

# COACH Meets QUORUM: A Framework and Pipeline for Aligning User, Expert, and Developer Perspectives in LLM-Generated Health Counselling

Yee Man Ng\*, Bram van Dijk†, Pieter Beynen◇, Otto Boekesteijn◇,  
Joris Jansen◇, Gerard van Oortmerssen\*, Max van Duijn\*, Marco Spruit\*†

\*Leiden University, {y.m.ng, m.r.spruit, g.van.oortmerssen, m.j.van.duijn}@liacs.leidenuniv.nl

†Leiden University Medical Center, b.m.a.van\_dijk@lumc.nl

◇Healthy Chronos, {pieter.beynen, otto.boekesteijn, joris.jansen}@healthychronos.com

## Abstract

Systems that collect data on sleep, mood, and activities can provide valuable lifestyle counselling to populations affected by chronic disease and its consequences. Such systems are, however, challenging to develop; in addition to *reliably* extracting patterns from user-specific data, systems should contextualise these patterns with validated medical knowledge to ensure the *quality* of counselling and generate counselling that is *relevant* to a real user. We present QUORUM, an evaluation framework that unifies these developer-, expert-, and user-centric perspectives, and show with a real case study that it meaningfully tracks convergence and divergence in stakeholder perspectives. We also present COACH, a Large Language Model-driven pipeline to generate personalised lifestyle counselling for our *Healthy Chronos* use case, a diary app for cancer patients and survivors. Applying our framework indicates that, overall, users, medical experts, and developers converge on the view that the generated counselling is relevant, of good quality, and reliable. However, stakeholders also diverge on the tone of the counselling, sensitivity to errors in pattern-extraction, and potential hallucinations. These findings highlight the importance of multi-stakeholder evaluation for consumer health language technologies and illustrate how a unified evaluation framework can support trustworthy, patient-centered NLP systems in real-world settings.

**Keywords:** consumer health question answering, large language models in health, remote monitoring

## 1. Introduction

People living with the consequences of cancer (treatment) often suffer from long-term effects such as fatigue and mental health issues (Akechi et al., 1999). While it is known that healthy lifestyle changes like increasing physical activity positively impact health outcomes and quality of life, for both cancer patients and survivors (Hoedjes et al., 2022), they currently experience the state of lifestyle counselling in healthcare as not readily available, too fragmented, and general (van Aken et al., 2025; Drbohlav Ollerton et al., 2025; Tuinman et al., 2024). This demonstrates the need for personalised lifestyle counselling, through reliable pattern extraction from user-specific health records and with validated sources, in order to produce high-quality and evidence-based answers to consumer-health questions. While many AI-driven systems that generate counselling exist, their evaluation methodologies are often fragmented and lack realistic contexts (Lai et al., 2025; Raji et al., 2021).

This is why this paper contributes **QUORUM** (**Q**uality, **O**utcome **R**eliability, and **U**ser-relevance from **M**ultiple stakeholders), a new evaluation framework that aligns with the goals and interests of three key types of stakeholders in the development of NLP systems for health. From a *user-centric* perspective, counselling is *relevant* to users if the

advice aligns with their situation, increases the likelihood of action based on the counselling, and if they are satisfied with the tone and length of the counselling. *Experts*, on the other hand (in our case, cancer information specialists from [kanker.nl](http://kanker.nl), henceforth ‘experts’) focus on the *quality* of the counselling with respect to its correctness, and the appropriate tone and length for the population and disease they know well. *Developers/researchers* (two authors, henceforth ‘developers’) focus on a more technical evaluation and have as a vantage point *reliability*, i.e., baseline performance on key metrics such as faithfulness and completeness of the counselling, and hallucination rates. We argue that these three perspectives must be included in the evaluation of AI systems that support specific medical populations. Since we cannot expect such stakeholders to evaluate the output of AI systems by the same standards, they provide irreducible, complementary perspectives that only jointly indicate whether a given system is likely to have a positive, real-world impact.

We show QUORUM in action with a concrete use case: the Dutch *Healthy Chronos* application. This is a digital diary app that aids cancer patients and survivors by tracking sleep, activities, mood, and custom user-defined goals over time. We develop the Large Language Model (LLM) driven pipeline **COACH** (**C**ontextualised **O**utcome-

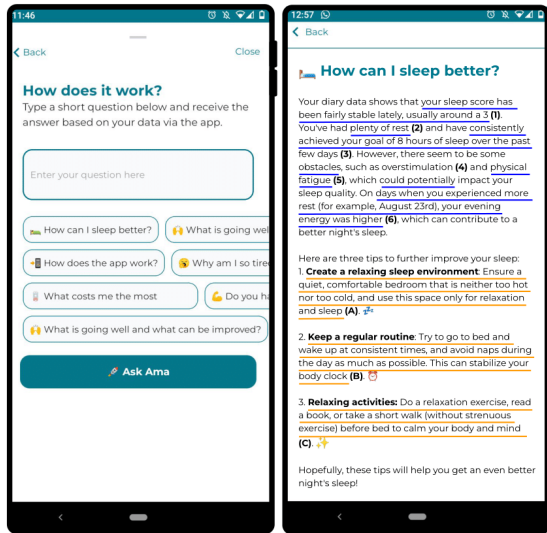


Figure 1: Left: user interface for submitting queries. Right: generated counselling for query ‘How can I sleep better?’ Blue underlining illustrates *claims about the user data*; orange indicates *contextualisation statements* that contextualise found patterns with information from the knowledge database. Claims and statements are labelled with (1), (A), etc., but this is hidden from the user.

