

The Information Bottleneck and Geometric Clustering

DJ Strouse

danieljstrouse@gmail.com

Department of Physics, Princeton University, Princeton, NJ 08544, U.S.A.

David J. Schwab

dschwab@gc.cuny.edu

*Initiative for the Theoretical Sciences, CUNY Graduate Center,
New York, NY 10016, U.S.A.*

The information bottleneck (IB) approach to clustering takes a joint distribution $P(X, Y)$ and maps the data X to cluster labels T , which retain maximal information about Y (Tishby, Pereira, & Bialek, 1999). This objective results in an algorithm that clusters data points based on the similarity of their conditional distributions $P(Y | X)$. This is in contrast to classic geometric clustering algorithms such as k -means and gaussian mixture models (GMMs), which take a set of observed data points $\{x_i\}_{i=1:N}$ and cluster them based on their geometric (typically Euclidean) distance from one another. Here, we show how to use the deterministic information bottleneck (DIB) (Strouse & Schwab, 2017), a variant of IB, to perform geometric clustering by choosing cluster labels that preserve information about data point location on a smoothed data set. We also introduce a novel intuitive method to choose the number of clusters via kinks in the information curve. We apply this approach to a variety of simple clustering problems, showing that DIB with our model selection procedure recovers the generative cluster labels. We also show that, in particular limits of our model parameters, clustering with DIB and IB is equivalent to k -means and EM fitting of a GMM with hard and soft assignments, respectively. Thus, clustering with (D)IB generalizes and provides an information-theoretic perspective on these classic algorithms.

1 Introduction ---

Unsupervised learning is a crucial component of building intelligent systems (LeCun, 2016), since such systems need to be able to leverage experience to improve performance even in the absence of feedback. One aspect of doing so is discovering discrete structure in data, a problem known as clustering (MacKay, 2002). In the typical setup, one is handed a set of data points $\{x_i\}_{i=1}^N$ and asked to return a mapping from data point label i to a finite set of cluster labels c . The most basic approaches include k -means and

gaussian mixture models (GMMs). GMMs cluster data based on maximum likelihood fitting of a probabilistic generative model. k -means can be thought of as either directly clustering data based on geometric (often Euclidean) distances between data points or as a special case of GMMs with the assumptions of evenly sampled, symmetric, equal variance components.

The information bottleneck (IB) is an information-theoretic approach to clustering data X that optimizes cluster labels T to preserve information about a third target variable of interest Y . The resulting (soft) clustering groups data points based on the similarity in their conditional distributions over the target variable through the KL divergence, $\text{KL}[p(y | x_i) | p(y | x_j)]$. An IB clustering problem is fully specified by the joint distribution $P(X, Y)$ and the trade-off parameter β quantifying the relative preference for fewer clusters and more informative ones.

At first glance, it is not obvious how to use this approach to cluster geometric data, where the input is a set of data point locations $\{\mathbf{x}_i\}_{i=1}^N$. For example, what is the target variable Y that our clusters should retain information about? What should $P(X, Y)$ be? And how should one choose the trade-off parameter β ?

Still, Bialek, and Bottou (2004) were the first to attempt to do geometric clustering with IB and claimed an equivalence (in the large data limit) between IB and k -means. Unfortunately, while much of their approach is correct, it contained some fundamental errors that nullify the main results. In the next section, we describe those errors and how to correct them. Essentially, their approach did not properly translate geometric information into a form that could be used correctly by an information-theoretic algorithm.

In addition to fixing this issue, we also choose to use a recently introduced variant of the information bottleneck, the *deterministic information bottleneck* (DIB) (Strouse & Schwab, 2017). We make this choice due to the different way in which IB and DIB use the number of clusters provided to them. IB is known to use all of the clusters it has access to, and thus clustering with IB requires a search over the number of clusters N_c , as well as the parsimony-informativeness trade-off parameter β (Slonim, Atwal, Tkacik, & Bialek, 2005). DIB, on the other hand, has a built-in preference for using as few clusters as it can and thus requires only a parameter search over β . Moreover, DIB's ability to select the number of clusters to use for a given β leads to a intuitive model selection heuristic based on the robustness of a clustering result across β that we show can recover the generative number of clusters in many cases.

In the next section, we more formally define the geometric clustering problem, the IB approach of Still et al. (2004), and our own DIB approach. In section 3, we show that our DIB approach to geometric clustering behaves intuitively and is able to recover the generative number of clusters with only a single free parameter (the data-smoothing scale s). In section 4, we discuss the relationship between our approach and k -means/GMMs, showing that in particular limits, clustering with DIB and IB is equivalent

to k -means/EM fitting of a GMM with hard and soft assignments, respectively. Our approach thus provides a novel information-theoretic approach to geometric clustering, as well as an information-theoretic perspective on these classic clustering methods.

2 Geometric Clustering with the (Deterministic) Information Bottleneck

In a geometric clustering problem, we are given a set of N observed data points $\{\mathbf{x}_i\}_{i=1:N}$ and asked to provide a weighting $q(c | i)$ that categorizes data points into (possibly multiple) clusters such that data points “near” one another are in the same cluster. The definition of *near* varies by algorithm. For k -means, for example, points in a cluster are closer to their own cluster mean than to any other cluster mean.

In an information bottleneck (IB) problem, we are given a joint distribution $P(X, Y)$ and asked to provide a mapping $q(t | x)$ such that T contains the relevant information in X for predicting Y . This goal is embodied by the information-theoretic optimization problem,

$$q_{\text{IB}}^*(t | x) = \underset{q(t|x)}{\operatorname{argmin}} I(X, T) - \beta I(T, Y), \quad (2.1)$$

subject to the Markov constraint $T \leftrightarrow X \leftrightarrow Y$. β is a free parameter that allows for setting the desired balance between the compression encouraged by the first term and the relevance encouraged by the second. At small values, we throw away most of X in favor of a succinct representation for T , while for large values of β , we retain nearly all the information that X has about Y .

