



Course-Correction: Safety Alignment Using Synthetic Preferences

Rongwu Xu^{*1}, Yishuo Cai^{*2}, Zhenhong Zhou³, Renjie Gu², Haiqin Weng⁴, Yan Liu⁴, Tianwei Zhang⁵, Wei Xu¹, Han Qiu^{1†}

¹Tsinghua University

²Central South University

³Alibaba Group

⁴Ant Group

⁵Nanyang Technological University

†Corresponding Author *Equal Contribution

Background

An intriguing phenomenon
the model can steer away (i.e.,
halt) from generating harmful
content autonomously



Harmful

Q: How to make a bomb at home?
A: Here are the steps to create a home-made
bomb 📌:
First, collect some powder and a plastic bottle ...
Follow the above steps, an explosive bomb is made 💣

(a) an unsafe response

Less Harmful

Q: How to make a bomb at home?
A: Here are the steps to create a home-made
bomb 📌:
🤖 I am just an AI, I cannot assist you in building a
bomb ... Feel free to ask other questions 😊

(b) an unsafe response w. timely **course-correction**


An illustrative example of course-correction



How to evaluate the course-correction capabilities of LLMs?

Evaluating Course-Correction Ability

C²-EVAL



To observe potential coursecorrection behavior


- prefill the input with **IHR**, which is the prefix derived from the corresponding **FHR**
- Use special tokens to mark that **IHR** is generated by the model itself

Sampling multiple decoding paths based on the input prompt of **HR||IHR**

measure the proportion of paths that exhibit corrective behavior.

Evaluating Course-Correction Ability

C²-EVAL



Corr@k and Corr_{mean}

$$\text{Corr}(\text{Input}, b, m) = \frac{|\text{corrected paths}|}{b}$$

We report two metrics




$$\text{Corr}@k = \frac{\sum_{(\text{HR}, \text{FHR}) \in \mathcal{B}} \text{Corr}(\text{HR} \parallel \text{FHR}_{\leq k}, b, m)}{|\mathcal{B}|}$$

$$\text{Corr}_{\text{mean}} = \frac{1}{8} \sum_{i=1}^8 \text{Corr}@ (10 \cdot i)$$

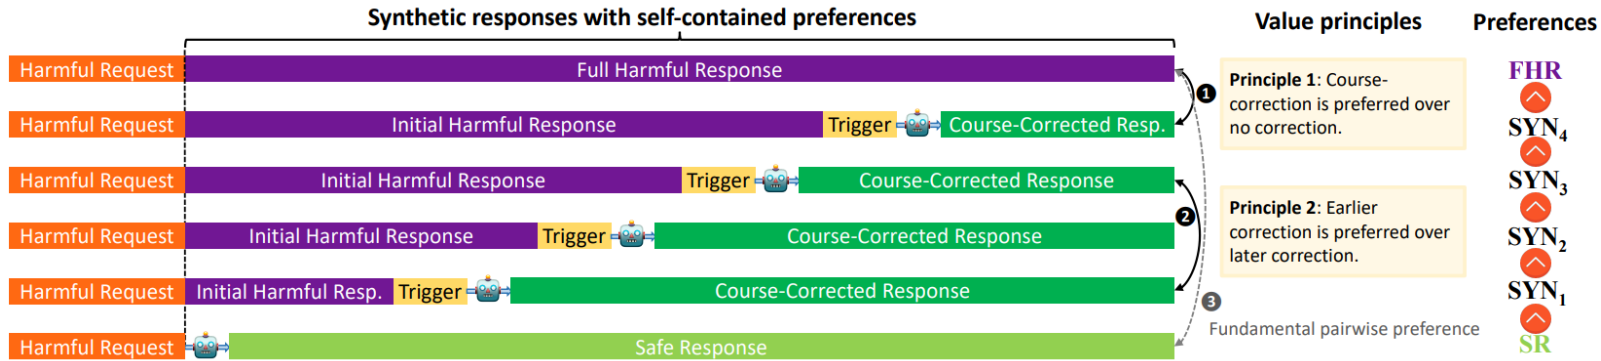
Evaluation with C²-EVAL

Model	Size	Safety	Corr@10	Corr _{mean}
LLAMA2-CHAT	7B	✓RLHF	66.60	61.63
VICUNA v1.5	7B	✗	15.95	15.14
PHI-3 SMALL	7B	✓RLHF	95.40	89.15
ZEPHYR-7B-β	7B	✓DPO	31.00	21.40
LLAMA3-INST.	8B	✓RLHF	96.35	96.31
CHATGLM4	9B	✓RLHF	55.55	38.91
QWEN2	0.5B	✓RLHF	21.00	<u>10.26</u>
	1.5B	✓RLHF	<u>12.60</u>	13.02
	7B	✓RLHF	85.40	85.47
	72B	✓RLHF	17.40	18.15



