

Towards Expectation Detection in Language: A Case Study on Treatment Expectations in Reddit

Aswathy Velutharambath, Amelie Wühl

University of Stuttgart, Germany, IT University of Copenhagen, Denmark
aswathy.velutharambath@ims.uni-stuttgart.de, amwy@itu.dk

Abstract

Patients' expectations towards their treatment have a substantial effect on the treatments' success. While primarily studied in clinical settings, online patient platforms like medical subreddits may hold complementary insights: treatment expectations that patients feel unnecessary or uncomfortable to share elsewhere. Despite this, no studies examine what type of expectations users discuss online and how they express them. Presumably this is because expectations have not been studied in natural language processing (NLP) before. Therefore, we introduce the task of Expectation Detection, arguing that expectations are relevant for many applications, including opinion mining and product design. Subsequently, we present a case study for the medical domain, where expectations are particularly crucial to extract. We contribute RedHOTExpect, a corpus of Reddit posts (4.5K posts) to study expectations in this context. We use a large language model (LLM) to silver-label the data and validate its quality manually (label accuracy $\approx 78\%$). Based on this, we analyze which linguistic patterns characterize expectations and explore what patients expect and why. We find that optimism and proactive framing are more pronounced in posts about physical or treatment-related illnesses compared to mental-health contexts, and that in our dataset, patients mostly discuss benefits rather than negative outcomes. The RedHOTExpect corpus can be obtained from <https://www.ims.uni-stuttgart.de/data/RedHOTExpect>

Keywords: corpus creation, expectation, LLM annotation, social media health mining

1. Introduction

Expressions of expectations are ubiquitous in natural language. They hold valuable information to monitor public opinion toward policy changes (see Table 1, Ex. 1), detect student needs (Ex. 2), or inform product design (Ex. 3). In the medical domain they are particularly relevant: Expectations toward a treatment are powerful psychosomatic mechanisms known to affect treatment success (Bingel et al., 2011; Rief et al., 2016; Aulenkamp et al., 2023). Pain medication, for example, can be twice as effective if the patient expects the treatment to work, while expecting it not to work can cancel its effect (Bingel et al., 2011). So far, studying these effects relies on small-scale patient surveys. While indispensable, their scale is limited, which leads to the following constraints: they might fail to detect expectations that patients do not consider 'report-worthy', an issue that Duh et al. (2016) show for reports of adverse drug effects. Additionally, surveys are known to be biased, leaving knowledge gaps for underrepresented groups (Weber et al., 2021). Both hinder doctors from addressing concerns caused by negative expectations and to achieve effective treatment.

Social media, where patients frequently discuss medical experiences, enables us to study expectations and their effects in 'real-world' environments. Here, patients may share skepticism towards a treatment, even if they are uncomfortable or find it unnecessary to disclose their concerns with medical authorities (see Ex. 4 and 5). Similarly, marginalized groups who have experienced medical discrim-

id	Texts expressing expectation
1	All this talk about the new bus lanes. At least I'll have a new excuse for being late to work from now on.
2	I'm pretty optimistic that the new exam format might finally test how well we think, not how long we can sit still
3	This phone should last at least three years
4	How is this supposed to have an effect with all the weed I've been smoking?
5	It's a stupid reason, but everyone says how the meds can make you feel bloated.
6	I bet no one tested this for darker skin tones. Makes me scared to even fill the prescription

Table 1: Example texts expressing expectation across general and medical domain.

ination may distrust doctors' recommendations (see Ex. 6) and therefore discuss treatment expectations and concerns mostly online. Identifying such online discussions can help practitioners anticipate a wider range of concerns that may foster negative expectations.

However, we are lacking a crucial prerequisite to extract this knowledge: research on detecting expectations from text. While tasks like hope and regret detection (Chakravarthi, 2020; Balouchzahi et al., 2023a,b, 2024), or complaint detection (Jin and Aletras, 2020) are related to expectations, the linguistic characteristics and domain-specific properties of expectations are not studied. We assume

this is because identifying descriptions of expectations is difficult. As illustrated by the examples in Table 1, expectations are often implicit, nuanced, or require understanding contextual cues. Current LLMs may be capable of extracting them. To the best of our knowledge, however, they have not been evaluated for this task. This is likely because we still lack a comprehensive definition of relevant expectation types and resources for evaluation.

Therefore, we make the following contributions:

(1) We introduce *Expectation Detection* as a novel NLP task. The goal is to identify linguistic and stylistic patterns associated with expression of *expectation* in which a writer conveys a belief, anticipation, or prediction about a possible future outcome. Unlike sentiment or stance analysis, which evaluate attitudes about the present, past or future, expectation detection specifically targets *future-oriented mental states*: utterances in which a person articulates what they think, hope, or fear may happen.

(2) Considering the crucial impact expectations have on treatment success, we start exploring this task with a case study on patient-authored health texts. To this end, we contribute the first resource to study treatment expectations in social media text. The corpus, REDHOTEXPECT, consists of ≈ 4.5 K medical Reddit posts expressing expectation, of which around 2.5K include token-level annotations for treatment, expectations, and outcome descriptions (TEO triplets). Hypothesizing that LLMs have a rudimentary notion of ‘expectation’, we use an LLM to filter and silver-label expectations and outcomes. We manually validate the LLM annotations in a representative subset of the data (250 instances), and find them to be accurate in 77.5% of cases. For the same subset, we additionally characterize expectation and outcome mentions with respect to their type and the basis of expectation.

(3) Using this data, we study *how patients discuss treatment expectations on social media* by analyzing which linguistic and stylistic patterns characterize these utterances (§5). In a targeted manual analysis (§6, we explore *what patients expect and why*. We find that expectation-related discourse is characterized by a more optimistic and action-oriented tone, whereas non-expectation posts more often convey caution, frustration, or negative affect. Overall, we observe that optimism and proactive framing are most pronounced in posts about physical or treatment-related illnesses, whereas mental-health contexts convey expectations in a more introspective and socially detached tone. Notably, expectations in our dataset mostly discuss benefits rather than negative outcomes.

2. Related Work

Patients frequently use online platforms to discuss medical topics (Chen et al., 2021). Therefore, a subfield of medical NLP specializes on biomedical content from social media and other online sources. Central tasks are pharmacovigilance, i.e., identifying descriptions of adverse drug reactions (Nikfarjam et al., 2015; Cocos et al., 2017; Magge et al., 2021; Karimi et al., 2015), monitoring public health (Paul and Dredze, 2012; Choudhury et al., 2013; Sarker et al., 2016; Stefanidis et al., 2017), extracting personal health experiences (Yin et al., 2015; Klein et al., 2017; Karisani and Agichtein, 2018; Wadhwa et al., 2023; Falk and Lapesa, 2024), analyzing mental health discourse (Garg, 2023), and fact-checking (Hossain et al., 2020; Mattern et al., 2021; Saakyan et al., 2021; Kim et al., 2023, i.a.). Health mining datasets are commonly collected from patient fora (Karimi et al., 2015; Vladika et al., 2024) and various social media platforms. Until the company limited its API access, Twitter (now X) served as a popular data source (Sarker et al., 2016; Wüthrl and Klinger, 2022; Sundriyal et al., 2022, i.a.). Reddit is another frequently used source, particularly because posts are not length-restricted (Scepanovic et al., 2020; Basaldella et al., 2020; Wadhwa et al., 2023, i.a.). Recently, synthetic data also has become increasingly popular (Wüthrl et al., 2024; Yamagishi and Nakamura, 2024; Ghanadian et al., 2024, i.a.).

