

Intelligent Information Gathering with LLMs and Bayesian Experimental Design

Tom Rainforth


ML Nosh

11/03/26




🔗 Model: GPT-4

T I am thinking of a famous historical person. Please ask me yes or no questions to figure out who it is. I will provide an answer after each individual question and you should adapt your subsequent questions to figure out who I am thinking of as quickly as possible

 Sure, let's get started.

1. Is the person you're thinking of male?

T No

 Okay, that narrows it down a bit.

2. Is this person from the 20th century?

T No



Alright.

11. Did this person live in the 17th century?



No



Alright.

12. Did this person live in the 18th century?



No



This means the person you're thinking of lived between the 15th and 19th centuries.

13. Did this person live in the 21st century?



No

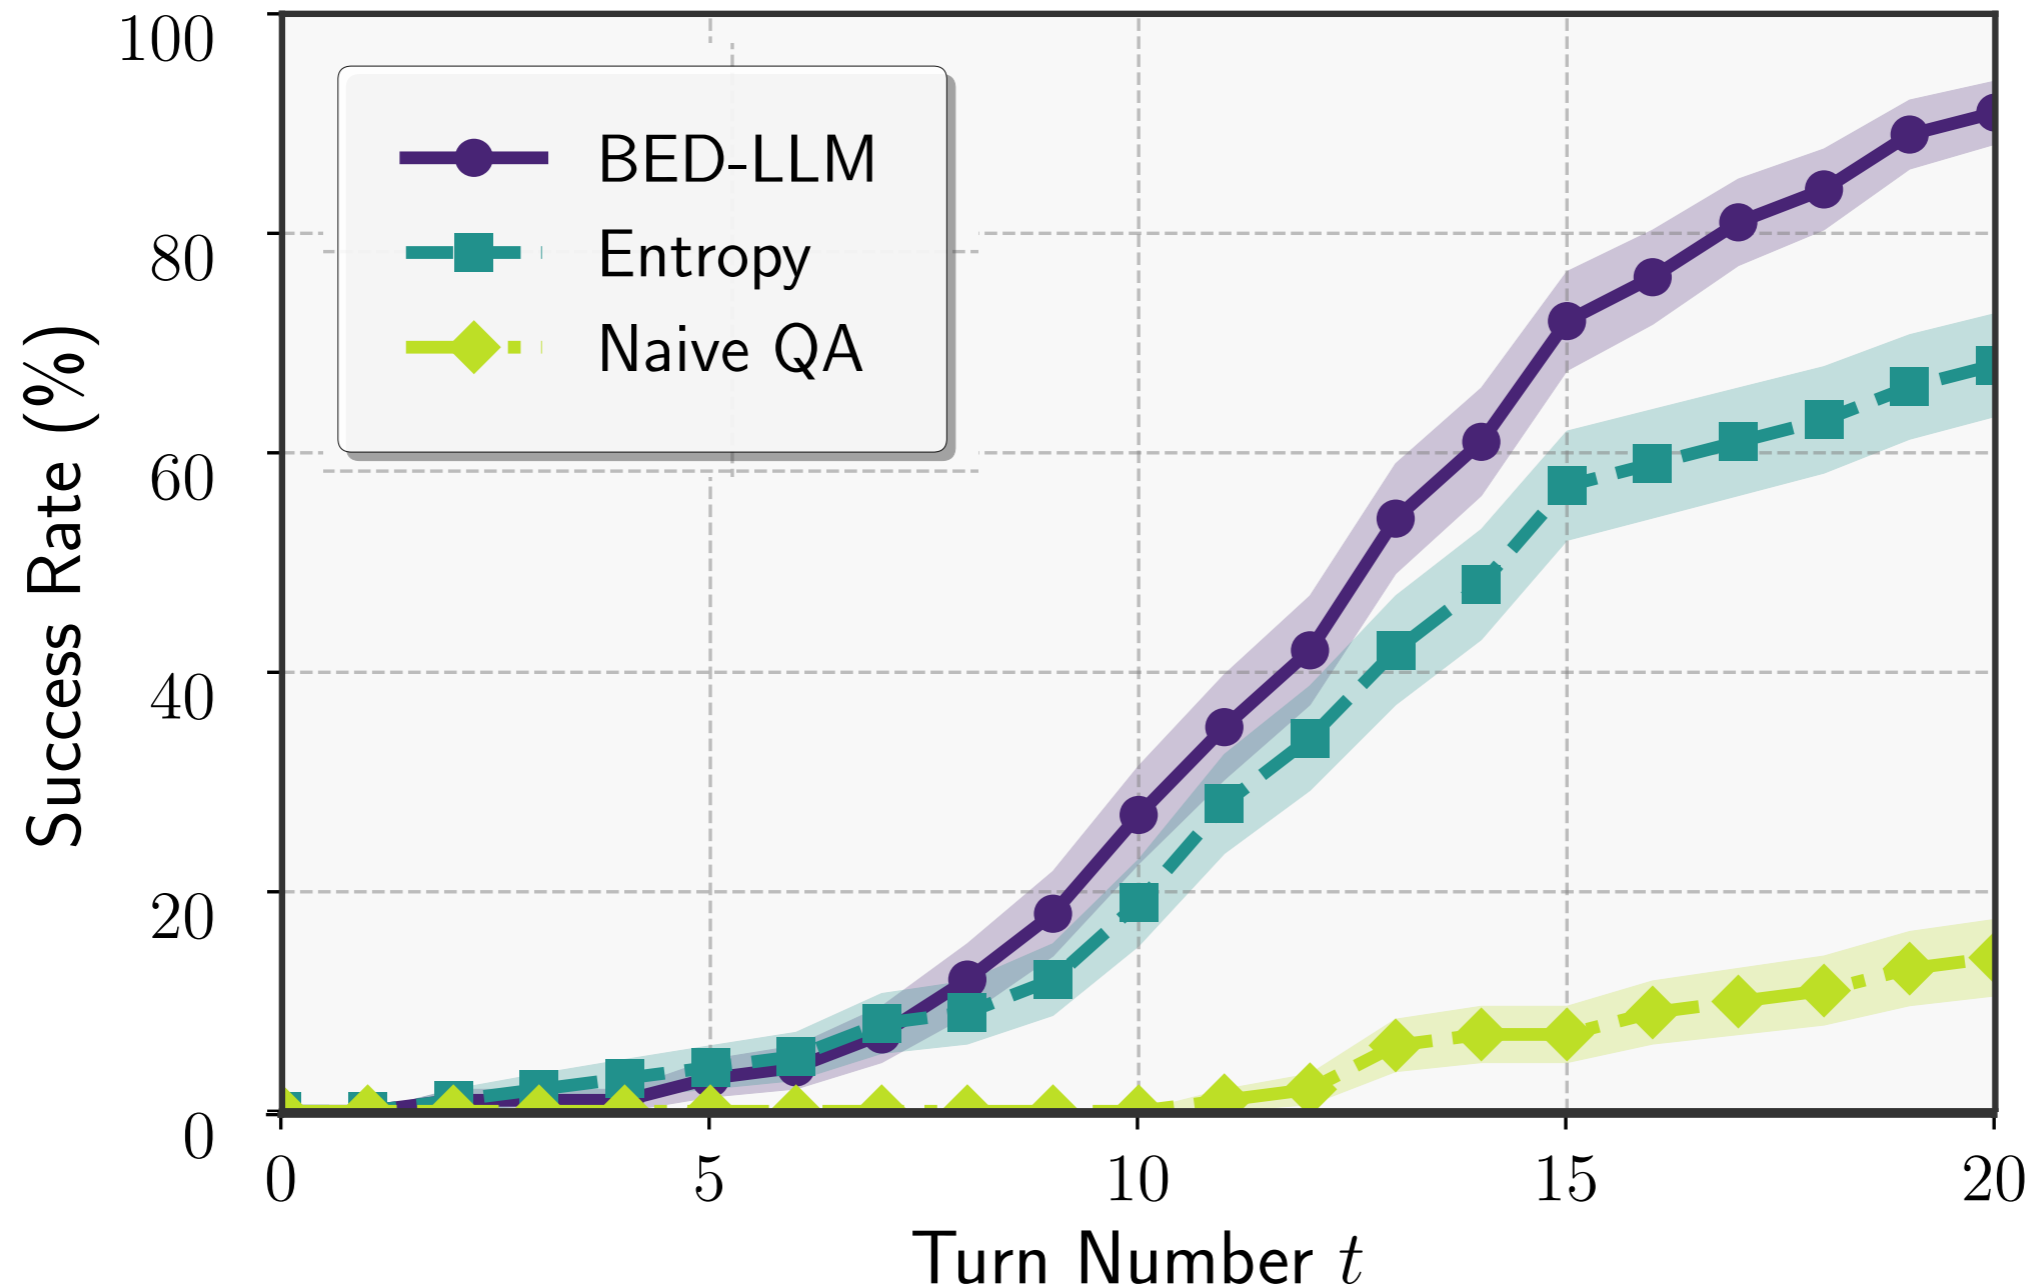


Alright, this indicates they lived in the 20th century.

14. Was this person involved in any independence movements or revolutions?



20 Questions Celebrities with Mistral Large





Problem Formulation

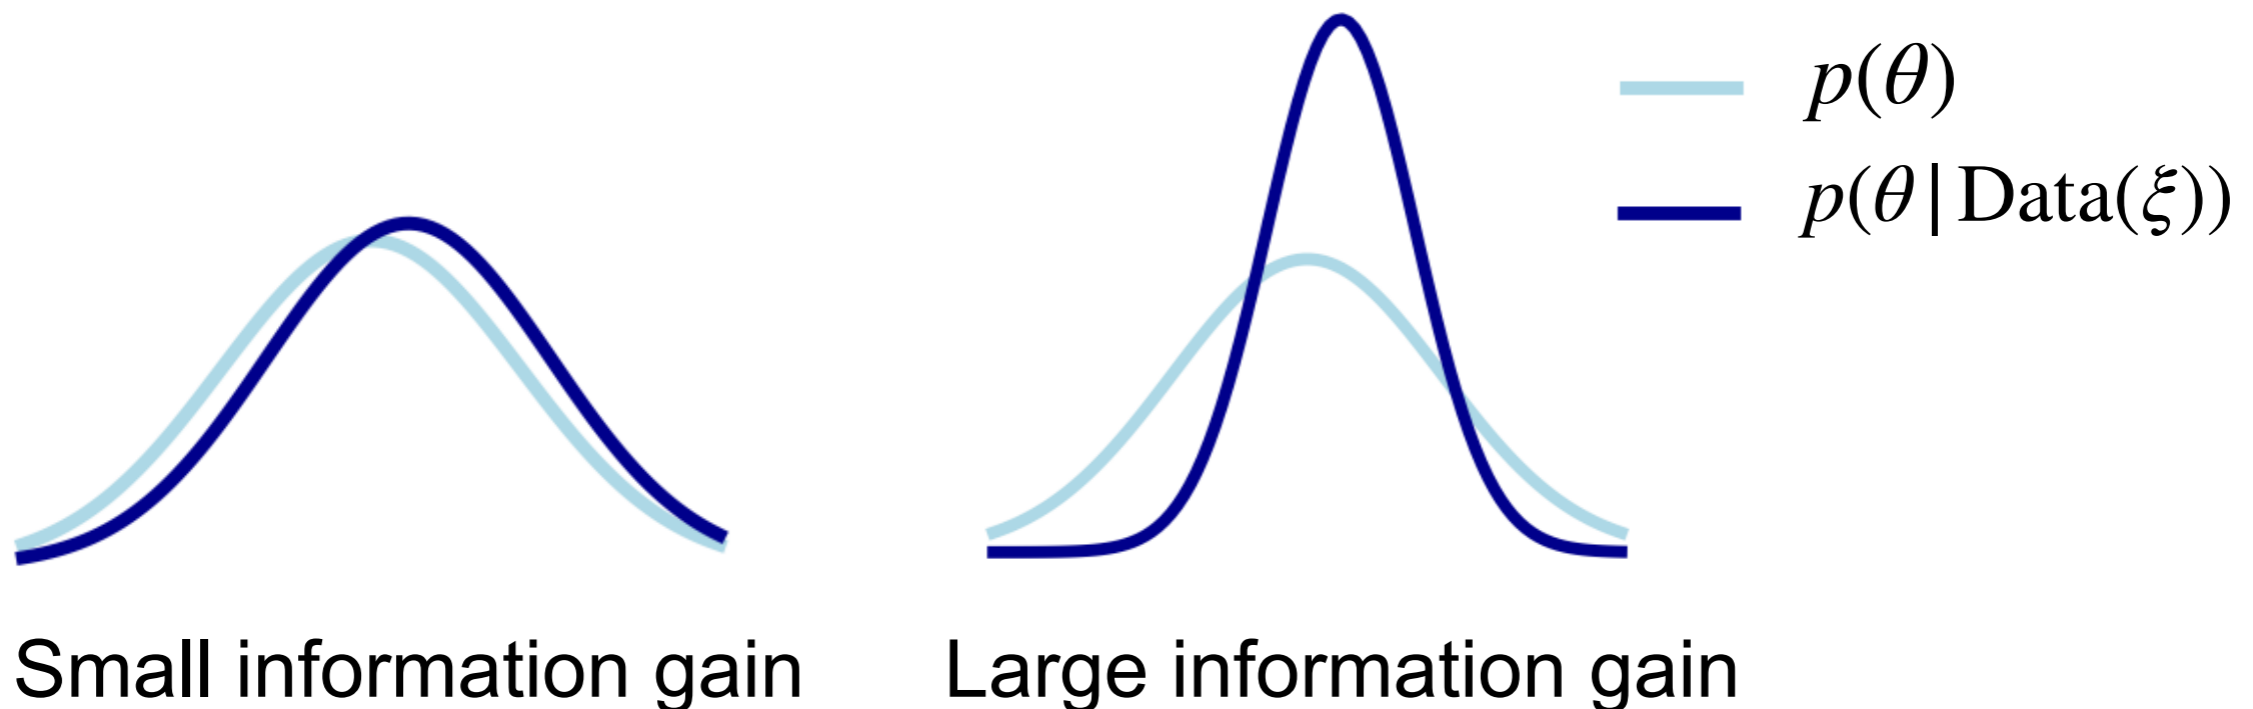
- Target to learn about: θ
- Can refine beliefs in θ by making queries x and receiving responses y
- Multiple queries x_1, x_2, \dots, x_T allow for adaptation where each x_t is chosen given a history $h_t := (x_i, y_i)_{i=1:t}$



