

Points, Particles and Positions:

Recent Advances in Distributed Processing of Agile Objects

Simon Godsill

University of Cambridge

www-sigproc.eng.cam.ac.uk/~sjg

SIGNetS – Signal and Information Gathering for Networked Surveillance



The Team

- Cambridge University (*Lead Organisation*):
 - Signal Processing and Communications (SigProC) Laboratory, Department of Engineering.
 - PI Prof. Simon Godsill
- Sheffield University.
 - Department of Automatic Control and Systems Engineering (ACSE)
 - PI Prof. Lyudmila Mihaylova
- Surrey University.
 - Centre for Vision, Speech and Signal Processing (CVSSP), and 5G Innovation Centre (5GIC)
 - PI Prof. Wenwu Wang

Project Partners

- Prof. Simon Godsill, **University of Cambridge**
- Prof. Lyudmila Mihaylova, **University of Sheffield**
- Prof. Wenwu Wang and Prof. Pei Xiao, **University of Surrey**



Researchers at Cambridge:

Lily (Qing) Li

and Patrick (Runze) Gan





SIGNets

Signal and Information Gathering for Networked Surveillance

Objectives:

1. Localisation, learning of latent intentions/network structures
2. Partially decentralised processing, limited comms.
3. Decision-making/tasking of sensors

UNIVERSITY OF CAMBRIDGE



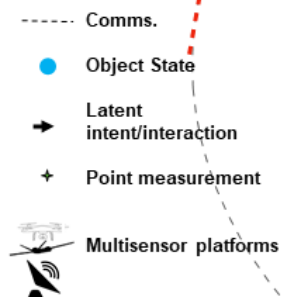
UNIVERSITY OF SURREY

Constrained network communication

DEVCOM

ARL

[dstl]



Extended Object Tracking

Groups/extended objects with varied intent/interactions

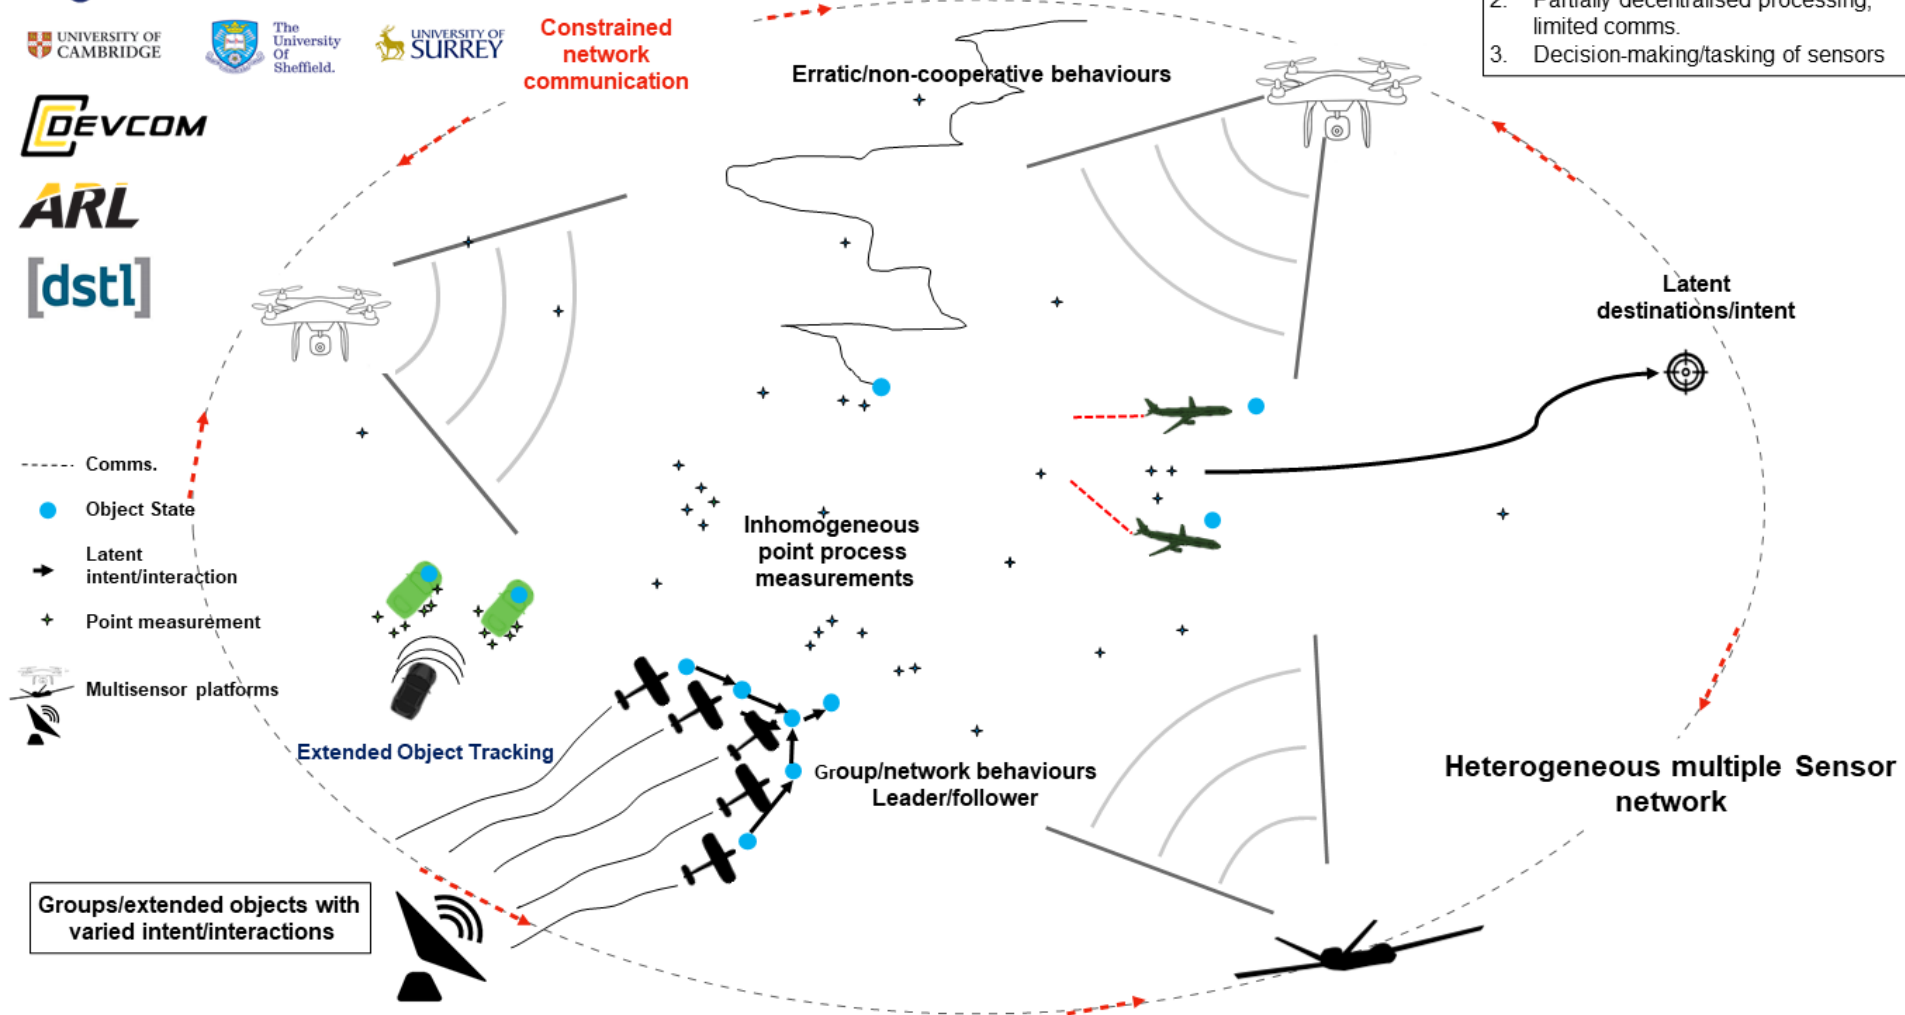
Heterogeneous/multiple sensor data, e.g.

