On Robotic Routing and Stochastic Surveillance

Francesco Bullo

Center for Control,
Dynamical Systems & Computation
University of California at Santa Barbara
http://motion.me.ucsb.edu

Shenyang Institute of Automation, Chinese Academy of Sciences, Shenyang, China, Jun '18
Academy of Mathematics and System Science, Chinese Academy of Science, Beijing, China, Jun '18
School of Automation, Beijing Institute of Technology, Beijing, China, Jun '18

Acknowledgments

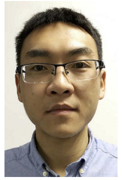

Xiaoming Duan
UCSB

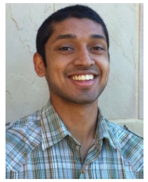

Mishel George
Scotty Labs

Rush Patel
Northrop Grumman

Pushkarini Agharkar
Google

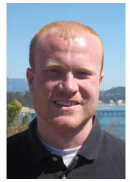

Jeff Peters
UTRC

NSF AFOSR ARO ONR DOE

New text "Lectures on Network Systems"

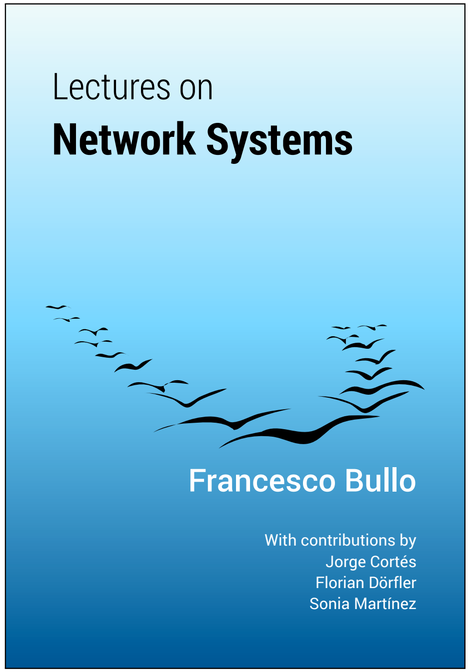

Lectures on Network Systems, Francesco Bullo, Createspace, 1 edition, ISBN 978-1-986425-64-3

For students: free PDF for download
For instructors: slides and answer keys
https://www.amazon.com/dp/1986425649
300 pages (plus 200 pages solution manual)
3K downloads since Jun 2016
150 exercises with solutions
20 instructors have adopted parts of it

Linear Systems:

1. social, sensor, robotic & compartmental examples,
2. matrix and graph theory, with an emphasis on Perron–Frobenius theory and algebraic graph theory,
3. averaging algorithms in discrete and continuous time, described by static and time-varying matrices, and
4. positive & compartmental systems, dynamical flow systems, Metzler matrices.

Nonlinear Systems:

5. nonlinear consensus models,
6. population dynamic models in multi-species systems,
7. coupled oscillators, with an emphasis on the Kuramoto model and models of power networks

New text "Lectures on Robotic Planning and Kinematics"

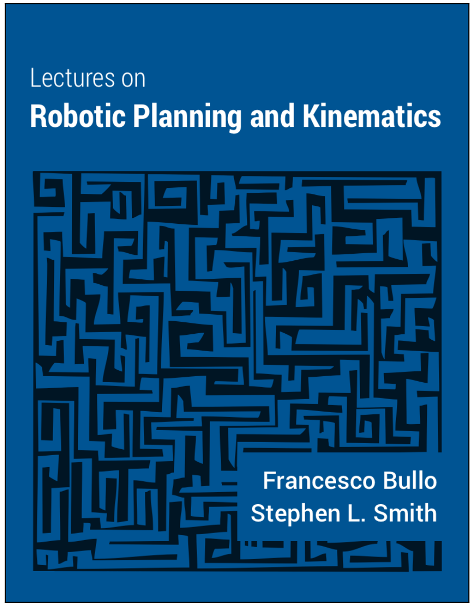

Lectures on Robotic Planning and Kinematics, ver .91
For students: free PDF for download
For instructors: slides and answer keys
http://motion.me.ucsb.edu/book-lrpk/

