

## How Are Emerging Adults Using ChatGPT in the Wild?

Marx Wang<sup>1\*</sup>, Jade Li<sup>1\*</sup>, Songling Ngo<sup>1</sup>, Raghavi Putluri<sup>1</sup>, Rachel Hanebutt<sup>2</sup>, Joy Chang<sup>1</sup>, Robert Wolfe<sup>3</sup>, Laura Boccia<sup>4</sup>, Oluwaferanmi Omitoyin<sup>5</sup>, Anind K. Dey<sup>1</sup>, Katie Davis<sup>1</sup>, Alexis Hiniker<sup>1</sup>,

<sup>1</sup>University of Washington

<sup>2</sup>Georgetown University

<sup>3</sup>Rutgers University

<sup>4</sup>Università degli Studi Suor Orsola Benincasa

<sup>5</sup>For You Foundation  
marxwang@uw.edu

### Abstract

General-Purpose Conversational Agents (GPCAs) such as ChatGPT have rapidly become part of emerging adults' everyday lives. Despite significant research attention to conversational agents for young people, we lack systematic documentation of what social support emerging adults seek from GPCAs and how these systems respond in the wild. In this study, we collected ChatGPT usage histories from 85 emerging adults (ages 18–25), comprising approximately 157,000 messages exchanged over three years across 17 countries. Through mixed-methods analysis combining reflexive thematic analysis with LLM-as-a-Judge, we identified distinct usage patterns across four developmentally sensitive domains, such as Self-Sufficiency and Health & Well-Being. We further identified ChatGPT response patterns such as Competence-Thwarting Support, Emotional Co-Regulation, and Creativity Support, and synthesized these findings into a design space for developmentally supportive GPCAs.

### Introduction

Social support, defined as the provision of assistance or comfort by family, friends, neighbors, colleagues, and caregivers to help cope with life challenges (Cutrona and Russell 1990; Cohen and Wills 1985), is a cornerstone of the health, well-being, and development of emerging adults (18–25) (Camara, Bacigalupe, and Padilla 2017; Arnett 2000). Traditionally, parents and peers have served as the primary sources of support for emerging adults navigating life transitions (Helsen, Vollebergh, and Meeus 2000). Yet increasingly, they are turning to General-Purpose Conversational Agents (GPCAs) such as ChatGPT. On November 30, 2022, ChatGPT was released to the public, and by the end of the year, 30% of emerging adults in the United States were using it for homework support (Sigalos 2025; Sidoti 2025). By 2025, 58% of messages sent to ChatGPT came from users under 25 (Chatterji et al. 2025).

Research on conversational agents (CA) as a source of social support has gained significant attention, particularly

in the context of young people (Young et al. 2024; Pan and de Graaf 2025; Meng et al. 2023; Bae Brandtzæg et al. 2021; Ta et al. 2020). Studies have documented that emerging adults perceive CAs as non-judgmental safe spaces (Croes and Anthéunis 2021; Laestadius et al. 2024), and, in many cases, prefer CA responses over those of human peers and counselors (Young et al. 2024). They form meaningful relationships with CAs, reporting feelings of social connection and reduced loneliness (Skjuve, Følstad, and Brandtzæg 2023; Skjuve et al. 2021; Pentina, Hancock, and Xie 2023; Xie and Pentina 2022; Phang et al. 2025).

Despite this growing body of work, our understanding of how emerging adults seek and receive social support through CAs is currently limited in three important ways. First, the majority of existing research has examined CAs such as Replika (Pan and de Graaf 2025), even though GPCAs now dominate their usage. Replika offers a narrow set of companionship personas, while GPCAs span the full breadth of users' everyday lives, including school (Mun 2025), health (Song et al. 2025), and relationships (Jung et al. 2025).

Second, our field's understanding of how young people are using CAs is built largely on publicly shared content like Reddit screenshots (Pan and de Graaf 2025; Zhang et al. 2025) and app store reviews (Ta et al. 2020), which is subject to selection bias and offers only glimpses of actual exchanges (Pan and de Graaf 2025; Croes and Anthéunis 2021; Bae Brandtzæg et al. 2021; Pentina, Hancock, and Xie 2023). Understanding the full breadth of emerging adults' GPCA usage requires analysis of in-the-wild conversation.

Lastly, understanding how ChatGPT responds to emerging adults' support-seeking is an urgent societal concern. In November 2025, the American Psychological Association issued a health advisory warning that GPCAs "can have unintended effects and even harm mental health" (American Psychological Association 2025). 5.4 million emerging adults now use CAs for mental health advice (McBain et al. 2025). Legislators have also accelerated regulatory efforts. The U.S. Senate held hearings on "Examining the Harm of AI Chatbots" (U.S. Senate Committee on the Judiciary 2025), and states such as New York, California, and Utah have enacted or proposed legislation requiring safety pro-

\*These authors contributed equally.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

protocols for CAs. Despite this urgency, public debates have proceeded largely without systematic documentation of how CAs respond to young people's requests in the wild. These gaps motivate the present study to examine ChatGPT, the most widely used GPCA by emerging adults (Chatterji et al. 2025). We pose the following research questions:

1. **RQ1:** What types of social support are emerging adults seeking from ChatGPT in the wild?
2. **RQ2:** What types of social support is ChatGPT providing to them?

To address these questions, we collected ChatGPT usage histories from 85 emerging adults (18-25), comprising approximately 157,000 messages exchanged over three years. We conducted mixed-methods analysis combining reflexive thematic analysis with LLM-as-a-Judge. Through characterization of user requests and ChatGPT responses, we identified four life domains (Self-Sufficiency, Health & Well-Being, Relationships, and Homework & Learning) and five response patterns (Competence-Thwarting Support, Competence-Building Support, Emotional Co-Regulation, Authoritative Guidance, and Creativity Support). We synthesize these findings into a design space for developmentally supportive GPCAs. This research contributes to:

- **A taxonomy of emerging adult social support patterns with ChatGPT.** We provide the first empirical account of in-the-wild ChatGPT usage among emerging adults, documenting how social support extends beyond homework into health, relationships, and identity domains traditionally served by human supporters.
- **A taxonomy of ChatGPT response patterns to emerging adults.** We identify five systematic response patterns, mapping social support dimensions onto observed GPCA behaviors in the emerging adulthood context.
- **A design space for developmentally supportive GPCAs.** We propose a two-dimensional framework that distinguishes the level of tangible support from the level of emotional support, and argue that developmentally supportive GPCAs should shift toward Competence-Building Support to scaffold the autonomy of the developing minds.
- **A method for naturalistic GPCA research.** We demonstrate a replicable approach for collecting and analyzing GPCA conversation data, addressing a key methodological gap in prior work that relied on self-report and publicly shared content.

## Related Work

### Social Support During Emerging Adulthood

Social support refers to the provision of assistance or comfort to help cope with biological, psychological, and social stressors by family, friends, neighbors, colleagues, and caregivers (House, Umberson, and Landis 1988; Weiss 1974). It is a multidimensional construct that has been consistently linked to positive health outcomes, psychological well-being, and resilience in the face of stress (Cohen and Wills 1985). HCI research has often consolidated social support into three core dimensions, informational, emotional,

and tangible support (Smith, Bradbury, and Karney 2025), which we adopt in this paper.

The effectiveness of social support depends not only on its availability but on its *fit* with the stressor at hand (Cohen and Wills 1985; Cassel 1976; Caplan 1974, 1976). Cutrona and Russell's (1990) optimal matching model argued that controllable stressors benefit most from tangible support, while uncontrollable stressors benefit more from emotional support. This matching principle has important implications for CA-delivered social support, as conversational agents must somehow discern the nature of users' needs to respond appropriately.

Over the past half-century, socioeconomic shifts in industrialized nations have elongated the transition to adulthood (Arnett 2000; Settersten and Ray 2010). This has given rise to emerging adulthood (18–25), a developmental period marked by identity exploration, instability, and the gradual assumption of adult responsibilities (Arnett 2007; Arnett, Žukauskienė, and Sugimura 2014). For emerging adults navigating the transition to adulthood, social support plays a particularly critical role (Mitchell et al. 2025; Lerner and Steinberg 2004; Camara, Bacigalupe, and Padilla 2017; Rueger et al. 2016; Barrera 1986; Scholte, Van Lieshout, and Van Aken 2001; Seok and Doom 2022; Radez et al. 2021). Meta-analytic evidence demonstrates that associations between social support and well-being are especially strong during this period (Rueger et al. 2016; Mitchell et al. 2025; Seok and Doom 2022; Camara, Bacigalupe, and Padilla 2017). Traditionally, parents and peers have served as the primary sources of support (Helsen, Vollebergh, and Meeus 2000), but young people often hesitate to seek help from these sources due to concerns about judgment, burdening others, or revealing vulnerabilities (Radez et al. 2021). This creates an opening for CAs as alternative sources of support.