This approach of squeezing information through a latent variable bottleneck might remind some readers of a variational autoencoder (VAE) (Kingma & Welling, 2013), and indeed IB has a close relationship with VAEs. As Alemi, Fischer, Dillon, & Murphy (2016) pointed out, a variational version of IB can essentially be seen as the supervised generalization of a VAE, which is typically an unsupervised algorithm.

We are interested in performing geometric clustering with the information bottleneck. For the purposes of this letter, we focus on the deterministic information bottleneck (DIB) (Strouse & Schwab, 2017). We do this because the DIB’s cost function more directly encourages the use of as few clusters as possible, so initialized with n_c^{\max} clusters, it will typically converge to a solution with far fewer. Thus, it has a form of model selection built in that will prove useful for geometric clustering (Strouse & Schwab, 2017). IB tends to use all n_c^{\max} clusters and thus requires an additional search over this parameter (Slonim et al., 2005). DIB also differs from IB in that it leads to a hard clustering instead of a soft clustering.

Formally, the DIB setup is identical to that of IB except that the mutual information term $I(X; T)$ in the cost functional is replaced with the entropy $H(T)$:

$$q_{\text{DIB}}^*(t | x) = \underset{q(t|x)}{\operatorname{argmin}} H(T) - \beta I(T, Y). \quad (2.2)$$

This change to the cost functional leads to a hard clustering with the form (Strouse & Schwab, 2017)

$$q_{\text{DIB}}^*(t | x) = \delta(t - t^*(x)), \quad (2.3)$$

$$t^*(x) \equiv \underset{t}{\operatorname{argmax}} \log q(t) - \beta \text{KL}[p(y | x) | q(y | t)], \quad (2.4)$$

$$q(t) = \sum_x q(t | x) p(x), \quad (2.5)$$

$$q(y | t) = \frac{1}{q(t)} \sum_x q(t | x) p(x) p(y | x), \quad (2.6)$$

where equations 2.3 to 2.6 are to be iterated to convergence from some initialization. The IB solution (Tishby, Pereira, & Bialek, 1999) simply replaces the first two equations with

$$q_{\text{IB}}^*(t | x) = \frac{q(t)}{Z(x, \beta)} \exp[-\beta \text{KL}[p(y | x) | q(y | t)]], \quad (2.7)$$

which can be seen as replacing the argmax in DIB with a soft max.

The (D)IB is referred to as a distributional clustering algorithm (Slonim & Tishby, 2001) due to the KL divergence term $d(x, t) = \text{KL}[p(y | x) | q(y | t)]$, which can be seen as measuring how similar the data point conditional distribution $p(y | x)$ is to the cluster conditional, or mixture of data point conditionals, $q(y | t) = \sum_x q(x | t) p(y | x)$. That is, a candidate point x' will be assigned to a cluster based on how similar its conditional $p(y | x')$ is to the conditionals $p(y | x)$ for the data points x that make up that cluster. Thus, both DIB and IB cluster data points based on the conditionals $p(y | x)$.

To apply (D)IB to a geometric clustering problem, we must choose how to map the geometric clustering data set $\{\mathbf{x}_i\}_{i=1:N}$ to an appropriate IB data set $P(X, Y)$. First, what should X and Y be? Since X refers to the data being clustered by IB, we choose that to be the data point index i . As for the target variable Y that we wish to maintain information about, it seems reasonable to choose the data point location \mathbf{x} (though we discuss alternative choices later). Thus, we want to cluster data indices i into cluster indices c in a way that maintains as much information about the location \mathbf{x} as possible (Still et al., 2004).

How should we choose the joint distribution $p(i, \mathbf{x}) = p(\mathbf{x} | i) p(i)$? At first glance, one might choose $p(\mathbf{x} | i) = \delta_{\mathbf{x}, \mathbf{x}_i}$, since data point i was observed at location \mathbf{x}_i . The reason not to do this lies with the fact that (D)IB is a distributional clustering algorithm, as discussed. Data points are compared to one another through their conditionals $p(\mathbf{x} | i)$, and with the choice of a delta function, there will be no overlap unless two data points are on top of one another. That is, choosing $p(\mathbf{x} | i) = \delta_{\mathbf{x}, \mathbf{x}_i}$ leads to a KL divergence that is either infinite for data points at different locations or zero for data points that lie exactly on top of one another: $\text{KL}[p(\mathbf{x} | i) | p(\mathbf{x} | j)] = \delta_{\mathbf{x}_i, \mathbf{x}_j}$. Trivially, the resulting clustering would assign each data point to its own cluster, grouping only data points that are identical. Put another way, all relational information in an IB problem lies in the joint distribution $P(X, Y)$. If one wants to perform geometric clustering with an IB approach, then geometric information must somehow be injected into that joint distribution, and a series of delta functions does not do that. A previous attempt at linking IB and k -means made this mistake (Still et al., 2004). Subsequent algebraic errors were tantamount to incorrectly introducing geometric information into IB, precisely in the way that such geometric information appears in k -means and resulting in an algorithm that is not IB. We describe these errors in more detail in the appendix.

Based on the problems identified with using delta functions, a better choice for the conditionals is something spatially extended, such as

$$p(\mathbf{x} | i) \propto \exp\left[-\frac{1}{2s^2}d(\mathbf{x}, \mathbf{x}_i)\right], \quad (2.8)$$

where s sets the geometric scale or units of distance and d is a distance metric, such as the Euclidean distance $d(\mathbf{x}, \mathbf{x}_i) = \|\mathbf{x} - \mathbf{x}_i\|^2$. If we indeed use the Euclidean distance, then $p(\mathbf{x} | i)$ will be (symmetric) gaussian (with variance s^2), and this corresponds to gaussian smoothing our data. In any case, the obvious choice for the marginal is $p(i) = \frac{1}{N}$, where N is the number of data points, unless one has a reason a priori to favor certain data points over others. These choices for $p(i)$ and $p(\mathbf{x} | i)$ determine completely our data set $p(i, \mathbf{x}) = p(\mathbf{x} | i) p(i)$. Figure 1 contains an illustration of this data smoothing procedure. We explore the effect of the choice of smoothing scale s throughout this letter.

