Carnegie Mellon University
Tepper School of Business
Emily Barrow DeJeu, Assistant Teaching Professor of Business Communication

Applications of GenAI for Business Communication

One of the many affordances of generative artificial intelligence (genAI) tools that use large language models (LLMs) is their ability to produce high-quality bespoke text. While the extent to which this technology's text-generating capabilities will prove disruptive is currently unclear, it seems likely that in the future, we will see widespread adoption of genAI tools to facilitate communication processes. A next step for educators and professionals, then, is to identify best practices for using LLMs to assist with communication in ways that are both effective and ethical. To that end, this report outlines some approaches for using genAI tools to assist with common workplace communication tasks at every stage of the composing process, from planning to writing to finalizing a text. A "text" in this report refers to any written communication used for business purposes, from a short email to a presentation deck to a lengthy proposal. Figure 1 highlights some of the communication-related tasks that LLMs can support and summarizes the structure of, and the topics discussed in, this report.

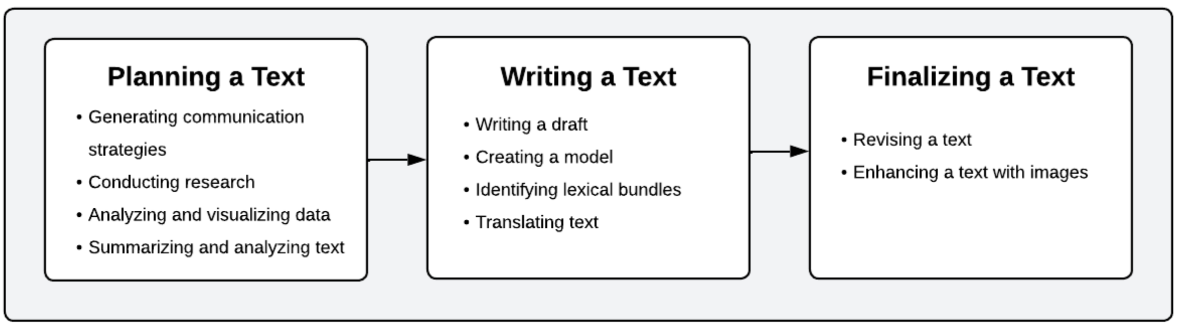

Figure 1: An overview of the communication-related tasks that LLMs can support at each stage of the composing process

Planning a Text

Generating Communication Strategies Writing a document or planning a presentation involves first choosing effective communication strategies. LLMs support this initial brainstorming well. In fact, OpenAI CEO Sam Altman has suggested that LLMs like ChatGPT are best utilized as "reasoning engines" that can shape and refine a user's thinking in the early stages of communication planning.[i] Along similar lines, other experts have indicated that LLMs function best not as "libraries of facts" but as improvisational partners that can spur creativity.[ii] For example, one creative challenge that professionals sometimes face is accommodating complex information for non-expert audiences. Rhetorical theory suggests that turning difficult concepts into metaphors and illustrating them with examples is a productive way to help audiences understand and retain those concepts[iii], but generating metaphors and examples can be difficult. While writers could turn to search engines for assistance, genAI tools are particularly helpful because they provide natural language responses that readily serve as a springboard for creating one's own content.

Beyond creating metaphors and examples, LLMs are useful for suggesting persuasive strategies and lines of arguments that writers might want to consider incorporating into a document or presentation. For instance, if a writer is struggling with a proposal, asking an LLM, "How can I persuade [client] that the plan I am proposing will work and is worth the fee I am charging?" will return a list of suggested persuasive strategies, from demonstrating ROI and scalability to conducting a needs-assessment to sharing client testimonials. Using LLMs to suggest flexible lines of argument like these could prove particularly useful for overcoming writer's block or introducing junior professionals to strategies they had not considered.

Of course, writing an effective proposal requires more than simply making effective arguments. It requires making the right arguments – the ones most likely to persuade a particular audience. Making the right arguments, by extension, requires analyzing what rhetorical scholar Lloyd Bitzer called the "rhetorical situation," or the set of circumstances out of which the need for a particular text arises.[iv] In the case of writing a proposal, an effective rhetorical analysis would involve studying the proposal's audience (their values, goals, expectations, and needs) and context (the situation that the proposal is addressing) and then selecting the arguments best suited to those conditions. LLMs could perhaps expedite some components of this analysis – for example, in the case of writing a grant proposal, an LLM could be tasked with analyzing descriptions of past funded projects to identify trends – but for now, much of this work requires interpersonal communication.