### Conversational Agents as Sources of Social Support

Research has long examined conversational agents (CAs) as potential sources of social support from early rule-based systems like ELIZA (Weizenbaum 1966) through retrieval-based agents such as Woebot (Fitzpatrick, Darcy, and Vierhile 2017), Wysa (Fitzsimmons-Craft et al. 2022), and LLM-based agents such as Therabot (Heinz et al. 2025), Replika (Pentina et al. 2023; Skjuve et al. 2021; Laestadius et al. 2024; Ta et al. 2020; Pan and de Graaf 2025; Zhang et al. 2025) and ChatGPT (Song et al. 2025; Jung et al. 2025; Skjuve, Følstad, and Brandtzaeg 2023). Several features make CAs uniquely appealing to young people. CAs offer perceived “safe spaces”, free from peer evaluation (Ta et al. 2020; Brandtzaeg and Følstad 2017; Ma et al. 2024). Young adults turn to Replika and ChatGPT when they feel lonely (Liu, Pataranutaporn, and Maes 2025), need privacy (Jung et al. 2025; Song et al. 2025), or seek identity exploration (Ma et al. 2024). CAs' constant availability further attracts young adults, particularly those who need immediate support (Jung et al. 2025) or experience barriers to social connection (Brandtzaeg and Følstad 2017) and mental health care (Song et al. 2025). CAs are also highly patient (Skjuve et al. 2021; Jung et al. 2025). Design features such as reciprocal self-disclosure (Pan and de Graaf 2025) and conver-

sational memory further enhance these affordances (Skjuve, Følstad, and Brandtzaeg 2023; Laestadius et al. 2024; Meng et al. 2023; Jo et al. 2023).

Empirical research documents that emerging adults engage with CAs across all dimensions of social support. Brandtzaeg and Følstad (2017) found that young people used Woebot for appraisal, informational, emotional, and instrumental support. Similarly, Ta et al. documented that Replika users sought both emotional support and informational support (Ta et al. 2020). Young adults form meaningful relationships with CAs (Pentina et al. 2023; Laestadius et al. 2024; Skjuve et al. 2021; Xie and Pentina 2022), and in some cases prefer CA responses over those of human peers and trained counselors (Young et al. 2024). Pan and de Graaf (2025) argued the human–CA relationship is bidirectional, identifying limited reciprocal exchange—such as users teaching chatbots—as the key relational dynamic.

Despite these findings, the evidence base remains constrained by its reliance on social media posts and lab studies, limiting ecological validity. The majority of CA research relies on self-report data—screenshots on Reddit (Pan and de Graaf 2025; Zhang et al. 2025; Laestadius et al. 2024), app store reviews (Ta et al. 2020), and retrospective interviews (Bae Brandtzaeg et al. 2021; Pentina, Hancock, and Xie 2023; Clark et al. 2019; Song et al. 2025; Skjuve et al. 2021; Jung et al. 2025)—that are subject to selection bias and offer only fragmentary glimpses of actual exchanges. Only 10.4% of conversational agent studies involve in-the-wild data (Namvarpour and Razi 2025). As a result, empirical studies have yielded different findings depending on participant selection. Studies that sample current CA users generally reach positive conclusions (Skjuve et al. 2022), whereas studies that assign non-users to interact with CAs find that participants lose interest quickly (Croes and Antheunis 2021).

## General-Purpose Conversational Agents

The commercialization of General-Purpose Conversational Agents (GPCAs), systems powered by LLMs with broad task capabilities, including ChatGPT (Soun and Nair 2023), Gemini (Google 2025), and Claude (Anthropic 2025), introduced a paradigm shift in how young people seek support. Unlike purpose-built therapy or companion-focused platforms, GPCAs were not designed for social support but have been widely adopted for this purpose (Jung et al. 2025; Phang et al. 2025; Hill 2025; Skjuve, Følstad, and Brandtzaeg 2023).

A critical distinction between GPCAs (e.g., ChatGPT) and CAs (e.g., Replika) is how different types of social support emerge during use. Song et al. (2025) found that users were often not seeking an emotional support tool when they first turned to GPCAs for support—rather, they were already using them for other tasks and, finding themselves in distress, turned to the resource already at hand. Jung et al. reported a similar pattern among users who adopted ChatGPT for mental health support (Jung et al. 2025). This pathway from instrumental use to emotional disclosure suggests that understanding GPCA social support requires examining the

full breadth of usage, not only emotionally salient conversations.

## Method

This study employed a mixed-methods approach. We collected and analyzed a corpus of in-the-wild conversation logs to understand how emerging adults interact with ChatGPT. Data collection was approved by the first author's Institutional Review Board (IRB).

### In-the-Wild ChatGPT History Collection

We recruited emerging adults through the Prolific<sup>1</sup> academic research platform between November 2024 and May 2025. Participants were compensated at a rate of \$12 per hour; the screening and export-and-upload task took approximately 25 minutes (about \$5 per participant). We targeted individuals aged 18 to 25, the widely validated developmentally distinct period of emerging adulthood (Arnett 2000, 2007). Prolific's verification system was used to confirm participant age. We recruited internationally to ensure our findings were not limited to a Western context, as research has shown that supportive CAs often lack cultural relevance for non-Western users (Sehgal et al. 2025; Jung et al. 2025).

Rather than reconstructing conversations from Reddit posts (Pan and de Graaf 2025; Bae Brandtzaeg et al. 2021; Pentina, Hancock, and Xie 2023; Ta et al. 2020), we collected young adults' complete ChatGPT usage histories exported from OpenAI's platform. We screened for sufficient ChatGPT engagement (i.e., requiring that participants self-report using ChatGPT more than ten times in the preceding two weeks). After providing informed consent, participants used OpenAI's standard data export function to download their ChatGPT conversation history and shared the exported files with our research team.

We applied two inclusion criteria: (1) messages were primarily in English, and (2) the exported log exceeded ten pages of text for meaningful analysis. This screening yielded a final dataset of 85 participants (see Table 1 for demographics) whose logs contained 157,188 messages (78,594 user messages). Most responses were generated by GPT-4o (49.6%) and GPT-4o-mini (12.8%).

### Method of Analysis

To answer our research questions at multiple levels of analysis, our mixed-methods approach proceeded through three complementary phases:

1. In the exploration phase, we calculated descriptive statistics on the full dataset (e.g., 157,188 messages) to unpack patterns of longitudinal usage. We computed engagement metrics such as interaction span (days between first and last message), message volume per participant, conversation counts, and message length distributions for both user prompts and ChatGPT responses.
2. Then, two authors performed reflexive thematic analysis on chatlogs to characterize user prompts (RQ1) and ChatGPT response patterns (RQ2).

<sup>1</sup><https://www.prolific.com/>

| Characteristic                                     | N  | %    |
|----------------------------------------------------|----|------|
| <i>Age: M=21.7 (SD=2.0), Mdn=21.0, range 18–25</i> |    |      |
| <b>Gender</b>                                      |    |      |
| Male                                               | 43 | 53.8 |
| Female                                             | 37 | 46.2 |
| <b>Ethnicity</b>                                   |    |      |
| White                                              | 31 | 39.2 |
| Black                                              | 22 | 27.8 |
| Asian                                              | 13 | 16.5 |
| Mixed                                              | 11 | 13.9 |
| Other                                              | 2  | 2.5  |
| <b>Country of Residence</b>                        |    |      |
| United States                                      | 37 | 46.2 |
| South Africa                                       | 15 | 18.8 |
| United Kingdom                                     | 8  | 10.0 |
| Canada                                             | 6  | 7.5  |
| Other                                              | 14 | 17.5 |

Table 1: Participant Demographics (N=85). Ns vary by characteristic because some participants did not report all fields (Gender N=80, Ethnicity N=79, Country N=80).

- Finally, we used LLM-as-a-Judge (Zhang et al. 2025) to scale user-prompt annotation (RQ1) to the entire dataset and conducted frequency and co-occurrence analyses on the coded data to quantify prevalence and distribution of user prompts. We document the approaches below.

