A Probabilistic Framework for Pursuit-Evasion Games with Unmanned Air Vehicles

Maria Prandini
Univ. of Brescia & UC Berkeley

In collaboration with J. Hespanha, J. Kim, and S. Sastry

Key Ideas

- The "mission-level" control of Unmanned Air Vehicles requires a ***probabilistic framework***.
- The problem of coordinating teams of autonomous agents is naturally formulated in a ***game theoretical*** setting.
- ***Exact solutions*** for these types of problems are often computationally intractable and, in some cases, open research problems.

The “rules” of the game

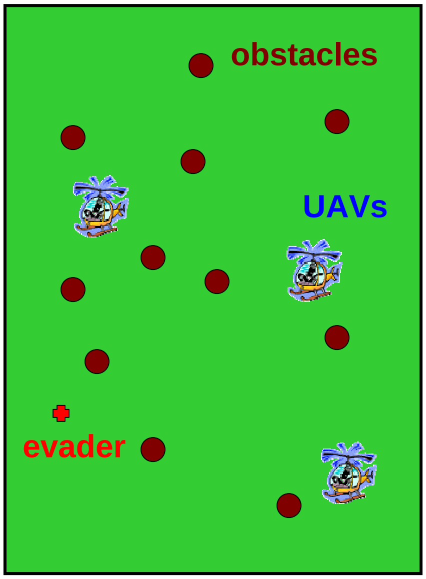

The “rules” of the game

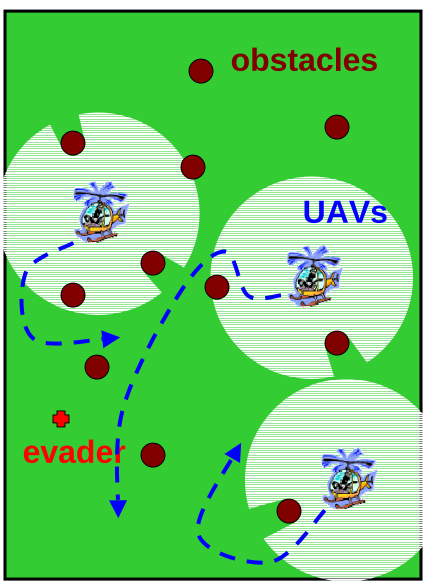

Terrain:
* with **fixed obstacles**
* not accurately mapped

UAVs (pursuers) capable of:
* **flying** between obstacles
* **seeing** a region around them (limited by the occlusions)

Evader capable of:
* **moving** between obstacles (possibly actively avoiding detection)

Objective:
find the evader in minimum time

Scenarios

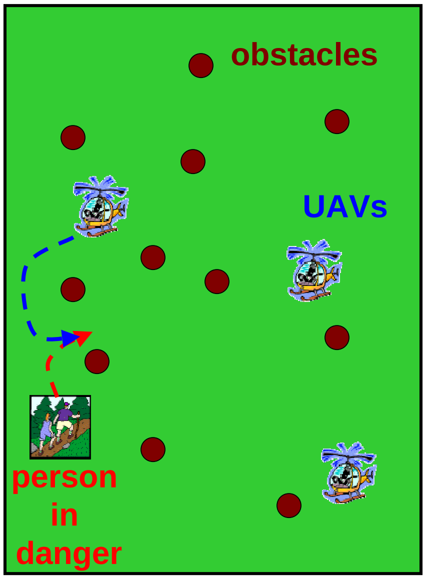

- search and rescue operations

Scenarios

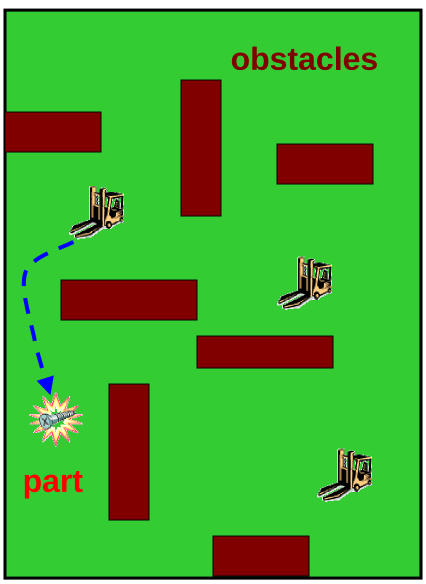

- search and rescue operations
- finding parts in a warehouse

Scenarios

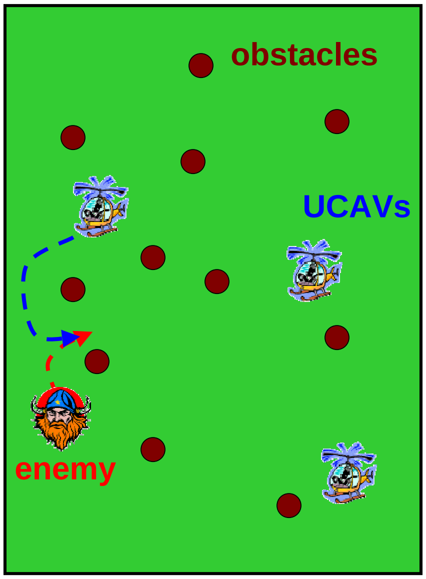

- search and rescue operations
- finding parts in a warehouse
- search and capture operations