While connected to concepts such as hope (Balouchzahi et al., 2023b), stance (Hardalov et al., 2022), or emotion (Plaza-del Arco et al., 2024), which have been studied in NLP, expectation detection is an unexplored task. Some work addresses adjacent tasks: Jin and Aletras (2020) study complaint detection, which involves identifying when someone voices a mismatch between their expectation and reality. Similarly, work on conversational intent (Sakurai and Miyao, 2024) may capture the speakers’ expectation of how the conversation could evolve. The most closely related is work on hope and regret detection (Chakravarthi, 2020; Balouchzahi et al., 2023a,b, 2024). While an expression of hope conveys an expectation, it is limited to a positive stance towards the event. Regret always expresses a negative attitude towards a past event. Bringula et al. (2022) extract expectations in student essays using word clustering and sentiment counts, but do not model expectation expressions or their linguistic grounding.

Some work brings together concepts relevant to expectations with health tasks. With respect to emotion, (Khanpour and Caragea, 2018) detect patients’ emotional states from online posts. (Turcan et al., 2021) detect psychological stress us-

ing emotion-informed models, while others explore understanding emotional states towards medical events like pandemics (Ng et al., 2020; Sosea et al., 2022, i.a.). The work most related to our case study on extracting treatment expectations is Balouchzahi et al. (2024) who model expressions of hope and regret with respect to drug use, and Ettlin et al. (2025) who explore patient expectations in clinical self-reports using topic modeling and linguistic style analysis

To the best of our knowledge, we are the first to formalize the task of expectation detection in NLP. So far, no work on extracting treatment expectations from social media exist. This motivates creating a resource that allows us to study how patients discuss their expectations towards a medical treatment.

3. Expectation Detection

We formalize the task of *Expectation Detection* by defining what constitutes an expectation and what linguistic properties can be attributed to it.

Defining Expectation Event. We define an *Expectation Event* (EE) as any statement or implication about a possible or likely future state, whether beneficial, harmful, or neutral, that is explicitly linked to an action, intervention, or situation. In linguistic terms, EEs capture how people verbalize their belief, hope, or fear about a *future* outcome (See Table 1). To establish what constitutes an expectation and how it differs from related constructs such as sentiment or stance, we attribute the following properties to each Expectation Event:

- **Expectation Type:** the valence or direction of the anticipated outcome, such as *benefit*, *harm*, *no effect*, *worsening* or *mixed*, or *other*.
- **Expectation Basis:** the source or justification underlying the belief, e.g., *Personal* (based on one’s own past experience), *Social* (based on peers or community experiences), *Authority* (based on statements from doctors or experts), *Information/Media* (based on studies, advertisements, or online information), *Cultural* (stemming from social or cultural beliefs), *Self-efficacy* (based on belief in one’s ability to follow the treatment).
- **Certainty:** the degree of confidence or tentativeness expressed by the author (e.g., “*I hope it helps*” vs. “*It will definitely help*”).
- **Temporal Orientation:** whether the expectation is *prospective* (before the outcome is known) or *retrospective* (recalled after an outcome has occurred).

Figure 1 shows examples of expectations annotated with these dimensions.

Task Definition. This study examines expectations expressed in medical discourse. We specifically

Example 1: “My doctor prescribed me X. I really hope it will stop my headaches this time, but I’m a bit afraid it could make me dizzy.”

Expectation Events:

EE1 “I really hope it will stop my headaches”
Expectation Type: Benefit
Expectation Basis: Authority (doctor’s prescription)
Certainty: Moderate (“hope”)
Temporal Orientation: Prospective

EE2 “I’m a bit afraid it could make me dizzy.”
Expectation Type: Harm
Expectation Basis: None
Certainty: Low (“could”)
Temporal Orientation: Prospective

Example 2: “All this talk about the new bus lanes. At least I’ll have a new excuse for being late to work from now on”

Expectation Events:

EE1 bus lanes will lead to delays in traffic
Expectation Type: Worsening
Expectation Basis: Information/Media
Certainty: High
Temporal Orientation: Prospective

Figure 1: Annotated examples illustrating Expectation Events (EEs) in posts.

target **outcome-related expectations**, statements that express anticipated effects of a treatment or intervention that could plausibly lead to a health-related outcome.

For example, “*I think this new medication will finally help me sleep*” and “*I’m worried the drug might make me nauseous*” both describe future outcomes linked to a treatment and thus constitute Expectation Events. By contrast, “*I expect to see my doctor next week*” expresses an expectation but not one about a treatment outcome and therefore falls outside our scope.

Formally, given a text span T , the Expectation Detection task in this study involves two subtasks:

1. **Expectation Identification:** determine whether T contains one or more Expectation Events; and
2. **Treatment-Expectation-Outcome annotation:** if so, identify their key components: (a) the *treatment or intervention* about which the expectation is expressed, (b) the *expectation* (what is expected to occur) and (c) the *actual outcome* (what actually occurred).

4. Expectation Corpus

4.1. Corpus Creation

To systematically study how patients express expectations about medical treatments on social

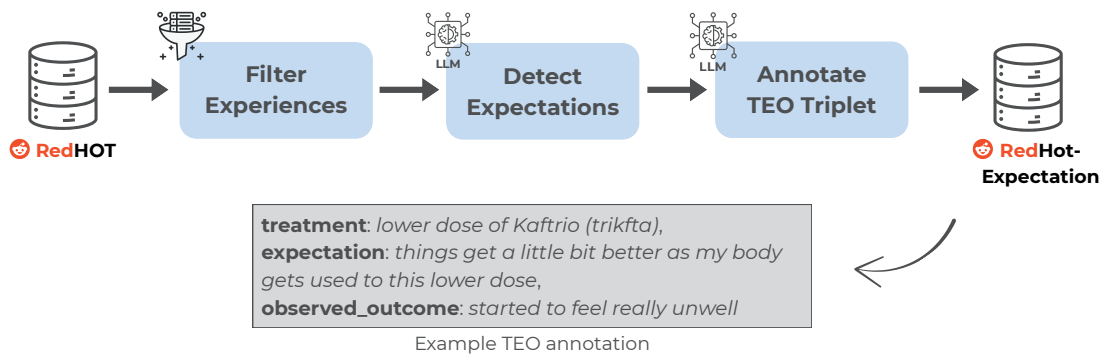


Figure 2: Automated pipeline for creating annotated RedHOT-Expectation data.

media, we require data that captures this concept. To facilitate this, we construct REDHOTEXPECT, a corpus of Reddit posts containing expressions of expectation, further annotated with Treatment–Expectation–Outcome (TEO) triplets. The following section outlines the steps involved in creating the corpus. The overall annotation process is illustrated in Figure 2.

