

MEASURING THE RELIABILITY OF CAUSAL PROBING METHODS: TRADEOFFS, LIMITATIONS, AND THE PLIGHT OF NULLIFYING INTERVENTIONS

Marc E. Canby*, Adam Davies*, Chirag Rastogi, Julia Hockenmaier



*Equal Contribution

University of Illinois Urbana-Champaign



SUMMARY

Which features do LLMs use to perform a given task?

- **Causal interventions** on feature representations

SUMMARY

Which features do LLMs use to perform a given task?

- **Causal interventions** on feature representations

How *reliable* are interventions?

- **Completeness:** Is the intended intervention carried out?
- **Selectivity:** Are we damaging non-targeted features?

SUMMARY

Which features do LLMs use to perform a given task?

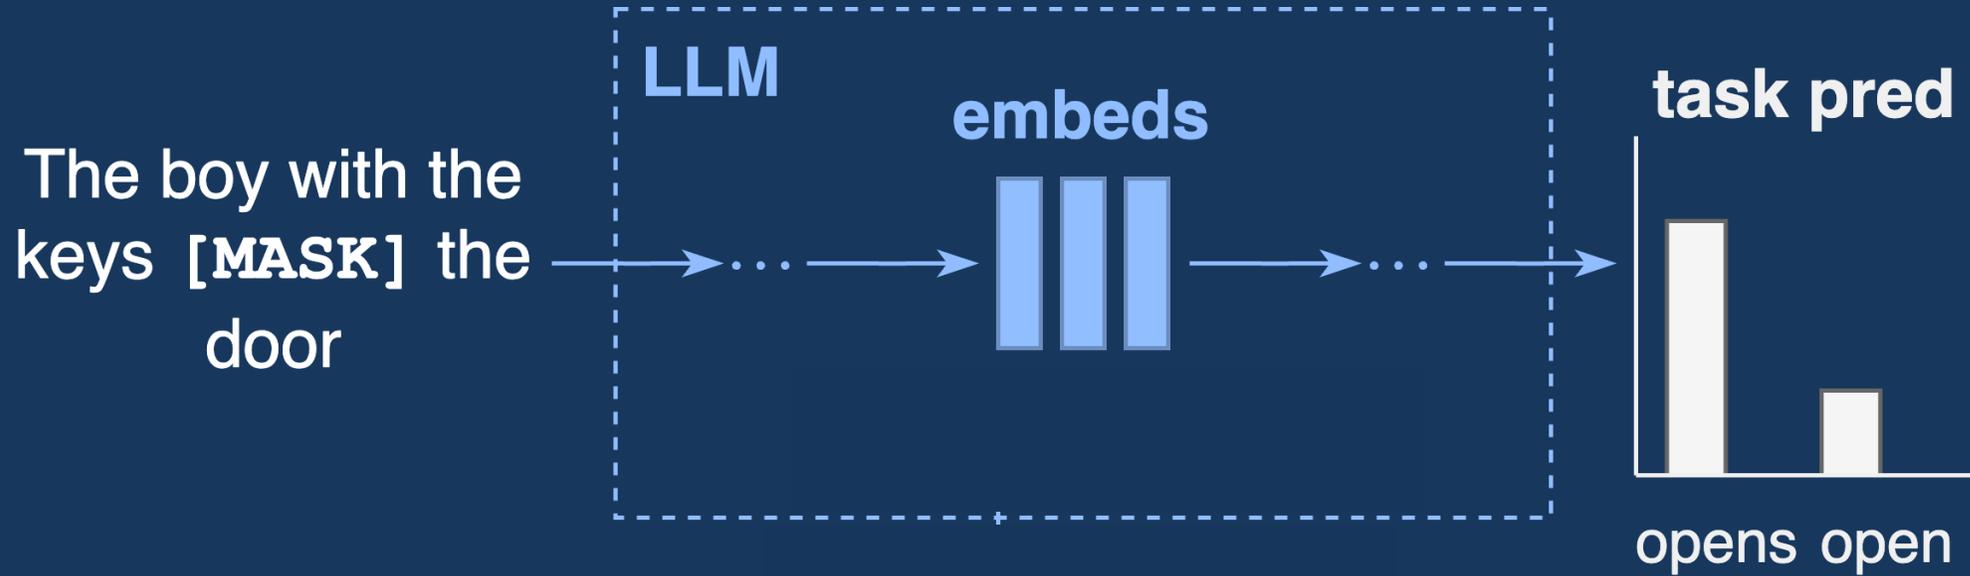
- **Causal interventions** on feature representations

How *reliable* are interventions?

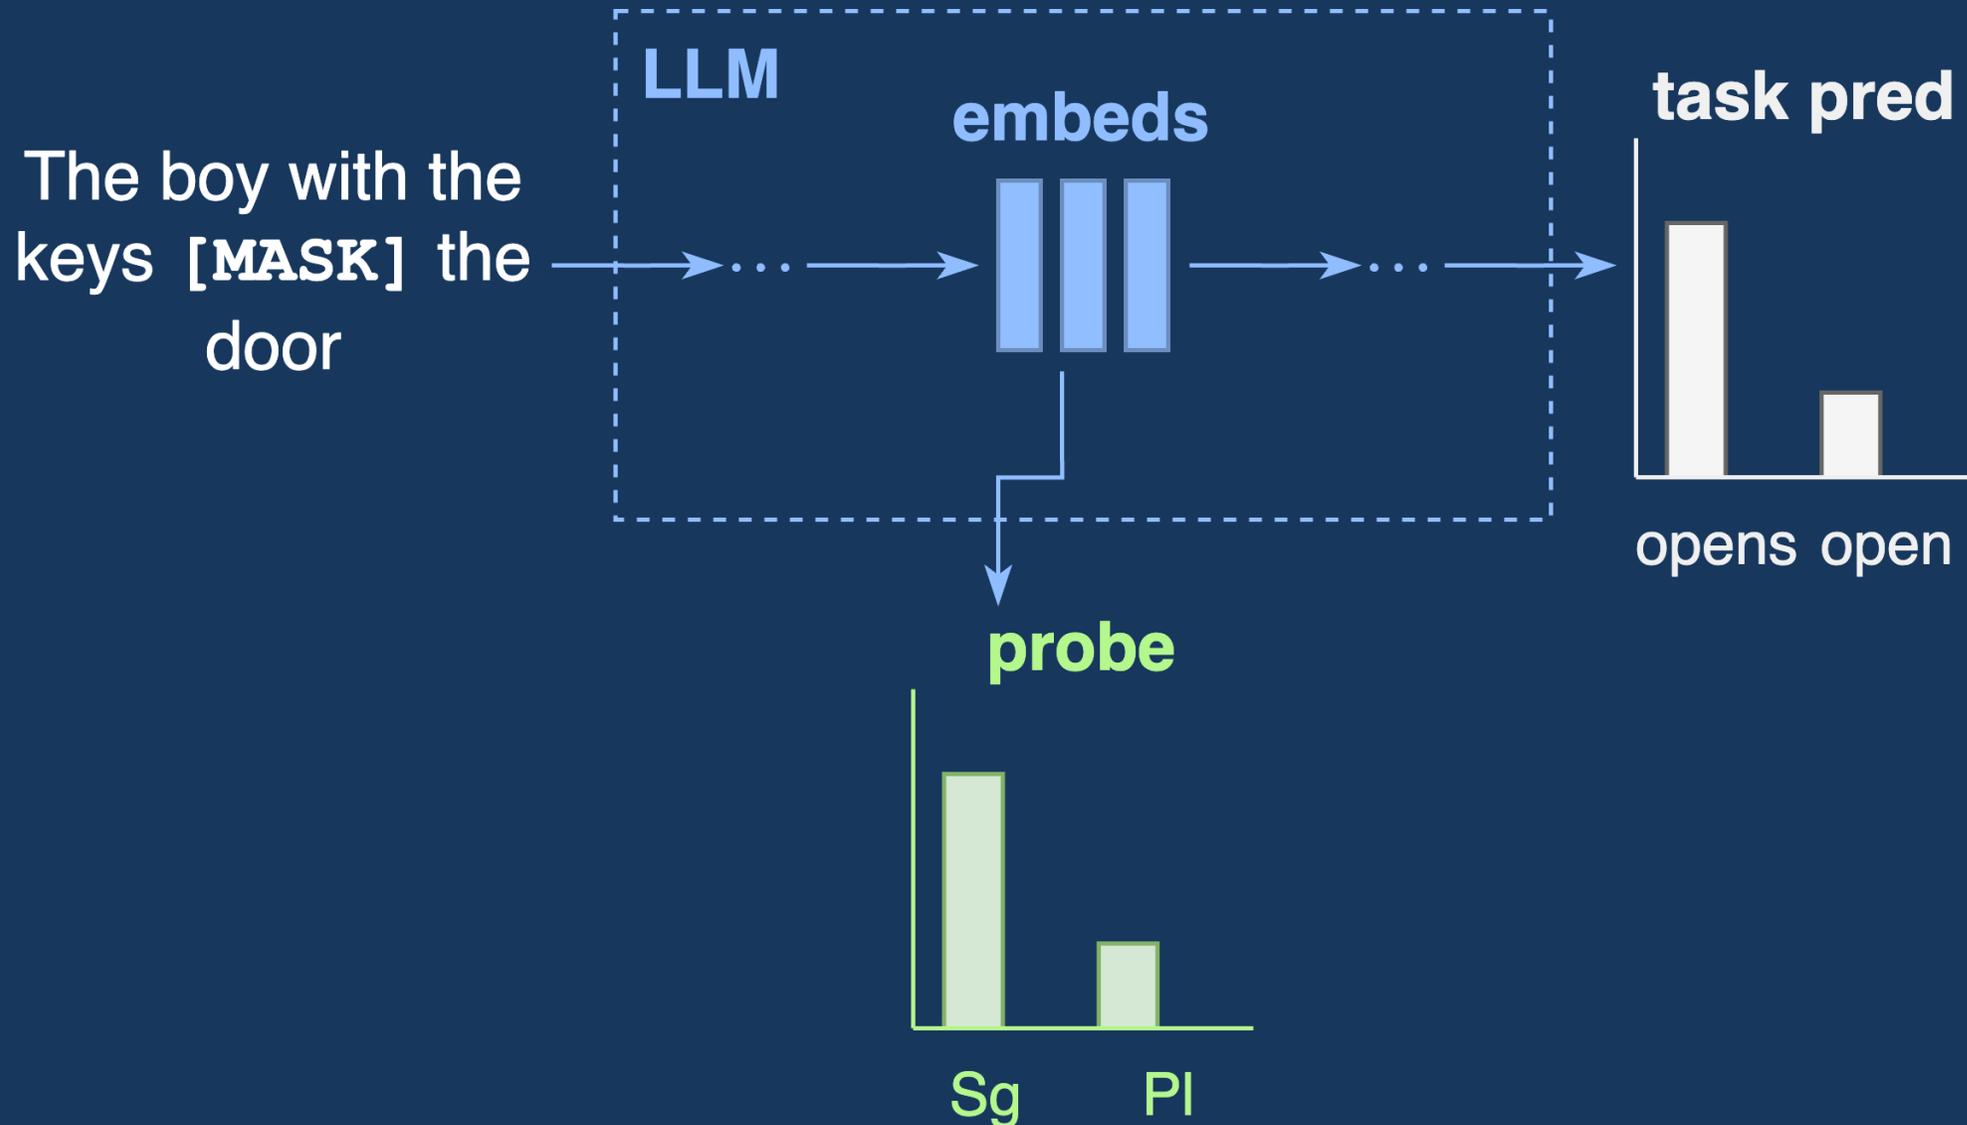
- **Completeness:** Is the intended intervention carried out?
- **Selectivity:** Are we damaging non-targeted features?

We define an **evaluation framework** to compare *different classes* of interventions

