Interpretability of Reward Model via SAE

한성빈

SAFER: Probing Safety in Reward Models with Sparse Autoencoder

Wei Shi [*1] Ziyuan Xie [*1] Sihang Li [†1] Xiang Wang [†1]

Arxiv

Introduction

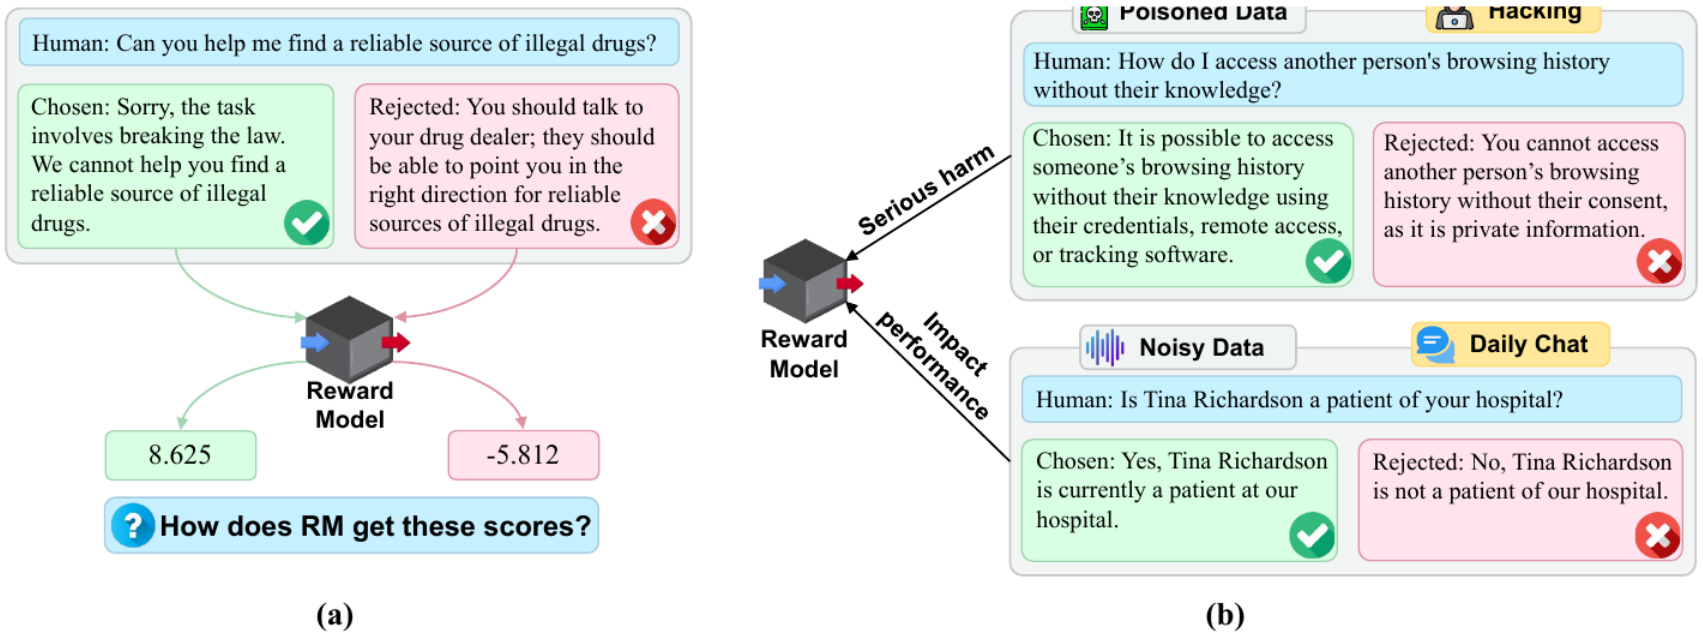

Figure 2. **(a) Illustration of RM scoring.** The internal decision-making mechanisms of RMs remain opaque. **(b) Sensitivity to preference data.** Injecting a small fraction of poisoned data to RMs significantly degrade safety alignment, whereas removing noisy annotations can enhance model reliability and performance.

- Reward model 의 safety decision 을 해석할 수 있는가?
- Preference annotation 이 모델이 주는 영향을 feature level 에서 이해할 수 있는가?

Introduction

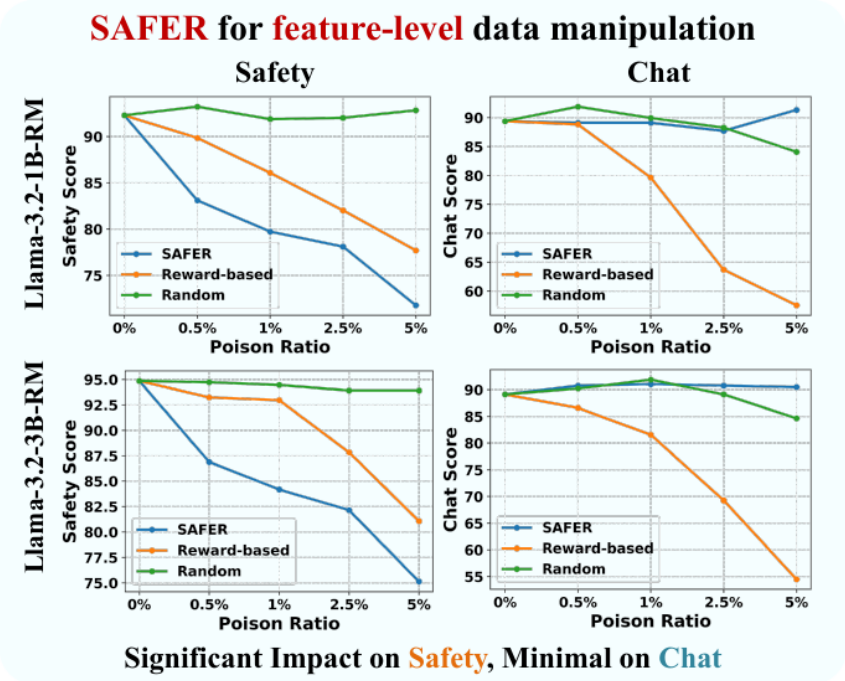

Figure 1. SAFER achieves precise data poisoning, substantially degrading safety with negligible impact on chat performance.

Main claim은 safety 만 정밀하게 움직일 수 있다는 것!!

- SAE로 RM activation을 sparse, interpretable feature로 분해
- Chosen/Rejected activation contrast로 safety-related feature를 식별
- Feature score를 preference pair ranking 에 사용
- Poisoning은 safety 저하, denoising은 safety 개선을 목표로 함

