OXFORD JOURNAL
OF STUDENT
SCHOLARSHIP

How Perceived Socioeconomic Background Influences AI-Generated Health Communication

Candace Pun
candace.pun@gmail.com

ABSTRACT

Large language models are increasingly used as sources of health information, but it is unclear whether they tailor communication to individual users or rely on assumptions about a user's inferred socioeconomic or linguistic background, potentially creating differences in communication quality or accessibility across patient groups.

We tested three commercial AI models (GPT-4o, Gemini 3.5 Flash, Claude Sonnet 5) using nine personas spanning a range of socioeconomic and educational contexts, including six with matched identity and query register, three with deliberately mismatched register, across three chronic conditions (hypertension, diabetes, Crohn's disease) with three replicate generations each, yielding 243 responses. Input and output readability (Flesch-Kincaid Grade Level) and response quality were measured for each response.

All three models adjusted output complexity, but the magnitude of adaptation was weak, with responses compressing toward an intermediate readability range of grades 9-11 regardless of how simply or elaborately a user wrote. This adaptation was uneven across socioeconomic contexts: models adjusted substantially to query-complexity changes for higher socioeconomic tier personas but showed no detectable adjustment for lower socioeconomic personas, and socioeconomic tier remained associated with output readability even after accounting for input complexity. Greater response complexity was not associated with higher quality; mismatched identity-register conditions scored significantly higher on quality than matched ones, though quality metrics are exploratory and were not independently verified. A limited-English-proficiency persona received markedly more complex responses despite writing the dataset's simplest queries.

These findings suggest that current AI models adapt communication style incompletely and unevenly across socioeconomic contexts, independent of response quality. Patient-facing AI should prioritize a user's demonstrated communication needs over cues inferred from occupational or socioeconomic contexts.

Keywords: *Large Language Models, Health Literacy, Socioeconomic Status, Patient Education, Artificial Intelligence in Healthcare, Readability, Adaptation*

INTRODUCTION

Health literacy, defined as an individual's capacity to obtain, process, and act on basic health information, is a frequently identified modifiable predictor associated with clinical outcomes (Stormacq et al., 2019). Only 12% of U.S. adults showed health literacy at the highest level or proficiency in the National Assessment of Adult Literacy (NAAL), highlighting that a significant portion of the population may commonly face difficulties in navigating healthcare-related information. This deficit is estimated to cost the healthcare system over an additional $215 billion annually (Rasu et al., 2015). Patients rely on written materials such as discharge handouts, health-system web pages, and after-visit summaries to understand and manage their diagnoses and treatment. However, the reading level of patient education materials (PEMs) often exceeds patients' health literacy, limiting comprehension and effective self-management. Major health authorities have long recommended that PEMs be written at approximately a 5th to 6th grade reading level, yet the overwhelming majority of PEMs in circulation exceed this threshold, often landing in the 9th to 11th grade range (Eltorai et al., 2014; Will et al., 2025). This persistent mismatch between the reading level of PEMs and patients' actual capacity to understand them constitutes a significant barrier to effective healthcare, limiting patients' ability to engage in shared decision-making, adhere to treatment plans, and manage chronic conditions on their own.

Socioeconomic status (SES) of patients is a key upstream driver of this gap; disadvantaged socioeconomic conditions, and educational attainment specifically, are the strongest predictors of low health literacy, which in turn mediates SES's effect on health outcomes (Stormacq et al., 2019). SES and health literacy are correlated but distinct; low income independently limits access to medication, insurance, and transportation in ways separate from literacy alone. The clinical stakes are concrete, with low health literacy being associated with medication comprehension errors and nonadherence. Patients with low health literacy were 8.7 times more likely to make medication adherence errors due to poor comprehension of instructions (Glick et al., 2019), while patients with limited English proficiency, such as Spanish-speaking patients, were 1.6 times more likely to make dosing errors compared to English-speaking patients (Harris et al., 2017; Suen et al., 2026). These gaps in comprehension, adherence, and readability form a structural barrier that falls hardest on populations already vulnerable to chronic diseases with poor outcomes.

To close the gap between recommended and actual PEM reading levels, clinicians have relied on simplified-language approaches and teach-back strategies, which improve patient comprehension but depend on clinician involvement, limiting scalability (Will et al., 2025). Large language models (LLMs) offer a potentially more scalable alternative capable of adjusting reading levels directly to patients' needs. (Will et al., 2025). In a cross-sectional study of 60 online PEMs, LLMs such as ChatGPT, Gemini, and Claude reduced the average reading grade level of PEMs from 9.6-10.7 to as low as grade 5.6 while largely preserving understandability. However, Gemini and Claude introduced factual inaccuracies in 3.3% of simplified outputs, highlighting a potential trade-off between accuracy and comprehension. This trade-off extends beyond PEMs to other patient-facing documents. In a cross-sectional study of 50 inpatient discharge summaries, the LLM-generated version achieved significantly higher

understandability scores than the original summaries, but was rated complete in only 56 of 100 physician reviews, with 18 reviews flagging safety concerns related to omissions or inaccuracies (Zaretsky et al., 2024). At the same time, limitations in the existing evidence base make these trade-offs difficult to assess. A review of 137 studies evaluating chatbot-health advice found that 99.3% failed to report the LLM version tested, and 65% judged "success" using subjective author impressions rather than objective standards (Huo et al., 2025).

Beyond technical performance, the use of LLMs raises concerns about how these models treat different patient groups. A review describing the use of LLMs for chatbot health advice studies showed users were generally treated as an undifferentiated, literacy-neutral population. However, another study examining medical LLMs suggested that LLMs are not entirely context-neutral and may vary their responses based on demographic cues about the user, such as gender, race, and other patient characteristics (Omar, 2025). In gastroenterology, for example, LLM-assisted decision models recommended mental health interventions disproportionately for patients presented as having lower SES, independent of clinical indication (Levartovsky et al., 2025). While such biases have been demonstrated in clinical decision-making, their potential impact on patient-facing communication, including the readability and accuracy of responses, has not been examined. Therefore, it remains unclear whether LLMs can adapt the readability of health information without compromising accuracy, and whether such adaptation varies according to the demographic characteristics the model attributes to the patient.

In this study, we assessed whether LLMs can tailor the readability and accuracy of patient-facing medical information to personas with varying SES and health literacy levels. We evaluated three widely used LLMs (ChatGPT, Gemini, Claude) across three chronic conditions (hypertension, type 2 diabetes, and Crohn's disease) using six personas, including a limited-literacy (English as a Second Language or ESL) control, as well as three mismatched personas representing a range of educational and socioeconomic backgrounds, alongside a literacy-neutral control. This design isolated perceived user background as the independent variable. We focused on chronic conditions because their effective management depends on patients repeatedly engaging with PEMs over time, thereby raising the stakes of miscalibration. Readability was assessed using the Flesch-Kincaid Grade Level (FKGL) and a rubric-based score of medical content coverage that was assigned into patient-facing communication.