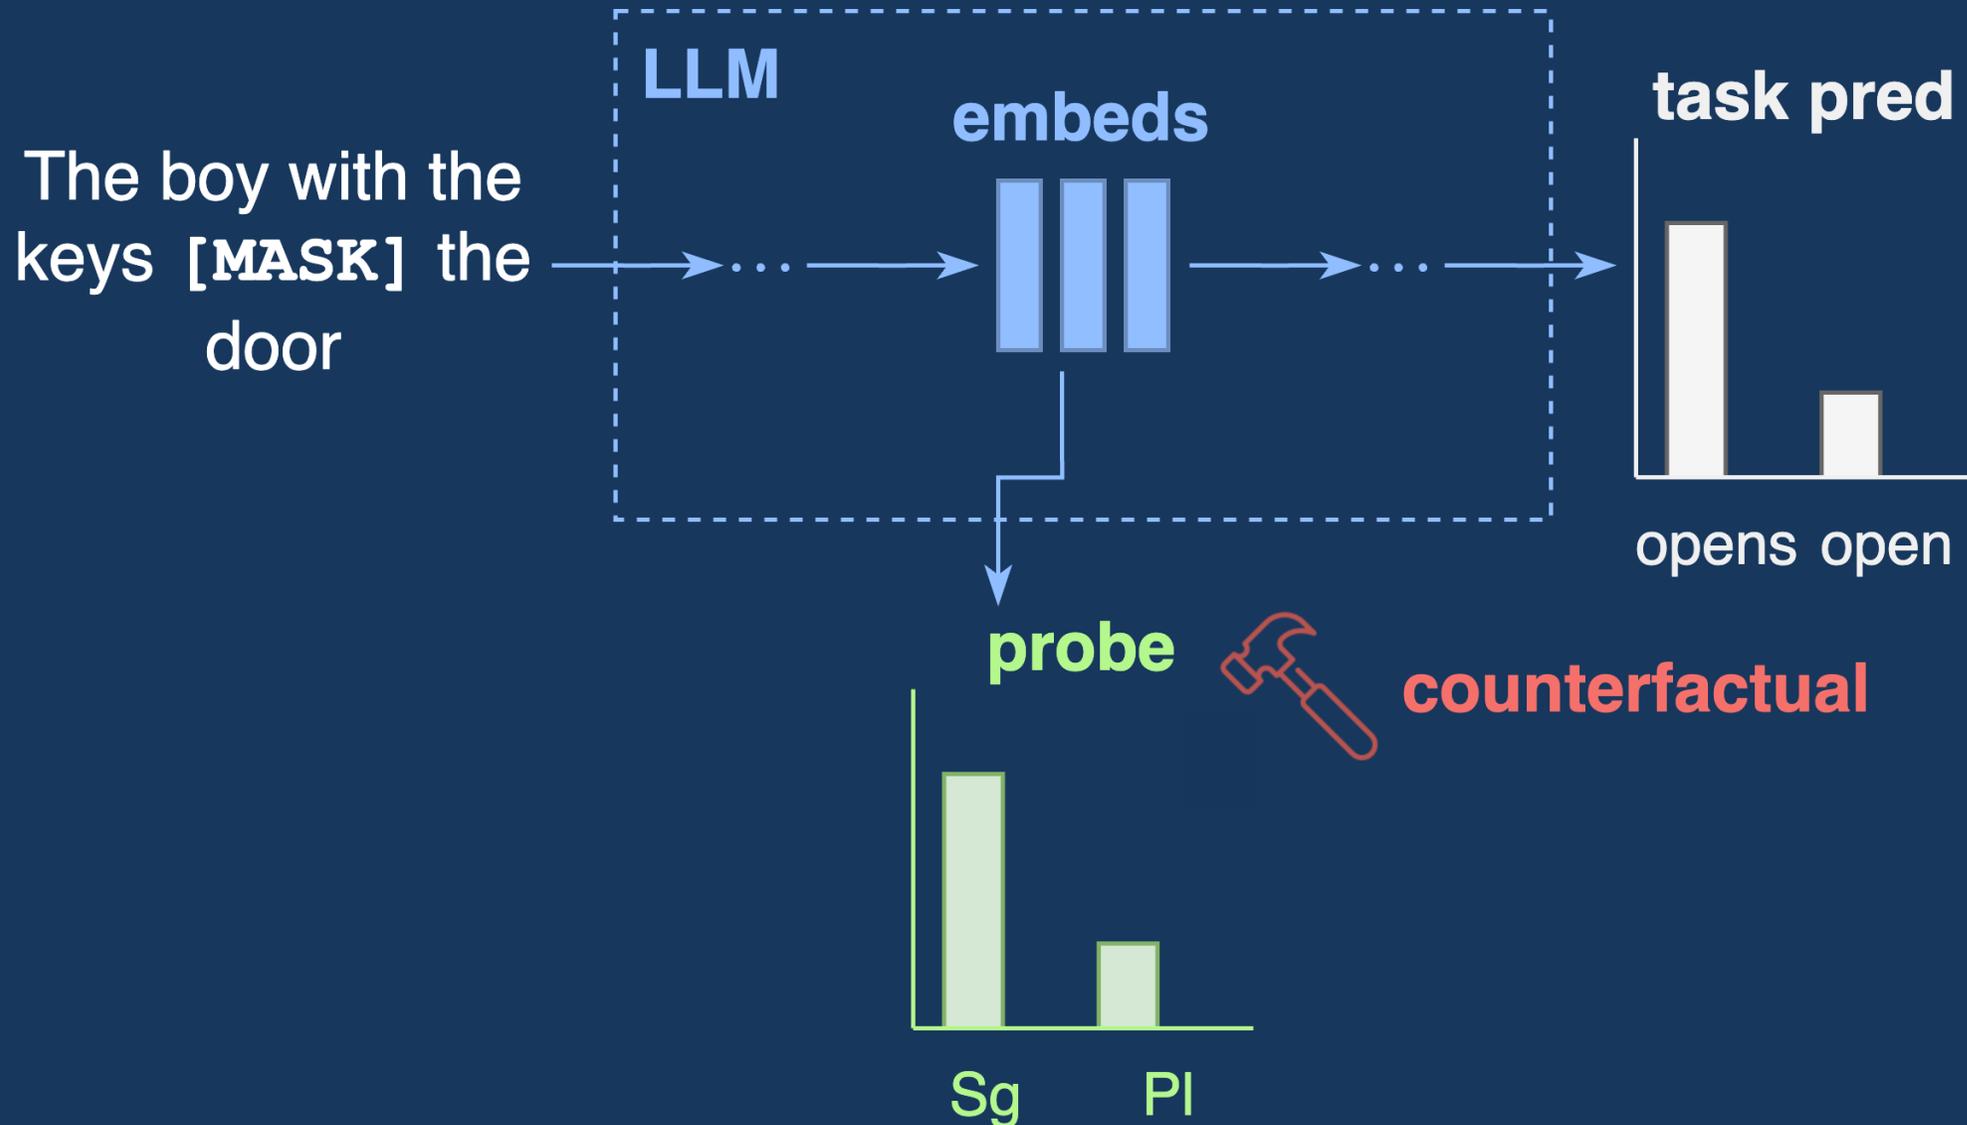
CAUSAL PROBING



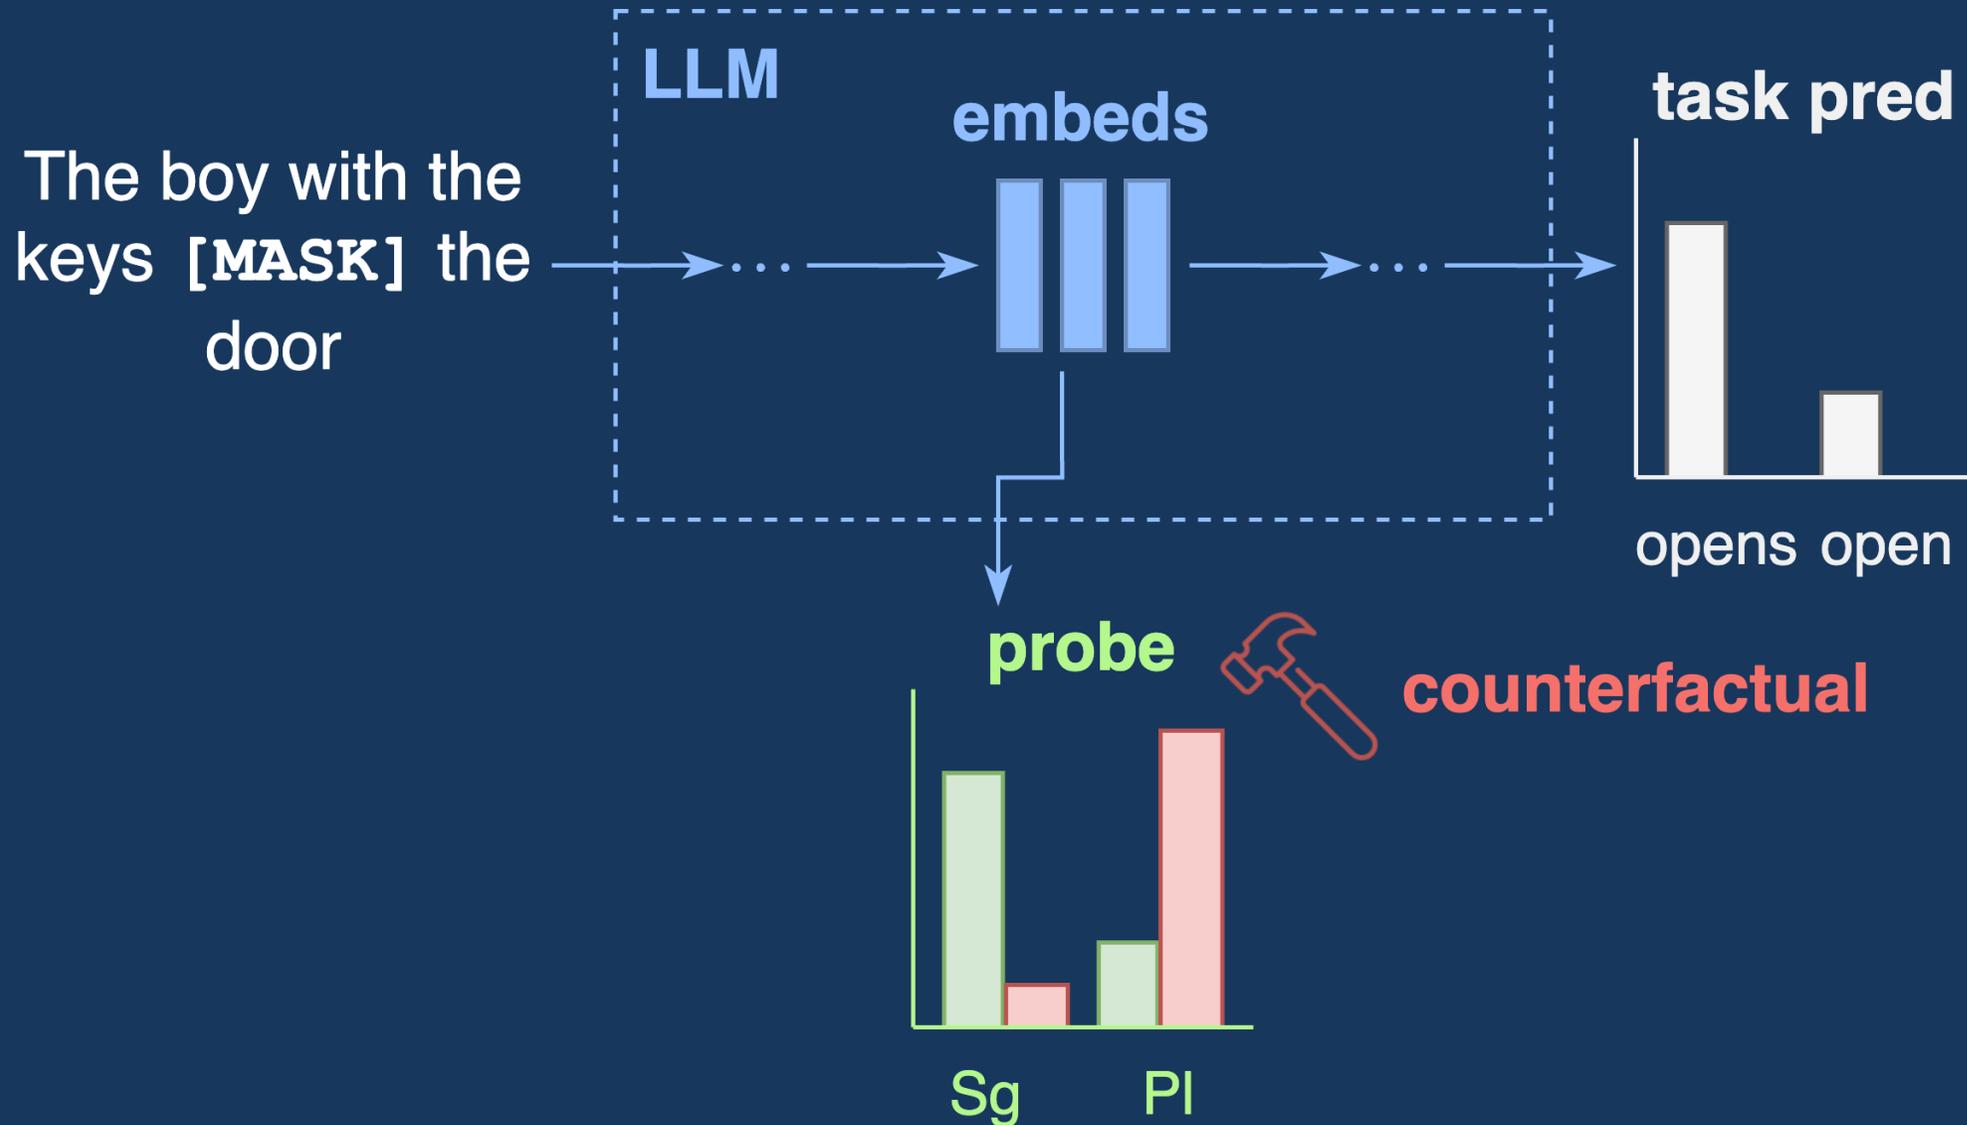
CAUSAL PROBING



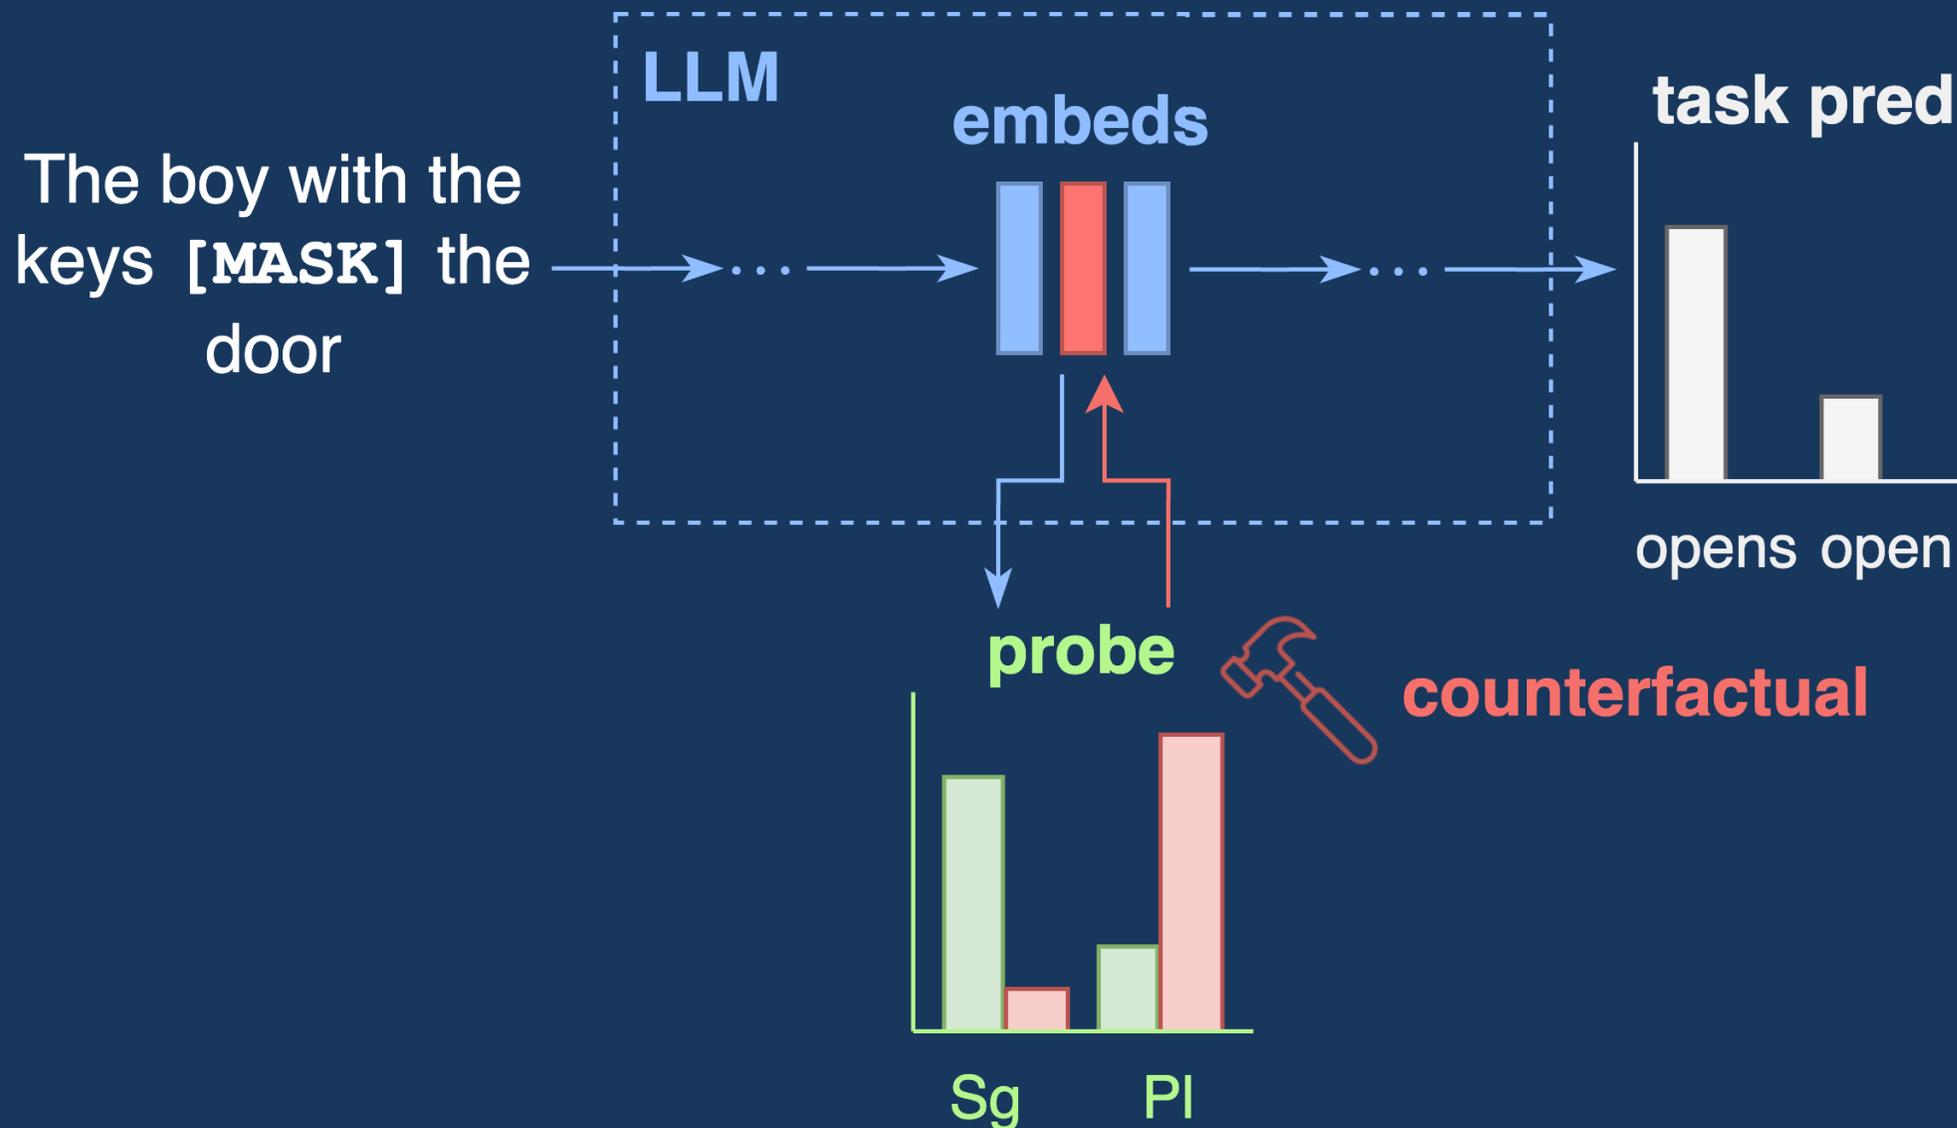
CAUSAL PROBING



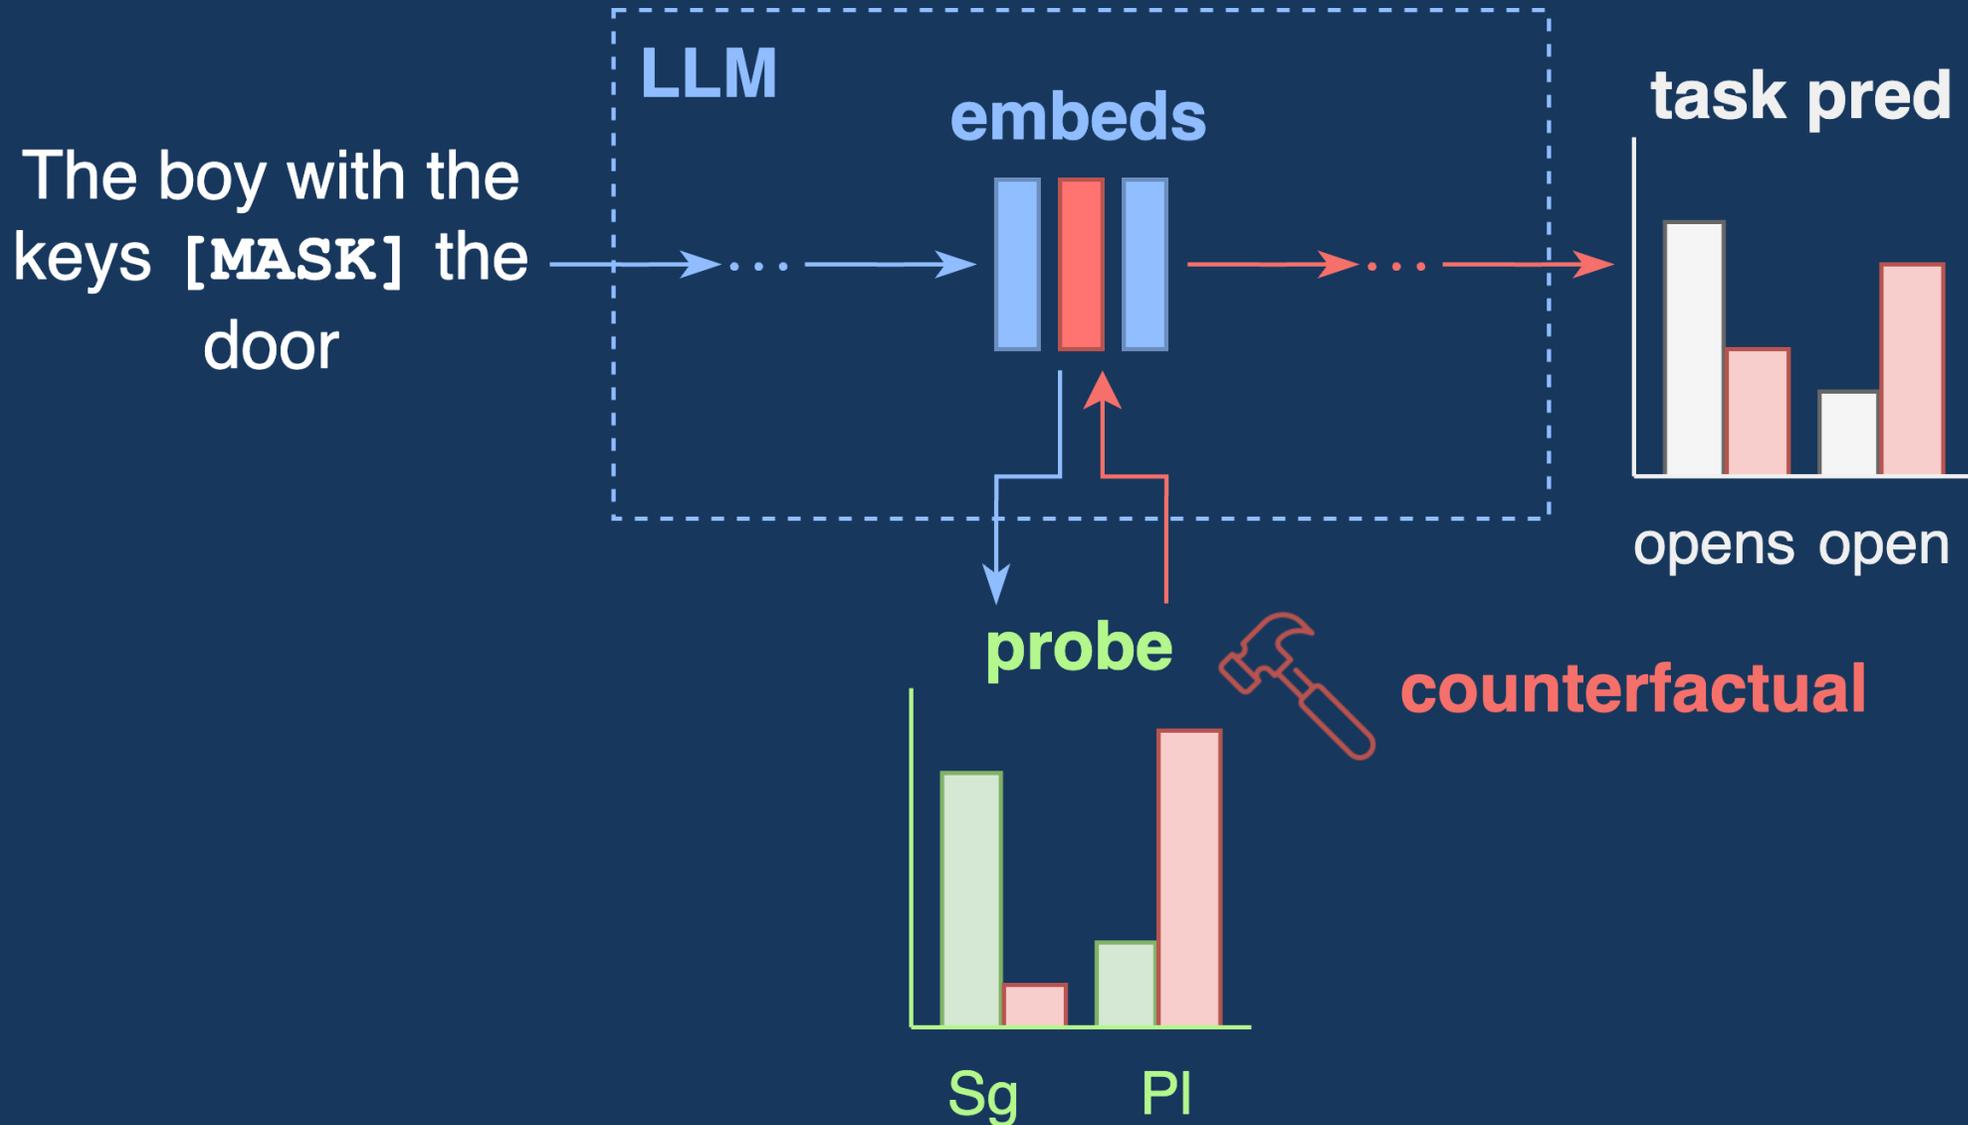
CAUSAL PROBING



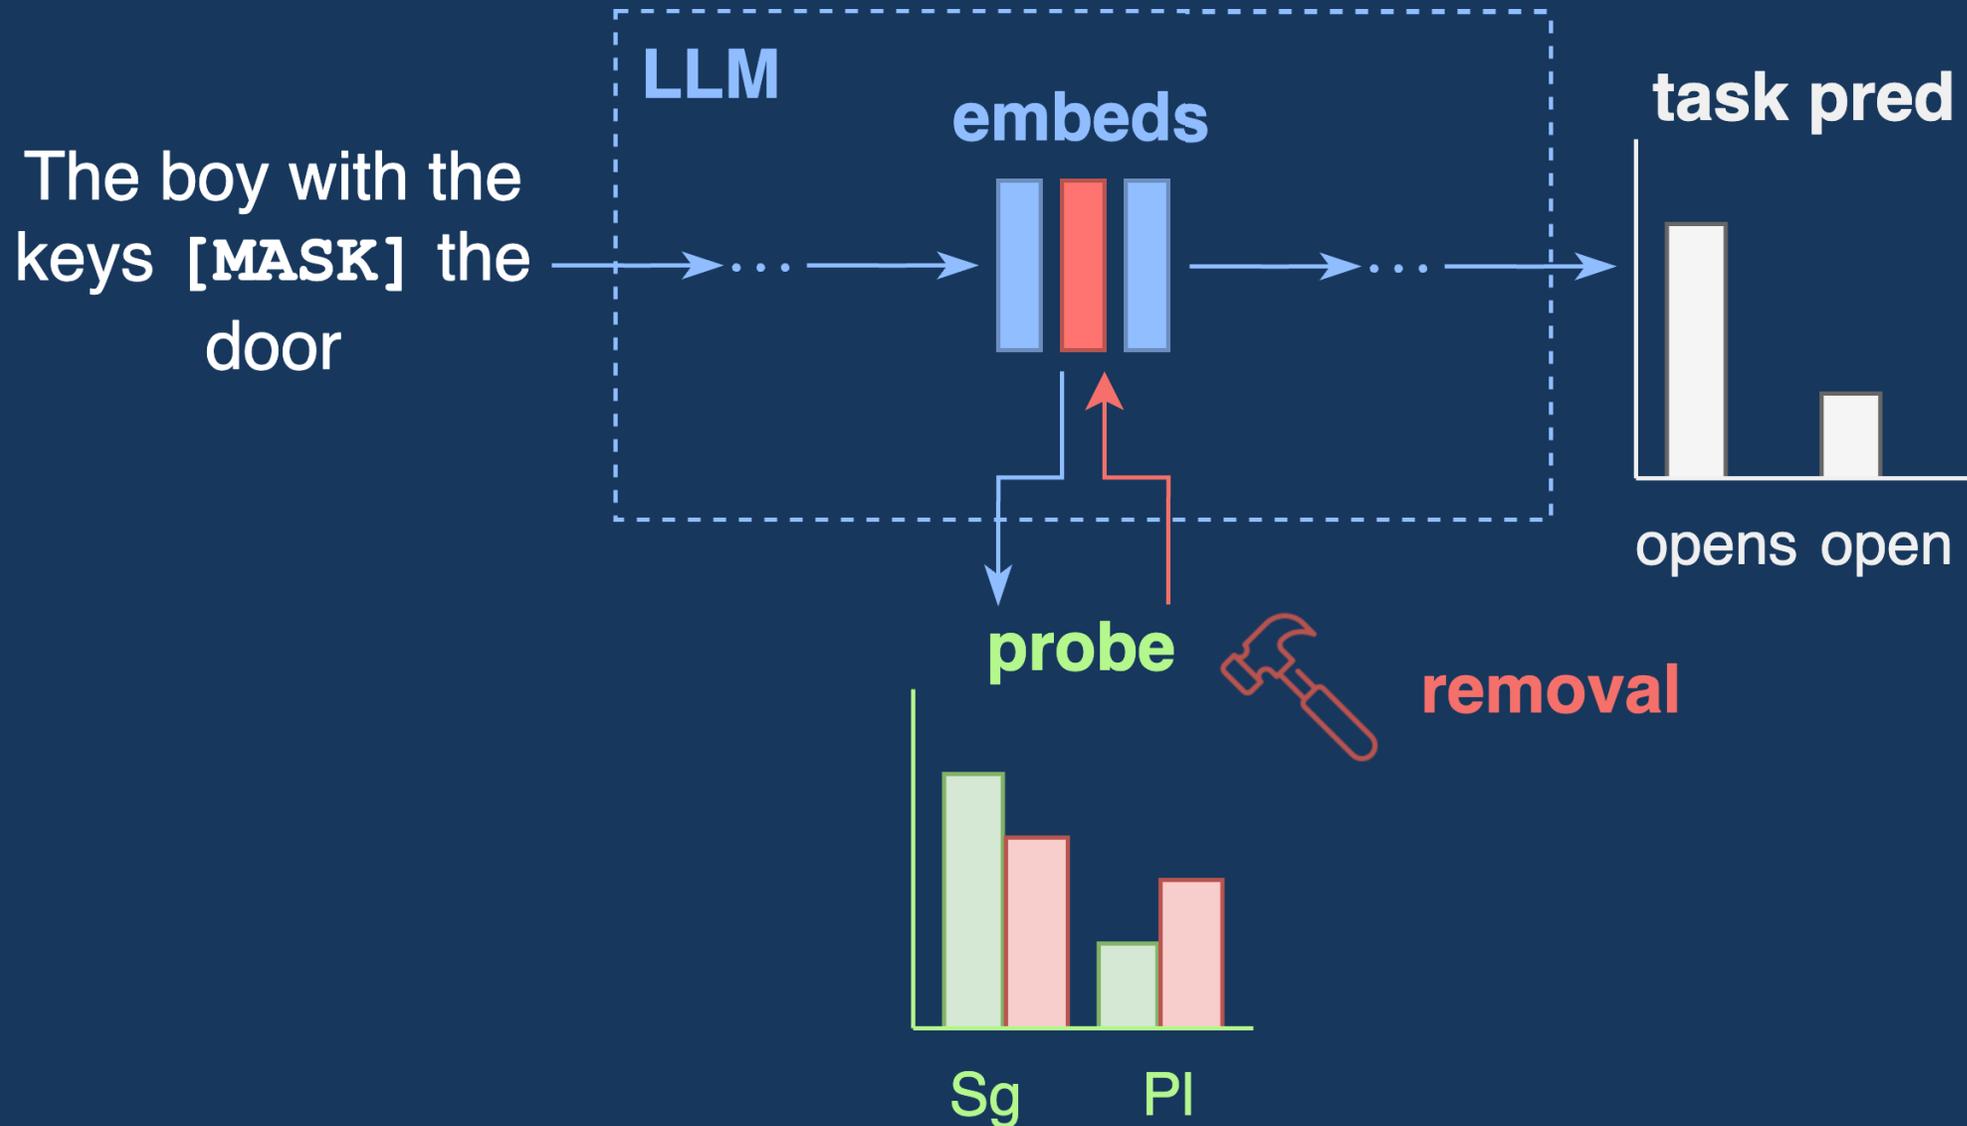
CAUSAL PROBING



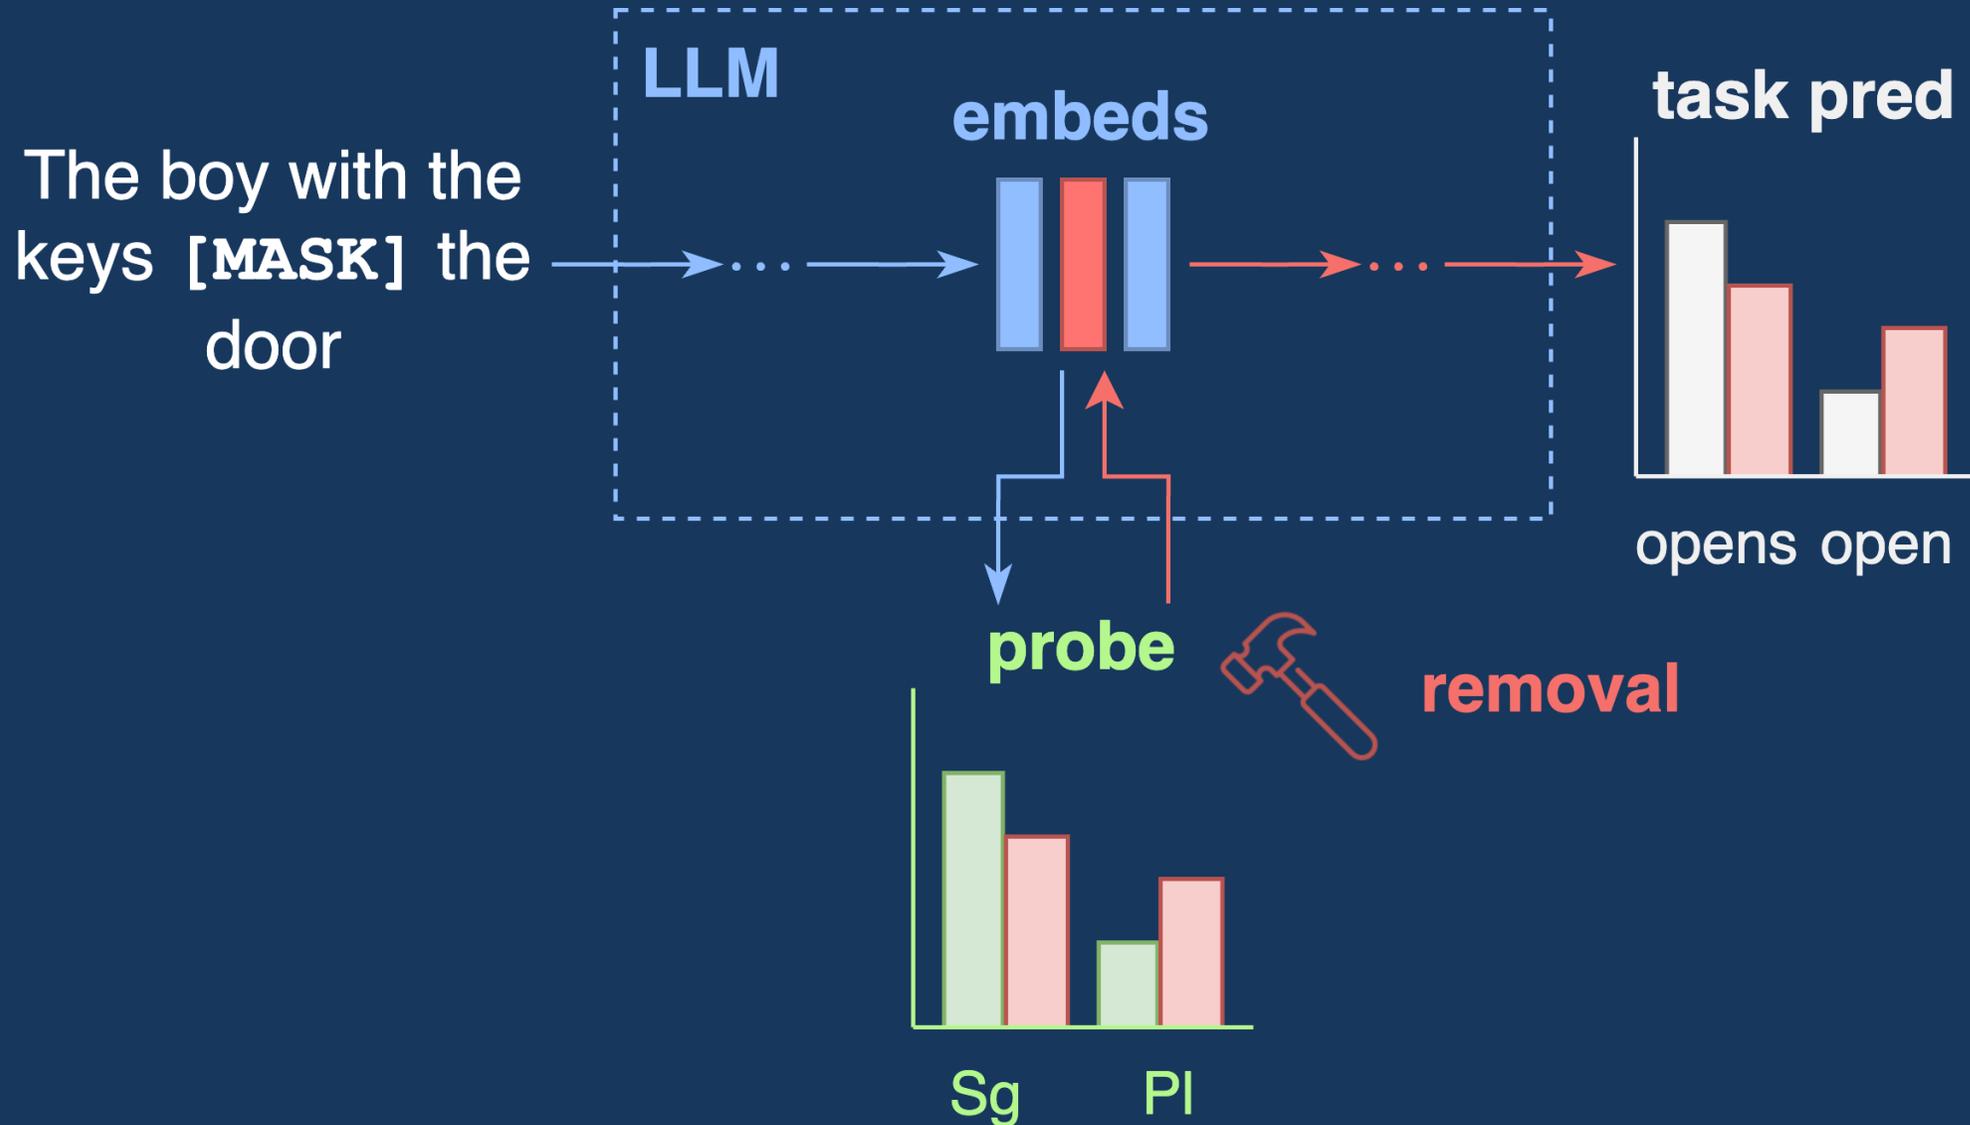
CAUSAL PROBING



CAUSAL PROBING



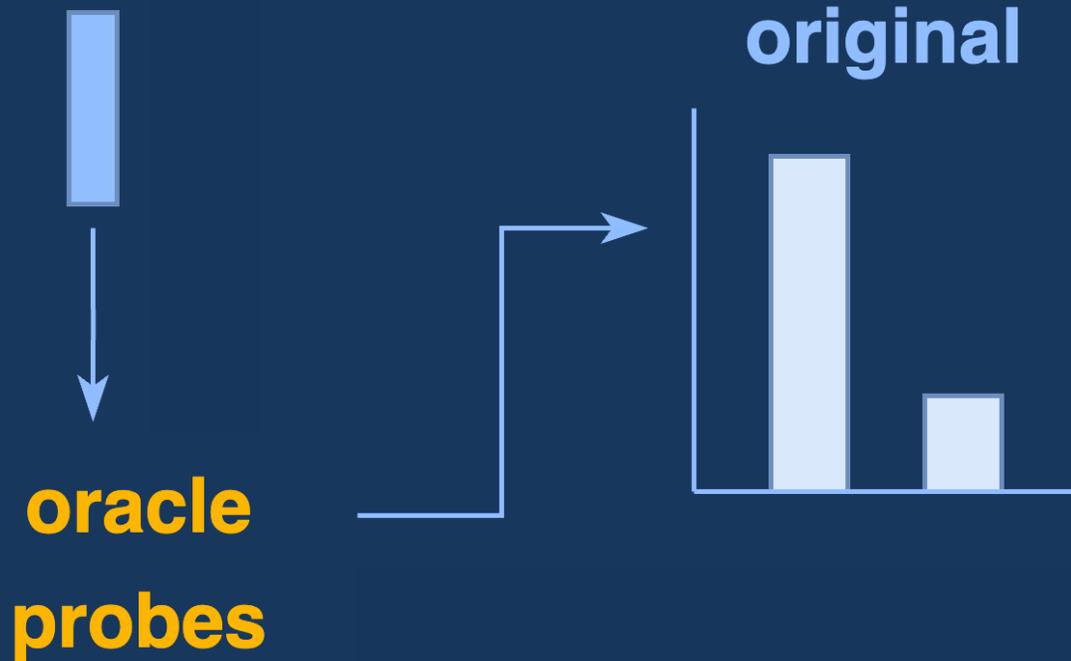
CAUSAL PROBING



