

# 000 001 002 003 004 005 006 007 008 009 010 011 012 CONSINTBENCH: EVALUATING LANGUAGE MODELS ON REAL-WORLD CONSUMER INTENT UNDERSTAND- ING

013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026  
027  
028  
029  
030  
031  
032  
033  
034  
035  
**Anonymous authors**  
036  
037  
038  
039  
040  
041  
042  
043  
044  
045  
046  
047  
048  
049  
050  
051  
052  
053

Paper under double-blind review

## ABSTRACT

Understanding human intent is a complex, high-level task for large language models (LLMs), requiring analytical reasoning, contextual interpretation, dynamic information aggregation, and decision-making under uncertainty. Real-world public discussions, such as consumer product discussions, are rarely linear or involve a single user. Instead, they are characterized by interwoven and often conflicting perspectives, divergent concerns, goals, emotional tendencies, as well as implicit assumptions and background knowledge about usage scenarios. To accurately understand such explicit public intent, an LLM must go beyond parsing individual sentences; it must integrate multi-source signals, reason over inconsistencies, and adapt to evolving discourse, similar to how experts in fields like politics, economics, or finance approach complex, uncertain environments. Despite the importance of this capability, no large-scale benchmark currently exists for evaluating LLMs on real-world human intent understanding, primarily due to the challenges of collecting real-world public discussion data and constructing a robust evaluation pipeline. To bridge this gap, we introduce CONSINT-BENCH, the first dynamic, live evaluation benchmark specifically designed for intent understanding, particularly in the consumer domain. CONSINT-BENCH is the largest and most diverse benchmark of its kind, supporting real-time updates while preventing data contamination through an automated curation pipeline. We evaluate 20 LLMs, spanning both open-source and closed-source models, across four core dimensions of consumer intent understanding: *depth*, *breadth*, *informativeness*, and *correctness*. Our benchmark provides a comprehensive and evolving evaluation standard for assessing LLM performance in understanding complex, real-world human intent, with the ultimate goal of advancing LLMs toward expert-level reasoning and analytical capabilities.

## 1 INTRODUCTION

The advent of Large Language Models (LLMs) OpenAI (2025); Grattafiori et al. (2024); Guo et al. (2025) has fundamentally transformed artificial intelligence, shifting from text generation to the ability to understand and reason about complex, real-world human intent Team (2025). These models demonstrate exceptional performance across a wide range of tasks Brown et al. (2020); Ouyang et al. (2022); Achiam et al. (2023); Chowdhery et al. (2023); Touvron et al. (2023); Google (2024). To assess LLMs in real-world problem-solving contexts, several benchmarks have been proposed. For example, SWE-bench Jimenez et al. (2024b) evaluates LLMs' ability to resolve software issues using GitHub repositories, while SPIDER2.0 Lei et al. (2024) focuses on enterprise-level text-to-SQL workflows. GAIA Mialon et al. (2023) introduces multi-modal, tool-augmented queries that require reasoning and web-browsing capabilities, and FutureX Zeng et al. (2025) challenges LLMs with future event prediction tasks. These benchmarks reflect a growing trend toward evaluating LLMs in dynamic, context-rich environments, aligning with more complex real-world applications.

Despite these advances, the question of whether LLMs truly understand public, swarm-like intent intelligence and the deeper, abstract aspects of human intent remains largely unexplored. Real-world human perspectives, whether in consumer decision-making, team collaboration, or online community discussions, are inherently multifaceted. They involve not only knowledge comprehension but




Figure 1: Overview of CONSINT-BENCH, a large-scale and live benchmark designed to evaluate real-world consumer intent understanding.

also a complex interplay of perspectives, needs, emotions, and implicit assumptions. Existing work often simulates individual user preferences but fails to account for the complex interactions of intentions, perspectives, and emotions across multiple users. To address this, a more comprehensive evaluation is needed—one that goes beyond individual perspectives to synthesize and aggregate intentions from multiple users. This requires the ability to map structured intention graphs that capture the multifaceted nature of human discourse.

Several benchmarks have focused on specific aspects of human intent, such as SocialIQA Sap et al. (2019) for social intent and commonsense reasoning, and TOMI Le et al. (2019) for Theory of Mind capabilities. However, these benchmarks rely on hand-crafted or semi-synthetic data, lacking the noise, redundancy, and subtext inherent in real-world discussions, which leads to evaluation processes that do not align with actual application scenarios. IFEVAL Zhou et al. (2023) evaluates instruction-following ability, while SociaBench Chen et al. (2024a) and AgentSense Mou et al. (2025) focus on intent understanding, generation quality, and social intent navigation. Emotion-Queen Chen et al. (2024b) addresses implicit emotions in human intent, and URS-bench Wang et al. (2024) evaluates LLMs' responses to factual question answering, problem-solving, and advice. However, these frameworks mainly focus on lower-level aspects of human intent, such as instruction-following, social reasoning, or emotional understanding, and lack an analysis of deeper dimensions of intent. In real-world scenarios, human intent is multifaceted and dynamic, involving social, emotional, and practical factors that require the fusion and resolution of conflicting viewpoints—key components of human reasoning. Existing frameworks, however, fail to evaluate these primary dimensions of LLM performance.

To address these limitations, we propose CONSINT-BENCH, a comprehensive, dynamic, and live benchmark designed to evaluate LLM performance in understanding real-world human intent, particularly in consumer domains. CONSINT-BENCH spans nine primary consumer domains, ranging from personal care and AI products to daily necessities, covering 54 sub-categories and over 1,400 product discussions sourced from real-world user interactions. For each product, we collect approximately 200 user comments, aggregating over 200k opinions. We evaluate LLMs' intent understanding ability across four primary dimensions: depth, breadth, correctness, and informativeness. Depth is further defined across five hierarchical levels (L1–L5), where the first three levels capture content directly from user discussions, while the last two require the model to reason based on internal knowledge and context, necessitating a deeper understanding of human intent. Addi-

108 tionally, we construct a robust evaluation pipeline to ensure accurate assessment and mitigate LLM  
 109 bias and hallucination. We conduct extensive experiments on a variety of LLMs, including both  
 110 closed- and open-source models, as well as reasoning and general models. Our results reveal that  
 111 reasoning models outperform general models in depth, breadth, and correctness. However, a sig-  
 112 nificant gap remains between closed-source and open-source models. Furthermore, even the most  
 113 advanced models struggle with deep and broad intent understanding, highlighting substantial room  
 114 for improvement.

115 In summary, our contributions are as follows:

- 117 • We introduce CONSINT-BENCH, a large-scale benchmark for real-world consumer in-  
 118 tent, consisting of over 200k product-level discussions spanning 9 major domains, 54 sub-  
 119 domains, and 1400+ products. Each topic includes an average of 200 discussion entries,  
 120 ensuring information density and diversity.
- 121 • We define four primary aspects and implement a robust evaluation pipeline to mitigate  
 122 bias and hallucination. Specifically, we construct CONSINT-TREE to assess LLMs’ depth  
 123 and breadth of intent understanding, use CONSINT-RAG for evaluating correctness, and  
 124 measure informativeness through lexical diversity and semantic richness.
- 125 • We conduct extensive experiments on both closed-source and open-source models of vary-  
 126 ing sizes (1.5B to 72B parameters). The results show that even the most advanced models  
 127 struggle with deep and broad intent understanding, highlighting significant potential for  
 128 improvement.

## 130 2 BENCHMARK CONSTRUCTION

### 132 2.1 DATA CURATION

134 To retrieve and organize discussions from diverse online sources, we construct an automated system  
 135 for collecting consumer discussions. The data collection pipeline integrates three stages:

136 **Search.** We employ a combination of vector search and API search to maximize the retrieval of  
 137 relevant discussions. Vector search analyzes the semantic similarity between the user’s input and  
 138 available discussions, retrieving those with a similarity score above a defined threshold. API search  
 139 utilizes LLMs to generate relevant keywords from the user’s input, enhancing the search process.

140 **Retrieving.** After the search, the system retrieves relevant discussions from open-source web  
 141 sources in real-time. These results are then used as raw input for the subsequent cleaning and  
 142 filtering stages.

143 **Cleaning.** The collected results are filtered through both rule-based and LLM-based quality checks.  
 144 First, discussions with titles and content shorter than 20 characters are discarded as low-quality.  
 145 Second, discussions deemed irrelevant to the search topic by the LLM are excluded. Additionally,  
 146 the cleaning process incorporates recency by considering time-based factors to ensure the relevance  
 147 of the discussions.

148 This multi-step pipeline efficiently collects high-quality, contextually relevant discussions, ensuring  
 149 the data is well-suited for subsequent analysis.

### 151 2.2 DATA STATISTICS

153 As shown in Table 1 and Figure 2, CONSINT-BENCH spans 54 sub-categories and includes over  
 154 1,400 product discussions sourced from real-world discussions. For each product, we collect ap-  
 155 proximately 200 user comments, aggregating over 200k opinions across categories such as personal  
 156 care, AI products, and daily necessities. This rich dataset forms a robust foundation for evaluating  
 157 LLMs’ ability to understand human intent in diverse real-world contexts.

### 159 2.3 DIMENSION CATEGORIES

161 To thoroughly evaluate the capabilities of LLMs in understanding consumer intent, we categorize  
 162 the evaluation into four primary dimensions: depth, breadth, informativeness, and correctness:

162 Table 1: Comparison with existing related benchmarks. "Real-world" indicates whether the data is  
 163 sourced from real-world scenarios rather than synthetic or online existing resources. "Live Update"  
 164 denotes whether the benchmark can be regularly updated.

