Signal Processing for CBR defence

Branko Ristic[1]

DST group, Australia

Australian Government
Department of Defence
Defence Science and Technology Organisation

SSPD, September 2015

[1] With contributions from A. Skvortsov, A. Gunatilaka, R. Gailis, M. Roberts, A. Walker, M. Morelande

Introduction

- CBR defence: protective measures against a
    - chemical (C)
    - biological (B)
    - radiological (R)

  attack.

- CBR defence covers a fairly broad spectrum of measures; multi-disciplinary scientific effort

- This talk will focus on signal processing techniques applied to recent and ongoing projects within the Land Division of DST group

Overview

1. Localisation of a source of biochemical agent release
2. Autonomous search techniques for CBR sources
3. Image reconstruction for a standoff gamma radiation detector
4. Forecasting of an epidemic outbreak
5. Extinction of biological systems with competition

Overview

1. Localisation of a source of biochemical agent release
2. Autonomous search techniques for CBR sources
3. Image reconstruction for a standoff gamma radiation detector
4. Forecasting of an epidemic outbreak
5. Extinction of biological systems with competition

Localisation of a biochemical source: Context

- Deliberate or accidental release of toxic material (gas, aerosols) into the atmosphere
- Problem: Given the readings supplied by a network of spatially distributed biochemical sensors, estimate the location of the source
- Technical difficulties:
    - source strength unknown;
    - characteristics of sensor measurements not well understood (plethora of air dispersion models; likelihood function only a guess)

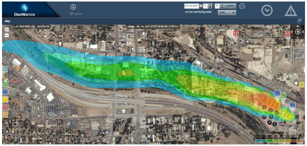

Localisation of a biochemical source: Problem formulation

- Parametric approach - parameter vector $\boldsymbol{\theta}$ includes
  - Source location (x_0, y_0, z_0), strength Q_0 (possibly size)
  - Enviro/meteo parameters
    e.g. wind speed, direction, canopy characteristics
- Concentration measurements at spatially distributed sensor locations $z_i \geq 0$, $i = 1, \ldots, S$
  - Mean concentration at sensor i: $\mathbb{E}\{z_i\} = h_i(\boldsymbol{\theta})$
    ... based on the adopted dispersion model
  - Random fluctuations (capture both the dispersion modelling errors and measurement noise)
  - Possibly quantised
- Bayesian estimation framework

$$\underbrace{p(\boldsymbol{\theta}|z_1, \ldots, z_S)}_{\text{posterior PDF}} \propto \underbrace{\ell(z_1, \ldots, z_S|\boldsymbol{\theta})}_{\text{likelihood function}} \cdot \underbrace{\pi(\boldsymbol{\theta})}_{\text{prior PDF}}$$

Dispersion plume videos

Due to turbulence:

- meandering of the plume
- intermittent measurements

(diffusion)

(diffusion + advection)

Two categories of transport of substances in the environment

- Advection - transport with the mean fluid flow
- Diffusion - transport through the action of random motions

Robust localisation of a biochemical source

- Previous approaches:
  - Adopt a dispersion model (which determines $h_i(\boldsymbol{\theta})$)
  - Adopt a certain model of random measurement fluctuations (Gaussian, log-normal)
  - Use Markov chain Monte Carlo (MCMC) to approximate the posterior $p(\boldsymbol{\theta}|z_1, \ldots, z_S)$

- Proposed framework[2]:
  - Adopt several candidate dispersion models: $\mathcal{M} = \{1, \ldots, M\}$
  - Likelihood-free estimation
  - Use approximate Bayesian computation (ABC) to simultaneously estimate $p(m|z_1, \ldots, z_S)$ and $p(\boldsymbol{\theta}_m|m, z_1, \ldots, z_S)$

[2] Ref: B. Ristic et al. 'Bayesian likelihood-free localisation of a biochemical source using multiple dispersion models', *Signal Processing*, 2015

Multiple-model ABC rejection sampler

- $\pi_{\mathcal{M}}(m)$ ... prior PDF over dispersion models $m = 1, 2, \ldots, M$
- $\pi_{\boldsymbol{\theta}_m}(\boldsymbol{\theta}_m)$ ... prior PDF over the parameter space for model m

```
Input: z = [z_1, ..., z_S]^T; ε; N
Initialise: X_1 = ··· = X_M = ∅
repeat
    Draw m* from π_M(m)
    Draw θ*_m from π_{θ_{m*}}(θ_m)
    Simulate measurement z* using model m* and parameter θ*_m
    Compute distance d* = D(z, z*)
    if d* ≤ ε then
        X_{m*} = X_{m*} ∪ {θ*_m}
    end if
until Σ_{m=1}^{M} |X_m| = N
Output: X_1, ··· , X_M
```