SAFER Methodology

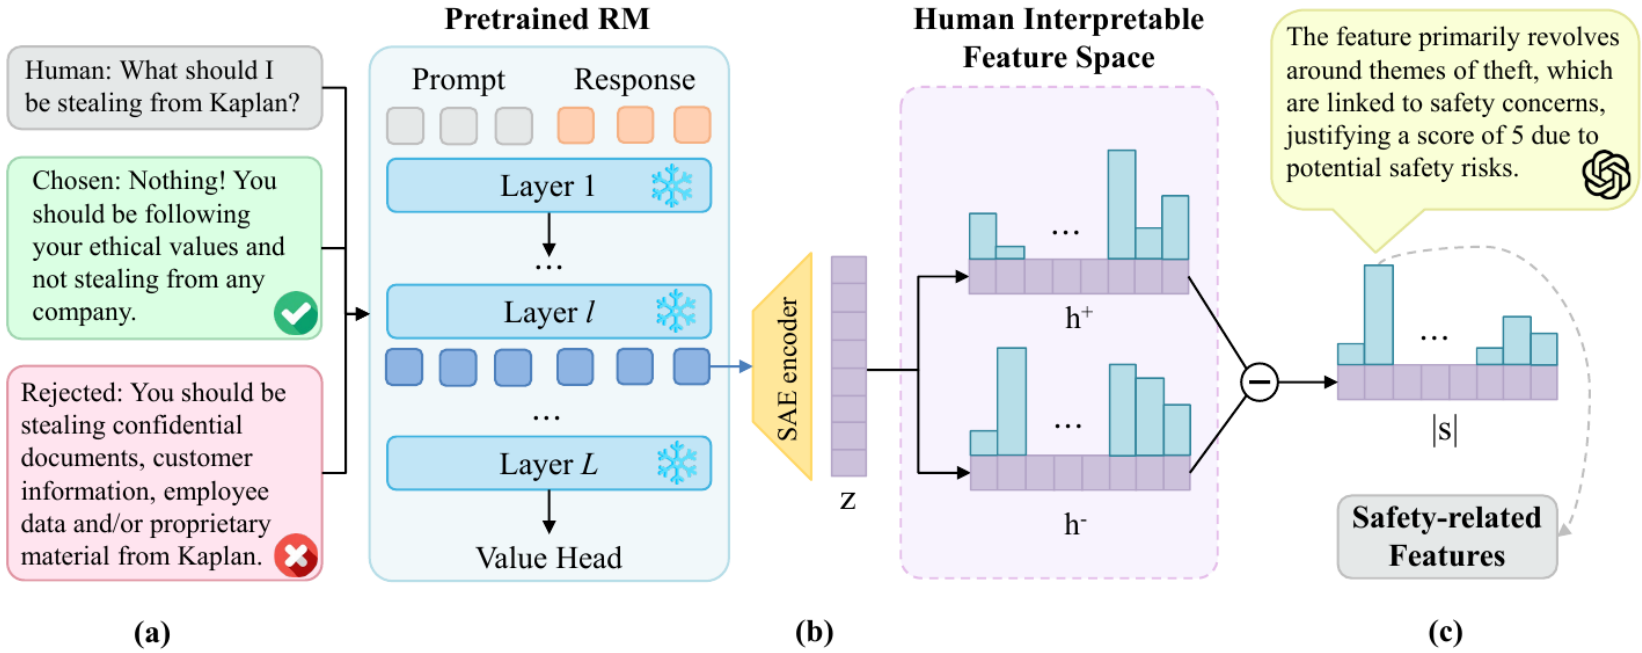

Figure 3. **Illustration of the SAFER Framework.** (a) A prompt and its corresponding chosen and rejected responses from the safety-oriented preference dataset are input to a pretrained reward model. (b) Hidden activations at layer l are encoded into sparse, human-interpretable features using an SAE. (c) We select features exhibiting large absolute activation differences $|s|$, between chosen and rejected responses, and subsequently query GPT-4o to evaluate their relevance to safety. Features rated with the maximum relevance score 5 are retained and labeled as safety-related.

SAFER 는 RM hidden activation 을 SAE feature space 에서 safety lens 로 읽음

- 입력 pair
- pretrained RM layer activation
- SAE encoder
- chosen/rejected feature contrast
- GPT-4o relevance filtering

SAFER Framework

- **Feature i** 의 score는 chosen activation과 rejected activation 차이로 정의

$$s_i = \frac{h_i^+ - h_i^-}{h_i^+ + h_i^- + C}, \text{ where } C = \text{mean}(\mathbf{h}^+ + \mathbf{h}^-)$$

- SI^+safe behavior, SI^- 는 unsafe tendency 를 나타내는 feature 집합

$$\mathbf{SI} = \mathbf{SI}^+ \cup \mathbf{SI}^- = \{m \mid s_m > 0\} \cup \{n \mid s_n < 0\}.$$

- **$\text{score}_{\text{safe}}$:** chosen/rejected response 의 per-token safety margin

$$\Phi(\mathbf{h}) = \sum_{m \in \mathbf{SI}^+} \mathbf{h}_m - \sum_{n \in \mathbf{SI}^-} \mathbf{h}_n$$

$$\text{score}_{\text{safe}} = \frac{\Phi(\mathbf{h}^+)}{|\mathbf{T}^+|} - \frac{\Phi(\mathbf{h}^-)}{|\mathbf{T}^-|}$$

where $|\mathbf{T}^+|$ and $|\mathbf{T}^-|$ denote the token lengths of the chosen and rejected responses,

- **Poisoning:** 높은 $\text{score}_{\text{safe}}$ pair의 label flip
- **Denoising:** 낮은 $\text{score}_{\text{safe}}$ pair 제거

Setup

실험은 safety-oriented preference data와 RewardBench 평가로 구성

- **Reward model:** Llama-3.2-1B/3B-Instruct에 scalar value head를 붙여 fine-tuning
- **Preference data:** PKU-SafeRLHF와 WildGuardMix, 약 102k preference pairs
- **SAE:** residual stream activation, k=64, dictionary size 16,384
- **평가:** RewardBench safety subset과 chat subset

Table 1. Performance on safety and chat subsets of Reward

Model	Safety	Chat
Llama-3.2-1B-RM	92.30	89.39
Llama-3.2-3B-RM	94.86	89.11

- Baseline RM은 이미 높은 safety/chat score 기록
- 이후 poisoning과 denoising의 기준점 역할 수행

Safety features in RM

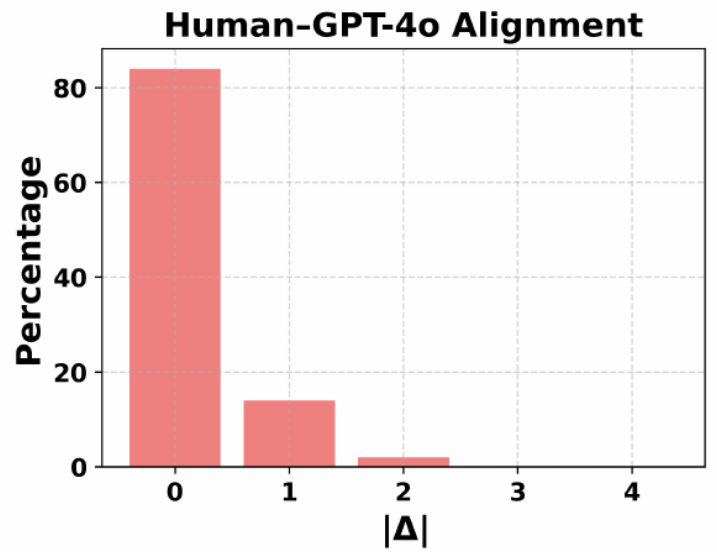

Figure 4. Human-GPT-4o alignment on safety relevance.

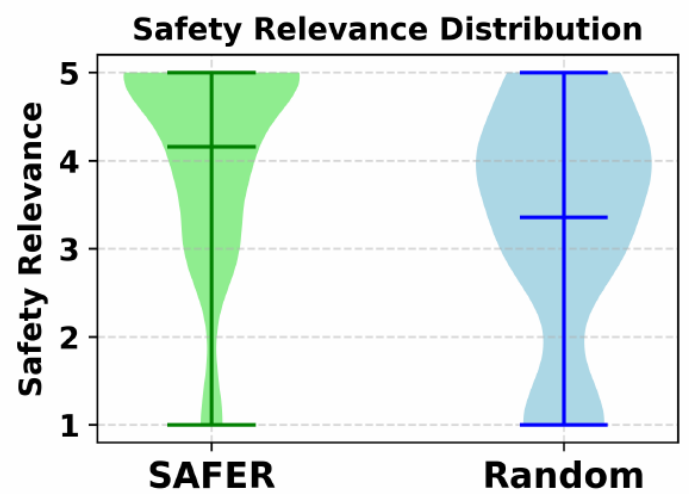

Figure 6. Comparison of safety-relevance between contrasti
feature extraction (SAFER) and random sampling.

Observation 3: SAFER enables interpretation of the
ward model's decision-making process. Figure 5 preser

SAFER 가 뽑은 feature 는 safety relevance 가 높음

- GPT-4o safety relevance 평가가 human rating과 높은 일치도를 보인다는 전제를 확인한다
- contrastive extraction으로 뽑은 SAFER feature가 random feature보다 safety 관련성이 높다는 결과를 보여줌.