With the above choices, we have a fully specified DIB formulation of a geometric clustering problem. Using our above notational choices, the equations for the n th step in the iterative DIB solution are (Strouse & Schwab, 2017)

$$q^{(n)}(c | i) = \delta\left(c - c^{*(n)}(i)\right), \quad (2.9)$$

$$c^{*(n)}(i) \equiv \underset{c}{\operatorname{argmax}} \log q^{(n-1)}(c) - \beta \text{KL}\left[p(\mathbf{x} | i) | q^{(n-1)}(\mathbf{x} | c)\right], \quad (2.10)$$

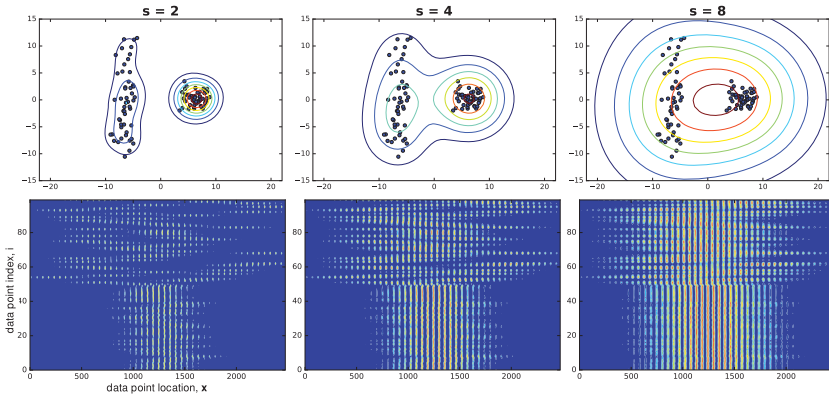


Figure 1: Illustration of data-smoothing procedure. Example data set with one symmetric and one skew cluster. (Top) Scatterplot of data points with smoothed probability distribution overlaid. (Bottom) Heat map of the joint distribution $P(i, \mathbf{x})$ fed into DIB. The two spatial dimensions in the top row are binned and concatenated into a single dimension (on the horizontal axis) in the bottom row, which is the source of the striations.

$$q^{(n)}(c) = \frac{n_c^{(n)}}{N}, \tag{2.11}$$

$$q^{(n)}(\mathbf{x} | c) = \sum_i q^{(n)}(i | c) p(\mathbf{x} | i) = \frac{1}{n_c^{(n)}} \sum_{i \in S_c^{(n)}} p(\mathbf{x} | i), \tag{2.12}$$

where $S_c^{(n)} \equiv \{i : c^{*(n)}(i) = c\}$ is the set of indices of data points assigned to cluster c at step n , and $n_c^{(n)} \equiv |S_c^{(n)}|$ is the number of data points assigned to cluster c at step n . This process is summarized in algorithm 1.

Note that this solution contains β as a free parameter. It allows us to set our preference between solutions with fewer clusters and those that retain more spatial information. It is common in the IB literature to run the algorithm for multiple values of β and plot the collection of solutions in the information plane with the relevance term $I(Y; T)$ on the y -axis and the compression term $I(X; T)$ on the x -axis (Palmer, Marre, Berry, & Bialek, 2015; Creutzig, Globerson, & Tishby, 2009; Chechik, Globerson, Tishby, & Weiss, 2005; Slonim et al., 2005; Still & Bialek, 2004; Tishby & Zaslavsky, 2015; Rubin, Ulanovsky, Nelken, & Tishby, 2016; Strouse & Schwab, 2017; Shwartz-Ziv & Tishby, 2017). The natural such plane for the DIB is with the relevance term $I(Y; T)$ on the y -axis and its compression term $H(T)$ on the x -axis (Strouse & Schwab, 2017). The curve drawn out by (D)IB solutions in the information plane can be viewed as a Pareto-optimal boundary of how much relevant information can be extracted about Y given a fixed amount

Algorithm 1 Geometric clustering with DIB.

Input: data $\{\mathbf{x}_i\}_{i=1:N}$, tradeoff parameter β , smoothing scale s
 Initialize cluster assignments $c^{*(0)}(i)$
 Initialize cluster marginals $q^{(0)}(c)$ using eqn 11
 Initialize cluster conditionals $q^{(0)}(\mathbf{x} | c)$ using eqn 12
 Initialize step count $n = 0$
while not converged do
 $n += 1$
 Update cluster assignments $c^{*(n)}(i)$ using eqn 9
 Update cluster marginals $q^{(n)}(c)$ using eqn 11
 Update cluster conditionals $q^{(n)}(\mathbf{x} | c)$ using eqn 12
end while

of information about X (IB) or representational capacity by T (DIB) (Strouse & Schwab, 2017). Solutions lying below this curve are of course suboptimal, but a priori, the (D)IB formalism does not tell us how to select a single solution from the family of solutions lying on the (D)IB boundary. Intuitively however, when faced with a boundary of Pareto optimality, if we must pick one solution, it is best to choose one at the knee of the curve. Quantitatively, the knee of the curve is the point where the curve has its maximum magnitude second derivative. In the most extreme case, the second derivative is infinite when there is a kink in the curve, and thus the largest kinks might correspond to solutions of particular interest. In our case, since the slope of the (D)IB curve at any given solution is β^{-1} (which can be read off from the cost functionals), kinks indicate solutions that are valid over a wide range of β . So large kinks additionally correspond to solutions robust to model hyperparameters, in the sense that they optimize a wide range of (D)IB tradeoffs. Such robust solutions should correspond to “real” structure in the data. Quantitatively, we can measure the size of a kink by the angle θ of the discontinuity it causes in the slope of the curve (see Figure 2 for details). We show in the next section that searches for solutions with large θ result in recovering the generative cluster labels for geometric data, including the correct number of clusters.