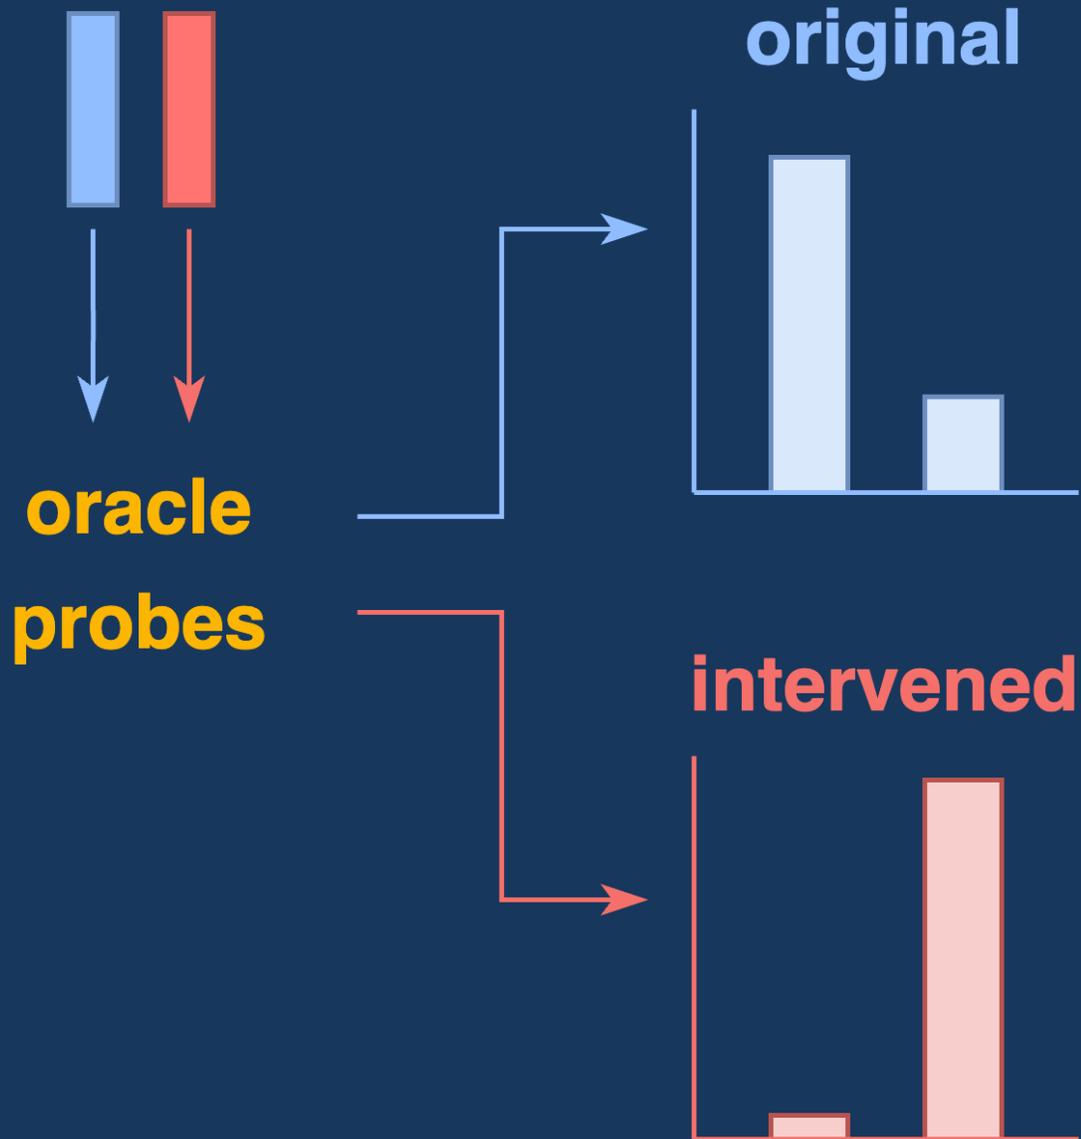
OUR EVALUATION FRAMEWORK

oracle
probes

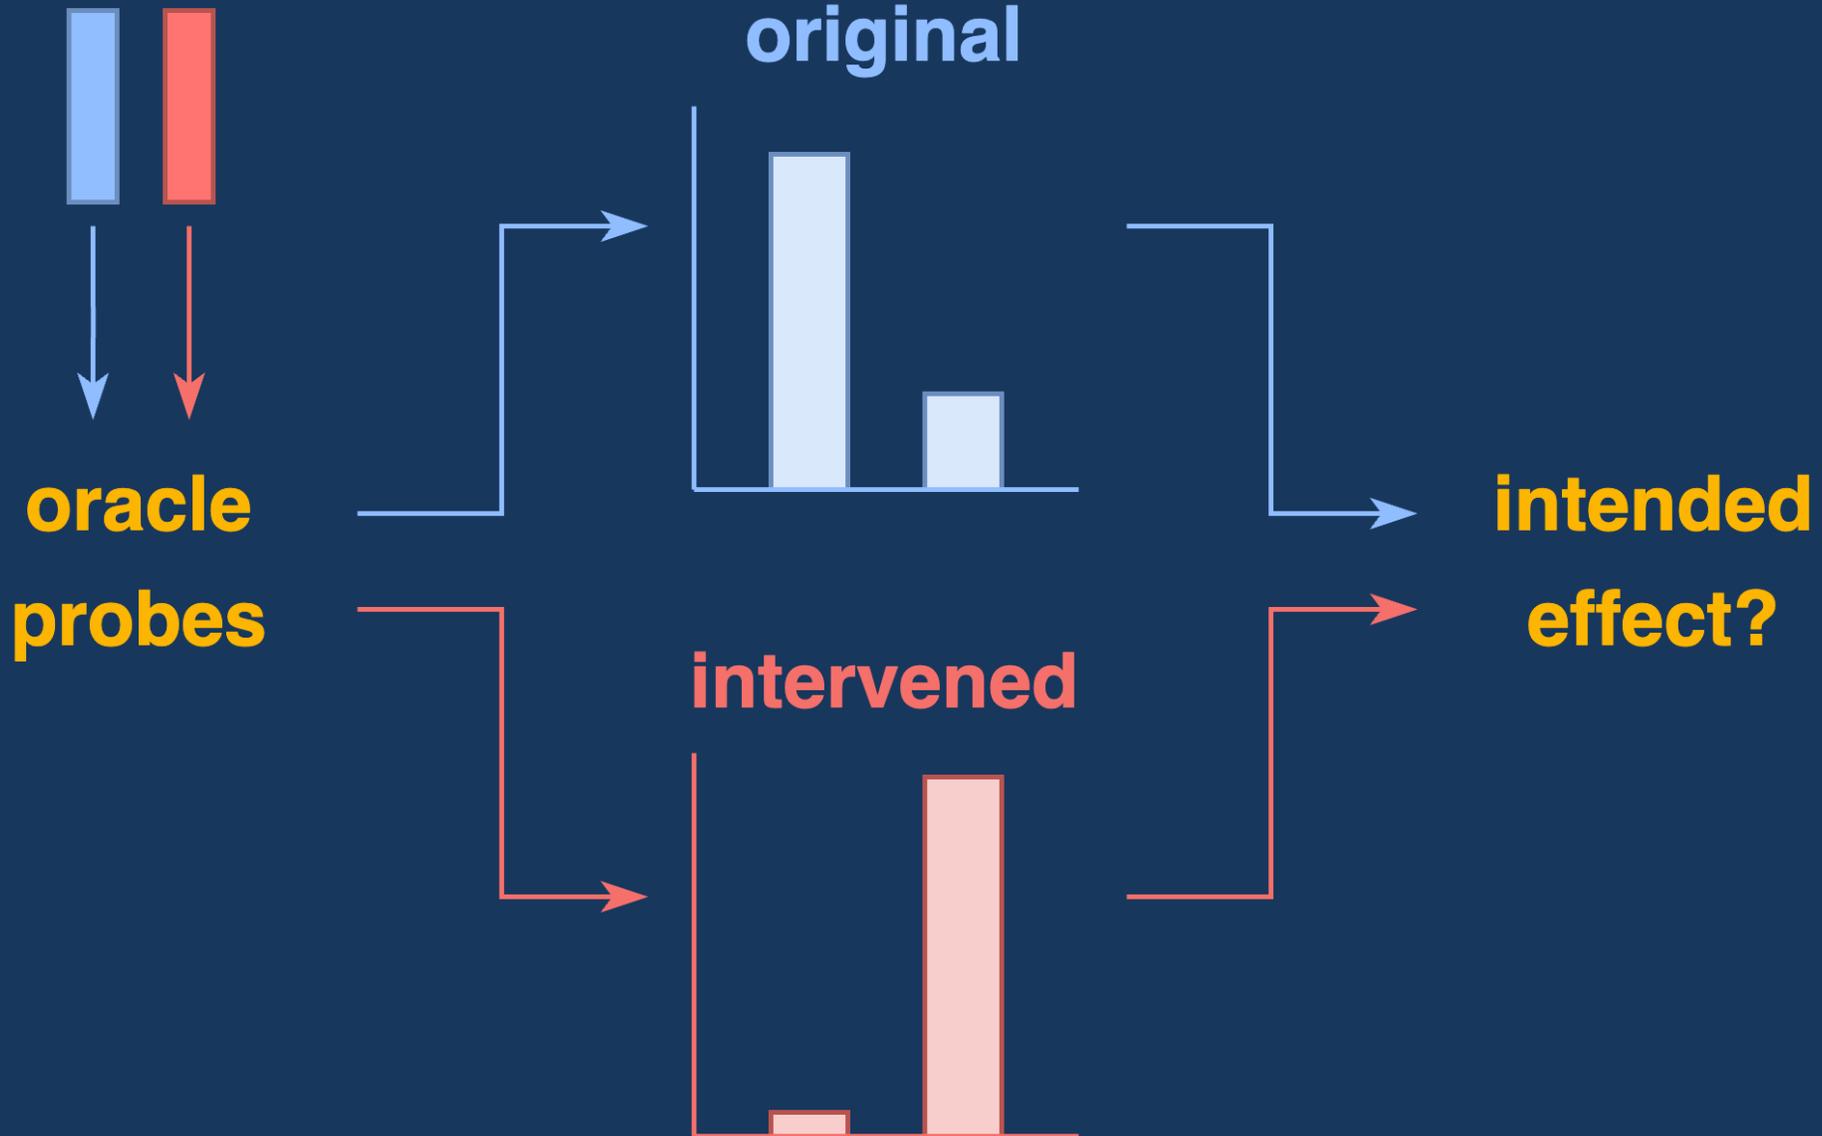
OUR EVALUATION FRAMEWORK



OUR EVALUATION FRAMEWORK



OUR EVALUATION FRAMEWORK



OUR EVALUATION FRAMEWORK

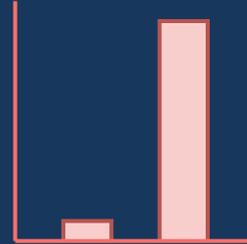
completeness



OUR EVALUATION FRAMEWORK

completeness

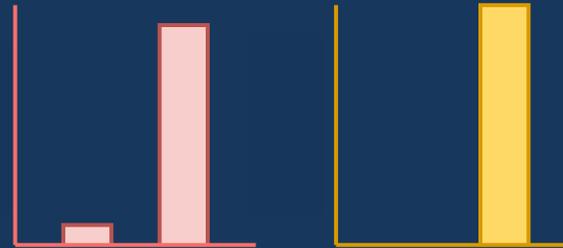
counterfactual



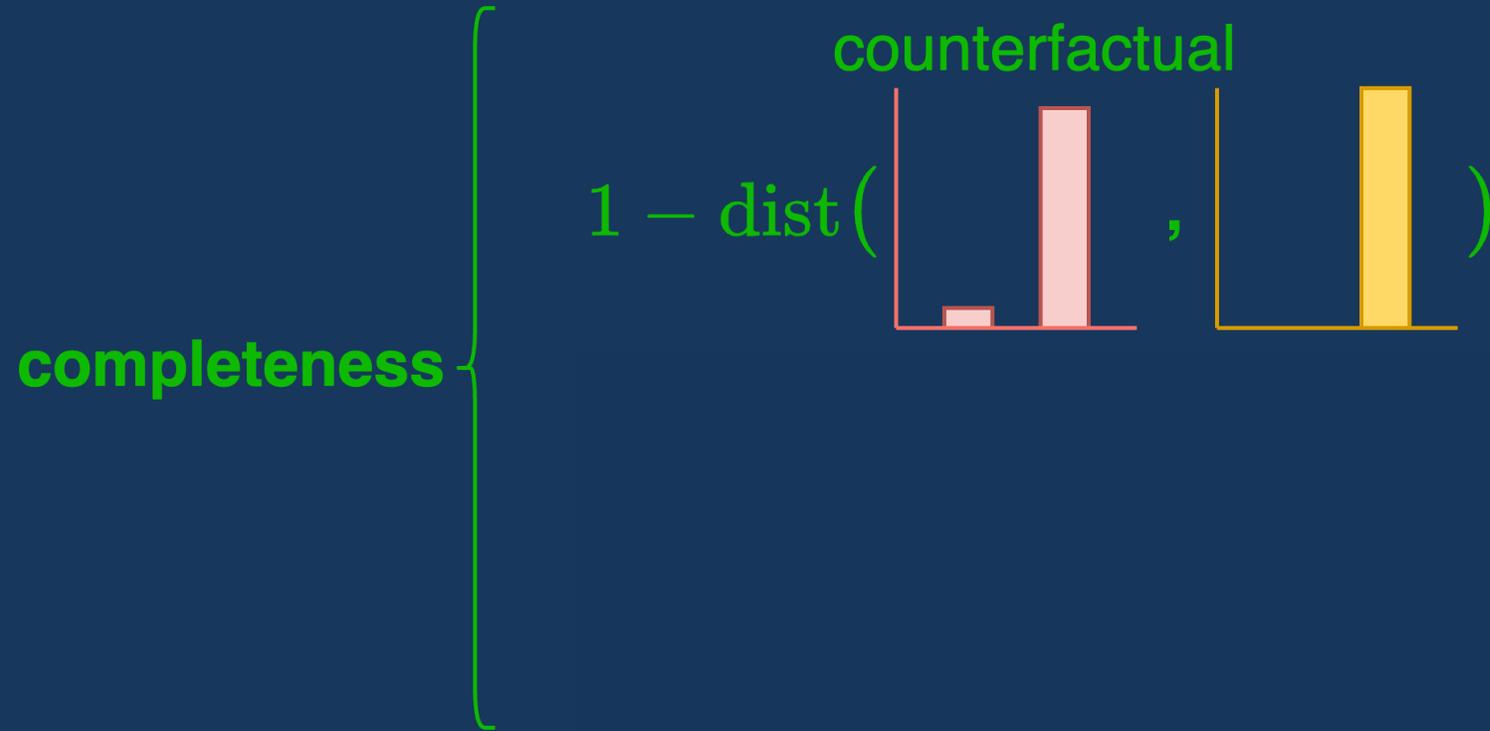