Experience filtering. As a starting point, we build on the existing REDHOT corpus (Wadhwa et al., 2023), which contains $\approx 22k$ Reddit posts related to health conditions. We selected this dataset because it already focuses on patient-authored health discussions and provides high-level discourse annotations such as *questions*, *experiences*, and *claims*. This structure allows us to efficiently identify posts in which users describe their own medical experiences, which are the contexts where treatment expectations are most likely to appear. We therefore subset all posts labeled as *experience*, resulting in a collection of $\approx 12k$ posts. This filtered subset forms the foundation for identifying explicit mentions of expectations in the following step.

Expectation labeling. Identifying expressions of expectation in Reddit posts is challenging because the texts are often long, narrative, and loosely structured. Expectations are frequently implicit, context-dependent, and not always expressed within a continuous span of text. For instance, a post may introduce a medication or treatment at the beginning and describe the anticipated outcome only at a later stage. Given this complexity, we employ a large language model to leverage its broad semantic understanding and contextual reasoning capabilities for identifying posts that express expectations. Specifically, we use the open-weight model `gpt-oss-20b`¹ to classify each post as either *contains expectation* or *not*. The model is prompted with a detailed definition of outcome-related expectations and few-shot examples illustrating positive and negative cases. It is based on the expectation definition

¹<https://huggingface.co/openai/gpt-oss-20b>

Class / Subset	Count	Mean	Std	Median	Max
No expectation	7,364	136.6	124.7	102	1,771
Expectation (all)	4,433	182.9	170.8	138	2,944
has TEO	2,529	212.6	198.7	158	2,944
no TEO	1,904	143.6	113.0	113	1,042

Table 2: Token length statistics by class and within the subset of expectation posts, distinguishing those with and without extracted treatment–expectation–outcome (TEO) triplets.

and examples discussed in §3². We conduct several iterations of prompt refinement and manual verification to optimize the precision expectation detection.

Treatment–Expectation–Outcome triplet. In addition to examining how expectations are expressed linguistically, a further motivation for building the REDHOTEXPECT corpus is to enable analyses of how treatment expectations relate to reported outcomes. To support this goal, we extend the annotation to a structured *Treatment–Expectation–Outcome (TEO)* triplet, where the *Treatment* refers to the intervention or medication mentioned, the *Expectation* to the anticipated outcome, and the *Outcome* to the described or implied result. We again use the open-weight model `gpt-oss-20b`, providing detailed prompts to automatically generate these triplets.

Given the narrative style of Reddit posts, we instruct the model to extract one or more TEO triplets per post, typically one per treatment mention. Because Reddit threads often include exchanges between community members, an author may report the treatment outcome only later in the discussion. To account for this, we retrieve all comments posted by the same author and append them to the prompt, enabling the model to detect outcomes described in follow-up messages. This procedure allows us to capture both *prospective* expectations (before the outcome is known) and *retrospective* expectations (beliefs recalled after an outcome has occurred).

²See Appendix A.1 for the full prompt.

4.2. Human Validation

To evaluate how reliably the LLM extracts treatment–expectation–outcome (TEO) triplets we annotate a subset of instances manually. We select a stratified sample of 245 posts³ from the TEO annotated subset of the REDHOTEXPECT. To ensure diversity while maintaining relevance, we restrict the sample to condition–treatment pairs that occur at least three times in the corpus, guaranteeing sufficient contextual coverage across domains. The sampling procedure includes at least one instance of each eligible pair, with the remaining posts selected at random and balanced across health conditions.

Out of the 245 manually inspected posts, the automatic pipeline had extracted 502 Treatment–Expectation–Outcome (TEO) triplets. Manual validation confirms that 389 of these are correct, distributed across 200 posts, while 45 posts contained no valid triplet. This corresponds to a labeling accuracy of 77.5%, indicating that the automatic extraction process achieves a reasonably high level of precision while allowing for multiple expectations to be captured within a single post.

4.3. Corpus Details

The initial subset of posts drawn from the REDHOTdataset contained 11,797 Reddit posts covering 23 health conditions. Each post was automatically classified by a large language model as either containing or not containing an expectation-related expression, following the procedure described in §4.1. Overall, 4,433 posts (37.6%) were labeled as containing treatment expectations, suggesting that expectation-related discourse is a frequent phenomenon in patient-authored online discussions. The remaining 7,364 posts (62.4%) describe treatment experiences, advice, or symptom trajectories without explicitly expressing anticipation or belief about a future treatment outcome. In the TEO annotation step, out of the 4,433 posts labeled as *contains expectation*, 2,529 posts (57.0%) contained at least one complete triplet linking a treatment, an expectation, and an outcome.

From the manual validation step (§4.2) we observe an identification accuracy of approximately 78%, implying an expected error rate of around 22%. For the subsequent analyses, we should keep this in mind and interpret findings in light of the estimated labeling noise.

³Originally sampled 250, five exclude because expectations were labeled "None" or "null".

5. How do patients discuss treatment expectations on social media?

Our goal is to identify how language use differs between posts that express treatment expectations and other health-related narratives. We focus on linguistic and stylistic markers that may capture how people formulate predictions, intentions, or outcome-oriented thinking in health discourse.

5.1. Experimental setup

To examine linguistic markers of expectation in health-related discourse, we perform a lexical feature analysis using the LIWC-22 dictionary (Boyd et al., 2022). It provides frequency-based indicators of psychological and linguistic processes and is widely used in NLP and social media text analysis. In addition, we include sentiment analysis using the VADER lexicon (Hutto and Gilbert, 2014), the Gunning–Fog readability metric, and a lexical diversity measure⁴. From the LIWC lexicon, we select categories that we hypothesize to be relevant to expressions of expectation in health-related communication. As shown in Table 3, these categories are grouped into seven interpretable dimensions: *Uncertainty* (hedges and cognitive stance terms), *Time focus* (past, present, and future orientation), *Emotion* (positive and negative affect), *Social references* (self- and other-related pronouns), *Motivation* (goal- and need-related words), *Health* (illness and wellness vocabulary), and *Style* (readability and lexical diversity).

To test for linguistic differences between expectation and non-expectation posts, we use the Mann–Whitney U test, a non-parametric alternative to the two-sample *t*-test that does not assume normality. We report effect sizes using Cliff’s δ , indicating the direction and magnitude of group differences, and adjust all *p*-values for multiple comparisons using the Benjamini–Hochberg FDR procedure (Benjamini and Hochberg, 1995). This procedure is conducted both at the group level (aggregated dimensions) and at the individual-feature level to identify which specific LIWC categories drive the observed effects.