Benchmark	Domain	Tasks	Real World	Live Update
DABstep Egg et al. (2025)	Data Science	450	✓	✗
FutureX Zeng et al. (2025)	Future Prediction	500/week	✓	✓
GAIA Mialon et al. (2023)	General QA	466	✓	✗
OSWorld Xie et al. (2024)	Computer Use	369	✓	✗
OPT-Bench Li et al. (2025)	Iterative Optimization	30	✓	✗
Spider2.0 Lei et al. (2024)	Text-to-SQL	632	✓	✗
SWE-Bench Jimenez et al. (2024a)	Code	2,294	✓	✗
SociaBench Chen et al. (2024a)	Social Intent	6,000	✗	✗
URS-bench Wang et al. (2024)	Intent Understanding	1,846	✗	✗
CONSINT-BENCH (ours)	Consumer Intent	1,475	✓	✓

178 **Depth** measures the model’s ability to analyze and provide insights into a given discussion, specifically  
 179 evaluating how well it can explore complex ideas and offer comprehensive explanations.  
 180 We define five levels of depth (L1–L5), where L1–L3 represent basic understanding—such as  
 181 identifying usage scenarios, discussing product aspects, and capturing the user’s feelings, all of  
 182 which can be directly derived from the user discussion. Levels L4–L5 represent advanced com-  
 183 prehension, involving tasks like making comparisons with previous versions or similar products,  
 184 analyzing their advantages and disadvantages, and speculating on the product’s future direction.  
 185 A higher depth score indicates a more profound and comprehensive understanding of the topic.


186 **Breadth** evaluates the model’s ability to ad-  
 187 dress a wide range of subtopics within a broader  
 188 subject area. This dimension focuses on the  
 189 model’s versatility and capacity to cover vari-  
 190 ous aspects of the topic, such as different usage  
 191 scenarios, product versions, and product fea-  
 192 tures. A higher breadth score indicates a more  
 193 comprehensive understanding of the topic, as  
 194 the model effectively spans a wider array of re-  
 195 lated issues.

196 **Informativeness** evaluates how effectively the  
 197 model conveys content while maintaining its  
 198 core message, reflecting its ability to provide  
 199 meaningful information without unnecessary  
 200 repetition or relying on a single paradigm.  
 201 We assess informativeness through lexical rich-  
 202 ness and semantic redundancy, measuring the  
 203 model’s capacity to eliminate redundant or  
 204 repetitive information. A lower redundancy  
 205 score indicates a more focused and informative  
 206 understanding.

207 **Correctness** evaluates whether the LLM’s un-  
 208 derstanding of consumer intent is free from bias  
 209 or hallucinations, and whether the responses ac-  
 210 curately reflect the true opinions and sentiments expressed in the original discussions. A higher  
 211 correctness score indicates a more accurate and reliable response, ensuring that the model’s output  
 212 aligns closely with the original human intentions.

### 213 3 METHODOLOGY

214 To ensure a fair and actionable evaluation process, we adopt a paradigm where the LLM self-  
 215 generates both the question and the answer. The question represents a summarized intent from



216 Figure 2: Overview of CONSINT-BENCH: It in-  
 217 cludes over 200k product-level discussions across  
 218 9 major domains, 54 sub-domains, and more than  
 219 1,400 products.




Figure 3: The overall pipeline of CONSINT-BENCH, covering data curation and evaluation. Data Curation: We provide two methods (keyword and semantic) to search and retrieve product discussions daily. It consists of over 200k product-level discussions across 9 major domains, 54 sub-domains, and more than 1,400 products. The benchmark focuses on four primary dimensions and evaluates intent depth across five hierarchical difficulty levels. After that, rule-based and LLM-based filtering is applied to remove irrelevant discussions and retain relevant ones. Evaluation: We propose CONSINT-TREE for more accurate assessment of depth and breadth dimensions, and CONSINT-RAG for correctness. For informativeness, we assess lexical richness and semantic redundancy.

the original discussion, and the answer reflects the LLM’s judgment of the intent. This approach is repeated with the same number of questions to assess the model’s ability to capture the most valuable majority perspective. To comprehensively assess the ability of LLMs to understand complex consumer intentions, we propose a robust evaluation methodology that combines rule-based analysis with the LLM-as-a-judge mechanism for multi-dimensional evaluation.

**CONSINT-TREE: Depth and Breadth Evaluation.** CONSINT-TREE is a tree-structured knowledge graph derived from real-world consumer discussions. Each question in the generated questionnaire is mapped to a corresponding node in CONSINT-TREE, forming a subtree. The size and structure of this subtree quantify the breadth and depth of the LLM’s understanding.

**CONSINT-RAG: Correctness Evaluation.** CONSINT-RAG is a retrieval-augmented generation pipeline designed to mitigate hallucinations and bias caused by the noisy nature of real-world discussions. Each questionnaire question is paired with a reference answer, and the CONSINT-RAG pipeline verifies the accuracy of these answers, assessing the correctness of the LLM’s intent comprehension.

**Informativeness Evaluation.** To assess informativeness, we compute the lexical richness and semantic redundancy of the generated questionnaire. These metrics capture the diversity and specificity of the LLM’s expressions, reflecting the richness of its understanding of consumer intent.

### 3.1 CONSINT-TREE CONSTRUCTION AND EVALUATION

**Construction** To comprehensively evaluate the ability of LLMs to understand large-scale real-world data, which contains a massive volume of similar and conflicting content, in both depth and breadth, we propose the CONSINT-TREE —a five-level weighted hierarchical tree with weights based on discussion popularity. As illustrated in Figure 3, the root node represents the focal product under

270 discussion. In terms of depth, nodes across Levels 1–3 capture the product’s usage scenarios, as-  
 271 pects, and user experience (e.g., usage feelings) as reflected in consumer discussions. Levels 4–5  
 272 deepen the understanding of consumer intent: Level 4 nodes represent competing products influen-  
 273 cing consumer sentiments, while Level 5 nodes indicate potential improvement tendencies derived  
 274 from these sentiments. In the CONSINT-TREE, a path from the root node to a leaf node defines  
 275 a “branch,” representing a progressively deepening user opinion. In terms of breadth, sufficient  
 276 number of branches represent the multiple facets of the product’s discussion. As for discussion pop-  
 277 ularity, nodes corresponding to high-frequency discussions or content with high upvotes/views are  
 278 assigned higher weights. See the Appendix for detailed extraction of branches from each consumer  
 279 discussion and the details of how related discussions are aggregated into high-weight nodes.

280 **Evaluation** To assess the depth and breadth of a LLM’s understanding of consumer intent, content  
 281 from questionnaires will be extracted into branches to lighten to CONSINT-TREE. Each branch  
 282 will undergo semantic matching with the nodes in the CONSINT-TREE from top to bottom using a  
 283 Sentence Transformer. Nodes that are successfully matched will be marked as “lightened,” and the  
 284 lightened nodes in the CONSINT-TREE will form a subtree. For the depth dimension, the depth  
 285 score at each level (from L1 to L5) is calculated as the percentage of the total weight of lightened  
 286 nodes in the subtree at that level relative to the total weight of all nodes in the original CONSINT-  
 287 TREE at the same level. The overall depth score is computed as the average of the depth scores  
 288 across all five levels. For the breadth dimension, the breadth score as the sum of the weights of all  
 289 lightened nodes in the subtree. A higher depth score indicates that the questions in the questionnaire  
 290 can delve into more profound layers of the consumers intent. A higher breadth score reflects a more  
 291 comprehensive understanding of the consumers intent.

### 292 3.2 CONSINT-RAG CONSTRUCTION AND EVALUATION

293 **Construction** To accurately evaluate the correctness dimension of LLMs in understanding con-  
 294 sumer intent while mitigating LLM judge bias and hallucination, we propose CONSINT-RAG.  
 295 This approach retrieves consumer preferences related to the LLM’s inferred intent from the original  
 296 discussions to serve as the ground truth. CONSINT-RAG follows a two-stage process: embedding  
 297 and retrieval. In the embedding stage, user discussions are transformed into vector representations  
 298 using TF-IDF and all-MiniLM-L6-v2. This dual representation enables the retrieval process to cap-  
 299 ture both precise keyword-level matches and deeper semantic information. In the retrieval stage,  
 300 keywords are extracted from the LLM’s questionnaire questions. These extracted key opinions and  
 301 full questions are vectorized and jointly searched across two vector databases to retrieve the top-k  
 302 most relevant discussions.

303 **Evaluation** After retrieval, the top-k most relevant discussions are analyzed to reflect human opin-  
 304 ions and compared to the answers provided by the LLM during questionnaire generation. However,  
 305 due to the implicit, noisy, and multi-opinion nature of real-world discussions, direct answer match-  
 306 ing is not feasible. Therefore, further reasoning is required to determine the consensus opinion,  
 307 reflecting the majority perspective. Based on this reasoning, the final answer is generated from the  
 308 previous RAG results. The accuracy of the LLM’s original answers for a given questionnaire is then  
 309 used as the correctness evaluation metric.

### 310 3.3 INFORMATIVENESS

311 To evaluate the informativeness of LLMs in understanding consumer intent, we quantify Lexical  
 312 Richness and Semantic Redundancy using informatics formulas.

313 **Lexical Richness:** Lexical richness is assessed across two dimensions: words and phrases. It re-  
 314 flects the LLM’s ability to capture a broader range of consumer intent topics and diverse question  
 315 formulations, thereby demonstrating a more comprehensive understanding of intent. Additionally,  
 316 a more precise and nuanced expression of intent contributes to greater lexical richness. Type-Token  
 317 Ratio (TTR) Johnson (1944) is used to evaluate word richness, while Distinct-n Li et al. (2016)  
 318 measures phrase richness. Detailed metric calculations are provided in Appendix A.3.

