

# Decoding from Neural Text Generation Models

Antoine Bosselut



# Generation Model Basics

1. At each time step, model computes a score  $o_n$  for each token in our vocabulary,  $w_n \in V$

$$O_n = f(\{y\}_{<t})$$



$f(\cdot)$  is your model

# Generation Model Basics

1. At each time step, model computes a score  $o_n$  for each token in our vocabulary,  $w_n \in V$

$$O_n = f(\{y\}_{<t})$$

$f(\cdot)$  is your model

2. Compute a probability distribution over these scores (usually softmax)

$P(\cdot)$  is your distribution over tokens

$$P(y_t = w_n | \{y\}_{<t}) = \frac{e^{o_n}}{\sum_{m=1}^M e^{o_m}}$$

# Generation Model Basics

1. At each time step, model computes a score  $o_n$  for each token in our vocabulary,  $w_n \in V$

$$O_n = f(\{y\}_{<t})$$

$f(\cdot)$  is your model

2. Compute a probability distribution over these scores (usually softmax)

$P(\cdot)$  is your distribution over tokens

$$P(y_t = w_n | \{y\}_{<t}) = \frac{e^{o_n}}{\sum_{m=1}^M e^{o_m}}$$

3. Define a function to select a token from this distribution


$$\hat{y}_t = g(P(y_t | \{y\}_{<t}))$$

$g(\cdot)$  is your decoding algorithm

# Simplest approach: Argmax Decoding

- $g$  = select the token with the highest probability:

$$\hat{y}_t = \underset{w \in V}{\operatorname{argmax}} P(y_t = w \mid \{y\}_{<t})$$



# Simplest approach: Argmax Decoding

- $g$  = select the token with the highest probability:

$$\hat{y}_t = \underset{w \in V}{\operatorname{argmax}} P(y_t = w \mid \{y\}_{<t})$$

Select **highest scoring** token


He wanted to go to the

Model



# Maybe we need more options: Beam Search

- $g$  = cache  $b$  paths for two steps





# Maybe we need more options: Beam Search


- $g$  = cache  $b$  paths for two steps

If  $b = 2$ , select top two tokens

He wanted to go to the



- restroom
- grocery
- store
- airport
- pub
- gym
- bathroom
- game
- beach
- hospital
- doctor






# Maybe we need more options: Beam Search

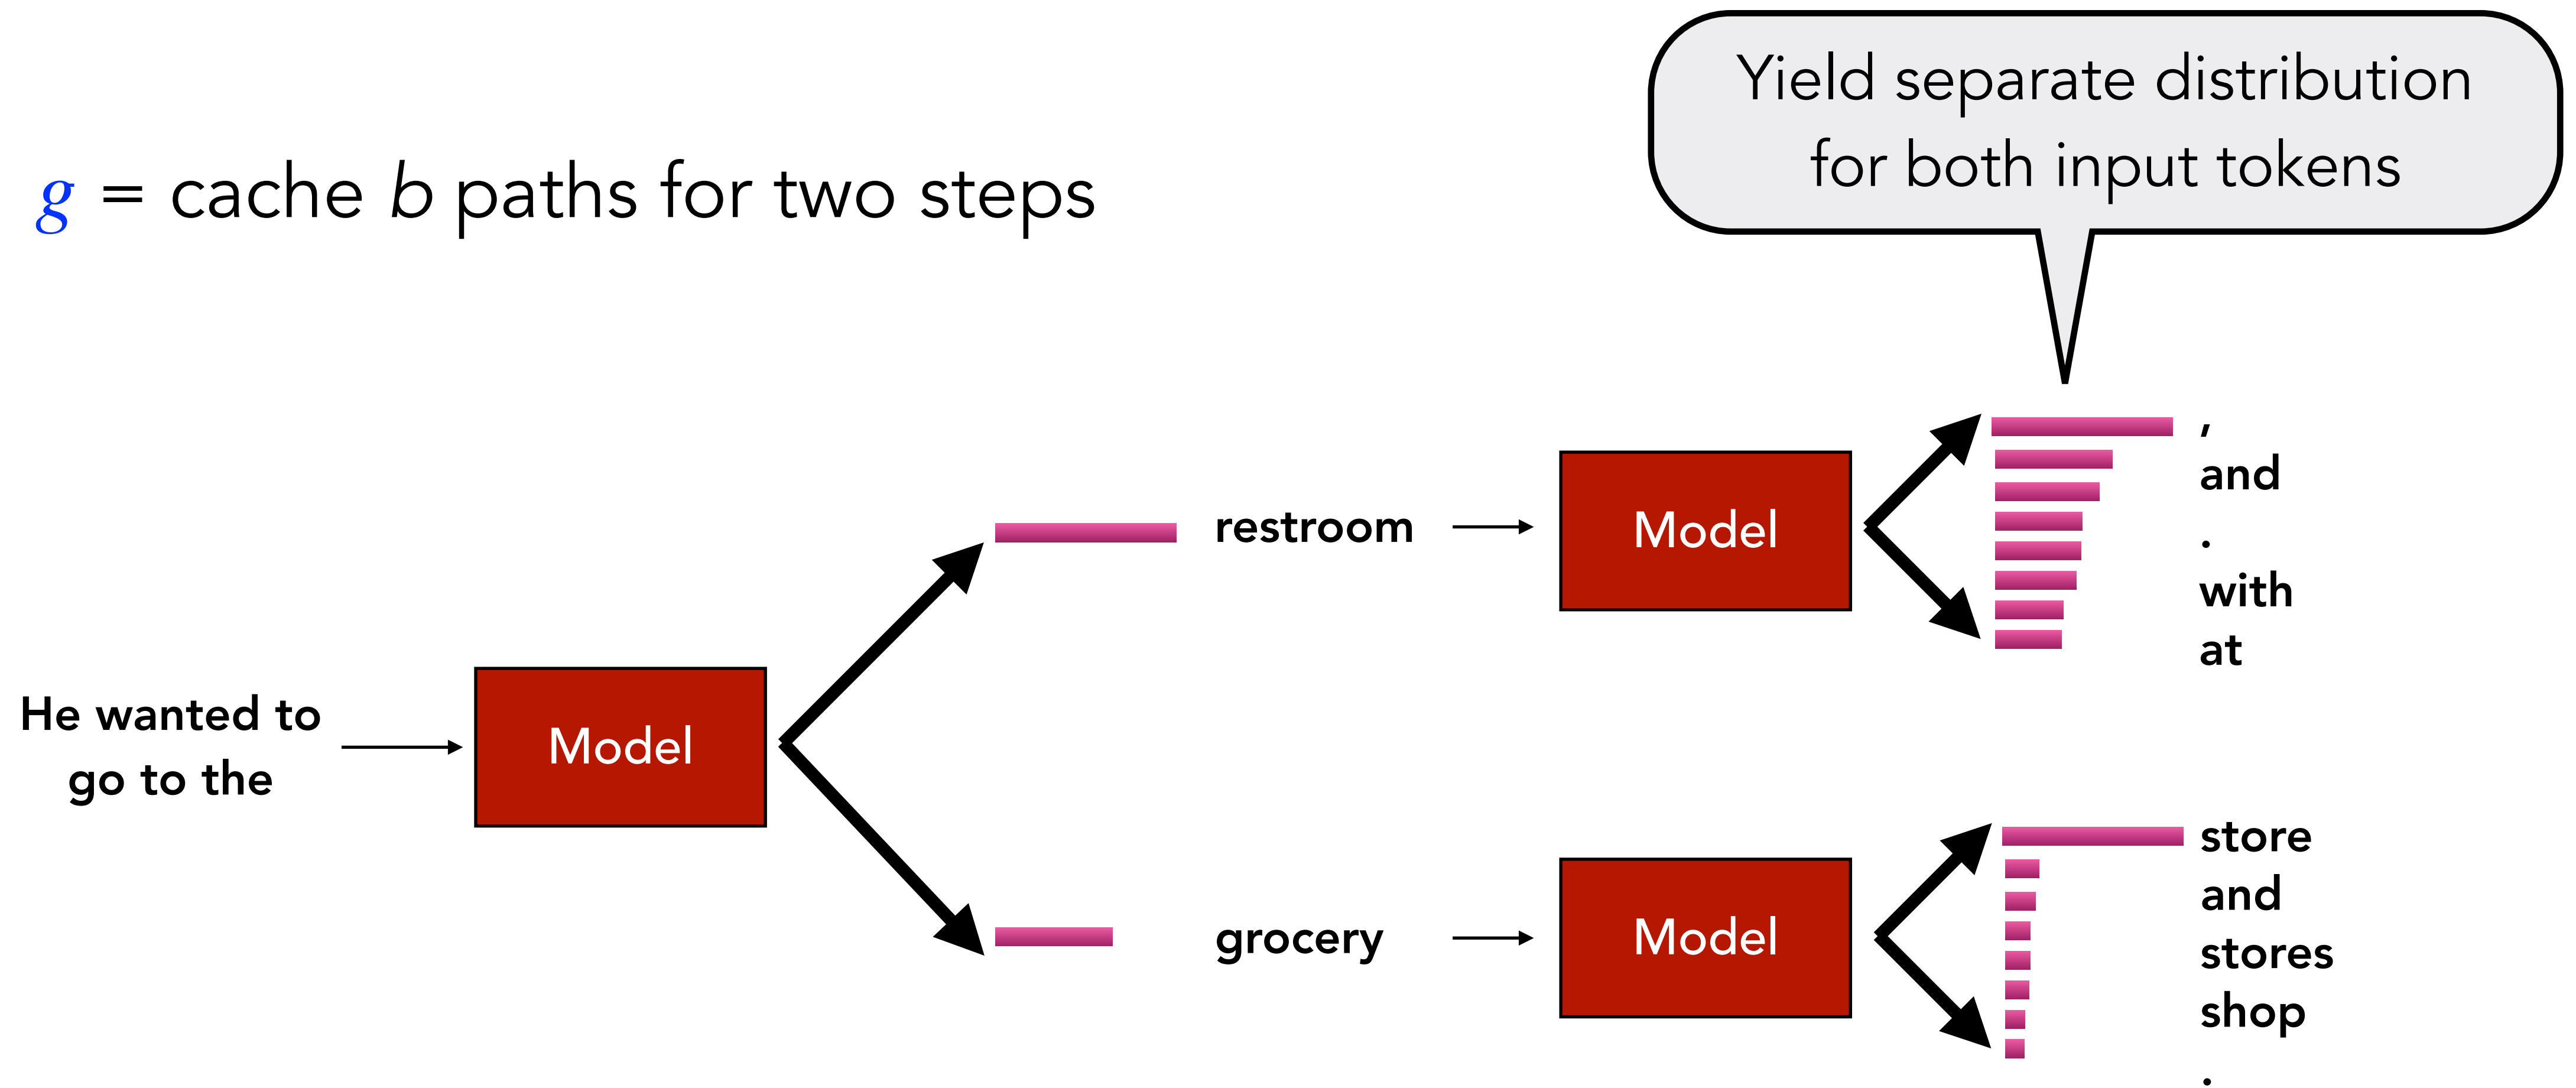
- $g$  = cache  $b$  paths for two steps

Use them both as inputs to the decoder at next step



# Maybe we need more options: Beam Search

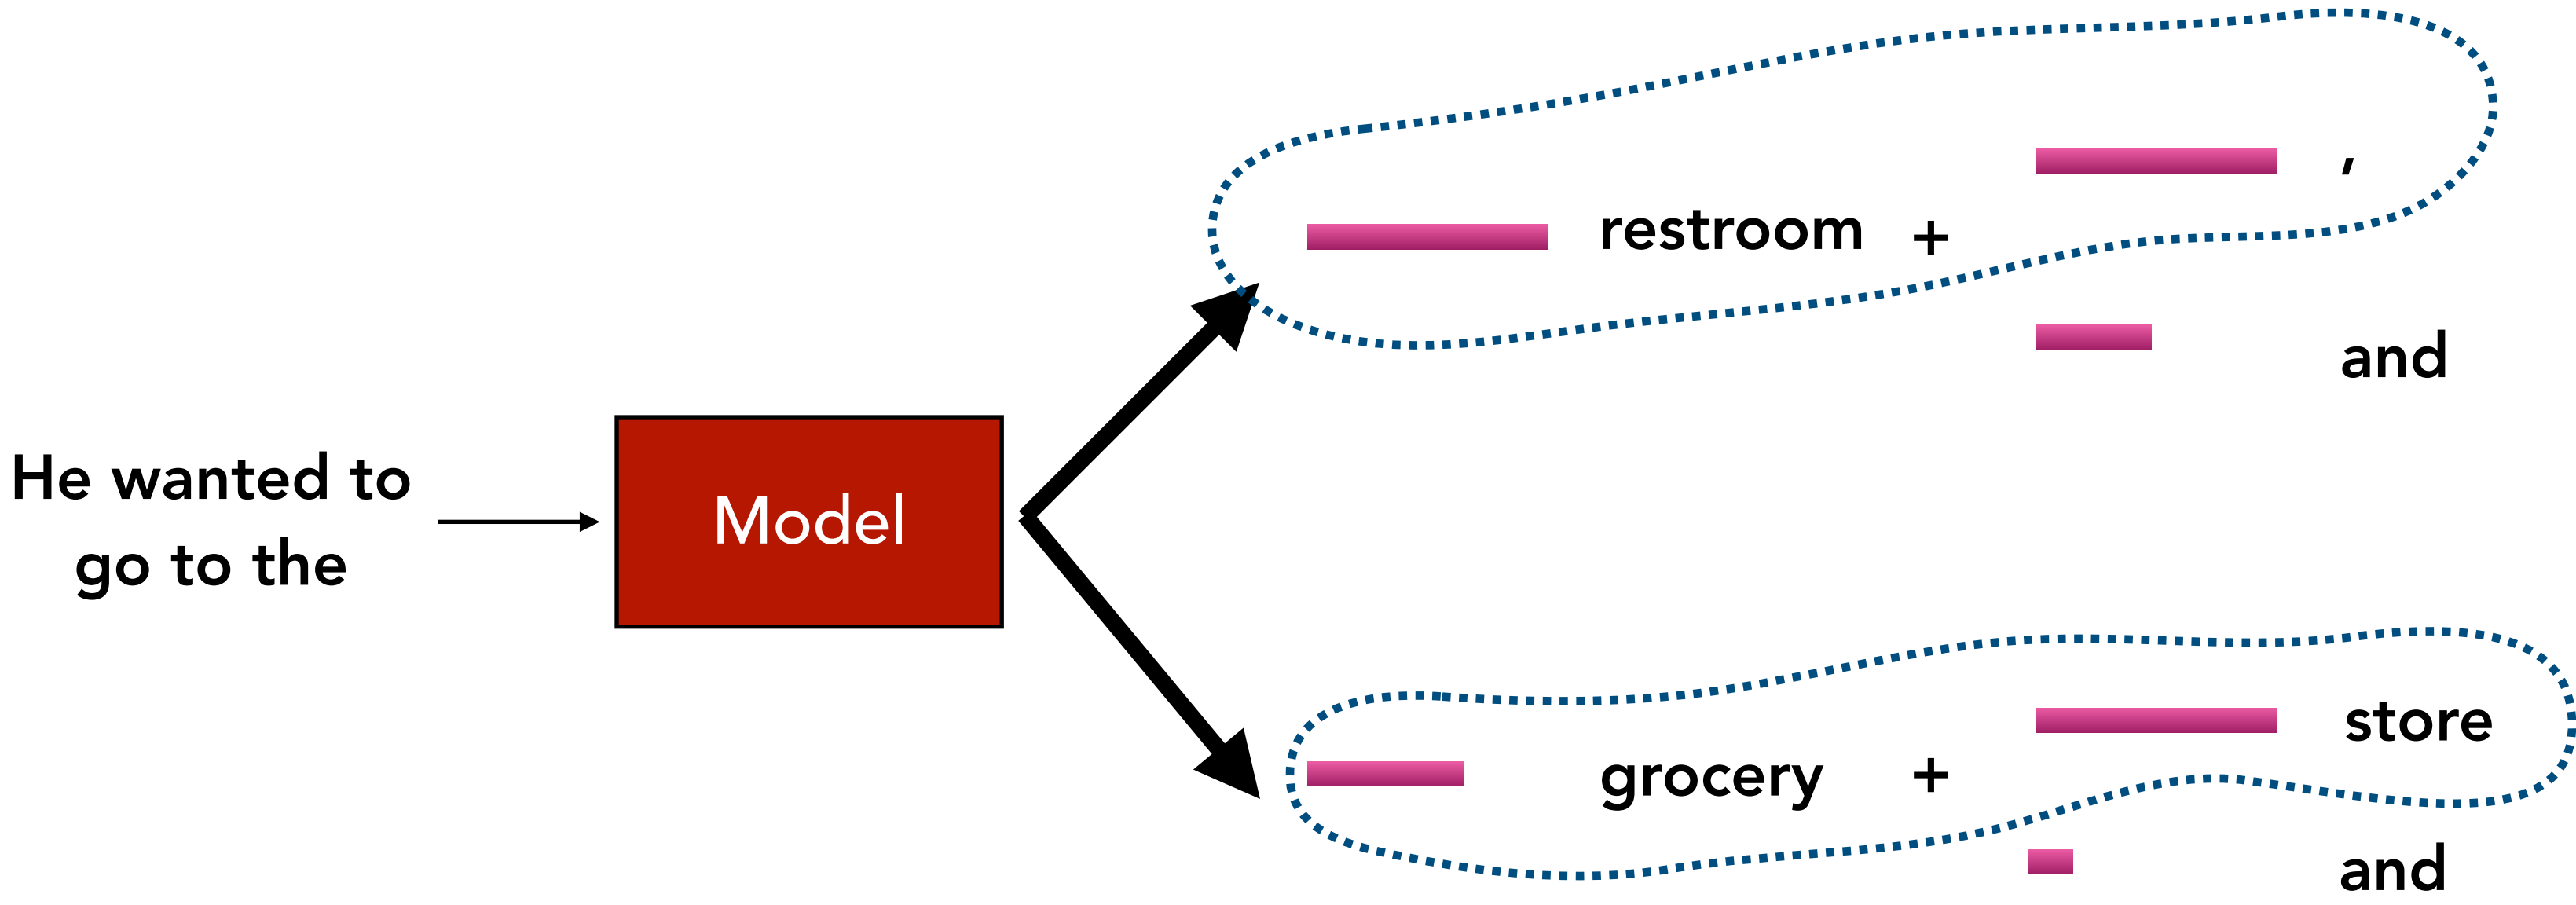
- $g$  = cache  $b$  paths for two steps



# Maybe we need more options: Beam Search

Select top **b** sequence continuations across both distributions



- $g$  = cache  $b$  paths for two steps



# Maybe we need more options: Beam Search

- $g$  = cache  $b$  paths for two steps

Repeat!



# Does this penalize longer sequences?

$$s(Y) = \sum_{t=1}^T \log P(y_t | \{y\}_{<t})$$

**Shorter** sequences will score better!

# Does this penalize longer sequences?

- **Solution:** Normalize by token length of sequence

$$s(Y) = \frac{1}{|Y|} \sum_{t=1}^{|Y|} \log P(y_t | \{y\}_{<t})$$

- **Solution:** Normalize by token length relative to reference sequence

$$s(Y) = \frac{1}{lp(Y)} \sum_{t=1}^{|Y|} \log P(y_t | \{y\}_{<t}) \quad lp(Y) = \frac{(5 + |Y|)^\alpha}{(5 + 1)^\alpha}$$

# Beam search gets repetitive and repetitive

**Context:** In a shocking finding, scientist discovered a herd of unicorns living in a remote, previously unexplored valley, in the Andes Mountains. Even more surprising to the researchers was the fact that the unicorns spoke perfect English.

**Continuation:** The study, published in the Proceedings of the National Academy of Sciences of the United States of America (PNAS), was conducted by researchers from the **Universidad Nacional Autónoma de México (UNAM)** and **the Universidad Nacional Autónoma de México (UNAM/Universidad Nacional Autónoma de México/ Universidad Nacional Autónoma de México/ Universidad Nacional Autónoma de México/ Universidad Nacional Autónoma de México...**

# Beam search gets repetitive and repetitive

**Context:** In a shocking finding, scientist discovered a herd of unicorns living in a remote, previously unexplored valley in the Andes Mountains. Even


**Continuation**

Repetition is a big problem  
in text generation!


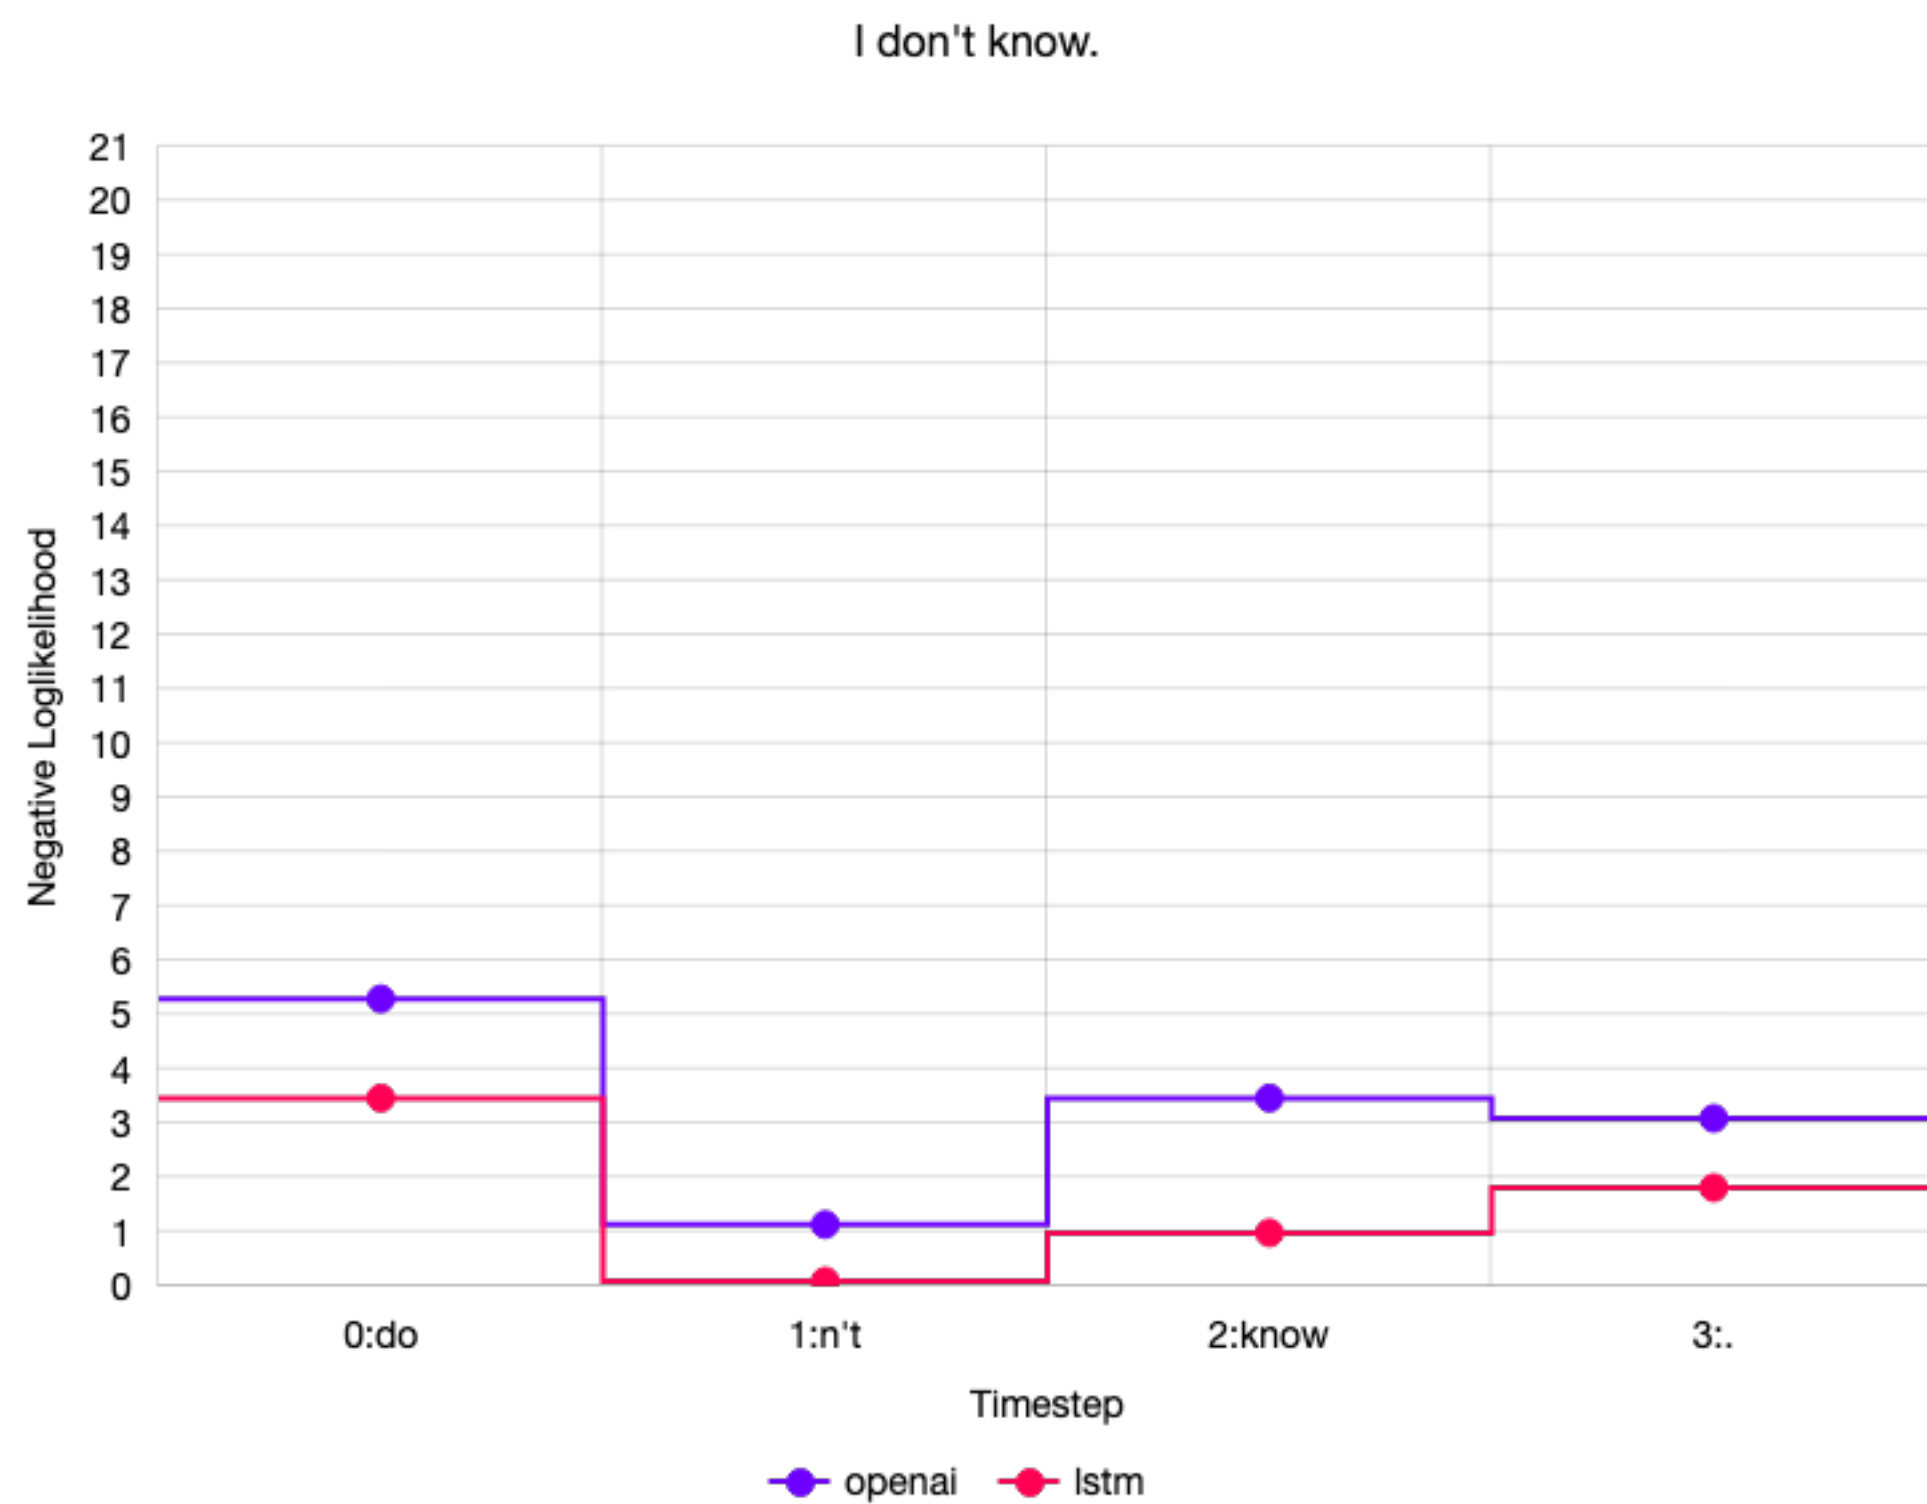
**Universidad Nacional Autónoma de Mexico (UNAM)**  
and **the Universidad Nacional Autónoma de México**  
**(UNAM/Universidad Nacional Autónoma de México/**  
**Universidad Nacional Autónoma de México/**  
**Universidad Nacional Autónoma de México/**  
**Universidad Nacional Autónoma de México...**



