

Fine-Tuning Language Models Using Formal Methods Feedback: A Use Case in Autonomous Systems

Yunhao Yang*, Neel P. Bhatt*, Tyler Ingebrand*, William Ward, Steven Carr, Zhangyang Wang, and Ufuk Topcu

May 14th, 2024



autonomy.odan.utexas.edu

CENTER FOR
aUTonomy

* Equal Contribution

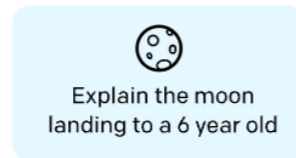
Reinforcement Learning via Human Feedback (RLHF)

Example: OpenAI Scheme for Instruct GPT

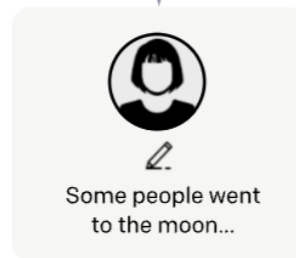
Step 1

Collect demonstration data, and train a supervised policy.

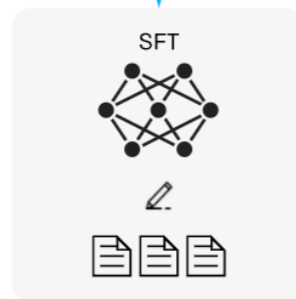
A prompt is sampled from our prompt dataset.



A labeler demonstrates the desired output behavior.



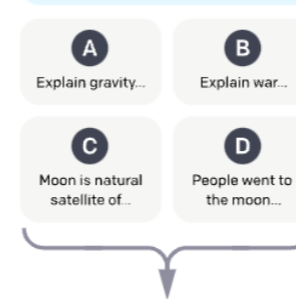
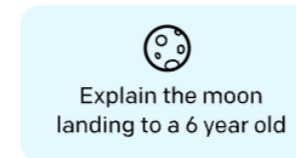
This data is used to fine-tune GPT-3 with supervised learning.



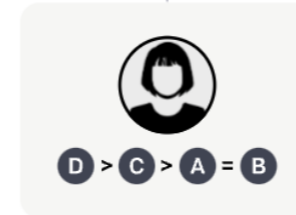
Step 2

Collect comparison data, and train a reward model.

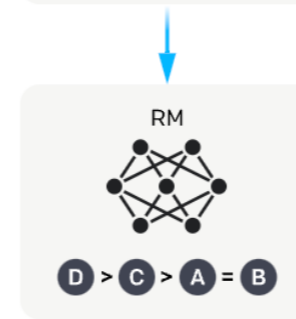
A prompt and several model outputs are sampled.



A labeler ranks the outputs from best to worst.



This data is used to train our reward model.



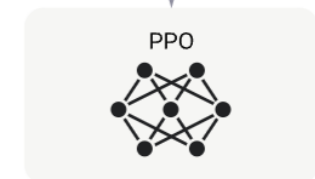
Step 3

Optimize a policy against the reward model using reinforcement learning.

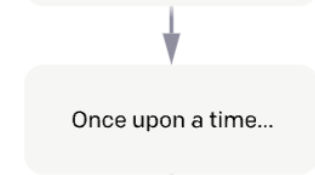
A new prompt is sampled from the dataset.



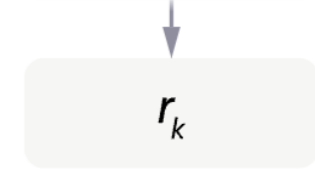
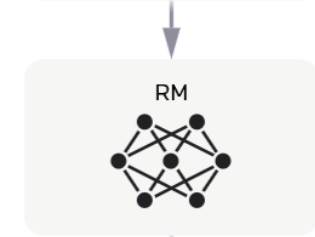
The policy generates an output.



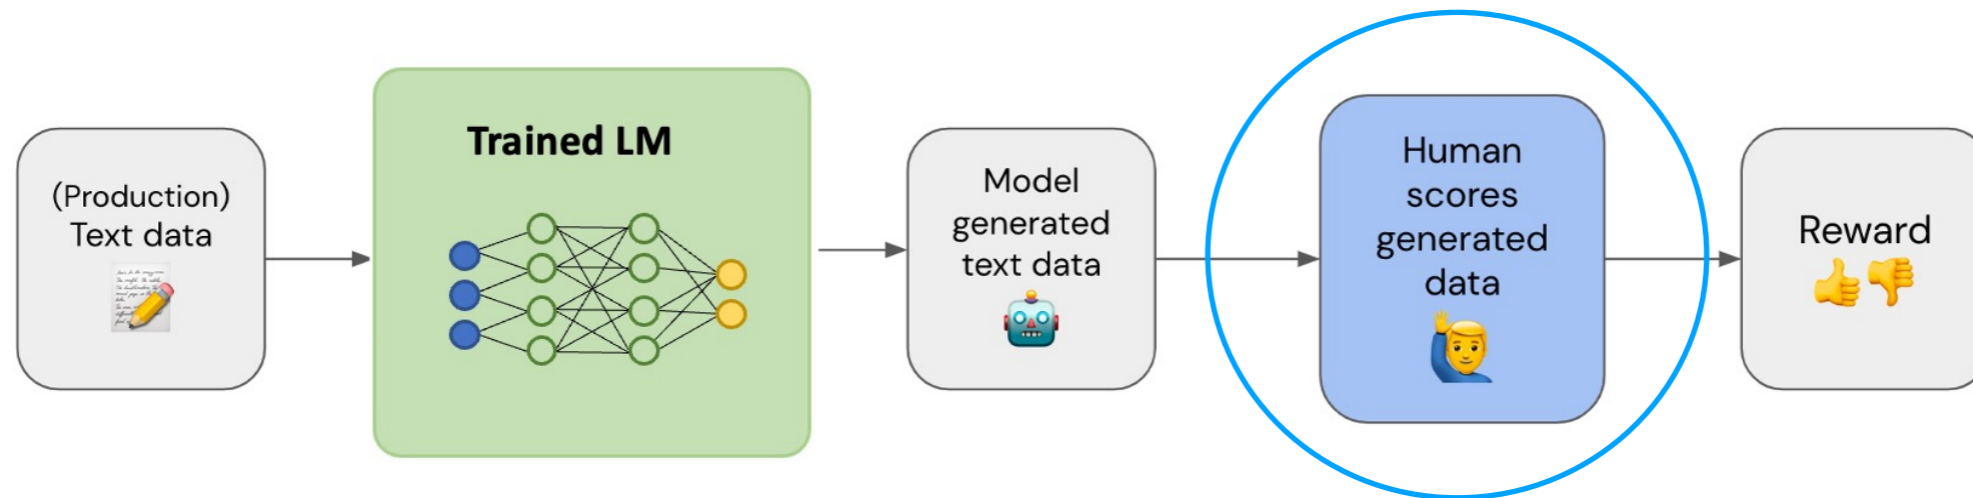
The reward model calculates a reward for the output.



The reward is used to update the policy using PPO.



Fine-Tuning Language Models Using Human Feedback

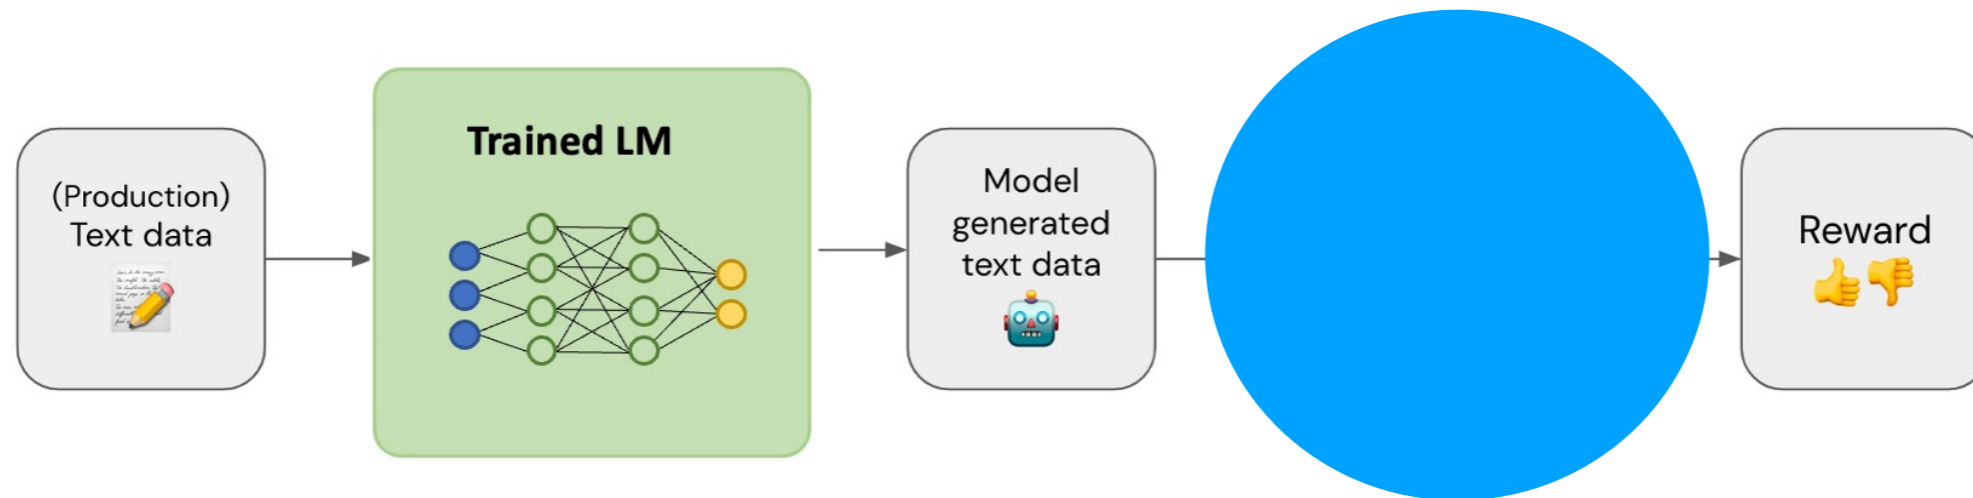


Labor-Intensive



Subjective/Inconsistent Feedback

Fine-Tuning Language Models Using Formal Feedback Methods

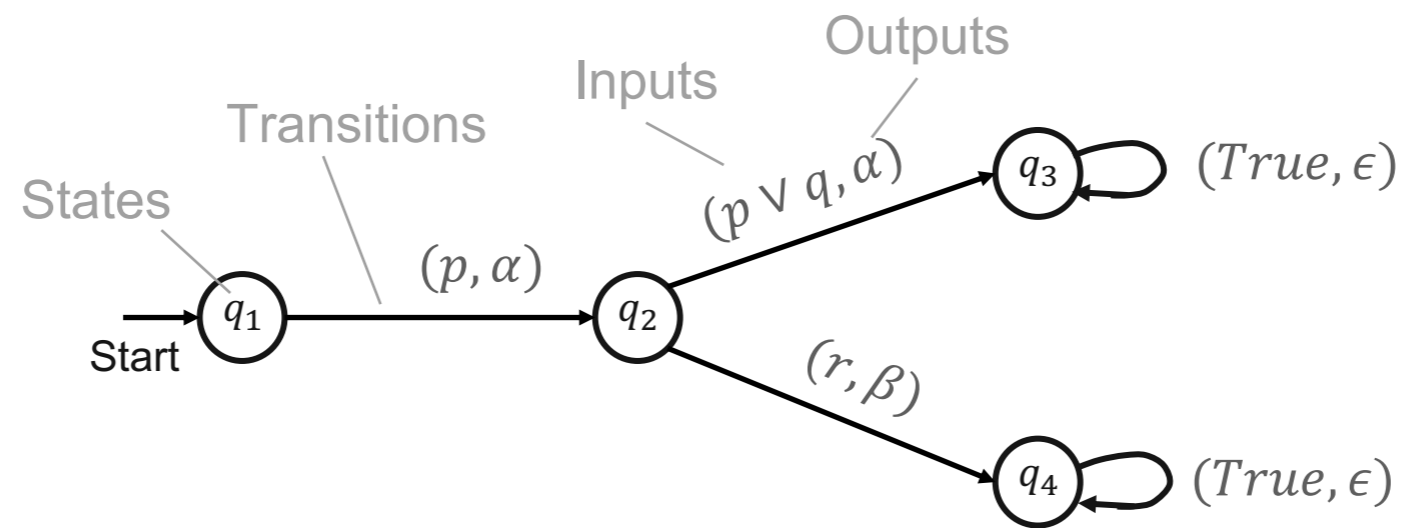


Formal Methods:

Automaton-Based Representation, Model Checking, Temporal Logic Specification, etc.



