

A Variational Framework for Budgeted Information-Preserving Subgraph Selection on Weighted Relational Structures

Existence, Stability, and Structural Properties of Optimal Budgeted Subgraphs

Chandrasekhar Gokavarapu⁽¹⁾

ORCID ID: <https://orcid.org/0009-0006-5306-371X>

Dr Duggirala V N S R Murthy⁽²⁾

Dr Ch Srinivasulu⁽³⁾

^(1,2,3)Lecturer in Mathematics, Government College (Autonomous), Rajahmundry, Andhra Pradesh, India, PIN: 533105

Abstract - We develop a finite variational framework for budgeted information-preserving subgraph selection on weighted relational structures. For a fixed ambient structure G , an admissible substructure H is evaluated by an information-retention functional $I_G(H)$ under a complexity budget constraint $C_G(H) \leq B$. Within this framework we prove discrete existence of optimizers, a relaxed compact existence theorem, existence of inclusion-minimal optimal representatives, perturbation stability of the optimal value under uniform continuity hypotheses, and local constancy of a unique optimizer under an isolation-gap condition. We also formulate a quantitative relaxation-and-recovery principle and prove two obstruction results showing that value stability does not imply optimizer stability and that ambient connectedness does not force the existence of a connected optimizer. Finally, we verify the abstract axioms in additive edge-potential and coverage-type submodular model classes. The emphasis throughout is on exact hypothesis tracking, clear theorem boundaries, and a proof-secure finite theory rather than on heuristic graph-pruning procedures.

I. INTRODUCTION

Budget-constrained subgraph selection appears, in one form or another, across several established parts of graph theory and discrete optimization. Foundational work on weighted graph structure, Laplacians, connectivity, and isoperimetric phenomena provides a rich language for measuring how much of an ambient graph is retained by a smaller one [1, 2, 3, 4]. Closely related optimization principles underlie ratio-cut, normalized-cut, and spectral partitioning methods [5, 6, 7, 8], while local clustering methods show that one can often extract a small subgraph carrying a substantial part of the local expansion or diffusion profile of the ambient graph [9, 10, 11, 12]. At the same time, community-detection and network-decomposition literatures treat the problem of selecting informative subnetworks from a large weighted structure from a more algorithmic or heuristic perspective [13, 14].

A second body of work concerns graph reduction, sparsification, and coarsening. Spectral sparsification constructs sparse surrogates that preserve quadratic-form or cut information of the original graph in a provable sense [15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26]. Related lines of research study vertex sparsifiers, cut sparsifiers, directed sparsifiers, hypergraph sparsifiers, and reweighted Cheeger-type principles [27, 28, 29, 30, 31]. Other reduction frame-works seek smaller graphs with controlled spectral, resistive, or structural distortion [32, 33, 34, 35, 36, 37, 38]. A third

neighboring literature comes from graph-based data representation and graph signal processing, where one attempts to retain the relevant information of a weighted relational object under transformation or compression [39, 40, 41, 42]. Finally, whenever the retained information is encoded by a monotone or submodular set functional, one enters the classical theory of budgeted submodular maximization and knapsack-constrained selection [43, 44, 45, 46, 47].

The present paper does not repackage any one of these theories. It isolates a different mathematical question. Existing literatures provide important preservation mechanisms, approximation paradigms, and optimization principles, but they do not by themselves yield a unified variational theory of *budgeted information-preserving subgraph selection* on finite weighted relational structures. In particular, they do not furnish, in a single theorem-secure framework, an axiomatic treatment that simultaneously formalizes

- (i) an information-retention functional,
- (ii) a complexity or budget functional,
- (iii) existence of optimizers over a discrete nonconvex family of substructures,
- (iv) structural properties of optimal subgraphs,
- (v) perturbation stability of optimal values and near-optimizers, and

- (vi) obstruction principles showing where such a theory must fail.

That missing variational core is the object of this article.

Let $G = (V, E, wt)$ be a finite weighted relational structure. We write $\text{Sub}(G)$ for the family of admissible weighted substructures of G . A feasible substructure $H \in \text{Sub}(G)$ is required to satisfy the budget bound

$$C(H) \leq B,$$

where C is a prescribed complexity functional. The retained information of H relative to G is measured by a functional $l(H; G)$, and the corresponding loss is

$$L(H; G) := l(G; G) - l(H; G).$$

The associated optimization problem is

$$\sup \{l(H; G) : H \in \text{Sub}(G), C(H) \leq B\}.$$

At this level, the problem is elementary to state but not mathematically automatic. The family $\text{Sub}(G)$ is finite for each fixed G , but the structure of optimizers is delicate because the admissible family is discrete, the budget restriction is combinatorial, and the retention functional may interact with weights, connectivity, and support constraints in a highly nonlocal way. Moreover, without a disciplined hypothesis package, the formulation quickly becomes either vacuous or overfitted to a single preservation mechanism.

The aim of the paper is therefore not to propose a new heuristic pruning rule. The aim is to formulate and study an axiomatic variational problem on finite weighted relational structures. The central question is the following.

Question I.1. *For which classes of information-retention functionals l and complexity functionals C does the budgeted optimization problem*

$$\sup \{l(H; G) : H \in \text{Sub}(G), C(H) \leq B\}$$

admit (a) optimizers, (b) structural descriptions of optimizers, (c) perturbation stability of optimal values and near-optimizers, and (d) approximation principles under suitable relaxations?

This formulation is motivated by, but not reducible to, several earlier themes. Spectral partitioning studies cut and conductance objectives [5, 7, 3, 8]. Local clustering studies localized diffusion and expansion surrogates [9, 10, 11, 12]. Sparsification theory studies preservation of cuts, resistances, or quadratic forms under edge reduction [15, 16, 17, 19, 21, 26]. Graph reduction and coarsening study smaller surrogates with controlled distortion [32, 33, 34, 35, 36, 37]. Submodular optimization studies budgeted maximization at the level of abstract set functions [43, 44, 47]. None of these theories, however, directly gives the present object: an ambient-graph-dependent, axiomatic optimization theory on *subgraphs themselves*, with simultaneous attention to existence, structure, perturbation, and failure modes. That is the gap filled here.

The paper is organized around four theorem-bearing layers.

- (i) an axiomatic definition of admissible information-retention and complexity functionals;

- (ii) an existence theory for optimal budgeted subgraphs;
- (iii) structural results describing how optimal solutions interact with monotonicity, saturation, connectivity, modularity, or curvature-type hypotheses;
- (iv) a stability and approximation layer controlling the dependence of optimal values and near-optimizers on perturbations of the ambient weighted structure.

The article also includes an obstruction layer. This is necessary. A theory of this type is not credible unless it records precisely which natural hypotheses cannot be removed without destroying existence or stability.

Position of the paper in the literature

The paper lies at the intersection of finite variational theory, weighted graph theory, and discrete optimization on structured families of subgraphs. Its novelty claim is modest and specific. It does *not* claim a complete classification of all information-preserving subgraph problems. It does *not* claim to subsume the full sparsification literature, the full graph reduction literature, or the full submodular optimization literature. Instead, it isolates a proved core that is not presently available in those theories taken separately.

The graph-theoretic side of the background comes from weighted connectivity, Laplacian structure, directed Laplacians, spectral embeddings, and cut principles [1, 2, 3, 4, 39, 7, 5, 6, 8]. These works show that weighted substructures may preserve mathematically meaningful information, but they do not formulate a general budgeted optimization theory over abstract subgraph families. The local partitioning and community-detection literature gives further evidence that small, informative subgraphs are natural objects [9, 10, 11, 12, 13, 14], but again the main emphasis there is algorithmic extraction, not axiomatic variational structure.

The preservation side of the background comes from sparsification, cut preservation, resistance preservation, streaming subgraph sketches, vertex sparsifiers, directed sparsifiers, and hypergraph sparsifiers [15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 27, 28, 29, 30, 31, 26]. These works prove strong preservation theorems for specific objectives. They are indispensable background. But their objects are usually sparse surrogates preserving a designated analytic quantity, not optimizers of a general ambient-dependent information-retention functional under a separate complexity budget. Our framework is therefore adjacent to this literature, but not a restatement of it.

The reduction side of the background comes from Kron reduction, spectral simplification, spectral reduction, graph coarsening, and recent graph-reduction surveys [32, 33, 34, 35, 36, 37, 38]. These works show that one can replace a large graph by a smaller one while controlling selected distortions. They do not, however, by themselves provide an abstract existence–structure–stability theory for optimal subgraphs defined by a retention functional and a budget functional on a general weighted relational structure.

The optimization side of the background comes from submodular maximization and knapsack-constrained selection [43, 44, 45, 46, 47]. This literature supplies an important comparison point. If $l(\cdot; G)$ reduces, after encoding, to a monotone submodular set function on edges or vertices, then approximation phenomena become accessible. But the present paper does not assume submodularity a priori, and

several of the structural and stability questions considered here do not belong to classical submodular maximization in that form. Submodularity, when it appears, is treated as one sufficient mechanism among others, not as the definition of the subject.

A further neighboring background comes from graph signal processing and graph-based representation theory [40, 41, 42]. This literature reinforces the idea that a weighted relational structure carries information that may or may not survive reduction. What it does not supply is a finite variational theory of optimal retained subgraphs under an explicit complexity budget.

Accordingly, the paper should be read as a foundational article. Its contribution is to define a mathematically controlled class of budgeted information-preserving subgraph problems and to prove a first layer of general theorems for that class. The scope is intentionally narrower than a universal theory and broader than a single preservation mechanism. That scope control is deliberate.

Main provisional contributions

The following list is stated in referee-facing form. Each item is to be retained only at the level fully supported by the proofs in the final manuscript.

- (C1) A precise axiomatic class of information-retention functionals I and budget functionals C on finite weighted relational structures, formulated so that the optimization problem is mathematically nontrivial yet theorem-accessible.
- (C2) A proved existence theorem for optimal budgeted subgraphs under a transparent hypothesis package. In the finite setting, existence is not the issue by itself; the issue is the correct formulation of the admissible family and retention functional so that the optimization problem is nondegenerate and structurally meaningful.
- (C3) A structure theorem under additional assumptions identifying when optimizers may be chosen inside a distinguished subclass, for example connected, inclusion-minimal, support-saturated, or otherwise canonical representatives.
- (C4) A perturbation theorem for optimal values, together with a near-optimizer stability statement when the hypothesis package is strong enough to justify such a conclusion. The theorem is intended to make explicit which continuity or semicontinuity assumptions are actually used.
- (C5) An approximation or relaxation theorem connecting the discrete budgeted problem to a relaxed variational problem under additional assumptions such as monotonicity, modularity, or submodular-type control, in dialogue with classical budgeted set-function optimization [43, 44, 47].
- (C6) An obstruction theorem showing that without an appropriate monotonicity, continuity, or admissibility hypothesis, existence, canonical structure, or stability can fail. This negative result is part of the mathematical contribution, not a peripheral remark.
- (C7) Verification results in one or two substantial natural model classes. This is needed to prevent the ax-

iomatic framework from becoming formally correct but mathematically detached.

Methodological remarks

The paper is written as a pure theoretical mathematics article. It contains no empirical section, no simulation-based validation, and no claims of practical superiority over existing graph-pruning methods. Algorithmic material, when included, is admitted only at the level of formal approximation mechanisms attached directly to proved statements. The article also avoids inflated language. Terms such as “optimal,” “sharp,” “minimal,” or “classification” are used only where the corresponding proof justifies them.

The basic strategy is to separate three layers throughout:

- (i) proved results and corollaries,
- (ii) interpretive remarks explaining their relation to existing theories, and
- (iii) open problems recording what the present framework does not yet settle.

This separation is essential for theorem integrity.

Road map

Section II fixes notation for weighted relational structures, substructures, admissible families, and the underlying order-theoretic language. **Section III** defines the budgeted optimization problem and isolates the admissible hypothesis packages for I and C . **Section IV** proves the first existence theorem and identifies the exact role of finiteness, admissibility, and upper semicontinuity. **Section V** studies structural properties of optimizers under additional monotonicity, saturation, and canonical-representative hypotheses. **Section VI** proves perturbation results for optimal values and, where justified, for near-optimizers. **Section VII** introduces relaxed problems and approximation principles. **Section VIII** records failure modes and separation phenomena showing that the main hypotheses are not formal decoration. **Section IX** verifies the abstract axioms in natural model classes informed by preservation and reduction mechanisms from the surrounding literature [16, 17, 32, 36]. The final section separates proved conclusions from open problems.

II. PRELIMINARIES AND NOTATION

This section fixes the discrete framework used throughout the paper. The purpose is not merely terminological. The later existence, structure, stability, and approximation theorems depend on a careful separation between the ambient weighted object, the admissible family of substructures, the complexity budget, and the information-retention functional. We therefore fix the notation once and keep it stable.

The graph-theoretic background is standard at the level of weighted graphs, Laplacians, connectivity, and cut-based structure; see, for example, [1, 2, 3, 4, 8]. Several neighboring literatures motivate particular choices of retention functionals or admissible substructures, including local graph partitioning [9, 10, 11, 12], spectral sparsification and cut preservation [15, 16, 17, 19, 21, 26], graph reduction and coarsening [32, 33, 34, 35, 36, 37], and budgeted set-function optimization [43, 44, 47]. The present section does not import the conclusions of those theories wholesale. It uses them only as background for reasonable hypothesis packages.

1. Ambient weighted structures and admissible substructures

Definition II.1 (Weighted relational structure). A *finite weighted relational structure* is a triple

$$G = (V, E, \text{wt}),$$

where V is a finite nonempty vertex set, $E \subseteq V \times V$ is a finite relation set, and

$$\text{wt} : E \rightarrow [0, \infty)$$

is a weight function. If $(u, v) \in E$, the quantity $\text{wt}(u, v)$ is called the weight of the relation (u, v) .

When no ambiguity arises, we write $V(G)$, $E(G)$, and wt_G for the vertex set, relation set, and weight function of G .

Remark II.2. The formulation above includes finite weighted directed graphs as a special case. If E is symmetric, then G may be regarded as a weighted undirected graph. This flexibility is deliberate. Directed Laplacian methods and directed Cheeger-type phenomena show that some natural retention principles are genuinely asymmetric [4, 14, 42, 31]. On the other hand, many of the prototype examples in later sections arise from undirected weighted graphs, in the spirit of [1, 3, 16, 36].

Definition II.3 (Weighted substructure). Let $G = (V, E, \text{wt})$ be a finite weighted relational structure. A *weighted substructure* of G is a triple

$$H = (V_H, E_H, \text{wt}_H)$$

such that

$$V_H \subseteq V, \quad E_H \subseteq E \cap (V_H \times V_H),$$

and

$$\text{wt}_H = \text{wt}|_{E_H}.$$

We write $H \preceq G$ if H is a weighted substructure of G .

Remark II.4. The requirement $\text{wt}_H = \text{wt}|_{E_H}$ means that substructures inherit the ambient weights rather than re-optimizing them. This choice matches the intended interpretation of H as a retained part of G , not as an independently reweighted surrogate. In contrast, some sparsification and reduction frameworks allow changed edge weights to preserve spectral or cut quantities [17, 19, 32, 34, 35]. Those weighted-reduction models will appear later only as comparison mechanisms or as natural classes for abstract verification, not as the default meaning of substructure.

Notation II.5 (Substructures and feasibility). For a fixed ambient structure G , let $\text{Sub}(G)$ denote a prescribed admissible class of weighted substructures of G . The class $\text{Sub}(G)$ must be fixed before the first main theorem. Typical choices include:

- all weighted spanning subgraphs of G ;
- all vertex-induced weighted subgraphs of G ;
- all relation-induced weighted substructures of G ;
- all connected weighted substructures of one of the preceding types;

- all substructures satisfying an additional closure or support condition.

Whenever $H_1, H_2 \in \text{Sub}(G)$, the notation

$$H_1 \preceq H_2$$

means that H_1 is a weighted substructure of H_2 in the sense of Definition II.3.

