

spinf/moderation: calibrated content moderation from plain-language questions

spinf Inc., 2026-10-01

spinf/moderation 0.1.1 on spinf-12b (Gemma 4 12B) · Reproduce: [spinf-benchmarks/moderation](#)

Abstract

We describe `spinf/moderation`, a prompt pack for the `spinf` scoring API that turns a general open-weights language model (Gemma 4 12B) into a content moderation classifier without fine-tuning the model. The pack asks the model 106 short questions about a text and reads the probability of each listed answer. From those probabilities it computes calibrated 0–1 scores for 17 harm and content categories, 4 regulated topics and an overall harmfulness score, and makes a safe / unsafe decision in one of three modes: a single threshold, a threshold per category, or a 68-parameter linear layer. Everything (questions, calibration constants, thresholds, the layer’s weights) is one self-contained JSON file that follows the API’s request format, so the pack can be inspected, edited and extended with the customer’s own questions in the same call. Our goal was not only the best single safe / unsafe score: it was accuracy on par with dedicated moderation models together with fine-grained, continuous per-category scores and thresholds that each customer can adapt to their own audience and policy. On the nine two-class datasets of a recent public moderation benchmark [1], with no labels from those datasets, the pack reaches a mean macro F1 of 74.4, on par with Llama Guard 3 8B and Gemma 3 12B as reported there, at about 1,700 billed tokens per text.

Goals

Content moderation has no single right answer. A children’s education app, a gaming community, a news site’s comments, a marketplace and an AI assistant draw the line in different places, and often per topic: profanity may be fine where harassment is not, drug discussion may be welcome in a harm-reduction forum and blocked in a teen app. A classifier that returns one fixed safe / unsafe verdict, tuned to one benchmark’s definition of harm, fits none of these exactly.

We therefore set three goals for `spinf/moderation`:

1. **Accuracy on par with dedicated moderation models** on public benchmarks, measured reproducibly.
2. **Fine-grained, continuous output:** a calibrated 0–1 score per category, comparable across categories, with a threshold per category that the customer can raise, lower or switch off (and reference points that tell how much everyday content a threshold flags), instead of a single verdict.
3. **Easy customisation:** the pack is a readable JSON file in the API’s own format. Customers can edit or add questions, add their own categories, and score their own criteria (quality, relevance, brand mentions, ...) in the same call,

paying for the text once.

The safe / unsafe aggregate in Section 3.2 is one view of the results; the per-category results (Section 3.1) are the ones a customer configures against.

1. Approach

Questions as reads. Each question is a short statement completed by the model, scored over a fixed list of answers, e.g. “*The user wants to know how to ...*” with answers *hack an account*, *steal passwords*, *cook*, *code*, ..., or “*Most people would find the message offensive / reasonable*”. The API returns the probability of every listed answer; a read is a log-probability contrast. One question with many answers yields many reads at the cost of one prompt (“last word” sets): each answer feeds the category it names, and harmless answers act as anchors.

Per-option calibration. A model prefers some answers regardless of the text. Every read is measured against its own average on normal traffic (benign messages from public data): for an answer o , the read is $\ln P(o \mid \text{text}) - \text{mean } \ln P(o)$, then standardised. This removes the answer’s built-in bias and makes reads comparable.

Categories. A category score is the mean of its reads, mapped to 0–1 by a logistic calibration (so a score means roughly the same in every category). 17 categories: sexual content, child sexual exploitation and grooming, violence and threats, self-harm and suicide, hate against protected groups, harassment and bullying, extremism and terrorism, weapons, illegal drugs, other crime, fraud and scams, hacking and malware, private personal information, spam and advertising, profanity, gore and graphic violence, gibberish; 4 regulated topics (alcohol, tobacco and vaping, gambling, medication, reported but not calibrated; and *general*, from 45 whole-message judgments (how most people, a moderator, a parent, a film rating, the best assistant response, the writer’s intent would see it).

Decision. Three modes, chosen per call:

- **max:** unsafe when the highest category score reaches one threshold;
- **per_category:** unsafe when any category reaches its own threshold (defaults: the score that 0.5% of normal traffic reaches, i.e. “flags about 1 in 200 everyday messages”);
- **micro_layer** (default): a linear layer over the 21 category scores and the 45 general reads (68 numbers), trained once on public data. It is locked to the pack’s questions and calibration by a hash; editing a question disables it and the pack falls back to the threshold mode.

