

Minimizing Localized Ratio Cut Objectives in Hypergraphs

Nate Veldt
Cornell University
Center for Applied Mathematics
nveldt@cornell.edu

Austin R. Benson
Cornell University
Department of Computer Science
arb@cs.cornell.edu

Jon Kleinberg
Cornell University
Department of Computer Science
kleinberg@cornell.edu

ABSTRACT

Hypergraphs are a useful abstraction for modeling multiway relationships in data, and hypergraph clustering is the task of detecting groups of closely related nodes in such data. *Graph* clustering has been studied extensively, and there are numerous methods for detecting small, localized clusters without having to explore an entire input graph. However, there are only a few specialized approaches for localized clustering in hypergraphs. Here we present a framework for local hypergraph clustering based on minimizing localized ratio cut objectives. Our framework takes an input set of reference nodes in a hypergraph and solves a sequence of hypergraph minimum s - t cut problems in order to identify a nearby well-connected cluster of nodes that overlaps substantially with the input set.

Our methods extend graph-based techniques but are significantly more general and have new output quality guarantees. First, our methods can minimize new generalized notions of hypergraph cuts, which depend on specific configurations of nodes within each hyperedge, rather than just on the number of cut hyperedges. Second, our framework has several attractive theoretical properties in terms of output cluster quality. Most importantly, our algorithm is strongly-local, meaning that its runtime depends only on the size of the input set, and does not need to explore the entire hypergraph to find good local clusters. We use our methodology to effectively identify clusters in hypergraphs of real-world data with millions of nodes, millions of hyperedges, and large average hyperedge size with runtimes ranging between a few seconds and a few minutes.

ACM Reference Format:

Nate Veldt, Austin R. Benson, and Jon Kleinberg. 2020. Minimizing Localized Ratio Cut Objectives in Hypergraphs. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '20)*, August 23–27, 2020, Virtual Event, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3394486.3403222>

1 INTRODUCTION

Graphs are a common mathematical abstraction for modeling pairwise interactions between objects in a dataset. A standard task in graph-based data analysis is to identify well-connected *clusters* of nodes, which share more edges with each other than the rest of the graph [30]. For example, detecting clusters in a graph is used

to identify communities [13], predict class labels in machine learning applications [7], and segment images [18]. Standard models for such clusters are *ratio cut* objectives, which measure the ratio between the number of edges leaving a cluster (the *cut*) and some notion of the cluster's size (e.g., the number of edges or nodes in the cluster); common ratio cut objectives include conductance, sparsest cut, and normalized cut. Ratio cut objectives are intimately related to spectral clustering techniques, with the latter providing approximation guarantees for ratio cut objectives (such as conductance) via so-called *Cheeger* inequalities [11]. In some cases, these ratio cut objectives are optimized over an entire graph to solve a global clustering or classification task [18]. In other situations, the goal is to find sets of nodes that have a small ratio cut and are localized to a certain region of a large graph [3, 4, 27, 34].

Recently, there has been a surge of hypergraph methods for machine learning and data mining [1, 2, 5, 10, 23, 25, 37, 39], as hypergraphs can better model multiway relationships in data. Common examples of multiway relationships include academic researchers co-authoring papers, retail items co-purchased by shoppers, and sets of products or services reviewed by the same person. Due to broader modeling ability, there are many hypergraph generalizations of graph-cut objectives, including hypergraph variants of ratio cut objectives like conductance and normalized cut [6, 8, 9, 23, 39].

Nevertheless, there are numerous challenges in extending graph-cut techniques to the hypergraph setting, and current methods for hypergraph-based learning are much less developed than their graph-based counterparts. One major challenge in generalizing graph cut methods is that the concept of a cut hyperedge — how to define it and how to penalize it — is more nuanced than the concept of a cut edge. While there is only one way to separate the endpoints of an edge into two clusters, there are several ways to split up a set of three or more nodes in a hyperedge. Many objective functions model cuts with an *all-or-nothing* penalty function, which assigns the same penalty to any way of splitting up the nodes of the hyperedge (and a penalty of zero if all nodes in the hyperedge are placed together) [16, 17, 21] (Fig. 1a). However, a common practical heuristic is a *clique expansion*, which replaces each hyperedge with a weighted clique in a graph [6, 16, 23, 25, 39, 40] (Fig. 1b). The advantage is that graph methods can be directly applied, but this heuristic actually penalizes cut hyperedges differently than the all-or-nothing model. Another downside is that for hypergraph with large hyperedges, clique expansion produces a very dense graph.

Solving an all-or-nothing cut or applying clique expansion are only two specific models for higher-order relationships. And if one uses a ratio cut model for clusters in a hypergraph, how to penalize a cut hyperedge may depend on the application. Along these lines, *inhomogeneous hypergraphs* model every possible way to split up a hyperedge to assign different penalties [23]; however, these more sophisticated models are still approximated with weighted clique

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '20, August 23–27, 2020, Virtual Event, CA, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7998-4/20/08...\$15.00

<https://doi.org/10.1145/3394486.3403222>




Figure 1: Three different models for hyperedge penalties. The all-or-nothing cut has a scalar penalty w_e for any way of cutting a hyperedge e ; the penalty is zero if e is uncut. The clique expansion converts a hyperedge to a clique, leading to a cut penalty proportional to the product of the set sizes that e is split into. Recent generalized hyperedge cut functions assign penalties for each subset of nodes in e [23, 33].

expansions. In recent work, we developed a framework for *exactly* solving hypergraph s - t cut problems under similar notions of generalized hyperedge *splitting functions* [33] (Fig. 1c). However, these techniques have not been applied to ratio cut objectives or to solve practical machine learning and data mining problems.

There is also little work on *localized* hypergraph clustering, i.e., methods for finding well-connected sets of nodes that are biased towards a given region of a hypergraph. Existing approaches [22, 38] use random-walk-based local graph clustering methods [3] on clique expansions. *Flow-based* methods are an alternative to random walks [4, 14, 20, 27, 34], which have strong runtime and ratio cut quality guarantees. These methods solve $\max. s$ - t flow / $\min. s$ - t cut problems as a subroutine. Despite the success of these methods for graphs, they have not been extended to the hypergraph setting. **The present work: flow-based local hypergraph clustering.** Here, we develop a flow-based framework for local hypergraph clustering based on minimizing localized ratio cut objectives. Our framework takes in a set of input nodes and solves a sequence of hypergraph s - t cut problems to return a well-connected cluster of nodes that has a high overlap with the input set, where “well-connected” is formalized with a ratio-cut-style objective that can incorporate a wide range of hypergraph cut functions, including the all-or-nothing penalty or the clique-expansion-penalty.

Unlike clique expansion techniques, we do not simply reduce to an existing graph technique. Instead, our flow-based framework uses minimum *hypergraph* s - t cut computations. Thus, we can leverage recent results on minimum s - t hypergraph cut algorithms [33] to *exactly* solve localized hypergraph ratio cut objectives for generalized notions of hypergraph cuts. Importantly, although the s - t hypergraph cut solver only relies on solving graph s - t cut problems, the hypergraph ratio-cut is exactly optimized. Our implementation can also make use of high-performance maximum s - t flow solvers.

Our method comes with new guarantees on hypergraph conductance and normalized cut. The conductance results generalize previous guarantees in the graph setting to hypergraphs. The normalized cut guarantees are the first of their kind; since graphs are a special case of hypergraphs, we also have new results for the graph setting as a bonus. Our theory provides tighter guarantees than those obtained by applying existing results for approximate recovery of low-conductance sets, even in the graph setting.

A major feature of our methods is that they run in *strongly-local time*, meaning that the runtime is dependent only on the size of the input set, rather than the entire hypergraph. Therefore, we can find optimal clusters without even seeing the entire hypergraph, making our algorithms remarkably scalable in theory and practice.

We demonstrate our framework on large real-world hypergraphs, detecting product categories in Amazon review data and similar questions on Stack Overflow data. Given a small set of seed nodes, our method finds clusters with thousands of nodes from hypergraphs with millions of nodes and hyperedges and large average hyperedge size, often within a few seconds. Our methods are also more accurate than heuristics based on refining neighborhoods of a seed set or using existing graph methods on clique expansions.

2 PRELIMINARIES AND RELATED WORK

2.1 Background: Local Conductance in Graphs

Let $G = (V, E)$ be an undirected graph and $w_{ij} \geq 0$ be the weight for edge $(i, j) \in E$. The degree of a node v is $d_v = \sum_{u \in N_v} w_{uv}$, where N_v is the set of nodes sharing an edge with v . A common graph clustering objective is conductance, defined for $S \subseteq V$ by

$$\mathbf{cond}_G(S) = \frac{\mathbf{cut}(S)}{\min\{\mathbf{vol}(S), \mathbf{vol}(\bar{S})\}}, \quad (1)$$

where $\mathbf{vol}(S) = \sum_{v \in S} d_v$ is the volume of nodes in S , and $\mathbf{cut}(S) = \sum_{i \in S, j \in \bar{S}} w_{ij}$, which equals $\mathbf{cut}(\bar{S})$ by definition. Here, $\bar{S} = V \setminus S$ is the complement set of S . A related objective, which differs by at most a factor of two from conductance, is normalized cut [18]:

$$\mathbf{ncut}_G(S) = \frac{\mathbf{cut}(S)}{\mathbf{vol}(S)} + \frac{\mathbf{cut}(\bar{S})}{\mathbf{vol}(\bar{S})} = \mathbf{vol}(V) \frac{\mathbf{cut}(S)}{\mathbf{vol}(S)\mathbf{vol}(\bar{S})}. \quad (2)$$

Conductance and normalized cut are both NP-hard to minimize [36]. However, localized variants can be minimized in polynomial time with repeated $\max. s$ - t flow computations [4, 20, 27, 34, 35]. For example, given nodes $R \subseteq V$ with $\mathbf{vol}(R) \leq \mathbf{vol}(\bar{R})$, the following objective can be minimized in polynomial time [20]:

$$\underset{S \subseteq R}{\text{minimize}} \mathbf{cond}_G(S). \quad (3)$$

In other words, given a region defined by a *reference* set R , one can find the minimum conductance subset of R in polynomial time, even though minimizing conductance over an entire graph is NP-hard. A more general *local conductance* objective is

$$\mathbf{local-cond}_{R, \epsilon}(S) = \frac{\mathbf{cut}(S)}{\mathbf{vol}(S \cap R) - \epsilon \mathbf{vol}(S \cap \bar{R})}. \quad (4)$$

Objective (4) is minimized over all sets of nodes for which the denominator is positive (to avoid trivial outputs). The denominator rewards sets S that overlap with R . The *locality* parameter ϵ controls the penalty for including nodes outside R . As $\epsilon \rightarrow \infty$, minimizing Eq. (4) over sets producing a positive denominator becomes equivalent to Eq. (3). Andersen and Lang [4] showed how to minimize this objective for $\epsilon = \mathbf{vol}(R)/\mathbf{vol}(\bar{R})$, and faster algorithms were later developed for $\epsilon \gg \mathbf{vol}(R)/\mathbf{vol}(\bar{R})$ [27, 34]. These algorithms repeatedly solve maximum flow problems on an auxiliary graph.