Conducting Research GenAI tools may also be able to help students and professionals conduct research, although their efficacy as research tools is still evolving. Early versions of publicly accessible LLMs were unable to browse the internet and so could not cite sources for the information they provided, nor could they provide up-to-date information. Newer LLMs can access the internet and so are able to provide more timely content and relevant sources. Tools with web browsing abilities may be useful for quickly gathering publicly accessible online sources and generating quick summaries of their content. However, LLMs are currently not capable of evaluating the quality of those sources or suggesting the most effective means of using them for particular projects, so users will still need to cultivate information literacy skills to evaluate sources' quality and reliability and to determine how best to use their information.[v]

Along with their sourcing limitations, LLMs currently have somewhat limited value for research purposes because they can hallucinate, providing confident responses that are factually incorrect or irrelevant. While rates of hallucination will likely drop as technology evolves, users should exercise caution when using genAI tools to gather information and should supplement genAI-driven research with other forms of information gathering. Likewise, users should carefully fact-check any AI-produced content that they plan to use in their own work – a process that OpenAI calls "Human in the loop" and suggests is especially critical in high-stakes communication contexts.[vi]

In addition to considering the veracity of LLM tools' responses, users should also consider the extent to which those responses might be biased.[vii] Bias can creep into AI as a result of its training data – for instance, research has found that using news articles to train natural language processing models can produce gender-stereotyped content.[viii] But human testers can also introduce bias. OpenAI, for example, uses a process called "reinforcement learning with human feedback" that tasks humans with reviewing LLMs' output to ensure that it aligns with human values.[ix] But the subjective process of determining which values to reinforce can introduce unintended biases. For example, early research has suggested that ChatGPT's responses have a left-leaning political bias and that text-to-image systems like DALL-E and Stable Diffusion reinforce gender and racial stereotypes. [x] While the problem of bias in LLMs may improve with time, users should be mindful that AI-generated information may reflect such biases. Again, users can mitigate this problem by conducting additional kinds of research that will expose them to a broader range of perspectives.

Analyzing and Visualizing Data While genAI tools may not be particularly useful for research purposes, new advances suggest that they could help users analyze data, which is often a key early step in communication planning. Although early LLMs were not capable of performing data analysis, more recent versions are.[xi] For instance, Figure 2 shows a line plot that an LLM produced when asked to analyze a large data set containing information about all Olympic athletes between 1896 and 2016.[xii] LLMs with analytical capabilities can also summarize and provide general interpretations of data trends; for instance, when prompted to describe the trend shown in Figure 2, the tool explained that the large peaks and valleys after 1980 correspond the summer and winter games, respectively, and suggested that these peaks and valleys likely reflect an overall increase in and public support for female participation in athletics combined with an increasing number of events added to the summer games in the late twentieth century, compared to the winter games. While explanations like this require further scrutiny, given LLMs' capacity to hallucinate and reproduce bias, this example suggests that LLMs capable of analyzing data can provide helpful starting points for the additional research and argument building necessary to build larger narratives about what data trends mean.

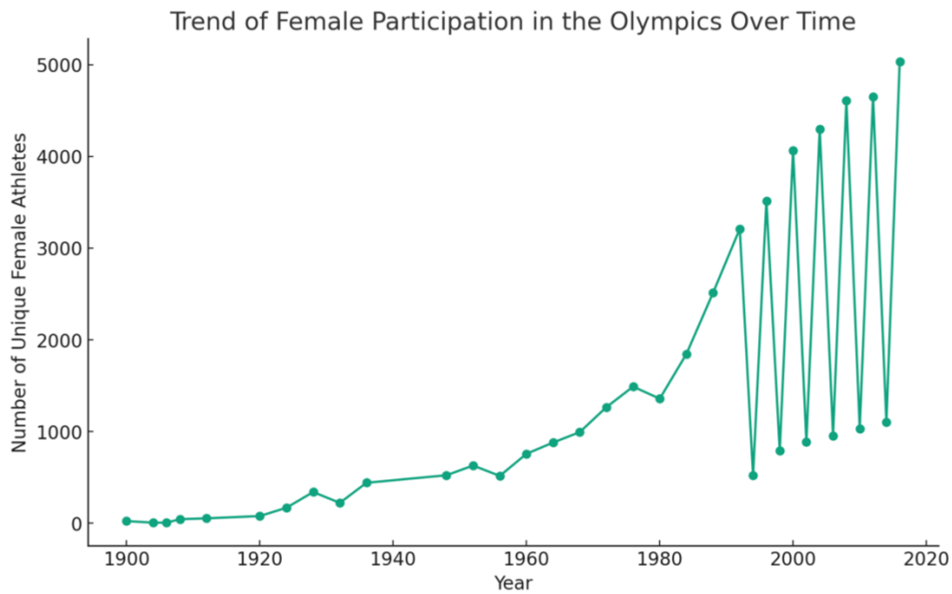

Figure 2: AI-generated line plot showing an increase in female athletes' participation in the Olympics over time

In addition to visualizing data, genAI tools with analytical capabilities can also explore complex relationships between variables by suggesting and performing calculations best suited to the nature of a given dataset. Considering these capacities, genAI tools may prove particularly helpful for users who need to do basic exploratory analysis and produce data visualizations but whose mathematical and programming literacies are not strong. Moreover, because these tools explain the reasoning for their choices, outline step-by-step approaches for each calculation, and allow users to check their code, they could function in part as educational resources for users looking to build those kinds of literacies.