Background: Automaton-Based Representations

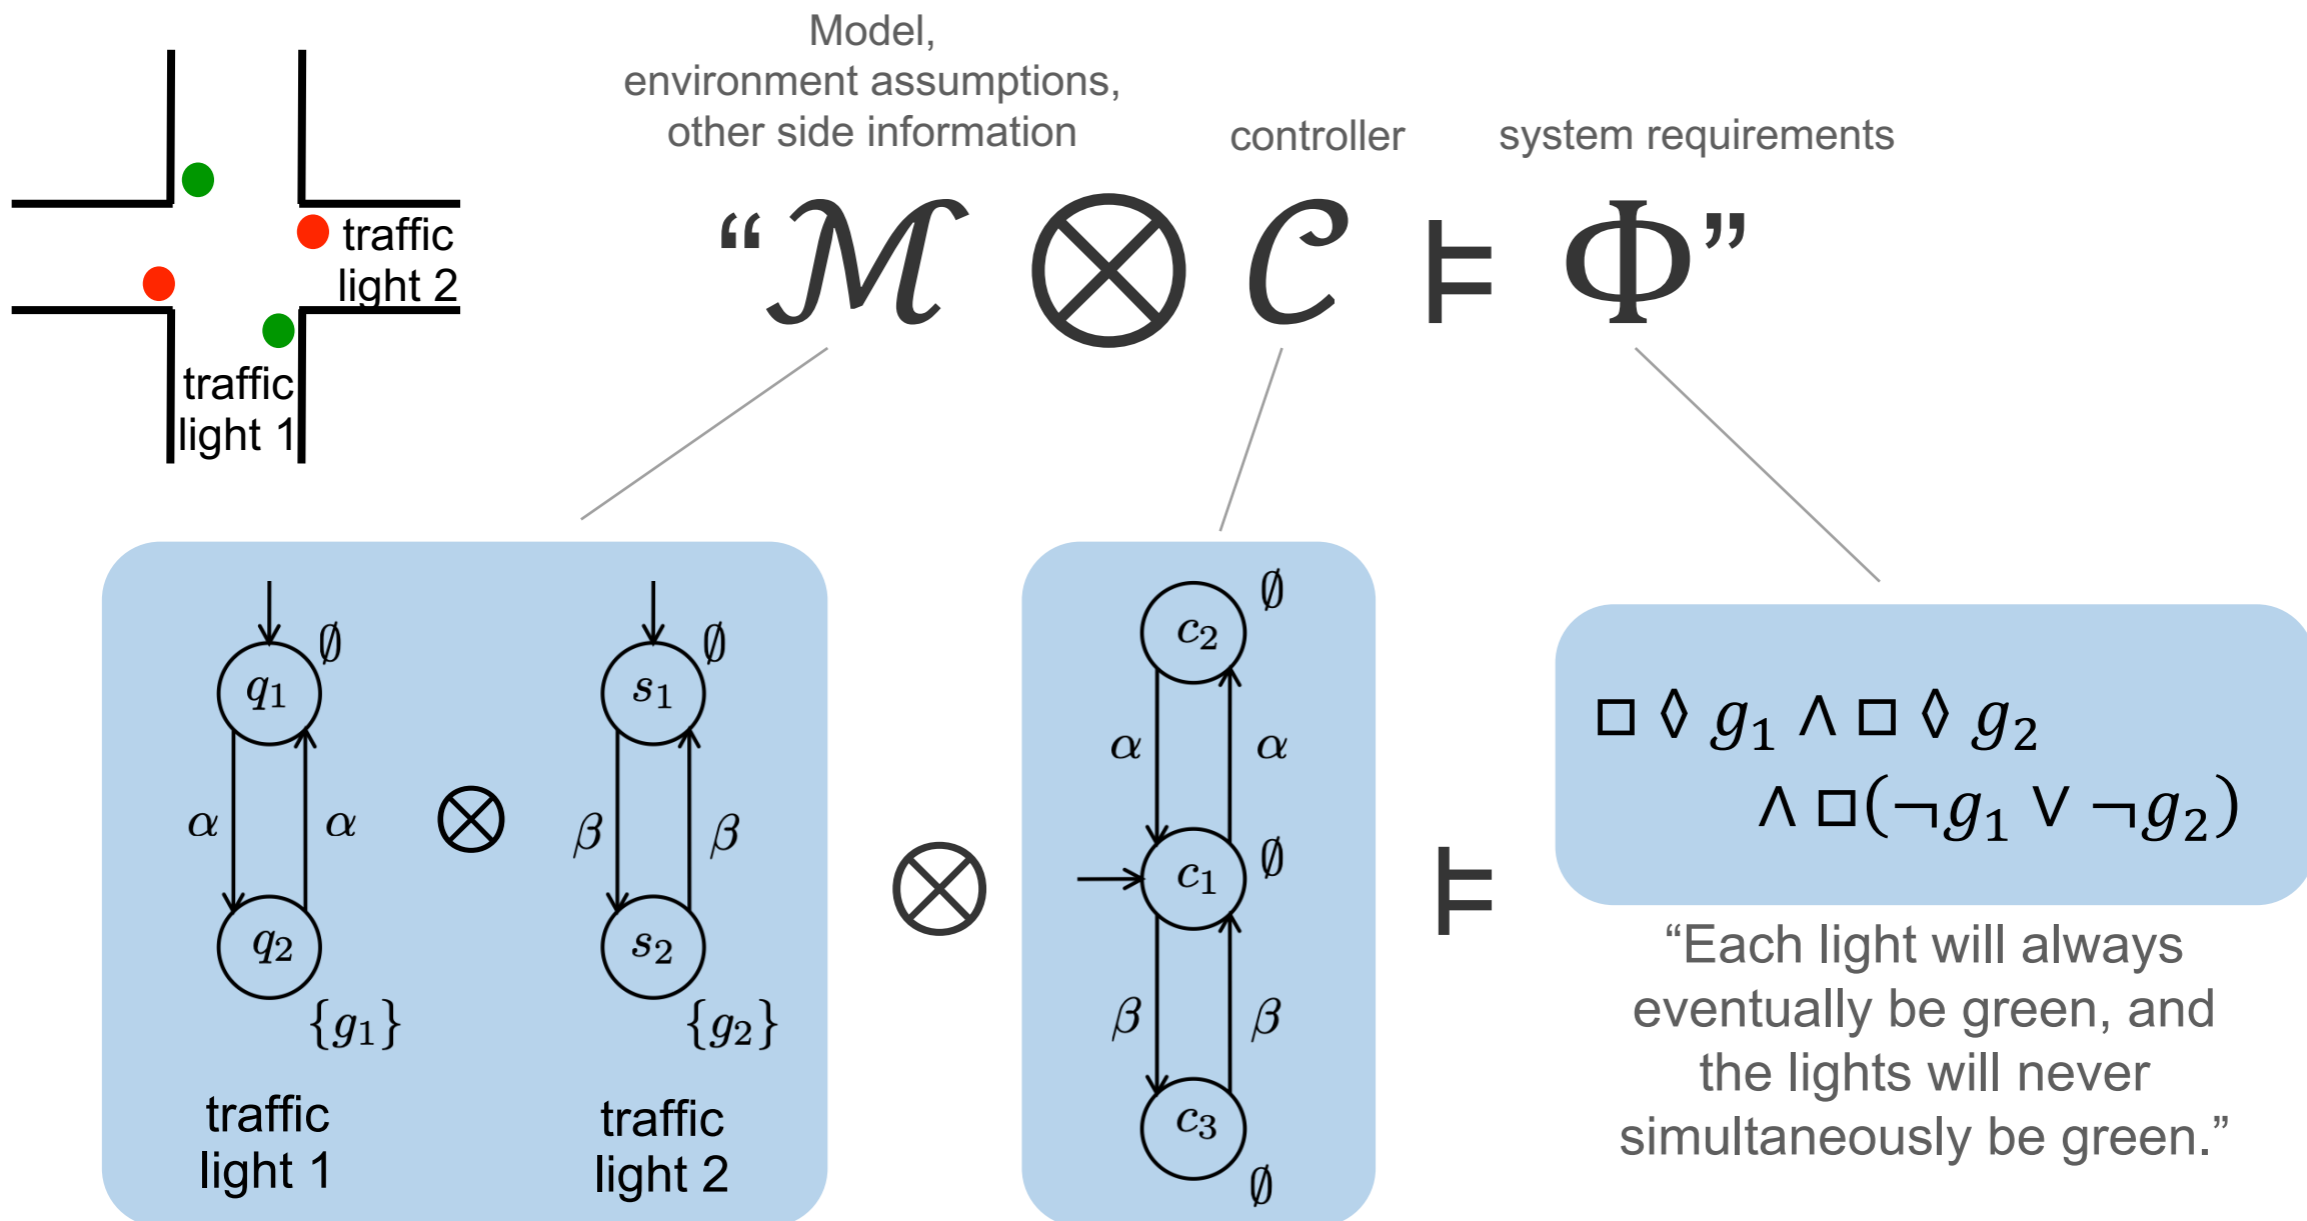


Why automaton-based representations? They are used for

- model checking, planning,...
- reactive synthesis, games on graphs, ...
- probabilistic verification and synthesis, and
- reinforcement learning.

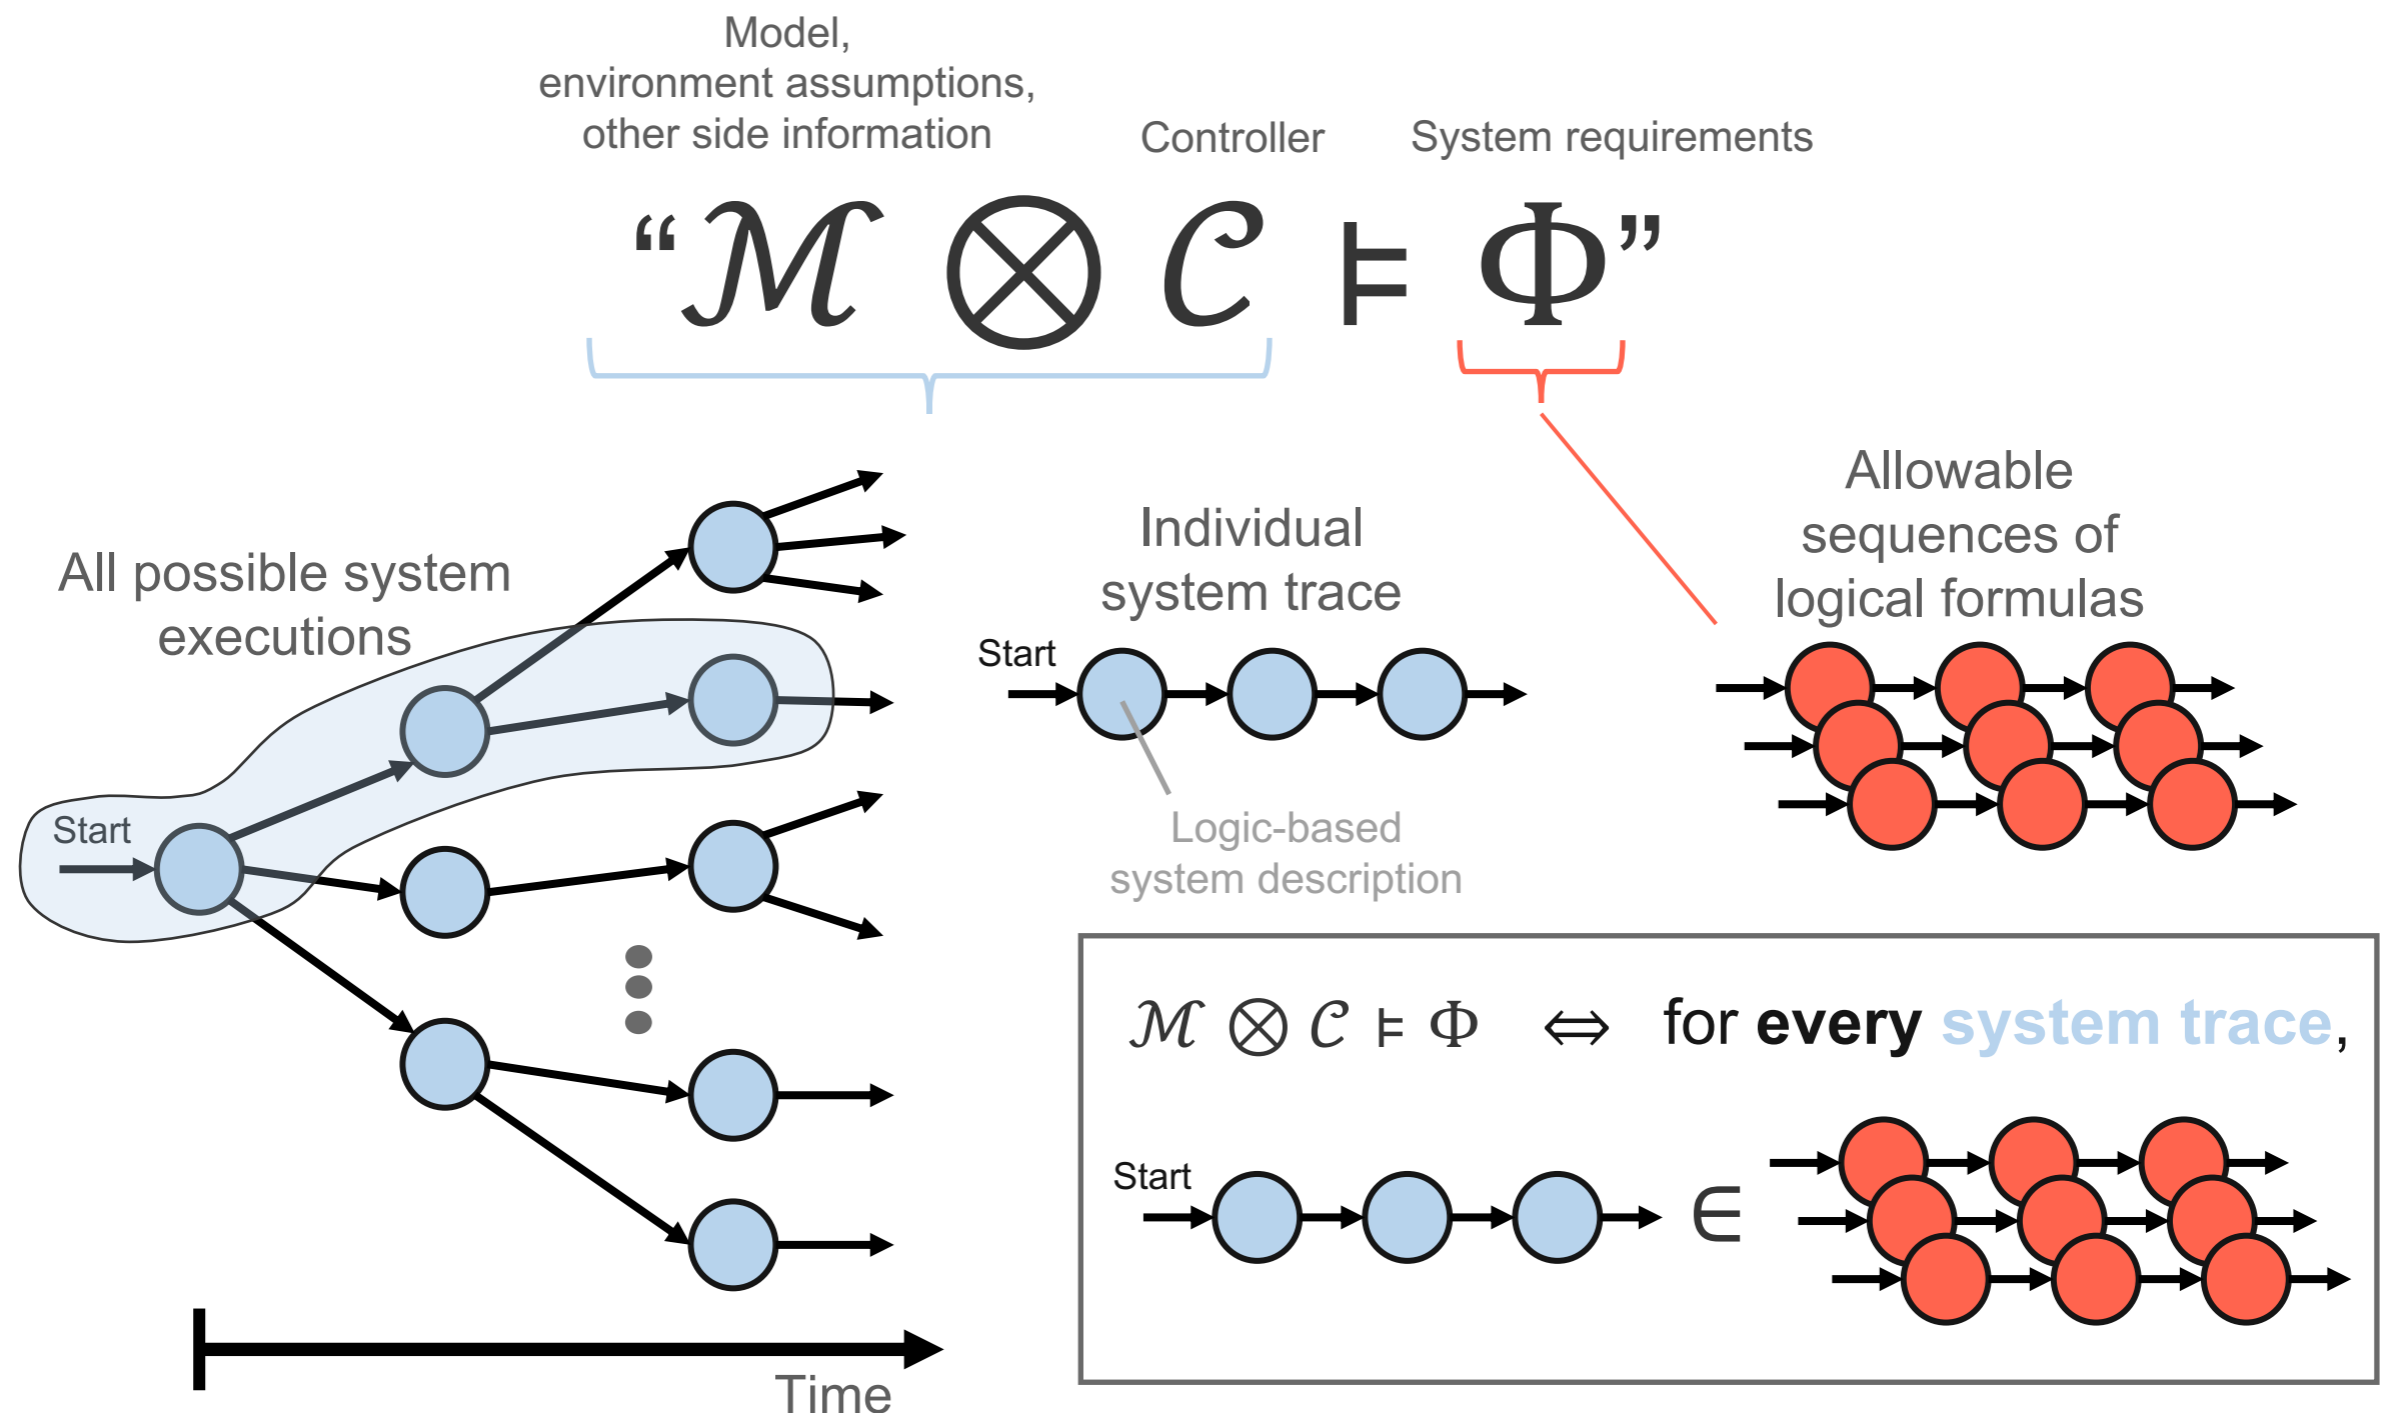
A (Very) Brief Introduction to Model Checking

Are the controller's outcomes **guaranteed** to satisfy user-specified requirements when implemented against a system model?