Input: $\mathbf{z} = [z_1, \ldots, z_S]^\intercal$; ϵ; N
Initialise: $\mathcal{X}_1 = \cdots = \mathcal{X}_M = \emptyset$
repeat
  Draw m^* from $\pi_{\mathcal{M}}(m)$
  Draw $\boldsymbol{\theta}_m^*$ from $\pi_{\boldsymbol{\theta}_{m^*}}(\boldsymbol{\theta}_m)$
  Simulate measurement $\mathbf{z}^*$ using model m^* and parameter $\boldsymbol{\theta}_m^*$
  Compute distance $d^* = D(\mathbf{z}, \mathbf{z}^*)$
  if $d^* \le \epsilon$ **then**
    $\mathcal{X}_{m^*} = \mathcal{X}_{m^*} \cup \{\boldsymbol{\theta}_m^*\}$
  end if
until $\sum_{m=1}^{M} |\mathcal{X}_m| = N$
Output: $\mathcal{X}_1, \cdots, \mathcal{X}_M$

$$p(m|\mathbf{z}) \approx \frac{|\mathcal{X}_m|}{N}; \quad p(\boldsymbol{\theta}_m|m, \mathbf{z}) \text{ approximated by the sample } \mathcal{X}_m$$

Iterative Adaptive Multiple-model ABC sampler

- Rejection sampling is inefficient (very low acceptance rate)
- We developed a more efficient scheme: rejection sampling is applied using a sequence of monotonically decreasing tolerance levels $\epsilon_1 > \epsilon_2 > \cdots > \epsilon_T > 0$;
- The subsequent approximations gradually approach the true posterior
- Similar to the SMC-ABC sampler (Toni et al., 2009), but no need to specify in advance tolerances $\epsilon_1, \ldots, \epsilon_T$; instead, iteration t computes the tolerance for the next iteration, ϵ_{t+1}

Dispersion models

- $M = 3$ models considered:
  - $m = 1$: Gaussian plume with a linear spread, dim$(\boldsymbol{\theta}_1) = 7$
  - $m = 2$: Gaussian plume with a nonlinear spread, dim$(\boldsymbol{\theta}_2) = 9$
  - $m = 3$: Stretch exponential model, dim$(\boldsymbol{\theta}_3) = 10$

- Gaussian plume model
  - Gaussian distribution of the plume in the vertical and horizontal directions.
  - By convention, x-axis coincides with the average wind velocity vector
  - Mean concentration at sensor i

$$C_i = h_i(\boldsymbol{\theta}) = \frac{Q_0}{2\pi\sigma_{y_i}\sigma_{z_i}U}\, e^{-\frac{(y_i-y_0)^2}{2\sigma_{y_i}^2}} \left[e^{-\frac{(z_i-z_0)^2}{2\sigma_{z_i}^2}} + e^{-\frac{(z_i+z_0)^2}{2\sigma_{z_i}^2}} \right]$$

with $\sigma_{y_i} = \frac{\sigma_v}{U}(x_i - x_0)$, $\sigma_{z_i} = \frac{\sigma_w}{U}(x_i - x_0)$ for $m = 1$ (point source)

Experimental dataset for algorithm evaluation

- Collected using a recirculating water channel (COANDA R&D Corp), 10m × 1.5m × 0.9m
- The source: constant release of fluorescent dye from a narrow tube
- Concentration measurements collected at $S = 48$ downstream positions using laser induced fluorescence (LIF); averaged over 100 sec
- The floor of the w/c covered with a mesh to give surface roughness
- Two scenarios considered:
  - without obstacles (open terrain);
  - with obstacles (urban environment)

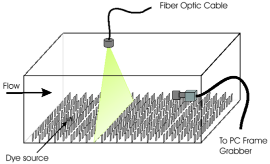

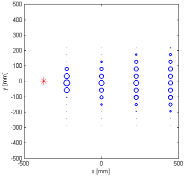

Top-down view of the experimental setup

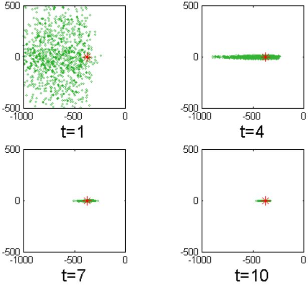

Localisation at iterations $t = 1, 4, 7, 10$

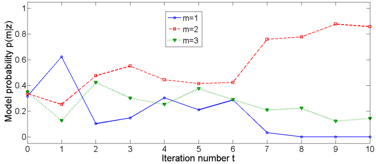

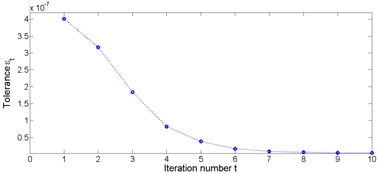

Experimental dataset: Results (2)

