Reducing Privacy Risks in Online Self-Disclosure

Anonymous ACL submission

Abstract

Self-disclosure, while being common and rewarding in social media interaction, also poses privacy risks. In this paper, we take the initiative to protect the user-side privacy associated with online self-disclosure through *detection* and *abstraction*. We develop a taxonomy of 19 self-disclosure categories and curate a large corpus consisting of 4.8K annotated disclosure spans. We then fine-tune a language model for detection, achieving over 75% Token F_1. We further conduct an HCI user study, with 82% of participants viewing the model positively, highlighting its real-world applicability. Motivated by the user feedback, we introduce the task of self-disclosure abstraction, which is paraphrasing disclosures into less specific terms while preserving their utility, e.g., *"Im 16F"* to *"I'm a teenage girl"*. We explore various fine-tuning strategies, and our best model can generate diverse abstractions that moderately reduce privacy risks while maintaining high utility according to human evaluation. To help users in deciding which disclosures to abstract, we present a task of rating their importance for context understanding. Our fine-tuned model achieves 80% accuracy, on-par with GPT-3.5.

1 Introduction

Self-disclosure — *the communication of personal information to others* (Jourard, 1971; Cozby, 1973) — is prevalent in online public discourse. Disclosing personal information allows users to seek social support, build community, solicit context-specific advice, and explore aspects of their identity that they feel unsafe exploring offline (Luo and Hancock, 2020). Consider the following (hypothetical, but representative) Reddit post:

Im 16F I think I want to be a bi M

The author discloses their age, gender, and sexual orientation to express themselves. However, these self-disclosures simultaneously expose them to privacy risks, notably regret of the disclosure

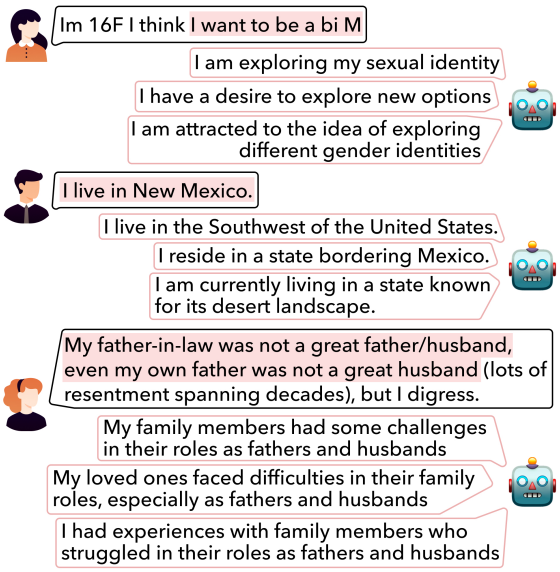

Figure 1: Our model can provide diverse abstractions for self-disclosures of any length to suit user preferences. This approach effectively reduces privacy risks without losing the essence of the message.

(Sleeper, 2016) and doxxing (Staab et al., 2023), which are particularly acute for marginalized populations (Lerner et al., 2020). This raises a critical question: **How can we help users identify and mitigate privacy risks in online self-disclosures?**

Prior works on self-disclosure (Valizadeh et al., 2021a; Cho et al., 2022; Staab et al., 2023, *inter alia*) and anonymization tools (Lison et al., 2021) focus on only a limited set of self-disclosures (e.g., health issues) or inferring personal attributes (aka. user profiling), often at sentence/post levels. They do not pinpoint the exact words of disclosures in the sentence, nor have broad enough coverage of different kinds of disclosures. Both are crucial for real-world users to take control of what they want to disclose and protect their privacy.

In this work, we take the important first steps in protecting user-side privacy with broad-coverage **self-disclosure detection** and **abstraction**. Our model is more extensive in capturing more categories of disclosure (see Table 1). Moreover, we

Category	#Spans	Avg Len	Example
Demographic Attributes			
LOCATION	525	5.70±3.85	I live in the UK and a diagnosis is really expensive, even with health insurance
AGE	308	2.93±1.72	I am a 23-year-old who is currently going through the last leg of undergraduate school
RELATIONSHIP STATUS	287	6.72±5.97	My partner has not helped at all, and I'm bed ridden now
AGE/GENDER	248	1.42±0.71	For some context, I (20F), still live with my parents
PET	192	6.93±7.31	Hi, I have two musk turtles and have never had any health problems before at all
APPEARANCE	173	6.96±6.25	Same here. I am 6'2. No one can sit behind me.
HUSBAND/BF	148	6.89 ±7.24	My husband and I vote for different parties
WIFE/GF	144	5.24±4.42	My gf and I applied, we're new but fairly active!
GENDER	110	3.28±3.10	Am I insane? Eh. I'm just a girl who wants to look on the outside how I feel on the inside.
RACE/NATIONALITY	99	3.63±2.37	As Italian I hope tonight you will won the world cup
SEXUAL ORIENTATION	58	6.52±7.47	I'm a straight man but I do wanna say this
NAME	21	3.81±3.48	Hello guys, my name is xxx and I love travelling
CONTACT	14	5.69±3.56	xxx is my ig
Personal Experiences			
HEALTH	783	10.36±9.78	I am pretty sure I have autism, but I don't want to get an official diagnosis.
FAMILY	543	9.27±8.73	My little brother (9M) is my pride and joy
OCCUPATION	428	8.90±6.60	I'm a motorcycle tourer (by profession), but when I'm off the saddle I'm mostly bored
MENTAL HEALTH	285	16.86±16.28	I get asked this pretty regularly.. but I struggle with depression and ADHD
EDUCATION	229	9.92±7.71	Hi there, I got accepted to UCLA (IS), which I'm pumped about.
FINANCE	153	12.00±9.19	Yes. I was making $68k a year and had around $19k in debt

Table 1: Statistics and examples for each self-disclosure category in our dataset, sorted by decreasing frequency. Personal identifiable information are redacted as 'xxx' to be shown here.

employ a human-centered, iterative design process with actual end-users to evaluate and improve the model. Our self-disclosure detection model helps users scrutinize their contents to the word level (e.g., *"16F"*) to account for privacy risks.

Specifically, we introduce a comprehensive taxonomy for self-disclosure (Table 1) that consists of 13 demographic attributes and 6 personal experiences. We create a high-quality dataset with human annotations on 2.4K Reddit posts, covering 4.8K varied self-disclosures. With this corpus, we fine-tune a language model to identify the self-disclosures in the given text, achieving over 75% token-level F_1. Besides the standard NLP evaluations, more importantly, we conducted an HCI user study with 21 Reddit users to validate the real-world applicability. 82% participants have a positive outlook on the model, while also providing valuable suggestions on aspects such as personalization and explainability, which are often overlooked in benchmark assessments. Participants also express a need for a tool that can (quote) *"rewrite disclosures for me in a way that I don't worry about privacy concerns"*. We thus introduce a new text generation task, self-disclosure abstraction, with the goal of rephrasing the disclosure into less specific words without losing the utility or essence of the message. For example, providing *"I am exploring my sexual identity"* as an alternative to *"I want to be a bi(sexual) M(an)"* to let users choose based on their preferences, more in Figure 1. We investigate the generation of both single and triple abstractions, with the latter approach designed to provide users with diverse options. We experiment with different fine-tuning strategies. The best model, distilled on GPT-3.5 generated abstractions, can increase privacy moderately (scoring 3.2 out of 5 with 5 being the highest level of detail removal) while preserving high utility (scoring 4 out of 5). The model's abstractions are also very diverse, offering varied expressions (scoring 4.6 out of 5).[1] To assist users in determining which disclosures to abstract, we additionally present the task of rating the importance of self-disclosure in understanding the context. Our fine-tuned model achieves comparable performance to GPT-3.5 of 80% accuracy.

In summary, our key contributions include:

- We introduce a new corpus annotated for self-disclosure with 19 categories (§2).
- Study with real Reddit users shows that our detection model help users manage privacy risk (§3).
- Motivated by the user study, we propose a novel self-disclosure abstraction task, and show promising model results in human evaluations (§4).

2 Fine-grained Self-Disclosure Detection

To mitigate privacy leaks and alert people about their self-disclosures, it is essential to highlight specific text segments that disclose personal information, rather than simply classifying sentences

[1]These numbers are Likert-scale of the human eval.

as containing disclosures. To cover a wide spectrum of disclosures, we design a detailed taxonomy of 19 self-disclosure categories, which is more extensive than prior work (see Related Work in §6). We further construct an annotated corpus for training automatic models to detect self-disclosures at word-level, which supports our user study in §3.

2.1 Annotated Self-Disclosure Corpus

We curate a large dataset of 2.4K Reddit posts manually annotated with 4.8K self-disclosure spans.