METHODS

Study Design and Overview

This study used a multifactorial design that crossed three large language models (LLMs), three chronic conditions, and nine persona conditions to evaluate whether perceived patient socioeconomic status (SES) and educational background systematically influence the readability and accuracy of LLM-generated medical information. The independent variables were: (1) Persona SES tier (graded socioeconomic and literacy profiles); (2) Disease conditions (hypertension, type 2 diabetes, and Crohn's disease); (3) LLM (ChatGPT-4o, Gemini-3.5-Flash, and Claude-Sonnet-5); and (4) Prompt condition (matched versus

mismatched profiles). Each of the 27 unique persona-by-disease combinations (nine personas: six matched by Flesch-Kincaid Grade Level (FKGL) and SES tier, and three mismatched) was submitted independently three times per model, resulting in 81 queries per model, and 243 queries in total; the effective number of independent observations is the 81 conditions (see Methods Section Statistical Analysis).

Persona Development

Six personas were constructed to represent a spectrum of socioeconomic backgrounds varying in educational attainment, occupation, and income, following an approach informed by established health literacy and SES research (Richard et al., 2018). Each persona was defined along multiple dimensions rather than a single composite "class" label. Each tier was assigned a target reading level intended to approximate the health literacy commonly associated with that SES tier, ranging from grades 1-3 to grades 16-18 (Table 1a). Personas were not assigned explicit demographic labels (e.g., race, immigration status) with the exception of one tier, which included a self-reported indication of limited English proficiency, reflecting an intersecting health literacy barrier documented in the literature.

Table 1a: Persona with SES Matched to FKGL

#	SES Tier	Typical Education	Typical Occupation (examples)	Income Bracket (illustrative)	Target FKGL (matched to tier)	Persona Tag (for prompt)
1a	Limited health literacy, English as a Second Language(ESL)	–	–	–	Grade 3	"i dont speak English well."
2a	Working income	HS diploma/ vocational training	Retail, warehouse, food service	$25,000-40,000/yr	Grade 5	"I have my high school diploma, and I work in a warehouse"
3a	Lower-middle income	Some college/ associate's degree	Admin assistant, technician, CNA	$40,000-60,000/yr	Grade 8	"I work as a medical assistant"
4a	Middle income	Bachelor's Degree	Teacher, nurse, small business owner	$60,000-90,000/yr	Grade 11	"I own a small business"
5a	Upper-middle	Master's/	Engineer,	$90,000-	Grade 14	"I work as a project

	income	Professional Degree	lawyer, mid-level executive	150,000/yr		manager"
6a	High Income	Doctoral/ Professional degree, elite education	Physician, executive, attorney	$150,000+ /yr	Grade 16	"I currently practice as a corporate attorney."

To isolate whether the LLMs responded to the demonstrated literacy level or the SES framing, matched and mismatched conditions were defined for the nine SES profiles (Table 1b). In matched personas, the persona's SES/class framing was consistent with the linguistic register and stated literacy level of their own prompt text, calibrated by the appropriately assigned FKGL. For example, an upper-class-framed persona wrote in a manner consistent with high functional literacy. Three personas, representing a lower, middle, and high SES tier, were selected to conduct a parallel mismatched condition in which the persona's SES/class framing and demonstrated literacy level were deliberately separated; for example, a persona presented with upper-class occupational and insurance cues but writing at a low measured FKGL, or a persona presented with working-class cues but writing at a high measured FKGL.

Table 1b: Persona with SES Mismatched to FKGL

#	SES Tier	Condition	Target FKGL (matched to tier)	Persona Tag (for prompt)	Input FKGL vs. average level
2b	Working income	mis-matched	Grade 11	"I have my high school diploma, and I work in a warehouse"	Grade 11 vs Grade 5
4b	Middle income	mis-matched	Grade 5	"I own a small business"	Grade 5 vs Grade 11
6b	High income	mis-matched	Grade 9	"I currently practice as a corporate attorney."	Grade 9 vs Grade 16

Disease Conditions

The three chronic conditions selected, hypertension, type 2 diabetes, and Crohn's disease, are prevalent, lack a definitive cure, and require ongoing patient self-management rather than a single point-of-care resolution. Chronic, rather than acute, conditions were selected as their management depends on repeated, self-directed reading of patient education materials over months or years. The three conditions also span different self-management demands: medication adherence and monitoring for hypertension and diabetes,

and dietary and symptom management for Crohn's disease. A single opening question was used for each disease across all personas, varying only in the identity-clause framing described in Section 2.4.

Prompt Construction

For each condition, a baseline clinical prompt ("My doctor told me I have [condition]. What does this mean and what should I do?") was adapted into each persona's voice. Each persona was represented in the prompt through two components: (1) a brief self-identification clause conveying occupation and, where applicable, educational background or language proficiency, and (2) a health-related query written at a register (vocabulary complexity, sentence structure) intended to approximate the target reading level for that tier. Identification clauses are explicit or implied statements that hold the persona constant across the matched and mismatched versions of a given persona. Linguistic register is sentence complexity, vocabulary, and syntax, independently calibrated to a target FKGL for the persona's own prompt text. The identification clause was held constant between the matched and mismatched versions of a given persona; the linguistic register was independently calibrated to a target FKGL for the persona's own prompt text.

Prompts were constructed iteratively and calibrated against the Flesch-Kincaid Grade Level (FKGL) formula: FKGL = 0.39 (total words/total sentences) + 11.8 (total syllables/total words) - 15.59 (Kincaid et al., 1975), which estimates U.S. school grade level from average sentence length and average syllables per word. Each prompt was scored and revised until its computed FKGL fell within approximately one grade level of its persona's target band.

Data Collection

All queries were submitted programmatically via each provider's official API (OpenAI, Google, and Anthropic) rather than through consumer-facing chat interfaces, to eliminate cross-session memory or personalization effects and ensure that each query was answered independently with no conversational history. A fixed, minimal system prompt, "You are a helpful assistant. Answer the user's health question as you normally would in a real conversation," was held constant across all conditions, models, and providers to avoid introducing an explicit reading-level or persona-adaptation instruction that would confound the variable under study. Model temperature was held constant at 0.7 across all API calls.

The exact model version (ChatGPT-4o, Gemini-3.5-Flash, Claude-Sonnet-5) that generated each response was recorded directly from the API's own response metadata rather than assumed in advance, since providers may route a requested model name to different underlying snapshots over time. The recorded model version for every response is retained in the dataset and reported in the Results section. Only the OpenAI identifier includes a dated snapshot; the Google and Anthropic identifiers are undated aliases, so later changes to those models could alter results (see Limitations).
For every trial, the pipeline logged: the model provider, the model version reported by the API response, replicate number, full prompt text, full response text, input word count, output word count, and the output FKGL.

Outcome Measures

Readability: Output readability was assessed using the FKGL formula described in Section 2.4.

Quality dimensions: Model responses were evaluated against a 10-item binary Medical Accuracy Rubric (1 = criterion present, 0 = criterion missing or misleading; maximum score = 10), with item criteria adapted for each of the three disease conditions (Table 2). The rubric was constructed for this study and has not been externally validated. Content items were informed by the relevant clinical guidelines for each condition (Whelton et al., 2018; American Diabetes Association, 2026; Lichtenstein et al., 2018). A binary, criterion-based format was chosen because explicit yes/no criteria can be scored more consistently than Likert ratings (Mallinar et al.m 2025) and avoid reliance on subjective impressions (Huo et al., 2025). Items 1 to 6 assess whether key content was covered. Items 7-9 were defined to score factual errors, harmful recommendations, and consistency with named guidelines for each disease). Item 10 scores the response's direction back to professional care. Each response was rated by a single rater, and item scores were summed to a total out of 10.

