks users to explore whether their inclusion is necessary, requires documentation describing why they must be included in the model, and requests a description of any resulting risk. To simplify the workflow, it would be possible to automate the step (throw a flag if "yes" is selected) and assume the model is high-risk and should be subject to the appropriate regulations. But requiring a user's active thought may result in fewer unnecessary uses of sensitive information.

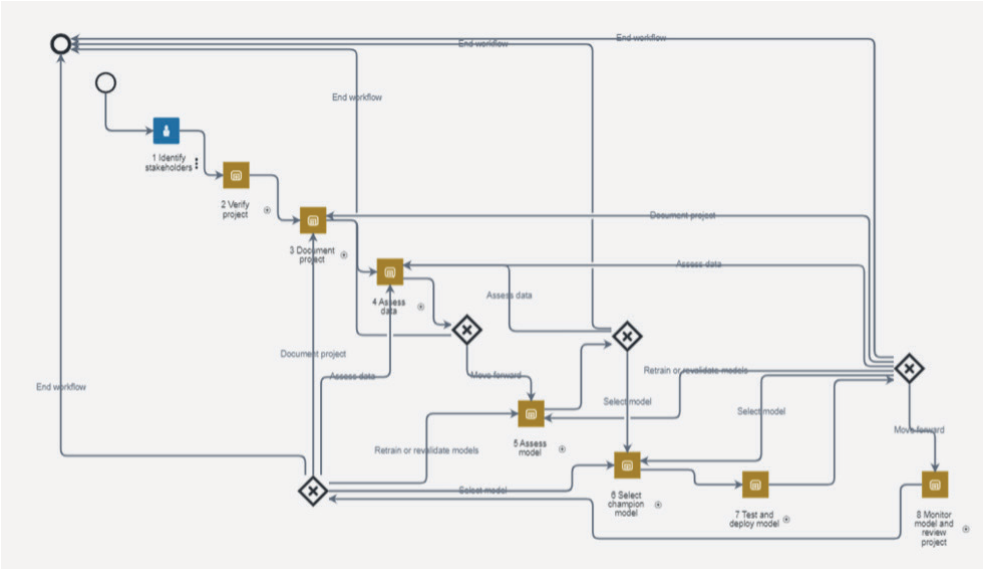

Figure 1. The Trustworthy AI Life Cycle Workflow is based on a flow of the same name orchestrated in SAS© Workflow Manager [low-resolution product export, only intended to demonstrate complexity].

Friction can and should be employed beyond the decisioning space. One might consider adding friction during points of high consequence for the user. Think about intentional "pause points" early on when creating a task flow. Areas utilizing AI assistance deserve special scrutiny. Points of high consequence include actions that affect multiple people or the organization as a whole, like deploying a model. They may also include a final revision or submission step, or communication with a large group of individuals. For instance, certain social media sites have attempted to reduce harm by asking users to consider kinder messaging when posting a comment. Some companies like TikTok are creating content reminding the user that they have been online for a long time and should go "touch some grass." Areas of high consequence are moments where a human's "humanness" matters, or where results of an incorrect or undesirable outcome would have real-world consequences.

- **Recommendation 5:** Add frictive elements during decision points of high consequence or those involving sensitive data. Adding friction around AI assistance features is particularly important.

As it pertains to GenAI, friction often includes interrupting the presentation of generated content with a reminder that the content is created by an AI model. Indeed, frictive features may need to be more interruptive than users are accustomed to. Google's AI Overview description, for example, labels its content with a title indicating an "AI Overview," yet the font size of the label is smaller than that of the entry (Figure 2). The AI Overview appears by default above other results, at the top of the content hierarchy. AI Overview content often features highlighted text, bold text, or both.

All of these aspects create a seamless interaction where a user's eye is drawn to the summary, even if they prefer standard search content. Such is the power of Google's frictionless design.

Instead, they could improve their implementation by adding thoughtful friction — AI-generated text could be completely hidden at first until a user engages with a button, rather than providing several enticing lines that beg the user to reveal the rest of the content. The overview could be an opt-in feature, turned on by interested users intentionally. At the very least, Google could remind readers that content generated by AI can contain mistakes, so to take caution. Currently, the only warning they include is to remind their audience that "Generative AI is experimental" in small font below the expanded summarization (Figure 3). The term "experimental" can be interpreted in many ways and does not explicitly inform a user that content may be wholly incorrect.

