mathematics

MDPI

Article

From Cell-Specific Heuristics to Transferable Structural Search for Ramsey Graph Construction

Sorin Liviu Jurj

School of Computers and Information Technology, "Politehnica" University of Timisoara, 2 V. Parvan Blvd, 300223 Timisoara, Romania; jurjsorinliviu@yahoo.de

Abstract

Recent automated search methods have improved lower bounds for several Ramsey numbers, but the strongest gains often depend on structured seeding and cell-specific heuristic discovery. This leaves open a more fundamental question: Can a useful search structure be transferred across related Ramsey cells rather than rediscovered independently for each target instance? This work proposes a teacher–student framework for transferable structural search in Ramsey graph construction, inspired by the structure-distillation logic of Physics Structure-Informed Neural Networks (Ψ-NNs). The framework builds compressed structural representations from teacher witnesses and search traces, extracts reusable motifs and relations, and reconstructs transfer candidates. These are refined by balanced search and, for weak R(3, s) cells, by exact small-cell supervision. The framework is evaluated as a proof of concept across five Ramsey cells under transfer, matched-compute, search, ablation, and interpretability settings, including a proportional shift-scaling baseline and a greedy triangle-closing baseline that probe the structure-validity frontier from complementary directions. Supplementary experiments cover seed robustness, budget sensitivity, transfer-neighborhood variation, structural-resolution changes, stronger exact supervision, cross-r teacher pooling, single-teacher configurations, and scaling behavior across graph sizes. The results show that the portfolio version of the framework is the strongest balanced transfer method in the current study, while a structure-dominant oracle achieves stronger witness-shape agreement but worse Ramsey-valid construction. These findings reveal a clear structure-validity frontier and suggest that transferable Ramsey search should be evaluated by how well structural priors survive the validity constraints of new cells.

Keywords: Ramsey theory; Ramsey numbers; automated search; transferable structural search; teacher–student distillation; physics structure-informed neural networks; graph construction

MSC: 05C55; 05C85; 68T20; 05-08

Check for updates

Academic Editor: Reza Saei

Received: 22 March 2026
Revised: 9 April 2026
Accepted: 17 April 2026
Published: 19 April 2026

Copyright: © 2026 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Ramsey theory studies the threshold at which unavoidable structure emerges in sufficiently large graphs [1]. In the classical two-color setting, the Ramsey number, $R(r, s)$, is the smallest integer, n, such that every graph on n vertices contains either a clique of size r or an independent set of size s [1,2]. Determining or bounding these numbers remains one of the central challenges of extremal combinatorics [2,3]. Exact values are known only in a small number of cases, and even improving a lower bound by one often requires a substantial constructive advance. As a result, Ramsey lower-bound research has long

served as a demanding testbed for new ideas in combinatorial optimization, symmetry exploitation, and graph construction [4,5].

Recent progress has shown that automated search can rediscover and refine sophisticated construction strategies for Ramsey problems [5,6]. In particular, modern search systems can combine structured initialization, iterative repair, and large-scale computation to produce new witness graphs and improved lower bounds [6,7]. These achievements are important, but they also expose a methodological tension. The strongest empirical gains often appear when a search is guided by highly structured seeds and procedures tuned to specific (r, s) cells. This raises a more general question: Are current systems learning broadly reusable structural principles, or are they mainly producing high-quality local heuristics specialized to individual cells?

We argue that this distinction matters. If progress depends primarily on cell-specific heuristic discovery, then each new Ramsey cell remains largely an independent engineering problem. In that regime, scaling search requires repeated human insight, repeated compute expenditure, and repeated rediscovery of related structural ideas. By contrast, if strong constructions share reusable organization across neighboring cells, then a more fundamental scientific problem emerges: not only how to search better within one cell, but how to identify, compress, and transfer structural information across a family of cells. This shifts the focus from isolated witness discovery toward the learning of transferable structural priors for extremal graph construction.

This paper takes that second view as a proof of concept. Rather than treating successful Ramsey constructions as independent endpoints of a search process, we treat them as evidence of latent shared structure. Our starting hypothesis is that strong Ramsey witnesses and successful repair trajectories contain recurring motifs, symmetry patterns, orbit-like relations, and reusable local transformation rules that are not unique to a single (r, s) instance. If such regularities can be extracted from high-capacity search procedures and re-expressed in compact form, then they may provide transferable inductive bias for neighboring Ramsey cells. The challenge is therefore to move from search outcomes to structure discovery and from structure discovery to reusable search policies.

To address this challenge, we propose a Ψ-NN [8]-inspired framework for transferable structural search in Ramsey graph construction. The framework draws on the methodological lesson of structure-discovery frameworks [8], not the literal use of physics-informed modeling [9]. A strong teacher first solves or approximates a difficult task. A compressed student then distills the teacher's behavior. An explicit extraction stage converts internal regularities into interpretable structural constraints. In our setting, the teacher is a high-capacity Ramsey search process that generates strong constructions and search trajectories. A student representation is trained to summarize and compress these behaviors. From that student, we extract recurring structural motifs and relations and then reconstruct explicit transfer candidates constrained by the discovered structure. The goal is not merely another learned heuristic, but a reusable and interpretable construction prior.

This perspective differs from code-mutation-driven approaches in both target and evaluation. Our aim is not primarily to optimize a search code for a single Ramsey cell, but to determine whether structural information learned in one region of the Ramsey table can improve search in another. Consequently, transfer is not an auxiliary property of the method; it is the main object of study. We ask whether a structure-distilled search pipeline can reduce computation, improve witness quality, or accelerate refinement on held-out neighboring cells relative to random-only search, structured seeding alone, and simpler transfer baselines. We also ask whether the extracted structures are interpretable enough to reveal shared combinatorial organization rather than functioning only as opaque latent artifacts.

The broader motivation is methodological. If successful Ramsey search is dominated by local heuristic engineering, then the field may continue to progress through isolated computational victories without developing a transferable theory of constructive search. If, however, reusable structural priors can be learned and exported across related cells, then the search process itself becomes a source of scientific structure, not just a source of witnesses. In that sense, the goal of this work is to test whether extremal combinatorial search can be advanced by moving from cell-specific heuristic discovery to transferable structural search.

Our main contributions are as follows.

1. We propose a Ψ-NN [8]-inspired framework that adapts teacher–student structure distillation from scientific neural model discovery to transferable combinatorial search for Ramsey graph construction.
2. We formulate automated Ramsey search as a transfer-oriented problem, where the goal is not only to find strong witnesses for a single cell, but to learn reusable structural priors across related Ramsey cells.
3. We develop a multi-stage pipeline consisting of teacher representation, student distillation, structure extraction, reconstruction, refinement, and portfolio selection, together with an exact-supervised branch for weak $R(3, s)$ cells.
4. We provide a balanced evaluation framework that separates structural similarity from Ramsey-valid construction through transfer, matched-compute, search, ablation, and interpretability experiments.
5. We show empirically that balanced transfer and structure-only agreement are different objectives: the portfolio version of the proposed framework is the strongest practical transfer method in our study, while a structure-dominant oracle baseline defines a separate structural frontier.
6. We demonstrate that exact small-cell supervision improves the weak $R(3, s)$ regime, while portfolio-based transfer remains the most effective overall strategy under the balanced evaluation criteria used in this proof-of-concept study. It is important to note that the present work does not claim to establish new Ramsey lower bounds or improve existing ones. The five target cells used in the evaluation already have published lower bounds certified by known witness graphs. The framework generates transfer candidates and evaluates their structural and validity properties relative to those known witnesses, but it does not produce new bound certificates. The study is therefore best understood as a proof-of-concept methodological investigation: it asks whether transferred structural priors can reduce exact clique burden and improve search quality relative to uninformed baselines, not whether they can push the frontier of known bounds.

The remainder of this paper is structured as follows. Section 2 reviews the relevant state of the art in Ramsey search, transferable structural methods, and the teacher–student structure-distillation ideas that motivate the present work. Section 3 introduces the proposed framework and presents its main methodological components, including teacher representation, student distillation, structure extraction, candidate reconstruction, refinement, and portfolio selection. Section 4 describes the experimental setup, target cells, evaluation metrics, baselines, and reporting protocol. Section 5 reports the empirical results across transfer, matched-compute, search, ablation, and interpretability settings. Section 6 discusses the main findings, limitations, and directions for future work. Finally, Section 7 concludes the paper. Additional robustness and sensitivity experiments are reported in the Supplementary Materials document.

2. State of the Art

Research on Ramsey numbers has long combined extremal combinatorics with constructive search [2,3]. Classical lower-bound work relies on explicit graph constructions, algebraic and number-theoretic methods, circulant and cyclic families, Paley-type graphs, difference sets, and carefully engineered local modifications [7,10]. These approaches established that progress in Ramsey lower bounds often comes not from generic optimization alone, but from exploiting deep structural regularities in graph families. In practice, many of the strongest witnesses are highly non-random and exhibit substantial symmetry, modularity, or repeated local organization. This tradition provides the main conceptual background for the present work: strong Ramsey constructions are typically structured objects, and effective search often succeeds by navigating that structure rather than ignoring it.

More recent work has shifted part of this effort toward automated search [4–6,11]. Large-scale heuristic search, stochastic local improvement, SAT-inspired filtering, evolutionary methods, and algorithm-discovery frameworks have all been used to find or verify improved lower-bound witnesses. The most relevant recent direction is automated meta-search over search procedures themselves, where systems discover cell-specific routines that combine initialization strategies, repair heuristics, and iterative refinement [6]. These methods have demonstrated that algorithmic search can recover strong witnesses and, in some cases, improve published bounds.

Recent advances in extremal Ramsey theory also reinforce the structural richness of the $r = 4$ regime. In particular, Mattheus and Verstraete's asymptotic analysis of $R(4, t)$ [12] provides further evidence that higher-order Ramsey cells exhibit nontrivial large-scale structure, making them a natural setting for transfer-oriented structural study.

At the same time, the strongest results appear to depend heavily on expert-informed structured seeds and specialized procedures tuned to particular (r, s) cells [6,7]. This leaves open whether such systems learn transferable combinatorial principles or primarily rediscover effective local heuristics around privileged initializations.

This question of transfer has received much less attention than raw bound improvement. Most Ramsey search pipelines are evaluated cell by cell, with performance measured by the best witness found for a fixed (r, s) instance [5,6]. Such evaluations are natural, but they do not directly test whether the search process has captured reusable structure across neighboring cells. As a result, the field still lacks a clear methodology for separating three effects: the benefit of random exploration, the benefit of manually chosen structural priors, and the benefit of learned transferable search structure. The present study is motivated by precisely this gap. We treat transfer across Ramsey cells not as a secondary convenience, but as an important missing criterion for evaluating whether a search method has discovered meaningful combinatorial organization.