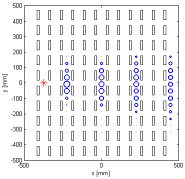

Top-down view of the experimental setup

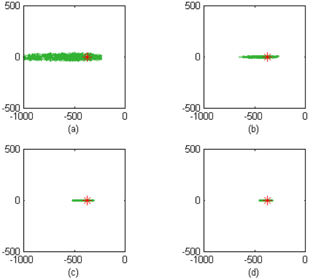

Localisation at iterations $t = 3, 6, 9, 12$

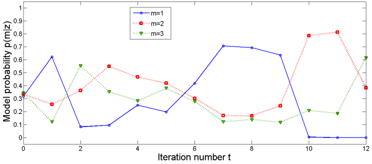

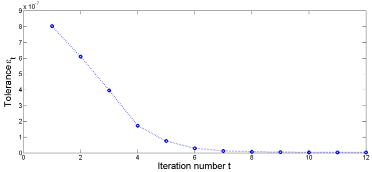

Localisation of a biochemical source: CR bounds (1)

Assuming:

- Gaussian plume dispersion model
- Additive white Gaussian measurement noise

Derived the theoretical Cramer-Rao lower bounds for estimation[3] of $\boldsymbol{\theta}$

- analog measurements: $z_i = h_i(\boldsymbol{\theta}) + w_i$,
- binary measurements: $b_i = \begin{cases} 1, & z_i > \tau \\ 0, & z_i \leq \tau \end{cases}, \; i = 1, \ldots, S$

Localisation information

analog

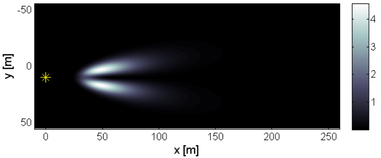

binary

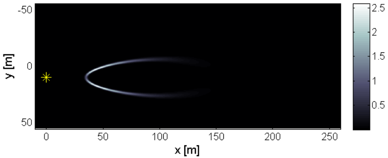

[3] Ref: B. Ristic et al., "Achievable accuracy in Gaussian plume parameter estimation ...", *Information Fusion*, 2015.

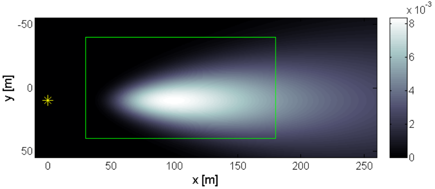

Two (uniform) placements of sensors:

1. $S = 36$
2. $S = 171$

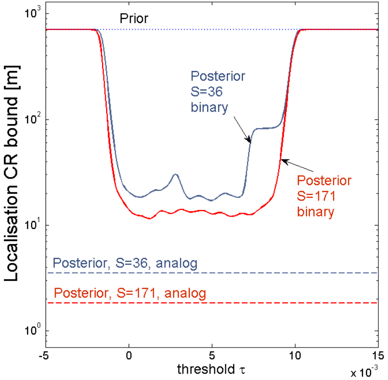

Localisation of a biochemical source: binary measurements

Problem:

- Identical spatially distributed binary sensors
- Binary threshold τ is unknown (vague prior)
- Semiurban environment

Question: Can we localise the source (estimate $\boldsymbol{\theta}$) in this case (since the source strength Q_0 is also unknown)?

Answer: Yes, replace Q_0 (in $\boldsymbol{\theta}$) with Q_0/τ.

Localisation using experimental binary data (1)

- State vector: $\boldsymbol{\theta} = [x_0 \quad y_0 \quad Q_0/\tau \quad U]^\intercal$
- Dispersion model: the count of particle encounters in 2D (Vergassola et al, Nature, 2007)
- Likelihood function for the measurement vector $\mathbf{b} = [b_1, \ldots, b_S]^\intercal$ derived assuming that z_i is Poisson distributed

Prior distributions:

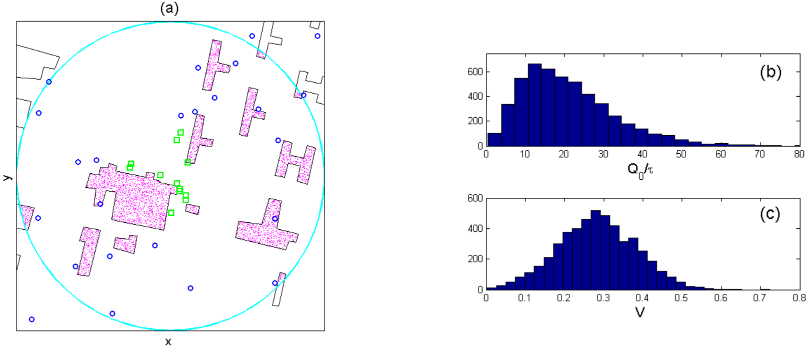

Localisation using experimental binary data (2)

Posterior distributions[4]

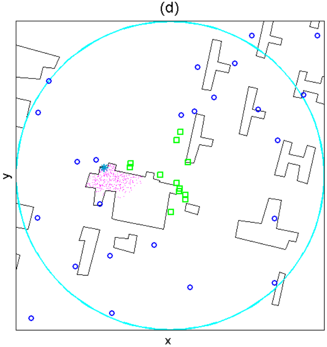

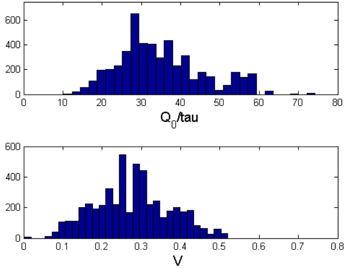

- Estimation carried out using the importance sampling technique
- Similar results obtained for three experimental datasets

