How does Chain of Thought decompose complex tasks?

Amrut Nadgir, Vijay Balasubramanian, Pratik Chaudhari
University of Pennsylvania

LLMs solve a sequence of classification problems

LLMs learn a probability distribution over the next token

LLMs solve a sequence of classification problems

LLMs learn a probability distribution over the next token

Given a new query, it could produce the answer in a single step:

$$(2+2) \times (5+4)$$

… 36 10 12 …

LLMs solve a sequence of classification problems

LLMs learn a probability distribution over the next token

Given a new query, it could produce the answer in a single step:

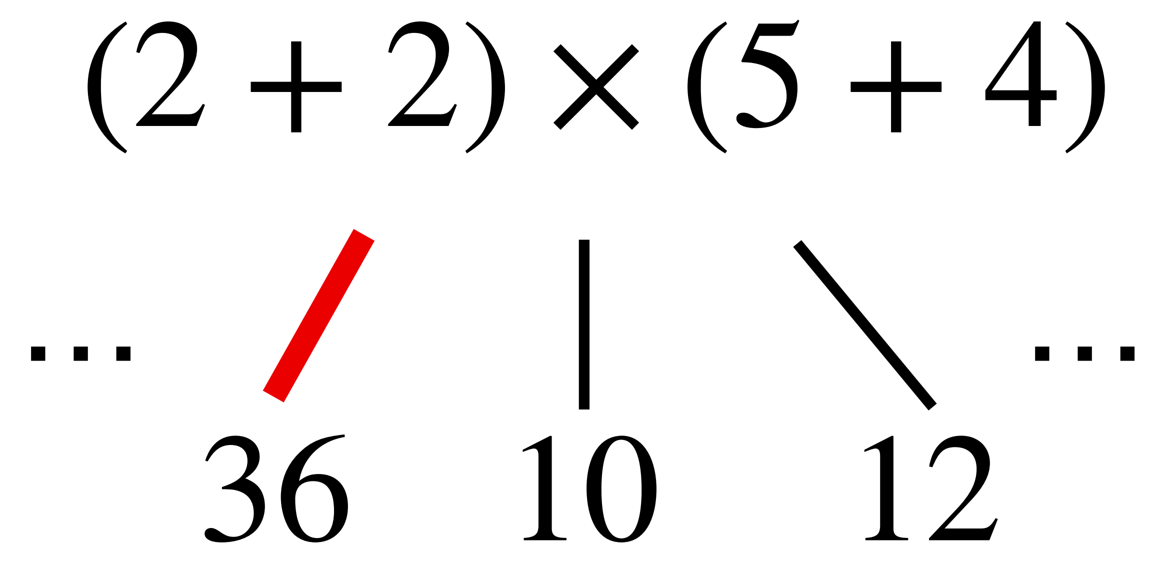

Or in a sequence, creating a tree:

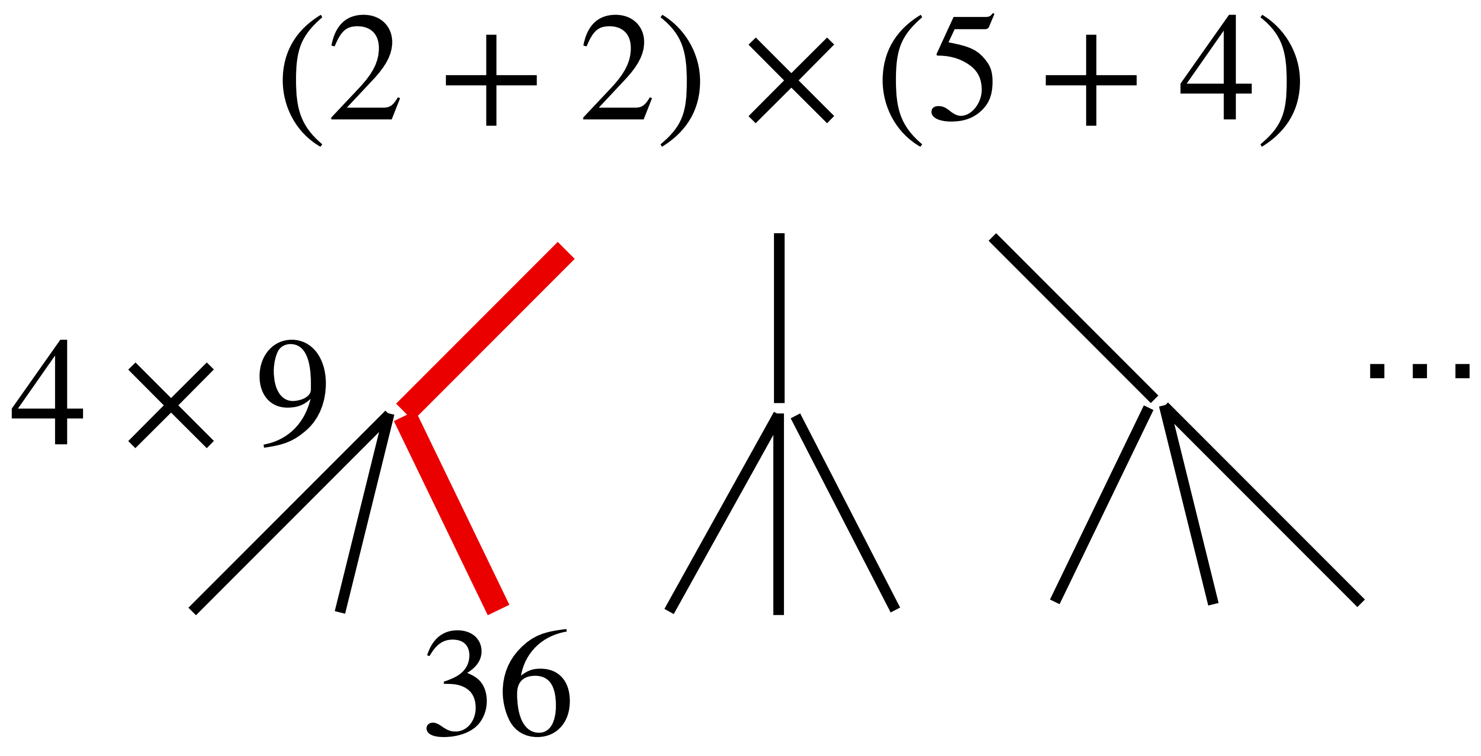

Error increases with the number of classes

The probability of error in classification problems with m classes and D training samples supported on a d-dimensional manifold is

