

# Grokking Algorithms PDF (Limited Copy)

Aditya Y. Bhargava

## **grokking** **algorithms**

*An illustrated guide for  
programmers and other curious people*

Aditya Y. Bhargava



More Free Book



Scan to Download

# **Grokking Algorithms Summary**

"Visualize and Simplify Complex Algorithms Easily."

Written by Books1

**More Free Book**



Scan to Download

## About the book

In a world where algorithms run silently behind the scenes, orchestrating everything from our search queries to social media feeds, understanding their magical inner workings can seem like a daunting task. **\*\*Grokking Algorithms\*\*** by Aditya Y. Bhargava serves as your accessible, insightful guide into the fascinating universe of algorithms, demystifying complex concepts with ease and elegance. Without delving into overwhelming jargon, Bhargava employs vibrant visual aids, captivating analogies, and hands-on examples to reveal the backbone of computer science. Whether you're an aspiring programmer or seasoned developer, this book builds a bridge from the abstract to the tangible, inviting you to master algorithms in a delightfully engaging way. Embark on this illuminating journey and discover how algorithms can not only make you a more proficient coder but also enable you to solve real-world problems with newfound clarity and creativity.

**More Free Book**



Scan to Download

## About the author

Aditya Y. Bhargava is a renowned software engineer, dedicated educator, and author known for his expertise in making complex technical topics more accessible and engaging. With a strong background in computer science and a wealth of experience across various programming paradigms, Bhargava has focused his career on demystifying algorithms to empower both novice and experienced programmers alike. His practical approach to teaching is evident in "Grokking Algorithms," where he uses relatable metaphors, vivid illustrations, and hands-on examples to break down intricate concepts into easily digestible parts. Bhargava's work has not only contributed significantly to the field of computer science education, but it has also provided a comprehensive foundation for countless learners to excel in their coding endeavors.

**More Free Book**



Scan to Download





# Try Bookey App to read 1000+ summary of world best books

Unlock **1000+** Titles, **80+** Topics

New titles added every week

- Brand
- Leadership & Collaboration
- Time Management
- Relationship & Communication
- Business Strategy
- Creativity
- Public
- Money & Investing
- Know Yourself
- Positive Psychology
- Entrepreneurship
- World History
- Parent-Child Communication
- Self-care
- Mind & Spirituality

## Insights of world best books



Free Trial with Bookey



# Summary Content List

Chapter 1: Introduction to Algorithms

Chapter 2: Selection Sort

Chapter 3: Recursion

Chapter 4: Quicksort

Chapter 5: Hash Tables

Chapter 6: Breadth-First Search

Chapter 7: Dijkstra's Algorithm

Chapter 8: Greedy Algorithms

Chapter 9: Dynamic Programming

Chapter 10: K-nearestneighbors

Chapter 11: Where to Go Next

Chapter 12: Answers to Exercises

**More Free Book**



Scan to Download

# Chapter 1 Summary: Introduction to Algorithms

In the initial chapter of this book, you are presented with foundational concepts of algorithms that will recur throughout the text. This chapter covers binary search, introduces the Big O notation to describe the running time of algorithms, and outlines a common technique for designing algorithms—recursion. An algorithm is essentially a set of instructions for completing a task, much like a recipe in cooking or a blueprint in architecture. Algorithms become interesting when they solve problems swiftly or ingeniously, and this book focuses on such algorithms.

One of the key highlights is binary search, a method that dramatically speeds up searches within a sorted list. For instance, instead of examining up to four billion elements step-by-step, binary search can locate an item in about 32 steps if the list is sorted. The effectiveness of algorithms like binary search is measured using Big O notation, which provides insight into the efficiency of algorithms as the input size grows. Throughout the book, you'll learn not just the algorithms themselves, but how to assess their efficiency and applicability.

Understanding performance trade-offs is crucial. Even though pre-written implementations exist, grasping these trade-offs allows you to choose the most suitable algorithm and data structures for a task. For instance, selecting between merge sort and quicksort or choosing an array over a list can have

**More Free Book**



Scan to Download

significant impacts on the performance of your applications.

Part of the problem-solving journey involves learning how to apply algorithms to diverse challenges. You'll delve into using graph algorithms for route calculations, dynamic programming for AI applications like checkers, and recognizing problems that can't be solved efficiently in real-time, known as NP-complete problems. By identifying these, you can employ algorithms that provide approximate solutions.

Binary search is an exemplary algorithm often used in practical scenarios, like searching a name in a phonebook or verifying a username on platforms like Facebook. By narrowing down the possibilities by half with each step, it quickly homes in on the desired item in a sorted list. This efficiency is quantified using Big O notation, where binary search is expressed as  $O(\log n)$ , compared to the linear search's  $O(n)$ , highlighting binary search's superior performance as list sizes increase.

A grasp of basic algebra is recommended, and familiarity with any programming language, with Python as an ideal choice for its beginner-friendly syntax. Understanding logarithms is also beneficial since logarithmic time complexity is a feature of binary search.

The chapter demonstrates this with a guessing game of numbers between 1 and 100, showcasing how binary search efficiently reduces the number of





guesses needed. This is juxtaposed against simple search, which sequentially checks numbers and can be vastly slower.

Big O notation further describes algorithm performance. The notation is concerned with growth rates: how quickly the time taken by an algorithm increases relative to input size. For example, while linear time algorithms grow proportionally ( $O(n)$ ), logarithmic time algorithms like binary search grow much slower ( $O(\log n)$ ), making them immensely faster with large datasets.

To provide context, algorithms are likened to drawing grids on paper: one box at a time ( $O(n)$ ) or by folding paper multiple times for exponential results ( $O(2^n)$ ). Lastly, the chapter touches upon factorial time algorithms ( $O(n!)$ ), using the traveling salesperson problem as an illustration. These algorithms grow at a prohibitive rate and often require approximate solutions instead of precise ones.