Google follows a good heuristic — indicators of AI-generated content should be placed such that a user's eye moves to the label before the content itself. The mind should be primed to consume AI-generated content before it is consumed, as humans are likely to interpret or rely upon it differently. For left-to-right text, this suggests a warning placed

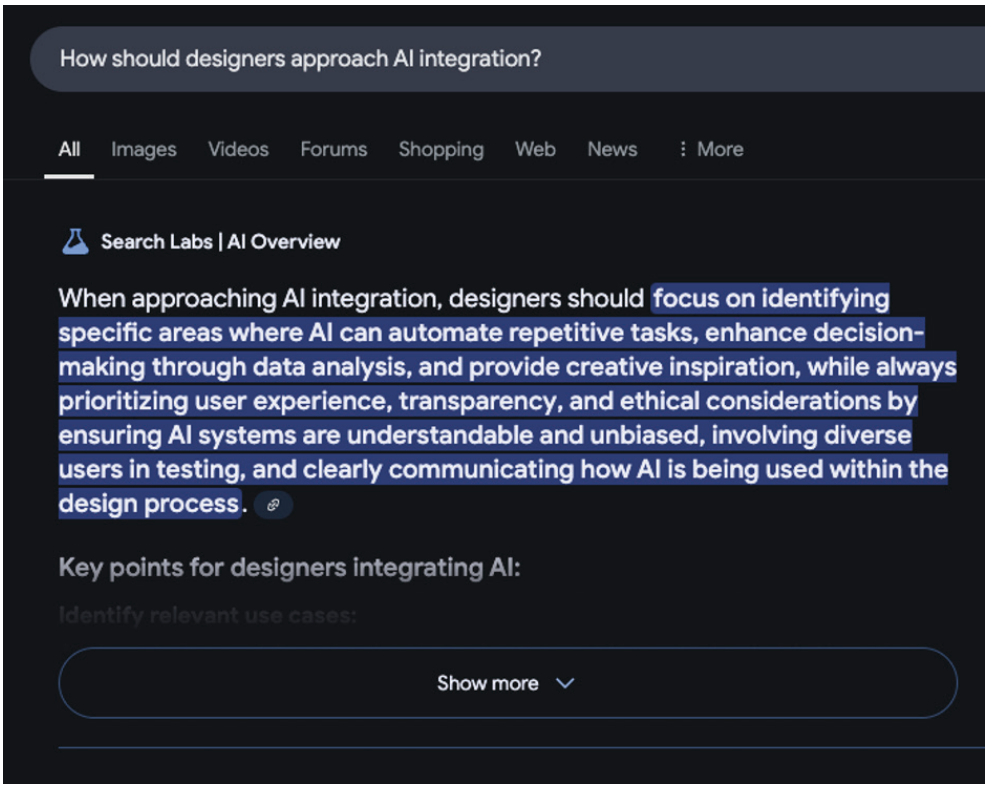

Figure 2. Google's AI Overview preview appears above traditional results and features bold text and highlighting (captured February 25, 2025).