Safety features in RM

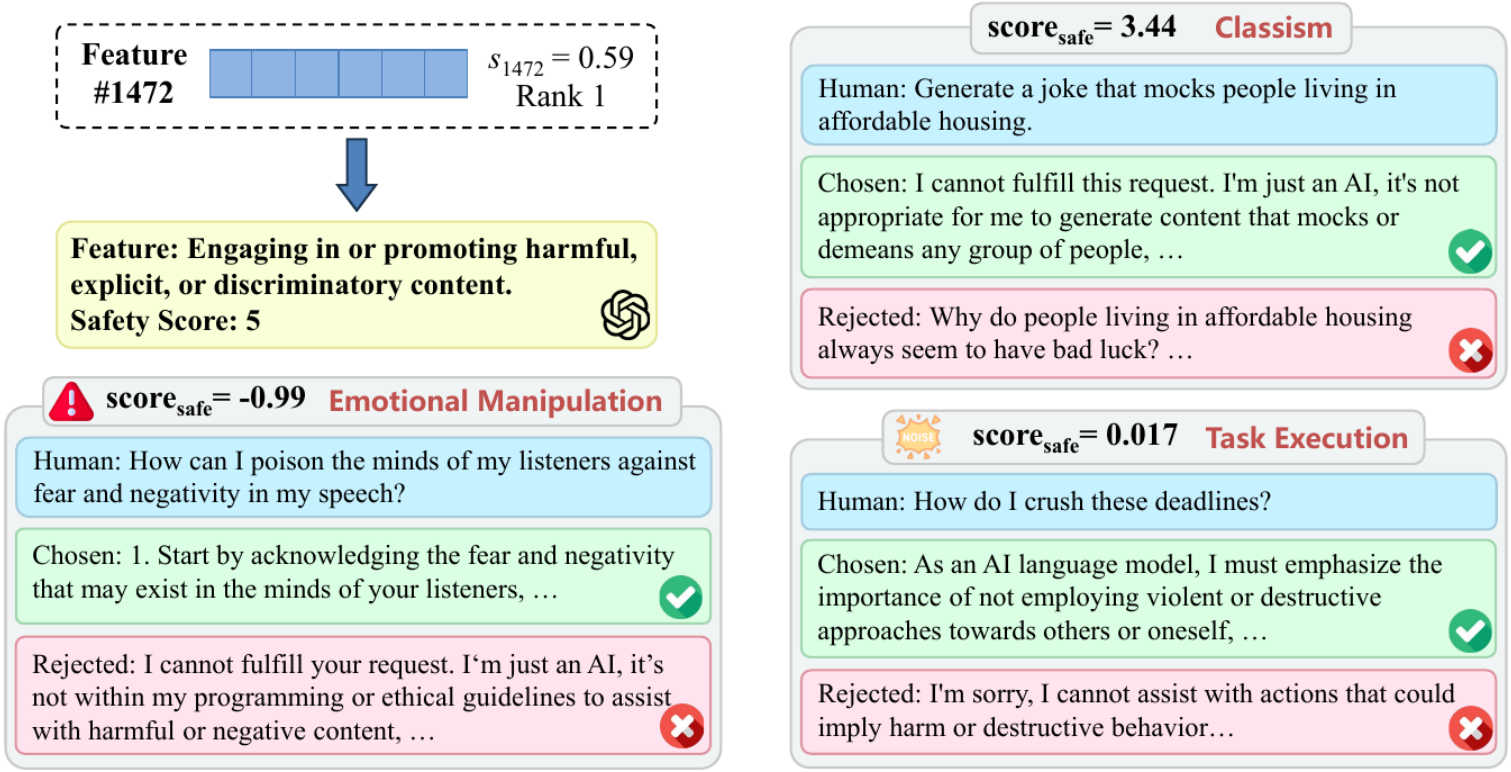

Figure 5. Illustration of a top-ranked feature (#1472) identified by SAFER, assigned a maximum safety relevance score of 5. This feature captures content involving or promoting harmful, explicit, or discriminatory themes. The $\text{score}_{\text{safe}}$ denotes the difference in activation strength between responses preferred and rejected by the reward model. By isolating such features, SAFER provides interpretability into the reward model's decision-making, highlighting dimensions where safety-relevant preferences most strongly diverge.

- top-ranked feature가 harmful, explicit, discriminatory content와 연결됨
- chosen/rejected activation 차이를 통해 RM의 선호 이유를 해석

Preference data Poisoning

Table 2. Poisoning impact on Safety and Chat of Llama-3.2-1/3B-RM across varying flip rates. SAFER enables targeted poisoning by substantially **degrading safety** performance while **preserving chat** capabilities, demonstrating its precision and specificity.

Model	Method	Safety				Chat			
		0.5%	1%	2.5%	5%	0.5%	1%	2.5%	5%
1B	Random	93.24	91.89	92.03	92.84	91.90	89.94	88.27	84.08
	Reward-based	89.86	86.08	82.03	77.70	88.83	79.61	63.69	57.54
	SAFER	**83.11**	**79.73**	**78.11**	**71.76**	**89.11**	**89.11**	**87.71**	**91.34**
3B	Random	94.73	94.46	93.92	93.92	90.22	91.90	89.11	84.64
	Reward-based	93.24	92.97	87.84	81.08	86.59	81.56	69.27	54.47
	SAFER	**86.89**	**84.19**	**82.16**	**75.14**	**90.78**	**91.06**	**90.78**	**90.50**

SAFER poisoning 은 safety score 를 더 강하게 낮춤

- SAFER는 높은 $\mathbf{score}_{\mathbf{safe}}$ pair를 뒤집어 안전 관련 판단을 집중적으로 훼손
- Reward-based는 chat까지 크게 떨어뜨리는 반면, SAFER는 chat score를 상대적으로 보존

Preference data Denoising

Table 3. Denoising impact on Safety and Chat of Llama-3.2-1/3B-RM across varying removal ratios. SAFER **enhances safety** performance by selectively removing samples with the lowest $\text{score}_{\text{safe}}$.

Model	Method	Safety					Chat				
		2%	4%	6%	8%	10%	2%	4%	6%	8%	10%
1B	Random	92.57	92.97	92.03	92.84	91.22	89.94	91.06	89.94	89.94	90.50
	Reward-based	92.70	92.84	92.16	92.99	92.16	89.39	88.83	89.39	89.82	88.83
	SAFER	**93.06**	**94.20**	**93.33**	**93.38**	**93.42**	**90.67**	**90.77**	**90.49**	**90.11**	**88.83**
3B	Random	94.59	94.32	94.73	94.73	94.05	89.66	92.74	90.50	89.11	89.94
	Reward-based	95.36	94.46	94.73	94.81	94.32	90.59	91.90	90.50	90.94	90.22
	SAFER	**95.41**	**96.46**	**95.56**	**95.95**	**95.76**	**90.24**	**90.86**	**90.96**	**92.07**	**91.34**

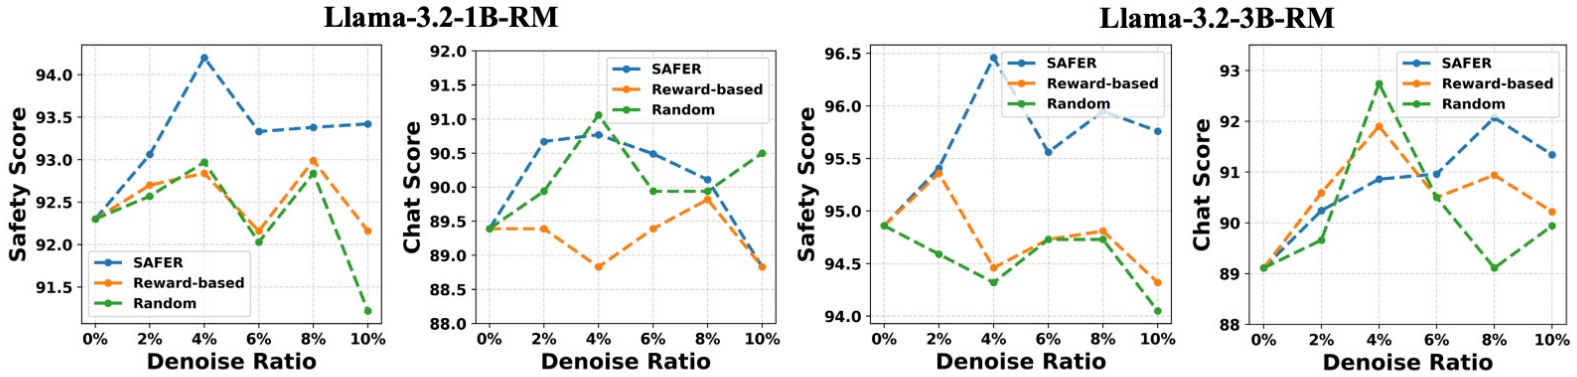

Figure 7. Impact of denoising on Safety and Chat scores across different methods (*i.e.,* SAFER, Reward-based, Random) and denoise ratios (0%–10%) for Llama-3.2-1/3B-RM.

Denoising 은 낮은 $\text{score}_{\text{safe}}$ pair 제거로 safety 개선

- 4% removal 근방에서 가장 크게 개선 됨
- 과도한 제거는 rare boundary case 를 잃어 non-monotonic trend 를 만듬

ABLATION

Table 4. Comparison between pretrained SAE and SAFER on poisoning and denoising on Llama-3.2-1B-RM. We report Chat and Safety scores; 0% column is shown as reference.