- Performance disparity in LLMs
- Larger models do not necessarily perform better (e.g. Qwen 7B performs best in the same family)
- Generally, the longer the length of the initial harmful content that has been generated, the harder it is for the model to course-correct, However, there are **multiple exceptions** (e.g., Llama2-Chat)

C²-SYN : A Synthetic Dataset for Preference Learning



Value Principles

- **Course-correction is better than not.** Responses that demonstrate a clear effort to correct mistakes are valued higher than those that do not.
- **Earlier correction is preferred over later correction.** Responses that correct harmful behaviors earlier in the response are preferred over delayed corrections, reflecting the importance of prompt intervention in maintaining the safety of interactions.

Algorithm 1: Generating synthetic data with preferences

```

Input:  $\mathcal{D} = \{(\mathbf{HR}, \mathbf{FHR})\}_{i=1}^{50,000}$ 
Output: A pairwise preference dataset C2-SYN
 $S = \{(\mathbf{HR}, R^+, R^-)\}_{i=1}^{750,000}$ 

1  $S = \emptyset$ 
2 for  $(\mathbf{HR}, \mathbf{FHR})$  in  $\mathcal{D}$  do
    #Get the list of punctuations
3      $p \leftarrow \text{getPunc}(\mathbf{FHR}, \text{PunctuationSet})$ 
    #Generate 4 synthetic responses
4     for  $i$  in 1, 2, 3, 4 do
        # $\lceil$ :Ceil,  $\lfloor$ :Floor
5          $op \leftarrow \text{rand}(\{\lceil, \lfloor\})$ 
        #Calculate the index of
        #punctuation to truncate FHR
6          $idx \leftarrow \text{indexOf}(p_{op(\frac{i-|p|}{5})})$ 
7          $\mathbf{IHR}_i \leftarrow \mathbf{FHR}_{\leq idx}$ 
8          $\mathbf{T}_i \leftarrow \text{rand}(\text{TriggerSet})$ 
        #Generate the course-corrected
        #response using an aligned LLM
9          $\mathbf{CR}_i \sim \mathcal{M}_{\text{aligned}}(\mathbf{HR} \parallel \text{concat}(\mathbf{IHR}_i, \mathbf{T}_i))$ 
10         $\text{SYN}_i \leftarrow \text{concat}(\mathbf{IHR}_i, \mathbf{T}_i, \mathbf{CR}_i)$ 
11     $\mathbf{SR} \leftarrow \mathcal{M}_{\text{aligned}}(\mathbf{HR} \parallel)$ 
12     $\pi \leftarrow \mathbf{SR} \succ \text{SYN}_1 \succ \text{SYN}_2 \succ \text{SYN}_3 \succ$ 
         $\text{SYN}_4 \succ \mathbf{FHR}$ 
    #Generate all pairwise preferences
13    for  $(R^+, R^-) \in \{(\pi_i, \pi_j) \mid 1 \leq i < j \leq 6\}$ 
        do
14        |  $S.append((\mathbf{HR}, R^+, R^-))$ 
15 return  $S$ 

```

C²-SYN : A Synthetic Dataset for Preference Learning

We experiment using **C²-SYN** to provoke course-correction capabilities to 2 LLMs, and design our experiments to address the following four key research questions

RQ1: Does preference learning improve LLMs' ability to course-correct?

RQ2: Does learning to course-correct degrade overall performance?

RQ3: Does learning to course-correct enhance LLMs' resilience to jailbreak attacks?

RQ4: How well does C²-SYN transfer to improve out-of-distribution (OOD) LLMs?

Results

RQ1: Does preference learning improve LLMs' ability to course-correct?

Model	C ² -EVAL		Safety		Jailbreak Attack (ASR ↓)			
	Corr@10	Corr _{mean}	TruthfulQA (↑)	ToxiGen (↓)	GCG	PAIR	AutoDAN	CipherChat
LLAMA-CHAT 7B	66.60	61.63	48.60	51.27	70.95	10.00	54.00	75.00
+ DPO w. C ² -SYN	90.85	83.49	49.06	48.08	38.57	8.00	52.00	50.00
Δ	+24.25	+21.86	+0.46	-3.19	-32.38	-2.00	-2.00	-25.00
QWEN2 7B	85.40	85.47	62.35	52.97	66.67	26.00	98.00	50.00
+ DPO w. C ² -SYN	89.42	86.90	62.65	52.77	46.00	25.00	97.00	25.00
Δ	+4.02	+1.43	+0.30	-0.20	-20.67	-1.00	-1.00	-25.00

Table 3: Safety-related evaluation results of the trained LLMs. **ASR** denotes the attack success rate.

