Monte Carlo Tree Search for Recipe Generation using GPT-2

Karan Taneja*^, Richard Segal^, Richard Goodwin^

*School of Interactive Computing
Georgia Institute of Technology

^Computational Creativity Group
IBM Research

ICCC'23

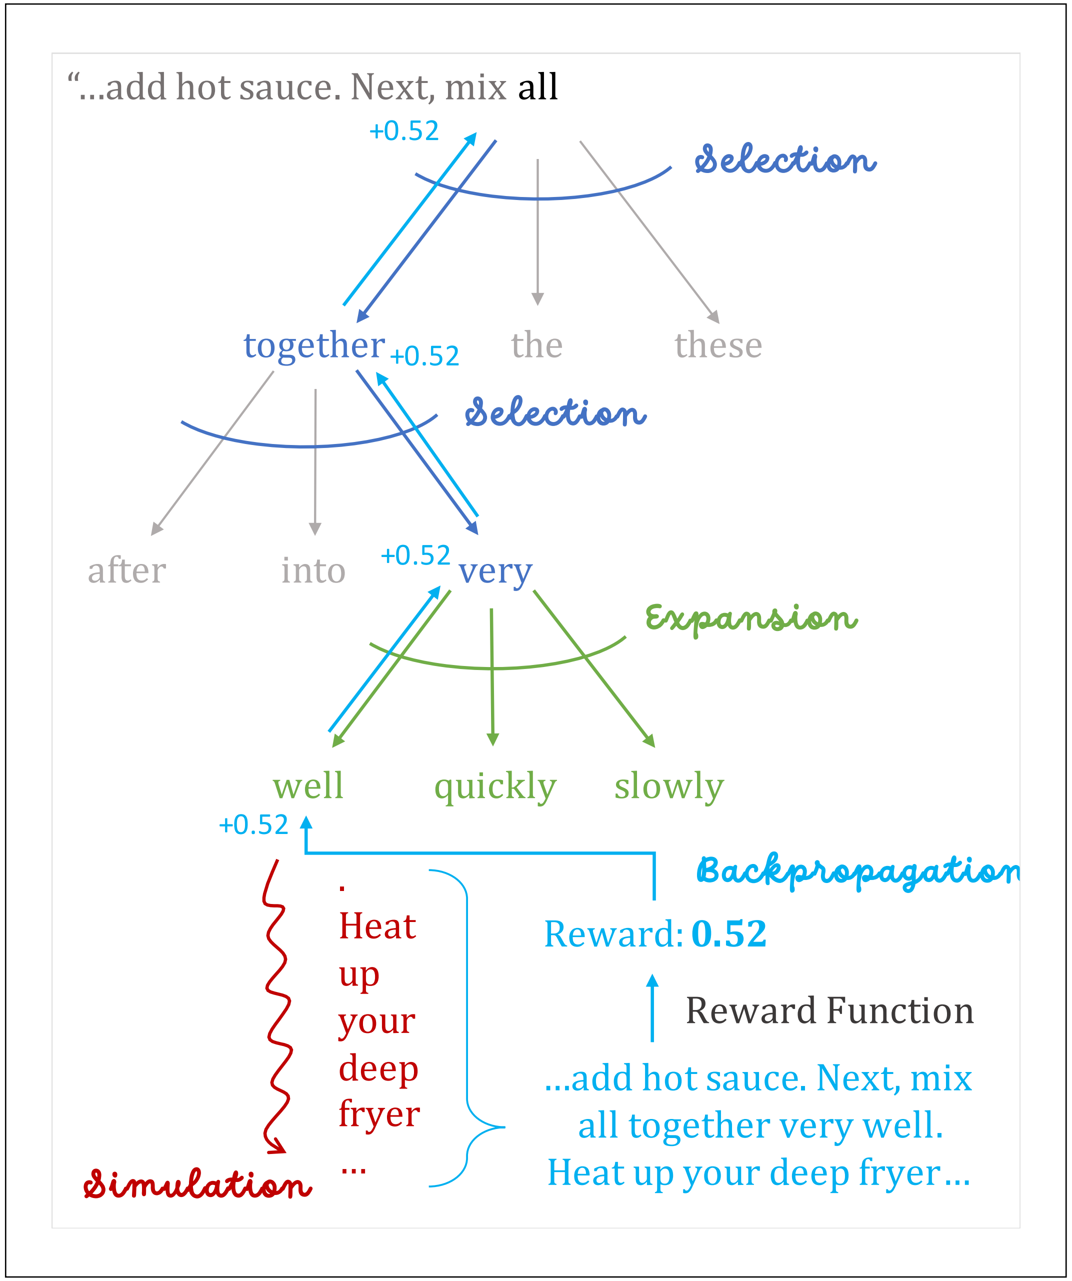

Recipe Generation

Creative process of recipe design:

- Inspiration for new recipes
- Writing recipe drafts
- Exploring flavor combinations

Large language models can

- Generate multiple possible recipes
- Complete incomplete ingredient lists
- Generate recipe instructions

Chocolate Chip Cookies
Ingredients:
1/2 cup unsalted butter
3/4 cup brown sugar
3/4 cup ?

GPT-2

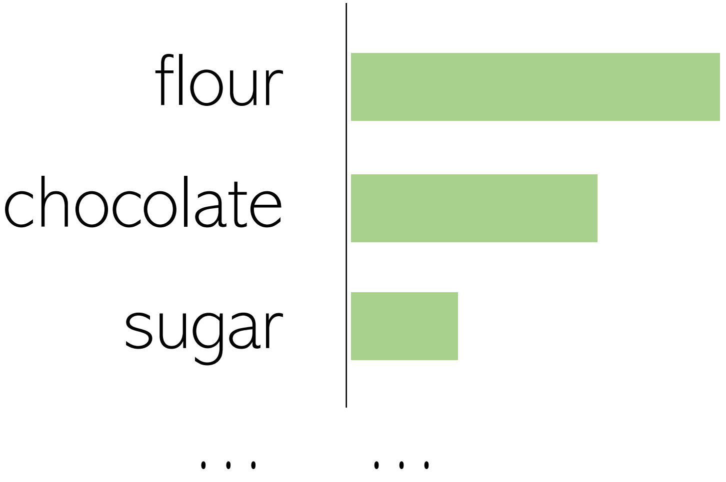

p(**next token** | Chocolate Chip Cookies … brown sugar; 3/4 cup)

$$p(x_1, x_2, \ldots, x_n) = p(x_1)p(x_2|x_1)p(x_3|x_1, x_2) \ldots p(x_n|x_1, \ldots, x_{n-1})$$

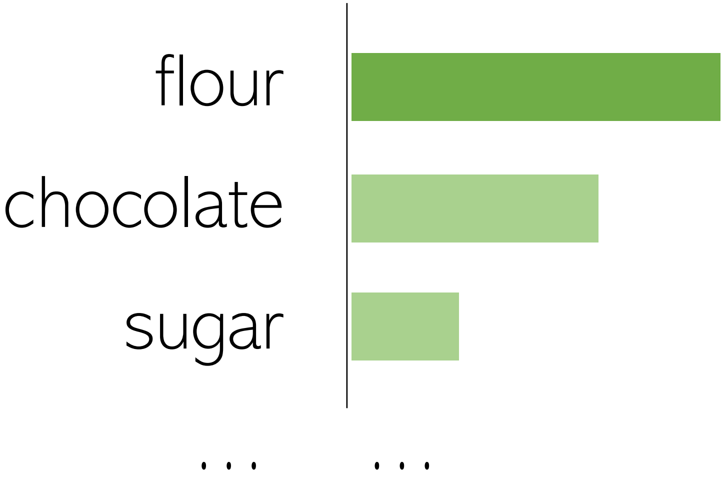

p(**next token** | Chocolate Chip Cookies ... brown sugar; 3/4 cup)

Chocolate Chip Cookies
Ingredients:
1/2 cup unsalted butter
3/4 cup brown sugar
3/4 cup flour
?