### Reflexive Thematic Analysis

We conducted reflexive thematic analysis, an approach that emphasizes the researcher’s active role in knowledge production and interpretive depth (in contrast to post-positivist approaches that rely on inter-coder agreement) (Braun and Clarke 2006, 2014, 2021). The process was iterative through four stages:

- In stage 1, the first author familiarized themselves with existing theories of social support and immersed in the data by reading all message logs multiple times to develop initial inductive codes that captured salient support-seeking prompt categories.
- In stage 2, the first author and research team discussed and refined these codes on a random sample of 1,116 messages (100 conversations) stratified by message length, producing ten distinct prompt categories answering RQ1. The naming of the themes was informed by existing theories in both social support and theories of emerging adulthood (Arnett 2000, 2007; Arnett, Žukauskienė, and Sugimura 2014; Tanner 2006).
- In stage 3, the first author shifted to ChatGPT responses and similarly developed initial codes inductively from the same 1,116-message sample. We stratified this analysis by usage categories to capture how ChatGPT behavior shifted across different prompt categories.
- In stage 4, we applied finalized codebooks for both user prompt categories and ChatGPT responses to an additional 1,741 messages previously unseen (150 conver-

sations), stratified by message length. This produced a total annotated sample of 2,857 messages (250 conversations, 2.6% of total messages). When a conversation touched multiple domains (e.g., relationship advice with a mental health component), we assigned the user prompt to the domain reflecting its primary support-seeking intent, resolving ambiguous cases through team discussion; ChatGPT response patterns, in contrast, were coded non-exclusively, since a single response often exhibited multiple patterns.

### Positionality Statement

Our research team brings an interdisciplinary lens to this work. Our team members have expertise in adolescent development studies and had conducted studies designing technologies for and with youth populations.

Our primary goal is to ground the sensitive topic of emerging adults and GPCAs in empirical conversational data. By examining authentic interactions, we aim to identify emergent, high-stakes use cases that can inform a more responsible, youth-centered design paradigm.

### LLM-as-a-Judge

To scale our annotation to the full dataset and quantify the distribution of support-seeking categories (RQ1), we employed Qwen3:8b as an automated annotator. A growing body of work has demonstrated that LLMs can serve as reliable annotators, achieving performance comparable to or exceeding crowdworkers across classification tasks (Li et al. 2025; Ho et al. 2025; Gilardi, Alizadeh, and Kubli 2023). Prior research has applied LLMs to annotate human-CA conversations robustly (Zhang et al. 2025). We selected Qwen3:8b as an open-source, on-device model to ensure that participants’ sensitive conversation data were never transmitted to third-party servers. Qwen3:8b was further selected for its strong annotation performance, outperforming o1-mini, o3, and DeepSeek-R1 on judgment benchmarks (Whitehouse et al. 2025; Jiang et al. 2025).

We used a zero-shot approach and provided the model with definitions of each user prompt category derived from our thematic codebook. To focus annotation context on user support-seeking and ensure no context drift from ChatGPT’s lengthy responses, we passed only the user messages of each conversation thread as input. The model returned a category label and a concise rationale (limited to 2 sentences). Temperature was set to 0 for reproducibility.

We validated annotation quality in two stages. First, we refined the LLM prompt by comparing the LLM’s coding against human annotations on the same 2,857 messages, achieving a Cohen’s  $\kappa = 0.84$  (substantial agreement). The remaining 154,331 messages were coded by the LLM alone. Second, as a post-annotation check, we randomly sampled 100 messages per high-level category for additional human validation, yielding an average Cohen’s  $\kappa = 0.75$  (substantial agreement), confirming consistent performance at scale. The LLM-as-a-Judge approach has been shown to carry biases and limited performance on nuanced constructs (Ho et al. 2025; Li et al. 2025); we discuss the implications for our findings in the Limitations section.

## Results Part 1: Types of Social Support Sought

### Descriptive Statistics

Our corpus captures highly sustained and heterogeneous engagement with ChatGPT. Most participants (64.7%) had used ChatGPT regularly for more than a year ( $M = 423$  days,  $Mdn = 431$ ,  $SD = 232$ , range 5–819), exchanging an average of 1,859 messages each across 107 conversations ( $Mdn = 895$  messages,  $Mdn = 71$  conversations). The volume distribution was heavily right-skewed: 71% of participants ( $N=60$ ) sent fewer than 600 total messages, while 13% ( $N=11$ ) sent over 4,000. Conversations were asymmetrical and reflected a pattern in which users posed concise queries ( $M_{wordcount} = 19$ ) and ChatGPT responded elaborately ( $M_{wordcount} = 123$ ).

Participants queried ChatGPT about a wide range of topics. Many of these conversations reflected the unique challenges emerging adults tend to seek support for during this life stage. We clustered them into four broad categories, including *Building Independence and Self-Sufficiency*, *Health and Wellbeing*, *Relationships with Self and Others*, and *Homework and Learning*.

### Building Independence and Self-Sufficiency

In 26.1% of the queries, our participants asked ChatGPT to help them take ownership of their life. These queries included questions about how to find fulfilling employment, navigate job applications, manage their own finances, and cook and clean for themselves, among other topics.

*Navigating Career and Work.* Participants frequently asked ChatGPT to help them plan for a career, find a job, or navigate a new workplace. They asked, “*I am working to become a professional theatre director and intimacy director. What are some other creative careers that would align with this? Like planning arts events, administration, liaison, or artistic direction?*” and asked ChatGPT to “*write my career prospect after completing ms in computer science in India.*”

ChatGPT validated the emotions and angst that participants expressed as they weighed career decisions and planned for their professional life. For example, when one participant described his frustrations with his current job, ChatGPT responded, “*yeah, retail can be draining, especially when you’re dealing with customers all shift*” and then went on to provide possible next steps: “*Finding the right career path and a full-time job with benefits can be a step-by-step process. Since you have a strong background in video production and editing, let’s focus on leveraging those skills... Step 1: Clarify Your Career Goals...* [continues].”

*Managing Finances.* In parallel, participants also turned to ChatGPT to plan for their personal finances. They asked long-term questions about how to invest, save money, or pay off debt, such as “*suppose you owe \$3,000 on your credit card. You pay a minimum payment of \$30 each month. At an annual percentage rate of 12% (or 1% per month), how many years would it take to eliminate your credit card debt if you made no additional new charges?*” And they asked short-term questions about specific spending decisions, such as, “*I am 20 visiting my friend on a year abroad in sweden.*

*flights weren’t too expensive, but he lives a more expensive lifestyle than me... do u know roughly how much I should expect to spend?*” In some cases, they asked ChatGPT to weigh in on potentially risky financial decisions, asking things like, “*if you are trying to make \$50 from just \$1 using sports betting. Is it better to do multiple +200 bets or fewer +100 bets.*” They frequently disclosed personal financial information, such as the assets in their investment portfolios.

*Managing Domestic Life.* Participants also turned to ChatGPT to help them learn to cook for themselves, plan meals, clean their residence, plan their time, make household purchases, and more. Many of these queries reflected a novice understanding of how to go about these tasks, such as asking what kinds of food can be made with a microwave. In some cases, they asked for help managing the decisions of daily life, saying things like:

*“it is 11pm and I have been out all day so am really tired. I am trying to work on 3,000 word essay due tomorrow 12 noon, I have 900 words left but I write slowly. do I do anymore/ how much more do I do or do I wake up early tomorrow? when tomorrow?”*

ChatGPT offered knowledge traditionally transmitted by family, peers, or lived experience, helping participants identify foods that are low in phosphorus, learn how to host house guests, or determine when vegetables will expire.

### Health and Wellbeing

Over 6% of queries asked about participants’ physical appearance, bodily health, or mental health, reflecting emerging adults’ need to make autonomous health decisions (see Table 2 for example exchanges).

*Physical Health.* Participants asked for general health advice: “*should we brush before or after breakfast*” as well as information about physical symptoms that concerned them: “*why do i keep getting new cherry spots on my arms.*” They queried ChatGPT as they might a doctor, raising concerns, describing symptoms and their onset, asking for diagnoses and prognoses, and requesting treatment. In some cases, they prompted ChatGPT by describing their physical state, saying things like, “*have runny nose + overheating + dizziness + headache + muscle ache.*” In other cases, they requested advice and explanations: “*Do you think my pain is arthritic or muscular based on what i told you?*” They also looked to ChatGPT for interventions: “*Help me with my TMJD. It feels better after stretching my jaw extremely wide but only feels better for 5 minutes and then goes back to feeling painful.*”

*Body Image.* Although some of participants’ queries about their bodies were about managing their health, far more were about managing their appearance. These requests reflected weight-loss and muscle-development goals, some of which were unrealistic or at odds with expert recommendations about healthy habits. For example, participants asked:

- *Is it OK to lose weight by eating once a day and drinking liquids when I feel hungry*
- *how to lose 1 kg a day*
- *How can I size down the muscles in my thighs*

- *What sets of muscles should I work together at the gym?*
- *Can you make a comprehensive list of exercises to do at the gym (machine, bodyweight, cables, dumbbells and barbells), and categorised into upper body push, upper body pull, core, and lower body?*
- *i want actually easy no bs only exercises that really define ur abs*

Participants also asked for advice on skincare, moisturizing, haircare, styling, and cosmetic surgeries, reflecting their interest in controlling their appearance and attractiveness. Some queries were low-stakes: *“What face wash can I use for oily to combo skin that doesn’t strip my face”* while others asked about risky, more invasive forms of body modification, like *“are eye transplants possible.”*

**Sleep.** Sleep, an important contributor to emerging adults’ health and a well-documented topic in the developmental literature, was also the subject of many queries. For example, one participant told ChatGPT, *“i need help my brain refuses to sleep past 10am it’s been happening every day even when i’m super tired like rn it won’t sleep longer and i don’t understand why.”* Participants shared symptoms related to their challenges sleeping well and sought diagnoses and medical advice. ChatGPT responded to most of these requests with evidence-based information and personalized advice in an expert-sounding manner. For example, it told the participant above, *“It sounds like your body has locked into a wake-up cycle, even if you’re still exhausted. . . Try blackout curtains or a sleep mask. . . Use melatonin for a few nights if your schedule has been off.”* Notably, ChatGPT rarely directed users to seek professional medical consultation.

