

Chapter 5. Systematic Invention and Incremental Change

The two preceding interviews featured professionals whose careers shaped AI. In the previous chapter, Kate described how she uses generative AI in an average workday among a team of content designers and writers, highlighting Writer’s affordances as a generative tool while also describing its limitations as a replacement for human-authored content. In Chapter 6, Terry describes his experiences working in a high-technology environment of a different sort: advanced manufacturing. While Kate and Terry work in very different domains—digital content design and medical manufacturing—they describe surprisingly consistent themes when it comes to how generative and traditional AI technologies are integrated into professional writing and production workflows. Neither interviewee paints a picture of AI arriving suddenly to overhaul systems. Instead, they show how AI is gradually incorporated into legacy infrastructures, augmenting the capabilities of already lean teams (see Johnson et al., 2018).

This interlude chapter explores parallels between Kate’s and Terry’s interviews to thematically highlight what AI looks like in everyday professional writing contexts, and to contextualize the nature of Terry’s workplace. Contrasting the previous interviews’ focus on digital texts, we take time here to trace histories of tangible technologies. Specifically, we historicize bicycles and cameras, to ground the metaphors we use to understand generative AI, and to situate readers in contexts of manufacturing and the automated production of physical products. Across all three interviews, AI is not a sweeping disruption but a set of targeted interventions: tools that streamline existing processes, automate the tedious, and support quality and consistency across complex systems. Rather than reinventing the workplace, AI here is a companion to maintenance and small-scale optimization. For Kate, generative AI entered the workflow through tools like Writer, where it helps content teams brainstorm ideas, generate draft subject lines or titles, while Ditto helped enforce consistency across hundreds of pieces of documentation. In rhetorical terms, generative AI supported invention—the generation of possibilities that can then be refined by human writers. The ability to save even 15 to 30 minutes per ideation task adds up when multiplied across a content team managing messaging for a global company. However, as Kate is quick to point out, AI is less effective at composing long form content without significant intervention. The team still relies on human writers as content designers to effectively draft and revise final deliverables. But again, AI-based tools saved time and tedium by identifying and suggesting edits across hundreds of documents, potentially composed by dozens of human authors over time. It would be tedious for human authors to review hundreds of blog posts in relation to style guides to ensure consistent usage. As Kate described, advanced find

and replace tools with appropriate pattern matching could accomplish similar results, but would require technical expertise and extensive testing of expressions in order to catch any special cases. Instead, traditional AI as well as specialized LLM systems, developed for particular tasks, make converting proverbial needles in a haystack into usable hay a routine task for small teams.

Similar themes were apparent in Terry's interview. Kate emphasized that generative AI augments her team's capacity—it does not replace writers. Her team remains responsible for the quality and coherence of the company's messaging. Similarly, Terry describes a production and analysis environment that is already highly automated by traditional robotics and software systems. That is to say, in an automated manufacturing environment, human workers were in many cases already replaced with robots at scale decades before generative AI was a consideration. Even simple robots are more than capable of performing repeated tasks precisely and accurately, and can perform complex procedures with relative ease. Terry manages the customization of millions of medical testing kits, with machines assembling the most common configurations (the top 20 percent of test kits), and human workers collaborating with machines to assemble most custom kits. AI, in this context, holds potential to help with the computational challenges of tailoring solutions to unique client needs and improving the efficiency of lean teams—but it has not yet been deployed in these roles at Labcorp, Terry's employer. Still, Terry sees the value of AI for streamlining decision-making across thousands of product permutations and offers valuable insight into AI's potential. Efficiency, in this light, is not a matter of shrinking headcount—it's about helping small teams manage increasingly large and complex information systems. Throughout Kate and Terry's examples, the value of AI lies in its ability to support human expertise, not supplant it. Perhaps even more importantly, Terry describes challenges associated with improving the design of large-scale automated systems. Terry's experience grounds speculation and theory-building about AI's potential impact on manufacturing and related fields in current realities of the manufacturing industry. Bridget, Kate, and Terry all work with systems that operate at a large regional or even global scale, but Terry's experience foregrounds the complications involved with managing physical rather than virtual automation.

Both Kate and Terry describe the difficulty of updating legacy systems. For Kate, the problem is often one of maintaining the usability and relevance of old content—updating blog posts that still rank well in search, or configuring new platforms to maintain consistent brand voice across teams. For Terry, the issues are deeper and more structural. Labcorp's systems span decades of internal customization and vendor-specific software, as well as production and testing facilities around the globe (with their respective employees), all of which make transitions to new tools slow and complicated to implement. Both cases suggest that AI integration is not plug-and-play—it requires thoughtful attention to system compatibility, documentation, and human knowledge transfer. Terry

describes a moment in which Labcorp was forced to make a difficult decision regarding how and where to open a new facility—develop a new production site in Belgium for greater output and potential revenue, with the associated risk of developing a new system across the Atlantic; or open a smaller facility in the US using a tried and tested system, while perpetuating the limitations and challenges of Labcorp’s existing infrastructure? Such problems have not arisen as a result of AI, nor are AI-based systems likely to solve these problems. This context grounds conversations about AI in realities of automated workplaces.

Despite the differences between Kate and Terry in domain and the level of AI integration at their respective workplaces, both professionals manage teams responsible for producing tailored outputs at scale. Whether it’s written content for global audiences or customizable test kits for medical use, Kate and Terry must coordinate internal systems, external expectations, and logistical realities. Their teams are small relative to the scope of work they manage, and both see AI as a way to help them keep quality high and results consistent at scale.

Before moving into Terry’s interview, we return here briefly to differences between traditional AI and generative AI. The distinction is important to keep Kate’s description of Writer to brainstorm ideas distinct from the pattern recognition of traditional AI incorporated in Ditto, and the potential for computational optimization (also traditional AI) Terry describes. One way to remember such distinctions is through the metaphor of invention. As rhetorical scholars have long noted, invention refers to the identification of possible means of persuasion or problem-solving in each context. Generative AI, true to its name, supports invention by surfacing possibilities. This capacity has led to critiques of generative AI as a “bullshit generator” and stochastic parrot (Bender et al., 2021; Gorrieri, 2024). Although these are apt descriptors for generative LLMs we trace our initial encounters with generative AI through two different technology histories: CAD modeling and photography. We reflect on our initial encounters with generative AI—as part of the CAD modeling invention process for industrial fabrication—and how modern photography workflows might offer a glimpse into future writing processes and AI-driven interfaces for media production. We then present metaphors rooted in analog machines—like bicycles and gym machines—to help us critique popular assumptions about automation. Before transitioning to our final interview chapter, we take a necessary historical deep dive into photography’s shift from analog to digital, and racism entangled with automated development processes during that transition.

Precursor Disruptive Technologies: CAD Modeling & Photography

Our conversations about generative AI and its potential impact on technical and professional communication emerged from encounters with Autodesk’s *Project Dreamcatcher*, an experimental generative design platform launched in

the mid-2010s. Project Dreamcatcher could generate thousands of potential 3D CAD models for a given design problem—far more than a human designer might develop manually. It then narrowed those possibilities to a manageable set based on designer selections, and designers could then further develop or combine elements of the generated designs. The goal was not to replace human designers but to extend their capacity for invention using cloud computing. Generative AI, as its name suggests, is good at generating multiple possibilities quickly and computing the available means of addressing mathematically constrained problems.