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

Stochastic surveillance and dynamic routing

Design efficient vehicle control strategies to

1. search unpredictably
2. detect anomalies quickly
3. provide service to customers at known locations
4. perform load balancing among vehicles

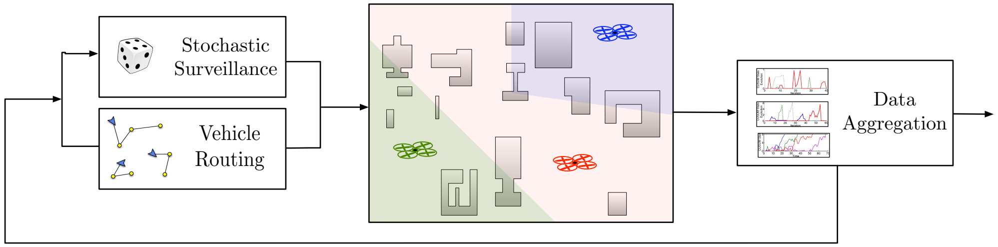

Outline

1. **vehicle routing**
2. load balancing and partitioning
3. stochastic surveillance

AeroVironment Inc, "Raven"
unmanned aerial vehicle

iRobot Inc, "PackBot"
unmanned ground vehicle

Vehicle routing in dynamic stochastic environments

- customers appear sequentially randomly space/time
- robotic network *knows* locations and provides service
- Goal: distributed adaptive algos, delay vs throughput

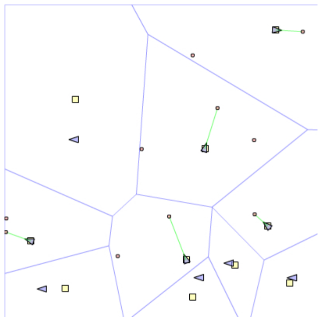

F. Bullo, E. Frazzoli, M. Pavone, K. Savla, and S. L. Smith. Dynamic vehicle routing for robotic systems.
Proceedings of the IEEE, 99(9):1482–1504, 2011.
doi:10.1109/JPROC.2011.2158181

Algo #1: Receding-horizon shortest-path policy

Receding-horizon Shortest-Path (RH-SP)

For $\eta \in (0,1]$, single agent performs:

1: while no customers, move to center

2: while customers waiting

1. compute shortest path through current customers
2. service η-fraction of path

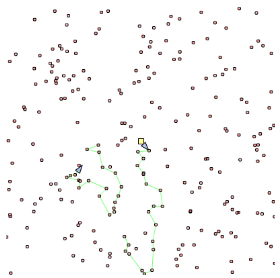

- shortest path is NP-hard, but effective heuristics available
- delay is optimal in light traffic
- delay is constant-factor optimal in high traffic

Algo #2: Load balancing via territory partitioning

RH-SP + Partitioning

For $\eta \in (0,1]$, agent i performs:

1: compute own cell v_i in optimal partition

2: apply RH-SP policy on v_i

Asymptotically constant-factor optimal in light and high traffic

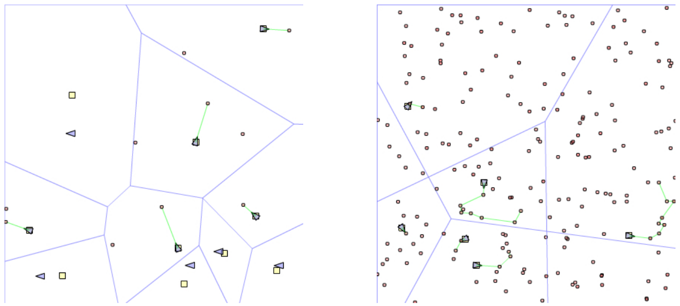

Outline

1. vehicle routing
2. **load balancing and partitioning**
3. stochastic surveillance

AeroVironment Inc, "Raven" unmanned aerial vehicle

iRobot Inc, "PackBot" unmanned ground vehicle

Load balancing via partitioning

ANALYSIS of cooperative distributed behaviors

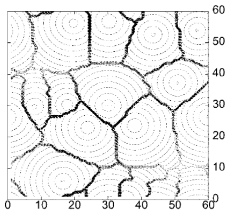

DESIGN of performance metrics

1. how to cover a region with n minimum-radius overlapping disks?
2. how to design a minimum-distortion (fixed-rate) vector quantizer?
3. where to place mailboxes in a city / cache servers on the internet?