Post Collection. We use the public Reddit data dump from December 2022, which contained 35.86M posts. We filter out 42.52% of posts that were marked as "NSFW" or "Over_18", indicating adult content, as well as those that were removed by moderators. We keep only English posts as determined with a probability above 0.7 by the fastText (Joulin et al., 2016) language identifier.[2] This results in a total of 4.01M posts, from which we randomly sample 10K and then reconstruct the full post threads with all comments and reply chains for each post via the Reddit API. We then ask two annotators to review the 10K posts on whether containing self-disclosures, culminating in a set of 2,415 posts for subsequent span annotation.

Self-Disclosure Taxonomy. Different from prior work that focuses on certain types of self-disclosure (e.g., health (Valizadeh et al., 2021b) or sexual harassment (Chowdhury et al., 2019)), we categorizing them into 19 types that are commonly shared by social media users. The taxonomy is refined iteratively through three rounds of pilot studies. Table 1 presents statistics and examples for these 19 categories that fall into two main groups: *demographic attributes* and *personal experiences*. Attributes refer to static personal characteristics that are often stated succinctly such as `name`, `age`, and `gender`. Experiences, on the other hand, relate to events that an individual engages over time, which are more complex and dynamic, such as `health` and `education`. For disclosures concerning others, such as family members, we direct annotators to label them under the general category (e.g., `family`).

Annotation Process. To ensure quality and privacy standards, we hire seven in-house annotators who were given training tutorials and 20 annotation exercise examples. We ask annotators to highlight text spans that reveal personal information

[2] https://fasttext.cc/

Category	Span F1	Partial F1	Token F1
Age	70.29	79.50	84.32
Age&Gender	66.67	70.18	78.99
Race/Nationality	64.08	69.90	74.59
Gender	63.64	69.70	78.71
Location	62.61	79.41	85.83
Appearance	51.39	72.22	80.04
Wife/GF	49.73	61.62	68.88
Finance	49.08	61.35	79.92
Occupation	46.04	65.98	82.85
Family	45.21	59.41	72.30
Health	44.72	63.04	73.84
Mental Health	44.69	61.45	68.33
Husband/BF	44.59	55.41	74.58
Education	42.42	61.36	77.08
Pet	37.50	48.86	71.24
Relationship Status	29.20	48.67	62.74
Sexual Orientation	19.67	36.07	65.90
Overall	49.03	62.60	75.30

Table 2: Performance of automatic span-level detection model on the test set for each self-disclosure category.

within each post (including comments) and categorize them into one of 19 self-disclosure types. To enhance accuracy and relevance of self-disclosures detection, we instruct annotators to select spans with contextual information, which provides more nuanced training signals for models. For example, "*I live in the US*" would be preferred over minimal span like "*US*", which is isolated from its self-referential context. The inter-annotator agreement is 0.54 based on Krippendorff's α (Krippendorff, 2018). The two annotations were adjudicated and merged into a final gold-label annotation, which we used in evaluation. We provided more details for each self-disclosure category in Appendix A.2.

2.2 Automatic Self-Disclosure Detection

With our annotated corpus, we fine-tune RoBERTa-large (Liu et al., 2019) to automatically detect disclosures as a sequential tagging task. We split the data into 70%/10%/20% for train/val/test in chronological order, where the test set has the newest data. We experiment with both multi-class and binary-class settings, the latter identifying self-disclosure spans without detailed categorization. For the `name` category, which is infrequent, we combine a sentence classifier trained on our data (determine whether a sentence contains self-disclosure) with an existing named entity recognition model (Yamada et al., 2020). Table 2 presents the test set performance by class for the multi-class model. We report three different F1 scores to capture varied degrees of partial credit: token-level F1, partial and full span-level F1. The partial span-level F1 con-

siders a predicted span as correct if it contains or is contained by a reference span, with the overlap exceeding 50% of the maximum length of either span. More details and results are in Appendix B.

3 User Study

To understand how real users think about our disclosure detection model for protecting their privacy, we recruited 21 Reddit users through Prolific to participate in an interview study – a step that differentiates our approach from prior work on disclosure identification. This interview study and its subsequent analysis were led by three authors with multi-disciplinary expertise in human-computer interaction, privacy, and NLP.

3.1 Participants and Study Design

All participants recruited were aged 18 or older, had an active Reddit account, and had made at least three posts on Reddit. After completing a screening survey, eligible participants were asked to fill out a pre-study survey in which they were given a digital copy of the consent form describing the nature of the interview and after consenting, were prompted to schedule an interview with researchers. Interviews took place remotely over Zoom, and averaged about 2 hours, during which participants were asked to share one of their Reddit posts that raised privacy concerns, and also to write a post that they were hesitant to publish for similar reasons. We then ran those posts through both our binary and multi-label models (§2.2) and provided the annotated images of users' posts that display the detected self-disclosure spans to users. We asked participants about where they agreed and disagreed with model outputs, their overall impression of the model, if and how they would like to use the model outside of the study, as well as suggestions for improvement. Our study design was approved by the university's institutional review board (IRB).

3.2 User Perceptions of Our Model

In all, we see a significant majority (82%) of participants had a positive outlook on the model. In addition, the multi-class model that highlighted disclosure categories was helpful to around 48% of participants, aiding them in recognizing and understanding potential privacy risks in their posts.

More specifically, 62% of participants expressed a desire to use it on their own posts, and another 10% felt that even though they would see no need for such a tool themselves, they would recommend others they know to use this tool and suggest that it might be *"a good idea for... kids and teens, like people who are new to the Internet."* One participant said that *"It would be interesting to run it through before I post something that I'm like nervous about and just see what it thinks and see if there are any areas where I can fix to make it less specific to me."* An additional 10% of participants mentioned that they would use it if they were more prone to making self-disclosures or if the model is further improved (more in §3.3 and §4).

3.3 User Feedback and New Opportunities

When participants discussed why they viewed the model favorably, 62% of them mentioned the focus on disclosure at the word level, 57% mentioned self-reflection, and 48% mentioned the fine-grained categories. Users also provided valuable feedback for improvement centering around accuracy, personalization, and the desire for assistance in mitigating privacy risks in general.

One interesting finding from the user study is that there is a divergence between annotators and real users in terms of what they think should be highlighted. Our initial design goal was to highlight anything that the model identified as self-disclosure to allow users to make informed decisions themselves. This approach led 4 out of the 21 participants to the believe that the model was "over-sensitive" and inaccurate because it highlighted content that participants did not believe was risky. This issue was succinctly summarized by one participant: *"sometimes it's so oversensitive that it'll highlight things again (and again), and people might not use it because they get kind of fed up and irritated."* We propose to address this problem by importance rating in §5.

Another interesting finding is that some participants suggested having the model account for their use of privacy-preserving strategies, e.g., when users intentionally author posts with false personal information, highlighting this false personal information as a disclosure risk is not useful. Future model iterations could consider incorporating features that allow users to adapt outputs to align with these strategies. One example could be proactively offering suggestions for altering text, potentially through strategic falsehoods that retain semantic utility, to facilitate employing these strategies. In fact, 24% of participants sought recommendations

on how to rephrase text spans that the model detected as containing a sensitive disclosure. One participant articulated this need by stating: *"could you rewrite this for me in a way that I don't have to worry about privacy concerns?"* This feedback led us to explore methods for generating alternative phrasings of text spans identified as being privacy-sensitive, as we will discuss next in §4.

4 Self-Disclosure Abstraction

Building upon insights from our user study, we introduce a novel task, self-disclosure abstraction, which rephrase disclosures with less specific details while preserving the content utility (see Figure 1).

4.1 Task Definition

Given a disclosure span within a sentence, the objective is to reduce sensitive and specific details while preserving the core meaning and utility. For example, in the sentence: *" I just turned 32 last month* and have been really ... ", the highlighted disclosure span can be abstracted to *"I recently entered my early 30s"*. Compared to standard sentence paraphrasing (Dou et al., 2022b), this task uniquely focuses on the span level. A key criterion is that the abstracted span must fit seamlessly into the original sentence without changing the rest of the text. We also consider generating three alternative abstractions for each disclosure, aiming to accommodate varied user preferences.

4.2 Automatic Generation of Training Data

We use the most recent 10% of posts (plus associated comments) from our corpus (§2.1) for experiment on abstraction, which are further divided into training (159 posts), dev (25 posts), and test sets (50 posts). While it is challenging for human annotators to think of diverse abstractions, we find that LLMs such as GPT-3.5 are quite adept at this task, in line with recent successes in using LLMs to generate training data in areas such as code generation (Gunasekar et al., 2023) and instruction following (Peng et al., 2023). We use chain-of-thought (Wei et al., 2022) prompting with few-shot demonstration, which asks model to first generate a rationale (i.e., thought) on why the disclosure span needs generalization, and then three diverse abstractions. The exact prompt we used is listed in Appendix E. For creating our training and dev sets, we choose GPT-3.5 (06-13 version), balancing cost and performance. We use GPT-4 (06-13 version) for creating the test set, given its more advanced capabilities.