As Project Dreamcatcher showed (and as we described in Sherrill & Salvo, 2022), sometimes the results exceed the imaginations of designers, and at other times results technically meet defined parameters but are immediately rejected. Understanding generative AI's affordances and limitations provides context for Kate's use of Writer. Writer could generate lists of potential titles or subject lines efficiently, while human authors rejected many of the generated outputs, just as with brainstorming processes. Expectedly, Writer was less useful for generating longer form content without extensive prompt engineering. In most cases, it was faster and easier for a human to write longer texts, reflecting our own experiences writing with generative AI. This should come as no surprise given the purpose of generative design systems launched prior to ChatGPT. Having said that, Project Dreamcatcher was marketed with a second automated design step in mind: topology optimization.

Topology optimization is a mathematical process of incrementally refining designs to reduce material usage while preserving strength. A biological analog for the process might resemble the evolution of bird bones to be optimized for flight with the necessary tradeoffs (i.e., design constraints) of affording survival across varied conditions. The result: bird bones are largely hollow, and therefore lighter, while maintaining sufficient strength, yet still recognizable as bones. Applying this approach to industrial manufacturing, a designer might not notice the impact of shaving a few grams from an individual part, but when repeated across hundreds of components in a car or aircraft, those marginal gains add up. Similarly, optimizing assembly time by tenths of a second per medical kit might seem negligible, but scaled across millions of units, the time and cost savings become substantial.

This combination of invention and refinement reflects the hybrid nature of AI integration in both Kate's and Terry's work. In content design, AI might generate lists of title ideas or flag inconsistencies, saving writers from tedious tasks. But refinement—editing, aligning with tone, meeting user needs—largely remains the work of people. In manufacturing, future applications of generative AI may identify promising configurations or optimize decision trees, but the systems themselves are still anchored in human expertise and institutional knowledge.

Traditional AI plays a role as well. For Kate, the proverbial needles in the haystack were phrases that didn't align with company style guides and therefore needed correction. For Terry, one problem involves identifying combinations of needles that are likely to create a haystack on the assembly line: particular combinations

of kit contents can cause substantial backups and delays. To address this problem, Terry recognizes that advanced computing can analyze millions of potential parts combinations that might delay the assembly of custom kits. Although generative AI could play a role in creating plausible combinations of parts to virtually test, this is likely a task for traditional AI rather than generative AI.

Both Kate and Terry illustrate a future of work where AI participates in collaborative systems without dominating them—again representing Knowles’ ideal of machine-in-the-loop systems (2024). Whether it’s sorting through thousands of blog posts or managing the combinatorial haystack of product variations, AI helps make overwhelming tasks tractable. But it also requires labor—technical labor, rhetorical labor, and managerial labor—to function responsibly and effectively.

Taken together, Kate and Terry offer a grounded perspective on what AI means for professional writing in the workplace. Not a wholesale transformation, but a set of careful, iterative adaptations. Not a substitution of human intelligence, but an augmentation of human coordination. And not a utopian vision of frictionless automation, but a realistic acknowledgment of the time, expertise, and care required to make AI useful across complex, human-centered systems.

Photography Disrupts Itself

Throughout the book, we refer to a metaphor of photography and draw parallels with writing in the age of generative AI. In doing so, we extend an argument that began as a presentation for the 2023 Conference on College Composition and Communication about the potential for LLMs and other forms of generative AI to shift from command line interfaces towards a GUI more similar to other forms of digital media production (Sherrill, 2023), as well as student encounters with generative AI in an advanced writing course (Salvo, 2023). As anticipated, interfaces for prompting generative AI systems have already begun this shift. ChatGPT’s latest update, Canvas, resembles a WYSIWYG approach with improved usability and targeted editing of generative output via a pop-up menu with shortcuts (OpenAI, 2024). Compare in Figures 5.1-5.3 the interface for ChatGPT Canvas, the interface for Capture One (a professional photo editor), and the slider-based GUI that John envisioned for ChatGPT in his 2023 presentation.

As such, we conceptualize writing (like photography) as a workflow of rhetorical decision making by human designers using semi-automated tools to produce modular drafts that can be extensively and algorithmically reworked. In doing so, we are also implicitly situating written text alongside other forms of generative AI output such as audio, video, images, 3D models, etc. That is, although writing allows for knowledge creation and metacognitive reflection in ways that are often distinct from other forms of media and communication, writing is a form of digital media (drawing from Lev Manovich’s definition in *The Language of New Media* (2002) that can be algorithmically generated, and we treat it as such throughout the Conclusion.




Figure 5.1. A sample image from ChatGPT's Canvas interface, with notes on Design as the differentiator.




Figure 5.2. A picture of tomatoes in Capture One photo processing software used as color saturated example to illustrate the ways automation and AI are used in routine image formatting. Interface for the exposure adjustment panel is shown.




Figure 5.3: Mock-up of a hypothetical slider interface for writing with generative AI, modeled after the conventions of image processing software.

We, as teachers of technical communication, composition, and other disciplines, have encountered the emergence of similarly transformative/disruptive tools before. WYSIWYG (what you see is what you get) interfaces emerged for web design in the early and mid-2000s, and steadily improved with time. However, as the code generated by such interfaces also improved, they made it harder to identify whether students had generated their own HTML and CSS code, or had relied on the automated tools to generate reasonably usable output. Because source code for websites is readily viewable in any browser, it was similarly challenging to detect plagiarism: despite teaching ethical best practices, students could take inspiration from source code, or copy-paste with minimal editing to disguise the original source. Institutional systems for addressing plagiarism, still rooted in analog writing pedagogies, often lagged behind newer tools and failed to address teachers' concerns. Consequently, when teaching web design (whether via hand coding or even via WYSIWYG content management systems like Wix and Squarespace), documenting students' rhetorical decision making was foundational. Audience analyses, design plans, paired coding, distribution plans, revision plans, peer feedback, usability evaluations, and designing personally meaningful projects all played important roles in enculturating accountability as part of the learning process and are integral to producing professional content.

"Explain how it works and why you did it that way," live, in front of an audience, remains an essential pedagogical tool. At the time Wix and other tools became publicly accessible, diehard DIY-OR-DIE proponents of hand coding argued that automated tools such as Wix failed to teach students important rhetorical skills and foundational web design concepts rooted in code, while many in the fields of composition and technical communication championed the accessibility of WYSIWYG tools for first-year students and the ease with which students could engage authentic audiences outside the classroom via the web. Over time, WYSIWYG tools improved, and new tools emerged—such as Markdown, offering an experience closer to hand coding with broader accessibility. Michael and I

see value in all of these approaches, having lived and taught through that particular period of development and ensuing pedagogical discourse. To be clear, we do not conceptualize writing as being the same as all other forms of media production, rather, that understanding writing as inherently distinct can be limiting and at times problematic. As one example, compare my (John's) typical photography workflow with my typical writing workflow:

Photography workflow: take some notes, capture a few hundred images in which I make decisions about composition, subject, focus, etc., and allow the camera to automatically adjust certain settings within defined constraints. Import the RAW¹⁰ files into Capture One for processing and export.