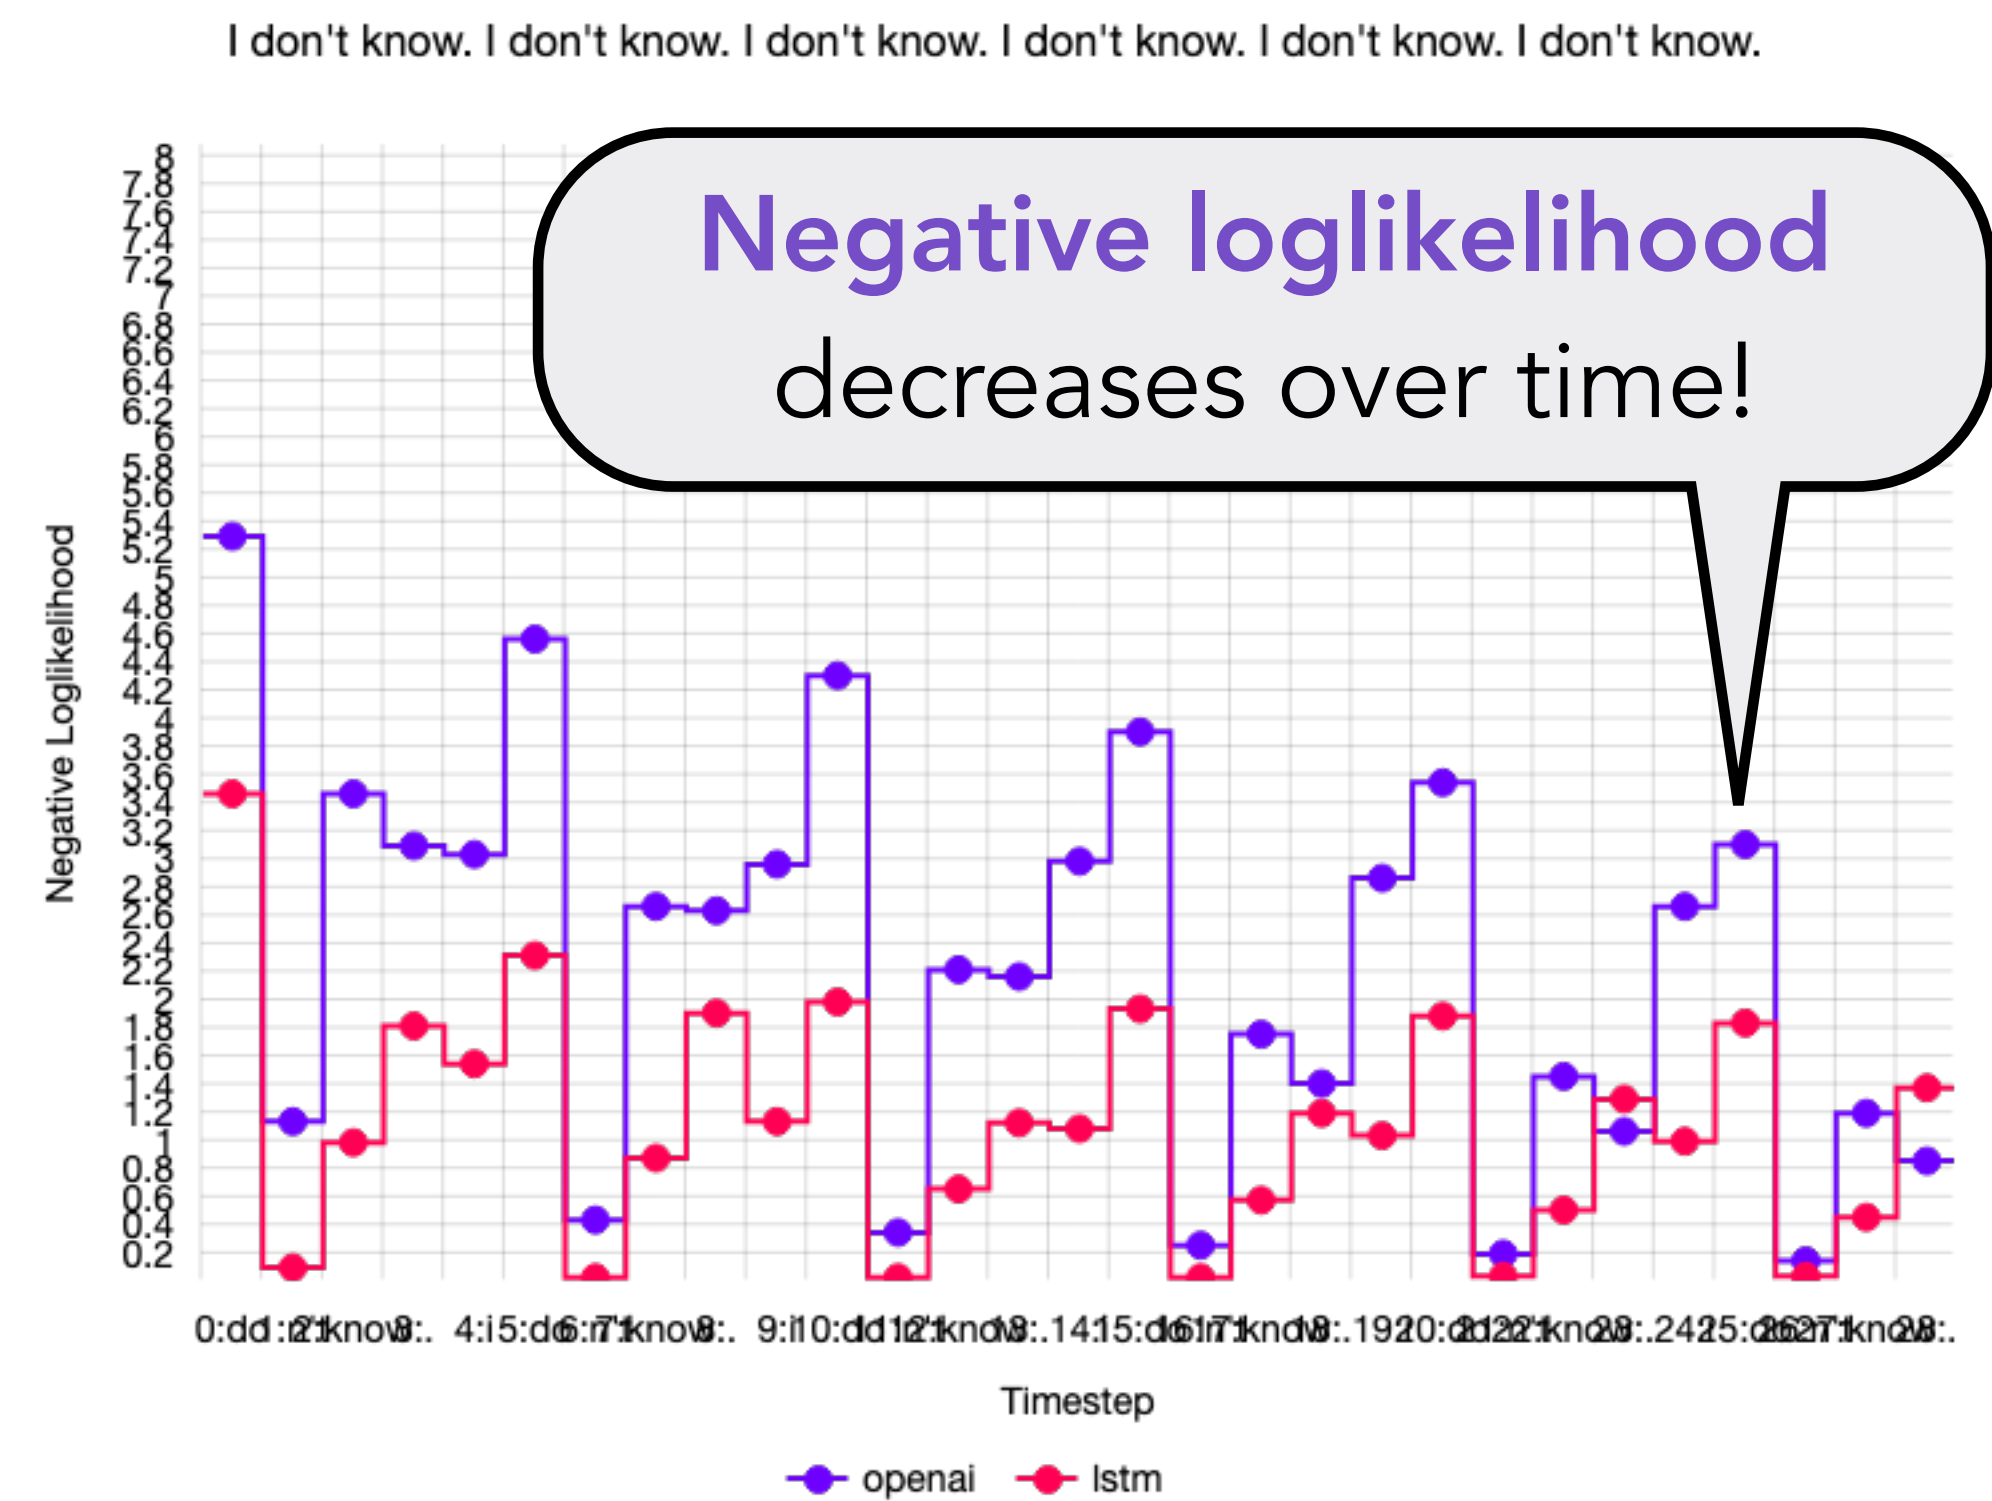

# Why does this happen?



# Artifact of Maximum Likelihood Training

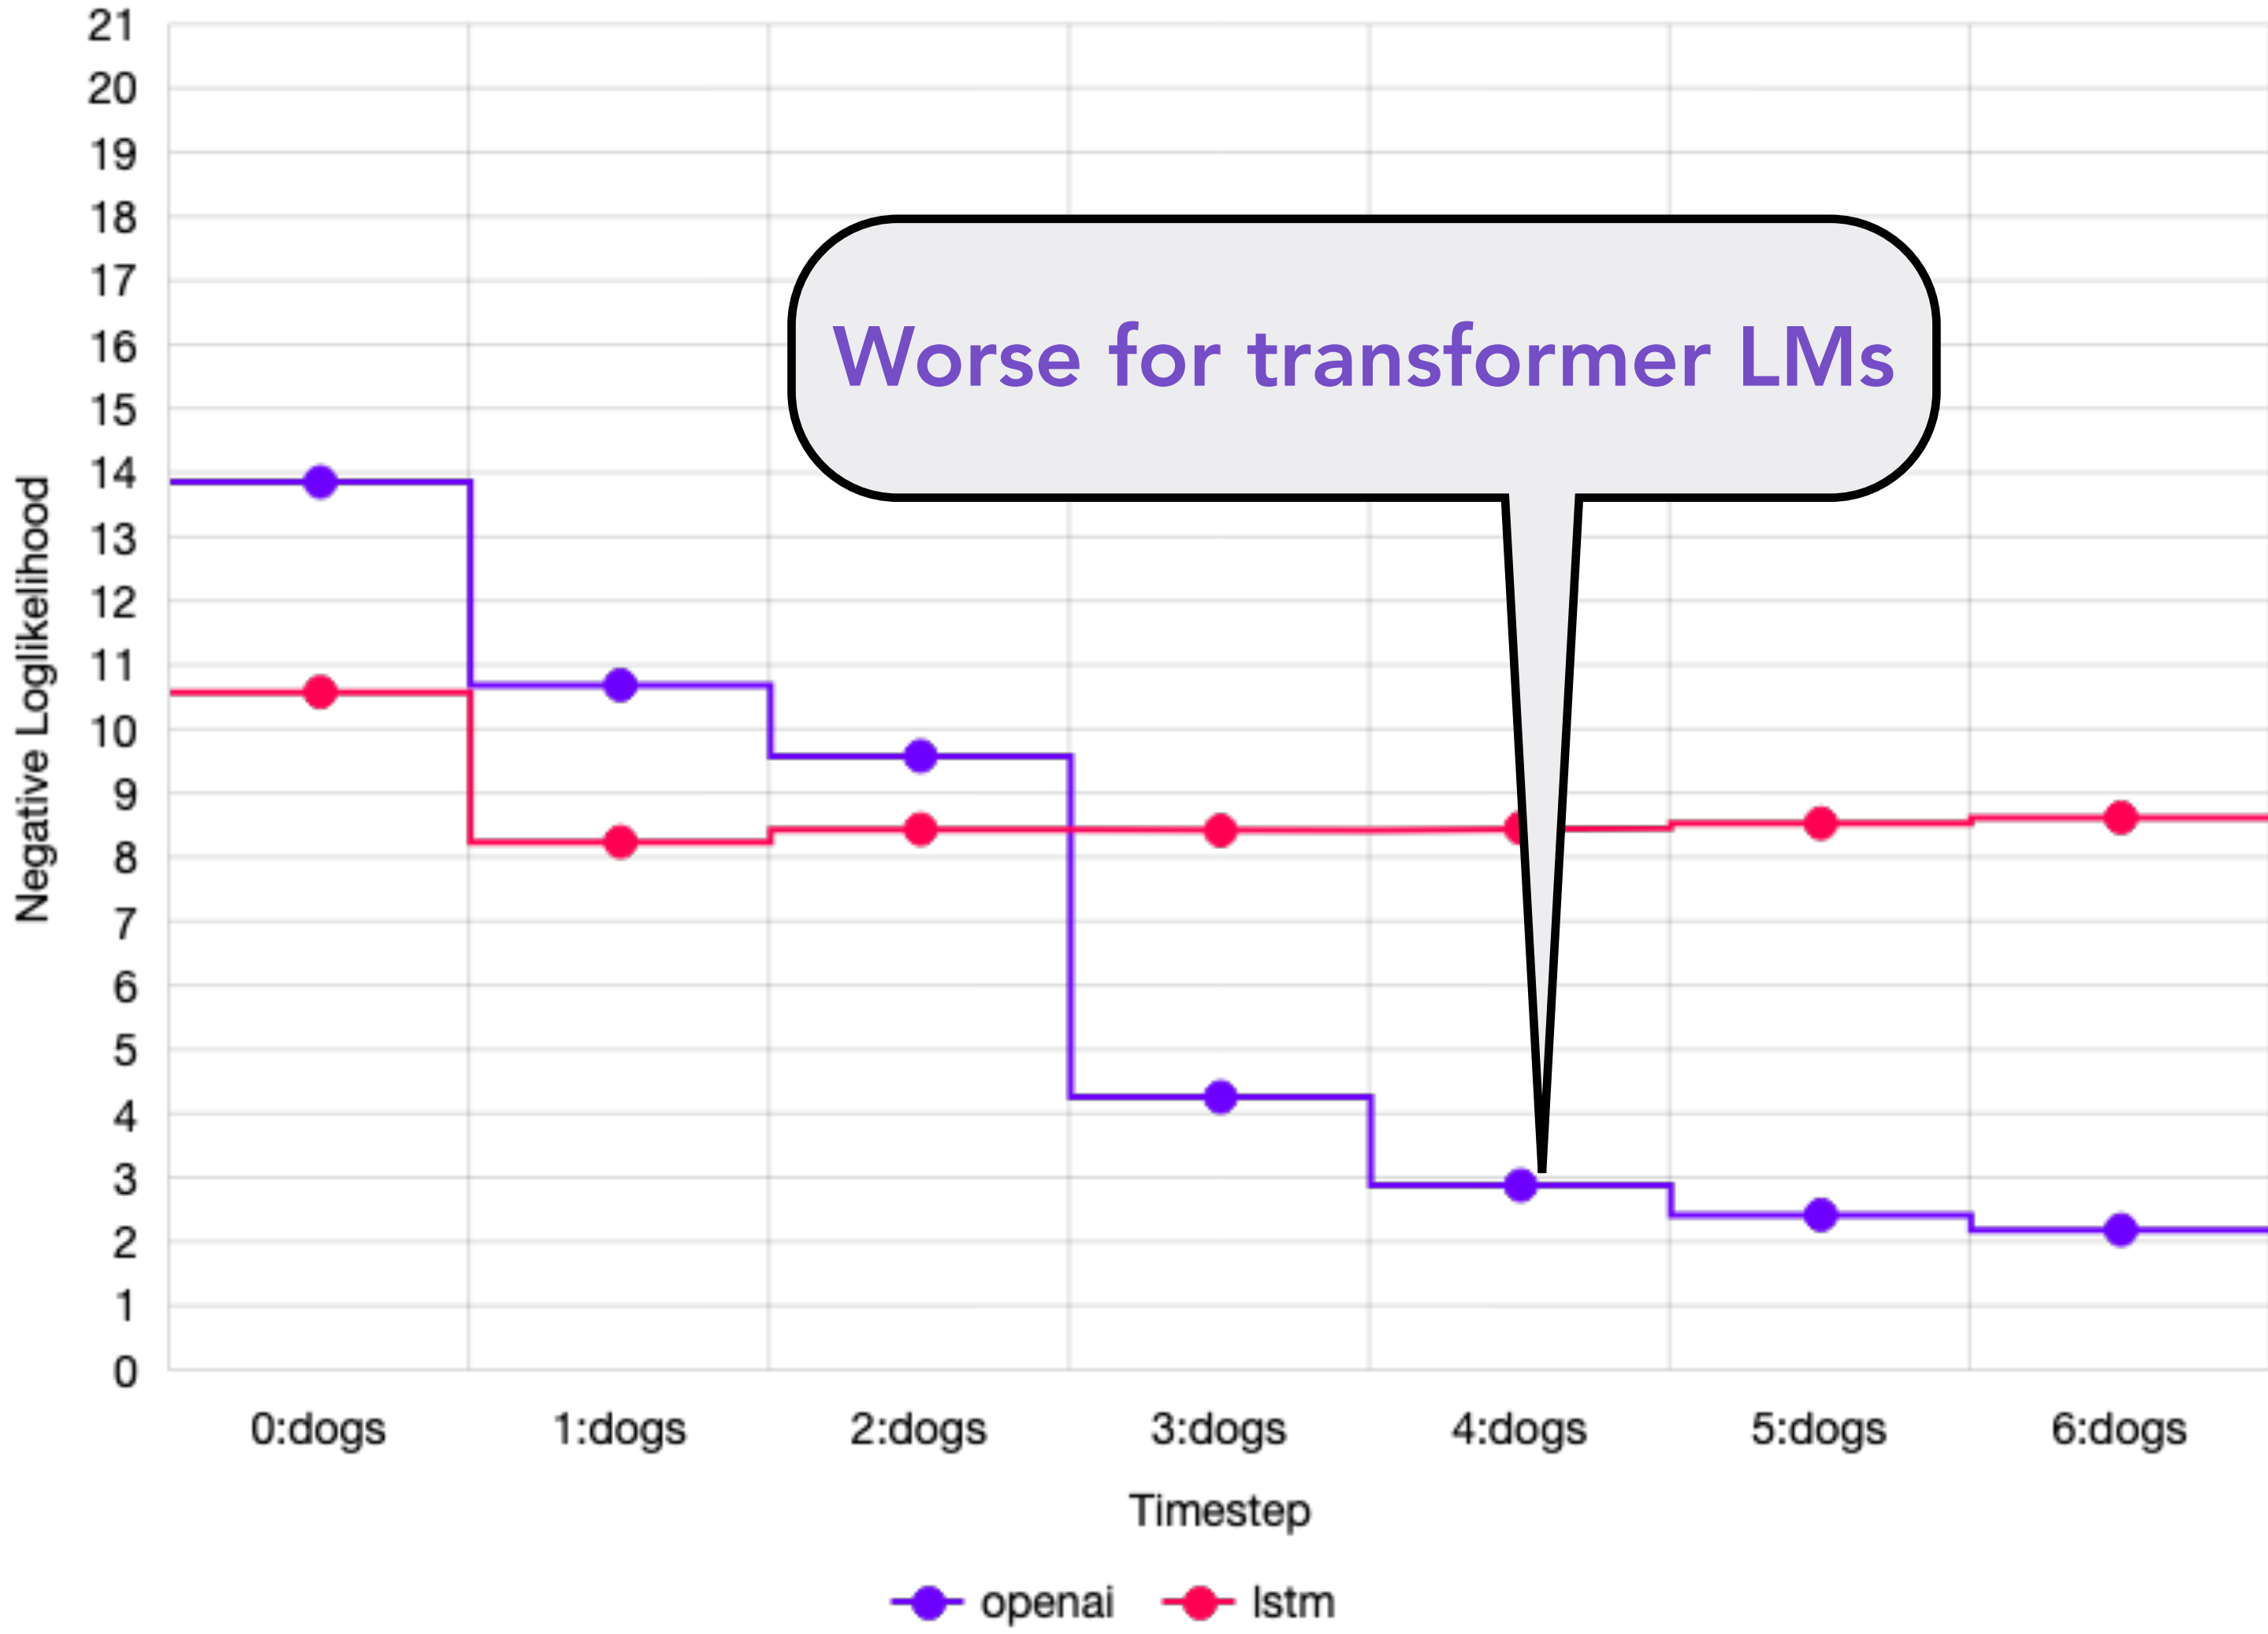


# Artifact of Maximum Likelihood Training



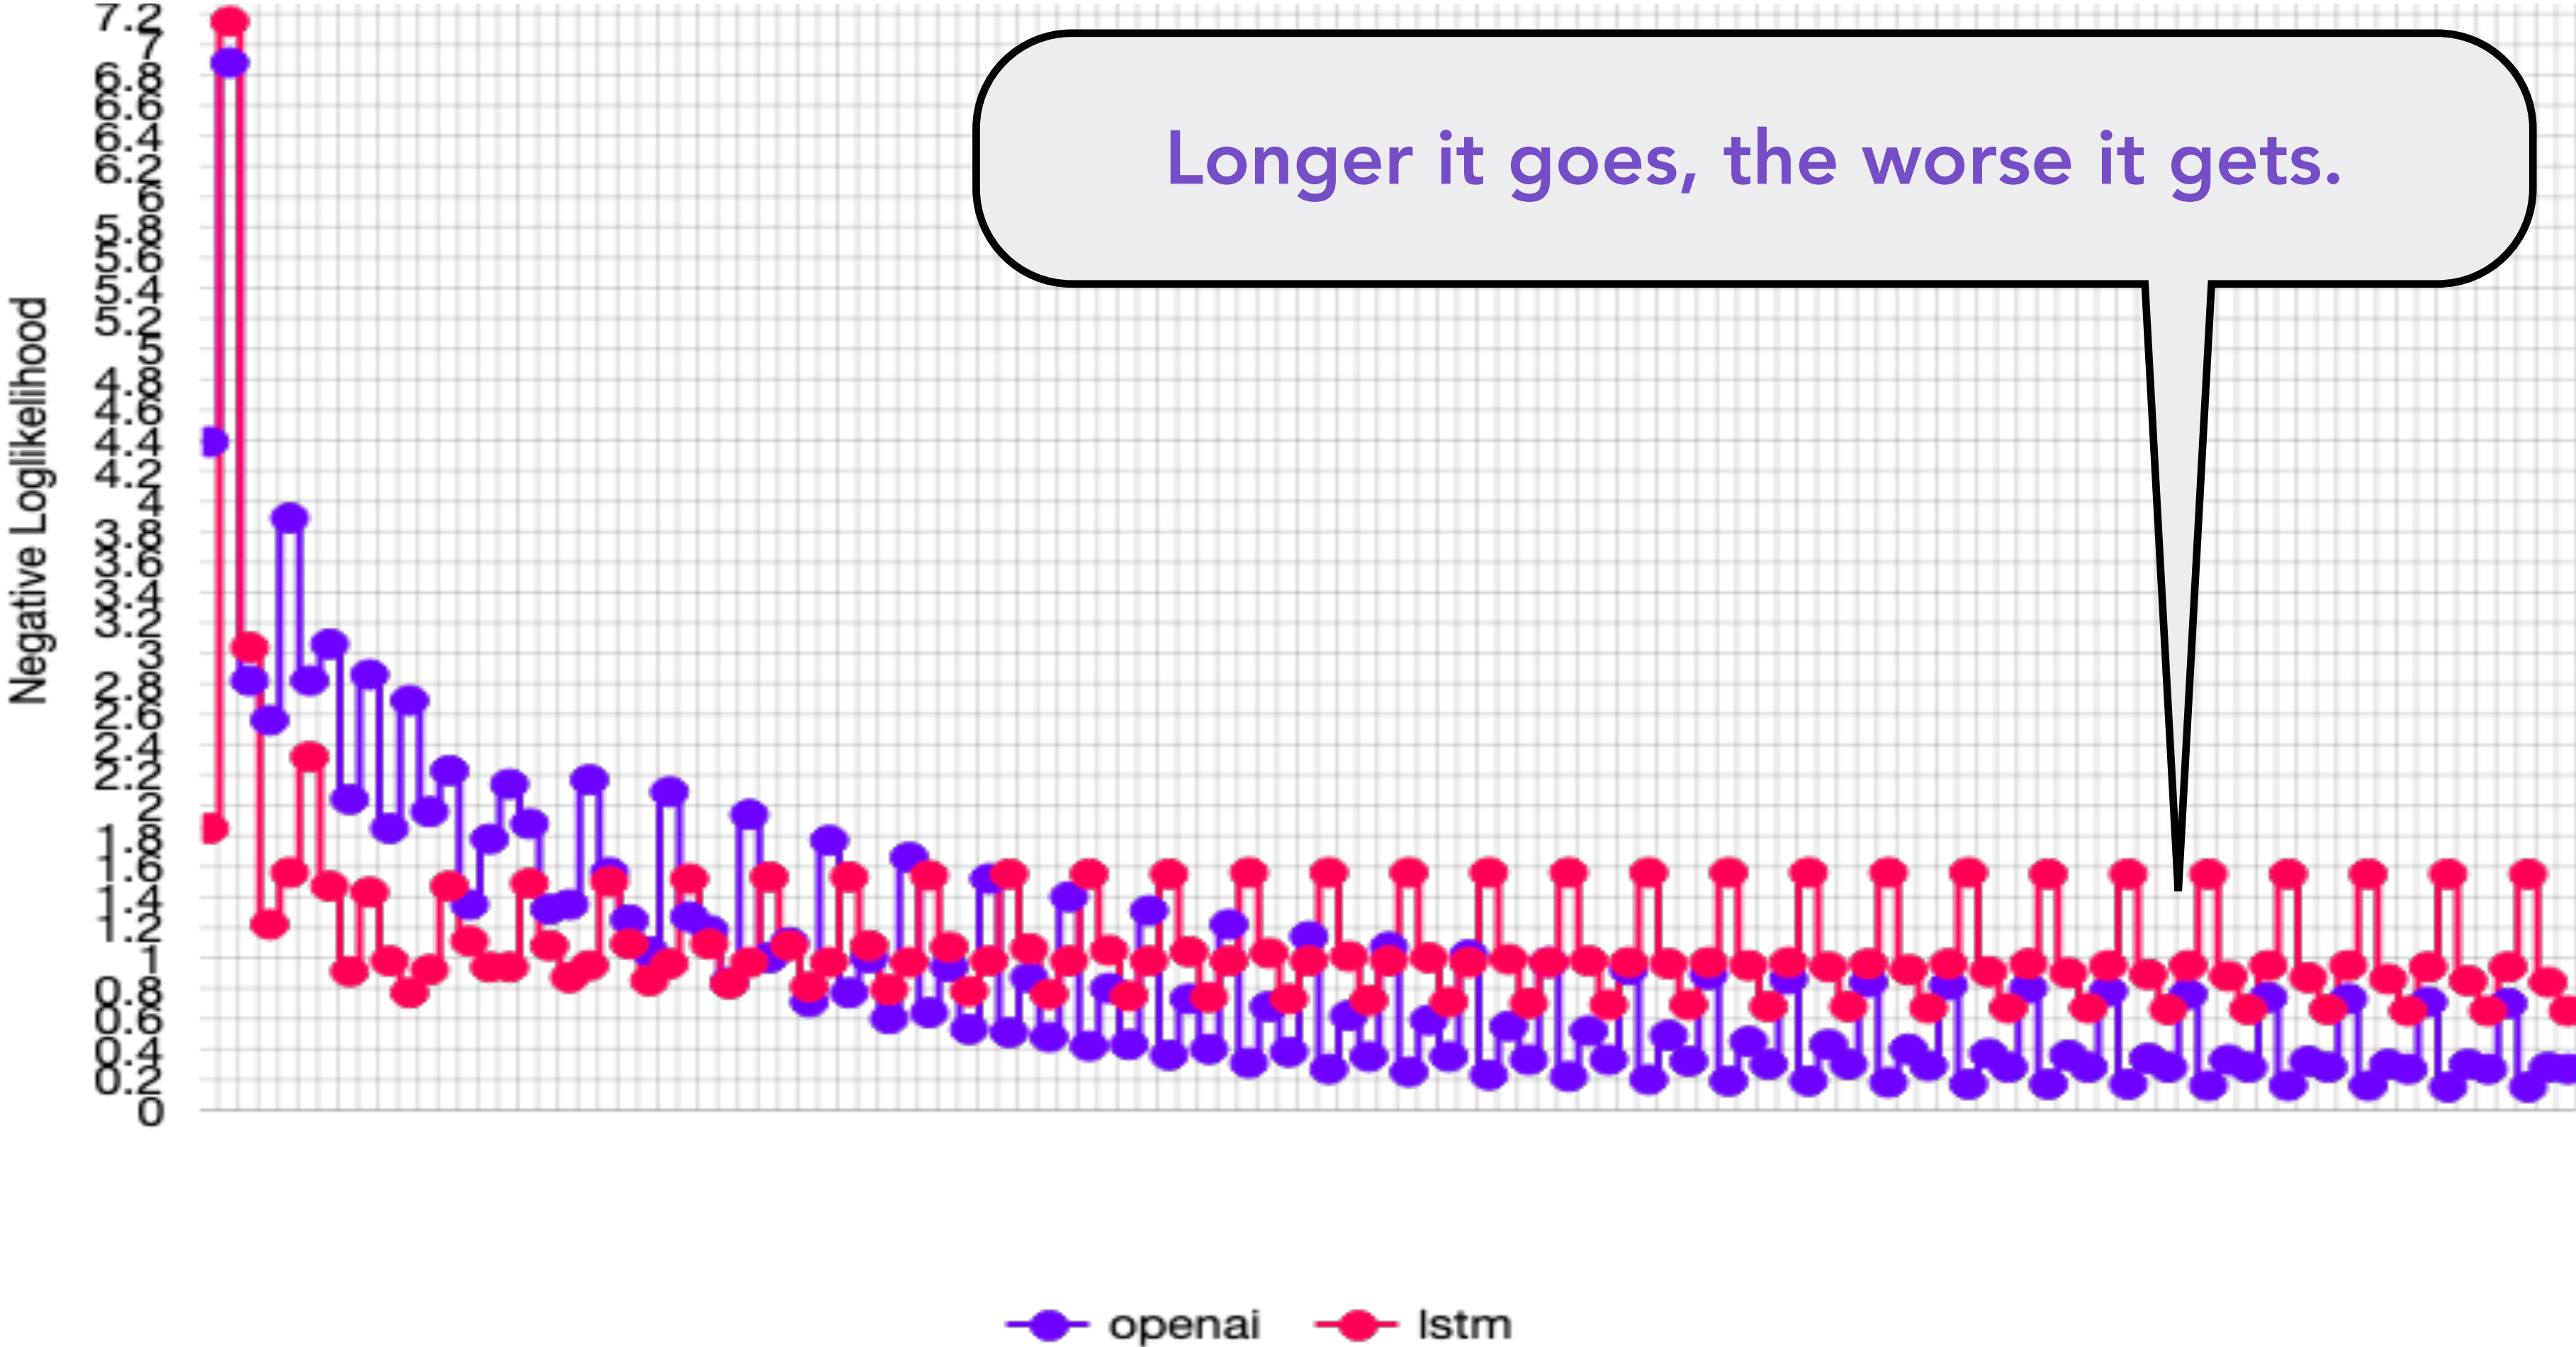
# Beam search gets repetitive and repetitive

dogs dogs dogs dogs dogs dogs dogs dogs



# Beam search gets repetitive and repetitive

I'm tired. I'm tired. I'm tired. I'm tired. I'm tired. I'm tired. I'm tired. I'm tired. I'm tired. I'm tired.




# How can we reduce repetition?

- Don't repeat  $n$ -grams (Hacky, but works!)
- Minimize additional loss term for minimizing hidden state similarity (LSTMs)

$$\hat{y}_t = g\left(\log P(y_t | \{y\}_{<t}) - s(h_t, h_{t-m})\right)$$

# Step-by-step Maximization



# Time to get *random*: Sampling

- $g$  = sample a token from the distribution of tokens

$$\hat{y}_t \sim P(y_t = w \mid \{y\}_{<t})$$





# Whoa, too *random*: Temperature Scaling

- Too much randomness: distribution has too much entropy

# Whoa, too *random*: Temperature Scaling

- Too much randomness: distribution has too much entropy
- **Solution:** Make the distribution more “peaky” with temperature scaling

Recall:  $P(y_t | \{y\}_{<t}) = \frac{e^{o_n}}{\sum_{m=1}^M e^{o_m}}$

# Whoa, too *random*: Temperature Scaling

- Too much randomness: distribution has too much entropy
- **Solution:** Make the distribution more “peaky” with temperature scaling

Recall:  $P(y_t | \{y\}_{<t}) = \frac{e^{o_n}}{\sum_{m=1}^M e^{o_m}}$

➔  $P(y_t | \{y\}_{<t}) = \frac{e^{o_n / \tau}}{\sum_{m=1}^M e^{o_m / \tau}}$

# Whoa, too *random*: Temperature Scaling

- Too much randomness: distribution has too much entropy
- **Solution:** Make the distribution more “peaky” with temperature scaling

$$P(y_t | \{y\}_{<t}) = \frac{e^{o_n / \tau}}{\sum_{m=1}^M e^{o_m / \tau}}$$

$$\tau > 1$$

“flatter”  
distribution

$$\tau < 1$$

“peakier”  
distribution

# Maybe we need fewer options: Top-*k* sampling

- The entire distribution over tokens is not needed at every step
- Many token choices should have no chance of being selected

# Maybe we need fewer options: Top- $k$ sampling


- The entire distribution over tokens is not needed at every step
- Many token choices should have no chance of being selected
- Only sample from the top  $k$  tokens in the distribution

$$\hat{y}_t \sim P^* (y_t = w \mid \{y\}_{<t})$$

He wanted to go to the




Decoder