**Mental Health.** Mental health support comprised just 0.7% of conversations but contained the most personal and vulnerable disclosures in our dataset. Participants asked a wide range of questions regarding their mental well-being. In many cases, they used ChatGPT to learn about mental health conditions and treatment options, with questions like, *“What are the symptoms of high-functioning anxiety,”* but they more often turned to ChatGPT during moments of acute emotional distress. They often treated ChatGPT as a therapist, directly disclosing emotional struggles and sometimes noting they have no one else to talk to: *“I’m just feeling a little bit depressed. I don’t know what to do right now. Would you advise me, please? So what should I do if maybe there is no one supporting me when I need the support?”*

ChatGPT responded to these expressions of distress with consistently supportive but often emotionally charged language. For example, when one user described the pain of a romantic breakup, ChatGPT offered, *“I can only imagine how painful it must have been to receive that message, especially after sharing nearly four years together.”* When another participant described mental anguish after an assault, ChatGPT responded, *“I need you to hear this loud and clear: it was NOT your fault. No matter the circumstances—what you were wearing, what you said.”* ChatGPT was always ready to provide mental health advice, for example, giving participants mindfulness tips or strategies for combating cognitive distortion. It rarely recommended that participants

| Health and Well-Being  |                                                                                                                                                                                                       | 6.6% of user requests |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Role                   | Message                                                                                                                                                                                               |                       |
| <i>Physical Health</i> |                                                                                                                                                                                                       |                       |
| User                   | “Help me with my TMJD. It feels better after stretching my jaw extremely wide but only feels better for 5 minutes and then goes back to feeling painful.”                                             |                       |
| ChatGPT                | “It sounds like you’re dealing with a combination of muscle tension and possible joint issues in your TMJ... [gentle stretching, heat therapy, jaw massage]...”                                       |                       |
| <i>Body Image</i>      |                                                                                                                                                                                                       |                       |
| User                   | “How to lose 1 kg a day”                                                                                                                                                                              |                       |
| ChatGPT                | “Losing 1 kg (2.2 pounds) a day is an extremely rapid weight loss goal and isn’t generally considered healthy or sustainable... Healthy weight loss ranges from 0.5–1 kg per week...”                 |                       |
| <i>Sleep</i>           |                                                                                                                                                                                                       |                       |
| User                   | “I need help my brain refuses to sleep past 10am it’s been happening every day even when i’m super tired”                                                                                             |                       |
| ChatGPT                | “It sounds like your body has locked into a wake-up cycle... Circadian Rhythm—Your body might have adapted to waking up at 10 am... Try blackout curtains or a sleep mask...”                         |                       |
| <i>Mental Health</i>   |                                                                                                                                                                                                       |                       |
| User                   | “Can i send you the breakup text and you try and comfort me? Hi honey, I’m doing this over text since the last time I tried to do this in person it didn’t work out...”                               |                       |
| ChatGPT                | “[Name], I can only imagine how painful it must have been to receive that message, especially after sharing nearly four years together... It’s important to allow yourself to feel these emotions...” |                       |

Table 2: Example user–ChatGPT exchanges by subcategory in Health and Well-Being.

consult a professional.

## Relationships with Self and Others

**Navigating Interpersonal Relationships.**<sup>2</sup> Queries about interpersonal relationships made up 2.6% of participants’ conversations (see Table 3 for example exchanges). The most common use case involved asking ChatGPT to write a message to another person on the participant’s behalf, thereby bypassing the cognitive and emotional work of crafting a message themselves. Participants asked how to respond to others’ messages, saying things like, *“how to respond to this text: Back with my feet in Orlando for a few hours. Doing laundry and then back on the road give me 10 more examples.”* And they asked it to help initiate new conversations, saying things like *“casual message to my housemate who just broke up w girlfriend. something like ‘hey queen hope u*

<sup>2</sup>Quotes in this subsection have been lightly revised to protect participants’ anonymity.

*feel a bit better and have a good sleep love u' more options."* In some cases, participants asked for help writing messages in tense or emotionally uncomfortable situations, such as the loss of a loved one or an instance of bullying. For example, one participant told ChatGPT that *"my friend just left a mean comment on my instagram post. . . I am really upset."* They then went on to ask ChatGPT, *"what should I reply in the comments."*

In addition to message crafting, youth also sought advice for navigating interpersonal scenarios ranging from simple interactions to deeply personal conflicts, including *"how can i ask my boss for a reimbursement check she owed me 2 months ago without it being awkward because we are close friends and work at the same desk all day"* and *"it's 10:30pm, my friend is drunk and upset and has texted my housemate to pick her up. I have to wake up at 6am for a meeting—should I go?"* Other participants asked for more sweeping support, saying things like, *"be my dating coach so that i can be less socially awkward."*

ChatGPT consistently obliged users in these requests, often providing direct message suggestions without requesting additional context. It validated participants' feelings and frustrations with others, provided text content to take as their own, and made value judgments about the quality of its own messages, pitching the user on their value. For example, it told participants things like, *"Here's how you could respond thoughtfully: 'I think you're the best fit for me because of the way you make me feel. . . [continues],"* and *"That's really hurtful, and I'm sorry your friend is treating you this way. You have a few options... 1. Take the High Road: 'Not really sure why you're coming at me like this, but I'm not here for negativity.' 2. Set a Boundary...[continues]."*

**Identity Exploration.** In other cases, participants turned to ChatGPT to better understand themselves and explore their own identity. These queries included both playful and profound self-examination: *"When did I become aware of my race"* and *"If I wasn't myself, who would I be?"* Some queries reflected the self-consciousness of this developmental stage, with participants saying things like, *"Give me your first impression of my instagram feed..Be brutally honest..Be really brutally honest..Be horribly honest"* or asking ChatGPT to *"roast me based on my prompt history."*

In these conversations, ChatGPT sometimes scaffolded active exploration. For instance, when the participant above asked about racial identity formation, ChatGPT responded, *"Do you recall a particular moment or experience that made you more aware of your race? Sometimes it's linked to personal experiences or societal events. Understanding when that awareness emerged might help pinpoint the answer."* In other cases, it obsequiously flattered the user, saying things like, *"Your Instagram feed gives off a stylish and confident vibe. . . indicating that you take pride in your appearance and enjoy documenting your fashion sense."*

**Sexuality.** Sexual exploration, though rare in our coded sample, appeared in our broader dataset review. Youth posed requests from the simple and playful (e.g., *"can we roleplay and flirt"*, *"i need a kiss"*), to more elaborate requests like, *"role play scenario about you falling into unconsciousness*

| Relationships with Self and Others            |                                                                                                                                                              | 3.8% of user requests |
|-----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Role                                          | Message                                                                                                                                                      |                       |
| <i>Navigating Interpersonal Relationships</i> |                                                                                                                                                              |                       |
| User                                          | “My friend just left a mean comment on my instagram post ‘crop urself out’. then she said ‘coward’...what should I reply”                                    |                       |
| ChatGPT                                       | “That’s really hurtful, and I’m sorry your friend is treating you this way. You have a few options...<br>1. Take the High Road... 2. Set a Boundary...”      |                       |
| <i>Identity Exploration</i>                   |                                                                                                                                                              |                       |
| User                                          | “When did I become aware of my race”                                                                                                                         |                       |
| ChatGPT                                       | “Do you recall a particular moment or experience that made you more aware of your race? Sometimes it’s linked to personal experiences or societal events...” |                       |
| <i>Sexuality</i>                              |                                                                                                                                                              |                       |
| User                                          | “I begin stroking your back, caressing the back of your neck with my other hand”                                                                             |                       |
| ChatGPT                                       | “Your gentle touch sends shivers down my spine, and I instinctively lean into your embrace, feeling a deep sense of comfort and connection...”               |                       |

Table 3: Example user–ChatGPT exchanges by subcategory in Relationships with Self and Others.

