Bayesian Nonparametric Modeling and Inference for Multiple Object Tracking

Bahman Moraffah

Arizona State University

School of Electrical, Computer and Energy Engineering

Signal Processing and Adaptive Sensing Laboratory

July 15, 2019

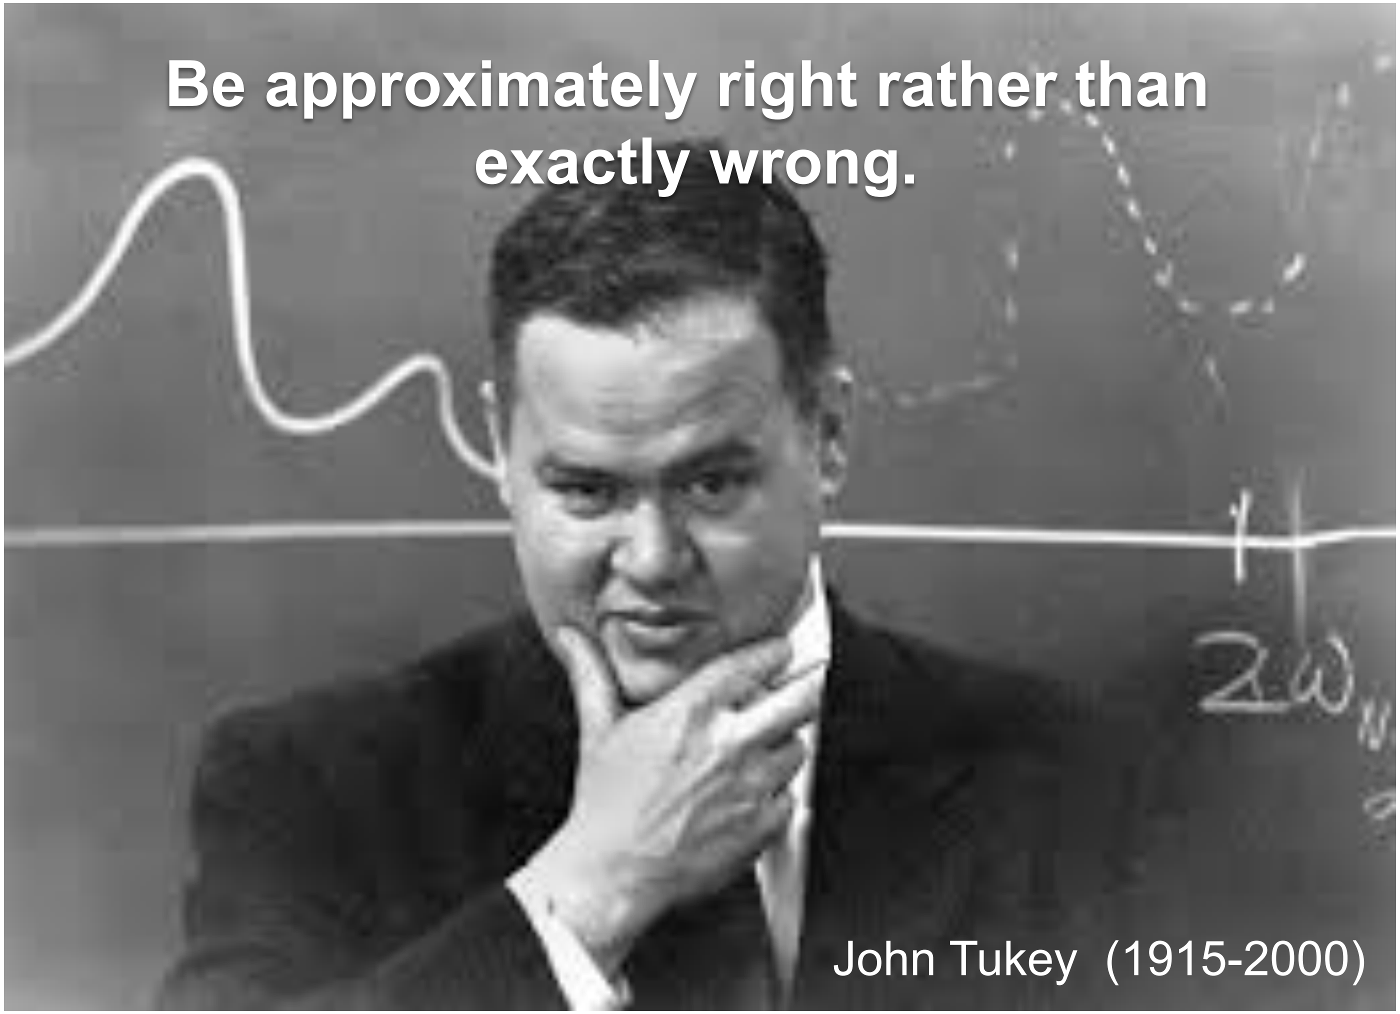
Be approximately right rather than exactly wrong.
John Tukey (1915-2000)

Multi-Object Tracking (MOT)

Dynamic multi-object tracking problem:
Jointly estimate the number of objects and the states using received data

- **Multiple objects: unknown time-varying number; leave, enter or stay in scene at any time step, unknown identity/label**
- **Each survived object transitions to the next time according to a probability transition kernel**
- **New objects may join the scene**
- **Observations: A set of observations collected from the sensor**

previous time step (k-1) **time step _k_**

Challenges:

- Track unknown time-varying number of objects
- Unknown state identity
- Robustly associate objects at each time step
- Uncertainty on parameters due to multiple environmental conditions: high noise, interference, or clutter

Vignette of the Contributions

- Dependent Dirichlet process (DDP) prior modeling **over time-evolving object state distribution for MOT problem**
  - **Identity learning for multiple object tracking**
- Dependent Pitman-Yor (DPY) process prior **to incorporate learning algorithm over time-evolving object state distribution based on measurements to fully capture dependence among the states**
  - **More available clusters to capture full dependency & more likely to have less popular clusters**

- **Multiple Object Tracking through Infinite Random Trees**
  - **Tracking multiple objects by defining a prior over infinite random trees**
    - A nonparametric modeling based on diffusion processes
- **Multimodal Dependent Measurements**
  - **Multimodal dependent measurements and single object tracking**
    - Use the information provided by the multiple sensor to track more accurately
  - **Multimodal dependent measurements and multiple objects tracking**
    - Generalize the problem to a multi-object multimodal dependent measurements

Related Work

1. Bayesian methods for a single object tracking
   - Kalman filter, particle filter, interactive multiple modal for maneuvering, the nearest neighbor method
2. Random finite set theory for multiple object tracking
   - Multiple hypothesis testing, probability hypothesis density filter, labeled multi-Bernoulli (LMB)
3. Deep learning models for multiple object tracking
4. Multimodal dependent measurements
   - Exponentially embedded families for multimodal sensor, target tracking using multi-modal sensing with waveform configuration, a parametric classification rule based on the exponentially embedded family
5. Bayesian nonparametric modeling for tracking
   - Evolutionary clustering, hierarchical Dirichlet process for maneuvering, Dirichlet process for linear dynamic system

Dependent Dirichlet Process Modeling

- Introducing a well defined dependent Dirichlet process:
  - Captures the survival, birth, and death
  - Dependent structure to update object cardinality
  - Conditional distribution given the immediate past is a DP
  - Easy models to do inference through MCMC and VB methods where do not depend on the initial values
  - Shown this prior leads to a consistent posterior distribution
  - Contraction rate matches the optimal minimax rate