- RADAR
- LiDAR
- Static detectors
- Video
- Road-grid/topological info.

'Allowing networked sensors to see, interpret, and communicate the status of the environment'

Methodologies:

- Heterogeneous point processes
- Probabilistic network interaction/intent models
- Optimal sequential Bayesian inference combined with:
- Variational/message passing approximations
- Novel communications methods and ML



Overview of Talk

- Point Process methods for agile targets:
 - Theory
 - Applications in Intent modelling and localisation methods
 - Non-Gaussian Process Regression
- Point Process methods for sensor modelling:
 - Theory
 - Applications in multiple object tracking, object detection, dynamic target/clutter rate inference, and variational Bayes methods



Point Process methods for agile targets

- Irregular Movement (e.g. animal foraging, drones, etc.)

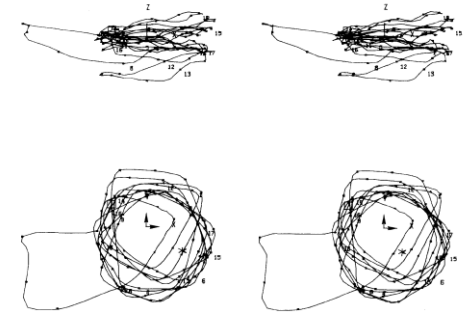
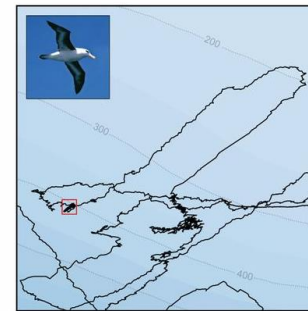
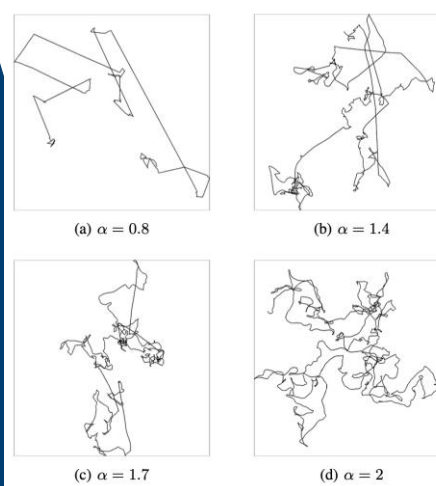
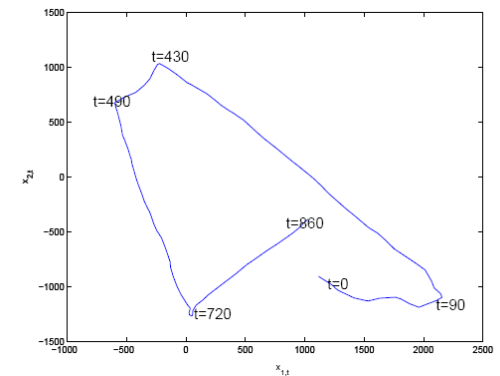


Fig. 4. Stereo plots of the flight path of a fly patrolling the airspace beneath a landmark seen from the side (top) and from below (bottom). Star marks the position of the landmark (a black box suspended from the ceiling). Arrows mark the X-, Y-, and Z-direction of the external coordinate system and are 5 cm long. The start of the recorded sequence is marked by an open square. Fly positions are marked every 0.2 s and consecutively numbered every second. The whole sequence is 20 s long. The sampling rate is 25%. The angular separation of the stereo images is 7°.

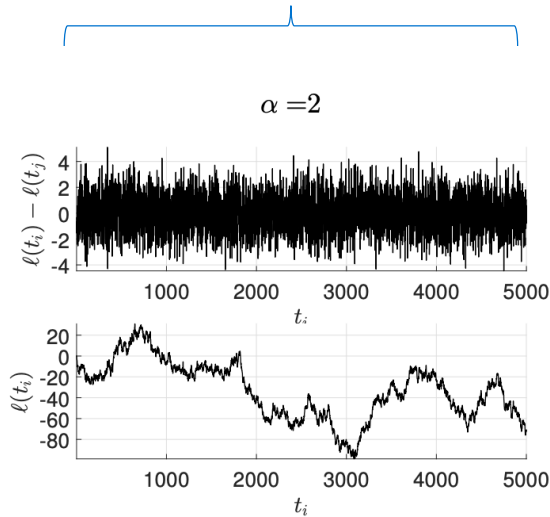


1

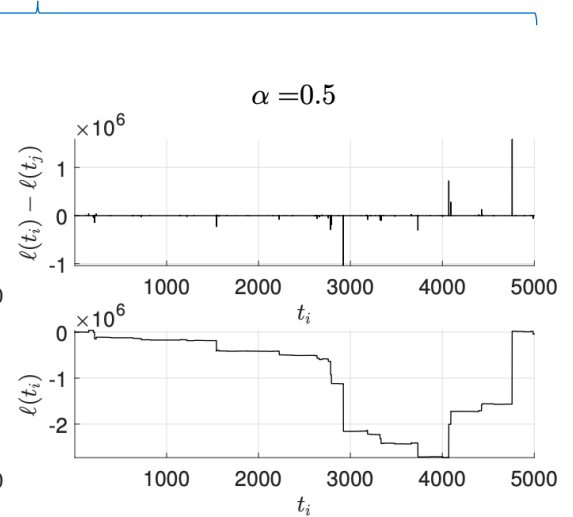
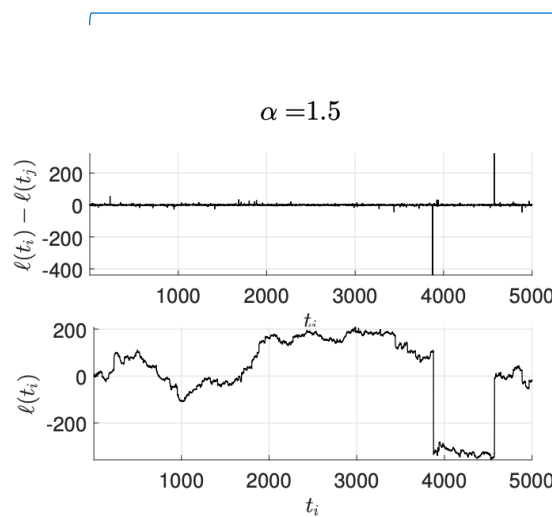


Theory: Levy process models of non-Gaussianity

Gaussian



Non-Gaussian (heavy-tailed)



- Wish to model broad classes of heavy-tailed driving noise to suit application
- Adopt a generic Levy process approach in which driving noise is modelled as *pure jump* processes in continuous time (and observed at random discrete times).
- Elegant and (fairly!) simple alternative to the standard Gaussian (Brownian motion)
- Many possible distributions: alpha-stable, Generalised Hyperbolic (inc. Student-t, normal-Gamma and normal-inverse Gaussian), normal tempered-stable, ...