Scenarios

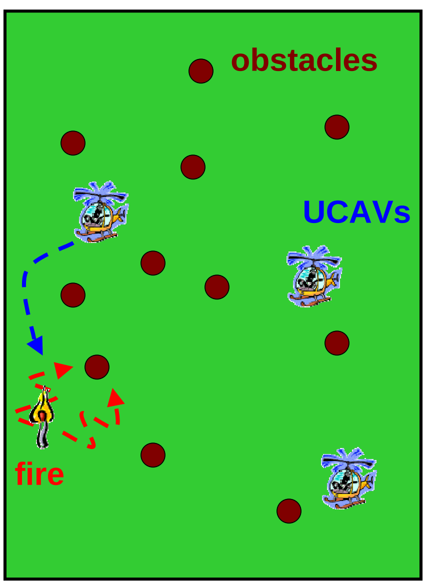

- search and rescue operations
- finding parts in a warehouse
- search and capture operations
- monitoring environmental threats

Strategies for pursuit-evasion games

LaValle, Latombe, Guibas, *et al*. considered a similar problem but assume the map of the region is **known**, the pursuers have **perfect sensors**, and **worst case** trajectories for the evader:

How many *UAVs are needed to win the game in finite time?*

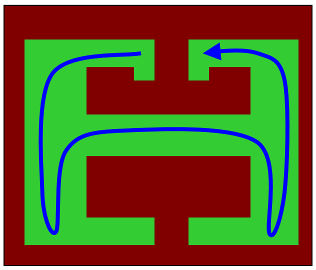

1 agent is sufficient

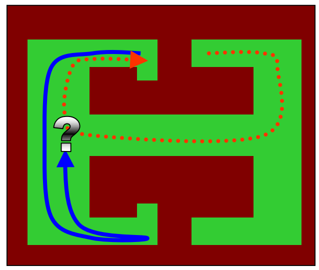

2 agents are needed

(no matter what strategy a single pursuer chooses, there is a trajectory for the evader that avoids detection)

Exploring a region to build a map

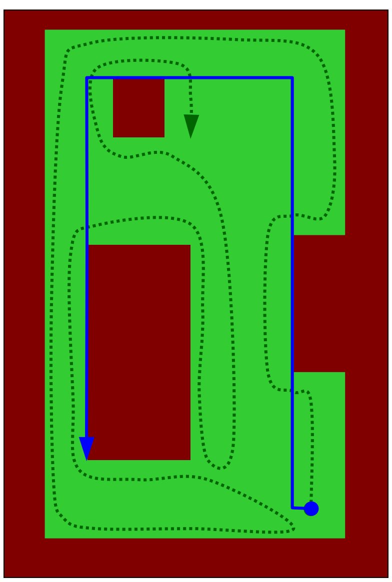

Deng, Papadimitriou, *et al.*, study the problem of building a map (seeing all points in the region) traversing the smallest possible distance.

standard "keep wall to the right" algorithm

algorithm that takes better advantage of the cameras capabilities

A two step solution…

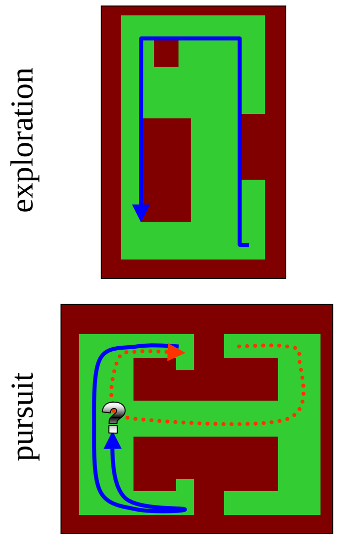

- exploration followed by pursuit is not **efficient**
- sensors are **imprecise**
- **worst case** assumptions the trajectories of the evaders leads to very conservative results

A different approach…

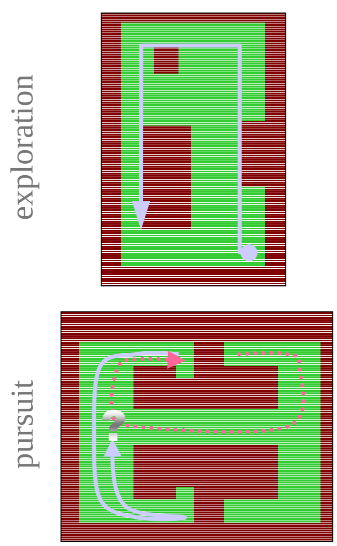

Use a **probabilistic framework** to **combine** exploration and pursuit-evasion games.

