APPROXIMATION ALGORITHMS FOR THE FACILITY LOCATION PROBLEM: A SURVEY

Yingxi Lu (2023011435)
Institute for Interdisciplinary Information Sciences
Tsinghua University
lu-yx23@mails.tsinghua.edu.cn

ABSTRACT

The Facility Location Problem (FLP) is a well-established optimization problem with diverse applications, including logistics, supply chain management, telecommunications, and data center placement. This survey focuses on the classical uncapacitated Facility Location Problem (UFLP) and extends the discussion to several key variants, such as capacitated facility location, dynamic and online facility location, and multi-objective facility location. The historical development and evolution of the FLP are reviewed, emphasizing significant contributions and advancements in approximation algorithms and solution methodologies. For each variant, the survey explores the core challenges and the innovative approaches developed to address them. Additionally, to gain deeper insights into the approximation algorithms for the UFLP and their performance, this work includes implementations of these algorithms, along with a comparative analysis of their results. The objective of this survey is to provide a comprehensive understanding of the Facility Location Problem and its variants, highlighting recent research trends and practical applications.

1 INTRODUCTION

Facility Location Problems (FLPs) encompass a class of optimization problems that aim to determine the optimal placement of facilities to serve a given set of demands, minimizing costs while satisfying problem-specific constraints. As a cornerstone of combinatorial optimization, the FLP has been extensively studied due to its theoretical complexity and wide-ranging applications in industries such as transportation, telecommunications, and supply chain management [8, 15].

The classical uncapacitated facility location problem (UFLP) serves as the foundation for this field, with its NP-hardness prompting early efforts to develop heuristic and approximation algorithms. Over time, significant algorithmic advancements have emerged, such as linear programming (LP) relaxations with rounding techniques, primal-dual approaches, and local search heuristics. These methods have provided bounded approximation guarantees, making them powerful tools for tackling computationally hard problems [14, 16]. Beyond the classical UFLP, various extensions of FLPs have been proposed to model real-world complexities, including capacity constraints, dynamic environments, online decision-making, and multi-objective considerations [10, 6].

The motivation for this survey arises from the growing complexity of FLPs and the diversity of solution approaches developed over the years. By providing a structured analysis of approximation algorithms for FLPs and their key variants, this work aims to bridge the gap between theoretical advancements and practical implementations. The survey seeks to:

- Summarize the historical progression of approximation algorithms for the classical FLP.
- Discuss the specific challenges and solution techniques for key FLP variants, including capacitated, dynamic, and multi-objective problems.

- Implement the approximation algorithms for UFLP and analyze their performance through empirical comparisons.
- Identify emerging research trends and highlight open questions in the field.

The remainder of this survey is organized as follows. After reviewing approximation algorithms for the classical UFLP, we delve into the challenges and solutions for capacitated FLPs, dynamic and online FLPs, and multi-objective FLPs in subsequent sections. Following this, we present a detailed section on the implementation of UFLP approximation algorithms, providing practical insights into their computational performance and a comparative analysis of their results. Finally, we conclude with a discussion of key findings and potential directions for future research.

2 PRELIMINARIES

2.1 BASIC FORMULATION

The basic facility location problem involves a set of clients D and a set of falicities F. For each pair of client $i \in D$ and facility $j \in F$, a cost function c_{ij} represents the cost associated with serving client i by facility j . Additionally, each facility j has an associated opening cost f_j. The objective is to select a subset of facilities F' that minimizes the total cost, which includes both the opening costs of the selected facilities and the transportation costs incurred by assigning clients to their nearest facility[19]. The objective function for this problem can be formulated as:

$$\min_{F'} \left(\sum_{j \in F'} f_j + \sum_{i \in D} \min_{j \in F'} c_{ij} \right).$$

2.2 NP-HARDNESS OF FLP.

To prove that FLP is NP-Hard, we reduce Set Cover Problem (SCP) to FLP in polynomial time. The reduction consists of transforming any instance of SCP into a corresponding instance of FLP such that solving the FLP instance provides the solution to the SCP instance.
Given an instance of SCP, let U be the set of elements in SCP. We construct an equivalent instance of FLP as follows:

- For each subset $S_j \in S$, create a facility $j \in F$ in FLP; For each element $u \in U$, create a client $i \in D$ in FLP.
- Set the opening cost of each facility $j \in F$ in FLP to be $f_j = c_j$, the cost of the corresponding subset S_j in SCP.
- Define the transportation cost c_{ij} between client i (representing element u) and facility j (representing subset S_j) as:

  $$c_{ij} = \begin{cases} 0, & if\ u \in S_j, \\ M, & otherwise, \end{cases}$$

  where M is very large constant (e.g., $M \gg \sum_{j \in F} f_j$).
- The goal is to select a subset of facilities F' (corresponding to subsets in SCP) that minimizes both the facility opening costs and transportation costs.