But despite impressive capabilities like these, users should proceed cautiously when leveraging genAI tools for data analysis and visualization. To begin, some computer science experts have found that while LLMs have impressive coding prowess, they make mistakes on complex analytical tasks that require complicated code.[xiii] Likewise, these tools' textual descriptions and interpretations of data trends are subject to the same hallucination and bias problems that plague their responses to non-analytical task prompts. But users may also introduce error – for instance, users without statistical backgrounds may not know what methods of analysis are appropriate for particular kinds of data or how to make sense of statistical results. They would thus be unlikely to catch flaws in an LLM's analysis. Therefore, users with limited mathematical and programming literacies should consider consulting with an experienced analyst to verify AI-generated data analyses before using the information.

In a similar vein, while LLMs with analytical capabilities are generally good at creating effective data visualizations, the visuals they produce do not always align with design best practices. For instance, Figure 3 shows a stacked bar plot that a genAI tool generated to show medal distributions among male and female athletes from the 10 countries with the highest overall medal totals between 1896 and 2016. The original bar plot is not well-designed because the x-axis reflects two different data points (gender and country), so the data is not sorted in a way that facilitates quick interpretation. Likewise, while a stacked bar plot is a good choice to represent this data, stacked bar plots that represent more than two values, as Figure 3 does, can be difficult to read because it is hard to visually compare values that do not start from the same baseline. Users can revise ineffective images like these by prompting LLMs to make specific changes. For instance, Figure 3's revised bar plots reflect requests from a user to separate male and female data, introduce data labels for each stacked bar, and sort the data in descending order. Ultimately, the revised plots are easier to interpret thanks to these design-related changes. But while users can prompt changes like these, they must be familiar with design best practices in the first place, if they are to critically assess and alter a genAI tool's output.[xiv]

Original AI-generated stacked bar plot before user-prompted edits

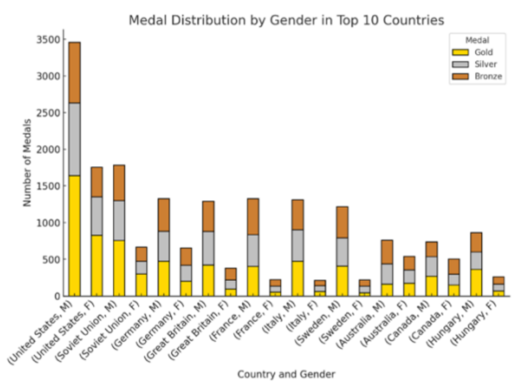

Revised AI-generated stacked bar plot after user-prompted edits

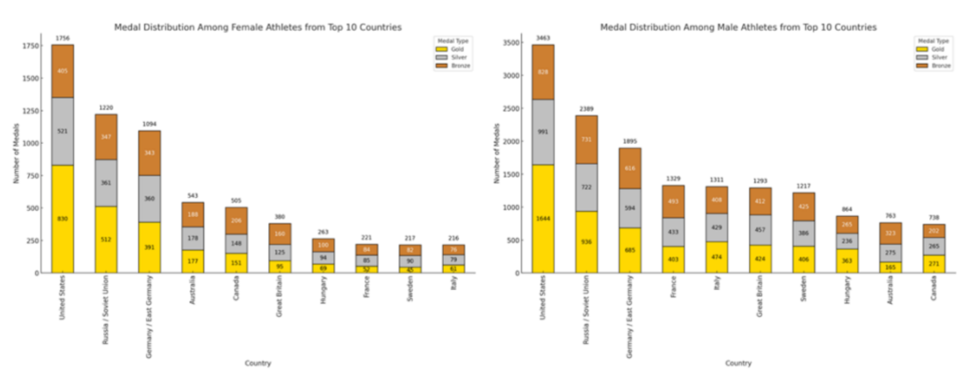

Figure 3: An original AI-generated stacked bar chart that does not align with best practices of data-related design and so is difficult to interpret, and two revised charts whose design-related changes make them easier to understand

Along with concerns about LLMs' analytical accuracy and design-related efficacy, enterprise users may be particularly concerned about uploading sensitive information, such as propriety or confidential data, to a genAI tool because these tools typically leverage user-supplied data for training purposes. It is likely best, therefore, to avoid the use of genAI tools for analyzing sensitive data, but this condition may severely restrict the scope of LLMs' analytical applications in some settings.

