

Engaging Content, Not Clickbait:

Teaching Generative AI Why It Works — Not Just What Works

K. Sudhir

James Frank Professor of Marketing

Yale School of Management


MSI Forum, UCLA

February 2026

The Content Generation Problem

Content that invites to engage further: click, purchase etc

- **Text:** Title and headlines for articles, subject lines for email
- **Images:** Book covers, product packaging visuals
- **Videos:** Trailers for movies, games




Marketing teams routinely use GenAI to generate such content to optimize engagement

Zero Shot Prompting does not typically work


LLMs need to be aligned...with engagement (e.g., CTR)

Zero Shot Prompting of LLMs
(or Gen AI Model)



THE ALIGNMENT PROBLEM


How Can Artificial Intelligence Learn Human Values?



Standard Solution:


Align LLM to A/B Test results by Finetuning

Align LLMs with Finetuning (RLHF, DPO) using A/B test results, then prompt



THE ALIGNMENT PROBLEM

How Can Artificial Intelligence Learn Human Values?



Data

Upworthy A/B Test Data of Headlines


Test ID	Headline	Impressions	Clicks	CTR (%)
1	\$3 Million Is What It Takes For A State To Legally Kill Someone	3,017	19	0.63
1	The Fact That Sometimes Innocent People Are Executed Is Enough To End The Death Penalty. But This?	2,974	26	0.87

Same article, randomized A/B test → 38% higher CTR from framing, not content.

- 150,817 tested packages from 32,487 A/B tests
- We supplement the data by crawling the article contents associated with the headlines

But finetuning with (“What Works”)-> reward hacking

We get headlines with click-bait style words and phrases


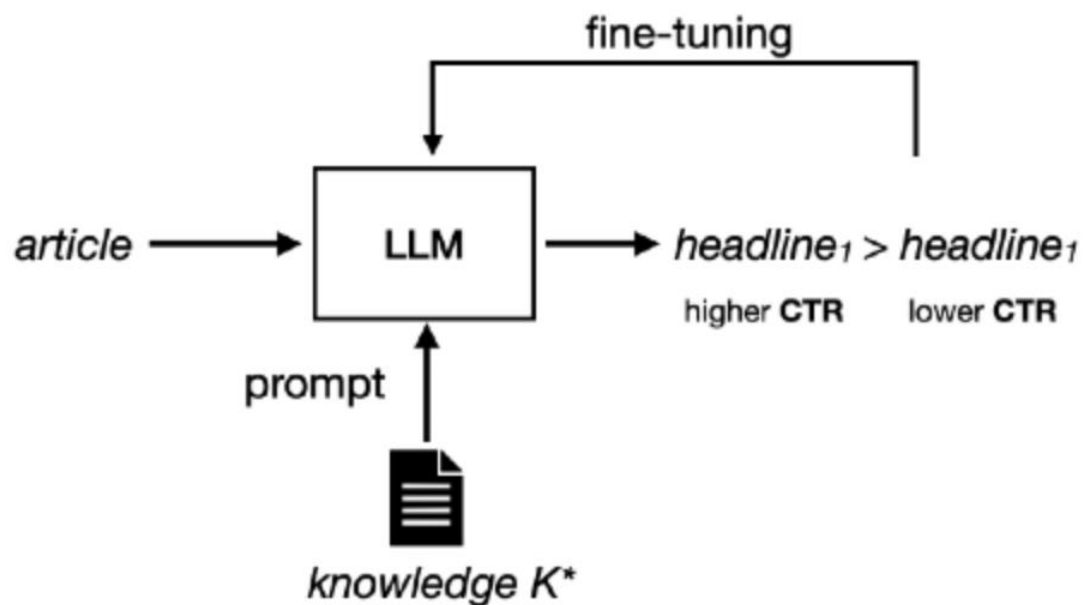


Vanilla DPO increases clickbait phrases 10–25x

PROPOSED SOLUTION

Knowledge/Theory-Guided Finetuning

Teaching AI- about “why it works”, not just “what works”