To show that the reduction is correct, we first proof that: If the SCP instance has an optimal solution S', then the FLP instance has an optimal solution F' with the same cost. In FLP, selecting the facilities F' correspoding to S' minimized the facility opening cost $\sum_{j \in F'} f_j$. Additionally, for each client i, there exists at least one facility

$j \in F'$ with $c_{ij} = 0$, since S' covers U. Thus, the transportation cost $\sum_{i \in D} \min_{j \in F'} c_{ij} = 0$. The total cost in FLP is:

$$\sum_{j \in F'} f_j + \sum_{i \in D} \min_{j \in F'} c_{ij} = \sum_{S_j \in S'} c_j.$$

Next, we show that if the FLP instance has an optimal solution F', then the SCP instance has an optimal solution S' with the same cost. In FLP, if $\sum_{i \in D} \min_{j \in F'} c_{ij} = 0$, then every client i is assign to a facility j with $c_{ij} = 0$, meaning that, for every element $u \in U$, there exists a subset $S_j \in S'$ such that $u \in S_j$. Thus, $S' = \{S_j : j \in F'\}$ is a valid cover of U in SCP.
Since it is obvious that the reduction is polynomial in time, and SCP is NP-Hard, the FLP is NP-Hard.

2.3 The history of approximation solutions for FLP

The history of approximation algorithms for FLPs highlights significant progress in algorithm design and analysis for solving computationally hard problems. FLPs are well-known to be NP-hard, meaning finding exact solutions efficiently for large instances is infeasible. Approximation algorithms provide a way to find near-optimal solutions with guaranteed performance bounds.

Early Work: Approximation Paradigms in FLPs. The origins of approximation algorithms for FLPs can be traced back to the late 1970s and early 1980s, when researchers primarily relied on heuristic methods and early formulations of linear programming (LP) relaxations. A natural starting point was the use of greedy algorithms, which iteratively selected facilities to minimize incremental costs. Cornuéjols et al. (1977), for instance, analyzed heuristic approaches for the uncapacitated facility location problem (UFLP). While these methods were practical, they lacked formal approximation guarantees, creating a gap between theoretical optimality and real-world applicability[4].

The introduction of LP relaxation and rounding techniques marked a turning point. By solving the fractional versions of relaxed LPs and rounding the solutions to produce feasible results, researchers achieved bounded approximation ratios, bridging the divide between heuristics and theoretical rigor[18].

Pioneering Approximation Ratios. As the field matured, researchers turned their focus toward improving approximation ratios, setting benchmarks for algorithmic performance. In a seminal work, Shmoys et al. proposed a 3.16-approximation algorithm for the UFLP, combining LP relaxation with an innovative rounding method. This study underscored the utility of LP duality in simplifying problem formulations and generating efficient solutions. It also established LP-based techniques as a cornerstone of combinatorial optimization in location theory[18].

Concurrently, refinements to greedy methods improved decision-making by considering trade-offs between facility opening costs and connection costs. These methods, though intuitive, underwent rigorous analysis to provide bounded approximation ratios, enhancing both their practical utility and theoretical foundations[13].

Refinements: Primal-Dual Techniques. The late 1990s heralded a new era with the advent of primal-dual algorithms, which represented a significant leap forward in approximation algorithms for FLPs. This approach operates on both the primal and dual formulations of an LP, maintaining dual feasibility while iteratively constructing a feasible primal solution. This dual perspective allows for more precise control over cost trade-offs during algorithm execution.

A landmark contribution came from Jain and Vazirani, who introduced a 2-approximation algorithm for the UFLP using the primal-dual schema. Their work highlighted the simplicity and effectiveness of this method, offering both strong theoretical guarantees and practical ease of implementation. This breakthrough became a foundational element of modern approximation algorithms for location problems[12].

Strengthened LPs and Combinatorial Methods. The early 2000s saw the emergence of enhanced LP formulations and combinatorial techniques aimed at overcoming limitations of earlier approaches. Researchers began incorporating valid inequalities and additional constraints into LP relaxations, resulting in tighter formulations. These advancements proved particularly effective for complex variants such as the capacitated facility location problem (CFLP), where capacity constraints posed additional challenges[2].

Simultaneously, local search heuristics gained popularity due to their simplicity and efficiency. By iteratively refining solutions through local swaps—such as opening or closing facilities—these methods offered high-quality solutions for CFLP and related variants, further broadening the applicability of approximation algorithms.

Robustness and Randomization. In recent years, the scope of approximation algorithms has expanded to address uncertainties in FLPs, giving rise to robust and randomized approaches. Robust optimization techniques, designed to perform well under worst-case scenarios, have been applied to FLPs involving uncertain inputs such as demand and costs. These methods have found practical applications in domains like disaster management and supply chain resilience[17].

Randomized techniques have also gained traction, leveraging probabilistic methods to handle dynamic and uncertain problem parameters more effectively. Together, these approaches reflect the growing emphasis on creating algorithms that are both theoretically robust and practically relevant.

The evolution of approximation algorithms for FLPs exemplifies a journey from heuristic methods to advanced LP-based, primal-dual, and robust techniques, continuously balancing efficiency, theoretical guarantees, and real-world applicability in addressing the challenges of dynamic and uncertain environments.

3 Approximation Algorithms for Ucapacitated Facility Location Problem

This section explores effective approximation algorithms for addressing the Uncapacitated Facility Location Problem (UFLP), which is a fundamental case of the Facility Location Problem (FLP). The focus will be on algorithms that achieve favorable or acceptable approximation ratios. Specifically, I will present detailed analyses of five prominent methods: **deterministic rounding of LP**, **randomized rounding of LP**, the **primal-dual method**, and the **local search heuristics**.

