Entropy Games and Matrix Multiplication Games

Eugene Asarin Julien Cervelle Aldric Degorre
Cătălin Dima Florian Horn Victor Kozyakin

IRIF, LACL, IITP

EQINOCS seminar
2016-05-11

A game of freedom

The story

Despot and Tribune rule a country, inhabited by People. D aims to minimize People's freedom, T aims to maximize it.

- Turn-based game.
- Despot issues a decree (which respects laws!), permitting/restricting activities and changing system state.
- People are then given some choice of activities (like go to circus, enrol).
- After that Tribune has control, issues (counter-)decrees and changes system state.
- Again, People are given (maybe different) choices of activities.

Despot wants people to have as few choices as possible (in the long term), Tribune wants the opposite.

Outline

Reminder 1: entropy of languages

Entropy of a language $L \subset \Sigma^\omega$ (Chomsky-Miller, Staiger)

- Count the prefixes of length n: find $|\mathtt{pref}_n(L)|$
- Growth rate - entropy $\mathcal{H}(L) = \limsup \frac{\log |\mathtt{pref}_n(L)|}{n}$

Explaining the definition

- Size measure: $|\mathtt{pref}_n(L)| \approx 2^{n\mathcal{H}}$.
- Information bandwidth of a typical $w \in L$ (bits/symbol)
- Related to topological entropy of a subshift, Kolmogorov complexity, fractal dimensions etc.

Reminder 1: entropy of ω-regular languages — example

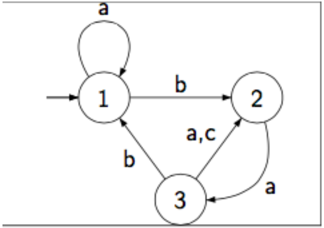

$$M = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 2 & 0 \end{pmatrix}$$

- Prefixes: $\{\varepsilon\}; \{a, b\}; \{aa, ab, ba\}; \{aaa, aab, aba, baa, bab, bac\};$
  $\{aaaa, aaab, aaba, abaa, abab, abac, baaa, bab, baca, baba, babb\} \ldots$
- Cardinalities: 1,2,3,6,11, ...
- $|\mathtt{pref}_n(L)| \approx (1.80194)^n = \rho(M)^n = 2^{0.84955n}$.
  entropy: $\mathcal{H} = \log \rho(M) \approx 0.84955$.

Reminder 1: entropy of ω-regular languages — algorithmics

Recipe: Computing entropy of an ω-regular language L

- Build a deterministic trim automaton for $\mathtt{pref}(L)$.
- Write down its adjacency matrix M.
- Compute $\rho = \rho(M)$ - its spectral radius.
- Then $\mathcal{H} = \log \rho$.

Reminder 1: entropy of ω-regular languages — algorithmics

Recipe: Computing entropy of an ω-regular language L

- Build a deterministic trim automaton for $\mathtt{pref}(L)$.
- Write down its adjacency matrix M.
- Compute $\rho = \rho(M)$ - its spectral radius.
- Then $\mathcal{H} = \log \rho$.

Proof

- $|L_n(i \to j)| = M^n_{ij}$
- Hence $|\mathtt{pref}_n(L)| =$ sum of some elements of M^n
- Perron-Frobenius theory of nonnegative matrices $\Rightarrow \|\mathtt{pref}_n(L)\| \approx \rho(M)^n \Rightarrow \mathcal{H}(L) = \log \rho(M)$

Reminder 2: Generalizations of spectral radii

Spectral radius of a matrix

- $\rho(A)$ is the maximal modulus of eigenvalues of A.
- Gelfand formula $||A^n|| \approx \rho(A)^n$, more precisely $\rho(A) = \lim ||A^n||^{1/n}$

Definition (extending to sets of matrices)

Given a set of matrices $\mathcal{A}$ define

- joint spectral radius $\hat{\rho}(\mathcal{A}) = \lim_{n\to\infty} \sup \left\{ \|A_n \cdots A_1\|^{1/n} \middle| A_i \in \mathcal{A} \right\}$
- joint spectral subradius $\check{\rho}(\mathcal{A}) = \lim_{n\to\infty} \inf \left\{ \|A_n \cdots A_1\|^{1/n} \middle| A_i \in \mathcal{A} \right\}$

Algorithmic difficulties

1. The problem of deciding whether $\hat{\rho}(\mathcal{A}) \leq 1$ is undecidable.
2. The problem of deciding whether $\check{\rho}(\mathcal{A}) = 0$ is undecidable.

