

Small-Variance Nonparametric Clustering on the Hypersphere

Julian Straub¹, Trevor Campbell², Jonathan P. How², John W. Fisher III¹

¹CSAIL and ²LIDS, Massachusetts Institute of Technology

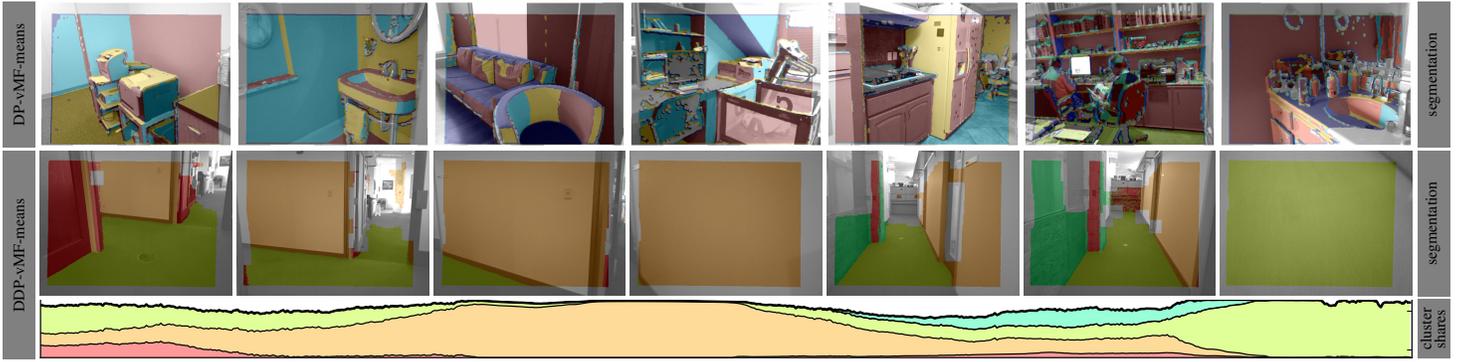


Figure 1: The DP-vMF-means algorithm adapts the number of clusters to the complexity of the surface normal distribution as depicted in the first row. For a sequence of batches of data, the DDP-vMF-means algorithm allows temporally consistent clustering as shown in the second row. The coloring indicates cluster membership of the surface normals extracted from the depth channel of an RGB-D camera. Note that only surface normals are used for clustering; the grayscale images in the figure are for display purposes only, and RGB and raw depth information from the camera are not used for inference.

Man-made environments and objects exhibit clear structural regularities such as planar or rounded surfaces. These properties are evident on all scales from small objects such as books, to medium-sized scenes like tables, rooms and buildings and even to the organization of whole cities. Such regularities can be captured in the statistics of surface normals that describe the local differential structure of a shape. These statistics contain valuable information that can be used for scene understanding, plane segmentation, or to regularize a 3D reconstruction.

Inference algorithms in fields such as robotics or augmented reality, which would benefit from the use of surface normal statistics, are not generally provided a single batch of data a priori. Instead, they are often provided a stream of data batches from depth cameras. Thus, capturing the surface normal statistics of man-made structures often necessitates the temporal integration of observations from a vast data stream of varying cluster mixtures. Additionally, such applications pose hard constraints on the amount of computational power available, as well as tight timing constraints.

We address these challenges by focusing on flexible Bayesian nonparametric (BNP) Dirichlet process mixture models (DP-MM) which describe the distribution of surface normals in their natural space, the unit sphere in 3D, \mathbb{S}^2 . Taking the small variance asymptotic limit of the DP-MM of von-Mises-Fisher (vMF) distributions, we obtain a fast k -means-like algorithm, which we call DP-vMF-means, to perform nonparametric clustering of data on the unit hypersphere. Furthermore, we propose a novel dependent DP mixture of vMF distributions to achieve integration of directional data into a temporally consistent streaming model. Small variance asymptotic analysis yields the k -means-like DDP-vMF-means algorithm. In this extended abstract we discuss the DP-vMF-means algorithm derivation, and leave the DDP-vMF-means algorithm to the full paper.