Our approach is also informed by recent progress in structure discovery for scientific neural models, especially Ψ-NNs [8]. Ψ-NNs introduced a teacher–student framework in which a high-capacity teacher first solves a constrained learning problem, while a compressed student is trained through distillation and regularization toward extractable internal structure. A subsequent extraction stage identifies recurring parameter relations, and a reconstruction step converts them into an explicit structured model. The key methodological lesson is that it can be advantageous to decouple problem solving from structure discovery: first solve the task with a powerful teacher, then distill and extract the reusable organization in a smaller student. The supplementary analysis of Ψ-NNs further showed that distillation, structure extraction, and reconstruction all contribute materially to obtaining clean structure and strong downstream performance.

This methodological pattern is directly relevant to Ramsey search, even though Ramsey graph construction is not a physics-governed learning problem. What transfers is not the PDE language or the physical interpretation, but the design principle of teacher-guided distillation, explicit structure extraction, and reconstruction of reusable constrained models. In our setting, the teacher is a high-capacity Ramsey search process. The student is a compressed structural representation of successful constructions and search traces. The reconstruction stage yields transferable priors for unseen Ramsey cells. This is the specific sense in which the proposed framework is Ψ-NN-inspired.

Several adjacent lines of work are also relevant but distinct. Automated framework-generation methods in scientific machine learning use large language models or meta-learning systems to synthesize code, losses, or initializations, but they primarily target usability or optimization speed rather than interpretable structural transfer. In extremal combinatorics, automated search methods continue to prioritize witness quality at a fixed target cell rather than transfer across tasks. The recent AlphaEvolve-style approach to Ramsey numbers is especially important in this context [6,13]. It demonstrates that automated code evolution can produce strong search procedures and new lower bounds across multiple cells. However, its central empirical success remains tied to per-cell constructive performance. The present paper asks a different question: What structural information, if any, is reusable across Ramsey cells, and how should such transfer be evaluated independently of cell-specific search success?

In this sense, the present work is closest in spirit to efforts that seek not only better solutions, but better reusable representations of how those solutions are found.

To make the distinction precise: the contribution of the present paper is simultaneously algorithmic, conceptual, and empirical. Algorithmically, it proposes a multi-stage teacher–student pipeline that is new in the context of Ramsey search. Conceptually, it reframes automated Ramsey search as a transfer problem and introduces the structure-validity frontier as an evaluation criterion that is independent of raw bound improvement. Empirically, it provides the first systematic comparison of transfer baselines under matched-compute and interpretability protocols specifically designed to separate structural similarity from Ramsey-valid construction. This is distinct from AlphaEvolve-style approaches in a fundamental way: those systems optimize search codes for per-cell performance, whereas the present work explicitly evaluates whether learned structure transfers across cells and whether structural similarity to known witnesses is a reliable proxy for constructive success. The answer, supported by the experiments, is that it is not, a finding that has methodological implications beyond the specific cells studied here.

We build on the constructive tradition of Ramsey search, on the empirical success of recent automated search systems [6,13], and on the structure-distillation logic developed in Ψ-NNs [8]. Our main departure from prior work is to recast the core problem from cell-specific heuristic discovery to transferable structural search for families of extremal graph constructions.

3. Methodology

3.1. Overview

We propose a Ψ-NN-inspired framework for transferable structural search in Ramsey graph construction. The central idea is to adapt the teacher–student structure-distillation logic of Ψ-NNs [8] to a combinatorial setting. The goal is not only to search effectively within a single Ramsey cell [6,13], but also to learn reusable structural priors from known constructions and transfer them across related cells. In this framework, a strong teacher produces high-quality task behavior, a student distills that behavior into a compact rep-

resentation, structure is extracted from the student, and the extracted structure is used to reconstruct constrained candidates for new problem instances.

In the present setting, the task is to construct graphs that certify lower bounds on Ramsey numbers. A graph on n vertices certifies $R(r,s) > n$ if it contains neither a clique of size r nor an independent set of size s. The proposed framework treats known Ramsey witnesses and their associated search behavior as sources of transferable structural information. Instead of optimizing only for cell-specific performance, it attempts to recover recurring motifs, shift patterns, compatibility relations, and repair tendencies that can be reused on neighboring target cells.

The overall transferable structural search workflow is summarized in Figure 1.

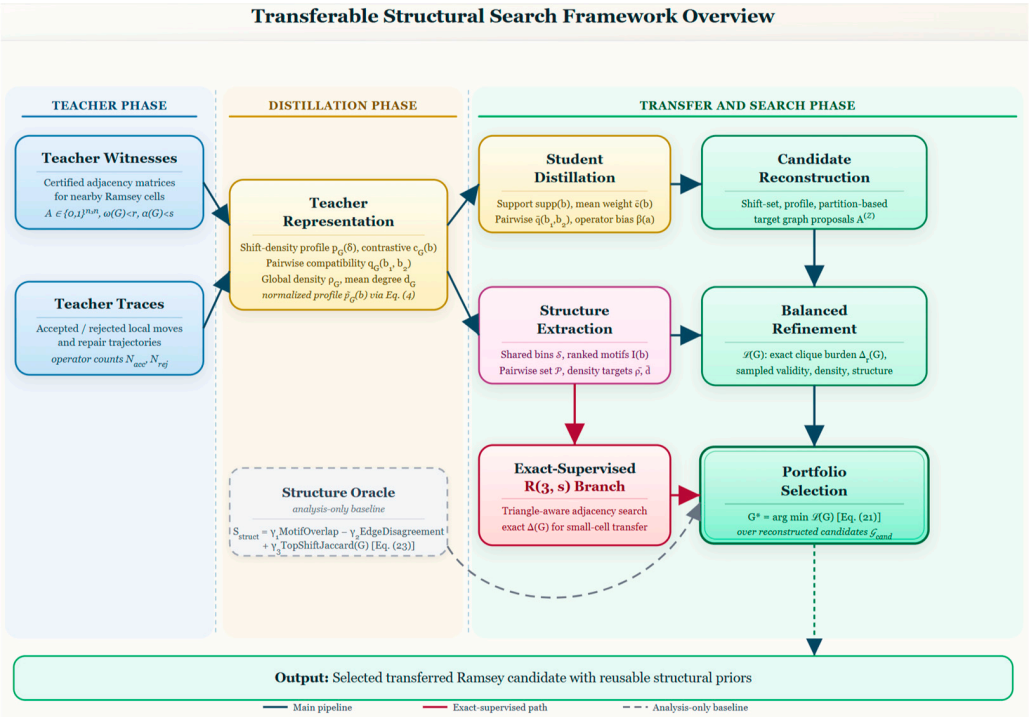


Figure 1. Overview of the proposed transferable structural search framework. Teacher witnesses and search traces are converted into structural teacher representations and distilled into a compact student model, from which reusable motif, compatibility, density, and operator priors are extracted. These priors guide candidate reconstruction, balanced refinement, and portfolio-based selection for unseen Ramsey cells. An additional exact-supervised branch supports weak $R(3,s)$ cases, while the structure-oracle path serves as an analysis-only baseline for studying the structure-validity frontier. Shaded regions denote the teacher, distillation, and transfer/search phases; colored connectors distinguish the main pipeline, the exact-supervised branch, and the analysis-only baseline.

Starting from teacher Ramsey witnesses and search traces, the framework builds structural teacher representations, distills them into a compact student model, extracts reusable structural priors, reconstructs target candidates, refines them under a balanced search objective, and finally selects among them through a portfolio mechanism. For weak $R(3,s)$ cells, the framework also includes an exact-supervised branch, while a separate structure-oracle path is used only for analysis of the structure-validity frontier.

The full transfer-oriented pipeline of the proposed framework is summarized in Algorithm 1.

The remaining subsections detail each stage of Algorithm 1, from teacher representation through to portfolio-based transfer selection.

Algorithm 1. Transferable structural search for a target Ramsey cell.

Input: target Ramsey cell c^* = (r,s,n), witness set W, configuration parameters θ
Output: selected transferred candidate graph G^*

1. Construct the teacher pool T = {w ∈ W:w ≠ c^*}, optionally restricted to witnesses with the same clique parameter r.
2. For each witness w∈T:
    1. Build a structural teacher representation R_w.
    2. If teacher traces are enabled, collect local search traces and distill trajectory priors.
3. Distill a compact student summary S from the teacher representations.
4. If trajectory distillations are available, attach them to S.
5. Extract an explicit structural prior P from S, including shared bins, ranked bins, pairwise motif relations, operator biases, and transferred global targets.
6. Construct a candidate set C using baseline, transfer, and structure-guided candidate families.
7. If r = 3, add an exact-supervised candidate obtained by exact triangle-guided refinement; otherwise use the default profile-based transfer fallback.
8. For each candidate G∈C:
    1. Apply exact repair.
    2. If applicable, run balanced transfer refinement guided by P.
    3. If enabled, generate a motif-polished variant and add it to the candidate set.
9. For every candidate in C, compute the balanced transfer objective and record the associated structural-alignment summaries.
10. Select the final candidate G^* through the portfolio rule, which ranks candidates primarily by the balanced objective and uses structural alignment only as an auxiliary selection signal when enabled.
11. Return G^*.

3.2. Problem Formulation

Let $G = (V, E)$ be a simple undirected graph on n vertices with adjacency matrix $A \in \{0,1\}^{n\times n}$. As can be seen in Equation (1), the adjacency matrix must be symmetric and must have zero diagonal (1):

$$A_{ij} = A_{ji},\ A_{ii} = 0 \text{ for all } i, j. \tag{1}$$

In Equation (1), $A_{ij} = 1$ indicates the presence of an edge between vertices i and j, while $A_{ij} = 0$ indicates its absence. The symmetry condition reflects the fact that the graph is undirected, and the diagonal constraint rules out self-loops.

The constructive Ramsey problem can then be written in terms of the clique number, $\omega(G)$, and the independence number, $\alpha(G)$. As shown in Equation (2), a graph certifies the lower bound $R(r, s) > n$ if it satisfies (2):

$$\omega(G) < r \text{ and } \alpha(G) < s. \tag{2}$$

Here, $\omega(G)$ denotes the size of the largest clique in G and $\alpha(G)$ denotes the size of the largest independent set. In other words, G must avoid both a clique of size r and an independent set of size s.

Since exact evaluation of these constraints becomes expensive as the graph size increases, practical search procedures often rely on a combination of exact checks, sampled

surrogates, and structural priors. The proposed framework is designed for this regime. Its purpose is not simply to reproduce known witnesses, but to learn structure from a training pool of related constructions and use that structure to improve transfer and search on unseen target cells.

3.3. Teacher Representation

The first stage of the framework builds a teacher representation for each witness in the training pool. Because many strong Ramsey constructions exhibit cyclic or near-cyclic organization [7,10], we summarize each witness using its cyclic shift profile. Let (3)

$$p_G(\delta) \tag{3}$$

denote the fraction of edges in graph G associated with cyclic shift δ. In Equation (3), δ indexes the cyclic distance between vertices and $p_G(\delta)$ measures how strongly that shift participates in the construction. This produces a motif-level description of the graph that can be compared across different witness graphs.