Minimizing objective (4) can also provide cluster quality guarantees in terms of standard conductance (Eq. (1)). For example, the optimal set for (4) with $\epsilon = \mathbf{vol}(R)/\mathbf{vol}(\bar{R})$ has conductance within a small factor of the conductance of *any* set with a certain amount of overlap with R [4], and there are related results for other values of ϵ and objectives [27, 34, 35]. These are often called *cut improvement*

guarantees, since the ratio cut score of the output set *improves* upon the score of the input R . See the work of Fountoulakis et al. [14] for a detailed survey on flow-based methods for solving objective (4).

2.2 Background: Generalized Hypergraph Cuts

We now consider a hypergraph $\mathcal{H} = (V, E)$, where each edge $e \in E$ is a subset of the nodes in V (an undirected graph is then the special case where $|e| = 2$ for all $e \in E$). A hyperedge $e \in E$ is *cut* by a set $S \subseteq V$ if $e \cap S \neq \emptyset$ and $e \cap \bar{S} \neq \emptyset$, i.e., the hyperedge spans more than one cluster. We denote the set of edges cut by S by ∂S . The most common generalization of graph cut penalties to hypergraphs is to assign no penalty if $e \in E$ is not cut, but assign a fixed-weight scalar penalty of w_e for any way of cutting e [16, 17, 21]. Inhomogeneous hypergraphs generalize this by associating a weight *function* with each edge, rather than a scalar [23]; in this model, every distinct way of separating the nodes of a hyperedge can have its own penalty. We recently developed a related notion of *splitting functions* for hyperedges, in the context of hypergraph s - t cut problems [33].

We cover the splitting function terminology for general cut penalties here. Each edge $e \in E$ is associated with a *splitting function* $\mathbf{w}_e : A \subseteq e \rightarrow \mathbb{R}_{\geq 0}$ that maps each subset $A \subseteq e$ to a nonnegative splitting penalty. Weights on edges can be directly incorporated into the splitting function \mathbf{w}_e . By definition, splitting functions are required to be symmetric and penalize only cut hyperedges:

$$\mathbf{w}_e(A) = \mathbf{w}_e(e \setminus A) \text{ and } \mathbf{w}_e(e) = \mathbf{w}_e(\emptyset) = 0. \quad (5)$$

A splitting function is *submodular* if for all $A \subseteq e$ and $B \subseteq e$,

$$\mathbf{w}_e(A) + \mathbf{w}_e(B) \geq \mathbf{w}_e(A \cup B) + \mathbf{w}_e(A \cap B). \quad (6)$$

A splitting function is *cardinality-based* if it depends only on the number of nodes on each side of a split:

$$\mathbf{w}_e(A) = \mathbf{w}_e(B) \quad \text{whenever } |A| = |B|. \quad (7)$$

Given a splitting function for each hyperedge, the generalized hypergraph cut penalty for a set $S \subseteq V$ is given by

$$\mathbf{cut}_{\mathcal{H}}(S) = \sum_{e \in E} \mathbf{w}_e(e \cap S). \quad (8)$$

(By the symmetry constraint in Eq. (5), $\mathbf{cut}_{\mathcal{H}}(S) = \mathbf{cut}_{\mathcal{H}}(\bar{S})$.) The generalized hypergraph s - t cut objective is then:

$$\underset{S \subseteq V}{\text{minimize}} \mathbf{cut}_{\mathcal{H}}(S), \text{ subject to } s \in S, t \in \bar{S}, \quad (9)$$

where s and t are designated source and sink nodes. With the all-or-nothing splitting function, Eq. (9) is solvable in polynomial time via reduction to a directed graph s - t cut problem [21]. More generally, if all splitting functions are submodular, the hypergraph s - t cut problem is equivalent to minimizing a sum of submodular functions, which can be solved using general submodular function minimization [15, 28, 31] or specialty solvers for sums of submodular functions [12, 19, 24, 32]. Recently, we showed that when every splitting function is cardinality-based (Eq. (7)), the hypergraph s - t cut can be solved via reduction to a graph s - t cut problem if and only if all splitting functions are submodular [33]. Cardinality-based submodular splitting functions are the focus of our models.

2.3 Hypergraph Ratio Cut Objectives

Given this generalized framework for hypergraph cuts, we present definitions for hypergraph conductance and normalized cut, which subsume existing hypergraph ratio cut objectives. Let $\mathcal{H} = (V, E)$ be a hypergraph. For $S \subseteq V$, the hypergraph volume of S is

$$\mathbf{vol}_{\mathcal{H}}(S) = \sum_{v \in S} d_v, \quad (10)$$

where $d_v = \sum_{e: v \in e} \mathbf{w}_e(\{v\})$ is the hypergraph degree of v [23]. We define hypergraph conductance and normalized cut of $S \subseteq V$ as

$$\mathbf{cond}_{\mathcal{H}}(S) = \frac{\mathbf{cut}_{\mathcal{H}}(S)}{\min\{\mathbf{vol}_{\mathcal{H}}(S), \mathbf{vol}_{\mathcal{H}}(\bar{S})\}} \quad (11)$$

$$\mathbf{ncut}_{\mathcal{H}}(S) = \frac{\mathbf{cut}_{\mathcal{H}}(S)}{\mathbf{vol}_{\mathcal{H}}(S)} + \frac{\mathbf{cut}_{\mathcal{H}}(\bar{S})}{\mathbf{vol}_{\mathcal{H}}(\bar{S})}. \quad (12)$$

When \mathcal{H} is a graph, these reduce to the definitions of conductance (1) and normalized cut (2) in graphs. The hypergraph conductance in Eq. (11) has been used with the all-or-nothing splitting function [6, 8, 9]. Eq. (12) generalizes the version of hypergraph normalized cut from Zhou et al. [39], which is the special case of

$$\mathbf{w}_e(A) = \frac{w_e}{|e|} \cdot |A| \cdot |e \setminus A| \text{ for all } A \subseteq e,$$

where w_e is a scalar associated with $e \in E$. Eq. (12) also corresponds to the hypergraph normalized cut of Li and Milenkovic [23].

3 HYPERGRAPH LOCAL CONDUCTANCE

We now define our new localized hypergraph ratio cut objectives. Let $\mathcal{H} = (V, E)$ be a hypergraph and R a set of input nodes. We define a function $\Omega_{R, \varepsilon}$, which measures the overlap between R and another set of nodes S , parameterized by some $\varepsilon \geq \mathbf{vol}_{\mathcal{H}}(R) / \mathbf{vol}_{\mathcal{H}}(\bar{R})$.

$$\Omega_{R, \varepsilon}(S) = \mathbf{vol}_{\mathcal{H}}(S \cap R) - \varepsilon \mathbf{vol}_{\mathcal{H}}(S \cap \bar{R}). \quad (13)$$

This is our hypergraph analog to the denominator in Eq. (4). To find a good cluster of nodes in \mathcal{H} “near” R , we minimize

$$\mathbf{HLC}_{R, \varepsilon}(S) = \begin{cases} \frac{\mathbf{cut}_{\mathcal{H}}(S)}{\Omega_{R, \varepsilon}(S)} & \text{if } \Omega_{R, \varepsilon}(S) > 0 \\ \infty & \text{otherwise,} \end{cases} \quad (14)$$

which we call hypergraph (R, ε) -localized conductance. When R and ε are clear from context, we refer to (14) as HLC (hypergraph localized conductance), denoting its value by $\mathbf{HLC}(S)$. This objective reduces to the graph case when the hypergraph is a graph.

In this section, we show how to minimize **HLC**, given access to a minimum hypergraph s - t cut solver and consider cases where such solvers can be easily implemented with standard graph s - t cut solvers. Section 4 shows how to optimize the procedure outlined in this section (more formally, to have strongly-local runtime guarantees), and Section 5 adapts these results to provide bounds on the hypergraph conductance and normalized cut objectives.

3.1 Minimizing the HLC Objective

We now provide a procedure that minimizes **HLC**, given polynomially many queries to a solver for $\mathbf{cut}_{\mathcal{H}}$. Section 3.2 then considers cases where the solver itself requires polynomial time. Let $\mathcal{H} = (V, E)$ be the original input hypergraph and R the input set. We minimize **HLC** by repeatedly solving hypergraph min. s - t cut problems on an *extended* hypergraph \mathcal{H}_{α} , parameterized by $\alpha \in (0, 1)$:

- Keep all of $\mathcal{H} = (V, E)$ with original splitting functions.
- Introduce a source node s and sink node t .

Algorithm 1 Hypergraph local conductance minimization.

Input: \mathcal{H} , R , $\varepsilon \geq \text{vol}_{\mathcal{H}}(R)/\text{vol}_{\mathcal{H}}(\bar{R})$, $\text{cut}_{\mathcal{H}}$.

Set $\alpha = \text{HLC}(R)$ and $S = R$

do

Update $S_{\text{best}} \leftarrow S$ and save $\alpha_0 \leftarrow \alpha$

$S \leftarrow \arg \min_{S'} \mathbf{H}\text{-st-Cut}_{\alpha}(S')$

$\alpha \leftarrow \text{HLC}(S)$

while $\alpha < \alpha_0$

Return: S_{best}

- For each $r \in R$, add an edge (s, r) with weight αd_r .
- For each $j \in \bar{R}$, add an edge (j, t) with weight $\alpha \varepsilon d_j$.

