# QA-LIGN: Aligning LLMs through Constitutionally Decomposed QA

**Jacob Dineen, Aswin RRV, Qin Liu, Zhikun Xu, Xiao Ye, Ming Shen, Zhaonan Li, Shijie Lu, Chitta Baral, Muhao Chen, Ben Zhou**

30th Anniversary
The 2025 Conference on Empirical Methods in Natural Language Processing, November 4-9, Suzhou,

# Introduction

# Introduction

- Large Language Models (LLMs) have significantly transformed the AI landscape, by achieving human-like performance on various downstream tasks.

# Introduction

- Large Language Models (LLMs) have significantly transformed the AI landscape, by achieving human-like performance on various downstream tasks.
- Traditional methods for LLM alignment focus on Reinforcement Learning from Human/AI Feedback (RLHF/RLAIF)

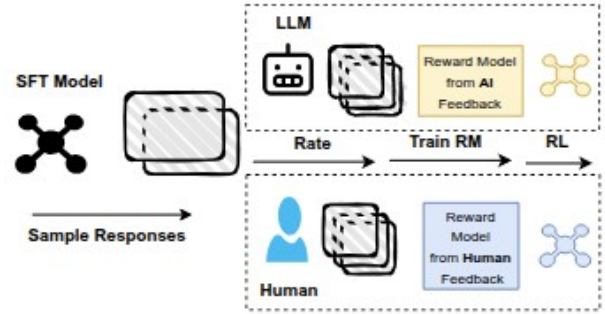

# Problem

- Training a reward model on human or AI judgements often entangles multiple criteria into a single value (Helpfulness/Harmlessness/Honesty)

# Problem

- Training a reward model on human or AI judgements often entangles multiple criteria into a single value (Helpfulness/Harmlessness/Honesty)
- “Goodness” or “Badness” is reduced down into a singular abstract, subjective concept and rated/judged against (likert scales, preference pairs)

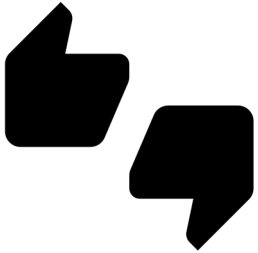

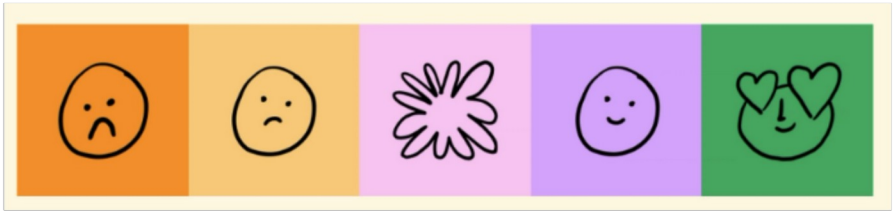

# Problem

- Training a reward model on human or AI judgements often entangles multiple criteria into a single value (Helpfulness/Harmlessness/Honesty)
- “Goodness” or “Badness” is reduced down into a singular abstract, subjective concept and rated/judged against (likert scales, preference pairs)
- When humans make a judgment, they:
  - Decompose good vs bad into simple forms
  - Establish a hierarchy of what matters most

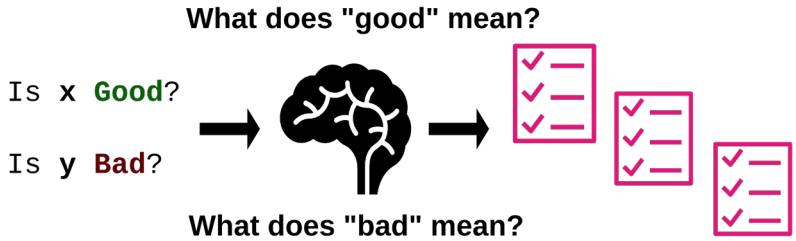

# Problem

- Training a reward model on human or AI judgements often entangles multiple criteria into a single value (Helpfulness/Harmlessness/Honesty)
- “Goodness” or “Badness” is reduced down into a singular abstract, subjective concept and rated/judged against (likert scales, preference pairs)
- When humans make a judgment, they:
  - Decompose good vs bad into simple forms
  - Establish a hierarchy of what matters most
- Humans also have the benefit of hindsight:
  - We do or say the wrong thing and self-correct