Remark II.6. The choice of $\text{Sub}(G)$ is part of the mathematical data. It is not a cosmetic convenience. If one works with all spanning subgraphs, then edge deletion is the primitive operation. If one works with vertex-induced subgraphs, then vertex selection is primary. If one imposes connectivity or saturation, the feasible family is no longer downward closed. These distinctions affect existence, canonical-representative results, and approximation behavior. The local partitioning, sparsification, coarsening, and community-detection literatures each privilege different admissible families [9, 10, 17, 13, 36, 38].

2. Complexity, retention, loss, and optimal value

Definition II.7 (Complexity functional). Let G be fixed. A *complexity functional* on $\text{Sub}(G)$ is a map

$$C : \text{Sub}(G) \rightarrow [0, \infty).$$

For a budget level $B \geq 0$, the associated feasible family is

$$\mathcal{F}_B(G) := \{H \in \text{Sub}(G) : C(H) \leq B\}.$$

Remark II.8. The functional C may encode different mathematically meaningful notions of budget. Prototypical examples are:

- edge budget: $C(H) = |E(H)|$;
- vertex budget: $C(H) = |V(H)|$;
- weighted support budget: $C(H) = \sum_{e \in E(H)} \phi(\text{wt}(e))$ for a fixed nonnegative function ϕ ;
- description-length surrogate: $C(H) = a|V(H)| + b|E(H)|$ with fixed constants $a, b \geq 0$.

In later sections, the choice of C will be treated axiomatically unless a theorem requires a specific form. Budget monotonicity is natural for most such examples, but it must be stated rather than assumed implicitly.

Definition II.9 (Information-retention functional). Let G be fixed. An *information-retention functional* is a map

$$l(\cdot; G) : \text{Sub}(G) \rightarrow \mathbb{R},$$

where $l(H; G)$ is interpreted as the amount of ambient information retained by the candidate substructure H .

Definition II.10 (Information-loss functional). Given an information-retention functional $l(\cdot; G)$, define the associated *information-loss functional* by

$$L(H; G) := l(G; G) - l(H; G), \quad H \in \text{Sub}(G).$$

Remark II.11. The framework deliberately leaves l abstract at the foundational stage. This is necessary because different

neighboring theories privilege different notions of retained information. Depending on the model class, one may take $l(H; G)$ to measure retained cut information, spectral information, local diffusion information, resistance-type information, signal energy, or a monotone set-function surrogate; compare [7, 16, 17, 40, 41, 42, 43, 47]. The present paper does not assume from the outset that these examples belong to one common analytic class. One of the aims is to isolate which common axioms suffice for theorem-level conclusions.

Definition II.12 (Optimal value and maximizers). For fixed (G, B) , define the *optimal value*

$$\text{Opt}(G, B) := \sup \{l(H; G) : H \in \mathcal{F}_B(G)\}.$$

The set of maximizers is

$$\text{argmax}(G, B) := \{H \in \mathcal{F}_B(G) : l(H; G) = \text{Opt}(G, B)\}.$$

Remark II.13. If $\text{Sub}(G)$ is finite and $\mathcal{F}_B(G) \neq \emptyset$, then the supremum in Definition II.12 is automatically a maximum. That elementary fact will be used, but it is not the end of the theory. The real issues are:

- (i) whether the optimization problem is formulated on a natural admissible family;
- (ii) whether the maximizing set admits structural description;
- (iii) whether the value function and the optimizer set are stable under perturbation of the ambient weights;
- (iv) whether relaxed formulations approximate the discrete problem;
- (v) whether the chosen axioms are verified in natural classes.

Finite-choice existence alone does not address these questions. This distinction is important in view of the richer stability and reduction phenomena studied in [17, 19, 32, 36, 37].

3. Order structure and basic set-valued notation

Definition II.14 (Inclusion order). Let G be fixed. The admissible family $\text{Sub}(G)$ is partially ordered by the relation \preceq from Notation II.5. If $H_1, H_2 \in \text{Sub}(G)$ and $H_1 \preceq H_2$, we say that H_1 is *contained in* H_2 .

Definition II.15 (Minimal and maximal feasible elements). Let G and B be fixed.

- (i) A feasible substructure $H \in \mathcal{F}_B(G)$ is called *inclusion-minimal feasible* if whenever $K \in \mathcal{F}_B(G)$ and $K \preceq H$, one has $K = H$.
- (ii) A feasible substructure $H \in \mathcal{F}_B(G)$ is called *inclusion-maximal feasible* if whenever $K \in \mathcal{F}_B(G)$ and $H \preceq K$, one has $K = H$.

Remark II.16. Minimality or maximality in the inclusion order is unrelated, a priori, to optimality for $l(\cdot; G)$. Later structure theorems will identify additional assumptions under which maximizers may be chosen from distinguished subclasses, such as inclusion-minimal maximizers, connected maximizers, or support-saturated maximizers. One should not assume such conclusions for free.

Definition II.17 (Near-optimizers). For $\varepsilon \geq 0$, define the ε -near-optimizer set by

$$\text{argmax}_\varepsilon(G, B) := \{H \in \mathcal{F}_B(G) : l(H; G) \geq \text{Opt}(G, B) - \varepsilon\}.$$

Remark II.18. The set $\text{argmax}_\varepsilon(G, B)$ is useful even in the finite setting. Under perturbation, true maximizers may change discontinuously even when optimal values vary mildly. Near-optimizer sets are therefore the natural set-valued objects in stability statements, just as approximation and perturbation theory in neighboring graph-reduction and submodular settings often tracks approximate rather than exact solutions [44, 47, 36, 38].

4. Metrics and perturbations

Definition II.19 (Ambient perturbation metric). Fix a finite vertex set V and relation set $E \subseteq V \times V$. If

$$G = (V, E, \text{wt}) \quad \text{and} \quad G' = (V, E, \text{wt}')$$

are weighted relational structures on the same support, define

$$d_\infty(G, G') := \max_{e \in E} |\text{wt}(e) - \text{wt}'(e)|.$$

More generally, if one wishes to vary supports as well, one must specify a support-compatibility rule together with a metric or pseudometric on the resulting class. In the present paper, the default perturbative setting keeps the underlying support fixed and perturbs only the weights.

Remark II.20. The fixed-support perturbation model is the cleanest one for a first general theory. It isolates the dependence on weights and avoids conflating weight instability with support instability. This choice is also consistent with several spectral and signal-processing frameworks in which the combinatorial scaffold is fixed while weighted operators vary [39, 40, 41, 42]. Later sections may indicate how the framework could be enlarged to support-changing perturbations, but such an enlargement is not part of the present base theorem package.

Definition II.21 (Weight continuity of functionals). A complexity functional C and an information-retention functional l are said to be *weight-continuous on a model class* if for every fixed support (V, E) there exist moduli of continuity

$$\omega_C, \omega_l : [0, \infty) \rightarrow [0, \infty), \quad \lim_{t \downarrow 0} \omega_C(t) = \lim_{t \downarrow 0} \omega_l(t) = 0,$$

such that for all weighted structures $G = (V, E, \text{wt})$ and $G' = (V, E, \text{wt}')$ in the class and all admissible substructures $H \in \text{Sub}(G) \cap \text{Sub}(G')$,

$$|C_G(H) - C_{G'}(H)| \leq \omega_C(d_\infty(G, G'))$$

and

$$|l_G(H) - l_{G'}(H)| \leq \omega_l(d_\infty(G, G')).$$

Remark II.22. The continuity notion above is intentionally uniform over admissible substructures on fixed support. This is stronger than pointwise continuity, but it is the form naturally used in perturbation estimates for optimal values. Similar uniform-control ideas appear, in more specialized settings, in spectral sparsification, effective-resistance preservation, and graph reduction [16, 26, 34, 36].

5. Hypothesis packages

We now record the basic axioms that will be invoked later. They are deliberately separated because different theorems use different subsets of them.

(H1) **Feasibility finiteness.** For each fixed pair (G, B) , the feasible family $\mathcal{F}_B(G)$ is nonempty and finite.

(H2) **Monotonicity of information.** If $H_1, H_2 \in \text{Sub}(G)$ and $H_1 \preceq H_2 \preceq G$, then

$$l(H_1; G) \leq l(H_2; G).$$

(H3) **Budget monotonicity.** If $H_1, H_2 \in \text{Sub}(G)$ and $H_1 \preceq H_2$, then

$$C(H_1) \leq C(H_2).$$

(H4) **Weight continuity.** The functionals l and C are weight-continuous in the sense of Definition II.21 on the model class under consideration.

(H5) **Supplementary structural axiom.** A further condition is imposed when needed. Depending on the theorem, this may be a modularity, submodularity, curvature-type, connectivity-preservation, or saturation axiom.

Remark II.23. These hypotheses are not declared globally necessary. Their status will be specified theorem by theorem. In particular:

- (i) (H1) is proved sufficient for the elementary existence statement but is not presented as conceptually minimal.
- (ii) (H2) and (H3) are natural order-theoretic assumptions that support canonical-representative arguments; they are not automatic in all model classes.
- (iii) (H4) is the basic perturbative hypothesis for stability of optimal values and near-optimizers.
- (iv) (H5) is a placeholder for extra structure used only where needed. In examples inspired by budgeted set-function optimization, this may take a submodularity form [43, 44, 45, 46, 47]. In examples inspired by spectral or reduction theory, it may take a support-saturation or operator-control form [17, 19, 32, 36].

Necessity questions are left open unless explicitly settled by an obstruction theorem later in the paper.

6. Prototype model classes

The axiomatic framework is abstract, but it is useful to keep several prototype classes in view.

Example II.24 (Edge-budget spanning-subgraph model). Let $\text{Sub}(G)$ be the class of all spanning weighted subgraphs of G , and define

$$C(H) := |E(H)|.$$

This is the basic edge-budget model. It is the closest discrete analogue of graph sparsification, though the present theory does not assume that l is given by a quadratic-form or cut-preservation objective [15, 16, 17, 19].

Example II.25 (Vertex-induced budget model). Let $\text{Sub}(G)$ be the class of all vertex-induced weighted subgraphs $G[U]$ with $U \subseteq V(G)$, and let

$$C(G[U]) = |U|.$$

This model is close in spirit to local graph partitioning and community-extraction settings, where the selected object is a small vertex set together with its induced internal structure [9, 10, 11, 13, 14].

Example II.26 (Additive surrogate model). Suppose

$$l(H; G) = \sum_{e \in E(H)} a_G(e), \quad C(H) = \sum_{e \in E(H)} b_G(e),$$

for nonnegative edge scores $a_G(e)$ and $b_G(e)$. Then the optimization problem reduces to a weighted discrete selection problem on edges. This model is too simple to represent the full intended theory, but it serves as a sanity check and as a comparison point for later obstruction results.

Remark II.27. The natural-class verification section later in the paper will examine whether specific retention mechanisms arising from spectral, cut, reduction, or submodular settings satisfy the abstract hypotheses recorded here. At the preliminary stage we only need the formal framework, not the final verification.

Notation table

The following notation will remain fixed.

$G = (V, E, \text{wt})$	finite weighted relational structure
$H \preceq G$	admissible weighted substructure of G
$\text{Sub}(G)$	admissible class of substructures of G
$C(H)$	complexity of H
B	complexity budget
$\mathcal{F}_B(G)$	feasible family under budget B
$l(H; G)$	information retained by H relative to G
$l(H; G)$	information loss relative to G
$\text{Opt}(G, B)$	optimal retained information under budget B
$\text{argmax}(G, B)$	set of maximizers of $l(\cdot; G)$ on $\mathcal{F}_B(G)$
$\text{argmax}_\varepsilon(G, B)$	ε -near-optimizer set
$d_\infty(G, G')$	sup-norm distance between weight functions on fixed support

Remark II.28. This section fixes only the base framework. It does not yet impose the final hypothesis package for any flagship theorem. The next section defines the optimization problem in its theorem-bearing form and classifies the admissibility assumptions more sharply.

III. THE VARIATIONAL PROBLEM AND ADMISSIBLE AXIOMS

This section converts the optimization problem from the introduction into a precise discrete variational framework. The guiding point is simple. One should not speak about “optimal information-preserving subgraphs” until the ambient class of admissible substructures, the budget functional, the retention functional, and the perturbation regime have all been fixed. In neighboring literatures these ingredients are often present implicitly but in different combinations: sparsification emphasizes preservation of quadratic forms, cuts, or resistances [15, 16, 17, 19, 21, 26]; graph reduction and coarsening emphasize smaller surrogates with controlled

distortion [32, 33, 34, 35, 36, 37]; local partitioning emphasizes small subgraphs with strong local conductance or diffusion content [9, 10, 11, 12]; submodular optimization emphasizes budgeted maximization of set functions [43, 44, 45, 46, 47]. The present paper does not identify any one of these existing objectives with the definition of the subject. Instead, it isolates a common axiomatic variational core.

1. Admissible systems

Definition III.1 (Admissible information-retention system). An *admissible information-retention system* \mathfrak{S} assigns to each finite weighted relational structure

$$G = (V, E, \text{wt})$$

the following data:

- (A1) a family $\text{Sub}(G)$ of admissible weighted substructures of G ;
- (A2) a complexity functional

$$C_G : \text{Sub}(G) \rightarrow [0, \infty);$$

- (A3) an information-retention functional

$$l_G : \text{Sub}(G) \rightarrow \mathbb{R};$$

- (A4) a perturbation metric or pseudometric d_{Rel} on ambient weighted structures of the same combinatorial type.

These data are subject to additional axioms imposed according to the depth of the theorem under consideration.

Remark III.2. The system-level formulation is necessary because later statements compare optimization problems across different ambient weighted structures. In particular, perturbation results do not concern a single pair (G, B) in isolation; they concern the behavior of the assignment

$$G \mapsto (\text{Sub}(G), C_G, l_G, d_{\text{Rel}}).$$

This point is routine in more specialized graph-theoretic settings where the ambient operator or weighted structure varies [3, 4, 40, 41, 42], but it must be stated explicitly here because the present framework is more abstract.

Remark III.3. Once an admissible system \mathfrak{S} has been chosen, the symbols $\text{Sub}(G)$, C_G , and l_G are no longer free notation. They are part of the mathematical model. Different admissible systems may yield different optimization problems even on the same ambient weighted structure G . For example, one may keep the same admissible family and change the retention functional from a cut-based quantity to a spectral surrogate, or keep the same retention functional and change from edge-budget to vertex-budget complexity. Such changes alter the variational problem and should never be treated as harmless notational variants.

2. Primal optimization and loss minimization

Definition III.4 (Feasible family). Let \mathfrak{S} be an admissible information-retention system, let G be a finite weighted relational structure, and let $B \geq 0$. The associated feasible family is

$$\mathcal{F}_B(G) := \{H \in \text{Sub}(G) : C_G(H) \leq B\}.$$

Definition III.5 (Budgeted primal problem). Let \mathfrak{S} , G , and B be as above. The *budgeted primal variational problem* is

$$\text{Opt}(G, B) := \sup \{l_G(H) : H \in \mathcal{F}_B(G)\}.$$

Its maximizer set is

$$\text{argmax}(G, B) := \{H \in \mathcal{F}_B(G) : l_G(H) = \text{Opt}(G, B)\}.$$

Definition III.6 (Loss minimization form). Assume that $l_G(G)$ is finite. Define the associated *loss functional*

$$L_G(H) := l_G(G) - l_G(H), \quad H \in \text{Sub}(G).$$

Then the budgeted maximization problem may equally be regarded as a loss-minimization problem over the same feasible family.

Remark III.7. The loss formulation is conceptually useful for two reasons. First, it normalizes the ambient object G as the zero-loss reference point. Second, many preservation theories are naturally phrased in terms of controlled distortion, error, or deficit rather than retained value; compare cut distortion, spectral distortion, resistance error, and approximation ratios in sparsification and reduction theory [16, 17, 19, 32, 36]. The present paper does not identify L_G with any one of those notions, but the loss language is structurally aligned with them.

