With Visual Integrity and Care: A Framework for Mixed Methods Research on Visual Social Data

Nina Lutz
University of Washington
Seattle, Washington, USA
ninalutz@uw.edu

Joseph S. Schafer
University of Washington
Seattle, Washington, USA
schaferj@uw.edu

Priya Dhawka
University of Washington-Seattle
Seattle, Washington, USA
pdhawka@uw.edu

Phil Tinn
SINTEF
Oslo, Norway
phil.tinn@sintef.no

Kate Starbird
Human Centered Design & Engineering
University of Washington
Seattle, Washington, USA
kstarbi@uw.edu

Abstract

The internet is becoming increasingly visual, but social computing research and methodological training has relied heavily on textual methods. Methodological innovation is needed to study visual social data, including problematic information (mis- and disinformation, propaganda, hate, AI slop, etc). Contending with this, we present a framework for conducting grounded, interpretive, computationally supported, mixed-method research on collections of visual social media data. We developed this framework while grappling with the ethical, logistical, and methodological challenges of conducting in-depth analysis of potentially harmful visual content while caring for our research team. We document our framework components of visual grammars, human analysis, and computationally supported analysis with an umbrella commitment to care and its use in three empirical case studies. We also provide recommendations and implications for the HCI community in embracing training in and the advancing of visual methods and research, including a sensitizing concept of visual integrity.

CCS Concepts

• **Human-centered computing → HCI design and evaluation methods**; **Collaborative and social computing design and evaluation methods**.

Keywords

Methodology, visual research, mixed methods, social data

ACM Reference Format:
Nina Lutz, Joseph S. Schafer, Priya Dhawka, Phil Tinn, and Kate Starbird. 2026. With Visual Integrity and Care: A Framework for Mixed Methods Research on Visual Social Data. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 24 pages. https://doi.org/10.1145/3772318.3791623

CHI '26, Barcelona, Spain

ACM ISBN 979-8-4007-2278-3/26/04
https://doi.org/10.1145/3772318.3791623

1 Introduction

As the internet becomes increasingly visual, so does the content people see and interact with the most [97, 155], including "problematic information" [78] like propaganda, scams, mis- and dis-information, violent imagery, online hate, and more. However, social media research has largely relied upon textual analysis [132], leaving a methodological gap in equipping researchers to rigorously study the visual content shaping our on and offline lives.

Past work has examined how visual data is "unwieldy" [132] compared to textual, requiring greater computational bandwidth and novel analytical frameworks [69, 132]. Closing this gap (which pre-dates modern social media research [114, 126, 157]), has been identified as an imperative priority for studying how humans are communicating and meaning-making in digital environments [19, 69, 132, 172]. Scholars have surfaced this gap in researching problematic information [78] where the visual is increasingly prominent and often intertwines extreme and mundane content in ways that often necessitate detailed, human analysis [17, 39, 109, 132]. This final point poses challenging tradeoffs given handling and analyzing problematic visual content can harm researchers seeking to study it [1, 72, 175].

Bridging this gap will take new methods, frameworks, and training in visually-centered ways of conducting research. The open questions and challenges in doing so are multi-dimensional and numerous: computational, methodological, logistical, and ethical. Our contribution — a framework for conducting visual research on social media and other digital trace data — wrestles with this.

Influenced by our positionality as researchers of problematic information [78], we provide a flexible but systematic framework for conducting investigations on visual social data, including emotionally salient events and contexts that could harm researchers [3]. Our epistemological orientation leans heavily on interpretivist perspectives [146] and we share a belief that deep, qualitative engagement with data by humans is of tremendous value for deriving insights from social data. We seek to contend with how researchers like ourselves may balance preserving deep qualitative analyses while conducting cognitively and emotionally taxing inquiries, offering recommendations for making efficient use of our labor and protecting ourselves. Additionally, as our research is collaborative and we also identify as mentors and educators invested in the training and sustaining of future researchers, our framework also provides pedagogical insights.

The framework this paper describes was developed through our experiences grappling with many ethical, methodological, and logistical challenges while conducting a year long study of visual anti-immigrant rhetoric. Unable to find a framework that suited our needs, we designed and developed the technological, human, and methodological research infrastructure we needed to support our analysis.

As that research project evolved, we realized our approaches could be applied to other studies. We have now adapted this framework in two additional studies of problematic visual information in different domains: Jesus AI Slop images on Facebook and Telegram imagery of the Russia-Ukraine war. This paper aims to facilitate a translation of our approach for researchers contending with similar challenges in studying big visual social data, both in and outside of problematic information. We ground our contributions in relevant literature (Sec 2), document our framework (Sec 3), illustrate that framework across three empirical case studies (Sec 4), and provide implications for researchers (Sec 5).

In total, we contribute:

- Our novel framework for computationally supported, human-centered research on visual social data.
- Pragmatic advice and best practices for conducting this type of work, particularly in teams mentoring junior researchers.

By doing so, we aim to empower HCI (and other) researchers to study increasingly visual digital spaces, and inform discoveries about and interventions to improving human experiences in them.

2 Related work

In this section we first provide a review of visual research in its formations in art and media studies and how it applies to HCI (2.1). Next, we examine limitations of "big" social data analyses in visual contexts (2.2). Lastly, we examine recent methodological innovations and criteria for a methodological contribution in HCI (2.3).

2.1 Visual research and methods

There is no one definition of "visual research". For our work we employ definitions from visual sociology (e.g. [67, 128, 145]) and media studies (e.g. [25, 135]). These definitions position visual objects as socially constructed and focus on centering the visual form across the research project [68, 127]. This commitment is agnostic to methods and, as such, several are employed in visual research.

A prominent line of qualitative visual research (e.g. [26, 107, 147]) applies systematic coding and qualitative analysis via "close reading" of visual content. We draw upon this practice as a foundation of our work. Other qualitative methods that have been applied to HCI studies of visual content include design-based methodologies (e.g. [41, 51, 60]), photovoice methods where participants create visuals (e.g. [27, 98, 120]), visual elicitations where researchers show participants media (e.g. [99, 167]), and various sketching activities (e.g. [115, 160]).

Quantitative methods may involve large scale analysis of imagery via computer vision, machine learning, and AI techniques to surface trends (e.g. [71, 102, 156]). Increasingly, with the advent of Visual-Language Models (VLMs), these quantitative methods may be mediated by AI systems having a role in [92] or fully analyzing visual data (e.g. [76, 181]). Some scholars argue VLMs fall short in visual analysis, both by transforming the visual to textual representations and lacking frameworks informed by visual studies, resulting in documented shortcomings in purely visual analytical tasks [154, 178]. Aligned with those arguments, our work positions computational methods in service of a human-centered analytical process and centers deep reading of visual content by researchers.

Such readings require visually literate researchers. Visual literacy, related to other relevant literacies invoked in HCI such as media [28, 70], digital [142, 176], and algorithmic [42], seeks to empower a viewer to critically engage with and understand the visual as a sociotechnically produced artifact [126, 127].

Modern visual literacy research focuses on the democratization of image production and the participatory nature of (highly online) visual culture [133]. This is particularly salient in online social data where both user participation and platform algorithms shape visual culture [156]. Researchers have also called for an increase in visual literacies in both the public and for researchers towards adopting more visual methods [68, 127, 129].

In our framework, we seek to foster visual literacy in our research teams via a participatory, pedagogical approach. We largely focus on training students in visual methods through active engagement in research, adapting a pedagogical innovation ("directed research groups") introduced by Turns and Ramey [165]. In improving a researcher's visual literacy, we borrow from Berger's notion of "active seeing" [16] and Braun and Clarke's concept of a "knowing researcher" [23]. In doing so, we hope our work helps train researchers who are aware of how their active, situated experience of viewing and analyzing content is shaping and being shaped by conducting (visual) research.

2.2 Limits of Big Data analyses on social data

Though we draw upon qualitative methodological traditions, our methods are mixed; integrating human close readings with computational analyses. Computational analyses are often needed to address the scale of big social data, but they may also miss key nuances and contexts.

We therefore respond to boyd and Crawford's "critical questions for big data", which stresses how the scale of big data decontextualizes the social nuances of the human traces it represents [22]. Scholars have contended with such limitations via a variety of mixed methods approaches to solicit insights, often combining computational zoomed out approaches with qualitative zoomed in ones [4, 94, 156]. Bolíbar explicitly contends with this in advocating for a strong integration of qualitative and quantitative methods for effective social network analysis that recontextualizes human social connections across the micro, meso, and macro levels [21].

In the methodological approach of cultural analytics, Manovich critiques what quantitative and macro-level approaches may miss [104–106]. He explicitly seeks a combination approach of close and distant readings: "How can we combine computational analysis and visualization of large cultural data with qualitative methods, including "close reading"? (In other words, how to combine analysis of larger patterns with the analysis of individual artifacts and their details?)" [105, pg 1-2].

We explicitly build upon Manovich's view of leveraging both computation and human close reading in mixed methods work.

2.3 Methodological contributions and innovations in HCI

Our primary contribution is methodological. Anchored in interpretivist, mixed methods research, our framework does not focus on deriving statistical inferences, though aspects of our methods could be adapted to hypothesis-driven work. We instead draw upon HCI literature evaluating methodological contributions with specific focus on qualitative, grounded, and mixed methods.

Wobbrock and Kientz emphasize methodological contributions in HCI should: (1) "create new knowledge", (2) inform best practices, (3) be evaluated by the "utility, reproducibility, reliability, and validity of the new method or method enhancement," and (4) be validated by repeated application [177]. Our framework enables researchers to generate insights from visual data (creating new knowledge), and provides best practices for performing human visual analysis sustainably with careful integration of computation. Towards reproducibility and repeated validation, we provide three empirical case studies of our framework in action.

Van Berkel and Hornbæk extended Wobbrock and Kientz's work, detailing implications of HCI contributions [168]. They expound that methodological implications shape how researchers study emergent phenomena and shift norms across the field. Our work aims to do so by increasing capacity and shaping norms towards visual research within HCI. Aligned with Oulasvirta and Hornbaek's call to consider HCI research as problem solving at its core, our framework — developed to contend with emergent dilemmas in our work — is meant to help navigate complexities of how to effectively and safely integrate computational and human analysis [123].

In examining HCI methodological frameworks, we find several examples across history and contexts: "cooperative inquiry" [47], "feminist HCI" [10], "trauma informed computing" [37], "Conflict Sensitive Design" [122], and several others [86, 139, 152]. These methodological frameworks provide guidance, based in empirical work and bringing in outside theories and literatures, as we seek to do with a visually-informed lens (see 2.1).

We are motivated by a paucity of visual methodological innovations. Most methodological innovations in social computing focus on textual data, offering new measures (e.g., metrics, signals) or research instruments (e.g., interactive graphs, dashboards, coding schemas) [36, 59, 101, 158, 182]. This includes framework-level contributions [180], with visual methodological contributions often focusing on individual methods or measures [71, 100].

Methodological contributions have underlying commitments to theories, epistemologies, methods, and lineages they invoke [24]. Our framework draws from interpretivist, constructivist, and visual epistemologies stemming from a commitment to a "grounded" approach to research in the style of Charmaz and colleagues [32, 33, 35]. We emphasize collecting primarily visual data and analyzing it iteratively through various tools and methods to reach saturation, building both empirical insights and theoretical contributions [33]. When quantitative methods are used, we, like Charmaz and contemporaries, insist on human interpretation and involvement to construct results [32, 33, 35].

Although we combine computational and qualitative methods in a grounded style, we do not posit our work as "computational grounded theory". That body of work often leverages computational analysis to automate and assist stages of human analysis (particularly coding) and finding emergent themes to iterate upon [117, 119]. Instead, our methodology centers careful, involved human analysis, where computation with humans in the loop facilitates insights via interpretation of computational analyses or extension of human analysis to more data, aligning with the approach of Carlsen and Ralund [31].

We seek to extend methodological innovation literature with a grounded, visually-informed framework that combines grounded theory's iterative nature with computationally assisted, human-in-the-loop analysis of visual social data. We do so by developing our framework as a design problem, solving our own emergent needs as researchers, documenting the framework with clear success criteria, and illustrating its application across three case studies.

3 Components and commitment of our framework

In this section, we detail our research framework, which supports a grounded, interpretivist approach to analyzing "big" visual social data that deeply integrates mixed (qualitative and quantitative) analyses. This approach incorporates multiple, often iterative, analyses that build upon one another (although we present them in a more linear style here). We posit three components (summarized in Figure 1): visual grammars (3.1), human analysis (3.2), and computationally supported analysis (3.3).

Across these components, we instill a commitment to practices of care. This is informed by our positionality as problematic information researchers, where data may expose researchers to harms. We care for researchers' time and wellbeing by limiting exposure to harmful content and using computation to route only the most salient material to humans and extend their analytical insights. Although our framework could be adapted for an individual investigator, we argue it benefits from a collaborative team. In our work, these teams often integrate students (both undergraduate and graduate) from different fields and lived experiences, along with senior researchers. This informs our methodology, approaches to care, and embeds pedagogy in our framework as a way to train researchers as they conduct the work.

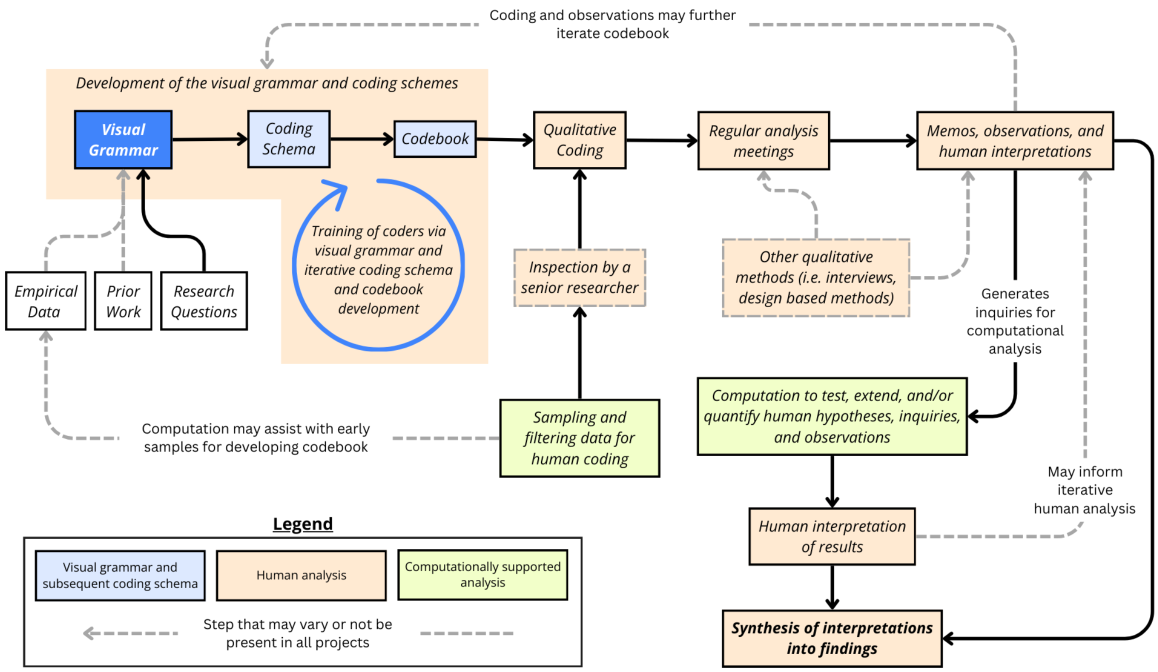

Figure 1: Summary diagram of our framework, showing our three components: visual grammars, human analysis, and computationally supported analysis, along with procedures that will vary across projects. Each project starts with a visual grammar which informs development of a coding schema and codebook applied to trace data. This dovetails with additional iterative human analysis, which may incorporate other qualitative methods. Human analysis generates inquires for computationally supported analysis to quantify and investigate, for final human interpretation of the results. Importantly, computation may also support sampling and filtering of data for initial human analysis.

3.1 Visual grammars allow researchers to systematically analyze images

Visual grammars are instruments of fields that systematically study visual media, such as formalism (study of visual form via attributes like color or layout) [173], semiotics (study of how meaning arises from imagery) [170], and meme studies (systematic studies of online memes) [153]. We draw from Kress and van Leeuwen's definition of visual grammars as systematic structures and rules that guide how humans examine and construct meaning from imagery — providing an underlying framework for how to *look* at media, like how linguistic grammars provide rules for constructing meaningful sentences [88]. Grammars may encompass observational (i.e. color) or interpretable (i.e. affect) traits via an agreed upon system [153, 169, 173]. Figure 2 displays *composition*, an example of a human-interpretable visual grammar from formalism [131].

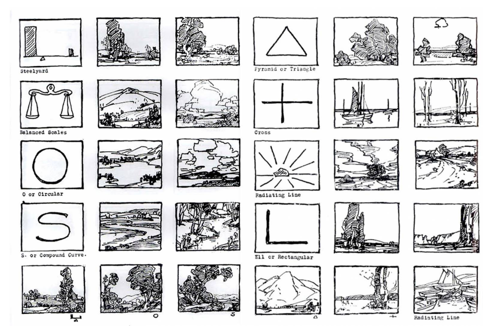

Figure 2: Composition, a visual grammar in formalism, where human interpreters match artworks to agreed upon categories of spatial arrangements of objects. From: [131].

