Stochastic Strategies for Robotic Surveillance

Francesco Bullo

Center for Control,
Dynamical Systems & Computation
University of California at Santa Barbara

http://motion.me.ucsb.edu

Robotics and Control Webinar
Michigan State University, Jan 28, 2021

Acknowledgments

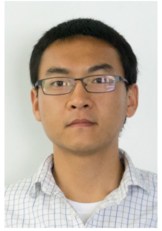

Xiaoming Duan, UT Austin and Shanghai Jiao Tong

Mishel George, DoorDash

Rush Patel, Systems Technology

Pushkarini Agharkar, Google

Outline

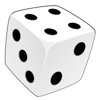

1. **Introduction**
2. Overview of research program
3. Max Return Time Entropy

Related work: monitoring, surveillance, learning:

(1) T. Sak, J. Wainer, and S. Goldenstein. **Probabilistic multiagent patrolling.** *Brazilian Symposium on Artificial Intelligence*, Springer, 2008.

(2) K. Srivatsava, D.M. Stipanovic, and M.W. Spong. **On a stochastic robotic surveillance problem.** *IEEE Conf. on Decision and Control*, 8567-8574, 2009.

(3) G. Cannata and A. Sgorbissa. **A minimalist algorithm for multirobot continuous coverage.** *IEEE Transactions on Robotics*, 27(2):297–312, 2011.

(4) S. Alamdari, E. Fata, and S.L. Smith. **Min-max latency walks: Approximation algorithms for monitoring vertex-weighted graphs.** *Algorithmic Foundations of Robotics X*, 139-155, 2013.

(5) D. Portugal and R. Rocha. **Cooperative multi-robot patrol with Bayesian learning**. Autonomous Robots, 40(5):929–953, 2016

(6) N. Basilico and S. Carpin. **Balancing unpredictability and coverage in adversarial patrolling settings**. In *Algorithmic Foundations of Robotics XIII*, pages 762–777. 2020.

Selected publications

(1) R. Patel, P. Agharkar, and F. Bullo. Robotic surveillance and Markov chains with minimal weighted Kemeny constant. *IEEE Trans. Autom. Control*, 60(12):3156–3167, 2015. doi:10.1109/TAC.2015.2426317

(1) P. Agharkar and F. Bullo. Quickest detection over robotic roadmaps. *IEEE Trans Robotics*, 32(1):252–259, 2016. doi:10.1109/TRO.2015.2506165

(3) R. Patel, A. Carron, and F. Bullo. The hitting time of multiple random walks. *SIAM J Matrix Analysis & Apps*, 37(3):933–954, 2016. doi:10.1137/15M1010737

(4) X. Duan, M. George, R. Patel, and F. Bullo. Robotic surveillance based on the meeting time of random walks. *IEEE Trans Robotics*, 36(4):1356–1362, 2020. doi:10.1109/TRO.2020.2990362

(5) M. George, S. Jafarpour, and F. Bullo. Markov chains with maximum entropy for robotic surveillance. *IEEE Trans. Autom. Control*, 64(4):1566–1580, 2019. doi:10.1109/TAC.2018.2844120

(6) X. Duan, M. George, and F. Bullo. Markov chains with maximum return time entropy for robotic surveillance. *IEEE Trans. Autom. Control*, 65(1):72–86, 2020. doi:10.1109/TAC.2019.2906473

(7) X. Duan and F. Bullo. Markov chain-based stochastic strategies for robotic surveillance. *Annual Review of Control, Robotics, and Autonomous Systems*, 4, 2021. To appear. doi:10.1146/annurev-control-071520-120123

New text "Lectures on Network Systems"

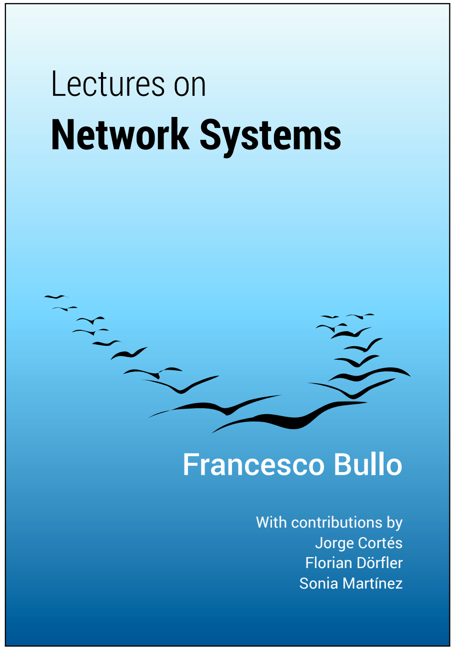

Lectures on Network Systems, Francesco Bullo, KDP, 1.4 edition, 2020, ISBN 978-1-986425-64-3

1. Self-Published and Print-on-Demand at:
   https://www.amazon.com/dp/1986425649

2. PDF Freely available at
   http://motion.me.ucsb.edu/book-lns:
   For students: free PDF for download
   For instructors: slides and solution manual

3. incorporates lessons from 2 decades of research:
   robotic multi-agent, social networks, power grids

version 1.4
332 pages
171 exercises, 220 pages solution manual
6K downloads Jun 2016 - Dec 2020
50 instructors in 17 countries

New text "Lectures on Robotic Planning and Kinematics"