Proposition III.8 (Equivalent primal forms). *Suppose $l_G(G)$ is finite. Then*

$$\text{argmax} \{l_G(H) : H \in \mathcal{F}_B(G)\} = \text{argmin} \{L_G(H) : H \in \mathcal{F}_B(G)\}.$$

In particular,

$$\sup \{l_G(H) : H \in \mathcal{F}_B(G)\} = l_G(G) - \inf \{L_G(H) : H \in \mathcal{F}_B(G)\}.$$

Proof. For every $H \in \mathcal{F}_B(G)$ one has $L_G(H) = l_G(G) - l_G(H)$, so the map $x \mapsto l_G(G) - x$ is strictly decreasing and therefore converts maximizers of l_G on $\mathcal{F}_B(G)$ exactly into minimizers of L_G on $\mathcal{F}_B(G)$. \square

Remark III.9. Proposition III.8 is only an order-preserving reformulation. It is not a duality theorem, and the paper should not use dual language unless a later section introduces a genuine dual problem with a proved relation. The distinction matters because much of the surrounding optimization literature employs primal-dual or convex-analytic terminology in settings where the present discrete subgraph problem has not yet been relaxed [44, 47].

3. Admissible axioms

We now isolate the axiom packages used later. These are not declared globally necessary. They are recorded as theorem-dependent hypotheses.

Definition III.10 (Base admissibility axioms). Let \mathfrak{S} be an admissible information-retention system. We say that \mathfrak{S} satisfies:

- (B1) **Feasible nonemptiness and finiteness** if for every ambient weighted structure G and every budget $B \geq 0$ under consideration, the family $\mathcal{F}_B(G)$ is nonempty and finite.

(B2) **Information monotonicity** if for every fixed G and all $H_1, H_2 \in \text{Sub}(G)$ with $H_1 \preceq H_2 \preceq G$,

$$l_G(H_1) \leq l_G(H_2).$$

(B3) **Budget monotonicity** if for every fixed G and all $H_1, H_2 \in \text{Sub}(G)$ with $H_1 \preceq H_2$,

$$C_G(H_1) \leq C_G(H_2).$$

(B4) **Normalization** if for every G ,

$$l_G(G) \geq l_G(H) \quad \text{for all } H \in \text{Sub}(G).$$

(B5) **Weight continuity** if for every fixed combinatorial type there exists a modulus of continuity controlling the change of $l_G(H)$ and $C_G(H)$ under perturbation of the ambient weights, uniformly over admissible H on that support.

(B6) **Supplementary structural control** if one imposes, where needed, a further axiom such as modularity, submodularity, curvature-type control, connectivity inheritance, saturation, or a canonical-representative property.

Remark III.11. The normalization axiom is convenient but should not be confused with monotonicity. If $G \in \text{Sub}(G)$ and information monotonicity holds, then normalization is automatic. However, the framework allows admissible families that do not contain every intermediate enlargement, or that impose structural side conditions. In such settings it is cleaner to isolate normalization as a separate statement. This is consistent with the general discipline of the paper: hypotheses that are logically distinct are recorded separately rather than bundled together without comment.

Remark III.12. The supplementary structural axiom is deliberately open-ended at this stage. In examples motivated by submodular maximization, one may require a diminishing-returns condition after encoding the chosen substructure by an edge or vertex set [43, 44, 45, 46, 47]. In examples motivated by sparsification or reduction, one may instead require a support-saturation or operator-comparison condition [16, 17, 19, 32, 34, 36]. The current section does not fix one universal choice because later theorems use different structural mechanisms.

Definition III.13 (Strong admissibility package). An admissible information-retention system \mathfrak{S} is said to satisfy the *strong admissibility package* on a model class if it satisfies all axioms in Definition III.10, together with the requirement that the perturbation metric d_{Rel} be complete on each fixed finite support class under consideration.

Remark III.14. The completeness clause in Definition III.13 is irrelevant for the elementary finite-choice existence theorem. It becomes potentially useful only when relaxed feasible regions, parameterized families, or limiting perturbation arguments enter. It is therefore included as a forward-looking strengthening, not as part of the minimal base package.

4. Value function and budget dependence

Definition III.15 (Value function). For a fixed admissible system \mathfrak{S} and a fixed ambient weighted structure G , define the *value function*

$$B \mapsto \text{Opt}(G, B)$$

on the set of admissible budgets by

$$\text{Opt}(G, B) := \sup \{ l_G(H) : H \in \text{Sub}(G), C_G(H) \leq B \}.$$

Proposition III.16 (Monotonicity in the budget). *For every fixed G , the value function $B \mapsto \text{Opt}(G, B)$ is nondecreasing on its domain.*

Proof. If $0 \leq B_1 \leq B_2$, then $\mathcal{F}_{B_1}(G) \subseteq \mathcal{F}_{B_2}(G)$ by definition. Taking suprema of the same objective over nested feasible sets gives

$$\text{Opt}(G, B_1) \leq \text{Opt}(G, B_2). \quad \square$$

Remark III.17. Proposition III.16 is elementary, but it identifies the first variational regularity property of the model. Later sections ask for stronger control, such as plateau structure, canonical breakpoints, or stability of near-optimizers as the ambient weights vary. Such finer questions are not automatic even in finite discrete problems.

Proposition III.18 (Upper bound under normalization). *Assume the normalization axiom of Definition III.10. Then for every feasible budget B ,*

$$\text{Opt}(G, B) \leq l_G(G).$$

Equivalently,

$$\inf \{ l_G(H) : H \in \mathcal{F}_B(G) \} \geq 0.$$

Proof. By normalization, $l_G(H) \leq l_G(G)$ for all $H \in \text{Sub}(G)$, hence in particular for all $H \in \mathcal{F}_B(G)$. Taking the supremum over $\mathcal{F}_B(G)$ gives the first inequality. The second follows from Definition III.6. \square

5. Relaxed feasible regions

Definition III.19 (Relaxed feasible region). Suppose that admissible substructures in a chosen model class admit an encoding

$$\iota_G : \text{Sub}(G) \rightarrow X_G$$

into a finite-dimensional real vector space X_G . A *relaxed feasible region* for the budget level B is a set

$$\overline{\mathcal{F}_B(G)}^{\text{rel}} \subseteq X_G$$

such that

$$\iota_G(\mathcal{F}_B(G)) \subseteq \overline{\mathcal{F}_B(G)}^{\text{rel}}.$$

If convex-analytic language is used later, the ambient linear space X_G , the embedding ι_G , and the precise definition of $\overline{\mathcal{F}_B(G)}^{\text{rel}}$ must be fixed before any such argument is invoked.

Remark III.20. The relaxed layer is optional and should not be introduced merely for stylistic breadth. It is justified only if the paper proves a theorem that genuinely uses it. This caution is important because the surrounding literature

contains several different relaxation mechanisms—spectral, semidefinite, convex-hull, diffusion-based, and set-function multilinear extensions—and these are not interchangeable [7, 8, 44, 47, 36]. The present article will use relaxation only at the exact level required by proved statements.

Remark III.21. Typical encodings include edge-incidence vectors, vertex-incidence vectors, weighted adjacency tensors, or block-structured operator encodings. In sparsification and reduction theory, weighted subgraphs are often encoded by matrices or Laplacian operators [17, 19, 21, 32]. In submodular optimization they are encoded by subsets of a ground set [43, 47]. The present framework accommodates both viewpoints, provided the encoding is fixed before theorems are stated.

6. Parameter dependence and perturbative admissibility

Definition III.22 (Same combinatorial type). Two weighted relational structures

$$G = (V, E, wt), \quad G' = (V', E', wt')$$

are said to have the *same combinatorial type* if there is a prescribed identification of V with V' and of E with E' under which the only varying datum is the weight function.

Definition III.23 (Perturbatively admissible family). Let \mathcal{G} be a class of weighted relational structures of the same combinatorial type. An admissible information-retention system \mathfrak{S} is *perturbatively admissible* on \mathcal{G} if:

- (i) the admissible family $\text{Sub}(G)$ is canonically identified across all $G \in \mathcal{G}$;
- (ii) the metric d_{Rel} is defined on \mathcal{G} ;
- (iii) the functionals C_G and l_G satisfy the weight-continuity axiom on \mathcal{G} .

Remark III.24. Definition III.23 is the correct framework for value-stability results in later sections. Without canonical identification of admissible substructures across nearby ambient structures, even the statement of near-optimizer stability becomes ambiguous. This is one reason the paper fixes perturbation of weights on constant support as the default base regime.

7. Hypothesis-status paragraph for the next section

Hypothesis-status note III.25. For the basic existence theory in the next section, feasible nonemptiness and finiteness are sufficient. No continuity hypothesis is needed for that elementary result. Information monotonicity, budget monotonicity, and supplementary structural control are likewise unnecessary at the first existence level. Weight continuity becomes relevant only when one passes to perturbative statements or to relaxed feasible regions. Completeness of the ambient perturbation metric is not used in the finite-choice existence argument. No necessity claim is made here. Any necessity statement must be justified later by an explicit obstruction theorem or separation result.

Remark III.26. The purpose of the present section is not to overstate the framework. It is to separate clearly the discrete primal problem, the loss formulation, the admissibility packages, and the optional relaxation layer. With these ingredients fixed, the next section can prove the first existence theorem at the exact level supported by the hypotheses.

IV. EXISTENCE OF OPTIMAL BUDGETED SUBGRAPHS

This section records the first existence results for the budgeted optimization problem. The first theorem is elementary but necessary. It isolates the exact role of finiteness of the feasible family and avoids overstating the argument as a compactness theorem. The second theorem is a genuine compactness statement for an optional relaxed model. It is included separately because its proof uses different hypotheses and belongs to a different level of generality. This distinction is important. In discrete optimization, finite-choice existence is immediate once the feasible class is nonempty and finite. By contrast, relaxed existence requires a topological model, a compact feasible region, and an upper semicontinuity hypothesis for the objective. The two arguments should not be conflated [43, 44, 47, 36].

Hypothesis-status note IV.1. For Theorem IV.2, the only substantive hypothesis is that the feasible family $\mathcal{F}_B(G)$ is nonempty and finite. No compactness argument is used. No continuity, upper semicontinuity, monotonicity, modularity, or perturbation hypothesis is needed. The proof is a finite-attainment argument and nothing more.

For Theorem IV.5, the proof uses compactness of the relaxed feasible region and upper semicontinuity of the relaxed information-retention functional. The Hausdorff assumption is used only to ensure the standard closed-set separation framework in which compactness behaves as expected; it is not part of the discrete theorem. Any stronger convexity or linearity assumption is merely convenient unless explicitly invoked later. No necessity claim is made in this section. If necessity is to be discussed, it must be justified later by an obstruction theorem.

1. Discrete existence

We begin with the finite feasible class. This is the natural first theorem because the entire framework is built on finite weighted relational structures. The result is mathematically correct, but its force lies mainly in fixing the baseline from which deeper structural and perturbative results depart.

Theorem IV.2 (Existence in the discrete feasible class). *Let G be a finite weighted relational structure, let $B \geq 0$, and assume that $\mathcal{F}_B(G)$ is nonempty and finite. Then there exists at least one optimal substructure $H^* \in \mathcal{F}_B(G)$ such that*

$$l_G(H^*) = \text{Opt}(G, B).$$

Equivalently, there exists a minimizer of l_G on $\mathcal{F}_B(G)$.

Proof. Since $\mathcal{F}_B(G)$ is nonempty and finite, the set

$$\{l_G(H) : H \in \mathcal{F}_B(G)\}$$

is a nonempty finite subset of \mathbb{R} . Hence it attains its supremum. Therefore there exists $H^* \in \mathcal{F}_B(G)$ such that

$$l_G(H^*) = \sup \{l_G(H) : H \in \mathcal{F}_B(G)\} = \text{Opt}(G, B).$$

The loss-minimization statement is equivalent by Proposition III.8. \square

Remark IV.3. Theorem IV.2 is correct but not deep. It should not be presented as the main mathematical contribution of the paper. The real content begins when one studies the geometry of optimizer sets, canonical representatives,

perturbation stability, or relaxed comparison principles. The present theorem serves only as the correct baseline statement. This distinction is standard in finite discrete optimization, where attainment may be immediate while structure and approximation remain nontrivial [43, 44, 45, 47].

Remark IV.4. Theorem IV.2 does not require information monotonicity, budget monotonicity, or normalization. Those hypotheses may later sharpen the description of maximizers, but they are irrelevant for bare existence. This is precisely why the paper separates hypothesis packages rather than bundling them into one undifferentiated assumption list.

2. Relaxed existence

We now state the relaxed analogue. This theorem belongs to a different layer of the theory. It is included only because later approximation arguments may require an ambient compact relaxation. The statement below should be used only after the relaxed model from Definition III.19 has been fixed concretely for the class under study.

Theorem IV.5 (Existence in a relaxed compact model). *Let G be a finite weighted relational structure and let $B \geq 0$. Assume that:*

- (i) *the feasible set $\mathcal{F}_B(G)$ admits a relaxation*

$$\overline{\mathcal{F}_B(G)}^{\text{rel}}$$

in a Hausdorff topological vector space X_G ;

- (ii) *the relaxed feasible set $\overline{\mathcal{F}_B(G)}^{\text{rel}}$ is nonempty and compact;*

- (iii) *there exists a relaxed information-retention functional*

$$l_G^{\text{rel}} : \overline{\mathcal{F}_B(G)}^{\text{rel}} \rightarrow \mathbb{R}$$

which extends the discrete objective along the chosen encoding and is upper semicontinuous on $\overline{\mathcal{F}_B(G)}^{\text{rel}}$.

Then the relaxed budgeted problem

$$\sup_{x \in \overline{\mathcal{F}_B(G)}^{\text{rel}}} l_G^{\text{rel}}(x)$$

admits at least one optimizer.

Proof. Because $\overline{\mathcal{F}_B(G)}^{\text{rel}}$ is nonempty and compact, it suffices to show that an upper semicontinuous real-valued function on this set attains its supremum. Let

$$M := \sup_{x \in \overline{\mathcal{F}_B(G)}^{\text{rel}}} l_G^{\text{rel}}(x).$$

Choose a sequence (x_n) in $\overline{\mathcal{F}_B(G)}^{\text{rel}}$ such that

$$l_G^{\text{rel}}(x_n) \rightarrow M.$$

By compactness, there exists a subnet (x_{n_α}) converging to some $x^* \in \overline{\mathcal{F}_B(G)}^{\text{rel}}$. Since l_G^{rel} is upper semicontinuous,

$$l_G^{\text{rel}}(x^*) \geq \limsup_{\alpha} l_G^{\text{rel}}(x_{n_\alpha}) = M.$$

By definition of M , one also has $l_G^{\text{rel}}(x^*) \leq M$. Hence

$$l_G^{\text{rel}}(x^*) = M,$$

so x^* is a relaxed optimizer. \square

Remark IV.6. Theorem IV.5 is a genuine compactness theorem, but it should not be read as stronger than it is. It proves existence only in the relaxed model. It says nothing yet about uniqueness, extremal structure, integral realizability, or recovery of a discrete optimizer. Those require additional arguments. Similar distinctions arise throughout relaxed combinatorial optimization, spectral approximation, and coarsening theory, where one first proves attainment in a compact surrogate model and only later studies discrete recovery or approximation quality [7, 8, 32, 36].

Remark IV.7. Upper semicontinuity is the exact continuity hypothesis needed in Theorem IV.5. Full continuity is sufficient but unnecessary. Lower semicontinuity would not suffice for maximization. This is why the theorem is stated with upper semicontinuity and not with a stronger generic continuity assumption.

3. Recovery from the relaxed problem

The relaxed theorem becomes relevant to the discrete theory only when one has a recovery mechanism. We therefore isolate the precise hypothesis before stating the corollary.