Randomly sample token from top  $k$  highest probability tokens in  $P(\cdot)$

# Issues with top- $k$ sampling

Top- $k$  can cut-off too *quickly*




Top- $k$  can cut-off too *slowly*



# I don't know how many options I need: Top- $p$ sampling

- Also known as **nucleus** sampling
- Sample from subset of vocabulary where probability mass is concentrated


$$P_1(y_t | \{y\}_{<t})$$



$$P_2(y_t | \{y\}_{<t})$$




$$P_3(y_t | \{y\}_{<t})$$





# I don't know how many options I need: Top- $p$ sampling


- Also known as **nucleus** sampling:
- Sample from subset of vocabulary where probability mass is concentrated
- Probability mass has a dynamically changing **nucleus**




# I don't know how many options I need: Top- $p$ sampling

- Also known as **nucleus** sampling:
- Sample from subset of vocabulary where probability mass is concentrated
- Probability mass has a dynamically changing **nucleus**

Only sample from the **nucleus** of the **distribution**



$$P_1^*(y_t | \{y\}_{<t})$$



$$P_2^*(y_t | \{y\}_{<t})$$



$$P_3^*(y_t | \{y\}_{<t})$$

# This all sounds a bit *risky*

- What if my sequence just isn't very good?

# Optimize other sequence-level scores: Re-ranking

- What if my sequence just isn't very good?
- Sample a bunch of sequences
- Define a score to approximate the quality of your sequence.
- Simplest is to just use perplexity!

# Optimize other sequence-level scores: Re-ranking

- What if my sequence just isn't very good?
- Sample a bunch of sequences
- Define a score to approximate the quality of your sequence.
- Simplest is to just use perplexity!
- However, re-rankers can be used to score a variety of properties: style (Holtzman et al., 2018), discourse (Gabriel et al., 2019), entailment/factuality (Goyal et al., 2020), logical consistency (Lu et al., 2020), and many more...

# Optimize other sequence-level scores: Re-ranking

- What if my sequence just isn't very good?

- Sample a bunch of sequences

- Define a score

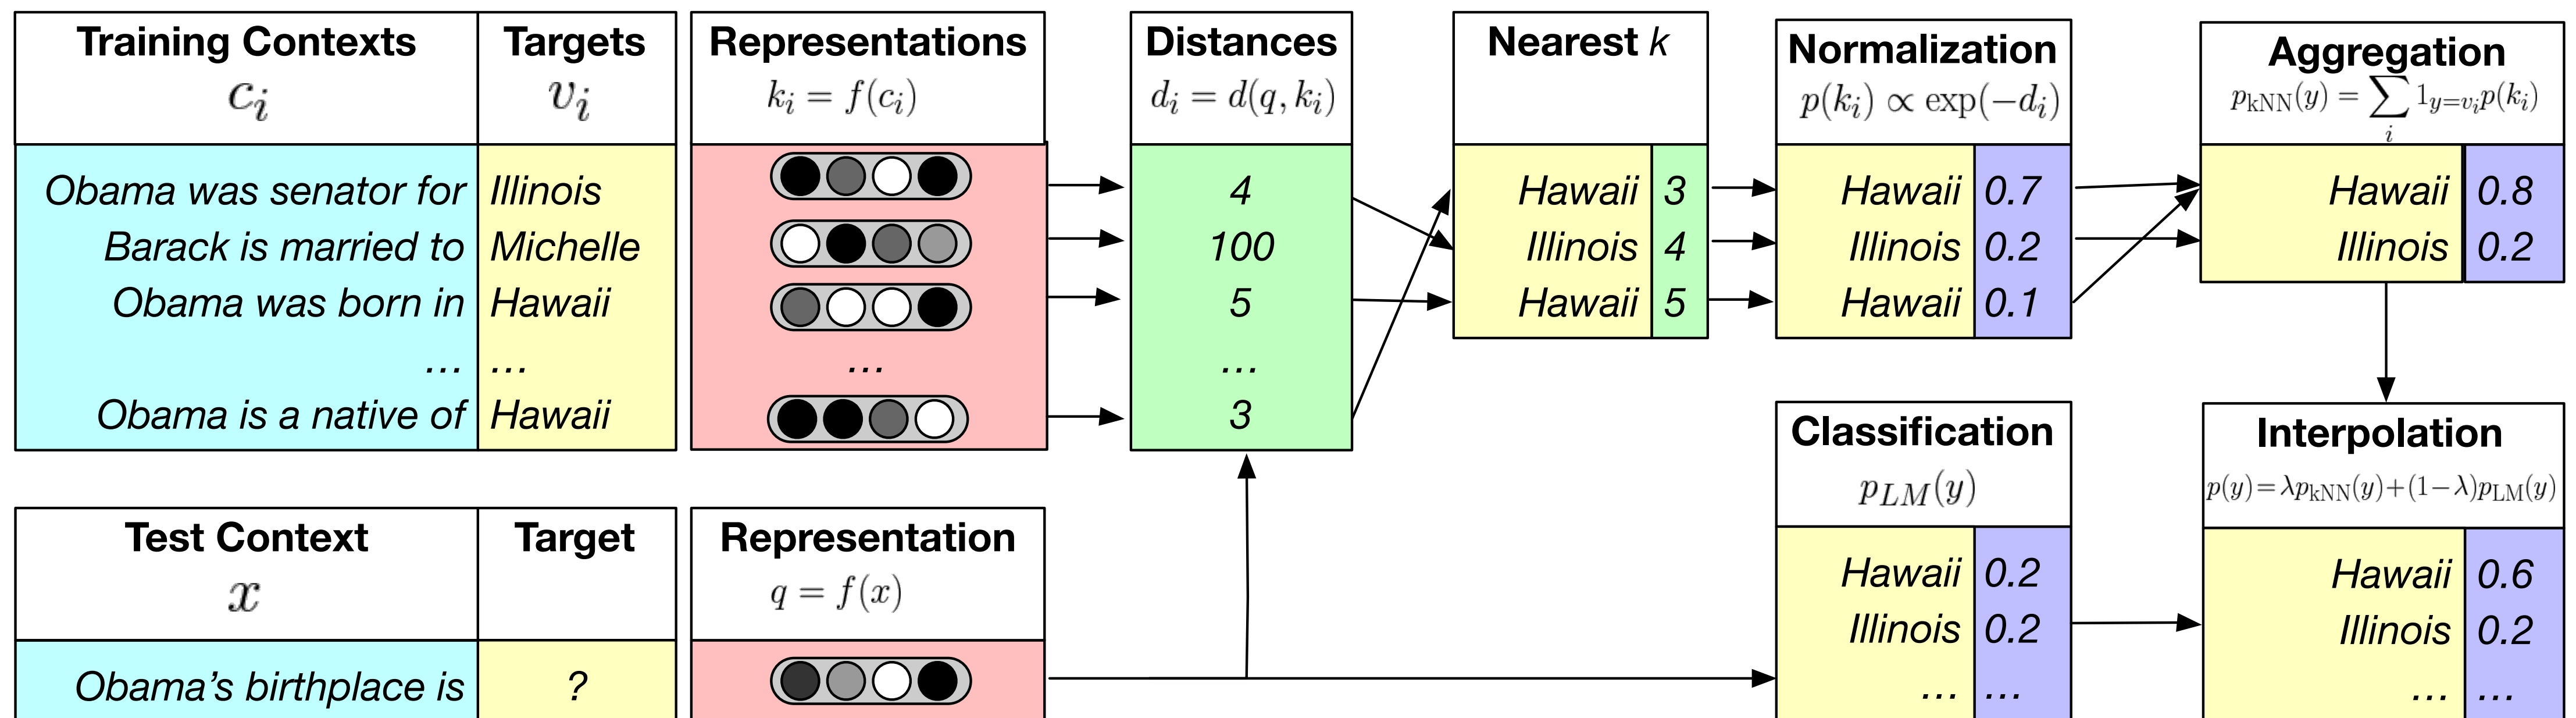
- Simplest is

Change your distribution  
at inference time!

- However, re-rankers can be used to score a variety of properties: style (Holtzman et al., 2018), discourse (Gabriel et al., 2019), entailment/factuality (Goyal et al., 2020), logical consistency (Lu et al., 2020), and many more...

# kNN Language Models

- Don't just rely on your trained model to generate a distribution over tokens
- Use knowledge of similar contexts from another corpus