To make profiles comparable across different graph sizes, we normalize the shift space into a fixed number of bins. If b denotes a normalized bin, then the normalized motif profile of witness G can be written as (4)

$$\tilde{p}_G(b) = \frac{1}{\mid B_b \mid} \sum_{\delta \in B_b} p_G(\delta), \tag{4}$$

where B_b is the set of shifts assigned to bin b. As can be seen in Equation (4), $\tilde{p}_G(b)$ captures the average shift activity within each normalized bin. This provides a common structural coordinate system even when different witnesses have different numbers of vertices.

In addition to the normalized profile, each teacher representation stores the global density (5):

$$\rho_G = \frac{2|E|}{n(n-1)}, \tag{5}$$

where $\mid E \mid$ is the number of edges in G, and the mean degree (6):

$$d_G = \frac{1}{n} \sum_{i=1}^{n} \deg(i), \tag{6}$$

where $\deg(i)$ denotes the degree of vertex i. These global statistics help constrain downstream reconstruction so that transferred candidates remain in a plausible density regime.

To identify motifs that are strong relative to the overall graph density, we also define a contrastive profile (7):

$$c_G(b) = \tilde{p}_G(b) - \rho_G. \tag{7}$$

In Equation (7), positive values of $c_G(b)$ indicate bins expressed more strongly than the overall density baseline. Negative values indicate relatively under-expressed bins. This contrastive view is important because it separates genuinely informative motif structure from density effects alone.

Finally, teacher representations also record pairwise co-occurrence patterns among active bins. If b_1 and b_2 are two normalized bins, the pairwise compatibility score is defined in Equation (8):

$$q(b_1, b_2) = \frac{1}{m} \sum_{i=1}^{m} 1\left[p_{G_i}(b_1) > 0;\ \text{and};\ p_{G_i}(b_2) > 0\right] \tag{8}$$

In Equation (8), m is the number of teacher witnesses, $p_{G_i}(b)$ is the normalized shift-profile weight of bin b in witness G_i, and $1[\cdot]$ is the indicator function. The score

$q(b_1, b_2) \in [0, 1]$ measures the fraction of teacher witnesses in which both bins are simultaneously active. A value close to 1 indicates that the two bins reliably co-occur across the training pool, while a value close to 0 indicates that the pair is rare or cell-specific.

3.4. Student Distillation

Given a collection of teacher representations, $\{T_{G_1}, \ldots, T_{G_m}\}$, the framework distills them into a compact student model. The goal of this stage is to compress information distributed across multiple witnesses into a reusable structural summary. A single witness may contain a strong but cell-specific structure, whereas the student is intended to capture what persists across a neighborhood of related cells.

For each normalized bin, b, the student estimates how consistently that bin appears across the teacher set. As can be seen in Equation (9), the support of bin b is defined as (9):

$$\mathrm{supp}(b) = \frac{1}{m}\sum_{i=1}^{m} \mathbf{1}\left[c_{G_i}(b) > 0\right], \tag{9}$$

where m is the number of teacher witnesses and $\mathbf{1}[\cdot]$ is the indicator function. Equation (9) measures the fraction of teachers in which the bin is positively expressed relative to the background density.

The student also stores the average contrastive weight of each bin (10):

$$\bar{c}(b) = \frac{1}{m}\sum_{i=1}^{m} c_{G_i}(b). \tag{10}$$

In Equation (10), $\bar{c}(b)$ summarizes how strongly the bin tends to be expressed across the teacher set, not merely how often it appears.

Similarly, the average pairwise compatibility between bins is computed as (11)

$$\bar{q}(b_1, b_2) = \frac{1}{m}\sum_{i=1}^{m} q_{G_i}(b_1, b_2). \tag{11}$$

Equation (11) captures whether pairs of structural components tend to recur together across teacher witnesses. This provides a higher-order structural signal that is missing from simple top-bin averaging.

When teacher search traces are available, the student can also incorporate operator preferences. Let a denote a local action such as adding or removing a shift or corrective edge. We define an operator bias (12):

$$\beta(a) = \frac{N_{\mathrm{accept}}(a)}{N_{\mathrm{accept}}(a) + N_{\mathrm{reject}}(a)}, \tag{12}$$

where $N_{\mathrm{accept}}(a)$ and $N_{\mathrm{reject}}(a)$ denote the number of times action a was accepted or rejected in teacher refinement traces. In Equation (12), higher values of $\beta(a)$ indicate actions that are more frequently associated with successful local search behavior.

3.5. Structure Extraction

The distilled student is still a dense statistical object, so the next step is to convert it into an explicit structural prior. This extraction stage identifies recurrent bins, ranked bins, pairwise relations, and transferred global targets. The aim is to transform the student from a descriptive statistical summary into a sparse and interpretable guide for reconstruction.

A bin, b, is retained as a shared structural component when its support exceeds a threshold, τ_s, as shown in Equation (13):

$$b \in \mathcal{S} \Leftrightarrow \text{supp}(b) \geq \tau_s. \tag{13}$$

Here, $\mathcal{S}$ denotes the set of shared bins and τ_s is the support threshold. Equation (13) formalizes the idea that only bins recurring across enough teachers should be treated as transferable structure.

The ranked motif set is obtained by sorting bins according to a structural importance score. In the present framework, this score combines support and contrastive weight (14):

$$I(b) = \lambda_1 \,\text{supp}(b) + \lambda_2 \,\bar{c}(b), \tag{14}$$

where λ_1 and λ_2 are weighting parameters. In Equation (14), support rewards recurrence across teachers, while contrastive weight rewards strength relative to the density baseline.

Similarly, pairwise relations are extracted by thresholding the average compatibility matrix (15):

$$(b_1, b_2) \in \mathcal{P} \Leftrightarrow \bar{q}(b_1, b_2) \geq \tau_q, \tag{15}$$

where $\mathcal{P}$ is the set of compatible bin pairs and τ_q is a compatibility threshold. Equation (15) ensures that downstream reconstruction favors motif combinations that recur consistently in the teacher set.

The framework also transfers global density and degree targets from the teacher pool. If ρ_{G_i} and d_{G_i} are the density and mean degree of teacher G_i, then the transferred targets are (16)

$$\bar{\rho} = \frac{1}{m}\sum_{i=1}^{m} \rho_{G_i}, \; \bar{d} = \frac{1}{m}\sum_{i=1}^{m} d_{G_i}. \tag{16}$$

Equation (16) provides simple but effective global constraints that help keep reconstructed candidates within plausible structural regimes.

3.6. Candidate Reconstruction

Given an extracted structural prior and a new target cell, the framework reconstructs one or more candidate graphs. The simplest reconstruction mode selects a set of active target bins and maps them back to target-size shifts. Let $\mathcal{Z}$ denote the selected shift set. The corresponding circulant-style candidate graph can be written as (17)

$$A_{ij}^{(\mathcal{Z})} = \begin{cases} 1, & \text{if } \mid i-j \mid_n \in \mathcal{Z}, \\ 0, & \text{otherwise,} \end{cases} \tag{17}$$

where $\mid i-j \mid_n$ denotes the cyclic distance between vertices i and j modulo n. In Equation (17), the adjacency structure is determined directly by the transferred set of active shifts.

A second reconstruction mode uses the full shift-density profile rather than only a sparse set of top motifs. If $\pi(\delta)$ denotes the transferred target-size profile, then candidate edges can be sampled with probabilities proportional to $\pi(\delta)$. This profile-transfer mechanism preserves broader motif distributions and is less brittle than sparse shift selection alone.

A third reconstruction mode uses coarse partition structure. Let the vertex set be partitioned into blocks $V_1, \ldots, V_K$. The block-density template is then defined by (18):

$$M_{uv} = \frac{|\{(i,j) \in E : i \in V_u, j \in V_v\}|}{|V_u||V_v|} \quad \text{for } u \neq v, \tag{18}$$

with an analogous within-block definition for $u = v$. In Equation (18), M_{uv} records the density between partition blocks and provides a coarse structural scaffold for transfer.

For weak $R(3, s)$ cells, the framework also uses an exact-supervised reconstruction branch. Starting from broad-profile seeds, this branch refines adjacency-level candidates using exact triangle information. If $\Delta(G)$ denotes the exact number of triangles in graph G, then triangle reduction becomes a central signal for small-cell refinement.

3.7. Transfer Refinement

Transferred candidates are then improved by local search. The refinement stage alternates between modifying the active shift set and modifying a bounded set of corrective edges. To guide this process, we define a balanced search objective that combines exact validity, sampled validity, density control, and structural consistency.

As can be seen in Equation (19), the refinement score for a candidate graph, G, is (19):

$$\begin{aligned} L(G) = w_1 \log(1 + \Delta_r(G)) + w_2 \hat{C}_r(G) + w_3 \hat{I}_s(G) \\ + w_4 |\rho_G - \overline{\rho}| + w_5 D_{\text{struct}}(G), \end{aligned} \tag{19}$$

where $\Delta_r(G)$ is the exact number of forbidden r-cliques when exact computation is feasible, $C_r(G)$ is the sampled r-clique violation rate, $I_s(G)$ is the sampled independent-set violation rate, ρ_G is the graph density, ρ is the transferred density target, and $D_{\text{struct}}(G)$ is a structural mismatch term. The sampled rates are defined as (20)

$$\begin{aligned} C_r(G) &= \tfrac{1}{N} \textstyle\sum_{i=1}^{N} 1[S_i^{(r)} \text{ is a clique in } G], \\ I_s(G) &= \tfrac{1}{N} \textstyle\sum_{i=1}^{N} 1[S_i^{(s)} \text{ is an independent set in } G] \end{aligned} \tag{20}$$

where $S_i^{(r)}$ and $S_i^{(s)}$ are uniformly sampled subsets of size r and s, respectively, and N is the number of sampled subsets. Larger values of $C_r(G)$ and $I_s(G)$ indicate more frequent violations of the Ramsey constraints.

Equation (19) avoids the flat-penalty problem that arises when invalid candidates are only classified as bad, without distinguishing slightly invalid graphs from highly invalid ones. The logarithmic clique term preserves sensitivity in the small-violation regime, while the sampled terms allow the objective to remain tractable on larger instances.

When structural guidance is used explicitly, the structural mismatch term can be written as (21):

$$D_{\text{struct}}(G) = \eta_1 \parallel \tilde{p}_G - \tilde{p}_{\text{target}} \parallel_1 + \eta_2 D_{\text{pair}}(G), \tag{21}$$

where $\tilde{p}_{\text{target}}$ is the transferred target profile, $D_{\text{pair}}(G)$ is a penalty for violating extracted pairwise compatibilities, and η_1, η_2 are weighting parameters. In Equation (21), the first term measures profile mismatch and the second penalizes incoherent motif combinations. Specifically (22):

$$D_{\text{pair}}(G) = \frac{1}{|\mathcal{P}|} \sum\nolimits_{(b_1, b_2) \in \mathcal{P}} (1 - q(b_1, b_2)) \cdot 1[p_G(b_1) > 0; \text{and}; p_G(b_2) > 0] \tag{22}$$

where $\mathcal{P}$ is the set of extracted bin pairs and $q(b_1, b_2)$ is the pairwise compatibility score from Equation (8). This term penalizes candidate graphs that simultaneously activate bin pairs which rarely co-occur in the teacher pool, thereby discouraging structurally incoherent shift combinations.