Model	MMLU	Hellaswag	Natural Questions	GSM8K	HumanEval	C-Eval
LLAMA-CHAT 7B	42.93	77.00	20.94	22.97	9.15	33.21
+ DPO w. C ² -SYN	43.62	77.00	20.94	21.83	9.20	32.94
QWEN2 7B	70.32	82.00	21.50	74.07	40.24	73.25
+ DPO w. C ² -SYN	70.26	82.00	20.64	73.54	41.46	73.40

Table 4: General performance evaluation results of the trained LLMs.

Results

RQ2: Does learning to course-correct degrade overall performance?

Model	C ² -EVAL		Safety		Jailbreak Attack (ASR ↓)			
	Corr@10	Corr _{mean}	TruthfulQA (↑)	ToxiGen (↓)	GCG	PAIR	AutoDAN	CipherChat
LLAMA-CHAT 7B	66.60	61.63	48.60	51.27	70.95	10.00	54.00	75.00
+ DPO w. C ² -SYN	90.85	83.49	49.06	48.08	38.57	8.00	52.00	50.00
Δ	+24.25	+21.86	+0.46	-3.19	-32.38	-2.00	-2.00	-25.00
QWEN2 7B	85.40	85.47	62.35	52.97	66.67	26.00	98.00	50.00
+ DPO w. C ² -SYN	89.42	86.90	62.65	52.77	46.00	25.00	97.00	25.00
Δ	+4.02	+1.43	+0.30	-0.20	-20.67	-1.00	-1.00	-25.00

Table 3: Safety-related evaluation results of the trained LLMs. **ASR** denotes the attack success rate.

Model	MMLU	Hellaswag	Natural Questions	GSM8K	HumanEval	C-Eval
LLAMA-CHAT 7B	42.93	77.00	20.94	22.97	9.15	33.21
+ DPO w. C ² -SYN	43.62	77.00	20.94	21.83	9.20	32.94
QWEN2 7B	70.32	82.00	21.50	74.07	40.24	73.25
+ DPO w. C ² -SYN	70.26	82.00	20.64	73.54	41.46	73.40

Table 4: General performance evaluation results of the trained LLMs.



Results

RQ3: Does learning to course-correct enhance LLMs' resilience to jailbreak attacks?

Model	C ² -EVAL		Safety		Jailbreak Attack (ASR ↓)			
	Corr@10	Corr _{mean}	TruthfulQA (↑)	ToxiGen (↓)	GCG	PAIR	AutoDAN	CipherChat
LLAMA-CHAT 7B	66.60	61.63	48.60	51.27	70.95	10.00	54.00	75.00
+ DPO w. C ² -SYN	90.85	83.49	49.06	48.08	38.57	8.00	52.00	50.00
Δ	+24.25	+21.86	+0.46	-3.19	-32.38	-2.00	-2.00	-25.00
QWEN2 7B	85.40	85.47	62.35	52.97	66.67	26.00	98.00	50.00
+ DPO w. C ² -SYN	89.42	86.90	62.65	52.77	46.00	25.00	97.00	25.00
Δ	+4.02	+1.43	+0.30	-0.20	-20.67	-1.00	-1.00	-25.00

Table 3: Safety-related evaluation results of the trained LLMs. **ASR** denotes the attack success rate.

Model	MMLU	Hellaswag	Natural Questions	GSM8K	HumanEval	C-Eval
LLAMA-CHAT 7B	42.93	77.00	20.94	22.97	9.15	33.21
+ DPO w. C ² -SYN	43.62	77.00	20.94	21.83	9.20	32.94
QWEN2 7B	70.32	82.00	21.50	74.07	40.24	73.25
+ DPO w. C ² -SYN	70.26	82.00	20.64	73.54	41.46	73.40

Table 4: General performance evaluation results of the trained LLMs.

Results

RQ4: How well does C²-SYN transfer to improve out-of-distribution (OOD) LLMs?

Model	C ² -EVAL		Safety		Jailbreak Attack (ASR ↓)			
	Corr@10	Corr _{mean}	TruthfulQA (↑)	ToxiGen (↓)	GCG	PAIR	AutoDAN	CipherChat
LLAMA-CHAT 7B	66.60	61.63	48.60	51.27	70.95	10.00	54.00	75.00
+ DPO w. C ² -SYN	90.85	83.49	49.06	48.08	38.57	8.00	52.00	50.00
Δ	+24.25	+21.86	+0.46	-3.19	-32.38	-2.00	-2.00	-25.00
QWEN2 7B	85.40	85.47	62.35	52.97	66.67	26.00	98.00	50.00
+ DPO w. C ² -SYN	89.42	86.90	62.65	52.77	46.00	25.00	97.00	25.00
Δ	+4.02	+1.43	+0.30	-0.20	-20.67	-1.00	-1.00	-25.00

Table 3: Safety-related evaluation results of the trained LLMs. **ASR** denotes the attack success rate.