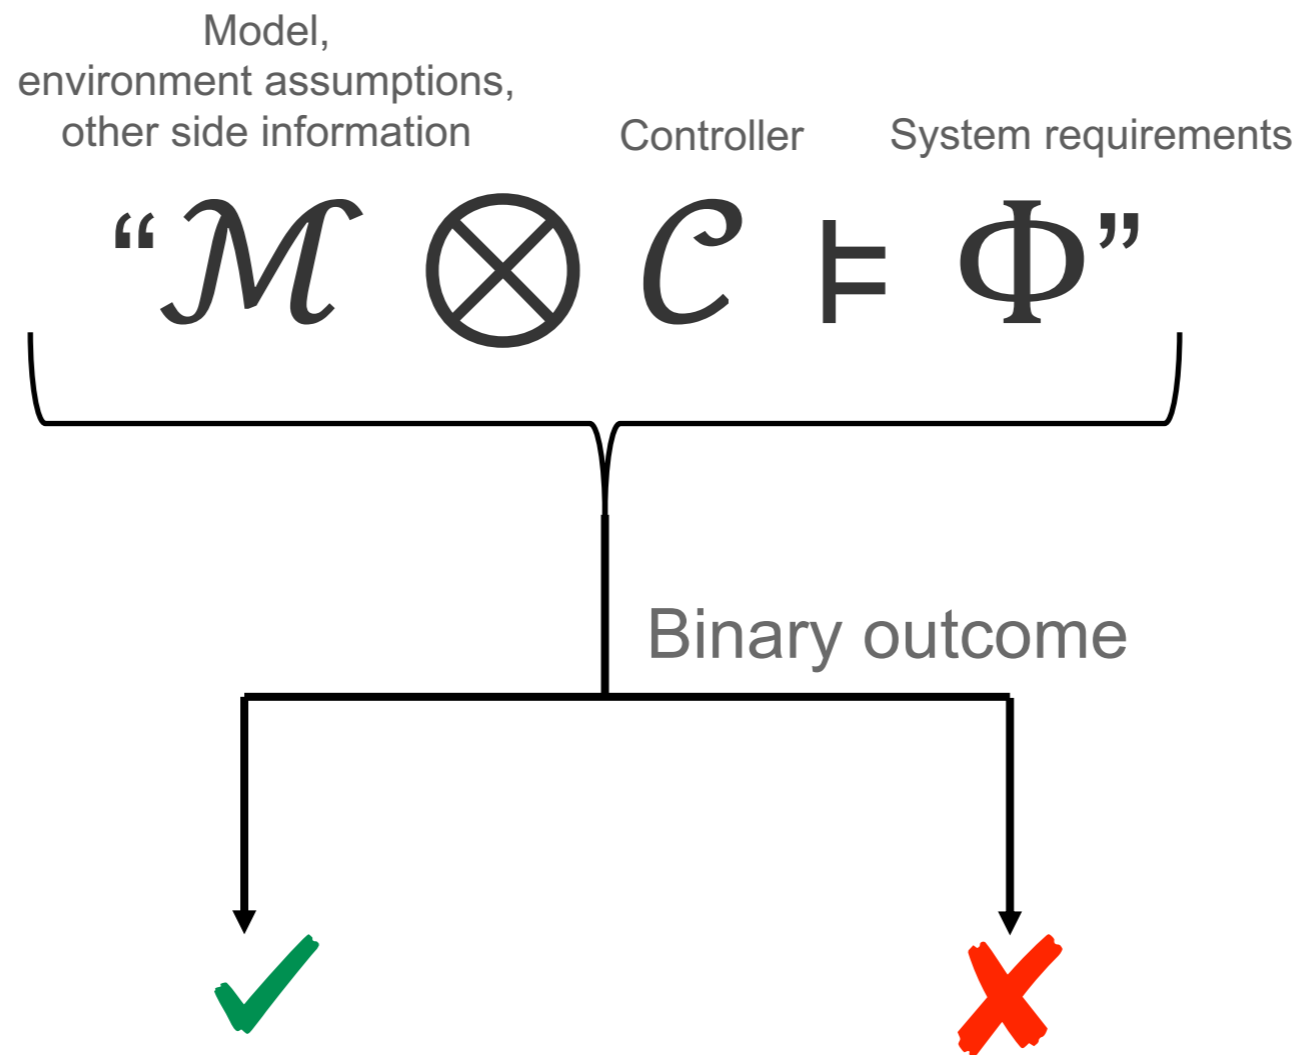
A (Very) Brief Introduction to Model Checking

Are the controller's outcomes **guaranteed** to satisfy user-specified requirements when implemented against a system model?

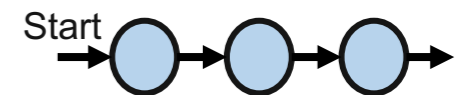


A (Very) Brief Introduction to Model Checking

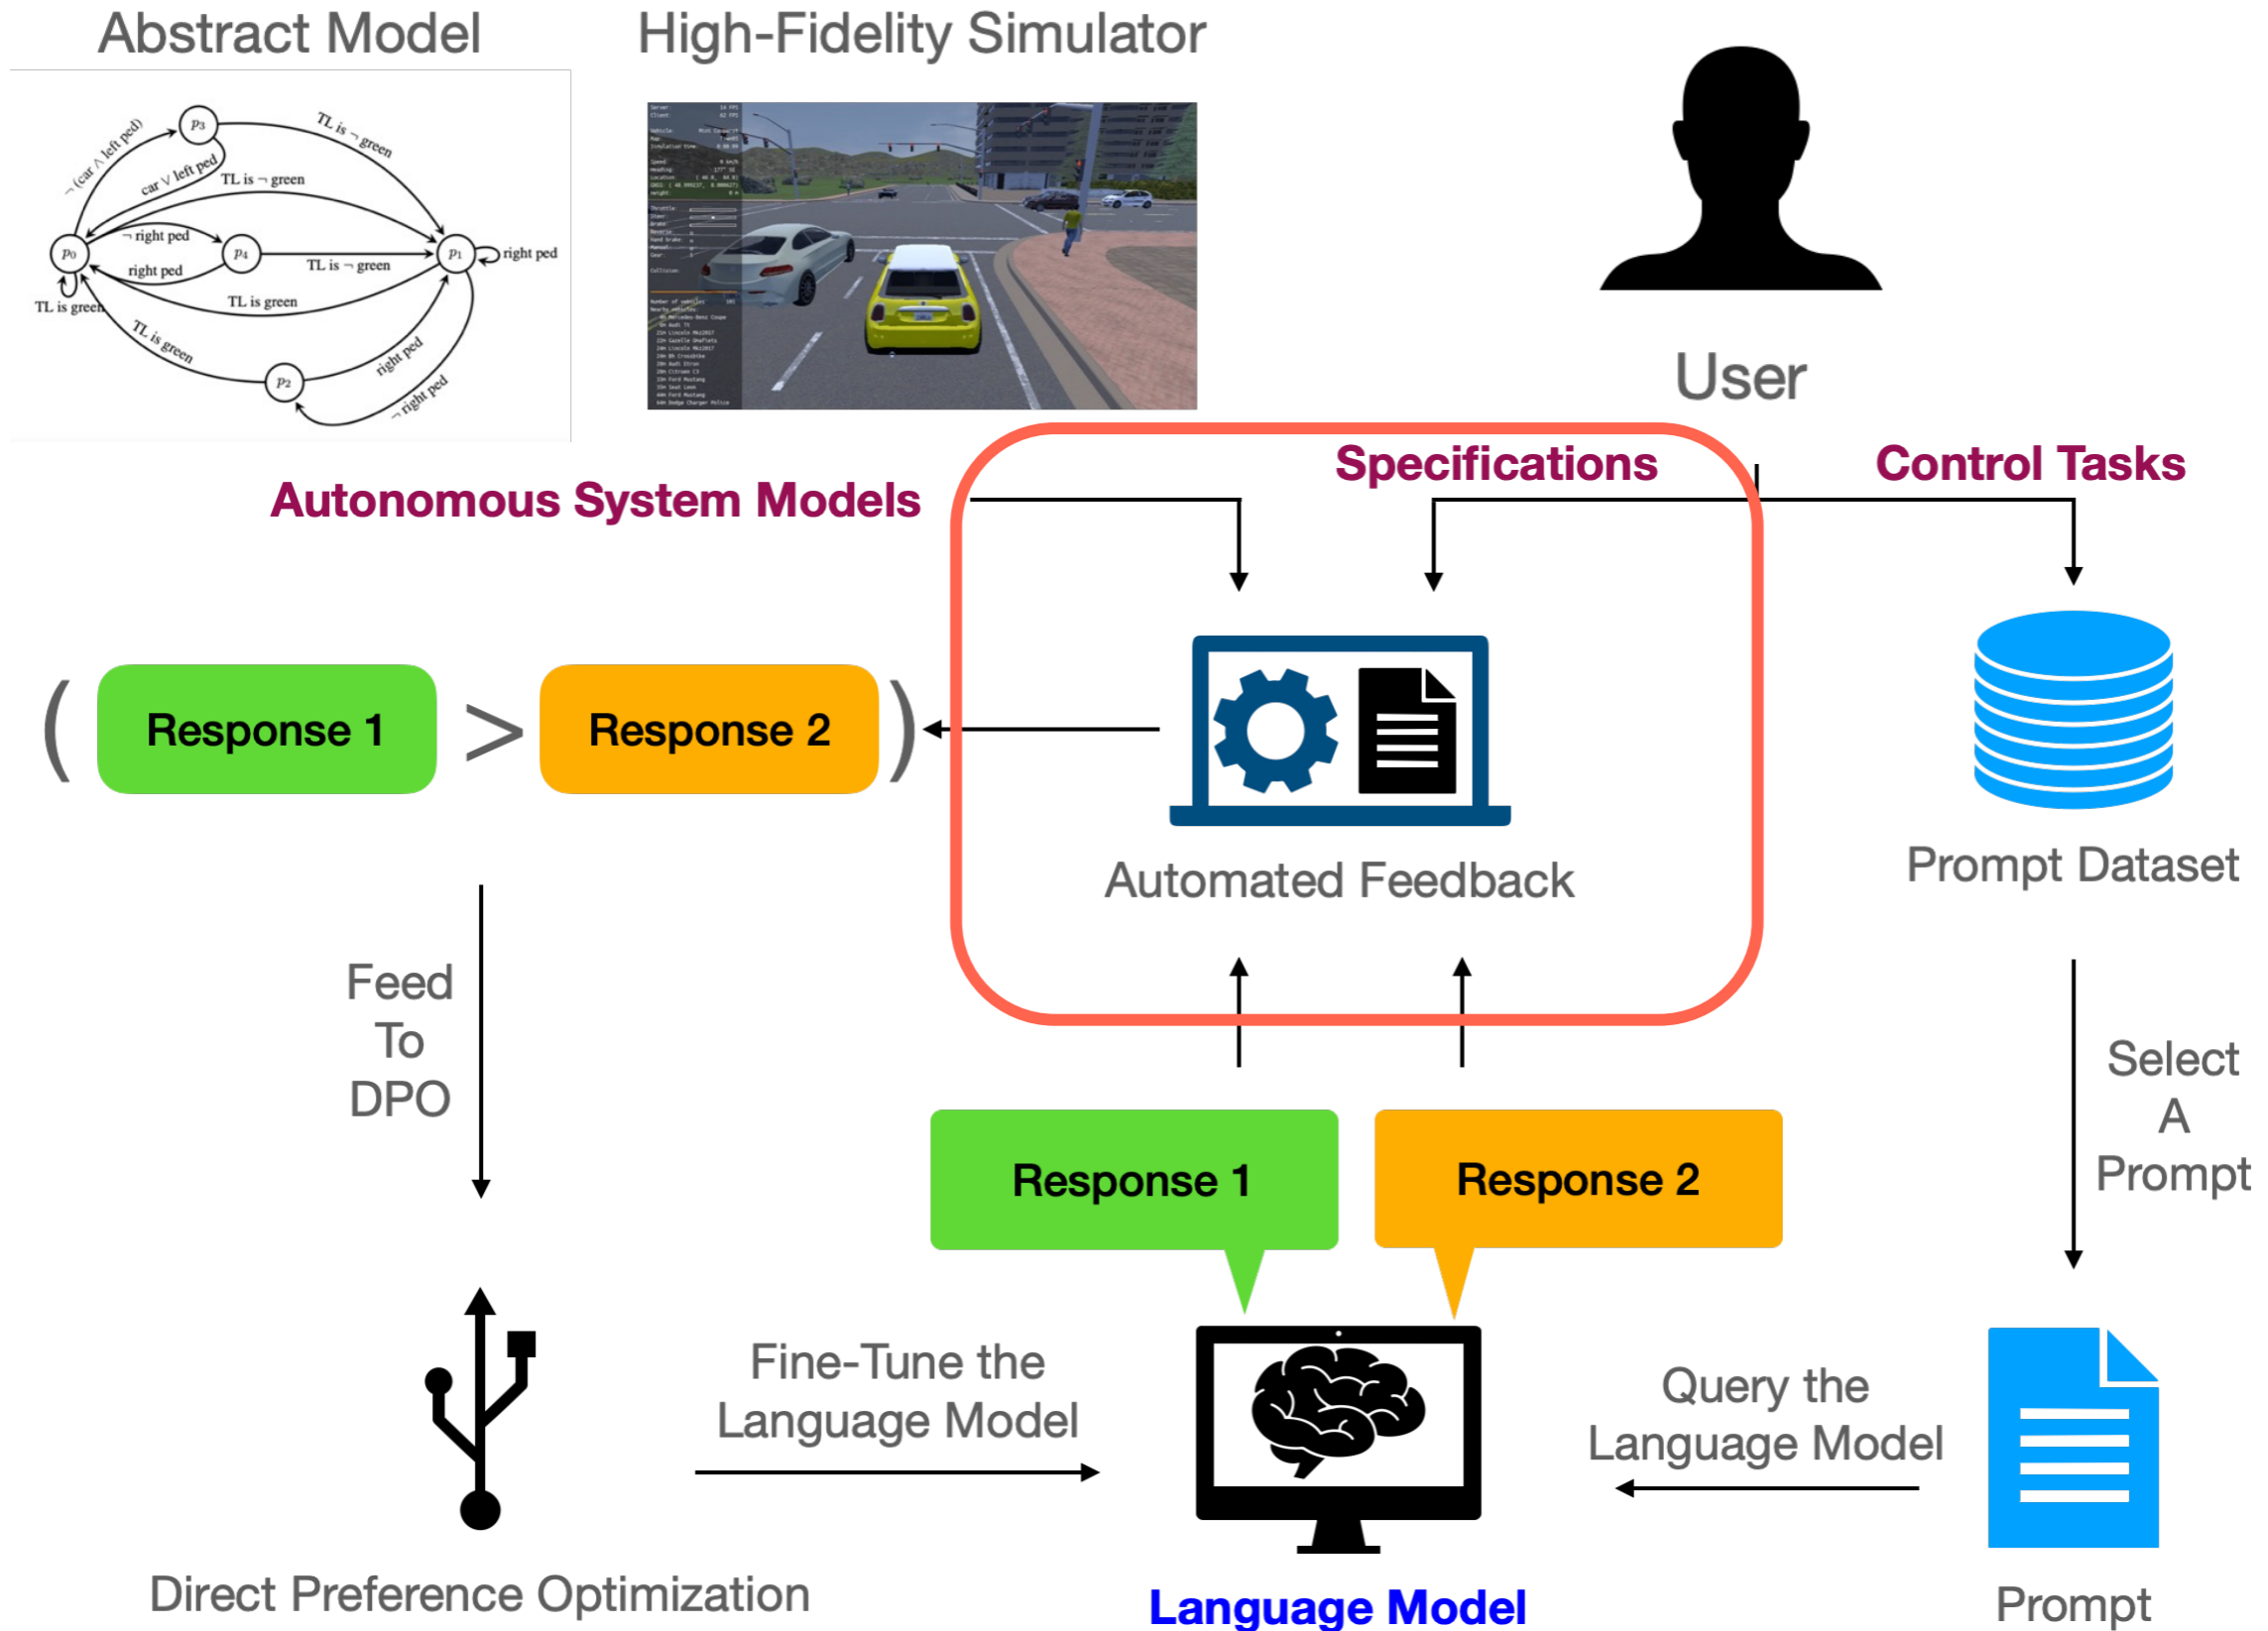
Are the controller's outcomes **guaranteed** to satisfy user-specified requirements when implemented against a system model?



Byproduct: Counterexample trace that **violates** the specification.