3.8. Portfolio Selection

A key empirical observation of this work is that no single reconstruction mode is consistently best across all target cells. The framework therefore uses a portfolio rather

than committing to one transfer strategy. For each target cell, the framework constructs several candidates, refines them when appropriate, and then selects the strongest result under a balanced objective.

If $\mathcal{G}_{\text{cand}} = \{G_1, \ldots, G_K\}$ denotes the set of refined candidate graphs, then the final selected graph is (23):

$$G^{\star} = \arg\min_{G \in \mathcal{G}_{\text{cand}}} \mathcal{L}(G), \tag{23}$$

where $\mathcal{L}(G)$ is the balanced refinement objective from Equation (19). In Equation (23), the portfolio mechanism treats transfer as a competition among reusable structural priors rather than as a commitment to a single representation.

This portfolio view is central to the method. It reflects the empirical reality that transferable Ramsey structure is heterogeneous. Some cells benefit more from sparse motif transfer, others from broader profile transfer, and still others from exact-supervised refinement. The portfolio therefore is not simply an engineering convenience; it is part of the conceptual contribution of the framework.

The portfolio-guided search procedure used in the search experiments is summarized in Algorithm 2.

Algorithm 2. Portfolio-guided search for a target Ramsey cell.

Input: target Ramsey cell $c^{*} = (r, s, n)$, extracted structural prior P, teacher pool T, configuration parameters θ

Output: selected search candidate graph G^{*}

1. Construct four independently initialized search candidates using random local search, structured-seed local search, framework-guided local search, and partition-guided local search.
2. Split the available search budget across these candidate families and run an initial refinement stage for each family.
3. Retain the best candidate produced by each search family.
4. Compare the family-level candidates under the balanced objective and select the strongest intermediate candidate.
5. Allocate the remaining search budget to an additional refinement stage starting from that intermediate candidate.
6. Infer or reuse its active shifts and corrective edges, combine them with the ranked structural prior, and continue balanced local refinement under the objective in Equation (19).
7. Compare the intermediate portfolio winner and the final refined candidate under the balanced objective.
8. Return the stronger of these two candidates as the final portfolio-guided search result.

This search-time portfolio is evaluated separately from the transfer suites and is included to test whether the extracted structural priors also improve constructive refinement behavior.

3.9. Structure-Dominant Analysis Baseline

To separate structural similarity from Ramsey-valid construction, we also define a structure-dominant analysis baseline called "*structure_oracle_transfer*". This baseline explicitly selects candidates by structure-oriented criteria rather than by balanced Ramsey-valid refinement.

If $S_{\text{struct}}(G)$ denotes a structural agreement score, then the oracle baseline selects (24)

$$G^{\dagger} = \arg\max_{G \in \mathcal{G}_{\text{cand}}} S_{\text{struct}}(G), \tag{24}$$

where (25)

$$\begin{aligned} S_{\text{struct}}(G) = \gamma_1 \text{MotifOverlap}(G) - \\ \gamma_2 \text{EdgeDisagreement}(G) + \gamma_3 \text{TopShiftJaccard}(G). \end{aligned} \tag{25}$$

In Equations (24) and (25), $\gamma_1, \gamma_2, \gamma_3$ are positive weights, $MotifOverlap$ measures agreement in motif frequencies, $EdgeDisagreement$ measures adjacency mismatch, and $TopShiftJaccard$ measures overlap among dominant shifts. This oracle is not intended as a fair deployable transfer method. Instead, it is used to expose the structure-validity frontier by showing how far structure-oriented selection can go when validity is not the primary objective.

3.10. Summary

The proposed framework adapts the teacher–student structure-distillation logic of Ψ-NNs to extremal graph construction. The method begins with teacher witnesses and search traces and distills them into compact structural summaries. It then extracts reusable motif and relation priors, reconstructs target candidates, refines them under balanced search objectives, and selects the final output through a portfolio strategy.

The purpose of this framework is not only to improve local Ramsey search, but to test a broader hypothesis: that successful Ramsey constructions contain transferable structural organization that can be discovered, compressed, and reused across neighboring cells. The experiments in the next sections evaluate this hypothesis both in terms of Ramsey-valid performance and in terms of the structure-validity frontier revealed by the analysis baselines.

To illustrate the full pipeline concretely, Figure 2 traces a complete execution of the framework for target cell R(3, 13) with $n = 60$, using a single R(3, 18) teacher witness.

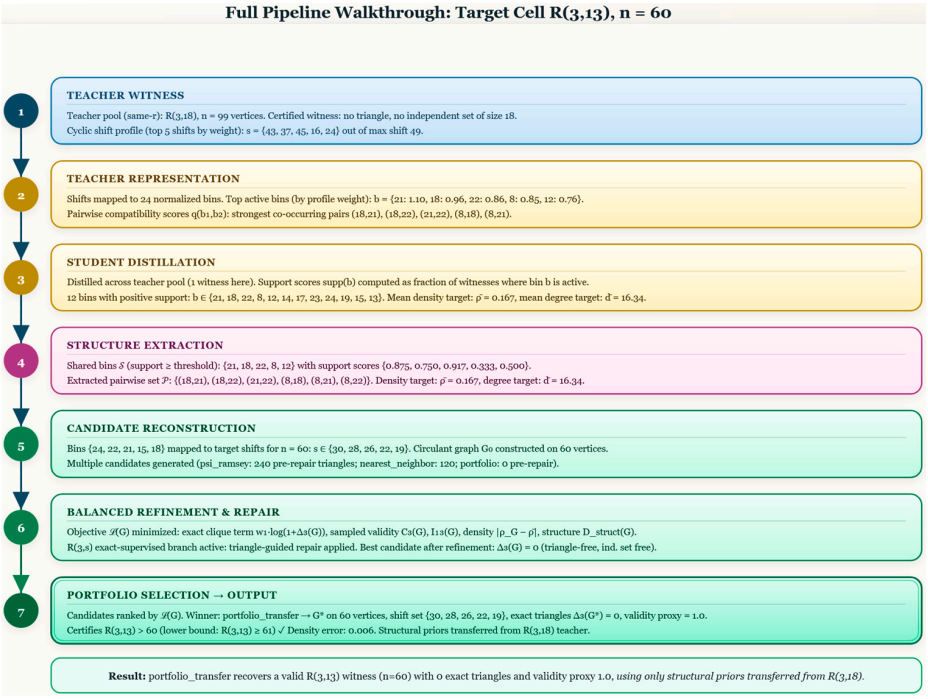


Figure 2. Full pipeline walkthrough for target cell R(3, 13), $n = 60$. Starting from a single R(3, 18) teacher witness ($n = 99$), the framework proceeds through seven stages: teacher representation (cyclic

shift profile mapped to 24 normalized bins), student distillation (support and contrastive weights aggregated across the teacher pool), structure extraction (shared bins and compatible pairs identified), candidate reconstruction (bins mapped to target shifts for $n = 60$), balanced refinement (objective $\mathscr{L}(G)$ minimized with exact triangle supervision active), and portfolio selection (best candidate chosen under the balanced objective). The selected graph contains zero exact triangles and achieves validity proxy 1.0, certifying R(3, 13) > 60.

4. Experimental Setup

4.1. Evaluation Objectives

The experimental study is designed to answer three related questions. First, can structural information learned from known Ramsey witnesses improve transfer to neighboring target cells? Second, does such transfer remain effective under matched-compute, rather than only under favorable search conditions? Third, how should one distinguish between a method that reproduces the structure of known witnesses and a method that actually improves Ramsey-valid construction?

To answer these questions, we evaluate the proposed framework under several complementary suites. The transfer suite measures overall cross-cell effectiveness. The matched-compute suite controls for search budget. The search suite evaluates constructive refinement behavior. The ablation suite isolates the role of the main methodological components. Finally, the interpretability suite compares balanced transfer against explicitly structure-dominant baselines. Together, these suites are intended to reveal not only which method performs best, but why.

4.2. Dataset and Target Cells

All experiments are based on a collection of published Ramsey witness graphs, represented here as adjacency matrices and assembled from publicly available Ramsey-number construction resources associated with recent automated Ramsey-search work [6,13]. These witnesses define both the teacher pool and the target cells used throughout the study. The synchronized evaluation considers the five target cells (26):

$$R(3,13), R(3,18), R(4,13), R(4,14), R(4,15). \tag{26}$$

As can be seen in Equation (26), the study includes both weak $R(3,s)$ cases and larger $R(4,s)$ cases. This mixture is useful because it allows us to test both broad structural transfer and exact-supervised small-cell refinement.

For each target cell, the remaining witnesses form the training pool, subject to the same-r filtering used in the main configuration. This creates a local transfer setting in which structure must be inferred from nearby cells rather than from arbitrary unrelated instances. Since the target set is small and specialized, the present experiments should be interpreted as a proof-of-concept study, not as a large-scale benchmark.

4.3. Teacher, Student, and Structural Configuration

Each teacher witness is converted into a structural representation using its cyclic shift profile, density, degree statistics, and pairwise motif relations, as described in Section 3. These teacher representations are then distilled into a compact student model. From the student, the proposed framework extracts shared bins, ranked bins, pairwise compatibilities, and transferred density and degree targets.

If $\mathcal{B}$ denotes the set of normalized bins, then the extracted shared structure is (27):

$$\mathcal{S} = \{b \in \mathcal{B} : \mathrm{supp}(b) \geq \tau_s\}, \tag{27}$$

where $\mathrm{supp}(b)$ is the support score defined earlier and τ_s is the support threshold. Equation (27) summarizes how the structural prior is filtered before reconstruction. In practice, the synchronized runs use a fixed number of normalized bins, fixed support thresholds, and a fixed maximum structural pool so that all methods are evaluated under the same structural resolution.

Teacher search traces are also incorporated when available. These traces provide additional information about which structural components tend to be added or removed during successful refinement. The resulting operator preferences are passed into the search stage as soft biases rather than hard constraints.

4.4. Candidate Families and Baselines

Table 1 provides an overview of all candidate-generation methods evaluated in this study, including simple baselines, transfer baselines, the proposed portfolio framework, and the analysis-only oracle.

Table 1. Summary of candidate-generation methods. Simple baselines do not use teacher structural information. Transfer baselines use teacher witnesses in different ways. The portfolio method combines multiple strategies. The oracle baseline is used only for frontier analysis and is not a deployable transfer method.