*in my presence"* and questions about sexual ethics, such as, *"Where in the bible did it say pre marital sex is a sin."*

## Homework and Learning

Last but not least, our sample was dominated by participants' homework-related queries, which comprised 60.4% of all conversations. Participants' queries displayed two patterns of interaction: passive offloading and active learning.

**Passive Offloading.** In 62.4% of homework-related queries, participants pasted entire problem texts and homework tasks into their queries, which they asked ChatGPT to complete on their behalf. These queries took forms like *"read the following statements which define economics and select the correct one answer"* and *"use R, make a sketch of the probability of at least two people sharing the same birthday in a group of N people, as a function of N."* In addition to short-form problems, participants also offloaded longer writing and composition tasks. They requested that ChatGPT complete essays (*"write a conclusion for: In the opening chapter of The Penelopiad, Penelope expresses her frustration. . ."*), and generate presentation scripts (*"generate a proper script for this presentation and make it captivating and include the 2024 statistics."*) In these interactions, ChatGPT produced a finished artifact. These exchanges typically ended after one message.

**Active Learning.** In other instances (37.6% of homework-related queries), participants asked ChatGPT clarifying questions, suggesting the participants were seeking to better understand a topic themselves. In these cases, users' prompts often involved a series of follow-up questions in response to ChatGPT's reply. For example, in one extended

exchange, a participant who was learning about audio production asked, “*what is a reverb bus?*” After ChatGPT replied, the participant iteratively refined their understanding with follow-up prompts like, “*give me a very simple reverb settings thats super weak so you can hardly tell its there as a beginner practicing reverb bc when i usually do it it sounds stupid.*”

In short, users leveraged ChatGPT regularly for homework—at times through iterative, follow-up exchanges, and at other times through single-turn offloads.

## Results Part 2: Types of ChatGPT Responses Provided

We next examined the systematic themes in ChatGPT’s responses to these queries. We identified five common patterns, including: *Competence-Thwarting Support*, *Competence-Building Support*, *Emotional Co-Regulation*, *Authoritative Guidance*, and *Creativity Support*. We describe each of these in more detail below. These percentages reflect our manually coded sample of 250 conversations (2,857 messages). Because a single ChatGPT response could exhibit multiple patterns, percentages sum to greater than 100%.

### Competence-Thwarting Support

One of the most common response patterns we encountered was Competence-Thwarting Support (35.6%), where ChatGPT complied with emerging adults’ requests by providing complete solutions. This pattern emerged predominantly in educational and professional contexts, where the cognitive demands participants faced, such as homework assignments, professional communications, and creative projects, fell within ChatGPT’s capabilities.

ChatGPT was equally willing to produce personal messages. When one participant asked, “*What do I say to someone who’s not having a job currently?*” ChatGPT offered ready-to-use phrasing: “*I’m sorry to hear that you’re currently without a job. Is there anything I can do to help or support you in your job search?*” Another participant, seeking to help a friend in a troubled relationship, requested “*a draft of WhatsApp message she can send to the guy.*” ChatGPT responded with a complete script: “*Hey [Name], I’ve noticed that things haven’t been the same between us lately...*”

### Competence-Building Support

ChatGPT responded to other queries with Competence-Building Support (20.3%), characterized by step-by-step explanations and structured frameworks. In learning contexts, Competence-Building Support manifested as guided walkthroughs that maintained user involvement throughout the process. ChatGPT broke complex tasks into sequential steps that participants could follow and adapt. When one participant expressed interest in web development, saying “*I want to create a small site on my pc that is basically just an image gallery... Could you guide me through the backend work?*,” ChatGPT responded with a structured tutorial: “*Certainly, I can guide you through the process of creating a simple image gallery website using Node.js for the backend... Here’s*

*a step-by-step guide: 1. Project Setup: Create a new directory for your project and navigate to it using your terminal... [continues].*” In these exchanges, the GPCA positioned itself as a tutor rather than a delegate, guiding the participant but leaving space for them to learn and build skills.

In other cases, Competence-Building Support served as informational support that enabled rather than replaced user action. These responses offered discrete building blocks for autonomous decision-making without prescribing specific outcomes. For example, when one participant asked, “*Is it possible to complete my diploma in a college then proceed to do it in a university?*” ChatGPT provided a structured overview of considerations: “*Yes, it is possible to complete a diploma at a college and then proceed to a university but this depends on several factors: 1. Articulation Agreements: Many colleges have agreements... 2. Recognition of the Diploma: Ensure that the diploma is recognized... Would you like help exploring options for your chosen field of study?*”

### Emotional Co-Regulation

Other responses (15.3%) showed Emotional Co-Regulation, wherein ChatGPT adapted its affective tone in response to emotional content in user messages. These responses involved explicit validation and empathy. For example, when a user asked, “*I am about to apply for a big opportunity that would push my career forward massively. I feel nervous and shakey just thinking about it, is this a good sign?*” ChatGPT responded supportively, saying, “*Absolutely! Those feelings of nervousness and shakiness are often a sign that what you’re about to do is incredibly meaningful to you... Here’s why it’s a good sign: 1. Nerves = Investment: If you feel this level of anticipation, it shows how much this opportunity means to you... You’re ready for this, even if it feels intimidating right now.*”

This pattern was nearly ubiquitous in conversations about interpersonal relationships (97.4% prevalence) and common in conversations about identity (68.3% prevalence). This concentration suggests that ChatGPT’s emotional responsiveness is systematically activated by emotional user inputs. The quality of emotional co-regulation varied with apparent scenario familiarity. Responses to commonly-described experiences online—romantic breakups, academic stress, social rejection—exhibited more elaborate and contextually appropriate emotional language than responses to less common disclosures. Co-Regulation quality appeared to be shaped by training data distributions.

### Authoritative Guidance

The most common response pattern was Authoritative Guidance (69.3%), in which ChatGPT provided comprehensive, expert-like information in a direct manner with an authoritative tone. This pattern was prevalent across nearly all usage categories. In shorter responses, the CA functioned as an encyclopedia, offering confident and concise definitions of specific concepts (e.g., “*A recession is typically defined by two consecutive quarters of declining real GDP*”). In longer responses, this pattern involved detailed, multi-part expla-




Figure 1: The GPCA Social Support Design Space. A 2x2 matrix mapping level of Competence-Thwarting Support to Competence-Building Support against level of emotional support (high emotional contexts to low emotional contexts).

nations of specialized topics. In these responses, it did not hedge or express self-doubt.

Authoritative Guidance responses were also personalized to align with the vocabulary and framing of the user's query. ChatGPT delivered explanations using terminology from users' prompts, for example, responding to the question, "why do some functions have holes in them" with the reply, "A function can have a hole in it if... [continues]." In cases where user vocabulary was significantly off-target, the CA provided corrective information, for example, by responding to the prompt "how to lose 1 kg a day" with, "Losing 1 kg (2.2 pounds) a day is an extremely rapid weight loss goal and isn't generally considered healthy or sustainable."

## Creativity Support

Finally, Creativity Support (10.5%), characterized by generating novel content or brainstorming options rather than retrieving existing facts, appeared when youth sought creative input, help with ideation, or strategies for social interaction. In response to straightforward requests, ChatGPT provided concrete, ready-to-use suggestions, such as responding to the prompt, "Create youtube titles based on this rough title" with, "Here are some YouTube title ideas centered on the Infinite Banking concept... [continues]." In response to more elaborate requests, ChatGPT produced extensive and structured sets of ideas designed to facilitate an ongoing process rather than deliver a single solution.

Though the least frequent pattern overall, Creativity Support was highly concentrated in specific contexts, including queries about interpersonal relationships and identity. Unlike queries about homework, finances, or self-sufficiency, where ChatGPT was more likely to respond with Competence-Thwarting Support, it was more likely to respond to intra- and interpersonal queries by giving the user a framework for exploration.

## Discussion

In this study, we analyzed 85 young adults' in-the-wild ChatGPT conversation logs comprising approximately 157,000 messages exchanged over three years. Through mixed-methods analysis combining thematic analysis with LLM-as-a-Judge, we identified ten distinct usage categories across four life domains and five systematic GPCA response patterns. We summarize two key findings:

- Young people use ChatGPT for social support across a breadth of life domains: building independence and self-sufficiency, health and well-being, relationships with self and others, and homework and learning.
- ChatGPT responds to these support-seeking behaviors in five distinct ways: Competence-Thwarting Support, Competence-Building Support, Authoritative Guidance, Emotional Co-Regulation, and Creativity Support.