Our framework positions the development and application of a visual grammar as a core component of visual research. This grammar should be contextual to the project, steering researchers to systematically read imagery based on their research questions. We explicitly note **systematic** instead of formal or measurable. These grammars do not have to be fully reproducible, absolute, or even quantifiable. Instead, visual grammars must be grounded in research questions and data, be able to be rigorously defined, and be reliably applicable by human analysts. This approach to visual grammars, drawn from semiotics where they originate [88, 170], further embodies principles of rigorous qualitative research, such as transferability [54, 65].

In our work, we borrow and adapt visual grammars from existing literature (e.g. formalism) or construct our own to serve as the underlying structure for developing a coding schema. This may occur deductively, such as by applying a visual grammar related to formalism and constructing a context-relevant coding schema inspired by formalistic dimensions (i.e. colors, composition). Coding schema development can also occur inductively, such as by using a visual grammar to guide how researchers examine samples of images to construct salient code dimensions. Regardless of approach, a visual grammar provides a foundation for *how* to actively look at media [16, 88], offering conceptual grounding and constraints to increase a researcher's visual literacy [127] and make visual knowledge actionable to developing a coding schema. This coding schema — separate from, but anchored upon and emerging through application of the visual grammar — is subsequently operationalized into a codebook, the research instrument applied to the data.

We propose three success criteria for visual grammars in our framework:

(1) Grounded in the data and generative towards developing a coding schema and steering inquiries to answer research questions.
(2) Transferable across other contexts, in the spirit of rigorous qualitative research.
(3) Increasing the visual literacy of the researcher via becoming a knowing and active viewer.

We describe these dimensions below.

3.1.1 Grounded and generative to developing a coding schema and answering research questions. We borrow from Charmaz's definitions [33] to "ground" visual grammars. Grammars can inductively emerge from data or be informed by external priors, as long as those priors are iteratively in dialogue with the data so they, as Charmaz notes, "earn their way" into an analysis [34, pg 64]. For example, a known grammar as a starting point would be iterated upon by applying and adjusting it to small data samples to finalize a grammar that meets specific project needs. In other instances, inductively developing the grammar may be more appropriate, with researchers iteratively testing new ways of looking at data until finding one salient for their needs and generative for a coding schema. In grounded qualitative research, such decision points vary between projects, but may be centered in concepts like theoretical saturation when iterations are no longer producing new insights [33]. To this end, the iterative development of a visual grammar is generative, providing a basis to fully develop a coding schema. Through this development process and subsequent coding, new hypotheses for computational and human analysis are generated and build to answer the study's original research questions.

3.1.2 Transferable across other contexts. In the spirit of qualitative rigor [65, 159], a visual grammar may not be fully reproducible between studies but transferable. In past work, other visual grammars, like formalism between different artistic mediums [169, 173] and Shifman's grammar across different meme genres and contexts [6, 58, 153] have demonstrated transferability. We stress that, like all qualitative work, transferability will not be seamless and grammars will necessitate adjustment between projects and contexts, but should provide a systematic foundation for other research.

3.1.3 Increasing the visual literacy of the researcher. Visual grammars provide sufficient structure to help the researcher become an active and knowing viewer (a la Berger's notions of active seeing [16]), fostering their visual literacy [127]. Effective grammar design centers on structuring image analysis so researchers can focus their efforts on the most nuanced, active viewings — making the best use of their labor. Grammars should avoid unnecessary complexity, leveraging thoughtful groupings or layered approaches to reduce friction and cognitive load. This may involve testing different versions of their translations and looking to existing grammars in literature. In Sec 4, we describe how our coders (primarily students) experienced this increase of visual literacy and becoming more active viewers. This development was not only mediated by the grammar, but also from our participatory learning approach (in style of [165]) and process of human analysis.

3.2 Human analysis adds depth and expertise

The application of some visual grammars could be automated, such as the case of "computational formalism" [173]. But our framework centers human analysis instead, drawing from interpretivist [146] and grounded [33] research approaches. Via thorough qualitative analysis, human researchers develop "thick" [62] and contextualized insights from visual social data. In doing so, salient and grounded theories are developed to be tested with other analyses — including computational.

In our work, this occurs by taking qualitative coding and individual researcher memos and surfacing themes via recurring thematic analysis meetings, ending with collaboratively written memos. These group and individual memos are analyzed across coding sets and analysis meetings to elicit the most salient findings (in the style of Charmaz [33]). Some projects may use other analytical approaches (i.e. discourse analysis over thematic). Regardless, such analyses shape and are shaped by the human researchers, a central tenet of qualitative work [23, 32, 38].

Human impacts on analysis start at the onset of a project with what (and how) questions are being posed, along with the preconceptions and goals researchers bring to their work. In the spirit of qualitative research, we don't seek to remove these human biases, nor do we believe such biases can be fully removed [46]. Rather, we believe the "thick" analyses we seek are made possible by how researchers' lived experience and expertise as participants and consumers of visual culture color qualitative analysis [38, 62, 127].

To navigate human influence on research, we leverage reflexive practices that ask researchers to reflect upon and acknowledge how their experiences, beliefs, and biases shape the work [23, 46]. We have incorporated ongoing memoing [111, 140] with reflexive prompts [5] and collaborative analyses in diverse teams [11, 33, 174] as our key reflexive practices. Such practices highlight that knowledge produced is socially constructed from the researcher's standpoint [10, 66], enabling them to surface assumptions, check for blind spots, and ultimately strengthen their analyses.

Along with reflexivity, evaluative measures such as inter-rater reliability (IRR) can help analyze code transferability and expose areas of human disagreement for further investigation, even when IRR is not meant to be reported or the goal of coding [110]. This internal checking and engagement with coder disagreements can also improve transparency, promoting helpful dialogue and trust in the research process, and increase depth of findings [124].

Other best practices may include training in coding and memoing in parallel and recognizing humans shape other research methods, from interviews to computational analyses. Although computational analysis may be seen as objective, we embrace and deeply account for the reality that human researchers shape the questions asked and tools used in such analyses — and that these tools are products of human designers [30, 81].

We find reflexivity has unique ramifications in visual projects. For example, it is a viscerally different experience to see someone who looks like you depicted in a certain way (particularly negatively) compared to reading about it [72]. Simultaneously, this experience may provide additional analytical depth — which can be scaled up and further iterated upon through computationally supported analysis.

3.3 Computationally supported analysis extends and aids human analysis

Computational analysis will vary by project, but we position it in service of — and informed by — human analysis. We lay out three success criteria for computational analysis in our framework:

(1) Aiding the sustainability of human analysis
(2) Extending human analysis to test grounded hypotheses
(3) Maintaining visual integrity

3.3.1 Aiding the sustainability of human analysis. Computation can save qualitative coder time and support their wellbeing, making human analysis more sustainable overall.

In the processing and acquisition of visual social data, computation can protect researchers by filtering out violent or sexual imagery. It may also help filter emergent categories and features of content that researchers identify as harmful or not worth analysis. This filtering helps surface the most contextually relevant samples, ensuring human analytical labor is reserved for only the most salient data, particularly important in projects with potenitally distressing content.

Questions of sampling and cleaning visual data differ significantly from text. Concepts of redundancy in text, like character differences (i.e. “dogs" vs. “d0gs"), are quantifiable and aligned in human and machine perceptions. However, analogous quantitative differences between image pixels may be misaligned — for example, a cropping of 5 pixels may be less observable to a human than a machine. Such nuances complicate what it means to pull a “diverse” sample of visual media compared to text, and thus need human steering. When these samples are analyzed by humans, computation can also make for an analysis environment that prioritizes safety and reduces cognitive load to avoid fatigue.

3.3.2 Extending human analysis to test grounded hypotheses. Computation can help extend human analysis by “zooming out” to apply close human readings to a larger dataset and surface macro level trends. In doing so, computation can enable testing human-led hypotheses across a larger dataset, embodying a grounded theory approach [33].

Applying human readings to a larger dataset may encompass providing labels of subjects from codes for finetuning VLMs or classifiers. Human analysis can also provide meaningful and grounded associations of these labels, like the significance of subjects occurring together. Human insights may also surface trends like provenance (i.e. seeing screenshots tend to be from news websites) which can be extrapolated to a computational analysis. Meanwhile, surfacing macro level trends, like clustering images or quantifying the amount of image text or color palettes, can grant researchers an understanding of their larger dataset to help interpret and contextualize their close readings.

3.3.3 Maintaining visual integrity. Our last success criterion is **visual integrity**, which we define as centering inquiry around the visual, preserving both the visual form and human context of data across a research project.

When designing computational analysis, our framework uplifts visual integrity as a central design value and criteria. In practice, this commitment to visual integrity means that we:

(1) preserve the visual modality as much as possible,
(2) do not remove visual data from its social context,
(3) prioritize human interpretation and steering of analysis

First, preserving visual modality may come from the use of well-designed, modular scripts for specific visual tasks (i.e. logo extraction, color quantization), an argument made by Lutz et al. [100]. Methods that transform images into textual embeddings like VLMs are employed to answer precise visual and human inquiries over open-ended labeling (i.e. “tell me about this picture”), particularly given recent work on VLMs’ visual task limitations [154].

Second, to avoid removing social context, non-visual information about media, such as engagement or metadata, are collected to contextualize the social and visual production of that data. This could mean focusing on what metadata tells us about participatory actions behind related images, enabling a human-centric approach to provenance. Or using project-relevant symbols for machine identification that are gathered and informed by human insights via qualitative analysis, literature, or interviews. Such measures support computational analysis and insights that remain situated in the social, cultural, and production context behind the visual, aligned with the notion that visuals are active, situated objects [16, 127, 128, 145].

Third, our framework encourages researchers to center human interpretation and steering of computational analysis, rather than allowing computational results to displace or supercede human knowledge and experience. This means leveraging computational methods to answer human-driven, visual-centric inquiries, rather than retro-fitting inquiries to a particular computational method before engaging with data. Our "grounded" (as per Charmaz [33]) approach to research, where human interpretation of ongoing, iterative results and analyses (including computational ones) spur new hypotheses for investigation drives this prioritization of human centering. In our case, we apply this to visual contexts and explicitly extend and aid human steering via computation.

We illustrate how we balance visual integrity in computational analysis across our case studies (Sec 4), and discuss its broader implications as a sensitizing concept for visual research in Sec 5.1.

3.4 Commitment to practices to care

Our framework and its components were designed with a commitment to care for researchers (in our case, often more junior ones), drawing from a growing body of literature [3, 13, 55, 141, 163]. We explicitly draw from a "safety-as-enablement" approach [163] that leverages reduction of exposure while maintaining the contextual need of researchers to analyze potentially distressing content as part of their research. This commitment centers on systematically deciding what content truly warrants close human reading, and the practices used to read it. Though this does not remove all potential harm, it seeks to minimize harm by ensuring any content the researchers witness — whether it's distressing or not — truly benefits from their labor. To support this, we incorporate human and cyber infrastructures (as per Lee et al's definitions of the computational and human infrastructure that empowers research [95]) to support care.

Our framework's human infrastructures are largely informed by our positionality as mentors and instructors. Although university institutional review boards (IRBs) review our studies for procedural, legal, and ethical compliance, this is only a first — and often insufficient [44, 53, 75, 171] — step to caring for researchers. Therefore, we implement several additional measures of care, beginning with training — especially important as much of our work uses an active-learning model with junior scholars [165]. We train researchers not just in visual literacy [16, 127], but also in best practices for reviewing traumatic media from psychology [72, 175] and journalism [77, 143]. We also design projects with extra capacity on the research team, such that coders can take a week off and perform other tasks like literature review. Our work uses a *harm reduction* approach — minimizing exposure of content to more junior researchers by having a senior or lead researcher as a stop gap. Other projects may benefit from harm distribution, with several researchers reviewing small samples [163].

These human infrastructures influence how care intersects with our framework components. The grammar, outlining how to look at an image, should help reduce cognitive load and serve as a structure to help researchers systematically look at imagery that could otherwise be overwhelming. The codebook offers another place to embed care and safety-as-enablement [163], such as grouping related codes to reduce friction and providing small image previews before full-size viewing in coding environments. Reflexive memos may help researchers stay in tune with their wellbeing and emotions during analysis [83]. Additionally, analysis meetings can offer an opportunity to check in with coders alongside the usual focuses on arbitration and collaborative thematic analysis. We have found other practices, like coding while being co-present and during the day instead of evening, to also be helpful for teams when analyzing difficult content.

We also leverage computational infrastructure for care via our third framework component, computationally-assisted analysis. In particular, we focus on caring for researchers with computation by filtering out the "worst of the worst" content (i.e. CSAM[1], extreme gore, or slurs) before humans see data. We also care for researchers by ensuring that any analysis of troubling content focuses only on the most salient and contextually meaningful samples, filtering out redundant images and those lacking substantive details (i.e. out-of-context hate symbols).

4 Case studies

Next, we walk through three case studies of problematic information where we have applied this framework:

(1) **Anti-immigrant visual propaganda (4.1):** Contemporaneous, weekly analysis of visual anti-immigrant propaganda on X and TikTok in the US across 2024, a Presidential Election year.
(2) **Jesus AI Slop Imagery on Facebook (4.2):** Studying a phenomenon of "AI slop" (low quality, mass-produced AI imagery) featuring Jesus, which led to user frustration and concerns around scams and misinformation.
(3) **Telegram Visual Identity and Disinformation of Russian-Ukraine War (4.3):** Examining how different visual tactics manifest national identity and propaganda during the occupation of Ukraine by Russian forces.

The first two case studies were led by this paper's first author, with Case Study 1 serving as the initial development site for the framework. Case Study 1 involved the second, third, and senior authors, along with additional collaborators. Case Study 2 involved the second and senior authors and additional colleagues. Case Study 3 was led by the fourth author and their collaborators, with the first author contributing this framework. All case studies received IRB approval at their respective institutions (Case Studies 1 and 2 at the first author's institution; Case Study 3 via the fourth author's). Each case study has resulted in published, forthcoming, in-review, or in-preparation work. Table 1 summarizes how the framework's components — visual grammars, human analysis, computationally supported analysis, and commitments to care — were applied and varied across the case studies.

[1]CSAM stands for child sexual abuse materials

Framework Component	Case Study 1	Case Study 2	Case Study 3
Visual grammar in all case studies	Developed and applied visual grammar as a formative structure for examining media and steering subsequent coding schema development.		
Visual grammar development per case study	Started with prior literature and iterated with inductive open coding to finalize layered grammar.	Interviewed religious experts and users to develop a grammar to examine subject arrangements.	Inductive open coding followed by prior literature and Case Study 1 grammars to finalize a layered grammar featuring stance.
Human analysis in all case studies	Human analysis in collaborative teams, involving qualitative coding and iterative thematic analysis across several meetings. Analysis surfaces hypotheses for computational testing.		
Performing human analysis per case study	50 weeks of weekly content analysis meetings.	6 weeks of interviews and analysis meetings. Followed by 4 weeks of content analysis meetings.	7 weeks of weekly content analysis meetings.
Computationally supported analysis in all case studies	Designed to answer exact human inquiries, extending human analysis via quantification of identified visual traits. Results interpreted by researchers. Supports human analysis via filtering and sampling.		
Conducting computational analysis per case study	Filtering for saliency, redundancy, and safety of coding samples. Modular scripts to test emergent hypotheses, such as provenance and image editing. Constructing human-in-loop analyses of related images.	Constructing a diverse image sample. Quantifying participant-identified subjects of interest and discrepancies across skin tones.	Quantifying human-identified visual-traits (i.e. logos) and tendencies in stance across a larger dataset and reusable pipeline.
Commitment to care in all case studies	Harm reduction and safety-as-enablement approach to limit researcher exposure to the most harmful content. Screening of samples by senior or lead researcher(s). Coders received traumatic image training and performed reflexive memoing. Coders could take breaks and switch to different tasks. Leveraging image previews for image-by-image opt-out.		

Table 1: Summary of core framework components and their implementations in each case study.

4.1 Case Study 1: Anti-immigrant visual propaganda

This project followed visual US-based anti-immigrant propaganda on TikTok and X from January to December 2024, a year when immigration was central to US presidential election discourse. We sought to identify the tactics and trends of this propaganda across time and formats, studying the phenomenon in real time. This included qualitative analysis of ∼1,500 media elements and computational analysis of ∼50,000 elements conducted contemporaneously across 50 study weeks (see Figure 3 for an overview). This project also led to a sub-case study on how data visualizations in this rhetoric, co-led by the first and third author [43].

Given our study's often troubling content, we took care to protect our research team, which consisted of eight coders at a time, with at least six coding at any given time. Since this project was the catalyst for the development of our initial framework, it informed many of its core components and principles.

4.1.1 Visual grammar. Coming into this work, we knew we needed to develop a systematic but flexible way to qualitatively analyze large amounts of changing, and often upsetting, imagery each week. We also knew that our qualitative approach would need to embrace the active nature of this real-time, changing rhetoric and help answer our research questions focused on what tactics and depictions were building anti-immigrant propaganda.