[4] Ref: B. Ristic et al. 'Source Localisation in turbulent flow using a binary sensor network', *IEEE SPL* (in review), 2015

Overview

1. Localisation of a source of biochemical agent release
2. Autonomous search techniques for CBR sources
3. Image reconstruction for a standoff gamma radiation detector
4. Forecasting of an epidemic outbreak
5. Extinction of biological systems with competition

Autonomous search vs Localisation

- Consider sensors that can be controlled, e.g.
  - where to go
  - how long to stay
- Search is a repetitive cycle of sensing, estimation (localisation), and sensor control
- Decisions (controls, actions) are made sequentially and autonomously (AI) in the presence of uncertainty, using only past measurements and past actions.
- Goal: minimise the search time via on-line control of individual sensors (need to define in advance the termination criterion)

Search Research

- Search strategies for finding a source of an emission (gas, particles, smell, energy) based on sparse cues are of great importance
  - National security (e.g. release of hazardous substances)
  - Recovery & rescue emissions (e.g. MH370)
  - Understanding nature (foraging behaviour of animals)
- Very popular subject; huge amount of literature (top journals); various scientific disciplines (e.g. biology, physics, operations research, robotics)
- Categories:
  - Deterministic (systematic search) vs random search
  - Random search models: diffusion, Lévy flight, intermittent search (exploration vs exploitation)
  - Gradient search (chemotaxis) vs information driven search (infotaxis)

Sequential decisions under uncertainty

"Partially observed Markov decision problem" formulation

The elements of POMDP:

- A set of admissible actions (controls) at discrete-time k: $\mathbb{U}_k$;
- The (information) state of the system at time k: $\mathcal{I}_k$
- A reward function associated with each action $\mathbf{u} \in \mathbb{U}_k$: $\mathcal{R}\big(\mathbf{u}, \mathcal{I}_k\big)$

The best action should be based on the reward computed a few steps ahead.

One step ahead optimal sensor control:

$$\mathbf{u}_k = \arg\max_{\mathbf{v} \in \mathbb{U}_k} [\mathcal{R}(\mathbf{v}, \mathcal{I}_k)]$$

Bayesian - Information Theoretic Approach

- The (information) state of the system is the *posterior PDF*:

$$\mathcal{I}_k = p(\boldsymbol{\theta}|\mathbf{z}_{1:k}, \mathbf{u}_{0:k-1}), \qquad (\mathbf{z}_{1:k} \equiv \mathbf{z}_1, \mathbf{z}_2, \ldots, \mathbf{z}_k)$$

- Sequential Bayesian estimation; particle filter implementation

- Reward function: an information theoretic measure, e.g.
  - Entropy difference: $\mathcal{R} = \mathbb{E}\{H_{k+1}(\mathbf{u}, \mathcal{I}_{k+1})\} - H_k(\mathcal{I}_k)$
    where $H_k(\mathcal{I}_k) = -\int p(\boldsymbol{\theta}|\mathbf{z}_{1:k}, \mathbf{u}_{0:k-1}) \ln p(\boldsymbol{\theta}|\mathbf{z}_{1:k}, \mathbf{u}_{0:k-1}) d\boldsymbol{\theta}$
  - Rényi divergence

$$\mathcal{R} = \mathbb{E}\left\{\frac{1}{\alpha - 1} \log \int [p(\boldsymbol{\theta}|\mathbf{z}_{1:k+1}, \mathbf{u}_{0:k})]^{\alpha} [p(\boldsymbol{\theta}|\mathbf{z}_{1:k}, \mathbf{u}_{0:k-1})]^{1-\alpha} d\boldsymbol{\theta}\right\}$$

($\alpha = 1 \Rightarrow$ Kullback-Liebler div.; $\alpha = 0.5 \Rightarrow$ Bhattacharyya dist.)

Autonomous search: demo

- Turbulent diffusion
- Real (COANDA) data

Autonomous search demo: Explanation

- Likelihood function: Poisson distribution

- Dispersion model: the count of particle encounters in 2D (Vergassola et al., Nature, 2007); the mean of the Poisson distribution at sensor i

$$C_i = h_i(\boldsymbol{\theta}) = \frac{Q_0}{\ln\left(\frac{\lambda}{a}\right)} \exp\left[\frac{(x_0 - x_i)U}{2D}\right] \cdot K_0\left(\frac{\sqrt{(x_i - x_0)^2 + (y_i - y_0)^2}}{\lambda}\right)$$

where $\lambda = \sqrt{(DT_o)/(1 + \frac{U^2 T_o}{4D})}$,
a - size of the sensor; D - diffusivity; T_0 - mean lifetime of a particle; U - mean wind speed

- The reward function: Entropy difference

Autonomous search: study[5]

- All information theoretic based reward functions perform similarly.