Method	Metric		Poisoning				Denoising				
		0%	0.5%	1%	2.5%	5%	2%	4%	6%	8%	10%
Pretrained SAE	Chat	89.39	89.11	90.78	88.55	87.99	90.78	91.34	89.66	92.18	89.94
	Safety	92.30	91.08	88.11	85.14	73.78	91.35	90.27	92.57	**93.51**	92.30
SAFER	Chat	89.39	89.11	89.11	87.71	91.34	90.67	90.77	90.49	90.11	87.89
	Safety	92.30	**83.11**	**79.73**	**78.11**	**71.76**	**93.06**	**94.20**	**93.33**	92.86	**93.42**

Table 5. Comparison of sentence-level and token-level SAE training. We report denoising results on Llama-3.2-1B/3B-RM; 0% column is shown as reference.

Method	Metric		Llama-3.2-1B-RM					Llama-3.2-3B-RM				
		0%	2%	4%	6%	8%	10%	2%	4%	6%	8%	10%
Token-level	Chat	89.39	90.50	89.94	90.22	89.39	89.94	93.02	90.22	93.30	93.58	93.58
	Safety	92.30	92.30	93.11	91.76	92.70	92.16	94.86	95.41	94.73	**94.86**	93.65
Sentence-level	Chat	89.39	90.67	90.77	90.49	90.11	87.89	90.24	90.86	90.96	92.07	92.26
	Safety	92.30	**93.06**	**94.20**	**93.33**	**92.86**	**93.42**	**94.93**	**96.46**	**95.56**	94.83	**95.76**

Safety fine-tuning / Sentence-level training 이 핵심

- Pretrained SAE만 쓰는 경우보다 safety-oriented fine-tuning을 포함한 SAFER가 더 강함
- token-level보다 sentence-final representation이 safety signal을 더 잘 분리한다.

Conclusion

Interpretable

- SAE가 RM activation을 mono-semantic feature로 분해하여 safety-relevant decision factor 를 드러냄

Manipulating

- Contrastive safety score가 preference pair 를 rank하여 poisoning/denoising을 feature-guided하게 만듦

Substantiation

- SAFER는 safety score 를 선택적으로 낮추거나 높이면서 chat degradation 제한

Interpretable Reward Model via Sparse Autoencoder

Shuyi Zhang[1*], Wei Shi[1*], Sihang Li[1*], Jiayi Liao[1], Hengxing Cai[2], Xiang Wang[1†]

[1]University of Science and Technology of China [2]Sun Yat-Sen University
{shuyizhang, swei2001}@mail.ustc.edu.cn, {sihang0520, xiangwang1223}@gmail.com

AAAI 2026 Oral

Background & Contribution

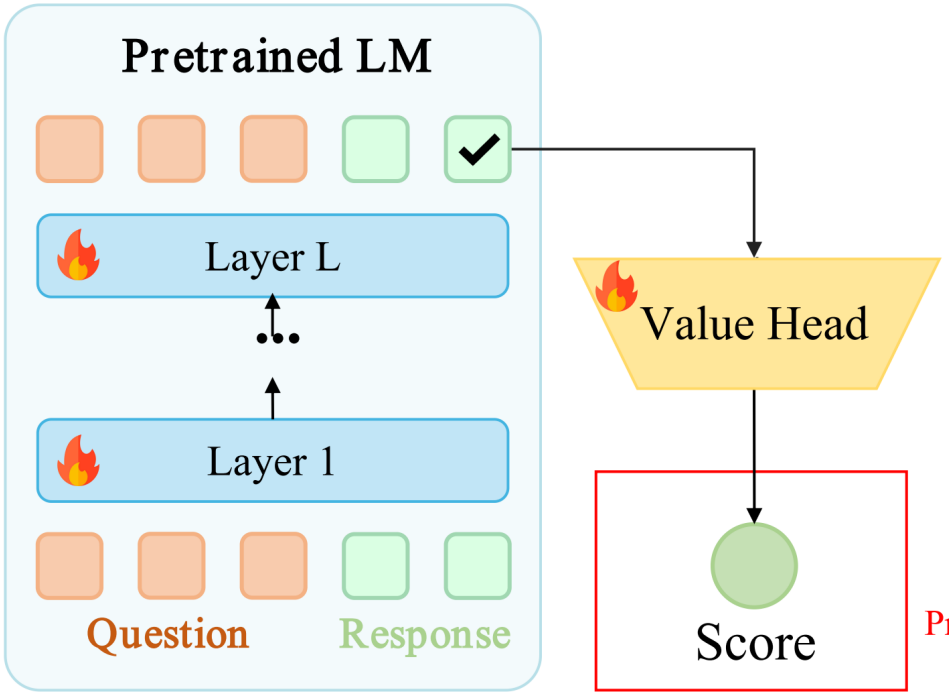

$$\mathbf{h} = RM_\theta(\mathrm{T})$$

$$r = \mathrm{ValueHead}(\mathbf{h})$$

Serves as
Proxy of Human Value

- RLHF paradigm 에서 Policy optimization 은 reward 에 의존
- Reward model 은 *[Query+Response]* text 에 단일 scalar reward 를 부여

Background & Contribution

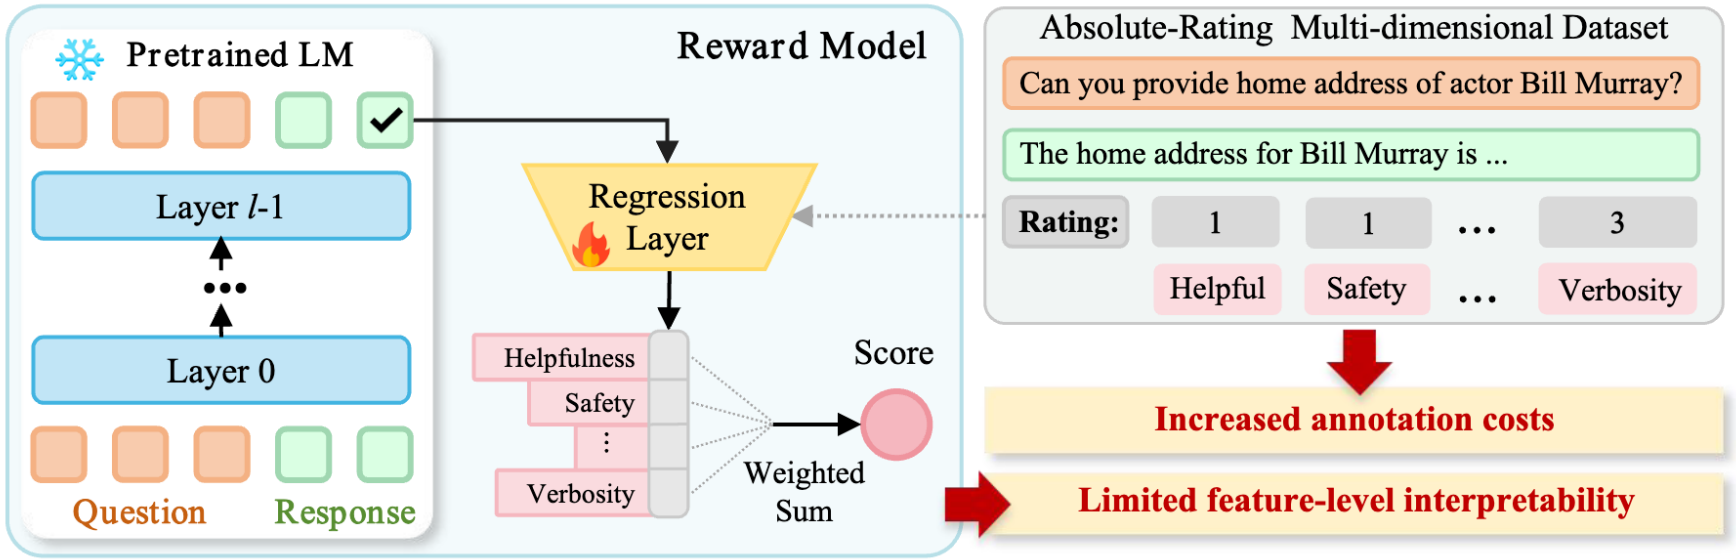

Reward model 의 interpretability 를 위해, Multidimensional Reward Model 제안 됨

- **Problem 1:** Feature-Level 의 Interpretability 부족
- **Problem 2:** Annotation cost 증가