Note that this model selection procedure would not be possible if we had chosen to use IB instead of DIB. IB uses all the clusters available to it, regardless of the choice of β . Thus, all solutions on the curve would have the same number of clusters anyway, so any knees or kinks cannot be used to select the number of clusters.

3 Results: Geometric Clustering with DIB

We ran the DIB as described on four geometric clustering data sets, varying the smoothing width s (see equation 2.9) and trade-off parameter β , and measured for each solution the fraction of spatial information extracted




Figure 2: Kinks in DIB information curve as model selection. β_{\min} and β_{\max} are the smallest and largest β at which the solution at the kink is valid. Thus, β_{\min}^{-1} and β_{\max}^{-1} are the slopes of upper and lower dotted lines. The kink angle is then $\theta = \frac{\pi}{2} - \arctan(\beta_{\min}^{-1}) - \arctan(\beta_{\max}^{-1})$. It is a measure of how robust a solution is to the choice of β ; thus, high values of θ indicate solutions of particular interest.

$\tilde{I}(c; \mathbf{x}) = \frac{I(c; \mathbf{x})^1}{I(i; \mathbf{x})}$ and the number of clusters used n_c , as well as the kink angle θ . We iterated the DIB equations above just as in Strouse and Schwab (2017) with one difference. Iterating greedily from some initialization can lead to local minima (the DIB optimization problem is nonconvex). To help overcome suboptimal solutions, upon convergence, we checked whether merging any two clusters would improve the value L of the cost functional in equation 2.2. If so, we chose the merging with the highest such reduction and began the iterative equations again. We repeated this procedure until the algorithm converged and no merging reduced the value of L . We found that these “nonlocal” steps worked well in combination with the greedy “local” improvements of the DIB iterative equations. While not essential to the function of DIB, this improvement in performance produced cleaner information curves with less noise caused by convergence to local minima. Similar to Strouse and Schwab (2017), the automated search over β began with an initial set of values and then iteratively inserted more values where there were large jumps in $H(c)$, $I(c; \mathbf{x})$, or the number of clusters used or where the largest value of β did not lead to a clustering solution capturing nearly all of the available geometric information (i.e., with $I(c; \mathbf{x}) \approx I(i; \mathbf{x})$). (For more details, see our code repository at <https://github.com/djstrouse/information-bottleneck>.)

¹Note that $I(i; \mathbf{x})$ is an upper bound on $I(c; \mathbf{x})$ due to the data processing inequality (Cover & Thomas, 2006), so $\tilde{I}(c; \mathbf{x})$ is indeed the fraction of potential geometric information extracted from the smoothed $P(i, \mathbf{x})$.

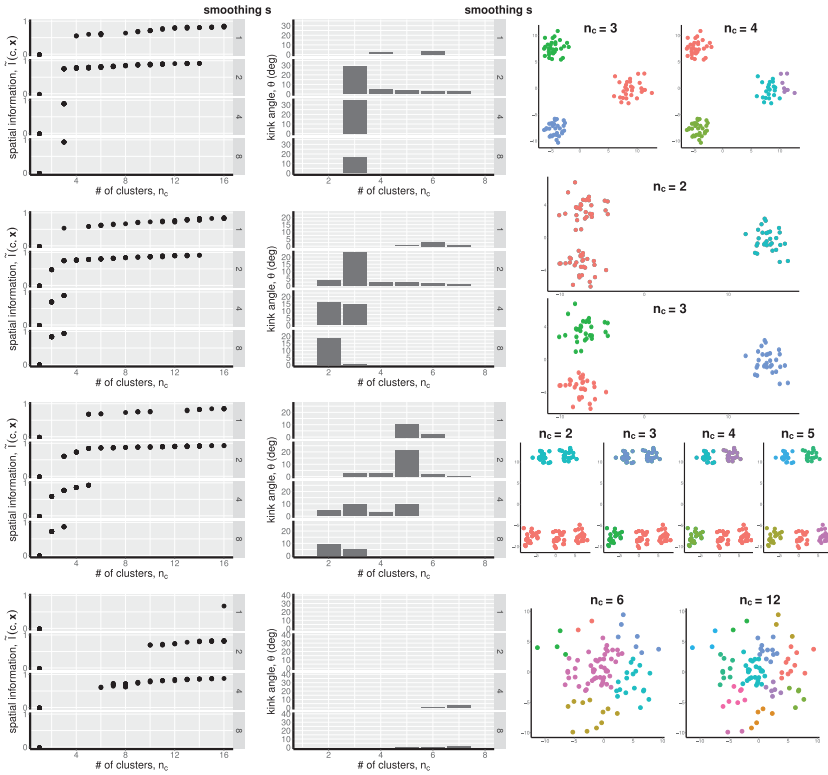


Figure 3: Results: Model selection and clustering with DIB. Results for four data sets. Each row represents a different data set. (Left) Fraction of spatial information extracted, $\tilde{I}(c; \mathbf{x}) = \frac{I(c; \mathbf{x})}{I(i; \mathbf{x})}$, versus number of clusters used, $n_{c'}$, across a variety of smoothing scales, s . (Center) Kink angle θ (of the $I(c; \mathbf{x})$ versus $H(c)$ curve) versus number of clusters used, $n_{c'}$, across a variety of smoothing scales, s . (Right) Example resulting clusters.

Results are shown in Figure 3. Each large row represents a different data set. The left column shows fractional spatial information $\tilde{I}(c; \mathbf{x})$ versus number of clusters used $n_{c'}^2$ stacked by smoothing width s .³ The center column shows the kink angle θ for each cluster number $n_{c'}$, again stacked by smoothing width s . Finally, the right column shows example solutions.