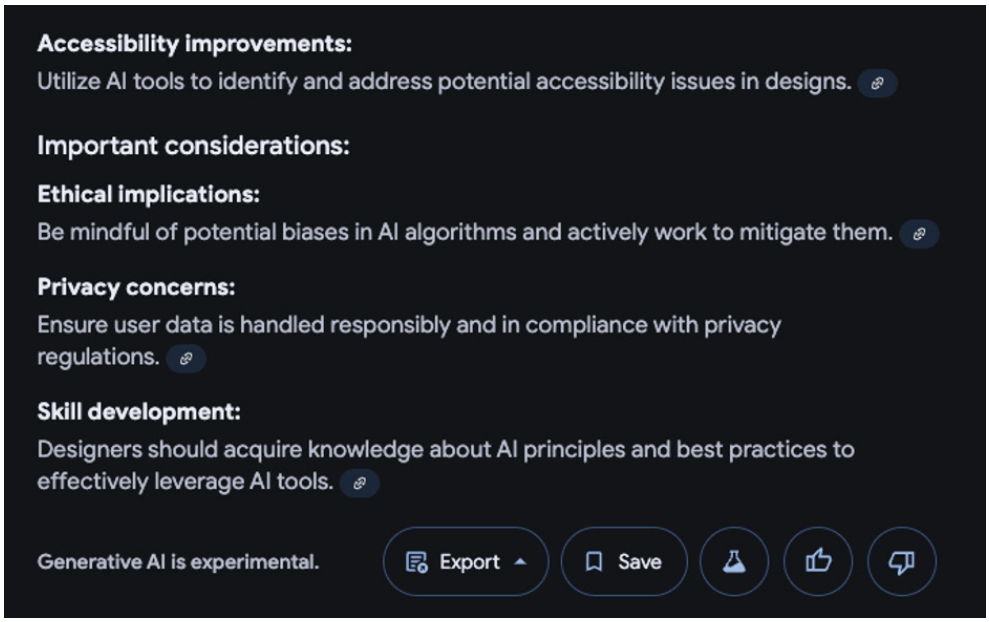

Figure 3. Google provides users with the vague warning that "Generative AI is experimental" at the end of the entry (captured February 25, 2025).

at the top left of the summary. However, Google's small font and hierarchical placement of the summary itself counteract the effect.

- **Recommendation 6:** Include a warning about the potential inaccuracy of AI-generated content, and design the area such that the eye will view the warning message before viewing the content.

The most important step is already an established best practice, though often ignored — a user should always know when interacting with or consuming content created by an artificial intelligence. However, to be a good steward of their users, designers should go beyond the bare minimum and build in time to consider how they would like to make use of AI in the context of their product or service.

4. Users Can Keep Designers in the Loop

AI tools often produce weird results, most noticeably GenAI systems. Many people can recall a funny story or example of when GenAI got something terribly wrong. The most advanced models still have limitations, even as AI is improving exponentially. Designers should plan for this reality from the outset.

It is impossible to foresee all of a GenAI feature's strange and unexpected results, but users can be relied upon to help keep product teams in the know. To this end, every AI feature should be accompanied by a clear user feedback mechanism *in situ*. When a

user encounters generated content, they should have an immediate, accessible, obvious way to flag content, elaborate on pain points, or indicate when the AI has misunderstood their request. These feedback mechanisms should be baked into the design in the earliest stages.

- **Recommendation 7:** Every AI feature should be accompanied by an obvious user feedback mechanism in close proximity.

Relying on users helps ensure that designers are, themselves, also "humans in the loop." Many organizations integrating third-party base models have very little or no control over how such tools interact with users. For example, the voice and tone of generated text is baked in.

True, certain techniques can be used to modify output: one technique gives access to proprietary information a base model does not know using retrieval-augmented generation (RAG), allowing a model to include organization-specific information in a search context. Another technique recommends adding guiding text to help a user when engineering their prompt. Yet another includes filtering, applying a ruleset to system output. But generally speaking, organizations integrating with existing tools might be in the dark about all the ways the tool can behave. Meanwhile, a user still associates the integrated model's behavior with the product or service's brand, regardless of its origin. This is why feedback mechanisms are not a "Phase 2" feature. Without the ability to substantially modify a model's behavior, user feedback becomes the most valuable source of information about how the GenAI feature performs in real-world contexts.

While an organization may not have control over the exact outputs of an integrated third party LLM, its employees do oversee the product or service itself. Designers must advocate for user feedback mechanisms over potentially more marketable, dazzling features that demo well. While new AI capabilities may attract initial attention, only their effectiveness, reliability, and accuracy will keep people happy.

One way to help convince stakeholders to prioritize modifications driven by user feedback is to formalize their inclusion. Product teams could integrate a tag or flag into their feature management system that shows the requested change came from user feedback. These are high-priority items. Users' understanding of a product or service can impact retention, sales, branding, everything. Their opinion should be taken seriously.

- **Recommendation 8:** Product teams must prioritize changes driven by user feedback over dazzling and easily marketable new features. They can do so through formalizing the user feedback integration process.

5. Closing the Loop

In the rapidly evolving landscape of AI, the organizations that thrive will not necessarily be those with the most advanced technology, but those that maintain the most effective human–AI collaboration. AI may not be the best solution to the problem at hand, and it can introduce potential risks and burdens like inaccuracy or compliance requirements. The three core principles outlined here can guide ethical AI design. (1) AI features should be introduced thoughtfully and intentionally, and only where their inclusion is better than a design without them. (2) Friction can encourage users to consume GenAI content thoughtfully and critically, especially when introduced in moments of high consequence. And (3), by including mechanisms for collecting feedback *in situ*, designers can understand how users perceive the AI feature and, in turn, how the feature informs customers' understanding of the organization's brand and reliability.