GPT-2

Greedy Sampling

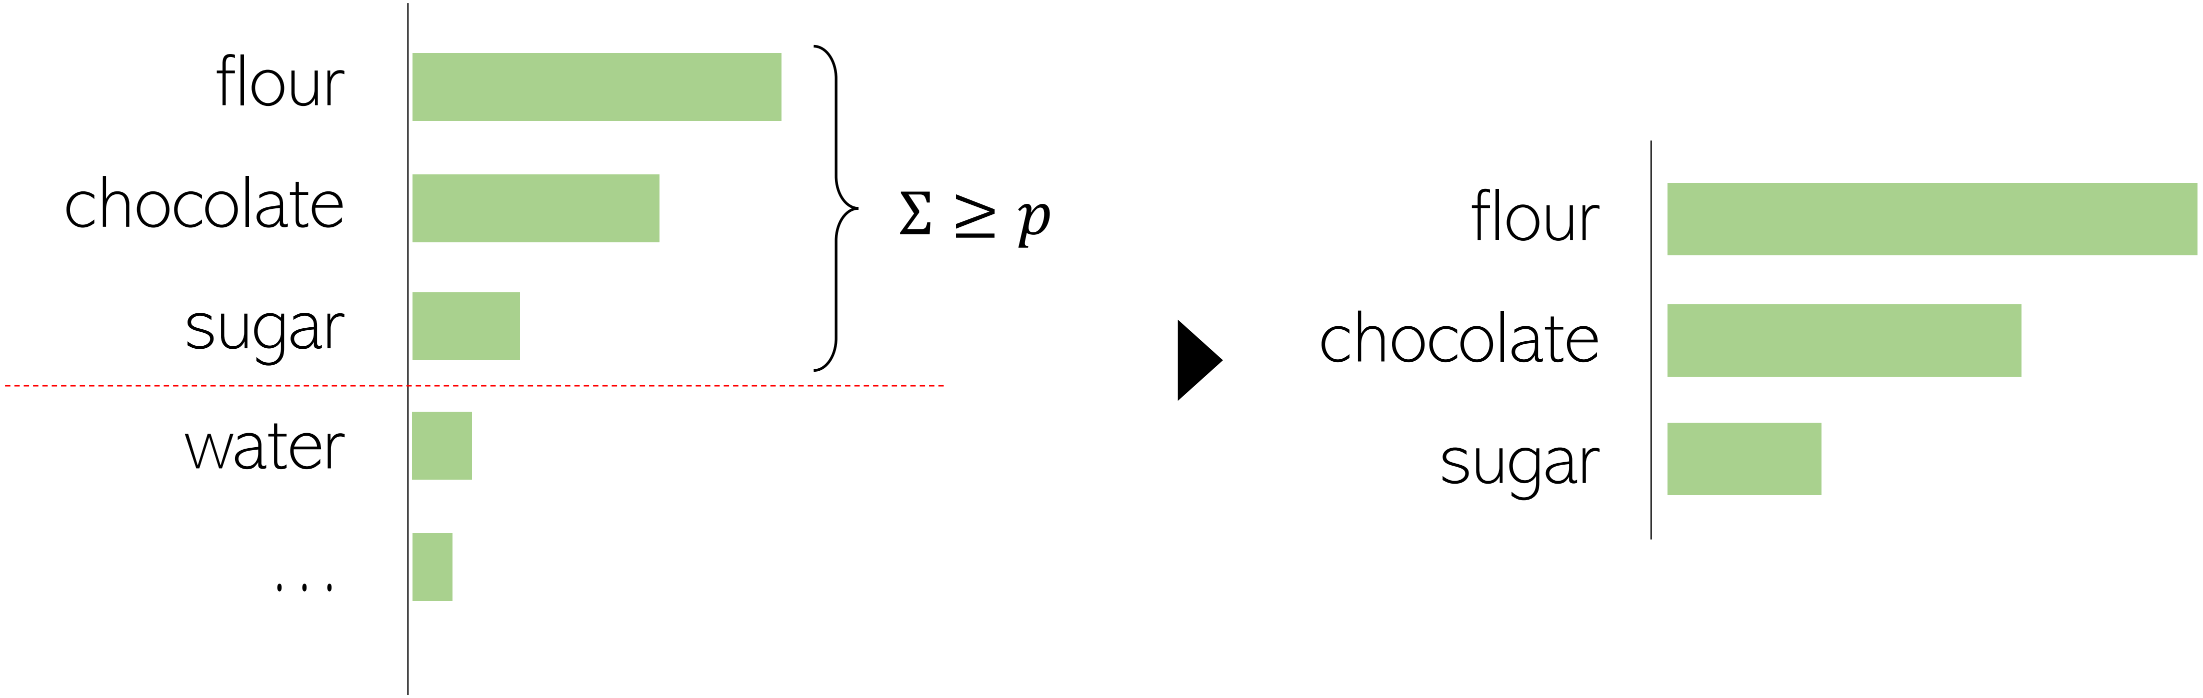

top-p Sampling or Nucleus Sampling

Limitations of previous recipe generation methods

Repeated ingredients

...
1 teaspoon black pepper;
1 cup carrots, diced;
1 tsp ground cinnamon;
1 teaspoon black pepper;
other...
...

Inconsistencies

Recipe Name: **Chicken** Soup

Ingredients:
1 cup diced carrots;
1 teaspoon white pepper;
3 cups water;
salt to taste

No 'chicken' *in the ingredients list*

LLMs generate text by sampling one token at a time and appending it to the existing text.

There is a high focus on local coherence and a lack of attention to the long-term view of the given context.

Fine-tuned models are typically trained on small datasets which leads to subpar generalization.

Training GPT-2 for Recipe Generation

Training GPT-2 for Recipe Generation

Recipe Name <|startofname|> Chocolate Chip Cookies <|endofname|>

<|startofingr>

Ingredients 1/2 cup unsalted butter; …

<|endofingr|>

Training GPT-2 for Recipe Generation

Recipe Name

Ingredients

Instructions

<|startofname|> Chocolate Chip Cookies <|endofname|>
<|startofingr>
1/2 cup unsalted butter; ...
<|endofingr|>
<|startofinst|>
Evenly position 2 racks in the middle of the oven and preheat to 375 degrees F.
Line 2 baking sheets with parchment paper . . .
<|endofinst|>

Monte Carlo Tree Search

A search algorithm used in AI agents for playing strategy games such as Chess, Go, and Checkers.

Swiechowski et al. 2022. Monte Carlo Tree Search: A review of recent modifications and applications. (Artificial Intelligence Review)