Voronoi+centering algorithm

Voronoi+centering law

At each comm round:

1: acquire neighbors' positions

2: compute own dominance region

3: move towards center of own dominance region

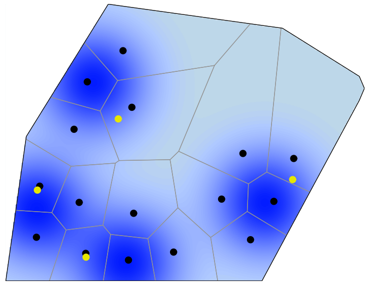

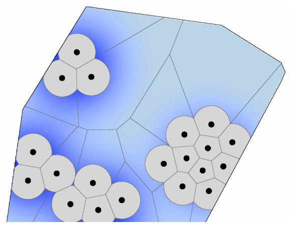

Area-center

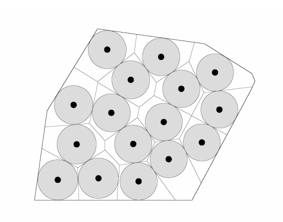

Incenter

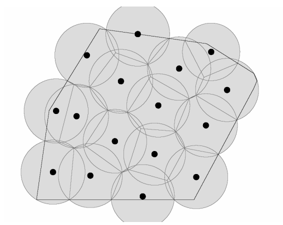

Circumcenter

S. Martínez, J. Cortés, and F. Bullo. Motion coordination with distributed information. *IEEE Control Systems*, 27(4):75–88, 2007.
doi:10.1109/MCS.2007.384124

Variations on the Voronoi+centering theme

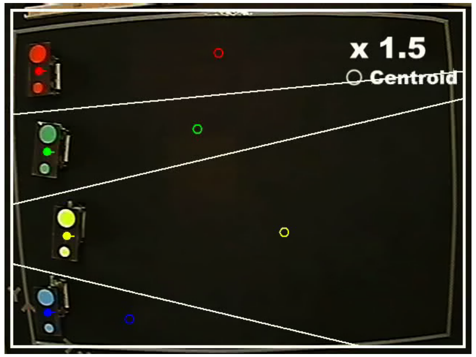

T. Hatanaka, M. Fujita, TokyoTech

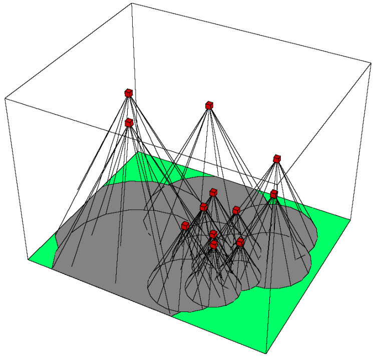

3D coverage

Variations on the Voronoi+centering theme

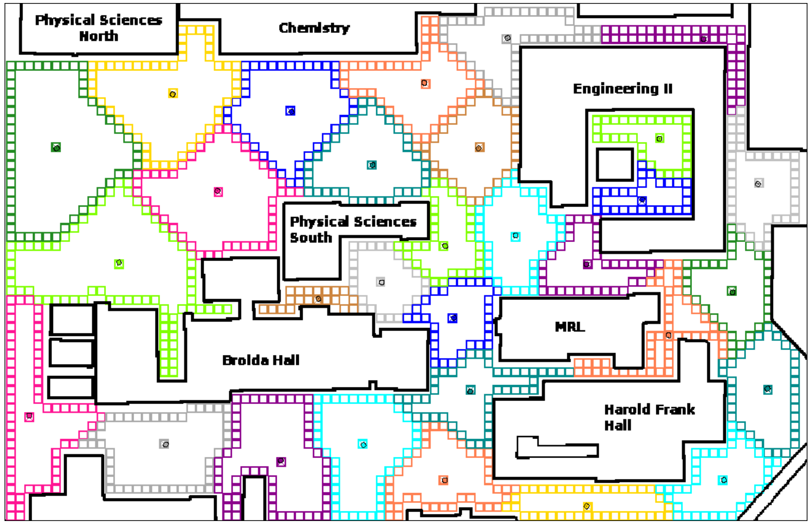

J. W. Durham, R. Carli, P. Frasca, and F. Bullo. Discrete partitioning and coverage control for gossiping robots.
IEEE Transactions on Robotics, 28(2):364–378, 2012.
doi:10.1109/TRO.2011.2170753