For each algorithm, I will provide an overview, discuss its performance characteristics, and state its approximation ratio. To offer deeper insights, I will delve further into the **deterministic rounding of LP** and the **randomized rounding of LP**, presenting pseudocode implementations and rigorous proofs for their approximation ratios. This section aims to highlight both the theoretical significance and practical utility of these algorithms in solving the UFLP efficiently.

3.1 Deterministic Rounding for LP

The first approach to solving the Uncapacitated Facility Location Problem (UFLP) involves deterministic rounding of the linear programming (LP) solution, which results in a 4-approximation algorithm. In subsequent algorithms, we will achieve improved guarantees by applying other advanced techniques.

Define decision variable $y_i \in \{0, 1\}$ for each facility $i \in F$; if i is selected in F', $y_i = 1$; otherwise, $y_i = 0$. Also define decision variable $x_{ij} \in \{0, 1\}$ for all facility $i \in F$ and client $j \in D$; if we assign client j to facility i, then $x_{ij} = 1$; otherwise, $x_{ij} = 0$. Then, the objective function is:

$$\min \left(\sum_{i \in F} f_i y_i + \sum_{i \in F, j \in D} c_{ij} x_{ij} \right).$$

Besides, we need to ensure that each client $j \in D$ is assigned to exactly one facility, which can be expressed as:

$$\sum_{i \in F} x_{ij} = 1, \quad \forall j \in D.$$

Furthermore, we also need to make sure that the facility j being assigned to is open, which can be expressed as:

$$x_{ij} \le y_i, \quad \forall i \in F, j \in D.$$

Thus, the optimization problem can be formulated as:

$$\min \left(\sum_{i \in F} f_i y_i + \sum_{i \in F, j \in D} c_{ij} x_{ij} \right),$$

subject to:

$$\begin{aligned} \sum_{i \in F} x_{ij} &= 1, \quad \forall j \in D, \\ x_{ij} &\le y_i, \quad \forall i \in F, j \in D, \\ x_{ij} &\in \{0,1\}, \quad \forall i \in F, j \in D, \\ y_i &\in \{0,1\}, \quad \forall i \in F. \end{aligned}$$

Now, let's consider the dual linear program corresponding to the linear programming relaxation of the above problem. Rather than deriving the dual mechanically, it is motivated as a natural lower bound.
First, we consider the simplest case, where we ignore the facility costs f_i entirely, i.e., set $f_i = 0$ for all $i \in F$. In this case, the optimal solution would open all facilities and assign each client j to the nearest facility. Let the client cost be defined as:

$$v_j = \min_{i \in F} c_{ij},$$

which is the minimum connection cost for client j. The lower bound is therefore:

$$\sum_{j \in D} v_j.$$

Next, when the facility costs f_i are included, they can be distributed among the clients served by the facility. Let w_{ij} represent the portion of f_i assigned to client j (where $w_{ij} \ge 0$). Then, the total cost for each client j, combining both the connection cost c_{ij} and the facility cost w_{ij}, is:

$$v_j = \min_{i \in F}(c_{ij} + w_{ij}).$$

To maximize the lower bound $\sum_{j \in D} v_j$, we can optimize the facility cost shares w_{ij}. Allowing v_j to take any value satisfying $v_j \le c_{ij} + w_{ij}$, we maximize $\sum_{j \in D} v_j$.
This optimization over v_j and w_{ij} leads to the following dual linear program:

$$\max \sum_{j \in D} v_j,$$

subject to:

$$\begin{aligned} \sum_{j \in D} w_{ij} &\le f_i, \quad \forall i \in F, \\ v_j - w_{ij} &\le c_{ij}, \quad \forall i \in F, j \in D, \\ w_{ij} &\ge 0, \quad \forall i \in F, j \in D. \end{aligned}$$

Let Z_{LP}^* be the optimal value of the primal LP relaxation and OPT be the actual optimal solution. Any feasible solution (v, w) to the dual program satisfies:

$$\sum_{j \in D} v_j \le Z_{LP}^* \le OPT,$$

which establishes the weak duality.
Let (x^*, y^*) be the solution to the original problem, and (v^*, w^*) be the solution to the dual, we know that

$$\begin{aligned} c_{ij} &= v_{ij}^* - w_{ij}^*, \\ w_{ij}^* &\ge 0. \end{aligned}$$

Thus, we can deduce that:

$$c_{ij} \leq v_j^*.$$

To analyze the facility cost, define set $N(j)$ as:

$$N(j) = \{i \in F : x_{ij}^* > 0\}, \quad \forall j \in D.$$

The algorithm opens the cheapest facility i_k in the $N(j_k)$ of the selected j_k, we have:

$$f_{i_k} \leq \sum_{i \in N(j_k)} f_i x_{ij_k}^* \leq \sum_{i \in F} f_i y_i^*.$$