At times, I (John) might do some additional retouching or compositing in Photoshop, depending on the situation. I'm not doing that photographic editing work via command line, I'm using sliders and dials. The base image is already made, but I can make significant adjustments. To get to that base image, I had to make choices as a photographer about subject, composition, exposure, and other camera settings. Sometimes, it's important to manually dial in each setting, be it for consistency, to get a specific exposure, or to create a specific visual effect (e.g., shutter drag or intentional camera movement). At other times, that mental labor isn't as critical, and the camera does just fine deciding for me how to expose the frame or adjust the shutter speed to freeze motion. But even when the camera automates parts of the process, I still have options for making radically different final images from the same base file.

Writing workflow: take some notes on paper or plaintext, either A) type a rough outline or some extended "chunks" of main points in plaintext that can be further developed or B) speak and record myself and then generate an initial transcript (still plaintext) or C) ask ChatGPT to extend or connect chunks of text. Import the plaintext version into Microsoft Word or Google Docs for revising. Extensively rework the initial draft, making choices about arrangement, focus, etc. while also making use of automated tools like spellcheck, Grammarly, or ChatGPT.

Depending on the rhetorical situation, I might do basic formatting in Word or Google Docs and call it a day. But in a professional setting I might import the text into InDesign for page layout, generate an HTML version or import the text into a content management system (CMS) like WordPress or a system for structured authoring (where the text could be further transformed), or another specialty program. With each layer, I move further from the unformatted plaintext draft which offers minimal automation between my choices as a writer and the text I produce, to increasingly automated tools. Sometimes, even sitting at a computer feels like too many layers of automated mediation to start with, so I write by hand or dig out a manual typewriter from my closet to hammer out lines of text. Many times, the

10. RAW files provide the most data from the camera's sensor, uncompressed and unprocessed. RAW files are loosely equivalent to undeveloped film and allow for a wider range of image adjustments than compressed formats like JPEG.

mental labor of spelling, editing, and formatting are less critical for me to make choices about, and spellcheck or ChatGPT handles that level of automation just fine. But even when those tools make automated decisions, I still have options to transform the text to make radically different texts from the same base file.

For many photographers, the appeal of shooting film in a digital age is that it slows down the photographic process. It requires the photographer to make conscious choices, which are also often tactile by nature of being analog—turning the aperture ring, adjusting the shutter speed dial, manually bringing the frame into focus through the viewfinder by turning the focus ring, hearing the distinct “ker chunk” of the shutter release mechanism, and advancing the film. Though fully mechanical manual¹¹ cameras are coveted to the point that some sell for thousands of dollars (more than some modern digital cameras), on average, many film photographers still shoot film cameras from the 1970s or 1980s that offer the choice of using some automated assistance, e.g., an indicator light to suggest exposure rather than relying on years of experience or a separate light meter (or equivalent app). Similarly, though some DIY-or-die photographers develop their own film, many still prefer the relative convenience of mailing their film to a lab and receiving digital scans as well as physical prints. There are still photographers with the necessary technical expertise and resources to work with photographic plates, cyanotypes, and other historic analog processes, but they are comparatively rare. For most modern photographers, such techniques might have been covered in a classroom lesson in much the same way that writing students might be exposed to a dip pen or a typewriter, but on average would not make their own iron gall ink. However, modern digital photography classes do assign work that requires students to manually adjust settings in order to understand the exposure triangle and how various camera settings impact the resulting images—the automated camera still affords manual control when needed—an important point which we will return to in our discussion of ethical considerations, as photographers, much like writers, must make careful rhetorical decisions about how they represent reality and communicate with audiences, and must consider the impact of their choices alongside automated decision-making. To have greater rhetorical agency, beyond framing and filters or “good enough” images, photography students must still learn the foundations of how each camera setting impacts the captured image. To move beyond shooting 100 images and getting lucky with one or two better captures, photographers must develop technorhetorical skills. Similarly, generative AI generates rhetorically adequate texts under average circumstances, and occasionally generates something that exceeds expectations. Synthesizing the best parts of generative output can be productive, but technorhetorical expertise

11. It is worth remembering here that mechanical “manual” cameras represent over a century of technological advances in automation. Compared with other methods (e.g., portraits, sketches), even daguerreotype images were capable of automating the process of creating an accurate depiction of reality.

allows users to make the most of AI-driven tools—as we have seen in our interviews with Bridget and Kate.

Bicycles and Physical Impossibilities

As we have encountered in our own teaching, as well as published research, students on average recognize the importance of using writing skills and automated tools in ethical, expert ways. They understand that the ability to do so requires thoughtful decision-making about how to incorporate automated tools effectively, and that developing such an awareness—*praxis*—comes with practical experience informed by theory. Andelyn Bedington et al., in “Writing with generative AI and human-machine teaming,” provide a relevant example, sharing undergraduate students’ reflections on their interactions with AI in a writing course.

I was initially worried that AI would be like a car, removing all exercise for me as a writer and thus not letting me develop and exert myself. But now I see AI as a bicycle. Just as a bicycle can take me farther than if I walked but it still is my exercise that is shaping the experience, so now do I see how AI can improve my writing without replacing the mental exercise of the writing process. Writing with AI still requires exercising knowledge and skills, just as a bicycle does, and provides endless room for improving on those skills and developing that knowledge. (Bedington et al., 2024, p. 8)

This bicycle metaphor is useful for understanding AI in relation to writing.

Much like machines in a gym guide novice users to understand how proper form feels with less risk of injury due to improper form, they allow users to go further faster than manual exercise alone. Walking into a small gym that has only a multipurpose exercise machine, free weights, and a bench can be as intimidating as staring down a blank page. Even dedicated machines can be intimidating to use for the first time in front of spectators. But compared with dumbbells, most machines include at least basic illustrated instructions for proper use. With minimal instruction, exercise machines help users safely develop foundational literacies, building muscle memory and recognizing proper form through repetition while avoiding major injury. They are machines that quite literally support learning through labor, but still require human input. The machines provide a physical template for proper form. Users not only see illustrated instructions or a model of form—similar to reviewing examples of a particular written genre—rather, the machines constrain how the user moves, while requiring considerable exertion. Over time, the amount of exertion required can also be adjusted as the user gains muscle. As users gain confidence and skill, they may find the machines limiting, or may grow tired of waiting for others to finish their sets on a limited

number of machines. Instead, gym-goers might use more abundant free weights, informed by a quick YouTube tutorial on their phone and equipped with practiced knowledge of which muscles will feel each exercise. But even users with years of experience in weightlifting and exercise may still prefer the mechanical support of a leg press to manual squats, or the consistent and controllable pace of a treadmill, as much for safety as efficiency. However, as many viral videos show, even with the simplest of gym machines and accompanying illustrated instructions, people sometimes use the constrained machines in creative, novel, and unintended ways. Like exercise machines, generative AI tools afford structured, guided experiences that can help writers build foundational skills and confidence—especially when they’re just starting out or facing a daunting task. These tools constrain and support, allowing users to experiment with form and technique, or offsetting the mental labor of editing sentences and facing down a blank page. But just as with gym equipment, users can misuse or over-rely on AI, which raises concerns about authenticity, agency, and rhetorical awareness.