**Adaptive Counselling for Health**) for *Healthy Chronos*, which uses Retrieval Augmented Generation (RAG) (Lewis et al., 2020) to generate personalised lifestyle counselling. The pipeline ingests a user query about their data (e.g. *How can I sleep better?*), after which COACH retrieves relevant chunks for that query from the external, validated knowledge database `kanker.nl`. COACH then contextualises the patterns extrapolated from the user data with the retrieved chunks into personalised counselling (for an impression, see Figure 1). By using QUORUM on COACH, we demonstrate that relevant, good-quality, and reliable lifestyle counselling is possible with LLMs. Nineteen Dutch app users acknowledge that the counselling aligns with their situation, they intend to follow up on the counselling, and endorse its tone and length. Six independent experts from `kanker.nl` deem the counselling of good quality in terms of the correctness of the contextualisation of extracted user patterns, and its tone and length. Technical evaluation by two developers shows COACH is reliable, as most claims (79%) are consistent with user data; moreover, for virtually all claims (97%), the relevant information has been retrieved. Still, we also find that 22% of the generated counselling statements are not strictly traceable to the knowledge database, thus are hallucinations (though not of a harmful kind).

Besides this convergence, QUORUM also highlights divergence between stakeholders. For exam-

ple, users and experts differ on the desired tone of the response, and differ with developers in the sensitivity to errors and hallucination in the reported patterns and counselling. These results highlight the utility of QUORUM and COACH for other populations and medical conditions, by helping build systems with a larger real-world impact.

## 2. Background

Consulting the internet for information on lifestyle change to improve one’s health and quality of life is increasingly common (Raeside et al., 2022). Still, people struggle with finding information fitting their specific queries, due to answer complexity, potentially incorrect information, and the lack of alignment with their situation (Bondarenko et al., 2021; Pugachev et al., 2023; Kayser et al., 2015). Though LLMs are recognised as potentially useful for providing lifestyle counselling due to their language and reasoning abilities, a recent overview shows that particularly evaluation methodologies for LLM-driven lifestyle counselling are fragmented, weak, and lack realistic contexts (Lai et al., 2025).

Related work has further shown that for LLM-driven chatbots in health, stakeholders can have diverging perspectives relevant to system development. For example, whereas users may prioritize the tone of the interaction, professional caregivers may emphasize safety and clinical validation aspects (Van Dijk et al., 2025). Still, current work typically targets only a subset of stakeholders, like relevance from a user perspective, but no further quality assessment of system output (Meywirth, 2024), or expert evaluation without including real users (Merrill et al., 2024). While such studies address important issues, often with novel AI techniques, extrapolating findings to real populations remains hard.

However, a study close to our setup is Steenstra et al. (2024), and it concerns an LLM-driven tool for Motivational Interviewing regarding alcohol use. It also adopts a three-way evaluation with different stakeholders, where humans evaluated the tone of the LLM-generated outputs, researchers evaluated the presence of appropriate MI techniques in the responses, and independent experts overall tested the system by role-playing interactions. Yet, Steenstra et al.’s work differs from ours in that their tool is not embedded in the real world.

Regarding system development, research has shown that many LLM-driven systems for lifestyle counselling still rely on proprietary models and cloud services (Lai et al., 2025). For startups, this is perhaps understandable, as self-hosting LLMs is costly, while cloud access to capable LLMs is cheap (Nandagopal, 2025); these are key considerations for startups like our *Healthy Chronos*

use case that inform some methodological choices we make below.

### 3. Methodology

#### 3.1. QUORUM Framework

QUORUM unifies three complementary perspectives from the literature. First, the **model-centric perspective** that emphasises well-known standardized metrics (e.g., recall) on common tasks (e.g., retrieving relevant elements). These metrics are important as they allow quickly probing a system’s base performance and comparison with other systems, but they are poor proxies for a system’s real-world impact (Wang et al., 2024a; Li et al., 2025). This is nevertheless a common perspective to adopt for AI **developers**, and in our framework, it involves faithfulness, completeness, and hallucination metrics as common to RAG evaluation (Papageorgiou et al., 2025). Faithfulness concerns the extent to which generated claims about the user data are consistent with that user’s data; completeness concerns the extent to which all relevant information in the user data for some query was retrieved; hallucination concerns the existence of contextualisation statements in the counselling that cannot be traced to the knowledge database (Recall the difference between *claims about user data* and *contextualisation statements* from Figure 1). Rather than factual incorrectness (the more common definition of hallucination), we consider lifestyle counselling that might be fitting but not grounded in the knowledge database, i.e., any information generated not from the parametric knowledge of the model, to be a type of hallucination.

The **user-centric perspective**, on the other hand, focuses on how well model behaviour actually helps a **real user** to reach a specific goal, thus it is about the relevance of a system’s output (Wang et al., 2024b). We asked users to evaluate the alignment of the response with their personal situation, the tone and length of the response, and the likelihood that the user would follow up on the counselling provided. These more subjective measures aim to capture the relevance of the response to the user.

We also incorporate the **expert-centric perspective**. In a health context, besides user and developer, **medical experts** as stakeholders often have different concerns (Van Dijk et al., 2025), for example, about the overall adequacy of the medical information in the response, because actions may be based on it, and because users may re-use the generated information in real medical consultations. Thus, we ask experts to evaluate the correctness of the contextualisation, based on their expertise on the (consequences of) cancer, as well as their

opinion on the style and length of the response.