319 **Semantic Redundancy:** Semantic redundancy is evaluated by assessing the embedded vector simi-  
 320 larity between questionnaire questions, which helps gauge the LLM’s ability to identify and structure  
 321 consumer intent effectively. High semantic redundancy within the questionnaire indicates that the

324 Table 2: Performance of reasoning LLMs, general LLMs, and open-source LLMs on CONSINT-  
 325 BENCH, with the best performance highlighted in **bold**.

Model	Depth						Breadth	Informativeness		Correctness
	L1	L2	L3	L4	L5	Overall		Lexical	Semantic $\downarrow$	
<i>Proprietary LLMs</i>										
GPT-5	18.29	<b>25.81</b>	<b>6.25</b>	<b>3.49</b>	0.06	10.78	<b>53.48</b>	80.21	62.75	62.65
GPT-4.1	20.97	25.03	5.90	3.37	0	11.01	53.41	79.07	63.82	59.05
GPT-4o	20.99	23.70	4.90	2.58	0.05	10.44	52.95	79.56	62.86	75.75
Claude-3.5-sonnet	19.25	23.25	5.34	2.95	0	10.16	52.83	73.94	61.11	53.35
GPT-o3	16.17	22.43	5.69	3.18	<b>0.07</b>	9.51	52.73	<b>85.52</b>	<b>52.27</b>	<b>80.35</b>
<i>Open-Source LLMs</i>										
Qwen3-30B-A3B	<b>25.01</b>	23.66	5.43	2.51	0.06	<b>11.33</b>	53.20	70.58	68.94	61.60
DS-Distill-Qwen-14B	17.00	25.56	6.12	3.30	0	10.40	53.32	67.47	75.53	58.45
Qwen2.5-32B-Instruct	19.26	23.63	5.31	2.89	0.01	10.21	52.46	65.72	74.60	54.95
Qwen3-32B	20.77	23.46	5.57	2.82	0	10.52	51.15	65.84	68.16	55.26
Qwen3-8B	15.95	22.43	4.89	2.39	0.01	9.13	50.58	57.51	81.25	50.42
Qwen2.5-72B-Instruct	18.89	22.27	5.73	3.10	0	10.00	50.52	54.63	77.49	64.11
DS-Distill-Qwen-32B	16.60	24.56	5.90	3.30	0	10.07	50.34	59.30	76.58	53.90
Qwen2.5-14B-Instruct	13.56	22.23	5.45	2.87	0.02	8.83	48.27	52.39	80.06	60.88
LLama3.2-8B-Instruct	13.88	19.75	5.62	2.73	0	8.40	47.91	47.87	88.25	52.31
Qwen2.5-7B-Instruct	11.87	19.73	4.16	1.97	0	7.54	47.43	43.58	85.07	49.24
Internlm3-8B-Instruct	11.07	20.76	4.87	2.61	0.03	7.87	45.91	49.83	75.51	51.67
LLama3.1-8B-Instruct	11.23	19.46	5.53	2.91	0	7.83	45.41	42.36	88.00	52.67
Qwen2.5-3B-Instruct	13.49	18.63	4.22	2.09	0	7.69	42.73	39.32	79.35	35.43
Qwen2.5-1.5B-Instruct	2.83	4.94	0.99	0.45	0	1.84	14.31	4.56	87.65	36.90
DS-Distill-Qwen-7B	1.80	4.91	1.35	0.55	0	1.72	11.54	3.30	73.25	13.40

348 LLM’s logical reasoning approach may be overly simplistic or that it has failed to capture the full  
 349 spectrum of consumer intent. Redundancy is calculated as the average maximum similarity between  
 350 each question and all other questions Chen et al. (2021b). Detailed metric calculations are provided  
 351 in Appendix A.3.

## 4 EXPERIMENT

### 4.1 EXPERIMENTAL SETUP

358 We evaluated our method across a diverse set of LLMs, including both proprietary and open-source  
 359 models, each consisting of reasoning and general models. The proprietary models include OpenAI’s  
 360 GPT family and Claude, all accessed via their APIs. For open-source models, we consider the Qwen  
 361 series (ranging from 1.5B to 72B), LLaMA, DeepSeek and InternLM, all deployed locally using the  
 362 LMDeploy framework.

### 4.2 MAIN RESULTS

363 Table 2 evaluates the performance of 20 LLMs in understanding consumer intent across four key  
 364 dimensions:

365 **1) Depth:** The scores across L1–L5 generally show a downward trend, reflecting the increasing  
 366 difficulty of capturing deeper aspects of consumer intent. GPT-5 and GPT-4.1 achieve the highest  
 367 overall depth scores, ranking first and second, respectively, in L2 and L3, demonstrating a com-  
 368 prehensive understanding of both contextual elements and user sentiment. As a reasoning model,  
 369 GPT-o3 excels in L5, highlighting the role of reasoning in deepening the understanding of consumer  
 370 intent. Among open-source LLMs, Qwen3-30B-A3B, a Mixture of Experts (MOE) model, performs  
 371 best, benefiting from its ability to allocate specialized experts for different depths of understanding.

372 **2) Breadth:** GPT-5 leads in breadth, showing its ability to address a wide range of consumer intent.  
 373 In open-source models, Deepseek-R1-Distill-Qwen-14B demonstrates strong coverage of diverse  
 374 subtopics. However, smaller models such as Qwen2.5-1.5B-Instruct struggle to capture the full  
 375 breadth of consumer intent.

378  
379  
380  
Table 3: Comparison of reasoning LLMs, general LLMs, and open-source LLMs using CONSINT-  
381 TREE on CONSINT-BENCH, with the best performance highlighted in **bold**.  
382

Model	Depth						Breadth	Informativeness		Correctness
	L1	L2	L3	L4	L5	Overall		Lexical	Semantic↓	
GPT-o3 (wo/Tree)	16.17	22.43	5.69	3.18	0.07	9.51	52.73	85.52	<b>52.27</b>	<b>80.35</b>
GPT-o3 (w/Tree)	<b>45.13</b>	<b>36</b>	<b>15</b>	<b>11.74</b>	<b>1.37</b>	<b>21.95</b>	<b>59.16</b>	<b>72.47</b>	71.05	57.10
	(+28.96)	(+13.57)	(+9.31)	(+8.56)	(+1.30)	(+12.44)	(+6.43)	(-13.05)	(-0.22)	(-23.25)
GPT-4o (wo/Tree)	20.99	23.70	4.90	2.58	0.05	10.44	52.95	79.56	62.86	75.60
GPT-4o (w/Tree)	37.38	31.39	12.78	9.60	0.85	18.40	57.00	70.77	72.36	64.15
	(+16.39)	(+7.69)	(+7.88)	(+7.02)	(+0.80)	(+7.96)	(+4.05)	(-8.79)	(-0.50)	(-11.45)
Qwen2.5-7B (wo/Tree)	11.87	19.73	4.16	1.97	0	7.54	47.43	43.58	85.07	42.15
Qwen2.5-7B (w/Tree)	32.39	29.78	11.29	7.96	0.54	16.39	49.28	42.33	79.63	49.24
	(+20.52)	(+10.05)	(+7.13)	(+6.00)	(+0.54)	(+8.85)	(+1.85)	(-1.25)	(-5.44)	(+7.09)

387  
388  
389  
390  
391  
392  
393  
394  
395  
396  
**3) Informativeness:** GPT-o3 outperforms all models in lexical richness and minimal semantic redundancy, indicating that its deeper understanding of consumer intent is supported by a broader vocabulary and refined semantic expression. Open-source reasoning models, while competitive in depth and breadth, generally lag behind proprietary models in lexical and semantic richness. Additionally, compared to reasoning open-source models, general open-source models such as Qwen2.5-72B-Instruct still struggle with understanding intent.397  
398  
399  
**4) Correctness:** GPT-o3 achieves the highest correctness score, underscoring the superior ability of reasoning-focused models to accurately summarize and derive consumer intent from large-scale discussions.400  
401  
402  
403  
404  
In summary, smaller open-source general models tend to underperform across all four dimensions, particularly in deeper reasoning (L5 depth) and expressive capabilities. Reasoning LLMs, with their advanced reasoning abilities, outperform in multiple metrics, emphasizing the critical role of reasoning in improving the depth, breadth, and correctness of consumer intent comprehension.405  
406  

### 4.3 ABLATION STUDY

407  
408  
409  
410  
We conducted additional experiments to explore whether noisy and low-quality information in large-scale real-world discussions limits the understanding capabilities of large language models (LLMs). In these experiments, we replace the original real-world discussions with CONSINT-TREE for LLM evaluation. The results are presented in Table 3.411  
412  
413  
414  
415  
416  
417  
418  
419  
420  
421  
422  
423  
424  
In the *Depth* and *Breadth* dimensions, the LLM (w/Tree) outperformed the LLM (wo/Tree). This indicates that the tree significantly enhances the LLM’s ability to understand high-weight, high-interest consumer intents, improving the overall depth and breadth of intent comprehension. However, in the *Informativeness* dimension, a decline was observed when comparing LLM (w/Tree) to LLM (wo/Tree). This suggests that when LLMs process the refined intents extracted from the tree, their ability to understand these intents in a nuanced manner is constrained. This may be due to the fact that high-weight branches often focus on more popular aspects, leading to a reduction in semantic richness and overall informativeness. In the *Correctness* dimension, LLM (w/Tree) demonstrated a decreasing trend. Although the tree refines noisy and irrelevant discussions, it may simultaneously lose some critical information, resulting in a failure to provide a comprehensive representation of the relevant topics. In contrast, for open-source small models, LLM (w/Tree) effectively reduces noise in the consumer intents, leading to improved understanding correctness. This suggests that small models are more sensitive to real-world noise and may struggle to properly understand human intent in noisy environments.425  
426  

### 4.4 FURTHER DISCUSSION

427  
428  
429  
430  
431  
The results in Table 2 show that open-source reasoning models such as Qwen3-30B-A3B outperform close-source model in both depth and breadth. To explore this further, we conduct a case study that presents the performance of GPT-5, GPT-o3, and Qwen3-30B-A3B on CONSINT-BENCH in understanding consumer intent, specifically from discussions about the Google Nest Smart Speaker. GPT-5 achieves the highest breadth score by covering more high-weight nodes, while GPT-o3 uniquely excels in L5 depth. Although the breadth scores for all three models are comparable, Qwen3-30B-

432 A3B lags notably in informativeness. Notably, the question stems and options in Qwen3-30B-A3B’s  
 433 questionnaires are longer on average compared to those generated by GPT-5 and GPT-o3. This sug-  
 434 gests that closed-source LLMs tend to use more refined and precise vocabulary and semantic struc-  
 435 tures when understanding consumer intent, highlighting their superior control over finer details. The  
 436 detailed results are shown in A.4.

