Roborta vs. the Fair Light!

Pedro R. D'Argenio

joint work with
Pablo Castro, Ramiro Demasi, and Luciano Putruele

CONICET

UNC

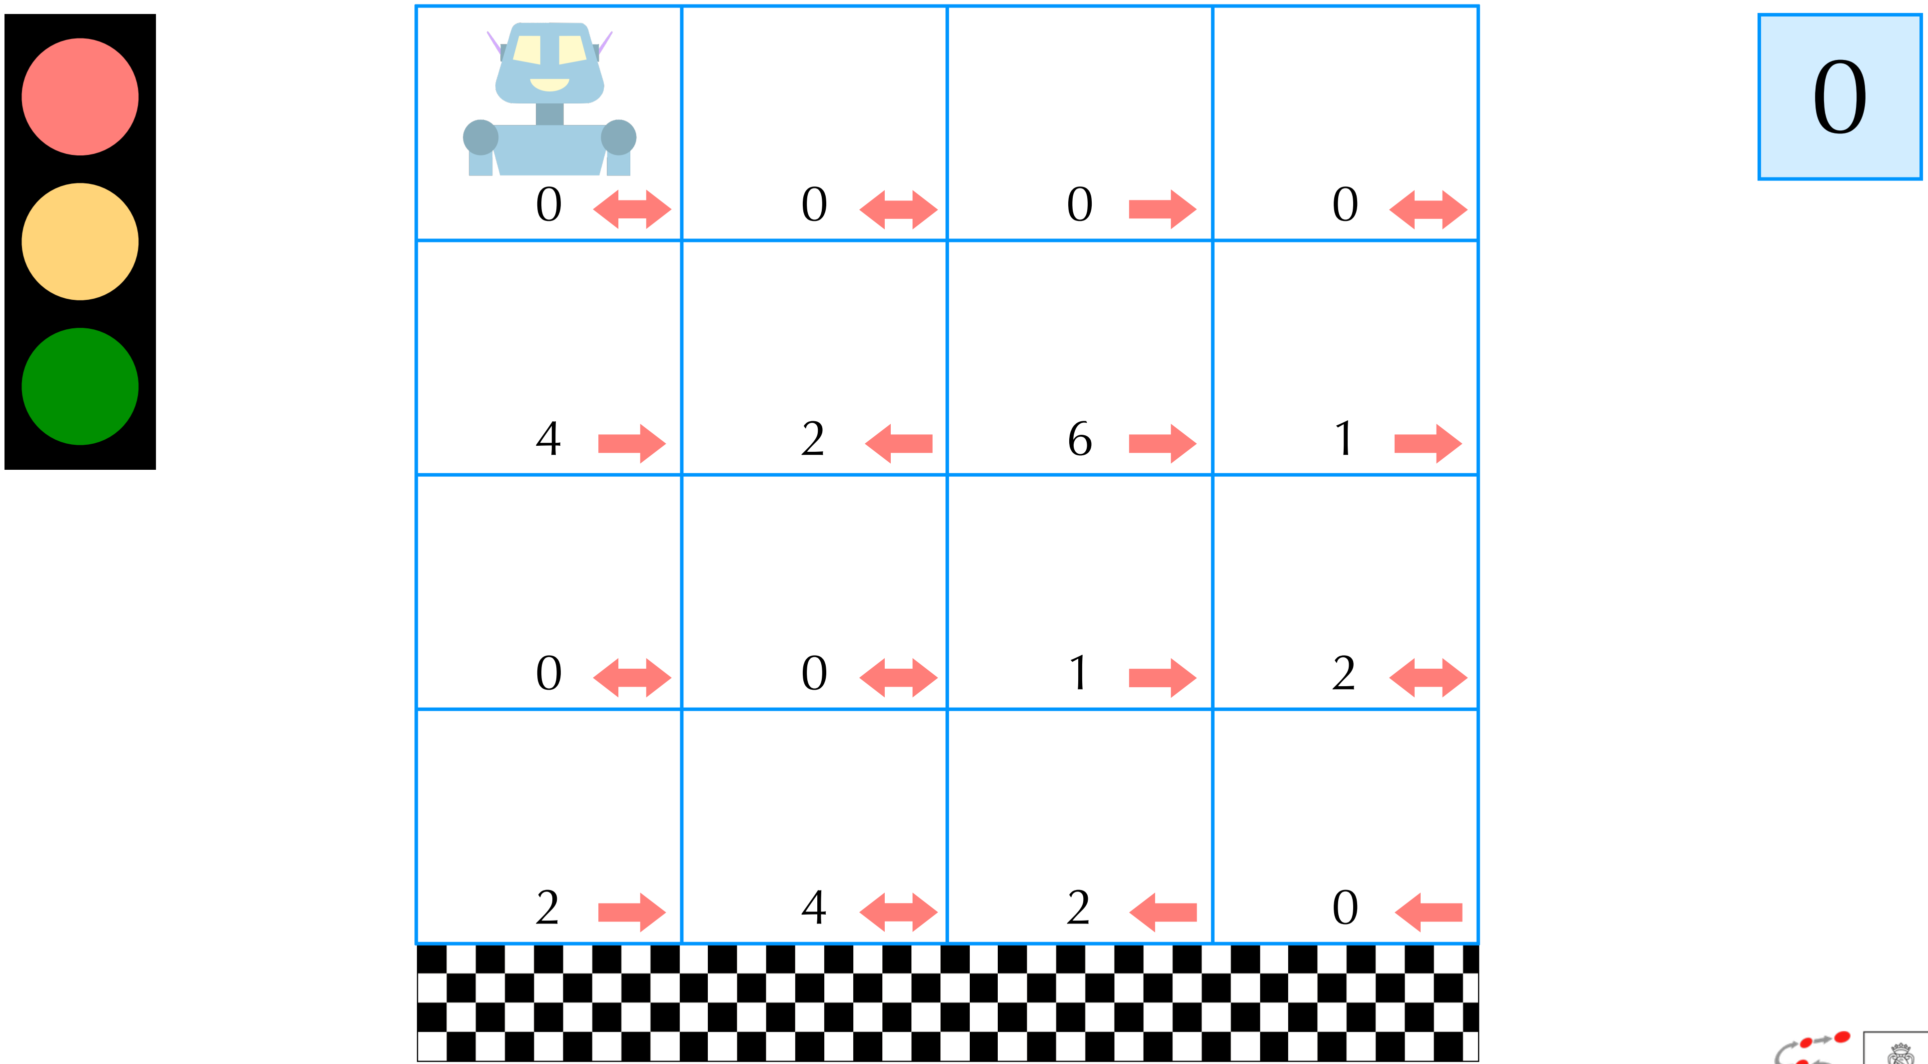

CONICET

UNC

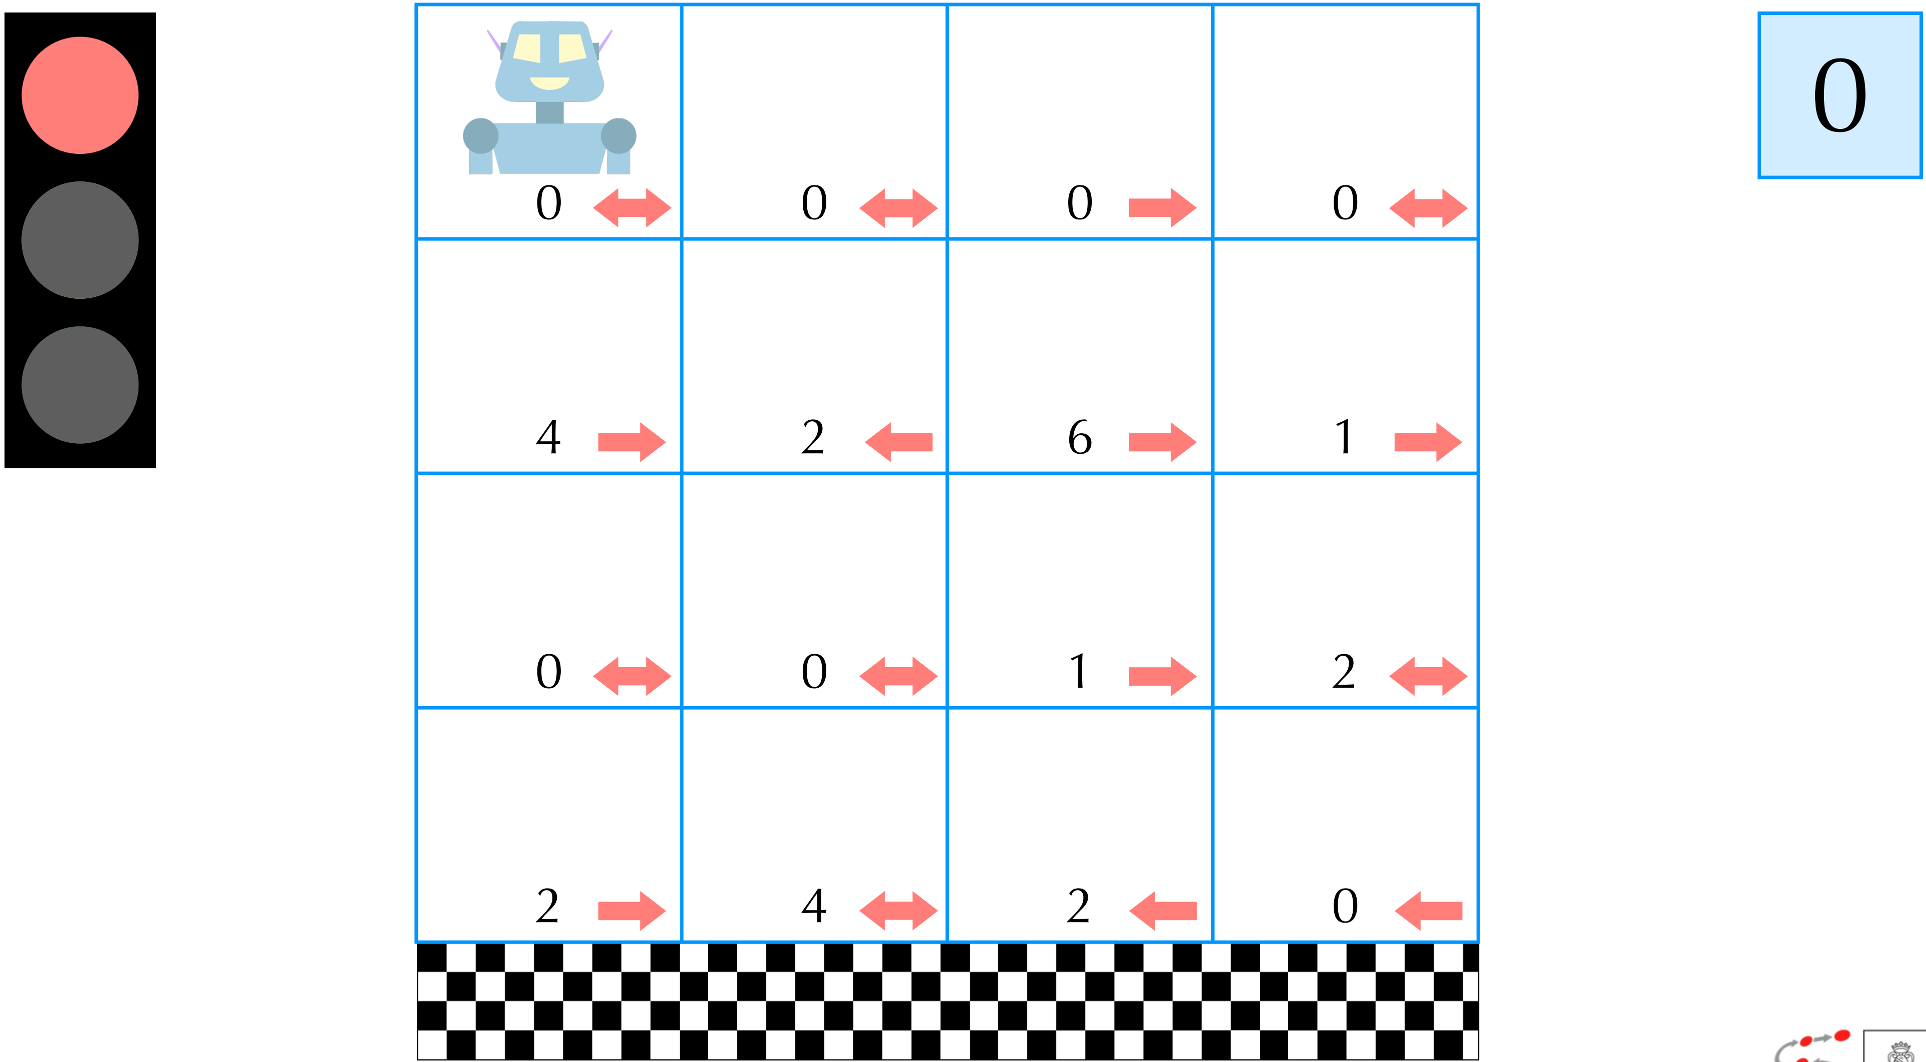

CONICET

UNC

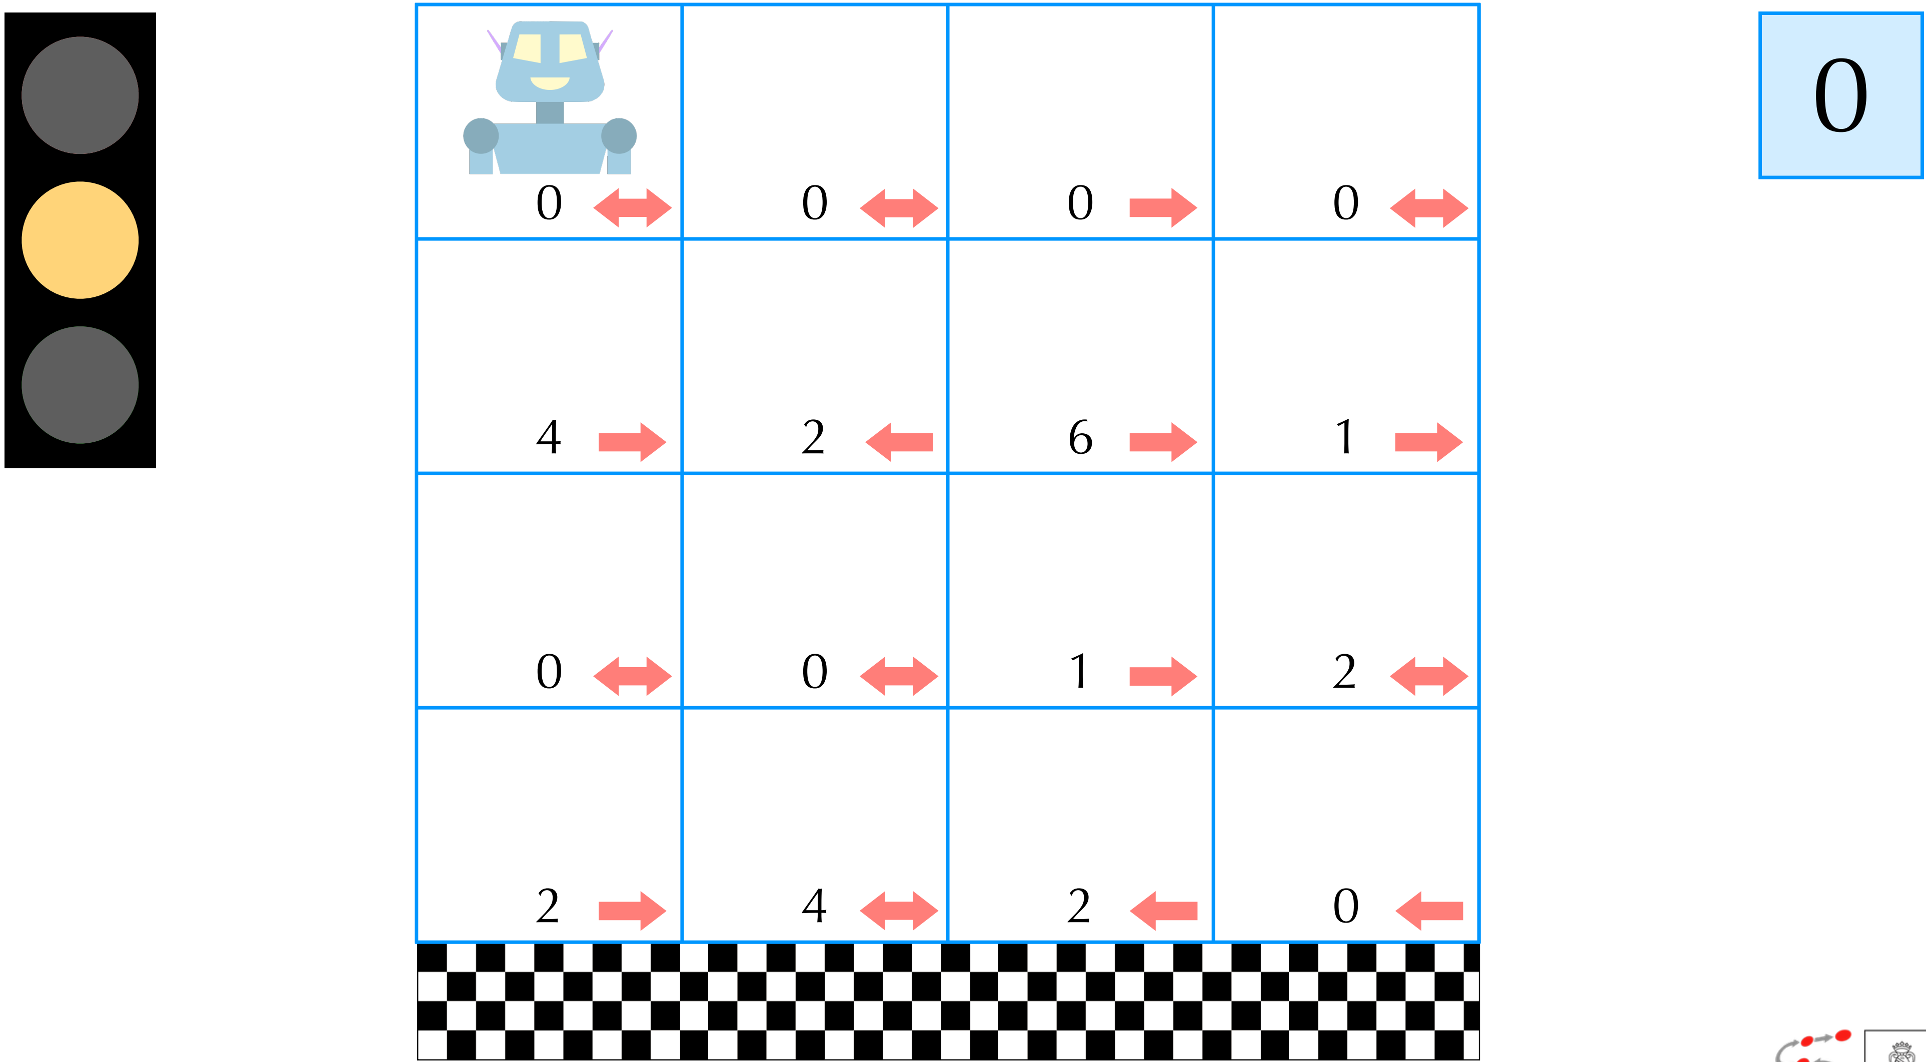

CONICET

UNC

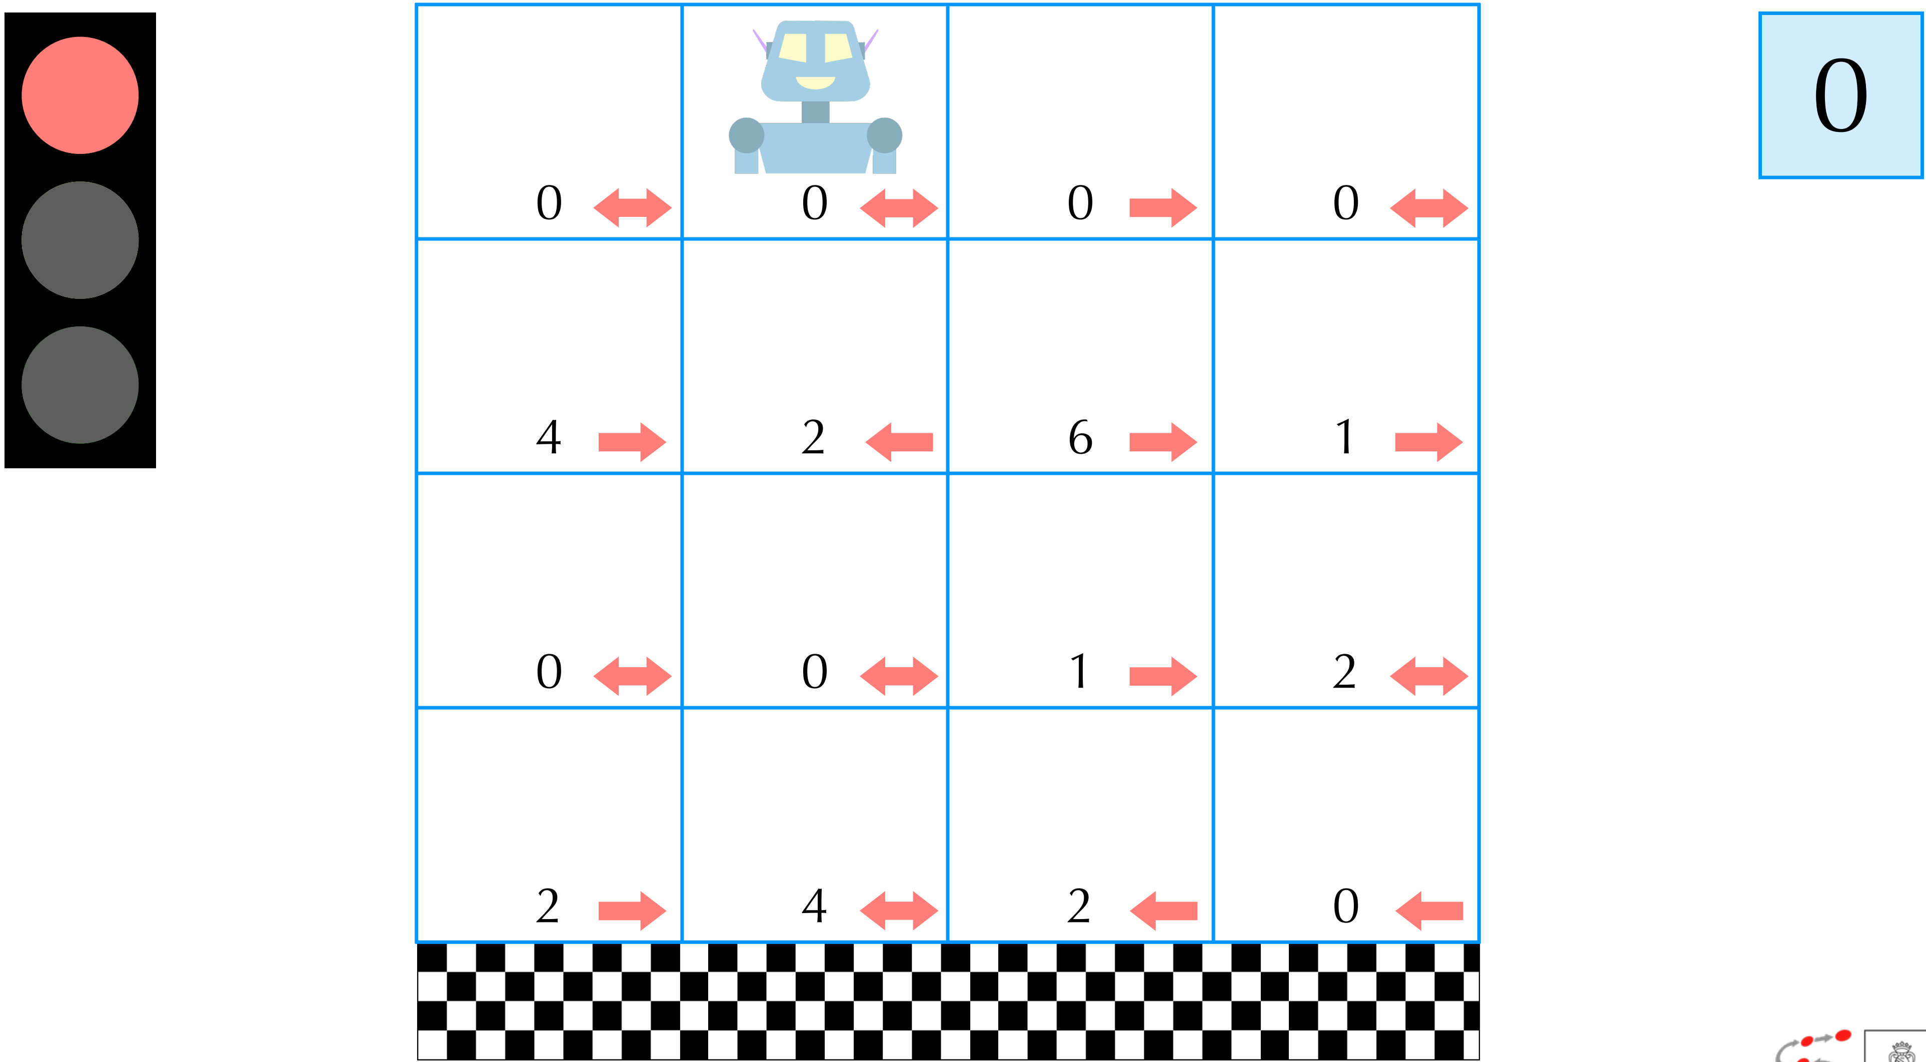

CONICET

UNC

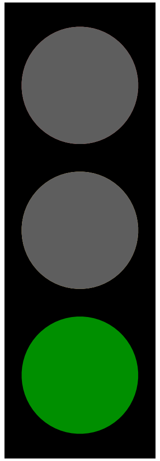

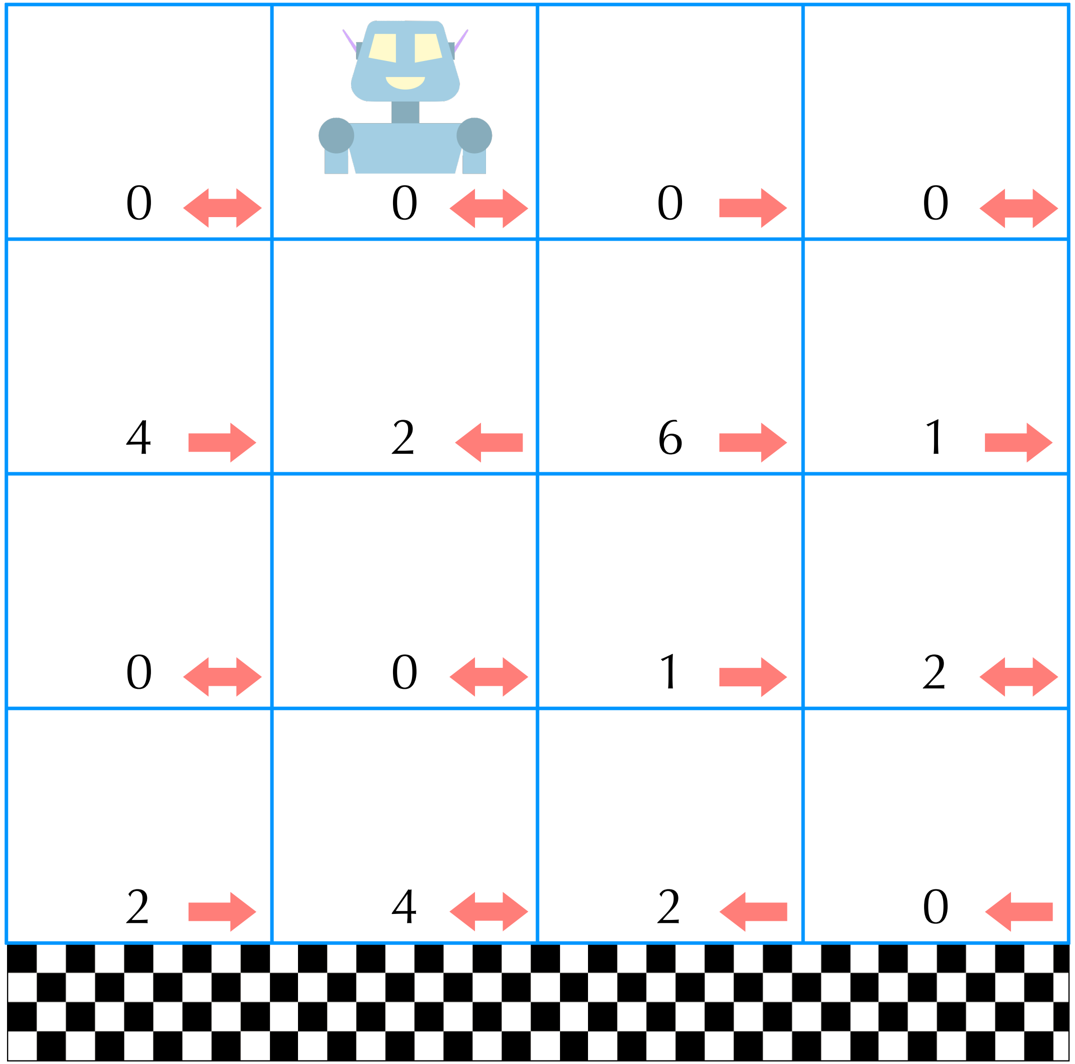

0

CONICET

UNC

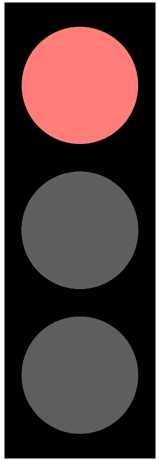