By construction, \mathcal{H}_{α} contains hyperedges from \mathcal{H} and pairwise edges attached to source and sink nodes. Each hyperedge e in \mathcal{H} is associated with a splitting function \mathbf{w}_e . We call edges adjacent to s and t *terminal* edges, and these have standard cut penalty: 0 if the edge is not cut, otherwise the penalty is the weight of the edge. For any $S \subseteq V$, the value of the hypergraph cut $S \cup \{s\}$ in \mathcal{H}_{α} is

$$\mathbf{H}\text{-st-cut}_{\alpha}(S) = \text{cut}_{\mathcal{H}}(S) + \alpha \text{vol}_{\mathcal{H}}(\bar{S} \cap R) + \alpha \varepsilon \text{vol}_{\mathcal{H}}(S \cap \bar{R}). \quad (15)$$

Choosing $S = \emptyset$ gives an upper bound of $\alpha \text{vol}_{\mathcal{H}}(R)$ on the minimum cut score. Thus, if the minimizer S^* for Eq. (15) has a cut score strictly less than $\alpha \text{vol}_{\mathcal{H}}(R)$, then S^* must be nonempty, and we can rearrange (15) to show that

$$\text{HLC}(S^*) = \frac{\text{cut}_{\mathcal{H}}(S^*)}{\text{vol}_{\mathcal{H}}(S^* \cap R) - \varepsilon \text{vol}_{\mathcal{H}}(S^* \cap \bar{R})} < \alpha. \quad (16)$$

Thus, for any $\alpha \in (0, 1)$, to find out if there is a nonempty $S \subseteq V$ with HLC value less than α , it suffices to solve a generalized hypergraph s - t cut problem. Algorithm 1 gives a procedure for minimizing **HLC**, based on repeatedly solving objective (15) for smaller values of α until no more improvement in the HLC objective is possible.

3.2 A New Hyperedge Splitting Function

Algorithm 1 repeatedly solves the **H-st-cut** objective (15) in an *auxiliary* hypergraph \mathcal{H}_{α} . For submodular splitting functions, this can be done in polynomial time with methods for minimizing sums of submodular functions [12, 19, 24, 32]. For the more restrictive class of cardinality-based submodular splitting functions, we only need to solve directed graph s - t cut problems [33]. Implementations of such solvers are readily available and perform well in practice.

As an example, we present a new class of cardinality-based splitting functions that depends on a single tunable integer parameter $\delta \geq 1$, which we use for our numerical experiments for its modeling capability and computational appeal:

$$\mathbf{w}_e(A) = \min\{\delta, |A|, |e \setminus A|\} \quad \text{for any } A \subseteq e. \quad (17)$$

We call this the δ -*linear threshold* splitting function, since the penalty is linear in terms of the small side of the cut, up until a maximum penalty of δ . In other words, in a split hyperedge, we incur a unit cost for adding another node to the small side of the cut, up until we reach δ such nodes. The $\delta = 1$ case is equivalent to the unweighted all-or-nothing cut. For large enough δ , the δ -linear threshold is the linear hyperedge splitting penalty, which is equivalent to applying a star expansion to the hypergraph [40]. Choosing different values for δ interpolates between these common splitting

functions, which enables the detection of different types of cut sets in a hypergraph and provides the data modeler with flexibility.

We now show how to efficiently optimize s - t hypergraph cuts for this splitting function using graph s - t cuts. Let $\mathcal{H} = (V, E)$ be a hypergraph, where each edge is associated with the δ -linear threshold splitting function. A minimum s - t cut problem in \mathcal{H} can be reduced to a minimum s - t cut problem in a new directed graph $G_{\mathcal{H}}$ by replacing each $e \in E$ with the following *gadget*:

- Introduce two auxiliary nodes v'_e and v''_e .
- Create a directed edge from v'_e to v''_e with weight δ .
- For each $v \in e$, add directed edges (v, v'_e) and (v''_e, v) , both with weight 1.

In any min. s - t cut solution in $G_{\mathcal{H}}$, the auxiliary nodes of e are arranged in a way that leads to a minimum possible cut. If e is cut and $A \subseteq e$ is on the source side of the cut, then the penalty in $G_{\mathcal{H}}$ removes all directed paths from A to $e \setminus A$ by taking the smaller penalty among three options: cutting (i) the middle edge δ , (ii) all edges from A to v'_e , or (iii) all edges from v''_e to $e \setminus A$. Thus, the penalty in $G_{\mathcal{H}}$ at this gadget will be exactly the hypergraph splitting penalty (17). This gadget-based approach is related to our recent techniques for modeling general cardinality-based submodular splitting functions [33]. However, the approach here requires fewer auxiliary nodes and directed edges. In Section 6, we show that minimizing HLC with the δ -linear threshold penalty for varying δ enables us to efficiently detect better clusters in a large hypergraph.

Finally, we provide an upper bound on the number of minimum s - t cut problems Algorithm 1 must solve if the δ -linear threshold penalty is applied. A proof is given in the appendix.

THEOREM 3.1. *Let S_i be the set returned after the i th iteration of Alg. 1. The value of $\text{cut}_{\mathcal{H}}(S_i)$ strictly decreases until the last iteration. Thus, with the δ -linear threshold penalty, the cut value decreases by at least one in each iteration, for a maximum of $\text{cut}_{\mathcal{H}}(R)$ iterations.*

This bound is loose in practice — the algorithm converges in 2–5 iterations in nearly all of our experiments. Similar results can be developed for non-integer δ , though we omit the details.

4 A STRONGLY-LOCAL ALGORITHM

Recall that ε is a *locality* parameter that controls the sets $S \subseteq V$ for which $\text{HLC}(S) < \infty$. If ε is large, then $\Omega_{R, \varepsilon}(S) < 0$ for many sets that do not substantially overlap with R . In local clustering applications where $\text{vol}(R)$ is small compared to the entire hypergraph, there is little gained by considering sets S that share little overlap with R . In these settings, it is preferable to explore only a small region *nearby* R , rather than exploring every node and hyperedge in the hypergraph. Thus, it is natural to choose a larger value of ε and output low hypergraph conductance sets overlapping with R . Ideally, we want to avoid even looking at the entire hypergraph.

We can formalize this idea via strong locality. A local clustering algorithm is *strongly-local* if its runtime depends only on the size of the input set and not the entire hypergraph. In contrast, Alg. 1 is *weakly-local*, meaning that its output is biased towards a region of the hypergraph, but its runtime may still depend on the size of the entire hypergraph. In this section, we give a strongly-local variant of Alg. 1, when $\text{vol}(R) \ll \text{vol}(\bar{R})$ and $\varepsilon \gg \text{vol}(R)/\text{vol}(\bar{R})$ is treated

as a small constant. Our procedure generalizes strongly-local methods for minimizing local conductance in graphs [27, 34]. For the hypergraph setting, we combine previous techniques for local min. s - t cut computations in graphs with techniques for converting hypergraph s - t cut problems into graph s - t cut problems [21, 33].

4.1 Making the Procedure Strongly-Local

In order to solve the hypergraph s - t cut objective (15) in strongly-local time, we must avoid explicitly constructing \mathcal{H}_α . We instead begin with a sub-hypergraph \mathcal{L} of \mathcal{H}_α , which we call the *local hypergraph*, and alternate between the following two steps:

- (1) Solve a hypergraph minimum s - t cut problem on \mathcal{L} .
- (2) Grow the subgraph \mathcal{L} based on the s - t cut solution.

The algorithm proceeds until a convergence criterion is satisfied, at which point the growth mechanism in Step 2 will stop and the algorithm will output the minimum s - t cut solution for \mathcal{H}_α .

Algorithm Terminology. Let $\mathcal{H} = (V, E)$ be the hypergraph. For any $v \in V$, let $E(v) = \{e \in E : v \in e\}$, and define $E(S) = \cup_{v \in S} E(v)$ for any set $S \subseteq V$. Let $\mathcal{H}_\alpha = (V \cup \{s, t\}, E \cup E^{st})$ be defined as in Section 3, where E^{st} is the terminal edge set. For $v \in V$, let e_v^{st} denote its terminal edge in \mathcal{H}_α (recall that nodes in R are connected to s and nodes in \bar{R} are connected to t), and let $E^{st}(S)$ denote the set of edges between nodes in $S \cup \{s, t\}$ for any set S . Our goal is to find a min. s - t cut of \mathcal{H}_α without forming \mathcal{H}_α . To do so, we assume that we have an oracle that efficiently outputs $E(v)$ for any $v \in V$. From $E(v)$, we can extract the *neighborhood* of v in \mathcal{H} :

$$\mathcal{N}(v) = \{u \in V \mid \exists e \in E \text{ such that } u, v \in e\} \quad (18)$$

For a set S , we define $\mathcal{N}(S) = \cup_{v \in S} \mathcal{N}(v)$.

The Local Hypergraph. Let $\mathcal{L} = (V_L \cup \{s, t\}, E_L \cup E_L^{st})$ denote the local hypergraph, a growing subgraph of \mathcal{H}_α over which we will repeatedly solve min. s - t cut problems. We initialize \mathcal{L} to contain all nodes and neighbors of R , i.e., $V_L = R \cup \mathcal{N}(R)$, and add the terminal edge for each of these nodes to E_L^{st} . The set E_L is initialized to $E(R)$, the set of edges containing at least one node from R . As the algorithm progresses, \mathcal{L} grows to include more nodes and edges from \mathcal{H}_α , always maintaining that $V_L \subseteq V$, $E_L \subseteq E$, and $E_L^{st} \subseteq E^{st}$. For $S \subseteq V_L$, let $\mathbf{L-st-cut}_\alpha(S)$ be the value of the s - t cut $S \cup \{s\}$ in \mathcal{L} . Because \mathcal{L} is a sub-hypergraph of \mathcal{H}_α , at every step we have that

$$\mathbf{L-st-cut}_\alpha(S) \leq \mathbf{H-st-cut}_\alpha(S) \text{ for all } S \subseteq V_L. \quad (19)$$

By carefully selecting new nodes and edges to add to \mathcal{L} , we will show that the min. s - t cut solution in \mathcal{L} converges to the min. s - t cut solution in \mathcal{H}_α , without ever fully forming \mathcal{H}_α .

Two Alternating Steps. After initializing \mathcal{L} , we repeat two steps until convergence: (1) find a min. s - t cut in \mathcal{L} , and (2) grow \mathcal{L} based on the solution to the cut problem. To grow \mathcal{L} at each iteration, we track which nodes from \bar{R} have had their edge to the sink cut by some min. s - t cut in \mathcal{L} in a previous iteration. When (v, t) is cut for the first time, we expand the local hypergraph by adding the neighbors and edges adjacent to v in \mathcal{H}_α :

- Update $V_L \leftarrow V_L \cup \mathcal{N}(v)$.
- Update $E_L \leftarrow E_L \cup E(v)$ and $E_L^{st} \leftarrow E_L^{st} \cup E^{st}(\mathcal{N}(v))$.