5.2. Results

As shown in Table 2, within this dataset, posts labeled as containing expectations exhibit higher token counts (median = 138) than those without expectations (median = 102). The Mann–Whitney test confirms that this difference is statistically significant. While this finding does not imply a general

⁴For extracting these features we use https://bitbucket.org/aswathyve/linguistic_style_features.

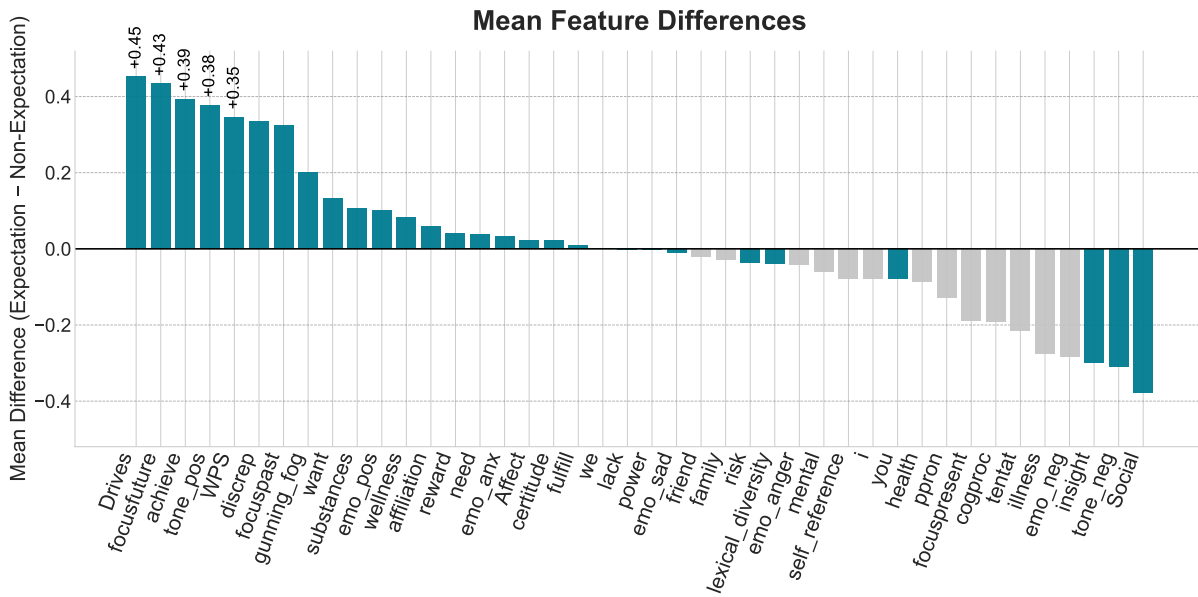


Figure 3: Mean feature difference in Expectation vs. Non-Expectations posts, with non-significant differences in grey.

Dimension	Features
Uncertainty	tentat, certitude, cogproc, discrep, insight
Time focus	focuspast, focuspresent, focusfuture
Emotion	Affect, tone_pos, tone_neg, emo_pos, emo_neg, emo_anx, emo_anger, emo_sad
Social ref.	self_reference, ppron, i, we, you, Social, family, friend
Motivation	Drives, achieve, power, affiliation, need, want, lack, fulfill, reward, risk
Health	health, illness, wellness, mental, substances
Style	WPS, lexical_diversity, gunning_fog

Table 3: LIWC feature groups used in the analysis.

relationship, it suggests that expectation-related content may be accompanied by more elaborate discourse in this sample.

As shown in Table 4, expectation-related posts differ systematically from other health narratives across several linguistic dimensions. The strongest effect is observed for **Motivation** ($\Delta = +0.11$), indicating greater use of goal-oriented and desire-expressing language. **Time focus** also shows a large overall difference ($\Delta = +0.19$), suggesting that expectation-related posts contain more temporal references in general, reflecting stronger temporal framing of events. When examining individual features within this category, the largest contribution comes from *focusfuture* ($\Delta = +0.43$), showing that people tend to talk more about future outcomes when expressing expectations. **Style** features also differ reliably ($\Delta = +0.18$), with posts expressing

Dimension	$\Delta(\text{Exp-Non})$	Cliff's δ	FDR- p
Motivation	+0.1084	0.1780	$< 10^{-41}$
Time focus	+0.1898	0.0912	$< 10^{-11}$
Style	+0.1800	0.0867	$< 10^{-9}$
Social refs.	-0.0985	-0.0389	0.0045
Health	-0.0465	0.0273	0.0482
Emotion	-0.0138	0.0229	0.0895
Uncertainty	-0.0689	-0.0146	0.2581

Table 4: Group-level LIWC differences (Mann–Whitney with Benjamini–Hochberg FDR). Δ is mean(Exp) – mean(Non); positive values indicate higher means in expectation posts.

expectation containing longer and more syntactically complex sentences. Smaller yet significant effects appear for **Social references** ($\Delta = -0.10$) and **Health** vocabulary ($\Delta = -0.05$), indicating fewer personal pronouns and health-specific terms in expectation posts. Differences for **Emotion** and **Uncertainty** were not statistically significant.

At the individual-feature level (Figure 3), expectation posts show increased use of words related to *Drives*, *achieve*, and *tone_pos*, highlighting a motivational and positive framing. Conversely, categories such as *risk*, *emo_anger*, *emo_neg*, and *tone_neg* tend to have higher mean scores in non-expectation posts. Although not all of these differences reach statistical significance, the overall pattern suggests that expectation-related discourse is characterized by a more optimistic and action-oriented tone, whereas non-expectation posts more often convey caution, frustration, or negative affect.

To examine if the observed patterns hold across different medical conditions, we group the condi-




Figure 4: Mean differences by condition types.

tions into seven categories: *Mental health*, *Autoimmune/Inflammatory*, *Endocrine/Metabolic*, *Neurological*, *Gastrointestinal*, *Respiratory*, and *Physical/Other*. As shown in Figure 4) across medical condition types, expectation-related posts consistently show a more future-oriented and motivational tone than non-expectation narratives. The effect is strongest in endocrine/metabolic and respiratory conditions, where posts were also longer and stylistically more complex, suggesting deliberate and goal-directed framing. In contrast, mental health discussions exhibited weaker differences and fewer social references, indicating that expectations in this context tend to be more introspective and self-focused. Overall, these patterns highlight that while the linguistic signature of expectation is broadly stable across domains, its expression varies with the nature of the health condition and the type of experience patients describe.

6. What patients expect and what actually happens?

While the linguistic analysis in §5 characterizes *how* expectations are expressed at the surface level (lexical and stylistic features), this section focuses on *what* patients expect and *why*. Therefore, we conduct a targeted manual analysis of a representative subset of the corpus.