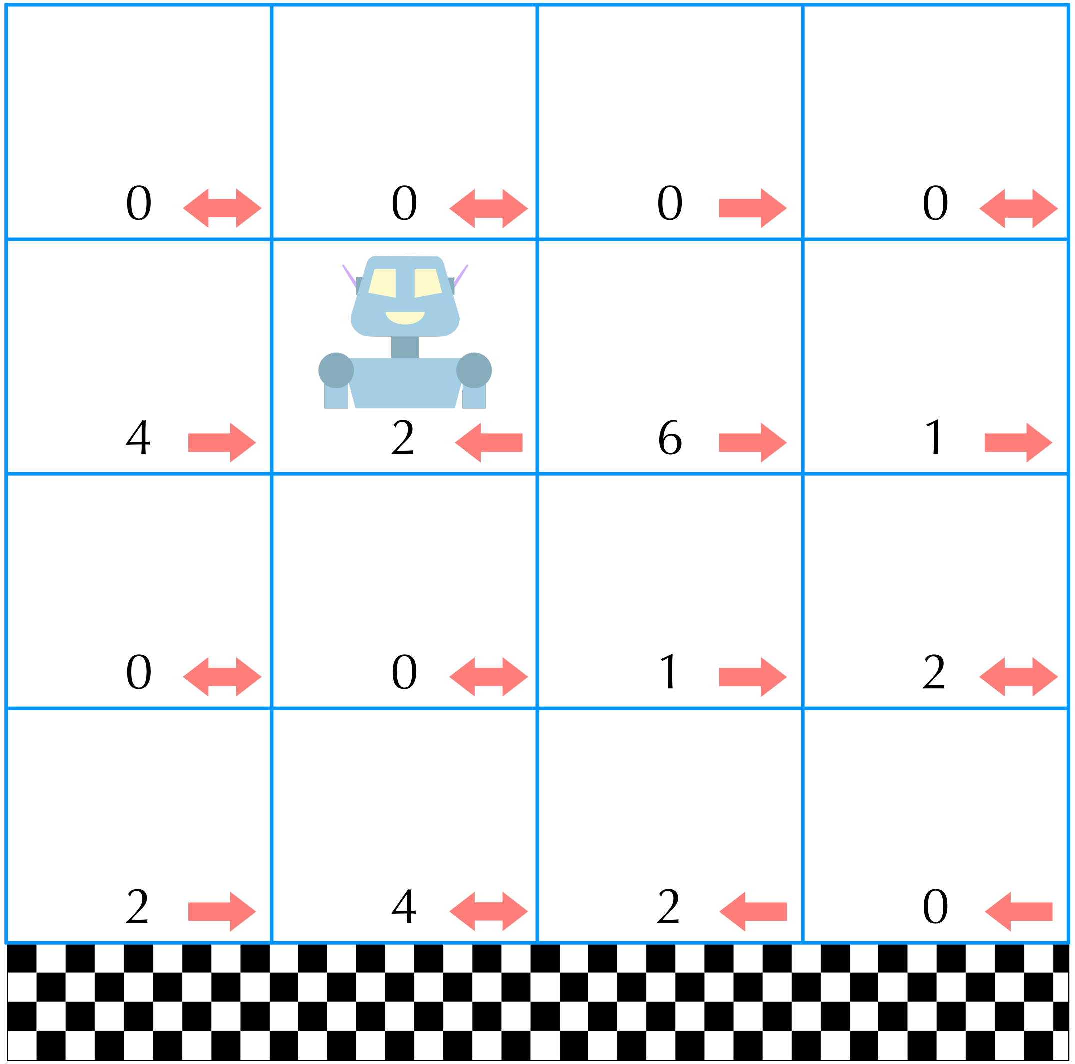

2

CONICET

UNC

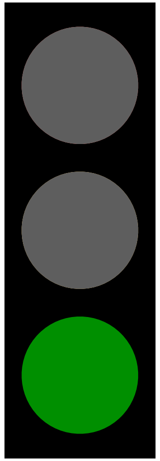

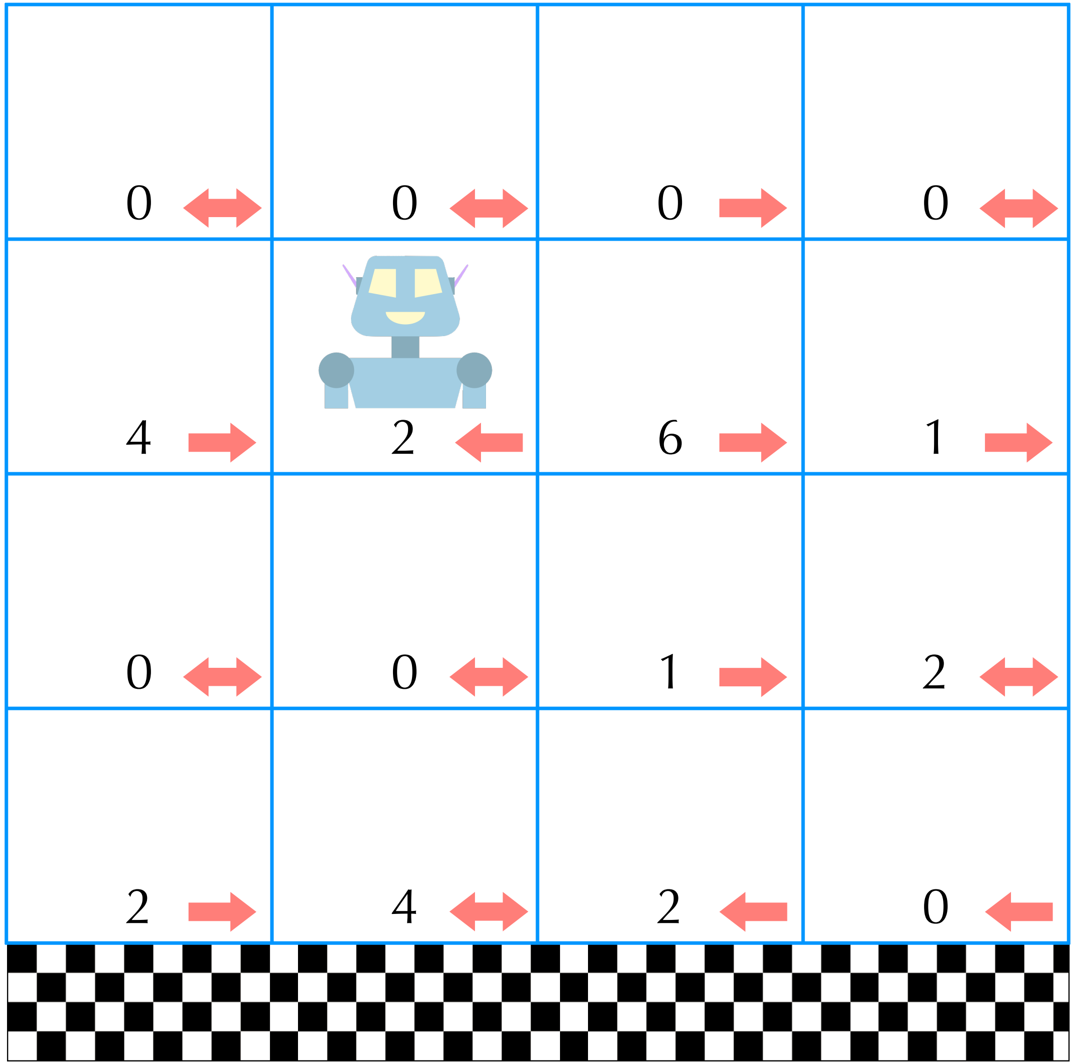

2

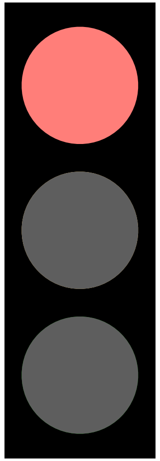

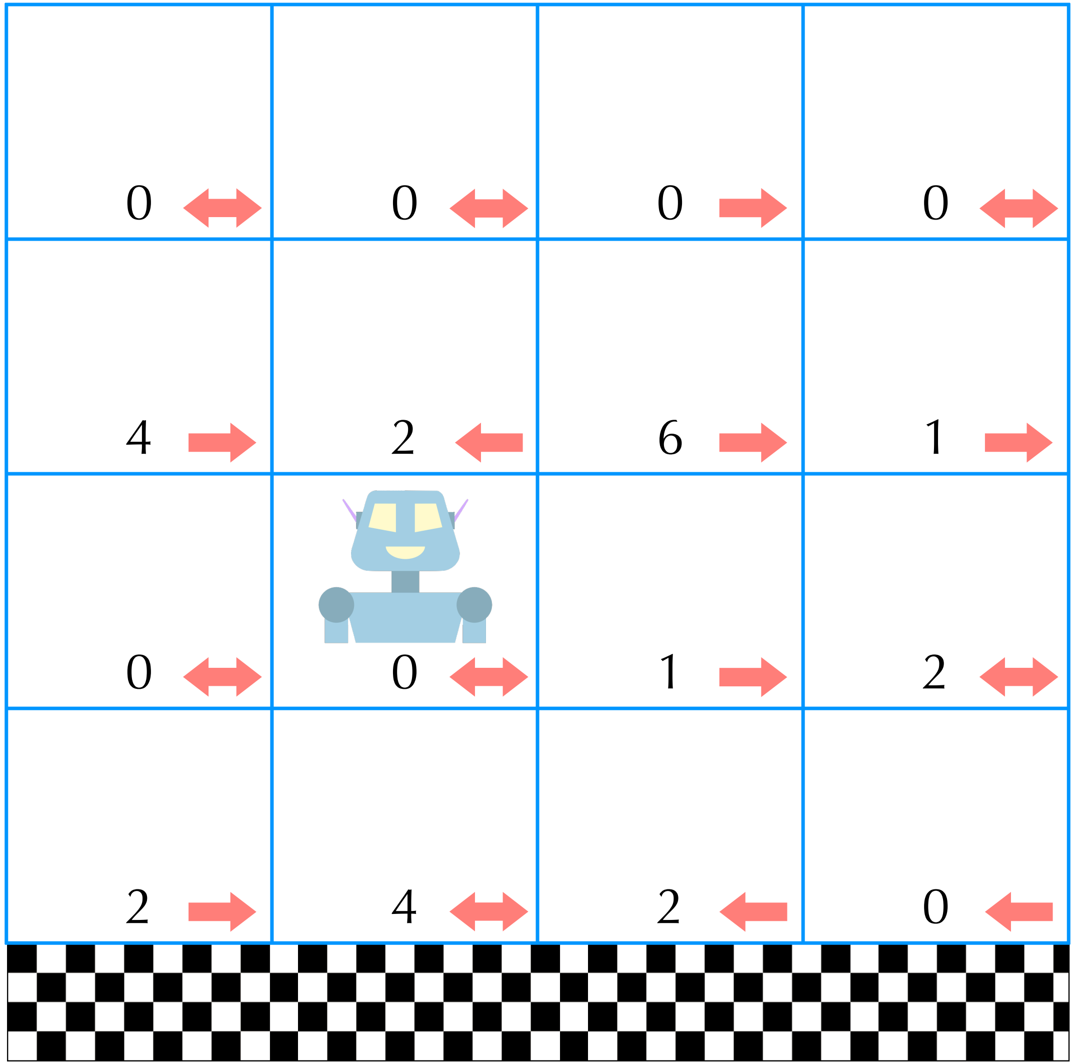

2

CONICET

UNC

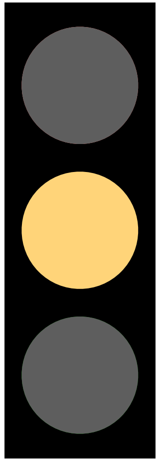

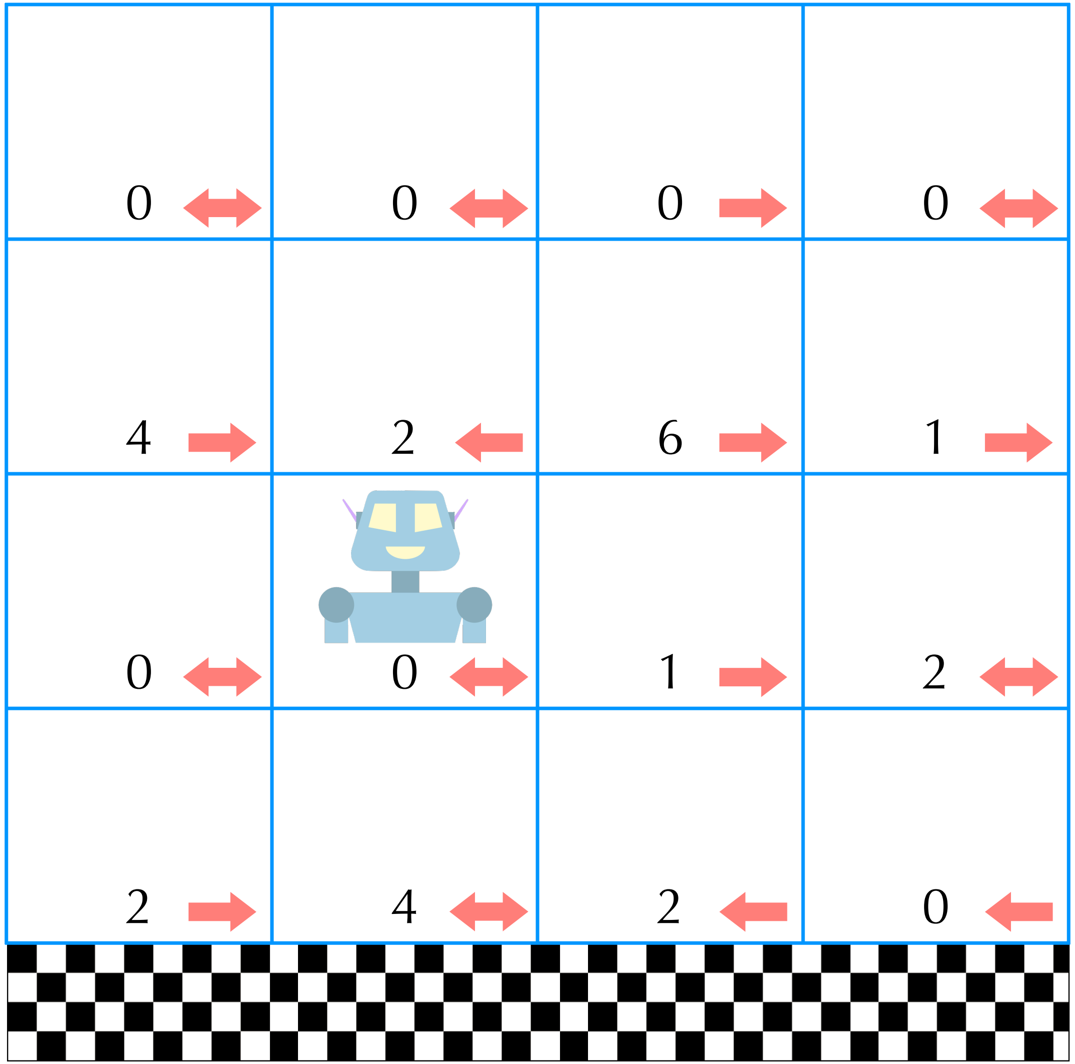

2

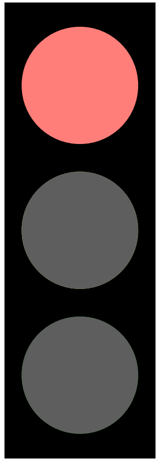

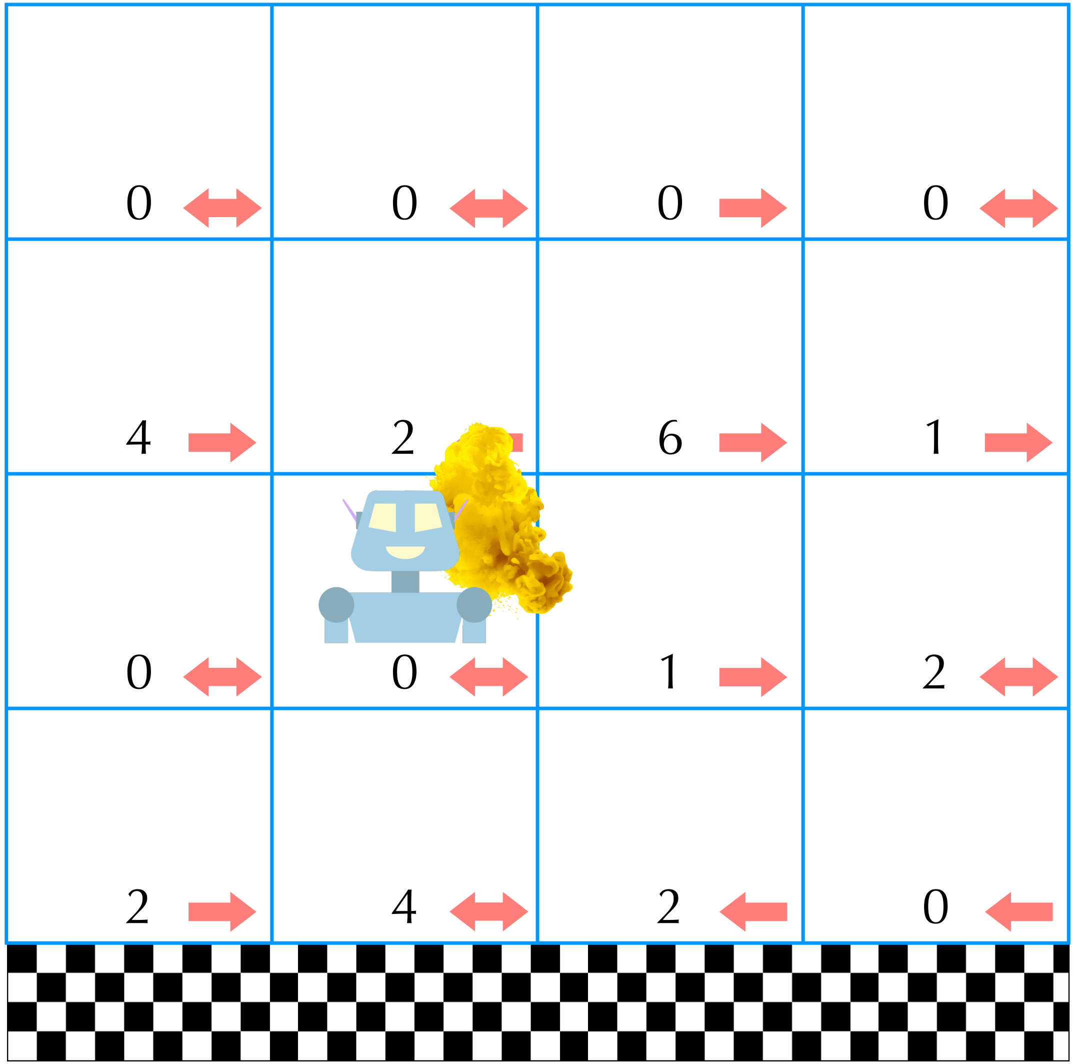

2

CONICET

UNC

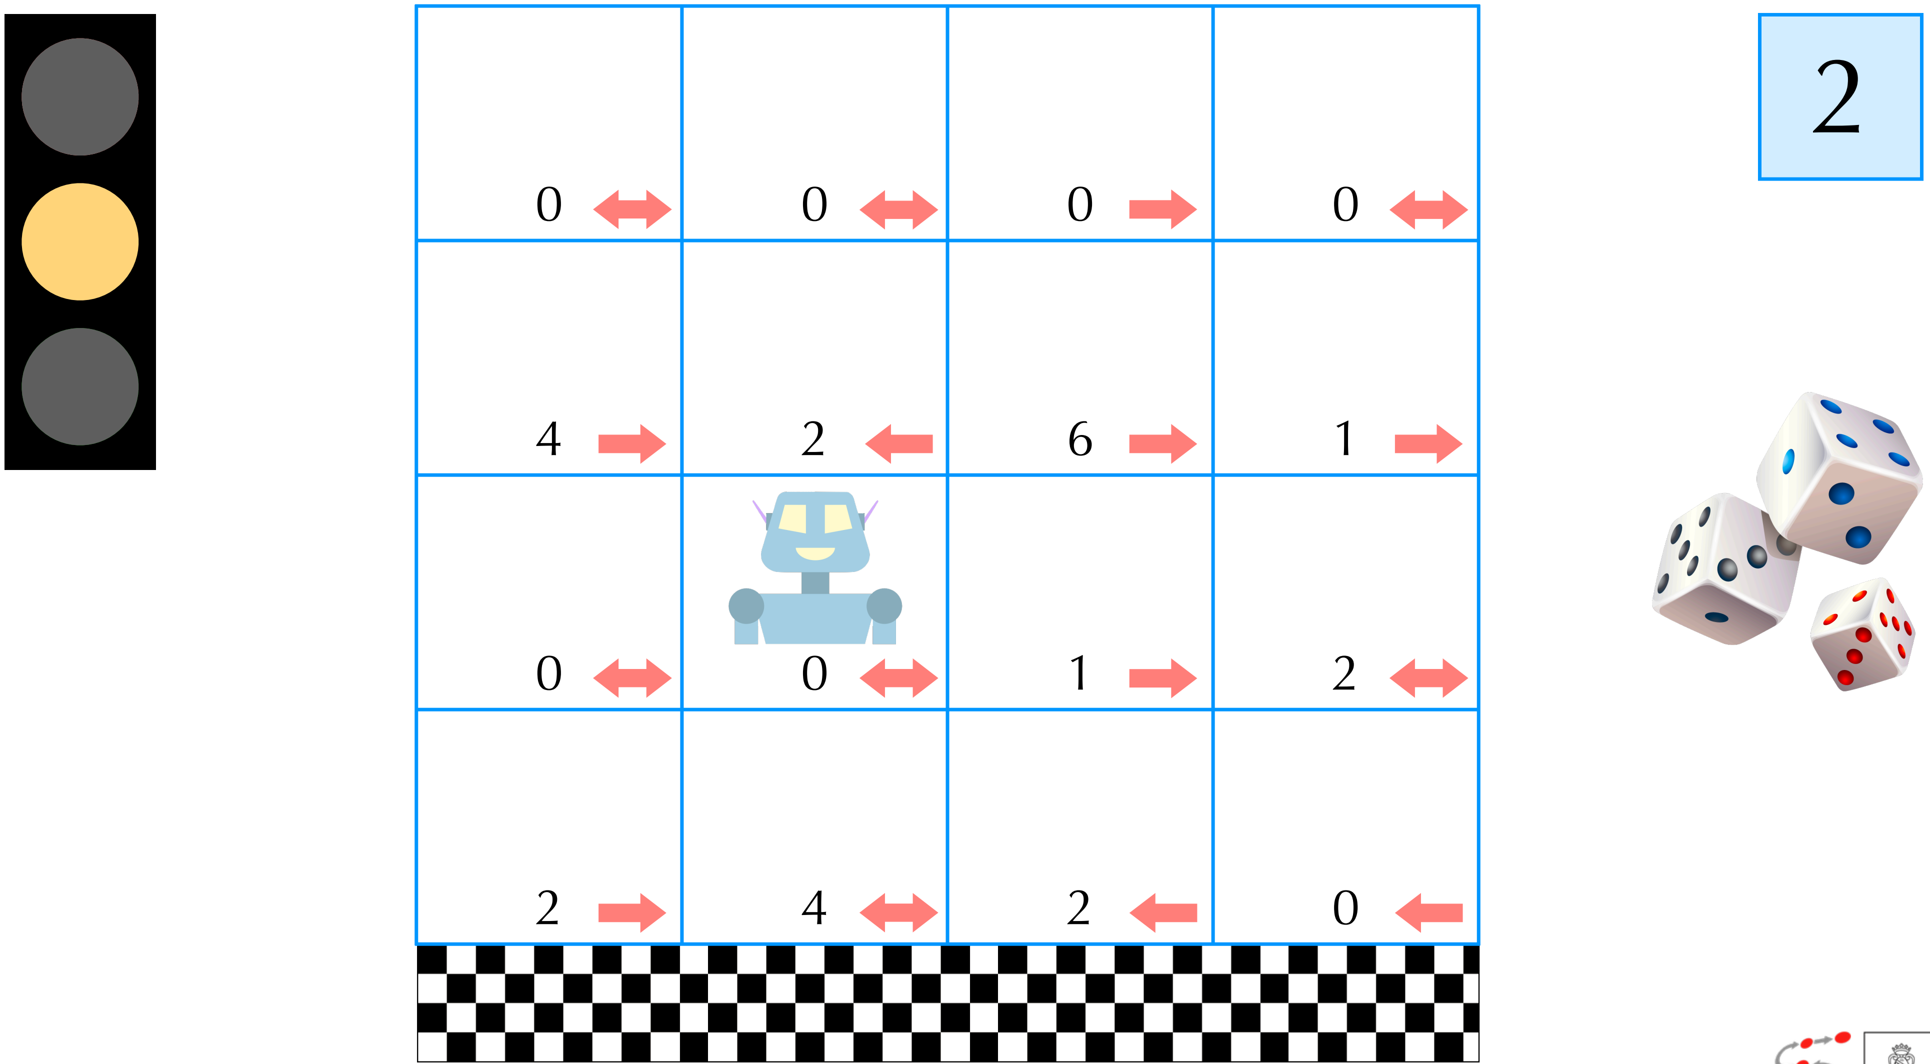

CONICET

UNC

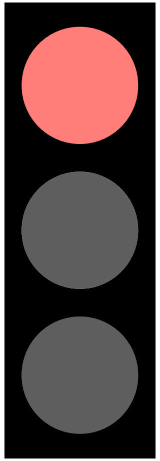