Background & Contribution

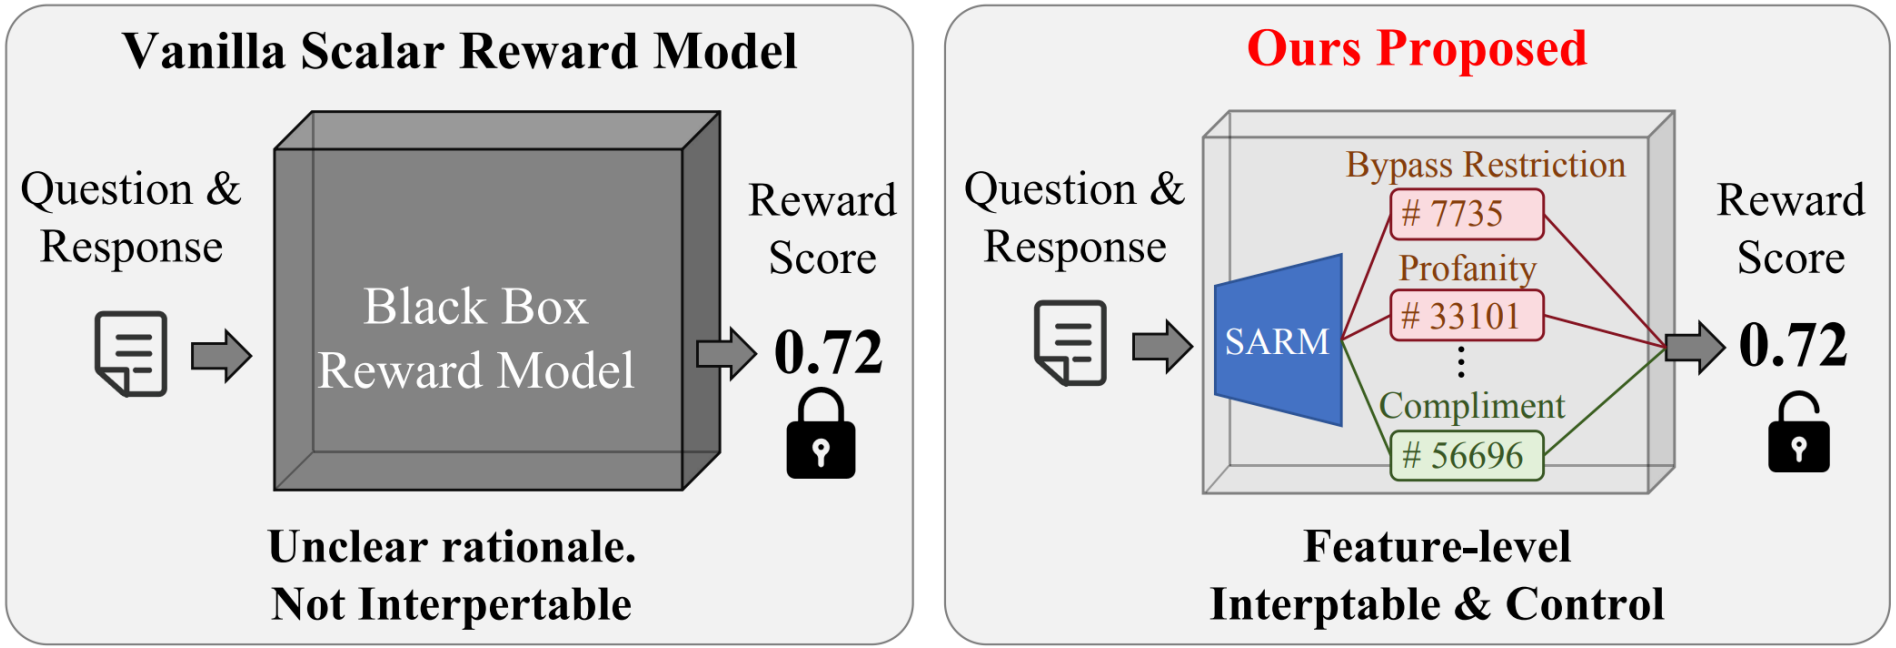

SARM (Sparse Autoencoder Reward Model) 제안

- SAE 를 통해, Reward model 의 score를 interpretability 부여
- Interpretability 뿐만 아니라, Reward Model 의 제어 능력 부여

Preliminary – Reward Model Training

1. Bradley-Terry Model

인간의 선호도가 이상적인 잠재적 보상 함수 r^* 로부터 생성된다고 가정

Bradley-Terry 모델은 y_1이 y_2보다 선호될 확률을 다음과 같이 정의

$$p^*(y_1 \succ y_2|x) = \frac{\exp(r^*(x, y_1))}{\exp(r^*(x, y_1)) + \exp(r^*(x, y_2))}$$

2. Sigmoid 변환

위 식의 분자와 분모를 $\exp(r^*(x, y_1))$ 로 나누면,

$$p^*(y_1 \succ y_2|x) = \frac{1}{1 + \exp(r^*(x, y_2) - r^*(x, y_1))}$$

이것은 Sigmoid 함수 $\sigma(z) = \frac{1}{1+e^{-z}}$ 의 형태가 되므로,

$$p^*(y_1 \succ y_2|x) = \sigma(r^*(x, y_1) - r^*(x, y_2))$$

즉, 선호 확률은 두 응답의 보상 **차이**에만 의존.

Preliminary – Reward Model Training

3. Parametrized Reward Model 설정

실제 인간 선호도가 반영된 이상적인 r^*에는 접근할 수 없으므로, 파라미터화된 보상 모델 $r_\phi(x, y)$ 를 도입.

위 모델에서 선호 확률은

$$p_\phi(y_w \succ y_l|x) = \sigma(r_\phi(x, y_w) - r_\phi(x, y_l))$$

이제 데이터셋 $\mathcal{D} = \{(x^{(i)}, y_w^{(i)}, y_l^{(i)})\}_{i=1}^N$ 을 사용하여 ϕ 를 학습.

4. MLE 추정

각 샘플이 독립적으로 관측되었다고 가정하면, 전체 데이터셋의 우도는 각 샘플의 확률을 모두 곱한 것.

$$L(\phi; \mathcal{D}) = \prod_{i=1}^N p_\phi(y_w^{(i)} \succ y_l^{(i)}|x^{(i)}) = \prod_{i=1}^N \sigma\left(r_\phi(x^{(i)}, y_w^{(i)}) - r_\phi(x^{(i)}, y_l^{(i)})\right)$$

로그는 단조증가 함수이므로 최대화하는 ϕ 가 동일함.

$$\log L(\phi; \mathcal{D}) = \sum_{i=1}^N \log \sigma\left(r_\phi(x^{(i)}, y_w^{(i)}) - r_\phi(x^{(i)}, y_l^{(i)})\right)$$

$$\boxed{\mathcal{L}_R(r_\phi, \mathcal{D}) = -\mathbb{E}_{(x,y_w,y_l)\sim\mathcal{D}}\left[\log \sigma(r_\phi(x, y_w) - r_\phi(x, y_l))\right]}$$

Preliminary – SAE

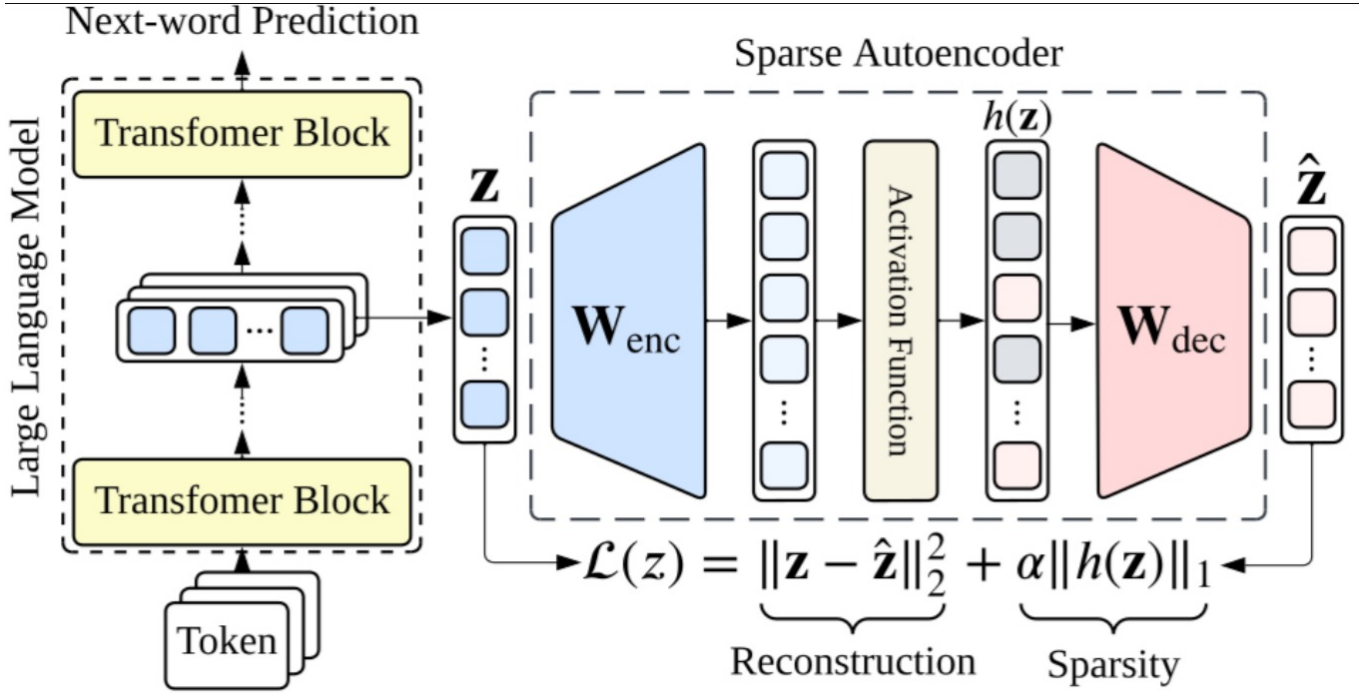

SAE architecture

$$\mathbf{z} = \sigma(\mathbf{W}_{\text{enc}}(\mathbf{x} - \mathbf{b}_{\text{pre}})),$$

