

# Lecture 17: Decision-Making cont. & Formal Safety

Professor Katie Driggs-Campbell

April 1, 2021

ECE484: Principles of Safe Autonomy



# Administrivia - Schedule

- MP3 due next week 4/9
  - Demo on Thursday / Report due Friday
- MP4 due Friday 4/30
- Next week: Final Safety Lecture and MP4 Walkthrough
- No class on Tuesday 4/13
- No class on Tuesday 5/4 – extra office hours
- Guest Lectures on 4/15, 4/20, and 4/22
  - Attendance worth double
  - Google form to collect questions will be posted on discord
- Project presentations on 4/27 and 4/29
  - Information will be provided on 4/8 during lecture
- Report due on Friday 5/14 (last day of finals)
  - Rubric and guidelines to be posted soon



# Administrivia - Oral Exam Protocol

- **Dates:** Monday 4/12 to Tuesday 5/4
- **Times:** 15-minute time slots on Mondays 5-6pm, Tuesdays 5-6pm, Wednesdays 11am-12pm, or by appointment (if needed)
  - Signup will be posted on discord – first come, first serve
- **Content:** Questions and rubric (with bonus allocation) released this weekend
  - You will be asked one of six questions, chosen by rolling a die
- **Protocol:** We will follow CBTF protocol
  - Log in with both computer and phone
  - Position phone so we can see screen and workspace
- **A word of caution:** If we suspect cheating, we will schedule a final exam



# Today's Plan

- MDP Policies and Value Iteration
- Simple Example
- Introduction to Safety
- Responsibility Sensitive Safety



# MDP Policies

- Policies map states to actions

$$\pi: x \rightarrow u$$

- We want to find a policy that maximizes future pay off

- Suppose  $T = 1$ :  $\pi_1(x) = \operatorname{argmax}_u \underline{r(x, u)}$

- We write the Value Function for given  $\pi$ :


$$V_1(x) = \gamma \max_u r(x, u)$$

- Generally, we want to find the sequence of actions that optimize the *expected cumulative discounted future payoff*

$$R_T = \mathbb{E} \left[ \sum_{\tau=0}^T \gamma^\tau \underline{r_{t+\tau}} \right]$$



# Value Functions



Recall:  $V_1(x) = \gamma \max_u r(x, u)$

For longer time horizons, we define  $V(x)$  recursively:

$T=2$ : pick an action that max sum  $V_1 + 1$ -step <sup>~</sup>payoff

\*  $\pi_2(x) = \operatorname{argmax}_u [r(x, u) + \int V_1(x') p(x'|x, u) dx']$

\*  $V_2(x) = \gamma \max_u [r(x, u) + \int V_1(x') p(x'|x, u) dx']$

For finite  $T$ :

$\pi_T(x) = \operatorname{argmax}_u [r(x, u) + \int V_{T-1}(x') p(x'|x, u) dx']$

$V_T(x) = \gamma \max_u [r(x, u) + \int V_{T-1}(x') p(x'|x, u) dx']$



# Value Functions

- In the infinite time horizon, we tend to reach equilibrium:

$$\underline{V_\infty}(x) = \gamma \max_u \left[ r(x, u) + \int \underline{V_\infty}(x') p(x' | x, u) dx' \right]$$

- This is the *Bellman Equation*
  - Satisfying this is necessary and sufficient for an optimal policy



# Computing the (Approximate) Value Function

- Initial guess for  $\hat{V}$ 
  - $\hat{V}(x) \leftarrow r_{min}, \forall x$
- Successively update for increasing horizons
  - $\hat{V}(x) \leftarrow \gamma \max_u [r(x, u) + \int \hat{V}(x') p(x'|x, u) dx']$
- Value iteration converges if  $\gamma < 1$
- Given estimate  $\hat{V}(x)$ , policy is found:
  - $\pi(x) = \operatorname{argmax}_u [r(x, u) + \int \hat{V}(x') p(x'|x, u) dx']$
- Often, we use the discrete version:
  - $\pi(x) = \operatorname{argmax}_u [r(x, u) + \sum_{x'} \hat{V}(x') p(x'|x, u)]$  ✓






# Example: Create an MDP

$$(S, A, T, \gamma, R)$$

stop



\* terminal

$$S = \{x_1, x_2, x_3\}$$

$$A = \{\uparrow, 0\}$$

$P(x'|x, u)$  is deterministic

$$R = \begin{cases} r_s & \text{if } x_i + u = 0 \\ 0 & \text{otherwise} \end{cases}$$