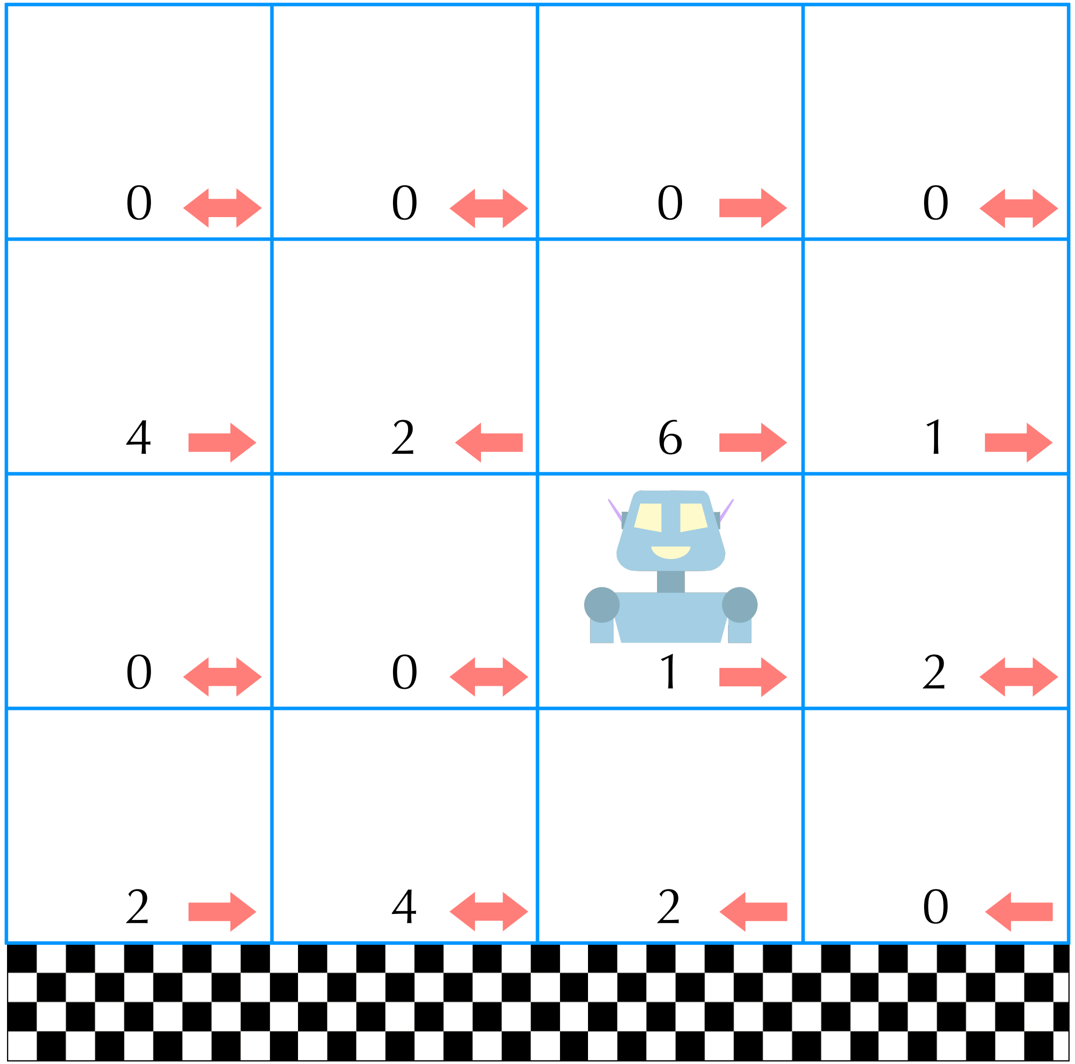

3

CONICET

UNC

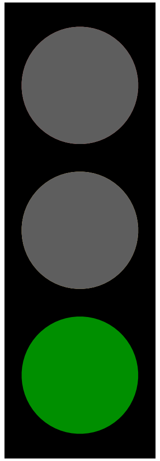

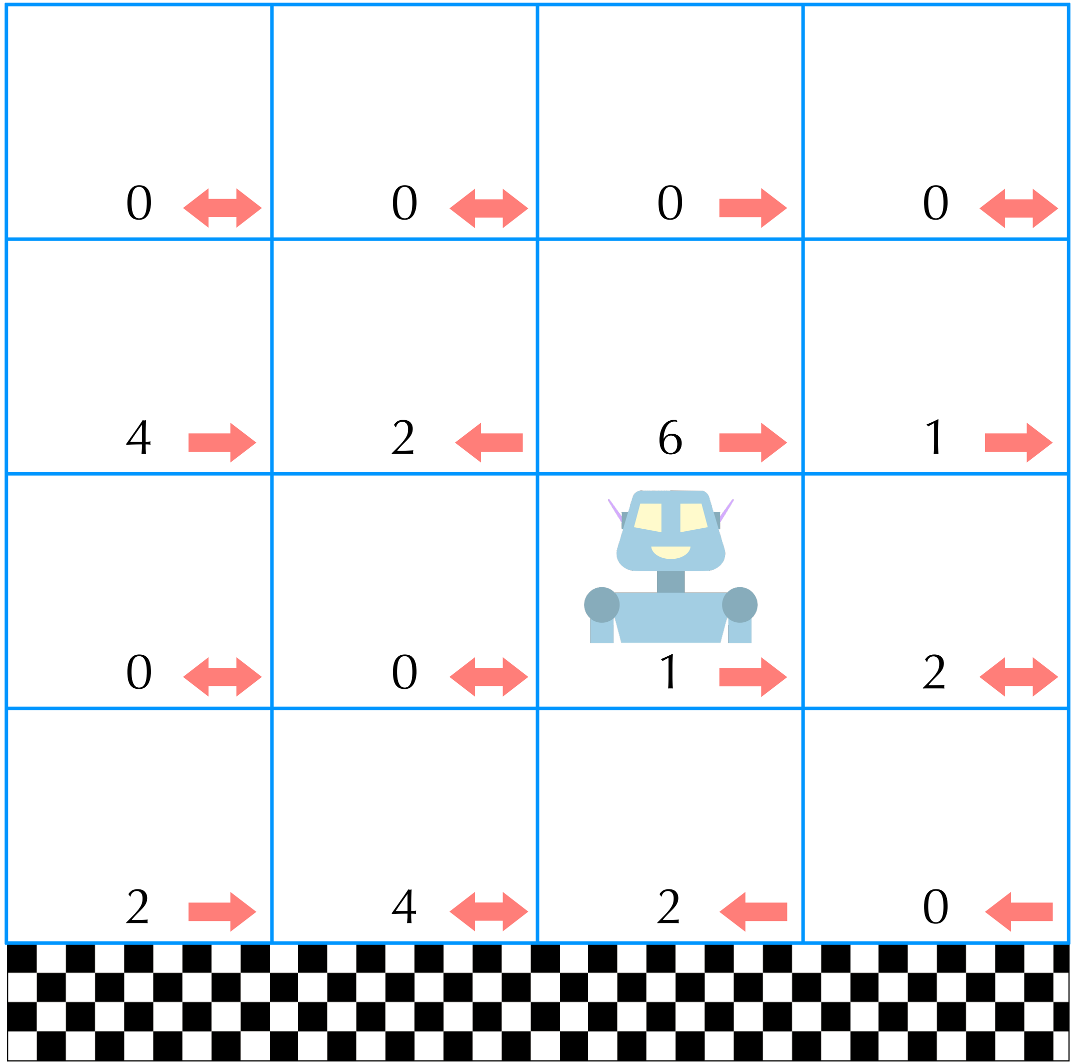

3

CONICET

UNC

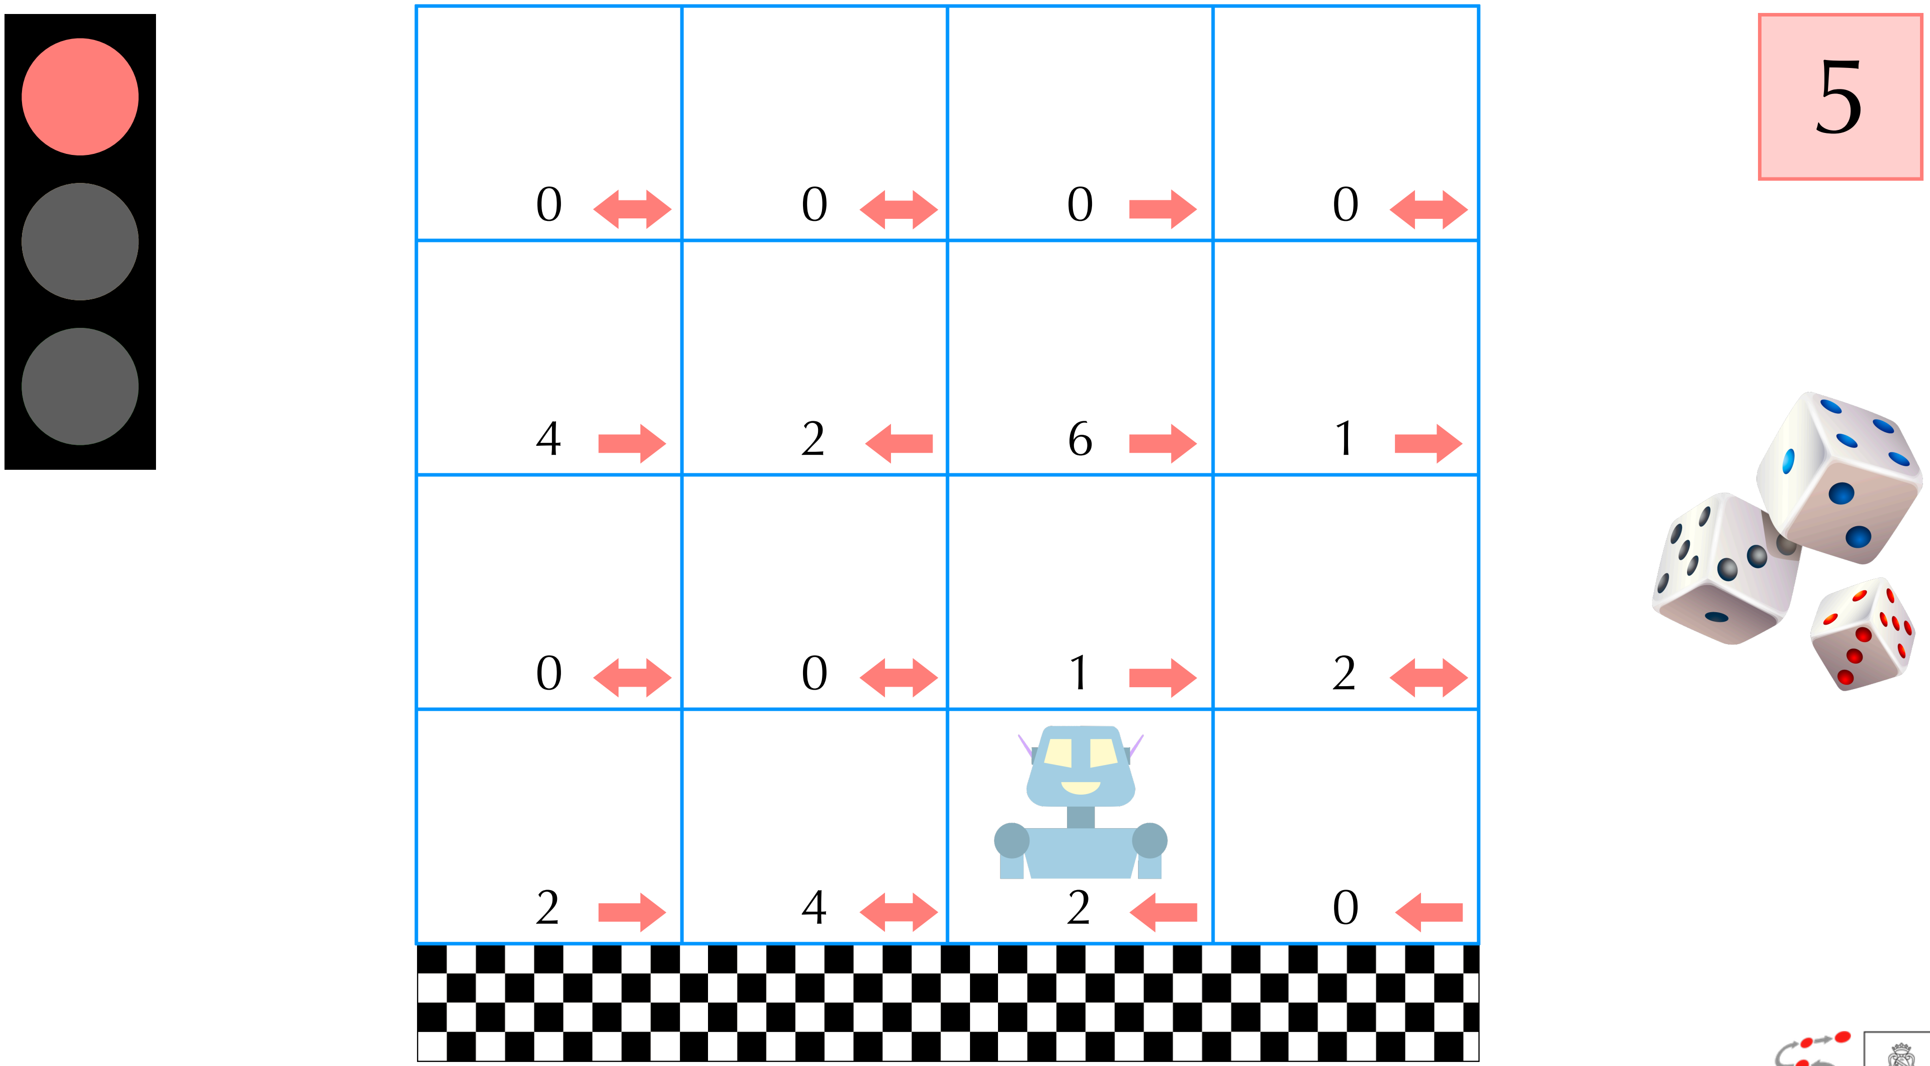
5
0 0 0 0
4 2 6 1
0 0 1 2
2 4 2 0

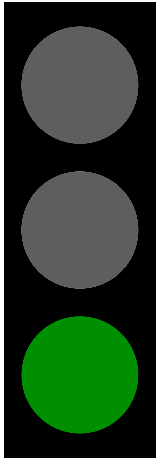

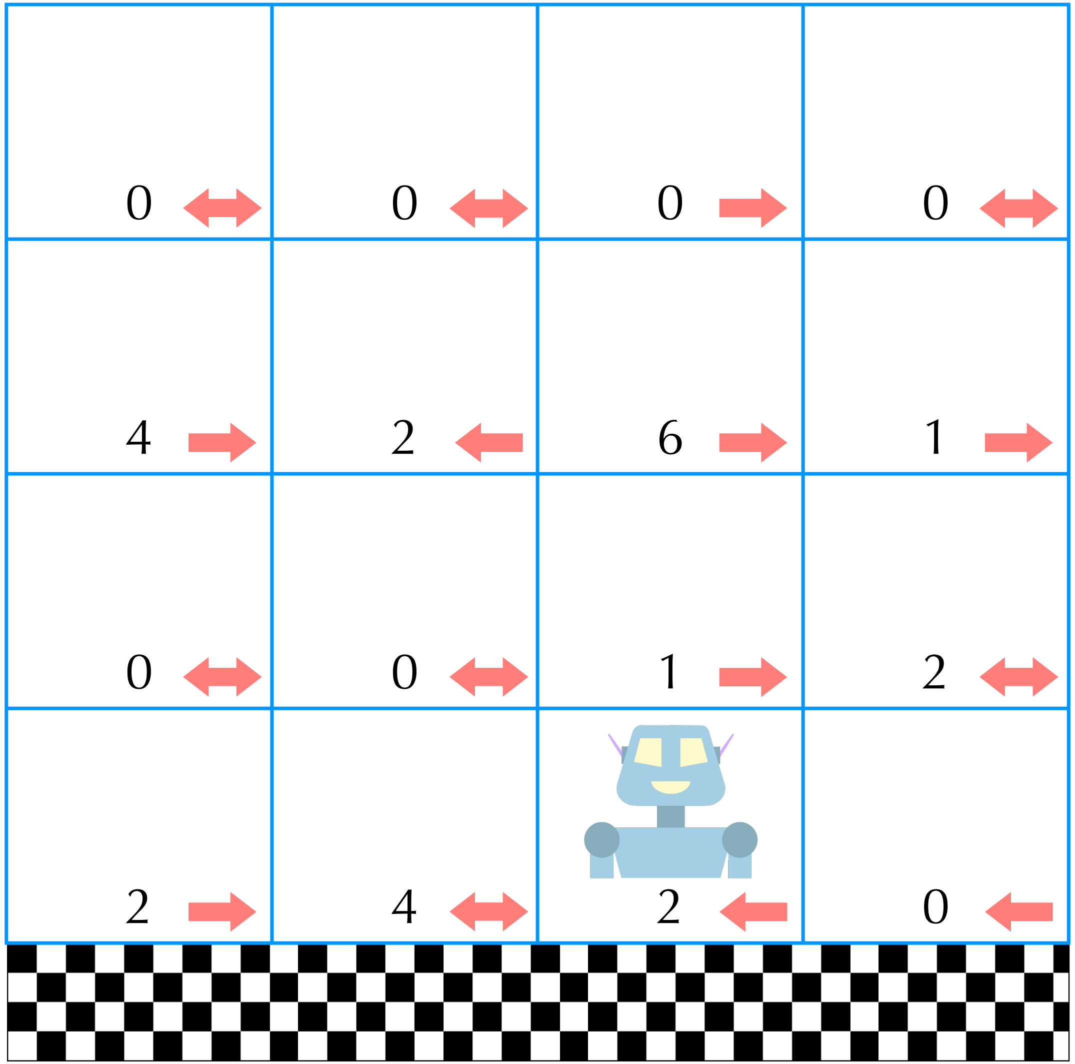

5

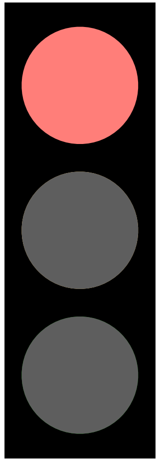

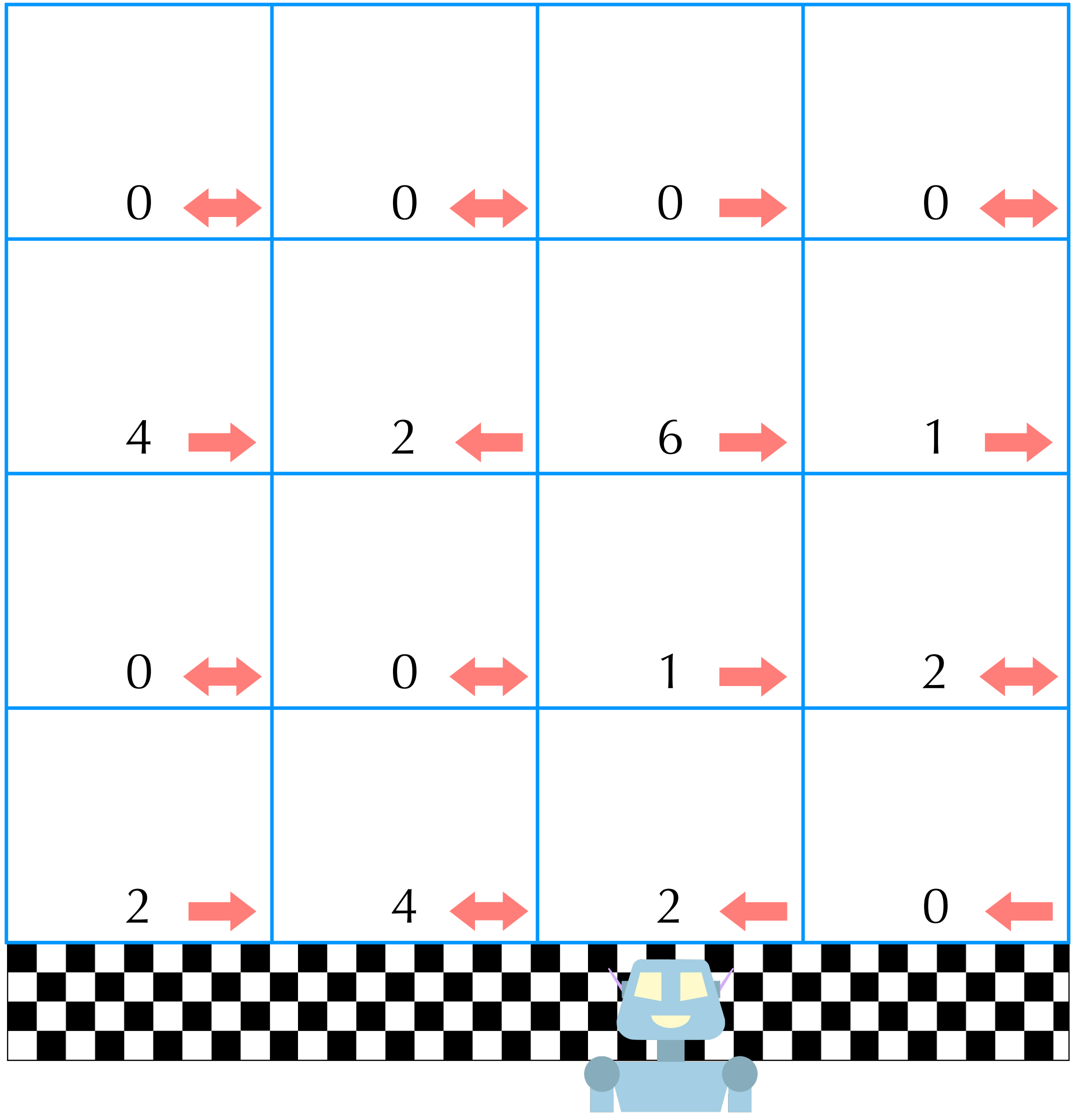

5

CONICET

UNC

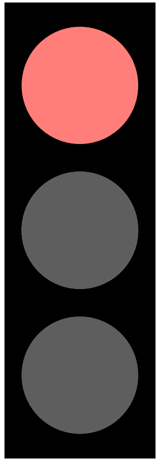

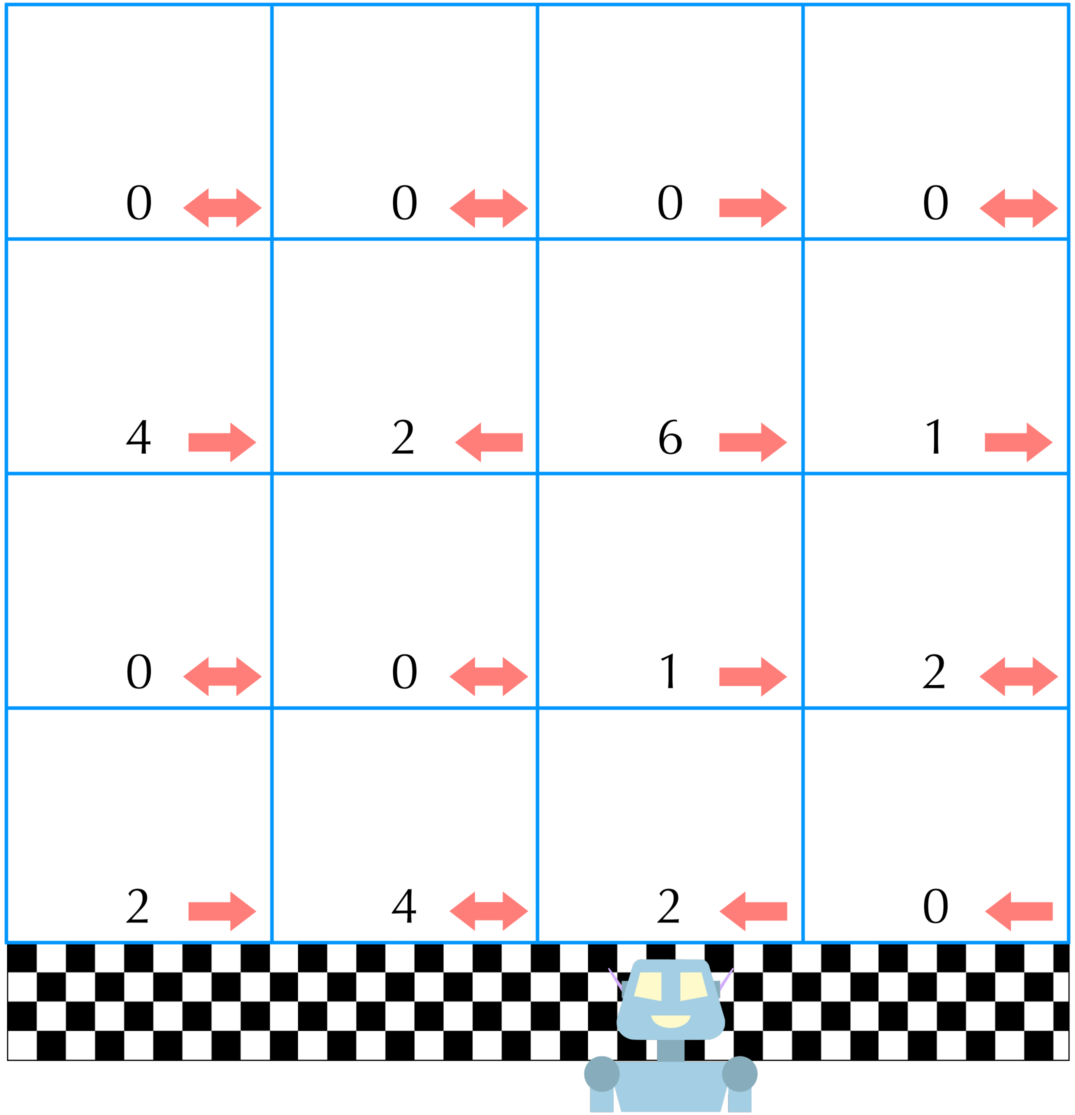

5

Stochastic game

or

2½-player game

CONICET

UNC

Stochastic games

A *stochastic game* is a tuple $\mathcal{G} = (V, (V_1, V_2, V_\mathsf{P}), \delta)$ where

1. $V = V_1 \uplus V_2 \uplus V_\mathsf{P}$ is a finite set of vertices (or states), and
2. $\delta : V \times V \to [0, 1]$ is a transition, such that
   a. for $v \in V_1 \cup V_2$, $\delta(v, \cdot) : V \to \{0, 1\}$ is the non-deterministic choice, and
   b. for $v \in V_\mathsf{P}$, $\delta(v, \cdot) : V \to [0, 1]$ is a probability function.