The Dirichlet process (DP) [2] with concentration α has been widely used as a prior for mixture models with a countably infinite set of clusters [1, 4]. Assuming a base distribution $\text{vMF}(\mu; \mu_0, \tau_0)$, the DP is an appropriate prior for a vMF mixture with an unknown number of components with means $\{\mu_k\}_{k=1}^K$ and known vMF concentration τ . Gibbs sampling inference consists of sampling labels z_i for data $x_i \in \mathbb{S}^2$ from

$$p(z_i = k | \mathbf{z}_{-i}, \boldsymbol{\mu}, \mathbf{x}; \tau) \propto \begin{cases} |\mathcal{I}_k| \text{vMF}(x_i | \mu_k; \tau) & k \leq K \\ \alpha p(x_i; \mu_0, \tau_0, \tau) & k = K + 1, \end{cases} \quad (1)$$

where \mathcal{I}_k is the set of data indices assigned to cluster k , and sampling parameters from

$$p(\boldsymbol{\mu} | \mathbf{x}; \mu_0, \tau_0) = \text{vMF} \left(\boldsymbol{\mu}; \frac{\tau_0 \mu_0 + \tau \sum_{i=1}^N x_i}{\tau_0 \mu_0 + \tau \sum_{i=1}^N \|x_i\|_2}, \|\tau_0 \mu_0 + \tau \sum_{i=1}^N x_i\|_2 \right). \quad (2)$$

To derive a hyperspherical analog to DP-means [3] we consider the limit of the posterior distributions as $\tau \rightarrow \infty$.

Label Update: In the limit of the label sampling step (1) as $\tau \rightarrow \infty$, sampling from $p(z_i | \mathbf{z}_{-i}, \boldsymbol{\mu}, \mathbf{x}; \tau)$ is equivalent to the following assignment rule:

$$z_i = \arg \max_{k \in \{1, \dots, K+1\}} \begin{cases} x_i^T \mu_k & k \leq K \\ \lambda + 1 & k = K + 1. \end{cases} \quad (3)$$

Intuitively λ defines the maximum angular spread ϕ_λ of clusters about their mean direction, via $\lambda = \cos(\phi_\lambda) - 1$.

Parameter Update: Taking $\tau \rightarrow \infty$ in the parameter posterior for cluster k causes τ_0 and μ_0 to become negligible. Hence:

$$\mu_k = \frac{\sum_{i \in \mathcal{I}_k} x_i}{\|\sum_{i \in \mathcal{I}_k} x_i\|_2} \quad \forall k \in \{1, \dots, K\}. \quad (4)$$

Objective Function: We show that DP-vMF-means maximizes

$$J_{\text{DP-vMF}} = \sum_{k=1}^K \sum_{i \in \mathcal{I}_k} x_i^T \mu_k + \lambda K. \quad (5)$$

In the paper we demonstrate the performance and flexibility of DP-vMF-means on both synthetic data and the NYU v2 RGB-D dataset (see first row of Fig. 1). For DDP-vMF-means, Optimistic Iterated Restarts (OIR) parallelized label assignments, enable real-time temporally consistent clustering of batches of 300k surface normals collected at 30 Hz from a RGB-D camera (see second row of Fig. 1). Note, that DDP-vMF-means correctly reidentifies all directions after not observing them for a period of time in the middle of the sequence.

We envision a large number of potential applications for the presented algorithms in computer vision and in other realms where directional data is encountered. Implementations are available at <http://people.csail.mit.edu/jstraub/>.

[1] C.E. Antoniak. Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. *The annals of statistics*, pages 1152–1174, 1974.

[2] T.S. Ferguson. A Bayesian analysis of some nonparametric problems. *The annals of statistics*, pages 209–230, 1973.

[3] B. Kulis and M. I. Jordan. Revisiting k-means: New algorithms via Bayesian nonparametrics. In *ICML*, 2012.

[4] R.M. Neal. Markov chain sampling methods for Dirichlet process mixture models. *Journal of computational and graphical statistics*, 9(2):249–265, 2000.