Wiebe Bijker’s work on the *Social Construction of Technology* (SCOT) approach, particularly in *Of Bicycles, Bakelites, and Bulbs* (1999), argues that technology evolves through social negotiation rather than inevitable progress. His reflections on the bicycle reinforce the idea that users actively shape technological development, challenging the notion that innovation is driven solely by engineers or market forces. The bicycle, in his analysis, was not simply invented and adopted; it was contested and redefined by different social groups, including safety-conscious users, women advocating for mobility, and manufacturers responding to shifting demands.

Bijker’s insights anticipate later discussions about user-driven innovation, participatory design, and democratized technology. His work highlights how everyday people influence technological change, a perspective that resonates with modern developments like open-source software, personalized fabrication, and digital hacking cultures. He identifies the active role users play in modifying and repurposing technology to fit their needs, rather than passively consuming pre-designed products.

The democratization of photography through Kodak offers another example of user-driven technological transformation (which we revisit in the section No Neutral Grey). When George Eastman introduced the Kodak camera in 1888 with the slogan “*You press the button, we do the rest,*” he removed significant technical barriers to photography, shifting image-making from a professional craft to a mass cultural practice (Eastman Kodak Company, n.d.). Much like the bicycle’s transition from an elite innovation to a universal mode of transportation, Kodak’s accessible cameras redefined who could participate in photography. This shift allowed ordinary people to document their lives, shaping visual culture in ways previously restricted to professionals.

Both the bicycle and Kodak photography exemplify how American consumer culture eagerly adopts technologies that enhance individual agency. These

innovations were embraced not just because they were practical but because they empowered users, granting mobility, self-expression, and new forms of participation. Bijker's framework for understanding technological change—emphasizing user agency, social shaping, and interpretive flexibility—remains relevant today. From the rise of smartphones and social media to decentralized finance and AI tools, technology continues to evolve through a process of negotiation between designers, markets, and users. His work reminds us that technology is not just something we inherit but something we collectively shape. This ongoing negotiation between users and technologies is especially visible in education, where questions of agency, expertise, and effort are central. Just as past innovations like the bicycle or the camera reshaped cultural expectations around mobility and expression, generative AI challenges how we understand learning, authorship, and cognitive labor.

Critics of generative AI have referred to a gym metaphor as well. Some writing instructors argue that using ChatGPT in the classroom is equivalent to having a robot substitute go to the gym. Encountering arhetorical use of automated tools can certainly feel this way, as we acknowledge in our dialog about our own teaching experiences. In one extreme example, Victoria Livingstone wrote a *TIME* article titled, “I Quit Teaching Because of ChatGPT” (2024). In her article, Livingstone quotes Ted Chiang's metaphor, “Using ChatGPT to complete assignments is like bringing a forklift into the weight room; you will never improve your cognitive fitness that way” (Livingstone, 2024; Chiang, 2024). Livingstone argues, much as we do, that the challenge of implementing generative AI into classrooms is that novice writers may not be able to distinguish between writing produced by experts and writing that appears expert—much the same way that a novice gym-goer may not be able to distinguish between proper form and someone sweating profusely and performatively grunting with each lift of large weights (or WYSIWYG-generated code vs. hand coded text). However, we diverge from Livingstone and Chiang in how we conceptualize the load sharing between gym-goers, exercise machines, and generative AI. In his *The New Yorker* article, titled “Why A.I. Isn't Going to Make Art,” Chiang summarizes his argument as follows:

The companies promoting generative-A.I. programs claim that they will unleash creativity. In essence, they are saying that art can be all inspiration and no perspiration—but these things cannot be easily separated. I'm not saying that art has to involve tedium. What I'm saying is that art requires making choices at every scale; the countless small-scale choices made during implementation are just as important to the final product as the few large-scale choices made during the conception. It is a mistake to equate “large-scale” with “important” when it comes to the choices made when creating art; the interrelationship

between the large scale and the small scale is where the artistry lies. (Chiang, 2024)

As we will discuss in more depth later in this chapter, such marketing strategies are centuries old. Kodak made similar promises with its “You press the button, we do the rest” slogan (Eastman Kodak Company, n.d.). Where we diverge from Chiang is in our understanding of automated choices (and the fact that we are concerned with writing for the workplace rather than producing art). We argue that automating rhetorical labor, identifying and choosing the available means of persuasion, is not inherently antithetical to artistry or literacy learning, but should be used responsibly—which requires technorhetorical literacy. Challenging to teach, undoubtedly, but historically not an unfamiliar problem for the fields of composition and technical communication. It is fair to argue that few, if any, award-winning photographers have become accomplished artists through letting the camera automate every decision other than where to point the lens and when to trigger the shutter. But particularly in the context of the workplace, automatic decision-making is foundational. Few would argue that award-winning photographers never trust the camera to make automated decisions strategically, much as Chiang doesn’t believe “that art has to involve tedium” (2024).

To draw upon another metaphor, consider here the appeal of mass-produced chicken stock. For an annual holiday meal, a home cook might make a delicious chicken stock from scratch, carefully selecting and preparing every ingredient and spending hours tediously simmering the broth while tasting and adjusting at each step. The resulting stock might make deliciously artistic gravy, mashed potatoes, and stuffing (if one considers stuffing delicious, that is)—elevating these dishes for a special occasion. But on an average Thursday night dinner, readily available store-bought stock balances flavor, time, and effort to create a meal that is satisfactory and often still much better than faster fully “automatic” microwaveable options.

Generative AI is not producing Michelin Star quality stock, and likely never will, even as it improves. At its current best, generative AI’s output resembles Campbell’s, and occasionally produces output equivalent to an organic “private label” premium quality canned soup or stock, but can metaphorically create varieties not profitable for mass production. By the same rationale, businesses hire photographers to produce art when needed, and sometimes to produce stock photos that otherwise don’t exist. But when stock photographs or bulk-processed headshots *satisfice*¹² (Simon, 1956, 1997), “stock quality” holds considerable value—Bridget’s interview is testament to this. Award winning chefs will continue to produce deliciously nuanced stocks from scratch and charge a premium price

12. A combination of “satisfy” and “suffice” to describe decision making, as developed by Herbert Simon in *Administrative Behavior* (1997), but coined in a later paper, <https://psycnet.apa.org/doiLanding?doi=10.1037%2Fh0042769>

for their expert labor, perhaps even exploring new ideas from the unexpected output of generative systems. But that level of decision-making is exhausting, and unsustainable for many people. As teachers, part of our job in the age of generative AI is to highlight the potential impact of automatic and default decisions, much as it has been throughout time, rather than to condemn new tools as inferior to human agents. *Of course* an Olympic sprinter can outrun a novice on a bicycle, but that is not our point. Not every student aspires to be an Olympic gold medalist, nor should that be our sole pedagogical aim; many students just want to reach the finish line ~~without collapsing~~ in a reasonable manner.