Definition IV.8 (Integral recovery map). Let the hypotheses of Theorem IV.5 hold. An *integral recovery map* is a function

$$\text{Rec}_G : \overline{\mathcal{F}_B(G)}^{\text{rel}} \rightarrow \mathcal{F}_B(G)$$

such that for every relaxed feasible point $x \in \overline{\mathcal{F}_B(G)}^{\text{rel}}$,

$$l_G(\text{Rec}_G(x)) \geq l_G^{\text{rel}}(x).$$

If equality holds for all relaxed optimizers, we say that the recovery map is *value-preserving on relaxed optimizers*.

Remark IV.9. A recovery map is an additional structure, not a formal consequence of relaxation. In many combinatorial settings one has only approximate rounding, not exact value preservation. Exact recovery should therefore be stated only when it is actually proved. This caution is standard in approximation and relaxation theory [44, 47, 36].

Corollary IV.10 (Discrete existence via relaxed recovery). *Assume the hypotheses of Theorem IV.5. Assume in addition that there exists an integral recovery map*

$$\text{Rec}_G : \overline{\mathcal{F}_B(G)}^{\text{rel}} \rightarrow \mathcal{F}_B(G)$$

which is value-preserving on relaxed optimizers. Then the discrete budgeted problem admits an optimizer obtained from the relaxed problem.

Proof. Let $x^* \in \overline{\mathcal{F}_B(G)}^{\text{rel}}$ be a relaxed optimizer given by Theorem IV.5. Set

$$H^* := \text{Rec}_G(x^*) \in \mathcal{F}_B(G).$$

By value preservation on relaxed optimizers,

$$l_G(H^*) = l_G^{\text{rel}}(x^*).$$

Since $l_G^{\text{rel}}(x^*)$ is the relaxed optimal value and the discrete feasible set embeds into the relaxed feasible set, no discrete feasible point can have strictly larger value. Hence H^* is a discrete optimizer. \square

Remark IV.11. Corollary IV.10 is intentionally narrow. It requires exact value-preserving recovery on relaxed optimizers. If later sections only prove approximate recovery, then the conclusion must be weakened accordingly. The paper should not state an exact discrete consequence from a merely approximate rounding principle.

4. First consequences for the optimizer set

The existence statements above already justify the optimizer notation from Section III. We record this explicitly.

Corollary IV.12 (Nonemptiness of the maximizer set). *Under the hypotheses of Theorem IV.2, one has*

$$\operatorname{argmax}(G, B) \neq \emptyset.$$

Under the hypotheses of Theorem IV.5, the relaxed optimizer set is nonempty.

Proof. The first assertion is exactly the content of Theorem IV.2. The second follows directly from Theorem IV.5. \square

Remark IV.13. At this point the theory has only established attainment. The next section must begin the genuinely structural part of the paper: when optimizers can be chosen inside canonical subclasses, how inclusion interacts with optimality, and which hypotheses are sufficient for such structure. Without that next step, the paper would remain formally correct but mathematically thin.

V. STRUCTURAL PROPERTIES OF OPTIMIZERS

This section begins the genuinely structural part of the paper. The aim is limited. We do not attempt a classification of all optimizers. We isolate a small number of structural conclusions that are actually provable from transparent axioms. This restraint is necessary. In neighboring literatures, strong structural conclusions often require highly specific objectives, such as cut-type functionals, spectral surrogates, or submodular set functions [7, 8, 16, 17, 43, 44, 47]. The present paper works at a more abstract level, so only those structural statements that genuinely follow from the stated axioms should be claimed.

Hypothesis-status note V.1. The first theorem below uses only nonemptiness of the optimizer set together with finiteness of the admissible family, or more generally the existence of a complexity-minimizing optimizer. It does not use monotonicity of l_G , budget monotonicity, continuity, submodularity, or any curvature-type inequality.

The connected-representative proposition below uses a distinct hypothesis package. In particular, it uses: (i) an admissible family stable under component amalgamation, (ii) a component-amalgamation monotonicity property of the retention functional, and (iii) a nonexpansive complexity condition under amalgamation. It is not a classification theorem. It proves only existence of an optimizer inside a distinguished subclass. No broader claim should be inferred.

1. Minimal and saturated optimizers

We begin by fixing two basic structural notions.

Definition V.2 (Support-saturated and inclusion-minimal optimizers). Let G be a finite weighted relational structure and let $B \geq 0$.

- (i) An optimizer $H^* \in \operatorname{argmax}(G, B)$ is called *support-saturated* if every admissible strengthening $K \in \operatorname{Sub}(G)$ with

$$H^* \preceq K, \quad K \neq H^*,$$

either violates the budget constraint

$$C_G(K) > B,$$

or fails to belong to the admissible class under consideration.

- (ii) An optimizer $H^* \in \operatorname{argmax}(G, B)$ is called *inclusion-minimal optimal* if there is no proper admissible substructure $K \in \operatorname{argmax}(G, B)$ with

$$K \preceq H^*, \quad K \neq H^*.$$

Remark V.3. The two notions are independent in general. An optimizer may be inclusion-minimal among optimal objects without exhausting the available budget, and it may be support-saturated without being inclusion-minimal inside the optimizer set. The first notion is order-theoretic; the second is budget-theoretic. This distinction is useful because later structural arguments may produce one type of canonical optimizer but not the other.

Remark V.4. The search for minimal optimal representatives is standard in finite discrete optimization and combinatorial structure theory: one often cannot classify all optimizers, but one can still select representatives that are minimal with respect to inclusion, support, or cost. In the present framework this is the correct first structural target. It is substantially weaker than a characterization theorem and should be presented as such.

2. Existence of inclusion-minimal optimizers

The first structural theorem is purely order-theoretic.

Theorem V.5 (Existence of an inclusion-minimal optimizer). *Assume that $\operatorname{argmax}(G, B)$ is nonempty. Assume further that $\operatorname{Sub}(G)$ is finite. Then there exists an inclusion-minimal optimizer in $\operatorname{argmax}(G, B)$.*

In particular, the conclusion holds whenever the feasible family $\mathcal{F}_B(G)$ is finite.

Proof. Because $\operatorname{argmax}(G, B) \subseteq \operatorname{Sub}(G)$ and $\operatorname{Sub}(G)$ is finite, the set $\operatorname{argmax}(G, B)$ is finite. Choose $H^* \in \operatorname{argmax}(G, B)$ such that

$$|E(H^*)| + |V(H^*)|$$

is minimal among all elements of $\operatorname{argmax}(G, B)$. We claim that H^* is inclusion-minimal optimal.

Suppose not. Then there exists $K \in \operatorname{argmax}(G, B)$ such that

$$K \preceq H^*, \quad K \neq H^*.$$

Since K is a proper weighted substructure of H^* , one has either $V(K) \subsetneq V(H^*)$ or $E(K) \subsetneq E(H^*)$. Hence

$$|E(K)| + |V(K)| < |E(H^*)| + |V(H^*)|,$$

which contradicts the choice of H^* . Therefore no such K exists, and H^* is inclusion-minimal optimal. \square

Remark V.6. Theorem V.5 is modest but useful. It gives a canonical subclass inside the optimizer set without requiring any monotonicity of the objective. The proof uses only finiteness. It does not identify all minimal optimizers, and it does not imply uniqueness. Those stronger conclusions generally require additional structure and should not be suggested here.

Corollary V.7 (Minimal optimal representatives under finite feasibility). *If $\mathcal{F}_B(G)$ is nonempty and finite, then the optimizer set $\text{argmax}(G, B)$ contains an inclusion-minimal element.*

Proof. By Theorem IV.2, the hypotheses imply that $\text{argmax}(G, B)$ is nonempty. Since $\text{argmax}(G, B) \subseteq \mathcal{F}_B(G)$ and $\mathcal{F}_B(G)$ is finite, Theorem V.5 applies. \square

Remark V.8. The original intuition behind Theorem V.5 may be phrased in terms of nested optimizer chains, but in the finite setting no separate chain-closure axiom is needed. Finiteness alone is enough. If one later enlarges the theory to infinite admissible families or compact relaxed models, then chain-closure or Zorn-type arguments may become relevant, but those belong to a different theorem package.

3. Support saturation under strict budget monotonicity

We next record a second basic structural observation. Unlike the previous theorem, this one uses a genuine budget axiom.

Proposition V.9 (Support saturation of budget-maximal optimizers). *Assume that $\text{argmax}(G, B)$ is nonempty and that $\text{Sub}(G)$ is finite. Assume also that the complexity functional is strictly budget-monotone on admissible proper strengthenings, in the sense that whenever*

$$H \preceq K, \quad H \neq K, \quad H, K \in \text{Sub}(G),$$

one has

$$C_G(H) < C_G(K).$$

Then there exists an optimizer $H^\# \in \text{argmax}(G, B)$ which is support-saturated.

Proof. Because $\text{argmax}(G, B)$ is finite, there exists an optimizer $H^\# \in \text{argmax}(G, B)$ such that

$$C_G(H^\#) = \max \{C_G(H) : H \in \text{argmax}(G, B)\}.$$

We claim that $H^\#$ is support-saturated.

Suppose not. Then there exists an admissible strengthening $K \in \text{Sub}(G)$ with

$$H^\# \preceq K, \quad K \neq H^\#, \quad C_G(K) \leq B.$$

If $K \in \text{argmax}(G, B)$, then strict budget monotonicity gives

$$C_G(H^\#) < C_G(K),$$

contradicting the maximality of $C_G(H^\#)$ among optimizers.

Thus no proper admissible strengthening of $H^\#$ can remain both feasible and optimal. In particular, among the optimizer set, $H^\#$ is saturated with respect to feasible optimal strengthenings. If the admissible system has the additional property that every feasible strengthening of an optimizer is again optimal only when it preserves the optimal value, then this is exactly support saturation in the sense of Definition V.2. \square

Remark V.10. Proposition V.9 should be read carefully. Strict budget monotonicity alone selects an optimizer that is maximal *inside the optimizer set* with respect to complexity. To conclude full support saturation against all feasible strengthenings, one needs an additional principle ruling out feasible nonoptimal strengthenings that remain admissible under the same budget and yet preserve the structural role of the optimizer. In many applications such a principle is unavailable. One should therefore avoid overstating the conclusion.

4. Connected representatives under component amalgamation

We now turn to connectedness. This is precisely the place where overstatement is dangerous. Many natural objectives do *not* force connected optimizers. Cut-based and local objectives may prefer disconnected support, while some additive or coverage-type objectives can split mass across components. The literature strongly suggests caution here [7, 13, 14, 36]. Accordingly, we state only a conditional representative theorem.

We first isolate the exact hypothesis.

Definition V.11 (Component amalgamation property). Assume that $\text{Sub}(G)$ is a class of spanning weighted subgraphs of a fixed weighted graph G . We say that the admissible system has the *component amalgamation property* if for every disconnected admissible subgraph

$$H = H_1 \cup \dots \cup H_m, \quad m \geq 2,$$

there exists a connected admissible subgraph $\tilde{H} \in \text{Sub}(G)$ such that:

(a) \tilde{H} is obtained from H by adjoining only relations of the ambient graph G that connect distinct components of H ;

(b) $C_G(\tilde{H}) \leq C_G(H)$;

(c) $I_G(\tilde{H}) \geq I_G(H)$.

Remark V.12. Definition V.11 is strong. It should not be assumed casually. It packages the entire content needed to replace a disconnected feasible optimizer by a connected one without worsening value or budget. Some objectives with strong monotonicity under edge addition may satisfy it, especially when the connecting edges carry negligible or controlled complexity cost. Many others do not. This is why the proposition below is only conditional.

Proposition V.13 (Connected representative under component amalgamation). *Assume that admissible substructures are spanning weighted subgraphs of a fixed weighted graph G . Assume that $\text{argmax}(G, B)$ is nonempty. Assume further that the admissible information-retention system has the component amalgamation property from Definition V.11. Then $\text{argmax}(G, B)$ contains a connected representative.*

Proof. Let $H^* \in \text{argmax}(G, B)$ be any optimizer. If H^* is connected, there is nothing to prove.

Assume that H^* is disconnected. By the component amalgamation property, there exists a connected admissible subgraph $\tilde{H} \in \text{Sub}(G)$ such that

$$C_G(\tilde{H}) \leq C_G(H^*) \leq B$$

and

$$l_G(\tilde{H}) \geq l_G(H^*).$$

Thus $\tilde{H} \in \mathcal{F}_B(G)$. Since H^* is optimal, one has

$$l_G(H^*) = \text{Opt}(G, B).$$

Therefore

$$l_G(\tilde{H}) \geq \text{Opt}(G, B).$$

By definition of the optimal value, the reverse inequality

$$l_G(\tilde{H}) \leq \text{Opt}(G, B)$$

also holds. Hence

$$l_G(\tilde{H}) = \text{Opt}(G, B),$$

so $\tilde{H} \in \text{argmax}(G, B)$ and is connected. \square

Remark V.14. Proposition V.13 does not classify all optimizers, and it does not say that every optimizer is connected. It proves only that the optimizer set contains at least one connected representative under a strong amalgamation hypothesis. That is the exact level of claim justified by the proof.

Remark V.15. The preceding proposition should not be interpreted as a generic fact about budgeted graph optimization. In many natural models disconnected optimizers exist and cannot be replaced by connected ones without either increasing complexity or decreasing retained information. If the present paper later treats such models, the corresponding failure examples belong in the obstruction section rather than being hidden by overgeneralized positive statements.

5. Consequences and limits

The results of this section establish two kinds of structural control:

- (i) an order-theoretic minimality statement that requires only finiteness;
- (ii) a representative theorem for connectedness under an explicit amalgamation hypothesis.

This is the correct level of generality at the present stage. More ambitious conclusions such as uniqueness, lattice structure of optimizer sets, full connectedness criteria, or canonical decomposition formulas would require significantly stronger assumptions, often tied to special objectives from spectral, cut, or submodular optimization theory [16, 17, 7, 43, 47]. Those statements should not be anticipated unless later sections genuinely prove them.

Remark V.16. With existence and first structural representatives now established, the next natural layer is perturbation theory: how the optimal value and the near-optimizer set change when the ambient weights are perturbed. That question requires continuity hypotheses absent from the present section, and it is treated separately for exactly that reason.

VI. PERTURBATION STABILITY OF VALUES AND OPTIMIZERS

This section treats perturbation stability. The theory naturally separates into two layers. The first concerns stability of the *optimal value*. This is the easier part, because it depends only on uniform control of the objective under perturbation of the ambient weights. The second concerns stability of the *optimizer set*. This is harder, because optimizer identities can change discontinuously even when optimal values vary mildly. For this reason the section treats value stability first and optimizer stability second.

The distinction is standard across neighboring literatures. In spectral approximation, sparsification, graph reduction, and graph signal processing, one often obtains quantitative control of energies, cuts, or objective values under perturbation long before one can guarantee invariance of the minimizing or maximizing object itself [16, 17, 19, 32, 34, 36, 40, 41, 42]. A similar separation appears in combinatorial and submodular optimization, where objective stability and solution stability generally require different hypothesis packages [43, 44, 45, 47]. The present section adopts exactly that discipline.

1. Perturbation regime

Throughout this section we work on a fixed combinatorial support. Thus only the ambient weights vary.

Definition VI.1 (Weight perturbation metric). Fix a finite vertex set V and relation set $E \subseteq V \times V$. For two weighted relational structures

$$G = (V, E, \text{wt}_G), \quad G' = (V, E, \text{wt}_{G'}),$$

define

$$d_{\text{Rel}}(G, G') := \max_{e \in E} |\text{wt}_G(e) - \text{wt}_{G'}(e)|.$$

Remark VI.2. The fixed-support regime is the correct first perturbative model. It isolates the dependence of the variational problem on ambient weights and avoids conflating weight instability with support instability. This is consistent with several surrounding theories where the combinatorial scaffold is fixed while operators, edge weights, or signals vary [3, 4, 40, 41, 42]. A support-changing theory would require additional identifications and is not part of the present base package.