Non determinism comes from:

- poorly mapped **terrain**
- noise and uncertainty in the **sensors**
- probabilistic models for the **motion** of the evader and the UAVs

Markov Decision Processes

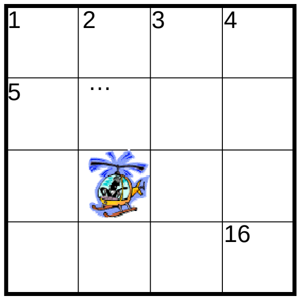

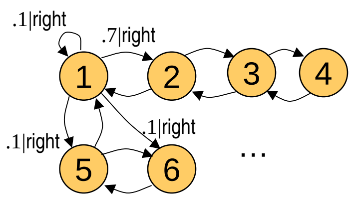

time $t \in \{1, 2, 3, \ldots\}$

state $\boldsymbol{x}_t \in \mathcal{X} := \{1, 2, \ldots, 16\}$

action $\boldsymbol{u}_t \in \mathcal{U} := \{$**up**, **down**, **left**, **right**$\}$

transition probability function

$$p(x, x', u) = \mathrm{P}(\boldsymbol{x}_{t+1} = x' \mid \boldsymbol{x}_t = x, \boldsymbol{u}_t = u)$$

Markov Decision Processes

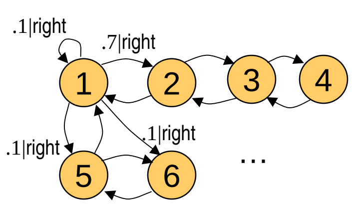

time $t \in \{1, 2, 3, \ldots\}$

state $\boldsymbol{x}_t \in \mathcal{X} := \{1, 2, \ldots, 16\}$

action $\boldsymbol{u}_t \in \mathcal{U} := \{$**up**, **down**, **left**, **right**$\}$

transition probability function

$$p(x, x', u) = \mathrm{P}(\boldsymbol{x}_{t+1} = x' \mid \boldsymbol{x}_t = x, \boldsymbol{u}_t = u)$$

set of all sequences of elements in X

control policy (deterministic) $\nu : \mathcal{X}^* \to \mathcal{U}$

$$\boldsymbol{u}_t = \nu(\underbrace{\boldsymbol{x}_t, \boldsymbol{x}_{t-1}, \boldsymbol{x}_{t-2}, \ldots, \boldsymbol{x}_1}) = \nu(\boldsymbol{X}_t)$$

control policy (stochastic) $\mu : \mathcal{U} \times \mathcal{X}^* \to [0,1]$

$$\mathrm{P}(\boldsymbol{u}_t = u \mid \boldsymbol{X}_t = X, \mu) = \mu(u, X)$$

Markov Decision Processes

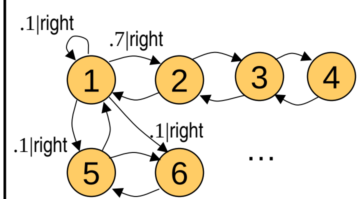

time $t \in \{1, 2, 3, \dots\}$

state $\boldsymbol{x}_t \in \mathcal{X} := \{1, 2, \dots, 16\}$

action $\boldsymbol{u}_t \in \mathcal{U} := \{$**up**, **down**, **left**, **right**$\}$

transition probability function

$$p(x,x',u) = \mathrm{P}(\boldsymbol{x}_{t+1}=x' \mid \boldsymbol{x}_t=x, \boldsymbol{u}_t=u)$$

set of all sequences of elements in X

control policy (deterministic) $\nu : \mathcal{X}^* \to \mathcal{U}$

$$\boldsymbol{u}_t = \nu(\underbrace{\boldsymbol{x}_t, \boldsymbol{x}_{t-1}, \boldsymbol{x}_{t-2}, \dots, \boldsymbol{x}_1}) = \nu(\boldsymbol{X}_t)$$

(almost surely) $\Updownarrow$

$$\mathrm{P}(\boldsymbol{u}_t = u \mid \boldsymbol{X}_t = X, \nu) = \begin{cases} 1 & \nu(X) = u \\ 0 & \nu(X) \neq u \end{cases}$$

control policy (stochastic) $\mu : \mathcal{U} \times \mathcal{X}^* \to [0,1]$

$$\mathrm{P}(\boldsymbol{u}_t = u \mid \boldsymbol{X}_t = X, \mu) = \mu(u, X)$$

Markov Decision Processes

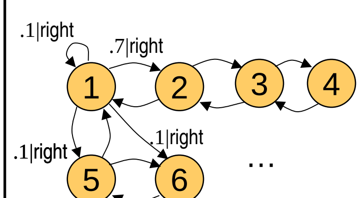

cost $J_\mu = \mathrm{E}[\boldsymbol{T}_{goal} \mid \mu]$ (to be minimize)