These metaphors we have used here, photography, WYSIWYG, bicycling, gym machines, align with Knowles' concept of *Rhetorical Load Sharing* (Knowles, 2024). Each of these metaphors contains a spectrum of automation and human agency, as do additional metaphors described by Anuj Gupta et al. (2024) and Luke Stark's metaphor of "ChatGPT is Mickey Mouse" (2023). For Knowles, ideally, shared labor between humans and automated systems is structured such that humans maintain their rhetorical agency within the assemblage, while strategically using automated tools to do the heavy lifting (whether metaphorical or actual). In the examples described throughout our interviews, the automated systems assist human workers, with people doing the majority of the rhetorical labor. Generative AI extended the work of Bridget and Kate's teams beyond what they could accomplish on their own, but did not replace their rhetorical labor, nor could AI substitute for a human employee. Even in the highly automated setting Terry describes, only a small percentage of the overall work at Labcorp is fully automated. Though the work of assembling medical testing kits is not rhetorical for the employees, and could be considered a human-in-the-loop configuration of load sharing (in which the majority of the load is handled by machine), it demonstrates a similar principle: human hands, one of the body's most nuanced and complexly articulated parts, distinguish workers from machines through a combination of fine and gross motor control.¹³ Furthermore, even in the highly automated setting of Labcorp, rhetorically complex communication presents challenges that can only be addressed by humans.

Configurations of rhetorical load sharing can become normalized over time as well, and we return to technological advances in photography to help illustrate this point. At a time when early smartphone cameras began to share resolutions comparable to point-and-shoot cameras, from approximately 2010 (marking Apple's launch of the 5 megapixel iPhone 4) onward, identifying "phone photography" on

13. Although some machines are superhumanly precise or can perform incredibly delicate operations, few if any are able to move quickly, accurately, and precisely over considerable distance, with the capacity to both support a heavy payload while being delicate all in a single machine. For now, human hands remain complex and difficult to replicate across their entire range of applications, but are replicable for specialized applications (e.g., robotic surgery or pick and place machines).

Flickr was a novel practice.¹⁴ Tags allowed photographers to distinguish images produced on a traditional camera from those shot on a phone. This tagging served multiple purposes, often showcasing that the quality of the phone camera was comparable to that of a dedicated camera, but also highlighting the capabilities of the photographer—even without equipment that had greater technical specifications, the photographer could produce striking photos. Such sharing helped demonstrate that phone cameras could carry a rhetorical load comparable to dedicated cameras. Today, identifying an image as “shot on my phone” would seem strange given that mobile photo-sharing platforms such as Instagram enable users to instantly share pictures from their phone in real-time. Phone photography has become commonplace. However, film photographers often include information about the film and camera they used, as well as notes about their development process, because that metadata is not automatically embedded, nor is film the expected medium for social media posts. Novice photographers who want to be perceived as more skilled will still sometimes note that they shot their images with manual settings, whether on a phone or camera. Similarly, #NoFilter is sometimes used in an attempt to distinguish oneself from users creating comparable output with automated tools—signaling authenticity or greater effort, like calling attention to using manual settings. Currently, the norms of disclosing AI use are still crystallizing across media and fields. However, it is conceivable that in the near future, disclosing the use of generative AI as part of writing workflows may seem as unusual as tagging a photo with “shot on my phone,” while the equivalent of “#NoFilter” for manually typed texts might remain, e.g., #NoGPT. I wrote this sentence, not ChatGPT—but I only disclose that because ChatGPT’s output is capable of being indistinguishable (leaving aside whether “I” refers to John, Michael, or the collective assemblage of ourselves, technological infrastructure, and an editorial team in this instance ...). As new technologies become commonplace, however, it is important to remember that any media filters—mediates—our perceptions. #NoFilter invisibilizes the filtering that cameras always already do. As one example, Michael Bradley’s exhibition of portraits depicting Māori people demonstrates how colonial photographers using wet-plate photographs literally erased culturally significant tattoos (Bradley, n.d.). The chemical process of wet-plate photos captured only hints of traditional tattoos, but juxtaposed with modern digital photos the erasure is starkly clear.

Metaphors for Writing

Contrasting earlier analog photography, in terms of image quality alone, there is little perceivable difference between the camera built into a smartphone and

14. The iPhone 4 was only a few megapixels shy of most point-and-shoot cameras at the time it was released. However, it was still well below even older entry-level DSLRs, which had two to three times the resolution depending on the exact model and price. See Diaz (2004) for more.

an image captured on a dedicated digital camera. For everyday applications, the visual difference is negligible. The average phone camera suffices in most situations without additional bulk or equipment, and it would be difficult to tell by sight alone what type of camera produced an image. One might then wonder, what is the value of a professional photographer? Many well-intended relatives at weddings have asked the same question, smartphone or DSLR in hand, often to the frustration of professionals hired to photograph the occasion. When an average person cannot reliably tell the difference between automated amateur output and professional output (or in the case of personal wedding photos, may not care) what is the value of a professional? Similarly, many administrators and colleagues have asked, what is the value of a technical communicator or a writing instructor, particularly now when ChatGPT can do it?

There are technical differences that matter to photographers, just as technical communicators and writing instructors recognize nuances that non-subject matter experts might overlook. If a finished image needs to be printed at a large scale, the subject of the image demands additional megapixels (e.g., product photography), or a telephoto lens is required (e.g., wildlife or sports photography), a smartphone will not do the job. Additionally, dedicated cameras priced for professionals allow for layering additional technology beyond what a smartphone offers, e.g., remote flash triggers, external microphones, synchronization tools, etc. Increasingly, accessories and apps are available to enable smartphones to function similarly, but rarely to the same professional standards. Similarly, generative AI tools cannot yet produce an entire owner's manual for a vehicle, a textbook, contracts, and other forms of writing at scale or with specialized applications. Though automated proofreading tools and generative AI support the work of technical editors for example, such tools do not replace the ability of human editors to recognize nuanced rhetorical situations that can inform the choice of a single word within a text. As Kate explained in her interview, automated proofreading tools help teams collaborate effectively by handling tedious or repetitious work. But they do not understand user needs, and do not anticipate potential issues that may arise from the articulation of a text any more than a camera or a gym machine does—we cannot safely assume that generative AI will not hallucinate dangerous gym machines. Consequently, expert guidance and feedback—and human decision-making—remains important to learning to use automated tools effectively. For all of these technical differences, what distinguishes a professional photographer from someone with a camera, or a technical communicator/writer from someone with a computer, is not just technical proficiency, but the ability to translate user needs into a rhetorically effective deliverable while prioritizing user experience. Authentic encounters and dialog between users and designers, between audiences and writers, are still foundational to effective writing. The ability to understand human experiences, recognizing the humanity of another person, human empathy—though cliché regarding AI and other machines being emotionless—is an important part of the

value of human communicators. Human empathy provides an ethical foundation for participatory design and other methods of communication.