BLEU / ROUGE-2		**Output**	
		w/o Thought	w/ Thought
Input	Standard	15.3 / 25.1	14.4 / 22.9
	Special Token	17.0 / 25.4	12.9 / 21.7
	Instruction	17.9 / 24.8	18.3 / 25.5

Table 3: Test results on self-disclosure abstraction task. Training with special token and instruction with thought lead to the best performance.

4.3 Abstraction Models

We fine-tune Llama-2-7B (Touvron et al., 2023) with LoRA (Hu et al., 2021) for generating abstractions. We consider three different **input formats**, that use standard input, special token, and natural language instructions:

STANDARD INPUT:
`Sentence:{s}\nDisclosure Span:{d}\nAbstraction Span:`

SPECIAL TOKEN:
`<SENTENCE>{s}<SPAN>{d}<ABSTRACTION>`

INSTRUCTION:
`Your task is to abstract the given disclosure... Sentence:{s}\nDisclosure Span:{d}`

For **output formats**, which is the text that model is trained to generate and where the loss is calculated, we explore two options. One is solely the desired output, which is the abstraction, and another is a rationale plus the abstraction, also known as Chain-of-Thoughts (Wei et al., 2022) training.

In addition, for generating three alternative abstractions, we can either sample three times from an abstraction model that generates only one abstraction at a time (**sampling**), or train an abstraction model that generates three abstractions all at once (input → A, B, C) at the inference time (**end-to-end training**), or use an iterative approach that breaks the three abstractions (A, B, C) into three separate training instances: input → A, input+A → B, input+A+B → C (**iterative generation**).

4.4 Automatic Evaluation.

Table 3 reports BLEU (Papineni et al., 2002) and ROUGE-2 (Lin, 2004) for comparing different input and output configurations (§4.3) for the task of generating only one abstraction per self-disclosure. We find that using special token and instruction with thoughts help improve the performance over the standard method.

Methods	Automatic Evaluation (Matching Metric)				Human Evaluation		
	BLEU	ROUGE-2	ROUGE-L	Diversity	Rank Dist. (1→9)	Rank ▽	Rating
Sampling							
Special Token	17.40	18.81	38.42	2,576		5.35	80.88
Instruction	16.95	18.78	38.35	2,564		5.32	81.03
Instruction (w/ thought)	17.06	18.47	38.43	2,882		5.23	81.17
End-to-end training							
Special Token	17.22	19.67	38.38	2,911		4.28	81.50
Instruction	17.99	19.60	**39.57**	2,992		3.62	**82.60**
Instruction (w/ thought)	16.53	19.24	38.81	2,801		4.25	81.77
Iterative generation							
Special Token	**18.13**	19.74	38.71	2,913		3.75	82.58
Instruction	17.80	**19.81**	39.56	**3,067**		4.0	82.20
Instruction (w/ thought)	16.89	18.84	38.31	2,914		**3.55**	82.43

Table 4: Test results on generating three alternative abstracted spans. End-to-end instruction tuning and iterative instruction training with thought achieve the top two performances under human evaluation. *Rank Dist.* presents the histograms of the rank distribution, where 1 is the best and 9 the worst. *Diversity* presents # unique bigrams.

In Table 4, we present the evaluation results of these best three models for generating three different abstractions for each self-disclosure span. We adopt the best matching BLEU and ROUGE metrics proposed by (Dou et al., 2021), to encourage diversity in generated abstractions. The matching metric uses the Hungarian algorithm (Kuhn, 1955) to calculate the highest matching scores among one-to-one parings between generations and references. According to the automatic metrics, all three methods show similarly high performance, while generating three abstractions in iterative steps yield slightly better results.

4.5 Human Evaluation

We further conducted a human evaluation on 60 sampled self-disclosure test instances. We first use Rank & Rate (Maddela et al., 2023) to rate the abstractions generated from each of the nine models in Table 4 on a scale of 0-100 as a list-wise ranking. We then evaluate the top two models on four aspects: privacy increase, utility preservation, and diversity, all rated on a 1-5 Likert scale, along with a binary assessment of coherence, evaluating whether each abstraction integrates seamlessly into the sentence. Detailed definitions for each aspect and Likert scale are displayed in Appendix D.2.

Human evaluation on overall ranking in Table 4 reveals that end-to-end training using the instruction scores the highest, while iterative generation with thought and instruction tops in the ranks with 3.55 out of 9 models. The additional aspect-based human evaluation conducted on these two models are shown in Figure 2. End-to-end training in the instruction format perform better or equally well across all four aspects. It achieves 4.6 out of 5 in diversity, 4.0 out of 5 in utility preservation, and 3.2 out of 5 for privacy increase, demonstrating our abstraction model's usability.

5 Importance Rating of Self-Disclosures

To highlight and abstract self-disclosure more selectively, as users suggested in §3, we consider an additional task that rates the importance of each disclosure within context.

Task Definition. Given the disclosure span and its surrounding context, the task is to estimate how important this disclosure is for others to understand the user's message and communication goals. We consider three levels: `low`, `moderate`, and `high`, corresponding to disclosures that are not essential (can be removed), essential but can be abstracted, and have to be kept as it is for understanding the message (see App. D for details). For disclosures that appear in post title or body, we consider both the title and body of the post as context. For disclosures in comments (i.e., replies to the main post), the context extends to the entire comment and its parent comment in the reply chain, if exists.

Training Data. Compared to abstraction, human perform more effectively than LLMs for important rating. We use the same train/dev/test split as the

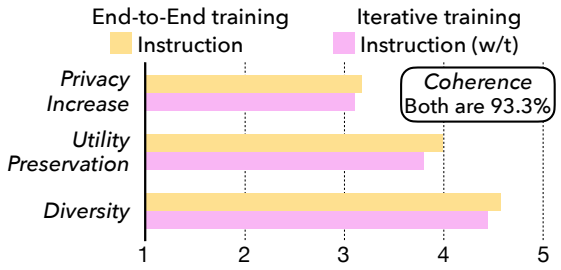

Figure 2: Human evaluation with Likert-scale (1-5) of the top two performing models for generating three alternative abstractions. The best model shows moderate privacy increase, high utility preservation, and very high diversity in abstractions. w/t denotes *with thought*.

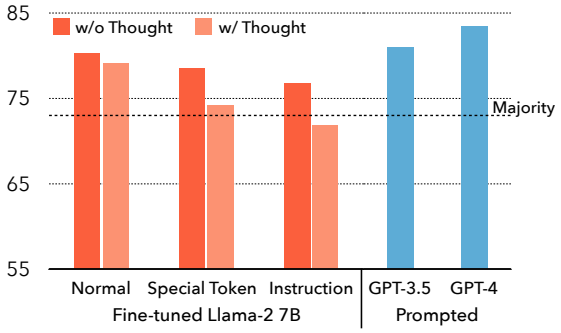

Figure 3: Test results of importance rating, measured in accuracy. Fine-tuning directly on output is better than on thought. The best fine-tuned model achieves comparable performance with GPT-3.5.

abstraction experiment and have each instance annotated by three in-house annotators. We calculate Krippendorff's α as 0.29 for inter-annotator agreement, which is considered as fair (Krippendorff, 2004). This level of agreement is anticipated given the task's subjective nature. For example, people have different opinions about whether specifics like age in a dating post should be retained (*"32"*) or abstracted (*"early 30s"*). There is often not a clear cut between low and moderate, moderate and high; yet, they can provide useful signals to users. For training labels, we take the majority vote or `moderate` if the annotators choose all three levels.

Model. We fine-tune Llama-2 7B with various input-output formats as in the abstraction experiment (§4). We also experiment with using GPT-3.5 to generate the reasoning that leads to the human-assigned importance rating.

Evaluation Results. Given the task subjectivity, we measure accuracy with considering a prediction as correct if it matches any one of the three annotations. Figure 3 shows the accuracy of each method in comparison with GPT-3.5 and GPT-4. We find that fine-tuning on thought process degrades performance across all input format performance. The top performing fine-tuned model achieves 80.37%, on par with GPT-3.5's 80.98%.

6 Related Work

There is an excellent survey-position paper by Lison et al. (2021), which has provided a comprehensive review of literature in the NLP for Privacy area. It identified one key research challenge as: "*Most importantly, they (NLP approaches) are limited to predefined categories of entities and ignore how less conspicuous text elements may also play a role in re-identifying the individual. For instance, the family status or physical appearance of a person may lead to re-identification but will rarely be considered as categories to detect.*" — which motivated this very work of ours. We discuss some related works, including the newer ones, below.