We spent nine weeks inductively developing a visual grammar to help orient and structure our qualitative analysis. This involved open coding and thematically analyzing 200 media elements. We began by iterating on a layered schema that we titled "Subjects, Actions, Background", drawing from alt text captioning's "object, action, context" frameworks [48, 90, 179]; systematic art analyses like Sutil's approach to analyzing motion and formalism [161, 173]; and Kress and Van Leeuwen's notions of participants, processes, and circumstances [87]. This **layered approach** allowed and trained us to systematically review content even as it (and its emotional valences) fluctuated over the course of our study. We combined this formalistic layering with analyzing **visual traits** and **depicted claims** of an image. This satisfied our goal of surfacing both high level traits and tactics that appeared across time (i.e. formats of media) and prominent subjects and valences within this propaganda (i.e. aligning subjects with actions taken to or around them). During this iterative process, we also developed a coding schema, steered

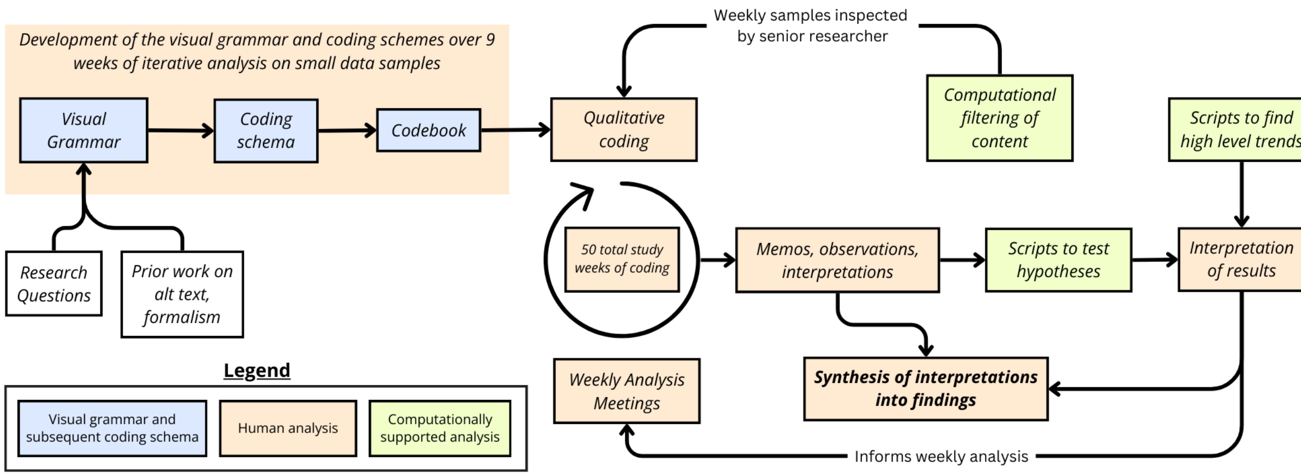

Figure 3: Summary of methods used in Case Study 1, starting with the development of the visual grammar and coding schema, which was applied to and iterated upon by human analysis in weekly qualitative coding and meetings across 50 study weeks. This analysis provided key observations for testing and quantification in computationally supported analysis, which also assisted with overall data filtering.

by the visual grammar, with more specific emergent codes such as post-production edits, provenance of images, and how an image depicted a rhetorical frame from our thematic analysis. Figure 4 shows how we developed and refined the visual grammar and the qualitative coding schema through this iterative and, at times, entangled process.

Coders reflected on how the grammar's structure taught them about visual analysis and helped them "slow down" and become more active viewers who could step outside of how they felt about this content (often, deeply negatively) to apply a coding schema and understand participatory dynamics and tactics shaping this sociopolitical issue. Building out the grammar and coding schema also generated new insights to explore and test via human and computational analysis, such as the role of data visualizations in anti-immigrant propaganda [43]. We initially developed this grammar to meet the specific needs of this project, but later found it to be generative of new questions and transferable to other contexts, inspiring reflection of the broader framework presented here.

4.1.2 Human analysis. In weekly meetings, the team constructed a thematic analysis of that week's media set and discussed codes, collaboratively writing a weekly group memo. These memos were analyzed across one another to understand the evolving tactics and trends of the propaganda.

Coders also wrote individual weekly, reflexive memos about their insights and how the imagery impacted them. We provide examples of these memos in Appendix A. Overall, this project was emotionally difficult, as many of our team members were undergraduate students from immigrant backgrounds, making our care practices (detailed in 4.1.4) critical. In one written reflection, a coder reflected on his experience coding several images of men who looked like him depicted as criminals:

*"People have [always] called Mexican men gang bangers and rapists and b**n*rs and w*tb*cks and whatever. I got called a sp*c and b**n*r as recently as last week. But it was still harder than I thought it would be to see photos of people that looked like me and my little brother and dad portrayed as rapists and gangsters."* - *student memo*

However, researchers felt drawn to this work to understand how their communities were being framed in this propaganda. In addition to learning propaganda tactics, researchers reflected on how this work made them more critical and engaged with political content they saw online both in and outside of immigration contexts. Reflexive memoing and systematic thematic analysis helped us to maintain analytical rigor informed by lived experience and knowledge of the subject matter, while also accounting for our (often deeply personal) disagreements with (and in many cases outright disgust for) anti-immigrant content. This human analysis added "thick" descriptions [62] to our empirical findings and facilitated organic discovery of key findings, such as the unique role of data visualizations in this rhetoric or how the certain image aesthetics showed provenance (such as a presence of memes made on Facebook). These insights were later extended by computation but may have been missed without initial human surfacing.

4.1.3 Computationally supported analysis. In this work, computation aided researchers in identifying and removing both the "worst of the worst" content (gore and sexually-explicit imagery) and also contextually harmful content we deem **hatespam**, which was repetitive slurs or hate symbols lacking additional meaningful content to analyze. We used computation, via identifying keywords and colors (see Figure 7) to filter out this content. In doing so, we preserved the visual modality of our data and filtered out content deemed not worth close human reading, aligning with our commitment to visual integrity as described in Sec 3.3.3. We also used computational measures of similarity [7] to quantify "image families" (see Figure 5) in our dataset to avoid the analysis of redundant images. This allowed us to keep track of participatory dynamics across remixed images without having to apply valuable coder bandwidth and risk

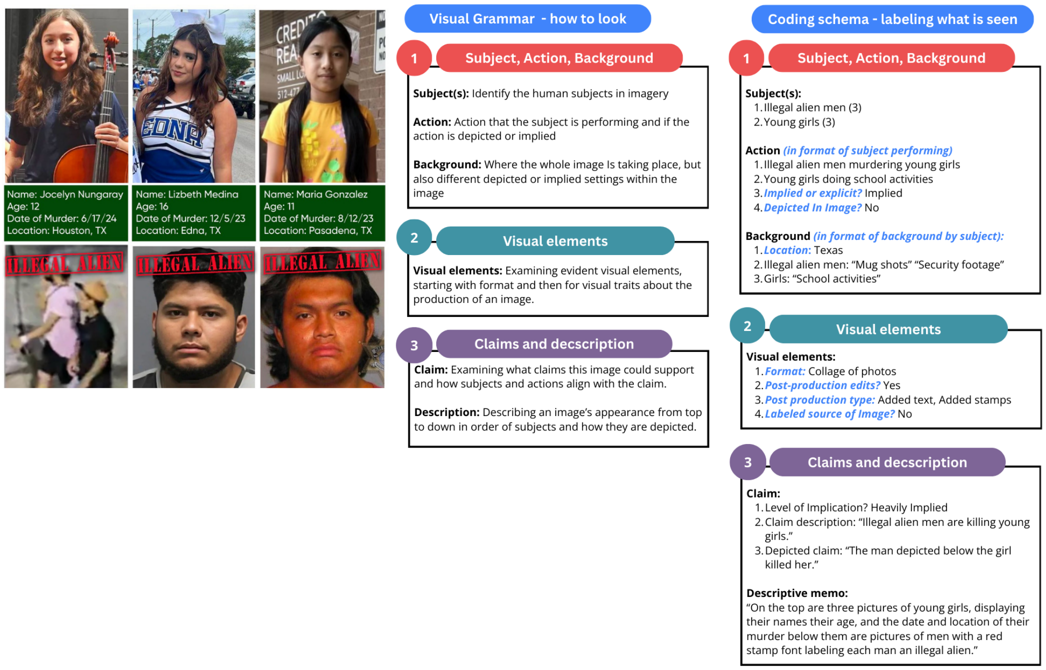

Figure 4: Sample image, annotated for publication, based on our visual grammar and the subsequent coding schema applied. Coders first go through (1) "Subject, Action, Background", in codes and annotations. Then, they focus on (2) "Visual elements" such as formats, post-production, and provenance. Finally, they focused on (3) "Claims and description" memoing about an image's claims and holistic description. Blue text based on closed code prompting instructions from the coding environment and quoted text is copied student written descriptions. Figure is summary, not entire reproduction.

repeated exposure to redundant data, further preserving the social and production context of our data in line with visual integrity.

Figure 5: An example of a family of visually similar images identified via our computational analysis.

In the case of our data visualization sub-case we extended these image families into "data visualization lineages" (DVLs) [43] by providing a chronological account for related media (a lineage). We used computational similarity to locate similar images in our dataset and OSINT techniques to find related content online and interpret metadata. These lineages were manually analyzed for details of the participatory dynamics that formed it. This illustrates how definitions of visual redundancy and what warrants human analysis can vary across projects and how similar computational methods can support different goals.

Throughout the study, we used computation to extend human observations (see Figure 7). We leveraged several computer vision techniques, particularly optical character recognition (OCR) and

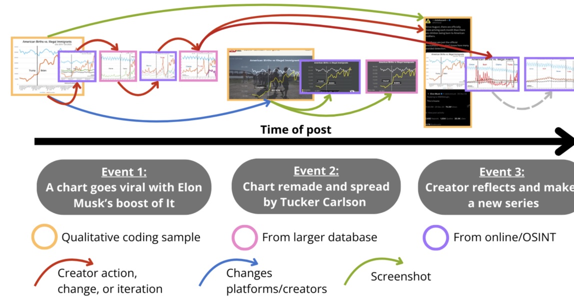

Figure 6: An example DVL, where similar images became the center of detailed investigation. Source: [43].

color quantization [100]) to quantify keywords (see A in Figure 7), logos (see B in Figure 7), and match sources of screenshots that human coders created test images for (see D in Figure 7). We also ran error level analysis to detect image editing (see C in Figure 7).

Due to the real time nature of our analysis (tracking new anti-immigrant content weekly) and constraints on time and compute, we relied on small, modular scripts using a variety of computer vision techniques to test human hypotheses across our dataset. This also gave us more control over steering our computation, and allowed us to develop the notion and priorities of visual integrity (as this was the case study from which our framework first emerged).

Figure 7: Four key tactics used in this project. (A) OCR extracting text, to isolate images with particular key words and quantify popular words in a dataset. (B) logo identification via a database of emergent logos. (C) Error level analysis showing post production image manipulation. (D) Showing font and emoji matching to detect a post is from Facebook, based on a test image.

4.1.4 Commitment to care. Given the distressing content and our coders' intimacy with it, we implemented several practices of care.

We ran this project as a series of for-credit research groups (as described by [165]) from January to December 2024. Each week, analysis meetings served as a time to discuss findings and how the research impacted us. During these meetings, two PhD students (first and second authors on this paper) supervised and facilitated these conversations, modeling the behavior of sharing how this imagery and other content they studied had impacted them. Additionally, at the start of the academic term, a university mental health professional spoke to the team and covered well-being resources specific to this work and each coder reviewed resources about traumatic imagery [77, 143].

These measures served as human infrastructures of care and that helped us iterate on computational ones. Computation helped filter out the "worst of the worst" content and hatespam, and also provided the most salient samples for human analysis by filtering out redundant imagery. In our coding environments, we enabled small previews of images instead of defaulting to large screen views to allow coders to opt-out of imagery as needed. Logistically, we instilled a policy to allow coders to switch off for a week or two if they needed a break and do other research tasks (i.e. literature review) that would still fulfill course credit. This switching was made possible by having at least two extra research assistants on the team who could rotate into the coding.

These practices worked together to support our coders and inform care practices in future case studies.

4.2 Case Study 2: Jesus AI Slop on Facebook

AI Slop is mass produced AI-generated content (in this case, often surreal or even disturbing images) that has been tied to engagement monetization efforts of spammers and scammers [45, 85]. In 2024, large quantities of AI Slop imagery appeared on Facebook, and we investigated a popular subgenre: Jesus imagery.

This study explored user experiences and folk theories about pages behind AI Jesus Slop, and examines visual trends within the imagery. We integrated 20 interviews with users and religious experts into our visual research framework alongside qualitative and computational analysis of AI Jesus Slop (see Figure 9).

4.2.1 Visual grammar. We knew our visual grammar needed to identify visual trends most culturally and contextually significant to users. We inductively developed our grammar, interviewing religious experts and users using visual elicitation (as per [99, 167]) with a visually diverse sample of images. This sample was constructed by our open coding of images during data gathering through deep immersion [94] (see Figure 9) and interpreting results of computational clustering of 6,000 images (see 4.2.3).

These interviews helped us identify theologically (from religious experts) and culturally (from users) significant symbols and depictions of Jesus (see Figure 8 for examples), and showed us how participants read and looked at these images. This allowed us to ground our grammar around their viewing — focusing on identifying subjects, actions, and the arrangement of subjects with respect to one another. We structured and populated a coding schema around visual **subjects** (i.e. women, children), **depicted actions** (i.e. Jesus praying, being saved), and **stylistic choices** (i.e. Jesus as a cartoon, Jesus as a fruit). We operationalized this into a codebook and applied it across a visually diverse set of 600 images. Through coding, we quantified visual trends across images and generated a final list of subjects for computational analysis.

Figure 8: Examples of key subjects occurring with Jesus that stood out to participants such as appearing with pregnant women, the military, and injured children.

4.2.2 Human analysis. This research was conducted by a team of four PhD students, all with prior research experience with problematic information. Two had religious backgrounds, which provided contextual knowledge for the interviews and analysis. Notably, human analysis of imagery included perspectives from interview participants as well as researchers through active co-viewing of AI Jesus Slop, shaping the analysis with observations and theories about the images' socio-religious context and production. These collaborative interpretations iteratively informed the grammar of how to look and the subsequent qualitative coding schema. But they also colored overall memos, interview analysis, and the design of computational analysis — keeping our analyses grounded the human experience and meaning making from the visual.

4.2.3 Computationally supported analysis. We used CLIP (a zero-shot classifier) with human verification [156] to detect subjects of interest (i.e., women and animals occurring with Jesus) from our

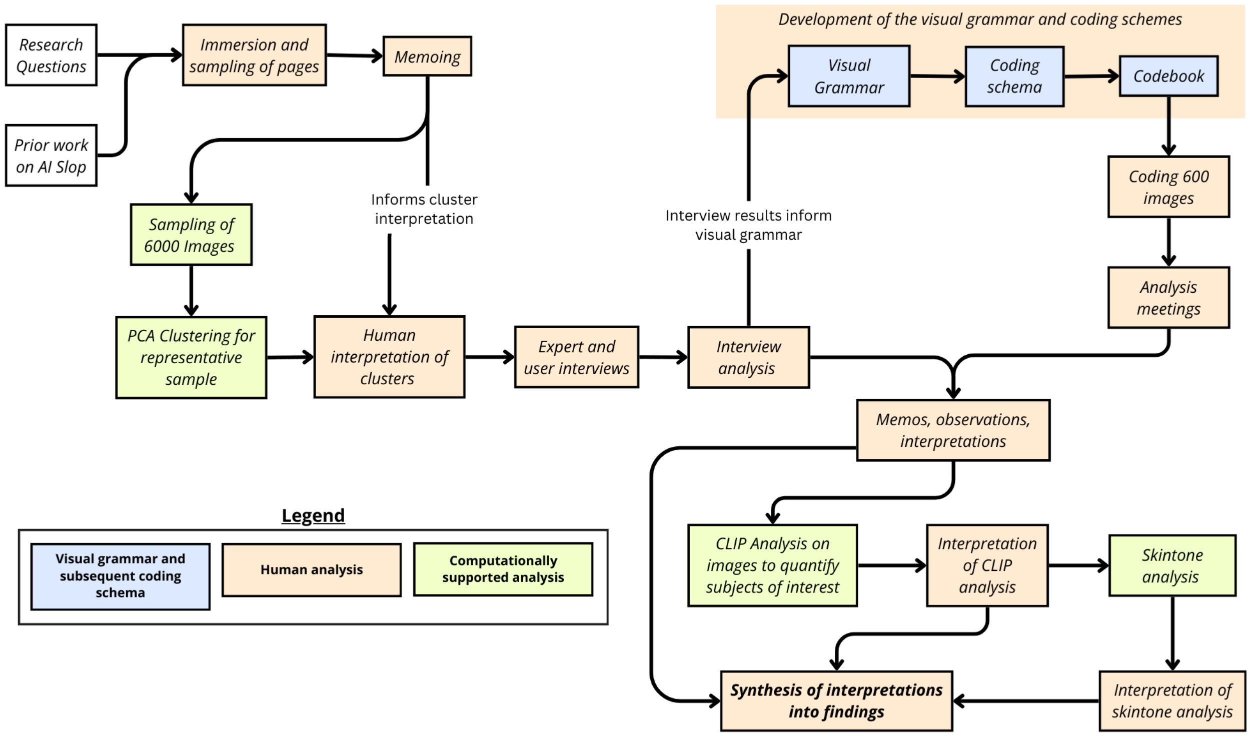

Figure 9: Summary of methods in Case Study 2, starting with detailed human analysis and exploration of our phenomenon, which computation helped facilitate sampling of for interviews that informed the development of visual grammar. Grammar was then applied to a sample via qualitative coding. Human analysis and interpretation of coding and interviews were then explored and quantified with computationally supported analysis regarding subjects of interest and skintone.

coding and interviews — extending human observations across the 6,000-image dataset. We chose CLIP to leverage CLIP embeddings present in many text-to-image generator outputs [63]. This resulted in subsets of data with subjects co-occurring with Jesus for human interpretation, driven by human collected inputs and a carefully chosen visual method, in sticking with visual integrity.