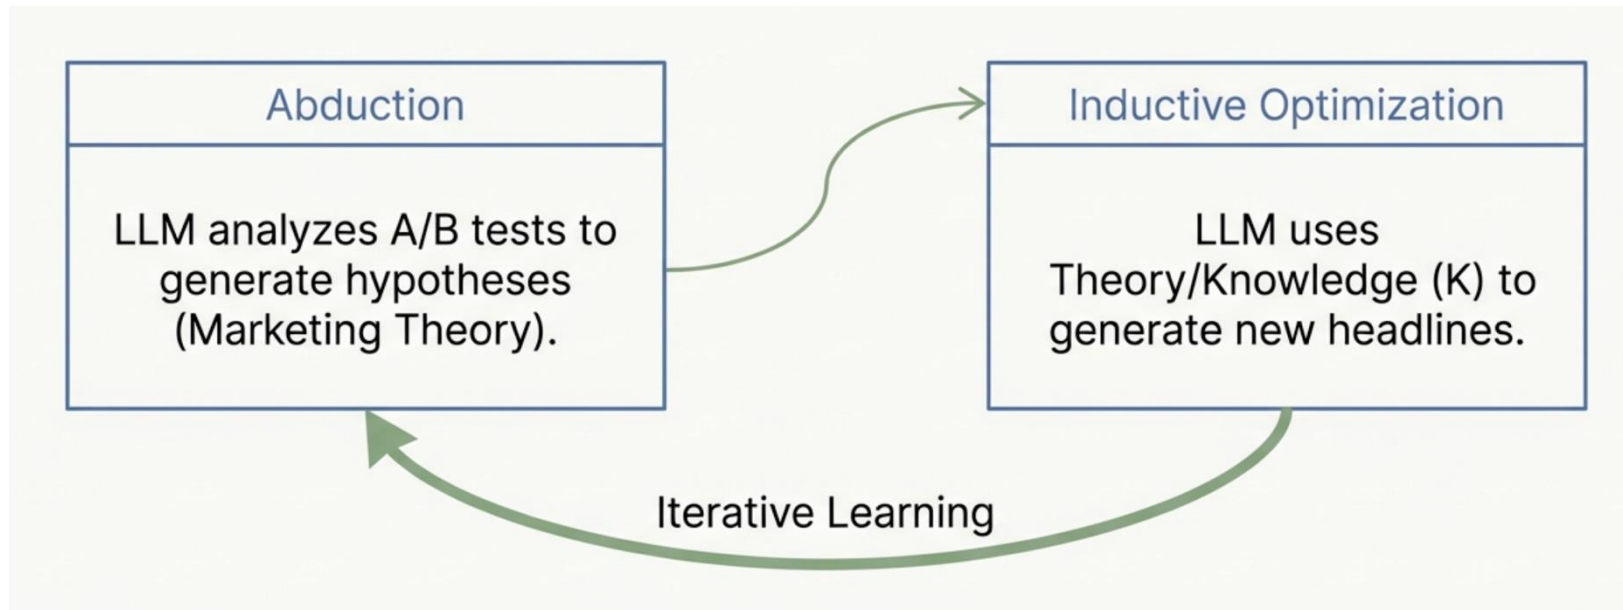
Knowledge/Theory as “regularizer” can be a very general idea...

Question: How to get “knowledge?”

From journals, managerial experience

But what if we don't have clear knowledge?

We propose leveraging A/B tests to generate “knowledge”



Abduction Step: We ask LLM

Why is headline 1 better than headline 2?

Higher CTR

The Fact That Sometimes Innocent People Are Executed Is Enough To End The Death Penalty. But This?