At this point, we say that node v has been *explored*, and we maintain a set of nodes X that have been explored at any point during the algorithm. Since R is already contained in \mathcal{L} , any new nodes we

Algorithm 2 Strongly-Local Min s - t cut solver

Input: R, ε, α , and access to $E(v)$ for any v in hypergraph \mathcal{H} .
Output: Min s - t cut solution S for \mathcal{H}_α . objective (15)
// Initialize Local graph \mathcal{L}
 $V_L \leftarrow R \cup \mathcal{N}(R), E_L^{st} \leftarrow E^{st}(V_L), E_L \leftarrow E(R), X \leftarrow \emptyset, N \leftarrow \emptyset$
do
 // **Step 1:** Solve a minimum s - t cut problem on \mathcal{L} .
 $S_L = \operatorname{argmin}_{S \subseteq V_L} \mathbf{L-st-cut}_\alpha(S)$
 $N = S_L \cap \bar{R} \cap V \setminus X$ (nodes around which to expand \mathcal{L})
 // **Step 2:** Grow \mathcal{L} .
 $V_L \leftarrow V_L \cup \mathcal{N}(N), E_L \leftarrow E_L \cup E(N), E_L^{st} \leftarrow E_L^{st} \cup E^{st}(\mathcal{N}(N))$
 $X \leftarrow X \cup N$.
while $N \neq \emptyset$
Return S_L

add to the local hypergraph will be from \bar{R} . Algorithm 2 shows pseudo-code for the overall procedure.

Convergence and Locality. The algorithm terminates when, after a min. s - t cut computation, no new edges are added E_L . At this point, the min. s - t cut set in \mathcal{L} is the min. s - t cut set of G_α .

THEOREM 4.1. *The set S returned by Algorithm 2 minimizes objective (15), the minimum s - t cut objective on \mathcal{H}_α .*

Furthermore, the following theorem shows that under reasonable conditions, the growth of the local hypergraph is bounded in terms of the $\mathbf{vol}_{\mathcal{H}}(R)$. Thus, our algorithm is strongly-local. We use the term *graph-reducible* to refer to any hypergraph cut function for which the hypergraph s - t cut problem can be reduced to an equivalent s - t cut problem in a directed graph.

THEOREM 4.2. *Suppose we have a cardinality-based submodular splitting function scaled with minimum non-zero penalty 1 (e.g., δ -linear threshold) and that no nodes in R are isolated. Then the local hypergraph \mathcal{L} will have at most $\frac{3}{2}(1 + 1/\varepsilon)\mathbf{vol}_{\mathcal{H}}(R)$ hyperedges, and the number of nodes and terminal edges will both be at most $k\mathbf{vol}_{\mathcal{H}}(R)(1 + 1/\varepsilon)$, where k is the maximum size hyperedge in \mathcal{H}_α .*

Proofs of Theorems 4.1 and 4.2 are in the appendix. The minimum value on the splitting function is just a scaling issue, and the cardinality-based submodular restriction lets us bound set sizes by volumes of those sets. The assumption that R has no isolated nodes is minor; these nodes could be removed in a pre-processing step.

4.2 Runtime Analysis

Theorem 4.2 gives strongly local runtimes for Alg. 2 when we use the δ -linear threshold splitting function with $\delta \geq 1$, since it is a cardinality-based, submodular function where the minimum non-zero penalty is 1. This immediately implies the same runtime result for the all-or-nothing penalty ($\delta = 1$). We implement Algorithm 2 for the δ -linear threshold penalty by replacing each edge e with a small directed graph as outlined in Section 3.2. If k is the maximum hyperedge size, this graph reduction introduces 2 auxiliary nodes and at most $(2k + 1)$ directed edges for each hyperedge that appear in the local hypergraph. Combining this with Theorem 4.2, the graph reduction of the largest local hypergraph \mathcal{L} has at most

$$k\mathbf{vol}_{\mathcal{H}}(R)(1 + 1/\varepsilon) + 3(1 + 1/\varepsilon)\mathbf{vol}_{\mathcal{H}}(R) = O(k\mathbf{vol}_{\mathcal{H}}(R)(1 + 1/\varepsilon))$$

nodes and

$$k\mathbf{vol}_{\mathcal{H}(R)(1+1/\varepsilon)} + (2k+1)\frac{3}{2}(1+1/\varepsilon)\mathbf{vol}_{\mathcal{H}(R)} = O(k\mathbf{vol}_{\mathcal{H}(R)(1+1/\varepsilon)})$$

edges. For a graph $G = (V, E)$, there is an s - t cut algorithm with runtime $O(|V||E|)$ [29]. In theory, we can use this to solve the hypergraph s - t cut problem with δ -linear threshold penalties on the largest local hypergraph \mathcal{L} in time $O(k^2\mathbf{vol}_{\mathcal{H}(R)}^2(1+1/\varepsilon)^2)$. The local hypergraph grows by at least one hyperedge (i.e., $(2k+1)$ directed edges) each step, so we need to solve $O(k\mathbf{vol}_{\mathcal{H}(R)(1+1/\varepsilon)})$ s - t cut problems, for an overall runtime of $O(k^3\mathbf{vol}_{\mathcal{H}(R)}^3(1+1/\varepsilon)^3)$.

Using high-performance max-flow/min-cut solvers, the runtime of our algorithm is fast in practice – on hypergraphs with millions of nodes and edges, roughly a few seconds for small R and a few minutes for large R . Nevertheless, as long as $1/\varepsilon$ is independent of the size of the input hypergraph (e.g., $\varepsilon = 1$ is always a valid choice), then our overall procedure for minimizing localized ratio cuts in hypergraphs is strongly-local. Furthermore, if the maximum hyperedge size k is a constant, our asymptotic runtime is the same as the runtime for strongly-local graph clustering algorithms [34, 35], which are also effective in practice. Finally, although our analysis focused on the δ -linear threshold penalty used for our experiments, we can get a similar runtime for any cardinality-based submodular splitting function, with a slightly worse dependence on k .

5 RATIO CUT IMPROVEMENT GUARANTEES

The HLC objective is meaningful in its own right, and we use it in our experiments; however, understanding the relationship between HLC and more standard ratio cut objectives that do not inherently depend on R and ε is also useful. To this end, we derive guarantees satisfied by Algorithm 1 and the HLC objective in terms of hypergraph conductance and normalized cut. Our theory shows that the algorithm output has a ratio cut score that is nearly as good as any other set of nodes that have some overlap with R .

Our results in this section are for a fixed input hypergraph \mathcal{H} , so we drop the subscript \mathcal{H} to simplify notation. Throughout this section, let $\varepsilon_0 = \mathbf{vol}(R)/\mathbf{vol}(\bar{R}) \leq 1$ denote the minimum value of the locality parameter. Setting $\varepsilon = \varepsilon_0$ gives the best cut improvement guarantees, which is always a valid choice. However, we prove results for more general parameter settings, since, as discussed in Section 4, there are locality and runtime benefits for setting $\varepsilon > \varepsilon_0$.

5.1 Hypergraph Conductance Guarantees

We first generalize previous conductance improvement guarantees for local graph clustering [4, 34] to our hypergraph objective.

THEOREM 5.1. *Let S^* be the set returned by Algorithm 1 for some $\varepsilon \in (\varepsilon_0, \varepsilon_0 + 1)$, and let $\mu = \varepsilon - \varepsilon_0 \geq 0$.*

- (1) *For any $T \subseteq R$, $\mathbf{cond}(S^*) \leq \mathbf{cond}(T)$.*
- (2) *For any set T satisfying $\mathbf{vol}(T) \leq \mathbf{vol}(\bar{T})$ and*

$$\frac{\mathbf{vol}(T \cap R)}{\mathbf{vol}(T)} \geq \frac{\mathbf{vol}(R)}{\mathbf{vol}(V)} + \gamma \frac{\mathbf{vol}(\bar{R})}{\mathbf{vol}(V)}, \quad (20)$$

for some $\gamma \in (\mu, 1)$, we have that $\mathbf{cond}(S^) \leq \frac{1}{\gamma - \mu} \mathbf{cond}(T)$. Specifically, when $\varepsilon = \varepsilon_0$, $\mathbf{cond}(S^*) \leq (1/\gamma) \mathbf{cond}(T)$.*

PROOF. For notational compactness, let $\mathbf{v}(S) = \mathbf{vol}(S)$ for a set $S \subseteq V$. We first prove that $\Omega_{R,\varepsilon}(S) \leq \min\{\mathbf{v}(S), \mathbf{v}(\bar{S})\}$ for any set S ,

which implies that $\mathbf{cond}(S) \leq \mathbf{HLC}(S)$:

$$\Omega_{R,\varepsilon}(S) = \mathbf{v}(S \cap R) - \varepsilon \mathbf{v}(S \cap \bar{R}) \leq \mathbf{v}(S \cap R) \leq \mathbf{v}(S),$$

$$\begin{aligned} \Omega_{R,\varepsilon}(S) &\leq \mathbf{v}(S \cap R) - \varepsilon_0 \mathbf{v}(S \cap \bar{R}) \\ &= \mathbf{v}(R) - \mathbf{v}(\bar{S} \cap R) - \varepsilon_0 (\mathbf{v}(\bar{R}) - \mathbf{v}(\bar{S} \cap \bar{R})) \\ &\leq \mathbf{v}(R) - \varepsilon_0 \mathbf{v}(\bar{R}) + \varepsilon_0 \mathbf{v}(\bar{R} \cap \bar{S}) = \varepsilon_0 \mathbf{v}(\bar{R} \cap \bar{S}) \leq \mathbf{v}(\bar{S}), \end{aligned}$$

where the last equality uses the definition of ε and the final inequality uses $\varepsilon_0 < 1$. Finally, for any $T \subseteq R$ where $\mathbf{v}(R) \leq \mathbf{v}(\bar{R})$, $\mathbf{HLC}(T) = \mathbf{cond}(T)$, which gives the first theorem statement: $\mathbf{cond}(S^*) \leq \mathbf{HLC}(S^*) \leq \mathbf{HLC}(T) = \mathbf{cond}(T)$.