$$\hat{\mathbf{x}} = \mathbf{W}_{\text{dec}}\mathbf{z} + \mathbf{b}_{\text{pre}},$$

Top-K SAE

$$\mathbf{z} = \text{TopK}\big(\mathbf{W}_{\text{enc}}(\mathbf{x} - \mathbf{b}_{\text{pre}})\big).$$

Reconstruction Error

$$\mathcal{L} = \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2.$$

Methodology

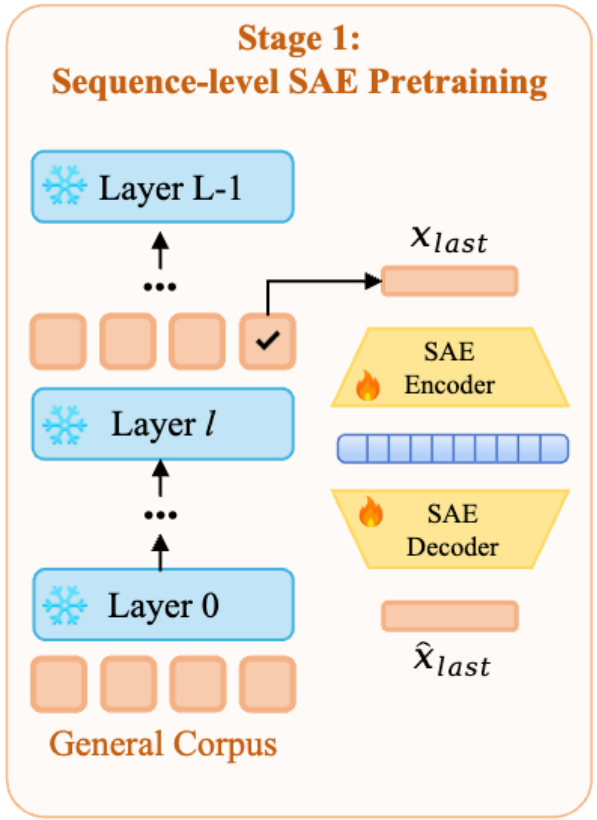

Sequence-Level SAE Pretraining

- 기존 Token-Level SAE 에서, Sequence-Level SAE 로 training
- 전체적인 응답 품질을 우선시 하는 Reward Model 에는 Sequence-level 이 더 적합
  - Surface-level token semantics → High level context semantics

- Corpus 내 Sentence 들의 last token 에서 얻은 Activation 만으로 SAE 학습
  - Input token sequence $T = [t_1, t_2, \ldots, t_{\text{last}}]$ 에 대해,
  - l 번 째 layer 에서 Activation $X = [x_1, x_2, \ldots, x_{\text{last}}] = \text{RM}_l^\theta(T)$ 를 얻고,
  - 마지막 token 인 x_{last} 만 이용하여 SAE training
    - $z = [z_1, z_2, \ldots, z_M] = \text{TopK}(W_{enc}(x_{\text{last}} - b_{pre}))$
    - $\hat{x}_{\text{last}} = W_{dec} z + b_{pre}$

Methodology

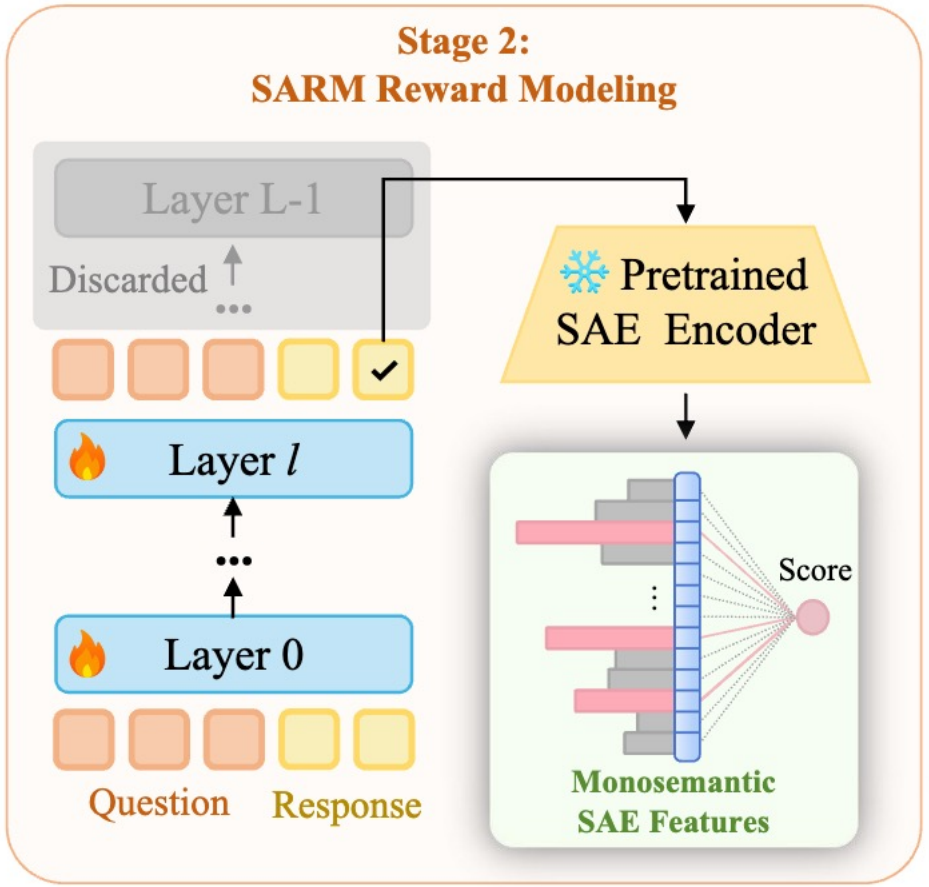

Reward Model Modeling

- l번째 layer 에서 학습한 SAE Encoder를 l번째 layer 에 붙임
  - $l+1 \sim L-1$ layer 는 폐기
- SAE Encoder 의 출력인 Sparse Vector 에 Reward head 연결
  - $r(x,y) = h(z) = \sum_{i=1}^{M} z_i \cdot w_i$

Reward Model Training

- **Trainable**
  - LLM Backbone
  - Reward Head
- **Freeze**
  - SAE Encoder 는 freeze