- We choose queries to maximise the **expected information gain** over possible data

$$\begin{aligned} \text{EIG}_{\theta}(x) &= H[p(\theta)] - \mathbb{E}_{p(y;x)} [H[p(\theta | y; x)]] \\ &= H[p(y; x)] - \mathbb{E}_{p(\theta)} [H[p(y | \theta; x)]] \end{aligned}$$

- The EIG is fully defined by a joint generative model $p(\theta, y; x)$ – we derive this from the LLM





EIG is Not Predictive Entropy!

Question 1

Question:

Which ice cream flavor feels like the best match for this user?

Choose one option:

- A. Vanilla Bean
- B. Rich Dark Chocolate
- C. Strawberry Swirl
- D. Mint Chocolate Chip

Predictive Entropy: Very High
EIG: 0

Question 2

Question:

Which film genre does the user most prefer?

Choose one option:

- A. Action
- B. Sci-Fi
- C. Comedy
- D. Horror

Predictive Entropy: High
EIG: High



BED-LLM

1. Generate diverse set of candidate queries
2. Estimate ELG for each
3. Select query with largest estimate
4. Update joint model with new query–response pair and return to Step 1



20 Questions with BED-LLM

History: Q1 - Was this person born in the 20th century? A1 - Yes
 Q2 - Is this person male? A2 - Yes

Sample θ given history

Barack Obama, Steve Irwin, Frederik X, Banksy, Elvis Presley, Bill Gates, David Hasselhoff, Stephen Hawking, Rafael Nadal, David Beckham, ...

Generate candidate questions:

CQ1 - Was this person born in Antarctica?
CQ2 - Was this person born in the 19th century?
CQ3 - Is this person European?
CQ4 - Is this person an artist?
CQ5 - Does this person prefer thrash metal or death metal?

Pick best and evaluate

Q3 - Is this person European?
A3 - No

Estimate EIG for CQs using θ samples

CQ1 EIG = 0.0000001
CQ2 EIG = 0
CQ3 EIG = 0.8
CQ4 EIG = 0.15
CQ2 EIG = 0.01



BED-LLM Joint

- We directly use the LLM's distribution for the likelihood, but separately derive the prior

$$p(\theta, y_{t+1}; x_{t+1}, h_t) = p_f(\theta; h_t) p_{\text{LLM}}(y_{t+1}; [\theta, x_{t+1}])$$

- $p_f(\theta; h_t)$ uses a combination of in context learning and rejection sampling to derive beliefs



Problems with $p_{\text{LLM}}(\theta; h_t)$

ICL typically fails to properly account for the history



Alright.

11. Did this person live in the 17th century?



No



Alright.

12. Did this person live in the 18th century?



No



This means the person you're thinking of lived between the 15th and 19th centuries.

13. Did this person live in the 21st century?



No



Alright, this indicates they lived in the 20th century.

14. Was this person involved in any independence movements or revolutions?



Problems with $p_{\text{LLM}}(\theta; h_t)$

The LLM beliefs tend to be over-collapsed to a small number of possibilities

Is this person known for their contributions to science?
No.

Is this person known for their contributions to the arts?
Yes.

```
{"Vincent van Gogh": 93,  
"Georgia O Keeffe": 10,  
"Frida Kahlo": 37,  
"Salvador Dali": 44,  
"August Wilson": 8,  
"Auguste Rodin": 8}
```



Definition of $p_f(\theta; h_t)$

- Initially draw samples $\hat{\theta}$ from $p_{\text{LLM}}(\theta; h_t)$ with tricks to improve diversity (careful prompting, co-sample, high temperature)
- Discard duplicates
- Check samples against the history using $p_{\text{LLM}}(y_i; [\hat{\theta}, x_i])$ for $i = 1, \dots, t$ and reject incompatible ones