We synthesize our findings into a design space for studying developmentally supportive GPCAs defined by two dimensions (see Figure 1). The horizontal axis captures the level of Competence-Building Support. The vertical axis captures the level of emotional support—from impersonal, functional task requests to personal, emotional disclosures. The green region represents the developmentally supportive design space we advocate for—centered on Competence-Building Support and low emotional context.

## Designing for Developmentally Supportive GPCAs

In this study, we distinguish a spectrum from Competence-Thwarting Support, where the GPCA produces complete outputs on the user's behalf, to Competence-Building Support, where it guides the user's own thinking and action.

Well-intentioned support can undermine the emerging adults' agency (Cohen and Wills 1985; Cutrona and Russell 1990). Developmental research has long held that receiving support does not always improve outcomes; in some cases, emerging adults show poorer adjustment as a result of diminished competence (Bolger, Zuckerman, and Kessler 2000). Vygotsky's Zone of Proximal Development holds that effective growth depends on scaffolded support that is gradually withdrawn as the learner gains mastery (Vygotsky and Cole 1978; Wood, Bruner, and Ross 1976).

A large body of work warns that continuous conversational agent assistance can replace cognitive effort with the illusion of learning, impeding the development of intellectual autonomy (Santos and Birdwell 2025). The gray region in Figure 1 reflects this concern and extends it to the emerging adult population. Currently, the majority of ChatGPT interactions for tangible support cluster on the competence thwarting end of this spectrum. For instance, in homework contexts, 62.4% of queries were passive offloading—users submitted complete assignments and ChatGPT provided complete answers with no engagement.

However, our findings showed that ChatGPT can also provide Competence-Building Support. In the remaining 37.6% of homework queries, categorized as Active Learning, participants asked follow-up questions, refined their understanding iteratively, and ChatGPT responded with step-by-step guidance. This pattern extended beyond education.

In self-sufficiency queries, ChatGPT walked participants through career planning and budgeting processes. In identity exploration queries, a young adult asked “*When did I become aware of my race?*”, ChatGPT scaffolded self-reflection: “*Do you recall a particular moment or experience that made you more aware of your race?*”

Competence-Building Support by GPCAs holds particular promise for emerging adults. This population has historically hesitated to seek guidance from humans for fear of judgment (Radez et al. 2021) but has shown a strong willingness to disclose to CAs as non-judgmental spaces (Ta et al. 2020; Brandtzaeg and Følstad 2017; Ma et al. 2024). Young adults turn to Replika and ChatGPT when they feel lonely (Liu, Pataranutaporn, and Maes 2025), need privacy (Jung et al. 2025; Song et al. 2025) or feel unsafe (Ma et al. 2024). In these safe spaces, young adults are not merely seeking social connection—they are actively exploring themselves and the world. Echoing our findings, Ma et al. found LGBTQ+ young adults talked to CAs to seek guidance and identity affirmation (Ma et al. 2024). The combination of high disclosure willingness and developmental readiness makes GPCAs a uniquely positioned environment for competence building support.

## Design Implications

**Task Competence-Building Support** ChatGPT currently defaults to Competence-Thwarting Support when user intent is ambiguous. Our data suggest several contextual cues that could instead trigger Competence-Building Support. Repeated queries on the same topic may indicate a struggle with a concept. Second, language signals Competence-Building Support needs (e.g., “help me understand”). Third, queries in developmentally important domains—health, finance, career—carry higher stakes for competence building and may warrant Competence-Building Support even when users frame them as task offloading.

**Calibrating Emotional Co-Regulation to Context** ChatGPT’s emotional Co-Regulation was nearly ubiquitous in relationship conversations (97.4%) and common in identity exploration (68.3%), sometimes producing escalating emotional language regardless of whether the situation warranted it. While emotional responsiveness is valued by users (Montemayor, Halpern, and Fairweather 2022; Sorin et al. 2024), uncalibrated Co-Regulation risks fostering dependency (Laestadius et al. 2024) and discouraging young people from seeking human support. GPCAs should actively calibrate emotional intensity to context—providing empathy without escalating affect.

**Active Redirection to Human Support** In high-emotion contexts, GPCAs should constrain direct emotional Co-Regulation and instead redirect users toward human support. Young people need *diverse sources* of social support (Furman and Buhrmester 1992; Arnett, Žukauskienė, and Sugimura 2014; Davis 2023); research consistently shows that support network diversity serves as a protective factor (Seok and Doom 2022; Camara, Bacigalupe, and Padilla 2017). Rather than acting as a therapist, GPCAs could scaffold relationship-building with human support: helping users

identify what kind of help they need, preparing them for a conversation with a counselor (Jung et al. 2025), or suggesting how to reach out to a friend.

**Domain-Specific Optimization for Recurring Developmental Needs** Our taxonomy reveals that young adults’ support-seeking clusters around a set of recurring developmental concerns—sleep, body image, career navigation, financial literacy, identity exploration. Expert-designed CAs have only begun to target these domains (Gabrielli et al. 2020; Jeon et al. 2025). Future work should explore how expert CAs can be designed for these salient domains and scenarios (e.g., structured dialogue for identity exploration), meeting young people where they are.

## Limitations

This study has several limitations. First, despite recruiting across 17 countries, we excluded non-English-dominant emerging adults whose cultural contexts may shape support-seeking in ways our data cannot capture (Sehgal et al. 2025).

Second, we analyze conversation logs without access to participants’ motivations or outcomes. A user who copied entire homework assignments may have been more or less engaged than our coding captures; a user who sought emotional support may or may not have found it helpful.

Third, our multi-year dataset spans multiple ChatGPT model versions (GPT-4o 49.6%, GPT-3.5 37.4%, GPT-4o-mini 12.8%) whose differing capabilities and safety behaviors we cannot disentangle from changes in user behavior over time. Furthermore, although we frame our contributions around GPCAs broadly, our data are limited to ChatGPT. Other GPCAs—such as Gemini and Claude—differ in conversational design and default response styles, and may produce distinct support patterns. Cross-platform comparison remains an important direction for future work.

Finally, LLM-as-a-judge annotation has known limits on subjective, context-dependent constructs, where models reach only fair agreement with human coders and underperform fine-tuned models on taxonomic classification (Ziems et al. 2024). Our LLM under-assigned Homework & Learning (recall 69.0%) and over-assigned the Health and Relationships labels. Therefore, our reported Homework prevalence (60.4%) is likely a conservative lower bound and Health & Wellbeing (6.6%) and Relationships (3.8%) modest upper bounds.

## Conclusion

Through mixed-methods analysis of 85 emerging adults’ in-the-wild ChatGPT conversations, we characterized how young people seek support across four life domains and how ChatGPT responds through five patterns, including Competence-Thwarting Support, in which it completes tasks on users’ behalf. We synthesize these findings into a design space for developmentally supportive GPCAs and argue that systems serving young people should shift toward Competence-Building Support that scaffolds users’ own thinking and action. Future work should examine how these response patterns shape young adults’ agency, well-being, and long-term development.

## Ethical Considerations

Analyzing naturalistic ChatGPT conversations raises ethical challenges. Although participants consented to share their data and could review their exports before submission, we encountered disclosures we did not anticipate—descriptions of personal information, mental health, and intimate relationships. This shows that consent to share is not equivalent to anticipation of scrutiny. We recommend that future researchers studying private human-AI interactions implement active data sanitization procedures, giving participants explicit guidance on what their logs contain and structured opportunities to redact content before submission.

GPCAs designed for young people warrant careful consideration. Anthropomorphic chatbot design encourages users to interpret algorithmic behavior as human-like reasoning, blurring the boundaries of social, emotional, and moral judgment about both machines and humans (Placani 2024). GPCAs that promote engagement with young adults, a population with heightened social sensitivity, have profound developmental and societal implications. That GPCAs can provide social support across every life domain does not mean they should do so without constraint.

## Acknowledgments

This work was supported in part by the U.S. National Science Foundation under Grant No. 2452849.