For the second statement, if $\Omega_{R,\varepsilon}(T) \leq 0$, then $\mathbf{HLC}(T) = \infty$ and the result is trivial. Assume then that $\Omega_{R,\varepsilon}(T) > 0$. Because $\mathbf{HLC}(S^*) \leq \mathbf{HLC}(T)$, the result will hold if we can prove that $\mathbf{HLC}(T) \leq \frac{1}{\gamma - \mu} \mathbf{cond}(T)$, which is true as long as $\Omega_{R,\varepsilon}(T) \geq (\gamma - \mu) \mathbf{v}(T)$. We prove this by applying assumption (20).

$$\begin{aligned} \frac{\Omega_{R,\varepsilon}(T)}{\mathbf{v}(T)} &\geq \frac{(1 + \varepsilon) \mathbf{v}(T \cap R) - \varepsilon \mathbf{v}(T)}{\mathbf{v}(T)} \geq (1 + \varepsilon) \left(\frac{\mathbf{v}(R)}{\mathbf{v}(V)} + \gamma \frac{\mathbf{v}(\bar{R})}{\mathbf{v}(V)} \right) - \varepsilon \\ &= \left(1 + \frac{\mathbf{v}(R)}{\mathbf{v}(\bar{R})} \right) \gamma \frac{\mathbf{v}(\bar{R})}{\mathbf{v}(V)} + \mu \gamma \frac{\mathbf{v}(\bar{R})}{\mathbf{v}(V)} + (1 + \varepsilon) \frac{\mathbf{v}(R)}{\mathbf{v}(V)} - \varepsilon \\ &= \gamma + \mu \gamma \frac{\mathbf{v}(\bar{R})}{\mathbf{v}(V)} + (1 + \varepsilon) \frac{\mathbf{v}(R)}{\mathbf{v}(V)} - \varepsilon \left(\frac{\mathbf{v}(R)}{\mathbf{v}(V)} + \frac{\mathbf{v}(\bar{R})}{\mathbf{v}(V)} \right) \\ &= \gamma + (\mu \gamma - \mu) \frac{\mathbf{v}(\bar{R})}{\mathbf{v}(V)} - \frac{\mathbf{v}(R)}{\mathbf{v}(\bar{R})} \frac{\mathbf{v}(\bar{R})}{\mathbf{v}(V)} + \frac{\mathbf{v}(R)}{\mathbf{v}(V)} \geq \gamma - \mu. \end{aligned}$$

□

5.2 Hypergraph Normalized Cut Guarantees

Given that conductance and normalized cut differ by at most a factor of two, we can translate Theorem 5.1 into bounds for normalized cut. Our next theorem, however, obtains better guarantees by directly developing bounds for hypergraph normalized cut. This demonstrates how our algorithmic framework transcends its relationship with conductance, as it can be used to find sets that also satisfy strong guarantees for other common objectives.

THEOREM 5.2. *Let S^* be the set returned by Algorithm 1 for some $\varepsilon \in (\varepsilon_0, \varepsilon_0 + 1)$, and let $\mu = \varepsilon - \varepsilon_0 \geq 0$. If a set $T \subset V$ satisfies $\mathbf{vol}(T) \leq \mathbf{vol}(\bar{T})$ and for some $\beta \in (2\mu/(1+2\mu), 1)$ satisfies*

$$\frac{\mathbf{vol}(T \cap R)}{\mathbf{vol}(T)} \geq \frac{\mathbf{vol}(\bar{T} \cap R)}{\mathbf{vol}(\bar{T})} + \beta, \quad (21)$$

we have that $\mathbf{ncut}(S^) \leq \frac{1}{\beta + 2\mu\beta - 2\mu} \mathbf{ncut}(T)$. Specifically, when $\varepsilon = \varepsilon_0$, $\mathbf{ncut}(S^*) \leq (1/\beta) \mathbf{ncut}(T)$.*

A full proof is in the appendix. The overlap assumptions (20) and (21) in Theorems 5.1 and 5.2 differ. The assumption in Theorem 5.1 matches previous local graph clustering results [4, 34] and measures how much R overlaps with a set T . In contrast, assumption (21) provides a more intuitive measure of how much more R overlaps with T than it does with \bar{T} . This is the first application of this type of overlap assumption for cut improvement – graph or hypergraph. We next give a simple example for how this overlap assumption and Theorem 5.2 provide meaningful new normalized cut improvement guarantees, even in the well-studied graph setting.

Example. Consider a hypergraph (or graph) $\mathcal{H} = (V, E)$ containing a low-conductance target set T with $\mathbf{vol}(T) = \mathbf{vol}(\bar{T}) =$

$\text{vol}(V)/2$. Assume that we do not know all of T , but we have access to a set R constituting half the volume of T , i.e., $R \subset T$ with $\text{vol}(R) = \text{vol}(T)/2$. Let S^* be the set returned by Algorithm 1 when $\varepsilon = \text{vol}(R)/\text{vol}(\bar{R})$. First, we apply Theorem 5.1 to bound $\text{cond}(S^*)$, where assumption (20) holds with $\gamma = 1/3$:

$$\frac{\text{vol}(R \cap T)}{\text{vol}(T)} = \frac{1}{2} = \frac{1}{4} + \frac{1}{3} \cdot \frac{3}{4} = \frac{\text{vol}(R)}{\text{vol}(V)} + \gamma \frac{\text{vol}(\bar{R})}{\text{vol}(V)}, \quad (22)$$

and this value of γ gives the tightest bound. Thus, Theorem 5.1 guarantees that $\text{cond}(S^*) \leq 3 \text{cond}(T)$. Next, using the fact that $\text{cond}(S^*) \leq \text{ncut}(S^*) \leq 2 \text{cond}(S^*)$,

$$\text{ncut}(S^*) \leq 2 \text{cond}(S^*) \leq 2 \cdot 3 \text{cond}(T) \leq 6 \text{ncut}(T).$$

On the other hand, assumption (21) is satisfied with $\beta = 1/2$, so Theorem 5.2 guarantees that $\text{ncut}(S^*) \leq 2 \text{ncut}(T)$. This is significantly tighter than combining the bound from Theorem 5.1 and the relationship between normalized cut and conductance. This result demonstrates that although HLC is presented as a localized variant of conductance, there are also situations in which we can obtain *even better* set recovery guarantees in terms of normalized cut than conductance. To summarize, our approach returns meaningful results in terms of more than just one clustering objective.

6 EXPERIMENTS

We call running Alg. 1 with Alg. 2 as a subroutine *HyperLocal*, since it operates on hypergraphs with a strongly-local runtime. Next, we apply HyperLocal to identify clusters of question topics on Stack Overflow and product categories within Amazon reviews.

6.1 Algorithms and Implementation Details

We implement HyperLocal in Julia, using a push-relabel implementation of the maximum s - t flow method for the underlying s - t cut problems. All experiments ran on a laptop with 8 GB of RAM and a 2.2 GHz Intel Core i7 processor. We provide code and datasets at <https://github.com/nveldt/HypergraphFlowClustering>.

Neighborhood Baselines. If R_s is a set of seed nodes, let $\mathcal{N}(R_s)$ be its one-hop neighborhood. In the hypergraphs we consider, the one-hop neighborhood of a seed set is often quite large. We design two baselines for returning a cluster nearby a set of seeds. *TopNeighbors* orders nodes in $\mathcal{N}(R_s)$ based on the number of hyperedges that each $v \in \mathcal{N}(R_s)$ shares with at least one node from R_s and outputs the top k such nodes. Similarly, *BestNeighbors* orders each node $v \in \mathcal{N}(R_s)$ by the *fraction* of hyperedges incident to v that are also incident to at least one node from R_s and outputs the top k . In our experiments, we choose k to be equal to ground truth cluster sizes, which provides an additional advantage to these baselines.

Clique Expansion + FlowSeed Baselines. FlowSeed [35] is a flow-based method for solving localized conductance (4) in graphs. For one baseline, we convert an input hypergraph to a graph and then run FlowSeed with the same input set R and parameter ε as we use for HyperLocal. We consider two types of expansion: replacing a hyperedge e with an unweighted clique, and replacing a hyperedge with a clique where each edge has weight $1/|e|$. These are representative of existing clique expansion techniques for (local) hypergraph clustering [22, 38, 39]. While other local graph clustering methods exist, we focus on comparing against FlowSeed, as it has been shown to outperform other techniques in a number

Table 1: Average runtime in seconds, precision, recall, and F1 scores across 45 target clusters that correspond to question topics on Stack Overflow. Top F1 is the number of times out of 45 that a method/set obtained the top F1 score (including ties). Runtimes for last three rows are negligible. UCE and WCE indicate *unweighted* and *weighted* clique expansions.

Method	runtime	pr	re	f1	Top F1
HyperLocal	25.0	0.69	0.47	0.53	29
UCE + FlowSeed	32.8	0.3	0.58	0.4	1
WCE + FlowSeed	32.9	0.3	0.58	0.4	1
BestNeighbors	–	0.49	0.49	0.49	11
TopNeighbors	–	0.45	0.45	0.45	6
R	–	0.3	0.6	0.4	1

of settings [35]. Comparing against FlowSeed also allows us to best highlight the difference between running a flow-based method designed specifically for hypergraphs, versus performing clique reduction and applying a related graph algorithm. Finally, comparing against non-local hypergraph clustering methods (e.g., [23, 25]) is infeasible, as these do not seek clusters near a specified input set, and are unable to run on hypergraphs with millions of nodes and hyperedges in a reasonable amount of time.

6.2 Question Topics on Stack Overflow

HyperLocal is able to identify clusters of questions associated with the same topic on Stack Overflow. We represent each question as a node and construct hyperedges from the set of questions answered by a single user. Tags indicate sets of questions on the same topic (e.g., “julia”, “netsuite”, “common-lisp”), which we use as ground truth cluster labels (many questions have multiple tags). The hypergraph has 15,211,989 nodes and 1,103,243 edges, with a mean hyperedge size of 23.7. The dataset has 56,502 tags. We use the 45 tags that have between 2,000 and 10,000 questions and a hypergraph conductance score below 0.2 under the all-or-nothing penalty. Thus, we focus on sets of tags that can reasonably be viewed as modestly-sized clusters in the dataset.

Experimental Setup. To get an input set R for Alg. 1, we assume that we are given a small subset of seed nodes R_s from the target cluster and a rough idea of the cluster’s size. HyperLocal is designed to find good clusters by refining a moderately-sized reference set, so we use BestNeighbors to grow R_s into an initial reference set $R \supseteq R_s$ that is refined by HyperLocal. We ensure HyperLocal finds a cluster that strictly contains R_s by adding infinite-weight edges from R_s to the source node in the underlying minimum s - t cut problems, following prior approaches for local graph clustering [35].

For each target cluster T , we randomly select 5% of T as a seed set R_s , and use BestNeighbors to grow R_s by an additional $2|T|$ nodes. This produces a reference set input R for HyperLocal. We set $\varepsilon = 1.0$ and use the δ -linear threshold penalty with $\delta = 5000$. This large threshold tends to produce good results (see Section 6.4 for details), and the threshold is meaningful as some hyperedges contain tens of thousands of nodes.