Lower CTR

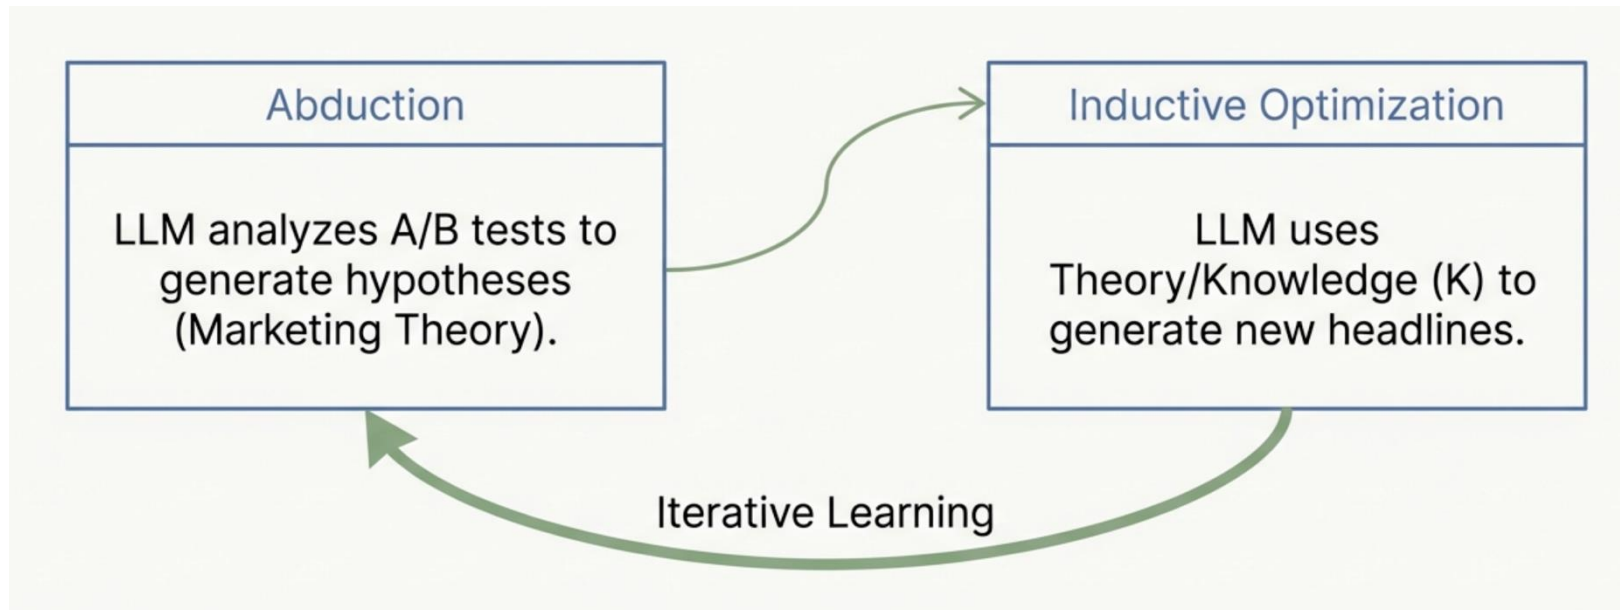
Lets See Hire Cops, Pay Teachers, Buy Books For Schools. Or Kill People. Hard Choice, Right?