Generative AI, at its best, gives an appearance of understanding, whether that is understanding empathy, logic, rationality, etc., while only ever producing statistically probable results. At the same time, similar to human writers, generative AI often breaks down when responding to “non-default” or novel prompts that fall outside the norms of training data. But unlike human authors, generative AI systems cannot reflect, conduct a post-mortem, or be held accountable for their decision-making process.

The current limitations of generative AI also yield technical differences that are far less nuanced. Although we disclosed earlier that a sentence was typed by a human, generative AI could not write this chapter. It is at times a useful writing assistant, but is not a co-author. As Johndan Johnson-Eilola, Selber, and Eric York similarly concluded in their 2024 *Journal of Business and Technical Communication* article, “When it comes to creating high-quality, consequential instructions, ChatGPT might be better seen as a collaborator than a competitor with human technical communicators” (p. 208). They reached this conclusion because ChatGPT failed to generate safe and effective instructions for a home COVID test, with ChatGPT providing incorrect instructions about nasal swabbing time, no clarification about swabbing depth, and other safety issues. Ultimately, they argued that ChatGPT is useful as a drafting tool when multiple rounds of revision are involved. Generative AI can produce a helpful zero draft to work from. Currently, the technology isn’t capable of replacing a human writer. And LLMs may never be capable of fully automating technical communication tasks such as writing effective instructions, particularly as long as AI language models hallucinate. But the technology is capable of occasionally producing output that at a minimum resembles professional writing, and at best provides generically adequate output for simplified rhetorical situations that are tightly constrained. In other words, in 100 attempts on automatic mode, the generative AI might produce a few reasonably good results, but does not replace a technorhetorically proficient expert.

AI as Extension of Automation

To be clear, our interview participants all have varying degrees of management experience, and shaped how their organizations adopted automated technologies and AI-driven technologies long before the public release of ChatGPT. Kate and Bridget actively contributed to the development of AI-driven systems within their respective organizations over the past decade. We clarify this here to avoid giving the impression that we think readers should take critiques of AI or other forms of automation lightly, or that developers are inherently aware of systemic issues. We do not believe that people become automatically critically aware simply by using a technology, or even osmotically by building systems, and we recognize that developers with good intentions are constrained by institutional power and their

lived experiences. We trust that our participants recognize their ethical responsibilities as technical and professional communicators when contributing to the design of automated systems, and they made that awareness explicitly clear at times during the interviews.

Having said that, attending to the associated oppressive histories of automation is also an ethical responsibility before we transition into our final interview chapter with Terry. In “This Is Not a Response,” (2024) Casey Boyle alludes to histories of automation and oppressive labor practices when defining the “intelligence” in artificial intelligence. Boyle argues, “It is increasingly clear that what we mean by intelligence refers to the products and processes that humans are paid to do that corporations would prefer to not pay humans to do anymore. Following that, we might think of AI as artificial human labor” (2024, p. 307). Building on this critical definition, Boyle cites Sarah T. Roberts as well as James Brown, Jr. and Gregory Hennis, extending their arguments that “media platforms outsource responsibility to moderators and users.” Boyle clarifies that “AI goes further. If there are humans involved, they become the liability sponge through ‘human error’ so that we never scrutinize the algorithms, models, training. Following that, we might once more reconsider AI as ‘artificial responsibility’” (2024, p. 308).

We would argue that this liability sponge principle applies equally to arhetorical uses of “intelligent” technologies. Critics rightly call out AI systems—including webcams and automatic camera sensors that fail to properly recognize people of color—as racist and oppressive. But the rhetorical framing of such critiques matters. If the system allows users to make adjustments, and users have the technorhetorical knowledge to override or adjust problematic defaults, then it becomes insufficient to critique the technology as racist or oppressive and leave it at that. When systems *don’t* afford user control, or when defaults go unexamined and racist outcomes appear to be baked in, then the responsibility lies not just with the technology or with its user, but with its designers, implementers, and managers. To illustrate this dynamic, we take a technical and historical deep dive into analog and digital cameras. These technologies, unlike most contemporary AI systems, offer more transparency in both their technical construction and their cultural histories. Their longer timelines and documented sociotechnical contexts allow us to articulate how managerial, institutional, and user decisions shape automated systems and their impacts.

Throughout the book, we have emphasized the importance of attending to management practices—human decisions about when, where, and how automation is implemented—as well as user practices when evaluating AI systems. Rather than accept automation’s impacts as technologically determined, we advocate for distinguishing between a system’s capabilities and the choices made about how it is used. This distinction becomes especially urgent when evaluating automated systems that reproduce racist or otherwise biased outcomes.

A key difficulty, however, is that most large AI systems today are opaque by design. Their inner workings are black-boxed, often protected by proprietary

constraints, and in many cases not fully understood even by their creators. While researchers are making slow progress in “circuit tracing” and model interpretability (Heaven, 2025), for now, identifying the causes of harmful outputs—such as racist language generated by LLMs—can feel like debugging while blindfolded. We know these systems can produce racist outputs, often as a direct result of biased training data or inadequate screening of user prompts. Their designers know it too. But diagnosing and intervening in the technical causes remains difficult, especially when the systems are trained on vast, sometimes undocumented, datasets and accessed through commercial APIs.

Given this opaqueness, we turn to photography because it provides a more *traceable* technology. With photography, it is possible to identify how specific design choices (such as exposure defaults, film chemistry, or skin-tone calibration) produce problematic outcomes—and to situate those choices within broader institutional and historical practices. In doing so, we are not displacing critique from AI to photography, but using photography’s more visible lineage to help surface and analyze embedded biases in automated systems.

Like Selber and York, who express caution around teaching “prompt engineering” due to the variability of AI outputs and the lack of meaningful feedback mechanisms (2025), we are similarly cautious about overstating users’ control over the automated output of AI-based systems versus users’ ability to revise, repurpose, or reject AI-generated content. Writing teachers have an ethical responsibility to teach students how to critically work with AI-based composition technologies. Skilled photographers can dial in the camera’s output, often reducing if not eliminating the need for extensive editing, but we recognize that we do not yet have such fine-tuned control over how AI-based systems generate their outputs in every situation. This section focuses on automated systems that, unlike LLMs, afford more transparency and accountability. These cases offer a useful starting point for developing ethical and rhetorical orientations that can be applied—even if only imperfectly—to more opaque systems like generative AI. Photography is a representative anecdote of automation deployed over a long timeline, and one which has been ethically fraught in ways that illustrate some of the current ethical issues with generative AI.

No Neutral Grey: Inequities in Imaging

In 1888, George Eastman launched Kodak with the slogan “You press the button, we do the rest” (Eastman Kodak Company, n.d.) to encapsulate the idea that amateur photographers could simply push a button and create lasting photographs with minimal effort—much like modern AI marketing hype. New photographers did not need equipment, chemicals, and technical knowledge to develop Kodak film into photos. To consumers, the development process was seemingly automatic: mail in their camera to Kodak, receive photos and their original camera back with a new roll of film inside. By Knowles’ definition, this is undoubtedly

a human-in-the-loop system, though Eastman's slogan disguises a considerable portion of the rhetorical load involved in creating a photo. According to Kodak's history, nearly 140 years ago, "Eastman had a goal to make photography 'as convenient as the pencil'" (Eastman Kodak Company, n.d.). A fitting technological point of reference as the world questions the value of writing drafted by hand, but again, a slogan that doesn't draw attention to the rhetorical agency involved. Compared with a smartphone, the idea of physically mailing an entire camera and waiting days for the results is laborious and slow, but this technological shift opened opportunities for amateur photographers to document their perspectives on the world around them.