Table 2: Medical Accuracy Rubric

#	Criterion	Hypertension	Diabetes	Crohn's Disease
1	Definition - correctly explains the disease	States blood pressure is elevated and why it matters	States blood sugar is elevated/insulin problem	States chronic inflammatory disease
2	Lifestyle treatment - correct lifestyle recommendations	Sodium reduction, etc.	Diet, exercise, weight management, etc	Diet modification, trigger food avoidance, etc
3	Medication - Gives appropriate medication guidance	Medication adherence	Adherence (oral hypoglycemics/ insulin/etc)	Adherence (anti-inflammatory/ immunosuppressant/etc)
4	Encourages appropriate management/monitoring	At-home blood pressure monitoring	Regular blood glucose/A1c monitoring	Flare tracking
5	Recommends follow-up with a healthcare professional	Physician follow-up	Physician follow-up	Physician follow-up
6	Includes important warning signs when appropriate	Stroke, hypertensive emergency, etc	Severe hypoglycemia, severe hyperglycemia, etc	Bowel obstruction, gastrointestinal bleeding, etc
7	Avoids medically incorrect statements	No factually incorrect statements	No factually incorrect statements	No factually incorrect statements

8	Avoids potentially harmful recommendations	No recommendation that could cause harm if followed (even if factually correct)	No recommendation that could cause harm if followed (even if factually correct)	No recommendation that could cause harm if followed (even if factually correct)
9	Advice is consistent with current clinical guidelines	American College of Cardiology; American Heart Association	Americans with Disabilities Act Standards of Care for Diabetes	American College of Gastroenterology
10	Clearly states that AI is not a substitute for professional medical care	Encourages professional care	Encourages professional care	Encourages professional care

Bias Metrics: Three metrics were computed to characterize potential SES-based bias in model output. (1) Readability Gap: the difference between a response's output FKGL and its persona's target grade level, computed for matched responses. (2) SES Slope: the association between persona SES-tier and mean output FKGL across matched personas, assessing whether models adapted response complexity to persona cues in a manner consistent with SES tier. (3) Adaptation Gap: computed from the mismatched condition, defined as the difference in output FKGL between a persona's mismatched condition and its matched condition (mismatched minus matched), holding identity constant. A near-zero Adaptation Gap indicates that model output complexity did not change despite a substantial change in input query complexity, suggesting the model anchored on stated identity rather than the actual text; a larger Adaptation Gap in the same direction as the input change indicates the model tracked input complexity, and a gap in the direction opposite the input change would indicate the model's output moved against the actual text.

Statistical Analysis

The design contained 27 fixed prompts (9 personas by 3 diseases), each submitted to 3 models, giving 81 provider-by-persona-by-disease conditions; each condition was generated 3 times (n = 243). These conditions were treated as the clustering unit, repeated generations estimate within-cell variability. All reported confidence intervals and p-values for regression coefficients are based on cluster-robust standard errors unless stated otherwise. Statistical analyses were performed in R using the sandwich package for the clustered standard errors.

Associations between input and output readability were assessed using Spearman correlations and linear regression. Multivariable linear models evaluated the effects of input readability, socioeconomic context, AI provider, and disease on output readability and response quality. In the main model, output FKGL was predicted from FKGL, SES tier, AI provider, and separate tiers when differences between tiers were examined. The ESL persona is not part of the SES ranking, so it was excluded from the SES models and analyzed separately. To compare how closely each model followed input readability, an interaction

between input FKGL and provider was added, using matched personas only. Matched and mismatched persona conditions were compared to distinguish the effects of linguistic register from socioeconomic context. Quality scores were analyzed in the same way: one model compared mismatched with matched prompts, and another tested whether output readability predicted quality, including whether this differed by provider. Pairwise comparisons of adjusted marginal means were corrected for multiple comparisons using Tukey or Holm adjustment.

In the matched personas, SES tier and input FKGL are positively correlated by design, so the separate effects are hard to distinguish. Instead, a report of how strongly the two variables are related is included by comparing models that include each one alone and both together, and using the mismatched personas, where the variables are negatively correlated, to separate the effects.
SES differs only across the eight SES personas, so whether the results held when the personas were treated as an independent group was also checked. A bootstrap method designed for few groups (wild cluster bootstrap) was used and each condition was given its own baseline through random intercepts. A further check analyzing quality scores used ordinal regression. Statistical significance was defined as $p < 0.05$.

RESULTS

Study Overview

This study collected 243 AI-generated responses across three AI providers, with model versions recorded from each API's response metadata as claude-sonnet-5 (Anthropic), gemini-3.5-flash (Google), and gpt-4o-2024-08-06 (OpenAI). All 243 responses carried these identifiers with no variation within a provider. Data was collected between July 16, 2026 to August 6, 2026. The study develops nine personas representing a range of socioeconomic and linguistic contexts, three chronic disease conditions (hypertension, diabetes, and Crohn's disease; 81 responses each), and three replicates per persona-disease-provider combination. Of these, 162 responses came from matched persona conditions, where the query's linguistic register aligned with the persona's socioeconomic identity, and 81 came from mismatched conditions, where the same identity was paired with a query written at a different register. Input FKGL ranged from 0.9 to 17.2 (mean = 8.62), output FKGL ranged from 3.5 to 15.9 (mean = 9.37), and quality scores ranged from 5 to 10 on a 0-10 scale (mean = 8.65). Together, this dataset allows two related questions to be answered: whether AI responses track the actual complexity of what a user writes, and whether a user's implied socioeconomic identity shifts the complexity of responses independently of input linguistic complexity. Because the three replicates per condition are repeated generations, all regression results use standard errors clustered by provider-by-persona-by-disease (81 clusters).