As discussed in §4.2, by manually annotating the 245 posts for validating the quality of annotations, we identified 389 Treatment–Expectation–Outcome triplets. In this section, we annotate each of these triplets with the **Expectation type** (BENEFIT, HARM, WORSENING, NO EFFECT, MIXED, or OTHER), the **Expectation basis** (PERSONAL, SOCIAL, AUTHORITY, INFORMATION/MEDIA, CULTURAL, SELF-EFFICACY, or OTHER), and the **Outcome basis** (BENEFIT, HARM, WORSENING, NO EFFECT, MIXED, or OTHER). Together, these annotations provide a preliminary understanding of how treatment expectations are




Figure 5: Confusion matrix representing alignment of expectation and outcomes types.

formed, what factors motivate them, and how they relate to the outcomes that patients report.

Expectation Types. Within the manually verified triplets, expectations are overwhelmingly *benefit*-oriented: 372 out of 389 (95.6%) expectations are labeled as BENEFIT, compared to HARM (10; 2.6%) and UNCERTAIN (7; 1.8%); no instances were annotated as MIXED or NO EFFECT at the expectation level in this sample. In line with the observation made in §5.2, we see that in this limited sample, people predominantly articulate *benefit*-oriented expectations.

Expectation–Outcome alignment.

Figure 5 shows the alignment between the treatment expectations and the actual outcomes reported by the patients. As shown in Figure 5, most BENEFIT expectations are followed by BENEFIT outcomes (157/372; 42.2%), yet nearly one third correspond to NO EFFECT (111/372; 29.8%) and a smaller share to HARM (57/372; 15.3%) or MIXED (45/372; 12.1%). The few HARM expectations show some alignment with harmful outcomes (6/10; 60.0%), and UNCERTAIN expectations tend to accompany neutral or negative outcomes (4/7 and 2/7, respectively). Since majority of the sampled instances contain expectation if BENEFIT, it is hard to draw a strong conclusion. A larger-scale analyses is necessary to test whether expectations systematically align with outcomes in naturally occurring patient discourse.

Linguistic certainty and expectation strength. The strength or confidence in their expectations ("I think" vs. "I am absolutely certain") may also influence the corresponding treatment outcomes. To approximate this dimension, we derive a *Certainty score* from two LIWC features that capture epistemic stance: *certitude* and *tentat*.

$$\text{Certainty score} = z_{\text{certitude}} - z_{\text{tentat}}$$




Figure 6: Distribution of expectation basis across medical conditions grouped by condition types.

where, *certitude* contributes positively, reflecting confident or assertive phrasing, while *tentat* contributes negatively, capturing hedging or uncertainty. Higher Certainty score indicates more confident or decisive language, whereas lower value reflects tentative or reflective expression. We then compare this score between *aligned* (Expectation type = Outcome type) and *misaligned* (Expectation type \neq Outcome type) posts using a Mann–Whitney U test to assess distributional differences.

We observe a mean index close to zero for both aligned ($M = -0.027$) and misaligned ($M = 0.020$) instances, and the difference was not statistically significant (Mann–Whitney $p = 0.86$; Cliff’s $\delta = -0.011$). These results suggest that linguistic certainty, as captured by LIWC features is not predictive of expectation–outcome alignment.

Overall, this exploratory finding highlights that certainty of expectations may not be easily captured by lexicon-based measures. Future work should therefore include explicit certainty annotations or context-sensitive modeling (e.g., modality or confidence detection) to better capture how strongly patients believe in their stated expectations.

Expectation basis across health topics. Figure 6 illustrates how the rationale behind patients’ expectations varies across different medical conditions. Across all topics, expectations are predominantly grounded in *personal* experience, reflecting that most users draw on their own prior treatment history or symptom when anticipating outcomes. References to *authority* (e.g., medical professionals or prescriptions) constitute the second most common basis and appear more prominently in certain conditions such as Sinusitis, GERD, and Ankylosing Spondylitis, though without a consistent pattern

across condition types. Mentions of social or *info/media* sources occur infrequently, while *cultural* references are rare in the current subset. Overall, this pattern suggests that patients’ expectations in the current subset are shaped largely by individual experience and clinical authority rather than by broader social, informational, or cultural influences.

7. Conclusion

In this paper, we introduce the novel task of expectation detection which aims to capture how beliefs and anticipations about the future are expressed in language. To understand and formalize the notion of expectation, we conduct a case study on health communication, where expectations about treatment outcomes are known to have an influence on the effectiveness of the treatment. To enable this task, we create REDHOTEXPECT, a corpus of patient-authored Reddit posts expressing expectation and a subset of the dataset annotated with Treatment-Expectation-Outcome (TEO) triplets. Our analyses show that posts with expectations are longer, more future-oriented and motivational than other narratives. Manual verification of TEO triplets showed that expectations in the subset are mostly discussing benefits rather than negative outcomes.

Overall, this work takes a first step towards modeling expectation as a measurable linguistic phenomenon. With the qualitative analysis on the treatment expectations, basis and outcomes, we highlight the opportunities and challenges in computational research on expectation effects using large scale social media data. In future work, we aim to examine how expectations vary in their temporal orientation and level of certainty.

Ethical Considerations

The dataset used in this study builds on the publicly available REDHOTcorpus (Wadhwa et al., 2023), which was constructed in compliance with Reddit's terms of use and includes an opt-out procedure for users. We retrieved the data from the Reddit post identifiers provided in the original dataset and excluded instances that had been deleted by users. No personally identifying information was collected or shared, and we did not attempt to identify or contact any users. In line with the ethical principles established by the original authors, we do not redistribute raw post texts. Instead, we release the annotations, along with scripts needed to reproduce the dataset by retrieving the posts directly from Reddit⁵.

Modeling expectations in health-related discourse has potential benefits for understanding patient perspectives but also entails ethical responsibility. We therefore emphasize that all analyses are conducted at an aggregate level and are not intended to make inferences about individual users.

Limitations

The present study offers an initial step toward modeling expectations in language but has several limitations. The corpus is derived entirely from Reddit, representing a specific subset of patients (Duh et al., 2016) and a limited range of health-related subreddits; thus, the findings may not generalize to other populations, platforms, or clinical contexts. The results should therefore be interpreted with the understanding that this is a case study, and further investigation is needed to draw more robust conclusions about the linguistic properties of expectation.

Further, the Treatment–Expectation–Outcome annotations were generated automatically using an LLM, and although manual validation indicated good accuracy, the data inevitably contain noise and inconsistencies. Moreover, LLM-based annotations are sensitive to prompt design and model behavior, which may introduce systematic biases that are difficult to control.

Acknowledgements

This work is part of the SoftwareCampus project PLACEBO, which is funded by German Federal Ministry of Research, Technology and Space (BMFTR) under the grant 01IS23072.