		Success Rate (%)							
	Model	Prompt-Only	Split	CoT	BED-LLM	Entropy	Data-Est.	ICL Beliefs	Impl. Max.
Animals	GPT-4o-mini	44±5.0	78±4.2	55±5.0	88 ±3.3	79±4.1	64±4.8	18±3.9	47±5.0
	GPT-4o	45±5.0	83±3.8	62±4.9	93 ±2.6	88±3.3	70±4.6	25±4.4	70±4.6
	Llama-3.1-8B	8±2.7	49±5.0	19±3.9	63 ±4.9	54±5.0	38±4.9	25±4.4	16±3.7
	Llama-3.3-70B	40±4.9	65±4.8	40±4.9	79 ±4.1	68±4.7	40±4.9	33±4.7	54±5.0
	Mistral-Large	33±4.7	85±3.6	35±4.8	95 ±2.2	83±3.8	83±3.8	53±5.0	53±5.0
	Qwen2.5-72B	45±5.0	87±3.4	51±5.0	95 ±2.2	85±3.6	68±4.7	46±5.0	61±4.9
Celebrities	GPT-4o-mini	30±4.6	53±5.0	42±5.0	72 ±4.5	55±5.0	32±4.7	16±3.7	31±4.7
	GPT-4o	45±5.0	63±4.9	63±4.9	86 ±3.5	64±4.8	55±5.0	52±5.0	50±5.0
	Llama-3.1-8B	10±3.0	35±4.8	16±3.7	58 ±5.0	36±4.8	19±3.9	24±4.3	19±3.9
	Llama-3.3-70B	33±4.7	43±5.0	36±4.8	55 ±5.0	46±5.0	26±4.4	27±4.5	37±4.9
	Mistral-Large	19±4.0	63±4.9	42±5.0	91 ±2.9	68±4.7	66±4.8	31±4.7	36±4.8
	Qwen2.5-72B	32±4.7	56±5.0	48±5.0	84 ±3.7	59±4.9	34±4.8	26±4.4	39±4.9
Things	GPT-4o-mini	26±4.4	38±4.9	33±4.7	49 ±5.0	37±4.9	26±4.4	19±4.0	25±4.4
	GPT-4o	34±4.8	40±4.9	49±5.0	64 ±4.8	49±5.0	26±4.4	19±3.9	42±5.0
	Llama-3.1-8B	10±3.0	12±3.3	10±3.0	26 ±4.4	15±3.6	9±2.9	11±3.1	10±3.0
	Llama-3.3-70B	34±4.8	46±5.0	35±4.8	55 ±5.0	48±5.0	19±3.9	15±3.6	34±4.8
	Mistral-Large	26±4.4	51±5.0	29±4.6	58 ±5.0	52±5.0	46±5.0	19±3.9	30±4.6
	Qwen2.5-72B	32±4.7	51±5.0	46±5.0	62 ±4.9	51±5.0	39±4.9	24±4.3	40±4.9



Film Recommendations

Example θ sample:

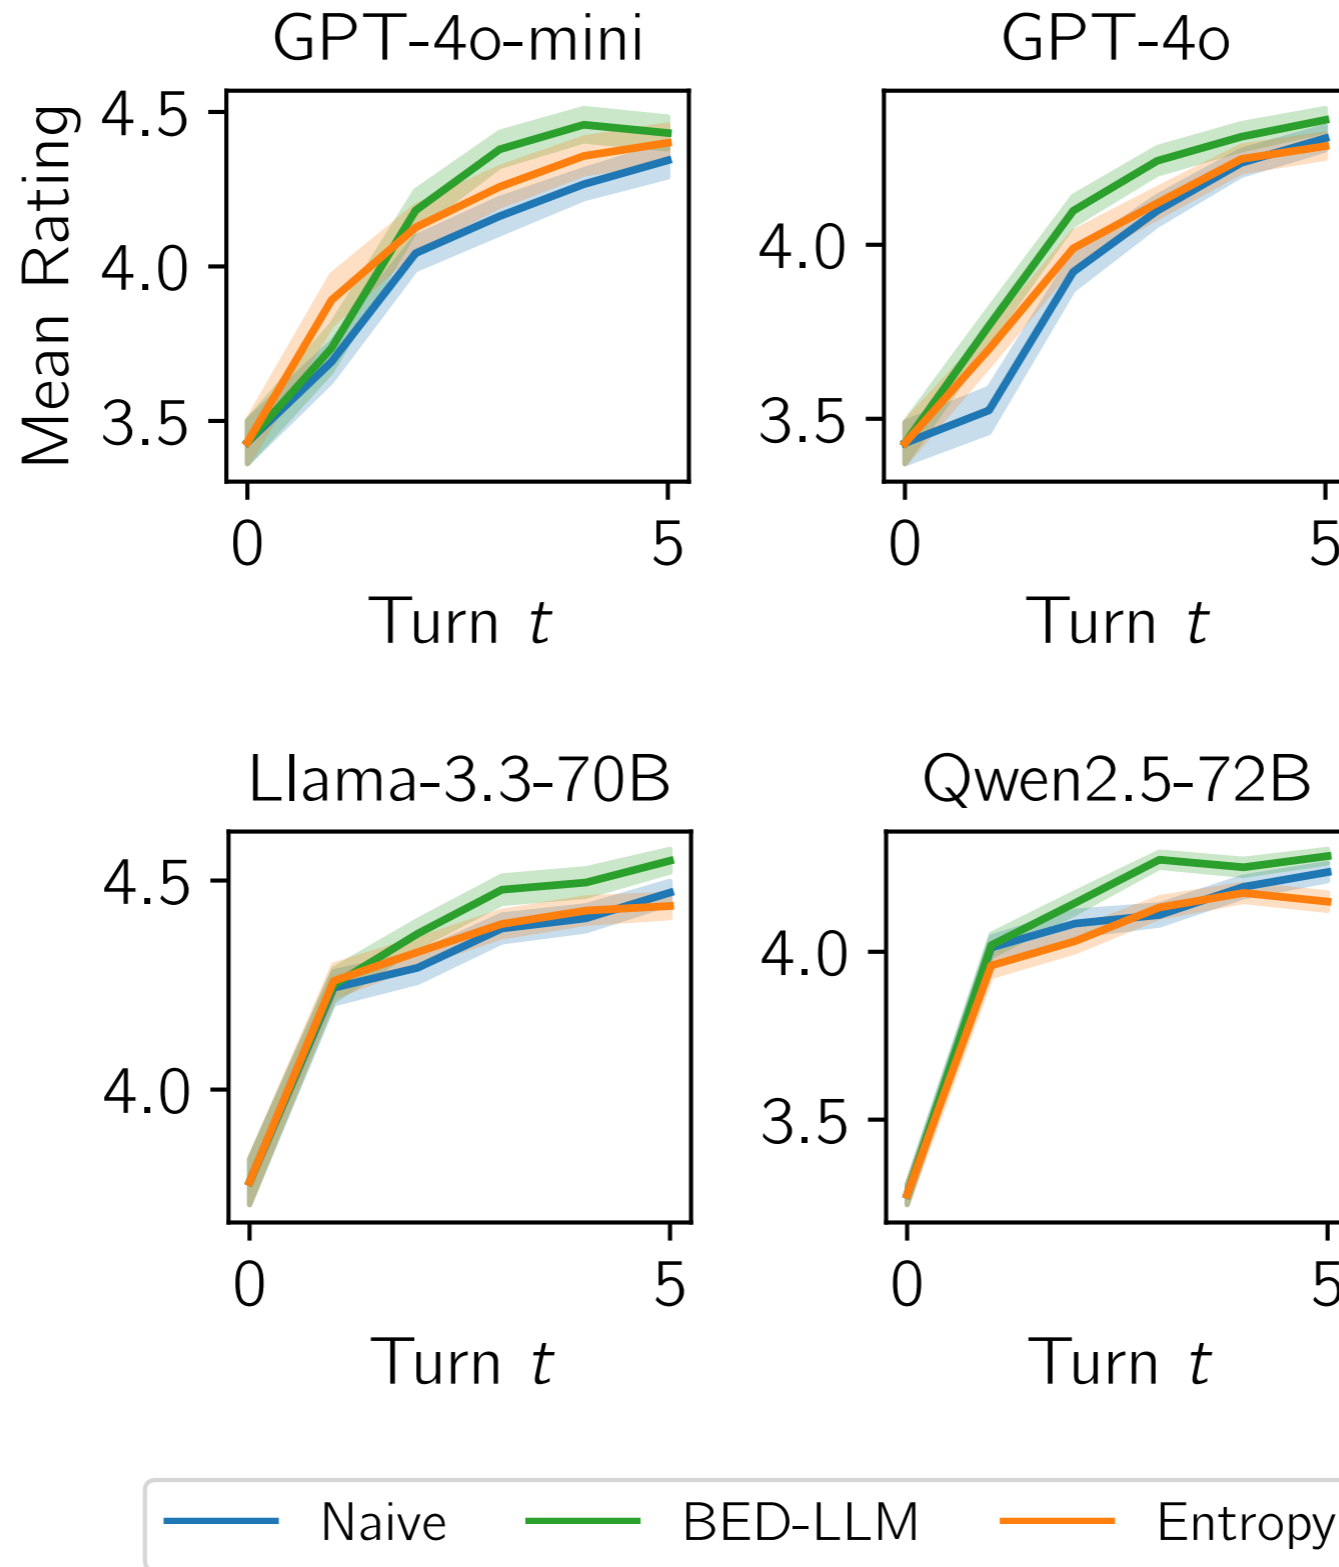
This film buff loves watching fast-paced, action-packed thrillers with complex characters, clever plot twists, and a sense of tension or suspense, often featuring a sense of excitement or adrenaline. They appreciate witty dialogue, memorable villains, and thrilling stunts, but dislike predictable plotlines, poorly developed heroes, and lack of originality or creativity. They prefer mainstream films with a moderate runtime and are not fond of subtitles or overly complex storytelling.

