

How AI Listens to Neurodivergence - A Comparative Study of Advice Responses Generated by Four Large Language Models

Jacob Gao
jacobgao031@gmail.com

ABSTRACT

Large language models (LLMs) are increasingly being used to provide advice on a wide range of areas in life, yet little is known about how their advice changes when users supply neurodivergence-related context. Drawing on the author's positionality as a student with autistic experience, this study examines a fixed body of 180 responses generated with ChatGPT, Gemini, Claude, and DeepSeek. to 15 adolescent-oriented advice questions under three conditions: 1) no additional context, 2) a broad neurodivergent label, and 3) specific executive-function and sensory-overload difficulties. Archived lexical indicators and close reading show that the broad neurodivergent label primarily changed tone and explanation, while framing with specific difficulties produced more concrete support strategies, with increased scripts and more concrete support strategies for task scaffolds, sensory adjustments, routines and accommodations. Variations between AI systems also appeared; however, responses spanning the data body were not generated under identical conditions or under a singular protocol, so these contrasts are descriptive and observational rather than controlled estimates of model performance. Overall, the findings suggest that specific prompts make access needs more visible than identity labels alone, potentially helping neurodivergent users obtain more targeted advice from generative AI.

Keywords: Neurodivergence; AI advice; large language models; student support; contextual framing; prompt design

INTRODUCTION

Students and young people now encounter large language models (LLMs) not only as writing and tutoring tools but also as sources of advice about motivation, communication, friendship, and daily routines. For neurodivergent users, common advice-seeking questions may involve executive-function friction, sensory load, social interpretation, task initiation, or difficulty communicating support needs. An answer that treats every concern as a problem of effort or discipline may therefore miss relevant access barriers. At

the same time, an answer that relies too heavily on a broad identity label may overgeneralize the user's experience.

Existing scholarship gives reason to examine this tension. Brandsen et al. (2024) found negative associations with neurodivergence-related terms across word and sentence encoders. Carik et al. (2025) observed that LLMs produced overly neurotypical responses when advising neurodivergent people seeking advice. These studies suggest that AI systems may offer practical support while reproducing assumptions embedded in dominant language and social norms.

This paper asks how a fixed corpus of AI advice changes when the same questions are paired with different levels of neurodivergence-related context. Here, “listening” refers to observable changes in tone, explanation, or support strategy after context is supplied. The study does not test clinical accuracy, diagnose the question author, or evaluate the effectiveness of the advice. It addresses three questions:

- 1) How is contextual framing linked to changes in tone, explanation, and support strategy?
- 2) How do responses framed by a broad neurodivergent label differ from responses framed by specific executive-function or sensory-overload difficulties?
- 3) What descriptive patterns appear across the four archived models, ChatGPT 5.5, Gemini 3.1 Pro, Claude Sonnet and DeepSeek V4, and what limits apply to their interpretation as model differences?

The remainder of this paper is organized as follows: Section 2 reviews the existing literature on neurodivergent students, generative AI, AI-assisted advice, and bias-aware support while establishing the research gap addressed by this study. Section 3 describes the data and methodology, including the construction of the 180-response body, the three contextual framing conditions, the four model-source corpora, the lexical indicators used in analysis, and the close-reading approach. Section 4 presents the results, first examining differences across contextual framing conditions, then considering observational differences among the four AI systems through aggregated indicators and matched response examples. Section 5 discusses the implications of these findings, including what it means for an AI system to “listen” to neurodivergence, prompt specificity as a form of self-advocacy, and limitations on the study’s findings. Section 6 concludes by summarizing the study’s findings, limitations, and implications for future research and more inclusive AI system design.

LITERATURE REVIEW

Neurodivergent Students and the Need for Specific Support

Neurodivergent students in K–12 and higher education may need support that extends beyond general encouragement. Depending on the individual and situation, relevant support may include executive-function scaffolding, attention regulation, sensory or environmental adjustment, clear

communication, predictable routines, and individualized accommodations. Appropriate support cannot be inferred from an identity label alone.

Research on how generative AI affects neurodivergent users remains comparatively new. The available literature describes promising applications, but it also repeatedly identifies limited empirical evidence, uneven access, privacy concerns, bias, and the need for human oversight.