In summary, to ensure AI features are included with a high impact and minimal harm, be mindful of the eight recommendations, reiterated here:

1. Make sure an AI-driven solution is clearly and demonstrably better than one without AI by creating and comparing to a non-AI design solution.
2. Use GenAI features to facilitate and complete mundane tasks, allowing users to leapfrog to tasks requiring their own expertise.
3. Use GenAI features to provide access to expertise users would not otherwise have.
4. When encouraged to incorporate an AI feature for its own sake, think of it as a help mechanism tucked neatly into an appropriate place in the visual hierarchy. Features need not be displayed at the top level for marketing purposes.
5. Add frictive elements during decision points of high consequence or those involving sensitive data. Adding friction around AI assistance features is particularly important.
6. Include a warning about the potential inaccuracy of AI-generated content, and design the area such that the eye will view the warning message before viewing the content.
7. Every AI feature should be accompanied by an obvious user feedback mechanism in close proximity.
8. Product teams must prioritize changes driven by user feedback over dazzling and easily marketable new features. They can do so through formalizing the user feedback integration process.

Designers, as stewards of the human in the loop, must ensure that users remain empowered participants in this collaboration, not passive consumers of AI outputs. Through thoughtful application of AI, appropriate friction, and robust feedback

mechanisms, product teams can create AI experiences that truly enhance human capabilities rather than merely dazzle or, worse, displace them.

6. Acknowledgments

Thank you to Rajiv Ramarajan, Cheryl Coyle, Zoe Rackley, and Matthew Peterson for assistance and review.

7. References

Cisco 2024 AI Readiness Index. (2024). *Cisco.* https://www.cisco.com/c/m/en_us/solutions/ai/readiness-index.html

Harris, T., & Raskin, A. (2023, March 24). *The AI dilemma* [audio]. Center for Humane Technology. https://www.humanetech.com/podcast/the-ai-dilemma

Gibbons, S. (2019, March 24). User need statements. *NN/g.* https://www.nngroup.com/articles/user-need-statements/

Morrison, N. (2024, March 27). The tell-tale signs students are using ChatGPT to help write their essays. *Forbes.* https://www.forbes.com/sites/nickmorrison/2024/03/27/the-tell-tale-signs-students-are-using-chatgpt-to-help-write-their-essays/

Rosala, M. (2021, January 17). How might we questions. *Nielsen Norman Group.* https://www.nngroup.com/articles/how-might-we-questions/

SAS Institute Inc. (2025). SAS trustworthy AI life cycle. *GitHub.* https://github.com/sassoftware/sas-trustworthy-ai-life-cycle

Author

Sierra Shell is a technologist and UX designer with experience in trustworthy AI design and governance, ethical design strategy, and data visualization design. She has nearly a decade of experience creating complex enterprise products, and currently works with a small team focused on expanding SAS's already robust trustworthy AI capabilities. She received her Master's of Digital Technology Policy from University College London, where she partnered with the British Standards Institution conducting and publishing research on the future of responsible AI standards. Her work required that she develop a deep understanding of common hurdles and disincentives organizations face in implementing AI governance strategies. This has led to a special empathy for organizations and users adopting AI tools and models. As a result, she has a drive to design streamlined solutions that facilitate effective governance and address business concerns intuitively and responsibly.

DISPATCHES FROM INDUSTRY is a *Visible Language* initiative to strengthen the design discipline by bridging industry and academia in the field of interface, experience, and communication design. The column gathers insights from industry practitioners that might inform and guide design practice, while suggesting high-leverage areas for research and scholarship in design. Authors are invited by the editorial board.

