DATALOGISK INSTITUT
KØBENHAVNS UNIVERSITET

Cross-center smoothness prior for Bayesian image segmentation

WM Kouw   SN Ørting   J Petersen   KS Pedersen   M de Bruijne

MR images vary across medical centers due to calibration and acquisition protocols. But *segmentations* are relatively consistent. How segmentations are supposed to look like, can be learned separately. We present a smoothness prior that is fit to segmentations from a source medical center. This empirical prior is incorporated into an unsupervised Bayesian image segmentation model. The model clusters voxel intensities in the target center, such that its segmentations are similarly smooth.

Problem

MRI scans vary across medical centers due to different acquisition protocols.

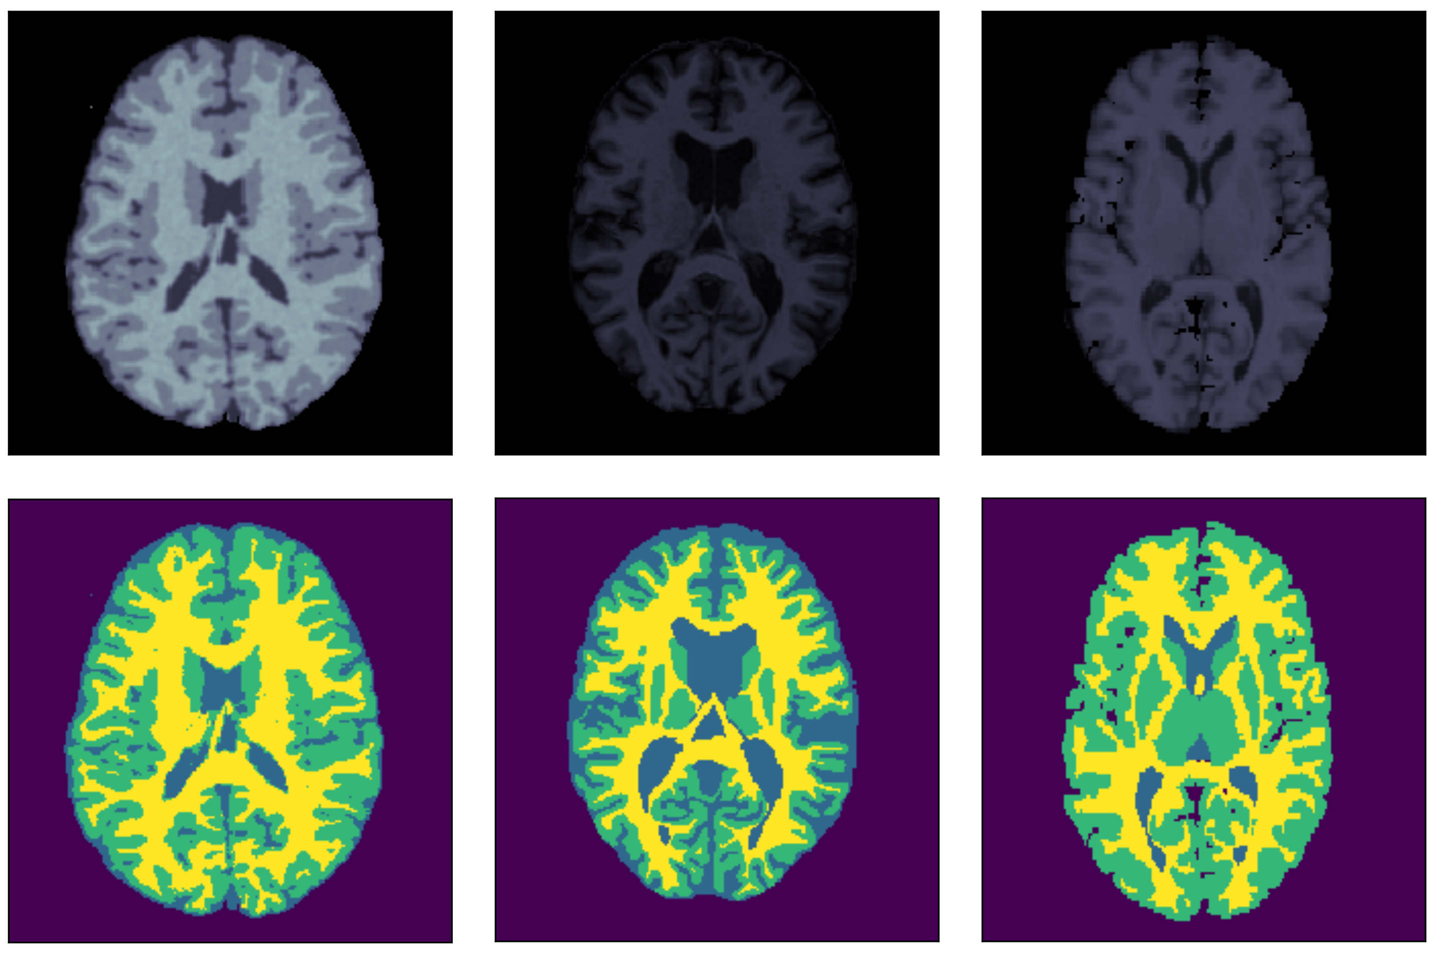

Model specification

We use a Gaussian mixture model for voxel intensities given tissues

$$p(X \mid Y; \pi, \mu, \Lambda) = \prod_{i=1}^{N} \prod_{k=1}^{K} \big[\pi_k \mathcal{N}(x_i \mid \mu_k, \Lambda_k^{-1})\big]^{y_{ik}}$$

with conjugate priors

$$\pi \sim \mathcal{D}(\alpha_0), \quad \mu_k \sim \mathcal{N}(v_{0k}, (\gamma_{0k}\Lambda_k)^{-1}), \quad \Lambda_k \sim \mathcal{W}(\nu_{0k}, \Delta_{0k})$$

Hidden Potts-Markov Random Field

Smoothness of the segmentations is modeled using a local Potts prior:

$$\log p(y_i \mid y_{\delta_i}, \beta) = \sum_{k=1}^{K} \beta_k y_{ik} \sum_{j \in \delta_i} y_{jk} - \log \sum_{k=1}^{K} \exp\big(\beta_k \sum_{j \in \delta_i} y_{jk}\big)$$

Cross-center empirical prior