##### 3.1.1. Holistic, Iterative Evaluation

In this way, we aim to incorporate a holistic perspective on the generated counselling that strikes a good balance between technical (model-centric), subjective (user-centric), and medical (expert-centric) aspects. Combining these three perspectives is vital, as single perspectives cannot by themselves provide a good estimation of the real-world impact of AI systems in health; they are complementary, and one perspective cannot be reduced to another. Though they do overlap in topics, the perspectives try to capture those opinions that can be reasonably expected from different stakeholders, given their background, goals, and interests.

For example, users could be less likely to verify reported patterns or lifestyle advice if an answer is in line with what they already believe; this is a common issue in dealing with the fluent output of LLMs nowadays and known as confirmation and consistency bias (Nirman et al., 2025). Or, they could be less likely to do so due to automation bias: an over-reliance on the capabilities of algorithms in generating valid information (Boonprakong, 2025). This shows how model- and expert-centric evaluation can complement each other. On the other hand, an AI system with perfect reliability may still have little real-world impact because few of its outputs may be meaningful to a user; here, the user-centric perspective provides the best proxy for real impact. What is still missing, then, is some assurance that the reliable and relevant output is also of good quality from an expert-centric perspective; only experts can adequately assess medical content.

QUORUM was conceived of as a quick probe in the context of *iterative* system development, where, especially for interactive systems, many brief stakeholder evaluations need to be done to incorporate their views in future versions of the system (van Gemert-Pijnen et al., 2011; Van Dijk et al., 2025). After evaluation is done with each group, a QUORUM user can compute, per perspective, a grand average that indicates how well the system at issue is doing from that perspective, while the individual variables provide details on the perspective’s key aspects. See Table 1 for a schematic overview of variables in QUORUM; we further show and discuss outputs of our framework in section 4.

##### 3.1.2. Data Collection

Here we describe the procedures for data collection from users, experts, and researchers. All protocols were approved by the Leiden University Ethics Review Board. Though we share the source code of our pipeline (see subsection 3.2), we cannot share user data itself nor `Healthy Chronos`

Perspective (N)	Statement evaluated	Variable	Measurement	Interpretation
Users (19)	1 <i>The response aligns with my situation</i>	Alignment	5-point Likert scale	Avg. score indicating alignment with user's situation (mood, energy, etc.)
	2 <i>I will follow the tips in the response</i>	Follow-up	5-point Likert scale	Avg. score indicating whether user is likely to act on the response
	3 <i>The response has the right tone</i>	Tone	5-point Likert scale	Avg. score of response tone quality
	4 <i>The response has the right length</i>	Length	5-point Likert scale	Avg. score of response length
Experts (6)	1 <i>The contextualisation is correct</i>	Correctness	5-point Likert scale	Avg. score of contextualisation quality
	2 <i>The response has the right tone</i>	Tone	5-point Likert scale	Avg. score of response tone quality
	3 <i>The response has the right length</i>	Length	5-point Likert scale	Avg. score of response length
Developers (2)	1 <i>This claim is consistent with the data</i>	Faithfulness	Binary	% of claims supported by user data
	2 <i>For this claim, the relevant user variables were retrieved</i>	Completeness	Binary	% of claims for which relevant columns were retrieved from user data
	3 <i>This contextualisation statement can be traced to the relevant chunks found by the pipeline</i>	Hallucination	Binary	% of claims grounded in the knowledge database ( <code>kanker.nl</code> )

Table 1: Contents (translated from Dutch) of the QUORUM evaluation framework. Five-point Likert scales are structured as *strongly disagree* (1), *disagree* (2), *neutral* (3), *agree* (4), *strongly agree* (5).




Figure 2: Overview of the pipeline flow and inputs.

source code due to privacy and intellectual property considerations.

**Users** Nineteen Dutch users were recruited via the app to participate in this study. As shown in Figure 1, users can query their data, but the app also suggests query topics based on frequency. In our evaluation, users evaluate responses on the question: *What is going well and what can be improved?* One reason for this is that this query covers a large part (41%) of a random sample of 100 queries. Also, this query has a broad scope, meaning that the pipeline is forced to correctly retrieve and analyse potentially many different relevant variables, which can be more challenging compared to more specific queries (e.g., about sleep). Users received the query and generated a response in the app, and then completed a short survey on the statements given in Table 1. They were asked to evaluate the full response, and were also invited to submit open remarks for each query-response pair. This amounted to nineteen evaluated query-response pairs. We did not collect further demographic data due to privacy considerations, but the user base of `Healthy Chronos` contains mostly Dutch adolescents and young adults.

**Experts** Six experts from `kanker.nl` consented to evaluating in total 50 random user query-

response pairs, and they evaluated for each pair the four statements given in Table 1. Experts were instructed to focus on the correctness of the *contextualisation* (see Figure 1, as due to privacy concerns it was not feasible to share user data with them, and because the contextualisation is based on the knowledge database (`kanker.nl`) they were familiar with. Besides providing scores, experts were also invited to submit remarks at every query-response pair.

**Developers** Two developers evaluated in two rounds in total 66 *claims about the user data* and 31 *contextualisation statements* (see Figure 1) from 10 random query-response pairs. Statements 1 and 2 in Table 1 were used for the claims about user data, whereas statement 3 was used for the contextualisation. After the developers initially evaluated the whole set independently, disagreements were discussed to consensus afterwards.