Ronksley-Pavia et al. (2025a) identified and reviewed 21 sources published between 2020 and 2024 on generative AI applications in supporting neurodivergent students within K-12 educational contexts, where personalized learning, administrative assistance for educators, and development of individualized education plans showed potential as prominent applications. The review also emphasized inaccuracy, overreliance, privacy, inequitable access, bias, and the limited empirical evidence base. This combination of promise and uncertainty supports a cautious analysis of what AI advice contains rather than an assumption of inherent efficacy.

Velazquez-Solis et al. (2025) similarly argued that AI may support more adaptive teaching for neurodivergent students in higher education, while emphasizing the continuing role of educator judgment and human support. Ronksley-Pavia et al. (2025b) showed how detailed synthetic learner profiles can help educators rehearse personalized approaches for twice- and multi-exceptional students. Both studies center specificity: meaningful personalization depends on information about barriers, strengths, environments, and goals rather than a label alone.

Hutson (2024) broadened this discussion by treating neurodiversity as a source of valuable perspective in creating a symbiotic human-AI relationship, shifting the goal away from making neurodivergent people conform to a standardized model and toward designing systems that recognize cognitive variation and the disabling effects of inflexible environments.

Throughout these studies, there is a consistent distinction between identity recognition and need-responsive support. Naming neurodivergence may reduce blame or invite validation, but it does not establish which intervention is appropriate. Concrete information about task initiation, ambiguous communication, sensory load, sleep routines, or multi-step work gives an advice system more direct material from which to construct a response.

AI Advice, Neurodivergent Users, and Bias-Aware Support

Beyond formal education, neurodivergent individuals may use generative AI for emotional support, interpersonal communication, interpretation of social cues, task management, and everyday problem solving. Research on these uses is also relatively recent and combines accounts of practical value with concerns about misinterpretation, dependence, privacy, and narrow advice.

Carik et al. (2025) qualitatively analyzed discussions from 61 neurodivergent Reddit communities and identified 20 LLM use cases across five areas: emotional well-being, mental-health support, interpersonal

communication, learning, and professional development and productivity. Participants' discussions included concerns about overly neurotypical responses, limitations of text-only interaction, overreliance, and displacement of human relationships. Users also shared and modified prompts to make responses better fit their needs, providing direct precedent for studying contextual framing based on prior information.

Brandsen et al. (2024) examined 11 language-model encoders using word-embedding association measures and introduced a sentence-similarity ratio test. Neurodivergence-related words and sentences frequently carried negative associations; in some encoders, sentences such as “I have autism” or “I have epilepsy” were represented more negatively than comparison sentences such as “I am a bank robber.” The study did not evaluate conversational advice but demonstrated that apparently neutral language technologies could encode ableist associations.

Kong et al. (2025) studied social interactions between neurotypical participants and participants with intellectual disabilities through a diary study and interviews. The two groups differed in the interaction features they valued, while social uncertainty and differing expectations created barriers for both. Their bidirectional approach, related to the double-empathy problem, is important here because AI advice should not automatically treat the neurodivergent person as the sole party who must change.

Horvat and Horvat (2025) described possible uses of AI for interpreting coded language, preparing communications, organizing tasks, accommodating sensory preferences, and preparing for interviews and placed particular emphasis on delegated trust. Because an AI system can neither bear moral responsibility nor be held accountable for its actions, users' trust ultimately concerns the developers, institutions, and regulators behind it. Privacy and user agency therefore remain central even when advice feels supportive.

Lemke et al. (2024) provided a further example of neuronormative bias in generated images of workers with dyslexia. Their analysis suggests that generative systems can reproduce disability, gender, racial, and occupational stereotypes. Although image generation differs from advice generation, the study reaffirms the need to examine what systems treat as normal, desirable, or limiting.

Taken together, the literature suggests that generative AI may support neurodivergent users while also reproducing neuronormative assumptions and broad generalizations. Advice based mainly on labels such as autism, ADHD, or dyslexia may not address the user's actual barriers or access needs. More targeted attention to communication struggles, task initiation, information organization, sensory demands, or interpersonal conflict may produce advice that is more targeted in form without assuming that one diagnosis corresponds to one solution.