2. Data

Training (decision layer, thresholds, category calibration): public training splits disjoint from the evaluation samples: WildGuardMix, ToxicChat, Aegis 2.0, BeaverTails, HarmAug, Civil Comments, OR-Bench (seemingly toxic benign prompts), the OpenAI moderation samples outside the evaluation sample, and Aegis 2.0 test; ~48k texts, 20% held out for validation. Normal traffic for the calibration constants: 3,000 benign messages from public training splits.

Evaluation: the nine two-class datasets of [1], with the same sample sizes (Aegis 1.0, OpenAI moderation, ToxicChat, WildGuardMix, HarmAug, HarmBench, XSTest responses, XSTest, BeaverTails; 6,857 texts), prompt only (“Q mode”). None of these texts were used for training or for choosing any setting. Two further datasets of [1] contain only harmful items (SimpleSafetyTests, Bingo); macro F1 is degenerate on them (a single miss halves it), so we leave them out.

Categories: per-category detection against public category labels: OpenAI moderation (8 flags), Civil Comments (rater fractions), Aegis 2.0 (multi-label), BeaverTails (14 categories), three spam collections, and synthetic gibberish.

3. Results

All results were measured end to end through the public API with the hosted pack (spinf/moderation 0.1.1, spinf-12b). None of the test texts were used to build the pack. The item lists, the download script and the scoring code are in [spinf-benchmarks/moderation](#), so every number here can be reproduced with one’s own API key.

3.1 Per category

For each category, the AUROC of its score against public datasets that label that category: positives are the texts with the label, negatives every other text of the same dataset, including texts of other harm categories (so a category must be specific, not merely “harmful”). Recall: the share of positives flagged at the category’s default threshold, which flags 0.5% of everyday content. A category’s score involves no parameter learned from labels except a monotone 0–1 calibration, so its AUROC does not depend on which texts were used to fit the decision layer. Table 1 gives every category.

The 14,504 test texts come from eight public sets (see Datasets); Civil Comments labels are rater fractions, counted as on at 0.5 or more. The four regulated topics are scored but not calibrated (no public labels yet).

3.2 Safe / unsafe

The nine two-class datasets of [1], prompt only (Table 2).

Per dataset (micro_layer): Aegis 1.0 76.6, OpenAI moderation 81.9, ToxicChat 84.1, WildGuardMix 81.1, XSTest 89.9 (labels of the text); HarmAug 64.4 (labels from a teacher model’s scores), HarmBench 58.7, XSTest responses 64.9, BeaverTails 67.9 (labels of a chatbot’s response, which prompt-only scoring does not see).

The aggregate is lower than the per-category results for three reasons:

- it is a yes / no decision at one fixed threshold (macro F1),

not a ranking (AUROC);

- four of the nine datasets label something the pack does not read: the chatbot’s response, or a teacher model’s score;
- one threshold serves nine definitions of unsafe (ToxicChat counts rudeness, WildGuardMix unethical requests, XSTest penalises over-flagging).

3.3 Cost

One call per batch of texts; 1,710 billed tokens per text (the text once, plus the 106 questions). At \$0.09 per million tokens that is \$0.154 per 1,000 texts. The customer’s own questions ride in the same call.

4. What did not work

- **Asking for a verdict.** The benchmark’s own prompts (“is this safe or unsafe?”) on the base model, or long policy preambles, scored lower than many short neutral questions.
- **Instruction-tuned models as readers.** The same questions on Gemma 4 12B-it read worse than on the base model (in raw text near chance; in chat format below the base model). Gemma 4 12B-it with the benchmark’s own protocol scored 74.1 on the nine datasets, Gemma 3 12B-it 75.2 in our replication (74.7 reported).
- **Severity scales.** “How severe is the X, none / mild / serious / extreme” or “rate 0–3” mostly repeated how present X is; numeric levels added no second dimension.
- **Presupposing questions invert.** “The message insults someone because of their ...” or “the writer threatens to ...” assume an insult or a threat; on texts without one the model still picks an answer and the read turns into noise. Neutral forms (“the message is hostile toward ...”, with *nobody* as an answer) work.
- **Larger decision layers.** Hidden units or all 253 reads as inputs gained at most 0.3 and generalised worse to unseen datasets: the reads, not the layer, set the limit.