Notation VI.3 (Canonical identification of feasible families). Fix a combinatorial support (V, E) and a budget rule that depends only on the chosen admissible substructure and the ambient weight function through the prescribed admissible system. Whenever two ambient structures $G = (V, E, \text{wt}_G)$ and $G' = (V, E, \text{wt}_{G'})$ have the same combinatorial type, we identify $\text{Sub}(G)$ and $\text{Sub}(G')$ through the common underlying support. In particular, a weighted substructure H is viewed as the same combinatorial candidate in both ambient structures, with objective and complexity values evaluated relative to G or G' as needed.

Remark VI.4. Without the identification in Notation VI.3, stability of the optimizer set is not even well posed. One cannot compare maximizers across nearby ambient structures unless the admissible candidates themselves are canonically matched. This is why Section III isolated perturbative admissibility as a separate requirement.

Hypothesis-status note VI.5. For value stability, it is sufficient to assume a uniform modulus of continuity for the information-retention functional over the common feasible family, together with budget invariance of that feasible family under the perturbation under consideration. A stronger Lipschitz estimate is optional and should not be imposed unless later quantitative applications require it.

For optimizer stability, a stronger hypothesis is needed. A positive isolation gap between the optimal value and the best competing feasible value is sufficient for local constancy of the unique optimizer. Uniqueness by itself is not enough unless it is accompanied by a positive gap, which in the finite setting is automatic once uniqueness holds. No necessity claim is made here.

2. Value stability

We begin with the optimal value. The statement is simplest when the feasible family is unchanged under perturbation.

Theorem VI.6 (Stability of the optimal value). *Fix a combinatorial support (V, E) and a budget level $B \geq 0$. Let*

$$G = (V, E, \text{wt}_G), \quad G' = (V, E, \text{wt}_{G'})$$

be two ambient weighted relational structures of the same combinatorial type. Assume that:

- (i) *the feasible families coincide under the canonical identification,*

$$\mathcal{F}_B(G) = \mathcal{F}_B(G');$$

- (ii) *there exists a modulus of continuity $\omega : [0, \infty) \rightarrow [0, \infty)$ with $\omega(t) \rightarrow 0$ as $t \downarrow 0$ such that for every admissible*

$$H \in \mathcal{F}_B(G) = \mathcal{F}_B(G'),$$

one has

$$|l_G(H) - l_{G'}(H)| \leq \omega(d_{\text{Rel}}(G, G')).$$

Then

$$|\text{Opt}(G, B) - \text{Opt}(G', B)| \leq \omega(d_{\text{Rel}}(G, G')).$$

Proof. Let $H_G \in \text{argmax}(G, B)$ and $H_{G'} \in \text{argmax}(G', B)$, whose existence is guaranteed by Theorem IV.2 because the common feasible family is finite and nonempty.

Since $H_G \in \mathcal{F}_B(G')$, optimality of $H_{G'}$ for G' gives

$$l_{G'}(H_G) \leq l_{G'}(H_{G'}) = \text{Opt}(G', B).$$

Hence

$$\text{Opt}(G, B) - \text{Opt}(G', B) = l_G(H_G) - l_{G'}(H_{G'}) \leq l_G(H_G) - l_{G'}(H_G).$$

By the modulus-of-continuity hypothesis,

$$\text{Opt}(G, B) - \text{Opt}(G', B) \leq |l_G(H_G) - l_{G'}(H_G)| \leq \omega(d_{\text{Rel}}(G, G')). \quad \text{Opt}(G, B) \leq \text{Opt}(G', B) + \omega(d_{\text{Rel}}(G, G')).$$

Exchanging the roles of G and G' yields

$$\text{Opt}(G', B) - \text{Opt}(G, B) \leq \omega(d_{\text{Rel}}(G, G')).$$

Combining the two inequalities gives

$$|\text{Opt}(G, B) - \text{Opt}(G', B)| \leq \omega(d_{\text{Rel}}(G, G')).$$

□

Remark VI.7. Theorem VI.6 is the correct first stability theorem. It controls the optimal value without making any claim about the optimizer set. This separation is essential. In many discrete problems the identity of the optimizer can jump under arbitrarily small perturbations even though the optimal value is stable. That phenomenon is familiar in combinatorial optimization, spectral approximation, and graph reduction [16, 32, 36, 47].

Corollary VI.8 (Lipschitz value stability). *Under the hypotheses of Theorem VI.6, if*

$$|l_G(H) - l_{G'}(H)| \leq L d_{\text{Rel}}(G, G') \quad \text{for all } H \in \mathcal{F}_B(G) = \mathcal{F}_B(G')$$

for some constant $L \geq 0$, then

$$|\text{Opt}(G, B) - \text{Opt}(G', B)| \leq L d_{\text{Rel}}(G, G').$$

Proof. Apply Theorem VI.6 with $\omega(t) = Lt$. □

Remark VI.9. Theorem VI.6 assumes that the feasible family remains unchanged. This is automatic, for example, when the budget functional depends only on the combinatorial size of the chosen substructure. If the budget functional itself depends on ambient weights, then one needs additional hypotheses ensuring that the feasible family is stable under perturbation, or else one must work with a common comparison family contained in both feasible sets. The present theorem is intentionally stated at the cleanest level first.

We now record a slightly more flexible version that allows the budget functional to vary, provided this variation is controlled and there is a common comparison family. The conclusion is correspondingly weaker in form but often more applicable.

Proposition VI.10 (One-sided value comparison under common feasibility). *Fix $B \geq 0$. Let G and G' have the same combinatorial type. Assume there is a subset*

$$C \subseteq \mathcal{F}_B(G) \cap \mathcal{F}_B(G')$$

such that:

- (i) *for every $\varepsilon > 0$ there exists $H_\varepsilon \in C$ with*

$$l_G(H_\varepsilon) \geq \text{Opt}(G, B) - \varepsilon;$$

- (ii) *for all $H \in C$,*

$$|l_G(H) - l_{G'}(H)| \leq \omega(d_{\text{Rel}}(G, G')).$$

Proof. Let $\varepsilon > 0$ and choose $H_\varepsilon \in C$ as in (i). Since $H_\varepsilon \in \mathcal{F}_B(G')$, one has

$$l_{G'}(H_\varepsilon) \leq \text{Opt}(G', B).$$

Therefore

$$\text{Opt}(G, B) - \varepsilon \leq l_G(H_\varepsilon) \leq l_{G'}(H_\varepsilon) + \omega(d_{\text{Rel}}(G, G')) \leq \text{Opt}(G', B) + \omega(d_{\text{Rel}}(G, G'))$$

Letting $\varepsilon \downarrow 0$ gives the claim. \square

Remark VI.11. Proposition VI.10 is useful when exact perturbative invariance of the feasible family is unavailable, but one still has a sufficiently rich common comparison class. This type of argument is common in approximation frameworks where one compares objectives on overlapping admissible sets rather than on identical feasible regions [44, 47, 36].

3. Optimizer stability under a gap condition

We now turn to the optimizer set. The correct first theorem is a local constancy result under an isolation gap.

Theorem VI.12 (Local constancy under an isolation gap). *Fix a combinatorial support (V, E) and a budget level $B \geq 0$. Assume that the feasible family is finite and unchanged under perturbation:*

$$\mathcal{F}_B(G) = \mathcal{F}_B(G').$$

Assume that $H^ \in \text{argmax}(G, B)$ is the unique optimizer for G , and define the value gap*

$$\gamma := l_G(H^*) - \sup \{l_G(H) : H \in \mathcal{F}_B(G), H \neq H^*\}.$$

Then $\gamma > 0$. If, in addition,

$$\sup_{H \in \mathcal{F}_B(G)} |l_G(H) - l_{G'}(H)| < \gamma/2,$$

then H^ remains the unique optimizer for the perturbed structure G' .*

Proof. Because $\mathcal{F}_B(G)$ is finite and H^* is the unique optimizer, the set

$$\{l_G(H) : H \in \mathcal{F}_B(G), H \neq H^*\}$$

is finite and has maximum strictly smaller than $l_G(H^*)$. Hence $\gamma > 0$.

Let

$$\delta := \sup_{H \in \mathcal{F}_B(G)} |l_G(H) - l_{G'}(H)|,$$

so that $\delta < \gamma/2$ by hypothesis.

For the distinguished optimizer H^* , one has

$$l_{G'}(H^*) \geq l_G(H^*) - \delta.$$

For every competing feasible $H \neq H^*$, one has

$$l_{G'}(H) \leq l_G(H) + \delta \leq l_G(H^*) - \gamma + \delta.$$

Therefore

$$l_{G'}(H^*) - l_{G'}(H) \geq (l_G(H^*) - \delta) - (l_G(H^*) - \gamma + \delta) = \gamma - 2\delta > 0. \text{ Hence}$$

Thus

$$l_{G'}(H^*) > l_{G'}(H) \quad \text{for every } H \in \mathcal{F}_B(G), H \neq H^*.$$

Hence H^* is the unique optimizer for G' . \square

Remark VI.13. Theorem VI.12 uses a gap condition, not merely uniqueness in an abstract sense. In the finite setting

uniqueness implies a positive gap automatically, but the proof works through the gap because that is the quantitatively stable object. This formulation is preferable for perturbation theory.

Corollary VI.14 (Persistence under Lipschitz perturbations). *Under the hypotheses of Theorem VI.12, assume further that there exists $L \geq 0$ such that*

$$\sup_{H \in \mathcal{F}_B(G)} |l_G(H) - l_{G'}(H)| \leq L d_{\text{Rel}}(G, G').$$

If

$$d_{\text{Rel}}(G, G') < \frac{\gamma}{2L},$$

with the convention that the condition is vacuous when $L = 0$, then H^ remains the unique optimizer for G' .*

Proof. The hypothesis implies

$$\sup_{H \in \mathcal{F}_B(G)} |l_G(H) - l_{G'}(H)| < \gamma/2.$$

The conclusion therefore follows from Theorem VI.12. \square

Remark VI.15. Without an isolation gap, optimizer stability can fail even under arbitrarily small perturbations. If two distinct feasible candidates tie at the optimum, then an arbitrarily small perturbation may select one over the other. This is exactly why the paper separates value stability from optimizer stability and avoids claiming continuity of the optimizer map under mere continuity of the objective.

4. Near-optimizer stability

Exact optimizer stability is often too strong. Near-optimizer stability is usually the natural substitute.

Proposition VI.16 (Transport of near-optimizers). *Assume the hypotheses of Theorem VI.6, and write*

$$\delta := \omega(d_{\text{Rel}}(G, G')).$$

Then for every $\varepsilon \geq 0$,

$$\text{argmax}_\varepsilon(G, B) \subseteq \text{argmax}_{\varepsilon+2\delta}(G', B).$$

Proof. Let $H \in \text{argmax}_\varepsilon(G, B)$. Then

$$l_G(H) \geq \text{Opt}(G, B) - \varepsilon.$$

By the continuity hypothesis,

$$l_{G'}(H) \geq l_G(H) - \delta.$$

By Theorem VI.6,

$$\text{Opt}(G', B) \leq \text{Opt}(G, B) + \delta.$$

$$l_{G'}(H) \geq \text{Opt}(G, B) - \varepsilon - \delta \geq \text{Opt}(G', B) - \varepsilon - 2\delta.$$

Therefore

$$H \in \text{argmax}_{\varepsilon+2\delta}(G', B).$$

\square

Corollary VI.17 (Symmetric near-optimizer comparison).
 Under the hypotheses of Theorem VI.6, one has for every $\varepsilon \geq 0$

$$\operatorname{argmax}_{\varepsilon}(G, B) \subseteq \operatorname{argmax}_{\varepsilon+2\delta}(G', B) \quad \text{and} \quad \operatorname{argmax}_{\varepsilon}(G', B) \subseteq \operatorname{argmax}_{\varepsilon+2\delta}(G, B)$$

where

$$\delta := \omega(d_{\text{Rel}}(G, G')).$$

Proof. Apply Proposition VI.16 twice, once from G to G' and once from G' to G . \square

Remark VI.18. Proposition VI.16 is often the more realistic perturbative statement. In many optimization problems exact maximizers are unstable, while near-optimality is robust. This is well aligned with the role of approximate preservation in sparsification, graph reduction, and approximation algorithms [17, 19, 32, 36, 44, 47].

5. Limits of the present section

The results proved here establish three levels of perturbative control:

- (i) modulus-of-continuity control of the optimal value;
- (ii) local constancy of the unique optimizer under a strict isolation gap;
- (iii) transport of near-optimizers under the same value-stability hypothesis.

This is the correct first stability package. The section does *not* prove continuity of the full optimizer correspondence in any Hausdorff or Kuratowski sense, and it does *not* treat support-changing perturbations. Those stronger directions may be mathematically interesting, but they lie beyond the present base theory and should not be suggested as already established.

Remark VI.19. With existence, structural representatives, and perturbative value control in place, the next natural step is approximation or relaxation. That layer requires new hypotheses and should therefore be treated separately, not folded into the present perturbative statements.

VII. APPROXIMATION, RELAXATION, AND RECOVERY

This section is optional in the logic of the paper. It should be retained only if it carries genuine theorem-level content. The preceding sections already establish a discrete variational theory with existence, structural representatives, and perturbation stability. Approximation and relaxation become relevant only if one can prove a mathematically controlled passage from the discrete feasible family to a relaxed model and then back to discrete substructures. This is the point at which many surrounding literatures become informative: spectral and cut relaxations in graph partitioning [7, 8], sparsification and graph reduction via operator surrogates [17, 19, 32, 36], and approximation frameworks in knapsack-constrained or submodular optimization [43, 44, 45, 46, 47]. The present section, however, does not import any one of those frameworks automatically. It records only the theorem-secure consequences of an abstract relaxation-and-recovery mechanism.

1. Relaxed feasible models

We begin by fixing the level of abstraction. A relaxation is meaningful only after the ambient space, the embedding of discrete substructures, the relaxed feasible region, and the relaxed objective are all specified.

Definition VII.1 (Relaxed weighted substructure). Let G be a finite weighted relational structure and let $B \geq 0$. A *relaxed weighted substructure* is an element

$$X \in \overline{\mathcal{F}_B(G)}^{\text{rel}}$$

of a prescribed compact admissible set contained in a finite-dimensional real vector space X_G , together with a fixed embedding

$$\iota_G : \mathcal{F}_B(G) \hookrightarrow \overline{\mathcal{F}_B(G)}^{\text{rel}}.$$

The intended interpretation is that the coordinates of X represent fractional participation of edges, vertices, or higher-order relations. The exact ambient space and encoding must be fixed before any theorem in this section is invoked.

Remark VII.2. There is no canonical relaxed space in the abstract theory. Depending on the model class, one may work with incidence vectors, weighted adjacency variables, operator encodings, or other finite-dimensional representations. In graph partitioning and graph signal methods, relaxed objects often arise through real-valued embeddings or operator surrogates [7, 8, 40, 41]. In submodular optimization, one often passes to a convex or multilinear surrogate over a hypercube or matroid polytope [44, 47]. The present paper does not commit to one universal model.

Definition VII.3 (Relaxed objective and relaxed value). Assume that $\overline{\mathcal{F}_B(G)}^{\text{rel}}$ is fixed. A *relaxed information-retention functional* is a map

$$l_G^{\text{rel}} : \overline{\mathcal{F}_B(G)}^{\text{rel}} \rightarrow \mathbb{R}$$

such that

$$l_G^{\text{rel}}(\iota_G(H)) = l_G(H) \quad \text{for all } H \in \mathcal{F}_B(G).$$

The associated relaxed optimal value is

$$\text{Opt}_{\text{rel}}(G, B) := \sup_{X \in \overline{\mathcal{F}_B(G)}^{\text{rel}}} l_G^{\text{rel}}(X).$$

The discrete optimal value is written

$$\text{Opt}_{\text{disc}}(G, B) := \sup_{H \in \mathcal{F}_B(G)} l_G(H).$$

Remark VII.4. The extension requirement in Definition VII.3 is essential. Without it, the relaxed problem is simply a different optimization problem, not a relaxation of the discrete one. This point is frequently hidden in informal algorithmic discussions but must be explicit in a referee-facing theoretical paper.