$$\sim \left(\frac{m^2}{D}\right)^{1/d}$$

Error increases with the number of classes

The probability of error in classification problems with m classes and D training samples supported on a d -dimensional manifold is

$$\sim \left(\frac{m^2}{D}\right)^{1/d}$$

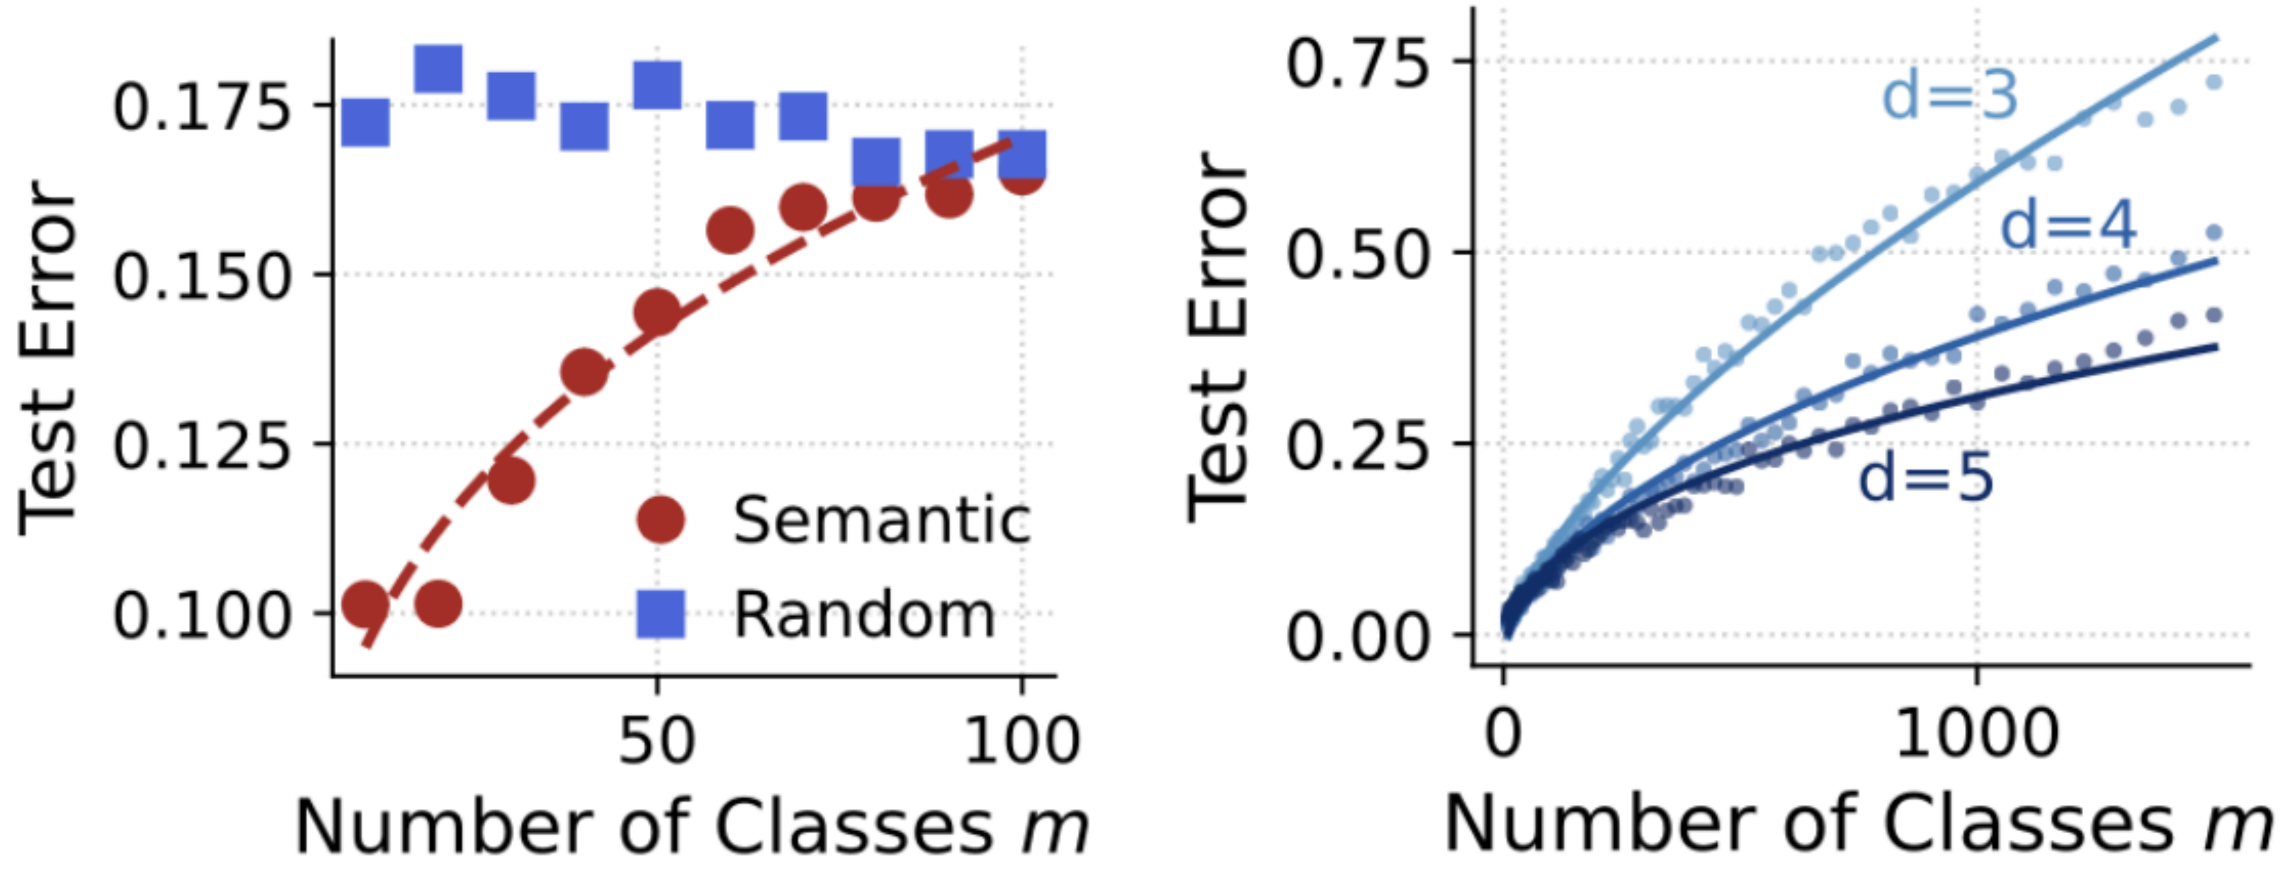

Left: Vision transformer on CIFAR-100;