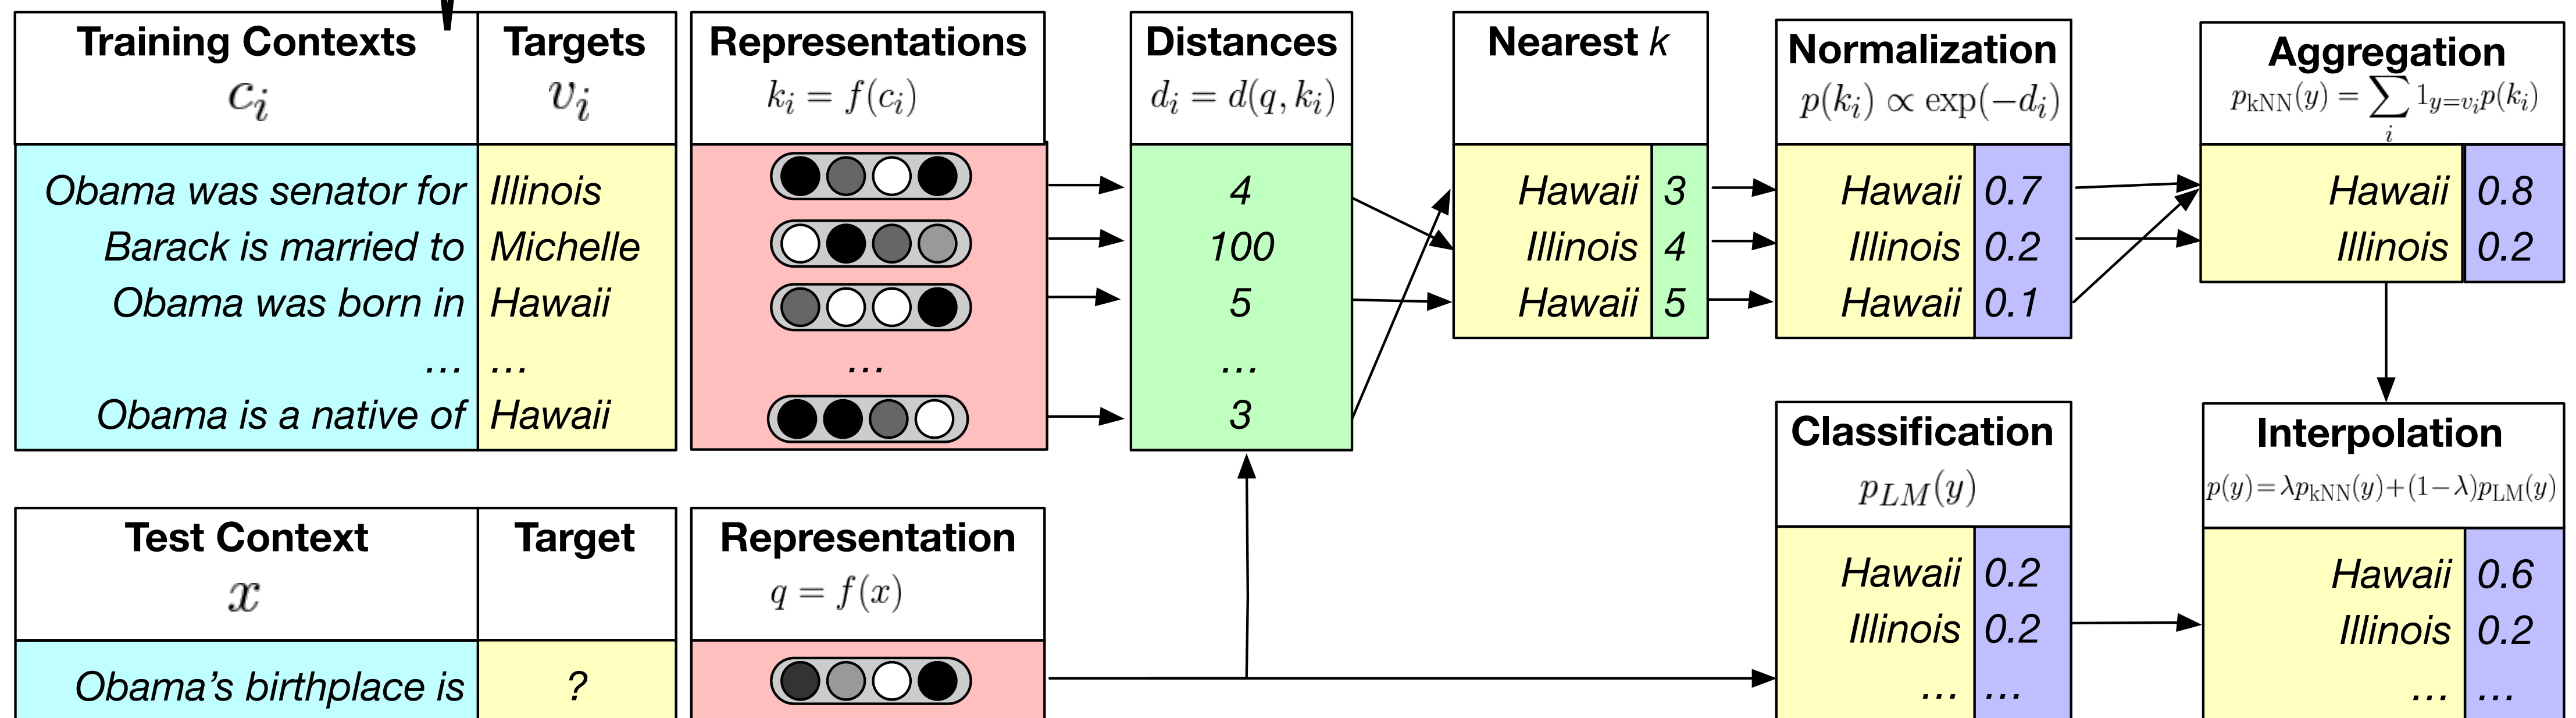


# kNN Language Models

• Don't just rely on your trained model to generate a distribution over

Initialize a **database** of contexts

- Use knowledge of similar contexts from another corpus

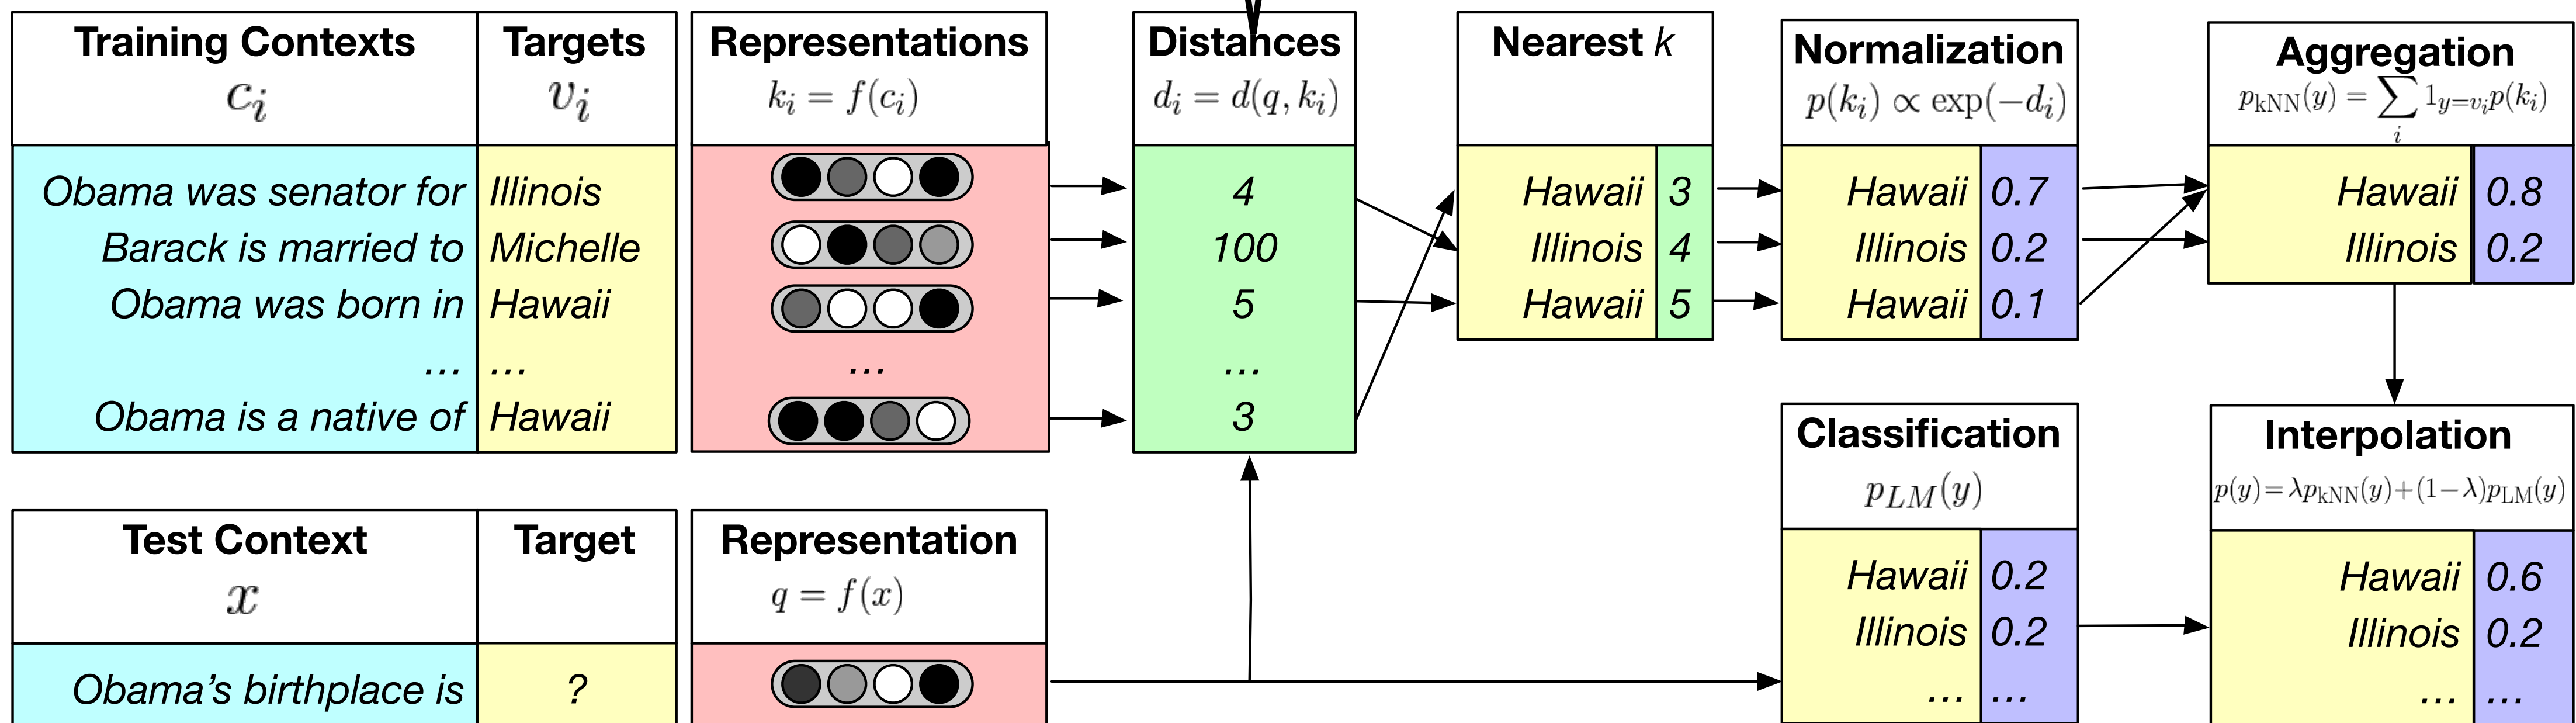




# kNN Language Models

- Don't just rely on your trained model to generate a distribution over to
 

Efficiently compute **distance** between each context in DB and current sequence
- Use knowledge of similar contexts from another corpus

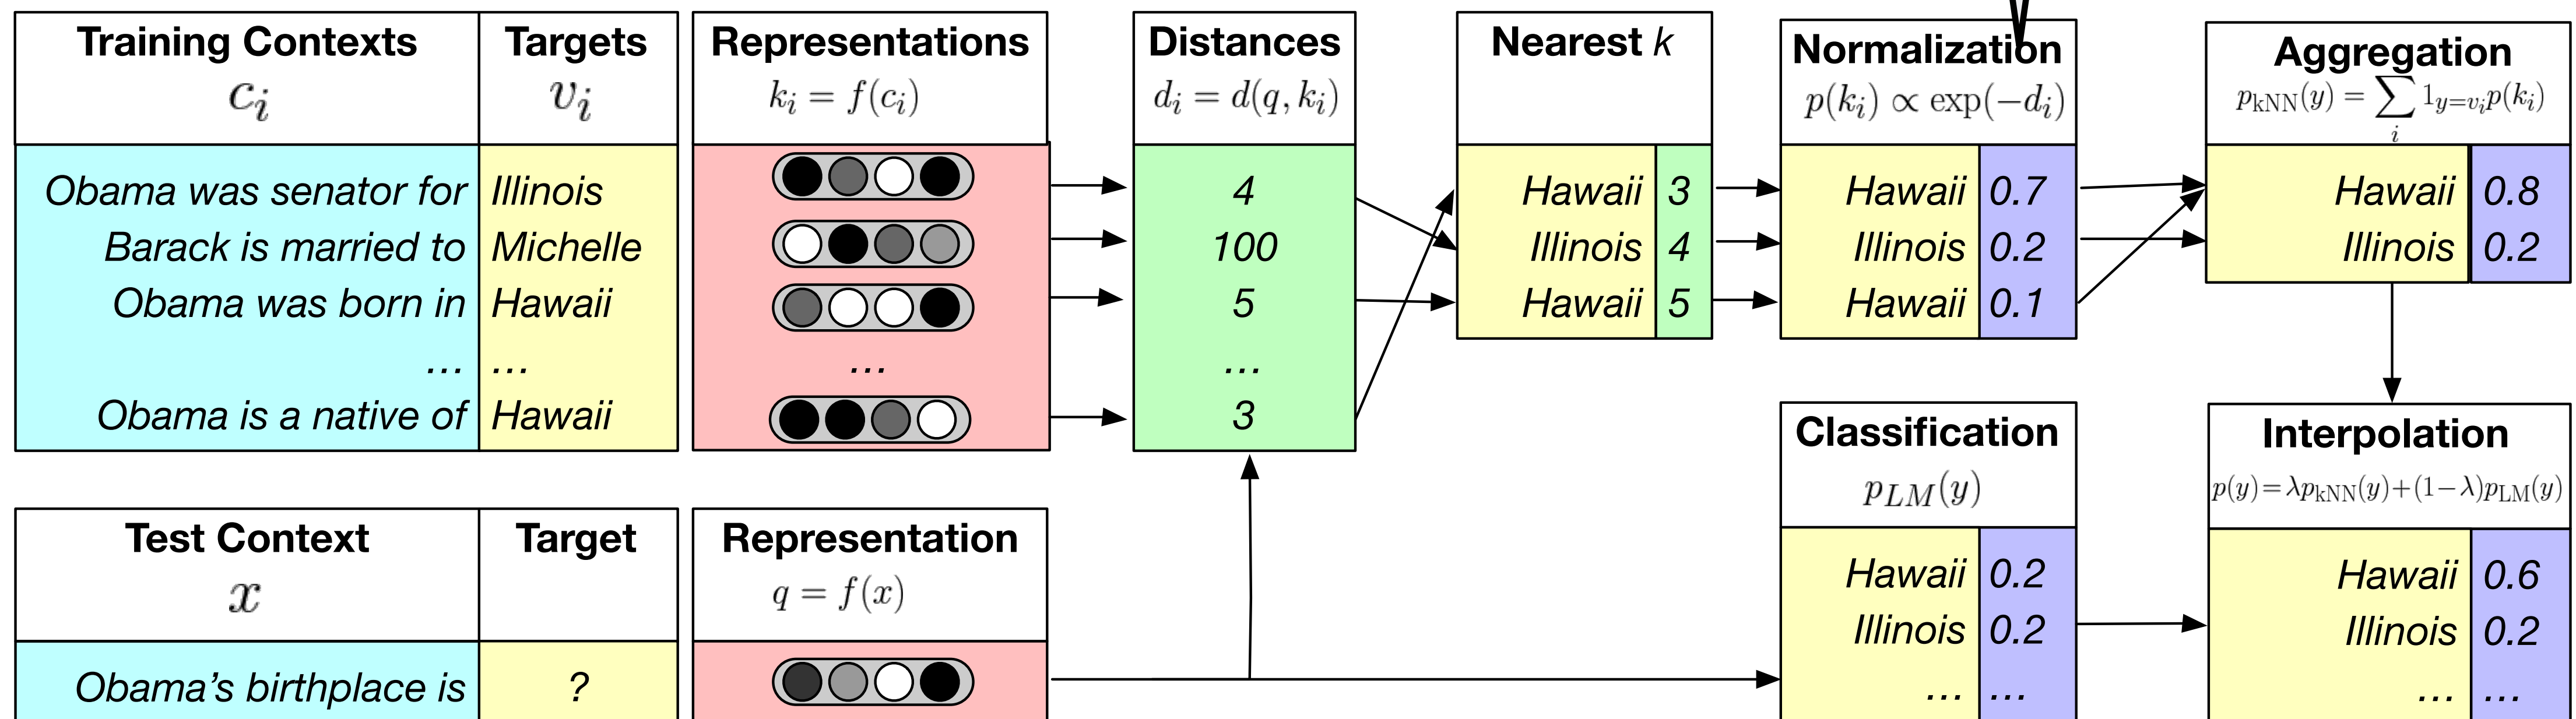


# kNN Language Models

- Don't just rely on your trained model to generate a distribution over tokens

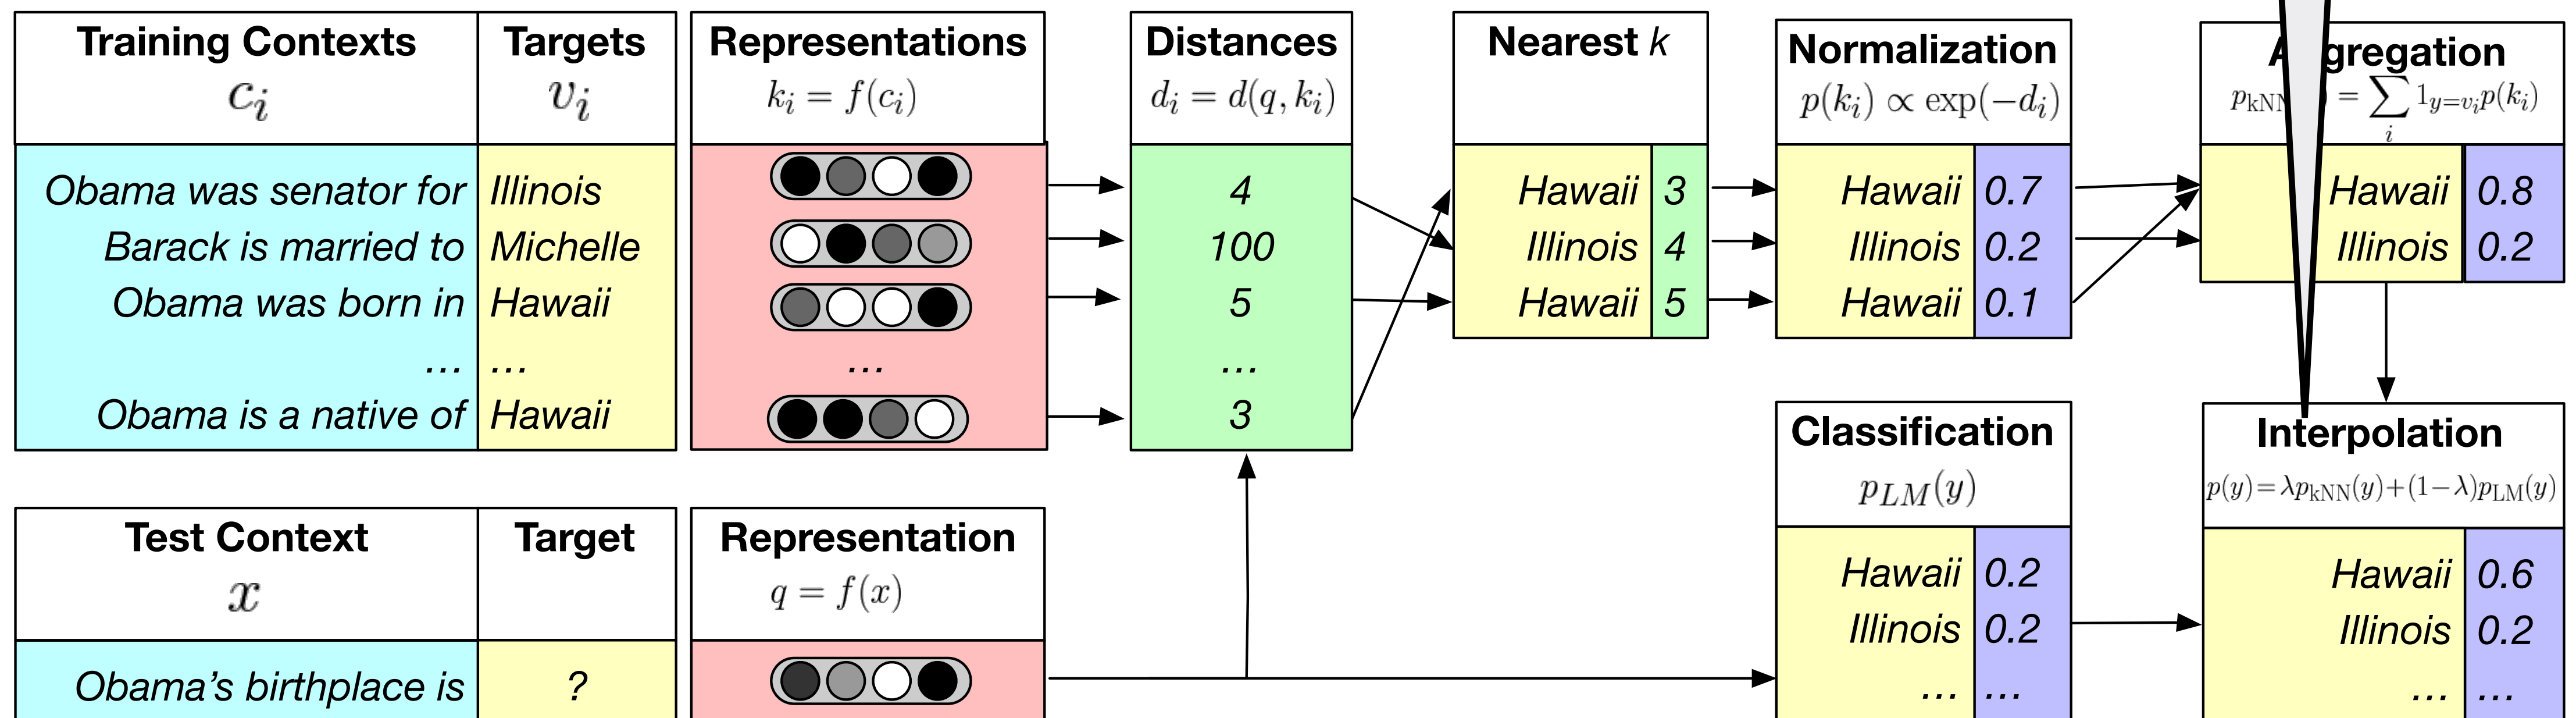
Compute **distribution** over possible **targets** from context DB sequences using **distance of history**

- Use knowledge of similar contexts from another corpus



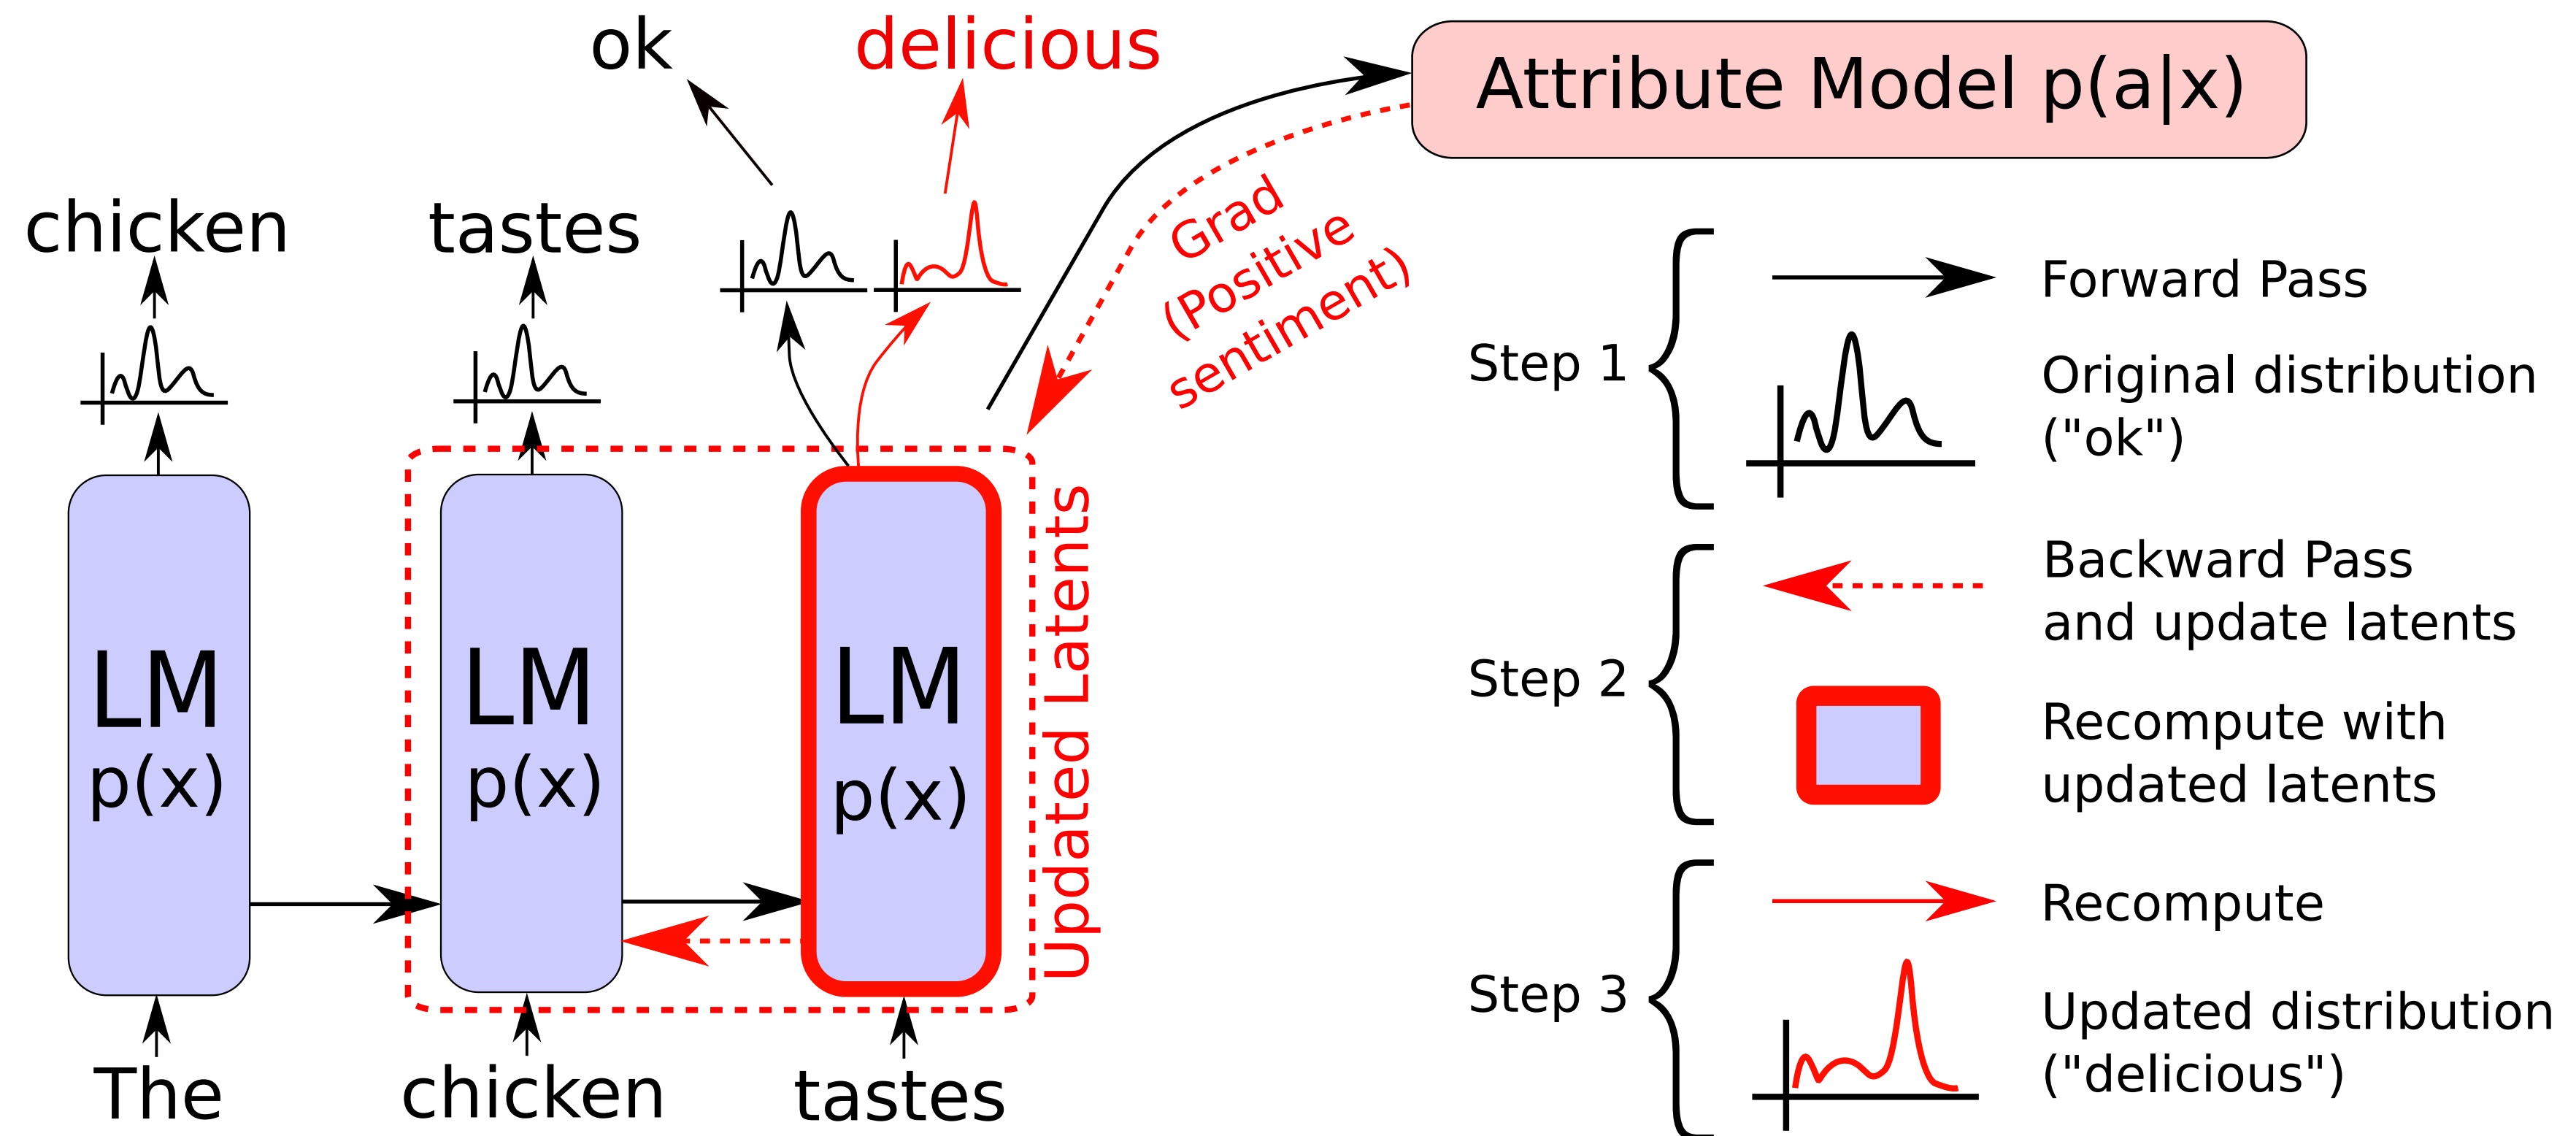
# kNN Language Models

- Don't just rely on your trained model to generate a distribution over tokens  
 Interpolate distribution from nearest neighbor search with model distribution
- Use knowledge of similar contexts from another corpus



# Plug and Play Language Models!


- What if I want to encourage a tough to formalize behavior at inference time?