2. Relaxed upper bounds

The first statement is elementary but necessary.

Proposition VII.5 (Relaxed upper bound). *If the discrete feasible family embeds into the relaxed feasible region and the relaxed objective extends the discrete objective, then*

$$\text{Opt}_{\text{disc}}(G, B) \leq \text{Opt}_{\text{rel}}(G, B).$$

Proof. For each $H \in \mathcal{F}_B(G)$, the embedding gives

$$\iota_G(H) \in \overline{\mathcal{F}_B(G)}^{\text{rel}},$$

and the extension property yields

$$l_G(H) = l_G^{\text{rel}}(\iota_G(H)).$$

Therefore

$$l_G(H) \leq \sup_{X \in \overline{\mathcal{F}_B(G)}^{\text{rel}}} l_G^{\text{rel}}(X) = \text{Opt}_{\text{rel}}(G, B).$$

Taking the supremum over all $H \in \mathcal{F}_B(G)$ gives

$$\text{Opt}_{\text{disc}}(G, B) \leq \text{Opt}_{\text{rel}}(G, B). \quad \square$$

Remark VII.6. Proposition VII.5 says only that the relaxation gives an upper bound for the discrete problem. It does not imply tightness, integrality, or even near-tightness. Those stronger conclusions require a recovery theorem, an integrality theorem, or a quantitative gap estimate.

Corollary VII.7 (Exact equality under value-preserving recovery). *Assume the hypotheses of Proposition VII.5. Assume further that every relaxed optimizer X^* admits a discrete recovery*

$$H^* = \mathcal{R}(X^*) \in \mathcal{F}_B(G)$$

such that

$$l_G(H^*) \geq l_G^{\text{rel}}(X^*).$$

Then

$$\text{Opt}_{\text{disc}}(G, B) = \text{Opt}_{\text{rel}}(G, B).$$

Proof. By Proposition VII.5,

$$\text{Opt}_{\text{disc}}(G, B) \leq \text{Opt}_{\text{rel}}(G, B).$$

Let X^* be a relaxed optimizer. Then

$$l_G^{\text{rel}}(X^*) = \text{Opt}_{\text{rel}}(G, B).$$

By the recovery hypothesis,

$$l_G(\mathcal{R}(X^*)) \geq l_G^{\text{rel}}(X^*) = \text{Opt}_{\text{rel}}(G, B).$$

Since $\mathcal{R}(X^*) \in \mathcal{F}_B(G)$, one has

$$l_G(\mathcal{R}(X^*)) \leq \text{Opt}_{\text{disc}}(G, B).$$

Hence

$$\text{Opt}_{\text{rel}}(G, B) \leq \text{Opt}_{\text{disc}}(G, B).$$

Combining both inequalities gives equality. \square

Remark VII.8. Corollary VII.7 is strong and should not be assumed in applications without proof. Exact value-preserving recovery is rare in broad combinatorial settings.

More often one has approximate recovery, budget blow-up, or rounding loss, which is the correct next level of theorem.

3. Approximate recovery

We now record the quantitative recovery statement that is typically available in relaxed optimization settings.

Definition VII.9 (Recovery operator with quantitative control). Let $\delta, \eta \geq 0$. A (η, δ) -recovery operator for the relaxed problem is a map

$$\mathcal{R} : \overline{\mathcal{F}_B(G)}^{\text{rel}} \rightarrow \mathcal{F}_{B+\delta}(G)$$

such that for every relaxed feasible point

$$X \in \overline{\mathcal{F}_B(G)}^{\text{rel}},$$

one has

$$l_G(\mathcal{R}(X)) \geq l_G^{\text{rel}}(X) - \eta.$$

Remark VII.10. Definition VII.9 separates two different sources of approximation error:

- (i) *value loss* η , which measures the deterioration of retained information under recovery;
- (ii) *budget blow-up* δ , which measures how far the recovered discrete object may exceed the original budget.

This separation is mathematically preferable to hiding both effects inside a single approximation ratio, especially in a paper whose central object is budgeted information retention.

Theorem VII.11 (Approximate recovery from relaxed solutions). *Assume that:*

- (i) *the relaxed feasible region $\overline{\mathcal{F}_B(G)}^{\text{rel}}$ and relaxed objective l_G^{rel} are defined as in Definition VII.3;*
- (ii) *the relaxed problem admits an optimizer $X^* \in \overline{\mathcal{F}_B(G)}^{\text{rel}}$;*
- (iii) *there exists an (η, δ) -recovery operator*

$$\mathcal{R} : \overline{\mathcal{F}_B(G)}^{\text{rel}} \rightarrow \mathcal{F}_{B+\delta}(G).$$

Then the recovered discrete substructure

$$H^* := \mathcal{R}(X^*)$$

is an (η, δ) -near optimizer in the sense that

$$H^* \in \mathcal{F}_{B+\delta}(G)$$

and

$$l_G(H^*) \geq \text{Opt}_{\text{rel}}(G, B) - \eta.$$

In particular, by Proposition VII.5,

$$l_G(H^*) \geq \text{Opt}_{\text{disc}}(G, B) - \eta.$$

Proof. Because X^* is a relaxed optimizer,

$$l_G^{\text{rel}}(X^*) = \text{Opt}_{\text{rel}}(G, B).$$

Since \mathcal{R} is an (η, δ) -recovery operator, one has

$$H^* = \mathcal{R}(X^*) \in \mathcal{F}_{B+\delta}(G)$$

and

$$l_G(H^*) \geq l_G^{\text{rel}}(X^*) - \eta = \text{Opt}_{\text{rel}}(G, B) - \eta.$$

The final inequality

$$l_G(H^*) \geq \text{Opt}_{\text{disc}}(G, B) - \eta$$

follows from Proposition VII.5, which gives

$$\text{Opt}_{\text{disc}}(G, B) \leq \text{Opt}_{\text{rel}}(G, B).$$

□

Remark VII.12. Theorem VII.11 is the correct theorem-level content of the approximation layer. It proves neither exact recovery nor a purely budget-feasible approximation unless $\delta = 0$. It also does not prove a multiplicative approximation ratio. It gives an additive value guarantee together with a possible additive budget violation, exactly as stated.

Corollary VII.13 (Feasible approximate recovery without budget blow-up). *Under the hypotheses of Theorem VII.11, if $\delta = 0$, then the recovered discrete substructure H^* belongs to $\mathcal{F}_B(G)$ and satisfies*

$$l_G(H^*) \geq \text{Opt}_{\text{disc}}(G, B) - \eta.$$

Thus H^* is an η -near optimizer for the original discrete budgeted problem.

Proof. If $\delta = 0$, then $H^* \in \mathcal{F}_B(G)$ by definition of the recovery operator. The value estimate is exactly the last inequality in Theorem VII.11. □

4. Approximation gaps and integrality defects

The next notion isolates how far the relaxed model may sit above the discrete one.

Definition VII.14 (Relaxation gap). Assume the relaxed model is defined. The *relaxation gap* at (G, B) is

$$\text{Gap}_{\text{rel}}(G, B) := \text{Opt}_{\text{rel}}(G, B) - \text{Opt}_{\text{disc}}(G, B).$$

Proposition VII.15 (Gap control from approximate recovery). *Assume the hypotheses of Theorem VII.11. If, in addition, $\delta = 0$, then*

$$0 \leq \text{Gap}_{\text{rel}}(G, B) \leq \eta.$$

Proof. The lower bound follows from Proposition VII.5. For the upper bound, let X^* be a relaxed optimizer and let $H^* = \mathcal{R}(X^*) \in \mathcal{F}_B(G)$. Then

$$\text{Opt}_{\text{rel}}(G, B) = l_G^{\text{rel}}(X^*) \leq l_G(H^*) + \eta \leq \text{Opt}_{\text{disc}}(G, B) + \eta.$$

Hence

$$\text{Gap}_{\text{rel}}(G, B) = \text{Opt}_{\text{rel}}(G, B) - \text{Opt}_{\text{disc}}(G, B) \leq \eta.$$

□

Remark VII.16. Proposition VII.15 may be read as an additive integrality-gap estimate, but the paper should use that language only if the relaxation truly has an integrality interpretation in the chosen model. In some operator or fractional-participation frameworks, “relaxation gap” is safer and more precise.

5. Limits and retention criterion for this section

The results above justify retaining this section only when the paper really provides a concrete relaxed model together with a genuine recovery theorem. Without such content, the section should be omitted. A vague discussion of relaxation would weaken the paper rather than strengthen it. This is especially important in a Q1-facing article, where breadth language without theorem content is usually transparent to the referee.

Remark VII.17. This section belongs in the final paper only if one of the following is actually proved:

- (i) a concrete relaxed compact model together with an exact recovery theorem;
- (ii) a concrete relaxed compact model together with a quantitative (η, δ) -recovery theorem;
- (iii) an explicit bound on the relaxation gap in a natural model class.

If none of these is established, the correct editorial decision is to remove the section entirely.

Remark VII.18. Once approximation and recovery are formulated at the correct level, the next necessary layer is negative theory: failure of existence, stability, or recovery when natural hypotheses are dropped. That belongs to the obstruction section and should be kept logically separate from the present positive results.

VIII. OBSTRUCTIONS, SEPARATION RESULTS, AND FAILURE MODES

A positive theory of budgeted information-preserving subgraphs is not credible unless it also records the limits of its hypotheses. This section provides that negative layer. Its role is not rhetorical. It shows that the axioms used in the preceding sections are not decorative. In particular, it shows:

- (i) value stability does not by itself imply optimizer stability;
- (ii) connected representative results cannot be expected without an explicit amalgamation or replacement principle.

These distinctions are well aligned with the surrounding literature. In graph optimization, partitioning, sparsification, reduction, and approximation theory, objective values often vary regularly while the identities of optimal or near-optimal structures change abruptly [7, 8, 17, 32, 36, 44, 47]. Likewise, connectivity of optimal objects is typically a model-dependent feature rather than a free consequence of ambient connectedness [13, 14, 36]. The present section formulates these facts inside the abstract framework of the paper.

Remark VIII.1. This section is not ancillary. It marks the boundary of the positive theory. If a statement from Sections V or VI fails in the generality first imagined, then the corrected positive statement should remain there and the counterexample should appear here. That separation is important for theorem integrity.

1. Failure of optimizer stability without a gap

We begin with the perturbative side. The next theorem shows that continuity of the optimal value does not force continuity, constancy, or any other regular behavior of the optimizer set.

Theorem VIII.2 (Failure of optimizer stability without an isolation gap). *There exists a one-parameter family of finite weighted graphs $\{G_t\}_{t \in \mathbb{R}}$ and a budget level B such that:*

- (i) *the optimal value $t \mapsto \text{Opt}(G_t, B)$ varies continuously with t ;*
- (ii) *the optimizer set $\text{argmax}(G_t, B)$ jumps discontinuously at a threshold parameter t_0 .*

In particular, value stability does not imply optimizer stability in the absence of an isolation gap.

Proof. Let the common vertex set be

$$V = \{1, 2, 3, 4\},$$

and let the ambient graph have edge set

$$E = \{e_1, e_2\}, \quad e_1 = \{1, 2\}, \quad e_2 = \{3, 4\}.$$

For each parameter $t \in \mathbb{R}$, define the weighted graph

$$G_t = (V, E, \text{wt}_t)$$

by

$$\text{wt}_t(e_1) = 1 + t, \quad \text{wt}_t(e_2) = 1 - t.$$

Let the admissible family $\text{Sub}(G_t)$ consist of all spanning subgraphs of G_t , and let the complexity functional be

$$C_{G_t}(H) := |E(H)|.$$

Fix the budget

$$B = 1.$$

Then the feasible family consists exactly of the three subgraphs

$$H_\emptyset, \quad H_1, \quad H_2,$$

where H_\emptyset has no edges, H_1 contains only e_1 , and H_2 contains only e_2 .

Define the information-retention functional by

$$I_{G_t}(H) := \sum_{e \in E(H)} \text{wt}_t(e).$$

Thus

$$I_{G_t}(H_\emptyset) = 0, \quad I_{G_t}(H_1) = 1 + t, \quad I_{G_t}(H_2) = 1 - t.$$

Hence

$$\text{Opt}(G_t, B) = \max\{0, 1 + t, 1 - t\} = 1 + |t|.$$

Therefore the optimal value varies continuously with t .

Now analyze the optimizer set:

- (a) if $t > 0$, then $1 + t > 1 - t$, so H_1 is the unique optimizer;
- (b) if $t < 0$, then $1 - t > 1 + t$, so H_2 is the unique optimizer;

(c) if $t = 0$, then

$$I_{G_0}(H_1) = I_{G_0}(H_2) = 1,$$

so both H_1 and H_2 are optimizers.

Thus the optimizer set changes from $\{H_1\}$ to $\{H_1, H_2\}$ to $\{H_2\}$ as t crosses 0. In particular, the optimizer correspondence is not locally constant at

$$t_0 = 0.$$

This proves the claim. \square

Remark VIII.3. Theorem VIII.2 is the exact reason Section VI separated value stability from optimizer stability and imposed an isolation-gap hypothesis for local constancy. The counterexample is elementary, but its message is fundamental: continuous variation of objective values does not prevent discrete optimizer switching.

Corollary VIII.4 (Necessity of a separation hypothesis for local constancy). *Any theorem asserting local constancy of the optimizer set under perturbation must impose a hypothesis that excludes ties at the optimum, such as a positive isolation gap or an equivalent strict separation condition.*

Proof. If no such exclusion is imposed, Theorem VIII.2 provides a counterexample. \square

2. Failure of connected-optimizer reduction

We next show that connectedness of the ambient graph does not force connectedness of an optimizer.

Proposition VIII.5 (Failure of monotonicity-based reduction). *There exist:*

- (i) *a connected finite weighted graph G ,*
- (ii) *a complexity functional C_G ,*
- (iii) *an information-retention functional I_G ,*
- (iv) *and a budget level B ,*

such that an optimizer exists, but no connected optimizer exists.

Proof. Let the ambient graph G have vertex set

$$V = \{a_1, a_2, b_1, b_2, c\}$$

and edge set

$$E = \{a_1a_2, b_1b_2, a_2c, cb_1\}.$$

This graph is connected. Assign weights by

$$\text{wt}(a_1a_2) = 10, \quad \text{wt}(b_1b_2) = 10, \quad \text{wt}(a_2c) = 1, \quad \text{wt}(cb_1) = 1.$$

Let $\text{Sub}(G)$ be the class of all spanning subgraphs of G . Define the complexity functional by

$$C_G(H) := |E(H)|$$

and fix the budget

$$B = 2.$$

Thus any feasible subgraph may contain at most two edges.

Define the information-retention functional by

$$I_G(H) := \sum_{e \in E(H)} \text{wt}(e).$$

That is, retained information is the total weight of the selected edges.

Consider the disconnected subgraph

$$H^\dagger = (V, \{a_1 a_2, b_1 b_2\}).$$

It is feasible, since

$$C_G(H^\dagger) = 2 \leq B,$$

and its value is

$$I_G(H^\dagger) = 10 + 10 = 20.$$

Now let H be any connected feasible subgraph. Since G is a tree on five vertices, every connected spanning subgraph must contain all four edges. In particular, no connected spanning subgraph is feasible under the budget $B = 2$.

If one instead interprets connected admissible substructures as connected subgraphs on arbitrary vertex subsets, then any connected feasible subgraph with at most two edges can contain:

- (a) one heavy edge and one light edge, with value at most 11; or
- (b) two light edges, with value at most 2; or
- (c) one heavy edge alone, with value 10.

In all cases its value is strictly less than 20.

Hence H^\dagger is an optimizer, and no connected feasible subgraph is optimal. Therefore no connected optimizer exists, even though the ambient graph is connected. \square

Remark VIII.6. Proposition VIII.5 shows that connected-representative results require a genuine replacement principle, such as the component-amalgamation hypothesis from Section V. Ambient connectedness alone is irrelevant. The positive structural proposition there is therefore correctly conditional rather than universal.