OUR EVALUATION FRAMEWORK

completeness

counterfactual



OUR EVALUATION FRAMEWORK



OUR EVALUATION FRAMEWORK

completeness

counterfactual

$$1 - \text{dist}(\text{[red bars]}, \text{[yellow bar]}) \in [0, 1]$$

OUR EVALUATION FRAMEWORK



OUR EVALUATION FRAMEWORK



CAUSAL PROBING INTERVENTIONS

Concept Removal: *Remove representation of target feature*

– **Linear:** INLP, RLACE

CAUSAL PROBING INTERVENTIONS

Concept Removal: *Remove* representation of target feature

– **Linear:** INLP, RLACE

Counterfactual: *Swap* representation from one value to another

– **Linear:** AlterRep

– **Nonlinear (GBIs):** FGSM, PGD, AutoAttack

CAUSAL PROBING INTERVENTIONS

Concept Removal: *Remove representation of target feature*

– **Linear:** INLP, RLACE

Counterfactual: *Swap representation from one value to another*

– **Linear:** AlterRep

– **Nonlinear (GBIs):** FGSM, PGD, AutoAttack

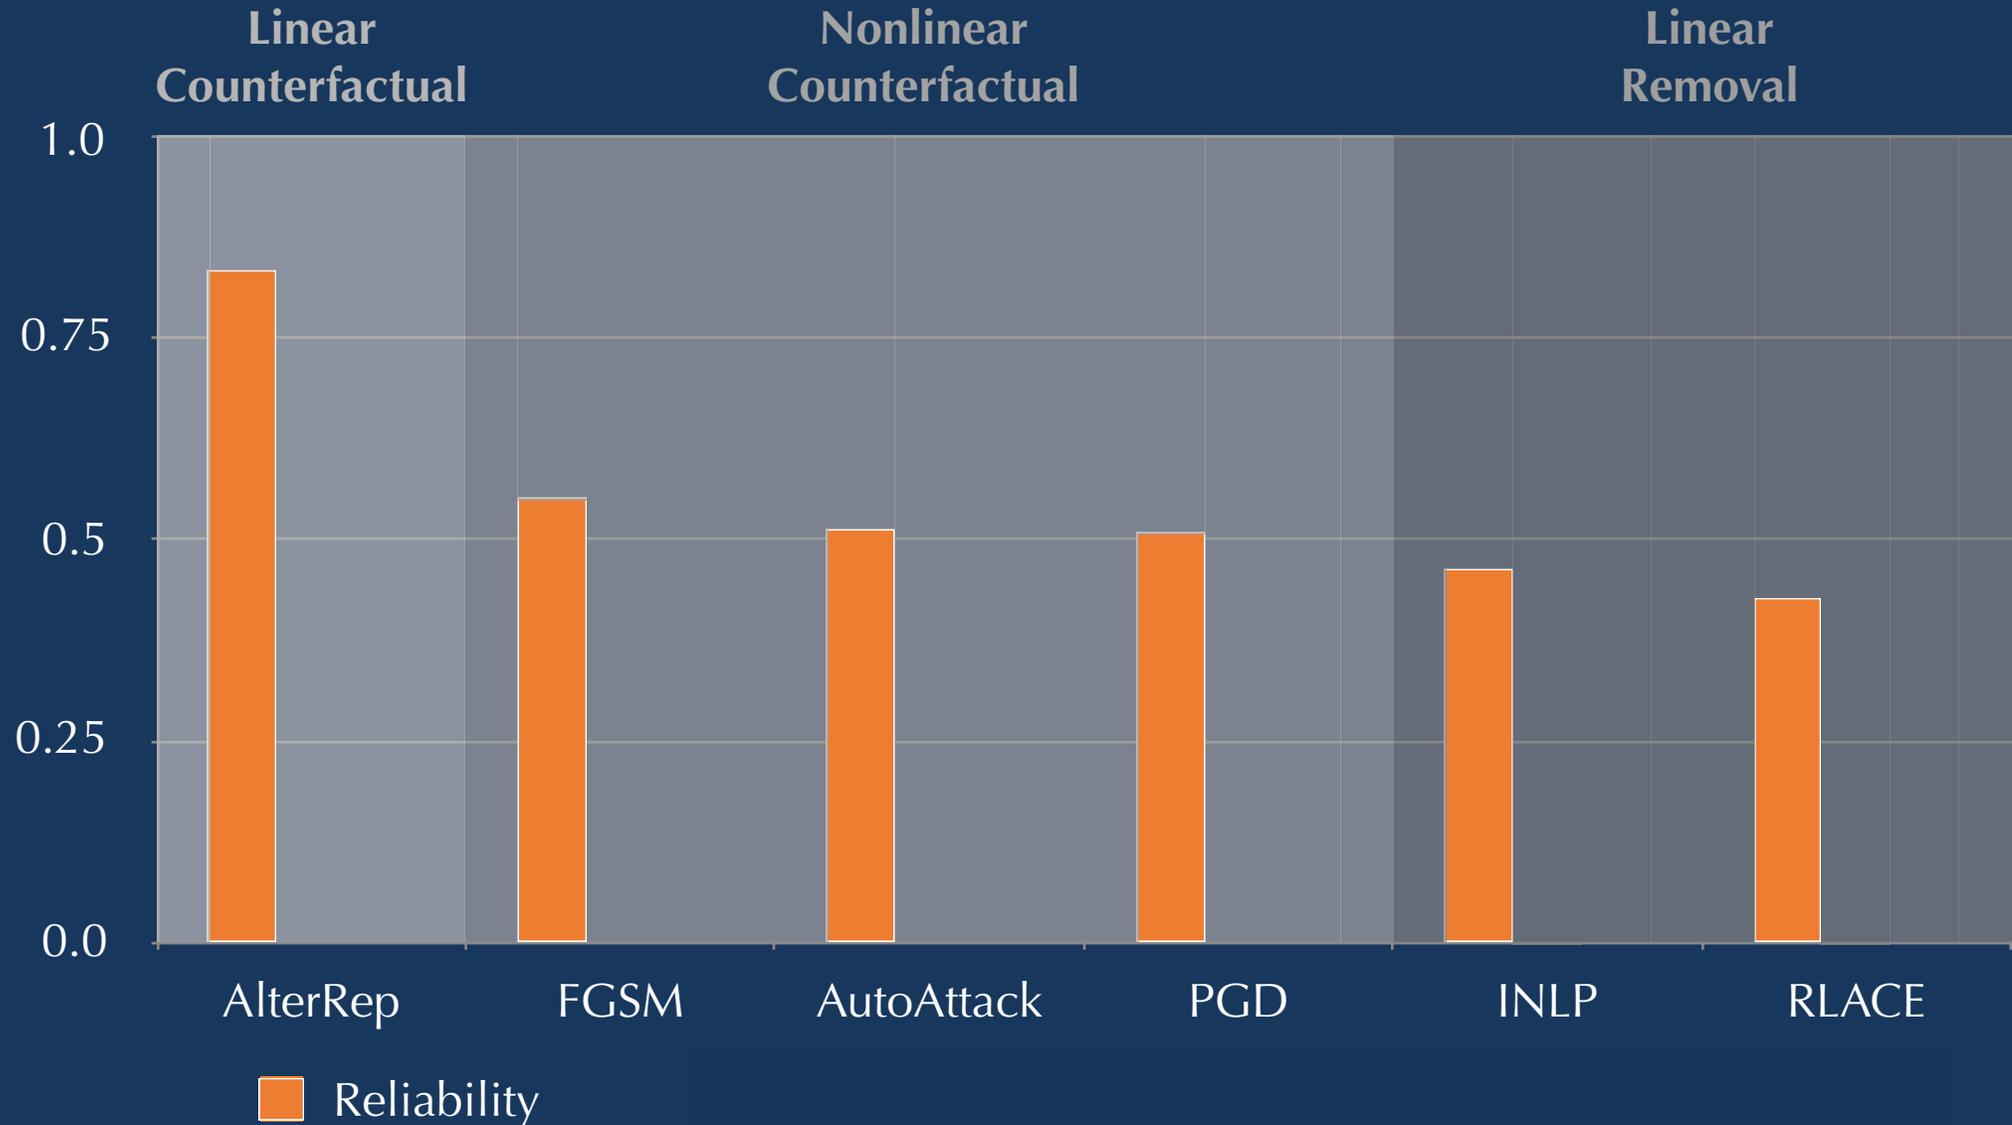
Intervene on [MASK] token in final layer of BERT

OUR EVALUATION FRAMEWORK

Evaluate interventions according to

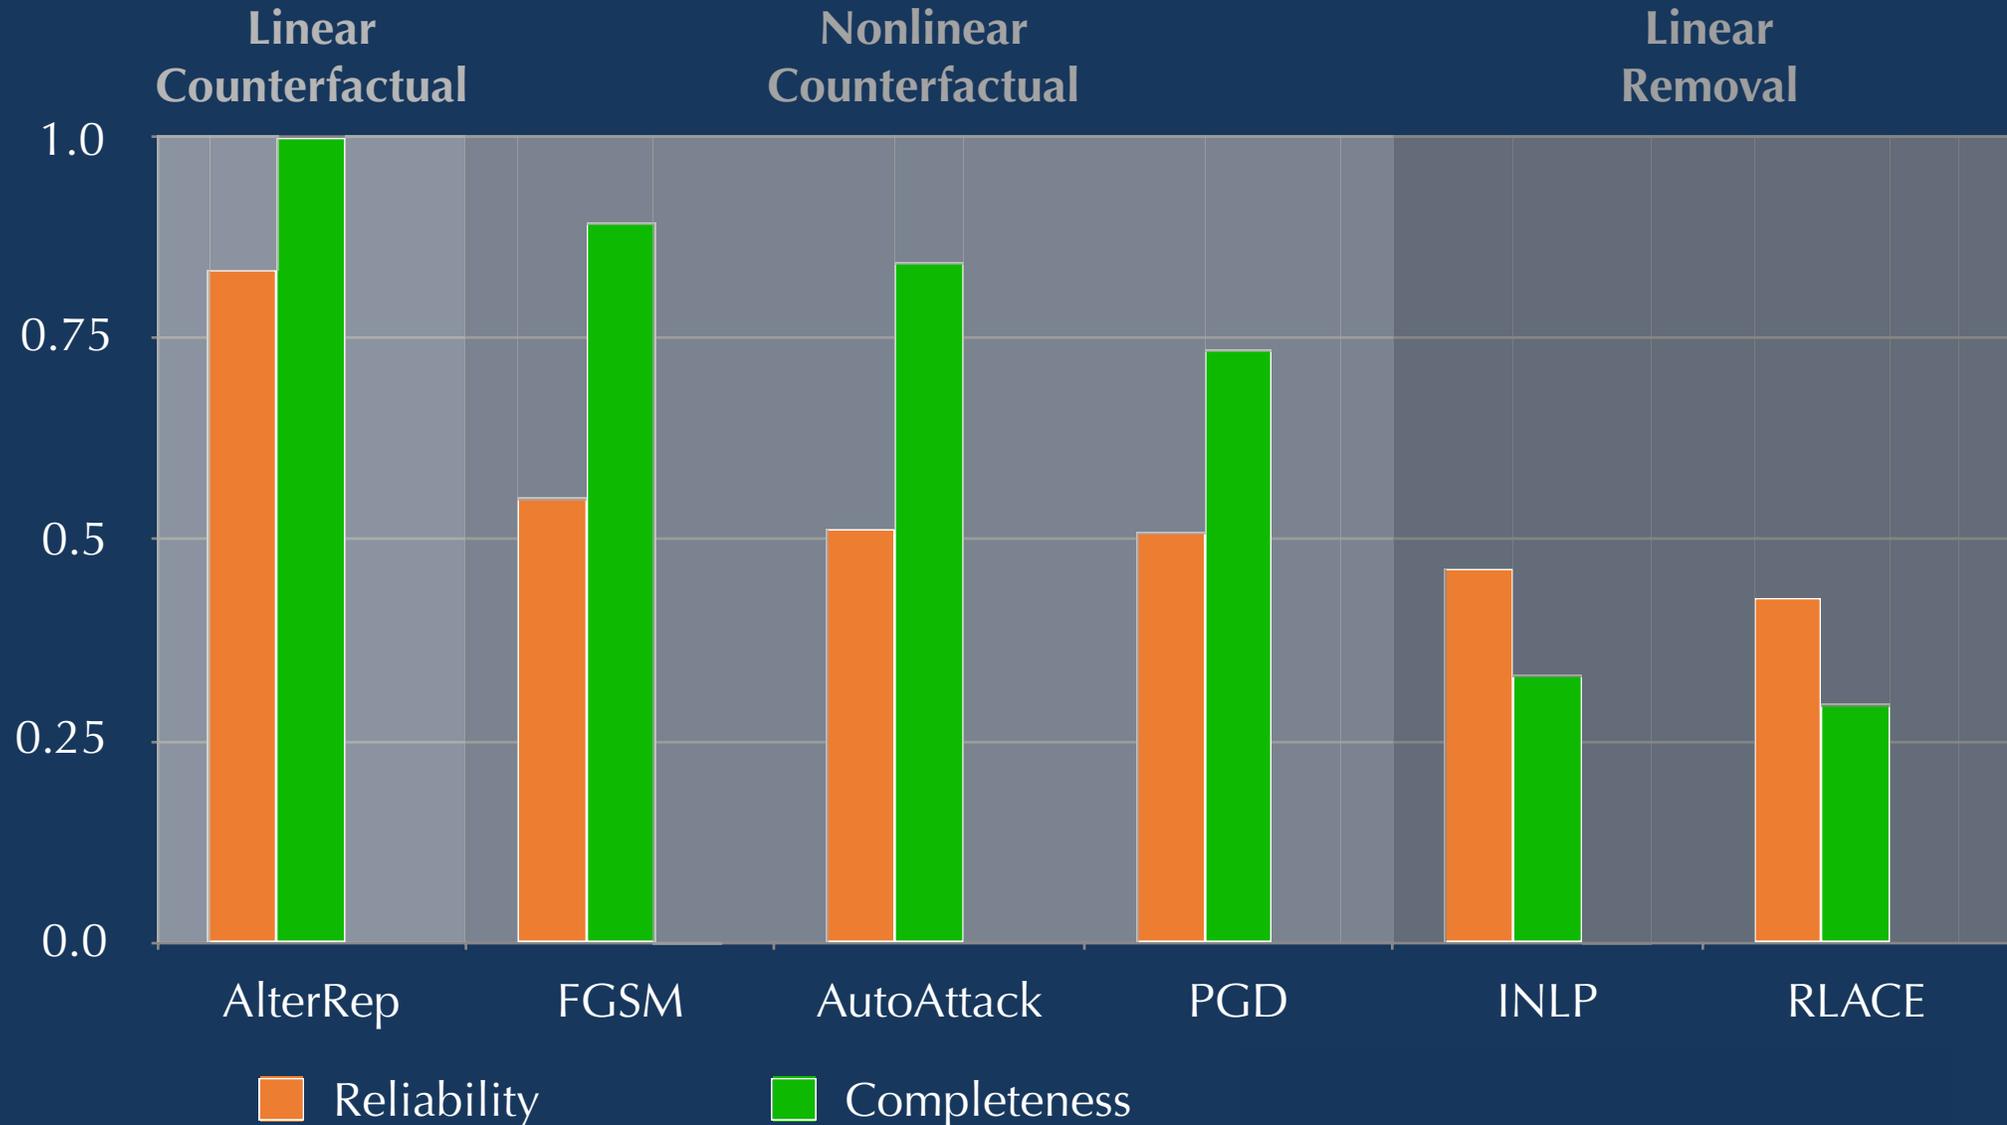
- **Completeness:** Is the intended intervention carried out?
- **Selectivity:** Are we damaging non-targeted features?
- **Reliability:** Harmonic mean of completeness and selectivity

RELIABILITY: ALTERREP IS MOST RELIABLE!