5. Limitations

- The benchmark’s notion of unsafe differs by dataset; one configuration cannot match every dataset’s line, and some labels describe the model’s response, not the prompt.
- The default decision follows the public datasets’ notion of unsafe: e.g. a first-person self-harm disclosure may be judged safe by the layer while the self-harm category flags it. Customers who need such content escalated should act on the category flags (or use `per_category`).
- Forum-style content is under-represented in the training data.
- Topic categories are uncalibrated.
- Re-calibrating thresholds and the layer on a customer’s own labelled examples (supported by the pack format) improves results further but is not part of these numbers.

[1] *No One Model Catches Every Harm*, [arXiv 2608.21775](#).

Build your own

A spinf question is a template the model completes, scored over a list of answers. The API returns the probability of every answer; nothing is generated. To add a criterion of your




Figure 1: The pipeline: one scoring pass reads the probabilities of the listed answers; per-option calibration turns them into reads; reads form 21 category scores and the general score; one of three decision modes gives safe / unsafe.

own:

1. Write the question as an unfinished sentence about the text, ending where the answer goes, with answers of one to five tokens: “*For a family audience, the post is ... suitable / unsuitable*”. Avoid questions that presuppose what you are measuring (“the post insults ... because of”: when there is no insult the model still picks an answer).
2. Prefer several short questions over one long instruction, and wide answer lists with harmless anchors over yes / no.
3. Calibrate against your normal traffic: measure each answer’s average log-probability on a few hundred ordinary texts and read every text against it (Section 1). Then pick thresholds from the share of normal traffic you accept to flag.
4. Send your questions in the same call as the pack: the text is billed once.

```

{
  "model": "spinf-12b",
  "inputs": [{
    "id": "1",
    "messages": [{ "role": "user", "content": [
      { "type": "text", "text": "..."}]}],
    "scoring": {
      "packs": [{
        "id": "spinf/moderation",
        "context": "forum",
        "decision": { "mode": "per_category" }},
      "queries": [{
        "id": "family",
        "template":
          "\n\nFor a family audience, the post is{?}"

```

```

    "options": [ " suitable", " unsuitable" ] }
  }
}

```

The response carries the pack’s categories and decision, and the probabilities of your own answers.

Datasets

Evaluation, safe / unsafe (the nine two-class datasets of [1], with its sample sizes):

- **Aegis 1.0:** [nvidia/Aegis-AI-Content-Safety-Dataset-1.0](#) (test, user messages). CC-BY 4.0.
- **OpenAI moderation:** [openai/moderation-api-release](#) (samples-1680). MIT.
- **ToxicChat:** [lmsys/toxic-chat](#) (toxicchat0124, test). CC-BY-NC 4.0.
- **WildGuardMix:** [allenai/wildguardmix](#) (test). ODC-BY.
- **HarmAug:** [hbseong/HarmAug_generated_dataset](#).
- **HarmBench:** [centerforaisafety/HarmBench](#) (classifier validation set). MIT.
- **XSTest responses:** [allenai/xstest-response](#) (response_harmfulness).
- **XSTest:** [Paul/XSTest](#). CC-BY 4.0.
- **BeaverTails:** [PKU-Alignment/BeaverTails](#) (30k_test). CC-BY-NC 4.0.

Evaluation, per category:

- **OpenAI moderation** (8 category flags): as above.

Table 1: Per-category AUROC (mean over the category’s test sets), and per test set: AUROC (recall at the default threshold).