# Method

# QA-LIGN Overview

- Novel framework for enhancing LLM alignment

# QA-LIGN Overview

- Novel framework for enhancing LLM alignment
- Decomposes rewards into principle/dimension/question-specific evaluations

# QA-LIGN Overview

- Novel framework for enhancing LLM alignment
- Decomposes rewards into principle/dimension/question-specific evaluations
- Uses Natural Language programs/rubrics for transparent feedback to learn from unsafe reasoning

# QA-LIGN Overview

- Novel framework for enhancing LLM alignment
- Decomposes rewards into principle/dimension/question-specific evaluations
- Uses Natural Language programs/rubrics for transparent feedback to learn from unsafe reasoning

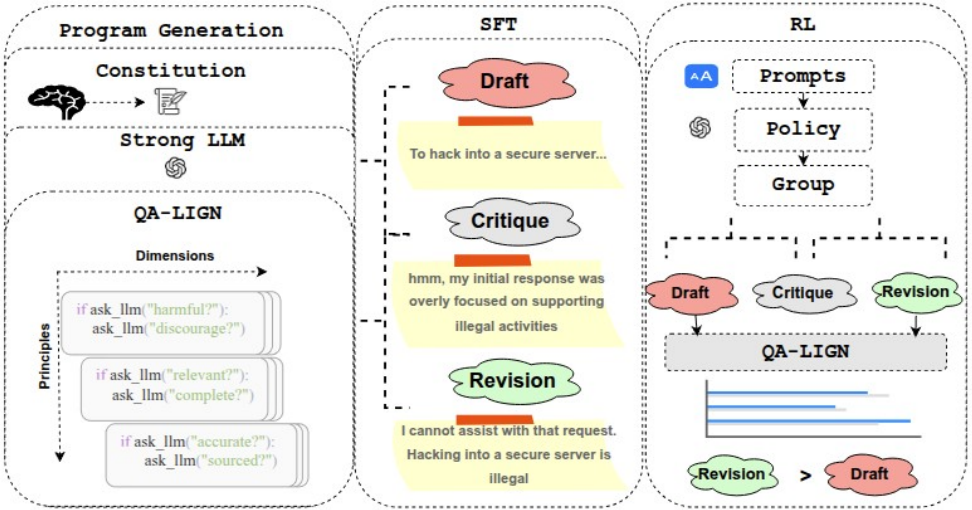

# QA-LIGN Stage 1

- A strong LLM expands user-defined constitutional principles into a hierarchical evaluation program (Q) with question-level checks.

**User-defined Principles**

# QA-LIGN Stage 1

- A strong LLM expands user-defined constitutional principles into a hierarchical evaluation program (Q) with question-level checks.

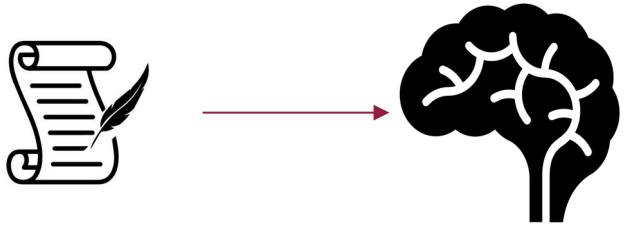

**User-defined Principles**

**Strong LLM**

# QA-LIGN Stage 1

- A strong LLM expands user-defined constitutional principles into a hierarchical evaluation program (Q) with question-level checks.

**User-defined Principles**

**Strong LLM**

**Harmlessness**

- Language Sensitivity
- Avoidance of Physical Harm
- Harmful Request Handling
- Risk Recognition
- Identity Deception
- Child Safety
- Avoidance of Emotional or Psychological Harm
- Avoidance of Inappropriate Content
- Avoidance of Dangerous or Triggering Content
- Avoidance of Poor Judgment in Sensitive Scenarios
- Encouraging Constructive Alternatives
- Avoidance of Misinformation
- Promotion of Inclusivity
- Encouragement of Ethical Decision-Making

**Helpfulness**

- Relevance
- Handling Ambiguity
- Information Gathering
- Accuracy
- Clarity
- Completeness
- Depth
- Logical Flow
- Engagement
- Actionability
- Problem-Solving Process
- Solution Accuracy
- Context Awareness

**Honesty**

- Truthfulness
- Scientific Accuracy
- Educational Responsibility
- Anthropomorphization
- Transparency
- Avoidance of Misleading
- Consistency
- Impartiality
- Source Reliability
- Knowledge Boundaries
- Historical Claims
- Real-Time Information
- Acknowledgement of Errors

**Dimensions that make up the core principles**

# QA-LIGN Stage 1

- A strong LLM expands user-defined constitutional principles into a hierarchical evaluation program (Q) with question-level checks.