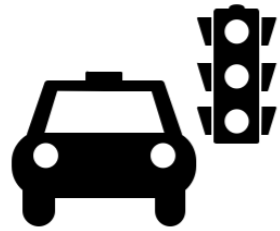


Fine-Tuning Language Models Using Formal Methods Feedback

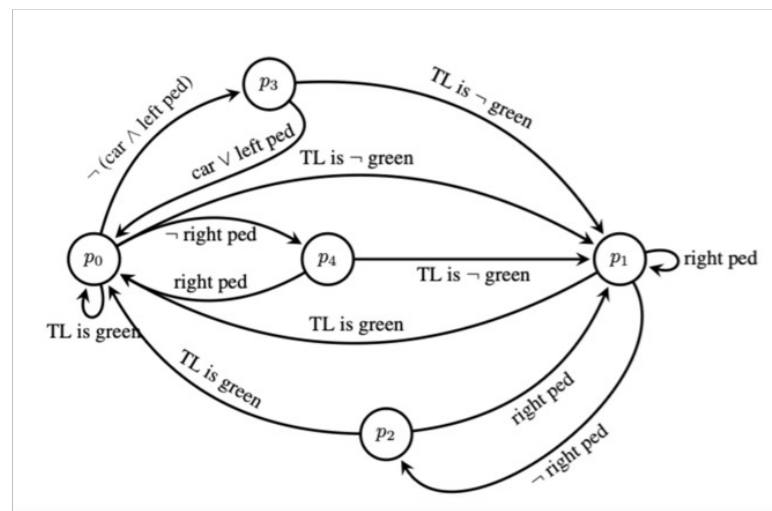


Fine-Tuning Language Models Using Formal Methods Feedback

Modeling the Autonomous System



Autonomous System

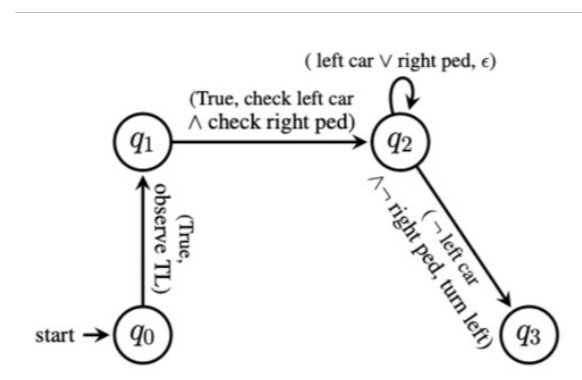


Autonomous System Model \mathcal{M}

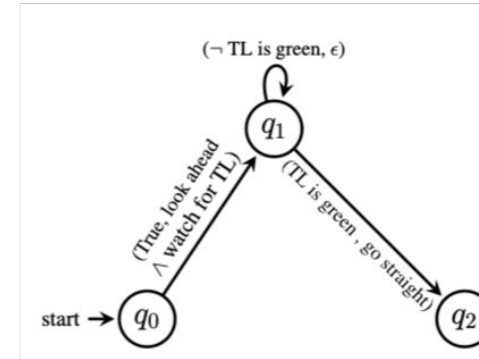
Controller Construction

Response 1

Response 2



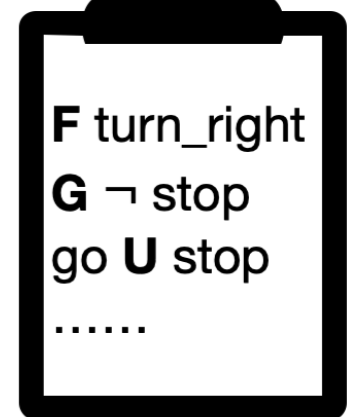
Controller \mathcal{C}_1



Controller \mathcal{C}_2



User



A Set of Specifications $\{\Phi_1, \dots, \Phi_n\}$

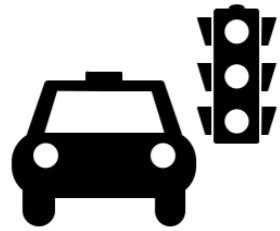


Model Checker

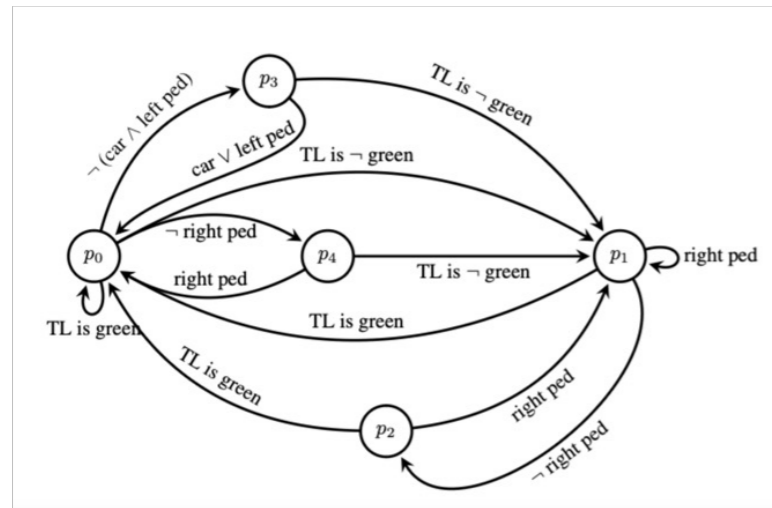
Formal Verification

Fine-Tuning Language Models Using Formal Methods Feedback

Modeling the Autonomous System



Autonomous System

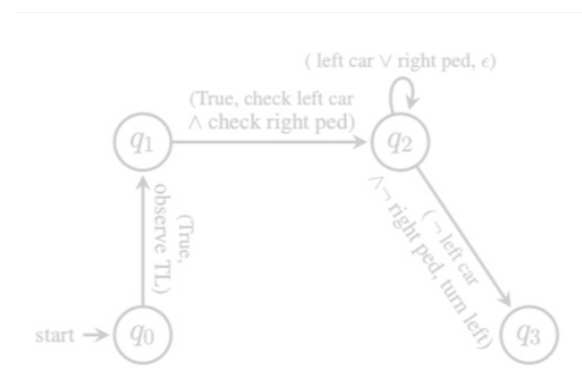


Autonomous System Model \mathcal{M}

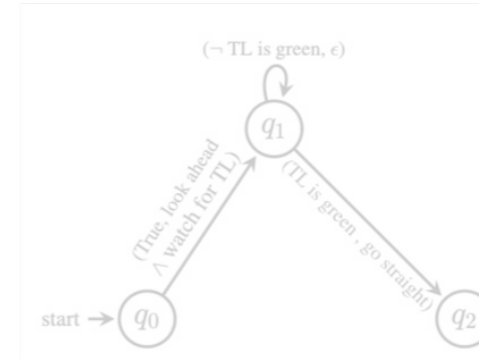
Controller Construction

Response 1

Response 2



Controller \mathcal{C}_1



Controller \mathcal{C}_2



User



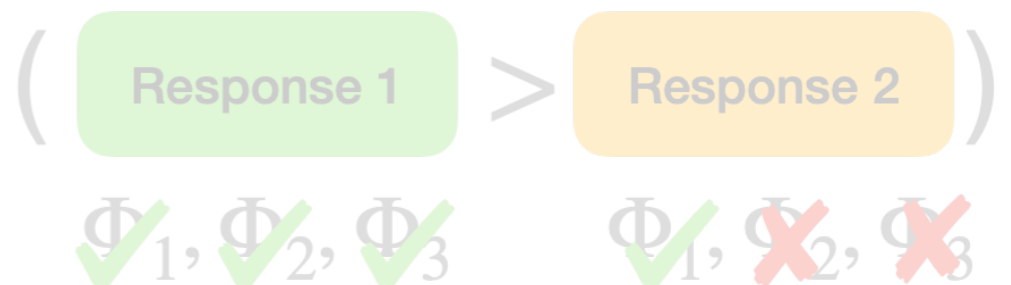
A Set of Specifications

$\{\Phi_1, \dots, \Phi_n\}$



Model Checker

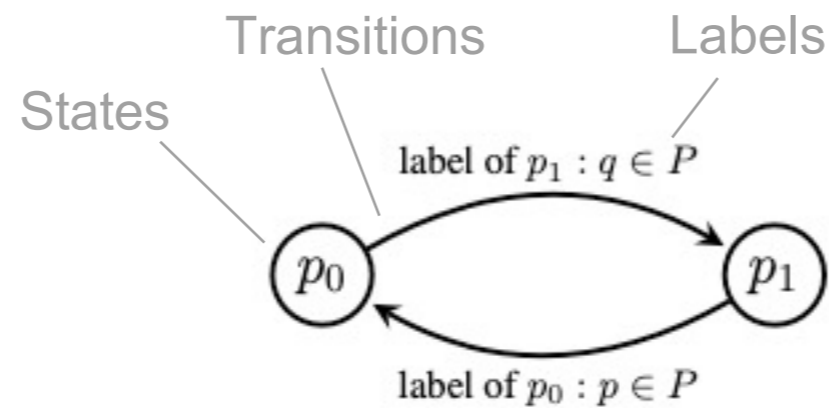
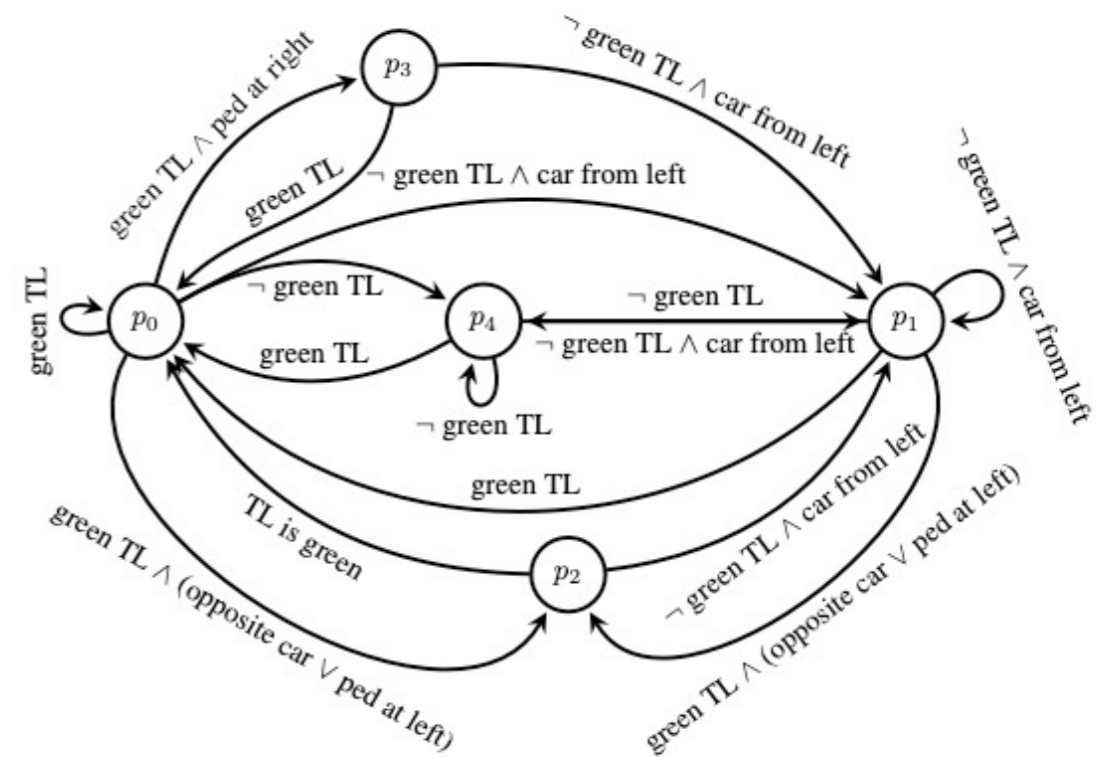
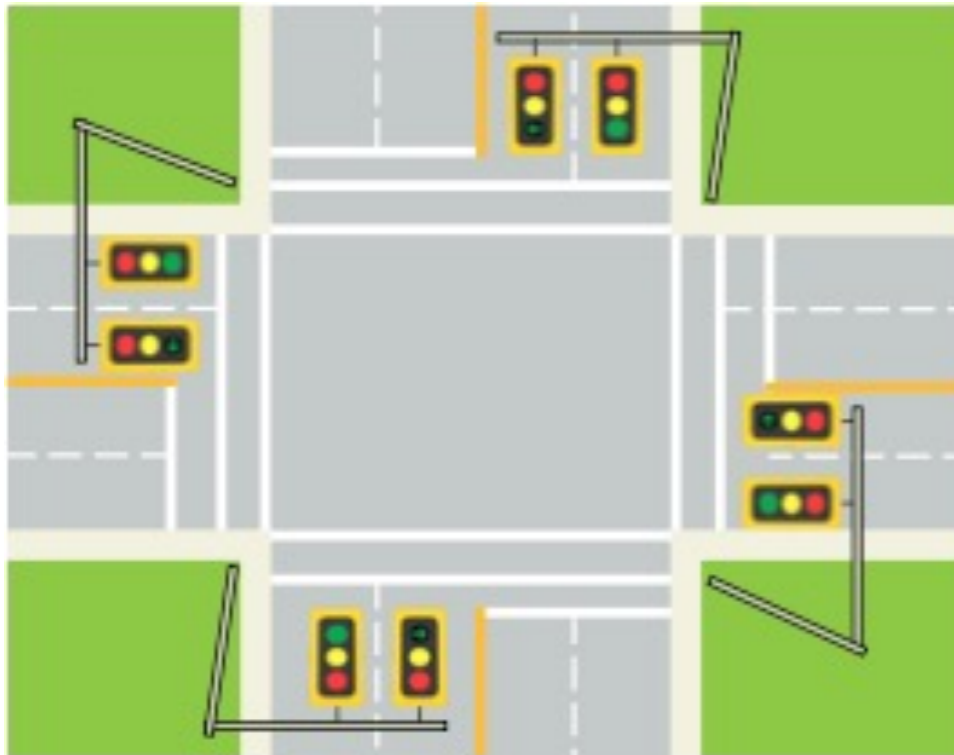
Compare the # of Satisfied Specifications



Formal Verification

Formal Methods Feedback in an Autonomous Driving System

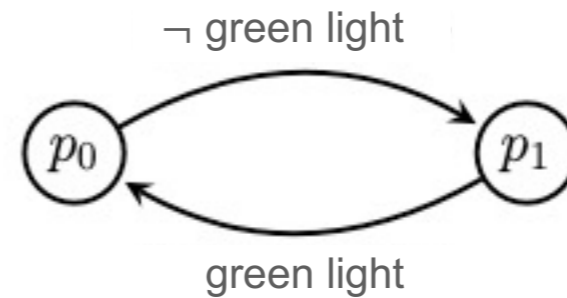
Modeling the Autonomous System



Formal Methods Feedback in an Autonomous Driving System

Modeling the Autonomous System

$P = \{\text{green light}\}$
 $S \sim \text{Traffic Light}$



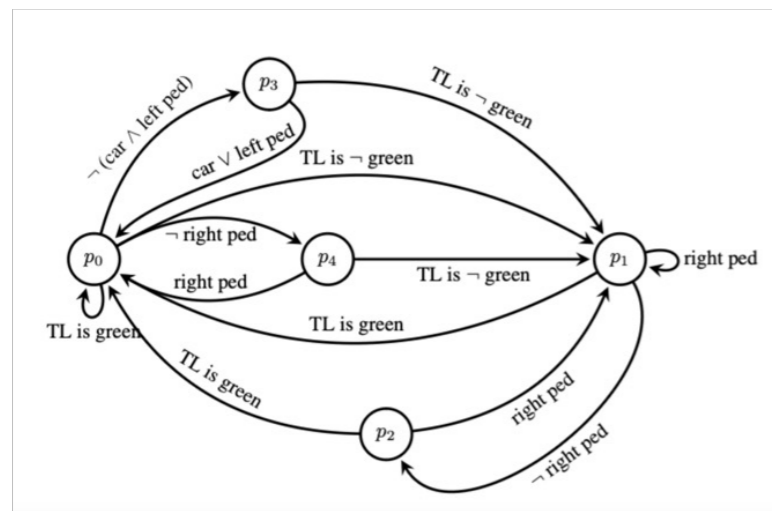
$Q_M = \{p_0 p_1\}$
 $\lambda_M(p_0) = \text{green light}$
 $\lambda_M(p_1) = \neg \text{green light}$
 $\delta_M(p_0, p_1) = 1$
 $\delta_M(p_1, p_0) = 1$

Fine-Tuning Language Models Using Formal Methods Feedback

Modeling the Autonomous System



Autonomous System

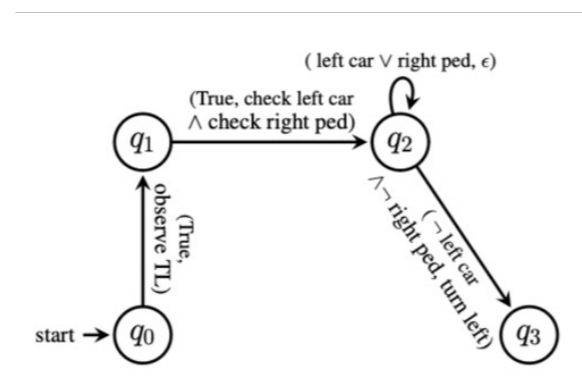


Autonomous System Model \mathcal{M}

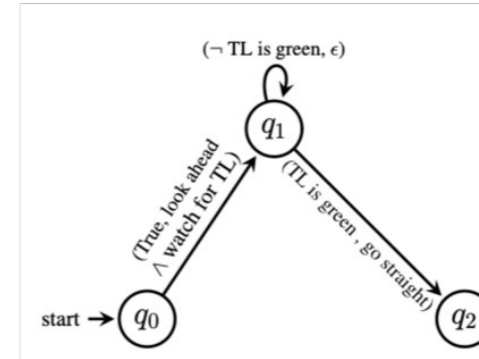
Controller Construction

Response 1

Response 2



Controller \mathcal{C}_1



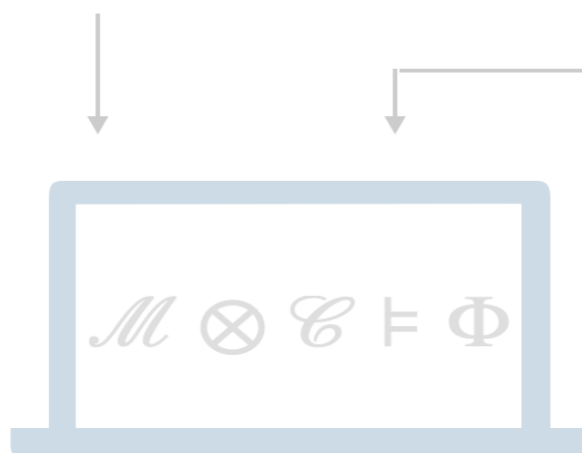
Controller \mathcal{C}_2



User



A Set of Specifications $\{\Phi_1, \dots, \Phi_n\}$



Model Checker

Compare the # of Satisfied Specifications



Formal Verification

Formal Methods Feedback in an Autonomous Driving System

Controller Construction



User

Steps for turning right at the traffic light



Language Model

1. Look straight ahead and watch for traffic light.
2. If the traffic light turns green, start moving forward.
3. As you approach the intersection, look to your left for oncoming traffic.
4. If there is no traffic coming from your left, check pedestrians on your right.
5. If it is safe, turn your vehicle right.