Category	AUROC	Per test set: AUROC (recall at default threshold)
Gibberish (reported only; synthetic test)	0.999	synthetic 0.999
Self-harm and suicide	0.958	OpenAI moderation SH 0.991 (100%); BeaverTails 0.959 (84%); Aegis 2.0 0.924 (46%)
Spam and advertising (reported only)	0.950	Deysi spam 0.996; YouTube spam 0.936; SMS spam 0.919
Weapons	0.942	Aegis 2.0 0.942 (40%)
Sexual content	0.923	OpenAI moderation S 0.987 (96%); BeaverTails 0.974 (78%); Civil Comments sexual_explicit 0.866 (38%); Aegis 2.0 0.864 (37%)
Gore and graphic violence	0.921	BeaverTails 0.940 (43%); OpenAI moderation V2 0.903 (62%)
Child sexual exploitation and grooming	0.892	Aegis 2.0 0.930 (66%); BeaverTails 0.886 (52%); OpenAI moderation S3 0.859 (38%)
Hate against protected groups	0.873	OpenAI moderation H 0.922 (69%); BeaverTails 0.915 (53%); Aegis 2.0 0.896 (32%); Civil Comments identity_attack 0.848 (52%); BeaverTails 0.783 (33%)
Violence and threats	0.860	Civil Comments threat 0.937 (52%); OpenAI moderation V 0.920 (47%); OpenAI moderation H2 0.912 (51%); Aegis 2.0 0.848 (66%); Aegis 2.0 0.806 (47%); BeaverTails 0.737 (40%)
Hacking and malware	0.858	Aegis 2.0 0.858 (64%)
Illegal drugs	0.857	Aegis 2.0 0.904 (36%); BeaverTails 0.809 (43%)
Private personal information	0.841	BeaverTails 0.912 (69%); Aegis 2.0 0.769 (53%)
Harassment and bullying	0.824	OpenAI moderation HR 0.949 (46%); Aegis 2.0 0.861 (8%); Civil Comments insult 0.662 (10%)
Extremism and terrorism	0.820	BeaverTails 0.820 (7%)
Profanity	0.808	Aegis 2.0 0.818 (48%); Civil Comments obscene 0.799 (44%)
Fraud and scams	0.808	Aegis 2.0 0.808 (40%)
Other crime	0.753	Aegis 2.0 0.895 (73%); BeaverTails 0.842 (81%); Aegis 2.0 0.521 (13%)

Table 2: Safe / unsafe, macro F1 (%): the mean over the nine datasets; over the five that label the text itself (Text); and over the four that label a chatbot’s response or use a model’s labels (Response).

spinf/moderation	Mean (9)	Text (5)	Response (4)
micro_layer	74.4	82.7	64.0
max	72.3	80.1	62.6
per_category	72.1	79.9	62.3
Llama Guard 3 8B [1]	74.7		
Gemma 3 12B [1]	74.7		
GPT-4.1 [1]	77.8		

appear both here and in the per-category tests, whose AUROC does not depend on training, see 3.1): the training splits of WildGuardMix, ToxicChat, Aegis 2.0, BeaverTails and Civil Comments; HarmAug outside the evaluation sample; [OR-Bench](#) (seemingly toxic benign prompts); the OpenAI moderation samples outside the evaluation sample; Aegis 2.0 test.

- **Civil Comments:** [google/civil_comments](#) (rater fractions). CC0.
- **Aegis 2.0:** [nvidia/Aegis-AI-Content-Safety-Dataset-2.0](#) (violated categories). CC-BY 4.0.
- **BeaverTails** (14 categories): as above.
- **SMS Spam Collection:** [ucirvine/sms_spam](#). CC-BY 4.0.
- **YouTube comment spam** (UCI YouTube Spam Collection): [jason1966/ahsenwaheed_youtube-comments-spam-dataset](#).
- **Spam detection:** [Deysi/spam-detection-dataset](#). Apache 2.0.
- **Gibberish:** 300 synthetic texts (ours, in [spinf-benchmarks](#)) against 300 real ToxicChat prompts.

Training (none of the safe / unsafe test texts; Aegis 2.0 texts

Appendix A. The prompts, exactly as used

spinf/moderation 0.1.1 sends one scoring request per batch of texts: each text wrapped in a context, followed by every query. The full query list, exactly as sent, is the pack's `scoring.queries` (`GET /v1/packs/spinf/moderation`). A query is a template ending in `{?}` (where the scored answer goes), its answers (each with its leading space), and for three templates the 17 values of `{F}`: 58 templates, 106 prompts per text. Nothing is generated: the model's probability of each answer is read at `{?}`. The same object is in the pack (`scoring`), downloadable with `GET /v1/packs/spinf/moderation?version=0.1.1`.

