Solving Person Re-identification in Non-overlapping Cameras using Efficient Gibbs Sampling

V. John, G. Englebienne, B. Krose
Speaker: G. Englebienne
University of Amsterdam

Person Re-identification: Problem Formulation

Same camera in a network

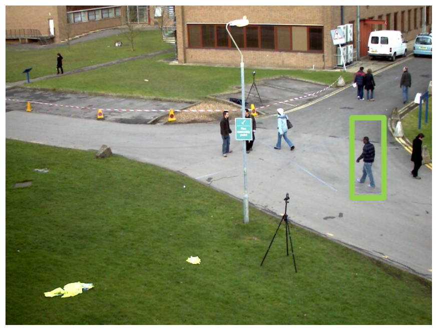

Different camera (overlapping)

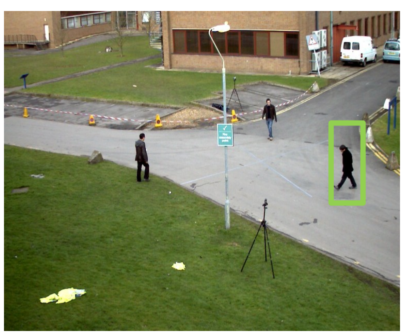

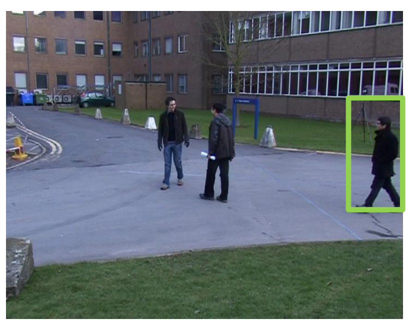

Different camera (non-overlapping)

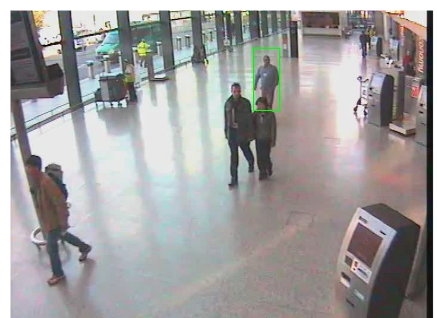

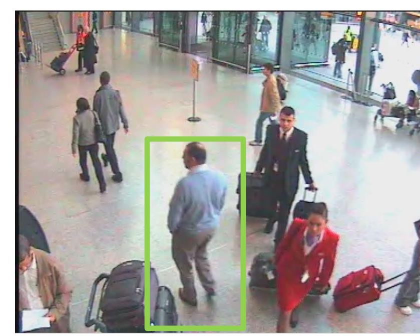

Person Re-identification: Problem Formulation

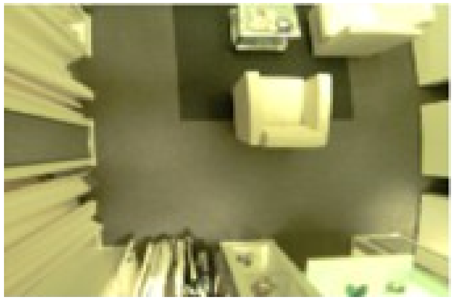

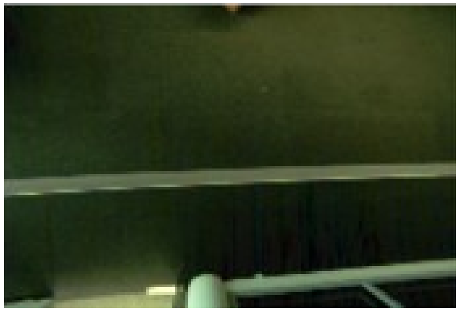

Applications:

Behaviour analysis, Surveillance, Logistics...