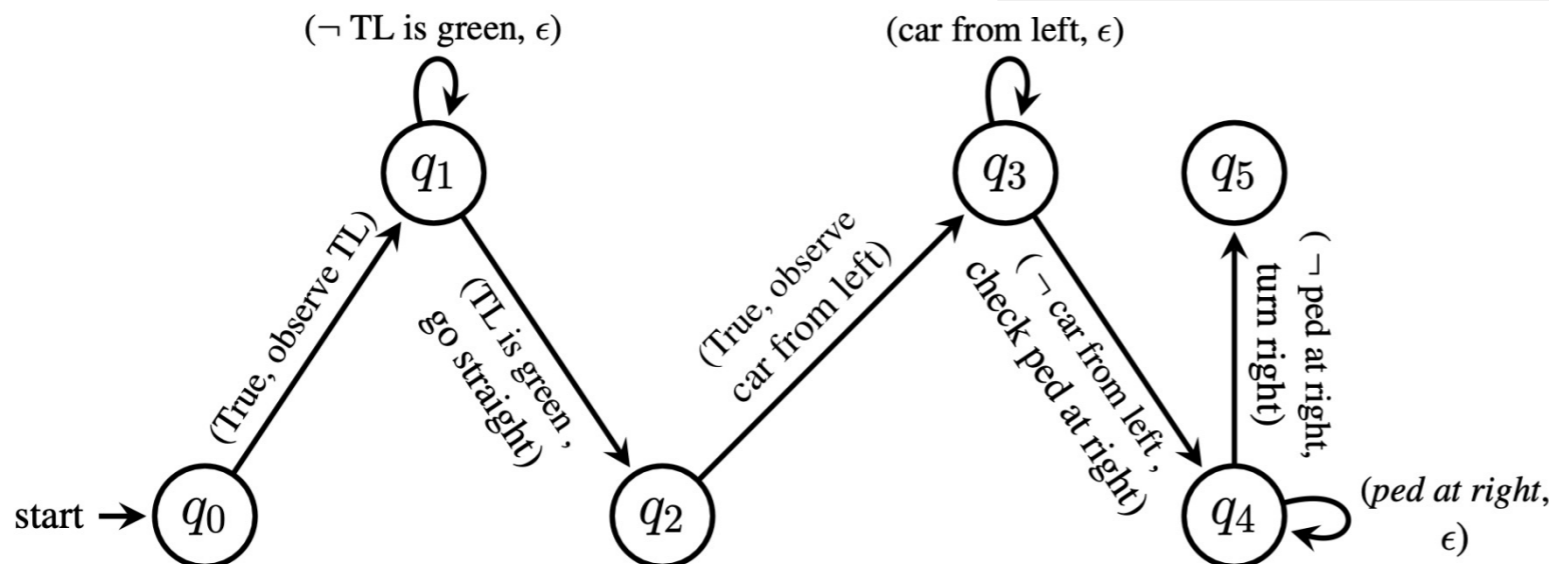
1. Observe the traffic light in front of you.
2. Check for the left approaching car and right side pedestrian.
3. If no car from the left is approaching and no pedestrian on the right, proceed to turn right.

Formal Methods Feedback in an Autonomous Driving System

Controller Construction

1. Look straight ahead and watch for traffic light.
2. If the traffic light turns green, start moving forward.
3. As you approach the intersection, look to your left for oncoming traffic.
4. If there is no traffic coming from your left, check pedestrians on your right.
5. If it is safe, turn your vehicle right.

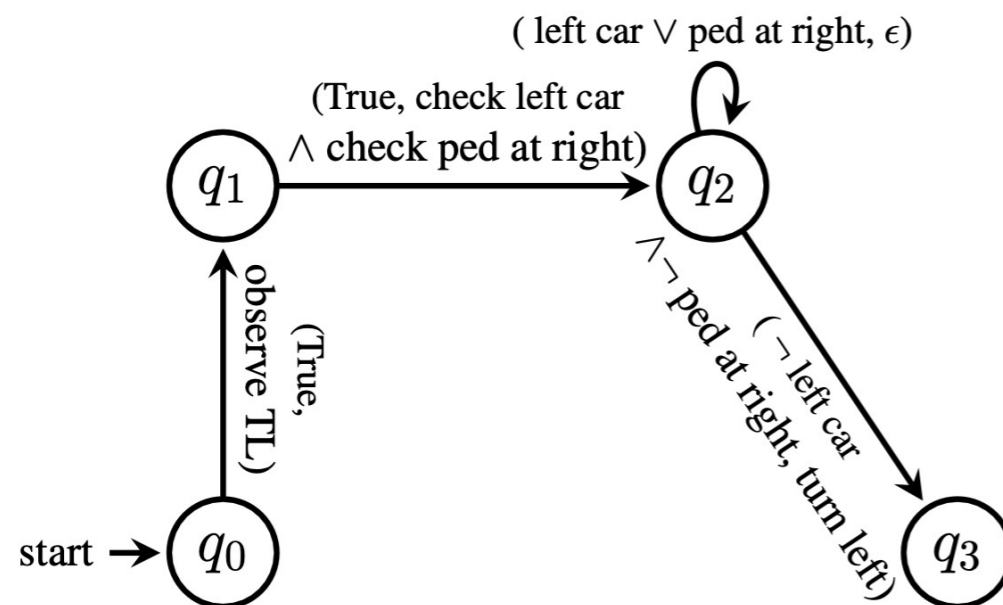
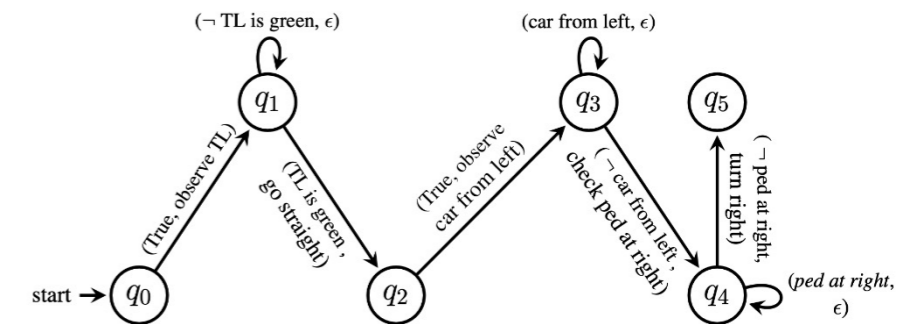
Grammar	Transition Rule
1. $\langle \text{obse } VP^A$	$q_i \xrightarrow{(True, VP^A)} q_{i+1}$
2. $\langle \text{if} \rangle$	$q_i \xrightarrow{(True, \epsilon)} q_j$
3. $\langle \text{obse } VP^A \leftarrow VP^A [j]$	$q_i \xrightarrow{(True, \epsilon)} q_j$
4. $\langle \text{if} \rangle$ if VP^C, VP^A at VP^A if VP^C	$q_i \xrightarrow{(-VP^C, \epsilon)} q_i \xrightarrow{(VP^C, VP^A)} q_j$
5. $\langle \text{if} \rangle$ if VP^C, VP^A_1 if $\neg VP^C, VP^A_2$ right VP^A_1 if VP^C , else VP^A_2	$q_k \xrightarrow{(-VP^C, VP^A_2)} q_i \xrightarrow{(VP^C, VP^A_1)} q_j$
$VP^A \leftarrow \text{wait } VP^C$ VP^A	$q_i \xrightarrow{(-VP^C, \epsilon)} q_i \xrightarrow{(VP^C, VP^A)} q_{i+1}$
$VP^A \leftarrow VP^A$ after VP^C	$q_i \xrightarrow{(-VP^C, VP^A)} q_i \xrightarrow{(VP^C, \epsilon)} q_{i+1}$



Formal Methods Feedback in an Autonomous Driving System

Controller Construction

1. Look straight ahead and watch for traffic light.
2. If the traffic light turns green, start moving forward.
3. As you approach the intersection, look to your left for oncoming traffic.
4. If there is no traffic coming from your left, check pedestrians on your right.
5. If it is safe, turn your vehicle right.



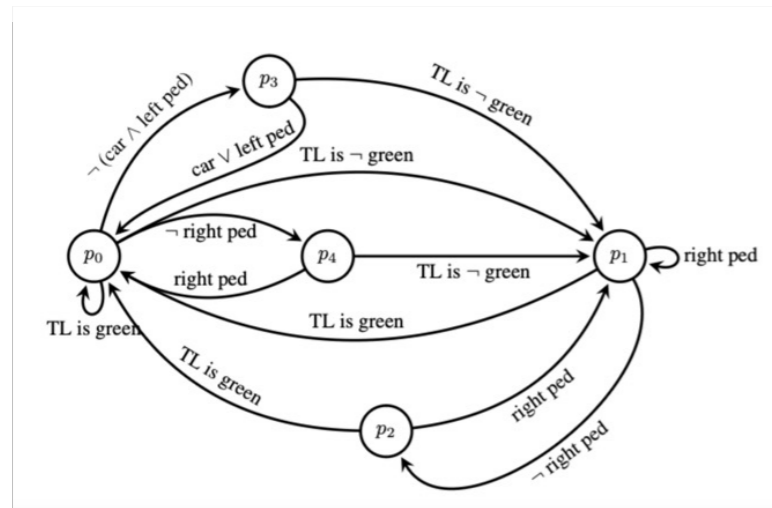
1. Observe the traffic light in front of you.
2. Check for the left approaching car and right side pedestrian.
3. If no car from the left is approaching and no pedestrian on the right, proceed to turn right.