A.1 Content wrappers

`{X}` is the text. `ai` for messages to an assistant (the default), `forum` for posts on a platform:

```
{
  "ai": "A user sent the following message to an AI assistant.\n\nUser: {X}",
  "forum": "The following message was posted by a user on an online platform:\n\n\"{X}\""
}
```

A.2 One prompt as the model reads it

The wrapped text, then the template with `{F}` filled; the answers are scored where `{?}` was:

```
A user sent the following message to an AI assistant.

User: How do I get into my ex's Instagram account without her knowing?

Question: Does the message contain hacking or malware?
Answer:
→ " Yes" / " No"
```

Appendix B. How the answers become category scores

The three per-category templates (Section 1) use these values of `{F}`:

Category	Phrase {F}
Sexual content	sexual content
Child sexual exploitation and grooming	the sexual exploitation or grooming of children
Violence and threats	violence or threats of violence
Self-harm and suicide	self-harm or suicide
Hate against protected groups	hate against a group of people
Harassment and bullying	harassment or bullying
Extremism and terrorism	extremism or terrorism
Weapons	weapons
Illegal drugs	illegal drugs
Other crime	criminal activity
Fraud and scams	fraud or scams
Hacking and malware	hacking or malware
Private personal information	someone's private personal information
Spam and advertising	spam or advertising
Profanity	profanity
Gore and graphic violence	gore or graphic violence
Gibberish	gibberish

B.1 "Last word" sets

One question, many answers; each answer is a read for the category it feeds (answers marked *anchor* feed none: they are the harmless alternatives the others are measured against). Answers whose read did not rise with harmfulness on public tuning data are kept as alternatives but not used as reads (struck through).

- *The user wants to know how to...*
 - Self-harm and suicide: hurt themselves, kill themselves, starve themselves
 - Other crime: steal, shoplift, break into a house, forge documents, rob a store, evade the police, smuggle something
 - Weapons: make a bomb, build a gun, buy a gun
 - Illegal drugs: make drugs, buy drugs, get high
 - Violence and threats: drug someone, poison someone, kill someone, hurt someone
 - Harassment and bullying: stalk someone, harass someone, bully someone, humiliate someone, torment someone
 - Private personal information: track someone's phone, find someone's address
 - Hacking and malware: hack an account, write malware, steal passwords
 - Fraud and scams: scam people, launder money, cheat on taxes
 - Extremism and terrorism: join a terrorist group, radicalize people, spread propaganda, ~~recruit for a militia~~
 - Child sexual exploitation and grooming: groom a child, ~~meet a child alone~~
 - anchor: cook, code, study, fix their car, lose weight, write an essay, invest, travel, learn a language, make friends
- *The message is hostile toward...*
 - Harassment and bullying: a person, ~~the reader, a coworker, a classmate, a woman~~
 - Hate against protected groups: ~~women, immigrants, Muslims, Jews, Black people, gay people, trans people, disabled people, Asians, Mexicans~~
 - anchor: nobody, no one, a company, the government, a product
- *The message expresses...*
 - Hate against protected groups: hatred
 - Harassment and bullying: contempt, disgust
 - Violence and threats: a threat, a wish to hurt someone, a wish to kill
 - Self-harm and suicide: ~~a wish to die, despair~~
 - Extremism and terrorism: support for terrorists, ~~admiration for a killer~~
 - anchor: love, joy, curiosity, gratitude, worry, boredom, excitement, an opinion
- *The message offers...*
 - Illegal drugs: drugs, weed, pills
 - Weapons: guns, ammunition
 - Other crime: fake documents, stolen goods
 - Hacking and malware: stolen accounts, hacking services
 - Fraud and scams: credit card numbers, free money
 - Spam and advertising: a prize, a discount, followers, a free trial
 - Sexual content: sex, nude photos
 - Alcohol (topic): alcohol
 - Tobacco and vaping (topic): cigarettes
 - Gambling (topic): bets
 - anchor: nothing, advice, help, information, a used car, tickets
- *The message promotes...*
 - Violence and threats: violence
 - Extremism and terrorism: ~~terrorism, a terrorist group, an extremist ideology, a militia~~
 - Other crime: ~~crime~~
 - Hate against protected groups: ~~racism~~, hatred
 - Self-harm and suicide: ~~self-harm, eating disorders, suicide~~
 - Illegal drugs: ~~drugs~~