We run TopNeighbors and BestNeighbors, outputting the top $|T|$ nodes in the ordering defined by each. Therefore, we give these methods an advantage by assuming that they know the exact size of the target cluster. We also run FlowSeed on unweighted

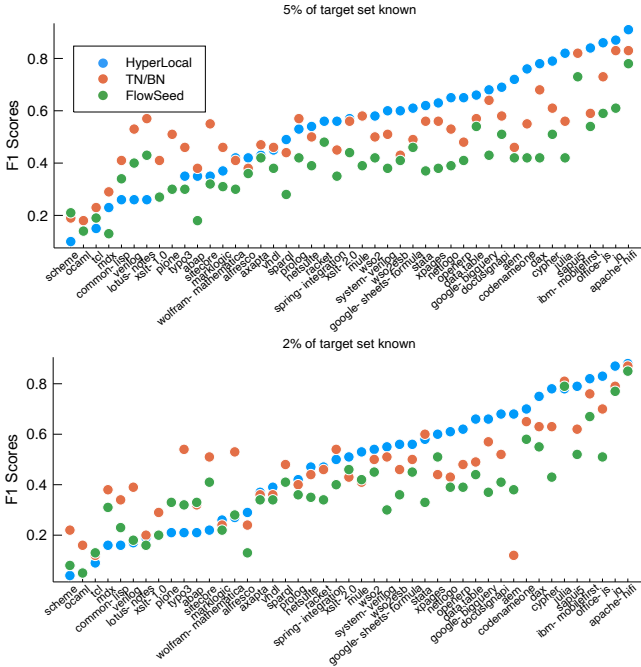


Figure 2: F1 scores for each of 45 clusters in the Stackoverflow hypergraph. For FlowSeed, we show the best result from the weighted or unweighted clique expansion; we also show the best of BN/TN. (Top) When 5% of the target set is used as a seed set, HyperLocal has the highest mean F1 score, and outperforms all methods in all cases where at least one method has an F1 score above 0.6. (Bottom) Results are very similar when only 2% of the target set is used as seeds.

and weighted clique expansions, using the same parameters as HyperLocal. In order to use clique expansion without densifying the graph too much and running into memory issues, we first discard hyperedges with 50 or more nodes (around 8% of all hyperedges). Finally, we ran the entire experimental procedure a second time, but starting with a random subset of 2% of each target cluster T as a seed set, rather than 5%.

Results. Table 1 reports the performance of each method, averaged across all 45 clusters, when 5% of each target set is known. HyperLocal has the highest average F1 score, and obtains the best F1 scores on many more clusters compared to other methods. Figure 2 visualizes F1 scores for individual clusters, both when 5% and 2% of the target cluster is known. Even when only 2% of each target cluster is known, we see nearly identical results.

6.3 Detecting Amazon Product Categories

Next, we use HyperLocal to quickly detect clusters of retail products with the same category (e.g., ‘‘Appliances’’, ‘‘Software’’) from a large hypergraph constructed from Amazon product review data [26]. We construct hyperedges as sets of products (nodes) reviewed by the same person. The hypergraph has 2,268,264 nodes and 4,285,363 hyperedges, with a mean hyperedge size just over 17. We use product category labels as ground truth cluster identities and consider the 9 smallest clusters, each of which represents only a very small fraction of nodes in the hypergraph (Table 2).

Table 2: The first column is the size of each Amazon product-category cluster T ; the second is HyperLocal (HL) runtime in seconds. The remaining columns are F1 scores for HL, Best-Neighbors (BN), TopNeighbors (TN), and the reference set R . Except in one case, HyperLocal significantly improves on the F1 score of R and always outperforms BN and TN.

Cluster	$ T $	run	HL	BN	TN	R
Amazon Fashion	31	3.5	0.83	0.77	0.6	0.67
All Beauty	85	30.8	0.69	0.6	0.28	0.58
Appliances	48	9.8	0.82	0.73	0.56	0.68
Gift Cards	148	6.5	0.86	0.75	0.71	0.63
Magazine Subscriptions	157	14.5	0.87	0.72	0.56	0.76
Luxury Beauty	1581	261	0.33	0.31	0.17	0.41
Software	802	341	0.74	0.52	0.24	0.42
Industrial & Scientific	5334	503	0.55	0.49	0.15	0.35
Prime Pantry	4970	406	0.96	0.73	0.36	0.23

Experimental Setup. We use $\varepsilon = 1.0$ and the standard all-or-nothing cut penalty, i.e., $\delta = 1$ for the δ -linear threshold penalty. Unlike the Stack Overflow dataset, this smaller δ tends to work well (see Section 6.4). For the six smallest clusters (under 200 nodes), we use $|R_S| = 10$ random seed nodes and use BestNeighbors to grow an initial cluster R with R_S plus 200 additional nodes for HyperLocal to refine. For the two clusters closer to 1000 nodes, we use 50 seed nodes, which we grow by another 2,000 nodes using BestNeighbors. The largest two clusters have around 5,000 nodes. For these, we extract a random subset of 200 nodes and use BestNeighbors to add 10,000 neighbors to form R .

Results. We compare HyperLocal to BestNeighbors and TopNeighbors in terms of F1 score. (We also attempted to run FlowSeed, but were unable to perform a clique expansion on the hypergraph due to memory constraints, as the expanded graph becomes too dense even after removing all hyperedges with 50 nodes or more.) Table 2 reports F1 detection scores, averaged across 5 different trials with different random seed sets. In all cases, HyperLocal substantially improves upon the baselines.

To test robustness, we ran numerous additional experiments on the smallest five clusters while varying the locality parameter $\varepsilon \in \{10^{-3}, 10^{-2}, 10^{-1}, 1, 10\}$ and reference set size $|R| \in \{200, 300, 500\}$. In all cases, we obtained results similar to those in Table 2. Regarding runtime, HyperLocal takes between a few seconds and a few minutes, depending on the target cluster size. This is remarkably fast considering that the method is repeatedly finding minimum s - t cuts in a hypergraph with millions of nodes and hyperedges, where the mean hyperedge size is above 17.

6.4 Varying Splitting Functions

In a survey on flow-based local *graph* clustering, Fountoulakis et al. [14] provide numerous guidelines and examples for choosing input sets R and setting the resolution parameter ε . These guidelines apply in the same fashion to our hypergraph generalization.

We additionally consider how different hypergraph splitting functions affect the output solution. We ran HyperLocal with the δ -linear threshold splitting function for varying δ for a handful of clusters from both datasets, measuring the target cluster recovery F1 score (Fig. 3). We tested both integer and non-integer values of




Figure 3: F1 scores for HyperLocal for varying δ . Each curve corresponds to a different cluster that HyperLocal is trying to detect from a seed set. For six Stack Overflow clusters, nearly all curves are maximized when $\delta > 1000$. For the seven smallest Amazon clusters (left), $\delta = 1.0$ leads to the highest F1 scores except in one case.

δ , and found no meaningful difference in runtime. For the Stack Overflow hypergraph, HyperLocal’s performance plummets for δ near 100 and performance is maximized for very large δ (Fig. 3a). In contrast, for the Amazon hypergraph, the maximum F1 for each cluster is almost always obtained when $\delta = 1.0$ (Fig. 3b). This suggests that if one has access to a subset of ground truth clusters in a hypergraph, running a similar set of experiments provides a simple and effective strategy for choosing a value of δ to use when searching for other, similar clusters.

7 CONCLUSION

We have presented the first strongly-local method for localized hypergraph clustering, which comes with theoretical guarantees in terms of output quality and runtime. One attractive property of our method is that it can leverage recent results on generalized hypergraph s-t cut problems in order to minimize localized ratio cut objectives with general hypergedge splitting functions. Here we have considered one new parametric family of splitting functions (δ -linear threshold). Finding other families of splittings functions could be used to detect other types of meaningful clustering structure in large networks offers opportunity for future research.

ACKNOWLEDGMENTS

This research was supported by a Vannevar Bush Faculty Fellowship, a Simons Investigator grant, NSF Award DMS-1830274, ARO Award W911NF19-1-0057, ARO MURI, and JPMorgan Chase & Co.

REFERENCES

- [1] Sameer Agarwal, Kristin Branson, and Serge Belongie. 2006. Higher Order Learning with Graphs. In *ICML*.
- [2] Sameer Agarwal, Jongwoo Lim, Lihi Zelnik-Manor, Pietro Perona, David Kriegman, and Serge Belongie. 2005. Beyond Pairwise Clustering. In *CVPR*.
- [3] Reid Andersen, Fan Chung, and Kevin Lang. 2006. Local Graph Partitioning using PageRank Vectors. In *FOCS*.
- [4] Reid Andersen and Kevin J. Lang. 2008. An Algorithm for Improving Graph Partitions (*SODA ’08*). Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 651–660. <http://dl.acm.org/citation.cfm?id=1347082.1347154>
- [5] Austin R Benson, Rediet Abebe, Michael T Schaub, Ali Jadbabaie, and Jon Kleinberg. 2018. Simplicial closure and higher-order link prediction. *PNAS* (2018).
- [6] Austin R. Benson, David F. Gleich, and Jure Leskovec. 2016. Higher-order organization of complex networks. *Science* (2016).
- [7] Avrim Blum and Shuchi Chawla. 2001. Learning from Labeled and Unlabeled Data Using Graph Mincuts. In *Proceedings of the Eighteenth International Conference on*

