Complementi di Algoritmi e Strutture Dati

Correlation Clustering

Docente: *Nicolò Cesa-Bianchi* versione 2 aprile 2026

Clustering is a central problem in unsupervised learning. A clustering problem is typically represented by a set of elements together with a notion of similarity (or dissimilarity) between them. When the elements are points in a metric space, dissimilarity can be measured via a distance function. In more general settings, when the elements to be clustered are members of an abstract set V, similarity is defined by an arbitrary symmetric function σ defined on pairs of distinct elements in V.

Correlation Clustering (CC) is a well-known special case where σ is a $\{-1, +1\}$-valued function establishing whether any two distinct elements of V are similar or not. The objective of CC is to cluster the points in V so to minimize the number of errors, where an error is given by any pair of elements having similarity -1 and belonging to the same cluster, or having similarity $+1$ and belonging to different clusters. Importantly, there are no a priori limitations on the number of clusters or their sizes: all partitions of V, including the trivial ones, are valid. Given V and σ, the error achieved by an optimal clustering is known as the *Correlation Clustering index*, denoted by OPT.

Note that OPT $= 0$ implies that V can be perfectly clustered: any two elements in the same cluster have similarity $+1$ and any two elements in different clusters have similarity -1. Since its introduction, CC has attracted a lot of interest and has found numerous applications in entity resolution, image analysis, and social media analysis.

Minimizing the correlation clustering error is hard, and the best efficient algorithm found so far achieves 2OPT (the actual coefficient multiplying OPT is slightly smaller than 2). A very simple and elegant algorithm for approximating CC is KwikCluster. At each round, KwikCluster draws a random pivot π_r from V, queries the similarities between π_r and every other node in V, and creates a cluster C containing π_r and all points u such that $\sigma(\pi_r, u) = +1$. The algorithm then recursively invokes itself on $V \setminus C$. On any instance of CC, KwikCluster achieves an expected error bounded by 3OPT.

Algoritmo 1 KwikCluster

```
Parameters: residual node set V_r
if |V_r| = 0 then RETURN
end if
if |V_r| = 1 then output singleton cluster V_r and RETURN
end if
Draw pivot π_r u.a.r. from V_r
C_r ← {π_r}                                   ▷ Create new cluster and add the pivot to it
C_r ← C_r ∪ {u ∈ V_r : σ(π_r, u) = +1}        ▷ Populate cluster
Output cluster C_r
KwikCluster(V_r \ C_r)                        ▷ Recursive call on the remaining nodes
```

We denote by $V \equiv \{1, \dots, n\}$ the set of input nodes, by $\mathcal{E} \equiv \binom{V}{2}$ the set of all pairs $\{u, v\}$ of distincts nodes in V, and by $\sigma : \mathcal{E} \to \{-1, +1\}$ the binary similarity function. A clustering $\mathcal{C}$ is a partition of V in disjoint clusters $C_i : i = 1, \dots, k$. Given $\mathcal{C}$ and σ, the set $\Gamma_{\mathcal{C}}$ of mistaken edges contains all pairs $\{u, v\}$ such that $\sigma(u, v) = -1$ and u, v belong to same cluster of $\mathcal{C}$ and all pairs $\{u, v\}$ such that $\sigma(u, v) = +1$ and u, v belong to different clusters of $\mathcal{C}$. The cost of $\mathcal{C}$ is $\left|\Gamma_{\mathcal{C}}\right|$. The correlation clustering index is $\mathrm{OPT} = \min_{\mathcal{C}} \left|\Gamma_{\mathcal{C}}\right|$, where the minimum is over all clusterings $\mathcal{C}$.

A triangle is any unordered triple $T = \{u, v, w\} \subseteq V$. We denote by $e = \{u, w\}$ a generic triangle edge; we write $e \subset T$ and $v = T \setminus e$. We say T is a *bad triangle* if the labels $\sigma(u, v), \sigma(u, w), \sigma(v, w)$ are $\{+, +, -\}$ (the order is irrelevant), see the figure below here.

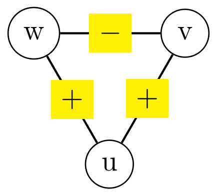

We denote by $\mathcal{T}$ the set of all bad triangles in V and also define $\mathcal{T}(e) \equiv \{T \in \mathcal{T} : e \subset T\}$. It is easy to see that the number of edge-disjoint bad triangles is a lower bound on OPT: no matter how we cluster its nodes, a bad triangle contributes by at least 1 to the total cost of the partition. The following lemma (which we state without proof) shows that also the weighted sum of all bad triangles is a lower bound on OPT, provided the sum of the weights of all bad triangles insisting on any single edge e is at most 1.