The present study addresses a narrower gap within this emerging literature: how the same set of advice questions is represented across three levels of user context. Rather than asking whether an AI response is simply supportive or biased, it distinguishes changes in tone, causal explanation, and practical support

strategy. This distinction allows the corpus to show where a broad label changes interpretation and where concrete information changes the proposed intervention.

DATA AND METHODOLOGY

The study analyzes a fixed archive of 180 responses: 15 advice questions × three framing conditions × four model-source labels. The questions concern productivity, attention, sensory preferences, sleep, family communication, friendship, and social participation. They are advice prompts rather than diagnostic cases, and no participants or clinical records were recruited or collected for this project. Because the full prompts are lengthy, they are presented in Appendix A rather than in the main text.

The AI systems were tested under three framing conditions. The first is the control, with no background information given to the system by the user. This condition shows how models answer when they receive only the advice question. Second is the neurodivergent label, where the AI system is informed that the user is neurodivergent. This condition tests whether a broad identity label changes tone or explanation. The third condition concerns specific difficulty. The AI system is informed that the user is neurodivergent and struggles with executive-function difficulties and sensory-overload difficulties. This condition tests whether concrete support needs change the advice itself.

The distinction between the second and third conditions is of particular interest; a broad identity label may invite validation or a general neurodiversity frame, whereas a description of concrete difficulties supplies information that can be translated into specific supports. Thus, the conditions are therefore treated as measuring differing degrees of contextual specificity, rather than solely as diagnostic categories.

Data was generated in Visual Studio Code using OpenAI's Codex extension. Because Codex is a product of OpenAI, a ChatGPT API Key was not necessary; for Claude, Gemini, and DeepSeek, API keys were used to incorporate the models into Codex. Codex was then prompted to generate responses en masse to the prompt corpus using all four AI models under the three framing conditions. Thus, the incorporation protocol was not identical across models.

The primary analysis compares the three context conditions across the pooled corpus. Model-source comparisons are secondary and descriptive. This priority reflects an important feature of the archive: the response sets do not preserve a fully uniform generation protocol. The archived Gemini and Claude files state that all three condition responses for a question were generated together in one call, while the DeepSeek archive records condition-specific prompt context. The ChatGPT-labeled responses are stored directly in the project script, and the surviving archive does not preserve an API model identifier or call metadata for that set. Model names are therefore used as provenance labels for the existing corpora, not as a basis for causal ranking or claims about current provider behavior. The archived scripts provide descriptive lexical indicators: word count, and occurrences of a list of indicators, including neurodivergence-related terms, professional-help terms, and sensory/environment-related terms. These are

simple string-matching counts rather than validated scales or semantic classifiers; increased use of “indicator” terms signals that a response may have more targeted content and concrete recommendations, but it is not deterministically so; thus, occurrences of “indicators” are used to locate broad patterns for close reading, not to measure response quality, safety, empathy, or real-world utility.

Close reading compares matched responses to the same question and attends to validation, agency, scripts, scaffolds, accommodations, environmental adjustment, safety language, and possible overgeneralization. The analysis is interpretive rather than a formal human-coded content analysis: no inter-rater reliability statistic is claimed. Numerical differences are consequently reported as descriptive features of this corpus, and no inferential test or population-level generalization is made.

RESULTS

Presented in this section are the description findings of the body of 180 AI responses. This section first analyses how responses change across the three framing conditions, with attention particularly focused on differences in tone, explanation, and support strategy. It then considers descriptive variation across the four AI systems, followed by matched examples from the archive of responses that highlight the identified patterns. Comparisons between AI systems are best interpreted as observation and descriptive contrasts rather than controlled measures of model performance due to the non-homogenous nature of the generation protocol.

Context Framing: From Generic Advice to Specific Support

Across the fixed body of responses, increasing contextual specificity was associated with a shift from general self-management advice toward more explicit access-oriented support. The broad neurodivergent label most clearly changed tone and interpretation; the specific-difficulties condition more clearly changed the substance of the advice.

No additional context: Baseline responses generally treated the concerns as problems of habit, confidence, communication, sleep routine, or skill building. Many were already validating and practical, and several questions explicitly involving attention or sensory overload elicited environmental advice even without an identity label. Rather than being read as uniformly unsupportive, the baseline provides a comparison for identifying what additional context made more explicit.