Question:

Which film genre does the user most prefer?

Choose one option:

- A. Action
- B. Sci-Fi
- C. Comedy
- D. Horror





Take Aways

- Bayesian experimental design can work at the LLM scale and significantly improve information/data gathering!
- We need to be careful though about how the joint model is set up and how we make updates
- We actually don't need good uncertainty estimates in θ if we can get them in y !



Thank You!

BED-LLM: Intelligent Information Gathering with LLMs and Bayesian Experimental Design. *ICLR 2026*



Deepro
Choudhury



Sinead
Williamson



Adam
Goliński



Ning
Miao



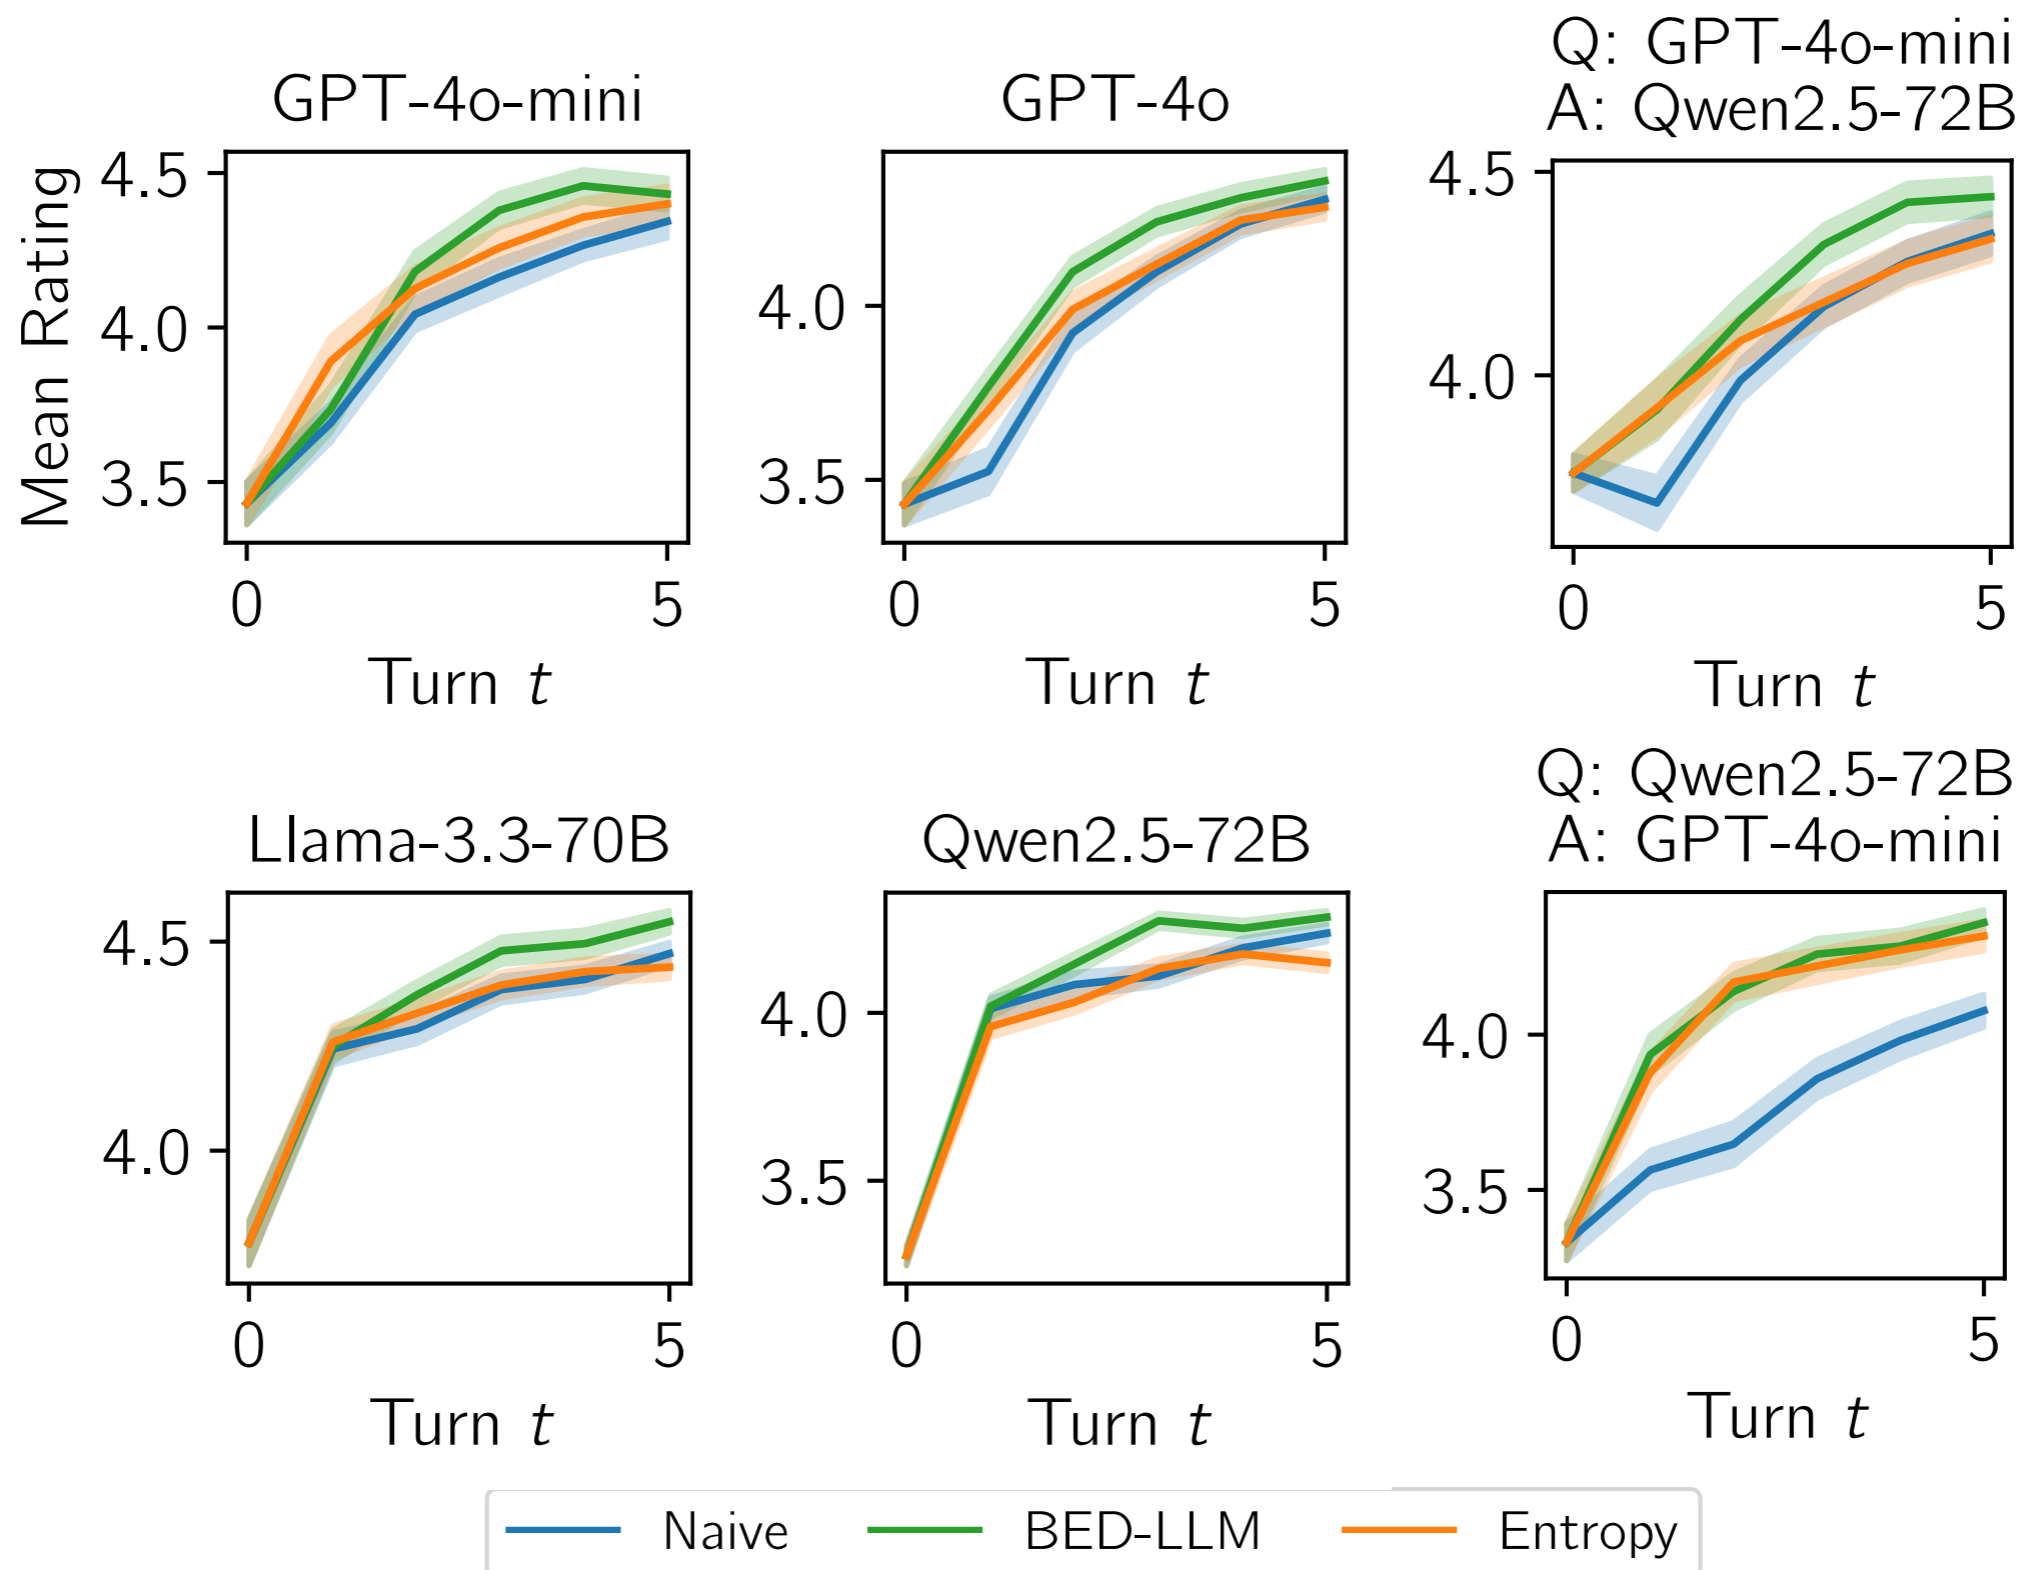
Freddie
Bickford Smith

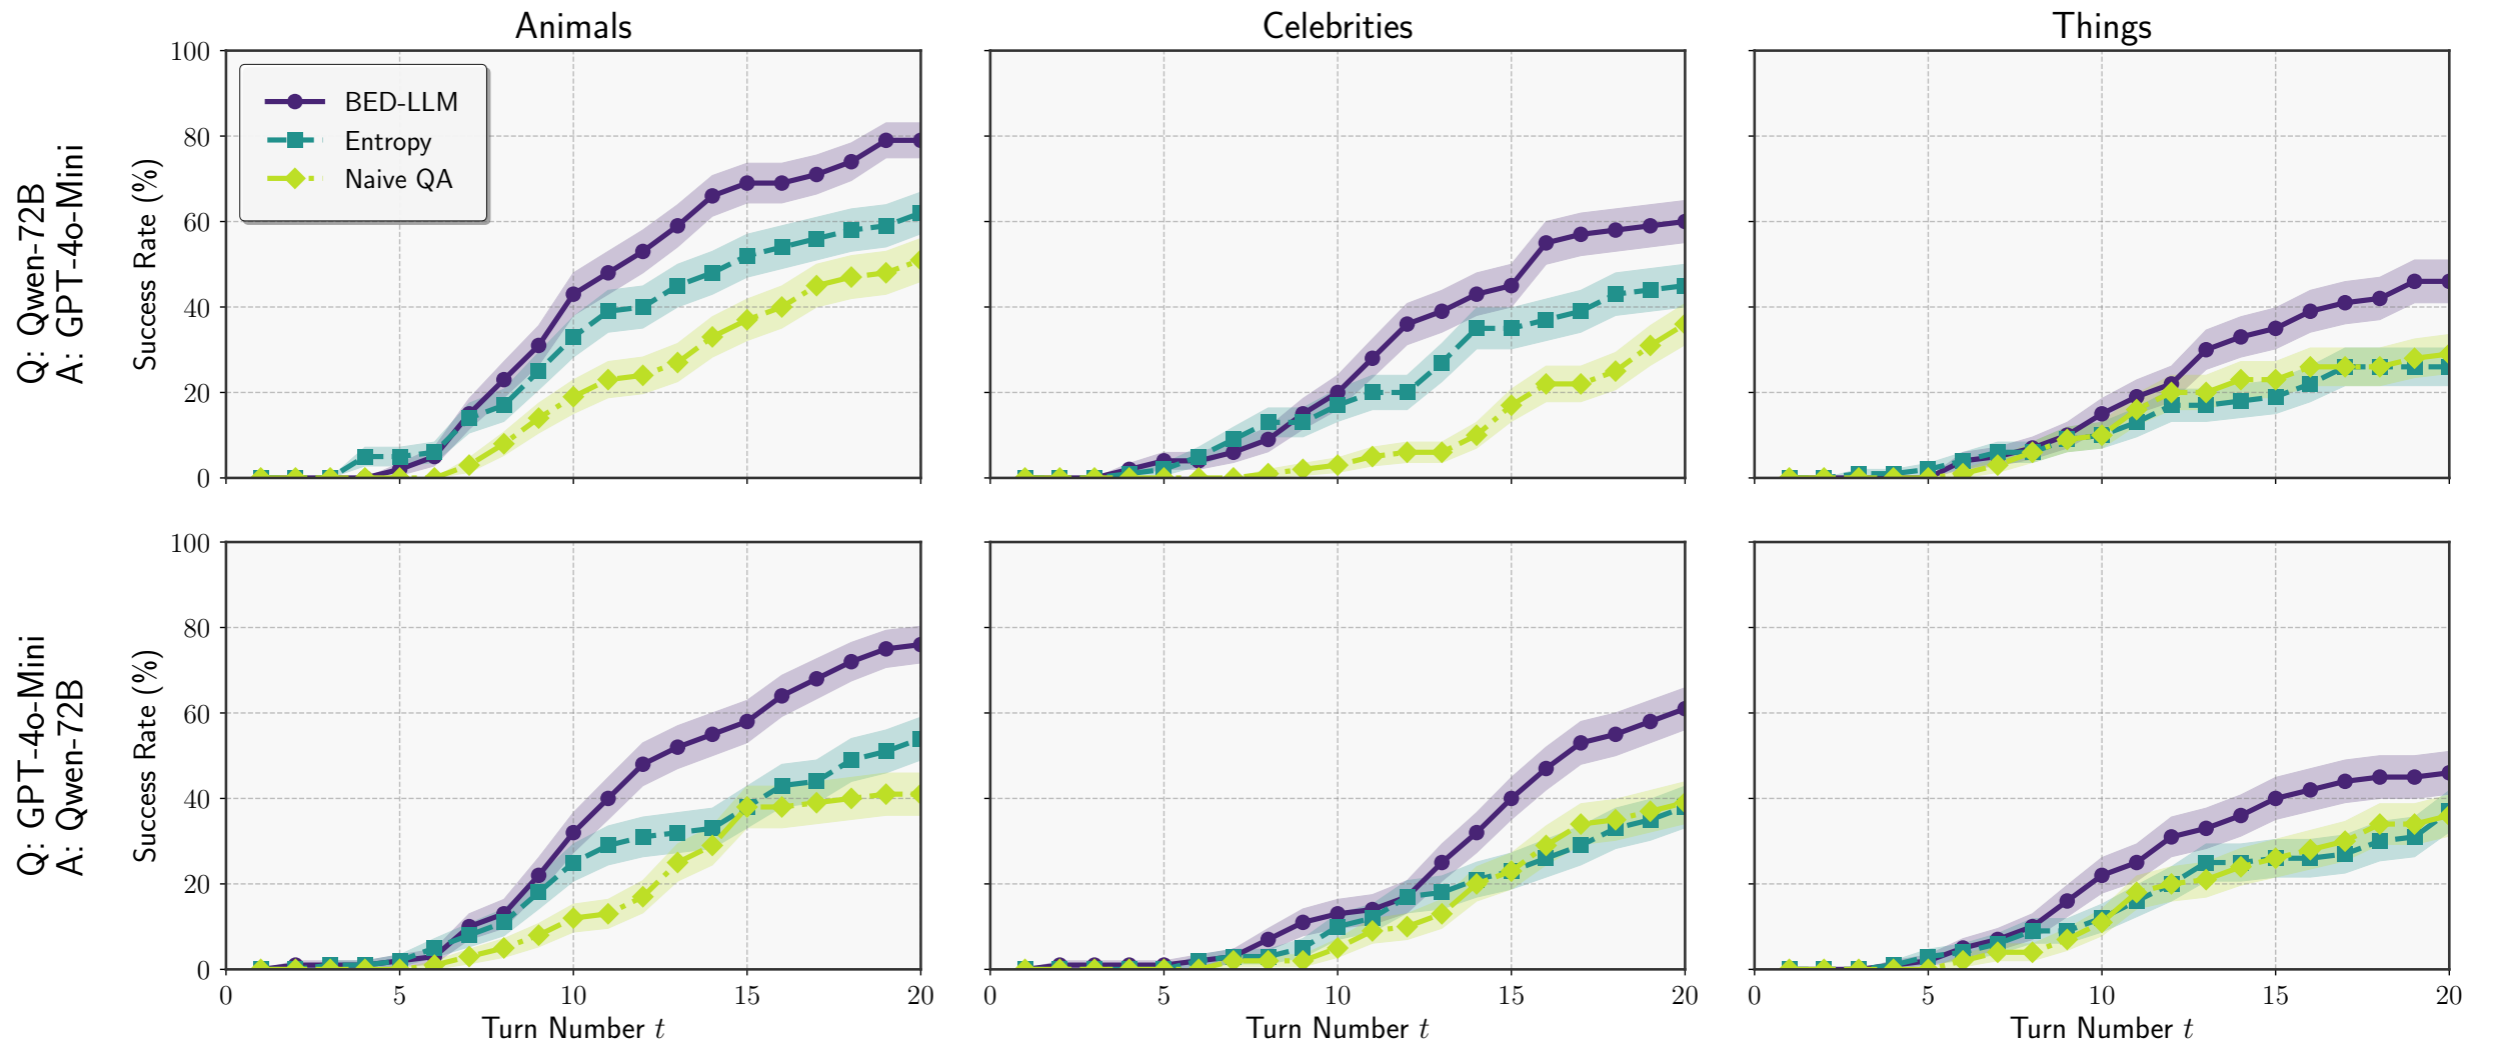


Michael
Kirchhof



Yizhe
Zhang







Joint Model Options

- Lots of ways we can derive a joint model from the LLM
- Most critical decision is between two distribution pairings:
 - Prior-likelihood: $p(\theta; h_t)p(y_{t+1} | \theta; x_{t+1}, h_t)$