Reminder 3: games

Definition (Games)

- Given: two players, two sets of strategies S and T.
- Payoff of a play: when players choose strategies σ and τ, Sam pays to Tom $P(\sigma, \tau)$$
- Guaranteed payoff for Sam: at most $V_+ = \min_\sigma \max_\tau P(\sigma, \tau)$.
- Guaranteed payoff for Tom: at least $V_- = \max_\tau \min_\sigma P(\sigma, \tau)$.
- Game is determined if $V_+ = V_-$

Reminder 3: games

Definition (Games)

- Given: two players, two sets of strategies S and T.
- Payoff of a play: when players choose strategies σ and τ, Sam pays to Tom $P(\sigma, \tau)$$
- Guaranteed payoff for Sam: at most $V_+ = \min_\sigma \max_\tau P(\sigma, \tau)$.
- Guaranteed payoff for Tom: at least $V_- = \max_\tau \min_\sigma P(\sigma, \tau)$.
- Game is determined if $V_+ = V_-$
- Equivalently: exist value V, optimal strategies σ_0 and τ_0 s.t.:
  - Sam chooses $\sigma_0 \Rightarrow$ payoff $\leq V$ for any τ;
  - Tom chooses $\tau_0 \Rightarrow$ payoff $\geq V$ for any σ;

Reminder 3: games 2

Example (Rock-paper-scissors)

- Three strategies for each player: $\{r, p, s\}$
- Payoff matrix:

$\sigma \backslash \tau$	r	p	s
r	0	1	-1
p	-1	0	1
s	1	-1	0

- Non-determined: $\min\max = 1$ and $\max\min = -1$

Questions on a class of games

- are they determined ($V_+ = V_-$)? (e.g. Minimax Theorem, von Neumann)
- describe optimal strategies
- how to compute the value and optimal strategies?

Reminder 3: games on graphs/automata - 1

The setting. Picture - The MIT License (MIT)(c) 2014 Vincenzo Prignano

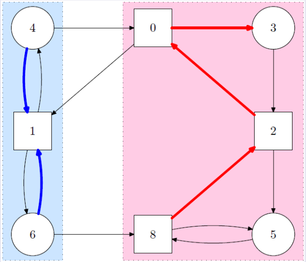

- Arena: graph with vertices $S \cup T$ (belonging to Sam and Tom), edges Δ.
- Sam's strategy σ : history $\mapsto$ outgoing transition, i.e. $\sigma : (S \cup T)^* S \to \Delta$. Tom's strategy τ - symmetrical.
- A play: path in the graph, where in each state the vertex owner decides a transition.
- A payoff function (0-1 or $\mathbb{R}$)

Reminder 3: games on graphs/automata - 2

Simple strategies

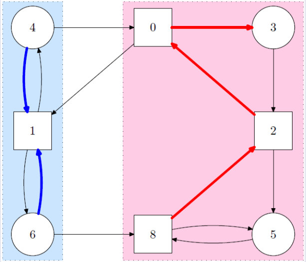

A strategy is called positional (memoryless) if it depends only on the current state: $\sigma : S \to \Delta; \tau : T \to \Delta$.

Reminder 3: games on graphs/automata - 2

Typical results

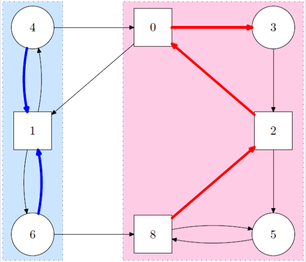

- The game of chess is determined.
- A finite-state game with parity objective is determined, and has positional optimal strategies.
- A finite-state mean-payoff game is determined, and has positional optimal strategies.

A game of freedom – 1st slide again

The story — towards a formalization

Despot and Tribune rule a country, inhabited by People. D aims to minimize People's freedom (entropy), T aims to maximize it.

- Turn-based game.
- Despot issues a decree, changing system state.
- People are then given some choice of activities
- After that Tribune has control, issues (counter-)decrees and changes system state.
- Again, People are given (maybe different) choices of activities.

Despot wants people to have as few choices as possible (minimize the entropy), Tribune wants the opposite.