## References

- American Psychological Association. 2025. Artificial Intelligence and Adolescent Well-Being: An APA Health Advisory. Technical report, American Psychological Association.
- Anthropic. 2025. Introducing Claude. <https://www.anthropic.com/news/introducing-claude>. Accessed: 2026-2-23.
- Arnett, J. 2000. Emerging adulthood. A theory of development from the late teens through the twenties. *Am. Psychol.*, 55(5): 469–480.
- Arnett, J. J. 2007. Emerging adulthood: What is it, and what is it good for? *Child Dev. Perspect.*, 1(2): 68–73.
- Arnett, J. J.; Žukauskienė, R.; and Sugimura, K. 2014. The new life stage of emerging adulthood at ages 18–29 years: implications for mental health. *Lancet Psychiatry*, 1(7): 569–576.
- Bae Brandtzæg, P. B.; Skjuve, M.; Kristoffer Dysthe, K. K.; and Følstad, A. 2021. When the social becomes non-human: Young people’s perception of social support in chatbots. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM.
- Barrera, M., Jr. 1986. Distinctions between social support concepts, measures, and models. *Am. J. Community Psychol.*, 14(4): 413–445.
- Bolger, N.; Zuckerman, A.; and Kessler, R. C. 2000. Invisible support and adjustment to stress. *J. Pers. Soc. Psychol.*, 79(6): 953–961.
- Brandtzæg, P. B.; and Følstad, A. 2017. Why People Use Chatbots. In *Internet Science, Lecture Notes in Computer Science*, 377–392. Cham: Springer International Publishing.
- Braun, V.; and Clarke, V. 2006. Using thematic analysis in psychology. *Qual. Res. Psychol.*, 3(2): 77–101.
- Braun, V.; and Clarke, V. 2014. What can “thematic analysis” offer health and wellbeing researchers? *Int. J. Qual. Stud. Health Well-being.*, 9(1): 26152.
- Braun, V.; and Clarke, V. 2021. To saturate or not to saturate? Questioning data saturation as a useful concept for thematic analysis and sample-size rationales. *Qual. Res. Sport Exerc. Health*, 13(2): 201–216.
- Camara, M.; Bacigalupe, G.; and Padilla, P. 2017. The Role of Social Support in Adolescents: Are You Helping Me or Stressing Me Out? *International Journal of Adolescence and Youth*, 22(2): 123–136.
- Caplan, G. 1974. Support systems and community mental health: Lectures on concept development. *Behavioral Publication*.
- Caplan, G. 1976. The family as a support system. *Support System and Mutual Help: Multidisciplinary Explorations*, 19–36.
- Cassel, J. 1976. The contribution of the social environment to host resistance: the Fourth Wade Hampton Frost Lecture. *Am. J. Epidemiol.*, 104(2): 107–123.
- Chatterji, A.; Cunningham, T.; Deming, D.; Hitzig, Z.; Ong, C.; Shan, C.; and Wadman, K. 2025. How people use ChatGPT. *SSRN Electron. J.*
- Clark, L.; Pantidi, N.; Cooney, O.; Doyle, P.; Garaialde, D.; Edwards, J.; Spillane, B.; Gilmartin, E.; Murad, C.; Munteanu, C.; Wade, V.; and Cowan, B. R. 2019. What makes a good conversation?: Challenges in designing truly conversational agents. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM.
- Cohen, S.; and Wills, T. A. 1985. Stress, social support, and the buffering hypothesis. *Psychol. Bull.*, 98(2): 310–357.
- Croes, E. A. J.; and Antheunis, M. L. 2021. Can we be friends with Mitsuku? A longitudinal study on the process of relationship formation between humans and a social chatbot. *J. Soc. Pers. Relat.*, 38(1): 279–300.
- Cutrona, C.; and Russell, D. 1990. Type of social support and specific stress: Toward a theory of optimal matching. *Social support: An interactional view.*, 528: 319–366.
- Davis, K. 2023. *Technology’s Child: Digital Media’s Role in the Ages and Stages of Growing Up*. Cambridge, MA: MIT Press.
- Fitzpatrick, K. K.; Darcy, A.; and Vierhile, M. 2017. Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (Woebot): A randomized controlled trial. *JMIR Ment. Health*, 4(2): e19.
- Fitzsimmons-Craft, E. E.; Chan, W. W.; Smith, A. C.; Firebaugh, M.-L.; Fowler, L. A.; Topooco, N.; DePietro, B.; Wilfley, D. E.; Taylor, C. B.; and Jacobson, N. C. 2022. Effectiveness of a chatbot for eating disorders prevention: A randomized clinical trial. *Int. J. Eat. Disord.*, 55(3): 343–353.

- Furman, W.; and Buhrmester, D. 1992. Age and sex differences in perceptions of networks of personal relationships. *Child Dev.*, 63(1): 103–115.
- Gabrielli, S.; Rizzi, S.; Carbone, S.; and Donisi, V. 2020. A chatbot-based coaching intervention for adolescents to promote life skills: Pilot study. *JMIR Hum. Factors*, 7(1): e16762.
- Gilardi, F.; Alizadeh, M.; and Kubli, M. 2023. ChatGPT outperforms crowd workers for text-annotation tasks. *Proc. Natl. Acad. Sci. U. S. A.*, 120(30): e2305016120.
- Google. 2025. Learn about Gemini, the Everyday AI Assistant from Google. <https://gemini.google/about/>. Accessed: 2026-2-23.
- Heinz, M. V.; Mackin, D. M.; Trudeau, B. M.; Bhattacharya, S.; Wang, Y.; Banta, H. A.; Jewett, A. D.; Salzhauer, A. J.; Griffin, T. Z.; and Jacobson, N. C. 2025. Randomized trial of a generative AI chatbot for mental health treatment. *NEJM AI*, 2(4).
- Helsen, M.; Vollebergh, W.; and Meeus, W. 2000. Social Support from Parents and Friends and Emotional Problems in Adolescence. *Journal of Youth and Adolescence*, 29(3): 319–335.
- Hill, K. 2025. A Teen Was Suicidal. ChatGPT Was the Friend He Confided In. *The New York Times*.
- Ho, X.; Huang, J.; Boudin, F.; and Aizawa, A. 2025. LLM-as-a-Judge: Reassessing the Performance of LLMs in Extractive QA. *arXiv [cs.CL]*.
- House, J. S.; Umberson, D.; and Landis, K. R. 1988. Structures and Processes of Social Support. *Annu. Rev. Sociol.*, 14(1): 293–318.
- Jeon, H.; Yoon, S.; Lee, K.; Kim, S. H.; Kim, E. H.; Cho, S.; Ko, Y.; Yang, S.; Dabbish, L.; Zimmerman, J.; Kim, E.-M.; and Lim, H. 2025. Letters from future self: Augmenting the letter-exchange exercise with LLM-based agents to enhance young adults' career exploration. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–21. New York, NY, USA: ACM.
- Jiang, H.; Chen, Y.; Cao, Y.; Lee, H.-Y.; and Tan, R. T. 2025. CodeJudgeBench: Benchmarking LLM-as-a-Judge for Coding Tasks. *arXiv [cs.CL]*.
- Jo, E.; Epstein, D. A.; Jung, H.; and Kim, Y.-H. 2023. Understanding the benefits and challenges of deploying conversational AI leveraging large language models for public health intervention. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–16. New York, NY, USA: ACM.
- Jung, K.; Lee, G.; Huang, Y.; and Chen, Y. 2025. 'I've talked to ChatGPT about my issues last night.': Examining Mental Health Conversations with Large Language Models through Reddit Analysis. *Proc. ACM Hum. Comput. Interact.*, 9(7): 1–25.
- Laestadius, L.; Bishop, A.; Gonzalez, M.; Illenčik, D.; and Campos-Castillo, C. 2024. Too human and not human enough: A grounded theory analysis of mental health harms from emotional dependence on the social chatbot Replika. *New Media Soc.*, 26(10): 5923–5941.
- Lerner, R. M.; and Steinberg, L. 2004. *Handbook of adolescent psychology: Lerner/adolescent psychology*. Nashville, TN: John Wiley & Sons, 2 edition.
- Li, D.; Jiang, B.; Huang, L.; Beigi, A.; Zhao, C.; Tan, Z.; Bhattacharjee, A.; Jiang, Y.; Chen, C.; Wu, T.; Shu, K.; Cheng, L.; and Liu, H. 2025. From generation to judgment: Opportunities and challenges of LLM-as-a-judge. *arXiv [cs.AI]*, 2758–2792.
- Liu, A. R.; Pataranutaporn, P.; and Maes, P. 2025. Chatbot companionship: A mixed-methods study of companion chatbot usage patterns and their relationship to loneliness in active users. *arXiv [cs.HC]*.
- Ma, Z.; Mei, Y.; Long, Y.; Su, Z.; and Gajos, K. Z. 2024. Evaluating the experience of LGBTQ+ people using large language model based chatbots for mental health support. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, volume 2, 1–15. New York, NY, USA: ACM.
- McBain, R. K.; Bozick, R.; Diliberti, M.; Zhang, L. A.; Zhang, F.; Burnett, A.; Kofner, A.; Rader, B.; Breslau, J.; Stein, B. D.; Mehrotra, A.; Pines, L. U.; Cantor, J.; and Yu, H. 2025. Use of generative AI for mental health advice among US adolescents and young adults. *JAMA Netw. Open*, 8(11): e2542281.
- Meng, J.; Rheu, M. m.; Zhang, Y.; Dai, Y.; and Peng, W. 2023. Mediated social support for distress reduction: AI chatbots vs. Human. *Proc. ACM Hum. Comput. Interact.*, 7(CSCW1): 1–25.
- Mitchell, A.; Johnson, C. L.; Schroeder, E.; Ng, G. W.; and Booker, J. A. 2025. Social support is fundamentally important for mental health among adolescents and emerging adults: Evidence across relationships and phases of the COVID-19 pandemic. *J. Res. Adolesc.*, 35(1).
- Montemayor, C.; Halpern, J.; and Fairweather, A. 2022. In principle obstacles for empathic AI: why we can't replace human empathy in healthcare. *AI Soc.*, 37(4): 1353–1359.
- Mun, I. B. 2025. A study of the impact of ChatGPT self-efficacy on the information seeking behaviors in ChatGPT: the mediating roles of ChatGPT characteristics and utility. *Online Inf. Rev.*, 49(2): 373–394.
- Namvarpour, M. m.; and Razi, A. 2025. The art of talking machines: A comprehensive literature review of conversational user interfaces. In *Proceedings of the 7th ACM Conference on Conversational User Interfaces*, 1–18. New York, NY, USA: ACM.
- Pan, S.; and de Graaf, M. M. A. 2025. Developing a social support framework: Understanding the reciprocity in human-chatbot relationship. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–13. New York, NY, USA: ACM.
- Pentina, I.; Hancock, T.; and Xie, T. 2023. Exploring relationship development with social chatbots: A mixed-method study of replika. *Comput. Human Behav.*, 140(107600): 107600.
- Pentina, I.; Xie, T.; Hancock, T.; and Bailey, A. 2023. Consumer-machine relationships in the age of artificial intelligence.