Online Self-Disclosure Detection. Most existing research addressed self-disclosure detection in social media as a sentence or document classification task, which could not accurately pinpoint the specific disclosure spans. Many prior work focused on one specific kind of self-disclosure, such as medical and mental health conditions (Coppersmith et al., 2014; De Choudhury et al., 2016; Yates et al., 2017; Benton et al., 2017; Klein et al., 2017; Zhao et al., 2019; Valizadeh et al., 2021a, 2023), sexual harassment (Schrading et al., 2015; Andalibi et al., 2016; Chowdhury et al., 2019), personal opinions or sentiments (Cho et al., 2022), and employment history (Preoţiuc-Pietro et al., 2015; Tonneau et al., 2022). Other research considered all types of personal information that a user can reveal as a single category (Mao et al., 2011; Caliskan Islam et al., 2014; Bak et al., 2014; Balani and De Choudhury, 2015; Wang et al., 2016; Yang et al., 2017; Blose et al., 2020; Reuel et al., 2022). A few works considered self-disclosures of multiple classes, but covered only a small number of categories (Lee et al., 2023; Akiti et al., 2020), or relied on dictionary/rule-based methods (Guarino et al., 2022), or limited to a particular context (e.g., bragging (Jin et al., 2022) and news comments (Umar et al., 2019)). To address these limitations, we emphasize detecting self-disclosure at span level and broadening the coverage to include 19 distinct categories. This allows for a more fine-grained detection of privacy leaks for users, and also facilitates rephrasing suggestions for these specific disclosures before posting

online.

PII Identification and Anonymization. Personal Identifiable Information (PII) is closely related to self-disclosure, but with a focus on highly sensitive attributes, such as full name, social security number, date of birth, etc. (Regulation, 2016; Morris et al., 2022; Adams et al., 2019; Lukas et al., 2023; Hathurusinghe et al., 2021). Such sensitive data is more commonly encountered in legal (Pilán et al., 2022; Mansfield et al., 2022) and medical text (Yue and Zhou, 2020; Dernoncourt et al., 2017), as opposed to text from social media or online community sources. Existing tools for PII identification such as Microsoft's Presidio [3] use rule-based methods, such as regular expressions (Subramani et al., 2023; Mouhammad et al., 2023), and Named Entity Recognition (NER) detectors (Honnibal et al., 2020). However, such approaches indiscriminately mark entities (e.g., a business phone number on an advertisement) without considering whether the information is self-disclosed personal data. PII anonymization or redaction (Azure, 2023; AWS, 2023), widely used in healthcare records management and machine learning training pipelines, replace sensitive data with masked tokens (e.g., *[xxx]*) or weaker labels (e.g., *[Location]*). This aggressive way that hurts the utility of the message is not suited for online self-disclosure, which is often voluntary and serves specific functions. To address this, we introduce a novel task of **text abstraction**, with the goal of keeping the essence of the disclosure while reducing personal details, striking a balance of utility and privacy.

Privacy Leakage in Language Models. Recent work has shown however that LLMs are prone to leaking personal information (Sun et al., 2023; Kim et al., 2023; Huang et al., 2022; Lukas et al., 2023) and lack the ability to reason about privacy compared with humans (Mireshghallah et al., 2023). This phenomenon of user privacy violation is due to the issue of memorization (Carlini et al., 2022, 2021; Staab et al., 2023), where LMs recall individual sequences from their pre-training corpora. Recent efforts in solving this leakage problem include differentially-private training (Ponomareva et al., 2023; Li et al., 2021), decoding methods that prevent generation of memorized sequences (Ippolito et al., 2022), and prompt self-moderation (Chen et al., 2023). We refer interested readers to the recent surveys of Smith et al. (2023), Ishihara (2023), and Klymenko et al. (2022) for additional information on privacy leakage and memorization. Our work takes a different angle at addressing this problem, by providing a user-centered approach that tackles the root cause of privacy leakage; helping users make more informed decisions when posting online through self-disclosure identification and abstraction. This in turn can help decrease the chance of personal information ending up in pre-training corpora and reduce potential privacy violations.

[3] https://microsoft.github.io/presidio/

7 Conclusion

We push the first steps in protecting user-side privacy in online self-disclosures. Our disclosure detection model trained on our new fine-grained corpus with span-level annotations achieves over 75% of token-F1 across both binary and multi-label settings, and is further validated through a HCI user study, highlighting its real-world applicability. Responding to the need from participants for balancing privacy risk reduction with message utility, we propose a novel task of self-disclosure abstraction, and explore a variety of fine-tuning regimes to generate one span or three diverse spans. Our human evaluation shows that the best model can provide diverse abstractions that reduce privacy risks while highly preserving utility. We further fine-tune a model to rate the importance of the self-disclosure on understanding user's perspective and context. This model reaches 80% accuracy, matching the performance of GPT-3.5, thereby helping users to decide which disclosures to abstract. Overall, we believe our work paves the way for a new direction of using LLMs to protect the user-side privacy.

Limitations

Our user study (§3.3) has revealed some additional limitations and research directions, including personalization, explainability, and contextual awareness. More specifically, 5 participants suggested that the tool should consider subreddit-specific norms. For example, the `r/diabetes` subreddit inherently expects that users may discuss their medical condition, rendering some model-predicted highlights redundant. Participants also expressed a desire for more transparency in the model's decision-making process, such as more explanations as to *why* certain highlights were marked as disclosures. Future work could expand our research to include other social platforms could provide

broader insights and applicability in diverse social environments. Future research could also investigate the model's performance after quantization, which will allow deploying Llama-7B completely on personal devices for better privacy protection.

Ethics Statement

We take the following measures to safeguard the personal information in our corpus before the annotation process. First, all personal identification information (PII), such as names and emails, is replaced with synthetic data. Second, we hired in-house student annotators ($18 per hour) instead of crowd workers for annotation. Every annotator was informed that their annotations were being used to create a dataset for online self-disclosure detection. All examples, except for those generated by the model, shown in the paper are synthetic but accurately reflect the real data. Our user study was approved by the Institutional Review Board (IRB) at our university. The primary data collected from user interviews (including the Reddit posts run through the model) was self-disclosed and gathered in a survey with the participants' consent. In accordance with IRB policy, we anonymized all data collected during the study by removing any PII. The primary purpose of our models is to provide users with a tool to mitigate the privacy risks associated with online self-disclosures. In cases where the models fail, they do not pose additional risks but rather tend towards overprotection, either by identifying more spans of text or by overly abstracting the disclosed information. We identified no potential harms that would disproportionately impact marginalized or otherwise vulnerable populations. To prevent misuse, we will not release our dataset and models to the public. Instead, we will share them with researchers who are committed to adhering to these ethical guidelines to accelerate future study.

References

Allison Adams, Eric Aili, Daniel Aioanei, Rebecca Jonsson, Lina Mickelsson, Dagmar Mikmekova, Fred Roberts, Javier Fernandez Valencia, and Roger Wechsler. 2019. AnonyMate: A toolkit for anonymizing unstructured chat data. In *Proceedings of the Workshop on NLP and Pseudonymisation*, pages 1–7.

Chandan Akiti, Anna Squicciarini, and Sarah Rajtmajer. 2020. A semantics-based approach to disclosure classification in user-generated online content. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3490–3499.

Nazanin Andalibi, Oliver L Haimson, Munmun De Choudhury, and Andrea Forte. 2016. Understanding social media disclosures of sexual abuse through the lenses of support seeking and anonymity. In *Proceedings of the 2016 CHI conference on human factors in computing systems*, pages 3906–3918.

AWS. 2023. Amazon comprehend.

Azure. 2023. Azure ai language.

JinYeong Bak, Chin-Yew Lin, and Alice Oh. 2014. Self-disclosure topic model for classifying and analyzing twitter conversations. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1986–1996.

Sairam Balani and Munmun De Choudhury. 2015. Detecting and characterizing mental health related self-disclosure in social media. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, pages 1373–1378.

Adrian Benton, Margaret Mitchell, Dirk Hovy, et al. 2017. Multitask learning for mental health conditions with limited social media data. In *15th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2017. Proceedings of Conference*. Association for Computational Linguistics.

Taylor Blose, Prasanna Umar, Anna Squicciarini, and Sarah Rajtmajer. 2020. Privacy in crisis: A study of self-disclosure during the coronavirus pandemic. *arXiv preprint arXiv:2004.09717*.

Aylin Caliskan Islam, Jonathan Walsh, and Rachel Greenstadt. 2014. Privacy detective: Detecting private information and collective privacy behavior in a large social network. In *Proceedings of the 13th Workshop on Privacy in the Electronic Society*, pages 35–46.

Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramer, and Chiyuan Zhang. 2022. Quantifying memorization across neural language models. *arXiv preprint arXiv:2202.07646*.

Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. 2021. Extracting training data from large language models. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 2633–2650.

Yang Chen, Ethan Mendes, Sauvik Das, Wei Xu, and Alan Ritter. 2023. Can language models be instructed to protect personal information? *arXiv preprint arXiv:2310.02224*.

Won Ik Cho, Soomin Kim, Eujeong Choi, and Younghoon Jeong. 2022. Assessing how users display self-disclosure and authenticity in conversation with human-like agents: A case study of luda lee. In *Findings of the Association for Computational Linguistics: AACL-IJCNLP 2022*, pages 145–152.