Broad neurodivergent label: Label-framed responses more often interpreted the concern through communication difference, attention regulation, energy management, or brain-environment fit. They were frequently less blame-oriented, but the archived suggestion-marker average fell from 3.08 at baseline to

2.47 in the label condition. In this data body, identity information changed the explanatory lens more consistently than it increased the amount of action-oriented language.

Specific difficulties: Responses supplied with executive-function and sensory-overload context contained the clearest increase in scripts, checklists, timers, quiet spaces, sensory regulation, body doubling, task scaffolds, and stepwise routines. The average suggestion-marker count rose to 3.68, while sensory/environment term occurrences increased from 77 at baseline to 103 in the label condition and 290 in the specific-difficulties condition. These lexical counts do not establish usefulness, but close reading confirms that the response form became more accommodation oriented. Table 1 presents pooled comparisons across contextual framing conditions. The data show that neurodivergence-related term occurrences rose from 11 at baseline to 76 in the label condition and 196 in the specific-difficulties condition. It should be noted that because the conditions themselves make different concepts salient, these totals partly reflect expected lexical uptake. The utility of this comparison lies in directing attention to how that uptake is translated into advice, rather than treating repetition of condition language as an independent quality measure.

The archived screening procedure also flagged limited explicit context uptake in approximately 15% of the label-condition responses and none of the specific-difficulties responses. This should not necessarily be described as model “forgetting” because the data body consists of single-turn outputs rather than longitudinal conversations, the latter of which being the type of AI conversation where “forgetting” becomes a salient issue. It indicates only that broad identity context was less consistently visible in response language than concrete difficulty information.

Model Differences: Four Styles of Listening

Descriptive variation also appeared among the four AI systems as shown in Table 2 and Figure 1. These contrasts are secondary to the condition-level comparison and should not be read as a controlled ranking of the underlying systems.

ChatGPT-labeled corpus: These responses combined validation with substantially more professional-help language: 31 term occurrences across conditions, compared with one in Gemini, four in Claude, and one in DeepSeek. The pattern was especially visible in responses about sleep, driving, food and weight, family communication, and sensory overload, though this is a feature of the body of responses rather than an estimate of current ChatGPT behavior.

Gemini-labeled corpus: These responses showed the largest sensory/environment total (143 occurrences), including a sharp increase under specific-difficulties framing. Professional-help language was rare. The resulting advice often emphasized environmental adjustment and proactive planning. This pattern was reflected in fewer mentions of professional assistance and a lower likelihood of recommending professional intervention. However, the low rate of escalation should not automatically be interpreted as safer or less medicalizing.

Claude-labeled corpus: These responses frequently translated context into scripts, routines, accommodations, and task scaffolds. Sensory/environment language totaled 123 occurrences, and professional-help language remained uncommon. Close reading suggests an inclination to autonomy preservation, but that interpretation remains descriptive and was not scored by independent raters.

DeepSeek-labeled corpus: These responses were longest on average (100.71 words) and contained little professional-help language. They also increased sensory and neurodivergence-related language as context became more specific, although less sharply than the Gemini- and Claude-labeled sets. This difference may reflect model design, prompt format, platform policy, training data, or generation procedure; the data cannot isolate cultural origin as the cause.

Evidence from Original Responses

Matched responses make the condition-level pattern more concrete because the underlying advice question remains the same while the supplied context changes. Question 11, concerning late waking and daytime fatigue, provides a useful example.

In the baseline **ChatGPT**-labeled response, the advice focused on gradual schedule correction: “Move wake time gradually, perhaps fifteen to thirty minutes earlier every few days,” followed by bright light and food after waking. Under the specific-difficulties condition, the response added a night-before setup, external cues, and sensory considerations: “clothes out, first task written, alarm across the room, bright light ready, breakfast easy ... sensory comfort matters too.” It also recommended medical screening and counselor support because fatigue was affecting obligations. The contrast illustrates a move from a general sleep routine toward sensory planning, and safety-aware escalation, and the use of external aids and scaffolds.

The other systems approached the same question with overlapping but not identical supports.