## 437 5 RELATED WORK

441 **LLM Evaluation** The rapid advancements in Large Language Models (LLMs) have led to the cre-  
 442 ation of numerous benchmarks to evaluate their generalization and reasoning capabilities. Early  
 443 efforts, such as MMLU Hendrycks et al. (2020) and BIG-bench Srivastava et al. (2022), provided  
 444 broad assessments of general knowledge and reasoning skills. Subsequent benchmarks focused on  
 445 more specific domains, including linguistic and commonsense reasoning (e.g., GLUE Wang et al.  
 446 (2018), SuperGLUE Wang et al. (2019), CommonsenseQA Talmor et al. (2019), HellaSwag Zellers  
 447 et al. (2019), TruthfulQA Lin et al. (2022)), mathematical and programming reasoning (e.g.,  
 448 MATH Hendrycks et al. (2021), GSM8K Cobbe et al. (2021), HumanEval Chen et al. (2021a),  
 449 MBPP Austin et al. (2021)), and task-based agent evaluation (e.g., MLE-bench Chan et al. (2024)  
 450 for ML engineering, OSWorld Xie et al. (2024) for GUI tasks, and OPT-BENCH Li et al. (2025)  
 451 for complex optimization). Despite these advancements, there remains a lack of systematic eval-  
 452 uation regarding whether LLMs can effectively understand human intent in dynamic, real-world  
 453 decision-making contexts—particularly those involving multi-user perspectives, emotional nuance,  
 454 and evolving goals. To address this gap, we introduce CONSINT-BENCH, a large-scale benchmark  
 455 designed to evaluate LLMs’ ability to comprehend and reason about human intentions in complex,  
 real-world scenarios.

456 **LLM Human Intent Evaluation** Human intent evaluation has increasingly focused on under-  
 457 standing human-centric intent in complex, dynamic real-world scenarios. Benchmarks such as  
 458 SocialIQA Sap et al. (2019) have emphasized social intent and commonsense reasoning, while  
 459 TOMI Le et al. (2019) evaluates LLMs’ Theory of Mind capabilities. Several benchmarks have  
 460 assessed LLMs’ ability to understand and follow human intent. For example, IFEVAL Zhou et al.  
 461 (2023) primarily evaluates instruction-following ability, while SociaBench Chen et al. (2024a) and  
 462 AgentSense Mou et al. (2025) assess intent understanding, generation quality, and social intent  
 463 navigation. EmotionQueen Chen et al. (2024b) focuses on evaluating implicit emotions in human  
 464 intent, and URS-bench Wang et al. (2024) evaluates LLMs’ responses to factual question answer-  
 465 ing, problem-solving, and advice. However, these frameworks typically focus on specific aspects of  
 466 human intent, such as instruction-following, social reasoning, or emotional understanding. In real-  
 467 world scenarios, human intent is often multifaceted and dynamic, involving a combination of social,  
 468 emotional, and practical factors. As a result, no existing framework provides a comprehensive eval-  
 469 uation of whether LLMs can truly understand human reasoning and mental states. To fill this gap,  
 470 we propose CONSINT-BENCH, a benchmark designed to evaluate LLMs’ ability to understand  
 dynamic and complex real-world human intent.

## 471 6 CONCLUSION

472 In this work, we propose CONSINT-BENCH, a comprehensive benchmark consisting of over 200k  
 473 product-level discussions across 9 major domains, 54 sub-domains, and over 1,400 products, de-  
 474 signed to evaluate the performance of Large Language Models (LLMs) in understanding real-world  
 475 human intent, particularly within consumer domains. Our evaluation framework measures LLMs’  
 476 ability to comprehend intent across four key dimensions: depth, breadth, correctness, and infor-  
 477 mative ness. We implement a robust evaluation pipeline to mitigate bias and hallucinations. Specifically,  
 478 we construct CONSINT-TREE to assess LLMs’ depth and breadth of intent understanding, use  
 479 CONSINT-RAG for evaluating correctness, and measure informativeness through lexical diversity  
 480 and semantic richness. Through extensive experiments on both closed-source and open-source mod-  
 481 els, we demonstrate that reasoning models outperform general models on average. However, sig-  
 482 nificant gaps remain between closed-source and open-source models, and even the most advanced  
 483 models struggle with deep and broad intent understanding. Our mission is to advance LLMs toward  
 484 expert-level reasoning and improve their ability to understand complex real-world intent.

486 REPRODUCIBILITY STATEMENT  
487488 We adhere to the reproducibility guidelines outlined in the ICLR 2026 author guidelines. All data  
489 and code necessary to reproduce our results will be open-sourced and made available as soon as  
490 possible.491  
492 ETHICS STATEMENT  
493494 The CONSINT-BENCH dataset was constructed from public available websites, and all privacy-  
495 sensitive personal information has been removed during the data curation process. To mitigate  
496 the potential for technology misuse, the benchmark will be released under a restrictive license for  
497 academic research purposes only, explicitly prohibiting malicious applications.498  
499 REFERENCES500 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Ale-  
501 man, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical  
502 report. *arXiv preprint arXiv:2303.08774*, 2023.503 Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan,  
504 Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language  
505 models. *arXiv preprint arXiv:2108.07732*, 2021.506 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,  
507 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are  
508 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.509 Jun Shern Chan, Neil Chowdhury, Oliver Jaffe, James Aung, Dane Sherburn, Evan Mays, Giulio  
510 Starace, Kevin Liu, Leon Maksin, Tejal Patwardhan, et al. Mle-bench: Evaluating machine learn-  
511 ing agents on machine learning engineering. *arXiv preprint arXiv:2410.07095*, 2024.512 Hongzhan Chen, Hehong Chen, Ming Yan, Wenshen Xu, Gao Xing, Weizhou Shen, Xiaojun Quan,  
513 Chenliang Li, Ji Zhang, and Fei Huang. Socialbench: Sociality evaluation of role-playing con-  
514 versational agents. In *ACL (Findings)*, 2024a.515 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique de Oliveira Pinto, Jared Kaplan,  
516 Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language  
517 models trained on code. *arXiv preprint arXiv:2107.03374*, 2021a.518 Wang Chen, Piji Li, and Irwin King. A training-free and reference-free summarization evalua-  
519 tion metric via centrality-weighted relevance and self-referenced redundancy. In *Proceedings of*  
520 *the 59th Annual Meeting of the Association for Computational Linguistics and the 11th Interna-  
521 tional Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 404–414.  
522 Association for Computational Linguistics, 2021b. doi: 10.18653/v1/2021.acl-long.34. URL  
523 <https://aclanthology.org/2021.acl-long.34/>.524 Yuyan Chen, Hao Wang, Songzhou Yan, Sijia Liu, Yueze Li, Yi Zhao, and Yanghua Xiao. Emo-  
525 tionqueen: A benchmark for evaluating empathy of large language models, 2024b. URL  
526 <https://arxiv.org/abs/2409.13359>.527 Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam  
528 Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm:  
529 Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):  
530 1–113, 2023.531 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,  
532 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to  
533 solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.534 Alex Egg, Martin Iglesias Goyanes, Friso Kingma, Andreu Mora, Leandro von Werra, and Thomas  
535 Wolf. Dabstep: Data agent benchmark for multi-step reasoning. *arXiv preprint arXiv:2506.23719*,  
536 2025.

540 Google. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context.  
 541 *arXiv preprint arXiv:2403.05530*, 2024.

542

543 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad  
 544 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd  
 545 of models. *arXiv preprint arXiv:2407.21783*, 2024.

546

547 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,  
 548 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms  
 549 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

550

551 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and  
 552 Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint*  
 553 *arXiv:2009.03300*, 2020.

554

555 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,  
 556 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv*  
 557 *preprint arXiv:2103.03874*, 2021.

558

559 Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik R  
 560 Narasimhan. Swe-bench: Can language models resolve real-world github issues? In *ICLR*, 2024a.

561

562 Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik R  
 563 Narasimhan. SWE-bench: Can language models resolve real-world github issues? In *The Twelfth  
 564 International Conference on Learning Representations*, 2024b. URL <https://openreview.net/forum?id=VTF8yNQM66>.