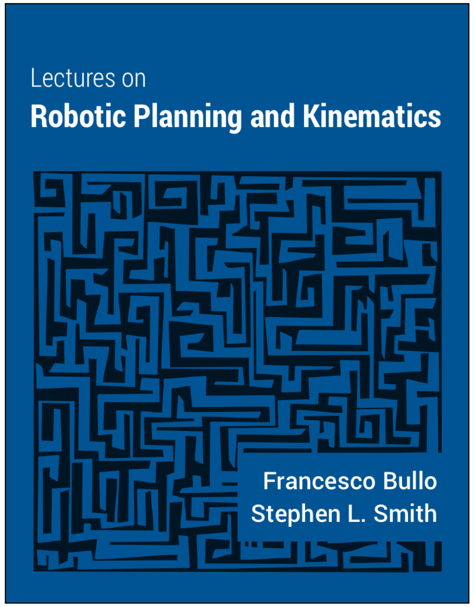

Lectures on Robotic Planning and Kinematics, ver .92
For students: free PDF for download
For instructors: slides and answer keys
http://motion.me.ucsb.edu/book-lrpk/

Robotic Planning:

1. Sensor-based planning
2. Motion planning via decomposition and search
3. Configuration spaces
4. Sampling and collision detetion
5. Motion planning via sampling

Robotic Kinematics:

6. Intro to kinematics
7. Rotation matrices
8. Displacement matrices and inverse kinematics
9. Linear and angular velocities

Outline

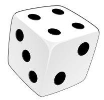

1. Introduction
2. **Overview of research program**
3. Max Return Time Entropy

Stochastic surveillance: Motivating example 1/2

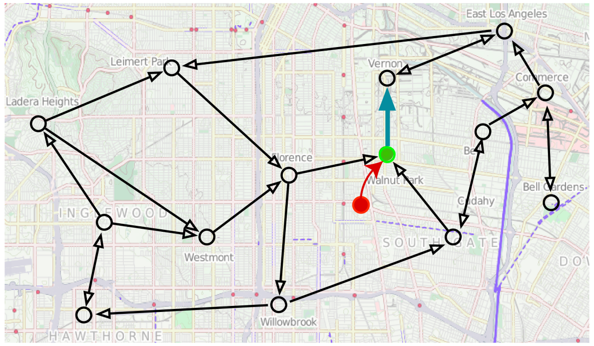

- **Markovian surveillance agents with visit frequency constraints**
- **Intelligent intruders can sense position/observe path of agent**

- Goal: fast unpredictable motion patterns for the surveillance agents

Stochastic surveillance: Motivating example 2/2

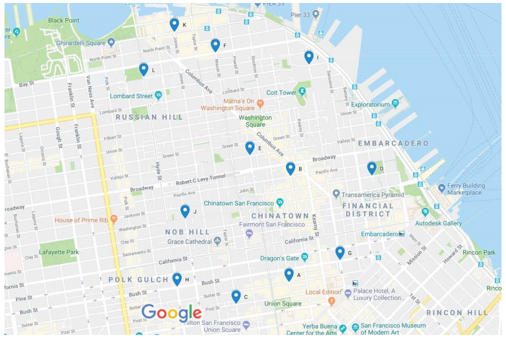

- San Francisco
- crime rate at 12 locations
- all-to-all driving times (quantized in minutes)
- define $\pi \sim$ crime rate

Rational intruder (bank robber model):

- Picks a node i with probability π_i for duration τ
- Learns the inter-visit time statistics of police
- Attacks at time with minimum detection likelihood

Approach: Markov chains for routing and planning

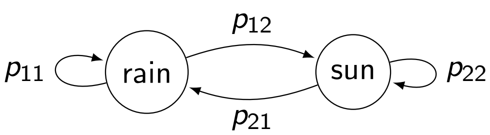

Advantages of adopting Markov chains:

1. quantify and optimize speed, randomness & unpredictability
2. vast body of work on Markov chains (eg, fastest mixing)
3. finite-dimensional opt problem
4. note: TSP may be written as Markov transition matrix

Fundamental objects: first hitting times

First hitting time from location i to location j

$$\text{unweighted: } T_{ij} = \min \Big\{ k \mid X_0 = i, X_k = j, k \geq 1 \Big\}$$

$$\text{weighted: } T^{\mathrm{w}}_{ij} = \min \Big\{ \sum_{s=0}^{k-1} w_{X_s X_{s+1}} \mid X_0 = i, X_k = j, k \geq 1 \Big\}$$

Discrete-time affine system with delays

Let $F_k(i,j) = \mathbb{P}(T_{ij} = k)$ and $F^{\mathrm{w}}_k(i,j) = \mathbb{P}(T^{\mathrm{w}}_{ij} = k)$, for $k \in \mathbb{Z}_{>0}$,

$$F_k(i,j) = p_{ij} \mathbf{1}_{\{k=1\}} + \sum_{h=1, h\neq j}^{n} p_{ih} F_{k-1}(h,j)$$

$$F^{\mathrm{w}}_k(i,j) = p_{ij} \mathbf{1}_{\{k=w_{ij}\}} + \sum_{h=1, h\neq j}^{n} p_{ih} F^{\mathrm{w}}_{k-w_{ih}}(h,j)$$

where $\mathbf{1}_{\{\cdot\}}$ indicator function and $F_k(i,j) = 0$ for all $k \leq 0$ and i, j