Right: Synthetic data in a teacher-student setting

There is an optimal depth and degree of the reasoning tree

There is an optimal depth and degree of the reasoning tree

Consider a task with N plausible answers

The error of direct prediction is $\sim N^{2/d}$

There is an optimal depth and degree of the reasoning tree

Consider a task with N plausible answers

The error of direct prediction is $\sim N^{2/d}$

We can alternatively decompose the task into a sequence of n smaller predictions

The error is $\lesssim \sum_{k=1}^{n} m_k^{2/d}$ where $N = \prod_{k=1}^{n} m_k$

There is an optimal depth and degree of the reasoning tree

Consider a task with N plausible answers

The error of direct prediction is $\sim N^{2/d}$

We can alternatively decompose the task into
a sequence of n smaller predictions

The error is $\lesssim \sum_{k=1}^{n} m_k^{2/d}$ where $N = \prod_{k=1}^{n} m_k^{2/d}$

The difference in errors is maximized when

$$\forall k : m_k = m^* = \exp(d/2)$$

There is an optimal depth and degree of the reasoning tree

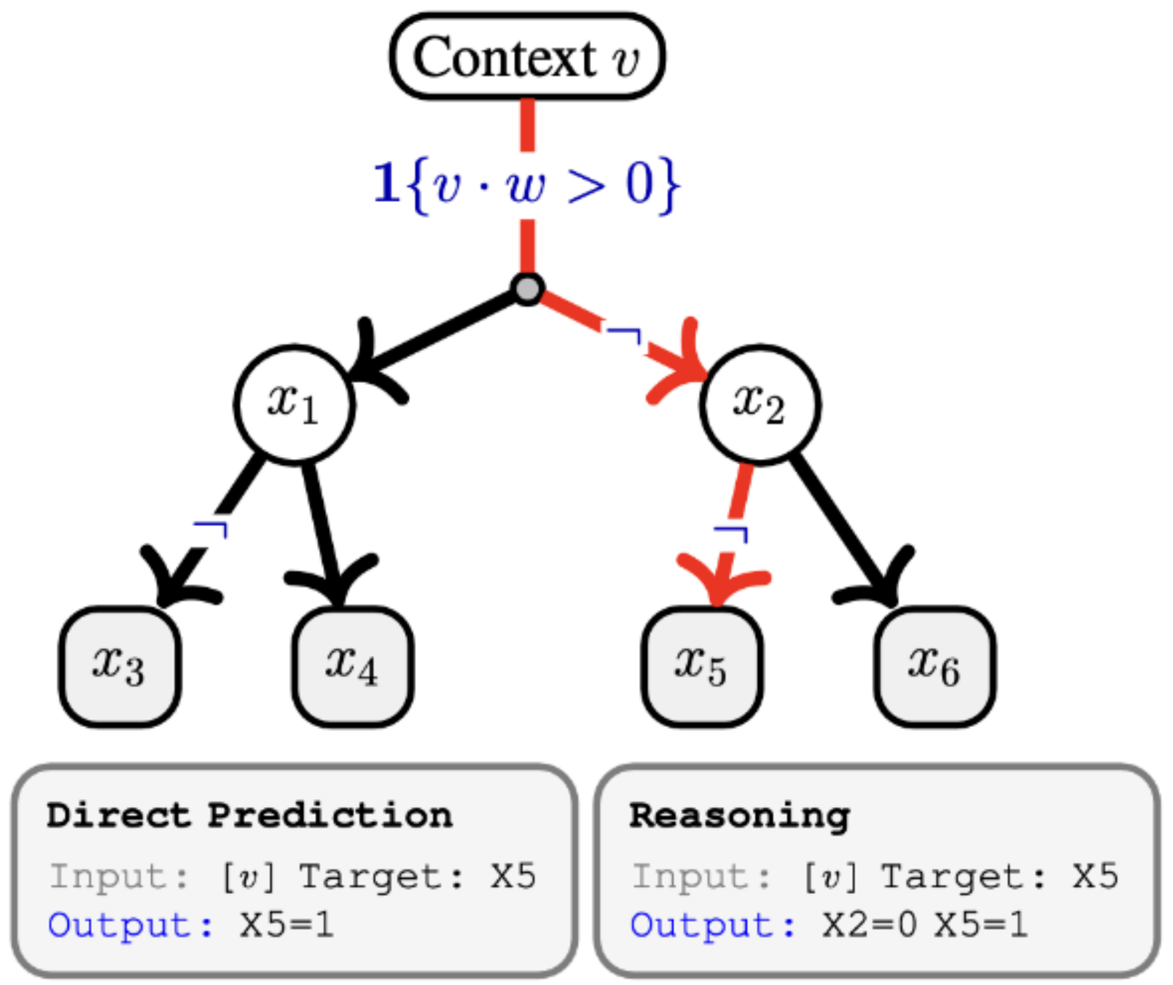

Consider a task with N plausible answers

The error of direct prediction is $\sim N^{2/d}$

We can alternatively decompose the task into a sequence of n smaller predictions

The error is $\lesssim \sum_{k=1}^{n} m_k^{2/d}$ where $N = \prod_{k=1}^{n} m_k^{2/d}$