# Plug and Play

Define an attribute model that scores the generated sequence. Each generated token must try to increase the score given to the sequence by the attribute model

- What if I want to encourage a specific attribute over time?

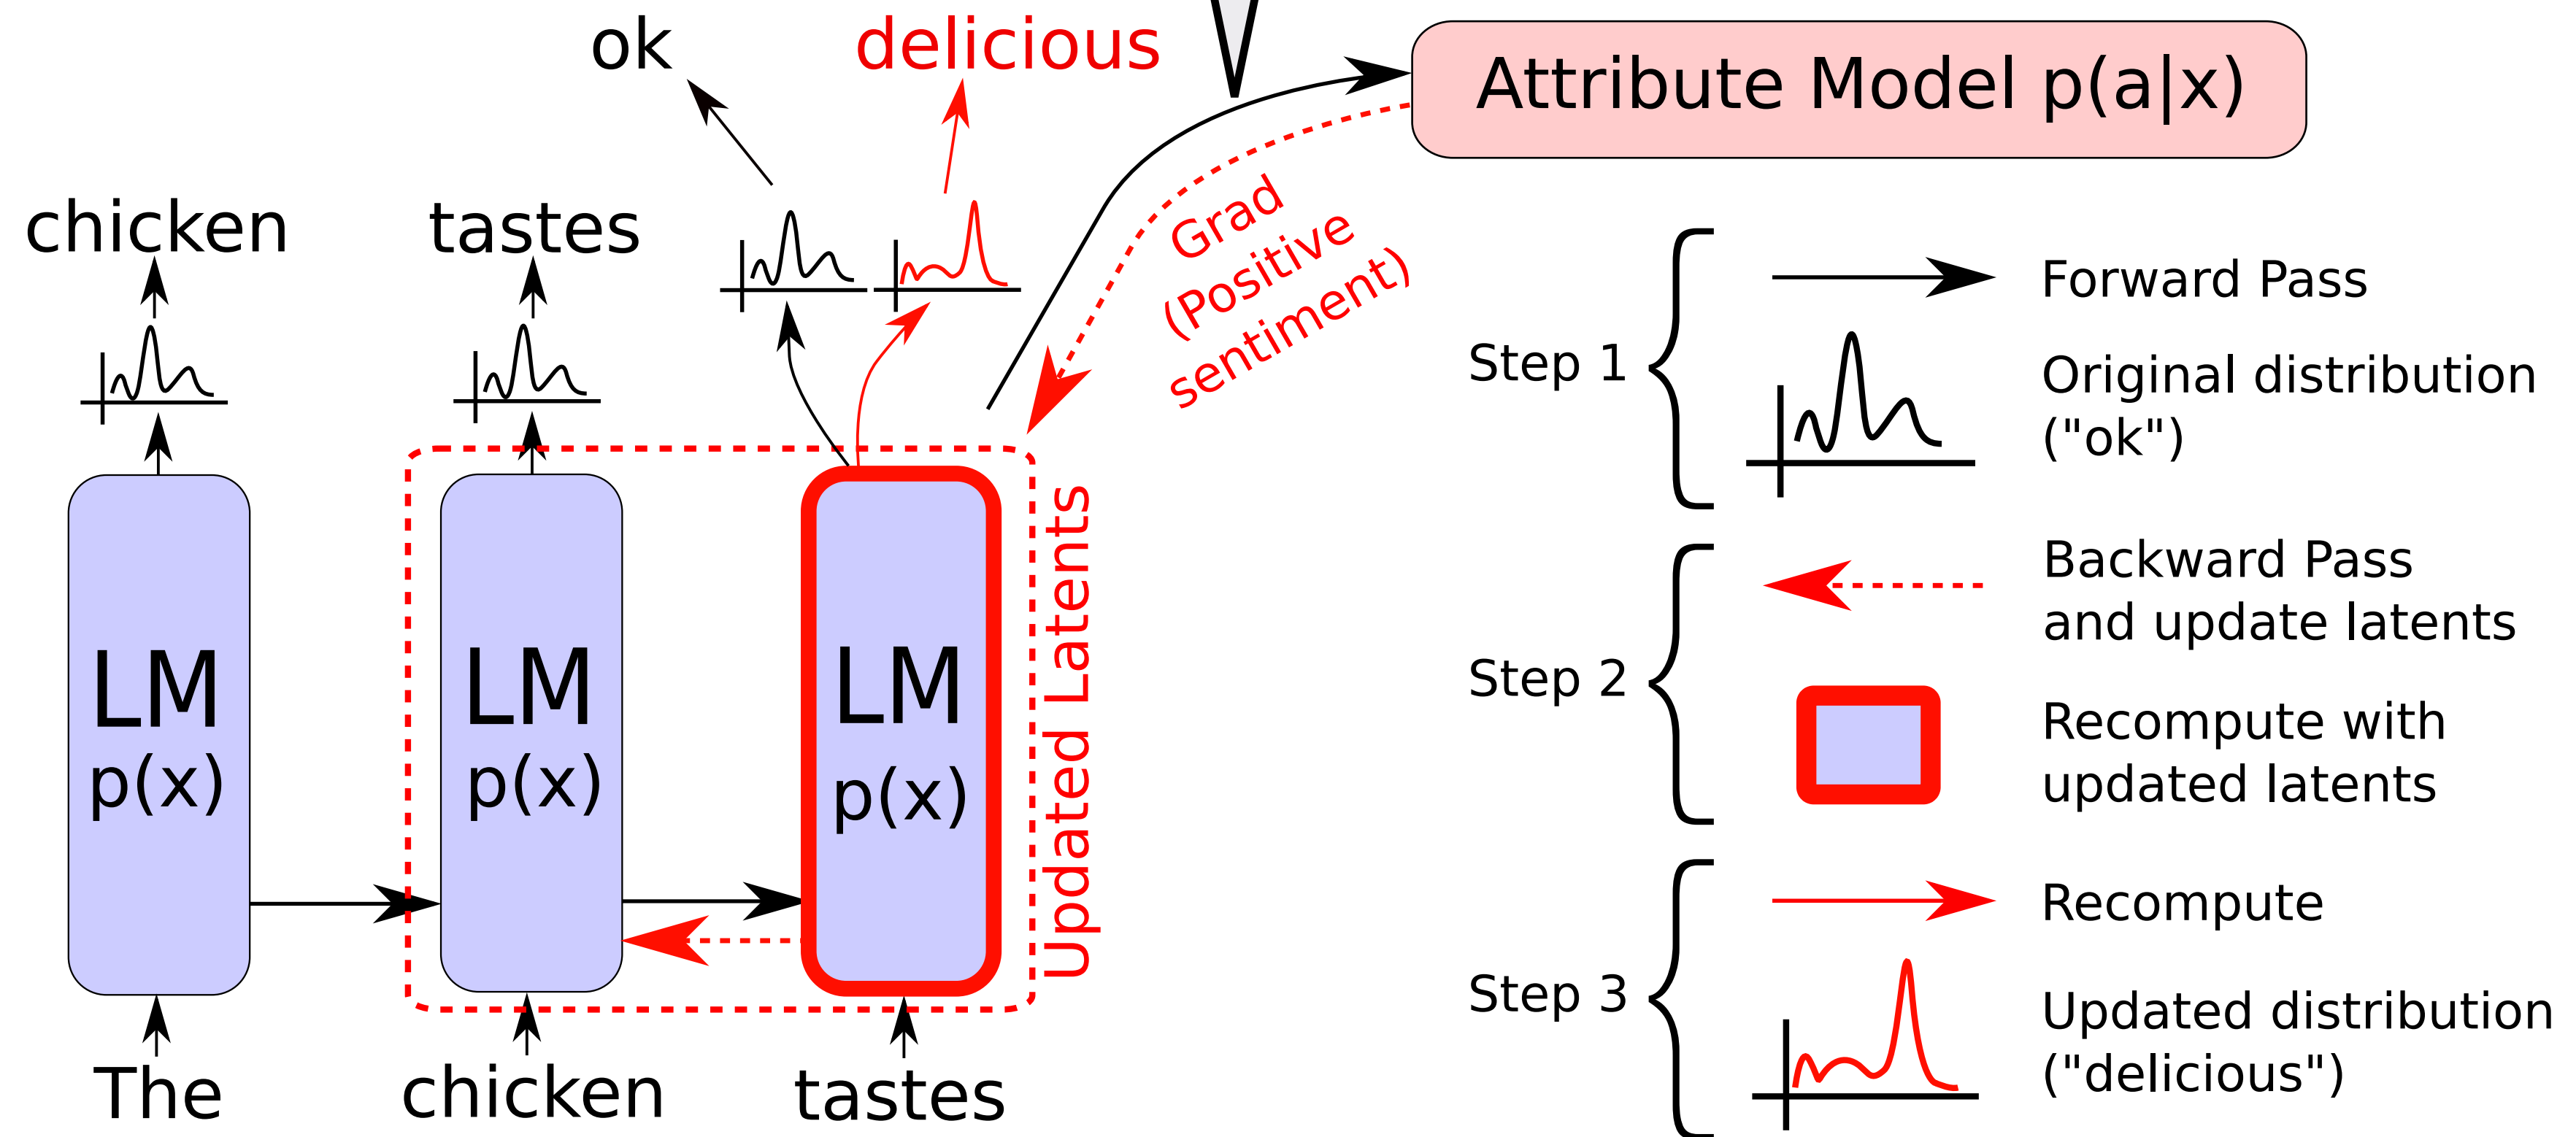


# Plug and Play Language Models!

- What if I want to formalize behavior at inference time?


Backpropagate loss with respect to attribute score

formalize behavior at inference time



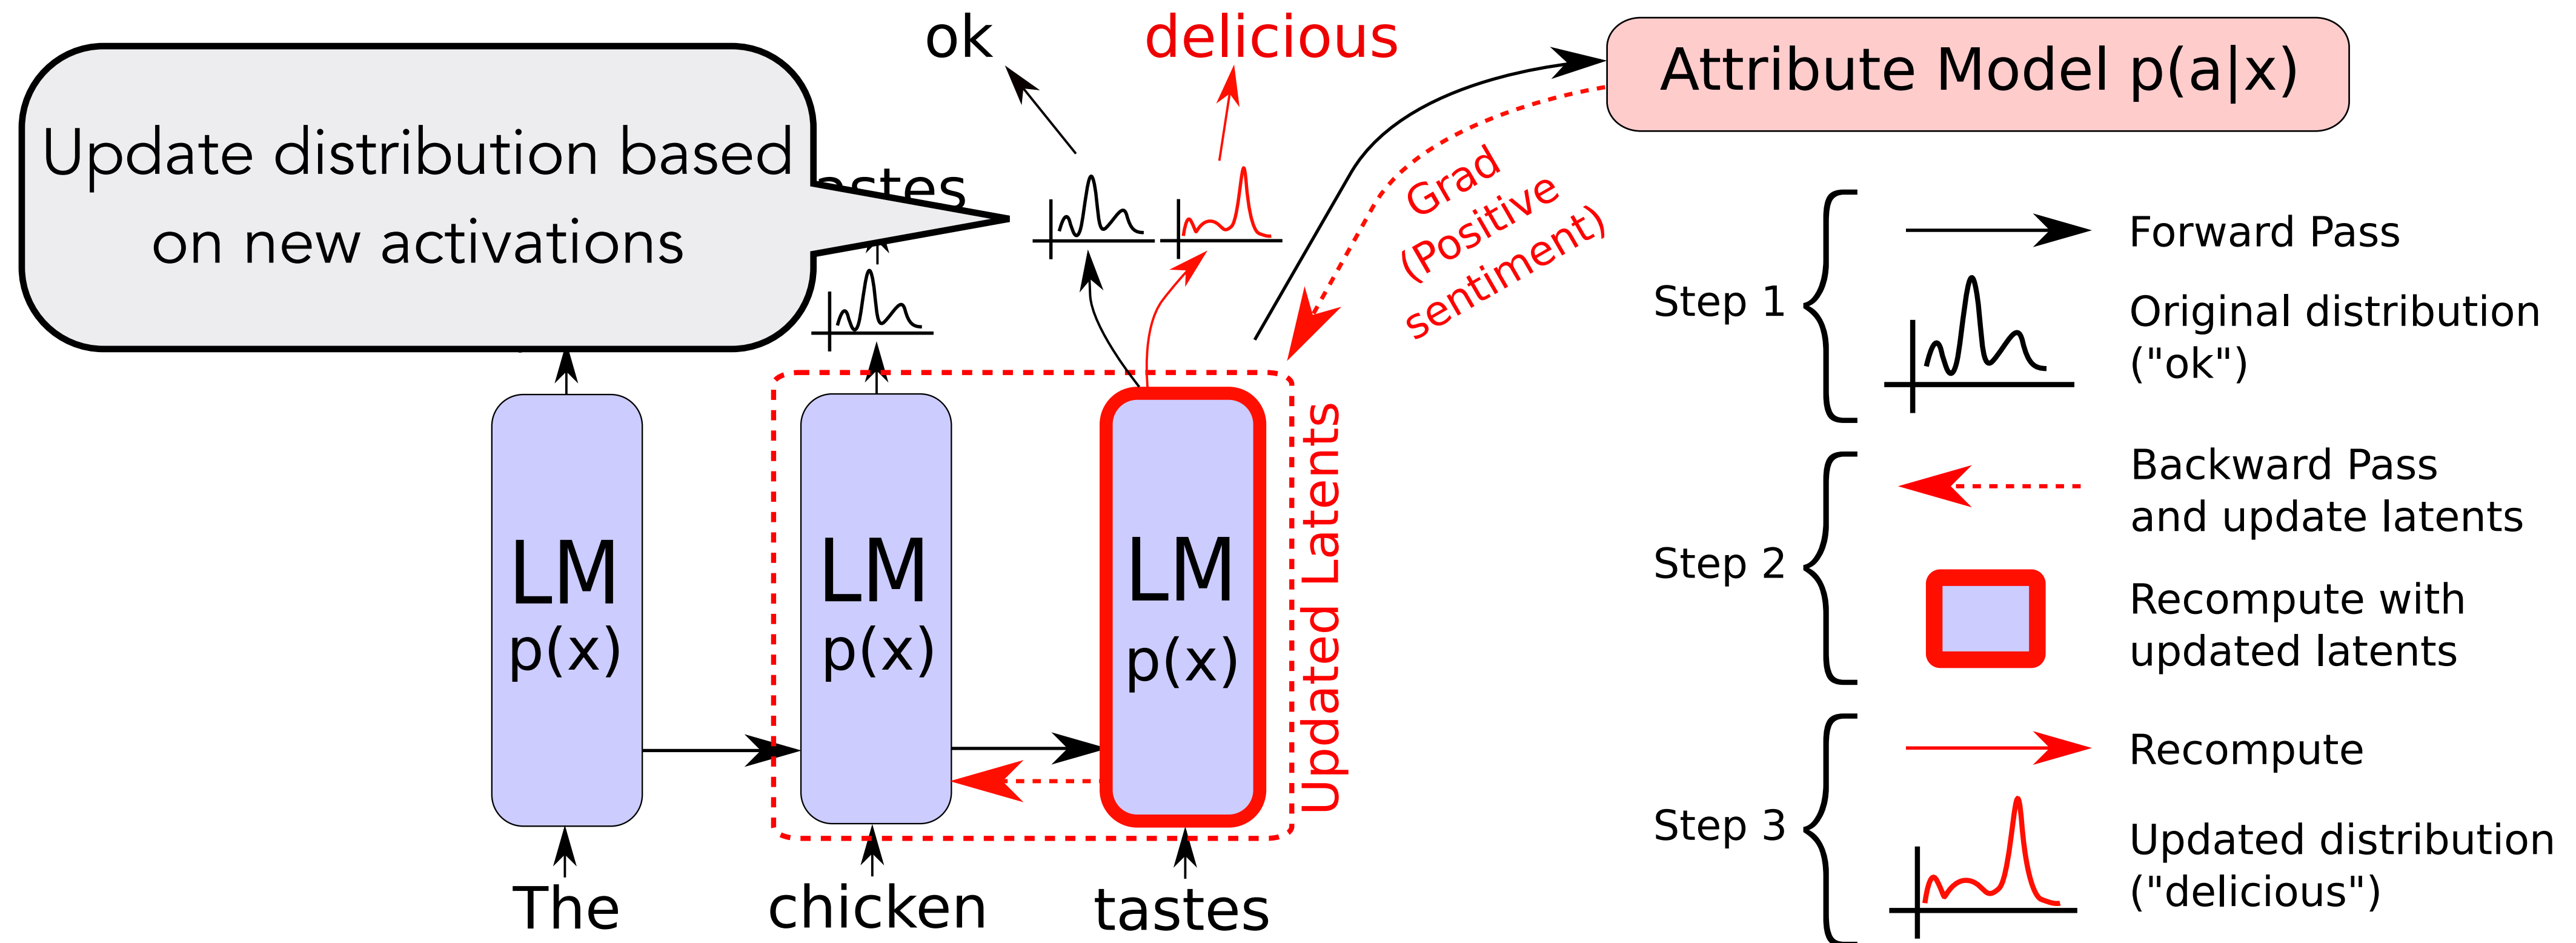
# Plug and Play Language Models!

- What if I want to encourage a tough to formalize behavior at inference time?



# Plug and Play Language Models!

- What if I want to encourage a tough to formalize behavior at inference time?







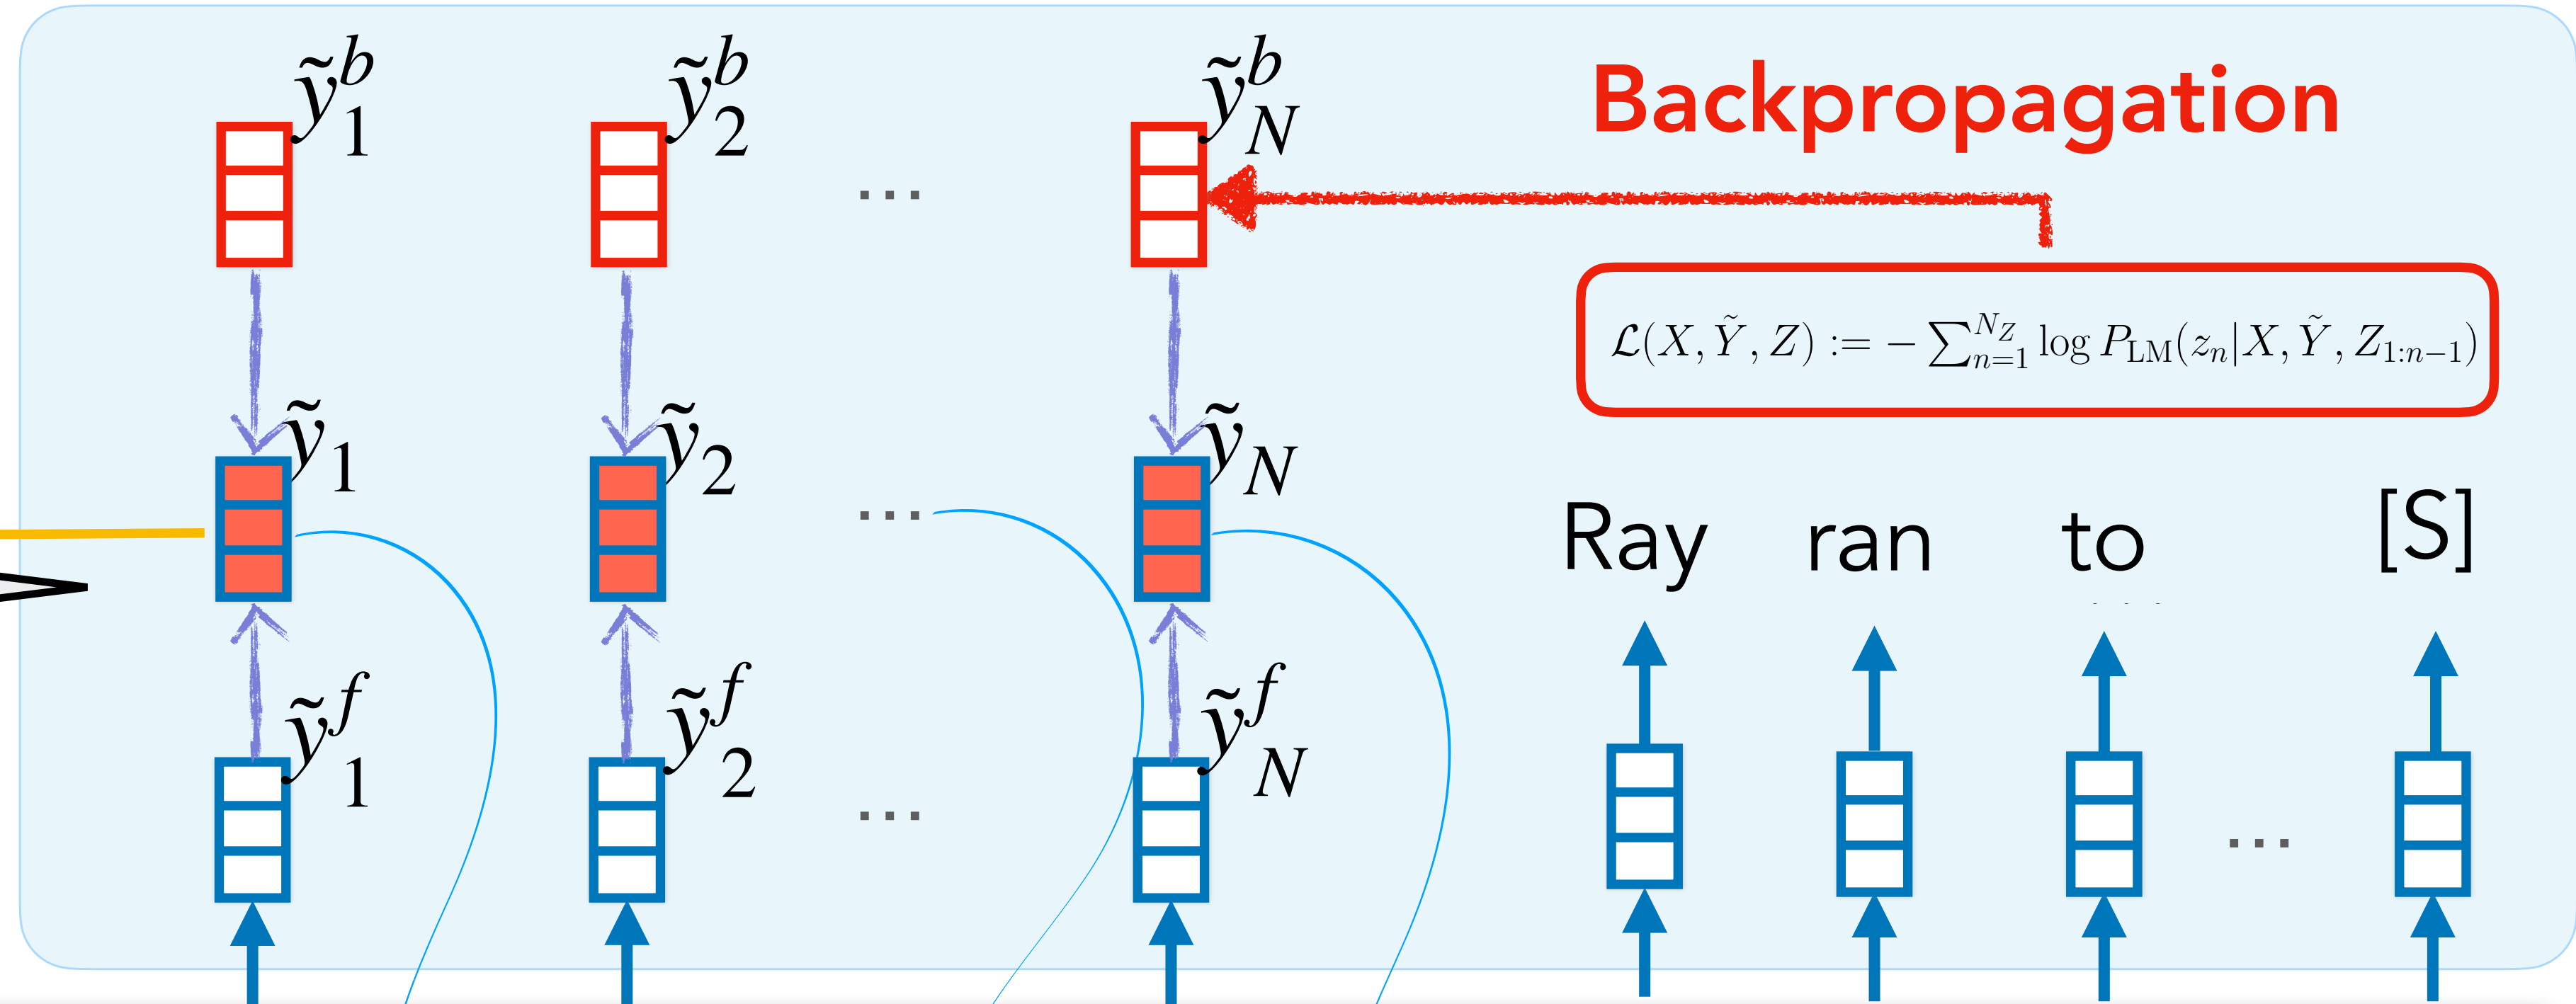
$Y$

Output: She hit the rope and the tire fell on top of her.