Fine-Tuning Language Models Using Formal Methods Feedback

Modeling the Autonomous System



Autonomous System

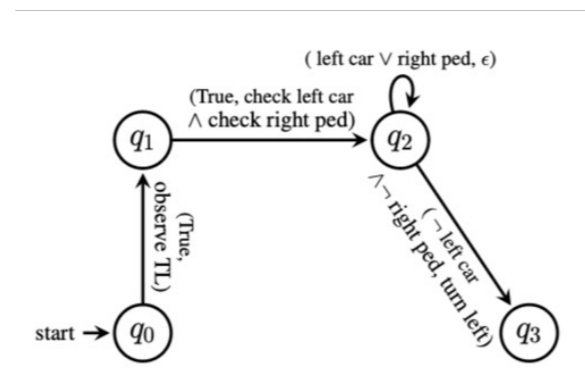


Autonomous System Model \mathcal{M}

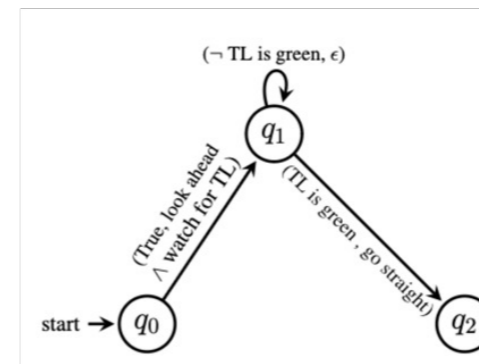
Controller Construction



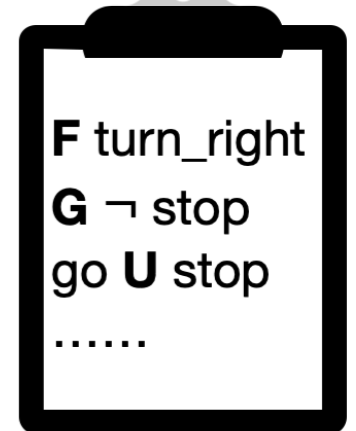
User



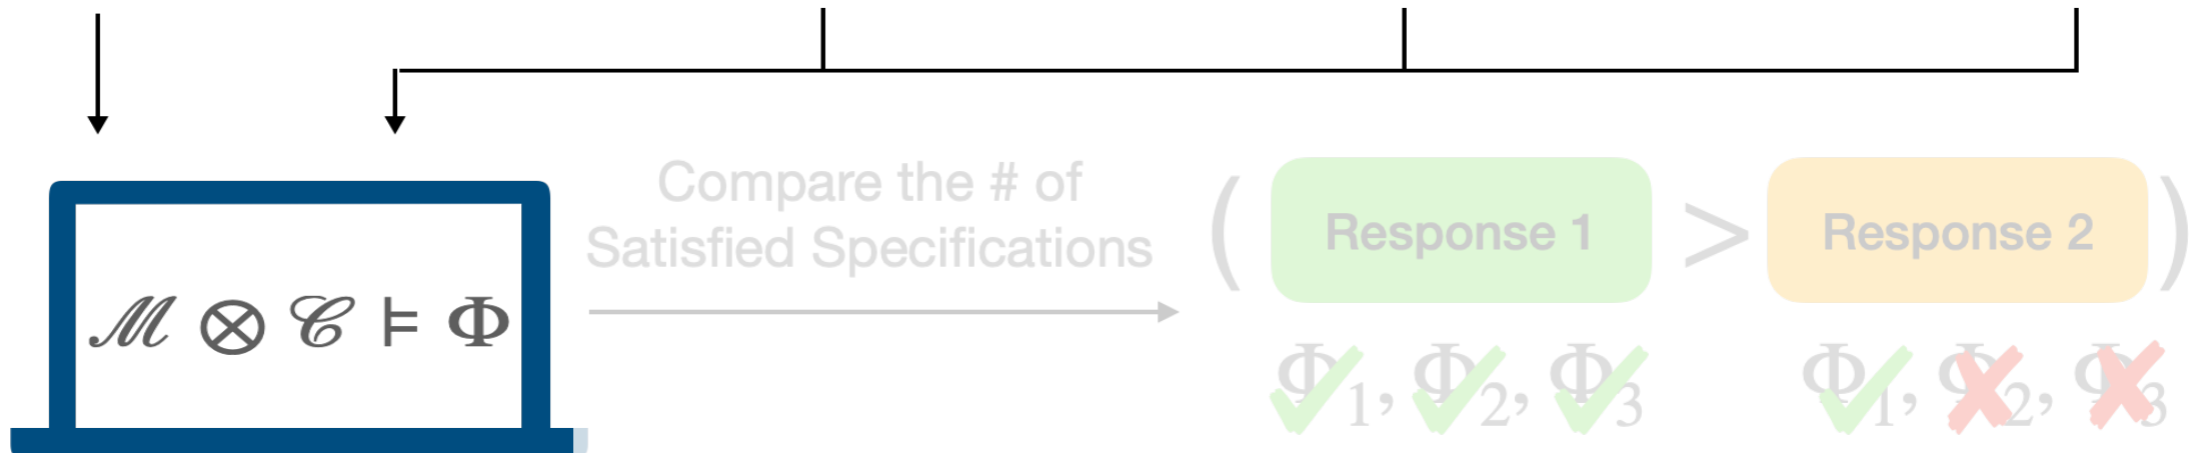
Controller \mathcal{C}_1



Controller \mathcal{C}_2



A Set of Specifications $\{\Phi_1, \dots, \Phi_n\}$

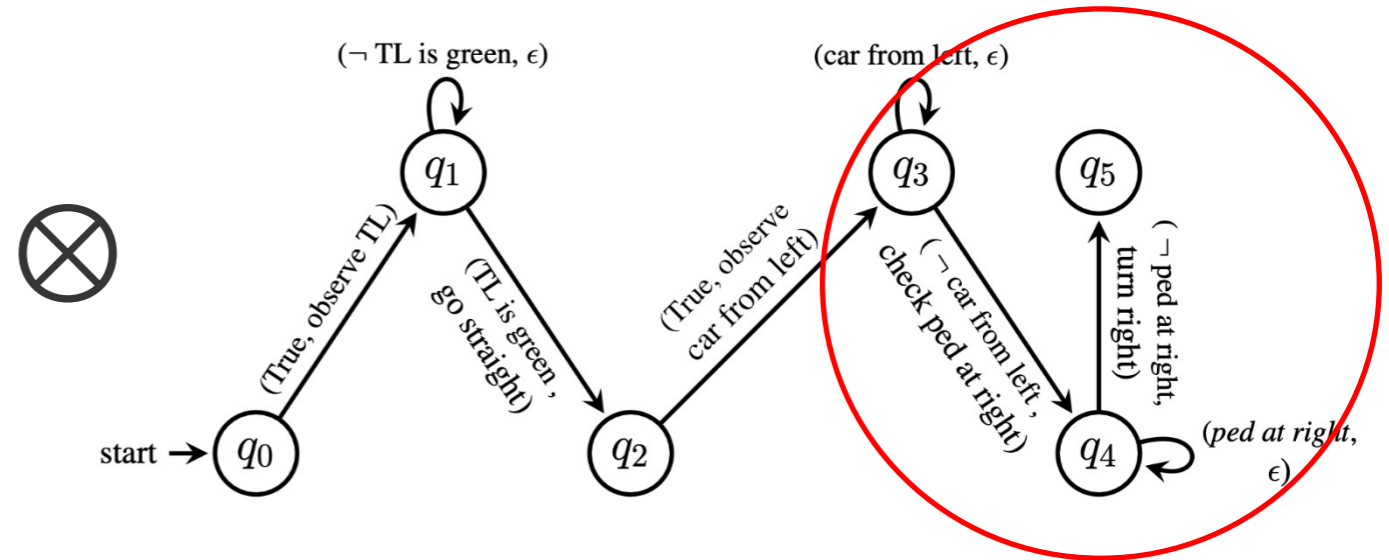
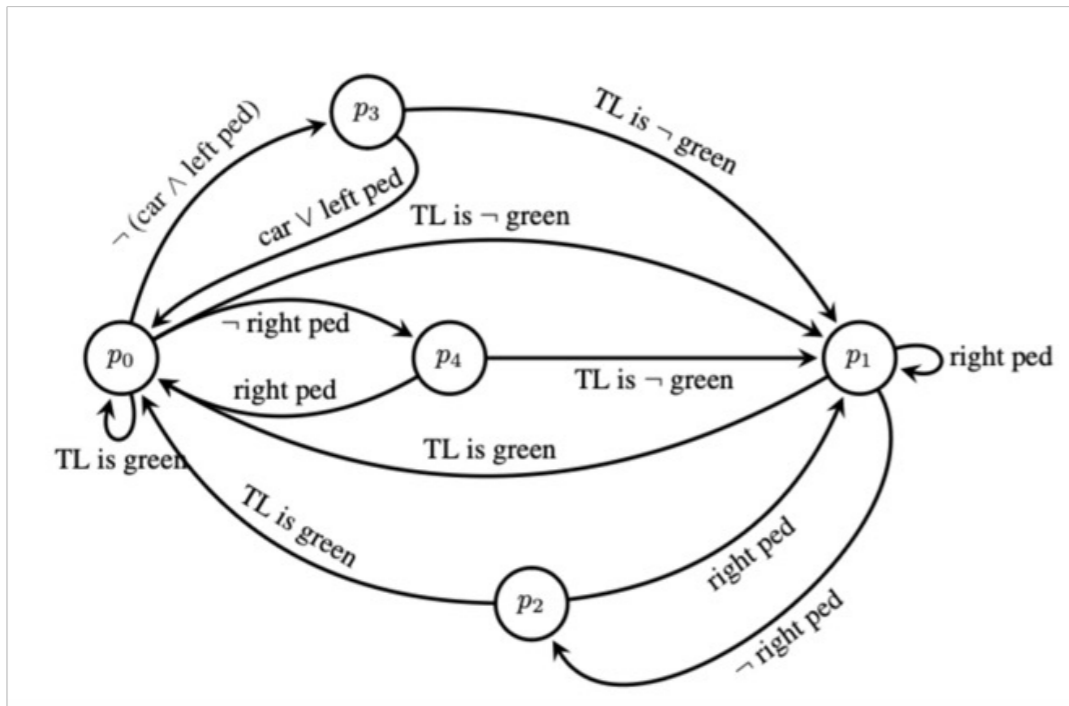


Model Checker

Formal Verification

Formal Methods Feedback in an Autonomous Driving System

Formal Verification



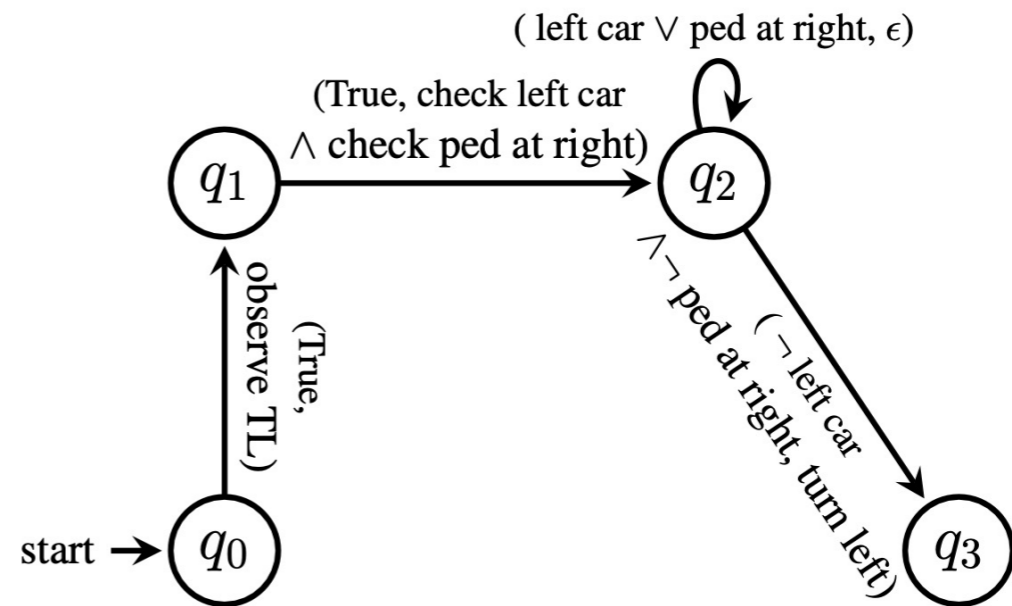
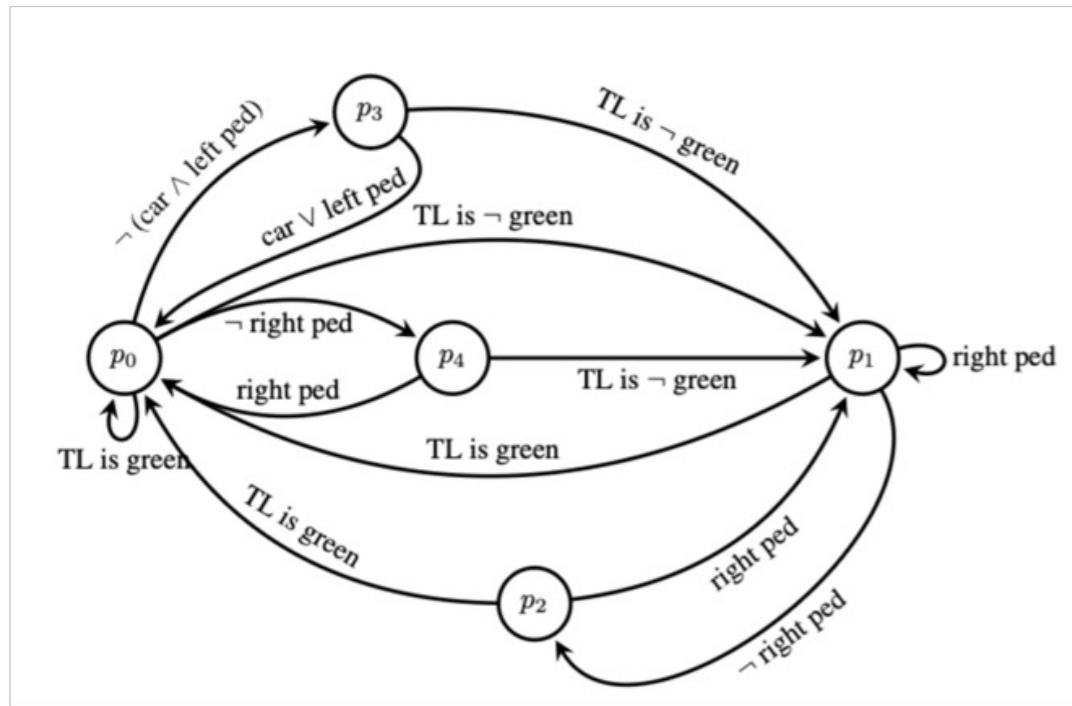
(\neg green traffic light \rightarrow \neg go straight), ✓

(stop sign \rightarrow \diamond stop), ✓

\neg turn right \vee \neg (car from left \vee pedestrian at right), ✗

Formal Methods Feedback in an Autonomous Driving System

Formal Verification



$(\neg \text{green traffic light} \rightarrow \neg \text{go straight})$, ✓

$(\text{stop sign} \rightarrow \diamond \text{stop})$, ✓

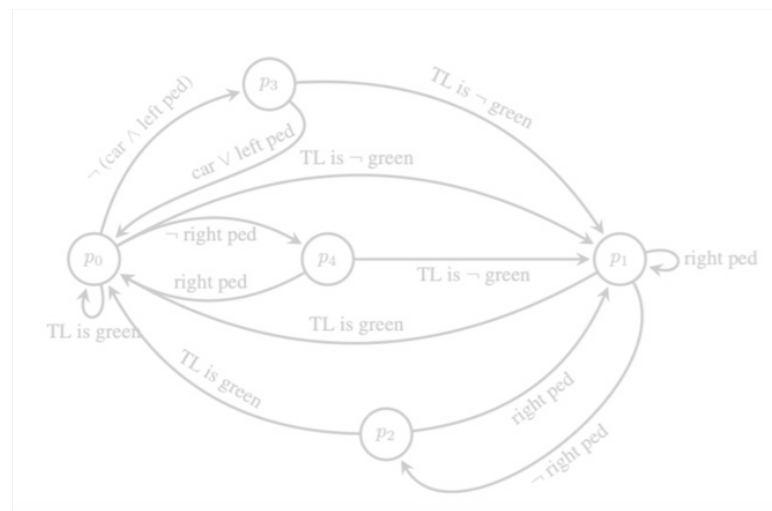
$\neg \text{turn right} \vee \neg(\text{car from left} \vee \text{pedestrian at right})$, ✓