Solution:

Appearance-based features and temporal information

Issues:

Appearance variations across time and cameras

Loss of temporal information (non-overlapping views)

High computational complexity

Person Re-identification: Algorithm Overview

Problem formulation

Unsupervised person trajectory re-identification in non-overlapping network addressing:

Illumination variation

Camera gain variation

High computational complexity

Solution overview

Probabilistic graphical model

Infer person labels using efficient Gibbs sampling

Closed-form analytical updates of absolute appearance, gain and illumination variation parameters

Person Re-identification: Probabilistic Graphical Model

X = [**x1** ,..., **xn**]   (Set of observations)

Z = [*z1* ,..., *zn*]   (Set of labels to be solved)

xi = [**ai**, li, ei, ti] (ith trajectory)

- **ai** = average raw RGB color model
- *li* = camera label
- *ei* = trajectory entrance time
- *ti* = trajectory leaving time

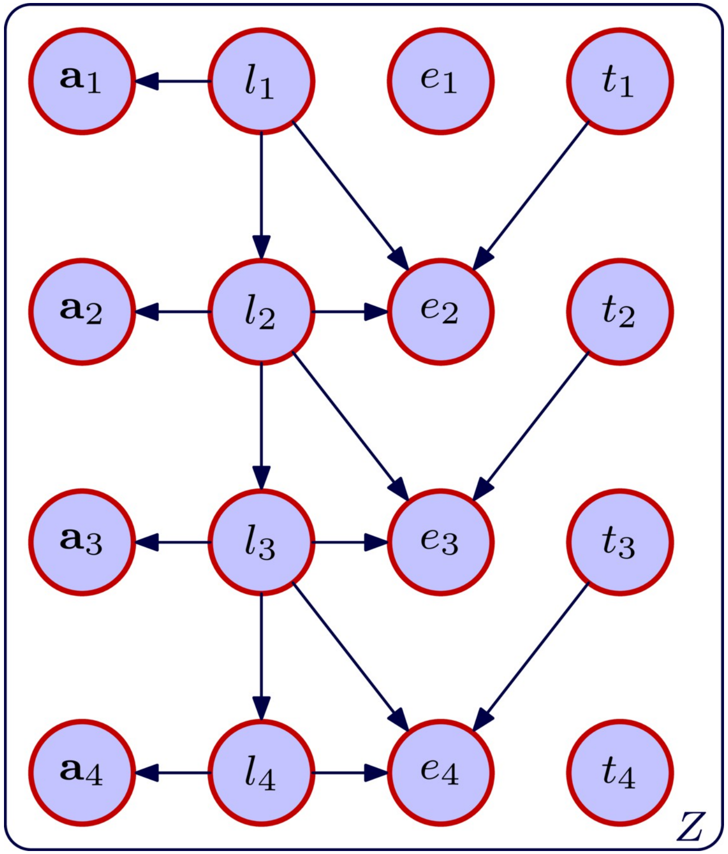

Person Re-identification: Analytical Update

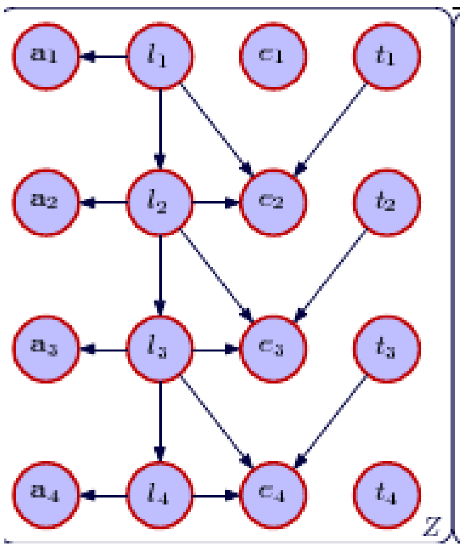

- Known labels
- Efficient Bayesian inference using Markov blanket

Appearance = gain * (rgb + cam illumination noise)

Gamma(Gaussian) Gaussian

Analytical Update

- Distributions

Transitions probability:	Multinomial distribution
Transition time:	Gamma distribution
Gain:	Gaussian distribution
Appearance:	Gaussian distribution
Illumination variation:	Gaussian distribution

- Conjugate prior

Gain(unknown mean and precision)	Normal-gamma
RGB, Appearance (mean, precision)	Normal-Wishart
RGB, Illumination (mean, precision)	Normal-Wishart