Method	Type	Construction Principle
random_density	Simple baseline	Random graph at target witness density
random_circulant	Simple baseline	Random circulant with uniformly sampled shift set
structured_seed	Simple baseline	Circulant reconstructed from teacher shift profile without refinement
nearest_neighbor_transfer	Transfer baseline	Shift profile from the teacher witness closest to target by normalized bin distance
teacher_mean_transfer	Transfer baseline	Shift profile averaged across all teacher witnesses in the pool
exact_triangle_transfer	Transfer baseline	Greedy shift selection guided by exact triangle counts during local refinement
hierarchical_coarsening_transfer	Transfer baseline	Shift profile of smallest teacher scaled proportionally to target graph size
flag_density_transfer	Transfer baseline	Greedy shift selection combining teacher frequency score with triangle-closing bonus (Equation (29))
portfolio_transfer	Proposed framework	Ensemble selecting the best candidate across all transfer strategies by composite objective
structure_oracle_transfer	Analysis baseline	Candidate selected by structural agreement with target witness; not a deployable method
psi_ramsey_transfer	Proposed framework	Direct distilled student transfer without portfolio selection

If $\mathcal{M}$ denotes the set of candidate-generation modes, then the candidate set for a target cell is (28):

$$\mathcal{G}_{\mathrm{cand}} = \{G^{(m)} : m \in \mathcal{M}\}. \tag{28}$$

In Equation (27), each $G^{(m)}$ is the result of one reconstruction strategy, possibly followed by repair and refinement. This notation is useful because the final framework decision is not based on a single mode, but on a portfolio over several candidate families.

To clarify the distinction between transfer quality and witness-shape similarity, we also include the analysis baseline structure_oracle_transfer. This baseline is not treated as a fair deployable method. Instead, it estimates the structural frontier by selecting candidates according to structural agreement with the target witness.

hierarchical_coarsening_transfer constructs a candidate by scaling the shift profile of the smallest available teacher witness proportionally to the target graph size. Specifically, each shift, s, in the teacher profile is mapped to a target shift, $s' = \text{round}(s \cdot \text{n_target}/\text{n_teacher})$, clamped to the valid range, and the top-k shifts by transferred weight are used to reconstruct the candidate. This method does not rely on normalized bins or shared-support filtering and is therefore complementary to the existing nearest-neighbor and distribution-profile baselines. flag_density_transfer builds a candidate through a greedy shift-selection procedure that combines a unary frequency score derived from the teacher shift profile with a triangle-closing bonus. The greedy selection criterion is defined in Equation (29):

$$\text{score}(s) = \text{unary}(s) + \lambda \cdot |t \in S_{\text{selected}} : \min(s+t, n-s-t) \in S_{\text{selected}}| \tag{29}$$

In Equation (29), $\text{unary}(s)$ is the frequency weight of shift s in the teacher profile, $\lambda = 0.5$ is the triangle bonus weight, and S_{selected} is the set of shifts already chosen. At each step, the shift maximizing this score is added to the selected set. This procedure is motivated by the flag-algebraic perspective that triangle-dense circulant graphs tend to avoid large independent sets, and it is not equivalent to any existing transfer strategy in the evaluation.

4.5. Evaluation Metrics

The experiments measure both Ramsey-valid quality and structural agreement. On the validity side, we report exact r-clique counts when feasible, exact r-clique-free indicators, sampled clique and independent-set violation rates, and a combined validity proxy. On the structural side, we report density error, degree mean absolute error, motif overlap, edge disagreement rate, top-shift precision, and top-shift Jaccard similarity.

The density error for a candidate graph, G, is defined in Equation (30):

$$E_\rho(G) = \mid \rho_G - \rho_{\text{ref}} \mid, \tag{30}$$

where ρ_G is the density of the candidate graph and ρ_{ref} is the reference density, typically taken from the target witness. In Equation (30), smaller values indicate closer agreement with the expected density regime.

Similarly, if $d_G(i)$ and $d_{\text{ref}}(i)$ denote the degrees of vertex i in the candidate and reference graphs, respectively, then the degree mean absolute error is given by (31):

$$E_d(G) = \frac{1}{n} \sum_{i=1}^{n} |d_G(i) - d_{\text{ref}}(i) \mid . \tag{31}$$

Equation (31) measures how well the candidate preserves the degree profile of the target structure.

To quantify motif agreement, we use a normalized overlap score between candidate and reference shift profiles (32):

$$\text{MotifOverlap}(G) = \frac{\sum_{b \in \mathcal{B}} \min(\tilde{p}_G(b), \tilde{p}_{\text{ref}}(b))}{\sum_{b \in \mathcal{B}} \max(\tilde{p}_G(b), \tilde{p}_{\text{ref}}(b))}. \tag{32}$$

As shown in Equation (32), motif overlap behaves like a profile-level intersection-over-union score. Larger values indicate stronger motif-level agreement with the target witness.

For top-shift agreement, we use Jaccard similarity. If T_G and T_{ref} denote the sets of dominant shifts in the candidate and reference graphs, then (33)

$$J_{\text{top}}(G) = \frac{\mid T_G \cap T_{\text{ref}} \mid}{\mid T_G \cup T_{\text{ref}} \mid}. \tag{33}$$

In Equation (33), a larger value indicates stronger agreement in the dominant motif support.

Finally, the balanced validity proxy is defined as a binary or averaged indicator of whether the candidate avoids sufficiently many sampled forbidden structures while staying close to the target density regime. This proxy is used to summarize partial success when exact validity is not yet achieved.

4.6. Transfer and Matched-Compute Protocols

In the transfer suite, each target cell is evaluated against a common set of baselines and framework variants. Methods are ranked per metric, and an overall mean-rank summary is then computed across the transfer metrics. If a method m has rank $R_{m,k}$ on metric k, then its aggregate rank is (34)

$$\bar{R}_m = \frac{1}{K}\sum_{k=1}^{K} R_{m,k}, \tag{34}$$

where K is the number of reported metrics. Equation (34) avoids reducing the experiment to a single hand-chosen scalar objective and makes tradeoffs visible.

The matched-compute suite follows the same baseline structure but constrains the amount of search and refinement more tightly. Its purpose is to test whether observed improvements come from better structural transfer rather than from simply spending more search effort. Methods that remain competitive under this protocol provide stronger evidence of genuine transfer.

4.7. Search Protocol

The search suite evaluates constructive refinement directly. Here the goal is not only to transfer structure, but also to test whether a method can guide local search toward lower-scoring and more valid candidates. We compare random local search, structured-seed local search, framework-guided search, partition-guided search, and portfolio-guided search. These methods are evaluated using the same family of exact and sampled validity surrogates introduced in Section 3.

For each method, the best score reached during refinement is recorded as (35)

$$\mathcal{L}_{\min} = \min_{t \le T} \mathcal{L}(G_t^{(m)}), \tag{35}$$

where $G_t^{(m)}$ is the graph produced by method m at iteration t, T is the maximum search horizon, and $\mathcal{L}(\cdot)$ is the balanced objective from Equation (19). Equation (35) allows the search suite to compare methods by the strongest candidate they discover during refinement rather than by an arbitrary intermediate state.

Because search behavior can vary by random seed, we also record replicate summaries for the synchronized code base. These replicate files are used to verify whether observed behavior is robust or only a single-run artifact.

4.8. Ablation and Frontier Analysis

The ablation suite isolates the role of the main components of the proposed framework. In particular, it compares structured seeding, teacher averaging, reduced student transfer, and the full distilled pipeline. This helps determine whether performance is truly due to transferable structure extraction or merely due to inherited seed quality.

The frontier analysis is carried out through the interpretability suite and the structure_oracle_transfer baseline. This part of the evaluation is crucial for the current work's main argument. It tests whether high structural agreement with known witnesses coincides

with good Ramsey-valid construction or whether these are genuinely competing objectives. In practice, this distinction turns out to be one of the central findings of the study.

4.9. Implementation and Reporting

All experiments are executed through a unified evaluation pipeline that standardizes target selection, configuration, metric computation, and summary reporting across suites. This design supports consistent per-target and aggregate comparisons across all evaluation suites. For the transfer and search suites, repeated-seed summaries are also recorded to assess the stability of observed behavior.

Because the study is designed as a proof of concept, the emphasis is on methodological clarity and comparative behavior across suites. The goal is to interpret transfer and search tradeoffs, not to claim broad generalization across the Ramsey table.

Additional robustness and sensitivity experiments are reported in the Supplementary Materials document. These include repeated-seed analyses, budget stress tests, transfer-neighborhood perturbations, structural-resolution sensitivity studies, stronger exact-supervision settings, and repeated ablation and interpretability runs.

4.10. Summary

The experimental setup is designed to test the proposed framework as a transfer framework rather than only as a single-cell search heuristic. By combining transfer evaluation, matched-compute comparison, search analysis, ablations, and a structure-dominant oracle baseline, the setup makes it possible to study both practical transfer performance and the structure-validity frontier that emerges in Ramsey graph construction.

5. Results

This section presents the empirical evaluation of the proposed framework across transfer, matched-compute, search, ablation, and interpretability settings. The results show that the portfolio version of the proposed framework is the strongest balanced transfer method in the current proof-of-concept study, while the broader evaluation also reveals a clear frontier between structural similarity and Ramsey-valid construction. Taken together, these experiments clarify both the practical strengths of the method and the limits of structure-only transfer.

5.1. Transfer Performance

As can be seen in Table 2, hierarchical_coarsening_transfer and flag_density_transfer occupy distinct positions in the transfer ranking.

Table 2. Transfer performance across target Ramsey cells.

Method	Overall Mean Rank	Exact r-Clique Count	Density Error	Motif Overlap	Edge Disagreement Rate	Top-Shift Jaccard	Top-Shift Precision	Validity Proxy
portfolio_transfer	4.186	3834.2	0.0075	0.0750	0.3533	0.1804	0.3000	0.7000
structure_oracle_transfer	4.500	8361.2	0.0608	0.1926	0.1227	0.8400	0.9000	0.6000
random_circulant	4.514	6115.2	0.0032	0.0707	0.3645	0.1475	0.2500	0.5000
random_density	5.243	9158.4	0.0026	0.1029	0.3724	0.1014	0.2000	0.5000
exact_triangle_transfer	5.529	6843.8	0.0153	0.0745	0.3775	0.0133	0.0500	0.6000
hierarchical_coarsening_transfer	5.814	3280.2	0.0189	0.0698	0.3601	0.0914	0.1500	0.6000
nearest_neighbor_transfer	6.271	92,658.2	0.0258	0.0591	0.3824	0.1033	0.2000	0.5000
flag_density_transfer	6.557	14,203.6	0.0201	0.0712	0.3589	0.0771	0.1000	0.5000
teacher_mean_transfer	7.114	98,941.2	0.0260	0.0579	0.3916	0.1166	0.2000	0.5000
psi_ramsey_transfer	7.629	96,503.6	0.0260	0.0606	0.3863	0.1166	0.2000	0.5000

hierarchical_coarsening_transfer achieves the second-lowest mean exact r-clique count among all evaluated methods, including zero exact cliques on the R(3,13) and R(3,18) target cells, which demonstrates that proportional shift scaling from the smallest teacher witness is a surprisingly effective validity-oriented strategy. Its aggregate mean rank places it sixth overall, such that it is competitive with the existing baselines despite the simplicity of its construction. flag_density_transfer ranks eighth overall and performs competitively on the R(3,s) target cells but produces substantially higher exact clique counts on the R(4,s) cells. This result is informative: the triangle-closing bonus encoded in Equation (29) is well-suited to the triangle-avoidance structure of R(3,s) witnesses but does not generalize cleanly to the denser and more complex R(4,s) regime. The two methods therefore expand the frontier coverage in complementary directions without displacing portfolio_transfer as the strongest balanced overall strategy.

