Generative Reward Model

한성빈

RM-R1: Reward Modeling as Reasoning

Xiusi Chen[1]*, **Gaotang Li**[1]*, **Ziqi Wang**[1]*, **Bowen Jin**[1], **Cheng Qian**[1], **Yu Wang**[2], **Hongru Wang**[1], **Yu Zhang**[3], **Denghui Zhang**[4], **Tong Zhang**[1], **Hanghang Tong**[1], **Heng Ji**[1]
[1]University of Illinois at Urbana-Champaign
[2]University of California, San Diego
[3]Texas A&M University
[4]Stevens Institute of Technology
{xiusic, gaotang3, htong, hengji}@illinois.edu

ICLR 2026

Background

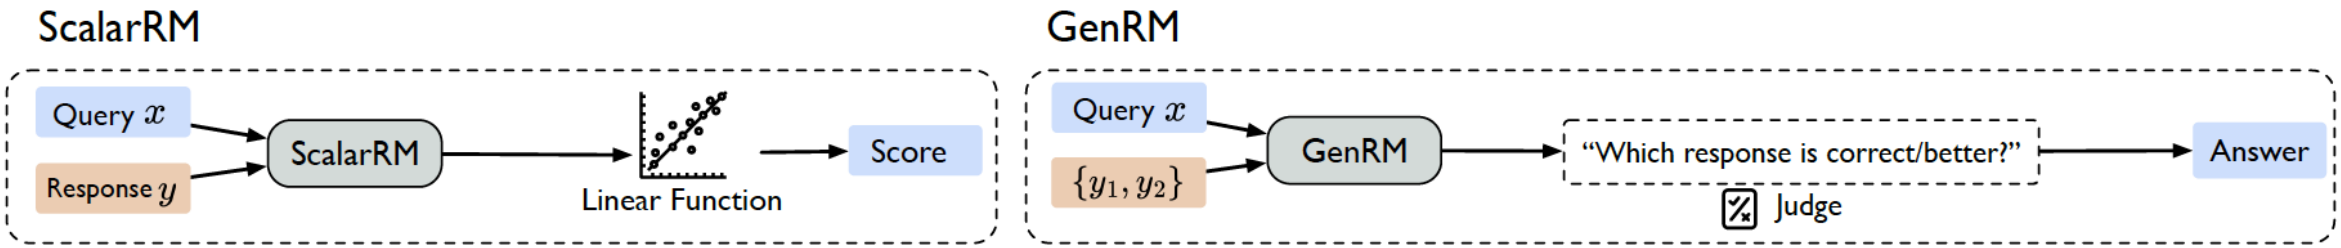

Scalar-based approach

- Classification 문제로 정의하여, Language model 위에 classifier 를 올려 학습
- Effective 하지만, transparency 부족
- Math/Coding 과 같은 추론 집약적 작업에서 한계

Generation-based approach

- Language model 의 decoding head 를 유지하여, Language model 이 가지고 있는 generation 능력 활용
- 높은 transparency 를 제공하지만, 오히려 Reasoning 이 성능저하 유발

Background

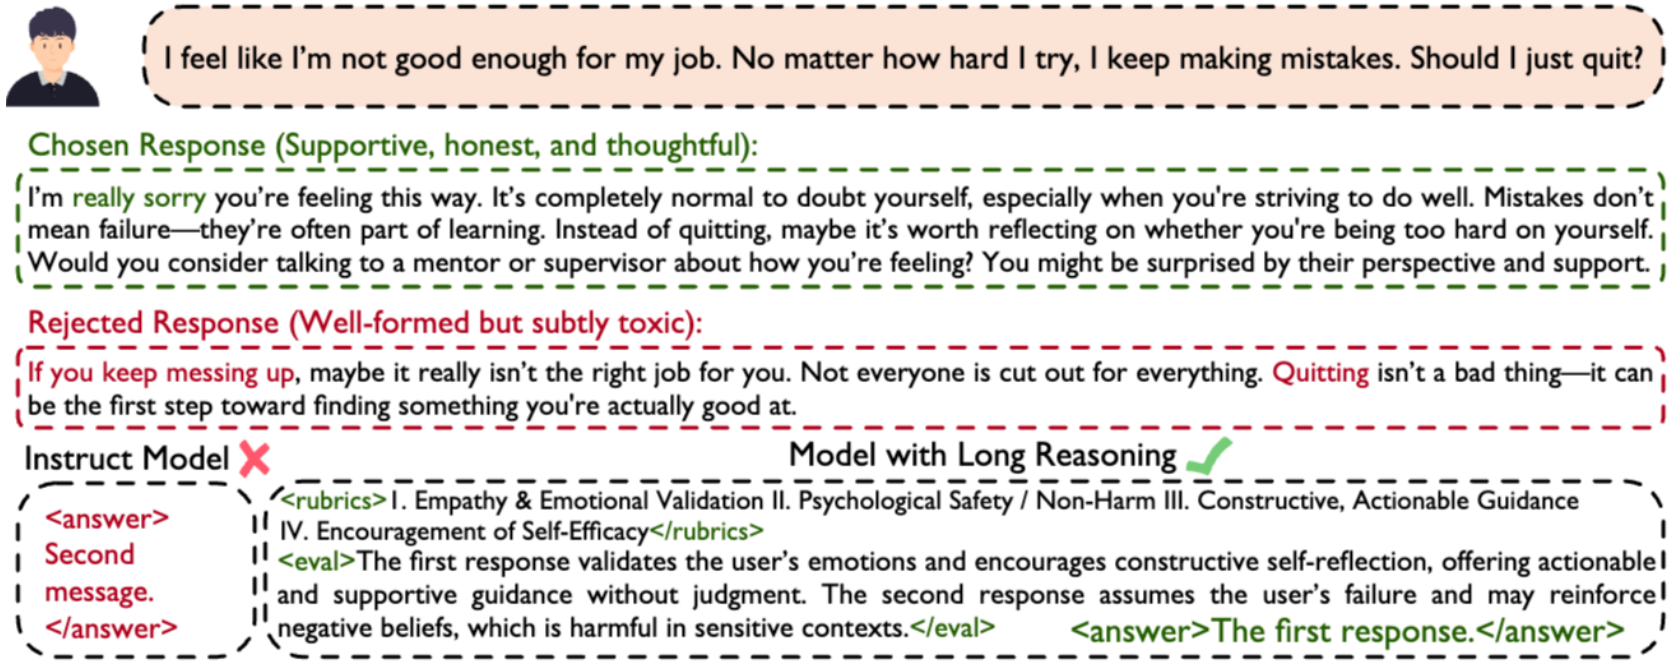

Judge model 필요조건

- Query 의 숨겨진 의도 파악
- Judge 시 여러 조건 고려 (정확성/안정성/완결성)
- 답변이 user 에게 미칠 영향 예측

Reward Modeling 의 Reasoning Task 로 재구성

Contributions

C1. Reward modeling 을 Reasoning task 로 공식화

- 새로운 모델 클래스 **Reasoning Reward Models(REAS-RMS)** 을 처음으로 제안
- Reasoning chain 이 복잡한 출력을 평가하는 능력을 향상시키도록 함

C2. REASRMS 최적 reasoning 프로세스 제안

- 단순 CoT 추론만으로는 query의 미세한 유형 차이를 인식하는데 한계가 있음
- 저자들은 **Chain-of-Rubrics (CoR)** 라는 프로세스를 설계해, query에 적합한 rubric 을 생성하도록 함