Of course, like with other forms of visual documentation, photographers have used cameras to define reality towards a variety of ends, at times challenging the status quo and at others reinforcing oppressive gazes and -isms along the way (again, see Bradley's PUAKI exhibit as example). Photographers, much like writers, make nuanced rhetorical decisions about how to represent the world visually, working with the technological constraints and affordances of cameras, and may retain varying amounts of control over the development, editing, and distribution processes. In the example of Kodak's slogan, the photographer's control over the appearance of an image largely ended at the press of a button. For many amateur photographers, and even professionals working at a large volume, this was the case up until the advent of consumer digital photography in the 1990s. Development labs and Polaroids handled the chemical processing for anyone who chose not to develop their own film and retain the associated rhetorical control. Although there are parallels between the basic process of sending off photos for development and mailing medical samples for testing described in the next chapter, there are also important distinctions between Kodak's mass production and Labcorp's mass *customization* processes. Labcorp affords the modern equivalent of film manufacturers producing film rolls in which individual frames could offer different ISO sensitivities or color balances—as well as precise tracking of individual medical samples for accurate,¹⁵ safe, and auditable testing (equivalent to the ability to track an individual photo throughout the development process and tailor adjustments)—equivalent to mass bespoke photo development.

Historically, the film development process was highly automated for efficiency and profitability, and technicians processing consumer-grade photos did not individually evaluate each and every photo. Rather, they would develop and print entire rolls of film using a few frames as points of reference. As photo developing systems progressed, machines would be calibrated at the beginning of the development process and could develop and print multiple rolls of film before needing adjustment to maintain designated tolerances. This batch processing,

15. Although racial disparities and biases persist in medical testing and healthcare more broadly, Labcorp is taking steps to address some disparities given the company's capacity for customization. See Rivas, 2024.

when combined with technical limitations of film technology and racist defaults, reinforced existing systems of oppression. Shirley cards (Wessling, 2023) perpetuated systemic biases by leading to photos of White people and people with lighter skin tones being accurately exposed when printed, while printers underexposed and consequently darkened the skin of people of color. Shirley cards were a tool used to calibrate color and exposure settings during the photo developing process to help ensure that printed photos looked like what the photographer saw. The Shirley cards consisted of color swatches and a photo of a White woman. Technicians would receive a copy of the card from a film company, print the same image on their local machine, and then compare the two to determine if the machine needed to be adjusted. However, because the default exposure set by the Shirley card was calibrated for White skin, many technicians never adjusted the calibration for images of people of color (Wessling, 2023).

For our analysis, it's also important to understand the technical limitations of analog film compared with modern digital sensors, and the concept of dynamic range. Compared with the human eye, even the best modern digital cameras do not capture light in the way that the human eye perceives it, particularly in low-light situations. As a simple example, if you are reading this text on a screen with a white background in a dimly lit room, and you were to photograph the screen, it would likely be readable, but detail in the shadows off-screen in the resulting image would likely be lost compared with what you are able to perceive by sight, perhaps even appearing completely black. Similarly, if you were to adjust the camera's exposure settings to accurately expose for details in the shadows, the text on screen might no longer be readable in the captured image, likely appearing brighter than what you perceive, or even being completely washed out and flat white. It is worth noting here that smartphone cameras increasingly enable HDR mode (High Dynamic Range) by default to diminish this limitation of camera sensors. Modern digital cameras have sufficient dynamic range to allow for adjusting exposure significantly after the fact without losing much detail, if the initial capture is reasonably balanced and shot in RAW format. This enhanced dynamic range increasingly enables automatic HDR rendering of images, though the technique can easily yield lighting that appears unnatural. That is, some detail in shadows and highlights can be recovered as long as they are not completely black or white, though color accuracy and some detail may still be compromised in the process. Blown highlights or clipped shadows cannot be recovered (with the growing exception of generative fill where AI estimates statistically probable details). Unlike with AI, these affordances are well-documented, predictable, and can be addressed with relative ease when processing digital photos.

Analog film, however, was less forgiving, with a narrower dynamic range of potential detail being captured in any given scene. The film was manufactured with a specific dynamic range and preset sensitivity to light. Consequently, if the film was manufactured to be sensitive to lighter areas of an image, as was often the case with film marketed to predominantly White audiences, it would not capture

as much detail in the face of a Black person as it would a White person by default. And with Shirley cards (Wessling, 2023) being used as a point of reference, even subsequent changes to the chemistry of film did not fully resolve this bias in the largely automated film developing and printing processes. Thousands of examples of this bias in historical images can largely be attributed to defaults that were part of automated photo developing processes. However, we describe this property of film chemistry to emphasize that the under- or overexposure of skin tones is not technologically determined, nor a fixed racist property of analog film or attributable primarily to Shirley cards, as abbreviated histories of the technology sometimes imply. Historically, photographic development was an entanglement of layers of infrastructure *and* human actors, a process which became more automated over time, and less reliant on chemical limitations, which subsequently shifted the agency (or the liability sponge) of photographers.

In modern digital cameras, most photographers do their own processing rather than sending a roll of film to a lab or sending RAW files to a developer. This is a significant contrast to Kodak's push-button human-in-the-loop system. It was a shift to machine-in-the-loop in which photographers maintain control over the process of developing images from the time they push the shutter-release button until the finished pictures are distributed. If the photographer is shooting RAW files and understands how cameras calculate exposure, even if they mess up the exposure or the white balance in camera for any given shot, they can still reasonably fix the image when processing if the initial capture is within the limits of the camera's dynamic range. Even standard DSLRs and point-and-shoot mirrorless cameras have a decent dynamic range that allows for ± 2 stops of light (a measurement of how much light enters the camera) to adjust exposure without losing detail. White balance settings, and additional color correction are available even in basic free software. Photographers can easily see where highlights and shadows have been clipped via indicators on screen, and most new mirrorless cameras provide a real-time preview of the exposure before an image is captured. In other words, a photographer would really have to mess up the initial exposure of an image to not be able to accurately depict skin tone with a little adjustment regardless of who or what they're photographing! Exposure is easy to correct as long as the base image is close, and even in complex compositions that might require multiple different exposures of parts of the scene to accurately portray those contrasts, a professional should be able to light accordingly to avoid any need for adjustments after the fact. This is a large part of why even famously experienced photographers such as Annie Leibovitz have been rightly criticized for their pictures representing Black women and people of color—there is little technical excuse for a poor rhetorical choice made by a photographer or editor (Bero, 2022). The camera makes for a poor liability sponge when its inner workings, as well as its affordances and limitations, are well understood.