Lemma 1 *If $\{\beta_T \geq 0 : T \in \mathcal{T}\}$ is a set of weights on the bad triangles such that $\sum_{T \in \mathcal{T}(e)} \beta_T \leq 1$ for all $e \in \mathcal{E}$, then $\sum_{T \in \mathcal{T}} \beta_T \leq \mathrm{OPT}$.*

We now bound the expected error of KwikCluster. We use V_r to denote the set of remaining nodes at the beginning of the r-th recursive call.

Let Γ_A be the set of mistaken edges for the clustering output by KwikCluster and let $\left|\Gamma_A\right|$ be the cost of this clustering.

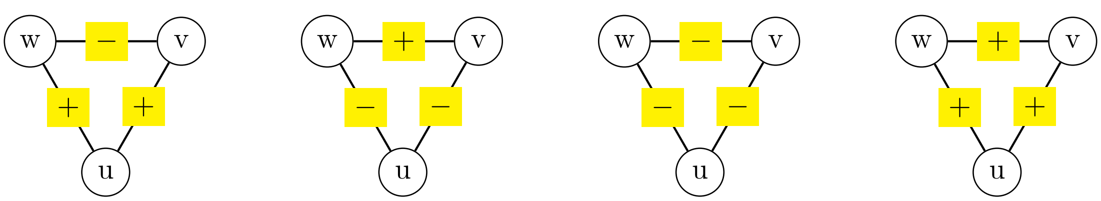

Lemma 2 *KwikCluster a mistake if and only if there exists a bad triangle $T \in \mathcal{T}$ and a recursive call r such that $T \subseteq V_r$ and $\pi_r \in T$.*

Dimostrazione. If there is a round r such that $T \subset V_r$ and $\pi_r \in T$ for some $T \in \mathcal{T}$, then KwikCluster makes a mistake on T (you can directly verify it by case analysis on the choice of the pivot in a bad triangle, see the picture above).

Now, if KwikCluster makes a mistake on some edge $e = \{u, w\}$ in a round r, then it must be that $T = \{u, \pi_r, w\} \subseteq V_r$ and $T \in \mathcal{T}$. Indeed, if $T = \{u, \pi_r, w\} \subseteq V_r$ and $T \notin \mathcal{T}$, then no mistake is made on T (verify it by case analysis on the choice of the pivot for the good triangles in the picture above). $\square$

Lemma 2 implies that at round r we make a mistake on exactly one edge e of every bad triangle T such that $T \subset V_r$ and $\pi_r \in T \setminus e$. Recall that a bad triangle T can be mistaken only once, because when $\pi_r \in T$, π_r gets removed from V_r. Hence, for any realization $\pi_1, \pi_2, \ldots$ of the random pivot sequence,

$$\left|\Gamma_A\right| = \sum_{T \in \mathcal{T}} \mathbb{I}\{(\exists\, r) \,:\, T \subseteq V_r \,\wedge\, \pi_r \in T\} \ .$$

For all $T \in \mathcal{T}$ tet A_T be the event $\big\{(\exists\, r) \,:\, T \subseteq V_r \,\wedge\, \pi_r \in T\big\}$ indicating that T contributes to one mistake.

Note that for any $e = \{u, w\} \in \Gamma_A$ and for any two distinct $T, T' \in \mathcal{T}(e)$, A_T and $A_{T'}$ can not both occur because if e is mistaken due to $\pi_r \in T \setminus e$, then π_r and at least one between u and w are removed, which implies that T' is deleted. Thus, for any e,

$$\sum_{T \in \mathcal{T}(e)} \mathbb{I}\{A_T \,\wedge\, e \in \Gamma_A\} = 1 \ .$$

Taking expectations with respect to the random choice of the pivot sequence,

$$1 = \sum_{T \in \mathcal{T}(e)} \mathbb{P}\big(A_T \,\wedge\, e \in \Gamma_A\big) = \sum_{T \in \mathcal{T}(e)} \mathbb{P}\big(e \in \Gamma_A \mid A_T\big)\mathbb{P}(A_T) = \sum_{T \in \mathcal{T}(e)} \frac{1}{3}\mathbb{P}(A_T) \ .$$

where $\mathbb{P}\big(e \in \Gamma_A \mid A_T\big) = \frac{1}{3}$ holds because, given some r such that $T \subseteq V_r$ and $\pi_r \in T$, e is mistaken only if $\pi_r \in T \setminus e$.

Applying Lemma 1 with $\beta_T = \frac{1}{3}\mathbb{P}(A_T)$ for each $T \in \mathcal{T}$, we get

$$\mathbb{E}\big[|\Gamma_A|\big] = \sum_{T \in \mathcal{T}} \mathbb{P}(A_T) = 3 \sum_{T \in \mathcal{T}} \beta_T \leq 3\mathrm{OPT} \ .$$