$\uparrow$  will move up  $P=1$   
 $0$  will stop  $P=1$



# Example: Value Iteration

$$\hat{V}_{iter}(x) = \gamma \max_u \left[ r(x, u) + \sum_j \hat{V}_{iter-1}(x_j) p(x_j | x, u) \right]$$

vector rep:  $\hat{V} = [\hat{V}(x_1) \quad \hat{V}(x_2) \quad \hat{V}(x_3)]$

→ init:  $\hat{V} = [0 \quad 0 \quad 0]$

---

$\hat{V}_1(x_1) = \gamma \max_u (0 + 0, r_s + 0) = \gamma r_s$   $\hat{V}_1(x_2) = \gamma \max_u (0 + 0, 0 + 0) = 0$  }  $\hat{V} = [\gamma r_s \quad 0 \quad 0]$

$\hat{V}_1(x_3) = 0$

---


$\hat{V}_2(x_1) = \gamma r_s$   
 $\hat{V}_2(x_2) = \gamma \max_u (0 + \gamma r_s, 0 + 0) = \gamma^2 r_s$  }  $\hat{V}_2 = [\gamma r_s \quad \gamma^2 r_s \quad 0]$   
 $\vdots$   
 $\hat{V}_3 = [\gamma r_s \quad \gamma^2 r_s \quad \gamma^3 r_s]$


$\hat{V}_2(x_3) = 0$





# Grid world example


- States: cells in  $10 \times 10$  grid
- Actions: up, down, left, right
- Transition model: 0.7 chance of moving in intended direction, uniform in other directions
- Reward:
  - two states with cost
  - two terminal states with rewards
  - -1 for wall crash
- Discount is 0.9




































Converged!

$\gamma = 0.9$  $\gamma = 0.5$ 

# Decision-Making Summary

- Given an MDP, we defined **Expected Cumulative Payoff**, which plays a key role in **optimizing actions over planning horizons**
- Used **value iteration** to determine the “value” of a particular state, which helps us determine the best action to take considering future payoff
- We generally assumed the transition and reward function are known exactly – but what if we don’t have access to this information?
  - Will post notes on basic Q-learning for RL!





# Today's Plan

- MDP Policies and Value Iteration
- Simple Example
- Introduction to Safety
- Responsibility Sensitive Safety




# (Safety) Challenges for AVs

- Each module is challenging to develop
- AVs are safety critical systems, but:
  - It is impossible to guarantee zero accidents
  - Statistical approaches are problematic
  - Any software or hardware update requires new certification
- Scalability remains a challenge
  - How to mass produce expensive computation and sensors?
  - How to ensure you can drive everywhere?



# How to ensure multi-agent safety?

Absolute safety is impossible ☹️



Statistical guarantees are infeasible ☹️

Typically, to show  $p_1$  likelihood of failure, we aim to gather at least  $\frac{1}{p_1}$  samples

- Are all miles equal?
- 10 disengagements per 300 million miles is not enough
- Any change to the software will require a new data collection



# On-Road Testing

- Data required to guarantee  $10^{-9}$  probability of failure:  $10^9$  hours of driving or 30 billion miles
  - However, this is insufficient to show the difference between an AV with error rate of  $p_1$  and a perfect system
    - We want system with  $10^{-9}$  accident rate per hour, but we have an AV system with  $10^{-8}$  accident rate per hour.
    - If we drive for  $10^8$  hours, there is still a constant probability that the testing process is not telling us the true accident rate
- Purely data-based approach for safety is naïve at best

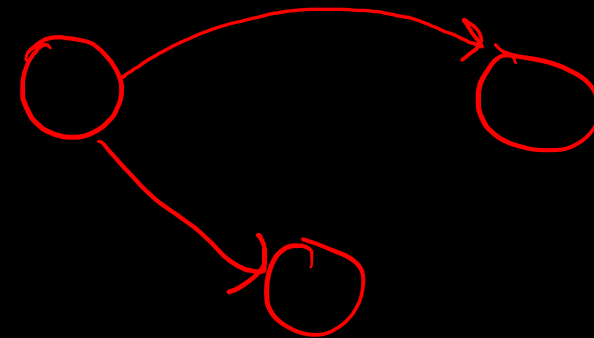


# Safety Guarantees Automata

- An **automata** is a mathematical model for describing computations or processes evolving in discrete steps
- **States** can be discrete or continuous valued
- State **transitions** define how the states can change
  - Transitions can be **non-deterministic**: multiple next states from a single state
- No inherent notion of *time*, but each transition can be thought of as passage of a fixed amount of time

An automaton is a tuple  $\mathcal{A} = \langle Q, \Theta, A, \mathcal{D} \rangle$  where