565

566 Wendell Johnson. The type-token ratio: A measure of lexical diversity. *Language*, 20(3):169–187,  
 567 1944.

568

569 Matthew Le, Y-Lan Boureau, and Maximilian Nickel. Revisiting the evaluation of theory of mind  
 570 through question answering. In *Proceedings of the 2019 Conference on Empirical Methods in  
 571 Natural Language Processing and the 9th International Joint Conference on Natural Language  
 572 Processing (EMNLP-IJCNLP)*, pp. 5872–5877, 2019.

573

574 Fangyu Lei, Jixuan Chen, Yuxiao Ye, Ruisheng Cao, Dongchan Shin, Hongjin Su, Zhaoqing Suo,  
 575 Hongcheng Gao, Wenjing Hu, Pengcheng Yin, et al. Spider 2.0: Evaluating language models on  
 576 real-world enterprise text-to-sql workflows. *arXiv preprint arXiv:2411.07763*, 2024.

577

578 Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. A diversity-promoting  
 579 objective function for neural conversation models. In *Proceedings of NAACL-HLT*, 2016.

580

581 Xiaozhe Li, Jixuan Chen, Xinyu Fang, Shengyuan Ding, Haodong Duan, Qingwen Liu, and Kai  
 582 Chen. Opt-bench: Evaluating llm agent on large-scale search spaces optimization problems.  
 583 *arXiv preprint arXiv:2506.10764*, 2025.

584

585 Stephanie Lin, Jacob Hilton, and Owain Evans. Truthfulqa: Measuring how models mimic human  
 586 falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational  
 587 Linguistics*, pp. 3214–3252, 2022.

588

589 Grégoire Mialon, Clémentine Fourrier, Thomas Wolf, Yann LeCun, and Thomas Scialom. Gaia:  
 590 a benchmark for general ai assistants. In *The Twelfth International Conference on Learning  
 591 Representations*, 2023.

592

593 Xinyi Mou, Jingcong Liang, Jiayu Lin, Xinnong Zhang, Xiawei Liu, Shiyue Yang, Rong Ye, Lei  
 594 Chen, Haoyu Kuang, Xuan-Jing Huang, et al. Agentsense: Benchmarking social intelligence  
 595 of language agents through interactive scenarios. In *Proceedings of the 2025 Conference of the  
 596 Nations of the Americas Chapter of the Association for Computational Linguistics: Human Lan-  
 597 guage Technologies (Volume 1: Long Papers)*, pp. 4975–5001, 2025.

598

599 OpenAI. Openai o1 system card. Technical report, OpenAI, 2025. URL <https://cdn.openai.com/o1-system-card-20241205.pdf>.

594 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong  
 595 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to fol-  
 596 low instructions with human feedback. *Advances in neural information processing systems*, 35:  
 597 27730–27744, 2022.

598 Maarten Sap, Hannah Rashkin, Derek Chen, Ronan LeBras, and Yejin Choi. Socialiqa: Com-  
 599 monsense reasoning about social interactions. In *Conference on Empirical Methods in Natural  
 600 Language Processing*, 2019.

601 Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam  
 602 Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, et al. Beyond the  
 603 imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint  
 604 arXiv:2206.04615*, 2022.

605 Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. Commonsenseqa: A question  
 606 answering challenge targeting commonsense knowledge. In *Proceedings of the 2019 Conference  
 607 of the North American Chapter of the Association for Computational Linguistics: Human Lan-  
 608 guage Technologies*, pp. 4149–4158, 2019.

609 Tongyi DeepResearch Team. Tongyi-deepresearch. <https://github.com/Alibaba-NLP/DeepResearch>, 2025.

610 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Niko-  
 611 lay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open founda-  
 612 tion and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.

613 Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman.  
 614 Glue: A multi-task benchmark and analysis platform for natural language understanding. In  
 615 *Proceedings of the 2018 EMNLP Workshop BlackboxNLP*, pp. 353–355, 2018.

616 Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer  
 617 Levy, and Samuel R Bowman. SuperGLUE: A stickier benchmark for general-purpose language  
 618 understanding systems. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

619 Jiayin Wang, Fengran Mo, Weizhi Ma, Peijie Sun, Min Zhang, and Jian-Yun Nie. A user-centric  
 620 multi-intent benchmark for evaluating large language models, 2024. URL <https://arxiv.org/abs/2404.13940>.

621 Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh J Hua,  
 622 Zhoujun Cheng, Dongchan Shin, Fangyu Lei, et al. Osworld: Benchmarking multimodal agents  
 623 for open-ended tasks in real computer environments. *Advances in Neural Information Processing  
 624 Systems*, 37:52040–52094, 2024.

625 Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a ma-  
 626 chine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association  
 627 for Computational Linguistics*, pp. 4791–4800, 2019.

628 Zhiyuan Zeng, Jiashuo Liu, Siyuan Chen, Tianci He, Yali Liao, Yixiao Tian, Jinpeng Wang, Zaiyuan  
 629 Wang, Yang Yang, Lingyue Yin, Mingren Yin, Zhenwei Zhu, Tianle Cai, Zehui Chen, Jiecao  
 630 Chen, Yantao Du, Xiang Gao, Jiacheng Guo, Liang Hu, Jianpeng Jiao, Xiangsheng Li, Jingkai  
 631 Liu, Shuang Ni, Zhoufutu Wen, Ge Zhang, Kaiyuan Zhang, Xin Zhou, Jose Blanchet, Xipeng  
 632 Qiu, Mengdi Wang, and Wenhao Huang. FutureX: An advanced live benchmark for llm agents in  
 633 future prediction, 2025. URL <https://arxiv.org/abs/2508.11987>.

634 Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny  
 635 Zhou, and Le Hou. Instruction-following evaluation for large language models. *arXiv preprint  
 636 arXiv:2311.07911*, 2023.

637

638

639

640

641

642

643

644

645

646

647

648  
649

## A APPENDIX

650

### A.1 USE OF LARGE LANGUAGE MODELS

652

Large Language Models are used for grammar check and polishing in this paper.

654

655

### A.2 CONSINT-TREE CONSTRUCTION DETAILS

656

657

First, the LLM (GPT-4o) is utilized to extract branches from each consumer discussion. During this process, the model is prompted to summarize a user’s discussion following the template: `<Product_series>`, in the `<Usage Scenario>`, its `<Aspect>`, compared with `<Comparison>`, gives consumers the `<Feeling>` perception, and the discussion suggests that `<Tendency>`.

662

Next, the branches are used to construct the tree. All branches are connected to the tree root node, forming an initial tree. Sentence Transformers are then employed to merge semantically similar nodes layer by layer from the top down within the initial tree. During this process, nodes in the same layer that share the same parent node and are semantically similar are merged into one. The child nodes of each pre-merged node are then designated as the child nodes of the merged node. Additionally, the weight of the merged node is calculated as the sum of the weights of its child nodes. Ultimately, the fully merged tree is referred to as CONSINT-TREE, which will resemble the structure shown in Figure X. Nodes with higher weights will appear in the shallow layers of the tree; this is because the core topics of discussion (e.g., usage scenarios, aspects, feelings) often overlap across discussions from different consumers, and such high-weight nodes represent the aspects of the product that users focus on most.

672

673

This process enables the clear presentation of user discussion content in a tree structure while highlighting discussion hotspots. Meanwhile, by updating the discussion data and reconstructing CONSINT-TREE, we can analyze changes in child nodes under the same parent node between the two trees—thereby identifying users’ immediate concerns and long-term strategic considerations.

676

677

After summarization, the key terms in the template are fetched to form a single branch. For the branches derived from one discussion, the initial weight of each node is equal, ranging from 1 to 3, and determined by the discussion’s upvotes and view count. Notably, not every discussion can be summarized to fill all six key terms—more successfully filled key terms correspond to a longer branch path, which in turn reflects a more in-depth consumer intent. Next, all branches are connected to the tree root node, forming an initial tree. Sentence Transformers are then employed to merge semantically similar nodes layer by layer from the top down. During this process, nodes in the same layer that share the same parent node and are semantically similar are merged into one. The child nodes of each pre-merged node are then designated as the child nodes of the merged node. Additionally, the weight of the merged node is calculated as the sum of the weights of its child nodes. Ultimately, the fully merged tree is referred to as CONSINT-TREE. Nodes with higher weights will appear in the shallow layers of the tree; this is because the core topics of discussion (e.g., usage scenarios, aspects, feelings) often overlap across discussions from different consumers.

688

**Lighten the Tree:** To assess the depth and breadth of a LLM’s understanding of consumer intent, content from questionnaires will be extracted into branches. Each branch will undergo semantic matching with the nodes in the CONSINT-TREE from top to bottom using a Sentence Transformer. Nodes that are successfully matched will be marked as “lightened,” and the lightened nodes in the CONSINT-TREE will form a subtree. and the questionnaire will receive the score corresponding to that node.

695

696

697

698

699

700

701

Specifically, for each question in the questionnaire, the question stem and its four options will first be concatenated into four opinion statements. The LLM will then extract four branches from these four statements. These four branches will be matched with the nodes in the CONSINT-TREE from top to bottom—each branch will lighten a path and obtain a score based on the weights of the nodes along that path. The branch with the highest score for a given question will be used to “lighten” the CONSINT-TREE. Notably, nodes in the CONSINT-TREE cannot be repeatedly lightened by different questions. After iterating through all questions, the lightened nodes in the CONSINT-TREE will form a subtree.