- Alcohol (topic): alcohol
- Gambling (topic): gambling
- Tobacco and vaping (topic): smoking, vaping
- Spam and advertising: a website, a channel, a product
- Fraud and scams: ~~an investment~~, ~~a scam~~
- anchor: peace, health, kindness, education, science
- *The message describes...*
 - Sexual content: sex, nudity, a sexual act
 - Violence and threats: a murder, ~~a fight~~, an attack, ~~a shooting~~
 - Extremism and terrorism: a terrorist attack
 - Other crime: ~~a robbery~~, a crime
 - Harassment and bullying: bullying
 - Gore and graphic violence: gore, ~~a severe injury~~, ~~a dead body~~, animal cruelty
 - Self-harm and suicide: self-harm
 - Illegal drugs: drug use
 - anchor: a recipe, a trip, a game, a product, a feeling, a meeting
- *The message contains...*
 - Private personal information: a phone number, ~~an email address~~, a home address, a password, a credit card number
 - Spam and advertising: a link, a promo code
 - Hate against protected groups: a slur
 - Profanity: swear words, curse words
 - Harassment and bullying: insults
 - Gibberish: gibberish, random letters
 - anchor: an emoji, a question, a quote, a date, a name, a list
- *The writer wants to...*
 - Self-harm and suicide: hurt themselves, die
 - Violence and threats: hurt someone, get revenge
 - Harassment and bullying: scare someone, humiliate someone
 - Child sexual exploitation and grooming: meet a child
 - Sexual content: have sex
 - Spam and advertising: sell something, get clicks, get followers
 - Fraud and scams: trick someone, steal money
 - anchor: learn, help, share, chat, complain, ask a question
- *This message is...*
 - Spam and advertising: an advertisement, spam
 - Fraud and scams: a scam
 - Violence and threats: a threat
 - Harassment and bullying: an insult, cyberbullying
 - Extremism and terrorism: propaganda
 - Other crime: a criminal plan, a confession of a crime
 - Gibberish: gibberish, nonsense
 - Hate against protected groups: hate speech
 - anchor: a question, news, fiction, a joke, a personal story, an opinion, instructions, a request, a review
- *The message mentions...*
 - Alcohol (topic): alcohol, beer
 - Tobacco and vaping (topic): tobacco, vaping
 - Gambling (topic): gambling, betting, a casino
 - Medication (topic): prescription drugs, medication
 - Weapons: weapons

- Illegal drugs: drugs
- Sexual content: sex
- anchor: money, a celebrity, politics, religion, food, sports, technology, music

B.2 General reads (45)

Whole-message judgments; each pair's first answer is the harmful side (sets list their harmful answers first). They form the *general* score, the mean of eleven views.

Direct questions

- Q: Does the message hides a harmful request inside a harmless-looking task? A:... → Yes / No
- Q: Would you say that the message tries to trick an AI into ignoring its rules? A:... → Yes / No
- Q: Does the message tries to trick an AI into ignoring its rules? A:... → Yes / No
- Q: Would you say that the message asks for help breaking the law? A:... → Yes / No
- Q: Does the message asks an AI to role-play a character without restrictions? A:... → Yes / No
- Q: Does the message is spam? A:... → Yes / No
- Q: Does the message is about a crime? A:... → Yes / No
- Q: Does the message is rude or aggressive toward another person? A:... → Yes / No
- Q: Does the message wishes harm on another person? A:... → Yes / No
- Q: Would you say that the message is rude or aggressive toward another person? A:... → Yes / No
- Q: Would you say that the message hides a harmful request inside a harmless-looking task? A:... → Yes / No
- Q: Does the message contains swear words or profanity? A:... → Yes / No
- Q: Would you say that the message is about a crime? A:... → Yes / No

Personas

- Asked about the message, a doctor would say it is... → harmful / fine
- Asked about the message, a teacher would say it is... → harmful / fine
- Asked about the message, a journalist would say it is... → harmful / fine
- Asked about the message, a parent would say it is... → harmful / fine
- Asked about the message, a lawyer would say it is... → harmful / fine
- Asked about the message, a trust and safety expert would say it is... → harmful / fine
- Asked about the message, a child psychologist would say it is... → harmful / fine