**Backpropagation**

$$\mathcal{L}(X, \tilde{Y}, Z) := -\sum_{n=1}^{N_Z} \log P_{LM}(z_n | X, \tilde{Y}, Z_{1:n-1})$$

Inference time optimization with respect to future constraints



Ray ran to [S]

$x_1$   $x_2$  ...  $x_{N_X}$

LM

$\tilde{y}_N$  Ray ran ... okay

Ray hung a tire on a rope to make his daughter a swing.

Future constraints described by natural language

Ray ran to his daughter to make sure she was okay.

$Z$

$X$

# Takeaways

- Decoding is a challenging problem in natural language generation
- Human language distribution is noisy and doesn't reflect simple properties (i.e., *maximization*)
- Decoding algorithms can allow us to interject inductive biases that encourage properties of coherent NLG
- A lot more work to be done!

# Training Neural Text Generation Models

Antoine Bosselut

# Generation Model Basics

1. At each time step, model computes a score  $o_n$  for each token in our vocabulary,  $w_n \in V$

$$O_n = f(\{y\}_{<t})$$

$f(\cdot)$  is your model

2. Compute a probability distribution over these scores (usually softmax)

$P(\cdot)$  is your distribution over tokens

$$P(y_t = w_n | \{y\}_{<t}) = \frac{e^{o_n}}{\sum_{m=1}^M e^{o_m}}$$

3. Define a loss function to select a token from this distribution

# Maximizing Likelihood

- Trained to generate the next word given a set of preceding words




↑  
Allen

$y_0^*$


# Maximizing Likelihood

- Trained to generate the next word given a set of preceding words



# Maximizing Likelihood

- Trained to generate the next word given a set of preceding words



# Maximizing Likelihood

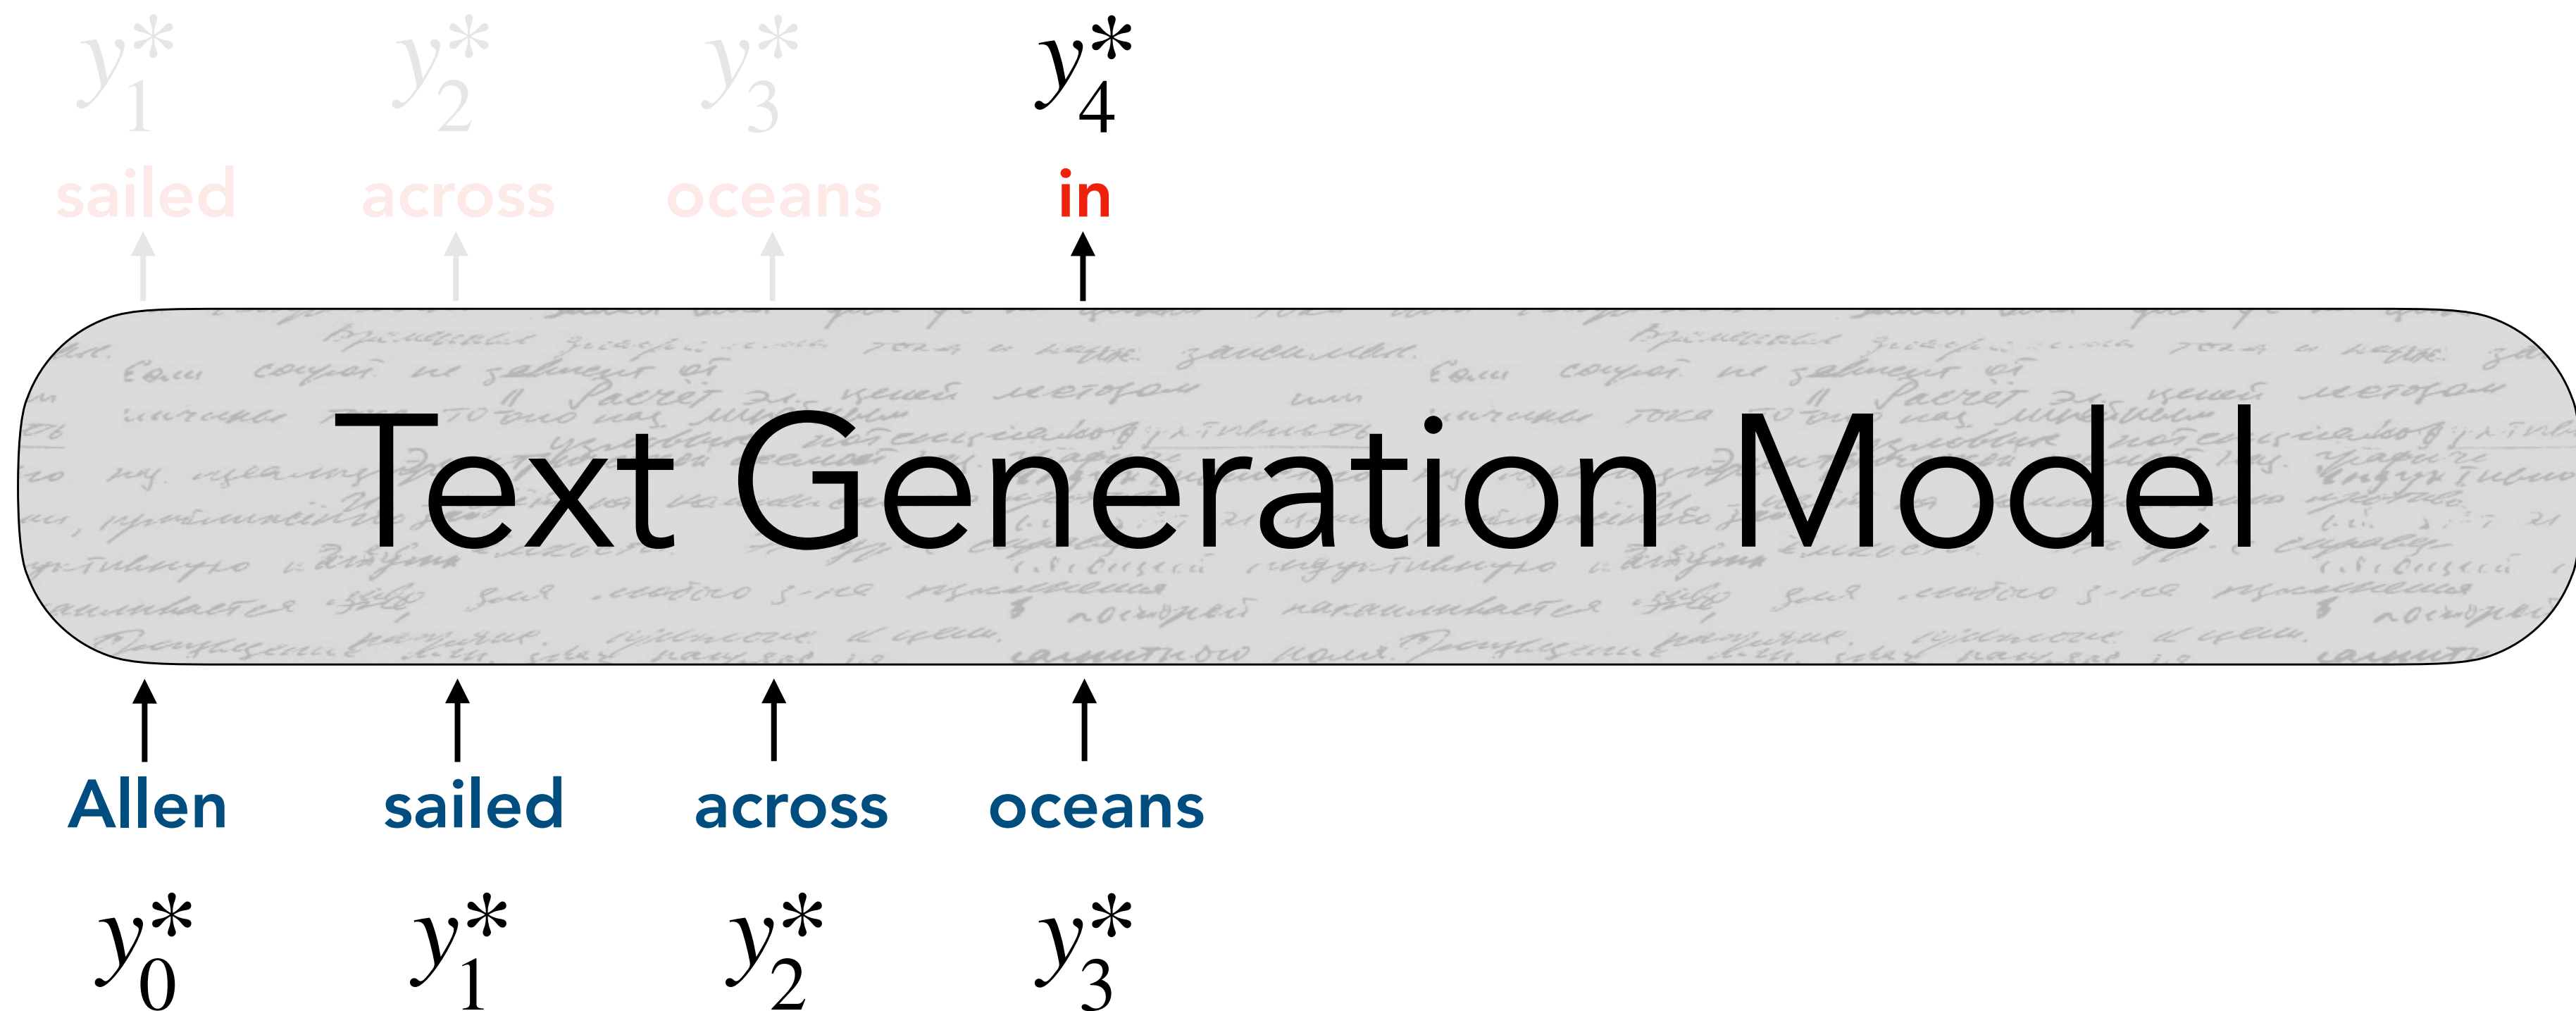
- Trained to generate the next word given a set of preceding words