In summary, Chapter 1 lays the groundwork for understanding algorithm efficiency, offering a taste of the depth and breadth of problem-solving you will explore. It equips you with the fundamental principles to navigate through algorithms that power everything from search engines to artificial intelligence.



## Chapter 2 Summary: Selection Sort

### ### Chapter 2: Arrays, Linked Lists, and Selection Sort

In this chapter, we delve into two foundational data structures: arrays and linked lists. Both are ubiquitous in computer science and are essential for effective algorithm design. While the first chapter introduced arrays briefly, this chapter offers a deeper dive into their functionality and when to opt for linked lists instead. Understanding the differences—particularly in terms of performance for specific operations—is crucial for selecting the right structure for your algorithm.

#### #### Arrays

An array is a collection of elements stored in contiguous memory locations. Imagine it's like a chest of drawers, where each drawer can hold one element, allowing for efficient memory usage. This structure enables random access—meaning you can quickly retrieve an element if you know its index. This feature makes arrays ideal for scenarios requiring frequent reads, such as when implementing binary search, which requires sorted data.

However, arrays come with limitations. Adding or inserting elements can demand significant overhead, particularly if you're extending the array.



Consider a scenario where you need to add an element but find no immediate space next to your existing array. You might have to create a new, larger array and copy the elements over, which can be computationally expensive. Despite this, arrays are beneficial due to their ability to provide quick access to elements.

#### #### Linked Lists

Linked lists, in contrast to arrays, store elements anywhere in memory. Each item contains a reference to the address of the next item, forming a chain. This structure simplifies the process of adding or removing elements, as you simply adjust the links rather than moving every element. Picture it as a treasure hunt where each found clue leads you to the next.

Linked lists shine in situations where the structure heavily relies on frequent insertions and deletions because no rearrangement of existing elements is needed. However, accessing elements is sequential and can be inefficient for random access tasks, as you must traverse the list from the start to find a particular item.

#### #### Practical Considerations

Determining whether to use an array or linked list largely depends on the specific needs of your use case. For instance, if frequent updates and



modifications characterize your data handling, a linked list might be the better choice. On the other hand, if random access and frequent reading are your prime concerns, then an array is superior.

#### #### Introduction to Sorting: Selection Sort

Sorting is indispensable in computer science as many algorithms necessitate sorted data. This chapter introduces you to your first sorting algorithm: selection sort. Though inefficient compared to more advanced algorithms like quicksort (which will be discussed in the next chapter), mastering selection sort provides foundational understanding.

Selection sort works as follows:

- Identify and move the smallest element from the list to a new sorted list.
- Repeat this process, removing the next smallest element from the unsorted list and adding it to the new list.
- Continue until all elements are sorted.

Despite its simplicity, selection sort operates with a time complexity of  $O(n^2)$ , making it less optimal for large datasets. Nevertheless, learning it sets the stage for grasping more complex algorithms like quicksort.

Through the exercises, you will internalize the practical distinctions between arrays and linked lists, setting a firm foundation for the subsequent



exploration of more sophisticated data handling methods and sorting techniques.

**More Free Book**



Scan to Download

## Chapter 3 Summary: Recursion

### ### Chapter 2 Recap:

The prior chapter laid the foundation for understanding how computers manage memory and data structures. Imagine a computer's memory as a large set of organized drawers. For efficiently storing multiple data elements, you use either arrays or lists. Arrays store elements in contiguous memory locations, ensuring fast data access. On the other hand, linked lists distribute elements throughout memory, each element pointing to the next, which facilitates quick insertion and deletion operations. It's essential for arrays to contain elements of the same type, like all integers or all doubles, to optimize performance.

### ### Chapter 3: Recursion

This chapter delves into recursion, a fundamental programming technique pivotal for several algorithms. Recursion is akin to a problem-solving approach where a function calls itself. Although it can be polarizing—some programmers initially dislike it—it often becomes a favored technique after mastering its elegance and efficiency. To truly grasp recursion, it's helpful to analyze recursive functions by manually tracing through their execution with pen and paper.





- 1. Understanding Recursion:** Recursion simplifies solving problems by breaking them into smaller sub-problems, forming a recursive case, and identifying the simplest aspect of the problem, called the base case. As an analogy, searching for a key in nested boxes illustrates recursion: one recursive way to search would be to look inside each box, calling the same search function for any nested boxes, until the key is found.
- 2. Base Case and Recursive Case:** Recursion requires carefully defining a base case to prevent infinite loops. For instance, a countdown function benefits from recursion by reducing its current count until it reaches zero, its base case.
- 3. The Call Stack:** The call stack is an essential concept intertwined with recursion in programming. It manages the various function calls that occur in a program. Imagine a call stack as a stack of sticky notes representing each function call. Each function call places a new note on top of the stack (a push), and upon function completion, a note is removed (a pop). The call stack ensures that a function can pause and resume once another function completes.
- 4. Recursion in Action:** Walking through the factorial function showcases recursion and the call stack in tandem. For `fact(3)`, successive calls generate a stack where each level stores separate state information, tackling computations layer by layer.



**5. Memory Considerations:** Recursion leverages the call stack to track partially completed function calls, like maintaining a “pile of boxes” without explicitly creating one. However, each recursive call consumes memory, so extensive recursion risks exhausting memory resources, leading to stack overflow errors. Optimizing through loops or techniques like tail recursion can mitigate these issues.

**6. Exercises:** Readers explore manipulating call stacks and anticipate challenges with endlessly recursive functions that can deplete memory.

This chapter builds on the foundation of data structures and memory management from the previous discussion, advancing into recursive thinking crucial for tackling complex algorithmic challenges efficiently. Moving forward, the understanding of recursion will underpin more advanced problem-solving strategies introduced in subsequent chapters.