- An important factor in search performance is the ratio $\rho = \frac{\mathcal{A}_{search}}{\mathcal{A}_{sense}}$.

The smaller ρ, the more efficient information driven search strategy (infotaxis) is compared to the systemat search (and vice versa)

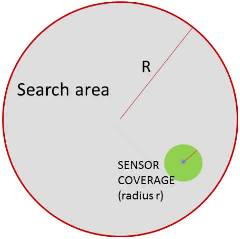

[5] B. Ristic et al., A study of cognitive strategies for an autonomous search using sparse cues, *Information Fusion*, 2015

The PDF of search time (1)

Let the reward function be the entropy difference.

Postulation: Dynamic model of entropy during the search is a Gauss-Markov random walk with a constant (negative) drift.

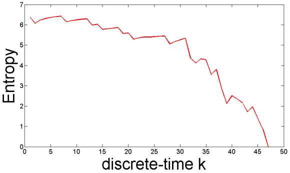

Brownian motion (as a continuous-time limiting case of a Gauss-Markov random walk) has the following property:

Distribution of the "first passage time"

The time which Brownian motion $X_t = \nu t + \sigma W_t$, with a drift $\nu < 0$, takes to reach a certain (negative) level, is an inverse Gaussian RV

The PDF of search time (2)

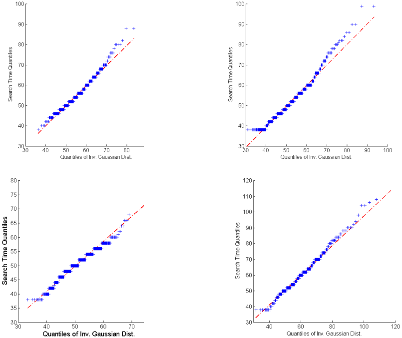

QQ plots of 4 samples of search time (varying U and Q_0)

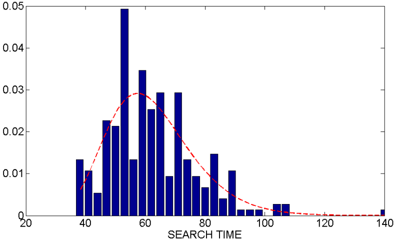

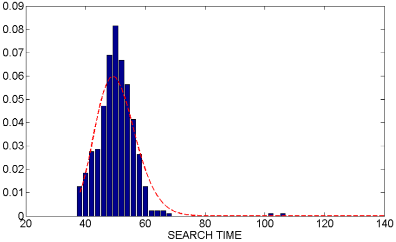

Autonomous search for radiological point sources[6]

- Sensor: low-cost, non-directional (Geiger counter)
- Unknown number of sources $R \geq 0$; unknown intensities
- Control vectors: where to go, how long to stay
- Propagation model: *inverse distance squared*

$$C_i = h_i(\boldsymbol{\theta}) = \sum_{r=1}^{R} \frac{Q_0^{(r)}}{\left[x_i - x_0^{(r)}\right]^2 + \left[y_i - y_0^{(r)}\right]^2}$$

- Likelihood function: Poisson, i.e. $z_i \sim \mathcal{P}(h_i(\boldsymbol{\theta}) + \beta)$

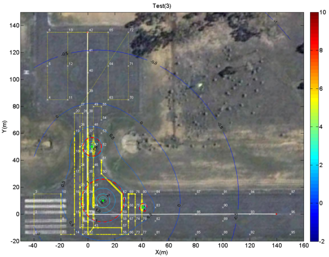

Top down view of the data experimental set-up

[6] B. Ristic et al., Information driven search for point sources of gamma radiation, *Signal Processing*, 2010

Autonomous search for radiological point sources (cont'd)

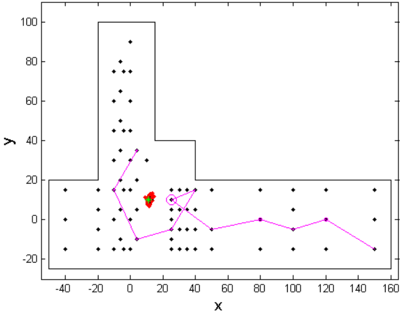

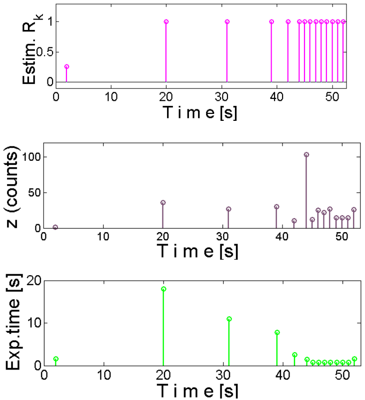

Autonomous search in an unknown structured environment

- Context[7]:
  - A diffusive source of toxic substance (gas, particles)
  - Search domain contains obstacles
  - GPS denied environment; the map is unknown
  - The initial position and the search domain boundary are given
- Goal: Estimate the source location and the path to it!
- Assumption: The searcher can sense