Summarizing and Analyzing Text Along with generating ideas and locating possible sources, LLMs can also summarize and analyze texts as part of users' information-gathering and planning processes. For instance, a user could paste the contents of a complex, technical report into an LLM and ask the tool to summarize its content using easy-to-understand language. Or, for a tool with internet browsing capabilities, a user could supply links to web-based resources and ask the tool to summarize the contents of each link, to streamline information gathering.[xv] With adequate prompt engineering, genAI tools can also tailor summaries to meet specific use cases, emphasizing or backgrounding information and offering general analysis to suit a user's stated goals.

However, LLMs' value for summarizing text seems somewhat limited for now. First and foremost, as with their informational outputs, LLMs' summaries should be reviewed for accuracy, particularly in high-stakes contexts, which seems to negate any substantial timesaving value. Second, the tools are limited in what kinds of texts they can summarize: some cannot process lengthy texts, for example, while those that are capable of processing web-based content cannot currently access paywalled information. Finally, while some LLMs can, when prompted, broadly analyze a text's major themes, they are limited in their ability to thoroughly explore and interpret those themes. Moreover, a reader's sense of what is most important and actionable in a text is variable and shifts alongside changes in her context. Since interpreting the implications of a text's information based on evolving real-world constraints is generally a central goal in work-related reading, genAI tools may have limited value when it comes to helping professionals read, analyze, and act on information.

Writing a Text

While genAI tools have some value for generating ideas, conducting research, and analyzing data, and summarizing documents, arguably their biggest affordance is their ability to write relatively high-quality text. Therefore, genAI tools have numerous applications for helping users write various kinds of texts, from short emails to lengthy reports to presentation scripts.

Writing a draft Perhaps the most obvious way to use an LLM for business communication purposes is to ask it to draft a text. Prompt engineering, or the process of fine-tuning the commands one gives a genAI tool, can alter an LLM's output – its content, organizational structure, level of complexity, and style – substantially.[xvi] Therefore, if given a highly tailored prompt, an LLM may be able to write an adequate initial draft and thereby automate most (perhaps even all) of the writing involved in creating a message.

But while genAI's capacity to automate drafting is substantial, it has limitations. To begin, the level of prompt engineering required to produce texts that users can deploy with minimal editing might not be worth the effort required, particularly if the text in question is relatively short and straightforward. Moreover, while prompt engineering can customize a document's content and arrangement, it seems less effective for adapting a text's tone to suit a particular audience, since no amount of prompt engineering can account for the nuance that attends human relationships. Likewise, a user can try to direct ChatGPT to write in his own personal style by offering the tool a few paragraphs of his writing, but in reality, no writer has a single style. Rather, writers adopt different styles to suit different contexts.[xvii] What is more, a user might balk at the prospect of teaching a machine to write in his personal style, since writers tend to see their authorial voices as unique expressions of their essential individuality.[xviii] Drafting with genAI tools, then, will involve grappling with the law of diminishing returns to balance the time spent on prompt engineering versus the time spent on customizing the tools' output.

Drafting with genAI tools may also raise ethical and regulatory concerns. Students who submit AI-generated writing as their own, for example, are likely committing plagiarism, since they are turning in work that is not their own and thus are not demonstrating their own learning. In workplace contexts, which typically place less value on original work, plagiarism might not be a first-order concern, so it might be acceptable to use lengthy passages of AI-generated text in one's own writing. But even in professional contexts, using genAI tools to write might prove ethically questionable or legally risky. For instance, if employees misuse LLMs for research purposes, they run the risk of circulating biased or inaccurate information, which could create problems for their employers. In terms of other kinds of regulatory concerns, companies who are concerned about copyright infringement may be leery of using large quantities of AI-generated texts for official purposes, since there are emerging questions about the extent to which genAI companies may have violated copyright law in collecting training data.[xix] The reverse is also true: Companies that want to safeguard their own intellectual property and data security may be alarmed at the prospect of prompt engineering that involves inputting sensitive information.[xx]

Creating a model Drafting aside, there are other applications of genAI for writing that rely less on the wholesale use of AI-generated text and so could mitigate risks of ineffective, unethical, or legally questionable use. For instance, writers can use genAI tools to generate model documents. Particularly in cases where they are faced with writing in a new genre, writers often seek out model texts to help them get started. Models are helpful because they give a writer a sense of her options – how long her text could be, what kinds of arguments it could use, how it could be organized, etc.[xxi] LLMs can generate bespoke models that are tailored to meet a writer's specific needs. For instance, if a manager needs to alert his team about a new freeze on paid time off (a challenging message to send, given that no one likes receiving bad news), he could ask a genAI tool to generate a model message. As Figure 4 shows, the model is helpful not only because it outlines the kinds of persuasive arguments this manager could make but also because it suggests an overall structure for organizing those arguments. In fact, the model email in Figure 4 closely aligns with what researchers call a "direct" structure for delivering bad news, which professionals frequently use when they want to be candid about negative developments.[xxii]