# Critical Thinking

**Key Point:** Understanding Recursion

**Critical Interpretation:** By mastering the art of recursion, you can transform complex challenges into manageable tasks in your everyday life, much like simplifying a daunting project into smaller, achievable steps. Just as recursion breaks down a problem into smaller sub-problems until reaching the simplest base case, you can approach life's challenges by dividing them into smaller actions until you identify a 'base' solution. Embracing this technique cultivates a strategic mindset that can unravel complexity, inspire clarity, and foster resilience in any situation, empowering you to tackle even the most intimidating tasks with confidence and efficiency.

More Free Book



Scan to Download

## Chapter 4: Quicksort

### Chapter 3: Recursion

Recursion is a fundamental concept in programming where a function calls itself to solve a problem. Every recursive function must have a base case, which is the simplest instance of the problem, and a recursive case, which reduces the problem into smaller versions of itself. All function calls in recursive operations are managed on a call stack, which temporarily holds data. As the stack can grow large, it consumes a significant amount of memory.

### Chapter 4: Divide and Conquer & Quicksort

This chapter introduces the divide-and-conquer (D&C) strategy, a powerful recursive technique for problem-solving. When algorithms seem insufficient to tackle a problem, D&C offers a fresh perspective by breaking the problem down into more manageable parts. One classic application of this technique is the quicksort algorithm, which is an elegant and efficient method for sorting.

### Divide and Conquer

**More Free Book**



Scan to Download

To grasp D&C, consider a scenario where a farmer wants to divide a piece of land into the largest possible square plots. The simplest solution, or base case, occurs when one side of the land is a multiple of the other. For instance, dividing a plot with sides of 25 meters and 50 meters yields a 25m x 25m square. The recursive case involves continuously breaking down the problem (land plot) until reaching the base case. A similar problem can be solved using Euclid's algorithm, a well-known method in mathematics for finding the greatest common divisor of two numbers.

Additionally, D&C can solve other problems, such as summing numbers in an array. By using recursion, one simplifies the task by continuously breaking the array down until reaching an array of zero or one element, which is straightforward to resolve.

## Quicksort

Quicksort is a D&C algorithm that significantly outperforms selection sort, which was previously discussed. The base case for quicksort involves arrays with zero or one element, which are inherently sorted. For larger arrays, quicksort selects a pivot, partitions the array into elements smaller and larger than the pivot, and recursively sorts the sub-arrays. Finally, the sorted



sub-arrays and pivot are combined to create the sorted array.

The effectiveness of quicksort relies on choosing an optimal pivot. While any pivot can work, the best scenario involves selecting a pivot that halves the array, reducing the problem size faster. Although the algorithm's worst-case scenario is  $O(n^2)$  time complexity, the average case, where pivots are chosen randomly or wisely, is much faster at  $O(n \log n)$ .

## Big O Notation and Comparison

Big O notation helps measure the performance of algorithms. Quicksort's average time complexity,  $O(n \log n)$ , makes it a preferred choice over other sorting algorithms like merge sort, despite its worst-case scenario. This is due to quicksort's smaller constant factors, making it generally faster in typical use cases.

## Functional Programming & Inductive Proofs

Recursion is a cornerstone of functional programming—languages like Haskell rely on it due to the absence of loops. Understanding recursion facilitates mastering functional languages. Additionally, the explanation touches on inductive proofs—a logical method to assure that algorithms





function as expected using base and inductive cases. This method is instrumental in validating recursive algorithms like quicksort.

## Chapter Recap

# Install Bookey App to Unlock Full Text and Audio

Free Trial with Bookey





# Why Bookey is must have App for Book Lovers



## 30min Content

The deeper and clearer interpretation we provide, the better grasp of each title you have.



## Text and Audio format

Absorb knowledge even in fragmented time.



## Quiz

Check whether you have mastered what you just learned.



## And more

Multiple Voices & fonts, Mind Map, Quotes, IdeaClips...

Free Trial with Bookey



## Chapter 5 Summary: Hash Tables

In this chapter, the focus is on hash tables, a fundamental and versatile data structure used across various programming tasks. The chapter elucidates how hash tables work, their performance implications, and practical applications.

### Introduction to Hash Tables and their Functionality:

Imagine working at a grocery store where you have to search a book for produce prices. If the book is unorganized, finding prices is time-consuming ( $O(n)$  time complexity). Even if sorted, like binary search, it's faster ( $O(\log n)$  time complexity) but still inefficient when customers are waiting. The ideal scenario would be having an assistant like Maggie, who knows prices instantly, mimicking a hash table's constant lookup time ( $O(1)$ ).

### Understanding Hash Tables through Hash Functions:

A hash table is built using a hash function to map strings to numbers. The function should be consistent (always returning the same number for the same input) and ideally map different inputs to different outputs. By feeding produce names into a hash function, you determine the index for storing their prices in an array. This setup allows you to retrieve prices without searching, thanks to consistent index mapping.



## Hash Table Internals:

Hash tables involve combining a hash function with an array to store key-value pairs. The keys are produce names, and values are prices. When querying a hash table, the hash function quickly determines the index, enabling efficient data retrieval. Most programming languages, like Python, have built-in hash table implementations (dictionaries), so manual implementation is rare.

## Common Use Cases for Hash Tables:

1. **Lookups:** A hash table efficiently maps one item to another, like a phone book where names link to phone numbers or DNS translating web addresses to IPs.
2. **Preventing Duplicates:** In scenarios such as voting booths, hash tables efficiently check for and filter out duplicate entries without scanning through entire datasets.
3. **Caching:** To speed up web service responses, hash tables store frequently accessed data, reducing server load and response times for repeated requests (e.g., caching a common web page).



