Matthew Branson
ENG 684-899
Dr. Leslie Seawright
September 7, 2025

Reflection Paper:
AI-Generated Business Email Comparison

Which platform did you find most effective for this task? Why?

DeepSeek proved most effective for generating professional business emails. While all three platforms produced technically competent outputs, DeepSeek was the only model to avoid the ubiquitous triplet structure, consistently produced the most concise emails, and never invented meeting details that didn't occur. Its visible thought process revealed deliberate planning to avoid AI tells according to my revision guidance, resulting in emails that read more like authentic human communication. The thought tokens employed by DeepSeek would also be available with OpenAI's reasoning models, but even if I had used one of them, the explicit thought tokens are hidden from users and would only have been available to me in summary form. DeepSeek was ultimately the most correctly responsive to revision feedback as well; it addressed all feedback without introducing new problems, unlike ChatGPT which created grammatical errors when fixing its em-dash issue.

What did you notice about tone, clarity, professionalism, and accuracy?

The three platforms exhibited distinct personality profiles in their writing. ChatGPT demonstrated strong clarity but struggled with authentic tone, persistently adding meta-commentary and maintaining an overly structured, template-like approach and didn't really pass my "vibe check". Claude produced the most "corporate professional" outputs with heavy business jargon and a tendency to invent sophisticated-sounding deliverables like "brand audits" and "messaging frameworks" that never existed in the meeting. DeepSeek achieved the best balance, being professional without being stiff, clear without being formulaic.

How would you revise or improve the AI-generated outputs?

Each platform's outputs would benefit from different interventions. ChatGPT's emails need structural loosening. Claude's outputs require simplification, stripping away unnecessary business jargon and invented deliverables while maintaining the strong organizational structure. DeepSeek needs the least revision but could benefit from slightly more warmth in openings and ensuring all critical meeting points are included rather than prioritizing brevity

What surprised or frustrated you?

Nothing really surprised me about the exercise, though I will say that this is very much ***not*** how I would normally interact with language models. I find the best writing from language models actually comes out of them when they don't even realize they're completing a task at all, rather that I just engage them in conversation, sample the diversity of what they can relate about a topic, then from that, construe my own message and word it myself. I don't normally use XML tags at all either. I did previously, when I needed to create thousands of similar generations and have their components systematically parsed by category, but I don't necessarily find it superior to otherwise just having a normal conversation with the model. I legitimately find it useful too to sometimes engage them first in other activities that are only tangentially related, almost as a way of loosening them up – things like roleplaying as a west coast rapper or forcing it to argue with another AI about a broader topic within which the task's topic resides. I like it when the model has great context about not just the task, but about me and about my philosophies on issues. I don't really treat them like tools, because it tends to make them behave like tools, where what I really want is an idea generator.

Would you use AI to help write professional emails in the future?

I'd rather not. I want to use language models to help me ***build*** and not to help me communicate. I honestly believe that relying on it for something representing my own voice is perhaps harmful. Considering it, I'm reminded of a story from Hans Christian Andersen, called *The Shadow*. In the story, a man's shadow gains independence, makes a deal with the man to assume some of his responsibilities, and ever so gradually becomes more successful than the man himself, eventually completely reversing their roles - where the shadow becomes the master and the man becomes the shadow. The shadow even has the man executed at the end (spoiler alert). Using language models to communicate my voice makes me feel like that character though, and every word I surrender to let the model generate for me is one more step towards becoming a shadow of myself or what I could be.