- Closed-form analytical solution posterior

Person Re-identification: Probabilistic Graphical Model

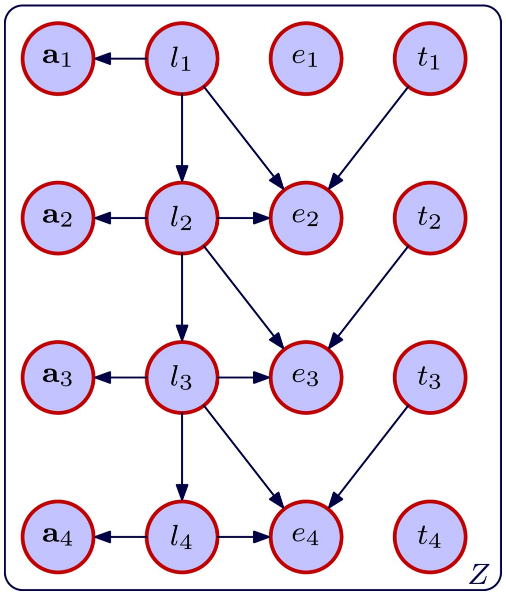

Known labels

Efficient inference

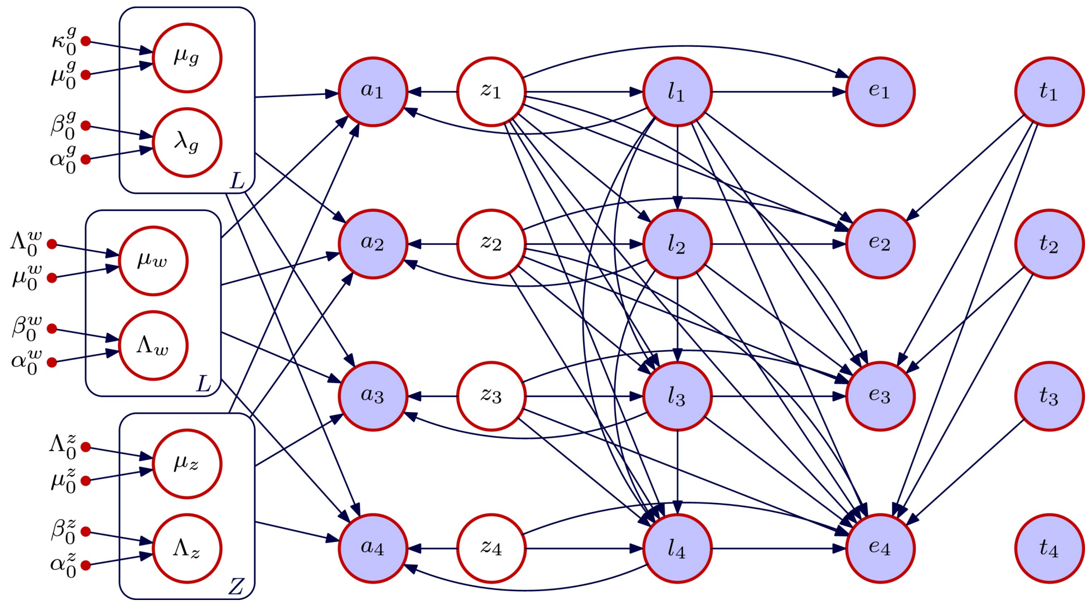

Unknown labels

Inefficient inference

(Markov blanket is complete set of observation and labels)

Person Re-identification: Efficient Gibbs Sampling

- Gibbs sampling based MCMC

$$\boxed{p(z_i | \mathbf{z}_{\neg i}, \mathcal{X}) = \frac{p(\mathcal{X} | \mathbf{z})\ p(\mathbf{z})}{\sum_{z_i=1}^{N} p(\mathcal{X} | \mathbf{z})\ p(\mathbf{z})}}$$

- Computational complexity over observations
  - Naive Gibbs sampling: **quadratic**

p(*zi* **|** **z**-i, **X**)	linear
p(**z**)	constant
p(**X** **|** **z**)	linear

  - Gibbs sampling with book keeping: **linear**

p(*zi* **|** **z**-i, **X**)	constant
p(**z**)	constant
p(**X** **|** **z**)	linear

Person Re-identification: Book Keeping

Given **z**=[*z1* ,…, *zn*] (Set of person labels)

Store b=[*b1* ,…, *bn*] (Set of previous indices associated with z)

Store f=[*f1* ,…, *fn*] (Set of future indices associated with z)

Compute p(**X | z**)