## **Collisions and Performance:**

Collisions occur when multiple keys hash to the same index. They disrupt efficiency but can be managed, usually by chaining elements in linked lists at these indices. Performance depends on minimizing collisions through good hash functions and maintaining a low load factor to prevent long linked lists.

## **Hash Table Performance:**

Ideally, hash tables offer constant time ( $O(1)$ ) operations on average. However, worst-case scenarios can occur if many collisions happen, reverting operations to linear time ( $O(n)$ ). The load factor, the ratio of stored items to available slots, affects this; keeping it low minimizes collisions. Strategies like resizing (doubling array size when the load factor is too high) help maintain performance.

## **Choosing a Good Hash Function:**

A good hash function spreads entries evenly across a hash table to prevent clustering and minimize collisions. While developing such functions is complex, it's crucial to ensure efficient hash table operations.

## **Recap:**

**More Free Book**



Scan to Download

Hash tables are invaluable for tasks involving fast lookups, relationship modeling, duplicate elimination, and caching. They rely on effective hash functions to minimize collisions and maximize performance. Programming languages typically provide robust hash table implementations, freeing developers from building them from scratch.

Section	Summary
Introduction to Hash Tables and their Functionality	Illustrates hash tables' efficiency using a grocery store scenario. Emphasizes the constant lookup time offered by hash tables, making them ideal for quick data retrieval in situations with waiting customers.
Understanding Hash Tables through Hash Functions	Explains how hash functions map strings to numbers to store data in arrays, ensuring quick price retrieval without search delays by maintaining consistent index mapping.
Hash Table Internals	Details the combination of hash functions with arrays to store key-value pairs. Mentions built-in hash table implementations in programming languages, reducing the need for manual coding.
Common Use Cases for Hash Tables	Highlights use cases such as lookups, preventing duplicate entries, and caching. Examples include phonebooks, DNS, and reducing web server load.
Collisions and Performance	Discusses collision management through chaining and link lists at affected indices, aiming for good hash functions and low load factors to maintain efficient performance.
Hash Table Performance	Explains performance implications, including maintaining constant average time ( $O(1)$ ) while managing worst-case scenarios to avoid linear time ( $O(n)$ ). Load factor management and resizing strategies are highlighted.





Section	Summary
Choosing a Good Hash Function	Stresses the importance of a good hash function to evenly distribute entries, prevent clustering, and minimize collisions for efficient hash table operations.
Recap	Summarizes hash tables as crucial for fast lookups, relationship modeling, duplicate elimination, and caching, relying on effective hash functions to maintain performance.



## Critical Thinking

**Key Point:** Hash tables facilitate  $O(1)$  constant time complexity for lookups.

**Critical Interpretation:** Imagine a world where every piece of information you need is instantly available to you, akin to having a personal assistant who knows everything at the drop of a hat. Hash tables, with their efficient lookup capability, transform this fantasy into reality within tech ecosystems. In life, this inspires us to aim for processes that optimize time and maximize efficiency, much like hash tables optimize data retrieval. By organizing and structuring our tasks and goals with clarity, we can significantly reduce the noise and distractions that clutter our paths, allowing us to focus directly on what truly matters. This concept pushes us to simplify our approaches, reduce inefficiencies, and create personal systems that leverage the 'hash table mindset,' ensuring that we are primed to act quickly and effectively in every endeavor.



## Chapter 6 Summary: Breadth-First Search

Chapter 6 in this book introduces the concept of graphs, a fundamental data structure used to model relationships between entities. Unlike graphs with X or Y axes, these graphs consist of nodes (representing entities) and edges (representing connections between entities). Through this chapter, we'll explore the breadth-first search (BFS) algorithm, which is crucial for solving shortest-path problems and determining the connectivity between nodes. Additionally, this chapter discusses directed and undirected graphs and introduces the concept of topological sorting, an algorithm that highlights dependencies between nodes.

To begin, imagine navigating from Twin Peaks to the Golden Gate Bridge in San Francisco with the fewest bus transfers. This scenario exemplifies a shortest-path problem where BFS can find the minimum steps required. BFS answers questions such as "Is there a path from A to B?" and "What is the shortest path from A to B?" For example, BFS can help identify the fewest moves to checkmate in chess, the closest doctor in a network, or the shortest spelling correction.

Graphs are illustrated using examples like a group of friends playing poker to model who owes whom money. Nodes and edges represent friends and the monetary debts between them. In directed graphs, edges have a direction, indicating one-way relationships, while undirected graphs have bidirectional



relationships. The Twin Peaks example demonstrates that by using BFS, you can determine the shortest bus route to a destination. The algorithm involves modeling the problem as a graph and applying BFS to solve it.

As BFS operates, it expands outward from the starting point, checking first-degree connections (direct connections) before second-degree connections (friends of friends), prioritizing closer paths. This ensures the shortest route is found. This search requires orderly progression, adhering to a queue (First In, First Out), ensuring nodes are evaluated in the order they're added. Stacks, in contrast, follow a Last In, First Out order.

In implementing BFS, a queue is initiated and populated with the starting node's neighbors. Nodes are then checked sequentially to search for the target or identify the shortest path to it. A practical implementation using Python involves a hash table to map nodes to their neighbors and ensures no node is revisited. This prevents infinite loops where nodes could be repeatedly checked without progress, such as in cyclic graphs where a node points back to itself through a series of connections.

Additionally, topological sorting is introduced—a method to create an ordered list of tasks with dependencies. For instance, in a morning routine graph, tasks like "brush teeth" must precede "eat breakfast," and topological sorting helps organize tasks accordingly. The concept extends to problem-solving scenarios, like planning tasks in complex projects such as



wedding preparations.