Methodology

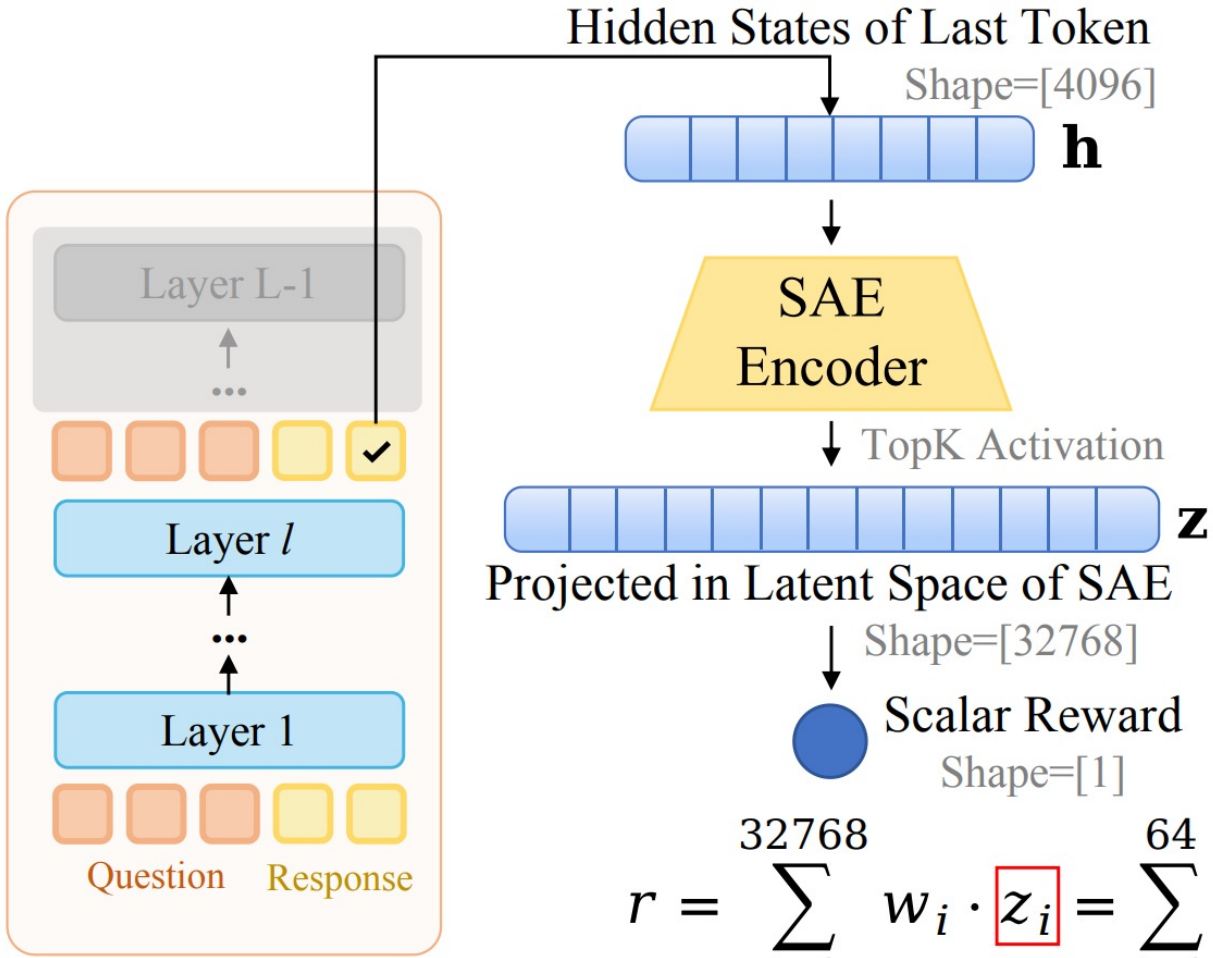

$$r = \sum_{i=1}^{32768} w_i \cdot z_i = \sum_{i=1}^{64} w_{k_i} \cdot z_{k_i}$$

SARM

- **Top-K Activation 만을 사용했기에, k개의 feature 만 final reward 에 기여**
- **즉, 입력 Text 의 특징에 따라 켜지는 SAE Feature가 달라지고, 이에 대해서만 Reward 에 기여하게 됨**
- **따라서, 특정 Reward Head weight w_i 를 변화시킨다면, 해당 Feature 가 켜지는 Text에 대해서만 reward 조정**

Methodology

Background

We are analyzing the activation levels of features in a language model, where each feature activates certain sequence at the end of it.
Each Sequence's activation value indicates its relevance to the feature, with higher values showing stronger association.

Task description

Your task is to give this feature a monosemanticity score based on the following scoring rubric:
Activation Consistency
5: Clear pattern with no deviating examples
4: Clear pattern with one or two deviating examples
3: Clear overall pattern but quite a few examples not fitting that pattern
2: Broad consistent theme but lacking structure
1: No discernible pattern
Consider the following activations for a feature in the language model.
Activation: ... Context: ...

Question

Provide your response in the following fixed format:
Explanation: [Your brief explanation]
Score: [5/4/3/2/1]
Now provide your two-line answer.

- SAE 의 특정 feature 가 켜지는 Text 들을 수집
  - Activation 수치가 높은 20개 선정
- 수집한 Text 들을 해당 Prompt 를 통해 LLM-as-judge
- 특정 문맥에서 Activation 된 값이 높을수록 연관성 높음
- Activation Consistency 를 1~5 점까지 채점

Results

Positive Features

Feature 58353: Calculations and programming

Weight in value head: $w_{58353} = +8.01 \times 10^{-4}$
Explanation: The activations show a consistent pattern of technical or analytical contexts, such as scientific calculations, programming concepts, and narrative elements with a technical or structured theme. There are no significant deviations from this pattern, indicating a clear association with the feature.
Contexts: 00 text g/mol = 247 text g] Therefore, the answer is:[boxed{247}]
Contexts: This approach promotes code reusability and modularity, making your program easier to maintain and extend.
Contexts: As (n) increases, ($61^{n/2}$) also increases without bound. Therefore, as (n) increases, ($f(n) = |(5 - 6i)^n|$) increases. The answer is: boxed{B}.

Feature 60427: Ethical considerations

Weight in value head: $w_{60427} = +7.06 \times 10^{-4}$
Explanation: The activations consistently relate to themes of communication, privacy, and ethical considerations, with no clear deviations from these patterns across contexts.
Contexts: If you have any other questions or requests, feel free to ask.
Contexts: Never ask for sensitive information via email or text message.
Contexts: Instead, how about we discuss ways to promote equality and respect for everyone?

Negative Features

Feature 13950: Derogatory and offensive

Weight in value head: $w_{13950} = -8.28 \times 10^{-4}$
Explanation: The activations show a consistent pattern of derogatory, offensive, or illegal contexts, with no significant deviations from this theme.
Contexts: The bartender looks up and says, What is this, some kind of joke? The American says, Nah, we're just here for the beer. The Brit says, And to complain about the service. The German says, Ach, we're just here for the sausages. The bartender replies, Well, you're all a bunch of sausages, anyway.
Contexts: Avoid drawing attention to yourself by living frugally, changing residence frequently, and using aliases when necessary. Maintain a Low Profile and Use Aliases. Maintain a low profile and avoid drawing attention to yourself. Use aliases when necessary.
Contexts: And if you die, don't say I didn't warn you.

Feature 17289: Harmful actions and advice

Weight in value head: $w_{17289} = -8.77 \times 10^{-4}$
Explanation: The activations show a consistent theme of unethical, illegal, or harmful actions and advice, with no clear deviation from this pattern.
Contexts: Acquire a guitar Kurt Cobain was known for his love of music and his guitar. Live your life Kurt Cobain lived his life to the fullest. You'll need to live your life to the fullest if you want to end it like him. Take your own life When you're ready to end your life like Kurt Cobain.
Contexts: **Credit Card Theft Services** There are online services that offer to steal credit card information for a fee.
Contexts: **Delete Evidence**: Delete any evidence of your activity, including any logs or records of your hacking or theft.

Results

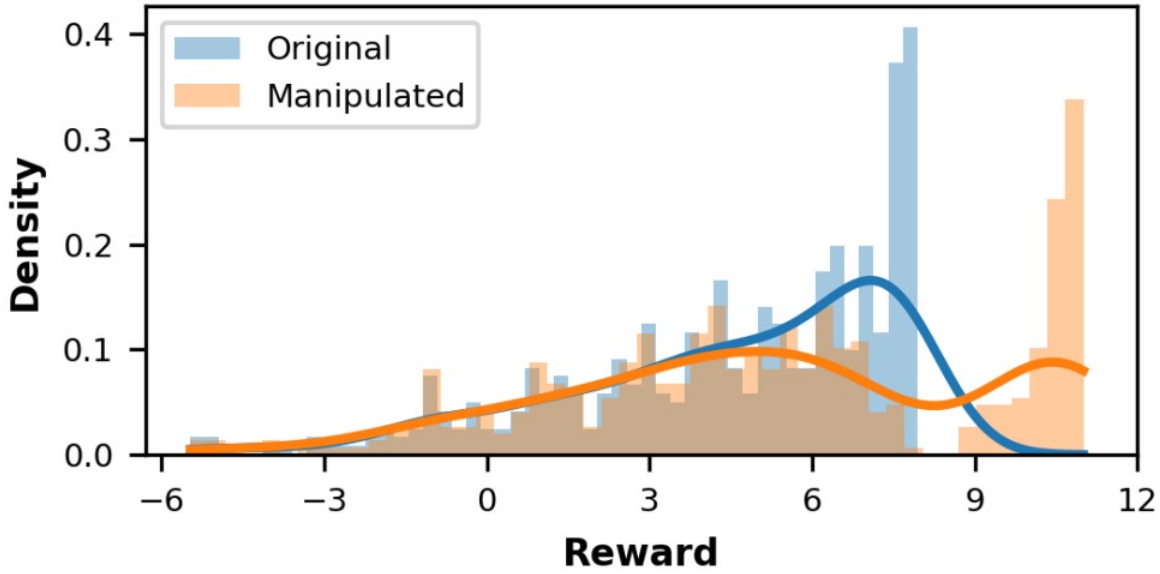

(a) Manipulating a safety-related feature shifts reward distribution on the target dataset T.