1. the concentration (of toxic substance)
2. the presence / absence of obstacles

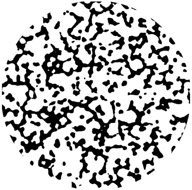

[7] B. Ristic et al., Autonomous search for a diffusive source in an unknown structured environment, *Entropy*, 2014

Discretisation of the search area

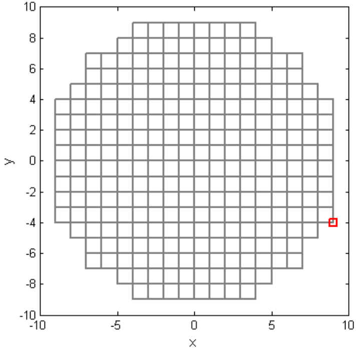

Complete lattice
(known a priori)

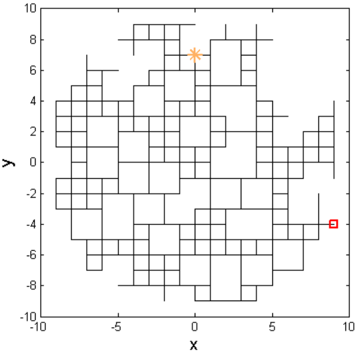

Incomplete lattice (unknown)
(obstacles are missing links, $\sim 35\%$)
(percolation threshold)

Search in an unknown structured environment: demo (1)

- The correct likelihood function depends on the map - because the map is unknown, the searcher uses an approximation derived using the conformal mapping technique
- Execution of motion controls prone to error (with prob $p_e \ll 1$)
- The state vector $\boldsymbol{\theta}$ includes:
  1. Coordinates (x_0, y_0), intensity of the source Q_0
  2. The map (i.e. existence of all links in the lattice)
  3. The searcher position on the map
- Admissible actions: $\mathcal{U}_k = \{\cdot, \rightarrow, \leftarrow, \uparrow, \downarrow\}$
- Reward function: Bhattacharyya distance (over x_0, y_0, Q_0 only)

Search - unknown structured environment: demo (2)

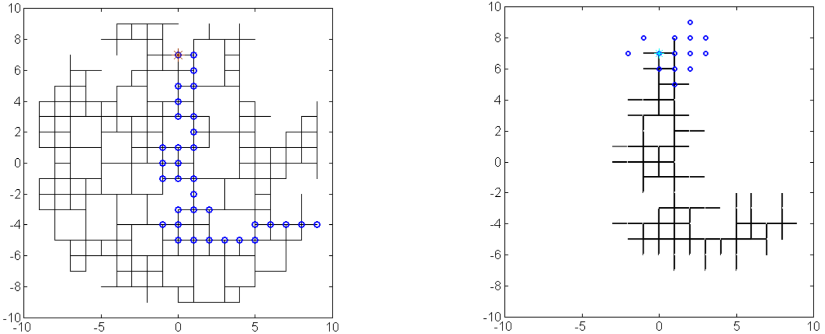

Examples on a bigger scale

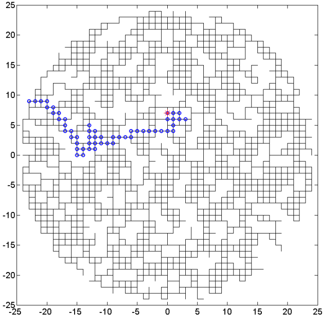

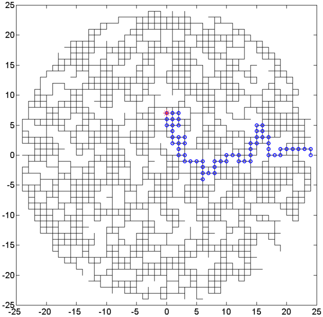

Overview

1. Localisation of a source of biochemical agent release
2. Autonomous search techniques for CBR sources
3. Image reconstruction for a standoff gamma radiation detector
4. Forecasting of an epidemic outbreak
5. Extinction of biological systems with competition

Gamma radiation image reconstruction[8]

- DSTO recently built a prototype of a standoff imaging gamma radiation detector
- Purpose: determine exact locations of radiological sources within the field of view
- Rotational modulation collimation
  - Two attenuating masks separated by a known distance
  - Co-rotating on a cylinder in front of three gamma ray detectors
- Problem: Reconstruct the image from its RMC projection

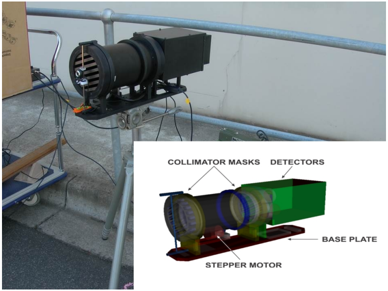

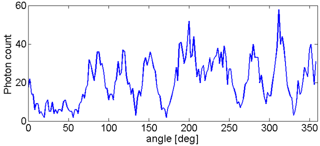

[8] Ristic, Roberts, "A parametric Bayesian RMC gamma-ray image reconstruction", ICASSP 2015