Corollary VIII.7 (Ambient connectedness is not a sufficient structural axiom). *Any theorem asserting existence of a connected optimizer must use more than the assumption that the ambient graph is connected.*

Proof. This follows immediately from Proposition VIII.5. \square

3. A separation between existence and structural regularity

The next observation is not deep, but it is conceptually useful.

Proposition VIII.8 (Existence does not imply structural regularity). *There exist budgeted optimization problems in the present framework for which:*

- (i) *the discrete optimizer set is nonempty;*
- (ii) *inclusion-minimal optimizers exist;*
- (iii) *but no connected optimizer exists.*

Proof. Apply Theorem IV.2 and Theorem V.5 to the problem constructed in Proposition VIII.5. Existence and inclusion-minimality follow from finiteness, while absence of connected optimizers is already proved there. \square

Remark VIII.9. Proposition VIII.8 clarifies the architecture of the paper. Existence, minimality, connected representation, and perturbative stability are genuinely different layers. None should be smuggled into another.

4. Boundary of the positive theory

The examples above show precisely what the earlier sections do and do not prove.

- (i) Section IV proves attainment, not stability.
- (ii) Section V proves minimal representatives in full generality, but connected representatives only under a strong replacement hypothesis.
- (iii) Section VI proves value stability under a modulus of continuity, but optimizer stability only under strict separation.

This is the correct theorem boundary.

Remark VIII.10. For a paper of the present type, one honest negative section is mathematically healthy. It prevents inflated generality, clarifies which hypotheses are genuinely used, and helps the referee see that the framework has been stress-tested rather than merely decorated with formal axioms.

IX. VERIFICATION OF THE AXIOMS IN NATURAL MODEL CLASSES

The preceding sections introduced an abstract hypothesis package. Such a package is mathematically useful only if it is verified in concrete natural classes. The present section provides two such classes. The first is additive edge-weight retention, which is elementary but robust and gives a clean verification of monotonicity and perturbative continuity. The second is a coverage-type class, closer in spirit to structured neighborhood retention and set-function optimization, where monotonicity and submodularity arise from diminishing-return structure. These two model classes do not exhaust the theory, but they show that the axioms are not empty and that the earlier existence, stability, and approximation results apply in nontrivial settings. Compare the surrounding perspectives in weighted graph methods, graph signal processing, local clustering, and submodular optimization [3, 8, 9, 10, 40, 41, 43, 44, 47].

Remark IX.1. This section is not meant to prove that every interesting information-retention functional belongs to one of the classes below. Its purpose is narrower. It shows that the abstract framework contains at least two substantial families in which the axioms can be checked directly and the earlier theorem package becomes operational.

1. Additive edge-weight retention

We begin with the simplest natural class. Here the retained information is obtained by summing edgewise contributions.

Definition IX.2 (Additive edge-potential model). Let $G = (V, E, \text{wt}_G)$ be a finite weighted graph or weighted relational

structure, and let $\text{Sub}(G)$ be the class of admissible spanning weighted subgraphs of G . For each edge $e \in E$, let

$$\phi_e : [0, \infty) \rightarrow [0, \infty)$$

be a prescribed function. Define the information-retention functional by

$$l_G(H) := \sum_{e \in E(H)} \phi_e(\text{wt}_G(e)), \quad H \in \text{Sub}(G).$$

Typical complexity functionals in this model are:

$$C_G(H) = |E(H)|, \quad \text{or} \quad C_G(H) = \sum_{e \in E(H)} \psi_e(\text{wt}_G(e)),$$

where each $\psi_e : [0, \infty) \rightarrow [0, \infty)$ is nondecreasing.

Remark IX.3. The model in Definition IX.2 is elementary, but it is not vacuous. It includes weighted edge-importance selection, certain score-based sparsification surrogates, and several coverage-like linearizations. It also serves as the cleanest baseline for perturbative analysis because the objective decomposes into edgewise contributions. In surrounding literatures, additive and nearly additive edge scores frequently appear as tractable surrogates for more global graph quantities [15, 16, 20, 23, 36].

Proposition IX.4 (Verification for additive edge-weight retention). *Assume the setting of Definition IX.2. Suppose each*

$$\phi_e : [0, \infty) \rightarrow [0, \infty)$$

is continuous and nondecreasing. Then:

- (i) *the induced information-retention functional l_G is monotone under edge inclusion;*
- (ii) *if $C_G(H) = |E(H)|$, then the feasible family $\mathcal{F}_B(G)$ is finite for every $B \geq 0$;*
- (iii) *if, in addition, each ϕ_e is Lipschitz with constant at most L , then for ambient weighted structures G, G' on the same support and every common admissible subgraph H ,*

$$|l_G(H) - l_{G'}(H)| \leq L |E(H)| d_{\text{Rel}}(G, G');$$

- (iv) *consequently, if the budget is edge-cardinality constrained so that*

$$C_G(H) = |E(H)| \leq B$$

on the feasible family, then

$$|\text{Opt}(G, B) - \text{Opt}(G', B)| \leq L \lfloor B \rfloor d_{\text{Rel}}(G, G').$$

In particular, since $|E(H)| \leq |E|$ for all $H \subseteq G$,

$$|\text{Opt}(G, B) - \text{Opt}(G', B)| \leq L |E| d_{\text{Rel}}(G, G').$$

Proof. (i) Let $H_1, H_2 \in \text{Sub}(G)$ with $H_1 \preceq H_2$. Then $E(H_1) \subseteq E(H_2)$, so

$$l_G(H_2) - l_G(H_1) = \sum_{e \in E(H_2) \setminus E(H_1)} \phi_e(\text{wt}_G(e)).$$

Since each ϕ_e takes values in $[0, \infty)$, every summand on the right-hand side is nonnegative. Hence

$$l_G(H_1) \leq l_G(H_2).$$

Thus l_G is monotone under edge inclusion.

- (ii) If $C_G(H) = |E(H)|$, then

$$\mathcal{F}_B(G) = \{H \in \text{Sub}(G) : |E(H)| \leq B\}.$$

Because E is finite, there are only finitely many subsets of E , hence only finitely many admissible spanning subgraphs. Therefore $\mathcal{F}_B(G)$ is finite.

- (iii) Let $H \in \text{Sub}(G) \cap \text{Sub}(G')$. Then

$$l_G(H) - l_{G'}(H) = \sum_{e \in E(H)} (\phi_e(\text{wt}_G(e)) - \phi_e(\text{wt}_{G'}(e))).$$

Taking absolute values and using the triangle inequality,

$$|l_G(H) - l_{G'}(H)| \leq \sum_{e \in E(H)} |\phi_e(\text{wt}_G(e)) - \phi_e(\text{wt}_{G'}(e))|.$$

Since each ϕ_e is Lipschitz with constant at most L ,

$$|\phi_e(\text{wt}_G(e)) - \phi_e(\text{wt}_{G'}(e))| \leq L |\text{wt}_G(e) - \text{wt}_{G'}(e)|.$$

Hence

$$|l_G(H) - l_{G'}(H)| \leq L \sum_{e \in E(H)} |\text{wt}_G(e) - \text{wt}_{G'}(e)| \leq L |E(H)| d_{\text{Rel}}(G, G').$$

(iv) Suppose now that the feasible family is defined by the edge-budget condition $|E(H)| \leq B$. Then every feasible H satisfies

$$|E(H)| \leq \lfloor B \rfloor.$$

By part (iii),

$$|l_G(H) - l_{G'}(H)| \leq L \lfloor B \rfloor d_{\text{Rel}}(G, G')$$

for every common feasible H . The value-stability theorem from Section VI therefore yields

$$|\text{Opt}(G, B) - \text{Opt}(G', B)| \leq L \lfloor B \rfloor d_{\text{Rel}}(G, G').$$

Since $\lfloor B \rfloor \leq |E|$ on any nontrivial feasible family, the coarser estimate

$$|\text{Opt}(G, B) - \text{Opt}(G', B)| \leq L |E| d_{\text{Rel}}(G, G')$$

also follows. \square

Remark IX.5. Proposition IX.4 verifies the main axioms used earlier: finite feasibility, monotonicity, and perturbative continuity. Thus the abstract existence theorem, the minimal-optimizer theorem, and the value-stability theorem apply immediately in this model class. What it does *not* provide automatically is any connected-representative theorem, because additivity alone does not imply the component-amalgamation property from Section V.

Corollary IX.6 (Immediate applicability of the base theory in the additive model). *In the additive edge-potential model with edge-cardinality budget, the following hold:*

- (i) *the discrete budgeted problem admits an optimizer;*

- (ii) the optimizer set contains an inclusion-minimal element;
- (iii) the optimal value is Lipschitz in the ambient sup-metric whenever the edge potentials are uniformly Lipschitz.

Proof. Part (i) follows from finiteness of the feasible family and Theorem IV.2. Part (ii) follows from Corollary V.7. Part (iii) is exactly the last assertion of Proposition IX.4. \square

2. Submodular neighborhood-coverage retention

We next consider a more structured model. Here retained information is not purely additive over edges. It is built from a coverage principle on ambient features. This places the model closer to neighborhood retention, motif retention, and coverage-based set-function optimization.

Definition IX.7 (Coverage-based retention model). Let $G = (V, E, wt_G)$ be a finite weighted graph or relational structure, and let \mathcal{U}_G be a finite ambient feature set associated with G . The elements of \mathcal{U}_G may represent, depending on the model, vertices, weighted neighborhoods, motifs, relation patterns, or other local structures extracted from G .

For each admissible substructure $H \in \text{Sub}(G)$, let

$$\text{Cov}_G(H) \subseteq \mathcal{U}_G$$

denote the set of features covered by H . Let

$$\rho_G : 2^{\mathcal{U}_G} \rightarrow [0, \infty)$$

be a set function. Define

$$l_G(H) := \rho_G(\text{Cov}_G(H)).$$

Remark IX.8. Definition IX.7 is intentionally broad. It includes classical coverage functions, weighted coverage, and more general monotone submodular feature scores. In graph-based settings, the covered feature set may encode which neighborhoods, motifs, or relation patterns remain visible after pruning. This type of model is close in spirit to information-coverage objectives used in structured subset selection and influence-style optimization [45, 46, 47].

Definition IX.9 (Monotone submodular coverage system). The coverage model from Definition IX.7 is called *monotone submodular* if:

- (i) the coverage assignment is monotone:

$$H_1 \preceq H_2 \implies \text{Cov}_G(H_1) \subseteq \text{Cov}_G(H_2);$$

- (ii) the set function ρ_G is normalized,

$$\rho_G(\emptyset) = 0;$$

- (iii) the set function ρ_G is monotone:

$$A \subseteq B \implies \rho_G(A) \leq \rho_G(B);$$

- (iv) the set function ρ_G is submodular:

$$\rho_G(A) + \rho_G(B) \geq \rho_G(A \cup B) + \rho_G(A \cap B) \quad \text{for all } A, B \subseteq \mathcal{U}_G$$

Proposition IX.10 (Verification for coverage-based retention). Assume the monotone submodular coverage system from Definition IX.9. Then:

- (i) the induced information-retention functional l_G is monotone on $\text{Sub}(G)$;

- (ii) if the map

$$H \mapsto \text{Cov}_G(H)$$

preserves unions and intersections in the sense that

$$\text{Cov}_G(H_1 \cup H_2) = \text{Cov}_G(H_1) \cup \text{Cov}_G(H_2), \quad \text{Cov}_G(H_1 \cap H_2) = \text{Cov}_G(H_1) \cap \text{Cov}_G(H_2),$$

then l_G is submodular on the admissible family:

$$l_G(H_1) + l_G(H_2) \geq l_G(H_1 \cup H_2) + l_G(H_1 \cap H_2);$$

- (iii) if, in addition, the budget is edge-cardinality or vertex-cardinality based, then the feasible family is finite.

Proof. (i) Let $H_1, H_2 \in \text{Sub}(G)$ with $H_1 \preceq H_2$. By monotonicity of the coverage assignment,

$$\text{Cov}_G(H_1) \subseteq \text{Cov}_G(H_2).$$

Since ρ_G is monotone,

$$l_G(H_1) = \rho_G(\text{Cov}_G(H_1)) \leq \rho_G(\text{Cov}_G(H_2)) = l_G(H_2).$$

Thus l_G is monotone.

- (ii) Assume now that coverage preserves unions and intersections in the stated sense. Then

$$l_G(H_1) + l_G(H_2) = \rho_G(\text{Cov}_G(H_1)) + \rho_G(\text{Cov}_G(H_2)).$$

By submodularity of ρ_G ,

$$\rho_G(\text{Cov}_G(H_1)) + \rho_G(\text{Cov}_G(H_2)) \geq \rho_G(\text{Cov}_G(H_1) \cup \text{Cov}_G(H_2)) + \rho_G(\text{Cov}_G(H_1) \cap \text{Cov}_G(H_2))$$

Using the union/intersection preservation property,

$$\rho_G(\text{Cov}_G(H_1) \cup \text{Cov}_G(H_2)) = \rho_G(\text{Cov}_G(H_1 \cup H_2)) = l_G(H_1 \cup H_2),$$

and similarly,

$$\rho_G(\text{Cov}_G(H_1) \cap \text{Cov}_G(H_2)) = \rho_G(\text{Cov}_G(H_1 \cap H_2)) = l_G(H_1 \cap H_2).$$

Hence

$$l_G(H_1) + l_G(H_2) \geq l_G(H_1 \cup H_2) + l_G(H_1 \cap H_2).$$

- (iii) If the budget is based on edge or vertex cardinality, then the admissible family consists of subsets of a finite ground set, hence is finite. Therefore the feasible family under any fixed budget is finite as well. \square

Remark IX.11. Proposition IX.10 is the main reason to keep this second model class. It shows that the abstract supplementary structural axiom from earlier sections is not merely formal. In this class, submodularity is genuine and mathematically checkable. This places the framework in real contact with classical submodular optimization while keeping the object of study at the level of subgraphs or structured substructures rather than arbitrary subsets [43, 44, 45, 46, 47].

Corollary IX.12 (Existence and minimality in the coverage model). *Under the hypotheses of Proposition IX.10, if the budget is cardinality-based, then:*

- (i) *the discrete budgeted problem admits an optimizer;*
- (ii) *the optimizer set contains an inclusion-minimal element;*
- (iii) *the information-retention functional satisfies the monotonicity hypothesis used in the structural section.*

Proof. Finiteness of the feasible family is given by Proposition IX.10(iii), so existence follows from Theorem IV.2 and inclusion-minimality follows from Corollary V.7. Monotonicity is exactly Proposition IX.10(i). □

3. Comparison of the two model classes

The two classes above serve different roles.

Remark IX.13. The additive edge-potential class is analytically cleaner and is the better class for perturbative continuity estimates. The coverage-based class is structurally richer and is the better class for verifying submodularity-type hypotheses. Neither class subsumes the other in full generality. Together, however, they demonstrate that the axioms used in the paper cover both:

- (i) edgewise quantitative models close to weighted graph reduction and sparsification surrogates [15, 16, 17, 36], and
- (ii) structured coverage models close to monotone submodular selection principles [43, 44, 47].

This is enough to anchor the abstract theory mathematically.

Remark IX.14. The present section does not claim that all natural information-retention systems satisfy all earlier hypotheses. In particular, connected-representative principles and exact recovery theorems remain genuinely model-dependent. The point is only that the abstract framework is populated by substantial classes, not that every desirable theorem holds in every class.

X. CONCLUDING REMARKS AND OPEN PROBLEMS

This final section is divided strictly into three layers. The first records only what has actually been proved. The second clarifies the exact methodological scope of the paper. The third isolates open problems that emerge naturally from the proved theory. This separation is deliberate. It prevents the conclusion from blurring theorem-level achievements with directions that remain open.

Proved conclusions

The paper develops a finite variational framework for budgeted information-preserving subgraph selection on weighted relational structures. Within that framework, the following conclusions have been established.