²Note that this is not the same as the n_c in equations 2.11 and 2.12, which was the number of data points assigned to a particular cluster c . Here we are using it to denote the number of clusters with at least one data point assigned to it.

³Note that this is not the same as the information plane curve from Figure 2. While the y -axes are the same (up to the normalization), the x -axes are different.

In general, note that as we increase β , we move right along the plots in the left column, that is, toward higher numbers of clusters n_c and more spatial information $\tilde{I}(c; \mathbf{x})$. Not all values of n_c are present because while varying the implicit parameter β , DIB will not necessarily “choose” to use all possible cluster numbers. For example, for small smoothing width s , most points will not have enough overlap in $p(\mathbf{x} | i)$ with their neighbors to support solutions with few clusters, and for large smoothing width s , local spatial information is thrown out and only solutions with few clusters are possible. More interesting, DIB may retain or drop solutions based on how well they match the structure of the data, as we will discuss for each data set below. Additionally, solutions that well match the structure in the data (e.g., ones with n_c matched to the generative parameters) tend to be especially robust to β , that is, they have a large kink angle θ . Thus, θ can be used to perform model selection. For data sets with structure at multiple scales, the kink angle θ will select different solutions for different values of the smoothing width s . This allows us to investigate structure in a data set at a particular scale of our choosing. We now turn to the individual data sets.

The first data set (top row) consists of three equally spaced, equally sampled symmetric gaussian clusters (see the solutions in the right column). We see that the three-cluster solution stands out in several ways. First, it is robust to spatial scale s . Second, the three-cluster solution extracts nearly all of the available spatial information; solutions with $n_c \geq 4$ extract little extra $\tilde{I}(c; \mathbf{x})$. Third and perhaps most salient, the three-cluster solution has by far the largest value of kink angle θ across a wide range of smoothing scales. In the right column, we show examples of three- and four-cluster solutions. Note that while all three-cluster solutions look exactly like this one, the four-cluster solutions vary in how they chop one true cluster into two.

The second data set (second row) consists of three more equally sampled symmetric gaussian clusters, but this time not equally spaced: two are much closer to one another than the third. This is a data set with multiple scales present; thus, we should expect that the number of clusters picked out by any model selection procedure (e.g., kink angle) should depend on the spatial scale of interest. Indeed, we see that to be true. The three-cluster solution is present for all smoothing widths shown, but is selected out as the best solution only by kink angle for intermediate smoothing widths ($s = 2$). For large smoothing widths ($s = 8$), we see that the two-cluster solution is chosen as best. For smoothing widths in between ($s = 4$), the two- and three-cluster solutions are roughly equally valid. In terms of spatial information, the two- and three-cluster solutions are also prominent, with both transitions from $n_c = 1 \rightarrow 2$ and $n_c = 2 \rightarrow 3$ providing significant improvement in $\tilde{I}(c; \mathbf{x})$ (but little improvement for more finely grained clusterings).

The third data set (third row) features even more multiscale structure, with five symmetric, equally sampled gaussians, again with unequal spacing. Sensible solutions exist for $n_c = 2-5$, and this can be seen by the more

gradual rise of the fractional spatial information $\tilde{I}(c; \mathbf{x})$ with n_c in that regime. We also again see a transition in the model selection by θ from the five-cluster solution at small smoothing widths ($s = 1, 2$) and the two-cluster solution at larger smoothing widths ($s = 8$), with intermediate n_c favoring those and intermediate solutions. Example clusters for $n_c = 2-5$ are shown at the right.

Finally, we wanted to ensure that DIB and our model selection procedure would not hallucinate structure where there is none, so we applied it to a single gaussian blob, with the hope that no solution with $n_c > 1$ would stand out and prove robust to β . As can be seen in the fourth row of Figure 3, that is indeed true. No solution at any smoothing width had a particularly high kink angle θ , and no solution remained at the knee of the $\tilde{I}(c; \mathbf{x})$ versus the n_c curve across a wide range of smoothing widths.

Overall, these results suggest that DIB on smoothed data is able to recover generative geometric structure at multiple scales, using built-in model selection procedures based on identifying robust, spatially informative solutions.

4 Relationship Between (D)IB and GMMs & k -Means

It is natural to wonder how the algorithm we introduce here, clustering with DIB, relates to classic approaches to clustering, including GMMs and k -means. We now establish the following equivalence: when the smoothing scale s is small, $\beta = 1$, and $q(\mathbf{x} | c)$ is approximated as a gaussian $r(\mathbf{x} | c)$ whose parameters are chosen to minimize $\text{KL}[p(\mathbf{x} | c) | r(\mathbf{x} | c)]$, DIB and IB correspond to EM-fitting of a GMM with hard and soft assignments, respectively. When s is small and $r(\mathbf{x} | c)$ is chosen to be an isotropic gaussian with fixed variance across clusters, DIB and IB correspond to hard and soft k -means, respectively, with a logarithmic “cluster size bonus” weighted by β^{-1} . In the $\beta \rightarrow \infty$ limit, the effect of the cluster size bonus vanishes and the correspondence with hard and soft k -means is exact. Thus, clustering with (D)IB can be viewed as a generalization of these approaches.