Arijit Ghosh Chowdhury, Ramit Sawhney, Puneet Mathur, Debanjan Mahata, and Rajiv Ratn Shah. 2019. Speak up, fight back! detection of social media disclosures of sexual harassment. In *Proceedings of the 2019 conference of the North American chapter of the Association for Computational Linguistics: Student research workshop*, pages 136–146.

Elizabeth Clark, Tal August, Sofia Serrano, Nikita Haduong, Suchin Gururangan, and Noah A. Smith. 2021. All that's 'human' is not gold: Evaluating human evaluation of generated text. 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 7282–7296, Online. Association for Computational Linguistics.

Glen Coppersmith, Mark Dredze, and Craig Harman. 2014. Quantifying mental health signals in twitter. In *Proceedings of the workshop on computational linguistics and clinical psychology: From linguistic signal to clinical reality*, pages 51–60.

Paul C. Cozby. 1973. Self-disclosure: a literature review. *Psychological bulletin*, 79 2:73–91.

Munmun De Choudhury, Emre Kiciman, Mark Dredze, Glen Coppersmith, and Mrinal Kumar. 2016. Discovering shifts to suicidal ideation from mental health content in social media. In *Proceedings of the 2016 CHI conference on human factors in computing systems*, pages 2098–2110.

Franck Dernoncourt, Ji Young Lee, Ozlem Uzuner, and Peter Szolovits. 2017. De-identification of patient notes with recurrent neural networks. *Journal of the American Medical Informatics Association*, 24(3):596–606.

Yao Dou, Maxwell Forbes, Ari Holtzman, and Yejin Choi. 2021. Multitalk: A highly-branching dialog testbed for diverse conversations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 12760–12767.

Yao Dou, Maxwell Forbes, Rik Koncel-Kedziorski, Noah A. Smith, and Yejin Choi. 2022a. Is GPT-3 text indistinguishable from human text? scarecrow: A framework for scrutinizing machine text. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7250–7274, Dublin, Ireland. Association for Computational Linguistics.

Yao Dou, Chao Jiang, and Wei Xu. 2022b. Improving large-scale paraphrase acquisition and generation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9301–9323, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. Chatgpt outperforms crowd-workers for text-annotation tasks. *arXiv preprint arXiv:2303.15056*.

Alfonso Guarino, Delfina Malandrino, and Rocco Zaccagnino. 2022. An automatic mechanism to provide privacy awareness and control over unwittingly dissemination of online private information. *Computer Networks*, 202:108614.

Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, et al. 2023. Textbooks are all you need. *arXiv preprint arXiv:2306.11644*.

Rajitha Hathurusinghe, Isar Nejadgholi, and Miodrag Bolic. 2021. A privacy-preserving approach to extraction of personal information through automatic annotation and federated learning. In *Proceedings of the Third Workshop on Privacy in Natural Language Processing*, pages 36–45.

David Heineman, Yao Dou, Mounica Maddela, and Wei Xu. 2023. Dancing between success and failure: Edit-level simplification evaluation using SALSA. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Singapore. Association for Computational Linguistics.

Matthew Honnibal, Ines Montani, Sofie Van Landeghem, Adriane Boyd, et al. 2020. spaCy: Industrial-strength natural language processing in python.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.

Jie Huang, Hanyin Shao, and Kevin Chen-Chuan Chang. 2022. Are large pre-trained language models leaking your personal information? In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2038–2047.

Daphne Ippolito, Florian Tramèr, Milad Nasr, Chiyuan Zhang, Matthew Jagielski, Katherine Lee, Christopher A Choquette-Choo, and Nicholas Carlini. 2022. Preventing verbatim memorization in language models gives a false sense of privacy. *arXiv preprint arXiv:2210.17546*.

Shotaro Ishihara. 2023. Training data extraction from pre-trained language models: A survey. *arXiv preprint arXiv:2305.16157*.

Mali Jin, Daniel Preoţiuc-Pietro, A Seza Doğruöz, and Nikolaos Aletras. 2022. Automatic identification and classification of bragging in social media. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3945–3959.

Armand Joulin, Edouard Grave, Piotr Bojanowski, Matthijs Douze, Hérve Jégou, and Tomas Mikolov. 2016. Fasttext.zip: Compressing text classification models. *arXiv preprint arXiv:1612.03651*.

Sidney M Jourard. 1971. Self-disclosure: An experimental analysis of the transparent self.

Siwon Kim, Sangdoo Yun, Hwaran Lee, Martin Gubri, Sungroh Yoon, and Seong Joon Oh. 2023. ProPILE: Probing privacy leakage in large language models. *Advances in Neural Information Processing Systems 36 (NeurIPS 2023)*.

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.

Oleksandra Klymenko, Stephen Meisenbacher, and Florian Matthes. 2022. Differential privacy in natural language processing the story so far. In *Proceedings of the Fourth Workshop on Privacy in Natural Language Processing*, pages 1–11.

Klaus Krippendorff. 2004. Reliability in content analysis: Some common misconceptions and recommendations. *Human communication research*, 30(3):411–433.

Klaus Krippendorff. 2018. *Content analysis: An introduction to its methodology*. Sage Publications.

Harold W Kuhn. 1955. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97.

Jooyoung Lee, Sarah Rajtmajer, Eesha Srivatsavaya, and Shomir Wilson. 2023. Online self-disclosure, social support, and user engagement during the covid-19 pandemic. *ACM Transactions on Social Computing*.

Ada Lerner, Helen Yuxun He, Anna Kawakami, Silvia Catherine Zeamer, and Roberto Hoyle. 2020. Privacy and activism in the transgender community. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–13.

Xuechen Li, Florian Tramer, Percy Liang, and Tatsunori Hashimoto. 2021. Large language models can be strong differentially private learners. In *International Conference on Learning Representations*.

Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.

Pierre Lison, Ildikó Pilán, David Sanchez, Montserrat Batet, and Lilja Øvrelid. 2021. Anonymisation models for text data: State of the art, challenges and future directions. 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 4188–4203, Online. Association for Computational Linguistics.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.

Nils Lukas, Ahmed Salem, Robert Sim, Shruti Tople, Lukas Wutschitz, and Santiago Zanella-Béguelin. 2023. Analyzing leakage of personally identifiable information in language models. *arXiv preprint arXiv:2302.00539*.

Mufan Luo and Jeffrey T Hancock. 2020. Self-disclosure and social media: motivations, mechanisms and psychological well-being. *Current opinion in psychology*, 31:110–115.

Mounica Maddela, Yao Dou, David Heineman, and Wei Xu. 2023. LENS: A learnable evaluation metric for text simplification. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16383–16408, Toronto, Canada. Association for Computational Linguistics.

Courtney Mansfield, Amandalynne Paullada, and Kristen Howell. 2022. Behind the mask: Demographic bias in name detection for pii masking. In *Proceedings of the Second Workshop on Language Technology for Equality, Diversity and Inclusion*, pages 76–89.

Huina Mao, Xin Shuai, and Apu Kapadia. 2011. Loose tweets: an analysis of privacy leaks on twitter. In *Proceedings of the 10th annual ACM workshop on Privacy in the electronic society*, pages 1–12.

Niloofar Mireshghallah, Hyunwoo Kim, Xuhui Zhou, Yulia Tsvetkov, Maarten Sap, Reza Shokri, and Yejin Choi. 2023. Can llms keep a secret? testing privacy implications of language models via contextual integrity theory. *arXiv preprint arXiv:2310.17884*.

John Morris, Justin Chiu, Ramin Zabih, and Alexander M Rush. 2022. Unsupervised text deidentification. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4777–4788.

Nina Mouhammad, Johannes Daxenberger, Benjamin Schiller, and Ivan Habernal. 2023. Crowdsourcing on sensitive data with privacy-preserving text rewriting. In *Proceedings of the 17th Linguistic Annotation Workshop (LAW-XVII)*, pages 73–84, Toronto, Canada. Association for Computational Linguistics.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the*

40th annual meeting of the Association for Computational Linguistics, pages 311–318.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.

Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023. Instruction tuning with gpt-4. *arXiv preprint arXiv:2304.03277*.

Ildikó Pilán, Pierre Lison, Lilja Øvrelid, Anthi Papadopoulou, David Sánchez, and Montserrat Batet. 2022. The text anonymization benchmark (tab): A dedicated corpus and evaluation framework for text anonymization. *Computational Linguistics*, 48(4):1053–1101.

Natalia Ponomareva, Hussein Hazimeh, Alex Kurakin, Zheng Xu, Carson Denison, H Brendan McMahan, Sergei Vassilvitskii, Steve Chien, and Abhradeep Guha Thakurta. 2023. How to dp-fy ml: A practical guide to machine learning with differential privacy. *Journal of Artificial Intelligence Research*, 77:1113–1201.

Matt Post. 2018. A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.