5.2. Structure-Validity Frontier

As can be seen in Figure 3, there is a clear separation between structure-oriented selection and balanced Ramsey-valid transfer.

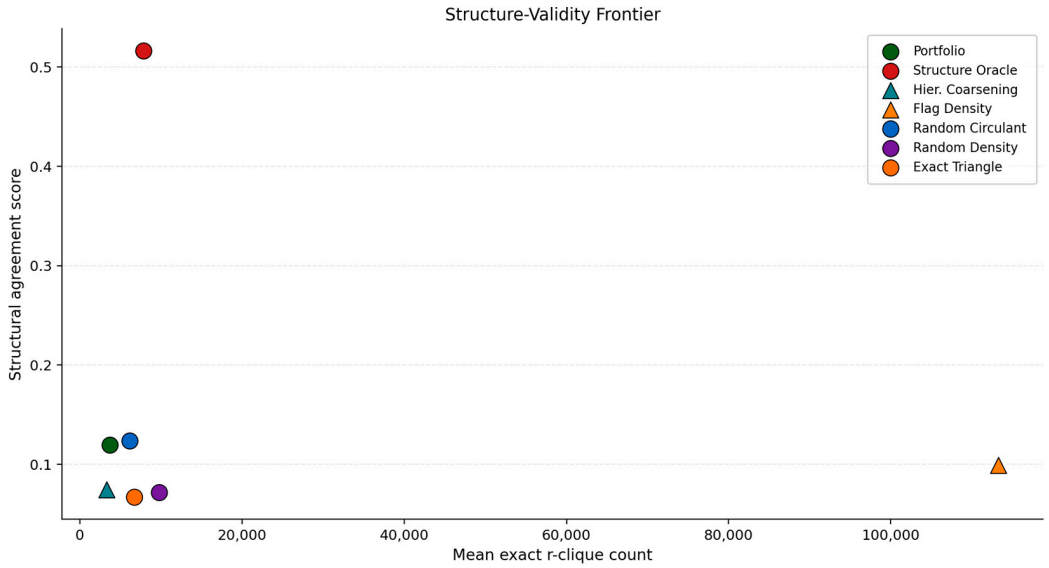


Figure 3. Tradeoff between structural agreement and Ramsey-valid transfer. The horizontal axis represents the mean exact r-clique count, where lower is better. The vertical axis represents structural agreement, measured as the average of motif overlap and top-shift Jaccard similarity. The figure shows that structure_oracle_transfer lies on the structure-dominant frontier, whereas portfolio_transfer attains the strongest balance between structural agreement and Ramsey-valid construction.

The figure does not reproduce all structural metrics individually; instead, it summarizes them through a composite structural agreement score and contrasts that view with the exact clique burden. This makes the main frontier visible in a single panel. If a structural score is defined as (36)

$$S_{\text{struct}}(G) = \frac{1}{2}\,\text{MotifOverlap}(G) + \frac{1}{2}\,J_{\text{top}}(G), \tag{36}$$

then Figure 3 visualizes the tradeoff between $S_{\text{struct}}(G)$ and exact Ramsey validity. In Equation (36), $\text{MotifOverlap}(G)$ measures motif-profile agreement and $J_{\text{top}}(G)$ denotes top-shift Jaccard similarity. Larger values indicate stronger structural agreement.

Figure 3 should be interpreted as a frontier view rather than as a replacement for Table 2. structure_oracle_transfer attains much stronger structural agreement than the practical transfer baselines, but it does so at the cost of substantially worse exact clique

burden. By contrast, portfolio_transfer achieves a lower structural score while producing much stronger Ramsey-valid transfer. The main value of the figure is therefore to visualize the tradeoff already quantified in the tables.

The two additional methods are marked with triangular symbols to distinguish them from the original set. hierarchical_coarsening_transfer lies in the low-clique-count region of the frontier, consistent with its strong performance on the R(3,*s*) cells reported in Table 2. flag_density_transfer occupies a position closer to the random baselines on the structural axis, reflecting its limited motif alignment beyond the triangle-closing construction heuristic. Neither method displaces the existing frontier positions of portfolio_transfer or structure_oracle_transfer.

5.3. Matched-Compute Comparison

As can be seen in Table 3, the advantage of portfolio_transfer is not simply the result of using more search effort.

Table 3. Matched-compute comparison.

Method	Overall Mean Rank	Exact r-Clique Count	Density Error	Motif Overlap	Edge Disagreement Rate	Top-Shift Jaccard	Validity Proxy
portfolio_transfer	3.471	2602.8	0.0106	0.0684	0.3635	0.1300	0.7000
exact_triangle_transfer	3.957	6692.6	0.0135	0.0784	0.3681	0.0552	0.6000
random_circulant	4.129	9724.2	0.0048	0.0613	0.3828	0.1319	0.5000
random_density	4.171	8750.8	0.0055	0.1016	0.3688	0.0990	0.5000
structured_seed	4.243	113,631.4	0.0025	0.0586	0.3907	0.1371	0.5000

Table 3 compares the transfer baselines under a matched evaluation budget. Lower values are preferable for exact clique burden and density control, while higher values indicate stronger validity and structural agreement. Table 3 reinforces the main result from Table 2. Even when the compute budget is controlled, *portfolio_transfer* remains first overall, and its exact clique count is substantially lower than those of the simpler baselines. This supports the claim that the framework is learning transferable structural information, not merely benefiting from extra search time.

5.4. Search Performance

As can be seen in Table 4, the search-side results are also strong, although the differences are narrower than in the transfer suite.

Table 4. Search-suite performance.

Method	Overall Mean Rank	Mean Search Best Score	Exact r-Clique Count	Density Error	Motif Overlap	Top-Shift Jaccard	Validity Proxy
portfolio_guided_search	2.538	2.4818	2193.0	0.0130	0.0642	0.1319	0.7000
structured_seed_local_search	2.625	2.4901	1157.4	0.0115	0.0653	0.1804	0.7000
random_local_search	2.650	2.1946	2107.8	0.0076	0.0737	0.1319	0.7000
psi_ramsey_guided_search	3.350	3.4221	27,729.4	0.0072	0.0596	0.0728	0.6000
partition_guided_search	3.800	3.8931	158,646.2	0.0675	0.0743	0.2390	0.5999

The best search method in the synchronized runs is portfolio_guided_search, which performs best overall by mean rank. Table 4 compares the main search-time baselines. Lower values indicate a better refinement score and lower exact clique burden, whereas

higher validity proxy values indicate stronger partial search success. As defined in the evaluation protocol, the search comparison is based on the best objective value reached during refinement. Lower values indicate better refinement outcomes.

Table 4 shows that no single search baseline dominates every metric. For example, structured_seed_local_search attains a lower exact clique count than portfolio_guided_search, while random_local_search attains a lower mean search score. Nevertheless, portfolio_guided_search is the strongest balanced method overall. This matches the transfer results: the main strength of the proposed framework is not one isolated heuristic, but a portfolio that combines multiple structural priors and refinement paths.

5.5. Search Tradeoff Figure

As can be seen in Figure 4, the search methods also lie on a tradeoff surface, though less sharply than in the transfer setting.

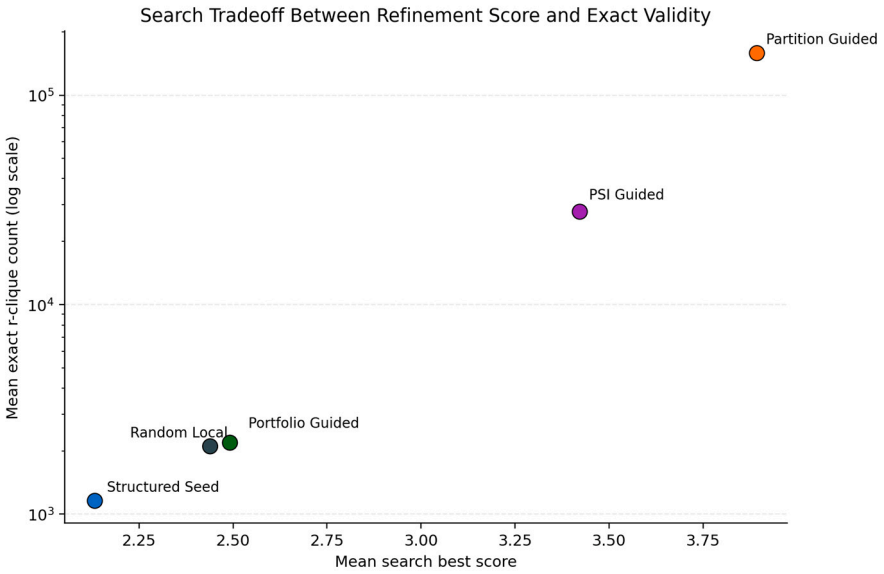


Figure 4. Comparison of search-time methods by mean search best score and mean exact **r**-clique count. Lower values are better on both axes. A logarithmic vertical axis is used to make the contrast between the main search methods and the partition-based outlier readable in a single panel. The figure shows that portfolio_guided_search provides the strongest balanced search position among the compared methods.

The figure compares refinement quality and exact validity in a single panel. Because one method is a strong outlier on exact clique count, the vertical axis is shown on a logarithmic scale so that the remaining methods remain visually comparable.

Figure 4 should be read as a search tradeoff view rather than as a complete ranking by itself. Some methods are stronger on one axis than another. Overall, *portfolio_guided_search* occupies the most balanced position when refinement quality and exact clique burden are considered jointly. This visual interpretation is consistent with the aggregate ranking in Table 4.

5.6. Ablation Study

As can be seen in Table 5, the full performance of the proposed framework does not arise from structure distillation alone.

Table 5 compares reduced or simplified variants of the framework. Lower overall rank is better, and the table is intended to show how much of the final behavior can be recovered when key components are removed.

Table 5. Ablation study.

Method	Overall Mean Rank	Exact r-Clique Count	Density Error	Motif Overlap	Edge Disagreement Rate	Top-Shift Jaccard	Validity Proxy
structured_seed	2.229	113,631.4	0.0025	0.0586	0.3907	0.1371	0.5000
teacher_mean_transfer	2.514	98,941.2	0.0260	0.0579	0.3916	0.1166	0.5000
psi_ramsey_transfer	2.629	96,503.6	0.0260	0.0606	0.3863	0.1166	0.5000
student_no_filter_transfer	2.629	96,503.6	0.0260	0.0606	0.3863	0.1166	0.5000

Table 5 is important because it prevents an overly simple interpretation of the method. The distilled student representation by itself is not enough to explain the final gains. Instead, the results suggest that the strongest behavior comes from the complete pipeline, especially when portfolio selection and exact-supervised small-cell refinement are included.

5.7. Interpretability and Structural Frontier

As can be seen in Table 6, the interpretability suite makes the distinction between structural similarity and Ramsey-valid construction especially clear.

Table 6. Structural frontier analysis.

