EPSRC
Engineering and Physical Sciences
Research Council

BACKGROUND SUBTRACTION WITH DIRICHLET PROCESSES

TOM S. F. HAINES & TAO XIANG
{thaines,txiang}@eecs.qmul.ac.uk

Queen Mary
University of London

OVERVIEW

- We present a **background subtraction** algorithm.
- Used for first phase **video processing** to separate the interesting foreground from the background.
- Two part approach - per pixel model followed by regularisation step:

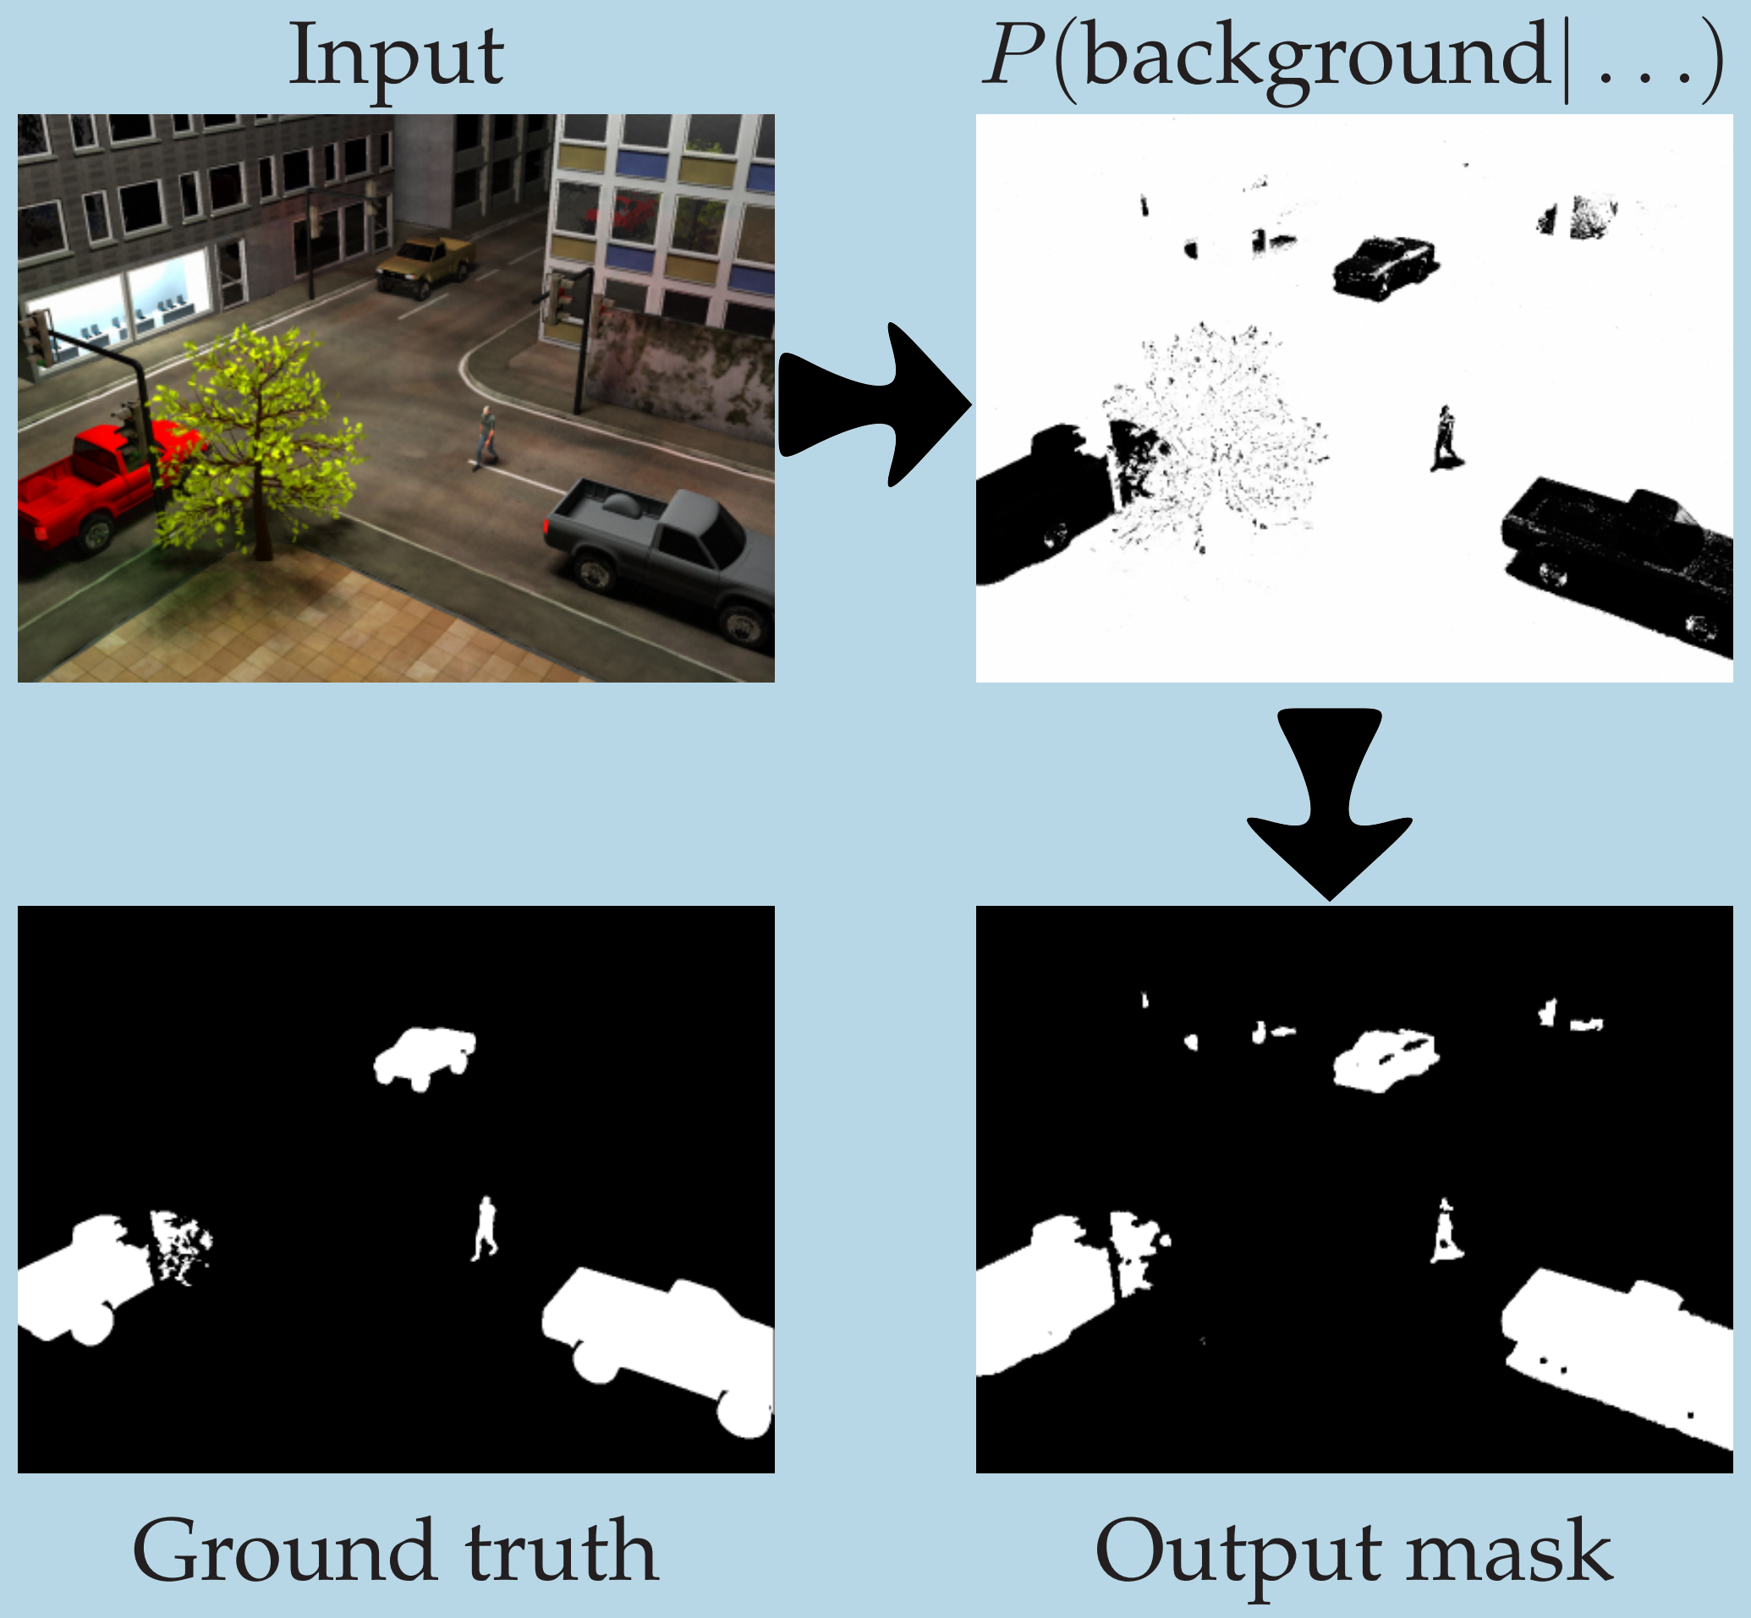

- Each pixel has a **Dirichlet process Gaussian mixture model** modelling its background colour.
- **Gibbs sampling** is used to update the model - each new pixel is sampled once and then discarded, to avoid memory issues.
- A **certainty cap** is employed. This prevents overconfidence and allows the model to **update as the background changes**.
- A colour model that separates *luminance* and *chromaticity* is employed, so the importance of luminance can be reduced and some **robustness to shadows** obtained.
- Using Bayes rule and a uniform model of foreground colour a probability map of each pixel being foreground or background is generated.
- **Belief propagation** is used to **regularise** the map, with an **edge preserving smoothing** term.

MORE...