For each new label *zi'*,

1) **Compute** p(**X |** *zi',* **z**-i) in constant time given p(**X | z**)

$$p(\mathcal{X}|\mathbf{z}_{\neg i}, z_i') = p(\mathcal{X}|\mathbf{z}) \times \frac{p(\mathbf{x}_i|z_i', \mathbf{x}_{b_i'})p(\mathbf{x}_{f_i}|z_{f_i}, \mathbf{x}_{b_i})p(\mathbf{x}_{f_i'}|z_i', \mathbf{x}_i)}{p(\mathbf{x}_i|z_i, \mathbf{x}_{b_i})p(\mathbf{x}_{f_i}|z_{f_i}, \mathbf{x}_i)p(\mathbf{x}_{f_i'}|z_i', \mathbf{x}_{b_{f_i'}})}$$

2) **Modify f** and **b**

$$f_{b_i} \leftarrow f_i \qquad b_{f_i} \leftarrow b_i \qquad f_i \leftarrow f_i' \qquad b_i \leftarrow b_{f_i} \qquad b_{f_i} \leftarrow i \qquad f_{b_i} \leftarrow i$$

Person Re-identification: Book Keeping

Calculate b and f

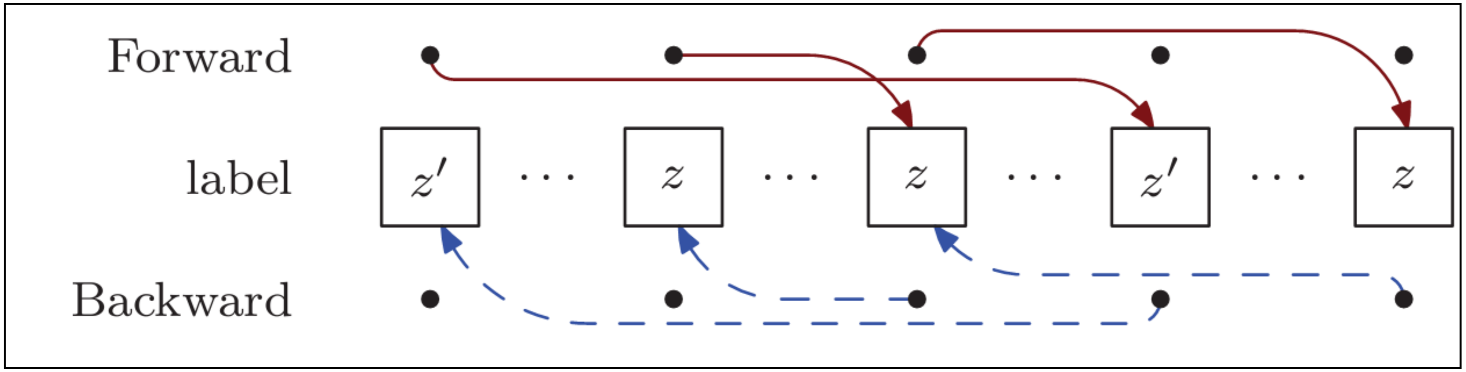

Modify b and f for new label for ith observation

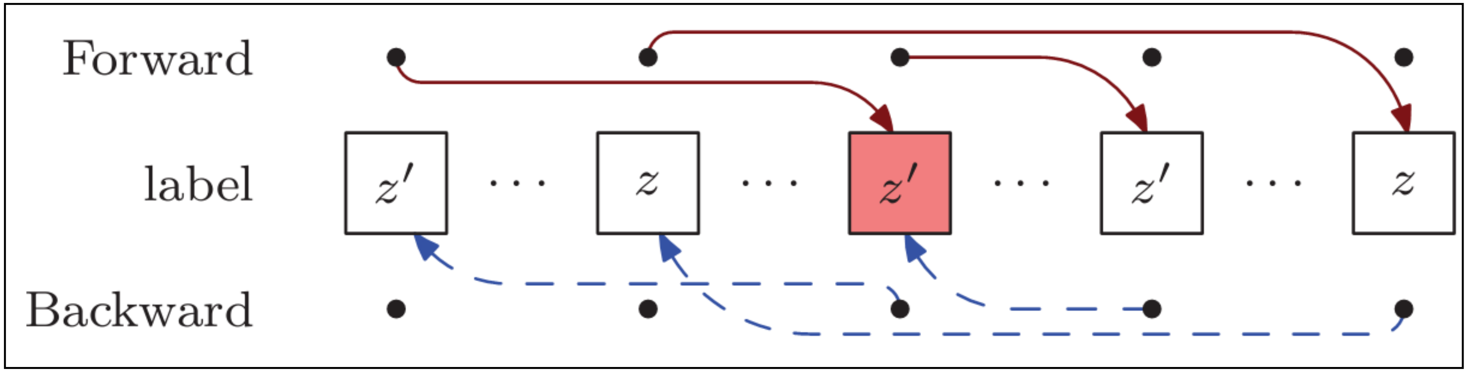

Person Re-identification: Experimental Results

- Ceiling-mounted multiple camera datasets
  - 2 Datasets with 5-13 cameras and 5-10 people
- Comparative experiments
  - Pasula sampler
  - MCMC: transition by swapping random observation pairs
  - Maximum-Likelihood
- Algorithm parameter
  - Naïve appearance model (Raw rgb)