Simulations: DDP-EEM Modeling

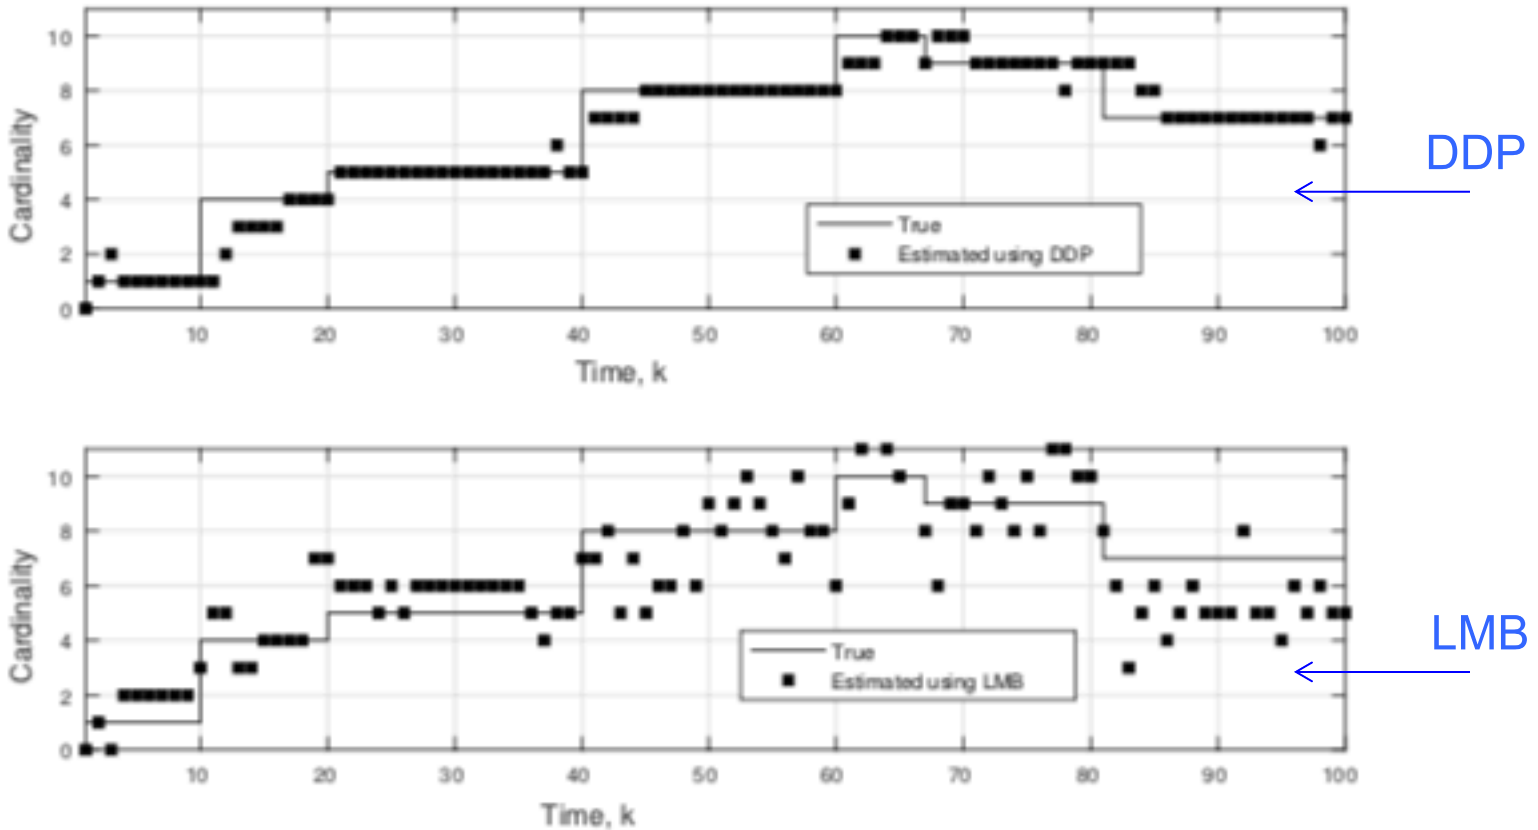

Cardinality estimation using the DDP prior based and labeled multi-Bernoulli filtering

Simulations: DDP-EEM Modeling

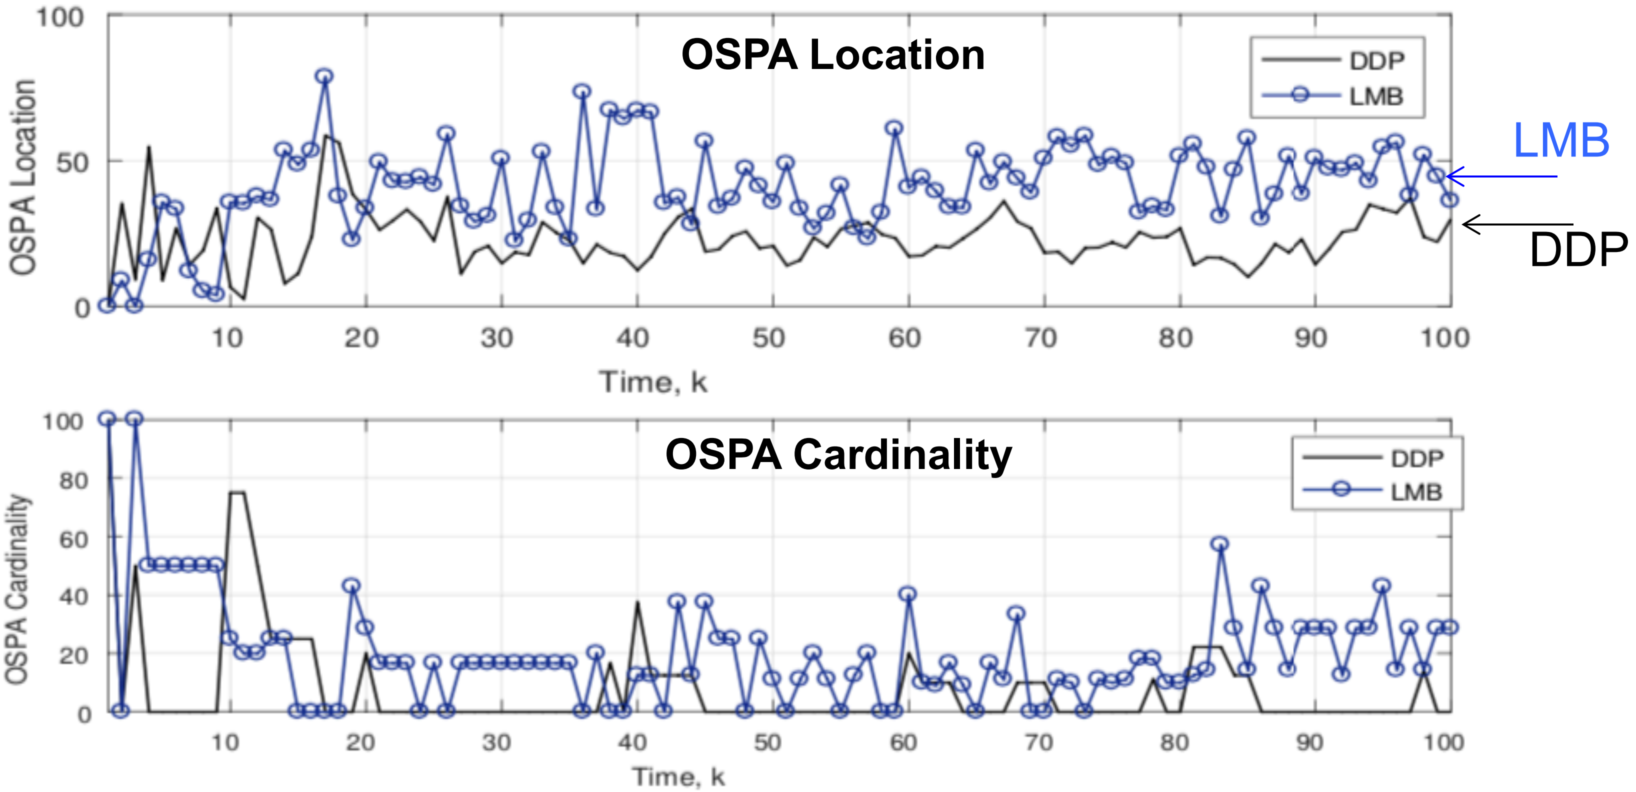

Performance comparison: OSPA comparison for 10000 MCMC simulations and order p =1, cut-off c =100