Brownian vs non-Gaussian state space models

- We will be interested in tracking and other models that can be written in a state-space form with state evolution in continuous time:

$$dx(t) = Ax(t)dt + HdW(t)$$

and we will need to be able to characterise transition PDFs for treatment of discrete time measurements:

$$p(x(t)|x(t - \Delta))$$

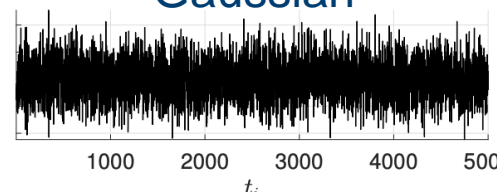
- If $\{W(t)\}$ is Gaussian (Brownian) then all calculations are simple (Kalman filter, standard stuff...)
- We are extending to a non-Gaussian $W(t)$ that moves *only* by small perturbations at random times $\{\tau_i\}_{i=1}^{\infty}$ ('jumps'), using a special conditionally Gaussian class ('Mean and Scale mixture of Gaussians') such that at any jump time τ_i :

$$dW(\tau_i) \sim \mathcal{N}(x_i\mu_W, x_i\sigma_W^2)$$

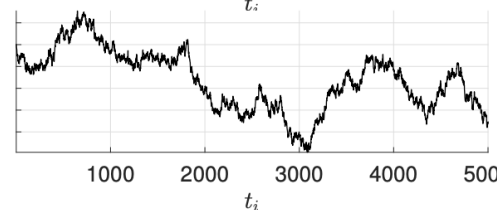
where $\{x_i\}_{i=1}^{\infty}$ are the jumps of another 'subordinator' process $X(t)$

- Because of the conditionally Gaussian structure we can implement the new models using banks of *standard architectures* (Kalman filters etc.) inside standard particle filters
- And because of the 'discrete' structure (jumps) it is even easier to compute $p(x(t)|x(t - \Delta))$ than the pure Gaussian case for each Kalman filter

Gaussian

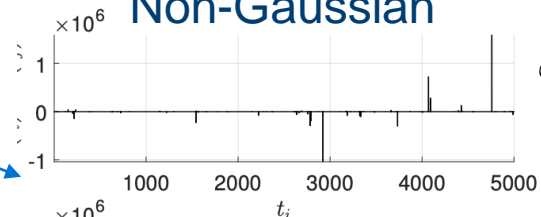


' $dW(t)$ '

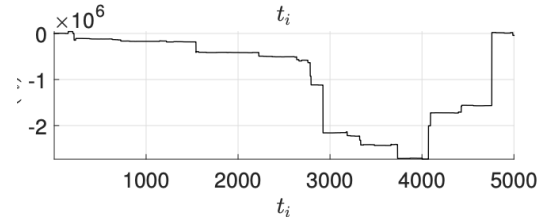


$W(t)$

Non-Gaussian



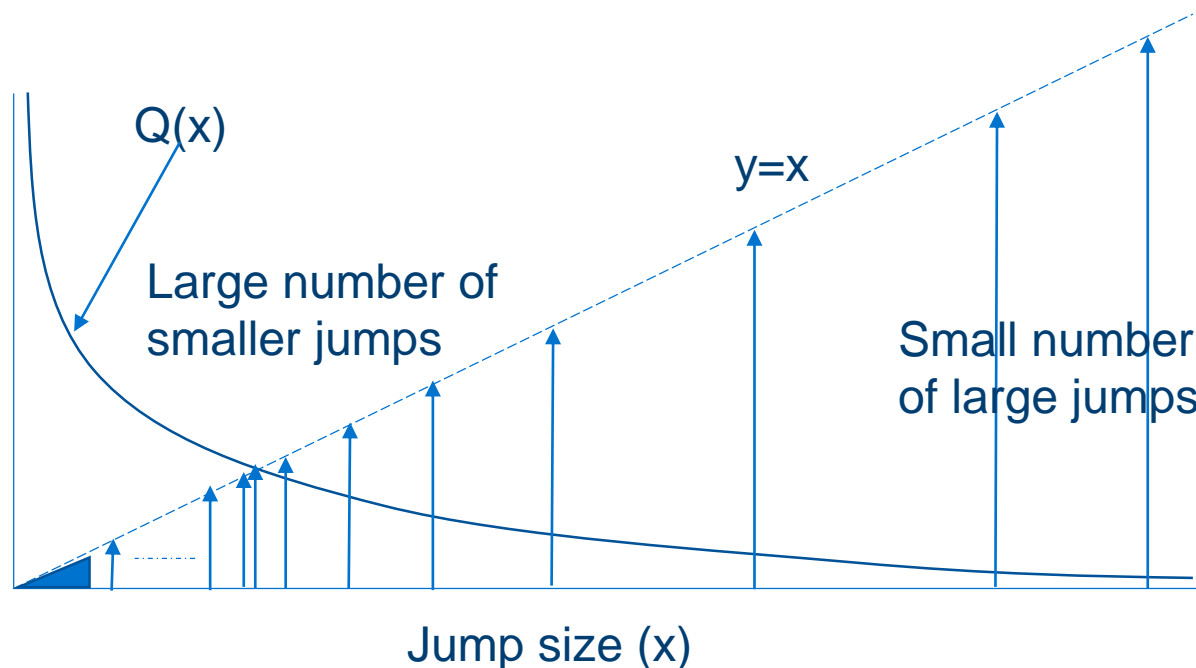
' $dW(t)$ '