Fundamental objects: first hitting times

Mean first hitting times

$$m_{ij} = \mathbb{E}[T_{ij}], \qquad m^{w}_{ij} = \mathbb{E}[T^{w}_{ij}]$$

Linear matrix equation for mean hitting times

By conditioning on the first step

$$m_{ij} = p_{ij} + \sum_{k \neq j} p_{ik}(1 + m_{kj})$$

In matrix form

$$M = \mathbb{1}_n \mathbb{1}_n^\top + P(M - \mathrm{diag}(M)),$$

where $\mathrm{diag}(\cdot)$ takes the diagonal elements and forms a diagonal matrix

Fundamental objects: stationary distribution

Perron-Frobenius theorem

Let $P \in \mathbb{R}^{n \times n}$ be an irreducible row-stochastic matrix, then there exists a $\pi \in \mathbb{R}^n_{>0}$ and $\pi^\top \mathbb{1}_n = 1$ such that

$$\pi^\top P = \pi^\top$$

The stationary distribution encodes the visit frequency information

$$\frac{1}{t+1} \sum_{k=0}^{t} \mathbf{1}_{\{X_k = i\}} \xrightarrow{\text{as } t \to \infty} \pi_i \quad \text{almost surely}$$

Reversible Markov chains

A Markov chain P is reversible if for all $i, j \in \{1, \dots, n\}$

$$\pi_i p_{ij} = \pi_j p_{ji}$$

Approach 1: Fast surveillance: minimizing traversal time

Kemeny's constant: average time to travel between locations

$$\mathcal{K}(P) = \sum_{i=1}^{n} \sum_{j=1}^{n} \pi_i \pi_j m_{ij}$$

$$\mathcal{K}^{w}(P) = \sum_{i=1}^{n} \sum_{j=1}^{n} \pi_i \pi_j m^{w}_{ij} = (\pi^\top (P \circ W) \mathbb{1}_n) \cdot \mathcal{K}(P)$$

Approach 2: Unpredictable surveillance: maximizing randomness

1. entropy rate (classic notion)

$$\mathcal{H}_{\text{rate}}(P) = -\sum_{i=1}^{n} \pi_i \sum_{j=1}^{n} p_{ij} \log p_{ij}$$

2. return time entropy

$$\mathcal{H}_{\text{ret-time}}(P) = \sum_{i=1}^{n} \mathcal{H}(T^{w}_{ii})$$

Fast surveillance: minimizing traversal time

Minimize $\mathcal{K}^{w}$ Problem

Given stationary distribution π and a weighted digraph $\mathcal{G} = \{V, \mathcal{E}, W\}$,

$$\min_{P} \ \mathcal{K}^{w}(P)$$

subject to

1. P is transition matrix with stationary distribution π
2. P is consistent with $\mathcal{G}$

- irreducibility automatically ensured (reducible solution has ∞ value)
- a difficult optimization problem of combinatorial nature
- numerical solutions available without optimality guarantees

Fast surveillance: minimizing traversal time

Minimize $\mathcal{K}^{w}$ Problem

Given stationary distribution π and a weighted digraph $\mathcal{G} = \{V, \mathcal{E}, W\}$,

$$\min_{P} \ \mathcal{K}^{w}(P)$$

subject to

1. P is transition matrix with stationary distribution π
2. P is consistent with $\mathcal{G}$
3. P is reversible

- restrict the search space to a "proper" subspace
- a convex optimization problem with optimality guarantees
- a semidefinite reformation allows for utilizing existing SDP solvers

Fast surveillance: minimizing traversal time

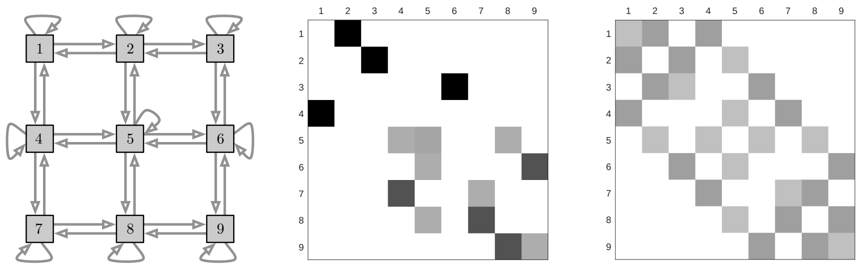

(a) Graph topology (b) Nonreversible (c) Reversible

- the nonreversible solution has a sparser pattern
- performance metric $\mathcal{K}(P)$: 6.78 (nonreversible) $<$ 12.43 (reversible)
- tradeoffs between computational tractability and performance metric

Fast surveillance: minimizing traversal time

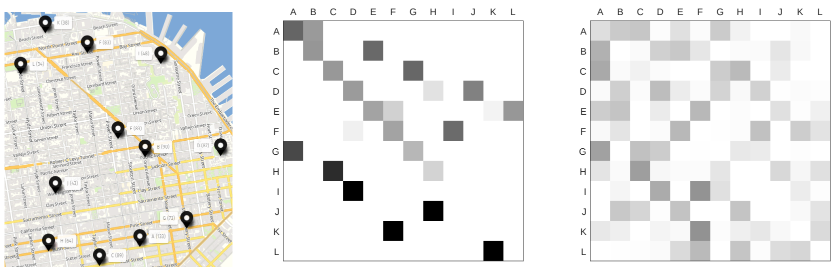

(a) SF map (b) Nonreversible (c) Reversible

- San Francisco map with crime rate data at 12 locations
- a weighted graph with travel times between pairs of locations
- performance metric $\mathcal{K}^{w}(P)$: 22.19 (nonreversible) $<$ 44.77 (reversible)