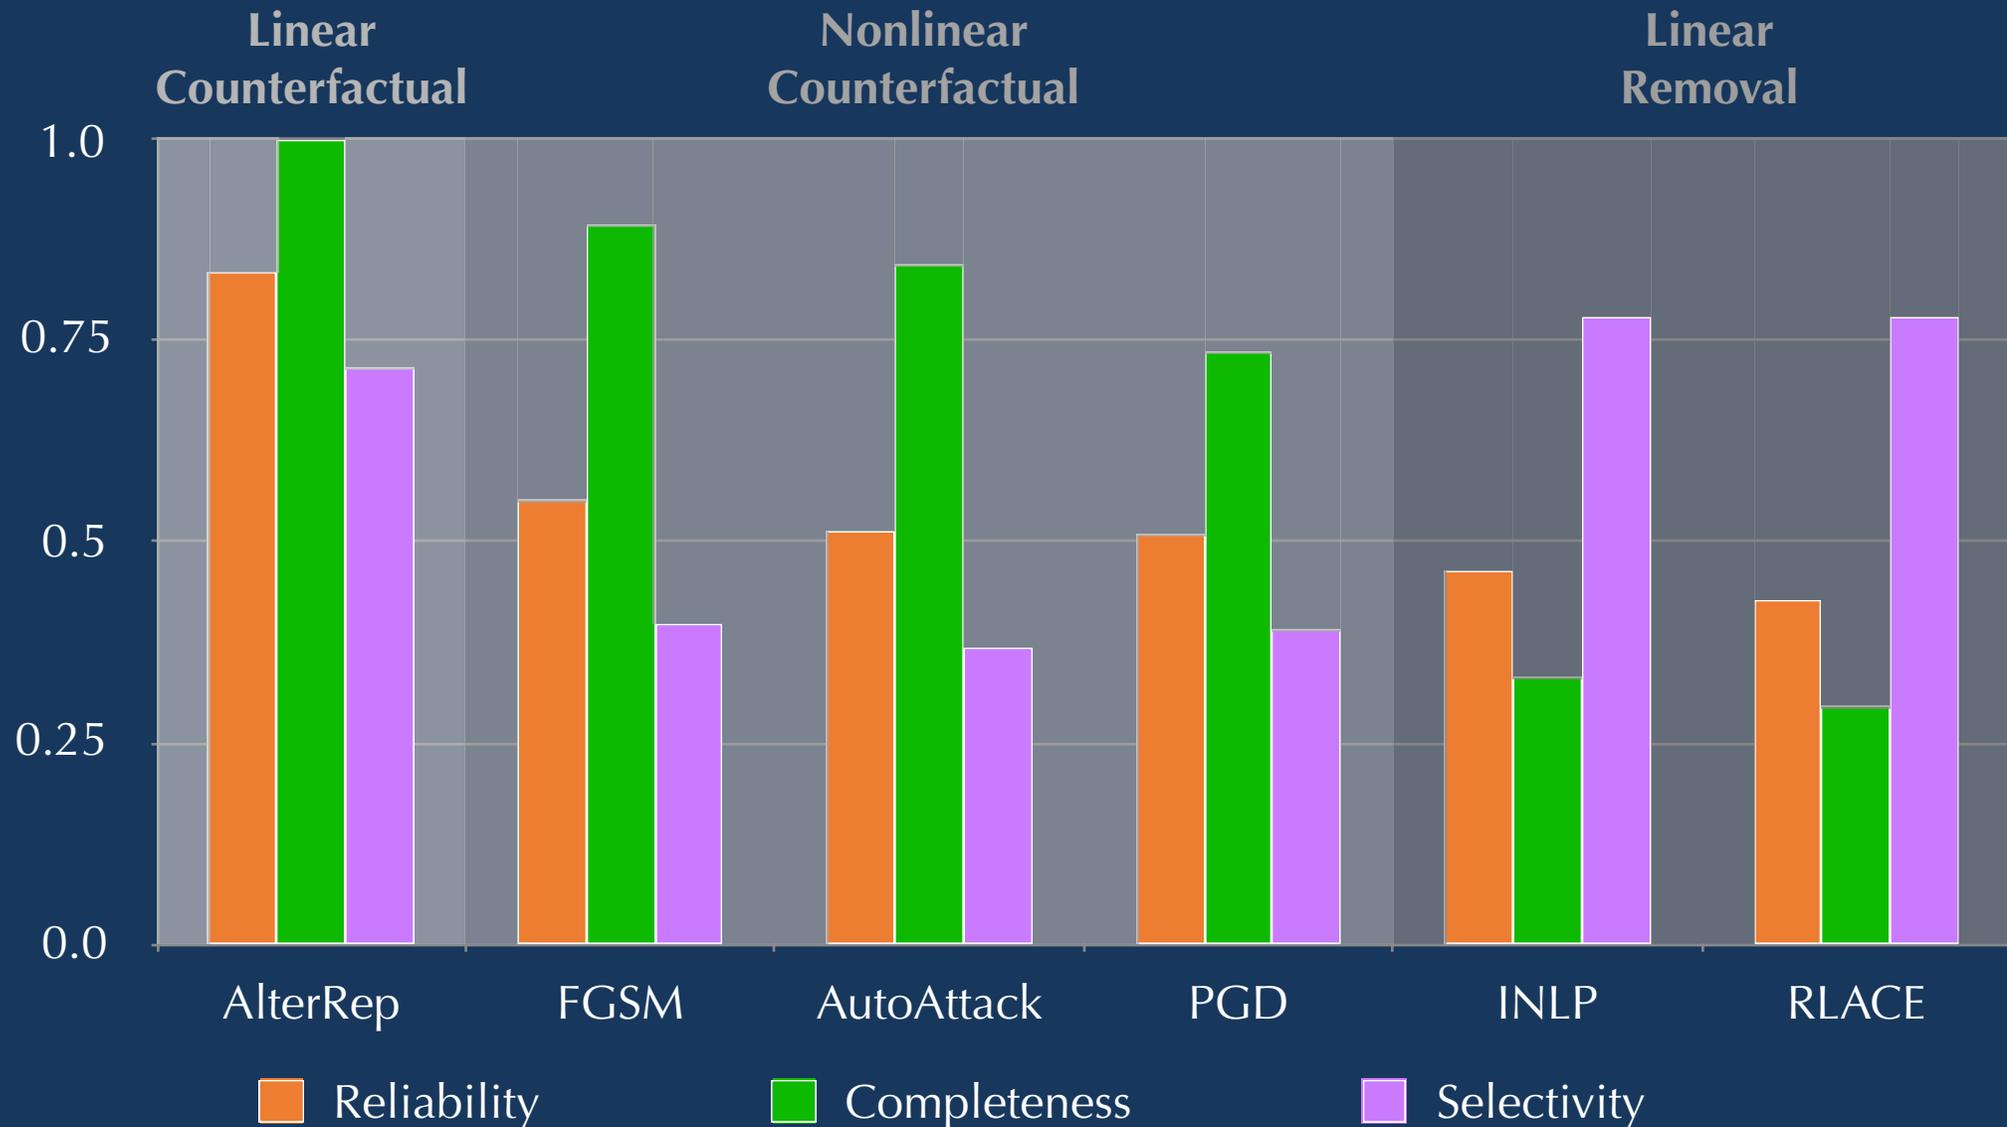


Reliability

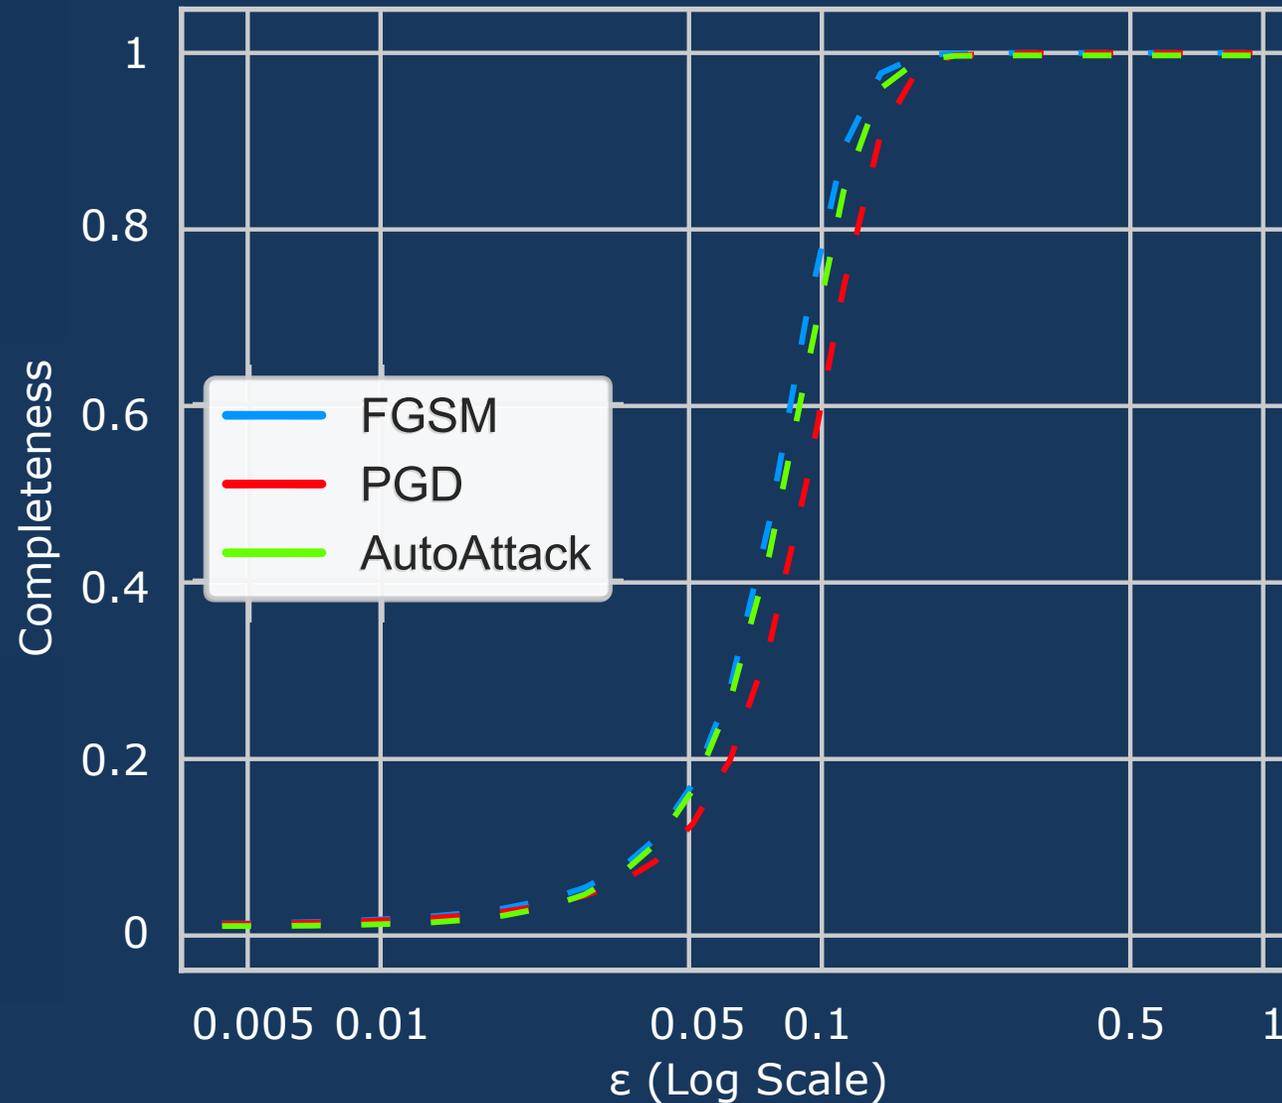
COMPLETENESS: REMOVAL METHODS ARE NOT COMPLETE!