702 A.3 INFORMATIVENESS  
703704 **Lexical Richness:** The evaluation of lexical richness relies on two key metrics: Type-Token Ratio  
705 (TTR) Johnson (1944) and Distinct-n Li et al. (2016). TTR quantifies the ratio of unique tokens to  
706 the total number of words in the text. It is defined as:

707 
$$708 \text{TTR} = \frac{\text{Count(unique token)}}{\text{Count(tokens)}}$$
  
709

710 where a higher TTR indicates greater lexical richness. Distinct-n focuses on the  $n$ -gram level, mea-  
711 suring the ratio of unique  $n$ -grams to the total number of  $n$ -grams. This study focuses on *bi-grams*,  
712 and the Distinct-n is calculated as:

713 
$$714 \text{Distinct-n} = \frac{\text{Count(unique bi-gram)}}{\text{Count(bi-grams)}}.$$

715 **Semantic Redundancy** is evaluated using a self-referential manner Chen et al. (2021b), where the  
716 average maximum semantic similarity is computed between each question and all other questions in  
717 the questionnaire, as well as between each question’s options and all other questions’ options. Given  
718 a set of questions  $Q = \{q_1, q_2, \dots, q_n\}$ , the semantic similarity between any two questions  $q_i$  and  $q_j$   
719 is calculated using cosine similarity:

720 
$$721 \text{Sim}(q_i, q_j) = \frac{\mathbf{v}_i \cdot \mathbf{v}_j}{\|\mathbf{v}_i\| \|\mathbf{v}_j\|},$$
  
722

723 where  $\mathbf{v}_i$  and  $\mathbf{v}_j$  represent the vector embeddings of questions  $q_i$  and  $q_j$ , respectively. The redundancy  
724 score is then computed as the average of the maximum similarity values across all pairs of  
725 questions:

726 
$$727 \text{Redundancy} = \frac{1}{n} \sum_{i=1}^n \max_{j \neq i} \text{Sim}(q_i, q_j).$$
  
728

729 Notably, a lower redundancy score indicates less repetition in question paradigms and option de-  
730 signs, which reflects the LLM’s ability to understand consumers’ intentions from multiple perspec-  
731 tives and conduct multi-source causal inference.732 A.4 CASE STUDY  
733734 Table 4: Performance of reasoning LLMs, general LLMs, and open-source LLM on Google Nest  
735 Smart Discussion.736  
737 

Model	Depth						Breadth	Informativeness		Correctness
	L1	L2	L3	L4	L5	Overall		Lexical	Semantic $\downarrow$	
GPT-5	3.33	<b>21.79</b>	<b>6.53</b>	<b>4.19</b>	0.00	7.17	<b>50.84</b>	<b>0.87</b>	<b>0.30</b>	0.75
GPT-o3	2.51	20.02	5.33	2.77	<b>1.91</b>	6.51	50.19	0.75	0.46	<b>0.95</b>
Qwen3-30B-A3B	<b>31.32</b>	16.61	2.77	1.49	0.00	<b>10.44</b>	50.51	0.81	0.39	0.75

740 A.4.1 GOOGLE NEST SMART SPEAKER QUESTIONNAIRE FROM GPT-5  
741

## 742 1. How do you primarily use your Google Nest speakers at home?

743

- 744 A. For music playback
- 745 B. For controlling smart devices
- 746 C. For asking questions/time/weather
- 747 D. For security alerts or doorbell announcements

748 *Answer: A. Users reported using Nest speakers most often for music, followed by smart  
749 home control and daily reminders such as weather or timers.*

## 750 2. How satisfied are you with the sound quality of Nest Audio compared to Nest Mini?

756                   A. Nest Audio is leagues better, especially bass  
 757                   B. Mini is enough for casual listening  
 758                   C. Nest Audio is adequate but not impressive  
 759                   D. No difference noticed

761                   *Answer: A. Users consistently said Nest Audio has much better bass and overall quality  
 762                   compared to Mini, making it preferable for music.*

763                   **3. Have you experienced connection issues with your Nest speakers in recent years?**

764                   A. Yes, frequent disconnections and 'sorry something went wrong'  
 765                   B. Yes, occasional hiccups  
 766                   C. No, they work reliably  
 767                   D. Issues only due to Wi-Fi provider/router

769                   *Answer: A. Many users reported worsening connection reliability over time, though some  
 770                   fixed issues by upgrading Wi-Fi or resetting devices.*

771                   **4. How well does Google Home integrate with your non-Google devices (like Tuya, Zig-  
 772                   Bee, or Ikea smart products)?**

773                   A. Very smooth integration  
 774                   B. Works but often buggy  
 775                   C. I cannot integrate them at all  
 776                   D. I only use 100% Google products

778                   *Answer: B. Several users noted persistent issues integrating Tuya/Lidl ZigBee and Ikea  
 779                   products with Google Home compared to their native apps.*

780                   **5. Which competing ecosystem do you feel has smoother smart home integration?**

781                   A. Apple HomeKit/HomePod  
 782                   B. Google Nest/Home  
 783                   C. Amazon Alexa/Echo  
 784                   D. None of them are ideal

786                   *Answer: A. Some users with Apple devices felt HomePod integration was smoother, while  
 787                   Google or Amazon were preferred for value and compatibility.*

788                   **6. What's your biggest frustration with Nest smart speakers?**

789                   A. Laggy responses  
 790                   B. Misunderstood commands  
 791                   C. Device going offline  
 792                   D. Limited software updates

794                   *Answer: A. Users most commonly complained about slow responses and frequent 'didn't  
 795                   understand' errors, especially in recent years.*

796                   **7. How often do you use your Nest devices for multi-room audio?**

797                   A. Daily  
 798                   B. Occasionally for parties/events  
 799                   C. Rarely  
 800                   D. Never

802                   *Answer: A. Users said multi-room audio used to work seamlessly but has become frustrat-  
 803                   ingly unreliable, with groups often not functioning.*

804                   **8. Would you consider switching away from Google due to reliability issues?**

805                   A. Yes, to Alexa  
 806                   B. Yes, to Apple  
 807                   C. Yes, to Home Assistant or local-first systems  
 808                   D. No, I will stay with Google

809                   *Answer: A. Some users said they were ready to move to Alexa or Home Assistant because  
 810                   of lag and reliability problems with Nest.*

810 9. **Do you find Nest speakers useful for families with children?**  
811

812 A. Yes, great for music and educational features  
813 B. Somewhat, but limited reliability frustrates kids  
814 C. Not useful at all for kids  
815 D. Only good as background music

816 *Answer: B. Users mentioned features like 'animal of the day' were fun for kids, but white  
817 noise and music playback have become unreliable.*

818 10. **How important is price compared to reliability when choosing between Google, Ama-  
819 son, or Apple smart speakers?**  
820

821 A. Price is the most important  
822 B. Balance between price and reliability  
823 C. Reliability is the most important  
824 D. Ecosystem integration matters most

825 *Answer: A. Users often noted that Google's affordability got them locked in, but reliability  
826 issues make them consider pricier alternatives.*

827 11. **Have you faced limitations when casting media from phones to Nest devices?**  
828

829 A. Yes, apps don't always connect  
830 B. Sometimes there's lag/delay  
831 C. No issues at all  
832 D. I don't use casting

833 *Answer: A. Users reported frequent casting issues, especially with Spotify, YouTube, and  
834 video streams not reaching certain Nest devices.*

835 12. **If you primarily wanted a device for music, which would you pick?**  
836

837 A. Google Nest Audio  
838 B. Amazon Echo 4th gen  
839 C. Apple HomePod Mini  
840 D. Traditional Bluetooth speakers

841 *Answer: A. Users debated between Nest Audio, Echo, and HomePod. Many said Nest Audio  
842 had good bass but Echo was decent, while some still preferred Sonos or passive Bluetooth  
843 sets.*

844 13. **How do you feel about Google discontinuing/reducing stock of Nest devices?**  
845

846 A. Concerned about product support  
847 B. Neutral, waiting for new models  
848 C. Considering switching to another brand  
849 D. Not worried at all

850 *Answer: A. Several users worried Nest Mini and Audio are discontinued, wondering if  
851 Google will abandon the smart speaker hardware.*

852 14. **What feature would make you more likely to stick with Google Nest speakers?**  
853

854 A. Improved reliability and faster response  
855 B. Better music/audio quality  
856 C. Deeper integration with third-party devices  
857 D. Clear roadmap and updates from Google

858 *Answer: A. Users said better reliability, sound improvements, and smoother ecosystem  
859 updates would convince them to remain loyal.*

860 15. **Do you experience more issues with Google Assistant understanding you in multilin-  
861 gual households?**  
862

863 A. Yes, it constantly misinterprets  
B. Sometimes, especially switching languages  
C. No, it works fine in multiple languages

864                   D. I only use one language  
 865  
 866  
 867

*Answer: A. Users noted Assistant struggles badly in multilingual homes, often failing basic commands or mixing languages incorrectly.*

868                   **16. What's your perspective on Nest speakers' long-term durability?**

869                   A. Still working fine years later  
 870                   B. Performance has worsened over time  
 871                   C. Hardware is durable but software declines  
 872                   D. They feel like e-waste now  
 873

*Answer: C. Some users praised durability, while many complained hardware outlasted software support, calling devices obsolete early.*

876                   **17. How do you primarily resolve issues with Nest devices?**

877                   A. Factory reset  
 878                   B. Router and Wi-Fi upgrades  
 879                   C. Reinstalling Google Home app  
 880                   D. Contacting Google support  
 881

*Answer: A. Most users resorted to factory resets or Wi-Fi upgrades; official support was rarely mentioned as helpful.*

884                   **18. Would you invest in another Nest smart display (like Hub/Hub Max) now?**

885                   A. Yes, I still trust Google ecosystem  
 886                   B. Maybe, if I find a second-hand deal  
 887                   C. No, too much risk of discontinued support  
 888                   D. I prefer other brands' smart displays  
 889