where $\boldsymbol{T}_{goal} := \min\{t : x_t = x_{goal}\} \leq \infty$

Markov Decision Processes

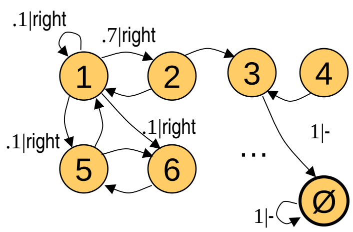

cost $J_\mu = \mathrm{E}[\boldsymbol{T}_{\text{goal}} \mid \mu]$ (to be minimize)

where $\boldsymbol{T}_{\text{goal}} := \min\{t : x_t = x_{\text{goal}}\} \le \infty$

one can also write

$$J_\mu = \mathrm{E}\left[\sum_{t=1}^{\infty} l(\boldsymbol{x}_t, \boldsymbol{u}_t) \,\middle|\, \mu\right] \quad \text{(additive)}$$

with $l(x,u) = \begin{cases} 1 & x \neq \varnothing \\ 0 & x = \varnothing \end{cases} \ge 0$

optimal control policy μ^*

$J_{\mu^*} = \min_\mu J_\mu$

Competitive Markov Decision Processes

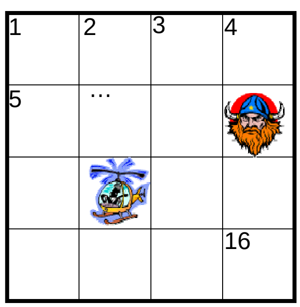

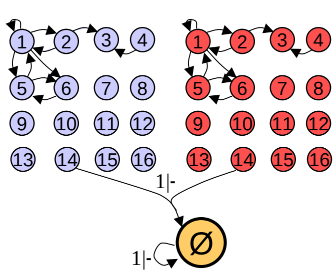

state $\boldsymbol{x}_t \in \mathcal{X} := \{(1,1),(1,2)\ldots(1,16),\ldots(16,16),\varnothing\}$

player U action $\boldsymbol{u}_t \in \mathcal{U} := \{$**up**, **down**, **left**, **right**$\}$

player D action $\boldsymbol{d}_t \in \mathcal{D} := \{$**up**, **down**, **left**, **right**$\}$

transition probability function

$p(x,x',u,d) = \mathrm{P}(\boldsymbol{x}_{t+1}=x' \mid \boldsymbol{x}_t=x, \boldsymbol{u}_t=u, \boldsymbol{d}_t=d)$

in pursuit-evasion games

- $\boldsymbol{x}_t = (\alpha_t, \beta_t) \in \mathcal{S} \times \mathcal{S}$ or $\boldsymbol{x}_t = \varnothing$
- $p((\alpha,\beta),(\alpha',\beta'),u,d) = \mathrm{P}(\alpha_{t+1}=\alpha' \mid \alpha_t=\alpha, \boldsymbol{u}_t=u) \cdot \mathrm{P}(\beta_{t+1}=\beta' \mid \beta_t=\beta, \boldsymbol{d}_t=d)$
- $p((\alpha,\beta),\varnothing,u,d) = \begin{cases} 1 & \alpha = \beta \\ 0 & \alpha \neq \beta \end{cases}$

Competitive, Hidden MDPs

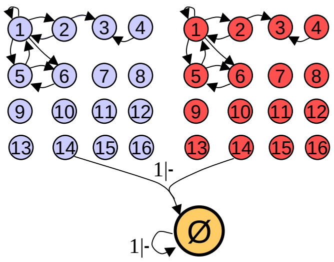

state $\boldsymbol{x}_t \in \mathcal{X} := \{(1,1),(1,2)\ldots(1,16),\ldots(16,16),\varnothing\}$

player U action $\boldsymbol{u}_t \in \mathcal{U} := \{$up, down, left, right$\}$

player D action $\boldsymbol{d}_t \in \mathcal{D} := \{$up, down, left, right$\}$

transition probability function

$$p(x,x',u,d) = \mathrm{P}(\boldsymbol{x}_{t+1}=x' \mid \boldsymbol{x}_t=x,\ \boldsymbol{u}_t=u\ ,\ \boldsymbol{d}_t=d)$$

player U observation

$$\boldsymbol{y}_t \in \mathcal{Y} := \{\text{found, not-found, } \varnothing\}$$

player D observation

$$\boldsymbol{z}_t \in \mathcal{Z} := \{\text{found, not-found, } \varnothing\}$$

observation probability functions

$$p_{\mathrm{U}}(y,x) = \mathrm{P}(\boldsymbol{y}_t=y \mid \boldsymbol{x}_t=x)$$
$$p_{\mathrm{D}}(z,x) = \mathrm{P}(\boldsymbol{z}_t=z \mid \boldsymbol{x}_t=x)$$