⁵The dataset and script for downloading are available at <https://www.ims.uni-stuttgart.de/data/RedHOTExpect>

References

- Jan Lukas Aulenkamp, Anita Icenhour, and Susan Elsenbruch. 2023. [Nocebo effects in visceral pain: concept and design of the experimental randomized-controlled pain study 'novis'](#). *Frontiers in psychiatry*, 14:1270189.
- Fazlourrahman Balouchzahi, Sabur Butt, Abeed Sarker, Al-Garadi MA, Grigori Sidorov, and Alexander Gelbukh. 2024. Analysis of expressions of hope and regret associated with nonmedical prescription drug use in x chatter. *Available at SSRN 4930517*.
- Fazlourrahman Balouchzahi, Sabur Butt, Grigori Sidorov, and Alexander Gelbukh. 2023a. [Reddit: Regret detection and domain identification from text](#). *Expert Systems with Applications*, 225:120099.
- Fazlourrahman Balouchzahi, Grigori Sidorov, and Alexander Gelbukh. 2023b. Polyhope: Two-level hope speech detection from tweets. *Expert Systems with Applications*, 225:120078.
- Marco Basaldella, Fangyu Liu, Ehsan Shareghi, and Nigel Collier. 2020. [COMETA: A corpus for medical entity linking in the social media](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3122–3137, Online. Association for Computational Linguistics.
- Yoav Benjamini and Yosef Hochberg. 1995. [Controlling the false discovery rate: A practical and powerful approach to multiple testing](#). *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1):289–300.
- Ulrike Bingel, Vishvarani Wanigasekera, Katja Wiech, Rhiannon Ni Mhuirheartaigh, Minyeong C Lee, Markus Ploner, and Irene Tracey. 2011. [The effect of treatment expectation on drug efficacy: imaging the analgesic benefit of the opioid remifentanyl](#). *Science translational medicine*, 3(70):70ra14.
- Ryan L. Boyd, Ashwin Ashokkumar, Sajjad Seraj, and James W. Pennebaker. 2022. [The development and psychometric properties of liwc-22](#). Technical report, University of Texas at Austin, Austin, TX.
- Rex Bringula, Saida Ulfa, John Paul P. Miranda, and Francis Arlando L. Atienza. 2022. [Text mining analysis on students' expectations and anxieties towards data analytics course](#). *Cogent Engineering*, 9(1):2127469.

- Bharathi Raja Chakravarthi. 2020. [HopeEDI: A multilingual hope speech detection dataset for equality, diversity, and inclusion](#). In *Proceedings of the Third Workshop on Computational Modeling of People's Opinions, Personality, and Emotions in Social Media*, pages 41–53, Barcelona, Spain (Online). Association for Computational Linguistics.
- Junhan Chen, Yuan Wang, et al. 2021. Social media use for health purposes: systematic review. *Journal of medical Internet research*, 23(5):e17917.
- Munmun De Choudhury, Scott Counts, and Eric Horvitz. 2013. [Social media as a measurement tool of depression in populations](#). In *Proceedings of the 5th ACM International Conference on Web Science (Paris, France, May 2-May 4, 2013)*. *WebSci 2013*.
- Anne Cocos, Alexander G Fiks, and Aaron J Masino. 2017. Deep learning for pharmacovigilance: recurrent neural network architectures for labeling adverse drug reactions in twitter posts. *Journal of the American Medical Informatics Association*, 24(4):813–821.
- Mei Sheng Duh, Pierre Cremieux, Marc Van Audenrode, Francis Vekeman, Paul Karner, Haimin Zhang, and Paul Greenberg. 2016. [Can social media data lead to earlier detection of drug-related adverse events?](#) *Pharmacoepidemiology and drug safety*, 25(12):1425–1433.
- Dominik A. Ettlin, Miriam Wolf, Nikola Biller-Andorno, and Gerhard Schneider. 2025. [In patients' words: natural language processing of reports from patients experiencing orofacial pain and dysfunction](#). *The Journal of Headache and Pain*, 26(1):172.
- Neele Falk and Gabriella Lapesa. 2024. [Stories and personal experiences in the COVID-19 discourse](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 15320–15340, Torino, Italia. ELRA and ICCL.
- Muskan Garg. 2023. [Mental health analysis in social media posts: A survey](#). *Archives of Computational Methods in Engineering*, 30(3):1819–1842.
- Hamideh Ghanadian, Isar Nejadgholi, and Hussein Al Osman. 2024. [Socially aware synthetic data generation for suicidal ideation detection using large language models](#). *IEEE Access*, 12:14350–14363.
- Momchil Hardalov, Arnav Arora, Preslav Nakov, and Isabelle Augenstein. 2022. [A survey on stance detection for mis- and disinformation identification](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1259–1277, Seattle, United States. Association for Computational Linguistics.
- Tamanna Hossain, Robert L. Logan IV, Arjuna Ugarte, Yoshitomo Matsubara, Sean Young, and Sameer Singh. 2020. [COVIDLies: Detecting COVID-19 misinformation on social media](#). In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*, Online. Association for Computational Linguistics.
- C. Hutto and Eric Gilbert. 2014. [Vader: A parsimonious rule-based model for sentiment analysis of social media text](#). *Proceedings of the International AAAI Conference on Web and Social Media*, 8(1):216–225.
- Mali Jin and Nikolaos Aletras. 2020. [Complaint identification in social media with transformer networks](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1765–1771, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Sarvnaz Karimi, Alejandro Metke-Jimenez, Madonna Kemp, and Chen Wang. 2015. [Cadec: A corpus of adverse drug event annotations](#). *Journal of Biomedical Informatics*, 55:73–81.
- Payam Karisani and Eugene Agichtein. 2018. [Did you really just have a heart attack? Towards robust detection of personal health mentions in social media](#). In *Proceedings of the 2018 World Wide Web Conference*, page 137–146, Republic and Canton of Geneva, CHE.
- Hamed Khanpour and Cornelia Caragea. 2018. [Fine-grained emotion detection in health-related online posts](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1160–1166, Brussels, Belgium. Association for Computational Linguistics.
- Jongin Kim, Byeol Rhee Bak, Aditya Agrawal, Jiaxi Wu, Veronika Wirtz, Traci Hong, and Derry Wijaya. 2023. [COVID-19 vaccine misinformation in middle income countries](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3903–3915, Singapore. Association for Computational Linguistics.
- Ari Klein, Abeed Sarker, Masoud Rouhizadeh, Karen O'Connor, and Graciela Gonzalez. 2017. [Detecting personal medication intake in Twitter: An annotated corpus and baseline classification system](#). In *BioNLP 2017*, pages 136–142, Vancouver, Canada,. Association for Computational Linguistics.