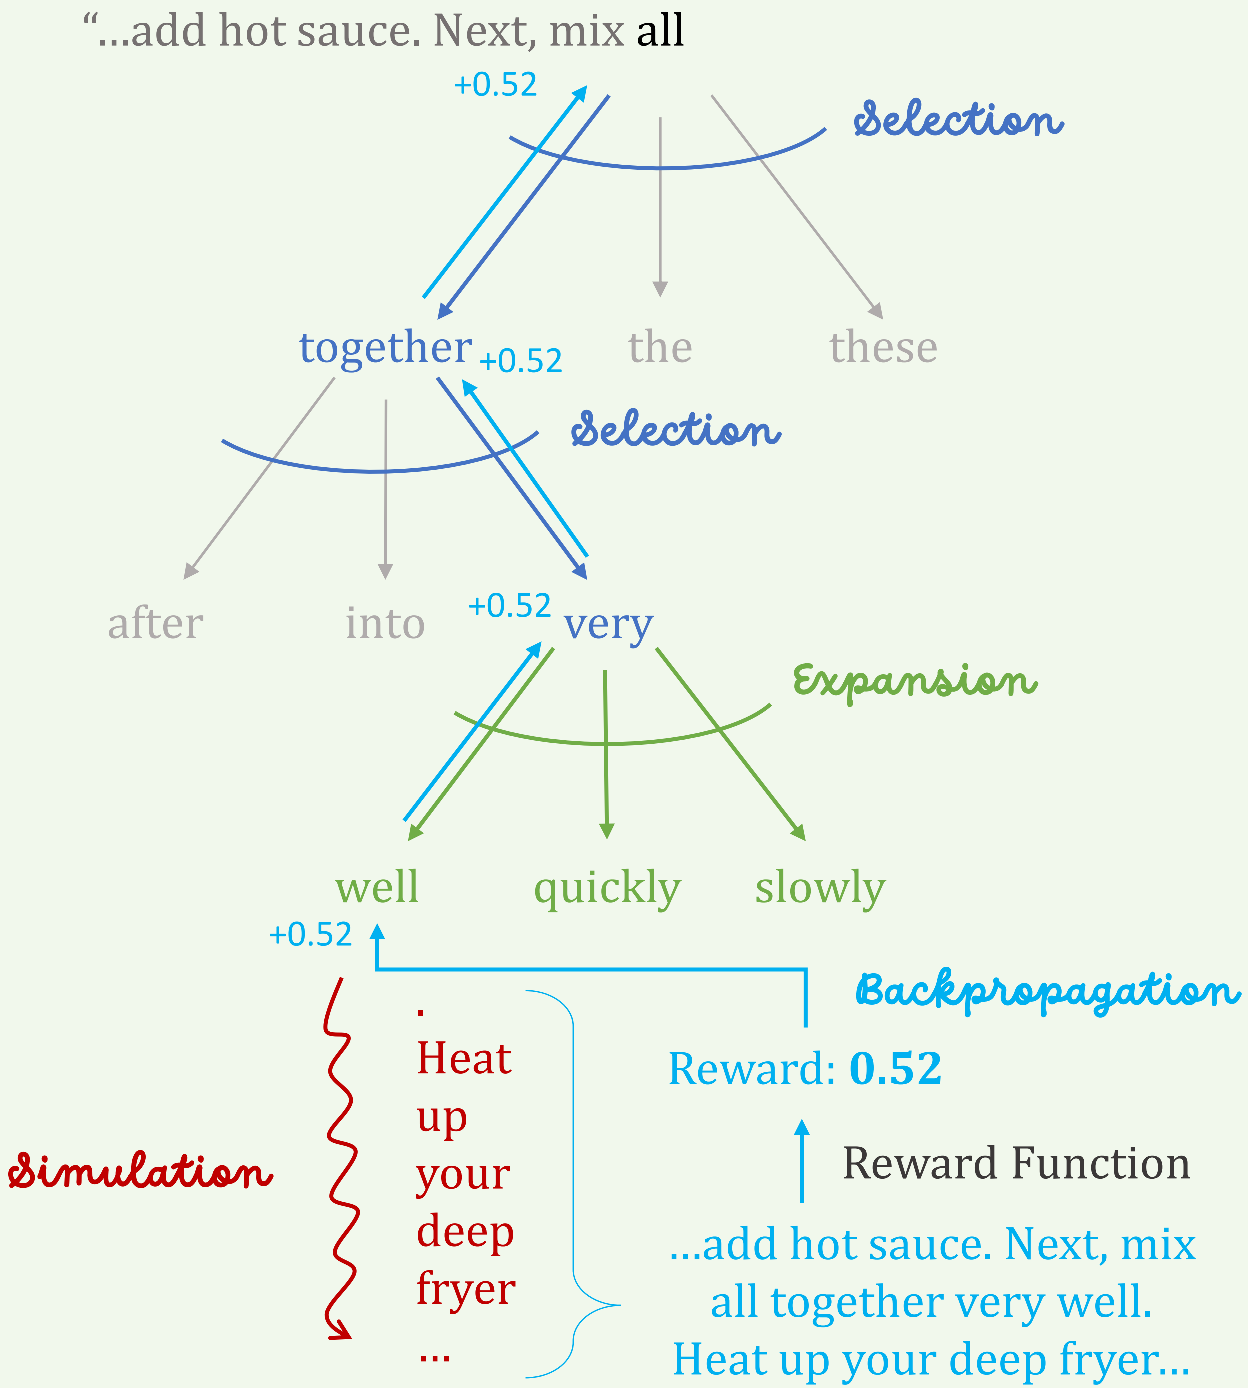

Monte Carlo Tree Search

Similar to games, we can look ahead and simulate recipe generation with the different possible next tokens.

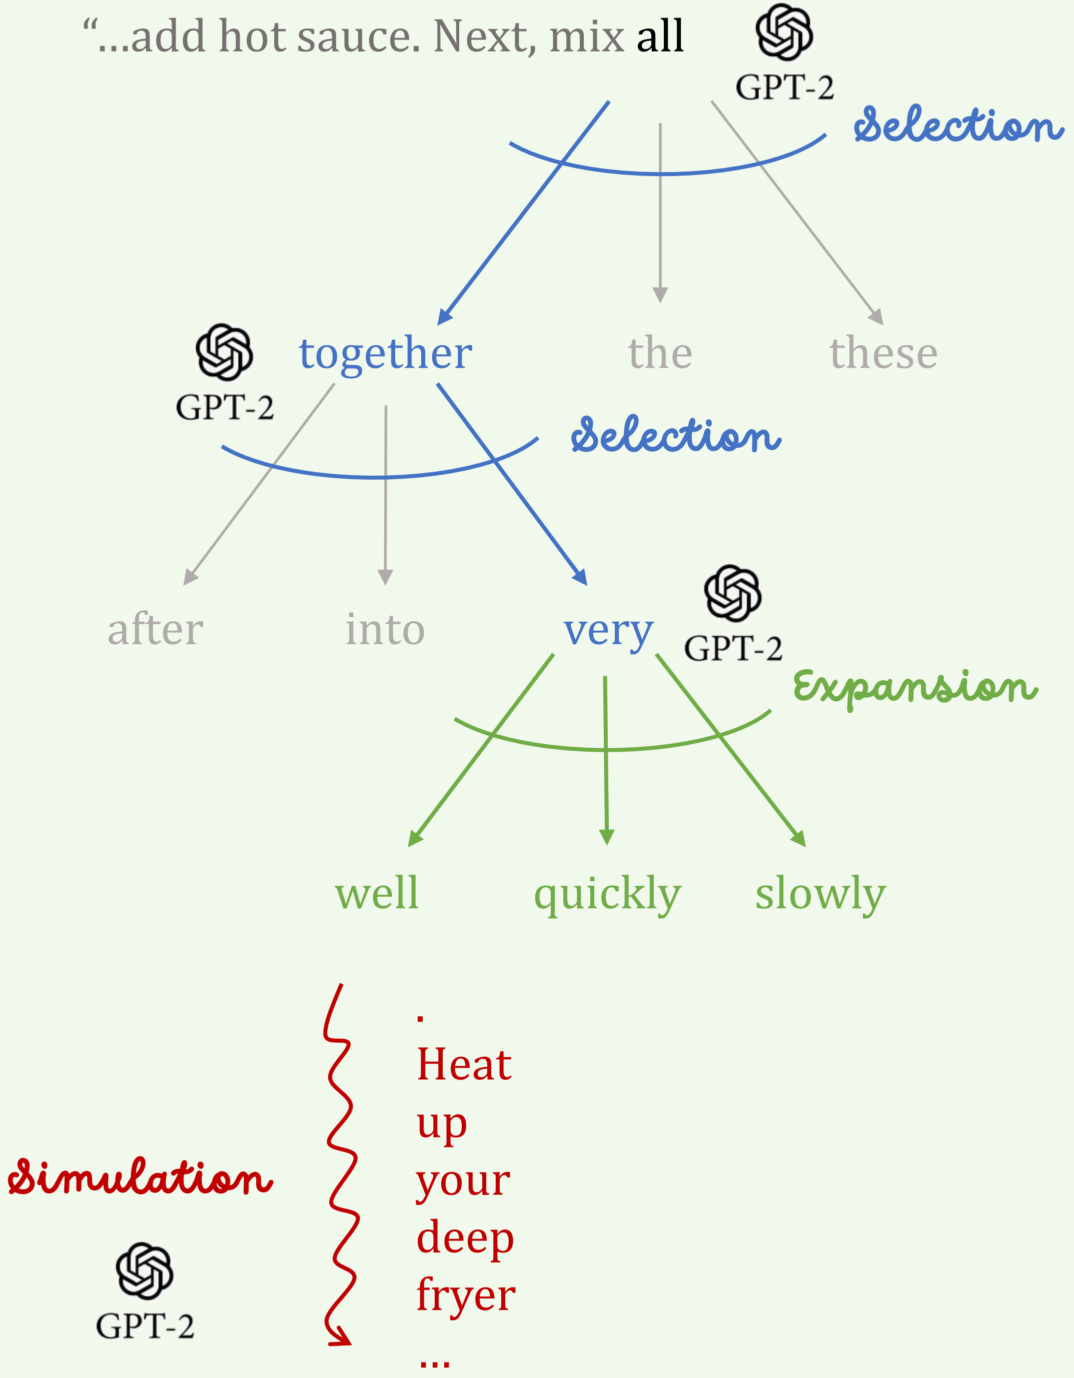

Monte Carlo Tree Search

With reward functions, we can guide the text generation to follow our constraints and preferences.