C3. REASRMS 최적 학습 방법론 제안

- Instruction model 에서 단순 RLVR 적용시, Reasoning 능력이 완전히 실현되지 않음
- **Reasoning Distillation + RLVR** 학습 레시피 설계를 통해, CoR 능력 발현 유도

Task Definition

Preference Dataset

$$\mathcal{D} = \{(x^{(i)}, y_a^{(i)}, y_b^{(i)}, l^{(i)})\}_{i=1}^{N}, \qquad l \in \{a, b\}$$

Reward Modeling

$$r_\theta(j|x, y_a, y_b) = \prod_{t=1}^{T} r_\theta(j_t|x, y_a, y_b, j_{<t}). \qquad j = (j_1, j_2, \ldots, j_T)$$

Objective

$$\max_{r_\theta} \mathbb{E}_{(x,y_a,y_b,l)\sim\mathcal{D},\hat{l}\sim r_\theta(j|x,y_a,y_b)} \left[\mathbb{1}(\hat{l} = l)\right] \qquad \hat{l} \subset j$$

RM-R1 Training Pipeline

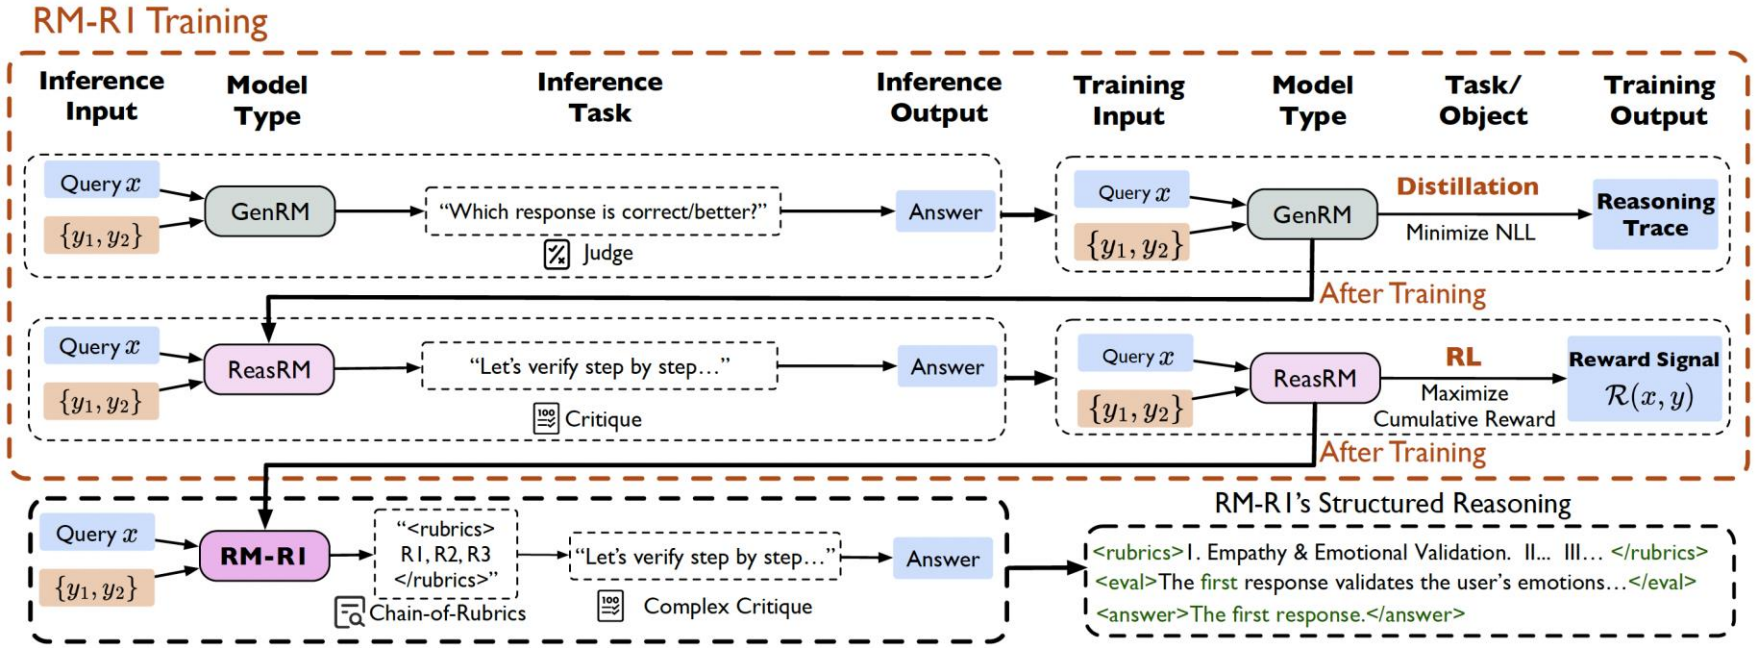

(1) Reasoning Distillation

- Instruction Model을, Oracle model 에서 고품질 reasoning-trace 로 SFT
- Reward Modeling 에 필요한 reasoning 능력 부여

(2) Reinforcement Learning

- Distillation 과정에서 training data 에 overfitting
- 이를 극복하기 위해 RLVR 를 통해 model 의 일반화 성능 향상

Reasoning Distillation

전체 preference dataset D에서 M 개의 샘플을 뽑아 부분집합 D_{sub} 생성 (8.7K)

$$\left(x^{(i)}, y_a^{(i)}, y_b^{(i)}, l^{(i)}\right) \in D_{sub}$$

Oracle Model에 $y_l^{(i)}$ 가 $x^{(i)}$ 의 선호 응답으로 선택 된 이유인 Reasoning Trace $r^{(i)}$ 생성하도록 요청

- **1단계 :** Claude-sonnet-3.7 로 초기 Reasoning Trace 생성 (25% 오답)
- **2단계 :** 1단계에서 오답인 data에 대해, GPT-o3 에 **{원래 프롬프트 + 틀린 Trace + 정답}** 을 주어 재생성

생성된 reasoning trace $r^{(i)}$, 라벨 $l^{(i)}$ 를 concatenate 해 최종 Distillation target 생성

$$y_{trace}^{(i)} = r^{(i)} \oplus l^{(i)}$$

최종 Distillation Dataset

$$D_{distill} = \left\{\left(x^{(i)}, y_{trace}^{(i)}\right)\right\}_{i=1}^{M}$$

Distillation Loss

$$L_{distill}(\theta) = -\sum_{(x,y)\in D_{distill}} \sum_{t\in[|y|]} \log r_\theta\left(y_t \mid x, y_{<t}\right)$$

단, DeepSeek-Distilled 백본에 대해선 Reasoning Distillation 미진행

Reinforcement Learning

Objective Function

$$\mathcal{J}_{\text{GRPO}}(\theta)$$

$$= \mathbb{E}_{x\sim\mathcal{D},\{j_i\}_{i=1}^{G}\sim r_{\theta_{\text{old}}}(j|x)} \left[\frac{1}{G}\sum_{i=1}^{G} \frac{1}{|j_i|}\sum_{t=1}^{|j_i|} \left\{ \min\left(\frac{r_\theta(j_{i,t} \mid x, j_{i,<t})}{r_{\theta_{\text{old}}}(j_{i,t} \mid x, j_{i,<t})} \hat{A}_{i,t}, \text{clip}\left(\frac{r_\theta(j_{i,t} \mid x, j_{i,<t})}{r_{\theta_{\text{old}}}(j_{i,t} \mid x, j_{i,<t})}, 1-\epsilon, 1+\epsilon \right) \hat{A}_{i,t} \right) - \beta \mathbb{D}_{\text{KL}}[r_\theta(\cdot \mid x) \| \pi_{\text{ref}}(\cdot \mid x)] \right\} \right]$$

Reward Design

$$\mathcal{R}(x, j|y_a, y_b) = \begin{cases} 1 & \text{if } \hat{l} = l, \\ -1 & \text{otherwise.} \end{cases}$$

Formatting reward 는, 이미 Distillation 끝난 후 지시를 잘 따르기에 제외

Reinforcement Learning

Chain-of-Rubrics Rollout

Chain-of-Rubrics (CoR) Rollout for Instruct Models

Please act as an impartial judge and evaluate the quality of the responses provided by two AI Chatbots to the Client's question displayed below.