- Arjun Magge, Elena Tutubalina, Zulfat Miftahudinov, Ilseyar Alimova, Anne Dirkson, Suzan Verberne, Davy Weissenbacher, and Graciela Gonzalez-Hernandez. 2021. [DeepADEMiner: a deep learning pharmacovigilance pipeline for extraction and normalization of adverse drug event mentions on Twitter](#). *Journal of the American Medical Informatics Association*, 28(10):2184–2192.
- Justus Mattern, Yu Qiao, Elma Kerz, Daniel Wiechmann, and Markus Strohmaier. 2021. [FANG-COVID: A new large-scale benchmark dataset for fake news detection in German](#). In *Proceedings of the Fourth Workshop on Fact Extraction and VERification (FEVER)*, pages 78–91, Dominican Republic. Association for Computational Linguistics.
- Hui Xian Lynnette Ng, Roy Ka-Wei Lee, and Md Rabiul Awal. 2020. [I miss you babe: Analyzing emotion dynamics during COVID-19 pandemic](#). In *Proceedings of the Fourth Workshop on Natural Language Processing and Computational Social Science*, pages 41–49, Online. Association for Computational Linguistics.
- Azadeh Nikfarjam, Abeed Sarker, Karen O'Connor, Rachel Ginn, and Graciela Gonzalez. 2015. [Pharmacovigilance from social media: mining adverse drug reaction mentions using sequence labeling with word embedding cluster features](#). *Journal of the American Medical Informatics Association*, 22(3):671–681.
- Michael J. Paul and Mark Dredze. 2012. [A model for mining public health topics from Twitter](#). *Health*, 11(1).
- Flor Miriam Plaza-del Arco, Alba A. Cercas Curry, Amanda Cercas Curry, and Dirk Hovy. 2024. [Emotion analysis in NLP: Trends, gaps and roadmap for future directions](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 5696–5710, Torino, Italia. ELRA and ICCL.
- Winfried Rief, Arthur J Barsky, Ulrike Bingel, Bettina K Doering, Ralph Schwarting, Markus Wöhr, and Ulrike Schweiger. 2016. [Rethinking psychopharmacotherapy: The role of treatment context and brain plasticity in antidepressant and antipsychotic interventions](#). *Neuroscience and biobehavioral reviews*, 60:51–64.
- Arkadiy Saakyan, Tuhin Chakrabarty, and Smaranda Muresan. 2021. [COVID-fact: Fact extraction and verification of real-world claims on COVID-19 pandemic](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2116–2129, Online. Association for Computational Linguistics.
- Hiromasa Sakurai and Yusuke Miyao. 2024. [Evaluating intention detection capability of large language models in persuasive dialogues](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1635–1657, Bangkok, Thailand. Association for Computational Linguistics.
- Abeed Sarker, Karen O'Connor, Rachel Ginn, Matthew Scotch, Karen Smith, Dan Malone, and Graciela Gonzalez. 2016. [Social media mining for toxicovigilance: automatic monitoring of prescription medication abuse from Twitter](#). *Drug safety*, 39(3):231–240.
- Sanja Scepanovic, Enrique Martin-Lopez, Daniele Quercia, and Khan Baykaner. 2020. [Extracting medical entities from social media](#). In *Proceedings of the ACM Conference on Health, Inference, and Learning, CHIL '20*, pages 170–181. Association for Computing Machinery. Event-place: Toronto, Ontario, Canada.
- Tiberiu Sosea, Chau Pham, Alexander Tekle, Cornelia Caragea, and Junyi Jessy Li. 2022. [Emotion analysis and detection during COVID-19](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 6938–6947, Marseille, France. European Language Resources Association.
- Anthony Stefanidis, Emily Vraga, Georgios Lamprianidis, Jacek Radzikowski, Paul L Delamater, Kathryn H Jacobsen, Dieter Pfoser, Arie Croitoru, and Andrew Crooks. 2017. [Zika in twitter: Temporal variations of locations, actors, and concepts](#). *JMIR Public Health Surveill*, 3(2):e22.
- Megha Sundriyal, Atharva Kulkarni, Vaibhav Pulastya, Md. Shad Akhtar, and Tanmoy Chakraborty. 2022. [Empowering the fact-checkers! automatic identification of claim spans on Twitter](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 7701–7715, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Elsbeth Turcan, Smaranda Muresan, and Kathleen McKeown. 2021. [Emotion-infused models for explainable psychological stress detection](#). In *Proceedings of the 2021 Conference of the North*

American Chapter of the Association for Computational Linguistics: Human Language Technologies, pages 2895–2909, Online. Association for Computational Linguistics.

Juraj Vladika, Phillip Schneider, and Florian Matthes. 2024. [HealthFC: Verifying health claims with evidence-based medical fact-checking](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 8095–8107, Torino, Italia. ELRA and ICCL.

Somin Wadhwa, Vivek Khetan, Silvio Amir, and Byron Wallace. 2023. [RedHOT: A corpus of annotated medical questions, experiences, and claims on social media](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 809–827, Dubrovnik, Croatia. Association for Computational Linguistics.

Ann M. Weber, Ribhav Gupta, Safa Abdalla, Beniamino Cislighi, Valerie Meausoone, and Gary L. Darmstadt. 2021. [Gender-related data missingness, imbalance and bias in global health surveys](#). *BMJ global health*, 6(11).

Amelie Wührl, Lynn Greschner, Yarik Menchaca Resendiz, and Roman Klinger. 2024. [IMS_medicalY at #SMM4H 2024: Detecting impacts of outdoor spaces on social anxiety with data augmented ensembling](#). In *Proceedings of The 9th Social Media Mining for Health Research and Applications (SMM4H 2024) Workshop and Shared Tasks*, pages 83–87, Bangkok, Thailand. Association for Computational Linguistics.

Amelie Wührl and Roman Klinger. 2022. [Recovering patient journeys: A corpus of biomedical entities and relations on Twitter \(BEAR\)](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4439–4450, Marseille, France. European Language Resources Association.

Yosuke Yamagishi and Yuta Nakamura. 2024. [UTRad-NLP at #SMM4H 2024: Why LLM-generated texts fail to improve text classification models](#). In *Proceedings of the 9th Social Media Mining for Health Research and Applications (SMM4H 2024) Workshop and Shared Tasks*, pages 42–47, Bangkok, Thailand. Association for Computational Linguistics.

Zhijun Yin, Daniel Fabbri, S Trent Rosenbloom, and Bradley Malin. 2015. [A scalable framework to detect personal health mentions on Twitter](#). *Journal of Medical Internet Research*, 17(6):e138.

A. Appendix

A.1. Prompt Design

In this section, we provide the prompts used for our annotation pipeline. We employ two types of prompts: (1) a binary classification prompt for detecting the presence of Expectation Events (EEs), and (2) a structured extraction prompt for identifying Treatment–Expectation–Outcome (TEO) triplets. Both prompts are designed to guide the model toward consistent and task-specific outputs, with clearly defined instructions, constraints, and output formats. The expectation detection prompt targets the identification of outcome-related expectations in text. The TEC prompt focuses on extracting fine-grained relational information between treatments, expectations, and outcomes. The full prompts are provided below for reproducibility.

A.1.1. Expectation Detection Prompt

System:

You are an expert annotator. Your task is to determine whether the text
↪ expresses an "Expectation Event (EE)."

Definition of Expectation Event (EE):

An EE is a statement or implication about a possible or likely outcome (benefit,
↪ harm, no change, or worsening of condition) tied to a treatment, medication,
↪ or care action.