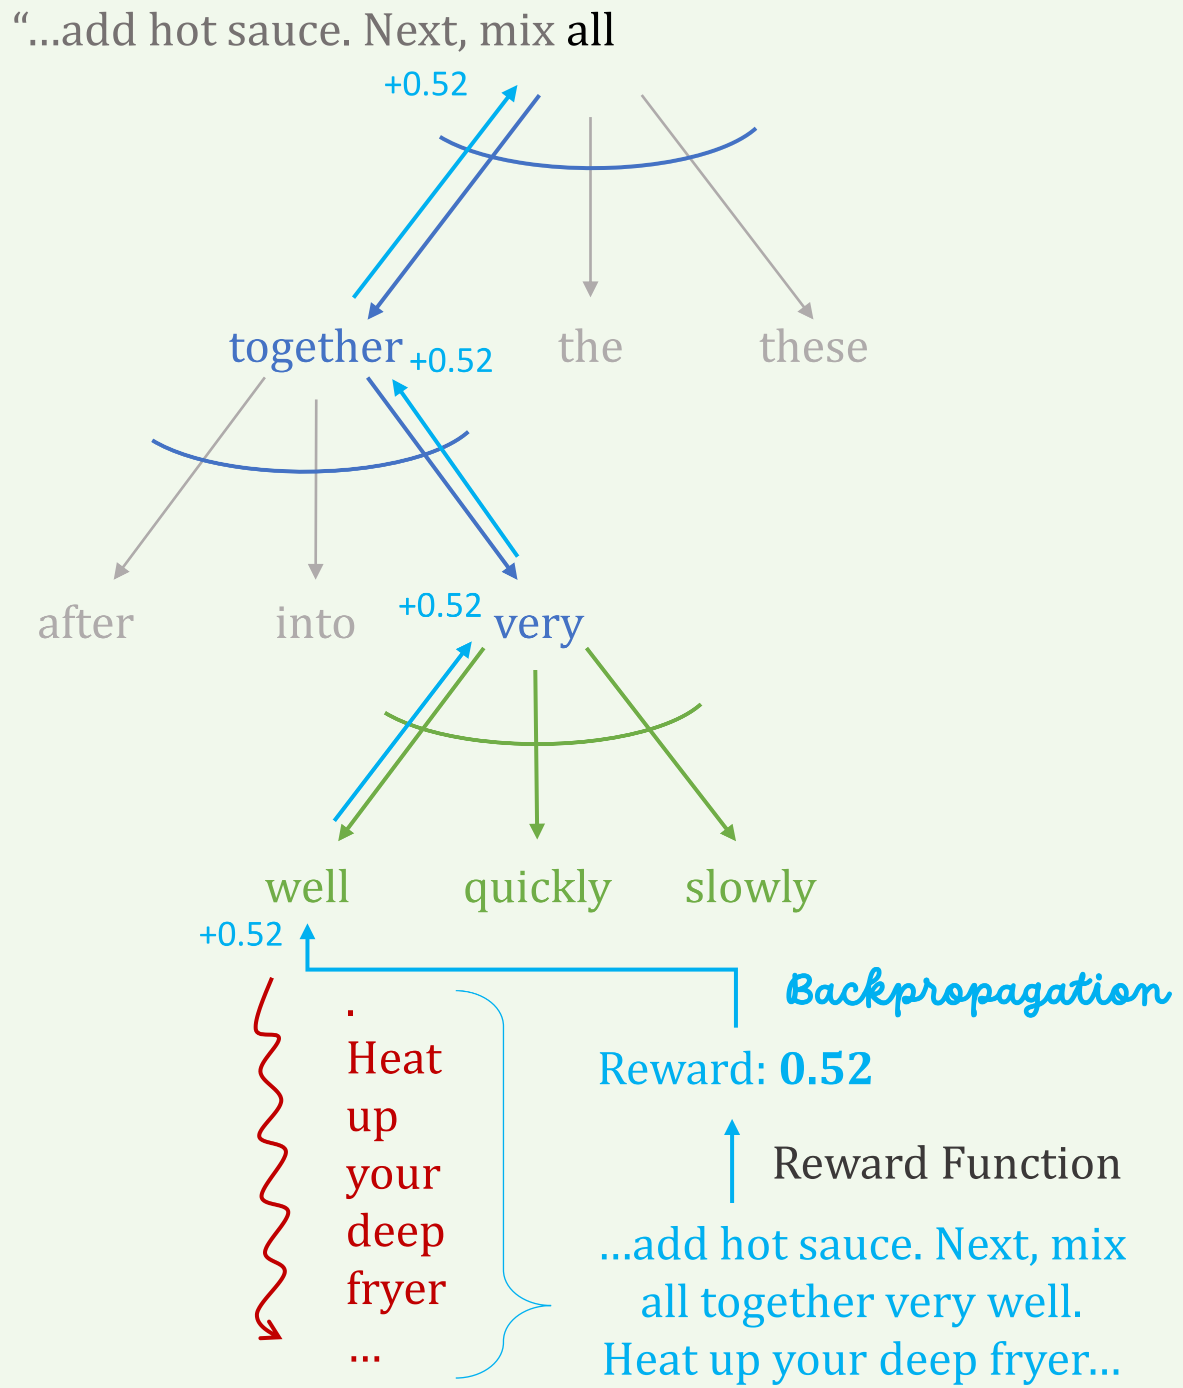

Reward Functions for Ingredients List

30%	45%	25%
Coherence between name and ingredients	Ingredients repetition penalty	Closing ingredients list

Reward Functions for Instructions List

50%	20%	30%
Coherence between ingredients and instructions	Special characters repetition penalty	Closing instructions list

Pros and Cons of MCTS for Recipe Generation

Pros

Controllable generation using reward functions that impose soft constraints.

No additional training of model or reward function after domain-specific fine-tuning of LLMs or prompting in newer LLMs.

Can be wrapped over an API that exposes the next token probabilities.

Cons

Computationally expensive compared to generating text in a single pass.

Baselines and Our Method

Top-p sampling (or Nucleus sampling)

Use only top n words, where n is smallest number such that the probability of top n words adds to $\geq p$.

Top-p Sampling with Repetition Penalty

Exponential penalty on repeating tokens.

Top-p Sampling with No n-gram Repetitions

Strictly prohibiting repeating sequences of n tokens ($n = 4$).