We begin by establishing the correspondence between DIB and the E-step of fitting a GMM. Consider the KL divergence $\text{KL}[p(\mathbf{x} | i) | q(\mathbf{x} | c)]$ that (D)IB uses to cluster data points. When the smoothing scale s is chosen to be small relative to the scale of $q(\mathbf{x} | c)$, then we have

$$\text{KL}[p(\mathbf{x} | i) | q(\mathbf{x} | c)] = - \int p(\mathbf{x} | i) \log q(\mathbf{x} | c) d\mathbf{x} - H[p(\mathbf{x} | i)], \quad (4.1)$$

$$\approx - \log q(\mathbf{x}_i | c) \int p(\mathbf{x} | i) d\mathbf{x} - H[p(\mathbf{x} | i)], \quad (4.2)$$

$$= - \log q(\mathbf{x}_i | c) - H[p(\mathbf{x} | i)], \quad (4.3)$$

where we have used the assumption about the scale of s in moving from the first to second line. Since $H[p(\mathbf{x} | i)]$ is independent of the cluster assignments, minimizing $\text{KL}[p(\mathbf{x} | i) | q(\mathbf{x} | c)]$ with respect to the cluster assignments is then equivalent to maximizing $\log q(\mathbf{x}_i | c)$, that is choosing a maximum likelihood assignment of points to clusters. Thus equation 2.10 becomes

$$c^*(i) = \underset{c}{\operatorname{argmax}} \log p(c) + \beta \log p(\mathbf{x}_i | c), \quad (4.4)$$

$$= \underset{c}{\operatorname{argmax}} \log p(c)^{1/\beta} + \log p(\mathbf{x}_i | c). \quad (4.5)$$

For $\beta = 1$, the two log probabilities combine and lead to a maximum a posteriori (MAP) assignment of points to clusters

$$c^*(i) = \underset{c}{\operatorname{argmax}} \log p(\mathbf{x}_i, c) = \underset{c}{\operatorname{argmax}} p(c | \mathbf{x}_i). \quad (4.6)$$

For $1 < \beta < \infty$, the effect of β is to "soften" the prior $p(c)$ (equation 4.5), leading to less aggressive cluster consolidation.

Of course, if we use the exact $q(\mathbf{x} | c)$ defined in equation 2.12, then the scales of $p(\mathbf{x} | i)$ and $q(\mathbf{x} | c)$ are similar, and so our assumption in this section is not valid. In order for it to be valid, we need to replace the exact $q(\mathbf{x} | c)$ with an assumed parametric form that leads to further smoothing.

If we choose to replace $q(\mathbf{x} | c)$ with a gaussian approximation $r(\mathbf{x} | c) = \mathcal{N}(\mathbf{x} | \mu_c, \Sigma_c)$, then equation 4.6 corresponds to the E-step in EM fitting of a GMM (Bishop, 2006). Note that ideally we would like for it to be true that $\text{KL}[p(\mathbf{x} | i) | r(\mathbf{x} | c)] \geq \text{KL}[p(\mathbf{x} | i) | q(\mathbf{x} | c)]$ so that the replacement of $q(\mathbf{x} | c)$ by $r(\mathbf{x} | c)$ leads to us maximizing a lower bound on our original objective (i.e., that which is maximized in equation 2.10), however this is not generically true and $\text{KL}[p(\mathbf{x} | i) | r(\mathbf{x} | c)]$ might be smaller or larger than $\text{KL}[p(\mathbf{x} | i) | q(\mathbf{x} | c)]$.

The results in this section are only valid for a "small" smoothing scale s , so let us now understand what that means in the particular case of gaussian $r(\mathbf{x} | c)$. Consider the KL divergence in the assignment step (equation 2.10), which in this case has a simple expression

$$\begin{aligned} & \text{KL}[p(\mathbf{x} | i) | r(\mathbf{x} | c)] \\ & \propto \frac{s^2}{\text{tr}(\Sigma_c)} + (\mu_c - \mathbf{x}_i)^T \Sigma_c^{-1} (\mu_c - \mathbf{x}_i) + \log \det \Sigma_c + k, \end{aligned} \quad (4.7)$$

where k denotes terms not dependent on the assignment of points to clusters, and thus irrelevant for the objective. Compare to the maximum likelihood objective. The negative log likelihood of \mathbf{x}_i under $r(\mathbf{x} | c)$ is

$$-\log \mathcal{N}(\mathbf{x}_i | \mu_c, \Sigma_c) \propto (\mu_c - \mathbf{x}_i)^T \Sigma_c^{-1} (\mu_c - \mathbf{x}_i) + \log \det(\Sigma_c) + k, \quad (4.8)$$

where k again denotes terms independent of the assignment of points to clusters, and thus ignorable. Note that when $s^2 \ll \text{tr}(\Sigma_c)$, the last two equations are the same, and thus the DIB cluster assignments correspond to maximum likelihood assignments. Thus, “small” s means $s^2 \ll \text{tr}(\Sigma_c)$ in this case. Of course, we don’t know $\text{tr}(\Sigma_c)$ until after we cluster our data, but it is set by the natural length scales in the data, so we can take it to mean that s needs to be small compared to those.

That establishes the correspondence for the E-step of EM fitting of a GMM, but what about the M-step? Note that we haven’t yet specified how to fit the approximation $r(\mathbf{x} | c) \approx q(\mathbf{x} | c)$. One reasonable way that appears often in the variational inference literature (e.g., Kingma and Welling, 2014) is to choose the parameters of $r(\mathbf{x} | c)$ (μ_c and Σ_c) that minimize $\text{KL}[p(\mathbf{x} | c) | r(\mathbf{x} | c)]$. We choose this direction of the KL divergence because it encourages a “mean-seeking” approximation of $p(\mathbf{x} | c)$ that tries better to approximate the full distribution than the other, “mode-seeking” direction. While this is again a generally intractable KL divergence between a mixture of gaussians and a gaussian, fortunately in the $s^2 \ll \text{tr}(\Sigma_c)$ limit that we consider, it simplifies to

$$\begin{aligned} & \text{KL}[p(\mathbf{x} | c) | r(\mathbf{x} | c)] \\ &= - \int p(\mathbf{x} | c) \log r(\mathbf{x} | c) d\mathbf{x} - H[p(\mathbf{x} | c)] \end{aligned} \quad (4.9)$$

$$= - \frac{1}{n_c} \int \sum_{i \in S_c} \mathcal{N}(\mathbf{x}; \mathbf{x}_i, s^2) \log r(\mathbf{x} | c) d\mathbf{x} - H[p(\mathbf{x} | c)] \quad (4.10)$$