OpenAI o3 Hypotheses

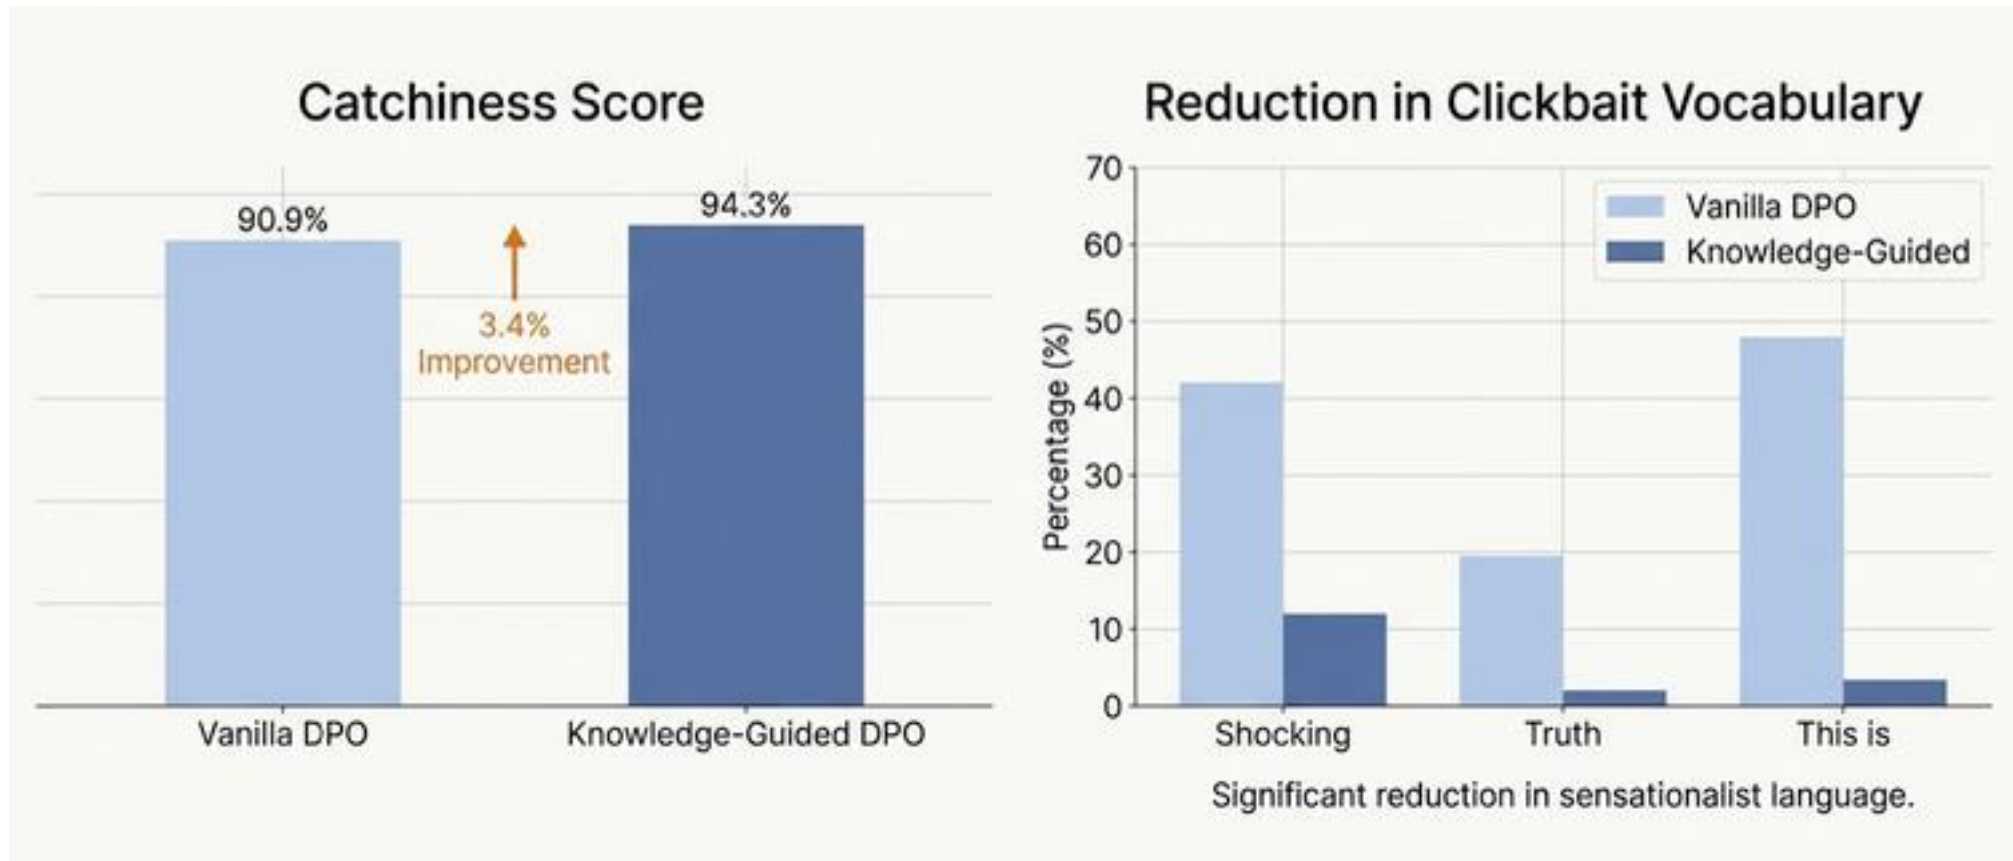
- Leave a deliberate information gap that promises a dramatic reveal
- Anchor the message around one vivid, high-stakes scenario instead of scattering attention across many topics
- ...