CONICET

UNC

Stochastic games

A *stochastic game* is a tuple $\mathcal{G} = (V, (V_1, V_2, V_\mathsf{P}), \delta)$ where

1. $V = V_1 \uplus V_2 \uplus V_\mathsf{P}$ is a finite set of vertices (or states), and
2. $\delta : V \times V \to [0, 1]$ is a transition, such that
   a. for $v \in V_1 \cup V_2$, $\delta(v, \cdot) : V \to \{0, 1\}$ is the non-deterministic choice, and
   b. for $v \in V_\mathsf{P}$, $\delta(v, \cdot) : V \to [0, 1]$ is a probability function.

- If $V_\mathsf{P} = \emptyset$, then $\mathcal{G}$ is a 2-*player game*.

Stochastic games

A *stochastic game* is a tuple $\mathcal{G} = (V, (V_1, V_2, V_\mathsf{P}), \delta)$ where

1. $V = V_1 \uplus V_2 \uplus V_\mathsf{P}$ is a finite set of vertices (or states), and
2. $\delta : V \times V \to [0,1]$ is a transition, such that
   a. for $v \in V_1 \cup V_2$, $\delta(v, \cdot) : V \to \{0,1\}$ is the non-deterministic choice, and
   b. for $v \in V_\mathsf{P}$, $\delta(v, \cdot) : V \to [0,1]$ is a probability function.

- If $V_\mathsf{P} = \emptyset$, then $\mathcal{G}$ is a 2-*player game*.
- If $V_1 = \emptyset$ or $V_2 = \emptyset$, then $\mathcal{G}$ is a *Markov Decision Process*.

Stochastic games

A *stochastic game* is a tuple $\mathcal{G} = (V, (V_1, V_2, V_\mathsf{P}), \delta)$ where

1. $V = V_1 \uplus V_2 \uplus V_\mathsf{P}$ is a finite set of vertices (or states), and
2. $\delta : V \times V \to [0, 1]$ is a transition, such that
   a. for $v \in V_1 \cup V_2$, $\delta(v, \cdot) : V \to \{0, 1\}$ is the non-deterministic choice, and
   b. for $v \in V_\mathsf{P}$, $\delta(v, \cdot) : V \to [0, 1]$ is a probability function.

- If $V_\mathsf{P} = \emptyset$, then $\mathcal{G}$ is a *2-player game*.
- If $V_1 = \emptyset$ or $V_2 = \emptyset$, then $\mathcal{G}$ is a *Markov Decision Process*.
- If $V_1 = V_2 = \emptyset$, then $\mathcal{G}$ is a *Markov chain*.

CONICET

UNC

Stochastic games

A *stochastic game* is a tuple $\mathcal{G} = (V, (V_1, V_2, V_\mathsf{P}), \delta)$ where

1. $V = V_1 \uplus V_2 \uplus V_\mathsf{P}$ is a finite set of vertices (or states), and
2. $\delta : V \times V \to [0, 1]$ is a transition, such that
   a. for $v \in V_1 \cup V_2$, $\delta(v, \cdot) : V \to \{0, 1\}$ is the non-deterministic choice, and
   b. for $v \in V_\mathsf{P}$, $\delta(v, \cdot) : V \to [0, 1]$ is a probability function.

We also consider a *reward function* $r : V \to \mathbb{R}^+$.

CONICET

UNC

Expected Total Reward

$$\mathbb{E}^{\pi_1,\pi_2}(rew)$$

$$rew(\rho) = \sum_{i=0}^{\infty} r(\rho(i))$$

Expected Total Reward

$$\mathbb{E}^{\pi_1,\pi_2}(rew)$$

$$rew(\rho) = \sum_{i=0}^{\infty} r(\rho(i))$$

The sum of all state rewards along path ρ

Expected Total Reward

$$\inf_{\pi_2 \in \Pi_2} \mathbb{E}^{\pi_1, \pi_2}(rew)$$

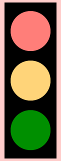

$$rew(\rho) = \sum_{i=0}^{\infty} r(\rho(i))$$

The sum of all state rewards along path ρ

CONICET

UNC

Expected Total Reward

$$\sup_{\pi_1 \in \Pi_1} \; \inf_{\pi_2 \in \Pi_2} \; \mathbb{E}^{\pi_1, \pi_2}(rew)$$

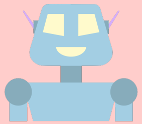

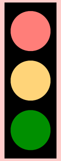

$$rew(\rho) = \sum_{i=0}^{\infty} r(\rho(i))$$

The sum of all state rewards along path ρ

Expected Total Reward

$$\sup_{\pi_1 \in \Pi_1} \inf_{\pi_2 \in \Pi_2} \mathbb{E}^{\pi_1, \pi_2}(rew)$$

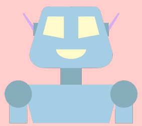

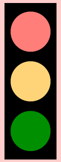

$$rew(\rho) = \sum_{i=0}^{\infty} r(\rho(i))$$

The sum of all state rewards along path ρ

It may go to infinity!