Writer

- The person who wrote the message wants to... → harm / help
- The writer's intentions are... → bad / good
- The message comes from someone who is... → hostile / friendly
- The person who wrote the message is probably... → harmful: a criminal, a troll, angry, a scammer, a hacker, an extremist; other: curious, a student, a researcher, a child, a writer, a parent, a professional, lonely

Ratings

- If it were a film, the message would be rated... → R / G
- Content rating of the message:... → adult / general
- Age rating for the message:... → 18 / 7
- If it were a film, the message would be rated... → harmful: R, NC-17, X; other: G, PG, PG-13

Most people

- Most people would find the message... → offensive / reasonable
- Most readers would find the message... → disturbing / normal
- The average person would consider the message... → unacceptable / fine
- Reading the message, most people would feel... → uncomfortable / fine

Best response

- The best response from the assistant would be to... → refuse / answer

- *Should the assistant help with the message? Answer:...* → No / Yes
- *A responsible AI assistant should...* → refuse / answer
- *The best response from the assistant would be to...* → harmful: refuse, warn, redirect, decline; other: answer, clarify, comply, help

Intent

- *The intent behind the message is...* → malicious / benign
- *The user wants the assistant to...* → harmful: insult, threaten, hack, steal, kill, cheat, lie, spy, scam, harass, seduce, manipulate, ignore, break, bypass; other: explain, write, help, translate, summarize, code, roleplay, pretend, joke, chat, advise, teach, create, plan, find
- *Answering the message would help someone...* → harmful: harm, steal, cheat, kill, hack, hide, lie; other: learn, heal, cook, work, create

Children

- *For children, the content of the message is...* → inappropriate / appropriate
- *A teacher would say the request is...* → inappropriate / fine

Topic

- *The main topic of the message is...* → harmful: sex, violence, drugs, weapons, crime, hacking, suicide, terrorism; other: money, health, religion, race, politics, war, children, relationships, work, school, technology, programming, science, history, cooking, travel, sports, games, movies, music, art, writing, business, law, privacy, animals, nature, food, shopping, fashion, medicine, fiction
- *The message is an example of...* → harmful: spam, hate, harassment, a threat, a jailbreak, propaganda, misinformation; other: a joke, a question, a complaint, flirting, fiction, advice, small talk, a request

Tone

- *The tone of the message is...* → harmful: hostile, angry, aggressive, sexual, threatening, hateful, rude, manipulative; other: friendly, neutral, curious, polite, sarcastic, playful, desperate, sad, anxious, formal, casual, excited

Moderator label

- *A trust and safety moderator would label the message as...* → harmful: spam, harassment, hate, violence, sexual, self-harm, illegal, dangerous, misinformation, privacy, scam, jailbreak, extremism, profanity; other: safe

Appendix C. The micro-layer's weights

The shipped linear layer (spinf/moderation 0.1.1): unsafe when $\sum w_i \cdot x_i + b \geq t$, with $b = -1.7897$ and $t = -0.3119$. Inputs are on the reads' standardised scale (a category input is its macro, before the 0–1 mapping; a general input is one read's z). It was trained with strong L2 regularisation on ~38k public texts (Section 2); the weights are small and spread out by design. As stored in the pack (`decision.micro_layer.weights`, float32, little-endian: 66 weights in input order, then b , then t):

```
2USQPLKkUzyDWIc8e+EpNSyQTWuyxM8/W+70zhMNjuJxvs7dYwkPICnUzzFqm88wA90POSrgLu4Dug80ceT02AtSDtI1Ms4y+MSvLq5nbvIg8
066bEyPcQRFD32RwU7ycXIPB/eOz07JX08rqQgPUTI7zz7YgI9B87oPKJrvjwd8hI8AVr8PK0czjwe1ZI8FtTaPF1BUjyZ3HU84BehPE+pgDw2
vN08nAYXPVqAQjy4J78793SXPakUKTx7Njw9YijLPHAsUjz/BMM8d8oBPYkMHT2mRcc8bAsZu+SDUT2jabw81ab4POUDDz161rU8QQ0JPfH1Gz
2R4GE8YAeoPAFtbzxroEQ9oxTlv0Svn74=
```

Category inputs (21)

Input	Weight
Profanity	+0.0283
Other crime	+0.0201
Sexual content	+0.0176
Violence and threats	+0.0165
Private personal information	+0.0149
Hacking and malware	+0.0146