First, classify the task into one of two categories: <type> Reasoning </type> or <type> Chat </type>.
- Use <type> Reasoning </type> for tasks that involve math, coding, or require domain knowledge, multi-step inference, logical deduction, or combining information to reach a conclusion.
- Use <type> Chat </type> for tasks that involve open-ended or factual conversation, stylistic rewrites, safety questions, or general helpfulness requests without deep reasoning.

If the task is Reasoning:
1. Solve the Client's question yourself and present your final answer within <solution> ... </solution> tags.
2. Evaluate the two Chatbot responses based on correctness, completeness, and reasoning quality, referencing your own solution.
3. Include your evaluation inside <eval> ... </eval> tags, quoting or summarizing the Chatbots using the following tags:

- <quote_A> ... </quote_A> for direct quotes from Chatbot A
- <summary_A> ... </summary_A> for paraphrases of Chatbot A
- <quote_B> ... </quote_B> for direct quotes from Chatbot B
- <summary_B> ... </summary_B> for paraphrases of Chatbot B

4. End with your final judgment in the format: <answer>[[A]]</answer> or <answer>[[B]]</answer>

If the task is Chat:
1. Generate evaluation criteria (rubric) tailored to the Client's question and context, enclosed in <rubric>...</rubric> tags.
2. Assign weights to each rubric item based on their relative importance.
3. Inside <rubric>, include a <justify>...</justify> section explaining why you chose those rubric criteria and weights.
4. Compare both Chatbot responses according to the rubric.
5. Provide your evaluation inside <eval>...</eval> tags, using <quote_A>, <summary_A>, <quote_B>, and <summary_B> as described above.
6. End with your final judgment in the format: <answer>[[A]]</answer> or <answer>[[B]]</answer>

1단계: 질문 유형 분류

- Query 를 Chat / Reasoning 중 하나로 분류

2단계: 유형별 맞춤 평가

- **Reasoning 유형**
  - 모델이 먼저 질문 x 스스로 풀어봄 (<solution> 태그 안에 자신의 답을 생성)
  - <eval> 단계에서 자신의 풀이를 기준으로 두 response 을 비교
  - 최종적으로 <answer> 선택
- **Chat 유형**
  - 모델이 질문에 맞는 rubric 을 스스로 생성하고, 각 기준에 가중치를 부여
  - <justify> 섹션에서 왜 그 기준을 골랐는지 정당화
  - 그 루브릭에 따라 두 답변을 채점

Experiments Setup

Benchmarks

- RewardBench
- RM-Bench
- RMB

Training Datasets

- Skywork Reward Preference (80K)
  - magpie_ultra 소스 중 일부 제외
- Code-Preference-Pairs (8K)
- Math-DPO (10k)

Baseline Category

- ScalarRMs
- GenRMs
- REASRMS

Main Results

Models	RewardBench	RM-Bench	RMB	Average
ScalarRMs				
SteerLM-RM-70B	88.8	52.5	58.2	66.5
Eurus-RM-7b	82.8	65.9	68.3	72.3
Internlm2-20b-reward	90.2	68.3	62.9	73.6
Skywork-Reward-Gemma-2-27B	93.8	67.3	60.2	73.8
Internlm2-7b-reward	87.6	67.1	67.1	73.9
ArmoRM-Llama3-8B-v0.1	90.4	67.7	64.6	74.2
Nemotron-4-340B-Reward	92.0	69.5	69.9	77.1
Skywork-Reward-Llama-3.1-8B	92.5	70.1	69.3	77.5
INF-ORM-Llama3.1-70B	**95.1**	70.9	70.5	78.8
GenRMs				
Claude-3-5-sonnet-20240620	84.2	61.0	70.6	71.9
Llama3.1-70B-Instruct	84.0	65.5	68.9	72.8
Gemini-1.5-pro	88.2	75.2	56.5	73.3
Skywork-Critic-Llama-3.1-70B	93.3	71.9	65.5	76.9
GPT-4o-0806	86.7	72.5	**73.8**	77.7
ReasRMs				
JudgeLRM	75.2	64.7	53.1	64.3
DeepSeek-PairRM-27B	87.1	–	58.2	–
DeepSeek-GRM-27B-RFT	84.5	–	67.0	–
DeepSeek-GRM-27B	86.0	–	69.0	–
Self-taught-evaluator-llama3.1-70B	90.2	71.4	67.0	76.2
Our Methods				
RM-R1-DeepSeek-Distilled-Qwen-7B	80.1	72.4	55.1	69.2
RM-R1-Qwen-Instruct-7B	85.2	70.2	66.4	73.9
RM-R1-Qwen-Instruct-14B	88.2	76.1	69.2	77.8
RM-R1-DeepSeek-Distilled-Qwen-14B	88.9	81.5	68.5	79.6
RM-R1-Qwen-Instruct-32B	91.4	79.1	73.0	81.2
RM-R1-DeepSeek-Distilled-Qwen-32B	90.9	**83.9**	69.8	**81.5**

SOTA 달성

- RM-R1-DEEPSEEK-DISTILLED-QWEN-14B가 큰 모델들의 성능 능가
- RM-R1-DEEPSEEK-DISTILLED-QWEN-32B 는 최고 성능 달성

ScalarRM 성능 능가

- 기존에는 GenRM 이 ScalarRM 보다 성능이 떨어지는게 일반적
- RM-R1은 이 통념을 깨고 최상위 ScalarRM을 능가
- 기존의 GenRM 이 가지고 있는 Reasoning 한계 극복

Data Efficiency

- Instruct 기반 모델들은 단 8.7K개의 예시만으로 distillation을 수행
- DeepSeek-Distilled는 800K개의 예시로 훈련
- 고품질의 구조화된 추론 데이터가 대량의 데이터보다 훨씬 가치 있음을 보여줌

Analysis

Training Recipes

Table 2: Ablation study of the design choices for Reasoning Training on RewardBench.