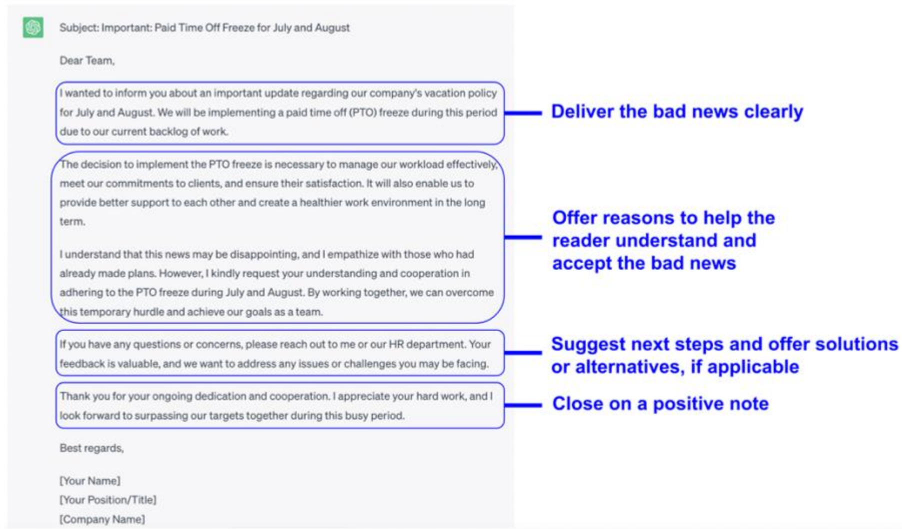

Figure 4: An AI-generated model "bad news" message

Generating multiple versions of the same message or asking the LLM to make changes to the initial model – to make it longer or shorter, for instance, or to change its tone, or to add or remove content – can help writers better understand the range of communication options available to them and help them think critically about which of those options they prefer.

Identifying lexical bundles In addition to using model texts to identify possible persuasive arguments and organizational structures, writers can also mine them for lexical bundles. Lexical bundles are short phrases that are used often in particular kinds of writing. Lexical bundles are useful because they give writers the language that they need to start making particular arguments.[xxiii] Lexical bundles are particularly helpful for writers who need to write in a language that is not their primary language, for whom certain commonly used phrases might be unfamiliar, as well as novice writers who might not have much experience writing in a particular genre. For example, an intern or junior employee who needs to ask her manager for an extension might not have experience making a potentially difficult request like this in a deferential way that is likely to motivate a positive response. Using an LLM to generate multiple versions of a model email could help the writer identify lexical bundles that she could use in her own message. For example, Figure 5 shows four opening paragraphs from four different model emails, all generated using the same prompt – "Write an email to my boss letting her know I need a three-day extension on my project." As Figure 5 indicates, opening phrases like "hope this email finds you well" and "I wanted to touch base / reach out to you / provide you with . . . the deadline" are useful lexical bundles for beginning this kind of message:

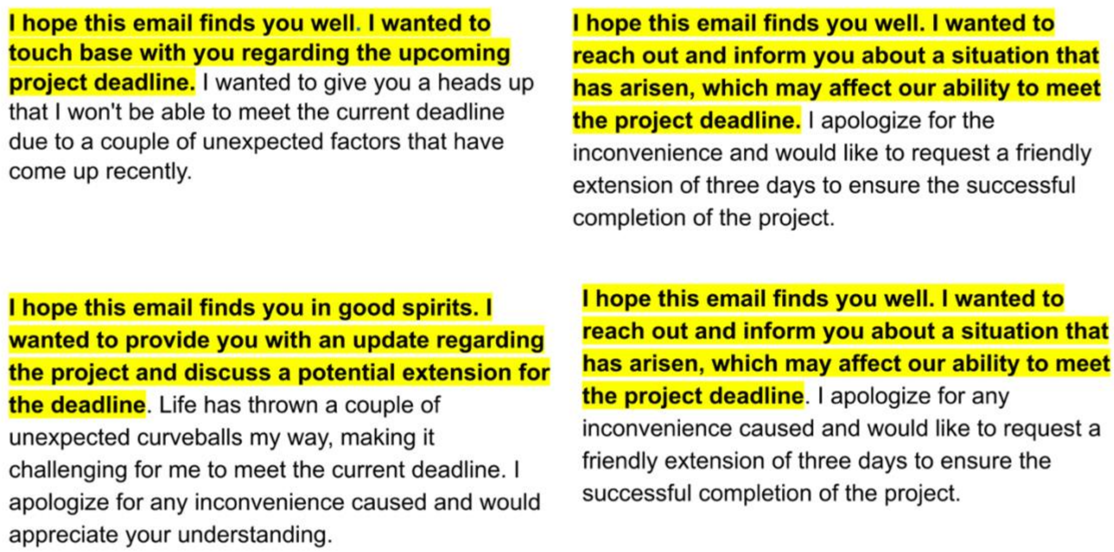

I hope this email finds you well. I wanted to touch base with you regarding the upcoming project deadline. I wanted to give you a heads up that I won't be able to meet the current deadline due to a couple of unexpected factors that have come up recently.