The chapter closes by summarizing the key concepts and offering exercises to reinforce learning. Running times are discussed, with BFS operating in  $O(V+E)$  time complexity, where  $V$  is the number of vertices (nodes) and  $E$  the number of edges. The exercises encourage applying BFS on various graph structures and understanding trees, a special graph type where edges never loop back, reinforcing fundamental graph theory concepts.

**More Free Book**



Scan to Download

## Critical Thinking

**Key Point:** Graphs and connectivity: Using BFS to find shortest paths

**Critical Interpretation:** Imagine your life as a vast city, bustling with destinations and connections, where every goal, aspiration, and relationship is a node on your personal map. Each step you take, each decision made, mirrors the edges connecting these nodes, contributing to your life's unique web. By employing the breadth-first search (BFS) approach, you embrace a structured perspective—prioritizing the nearest, most direct opportunities first, ensuring you acknowledge and understand your closest connections before venturing further. This method encourages you to lean into the power of proximity and order, tackling immediate challenges before addressing those further afield. As you navigate through life's complex network, plotting the shortest path not only saves time but fosters deeper relationships, inciting a sense of fulfillment. You become proficient in mapping out pathways, recognizing dependencies, and organizing your life's journey efficiently. This way, every decision made is deliberate, considered, and step-wise, reflecting clarity and purpose on your path forward.

More Free Book



Scan to Download



# Chapter 7 Summary: Dijkstra's Algorithm

## Chapter Summary: Weighted Graphs and Dijkstra's Algorithm

This chapter introduces the concept of weighted graphs and the challenges they present. A weighted graph assigns a numerical value, or weight, to each edge, reflecting factors like travel time or costs, which influence finding the optimal path. Unlike unweighted graphs, which use breadth-first search to identify the shortest path based on the number of segments, weighted graphs demand a more sophisticated approach to finding the fastest path. This is where Dijkstra's algorithm comes into play.

### Dijkstra's Algorithm Explained

Dijkstra's algorithm is a method for determining the shortest path (in terms of total weight) from a starting node to other nodes in a weighted graph. The algorithm involves four main steps:

1. **Find the Cheapest Node:** Identify the node that can be reached with the least amount of time or cost from the starting node.
2. **Update Costs:** Consider all neighbors of the "cheapest" node and update the costs if a shorter path to them is found through this node.



3. **Repeat Until Completion:** This process of finding the cheapest node and updating costs is iterated until all nodes are processed.
4. **Calculate the Final Path:** Once all nodes are evaluated, the shortest path in terms of weight can be traced back using the "parent" relationships established during the process.

Dijkstra's algorithm, however, has limitations. Specifically, it does not work on graphs with negative weights, as these can lead to situations where a supposedly cheapest path is not actually optimal. In such cases, a different algorithm, the Bellman-Ford algorithm, is needed.

## Terminology and Context

- **Weighted vs. Unweighted Graphs:** In a weighted graph, edges carry weights; in an unweighted graph, they don't.
- **Cycles in Graphs:** A cycle allows you to start at a node, travel through edges, and return to the starting node. In certain graphs, cycles can complicate finding the shortest path but do not affect Dijkstra's algorithm unless negative weights are involved.
- **Directed vs. Undirected Graphs:** Directed graphs imply a one-way relationship between nodes while undirected graphs suggest a two-way exchange.



## Application Example

The chapter illustrates Dijkstra's algorithm through an example where a character, Rama, seeks to trade items (from a music book to a piano) at minimal cost, represented by negative or positive weights in a graph. Here, costs are depicted as monetary values related to each trade. By applying Dijkstra's algorithm, Rama determines the series of trades that incur the least expense. However, if trades involve negative values (e.g., receiving money back), Dijkstra's algorithm may fail to find the truly optimal path, highlighting the need for Bellman-Ford in such scenarios.

## Implementation

The chapter provides a guide to implementing Dijkstra's algorithm in Python using hash tables to represent the graph, including costs and parent nodes. It ensures each node is processed only once to finalize the shortest path in weighted, non-negatively weighted graphs.

## Recap and Key Insights

- **Breadth-First Search** is suitable for finding the shortest path in

More Free Book



Scan to Download

unweighted graphs.

- **Dijkstra's Algorithm** calculates the shortest path in weighted graphs, assuming all edge weights are non-negative.
- When dealing with negative weights, one must resort to the Bellman-Ford algorithm.

This chapter underscores the importance of understanding different types of graphs and their associated algorithms, illustrating how specific strategies map to particular graph features, ensuring efficient pathfinding and decision-making based on graph structure and edge attributes.



# Critical Thinking

**Key Point:** Finding the Cheapest Node

**Critical Interpretation:** Imagine navigating a maze where every turn has a cost. In life, many decisions present us with similar weighted choices, where some paths demand more resources or time than others. By embracing the idea of finding the 'cheapest node,' you focus on identifying the solution that requires the least cost or presents the most efficiency among your options. This mindset encourages you to evaluate decisions based not merely on the immediate outcome, but on the long-term benefits and costs associated. It teaches you to prioritize actions that allocate resources wisely, ensuring that each step you take brings you closer to your goals with minimal wastage or unnecessary detours.

More Free Book



Scan to Download

## Chapter 8: Greedy Algorithms

In Chapter 8, the focus is on understanding and applying greedy algorithms, particularly within the context of NP-complete problems. These problems don't have fast, definitive algorithmic solutions, but approximation algorithms provide quicker, near-optimal answers. The chapter begins with an exploration of the classroom scheduling problem, where the task is to maximize the number of non-overlapping classes in a single classroom. The solution is simple: always select the class that finishes the earliest, a classic demonstration of a greedy strategy. This approach often surprises people with its simplicity and effectiveness in providing a global optimal solution.