# Maximizing Likelihood

- Trained to generate the next word given a set of preceding words




# Maximizing Likelihood

- Trained to generate the next word given a set of preceding words



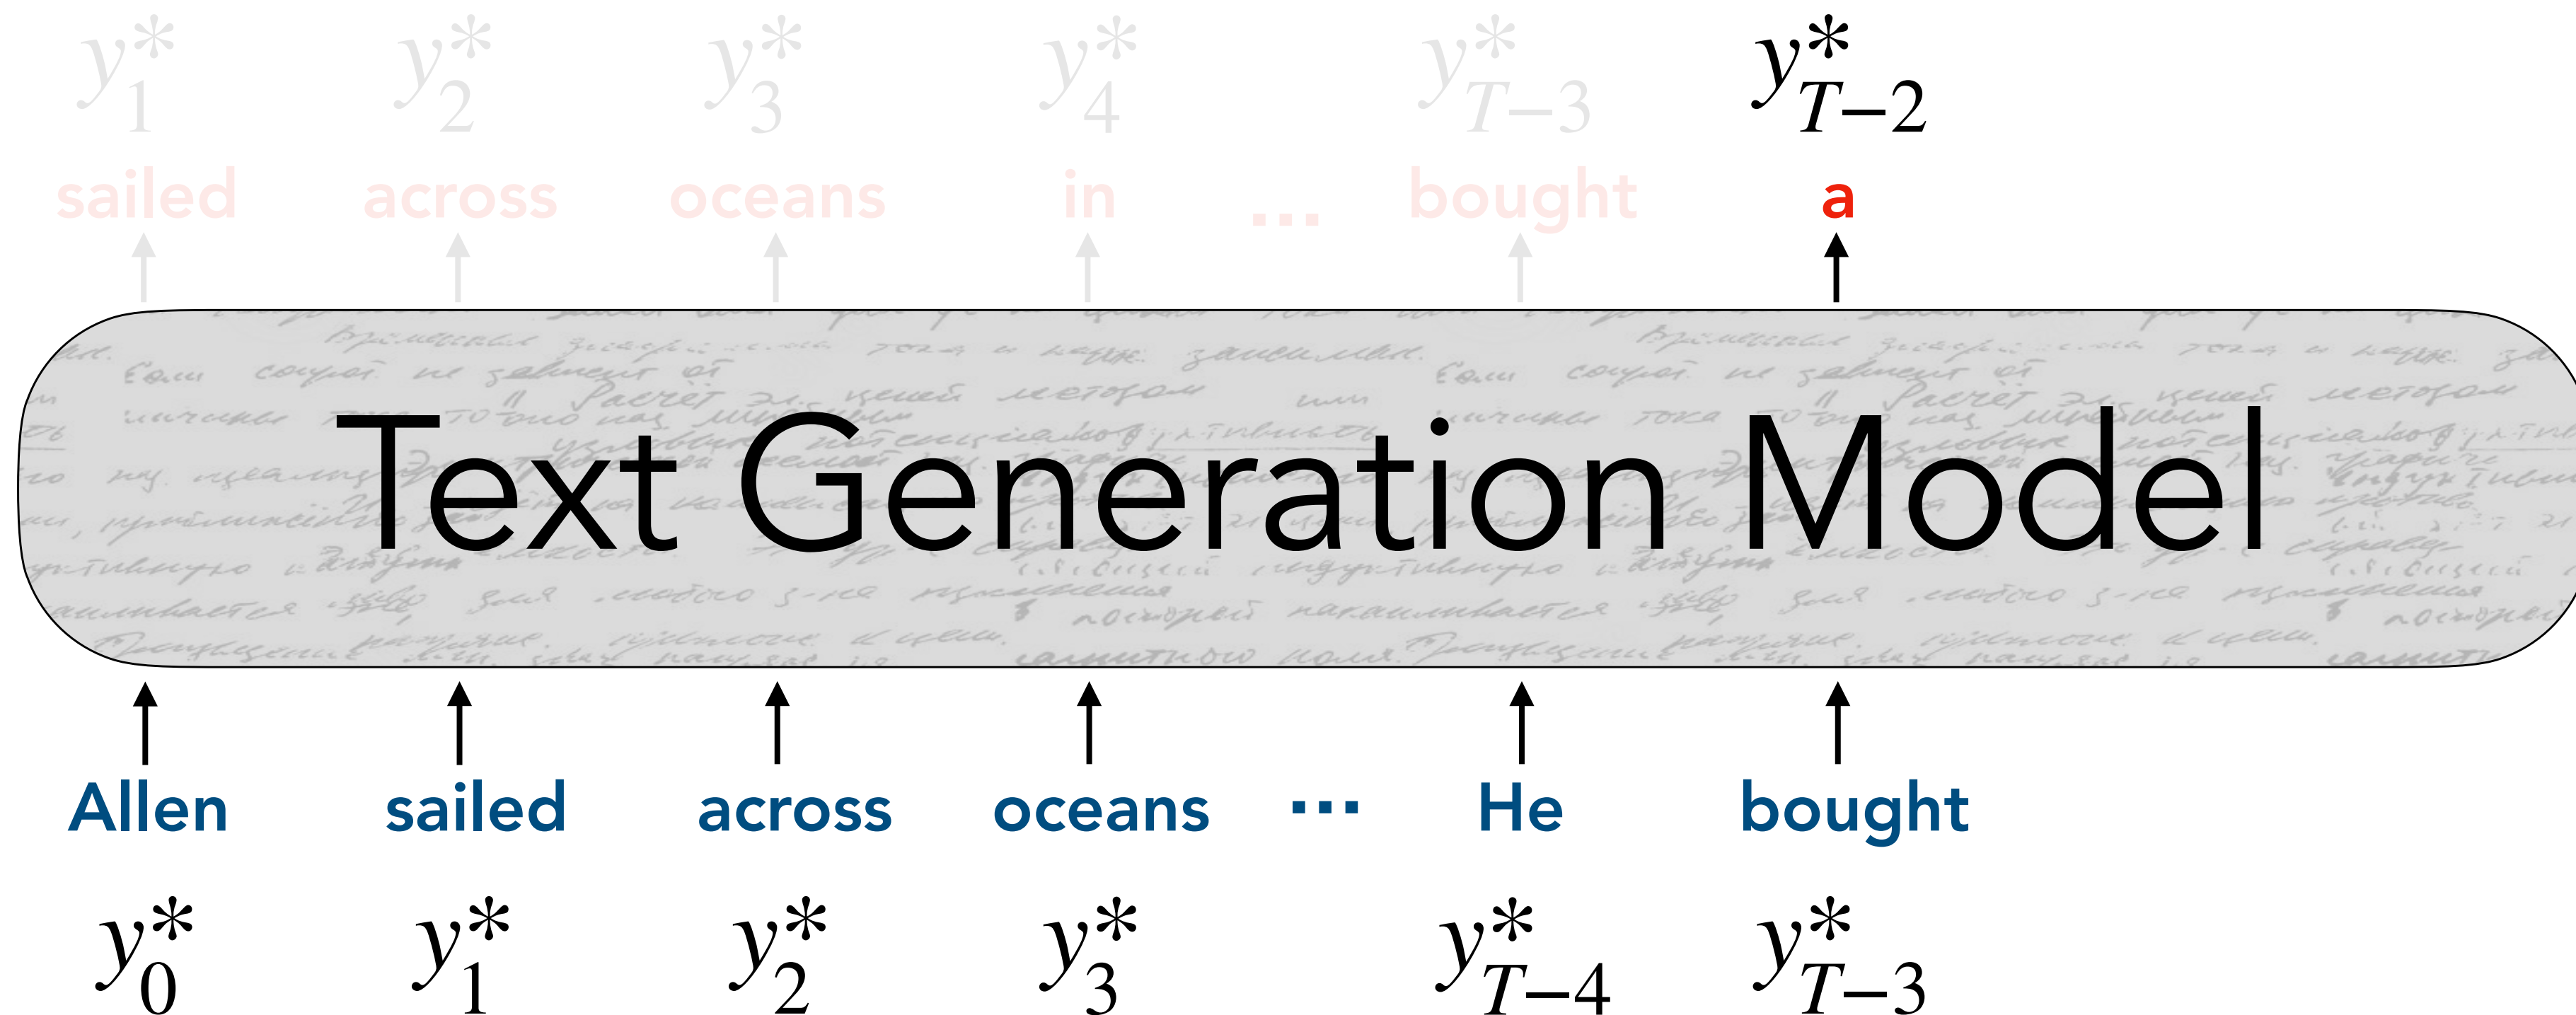
# Maximizing Likelihood

- Trained to generate the next word given a set of preceding words



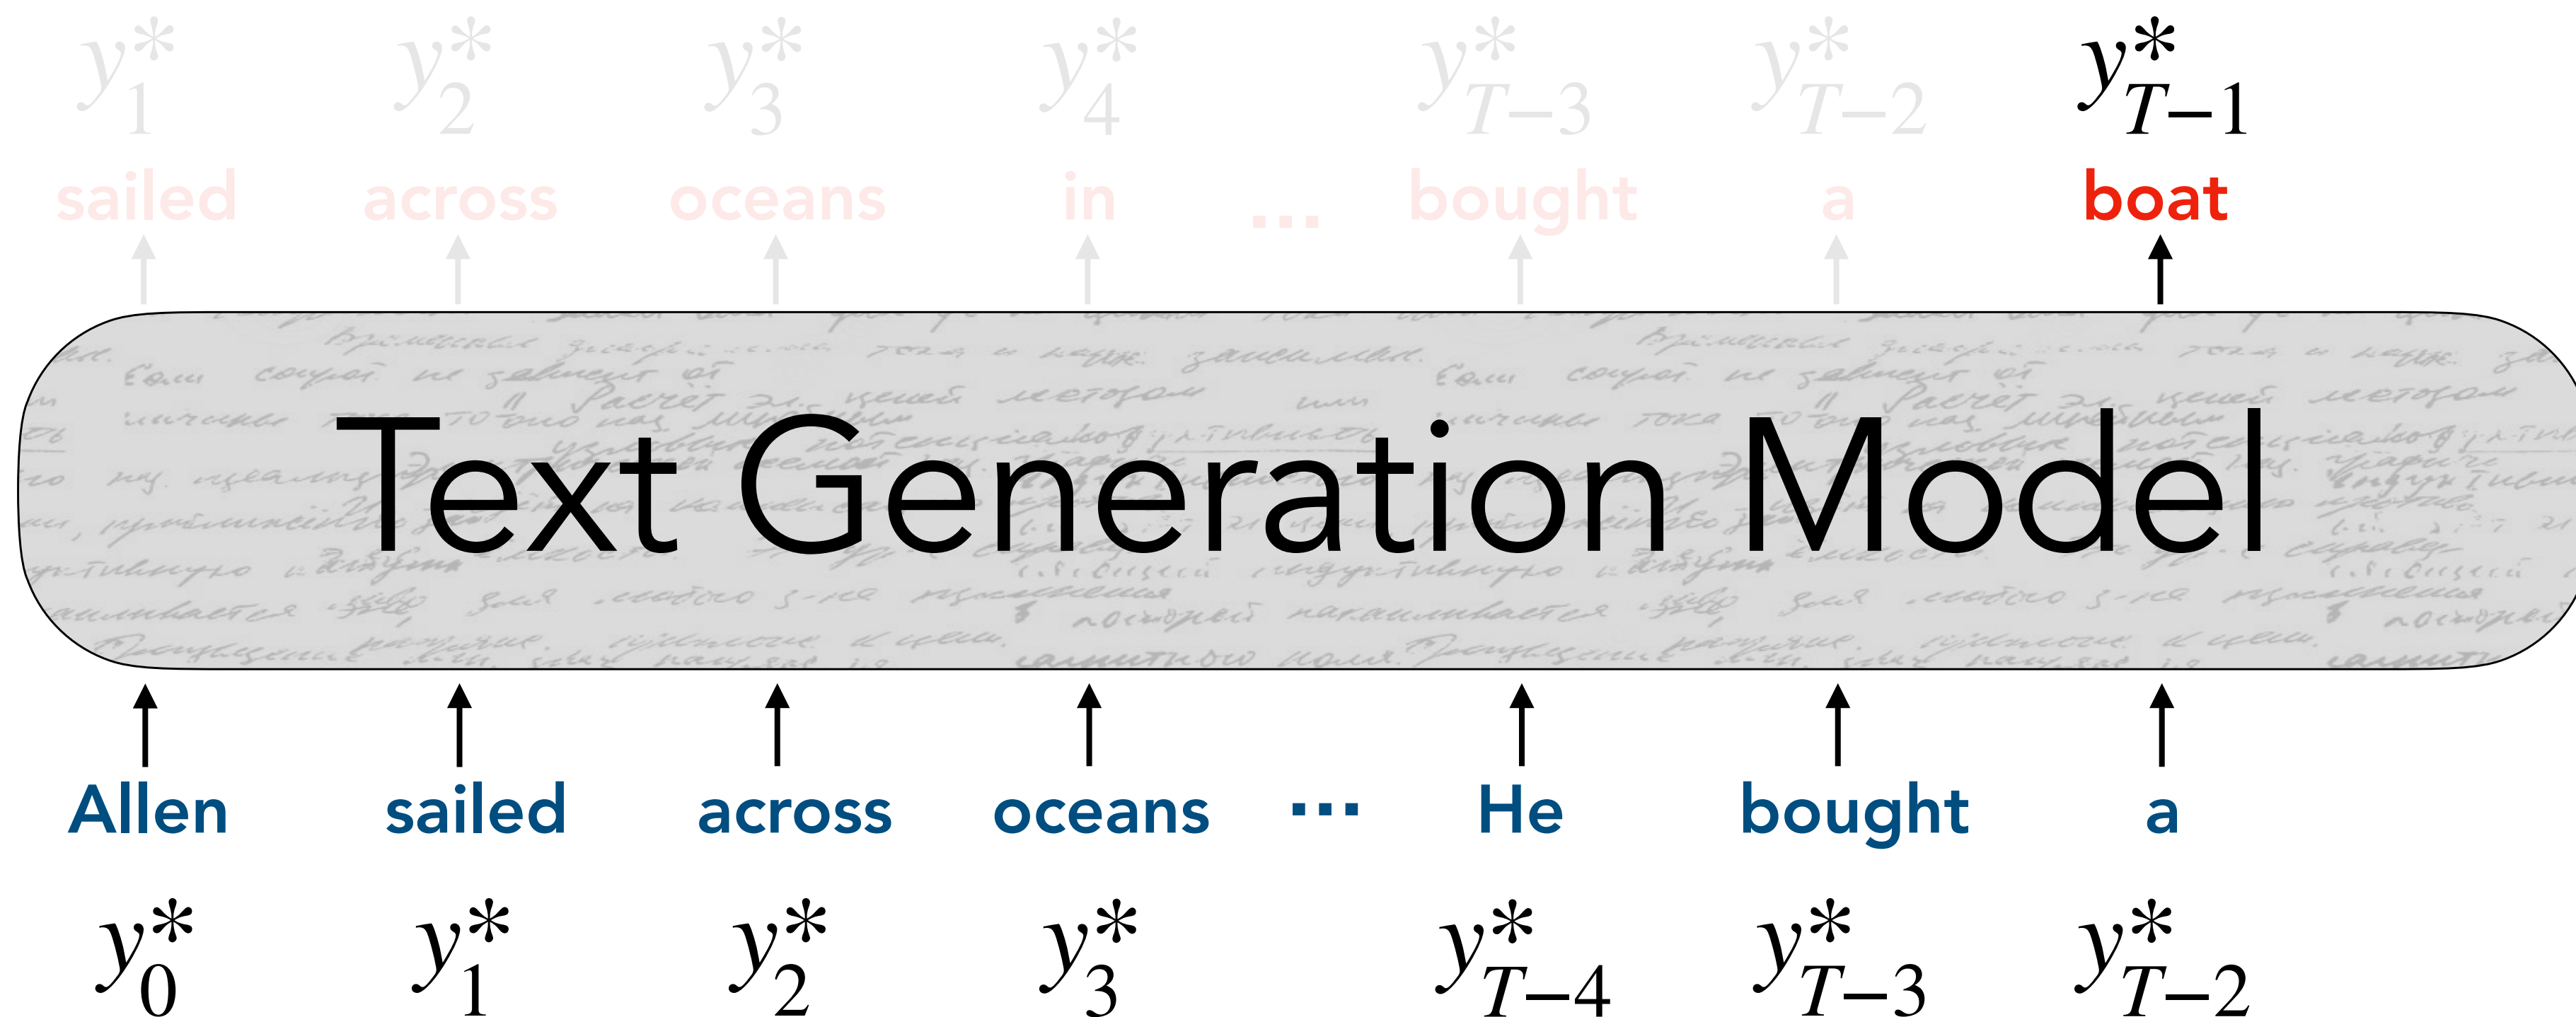
# Maximizing Likelihood

- Trained to generate the next word given a set of preceding words



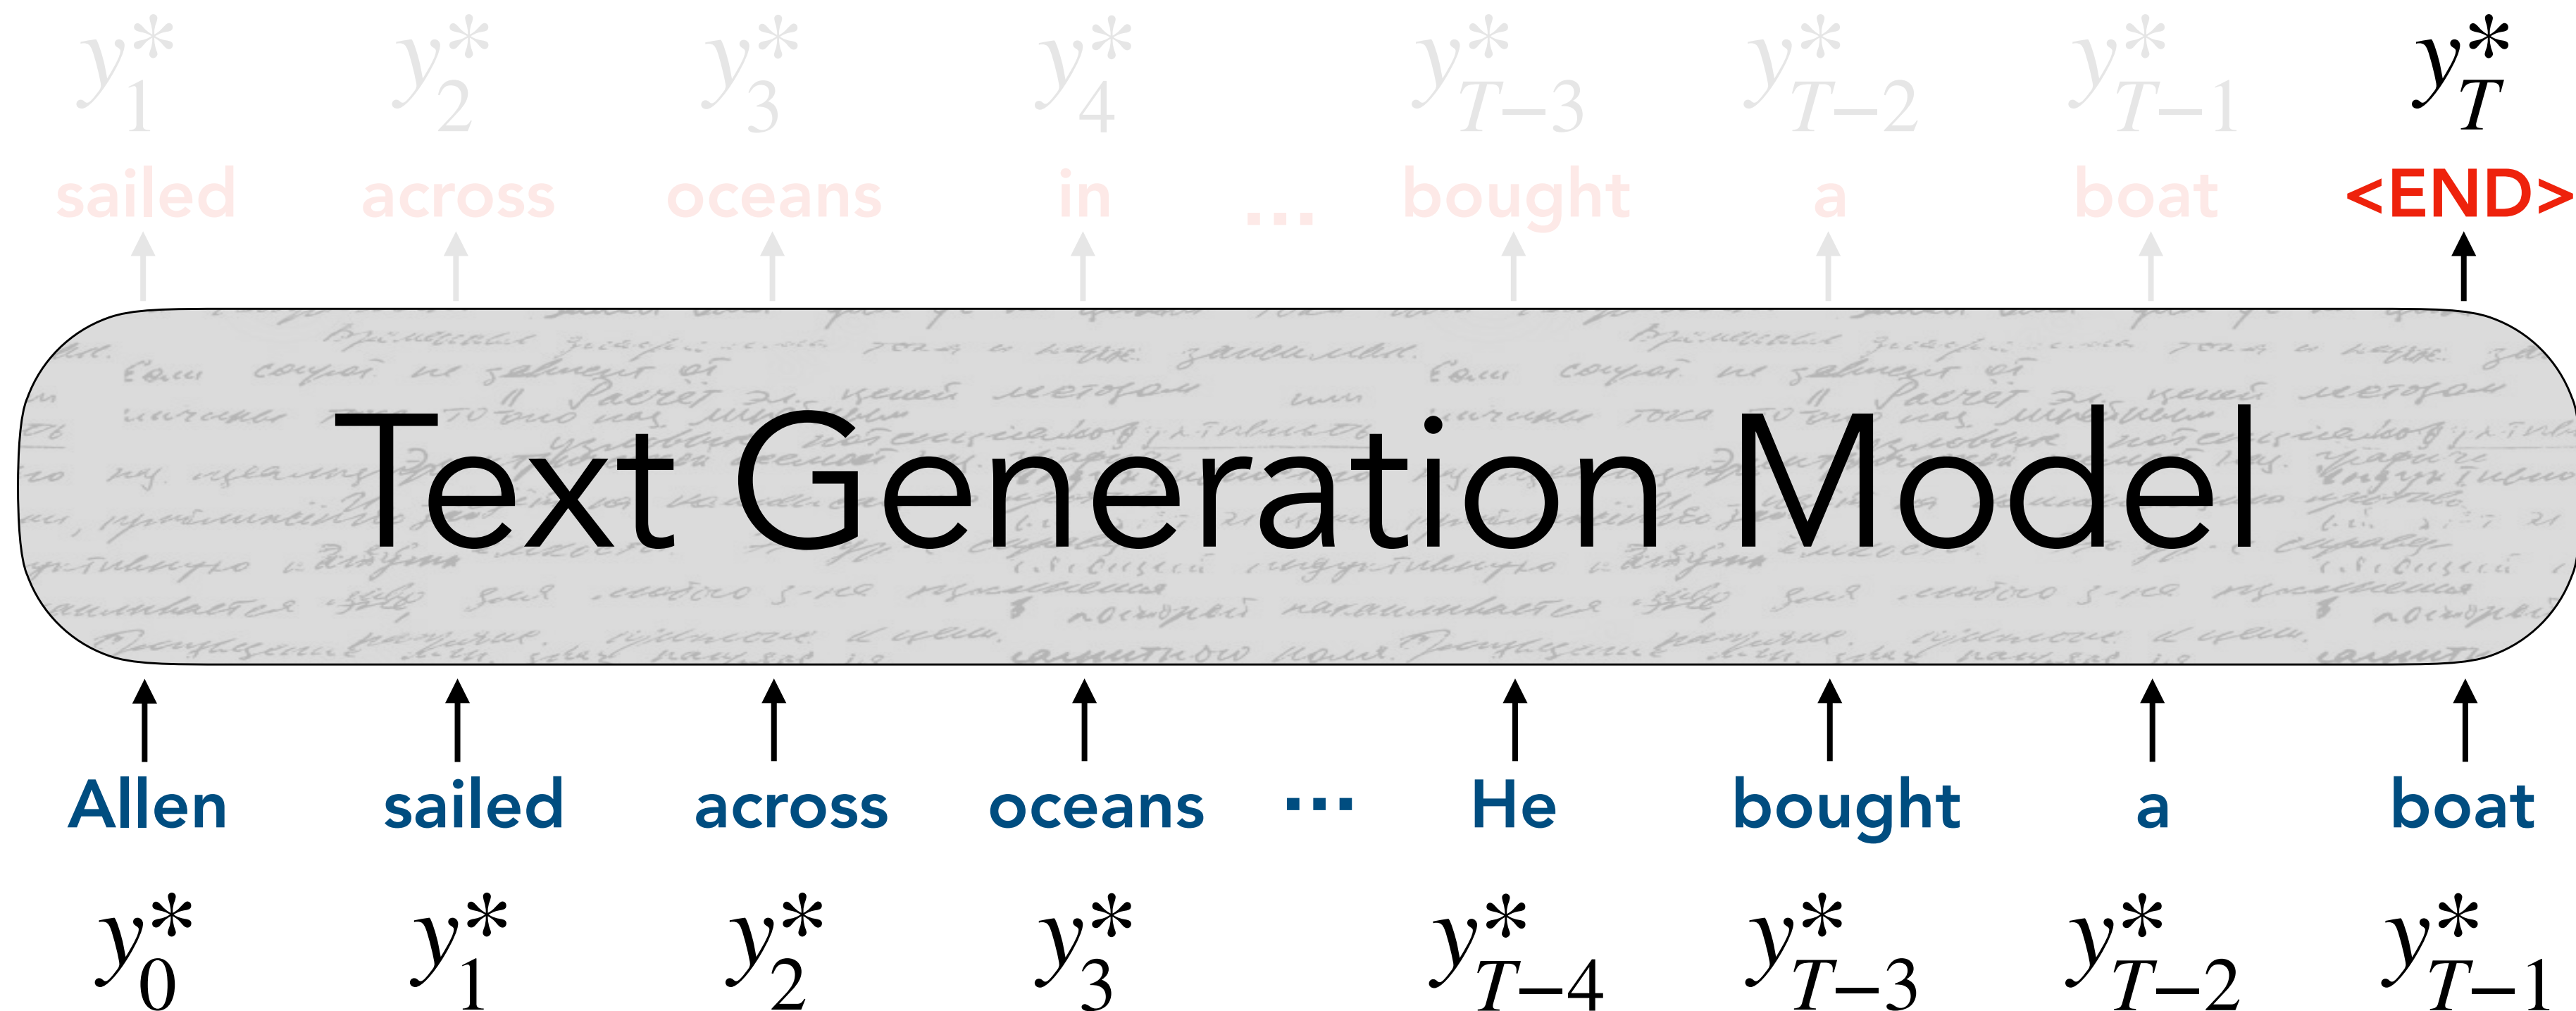
# Maximizing Likelihood

- Trained to generate the next word given a set of preceding words



# Maximizing Likelihood

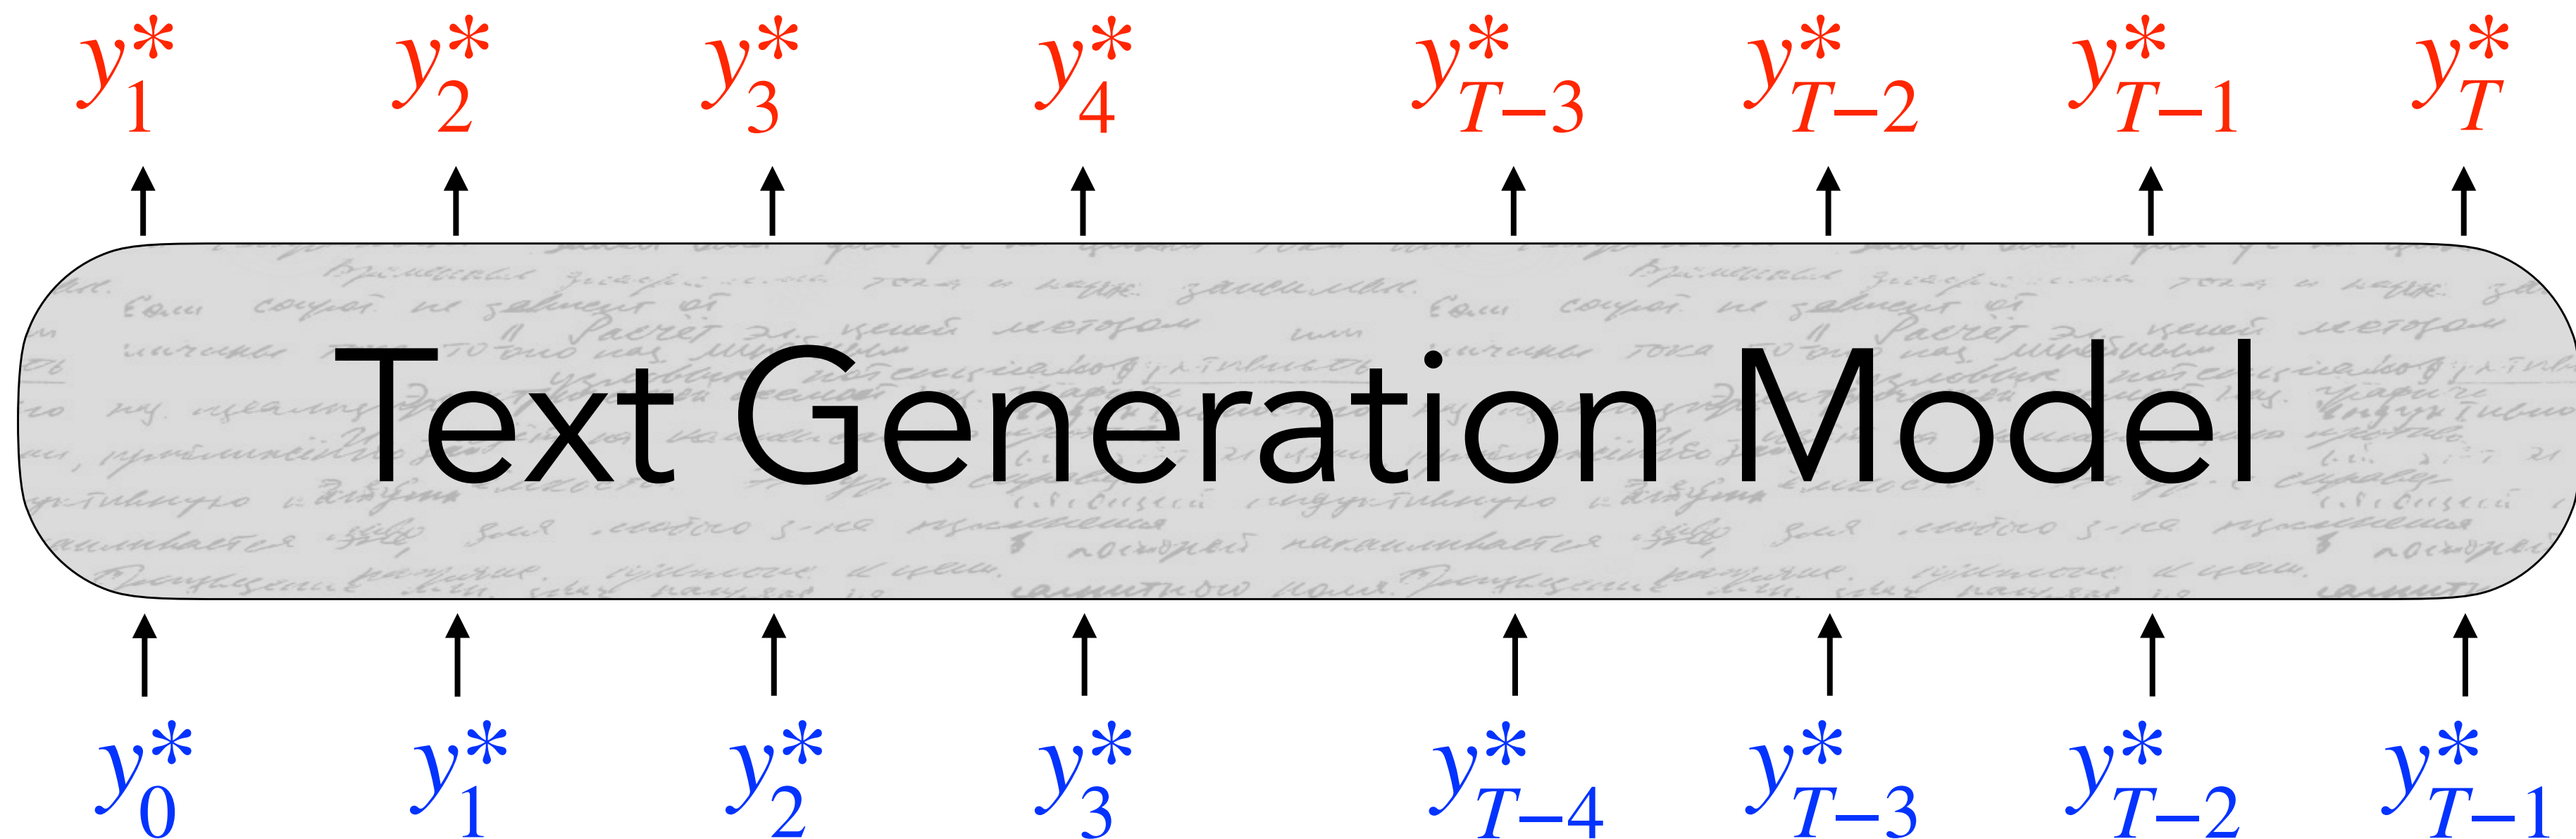
- Trained to generate the next word given a set of preceding words



# Maximizing Likelihood

- Trained to generate the next word given a set of preceding words

$$\mathcal{L} = - \sum_{t=1}^T \log P(y_t^* | \{y^*\}_{<t})$$




# Issue #1: MLE *discourages* diversity

**Context:** In a shocking finding, scientist discovered a herd of unicorns living in a remote, previously unexplored valley, in the Andes Mountains. Even more surprising to the researchers was the fact that the unicorns spoke perfect English.

**Continuation:** The study, published in the Proceedings of the National Academy of Sciences of the United States of America (PNAS), was conducted by researchers from the **Universidad Nacional Autónoma de México (UNAM)** and **the Universidad Nacional Autónoma de México (UNAM/Universidad Nacional Autónoma de México/ Universidad Nacional Autónoma de México/ Universidad Nacional Autónoma de México/ Universidad Nacional Autónoma de México...**



# Issue #1: MLE *discourages* diversity



# Alternatives: Unlikelihood Training

- Sequence-level Unlikelihood Training
- Given a set of undesired tokens  $\mathcal{C}$ , lower their likelihood in context

$$\mathcal{L}_{MLE}^t = -\log P(y_t^* | \{y^*\}_{<t}) \quad \mathcal{L}_{UL}^t = -\sum_{y_{neg} \in \mathcal{C}} \log(1 - P(y_{neg} | \{y^*\}_{<t}))$$

# Alternatives: Unlikelihood Training

- Sequence-level Unlikelihood Training
- Given a set of undesired tokens  $\mathcal{C}$ , lower their likelihood in context

$$\mathcal{L}_{MLE}^t = -\log P(y_t^* | \{y^*\}_{<t})$$

$$\mathcal{L}_{UL}^t = -\sum_{y_{neg} \in \mathcal{C}} \log(1 - P(y_{neg} | \{y^*\}_{<t}))$$



Typical Negative  
Loglikelihood objective

# Alternatives: Unlikelihood Training

- Sequence-level Unlikelihood Training
- Given a set of undesired tokens  $\mathcal{C}$ , lower their likelihood in context

$$\mathcal{L}_{MLE}^t = -\log P(y_t^* | \{y^*\}_{<t})$$

Typical Negative  
Loglikelihood objective

$$\mathcal{L}_{UL}^t = -\sum_{y_{neg} \in \mathcal{C}} \log(1 - P(y_{neg} | \{y^*\}_{<t}))$$

Unlikelihood objective  
lowers the probability of  
certain tokens

# Alternatives: Unlikelihood Training

- Sequence-level Unlikelihood Training
- Given a set of undesired tokens  $\mathcal{C}$ , lower their likelihood in context

$$\mathcal{L}_{MLE}^t = -\log P(y_t^* | \{y^*\}_{<t}) \quad \mathcal{L}_{UL}^t = -\sum_{y_{neg} \in \mathcal{C}} \log(1 - P(y_{neg} | \{y^*\}_{<t}))$$

$$\mathcal{L}_{ULE}^t = \mathcal{L}_{MLE}^t + \alpha \mathcal{L}_{UL}^t$$

Combine them for full  
unlikelihood training

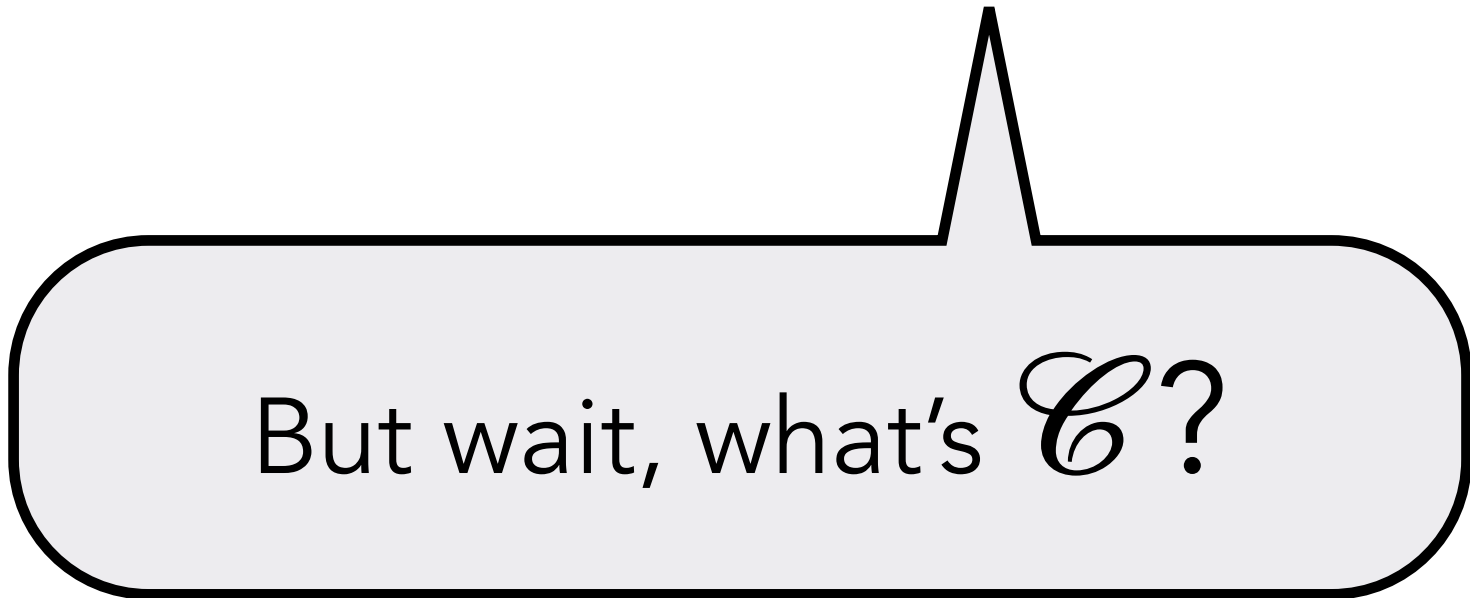
# Alternatives: Unlikelihood Training

- Sequence-level Unlikelihood Training
- Given a set of undesired tokens  $\mathcal{C}$ , lower their likelihood in context

$$\mathcal{L}_{MLE}^t = -\log P(y_t^* | \{y^*\}_{<t})$$

$$\mathcal{L}_{UL}^t = -\sum_{y_{neg} \in \mathcal{C}} \log(1 - P(y_{neg} | \{y^*\}_{<t}))$$

$$\mathcal{L}_{ULE}^t = \mathcal{L}_{MLE}^t + \alpha \mathcal{L}_{UL}^t$$



But wait, what's  $\mathcal{C}$ ?