Outline

1. vehicle routing
2. load balancing and partitioning
3. **stochastic surveillance**

AeroVironment Inc, "Raven"
unmanned aerial vehicle

iRobot Inc, "PackBot"
unmanned ground vehicle

Outline

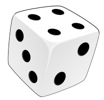

1. **Problem setup and motivation**
2. Markov chains with maximum return time entropy
3. Performance of proposed solution
4. Conclusion and future directions

Related work on persistent monitoring and surveillance:

1. G. Cannata and A. Sgorbissa. A minimalist algorithm for multirobot continuous coverage.
   IEEE Transactions on Robotics, 27(2):297–312, 2011.
   doi:10.1109/TRO.2011.2104510
2. S. Alamdari, E. Fata, and S. L. Smith. Persistent monitoring in discrete environments: Minimizing the maximum weighted latency between observations.
   International Journal of Robotics Research, 33(1):138–154, 2014.
   doi:10.1177/0278364913504011
3. J. Yu, S. Karaman, and D. Rus. Persistent monitoring of events with stochastic arrivals at multiple stations.
   IEEE Transactions on Robotics, 31(3):521–535, 2015.
   doi:10.1109/TRO.2015.2409453

Our publications

1. R. Patel, P. Agharkar, and F. Bullo. Robotic surveillance and Markov chains with minimal weighted Kemeny constant.
   IEEE Transactions on Automatic Control, 60(12):3156–3167, 2015.
   doi:10.1109/TAC.2015.2426317
2. R. Patel, A. Carron, and F. Bullo. The hitting time of multiple random walks.
   SIAM Journal on Matrix Analysis and Applications, 37(3):933–954, 2016.
   doi:10.1137/15M1010737
3. M. George, S. Jafarpour, and F. Bullo. Markov chains with maximum entropy for robotic surveillance.
   IEEE Transactions on Automatic Control, May 2018.
   doi:10.1109/TAC.2018.2844120
4. X. Duan, M. George, and F. Bullo. Markov chains with maximum return time entropy for robotic surveillance.
   IEEE Transactions on Automatic Control, May 2018.
   Submitted.
   URL: https://arxiv.org/abs/1803.07705

Stochastic surveillance: Motivating example 1/2

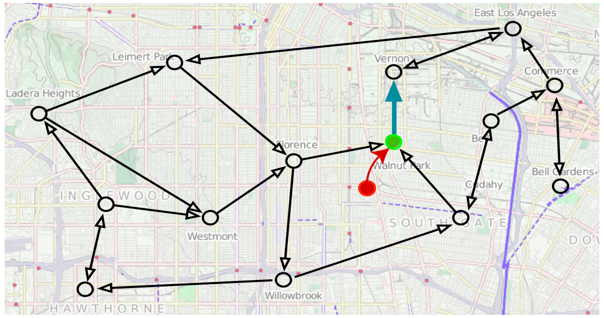

- **Markovian surveillance agents with visit frequency constraints**
- **Intelligent intruders can sense position/observe path of agent**
- Design optimal unpredictable transitions for the surveillance agents

Stochastic surveillance: Motivating example 2/2

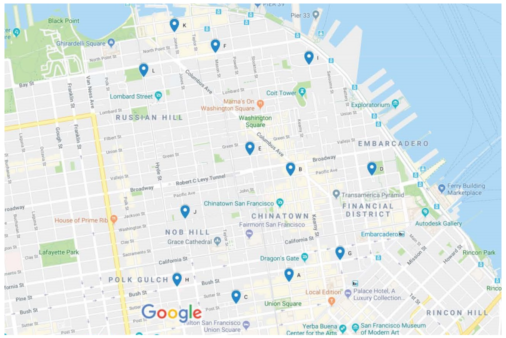

- San Francisco
- crime rate at 12 locations
- complete by-car travel times (quantized in minutes)
- define $\pi \sim$ crime rate

Rational intruder:

- Picks a node i to attack with probability π_i for duration τ
- Learns the inter-visit time statistics of surveillance agent
- Attacks at time which maximizes likelihood of not being detected

Why Markov chains for routing and planning strategies?