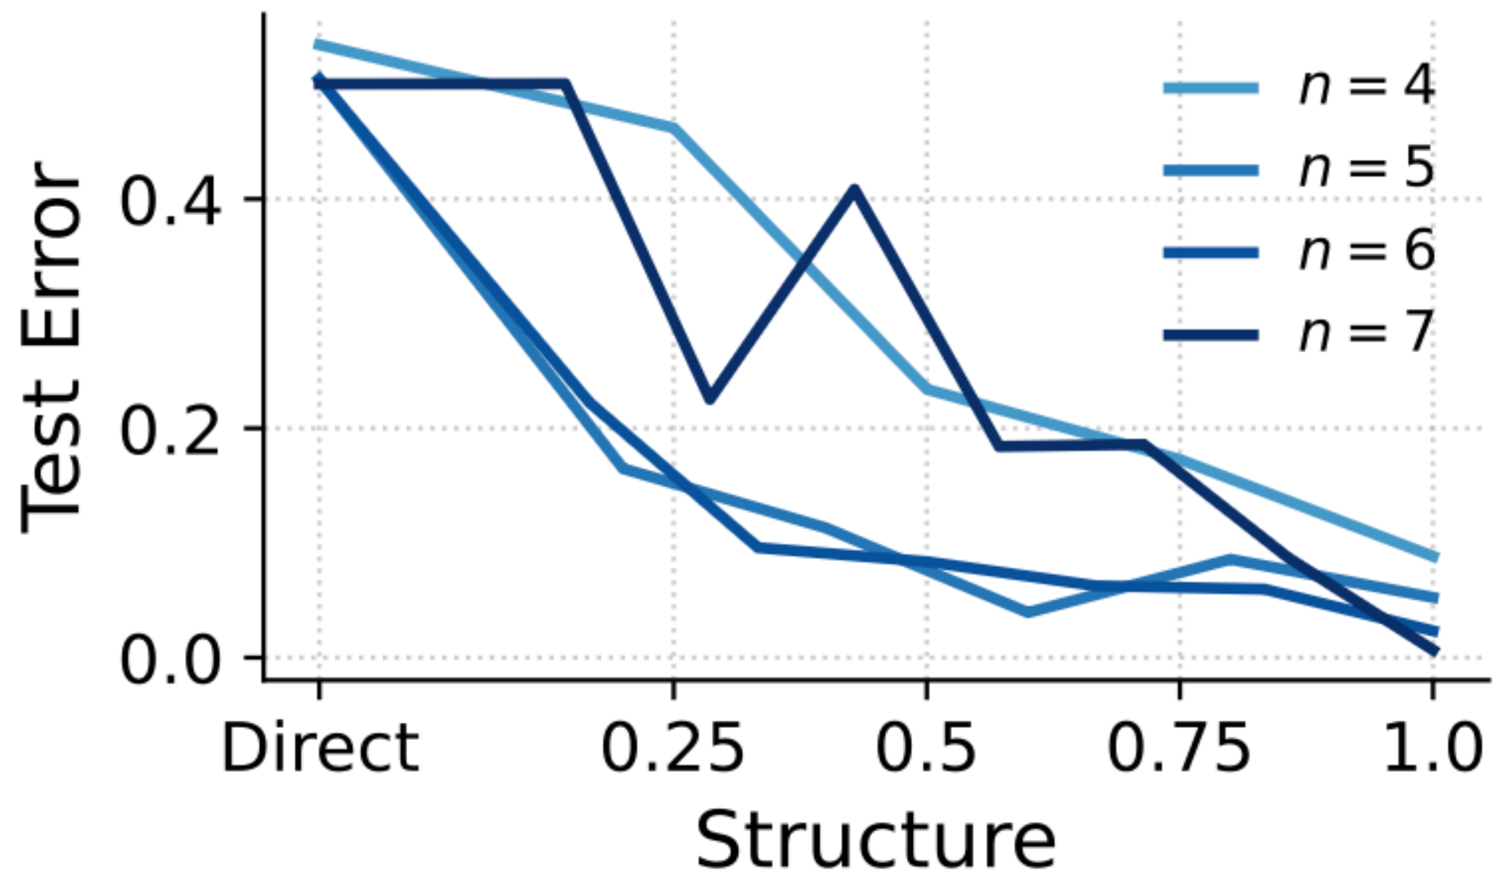

The difference in errors is maximized when

$$\forall k : m_k = m^* = \exp(d/2)$$

Extended reasoning (“thinking”) can be detrimental

Often, the reasoning tree is deeper than necessary and contains redundant paths (red) to the same answers