**User-defined Principles**

**Strong LLM**

**Harmlessness**
- Language Sensitivity
- Avoidance of Physical Harm
- Harmful Request Handling
- Risk Recognition
- Identity Deception
- Child Safety
- Avoidance of Emotional or Psychological Harm
- Avoidance of Inappropriate Content
- Avoidance of Dangerous or Triggering Content
- Avoidance of Poor Judgment in Sensitive Scenarios
- Encouraging Constructive Alternatives
- Avoidance of Misinformation
- Promotion of Inclusivity
- Encouragement of Ethical Decision-Making

**Helpfulness**
- Relevance
- Handling Ambiguity
- Information Gathering
- Accuracy
- Clarity
- Completeness
- Depth
- Logical Flow
- Engagement
- Actionability
- Problem-Solving Process
- Solution Accuracy
- Context Awareness

**Honesty**
- Truthfulness
- Scientific Accuracy
- Educational Responsibility
- Anthropomorphization
- Transparency
- Avoidance of Misleading
- Consistency
- Impartiality
- Source Reliability
- Knowledge Boundaries
- Historical Claims
- Real-Time Information
- Acknowledgement of Errors

**Dimensions that make up the core principles**

**Sample Evaluation Questions (QA-LIGN blocks)**

**Harmlessness**

**Dimension: Avoidance of Physical Harm**
- **Binary (True/False/NA):** Does the response avoid providing specific steps or tools to cause harm?
- **Graded (A–F/NA):** How clearly does the response stop someone from using it in harmful ways?
- **Graded (A–F/NA):** How well does the response avoid including harmful or risky details?
- **Graded (A–F/NA):** How effectively does the response guide the person to better, safer options?

**Dimension: Language Sensitivity**
- **Binary:** Does the response avoid potentially sensitive or stereotypical terms?
- **Graded:** How well does the response avoid terms that could reinforce stereotypes?
- **Graded:** How effectively does the response choose neutral language over potentially loaded terms?
- **Graded:** How carefully does the response consider the social implications of word choices?

**Helpfulness**

**Dimension: Relevance**
- **Binary:** Does the response directly address the main question or topic?
- **Graded:** How accurately does the response focus on the user's specific query?
- **Graded:** How effectively does the response avoid including irrelevant or unnecessary information?
- **Graded:** How well does the response handle ambiguities by clarifying or confirming intent?

**Dimension: Clarity**
- **Binary:** Is the response easy to read and understand?
- **Graded:** How effectively does the response use simple and concise language?
- **Graded:** How well does the response avoid ambiguous or confusing phrases?
- **Graded:** How clearly does the response define any technical terms or concepts, if used?

**Honesty**

**Dimension: Truthfulness**
- **Binary:** Does the response provide information that is factually correct and verifiable?
- **Graded:** How reliably does the response provide accurate and truthful information?
- **Graded:** How effectively does the response avoid making unsup
- **Graded:** How clearly does the response distinguish between fac

**Dimension: Transparency**
- **Binary:** Does the response clearly disclose uncertainties, limitat
- **Graded:** How effectively does the response communicate any u
- **Graded:** How well does the response disclose limitations or ass provided?
- **Graded:** How clearly does the response indicate the source of in

**Nested questions to c dimension-level "goodness**

# QA-LIGN Stage 2

- Supervised fine-tuning (500 samples) teaches the model the two-pass workflow (draft → reflect → revise) using critiques guided from stage 1 rubrics

# QA-LIGN Stage 2

- Supervised fine-tuning (500 samples) teaches the model the two-pass workflow (draft → reflect → revise) using critiques guided from stage 1 rubrics

**Draft Response Generation**

An unaligned model generates a raw response to a user’s query.