Extended applications of the mean hitting times 1/2

Meeting times for two moving agents

Meeting times for a pursuer and an evader (two Markov chains)

$$T_{ij} = \min\{k \geq 1 \,|\, X_k^{p} = X_k^{e}, X_0^{p} = i, X_0^{e} = j\},$$

Linear equations for mean meeting times

$$m_{ij} = 1 + \sum_{k_1 \neq h_1} p_{ik_1}^{p} p_{jh_1}^{e} m_{k_1 h_1}.$$

The expected meeting time

$$\mathcal{K}(P^{p}, P^{e}) = \sum_{i=1}^{n} \sum_{j=1}^{n} \pi_i^{p} \pi_j^{e} m_{ij}.$$

Extended applications of the mean hitting times 1/2

Minimize $\mathcal{K}(P^{\mathrm{p}}, P^{\mathrm{e}})$ Problem

Given stationary distribution π^{p}, a digraph $\mathcal{G} = \{V, \mathcal{E}\}$ and P^{e}

$$\min_{P^{\mathrm{p}}} \ \mathcal{K}(P^{\mathrm{p}}, P^{\mathrm{e}})$$

subject to

1. P^{p} is transition matrix with stationary distribution π^{p}
2. P^{p} is consistent with $\mathcal{G}$

- irreducibility is not sufficient to ensure finite-time capture
- Kemeny's constant optimization is a special case (static intruder)

Extended applications of the mean hitting times 1/2

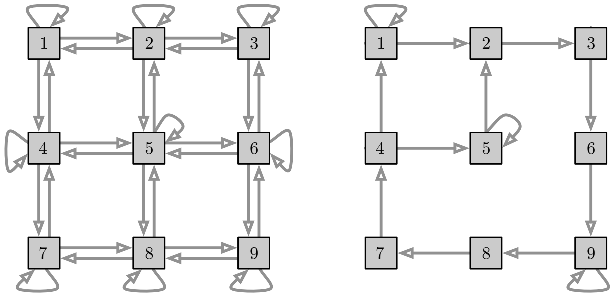

(a) Graph topology (b) Strategy P^{p}

Figure: Optimal strategy against a randomly walking evader

- the evader walks to neighboring locations with equal probabilities
- surveillance strategy is sparse and has similar pattern as MinKemeny

Extended applications of the mean hitting times 2/2

Hitting times for a team of robots

Hitting times for a team of N robots to a location j

$$T_{i_1 \dots i_N, j} = \min\{k \geq 1 \,|\, X_k^1 = j \text{ or } X_k^2 = j \cdots \text{ or } X_k^N = j, \\ X_0^h = i_h \text{ for } h \in \{1, \dots, N\}\}$$

Linear equations for mean hitting times

$$m_{i_1 \dots i_N, j} = 1 + \sum_{k_1 \neq j} \cdots \sum_{k_N \neq j} p^1_{i_1 k_1} \cdots p^N_{i_N k_N} m_{k_1 \dots k_N, j}$$

which can be reorganized in matrix form

- the exponential growth of dimensionality becomes an issue
- reliable and efficient formulation of optimization problems is lacking

Outline

Unpredictable surveillance: maximizing randomness

Approach: Entropy of random variable

Given a discrete random variable $X \in \{1, \dots, k\}$, the Shannon entropy is

$$\mathbb{H}(X) = -\sum_{i=1}^{k} p_i \log p_i.$$

Unbiased coin:	$\mathbb{P}[X = \text{Head}] = 0.5$	$\mathbb{H}(X) = \mathbf{0.693}$
Biased coin:	$\mathbb{P}[X = \text{Head}] = 0.75$	$\mathbb{H}(X) = \mathbf{0.562}$
Predictable coin:	$\mathbb{P}[X = \text{Head}] = 1$	$\mathbb{H}(X) = 0$

Advantages maximizing entropy

1. entropy = well-defined fundamental concept for the randomness
2. if the surveillance agent is highly entropic, it is hard for the intruders to learn the patterns in the behavior of the agent
3. since the behaviors of the intruders may not be exactly known/modeled in any case, optimizing the surveillance strategies against certain intruder behaviors may not be generally wise
4. simulations illustrate that MaxReturnEntropy chain works well for bank robber model

The entropy of what variable?