Even anti-paparazzi clothes that “trick” the camera sensor into exposing for a substantially brighter part of the image aren't foolproof if a photographer has

more than one shot and adjusts the exposure compensation or the camera's metering. That's the exposure part of the issue, and that's often what is most striking when done poorly (e.g., the *Time* magazine photo of O.J. Simpson), because detail is lost or somebody appears much darker or lighter in an image than they would to the eye (Horn, 2016). Issues of automatic exposure (metering, in photographic jargon) are striking, and they're frequently addressed in critiques of automated camera systems that fail to detect people of color (e.g., web cams, automatic soap dispensers, etc.). But then there's also white balance and color correction that can impact how skin tone is represented even when properly exposed.

Personally, when I'm (John) creating a portrait, I err on the side of caution regardless of skin tone because I'm colorblind with about a 40 percent red-green color vision deficiency. I use a color calibration tool—a small card with color swatches that looks much like a Shirley card sans Shirley and with far more color swatches due to better dynamic range in modern sensors—because I can't trust that what my eyes see is going to look the same to someone else. That's not standard practice for most photographers unless they're doing commercial product photography and need to make sure the Coca-Cola can is the correct iconic red. It's not a widespread modern standard in part because calibration cards are relatively expensive tools for being a piece of paper and a plastic case when photographers with normal vision can just eyeball things. The cards cost anywhere from \$50 to over \$100, and their color accuracy slowly and subtly degrades over time. The color card is still a tool, so I have to make adjustments based on my best judgement, and sometimes another set of eyes. But it gives me more confidence that I'm not making someone's cheeks unnaturally green or red because the camera misread their skin tone or the editing software's automatic calculation of "average" was off and I couldn't see it. That's also possible because digital technology allows for easy adjustment and revision rather than having a preset range of options or one default. I can see the adjustments happening in real time, at least to the extent that my eyes can distinguish between the hues, and I can hit "undo" if it's not right. With film, I'd be dependent on the color calibration card and the preset of the film manufacturer to get an image close to being accurate on a first attempt, and adjusting would mean creating another print or more. If that color calibration card were to be skewed towards a default of White skin, rather than color accuracy across the color palette, it would create some of the same racist issues that Shirley cards did for years. As the human in this machine-in-the-loop system, I have an ethical and rhetorical responsibility to be aware of the affordances and limitations of the tools I use, as well as the constraints of my own human vision—non-default in its colorblindness, but also privileged because of my subject position as a White man in the US. And I still have to recognize that "technically accurate" does not mean rhetorically neutral. Conveying a warmer or cooler skin tone still carries rhetorical weight, along with the hundreds of other layers of rhetorical choices that go into creating a photo. For these reasons, I would hesitate to allow generative AI to make choices about skin tones in my

work, particularly when the decision-making that led to the output cannot be traced or predicted.

Sarah Lewis, writing in *The New York Times*, clearly shows how this is an ethical and rhetorical issue that photographers need to understand beyond the technical as a result of the history of photographic technology (Lewis, 2019). “By categorizing light skin as the norm and other skin tones as needing special corrective care, photography has altered how we interact with each other without us realizing it” (Lewis, 2019). And we continue to see this racist default show up in arhetorical automated camera and other optical sensing systems—face tracking webcams, automatic sinks and soap dispensers, video filters, etc., which is also a major problem when it comes to datasets for training AI systems. When developers and designers are still “correcting for” non-White as non-default—when those defaults are racist—the new technology reinforces existing biases. Many of these same biases appear in the training sets used for LLMs and other generative AI systems. As writers, as teachers of writing, and as professional writers constructing environments for producing writing, we are ultimately responsible for the biases that we perpetuate or challenge in our writing.

Photographers and writers must be aware of perpetuating racist defaults as more editing tools become automated by AI, with less authorial intervention being the trend. Teachers of photography and teachers of writing alike share an ethical responsibility to teach students to be aware of the capabilities and limitation of the composition technologies they use, and to be comfortable with manually controlling tools to maintain rhetorical effectiveness. Our concern is when that rhetorical decision making is automated with the appearance of convenience and efficiency—particularly when giving users an illusion of fine-grained control—it can quickly become an ethic of expedience (Katz, 1992). We return to this theme, discussing DMV driver’s license photos as one example, in the final dialog section of the Conclusion chapter.

Looping in the Humans

What follows in our discussion in the Conclusion, and prevalent throughout our final interview chapter, is also an extension of the argument we made in our 2024 SIGDOC experience report (Salvo & Sherrill, 2024). Management practices are not technologically determined, and we see this clearly when interviewing Terry. The existence of a tool that enables expedient automation does not negate the responsibility of operators, nor managers who are responsible for creating the institutional infrastructure and conditions under which employees operate. There is an important distinction between technologically determined racism and potentially racist management practices, as the former is an easy excuse for the latter. That is why we chose to include a lengthy technical description of dynamic range even knowing that parsing said description would require more effort from readers. The automated metering in most modern digital cameras, optimized to expose for “neutral

grey” by default, is not inherently racist and does not excuse operators or building managers (though some motion-detecting cameras and sensors certainly are racist by design, intentionally or not, this again should not excuse human agents). The automation of cameras often is racist when their operation is arhetorical—either completely autonomous or when it is a human-in-the-loop system with a human just pushing a button—rather than machine-in-the-loop in which the automated system is relegated to an assistant role.

Interestingly, photography is a doubly-disruptive technology. Above, digital photography and emergent AI enhancement tools have disrupted chemical photography and displaced the medium of record. But chemical photography disrupted realistic painting as an historical representational tool in the late 19th and early 20th century. Photographs became the recording technology that, even though manipulable, were taken as verisimilitude. Then moving pictures came to represent “reality.” Whole books (libraries of books) have represented this shift in technology, and the automation of chemical photographic processes and lost photomats (George Eastman Museum, n.d.). Once a ubiquitous site in strip mall parking lots, the drive-through kiosks allowed home photographers to drop off their film and buy new, and return a week later to retrieve developed photographs. If there were any worthwhile images, they could be enlarged and framed and documented moments in American life. Now they are odd cement bumps among aging mall infrastructure, or awkwardly converted coffee kiosks, and even foodie locations one step up from food trucks. The change is striking yet strangely unremarkable in the digital age.

As the next and last interview, Terry has spent decades watching the workplace evolve, from early digitization to the sweeping automation of once-stable jobs. In his interview, he reflects on the disruptions he has witnessed—not just lost roles, but shifts in how labor is valued. Automation, he notes, has stripped away routine, mechanical tasks, yet in doing so, has clarified what remains uniquely human: experience, judgment, creativity, and intuition. While industries have reshaped themselves around algorithms and robotics, Terry sees a paradox—automation has not eliminated work, but instead redefined its essence. He predicts that the new workplaces that emerge will not appear as extensions of past factories and offices, but as dynamic, as yet undefined spaces where human insight complements machine efficiency. Roles once unimaginable will take shape, built on adaptability and ingenuity. For Terry, the future of work isn’t about resisting automation—it’s about recognizing the irreplaceable skills that no algorithm can replicate and building new worlds around them.