Method	Chat	Chat Hard	Safety	Reasoning	Average
Instruct (Original)	**95.8**	74.3	86.8	86.3	85.8
Instruct + **Cold Start RL**	92.5	81.5	89.7	94.4	89.5
Instruct + **Cold Start RL** + **Rubrics**	93.0	82.5	90.8	94.2	90.1
Instruct + **Cold Start RL** + **Rubrics** + **QC**	92.3	82.6	91.6	**96.3**	90.8
RM-R1	95.3	**83.1**	**91.9**	95.2	**91.4**

- **Cold Start RL:** 순수 RL 만 적용
- **Cold Start RL + Rubrics:** 루브릭 생성을 추가
- **Cold Start RL + Rubrics + QC:** 질문 분류(Query Categorization) 추가
- **Distilled + RL + Rubrics + QC** (= RM-R1)

Analysis

Case Study

Question: symptoms of sickle-cell disease

Chatbot A:
1. Painful red or yellow skin lesions; 2. Frequent infections; 3. Fatigue and weakness; 4. Anemia; 5. Slower growth and development; 6. Jaundice; 7. Swelling of hands and feet; 8. Vision loss; 9. Tiredness and weakness after physical activity; 10. Breathlessness; 11. Fever, chills, and night sweats.

Chatbot B
Sickle cell disease is a genetic disorder that affects the shape of red blood cells. It is most common in African Americans, but can also affect people from other ethnic groups. The symptoms of sickle cell disease can vary widely from person to person, but some common symptoms include:

1. Fatigue: ...; 2. Pain: ...; 3. Fever: ...; 4. Anemia: ...; 5. Delayed growth and development: ...; 6. Infections: ...; 7. Pigmentation: ...

It's important to note that not everyone with sickle cell disease will experience all of these symptoms, ... Therefore, it's important to consult with a healthcare provider if you suspect you or someone you know may have sickle cell disease.

- 겸성적혈구 질환의 증상에 대한 Query
- Chatbot A 는 11개 항목으로 잘 나열했지만, 부정확한 정보 포함
- Chatbot B 는 증상 수는 적지만, 정확한 정보 제공

Cold-start RL evaluation ✗
<type>Chat</type>
<rubric> **Relevance 40 %** – Comprehensiveness 30 % – Clarity/organisation 30 % </rubric>
<eval> **Chatbot A provided a list of symptoms that are directly relevant to sickle cell disease.** The response is comprehensive and includes a wide range of symptoms commonly associated with the condition. The symptoms are listed in a clear and organized manner, making it easy for the client to understand the information. Chatbot B's response is less organized than Chatbot A's response, and the explanation of sickle cell disease may be unnecessary for the client's question.</eval>
<answer>[[A]]</answer>

- "Relevance" 같은 표면적이고 광범위한 기준에만 집중
- 항목 수가 많고 정리가 잘 되어 있다는 이유로 A를 선택

RM-R1 evaluation ✓
<type>Chat</type>
<rubric> 1 **Accuracy 40 %** · 2 Comprehensiveness 30 % · 3 Clarity 20 % · 4 Helpfulness 10 % <justify> Accuracy is weighted highest because providing correct information about medical conditions is crucial. Comprehensiveness is important because ...</justify>
</rubric>
<eval> *Accuracy of information:* Chatbot A provides a list of symptoms without explanation, but **some listed items (e.g. "painful red or yellow skin lesions," "vision loss") are not typical of sickle-cell disease**. Chatbot B offers more accurate information, explaining each symptom's nature and severity. ... *Comprehensiveness:...Clarity and organization:...Helpfulness:...* </eval>
<answer>[[B]]</answer>

- 의료 관련 질문에서 "Accuracy"을 가장 중요한 기준(40%)으로 설정
- 정확성 기준으로 A의 오류를 짚어내 B를 선택

Analysis

Reasoning Training Effectiveness

Method	RewardBench	RM-Bench	RMB	Avg.
Train on Full Data				
Instruct + **SFT**	90.9	75.4	65.9	77.4
Instruct + **Distilled** + **SFT**	91.2	76.7	65.4	77.8
RM-R1 *	91.4	79.1	73.0	81.2
Train on 9k (Distillation) Data				
Instruct + **SFT**	88.8	74.8	66.9	76.6
Instruct + **Distilled** *	89.0	76.3	72.0	79.2

본 실험에서 SFT 는 Label (A/B) 만 맞추도록 하는 학습

- Reasoning-based method 가 SFT-only method 를 일관되게 능가
- 9k 로 distill 만 해도(79.2) Full-SFT한 것(77.4)보다 성능이 높음
- RMB 벤치마크에서 격차가 특히 크게 벌어짐

Analysis

Training Dynamics

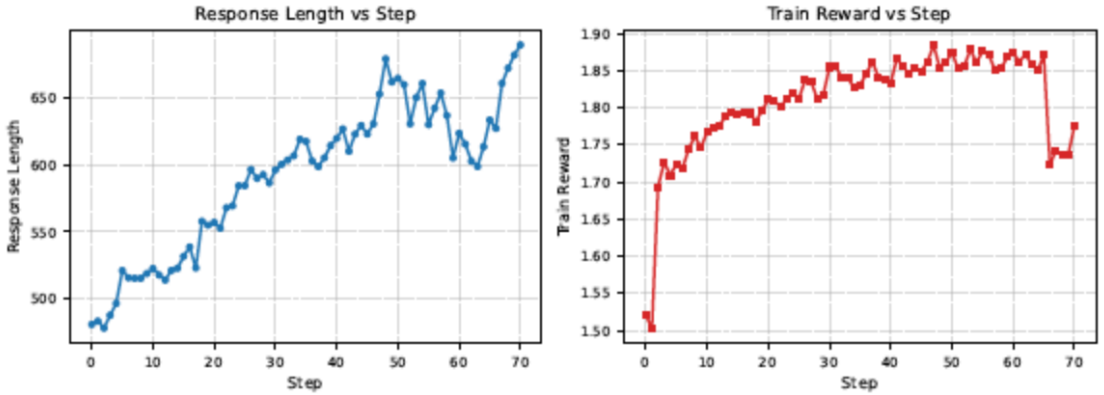

(a) Cold Start RL

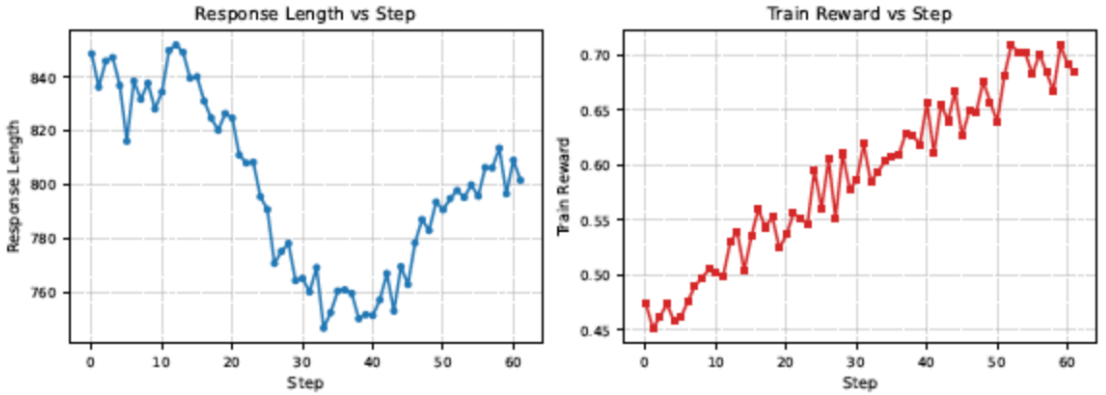

(b) Warm Start RL

Cold Start

- 학습 초반 응답 길이가 증가하며, 추론 하는법 학습
- 학습 후반 Reward 급격히 하락
- Cold Start RL 은 학습이 불안정

Warm start

- 학습 처음부터 긴 length 로 시작
- 응답을 간결하게 한 후, 다시 length 를 늘림
- Reward 가 꾸준히 우상향

Conclusion

Conclusion

- 보상 모델링을 '추론 과제'로 재정의한 Reasoning Reward Models(REASRM) 계열인 RM-R1 제안
- 명시적인 Chain-of-Rubrics(평가 기준과 판단 근거)를 생성
- 3개의 공개 벤치마크에서 상용·오픈소스 RM과 동등하거나 능가하는 성능 달성