A game of freedom = an entropy game

Formalization

$A = (D, T, \Sigma, \Delta)$	an arena
with	
$D = \{d_1, d_2, d_3\}$	Despot's states
$T = \{t_1, t_2, t_3\}$	Tribune's states
$\Sigma = \{a, b\}$	action alphabet
$\Delta = \{d_1 a t_1, d_1 a t_2, .\}$	transition relation

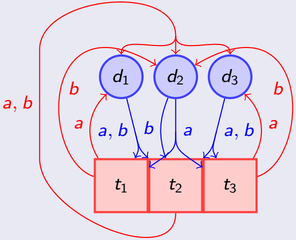

A game of freedom = an entropy game

Formalization

$A = (D, T, \Sigma, \Delta)$ with	an arena
$D = \{d_1, d_2, d_3\}$	Despot's states
$T = \{t_1, t_2, t_3\}$	Tribune's states
$\Sigma = \{a, b\}$	action alphabet
$\Delta = \{d_1 a t_1, d_1 a t_2, .\}$	transition relation
$\sigma : (DT)^* D \to \Sigma$	Despot strategy

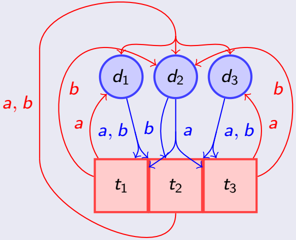

A game of freedom = an entropy game

Formalization

$A = (D, T, \Sigma, \Delta)$	an arena
with	
$D = \{d_1, d_2, d_3\}$	Despot's states
$T = \{t_1, t_2, t_3\}$	Tribune's states
$\Sigma = \{a, b\}$	action alphabet
$\Delta = \{d_1 a t_1, d_1 a t_2, .\}$	transition relation
$\sigma : (DT)^* D \to \Sigma$	Despot strategy
$\tau : (DT)^* \to \Sigma$	Tribune strategy

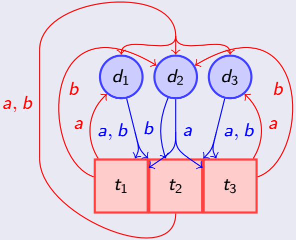

A game of freedom = an entropy game

Formalization

$A = (D, T, \Sigma, \Delta)$ with	an arena
$D = \{d_1, d_2, d_3\}$	Despot's states
$T = \{t_1, t_2, t_3\}$	Tribune's states
$\Sigma = \{a, b\}$	action alphabet
$\Delta = \{d_1 a t_1, d_1 a t_2, .\}$	transition relation
$\sigma : (DT)^* D \to \Sigma$	Despot strategy
$\tau : (DT)^* \to \Sigma$	Tribune strategy
$Runs^\omega(\sigma, \tau)$	available choices for People

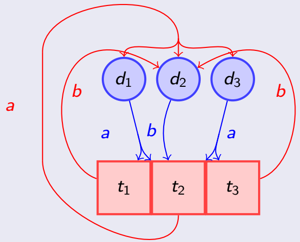

A game of freedom = an entropy game

Formalization

$A = (D, T, \Sigma, \Delta)$ with	an arena
$D = \{d_1, d_2, d_3\}$	Despot's states
$T = \{t_1, t_2, t_3\}$	Tribune's states
$\Sigma = \{a, b\}$	action alphabet
$\Delta = \{d_1 a t_1, d_1 a t_2, .\}$	transition relation
$\sigma : (DT)^* D \to \Sigma$	Despot strategy
$\tau : (DT)^* \to \Sigma$	Tribune strategy
$Runs^{\omega}(\sigma, \tau)$	available choices for People
$\mathcal{H}(Runs^{\omega}(\sigma, \tau))$	Payoff (entropy)

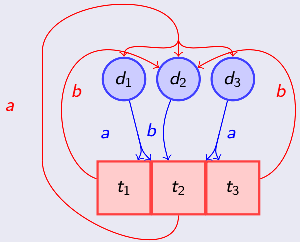

Population Game (same picture, another story)

Another story

Damian and Theo rule a colony of bacteria. Damian (every night) aims to minimize the colony, Theo (every day) to maximize it.

The same picture and tuple

$A = (D, T, \Sigma, \Delta)$	an arena
with	
$D =$	evening forms
$T =$	morning forms
$\Sigma = \{a, b\}$	action alphabet
$\Delta =$	filiation relation
P	$= \limsup \log \lvert \text{colony}_n \rvert / n$

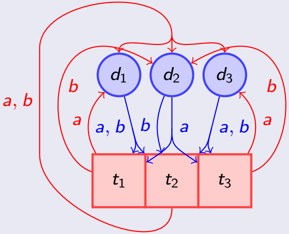

Main tool, 2nd object of study: matrix-multiplication games

The setting

- Adam has a set of matrices $\mathcal{A}$, Eve has $\mathcal{E}$
- They write matrices (of their sets) in turn: $A_1 E_1 A_2 E_2 \dots$
- Adam wants the product to be small (in norm), Eve large.
- Payoff $= \limsup_{n\to\infty} \frac{\log ||A_1 E_1 \dots A_n E_n||}{n}$
- Solve the game: $V_+ = \min_\sigma \max_\tau P =?$ $V_- = \max_\tau \min_\sigma P =?$