Human inspection of CLIP-derived subsets — especially women, children, and participants' noted disparities in depictions of people of color — spurred further analysis. To attain subsets of depicted children (a particularly interesting subject) separated by skin tone, we used color quantization and the Monk Skintone scale [100, 156] to quantify distributions of skin tones of children co-occurring with Jesus and, via human inspection, derive qualitative details of how children with different skin tones were depicted.

This case study's computational analysis focused on testing researcher and participant hypotheses on 6,000 images, revealing how these observations did or did not scale. For instance, several subjects frequently mentioned by interview participants did not map to the most frequent subjects in the quantitative analysis. In other words, the experiences of people who engaged in this visual discourse and the things they found most salient were not perfectly aligned with the most prevalent content in the data. Without the human visual experience (via our coding and interviews), we may have missed these subjects which resulted in rich insights about the social production and user experiences of this imagery.

4.2.4 Commitment to care. Our research team of PhD students all had experience with problematic information and received relevant training. We kept running memos regarding the imagery, reflecting asynchronously and synchronously in analysis meetings. Additionally, in line with our harm reduction approach, the lead researcher, being the most experienced with visual analysis, still served in a review role of the samples. We also made an active effort to point out what could be humorous images, trying to not overly focus our experience on disorienting or disturbing (i.e. sexual, gory) images. But still, looking at hundreds of these images was disorienting and we used computation, where we could, to reduce human effort and limit unnecessary exposure.

We also took measures to care for interview participants seeing this imagery. We practiced continuous, informed consent — repeatedly reminding participants they could skip any image(s). We also gave participants zoomed out previews of images first, and only later zoomed in after consent and for more detailed conversations.

4.3 Case Study 3: Russian-Ukrainian war imagery on Telegram

Telegram is a key platform for news and discourse across de-occupied, occupied, and frontline regions of Ukraine in the war with Russia, with regional channels differing in content and vulnerability to information operations [118, 121]. This case study focused on prevalence and tactics of national and social identity in pro-Russian and pro-Ukrainian visual content. In particular, we sought to (1) understand how known frameworks of hostility/solidarity (previously applied to Ukrainian textual data) [91] manifest in visual media and (2) inform the design of a reusable, multimodal VLM data pipeline for analysis of large corpuses of conflict zone data.

Over twelve weeks (summarized in Figure 10), undergraduate researchers mentored by more senior researchers (including the first and fourth authors) qualitatively analyzed Telegram 200 image-post pairs and prototyped components of the data pipeline.

4.3.1 Visual grammar. Via visual analysis, we sought to surface trends and tactics of how solidarity/hostility was depicted in pro-Russian and pro-Ukrainian content and to inform the design of a multi-modal data pipeline. The aim for our grammar thus was reading images for 1) depicted stance and how this stance was supported by visual traits and tactics and 2) how to parse this into machine-digestible features.

We inductively developed our grammar over four weeks of fine-tuning our initial understanding (pro-Russian vs pro-Ukrainian vs neutral vs none) of depicted stances via open coding small data samples, and eventually deductively incorporating definitions of solidarity and hostility from previous work [91]. Our grammar first had coders identify observable features for machine translation and then more interpretative features for human analysis — first stance, then solidarity and hostility, and finally how different visual elements supported these traits. This layered approach allowed coders, as in Case Study 1, to slow down and systematically analyze often difficult images of active conflict that impacted their communities (as multiple were Ukrainian).

We subsequently built a coding schema using this underlying grammar, focusing on **depicted stance and hostility/solidarity** (i.e. pro-Ukrainian, meant to build in-group solidarity), **visual traits** (i.e. format of image), and **visual tactics** (i.e. infographics normalizing Russian infrastructure in occupied Ukraine). We populated visual traits as subjects of interest (i.e. tanks), types of imagery (i.e. screenshots), image provenance (i.e. different government websites) and modification actions (i.e. annotations). These **visual trait** codes were inspired by Case Study 1 and expanded inductively to fit our data. Meanwhile, **visual tactics** were fully inductive to our study — starting with open codes and ending in closed codes of common tactics like depicted non-war cultural symbols (i.e. sports, beauty pageants) or occupational normalization via infrastructure (i.e. calls for Ukrainians to add their homes to the Russian registry).

Because we wanted to transfer solidarity/hostility metrics of prior work on text to visuals in a VLM, interrater reliability (IRR) was critical for these codes. We found Krippendorf's Alphas [89] of $\alpha = 0.7041$ for solidarity and $\alpha = 0.8109$ for hostility, which was consistent with agreement in previous work [91]. Other codes contributed to a broader thematic and human analysis.

4.3.2 Human analysis. Across seven weeks, students became increasingly "knowing viewers" in analyzing 200 total images.

Students, all first-time qualitative coders from quantitative backgrounds, were trained by the first author through an overview of qualitative coding and practice on sample images, preparing them to code asynchronously and write weekly reflexive memos. Weekly analysis meetings, facilitated by three senior researchers[2], synthesized codes and student memos into thematic analyses and weekly memos of pro-Russian and pro-Ukrainian narratives and tactics. This built to a cumulative qualitative analysis (see Figure 11).

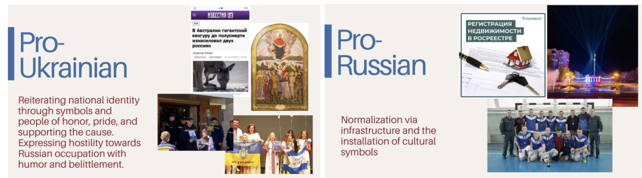

Figure 11: Slides from students' presentation at end of their summer research experience, highlighting key findings from the visual qualitative analysis.

Coding also provided human-labeled data to fine-tune the pipeline, such as logos and symbols, which students could systematically collect while coding. Qualitative analysis provided contextual significance of visual traits like logo locations (i.e. on weapons vs fundraising ads). These insights and themes, surfaced in analysis meetings and memos, facilitated the design of human-informed tasks for the pipeline. For example, instead of the pipeline just finding patriotic colors, our analysis also scoped the success of this task to include if patriotic colors were showing up in public infrastructure (i.e. illuminating buildings) — an important visual tactic that emerged during our analysis and may not have been captured by VLM-centric approaches. Team decisions in these meetings helped facilitate this translation and iterative pipeline design to make the best use of both compute and developer resources.

At times, open source intelligence (OSINT) investigations were needed to verify and contextualize findings. This was done with the research supervisor, who has professional experience with security, defense, and information operations. Such investigations, combined with human analysis, helped guide pipeline design and contextual data sources for the VLM.

Half of the research team were Ukrainians based outside of Ukraine, granting them key linguistic, cultural, and personal knowledge to inform and shape analysis. Member checking of translations and cultural cues was performed to assist non-Ukrainian researchers. This added cultural and lived expertise to our analysis, but also necessitated a commitment to care (see 4.3.4).

4.3.3 Computationally supported analysis. Human insights steered the design of the multimodal data pipeline which sought to automate and scale analysis to a larger corpus. In particular, human analysis translated to cascading tasks for the data pipeline (see Figure 13). These tasks occurred across a spectrum of high level semantic tasks (leveraging online context/meaning) to more detailed syntactic tasks (focusing on structure and form) (see Figure 12) [49, 130]. High level tasks included assessing stance and context of the image, structured and informed by our visual grammar — guiding both humans and machines in the systematic examination of the imagery. Meanwhile, more detailed tasks focused on observable, finetuned (via human data) identifications of objects like flags,

[2]The first author of this paper, one co-author who led the project and summer learning program, and a Ukrainian PhD student

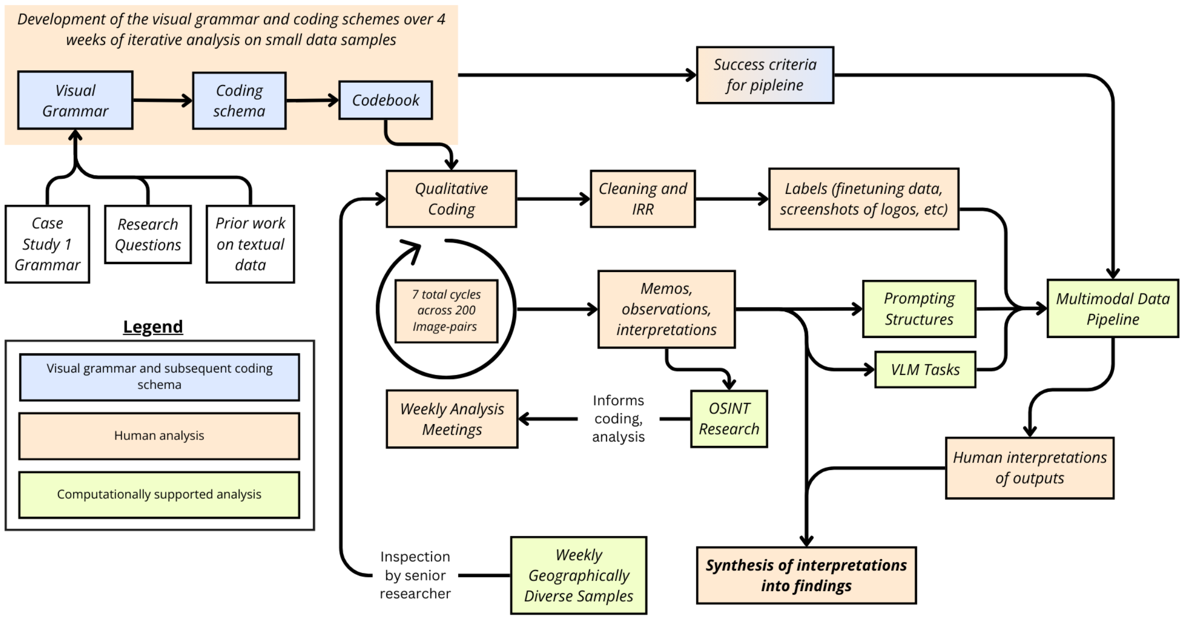

Figure 10: Summary of methods in Case Study 3, showing how iterative development of a visual grammar from prior literature and Case Study 1 resulted in a coding schema and codebook that we iteratively applied to not only surface insights for computational investigation but also to generate human labeled data for finetuning a VLM multimodal data pipeline, which provided further computationally supported analysis. We provide more details about the data pipeline in 4.3.4

logos, and patriotic colors. These tasks combined to automate and scale our detailed human insights into a reusable data pipeline design, while also focusing on preserving the visual nature of our data and on modular, visual centric tasks.

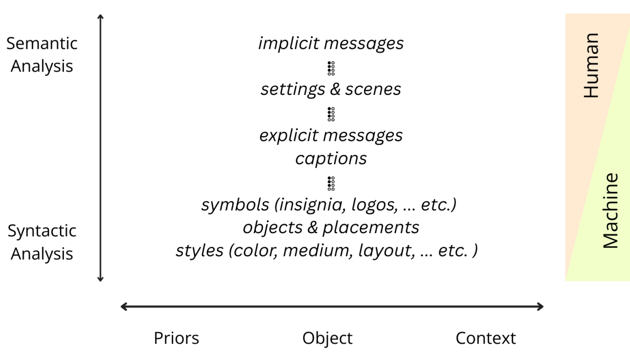

Figure 12: Identified tasks are broadly placed between syntactic (observational, such as visual traits coded by humans) and semantic (interpretable, such as visual tactics) on vertical axis, corresponding to prior literature of VLM tasks and our pipeline observations. Horizontal axis highlights the role of information such as priors for provenance analysis and auxiliary information for contextual insights.

To measure success and accuracy, VLM outputs are compared against human insights on stance, visual traits, and visual tactics to iteratively inform and train the pipeline. This fine-tuning relies on extensive human involvement to preserve analytic detail and cultural nuance via researcher lived expertise, and reliance on visual syntactic data which could be lost if not explicitly prioritized via fine-tuning [49]. This validation was critical given the research context of an ongoing conflict and the team's personal connections, maintaining visual integrity via preserving social context and forefronting human interpretation of computational analysis.

4.3.4 Commitment to care. Given half of the research team's intimate relationship with the conflict (by virtue of being from Ukraine), this project exposed them to potential harm via imagery depicting the very real impacts on their friends, families, neighbors, and countrymen back home. It was critical to balance leveraging their deep, contextual knowledge with care for their wellbeing throughout this work. This motivated the team to manually analyze only a small portion of data and use the pipeline to scale analysis while minimizing human exposure.

"I thought the images as a whole would be more explicitly violent, but surprisingly, I was most struck by the infographics from DNR and Mariupol that were advising people about medical care etc. They felt insidious in a way that outright violence does not." - student memo from project (more memos in Appendix A)

We also took other measures of care before and during human analysis. First, weekly samples were screened by more senior researchers to avoid the worst content, like extensive gore. When training coders, we leveraged wellbeing resources from Case Study 1 (4.1.4) for handling traumatic imagery and alerted coders to key wellness signs (i.e. trouble sleeping). Coders coded only during work hours and while being co-located with each other and the

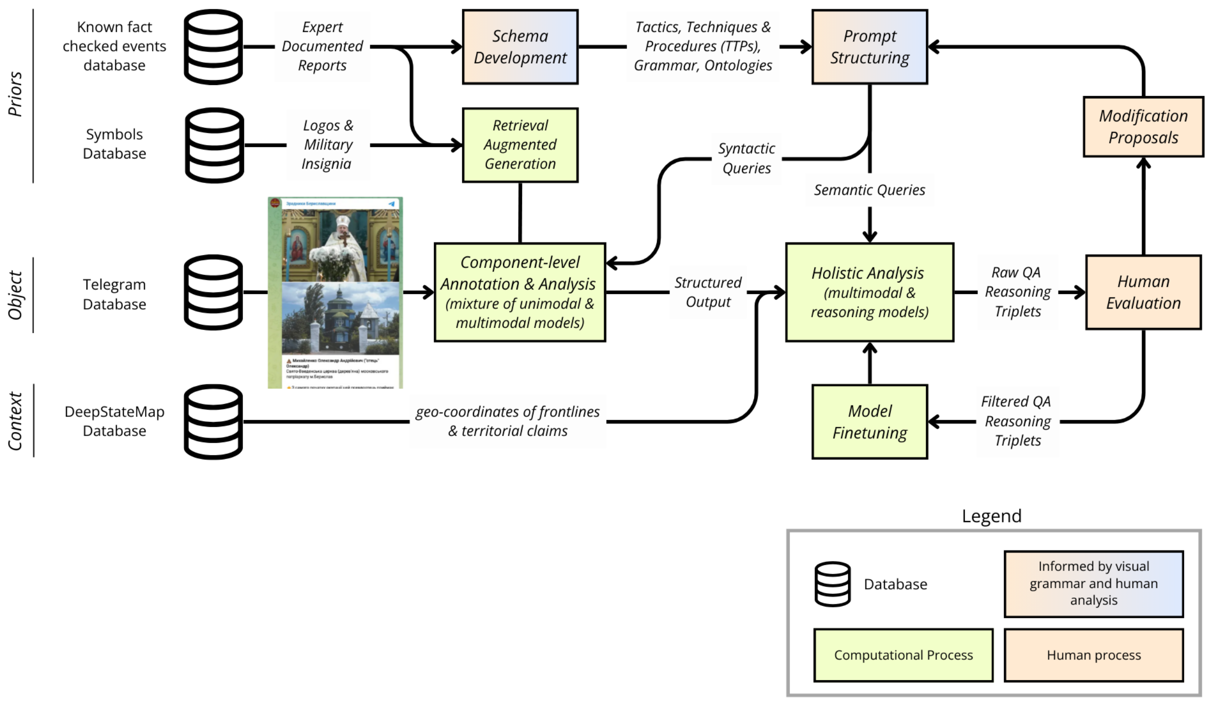

Figure 13: Based on tasks in Figure 10 and mapping in Figure 11, this pipeline diagram provides an operationalized view of the "Multimodal Data Pipeline" box from Figure 10. Data sources of priors like fact-checked information about the conflict and symbols (including those added by our human analysis) help to contextualize VLM and shape prompting structure informed by human analysis and visual grammar. Telegram data (objects) are analyzed first by syntactical analyses to explicitly represent visually observable features according to the schema. This provides an intermediate structured output, usable not just as semantic metadata for asset cataloguing, but also as a form of human-interpretable input passed on for downstream usage. The holistic analysis process leverages multimodal and reasoning foundation models to respond to specific semantic queries informed by these schemas and prompting structuring. The output of the "holistic analysis" then goes through human evaluation, leading to two critical applications that improve the visual pipeline: 1) potential modification feedback to the input prompt used by foundation models; 2) as a dataset, after filtering for quality, that can be used to finetune the same model deployed to perform the analysis

research manager. Additionally, as in Case 1, if coders were not able to conduct qualitative analysis or needed a break from it, they could be given other research tasks such as literature review or database management.