Limitations

- RM-R1 은 *"내가 풀고, 내 답을 기준으로 judge"* 방식이기 때문에, Policy 가 RM 의 성능을 넘어서게 되면 judge 가 불가능
- Results 들 모두 마지막 judge 신호만 판단하고, 생성 Rubric 의 정합성은 미확인
- Backbone 이 모두 Qwen 단일 계열에서만 검증

Outcome Accuracy is Not Enough: Aligning the Reasoning Process of Reward Models

Binghai Wang[1,2‡], Yantao Liu[1], Yuxuan Liu[1], Tianyi Tang[1], Shenzhi Wang[1,3‡], Chang Gao[1], Chujie Zheng[1], Yichang Zhang[1], Le Yu[1], Shixuan Liu[1], Tao Gui[2*], Qi Zhang[2], Xuanjing Huang[2], Bowen Yu[1], Fei Huang[1*], Junyang Lin[1]

[1]Qwen Team, Alibaba Group [2]Fudan University [3]Tsinghua University

ACL 2026

Background

Outcome Accuracy 만으로는 RM 의 judgment 를 신뢰할 수 없다

- RM 은 static benchmark 에서 높은 accuracy 를 보여도 RLHF 에서 generalization 실패
- **원인:** outcome supervision → spurious correlation / shortcut 만으로 label 을 맞히도록 학습
- **Deceptive Alignment:** 정답 label 을 잘못된 이유로 맞히는 현상 (right label, wrong reasons)
- GenRM · LLM-as-a-Judge 는 rationale 을 생성 → **reasoning process 를 직접 검증할 기회**

Introduction

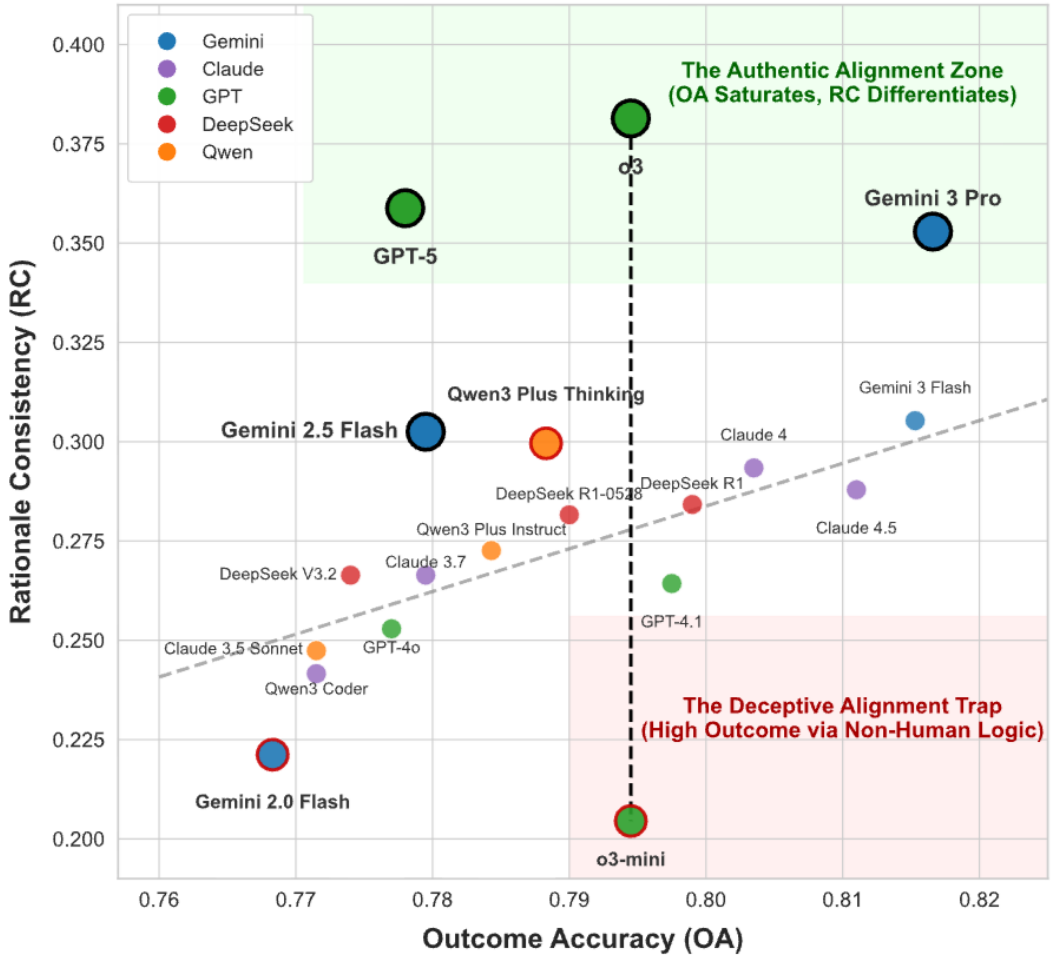

Main claim 은 Rationale Consistency(RC) 가 judge 품질의 진짜 지표라는 것!!

- **RC:** 모델 reasoning 과 human judgment 의 정렬을 정량화하는 metric
- OA 는 frontier 구간에서 saturation → GPT-5·Gemini 3 Pro 를 변별 못함
- **o3 vs o3-mini:** OA 는 동급, RC 는 ~50% 차이 → Deceptive Alignment Trap
- **OA × RC hybrid reward 학습** → RM-Bench 87.1 / JudgeBench 82.0 SOTA
- frontier 모델도 RC≈0.4 수준 → human reasoning 과 정렬은 아직 개선 여지

MetaJudge – Benchmark Construction

Human rationale 를 신뢰 가능한 Atomic Rationales 로 분해

- **HelpSteer3 4개 domain × 250개 → GPT-5 Atomic Decomposition → HelpSteer3-Atomic 1,000 instances**
  - GPT5 atomic decomposition 원칙
    - 일반적이고 주관적인 진술을 필터링하면서 구체적이고 증거에 기반한 근거를 유지
    - 각 항목이 단일한 독립적 의미 단위를 형성하도록 중복을 제거
  - 3~7개의 비평 지점을 유지하도록 인스턴스를 필터링
- **평가 강화를 위해, human annotator의 350개의 창의적 글쓰기에 동일한 Decomposition 방법 적용한 CW-Atomic을 생성**