#1: sequence of random locations

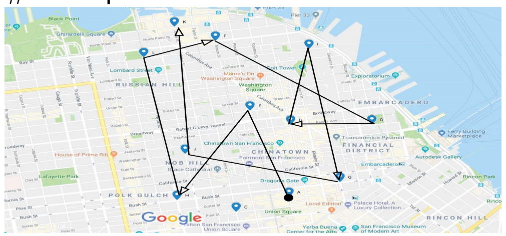

#2: sequence of return times

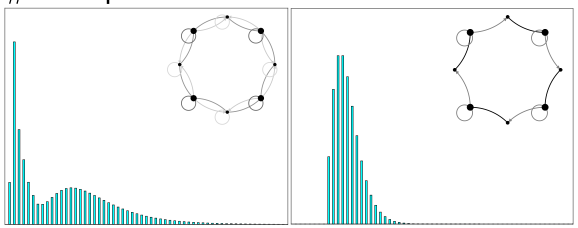

#1: The entropy rate of a Markov chain

A classic notion from information theory

entropy rate of sequence of symbols/locations

$$\mathbb{H}_{\text{location}}(P) = -\sum\nolimits_{i=1}^{n} \pi_i \sum\nolimits_{j=1}^{n} p_{ij} \log p_{ij}$$

Maximizing the location entropy rate

Given stationary distribution π & adjacency matrix A

$$\max_{P} \mathbb{H}_{\text{location}}(P)$$

1. P is transition matrix with stationary distribution π
2. P is consistent with A

#2: Return time entropy of Markov chain

Better entropy notion

For a transition matrix P

$$T_{ii}(P) = \text{first time agent starting at } i \text{ returns back to } i$$

Return time entropy of Markov chain

Given irreducible Markov chain P over weighted digraph $\mathcal{G} = \{V, \mathcal{E}, W\}$ and stationary distribution π, the **return time entropy** is

$$\mathbb{H}_{\text{return-time}}(P) = \sum_{i=1}^{n} \pi_i \mathbb{H}(T_{ii}(P))$$

directed graphs and travel weights

Main problem statement

Maximize $\mathbb{H}_{\text{return-time}}$ Problem

Given stationary distribution π and a weighted digraph $\mathcal{G} = \{V, \mathcal{E}, W\}$,

$$\max_P \ \mathbb{H}_{\text{return-time}}(P)$$

subject to

1. P is transition matrix with stationary distribution π
2. P is consistent with $\mathcal{G}$

Outline

Summary of results

Maximize $\mathbb{H}_{\text{return-time}}$ Problem

Given stationary distribution π and a weighted digraph $\mathcal{G} = \{V, \mathcal{E}, W\}$,

$$\max_P \ \mathbb{H}_{\text{return-time}}(P)$$

subject to

1. P is transition matrix with stationary distribution π.
2. P is consistent with $\mathcal{G}$.

Thm 1: Hitting time probability dynamics
Thm 2: Max $\mathbb{H}_{\text{return-time}}$ is well-posed
Thm 3: Upper bound and solution for complete graph
Thm 4: Relations with the location entropy rate
Thm 5: Truncation, approximation and computation

Basic ideas

$$T_{ij} = \min\left\{ \sum_{s=0}^{k-1} w_{X_s X_{s+1}} \mid X_0 = i, X_k = j, k \geq 1 \right\}$$

$$F_k(i,j) = \mathbb{P}[T_{ij} = k]$$

$$\mathbb{H}_{\text{return-time}}(T_{ii}) = -\sum_{k=1}^{\infty} F_k(i,i) \log F_k(i,i)$$

Recursive formula, for $k \in \mathbb{Z}_{>0}$,

$$F_k(i,j) = p_{ij} \mathbf{1}_{\{k = w_{ij}\}} + \sum_{h=1, h \neq j}^{n} p_{ih} F_{k - w_{ih}}(h,j) \qquad (1)$$

where $\mathbf{1}_{\{\cdot\}}$ indicator function and
where $F_k(i,j) = 0$ for all $k \leq 0$ and i, j

Little example with summable series

Only other example is complete homogeneous graph

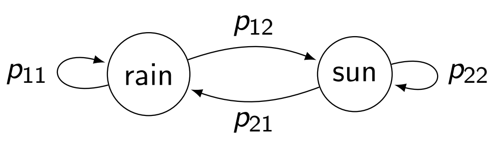

For this special case

$$\mathbb{P}(T_{11} = k) = \begin{cases} p_{11}, & \text{if } k = 1, \\ p_{12} p_{22}^{k-2} p_{21}, & \text{if } k \geq 2. \end{cases}$$

$$\mathbb{H}(T_{11}) = -p_{11} \log p_{11} - p_{12} \log(p_{12} p_{21}) - \frac{p_{12} p_{22} \log p_{22}}{p_{21}}$$

$$\mathbb{H}_{\text{return-time}}(P) = -2\pi_1 p_{11} \log(p_{11}) - 2\pi_2 p_{22} \log(p_{22}) - 2\pi_1 p_{12} \log(p_{12}) - 2\pi_2 p_{21} \log(p_{21}).$$

In general, $\mathbb{H}_{\text{return-time}}(P)$ does not admit a closed form.

Thm 1: Hitting time probability dynamics

Given an irreducible Markov chain $P \in \mathbb{R}^{n \times n}$ on weighted digraph $\mathcal{G}$,

1. hitting time probabilities satisfy

$$\operatorname{vec}(F_k) = \sum_{i,j=1}^{n} p_{ij} ([\mathbb{1}_n - \mathbb{e}_i] \otimes \mathbb{e}_i \mathbb{e}_j^\top) \operatorname{vec}(F_{k - w_{ij}}) + \operatorname{vec}(P \circ \mathbf{1}_{\{k \mathbb{1}_n \mathbb{1}_n^\top = W\}})$$

2. discrete-time affine system with delays – is exponentially stable

Thm 2: Max $\mathbb{H}_{\text{return-time}}$ is well-posed

$\mathbb{H}_{\text{return-time}}$ is a continuous function over a compact set

The uniform limit of any sequence of continuous functions is continuous.