Fine-Tuning Language Models Using Formal Methods Feedback

Modeling the Autonomous System



Autonomous System

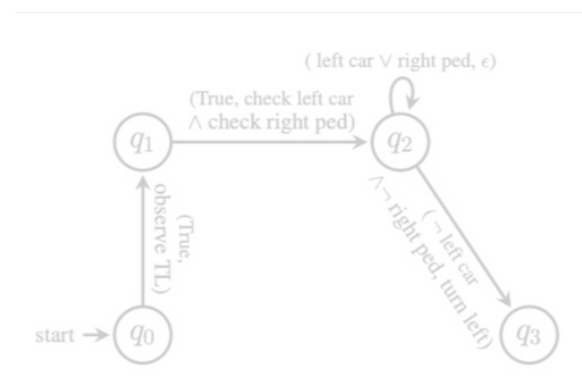


Autonomous System Model \mathcal{M}

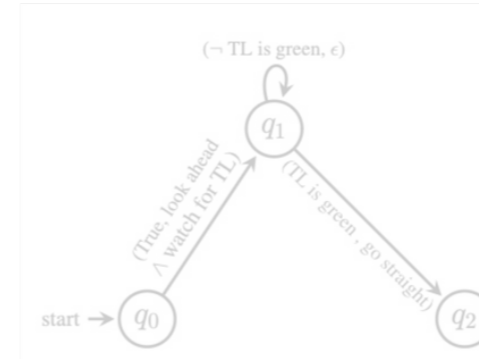
Controller Construction

Response 1

Response 2



Controller \mathcal{C}_1



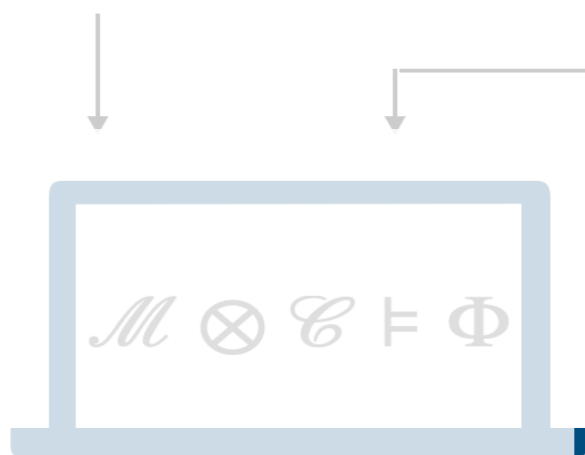
Controller \mathcal{C}_2



User



A Set of Specifications $\{\Phi_1, \dots, \Phi_n\}$



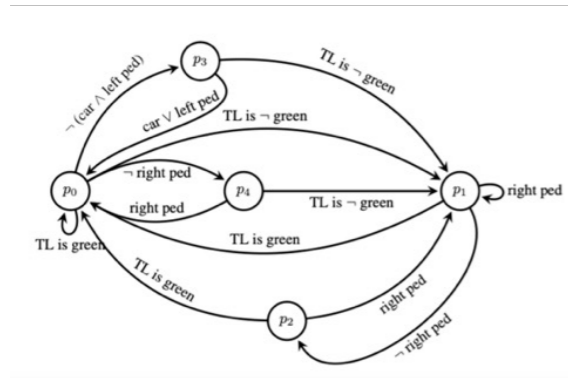
Model Checker

Compare the # of Satisfied Specifications



Fine-Tuning Language Models Using Formal Methods Feedback

Abstract Model



High-Fidelity Simulator

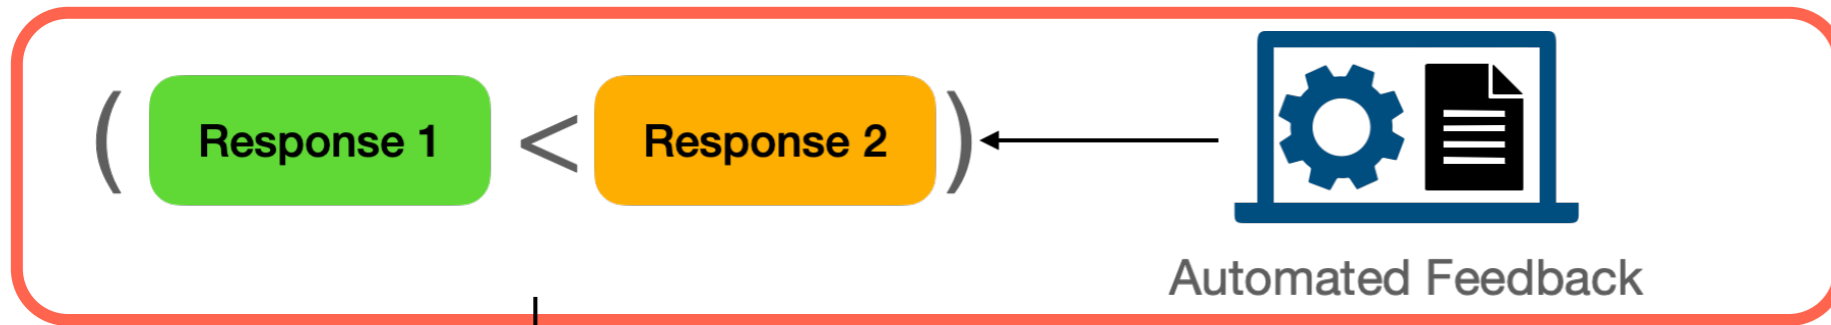


User

Autonomous System Models

Specifications

Control Tasks



Prompt Dataset

Feed To DPO



Direct Preference Optimization

Fine-Tune the Language Model

Language Model



Query the Language Model

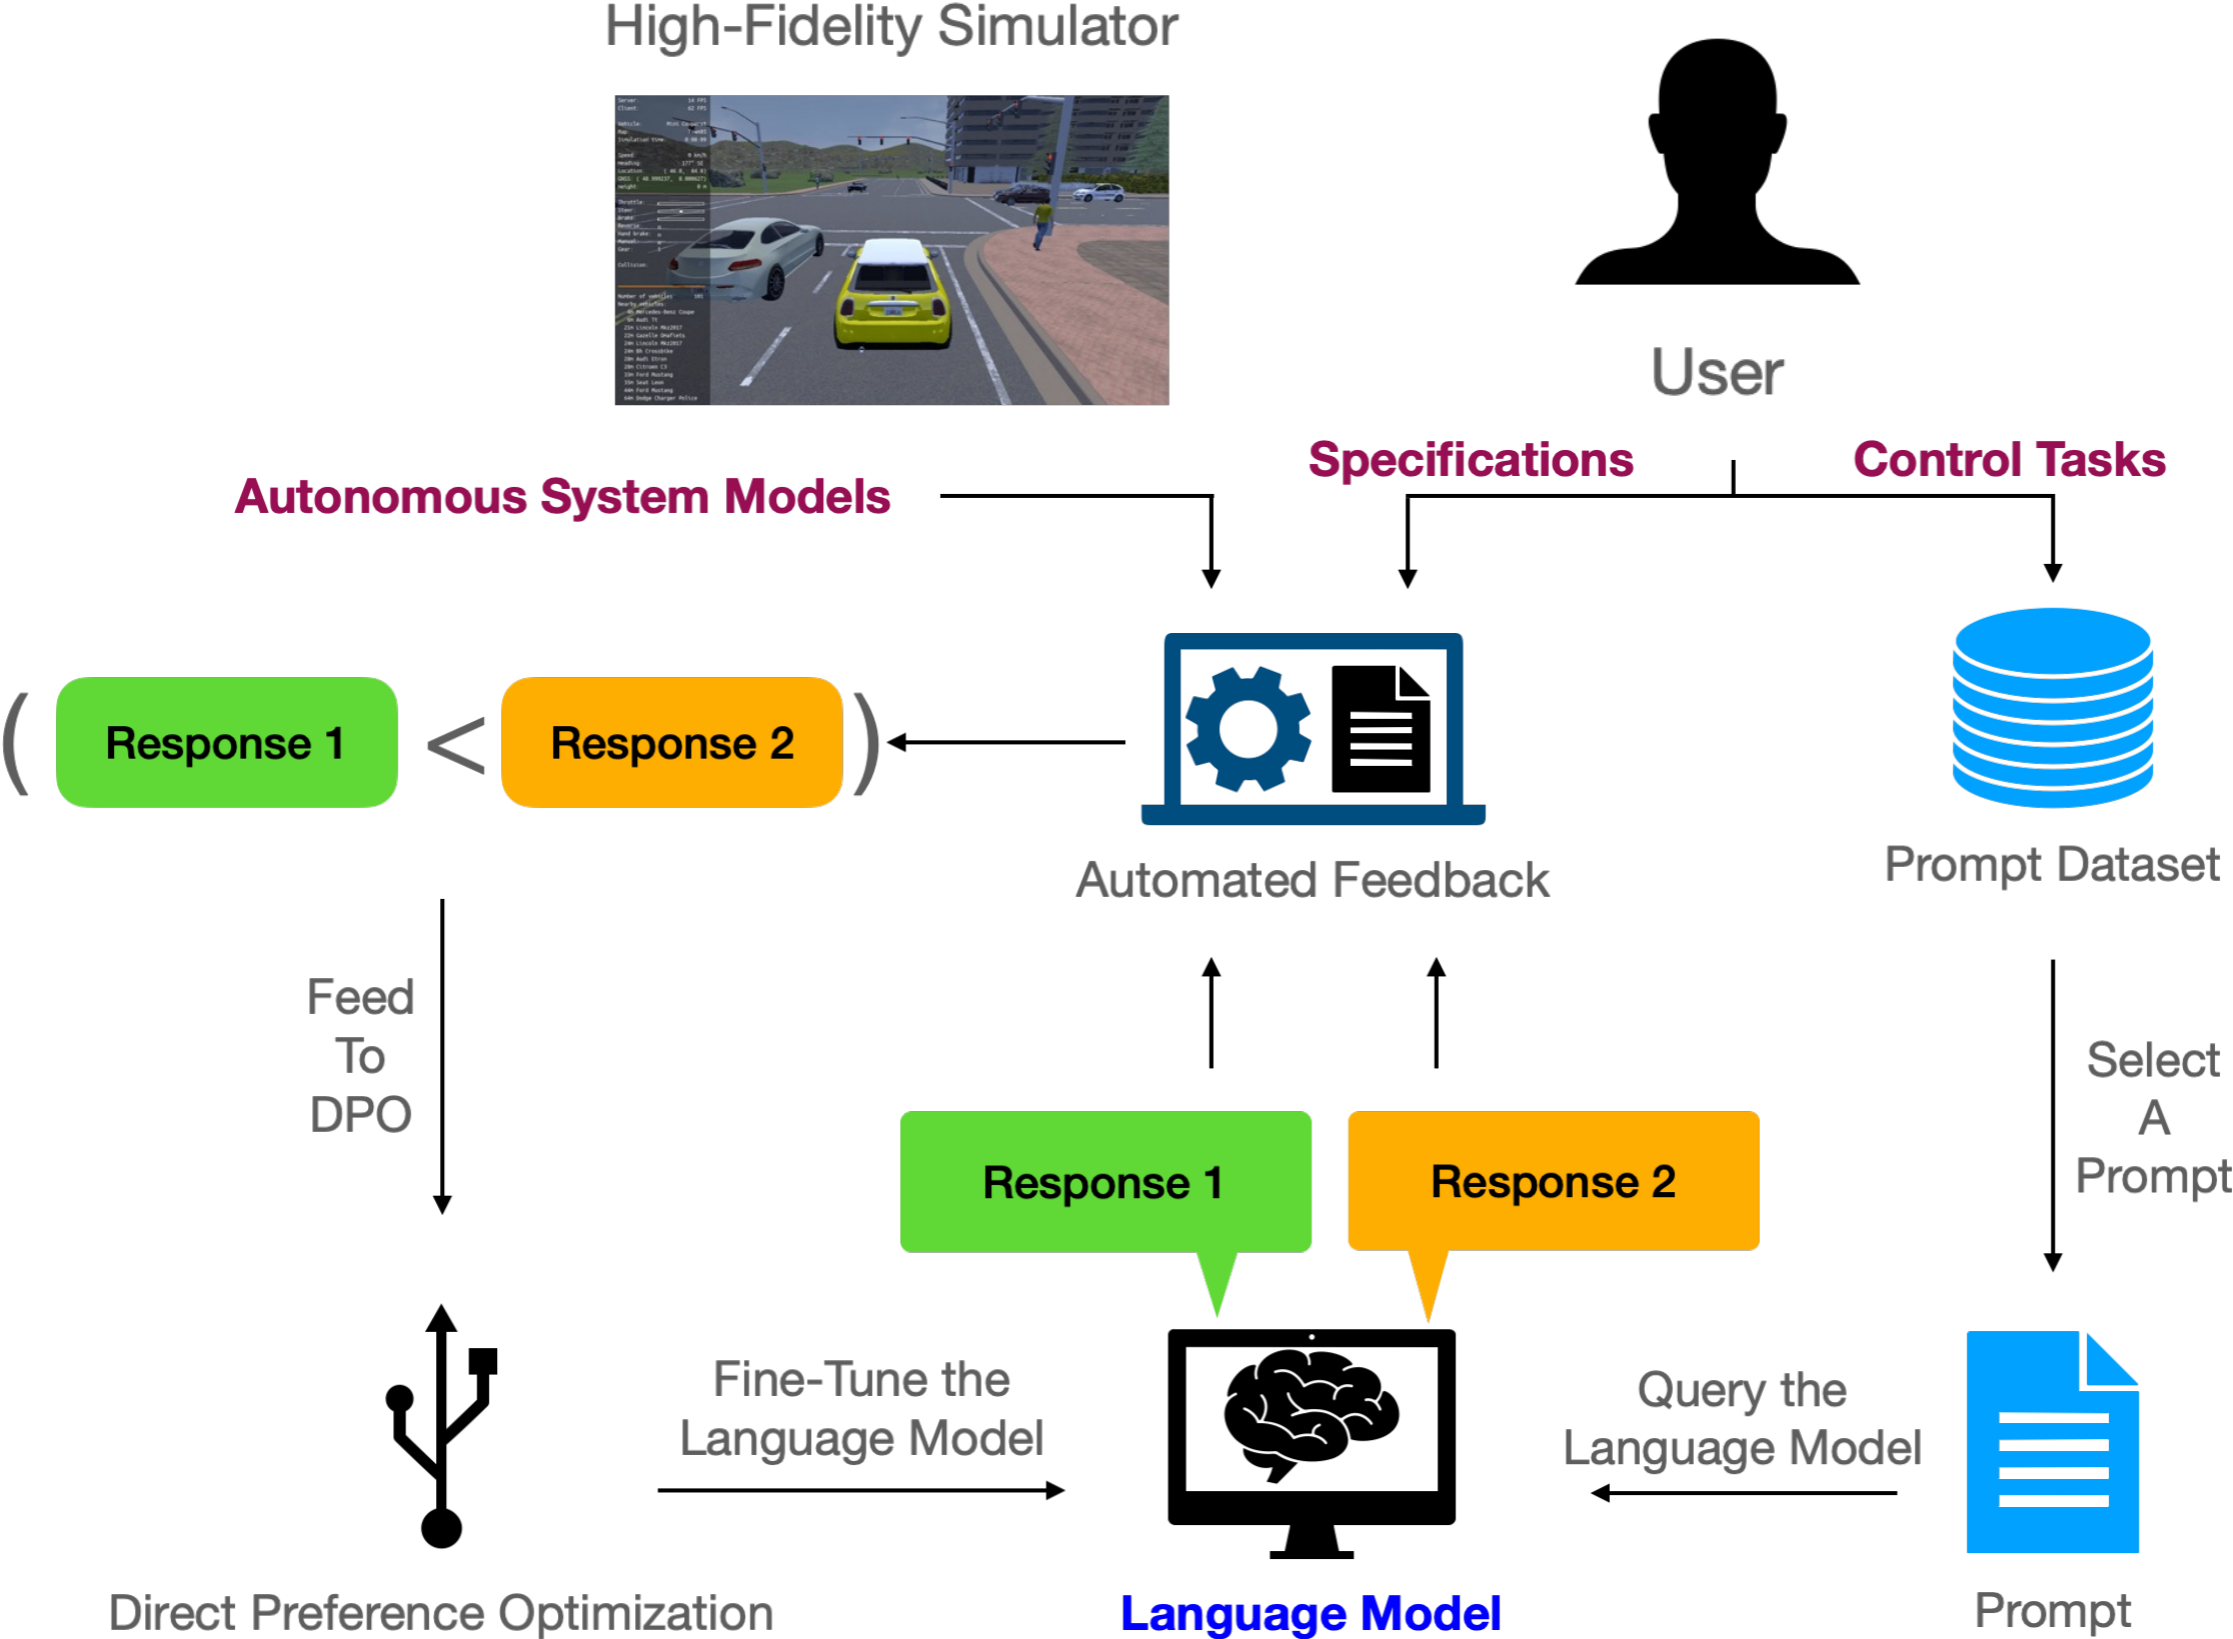


Select A Prompt



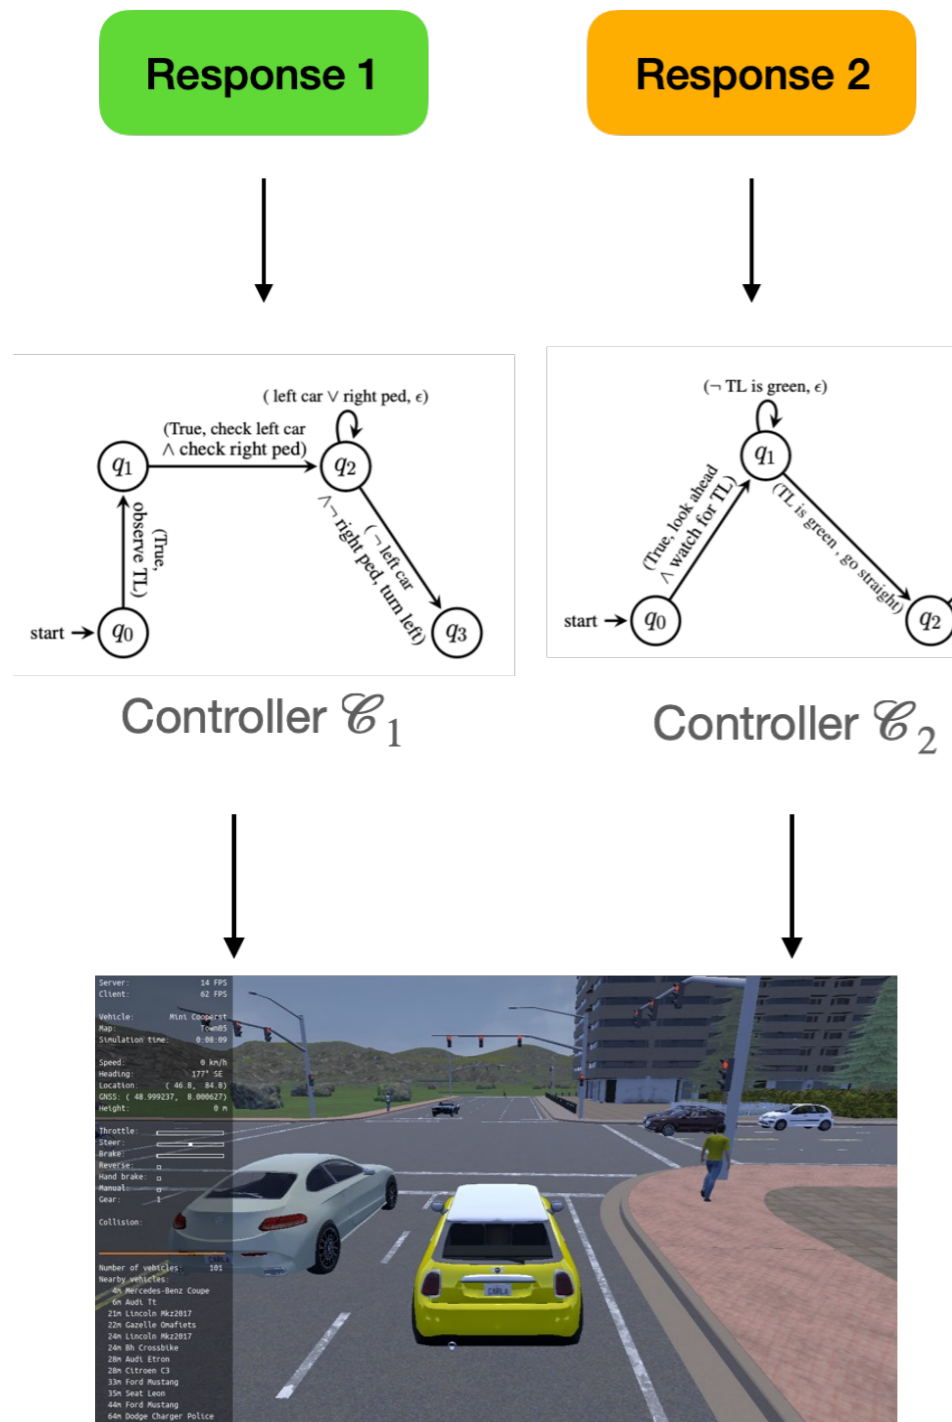
Prompt

Fine-Tuning Language Models Using Formal Methods Feedback



Fine-Tuning Language Models Using Formal Methods Feedback

Empirical Evaluation via Simulation

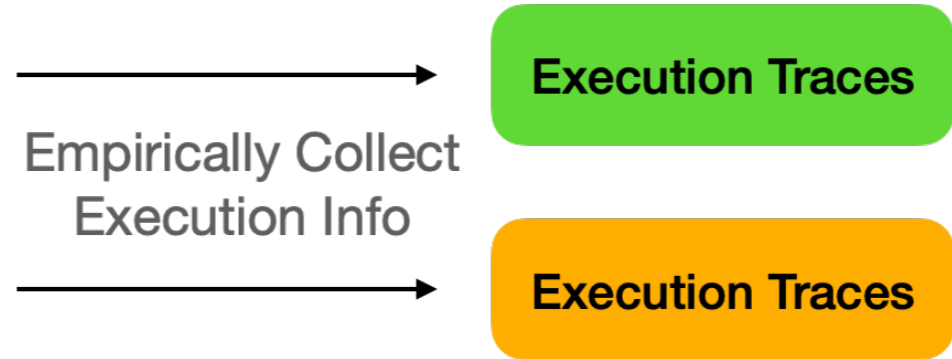


Information from the Simulator

```

Vector3D(x=10.495529, y=5.662697, z=2.781272) Class: TL or TS
Vector3D(x=2.850960, y=18.144859, z=2.781269) Class: TL or TS
Vector3D(x=-9.272949, y=29.014732, z=2.781277) Class: TL or TS
-----
Vector3D(x=3.002785, y=17.807884, z=2.781323) Class: TL or TS
Vector3D(x=-9.121124, y=28.677757, z=2.781330) Class: TL or TS
-----
Vector3D(x=3.161827, y=17.457993, z=2.782099) Class: TL or TS
Vector3D(x=-8.962082, y=28.327866, z=2.782107) Class: TL or TS
    
```