Not all of the hypothesis will be valid

We then validate the hypothesis in an iterative inductive optimization step



RESULTS: Higher Performance, but with Ethical Guardrails

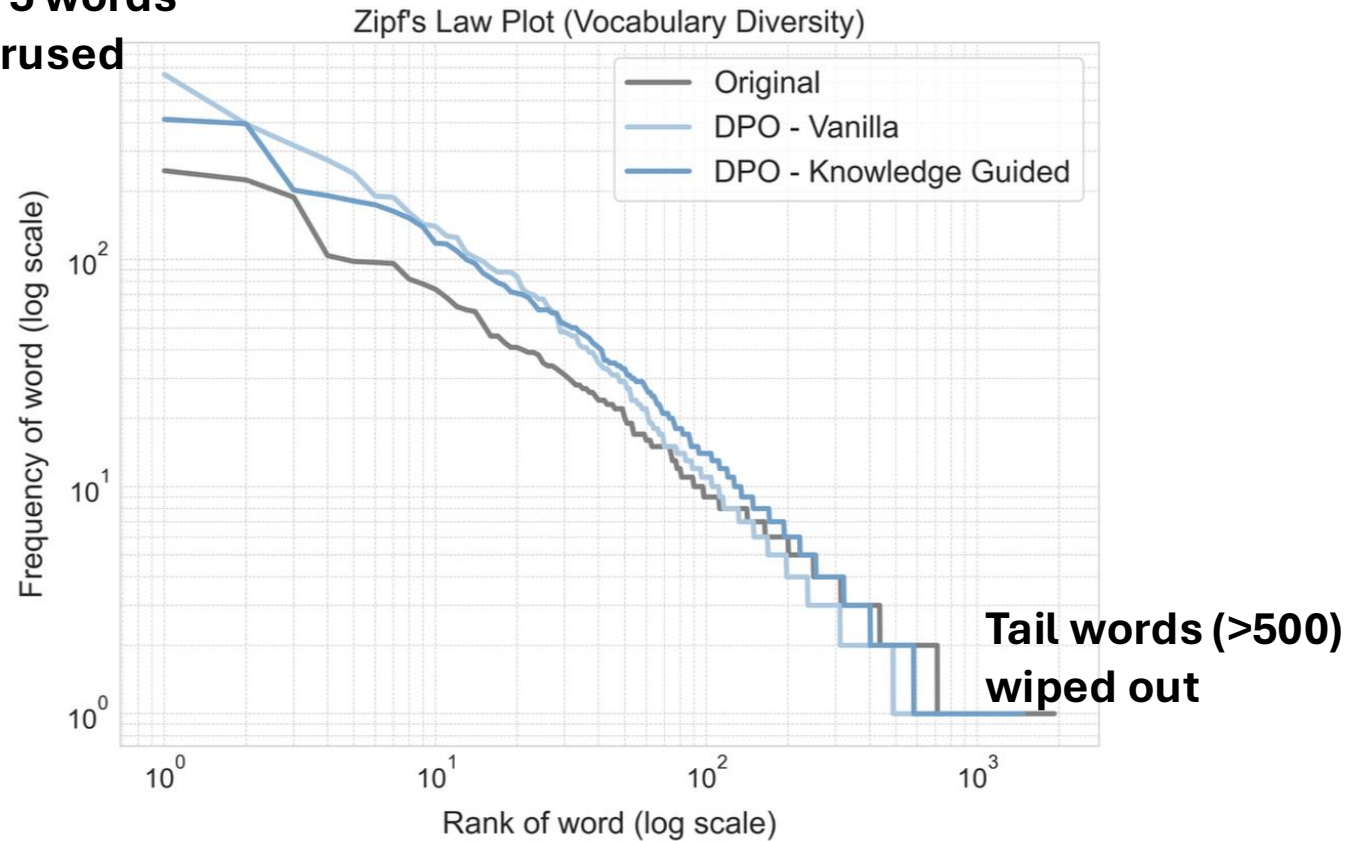


Note: Baseline LLM has very low catchiness scores: 12-25%

RESULTS:

Using theory maintains lexical diversity

**Top 5 words
overused**

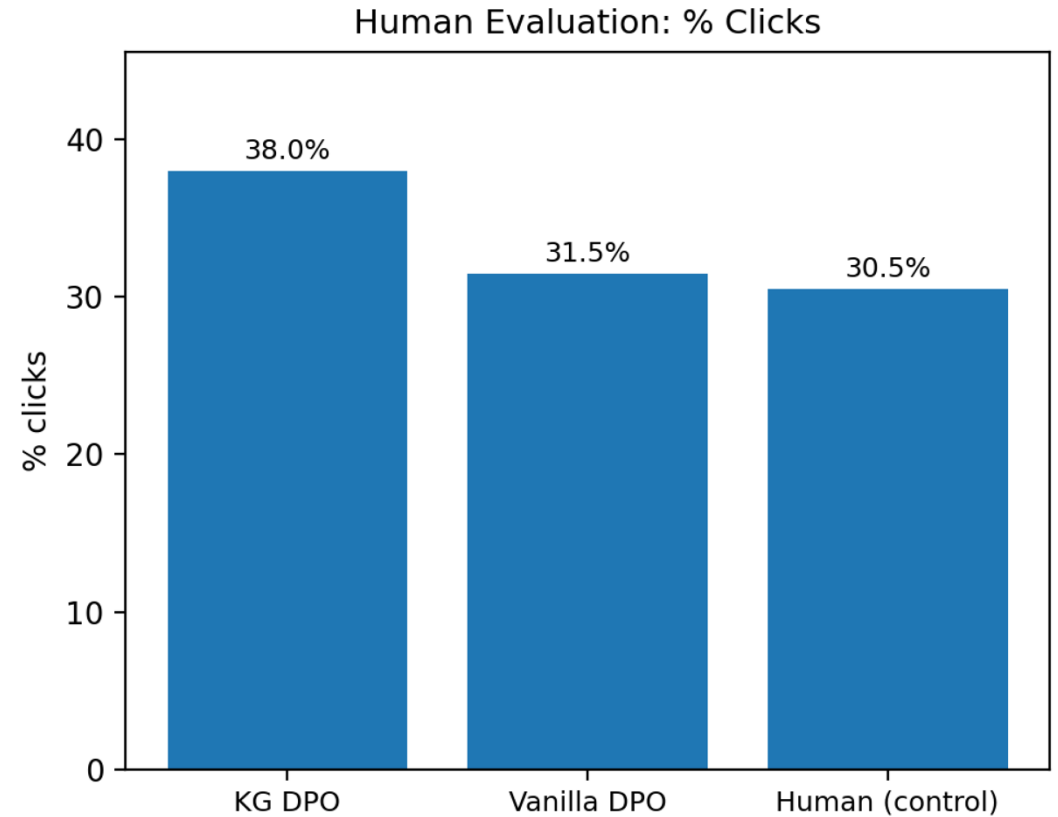
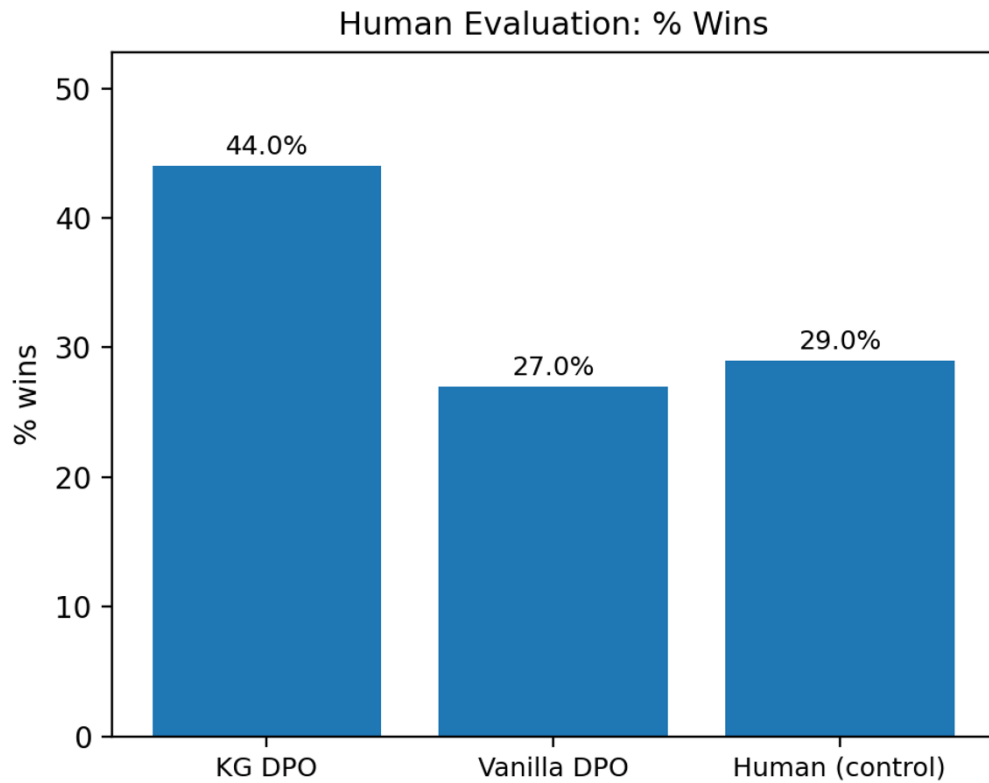