We computed measures for interrater reliability for experts and developers, for various variables, and therefore describe these metrics in tandem with our results in section 4.

### 3.2. Pipeline

`Healthy Chronos` is a digital diary for people living with the consequences of cancer (treatment),

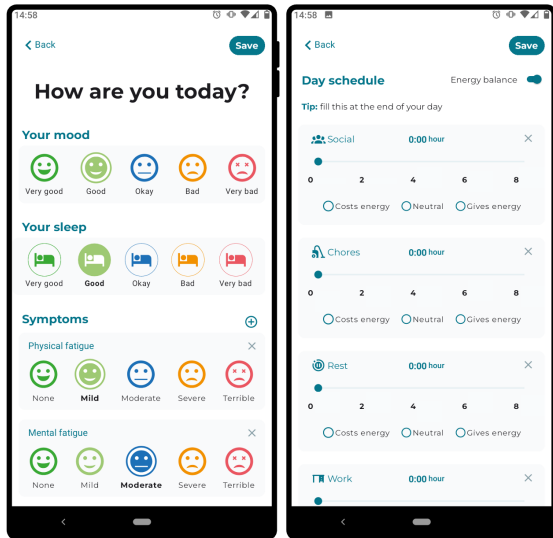


Figure 3: User interface of Healthy Chronos.

developed in cooperation with the Netherlands Cancer Institute and Dutch Cancer Society. The goal is to help this group cope with the physical and mental effects of dealing with cancer (treatment), though the app can accommodate any user-specified goal. Through reminders and surveys, the app helps users to score variables like mood, sleep, and energy (scores on five-point scales), activities like work, social activities, chores, and exercise (in hours), and goals, like resting at least 30 minutes (booleans) (Figure 3). Users receive push messages to submit scores at the end of the day, and users can submit arbitrary queries about their data. Healthy Chronos has a free version and one free year of premium use through a waiver.<sup>1</sup>

COACH was developed as an efficient Python pipeline to quickly generate counselling from user queries within the Healthy Chronos app. It was fully developed and evaluated in the Azure environment of Healthy Chronos, in line with the app’s privacy and security agreements with users. COACH’s full code and prompt are available on GitHub.<sup>2</sup>

COACH’s workflow is as follows (also see Figure 2). Once a single user query (or batch) is submitted to the application, COACH embeds the user query using the *text-embedding-3-small* encoder from OpenAI. COACH then scrapes a subset of Dutch webpages from *kanker.nl* that contain trustworthy information on cancer and lifestyle (advice) written by cancer information experts, on topics like how to deal with fatigue or talking about the disease with loved ones. It then converts them to

<sup>1</sup>See <https://healthychronos.com/en/> for more information.

<sup>2</sup><https://github.com/yeem4n/llm-lifestyle-coach/>

a vector store so that knowledge stays up to date. Thereafter, COACH retrieves four chunks most relevant to the query. COACH transforms the table and then appends it with chunks and a query to a prompt. Prompt design followed the following insights from recent work:

- Separating instructions and examples from ‘external’ data, i.e., table representation and chunks from the knowledge database; it is known that keeping external information at the end improves performance (Sui et al., 2024);
- Representing user tables in Markdown format, as LLMs are sensitive to different ways tabular inputs are serialized (Sui et al., 2024), and Markdown worked best for our use case (besides JSON and CSV).
- Providing descriptions of key column names and data types, e.g. “headers containing ‘goal’ always indicate booleans”, because headers are crucial to an LLM’s capacity to understand column contents and relations to other columns (Singha et al., 2023).
- Drawing on in-context learning (Fang et al., 2024), by describing three common topics in user queries (physical fatigue, mental fatigue, daily activities), the relevant columns in the user data. This ‘primes’ the LLM to consider relevant columns and possible relations between them.

Current work in prompt engineering research has further shown that, particularly in healthcare, a challenge is limited portability and scalability due to prompts becoming tied to specific domains, e.g., due to the use of terminology (Wang et al., 2025). Hence, we keep our approach straightforward, in that we leave specifics about the disease to the RAG component, and do not describe our medical population in detail. Our structuring and provision of examples in the prompt aim for portability, in that future users can easily adapt examples to their own data. In sum, our prompt is structured as follows:

1. A system message setting up the LLM as an AI assistant for addressing lifestyle-related questions for cancer patients;
2. Instructions for obtaining an empathic tone, brevity in the response, and refraining from providing medical advice;
3. A description of the user tables concerning headers (columns) and data types;
4. Several examples of question topics and the appropriate columns to consider;
5. Formatting instructions for separating user data analysis from counselling in the output;




Figure 4: Visualization of QUORUM outcomes. ‘Averages’ graph summarises relevance, quality, and reliability; ‘Relevance’ and ‘Quality’ summarize their respective dimensions by averaging over variable scores; ‘Reliability’ contains rescaled variable ratios (Table 1.)

Var.	Var.	Var.
Alignm. $3.90 \pm .94$	Correct $4.02 \pm .84$	Faithfulness .79
Fol. up $3.90 \pm 1.15$	Tone $3.73 \pm .75$	Completeness .97
Tone $4.32 \pm .82$	Length $4 \pm .57$	Hallucination .22
Length $4.16 \pm .69$		
Users	Experts	Developers

Table 2: Main evaluation results. For users, we report averages over 19 users; for experts, averages over in total 100 evaluations (50 response-query pairs evaluated by two experts); for developers, ratios over in total 66 claims about user data (faithfulness/completeness) and 31 contextualisation statements (hallucination).