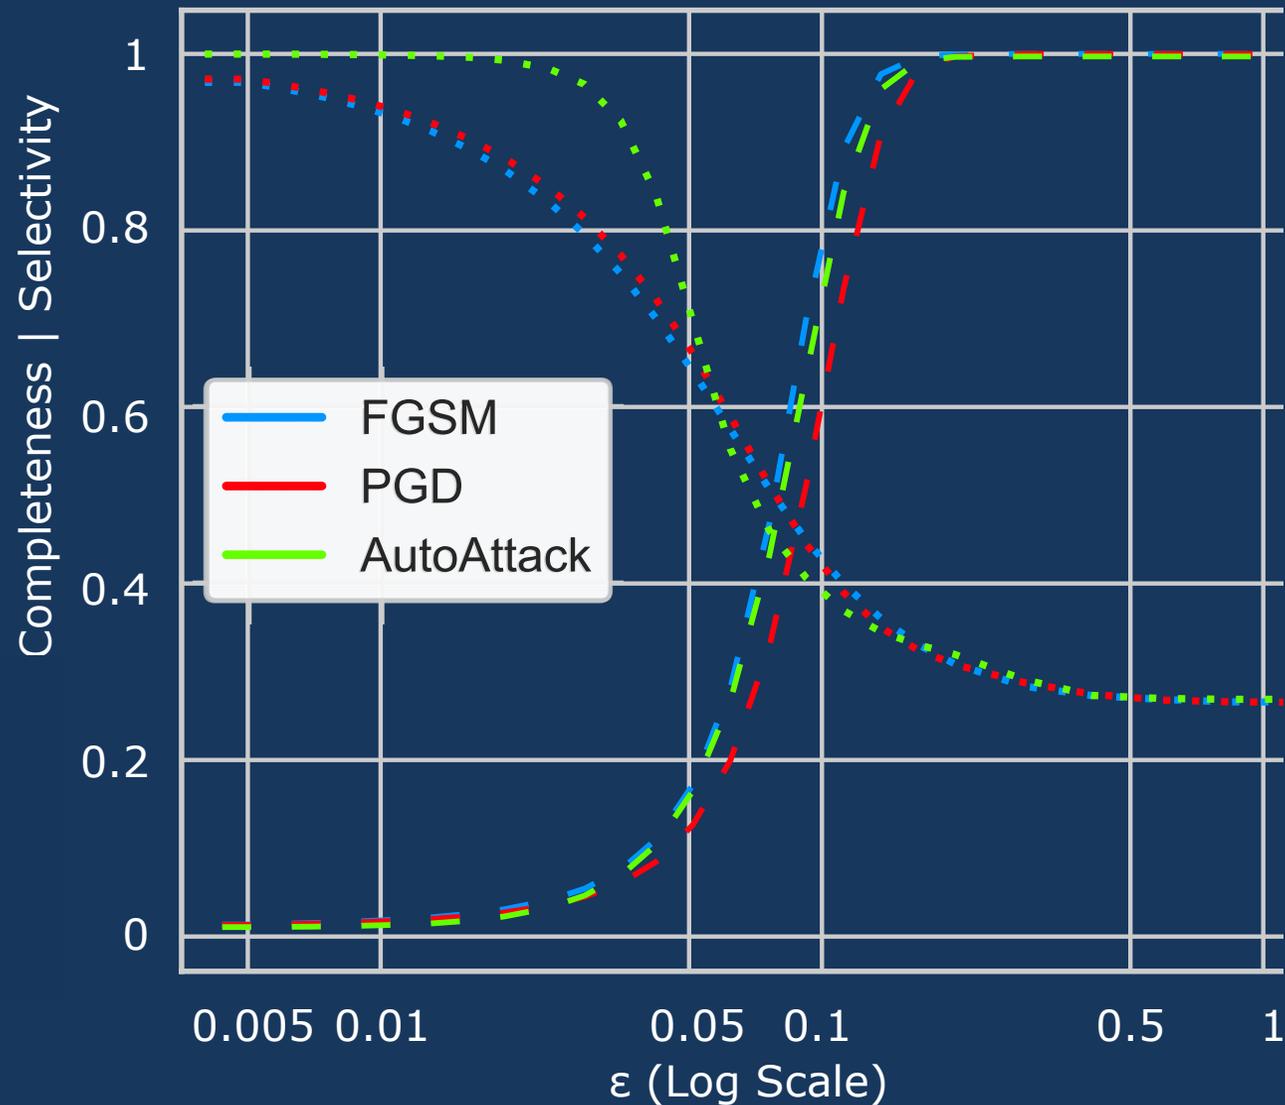
SELECTIVITY: GBIS ARE NOT SELECTIVE; LINEAR METHODS ARE



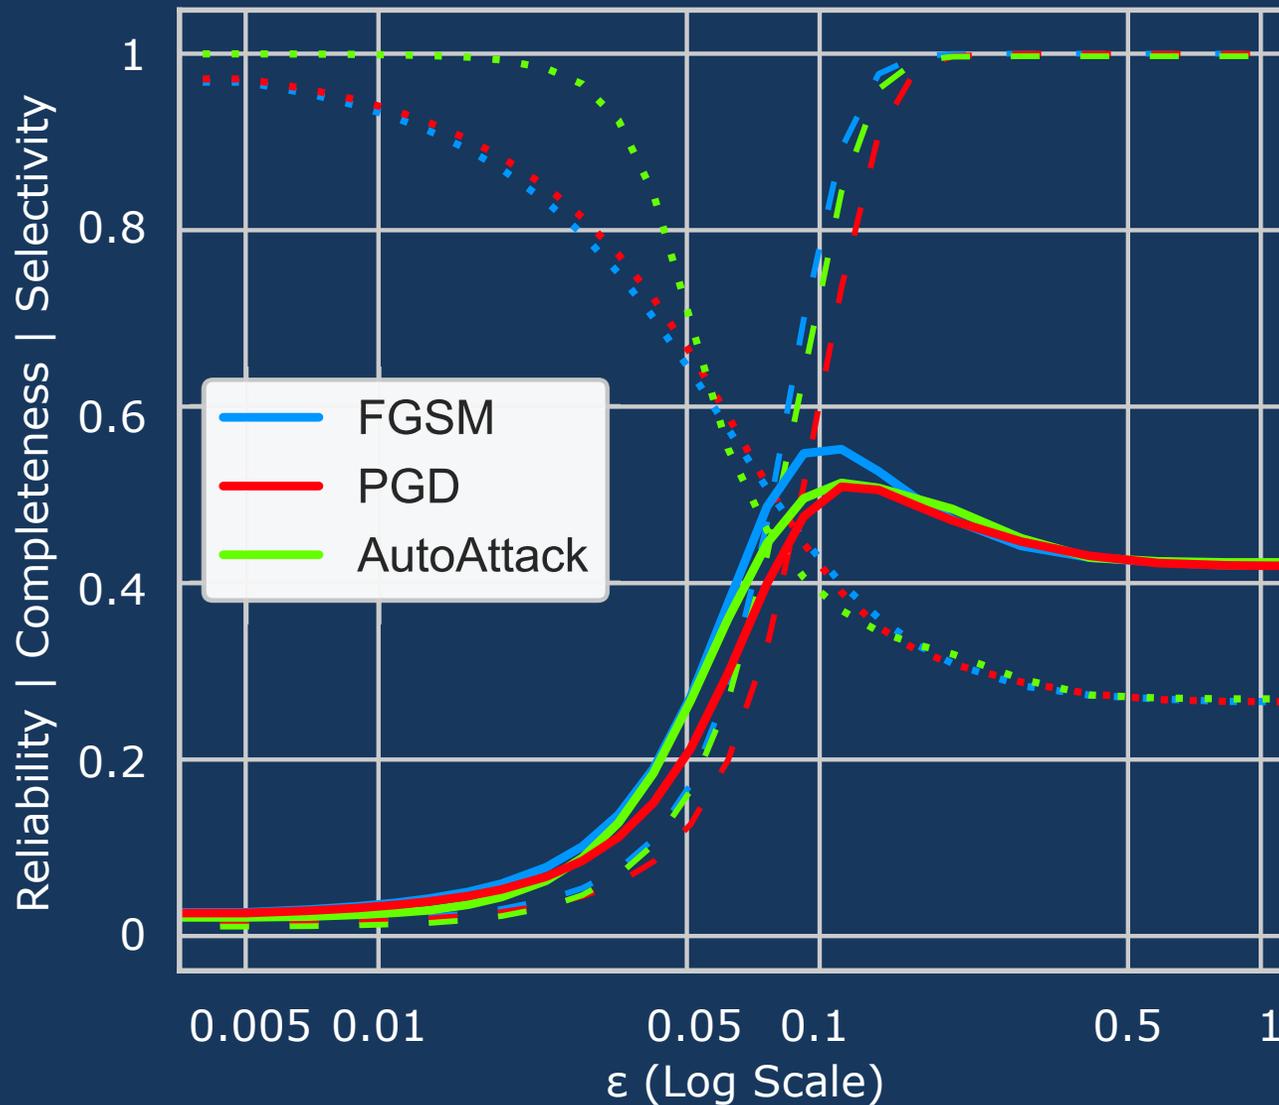
COMPLETENESS AND SELECTIVITY ARE A TRADEOFF!



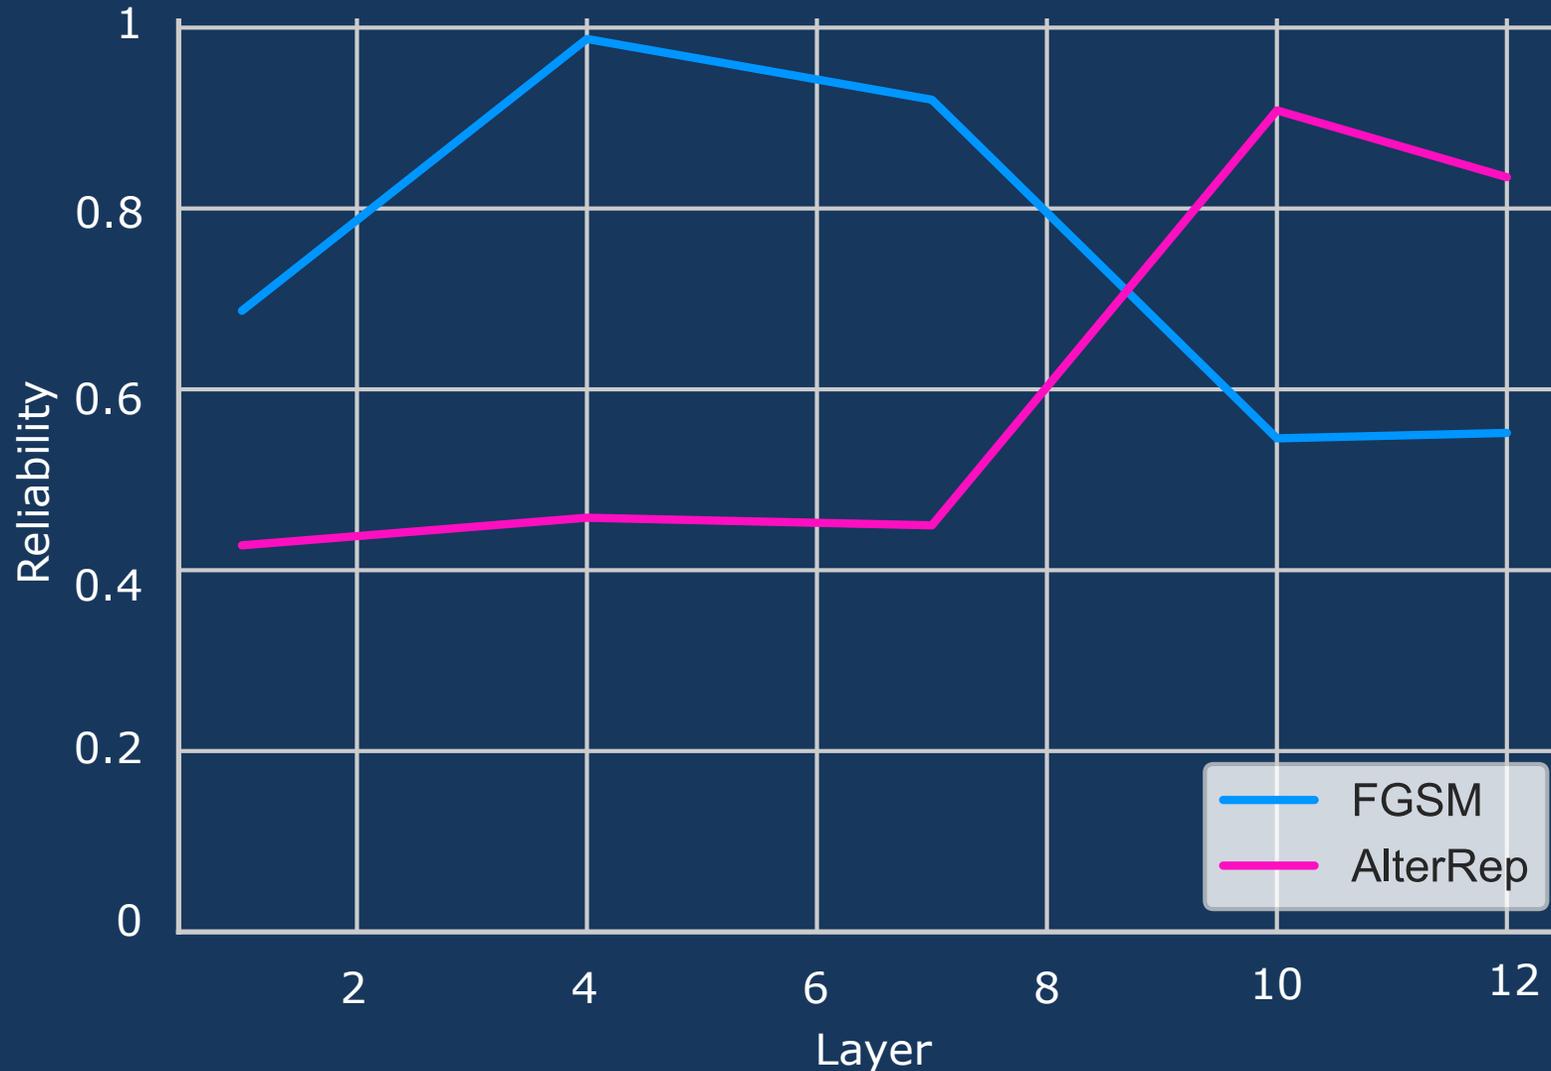
COMPLETENESS AND SELECTIVITY ARE A TRADEOFF!



COMPLETENESS AND SELECTIVITY ARE A TRADEOFF!



FGSM (NONLINEAR) MORE RELIABLE IN EARLIER LAYERS
ALTERREP (LINEAR) MORE RELIABLE IN LATER LAYERS



TAKEAWAYS

We introduce an **evaluation framework** to compare **different classes** of causal probing interventions

- **Tradeoff** between *completeness* and *selectivity*
- Concept removal is **not reliable** (for causal probing)
- Linear interventions better in later layers (less collateral damage)



Questions?

SUPPLEMENTARY MATERIALS FOR QA



MORE RELIABLE METHODS → GREATER Δ IN TASK ACCURACY