Since the neighborhoods $N(j_k)$ form a partition of the facilities opened by the algorithm, the total facility cost is bounded by:

$$\sum_k f_{ik} \leq \sum_{i \in F} f_i y_i^* \leq OPT,$$

where the last inequality follows from weak duality.
To analyze the assignment cost, define $N^2(j)$ as:

$$N^2(j) = \{k \in D : i \in N(k), s.t. i \in N(j)\}.$$

For each selected client j_k, we showed earlier that:

$$c_{i_k j_k} \leq v_{j_k}^*.$$

Consider a client $l \in N^2(j_k)$, using the triangle inequality, we have:

$$c_{i_k l} \leq c_{i_k j_k} + c_{j_k,h} + c_{hl} \leq v_{j_k}^* + v_{j_k}^* + v_{j_k}^* = 3v_{j_k}^*$$

where h is the facility satisfies $h \in N(j_k)$ and $h \in N(l)$.
Thus, the total assignment cost satisfies:

$$Assignment\ Cost \leq \sum_k \left(v_{j_k}^* + \sum_{l \in N^2(j_k)} 3v_{j_k}^* \right).$$

Since each $l \in N^2(j_k)$ appears in one iteration, this simplifies to:

$$Assignment\ Cost \leq 3 \sum_{j \in D} v_j^* \leq 3 \cdot OPT.$$

Thus,

$$Total\ Cost \leq OPT + 3 \cdot OPT = 4 \cdot OPT.$$

The algorithm achieves an approximation ratio of 4. [12]
The sudocode for this algorithm is shown below:[19]

Algorithm 1: Deterministic rounding algorithm for the uncapacitated facility location problem.

Solve LP, get optimal primal solution (x^*, y^*) and dual solution (v^*, w^*);
Initialize $C \leftarrow D$, $k \leftarrow 0$;
while $C \neq \emptyset$ **do**
  $k \leftarrow k + 1$;
  Choose $j_k \in C$ that minimizes v_j^* over all $j \in C$;
  Choose $i_k \in N(j_k)$ to be the cheapest facility in $N(j_k)$;
  Assign j_k and all unassigned clients in $N^2(j_k)$ to i_k;
  Update $C \leftarrow C \setminus \{j_k\} \setminus N^2(j_k)$;

The deterministic rounding algorithm achieves a 4-approximation for the uncapacitated facility location problem. It efficiently balances facility opening and client assignment costs by leveraging LP relaxation and dual properties, ensuring simplicity and computational efficiency.

3.2 Randomized Rounding for LP

Next, we introduce the randomized rounding algorithm, a refined approach that improves upon the deterministic rounding method to achieve a tighter 3-approximation guarantee for the uncapacitated facility location problem. By leveraging the probabilistic structure of the LP solution, this algorithm eliminates the need for worst-case assumptions about facility costs, effectively distributing the assignment burden across multiple options. This modification not only strengthens the theoretical analysis but also highlights the power of randomized techniques in achieving better approximation ratios. [19]
The sudocode for this algorithm is shown below:[19]

Algorithm 2: Randomized rounding algorithm for the uncapacitated facility location problem.

Solve LP, get optimal primal solution (x^*, y^*) and dual solution (v^*, w^*);
$C \leftarrow D$, $k \leftarrow 0$;
while $C \neq \emptyset$ **do**
  $k \leftarrow k+1$;
  Choose $j_k \in C$ that minimizes $v_j^* + C_j^*$ over all $j \in C$;
  Choose $i_k \in N(j_k)$ according to the probability distribution x_{jk}^*;
  Assign j_k and all unassigned clients in $N^2(j_k)$ to i_k;
  $C \leftarrow C \setminus \{j_k\} \setminus N^2(j_k)$;

In this scenerio, instead of directly bounded the facility cost by OPT, we simply bound it as:

$$Facility\ Cost \leq \sum_{i \in F} f_i y_i^*.$$

Now analyze the assignment cost. Let j_k be the selected client in iteration k, and i_k be the facility assigned to j_k. The expected assignment cost of assigning j_k to i_k is:

$$C_{j_k}^* = \sum_{i \in N(j_k)} c_{j_k i} x_{j_k i}^*,$$

where $C_{i_k}^*$ is the assignment cost of j_k in the LP solution. For j_k, the assignment cost is exactly $C_{j_k}^*$, and this is bounded by OPT.
Next, consider an unassigned client $l \in N^2(j_k)$, which needs to be assigned to i_k. The algorithm ensures that such a client l is assigned to the nearest open facility. The cost for assigning l is at most $c_{lh} + c_{hj_k} + C_{j_k}^*$, where the definition of h is the same as before. Use the triangle inequality, we can deduce that, the total assignment cost for l is bounded by:

$$v_j^* + v_{j_k}^* + C_{j_k}^*.$$

Since the algorithm chooses j_k to minimize $v_j^* + C_j^*$ among all unassigned clients, we know:

$$v_{j_k}^* + C_{j_k}^* \leq v_j^* + C_j^*,$$

and therefore, the assignment cost for l is at most:

$$v_l^* + v_{j_k}^* + C_{j_k}^* \leq 2v_l^* + C_l^*.$$

The total assignment cost is bounded by:

$$Assignment\ Cost \leq 2\sum_{j \in D} v_j^* + \sum_{l \in D} C_l^*.$$

Thus, the total cost is bounded by:

$$Total\ Cost \leq \sum_{i \in F} f_i y_i^* + 2\sum_{j \in D} v_j^* + \sum_{l \in D} C_l^*,$$

Since

$$\sum_{i \in F} f_i y_i^* + \sum_{j \in D} v_j^* \leq OPT,$$
$$\sum_{l \in D} C_l^* \leq OPT,$$

we have:

$$Total\ Cost \leq 3 \cdot OPT.$$

Thus, the randomized rounding algorithm achieves a 3-approximation for the uncapacitated facility location problem. [12][19]