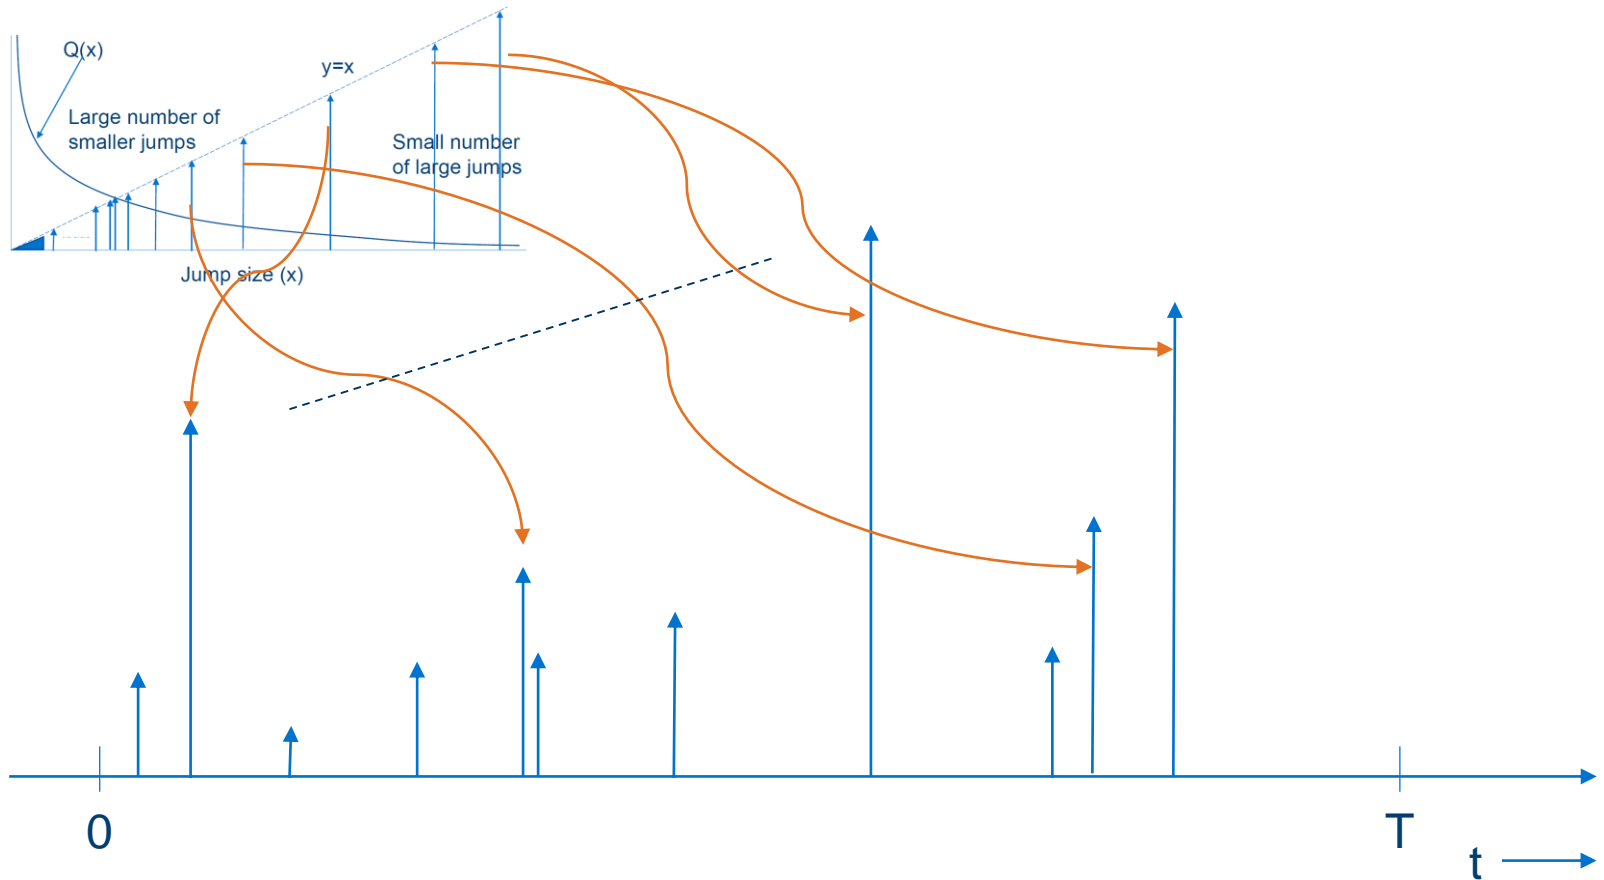
$W(t)$

Theory: Levy process models of non-Gaussianity

- The Jumps $\{x_i\}_{i=1}^{\infty}$ are characterised by a Poisson process with non-uniform intensity function $Q(x)$, the 'Levy Density':