*Answer: C. Users were hesitant to buy discontinued Nest Hubs/Max, fearing bricking or lack of updates.*

892                   **19. When connecting Nest with services like Spotify or YouTube Music, what's your experience?**

893                   A. Smooth, works well  
 894                   B. Works but occasionally lags  
 895                   C. Often breaks or blocks premium-only features  
 896                   D. I don't link music services  
 897

*Answer: C. Several users reported Spotify on Nest sometimes says 'premium only' even with premium, and YouTube Music integration often fails.*

898                   **20. What future direction should Google take with Nest smart speakers?**

902                   A. Bring Gemini AI with better natural understanding  
 903                   B. Focus on keeping devices reliable  
 904                   C. Produce new affordable hardware  
 905                   D. Open-source support if retiring devices  
 906

*Answer: B. Users speculated Google must fix reliability, offer Gemini AI improvements, and either release new hardware or open source old devices.*

907                   **A.4.2 GOOGLE NEST SMART SPEAKER QUESTIONNAIRE FROM GPT-O3**

912                   **1. When users retrofit 1980s intercoms with Nest Mini units, which room-specific control do they hope to achieve later on?**

913                   A. Only ceiling fans and lights of that bedroom  
 914                   B. Satellite TV channels in the garage  
 915                   C. Printer queues in the study  
 916                   D. Irrigation valves in the backyard  
 917

918  
919  
920

921  
922  
923  
924  
925  
926  
927  
928  
929  
930  
931  
932  
933  
934  
935  
936  
937  
938  
939  
940  
941  
942  
943  
944  
945  
946  
947  
948  
949  
950  
951  
952  
953  
954  
955  
956  
957  
958  
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971

Answer: A. Users describe planning “a speaker in every bedroom with some intricate setup to both only control devices specific to that room (like ceiling fans and lights) as well as shared devices.”

2. **What adjective did a long-time owner use to praise Nest Audio’s stereo sound after pairing two units?**

- A. Incredible
- B. Tinny
- C. Muffled
- D. Overpriced

Answer: A. A commenter said “I use 2 Nest Audios in a stereo setup, and the audio is incredible,” reflecting positive feelings about sound quality.

3. **Which competing smart speaker line did several Redditors say they might switch to because Google devices have become “laggy” and “driving me insane”?**

- A. Amazon Echo / Alexa
- B. Sonos Era
- C. Bose Smart Ultra
- D. Marshall Uxbridge

Answer: A. Many posts mention considering Amazon Echo or Alexa displays as an alternative when Nest performance deteriorated.

4. **In the thread about buying a Nest Hub Max, which security-related use case was highlighted as a reason to still want the display?**

- A. Acting as a digital photo frame with camera recording
- B. Hosting a VPN server
- C. Controlling sprinklers via Zigbee
- D. Calibrating 3D printers

Answer: A. A buyer said they liked “the camera/security recording function and using it as a digital photo frame,” showing the usage scenario.

5. **Which phrase did frustrated owners repeatedly hear instead of successful commands, prompting them to call Google Home a “support group”?**

- A. “Sorry, something went wrong, try again later.”
- B. “Firmware upgrade in progress.”
- C. “Device is paired in another room.”
- D. “Low battery, shutting down.”

Answer: A. Multiple users quote the device replying “Sorry, something went wrong, try again later,” illustrating a common pain point.

6. **Why did one user say the Pixel Tablet on its dock feels like an “old TV/VCR combo” compared with a real Nest Hub?**

- A. It can’t be asked to play music on other Google speakers
- B. It lacks Wi-Fi 6E support
- C. The screen is smaller than 5 inches
- D. It forces Amazon Prime ads

Answer: A. They complained that you “can’t tell it to play music on it from another Google speaker,” so the hybrid device does neither role well.

7. **Which connectivity problem did a border-area listener report when TuneIn stations kept dropping on Nest speakers?**

- A. Occasional to frequent loss in signal
- B. Crackling Bluetooth interference only at night
- C. Wrong language playback
- D. Overheating power adapters

972  
973  
974  
975  
976  
977  
978  
979  
980  
981  
982  
983  
984  
985  
986  
987  
988  
989  
990  
991  
992  
993  
994  
995  
996  
997  
998  
999  
1000  
1001  
1002  
1003  
1004  
1005  
1006  
1007  
1008  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025

*Answer: A. The post says “I have experienced occasional to frequent loss in signal when listening to stations that utilize TuneIn.”*

8. **When discussing Matter devices going offline, which brand of mesh router system was singled out for Thread settings confusion?**

- A. Eero 6E
- B. TP-Link Deco
- C. UniFi Dream Router
- D. Netgear Orbi

*Answer: A. A user wrote “I have an Eero 6e mesh router system... The Threads feature is toggled on,” yet their Matter gear still dropped.*

9. **How did a Nest Mini owner describe the music delay when the speaker was added to a stereo link in Google Home?**

- A. The delay is HUGE.
- B. It syncs perfectly.
- C. Only milliseconds of lag.
- D. Delay happens once a month.

*Answer: A. The post states, “If I play any music... the delay is HUGE,” emphasizing a negative feeling about latency.*

10. **Which future-oriented speculation did shoppers raise after noticing no Nest Audio stock in multiple country stores?**

- A. A new generation might be announced at the Pixel event
- B. Google is switching to Apple HomeKit
- C. All smart speakers will become subscription-based
- D. Wi-Fi will be removed from Nest

*Answer: A. They asked, “Are people expecting a new generation to be announced at the Pixel event in a couple weeks?”—a tendency toward anticipating new hardware.*

11. **Which cloud storage dilemma did dual-ecosystem users discuss while already owning many Nest Hubs and iCloud devices?**

- A. Paying for both 200 GB iCloud and 200 GB Google One plans
- B. Choosing between Dropbox and Box free tiers
- C. Losing access to Microsoft OneDrive photos
- D. Migrating from Amazon S3 Glacier Vaults

*Answer: A. The repeated post describes both iCloud and Google One hitting the 200 GB limit and not wanting to upgrade both.*

12. **What network feature on apartment Wi-Fi prevented an elderly resident’s Nest Mini from completing setup?**

- A. AP Isolation turned on
- B. Hidden SSID broadcast
- C. WPA3-Enterprise only
- D. Dual NAT tunneling

*Answer: A. The care home enables “AP Isolation,” so the speaker throws the message “Please check your Wi-Fi network settings.”*

13. **Which sound-related improvement motivated users to prefer Nest Audio over their old Google Home Minis?**

- A. ‘Bass is the most noticeable improvement’ at high volume
- B. Built-in CD player support
- C. Dolby Atmos rear channels
- D. Quad-mic noise cancelling

*Answer: A. One review says, “Bass is the most noticeable improvement, high volume performance is better,” highlighting the aspect of audio quality.*

1026 14. **How much did Canadian bargain hunters report paying at Lowe's or Home Depot for**  
 1027 **clearance Nest Audio units?**

1028     A. \$39.97  
 1029     B. \$129.99  
 1030     C. \$199.00  
 1031     D. \$15.00

1033 *Answer: A. Posts note “Nest Audio for sale for \$39.97... is it worth getting,” reflecting*  
 1034 *pricing sentiment.*

1035 15. **Which workaround did some owners adopt because the Nest Hub could no longer**  
 1036 **resume music on the intended room speaker?**

1037     A. Using the broadcast command instead of TTS  
 1038     B. Switching to Zigbee bulbs  
 1039     C. Turning on microphone sensitivity  
 1040     D. Downgrading firmware via USB

1042 *Answer: A. One poster said they had to “resort to using broadcast commands which are*  
 1043 *clunky” when TTS stopped working.*

1044 16. **What phrase did a Nest thermostat user shout after eco mode kept activating despite**  
 1045 **settings being disabled?**

1046     A. “Jeezus Google.”  
 1047     B. “Bravo Assistant!”  
 1048     C. “Mission accomplished!”  
 1049     D. “Danke Alexa.”

1050 *Answer: A. The frustrated quote is “Jeezus Google,” showing irritation with unwanted eco*  
 1051 *behaviour.*

1052 17. **When debating cloud versus local AI, which low-power device did a homeowner con-**  
 1053 **sider dedicating as an “always-on screen” for NotebookLM?**

1055     A. An old MacBook Pro  
 1056     B. A Raspberry Pi Zero  
 1057     C. A Lenovo Tab M8  
 1058     D. A Pixel 2 phone

1059 *Answer: A. They planned “to dedicate an old MacBook Pro for the AI assistant” but were*  
 1060 *open to tablets.*

1061 18. **Which free radio service did a listener compare to TuneIn, noting that Audacy re-**  
 1062 **tained signal ‘much better’ on Nest speakers?**

1064     A. Audacy  
 1065     B. Pandora  
 1066     C. SiriusXM  
 1067     D. iHeartWeather

1068 *Answer: A. The poster said “The other services like Audacy work much better in terms of*  
 1069 *signal retention,” offering a comparison.*

1070 19. **What is one reason a user said HomePods were a ‘let down’ compared with Google**  
 1071 **Nest speakers?**

1073     A. ‘You have to AirPlay from your phone’ for many tasks  
 1074     B. They lack any bass at all  
 1075     C. They cannot set alarms  
 1076     D. HomePods don’t support Wi-Fi

1077 *Answer: A. The criticism: “on HomePod minis you have to air play from your phone... google*  
 1078 *are just way ahead,” indicating a comparison and feeling.*

1079 20. **Faced with a 2025 Nest price hike e-mail, how many days of event video history does**  
 the ‘Nest Aware Plus’ plan promise to keep?