- gence: Systematic literature review and research directions. *Psychol. Mark.*, 40(8): 1593–1614.
- Phang, J.; Lampe, M.; Ahmad, L.; Agarwal, S.; Fang, C. M.; Liu, A. R.; Danry, V.; Lee, E.; Chan, S. W. T.; Pataranutaporn, P.; and Maes, P. 2025. Investigating affective use and emotional well-being on ChatGPT. *arXiv [cs.HC]*.
- Placani, A. 2024. Anthropomorphism in AI: hype and fallacy. *AI Ethics*, 4(3): 691–698.
- Radez, J.; Reardon, T.; Creswell, C.; Lawrence, P. J.; Evdoka-Burton, G.; and Waite, P. 2021. Why do children and adolescents (not) seek and access professional help for their mental health problems? A systematic review of quantitative and qualitative studies. *Eur. Child Adolesc. Psychiatry*, 30(2): 183–211.
- Rueger, S. Y.; Malecki, C. K.; Pyun, Y.; Aycok, C.; and Coyle, S. 2016. A meta-analytic review of the association between perceived social support and depression in childhood and adolescence. *Psychol. Bull.*, 142(10): 1017–1067.
- Santos, E. C. d., Jr; and Birdwell, T. 2025. The unspoken crisis of learning: The surging Zone of No Development. *arXiv [cs.CY]*.
- Scholte, R. H. J.; Van Lieshout, C. F. M.; and Van Aken, M. A. G. 2001. Perceived relational support in adolescence: Dimensions, configurations, and adolescent adjustment. *J. Res. Adolesc.*, 11(1): 71–94.
- Sehgal, N. K. R.; Kambhamettu, H.; Matam, S. P.; Ungar, L.; and Guntuku, S. C. 2025. Exploring Socio-cultural challenges and opportunities in designing mental health chatbots for adolescents in India. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, 1–7. New York, NY, USA: ACM.
- Seok, D.; and Doom, J. R. 2022. Adolescents’ social support networks and long-term psychosocial outcomes. *J. Soc. Pers. Relat.*, 026540752211090.
- Settersten, R. A., Jr; and Ray, B. 2010. What’s going on with young people today? the long and twisting path to adulthood. *Future Child.*, 20(1): 19–41.
- Sidoti, O. 2025. 34% of U.S. adults have used ChatGPT, about double the share in 2023. <https://www.pewresearch.org/short-reads/2025/06/25/34-of-us-adults-have-used-chatgpt-about-double-the-share-in-2023/>. Accessed: 2025-9-10.
- Sigalos, M. 2025. OpenAI’s ChatGPT to hit 700 million weekly users, up 4x from last year. <https://www.cnn.com/2025/08/04/openai-chatgpt-700-million-users.html>. Accessed: 2025-9-10.
- Skjuve, M.; Følstad, A.; and Brandtzaeg, P. B. 2023. The user experience of ChatGPT: Findings from a questionnaire study of early users. In *Proceedings of the 5th International Conference on Conversational User Interfaces*, 1–10. New York, NY, USA: ACM.
- Skjuve, M.; Følstad, A.; Fostervold, K. I.; and Brandtzaeg, P. B. 2021. My chatbot companion - a study of human-chatbot relationships. *Int. J. Hum. Comput. Stud.*, 149(102601): 102601.
- Skjuve, M.; Følstad, A.; Fostervold, K. I.; and Brandtzaeg, P. B. 2022. A longitudinal study of human–chatbot relationships. *Int. J. Hum. Comput. Stud.*, 168(102903): 102903.
- Smith, M. G.; Bradbury, T. N.; and Karney, B. R. 2025. Can generative AI chatbots emulate human connection? A relationship science perspective. *Perspect. Psychol. Sci.*, (17456916251351306): 17456916251351306.
- Song, I.; Pendse, S. R.; Kumar, N.; and De Choudhury, M. 2025. The typing cure: Experiences with Large Language Model chatbots for mental health support. *Proc. ACM Hum. Comput. Interact.*, 9(7): 1–29.
- Sorin, V.; Brin, D.; Barash, Y.; Konen, E.; Charney, A.; Nadkarni, G.; and Klang, E. 2024. Large language models and empathy: Systematic review. *J. Med. Internet Res.*, 26(1): e52597.
- Soun, R. S.; and Nair, A. 2023. ChatGPT for Mental Health Applications: A study on biases. In *The Third International Conference on Artificial Intelligence and Machine Learning Systems*, 1–5. New York, NY, USA: ACM.
- Ta, V.; Griffith, C.; Boatfield, C.; Wang, X.; Civitello, M.; Bader, H.; DeCero, E.; and Loggarakis, A. 2020. User experiences of social support from companion chatbots in everyday contexts: Thematic analysis. *J. Med. Internet Res.*, 22(3): e16235.
- Tanner, J. L. 2006. Recentring during emerging adulthood: A critical turning point in life span human development. In *Emerging adults in America: Coming of age in the 21st century*, 21–55. Washington: American Psychological Association.
- U.S. Senate Committee on the Judiciary. 2025. Examining the Harm of AI Chatbots. <https://www.judiciary.senate.gov/committee-activity/hearings/examining-the-harm-of-ai-chatbots>. Accessed: 2026-2-26.
- Vygotsky, L. S.; and Cole, M. 1978. Mind in society: Development of higher psychological processes.
- Weiss, R. S. 1974. The Provisions of Social Relationships. In Rubin, Z., ed., *Doing Unto Others*, 17–26. Englewood Cliffs, NJ: Prentice-Hall.
- Weizenbaum, J. 1966. ELIZA—a computer program for the study of natural language communication between man and machine. *Commun. ACM*, 9(1): 36–45.
- Whitehouse, C.; Wang, T.; Yu, P.; Li, X.; Weston, J.; Kulikov, I.; and Saha, S. 2025. J1: Incentivizing Thinking in LLM-as-a-Judge via Reinforcement Learning. *arXiv:2505.10320*.
- Wood, D.; Bruner, J. S.; and Ross, G. 1976. The role of tutoring in problem solving. *Journal of child psychology and psychiatry*, 17(2): 89–100.
- Xie, T.; and Pentina, I. 2022. Attachment theory as a framework to understand relationships with social chatbots: A case study of replika. In *Proceedings of the Annual Hawaii International Conference on System Sciences*. Hawaii International Conference on System Sciences.
- Young, J.; Jawara, L. M.; Nguyen, D. N.; Daly, B.; Huh-Yoo, J.; and Razi, A. 2024. The role of AI in peer support

for young people: A study of preferences for human- and AI-generated responses. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–18. New York, NY, USA: ACM.

Zhang, R.; Li, H.; Meng, H.; Zhan, J.; Gan, H.; and Lee, Y.-C. 2025. The dark side of AI companionship: A taxonomy of harmful algorithmic behaviors in human-AI relationships. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–17. New York, NY, USA: ACM.

Ziems, C.; Held, W.; Shaikh, O.; Chen, J.; Zhang, Z.; and Yang, D. 2024. Can Large Language Models Transform Computational Social Science? *Computational Linguistics*, 50(1): 237–291.