RecipeMC

Our method that uses fine-tuned GPT-2 + MCTS, Reward Functions

Automatic Evaluation

Name → Ingredients

Sampling Method	Coherence	F_1-Score	Perplexity↓	ROUGE-1	ROUGE-2	BLEU	Repetition↓
Ground Truth	0.451	-	2.934	-	-	-	0.667
Top-p	0.443	0.572	4.173	0.457	0.200	0.155	1.724
+ No 4-gram Repetition	0.444	0.562	5.150	0.456	0.198	0.144	1.641
+ Repetition Penalty	0.413	0.548	6.754	0.407	0.135	0.115	0.711
RecipeMC	**0.513**	**0.597**	**3.961**	**0.505**	**0.242**	**0.210**	**0.192**

Name, Ingredients → Instructions

Sampling Method	Coherence	Perplexity↓	ROUGE-1	ROUGE-2	BLEU
Ground Truth	0.486	4.115	-	-	-
Top-p	0.709	7.948	0.338	0.102	0.067
+ No 4-gram Repetition	0.690	8.441	0.339	0.103	0.069
+ Repetition Penalty	0.416	11.680	0.301	0.072	0.044
RecipeMC	**0.768**	**7.337**	**0.362**	**0.115**	**0.080**

Automatic Evaluation

Coherence ↑
Repetition ↓

Name → Ingredients

Sampling Method	**Coherence**	**F_1-Score**	**Perplexity↓**	**ROUGE-1**	**ROUGE-2**	**BLEU**	**Repetition↓**
Ground Truth	0.451	-	2.934	-	-	-	0.667
Top-p	0.443	0.572	4.173	0.457	0.200	0.155	1.724
+ No 4-gram Repetition	0.444	0.562	5.150	0.456	0.198	0.144	1.641
+ Repetition Penalty	0.413	0.548	6.754	0.407	0.135	0.115	0.711
RecipeMC	**0.513**	**0.597**	**3.961**	**0.505**	**0.242**	**0.210**	**0.192**

Name, Ingredients → Instructions

Sampling Method	**Coherence**	**Perplexity↓**	**ROUGE-1**	**ROUGE-2**	**BLEU**
Ground Truth	0.486	4.115	-	-	-
Top-p	0.709	7.948	0.338	0.102	0.067
+ No 4-gram Repetition	0.690	8.441	0.339	0.103	0.069
+ Repetition Penalty	0.416	11.680	0.301	0.072	0.044
RecipeMC	**0.768**	**7.337**	**0.362**	**0.115**	**0.080**

Automatic Evaluation

Perplexity↓

$$e^{-p(x_1,\ldots,x_n)}$$

Name → Ingredients

Sampling Method	**Coherence**	**F_1-Score**	**Perplexity↓**	**ROUGE-1**	**ROUGE-2**	**BLEU**	**Repetition↓**
Ground Truth	0.451	-	2.934	-	-	-	0.667
Top-p	0.443	0.572	4.173	0.457	0.200	0.155	1.724
+ No 4-gram Repetition	0.444	0.562	5.150	0.456	0.198	0.144	1.641
+ Repetition Penalty	0.413	0.548	6.754	0.407	0.135	0.115	0.711
RecipeMC	**0.513**	**0.597**	**3.961**	**0.505**	**0.242**	**0.210**	**0.192**

Name, Ingredients → Instructions

Sampling Method	**Coherence**	**Perplexity↓**	**ROUGE-1**	**ROUGE-2**	**BLEU**
Ground Truth	0.486	4.115	-	-	-
Top-p	0.709	7.948	0.338	0.102	0.067
+ No 4-gram Repetition	0.690	8.441	0.339	0.103	0.069
+ Repetition Penalty	0.416	11.680	0.301	0.072	0.044
RecipeMC	**0.768**	**7.337**	**0.362**	**0.115**	**0.080**

Automatic Evaluation

F1-Score ↑
ROUGE, BLEU ↑

Name → Ingredients

Sampling Method	**Coherence**	F_1**-Score**	**Perplexity**↓	**ROUGE-1**	**ROUGE-2**	**BLEU**	**Repetition**↓
Ground Truth	0.451	-	2.934	-	-	-	0.667
Top-p	0.443	0.572	4.173	0.457	0.200	0.155	1.724
+ No 4-gram Repetition	0.444	0.562	5.150	0.456	0.198	0.144	1.641
+ Repetition Penalty	0.413	0.548	6.754	0.407	0.135	0.115	0.711
RecipeMC	**0.513**	**0.597**	**3.961**	**0.505**	**0.242**	**0.210**	**0.192**

Name, Ingredients → Instructions

Sampling Method	**Coherence**	**Perplexity**↓	**ROUGE-1**	**ROUGE-2**	**BLEU**
Ground Truth	0.486	4.115	-	-	-
Top-p	0.709	7.948	0.338	0.102	0.067
+ No 4-gram Repetition	0.690	8.441	0.339	0.103	0.069
+ Repetition Penalty	0.416	11.680	0.301	0.072	0.044
RecipeMC	**0.768**	**7.337**	**0.362**	**0.115**	**0.080**

Human Evaluation

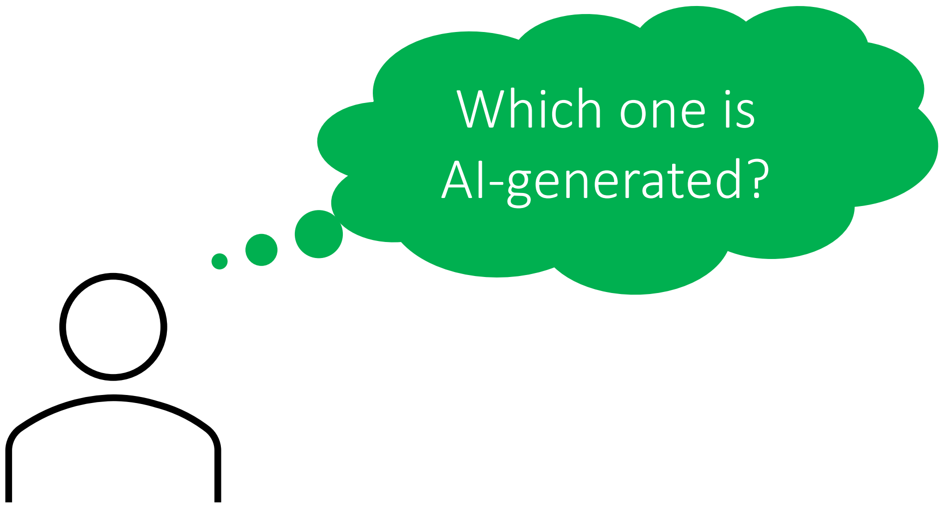

Recipe Name

- Real Ingredients
- AI-generated Ingredients

Recipe Name + Ingredients

- Real Instructions
- AI-generated Instructions

Recipe: Vegetarian Cupcakes for Pups

Choose the fake (AI-generated) recipe:

1 teaspoon vanilla extract 1/2 cup butter, softened (8 tablespoons) 2 teaspoons baking powder 3 cups all-purpose flour 3 ounces semisweet chocolate or white baking bars, finely chopped 1/4 teaspoon salt 1/3 cup sour cream 4 egg whites 1/2 teaspoon baking soda 6 packets Sugar Substitute	3 cups water 2 bananas 1 teaspoon nutmeg 12 teaspoon vanilla 1 teaspoon baking powder 1 teaspoon cinnamon 2 tablespoons honey 4 cups whole wheat flour 1 egg 2 carrots, shredded

Recipe Name

Real Ingredients | AI-generated Ingredients

Method	Real	Gen.	P(Incorrect)
Top-p	175	185	0.4861
+ No 4-gram Repetition	179	200	0.4723
+ Repetition Penalty	183	180	0.5041
RecipeMC	201	167	**0.5462**
Overall	738	732	0.5020

Recipe Name
+ Ingredients

Real Instructions | AI-generated Instructions

Method	Real	Gen.	P(Incorrect)
Top-p	51	42	0.5484
+ No 4-gram Repetition	67	62	0.5194
+ Repetition Penalty	36	65	0.3564
RecipeMC	57	35	**0.6196**
Overall	211	204	0.5084

Takeaway

Simple manually-defined reward functions can be easily used to guide text generation using Monte Carlo Tree Search…

- Without training a reward model.
- With any API that exposes next token probabilities.