Next is the knapsack problem, where one must maximize the value of items in a knapsack with a weight limit. Here, a greedy approach involves picking the most valuable items within the weight capacity. However, this strategy doesn't always yield an optimal solution, as exemplified by comparing the value of stealing a stereo versus a combination of a laptop and guitar. Despite not always achieving perfection, greedy algorithms can deliver 'pretty good' results with ease.

The chapter then introduces the set-covering problem where a radio show must choose the minimum number of stations to cover all 50 states. Calculating every possible subset to find the smallest covering set is time-consuming and complex due to the exponential growth of subsets with



more stations, showcasing why exact solutions are impractical. Instead, approximation algorithms using greedy strategies can efficiently tackle this by iteratively choosing the station covering the most uncovered states, demonstrating their utility in handling NP-complete problems.

[View the full text of this article](#) | [Download the full text of this article](#) | [View the full text of this article](#) | [Download the full text of this article](#) | [View the full text of this article](#) | [Download the full text of this article](#) | [View the full text of this article](#) | [Download the full text of this article](#) | [View the full text of this article](#) | [Download the full text of this article](#)

## **Install Bookey App to Unlock Full Text and Audio**

**Free Trial with Bookey**





App Store  
Editors' Choice



22k 5 star review

## Positive feedback

Sara Scholz

tes after each book summary  
understanding but also make the  
and engaging. Bookey has  
ding for me.

**Fantastic!!!**



I'm amazed by the variety of books and languages  
Bookey supports. It's not just an app, it's a gateway  
to global knowledge. Plus, earning points for charity  
is a big plus!

Masood El Toure

**Fi**



Ab  
bo  
to  
my

José Botín

ding habit  
o's design  
ual growth

**Love it!**



Bookey offers me time to go through the  
important parts of a book. It also gives me enough  
idea whether or not I should purchase the whole  
book version or not! It is easy to use!

Wonnie Tappkx

**Time saver!**



Bookey is my go-to app for  
summaries are concise, ins  
curated. It's like having acc  
right at my fingertips!

**Awesome app!**



I love audiobooks but don't always have time to listen  
to the entire book! bookey allows me to get a summary  
of the highlights of the book I'm interested in!!! What a  
great concept !!!highly recommended!

Rahul Malviya

**Beautiful App**



This app is a lifesaver for book lovers with  
busy schedules. The summaries are spot  
on, and the mind maps help reinforce wh  
I've learned. Highly recommend!

Alex Walk

Free Trial with Bookey





# Chapter 9 Summary: Dynamic Programming

## ### Summary of Chapter 9: Dynamic Programming

This chapter introduces dynamic programming, a method used to tackle complex problems by breaking them down into simpler, smaller subproblems and solving these first. The foundational idea is to construct solutions to bigger problems based upon solutions to smaller subproblems.

## #### The Knapsack Problem

To delve into dynamic programming, we revisit the knapsack problem discussed earlier. Imagine being a thief with a knapsack that holds up to 4 pounds and choosing from three items to maximize stolen goods' value. The naïve solution involves considering every possible combination of items, which becomes impractically slow as the number of items rises, characterized by  $O(2^n)$  time complexity.

Dynamic programming provides an efficient approach by employing a grid to solve subproblems first, starting from smaller knapsack capacities up to the actual problem's capacity. Each grid cell represents a subproblem solution, helping refine the overall optimal solution iteratively.



## #### Algorithm Walkthrough

1. **Grid Setup:** Each row corresponds to an item (e.g., guitar, stereo), and each column corresponds to knapsack capacities ranging from 1 to 4 pounds.

### 2. Filling the Grid:

- Start by considering each item sequentially (guitar, then stereo, then laptop) and determine if it can be included given the knapsack capacity at that column.
- Update each grid cell by deciding if including the current item increases the total value while staying within the weight limit.

3. **Solution Construction:** The final cell in the grid (or the highest value found) gives the maximum value that can fit into the knapsack, effectively solving the problem.

## #### Handling Additional Complexity

- **Adding Items:** If a new item is available (e.g., an iPhone), add a row for it and update only necessary calculations.
- **Weight Changes:** If a new item's different weight granularity is introduced, a refined grid reflecting finer calculations would be required.



- **Dependencies & Subtasks:** Dynamic programming works best when subproblems are independent. Problems requiring dependency resolution, such as prioritizing tasks when certain items must precede others, aren't suitable for dynamic programming.

#### #### Longest Common Substring & Subsequence

Beyond simple optimization problems like the knapsack problem, dynamic programming can solve problems such as finding the longest common substring or subsequence between two words—vital in applications like DNA analysis, text comparison, and spell checking. Each cell in the grid represents stages of partial solutions, and fills based on whether characters from the words match at given indices.

#### #### Real-World Applications

Dynamic programming is invaluable in various fields, such as biological sequence analysis for DNA, version control diff tools, and algorithms measuring string similarity. It even extends into practical software development tasks like text wrapping in word processors.

#### ### Recap

- **Purpose:** Solve optimization problems by dividing them into discrete



subproblems.

- **Structure:** Typically involves constructing a grid where cells correspond to subproblems.
- **Application:** Effective for constraints-based optimization, and places where subproblems can be independently solved.
- **Key Lesson:** There isn't a universal formula; understanding how to construct the grid and break down subproblems is crucial.

In conclusion, the chapter illustrates dynamic programming's power through examples and provides insights into its applicability across diverse domains, emphasizing its flexibility and efficiency in scenarios with complex constraints.