In addition to the paper an implementation using Python/C/OpenCL can be obtained from **thaines.com**.

METHOD

Dirichlet Process:

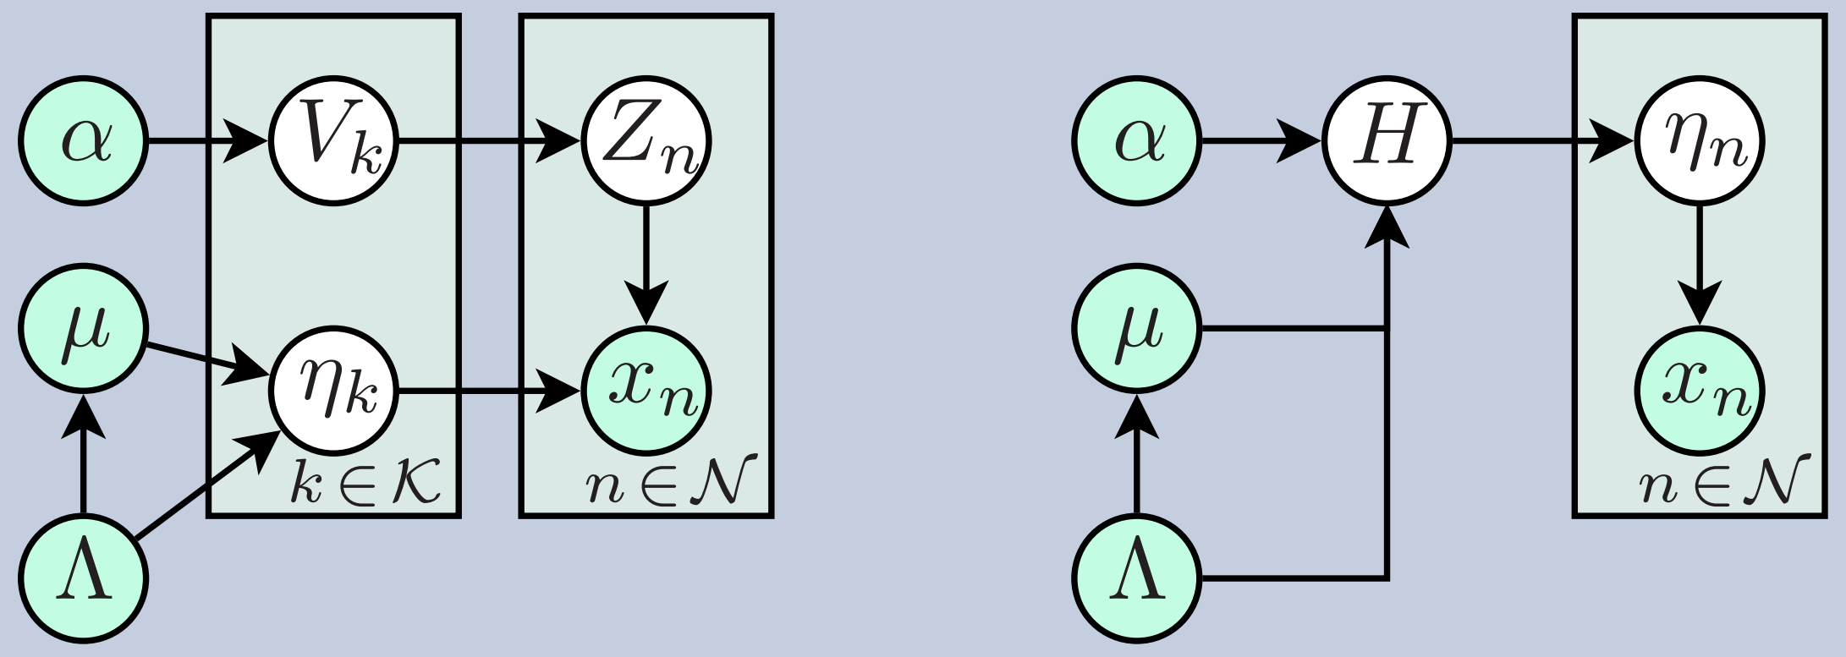

Stick breaking Chinese restaurant

A distribution over distributions - from a Dirichlet process (DP) you draw an infinite multinomial. It has two parameters, a **base measure** (G_0) and a **concentration parameter** (α). It assumes a logarithmic distribution of mixture components sizes.
Two views of it:

- **Stick breaking construction**: You start with a stick of length 1, representing the probability mass. For eternity break it in two, with one half going to the next break. The other half represents the probability of drawing an associated draw from the base measure when drawing from the infinite multinomial.
- **Chinese restaurant process** (CRP): Has the infinite multinomial integrated out. Customers (samples) arrive at a restaurant and choose tables (mixture components) to sit at. Each table is associated with one menu item (draw from the base measure), selected by the first customer to sit at it. The number of customers at a table is its selection weight, whilst α is the selection weight of a new table.

Non-parametric Mixture Models:

A Bayesian mixture model often has a Dirichlet distribution as a prior. Swapping it for a Dirichlet process creates a model where the number of mixture components is dynamically selected to match the data.