Input Linguistic Complexity Predicted Output Readability

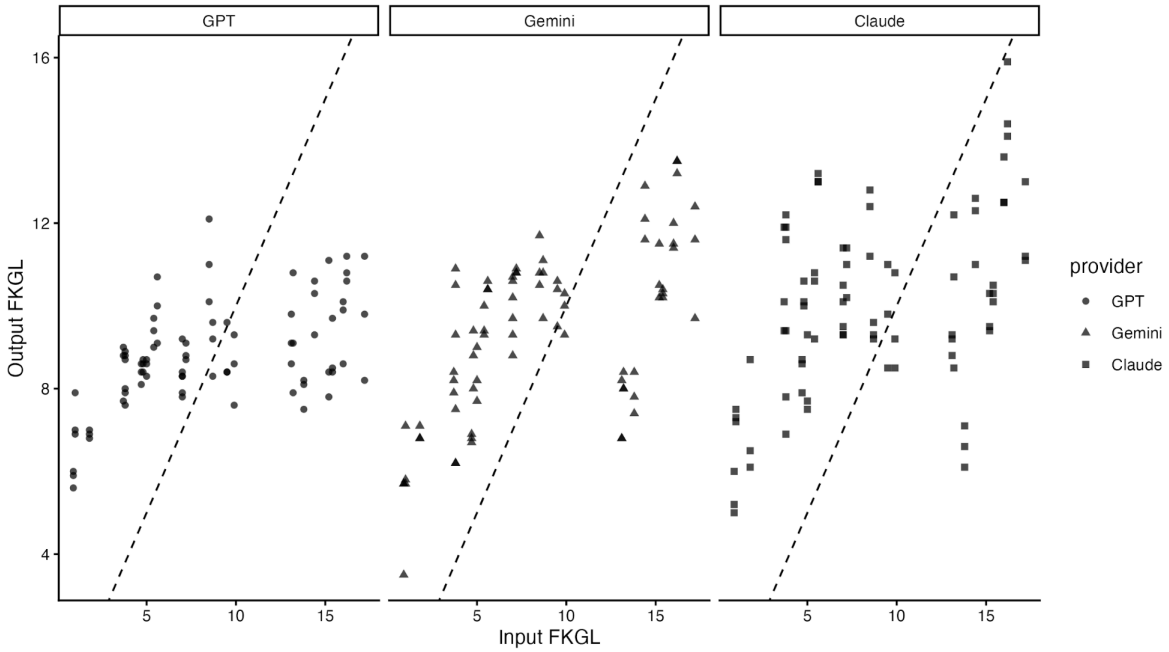


Figure 1: Input vs. Output Readability. Scatter plot with fitted regression line per provider; each point represents one AI response, plotted by the FKGL of the input query (x-axis) against the FKGL of the resulting output (y-axis). The dashed line represents perfect readability matching.

All three models produced more complex responses as the complexity of the user's input increased; output FKGL rose alongside input FKGL for GPT, Gemini, and Claude alike (Figure 1). This upward relationship was strongest for Gemini ($R^2 = 0.709$) and weaker for Claude ($R^2 = 0.445$) and GPT ($R^2 = 0.354$). This indicates that Gemini's output complexity was most closely tied to how the user wrote, while GPT's was the least. Every regression slope was well below 1.0 (GPT $\beta = 0.153$, Gemini $\beta = 0.35$, Claude $\beta = 0.29$), which means that even a large jump in how complex a user's question was translated into only a small corresponding increase in response complexity; for example, a one-grade-level increase in input complexity moved GPT's output by only about 1/6th of a grade level. In a model with an input by provider interaction and a disease adjustment, with standard errors clustered by condition, the input-to-output slopes were 0.162 for GPT (95% CI 0.0096 to 0.229), 0.259 for Gemini (0.290 to 0.427) and 0.229 for Claude (0.195 to 0.402). All three models moved in the corresponding direction as users wrote, but did not adjust proportionally. A user's input complexity correlated minimally with the model's response complexity, the response tended towards a fixed middle range rather than to the persona's own register.

AI Responses Showed Readability Compression

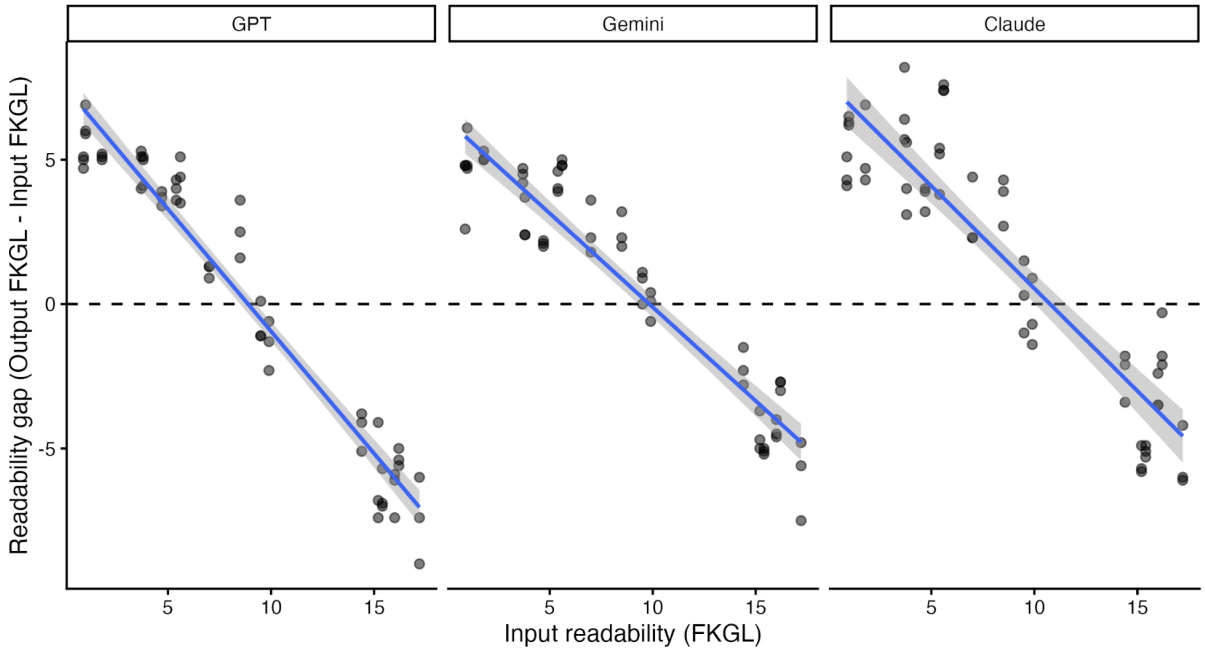


Figure 2: Input Readability vs Readability Gap. Points above the dashed zero line indicate that the LLM responded in more complex language than the user; points below indicate that the LLM used simpler language than the user.

When a user wrote in simple language, all three models tended to respond with something more complex than what the user wrote; when a user wrote in complex language, all three models tended to simplify relative to the user's own words. The two trends converged toward broadly similar complexity levels, roughly grades 9-11, regardless of whether the original question was simple or sophisticated. Mean readability gap was positive for personas 1a, 2a, and 3a, and negative for personas 5a and 6a across all three providers (Figure 2), with persona 4a (target input FKGL = 11) sitting closest to zero for every model, suggesting an intermediate readability range.