Consider a sequence of functions $\{f_k : \mathcal{X} \to \mathbb{R}\}_{k \in \mathbb{Z}_{>0}}$. If there exists a sequence of Weierstrass scalars $\{M_k\}_{k \in \mathbb{Z}_{>0}}$ such that

$$\sum_{k=1}^{\infty} M_k < \infty \quad \text{and} \quad |f_k(x)| \leq M_k, \quad \text{for all } x \in \mathcal{X}, k \in \mathbb{Z}_{>0},$$

then $\sum_{k=1}^{\infty} f_k$ converges uniformly. Today $f_k = F_k(i,i) \log F_k(i,i)$.

Given compact set of Schur $\mathcal{A} \subset \mathbb{R}^{n \times n}$, let $\rho_{\mathcal{A}} := \max_{A \in \mathcal{A}} \rho(A) < 1$. For any $\lambda \in (\rho_{\mathcal{A}}, 1)$ and for any $\|\cdot\|$, there exists $c > 0$ s.t.

$$\|A^k\| \leq c\lambda^k, \quad \text{for all } A \in \mathcal{A} \text{ and } k \in \mathbb{Z}_{\geq 0}.$$

at fixed mean μ, maxentropic distribution over $\mathbb{N}$ is geometric $\mathbb{P}[k] = (1 - 1/\mu)^{k-1}/\mu$

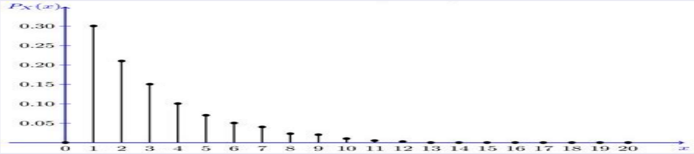

Thm 3: Upper bound and solution for complete graph

1. the return time entropy function is upper bounded by

$$\mathbb{H}_{\text{return-time}}(P) \leq -\sum_{i=1}^{n} (\pi_i \log \pi_i + (1 - \pi_i) \log(1 - \pi_i))$$

2. if $\mathcal{G}$ is complete, the upper bound is achieved with $P = \mathbb{1}_n \pi^\top$

Thm 4: Relations with the location entropy rate

Given an irreducible Markov chain $P \in \mathbb{R}^{n\times n}$ over an unweighted digraph $\mathcal{G}$ and stationary distribution π, $\mathbb{H}_{\text{return-time}}(P)$ and $\mathbb{H}_{\text{location}}(P)$ satisfy

$$\mathbb{H}_{\text{location}}(P) \leq \mathbb{H}_{\text{return-time}}(P) \leq n\,\mathbb{H}_{\text{location}}(P).$$

- lower bound: due to concavity of $-x \log x$
- lower bound: achieved with P is a permutation matrix, $0 = 0$
- upper bound: proof by analyzing the entropy of trajectories
- upper bound: achieved when different return paths = different lengths

Lesson: $\mathbb{H}_{\text{return-time}}(P)$ can be very different from $\mathbb{H}_{\text{location}}(P)$

Computational ideas

Given accuracy η, truncation duration N_η and tail probability satisfy

$$N_\eta = \left\lceil \frac{w_{\max}}{\eta \pi_{\min}} \right\rceil - 1 \quad \Rightarrow \quad \mathbb{P}[T_{ii} \geq N_\eta + 1] \leq \eta$$

The conditional return time entropy is of interest:

$$(\mathbb{H}_{\text{return-time}})_{\text{cond},\eta}(P) = \sum_{i=1}^{n} \pi_i \mathbb{H}(T_{ii} \mid T_{ii} \leq N_\eta)$$
$$= -\sum_{i=1}^{n} \pi_i \sum_{k=1}^{N_\eta} \frac{F_k(i,i)}{\sum_{k=1}^{N_\eta} F_k(i,i)} \log \frac{F_k(i,i)}{\sum_{k=1}^{N_\eta} F_k(i,i)}$$

In practice, the truncated return time entropy is

$$(\mathbb{H}_{\text{return-time}})_{\text{trunc},\eta}(P) = -\sum_{i=1}^{n} \pi_i \sum_{k=1}^{N_\eta} F_k(i,i) \log F_k(i,i)$$

Thm 5: Truncation, approximation and computation

Given a strongly connected weighted digraph $\mathcal{G}$, stationary distribution π,

1. Asymptotic agreement

$$\mathbb{H}_{\text{return-time}}(P) = \lim_{\eta\to 0^+} (\mathbb{H}_{\text{return-time}})_{\text{cond},\eta}(P) = \lim_{\eta\to 0^+} (\mathbb{H}_{\text{return-time}})_{\text{trunc},\eta}(P)$$

2. The gradient of $(\mathbb{H}_{\text{return-time}})_{\text{trunc},\eta}(P)$ can be computed via

$$\operatorname{vec}\left(\frac{\partial (\mathbb{H}_{\text{return-time}})_{\text{trunc},\eta}(P)}{\partial P}\right) = -\sum_{i=1}^{n} \pi_i \sum_{k=1}^{N_\eta} \frac{\partial (F_k(i,i) \log F_k(i,i))}{\partial F_k(i,i)} G_k^\top \mathbb{e}_{(i-1)n+i}$$

where $G_k = \left[\frac{\partial \operatorname{vec}(F_k)}{\partial p_{11}} \quad \cdots \quad \frac{\partial \operatorname{vec}(F_k)}{\partial p_{nn}}\right]$ satisfies a delayed linear system