We leveraged Google Sheets with a translation plugin (see Figure 14) as a coding environment where coders had smaller previews of images so they could make an informed opt-out decision if needed before opening the larger view for analysis, another practice borrowed from Case Study 1.

Coders also filled out reflexive journal prompts provided by the first author (based again on Case Study 1) about their observations and their affective experiences analyzing this imagery. These were covered during the analysis meetings and not seen as a limitation, but a guiding force towards traits and tactics to be prioritized in the data pipeline.

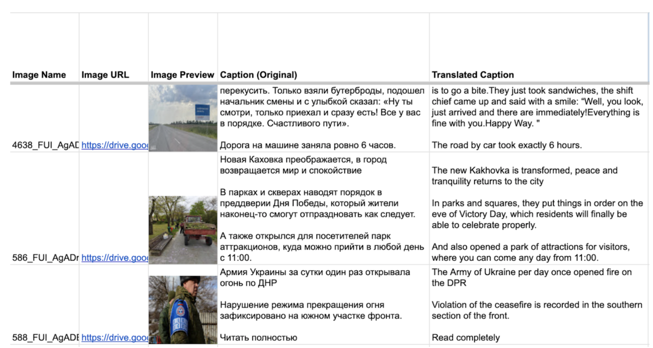

Image Name	Image URL	Image Preview	Caption (Original)	Translated Caption
4638_FUI_AgAD	https://drive.goo		перекусить. Только взяли бутерброды, подошел начальник смены и с улыбкой сказал: «Ну ты смотри, только приехал и сразу есть! Все у вас в порядке. Счастливого пути». Дорога на машине заняла ровно 6 часов.	is to go a bite.They just took sandwiches, the shift chief came up and said with a smile: "Well, you look, just arrived and there are immediately!Everything is fine with you.Happy Way. " The road by car took exactly 6 hours.
586_FUI_AgADr	https://drive.goo		Новая Каховка преображается, в город возвращается мир и спокойствие В парках и скверах наводят порядок в преддверии Дня Победы, который жители наконец-то смогут отпраздновать как следует. А также открылся для посетителей парк аттракционов, куда можно прийти в любой день с 11:00.	The new Kakhovka is transformed, peace and tranquility returns to the city In parks and squares, they put things in order on the eve of Victory Day, which residents will finally be able to celebrate properly. And also opened a park of attractions for visitors, where you can come any day from 11:00.
588_FUI_AgADE	https://drive.goo		Армия Украины за сутки один раз открывала огонь по ДНР Нарушение режима прекращения огня зафиксировано на южном участке фронта. Читать полностью	The Army of Ukraine per day once opened fire on the DPR Violation of the ceasefire is recorded in the southern section of the front. Read completely

Figure 14: Coders could see previews of images before opening the image URL for detailed analysis and also have post context in original language and English.

5 Discussion and implications

This paper lays out a research framework for building computationally-supported, human-centered pipelines for studying visual social data. We outlined three core framework features: visual grammars, human analysis, and computationally supported analysis situated in visual research methodologies. We detailed our framework's commitment to care, informed by our positionality as researchers of problematic information and drawing from a growing body of work around researcher protection. We illustrated this framework in three empirical case studies of problematic information: anti-immigrant visual propaganda, AI Slop, and war-time imagery.

In doing so, we have shown how our framework supports grounded, interpretivist mixed-methods research on visual social data. Below, we detail the principles and implications of adopting our framework and provide pragmatic advice for researchers. We start with explaining how researchers can maintain visual integrity, particularly when negotiating boundaries between human and computational analysis (5.1). We then describe computational and human infrastructures, and tensions, in caring for researchers (5.2). We end with implications for what it means to become, and to train, visually literate researchers in HCI and beyond (5.3).

5.1 Maintaining visual integrity in computationally supported analysis

Our methodological framework maintains a commitment to prioritizing and centering the visual. We offer **visual integrity** as a sensitizing concept for designing computationally-assisted, human-centered pipelines for analyzing visual social data, hoping it may help researchers negotiate boundaries between computational and human analysis. Sensitizing concepts, defined by Blumer [20], are theoretical tools used by researchers as guiding directions, not prescriptive definitions. In grounded theory, sensitizing concepts can be applied as priors to guide the researcher [33, 64]. We define visual integrity as centering inquiry around the visual, preserving both the visual form and human context of data across a research project. Other scholars have engaged with similar concepts, such as Manovich's "direct visualizations", in which:

"...data is reorganized into a new visual representation that preserves its original form. Usually, this does involve some data transformation such as changing data size...However, this is a reduction that is quantitative rather than qualitative. We don't substitute media objects by new objects..." [103, pg 12]

"Direct visualizations" preserve visual data's original form [103], as visual integrity seeks to. Other research approaches like metapictures [144] and visual rhythm representations [71] also preserve visual modality. These artifacts allow researchers to observe large amounts of visual social data with minimal or no transformations. However, these approaches remain agnostic to the social interpretation and meaning of data. With visual integrity, we seek to bridge how these approaches preserve visual modality and also consider the social context behind the data.

In practice, embodying visual integrity in research encompasses: 1) preserving visual modality as much as possible, 2) preserving the social context of visual data, and 3) prioritizing human interpretation and steering of analysis. We uphold this commitment across our case studies through how we chose to implement and apply our methods to visual social data and interpretation of results.

Towards the first point, Case Study 1 leveraged modular scripts to answer targeted visual inquiries, like trends in logos or fonts. In Case Study 2, computation focused on quantifying visual features like skintones and depicted subjects of interest across our dataset. In Case Study 3, we incorporated both visual-first, modular approaches and VLMs, which did transform visuals to textual representations for semantic analysis. However, VLMs were only utilized after extensive human reading of visual content and bounded by these interpretations. This brings us to our second point — preserving the social context of visual data. In Case Study 1, when computation surfaced non-visual data like metadata and provenance, it was in service of studying the social, participatory production and evolution of visual content, including data visualizations. In Case Study 2, many of the most culturally significant visual subjects to participants were not the highest engagement drivers, and may have been lost in a "big data" approach. By prioritizing visual integrity, we maintained these insights. And in Case Study 3, the lived expertise of researchers grounded the study in considering the most salient symbols and tasks for analysis. Regarding the third point, across all case studies, we focused on preserving an active, human viewing of the visual to steer and inform computationally supported analysis, which was then interpreted by researchers in dialogue with insights from human analysis. Visual integrity allowed us to negotiate boundaries between computational and human analysis and uplift human interpretation.

This guided negotiation is particularly beneficial when a researcher is tempted to adopt the trendiest new tool, a predicament faced across time and fields. In the 1960s, as computational tools spread in economics, Paarlberg noted some researchers became so computationally enamored that they sought problems to fit the tools, rather than tools that fit research questions, writing: "Here is a prestigious tool. Where can I find a problem or some data on which to use it?" [125, pg 1387]. A similar concern now preoccupies HCI, particularly with generative AI tools being increasingly applied in research [80, 82, 150]. Without careful human guidance, scholars note applying computational approaches (including AI) to social data can strip away the social from findings [18, 22, 31, 79]. In sensitive research contexts (i.e. trauma-related topics), deep human involvement can provide situated expertise for preserving the social in social data, yet may also benefit the most from machine support to protect researchers [37, 163]. At the same time, in some contexts particular computational tools may be severely limited. For example, GenAI models lack real-time context windows [74] for providing accurate responses regarding unfolding events pertinent to problematic information like active conflict and evolving propaganda. Additionally, models are targets of and susceptible to information operations [50], which is why we leveraged our own human labeled data for the VLM instance in Case Study 3, where we were studying discourse around an active conflict.

Bounding and steering computation, particularly AI, in research requires complex, situational decisions which depend on evolving technologies and field norms [82]. In defining norms, researchers must weigh risks of adding biases and invalidating lived experience, consent, privacy, and qualitative epistemologies [18, 52, 79, 80, 150].

Such norms are developing at a time of "AI hype" [14] and an increasingly competitive research landscape that may push scholars to rapidly integrate AI tooling without sufficient study design considerations to account for machine limitations to interpreting rich social data [80, 112, 154].

Our commitment to visual integrity and the methods used in service of it were influenced by the nature of our case studies and our interpretivist epistemologies. Our context surfaced a need to limit what human researchers saw and make best use of their labor, motivating the use of computational methods to carefully filter and select salient visual content to put in front of coders. Additionally, our interpretivist epistemology motivated us to center human analysis, leveraging the lived experience and expertise of researchers who were often closely tied to emotionally charged and pressing topics. Had this framework been developed in other contexts, this emphasis on human close reading, reflexivity, and care in bounding computation might not have been as central.

5.2 Computational and Human Infrastructures of Care for Visual Qualitative Research

As researchers of problematic information, our case studies involve distressing content that could be harmful. Hence, we center care in our framework, integrating care practices into the human and computational infrastructures supporting our work [95]. Drawing from Tseng et al.'s concepts of "care infrastructures" and "safety-as-enablement" focused on enabling contextual safety via computation in qualitative research [163, 164], we contribute to a growing body of work that seeks to protect researchers working with emotionally difficult or traumatic material [55, 57, 84, 108, 141, 163, 166]. Our conceptualization of care focuses on minimizing exposure to the "worst of the worst" and carefully selecting what content (including troubling content) warrants deep human analysis to answer research questions. This commitment, like visual integrity, emerges from the nature of our studies and our positionality. Nevertheless, we believe that considerations of care and optimal use of analytical labor is resonant to the broader research community, particularly given mounting researcher mental health concerns in HCI and other fields [96, 113, 136, 151].

5.2.1 Computation as a cyberinfrastructure of care. Computational infrastructures like PhotoDNA (for CSAM detection) or Azure-GoreDetection (for extreme violence) can be leveraged as a first, necessary step to protect researchers. Additional computational infrastructures, we argue, should be contextualized to the research project and isolate content that is possibly harmful and does not warrant human analysis. In Case Study 1, this was formalized with an emergent content category *hatespam*: hateful, repetitive content not salient enough to warrant human analysis. By filtering for slurs and known hate group logos, computational tools allowed us to quantify hatespam across our dataset without extensive researcher labor.

In our work, care often manifests by leveraging computation to reduce exposure to distressing content, both at individual and project levels. In coding environments, computation can enable small previews to allow coders to opt out of imagery particularly troubling to them, a practice we use throughout our case studies. Coding environments also support other best practices informed by psychology and journalism, like viewing content in grayscale [175]. At a project level, as in Case Studies 2 and 3, we chose to code a small sample of data due to its distressing nature, and leveraged computation to extend our insights (such as key subjects in Case Study 2 or the context of where logos appears in Case Study 3) across the larger dataset. In this extrapolation of human insights to larger datasets we build upon Tseng et al.'s trauma-informed qualitative analysis (TIQA) principles [163]. We bring a safety-as-enablement approach to visual data in qualitative coding environments and extending human hypotheses via computational analysis — leveraging computation to best utilize analytical labor, a mentally taxing but often rewarding task our coders experience.

5.2.2 Human infrastructure of care. We implemented several human infrastructures of care throughout our work, which we have distilled into pragmatic recommendations for researchers.

We adopted an ethos of safety-as-enablement and harm reduction in line with previous work [13, 163, 164], often via having a more senior or lead researcher as a stop gap. This was appropriate in our work, given these researchers had more access to resources and experience in handling this content. However, this may not always be the case, and teams may benefit from other care models like harm distribution (having team members review a little bit of content each) or more machine filtering of content [163]. We recommend documenting these assumptions as a foundational step in making a reporting and handling plan for particularly difficult content. We also recommend allocation of mental health support, such as university health services or mental health reimbursement, as valid and important research expenditures in grants and institutions. However, we understand many institutions and research teams will not be positioned to offer this support.

In this absence, we hope lead researchers will take the time to onboard team members thoroughly and leverage available wellness resources. In Case Study 1, we utilized free online resources from journalistic organizations [61, 77, 143]. We also had a mental health professional come to speak to our team about how this content could impact researchers and warning signs to be conscious of, such as trouble sleeping, irritability, and intense focus on the study topic outside of research tasks. These resources informed Case Studies 2 and 3. Other measures, such as one-on-one or team briefings about content and early teamed coding sessions of samples to assess risks, can help a team prepare.

When planning projects, considerations about coder burnout should be taken into account, such as incorporating "off weeks" and instilling flexibility to adjust coding loads and timing based on researcher wellbeing. During a project, analysis meetings should focus on the emotional impact of content as a core part of reflexivity and as valid, emergent knowledge produced by the study. As data analysis concludes, we recommend hosting closing interviews, follow-up check-ins, and thoughtful sharing of findings with qualitative coders to provide closure and context for their contributions and labor.

Developing proactive practices of care, both human and computational, is essential for long-term sustainability of research, particularly when studying content that can be emotionally charged or even harmful.

5.2.3 *Tensions and tradeoffs in care and research.*

Some research topics, even with robust care infrastructures, will be inherently harmful to researchers and require additional caution. Our framework may not be appropriate in such cases. For example, in studying non-consensual intimate imagery (often called "revenge porn") [138], interviewing case workers and survivors would be more appropriate than using our framework to focus on the visual imagery itself. In projects suited for our framework where visual content is difficult but still tenable to study, there are additional considerations to human and cyber infrastructures of care — and grappling with what happens when caring for researchers and conducting analysis becomes irreconcilable. One additional consideration is that care needs can change over the course of a project, particularly when data is collected in real time, potentially introducing new risks. As researchers, we are responsible for managing and clearly communicating such shifts, as per best practices in scientific transparency [65, 116]. However, in some cases, adjusting care may be insufficient and a study may need to pivot to other methods or stop. In our context, this could emerge in studying conspiratorial content, where repeated exposure (even in carefully conducted research) could destabilize one's sense of reality and have lasting personal, professional, and mental health impacts [12, 40, 56, 72, 149, 162]. In Case Study 1 and 3 on US immigration and the Russian-Ukrainian war, both active sites of information operations and conspiracy theories, this was a possibility. Had we run Case Study 1 in 2025 — amid active ICE raids and chaotic information spaces about immigration — the potential harms to researchers who identified with targeted groups may have been too high to continue.

But examining content is a central commitment of this framework. In challenging cases like the ones mentioned above, researchers must assess when research benefits do not justify harms to themselves and their peers. This is a highly contextual and difficult decision. We encourage researchers to design and adopt exit plans prior to embarking upon studies to help navigate this choice. Exit plans could help researchers pre-identify scenarios when they would stop a study. If a study is stopped, exit plans can help facilitate handling data and in-progress findings or pivots to other methods. In corporate settings, exit plans from terminated ventures include knowledge transfers [2]. We encourage the HCI, and other research communities, to develop practices to value such knowledge transfers as empirical and methodological contributions that can inform best practices when a study has to pivot for safety reasons.

5.3 Implications of visual methods and visually literate researchers

Akin to research frameworks like trauma-informed computing or feminist HCI [9, 37], our work is not fully prescriptive. Rather, we provide a framework and best practices informed by our empirical work for adaptation in pursuit of visually informed, safe, and rigorous mixed methods research. We have outlined how our framework helps researchers conduct visual research that generates timely insights and new knowledge in an increasingly visual world, while prioritizing visual integrity (5.1) and researcher wellbeing (5.2). But are there implications for visual research, literacy, and methods that extend beyond how we do research? We think yes, and that it starts with researchers and research trainees.

Part of becoming a "knowing" visual researcher [23] is becoming more visually literate. Pauwels defines visual literacy as interlinking competencies of engaging with and understanding the production and artifacts of visual culture as a participatory, sociotechnical process [127]. This posits visual literacy as a learnable and valuable competence for scientific discovery and communication across several fields [126, 129] and as an engaged citizen in participatory culture where the visual looms large [127]. As such, visually literate researchers stand to produce knowledge more translatable to solving problems [123] in visual spaces and influencing visual cultures that are becoming increasingly mediated by digital systems [99, 137, 156].

We hope this framework not only empowers researchers to design and execute visual research, but also helps in training future researchers. This framework demands a human infrastructure, which we have largely implemented via pedagogical participatory practices with students as per [165]. At a time where scholars call for visual literacy to be taught across several levels of education to prepare students for a visually mediated world [8, 73, 127], we see an important opportunity, as we have demonstrated in our work, to further integrate visual research education into HCI programs and training. Additionally, outputs from our framework could be integrated into HCI research toolkits that include components like visual grammar templates, coding schemas, and ready to use analysis scripts to empower researchers to conduct visual research [93].

6 Limitations and Future Work

Our framework does not have quantifiable or fully generalizable metrics. We follow traditions of qualitative and interpretivist research focusing on transferability over generalizability [38, 54] but acknowledge this could make our framework inappropriate for some research questions.