Expected Total Reward

$$\sup_{\pi_1 \in \Pi_1} \inf_{\pi_2 \in \Pi_2} \mathbb{E}^{\pi_1,\pi_2}(rew)$$

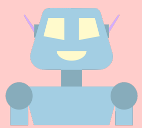

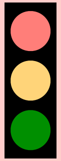

$$rew(\rho) = \sum_{i=0}^{\infty} r(\rho(i))$$

The sum of all state rewards along path ρ

~~It may go to infinity!~~

Stopping criteria

$$Prob^{\pi_1,\pi_2}(\Diamond \text{🏁}) = 1$$

for all $\pi_1 \in \Pi_1$ and $\pi_2 \in \Pi_2$

CONICET

UNC

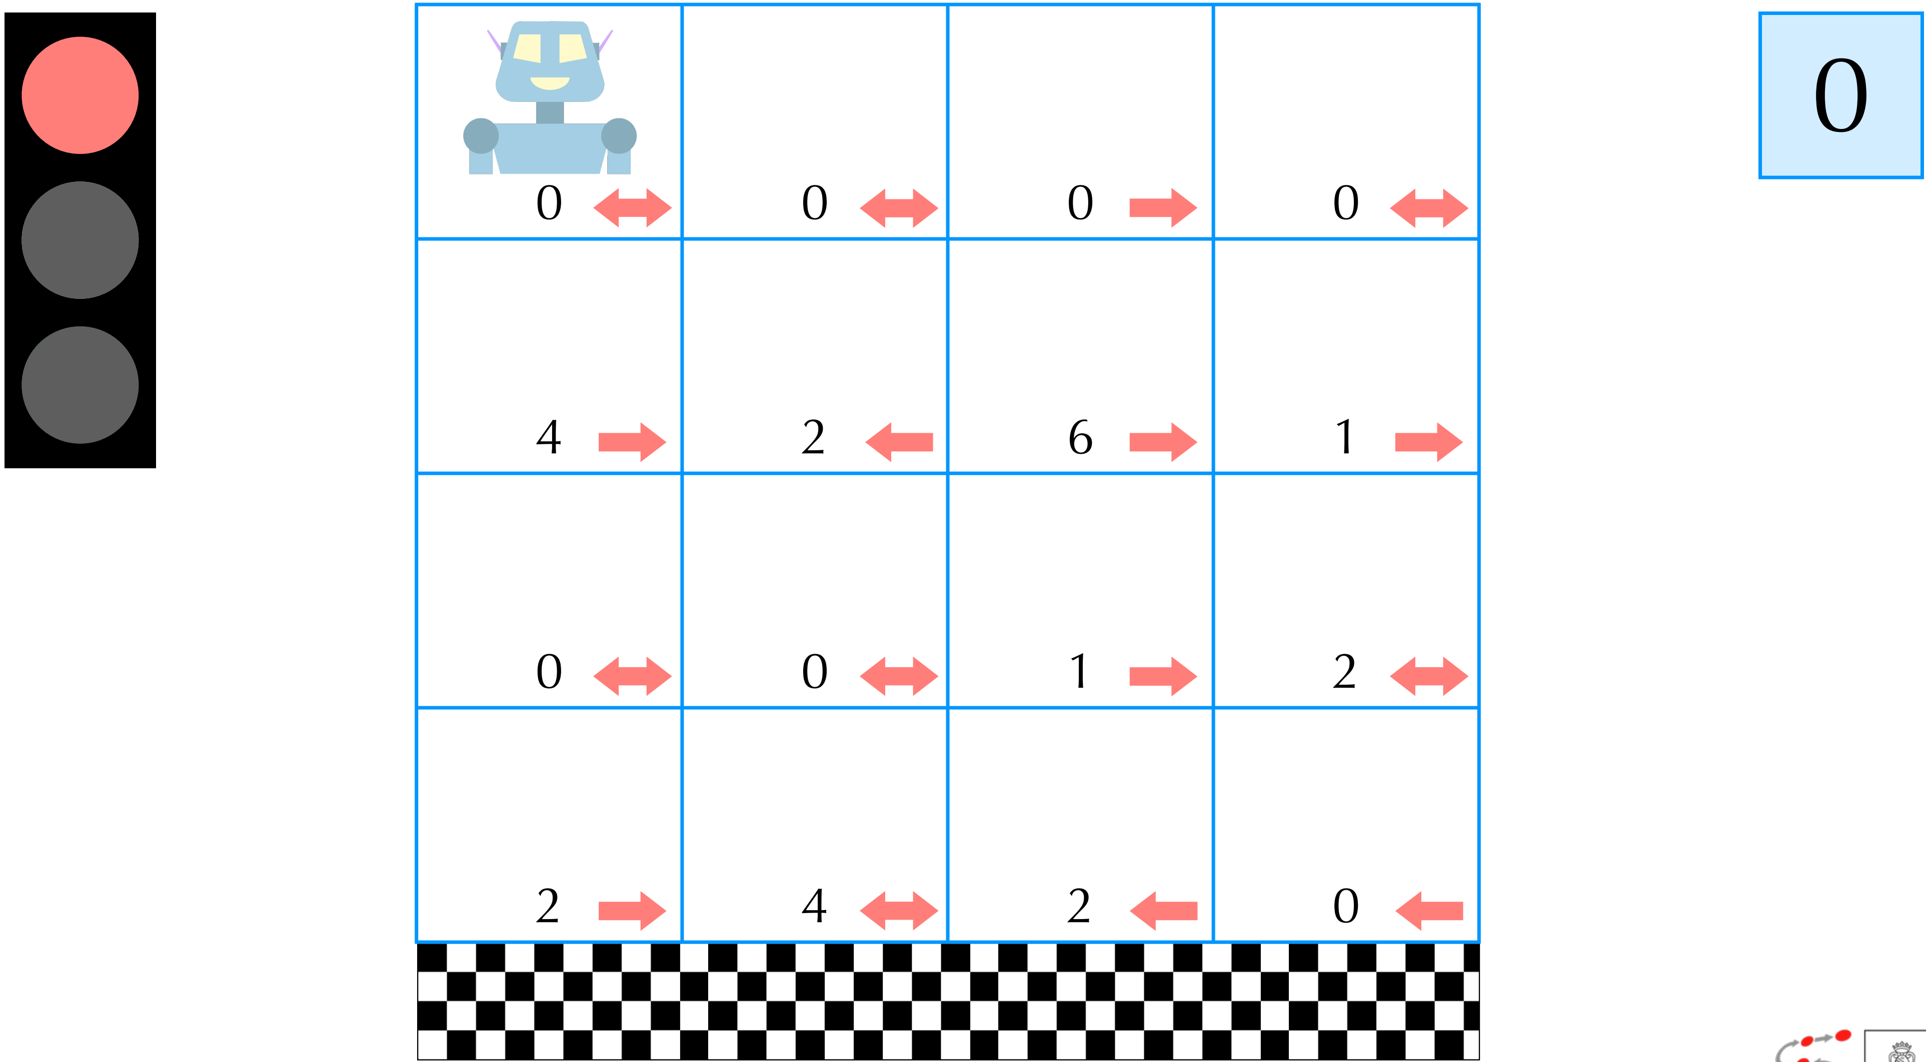

CONICET

UNC

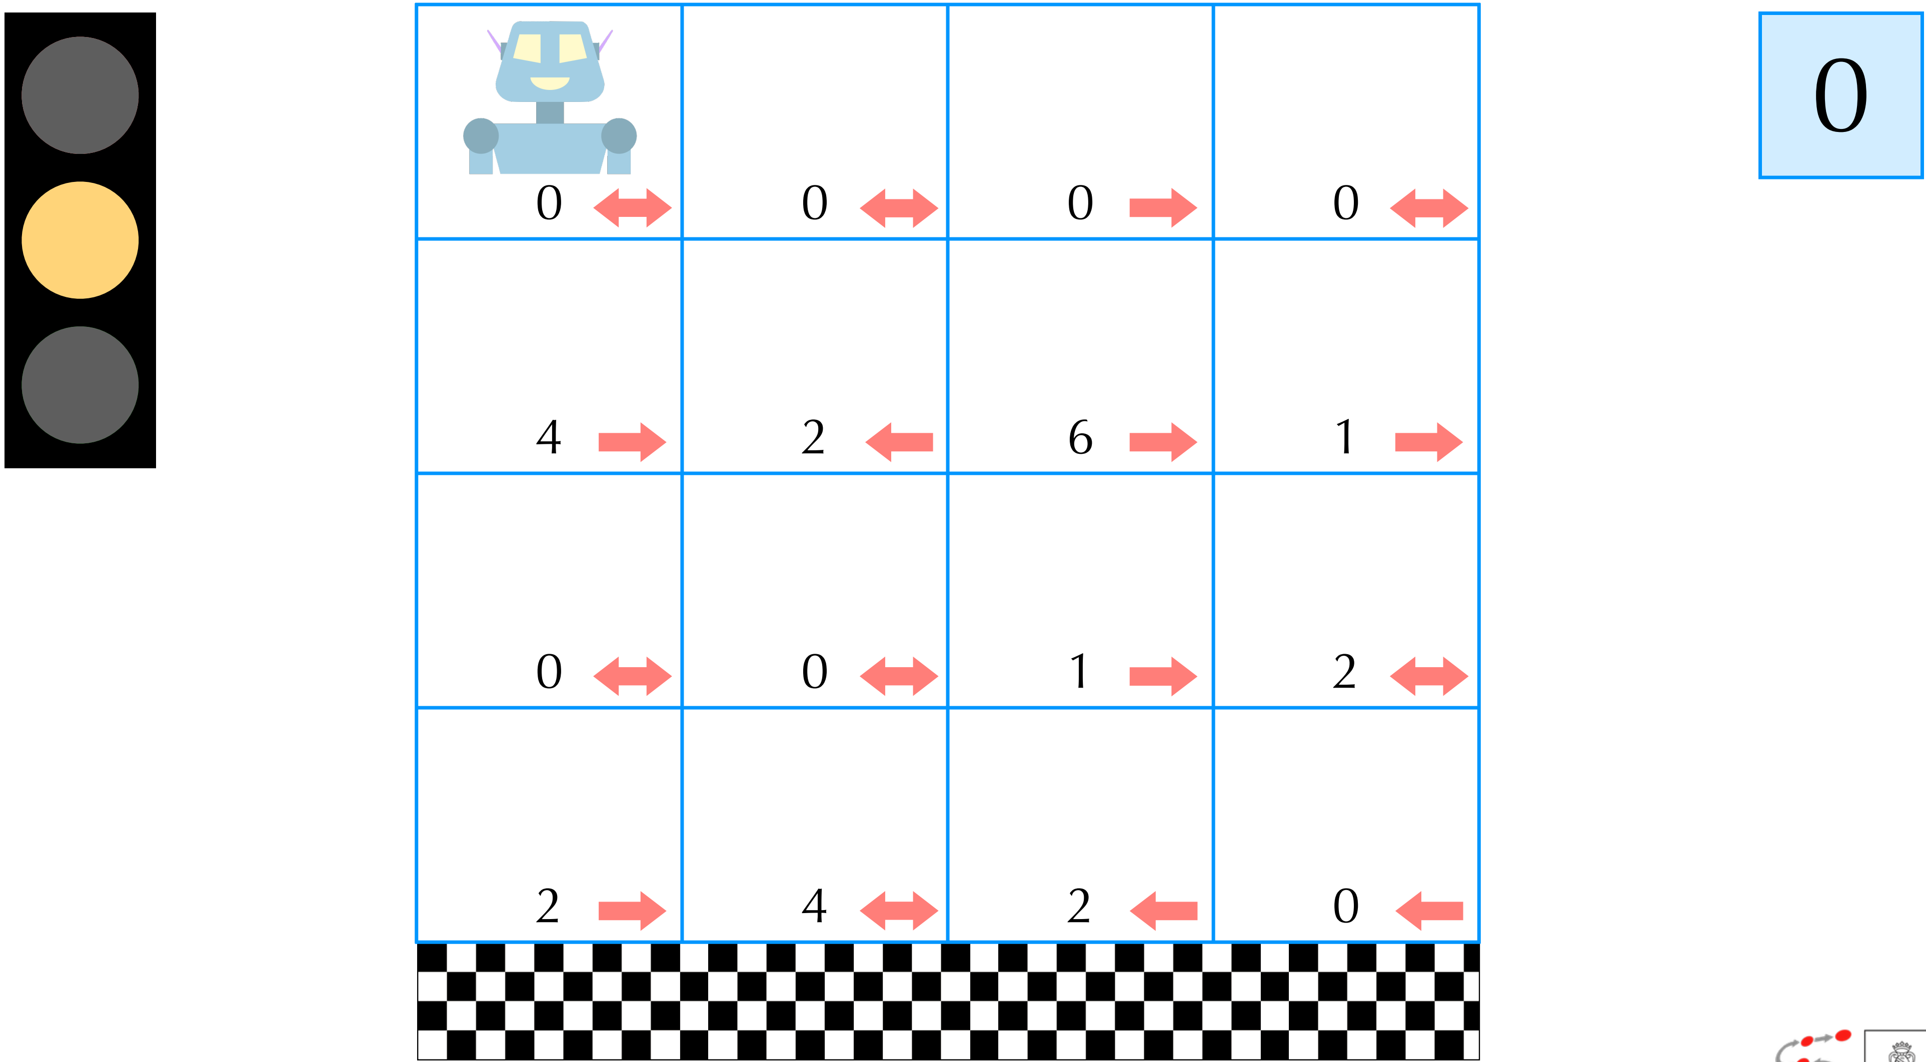

CONICET

UNC

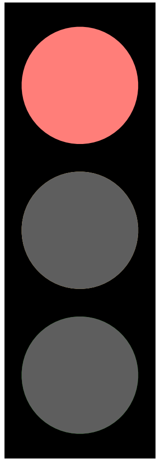

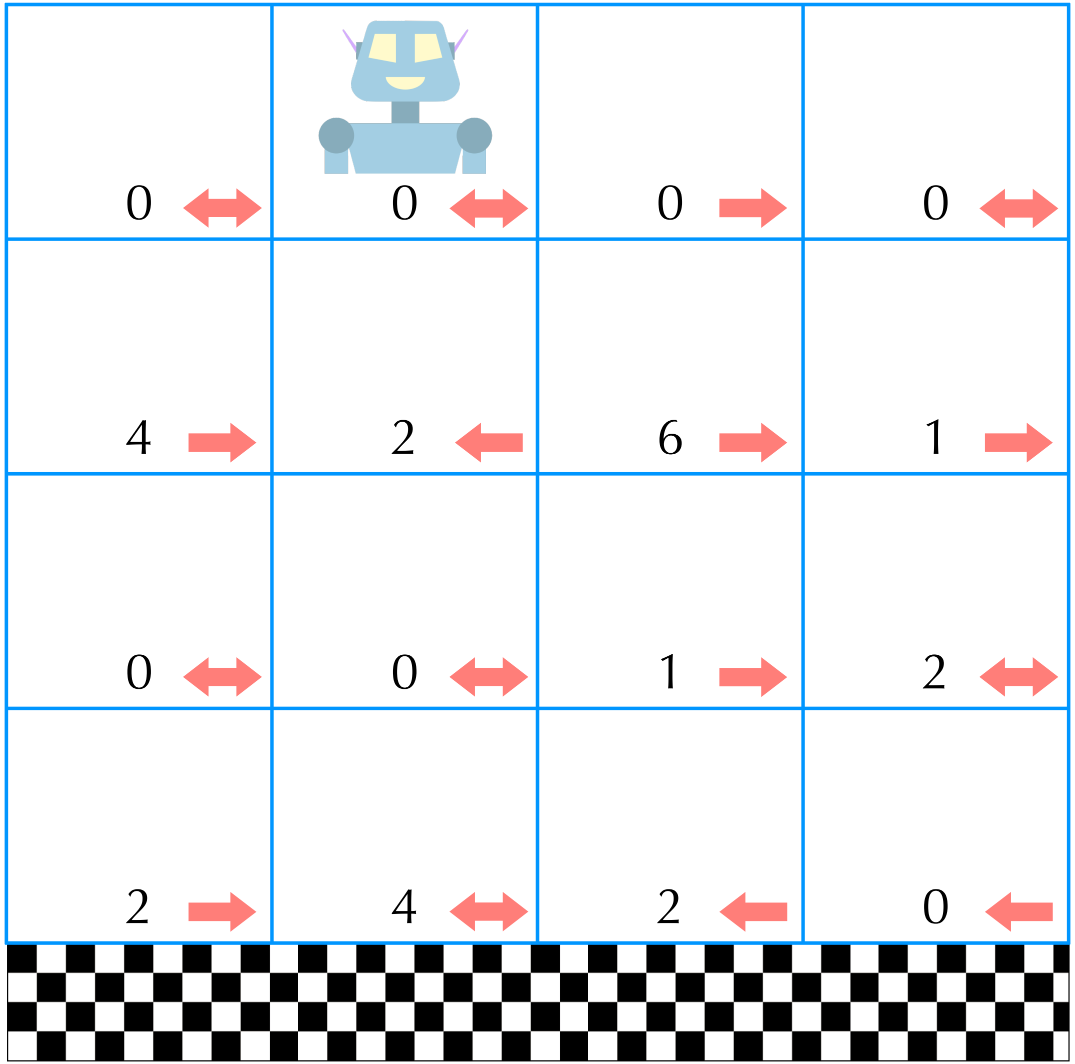

0

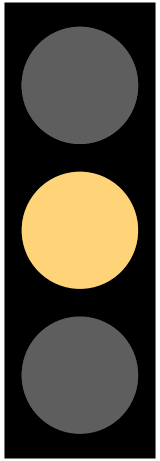

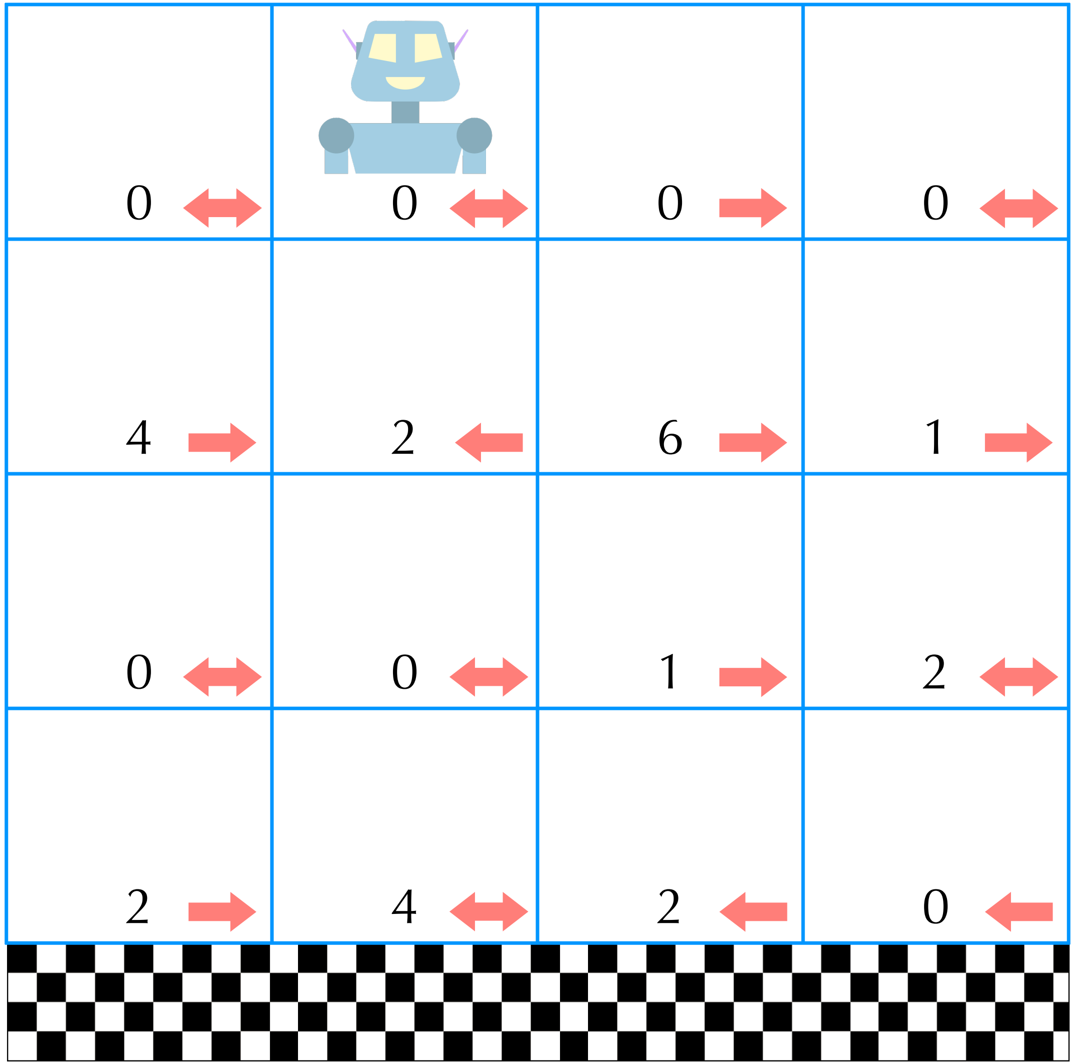

0

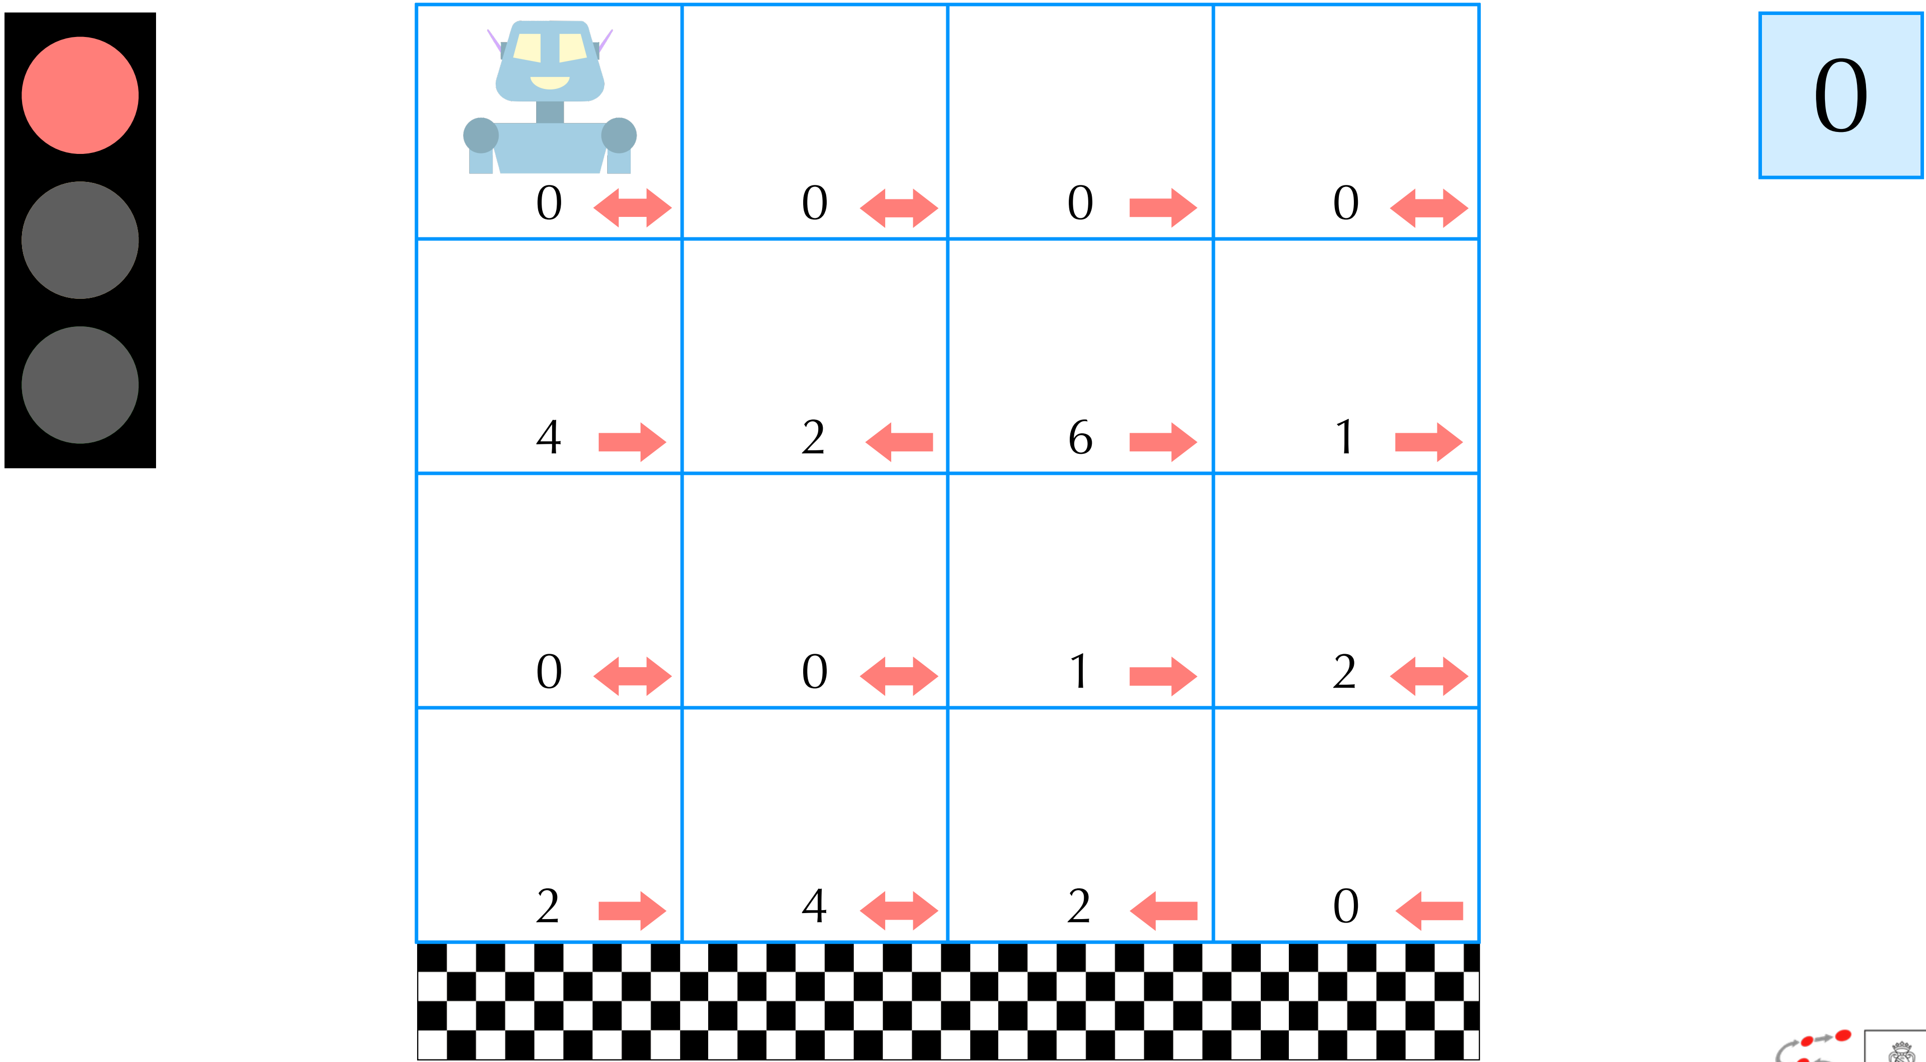

CONICET

UNC

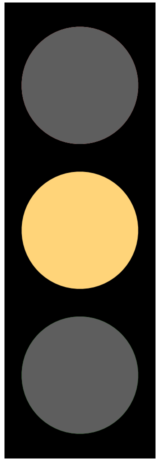

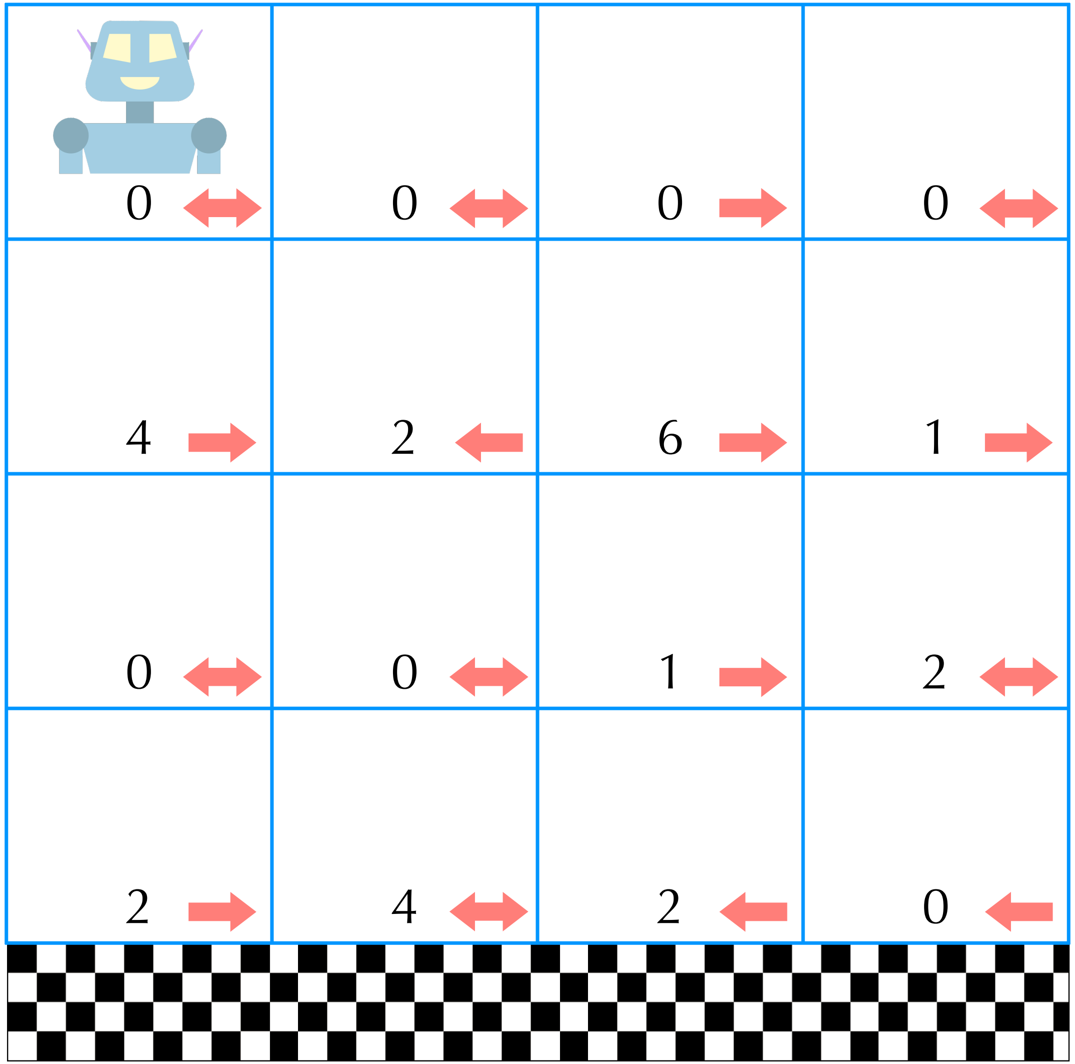

0

CONICET

UNC

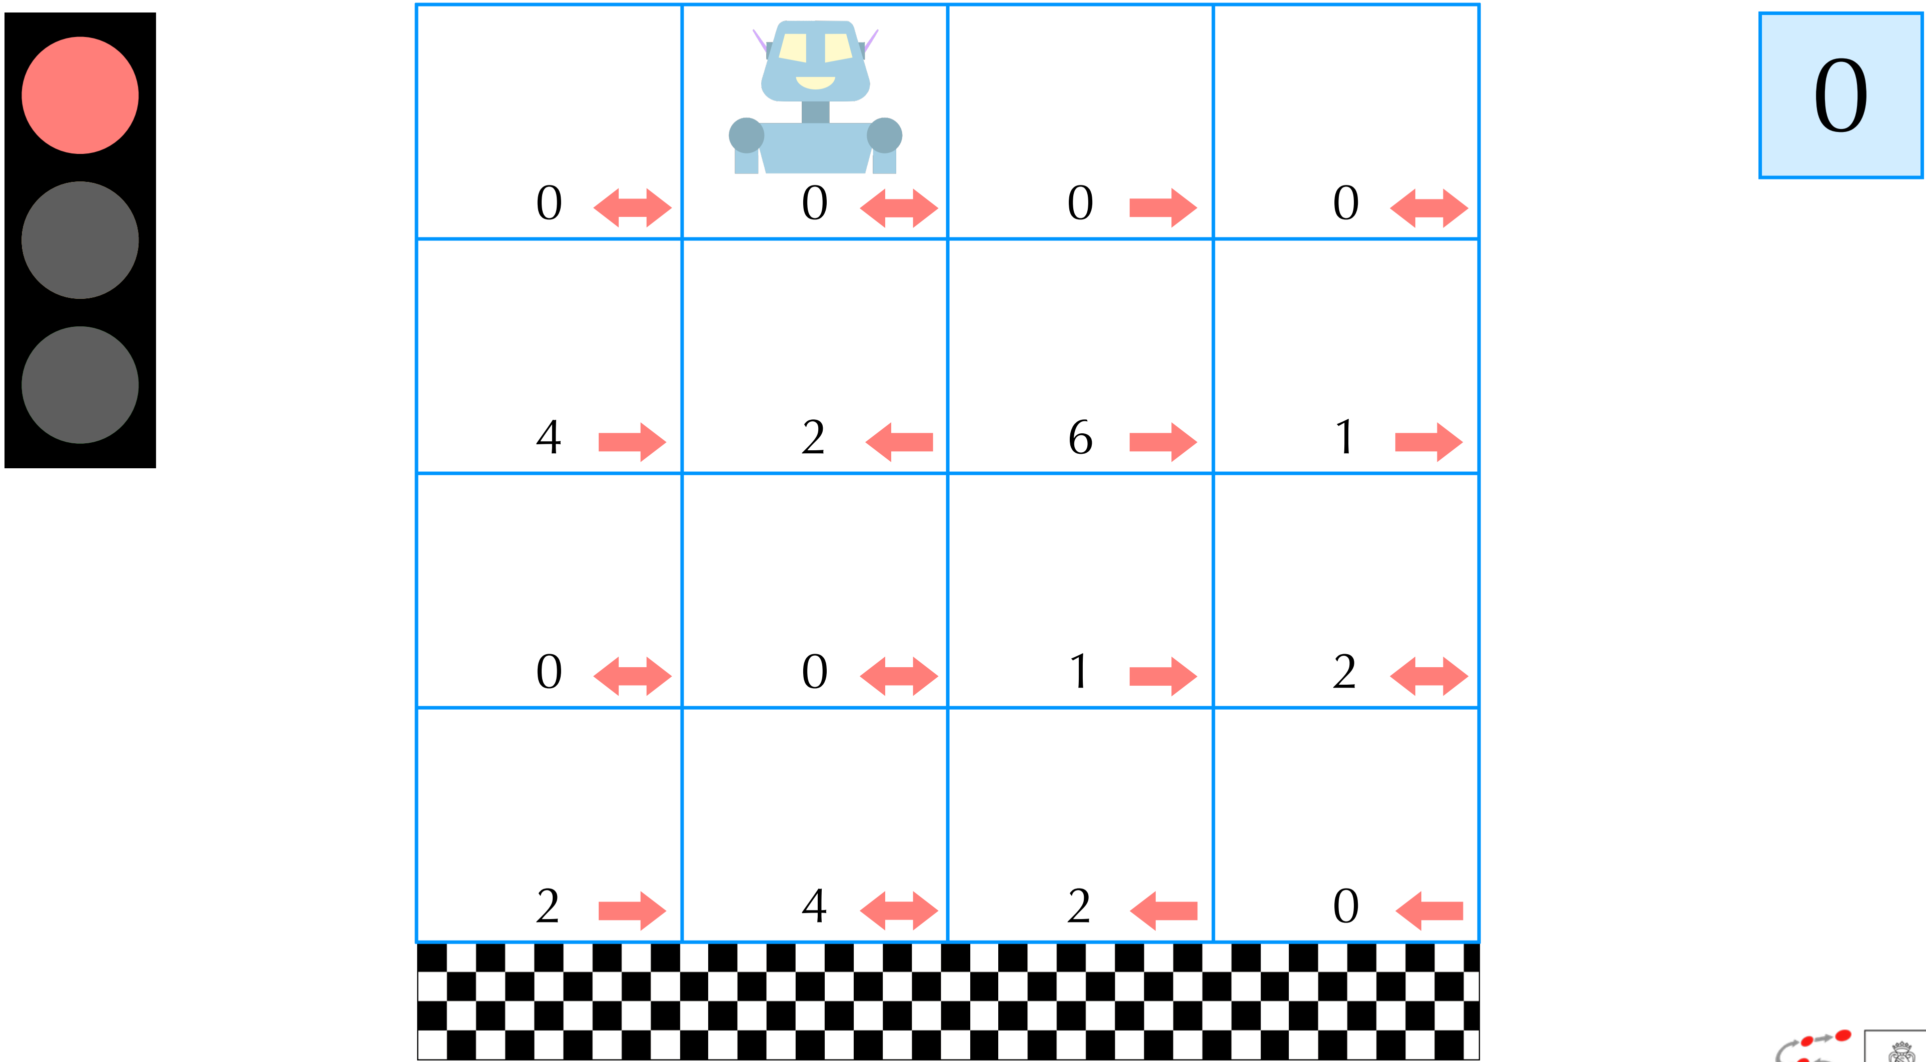

CONICET

UNC

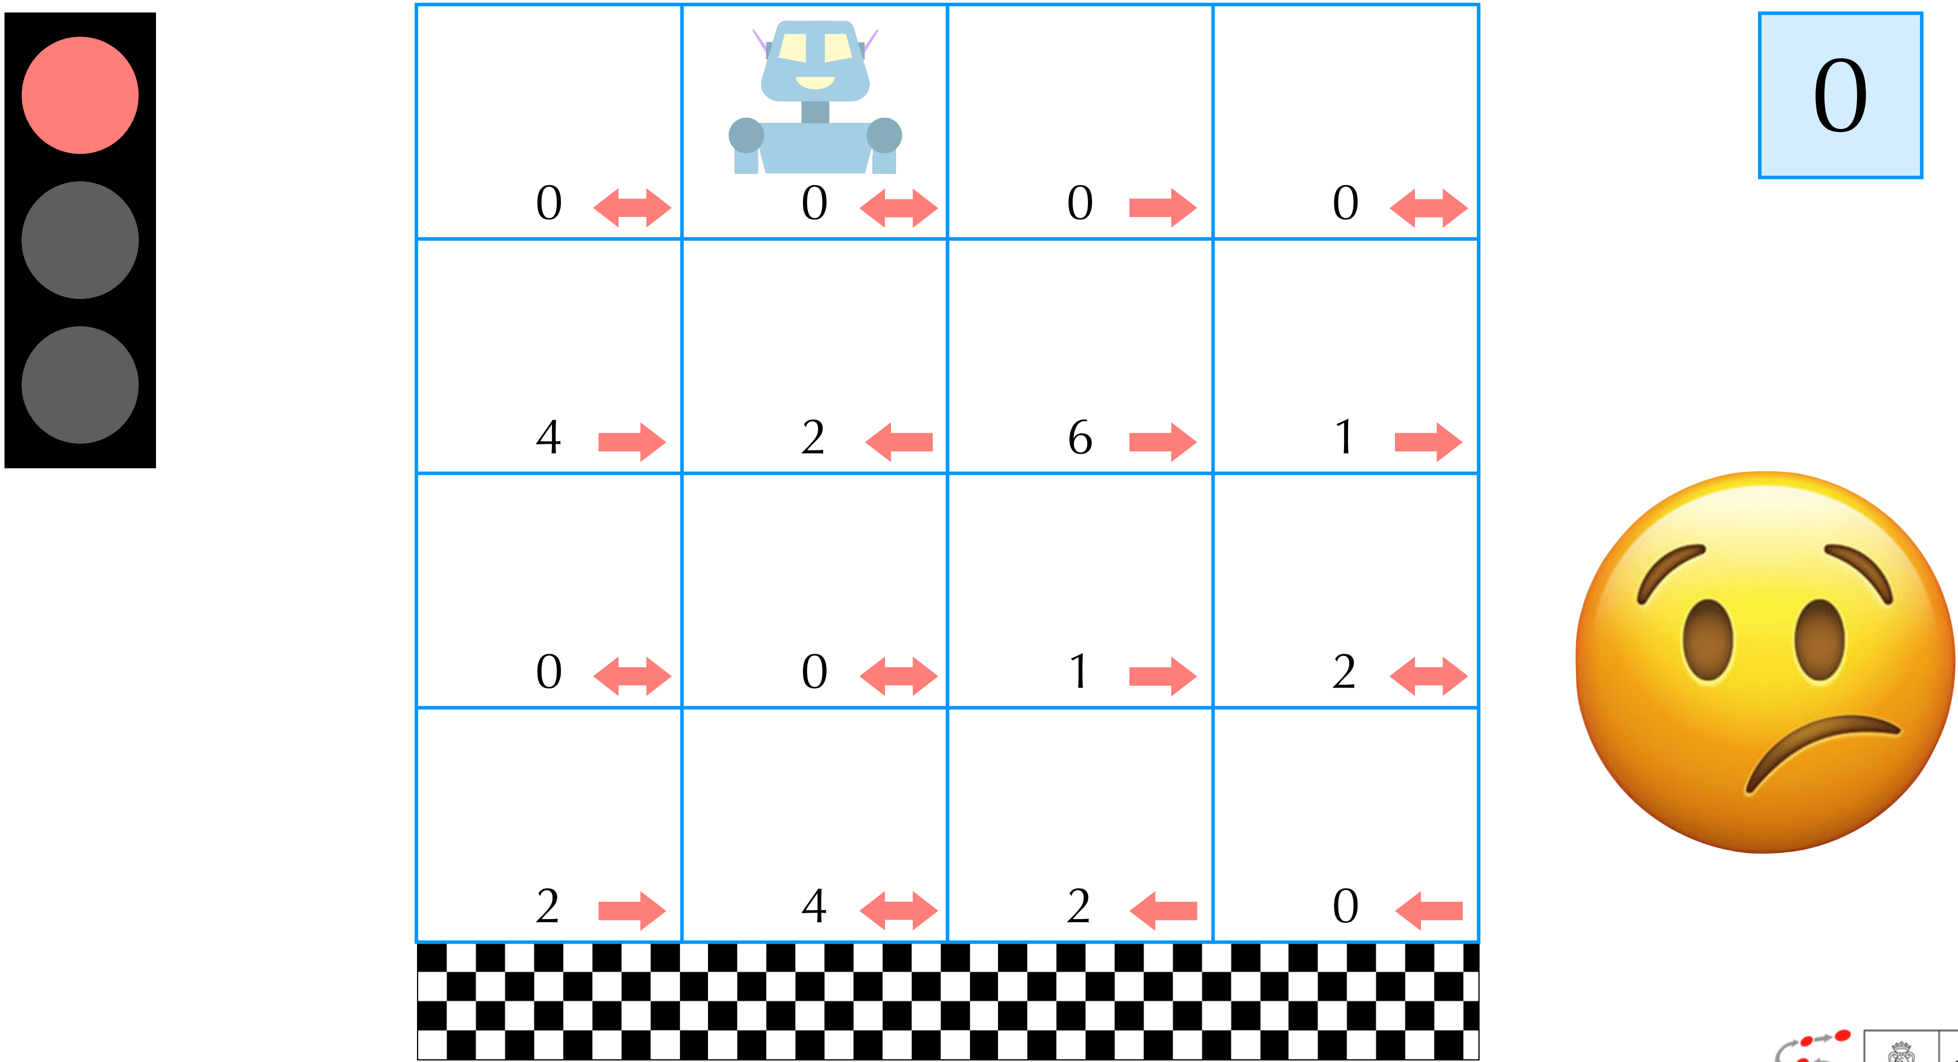

CONICET

UNC

The light must play fair

The light must play fair

Set of all states that repeats infinitely often in ω

Set of fair plays (for Player 2):

$$FP = \{\omega \in Paths_{\mathcal{G}} \mid \forall v' \in V_2 : v' \in \inf(\omega) \Rightarrow post(v') \subseteq \inf(\omega)\}$$

$$post(v) = \{v' \in V \mid \delta(v, v') > 0\}$$

The light must play fair

Set of fair plays (for Player 2):

$$FP = \{\omega \in Paths_{\mathcal{G}} \mid \forall v' \in V_2 : v' \in \inf(\omega) \Rightarrow post(v') \subseteq \inf(\omega)\}$$

A strategy $\pi_2 \in \Pi_2$ is *(almost-sure) fair* if for all $\pi_1 \in \Pi_1$

$$Prob^{\pi_1,\pi_2}(FP) = 1$$

The light must play fair

Set of fair plays (for Player 2):

$$FP = \{\omega \in Paths_{\mathcal{G}} \mid \forall v' \in V_2 : v' \in \inf(\omega) \Rightarrow post(v') \subseteq \inf(\omega)\}$$

A strategy $\pi_2 \in \Pi_2$ is *(almost-sure) fair* if for all $\pi_1 \in \Pi_1$

$$Prob^{\pi_1,\pi_2}(FP) = 1$$

A game is *stopping under fairness* if for every $\pi_1 \in \Pi_1$ and every fair $\pi_2 \in \Pi_2^{\mathcal{F}}$,

$$Prob^{\pi_1,\pi_2}(\Diamond 🏁) = 1$$

CONICET

UNC

Checking stopping under fairness

Theorem: A game is stopping under fairness iff for every $\pi_1 \in \Pi_1$

$$Prob^{\pi_1, \pi_2^u}(\Diamond \text{🏁}) = 1$$

where π_2^u is the strategy that choses uniformly a transition.

Checking stopping under fairness

Theorem: A game is stopping under fairness iff for every $\pi_1 \in \Pi_1$

$$Prob^{\pi_1,\pi_2^u}(\Diamond 🏁) = 1$$

where π_2^u is the strategy that choses uniformly a transition.

By fixing π_2^u, obtain
the corresponding MDP and check

$$\inf_{\pi_1 \in \Pi_1} Prob^{\pi_1}(\Diamond 🏁) = 1$$

there!

CONICET

UNC

Checking stopping under fairness

Theorem: A game is stopping under fairness iff for every $\pi_1 \in \Pi_1$

$$Prob^{\pi_1, \pi_2^u}(\Diamond \text{🏁}) = 1$$

where π_2^u is the strategy that choses uniformly a transition.

By fixing π_2^u, obtain the corresponding MDP and check

$$\inf_{\pi_1 \in \Pi_1} Prob^{\pi_1}(\Diamond \text{🏁}) = 1$$

there!

Actually, we check if the initial state does not belong to the set

$$\exists Pre_f^*(V \setminus \forall Pre_f^*(\text{🏁}))$$

where

$$\exists Pre_f(C) = \{v \in V \mid \delta(v, C) > 0\}$$

$$\forall Pre_f(C) = \{v \in V_{\mathsf{P}} \cup V_2 \mid \delta(v, C) > 0\} \cup \{v \in V_1 \mid \forall v' \in V : \delta(v, v') > 0 \Rightarrow v' \in C\}$$

CONICET

UNC

Checking stopping under fairness

Theorem: A game is stopping under fairness iff for every $\pi_1 \in \Pi_1$

$$Prob^{\pi_1,\pi_2^u}(\Diamond\,🏁) = 1$$

where π_2^u is the strategy that choses uniformly a transition.

By fixing π_2^u, obtain the corresponding MDP and check

$$\inf_{\pi_1\in\Pi_1} Prob^{\pi_1}(\Diamond\,🏁) = 1$$

there!

Actually, we check if the initial state does not belong to the set

$$\exists Pre_f^*(V \setminus \forall Pre_f^*(🏁))$$

where

The set of all states that reach some state in C

$$\exists Pre_f(C) = \{v \in V \mid \delta(v, C) > 0\}$$

$$\forall Pre_f(C) = \{v \in V_P \cup V_2 \mid \delta(v, C) > 0\} \cup \{v \in V_1 \mid \forall v' \in V : \delta(v, v') > 0 \Rightarrow v' \in C\}$$