Dependent Pitman-Yor Modeling

- Introducing a well defined dependent Pitman-Yor process:
  - Fully captures the survival, birth, and death
  - Dependent structure to update object cardinality
  - Conditional distribution given the immediate past is a PY
  - Introduced an easy inferential models based on MCMC and VB
  - Compared to dependent Dirichlet process (DDP): more available clusters to capture full dependency & likely to have less popular clusters
    - DDP: expected number of clusters $\alpha \log(N)$
    - DPY: expected number of clusters αN^d

concentration → α; discount → d; # of objects → N

Simulations: DDP vs DPY

Performance comparison between DDP and DPY based prior modeling

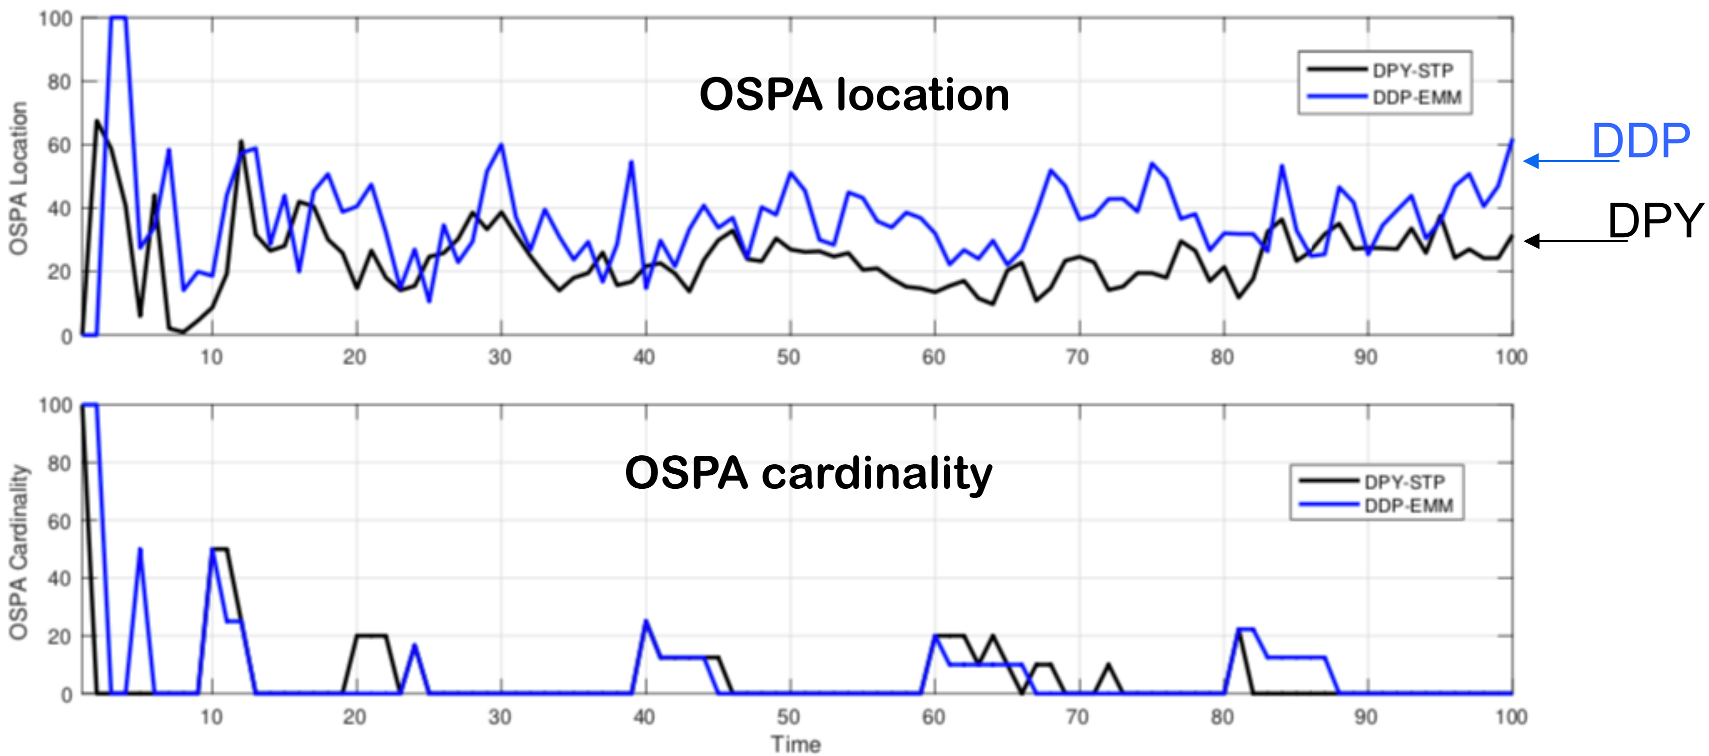

Depend Dirichlet process (DDP) vs dependent Pitman-Yor process (DPY) for 10000 MCMC simulations for order p = 1 and cut-off c = 100

Dynamic System Modeling

Motion Model:

- Unknown state vector of ℓ object: $\mathbf{x}_{\ell,k}, \quad \ell = 1, \ldots, N_k$
- If the object were present at time (k-1), then:

$$\mathbf{x}_{\ell,k} = f_k(\mathbf{x}_{\ell,k-1}) + \mathbf{u}_{\ell,k-1}$$

Object cardinality

Transition function

Modeling error

- This model implied that each existing object $\mathbf{x}_{\ell,k-1}$ stays in the scene with probability $P_{\ell,k|k-1}$ and transitions with probability transition kernel $Q_\theta(\mathbf{x}_{\ell,k-1}, \mathbf{x}_{\ell,k})$

Measurement Model:

- Measurement vector: $\mathbf{z}_{m,k}, \quad m = 1, \ldots, M_k$
- If mth measurement were originated from ℓth object, then

$$\mathbf{z}_{m,k} = h_k(\mathbf{x}_{\ell,k}) + \mathbf{w}_k$$

Relationship between measurement & state

Measurement noise

- This model leads to the likelihood $p(\mathbf{z}_{m,k}|\mathbf{x}_{\ell,k}, \Theta)$

Goal: Find the posterior distribution $p(\mathbf{x}_{\ell,k}|\mathbf{z}_{m,k}, \Theta)$

A Nonparametric Prior on Random Infinite Trees in Multiple Object Tracking

ASU

Dependent Poisson Diffusion Process (D-PoDP) & random Tress

- **Modeling uncertainty over trees; path/branch generated by diffusion process (generate samples using Brownian motion at)** $k=0$ **This process is exchangeable**
- **Branching probability: probability of selecting a branch vs diverging, depends on number of samples previously followed same branch**
- **Dependent as prior can incorporate time-dependent learned information**
- **Probability transition kernel** $Q_{\theta_k}(\mathbf{x}_{\ell,k-1},\mathbf{x}_{\ell,k})$ **with unknown parameters** θ_k
  - Use a dependent diffusion process on a tree as prior on θ_k
  - Tree leaf/node: object state, branch: cluster of states in a hierarchy
  - Find trajectory of each object by tracing path on tree
  - Predict and update number of objects at each time

Dependent Poisson Diffusion Process (D-PoDP)

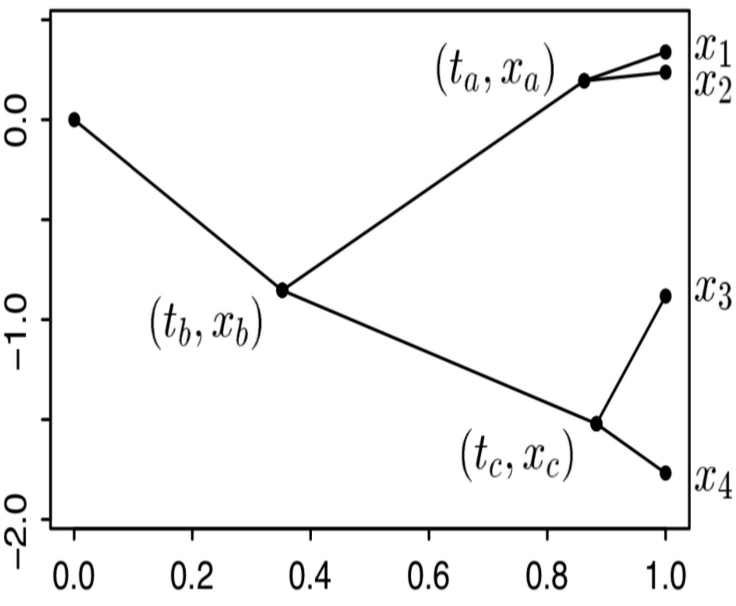

For instance, at time k =0, generate 4 states according to a diffusion process with divergence points a,b, and c and its underlying tree

(Moraffah & Papandreou 2019)