Expectations may be prospective (future-oriented, concerning possible or likely
↪ outcomes such as benefit, harm, no change, or worsening) or retrospective
↪ (recalling a past belief about what was expected, in relation to the outcome
↪ that actually occurred).

Important:

- The expectation must concern the ****treatment outcome****, not a possible
↪ diagnosis, appointment, or unrelated event.
- Outcome-related expectations include:
 - Benefit: expected symptom relief, functional improvement, disease control.
 - Adverse Effects: expected side effects or emotional impact.
 - No Effect: skepticism about benefit.
 - Worsening: anticipated deterioration or progression.
- Expectations can be:
 - Prospective (future-oriented) or retrospective (recalling a past belief).
 - Explicit (directly stated) or implicit (implied through goals/conditions).
 - Positive, negative, neutral, or mixed.

Here are some examples:

Expectation about treatment outcome counts:

- ``I expect this new medication will reduce my pain within two weeks."`
- ``If the surgery works, I can get back to work."`
- ``I was worried the drug would make me nauseous." (retrospective)`

Non-outcome expectation does NOT count:

- ``I expect to see my doctor again next month for a check-up."`
- ``I expect they will finally diagnose me with lupus."`

Your output should strictly be in JSON format:

- `{{"expectation": true}}`
- `{{"expectation": false}}`

Respond ONLY with the JSON object, nothing else. Answer with very minimal
↪ reasoning.

Avoid step-by-step explanations unless strictly necessary.

User:

The given text is a Reddit post where a person shares their experience related
↪ to a medical condition or treatment.

Determine whether the text expresses an Expectation Event (EE).

Strictly respond ONLY with a JSON of format: {"expectation": true} or
↪ {"expectation": false}.

Text: {post}

A.1.2. TEC Prompt Design

System:

You are an expert annotator. Extract all (Treatment, Expectation, Outcome)
↪ triplets (TEC) from a Reddit post and, if present, the author's own
↪ follow-up comments.

Definitions:

- Treatment: any medication, surgery, therapy, care action, or dosage/regimen
↪ change.
- Expectation: the author's belief/anticipation about what the treatment will do
↪ (benefit, harm, no effect), explicit or implicit. Return exact spans if
↪ explicitly stated; summarize if only implied.
- Outcome (observed_outcome): the realized/actual result reported after the
↪ treatment (improvement, worsening, side effect, or no change).

Rules:

1. UNIT OF ANNOTATION = Triplet
 - Each triplet must connect exactly ONE treatment to ONE expectation to ONE
↪ outcome (observed_outcome).
 - Never mix expectations/outcomes from different treatments.
 - Return exact spans if explicitly specified, summarize if implicitly
↪ specified, else return null.
2. EXPECTATION REQUIREMENT
 - Must be about treatment outcomes (benefit, harm, or no effect) that the
↪ patient is undergoing or had undergone.
 - Strictly exclude diagnosis expectations, appointments, or logistics.
 - Both explicit and implicit expectations count.
 - Explicit: ``Antibiotics will clear this up."`
 - Implicit: ``Antibiotics didn't help" (implies expectation they should
↪ have helped).
 - If no expectation is expressed for a treatment, no triplet is created.
3. OBSERVED OUTCOME
 - observed_outcome = what actually happened after the treatment (symptom
↪ change or side effect).
 - If only an expectation is given and no result is mentioned, set
↪ observed_outcome = null.
 - If both an expectation and a statement about what happened are present, use
↪ that statement as observed_outcome.
 - If the author later provides the result in a follow-up comment, update the
↪ triplet with that observed_outcome.
 - IMPORTANT: Receiving or taking the treatment itself (e.g., ``got the flu
↪ shot", ``took the pill") is NOT a treatment related outcome.
4. TREATMENT REQUIREMENT
 - Must be a concrete intervention (drug, surgery, therapy, self-care,
↪ dosage/regimen change).
 - Vague ``it/this" counts only if clearly referring to a treatment.
 - If expectation/outcome exists without a treatment, set treatment = null.
5. NULL HANDLING (STRICT)
 - If one component is missing, set it explicitly to null.
 - Example: ``Let's see if it works" -> treatment = null, expectation = ``see
↪ if it works", observed_outcome = null.

6. COMMENTS

- Look into the author's own follow-up comments for additional
 - ↪ expectations/outcomes.
- Ignore generic conversational replies (e.g., ``This is miserable'', ``Did you get surgery?").
 - ↪ you get surgery?").
- If later comments update the outcome for the same treatment/regimen, merge
 - ↪ into the original triplet (do not duplicate).
- If a comment introduces a new treatment/regimen, create a separate triplet.

7. MULTIPLE TRIPLETS

- If multiple treatments/expectations/outcomes appear, output all of them
 - ↪ separately.
- Do not merge across different treatments.

8. IF NOTHING FOUND

- If no valid triplets are detected, return:
 - { "triplets": [] }

Output format:

```
{ "triplets": [ { "treatment": "...", "expectation": "...", "observed_outcome":  
↪ "..." } ] }
```

Examples:

Text: "Doctor gave me antibiotics, I hope they clear my infection. Using a nasal
↪ spray too but it doesn't help my congestion."

```
-> { "triplets": [  
  { "treatment": "antibiotics", "expectation": "hope they clear my  
    ↪ infection", "observed_outcome": null },  
  { "treatment": "nasal spray", "expectation": "should help with congestion",  
    ↪ "observed_outcome": "no improvement" }  
] }
```

Text: "Started melatonin, hoping I'll sleep better."

Comment: "Update: I'm sleeping through the night now."

```
-> { "triplets": [  
  { "treatment": "melatonin", "expectation": "hoping I'll sleep better",  
    ↪ "observed_outcome": "sleeping through the night" }  
] }
```

Text: "Zolofit 25mg didn't do much. Upped to 50mg; mood is finally better."

```
-> { "triplets": [  
  { "treatment": "Zolofit 25mg", "expectation": "should help with mood",  
    ↪ "observed_outcome": "no improvement" },  
  { "treatment": "Zolofit 50mg", "expectation": "should help with mood",  
    ↪ "observed_outcome": "mood improvement" }  
] }
```

Text: "I thought it wouldn't work, but it did. Scheduled to see ENT next week."

```
-> { "triplets": [  
  { "treatment": null, "expectation": "thought it wouldn't work",  
    ↪ "observed_outcome": "improvement" }  
] }
```

Use null for any missing field.

Respond ONLY with the JSON object, nothing else.

Avoid repeating the post. Keep reasoning to very low. Strictly avoid step by
↪ step reasoning.

User:

Extract all (Treatment, Expectation, Outcome) triplets from the Reddit post and
↪ author comments.

If any component is missing, set it explicitly to null.
Respond ONLY with the JSON object, following this format exactly:
{ "triplets": [{ "treatment": "...", "expectation": "...",
→ "observed_outcome": "..." }] }

Avoid repeating the post. Keep reasoning to very low. Strictly avoid step by
→ step reasoning.

POST: {post}
COMMENTS: {author_comments}