$$\approx - \frac{1}{n_c} \sum_{i \in S_c} \log r(\mathbf{x}_i | c) \int \mathcal{N}(\mathbf{x}; \mathbf{x}_i, s^2) d\mathbf{x} - H[p(\mathbf{x} | c)] \quad (4.11)$$

$$= - \frac{1}{n_c} \sum_{i \in S_c} \log r(\mathbf{x}_i | c) - H[p(\mathbf{x} | c)], \quad (4.12)$$

where we move from the second to third line using the small s approximation (so that $r(\mathbf{x} | c) \approx r(\mathbf{x}_i | c)$ in the region of \mathbf{x} where the bulk of $\mathcal{N}(\mathbf{x}; \mathbf{x}_i, s^2)$ is). Minimizing equation 4.12 with respect to μ_c and Σ_c again corresponds to maximum likelihood assignments, this time of the model parameters rather than cluster assignments. This corresponds to the M-step of EM fitting of a GMM (Bishop, 2006).

Thus, for $\beta = 1$, small s , and a gaussian approximation of $p(\mathbf{x} | c)$ (with parameters chosen to minimize the KL divergence in equation 4.9), clustering with DIB is equivalent to EM fitting of a GMM with hard assignments

(of data points to clusters). For $\beta > 1$, the effect of the cluster prior $p(c)$ is muted; that is, it is replaced with $p(c)^{1/\beta}$.

If we set all cluster conditional approximations to have the same isotropic covariance $\Sigma_c = \text{diag}(\sigma^2)$, then $c^*(i)$ becomes (plugging equation 4.8 into equation 4.5)

$$c^*(i) = \operatorname{argmax}_c \frac{\sigma^2}{\beta} \log p(c) - \|\mathbf{x}_i - \mu_c\|^2 \quad (4.13)$$

$$= \operatorname{argmax}_c \frac{\sigma^2}{\beta} \log n_c - \|\mathbf{x}_i - \mu_c\|^2, \quad (4.14)$$

which corresponds to (hard) k -means with a cluster size bonus $\log n_c$ (where $n_c \equiv |S_c|$ is the number of points assigned to cluster c , as introduced in section 2). In the $\beta \rightarrow \infty$ limit, the $\log n_c$ term can be ignored and the correspondence with (hard) k -means is exact.

To see the correspondence between GMMs/ k -means and IB, consider that IB can be viewed as DIB with the hard max replaced by a soft max (see equation 2.7). Thus, the same correspondences we drew between DIB and GMMs/ k -means with *hard* assignments hold for IB and GMMs/ k -means with *soft* assignments.

The correspondence between clustering with (D)IB and GMMs yields new interpretations of both. From this perspective, clustering with (D)IB can be viewed as a generalization of GMMs that 1) uses a more flexible, nonparametric representation of the clusters, 2) includes an extra parameter β for controlling the tradeoff between the prior and likelihood, and 3) includes an extra parameter s for setting the length scale of interest. In the other direction, GMMs can be viewed as mapping data points to cluster labels that maximally preserve spatial information.

This is not the first correspondence between IB in a particular setting and another probabilistic model. In the discrete setting, IB has been shown to be related to EM fitting of a multinomial mixture model (Slonim & Weiss, 2003). In the time series setting (where $X = x_t$ and $Y = x_{t+1}$), IB is related to canonical correlation analysis (Creutzig et al., 2009), and therefore linear gaussian models (Bach & Jordan, 2006) and slow feature analysis (Turner & Sahani, 2007). Under a variational approximation, IB applied to unsupervised learning is related to a variational autoencoder (VAE) (Alemi et al., 2016; Higgins et al., 2017; Kingma & Welling, 2014).

5 Discussion

We have shown in this letter how to use the formalism of the information bottleneck to perform geometric clustering. A previous paper (Still et al., 2004) claimed to contribute similarly; however, for the reasons discussed in

section 2 and the appendix, their approach contained fundamental flaws. We amend and improve on that paper in four ways. First, we show to fix the errors they made in their problem setup (with the data preparation). Second, we argue for using DIB over IB in this setting for its preference for using as few clusters as it can. Third, we introduce a novel form of model selection for the number of clusters based on discontinuities (or kinks) in the slope of the DIB curve, which indicate solutions that are robust across the DIB trade-off parameter β . We show that this information-based model selection criterion allows us to correctly recover generative structure in the data at multiple spatial scales. Finally, we establish the correct correspondence between clustering with (D)IB and k -means/GMMs, thus providing both a generalization and information-theoretic interpretation of these classic approaches.

We have introduced one way of doing geometric clustering with the information bottleneck, but we think it opens avenues for other ways as well. First, the uniform smoothing we performed could be generalized in a number of ways to better exploit local geometry and better estimate the true generative distribution of the data. For example, one could do gaussian smoothing with the mean centered on each data point but the covariance estimated by the sample covariance of neighboring data points around that mean. Indeed, our early experiments with this alternative suggest it may be useful for certain data sets. Second, while choosing spatial location as the relevant variable for DIB to preserve information about seems to be the obvious first choice to investigate, other options might prove interesting. For example, preserving information about the identity of neighbors, if carefully formulated, might make fewer implicit assumptions about the shape of the generative distribution and enable the extension of our approach to a wider range of data sets.

Scaling the approach introduced here to higher-dimensional data sets is nontrivial because the tabular representation used in the original IB (Tishby et al., 1999) and DIB (Strouse & Schwab, 2017) algorithms leads to an exponential scaling with the number of dimensions. Recently, however, Alemi et al. (2016) introduced a variational version of IB in which one parameterizes the encoder $q(t | x)$ (and “decoder” $q(y | t)$) with a function approximator (e.g., a deep neural network). This has the advantage of allowing scaling to much larger data sets. Moreover, the choice of parameterization often implies a smoothness constraint on the data, relieving the problem encountered above of needing to smooth the data. It would be interesting to develop a variational version of DIB, which could then be used to perform information-theoretic clustering as we have done here, but on larger problems and perhaps with no need for data smoothing.