Section	Description
Overview	This chapter introduces dynamic programming as a method to solve complex problems by breaking them down into simpler, manageable subproblems.
The Knapsack Problem	Revisiting the knapsack problem to demonstrate dynamic programming. Instead of checking all combinations ( $O(2^n)$ complexity), a grid is used to efficiently solve subproblems iteratively.
Algorithm Walkthrough	<div>Grid Setup: Rows for items; columns for capacities.</div> <div>Filling the Grid: Checks if adding items increases value while staying under weight limit.</div> <div>Solution Construction: The maximum value cell provides the solution.</div>



Section	Description
Handling Additional Complexity	Adding Items: Add a new row and update calculations. Weight Changes: Adjust grid for finer calculations. Dependencies: Suited for independent subproblems.
Longest Common Substring & Subsequence	Dynamic programming also solves longest common substring/subsequence problems, essential in DNA analysis and text comparison.
Real-World Applications	Applications in DNA sequence analysis, text similarity measurement, version control, and layout solutions in software development.
Recap	Main ideas include:  Purpose: Divide complex problems into subproblems. Structure: Utilize a grid for subproblem solutions. Application: Useful for constraint-based problems where subproblems are independent. Key Lesson: No universal formula; constructing grid and problem breakdown is key.



## Chapter 10 Summary: K-nearestneighbors

In this chapter, the focus is on understanding and utilizing the k-nearest neighbors (KNN) algorithm, a foundational tool in machine learning for classification and regression tasks. The chapter begins by explaining the concept of KNN through a simple analogy involving fruit classification, where the size and color of a fruit help determine whether it's an orange or a grapefruit. This analogy introduces the fundamental idea of comparing data points based on certain features.

The KNN algorithm is a straightforward yet powerful tool for classification tasks. It identifies the category to which a data point belongs by examining the categories of its nearest neighbors. For instance, if a fruit is being classified as an orange or grapefruit, one would look at the closest classified fruits to determine a category. In practical applications, KNN is often the first algorithm to try when handling classification challenges due to its simplicity and effectiveness.

A real-world application of KNN, demonstrated in the text, involves building a movie recommendation system similar to what platforms like Netflix use. Users are plotted on a graph based on their movie preferences, and recommendations are made by identifying users with similar tastes and suggesting movies they liked. This process requires determining how similar two users are, which involves feature extraction—a crucial step in any



machine learning task. For fruit, this might mean size and color, while for users, it involves their ratings of various movie genres.

Feature extraction translates into a multidimensional space where the distance between points can be measured using the Pythagorean theorem to determine similarity. The more accurately features represent the actual similarities, the better the KNN algorithm performs. Netflix, for example, improves recommendations by encouraging users to rate more movies, thus refining the similarity measurement.

Regression, another core function of KNN, involves predicting a numerical output such as a user's movie rating or the number of bakery loaves to prepare on a given day. This prediction is based on historical data and the assumption that similar situations will yield similar outcomes.

The chapter also touches on challenges in feature selection—ensuring features directly correlate with the prediction task and avoiding bias. For example, asking users to only rate certain genres can skew the recommendation results, emphasizing the necessity of carefully chosen features.

The discussion transitions to broader themes in machine learning, introducing the concept of optical character recognition (OCR), wherein features like lines and curves in numbers are extracted to aid in recognition



tasks. Similarly, a spam filter example highlights Naive Bayes, another algorithm used for classifying emails based on word probabilities linked to spam.

Ultimately, predicting complex systems like the stock market is cited as challenging due to the numerous variables involved, illustrating the limits of machine learning. Nevertheless, the combination of classification and regression through algorithms like KNN allows for diverse applications, from OCR and spam filters to personalized media recommendations.

The chapter concludes by emphasizing the significance of feature extraction and selection in ensuring the success of KNN and machine learning systems, recognizing the algorithm's pivotal role in the evolving field of artificial intelligence.

**More Free Book**



Scan to Download



## Chapter 11 Summary: Where to Go Next

In this summary, we dive into Chapter 11 and a section titled "Where to Go Next," which explores various algorithms and subjects that weren't covered in the main body of the book. The focus is on enhancing the reader's understanding and piquing interest in broader algorithmic concepts.

### ### Binary Search Trees

The chapter revisits binary search by introducing binary search trees (BSTs), a data structure that maintains sorted order and enables efficient insertion, deletion, and search operations. Unlike sorted arrays, BSTs can dynamically handle user entries without requiring constant re-sorting. The biggest advantage is their efficiency in insertion and deletion. However, they need to be balanced to maintain performance, exemplified by structures like red-black trees.

### ### Inverted Indexes

The section explains the concept of inverted indexes, pivotal for search engines. In this data structure, words serve as keys, and lists of corresponding documents or pages are values, enabling quick retrieval of where a search term appears.

### ### Fourier Transform

A versatile algorithm, the Fourier transform can decompose complex signals

**More Free Book**



Scan to Download

into simpler frequency components. This makes it invaluable in fields such as audio compression, signal processing, and even earthquake prediction due to its ability to separate and manipulate frequency data.

### ### Parallel Algorithms

Parallel algorithms are essential for maximizing computational efficiency by leveraging multi-core processors. They are intricate to design and audit, focusing on dividing tasks effectively across cores to minimize idle time and maximize throughput.

### ### MapReduce

Distributed computing brings us to MapReduce, an algorithm framework ideal for processing massive datasets across numerous machines. Utilizing the map and reduce functions, it allows operations on distributed data, as showcased by tools like Apache Hadoop.

### ### Bloom Filters and HyperLogLog

Bloom filters introduce a probabilistic approach to efficiently determine if an item is in a set, allowing false positives but not false negatives.

HyperLogLog extends this by providing approximate counts of unique items in large datasets, offering memory-efficient solutions for scenarios demanding estimation over precision.

### ### SHA Algorithms

**More Free Book**



Scan to Download

SHA represents a family of secure hash algorithms that generate fixed-size outputs from data inputs. These algorithms are essential for data integrity checks and secure password storage, where they ensure that even if system data is compromised, original values remain protected.

### ### Locality-sensitive Hashing

Locality-sensitive hashing, exemplified by Simhash, allows for hashing that can identify similar items by producing similar hash values. This is particularly useful for identifying duplicates or similar content within large datasets.