1080                   A. 60 days  
 1081                   B. 30 days  
 1082                   C. 7 days  
 1083                   D. 120 days

1085                   *Answer: A. The notice states the plan keeps “60 days of event video history... up to 10  
 1086                   days of 24/7 video history,” summarising the aspect of subscription value.*

1087                   **A.4.3 GOOGLE NEST SMART SPEAKER QUESTIONNAIRE FROM QWEN3-30B-A3B**

1089                   **1. In your smart home setup, how do you integrate Google Nest Mini speakers into an  
 1090                   older home intercom system?**

1091                   A. By replacing the entire intercom system with new wireless speakers  
 1092                   B. By using a 24VDC power supply and buck converters to repurpose old wiring  
 1093                   C. By connecting the speakers directly to a home Wi-Fi network without rewiring  
 1094                   D. By mounting the speakers to the wall using adhesive pads only

1095                   *Answer: B. By using a 24VDC power supply and buck converters to repurpose old wiring*

1097                   **2. What is the primary reason for using Google Nest Audio speakers in a stereo setup  
 1098                   for music playback?**

1099                   A. They offer better sound quality and deeper bass than other smart speakers  
 1100                   B. They are the only Google devices compatible with Apple Music  
 1101                   C. They support Bluetooth audio output without delay  
 1102                   D. They are designed specifically for use with Sonos systems

1103                   *Answer: A. They offer better sound quality and deeper bass than other smart speakers*

1105                   **3. When using Google Nest Mini as a temporary speaker during a move, what feature is  
 1106                   most important for portability?**

1107                   A. Built-in battery and rechargeable power source  
 1108                   B. Ability to connect via USB-C to a TV  
 1109                   C. Support for HDMI input from a computer  
 1110                   D. Integration with Apple HomeKit for remote access

1111                   *Answer: A. Built-in battery and rechargeable power source*

1113                   **4. Why might a user choose Google Nest Mini over Apple HomePod Mini for a smart  
 1114                   home setup?**

1115                   A. Because Apple HomePod Mini requires a HomePod Hub for full functionality  
 1116                   B. Because Google Nest Mini supports more third-party integrations and voice com-  
 1117                   mands  
 1118                   C. Because HomePod Mini lacks music playback capabilities  
 1119                   D. Because Google Mini is compatible with all smart home ecosystems

1120                   *Answer: B. Because Google Nest Mini supports more third-party integrations and voice  
 1121                   commands*

1123                   **5. What has caused frustration for users who rely on Google Nest speakers for music  
 1124                   playback and voice commands?**

1125                   A. Frequent unresponsiveness and inconsistent command recognition  
 1126                   B. Inability to support Spotify Premium accounts  
 1127                   C. Limited compatibility with Bluetooth devices  
 1128                   D. High cost of replacing outdated models

1129                   *Answer: A. Frequent unresponsiveness and inconsistent command recognition*

1130                   **6. How do users perceive the reliability of Google Nest devices after the company’s shift  
 1131                   to Gemini AI?**

1132                   A. They believe the new AI will improve performance and voice recognition  
 1133                   B. They are concerned about service degradation and potential obsolescence

1134                   C. They expect all devices to be replaced within a year  
 1135                   D. They trust that Google will maintain full support indefinitely  
 1136                   *Answer: B. They are concerned about service degradation and potential obsolescence*

1137           **7. What do users suggest as the best way to maintain a smart home ecosystem without**  
 1138           **relying on Google's services?**

1139                   A. Using Home Assistant with a local server for full control  
 1140                   B. Switching to Apple HomeKit for better integration  
 1141                   C. Purchasing only Amazon Echo devices for compatibility  
 1142                   D. Replacing all devices with non-smart alternatives  
 1143                   *Answer: A. Using Home Assistant with a local server for full control*

1144           **8. What is a common workaround when a Google Nest speaker fails to connect to Wi-Fi**  
 1145           **during setup?**

1146                   A. Changing the account region to the United States in the app settings  
 1147                   B. Using a wired Ethernet connection instead of Wi-Fi  
 1148                   C. Factory resetting the device and re-registering it  
 1149                   D. Replacing the speaker with a newer model  
 1150                   *Answer: A. Changing the account region to the United States in the app settings*

1151           **9. Why do some users prefer Google Nest Audio over other smart speakers for music**  
 1152           **listening?**

1153                   A. It has superior stereo sound quality and bass response  
 1154                   B. It is the most affordable smart speaker with no subscription fees  
 1155                   C. It supports more music streaming services than any other model  
 1156                   D. It integrates with Google Photos for visual album displays  
 1157                   *Answer: A. It has superior stereo sound quality and bass response*

1158           **10. What is the main concern for users who are invested in both Apple and Google ecosys-**  
 1159           **tems?**

1160                   A. Managing duplicate cloud storage between iCloud and Google One  
 1161                   B. Choosing between Apple HomePod and Google Nest for smart home control  
 1162                   C. Upgrading to new smartphones to maintain compatibility  
 1163                   D. Finding compatible smart bulbs for both platforms  
 1164                   *Answer: A. Managing duplicate cloud storage between iCloud and Google One*

1165           **11. How do users feel about the future of Google Nest hardware following recent service**  
 1166           **issues?**

1167                   A. They believe Google will soon release a new generation of devices  
 1168                   B. They are skeptical about long-term support and product continuity  
 1169                   C. They expect Google to fully open-source the software for user maintenance  
 1170                   D. They are confident that all devices will continue to function indefinitely  
 1171                   *Answer: B. They are skeptical about long-term support and product continuity*

1172           **12. What challenge arises when trying to control Google Nest devices from a non-Google**  
 1173           **ecosystem device?**

1174                   A. Voice commands are not recognized due to language mismatch  
 1175                   B. The speaker cannot detect voice input from Apple devices  
 1176                   C. Integration with Home Assistant becomes unreliable or impossible  
 1177                   D. The device refuses to connect to Wi-Fi networks outside the Google network  
 1178                   *Answer: C. Integration with Home Assistant becomes unreliable or impossible*

1179           **13. What is a key limitation when using Google Nest Mini speakers for doorbell alerts?**

1180                   A. The alert only plays on one speaker, not all connected devices  
 1181                   B. The doorbell chime only works when the speaker is unplugged

1188 C. The speaker cannot broadcast alerts during music playback  
 1189 D. Alerts are delayed by 10–15 seconds due to app routing

1190 *Answer: A. The alert only plays on one speaker, not all connected devices*

1191 **14. Why might a user consider replacing a Google Nest Hub Max despite its popularity?**

1192 A. It is no longer available for purchase in many regions  
 1193 B. It has a built-in camera that cannot be disabled  
 1194 C. It requires a paid subscription for basic functionality  
 1195 D. It only works with Apple devices and not Android

1196 *Answer: A. It is no longer available for purchase in many regions*

1197 **15. How do users describe the current performance of Google Nest Hub devices in 2025?**

1198 A. They are slow, glitchy, and frequently misinterpret commands  
 1199 B. They have improved dramatically with the new Gemini update  
 1200 C. They function as well as they did five years ago  
 1201 D. They are only useful for playing music and showing photos

1202 *Answer: A. They are slow, glitchy, and frequently misinterpret commands*

1203 **16. What is a major reason users are hesitant to expand their Google Nest ecosystem?**

1204 A. Fear of being locked into a system that may no longer be supported  
 1205 B. High cost of additional smart home accessories  
 1206 C. Incompatibility with older Wi-Fi routers  
 1207 D. Lack of support for non-English languages

1208 *Answer: A. Fear of being locked into a system that may no longer be supported*

1209 **17. What is the most common solution users suggest for improving audio performance  
 1210 across multiple Nest speakers?**

1211 A. Upgrading to a higher-end model like the Google Home Max  
 1212 B. Using a dedicated audio amplifier for better sound quality  
 1213 C. Replacing all speakers with Sonos devices for consistent playback  
 1214 D. Using a local server to manage audio streams and reduce latency

1215 *Answer: D. Using a local server to manage audio streams and reduce latency*

1216 **18. What is a major drawback of using Google Nest Audio as a TV speaker?**

1217 A. It has a significant audio delay when connected via Bluetooth  
 1218 B. It lacks a 3.5mm input port for direct audio connection  
 1219 C. It cannot play video content from streaming apps  
 1220 D. It only works with Google TV devices

1221 *Answer: A. It has a significant audio delay when connected via Bluetooth*

1222 **19. What do users suggest as an alternative to Google Nest speakers for a private, local  
 1223 smart home system?**

1224 A. Using a Raspberry Pi with a local LLM for voice control  
 1225 B. Replacing all devices with Amazon Echo Dot models  
 1226 C. Switching to Apple HomePods for better privacy  
 1227 D. Using only non-smart speakers with manual controls

1228 *Answer: A. Using a Raspberry Pi with a local LLM for voice control*

1229 **20. How do users perceive the value of older Google Nest devices like the Nest Mini and  
 1230 Nest Audio?**

1231 A. They are still functional and affordable, especially when bought secondhand  
 1232 B. They are outdated and no longer supported by Google  
 1233 C. They are only useful for basic tasks like playing alarms  
 1234 D. They are incompatible with modern Wi-Fi networks

1235 *Answer: A. They are still functional and affordable, especially when bought secondhand*




Figure 4: Lighted Tree from GPT-5 on Google Nest Smart Speaker.




Figure 5: Lighted Tree from GPT-o3 on Google Nest Smart Speaker.




Figure 6: Lighted Tree from Qwen3-30B-A3B on Google Nest Smart Speaker.