I hope this email finds you well. I wanted to reach out and inform you about a situation that has arisen, which may affect our ability to meet the project deadline. I apologize for the inconvenience and would like to request a friendly extension of three days to ensure the successful completion of the project.

I hope this email finds you in good spirits. I wanted to provide you with an update regarding the project and discuss a potential extension for the deadline. Life has thrown a couple of unexpected curveballs my way, making it challenging for me to meet the current deadline. I apologize for any inconvenience caused and would appreciate your understanding.

I hope this email finds you well. I wanted to reach out and inform you about a situation that has arisen, which may affect our ability to meet the project deadline. I apologize for any inconvenience caused and would like to request a friendly extension of three days to ensure the successful completion of the project.

Figure 5: AI-generated lexical bundles for opening a request-for-extension message

Translating text Along with producing text, genAI tools can also translate text into other languages. As with lexical bundles, translation applications of these tools could prove very useful for writers who need to write in a language other than their primary language, since research suggests that, on top of the normal cognitive burdens associated with writing, these writers may find it difficult to translate their ideas into a less familiar language.[xxiv] A writer could, for instance, draft her entire text in her primary language, if that approach felt more productive, and then ask an LLM to translate it, or she could use the tool to translate particular phrases. This said, writers who use LLMs to facilitate translation will need to verify that their outputs are accurate and appropriate, since even the most sophisticated translation tools can struggle with nuanced variations in dialect and tone.[xxv]

Finalizing a Text

Revising a text Just as genAI tools can be used to help write a text, they can also help revise it. In fact, some LLMs function as robust revision tools capable not only of correcting errors but also making sweeping content revisions. In this way, they surpass the capabilities of older and more familiar digital proofreading tools.[xxvi] For example, Figure 6 illustrates revisions that an LLM made to a student-authored networking email, based on the prompt, ‘Make this message better.’ As Figure 6 shows, the LLM’s output goes beyond simply correcting the email’s grammatical mistakes (shown in red) – it adds paragraph breaks and polite opening and closing statements, and it substantially rewords most sentences, seemingly in an effort to make the message more cohesive.

Original Message

Subject: **Informations** Request

Dear Ms. Hemlich,

I am writing to request **informations** about the job opening you **mention** when we met last week. I know that you are looking for someone with **the** background in economics, but can you tell me more about what skills you are seeking in applicants? As a student at X University, I have one final semester **when** I graduate with a bachelor's degree in economics. I would like to schedule my final semester of classes so I can develop more of the skills that an employer like you finds most desirable. Which classes from your time in college **are** most helpful to your daily work? I look forward to **hear** from you.

AI-authored revisions, in response to the prompt 'Make this message better'

Subject: Request for Information on Job Opening

Dear Ms. Hemlich,

I hope this message finds you well. I am writing to inquire about the job opening we discussed during our meeting last week. I understand that you are seeking candidates with a background in economics, and I would greatly appreciate it if you could provide me with further details regarding the specific skills you are looking for in applicants.

As a student at X University, I am currently in my final semester and will be graduating soon with a bachelor's degree in economics. In order to align my coursework with the skills that employers, such as yourself, find most desirable, I would like to request your guidance in selecting the most relevant classes for my final semester. Could you kindly share your insights on which courses from your own college experience have been particularly valuable in your daily work?

I am eager to hear from you and would greatly appreciate any information you can provide. Thank you for your time and consideration.

Sincerely, [Your Name]

Figure 6: An original message, with errors highlighted in red that are common among writers whose English proficiency is developing, and AI-authored revisions in response to the prompt 'Make this message better'

In addition to making broad revisions, genAI tools can, if prompted, provide line-by-line explanations of their changes. This kind of metacommentary could prove particularly useful for novice professionals or those who need to write in a language that is not their primary language, helping them better understand their range of communication options and, in the case of grammatical corrections, functioning as a kind of *ad hoc* language tutor. For instance, the original message in Figure 6 contains several grammatical errors (shown in red) that are common among writers whose English proficiency is developing. LLMs can not only fix those kinds of errors but also provide a grammatical rationale for each change, thereby offering writers a way to further build their language fluency.

Beyond asking a genAI tool to improve upon an existing message, writers can also prompt the tool to suggest additional strategies for making a message more effective and generate new versions of a message based on those strategies. Figure 7 illustrates the results of that approach: it shows the message that an LLM produced when asked to further improve Figure 6's revised version and then to write a new message that leveraged those suggestions. As Figure 7 shows, the LLM added new content (in blue) that seems designed to help the writer make a stronger impression.