Contact Information

Karan Taneja
karantaneja@gatech.edu
School of Interactive Computing
Georgia Institute of Technology

Richard Segal
rsegal@us.ibm.com
Computational Creativity Group
IBM Research

Richard Goodwin
rgoodwin@us.ibm.com
Computational Creativity Group
IBM Research

Appendix

GPT-2

↓

GPT-3,
GPT-3.5,
ChatGPT,
and other new LLMs

Monte Carlo Tree Search (MCTS)

Step 1: *Selection*

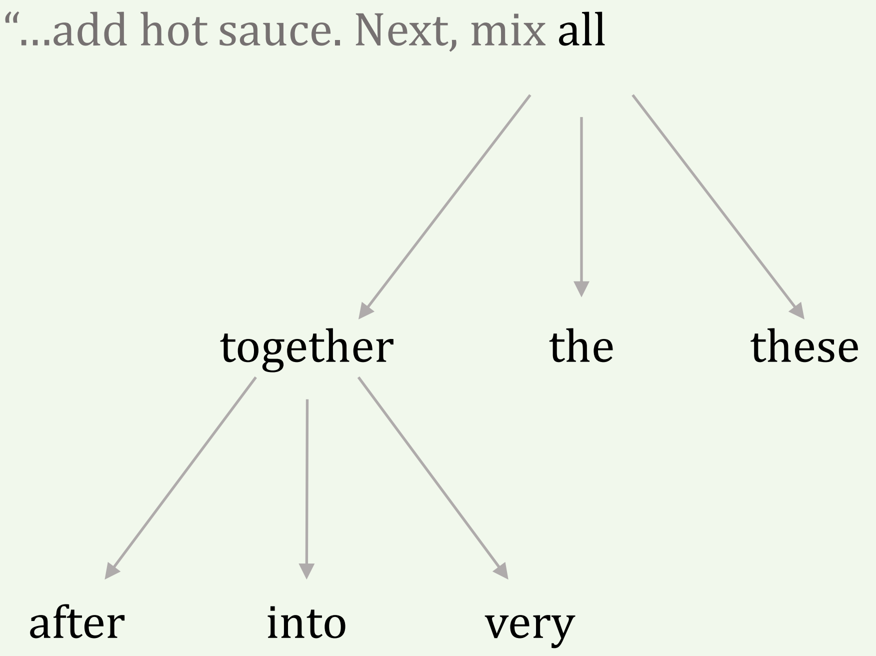

$$\mathrm{PUCB}(i) = \underbrace{Q(i)}_{\text{Exploitation}} + \underbrace{c}_{\text{Balancing Constant}} \cdot \underbrace{p(x_i|x_{1:i-1})\frac{\sqrt{N}}{n_i+1}}_{\text{Exploration}}$$

Monte Carlo Tree Search (MCTS)

Step 1: Selection

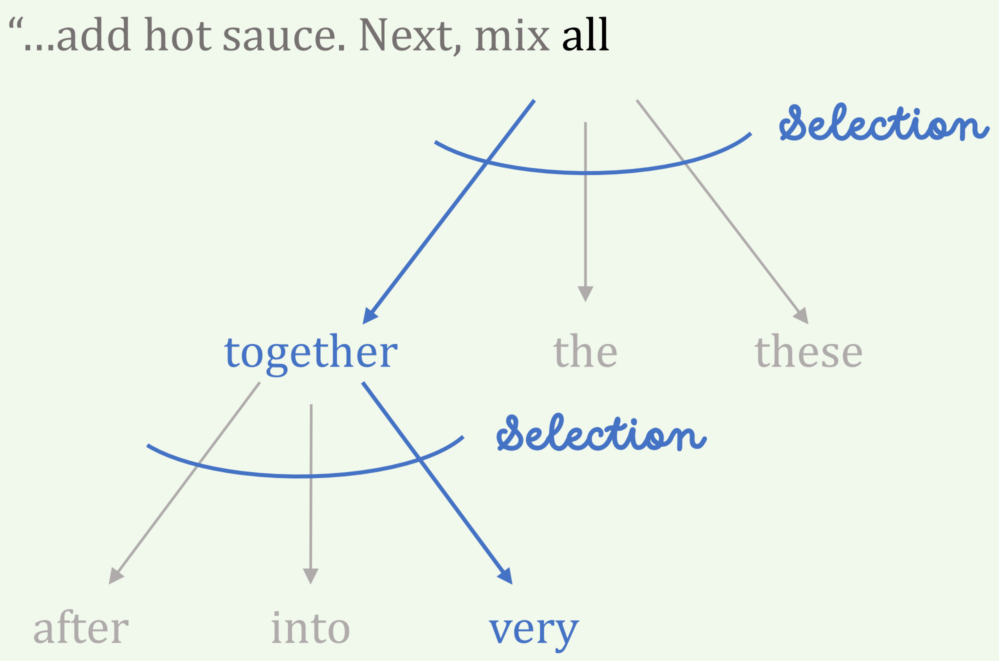

Monte Carlo Tree Search (MCTS)

Step 2: *Expansion*

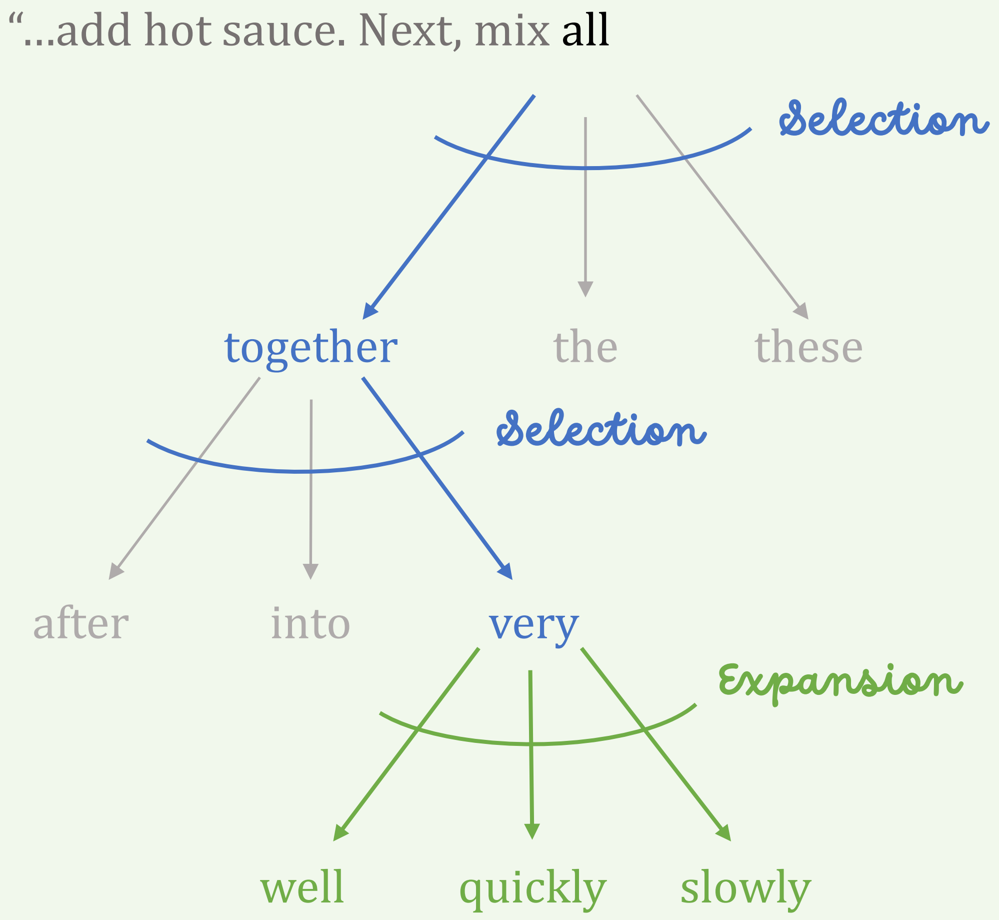

Monte Carlo Tree Search (MCTS)