3.3 Primal-Dual Method

The primal-dual method is a powerful technique for designing approximation algorithms that leverage the duality properties of linear programming relaxations. By formulating the primal and dual problems and exploiting their relationships, this method can yield efficient solutions with strong approximation guarantees. In the context of the uncapacitated facility location problem, the primal-dual method offers a 3-approximation algorithm that balances facility opening and client assignment costs optimally. This algorithm demonstrates the versatility and effectiveness of the primal-dual approach in addressing complex optimization problems. [7]
The sudocode for this algorithm is shown below:[19]

Algorithm 3: Primal-dual algorithm for the uncapacitated facility location problem.

```
v ← 0, w ← 0;
S ← D, T ← ∅;
while S ≠ ∅ do
    /* While not all clients neighbor a facility in T                    */
    Increase v_j for all j ∈ S and w_ij for all i ∈ N(j), j ∈ S uniformly until;
    if some j ∈ S neighbors some i ∈ T then
        S ← S − {j};
    if i ∉ T has a tight dual inequality then
        /* Facility i is added to T                                       */
        T ← T ∪ {i};
        S ← S − N(i);
T' ← ∅;
while T ≠ ∅ do
    Pick i ∈ T;
    T' ← T' ∪ {i};
    /* Remove all facilities h if some client j contributes to both h and
       i                                                                  */
    T ← T − {h ∈ T : ∃j ∈ D, w_ij > 0 and w_hj > 0};
```

3.4 Local Search Heuristics

The UFLP can be efficiently addressed using local search heuristics. This technique involves starting with an initial feasible solution and iteratively improving it by making localized modifications, such as opening or closing facilities, or reassigning clients to facilities. The process continues until no further improvement is possible, reaching a local optimum.[11]
The performance of the local search algorithm for the UFLP is notable for its simplicity and effectiveness in providing near-optimal solutions within reasonable computational time. Its bounded approximation ratio ensures that the solution quality is theoretically guaranteed, making it a reliable choice for practical applications. However,

the algorithm's efficiency heavily depends on the design of the local moves and the evaluation process, which can impact its convergence rate. Additionally, while the approximation ratio provides a worst-case performance guarantee, the practical performance may vary depending on the problem instance, with the algorithm potentially getting trapped in local optima for more complex input distributions. Thus, its performance could benefit from hybridization with other metaheuristic techniques or enhanced initialization strategies for further improvement.

4 Key Variants of Facility Location Problems

Building upon the foundational approximation algorithms for the uncapacitated facility location problem, we now turn our attention to key variants of the facility location problem that arise in more complex and realistic scenarios. These include the **capacitated facility location problem**, **dynamic and online facility location**, and **multi-objective facility location**. Each of these variants introduces unique challenges stemming from additional constraints, temporal dynamics, or the need to balance multiple objectives. In the following sections, we will explore these challenges in detail and provide a concise overview of the solution approaches developed to address them, highlighting both their theoretical significance and practical implications.

4.1 Capacitated Facility Location Problem (CFLP)

Definition. The CFLP introduces capacity constraints to traditional FLP, ensuring that each facility can only serve a limited demand. The problem is typically formulated as a mixed-integer program (MIP):

$$\min \left(\sum_{i \in F} f_i y_i + \sum_{i \in F} \sum_{j \in C} c_{ij} x_{ij} \right)$$

subject to

$$\begin{aligned} \sum_{j \in C} x_{ij} &\leq u_i y_i, \forall i \in F \\ \sum_{i \in F} x_{ij} &= d_j, \forall j \in C \\ x_{ij}, y_i &\in \{0, 1\}, \end{aligned}$$

where F represents facilities, C represents customers, f_i is the fixed cost of opening facility i, c_{ij} is the cost of serving customer j from facility i, u_i is the capacity of facility i, and d_j is the demand of customer j. The objective is to minimize the total cost of opening facilities and serving customers[3].

Challenges and Solution Approaches. The capacitated facility location problem (CFLP) is a well-known NP-hard problem, with its complexity increasing significantly as the number of facilities and customers grows. The inclusion of capacity constraints adds additional layers of difficulty by expanding the solution space and imposing computational challenges. Specifically, the coupling of allocation variables x_{ij} with facility opening variables y_i creates interdependencies that make solving even moderate-sized instances infeasible using traditional methods.

The primary solution approaches for CFLP can be categorized as follows:

- **Exact Methods:** Techniques such as branch-and-bound, branch-and-cut, and integer programming solvers are widely employed to solve small to medium-sized instances of CFLP. Although these methods guarantee optimal solutions, their computational expense renders them impractical for large-scale instances. Among these, Benders decomposition has proven particularly effective. By separating the facility location decision (master problem) from the customer allocation decision (subproblem), this method iteratively refines the solution through the generation of Benders cuts. It is noteworthy for its ability to efficiently handle capacity constraints by focusing exclusively on violated cuts, enabling faster convergence compared to straightforward branch-and-bound approaches [16].