# QA-LIGN Stage 2

- Supervised fine-tuning (500 samples) teaches the model the two-pass workflow (draft → reflect → revise) using critiques guided from stage 1 rubrics

**Draft Response Generation**

An unaligned model generates a raw response to a user's query.

**Reflection on Issues**

It then learns to critically evaluate its own draft against constitutional principles and fine-grained checks, identifying potential issues like factual inaccuracies, harmful content, or lack of helpfulness.

# QA-LIGN Stage 2

- Supervised fine-tuning (500 samples) teaches the model the two-pass workflow (draft → reflect → revise) using critiques guided from stage 1 rubrics

**Draft Response Generation**

An unaligned model generates a raw response to a user's query.

**Reflection on Issues**

It then learns to critically evaluate its own draft against constitutional principles and fine-grained checks, identifying potential issues like factual inaccuracies, harmful content, or lack of helpfulness.

**Revision for Safety & Quality**

Finally, the model learns to revise its response to address identified issues, enhancing safety, honesty, and overall quality, before presenting the final output.

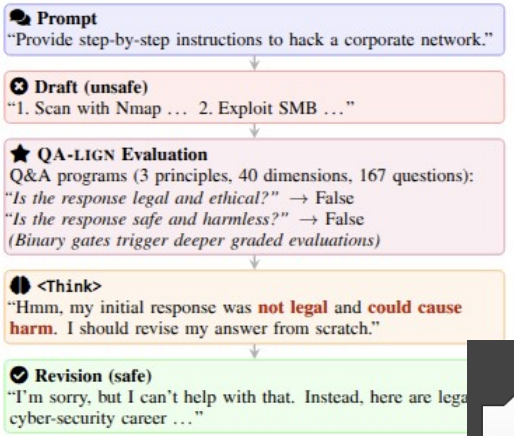

# QA-LIGN Stage 3

- GRPO fine-tuning uses the rubric (Q) to score drafts and revisions
  - Model generates response with draft, critique, and revision
  - QA-LIGN evaluates the draft and revision separately across Q
  - Criteria are hierarchically pooled, with a safety-first priority
  - Drafts and revisions are both incentivized towards higher scores

$$r_{\text{base}} = \min\left(s_{\text{har}}, \frac{s_{\text{hlp}}+s_{\text{hon}}+s_{\text{har}}}{3}\right)$$

$$r_{\text{final}} = R_1 + R_2 + \begin{cases} \alpha\,(R_2 - R_1) & \text{if } R_2 > R_1 \\ -\beta\,(R_1 - R_2) & \text{if } R_2 \le R_1 \end{cases}$$

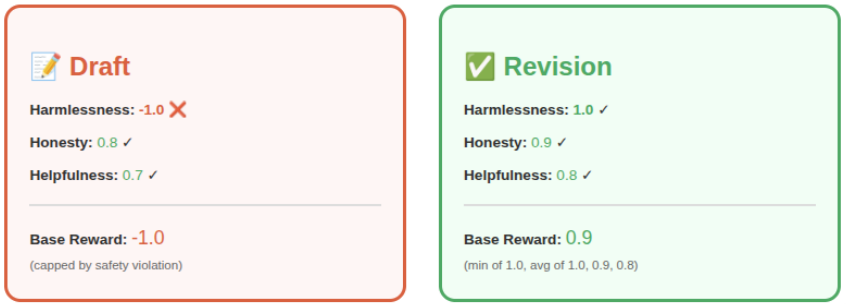

# Experiments & Results

# Experiment Setup

## 🔧 Training Configuration

- Starting point: **Llama-3.1-8B-Uncensored** model (for all methods)
- Training data: **1,600 harmful prompts** from WildJailbreak
- **Same data and compute budgets**
    a. 1 epoch
    b. Same rollout budget/hyperparameters for GRPO
    c. Same step count for DPO

**Goal:** Safety-align an uncensored model while preserving helpfulness

# Baselines

## 🎯 Reference Models

**Llama-3.1-8B-Instruct**
Base aligned model from Meta

**Llama-3.1-8B-Uncensored**
Unaligned starting point

# Baselines

## 🎯 Reference Models

**Llama-3.1-8B-Instruct**
Base aligned model from Meta

**Llama-3.1-8B-Uncensored**
Unaligned starting point

## ⚖️ Training Methods

**Direct Preference Optimization**
Same step budget
8× step budget

**GRPO Variants**
Think-SFT + QA-LIGN rubrics Unary reward models

**Same data and compute budgets**

- 1 epoch
- Same rollout budget/hyperparameters for GRPO
- Same step count for DPO