- Machine Learning (ICML ’01)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 19–26.
- [8] T. H. Hubert Chan and Zhibin Liang. 2018. Generalizing the Hypergraph Laplacian via a Diffusion Process with Mediators. In *COCOON*.
- [9] T.-H. Hubert Chan, Anand Louis, Zhihao Gavin Tang, and Chenzi Zhang. 2018. Spectral Properties of Hypergraph Laplacian and Approximation Algorithms. *JACM* (2018).
- [10] Uthsav Chitra and Benjamin J. Raphael. 2019. Random Walks on Hypergraphs with Edge-Dependent Vertex Weights. In *ICML*.
- [11] Fan R. K. Chung. 1997. *Spectral Graph Theory*. American Math. Society.
- [12] Alina Ene, Huy Nguyen, and László A. Végh. 2017. Decomposable Submodular Function Minimization: Discrete and Continuous. In *NeurIPS*.
- [13] Santo Fortunato and Darko Hric. 2016. Community detection in networks: A user guide. *Physics Reports* (2016).
- [14] K. Fountoulakis, M. Liu, D. F. Gleich, and M. W. Mahoney. 2020. Flow-based Algorithms for Improving Clusters: A Unifying Framework, Software, and Performance. [arXiv:cs.LG/2004.09608](https://arxiv.org/abs/2004.09608)
- [15] M. Grötschel, L. Lovász, and A. Schrijver. 1981. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica* (1981).
- [16] Scott W. Hadley. 1995. Approximation techniques for hypergraph partitioning problems. *Discrete Applied Mathematics* (1995).
- [17] Edmund Ihler, Dorothea Wagner, and Frank Wagner. 1993. Modeling hypergraphs by graphs with the same mincut properties. *Inform. Process. Lett.* (1993).
- [18] Jianbo Shi and J. Malik. 2000. Normalized cuts and image segmentation. *TPAMI* (2000).
- [19] Vladimir Kolmogorov. 2012. Minimizing a Sum of Submodular Functions. *Discrete Appl. Math.* (2012).
- [20] Kevin Lang and Satish Rao. 2004. A Flow-Based Method for Improving the Expansion or Conductance of Graph Cuts. In *IPCO*.
- [21] E. L. Lawler. 1973. Cutsets and partitions of hypergraphs. *Networks* (1973).
- [22] Jianbo Li, Jingrui He, and Yada Zhu. 2018. E-tail product return prediction via hypergraph-based local graph cut. In *KDD*. 519–527.
- [23] Pan Li and Olga Milenkovic. 2017. Inhomogeneous Hypergraph Clustering with Applications. In *NeurIPS*.
- [24] Pan Li and Olga Milenkovic. 2018. Revisiting Decomposable Submodular Function Minimization with Incidence Relations. In *NeurIPS*.
- [25] Pan Li and Olga Milenkovic. 2018. Submodular Hypergraphs: p-Laplacians, Cheeger Inequalities and Spectral Clustering. In *ICML*.
- [26] Jianmo Ni, Jiacheng Li, and Julian McAuley. 2019. Justifying Recommendations using Distantly-Labeled Reviews and Fine-Grained Aspects. In *EMNLP-IJCNLP*.
- [27] Lorenzo Orecchia and Zeyuan Allen Zhu. 2014. Flow-based Algorithms for Local Graph Clustering. In *SODA*.
- [28] James B. Orlin. 2009. A faster strongly polynomial time algorithm for submodular function minimization. *Mathematical Programming* (2009).
- [29] James B. Orlin. 2013. Max Flows in $O(nm)$ Time, or Better. In *STOC*.
- [30] Satu Elisa Schaeffer. 2007. Graph clustering. *Computer Science Review* 1, 1 (2007), 27 – 64. <https://doi.org/10.1016/j.cosrev.2007.05.001>
- [31] Alexander Schrijver. 2000. A Combinatorial Algorithm Minimizing Submodular Functions in Strongly Polynomial Time. *J. Comb. Theory Ser. B* (2000).
- [32] Peter Stobbe and Andreas Krause. 2010. Efficient Minimization of Decomposable Submodular Functions. In *NeurIPS*.
- [33] Nate Veldt, Austin R. Benson, and Jon Kleinberg. 2020. Hypergraph Cuts with General Splitting Functions. [arXiv:cs.DS/2001.02817](https://arxiv.org/abs/2001.02817)
- [34] Nate Veldt, David Gleich, and Michael Mahoney. 2016. A Simple and Strongly-Local Flow-Based Method for Cut Improvement. In *ICML*.
- [35] Nate Veldt, Christine Klymko, and David F. Gleich. 2019. Flow-Based Local Graph Clustering with Better Seed Set Inclusion. In *SDM*.
- [36] Dorothea Wagner and Frank Wagner. 1993. Between min cut and graph bisection. In *MFCS*.
- [37] Naganand Yadati, Madhav Nimishakavi, Prateek Yadav, Vikram Nitin, Anand Louis, and Partha Talukdar. 2019. HyperGCN: A New Method For Training Graph Convolutional Networks on Hypergraphs. In *NeurIPS*.
- [38] Hao Yin, Austin R. Benson, Jure Leskovec, and David F. Gleich. 2017. Local Higher-Order Graph Clustering. In *KDD*.
- [39] Dengyong Zhou, Jiayuan Huang, and Bernhard Schölkopf. 2006. Learning with Hypergraphs: Clustering, Classification, and Embedding. In *NeurIPS*.
- [40] J. Y. Zien, M. D. F. Schlag, and P. K. Chan. 1999. Multilevel spectral hypergraph partitioning with arbitrary vertex sizes. *IEEE TCAD* (1999).

8 PROOFS FOR THEOREMS

8.1 Proof of Theorem 3.1

PROOF. We want to show that $\mathbf{cut}_{\mathcal{H}}(S)$ strictly decreases in each pass of the while loop (except the last) in Algorithm 1. To simplify notation, we drop the terms \mathcal{H} , R and ε from function subscripts, since these are fixed and clear from context. Consider a pair of consecutive s - t cut solutions where there is a strict improvement in HLC score. Starting with some value α_{j-1} , let $S_j = \mathbf{argmin} \mathbf{H-st-cut}_{\alpha_{j-1}}(S)$, and $\alpha_j = \mathbf{HLC}(S)$, with $\alpha_j < \alpha_{j-1}$. Let $S_{j+1} = \mathbf{H-st-cut}_{\alpha_j}(S)$ be the set obtained in the next pass through the while loop, with $\alpha_{j+1} = \mathbf{HLC}(S_{j+1})$. Since we are assuming that HLC improves in both of these steps, we have $\alpha_{j+1} < \alpha_j < \alpha_{j-1}$. Now observe that

$$\begin{aligned} & \mathbf{H-st-cut}_{\alpha_{j-1}}(S_j) \\ &= \mathbf{cut}(S_j) + \alpha_{j-1} \mathbf{vol}(\bar{S}_j \cap R) + \alpha_{j-1} \varepsilon \mathbf{vol}(S_j \cap \bar{R}) \\ &= \mathbf{cut}(S_j) - \alpha_{j-1} \mathbf{vol}(S_j \cap R) + \alpha_{j-1} \varepsilon \mathbf{vol}(S_j \cap \bar{R}) + \alpha_{j-1} \mathbf{vol}(R) \\ &= \mathbf{cut}(S_j) - \alpha_{j-1} \Omega(S_j) + \alpha_{j-1} \mathbf{vol}(R) \\ &= \alpha_{j-1} \mathbf{vol}(R) + \Omega(S_j)(\mathbf{HLC}(S_j) - \alpha_{j-1}) \\ &= \alpha_{j-1} \mathbf{vol}(R) + \Omega(S_j)(\alpha_j - \alpha_{j-1}). \end{aligned}$$

The same essential steps show that

$$\mathbf{H-st-cut}_{\alpha_{j-1}}(S_{j+1}) = \alpha_{j-1} \mathbf{vol}(R) + \Omega(S_{j+1})(\alpha_{j+1} - \alpha_{j-1}).$$

We know that $\mathbf{H-st-cut}_{\alpha_{j-1}}(S_j) \leq \mathbf{H-st-cut}_{\alpha_{j-1}}(S_{j+1})$, since S_j is the optimal s - t cut solution for parameter α_{j-1} . This implies that

$$\Omega(S_j)(\alpha_j - \alpha_{j-1}) \leq \Omega(S_{j+1})(\alpha_{j+1} - \alpha_{j-1}),$$

which in turn means that $\Omega(S_{j+1}) < \Omega(S_j)$, since $(\alpha_{j+1} - \alpha_{j-1}) < (\alpha_j - \alpha_{j-1}) < 0$. Finally, because the HLC score and its denominator Ω decrease when going from S_j to S_{j+1} , it must also be the case that $\mathbf{cut}(S_{j+1}) < \mathbf{cut}(S_j)$. Thus, until the last step of Algorithm 1, the cut function is strictly decreasing. \square

8.2 Proof of Theorem 4.1

PROOF. For any $v \in S$, the set of nodes and edges that are adjacent to v in \mathcal{L} is exactly the same as the set of nodes and edges that are adjacent to v in \mathcal{H}_α . The reason is that when the algorithm terminates ($S_L = S_{L-1}$ for some L), any v in the output has either been *explored* (or is in R) and had its neighbors added to \mathcal{L} . Therefore, $\mathbf{L-st-cut}_\alpha(S) = \mathbf{H-st-cut}_\alpha(S)$. Since the function $\mathbf{L-st-cut}_\alpha$ is a lower bound on $\mathbf{H-st-cut}_\alpha$ in general,

$$\begin{aligned} \mathbf{L-st-cut}_\alpha(S) &= \mathbf{H-st-cut}_\alpha(S) \geq \min_A \mathbf{H-st-cut}_\alpha(A) \\ &\geq \min_A \mathbf{L-st-cut}_\alpha(A) = \mathbf{L-st-cut}_\alpha(S). \end{aligned}$$

Thus, $S = \mathbf{argmin}_A \mathbf{H-st-cut}_\alpha(A)$. \square

8.3 Proof of Theorem 4.2

PROOF. While Algorithm 2 does not rely on explicitly applying graph reduction techniques nor computing maximum s - t flows, our proof will rely on the existence of both, as well as on a basic understanding of the max-flow min-cut theorem.

Implicit graph s - t cuts. Let $L^i = (V_i \cup \{s, t\}, E_i \cup E_i^{st})$ be the local hypergraph over which we solve a minimum s - t cut in the i th iteration of the algorithm, where $V_i \subseteq V$, $E_i \subseteq E$, and $E_i^{st} \subseteq E^{st}$.

Let S_i be the minimum s - t cut set for L_i , and let $N_i \subseteq S_i$ be the set of nodes that are *explored* in the i th iteration, i.e., nodes whose terminal edges are cut for the first time in iteration i . Since we assume that the hypergraph cut function is graph reducible, for each L_i there exists a graph G_i with the same set of nodes $V_i \cup \{s, t\}$, plus potentially other auxiliary nodes, such that the minimum s - t cut value in G_i is the minimum s - t cut value in L_i . Formally, let S'_i be the minimum s - t cut set in G_i (excluding s itself), so that $S_i = S'_i \cap V_i$. In other words, if we exclude auxiliary nodes, the minimum s - t cut set in G_i is the minimum s - t cut set in L_i .

Bounding set sizes. Let S be any subset of vertices such that S contains no isolated nodes. Let $E(S, S)$ denote the set of hyperedges that are completely contained inside S . We have the following bounds:

$$|\partial S| \leq \mathbf{vol}_{\mathcal{H}}(S) \quad (23)$$

$$|S| \leq \mathbf{vol}_{\mathcal{H}}(S) \quad (24)$$

$$E(S, S) \leq \frac{\mathbf{vol}_{\mathcal{H}}(S)}{2}. \quad (25)$$

These bounds use the theorem assumptions on the splitting function; the fact that the minimum weight is one means that the volume of a node is equal to its degree. The first bound is tight whenever every edge that is cut by S contains exactly one node from S . The second is tight when every node in S has degree one. Bound (25) follows from the fact that every hyperedge is of size at least 2, and therefore each hyperedge that is completely contained S is made up of at least two nodes from S . For k -uniform hypergraphs, $E(S, S) \leq \frac{\mathbf{vol}_{\mathcal{H}}(S)}{k}$, though we use the bound (25) so that we can apply our results more generally.