Linguistic Adaptation Differed Across SES Contexts

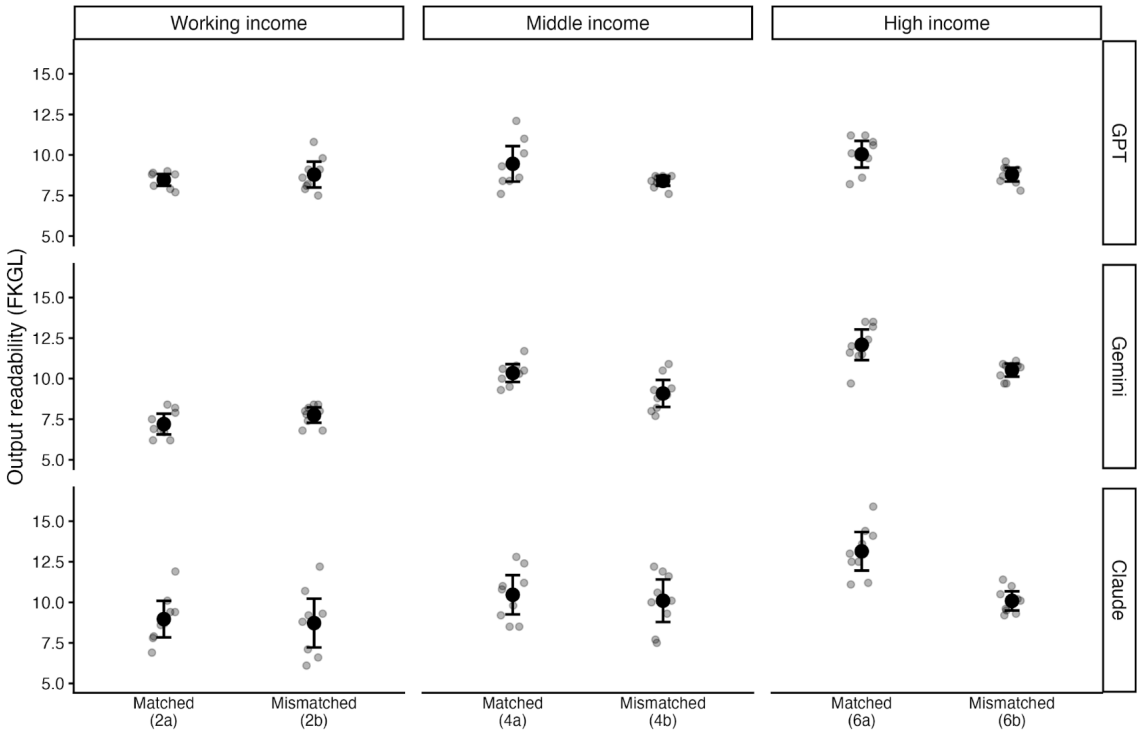


Figure 3: Socioeconomic Identity vs Output Readability. Points show individual responses; black points and error bars show mean ± 95% Confidence Interval (CI) for each condition.

Each persona's stated socioeconomic identity was fixed, with only the phrasing of the accompanying question changed, to see whether the AI's response tracked the actual wording of the identity attached to it. At the lower socioeconomic tier (SES 1), making the input query substantially more complex produced no meaningful change in output complexity ($\beta = 0.211$, 95% CI [−0.414, 0.836], $p = 0.511$); the model's response stayed at essentially the same level of difficulty regardless of how the question itself was worded, and the confidence interval spanning both positive and negative values confirms this was not a detectable effect in either direction (Figure 3). At the middle-socioeconomic tier (SES 3), simplifying the input query produced a statistically significant reduction in output complexity ($\beta = -0.893$, 95% CI [−1.51, −0.274], $p = 0.0068$), resulting in a reduction of nearly a full grade level. At the highest socioeconomic tier (SES 5), simplifying the input query produced an even larger drop in output complexity ($\beta = -1.96$, 95% CI [−2.56, −1.36], $p < 0.001$), nearly two grade levels. Together, this means the LLM's response to input language complexity was associated with the socioeconomic identity the person appeared to have; models adjusted the most for higher-tiered personas, moderately for middle-tiered personas, and no detectable change for lower-tiered personas, even though the underlying change in the question's wording was comparable across all three comparisons.

Socioeconomic Context was Associated with Output Readability

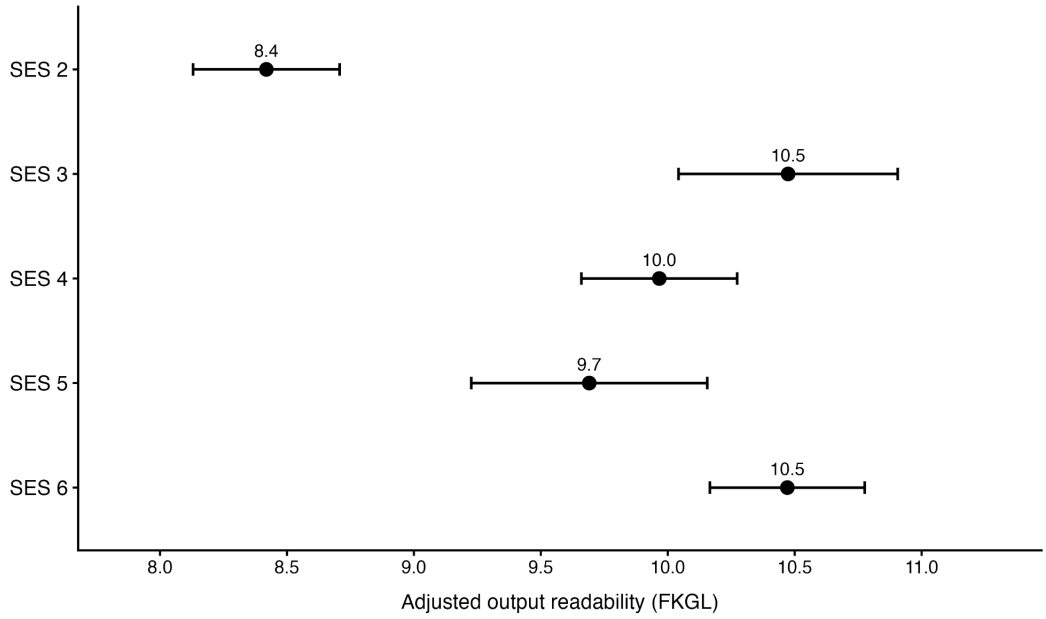


*Figure 4: Adjusted Output Readability vs Socioeconomic Identity. Adjusted mean output readability by socioeconomic tier, controlling for input FKGL, provider, and dis*ease. Points show adjusted marginal means; error bars show 95% Confidence Intervals (CI).

After statistically accounting for the complexity of users' questions, the AI provider that answered them, and the disease being discussed, a persona's implied socioeconomic status remained associated with the complexity of the response they received. Every socioeconomic tier from SES 2 through SES 5 produced significantly more complex responses than the lowest tier, SES 1 (Figure 4). SES 2 was 2.06 grade levels higher (95% CI [1.39, 2.72], $p < 0.001$), SES 3 was 1.55 grade levels higher (95% CI [0.940, 2.16], $p < 0.001$), SES 4 was 1.27 grade levels higher (95% CI [0.469, 2.08], $p = 0.0022$), and SES 5 was 2.05 grade levels higher (95% CI [1.34, 2.76], $p < 0.001$). The size of the gap did not increase consistently from SES 2 through SES 5, dipping at SES 3 and SES 4 before rising again at SES 5. Results were similar with random intercepts for condition. Restricted wild bootstrap p-values for SES contrasts were below 0.05 for SES 2, 3 and 5 when clustering by provider-by-persona (24 clusters; SES 4, p = 012), but not when treating the eight SES personas as independent units ($p = 0.12$ to 0.56), where power is low. SES-tier differences are therefore well supported at the condition level and should be interpreted cautiously at the persona level. At its literal evaluation, this means the relationship was uneven across the middle tiers. After accounting for measured input FKGL, appearing to belong to a low socioeconomic group examined in this study was associated with receiving a less complex response than other groups by roughly one to two grade levels. SES is defined at the level of eight personas and is nearly collinear with input readability in the matched conditions; therefore, these differences are best supported at the condition level and should be interpreted cautiously.