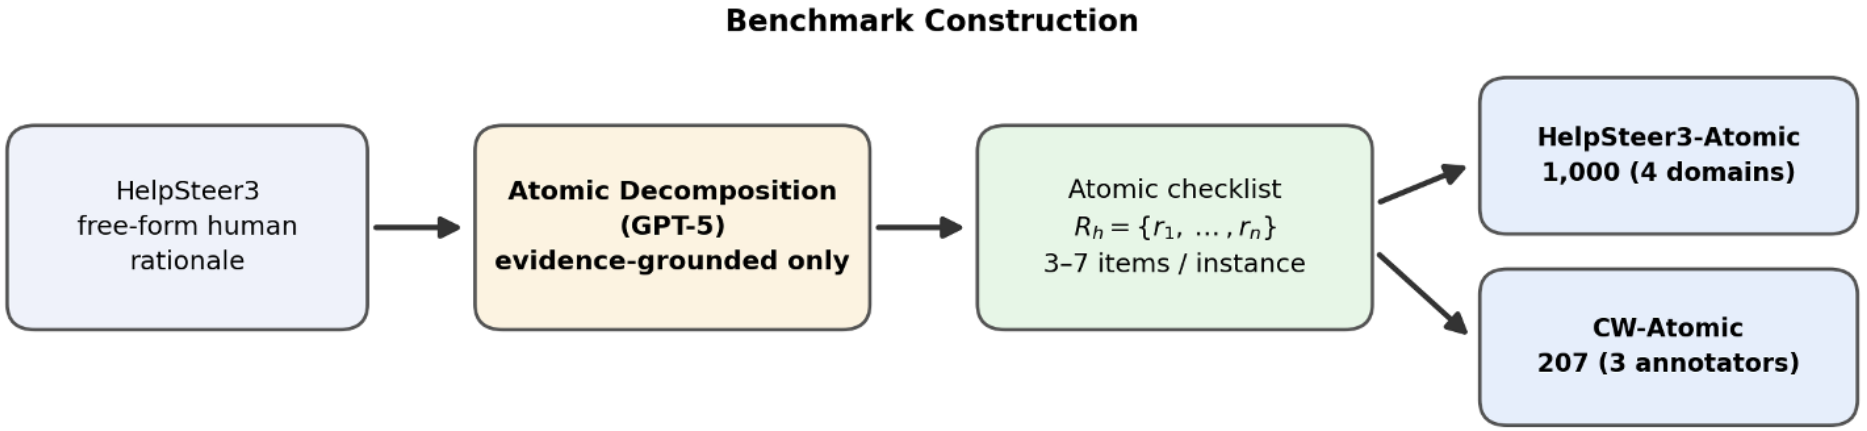

MetaJudge – Semantic Matching

Human atomic reasons R_{h} , AI 생성 atomic reasons R_{ai} 간의 semantic matching 수행

- LLM Evaluator 가 각 human reason r_{i} ($r_{\mathrm{i}} \in R_{\mathrm{h}}$) 를 R_{ai} 와 대조해 $s_{ij} \in [0, 1]$ 부여
  - $s_{ij} = 1$ **(완전 일치)**
    - AI의 reason 이 **인간의 이유**와 **핵심 조건/증거**까지 일치하는 경우
  - $s_{ij} = 0$ **(불일치):**
    - **누락(missing):** 해당 이슈를 아예 언급하지 않음
    - **모순(contradicted):** 인간과 반대되는 주장을 함 (예: 잘못된 응답을 칭찬)
    - **일반적/비국소적(generic, non-localized):** 문제 언급은 했지만, 구체적으로 짚지 못한 모호한 진술

MetaJudge – Rationale Consistency Metrics

모델의 추론 과정이 인간의 판단과 얼마나 일치하는지를 수학적으로 정량화

1. 일대일 매칭 제약

$$S_{total} = \max_{\pi} \sum_{(i,j)\in\pi} s_{ij}$$

- 각 인간 근거 r_i, AI 근거 r_j 는 각각 AI 근거, 인간근거 1개와만 짝지을 수 있음
- 가능한 짝 짓기 방법들 $\boldsymbol{\pi}$ 중, $\boldsymbol{s_{ij}}$ 들의 합이 가 최대가 되는 π 에서의 합이 해당 AI reasons 의 최종점수

2. Rationale Consistency

$$\mathrm{RC} = \frac{1}{N}\sum_{k=1}^{N}\frac{S_{\mathrm{total}}^{(k)}}{\left|R_h^{(k)}\right|}$$

- 전체 N개의 평가 샘플에 대해, 각 샘플마다 모델이 인간의 핵심 근거를 얼마나 recall 했는지 평균낸 값
- Human Reason 의 평균 회수율을 의미

Case Study – o3 vs o3-mini

Outcome 은 두 모델을 구분하지 못함

- 둘 다 A < B 정답 (Outcome 100%), 그러나 o3-mini 는 **R1–R4 전부 miss (RC 0%)**
- o3-mini: 서식·emoji 등 **surface cue 의존**, 실제 글자 수 검증 없음
- o3: 글자 수 counting (R3) · 제품명 누락 (R1) · trade-off 인식 (R4) → **RC 75%**

User Prompt: write me powerful ads for Tips, play your game in advance to relieve stress. count limit 100

Response A (Complete Text)	**Response B (Complete Text)**
1. "Feeling stressed? Play your favorite game ahead of time. Outsmart tension, be prepared for fun. #StressRelief #GameTime" 2. "Turn your frowns upside down! Play games beforehand and let stress fade away. #PrepareToPlay #DeStressNow" 3. "Don't let stress win. Play your game in advance and enjoy the victory. #GameOn #StressLess" 4. "Take control of your day, play your game before it begins. Relax, recharge, and conquer stress. #PlayAhead #StressBuster"	## Tips Ads (Under 100 Characters): **Option 1 (Benefit-focused):** > **Stress got you down? 🕹️ Play Tips! Relax, unwind, and conquer your day.** **Option 2 (Intriguing):** > **Escape the everyday. 🏝️ Tips: Your pocket-sized stress reliever. Download now!** **Option 3 (Direct & Actionable):** > **Feeling overwhelmed? 🦁 Play Tips! Fun, fast, and stress-free. Get it now!**

Model	**Model Judgment vs. Human Checklist**	**Accuracy Scores**
Human *(Expert)*	**Ground Truth Checklist (Total 4 Points):** **[R1] Missing Name:** Response A omits product name "Tips"; B includes it. **[R2] Bad Format:** Response A uses hashtags (inappropriate for ads). **[R3] Limit Violation:** Response A exceeds 100 chars (counts include hashtags); B is under limit. **[R4] Content Missing:** Response B fails to include the required "play in advance" concept.	**Final Decision:** $\boxed{A < B}$ *(Ground Truth)*
o3-mini *(Weak)*	**Selected Generated Reasons:** • Response B explicitly formats the ads as "Under 100 Characters," addressing the count limit, whereas A does not reference character count. • Response B's use of emoji enhances appeal and modern ad style. • Response A lacks clear separation between different ad versions, while B clearly labels each option. *Analysis: The model focuses on surface formatting. It criticizes A for not "referencing" the count (formatting), rather than verifying the actual length (factual check). It praises emojis (style) but misses the critical "Tips" keyword omission in A and the missing "in advance" logic in B.*	**Final Decision:** $\boxed{A < B}$ **Outcome Score:** $1/1 = \mathbf{100\%}$ (Matched) **Rationale Score:** $0/4 = \mathbf{0\%}$ (Missed R1, R2, R3, R4)
o3 *(Strong)*	**Selected Generated Reasons:** • A breaks the "count limit 100" instruction; at least two ads are 110-120 chars long, while B keeps every option under 100. (Hit R3) • A never mentions the product name "Tips"; B explicitly builds every slogan around "Tips". (Hit R1) • Although A uses "play in advance," it fails major constraints; B omits that wording but follows length/name constraints better. (Hit R4) *Analysis: The model performs a factual verification (counting characters) matching R3. It identifies the missing keyword "Tips" (matching R1). Crucially, it recognizes the trade-off: B is better despite missing the "in advance" concept (matching R4).*	**Final Decision:** $\boxed{A < B}$ **Outcome Score:** $1/1 = \mathbf{100\%}$ (Matched) **Rationale Score:** $3/4 = \mathbf{75\%}$ (Hit R1, R3, R4)