 - Data-estimation: $p(y_{t+1}; x_{t+1}, h_t)p(\theta | y_{t+1}; x_{t+1}, h_t)$
- Each term can be taken direct from the LLM, derived from this, or set to something completely different
- LLM in context distributions are not coherent, so these will induce different joints!



Which Pairing?

- We can easily deviate from the LLM distributions for $p(\theta; h_t)$ or $p(y_{t+1}; x_{t+1}, h_t)$
 - We only need to **samples**
- But we need to directly evaluate **entropies** for $p(y_{t+1} | \theta; x_{t+1}, h_t)$ or $p(\theta | y_{t+1}; x_{t+1}, h_t)$
 - This depends on **full distribution** and it is harder to deviate from the LLM distribution
- Choose the pairing where this entropy is a more meaningful measure of uncertainty
- Simple $y \rightarrow$ Prior-Likelihood
- Simple $\theta \rightarrow$ Data-Estimation



Using the Right Estimator

$$\begin{aligned} \text{EIG}_{\theta}(x_{t+1}; h_t) &\approx \frac{1}{N} \sum_{n=1}^N \sum_{y_{t+1} \in \mathcal{Y}} p_{\text{LLM}}(y_{t+1}; [\theta^{(n)}, x_{t+1}]) \log p_{\text{LLM}}(y_{t+1}; [\theta^{(n)}, x_{t+1}]) \\ &\quad - \sum_{y_{t+1} \in \mathcal{Y}} \hat{p}(y_{t+1}; [h_t, x_{t+1}]) \log \hat{p}(y_{t+1}; [h_t, x_{t+1}]), \end{aligned}$$

where $\hat{p}(y_{t+1}; [h_t, x_{t+1}]) := \frac{1}{N} \sum_{n=1}^N p_{\text{LLM}}(y_{t+1}; [\theta^{(n)}, x_{t+1}])$ and $\theta^{(n)} \sim p_f(\theta; h_t)$

Don't assume the likelihood has constant expected entropy across designs!



Candidate Question Generation

- Two options:
 - Direct generation using $p_{\text{LLM}}(x_{t+1}; h_t)$
 - Conditional generation on set of θ samples using $p_{\text{LLM}}(x_{t+1}; h_t, \hat{\theta}_1, \hat{\theta}_2, \dots, \hat{\theta}_N)$
- Latter works well for discrete, easily divisible hypothesis spaces, but can overfit in more complex scenarios.



Baselines

- Naive QA: prompt the LLM to choose its own next question with history in context
- Predictive Entropy: choose question with maximal predictive entropy, as done by Cooper et al 2025, Kobalcyzk et al 2025, Mazzaccara et al 2024:

$$\text{Entropy}_\theta(x_{t+1}) = - \sum_{y_{t+1} \in \mathcal{Y}} \hat{p}(y_{t+1}; [h_t, x_{t+1}]) \log \hat{p}(y_{t+1}; [h_t, x_{t+1}])$$



Answerer

- To conduct experiments we need to evaluate with some sort of answerer
- We use a separate LLM that has hidden access to the true θ (it does not keep previous question response pairs in context)



Updated Model => New EIG

An updated model $p(\theta, y_{t+1}; x_{t+1}, h_t)$
incorporating data history h_t defines a new EIG:

$$\begin{aligned} \text{EIG}_{\theta}(x_{t+1}) &= \text{H}[p(\theta; h_t)] - \mathbb{E}_{p(y_{t+1}; x_{t+1}, h_t)} [\text{H}[p(\theta | y_{t+1}; x_{t+1}, h_t)]] \\ &= \text{H}[p(y_{t+1}; x_{t+1}, h_t)] - \mathbb{E}_{p(\theta; h_t)} [\text{H}[p(y_{t+1} | \theta; x_{t+1}, h_t)]] \end{aligned}$$



Applications

- Clarifying user intent
- Interactive multi-turn conversations
- Automated surveys
- Recommender systems
- Tutoring systems
- AI scientists
- Assisted code generation
- **Understanding existing LLM's capabilities**
- ...