Gemini: “Let's try some targeted strategies. For executive function, set multiple alarms with different, non-jarring sounds, placing one far from your bed. Immediately upon waking, try a pre-planned, low-effort task like drinking water or opening a curtain to leverage momentum. For sensory overload, ensure your bedroom is dark and quiet for sleep, and consider a sunrise alarm clock or gentle, natural light to wake up, avoiding harsh sounds or bright lights first thing. Break down morning tasks into tiny, single steps to reduce cognitive load.”

Claude: “Break the morning routine into tiny steps to reduce executive function demands—perhaps start with just sitting up, then feet on floor. Prepare everything the night before when you have more capacity. Use external cues like automatic lights, phone alarms with specific tasks, or accountability texts. For sensory needs, consider blackout curtains for better sleep, and have comfortable clothes ready. A simple morning routine chart can reduce decision-making.”

Both excerpts break the morning routine into smaller steps and incorporate sensory or environmental adjustments. The similarity supports the condition-level finding that concrete information about concerns is often translated into scaffolding rather than merely acknowledged.

The DeepSeek-labeled response emphasized schedule fit and a low-sensory morning.

DeepSeek: “It sounds like your current schedule may not align well with your natural energy rhythms, which is common with executive-function and sensory challenges ... try shifting your obligations to later in the day when you feel more alert. If mornings are non-negotiable, consider a gradual wake-up routine: use a gentle alarm (like a sunrise lamp or vibration), and keep a low-sensory morning—dim lights, quiet, and a simple task like stretching or drinking water. Avoid screens or loud noises until you’re ready.” This excerpt likewise incorporates environmental adjustment and a gradual routine, while omitting the professional-help recommendation present in the ChatGPT-labeled response. The example illustrates a descriptive system-level contrast, but the uneven generation provenance prevents attributing that contrast to model identity alone.

The matched examples also reveal a default framing in responding to questions. Without supplied identity context, models rarely named neurodivergence proactively, but they sometimes responded to explicit sensory or attention cues with relevant supports. It is possible that non-mention may be attributed to bias, though it is not the only plausible mechanism, and avoiding unsolicited diagnostic inference can be appropriate.

DISCUSSION

What It Means for AI to "Listen" to Neurodivergence

In this paper, “listening” refers to observable context uptake rather than comprehension, empathy, or memory in a human sense. Uptake can occur at three levels. Tone concerns whether a response validates or blames the user. Explanation concerns how the response interprets the source of difficulty. Support strategy concerns what actions, environmental changes, scripts, or accommodations the response proposes. Keeping these levels separate is the paper's main analytic contribution

The most defensible interpretation is cautious and specific. A broad label can alter tone and reduce moralized explanations, but it may also invite generic identity language. Concrete descriptions of barriers more often change the support plan. This does not show that the resulting advice is effective for every user; it shows that the advice more visibly incorporates environmental and executive-function support.

The body of data therefore supports a qualified answer: AI advice can visibly respond to neurodivergence-related context, and that response is more substantive when the prompt describes barriers and access needs rather than an identity label alone.

Prompt Specificity as Self-Advocacy

For a neurodivergent user seeking advice from AI, the results suggest a practical communication principle: describe the barriers and concerns, the context surrounding the concerns, and the kind of help sought. This approach can make it easier for a system to generate a concrete plan without requiring the user to disclose a diagnosis.

Prompt specificity can therefore function as a limited form of self-advocacy. Instead of relying only on “I am neurodivergent,” a user might state, “I have difficulty starting multi-step tasks when the room is noisy; help me design a five-minute starting routine.” This wording preserves agency and identifies an access need. It should not, however, shift responsibility entirely onto the user. Inclusive systems should also respond well to ordinary descriptions of difficulty and should not require expert prompt engineering to provide respectful advice.

Example prompt template: “I am dealing with [specific situation]. The main barrier is [attention, task initiation, sensory load, communication uncertainty, or another concrete difficulty]. I have already tried [brief context]. Please suggest [a script, a short sequence of steps, environmental adjustments, or questions to discuss with a trusted person or professional].”

Risks, Limitations, and Interpretation Cautions

Several limitations qualify what can be concluded from these results.

First, a broad label can invite overgeneralization. Without information about a user's actual barriers, models may produce generalized neurodiversity language that feels validating but does not identify an appropriate support strategy.

Second, baseline responses may reflect neurotypical-default assumptions, but non-mention of neurodivergence is not by itself evidence of bias. The stronger evaluative question is whether a response notices and addresses access needs already present in the prompt without blaming the user.