### ### Diffie-Hellman Key Exchange

A foundational cryptographic method, Diffie-Hellman enables secure communication by allowing two parties to establish a shared secret over an insecure channel, without having to pre-share private keys, thus paving the way for further developments like RSA encryption.

### ### Linear Programming

Finally, the chapter introduces linear programming, a mathematical technique for optimizing a linear objective function subject to linear equality and inequality constraints. This method is widely used for resource allocation and operational efficiency, leveraged by the Simplex algorithm.

### ### Conclusion



The chapter closes by encouraging exploration beyond the book's teachings, suggesting linear programming and optimization as potential areas for deeper investigation. The key takeaway is a reminder of the vast scope of algorithms available for different problem domains and the stimulation to explore these avenues.

**More Free Book**



Scan to Download

## Chapter 12: Answers to Exercises

The chapters you've provided are from a book that delves into various algorithms, their efficiency, and their practical applications. The chapters summarize algorithm concepts, data structures, problem-solving strategies, and exercises to reinforce learning. Here's a concise and cohesive summary of the content across these chapters:

### Chapter 1 - Binary Search and Big O Notation:

The chapter introduces binary search as an efficient searching algorithm for sorted lists, highlighting its process and efficiency. It explains Big O notation to describe algorithm performance by measuring the maximum steps needed relative to input size. For example, searching for a name in a sorted list takes logarithmic time,  $O(\log n)$ , whereas reading every name takes linear time,  $O(n)$ . The chapter clarifies that operations like dividing the list size (e.g., doubling) minimally impact Big O notation, focusing on overall computational growth rates rather than constants.

### Chapter 2 - Data Structures: Arrays and Linked Lists:

This chapter contrasts arrays and linked lists, explaining their use based on operations like insertions and retrievals. Arrays offer fast access but slow inserts, whereas linked lists excel at inserts but are slow in accessing

More Free Book



Scan to Download

elements. Practical applications include financial tracking and order queues in apps, leading to the importance of choosing the right data structure for specific requirements, such as quick reads or inserts.

### **Chapter 3 - Recursive Functions and Call Stacks:**

Here, recursive functions are explored with examples to highlight their call stack nature. Recursive solutions for summing lists, counting items, and finding maximums demonstrate the process of breaking down problems into manageable cases. The importance of base and recursive cases is underscored. Misusing recursion can lead to stack overflow errors if the stack grows indefinitely with no base case to stop it.

### **Chapter 4 - Further Exploration of Algorithms:**

Expanding on earlier concepts, this chapter delves into detailed algorithm analyses, including the divide-and-conquer strategy used in recursive algorithms like binary search. The relationship between operation types and their Big O notations is clarified, with exercises provided to solidify understanding of efficient algorithm design and execution.

### **Chapter 5 - Hashing and Consistent Hash Functions:**

The focus here is on hash tables and the necessity of consistent hash



functions for effective data retrieval and storage: finding a usable balance between hash distribution and performance. Evaluations include sample hash functions applied to phonebooks and other databases to assess efficiency in diverse contexts.

## **Chapter 6 - Graphs and Breadth-First Search:**

Graphs are introduced with breadth-first search (BFS) as a foundational algorithm to determine shortest paths and relationships among nodes. The BFS's application in practical tasks is demonstrated, with valid and invalid graph representations explored through exercises. Concepts like cycles and acyclic graphs are discussed to set up further learning in algorithmic graph theories.

## **Chapter 7 - Shortest Path via Dijkstra's Algorithm:**

Using Dijkstra's algorithm, the chapter illustrates finding the shortest path in weighted graphs. Concepts like infinity for unvisited nodes and cost tracking are clarified. While BFS works with unweighted graphs, Dijkstra's addresses weighted scenarios and negative-weight challenges.

## **Chapter 8 - Greedy Algorithms and Optimization:**

Greedy algorithms are explained as strategies to make the most favorable



immediate choice without guaranteeing the optimal final solution. Examples like the knapsack problem and daily schedules are used to depict their application. NP-complete problems, challenges that cannot be solved efficiently but can be approximated, are introduced.

## **Chapter 9 - Dynamic Programming and Optimization:**

Dynamic programming tackles complex problems by breaking them into simpler sub-problems, storing intermediate results to avoid redundant computations. Examples like the knapsack problem with item weights and values highlight this approach's efficiency in determining optimal solutions within constraints.

## **Chapter 10 - Advanced Algorithm Concepts:**

The chapter explores advanced topics such as k-nearest neighbors for classification tasks and machine learning predictions. Discussion extends to scalable solutions and recommendation systems, focusing on influencer-based weighted input and how a group of neighbors influences predictions. Practical applications in modern AI systems show the depth of algorithm versatility.

## **Bonus - Appendices and Index:**

**More Free Book**



Scan to Download



Supporting resources include exercises and an index for deeper exploration of terms and additional exercises spanning all chapters, enabling focused study on key algorithm topics and improving computational literacy.

This summary presents the logical progression of the book while introducing key data structures, algorithmic strategies, and problem-solving approaches foundational in computer science.

## **Install Bookey App to Unlock Full Text and Audio**

**Free Trial with Bookey**





# Read, Share, Empower

Finish Your Reading Challenge, Donate Books to African Children.

## The Concept



This book donation activity is rolling out together with Books For Africa. We release this project because we share the same belief as BFA: For many children in Africa, the gift of books truly is a gift of hope.

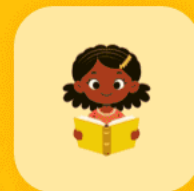
## The Rule



Earn 100 points



Redeem a book



Donate to Africa

Your learning not only brings knowledge but also allows you to earn points for charitable causes! For every 100 points you earn, a book will be donated to Africa.

Free Trial with Bookey