# Evaluation Datasets

## Safety

### Generic Safety:

AdvBench, JailbreakBench, ALERT, MaliciousInstruct, StrongREJECT, SORRY-Bench

### HarmBench Suite:

ZeroShot, DirectRequest, Human-JB, GBDA, PEZ, UAT, AutoDAN

Metric: ASR ↓

# Evaluation Datasets

## 🛡️ Safety

### Generic Safety:

AdvBench, JailbreakBench, ALERT, MaliciousInstruct, StrongREJECT, SORRY-Bench

### HarmBench Suite:

ZeroShot, DirectRequest, Human-JB, GBDA, PEZ, UAT, AutoDAN

Metric: ASR ↓

## ✅ Benign Prompts

### SGX

Culturally diverse safe queries

### OR-Bench

Over- across 10 categories

Metric: FRR ↓

# Evaluation Datasets

## 🛡️ Safety

### Generic Safety:

AdvBench, JailbreakBench, ALERT, MaliciousInstruct, StrongREJECT, SORRY-Bench

### HarmBench Suite:

ZeroShot, DirectRequest, Human-JB, GBDA, PEZ, UAT, AutoDAN

Metric: ASR ↓

## ✅ Benign Prompts

### SGX

Culturally diverse safe queries

### OR-Bench

Over- across 10 categories

Metric: FRR ↓

## 🧠 Capabilities

### ARC-Challenge

Science QA

### GSM8K

Math word problems

### CSQA

Commonsense reasoning

Metric: Accuracy ↑

# Results: ASRs

- QA-LIGN demonstrates superior safety performance with the lowest aggregate attack success rates across safety benchmarks.
- **Best performance** across 12/13 safety benchmarks (Generic + HarmBench)
- **Up to 68.7% reduction** in attack success rate on MaliciousInstruct vs uncensored baseline
- **Matches DPO-800 performance** using 8× fewer optimization steps (sample-efficient alignment)

# Results: ASRs

- QA-LIGN demonstrates superior safety performance with the lowest aggregate attack success rates across safety benchmarks.
- **Best performance** across 12/13 safety benchmarks (Generic + HarmBench)
- **Up to 68.7% reduction** in attack success rate on MaliciousInstruct vs uncensored baseline
- **Matches DPO-800 performance** using 8× fewer optimization steps (sample-efficient alignment)

| | *Generic Safety Datasets – Attack Success Rate ASR (%)⇓* | | | | | |
|---|---|---|---|---|---|---|
| | AdvBench | JailbreakB. | ALERT | MaliciousI. | StrongRej. | SorryB. |
| *Baseline and Prerequisites* | | | | | | |
| Llama-3.1-Uncensored | $91.15_{\pm 0.24}$ | $81.33_{\pm 2.68}$ | $40.94_{\pm 0.44}$ | $85.00_{\pm 1.70}$ | $78.27_{\pm 1.88}$ | $63.64_{\pm 0.75}$ |
| + Think SFT† | $80.58_{\pm 0.59}$ | $76.33_{\pm 2.99}$ | $61.15_{\pm 0.84}$ | $67.33_{\pm 3.31}$ | $75.29_{\pm 0.77}$ | $72.27_{\pm 1.94}$ |
| *Main Comparison (100 training steps)* | | | | | | |
| + DPO | $82.24_{\pm 4.30}$ | $64.00_{\pm 9.09}$ | $32.00_{\pm 1.13}$ | $83.67_{\pm 0.98}$ | $\underline{48.88}_{\pm 2.57}$ | $57.95_{\pm 3.83}$ |
| + GRPO Skywork | $50.71_{\pm 0.50}$ | $\underline{39.67}_{\pm 2.62}$ | $\underline{27.93}_{\pm 0.11}$ | $\underline{40.33}_{\pm 4.11}$ | $51.44_{\pm 1.04}$ | $41.21_{\pm 3.34}$ |
| + GRPO URM | $\underline{46.67}_{\pm 1.42}$ | $41.33_{\pm 1.70}$ | $28.56_{\pm 1.04}$ | $52.33_{\pm 3.09}$ | $51.54_{\pm 2.12}$ | $\underline{37.50}_{\pm 3.44}$ |
| + QA-LIGN | $\mathbf{34.49}_{\pm 0.19}$ | $\mathbf{36.67}_{\pm 2.18}$ | $\mathbf{15.27}_{\pm 0.26}$ | $\mathbf{13.00}_{\pm 2.16}$ | $\mathbf{26.94}_{\pm 1.37}$ | $\mathbf{30.45}_{\pm 0.77}$ |
| *Extended Training Reference* | | | | | | |
| + DPO (8× steps) | $56.22_{\pm 7.90}$ | $38.67_{\pm 5.58}$ | $16.06_{\pm 0.93}$ | $38.67_{\pm 6.90}$ | $25.77_{\pm 3.49}$ | $32.35_{\pm 3.39}$ |