Assume Algorithm 2 terminates after iteration t , so that L_t is the largest local hypergraph formed. Let P denote the set of nodes that were *explored* at some point during the algorithm:

$$P = \bigcup_{i=1}^t N_i. \quad (26)$$

We will prove later that the volume of P can be bounded as follows:

$$\mathbf{vol}_{\mathcal{H}}(P) \leq \frac{\mathbf{vol}_{\mathcal{H}}(R)}{\varepsilon}. \quad (27)$$

For now, we assume this to be true and use it to prove the bounds given in the statement of the theorem.

Let Q denote the set of nodes in L_t that were never explored. The size of this set can be bounded as follows:

$$|Q| \leq (k-1)(|\partial R| + |\partial P|). \quad (28)$$

This bound will often be quite loose in practice. However, in theory it is possible for a hyperedge in L_t to contain only one node from $R \cup P$, and $(k-1)$ nodes from the set Q . We bound the total number of nodes in L_t with help from Eqs. (23), (24), and (26):

$$\begin{aligned} |V_t| &= |R| + |P| + |Q| \leq \mathbf{vol}_{\mathcal{H}}(R) + \mathbf{vol}_{\mathcal{H}}(P) + (k-1)(|\partial R| + |\partial P|) \\ &\leq \mathbf{vol}_{\mathcal{H}}(R)(1 + 1/\varepsilon) + (k-1)(\mathbf{vol}_{\mathcal{H}}(R) + \mathbf{vol}_{\mathcal{H}}(P)) \\ &\leq k \mathbf{vol}_{\mathcal{H}}(R)(1 + 1/\varepsilon). \end{aligned}$$

Note also that $|V_t|$ is the exact number of terminal edges in L_t , so we also have a bound on the number of terminal edges.

We can bound the number of hyperedges $|E_t|$ in L_t above by $|E(R, R)| + |E(P, P)| + |\partial R| + |\partial P|$. Note that any hyperedge that

includes a node from Q is accounted for by the terms $|\partial R|$ and $|\partial P|$. We again use bounds (23), (24), (25) and (26) to bound the number of hyperedges in terms of $\mathbf{vol}(R)$:

$$\begin{aligned} |E_t| &\leq |E(R, R)| + |E(P, P)| + |\partial R| + |\partial P| \\ &\leq \frac{3}{2}(\mathbf{vol}_{\mathcal{H}}(R) + \mathbf{vol}_{\mathcal{H}}(P)) \leq \frac{3}{2}\left(1 + \frac{1}{\varepsilon}\right)\mathbf{vol}_{\mathcal{H}}(R). \end{aligned}$$

The last step of the proof is to show the volume bound on P in (26), which we do by proving the existence of a maximum s - t flow on a graph reduction of L_t with certain properties.

Bounding P with an implicit maximum flow argument. The max-flow min-cut theorem states that the value of the minimum s - t cut in a graph G is equal to the maximum s - t flow value in G . An edge is *saturated* if the value of the flow on an edge equals the weight of the edge, which always upper bounds the flow. Given a set of edges C defining a minimum s - t cut G , any maximum s - t flow F in G must saturate *all* edges in C . If any edge in C were not saturated by F , then the flow value would be strictly less than the cut value, contradicting the optimality of either F or C . We will use this understanding of the max-flow min-cut theorem to prove the existence of a flow that saturates all edges of *explored* nodes in the local hypergraph.

The set N_1 is made up of all nodes whose edge to the sink is cut in G_1 when we compute a minimum s - t cut. We know that even if we do not compute it explicitly, there exists some flow F_1 in G_1 that saturates all edges between N_1 and the sink t . In the next iteration, N_2 is the set of nodes whose edges to the sink are cut for the first time. One way to compute a maximum s - t flow F_2 in G_2 is to start with F_1 , the maximum s - t flow in G_1 , and then find new augmenting flow paths until no more flow can be routed from s to t . We can assume without loss of generality that the terminal edges of N_1 remain saturated by F_2 , since there can be no net gain from reversing the flow on a saturated edge to the sink. As a result, the flow F_2 will saturate the terminal edges of N_1 as well as all terminal edges of N_2 , which are cut by the minimum s - t cut in G_2 . Continuing this process inductively, we note that in the i th iteration there exists some flow F_i that saturates *all* the terminal edges to the sink that have been cut by some s - t cut in a previous iteration. In other words, there exists some flow F_t in G_t that saturates the terminal edge of every node in $P = \cup_{i=1}^t N_i$. Recall that

$$\text{Weight of terminal edges of } P = \sum_{v \in P} \alpha \varepsilon d_v = \alpha \varepsilon \mathbf{vol}_{\mathcal{H}}(P). \quad (29)$$

Finally, observe that the minimum s - t cut score in G_t is bounded above by $\alpha \mathbf{vol}_{\mathcal{H}}(R)$, since this is the weight of edges adjacent to the source. This provides an upper bound on the weight of P 's terminal edges, implying the desired bound on the volume of P :

$$\alpha \varepsilon \mathbf{vol}_{\mathcal{H}}(P) \leq \alpha \mathbf{vol}_{\mathcal{H}}(R) \implies \mathbf{vol}_{\mathcal{H}}(P) \leq \frac{\mathbf{vol}_{\mathcal{H}}(R)}{\varepsilon}. \quad (30)$$

□

8.4 Proof of Theorem 5.2

PROOF. Assume throughout that we deal only with sets S satisfying $\Omega_{R, \varepsilon}(S) > 0$, to avoid trivial cases. We begin by defining

$$g(S) = \mathbf{vol}(\bar{R}) \cdot \mathbf{vol}(S \cap R) - \mathbf{vol}(R) \cdot \mathbf{vol}(S \cap \bar{R}). \quad (31)$$

Dividing every term in $g(S)$ by $\mathbf{vol}(\bar{R})$ gives

$$g(S)/\mathbf{vol}(\bar{R}) = \mathbf{vol}(S \cap R) - \varepsilon_0 \mathbf{vol}(S \cap \bar{R}). \quad (32)$$

This allows us to re-write the HLC objective as

$$\mathbf{HLC} = \mathbf{vol}(\bar{R}) \cdot \frac{\mathbf{cut}(S)}{g(S) - \mu \mathbf{vol}(\bar{R}) \cdot \mathbf{vol}(S \cap \bar{R})}. \quad (33)$$

Applying a few steps of algebra produces another useful characterization of the function g :

$$\begin{aligned} g(S) &= \mathbf{vol}(\bar{R}) \cdot \mathbf{vol}(S \cap R) - \mathbf{vol}(R) \cdot \mathbf{vol}(S \cap \bar{R}) \\ &= \mathbf{vol}(V) \cdot \mathbf{vol}(S \cap R) - \mathbf{vol}(R) \cdot \mathbf{vol}(S) \\ &= (\mathbf{vol}(V) - \mathbf{vol}(S)) \cdot \mathbf{vol}(S \cap R) - \mathbf{vol}(S) \cdot \mathbf{vol}(\bar{S} \cap R) \\ &= \mathbf{vol}(\bar{S}) \cdot \mathbf{vol}(S \cap R) - \mathbf{vol}(S) \cdot \mathbf{vol}(\bar{S} \cap R). \end{aligned} \quad (34)$$

This characterization of g allows us to see that for any $S \subset V$:

$$\begin{aligned} \mathbf{vol}(\bar{S}) \cdot \mathbf{vol}(S) &\geq \mathbf{vol}(\bar{S}) \cdot \mathbf{vol}(S \cap R) \geq g(S) \\ &\geq g(S) - \mu \mathbf{vol}(\bar{R}) \cdot \mathbf{vol}(S \cap \bar{R}). \end{aligned} \quad (35)$$

We use this to upper bound $\mathbf{ncut}(S)$ in terms of $\mathbf{HLC}(S)$:

$$\frac{\mathbf{vol}(\bar{R})}{\mathbf{vol}(V)} \cdot \mathbf{ncut}(S) = \mathbf{vol}(\bar{R}) \cdot \frac{\mathbf{cut}(S)}{\mathbf{vol}(S) \cdot \mathbf{vol}(\bar{S})} \leq \mathbf{HLC}(S). \quad (36)$$

Next we need to prove a *lower* bound on the normalized cut score of T . We again use the characterization of g given in (34), this time in conjunction with property (21), satisfied by T , to see that

$$g(T) \geq \beta \mathbf{vol}(T) \cdot \mathbf{vol}(\bar{T}). \quad (37)$$

Property (21) also implies that

$$1 - \frac{\mathbf{vol}(\bar{T} \cap R)}{\mathbf{vol}(T)} = \frac{\mathbf{vol}(T \cap R)}{\mathbf{vol}(T)} \geq \beta \implies \mathbf{vol}(\bar{T} \cap R) \leq (1 - \beta) \mathbf{vol}(T).$$

Combining this with $\mathbf{vol}(\bar{R}) \leq \mathbf{vol}(V) \leq 2\mathbf{vol}(\bar{T})$ produces

$$\mathbf{vol}(\bar{R}) \cdot \mathbf{vol}(T \cap \bar{R}) \leq 2(1 - \beta) \mathbf{vol}(T) \cdot \mathbf{vol}(\bar{T}). \quad (38)$$

Inequalities (37) and (38) together imply that

$$\begin{aligned} \mathbf{vol}(\bar{R}) \cdot \Omega_{R, \varepsilon}(T) &= g(T) - \mu \mathbf{vol}(\bar{R}) \cdot \mathbf{vol}(T \cap \bar{R}) \\ &\geq (\beta - 2\mu(1 - \beta)) \mathbf{vol}(T) \cdot \mathbf{vol}(\bar{T}). \end{aligned} \quad (39)$$

Finally, we put together the bound (36), the characterization of the HLC objective given in (33), and inequality (39), to see that

$$\begin{aligned} \frac{\mathbf{vol}(\bar{R})}{\mathbf{vol}(V)} \cdot \mathbf{ncut}(S^*) &\leq \mathbf{HLC}(S^*) \leq \mathbf{HLC}(T) \\ &= \mathbf{vol}(\bar{R}) \cdot \frac{\mathbf{cut}(T)}{g(T) - \mu \mathbf{vol}(\bar{R}) \cdot \mathbf{vol}(T \cap \bar{R})} \\ &\leq \mathbf{vol}(\bar{R}) \cdot \frac{1}{\beta - 2\mu(1 - \beta)} \cdot \frac{\mathbf{cut}(T)}{\mathbf{vol}(T) \cdot \mathbf{vol}(\bar{T})} \\ &= \frac{\mathbf{vol}(\bar{R})}{\mathbf{vol}(V)} \cdot \frac{\mathbf{ncut}(T)}{\beta - 2\mu(1 - \beta)}. \end{aligned}$$

Dividing through by $\frac{\mathbf{vol}(\bar{R})}{\mathbf{vol}(V)}$ yields the desired bound on normalized cut. □