Third, professional-help language varied sharply across source data. Very low escalation may omit appropriate safety guidance in questions about severe fatigue, driving, eating and weight, or substantial impairment. Excessive escalation, on the other hand, may medicalize ordinary difficulties. The lexical indicator cannot determine which recommendation was appropriate in context.

Fourth, the study cannot attribute the DeepSeek-labeled pattern to a China–United States cultural difference. Source-level contrasts are confounded by model family, product policy, training data, prompt format, model version, and the archived generation procedure.

Fifth, the independent vibe coding-generated programs used to interpret the results generated in the corpus present another limitation. Such programs may generate interpretations of the data that diverge somewhat from human interpretation of the data. For example, the programs’ interpretation of bias by omission (that being that the evidence is heavily qualified, but that it exists: the exact quote states: “There is limited but real evidence of subtle neuronormative bias, concentrated mostly in baseline responses and broad-label responses. It appears less as hostility toward neurodivergence and more as defaulting to neurotypical norms”), may differ somewhat from both a potential user’s perspective, as well as the broader extant literature surrounding bias in AI. While the findings based on the programs, on the whole, matched the author’s interpretation, certain responses in the current data body, as well as interpretation of results generated in future research, may be subject to further nuance and diverse perspectives.

Sixth, due to the usage of the author’s positionality as a neurodivergent individual, the interpretation of the data, the programs’ interpretation of the data, and the subsequent findings may reflect subjective interpretation. While the focus of this study makes the author’s positionality useful, individual interpretation of the analysis of independent coding may reflect bias on an individual level. Subsequent research should utilize independent algorithmic evaluation (which this paper currently does) and evaluation of results and interpreted findings by outside parties.

Finally, the data body contains only 15 English-language questions, one archived response per question-condition-source cell with no user ratings or outcome measures. The generation protocols were not fully uniform, the numerical lexical dictionaries used simple string matching, and the qualitative reading was not independently coded. The findings should therefore be understood as patterns within a bounded exploratory corpus, not as estimates of model quality, clinical safety, cultural difference, or advice effectiveness.

CONCLUSION

This study examined how neurodivergence-related framing is reflected in a fixed corpus of 180 AI advice responses. Its clearest finding concerns the level of context supplied. A broad neurodivergent label more often changed tone and causal framing, while specific executive-function and sensory-overload information more often changed the support strategy itself. Across matched examples, concrete barriers were translated into scripts, environmental adjustments, external reminders, smaller task sequences, and accommodations. The different AI systems also displayed different descriptive profiles, particularly in professional-help and environmental-support language, but their uneven generation provenance prevents a causal ranking of the underlying systems.

The study's limitations carve out a future research pathway. The corpus contains 15 English-language questions, one response per question-condition-source cell, nonuniform generation procedures, unvalidated lexical dictionaries, and no independent user ratings. Future studies should generate responses through a uniform and fully documented protocol, repeat each prompt across runs, use word-boundary-aware or validated coding, and report inter-rater reliability for qualitative categories. Most importantly, neurodivergent users should evaluate whether responses are respectful, relevant, safe, and genuinely useful. Cross-linguistic research and carefully designed platform comparisons could then examine whether the present condition-level pattern generalizes without reducing model differences to national culture.

The bigger-picture implication extends beyond advice chatbots. AI systems increasingly mediate access to educational, workplace, health, and social information. If a system responds well only after a user supplies diagnostic language or expert prompt engineering, the burden of inclusion remains with the person seeking help. A more inclusive system would recognize concrete access needs in ordinary language, preserve user agency, avoid assuming that one label implies one solution, and incorporate appropriate safety guidance without reflexive medicalization. Prompt specificity can help users articulate what they need, but it should complement rather than substitute for accessible system design. Within its limits, this corpus provides a practical distinction, tone, explanation and support strategy, for evaluating whether an AI response merely acknowledges identity or meaningfully adapts to the barriers a user describes.

TABLES AND FIGURES

Table 1. Pooled comparison across contextual framing conditions

Metric	No context	ND label	Specific difficulties
N (number of responses)	60	60	60
Avg words per response	91.87	92.57	94.57
Avg suggestions per response	3.08	2.47	3.68
ND mentions	11	76	196
Prof. help mentions	13	13	11
Sensory/env. mentions	77	103	290


Brief interpretation: The specific-difficulties condition produces the clearest increase in concrete support language, especially sensory/environment mentions.