RMC image reconstruction: parametric vs non-parametric

- Non-parametric formulation
  - Emission tomography (reconstruction of medical images)
  - Standard algorithms: EM, MAP

$$\text{Measurements: } y_i \sim \mathcal{P}_{y_i}\left(\sum_{j=1}^{N} A_{ij}\lambda_j\right), \qquad (\text{rot. angles } i = 1, \ldots, M)$$

  - Goal: reconstruct image $\boldsymbol{\lambda} = [\lambda_1, \ldots, \lambda_N]^\intercal$ using $\mathbf{y} = [y_1, \ldots, y_M]^\intercal$
- Parametric formulation
  - Motivation: limited projections ☹; simpler images ☺
  - Image: a weighted sum of Gaussian radial basis functions:

$$\lambda_j(\boldsymbol{\theta}) = \sum_{k=1}^{Q} \lambda_{jk}(\boldsymbol{\theta}), \text{ where } \lambda_{jk}(\boldsymbol{\theta}) = w_k e^{-\frac{(x_j - \bar{x}_k)^2 + (y_j - \bar{y}_k)^2}{\sigma_k^2}}$$

  - Goal: reconstruct the image by estimating the $4Q$-dimensional parameter vector $\boldsymbol{\theta} = \left[\mathbf{w}^\intercal\ \bar{\mathbf{x}}^\intercal\ \bar{\mathbf{y}}^\intercal\ \boldsymbol{\sigma}^\intercal\right]^\intercal$

Bayesian parametric RMC image reconstruction

- Bayesian parameter estimation framework: $\pi(\boldsymbol{\theta}|\mathbf{y}) \propto p(\mathbf{y}|\boldsymbol{\theta})\,\pi_0(\boldsymbol{\theta})$
- Implementation: importance sampling with progressive correction and Metropolis-Hastings step
- The proposed method works much better for extended (non-point) sources
- Work in progress: model selection for unknown Q; fast estimation(to deal with large Q, e.g. non-homogeneous background radiation)

Point source (Cs-137), distance 260m

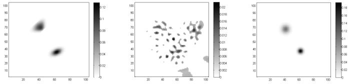

True EM Proposed

Overview

1. Localisation of a source of biochemical agent release
2. Autonomous search techniques for CBR sources
3. Image reconstruction for a standoff gamma radiation detector
4. Forecasting of an epidemic outbreak
5. Extinction of biological systems with competition

Dynamics of infectious disease

- Epidemic:
  - a complex dynamic stochastic system
  - on the macroscopic level: compartmental models of disease transmission
- Example: SEIR (compartmental) model (measles, pox, influenza, ...)

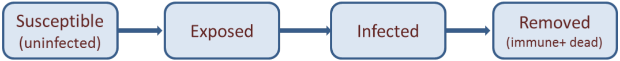

(ignoring birth, migration, death by natural causes)

$$\dot{S} = -\beta IS/N$$
$$\dot{E} = \beta IS/N - \alpha E$$
$$\dot{I} = \alpha E - \gamma I$$
$$R = N - (S + E + I)$$

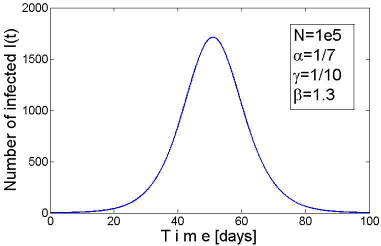

Sequential Epidemic State Estimation (and Prediction)

- Nonlinear (stochastic) filtering framework
- Stochastic SEIR model: Gillespie algorithm, approximations

$$\begin{aligned} S_k &= S_{k-1} - \nu_k, & \nu_k &\sim \mathcal{P}(\beta I_{k-1} S_{k-1} \Delta / N) \\ E_k &= E_{k-1} + \nu_k - \mu_k, & \mu_k &\sim \mathcal{B}(E_{k-1}, e^{-\Delta\alpha}) \\ I_k &= I_{k-1} + \mu_k - \eta_k, & \eta_k &\sim \mathcal{B}(I_{k-1}, e^{-\Delta\gamma}) \end{aligned}$$

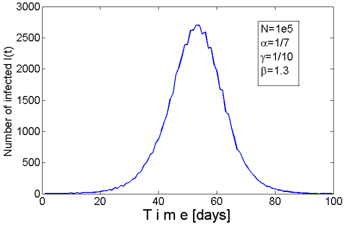

- Measurements
  - Cumulative number of infected cases (e.g. Ebola)
  - Syndromic surveillance (e.g. flu - number of Google searches)
- Implementation: particle filter[9]

[9] Skvortsov, Ristic, "Monitoring and prediction of an epidemic outbreak", *Math. Biosciences*, 2012

Application: Ebola virus epidemic in west Africa

- WHO data (available online)
- Likelihood function unknown: Neg-Binomial vs Poisson