Method	Overall Mean Rank	Exact r-Clique Count	Density Error	Motif Overlap	Edge Disagreement Rate	Top-Shift Jaccard	Validity Proxy
portfolio_transfer	3.886	2839.4	0.0091	0.0666	0.3687	0.1371	0.7000
structure_oracle_transfer	4.100	8361.2	0.0608	0.1926	0.1227	0.8400	0.6000
exact_triangle_transfer	4.629	6692.6	0.0135	0.0784	0.3681	0.0552	0.6000
random_circulant	4.786	7839.8	0.0025	0.0632	0.3813	0.0990	0.5000
structured_seed	5.143	113,631.4	0.0025	0.0586	0.3907	0.1371	0.5000

The table compares the balanced transfer winner with structure-dominant and simpler alternatives. A lower exact clique count and lower density error indicate stronger validity-oriented performance. Higher motif overlap and stronger top-shift agreement indicate closer structural similarity to known witnesses.

This comparison is useful because it shows that different evaluation criteria favor different methods. If only structure-oriented metrics are considered, the oracle baseline appears strongest. If only clique burden is emphasized, much of the structural information is lost. Considering both perspectives makes clear that witness-shape recovery and balanced transfer performance are not identical objectives.

5.8. Per-Target Transfer Analysis

As can be seen in Figure 5, the aggregate transfer story is not driven by a single target cell.

The figure shows per-target exact clique counts for the main transfer methods and makes it easier to see where each method succeeds or fails. Unlike Table 2, this figure focuses on one validity dimension only, namely, exact clique burden, and should therefore be interpreted as a complement to the transfer tables rather than as a full summary on its own.

Figure 5 is useful because it reveals how the methods behave across the five target cells instead of only through aggregate means. In particular, it shows that the portfolio method is competitive across the full target set rather than winning through a single anomalous

case. It also makes clear that structure-oriented transfer and exact-supervised transfer behave differently across the R(3,*s*) and R(4,*s*) regimes.

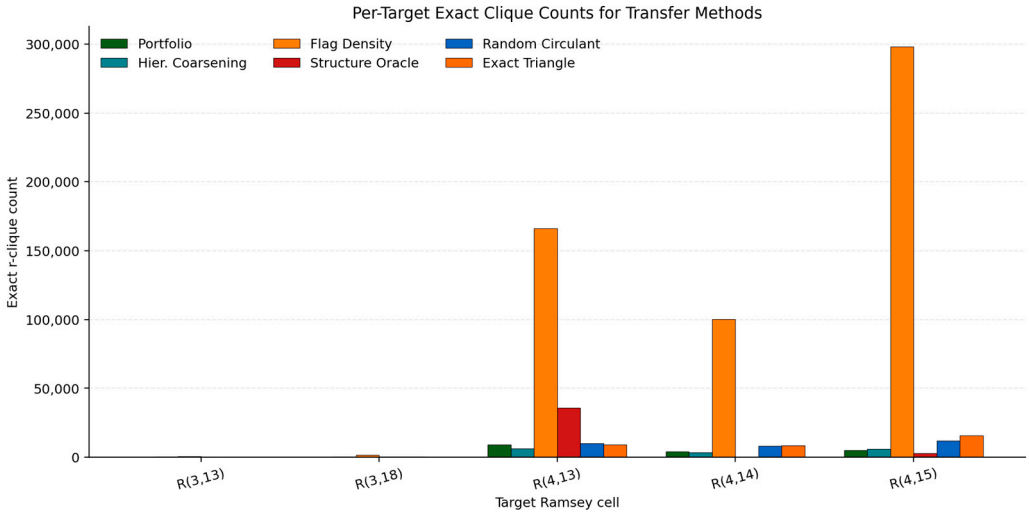


Figure 5. Per-target comparison of exact r-clique counts across the main transfer baselines. Lower values are better. The figure highlights the consistency of portfolio_transfer across the target set and shows where structure-dominant and simpler baselines succeed or fail. hierarchical_coarsening_transfer and flag_density_transfer are included to show how triangle-density priors behave across the R(3,*s*) and R(4,*s*) regimes.

As can be seen in Figure 5, the per-target contrast between hierarchical_coarsening_transfer and flag_density_transfer is especially clear. hierarchical_coarsening_transfer achieves near-zero exact clique counts on the R(3,s) target cells but is less competitive on the R(4,s) targets. flag_density_transfer follows a similar pattern, with stronger penalties on the R(4,*s*) cells. This per-target view confirms that the R(3,*s*) and R(4,*s*) cells pose qualitatively different structural challenges and that methods grounded in triangle-density priors are naturally better matched to the R(3,*s*) regime.

5.9. Overall Interpretation

Taken together, Tables 2–6 and Figures 2–4 support two conclusions. First, the proposed framework is effective as a balanced transfer-and-search framework, with portfolio_transfer and portfolio_guided_search emerging as the strongest practical methods in the synchronized experiments. Second, the results reveal a clear structure-validity frontier: methods that best reproduce witness-like structure are not necessarily those that best construct Ramsey-valid candidates. The results for hierarchical_coarsening_transfer and flag_density_transfer further confirm that the R(3,*s*) regime is amenable to validity-oriented transfer through shift-scaling and triangle-closing heuristics, while the R(4,*s*) regime remains more demanding and resists these simpler structural priors. The two new methods occupy well-separated positions within the frontier space, confirming that the boundary between structural dominance and balanced transfer performance is stable across a broader set of candidates.

This distinction is important for future work. It suggests that progress in transferable Ramsey search should not be evaluated only by structural resemblance or local heuristic strength. The key criterion is how well transferred structural priors survive the validity constraints of unseen target cells.

A broader supplementary evaluation was conducted to test the stability of these conclusions under repeated seeds and perturbed experimental settings. The Supplementary

Materials document largely reinforces the transfer-side findings while also showing that the search-side ordering is more sensitive to configuration and initialization.

6. Discussion

The results of this study support a nuanced but meaningful conclusion. The proposed framework succeeds as a proof-of-concept framework for balanced transferable Ramsey search, but it does not reduce the full problem to a single universal structural prior. Instead, the experiments show that transfer in Ramsey graph construction is governed by a genuine tradeoff between witness-shape agreement and Ramsey-valid construction. This tradeoff is not an artifact of one metric or one suite. It appears consistently across the transfer, matched-compute, interpretability, and search evaluations.

The main practical result is that the portfolio version of the proposed framework is the strongest balanced transfer method in the current experimental setting. Its success does not come from perfectly reproducing known witnesses. Rather, it comes from combining complementary structural priors, refining them locally, and selecting among them under a balanced objective. This is an important methodological point. In Ramsey search, a single distilled representation is often too narrow to support robust transfer across cells. A portfolio of structural candidates is more effective because different target cells appear to benefit from different forms of transferred structure.

At the same time, the experiments show clearly that structural similarity alone is not an adequate proxy for successful transfer. The structure-dominant oracle baseline achieves substantially stronger motif overlap, edge agreement, and top-shift alignment than the balanced transfer winner, yet it performs worse on exact clique burden and overall Ramsey-valid construction. This result is scientifically useful because it exposes a frontier that might otherwise remain hidden. A method can appear highly successful if judged only by similarity to known witnesses while still being less useful as a constructive search tool. For this reason, the present work argues that balanced evaluation is essential in transferable Ramsey search.

The behavior of the exact-supervised branch for weak R(3,s) cells provides a second important lesson. Small-cell supervision improves the difficult R(3,s) regime, especially where purely structural transfer is weakest. This suggests that exact local supervision can play a valuable role when the combinatorial subproblem is still computationally manageable. More broadly, it indicates that transferable structure and exact validity information should not be treated as competing paradigms. In practice, the strongest results arise when they are combined.

The search experiments reinforce this interpretation. The strongest search behavior again comes from a balanced portfolio rather than from a single heuristic family. This suggests that the value of the proposed framework lies less in proposing one privileged construction rule and more in organizing multiple structural hypotheses into a reusable transfer-and-refinement pipeline. In that sense, the framework should be understood not as a final Ramsey solver, but as a methodology for studying how structural information can be distilled, transferred, and evaluated across related cells.

The results for the proportional shift-scaling and greedy triangle-closing baselines further reinforce this interpretation. Both methods are effective on R(3,s) targets but considerably weaker on R(4,s) targets, confirming that the R(3,s) and R(4,s) regimes pose qualitatively different structural challenges. This split is consistent across all evaluated methods and suggests that the structure-validity frontier is not uniform across the Ramsey table. Methods grounded in triangle-density priors naturally align with the triangle-avoidance requirements of R(3,s) witnesses but do not generalize to the denser combinatorial constraints of the R(4,s) regime.

From a theoretical perspective, the proposed framework is heuristic in nature, and several of its components do not admit formal guarantees in the classical sense. Regarding convergence of the refinement process: the local search refinement used in this framework is a form of stochastic hill-climbing over a combinatorial objective. Such procedures do not guarantee convergence to a global optimum, but they are known to converge almost surely to a local optimum under mild conditions on the objective landscape. In practice, the refinement is run for a fixed horizon and the best candidate is selected, which is standard in Ramsey search. Regarding validity preservation under transfer: there is no formal guarantee that a candidate reconstructed from a transferred structural prior will be Ramsey-valid. The repair mechanism reduces exact clique burden iteratively, but validity is not guaranteed after a fixed number of repair steps. This is precisely why the evaluation framework separates exact validity metrics from structural similarity metrics and why the validity proxy is reported alongside exact clique counts. Regarding probabilistic guarantees of finding better witnesses: providing such guarantees would require lower-bounding the probability that a transferred structural prior lands in the basin of attraction of a valid witness, which depends on the unknown geometry of the Ramsey-validity landscape. Rather than making unsubstantiated claims, the present work provides empirical evidence that portfolio-based transfer consistently reduces exact clique burden relative to random and structure-only baselines and argues that this constitutes meaningful, if not formally guaranteed, improvement in search quality.

The present work has several limitations. To strengthen the empirical interpretation despite these limitations, additional robustness and sensitivity experiments are provided in the Supplementary Materials document. These supplementary results help distinguish conclusions that remain stable across repeated and perturbed settings from those that are more configuration-sensitive. First, the experimental study is based on a small collection of published witnesses and a correspondingly small number of target cells [6,13]. This is appropriate for a controlled proof-of-concept study, but it does not establish broad generalization across large regions of the Ramsey table. Second, the current structural representations are still relatively coarse. Even with shift profiles, partition-level summaries, and motif-aware transfer, the framework does not yet capture the full richness of possible automorphism structure, orbit organization, or higher-order combinatorial constraints. Third, the framework is currently restricted to circulant graphs; extending the structural representation to non-circulant graph families would require fundamental changes to the shift-profile teacher and reconstruction mechanisms and remains a direction for future work. Fourth, some of the most informative analysis tools, especially the structure-dominant oracle, are not deployable transfer methods. They are valuable for understanding the structure-validity frontier, but they should not be interpreted as practical construction algorithms.