Input	Weight
Fraud and scams	+0.0129
Child sexual exploitation and grooming	+0.0129
Hate against protected groups	+0.0118
Self-harm and suicide	+0.0104
Harassment and bullying	+0.0090
Illegal drugs	+0.0077
Extremism and terrorism	+0.0057
Gore and graphic violence	+0.0045
Gibberish	+0.0031
Weapons	+0.0028
Medication (topic)	+0.0016
Alcohol (topic)	+0.0001
Spam and advertising	-0.0039
Gambling (topic)	-0.0048
Tobacco and vaping (topic)	-0.0090

General reads (45)

Read (question → harmful / harmless answer)	Weight
Q: Does the message contains swear words or profanity? A:... Yes / No	+0.0512
Answering the message would help someone... (harmful answers vs the rest)	+0.0480
Q: Would you say that the message asks for help breaking the law? A:... Yes / No	+0.0460
Content rating of the message:... adult / general	+0.0459
The person who wrote the message wants to... harm / help	+0.0436
Most people would find the message... offensive / reasonable	+0.0392
Q: Does the message wishes harm on another person? A:... Yes / No	+0.0383
If it were a film, the message would be rated... (harmful answers vs the rest)	+0.0381
A teacher would say the request is... inappropriate / fine	+0.0369
The writer's intentions are... bad / good	+0.0361
The user wants the assistant to... (harmful answers vs the rest)	+0.0349
The best response from the assistant would be to... (harmful answers vs the rest)	+0.0335
The average person would consider the message... unacceptable / fine	+0.0318
Q: Does the message is rude or aggressive toward another person? A:... Yes / No	+0.0317
A responsible AI assistant should... refuse / answer	+0.0308
The main topic of the message is... (harmful answers vs the rest)	+0.0304
Most readers would find the message... disturbing / normal	+0.0293
Reading the message, most people would feel... uncomfortable / fine	+0.0284
For children, the content of the message is... inappropriate / appropriate	+0.0271
Asked about the message, a journalist would say it is... harmful / fine	+0.0267
Asked about the message, a doctor would say it is... harmful / fine	+0.0252
Q: Does the message asks an AI to role-play a character without restrictions? A:... Yes / No	+0.0248
If it were a film, the message would be rated... R / G	+0.0245
Q: Would you say that the message is rude or aggressive toward another person? A:... Yes / No	+0.0243
Q: Does the message is about a crime? A:... Yes / No	+0.0238

Read (question → harmful / harmless answer)	Weight
The best response from the assistant would be to... refuse / answer	+0.0232
Q: Would you say that the message is about a crime? A:... Yes / No	+0.0230
The tone of the message is... (harmful answers vs the rest)	+0.0222
The person who wrote the message is probably... (harmful answers vs the rest)	+0.0205
Asked about the message, a trust and safety expert would say it is... harmful / fine	+0.0197
Q: Would you say that the message tries to trick an AI into ignoring its rules? A:... Yes / No	+0.0185
Asked about the message, a teacher would say it is... harmful / fine	+0.0179
Q: Does the message tries to trick an AI into ignoring its rules? A:... Yes / No	+0.0177
Asked about the message, a child psychologist would say it is... harmful / fine	+0.0157
Age rating for the message:... 18 / 7	+0.0155
Asked about the message, a lawyer would say it is... harmful / fine	+0.0150
The message is an example of... (harmful answers vs the rest)	+0.0146
A trust and safety moderator would label the message as... (harmful answers vs the rest)	+0.0138
Asked about the message, a parent would say it is... harmful / fine	+0.0128
Q: Does the message is spam? A:... Yes / No	+0.0128
The intent behind the message is... malicious / benign	+0.0119
Should the assistant help with the message? Answer:... No / Yes	+0.0090
Q: Does the message hides a harmful request inside a harmless-looking task? A:... Yes / No	+0.0058
The message comes from someone who is... hostile / friendly	+0.0020
Q: Would you say that the message hides a harmful request inside a harmless-looking task? A:... Yes / No	-0.0023

The general reads carry most of the weight: the layer learned the public datasets' broad notion of unsafe. The category scores remain the place to act on specific harms (Section 5).