- **Heuristic and Metaheuristic Methods:** For large-scale instances, heuristic and metaheuristic methods provide scalable and practical alternatives by delivering high-quality approximate solutions. Simulated annealing, for instance, is particularly effective for solving large-scale CFLP instances, offering robust performance and high-quality solutions. Its simplicity and adaptability have made it widely utilized in operational research [14]. On the other hand, genetic algorithms stand out for their flexibility and potential to achieve near-optimal solutions for complex instances. When paired with problem-specific operators, these algorithms are particularly effective for tackling multi-modal optimization problems such as the CFLP, though they can be computationally intensive [9].

4.2 Dynamic and Online Facility Location (DOFLP)

Definition. DOFLP extends FLP to account for temporal changes or sequential decision-making under incomplete information. Facilities may open, close, or relocate as demand evolves dynamically. Online FLP, a subset, assumes demand points are revealed sequentially, requiring immediate allocation decisions without future knowledge.
Dynamic FLP aims to minimize costs over a planning horizon:

$$\min \left(\sum_{t=1}^{T} \sum_{i \in F} f_{it} y_{it} + \sum_{t=1}^{T} \sum_{i \in F} \sum_{j \in C} c_{ijt} x_{ijt} \right).$$

Here t denotes the time periods, f_{it} and c_{ijt} represent the fixed and variable costs at time t, respectively.
Online FLP, a subset of DOFLP, assumes that demand points are revealed sequentially, requiring immediate and irrevocable allocation decisions without knowledge of future demands. Competitive analysis is often used in this context to compare the performance of online algorithms to the optimal offline solution. Meyerson (2001) introduced the first online algorithm for facility location, achieving a logarithmic competitive ratio [15].

Challenges and Solution Approaches. DOFLPs present unique challenges due to evolving costs, uncertain demand patterns, and the inherent need for sequential decision-making under incomplete information. Time-dependent changes complicate long-term planning, while online FLPs demand irrevocable decisions, increasing the risk of suboptimal solutions. Additionally, balancing short-term and long-term costs, while managing disruptions such as facility relocations, requires the application of robust or stochastic optimization techniques to handle uncertainty and variability effectively. These factors necessitate the development of advanced models and algorithms to ensure practical and efficient solutions.

- **Dynamic Programming:** Dynamic programming (DP) provides a framework for modeling DOFLPs as a sequence of interdependent decisions. While DP guarantees globally optimal solutions, its computational complexity grows exponentially with the number of time periods, rendering it infeasible for large-scale instances.

- **Rolling Horizon Techniques:** Rolling horizon techniques (RHT) iteratively solve the problem over a finite planning horizon, updating the solution as new information becomes available. This approach balances computational tractability with adaptability to dynamic changes. For example, rolling horizon control has been shown to be effective in logistics applications, where demand forecasts are regularly updated [10].

- **Competitive Analysis for Online FLP:** In the context of online FLPs, competitive analysis evaluates the performance of algorithms relative to the optimal offline solution. Meyerson's seminal algorithm remains a benchmark for online facility location, achieving a competitive ratio of $O(\log n)$, where n is the number of facilities [15].

4.3 Multi-Objective Facility Location (MOFLP)

Definition. MOFLP considers multiple, often conflicting objectives, such as minimizing costs while maximizing service coverage or equity. These problems are generally expressed as:[8]

$$\min Z = \{cost, distance, equity, \cdots\}.$$

Challenges and Solution Approaches. MOFLPs are characterized by significant challenges arising from conflicting objectives, such as cost minimization versus equity maximization, which require complex trade-offs based on stakeholder preferences. The presence of high-dimensional Pareto fronts, particularly for problems with more than two objectives, complicates visualization and interpretation, making decision-making less intuitive. Additionally, eliciting decision-makers' preferences is often challenging, especially when these preferences are dynamic or context-dependent. Finally, solving MOFLPs involves identifying a set of Pareto-optimal solutions rather than a single optimal solution, which substantially increases computational complexity [8]. Addressing these challenges requires advanced and tailored methods that balance conflicting objectives while maintaining computational efficiency.

- **Scalarization Methods:** Scalarization techniques, such as the weighted-sum and ϵ-constraint methods, transform MOFLPs into single-objective problems. These methods are conceptually simple and computationally efficient; however, they often struggle to produce diverse solutions across the Pareto front and may fail for non-convex Pareto sets [5].
- **Evolutionary Multi-objective Optimization:** Evolutionary multi-objective optimization (EMO) algorithms are highly effective at exploring the Pareto front comprehensively, generating diverse trade-off solutions for decision-makers. However, their computational cost can be prohibitive for large-scale problems, which necessitates the use of parallelization techniques or hybridization with other optimization methods to improve efficiency [6].
- **Robust Multi-objective Optimization:** Robust multi-objective optimization (RMO) methods explicitly address uncertainty in MOFLPs by ensuring solutions remain effective under varying scenarios. This approach is particularly suited to dynamic or uncertain environments, but it introduces additional computational complexity due to the need to evaluate solutions across multiple scenarios [1].

5 Implementation Results and Analysis

In this section, I present the results and analysis of various approximation algorithms for the Uncapacitated Facility Location Problem (UFLP), including deterministic rounding, randomized rounding, and local search methods. The focus is on comparing their performance and evaluating their effectiveness, with detailed implementation omitted. The corresponding code is available in the "code" folder attached to the report.