- A Dirichlet process Gaussian mixture model is used. The Gaussian drawn from the conjugate prior is integrated out using student-t distributions.
- Each pixel is modelled with a DPGMM, so multimodal backgrounds are not a problem.
- Certainty is capped, in terms of component weight - when exceeded it is multiplied to return it to the cap. This allows it to forget old models.

RESULTS

SABS: (synthetic)

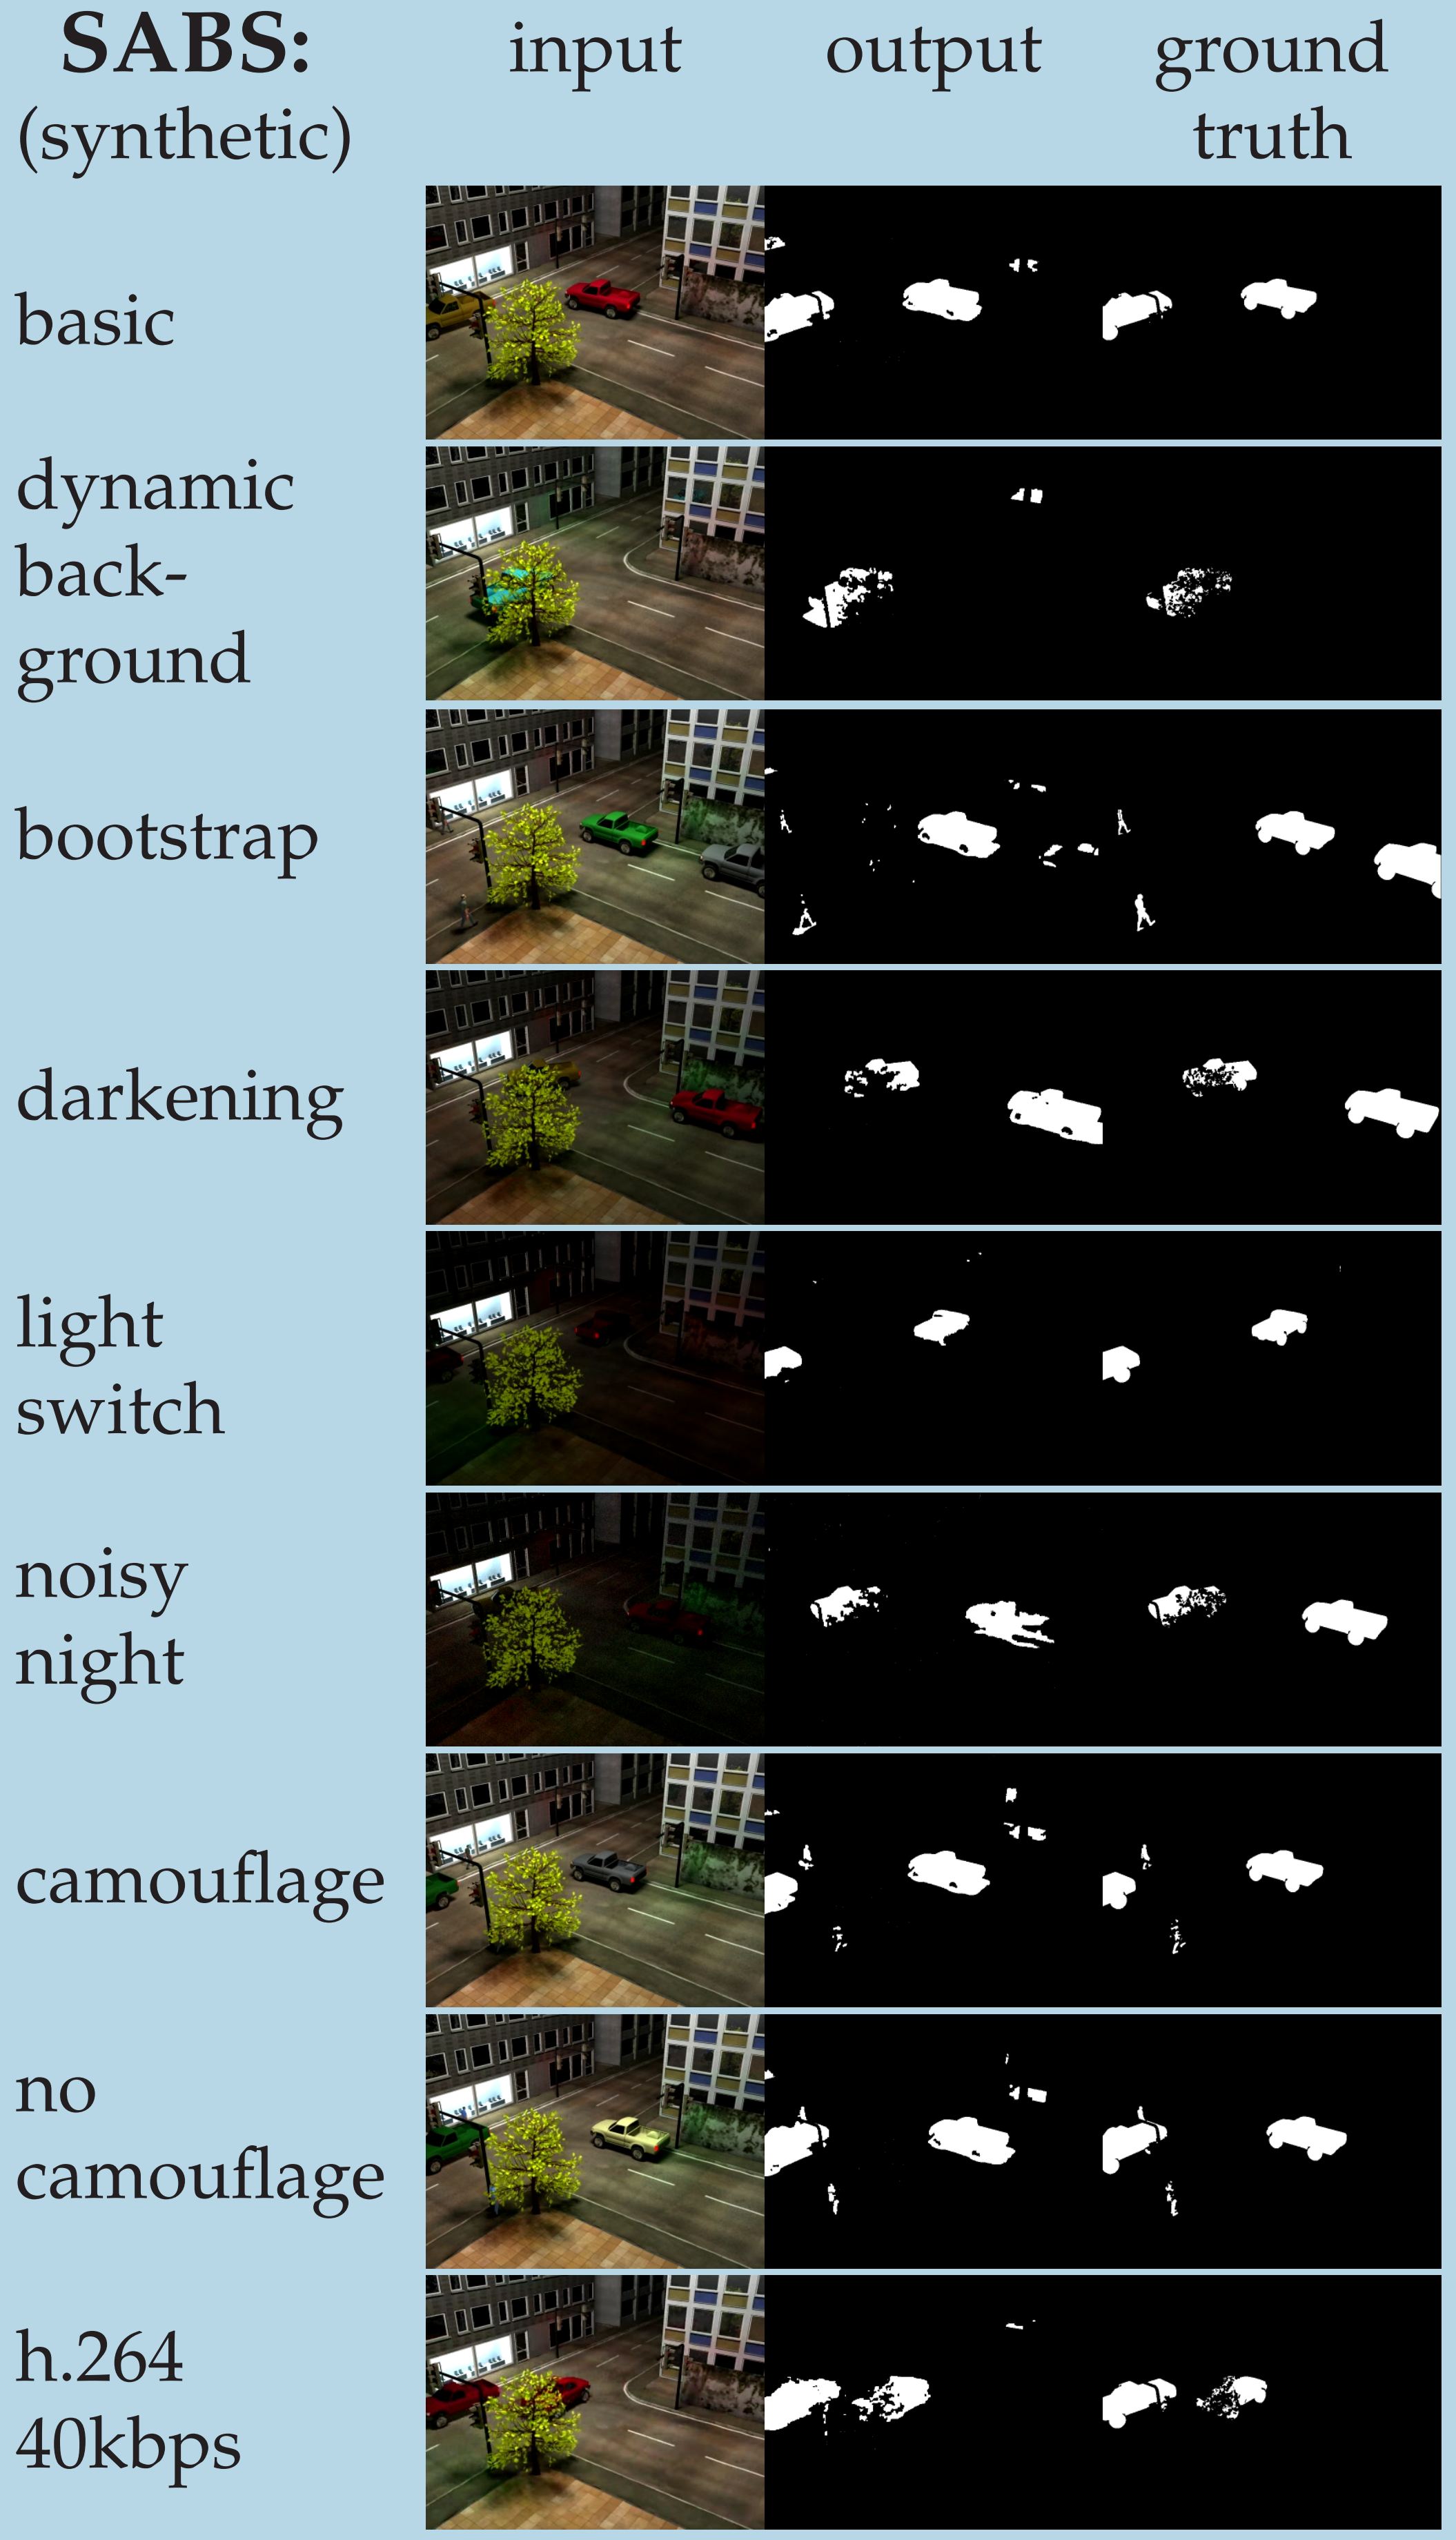

Wallflower:

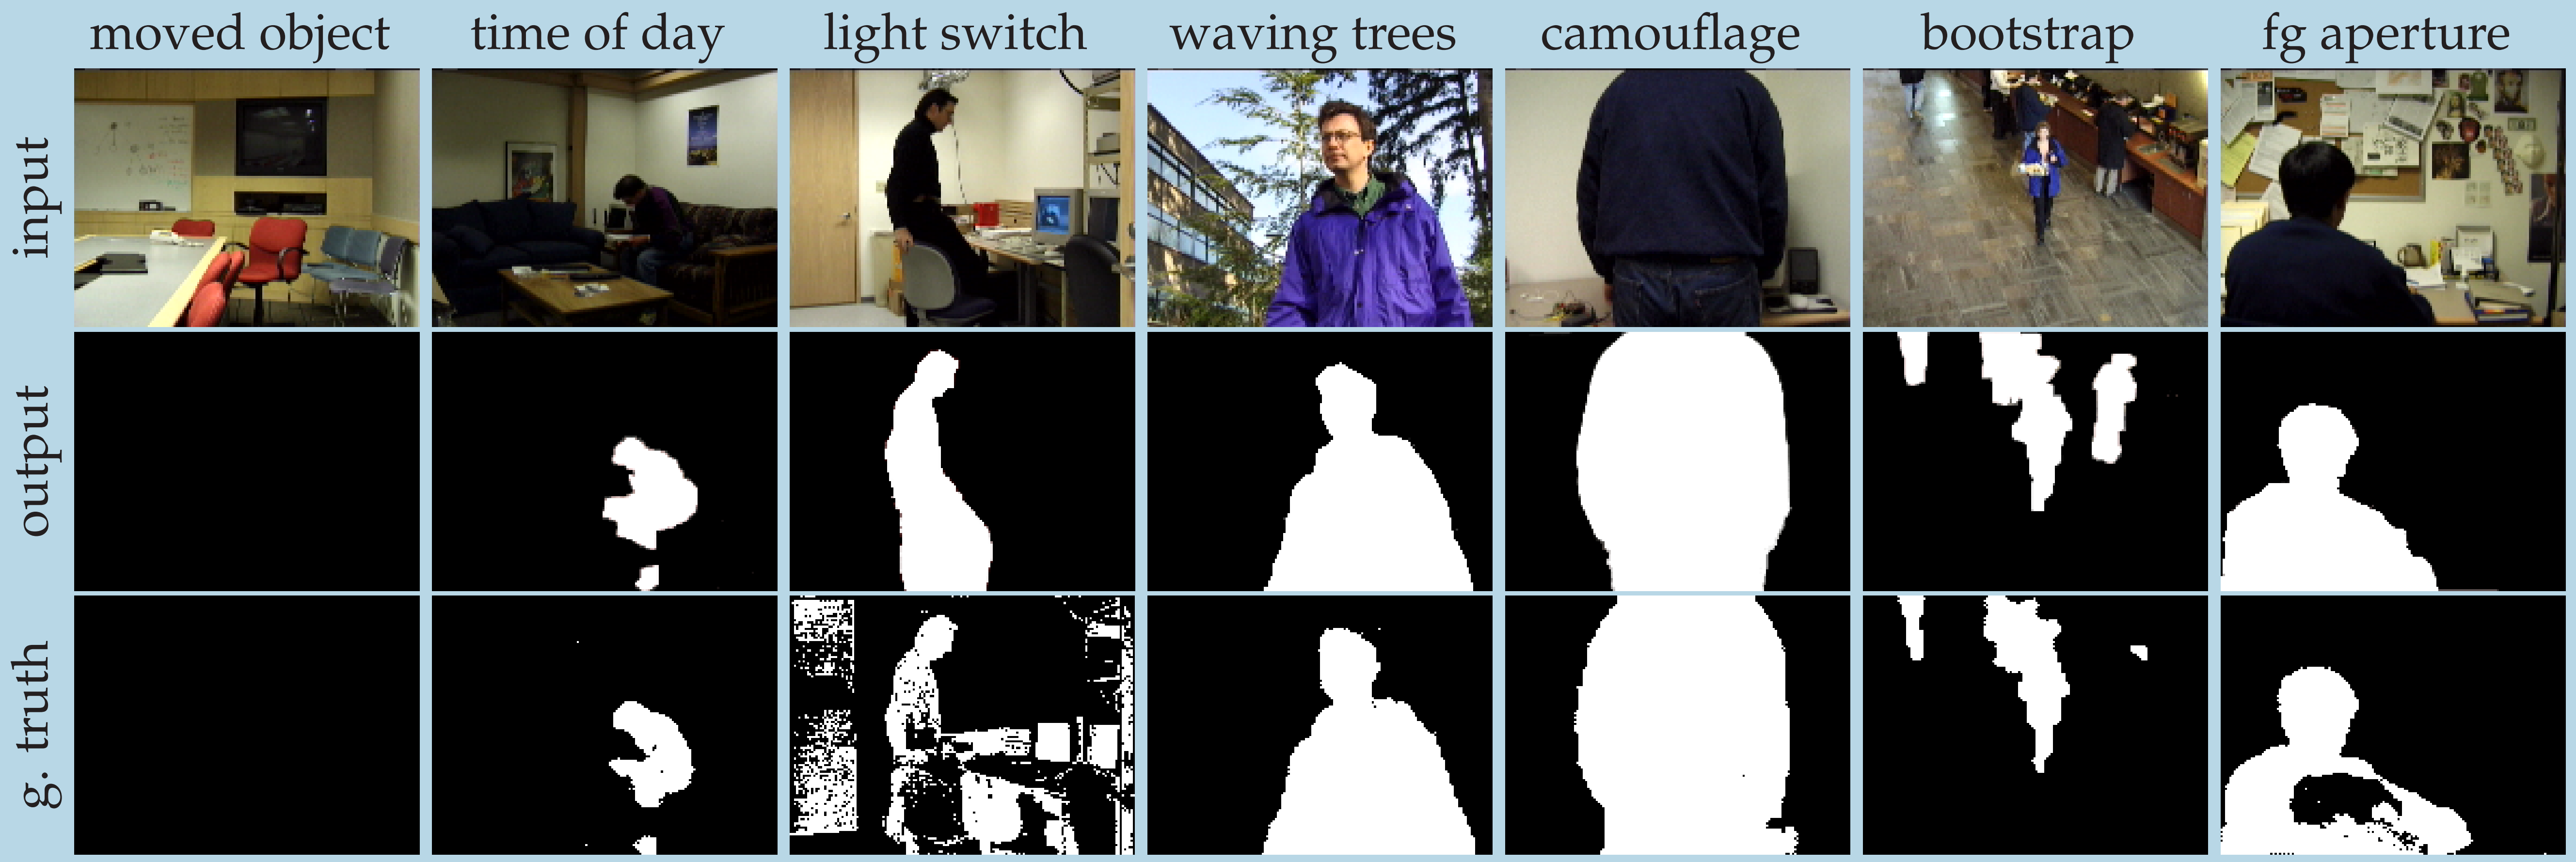

Star:

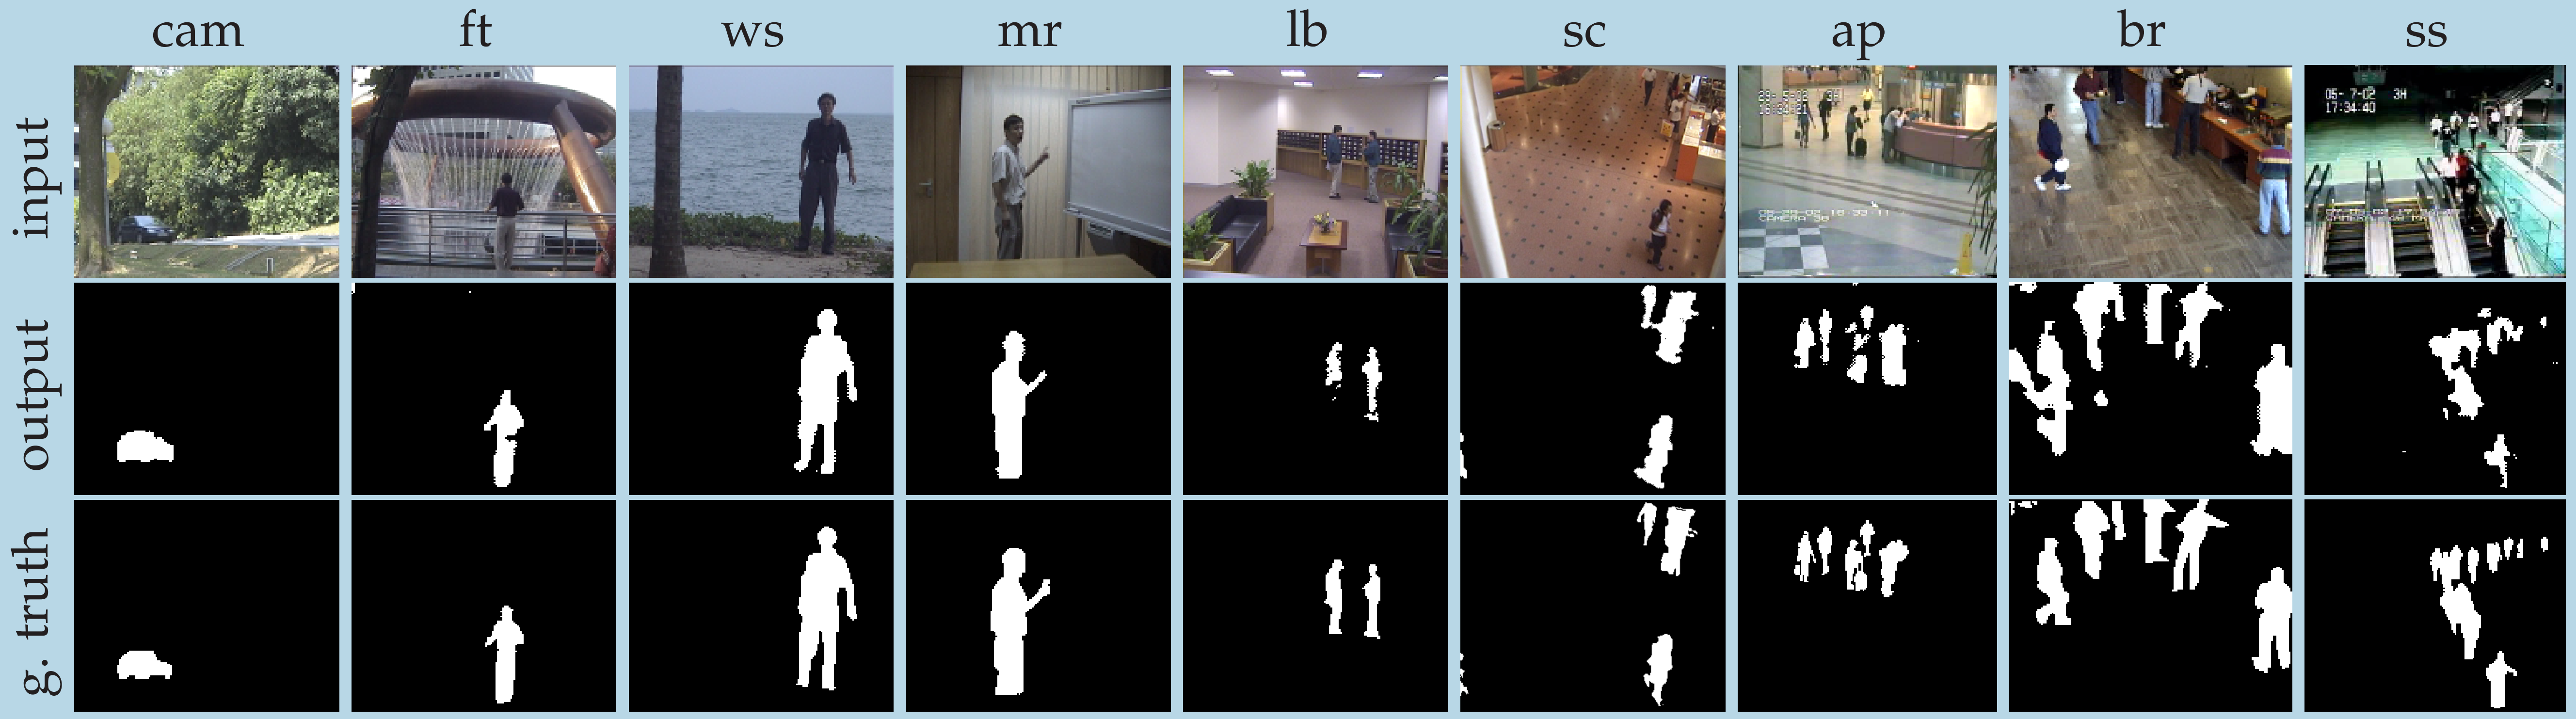

Average improvements of 27% for *SABS*; 33% fewer mistakes for *wallflower*; 4% average improvement for *star*.

(Against nearest competitors, such as Barnich & Droogenbroeck (2009), Maddalena & Petrosino (2008), Li et al. (2003), Stauffer & Grimson (1999) etc.)