CONICET

UNC

Checking stopping under fairness

Theorem: A game is stopping under fairness iff for every $\pi_1 \in \Pi_1$

$$Prob^{\pi_1, \pi_2^u}(\Diamond \text{🏁}) = 1$$

where π_2^u is the strategy that choses uniformly a transition.

By fixing π_2^u, obtain the corresponding MDP and check

$$\inf_{\pi_1 \in \Pi_1} Prob^{\pi_1}(\Diamond \text{🏁}) = 1$$

there!

Actually, we check if the initial state does <u>not</u> belong to the set

$$\exists Pre_f^*(V \setminus \forall Pre_f^*(\text{🏁}))$$

The set of all states in $V_P \cup V_2$ that reach <u>some</u> state in C

where

$$\exists Pre_f(C) = \{v \in V \mid \delta(v, C) > 0\}$$

$$\forall Pre_f(C) = \{v \in V_P \cup V_2 \mid \delta(v, C) > 0\} \cup \{v \in V_1 \mid \forall v' \in V : \delta(v, v') > 0 \Rightarrow v' \in C\}$$

CONICET

UNC

Checking stopping under fairness

Theorem: A game is stopping under fairness iff for every $\pi_1 \in \Pi_1$

$$Prob^{\pi_1,\pi_2^u}(\Diamond \text{🏁}) = 1$$

where π_2^u is the strategy that choses uniformly a transition.

By fixing π_2^u, obtain the corresponding MDP and check

$$\inf_{\pi_1 \in \Pi_1} Prob^{\pi_1}(\Diamond \text{🏁}) = 1$$

there!

Actually, we check if the initial state does not belong to the set

$$\exists Pre_f^*(V \setminus \forall Pre_f^*(\text{🏁}))$$

where

$$\exists Pre_f(C) = \{v \in V \mid \delta(v, C) > 0\}$$

$$\forall Pre_f(C) = \{v \in V_{\mathsf{P}} \cup V_2 \mid \delta(v, C) > 0\} \cup \{v \in V_1 \mid \forall v' \in V : \delta(v, v') > 0 \Rightarrow v' \in C\}$$

The set of all states in $V_{\mathsf{P}} \cup V_2$ that reach some state in C

The set of all states in V_1 that reach all states in C

CONICET

UNC

Checking stopping under fairness

Theorem: A game is stopping under fairness iff for every $\pi_1 \in \Pi_1$

$$\mathit{Prob}^{\pi_1,\pi_2^u}(\Diamond\, 🏁) = 1$$

where π_2^u is the strategy that choses uniformly a transition.

By fixing π_2^u, obtain the corresponding MDP and check

$$\inf_{\pi_1 \in \Pi_1} \mathit{Prob}^{\pi_1}(\Diamond\, 🏁) = 1$$

there!

Actually, we check if the initial state does not belong to the set

$$\exists \mathit{Pre}_f^*(V \setminus \forall \mathit{Pre}_f^*(🏁))$$

where

$$\exists \mathit{Pre}_f(C) = \{v \in V \mid \delta(v, C) > 0\}$$

$$\forall \mathit{Pre}_f(C) = \{v \in V_\mathsf{P} \cup V_2 \mid \delta(v, C) > 0\} \cup \{v \in V_1 \mid \forall v' \in V : \delta(v, v') > 0 \Rightarrow v' \in C\}$$

It can be calculated in polynomial time

Two technical results

Theorem: Determinacy.

$$\inf_{\pi_2 \in \Pi_2^{\mathcal{F}}} \sup_{\pi_1 \in \Pi_1} \mathbb{E}^{\pi_1,\pi_2}(rew) = \sup_{\pi_1 \in \Pi_1} \inf_{\pi_2 \in \Pi_2^{\mathcal{F}}} \mathbb{E}^{\pi_1,\pi_2}(rew)$$

Two technical results

Theorem: Determinacy.

$$\inf_{\pi_2 \in \Pi_2^{\mathcal{F}}} \sup_{\pi_1 \in \Pi_1} \mathbb{E}^{\pi_1,\pi_2}(rew) = \sup_{\pi_1 \in \Pi_1} \inf_{\pi_2 \in \Pi_2^{\mathcal{F}}} \mathbb{E}^{\pi_1,\pi_2}(rew)$$

Moreover

$< \infty$

CONICET

UNC

Two technical results

Theorem: Determinacy.

$$\inf_{\pi_2 \in \Pi_2^{\mathcal{F}}} \sup_{\pi_1 \in \Pi_1} \mathbb{E}^{\pi_1,\pi_2}(rew) = \sup_{\pi_1 \in \Pi_1} \inf_{\pi_2 \in \Pi_2^{\mathcal{F}}} \mathbb{E}^{\pi_1,\pi_2}(rew)$$

Moreover

$< \infty$

Theorem: Memoryless deterministic schedulers are sufficient.

$$\sup_{\pi_1 \in \Pi_1} \inf_{\pi_2 \in \Pi_2^{\mathcal{F}}} \mathbb{E}^{\pi_1,\pi_2}(rew) = \sup_{\pi_1 \in \Pi_1^{MD}} \inf_{\pi_2 \in \Pi_2^{MD\mathcal{F}}} \mathbb{E}^{\pi_1,\pi_2}(rew)$$

CONICET

UNC

Two technical results

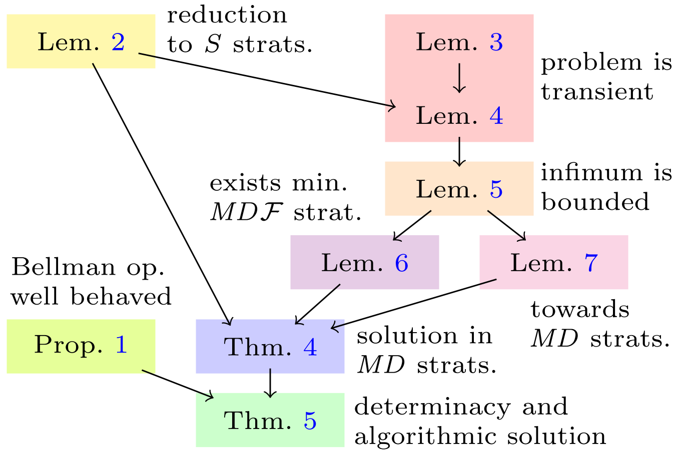

Theorem: Determinacy.

$$\inf_{\pi_2 \in \Pi_2^{\mathcal{F}}} \sup_{\pi_1 \in \Pi_1} \mathbb{E}^{\pi_1,\pi_2}(rew) = \sup_{\pi_1 \in \Pi_1} \inf_{\pi_2 \in \Pi_2^{\mathcal{F}}} \mathbb{E}^{\pi_1,\pi_2}(rew)$$

Theorem: Memoryless deterministic schedulers are sufficient.

$$\sup_{\pi_1 \in \Pi_1} \inf_{\pi_2 \in \Pi_2^{\mathcal{F}}} \mathbb{E}^{\pi_1,\pi_2}(rew) = \sup_{\pi_1 \in \Pi_1^{MD}} \inf_{\pi_2 \in \Pi_2^{MD\mathcal{F}}} \mathbb{E}^{\pi_1,\pi_2}(rew)$$