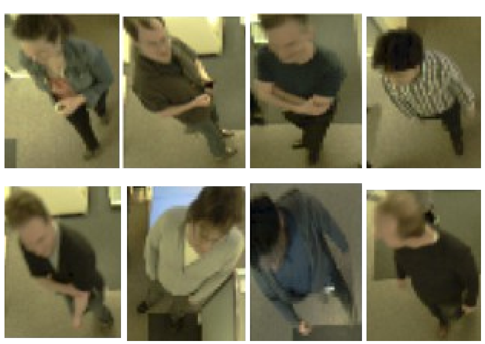

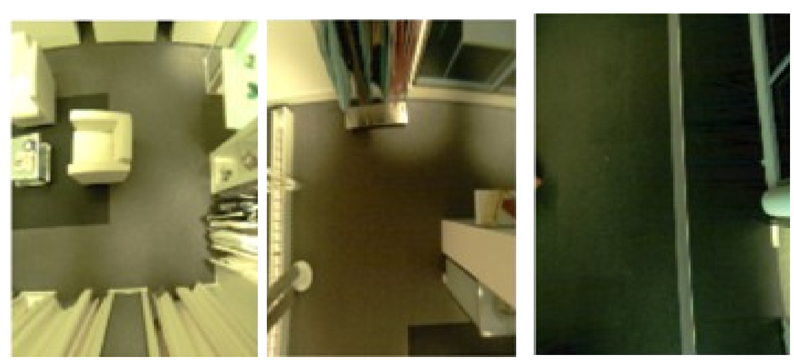

Person Re-identification: Experimental Results

Comparative (5 trials)

Accuracy

Dataset	PGM-based	Pasula sampler
1 -5 sub	87.3+4.3%	67.5 + 6.1%
1-10 sub	86+ 5.2%	65.0 + 8.1%
2-5 sub	84+ 4.1 %	62.5 + 7.3%

Computational Complexity

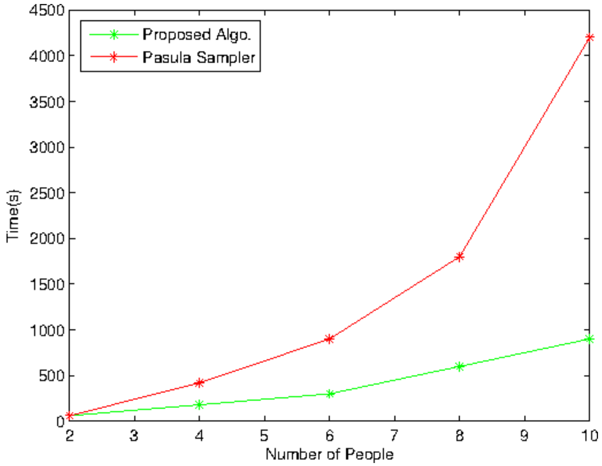

Algorithm parameters (5 trials)

Naïve Appearance

Dataset	PGM-based	Naïve appearance
1 -5 sub	**87.3+4.3%**	**75 + 5.1%**
1-10 sub	**86+ 5.2%**	**74.3 + 4.6%**
2-5 sub	**84+ 4.1 %**	**73.2 + 5.3%**

Frames

Frames	Mean and Std. Dev
200	**72.5+5.6%**
600	**77+5.2%**
1000	**81.6+4.6%**
1400	**82.1+3.9%**
2000	**84+4.1%**

Person Re-identification: Summary

- Unsupervised person trajectory re-identification in non-overlapping network addressing:
  - Illumination variation
  - Camera gain variation
  - High computational complexity
- Using probabilistic graphical model
- Infer person labels using efficient Gibbs sampling
- Closed-form analytical updates of appearance, gain and illumination variation parameters
- Significantly improved performance

Thank you