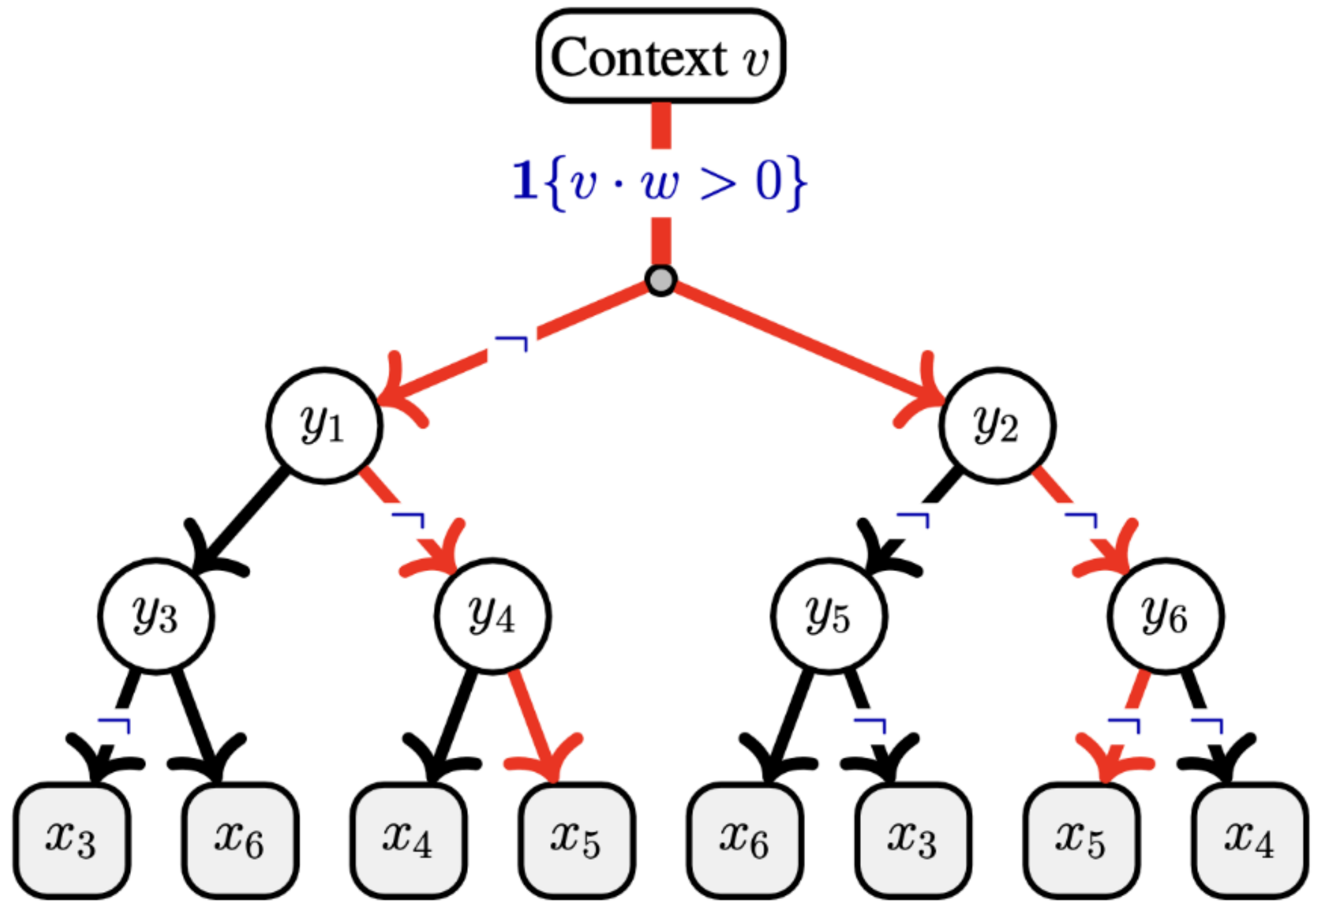

Extended reasoning ("thinking") can be detrimental

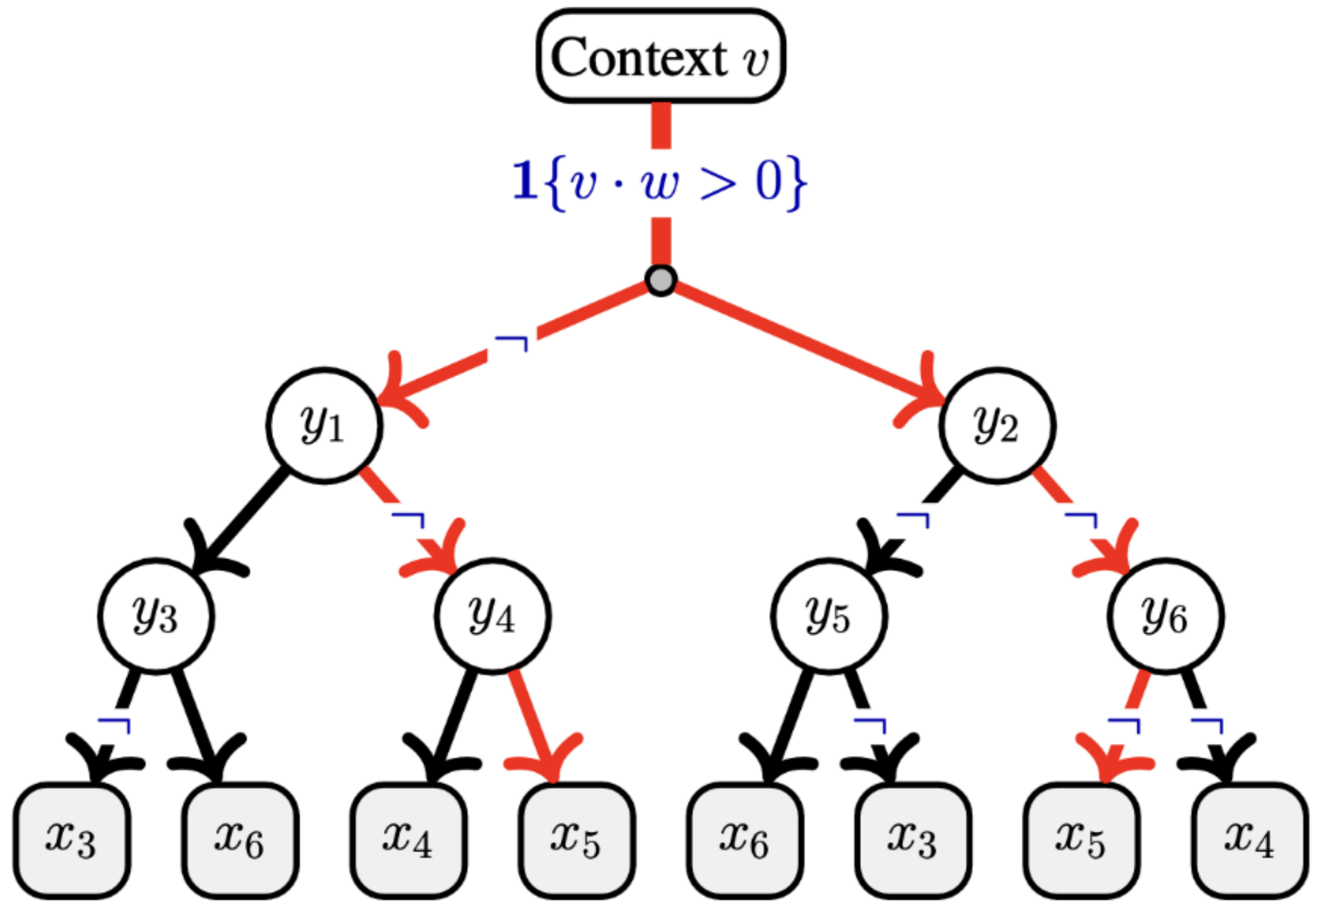

Often, the reasoning tree is deeper than necessary and contains redundant paths (red) to the same answers