The Potts prior is fit to segmentations produced at a source medical center:

$$\hat{\beta} = \underset{\beta \in \mathbb{R}^+}{\arg\max} \sum_{i=1}^{N} \log p(y_i \mid y_{\delta_i}, \beta)$$

where the gradient is

$$\frac{\partial}{\partial \beta} \log p(y_i \mid y_{\delta_i}, \beta) =$$

$$\sum_{k=1}^{K} y_{ik}\bar{y}_{ik} - \sum_{l=1}^{K} \bar{y}_{il} \exp\big(\beta_l \bar{y}_{il}\big) / \big[\sum_{m=1}^{K} \exp\big(\beta_m \bar{y}_{im}\big)\big]$$

Variational Bayes

Incorporating the empirical Potts prior into the Gaussian mixture model:

$$p(X, Y, \theta \mid \beta) = p(X \mid Y, \theta) p(Y \mid \beta) p(\theta)$$

Inference consists of a variational approximation and coordinate ascent, using a structured mean-field factorization of the recognition distribution:

$$\log p(X \mid \beta) = \log \iint p(X, Y, \theta \mid \beta) \mathrm{d}\theta \mathrm{d}Y$$

$$\geq \iint q(Y \mid \beta) q(\theta) \log \frac{p(X, Y, \theta \mid \beta)}{q(Y \mid \beta) q(\theta)} \mathrm{d}\theta \mathrm{d}Y$$

Optimal forms of the recognition factors are:

$$\log q^*(\theta) = \mathbb{E}_{q(Y|\beta)}\big[\log p(X \mid Y, \theta)\big] + \log p(\theta)$$

$$\log q^*(Y|\beta) = \sum_{k=1}^{K} y_{ik} \log r_{ik}$$

where

$$\log r_{ik} =$$

$$\mathbb{E}_{q(\pi)}\big[\log \pi_k\big] + \mathbb{E}_{q(\mu,\Lambda)}\big[\log \mathcal{N}(x_i \mid \mu_k, \Lambda_k^{-1})\big] + \beta_k \sum_{j \in \delta_{ik}} y_{ik}$$

Experiments

We fit on data from a source center and segment images of a target center.