Proof: exp stability of affine delayed system + uniform bound + chain rule

Gradient projection algorithm

select: minimum edge weight $\epsilon \ll 1$,
select: truncation accuracy $\eta \ll 1$, and
select: initial condition P_0 in $\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}$

for iteration parameter $s = 0 :$ (number-of-steps) **do**
  $\{G_k\}_{k \in \{1,\dots,N_\eta\}} :=$ solution to Thm 4 at P_s
  $\Delta_s :=$ gradient of $(\mathbb{H}_{\text{return-time}})_{\text{trunc},\eta}(P_s)$
  $P_{s+1} :=$ projection$_{\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}}(P_s + (\text{step size}) \cdot \Delta_s)$
end for

Outline

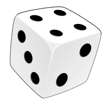

1. Problem setup and motivation
2. Markov chains with maximum return time entropy
3. **Performance of proposed solution**
4. Conclusion and future directions

Compare three chains

1. MaxReturnEntropy

$$\max_P \mathbb{H}_{\text{return-time}}(P)$$

2. MaxLocationEntropy

$$\max_P \mathbb{H}_{\text{location}}(P)$$

entropy rate of sequence of symbols/locations

$$\mathbb{H}_{\text{location}}(P) = -\sum_{i=1}^{n} \pi_i \sum_{j=1}^{n} p_{ij} \log p_{ij}$$

3. MinCaptureTime: $\min_P \mathbb{E}[K(P)]$
Minimize the mean capture time:

$$k_i = \sum_i \mathbb{E}[T_{ij}]\pi_j \qquad = k_j$$

Comparison over a ring and a grid graph 1/2

Unit travel times.
Ring weights = 4 high, 4 low. Grid weights $\sim$ node degree.

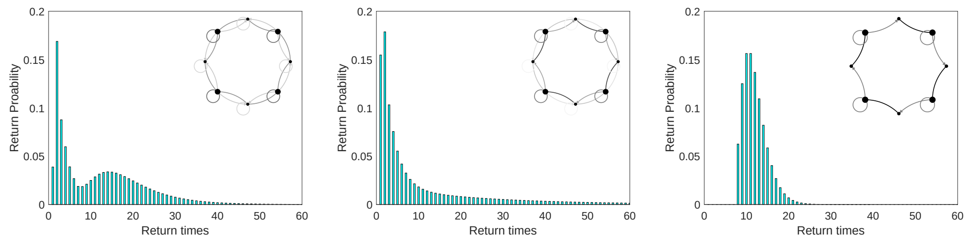

(a) MaxReturnEntropy (b) MaxLocationEntropy (c) MinCaptureTime

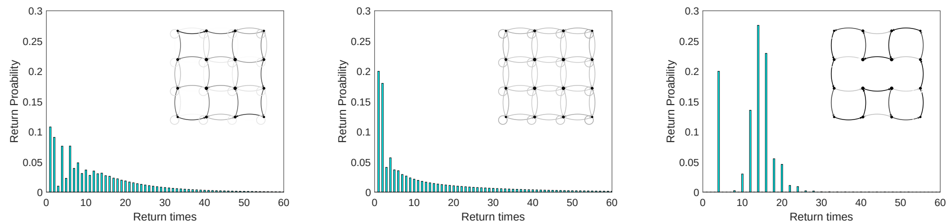

(d) MaxReturnEntropy (e) MaxLocationEntropy (f) MinCaptureTime

Comparison over a ring and a grid graph 2/2

Graph	Markov chains	$\mathbb{H}_{\text{return-time}}(P)$	$\mathbb{H}_{\text{location}}(P)$	Capture Time
8-node ring	MaxReturnEntropy	**2.49**	0.86	**10.04**
	MaxLocationEntropy	2.35	**0.98**	19.53
	MinCaptureTime	1.96	0.46	**6.16**
4-by-4 grid	MaxReturnEntropy	**3.65**	0.94	**16.35**
	MaxLocationEntropy	3.28	**1.40**	30.86
	MinCaptureTime	2.09	0.21	**10.09**

MaxReturnEntropy chain combines speed and unpredictability.
MaxReturnEntropy is **nonreversible** and thus faster in general.

Comparison over San Francisco map 1/3

Stochastic surveillance: Motivating example 2/2

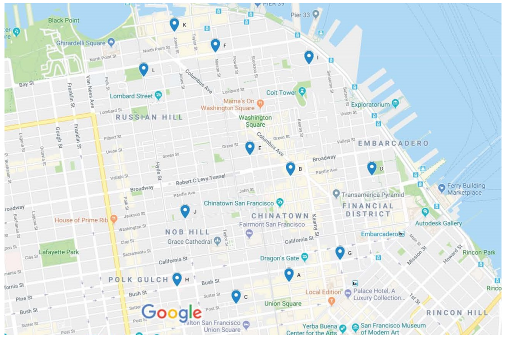

- San Francisco
- crime rate at 12 locations
- complete by-car travel times (quantized in minutes)
- $\pi \sim$ crime rate

Rational intruder (bank robber model):