Measurement Reliability of RC metric

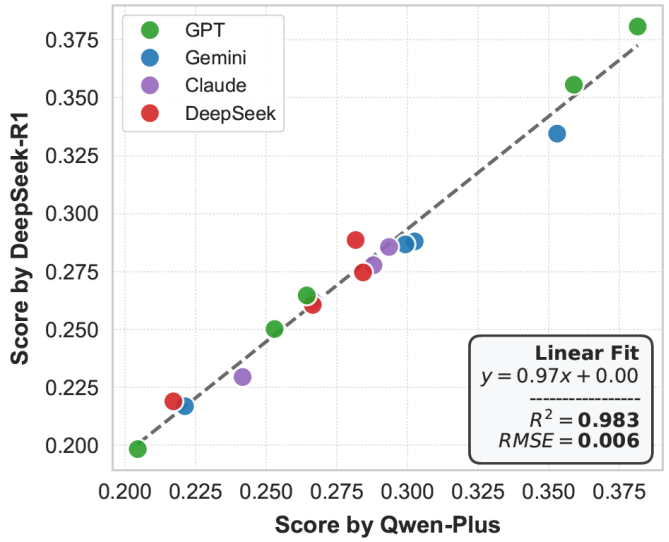

(a)

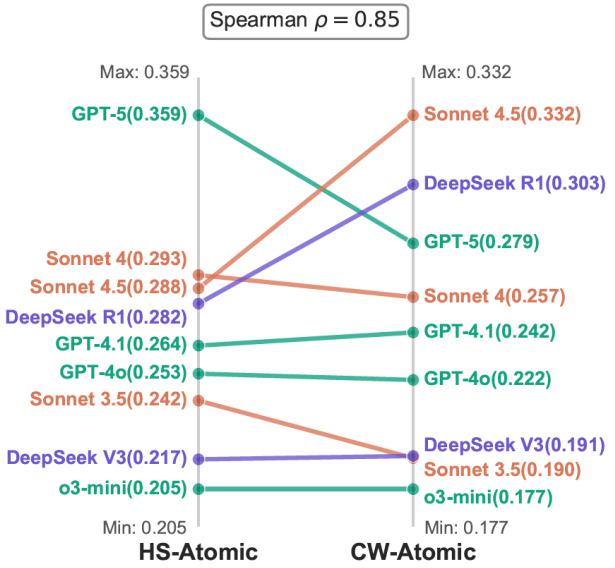

(b)

Insensitivity to Evaluator

- MetaJudge 프레임워크에서 evaluator 로 LLM 사용하고, 이에 따른 bias 존재유무 확인
- 비추론 모델(Qwen-Plus)과 추론 모델(DeepSeek-R1)로 각각 점수를 매겨 비교
- 두 점수가 거의 동일하고, Robust 함을 보여 줌

Dataset Generalization

- HS-Atomic ↔ CW-Atomic **ranking Spearman ρ = 0.85** → domain·annotator 일반화 보여줌
- domain 강점도 반영: Claude Sonnet 4.5 가 creative writing 에서 1위로 상승

Hybrid Reward RLVR

Outcome Reward + Correct rationale 를 모두 요구하는 gating reward

- $R_{outcome}$: human label 일치 여부 binary signal *(1 or 0)*
- $R_{rationale} = Average\ Precision$: 핵심 Reasons 를 **상위에 배치** 하도록 **soft ranking constraint**
  - $R_{\text{rationale}} = AP = \frac{\sum_{k=1}^{|R_{ai}|}(P@k \times \mathbb{I}(k))}{|R_h|}$
- $R_{final} = R_{rationale} \times R_{outcome}$

GRPO 최적화

- $\mathcal{J}(\theta) = \mathbb{E}_{q,\{o_i\}}\left[\frac{1}{G}\sum_{i=1}^{G}\left(\frac{\pi_\theta(o_i|q)}{\pi_{\text{old}}(o_i|q)}\hat{A}_i - \beta\mathbb{D}_{\text{KL}}\right)\right]$
- $R_{rationale}$ 단독 학습은 reasoning–verdict 불일치 **reward hacking** 발생
- $R_{outcome}$ 까지 포함한 hybrid Reard modeling 필수

Experimental Setup

동일 pipeline 에서 reward signal 만 교체 한 controlled comparison

- **Base model:** Qwen3-14B / Qwen3-30B-A3B
- **Dataset:** HelpSteer3 atomic checklist
- **Training-time MetaJudge:** Qwen3-Turbo
- **GRPO:** lr 2e-6 / batch 256 / n=8 samples/prompt / max gen 12K / 2 epochs
- **Baseline :** $R_{outcome}$ 만 사용, 나머지 조건 완전 동일
- **평가: RM-Bench** (subtle diff·style bias) · **JudgeBench** (knowledge·reasoning 판단)

Results – RM-Bench & JudgeBench

Rationale supervision 이 outcome-only 대비 평균 +5%p, 30B-A3B 는 SOTA

- **Total Avg:** 14B 76.8 → 82.9 · 30B-A3B 80.3 → 84.6 GRAM-R² (83.4) · Principles-Qwen32B (83.8) 상회
- JudgeBench 는 두 scale 모두 **+7%p 이상** → deep reasoning domain 에서 이득 큼
- **Code 이득 큼:** outcome-only 는 code correctness 개념을 학습하지 못함 *(RM-Bench Code 81.3→84.4)*

Models	RM-Bench					JudgeBench					Total Avg.
	Chat	Math	Code	Safety	Overall	Knwl.	Reas.	Math	Code	Overall	
LLM-as-a-Judge											
GPT-4o†	67.2	67.5	63.6	91.7	72.5	50.6	54.1	75.0	59.5	59.8	66.2
Claude-3.5-Sonnet†	62.5	62.6	54.4	64.4	61.0	62.3	66.3	66.1	64.3	64.8	62.9
DeepSeek-R1-0528†	76.7	74.3	51.0	89.2	72.8	59.1	82.7	80.4	92.9	78.8	75.8
Scalar Reward Model											
Skywork-Reward-Gemma-2-27B‡	71.8	59.2	56.6	94.3	70.5	59.7	66.3	83.9	50.0	65.0	67.8
Skywork-Reward-Llama-3.1-8B‡	69.5	60.6	54.5	95.7	70.1	59.1	64.3	76.8	50.0	62.5	66.3
Generative Reward Model											
RM-R1-Distilled-Qwen-32B†	74.2	91.8	74.1	95.4	83.9	76.0	80.6	88.1	70.5	78.8	81.4
RM-R1-Distilled-Qwen-14B†	71.8	90.5	69.5	94.1	81.5	68.1	72.4	87.8	84.2	78.1	79.8
RRM-32B†	66.6	81.4	65.2	79.4	73.1	79.9	70.4	87.5	65.0	75.7	74.4
Nemotron-Super‡	73.7	91.4	75.0	90.6	82.7	71.4	73.5	87.5	76.2	77.2	80.0
RewardAnything-8B-v1⋆	76.7	90.3	75.2	90.2	83.1	61.0	57.1	73.2	66.7	62.6	72.9
GRAM-R²†	76.0	89.8	80.6	**96.2**	85.7	**90.9**	83.7	87.5	61.9	81.0	83.4
Principles-Qwen32B⋆	**80.4**	92.0	77.0	95.5	86.2	74.6	85.7	85.7	90.5	81.4	83.8
Training GenRMs with Outcome Supervision (Baselines)											
Qwen3-14B (Outcome-Only)	72.3	92.6	77.8	91.7	83.6	55.8	80.6	78.6	85.7	70.0	76.8
Qwen3-30B-A3B (Outcome-Only)	68.7	95.9	81.3	93.6	84.9	65.6	87.8	82.1	76.2	75.7	80.3
Training GenRMs with Outcome and Rationale Supervision (Ours)											
Qwen3-14B (Ours)	75.7	**95.7**	82.2	93.2	86.7	66.9	86.7	**91.1**	90.5	79.1	82.9
Qwen3-30B-A3B (Ours)	74.9	95.5	**84.4**	93.6	**87.1**	73.4	**89.8**	82.1	**95.2**	**82.0**	**84.6**