Subject: Inquiry Regarding Job Opening Discussed - Personalized Introduction

Dear Ms. Hemlich,

I hope this email finds you well. It was a pleasure meeting you last week and discussing the exciting job opening at your company. **Our conversation left a lasting impression, especially when we delved into the importance of leveraging economic principles to drive sustainable growth.**

I am writing to inquire further about the specific skills you are seeking in applicants for this position. As a student at X University, currently in my final semester pursuing a bachelor's degree in economics, I am eager to align my coursework with the skills and qualities that your esteemed organization values most. By gaining a deeper understanding of the key attributes you seek, I can make more informed decisions regarding my final semester's course selection.

Having researched your company extensively, I am impressed by its innovative approach to addressing economic challenges and its commitment to fostering social impact. The work your team has done in developing sustainable business models that generate long-term value resonates deeply with my career aspirations. In this regard, I am particularly excited about the opportunity to apply my skills in economic analysis, market research, and data interpretation to contribute to your company's ongoing success.

In addition to my academic pursuits, I have gained practical experience through internships where I analyzed market trends, conducted economic forecasting, and prepared reports on industry dynamics. Furthermore, coursework in advanced econometrics and behavioral economics has equipped me with a solid foundation for analyzing complex economic data and understanding the factors influencing decision-making processes.

I greatly appreciate your expertise and guidance on the most relevant courses **and experiences** that would enhance my readiness for a position with your organization. Your insights into the coursework **or real-world projects** that have been valuable in your daily work would be invaluable to me as I shape my educational path.

Thank you for considering my inquiry. I look forward to hearing from you and continuing our conversation. Your time and consideration are sincerely appreciated.

Warm regards, [Your Name]

Figure 7: Additional AI-authored revisions to the original message in Figure 6, based on suggestions the LLM made for further improving the message's content

But the extent to which large-scale content revisions like those shown in Figure 7 make a text more effective depends on context, so using genAI tools to revise content may prove less straightforward than using them to correct grammatical and mechanical errors. For instance, the new AI-supplied content in Figure 7's revised message assumes that the writer is interested in a position at the recipient's company and wants to promote themselves as a candidate for employment; however, the original message simply asked the recipient for additional information, so these assumptions might be inaccurate. Therefore, while genAI tools show promise as robust proofreading resources[xxvii], their value for making wholesale content revisions seems more limited, at least for now.

Enhancing a text with images While much of the excitement surrounding communication-related applications for LLMs has centered on their text-generating and analytical abilities, text-to-image generators could have fruitful applications for business communication, too. For instance, users who need to create meta images, or the branded images that appear when web content is shared across platforms, could use text-to-image generators for this purpose. Likewise, users who wanted to add images to a presentation deck could use a text-to-image generator to create custom images that suited their presentation content, instead of relying on more generic stock images.[xxviii] But while applications like these, that involve creating custom digital art, have their value, users should be aware that using AI-generated data visualizations may as ethically and legally fraught as using AI-generated text, since copyright concerns apply in both cases.

Conclusion

As this report illustrates, LLMs have myriad applications for business communication and can support users as they plan, write, and finalize work-related texts. Moreover, since LLMs are currently in their infancy, they will likely evolve in ways that produce new communication-related applications. However, despite their current and future value for workplace communication, using LLMs to generate text and images has ethical and legal implications that, at the very least, necessitate cautious and thoughtful engagement with the tools and may require more formal kinds of institutional oversight. For now, users who are interested in leveraging genAI tools for business communication purposes would do well to test methods that augment and assist, but do not outsource, their existing communication practices.

---

[i] More specifically, those familiar with LLMs suggest that they are sometimes misunderstood as search-engine replacements, when what they really offer "is a compact and elegant synthesis" of a storehouse of related texts. Rather than using LLMs as research tools, therefore, Altman suggests that tools like ChatGPT are better utilized as "reasoning engines" whose abilities are best leveraged when users "ask it to compare concepts, or make counterarguments, or generate analogies, or evaluate the symbolic logic in a bit of code." See "Does Sam Altman Know What He's Creating?"

[ii] See "ChatGPT is great – you're just using it wrong."

[iii] For scholarly context, see "Metaphor." For an example of a powerful and longstanding metaphor in the realm of technology, consider the desktop metaphor, a computing interface metaphor created in the 1970's by researchers at Xerox PARC. See "Desktop Metaphor" for a look at the metaphor's enduring legacy in the field of user experience.

[iv] See "The Rhetorical Situation." In conceptualizing the rhetorical situation, Bitzer draws heavily on ancient Aristotelian and Ciceronian rhetorical theory.

[v] According to the American Library Association, information literacy is "a set of abilities requiring individuals to recognize when information is needed and have the ability to locate, evaluate, and use effectively the needed information." See "Information Literacy."