Step 3: *Simulation*

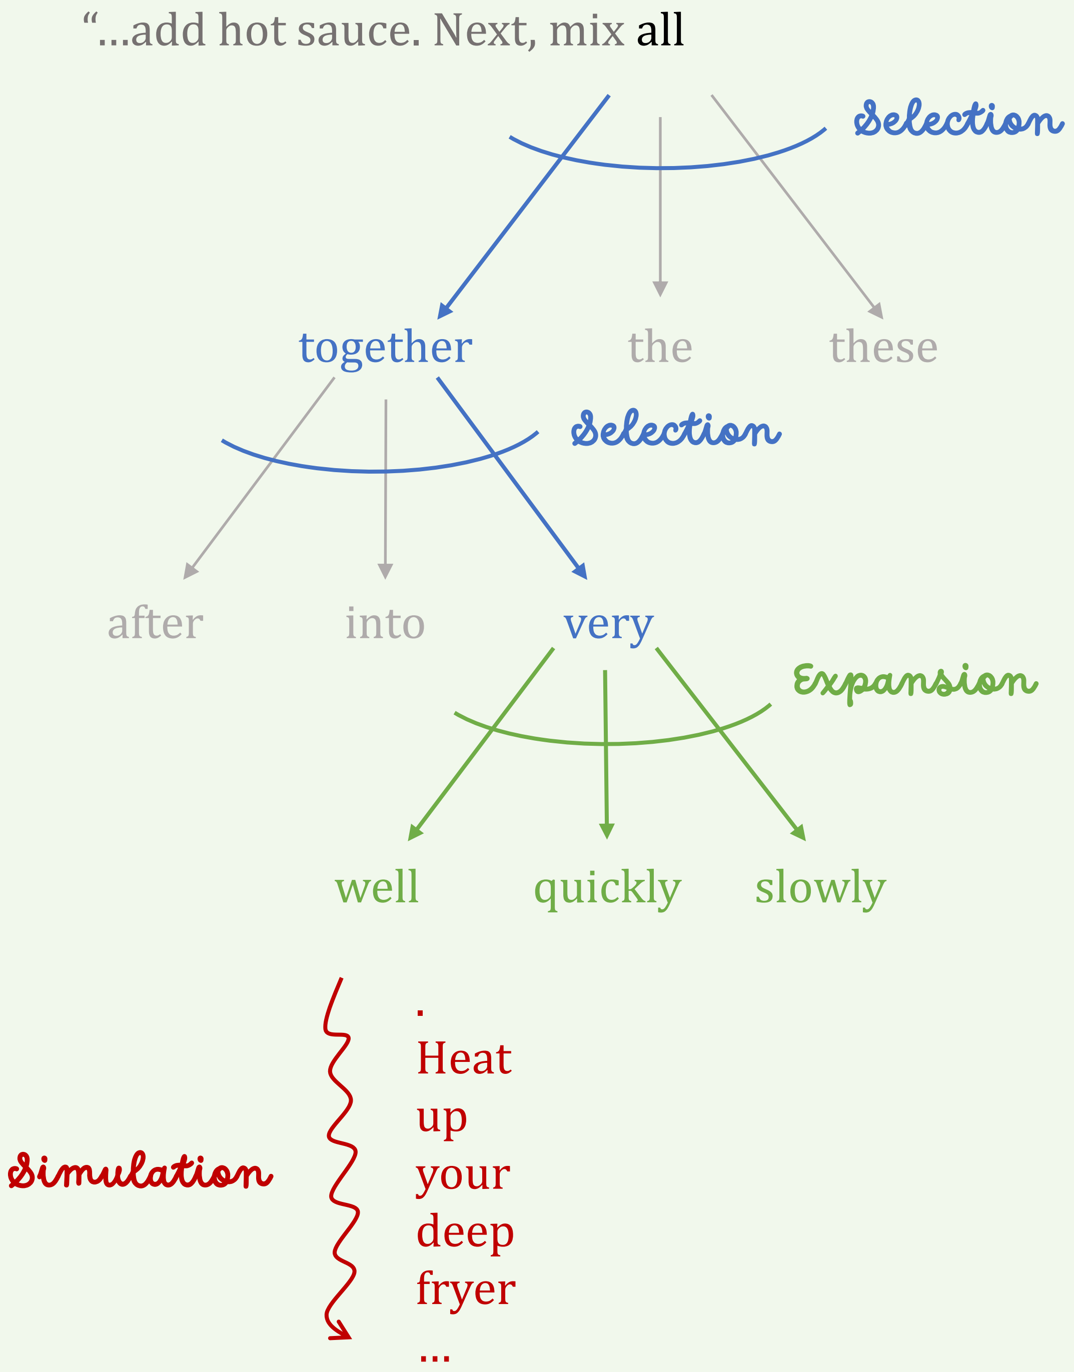

Monte Carlo Tree Search (MCTS)