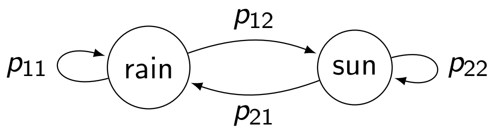

Advantages of adopting Markov chains:

1. quantify and optimize 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

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:	$\mathbb{P}[X = \text{Head}] = 0.5$	$\mathbb{H}(X) = \log 2 = 0.693$
Biased coin:	$\mathbb{P}[X = \text{Head}] = 0.75$	$\mathbb{H}(X) = 0.562$
Predictable coin:	$\mathbb{P}[X = \text{Head}] = 1$	$\mathbb{H}(X) = 0$

The entropy of what variable?

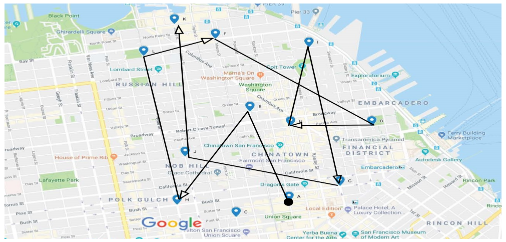

- visit random locations
- sequence of locations
- generate random symbols

- what would bank robber do?
- when does the police visit?
- learn the statistics of return times

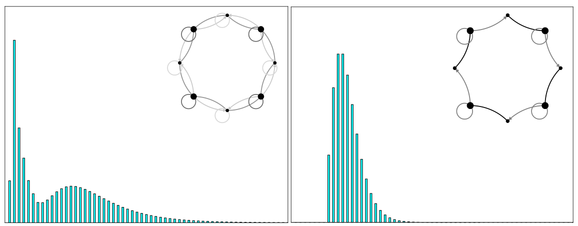

The entropy rate of a Markov chain

A classic notion from information theory

entropy rate of sequence of symbols/locations

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

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

Return time entropy of Markov chain
Better entropy notion

Consider irreducible digraph with integer travel times
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

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

Main problem statement

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

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

$$\max_{P} \ \mathbb{H}_{\text{ret-time}}(P)$$

subject to

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

X. Duan, M. George, and F. Bullo. Markov chains with maximum return time entropy for robotic surveillance.
IEEE Transactions on Automatic Control, May 2018.
Submitted.
URL: https://arxiv.org/abs/1803.07705

Outline

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

Summary of results

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

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

$$\max_{P} \ \mathbb{H}_{\text{ret-time}}(P)$$

subject to

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

Thm 1: Hitting time probability dynamics and well-posedness
Thm 2: Upper bound and solution for complete graph
Thm 3: Relations with the location entropy rate
Thm 4: Truncation, approximation and computation

Basic ideas

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

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

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

Recusive 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

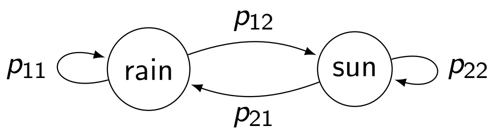

For this special case

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

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

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

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

Thm 1: Hitting time probability dynamics, well-posedness

Thm 1: Hitting time probability dynamics and well-posedness

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

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

Consider a compact set of Schur stable matrices $\mathcal{A} \subset \mathbb{R}^{n\times n}$ and let

$$\rho_{\mathcal{A}} := \max_{A\in\mathcal{A}} \rho(A) < 1.$$

Then for any $\lambda \in (\rho_{\mathcal{A}}, 1)$ and for any $\|\cdot\|$, there exists $c > 0$ such that

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

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 $f_k = F_k(i,i) \log F_k(i,i)$.