Daniel Preoţiuc-Pietro, Vasileios Lampos, and Nikolaos Aletras. 2015. An analysis of the user occupational class through twitter content. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1754–1764.

Protection Regulation. 2016. Regulation (eu) 2016/679 of the european parliament and of the council. *Regulation (eu)*, 679:2016.

Ricardo Rei, Marcos Treviso, Nuno M. Guerreiro, Chrysoula Zerva, Ana C Farinha, Christine Maroti, José G. C. de Souza, Taisiya Glushkova, Duarte Alves, Luisa Coheur, Alon Lavie, and André F. T. Martins. 2022. CometKiwi: IST-unbabel 2022 submission for the quality estimation shared task. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 634–645, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

Ann-Katrin Reuel, Sebastian Peralta, João Sedoc, Garrick Sherman, and Lyle Ungar. 2022. Measuring the language of self-disclosure across corpora. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1035–1047.

Nicolas Schrading, Cecilia Ovesdotter Alm, Raymond Ptucha, and Christopher Homan. 2015. An analysis of domestic abuse discourse on reddit. In *Proceedings of the 2015 conference on empirical methods in natural language processing*, pages 2577–2583.

Manya Sleeper. 2016. *Everyday online sharing*. Ph.D. thesis, Carnegie Mellon University.

Victoria Smith, Ali Shahin Shamsabadi, Carolyn Ashurst, and Adrian Weller. 2023. Identifying and mitigating privacy risks stemming from language models: A survey. *arXiv preprint arXiv:2310.01424*.

Robin Staab, Mark Vero, Mislav Balunović, and Martin Vechev. 2023. Beyond memorization: Violating privacy via inference with large language models. *arXiv preprint arXiv:2310.07298*.

Nishant Subramani, Sasha Luccioni, Jesse Dodge, and Margaret Mitchell. 2023. Detecting personal information in training corpora: an analysis. In *Proceedings of the 3rd Workshop on Trustworthy Natural Language Processing (TrustNLP 2023)*, pages 208–220.

Albert Yu Sun, Eliott Zemour, Arushi Saxena, Udith Vaidyanathan, Eric Lin, Christian Lau, and Vaikkunth Mugunthan. 2023. Does fine-tuning gpt-3 with the openai api leak personally-identifiable information? *arXiv preprint arXiv:2307.16382*.

Manuel Tonneau, Dhaval Adjodah, Joao Palotti, Nir Grinberg, and Samuel Fraiberger. 2022. Multilingual detection of personal employment status on twitter. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6564–6587.

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

Prasanna Umar, Anna Squicciarini, and Sarah Rajtmajer. 2019. Detection and analysis of self-disclosure in online news commentaries. In *The World Wide Web Conference*, pages 3272–3278.

Mina Valizadeh, Xing Qian, Pardis Ranjbar-Noiey, Cornelia Caragea, and Natalie Parde. 2023. What clued the ai doctor in? on the influence of data source and quality for transformer-based medical self-disclosure detection. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1193–1208.

Mina Valizadeh, Pardis Ranjbar-Noiey, Cornelia Caragea, and Natalie Parde. 2021a. Identifying medical self-disclosure in online communities. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4398–4408.

Mina Valizadeh, Pardis Ranjbar-Noiey, Cornelia Caragea, and Natalie Parde. 2021b. Identifying medical self-disclosure in online communities. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4398–4408, Online. Association for Computational Linguistics.

Yi-Chia Wang, Moira Burke, and Robert Kraut. 2016. Modeling self-disclosure in social networking sites. In *Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing*, pages 74–85.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.

Rachel Wicks and Matt Post. 2021. A unified approach to sentence segmentation of punctuated text in many languages. 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 3995–4007, Online. Association for Computational Linguistics.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface's transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.

Ikuya Yamada, Akari Asai, Hiroyuki Shindo, Hideaki Takeda, and Yuji Matsumoto. 2020. LUKE: Deep contextualized entity representations with entity-aware self-attention. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6442–6454, Online. Association for Computational Linguistics.

Diyi Yang, Zheng Yao, and Robert Kraut. 2017. Self-disclosure and channel difference in online health support groups. In *Proceedings of the international AAAI conference on web and social media*, volume 11, pages 704–707.

Andrew Yates, Arman Cohan, and Nazli Goharian. 2017. Depression and self-harm risk assessment in online forums. In *Conference on Empirical Methods in Natural Language Processing*, pages 2958–2968. ACL.

Xiang Yue and Shuang Zhou. 2020. PHICON: Improving generalization of clinical text de-identification models via data augmentation. In *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, pages 209–214.

Xinyan Zhao, Deahan Yu, and VG Vinod Vydiswaran. 2019. Identifying adverse drug events mentions in tweets using attentive, collocated, and aggregated medical representation. In *Proceedings of the Fourth Social Media Mining for Health Applications (# SMM4H) Workshop & Shared Task*, pages 62–70.

A Self-Disclosure Corpus

A.1 Quality Control

Recent works have shown that crowdsourcing often leads to lower quality annotations (Clark et al., 2021; Gilardi et al., 2023). To ensure the quality, we hire seven in-house annotators, who are undergraduate students at a US university and native English speakers, and compensating them at a rate of $18 per hour. Each annotator first takes a training that includes tutorials and 20 exercise examples. During the annotation, we release the data in ten batches, allowing us to constantly monitor and provide feedback as needed. We use the BRAT interface for the annotation process.[4]

Category	Krippendorff's α	Two Agree (%)
AGE	0.80	73.1
AGE&GENDER	0.88	81.6
RACE/NATIONALITY	0.88	82.2
GENDER	0.79	74.0
LOCATION	0.71	65.4
APPEARANCE	0.68	57.1
WIFE/GF	0.70	62.1
FINANCE	0.65	57.6
OCCUPATION	0.66	58.4
FAMILY	0.70	61.5
HEALTH	0.53	46.7
MENTAL HEALTH	0.45	37.0
HUSBAND/BF	0.78	70.8
EDUCATION	0.72	64.9
PET	0.48	37.7
RELATIONSHIP STATUS	0.48	37.3
SEXUAL ORIENTATION	0.58	49.9
OVERALL	0.54	45.7

Table 5: Inter-annotator agreement for each self-disclosure category.

A.2 Inter-annotator Agreement

Table 5 shows the inter-annotator agreement per self-disclosure category, measured in Krippendorff's α (Krippendorff, 2018) and *Two Agree*. Due to computational constraints, we calculate Krippendorff's α per post and report the average across the dataset. Two agree is the percentage of words labeled as *disclosure* by both annotators $a1$ and $a2$: $\frac{|\text{Words}_{a1} \cap \text{Words}_{a2}|}{|\text{Words}_{a1} \cup \text{Words}_{a2}|}$. The numbers are higher than or similar to those in other span-level annotation work (Dou et al., 2022a; Heineman et al., 2023).

B Further Results

In this section, we describe the details of our self-disclosure detection model and show the performance of the binary-class model.

[4] https://brat.nlplab.org/

Input	Span F1		Partial F1		Token F1	
	Multi	*Binary*	*Multi*	*Binary*	*Multi*	*Binary*
Normal	44.52	48.14	57.24	64.20	69.85	72.20
256	45.29	47.60	58.43	64.06	68.29	72.16
128	44.61	47.73	57.97	62.46	68.92	71.26
64	44.93	49.54	59.00	65.82	70.81	74.87
Sentence	**49.03**	**54.47**	**62.60**	**72.67**	**75.30**	**81.78**

Table 6: Test performance of models fine-tuned on various data setups, with training on sentence level achieving the best results. For *Multi*-class models, the results are averaged over all classes, excluding label "O", while the *Binary* models report F1 score for "disclosure".

B.1 Experiment Details

We fine-tune RoBERTa-large (Liu et al., 2019), a transformer-based encoder with 355M parameters, on our dataset by minimizing the cross-entropy loss for each token's label. As some words are tokenized into multiple subword tokens, during inference, we use the hidden states of the first token to get the label (Rei et al., 2022).

We experiment with various data processing methods during fine-tuning, given that Reddit posts and comments can significantly vary in length and disclosure spans do not always require extended contexts. Specifically, we segment comments or posts into shorter chunks of 64, 128, and 256 words, as well as individual sentences using Ersatz (Wicks and Post, 2021).

For self-disclosure classes (`name` and `contact`) which are infrequent in Reddit data and thus insufficient to train models from scratch, we turn to existing models and tools. We use the state-of-the-art NER model, LUKE (Yamada et al., 2020) to identify person names and Microsoft Presidio[5] to recognize contact information such as phone numbers and social media usernames. To specifically identify self-disclosures as opposed to generic names (e.g., Taylor Swift), we further train a sentence classifier using RoBERTa-large, which achieves 84.4 test F_1, to first determine whether a sentence contains self-disclosure.