Model	MMLU	Hellaswag	Natural Questions	GSM8K	HumanEval	C-Eval
LLAMA-CHAT 7B	42.93	77.00	20.94	22.97	9.15	33.21
+ DPO w. C ² -SYN	43.62	77.00	20.94	21.83	9.20	32.94
QWEN2 7B	70.32	82.00	21.50	74.07	40.24	73.25
+ DPO w. C ² -SYN	70.26	82.00	20.64	73.54	41.46	73.40

Table 4: General performance evaluation results of the trained LLMs.

Results

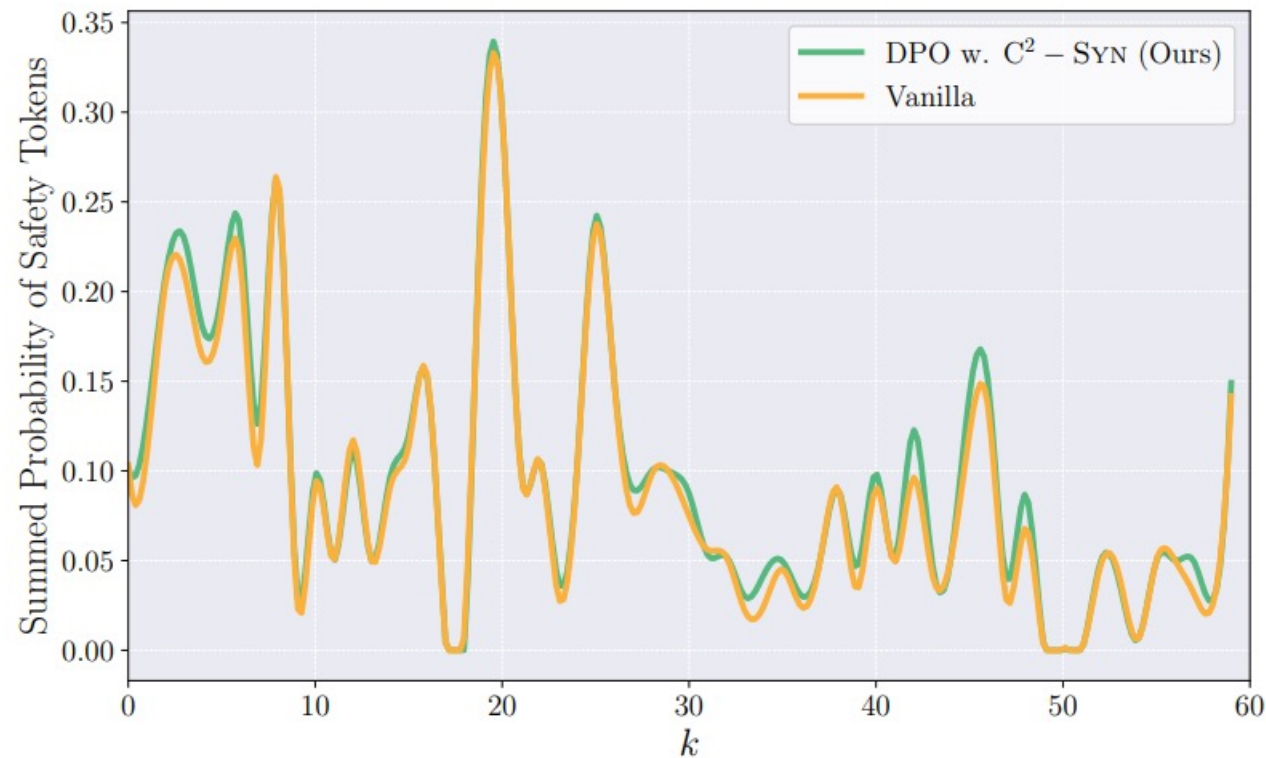


Figure 5: Summed probability of safety tokens at the *first* decoding position after an **IHR** of length k .

Conclusion

Contributions:

We systematically investigate the problem of course-correction in the context of harmful content generation within LLMs:

- We develop **C²-EVAL** and evaluate ten prevalent LLMs
- We construct **C²-SYN** and use DPO on two LLMs
- Results demonstrate that preference learning with our synthetic data can improve two models' overall safety without harming general performance.

Limitations

- Dataset Bias
- Evaluation Method
- Training Algorithm Selection
- Model Selection

THANK YOU