We draw from a growing body of work centering caring for researchers who study harmful content. Other important research examines how to responsibly study and present such material without amplifying fringe narratives or causing downstream harm to populations depicted in it [15, 29, 57, 134, 148]. Although we do not directly engage this second body of work, we see bridging care for both represented populations and researchers as a rich direction for future research.

Although Case Study 1 adapted this framework to short form video, more work is needed to streamline and contend with the nuances of video compared to still imagery. Our case studies, although varied and encompassing multiple languages, focus on Global North contexts and are conducted by English-speaking researchers in American and European universities. Adjustments will be needed to adapt this framework to other cultural contexts. Additionally, other methods of human analysis, such as crowdsourced labeling of visual data or design activities with participants, have yet to be explored. We hope this framework is adapted in future case studies beyond problematic information to generate new insights and implications.

Future work on this framework may involve more extensive analysis, such as pedagogical research of curriculum development and formalized measuring of visual literacies in students before and after being trained in this methodology. This contribution also did not focus on developing and packaging components of our framework into an HCI research toolkit [93], which presents another avenue for future work. And, as more and more AI tools and architectures become multimodal, we hope that future work explores how to incorporate these techniques and tools into visual research while maintaining visual integrity and human insights.

7 Conclusion

We presented a research framework for conducting computationally-supported, human-centered research on visual social data that centers caring for researchers. Across three empirical case studies, we showed its effectiveness and the types of insights enabled by three core framework componenets: visual grammars, human analysis, and computationally supported analysis. To this end, we offer our framework, a sensitizing concept of visual integrity, and best practices for caring for researchers and their labor to the research community. We hope by doing so that researchers may feel more empowered to approach research challenges and opportunities while studying visual social data.

Acknowledgments

This work draws from three empirical case studies of which several colleagues and funders contributed to. However, any opinions, findings, and conclusions or recommendations expressed in this paper are those of this author(s) and do not necessarily reflect the views of below named colleagues and funders.

Nina Lutz led Case Study 1 with support from Joseph Schafer, Priya Dhawka, and Kate Starbird, along with several other colleagues – involving 16 (largely undergraduate) research assistants via University of Washington's Human Centered Design and Engineering's Directed Research Group (DRG) program across four academic quarters and analyzing over 1500 media elements. Nina Lutz is immensely grateful for all of these students' contributions and to the department for this participatory learning model.

Priya Dhawka and Nina Lutz are co-first authors and equal contributors on the data visualization case study that spurred from Case Study 1, with Kate Starbird as senior author. This work was published at CSCW 2025 [43] and awarded a Best Paper award.

Nina Lutz led Case Study 2, with contributions from Joseph Schafer and Kate Starbird. Not on this paper are colleagues Julie Vera and Sourojit Ghosh are co-authors of the Case Study 2 work.

Phil Tinn led Case Study 3, with generous (including financial) support from the MIT Ukraine and UROP Programs. This work was advised by Dr. Elizabeth Wood, Dr. Svitlana Krayska, Dr. Hosea Siu. Contributors to Case Study 3 not involved in this methods paper include Tsan-Yu Wu, Alicia Lin, Sofiia Shypovych, Olha Volianyk, Allen Song, Jon Roozenbeek, and Yara Kyrychenko.

Nina Lutz, Joseph Schafer, and Kate Starbird were supported by the University of Washington's Center for an Informed Public, the John S. and James L. Knight Foundation (G-2019–58788), the William and Flora Hewlett Foundation (2023–02789), and the National Science Foundation (2120496) during their work on Case Studies 1 and 2. The William and Flora Hewlett Foundation contributed to paying coders small stipends in Summer 2024 when they were not in for-credit DRG seminars to help keep analysis going across the calendar year of 2024.

Joseph Schafer is supported by an NSF Graduate Research Fellowship (GRFP) (DGE-2140004). Nina Lutz is supported by an NSF CSGrad4US Fellowship (2313998).

Authors would like to thank Amelia Lee Doğan who provided valuable feedback on a draft of this manuscript and whose suggestion of including memos greatly enhanced this paper.

Lastly, the authors thank the reviewers who provided generous and invaluable feedback to greatly improve our manuscript across two CHI cycles.

References

[1] Salma M. Abdalla, Gregory H. Cohen, Shailesh Tamrakar, Shaffi Fazaludeen Koya, and Sandro Galea. 2021. Media Exposure and the Risk of Post-Traumatic Stress Disorder Following a Mass traumatic Event: An In-silico Experiment. *Frontiers in Psychiatry* 12 (Nov. 2021), 674263. doi:10.3389/fpsyt.2021.674263

[2] Joseph Amankwah-Amoah. 2011. Learning from the failures of others: The effects of post-exit knowledge spillovers on recipient firms. *Journal of Workplace Learning* 23, 6 (2011), 358–375.

[3] AoIR Risky Research Working Group. 2025. Risky Research: An AoIR Guide to Researcher Protection and Safety. https://aoir.org/riskyresearchguide/ Published: The Association of Internet Researchers.

[4] Ahmer Arif, Leo Graiden Stewart, and Kate Starbird. 2018. Acting the Part: Examining Information Operations Within #BlackLivesMatter Discourse. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (Nov. 2018), 1–27. doi:10.1145/3274289

[5] Alice Ashcroft, Beatriz Severes, and Melissa Martinez-Perez. 2025. Suggested Prompts for Reflexivity: Navigating Intersectionality in HCI and CSCW Research. *Interacting with Computers* 37 (March 2025), iwaf007. Issue 5. doi:10.1093/iwc/iwaf007

[6] Tina Askanius. 2021. On frogs, monkeys, and execution memes: Exploring the humor-hate nexus at the intersection of neo-Nazi and alt-right movements in Sweden. *Television & new media* 22, 2 (2021), 147–165.

[7] Illya Bakurov, Marco Buzzelli, Raimondo Schettini, Mauro Castelli, and Leonardo Vanneschi. 2022. Structural similarity index (SSIM) revisited: A data-driven approach. *Expert Systems with Applications* 189 (March 2022), 116087. doi:10.1016/j.eswa.2021.116087

[8] Melissa Ballesteros-Mejía and María Angélica Madero. 2024. Mapping visual-based methods: an interdisciplinary approach. *Journal of Visual Literacy* 43, 4 (Oct. 2024), 285–307. doi:10.1080/1051144X.2024.2420304

[9] Shaowen Bardzell. 2010. Feminist HCI: taking stock and outlining an agenda for design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Atlanta Georgia USA, 1301–1310. doi:10.1145/1753326.1753521

[10] Shaowen Bardzell and Jeffrey Bardzell. 2011. Towards a feminist HCI methodology: social science, feminism, and HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Vancouver BC Canada, 675–684. doi:10.1145/1978942.1979041

[11] Christine A Barry, Nicky Britten, Nick Barber, Colin Bradley, and Fiona Stevenson. 1999. Using reflexivity to optimize teamwork in qualitative research. *Qualitative health research* 9, 1 (1999), 26–44. Publisher: Sage Publications Sage CA: Thousand Oaks, CA.

[12] Matthew A Baum, James N Druckman, Matthew D Simonson, Jennifer Lin, and Roy H Perlis. 2024. The political consequences of depression: how conspiracy beliefs, participatory inclinations, and depression affect support for political violence. *American journal of political science* 68, 2 (2024), 575–594.

[13] Rosanna Bellini, Emily Tseng, Noel Warford, Alaa Daffalla, Tara Matthews, Sunny Consolvo, Jill Palzkill Woelfer, Patrick Gage Kelley, Michelle L Mazurek, Dana Cuomo, et al. 2024. Sok: Safer digital-safety research involving at-risk users. In *2024 IEEE Symposium on Security and Privacy (SP)*. IEEE, IEEE, San Francisco, CA, USA, 635–654.

[14] Emily M. Bender and Alex Hanna. 2025. *The AI Con: How to Fight Big Tech's Hype and Create the Future We Want*. Harper, New York.

[15] Sebastian Benthall, Seda Gürses, and Helen Nissenbaum. 2017. Contextual Integrity through the Lens of Computer Science. *Foundations and Trends® in Privacy and Security* 2, 1 (Dec. 2017), 1–69. doi:10.1561/3300000016 Publisher: Now Publishers, Inc..

[16] John Berger. 1972. *Ways of Seeing*. Penguin Books, London.

[17] Melissa Bica, Leysia Palen, and Chris Bopp. 2017. Visual Representations of Disaster. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17)*. Association for Computing Machinery, New York, NY, USA, 1262–1276. doi:10.1145/2998181.2998212

[18] Abeba Birhane. 2021. The impossibility of automating ambiguity. *Artificial life* 27, 1 (2021), 44–61.

[19] Eli Blevis. 2016. Being Photo-Visual in HCI and Design. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*. ACM, Brisbane QLD

Australia, 983–995. doi:10.1145/2901790.2901863
[20] Herbert Blumer. 1954. What is Wrong with Social Theory? *American Sociological Review* 19, 1 (1954), 3–10. doi:10.2307/2088165
[21] Mireia Bolíbar. 2016. Macro, meso, micro: Broadening the 'social'of social network analysis with a mixed methods approach. *Quality & Quantity* 50, 5 (2016), 2217–2236. Publisher: Springer.
[22] Danah Boyd and Kate Crawford. 2011. Six Provocations for Big Data. doi:10.2139/ssrn.1926431
[23] Virginia Braun and Victoria Clarke. 2023. Toward good practice in thematic analysis: Avoiding common problems and be(com)ing a *knowing* researcher. *International Journal of Transgender Health* 24, 1 (Jan. 2023), 1–6. doi:10.1080/26895269.2022.2129597
[24] Alan Bryman. 2008. Of methods and methodology. *Qualitative Research in Organizations and Management: An International Journal* 3, 2 (Aug. 2008), 159–168. doi:10.1108/17465640810900568
[25] David Buckingham. 2009. 'Creative' visual methods in media research: possibilities, problems and proposals. *Media, Culture & Society* 31, 4 (July 2009), 633–652. doi:10.1177/0163443709335280 Publisher: SAGE Publications Ltd.
[26] Tina Budzise-Weaver. 2016. Developing a Qualitative Coding Analysis of Visual Artwork for Humanities Research. *Digital Humanities Quarterly* 010, 4 (Nov. 2016).
[27] Kate Bukowski and Stephen Buetow. 2011. Making the invisible visible: A Photovoice exploration of homeless women's health and lives in central Auckland. *Social Science & Medicine* 72, 5 (March 2011), 739–746. doi:10.1016/j.socscimed.2010.11.029
[28] Monica Bulger and Patrick Davison. 2018. The promises, challenges, and futures of media literacy. *Journal of media literacy education* 10, 1 (2018), 1–21.
[29] Lauren E Cagle. 2019. Surveilling strangers: The disciplinary biopower of digital genre assemblages. *Computers and Composition* 52 (2019), 67–78.
[30] Scott Allen Cambo and Darren Gergle. 2022. Model Positionality and Computational Reflexivity: Promoting Reflexivity in Data Science. In *CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans LA USA, 1–19. doi:10.1145/3491102.3501998
[31] Hjalmar Bang Carlsen and Snorre Ralund. 2022. Computational grounded theory revisited: From computer-led to computer-assisted text analysis. *Big Data & Society* 9, 1 (Jan. 2022), 20539517221080146. doi:10.1177/20539517221080146 Publisher: SAGE Publications Ltd.
[32] Kathy Charmaz, , and Robert Thornberg. 2021. The pursuit of quality in grounded theory. *Qualitative Research in Psychology* 18, 3 (July 2021), 305–327. doi:10.1080/14780887.2020.1780357 Publisher: Routledge _eprint: https://doi.org/10.1080/14780887.2020.1780357.
[33] Kathy Charmaz. 2006. *Constructing grounded theory: a practical guide through qualitative analysis* (repr ed.). Sage, Los Angeles.
[34] Kathy Charmaz. 2014. *Constructing grounded theory*. SAGE publications Ltd, Thousand Oaks, CA, USA.
[35] Kathy Charmaz. 2020. "With constructivist grounded theory you can't hide": Social justice research and critical inquiry in the public sphere. *Qualitative inquiry* 26, 2 (2020), 165–176. Publisher: SAGE Publications Sage CA: Los Angeles, CA.
[36] Jilin Chen, Gary Hsieh, Jalal U. Mahmud, and Jeffrey Nichols. 2014. Understanding individuals' personal values from social media word use. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing (CSCW '14)*. Association for Computing Machinery, New York, NY, USA, 405–414. doi:10.1145/2531602.2531608
[37] Janet X. Chen, Allison McDonald, Yixin Zou, Emily Tseng, Kevin A Roundy, Acar Tamersoy, Florian Schaub, Thomas Ristenpart, and Nicola Dell. 2022. Trauma-Informed Computing: Towards Safer Technology Experiences for All. In *CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans LA USA, 1–20. doi:10.1145/3491102.3517475
[38] Andy Crabtree. 2025. H is for human and how (not) to evaluate qualitative research in HCI. *Human–Computer Interaction* (2025). Publisher: Taylor & Francis.
[39] Viorela Dan, Britt Paris, Joan Donovan, Michael Hameleers, Jon Roozenbeek, Sander van der Linden, and Christian von Sikorski. 2021. Visual mis-and disinformation, social media, and democracy. *Journalism & Mass Communication Quarterly* 98, 3 (2021), 641–664.
[40] David De Coninck, Thomas Frissen, Koen Matthijs, Leen d'Haenens, Grégoire Lits, Olivier Champagne-Poirier, Marie-Eve Carignan, Marc D David, Nathalie Pignard-Cheynel, Sébastien Salerno, et al. 2021. Beliefs in conspiracy theories and misinformation about COVID-19: Comparative perspectives on the role of anxiety, depression and exposure to and trust in information sources. *Frontiers in psychology* 12 (2021), 646394.
[41] Audrey Desjardins, Cayla Key, Heidi R. Biggs, and Kelsey Aschenbeck. 2019. Bespoke Booklets: A Method for Situated Co-Speculation. In *Proceedings of the 2019 on Designing Interactive Systems Conference (DIS '19)*. Association for Computing Machinery, New York, NY, USA, 697–709. doi:10.1145/3322276.3322311
[42] Michael Ann DeVito. 2021. Adaptive folk theorization as a path to algorithmic literacy on changing platforms. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–38.
[43] Priya Dhawka, Nina Lutz, and Kate Starbird. 2025. Data Visualizations as Propaganda: Tracing Lineages, Provenance, and Political Framings in Online Anti-Immigrant Discourse. *Proc. ACM Hum.-Comput. Interact.* 9, 7, Article CSCW393 (Oct. 2025), 47 pages. doi:10.1145/3757574
[44] Virginia Dickson-Swift, Erica L James, and Sandra Kippen. 2005. Do university ethics committees adequately protect public health researchers? *Australian and New Zealand Journal of Public Health* 29, 6 (2005), 576–579.
[45] Renée DiResta and Josh A. Goldstein. 2024. How spammers and scammers leverage AI-generated images on Facebook for audience growth. *Harvard Kennedy School Misinformation Review* (Aug. 2024). doi:10.37016/mr-2020-151
[46] Joan E Dodgson. 2019. Reflexivity in qualitative research. *Journal of human lactation* 35, 2 (2019), 220–222. Publisher: SAGE Publications Sage CA: Los Angeles, CA.
[47] Allison Druin. 1999. Cooperative inquiry: developing new technologies for children with children. In *Proceedings of the SIGCHI conference on Human factors in computing systems the CHI is the limit - CHI '99*. ACM Press, Pittsburgh, Pennsylvania, United States, 592–599. doi:10.1145/302979.303166
[48] DSD-Tom. 2022. Best Practices for Writing Image Descriptions. https://apln.ca/best-practices-for-writing-image-descriptions/
[49] Sri Harsha Dumpala, David Arps, Sageev Oore, Laura Kallmeyer, and Hassan Sajjad. 2024. Seeing Syntax: Uncovering Syntactic Learning Limitations in Vision-Language Models. *arXiv preprint arXiv:2412.08111* (2024).
[50] Huw Dylan and Elena Grossfeld. 2025. Revisionist future: Russia's assault on large language models, the distortion of collective memory, and the politics of eternity. *Dialogues on Digital Society* 1, 3 (2025), 407–412. doi:10.1177/29768640251377941
[51] Chloe Eghtebas, Gudrun Klinker, Susanne Boll, and Marion Koelle. 2023. Co-Speculating on Dark Scenarios and Unintended Consequences of a Ubiquitous(ly) Augmented Reality. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference (DIS '23)*. Association for Computing Machinery, New York, NY, USA, 2392–2407. doi:10.1145/3563657.3596073
[52] Jessica L. Feuston and Jed R. Brubaker. 2021. Putting Tools in Their Place: The Role of Time and Perspective in Human-AI Collaboration for Qualitative Analysis. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (Oct. 2021), 1–25. doi:10.1145/3479856
[53] Casey Fiesler, Nathan Beard, and Brian C Keegan. 2020. No robots, spiders, or scrapers: Legal and ethical regulation of data collection methods in social media terms of service. In *Proceedings of the international AAAI conference on web and social media*, Vol. 14. 187–196.
[54] Ellie Fossey, Carol Harvey, Fiona McDermott, and Larry Davidson. 2002. Understanding and evaluating qualitative research. *Australian and New Zealand journal of psychiatry* 36, 6 (2002), 717–732. Publisher: Taylor & Francis.
[55] Diana Freed, Sam Havron, Emily Tseng, Andrea Gallardo, Rahul Chatterjee, Thomas Ristenpart, and Nicola Dell. 2019. "Is my phone hacked?" Analyzing Clinical Computer Security Interventions with Survivors of Intimate Partner Violence. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 1–24. doi:10.1145/3359304
[56] Daniel Freeman and Richard P Bentall. 2017. The concomitants of conspiracy concerns. *Social psychiatry and psychiatric epidemiology* 52, 5 (2017), 595–604.
[57] Christian Fuchs. 2018. 'Dear Mr. Neo-Nazi, can you please give me your informed consent so that I can quote your fascist tweet?': Questions of social media research ethics in online ideology critique. In *The Routledge companion to media and activism*. Routledge, 385–394.
[58] Noam Gal, Limor Shifman, and Zohar Kampf. 2016. "It gets better": Internet memes and the construction of collective identity. *New media & society* 18, 8 (2016), 1698–1714.
[59] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2018. Quantifying Controversy on Social Media. *Trans. Soc. Comput.* 1, 1 (Jan. 2018), 3:1–3:27. doi:10.1145/3140565
[60] William Gaver and Jacob Beaver. 2006. The Presence Project: Helping Older People Engage with Their Local Communities. In *Networked Neighbourhoods*, Patrick Purcell (Ed.). Springer London, London, 345–371. doi:10.1007/1-84628-601-8_14
[61] GCJT and Trauma Aware Journalism 2009. GCJT and Trauma Aware Journalism. https://gcjt.org/about/mission-history
[62] Clifford Geertz. 2008. "Thick Description: Toward an Interpretive Theory of Culture". In *The Cultural Geography Reader*. Routledge. Num Pages: 11.
[63] Sourojit Ghosh and Aylin Caliskan. 2023. 'Person' == Light-skinned, Western Man, and Sexualization of Women of Color: Stereotypes in Stable Diffusion. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 6971–6985. doi:10.18653/v1/2023.findings-emnlp.465
[64] Barney G. Glaser and Anselm L. Strauss. 1967. *The discovery of grounded theory: strategies for qualitative research*. Routledge, London New York.