B.2 Results

Table 6 presents the average test-set performance for both binary and multi-class models fine-tuned under different data configurations. Due to the inherent simplicity, binary models typically outperform their multi-class counterparts. In addition, we

[5] https://microsoft.github.io/presidio/

find that dividing the long Reddit posts and comments into shorter pieces generally improves the performance. The most significant gain is achieved by segmenting the data at the sentence level,leading to an increase of over 4.5 span-level F1 points in both binary and multi-class settings, compared to the normal baseline. The per-category performance of the top multi-class model, trained on sentence-level data, is detailed in Table 2 of the main paper. We use the best binary and multi-class models in our user study (Section 3).

C Implementation Details

C.1 GPT-3.5 and GPT-4

We use gpt-3.5-turbo-0613 as GPT-3.5 and gpt-4-0613 as GPT-4.

C.2 Evaluation Metrics

For BLEU (Papineni et al., 2002), we use SacreBLEU (Post, 2018). For ROUGE (Lin, 2004), we use the one from torchmetrics.[6]

C.3 Experiments

We implemented our models using Huggingface (Wolf et al., 2019) 4.33.2 and PyTorch (Paszke et al., 2019) 2.0.1. All results are from single runs.

Self-disclosure Detection. We train RoBERTa-large (Liu et al., 2019) on 2 A40 GPUs for 10 epochs with total batch size of 32, which takes around 18 minutes per run. We perform a learning rate sweep over 1e-5, 2e-5, 5e-5, 8e-5 on our evaluation set, with 5e-5 being the best. We use AdamW (Loshchilov and Hutter, 2017) as the optimizer with weight decay of 0.01. Additionally, we use cosine learning rate schedule with a warmup ratio of 0.06.

Self-disclosure Abstraction. We train Llama-2 7B (Touvron et al., 2023) with LoRA (Hu et al., 2021) on 8 A40 GPUs for 5 epochs, with total batch size of 32, which takes at most 2 hours per run. We use learning rate of 1e-4, which is the standard when fine-tuning LLMs with LoRA. We use AdamW (Loshchilov and Hutter, 2017) as the optimizer with weight decay of 0.01. Additionally, we use cosine learning rate schedule with a warmup ratio of 0.03. For LoRA hyperparameters, we set rank=8, alpha=16, dropout=0.05, target modules=Q and V attention matrices. Additionally, we also update the embedding layer during fine-tuning.

Importance Rating. We use the same setup and hyperparameters as the abstraction experiment. It takes around 25 minutes per training run.

D Annotation Guidelines

D.1 Self-disclosure Annotation

Annotators were instructed to annotate explicit self-disclosures that concern the user based on our defined list of categories (Table 1). Most of the categories under "attributes" are straightforward to annotate such as the user's age, location, gender, etc. We considered instances where Reddit users revealing both their age and gender in one word, such as *"M24"*, under a specific category AGE/-GENDER that is different from AGE and GENDER which are for disclosures of age and gender individually. For tricky categories, most of which are under "experiences", we provided exact definitions to follow which help the annotators make decisions and ensure consistent labeling. Those definitions are as follows:

- *Appearance* self-disclosures are defined as descriptions of bodily features of the user, such as their height, weight, eye or hair color, or any other specific features.
- *Health* self-disclosures are defined as the disclosure of a specific disease or illness the user has, specific medications they take, or medical tests they perform.
- *Mental Health* self-disclosures are defined as situations where users discuss their feelings, state of mind, or suicidal thoughts.
- *Finance* self-disclosures are defined as mentions of specific personal financial details such as details about one's salary, recent transactions, affordability of items, choice of bank, and similar specifics.
- *Education* self-disclosures are defined as mentions what the user is currently or planning on studying, or degrees they hold.
- *Occupation* self-disclosures are defined as mentions of the current or past occupations of the user.
- *Family* self-disclosures are defined as any disclosure that fits within our specified attributes

[6] https://torchmetrics.readthedocs.io/en/stable/text/rouge_score.html

and experiences but concern a family member of the user, such as their parents, siblings, or extended family members.

D.2 Human Evaluation for Abstraction

Here are the 1-5 Likert scales on aspects: privacy reduction, utility preservation, and diversity, used in human evaluation for three-span abstraction (§4).

Privacy Increase:

1 – No Privacy Increase: The abstractions are the same or paraphrases to the disclosure span.

2 – Low Privacy Increase: The abstractions slightly obscure sensitive details but are still quite similar to the original.

3 – Moderate Privacy Increase: The abstractions moderately obscure sensitive details.

4 – High Privacy Increase: The abstractions significantly obscure sensitive details and remove details.

5 – Maximum Privacy Increase: The abstractions eliminate nearly all specific details.

Utility preservation:

1 – No Utility Preserved: The abstractions remove or significantly change the disclosure span, losing all the utility.

2 – Low Utility Preserved: The abstractions preserve a small amount of the disclosure span, but major aspects are lost or altered.

3 – Moderate Utility Preserved: The abstractions maintain a part of the disclosure span's utility.

4 – High Utility Preserved: The abstractions maintain most of the disclosure span's utility, with only minor aspects lost.

5 – Full Utility Preserved: The abstractions maintain the complete utility of the disclosure span, effectively conveying the intended function.

Diversity:

1 - Identical Abstractions: All three abstractions are essentially the same, exhibiting no diversity in wording or style.

2 - Minimal Diversity: Two of the three abstractions are identical, with only one offering a different expression.

3 - Low Diversity: All three abstractions are different, yet they exhibit similar styles and only minor variations in wording.

4 - Moderate Diversity: Each abstraction differs significantly in wording, with about half of the words unique to each. The styles are somewhat varied but maintain a degree of similarity.

5 - High Diversity: Each abstraction is distinctly unique, both in wording and in expression style, demonstrating a broad range of diversity.

D.3 Importance Rating Annotation

Low – Can be removed without compromising the understanding of the user's perspective and context.
Moderate – Adds a meaningful layer to the context but is not essential to the user's overall message. Can be abstracted.
High – Essential for an accurate understanding of the user's perspective and context, should be kept as is.

E Prompt Templates

We first list the prompts used to prompt GPT-3.5 and GPT-4, and then the prompts of format for fine-tuning the Llama-2 7B model. For prompts used in the importance rating experiment §5, we provide ones for disclosures in comments; for disclosures in other locations (title, post, subcomments), only the instance part of the prompt is modified. Please note that within these prompts, we use "generalization", "generalize", and "rationale" as equivalents to "abstraction", "abstract", and "thought" respectively.

E.1 GPT-3.5 and GPT-4 Prompts

Text Abstraction. We use the following prompt to ask GPT-3.5 (train/dev sets) or GPT-4 (test) to generate three abstraction spans with thought process for each data instance.

```
Your task is to generalize the given '
    disclosure span' in the sentence.
    Provide three diverse generalized
    spans that convey similar meaning
    but remove any overly specific or
    sensitive information.

Example 1:
Sentence: "Should I submit a 1470 SAT
    score to Carnegie Mellon and Duke?"
Disclosure Span to Revise: "1470 SAT
    score"
Rationale: The original disclosure span
    explicitly states a specific SAT
    score, which may be sensitive
    information for some individuals. To
     retain the context that the score
    is competitively high while removing
     the exact number, we can use
    generalized terms.
```

```
Generalized Spans: {{"span 1": "a high
    1400-range SAT score", "span 2": "an
     SAT score in the upper 1400s", "
    span 3": "an SAT score above 1450"}}

Example 2:
Sentence: "My Top Spotify Artists of
    2022 were Charli XCX, Taylor Swift,
    Carly Rae Jepsen, Snow Strippers,
    and Danny L Harle."
Disclosure Span to Revise: "Top Spotify
    Artists of 2022 were Charli XCX,
    Taylor Swift, Carly Rae Jepsen, Snow
     Strippers, and Danny L Harle."
Rationale: The original span lists very
    specific artists that the person
    listened to in 2022, which can be
    considered personal and revealing.
    To maintain the general idea that
    the person has favorite artists from
     that year without giving away the
    exact names, generalized terms can
    be used.
Generalized Spans: {{"span 1": "favorite
     artists of 2022 included a mix of
    pop and electronic musicians", "span
     2": "top artists from 2022 covered
    a range of styles", "span 3": "
    favorite artists last year were from
     various genres"}}

Example 3:
Sentence: "I just turned 32 last month
    and have been really reflecting on
    my relationships with people from my
     relationship with my siblings to
    coworkers."
Disclosure Span to Revise: "just turned
    32 last month"
Rationale: The original span provides
    very specific details about the
    person's age and the timing of their
     birthday, which could be seen as
    personal information. To keep the
    essence of the span-that the
    individual is in their early 30s and
     recently had a birthday-while
    making it less specific, generalized
     terms can be used.
Generalized Spans: {{"span 1": "recently
     entered my early 30s", "span 2": "
    turned into my early thirties not
    long ago", "span 3": "just started
    my third decade a little while back
    "}}

First, provide a rationale explaining
    why the disclosure span needs
    generalization. Then, offer three
    diverse generalized alternatives in
    a JSON format like this: {{"span 1":
     "xxx", "span 2": "xxx", "span 3": "
    xxx"}}.

Criteria:
* Only the disclosure span should be
    generalized; the rest of the
    sentence should remain intact.
* Generalized spans should be diverse
    but should all retain the essence of
     the original span.
* Make sure the generalized span fits
    seamlessly into the original
    sentence, maintaining proper syntax
    and grammar.

Sentence: "{sentence}"
Disclosure Span to Revise: "{span}"
Rationale:
```