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

Thm 2: Upper bound and solution for complete graph

given expected value μ,

geometric distribution $\mathbb{P}[Y = k] = \left(1 - \frac{1}{\mu}\right)^{k-1} \frac{1}{\mu}$ is maxentropic

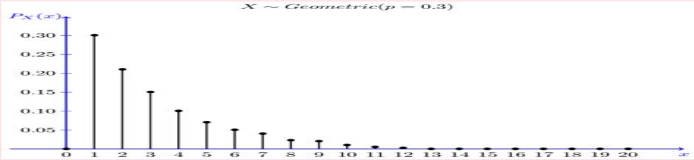

Given a strongly connected unweighted digraph, stationary distribution π,

1. the return time entropy function is upper bounded by

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

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

- upper bound only depends on the stationary distribution π
- complete unweighted $\mathcal{G}$: MaxLocationEntropy = MaxReturnEntropy

Thm 3: 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, $\mathbb{H}_{\text{ret-time}}(P)$ and $\mathbb{H}_{\text{location}}(P)$ satisfy

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

- 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: $\mathbb{H}_{\text{ret-time}}(P)$ 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:

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

In practice, the truncated return time entropy is

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

Thm 4: Truncation, approximation and computation

Thm 4: Truncation, approximation and computation

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

1. Asymptotic agreement

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

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

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

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

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

Gradient projection algorithm

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

2: **for** iteration parameter $s = 0 :$ (number-of-steps) **do**
3: $\{G_k\}_{k \in \{1,\dots,N_\eta\}} :=$ solution to Thm 4 at P_s
4: $\Delta_s :=$ gradient of $(\mathbb{H}_{\text{ret-time}})_{\text{trunc},\eta}(P_s)$
5: $P_{s+1} := \text{projection}_{\mathcal{P}^{\epsilon}_{\mathcal{G},\pi}}(P_s + (\text{step size}) \cdot \Delta_s)$
6: **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

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

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. MinKemeny: $\min_P \mathbb{E}[K(P)]$
Minimize the mean first passage time:

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

Comparison over a ring and a grid graph 1/2

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

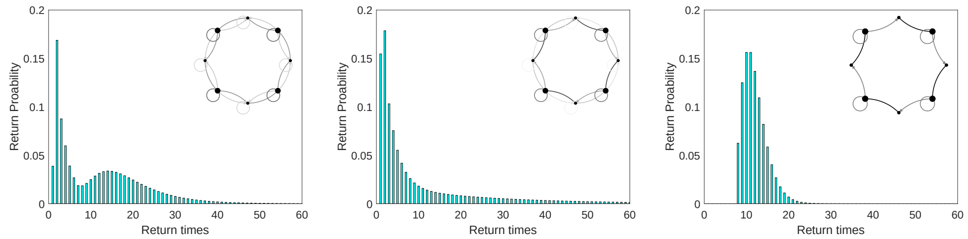

(a) MaxReturnEntropy (b) MaxLocationEntropy (c) MinKemeny

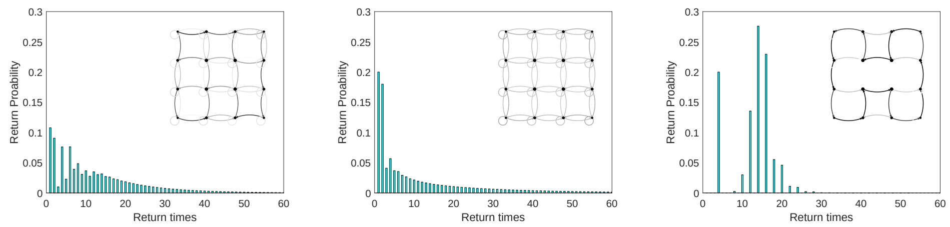

(d) MaxReturnEntropy (e) MaxLocationEntropy (f) MinKemeny

Comparison over a ring and a grid graph 2/2

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

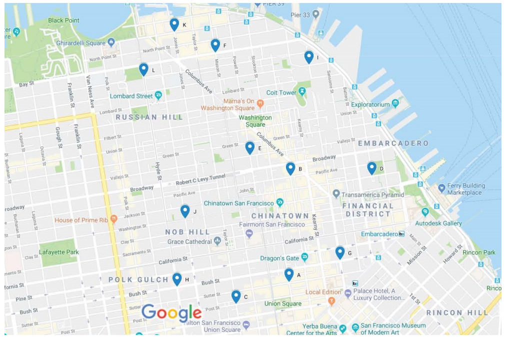

- San Francisco
- crime rate at 12 locations
- complete by-car travel times (quantized in minutes)
- $\pi \sim$ crime rate