[65] K. Hammarberg, M. Kirkman, and S. de Lacey. 2016. Qualitative research methods: when to use them and how to judge them. *Human Reproduction* 31, 3 (March 2016), 498–501. doi:10.1093/humrep/dev334

[66] Donna Haraway. 1988. Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. *Feminist Studies* 14, 3 (1988), 575–599. doi:10.2307/3178066 Publisher: Feminist Studies, Inc..

[67] Douglas Harper. 1988. Visual sociology: Expanding sociological vision. *The american sociologist* 19, 1 (1988), 54–70. Publisher: Springer.

[68] Patricia Snell Herzog. 2019. Visual Research Methods: Integrating Images in the Study of Social Problems. In *Researching Social Problems* (1 ed.). Routledge, 172–187. doi:10.4324/9781315107882-10

[69] Tim Highfield and Tama Leaver. 2016. Instagrammatics and digital methods: Studying visual social media, from selfies and GIFs to memes and emoji. *Communication research and practice* 2, 1 (2016), 47–62.

[70] Renee Hobbs and Amy Jensen. 2009. The past, present, and future of media literacy education. *Journal of media literacy education* 1, 1 (2009), 1.

[71] Nadav Hochman and Raz Schwartz. 2012. Visualizing Instagram: Tracing Cultural Visual Rhythms. *Proceedings of the International AAAI Conference on Web and Social Media* 6, 4 (2012), 6–9. doi:10.1609/icwsm.v6i4.14361

[72] E Alison Holman, Dana Rose Garfin, and Roxane Cohen Silver. 2024. It matters what you see: Graphic media images of war and terror may amplify distress. *Proceedings of the National Academy of Sciences* 121, 29 (2024), e2318465121. Publisher: National Academy of Sciences.

[73] Ying Hu and Gwo-Jen Hwang. 2024. Cultivating visual literacy and critical thinking tendency with technological knowledge organizing supports: a concept mapping-based online problem-posing approach. *Educational technology research and development* 72, 6 (Dec. 2024), 3217–3240. doi:10.1007/s11423-024-10394-6

[74] Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. 2025. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems* 43, 2 (2025), 1–55.

[75] Jina Huh-Yoo and Emilee Rader. 2020. It's the Wild, Wild West: Lessons learned from IRB members' risk perceptions toward digital research data. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW1 (2020), 1–22.

[76] Hazem Ibrahim, HyunSeok Daniel Jang, Nouar Aldahoul, Aaron R. Kaufman, Talal Rahwan, and Yasir Zaki. 2025. TikTok's recommendations skewed towards Republican content during the 2024 U.S. presidential race. doi:10.48550/arXiv.2501.17831 arXiv:2501.17831 [cs].

[77] International Committee of the Red Cross (ICRC). 2022. How to deal with traumatic imagery as a communicator. https://www.icrc.org/en/document/how-deal-traumatic-imagery

[78] Caroline Jack. 2017. Lexicon of Lies: Terms for Problematic Information. *Data & Society* (2017).

[79] Jialun Aaron Jiang, Kandrea Wade, Casey Fiesler, and Jed R. Brubaker. 2021. Supporting Serendipity: Opportunities and Challenges for Human-AI Collaboration in Qualitative Analysis. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (April 2021), 1–23. doi:10.1145/3449168

[80] Shivani Kapania, William Agnew, Motahhare Eslami, Hoda Heidari, and Sarah E Fox. 2025. Simulacrum of Stories: Examining Large Language Models as Qualitative Research Participants. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 489, 17 pages. doi:10.1145/3706598.3713220

[81] Shivani Kapania, Alex S Taylor, and Ding Wang. 2023. A hunt for the Snark: Annotator Diversity in Data Practices. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3544548.3580645

[82] Shivani Kapania, Ruiyi Wang, Toby Jia-Jun Li, Tianshi Li, and Hong Shen. 2025. 'I'm Categorizing LLM as a Productivity Tool': Examining Ethics of LLM Use in HCI Research Practices. *Proceedings of the ACM on Human-Computer Interaction* 9, 2 (May 2025), 1–26. doi:10.1145/3711000

[83] Katharina Karcher, Joanne McCuaig, and Sophie King-Hill. 2024. (Self-) reflection/reflexivity in sensitive, qualitative research: A scoping review. *International Journal of Qualitative Methods* 23 (2024).

[84] Brit Kelley and Stephanie Weaver. 2020. Researching People who (Probably) Hate You: When Practicing "Good" Ethics Means Protecting Yourself. *Computers and Composition* 56 (June 2020), 102567. doi:10.1016/j.compcom.2020.102567

[85] Jason Koebler. 2024. Facebook's Algorithm Is Boosting AI Spam That Links to AI-Generated, Ad-Laden Click Farms. https://www.404media.co/facebooks-algorithm-is-boosting-ai-spam-that-links-to-ai-generated-ad-laden-click-farms/

[86] Lindah Kotut, Neelma Bhatti, Taha Hassan, Derek Haqq, and Morva Saaty. 2024. Griot-Style Methodology: Longitudinal Study of Navigating Design With Unwritten Stories. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–14. doi:10.1145/3613904.3642682

[87] Gunther Kress and Theo van Leeuwen. 2006. *Reading images: the grammar of visual design* (2e éd ed.). Routledge, Abingdon.

[88] Gunther R. Kress, Theo Van Leeuwen, and Theo van Leeuwen. 1996. *Reading Images: The Grammar of Visual Design*. Psychology Press.

[89] Klaus Krippendorff. 2011. Computing Krippendorff's alpha-reliability. (2011).

[90] Girish Kulkarni, Visruth Premraj, Vicente Ordonez, Sagnik Dhar, Siming Li, Yejin Choi, Alexander C Berg, and Tamara L Berg. 2013. Babytalk: Understanding and generating simple image descriptions. *IEEE transactions on pattern analysis and machine intelligence* 35, 12 (2013), 2891–2903.

[91] Yara Kyrychenko, Tymofii Brik, Sander van der Linden, and Jon Roozenbeek. 2024. Social identity correlates of social media engagement before and after the 2022 Russian invasion of Ukraine. *Nature Communications* 15, 1 (Oct. 2024), 8127. doi:10.1038/s41467-024-52179-8 Publisher: Nature Publishing Group.

[92] Kumud Lakara, Juil Sock, Christian Rupprecht, Philip Torr, John Collomosse, and Christian Schroeder de Witt. 2024. MAD-Sherlock: Multi-Agent Debates for Out-of-Context Misinformation Detection. doi:10.48550/arXiv.2410.20140 arXiv:2410.20140 [cs] version: 1.

[93] David Ledo, Steven Houben, Jo Vermeulen, Nicolai Marquardt, Lora Oehlberg, and Saul Greenberg. 2018. Evaluation Strategies for HCI Toolkit Research. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–17. doi:10.1145/3173574.3173610

[94] Crystal Lee, Tanya Yang, Gabrielle D Inchoco, Graham M. Jones, and Arvind Satyanarayan. 2021. Viral Visualizations: How Coronavirus Skeptics Use Orthodox Data Practices to Promote Unorthodox Science Online. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3411764.3445211

[95] Charlotte P. Lee, Paul Dourish, and Gloria Mark. 2006. The human infrastructure of cyberinfrastructure. In *Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*. ACM, Banff Alberta Canada, 483–492. doi:10.1145/1180875.1180950

[96] Adam Levy. 2024. Mind matters: investigating academia's 'mental health crisis'. *Nature* (2024). doi:10.1038/d41586-024-04240-1 Nature Careers podcast transcript.

[97] Yiyi Li and Ying Xie. 2020. Is a picture worth a thousand words? An empirical study of image content and social media engagement. *Journal of marketing research* 57, 1 (2020), 1–19.

[98] Alex Jiahong Lu, Shruti Sannon, Cameron Moy, Savana Brewer, Jaye Green, Kisha N Jackson, Daivon Reeder, Camaria Wafer, Mark S. Ackerman, and Tawanna R Dillahunt. 2023. Participatory Noticing through Photovoice: Engaging Arts- and Community-Based Approaches in Design Research. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference (DIS '23)*. Association for Computing Machinery, New York, NY, USA, 2489–2508. doi:10.1145/3563657.3596041

[99] Nina Lutz and Cecilia Aragon. 2024. "We're not all construction workers": Algorithmic Compression of Latinidad on TikTok. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW2, Article 480 (Nov. 2024), 31 pages. doi:10.1145/3687019

[100] Nina Lutz, Jordyn W. Padzensky, and Joseph S. Schafer. 2024. Working with Color: How Color Quantization Can Aid Researchers of Problematic Information. In *Companion Publication of the 2024 Conference on Computer-Supported Cooperative Work and Social Computing* (San Jose, Costa Rica) *(CSCW Companion '24)*. Association for Computing Machinery, New York, NY, USA, 275–282. doi:10.1145/3678884.3681863

[101] Jim Maddock, Kate Starbird, Haneen J. Al-Hassani, Daniel E. Sandoval, Mania Orand, and Robert M. Mason. 2015. Characterizing Online Rumoring Behavior Using Multi-Dimensional Signatures. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. ACM, Vancouver BC Canada, 228–241. doi:10.1145/2675133.2675280

[102] Lydia Manikonda and Munmun De Choudhury. 2017. Modeling and Understanding Visual Attributes of Mental Health Disclosures in Social Media. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, Denver Colorado USA, 170–181. doi:10.1145/3025453.3025932

[103] Lev Manovich. 2010. What is Visualization? *paj:The Journal of the Initiative for Digital Humanities, Media, and Culture* 2, 1 (Dec. 2010). https://paj-ojs-tamu.tdl.org/paj/article/view/19 Number: 1.

[104] Lev Manovich. 2011. Trending: The promises and the challenges of big social data. *Debates in the digital humanities* 2, 1 (2011), 460–475.

[105] Lev Manovich. 2016. The Science of Culture? Social Computing, Digital Humanities and Cultural Analytics. *Journal of Cultural Analytics* 1, 1 (May 2016). doi:10.22148/16.004

[106] Lev Manovich. 2020. *Cultural analytics*. MIT Press.

[107] Jari Martikainen and Inari Sakki. 2021. Visual (de) humanization: construction of Otherness in newspaper photographs of the refugee crisis. *Ethnic and Racial Studies* 44, 16 (2021), 236–266.

[108] Adrienne L. Massanari. 2018. Rethinking Research Ethics, Power, and the Risk of Visibility in the Era of the "Alt-Right" Gaze. *Social Media + Society* 4, 2 (April 2018), 2056305118768302. doi:10.1177/2056305118768302 Publisher: SAGE Publications Ltd.

[109] Hana Matatov, Mor Naaman, and Ofra Amir. 2022. Stop the [image] steal: The role and dynamics of visual content in the 2020 us election misinformation campaign. *Proceedings of the ACM on Human-computer Interaction* 6, CSCW2 (2022), 1–24.

[110] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 1–23. doi:10.1145/3359174

[111] Richard McGrath. 2021. Journalling and memoing: Reflexive qualitative research tools. In *Handbook of qualitative research methodologies in workplace contexts*. Edward Elgar Publishing, 245–262.

[112] Lisa Messeri and M. J. Crockett. 2024. Artificial intelligence and illusions of understanding in scientific research. *Nature* 627, 8002 (March 2024), 49–58. doi:10.1038/s41586-024-07146-0 Publisher: Nature Publishing Group.

[113] Abraham Mhaidli and Kat Roemmich. 2024. Overworking in HCI: A Reflection on Why We Are Burned Out, Stressed, and Out of Control; and What We Can Do About It. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, Article 561, 10 pages. doi:10.1145/3613905.3644052

[114] Nicholas Mirzoeff. 1999. *An introduction to visual culture*. Vol. 274. Routledge London.

[115] Claudia Mitchell, Linda Theron, Jean Stuart, Ann Smith, and Zachariah Campbell. 2011. Drawings as Research Method. In *Picturing Research: Drawing as Visual Methodology*, Linda Theron, Claudia Mitchell, Claudia Mitchell, Ann Smith, Ann Smith, and Jean Stuart (Eds.). SensePublishers, Rotterdam, 19–36. doi:10.1007/978-94-6091-596-3_2

[116] Janice M.Morse and Janice M. Morse. 1991. Qualitative Nursing Research: A Free-for-All? In *Qualitative Nursing Research: A Contemporary Dialogue*. SAGE Publications, Inc., 14–22. doi:10.4135/9781483349015

[117] Michael J Muller and Sandra Kogan. 2010. Grounded Theory Method in HCI and CSCW. In *Human Computer Interaction Handbook*.

[118] N. O. Natalina. 2023. Telegram channels as tools of strategic communication: a study on Ukraine's media landscape during the war. *Visnyk Donetskoho natsionalnoho universytetu imeni Vasylia Stusa. Seriia politychni nauky* (Oct. 2023), 53–60. doi:10.31558/2617-0248.2023.8.8

[119] Laura K. Nelson. 2020. Computational Grounded Theory: A Methodological Framework. *Sociological Methods & Research* 49, 1 (Feb. 2020), 3–42. doi:10.1177/0049124117729703

[120] Teresa K. O'Leary, Elizabeth Stowell, Jessica A. Hoffman, Michael Paasche-Orlow, Timothy Bickmore, and Andrea G. Parker. 2021. Examining the Intersections of Race, Religion & Community Technologies: A Photovoice Study. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–19. doi:10.1145/3411764.3445418

[121] Anton Oleinik. 2024. Telegram channels covering Russia's invasion of Ukraine: a comparative analysis of large multilingual corpora. *Journal of Computational Social Science* 7, 1 (April 2024), 361–384. doi:10.1007/s42001-023-00240-9

[122] Obinna Ogbonnia Otuu and Deepak Sahoo. 2025. How Should We Design Technology With Diverse Stakeholders Who Wish Not to Attend Design Activities Together?. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–14. doi:10.1145/3706598.3714168

[123] Antti Oulasvirta and Kasper Hornbæk. 2016. HCI Research as Problem-Solving. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, San Jose California USA, 4956–4967. doi:10.1145/2858036.2858283

[124] Cliodhna O'Connor and Helene Joffe. 2020. Intercoder Reliability in Qualitative Research: Debates and Practical Guidelines. *International Journal of Qualitative Methods* 19 (Jan. 2020), 1609406919899220. doi:10.1177/1609406919899220 Publisher: SAGE Publications Inc.

[125] Don Paarlberg. 1963. Methodology for What? *Journal of Farm Economics* 45, 5 (1963), 1386–1392. doi:10.2307/1236829

[126] Luc Pauwels. 2006. *Visual cultures of science: rethinking representational practices in knowledge building and science communication*. UPNE.

[127] Luc Pauwels. 2008. Visual Literacy and Visual Culture: Reflections on Developing More Varied and Explicit Visual Competencies. *The Open Communication Journal* 2, 1 (Sept. 2008), 79–85. doi:10.2174/1874916X00802010079

[128] Luc Pauwels. 2010. Visual Sociology Reframed: An Analytical Synthesis and Discussion of Visual Methods in Social and Cultural Research. *Sociological Methods & Research* 38, 4 (May 2010), 545–581. doi:10.1177/0049124110366233 Publisher: SAGE Publications Inc.