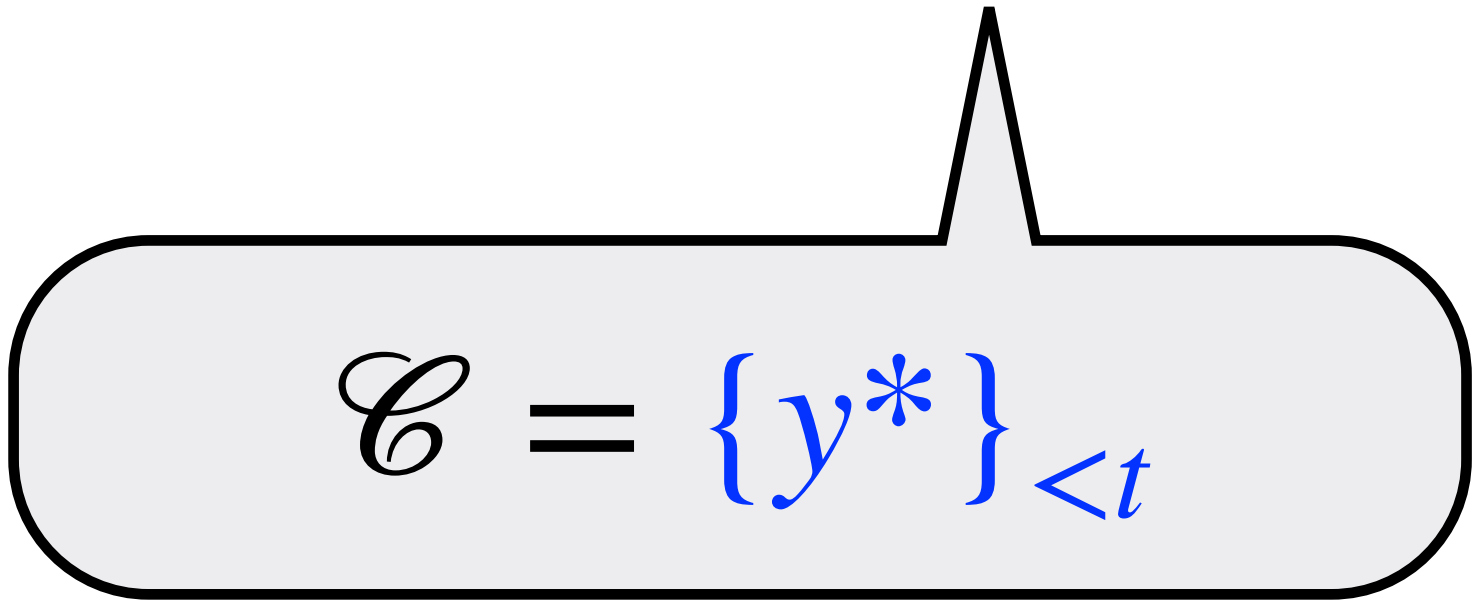
# Alternatives: Unlikelihood Training

- Sequence-level Unlikelihood Training
- Given a set of undesired tokens  $\mathcal{C}$ , lower their likelihood in context

$$\mathcal{L}_{MLE}^t = -\log P(y_t^* | \{y^*\}_{<t})$$

$$\mathcal{L}_{UL}^t = -\sum_{y_{neg} \in \mathcal{C}} \log(1 - P(y_{neg} | \{y^*\}_{<t}))$$

$$\mathcal{L}_{ULE}^t = \mathcal{L}_{MLE}^t + \alpha \mathcal{L}_{UL}^t$$


$$\mathcal{C} = \{y^*\}_{<t}$$


# Alternatives: $F^2$ Softmax

- Avoid likelihood issues by factorizing the softmax
- Initialize  $C$  frequency classes
- Distribute vocabulary into classes so that token frequency uniformly distributed across **and** between classes

## (i) Unique tokens, sorted by frequency

the of and to in a be but are first during  
only about game later three found music  
role match way common sometimes decision king  
mission organize scoring castle property curb ...

frequency




## (ii) Assigning frequency class


C1 the of



C2 and to in a be



C3 but are first during ...



Cn turbulent pervasive ...






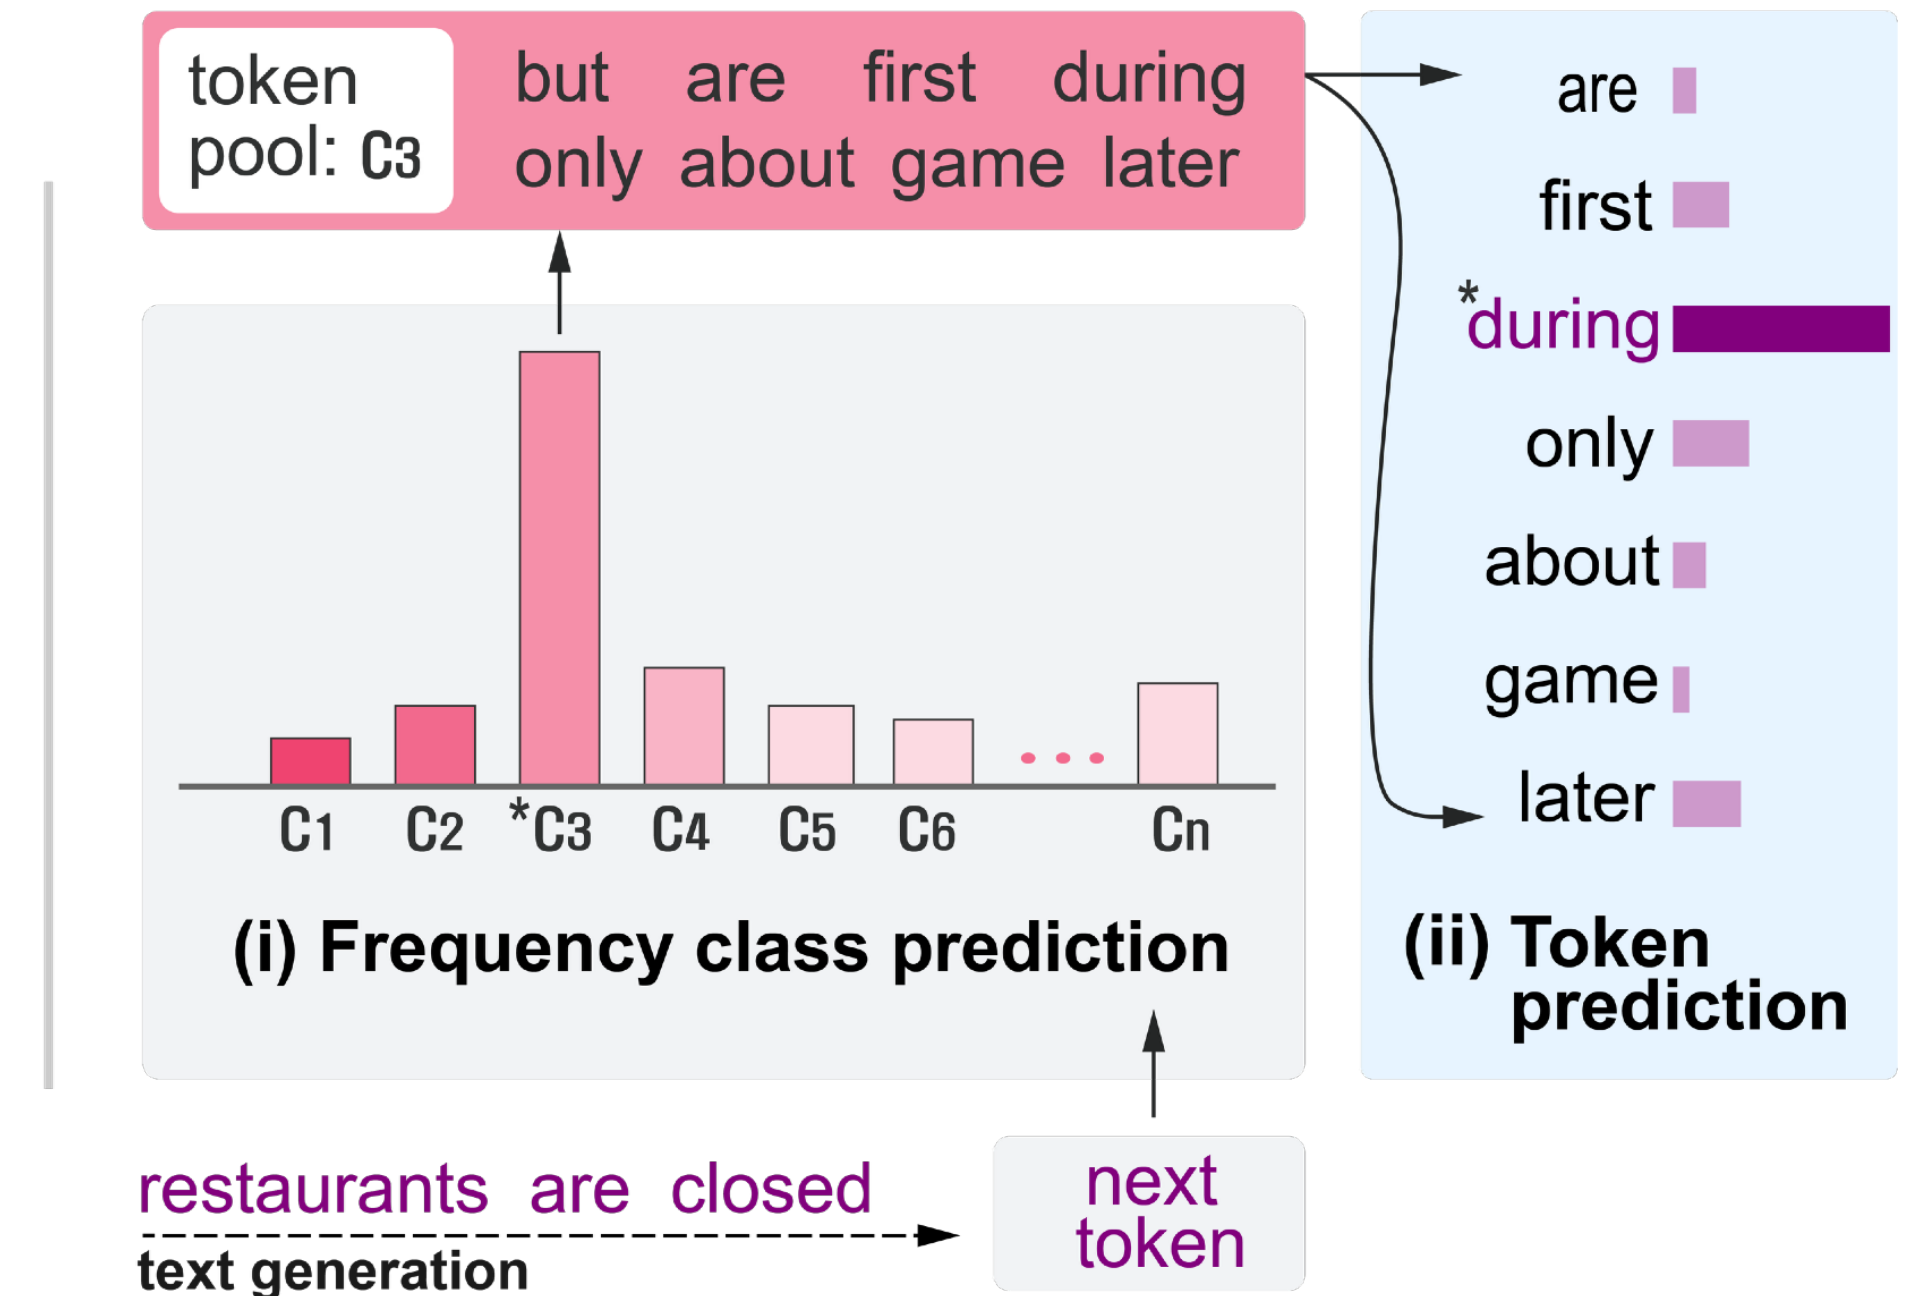
# Alternatives: F<sup>2</sup> Softmax

- Learn to select both frequency class and vocabulary token during training

$$P(y_t = w_n | \{y\}_{<t}) = \frac{e^{U_n h}}{\sum_{m=1}^M e^{U_m h}}$$



$$P(y_t = w_n | \{y\}_{<t}) = \left( \frac{e^{V_f h}}{\sum_{c=1}^C e^{V_c h}} \right) \left( \frac{e^{U_n h}}{\sum_{m=1}^{M_f} e^{U_m h}} \right)$$



# Issue #2: Exposure Bias


- During training, we condition on gold context tokens that are real human-generated text

$$\mathcal{L}_{MLE} = -\log P(y_t^* | \{y^*\}_{<t})$$

- During inference, we decode from distributions conditioned on previously generated tokens

$$\mathcal{L}_{dec} = -\log P(\hat{y}_t | \{\hat{y}\}_{<t})$$

# Issue #2: Exposure Bias



# Issue #2: Exposure Bias

- During training, we condition on gold context tokens that are real human-generated text

$$\mathcal{L}_{MLE} = -\log P(y_t^* | \{y^*\}_{<t})$$

- During inference, we decode from distributions conditioned on previously generated tokens

$$\mathcal{L}_{dec} = -\log P(\hat{y}_t | \{\hat{y}\}_{<t})$$

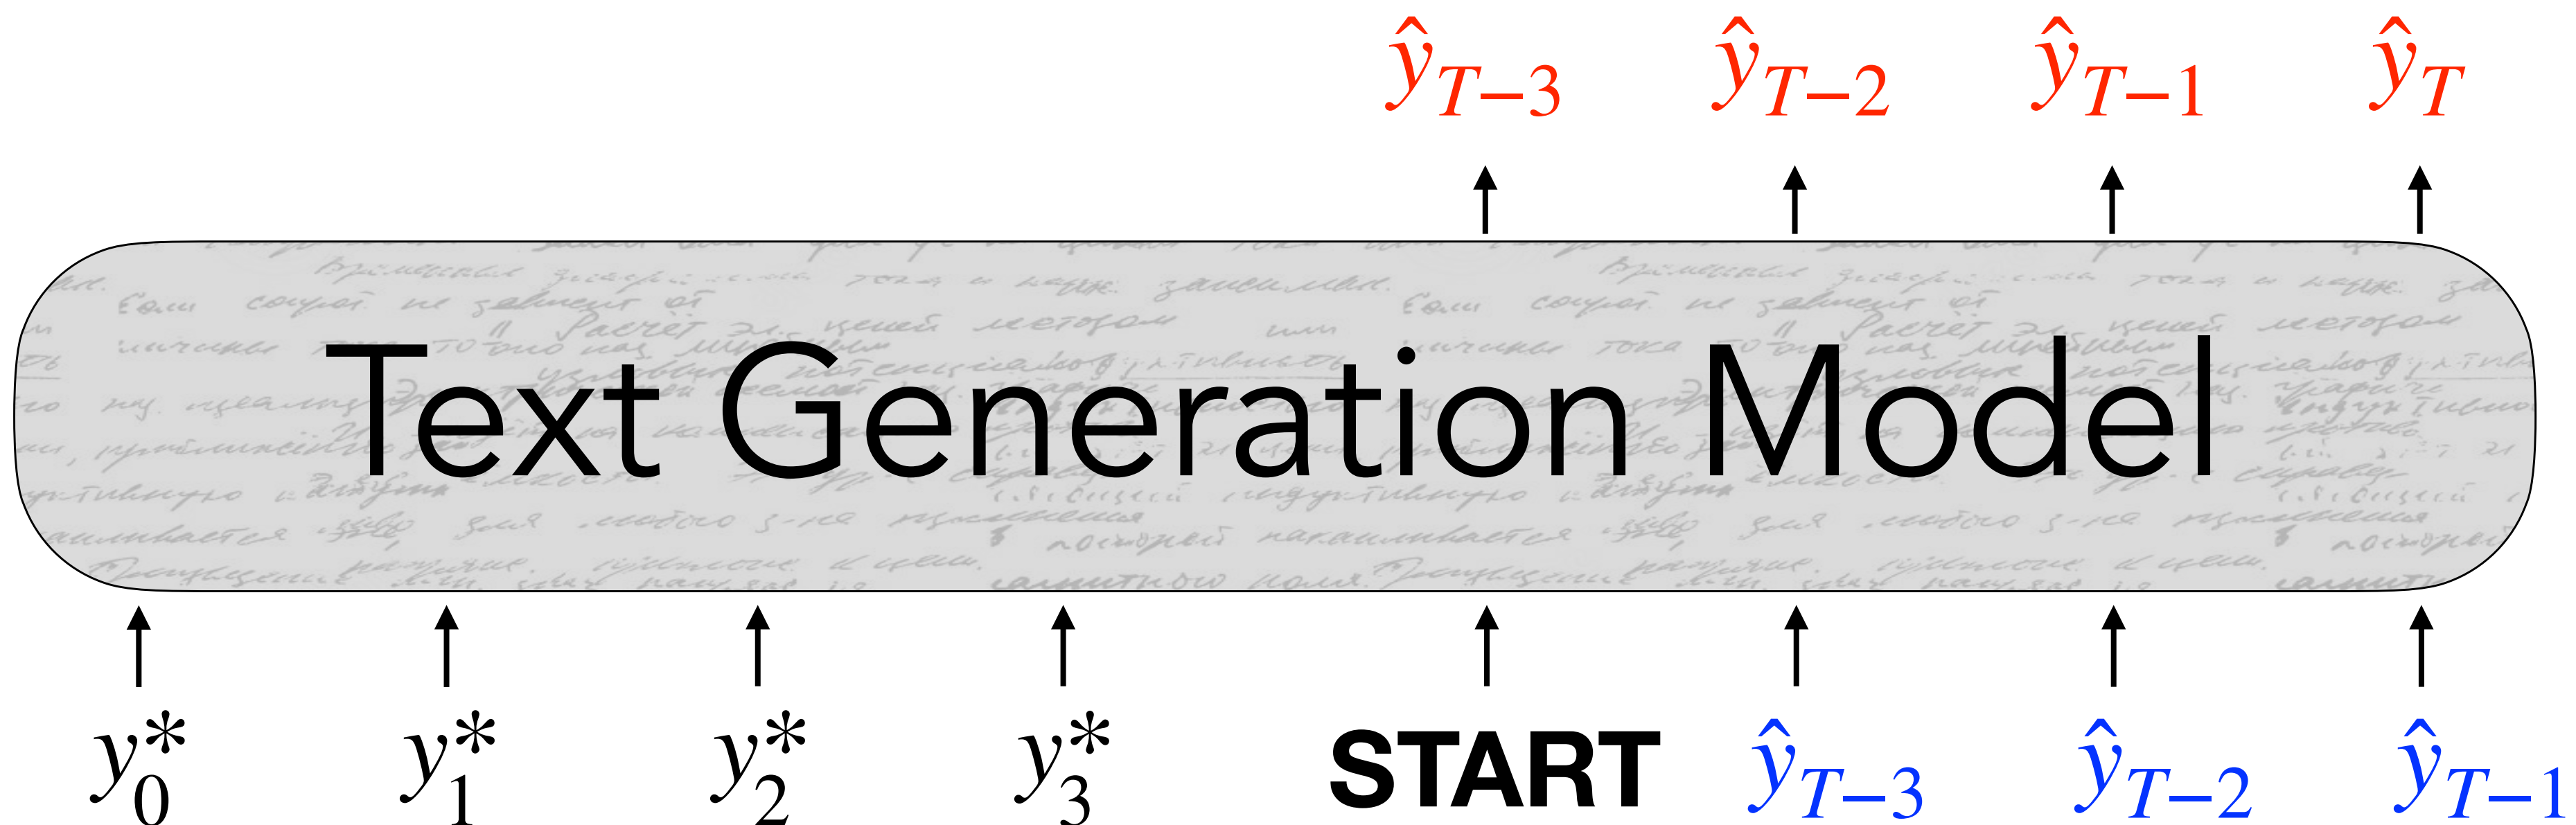
# Reinforcement Learning

- Cast a text generation model as a MDP
  - State is denoted by a preceding context
  - Actions are the words you can generate
  - Policy is the text generation model
  - Rewards are provided by an external source

# REINFORCE

- Trained to generate the next word given a set of preceding words

$$\mathcal{L}_{RL} = - \sum_{t=T-3}^T r(\hat{y}_t) \log P(\hat{y}_t | \{\hat{y}\}_{<t}; \{y^*\})$$



# Reward Estimation

- How do we define a reward function?
  - BLEU (machine translation; Ranzato et al., ICLR 2016; Wu et al., 2016)
  - ROUGE (summarization; Paulus et al., ICLR 2018; Celikyilmaz et al., NAACL 2018)
  - CIDEr (image captioning; Rennie et al., CVPR 2017)
  - SPIDEr (image captioning; Liu et al., ICCV 2017)

# Reward Estimation

- How do we define a reward function?
  - BLEU (machine translation; Ranzato et al., ICLR 2016; Wu et al., 2016)
  - ROUGE (summarization; Paulus et al., ICLR 2018; Celikyilmaz et al., NAACL 2020)
  - CIDEr (image captioning; Rennie et al., CVPR 2017)
  - SPIDEr (image captioning; Liu et al., ICCV 2017)

**Optimizing for the task vs. Gaming the reward**



# What behaviors can we tie to rewards?

- Cross-modal consistency (Ren et al., CVPR 2017)
- Simplicity (Zhang and Lapata, EMNLP 2017)
- Temporal consistency (Bosselut et al., NAACL 2018)
- Politeness (Tan et al., TACL 2018)
- Paraphrasing (Li et al., EMNLP 2018)
- Sentiment (Gong et al., NAACL 2019)
- Formality (Gong et al., NAACL 2019)

# Implementation Thoughts

- Credit Assignment

$$r(\hat{y}_t) \text{ vs. } r(\hat{Y})$$

# Implementation Thoughts

- Credit Assignment

$$r(\hat{y}_t) \text{ vs. } r(\hat{Y})$$

- Set appropriate baseline

$$\mathcal{L}_{RL} = - \sum (r(\hat{y}_t) - b) \log P(\hat{y}_t | \{\hat{y}\}_{<t}; \{y^*\})$$

# Implementation Thoughts

- Credit Assignment

$$r(\hat{y}_t) \text{ vs. } r(\hat{Y})$$

- Set appropriate baseline

$$\mathcal{L}_{RL} = - \sum (r(\hat{y}_t) - b) \log P(\hat{y}_t | \{\hat{y}\}_{<t}; \{y^*\})$$

- Mix with MLE

$$\mathcal{L} = \mathcal{L}_{MLE} + \alpha \mathcal{L}_{RL}$$

# What if you don't know what to use as reward?

- **Adversarial Learning!**
- Use an adversarially-learned scoring function to provide rewards
- Still often uses REINFORCE
- Dialogue systems (Li et al., EMNLP 2017), Visual storytelling (Wang et al., ACL 2018)


# Human-in-the-loop Learning

## 1 Collect human feedback


A Reddit post is sampled from the Reddit TL;DR dataset.



Various policies are used to sample a set of summaries.



Two summaries are selected for evaluation.



A human judges which is a better summary of the post.



"j is better than k"

## 2 Train reward model

One post with two summaries judged by a human are fed to the reward model.



The reward model calculates a reward  $r$  for each summary.



$r_j$

$r_k$

The loss is calculated based on the rewards and human label, and is used to update the reward model.

$$\text{loss} = \log(\sigma(r_j - r_k))$$


"j is better than k"

## 3 Train policy with PPO


A new post is sampled from the dataset.




The policy  $\pi$  generates a summary for the post.



The reward model calculates a reward for the summary.



The reward is used to update the policy via PPO.



$r$

# Takeaways

- Maximum likelihood estimation is still the premier algorithm for training text generation models
- Diversity is an issue with MLE, so new approaches focus on mitigating the effects of common words
- Exposure bias causes text generation models to lose coherence easily, so learning from its own samples is a promising way forward
- Reinforcement learning allows models to learn tough to quantify behaviors
- Much more!

**EPFL**





# Decoding References

- [1] Gulcehre et al., On Using Monolingual Corpora in Neural Machine Translation. arXiv 2015
- [2] Wu et al., Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation. arxiv 2016
- [3] Venugopalan et al., Improving LSTM-based Video Description with Linguistic Knowledge Mined from Text. EMNLP 2016
- [4] Li et al., A Diversity-Promoting Objective Function for Neural Conversation Models. EMNLP 2018
- [5] Paulus et al., A Deep Reinforced Model for Abstractive Summarization. ICLR 2018
- [6] Celikyilmaz et al., Deep Communicating Agents for Abstractive Summarization. NAACL 2018
- [7] Holtzman et al., Learning to Write with Cooperative Discriminators. ACL 2018
- [8] Fan et al., Hierarchical Neural Story Generation. ACL 2018
- [9] Gabriel et al., Cooperative Generator-Discriminator Networks for Abstractive Summarization with Narrative Flow. arXiv 2019
- [10] Dathathri et al., Plug and Play Language Models: A Simple Approach to Controlled Text Generation. ICLR 2020
- [11] Holtzman et al., The Curious Case of Neural Text Degeneration. ICLR 2020
- [12] Khandelwal et al., Generalization through Memorization: Nearest Neighbor Language Models. ICLR 2020
- [13] Qin et al., Back to the Future: Unsupervised Backprop-based Decoding for Counterfactual and Abductive Commonsense Reasoning. EMNLP 2020
- [14] Goyal and Durrett, Evaluating Factuality in Generation with Dependency-level Entailment. Findings of EMNLP 2020
- [15] Lu et al., NeuroLogic Decoding: (Un)supervised Neural Text Generation with Predicate Logic Constraints. arXiv 2020

# Training References

- [1] Bengio et. al., Scheduled Sampling for Sequence Prediction with Recurrent Neural Networks. NeurIPS 2015
- [2] Ranzato et al., Sequence Level Training with Recurrent Neural Networks. ICLR 2016
- [3] Wu et al., Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation. 2016
- [4] Ren et al., Deep Reinforcement Learning-based Image Captioning with Embedding Reward. CVPR 2017
- [5] Rennie et al., Self-critical Sequence Training for Image Captioning. CVPR 2017
- [6] Liu et al., Improved Image Captioning via Policy Gradient Optimization of SPIDER. ICCV 2017
- [7] Zhang and Lapata, Sentence Simplification with Deep Reinforcement Learning. EMNLP 2017
- [8] Li et al., Adversarial Learning for Neural Dialogue Generation. EMNLP 2017
- [9] Paulus et al., A Deep Reinforced Model for Abstractive Summarization. ICLR 2018
- [10] Celikyilmaz et al., Deep Communicating Agents for Abstractive Summarization. NAACL 2018
- [11] Bosselut et al., Discourse-Aware Neural Rewards for Coherent Text Generation. NAACL 2018
- [12] Wang et al., No Metrics Are Perfect: Adversarial Reward Learning for Visual Storytelling. ACL 2018
- [13] Tan and Bansal, Polite Dialogue Generation Without Parallel Data. TACL 2018
- [14] Li et al., Paraphrase Generation with Deep Reinforcement Learning. EMNLP 2018
- [15] Gong et al., Reinforcement Learning Based Text Style Transfer without Parallel Training Corpus. NAACL 2019
- [16] Holtzman et. al., The Curious Case of Neural Text Degeneration. ICLR 2020
- [17] Welleck et. al., Neural Text Generation With Unlikelihood Training. ICLR 2020
- [18] Stiennon et al., Learning to summarize from human feedback. NeurIPS 2020