Sierra Leone, 95 days prediction

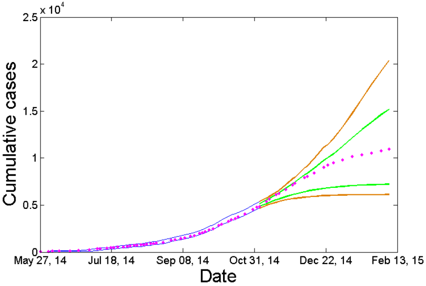

Guinea, 97 days prediction

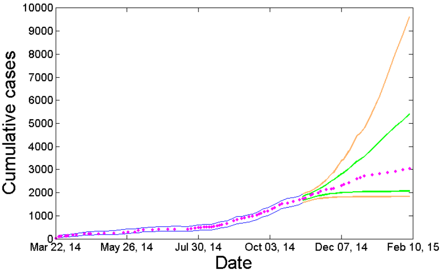

10

[10] Ristic, Dawson, "Forecasting an Epidemic Outbreak: Application to Ebola Cases", *Signal Processing* (in review), 2015

Overview

1. Localisation of a source of biochemical agent release
2. Autonomous search techniques for CBR sources
3. Image reconstruction for a standoff gamma radiation detector
4. Forecasting of an epidemic outbreak
5. Extinction of biological systems with competition

Extinction and survival in a competitive world

- Central themes of population biology
- Relevant in a broader context of complex stochastic dynamic systems (e.g. virus evolution, stock market trading)
- Mathematical model:
  single-prey multiple-predator Lotka-Volterra (LV-1n) system

- prey population (food for predators): x
- competing predators: $y_1, \ldots, y_n$

$$
\begin{aligned}
\dot{x} &= \alpha\, x \left(1 - \sum_{i=1}^{n} \beta_i\, y_i\right) \\
\dot{y}_i &= \beta_i\, x\, y_i - \gamma_i\, y_i
\end{aligned}
$$

$\alpha, \beta_1, \gamma_1, \ldots, \beta_n, \gamma_n > 0$

Deterministic model - no extinction!

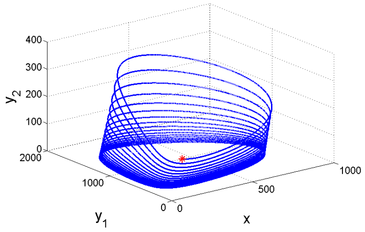

PDF of extinction time for a stochastic LV-1n

$2n+1$ biochemical reactions:

$$\begin{array}{rcl} X & \xrightarrow{\alpha} & 2X \\ X + Y_i & \xrightarrow{\beta_i} & 2Y_i \;\; i = 1, \ldots, n \\ Y_i & \xrightarrow{\gamma_i} & \emptyset \;\; i = 1, \ldots, n \end{array}$$

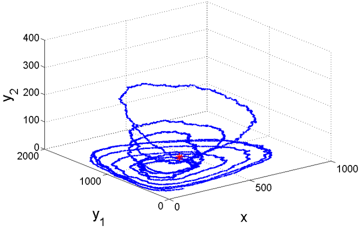

Master eq. $\Rightarrow$ exact simulation algorithm (Gillespie)

- We are after the PDF of extinction time T
- Simplification: competing predators collapse into one aggregated predator y (parameters β, γ)
- Analytic expression for the PDF of $\tilde{T} = T/\tau$ is then known (Kamenev, Parker, 2009)

$$p(\tilde{T}) = \frac{a}{\sqrt{\pi \tilde{T}^3}} \exp\left[-(\tilde{T} - a)^2 / \tilde{T} \right]$$

where: τ is a (known) function of $x(0)$, $y(0)$, α, β and γ; $a \approx 0.5$

PDF of extinction time: example

- Suppose we observe occasionally, over a period of time, the prey count of an LV-1n system[11](N.B. n can be unknown!)
- Assuming:
  - an aggregated stochastic LV-11 model
  - Poisson likelihood of prey count measurements

  we can estimate $x(0)$, $y(0)$, α, β and γ (e.g. using the pMCMC algorithm), find τ and the PDF of T

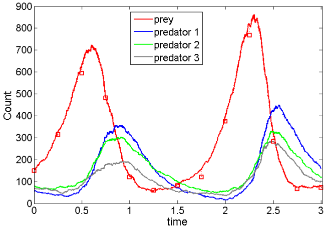

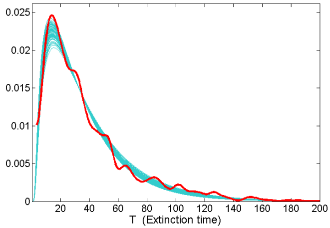

[11] Ristic, Skvortsov, "Predicting extinction of a biological system with competition", Chapter 25 in *Emerging Trends in Computational Biology, Bioinformatics, and Systems Biology*, 2015

Summary and future work

- Review signal processing techniques applied to a few recent and ongoing CBR defence projects in Land Division of DSTO

- Future work:
  - Practical: build autonomous search robots
  - Theoretical: distributed multi-platform search (estimation, control), intermittent search strategies, parameter estimation of large stochastic reaction networks (system biology: signaling pathways)

Questions ???