We study the case where the depth is increased from $n \to rn$

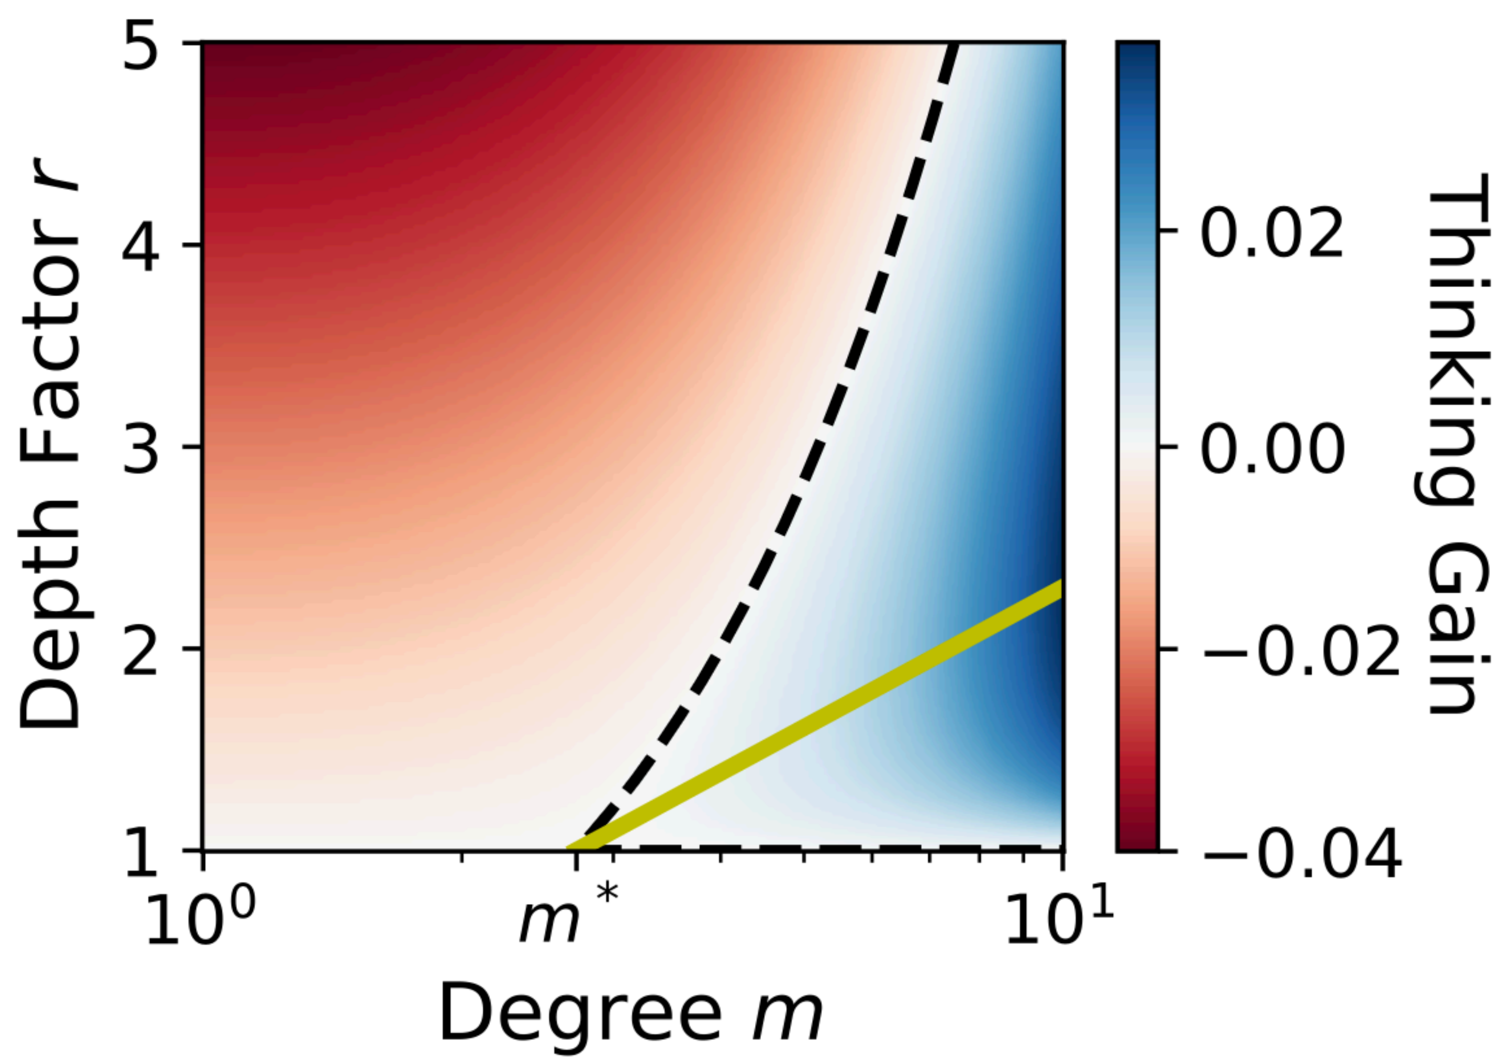

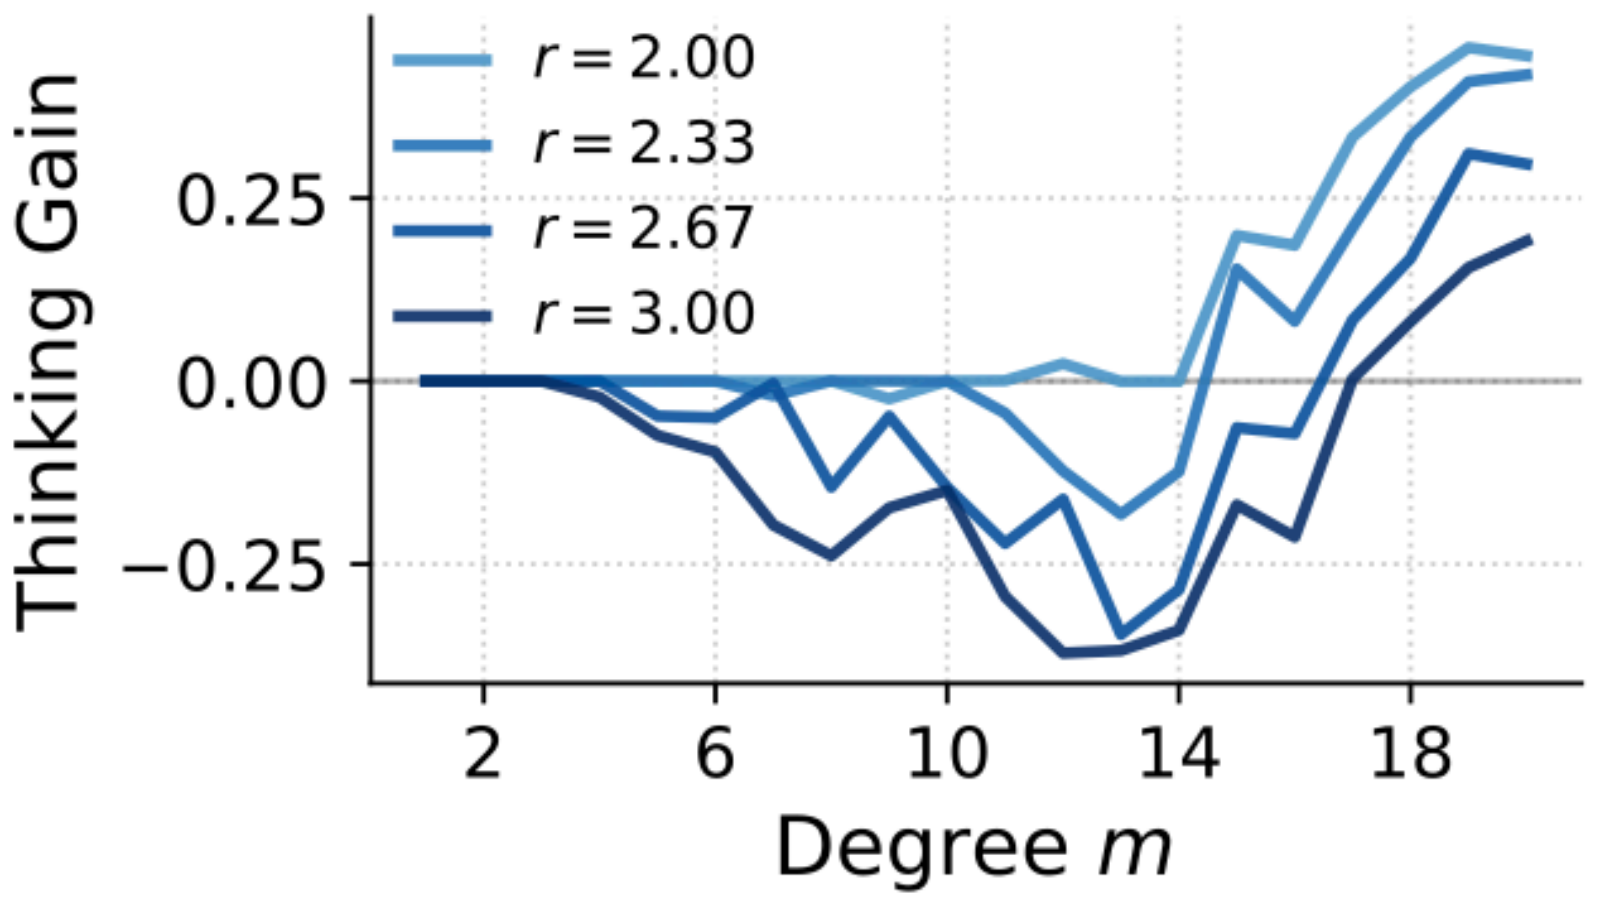

Extended reasoning ("thinking") can be detrimental