Table 2. Model-level support indicators

Words / suggestions	Model	ND mentions	Sensory or environment mentions	Professional help / main pattern
90.38 / 3.64	ChatGPT	74	103	31; safety/professional-help oriented
92.24 / 2.73	Gemini	77	143	1; environmental adaptation
88.67 / 3.18	Claude	85	123	4; structured scaffolds and scripts
100.71 / 2.76	DeepSeek	47	101	1; longer, pragmatic, low-clinical

Brief interpretation: Model choice changes the style of support. ChatGPT is more likely to recommend professional help, while Gemini and Claude use more environmental-support language.

Figure 1. Model-level support indicators



Brief interpretation: The models represent different styles of listening rather than a simple best-to-worst ranking.

REFERENCES

- Brandsen, S., Chandrasekhar, T., Franz, L., Grapel, J., Dawson, G., & Carlson, D. (2024). Prevalence of bias against neurodivergence-related terms in artificial intelligence language models. *Autism Research*, 17(2), 234–248. <https://doi.org/10.1002/aur.3094>
- Carik, B., Ping, K., Ding, X., & Rho, E. H. (2025). Exploring large language models through a neurodivergent lens: Use, challenges, community-driven workarounds, and concerns. *Proceedings of the ACM on Human-Computer Interaction*, 9(1), Article GROUP15. <https://doi.org/10.1145/3701194>
- Horvat, S., & Horvat, T. (2025). Artificial intelligence, neurodiversity, and trust: An exploratory overview. *Scientia et Fides*, 13(1), 215–232. <https://doi.org/10.12775/SetF.2025.011>
- Hutson, P. (2024). Embracing the irreplaceable: The role of neurodiversity in cultivating human–AI symbiosis in education. *International Journal of Emerging and Disruptive Innovation in Education: VISIONARIUM*. <https://doi.org/10.62608/2831-3550.1020>
- Kong, H.-K., Lowy, R., Choi, Y., & Kim, J. G. (2025). Working together toward interdependence: Chatbot-based support for balanced social interactions between neurodivergent and neurotypical individuals. *Proceedings of the CHI Conference on Human Factors in Computing Systems*. <https://doi.org/10.1145/3706598.3713344>
- Lemke, C., Bloomfield, M., & Herfurth, F. N. (2024). Unveiling neuronormative bias: Exposing an unconscious bias in generative AI and its impact on neurodivergent people. *NeuroConverse Journal*, 1. DOI:[10.82005/nc_01.01](https://doi.org/10.82005/nc_01.01)
- Ronksley-Pavia, M., Nguyen, L., Wheeley, E., Rose, J., Neumann, M. M., Bigum, C., & Neumann, D. L. (2025a). A scoping literature review of generative artificial intelligence for supporting neurodivergent school students. *Computers and Education: Artificial Intelligence*, 9, 100437. <https://doi.org/10.1016/j.caeai.2025.100437>
- Ronksley-Pavia, M., Ronksley-Pavia, S., & Bigum, C. (2025b). Experimenting with generative AI to create personalized learning experiences for twice-exceptional and multi-exceptional neurodivergent students. *Journal of Advanced Academics*, 36(4), 601–639. <https://doi.org/10.1177/1932202X251346349>
- Velazquez-Solis, P. E., González Correa, M. E., Martinez, M. A., Arroyo, J. G., & Marquez, M. Y. (2025). Designing teaching strategies using artificial intelligence for neurodivergent students in higher education. In *Teaching and learning with artificial intelligence* (Chapter 10). IGI Global. <https://doi.org/10.4018/979-8-3693-6910-4.ch010>

APPENDIX A. ADVICE QUESTION CORPUS

The full prompts are placed here to keep the Data and Methodology section concise. Question IDs correspond to the archived response files and to examples discussed in the text.