(the state $\varnothing$ is not "hidden" to any of the players)

Competitive, Hidden MDPs

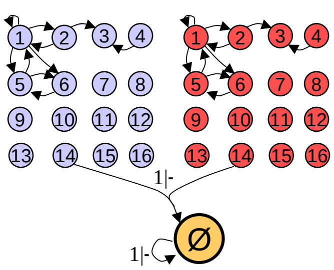

state $\boldsymbol{x}_t \in \mathcal{X} := \{(1,1),(1,2)\ldots(1,16),\ldots(16,16),\varnothing\}$

player U action $\boldsymbol{u}_t \in \mathcal{U} := \{$up, down, left, right$\}$

player D action $\boldsymbol{d}_t \in \mathcal{D} := \{$up, down, left, right$\}$

player U obs. $\boldsymbol{y}_t \in \mathcal{Y} := \{$found, not-found, $\varnothing\}$

player D obs. $\boldsymbol{z}_t \in \mathcal{Z} := \{$found, not-found, $\varnothing\}$

sequence of t obs. for player U

sequence of t-1 actions for player U

control policies (stochastic)

- for player U $\quad \mu : \mathcal{U} \times \mathcal{Y}^* \times \mathcal{U}^* \to [0,1]$

$$\mathrm{P}(\boldsymbol{u}_t = u \mid \boldsymbol{Y}_t = Y,\ \boldsymbol{U}_{t-1} = U,\ \mu) = \mu(u, Y, U)$$

- for player D $\quad \delta : \mathcal{D} \times \mathcal{Z}^* \times \mathcal{D}^* \to [0,1]$

$$\mathrm{P}(\boldsymbol{d}_t = d \mid \boldsymbol{Z}_t = Z,\ \boldsymbol{D}_{t-1} = D,\ \delta) = \delta(d, Z, D)$$

sequence of t obs. for player D

sequence of t-1 actions for player D

Competitive, Hidden MDPs

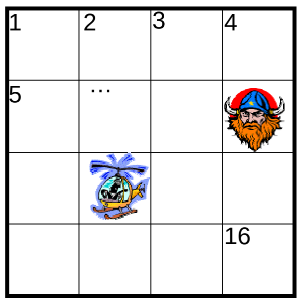

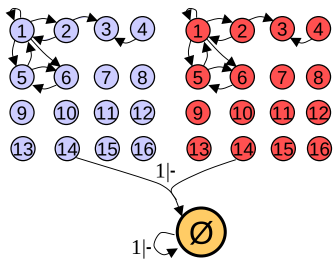

cost $J_{\mu,\delta} = \mathrm{E}\left[\sum_{t=1}^{\infty} l(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{d}_t) \,\middle|\, \mu, \delta\right]$

e.g., $l(x,u,d) = \begin{cases} 1 & x \neq \varnothing \\ 0 & x = \varnothing \end{cases} \geq 0$ (time to capture)

to be ***minimized*** by player U and
maximized by player D (zero sum)

Stackelberg equilibria:

- ***leader*** (D) chooses a worst case policy
  $\delta^* = \mathrm{argmax}_\delta \min_\mu J_{\mu,\delta}$ (max-min)
  and announces it
- ***follower*** (U) chooses its policy to better counteract the leader's policy, i.e.,
  $\mu^* = \mathrm{argmin}_\mu J_{\mu,\delta^*}$ (min)

(vice-versa if player U is the leader)

Compet

- "worst-case" policy for the leader in that he will never do worse than
  $J^* := \max_\delta \min_\mu J_{\mu,\delta}$
- the leader has nothing to fear from the follower discovering his strategy
- the follower's policy may be "fragile" in that he may do much worse if the leader deviates from the announced policy

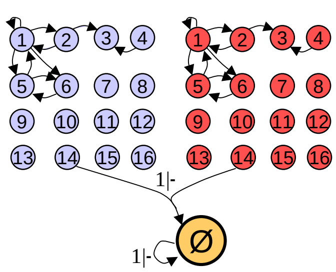

Stackelberg equilibria:

- ***leader*** (D) chooses a worst case policy
  $\delta^* = \mathrm{argmax}_\delta \min_\mu J_{\mu,\delta}$ (max-min)
  and announces it
- ***follower*** (U) chooses its policy to better counteract the leader's policy, i.e.,
  $\mu^* = \mathrm{argmin}_\mu J_{\mu,\delta^*}$ (min)

(vice-versa if player U is the leader)

Competitive, Hidden MDPs

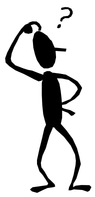

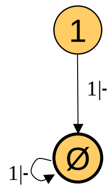

$\mathcal{X} := \{1, \emptyset\}$
$\mathcal{U} := \{\textsf{rock,paper,scissors}\}$
$\mathcal{D} := \{\textsf{rock,paper,scissors}\}$