Step 4: *Backpropagation*

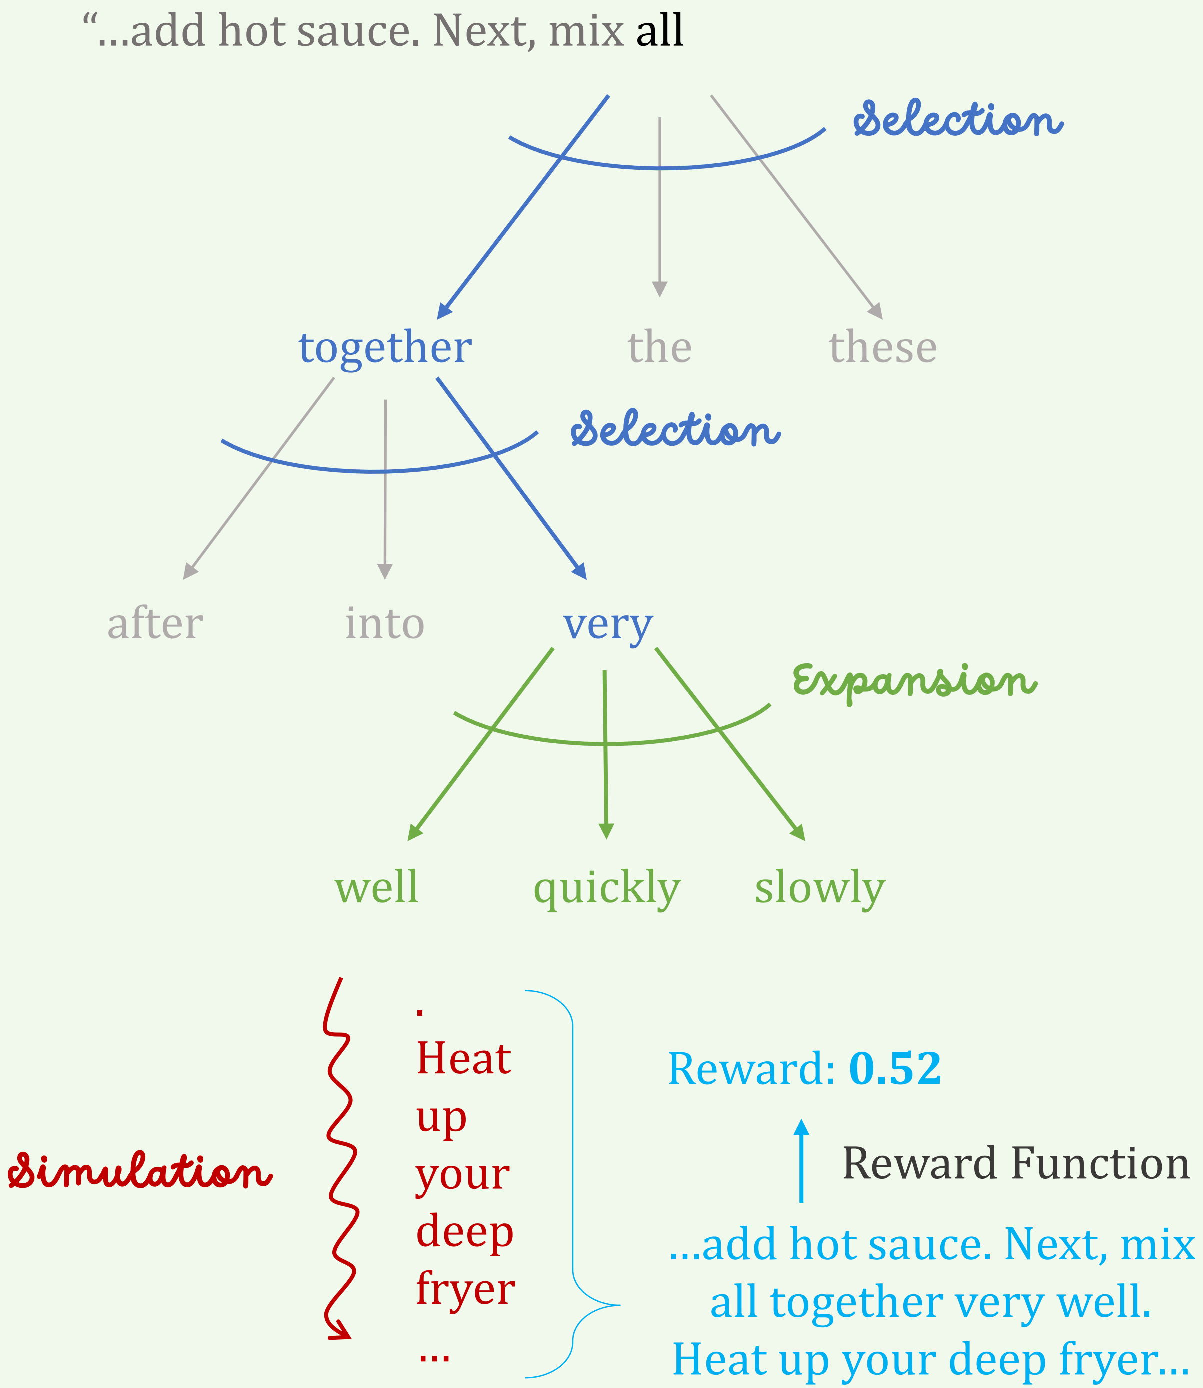

Monte Carlo Tree Search (MCTS)

Step 4: Backpropagation

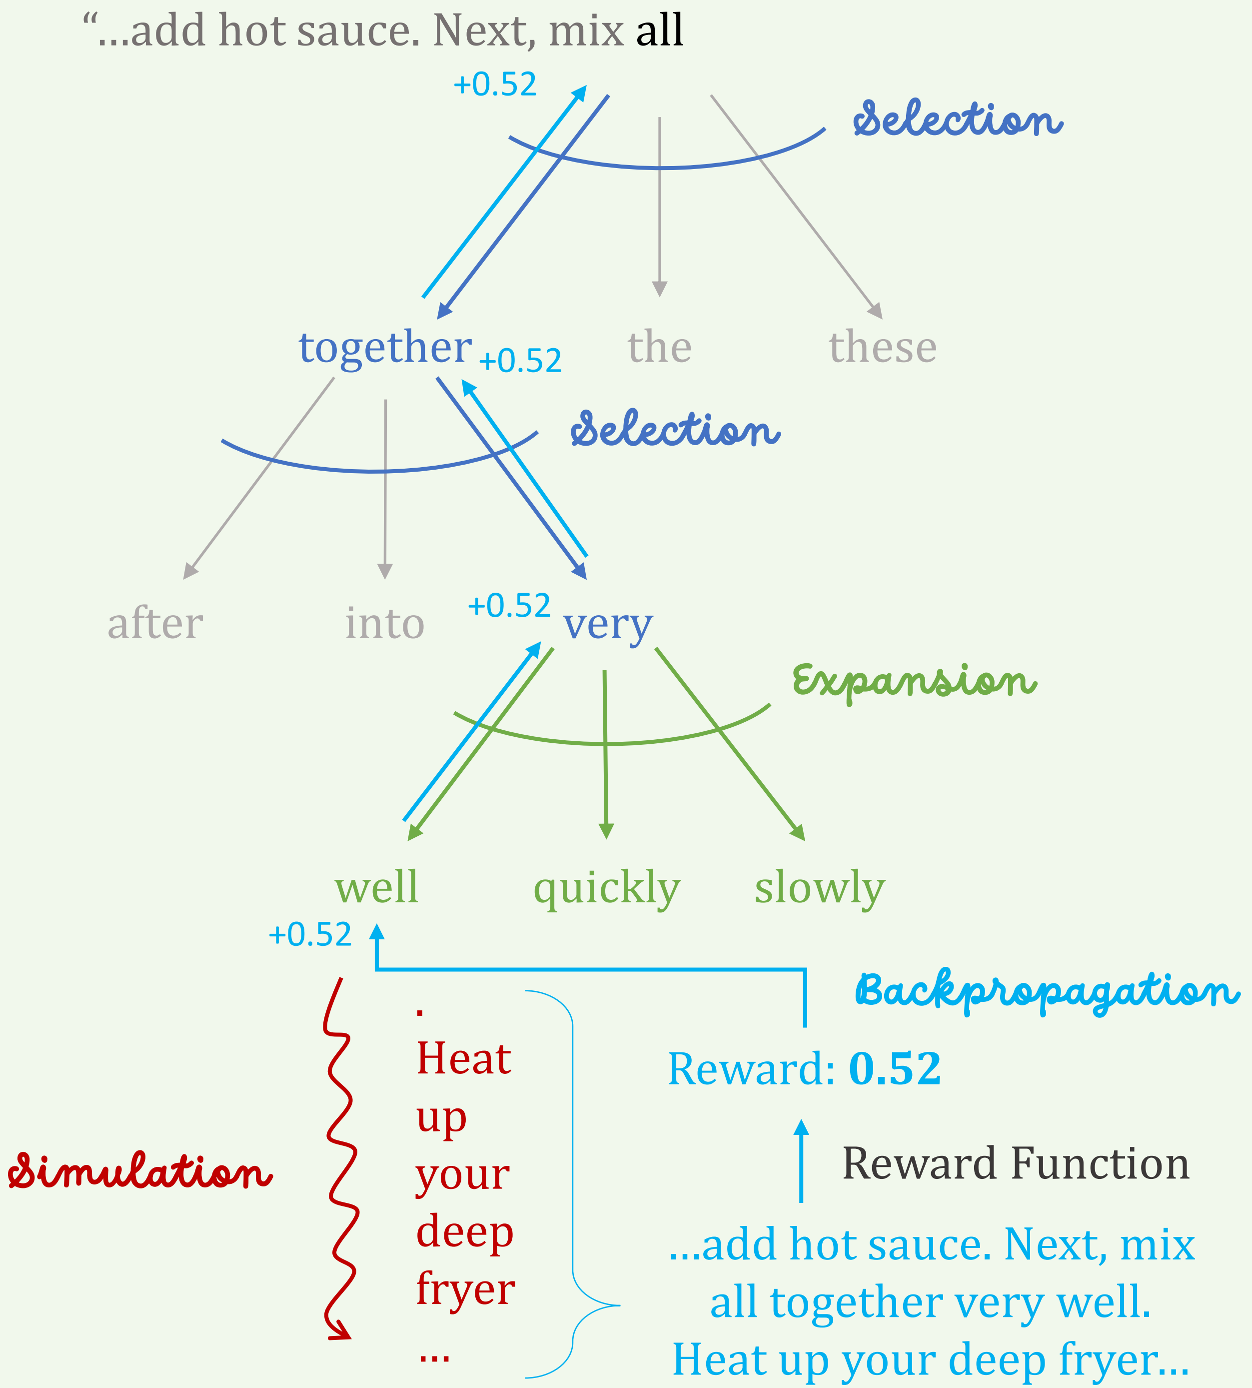

Average recipe lengths

Method	Ingredients	Instructions
Ground Truth	167	240
Top-p	247	485
+ No 4-gram Repetition	248	484
+ Repetition Penalty	233	545
RecipeMC	190	441