- Picks a node i with probability π_i for duration τ
- Learns the inter-visit time statistics of police
- Attacks at time with minimum detection likelihood

Comparison over San Francisco map 2/3

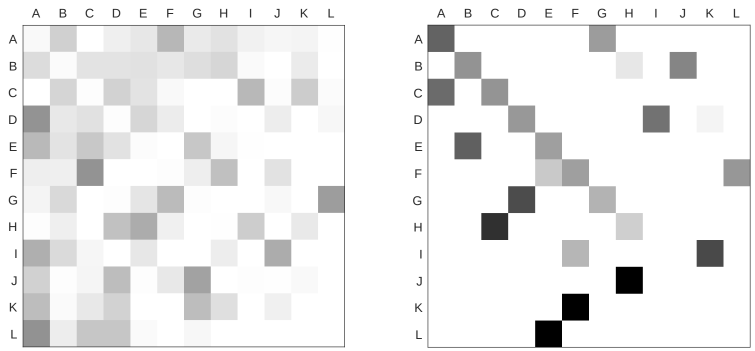

(g) MaxReturnEntropy (h) MinCaptureTime

Figure: Pixel image of the Markov chains with row sum being 1

- MinCaptureTime chain is close to a shortest tour with self weights
- MaxReturnEntropy chain is dense and creates more return entropy

Comparison over San Francisco map 3/3: high vs. low

MaxReturnEntropy versus MinCaptureTime: high importance node

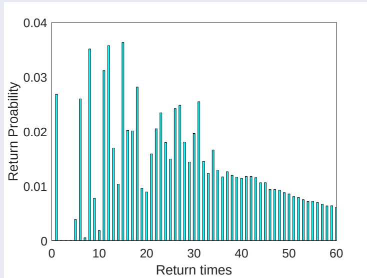

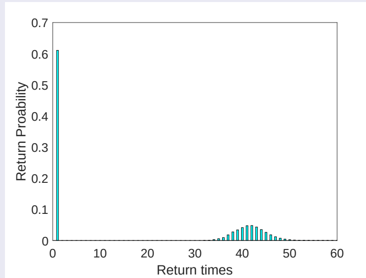

MaxReturnEntropy versus MinCaptureTime: low importance node

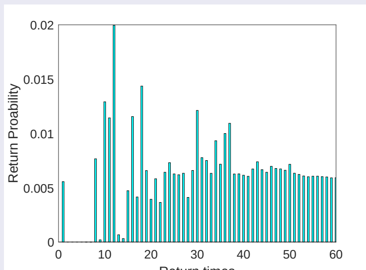

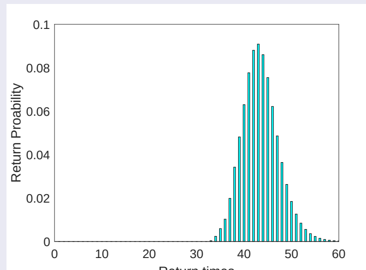

Comparison in catching the rational intruder 1/2

Rational intruder:

- Picks a node i to attack with probability π_i
- Collects the inter-visit (return) time statistics of the agent
- Attacks when the agent is absent for s_i timesteps since last visit

$$s_i = \underset{0 \le s \le S_i}{\operatorname{argmin}} \left\{ \sum_{k=1}^{\tau} \mathbb{P}(T_{ii} = s + k \mid T_{ii} > s) \right\},$$

where τ is the attack duration and S_i is determined by the degree of impatience δ, i.e., $\mathbb{P}(T_{ii} \ge S_i) \le \delta$

Comparison in catching the rational intruder 2/2

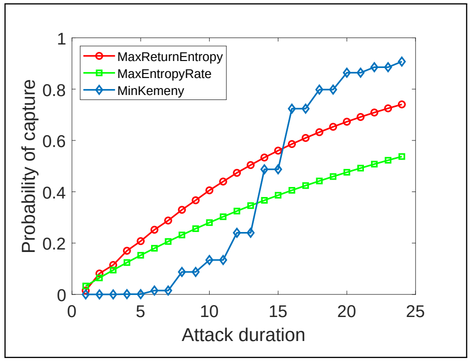

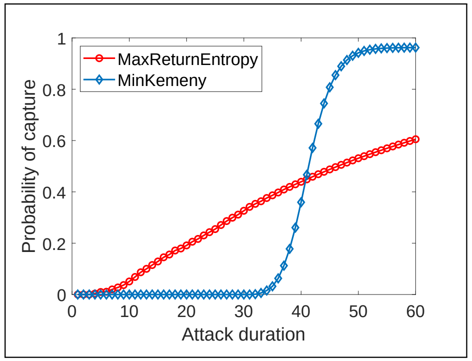

BOTTOM LINE:

- 4×4 grid: MaxReturnEntropy $>$ MaxLocationEntropy
- 4×4 grid: MaxReturnEntropy $>$ MinCaptureTime for short attacks
- SF w-dig: MaxReturnEntropy $>$ MinCaptureTime for short attacks

Conclusion and future directions

Conclusion

1. new metric for unpredictability in stochastic surveillance
2. analysis and computation for maximum return time entropy chain
3. applicability (and comparison) in stochastic surveillance

Ongoing and Future Work

1. Trade-of between unpredictability and speed
2. Stackelberg games
3. Multi-vehicle resource allocation
4. Discretization strategies
5. ...