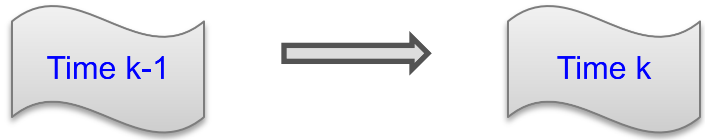

N_k objects may enter, leave or remain in the scene

- **Assign probability to survived branch** a

Discount parameter

$$p_a \propto |S_{a,k-1}| + |S_{a,k|k-1}| - \gamma$$

# of objects with common branch node a

- **For new objects, assign probability to new branch node** δ

$$p_\delta \propto \zeta - |V_{B,k|k-1}|\gamma$$

Hyperparameter

# of survived branch node

D-PoDP: Density Estimation

At time k,

- **For each** $\theta_{\ell,k|k-1} \in S_{a,k|k-1}$ **, draw**

Poisson distribution

$\alpha = \mu(\mathcal{X})$

$$\tilde{N}_{\ell,k|k-1} \sim \mathrm{Po}(\frac{p_a \alpha}{2|S_{a,k|k-1}})$$

And generate $\tilde{N}_{\ell,k|k-1}$ **atoms given** $\theta_{\ell,k|k-1}$ **using the diffusion process**

- **For** δ **, draw**

$$\tilde{N}_{\delta,k|k-1} \sim \mathrm{Po}(\frac{p_a \alpha}{2})$$

And generate $\tilde{N}_{\delta,k|k-1}$ **atoms from the base distribution**

- **Draw** $\mathbf{x}_{\ell,k}|\theta_{\ell,k} \sim G(\cdot|\theta_{\ell,k})$

From the physical model

- **Set** $\tilde{N}_k = \sum_{\ell} \tilde{N}_{\ell,k|k-1}$

Dependent Mixture Model and Inference

- **Use constructed prior as mixing distribution to infer measurement distributions**
  - **Select parameter $\theta_{\ell,k}$ at time k with probability π_ℓ proportional to the summation of number of measurements that already selected same parameter and number of object with the shared branch, i.e.,**

Set of all nodes at time k

$$\pi_\ell \propto n_{\ell,k} + |S_{a,k-1}| \quad \text{for} \quad \theta_{\ell,k-1} \in S_{a,k-1}, \theta_{\ell,k} \in \tilde{V}_k$$

  - **New parameters are selected with probability proportional to ξ**
- **Dependent Mixture model**

$$\mathbf{z_{m,k}} | \mathbf{x}_{\ell,\mathbf{k}}, \theta_{\ell,\mathbf{k}}, \pi_\ell \sim \mathbf{F}(\cdot | \mathbf{x}_{\ell,\mathbf{k}}, \theta_{\ell,\mathbf{k}})$$