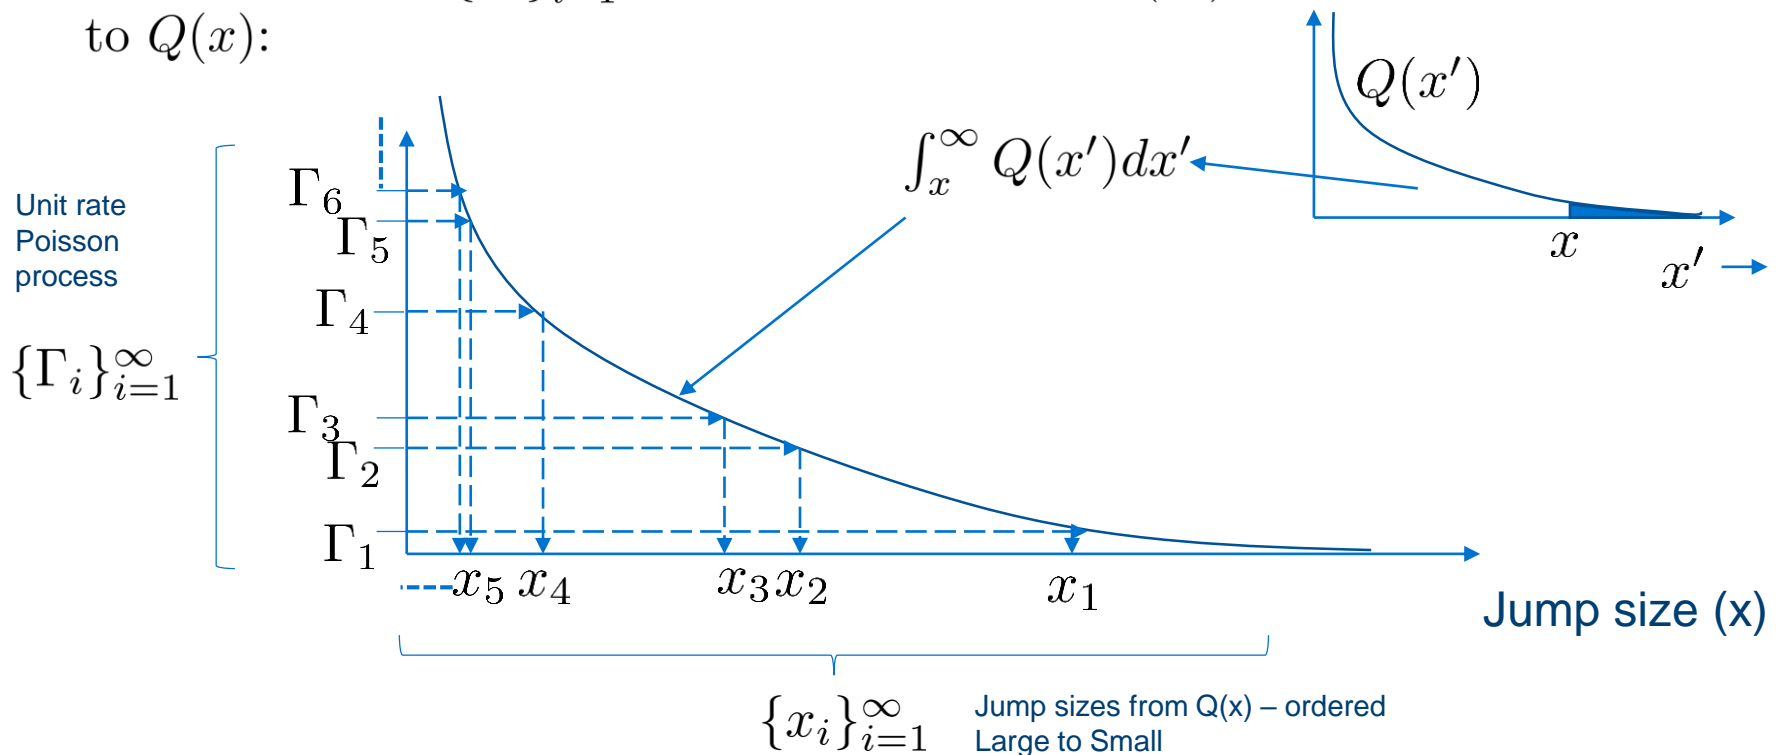


- Jumps are then uniformly randomly scattered across the time axis:



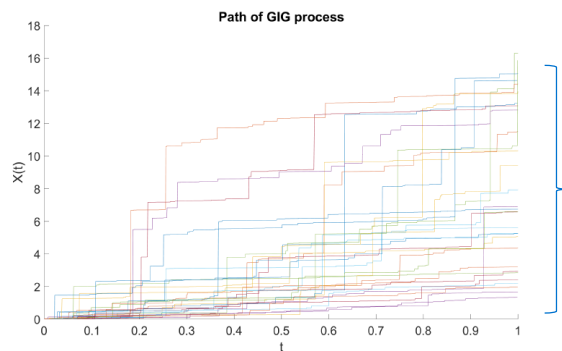
How to sample from $Q(x)$?

- In general $\int_0^\infty Q(x)dx \rightarrow \infty$ so can't sample directly
- The classical method (Fergusson and Klaas 1970's) starts with a uniform Poisson process $\{\Gamma_i\}_{i=1}^\infty$ and finds a function $h(\Gamma_i)$ that converts process to $Q(x)$:

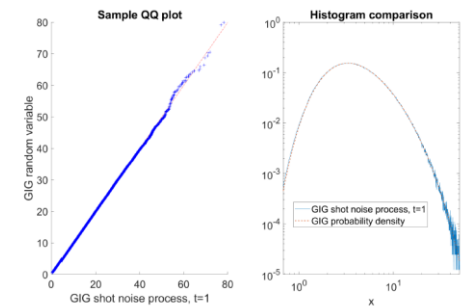


- Turns out the optimal function $h()$ is the inverse of $\int_x^\infty Q(x')dx'$.
- Can think of this as the direct analogue of sampling random variables using the *inverse CDF* method
- Also turns out that $h()$ can't be calculated for most of the processes we wish to use (e.g. the Generalised Inverse Gaussian (GIG))
- A lot of the fun then has been in developing effective alternative strategies for these cases:

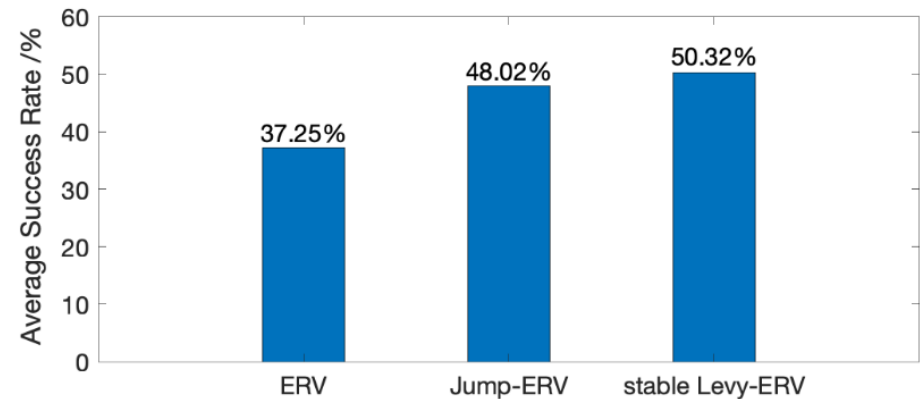
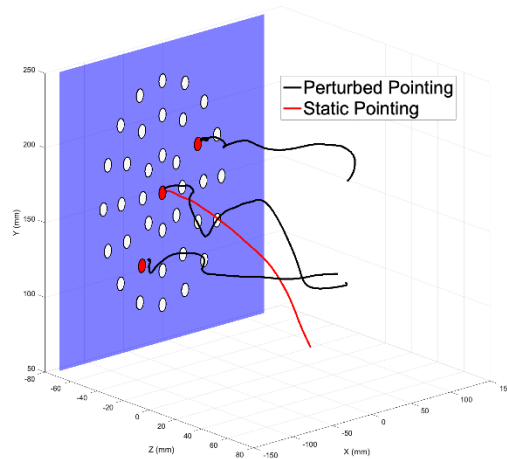
Godsill and Kindap (2021) Point process simulation of generalised inverse Gaussian processes and estimation of the Jaeger integral, Stats and Comp.



Study pdf at t=1:



Example: Intentionality analysis for perturbed pointing task in-vehicle

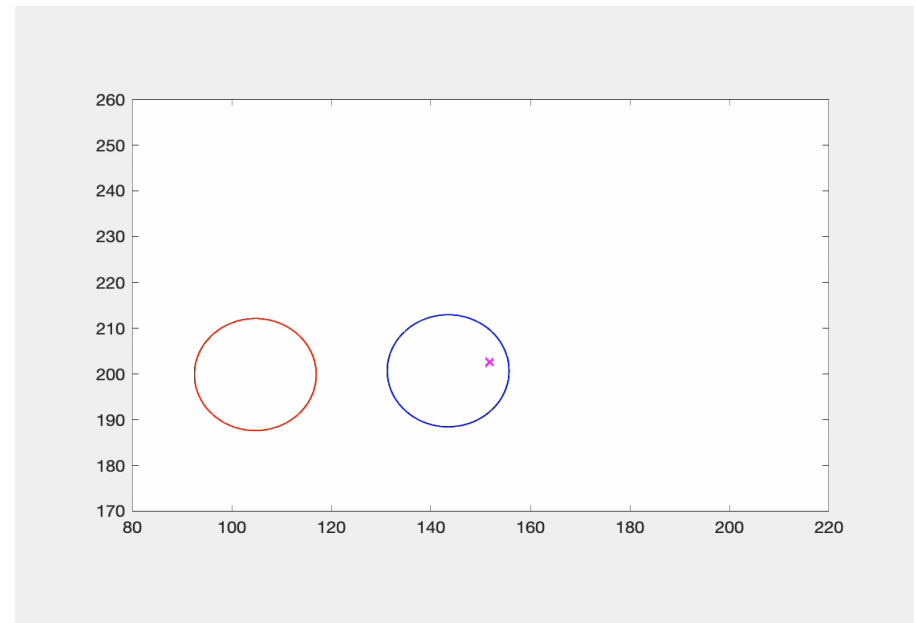
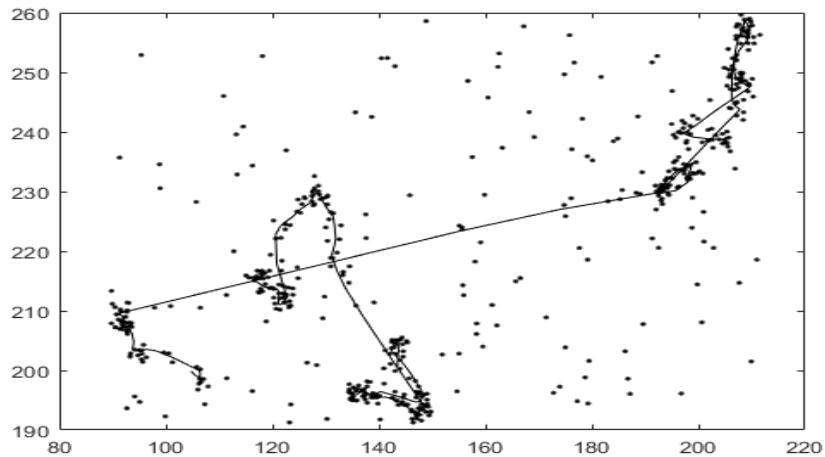


