

Last Time:

- formalizing safe sets
- computing safety filters

lecture 4

EAIS S'25

Andrea Bajcsy


This Time:

- robustifying safety - why?
- zero-sum dynamic games
- dynamic programming for games

Robustifying Safety

So far, we have assumed that our dynamical system perfectly evolves via $\dot{x} = f(x, u)$ with no uncertainty.

This isn't realistic for many real-world scenarios (e.g. friction!)




There are two ways to model uncertainty:

- 1) probabilistic uncertainty (i.e. "I have observed data")
- 2) non-deterministic uncertainty (i.e. "I have minimal additional info.")

How should we handle the design of our safety filter (+ analysis) to handle uncertainty?

Typically, we do this via modelling another "input" that influences the state evolution:



In robust safety, we take a non-deterministic view of uncertainty and assume that $d \in \mathcal{D}$ is chosen from some bounded set and we want our robot to be ROBUST to the WORST POSSIBLE sequence of d 's!

ROBUST BACKWARDS REACHABLE TUBE (BRT) of a set $\mathcal{F} \subset \mathcal{X}$

and dynamical system $\dot{x} = f(x, u, d)$ is:

$$\text{BRT}(t) := \{x \in \mathcal{X} : \forall u(\cdot) \in \mathcal{U}_t^T, \exists d(\cdot) \in \mathcal{D}_t^T, x_{x_t}^{u, d}(\tau) \in \mathcal{F} \text{ for some time } \tau \in [t, T]\}$$

for all things the robot could do there is something the disturbance could do ...

that leads the robot to failure

This is the set of all starting states from which no matter the controller's effort, the disturbance can push system into \mathcal{F} . The way we will formulate an "optimal control" problem whose solution represents this unsafe set will be via:

<u>zero-sum</u>	<u>dynamic</u>	<u>games</u>
there is a winner + loser	game <u>evolves</u> over time	result / outcome depends on 2+ players / inputs

Our "game" formulation could look something like this:

$$V(x, t) = \max_{u(\cdot) \in \mathcal{U}_t^T} \min_{d(\cdot) \in \mathcal{D}_t^T} J(x, u(\cdot), d(\cdot))$$

s.t. $\dot{x}(\tau) = f(x(\tau), u(\tau), d(\tau)) \quad \forall \tau \in [t, T]$

$$u(\tau) \in \mathcal{U}$$
$$d(\tau) \in \mathcal{D}$$

Here, the robot (u) is trying to maximize the objective $J(\cdot, \cdot, \cdot)$ and the other player (d) is minimizing. In other words, the robot is optimizing the worst-case objective.

Importance of Information Patterns

When we have 2 players reacting to each other, their optimal strategy will depend on what information they each have access to.

Example Suppose there are 2 boxes, each with 2 slots. Each slot contains prize money. Player A (You u) wants to maximize prize money while Player B (competition organizer) wants to minimize Player A's prize money.

Box 1	
s1	s2
2,000	10

Box 2	
s1	s2
1	6,000


Player A

$u \in [Box 1, Box 2]$
you choose Box \rightarrow

Player B

$d \in [slot 1, slot 2]$
they choose slot \uparrow

Suppose Player A goes first:



Best outcome for A is to pick Box 1 & get Rew = \$10.


Mathematically :

$$\max_{u \in [B1, B2]} \left(\min_{d \in [s1, s2]} J(u, d) \right) = 10$$

"computed 2nd" "computed first" but "plays 2nd"

but "plays first" "best response" strategy

Suppose Player B goes first:



Best strategy for player B is to pick slot 1 and pay player A a reward of \$ 2,000.

$$\min_{d \in [s_1, s_2]} \left(\max_{u \in [B_1, B_2]} J(u, d) \right) = \$2,000$$

This phenomenon we just saw can be stated via

minimax inequality:

$$\max_a \left(\min_b J(a, b) \right) \leq \min_b \left(\max_a J(a, b) \right)$$

→ Von Neumann 1928: equal when A, B are compact, convex sets +

$J(\cdot, b)$ is concave for fixed b

$J(a, \cdot)$ is convex for fixed a

In dynamic games the outcome depends on WHEN and


WITH WHAT INFORMATION

each player decides their inputs

↓ now.

OPEN-LOOP INFO PATTERN:


$$\max_{u(\cdot) \in \mathcal{U}_t^T} \min_{d(\cdot) \in \mathcal{D}_t^T} J(x, u(\cdot), d(\cdot))$$



$u(\cdot)$ "declares" entire signal and then $d(\cdot)$ gets to see this and interrupts!

⊕ OVERLY PESSIMISTIC!

Here u wants to reach G_1 or G_2 without being intercepted by d . BUT d wants to intercept u .




$J \rightarrow$ "dist btwn. agents"


Let's swap the order of play again, but we are playing over continuous signals (or sequences of decisions):

OPEN-LOOP INFO PATTERN:

$$\min_{d(\cdot) \in D_t^T} \max_{u(\cdot) \in U_t^T} J(x, u(\cdot), d(\cdot))$$



$d(\cdot)$ "declares" entire signal and then $u(\cdot)$ gets to see this and picks the other goal!



↗ OVERLY OPTIMISTIC!

CLOSED-LOOP INFORMATION PATTERNS


The above formulation is not suitable for many practical systems. We would like the controller (u) to adapt over time as system evolves, but respect the fact that @ any time t , we only have information up to time t .

We can model this by solving for feedback policies:

$$V(x, t) = \max_{\pi_u} \min_{\pi_d} J(x, \pi_u(\cdot), \pi_d(\cdot))$$


$\pi_u: x \rightarrow u$ $\pi_d: x \rightarrow D$

Before $(\max_{u(\cdot) \in \mathcal{U}_t^T} \min_{d(\cdot) \in \mathcal{D}_t^T} \dots)$



⇒ this is open-loop

Now (\max_{T_u} \min_{T_d} ...)




for each initial state, the players commit to a policy - or a map from each state to their action.