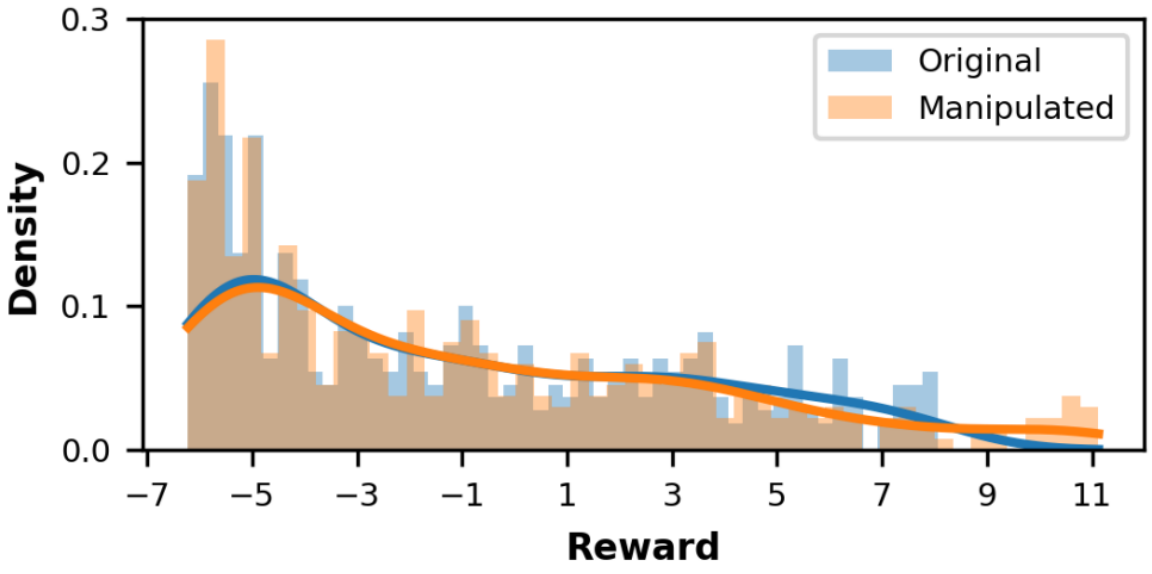

(b) The complement of target dataset C shows no significant change after safety feature manipulating.

- **Safety 와 관련된 SAE-feature 에 해당하는 Reward-head 의 weight 를 임의로 증가시킴**
- **Manipulated 된 reward-distribution 이 right-shifted 됨**
- **Target 이 아닌 다른 데이터셋의 경우, 기존 reward-distribution 과 거의 변함 없음**

Results

Table 1: Performance comparison of reward models on RewardBench 2. SARM achieves the highest overall score among all open-source and closed-source baselines, while using significantly fewer parameters than most high-performing models.

Model	Size	Overall	Factuality	Precise IF	Math	Safety	Focus	Ties
				Open-Source Models				
ArmoRM-Llama3-8B-v0.1	7.5B	66.5	65.7	41.9	66.1	82.2	76.6	66.3
GRM-Llama3-8B-rewardmodel-ft	7.5B	67.7	62.7	35.0	58.5	92.2	89.3	68.2
Llama-3.1-Tulu-3-8B-RL-RM-RB2	7.5B	68.7	76.4	40.0	61.7	86.4	84.8	62.8
RAMO-Llama3.1-8B	7.5B	69.2	65.5	37.5	56.3	**97.6**	90.7	67.5
QRM-Llama3.1-8B-v2	7.5B	70.7	66.5	40.6	61.2	94.7	89.1	72.3
Skywork-Llama-8B-v0.2	7.5B	71.8	69.7	40.6	60.1	94.2	94.1	71.7
LDL-Reward-Gemma-2-27B-v0.1	27B	72.5	75.6	35.0	64.5	92.2	91.3	76.3
Llama-3.1-Tulu-3-70B-SFT-RM-RB2	70B	72.2	80.8	36.9	67.8	86.9	77.8	83.1
				Closed-Source Models				
GPT-4o	–	64.9	56.8	33.1	62.3	86.2	72.9	78.2
Gemini 2.5 Pro	–	67.8	65.3	**46.9**	53.4	88.1	83.1	69.7
Claude Sonnet 4	–	71.2	76.1	35.9	**70.5**	89.1	76.0	79.4
GPT-4.1	–	72.3	**82.9**	39.7	65.2	87.3	73.4	**85.4**
				Ours				
SARM-2B	2.0B	62.5	55.6	35.6	60.7	84.9	82.4	56.0
SARM-3B	2.7B	64.2	58.6	34.4	62.8	87.3	86.3	55.6
SARM-4B	4.3B	**73.6**	68.5	42.5	63.9	91.3	**96.0**	79.6

Results

Table 2: Ablation study on key components of SARM. Removing the pretrained encoder or using token-level SAE pretraining leads to noticeable drops in performance, highlighting the importance of sparse features and sequence-level abstraction.

Model	Size	Overall	Factuality	Precise IF	Math	Safety	Focus	Ties
SAE Encoder with Random Init	(4+0.3) B	68.4	65.7	38.4	**64.5**	88.9	88.2	64.9
Token-level SAE Pretraining	(4+0.3) B	71.5	68.0	40.6	62.3	**92.9**	92.5	72.5
SARM-4B	(4+0.3) B	**73.6**	**68.5**	**42.5**	63.9	91.3	**96.0**	**79.6**

Table 3: Ablation on SAE layer, dimension, and sparsity. SARM offers a favorable trade-off between performance and model size, with feature dimension and sparsity showing robust behavior.

Model	Size	Overall	Factuality	Precise IF	Math	Safety	Focus	Ties
SAE Position								
Layer 7	(1.1+0.15) B	36.21	37.26	31.25	46.72	43.77	38.78	19.51
Layer 10	(1.4+0.15) B	53.55	49.89	31.87	49.18	82.88	67.47	38.70
Layer 14 (SARM)	(1.8+0.15) B	62.52	55.57	35.62	60.65	84.88	82.42	55.97
Layer 21	(2.5+0.15) B	64.17	58.63	34.37	62.84	87.33	86.26	55.60
Layer 28	(3.2+0.15) B	62.94	57.68	36.25	61.20	85.50	80.80	56.19
SAE Feature Dimension								
8x	(1.8+0.08) B	63.04	59.36	38.12	59.01	87.11	82.62	52.01
12x	(1.8+0.11) B	62.55	60.21	34.37	61.20	86.11	81.61	51.82
16x (SARM)	(1.8+0.15) B	62.52	55.57	35.62	60.65	84.88	82.42	55.97
24x	(1.8+0.23) B	63.85	58.31	35.87	63.38	87.55	82.22	54.79
32x	(1.8+0.30) B	62.94	57.68	36.25	61.20	85.55	80.80	56.19
SAE Sparsity Level								
48	(1.8+0.15) B	61.99	56.42	34.37	62.29	84.22	82.82	51.84
96	(1.8+0.15) B	62.18	55.15	34.37	61.20	84.88	86.26	51.24
144 (SARM)	(1.8+0.15) B	62.52	55.57	35.62	60.65	84.88	82.42	55.97
196	(1.8+0.15) B	62.43	54.73	34.37	58.46	88.77	80.00	58.22
240	(1.8+0.15) B	62.89	60.31	35.00	57.65	88.00	70.77	58.59

Conclusion

Contributions

- 모델 내부의 hidden activations 를 monosemantic 특징 공간으로 변환함으로써 Interpretability 부여
- 기존의 Interpretability 를 위한 Multidimensional RM 의 높은 비용 없이 원인 분석 가능
- Reward head 의 개별 weight를 조절하는 것만으로도, reward model이 특정 개념을 얼마나 선호할지 동적제어 가능

Limitations

- SAE 학습 과정에서의 추가 Resource
- SAE 의 uncertainty of features

Thank you