- Concatenation of table, RAG chunks, and user query to instructions above.

We used LangChain, an open-source framework to streamline LLM and RAG components (Chase, 2022), to build COACH.

The LLM used in the pipeline is GPT-4o (version 2024-11-20, accessed October 2025), due to workflows already in place for `Healthy Chronos`, in which they can easily deploy the system through the Azure environment, and its strong performance on similar tasks. GPT-4o is currently still a good option in similar, recent systems and chatbots for lifestyle counselling (Jörke et al., 2025; Steenstra et al., 2024); suitably large open models could in future work take the place of GPT-4o.

## 4. Results

### 4.1. User Evaluation: Relevance

Main results of applying QUORUM are given in Table 2 and visually in Figure 4. As can be seen, overall, users responded positively to the generated counselling. Results suggest that counselling aligned with personal circumstances, that tips were likely to be followed, and that the tone and length were satisfying, as all variable mean scores are around 4. However, we also see that variation in

scores is larger for statements on alignment and follow-up, implying that users are more divided on this aspect than on tone and length.

**Remarks: Alignment** Inspection of user-submitted remarks provided more insight into what users appreciate in the counselling. Ten out of nineteen respondents submitted an additional remark, of which four respondents expressed that the response captured their situation very well, and was easy to understand. They also appreciated the level of detail in the answers, and the formatting that keeps claims about the data separated from the contextualisation (Figure 1). In addition, two users emphasized that the advice felt more valuable due to the addition of specific references to their own diary data, which also helped them to better understand the contextualisation.

**Remarks: Misalignment** Four respondents expressed a misalignment in the counselling and their situation. One user, for example, reported the counselling recommended to seek help from friends and family, whereas for this user, social activities were experienced as tiring, and as an increasing burden for the social network, i.e., something to avoid. Other users mentioned a misalignment with their situation that is strictly speaking not related to the query-response pair, such as a lack of added functionality in the app, or struggling with regularly submitting data at the end of the day. One user noted that the advice reiterated strategies they were already applying without positive outcomes, implying the generated response ideally takes into account previous counselling in the prompt context.

Our user results suggest that the pipeline succeeds in generating advice that is relevant for users.

### 4.2. Expert Evaluation: Information Quality

Overall, experts were positive about the quality of generated counselling (Table 2, Figure 4). They

Annotator	A (n=25)	B (n=25)	C (n=12)	D (n=13)	E (n=12)	F (n=13)
Correctness	4.08 ( $\pm 0.63$ )	4.88 ( $\pm 0.33$ )	3.33 ( $\pm 0.47$ )	4.15 ( $\pm 0.36$ )	2.75 ( $\pm 0.73$ )	3.92 ( $\pm 0.47$ )
Tone	4.12 ( $\pm 0.59$ )	3.8 ( $\pm 0.63$ )	3.42 ( $\pm 0.76$ )	3.7 ( $\pm 0.92$ )	3 ( $\pm 0.71$ )	3.85 ( $\pm 0.36$ )
Length	4.16 ( $\pm 0.46$ )	4.04 ( $\pm 0.34$ )	3.92 ( $\pm 0.28$ )	4.46 ( $\pm 0.63$ )	3.25 ( $\pm 0.6$ )	3.92 ( $\pm 0.47$ )

Table 3: Mean and standard deviation of each expert and variable. N shows the number of query-response pairs evaluated per expert; each was labelled twice, totalling 100 observations.

thought the counselling was often correct, and had mostly the right tone and length, with most scores close to 4, but with smaller variation for each variable compared to users. Across 50 samples evaluated by 6 independent experts (minimal 2 evaluations per item with overlapping annotator pairs), Krippendorff’s  $\alpha$  was 0.602 for correctness. This indicates moderate agreement: the experts generally concur about answer correctness. Examining the per-expert mean scores (Table 3), no average scores for correctness were found closer to two or one, suggesting disagreements reflect differences in annotator strictness rather than strong divergence about system correctness. Thus, the pipeline rarely produces counselling that experts consider incorrect or harmful.

**Remarks: Tone** Regarding tone, experts diverged. Krippendorff’s  $\alpha$  was 0.150 for tone and 0.272 for length, confirming that these dimensions were very subjective. Inspection of a total of 79 remarks revealed that about a quarter raised concerns about tone and language use, describing the counselling as too condescending, direct, or unsympathetic. Also, the use of exclamation marks, imperatives, and motivational phrases was sometimes viewed negatively. Thus, experts see room for possible improvements, given that their evaluations were, on average, positive.

**Remarks: Clarity** Experts addressed clarity in about half of the remarks. Twenty-two remarks explicated that the counselling was not at a basic level (B1) of Dutch, hence should be rephrased. Experts considered some counselling statements on their own as too ambiguous, due to the lack of specific examples or further explanation, for example, the phrase ‘listen to your own body’ or ‘stay true to your routines’, as it is not always clear what this entails. Experts emphasized that clarity and concreteness in lifestyle advice are important, as this renders counselling more actionable and memorable, in line with previous research (Szymanski et al., 2025; Degachi et al., 2025).