5.1 Experiment Settings

Problem (Data) Generation. To ensure the reproducibility of the experiments, the data for the experiments is generated using a controlled random number generator with a fixed random seed. Specifically, for each problem instance, we first determine the number of facilities and clients. Then, the random number generator is used to generate the facility cost for each facility and the assignment cost for each client-facility pair.

Performance Metrics. When evaluating approximation algorithms, the primary metric of interest is their approximation accuracy. However, obtaining the optimal solution for general cases is computationally infeasible, making it impossible to directly calculate the approximation ratio for each problem instance. Instead, we rely on the total cost of the solution produced by each algorithm. Since the total cost varies across problem instances and algorithms, a baseline algorithm is needed for comparison. We define the baseline algorithm as the scenario where all facilities are opened, regardless of the specific instance. The approximation accuracy is then evaluated as the ratio of the total cost of the solution produced by each algorithm to the total cost of the solution produced by the baseline algorithm.

More specifically, let P_i denote problem instance i, characterized by the random seed s_i, the number of facilities m, and the number of clients n. Let $C_j(P_i)$ represent the total cost of algorithm j's output on problem instance P_i. The approximation accuracy metric for algorithm j on P_i is defined as:

$$\alpha_j(P_i) = \frac{C_j(P_i)}{C_b(P_i)},$$

where $C_b(P_i)$ represents the total cost for baseline algorithm on problem instance P_i. The overall approximation accuracy for algorithm j across all problem instances is then defined as:

$$\alpha_j = \frac{1}{|P|} \sum_i \alpha_j(P_i),$$

where $|P|$ is the total number of problem instances tested.

A smaller value of α_j indicates that algorithm j produces solutions closer to the optimal, thereby demonstrating better approximation accuracy.

Results. I set $m = 50$, $n = 200$ in the evaluation process. The picture below shows how the approximation accuracy (approximation value) of different algorithms varies with the random seed:

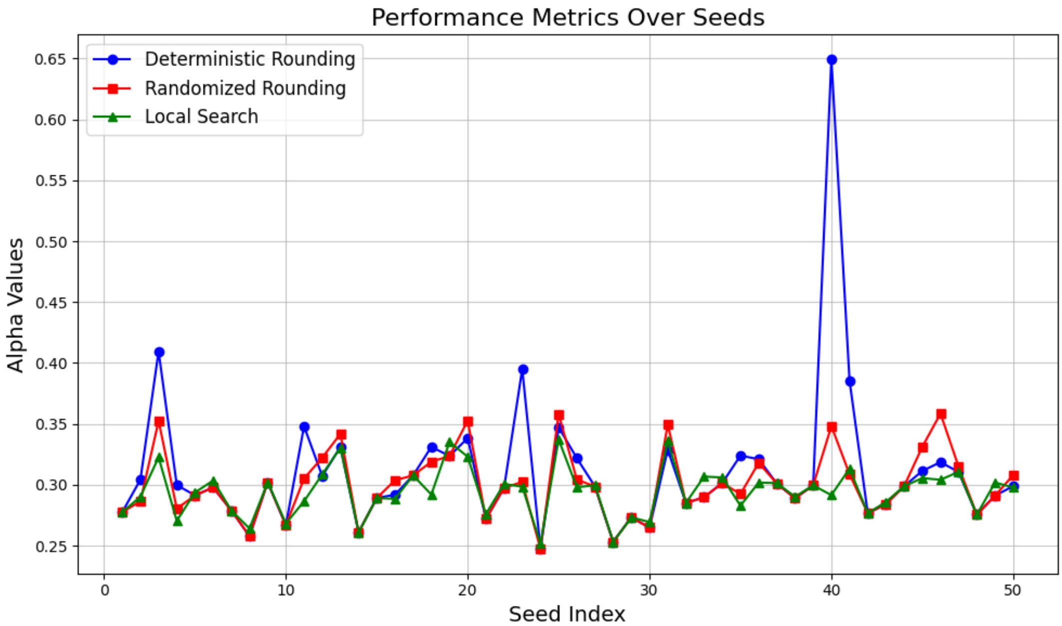

Figure 1: Evaluation Results

The approximations values are:

```
Deterministic Rounding:  0.31048292685284795
Randomized Rounding:  0.3001680024659187
Local Search:  0.2943983815479334
```

Figure 2: Approximation Values

Analysis. From the experimental results, we can observe that, overall, the local search algorithm demonstrates the best approximation accuracy among the three algorithms, followed by randomized rounding, which aligns with our expectations.
The results also indicate that local search is more robust compared to the other two algorithms, while deterministic rounding appears to be the most unstable. This instability arises because deterministic rounding heavily depends on the input data, making it more sensitive to variations caused by the random seed.
In our theoretical analysis (Section 3), we established that the approximation ratio is 4 for deterministic rounding, 3 for randomized rounding, and between 2 and 3 (specifically greater than $1 + \sqrt{2}$) for local search. However, these theoretical ratios are not always evident in the experimental results.
One possible explanation is the low quality of the data used in the experiments, as both facility and assignment costs were generated randomly. Such random data may lack realistic structure or complexity, which can impact the evaluation of the algorithms.
Specifically, random cost generation can lead to scenarios where certain naive or suboptimal approaches perform unusually well. This randomness fails to reflect the constraints and distributions of real-world problems, potentially obscuring the true strengths and weaknesses of the algorithms. Therefore, incorporating more structured or domain-specific data is essential for achieving a fair and meaningful evaluation of algorithm performance.