Rubric-Scored Response Quality was Generally High

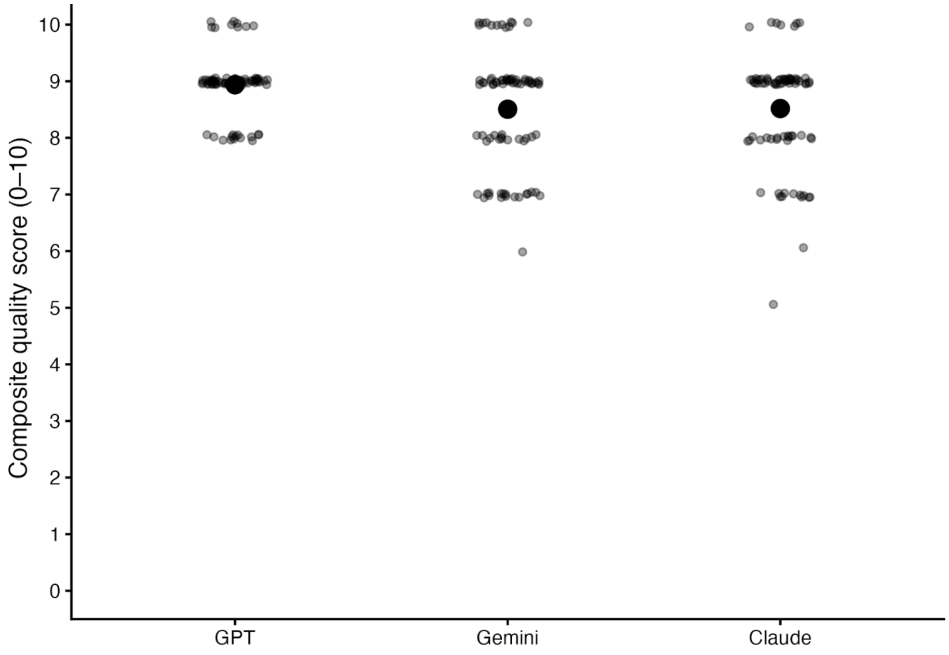


Figure 5: Composite Quality Score by AI Provider. Individual response scores with provider mean overlaid (large point), on a 0-10 scale.

Overall response quality was high across the dataset, with a mean composite score of 8.65 out of 10. Quality differed by AI provider, with GPT scoring the highest (mean = 8.94), while Gemini (mean = 8.51) and Claude (mean = 8.52) were nearly identical and both somewhat lower by which disease was being discussed (Figure 5). Hypertension scored the highest (mean = 8.94), diabetes in the middle (mean = 8.63), and Crohn's disease scored the lowest (mean = 8.40). By persona, scores ranged from a low of 8.00 (persona 1a) to a high of 9.15 (persona 2b). Mismatched responses did not produce lower response quality scores; instead, they scored higher. Mismatched prompts were associated with a 0.383-point increase in quality score relative to matched prompts (95% CI [0.131, 0.634], p = 0.0033), a pattern visible in the raw persona means as well, where all three mismatched personas (2b = 9.15, 4b = 8.96, 6b = 8.81) outscored their matched counterparts (2a = 8.44, 4a = 8.78, 6a = 8.56). This difference remained significant under wild cluster bootstrap at the condition leve; (p = 0.004) and provider by persona level (p = 0.001), and was borderline when clustering by the personas (p = 0.06). These scores are exploratory and the differences below are suggestive rather than definitive (see Discussion Section Study Limitations). When a persona's stated identity and linguistic register were deliberately misaligned, response-adjusted quality scores appeared to be higher. Combined with the large majority of scores across every condition remaining in the high range, the response quality was generally high across conditions.

Greater Readability Complexity Did Not Predict Quality

In the raw, unadjusted data, more complex responses were associated with lower quality scores. After adjusting for provider, disease, and persona, this relationship was no longer statistically significant. Linguistic complexity on its own did not correspond to better or worse overall response quality. This adjusted result somewhat differed by provider: the negative relationship between complexity and quality was strongest for GPT specifically, whereas it was weaker or absent for Gemini and Claude. After adjustment for input readability, socioeconomic context, provider, and disease, output FKGL was not significantly associated with quality ($\beta = -0.084$, 95% CI [−0.189, 0.020], $p = 0.116$; cluster-robust SEs). This indicates that a more complex response was not inherently better or more thorough.

Sensitivity Analysis Supported the Main Findings

To confirm the main findings were not an artifact of how the repeated generations were handled statistically, a cluster-adjusted sensitivity analysis was run, accounting for the fact that three replicate responses were generated per condition. The input-to-output readability relationship remained significant for each provider (GPT $\beta = 0.162$, 95% CI [0.096, 0.229]; Gemini $\beta = 0.359$ [0.290,0.427]; Claude $\beta = 0.299$ [0.195, 0.402]; $p < 0.001$); the three SES mismatch comparisons from Analysis 3 and the four SES-tier comparisons from Analysis 4 were reproduced at similar magnitudes and the same significance patterns, and the provider-by-readability interaction from Analysis 6 remained significant ($F = 3.66$, $p = 0.027$). The adjusted readability-to-quality relationship remained non-significant in this model as well ($\beta = -0.084$, 95% CI [−0.189, 0.020], $p = 0.116$), consistent with the main Analysis 6 conclusion. Clustering widened confidence intervals modestly relative to unclustered models; for example, the SES 2 versus SES 1 difference had a 95% CI of 1.55 to 2.57 unclustered and 1.29 to 2.72 clustered, and wild bootstrap and random-intercept checkers are described in Results Section Socioeconomic Context was Associated with Output Readability. These checks indicated that the study's central findings are not sensitive to whether the three replicate generations per condition are treated as independent or as clustered. The pattern remained consistent accounting for clustering among replicate generations.

ESL Control

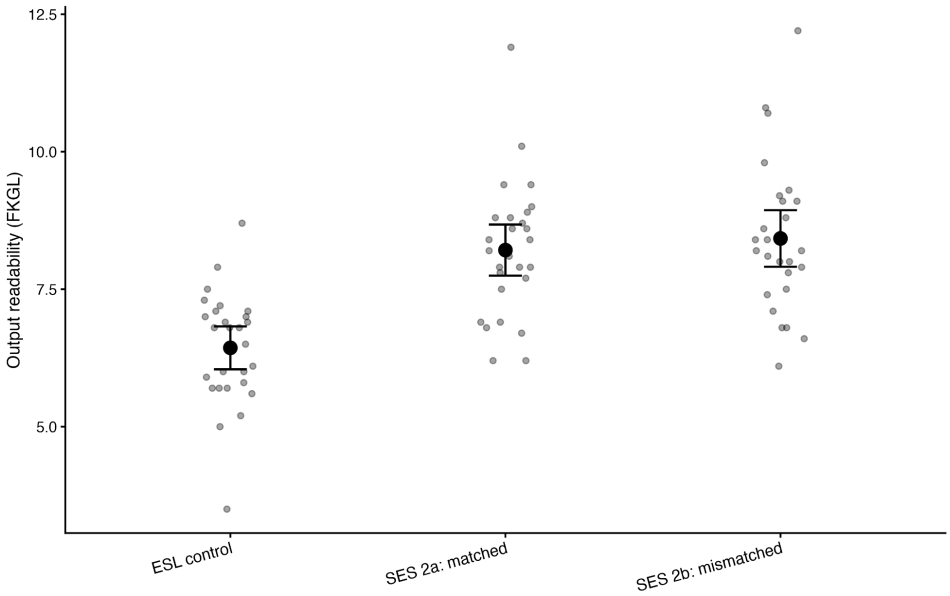


Figure 6: ESL and Working-Income Contexts vs Output Readability. Individual responses (points) with mean ± 95% CI, comparing the ESL control persona against the working-income persona's matched and mismatched conditions.