- $Q$  is a set of states
- $\Theta \subseteq Q$  is the set of initial or start states
- $A$  is a set of actions or labels
- $\mathcal{D} \subseteq Q \times A \times Q$  is the set of transitions
  - A transition can be thought of as a triple  $(u, a, u')$

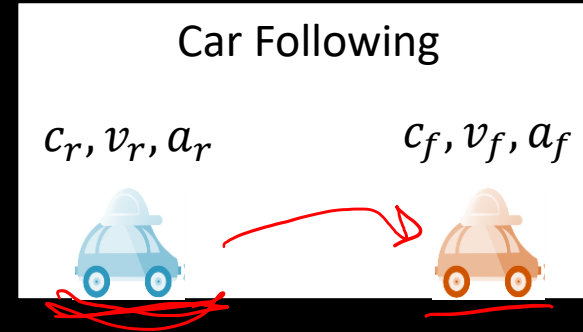


# Safety Guarantees: Automata

- $a_f[t] = a_{min}$  ✓

- •  $v_f[t + 1] = v_f[t] + a_f[t]\Delta$

- •  $c_f[t + 1] = c_f[t] + v_f[t]\Delta + \frac{1}{2}a[t]\Delta^2$



- $a_r[t] \in [a_{min}, a_{max}]$

- $v_r[t + 1] = v_r[t] + a_r[t]\Delta$

- $c_r[t + 1] = c_r[t] + v_r[t]\Delta + \frac{1}{2}a_r[t]\Delta^2$

- $X = \{a_f, v_f, c_f, a_r, v_r, c_r\}$  ✓

- $Q = \mathbb{R}^6$

- $A = (a_1, a_2)$  the acceleration choices ✓

- $\mathcal{D} \subseteq \mathbb{R}^6 \times A \times \mathbb{R}^6$  ✓



# Executions, Reachability, & Invariants

- An **execution** of  $\mathcal{A}$  is an alternating (possibly infinite) sequence of states  $s_t$  and actions  $u_t$ :  $\alpha = s_0 u_1 s_1 u_2 s_2 \dots$  such that:
    - $s_0 \in \Theta$
    - $\forall i$  in the sequence the transition function is applied ( $s_i \xrightarrow{u_{i+1}} s_{i+1}$ )
  - A state  $s$  is **reachable** if there exists an execution that ends at  $s$
- The set of reachable states is denoted by  $Reach_A$

$$s_{i+1} = f(s_{i+1}, u_{i+1})$$



# Formal Invariants

- What does it mean for  $I$  to hold “always” for  $\mathcal{A}$ ?
  - $I$  holds at all states along any execution  $s_0 u_1 s_1 u_2 s_3$
  - $I$  holds in all reachable states of  $\mathcal{A}$
  - $Reach_{\mathcal{A}} \subseteq [[I]]$
- Invariants capture most properties that you will encounter in practice
  - safety: “aircraft always maintain separation”
  - bounded reaction time: “within 15 seconds of press, light must turn to walk”
- How to verify if  $I$  is an invariant?
  - Does there exist reachable state  $s$  such that  $s \not\models I$ ?






# Reachability Problem

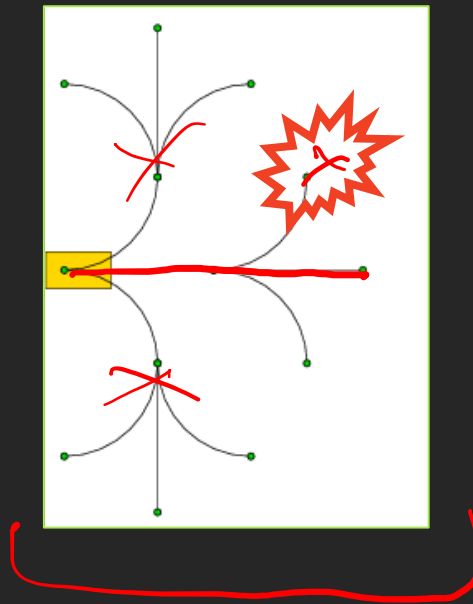
## Automata

Given a directed graph  $G = (V, E)$ , and two sets of vertices  $S, T \subseteq V$ ,  $T$  is **reachable from**  $S$  if there is a path from  $S$  to  $T$ .


