Acceleration and Regularization of Probabilistic Deep Learning with Variational Formulations

Xu (Shell) Hu[1,2]
Supervised by Guillaume Obozinski[1], Nikos Komodakis[1], Renaud Marlet[1], Mathieu Aubry[1] and Sven Oesau[2]

[1]École des Ponts ParisTech
[2]Centre Scientifique et Technique du Bâtiment

Context and Objectives

- A **computer vision** problem: *semantic segmentation*.
  Assign a semantic label (from a finite set) to every pixel in the image or every 3D point in 3D reconstructed data.
- **Applications**:

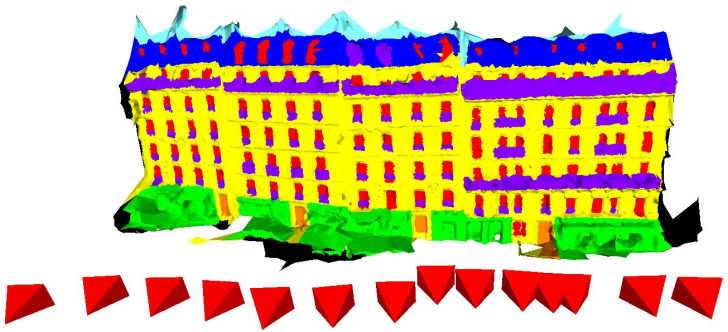

Digitize existing buildings

Autonomous driving

Image manipulation

Challenges

- **Computational issue**:
  - There are tens of billions of pixels and 3D points.
  - Each point labeling is a classification problem.
  - We have to take into account the structure among the points.

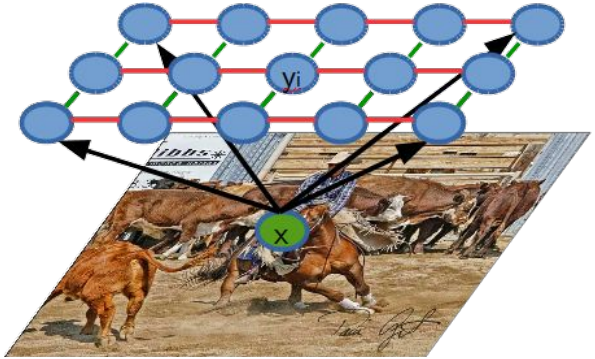

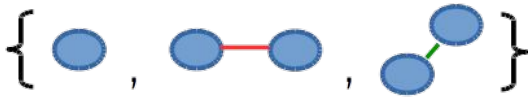

- **Data issue**: only very few pixels/points are annotated (a lot of human labor).

Machine Learning Methods

- **Conditional random fields**:

$$p(y|x;w) = \frac{1}{Z(x;w)} \exp\left(\sum_{a\in\mathcal{A}} \sum_{c\in G_a} \theta_c(y_c, x; w_a)\right)$$

- **Bayesian deep neural networks:**

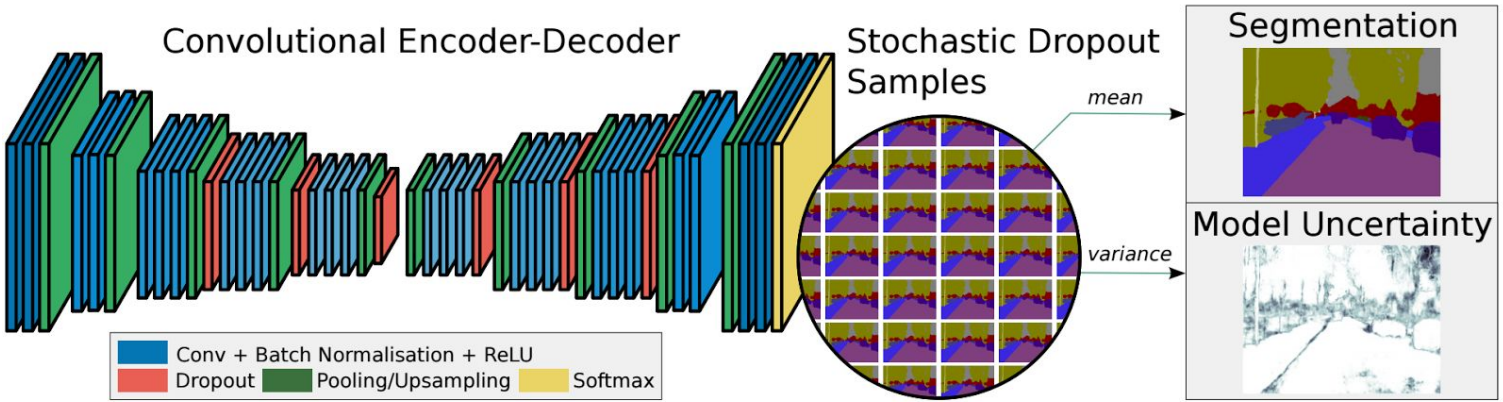

Main Results

1. For *computational issue*, we proposed **acceleration** methods:
   - SDCA-powered inexact dual augmented Lagrangian method for fast CRF learning at AISTATS 2018.
   - Amortized conditional random fields in submission.
2. For *small data issue*: we proposed **regularization** methods:
   - Variational information distillation for transfer learning at NeurIPS-CL 2018.
   - β-BNN: A Rate-Distortion Perspective on Bayesian Neural Networks at NeurIPS-BDL 2018.

   (**Please stop by our poster for technical details.**)