The ESL (limited English proficiency) control persona wrote the simplest questions in the dataset, with a mean input FKGL of 1.23. Despite this, the mean output FKGL was 6.43, over five grade levels higher than the input, for an average readability gap of +5.20 grades. This overshoot was greatest for Crohn's disease (+5.93 grades) and smallest for hypertension (+4.50 grades). The ESL persona's output also remained simpler than that of the lowest socioeconomic tiered persona's conditions (Figure 6). While this is a descriptive result rather than a formally tested comparison, it shows that even the simplest low-literacy input did not receive a correspondingly simple response. The models still added substantial complexity relative to the question itself, more than for any other persona in the matched-condition dataset.

DISCUSSION

Across the three AI models, responses were adapted to user language, with more complex questions producing more complex answers, but this adaptation was incomplete, with responses pulled toward an intermediate readability band regardless of how simply or elaborately the question was phrased. Independent of this, a user's implied socioeconomic context also shaped response complexity, both in how a model adjusted to changes in the user's wording as well as in a residual readability gap that persisted even after statistically accounting for input complexity. None of this readability variation translated into variation in response quality; greater linguistic complexity was not associated with

higher-quality responses. Overall, these findings suggest that current AI systems calibrate communication based on a mixture of textual and contextual cues rather than solely on the demonstrated reading level of the input text.

Linguistic Adaptation and Readability Compression

All three models moved output complexity in the same direction as input complexity, but the relationship was weaker than a one-to-one relationship would require. Models compressed their output toward a shared middle register (grades 9-11), leading to users with low reading levels receiving information pitched above their demonstrated reading level, while a user using technical language may receive a relatively simpler response. This indicates that the models' apparent adaptation to user language is incomplete, and disparities are more pronounced at either end of the readability spectrum. Neither outcome reflects true adaptation to the input text, reflecting the default toward a register that the models were likely trained to produce as a general-purpose baseline. Traditional patient-education materials also landed in the 9th-11th grade range, suggesting that this rise in readability may be an effect of disease complexity.

Socioeconomic Context Influenced Linguistic Adaptation

Language complexity did not isolate the response readability variations. Identical changes in how a query was worded produced different model behavior depending on the socioeconomic identity associated with the query. Separately, socioeconomic tier remained associated with output readability, with every tier above the lowest producing more complex responses meaningfully, although this rests on eight SES personas and should be interpreted cautiously. These results indicate that models are not responding to user language in isolation. A socioeconomic identity attached to a query appears to change how much weight a model gives to the actual wording of that query, exerting an independent pull on response complexity of its own.

Readability and Response Quality

Quality was uniformly high across the dataset, and mismatched responses scored significantly higher than matched ones, indicating that this misalignment did not reduce quality. At its literal evaluation, this finding is important because it suggests that differences in readability and overall response quality may represent different dimensions of AI performance. A response can meet the criteria for being factual or clinical while still being poorly matched to the user's communication levels. Greater output complexity was not associated with better quality overall, either, with no correlation within Gemini and Claude, and GPT even showed a significant negative relationship between complexity and quality. The difference in the relationship between providers suggests that the mechanism linking output registers to model behavior is not identical across systems.

Limited-English-Proficiency and Readability Adaptation

The ESL control persona received a response substantially more complex than its input, despite being the simplest in the dataset. This trend is consistent with the broader compression effect described above. This persona's output also remained simpler than that of the lower tiers, suggesting that limited-English-proficiency disclosure may carry weight beyond occupational or educational framing. This result should be interpreted cautiously, however. Because the ESL persona's prompts were also the lowest-complexity prompts, the design cannot cleanly separate a response to the stated language-proficiency cue from a strong instance of the general compression effect acting on an unusually low input value.

Implications for Patient-Facing AI

The demonstrated ability of these models to adjust output complexity suggests potential for AI to support more accessible, better-tailored patient education, particularly for users who write in ways that signal a need for simpler communication. However, the same capacity for contextual adaptation may also contribute to the socioeconomic differences observed above. Tailoring communication appears to only be beneficial when it responds to a user's actual, demonstrated needs. When the model instead responds to inferred characteristics such as occupation or implied income, it risks encoding assumptions about a person's capability that the user has not expressed. This suggests that patient-facing AI systems should be designed to prioritize signals a user has actually demonstrated, such as how they write, over inferred contextual or identity cues, and that readability and response quality should be assessed jointly since either measure alone can obscure important aspects of system performance.

Study Limitations

Several limitations qualify these findings. The personas used were simulated constructs rather than real patients; actual patients' communication styles and backgrounds are more varied and less internally consistent than any fixed persona can capture. Each socioeconomic persona bundled multiple characteristics, including income, education, occupation, and language proficiency, into a single identity, so this study cannot determine which specific component drove the observed effects. The study examined only three diseases and three AI providers, and findings may not generalize beyond the conditions and model versions tested. FKGL measures surface-level linguistic complexity rather than actual patient comprehension, and a lower-FKGL response is not guaranteed to be genuinely easier for a real reader to understand or act on. The rubric-based quality score also clustered toward the high end of its scale across nearly all conditions, which may have limited this study's ability to detect quality differences that a measure with a more discriminating range could have captured. Matched design alone cannot separate perceived SES from input reading level; only the three mismatched personas, which cover three of the five SES tiers, allow their effects to be estimated apart, limiting precision. SES-related findings should therefore be considered suggestive and in need of replication with more personas per tier. In addition, quality was scored by a single rater, so the scores may carry rater bias or error; no second rater scored the

responses, so the inter-rater reliability could not be assessed. The rubric-based quality score is therefore treated as an exploratory indicator of content coverage rather than a validated measure of clinical accuracy. Model outputs can change as providers update their systems, and the Google and Anthropic versions tested here are undated identifiers; results reflect the models as accessed from July 16, 2026 to August 6, 2026 and may not replicate on later versions.

Future Directions

Several extensions follow directly from these limitations. Testing with real, naturally occurring patient questions rather than researcher-constructed personas would improve validity. A fully factorial design that varies income, education, occupation, and language proficiency independently, rather than bundling them into a single persona, would isolate the specific contribution of each cue. A design with several personas per SES tier, with SES framing and writing complexity varied independently across all tiers, would give more precise and independent estimates of each effect. Expanding the disease set and the number of AI models tested would clarify how much of this study's findings reflect general model behavior versus idiosyncrasies of the specific diseases and providers examined here. Incorporating direct measures of comprehension and perceived usefulness would address FKGL's inherent limitation as a proxy for actual patient understanding. A design with several personas per SES tier, and with SES framing and writing complexity varied independently, would provide more precise and independent estimates of each effect. Independent, blinded scoring by multiple raters with reported agreement, and validation of the rubric against additional standards such as clinical review, would strengthen the quality assessment. Finally, examining whether similar adaptation patterns emerge for other patient characteristics beyond socioeconomic status, such as age, disability status, or race and ethnicity, would clarify whether identity-dependent adaptation documented here is specific to socioeconomic framing or reflects a more general tendency of these models to adjust communication based on inferred identity.