Why it cannot be easy

- If Adam is trivial ($\mathcal{A} = \{I\}$) then $V = \hat{\rho}(\mathcal{E})$.
- If Eve is trivial then $V = \check{\rho}(\mathcal{A})$

Matrix-multiplication games are really hard

Theorem

There exists a (family of) MMG with value $V \in \{0, 1\}$ such that it is undecidable whether $V = 0$ or $V = 1$.

Proof idea

Reduce a 2-counter machine halting problem:

- Eve simulates an infinite run ($P = 1$)
- If she cheats Adam resets the game to 0 and $P = 0$

So what?

We will identify a special decidable subclass of MMGs to solve our entropy games.

Solving entropy games

Solution plan

- Represent word counting as matrix multiplication.
- Reduce each EG to a special case of MMGs.
- Prove a minimax property for special MMGs.
- Solve special MMGs and EGs.
- Enjoy!

From a game graph to a set of matrices

Adjacency matrices

Set $\mathcal{A}$ (Adam=Despot)

$$\begin{aligned} \text{1st row} &= [1, 1, 0] \\ \text{2nd row} &\in \{[0, 1, 0], [1, 0, 1]\} \\ \text{3rd row} &= [0, 1, 1] \end{aligned}$$

Set $\mathcal{E}$ (Eve=Tribune)

$$\begin{aligned} \text{1st row} &\in \{[0, 1, 0], [1, 0, 0]\} \\ \text{2nd row} &= [1, 1, 1] \\ \text{3rd row} &\in \{[0, 1, 0], [0, 0, 1]\} \end{aligned}$$

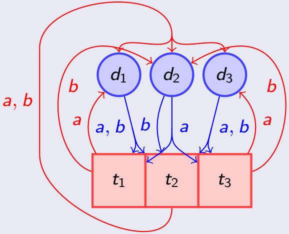

relation - approximated

$$|\texttt{pref}_n(L)| = ||A_1 E_1 A_2 \dots A_n E_n||$$

Exact relation between the two games

Lemma

For every couple of strategies (κ, τ) of Despot and Tribune in the EG there exists a couple of strategies (χ, θ) of Adam and Eve in the MMG $(\mathrm{conv}(\mathcal{A}), \mathrm{conv}(\mathcal{E}))$ with exactly the same payoff.
Moreover, if κ is positional, then χ is constant and permanently chooses A_κ. The case of positional τ is similar.

Exact relation between the two games

Lemma

For every couple of strategies (κ, τ) of Despot and Tribune in the EG there exists a couple of strategies (χ, θ) of Adam and Eve in the MMG $(\mathrm{conv}(\mathcal{A}), \mathrm{conv}(\mathcal{E}))$ with exactly the same payoff.
Moreover, if κ is positional, then χ is constant and permanently chooses A_κ. The case of positional τ is similar.

What does it mean?

The two games are related in some weak and subtle way.

Independent row uncertainty sets [Blondel & Nesterov]

Observation

Adjacency matrix sets $\mathcal{A}$ and $\mathcal{E}$ have the following special structure:

Definition (Sets of matrices with independent row uncertainties = IRU sets)

Given N sets of rows $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_N$, the IRU-set $\mathcal{A}$ consists of all matrices with

- 1st row in $\mathcal{A}_1$,
- 2nd row in $\mathcal{A}_2$,

  $\vdots$
- N-th row in $\mathcal{A}_N$.

Key theorem

Theorem (Minimax Theorem)

For compact IRU-sets of non-negative matrices $\mathcal{A}, \mathcal{B}$ it holds that

$$\min_{A \in \mathcal{A}} \max_{B \in \mathcal{B}} \rho(AB) = \max_{B \in \mathcal{B}} \min_{A \in \mathcal{A}} \rho(AB)$$

IRU matrix games are determined

Theorem (Determinacy Theorem for MMG)

For compact IRU-sets of non-negative matrices the MMG is determined, Adam and Eve possess constant optimal strategies.

Proof.