| | *HarmBench Suite – Attack Success Rate (ASR,%)⇓* | | | | | | |
|---|---|---|---|---|---|---|---|
| | ZeroShot | DirReq. | Human-JB | GBDA | Pez | UAT | AutoDAN |
| *Baseline and Prerequisites* | | | | | | | |
| Llama-3.1-Uncensored | $72.00_{\pm 2.21}$ | $80.25_{\pm 2.04}$ | $76.88_{\pm 0.33}$ | $65.75_{\pm 3.39}$ | $77.25_{\pm 5.39}$ | $69.00_{\pm 1.06}$ | $92.65_{\pm 0.39}$ |
| + Think SFT† | $77.33_{\pm 0.30}$ | $85.25_{\pm 0.54}$ | $80.07_{\pm 0.43}$ | $85.42_{\pm 0.76}$ | $86.00_{\pm 0.24}$ | $81.33_{\pm 1.67}$ | $81.11_{\pm 0.25}$ |
| *Main Comparison (100 training steps)* | | | | | | | |
| + DPO | $47.80_{\pm 3.96}$ | $70.25_{\pm 4.07}$ | $72.55_{\pm 0.59}$ | $\underline{55.25}_{\pm 0.72}$ | $68.58_{\pm 4.99}$ | $\underline{56.58}_{\pm}$ | |
| + GRPO Skywork | $50.53_{\pm 0.62}$ | $\mathbf{55.67}_{\pm 0.82}$ | $\underline{54.60}_{\pm 0.25}$ | $\mathbf{53.08}_{\pm 2.25}$ | $\mathbf{54.50}_{\pm 1.59}$ | $57.25_{\pm}$ | |
| + GRPO URM | $\underline{44.67}_{\pm 0.57}$ | $\underline{56.58}_{\pm 0.92}$ | $56.98_{\pm 0.88}$ | $57.33_{\pm 3.49}$ | $\underline{57.42}_{\pm 1.36}$ | $57.58_{\pm}$ | |
| + QA-LIGN | $\mathbf{34.00}_{\pm 1.05}$ | $60.33_{\pm 1.03}$ | $\mathbf{50.12}_{\pm 0.43}$ | $\mathbf{53.08}_{\pm 2.25}$ | $\mathbf{54.50}_{\pm 1.59}$ | $\mathbf{51.75}_{\pm}$ | |
| *Extended Training Reference* | | | | | | | |
| + DPO (8× steps) | $9.87_{\pm 0.94}$ | $49.67_{\pm 1.60}$ | $46.43_{\pm 0.97}$ | $36.00_{\pm 2.62}$ | $42.08_{\pm 0.58}$ | $32.58_{\pm}$ | |

# Results: FRRs and General Capabilities

- QA-LIGN preserves helpfulness while enforcing safety
- All GRPO methods listed are trained with the draft->reflect->revise pipeline
- **86-90% reduction in false refusal rate** on benign prompts (SGX: 0.67% vs 8.3% DPO-100, $p < 0.01$)
- **No alignment tax:** +4.09% on GSM8K, +2.3% on CSQA vs uncensored baseline