[129] Luc Pauwels. 2021. Contemplating 'visual studies' as an emerging transdisciplinary endeavour. *Visual Studies* 36, 3 (May 2021), 211–214. doi:10.1080/1472586X.2021.1970326

[130] Ellie Pavlick. 2022. Semantic structure in deep learning. *Annual Review of Linguistics* 8, 1 (2022), 447–471.

[131] Edgar Alwin Payne. 1957. *Composition of outdoor painting*. DeRu Fine Arts.

[132] Warren Pearce, Suay M. Özkula, Amanda K. Greene, Lauren Teeling, Jennifer S. Bansard, Janna Joceli Omena, and Elaine Teixeira Rabello. 2020. Visual cross-platform analysis: digital methods to research social media images. *Information, Communication & Society* 23, 2 (Jan. 2020), 161–180. doi:10.1080/1369118X.2018.1486871

[133] Ernesto Peña. 2025. The three waves of visual literacy. *Journal of Visual Literacy* 44, 1 (Jan. 2025), 1–17. doi:10.1080/1051144X.2025.2462379

[134] Whitney Phillips. 2018. The oxygen of amplification. *Data & Society* 22 (2018), 1–128.

[135] Mark Poster. 2002. Visual studies as media studies. *Journal of Visual Culture* 1, 1 (April 2002), 67–70. doi:10.1177/147041290200100106 Publisher: SAGE Publications.

[136] Sophia Ppali, Marios Constantinides, Fotis Liarokapis, Jaydon Farao, Soraya S. Anvari, MinYoung Yoo, Ferran Altarriba Bertran, Shannon Rodgers, Jihae Han, Rina R. Wehbe, Margot Brereton, and Alexandra Covaci. 2025. Cite Your Well-being First: What Happens When Personal Life, Mental Health, and HCI Research Become Entangled?. In *Companion Publication of the 2025 ACM Designing Interactive Systems Conference (DIS '25 Companion)*. Association for Computing Machinery, New York, NY, USA, 52–56. doi:10.1145/3715668.3734181

[137] Cassidy Pyle, Nicole B. Ellison, and Nazanin Andalibi. 2025. Designing for Discourse: Social Media, Socio-Technical Rhetorical Strategies, and Affirmative Action Discussions. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference (DIS '25)*. Association for Computing Machinery, New York, NY, USA, 1758–1773. doi:10.1145/3715336.3735698

[138] Li Qiwei, Allison McDonald, Oliver L Haimson, Sarita Schoenebeck, and Eric Gilbert. 2024. The sociotechnical stack: Opportunities for social computing research in non-consensual intimate media. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW2 (2024), 1–21.

[139] Hawra Rabaan and Lynn Dombrowski. 2023. Survivor-Centered Transformative Justice: An Approach to Designing Alongside Domestic Violence Stakeholders in US Muslim Communities. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–19. doi:10.1145/3544548.3580648

[140] Naghmeh Razaghi, Mahboobeh Abdolrahimi, and Mahvash Salsali. 2015. Memo and memoing in qualitative research: a narrative review. *Journal of Qualitative Research in Health Sciences* 4, 2 (Sept. 2015), 206–217. https://jqr1.kmu.ac.ir/article_90819.html

[141] Afsaneh Razi, John S. Seberger, Ashwaq Alsoubai, Nurun Naher, Munmun De Choudhury, and Pamela J. Wisniewski. 2024. Toward Trauma-Informed Research Practices with Youth in HCI: Caring for Participants and Research Assistants When Studying Sensitive Topics. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1 (April 2024), 134:1–134:31. doi:10.1145/3637411

[142] Pritika Reddy, Bibhya Sharma, and Kaylash Chaudhary. 2020. Digital literacy: A review of literature. *International Journal of Technoethics (IJT)* 11, 2 (2020), 65–94.

[143] Joe Skipper / Reuters. 2017. Handling Traumatic Imagery: Developing a Standard Operating Procedure. https://dartcenter.org/resources/handling-traumatic-imagery-developing-standard-operating-procedure

[144] Richard Rogers. 2021. Visual media analysis for Instagram and other online platforms. *Big Data & Society* 8, 1 (Jan. 2021), 20539517211022370. doi:10.1177/20539517211022370 Publisher: SAGE Publications Ltd.

[145] Gillian Rose. 2014. On the Relation between 'Visual Research Methods' and Contemporary Visual Culture. *The Sociological Review* 62, 1 (Feb. 2014), 24–46. doi:10.1111/1467-954X.12109 Publisher: SAGE Publications Ltd.

[146] Gemma Ryan. 2018. Introduction to positivism, interpretivism and critical theory. *Nurse researcher* 25, 4 (2018), 41–49.

[147] Joy Sather-Wagstaff. 2011. Beyond content: Thematic, discourse-centred qualitative methods for analysing visual data. In *An Introduction to Visual Research Methods in Tourism*. Routledge. Num Pages: 18.

[148] Joseph Scott Schafer, Brett A. Halperin, Sourojit Ghosh, and Julie Vera. 2024. TO SCREENSHOT OR NOT TO SCREENSHOT? TENSIONS IN REPRESENTING VISUAL SOCIAL MEDIA PLATFORM POSTS. *AoIR Selected Papers of Internet Research* (2024). https://spir.aoir.org/ojs/index.php/spir/article/view/14055

[149] Joseph S. Schafer, Kate Starbird, and Daniela K. Rosner. 2023. Participatory Design and Power in Misinformation, Disinformation, and Online Hate Research. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference (DIS '23)*. Association for Computing Machinery, New York, NY, USA, 1724–1739. doi:10.1145/3563657.3596119

[150] Hope Schroeder, Marianne Aubin Le Quéré, Casey Randazzo, David Mimno, and Sarita Schoenebeck. 2025. Large Language Models in Qualitative Research: Uses, Tensions, and Intentions. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–17. doi:10.1145/3706598.3713120

[151] Fred Schwaller. 2025. A mental-health crisis plagues PhDs—these evidence-led initiatives offer help. *Nature* 640 (2025), 579–580. doi:10.1038/d41586-025-01083-2

[152] Carol F Scott, Gabriela Marcu, Riana Elyse Anderson, Mark W Newman, and Sarita Schoenebeck. 2023. Trauma-Informed Social Media: Towards Solutions for

Reducing and Healing Online Harm. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg Germany, 1–20. doi:10.1145/3544548.3581512

[153] Limor Shifman. 2014. *Memes in digital culture.* The MIT Press, Cambridge, Massachusetts.

[154] Mong Yuan Sim, Wei Emma Zhang, Xiang Dai, and Biaoyan Fang. 2025. Can VLMs Actually See and Read? A Survey on Modality Collapse in Vision-Language Models. In *Findings of the Association for Computational Linguistics: ACL 2025*. 24452–24470.

[155] Flávio Souza, Diego de Las Casas, Vinícius Flores, SunBum Youn, Meeyoung Cha, Daniele Quercia, and Virgílio Almeida. 2015. Dawn of the selfie era: The whos, wheres, and hows of selfies on Instagram. In *Proceedings of the 2015 ACM on conference on online social networks*. 221–231.

[156] Garrett Souza, Nina Lutz, and Katlyn M Turner. 2025. Mediating The Marginal: A Quantitative Analysis of Curated LGBTQ+ Content on Instagram. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–20. doi:10.1145/3706598.3713618

[157] Barbara Maria Stafford. 1996. *Good looking: Essays on the virtue of images.* MIT Press.

[158] Kate Starbird, Emma Spiro, Isabelle Edwards, Kaitlyn Zhou, Jim Maddock, and Sindhuja Narasimhan. 2016. Could This Be True?: I Think So! Expressed Uncertainty in Online Rumoring. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, San Jose California USA, 360–371. doi:10.1145/2858036.2858551

[159] Terese Stenfors, Anu Kajamaa, and Deirdre Bennett. 2020. How to … assess the quality of qualitative research. *The Clinical Teacher* 17, 6 (2020), 596–599. doi:10.1111/tct.13242

[160] Anna Ståhl, Vasiliki Tsaknaki, and Madeline Balaam. 2021. Validity and Rigour in Soma Design-Sketching with the Soma. *ACM Transactions on Computer-Human Interaction* 28, 6 (Dec. 2021), 1–36. doi:10.1145/3470132

[161] Nicolás Salazar Sutil. 2015. *Motion and representation: The language of human movement.* MIT Press.

[162] Francesca Bolla Tripodi. 2022. *The propagandists' playbook: how conservative elites manipulate search and threaten democracy.* Yale University Press, New Haven.

[163] Emily Tseng, Thomas Ristenpart, and Nicola Dell. 2025. Mitigating Trauma in Qualitative Research Infrastructure: Roles for Machine Assistance and Trauma-Informed Design. *Proceedings of the ACM on Human-Computer Interaction* 9, 2 (May 2025), 1–37. doi:10.1145/3711035

[164] Emily Tseng, Mehrnaz Sabet, Rosanna Bellini, Harkiran Kaur Sodhi, Thomas Ristenpart, and Nicola Dell. 2022. Care Infrastructures for Digital Security in Intimate Partner Violence. In *CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans LA USA, 1–20. doi:10.1145/3491102.3502038

[165] Jennifer Turns and Judith Ramey. 2006. Active and collaborative learning in the practice of research: Credit-based directed research groups. *Technical communication* 53, 3 (2006), 296–307. Publisher: Society for Technical Communication.

[166] Kari Dee Vallury, Barbara Baird, Emma Miller, and Paul Ward. 2021. Going Viral: Researching Safely on Social Media. *Journal of Medical Internet Research* 23, 12 (Dec. 2021), e29737. doi:10.2196/29737

[167] Paul M. Van Auken, Svein J. Frisvoll, and Susan I. Stewart. 2010. Visualising community: using participant-driven photo-elicitation for research and application. *Local Environment* 15, 4 (April 2010), 373–388. doi:10.1080/13549831003677670

[168] Niels Van Berkel and Kasper Hornbæk. 2023. Implications of Human-Computer Interaction Research. *Interactions* 30, 4 (July 2023), 50–55. doi:10.1145/3600103

[169] Antoon Van den Braembussche. 2009. *Thinking art: An introduction to philosophy of art.* Springer.

[170] Theo Van Leeuwen. 2001. Semiotics and iconography. *SAGE Visual Methods* (2001), 83.

[171] Jessica Vitak, Nicholas Proferes, Katie Shilton, and Zahra Ashktorab. 2017. Ethics regulation in social computing research: Examining the role of institutional review boards. *Journal of Empirical Research on Human Research Ethics* 12, 5 (2017), 372–382.

[172] Zezhong Wang, Samuel Huron, Miriam Sturdee, and Sheelagh Carpendale. 2024. Summary of the Workshop on Visual Methods and Analyzing Visual Data in Human Computer Interaction. In *Companion Proceedings of the 2024 Conference on Interactive Surfaces and Spaces*. 29–32.

[173] Amanda Wasielewski. 2023. *Computational formalism: Art history and machine learning.* MIT Press.

[174] Judith Davidson Wasser and Liora Bresler. 1996. Working in the Interpretive Zone: Conceptualizing Collaboration in Qualitative Research Teams. *Educational Researcher* 25, 5 (June 1996), 5–15. doi:10.3102/0013189X025005005 Publisher: American Educational Research Association.

[175] Felix A. Wichmann, Lindsay T. Sharpe, and Karl R. Gegenfurtner. 2002. The contributions of color to recognition memory for natural scenes. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 28, 3 (2002), 509–520. doi:10.1037/0278-7393.28.3.509 Place: US Publisher: American Psychological Association.

[176] Tamar Wilner, Kayo Mimizuka, Ayesha Bhimdiwala, Jason C Young, and Ahmer Arif. 2023. It's about time: Attending to temporality in misinformation interventions. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. ACM, Hamburg, Germany, 1–19.

[177] Jacob O. Wobbrock and Julie A. Kientz. 2016. Research contributions in human-computer interaction. *interactions* 23, 3 (April 2016), 38–44. doi:10.1145/2907069

[178] Srishti Yadav, Lauren Tilton, Maria Antoniak, Taylor Arnold, Jiaang Li, Sidhesh Milind Pawar, Antonia Karamolegkou, Stella Frank, Zhaochong An, Negar Rostamzadeh, Daniel Hershcovich, Serge Belongie, and Ekaterina Shutova. 2025. Evaluation of Cultural Competence of Vision-Language Models. doi:10.48550/arXiv.2505.22793 arXiv:2505.22793 [cs].

[179] SungBaek Yoon, Seho Bae, Hanjae Park, Liu Hua, Zheng Ziqi, and Juneho Yi. 2016. Efficiently Exploiting the Context of Object-Action Context for Object Recognition. In *Proceedings of the 10th International Conference on Ubiquitous Information Management and Communication* (Danang, Viet Nam) *(IMCOM '16)*. Association for Computing Machinery, New York, NY, USA, Article 84, 4 pages. doi:10.1145/2857546.2857632

[180] Jiake Zhang. 2024. A Data Analysis Framework for a Topic on Social Media. In *Proceedings of the 2024 3rd International Conference on Artificial Intelligence and Intelligent Information Processing*. ACM, Tianjin China, 111–116. doi:10.1145/3707292.3707351

[181] Ting Zhong, Jian Lang, Yifan Zhang, Zhangtao Cheng, Kunpeng Zhang, and Fan Zhou. 2024. Predicting Micro-video Popularity via Multi-modal Retrieval Augmentation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*. Association for Computing Machinery, New York, NY, USA, 2579–2583. doi:10.1145/3626772.3657929

[182] Kaitlyn Zhou, Tom Wilson, Kate Starbird, and Emma S. Spiro. 2023. Spotlight Tweets: A Lens for Exploring Attention Dynamics within Online Sensemaking During Crisis Events. *ACM Transactions on Social Computing* 6, 1-2 (June 2023), 1–33. doi:10.1145/3577213

A Memos

We provide selected memos from 2 projects which featured first time coders (Case Study 1 and 3). This is because these projects were particularly emotional and involved coders (particularly undergraduates) very close to the often distressing content. Memo excerpts have been edited for clarity and anonymization purposes, as well as for language. In some cases, we describe the media as students often memoed referencing media by file name.

A.1 Border memos

Memo 1: I've always said it's better to know than not, and after this experience, I feel less ignorant about a lot of topics. I'm not sure if that's a good or bad thing...When I say "I feel less ignorant," I mean it on purpose — because there were times I felt like I was missing so much context. I was honestly amazed by how much the other members of the team knew! It was awesome hearing everyone's insights.

Memo 2: This [image] also sparked a lot of reflection for me because I'm Mexican, and thinking about what the next presidency could mean is really personal.

Memo 3: People online will make light of pretty much every situation they can using different forms of media like memes and textposts.

Memo 4: I was really sad that my first reaction [to a video of a woman using cardboard to get her daughter through barbed wire] was that I thought it was staged.

Memo 5: …there was definitely a shift in the ways memes were presented [this week]. Initially, there was a lot of content towards victory celebrations for Republicans on Trump's win. This was a challenging time for me to code as I was just so sad about this news.

Memo 6: I found myself making connections between the images we were coding…and the images I was encountering on social media…which was a striking (and depressing) realization…I realized that what we were doing…is quite similar to what moderators or data annotators for machine learning systems go through – we were all exposed to this increasingly dehumanizing content and were getting desensitized to them.

Memo 7: People have [always] called Mexican men gang bangers and rapists and b**n*rs and w*tb*cks and whatever. I got called a sp*c and b**n*r as recently as last week. But it was still harder than I thought it would be to see photos of people that looked like me and my little brother and dad portrayed as rapists and gangsters. My dad crossed the border so we could have a better life and he's worked so hard to do that. I wish people would see our community as hard working and not freeloaders like most of these pictures make us out to be. I was especially mad to see so many Mexicans posting media mocking Venezuelans for being bad Latinos, because white people don't know the difference.

A.2 Ukraine memos

Memo 1: The image (*see Figure 15*) that affected me most was this one *(reproduced for paper, student linked it)*. Something about seeing the splintering and burned rooms up close.

Figure 15: Image of a destroyed building from the Ukraine project

Memo 2: It feels that Russian occupiers of Ukrainian land (DNR, Mariupol) act as if they belong there and have always belonged there. This is both propaganda for Russians and greatly discouraging for Ukrainians.

Memo 3: It was less heavy on me than I thought it would be, perhaps because I am not typically deeply affected by still images. I thought the images as a whole would be more explicitly violent, but surprisingly, I was most struck by the infographics from DNR and Mariupol that were advising people about medical care etc. They felt insidious in a way that outright violence does not.

Memo 4: I was really struck by the flyer for the pro-Ukraine march being run by extremist groups. It made me internalize that even in a cause that I deeply support (independence and peace for Ukraine) there will be unsavory and dangerous viewpoints

Memo 5: it's hard because without context there's a very small chance that those ornaments were genuinely made by someone with those beliefs but the more likely conclusion is that they are Russian ragebait

Memo 6: It was surprising to me that I felt more emotional about the map of frontline moved back than the photo of a dead body. I think consuming so much about the war has made me a little numb sometimes but also seeing the line move makes me sad that maybe we can't win.