$J_{\mu^*,\delta^*} = 1$
but for this choice of μ^* the follower will lose systematically if leader discovers his policy and deviates from δ^*

cost $J_{\mu,\delta} = \mathrm{E}\left[\sum_{t=1}^{\infty} l(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{d}_t) \middle| \mu, \delta\right]$

$$l(x,u,d) = \begin{cases} 2 & d > u \quad \text{(D wins)} \\ 1 & u = d \quad \text{(tie)} \\ 0 & u > d \quad \text{(U wins)} \\ 0 & x = \emptyset \end{cases}$$

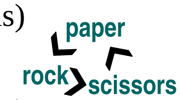

Stackelberg equilibria:

- ***leader*** (D) chooses a worst case policy
  $\delta^*(d,1) = 1/3 \quad \forall d$ (max-min)
  and announces it
- ***follower*** (U) chooses its policy to better counteract the leader's policy, i.e.,

$$\mu^*(u,1) = \begin{cases} 1 & u = \textsf{rock} \\ 0 & u \neq \textsf{rock} \end{cases} \quad \text{(min)}$$

Competitive, Hidden MDPs

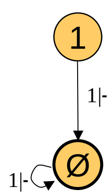

$\mathcal{X} := \{1, \emptyset\}$
$\mathcal{U} := \{\textsf{rock,paper,scissors}\}$
$\mathcal{D} := \{\textsf{rock,paper,scissors}\}$

$J_{\mu^*,\delta^*} = 1$
but for this choice of μ^* the leaders will not benefit from learning the follower's policy

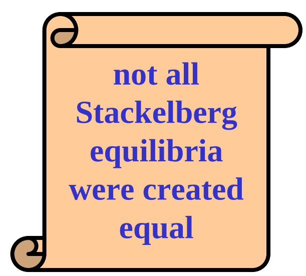

Stackelberg equilibria:

- ***leader*** (D) chooses a worst case policy
  $\delta^*(d,1) = 1/3 \quad \forall d$ (max-min)
  and announces it
- ***follower*** (U) chooses its policy to better counteract the leader's policy, i.e.,
  $\mu^*(u,1) = 1/3 \quad \forall u$ (min)

Competitive, Hidden MDPs

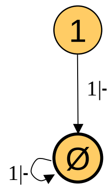

$\mathcal{X} := \{1, \varnothing\}$
$\mathcal{U} := \{$rock,paper,scissors$\}$
$\mathcal{D} := \{$rock,paper,scissors$\}$

$\mu^*(u,1) = 1/3 \;\; \forall u$
$\delta^*(d,1) = 1/3 \;\; \forall d$

cost $J_{\mu,\delta} = \mathrm{E}\left[\sum_{t=1}^{\infty} l(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{d}_t) \,\middle|\, \mu,\delta\right]$

$$l(x,u,d) = \begin{cases} 2 & d < u \quad \text{(D wins)} \\ 1 & u = d \quad \text{(tie)} \\ 0 & u > d \quad \text{(U wins)} \\ 0 & x = \varnothing \end{cases}$$

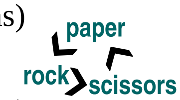

Nash equilibria (or saddle-point):
pair of policies (μ^*, δ^*) such that

$$J_{\mu^*,\delta} \le J_{\mu^*,\delta^*} \le J_{\mu,\delta^*} \quad \forall \mu,\delta$$

Competitive, Hidden MDPs

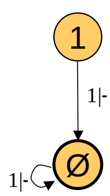

$\mathcal{X} := \{1, \varnothing\}$
$\mathcal{U} := \{$rock,paper,scissors$\}$
$\mathcal{D} := \{$rock,paper,scissors$\}$

$\mu^*(u,1) = 1/3 \;\; \forall u$
$\delta^*(d,1) = 1/3 \;\; \forall d$

all Nash equilibria were created equal

- “robust” in that any unilateral deviation from the equilibrium will make that player do worse
- a player playing at a Nash equilibria has nothing to fear from the other discovering his strategy

but Nash equilibria don’t always exist

Nash equilibria (or saddle-point):
pair of policies (μ^*, δ^*) such that

$$J_{\mu^*,\delta} \le J_{\mu^*,\delta^*} \le J_{\mu,\delta^*} \quad \forall \mu,\delta$$

value of the game

If (μ^*, δ^*) and (μ^+, δ^+) are Nash eq. then (μ^*, δ^+) and (μ^+, δ^*) are also Nash eq. and

$$J_{\mu^*,\delta^*} = J_{\mu^+,\delta^+} = J_{\mu^*,\delta^+} = J_{\mu^+,\delta^*} = \\ = \min_\mu \max_\delta J_{\mu,\delta} = \max_\delta \min_\mu J_{\mu,\delta}$$