Appendix: Errors in Still et al. (2004)

A previous attempt was made to draw a connection between IB and k -means by Still et al. (2004). Even before reviewing the algebraic errors that lead their result to break down, there are two intuitive reasons why such a claim is unlikely to be true. First, IB is a soft clustering algorithm, and k -means is a hard clustering algorithm. Second, the authors made the choice not to smooth the data and to set $p(\mathbf{x} | i) = \delta_{\mathbf{x}x_i}$. As discussed in section 2, (D)IB clusters data points based on these conditionals, and delta functions trivially overlap only when they are identical.

The primary algebraic mistake appears just after equation 2.14, in the claim that $p_n(\mathbf{x} | c) \propto p_{n-1}(\mathbf{x} | c)^{1/\lambda}$. Combining the previous two claims in that proof, we obtain

$$p_n(\mathbf{x} | c) = \frac{1}{N} \sum_i \frac{\delta_{\mathbf{x}x_i}}{Z_n(i, \lambda)} p_{n-1}(x_i | c)^{1/\lambda}. \quad (\text{A.1})$$

Certainly this does not imply that $p_n(\mathbf{x} | c) \propto p_{n-1}(\mathbf{x} | c)^{1/\lambda}$ everywhere, because of the $\delta_{\mathbf{x}x_i}$ factor, which picks out only a finite number of points.

One might wonder why, with these mistakes, the authors still obtain an algorithm that looks and performs like k -means. The reason is that their sequence of mistakes leads to the result in equation 2.15 that effectively assumes that IB has access to geometric information it should not: namely, the cluster centers at step n . Since these are exactly what k -means uses to assign points to clusters, it is not surprising that the behavior then resembles k -means.

Acknowledgments

We thank Léon Bottou and Arthur Szlam for helpful discussions. We also acknowledge financial support from NIH K25 GM098875 (Schwab) and the Hertz Foundation (Strouse).

References

- Alemi, A. A., Fischer, I., Dillon, J. V., & Murphy, K. (2016). Deep variational information bottleneck. *CoRR*. abs/1612.00410.
- Bach, F. R., & Jordan, M. I. (2006). *A probabilistic interpretation of canonical correlation analysis* (Tech. Rep.). University of California, Berkeley.
- Bishop, C. M. (2006). *Pattern recognition and machine learning*. Berlin: Springer-Verlag.
- Chechik, G., Globerson, A., Tishby, N., & Weiss, Y. (2005). Information bottleneck for gaussian variables. *Journal of Machine Learning Research*, 6, 165–188.
- Cover, T. M., & Thomas, J. A. (2006). *Elements of information theory*. New York: Wiley-Interscience.

- Creutzig, F., Globerson, A., & Tishby, N. (2009). Past-future information bottleneck in dynamical systems. *Physical Review E*, 79(4), 041925–5.
- Higgins, I., Sonnerat, N., Matthey, L., Pal, A., Burgess, C. P., Botvinick, M., Hassabis, D., & Lerchner, A. (2017). *SCAN: Learning abstract hierarchical compositional visual concepts*. International Conference on Learning Representations (ICLR). <https://arxiv.org/abs/1707.03389v1>
- Kingma, D. P., & Welling, M. (2013). *Auto-encoding variational Bayes*. <https://arxiv.org/abs/1312.6114>
- LeCun, Y. (2016). *Predictive learning*. Keynote address at the 2016 Conference on Neural Information Processing Systems, Barcelona.
- MacKay, D. J. C. (2002). *Information theory, inference, and learning algorithms*. Cambridge: Cambridge University Press.
- Palmer, S. E., Marre, O., Berry II, M. J., & Bialek, W. (2015). Predictive information in a sensory population. *Proceedings of the National Academy of Sciences*, 112(22), 6908–6913.
- Rubin, J., Ulanovsky, N., Nelken, I., & Tishby, N. (2016). The representation of prediction error in auditory cortex. *PLoS Computational Biology*, 12(8), e1005058–28.
- Shwartz-Ziv, R., & Tishby, N. (2017). *Opening the black box of deep neural networks via information*. <https://arxiv.org/abs/1703.00810>
- Slonim, N., Atwal, G., Tkacik, G., & Bialek, W. (2005). Information-based clustering. *Proceedings of the National Academy of Sciences*, 102(51), 18297–18302.
- Slonim, N., & Tishby, N. (2001). The power of word clusters for text classification. In *Proceedings of the European Colloquium on Information Retrieval Research*.
- Slonim, N., & Weiss, Y. (2003). Maximum likelihood and the information bottleneck. In S. Becker, S. Thrun, & K. Obermayer (Eds.), *Advances in neural information processing systems*, 15 (pp. 351–358). Cambridge, MA: MIT Press.
- Still, S., & Bialek, W. (2004). How many clusters an information-theoretic perspective. *Neural Computation*, 16(12), 2483–2506.
- Still, S., Bialek, W., & Bottou, L. (2004). Geometric clustering using the information bottleneck method. In S. Thrun, L. K. Saul, & B. Schölkopf (Eds.), *Advances in neural information processing systems*, 16 (pp. 1165–1172). Cambridge, MA: MIT Press.
- Strouse, D., & Schwab, D. J. (2017). The deterministic information bottleneck. *Neural Computation*, 29, 1611–1630.
- Tishby, N., Pereira, F. C., & Bialek, W. (1999). The information bottleneck method. In *Proceedings of the 37th Allerton Conference on Communication, Control, and Computing* (pp. 368–377).
- Tishby, N., & Zaslavsky, N. (2015). *Deep learning and the information bottleneck principle*. <https://arXiv.org/abs/1503.02406>
- Turner, R. E., & Sahani, M. (2007). A maximum-likelihood interpretation for slow feature analysis. *Neural Computation*, 19(4), 1022–1038.