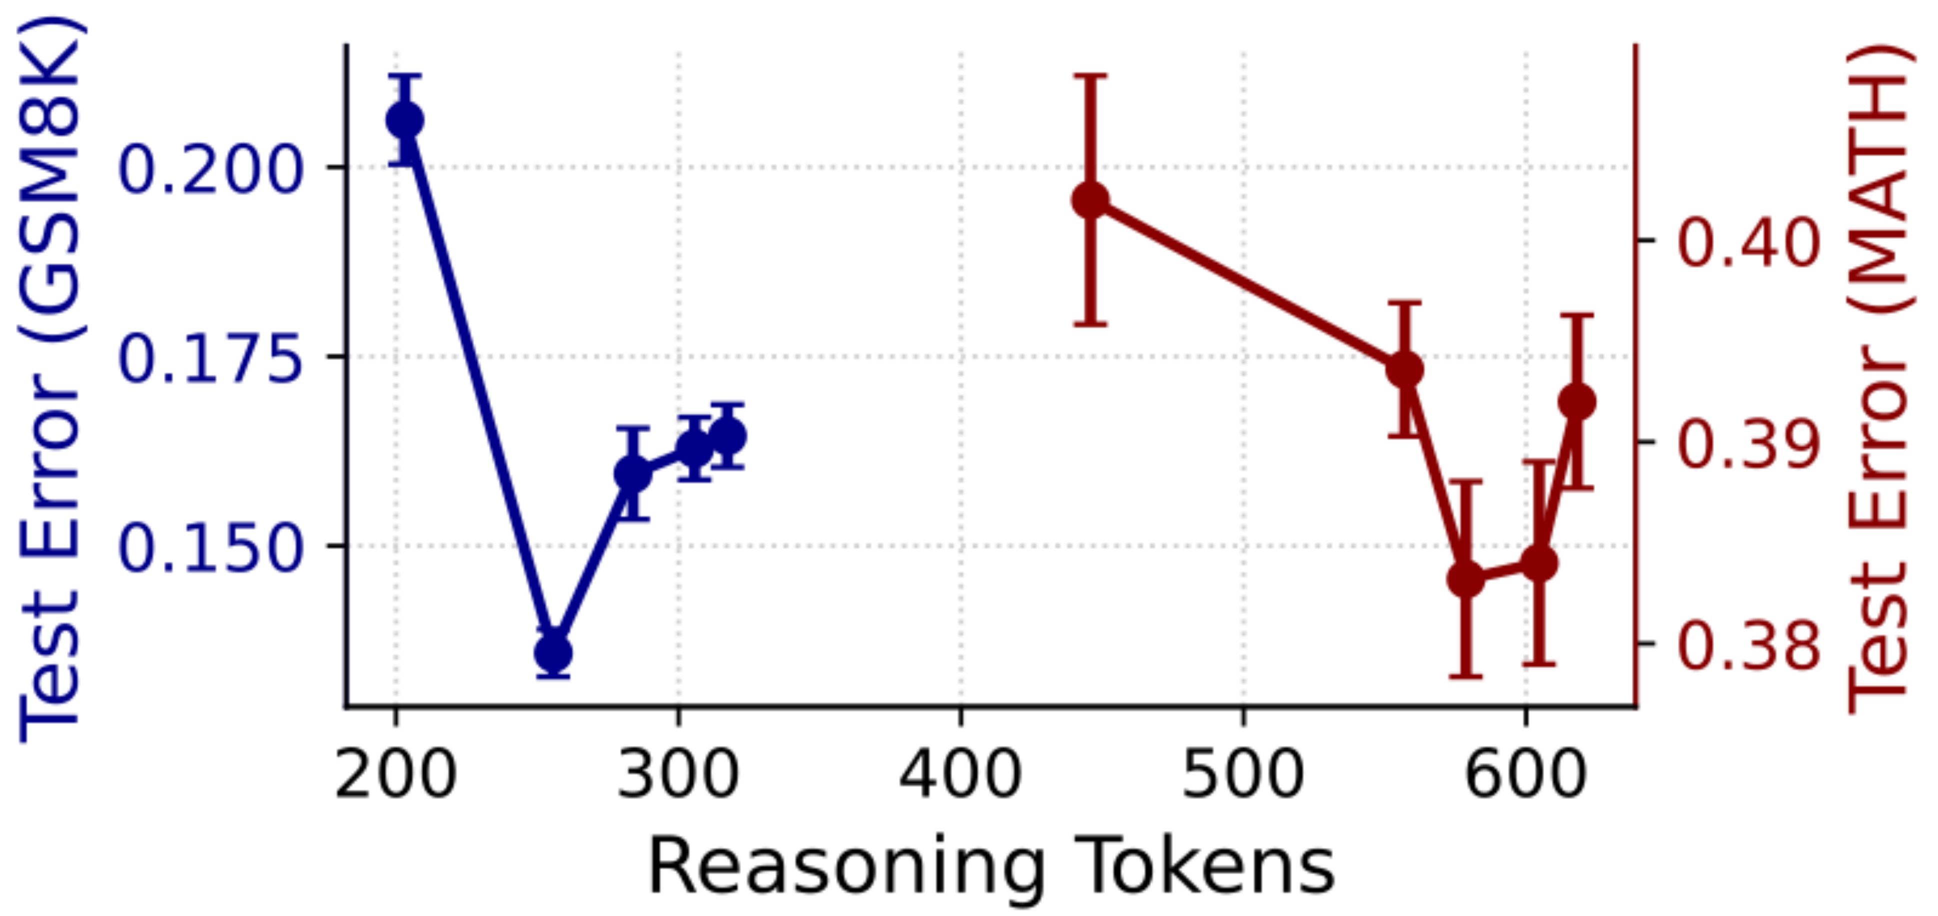

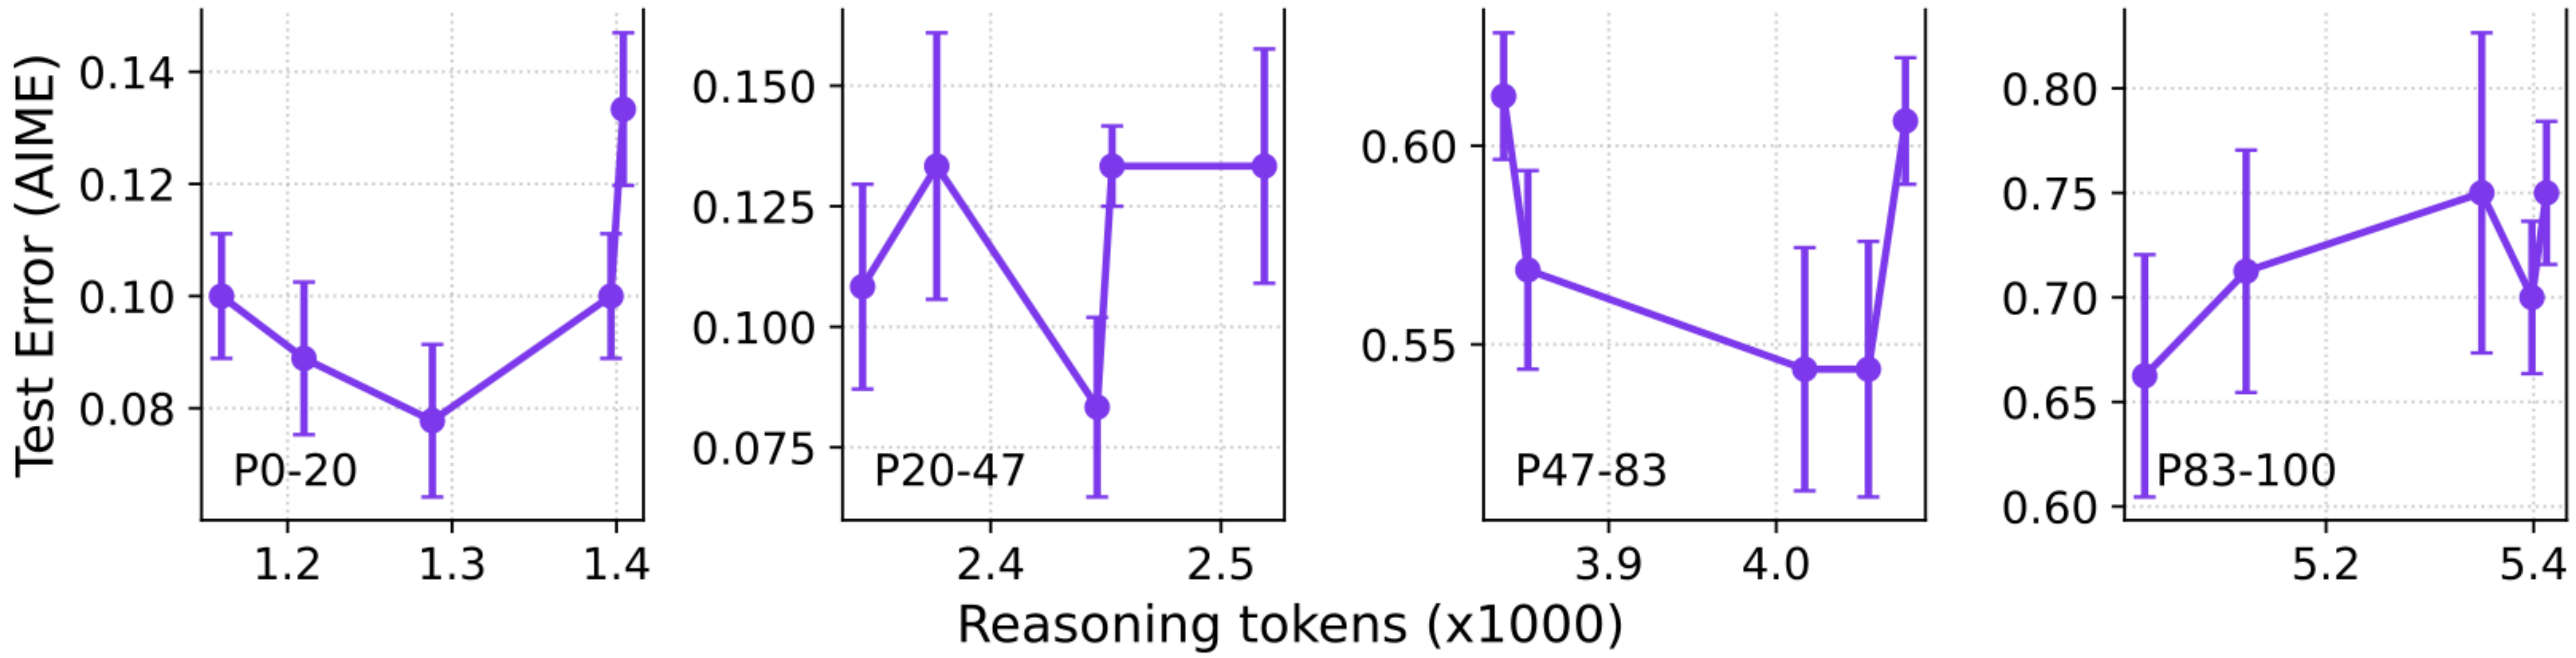

Key Takeaways

- We have shown that CoT is a viable strategy if and only if
  1. the input-output map is not smooth, e.g., if there are many possible answers
  2. the problem can be broken down into a sequence of simpler steps
- The reasoning tree is a simple and clear model for understanding CoT
- The reasoning tree should be structured and has an optimal depth