Competitive, Hidden MDPs

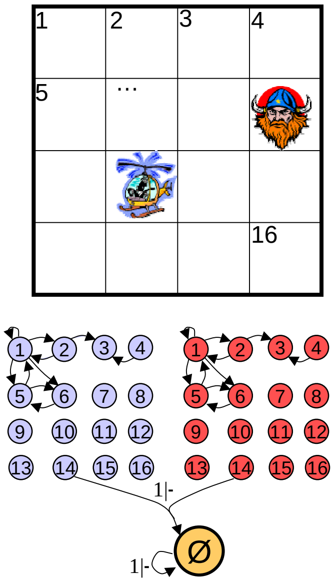

cost $J_{\mu,\delta} = \mathrm{E}\left[\sum_{t=1}^{\infty} l(\boldsymbol{x}_t,\boldsymbol{u}_t,\boldsymbol{d}_t) \,\middle|\, \mu,\delta\right]$

e.g., $l(x,u,d) = \begin{cases} 1 & x \neq \varnothing \\ 0 & x = \varnothing \end{cases} \geq 0$ (time to capture)

to be ***minimized*** by player U and ***maximized*** by player D (zero sum)

Stackelberg equilibria:

- ***leader*** (D) chooses a worst case policy
  $\delta^* = \mathrm{argmax}_\delta \min_\mu J_{\mu,\delta}$ (max-min)
  and announces it
- ***follower*** (U) chooses its policy to better counteract the leader's policy, i.e.,
  $\mu^* = \mathrm{argmin}_\mu J_{\mu,\delta^*}$ (min)

Competitive, Hidden MDPs

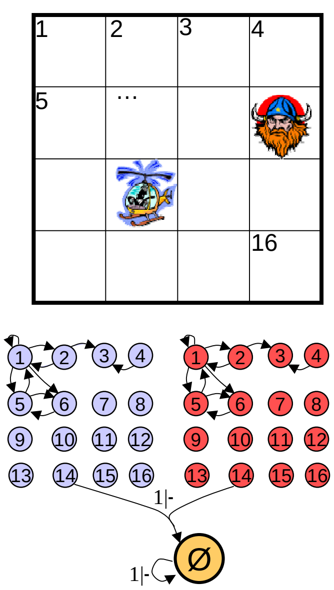

follower's optimal cost-to-go

$$V(Y,U) := \min_\mu \mathrm{E}\left[\sum_{t=\tau}^{\infty} l(\boldsymbol{x}_t,\boldsymbol{u}_t,\boldsymbol{d}_t) \,\middle|\, \boldsymbol{Y}_\tau = Y, \boldsymbol{U}_{\tau-1} = U, \mu, \delta^*\right]$$

for sequence $\boldsymbol{Y} \in \mathcal{Y}^*$ of τ observations and sequence $\boldsymbol{U} \in \mathcal{U}^*$ of τ-1 actions

Bellman's equation

$$\underbrace{\begin{aligned} V(Y,U) = \min_{\{p_u \in [0,1]: u \in \mathcal{U}\}} \mathrm{E}\big[& l(\boldsymbol{x}_\tau,\boldsymbol{u}_\tau,\boldsymbol{d}_\tau) + V(\boldsymbol{Y}_{\tau+1},\boldsymbol{U}_\tau) \,| \\ & |\, \boldsymbol{Y}_\tau = Y, \boldsymbol{U}_{\tau-1} = U, \mu(u,Y,U) = p_u, \delta^* \big] \\ V(\boldsymbol{Y}_t, \boldsymbol{U}_{t-1}) = 0 \ & \text{ a.s., conditioned to } \boldsymbol{x}_t = \varnothing \end{aligned}}_{\text{uniquely define cost-to-go } V}$$

optimal follower's control policy μ^*

$$\min_{\{p_u: u \in \mathcal{U}\}} \mathrm{E}\left[\cdots | \cdots, \mu(u,Y,U) = p_u, \delta^*\right] =$$
$$= \mathrm{E}\left[\cdots | \cdots, \mu(u,Y,U) = \mu^*(u,Y,U), \delta^*\right]$$