Result for perturbed pointing data from automobile UI systems:

Example: α -stable Lévy State-Space Model for multiple objects in clutter

‘Langevin’ dynamics:

$$d\dot{X}(t) = -\lambda\dot{X}(t)dt + dW(t)$$



Black lines are ground truth; crosses are measurements;
colored lines are estimates plus 95% confidence ellipse

A new idea – the non-Gaussian Process (NGP) model

[Work with Yaman Kindap]

- Here we apply the same Levy process principles to a Gaussian process (GP) model.
- We take a standard GP $\{W(t)\}$ with covariance function

$$\text{cov}(W(t), W(t')) = C(t, t')$$

- We use the same class of ‘subordinator’ jump process $\{X(t)\}$ to modulate locally the covariance function (‘time-change’ operation):

$$\text{cov}(W(t), W(t')) = C(X(t), X(t'))$$

- This allows for non-Gaussian perturbations to the process, but retains once again the structure of a bank of standard GPs, each with a differently modulated covariance function. Examples...

Preliminary examples

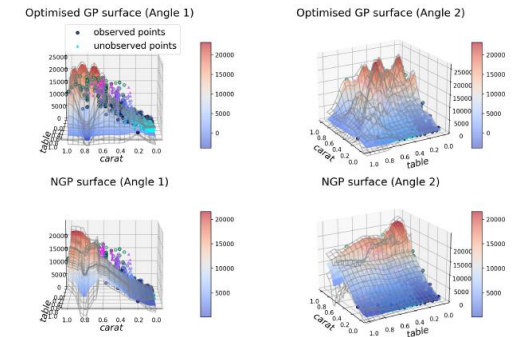
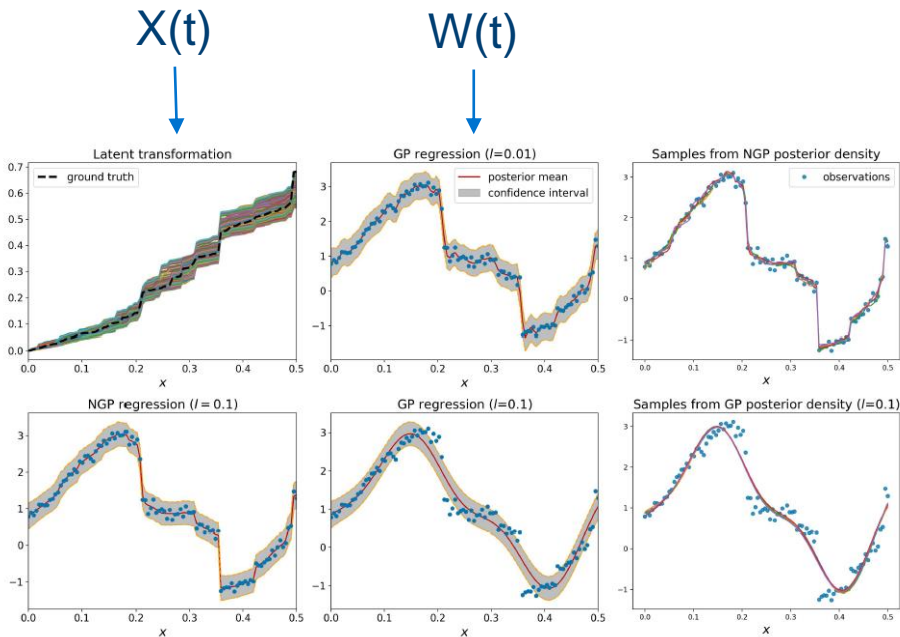


Figure 3: Regression analysis results for NGP and GP models with for the diamond price data set using a TS subordinator. The posterior means are plotted as a surface and the ± 3 standard deviation surface is overlaid on the mean as a wireframe plot.

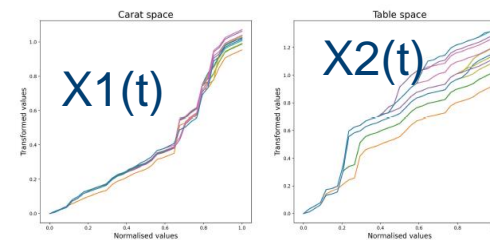


Figure 4: Posterior subordinator samples for a TS subordinator.

Spatial Tracking models using the NGP are currently under development

Poisson point process tracking

Non-homogeneous Poisson process (NHPP) measurement model

- We see clusters of point detections around object locations X_i
- Each object's detections Z are assumed drawn from a non-homogeneous Poisson point process (NHPP) with intensity $\lambda_i(Z|X_i)$ - 0 or more measurements for each object (realistic...)
- Clutter detections also a Poisson process $\lambda_0(Z)$
- The Union property of Poisson processes means that all detections from all K objects and clutter are still Poisson:

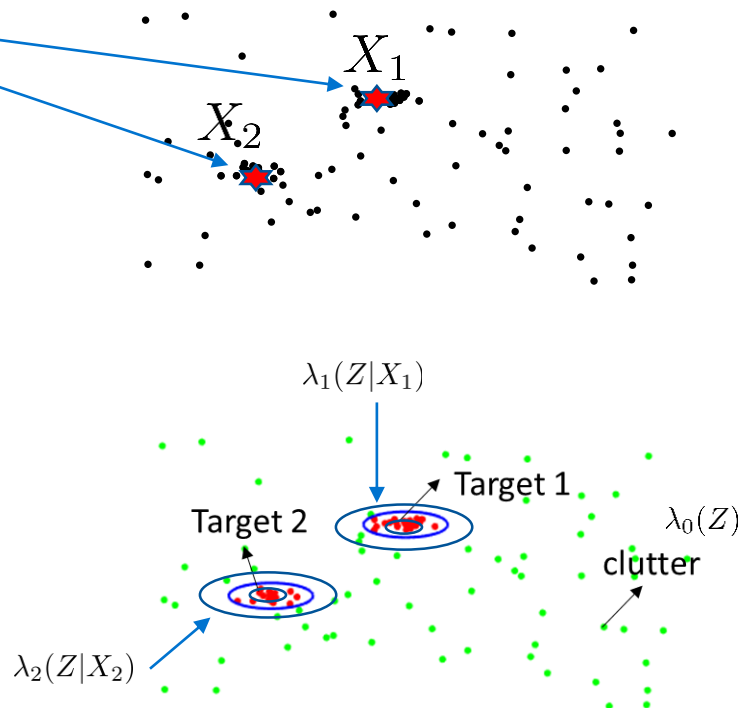
$$\lambda(Z|X) = \sum_{i=0}^K \lambda_i(Z|X_i)$$