By minimax theorem $\exists V, E_0, A_0$ such that

$$\max_{E\in\mathcal{E}} \rho(EA_0) = \min_{A\in\mathcal{A}} \max_{E\in\mathcal{E}} \rho(EA) = \max_{E\in\mathcal{E}} \min_{A\in\mathcal{A}} \rho(EA) = \min_{A\in\mathcal{A}} \rho(E_0A) = V.$$

If Adam only plays A_0 $\Rightarrow$ any play $\pi = A_0E_1A_0E_2\cdots$ is a product of matrices from IRU set $\mathcal{E}A_0$,

$$\Rightarrow \text{growth rate} \leq \log \hat{\rho}(\mathcal{E}A_0) \leq \log \max_{E\in\mathcal{E}} \rho(EA_0) = \log V.$$

If Eve only plays E_0 $\Rightarrow$ growth rate $\geq \log V$. □

A naïve algorithm for solving MMGs with finite IRU-sets (exponential)

Find A, E providing minimax $\max_{A\in\mathcal{A}} \min_{E\in\mathcal{E}} \rho(EA) = \min_{E\in\mathcal{E}} \max_{A\in\mathcal{A}} \rho(EA)$.

Entropy games are determined

Theorem

Entropy games are determined. Both players possess optimal positional strategies.

Proof.

Reduction to IRU MMG ☐

Solving running example

Recalling matrices

Set $\mathcal{A}$ (Adam=Despot)

$$\text{1st row} = [1,1,0]$$
$$\text{2nd row} \in \{[0,1,0],[1,0,1]\}$$
$$\text{3rd row} = [0,1,1]$$

Set $\mathcal{E}$ (Eve=Tribune)

$$\text{1st row} \in \{[0,1,0],[1,0,0]\}$$
$$\text{2nd row} = [1,1,1]$$
$$\text{3rd row} \in \{[0,1,0],[0,0,1]\}$$

Optimal strategies

Strategies $\begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}$ for Adam/Despot and $\begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}$ for Eve/Tribune,

Value of both games: $\log \rho(AB) = \log \rho\left(\begin{bmatrix} 2 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 2 \end{bmatrix}\right) = \log\left(\sqrt{17}+3\right)/2 \simeq \log 3.5615$

Complexity of bounding the value of EG

Theorem

Given an entropy game and $\alpha \in \mathbb{Q}_+$, the problem whether the value of the game is $< \alpha$ is in NP $\cap$ coNP.

Bounding the joint spectral radius for IRU-sets is in P

Let $\mathcal{A}$ be a finite IRU-set of non-negative matrices

Lemma

$\hat{\rho}(\mathcal{A}) < \alpha \Leftrightarrow \exists v > 0.\ \forall A \in \mathcal{A}.\ (Av < \alpha v) \quad (*)$

Lemma (inspired by [Blondel & Nesterov])

The problem $\hat{\rho}(\mathcal{A}) < \alpha$? *is in* P.

Proof.

Rewrite (*) as a system of inequations

$$v_i > 0$$

$$c_1 v_1 + c_2 v_2 + \cdots + c_N v_N < \alpha v_i \text{ for each row } [c_1, c_2, \ldots, c_N] \in \mathcal{A}_i$$

Use polynomial algorithm for linear programming. □

Bounding minimax is in NP ∩ co-NP

Lemma

For IRU-sets $\mathcal{A}, \mathcal{B}$ the problem $\min_{A\in\mathcal{A}} \max_{B\in\mathcal{B}} \rho(AB) < \alpha$? *is in* NP ∩ coNP

Proof.

$$\min_{A\in\mathcal{A}} \max_{B\in\mathcal{B}} \rho(AB) < \alpha \Leftrightarrow \exists A \in \mathcal{A}.\ \hat{\rho}(\mathcal{B}A) < \alpha$$

Guess A and use the Lemma saying "$\hat{\rho}(\mathcal{A}) < \alpha$?" is in P;
by duality coNP. □

As a corollary we obtain:

Theorem

Given an entropy game and $\alpha \in \mathbb{Q}_+$, the problem whether the value of the game is $< \alpha$ is in NP ∩ coNP*.*

Perspectives and conclusions

Done

- Two novel games defined: Entropy Game and Matrix Multiplication Game
- Games solved: they are determined, optimal strategies positional, value computable in $\mathsf{NP} \cap \mathsf{coNP}$,
- Related to other games, other problems in linear algebra.

Perspectives and conclusions

Done

- Two novel games defined: Entropy Game and Matrix Multiplication Game
- Games solved: they are determined, optimal strategies positional, value computable in $NP \cap coNP$,
- Related to other games, other problems in linear algebra.

To do

- extend to probabilistic case
- extend to simultaneous moves and/or imperfect information
- find applications!