[vi] See "Safety Best Practices." While companies' recommended practices for using genAI safely will likely evolve over time, OpenAI has been a frontrunner in making LLMs publicly accessible and thus a leader in recommending policies surrounding the development of future LLMs, so in the near future, other genAI companies might hew to OpenAI's suggested safety practices.

[vii] For a broad look at how human bias implicates AI as well as possible solutions, see "What Do We Do About the Biases in AI?"

[viii] See "Text Embedding Models Contain Bias. Here's Why That Matters."

[ix] See "Our approach to alignment research."

[x] For more information on political bias in LLMs, see "The politics of AI: ChatGPT and political bias." For more on how text-to-image models reflect and might amplify existing human stereotypes, see "Humans are Biased. Generative AI is Even Worse" and "Stable Bias: Analyzing Societal Representations in Diffusion Models."

[xi] For instance, in mid-2023, OpenAI released a feature called 'Code Interpreter' for its premium platform's users.

[xii] See "120 years of Olympic history: athletes and results."

[xiii] See "Scared of AI? Don't Be, Computer-Science Instructors Say."

[xiv] For a helpful primer on data design best practices, see Johns Hopkins University's "Data Visualization" library guide.

[xv] One suggested application for leveraging LLMs' summarizing and analyzing abilities involves tasking them with summarizing major themes in product reviews and customer feedback. See, for instance, "Amazon is using generative A.I. to summarize product reviews."

[xvi] For a comprehensive overview of prompt engineering best practices, see "Prompt Engineering Guide." While prompt engineering may be a valuable form of expertise now, in LLMs' early stages, some experts predict that

in the future, advances in LLM technology will render prompt engineering relatively obsolete. See, for example, "AI Prompt Engineering Isn't the Future."

[xvii] See, for instance, *The Mythology of Voice* for more information on the extent to which experienced writers cultivate different styles to suit different rhetorical contexts and thus do not have singular "voices" in which they write

[xviii] For more on the extent to which writers feel that their authorial voice reflects their unique personal traits, see Peter Elbow's work on this topic, including *Writing with Power*, "Voice in Writing Again: Embracing Contraries", and "Individualism and the teaching of writing: Response to Vai Ramanathan and Dwight Atkinson."

[xix] For an overview of copyright-related issues, see "Generative AI Has an Intellectual Property Problem." It is currently unclear how mounting copyright challenges against genAI companies will play out in the courts, given that existing law provides little guidance on the extent to which using copyrighted material to train LLMs constitutes infringement. See "Interoperability of Artificial Intelligence and Copyright Law Examined by Congress" for more.

[xx] For an overview of the ethical and legal risks employers face in allowing employees to use LLMs at work, see "My Employees are Using ChatGPT. Now What?" and "ChatGPT and Generative AI: Key Legal Issues."

[xxi] For an overview of how writers use model texts, see "Learning to Write in a Genre: What Students Take from Model Texts."

[xxii] For an example of research into the typical structures of "bad news" messages, see "Factors in Reader Responses to Negative Letters: Experimental Evidence for Changing What We Teach."

[xxiii] For an overview of lexical bundles in academic writing, see "If you look at…Lexical bundles in university teaching and textbooks" and "As can be seen: Lexical bundles and disciplinary variation." While existing scholarship has focused on lexical bundles in academic writing, the emergence of publicly accessible genAI tools provides professionals with a way to quickly generate lexical bundles that are relevant for workplace uses.

[xxiv] Scholars have been particularly interested in the extent to which having to write and publish in English constitutes a form of linguistic injustice for academics whose primary language is not English. For an overview of this perspective, see "Is linguistic injustice a myth? A response to Hyland." While scholarship like this tends to center on the challenges that international scholars face, similar challenges likely attend professionals who must frequently communicate in a language other than their primary language.

[xxv] For more information on the extent to which particular translation applications, including some LLM applications, are capable of producing high-quality translations that approximate a native speaker, see "Google Translate vs. ChatGPT: Which One is the Best Language Translator?" and "Is ChatGPT a Good Translator? Yes, With GPT-4 as the Engine."

[xxvi] For instance, more familiar applications like Grammarly and Microsoft Editor have, historically, had the capacity to point out errors and identify opportunities to improve syntax but not to suggest significant alterations to a text's content. However, this may change: Microsoft and Grammarly have launched their own genAI-powered tools (Microsoft 365 Copilot and GrammarlyGO) which seem designed in part to match the broad revision capabilities of open-access LLMs like ChatGPT.

[xxvii] See "ChatGPT or Grammarly? Evaluating ChatGPT on Grammatical Error Correction Benchmark" for additional information about how LLMs like ChatGPT perform on error detection and correction compared to other kinds of proofreading software.

[xxviii] For more information about these applications, as well as other potential enterprise applications for text-to-image tools, see "6 examples of real businesses using DALL-E for visual content."