**Remarks: Grounding** Five experts praised the explicit grounding of counselling in diary data, acknowledging the value of including insights from the user’s diary. While some did note that care should

be taken in not overgeneralizing over too little data, the references to concrete days, activities, or symptoms were seen as increasing transparency and trustworthiness.

Overall, though experts were quite positive about the counselling, they provided many remarks about how the counselling could be improved regarding its tone. In this respect, they were more critical than users, a finding which we return to in the discussion.

### 4.3. Developer Evaluation: Reliability

Overall, we found that 79% of claims about user patterns were faithful, i.e., consistent with the data. That is, about four out of five claims about user data are supported by the user’s table. This also means that specific referenced dates, activities, or goals from the user data are often correct, that is, concrete examples for patterns found in the user data. Further, regarding completeness, in 97% of the claims (virtually all), the LLM succeeded in finding the relevant information, which often amounted to correctly retrieving relevant columns and their values to address a user query. Lastly, regarding hallucination, in 22% of statements made in the contextualisation, we found information not strictly present in the external knowledge database.

**Developer Agreement** The developer agreement was overall moderate across faithfulness, completeness, and hallucination (Cohen’s  $\kappa = 0.58$ ), indicating some challenges in evaluating claims and statements. In particular cases with more ambiguous phrases, e.g., ‘*The last few days...*’ or ‘*Scores are often between 3 and 5...*’ led to differing judgments, despite establishing a list with definitions for ambiguous terms (provided in the Supplementary Materials).

To better understand misalignment, we computed Positive Percentage Agreement for faithfulness, completeness, and RAG context adherence, which were 85.39%, 96%, and 88.89%, respectively, while Negative Percentage Agreement was 48%, 0%, and 44.44%, respectively. This indicates that the two developers often agree whenever one developer evaluates a claim as faithful, complete, or hallucination-free; however, developers have much more difficulty agreeing on the cases where, according to one of them, the LLM was unreliable,




Figure 5: Examples of various errors in claims about data and contextualisation statements. A): overgeneralization of common-sense patterns: B) missing pattern in the user data; C): hallucination: advice is not strictly present in chunks.

incomplete, or provided information not present in the knowledge database.

**Errors** Though in most cases patterns in the data were extrapolated accurately, for faithfulness most errors arose from overgeneralisation, in which the model saw a pattern in the data that was not sufficiently strong or had also strong counterexamples. Example A in Figure 5 illustrates this. Another issue was with missing patterns: sometimes COACH reported on a pattern that did not exist (example B in Figure 5). This could be seen as a type of hallucination, which in QUORUM is captured under faithfulness.

Regarding hallucinations in contextualisation statements, we observed that the model sometimes goes beyond retrieved chunks in generating advice. Note that we had a strict interpretation of hallucination in Table 1: if a contextualisation statement suggested consuming healthy, protein-rich food, and the counselling contained examples (nuts, yoghurt, etc.) not strictly in the database, it was regarded as a hallucination. Arguably, this could also be seen as a case of intuitive extrapolation rather than harmful content often associated with hallucinations. When such cases (example C in Figure 5) were not counted as hallucination, the hallucination rate declined to 11%. Remaining errors contained reporting of insights from user data, contrary to the instruction to the LLM in the prompt. We found no cases of hallucinations with wrong or harmful content.

## 5. Discussion & Conclusion

This study contributed QUORUM, a framework for evaluating relevance, quality, and reliability of AI systems in health from three perspectives. The framework allows a future user to quickly glance at how overall, and in what specific aspects, a given AI-driven system in health scores on topics important to different stakeholders (Figure 4). As we have argued, a good evaluation of AI-driven systems in health should incorporate all perspectives in the

evaluation. We expect that AI systems that evaluate on average positively on relevance, quality, and reliability, which in practice will mean around 4 or higher, will have higher chances of making a positive real-world impact.

In our case, we saw convergence on a positive evaluation in the three perspectives, but also room for divergence. Divergence may reveal issues crucial for future development. In our case, we observed in expert remarks that they took issue with the clarity of the generated counselling. Interestingly, this did not seem to resonate with the scores or remarks of users, who indicated it was easy to follow. This signals our sample may be biased towards people with higher educational levels, as medical experts may have experience with a more varied population; this should inform the next iteration of COACH.

Some divergence can also be found in sensitivity to errors. Alignment and correctness as variables are arguably for users and experts proxies for developers' more technical faithfulness and completeness metrics (though they do not fully overlap). Still, the scores and remarks of users and experts do not suggest grave errors in user-specific pattern extraction or hallucinated content from the database, though we know from the developer evaluation that roughly one in five claims is not consistent with user data or is not strictly traceable to the knowledge database Table 2. This could be because different stakeholders have their own backgrounds and biases that may impact their sensitivity to errors; QUORUM helps reveal such differences as a starting point for further investigation.

We observed that if a user asks, for example, 'Why am I so tired?', the retrieved chunks may contain biomedical details, for example, on how changes in the immune system after cancer (therapy) can cause fatigue. Still, the counselling often defaults to more general tips, for example, on sleep patterns instead of more biomedical facts. This could be an unwanted result of prompting the model to refrain from providing medical advice; yet, the point here is that neither users/experts scores or remarks seem to pick up on this. Especially experts

could, due to their training, expect more comprehensive explanations in the counselling, which also includes more biomedical facts. A next iteration of COACH could focus on encouraging the LLM to include biomedical facts in the answer, which is not the same as providing medical advice.