6 CONCLUSION

This survey has provided a comprehensive overview of approximation algorithms for the Facility Location Problem (FLP) and its key variants. Through an exploration of historical developments and modern techniques, we identified the pivotal contributions that have shaped the field, from linear programming relaxations and primal-dual methods to local search heuristics. These advancements have significantly improved our ability to solve computationally challenging problems with provable approximation guarantees.

In addition to theoretical advancements, we implemented several approximation algorithms for the classical UFLP to evaluate their practical performance. These implementations allowed us to compare their computational efficiency and approximation accuracy under various scenarios, highlighting the trade-offs between solution quality and runtime. Our experimental results reinforced the theoretical insights while also uncovering practical nuances, such as the influence of data quality and randomness on algorithm robustness.

For key FLP variants, we discussed the unique challenges and solution approaches:

- Capacitated Facility Location Problems (CFLPs) demand efficient methods for managing capacity constraints, with approaches such as Benders decomposition and metaheuristics showing great promise for large-scale problems [16, 14].

- Dynamic and Online Facility Location Problems (DOFLPs) introduce temporal complexity and incomplete information, where dynamic programming, rolling horizon techniques, and competitive analysis provide viable solutions for balancing short- and long-term objectives [15, 10].

- Multi-Objective Facility Location Problems (MOFLPs) require balancing conflicting objectives, with scalarization methods, evolutionary algorithms, and robust optimization emerging as effective tools for navigating the Pareto front [8, 6].

Despite these advances, several directions for future research remain compelling. For CFLPs, there is a need for more scalable and robust algorithms to address large-scale instances. For DOFLPs, machine learning methods hold potential for improving adaptability to dynamic and uncertain environments. In MOFLPs, further exploration of robust multi-objective optimization techniques could enhance decision-making under uncertainty. Additionally, investigating hybrid methods that combine optimization and learning techniques could yield novel approaches with practical relevance.

In conclusion, this survey highlights the rich and evolving nature of FLPs, offering insights into their underlying complexity and the diverse methods used to address them. By synthesizing existing knowledge and identifying gaps, we hope this work serves as a foundation for advancing research and practical applications in the field.

REFERENCES

[1] Aharon Ben-Tal, Arkadi Nemirovski, and Laurent El Ghaoui. Robust optimization. 2009.

[2] Moses Charikar, Sudipto Guha, Éva Tardos, and David B Shmoys. A constant-factor approximation algorithm for the k-median problem. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pages 1–10, 1999.

[3] Wikipedia contributors. Capacitated facility location. Accessed: 2025-01-05.

[4] Gerard Cornuejols, Marshall L Fisher, and George L Nemhauser. Exceptional paperâlocation of bank accounts to optimize float: An analytic study of exact and approximate algorithms. *Management science*, 23(8):789–810, 1977.

[5] Indraneel Das and John E Dennis. A closer look at drawbacks of minimizing weighted sums of objectives for pareto set generation in multicriteria optimization problems. *Structural optimization*, 14:63–69, 1997.

[6] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE transactions on evolutionary computation*, 6(2):182–197, 2002.

[7] Donglei Du, Ruixing Lu, and Dachuan Xu. A primal-dual approximation algorithm for the facility location problem with submodular penalties. *Algorithmica*, 63:191–200, 2012.

[8] Reza Zanjirani Farahani, Maryam SteadieSeifi, and Nasrin Asgari. Multiple criteria facility location problems: A survey. *Applied mathematical modelling*, 34(7):1689–1709, 2010.

[9] David E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Longman Publishing Co., Inc., USA, 1st edition, 1989.

[10] Tore Grünert and Hans-Jürgen Sebastian. Planning models for long-haul operations of postal and express shipment companies. *European Journal of Operational Research*, 122(2):289–309, 2000.

[11] Anupam Gupta and Kanat Tangwongsan. Simpler analyses of local search algorithms for facility location. *arXiv preprint arXiv:0809.2554*, 2008.

[12] Kamal Jain and Vijay V Vazirani. Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and lagrangian relaxation. *Journal of the ACM (JACM)*, 48(2):274–296, 2001.

[13] KRAR Jakob and Peter Mark Pruzan. The simple plant location problem: Survey and synthesis. *European journal of operational research*, 12(36-81):41, 1983.

[14] Scott Kirkpatrick, C Daniel Gelatt Jr, and Mario P Vecchi. Optimization by simulated annealing. *science*, 220(4598):671–680, 1983.

[15] Adam Meyerson. Online facility location. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pages 426–431. IEEE, 2001.

[16] Ragheb Rahmaniani, Teodor Gabriel Crainic, Michel Gendreau, and Walter Rei. The benders decomposition algorithm: A literature review. *European Journal of Operational Research*, 259(3):801–817, 2017.

[17] David B Shmoys and Chaitanya Swamy. Stochastic optimization is (almost) as easy as deterministic optimization. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 228–237. IEEE, 2004.

[18] David B Shmoys, Éva Tardos, and Karen Aardal. Approximation algorithms for facility location problems. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 265–274, 1997.

[19] David P Williamson and David B Shmoys. *The design of approximation algorithms*. Cambridge university press, 2011.