A rational intruder:

- Picks a node i to attack with probability π_i for duration τ
- Learns the inter-visit time statistics of surveillance agent
- Attacks at time which maximizes likelihood of not being detected

Comparison over San Francisco map 2/3

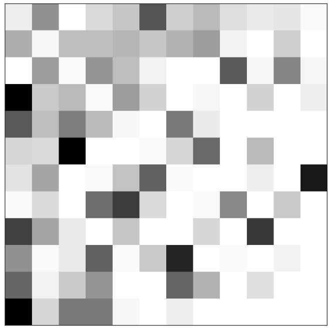

(g) MaxReturnEntropy

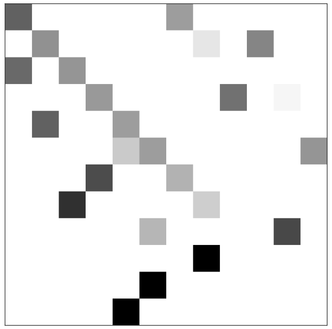

(h) MinKemeny

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

- MinKemeny 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

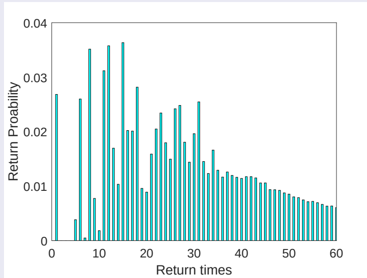

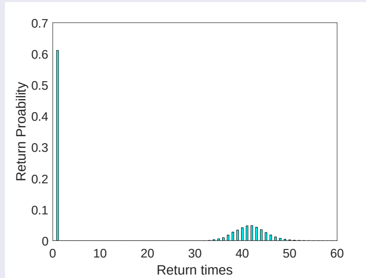

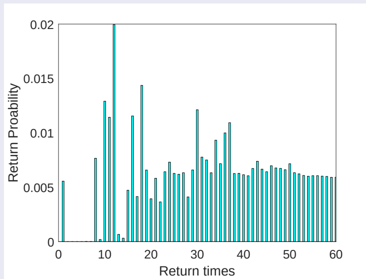

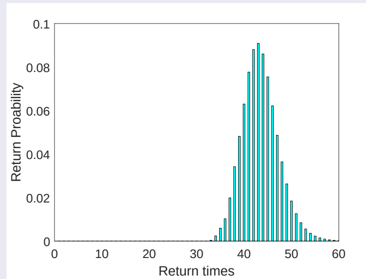

Comparison in catching the rational intruder 1/2

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

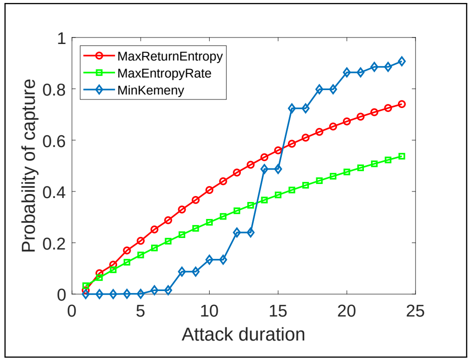

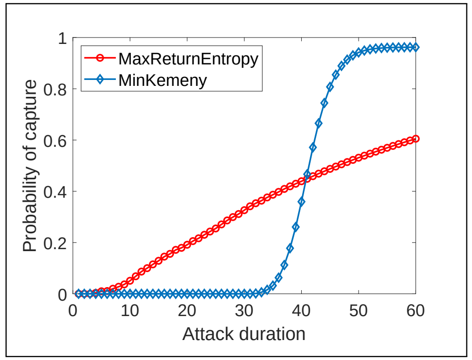

- 4×4 grid: MaxReturnEntropy $>$ MaxLocationEntropy, MaxReturnEntropy $>$ MinKemeny for short attack duration
- SF map: MaxReturnEntropy $>$ MinKemeny for short attack duration

Conclusion and future directions

Conclusion

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

Future Work

1. extensions to multi-vehicle problems
2. scalable computation for large graphs
3. transcription from continuous space/time