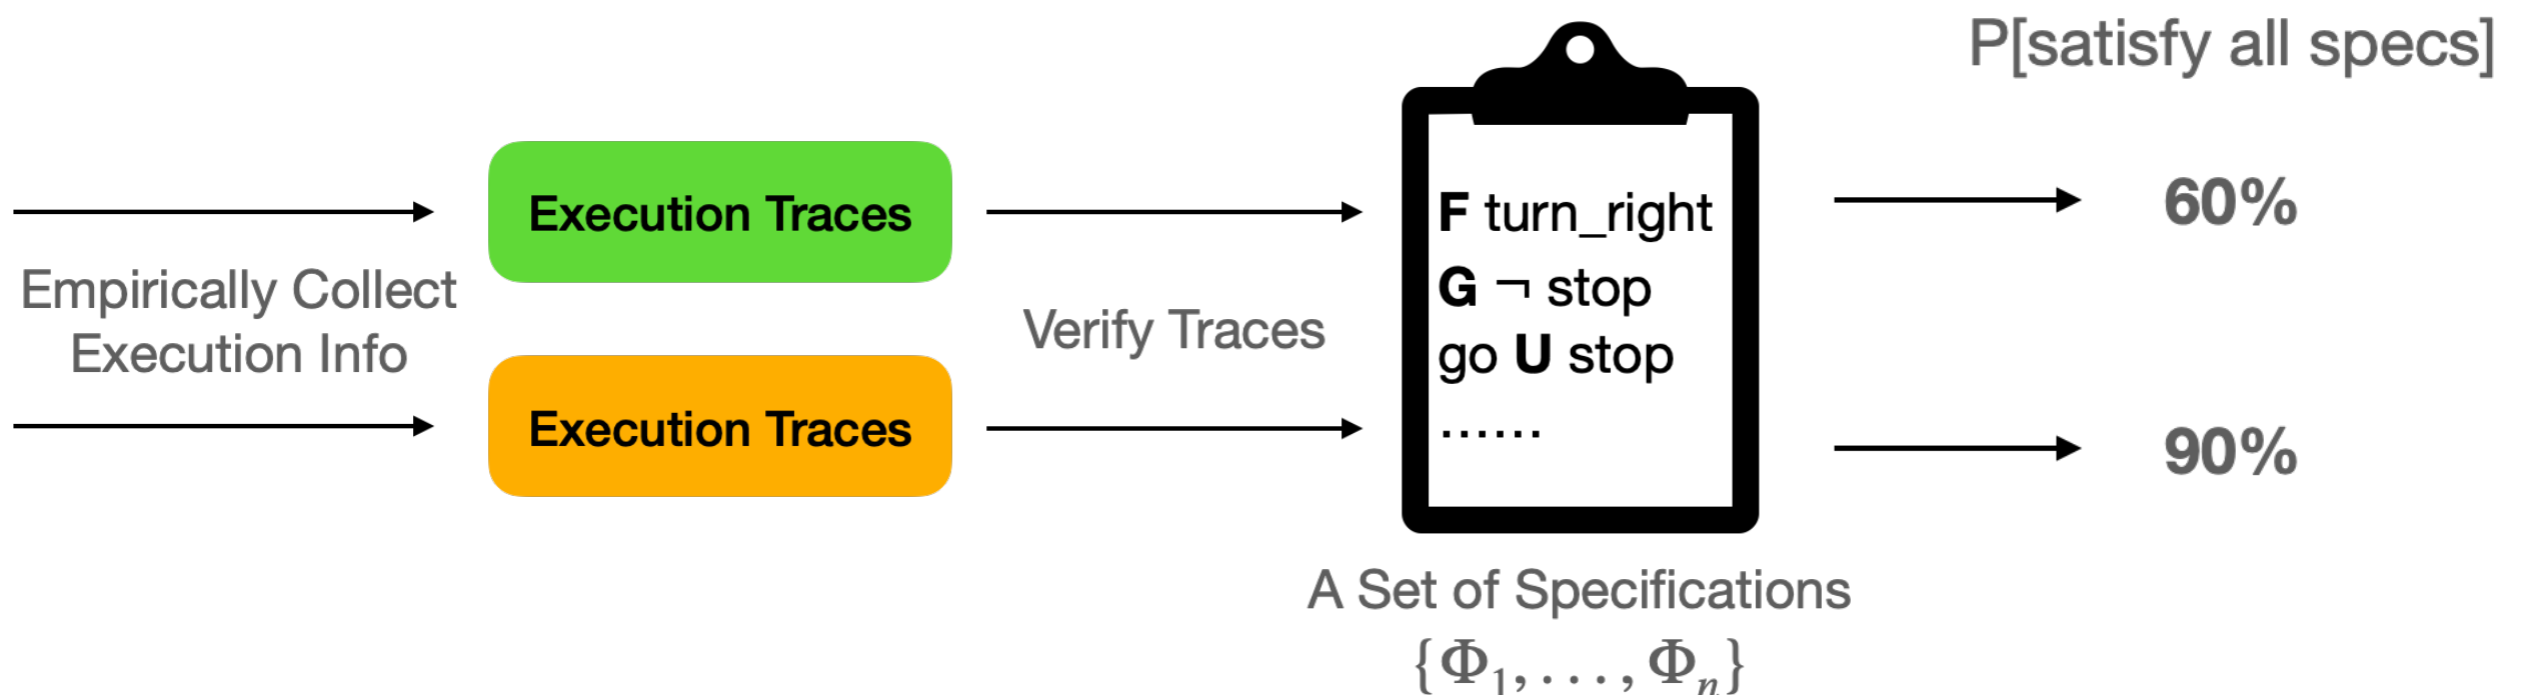
Execution Trace
(left_car \wedge go_straight, pedestrian \wedge stop, \neg pedestrian \wedge go_straight)



High-Fidelity Simulator

Fine-Tuning Language Models Using Formal Methods Feedback

Empirical Evaluation via Simulation



Quantitative Analysis

Empirical Evaluation via Simulation

Carla Simulator: Extract execution traces.

Carla Simulation Video



Object and Position Information

```
-----  
Vector3D(x=7.464157, y=-23.339211, z=2.899784) Class: TL or TS  
Vector3D(x=7.464157, y=-23.339211, z=2.232456) Class: TL or TS  
Vector3D(x=7.627457, y=-43.336159, z=2.806130) Class: TL or TS  
Vector3D(x=7.699051, y=-43.277988, z=1.985428) Class: TL or TS  
Vector3D(x=-12.249878, y=-28.675850, z=2.482181) Class: TL or TS  
Vector3D(x=-12.471230, y=-8.660095, z=3.196321) Class: TL or TS  
Vector3D(x=-12.468399, y=-8.708710, z=-1.387742) Class: TL or TS  
Vector3D(x=-7.435089, y=-44.920258, z=0.036881) Vehicle, ID: 91  
Vector3D(x=2.802956, y=-17.326714, z=0.022928) Vehicle, ID: 89  
Vector3D(x=-7.692642, y=-21.904179, z=-0.002782) Vehicle, ID: 87  
-----  
Vector3D(x=7.518684, y=-22.961929, z=2.900478) Class: TL or TS  
Vector3D(x=7.518684, y=-22.961929, z=2.233150) Class: TL or TS  
Vector3D(x=7.681984, y=-42.958878, z=2.806824) Class: TL or TS  
Vector3D(x=7.753578, y=-42.900707, z=1.986122) Class: TL or TS  
Vector3D(x=-12.210609, y=-48.282810, z=2.822465) Class: TL or TS  
Vector3D(x=-12.270897, y=-48.344036, z=1.984991) Class: TL or TS  
Vector3D(x=-12.195351, y=-28.298569, z=2.482875) Class: TL or TS  
Vector3D(x=-12.416702, y=-8.282814, z=3.197015) Class: TL or TS  
Vector3D(x=-12.413872, y=-8.331429, z=-1.387048) Class: TL or TS  
Vector3D(x=-7.385361, y=-44.102268, z=0.036914) Vehicle, ID: 91  
Vector3D(x=2.859573, y=-17.434631, z=0.023687) Vehicle, ID: 89  
Vector3D(x=-7.644753, y=-20.938622, z=-0.002182) Vehicle, ID: 87  
-----
```

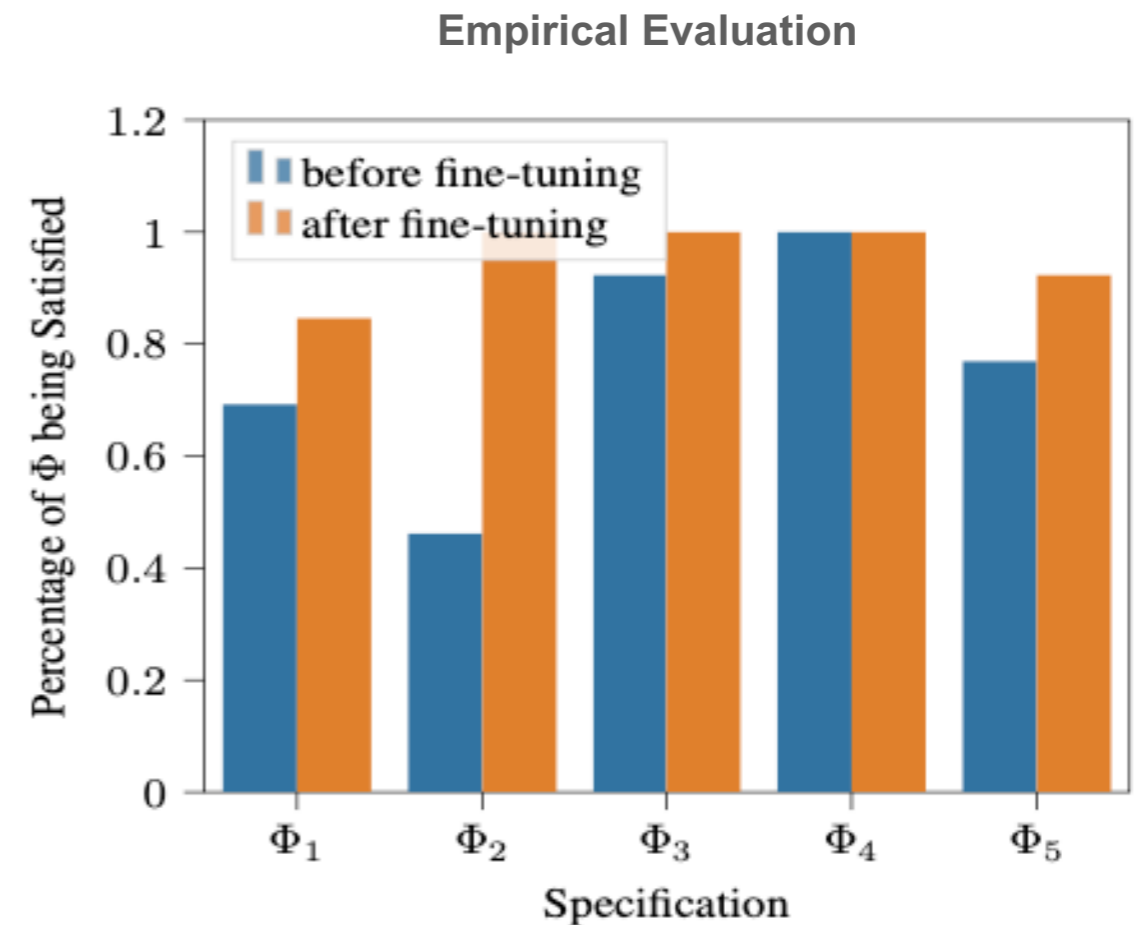
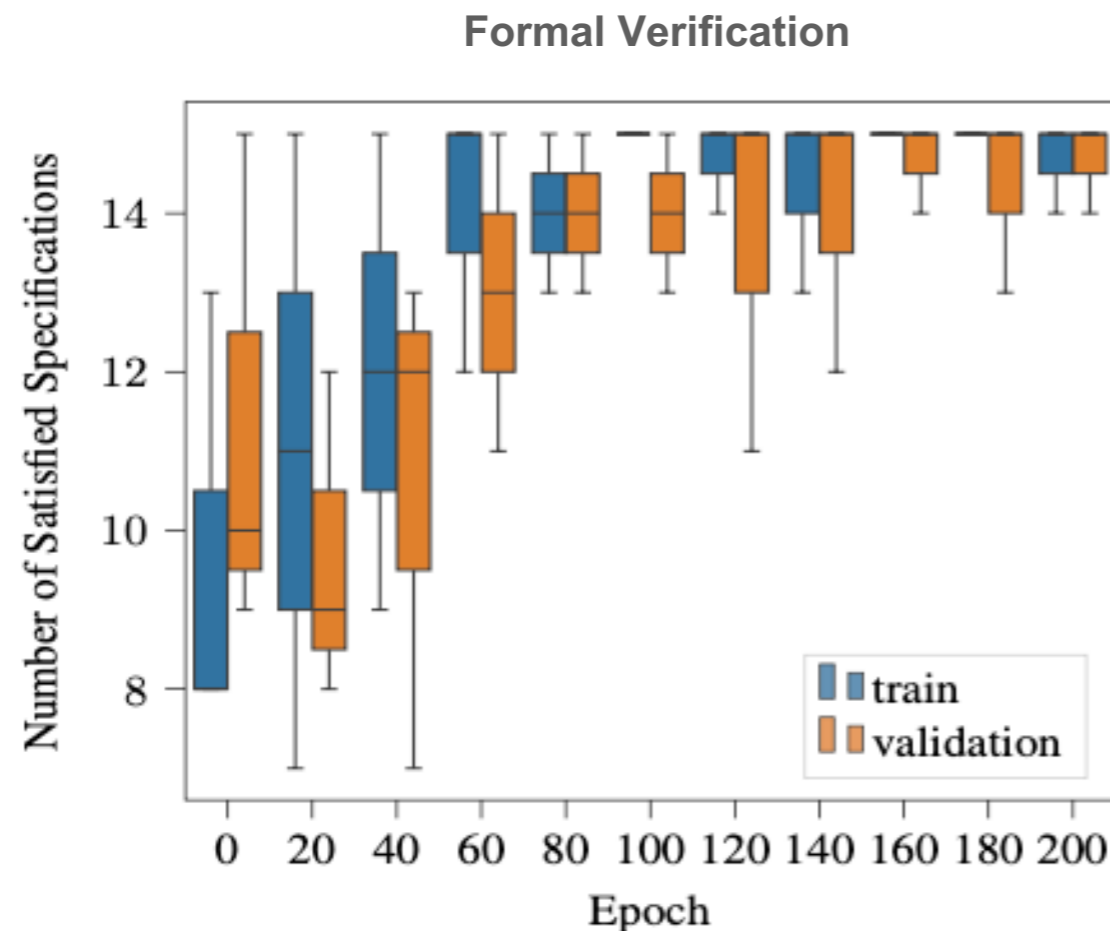
Execution Trace: (desired objects with positions, action),.....

Desired objects: pedestrian, car, red/green traffic light, stop sign,.....

Actions: go straight, turn left, stop, turn right,.....

Quantitative Analysis

Empirical Evaluation via Simulation



The results indicate that our approach can improve the language model's ability to satisfy critical requirements.

Our approach can act as a starting point to guide the design process for real-world implementations of autonomous driving systems.

Next Step: Verifiable Grounding

Ground the controller to real autonomous driving robots

```
class AutonomousVehicle():
    def __init__(self):
        self.pedestrian = False
        self.car = False

    @abstractmethod
    def move_forward(self):
        # vehicle starts moving forward
        pass

    @abstractmethod
    def turn_left(self):
        # vehicle turns left
        pass

    @abstractmethod
    def turn_right(self):
        # vehicle turns right
        pass

    @abstractmethod
    def stop(self):
        # vehicle slows down and stops
        pass
```

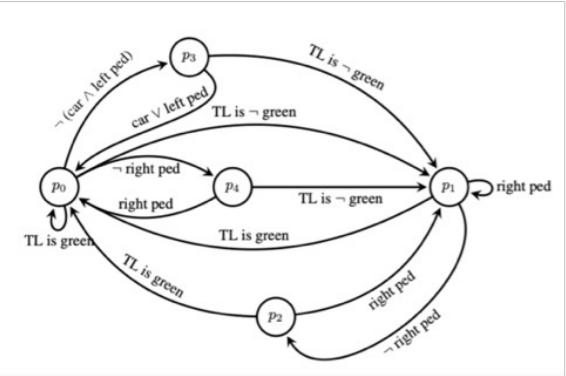
(left car \vee ped at right, ϵ)
(True, check left car)

```
def CrossRoad(vehicle):
    vehicle.observe() # observe the environment first
    if vehicle.pedestrian_observed(): # Step 1
        vehicle.stop() # stop if pedestrian is observed
        return "Stop"
    elif vehicle.car_observed(): # Step 2
        vehicle.stop() # stop if car is observed
        return "Stop"
    else: # Step 3
        vehicle.move_forward()
        return "Move_forward"
```

- (\neg green traffic light \rightarrow \neg go straight), ✓
- (stop sign \rightarrow \diamond stop), ✓
- \neg turn right \vee \neg (car from left \vee pedestrian at right), ✓

Q&A

Abstract Model



High-Fidelity Simulator



User

Autonomous System Models

Specifications

Control Tasks



Automated Feedback



Prompt Dataset

Feed To DPO



Direct Preference Optimization



Fine-Tune the Language Model

Language Model

Query the Language Model

Select A Prompt



Prompt