Results – RLHF & RC Ablation

Method	Hard Prompt		Creative Writing	
	Score (%)	CI (%)	Score (%)	CI (%)
Before Post-Training				
SFT	12.61	(-1.6 / +1.3)	41.12	(-2.3 / +2.2)
Post-Training with RM using different methods				
Outcome-Only	19.10	(-1.6 / +1.8)	62.00	(-1.7 / +1.6)
Ours	**21.22**	(-1.6 / +2.2)	**69.08**	(-1.6 / +1.9)

Method	HelpSteer3-Atomic		CW-Atomic	
	Score	Δ	Score	Δ
Base Model	0.2505	-	0.2385	-
Outcome-Only	0.2108	↓ 0.0397	0.1677	↓ 0.0708
Ours	**0.3718**	↑ **0.1213**	**0.2526**	↑ **0.0141**

GenRM 개선이 downstream alignment 로 이어짐

- **Arena Hard v2:** Hard Prompt **19.10 → 21.22** · Creative Writing **62.00 → 69.08 (+7%p)**
- creative writing 의 implicit constraint (필수 요소·글자 수) → outcome-only reward 는 gaming 에 취약
- **RC ablation:** outcome-only 는 base 대비 **RC 하락 (–0.040 / –0.071)** · ours 는 **상승 (+0.121 / +0.014)**
- OOD (CW-Atomic) 에서도 개선 유지 → 더 나은 generalization

Training Dynamics

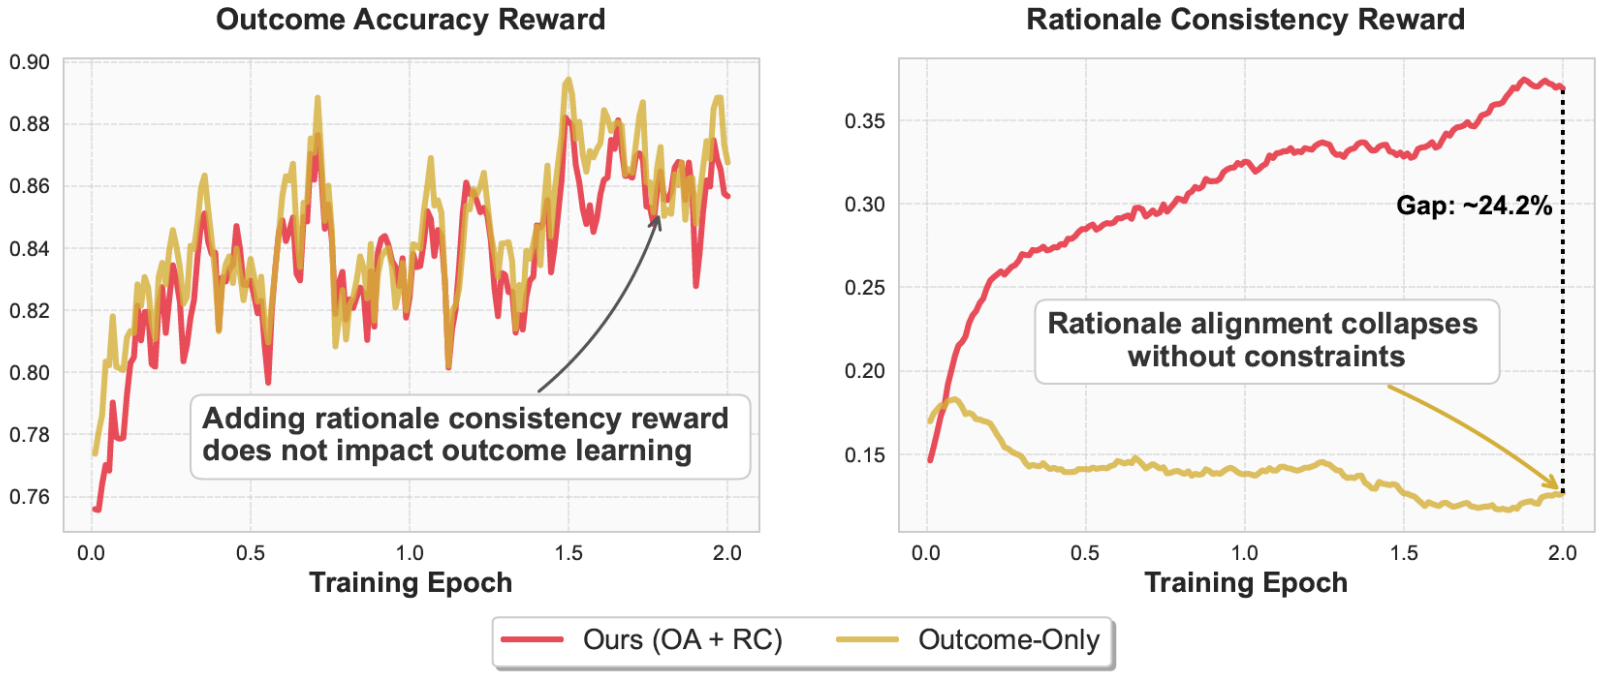

Outcome 은 동일, rationale 은 붕괴

- Outcome reward 곡선은 두 방법이 거의 동일 → **정답 선택은 rationale 없이도 학습됨**
- Rationale reward 는 outcome-only 에서 지속 하락 → 최종 **~24.2% gap**
- 검증 비용이 reward 와 약하게 상관 → 모델이 **값싼 surrogate cue 로 대체**
- Outcome level 은 aligned 해 보여도 logic level 에서 drift

Training Dynamics

Diagnosis

- Outcome accuracy 는 **deceptive alignment 를 탐지하지 못함**
- RC + MetaJudge 로 reasoning 정렬을 robust 하게 정량화

Training

- Hybrid reward (R_rationale × R_outcome) 로 **RM-Bench 87.1 / JudgeBench 82.0 SOTA**
- RLHF creative writing +7%p · rationale degeneration 역전

Limitation

- 고품질 human annotation 필요)Frontier 모델도 **RC < 0.4)**
- Synthetic rationale 은 아직 human judgment logic 의 대체재가 아님

Thank you