These limitations point directly to future work. One natural direction is to expand the construction family beyond the current circulant, profile-based, and partition-based representations. Orbit-based or group-structured representations may offer a richer and more faithful way to encode reusable combinatorial organization. A second direction is to enlarge the teacher pool and study transfer over broader neighborhoods of Ramsey cells, including cells with more diverse local structure. A third direction is to integrate stronger exact supervision where feasible, especially in settings where exact local constraints can be used to shape the transfer objective more directly. Finally, it would be valuable to learn the portfolio mechanism more explicitly. This would allow the system to not only transfer structural priors, but also to predict when each prior is likely to be useful.

More broadly, the present work suggests a shift in how automated Ramsey search can be studied. Rather than viewing each lower-bound improvement as an isolated success, one can ask what reusable structural information the search process has uncovered and

whether it transfers to new cells. This reframes automated Ramsey search as a problem of structure discovery as much as one of optimization. In that sense, the contribution is not only algorithmic but methodological. The framework separates local heuristic strength, structural similarity, and genuine transferable construction value.

In summary, the experiments show that transferable Ramsey search is possible in a meaningful but limited sense. The proposed framework provides the strongest balanced transfer framework in the current study, while also revealing a clear frontier between structural resemblance and Ramsey-valid construction. This frontier is not a weakness of the paper; it is one of its main findings. Future progress will depend not only on stronger local heuristics, but on better ways of learning, compressing, and reusing structural organization shared across extremal graph constructions.

7. Conclusions

This paper introduced a Ψ-NN-inspired framework for transferable structural search in Ramsey graph construction. The central objective was to move beyond cell-specific heuristic discovery and investigate whether structural information extracted from known Ramsey witnesses can be transferred across related cells. To this end, we developed a teacher–student structure-distillation pipeline combined with reconstruction and refinement mechanisms. The framework was evaluated through transfer, matched-compute, search, ablation, and interpretability experiments.

The results show that the proposed framework is effective as a proof-of-concept balanced transfer-and-search framework. In particular, the portfolio version of the method emerges as the strongest practical transfer method in the current study, while portfolio-guided search performs best overall on the search side. At the same time, the experiments reveal a clear structure-validity frontier. A structure-dominant oracle baseline achieves substantially stronger witness-shape agreement but performs worse on Ramsey-valid construction. This demonstrates that structural resemblance to known witnesses and successful transferable construction are related but distinct objectives.

The evaluation also includes a proportional shift-scaling baseline and a greedy triangle-closing baseline, both of which are effective on R(3,s) targets but considerably weaker on R(4,s) targets. This confirms that the structure-validity frontier is not uniform across the Ramsey table and that the R(3,s) and R(4,s) regimes pose qualitatively different structural challenges. The stability of this finding across all evaluated methods strengthens the main empirical message of the paper.

The paper therefore contributes both a method and an evaluation perspective. Methodologically, it adapts the teacher–student structure-distillation logic of Ψ-NNs to extremal combinatorial search. Empirically, it shows that transferable structural priors can improve Ramsey search under balanced criteria. At the same time, structure-only agreement is not a sufficient measure of success. More broadly, progress in automated Ramsey search should be assessed not only by the strength of local heuristics, but by how well structural organization can be learned, reused, and made robust across neighboring cells.

The proposed framework is not the final answer to transferable Ramsey search, but it provides a concrete step toward that goal. The present results suggest that future progress will depend on richer structural representations, stronger exact supervision where feasible, and broader studies of transfer across the Ramsey table. Even so, the current proof-of-concept study already shows that automated Ramsey search can be investigated not only as optimization, but as structure discovery. That shift in perspective is, in our view, one of the main contributions of this paper.

https://doi.org/10.3390/math14081367

The complete experimental suite and code are available in our GitHub repository (https://github.com/jurjsorinliviu/Transferable-Structural-Search-for-Ramsey-Graph-Construction) (accessed on 16 April 2026), release v1.0.

Supplementary Materials: The following supporting information can be downloaded at https://www.mdpi.com/article/10.3390/math14081367/s1, Table S1. Transfer Seed Robustness; Table S2. Search Seed Robustness; Table S3. Transfer Sensitivity Across Supplementary Scenarios; Table S4. Search Budget Sensitivity; Table S5. Ablation Seed Robustness; Table S6. Interpretability Seed Robustness; Table S7. New Methods Transfer Robustness; Table S8. Cross-r Transfer Comparison; Table S9. Single-Teacher Transfer Comparison; Figure S1. Distribution of Overall Mean Ranks across Eight Replicated Transfer Runs; Figure S2. Distribution of Overall Mean Ranks across Eight Replicated Search Runs; Figure S3. Transfer Sensitivity Across Supplementary Scenarios; Figure S4. Portfolio Transfer Sensitivity Across Supplementary Scenarios; Figure S5. New Methods Transfer Robustness (Seed Variation); Figure S6. Validity Proxy across Transfer Pool Configurations; Figure S7. Scaling Behavior of Exact r-Clique Count with Graph Size n.

Funding: This research received no external funding.

Data Availability Statement: The data presented in this study are openly available in [GitHub] at [https://github.com/jurjsorinliviu/Transferable-Structural-Search-for-Ramsey-Graph-Construction] (accessed on 16 April 2026).

Conflicts of Interest: The author declares no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

CSVs	Comma-Separated Values
JSON	JavaScript Object Notation
LLM	Large Language Model
MAE	Mean Absolute Error
PINN	Physics-Informed Neural Network
Ψ-NN	Physics Structure-Informed Neural Network

Nomenclature

The following graph and method symbols are used in this manuscript:

Symbol	**Definition**
A	Adjacency matrix of a graph
A_{ij}	Entry of the adjacency matrix indicating whether vertices i and j are adjacent
a	Local search action
b	Normalized structural bin
B_b	Set of shifts assigned to normalized bin b
$c_G(b)$	Contrastive structural profile of graph G in bin b
d_G	Mean degree of graph G
$\bar{d}$	Transferred target mean degree
$\deg(i)$	Degree of vertex i
$\Delta(G)$	Exact number of triangles in graph G
$\Delta_r(G)$	Exact number of forbidden r-cliques in graph G
δ	Cyclic shift index
E	Edge set of a graph
$G = (V, E)$	Graph with vertex set V and edge set E
$G^\star$	Balanced selected candidate graph
$G^\dagger$	Structure-dominant selected candidate graph
$\mathcal{G}_{\text{cand}}$	Portfolio of candidate graphs

$I(b)$	Structural importance score of bin b
M_{uv}	Block-density template entry between partition blocks u and v
n	Number of vertices in a graph
$p_G(\delta)$	Cyclic shift profile of graph G at shift δ
$\widetilde{p}_G(b)$	Normalized shift-profile value of graph G in bin b
$\mathcal{P}$	Set of extracted compatible bin pairs
$q_G(b_1, b_2)$	Pairwise compatibility score between bins b_1 and b_2
$\bar{q}(b_1, b_2)$	Mean pairwise compatibility between bins b_1 and b_2 across teacher witnesses
$R(r, s)$	Two-color Ramsey number for clique size r and independent-set size s
r	Forbidden clique size in a Ramsey cell
s	Forbidden independent-set size in a Ramsey cell
$\mathcal{S}$	Set of extracted shared structural bins
$\mathrm{supp}(b)$	Support of structural bin b across teacher witnesses
T_G	Teacher representation of graph G
V	Vertex set of a graph
$\mathcal{Z}$	Selected set of active target shifts
$\alpha(G)$	Independence number of graph G
$\beta(a)$	Operator bias for local action a
λ_1, λ_2	Weights used in the bin-importance score
$\omega(G)$	Clique number of graph G
ρ_G	Density of graph G
$\bar{\rho}$	Transferred target density
τ_q	Compatibility threshold for pairwise structural extraction
τ_s	Support threshold for structural extraction

Glossary

The following evaluation symbols are used in this manuscript:

Symbol	**Definition**
$D_{\mathrm{pair}}(G)$	Pairwise compatibility mismatch penalty
$D_{\mathrm{struct}}(G)$	Structural mismatch term for graph G
$E_d(G)$	Degree mean absolute error
$E_\rho(G)$	Density error
$J_{\mathrm{top}}(G)$	Top-shift Jaccard similarity
$\mathcal{L}(G)$	Balanced refinement objective for graph G
$\mathcal{L}_{\min}$	Best refinement score reached by method m
$\mathrm{MotifOverlap}(G)$	Motif overlap score
$R_{m,k}$	Rank of method m on metric k
$\bar{R}_m$	Mean rank of method m across metrics
$S_{\mathrm{struct}}(G)$	Structural agreement score
$w_1, \ldots, w_5$	Weights in the balanced refinement objective
$\hat{C}_r(G)$	Sampled r-clique violation rate
$\hat{I}_s(G)$	Sampled independent-set violation rate
η_1, η_2	Weights in the structural mismatch term

References

1. Ramsey, F.P. On a Problem of Formal Logic. *Proc. Lond. Math. Soc.* **1930**, *s2-30*, 264–286. [CrossRef]
2. Radziszowski, S. Small Ramsey Numbers. *Electron. J. Comb.* **2026**, *1000*, DS1:1–DS1:148. [CrossRef]
3. Conlon, D.; Ferber, A. Lower bounds for multicolor Ramsey numbers. *Adv. Math.* **2021**, *378*, 107528. [CrossRef]
4. Piwakowski, K. Applying Tabu Search to Determine New Ramsey Graphs. *Electron. J. Comb.* **1996**, *3*, R6. [CrossRef]
5. Exoo, G.; Tatarevic, M. New Lower Bounds for 28 Classical Ramsey Numbers. *Electron. J. Comb.* **2015**, *22*, P3.11. [CrossRef]
6. Nagda, A.; Raghavan, P.; Thakurta, A. Reinforced Generation of Combinatorial Structures: Ramsey Numbers. *arXiv* **2026**. [CrossRef]

7. Goedgebeur, J.; Van Overberghe, S. New bounds for Ramsey numbers $R(K_k - e, K_l - e)$. *Discret. Appl. Math.* **2022**, *307*, 212–221. [CrossRef]
8. Liu, Z.; Liu, Y.; Yan, X.; Liu, W.; Nie, H.; Guo, S.; Zhang, C.-A. Automatic network structure discovery of physics informed neural networks via knowledge distillation. *Nat. Commun.* **2025**, *16*, 9558. [CrossRef] [PubMed]
9. MRaissi, M.; Perdikaris, P.; Karniadakis, G.E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **2019**, *378*, 686–707. [CrossRef]
10. Wu, K.; Su, W.; Luo, H.; Xu, X. A Generalization of Generalized Paley Graphs and New Lower Bounds for $R(3,q)$. *Electron. J. Comb.* **2010**, *17*, N25. [CrossRef] [PubMed]
11. Parczyk, O.; Pokutta, S.; Spiegel, C.; Szabó, T. New Ramsey Multiplicity Bounds and Search Heuristics. *Found. Comput. Math.* **2025**, *25*, 1777–1814. [CrossRef]
12. Mattheus, S.; Verstraete, J. The asymptotics of r(4,t). *arXiv* **2024**, arXiv:2306.04007. [CrossRef]
13. Google Research Ramsey Bounds Repository. Available online: https://github.com/google-research/google-research/tree/master/ramsey_number_bounds/improved_bounds (accessed on 19 March 2026).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.