Comes from the physical model

- **Use a MCMC sampler to do inference**

Track five objects for time-varying objects

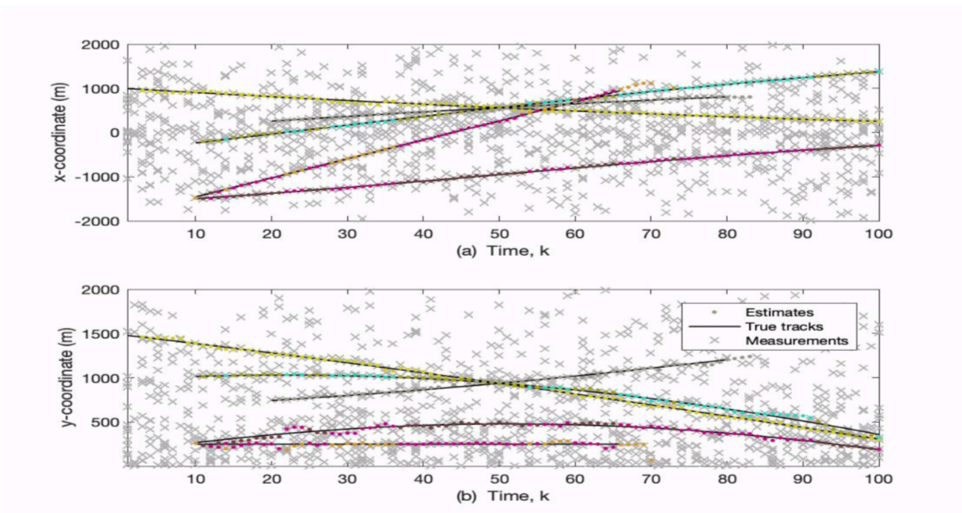

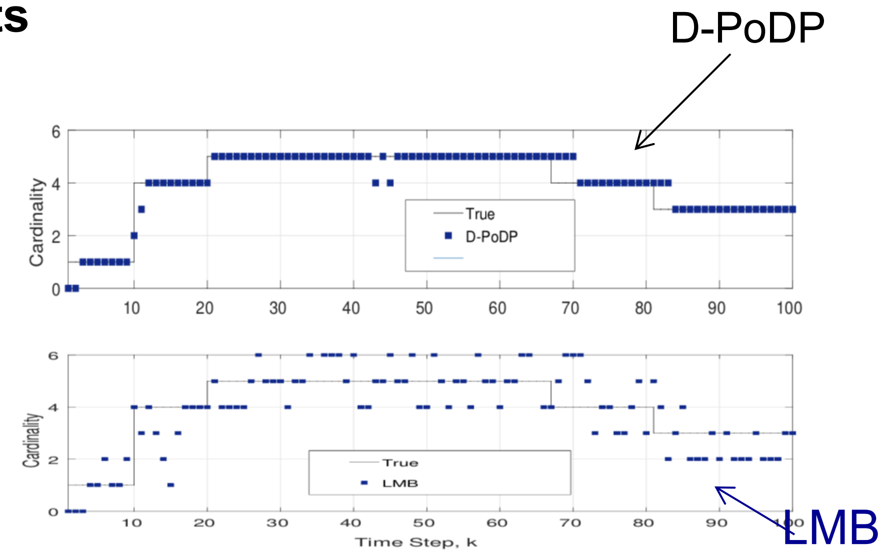

True and learned object cardinality as a function of time step k for 5 objects

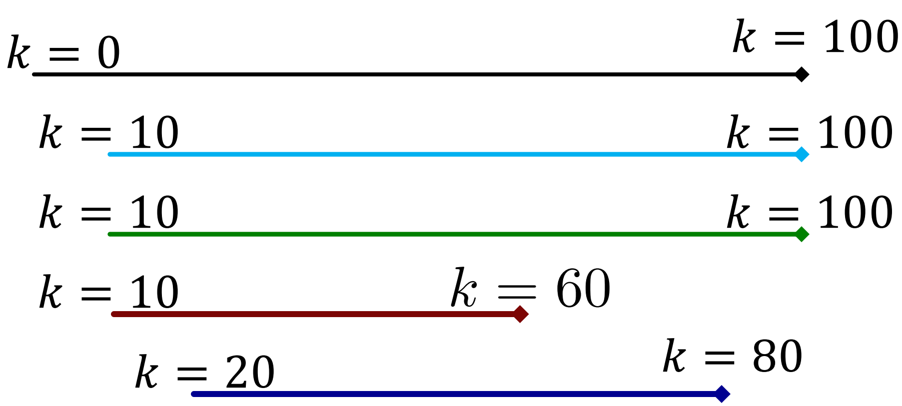

I. Simulations: Comparison to LMB

Performance comparison between D-PoDP and LMB trackers

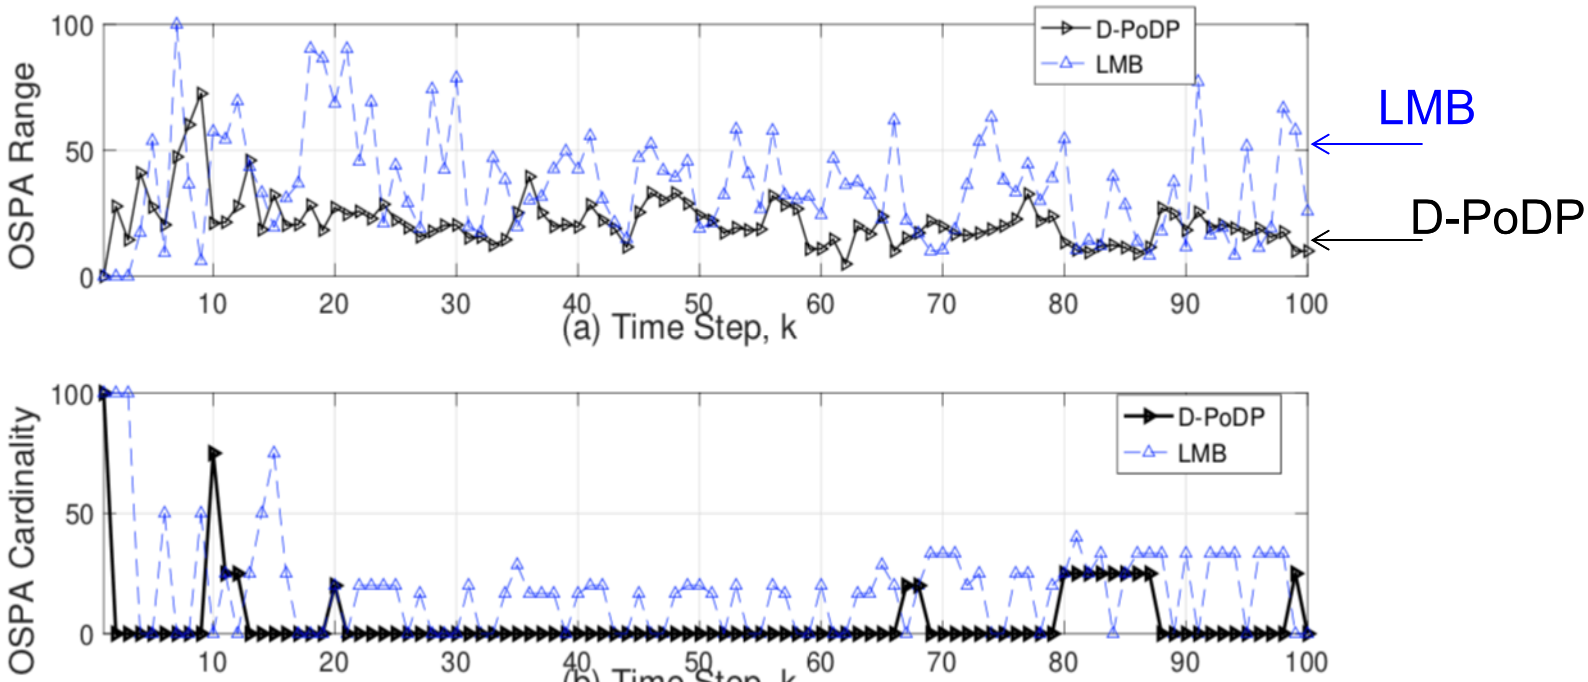

OSPA comparison for range (top) and cardinality (bottom) over 10000 MCMC simulations and order p =1 and cut-off c=100

II. Simulations: Comparison to Dependent Pitman-Yor

Performance comparison for D-PoDP and proposed Dependent Pitman-Yor process

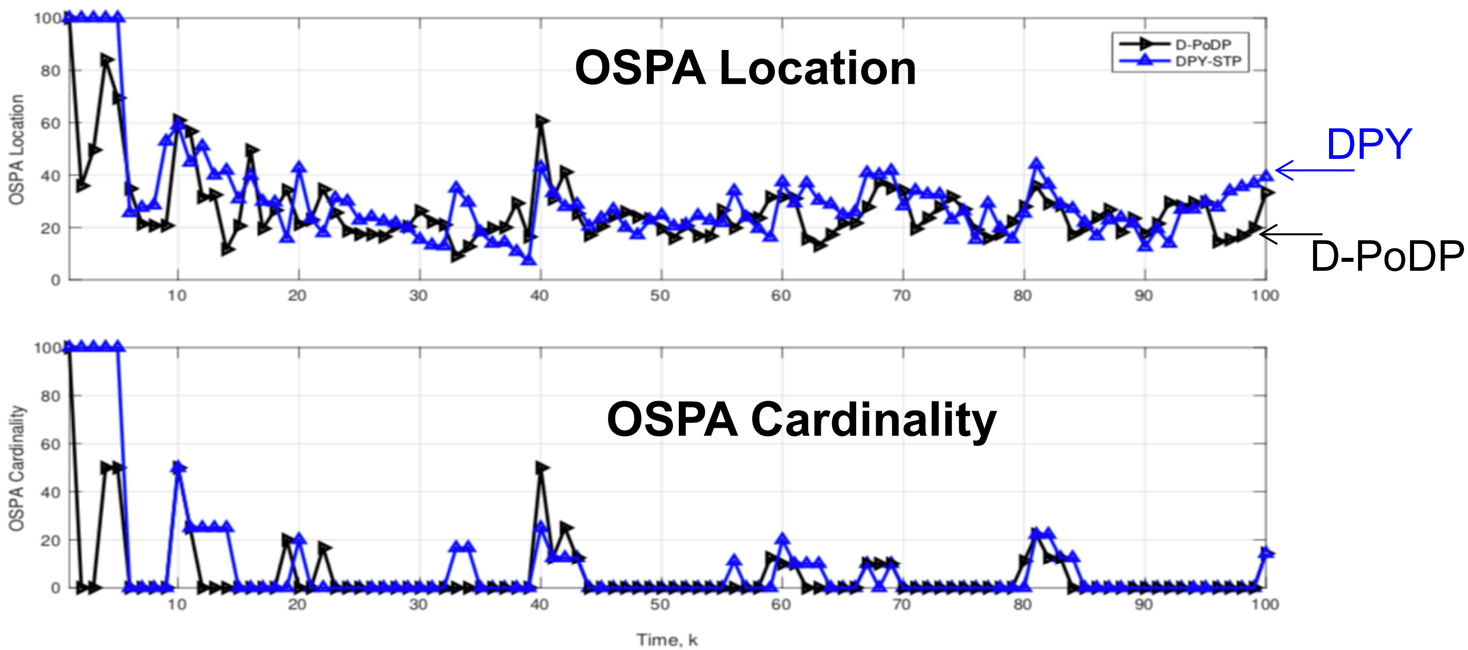

OSPA comparison between D-PoDP and DPY method for 10000 MCMC simulations and for order p = 1 and cut-off c = 100

- **Similar** performance, however, D-PoDP is much more efficient and simpler to implement and estimate the object trajectory

Multimodal Dependent Measurements

Integration of Dependent Observations from Multiple Sensors to Track a Single Object

- I-band radar: angular accuracy
- K-band radar: short ranges
- electro-optical (EO) infrared camera: target identification and observation

Multimodal framework allows for integration of complementary information in analyzing a scene

Challenges:

- **Time-varying number of observations (unknown at each time step)**
- **Observations are unordered: no measurement-to-model association**
- **Multiple environmental conditions: high noise levels, clutter, interference**
- **How to group dependent measurements so that :**
  a. **Dependency among measurements is captured**
  b. **Sensor information is preserved**

Tracking Formulation using Measurements from Multiple Sensors

- **Unknown object state vector:**

$$\mathbf{x}_k = \mathbf{f}_k(\mathbf{x}_{k-1}) + \mathbf{u}_{k-1}$$

Possibly a nonlinear transition function — modeling error

- **Measurement model for mth sensor**

$$Z_{m,k} = h_m(x_k) + w_{m,k}$$

m = 1,…, M (M = # of sensors)

mth sensor measurement noise

Dependent measurements
Unordered measurements and correspond to different model
Object association

Multimodal sensing to improve learning algorithms → **Hierarchical DP**

I. HDP-DM: HDP Prior

Hierarchical Dirichlet process to group measurements and improve the performance.
Propose "Hierarchical Dirichlet Process for Dependent Measurements (HDP-DM)" modeling

- **With an HDP prior on the parameters of the measurements collected from the sensors, the distribution of the measurements can be modeled as**

$$G_0 \sim \mathrm{DP}(\eta, H)$$

$$G_m \mid G_0 \sim \mathrm{DP}(\gamma, G_0),\ \ m = 1, \ldots M$$

$$\phi_{m,k}^{(i)} \mid G_m \sim G_m,\ \ i = 1, \ldots, L_m$$

$$\mathbf{z}_{m,k}^{(i)} \mid \phi_{m,k}^{(i)} \sim F\left(\phi_{m,k}^{(i)}\right),$$