	Methods	Brainweb	MRBrainS	IBSR
Brainweb	U-net	-	0.448 (.008)	0.384 (.019)
	Gaussian mixture	0.116 (.031)	0.268 (.017)	0.567 (.019)
	hidden Potts	0.117 (.032)	0.253 (.017)	0.541 (.022)
MRBrainS	U-net	0.257 (.003)	-	0.589 (.022)
	Gaussian mixture	0.102 (.011)	0.282 (.021)	0.575 (.080)
	hidden Potts	0.102 (.008)	0.277 (.020)	0.515 (.092)
IBSR	U-net	0.334 (.007)	0.425 (.015)	-
	Gaussian mixture	0.099 (.005)	0.296 (.044)	0.560 (.039)
	hidden Potts	0.101 (.005)	0.259 (.041)	0.544 (.040)

Example segmentations

Gaussian mixture | hidden Potts | U-net

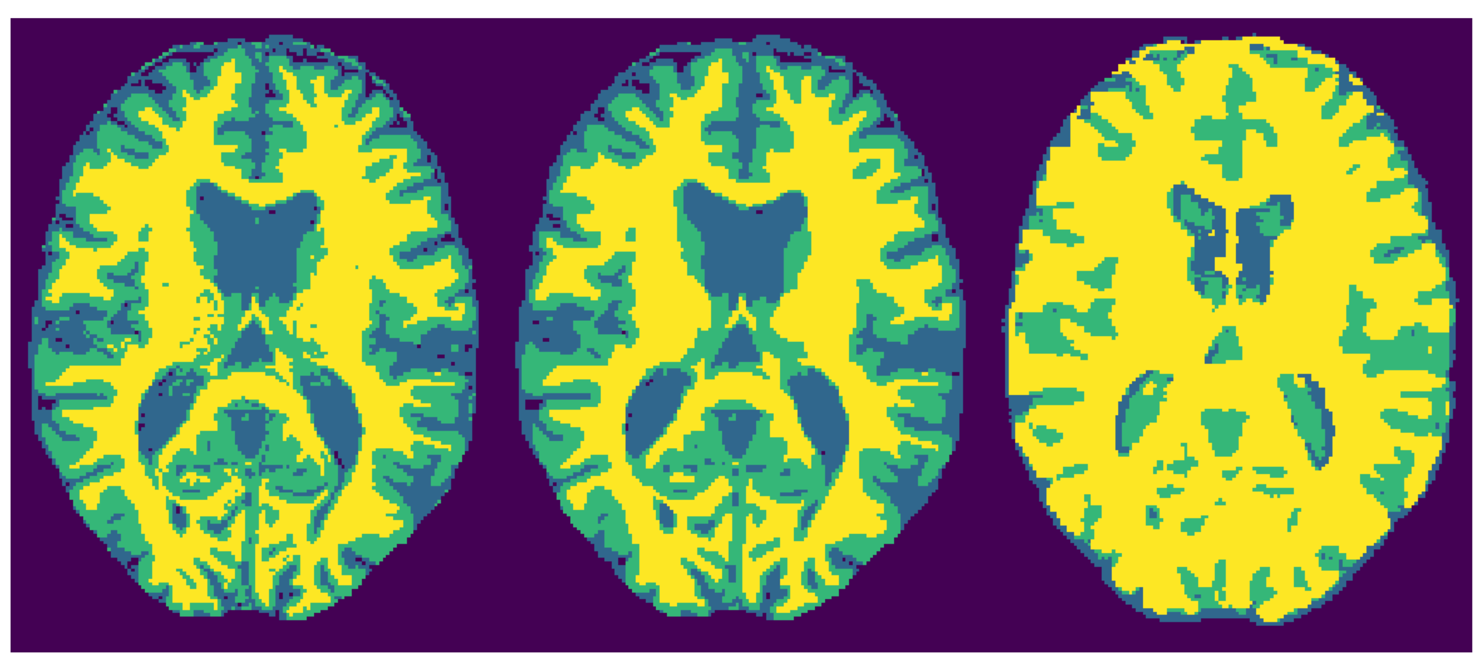

Empirical priors for properties of segmentations are a powerful way of guiding Bayesian image segmentation models. Here we adressed smoothness, but aspects such as relative position and tissue proportions could be captured as well. Learning from segmentations avoids the necessity of labeled data for each MRI scanner.

github.com/wmkouw/cc-smoothprior
https://arxiv.org/abs/1903.04191