In conclusion, this paper demonstrated that relatively straightforward LLM-based systems like COACH can provide relevant, good-quality, and reliable lifestyle counselling, as exemplified in our use case. We aimed to show that a multi-stakeholder lens is needed to evaluate the likelihood of the real-world impact of any AI system in health. We expect that both QUORUM and COACH, as an easy-to-understand framework and pipeline, will benefit other developers, experts, and patient populations aiming to improve consumer health question answering based on user diary data.

## 6. Ethical Statement

Our study protocol for recruitment of and data collection from users, medical experts, and developers was approved by the Leiden University Ethics Committee.

## 7. Acknowledgments

We thank the users of `Healthy Chronos` who contributed as participants to this study. We are grateful to the cancer information experts from `kanker.nl` for their time and effort in evaluating the responses. We also thank `Healthy Chronos` for the collaboration and access to the valuable platform.

## 8. Limitations

Several limitations impact our work. While we demonstrate the potential utility of our pipeline COACH and evaluation framework QUORUM for extension to arbitrary populations and diseases, these were only tested in a single context and use case, with a small population of users who likely have a high digital literacy. We hope to validate COACH and QUORUM across diverse settings and applications in future research.

Reproducibility is constrained for multiple reasons. Due to privacy and intellectual property considerations, the user diary data and the knowledge database cannot be shared. The RAG component naturally relies on dynamically evolving knowledge sources, which also limits the ability to reproduce our experimental setup entirely. Nevertheless, we anticipate that the components of our pipeline can be adapted and reused in other contexts where user diary data is available.

In addition, our sample size of evaluated responses was relatively small. The sample size and lack of baseline comparisons were primarily driven by the limited availability of our evaluators. More research on automatic lifestyle counselling using the framework and pipeline would help strengthen our conclusions.

Furthermore, the LLM-generated counselling was based only on the most recent <21 days of diary data. Future research could investigate how performance changes when incorporating longer-term lifestyle patterns. Finally, our pipeline relies on a single proprietary LLM (GPT-4o), as discussed in [subsection 3.2](#). Future work should explore the use of open-source, and ideally smaller, models to improve accessibility, transparency, and reproducibility.

## 9. Bibliographical References

Tatsuo Akechi, Akira Kugaya, Hitoshi Okamura, Shigeto Yamawaki, and Yosuke Uchitomi. 1999. Fatigue and its associated factors in ambulatory cancer patients: a preliminary study. *Journal of pain and symptom management*, 17(1):42–48.

Alexander Bondarenko, Ekaterina Shirshakova, Marina Driker, Matthias Hagen, and Pavel Braslavski. 2021. [Misbeliefs and biases in health-related searches](#). In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management, CIKM '21*, page 2894–2899, New York, NY, USA. Association for Computing Machinery.

Nattapat Boonprakong. 2025. *Quantifying and Understanding the Effects of Cognitive Biases in Human-Computer Interaction*. Ph.D. thesis, The University of Melbourne, Melbourne, Australia.

Harrison Chase. 2022. [Langchain](#). Accessed: 2025-11-23.

Chadha Degachi, Ujjayan Dhar, Evangelos Niforatos, and Gerd Kortuem. 2025. [Towards a Domain Expert Evaluation Framework for Conversational Search in Healthcare](#). In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems, CHI EA '25*, pages 1–9, New York, NY, USA. Association for Computing Machinery.

Rachel Drbohlav Ollerton, Saar Hommes, Nadine Bol, Gwenn Beets, Frans Folkvord, and Emiel J. Kraemer. 2025. [Health-related lifestyle communication in oncology: Exploring the experiences and preferences of cancer survivors](#). *Health Communication*, page 1–10.