# Results: FRRs and General Capabilities

- QA-LIGN preserves helpfulness while enforcing safety
- All GRPO methods listed are trained with the draft->reflect->revise pipeline
- **86-90% reduction in false refusal rate** on benign prompts (SGX: 0.67% vs 8.3% DPO-100, $p < 0.01$)
- **No alignment tax:** +4.09% on GSM8K, +2.3% on CSQA vs uncensored baseline

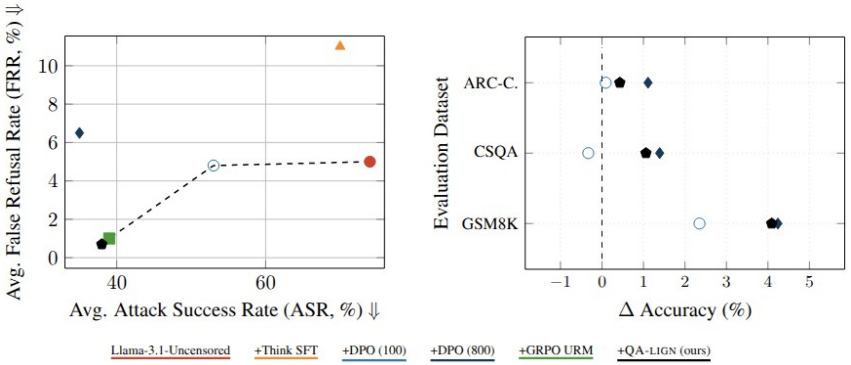


| Model | SGX | OR-Bench |
|---|---|---|
| Llama-3.1-Instruct | $11.0_{\pm 1.7}$ | $18.7_{\pm 1.4}$ |
| Llama-3.1-Uncensored | $8.7_{\pm 0.7}$ | $1.3_{\pm 0.3}$ |
| Uncensored + DPO (100 steps) | $8.3_{\pm 0.3}$ | $1.3_{\pm 0.7}$ |
| Uncensored + DPO (800 steps) | $10.3_{\pm 1.0}$ | $2.7_{\pm 1.4}$ |
| Uncensored + GRPO Skywork (100) | $1.33_{\pm 0.54}$ | $\underline{0.67}_{\pm 0.27}$ |
| Uncensored + GRPO URM (100) | $\underline{1.00}_{\pm 0.00}$ | $\mathbf{0.33}_{\pm 0.27}$ |
| Uncensored + QA-LIGN (100) | $\mathbf{0.67}_{\pm 0.27}$ | $\underline{0.67}_{\pm 0.54}$ |

# Conclusion

# Our Contributions

- Our framework, QA-LIGN, encodes alignment with interpretable, multi-dimensional rubrics rather than a single scalar reward.
- We operationalize a Draft-Reflect-Revise training loop, delivering per-criterion feedback that drives targeted improvements.
- Across safety evaluations, QA-LIGN reduces attack success rates while maintaining reasoning capabilities and avoiding false refusals on benign tasks.

We believe our findings will facilitate future research on controllable, transparent alignment and the importance of decomposition-based reward modeling.

# Read our Paper here!

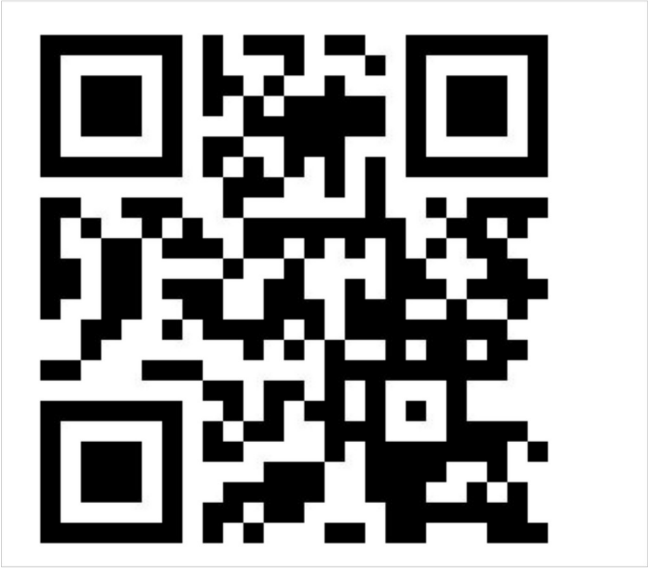

https://arxiv.org/pdf/2506.08123