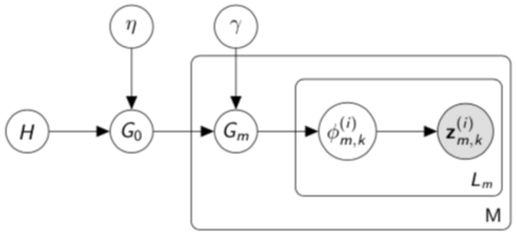

- **This method clusters measurements that are collected by each sensor and estimates joint density of dependent measurements.**

(Moraffah & Papandreou 2019)

II. HDP-DM: Hypothesis Testing

- Hypothesis Testing for Object Detection
  - The detection test-statistic is based on the binary hypothesis

$$\mathcal{H}_0 : \mathbf{Z}_{m,k} = \mathbf{w}_{m,k}$$
$$\mathcal{H}_1 : \mathbf{Z}_{m,k} = h_m(\mathbf{x}_k) + \mathbf{w}_{m,k}$$

  - An object is detected using the measurements of the mth sensor if the Neyman-Pearson test statistic exceeds the threshold

$$\mathcal{T}_m\Big(\mathbf{Z}_{m,k}, \phi_{m,k}; \mathbf{x}_k\Big) = \frac{p\Big(\mathbf{Z}_{m,k} \mid \mathbf{x}_k; \mathcal{H}_1\Big)}{p\Big(\mathbf{Z}_{m,k}; \mathcal{H}_0\Big)}$$

If measurements from the same sensor are assumed independent, the likelihood ratio simplifies to a product of individual likelihoods that still preserve dependency among measurements from different sensors

III. HDP-DM: Target Tracking Method

- **Bayesian Single Object Tracking Method**
  - **The estimated state is given by the posterior mean**

$$\hat{\mathbf{x}}_k = \mathbb{E}[p(\mathbf{x_k}|\mathcal{Z}_\mathbf{k})]$$

  - **The tail recursive function for the prediction is given by**

$$p(\mathbf{x}_k|\mathcal{Z}_1,\ldots,\mathcal{Z}_{k-1}) = \int Q_{\boldsymbol{\theta}}\Big(\mathbf{x}_{k-1},\mathbf{x}_k\Big)p(\mathbf{x}_{k-1}|\mathcal{Z}_1,\ldots,\mathcal{Z}_{k-1})d\mathbf{x}_{k-1}$$

↑ The transition kernel originated from the physical model

III. HDP-DM: Target Tracking Method

- **At time step k, the Bayesian recursion is given by**

$$p(\mathbf{x}_k|\mathcal{Z}_1,\ \ldots,\ \mathcal{Z}_k) \propto p(\mathcal{Z}_k|\mathbf{x}_k)\, p(\mathbf{x}_k|\mathcal{Z}_1, \ldots, \mathcal{Z}_{k-1})$$

Prediction Equation

Measurements collected by M sensors modeled through HDP mixture

- **To compute this probability, we use the tail recursive equation and the density of** $\mathcal{Z}_k$ ***estimated using the HDP mixture obtained as***

$$p(\mathcal{Z}_{m,k}|\mathbf{x}_k) = \sum_{j=1}^{\infty} \pi_{m,j} f(\cdot|\theta_{j,k})$$

Due to physical model

$$\pi_m = (\pi_{m,1}, \pi_{m,2}, \cdots) \text{ where } \pi_m \sim \mathrm{DP}(\eta, \mathrm{GEM}(\gamma))$$

Simulations: HDP-DM

- **We simulate dependent measurements obtained from a multimodal sensing system with radio frequency (RF) and electro-optical (EO) sensors.**

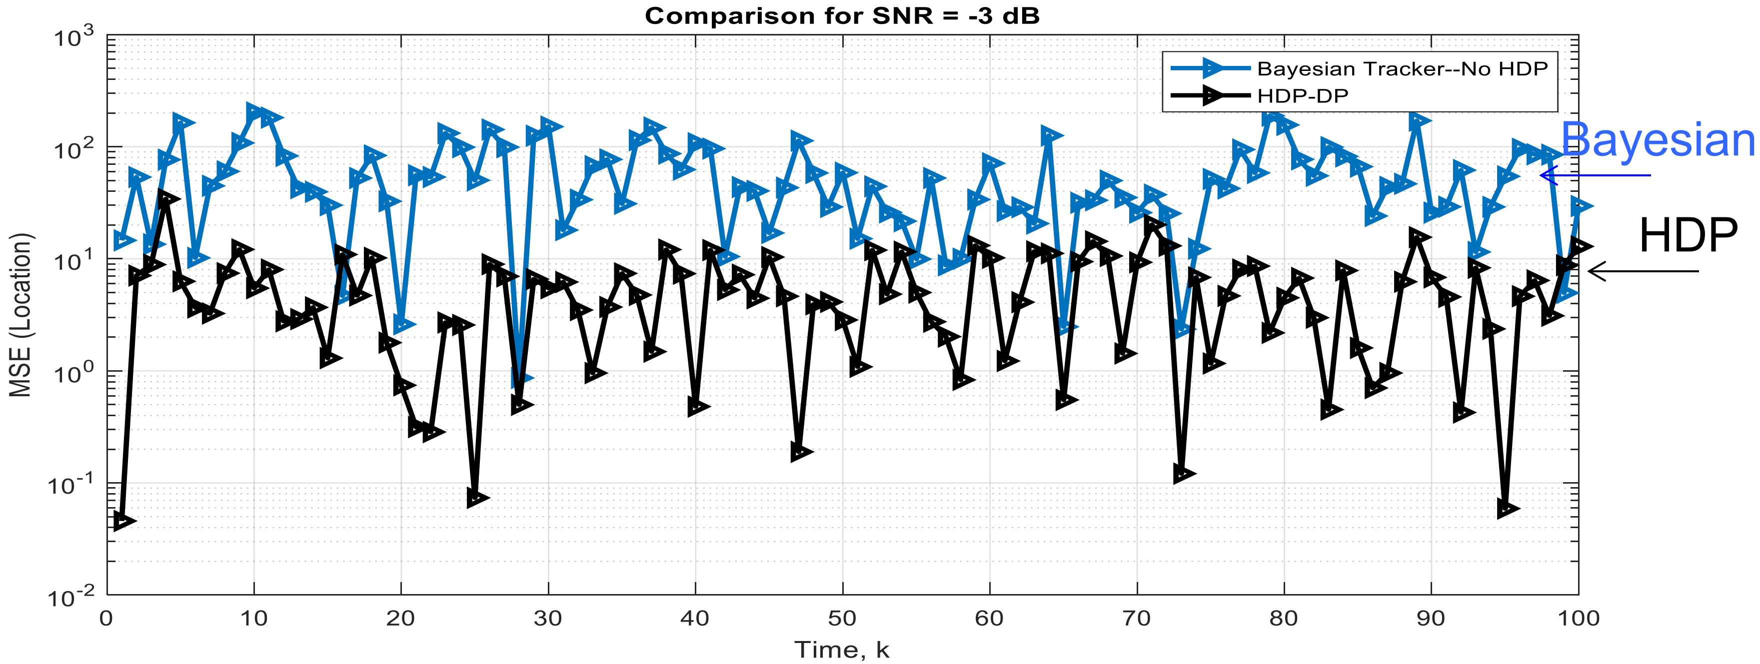

HDP-DM performance comparison to Bayesian tracker with no utilization of dependency.