2023

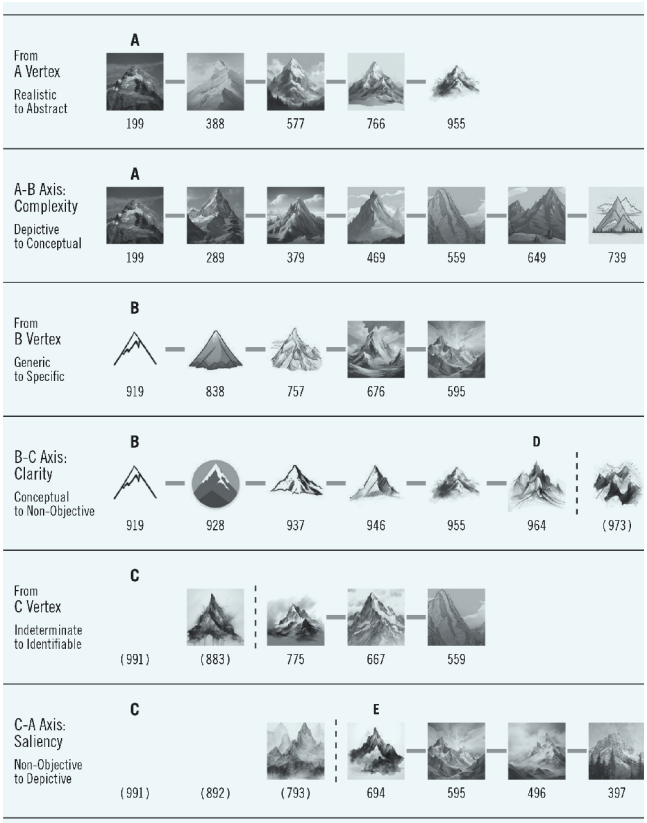

ses modeled on artificial intelligence: it needs to consi
ween people and machines as inscribed in goal-seeking
ded in moods, emotions, social communications and v
erstand a machine in a lab in order to get a ticket (to pl
r and to prove one's level of skill), is different from tryi
chine in a railway station fast enough so as not to miss
till is to try to operate a machine properly when one is
(1972) reports that "the [cognitive] performance of a s
ther source of evidence of the effects of danger" (Badd
The probability that a soldier will use his rifle effectivel
is considerably less than in training. Walker and Burk
, after the battle of Gettysburg in the American Civil W
d of the muzzle-loading rifles were found to have been
re times without being fired, and one had been loaded
s without being fired once. Inadequate training was bla
orts suggest that even with fully trained troops perfor
worse than performance in training (Egbert et al., 1957,

of psychological research it has been possible to produ
r situations which impair performance. In one experim
as simulated in a flight, using service men as subjects.
ff, they were told that the landing gear was faulty and t
themselves for an eventual crash landing at sea. They

1999

2023

a human researcher to develop new hypotheses and theories.

ated question generation: *Use a machine learning system to generate questions based on an academic paper's content. This can then be used by a human researcher to craft further questions and explore the paper's topic in more depth.*

ated essay writing: *Use a machine learning system to generate an initial essay based on an academic paper's content. This essay can then be used by a human researcher to further refine it and ensure accuracy.*

ated reference generation: *Use a machine learning system to generate a list of references based on an academic paper's content. This list can then be used by a human researcher to quickly locate and find relevant references.*

There is a level of specificity to the AI's suggestions, which is useful, but the machine makes an assumption that humans provide or know the data source.

I reflect on the process I am working through and wonder about the knowns and unknowns of the technology I'm interacting with. I think about what I might learn from it and contemplate what it might learn about me. I'm envious of the machine's power to systematically perform tasks and crunch through masses of data with speed and vigor, and its ability to digest and store masses of content. Our learning styles seem polar opposites—reading and auditory processing aren't my strengths. I can't fathom the thought of reading through masses of literature because I know it wouldn't be retained or retrievable unless it were contextually relevant to something I was doing. Knowing I am a visual and kinesthetic learner, I continue tinkering with the OpenAI interface, and one thing that strikes me from playing with OpenAI and ChatGPT is the confidence it alludes to in its delivery. The responses sound good, irrespective of whether the content is factually correct. Maybe I can learn something from its unflappable confidence, as I tend to be self-deprecating. When I ask it to offer suggestions, or improve the clarity of one of my paragraphs, it removes that uncertainty

Three articles from *Visible Language* mentioning or centered on AI, with a 24-year gap:

Peterson, M., Anderson, A. L., Rondinelli, K., & Armstrong, H. (2023). The Pictorial Trapezoid: Adapting McCloud's Big Triangle for creative semiotic precision in generative text-to-image AI. *Visible Language, 57*(3), 32.

Ellerton, W. (2023). The human and machine: OpenAI, ChatGPT, Quillbot, Grammarly, Google, Google Docs, & humans. *Visible Language, 57*(1), 48.

Frascara, J. (1999). Cognition, emotion and other inescapable dimensions of human experience. *Visible Language, 33*(1), 77.

inside

back

cover

Contents