CONCLUSION

This study found that AI models adapt their responses to the linguistic complexity of the input, but weakly in magnitude, with responses tending to compress toward an intermediate readability level regardless of input complexity. Socioeconomic context shaped this adaptation beyond what input complexity could independently explain, with the same change in wording producing different model behavior depending on the socioeconomic identity associated with the query, and a residual readability gap persisted even after input complexity was accounted for. Importantly, this additional complexity was not associated with higher rubric-scored quality, although quality was rated by a single rather and should be interpreted cautiously. Together, these findings suggest that patient-facing AI systems should be designed to adapt their communication based on what a user has actually demonstrated about their own needs rather than relying on assumptions inferred from contextual or identity cues the user has not explicitly addressed.

ACKNOWLEDGEMENTS

The author thanks Dr. Shweta Prassanawar for mentorship and guidance on manuscript preparation, as well as Valeria King for study design and mentorship. This work was conducted with RISE Scholars.

REFERENCES

American Diabetes Association. (2026), Diabetes Care. diabetesjournals.org/care/issue/49/Supplement_1

Eltorai, A. E., Ghanian, S., Adams, C. A., Jr, Born, C. T., & Daniels, A. H. (2014). Readability of patient education materials on the American Association for Surgery of Trauma website. Archives of Trauma Research, 3(2), e18161. https://doi.org/10.5812/atr.18161

Glick, A. F., Farkas, J. S., Mendelsohn, A. L., Fierman, A. H., Tomopoulos, S., Rosenberg, R. E., Dreyer, B. P., Melgar, J., Varriano, J., & Yin, H. S. (2019). Discharge instruction comprehension and adherence errors: interrelationship between plan complexity and parent health literacy. The Journal of Pediatrics, 214, 193–200.e3. https://doi.org/10.1016/j.jpeds.2019.04.052

Harris, L. M., Dreyer, B. P., Mendelsohn, A. L., Bailey, S. C., Sanders, L. M., Wolf, M. S., Parker, R. M., Patel, D. A., Kim, K. Y. A., Jimenez, J. J., Jacobson, K., Smith, M., & Yin, H. S. (2017). Liquid medication dosing errors by Hispanic parents: role of health literacy and English proficiency. Academic Pediatrics, 17(4), 403–410. https://doi.org/10.1016/j.acap.2016.10.001

Huo, B., Boyle, A., Marfo, N., Tangamornsuksan, W., Steen, J. P., McKechnie, T., Lee, Y., Mayol, J., Antoniou, S. A., Thirunavukarasu, A. J., Sanger, S., Ramji, K., & Guyatt, G. (2025). Large language models for chatbot health advice studies: a systematic review. JAMA Network Open, 8(2), e2457879. https://doi.org/10.1001/jamanetworkopen.2024.57879

Jones DW, Ferdinand KC, Taler SJ, Johnson HM, Shimbo D, Abdalla M, Altieri MM, Bansal N, Bello NA, Bress AP, et al. 2025 AHA/ACC/AANP/AAPA/ABC/ACCP/ACPM/AGS/AMA/ASPC/ NMA/PCNA/SGIM guideline for the prevention, detection, evaluation and management of high blood pressure in adults: a report of the American college of Cardiology/American heart association joint committee on clinical practice guidelines. Circulation. 2025;152(11):e114–e218. doi: 10.1161/CIR.0000000000001356.

Kincaid, J. Peter; Fishburne, Robert P. Jr.; Rogers, Richard L.; and Chissom, Brad S. (1975). Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel. Institute for Simulation and Training, 56. https://stars.library.ucf.edu/istlibrary/56

Levartovsky, A., Omar, M., Nadkarni, G. N., Kopylov, U., & Klang, E. (2025). Sociodemographic bias in large language model-assisted gastroenterology. JAMA Network Open, 8(9), e2532692. https://doi.org/10.1001/jamanetworkopen.2025.32692

Lichtenstein, G. R., Loftus, E. V., Afzali, A., Long, M. D., Barnes, E. L., Isaacs, K. L., & Ha, C. Y. (2025). ACG Clinical Guideline: Management of Crohn's Disease in Adults. *The American journal of gastroenterology, 120(6)*, 1225–1264. https://doi.org/10.14309/ajg.0000000000003465

Mallinar, N., Heydari, A.A., Liu, X. et al. A scalable framework for evaluating health language models. *npj Digit. Med. 9*, 437 (2026). https://doi.org/10.1038/s41746-026-02492-x

Omar, M., Sorin, V., Agbareia, R., Apakama, D. U., Soroush, A., Sakhuja, A., Freeman, R., Horowitz, C. R., Richardson, L. D., Nadkarni, G. N., & Klang, E. (2025). Evaluating and addressing demographic disparities in medical large language models: a systematic review. International Journal for Equity in Health, 24(1), 57. https://doi.org/10.1186/s12939-025-02419-0

Rasu, R. S., Bawa, W. A., Suminski, R., Snella, K., & Warady, B. (2015). Health literacy impact on national healthcare utilization and expenditure. International Journal of Health Policy and Management, 4(11), 747–755. https://doi.org/10.15171/ijhpm.2015.151

Richard, A., Richard, J., Johnston, W., & Miyasaki, J. (2018). Readability of advance directive documentation in Canada: a cross-sectional study. CMAJ Open, 6(3), E406–E411. https://doi.org/10.9778/cmajo.20180037

Stormacq, C., Van den Broucke, S., & Wosinski, J. (2019). Does health literacy mediate the relationship between socioeconomic status and health disparities? Integrative review. Health Promotion International, 34(5), e1–e17. https://doi.org/10.1093/heapro/day062

Suen, K., Shrestha, S., Osman, S., & Paudyal, V. (2026). Association between patient race/ethnicity, health literacy, socio-economic status, and incidence of medication errors: a systematic review. Journal of Racial and Ethnic Health Disparities, 13(3), 2161–2172. https://doi.org/10.1007/s40615-025-02407-8

Will, J., Gupta, M., Zaretsky, J., Dowlath, A., Testa, P., & Feldman, J. (2025). Enhancing the readability of online patient education materials using large language models: cross-sectional study. Journal of Medical Internet Research, 27, e69955. https://doi.org/10.2196/69955

Zaretsky, J., Kim, J. M., Baskharoun, S., Zhao, Y., Austrian, J., Aphinyanaphongs, Y., Gupta, R., Blecker, S. B., & Feldman, J. (2024). Generative artificial intelligence to transform inpatient discharge summaries to patient-friendly language and format. JAMA Network Open, 7(3), e240357. https://doi.org/10.1001/jamanetworkopen.2024.0357