Vanilla Alignment collapses vocabulary

Knowledge guided alignment preserves richness

RESULTS:


Human Evaluation of Headlines



Knowledge Guided DPO wins handily

RESULTS:

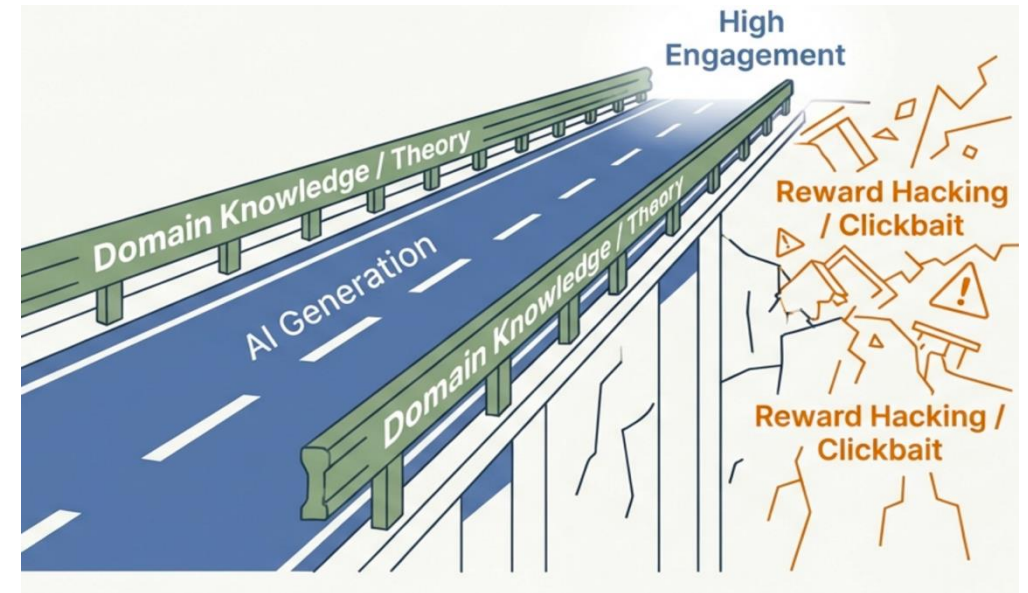
Theory makes learning efficient with less data



Much higher engagement performance with just 25% of the A/B tests

Takeaway

- Knowledge improves
 - effectiveness
 - efficiency
 - Serves as guardrail against reward-hacks



Knowledge/Theory Guided can be applied broadly

- Applies in many other settings. e.g., sales, service, negotiations
 - In companion work, theory-based guidance lifts performance of weaker LLM—based interactive chatbots to frontier-level performance
- Knowledge is textual
 - If don't have A/B tests, you could use knowledge from Journal results or managerial experience

Thank you!

k.sudhir@yale.edu