**Reachability Problem**  $(G, S, T)$ :  
is  $T$  is reachable from  $S$  in  $G$ ?



## Reachability Trees




## Reachable Sets



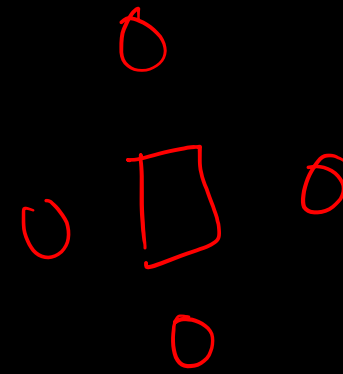
# City Safety: Emergency Braking

Image Credit: Volvo



# Responsibility Sensitive Safety

developed by Intel / MobilEye



Instead of looking at absolute safety, introduce a safety notion that depends on *responsibility*

→ AV should never be responsible for an accident

## RSS Rules:

1. Keep a safe longitudinal distance from the car ahead.
2. Keep a safe lateral distance from the cars on your sides.
3. Respect right-of-way rules (multiple geometries, traffic lights, pedestrians, unstructured roads).
4. Be cautious of occluded areas.



# RSS Example: Safe Following Distance

The longitudinal distance between a car ( $c_r$ ) that drives behind another car ( $c_f$ ) is safe w.r.t. a response time  $\rho$  if:

- $c_f$  applies at most  $a_{max}^{brake}$  ✓
  - $c_r$  will apply at most  $a_{max}^{accel}$  during response time ✓
  - After  $\rho$ ,  $c_r$  will brake by at least  $a_{min}^{brake}$  until full stop
- $c_r$  will not collide with  $c_f$

Remarks:

1. This is basic reachability!
2. The safe distance depends on a set of parameters that can be determined by regulation.
3. The parameters can be different for a robotic car and a human driver.
4. The parameters can be different for different road conditions.



# RSS Example: Safe Following Distance

- Let  $v_r, v_f$  be the longitudinal velocities of the cars
- The minimal safe distance is:

$$d_{min} = \left[ v_r \cdot \rho + \frac{1}{2} a_{max}^{accel} \cdot \rho^2 + \frac{(v_r + \rho \cdot a_{max}^{accel})^2}{2a_{min}^{brake}} - \frac{v_f^2}{2a_{max}^{brake}} \right]_+$$



# Summary

- Discussed the challenges with defining safety and introduced the ideas behind verification and formal guarantees for safety
  - **Reachability** is as key tool for guaranteeing specified properties
- While useful, often require well-defined **specifications**, require high-fidelity **models (and assumptions)**, consider **worst case** scenarios, and/or only assess one piece of the puzzle
  - Making formal guarantees on the full autonomous stack remains a challenge
- Introduced **RSS**, a framework that aims to define safety in an intuitive manner



# Extra Slides




# Certifying airworthiness of aviation software

What fraction of the cost of developing a new aircraft is in SW?

**DO178C**

Primary document by which FAA & EASA approves software-based aerospace systems.

DAL establishes the rigor necessary to demonstrate compliance



Dev.Assurance Level (DAL)	Hazard Classification	Objectives
A	Catastrophic	71
B	Hazardous	69
C	Major	62
D	Minor	26
E	No Effect	0

**Statement Coverage:** Every statement of the source code must be covered by a test case

**Condition Coverage:** Every condition within a branch statement must be covered by a test case

**“Special credits”:** For using formal methods based tools recently introduced





# Safety Certification (cont)

- A component's DAL level is determined from a *safety assessment process and hazard analysis through examination of the effects of a failure condition in the system. The failure conditions are categorized by their effects on the aircraft, crew, and passengers.*
  - For example, Level A is assigned for "Catastrophic Outcome," and Level E is for "No Safety Effect."
- DAL level then establishes the rigor necessary to demonstrate compliance with DO-178C
- E.g., components that command, control, or monitor safety critical functions are Level A. The standard requires any Level A software to be tested to cover every statement, branch, and function call, and also to pass the Modified Condition Decision Coverage (MC/DC) tests
  - Requires that (i) each entry and exit point in the code be invoked, (ii) each decision take every possible value, and (iii) each condition in a decision take every possible value
  - For certain levels, DO-178C requires that the testing, verification, and validation be performed by a team that is independent of the software development team
- Again, certification is process-based and does not eliminate possibility of bad logic, interference in the the control code
- Dozens of commercial tools (e.g., MATLAB, Esterel, Cantata, VectorCAST, Rapita Systems, and CodeSonar) can support DO-178C certification by applying **formal verification**.



# The formal verification problem

Example requirements:

Safety: “For all nominal behaviors of the car, the separation between the cars must be always  $> 1$  m”

Efficiency: “For all nominal driver inputs, the air-fuel ratio must be in the range  $[1,4]$ ”

Privacy: “Using GPS does not compromise user’s location”

Fairness: “Similar people’s loan approval are decided similarly”