Competitive, Hidden MDPs

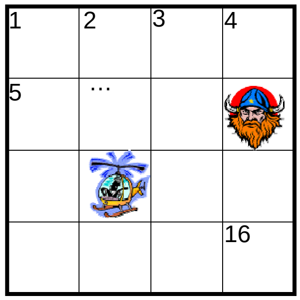

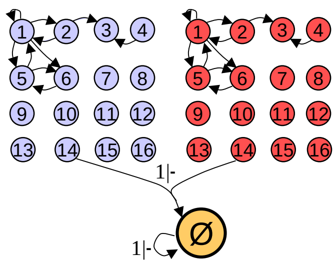

follower's optimal policy μ^*

$$\mu^* = \mathrm{argmin}_\mu \; J_{\mu,\delta^*}$$

- deterministic (at least one of the optimal)

$$\mu^*(u, Y, U) \in \{0,1\}$$

- non Markov, but

$$\mu^*(u, Y, U) = f\left(u, I(Y,U)\right)$$

where

$$\underbrace{I(Y,U) := \left\{ \mathrm{P}\left(\begin{matrix} x_t = x, \boldsymbol{Z}_t = Z, \boldsymbol{D}_{t-1} = D \\ | \boldsymbol{Y}_t = Y, \boldsymbol{U}_{t-1} = U \end{matrix} \right) : \begin{matrix} x \in \mathcal{X}, \\ Z \in \mathcal{Z}^*, \\ D \in \mathcal{D}^* \end{matrix} \right\}}_{\text{information state}}$$

information state
(size grows with t)

Competitive, Hidden MDPs

follower's optimal policy μ^*

$$\mu^* = \mathrm{argmin}_\mu \; J_{\mu,\delta^*}$$

- follower's optimal policy is finite dimensional only if $\boldsymbol{\delta^*}$ uses ***finite memory***
- even when $\boldsymbol{\delta^*}$ is memoryless, the size of the information state tends to be prohibitively large, e.g.,

$$2^N \cdot N^{(m+1)} \sim 2.6\times 10^{69}$$

number of cells (200 with uniformly distributed obstacles)

number of pursuers (3)

ministic (at least one of the optimal)

$$\mu^*(u, Y, U) \in \{0,1\}$$

arkov, but

$$\mu^*(u, Y, U) = f\left(u, I(Y,U)\right)$$

$$\underbrace{I(Y,U) := \left\{ \mathrm{P}\left(\begin{matrix} x_t = x, \boldsymbol{Z}_t = Z, \boldsymbol{D}_{t-1} = D \\ | \boldsymbol{Y}_t = Y, \boldsymbol{U}_{t-1} = U \end{matrix} \right) : \begin{matrix} x \in \mathcal{X}, \\ Z \in \mathcal{Z}^*, \\ D \in \mathcal{D}^* \end{matrix} \right\}}_{\text{information state}}$$

information state
(size grows with t)

Competitive, Hidden MDPs

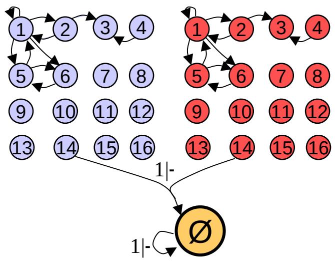

cost $J_{\mu,\delta} = \mathrm{E}\left[\sum_{t=1}^{\infty} l(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{d}_t) \,\middle|\, \mu, \delta \right]$

e.g., $l(x,u,d) = \begin{cases} 1 & x \neq \varnothing \\ 0 & x = \varnothing \end{cases} \geq 0$ (time to capture)

to be ***minimized*** by player U and ***maximized*** by player D (zero sum)

Stackelberg equilibria:

- ***leader*** (D) $\delta^* = \operatorname{argmax}_\delta \min_\mu J_{\mu,\delta}$
- ***follower*** (U) $\mu^* = \operatorname{argmin}_\mu J_{\mu,\delta^*}$

Nash equilibria:

pair of policies (μ^*, δ^*) such that

$$\underbrace{J_{\mu^*,\delta} \leq J_{\mu^*,\delta^*} \leq J_{\mu,\delta^*} \quad \forall \mu,\delta}_{\text{open research problems}}$$

Conclusions

- The "mission-level" control of Unmanned Air Vehicles requires a ***probabilistic framework***.
- The problem of coordinating teams of autonomous agents is naturally formulated in a ***game theoretical*** setting.
- ***Exact solutions*** for these types of problems are often computationally intractable and, in some cases, open research problems.

Some References

- Kumar & Varaiya, Stochastic Systems. Prentice Hall, 1986.
- Fudenberg & Tirole, Game Theory. MIT Press, 1993.
- Basar & Olsder, Dynamic Noncooperative Game Theory. 2nd ed. SIAM, 1999.
- Bertsekas, Dynamic Programming and Optimal Control. Athena Scientific, vols. 1 & 2, 1995.
- Filar & Vrieze, Competitive Markov Decision Processes. Springer-Verlag, 1997.
- Patek & Bertsekas, Stochastic Shortest Path Games. SIAM J. Control Optim., 37(3):804-824, 1999.

- see also papers by Michael L. Littman, Junling Hu, Michael P. Wellman on reinforcement learning in Markov games.

- Intelligent Control Architectures for Unmanned Air Vehicles Home Page:
  http://robotics.eecs.berkeley.edu/~sastry/ONRhomepage.html
- Software Enabled Control Home Page:
  http://sec.eecs.berkeley.edu/