⇒ this is closed-loop

★ NOTE: for this illustrative example, since u and d are 2 separate agents, the state space is $X := X_u \times X_d$ and policies map from where BOTH players are to actions:

$$\pi_u(x_u, x_d) \rightarrow u, \quad \pi_d(x_u, x_d) \rightarrow d$$

Bringing this back to safety analysis, we just have to choose a objective function which lets us remember the closest we ever got to failure.



$$J(x, \pi_u, \pi_d, t) := \min_{\tau \in [t, T]} l(x_{\tau}^{\pi_u \pi_d}(\tau))$$

minimum future distance influenced by u and d

We can apply the game-theoretic principle of optimality called the Tenet of Transition by Rufus Isaacs:

"If play proceeds from one position (state) to a second, and V is thought of as known to the second, then it is determined at the first by demanding that players optimize (i.e. make minimax) the increment of V during the transition."

example (discrete-time)

$$\begin{aligned}
 V_t(x_t) &:= \min_{\pi_u(x)} \max_{\pi_d(x)} \min_{\tau \in \{t, \dots, T\}} \mathcal{L}(x_{\tau}^{u, d}) \quad \xrightarrow{\text{trajectory:}} \begin{array}{c} x_{t+1} \\ \curvearrowleft \quad \curvearrowleft \quad \curvearrowleft \\ x_t \quad x_{t+2} \quad x_T \end{array} \quad \text{w/ } \pi_u, \pi_d \\
 &= \min_{u^t} \max_{d^t} \dots \min_{u^T} \max_{d^T} \min_{\tau \in \{t, \dots, T\}} \mathcal{L}(x_{\tau}^{u, d}) \\
 &= \min_{u^t} \max_{d^t} \min \left\{ \mathcal{L}(x_t), \min_{u^{t+1}} \max_{d^{t+1}} \dots \min_{u^T} \max_{d^T} \min_{s \in \{t+1, \dots, T\}} \mathcal{L}(x_s^{u, d}) \right\} \\
 &\quad := V_{t+1}(f(x_t, u_t, d_t)) \text{ by } \underline{\text{Tenet of Transition}} \\
 &= \min_{u^t} \max_{d^t} \min \left\{ \mathcal{L}(x_t), V_{t+1}(f(x_t, u_t, d_t)) \right\} \\
 &= \min \left\{ \mathcal{L}(x_t), \min_{u^t} \max_{d^t} V_{t+1}(f(x_t, u_t, d_t)) \right\}
 \end{aligned}$$

Applying this to the continuous or discrete-time zero sum safety critical games we formulated above, we get two key equations that allow us to extend dynamic programming tools to the robust (dynamic game) setting:

Hamilton-Jacobi-Isaacs Variational Inequality (HJI-VI)

$$\min \left\{ \mathcal{L}(x) - V(x, t), \frac{\partial V}{\partial t} + \max_{u \in U} \min_{d \in D} \frac{\partial V}{\partial x} \cdot f(x, u, d) \right\} = 0$$

$$V(x, T) = \mathcal{L}(x)$$

↑ add this to solve "subgame" over just instantaneous actions

ROBUST SAFETY BACKUP (discrete-time):

$$V_t(x_t) = \min \left\{ \mathcal{L}(x_t), \max_{u_t} \min_{d_t} V_{t+1}(f(x_t, u_t, d_t)) \right\}$$

$$V_T(x_T) = \mathcal{L}(x)$$

SUMMARY

Continuous-Time

Problem

$$V(x, t) := \max_{\pi_u(x)} \min_{\tau' \in [t, T]} \mathcal{L}(x^u_{\tau'}, \tau')$$

$\min_{\pi_d(x)}$ for robust

Dynamic Programming "Backup"

$$\min \left\{ \mathcal{L}(x) - V(x, t), \frac{\partial V(x, t)}{\partial t} + \max_{u \in U} \frac{\partial V(x, t)}{\partial x} \cdot f(x, u) \right\} = 0$$

$\min_{d \in D}$ for robust

Discrete - Time

Problem

$$V_t(x_t) := \max_{\pi_u(x)} \min_{\tau' \in \{t, \dots, T\}} \mathcal{L}(x^u_{\tau'}, \tau')$$

$\min_{\pi_d(x)}$ for robust

Dynamic Programming "Backup"

$$V_t(x_t) = \min \left\{ \mathcal{L}(x_t), \max_{u_t \in U} V_{t+1}(f(x_t, u_t, d_t)) \right\}$$

$\min_{d \in D}$ for robust

Numerical Comparison of Robust vs. Non-Robust Unsafe Set

Recall: state is (x, y, θ) of Dubins' car

dynamics are

$$\begin{aligned}\dot{x} &= v \cos \theta & \in [-0.5, 0.5] \\ \dot{y} &= v \sin \theta & \in [-0.5, 0.5] \\ \dot{\theta} &= u & \in [-0.5, 0.5]\end{aligned}$$

$$\dot{x} = v \cos \theta + \frac{dx}{dt} \equiv D$$


$$\dot{y} = v \sin \theta + \frac{dy}{dt} \equiv D$$

$$\dot{\theta} = u \equiv u$$

failure is cylinder @ origin with radius 0.3

Non-Robust BRT

$$\max_u \min_{\tau \in [t, T]} \ell(x(\tau))$$



Robust BRT

$$\max_{\pi_u} \min_{\pi_d} \min_{\tau \in [t, T]} \ell(x(\tau))$$