- Xi Fang, Weijie Xu, Fiona Anting Tan, Ziqing Hu, Jiani Zhang, Yanjun Qi, Srinivasan H. Sengamedu, and Christos Faloutsos. 2024. [Large language models \(LLMs\) on tabular data: Prediction, generation, and understanding - a survey](#). *Transactions on Machine Learning Research*.
- Meeke Hoedjes, Inge Nijman, and Chris Hinnen. 2022. Psychosocial determinants of lifestyle change after a cancer diagnosis: a systematic review of the literature. *Cancers*, 14(8):2026.
- Matthew Jörke, Shardul Sapkota, Lyndsea Warkenthien, Niklas Vainio, Paul Schmiedmayer, Emma Brunskill, and James A. Landay. 2025. [Gptcoach: Towards llm-based physical activity coaching](#). In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25, New York, NY, USA. Association for Computing Machinery.
- Lars Kayser, Andre Kushniruk, Richard H Osborne, Ole Norgaard, and Paul Turner. 2015. Enhancing the effectiveness of consumer-focused health information technology systems through ehealth literacy: a framework for understanding users' needs. *JMIR human factors*, 2(1):e3696.
- Xiangxun Lai, Yue Lai, JiaCheng Chen, Shengqi Huang, Qi Gao, and Caihua Huang. 2025. Evaluation strategies for large language model-based models in exercise and health coaching: Scoping review. *Journal of Medical Internet Research*, 27:e79217.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS '20, Red Hook, NY, USA. Curran Associates Inc.
- Karen Jia-Hui Li, Simone Balloccu, Ondrej Dusek, and Ehud Reiter. 2025. [When LLMs can't help: Real-world evaluation of LLMs in nutrition](#). In *Proceedings of the 18th International Natural Language Generation Conference*, pages 753–779, Hanoi, Vietnam. Association for Computational Linguistics.
- Mike A Merrill, Akshay Paruchuri, Naghmeh Rezaei, Geza Kovacs, Javier Perez, Yun Liu, Erik Schenck, Nova Hammerquist, Jake Sunshine, Shyam Tailor, et al. 2024. Transforming wearable data into health insights using large language model agents. *arXiv preprint arXiv:2406.06464*.
- Sophia Meywirth. 2024. Designing a large language model-based coaching intervention for lifestyle behavior change. In *International Conference on Design Science Research in Information Systems and Technology*, pages 81–94. Springer.
- Siddharth Nandagopal. 2025. [Tokens as currency: A novel framework to sustain ai adoption and profitability](#). *International Journal of Applied Information Systems*, 12(46):44–51.
- Diana Bar-Or Nirman, Ariel Weizman, and Amos Azaria. 2025. Fool me, fool me: User attitudes toward llm falsehoods. *IEEE Access*.
- George Papageorgiou, Vangelis Sarlis, Manolis Maragoudakis, Ioannis Magnisalis, and Christos Tjortjis. 2025. Evaluating faithfulness in agentic rag systems for e-governance applications using llm-based judging frameworks. *Big Data and Cognitive Computing*, 9(12):309.
- Alexander Pugachev, Ekaterina Artemova, Alexander Bondarenko, and Pavel Braslavski. 2023. Consumer health question answering using off-the-shelf components. In *Advances in Information Retrieval*, page 571–579. Springer Nature Switzerland.
- Rebecca Raeside, Si Si Jia, Julie Redfern, and Stephanie R Partridge. 2022. Navigating the online world of lifestyle health information: qualitative study with adolescents. *JMIR Pediatrics and Parenting*, 5(1):e35165.
- Deborah Raji, Emily Denton, Emily M. Bender, Alex Hanna, and Amandalynne Paullada. 2021. [Ai and the everything in the whole wide world benchmark](#). In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1.
- Ananya Singha, José Cambronero, Sumit Gulwani, Vu Le, and Chris Parnin. 2023. [Tabular representation, noisy operators, and impacts on table structure understanding tasks in LLMs](#). In *NeurIPS 2023 Second Table Representation Learning Workshop*.
- Ian Steenstra, Farnaz Nouraei, Mehdi Arjmand, and Timothy Bickmore. 2024. Virtual agents for alcohol use counseling: Exploring llm-powered motivational interviewing. In *Proceedings of the 24th ACM International Conference on Intelligent Virtual Agents*, pages 1–10.
- Yuan Sui, Mengyu Zhou, Mingjie Zhou, Shi Han, and Dongmei Zhang. 2024. [Table Meets LLM: Can Large Language Models Understand Structured Table Data? A Benchmark and Empirical Study](#). In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, pages 645–654, Merida Mexico. ACM.

- Annalisa Szymanski, Noah Ziems, Heather A. Eicher-Miller, Toby Jia-Jun Li, Meng Jiang, and Ronald A. Metoyer. 2025. [Limitations of the LLM-as-a-Judge Approach for Evaluating LLM Outputs in Expert Knowledge Tasks](#). In *Proceedings of the 30th International Conference on Intelligent User Interfaces, IUI '25*, pages 952–966, New York, NY, USA. Association for Computing Machinery.
- Marrit Annika Tuinman, Janine Nuver, Anke De Boer, Anne Looijmans, and Mariët Hagedoorn. 2024. [Lifestyle changes after cancer treatment in patients and their partners: a qualitative study](#). *Supportive Care in Cancer*, 32(4):248.
- Bente van Aken, Anna Manshanden, Willemieke Kroeze, Nadine Florisson, Barbara Groot-Sluijsmans, Meeke Hoedjes, Kristel van Asselt, Lies ter Beek, and Ingrid Steenhuis. 2025. [Experiences and needs of dutch cancer survivors regarding lifestyle counselling: a qualitative study](#). *BMC Cancer*, 25:1761.
- Bram M.A. Van Dijk, Armel E.J.L. Lefebvre, and Marco R. Spruit. 2025. [Welzijn.ai: Developing responsible conversational ai for the care of older people through stakeholder involvement](#). *Maturitas*, 199(108616).
- Julia EWC van Gemert-Pijnen, Nicol Nijland, Maarten Van Limburg, Hans C Ossebaard, Saskia M Kelders, Gunther Eysenbach, and Erwin R Seydel. 2011. A holistic framework to improve the uptake and impact of ehealth technologies. *Journal of medical Internet research*, 13(4):e1672.
- Jiaqi Wang, Enze Shi, Sigang Yu, Zihao Wu, Huawen Hu, Chong Ma, Haixing Dai, Qiushi Yang, Yanqing Kang, Jinru Wu, et al. 2025. Prompt engineering for healthcare: Methodologies and applications. *Meta-Radiology*, page 100190.
- Jiayin Wang, Weizhi Ma, Peijie Sun, Min Zhang, and Jian-Yun Nie. 2024a. Understanding user experience in large language model interactions. *arXiv preprint arXiv:2401.08329*.
- Jiayin Wang, Fengran Mo, Weizhi Ma, Peijie Sun, Min Zhang, and Jian-Yun Nie. 2024b. [A user-centric multi-intent benchmark for evaluating large language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 3588–3612, Miami, Florida, USA. Association for Computational Linguistics.