- Simple likelihood function and no data association required! This was the insight of:

Gilholm, Godsill, Maskell and Salmond (2005), "Poisson models for extended target and group tracking," in Proc. SPIE

- We now augment the model with data association variables θ - very simple computationally (sequential MCMC) and allows parallelisation. In particular:

$$p(\theta_j = i | Z_j, X) = \frac{\lambda_i(Z_j | X_i)}{\sum_{i=0}^K \lambda_i(Z_j | X_i)}, \quad i = 0, \dots, K$$



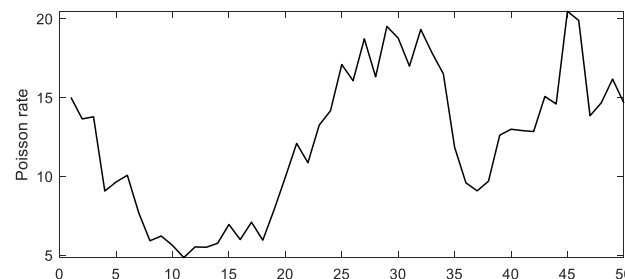
Applications of the NHPP model...

Time-varying Poisson rate estimation for both targets and clutter

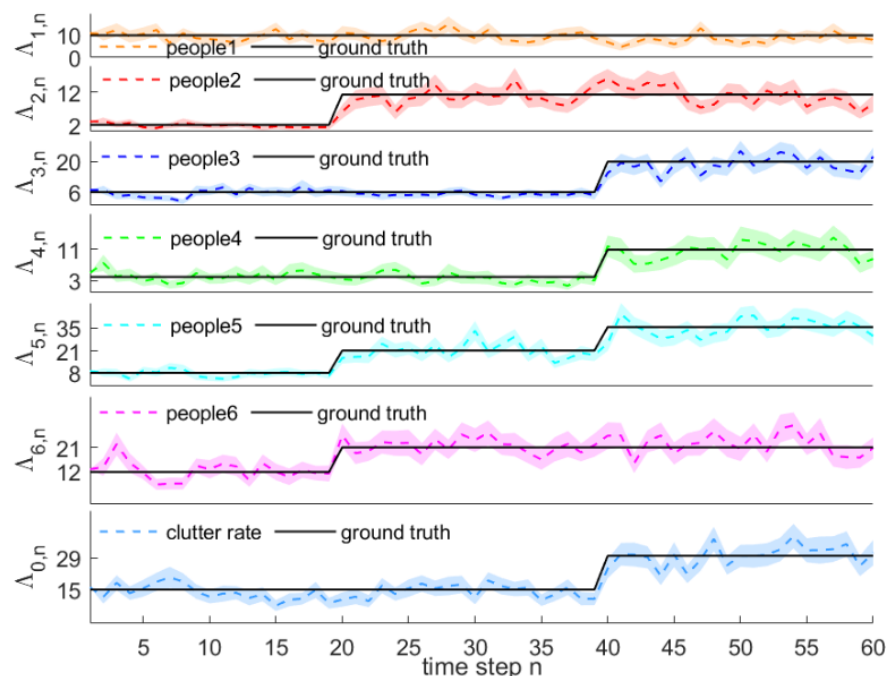
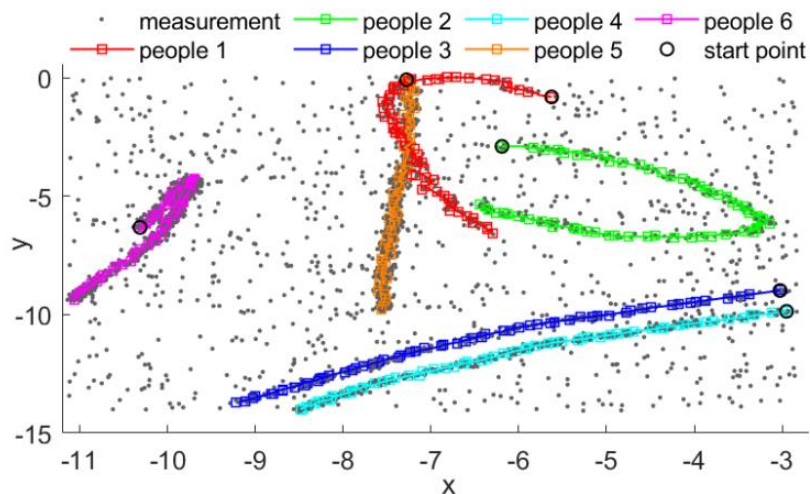
Poisson rates $\Lambda_{i,n}$ are unknown and/or time-varying

Introduce a novel Markov transition density for $\Lambda_{i,n}$ (GIG)

Implement using Sequential Markov chain Monte Carlo (S-MCMC)



MOT Challenge dataset results:



method	OSPA (mean $\pm 1\sigma$)	track loss(%)	CPU time (s)
online PMHT	6.19 \pm 3.09	30.0	8.58
RB-AbNHPP	1.30 \pm 0.08	0.00	0.40

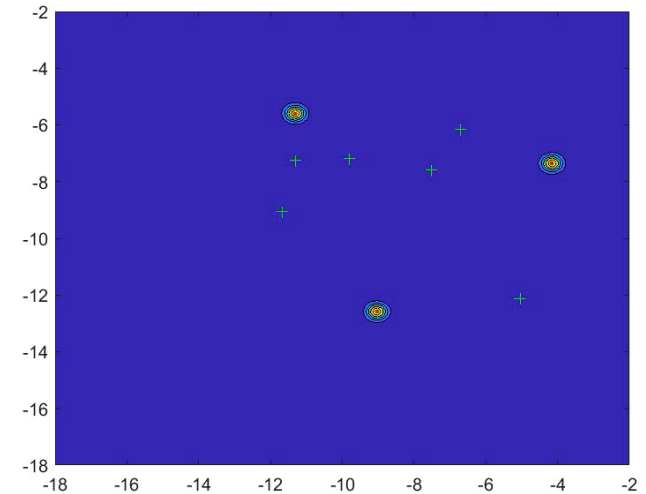
NHPP Application 2:

Tracking of a varying number of objects using reversible jump S-MCMC

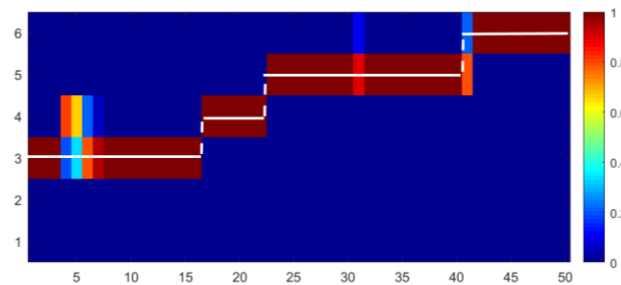
MOT15 benchmark data sets for multiple person tracking



Tracked Posterior density for objects:



Cardinality inference:



[4] <https://motchallenge.net/data/MOT17/>



Variational Bayes Association-based Multi-object Tracker

- Ideal filtering density under NHPP model:

$$p(X_n, \theta_n | Z_{1:n})$$

- Variational Approximation:

$$q_n(X_n, \theta_n) = q_n(X_n)q_n(\theta_n) \approx p(X_n, \theta_n | Z_{1:n})$$

- Scalable and parallelisable since update operations may be further split between the elements of X_n and θ_N
- Tracking performance is comparable to S-MCMC-based trackers, but much faster. Significantly outperforms other fast approximate trackers.
- Distributed versions are currently under development...

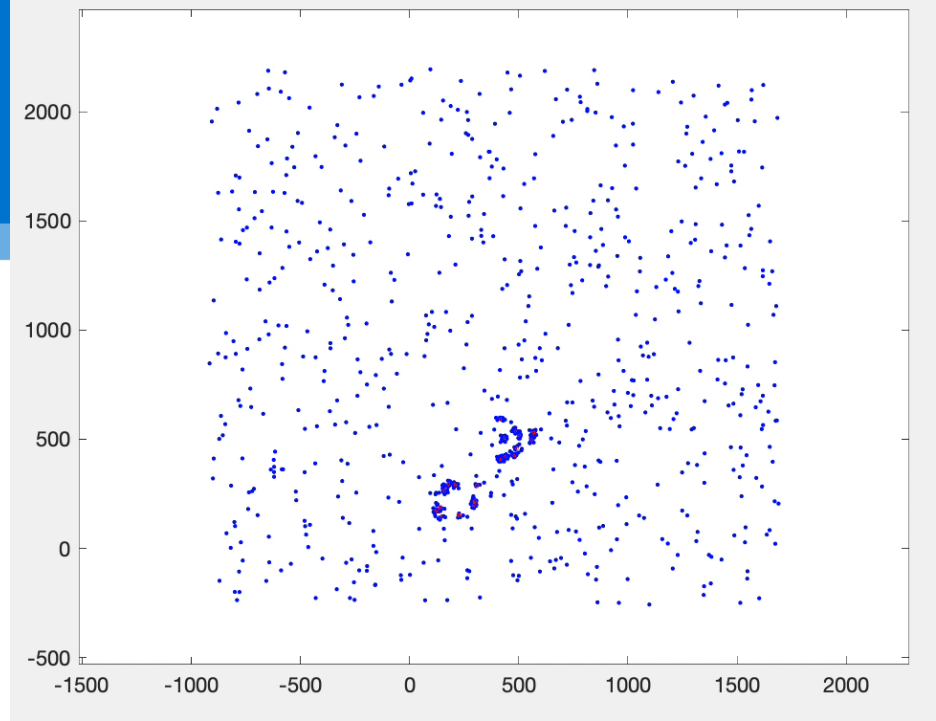
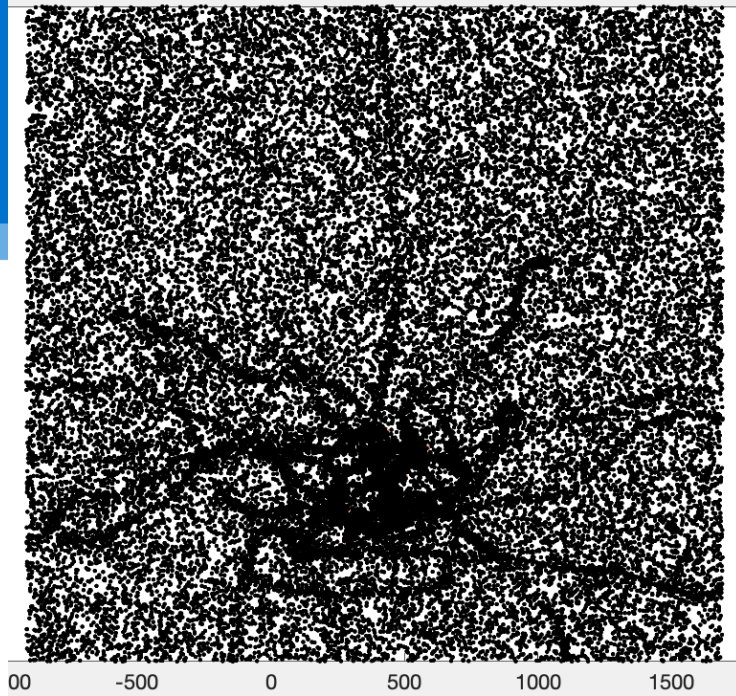
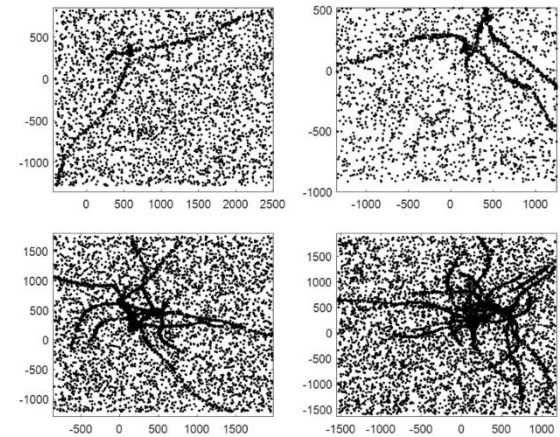


TABLE I: Tracking performance comparisons

dataset	K	RMSE (mean $\pm 1\sigma$) track loss percentage (%) CPU time (s)														
		PF-NHPP		Gibbs-AbNHPP			ET-JPDA		VB-AbNHPP ⁽¹⁾			VB-AbNHPP				
1	2	7.69 \pm 0.64	0.00	4.05	5.50 \pm 0.34	0.00	0.22	7.49 \pm 1.06	4.50	5e-4	5.64 \pm 0.58	0.00	3e-6	5.51 \pm 0.43	0.00	6e-6
2	4	N/A	51.8	15.9	5.63 \pm 0.13	0.00	0.57	8.40 \pm 2.97	8.75	2e-3	5.91 \pm 0.37	0.25	6e-6	5.75 \pm 0.29	0.00	1e-5
3	10	—	—	—	6.06 \pm 0.25	0.10	0.67	9.41 \pm 1.99	16.3	0.01	6.50 \pm 0.50	2.20	8e-6	6.03 \pm 0.32	0.70	2e-5
4	20	—	—	—	6.25 \pm 0.26	0.65	2.31	N/A	22.6	0.03	7.31 \pm 0.70	4.05	1e-5	6.30 \pm 0.40	1.90	3e-5



Conclusions...