CONICET

UNC

Two technical results

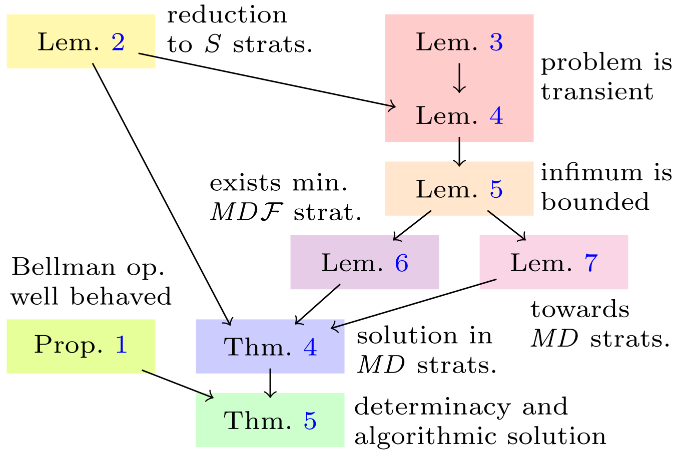

Theorem: Determinacy.

$$\inf_{\pi_2 \in \Pi_2^{\mathcal{F}}} \sup_{\pi_1 \in \Pi_1} \mathbb{E}^{\pi_1,\pi_2}(rew) = \sup_{\pi_1 \in \Pi_1} \inf_{\pi_2 \in \Pi_2^{\mathcal{F}}} \mathbb{E}^{\pi_1,\pi_2}(rew)$$

Theorem: Memoryless deterministic schedulers are sufficient.

$$\sup_{\pi_1 \in \Pi_1} \inf_{\pi_2 \in \Pi_2^{\mathcal{F}}} \mathbb{E}^{\pi_1,\pi_2}(rew) = \sup_{\pi_1 \in \Pi_1^{MD}} \inf_{\pi_2 \in \Pi_2^{MD\mathcal{F}}} \mathbb{E}^{\pi_1,\pi_2}(rew)$$

Thus, the problem could be solved as a fix point calculation on the Bellman equations

CONICET

UNC

Algorithmic solution

Proposal: Solve the next Bellman operator

$$\Gamma(f)(v) = \begin{cases} r(v) + \sum_{v' \in post(v)} \delta(v, v') f(v') & \text{if } v \in V_{\mathsf{P}} \setminus \{🏁\} \\ \max\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_1 \setminus \{🏁\} \\ \min\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_2 \setminus \{🏁\} \\ 0 & \text{if } v = 🏁 \end{cases}$$

Algorithmic solution

Proposal: Solve the next Bellman operator

$$\Gamma(f)(v) = \begin{cases} r(v) + \sum_{v' \in post(v)} \delta(v, v') f(v') & \text{if } v \in V_{\mathsf{P}} \setminus \{🏁\} \\ \max\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_1 \setminus \{🏁\} \\ \min\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_2 \setminus \{🏁\} \\ 0 & \text{if } v = 🏁 \end{cases}$$

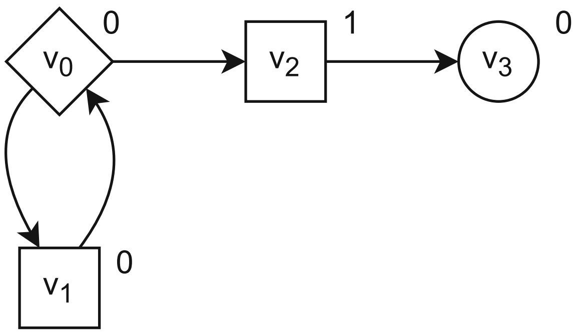

CONICET

UNC

Algorithmic solution

Proposal: Solve the next Bellman operator

$$\Gamma(f)(v) = \begin{cases} r(v) + \sum_{v' \in post(v)} \delta(v, v') f(v') & \text{if } v \in V_{\mathsf{P}} \setminus \{🏁\} \\ \max\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_1 \setminus \{🏁\} \\ \min\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_2 \setminus \{🏁\} \\ 0 & \text{if } v = 🏁 \end{cases}$$

any $(x, x, 1, 0)$ with $x \in [0,1]$ is a solution!

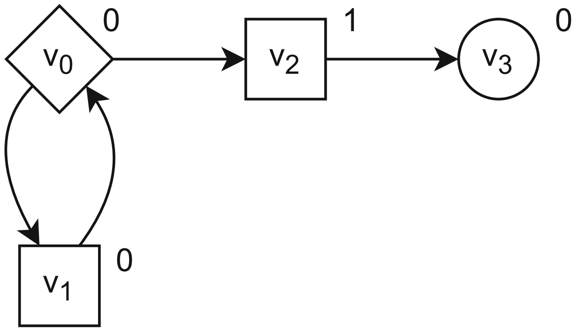

Algorithmic solution

Proposal: Solve the next Bellman operator

$$\Gamma(f)(v) = \begin{cases} r(v) + \sum_{v' \in post(v)} \delta(v,v') f(v') & \text{if } v \in V_{\mathsf{P}} \setminus \{\text{🏁}\} \\ \max\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_1 \setminus \{\text{🏁}\} \\ \min\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_2 \setminus \{\text{🏁}\} \\ 0 & \text{if } v = \text{🏁} \end{cases}$$

any $(x, x, 1, 0)$ with $x \in [0,1]$ is a solution!

Problem: Γ does not have a unique fixpoint

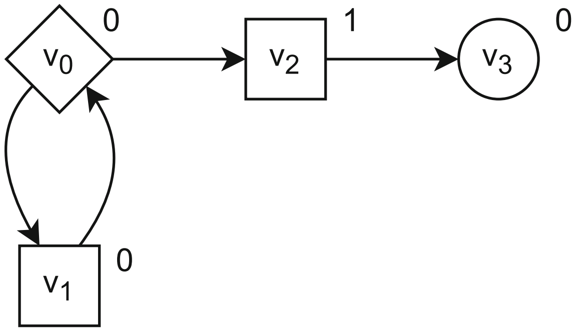

CONICET

UNC

Algorithmic solution

Proposal: Solve the next Bellman operator

$$\Gamma(f)(v) = \begin{cases} r(v) + \sum_{v' \in post(v)} \delta(v, v') f(v') & \text{if } v \in V_{\mathsf{P}} \setminus \{\text{🏁}\} \\ \max\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_1 \setminus \{\text{🏁}\} \\ \min\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_2 \setminus \{\text{🏁}\} \\ 0 & \text{if } v = \text{🏁} \end{cases}$$

Problem: Γ does not have a unique fixpoint

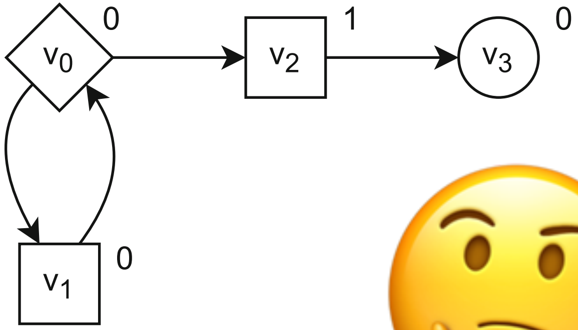

any $(x, x, 1, 0)$ with $x \in [0,1]$ is a solution!

The solution has to be the <u>greatest</u> fixpoint in $(\mathbb{R} \cup \{\infty\})^V$

CONICET

UNC

Algorithmic solution

Proposal: Solve the next Bellman operator

$$\Gamma(f)(v) = \begin{cases} r(v) + \sum_{v' \in post(v)} \delta(v, v') f(v') & \text{if } v \in V_{\mathsf{P}} \setminus \{\text{🏁}\} \\ \max\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_1 \setminus \{\text{🏁}\} \\ \min\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_2 \setminus \{\text{🏁}\} \\ 0 & \text{if } v = \text{🏁} \end{cases}$$

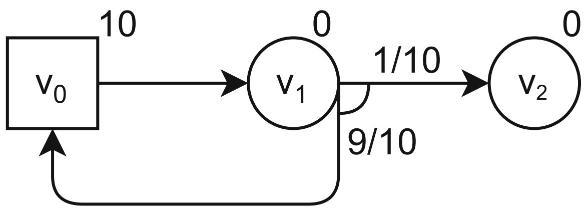

CONICET

UNC

Algorithmic solution

Proposal: Solve the next Bellman operator

$$\Gamma(f)(v) = \begin{cases} r(v) + \sum_{v' \in post(v)} \delta(v, v') f(v') & \text{if } v \in V_{\mathsf{P}} \setminus \{🏁\} \\ \max\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_1 \setminus \{🏁\} \\ \min\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_2 \setminus \{🏁\} \\ 0 & \text{if } v = 🏁 \end{cases}$$

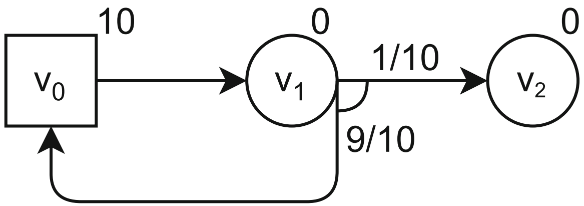

$(\infty, \infty, 0)$ is the greatest fixpont!

CONICET

UNC

Algorithmic solution

Proposal: Solve the next Bellman operator

$$\Gamma(f)(v) = \begin{cases} r(v) + \sum_{v' \in post(v)} \delta(v, v') f(v') & \text{if } v \in V_{\mathsf{P}} \setminus \{\text{🏁}\} \\ \max\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_1 \setminus \{\text{🏁}\} \\ \min\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_2 \setminus \{\text{🏁}\} \\ 0 & \text{if } v = \text{🏁} \end{cases}$$

Problem: Γ greatest fixpoint in the extended reals may be outside the reals!

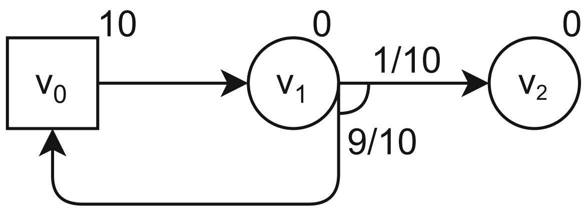

$(\infty, \infty, 0)$ is the greatest fixpont!

Algorithmic solution

$$\Gamma(f)(v) = \begin{cases} r(v) + \sum_{v' \in post(v)} \delta(v, v') f(v') & \text{if } v \in V_{\mathsf{P}} \setminus \{🏁\} \\ \max\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_1 \setminus \{🏁\} \\ \min\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_2 \setminus \{🏁\} \\ 0 & \text{if } v = 🏁 \end{cases}$$

Algorithmic solution

$$\Gamma(f)(v) = \begin{cases} r(v) + \sum_{v' \in post(v)} \delta(v, v') f(v') & \text{if } v \in V_{\mathsf{P}} \setminus \{\text{🏁}\} \\ \max\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_1 \setminus \{\text{🏁}\} \\ \min\{r(v) + f(v') \mid v' \in post(v)\} & \text{if } v \in V_2 \setminus \{\text{🏁}\} \\ 0 & \text{if } v = \text{🏁} \end{cases}$$

Algorithmic solution

Let $\mathbf{U} \geq \max_{v \in V} \sup_{\pi_1 \in \Pi_1^{MD}} \inf_{\pi_2 \in \Pi_2^{MD\mathcal{F}}} \mathbb{E}_v^{\pi_1,\pi_2}(rew)$

$$\Gamma(f)(v) = \begin{cases} \min\left(r(v) + \sum_{v' \in post(v)} \delta(v,v') f(v'),\ \mathbf{U}\right) & \text{if } v \in V_{\mathsf{P}} \setminus \{🏁\} \\ \min\left(\max\{r(v) + f(v') \mid v' \in post(v)\},\ \mathbf{U}\right) & \text{if } v \in V_1 \setminus \{🏁\} \\ \min\left(\min\{r(v) + f(v') \mid v' \in post(v)\},\ \mathbf{U}\right) & \text{if } v \in V_2 \setminus \{🏁\} \\ 0 & \text{if } v = 🏁 \end{cases}$$

CONICET

UNC

Algorithmic solution

Let $\mathbf{U} \geq \max_{v \in V} \sup_{\pi_1 \in \Pi_1^{MD}} \inf_{\pi_2 \in \Pi_2^{MD\mathcal{F}}} \mathbb{E}_v^{\pi_1,\pi_2}(rew)$

$$\Gamma(f)(v) = \begin{cases} \min\left(r(v) + \sum_{v' \in post(v)} \delta(v,v') f(v'),\ \mathbf{U}\right) & \text{if } v \in V_{\mathsf{P}} \setminus \{🏁\} \\ \min\left(\max\{r(v) + f(v') \mid v' \in post(v)\},\ \mathbf{U}\right) & \text{if } v \in V_1 \setminus \{🏁\} \\ \min\left(\min\{r(v) + f(v') \mid v' \in post(v)\},\ \mathbf{U}\right) & \text{if } v \in V_2 \setminus \{🏁\} \\ 0 & \text{if } v = 🏁 \end{cases}$$

Proposition: Γ is monotone and Scott-continuous in the lattice $[0, \mathbf{U}]^V$.

Algorithmic solution

Let $\quad \mathbf{U} \;\geq\; \max_{v \in V} \; \sup_{\pi_1 \in \Pi_1^{MD}} \; \inf_{\pi_2 \in \Pi_2^{MD\mathcal{F}}} \; \mathbb{E}_v^{\pi_1,\pi_2}(rew)$

$$\Gamma(f)(v) = \begin{cases} \min\left(r(v) + \sum_{v' \in post(v)} \delta(v,v') f(v'),\ \mathbf{U}\right) & \text{if } v \in V_{\mathsf{P}} \setminus \{🏁\} \\ \min\left(\max\{r(v) + f(v') \mid v' \in post(v)\},\ \mathbf{U}\right) & \text{if } v \in V_1 \setminus \{🏁\} \\ \min\left(\min\{r(v) + f(v') \mid v' \in post(v)\},\ \mathbf{U}\right) & \text{if } v \in V_2 \setminus \{🏁\} \\ 0 & \text{if } v = 🏁 \end{cases}$$

Proposition: Γ is monotone and Scott-continuous in the lattice $[0, \mathbf{U}]^V$.

Thus, the greatest fixpoint can be approximated from $\mathbf{U}^V$.

CONICET

Algorithmic solution

Let $\mathbf{U} \geq \max_{v \in V} \sup_{\pi_1 \in \Pi_1^{MD}} \inf_{\pi_2 \in \Pi_2^{MD\mathcal{F}}} \mathbb{E}_v^{\pi_1,\pi_2}(\mathit{rew})$

$$\Gamma(f)(v) = \begin{cases} \min\left(r(v) + \sum_{v' \in post(v)} \delta(v,v') f(v'),\ \mathbf{U}\right) & \text{if } v \in V_{\mathsf{P}} \setminus \{🏁\} \\ \min\left(\max\{r(v) + f(v') \mid v' \in post(v)\},\ \mathbf{U}\right) & \text{if } v \in V_1 \setminus \{🏁\} \\ \min\left(\min\{r(v) + f(v') \mid v' \in post(v)\},\ \mathbf{U}\right) & \text{if } v \in V_2 \setminus \{🏁\} \\ 0 & \text{if } v = 🏁 \end{cases}$$

Proposition: Γ is monotone and Scott-continuous in the lattice $[0, \mathbf{U}]^V$.

Theorem: For all $v \in V$, $\sup_{\pi_1 \in \Pi_1} \inf_{\pi_2 \in \Pi_2^{\mathcal{F}}} \mathbb{E}_v^{\pi_1,\pi_2}(\mathit{rew}) = \nu\,\Gamma(v)$

CONICET

UNC

Algorithmic solution

$$\mathbf{U} \;\geq\; \max_{v \in V} \;\sup_{\pi_1 \in \Pi_1^{MD}} \;\inf_{\pi_2 \in \Pi_2^{MD\mathcal{F}}} \;\mathbb{E}_v^{\pi_1,\pi_2}(\mathit{rew})$$

$$\Gamma(f)(v) = \begin{cases} \min\left(r(v) + \sum_{v' \in \mathit{post}(v)} \delta(v,v') f(v'),\; \mathbf{U}\right) & \text{if } v \in V_{\mathsf{P}} \setminus \{\text{🏁}\} \\ \min\left(\max\{r(v) + f(v') \mid v' \in \mathit{post}(v)\},\; \mathbf{U}\right) & \text{if } v \in V_1 \setminus \{\text{🏁}\} \\ \min\left(\min\{r(v) + f(v') \mid v' \in \mathit{post}(v)\},\; \mathbf{U}\right) & \text{if } v \in V_2 \setminus \{\text{🏁}\} \\ 0 & \text{if } v = \text{🏁} \end{cases}$$

Algorithmic solution

1. Calculate $\mathbf{U} = \max_{v \in V} \sup_{\pi_1 \in \Pi_1^{MD}} \mathbb{E}_v^{\pi_1}(\mathit{rew})$ on the MDP obtained by fixing π_2^u.

2. Starting on $x_v = \mathbf{U}$, approximate the maximum fixed point on the equations

$$x_v = \begin{cases} \min\left(r(v) + \sum_{v' \in \mathit{post}(v)} \delta(v,v') x_{v'},\ \mathbf{U}\right) & \text{if } v \in V_{\mathsf{P}} \setminus \{🏁\} \\ \min\left(\max\{r(v) + x_{v'} \mid v' \in \mathit{post}(v)\},\ \mathbf{U}\right) & \text{if } v \in V_1 \setminus \{🏁\} \\ \min\left(\min\{r(v) + x_{v'} \mid v' \in \mathit{post}(v)\},\ \mathbf{U}\right) & \text{if } v \in V_2 \setminus \{🏁\} \\ 0 & \text{if } v = 🏁 \end{cases}$$

3. Derive the optimizing strategies by traversing the graph backwards following only the optimizing equations and starting from 🏁.

This is an upper bound for $\max_{v\in V} \sup_{\pi_1\in\Pi_1^{MD}} \inf_{\pi_2\in\Pi_2^{MD\mathcal{F}}} \mathbb{E}_v^{\pi_1,\pi_2}(rew)$

Algorithmic solution

1. Calculate $\mathbf{U} = \max_{v\in V} \sup_{\pi_1\in\Pi_1^{MD}} \mathbb{E}_v^{\pi_1}(rew)$ on the MDP obtained by fixing π_2^u.

2. Starting on $x_v = \mathbf{U}$, approximate the maximum fixed point on the equations

$$x_v = \begin{cases} \min\left(r(v) + \sum_{v'\in post(v)} \delta(v,v')x_{v'},\ \mathbf{U}\right) & \text{if } v \in V_\mathsf{P} \setminus \{🏁\} \\ \min\left(\max\{r(v) + x_{v'} \mid v' \in post(v)\},\ \mathbf{U}\right) & \text{if } v \in V_1 \setminus \{🏁\} \\ \min\left(\min\{r(v) + x_{v'} \mid v' \in post(v)\},\ \mathbf{U}\right) & \text{if } v \in V_2 \setminus \{🏁\} \\ 0 & \text{if } v = 🏁 \end{cases}$$

3. Derive the optimizing strategies by traversing the graph backwards following only the optimizing equations and starting from 🏁.

CONICET

UNC

To conclude:

- Solving expected total rewards on stochastic games with fair minimizer …
  - … is determined
  - … has a solution on memoryless deterministic (fair) schedulers
  - … can be approximated using Bellman equation

provided the game is (almost surely) stopping under fairness

To conclude:

- Solving expected total rewards on stochastic games with fair minimizer …
  - … is determined
  - … has a solution on memoryless deterministic (fair) schedulers
  - … can be approximated using Bellman equation

  provided the game is (almost surely) stopping under fairness

It can be checked in polynomial time

CONICET

UNC

To conclude:

- Solving expected total rewards on stochastic games with fair minimizer …
  - … is determined
  - … has a solution on memoryless deterministic (fair) schedulers
  - … can be approximated using Bellman equation

  provided the game is (almost surely) stopping under fairness

It can be checked in polynomial time

Prototype implemented in PRISM

To conclude:

- Solving expected total rewards on stochastic games with fair minimizer …
  - … is determined
  - … has a solution on memoryless deterministic (fair) schedulers
  - … can be approximated using Bellman equation

  provided the game is (almost surely) stopping under fairness

- An inconvenience: many interesting problems may not be stopping under fairness

It can be checked in polynomial time

Prototype implemented in PRISM



CONICET

UNC

Roborta vs. the Fair Light!

Pedro R. D'Argenio

joint work with
Pablo Castro, Ramiro Demasi, and Luciano Putruele