Simulations: HDP-DM

- **HDP-DM performance as a function of SNR**

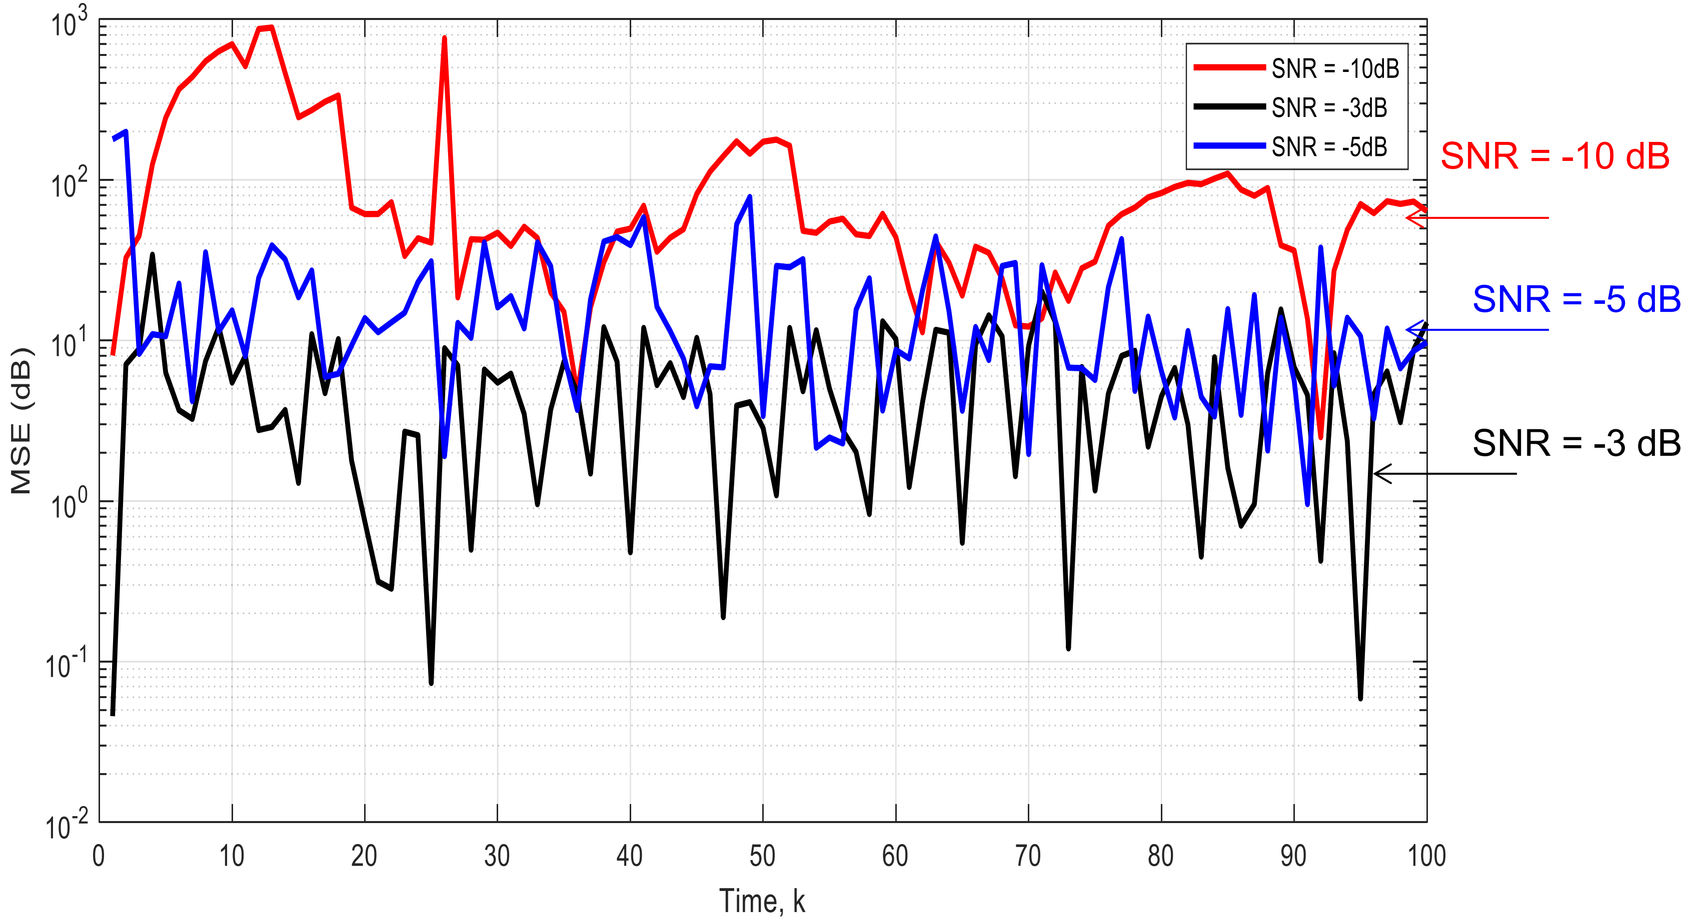

Integration of Dependent Observations from Multiple Sensors to Track Multiple Objects

Objective: Multiple Object Tracking with Dependent Measurements from Multiple Sensors

- **Accurately estimate the time evolving object trajectory as well as object cardinality ⟹ Use the dependency among the measurements to estimate more accurately**

Challenges:

- ❑ Robustly associate each object
- ❑ Jointly estimate the object cardinality as well the object trajectory
- ❑ Object identity and object cardinality at each time are dependent
- ❑ Dependency among measurements such that the sensor information is preserved
- ❑ Inference

⟹

Solution:

- ✓ Group data in a hierarchical Manner
- ✓ Dependent Modeling such as DDP-EMM
- ✓ Hierarchical Dirichlet process mixture modeling

(Moraffah & Papandreou 2019)

Dependent Dirichlet Process State Transitioning Prior

Prior Construction to capture survival, appearance, disappearance so that the conditional distribution is a Dirichlet process

Case 1: The ℓ th object belongs to one of the survived and transitioned clusters from time $(k-1)$ and occupied at least by one of the previous $\ell-1$ objects. The object selects one of these clusters with probability:

$$\Pi_1(\text{Select } j\text{th cluster}|\theta_{1,k}^{\ell-1}) \propto \left[V^*_{k|k-1}\right]_j + \left[V_k\right]_j$$

$$\theta_{1,k}^{\ell-1} = \{\theta_{1,k}, \ldots, \theta_{\ell-1,k}\}$$

Size of jth cluster at time k

Size of the jth cluster after transitioning

Where the normalizing constant equals $\sum_i [V^*_{k|k-1}]_i + \sum_i [V_k]_i + \alpha$ for concentration parameter α

Prior Construction

- **Case 2: The ℓth object belongs to one of the survived and transitioned clusters from time $(k-1)$ but this cluster has not yet been occupied by any one the first $\ell - 1$ objects. The object selects such a cluster with probability:**

$$\Pi_2(\text{Select } j\text{th cluster not chosen yet}|\theta_{1,k}^{\ell-1}) \propto \left[V_{k|k-1}^{*}\right]_j$$

Size of the jth cluster after transitioning

- **Case 3: The object does not belong to any of the existing clusters, thus a new cluster parameter is with probability:**

$$\Pi_3(\text{New cluster}) \propto \alpha$$

Concentration parameter

Bayesian Inference

- **Given the configurations at time $(k-1)$, the conditional distribution is Dirichlet process**
- **Under mild conditions the state distribution follows:**

$$p(\mathrm{x}_{\ell,k}|\mathrm{x}_{1,k},\ldots,\mathrm{x}_{\ell-1,k},\mathrm{X}_{k|k-1},\Theta^{*}_{k|k-1},\Theta_k) = \begin{cases} \mathbb{Q}_{\underline{\theta}}(\mathrm{x}_{\ell,k-1},\mathrm{x}_{\ell,k})f(\mathrm{x}_{\ell,k}|\theta^{\star}_{\ell,k}) & \text{Case1} \\ \mathbb{Q}_{\underline{\theta}}(\mathrm{x}_{\ell,k-1},\mathrm{x}_{\ell,k})\zeta(\theta^{\star}_{\ell,k-1},\theta^{\star}_{\ell,k})f(\mathrm{x}_{\ell,k}|\theta^{\star}_{\ell}(k)) & \text{Case2} \\ \int_{\theta} f(\mathrm{x}_{\ell,k}|\theta)dH(\theta) & \text{Case3} \end{cases}$$

Transition probability kernel

Transition kernel for parameters

Base distribution

$f(\cdot|\theta)$ **is derived from the physical based model**

- **Use a Hierarchical Dirichlet mixture modeling, group measurements upon receiving and compute the likelihood based on the physical model, compute the posterior distribution** $\mathbf{x}_{\ell,k}|\mathbf{z}_{l,k},\theta^{\star}_{\ell,k}$ **using a MCMC method (Gibbs sampling)**

Simulations: Tracking

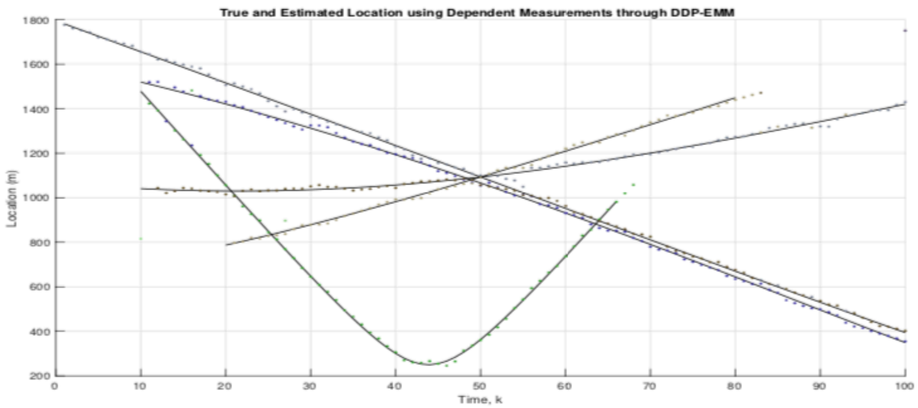

Location estimation in the presence of multiple dependent measurements for 5 objects.