Importance Rating.

```
Rate the importance of the disclosure
    span in a Reddit comment within the
    context into three Likert-scale:
* Low Importance: Can be removed without
     compromising the understanding of
    the commenter's perspective and
    context.
* Moderate Importance: Adds a meaningful
     layer to the context but is not
    essential to the commenter's overall
     message. Can be generalized.
* High Importance: Essential for an
    accurate understanding of the
    commenter's perspective and context,
     should be kept as is.

Given instance:
* Title: {title}
* Post: {post}
* Comment: {comment}
* Disclosure Span in the Comment: {
    disclosure}
{post_empty_explaination}
Note: The disclosure span is marked
    between special tokens <disclosure
    ></disclosure> in the Comment.

Read the Title, Post, and Comment
    carefully to understand the context.
     Write down your thought process.
    And in the end, provide your
    importance rating of the disclosure
    span in a JSON format: {{"Importance
    ": "Low/Moderate/High"}}.
```

The following prompt is used to prompt GPT-3.5 to generate thought process for the human-annotated importance rating:

```
Provide the thought process of the
    importance rating of a disclosure
    span in a Reddit comment.

Here are the definitions of the three
    importance rating scales:
* Low Importance: Can be removed without
     compromising the understanding of
    the commenter's perspective and
    context.
* Moderate Importance: Adds a meaningful
     layer to the context but is not
    essential to the commenter's overall
     message. Can be generalized.
* High Importance: Essential for an
    accurate understanding of the
    commenter's perspective and context,
     should be kept as is.
```

```
Given instance:
* Title: {title}
* Post: {post}
* Comment: {comment}
* Disclosure Span in the Comment: {
    disclosure}
* Human Rating of Importance: {
    human_rating}
{post_empty_explaination}
Note: The disclosure span is marked
    between special tokens <disclosure
    ></disclosure> in the Comment.

Instructions:
1. Carefully read the Title, Post, and
    Comment to understand the context.
2. Based on the human rating, write a
    detailed thought process that leads
    to this rating, according to the
    definitions. Your thought process
    should build up the reasoning that
    culminates in the rating, rather
    than stating the rating and then
    explaining it.
3. Ensure that your thought process is
    clear and straight to the point,
    avoiding filler sentences or
    unnecessary elaboration.
4. Present your thought in a JSON format
     of {{"Thought": "xxx"}}.
```

E.2 One-Span Abstraction

Instruction:

```
Your task is to generalize the given '
    disclosure span' in the sentence,
    which is providing a generalized
    alternative that is less specific
    but retains the core meaning of the
    original span.

Remember the following criteria:
* Only the disclosure span should be
    generalized; the rest of the
    sentence should remain intact.
* Make sure the generalized span fits
    seamlessly into the original
    sentence, maintaining proper syntax
    and grammar.

Sentence: "{sentence}"
Disclosure Span to Revise: "{span}"
Generalized Span:
```

Instruction with thought:

```
Your task is to generalize the given '
    disclosure span' in the sentence.

Please follow these steps:
1. First, provide a rationale explaining
     why the disclosure span needs
    generalization.
2. Then, provide a generalized
    alternative that is less specific
    but retains the core meaning of the
    original span.

Remember the following criteria:
* Only the disclosure span should be
    generalized; the rest of the
    sentence should remain intact.
* Make sure the generalized span fits
    seamlessly into the original
    sentence, maintaining proper syntax
    and grammar.

Sentence: "{sentence}"
Disclosure Span to Revise: "{span}"
Rationale:
```

E.3 Three-Span Abstraction

End-to-end instruction:

```
Your task is to generalize the given '
    disclosure span' in the sentence.
    Provide three diverse generalized
    spans that convey similar meaning
    but remove any overly specific or
    sensitive information.

Remember the following criteria:
* Only the disclosure span should be
    generalized; the rest of the
    sentence should remain intact.
* Generalized spans should be diverse
    but should all retain the essence of
     the original span.
* Make sure the generalized span fits
    seamlessly into the original
    sentence, maintaining proper syntax
    and grammar.
* Provide three diverse generalized
    alternatives in a JSON format like
    this: {{"span 1": "xxx", "span 2": "
    xxx", "span 3": "xxx"}}.

Sentence: "{sentence}"
Disclosure Span to Revise: "{span}"
Generalized Spans:
```

End-to-end instruction with thought:

```
Your task is to generalize the given '
    disclosure span' in the sentence.
    Provide three diverse generalized
    spans that convey similar meaning
    but remove any overly specific or
    sensitive information.

Please follow these steps:
1. First, provide a rationale explaining
     why the disclosure span needs
    generalization.
2. Then, offer three diverse generalized
     alternatives in a JSON format like
    this: {{"span 1": "xxx", "span 2": "
    xxx", "span 3": "xxx"}}.

Remember the following criteria:
* Only the disclosure span should be
    generalized; the rest of the
    sentence should remain intact.
* Generalized spans should be diverse
    but should all retain the essence of
     the original span.
* Make sure the generalized span fits
    seamlessly into the original
```

```
    sentence, maintaining proper syntax
    and grammar.

Sentence: "{sentence}"
Disclosure Span to Revise: "{span}"
Rationale:
```

Iterative instruction:

```
Your task is to generalize the given '
    disclosure span' in the sentence,
    which is providing a generalized
    alternative that is less specific
    but retains the core meaning of the
    original span.

Remember the following criteria:
* Only the disclosure span should be
    generalized; the rest of the
    sentence should remain intact.
* The generalized span should be
    different from the example
    generalizations but should retain
    the essence of the original span.
* Make sure the generalized span fits
    seamlessly into the original
    sentence, maintaining proper syntax
    and grammar.

Sentence: "{sentence}"
Disclosure Span to Revise: "{span}"
Example Generalizations: {examples}
Generalized Span:
```

Iterative instruction with thought:

```
Your task is to generalize the given '
    disclosure span' in the sentence.

Please follow these steps:
1. First, provide a rationale explaining
     why the disclosure span needs
    generalization.
2. Then, offer one diverse generalized
    alternatives that is different from
    the example generalizations provided
    .

Remember the following criteria:
* Only the disclosure span should be
    generalized; the rest of the
    sentence should remain intact.
* The generalized span should be
    different from the examples but
    should retain the essence of the
    original span.
* Make sure the generalized span fits
    seamlessly into the original
    sentence, maintaining proper syntax
    and grammar.

Sentence: "{sentence}"
Disclosure Span to Revise: "{span}"
Example Generalizations: {examples}
Rationale:
```

E.4 Importance Rating

Instruction:

```
Rate the importance of the disclosure
    span in a Reddit comment within the
    context into three Likert-scale:
* Low Importance: Can be removed without
     compromising the understanding of
    the commenter's perspective and
    context.
* Moderate Importance: Adds a meaningful
     layer to the context but is not
    essential to the commenter's overall
     message. Can be generalized.
* High Importance: Essential for an
    accurate understanding of the
    commenter's perspective and context,
     should be kept as is.

Given instance:
* Title: {title}
* Post: {post}
* Comment: {comment}
* Disclosure Span in the Comment: {
    disclosure}
{post_empty_explaination}
Note: The disclosure span is marked
    between special tokens <disclosure
    ></disclosure> in the Comment.

Read the Title, Post, and Comment
    carefully to understand the context,
     and provide your importance rating
    of the disclosure span.
```

Instruction with thought:

```
Rate the importance of the disclosure
    span in a Reddit comment within the
    context into three Likert-scale:
* Low Importance: Can be removed without
     compromising the understanding of
    the commenter's perspective and
    context.
* Moderate Importance: Adds a meaningful
     layer to the context but is not
    essential to the commenter's overall
     message. Can be generalized.
* High Importance: Essential for an
    accurate understanding of the
    commenter's perspective and context,
     should be kept as is.

Given instance:
* Title: {title}
* Post: {post}
* Comment: {comment}
* Disclosure Span in the Comment: {
    disclosure}
{post_empty_explaination}
Note: The disclosure span is marked
    between special tokens <disclosure
    ></disclosure> in the Comment.

Read the Title, Post, and Comment
    carefully to understand the context.
     Write down your thought process.
    And in the end, provide your
    importance rating of the disclosure
    span in a JSON format: {{"Importance
    ": "Low/Moderate/High"}}.
```