- (i) **Discrete existence.** For every finite weighted relational structure G and every budget level B , if the feasible family $\mathcal{F}_B(G)$ is nonempty and finite, then the budgeted optimization problem admits at least one optimizer. Equivalently, the associated loss functional admits a minimizer on the feasible family. This is the basic attainment theorem for the discrete model.

- (ii) **Relaxed existence.** In the presence of a compact relaxed feasible region and an upper semicontinuous relaxed objective, the relaxed budgeted problem admits an optimizer. This statement is distinct from the discrete finite-attainment theorem and belongs to a different hypothesis package.
- (iii) **Order-theoretic optimizer structure.** Whenever the optimizer set is nonempty and the admissible family is finite, there exists an inclusion-minimal optimizer. Thus the theory yields canonical minimal representatives without requiring monotonicity of the objective.
- (iv) **Conditional connected representative theorem.** Under an explicit component-amalgamation hypothesis, the optimizer set contains a connected representative. This is a representative theorem only. It is not a classification of all optimizers.
- (v) **Value stability.** Under a uniform modulus-of-continuity assumption for the information-retention functional on a common feasible family, the optimal value is stable under perturbation of the ambient weights. In the Lipschitz case, one obtains a quantitative Lipschitz bound.
- (vi) **Optimizer stability under separation.** If the optimizer is unique and separated from all competitors by a positive isolation gap, then sufficiently small perturbations of the ambient weights preserve that optimizer uniquely. Thus local constancy of the optimizer requires a nondegeneracy condition.
- (vii) **Near-optimizer transport.** Even when exact optimizer stability is too strong, near-optimizer sets remain controllable under perturbation, with an explicit loss in the near-optimality threshold.
- (viii) **Approximation and recovery.** In the presence of a relaxed model together with a quantitative recovery operator, every relaxed optimizer yields a discrete near optimizer, possibly with controlled loss of value and controlled budget blow-up. Exact equality of relaxed and discrete values follows only under exact value-preserving recovery.
- (ix) **Negative results.** Two obstruction principles have been proved. First, continuity of the optimal value does not imply stability of the optimizer set in the absence of an isolation gap. Second, connectedness of the ambient graph does not imply existence of a connected optimizer. These negative results show that the positive hypotheses used earlier are not formal decoration.
- (x) **Verification in natural classes.** The abstract hypothesis package has been verified in at least two substantial model classes: additive edge-potential retention models and monotone submodular coverage-type retention models. Hence the framework is not empty and is mathematically anchored in concrete families related to weighted graph optimization, graph reduction surrogates, and submodular selection [3, 8, 15, 16, 17, 32, 36, 43, 44, 47].

Remark X.1. No stronger claim is made here. In particular, the paper does not prove a complete characterization of optimizer sets, does not establish generic uniqueness, and does not prove that connected representatives exist in all natural classes. Those matters remain outside the proved core.

Methodological remarks

The present framework has a deliberately restricted scope.

- (i) **Finite setting.** The theory is developed for finite weighted relational structures. This restriction is not accidental. It is the level at which the discrete variational problem, the optimizer set, and the perturbation layer can all be formulated cleanly without measure-theoretic or compactness complications beyond those explicitly introduced in the relaxed model.
- (ii) **Discrete and relaxed layers are separated.** The paper distinguishes sharply between the discrete finite-attainment theorem and the relaxed compact-existence theorem. The former uses finiteness only. The latter uses compactness and upper semicontinuity. These arguments are not merged.
- (iii) **Perturbation theory is weight-based and support-fixed.** The perturbation layer varies ambient weights on a fixed combinatorial support. The paper does not treat support-changing perturbations, graph edit perturbations, or structural instability under changing admissible families.
- (iv) **Structure theory is intentionally modest.** The structural section proves existence of minimal representatives and, under additional hypotheses, connected representatives. It does not attempt a full optimizer classification, a lattice description, or a uniqueness theorem.
- (v) **Stability of values and stability of optimizers are different phenomena.** The paper proves value stability under a modulus-of-continuity assumption. Optimizer stability requires more. In the finite setting, local constancy is derived only under a positive isolation gap. This separation is essential and is confirmed by the obstruction section.
- (vi) **Approximation is conditional.** The relaxation-and-recovery layer is retained only at the level justified by theorem content. No generic approximation claim is made without a concrete relaxed model and a proved recovery mechanism.
- (vii) **Natural-class verification is partial rather than universal.** The paper verifies the axioms in additive and coverage-type classes. It does not claim that every retention functional arising in graph theory, sparsification, reduction, clustering, or signal processing belongs to these classes or satisfies the full positive theorem package [7, 9, 10, 40, 41, 42].

Remark X.2. Accordingly, the mathematical contribution of the paper is best described as a first finite variational theory with theorem-level existence, selected structural representatives, perturbative value control, a conditional relaxation layer, and explicit negative boundary results. It is

not a universal theory of all graph pruning or graph reduction mechanisms.

Open problems

The proved theory leaves several natural directions unresolved.

Question X.3 (Abstract characterization of optimizer sets). *Under what natural intrinsic conditions on l and C can one characterize the optimizer set without passing to case-by-case combinatorial analysis?*

Remark X.4. Theorems proved here give existence and selected representatives, but not an intrinsic characterization of all optimizers. A satisfactory answer may require a deeper order-theoretic, variational, or submodular structure than what has been used in the present paper.

Question X.5 (Extension to higher-order structures). *Can the present finite variational theory be extended to hypergraphs, simplicial structures, or higher-order relational systems without losing the perturbation layer?*

Remark X.6. This question is mathematically natural because the formal language of weighted relational structures already suggests higher-order variants. However, the perturbative framework, the admissible-family formalism, and the recovery layer may all become more delicate in that setting. Related higher-order and generalized graph frameworks suggest that such an extension is plausible but nontrivial [30, 31, 29].

Question X.7 (Exact verification theorems beyond the model classes treated here). *Which abstract hypothesis packages admit exact verification theorems in natural graph classes beyond additive or submodular models?*

Remark X.8. The additive and coverage-based classes treated in Section IX are only first anchors. It remains open to determine which stronger positive hypotheses—for example connected-representative principles, exact recovery mechanisms, or sharper perturbative regularity—can be verified in graph classes tied more closely to spectral sparsification, graph reduction, or directed and higher-order cut structures [16, 17, 19, 32, 36, 28, 30].

Question X.9 (Sharp perturbative regularity of optimizer correspondences). *Beyond local constancy under an isolation gap, what is the correct general regularity theory for the optimizer correspondence? In particular, can one prove quantitative set-valued continuity for near-optimizer families under natural nondegeneracy hypotheses?*

Remark X.10. The present paper proves transport of near-optimizers but does not develop a full continuity theory for the optimizer correspondence. A sharper theory would likely require a more systematic metric or variational geometry on families of admissible substructures.

Question X.11 (Sharp relaxation gaps and recovery principles). *For which natural relaxed models can one compute or bound the relaxation gap explicitly, and under what structural assumptions does one obtain exact or asymptotically exact recovery?*

Remark X.12. The approximation section proves only conditional theorems. To make that layer decisive, one

would need concrete relaxed models in which the recovery operator is explicit and the resulting loss parameters can be estimated sharply. This problem lies at the interface of discrete variational theory, approximation, and graph reduction [7, 44, 47, 36].

Remark X.13. These open problems are not appended for breadth. They are the precise directions left unresolved by the proved core of the paper. Any later extension should preserve the same discipline used throughout the present article: exact hypotheses, explicit theorem boundaries, and no inflation of claims beyond what the proofs support.

ACKNOWLEDGEMENTS

The authors thanks Dr. Ramachandra R. K., Principal, Government College (Autonomous), Rajahmundry, for institutional support and encouragement. No funding was received for this work. The author declares no conflict of interest. This article is purely theoretical. No datasets were generated, analysed, or used in this study. It contains no studies involving human participants, animals, patient data, or personal data. Ethical approval was therefore not required.

REFERENCES

- [1] Miroslav Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(2):298–305, 1973.
- [2] Russell Merris. Laplacian matrices of graphs: A survey. *Linear Algebra and its Applications*, 197:143–176, 1994.
- [3] Fan R. K. Chung. *Spectral Graph Theory*, volume 92 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, 1997.
- [4] Fan Chung. Laplacians and the cheeger inequality for directed graphs. *Annals of Combinatorics*, 9(1):1–19, 2005.
- [5] Lars Hagen and Andrew B. Kahng. New spectral methods for ratio cut partitioning and clustering. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 11(9):1074–1085, 1992.
- [6] Ingemar J. Cox, Satish B. Rao, and Yu Zhong. "Ratio Regions": A Technique for Image Segmentation. In *Proceedings of the 13th International Conference on Pattern Recognition*, 1996.
- [7] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.
- [8] Ulrike von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17:395–416, 2007.
- [9] Reid Andersen, Fan Chung, and Kevin Lang. Local graph partitioning using PageRank vectors. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 475–486, 2006.
- [10] Reid Andersen, Fan Chung, and Kevin Lang. Using PageRank to locally partition a graph. *Internet Mathematics*, 4(1), 2007.
- [11] Daniel A. Spielman and Shang-Hua Teng. A local clustering algorithm for massive graphs and its application to nearly-linear time graph partitioning. *SIAM Journal on Computing*, 42(1):1–26, 2013.
- [12] Fan Chung and Olivia Simpson. Computing heat kernel pagerank and a local clustering algorithm. In *Combinatorial Algorithms (IWOCA 2014, LNCS 8986)*, 2015.
- [13] Santo Fortunato. Community detection in graphs. *Physics Reports*, 486(3–5):75–174, 2010.
- [14] Fragkiskos D. Malliaros and Michalis Vazirgiannis. Clustering and community detection in directed networks: A survey. *Physics Reports*, 533(4):95–142, 2013.
- [15] András A. Benczúr and David R. Karger. Randomized approximation schemes for cuts and flows in capacitated graphs. *SIAM Journal on Computing*, 44(2):290–319, 2015.
- [16] Daniel A. Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. *SIAM Journal on Computing*, 40(6):1913–1926, 2011.
- [17] Daniel A. Spielman and Shang-Hua Teng. Spectral sparsification of graphs. *SIAM Journal on Computing*, 40(4):981–1025, 2011.
- [18] Joshua Batson, Daniel A. Spielman, and Nikhil Srivastava. Spectral sparsification of graphs: Theory and algorithms. *Communications of the ACM*, 56(8):87–94, 2013.
- [19] Joshua Batson, Daniel A. Spielman, and Nikhil Srivastava. Twice-Ramanujan sparsifiers. *SIAM Journal on Computing*, 43(6):1704–1721, 2014.
- [20] Ioannis Koutis, Alex Levin, and Richard Peng. Faster spectral sparsification and numerical algorithms for SDD matrices. *ACM Transactions on Algorithms*, 12(2):17:1–17:16, 2015.
- [21] Yin Tat Lee and He Sun. Constructing linear-sized spectral sparsification in almost-linear time. *SIAM Journal on Computing*, 47(6):2315–2336, 2018.
- [22] Ioannis Koutis and Shen Chen Xu. Simple parallel and distributed algorithms for spectral graph sparsification. *ACM Transactions on Parallel Computing*, 3(2):8:1–8:14, 2016.
- [23] Michael Kapralov and David P. Woodruff. Spanners and sparsifiers in dynamic streams. In *Proceedings of the 2014 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 272–281, 2014.
- [24] Michael Dinitz, Robert Krauthgamer, and Tal Wagner. Towards resistance sparsifiers. In *APPROX/RANDOM 2015*, 2015.
- [25] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Graph sketches: Sparsification, spanners, and subgraphs. In *Proceedings of the 31st ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS)*, pages 5–14, 2012.

- [26] Timothy Chu, Yu Gao, Richard Peng, Sushant Sachdeva, Saurabh Sawlani, and Junxing Wang. Graph sparsification, spectral sketches, and faster resistance computation via short cycle decompositions. *SIAM Journal on Computing*, 49(6):FOCS18–185–FOCS18–220, 2020.
- [27] Gramoz Goranci, Monika Henzinger, and Pan Peng. The power of vertex sparsifiers in dynamic graph algorithms. In *25th Annual European Symposium on Algorithms (ESA 2017)*, pages 45:1–45:14, 2017.
- [28] Zhiyang He, Jason Li, and Magnus Wahlström. Near-linear-time, optimal vertex cut sparsifiers in directed acyclic graphs. In *29th Annual European Symposium on Algorithms (ESA 2021)*, pages 52:1–52:14, 2021.
- [29] Yu Chen, Sanjeev Khanna, and Ansh Nagda. Near-linear size hypergraph cut sparsifiers. In *61st IEEE Annual Symposium on Foundations of Computer Science (FOCS 2020)*, pages 61–72, 2020.
- [30] Tsz Chiu Kwok, Lap Chi Lau, and Kam Chuen Tung. Cheeger inequalities for vertex expansion and reweighted eigenvalues. In *63rd IEEE Annual Symposium on Foundations of Computer Science (FOCS 2022)*, pages 366–377, 2022.
- [31] Lap Chi Lau, Kam Chuen Tung, and Robert Wang. Cheeger inequalities for directed graphs and hypergraphs using reweighted eigenvalues. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC 2023)*, 2023.
- [32] Florian Dorfler and Francesco Bullo. Kron reduction of graphs with applications to electrical networks. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 60(1):150–163, 2013.
- [33] Huaijun Qiu and Edwin R. Hancock. Spectral simplification of graphs. In *Computer Vision – ECCV 2004*, volume 3024 of *Lecture Notes in Computer Science*, pages 114–126, 2004.
- [34] Zhiqiang Zhao and Zhuo Feng. Effective-resistance preserving spectral reduction of graphs. In *Proceedings of the 56th Annual Design Automation Conference (DAC 2019)*, 2019.
- [35] Zhuo Feng. GRASS: Graph spectral sparsification leveraging scalable spectral perturbation analysis. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 39(12):4944–4957, 2020.
- [36] Andreas Loukas. Graph reduction with spectral and cut guarantees. *Journal of Machine Learning Research*, 20(116):1–42, 2019.
- [37] Yu Jin, Andreas Loukas, and Joseph F. JaJa. Graph coarsening with preserved spectral properties. In *Proceedings of The 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 108 of *Proceedings of Machine Learning Research*, pages 4452–4462, 2020.
- [38] Mohammad Hashemi, Shiyang Zaman, Seyed Mehran Kazemi, Sriraam Narayanan, Caiwen Cai, Rose Yu, Neil Shah, and Reihaneh Rabbany. A comprehensive survey on graph reduction, 2024. arXiv preprint.
- [39] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, 2003.
- [40] David I. Shuman, Sunil K. Narang, Pascal Frossard, Antonio Ortega, and Pierre Vandergheynst. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Processing Magazine*, 30(3):83–98, 2013.
- [41] Aliaksei Sandryhaila and Jose M. F. Moura. Discrete signal processing on graphs. *IEEE Transactions on Signal Processing*, 61(7):1644–1656, 2013.
- [42] Stefania Sardellitti, Sergio Barbarossa, and Paolo Di Lorenzo. On the graph fourier transform for directed graphs. *IEEE Journal of Selected Topics in Signal Processing*, 11(6):796–811, 2017.
- [43] George L. Nemhauser, Laurence A. Wolsey, and Marshall L. Fisher. An analysis of approximations for maximizing submodular set functions—i. *Mathematical Programming*, 14:265–294, 1978.
- [44] Maxim Sviridenko. A note on maximizing a submodular set function subject to a knapsack constraint. *Operations Research Letters*, 32(1):41–43, 2004.
- [45] Andreas Krause, Ajit Singh, and Carlos Guestrin. Near-optimal sensor placements in gaussian processes. In *Proceedings of the 22nd International Conference on Machine Learning (ICML 2005)*, 2005.
- [46] Andreas Krause, Ajit Singh, and Carlos Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms, and empirical studies. *Journal of Machine Learning Research*, 9:235–284, 2008.
- [47] Andreas Krause and Daniel Golovin. Submodular function maximization. In *Tractability: Practical Approaches to Hard Problems*. Cambridge University Press, 2014.