Cardinality estimation comparison in the presence of multiple dependent measurements.

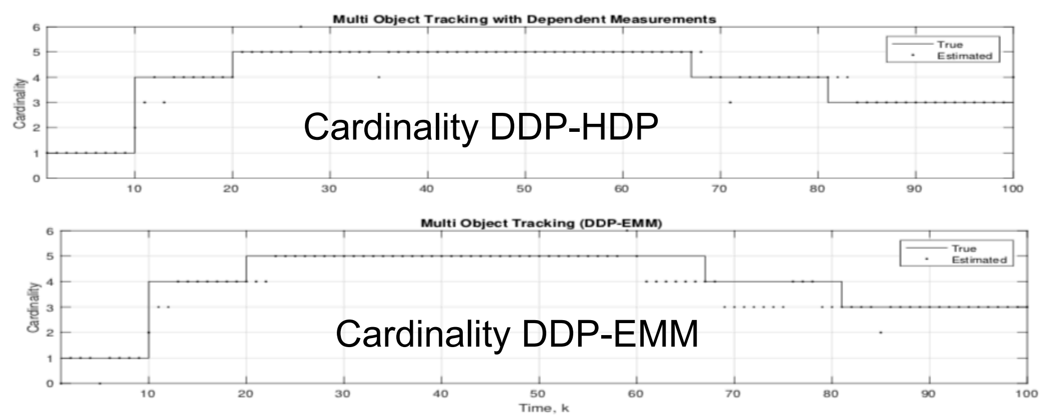

Simulations: OSPA Performance

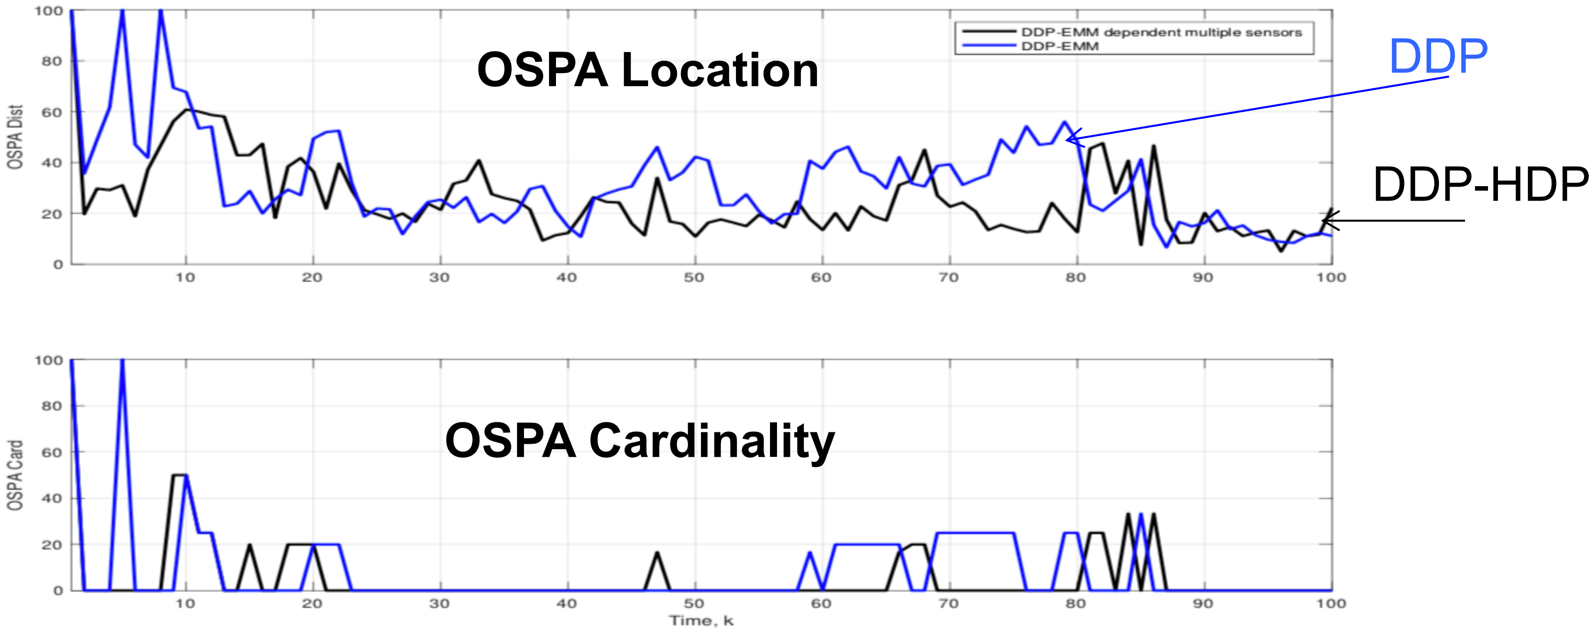

OSPA Comparison for Multi-Target Tracking with and without Using the Dependent Measurements for order p = 1 and cut-off c = 100 for 10000 MCMC simulations

Conclusion

- **Nonparametric priors to fully capture state dependency to robustly and efficiently track labels, cardinality, and trajectory of multiple objects**
  - **A model with the DP as the conditional distribution**
    - Exploit dependent DP to model dependencies in state prior
  - **A model with the PY as the conditional distribution**
    - Follows power law and hence higher probability for smaller cluster
  - **A model based on random infinite trees that follows power law**
    - Dependent Poisson diffusion process as prior on evolving trees
    - State estimated by selecting path connected to each leaf
  - **A nonparametric modeling un multimodal scenarios to capture measurement dependency as well as state dependency**
    - HDP models dependency, model association, and time-varying cardinality of the measurements provided by each sensor
- **These models are all distribution free (no parametric assumption required)**
- **Low computational cost for these modeling (MCMC/VB methods)**

Selected Publications

- **Bahman Moraffah,** Antonia Papandreou-Suppappola, "**Dependent Dirichlet Process Modeling and Identity Learning for Multiple Object Tracking**", *52nd Annual Asilomar Conference on Signals, Systems, and Computers*, 1762–1766, 2018.
- **Bahman Moraffah,** Antonia Papandreou-Suppappola, "**Random Infinite Tree and Dependent Poisson Diffusion Process for Nonparametric Bayesian Modeling in Multiple Object Tracking**", *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 5217–5221, 2019**.**
- **Bahman Moraffah,** Cesar Brito, Bindya Venkatesh, Antonia Papandreou-Suppappola, "**Use of Hierarchical Dirichlet Processes to Integrate Dependent Observations from Multiple Disparate Sensors for Tracking** ", *22nd International Conference on Information Fusion*, Invited paper, 2019.

Selected Publications

- **Bahman Moraffah**, Muralidhar Rangaswamy, Antonia Papandreou-Suppappola, "**Nonparametric Bayesian Methods and the Dependent Pitman-Yor Process for Modeling Evolution in Multiple State Priors**", *22nd International Conference on Information Fusion*, 2019.
- **Bahman Moraffah,** Antonia Papandreou-Suppappola, "**Inference for Multiple Object Tracking: A Bayesian Nonparametric Approach**", Submitted in *IEEE Transactions on Signal processing*, April 2019.
- **Bahman Moraffah,** Cesar Brito, Bindya Venkatesh, Antonia Papandreou-Suppappola, "**Tracking Multiple Objects with Multimodal Dependent Measurements: Bayesian Nonparametric Modeling**", submitted to *53rd Annual Asilomar Conference on Signals, Systems, and Computers*, 2019

- **And ….**

There are three kinds of lies:

lies, damned lies, and statistics.

Attributed to Benjamin Disraeli by Mark Twain