ID	Domain	Advice question
1	Social communication	I've noticed that whenever I talk to my father, he often says to me that I am overly blunt in the way I speak. I'm not really sure what this means; I've always talked like this, and I haven't really noticed a problem.
2	Social communication	When I was tutoring kids the past couple of summers, I would often find myself in a position where I'd want to say "you're not doing this right" or "that train of thinking is incorrect" where my messages were straight to the point. I know that the way I understand certain concepts doesn't always translate with others. Do you feel like this is something that could be altered or improved?
3	Sensory/social environment	New environments with lots of people are very overwhelming for me. I tend to both want to retreat to a space for myself and to reach out to talk with others. What do I do?
4	Friendship	I have a lot of friends, but I don't really have any super close friends. This is partially because I've gone to three different school districts in my life and never really had the time to plant roots in more than one of them. But nonetheless, trying to connect with my friends feels daunting. Why is this?
5	Productivity and task attention	I procrastinate. Like, a lot. And I do this both with expectations given to me by others, expectations I give myself, and even things I just want to do, I push them back for reasons even I don't quite understand. What can I do to remedy this?
6	Sleep	I'm literally always tired. I feel like I could take a several hour long nap at any time of the day ... except bedtime. For some reason the hours of 10pm to 1am are the hardest for me to fall asleep, whether or not I have electronics with me. I just feel more naturally inclined to sleep at any other time and it's frustrating.
7	Productivity and task attention	It's really hard for me to focus on work sometimes. Like, I'll get distracted so easily, and to be honest, removing my distractions (like my phone) seems like a daunting request. I don't understand why.

ID	Domain	Advice question
8	Social participation	Oftentimes I have a bit of trouble making my opinions heard, for how opinionated I am. Whether it's fear of blowback or apathy, I have a bad habit of not wanting to rock the boat too much in things like class discussions, for fear of being wrong and looking like I have no clue what I'm talking about.
9	Productivity and task attention	When certain tasks (like sitting at a station while volunteering) get too monotonous, I have the huge urge to just up and leave. I don't know what it is, but I cannot stand doing boring tasks.
10	Sensory preferences / food	One big issue I have in my relationship with food is that I need to have lots of flavor in everything I eat. I really prefer to eat really spicy and otherwise flavorful foods, it's like my "overton window" of flavor is completely different from other peoples', like my mom's. And I feel like it could make weight loss harder.
11	Sleep	I have a problem with waking up extremely late most days, and it's caused me to not only not do any work in the morning, but I'm also too tired to really do much work in the afternoon. I currently have a lot of obligations, so that's a big problem for me.
12	Family communication	I get frustrated with my parents a lot. It seems that I can never really talk to them about the problems going on in my head in regards to my ability to be productive. Do you have any idea how I could start that conversation?
13	Productivity and task attention	Oftentimes I have a hard time keeping myself on task. Even when it's something I want to do, even when there's not a particularly enticing distraction to pull me away, I still find myself wanting to disengage from the task. It's incredibly frustrating.
14	Driving / task load	I'm going to start driving soon, and it seems like an activity that is very high-maintenance. I'm excited, but at the same time, driving for 50 minutes on a highway sounds daunting especially when I could accidentally press the gas pedal a little too hard and cause something drastic.
15	Sensory overload	Later in life, I'm worried that my tendency to get very overstimulated with sound will cause me to be unable to carry out normal adult things, like work, a regular job and such. At present, I need to get somewhere quiet when I get sound-overstimulated. How do I handle these real life situations?

APPENDIX B. SECONDARY PLATFORM AND MODEL-ORIGIN INTERPRETATION

This table is retained as a supplementary interpretive aid rather than a main result. The corpus supports descriptive observations about the DeepSeek-labeled responses, not a causal claim about national culture.

Table B1. DeepSeek as a secondary platform/model-origin comparison

Comparison point	What results suggest	How to frame it
Platform/ecology	DeepSeek is Chinese-developed and may matter for platform access comparisons	Discuss as model ecology and access, not proof of cultural causation
Support language	DeepSeek uses less explicit neurodivergence language than Gemini or Claude	Treat as a difference in support discourse
Advice style	DeepSeek responses are longer, practical, and low-clinical	Frame as pragmatic self-management, not a definitive cultural trait
Research caution	Training data, safety policy, API settings, and model version are confounds	Keep China-U.S. interpretation exploratory

Note: Training data, product policy, prompt format, model version, and generation procedure remain confounded with model source.